HP-UX 11i March 2002 Release Notes

HP-UX Servers and Workstations

Edition 5

**HEWLETT®
PACKARD**

United States

Legal Notices

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

Trademark Notices

ActivePerl ® is a registered trademark of ActiveState Tool Corporation.

Apple® and Macintosh® are trademarks of Apple Computer, Inc., registered in the United States and other countries.

AppleShare® is a registered trademark of Apple Computer, Inc.

CHAMELEON™ is a trademark of NetManage, Inc.

DIGITAL™ and PATHWORKS™ are trademarks of Digital Equipment Corporation.

DiskAccess® is a registered trademark of Intergraph.

EXCURSION™ is a trademark of Digital Equipment Corporation.

Exeed® is a registered trademark of Hummingbird Communications Ltd.

eXodus™ is a trademark of White Pine Software, Inc.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. Hewlett-Packard is independent of Sun Microsystems.

MS-DOS® and Microsoft® are U.S. registered trademarks of Microsoft Corporation.

NTRIGUE™ is a trademark of Insignia Solutions, Inc.

NetMeeting ® is a registered trademark of Microsoft Corporation.

Netscape ® is a registered trademark of Netscape Communications Corporation.

OpenGL ® is a registered trademark of Silicon Graphics, Inc.

Oracle ® is a registered trademark of Oracle Corporation.

Oracle8 ™ is a trademark of Oracle Corporation.

OSF/Motif™ is a trademark of the Open Software Foundation, Inc. in the U.S. and other countries.

PC_Xware™ is a trademark, and WinCenter® is a registered trademark of Network Computing Devices, Inc.

REFLECTION® and WRQ® are registered trademarks of WRQ, Inc.

SGImeeting is a trademark of Silicon Graphics, Inc.

SunForum ® is a registered trademark of Sun Microsystems, Inc. in the United States and other countries.

VERITAS® is a registered trademark of VERITAS Software Corporation.

VERITAS File System™ is a trademark of VERITAS Software Corporation.

WinDD™ is a trademark of Tektronix, Inc.

X Window System™ is a trademark of the Massachusetts Institute of Technology.

This product includes software developed by the Apache Software Foundation (`http://www.apache.org/`).

Publication History

The manual publication date and part number indicate its current edition. The publication date will change when a new edition is released. The manual part number will change when extensive changes are made.

To ensure that you receive the new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

- First Edition: December 2000,   B3920-90091,   (HP-UX 11i),
  CD-ROM (Instant Information), and Web (`http://docs.hp.com`)

- Second Edition: June 2001,   5971-2229,   (HP-UX 11i),
  CD-ROM (Instant Information), and Web (`http://docs.hp.com`)

- Third Edition: September 2001,   5185-6522,   (HP-UX 11i),
  CD-ROM (Instant Information), and Web (`http://docs.hp.com`)

- Fourth Edition: December 2001,   5185-4304,   (HP-UX 11i),
  CD-ROM (Instant Information), and Web (`http://docs.hp.com`)

- Fifth Edition: March 2002,   5185-4391,   (HP-UX 11i),
  CD-ROM (Instant Information), and Web (`http://docs.hp.com`)

Contents

Contents

Contents

Contents

Contents

Contents

Contents

Contents

Contents

Contents

Contents

Contents

1 Overview of the HP-UX 11i Release

What's in This Chapter?

The purpose of this chapter is to introduce you to the HP-UX 11i release, as well as to help you to use these *Release Notes* most effectively. The following topics are covered:

- Where Do I Begin? (see page 18)
- What's in This Document? (see page 19)
- What's HP-UX 11i? (see page 21)
- What's New or Changed in the March 2002 11i OEs? (see page 24)
- What's New or Changed in the December 2001 11i OEs? (see page 25)
- What Was New or Changed in the September 2001 11i OEs? (see page 26)
- What Was New or Changed in the June 2001 11i OEs? (see page 27)

Where Do I Begin?

The *HP-UX 11i Release Notes* describe what is new, has changed, or has become obsolete since the initial release of HP-UX 11.0.

As with other HP-UX release notes, *HP-UX 11i Release Notes* does not completely document all the features of the latest 11i release (in this case, March 2002). Instead, it contains high-level information and pointers to more detailed operating system and product-specific documentation. Where appropriate, it also notes changes in the support of products.

If you need to . . .	Go to . . .	Located at . . .
check changes from HP-UX 11.0 to 11i	• *HP-UX 11i Release Notes* (you are here) These include all changes from HP-UX 11.0 to date.	• HP Instant Information CD-ROM • `http://docs.hp.com` • `/usr/share/doc/` directory
check changes from HP-UX 10.20 to 11.0	• *HP-UX 11.x Release Notes*	• `http://docs.hp.com` • `/usr/share/doc/` directory
install HP-UX 11i or update from 11.0[*]	• *Read Before Installing or Updating to HP-UX 11i* • *HP-UX 11i Installation and Update Guide*	• CD-booklet with the Operating Environment (OE) CD-ROM • Media Kit (supplied with the OE) • HP Instant Information CD-ROM • `http://docs.hp.com`
update to HP-UX 11i from 10.20[*]	• *Read Before Installing or Updating to HP-UX 11i* • *HP-UX 11i Installation and Update Guide*	• CD-booklet with the Operating Environment (OE) CD-ROM • Media Kit (supplied with the OE) • HP Instant Information CD-ROM • `http://docs.hp.com`

administer an HP-UX Operating Environment	• *Managing Systems and Workgroups: A Guide for HP-UX System Administrators*
	• *Managing Superdome Complexes: A Guide for HP-UX System Administrators*
		• Both *Managing* books are available from the following sources:
		  — HP Instant Information CD-ROM
		  — `http://docs.hp.com`

find a white paper	• See listings in `docs.hp.com` and `/usr/share/doc/` for the specific topic you want to read about.	• `http://docs.hp.com` (most white papers can be found here) • `/usr/share/doc/` directory (a few white papers can be found here)
develop code on HP-UX	• *HP-UX Software Transition Kit*	• `http://devresource.hp.com/STK`

[*]If you are *updating* to HP-UX 11i (as opposed to *cold-installing*), your system must first be running either HP-UX 11.0 or 10.20.

What's in This Document?

HP-UX 11i Release Notes contains information about the release of HP-UX 11i as of March 2002. It applies only to features that are included in the HP-UX operating system or one of the five Operating Environments (discussed later in this chapter under "What's HP-UX 11i?" on page 21).

In addition to changes made since the initial release of HP-UX 11i, this document includes information on changes to HP-UX that occurred *after* the last major system release, HP-UX 11.0, and which have been incorporated in HP-UX 11i. These latter changes were made as part of less comprehensive releases known as Extension Pack releases (for servers) and ACE (Additional Core Enhancements) releases (for servers and workstations). If your organization received one or more of these releases, then you may already be familiar with some of the information. (If, on the other hand, your organization did not receive any of the Extension Pack or ACE releases, then *all* of the information in this document will be new to you.)

Additional product-specific release notes files are located in the `/opt` directory, in sub-directories named `/opt/`*`product_name`*`/newconfig/RelNotes` (where *`product_name`* represents the name of the product for which those release notes apply). (For example, Distributed Computing Environment [DCE] release notes are located in the `/opt/dce/newconfig/RelNotes` directory.)

How Is Release Information Identified?

The HP-UX 11i releases covered in this document are identified as follows:

new (*or* updated) for March 2002
This material covers features that have been introduced or updated with the current release. (Throughout the *Release Notes*, the current release might be referred to as "HP-UX 11i as of March 2002," "HP-UX 11i," or just "11i.")

new (*or* updated) for December 2001
This material covers features that were introduced or updated in the December 2001 version of HP-UX 11i.

new (*or* updated) for September 2001
This material covers features that were introduced or updated in the September 2001 version of HP-UX 11i.

new (*or* updated) for June 2001
This material covers features that were introduced or updated in the June 2001 version of HP-UX 11i.

new at 11i original release
This material covers features that were newly introduced with the original release of HP-UX in December 2000.

The changes made for the June 2001, September 2001, December 2001, and March 2002 releases are summarized later in this chapter.

What's in the Remaining Chapters?

Here is a listing of the remaining chapters of these *Release Notes*:

- Chapter 2, "nPartition (Hard Partition) Systems," describes Hewlett-Packard's new high-performance HP-UX server environment.

- Chapter 3, "Workstation/Server Specific Information," presents information on which platforms support the 11i release, as well as other platform-specific information.

- Chapter 4, "HP-UX 11i Operating Environment Applications," presents information on each of the five Operating Environments, including the new "Minimal Technical Operating Environment."

- Chapter 5, "I/0 and Networking Cards," describes cards and drivers.

- Chapter 6, "Installation," describes new and changed aspects of installation.

- Chapter 7, "General System Administration and Performance Monitoring," describes changes which may be of particular interest to system administrators.

- Chapter 8, "Process, Threads, Memory, and Kernel Parameters," describes a variety of topics in these areas.

- Chapter 9, "New and Changed Disk and File Management Features," includes changes to striping, mirroring, JFS, and NFS.

- Chapter 10, "Internet and Networking Services," includes changes to sendmail, BIND, FTPD, STREAMS/UX, and other services.

- Chapter 11, "Security," presents new security features such as GSS API, executing protected stacks, and Kerberos Client software.

- Chapter 12, "Compatibility," describes various compatibility issues between HP-UX 11.0 and 11i.

- Chapter 13, "Programming," covers a variety of changes which will be of particular interest to programmers.

- Chapter 14, "Licensing Products," covers impending changes to LicensePower/iFOR and LSSERV.

- Chapter 15, "New and Changed Internationalization Features," describes information that will be of interest to localizers or international users of HP-UX.

What's HP-UX 11i?

HP-UX 11i provides new hardware enablement, additional software functionality, and various HP-UX applications bundled into Operating Environments. HP-UX 11i is the recommended next-level enterprise release for all HP-UX systems currently running HP-UX 10.x or 11.0.

See the *HP-UX 11i Installation and Update Guide* for information on disk and memory requirements.

Performance Considerations

The performance of HP servers and applications with the HP-UX 11i Operating Environments is as good or better as compared to HP-UX 11.0 installed in the same configuration. It must be kept in mind, however, that in the vast majority of cases, HP-UX is likely to be installed as part of one of the predefined HP-UX 11i Operating Environments.

The HP-UX 11i Operating Environments provide numerous benefits, including lower cost as well as simplification of product ordering, installation, integration, and support. While the HP-UX 11i Operating Environments provide these desirable benefits, they all involve more software components and daemons than a base HP-UX 11.0 or HP-UX 11i operating system. Because of this, an Operating Environment may have somewhat greater system overhead compared to a base operating system installation. This increase in system overhead, while modest, can nevertheless make it erroneously appear that system and application performance are degraded with HP-UX 11i as compared to HP-UX 11.0.   In fact, performance with HP-UX 11i is generally better than HP-UX 11.0 when the software configurations are the same.

While not recommended, one of the options for installing HP-UX 11i is as a base operating system only, without an Operating Environment. Installing HP-UX 11i as a base operating system will give the greatest performance and the lowest system overhead but eliminates the advantages provided by the HP-UX 11i Operating Environments.

The HP-UX 11i Operating Environments (OE)

**new at 11i
original release**

Beginning with HP-UX 11i, the operating system is delivered as part of an HP-UX Operating Environment (OE), an integrated and tested software solution containing the operating system and selected applications.

The following software bundles are always delivered with an Operating Environment. Thus, if you do a minimum install, these bundles, plus the HP-UX applications within each OE, will be loaded:

- `HPUXBase32` or `HPUXBase64`, which consist of operating system commands and libraries bundled for either 32- or 64-bit systems.

- `HPUXBaseAux`, which includes system manageability software such as Software Distributor (SD) and the Partition Manager (parmgr). This additional required core software is also referred to as the Auxiliary OS.

- HWEnable11i patches, installed by default, which enable the latest hardware on HP-UX 11i. Only a few patches are added at each quarterly release to minimize the number of patches needed to support new systems and I/O devices.

- Network and I/O drivers, installed by default, that are required by the operating system and other selectable drivers. The default drivers include the following:

  — `FibrChanl-00`: drivers for the PCI Fibre Channel HB adapter (64-bit OS)
  — `RAID-00`: driver for the PCI RAID 4Si controller (64-bit OS)
  — `FDDI-00`: FDDI drivers (32-bit OS)
  — `GigEther-00`: 1000Base-T/SX (Gigabit Ethernet) drivers
  — `GigEther-01`: 1000Base-T driver for HP server rp7410 core (64-bit only)

  (See "Selectable Network Drivers" on page 104 in Chapter 4 for a list of the selectable drivers.)

- `OnlineDiag`, installed by default, which provides the HP-UX 11i Diagnostics that include STM, ODE, and EMS hardware monitors.

- `CDE-English`: CDE language (such as for English or alternate languages).

Operating Environments for HP Commercial Servers

The following lists each 11i Operating Environment available for HP servers. These Operating Environments are supported on all HP 9000 servers, including technical servers. They are not supported on technical workstations.

- HP-UX 11i OE is the standard Internet server environment for HP-UX systems. It includes Java, the Apache-based Web Server, Netscape Communicator, WebQoS, and other applications.[*] This OE is included without additional charge.

- HP-UX 11i Enterprise OE contains the HP-UX 11i OE and additional applications to enable an enterprise-level server. Products include OnLineJFS 3.3, GlancePlus, MirrorDisk/UX, and other applications.[*]

- HP-UX 11i Mission Critical OE contains both the HP-UX 11i OE and the HP-UX 11i Enterprise OE plus applications to enable a mission-critical server, such as MC/ServiceGuard and HP-UX Workload Manager.[*]

NOTE   The 11i, Enterprise, and Mission Critical OEs are also supported on technical servers. (They are not supported on technical workstations, however.)

Operating Environments for HP Technical Workstations and Technical Servers

updated for June 2001

The following 11i Operating Environments are available for HP 9000 technical workstations and technical servers:

- HP-UX 11i Minimal Technical OE (new for June 2001) contains all the base functionality that is common to the other four OEs, including the base 32/64 bit HP-UX Operating System, network drivers, and other always-installed functionality.

---

*All products included with each OE are described in Chapter 4. Some applications may require post-installation configuration; see the current *HP-UX 11i Installation and Update Guide* for details.

However, it includes only a sparse set of high-demand applications so as not to increase purchase costs, support costs, or license costs over the base OE.[*]

- HP-UX 11i Technical Computing OE contains applications to enable a technical workstation or technical server.[*]

What's New or Changed in the March 2002 11i OEs?

- HP now supports hard partitions on the HP rp7410 server (model 9000/800/rp7410). (See Chapter 2, "nPartition (Hard Partition) Systems," on page 29.)

- HWE Patch Bundle updated to support new servers, I/O adapters, and peripherals. (See "Hardware Enablement Patch Bundle (HWEnable11i)" on page 46.)

- SCSI Driver `c8xx` updated to support two new SCSI PCI Ultra160 Host Bus Adapters. (See "SCSI Driver c8xx" on page 50.)

- HP Apache-based Web server updated to version 1.3.19.23 with enhancement of `chroot` feature. (See "HP Apache-based Web Server for HP-UX" on page 70.)

- Version 1.3 of HP-UX Runtime Environment for the Java 2 (RTE) Platform added. (See "HP-UX Runtime Environment for the Java 2 Platform" on page 73.)

- GlancePlus Pak, updated to version C.03.55, no longer supports HP-UX 10.20. (See "GlancePlus Pak" on page 80.)

- MC/ServiceGuard updated to version A.11.14. (See "MC/ServiceGuard" on page 88.)

- FirstSpace VRML obsoleted. (See "FirstSpace VRML Viewer" on page 92.)

- HP MPI updated to version 1.7.2. (See "HP Message-Passing Interface (MPI)" on page 94.)

- Ignite-UX updated to version B.3.6. (See "Ignite-UX (IUX)" on page 98.)

- IPFilter/9000 updated to version A.03.05.03. (See "IPFilter/9000" on page 101.)

- Java Out-of-Box (JAVAOOB) has been added to allow large server-side Java applications to run. (See "Java Out-of-Box (JAVAOOB)" on page 102.)

- Networking and I/O card drivers updated. (See "Networking and I/O Card Drivers Pre-Installed" on page 109.)

- Gigabit Ethernet now supports the HP rp7410 server core 1000Base-T product. (See "HSC and PCI 1000Base-SX/T (Gigabit Ethernet)" on page 112.)

- RAID 4Si controller available as a patch in a new firmware version (U.01.06). (See "HP RAID 4Si Driver" on page 116.)

- The new Update-UX product now enables customers to install the `update-ux` script and its accompanying `update-ux` manpage on a 10.20 or 11.00 system. (See "New Update-UX Product" on page 124.)

- BIND updated to version 9.2. (See "BIND 8.1.2" on page 171.)

- Support for CXperf discontinued. (See "Changes to the linker/dld Interface" on page 222.)

What's New or Changed in the December 2001 11i OEs?

- The base VERITAS Volume Manager (VxVM) added as an always-installed application for all OEs. (See "Base VERITAS Volume Manager (VxVM)" on page 68.)

- The HP Apache-based Web server updated with version 1.3.19.21. (See "HP Apache-based Web Server for HP-UX" on page 70.)

- Support Tools updated with version A.29.00 (`swlist` ID B.11.11.05.10) to incorporate fixes and support for new devices. (See "HP-UX Support Tools: STM, ODE, & EMS Hardware Monitors" on page 73.)

- The Judy libraries now feature the ability to search for empty locations in the array (search empty). (See "Judy Libraries" on page 212.)

- Servicecontrol Manger updated with version A.02.05. (See "Servicecontrol Manager (SCM)" on page 77.)

- Enterprise Cluster Manager Toolkit updated with version B.01.07. (See "Enterprise Cluster Master (ECM) Toolkit" on page 85.)

- MC/ServiceGuard updated with version A.02.00. (See "MC/ServiceGuard" on page 88.)

- MC/ServiceGuard NFS Toolkit updated. (See "MC/ServiceGuard NFS Toolkit" on page 89.)

- HP Intrusion Detection System/9000 added as a selectable application for the three commercial server OEs. (See "HP Intrusion Detection System/9000 (IDS/9000)" on page 97.)

- IPFilter/9000 added as a selectable application for the three commercial server OEs. (See "IPFilter/9000" on page 101.)

- The always-installed Fibre Channel drivers updated. (See "Fibre Channel Tachlite Driver" on page 114.)

- New always-installed Gigabit Ethernet (`GigEther-01`) driver to support the 1000Base-T/SCSI card (A6794A) for HP-UX 11i, 64-bit only. (See "HSC and PCI 1000Base-SX/T (Gigabit Ethernet)" on page 112.)

- Additional hardware enablement. (See "Hardware Enablement Patch Bundle (HWEnable11i)" on page 46.)

What Was New or Changed in the September 2001 11i OEs?

- The Judy Libraries product moved from the selectable applications and now installed with all Operating Environments. (See "Judy Libraries" on page 212.)

- Perl programming language installed with all Operating Environments. (See "Perl Programming Language" on page 75.)

- The base VERITAS Volume Manager (VxVM) added as a selectable application for the three commercial server OEs. (See "Base VERITAS Volume Manager (VxVM)" on page 68.)

- GlancePlus Pack updated with version C.03.35.00.

- HP Visualize Conference updated with version 1.4.

- Ignite-UX updated with version 3.4 to incorporate defect fixes.

- HP MPI updated with version 1.7. (See "HP Message-Passing Interface (MPI)" on page 94.)

- HP MLIB updated with version B.08.01. (See "High Performance Math Libraries (HP MLIB)" on page 92.)

- Support Tools updated with version A.28.00 (`swlist` ID B.11.11.04.09) to incorporate fixes and support for new devices.

- MC/ServiceGuard updated with version A.11.13. (See "MC/ServiceGuard" on page 88.)

- MC/ServiceGuard NFS Toolkit updated. (See "MC/ServiceGuard NFS Toolkit" on page 89.)

- OpenGL 3.0 Graphics Developers kit updated. (See "OpenGL 3D Graphics Developers Kit and Runtime Environment" on page 91.)

- HP-UX Workload Manager updated with version A.01.02. (See "HP-UX Workload Manager" on page 86.)

- HP-UX Workload Manager Oracle Database Toolkit added to the Mission Critical Operating Environment. (See "HP-UX Workload Manager Oracle , Database Toolkit" on page 87.)

- HP PRM updated with version C.02.00. (See "Process Resource Manager (PRM)" on page 82.)

- CIFS/9000 Server updated with version A.01.07 and CIFS/9000 Client updated with version A.01.06. (See "CIFS/9000 Client and CIFS/9000 Server" on page 69.)

- PCI RAID driver updated to support A400 and A500 servers. (See "Selectable Network Drivers" on page 104.)

- `ATM` drivers updated to repair defects and improve performance for PCI and HSC.

- `TermIO` drivers updated to include 3.3Volt support and defect repairs for PCI MUX and EISA MUX.

- Apache-based Web Server updated. See "HP Apache-based Web Server for HP-UX" on page 70.

- Java™ Runtime Environment and Java™ Plug-In updated.
  (See "Plug-In for the Java 2 Platform for Netscape Communicator" on page 77.)

- iPlanet Directory Server updated with version B.04.13.
  (See "Netscape Directory Server (J4258BA)" on page 103.)

- Additional hardware enablement.
  (See "Hardware Enablement Patch Bundle (HWEnable11i)" on page 46.)

What Was New or Changed in the June 2001 11i OEs?

- New Minimal Technical Operating Environment now available.
  (See "HP-UX 11i Minimal Technical Operating Environment" on page 91.)

- Judy Libraries added.
  (See "Judy Libraries" on page 212.)

- iCOD software removed from the Operating Environments.

- SD/Update UX updated to enable updating the Operating Environments and upgrading from the Enterprise OE to the Mission Critical OE.

- Diagnostics updated to support new hardware.
  (See "Diagnostics: EMS Hardware Monitors" on page 139.)

- Partition Manager updated to support new hardware.
  (See "Partition Manager (parmgr)" on page 35.)

- PCI TachyonTL driver updated to support the A6684A and A6685A.
  (See "Fibre Channel Tachlite Driver" on page 114.)

- HyperFabric driver updated to support the A6092A and A6386A network cards.
  (See "HyperFabric2 PCI Fiber Adapter" on page 117.)

- PCI RAID driver updated to fix defects.
  (See "HP RAID 4Si Driver" on page 116.)

- Apache-based Web Server updated to fix defects and enhance performance.

- Java™ Runtime Environment and Java™ Plug-In updated to fix defects and enhance performance.

- Additional hardware enablement.

New Software Pack -- Optional HP-UX 11i Core Enhancements

**New for
December 2001**

The new "Software Pack -- Optional HP-UX 11i Core Enhancements" CD is included in the HP-UX 11i media kit for December 2001. You can also download Software Pack from HP's Software Depot (`http://software.hp.com`). Select **enhancement releases**, and then **hp-ux software pack**. The December 2001 Software Pack includes Processor Sets and IPv6.

(See "Processor Sets Available on Software Pack" on page 131 and "IPv6 Available on Software Pack" on page 167.

2　nPartition (Hard Partition) Systems

Introduction

Hewlett-Packard's new nPartition (hard partition) servers provide highly configurable, high-performance HP-UX system environments.

updated for March 2002

In addition to the servers listed below, HP now supports hard partitions on the HP rp7410 server (model 9000/800/rp7410).

new at 11i original release

Currently, HP supports hard partitions on the following servers:

- HP Superdome 16-way (model 9000/800/SD16000)
- HP Superdome 32-way (model 9000/800/SD32000)
- HP Superdome 64-way (model 9000/800/SD64000)
- HP rp8400 server (model 9000/800/S16K-A)

The ability to create hard partitions allows you to configure a single nPartition server as either one large system or as multiple smaller systems. Because hard partitions are managed through software, you can reconfigure a server without physically modifying the server's hardware configuration.

As a result, an nPartition server can run multiple instances of the 11i operating system on a single server. This capability is accomplished by defining multiple partitions within an nPartition server.

Each partition definition establishes a subset of a server's hardware resources that are to be used as a system environment for booting a single instance of HP-UX.

All processors, memory, and I/O in a partition are available exclusively to the software running in the partition. Thus, each partition runs a single instance of the Boot Console Handler (BCH) interface and HP-UX.

You can reconfigure partitions to include more, fewer, and/or different hardware resources, but this will require shutting down the operating system running in the partition and resetting the partition as part of reconfiguring it.

Two books provide details on performing nPartition administration tasks. The *HP System Partitions Guide*, part number B2355-90746, has current details on managing all nPartition sever models. *Managing Superdome Complexes: A Guide for HP-UX System Administrators*, part number B2355-90702, covers using the December 2000 HP-UX 11i release on HP Superdome servers. Both books are available at HP's documentation web site, `http://docs.hp.com`, and on the HP-UX 11i Instant Information CD.

Superdome Systems at HP-UX 11i

**new at 11i
original release**

Superdome models differ with regard to the characteristics shown in Table 2-1.

Table 2-1   **Maximums for Superdome Models**

	Superdome 16-way	Superdome 32-way	Superdome 64-way
Number of Cells	4	8	16
Number of Processors	16	32	64
Amount of Memory (using 512 MB DIMMs)	64GB	128GB	256GB
Number of I/O Slots (internal chassis)	48	48	96

Note: Each I/O expansion (IOX) cabinet can provide an additional six I/O chassis containing 12 slots each, providing a total of 72 slots for each I/O expansion cabinet.

Machine Identifier

The `uname -i` command on your nPartition systems may not return a unique value for each system. To guarantee compatibility on current and future platforms, use the new interfaces to *getconf* (1) and *confstr* (3C) to retrieve unique machine identifiers.

For example, use the following `getconf` commands instead of `uname -i` in order to obtain a unique identifier for an HP Superdome server complex (`getconf _CS_MACHINE_IDENT`) or a hard partition (`getconf _CS_PARTITION_IDENT`) within a complex:

```
# uname -i
1945761737
# getconf _CS_PARTITION_IDENT
Z3e02955673f9f7c9_P1
# getconf _CS_MACHINE_IDENT
Z3e02955673f9f7c9
#
```

As shown in the above example, the `getconf` commands return partition-unique and machine-unique (complex-unique) identifiers on HP PA-RISC nPartition servers.

The output returned by `uname` is *not necessarily a unique identifier* and should not be used for licensing or other purposes that require unique IDs.

These interfaces are documented in the manpages and in Chapter 13, "Programming," in the section "Machine Identifier Changes to confstr" on page 227.

Hard Partition Hardware Path Format

new at 11i original release

The HP-UX command `ioscan` reports the hardware paths for components within the partition in which the command is issued. (Note that the `ioscan` command reports information only for the *currently active* hardware components *in the local partition*. It does not report details for hardware not assigned to the local partition.)

On nPartition systems, HP-UX hardware paths conform to the following format:

a/b/c/d/e.f.g

The components of a hardware path are as follows:

a   The global cell number

b   A processor (10-13), memory (5), or system bus adapter (0). Each nPartition server I/O chassis has a single system bus adapter.

c   A local bus adapter (the LBA, one for each PCI card slot in the chassis). The LBA connects its corresponding PCI card slot with the system bus adapter.

                (Note that the LBA number *is not necessarily the same* as the PCI slot number. Use the `rad -q` command to list all PCI slots and their corresponding hardware paths.)

d   The card's address on the slot's PCI bus.

                Typically this is 0 (zero), although the Superdome core I/O card has multiple devices in a single card.

e   The function for the I/O card. Typically this is 0 (zero) for single-function cards.

f   The target of the I/O device, or SCSI ID.

g   The device-specific address such as a SCSI controller (initiator).

Refer to the *ioscan* (1M) manpage for details on using `ioscan` to list hardware path information.

This hardware path format is used only on nPartition systems.

New and Modified Hard Partition Commands

**new at 11i
original release**

Several system administration commands are provided with HP-UX 11i for creating and maintaining partitions on all nPartition systems.

Additionally, the existing commands `shutdown`, `reboot`, and `setboot` have been modified to support nPartition platforms.

New Commands

Below are brief descriptions of the seven new commands used to manage an nPartition server complex:

Command	Description
parcreate (1M)	Creates a new partition.
parmodify (1M)	Modifies an existing partition.
parstatus (1)	Provides information about an entire Superdome complex, including partition information and available resources in the complex.
parremove (1M)	Removes an existing partition.
parunlock (1M)	Unlocks the Stable Complex Configuration Data or Partition Configuration Data.
fruled (1)	Turns locator LEDs on/off for cells, cabinets and I/O chassis.
frupower (1M)	Enables or disables power to a cell or I/O chassis, or displays the power status of cells or I/O chassis.

Modified Commands

The commands `shutdown` and `reboot` have been modified for systems that support hard partitions. Two new options, `-R` and `-H`, have been added to both these commands.

The `-R` option is used to shut down the system to a "ready to reconfig" state and reboot automatically. The `-R -H` option is used to shut down the system to a "ready to reconfig" state but without rebooting. A partition must be put into the "ready to reconfig" state before its newly assigned cells can be activated. (A normal reboot will not activate newly assigned cells.) The "ready to config" state is also necessary to complete the removal of active cells that have been unassigned from the partition. The commands' default behaviors, or the behaviors of all the other existing options, remain unchanged.

See *Managing Superdome Complexes: A Guide for HP-UX System Administrators*, part number B2355-90702, and *HP Systems Partitions Guide*, part number B2355-90746, for more information.

The interpretation of Autoboot and Autosearch in the command `setboot` has changed for systems that support hard partitions. The firmware interprets the bits in combination and not individually as done before.

In order to approximate the traditional behavior of `setboot`, the user input for the Autoboot and Autosearch flags is internally mapped to the right combination to achieve the desired behavior. This mapping should be transparent to the user of `setboot`, but it might not be so apparent to those accessing the firmware using other means. For the primary path, the boot action corresponds to the Autoboot and Autosearch flags in the following manner:

AutoBoot (setboot -b)	AutoSearch (setboot -s)	nPartition Boot Action
off	off	Go to the Boot Console Handler (BCH) and prompt the user.
on	off	Attempt the primary path; on failure go to BCH.
on	on	Attempt to boot the PRI boot path; if fail to boot PRI, attempt to perform the HAA boot path's boot action.
off	on	Skip the PRI boot path; attempt to perform the HAA boot path's boot action.

Additionally, systems with hard partitions support a boot action for each path. However, boot actions for paths other than the primary path cannot be set using `setboot`. (In this case, the default boot action for boot paths on nPartitions is "skip this device and try next path," [a path flag value of 2.]) Instead, setting both the Autosearch and Autoboot flags to "on" (`setboot -b on -s on`) means that if the PRI path fails to boot, then booting the HAA boot path will not be attempted unless the HAA path flag is set to a non-default value (such as 1 or 2, which attempt booting HAA).

In the default case (when the HAA boot path is not set to "boot from the path"), specifying `setboot -b on -s on` will not cause an alternate path to be automatically booted when the primary path fails. Instead the user will be prompted.

Documentation Changes

Eight new manpages document the syntax and usage of the new commands. The *reboot* (1M) and *shutdown* (1M) manpages have been updated to document the `-R` and `-H` options, and the *setboot* (1M) manpage has been updated as well.

Partition Manager (parmgr)

new at 11i original release

A new system administration tool, Partition Manager (parmgr), supports the initial and ongoing configuration of systems that support hard partitions, (such as Superdome systems), and also provides extensive information about the status of a Superdome complex. Partition Manager can be installed on any HP computer system that supports hard partitions, subject to the system requirements listed in "System Requirements" on page 36.

Partition Manager can be launched from SAM or directly from the command line. It can also be launched from a web browser running on a PC; this requires that the Apache-based Web Server (described under "HP-UX 11i Operating Environment" on page 67) be running on a partition. The URL for the parmgr launch page is **http://*hostname*:1188/parmgr/**. (See the parmgr online help or *Managing Superdome Complexes: A Guide for HP-UX System Administrators*, for details.)

The Partition Manager's functionality includes the ability to do the following:

- create, modify, and delete partitions. (Modifying partitions includes adding cells to a partition, removing cells from a partition, changing the name of a partition, and changing the per cell "Use On Next Boot" attribute. For more detail, see *HP System Partitions Guide*. Also see *Managing Superdome Complexes: A Guide for HP-UX System Administrators* for details on using the original HP-UX 11i release on Superdome servers.)

- display a complete hardware inventory of an entire complex

- display the status of key complex components

- check for problems or unusual conditions in the complex

- manage power to cells and I/O chassis

- turn on/off attention indicators associated with cells, I/O chassis, I/O cards, and cabinets

Selected configuration screens of the Partition Manager can also be launched via the command line by use of the -t *task* option. See the *parmgr* (1M) manpage for more details.

NOTE

Partition Manager includes online help that is displayable within a web browser. An appropriate version of Netscape is included in the Operating Environment (OE) bundle that is shipped with this release of HP-UX. Please install the OE bundle on any machine running Partition Manager for full access to the online help.

updated for December 2001

Partition Manager version B.11.11.01.03 includes updates to both the PARMGR and PARMGR-HELP file sets. Beginning with this release, a new fileset is also included: PARMGR-MAN, containing the *parmgr* (1M) manual page. This fileset supersedes earlier versions of the manual page that were previously included with the HP-UX core manual pages.

This update to Partition Manger includes the following major improvements:

- support for new HP server systems

- operability on partitions with HP Virtual Partitions (vPars) installed

- improvements to the online help

- various bug fixes and enhancements, including the following:

  — Messages and online help associated with card OLAR (online addition and replacement of I/O cards) have been improved for readability and clarity.

  — Messages and online help associated with partition shutdown and reboot procedures have been improved for readability and clarity, and to accommodate the special needs of partitions that are running vPars.

  — An error was fixed that could cause a hang on systems with a kernel that is not properly configured for OLAR.

  — Partition Manager now displays limited information about PCI cards that are installed in remote partitions and inactive I/O chassis, instead of reporting "unknown card."

  — The I/O Details Notes tab now shows additional messages when I/O information is unavailable or incomplete.

  — The dependency on library patches has been removed. Therefore, the current release of Partition Manager can be installed on any HP-UX 11i system, regardless of the patch level.

  — The **Empty Cell/Chassis Slots** container object has been renamed to **Empty Cell/IOChassis Slots**.

For more information, see the online Partition Manager B.11.11.01.03 Release Notes included with the Partition Manager online help in `/opt/webadmin/parmgr/help/C/release-notes.txt`.

updated for June 2001

Partition Manager version B.11.11.01.01 includes updates to both the `PARMGR` and the `PARMGR-HELP` file sets. This update includes improvements to the online help system and various bug fixes and enhancements, summarized below.

System Requirements

Partition Manager requires HP-UX Release 11i.

Partition Manager B.11.11.01.01 depends on the following additional products:

Product	Notes
PHCO_23510 `libfab.1` library patch [no longer needed with version B.11.11.01.03 (December 2001)]	Included with the installation media.
OBAM B.11.00.05.2.06 Object Action Manager user interface	Delivered with HP-UX 11i.
SystemAdmin B.11.11 HP-UX System Administration Tools	Delivered with HP-UX 11i.

PARMGR(1M) Manual Page Error

The HP-UX 11i manual page for *parmgr* (1M) contains the error described below. Note that this error only occurs in the PDF version of the *HP-UX Reference (Volume 4 of 9)*, "System Administration Commands," published on the Hewlett-Packard Technical Documentation web site (`http://docs.hp.com`) in December 2000.

In the SYNOPSIS section, the pathname of the `parmgr` command was incorrect. The correct pathname is:

`/opt/parmgr/bin/parmgr`

Additional Information

For detailed information about using Partition Manager, consult the online help installed from the `PARMGR-HELP` file set. Within Partition Manager, the help can be accessed through the **Help** menu on the tool bar. You can also view the online help simply by pointing your web browser to:

`file:/opt/webadmin/parmgr/help/C/overview.html`

Summary of Changes

- Empty cell slots and empty or blocked I/O chassis slots are no longer listed as Available Resources; they are now listed under a separate category: "Empty Cell/Chassis Slots."

- The printed output of I/O Details and I/O Slot property sheets has been reformatted for improved readability.

- A **Print** button was added to the Card Critical Resource Analysis dialog.

- The HA status of the I/O fans in Superdome CPU cabinets is now reported as "N/A," since this information is not made available by system firmware.

- Complex profile locks are now freed after an unlock failure.

- Additional PCI class codes were made available for use in I/O details filters.

- Partition Manager is now able to work around certain errors that used to be fatal (for example, a cell reports a fatal firmware error). Warning messages are displayed, but you can choose to continue despite the error.

- A **Show Details** button was added to the Delete Partition dialog, allowing you to review, save, or print the partition details prior to deletion.

- Partition Manager now flashes the amber attention LEDs, instead of setting them to a steady-on state (in accordance with the new PCI standard specification). See "New Attention Indicator Behavior" on page 40.

- Double-clicking on a cell or I/O chassis in the Create Partition Summary window now displays details (equivalent to clicking **Show Details**).

- CPU Type is now formatted as a hex value in the Cell Details display.

- When attempting to power-on an I/O chassis, the Modify Partition task now power-cycles the cell connected to the chassis, instead of attempting (and failing) to power-on the chassis directly.

- A more meaningful error message is displayed when a non-root user attempts to start Partition Manager.

- When launching Partition Manager from a PC web browser, a graphic displayed in the initial launch page has been corrected.

- If all cells currently assigned to a partition are removed, and others are added, an incorrect message is no longer displayed.

- Partition Manager no longer warns about the state of the "Failure Usage" flag when removing a cell from a partition.

- The Modify Partition dialog now checks for valid partition data and refreshes the display, ensuring that compatibility checks are accurate.

- The Cell Details dialog now lists DIMMs using the same identification codes that are printed on the cell board.

- The following have been fixed:

  — an error preventing the retrieval of the list of disks attached to an I/O card

  — an error preventing the deletion of the local partition

- New problem checkers report:

  — active partitions with inactive cells that do not have "Use on Next Boot" set to "no."

  — any HDCI errors that have occurred

- Messages and online help have been improved for the following:

  — deleting a partition

  — deleting cells from a partition

  — when the GSP is locked or has pending actions

  — Confirm Modify Partition dialog

  — Online Addition/Replacement of an I/O card

Service Processor (GSP or MP)

Each of HP's nPartition servers include a service processor that provides server-wide status, maintenance, and console interfaces.

NOTE

The service processor in HP servers is sometimes called the Management Processor (MP) and sometimes the Guardian Service Processor (GSP).

Regardless of the name, the service processor in these servers provides approximately the same features and performs essentially the same role.

Throughout this document the term "service processor" refers to both the MP and GSP service processors.

hd_fabric Driver

new at 11i original release

Although the HP-UX 11i install kernel does not include a `/stand/system` entry for the new 11i driver, `hd_fabric`, an entry is added to the system file as part of the installation process on nPartition systems. When the `/usr/sbin/mk_kernel` command builds a kernel, the `hd_fabric` driver is then built into the HP-UX kernel.

This only affects nPartition systems. HP-UX kernels built with `hd_fabric` present in the system file will not cause compatibility issues on non-nPartition systems.

Although it is not required for booting HP-UX on hard partitions, the `hd_fabric` driver supports partition configuration commands (`/usr/sbin/parstatus` and others) and related functionality in the Partition Manager.

New Attention Indicator Behavior

**updated for
September 2001**

The `fruled` command and Partition Manager now blink the attention indicators (LEDs) on HP Superdome cells and I/O chassis rather than illuminating them in a steady-on state, as was the case with the original 11i release.

Likewise, on HP rp7410 and rp8400 servers, the `fruled` command and Partition Manager behaviors are to blink cell attention indicators. (HP rp7410 and rp8400 servers do not provide user-controllable I/O chassis or cabinet-level attention indicators.)

In addition, the SAM (`/usr/sbin/sam`) and Partition Manager utilities now blink PCI slot attention LEDs for procedures that involve locating PCI slots. Both HP Superdome and HP rp8400 servers support this PCI slot attention indicator usage.

For example, the **`fruled -o -c2`** command now causes cell 2's attention indicator to blink. On HP Superdome servers, the **`fruled -o -i0/1/3`** command blinks the attention indicator for cabinet 0, I/O bay 1, chassis 3.

The `/usr/sbin/rad` command still provides the same functionality for managing PCI card slot attention indicators. You should use the **`rad -f attention #-#-#-#`** command to blink the attention LED when locating the specified slot (*#-#-#-#*). See *rad* (1M) for options.

Benefits

These changes conform to the PCI Hot-Plug specifications for attention indicator states and meanings, as detailed in the following table:

Table 2-2 **Attention Indicator (LED) States and Meanings**

Attention LED State	Meaning
OFF	Normal; not selected, normal operation.
ON	Attention; operational problem. *Supported for PCI card slot LEDs only.* Service is required; problems have been identified with the component.
BLINKING	Locate; user requested to locate hardware. User has selected the hardware for use in a service operation.

Impact

These changes affect all current HP hard-partitionable servers and will be implemented for all future HP hard-partitionable servers, as well. (Note that only Superdome servers were supported by HP-UX releases with the old behavior.)

However, it wasn't until the HP-UX 11i June 2001 release that SAM and Partition Manager started to blink PCI slot attention indicators during tasks involving PCI cards

or card slots. As a result, sites that have HP Superdome servers running both the original HP-UX 11i release and the HP-UX 11i June 2001 (or later) release can have attention indicators at different states (steady-on and blinking).

NOTE   These changes do not affect the behavior of the HP rp7410, rp8400 or HP Superdome cabinet "attention" lights, which are not user-controllable.

These changes also do not affect HP Superdome cabinet number LCD behavior.

Documentation

Details about how the new `fruled` and Partition Manager affect cell and I/O chassis attention indicators are provided in the *fruled* (1) manpage, the *HP System Partitions Guide*, and the Partition Manager online help.

The `rad` command, whose functionality has not changed, is documented in the *rad* (1M) manpage.

3 Workstation/Server Specific Information

Supported Systems

updated for September 2001

HP-UX 11i continues to support both a 32-bit and 64-bit version of the HP-UX kernel.

The tables below outline the supported HP-UX 11i configurations for HP servers and workstations.

NOTE   The information in the following tables is subject to change. For the most up-to-date data, refer to `http://devresource.hp.com/STK/hpux_faq.html`.

Table 3-1   **Servers**

Bezel or Model	32-bit Support	64-bit Support	Comments
A-Class: A180, A180C	Yes		PA-7300LC
A-Class: A400, A500		Yes	PA-8500 and newer processors
bp2200 blade (in an HP bh3710 chassis)		Yes	PA-8600 processor only; HP-UX `model` string: 900/800/HPServerBlade8600/3xx2[a]
D-Class: D270/370, D280/380, D390	Yes	Yes	PA-8*xxx* (See Table 3-2 for firmware revision information.)
D-Class: All other	Yes		PA-7*xxx*
K-Class: Kx50, Kx60, Kx70, Kx80	Yes	Yes	PA-8*xxx*
K-Class: Kx00, Kx10, Kx20	Yes		PA-7*xxx*
L-Class: L1000, L2000, L3000		Yes	PA-8500 and newer processors
L-Class: L1500, L3000		Yes	PA-8700 and newer processors
N-Class: N4000		Yes	PA-8500 and newer processors
R-Class: R380, R390	Yes	Yes	PA-8200 and newer processors
rp2400		Yes	PA-8500 processor; HP-UX `model` string: A400-44[a]
rp2430		Yes	PA-8700 processor; HP-UX `model` string: A400-6X[a]
rp2450		Yes	PA-8500 and newer processors; HP-UX `model` string: A500-44 or A500-5X[a]
rp2470		Yes	PA-8700 processors; HP-UX `model` string: A500-6X or A500-7X[a]
rp5400		Yes	PA-8500 and PA-8600 processors; HP-UX `model` string: L1000-36, -44, -5X[a]

Table 3-1 **Servers**

Bezel or Model	32-bit Support	64-bit Support	Comments
rp5430		Yes	PA-8700 processors; HP-UX `model` string: L1500-5X, -6X, or 7X[a]
rp5450		Yes	PA-8500 and PA-8600 processors; HP-UX `model` string: L2000-36, -44, 5X[a]
rp5470		Yes	PA-8700 processors; HP-UX `model` string: L3000-5X, -6X, or -7X[a]
rp7400		Yes	PA-8700 processors; HP-UX `model` string: N4000-65 or N4000-75[a]
rp7410		Yes	(See Chapter 2, "nPartition (Hard Partition) Systems.")
rp8400		Yes	PA-8700 and newer processors; HP-UX `model` string: S16K-A[a] (See Chapter 2, "nPartition (Hard Partition) Systems.")
Superdome systems: Superdome 16-way, Superdome 32-way, Superdome 64-way		Yes	(See Chapter 2, "nPartition (Hard Partition) Systems.")
T-Class: T5*xx*	Yes		PA-7*xxx*
T-Class: T6*xx*	Yes	Yes	PA-8*xxx*
V-Class: V22*xx*, V2500, V2600		Yes	PA-8200 and newer processors

a. Use the `model` command to display the model string.

The following D-Class systems support HP-UX 11i 64-bit operation:

Table 3-2 **D-Class Support for 11i 64-bit**

Model	CPU	Speed (MHz)	32/64 bit	Minimum Firmware Revision
D270/370	PA8000	160	Both	38.27 or later
D280/380	PA8000	180	Both	38.27 or later
D390	PA8200	240	Both	38.28 or later

Table 3-3 **Workstations**

Model(s)	32-bit Support	64-bit Support	Comments
Series 700: PA-7xxx	Yes		All 712, 715/64/80/100/100XC, 725/100, 743, 744, 745, 748i, 748

Table 3-3   **Workstations**

Model(s)	32-bit Support	64-bit Support	Comments
B-Class: PA-7300LC	Yes		B132L, B132L+, B160L, B180L
B-Class: PA-8500 and forward		Yes	Bx000
B-Class: PA-8600		Yes	B2600
C-Class: PA-7*xxx*	Yes		C100, C110, C160L
C-Class: PA-8*xxx*	Yes	Yes	C160, C180, C180-XP,C200, C240, C360
C-Class: PA-8500 and forward		Yes	C3x00
C-Class: PA-8700		Yes	C3650, C3700
J-Class: PA-7*xxx*	Yes		J200, J210, J210XC
J-Class: PA-8000/8200	Yes	Yes	J280, J282, J2240
J-Class: PA-8500 and forward		Yes	J5x00, J6000, J7000
J-Class: PA-8700		Yes	J6700

The following servers, workstations, and graphics adapters are no longer supported:

- Servers: E-, F-, G-, H-, and I-Class

- Workstations: 705, 710, 715/33, 715/50, 715/75, 720, 725/50, 725/75, 730, 735, 750, 755

- Graphics adapters: GRX, CRX, CRX-24, CRX-48Z

Hardware Enablement Patch Bundle (HWEnable11i)

The HWE Patch Bundle (`HWEnable11i`) delivers support for new and future SPU, I/O, and mass-storage products. This bundle is labeled as required, ignited on all systems in manufacturing, and automatically loaded as part of any OE update. It is included on the March 2002 11i OE media and on the Support Plus CD.

Some of the bundles' patches include the following:

updated for March 2002

- Support for new servers:

  — rp2430 (`model` string: 9000/800/A400-6X)
  — rp2470 (`model` string: 9000/800/A500-6X or 9000/800/A500-7X)
  — rp7410 (`model` string: 9000/800/rp7410)

- Support for new I/O adapters:

  — A7784A PCI-Universal Audio Card
  — A7789A PCI Graphics Card for PA-8700 Workstations
  — A6684A HSC Tachlite FC Mass-Storage I/O Adapter for D- and R-Class servers (HP-UX 11.x boot only supported on D390/R390 with PDC version 41.35)
  — A6794A 10/100/1000Base-T Product for HP rp7410 server
  — A6828A PCI Single Channel Ultra160 SCSI Adapter[1]
  — A6829A PCI Dual Channel Ultra160 SCSI Adapter[1]

- Support for new peripherals:

  — A6188A VA7100 Disk Array
  — A6189A VA7405 Disk Array
  — A6189B (2Gb-FC) VA7405 Array
  — A6218A (2Gb-FC) VA7410 Array

updated for December 2001

- Support for new servers:

  — rp2400 (`model` string: 9000/800/A400-6X)
  — rp2450 (`model` string: 9000/800/A500-6X or 9000/800/A500-7X)
  — rp5400 (`model` string: 9000/800/L1000)
  — rp5430 (`model` string: 9000/800/L1500-7X or 9000/800/L1500-8X)
  — rp5450 (`model` string: 9000/800/L2000)
  — rp5470 (`model` string: 9000/800/L3000-7X or 9000/800/L3000-8X)
  — rp7400 (`model` string: 9000/800/N4000-7X or 9000/800/N4000-8X)
  — rp8400 (`model` string: 9000/800/S16K-A)

- Support for new workstations:

  — C3650
  — C3700
  — J6700

---

1. See "SCSI Driver c8xx" on page 50 for further information about supported systems and boot support.

- Support for Superdome IO Expansion (IOX) cabinets

- Support for all PA 8700-based products

- Support for changes of SAM and `fruled` commands to modify the behavior of PCI Attention light (LED), so that it conforms to PCI Hotplug Specification (See "New Attention Indicator Behavior" on page 40 in Chapter 2.)

- Support for modified DIMM labels so that user interface to `parstatus` command more accurately matches HW labeling

- Support for 2D graphics on L-class systems

- Support for add-on workstation audio card/driver for B-Class

- Improved USB driver supports Update-UX tool for migration from HP-UX 10.20 to 11i

- The `ups_mond` software now supports longer shutdown values for PowerTrust II UPS units

- Added Ultrium tape drive support to `fbackup` and `frecover` commands

The `HWEnable11i` bundle also includes support for the following peripherals:

- A4982B PCI Visualize-fxe Graphics Adapter

- A6077A PCI 128 Audio Card for B2600 workstations

- A6795A PCI Tachlite 2Gb FC Mass-Storage I/O Adapters
  (HP-UX 11.*x* boot supported with PDC version 41.46 on L3000, N4000 servers)

- A6685A HSC Tachlite FC Mass-Storage I/O Adapter for Kx60, Kx70, Kx80 servers
  (HP-UX 11.*x* boot supported with PDC version 41.33/41.34)

- A6188A VA7100 Disk Array

- A6189A VA7405 Disk Array

- A6092A HyperFabric Adapter (PCI 4X)

- A6386A PCI HyperFabric II Adapter

HP-UX V-Class Changes

Single-Bit Memory Error Handling Enhancement

**new at 11i
original release**

Single-bit memory errors are now handled exclusively by `memlogd`. This allows the system to remove lockable pages that experience repeated single-bit memory errors. At boot time, the system uses the Page Deallocation Table to remove these pages dynamically from the kernel's list of free pages.

Single-bit memory error logging information can be viewed using the Support Tools Manager (STM). This information can no longer be found in `/var/adm/syslog/syslog.log`.

SCSI Drivers scsi3 and c720

updated for September 2001

The table below lists information about the SCSI Host Bus Adapters (HBAs) that the `scsi3` and `c720` SCSI drivers support for HP-UX 11i. Both the `scsi3` and `c720` drivers are included in the base 11i Operating Environment.

For the most recent information about the SCSI HBAs, you should periodically check the SCSI HBA information at the HP IT Resource Center (ITRC), at `http://itrc.hp.com`. Go to the area for knowledge trees. Then, under "Hardware," select "Networks," then "mass storage connectivity," then "scsi hba's," and then "technical support." This takes you to a page titled "scsi hba's." (You might want to bookmark this page for future reference.)

Table 3-4   Supported HBAs for HP-UX 11i

SCSI HBA	Bus Type	Technology	Supported HP-UX Systems
28655A (single channel)	HP-PB (NIO)	SE SCSI-2	K- and T-Class servers
28696A (single channel)	HP-PB (NIO)	FWD SCSI-2 (HVD)	K- and T-Class servers
A2679A (single channel)	EISA	SE SCSI-2	D- and R-Class servers (32-bit only) 715/64/80/100, 725/100, B132L+, B180L, C200, C240, J200, J210, and J210XC workstations
A2969A (single channel)	HSC	FWD SCSI-2 (HVD)	K-Class servers
A3644A (single channel)	HSC	FWD SCSI-2 (HVD)	T600 servers
A4107A (single channel)	HSC (EISA FF)	FWD SCSI-2 (HVD)	A180, A180C, and D-, and R-Class servers 715/64/80/100, 725/100, B132L, B132L+, B160L, B180L, C100, C110, C160, C160L, C180, C180-XP, C200, C240, J200, J210, J210XC, J280, H282, and J2240 workstations
A4800A (single channel)	PCI	FWD SCSI-2 (HVD)	A-, L-, N-, and V-Class, rp24xx, rp54xx, rp74xx, rp8400, and Superdome servers
A5149A (single channel)	PCI	Ultra2 LVD/SE SCSI-3	A-, L-, and N-Class, V2500, V2600, rp24xx, rp54xx, rp74xx, rp8400, and Superdome servers
A5150A (dual channel)	PCI	Ultra2 LVD/SE SCSI-3	A400, A500, L- and N-Class, and rp24xx, rp54xx, rp74xx
A5159A (dual channel)	PCI	FWD SCSI-3 (HVD)	A-, L-, and N-Class, rp24xx, rp54xx, rp74xx, rp8400, and Superdome servers

SCSI Driver c8xx

**new for
March 2002**

The `c8xx` driver is being updated to support two new SCSI PCI Ultra160 Host Bus Adapters (HBAs) that are available beginning in March 2002. The driver is included in the March 2002 `HWEnable11i` bundle, which is available on the March 2002 11i OE media and on the Support Plus CD.

Four patches (`PHKL_25712`, `PHCO_25831`, `PHKL_24854`, and `PHKL_25218`) are required for the `c8xx` driver. Note that all four patches are included in the `HWEnable11i` bundle. So, if you install the entire bundle, you will automatically install the four patches.

Table 3-5   **HBAs Supported by c8xx SCSI Driver for HP-UX 11i**

SCSI HBA	Bus Type	Technology	Supported HP-UX Systems
A6828A (single channel)	PCI	Ultra160 LVD/SE SCSI-3	A400, A500, and L- and N-Class, rp24xx, rp54xx, rp74xx, rp8400, and Superdome[a] servers
A6829A (dual channel)	PCI	Ultra160 LVD/SE SCSI-3	A400, A500, and L- and N-Class, rp24xx, rp54xx, rp74xx, rp8400, and Superdome[a] servers

a. NOTE: Check the HP IT Resource Cenger (ITRC) periodically for the availability of OLAR support on Superdome servers. Information for accessing the ITRC is given below.

For the most recent information about the SCSI HBAs, you should periodically check the SCSI HBA information at the HP IT Resource Center (ITRC), at `http://itrc.hp.com`. Go to the area for knowledge trees. Then, under "Hardware," select "Networks," then "mass storage connectivity," then "scsi hba's," and then "technical support." This takes you to a page titled "scsi hba's." (You might want to bookmark this page for future reference.)

Service Processor (GSP or MP)

NOTE   The service processor in HP servers is sometimes called the Management Processor (MP) and sometimes the Guardian Service Processor (GSP).

Regardless of the name, the service processor in these servers provides approximately the same features and performs essentially the same role.

Throughout this document the term "service processor" refers to both the MP and GSP service processors.

The service processor is a service and console subsystem on the following servers:

- N4000
- all L-Class
- A-Class (the A400 and A500)
- Superdome systems
- rp2400, rp2450
- rp5400, rp5430, rp5450, rp5470
- rp7400, rp7410
- rp8400
- all new servers introduced starting with the N-Class

The GSP console driver, the software component of the GSP, provides the following features on HP-UX:

Local console port   provides system console while HP-UX is running.

Remote session port   establishes an HP-UX login session on the remote console.

Local session port   establishes an HP-UX login session on the local console.

Internal console port   supports firmware upgrade and diagnostics on GSP.

UPS port   establishes a communication channel between the UPS daemon and UPS.

SAM provides configuration support (that is, modem and UPS connections) over the GSP serial ports. The *insf* (1M) and *mksf* (1M) commands create device files for the GSP serial ports.

The following commands have been changed to provide additional support for the GSP console:

- `ttytype` can determine the ID of the terminal connected to the local console port.

- `stty` supports the status query and reset function of the GSP.

The GSP console driver is based on the existing built-in serial port driver (`asio0`). Every serial port on the GSP adheres strictly to the `termio` feature set; these features are described in the *termio* (7) and *modem* (7) manpages.

GSP Logging Capabilities

The introduction of GSP to the above platforms dramatically changes the way chassis operations and diagnostic evaluations are performed on a running system.

The new subsystem requires HP-UX to provide more information than was provided on previous platforms. HP-UX will continue to output the same chassis-codes and forward-progress indicators that have been provided in previous releases. On the above and subsequent systems, however, the codes are displayed on the Virtual Front Panel (VFP) of the system. Most of the existing four-hex digit chassis codes are enclosed in GSP-specific encoding.

The GSP subsystem interprets various forms of logging information from both firmware and software. Several new software events are now logged, including:

- "Boot Complete" indicator

- Timestamp

- Periodic heartbeat, with:

    — timeout value (a time-limit within which another event must be logged before the system is declared "dead")

    — activity level indicating system usage

- Minimal LED control

In addition to existing four-hex digit chassis codes, the following information is sent with each event:

- Alert level

- CPU number

NOTE

The GSP will not store codes of alert level 0 after PDC's "boot complete" code. All incoming codes will display on the VFP, but level 0's will not be stored for later retrieval. This is so the log won't fill up with heartbeat entries.

PDC_CHASSIS, the old firmware call for old-style, four-hex digit chassis codes, always produces codes of alert level 0. In order to create new-style chassis codes, the PAT_ call for CHASSIS must be used.

N4000 and rp7400 Server Functionality

This section describes 11i functionality to enable HP 9000 model's HP N4000 and rp7400 mid-range servers. Related operating system changes can be found in the following sections of this document:

- "Changes to System Administration Manager (SAM)" on page 133
- "Improved ioscan Description Field for PCI Devices" on page 140
- "ttytype Support for the N4000 and rp7400 Console" on page 57
- "New stty Options" on page 57
- "SCSI Drivers scsi3 and c720" on page 49
- "Service Processor (GSP or MP)" on page 51

With the exception of some new system build options, changes to HP-UX 11i for these servers will have little, if any, bearing on customers using legacy PA-RISC systems.

NOTE	For the purpose of this document, all systems prior to N4000 are termed "legacy" (including B-, C-, and D-Class low-end systems, K-Class mid-range systems, and T- and V-Class high-end systems).

Platform Infrastructure

The HP N4000 and rp7400 servers are the first HP systems based on PA-RISC processors with IA-64 Core Electronic Complex (CEC) components. This "hybrid" system contains a new modular platform infrastructure. Whether PA-based, IA-64-based, or hybrid, new kernel interfaces and platform modules are being provided to support all current platforms.

Subsequent sections describe the following new platform architecture components:

- Platform Support Modules (PSM)

  The PSMs control specific hardware or the functions of a given platform. PSMs designed for the new functionality include the following:

  — PAT PSM
  — SBA PSM
  — SAPIC PSM

- Context Dependent I/O module (CDIO)

  Because of the hierarchal dependency requirements of some platform modules, not all new platform code is handled by PSMs. The following CDIOs are included in HP-UX:

  — CB CDIO
  — LBA CDIO
  — PCI CDIO
  — PCItoPCI CDIO

Impact on Legacy Systems

Legacy system users will see minimal impact in their applications or system administration tools due to the changes in the platform infrastructure.

Although the configuration files on 64-bit systems (for example, `/stand/system` and `master.d/core-hpux`) and SAM will refer to CB-CDIO, PSMs and new CDIOs now included in the system, these components may coexist in the configuration files and be loaded into the kernel at the same time, even if they are inactive on a particular platform. Run-time checks evaluate which components are activated.

For legacy systems, end users might see new entries (`sapic`, `lba`, and `sba`) in the `/stand/system` file. In addition, some new lines have been added to CDIO and `DRIVER_DEPENDENCY` tables of the `/usr/conf/master.d/core-hpux` file to include the new central bus (`cb`) and the various new PSMs (for example, `pa_psm` or `pa_generic_psm`).

Configuration Changes

N4000 and rp7400 users must have the following modules in the kernel (via the master file entries) for the PAT, SBA, and Lower Bus Adapter (LBA) components to be detected and properly configured—and for the HP-UX kernel to boot. Without these modules, the HP-UX kernel will be unable to detect the hardware CEC components on a N4000 pr rp7400 system and the kernel will not boot.

The master file, `/usr/conf/master.d/core-hpux`, contains the following entries for all systems.

$CDIO Table:

```
cb 1
lba 0
PCItoPCI 0
pa_generic_psm 0
pa_psm 0
pat_psm
sapic 0
sba 0
```

$DRIVER_DEPENDENCY table:

```
core   pa_psm   pa_generic_psm   asp lasi
sio    pa_psm   pa_generic_psm
wsio   pat_psm core DlkmDrv
lba    pci sapic PCItoPCI
GSCtoPCI pci PCItoPCI
```

The `/stand/system` file contains the following entries:

```
*************************************************
* Bus-Dependent subsystems
*************************************************
* lba
*************************************************
* PSMs
*************************************************
* sapic
* sba
```

PAT PSM

This software module interacts with N-Class and rp7400 firmware to discover and keep track of the CEC components configured on the N4000 and rp7400. The PAT PSM also provides access to platform-specific hardware components at runtime.

Although it may be included and linked into all 64-bit kernels, the PAT PSM is useful only to N4000 and rp7500 systems. As of HP-UX 11.0 Extension Pack, May 1999, a run-time test determines whether the linked-in PAT PSM is installed on the system.

Since PAT functionality is only supported on 64-bit systems, 32-bit kernels do not have the PAT PSM built into them.

SBA PSM

The SBA PSM detects and configures the system bus adapter hardware and translates addresses between the Merced bus and the underlying LBA.

The SBA PSM supports system bus adapters on all N4000 and rp7400 systems, and is active and visible to N4000 and rp7400 users.

SAPIC PSM

The SAPIC PSM manages line-based interrupts. This configurable software module handles interrupts routed through the I/O SAPICs.

The SAPIC PSM conforms to the Central Bus CDIO platform infrastructure. It maintains the SAPIC redirection table.

CB CDIO

The CB CDIO contains interfaces that isolate platform-specific code from the rest of the kernel. These interfaces allow generic access to the platforms, regardless of which platform-specific PSMs are active in the kernel. The Central Bus framework interconnects the different PSMs that control the hardware.

For backward compatibility, the PA-CDIO has been restructured into a PA-generic PSM and PA-legacy PSM.

LBA CDIO

The LBA CDIO provides bus translation for all activity between the System Bus Adapter and the PCI bus. The LBA CDIO is the hardware-enabling HP-UX kernel module that controls the lower bus adapter and, therefore, all the intricacies of the dependent hardware. This CDIO also resolves any overlapping configuration issues with LBA, and interacts directly with the PCI CDIO.

PCI CDIO

The PCI subsystem has been redesigned to support PCI Card Online Addition and Replacement (OLAR) and to support a new interrupt line routing architecture.

On legacy systems (B-, C-, and V-Class), platform firmware had complete responsibility for configuring all devices. In contrast, the PCI CDIO detects unconfigured PCI devices and programs the base address registers in order to support PCI Card OLAR.

On N4000 and rp7400 systems, the firmware programs only the boot and console devices. The PCI CDIO programs the remaining devices, using information provided by firmware to the operating system (PAT PSM gets this for PCI).

The N4000 and rp7400 disassociates interrupt routing/handling from the platform-specific bus adapter. On legacy PCI systems, the interrupt lines are routed to the PCI host bus-adapter chip and handled by the same driver (for example, GSCtoPCI and EPIC CDIOs). On N4000 and rp7400 systems, though the interrupt lines are routed to the LBA (PCI bus interface chip), SAPIC PSM handles the interrupt support instead of the LBA CDIO.

LBA CDIO provides N4000 and rp7400 specific services to support PCI drivers and access to the PCI bus. Legacy PCI bus adapter drivers have been modified to be compatible with the new PCI CDIO.

PCItoPCI CDIO

The restructuring of the PCI subsystem permits PCItoPCI configuration of devices to more than two bridges deep. There is no new functionality for this release.

ttytype Support for the N4000 and rp7400 Console

**new at 11i
original release**

Although the `ttytype` command has been enhanced to support the N4000 and rp7400 console, there are no user-visible changes in the behavior of the command.

However, a new `ioctl()` call has been added to the command to query the Guardian Service Processor (GSP) console driver for the TERM identity. If the `ioctl()` call fails, `ttytype` will continue with the existing terminal identification process.

For more information on the GSP, see "Service Processor (GSP or MP)" earlier in this chapter. For information on the ttytype command, see the *ttytype* (1) manpage.

New stty Options

**new at 11i
original release**

Two new options have been added to the `stty` command to support the console on the following systems:

- N4000
- all L-Class
- all N-Class
- rp2400, rp2450
- rp5400, rp5430, rp5450, rp5470
- rp7400

`+queryGSP`   queries the status of the GSP (Guardian Service Processor)

`+resetGSP`   resets the GSP of the console

Typically, you might use `+queryGSP` if you are getting no response at the console or `+resetGSP` if the console locks up. Here is an example of the latter, which runs the command elsewhere from the console but directs the command at the console device:

stty +resetGSP < /dev/GSPdiag1

Note that these options require superuser status.

For information on the stty command, see the *stty* (1) manpage.

Workstations

Workstation Graphics Support

Graphics Software Support

The OpenGL, Starbase, HP PEXlib, and HP-PHIGS 3D APIs are fully supported on HP 9000 PA-RISC workstations and selected servers. HP's workstation graphics software support includes the run-time and programming environment packages for the 3D graphics APIs named above, plus additional software for technical computing environments.

The Graphics and Technical Computing Software is supported on all PA-RISC workstations and servers. However, some functionality (notably accelerated 3D graphics) requires special hardware.

Graphics Hardware Support

In addition to the many existing graphics cards, HP-UX 11i now supports the HP VISUALIZE-fxe card. This new entry-level, low cost, full-featured 3D graphics card replaces the VISUALIZE-fx^2 card for 3D applications and the VISUALIZE-EG card for 2D applications.

HP VISUALIZE-fxe provides 3D support for OpenGL, Starbase, HP PEXlib and HP-PHIGS 3D APIs, with a full VISUALIZE-fx^2-like feature set. It also provides 2D features via the X server and Xlib comparable to those of the VISUALIZE-fx^2 and VISUALIZE-EG products.

HP VISUALIZE-fxe is supported on these systems:

*   B180L (2D X libraries only; 3D supported only via VMX/VMD)
*   B1000, B2000
*   C3000, C3600
*   J5000, J5600, J6000, J7000

HP VISUALIZE-fxe is not supported on the C360 workstation.

For a complete list of supported systems and graphics combinations on HP-UX 11i, consult `http://www.hp.com/workstations`.

Workstation Firmware Requirements

updated for March 2002

PHSS_25983 was released on January 11, 2002 to correctly update workstation firmware when using an Ignite-UX server. To get the patch (and for further instructions), go to the HP IT Resource Center (ITRC), at `http://itrc.hp.com`.

updated for September 2001

Note that PHSS_24106 will automatically update PDC firmware on your workstation. This HP-UX patch is included in the HWEnable11i bundle, and is always installed on the workstation during a cold-install of any 11i OE.

**updated for
June 2001**

Most workstations supporting 11.0 do not require a firmware upgrade to run HP-UX 11i. For 64-bit operation, however, please note the following recommendations:

System	Minimum Firmware Revision	Latest Firmware Revision
B1000, B2000	2.3	5.0
C200, C240, C200, C240	4.3	6.3
C360	1.0	1.5
C3000, C3600	2.3	4.9
J280, J282	2.1	2.4
J2240	1.2	2.1
J5000, J5600, J6000, J7000	2.3	4.9

B1000, C3000 and J5000 systems manufactured before September 1999 require a firmware upgrade *prior* to updating to HP-UX 11.*x* or 11i. For more detailed information, see *HP-UX 11i Installation and Update Guide*.

Workstation Tuned Kernel Parameters

**updated for
June 2001**

As of the June 2001 release, kernel parameters for CAE and EE Engineering workstation kernels will be optimized during the installation or update. If the system is installed or updated using Ignite-UX (IUX), this occurs automatically. If you install or update manually, the optimization only occurs after you select one of the new engineering workstation kernel sets via SAM.

On factory Instant Ignition, IUX will install workstation systems with optimized default kernel parameter settings as long as the system has at least 64MB of RAM. The new defaults are optimized for general performance and are tailored appropriately for a 32-bit or 64-bit kernel. (The default parameter settings for each set are listed at the end of this section.) The larger `maxdsiz` limit for 64-bit installations now allows users to take advantage of the increased (approximately 3GB) process data space available with the June 2001 release.

Also on factory Instant Ignition, IUX automatically configures the kernel with the appropriate new CAE kernel parameter set:

- CAE/ME/General Eng. Workstation 64-bit Kernel, or
- CAE/ME/General Eng. Workstation 32-bit Kernel

Via SAM, you can apply tuned kernel parameter settings by selecting one of these new sets:

- CAE/ME/General Eng. Workstation 64-bit Kernel
- CAE/ME/General Eng. Workstation 32-bit Kernel
- EE Engineering Workstation 64-bit Kernel

- EE Engineering Workstation 32-bit Kernel

The 64-bit versions of these parameter sets configure the kernel to use the increased process data space. The CAE/ME/General Engineering Workstation sets are for general workstation use, which includes running typical MDA applications. The EE Engineering Workstation sets are for compute-intensive applications that do not perform large amounts of disk I/O. Many EDA applications fall into this category. Be sure to select the 64-bit or 32-bit versions depending on the "bitness" of your installed kernel.

CAE/ME/General Eng. Workstation 64-bit Kernel Parameter Defaults

```
maxusers        128
maxfiles        200
maxfiles_lim    2048
maxdsiz         0xC0000000
maxdsiz_64bit   0x400000000
maxtsiz         0x40000000
maxtsiz_64bit   0x100000000
maxssiz         0x04FB3000
maxssiz_64bit   0x10000000
shmmax          0x40000000
ninode          4000
maxuprc         256
npty            200
nstrpty         200
maxswapchunks   4096
create_fastlinks 1
fs_async        1
```

CAE/ME/General Eng. Workstation 32-bit Kernel Parameter Defaults

```
maxusers        128
maxfiles        200
maxfiles_lim    2048
maxdsiz         0x7b03a000
maxtsiz         0x40000000
maxssiz         0x04FB3000
shmmax          0x40000000
ninode          4000
maxuprc         256
npty            200
nstrpty         200
maxswapchunks   4096
create_fastlinks 1
fs_async        1
```

EE Engineering Workstation 64-bit Kernel Parameter Defaults

```
maxusers        128
maxfiles        200
maxfiles_lim    2048
maxdsiz         0xC0000000
maxdsiz_64bit   0x400000000
maxtsiz         0x40000000
maxtsiz_64bit   0x100000000
maxssiz         0x04FB3000
maxssiz_64bit   0x10000000
shmmax          0x40000000
ninode          4000
maxuprc         256
npty            200
nstrpty         200
maxswapchunks   4096
```

```
create_fastlinks 1
fs_async        1
vps_ceiling     64
dbc_max_pct     15
dbc_min_pct     15
```

EE Engineering Workstation 32-bit Kernel Parameter Defaults

```
maxusers        128
maxfiles        200
maxfiles_lim    2048
maxdsiz         0x7b03a000
maxtsiz         0x40000000
maxssiz         0x04FB3000
shmmax          0x40000000
ninode          4000
maxuprc         256
npty            200
nstrpty         200
maxswapchunks   4096
create_fastlinks 1
fs_async        1
vps_ceiling     64
dbc_max_pct     15
dbc_min_pct     15
```

X Window System (X11 R6) Run-Time Libraries on Workstations

The HP-UX 11i provides workstation support for the 64-bit X Window System shared library (stack).

The following X and Motif libraries are available in 64-bits:

```
libMrm.a
libXm.4
libICE.2
libSM.2
libX11.3
libXIE.2
libXext.3
libXhp11.3
libXi.3
libXp.2
libXmu
libXaw
```

To date, these libraries are only found in release 6 of the X libs (X11 R6) and Motif version 2.1. No 64-bit versions of the `tooltalk` libraries, `libtt`, or `libDtSvc` are available.

The 64-bit X Window System (X11 R6) run-time libraries are usable only on systems that support the 64-bit operating system. To use the 64-bit run-time libraries, you must specify that the application will run (compile) in 64-bit mode. The 64-bit libraries are then used automatically.

4   HP-UX 11i Operating Environment Applications

The HP-UX 11i Operating Environments

**new at 11i
original release**

The HP-UX 11i Operating Environments consist of the HP-UX operating system, a set of "always installed" applications (the "environment"), and a collection of additional applications and drivers that you can selectively install. The HP-UX 11i release is available in one of the following Operating Environments (OEs):

- HP-UX 11i **OE** (see page 67)
- HP-UX 11i **Enterprise** OE (see page 80)
- HP-UX 11i **Mission Critical** OE (see page 85)
- HP-UX 11i **Minimal Technical** OE (see page 91)
- HP-UX 11i **Technical Computing** OE (see page 92)

You can choose the HP-UX 11i Operating Environment that is best suited for your computing environment. Although there are five OEs available, only one can be installed and operate on your HP 9000 server or workstation. The application contents of the HP-UX 11i Operating Environments are shown below:

Table 4-1   Operating Environment Contents

Application	HP-UX 11i OE (commercial servers)	HP-UX 11i Enterprise OE (commercial servers)	HP-UX 11i Mission Critical OE (commercial servers)	HP-UX 11i Minimal Technical OE (workstations)	HP-UX 11i Technical Computing OE (technical servers & workstations)
Always-Installed Network Drivers (see page 67)	YES	YES	YES	YES	YES
Base VERITAS Volume Manager (VxVM) (see page 68)	YES	YES	YES	YES	YES
CIFS/9000 Client and CIFS/9000 Server (see page 69)	YES	YES	YES	NO	YES
Enterprise Cluster Master (ECM) Toolkit (see page 85)	NO	NO	YES	NO	NO
Event Monitoring Service (EMS) (see page 70)	YES	YES	YES	YES	YES
GlancePlus Pak (see page 80)	NO	YES	YES	NO	NO
High Availability Monitors (see page 80)	NO	YES	YES	NO	NO

Table 4-1   **Operating Environment Contents**

Application	HP-UX 11i OE (commercial servers)	HP-UX 11i Enterprise OE (commercial servers)	HP-UX 11i Mission Critical OE (commercial servers)	HP-UX 11i Minimal Technical OE (workstations)	HP-UX 11i Technical Computing OE (technical servers & workstations)
High Performance Math Libraries (HP MLIB) (see page 92)	NO	NO	NO	NO	YES
HP 3D Technology for the Java Platform (see page 94)	NO	NO	NO	NO	YES
HP Apache-based Web Server for HP-UX (see page 70)	YES	YES	YES	NO	YES
HP Message-Passing Interface (MPI) (see page 94)	NO	NO	NO	NO	YES
HP OnLineJFS 3.3 (see page 81)	NO	YES	YES	NO	NO
HP Visualize Conference (see page 96)	NO	NO	NO	NO	YES
HP-UX Runtime Environment for the Java 2 Platform (see page 73)	YES	YES	YES	NO	YES
HP-UX Support Tools: STM, ODE, & EMS Hardware Monitors (see page 73)	YES	YES	YES	YES	YES
HP-UX Workload Manager (see page 86)	NO	NO	YES	NO	NO
HP-UX Workload Manager Oracle , Database Toolkit (see page 87)	NO	NO	YES	NO	NO
MC/ServiceGuard (see page 88)	NO	NO	YES	NO	NO
MC/ServiceGuard NFS Toolkit (see page 89)	NO	NO	YES	NO	NO
MirrorDisk/UX (see page 81)	NO	YES	YES	NO	NO
Netscape Communicator (see page 75)	YES	YES	YES	YES	YES

Table 4-1 **Operating Environment Contents**

Application	HP-UX 11i OE (commercial servers)	HP-UX 11i Enterprise OE (commercial servers)	HP-UX 11i Mission Critical OE (commercial servers)	HP-UX 11i Minimal Technical OE (workstations)	HP-UX 11i Technical Computing OE (technical servers & workstations)
OpenGL 3D Graphics Developers Kit and Runtime Environment (see page 91)	NO	NO	NO	YES	YES
Perl Programming Language (see page 75)	YES	YES	YES	YES	YES
Pluggable Authentication Modules (PAM) Kerberos (see page 75)	YES	YES	YES	NO	YES
Plug-In for the Java 2 Platform for Netscape Communicator (see page 77)	YES	YES	YES	NO	YES
Process Resource Manager (PRM) (see page 82)	NO	YES	YES	NO	NO
Servicecontrol Manager (SCM) (see page 77)	YES	YES	YES	NO	NO

NOTE   The iCOD product was removed as an "always installed" component from the HP-UX Operating Environments in the June 2001 release. You can find this product on the Support Plus media and at the following Web site:

`http://software.hp.com.`

The following applications are selectable and not automatically installed with the Operating Environments. They are available as noted in the following table:

Table 4-2   **Selectable Applications**

Application	HP-UX 11i OE (commercial servers)	HP-UX 11i Enterprise OE (commercial servers)	HP-UX 11i Mission Critical OE (commercial servers)	HP-UX 11i Minimal Technical Computing OE (workstations)	HP-UX 11i Technical Computing OE (technical servers & workstations)
HP Intrusion Detection System/9000 (IDS/9000) (see page 97)	YES	YES	YES	NO	NO
Ignite-UX (IUX) (see page 98)	YES	YES	YES	YES	YES
IPFilter/9000 (see page 101)	YES	YES	YES	NO	NO
Java Out-of-Box (JAVAOOB) (see page 102)	YES	YES	YES	YES	YES
Netscape Directory Server (J4258BA) (see page 103)	YES	YES	YES	YES	YES
Selectable Network Drivers (see page 104)	YES	YES	YES	YES	YES
WebQoS Peak on HP-UX Packaged Edition (J4274AA) (see page 104)	YES	YES	YES	NO	NO

HP-UX 11i Operating Environment

**new at 11i
original release**

The HP-UX 11i Operating Environment is the standard OE from which the Enterprise OE and Mission Critical OE have been derived by adding appropriate applications. The HP-UX 11i OE includes the base 32/64-bit HP-UX Operating System, network drivers, other always-installed functionality, and the following applications:

Always-Installed Applications

- Always-Installed Network Drivers (see page 67)
- Base VERITAS Volume Manager (VxVM) (see page 68)
- CIFS/9000 Client and CIFS/9000 Server (see page 69)
- Event Monitoring Service (EMS) (see page 70)
- HP Apache-based Web Server for HP-UX (see page 70)
- HP-UX Runtime Environment for the Java 2 Platform (see page 73)
- HP-UX Support Tools: STM, ODE, & EMS Hardware Monitors (see page 73)
- Netscape Communicator (see page 75)
- Perl Programming Language (see page 75)
- Pluggable Authentication Modules (PAM) Kerberos (see page 75)
- Plug-In for the Java 2 Platform for Netscape Communicator (see page 77)
- Servicecontrol Manager (SCM) (see page 77)

Selectable Applications

- HP Intrusion Detection System/9000 (IDS/9000) (see page 97)
- Ignite-UX (IUX) (see page 98)
- IPFilter/9000 (see page 101)
- Java Out-of-Box (JAVAOOB) (see page 102)
- Netscape Directory Server (J4258BA) (see page 103)
- Selectable Network Drivers (see page 104)
- WebQoS Peak on HP-UX Packaged Edition (J4274AA) (see page 104) [1]

Always-Installed Network Drivers

The following network drivers are installed by default in the Operating Environment:

Table 4-3   **Always-Installed I/O Drivers on 11i OE Media**

I/O Driver	Description	I/O Adapter Product Number
FDDI-00	PCI FDDI PCI FDDI	A3739A (V, L, N, A180) A3739B
FibrChanl-00	PCI FibreChannel HSC FibreChannel	A5158A A6684A, A6685A (D, R, A180, K)
GigEther-00	PCI GigEther HSC GigEther	A4926A, A4929A A4924A, A4925A

---

1. As of December 2001, WebQoS Peak has been slated for future obsolescence.

Table 4-3 **Always-Installed I/O Drivers on 11i OE Media**

I/O Driver	Description	I/O Adapter Product Number
GigEther-01	PCI/PCI-X GigEther	A6794A
RAID-00	PCI RAID	A5856A (L, N, A4, A5)

Base VERITAS Volume Manager (VxVM)

updated for December 2001

The Base VERITAS Volume Manager 3.2 for HP-UX (`swlist` version B.03.20) is a new version of the VERITAS Volume Manager and Cluster Volume Manager with enhanced features that are currently only available on Solaris, but will soon be available on Linux and AIX.

Similar to the Base VERITAS Volume Manager 3.1 for HP-UX, this new version includes the following add-on products:

- B9116AA: VERITAS Volume Manager 3.2 for HP-UX provides a full set of enhanced volume manager capabilities, including mirroring, RAID-5, and DMP for active/active devices.

- B9117AA: VERITAS Volume Manager 3.2 Cluster Volume Manager for HP-UX provides enhanced volume manager functionality for clustered environments, and is integrated with MC/ServiceGuard 11.13 and ServiceGuard OPS Edition 11.13. (The Cluster Volume Manager for HP-UX requires that the VERITAS Volume Manager be installed first.)

All the software for both the base and add-on products is included in the Base VxVM product, but the enhanced features are unavailable without a license. When you purchase an add-on product, you will get a license that enables the enhanced features.

What's Changed

The Base VERITAS Volume Manager 3.2 for HP-UX will be installed by default on all of the Operating Environments in the December 2001 release. Customers who purchase one of these OEs will no longer need to install Base VxVM. However, none of the VxVM add-on products are available as part of any OEs.

The VERITAS Volume Manager 3.2 includes the following new features:

- enhancements to the Dynamic MultiPathing (DMP) feature via a new Device Discovery Layer, which makes it easier to add new array support without requiring a kernel rebuild or a system reboot

- 4-node CVM support for MC/ServiceGuard and ServiceGuard OPS Edition

- CVM and shared disk group support for striped mirrors, online relayout, and Oracle resilvering

- full support for HP-UX workstations

- performance improvements

Compatibility

Disk groups created with the Base VERITAS Volume Manager 3.1 for HP-UX product must be upgraded to a new disk group version supported with VERITAS Volume

Manager 3.2. This upgrade will be transparent to customers when updating to the December 2001 release. See the *vxdg* (1M) manual page for additional information about upgrading disk groups.

Documentation

All of the VERITAS Volume Manager 3.2 documents are available on the Instant Information CD, as well as on the `docs.hp.com` web site in both HTML and PDF formats.

The following general VERITAS Volume Manager documents include information about the VxVM products (base and add-ons):

- *VERITAS Volume Manager 3.2 Administrator's Guide*
- *VERITAS Volume Manager 3.2 Storage Manager Administrator's Guide*
- *VERITAS Volume Manager 3.2 Troubleshooting Guide*
- *VERITAS Volume Manager 3.2 Migration Guide*

For release notes on the Base VxVM 3.2 and add-on products for HP-UX 11i, see *VERITAS Volume Manager 3.2 for HP-UX Release Notes.*

A full set of manpages for both the base VxVM 3.2 and add-on products is included with the Operating Environments.

CIFS/9000 Client and CIFS/9000 Server

With CIFS/9000 Client and CIFS/9000 Server, Hewlett-Packard provides a Common Internet File System (CIFS), the Microsoft protocol for remote file access. CIFS is built into all recent Windows operating systems, including Windows 95, 98, NT 4.0, and 2000. By providing both server and client, CIFS/9000 enables file and print interoperability for environments with a mix of UNIX and Windows platforms.

For the most recent information on CIFS/9000 Client and CIFS/9000 Server, see the release notes and other documents listed below under "Documentation."

updated for September 2001

CIFS/9000 Server version A.01.07 incorporates Samba 2.09, allows Windows 2000 clients to modify POSIX ACLs, and includes an option that can translate open mode locks to HP-UX advisory locks. This version also contains three fixes, including one which allows you to view, and therefore delete, broken (symbolic) links.

CIFS/9000 Client version A.01.06 provides improved file attribute caching and updated PAM-NTLM troubleshooting information.

new at 11i original release

The following changes were made for HP-UX 11i:

- The CIFS/9000 product consists of only two products instead of four. The product numbers are B8724 and B8725.

- CIFS/9000 product documentation is provided as `.pdf` files located in the `/opt/samba/HP_docs` directory. (Product documentation is also available on the HP-UX 11i Instant Information CD and on the Web at: `http://docs.hp.com`.)

Documentation

Available documents include:

- *Installing and Administering the CIFS/9000 Server* (B8725-90016)
- *Installing and Administering the CIFS/9000 Client* (B8723-90011)
- *CIFS/9000 Server Release Note* (B8725-90017)

- *CIFS/9000 Client Release Note Version A.01.06* (B8724-90012)
- *CIFS/9000 Client Release Note Version A.01.07* (B8724-90013)

NOTE   The CIFS/9000 documentation files require 2MB of disk space.

Event Monitoring Service (EMS)

The Event Monitoring Service (EMS) version A.03.20.01 is a framework used to monitor various system resources. In addition to the basic monitoring framework, the EMS product includes a set of general monitors for basic network interfaces, system resources, and ServiceGuard cluster objects. EMS is being released for use with the HP-UX 11.0 and HP-UX 11i operating systems, and is included in the HP-UX 11i Operating Environment. This release has all the features found in earlier versions in addition to new functionality, defect repairs, and support for new hardware configurations.

The EMS version A.03.20.01 is a minor release, with minor changes and defect fixes. The contents of EMS releases A.03.00 through A.03.10 have been incorporated, together with all A.03.10 patches.

Size Requirements

- The disk space requirement is 2.75MB. An additional 13MB of disk space should be allocated in `/etc/opt` to support EMS logging facilities.

- The memory requirement is 3MB.

Impact

With HP-UX 11i, EMS adds a new state to the package monitor: `UNAVAIL`. If the monitor does not have sufficient information to determine status, the current value for the resource is set to `UNAVAIL`.

NOTE   The Event Monitoring Service version A.03.20.01 does not provide Native Language Support.

Documentation

In addition to the current user's manual (*Using the Event Monitoring Service* [B7612-90015]), consult the *Event Monitoring Service version A.03.20.01 Release Notes for HP-UX 11i* (B7609-90015) for further information. Both publications are available on the HP-UX 11i Instant Information CD and on the Web at:

`http://docs.hp.com/hpux/diag`

HP Apache-based Web Server for HP-UX

The HP Apache-based Web Server for HP-UX is an HTTP/1.1 compliant server that implements the latest protocols. The server includes software developed by the Apache Software Foundation for use in the Apache HTTP server project (see `http://www.apache.org`). The HP Apache-based Web Server can be customized by writing software modules using the Apache module API.

The HP-UX release of the HP Apache-based Web Server includes pre-compiled binaries that have been preconfigured to run on HP-UX 11.0 and later releases. It is supported on 32-bit and 64-bit systems. It runs as a 32-bit binary on 64-bit HP-UX 11.0 and 11i. It is *not* supported on HP-UX 10.20.

This version of the HP Apache-based Web Server includes 128-bit strong encryption. Apache SSL connections use a default dummy certificate provided by HP. See the release notes for the PEM pass phrase associated with the certificate.

NOTE   You must obtain certificates from authorized agencies for commercial purposes and make the necessary changes in `httpd.conf` to use the desired certificates.

---

updated for March 2002

The following new feature is included in version 1.3.19.23:

- With this release, the `chroot` feature of the HP Apache-based Web Server has been enhanced. Customers who are currently using `chroot` (or are planning to use it in the future) are strongly encouraged to upgrade to this release. Current `chroot` users should recreate their `chroot` environment by running `/opt/apache/bin/chroot_os_cp.sh`, making sure to back up their original `chroot` environment beforehand.

updated for December 2001

The following new features are included in version 1.3.19.21:

- Fix for the Automatic restart on reboot
- `Certmig` utility re-built on PA1.1
- Fix for support of C++ shared modules.

updated for September 2001

The following new features are included in version 1.3.19.20:

- mod_perl now supports Perl version 5.6.1 (included in the OE).
- Apache/Tomcat/Webmin can now be automatically restarted on reboot. More information on customizing and configuring this feature can be found at `/opt/apache/htdocs/doc/config.notes`.
- The security module `mod_ssl` updated to v.2.8.3.
- Tomcat now starts with a `nohup` command, so that even if the parent shell is terminated, Tomcat will not die. The output from Tomcat startup is kept in `/opt/tomcat/logs/tomcat_startup.log`.
- `Chroot` now causes the named directory to become the root directory, the starting point for path searches. A malicious user cannot get to the root file system. The HP Apache-based Web Server's `chroot` includes SSL enhancements. Under the `chroot` directory, there is a script for copying OS files.
- Now included is Memory Management (MM), a 2-layer abstraction library that simplifies the usage of shared memory between forked processes under UNIX platforms. MM support allows the use of `httpd.conf` SSLSessionCashe directives `shm:/opt/apache/logs/ssl_scache(512000)`.
- `certmig` now makes possible the sharing of certificates between the iPlanet Web Server 4.x and any server that supports PKCS#12 formats. The `certmig` utility is an extension of the `pk12util` utility, provided by the Mozilla community. In addition to

the `pk12util` functionality, `certmig` lists and extracts certificates from iPlanet certificate databases.

updated for June 2001

The following new features are included in this version:

- `mod_perl` is configured for Perl 5.005_03.

- Tomcat v.3.1.1 compliant with Java Servlets 2.2 and JavaServer Pages 1.1 and contains defect fixes.

- PHPv.4.0.4pl1, a popular, server-side, cross-platform, HTML embedded full-featured language with a Java/C++ syntax. Supports many databases.

- Support for loading customized C++ shared libraries.

- Third Party Support: BroadVision plug-in provides out-of-the-box support for the BroadVision e-commerce application suite.

- Certificate Migration Tool v.3.1.1 to migrate iPlanet certificates to Apache seamlessly. License and export information are included.

- Webmin v. 0.84, a Web-based administration and configuration tool from Webmin, enhanced to handle administration and configuration for the HP Apache-based Web Server.

- Includes `mod_ssl` v.2.7.1.

Installation Requirements

For HP-UX 11i, the HP Apache-based Web Server is perl-enabled but not configured; therefore, it will not allow the execution of perl scripts unless `mod_perl` is configured. For detailed configuration and installation instructions, see the release notes at **/opt/apache/htdocs/doc/apache.release.notes** or once the Web Server has been started, at **http://*machinename*/doc**.

CAUTION

If you are receiving the HP Apache-based Web Server as part of the HP-UX 11i Operating Environment (OE), the software will be installed automatically as part of the OE bundle *unless there is a non-HP version of Apache already on the system*. If there is a non-HP version of Apache already on your system, the HP Apache-based Web Server will *not* install. In this case, you will need to install Apache separately from the rest of the OE bundle, as described in the following section.

Installing Apache Separately

If you are installing the HP Apache-based Web Server for HP-UX separately from the rest of the OE bundle, or if you obtained the Apache product (B9415AA) independently of the HP-UX 11i Operating Environment, follow these steps:

1. From CD2 of the HP-UX 11i Operating Environment CDs, run **/usr/sbin/swinstall&**.

2. Select the appropriate depot.

3. From the `View` menu, go to `Change Software View` and select `Start With Products`. The products that are available will display for your selection.

4. Select `Apache`.

5. Go to the `Actions` menu and select `Install`. The installation paths are
   `/opt/apache` and `/opt/tomcat`.

NOTE   By default, `swinstall` does not reinstall filesets if the same revision already exists on
your system. If you want to reinstall the same revision (for example, if some files are
lost), you can change the installation options by choosing `Options/Change Option`.

Installing a product or a fileset may automatically install dependent filesets necessary to
run the selected items.

If an HP or non-HP version of Apache is already on the system, `swinstall` preserves the
existing configuration files under `/opt/apache/conf`, `/opt/apache/conf/jserv`, and
`/opt/tomcat/conf` by renaming `<file>` to `<file>.save`. It also preserves certificates
and certificate-related files under `/opt/apache/conf/ssl.*` directories by renaming
`<file>` to `<file>.save`. In this way, you will not lose previous configuration
information. However, the original configuration file (`<file>.save`) will be over-written
if you re-install Apache.

Upon successful installation, `swinstall` runs the HP Apache-based Web Server
automatically.

Documentation

All HP-specific documentation included in the HP Apache-based Web Server software
can be found online after installation at **http://*machinename*/doc** or in the
`/opt/apache/htdocs/doc` directory.

For the latest information on the HP Apache-based Web Server, go to the Web at
http://www.hp.com/go/webservers.

HP-UX Runtime Environment for the Java 2 Platform

**updated for
March 2002**   Both version 1.2 and version 1.3 will be installed unless the customer chooses otherwise.
(See the following documentation section for the location of further information.)

**updated for
December 2001**   The HP-UX Runtime Environment for the Java 2 (RTE) Platform has been updated to
version 1.3.

**updated for
September 2001**   The HP-UX Runtime Environment for the Java 2 (RTE) Platform contains the basic
components for executing a Java application on HP servers and workstations with
HP-UX 11i.

Documentation

For further information, please read the release notes in the RTE software. Or, for the
most up-to-date information, go to the Web at:

http://www.hp.com/go/java

HP-UX Support Tools: STM, ODE, & EMS Hardware Monitors

HP-UX Support Tools provide a complete set of tools for verifying, troubleshooting, and
monitoring HP 9000 system hardware, including CPUs, memory, interface cards, and
mass storage devices.

updated for December 2001

Support Tools Manager (STM) is the platform for executing online diagnostics. The commands to start it are `xstm` (GUI interface), `mstm` (menu-driven interface), `cstm` (command line interface), or `stm` (general).

With the December 2001 release, STM has been updated to version A.29.00.

Offline Diagnostic Environment (ODE) is the platform for executing offline diagnostics. Normally it is run from the Support Plus CD with the system offline.

EMS Hardware Monitors allow you to monitor the operation of a wide variety of hardware products and be alerted immediately if any failure or other unusual event occurs. The EMS Hardware Monitors are started automatically with no user intervention.

NOTE   Since the initial release of HP-UX 11i, there have been numerous changes to the Support Tools, including support for new devices, bug fixes, and enhancements. For the latest information on these changes, see the "Documentation" section below.

The rest of this section describes the major differences between Support Tools on HP-UX 11i and Support Tools on previous releases (HP-UX 10.20 and 11.0).

Impact

The HP-UX Support Tools have been modified to support new products, such as Superdome systems.

With HP-UX 11i, the Support Tools are automatically installed with the HP-UX 11i Operating Environment CD. It is no longer necessary to load the Support Tools from the Support Plus media. (The Support Plus media, however, still contains the Support Tools, and will continue to be distributed. Offline tools are run from the Support Plus CD and cannot be run from the HP-UX 11i Operating Environment CD.)

NOTE   As of HP-UX 11i, Predictive Support is no longer distributed with the Support Tools.

Disk space required by the HP-UX 11i Support Tools is comparable to the disk space required for previous releases (in the range of 60-70MB).

Compatibility

There are minor changes in `monconfig`, the user interface for configuring EMS Hardware Monitors. These changes relate to the client configuration files which have been added to support the multiple-view (Predictive-enabled) feature.

If you have scripts which invoke `monconfig`, they may have to be modified.

For more information on these changes, refer to "Adding a Monitoring Request" in Chapter 2 of the *EMS Hardware Monitors User's Guide* (June 2000 or later edition) available by searching on the manual title at:

http://docs.hp.com/hpux/diag/

Alternately, you can just run `monconfig` on HP-UX 11i to see the revised dialog.

Documentation

The `http://docs.hp.com/hpux/diag/` web site also has tutorials, FAQs, Release Notes, and manuals documenting the Support Tools. Although some documentation is also available through other means, such as through the Support Plus CD, the web pages provide the latest information.

For changes since the initial release of HP-UX 11i, see the March 2001, June 2001, December 2001, or March 2002 documents at the following URLs:

For *EMS Release Notes*, see
`http://docs.hp.com/hpux/onlinedocs/diag/ems/ems_rel.htm`

For *STM Release Notes*, see
`http://docs.hp.com/hpux/onlinedocs/diag/stm/stm_rel.htm`

Netscape Communicator

Netscape Communicator version 4.7x (B.11.11.05) includes Netscape's popular Web browser, Navigator, as well as Messenger and Composer. Communicator offers the complete set of tools for browsing dynamic Web content, plus complete e-mail capability.

Netscape provides periodic maintenance releases for enterprise customers that include minor feature enhancements as well as improvements to overall stability.

Installation Requirements

Netscape Communicator requires 25MB of disk space.

Perl Programming Language

**new for
September 2001**

Included as of the September 2001 release, Perl programming language version 5.6.1 is a release of ActivePerl, a product of ActiveState Tool Corporation.

Documentation

For more information see the following:

- *Perl Programming*, Third Edition, by Larry Wall, Tom Christiansen, and Jon Orwant. O'Reilly and Associates, Inc. USBN 0-596-00027-8

- the *perl* (1) manpage (points you to related perl manpages)

- the `/opt/perl/bin/perldoc` program

For further information, see the following URLs:

`http://www.perl.com`

`http://www.perl.org`

Pluggable Authentication Modules (PAM) Kerberos

Pluggable Authentication Modules (PAM) Kerberos version B.11.11 is a service for authenticating users or services across an open network. HP-UX 11i provides Kerberos authentication through a Kerberos-Client product which is a part of the HP-UX base

operating system. Kerberos, the primary authentication mechanism for Windows 2000, is integrated with Active Directory Service to provide enterprise-wide account management. This necessitates the implementation of the Kerberos authentication mechanism on HP-UX as a Pluggable Authentication Module.

Pluggable Authentication Modules (PAM) [OSF RFC 86] is the standard authentication mechanism, and is easily configurable to support multiple authentication technologies on HP-UX.

PAM Kerberos provides the PAM mechanism and encryption support.

The PAM service modules were implemented as a shared library, `libpam_krb5.1`. This library is built by linking with `libkrb5.1`, and is therefore not dependent on the `libsys.sl` library.

The HP-UX 11i implementation of Kerberos version 5 protocol provides enterprise-wide strong user authentication. Using encryption during the user authentication process, Kerberos infrastructure provides privacy and integrity of user login information since passwords are no longer communicated in clear text over the network.

HP-UX system entry services can work with any Kerberos v5 Server, namely, MIT Kerberos and Microsoft Windows 2000. Thus, passwords can be effectively unified in an Intranet with heterogeneous systems such as UNIX and Microsoft Windows 2000. Furthermore, support of password change protocol automates propagation of password changes. These two features can significantly reduce user administration complexity in a heterogeneous environment.

The HP-UX applications using PAM include `telnet`, `login`, `remsh`, `ftp`, `rexec`, `rlogin`, `dtlogin`, and `rcp`. PAM Kerberos interoperates with a Key Distribution Center (KDC) operating on either a UNIX or a Microsoft Windows 2000 server.

The PAM Kerberos module is compliant with IETF RFC 1510 and Open Group RFC 86. PAM Kerberos is also available under the product number J5849AA on the Applications Software CD. This product provides a `libpam_krb5.1` library, a *pam_krb5* (1) manpage, and a release note document.

Installation Requirements

The minimum disk space required to install the product is 1MB. Additional disk space of about 1KB per user in the system `/tmp` file is required to store initial Ticket Granting Tickets in the credential cache file.

Impact

HP-UX PAM Kerberos is implemented under the PAM framework, which allows new authentication service modules to be plugged in and made available without modifying the application or rebooting the system.

PAM Kerberos works on HP 9000 servers and workstations with a minimum of 32MB of memory and sufficient swap space (a minimum of 50MB is recommended).

NOTE   PAM Kerberos is not thread safe.

Coexistence Issues

PAM Kerberos (`libpam_krb5.1`) and PAM DCE (`libpam_dce.1`) plug-in modules can not be stacked together in the `pam.conf` file because of different principal styles and credential file paths. If so stacked, the results will be unpredictable.

The Kerberos system `ftp` service may list the `/etc/issue` file before the expected output. The *sis* (5) manpage provides detailed information. If the password has expired on a Microsoft Windows 2000 KDC, you will be asked for a new password but will not be allowed to log in. This is a known problem in Windows 2000.

When changing passwords on a MIT KDC with a version prior to 1.1, up to 45 seconds may elapse before the password is actually changed due to the selection mechanism of the change password protocol.

Documentation

The following documentation is available:

- The newly created manpage for `pam_kerberos` is available at `/usr/share/man/man5.Z/pam_krb5.5`.

- The white paper, *Network Security Features of HP-UX 11i*, is available at `http://www.unix.hp.com/operating/hpux11i/infolibrary/`.

- The *PAM Kerberos Release Notes for HP-UX 11i* is available at `ttp://docs.hp.com`.

Plug-In for the Java 2 Platform for Netscape Communicator

updated for September 2001

The Runtime Plug-in (JPI) for the Java 2 platform allows you to use a version of the runtime environment that differs from the runtime environment embedded with Netscape Navigator 4.61 or later.

Documentation

For prerequisites, installation information, and documentation, read the release notes included in the Plug-in software. Or, for the most up-to-date information, go to the Web at `http://www.hp.com/go/java`.

Servicecontrol Manager (SCM)

Servicecontrol Manager provides a convenient, single point of administration for HP-UX systems and Linux managed nodes. Now bundled with every HP-UX 11 release (as well as being available for download), Servicecontrol Manager integrates the key products of the HP-UX Servicecontrol systems management tools suite. You can now manage Linux managed nodes from an HP-UX central management server.

updated for December 2001

Servicecontrol Manager version A.02.05 provides the following new features:

- You can now manage Linux nodes from the HP-UX central management station (CMS). Supported distributions for Linux managed nodes are Debian 2.2 and Red Hat 6.2.

- You now have new GUI tools to reconfigure managed nodes after software upgrades.

- Integration with Partition Manager is now available for Superdome users.

Compatibility

As of Servicecontrol version A.02.05, the following Operating Systems are supported:

- Central Management Servers

  — HP-UX 11.00
  — HP-UX 11i Operating Environment
  — HP-UX 11i Enterprise Operating Environment
  — HP-UX 11i Mission Critical Operating Environment

- Managed Nodes

  — HP-UX 10.20 and 11.00
  — HP-UX 11i Operating Environment
  — HP-UX 11i Enterprise Operating Environment
  — HP-UX 11i Mission Critical Operating Environment
  — Linux, Debian 2.2 and Red Hat 6.2

- PC Web Access to Central Management Server

  — Microsoft Windows 95, 98, and NT 4.0

Note that the Java Plug-in, required software for a PC Web Client, does not support Windows 2000.

updated for September 2001

Servicecontrol Manager version A.02.03 provides the following new management and security features:

- Integration with HP TopTools 5.5 is now included. TopTools is a web-based tool that helps you to manage your computer assets, network devices, and HP printers, as well as track your network resources and performance.

- When selecting nodes for managed clusters, you can now manage up to 999 nodes in one cluster instead of the previous limit of 64.

- The Distributed Task Facility (DTF) now has a theoretical limit of 500 task executions, instead of the previous limit of 10, and a theoretical limit of 1000 agent connections, up from the previous limit of 16. Both limits may vary depending on such factors as the machine's kernel settings, memory capacity, tasks running, and the number of threads the system is set up to allow.

- A Trusted User can now assign users' roles on node groups as well as individual nodes.

- New public/private key authentication between the Central Management Server and the managed node is now available.

- Digital signatures now provide tamper resistance between the CMS and the managed nodes; however, digital signatures do not provide encryption.

- "On-the-wire" encryption with HP Praesidium IPSec/9000 cryptographic suites can be configured.

- Servicecontrol Manager now provides an automated way to start, stop, and restart daemons, instead of having to kill daemons manually.

new at 11i original release

For HP-UX 11i, SCM includes the following enhancements:

- HP-UX 11i support for the central management station (CMS) and managed nodes
- HP-UX 10.20 support for managed nodes

- support for workstations as CMS/Managed nodes

Documentation

For more information, see the *mxtool* (1) manpage and the following documents available through Servicecontrol's web site at **www.software.hp.com/products/SCMGR**:

- *Servicecontrol Manager Technical Reference v.1.3* (MPN: B8339-90030)
- *Planning, Installing, and Updating Servicecontrol Manager Guide* (MPN: B8339-90029)
- Readme documents and Release Notes

HP-UX 11i Enterprise Operating Environment

**new at 11i
original release**

The HP-UX 11i **Enterprise** Operating Environment provides a superset of the features available in the HP-UX 11i Operating Environment described in "HP-UX 11i Operating Environment" on page 67. Targeted especially for database servers, the Enterprise OE includes these additional applications:

- GlancePlus Pak (see page 80)
- High Availability Monitors (see page 80)
- HP OnLineJFS 3.3 (see page 81)
- MirrorDisk/UX (see page 81)
- Process Resource Manager (PRM) (see page 82)

See also "HP-UX 11i Operating Environment" on page 67.

GlancePlus Pak

**updated for
March 2002**

GlancePlus Pak, now version C.03.55, supports HP-UX 11.0 and 11i. (As of March 2002, it no longer supports HP-UX 10.20.)

**updated for
December 2001**

GlancePlus Pak version C.03.50.00 integrates GlancePlus, HP OpenView Performance Agent for HP-UX (OVPA), and IT/Operations Special Edition (ITO-SE) into a single tool to help customers better manage the performance and availability of their servers.

For more information, please refer to each product's Release Notes at the following site:

> `http://ovweb.external.hp.com/lpe/doc_serv/`

NOTE

HP OpenView Performance Agent for HP-UX (OVPA) was previously named MeasureWare Agent for HP-UX. However, the software components and process names operationally remain MeasureWare Agent (MWA).

High Availability Monitors

The High Availability (HA) Monitors version A.03.20.01 product includes database monitors, disk monitors, and Management Information Base (MIB) monitors that can be used to set up notifications of changes in status for the important objects in a high availability cluster environment.

High Availability Monitors A.03.20.01 is being released for use with the HP-UX 11.0 and HP-UX 11i. This release has all the features found in earlier versions in addition to new functionality, defect repairs, and support for new hardware configurations.

This version is a minor release, with minor changes and defect fixes. The contents of HA Monitors releases A.03.00 through A.03.10 have been incorporated, together with all A.03.10 patches.

Since Event Monitoring Service (EMS) has added a new state, UNAVAIL, to the package monitor, this version of HA Monitors is provided to ensure compatibility with the change in EMS.

NOTE	The HA Monitors product does not provide Native Language Support.

Installation Requirements

The HA Monitors software requires a minimum of 4.45MB of disk space and 32MB of memory.

Documentation

- *Using High Availability Monitors* (B5736-90025)

- *High Availability Monitors version A.03.20.01 Release Notes for HP-UX 11i* (B5736-90036)

These documents are available on the HP-UX 11i Instant Information CD and on the Web at:

`http://docs.hp.com/`

HP OnLineJFS 3.3

HP OnLineJFS 3.3 is the advanced optional product for JFS 3.3, which is the latest version of JFS, the Journaled File System. (JFS is also known as the VERITAS File System or VxFS.)

You can use the capabilities of HP OnLineJFS to perform certain key administrative tasks on mounted JFS file systems; this allows users on the system to perform their work uninterrupted. These tasks include:

- Defragmenting a file system to regain performance

- Resizing a file system

- Creating a snapshot file system for backup purposes

Documentation

See "New Version of Journaled File System (JFS)" on page 158 for information about new features in JFS 3.3.

For more information on JFS 3.3 and OnLineJFS 3.3, see the following books:

- *HP JFS 3.3 and HP OnLineJFS 3.3 VERITAS File System 3.3 System Administrator's Guide*

- *Managing Systems and Workgroups: A Guide for HP-UX System Administrators*

Both are available on the Web at:

`http://docs.hp.com/`

MirrorDisk/UX

Prior to HP-UX 11i, Logical Volume Manager (LVM) mirroring supported the non-Shared Logical Volume Manager (non-SLVM) environment only. In other words, the disks were only accessible by a single system and could not be shared by multiple hosts.

Beginning with HP-UX 11i, LVM mirroring now automatically enables SLVM for a two-node environment supporting both non-SLVM and SLVM environments. All LVM systems can mirror their data on disk, and the mirrored copy of the data can also be accessed from a two-node cluster.

Impact

There have been no changes to the LVM command interface to enable LVM mirroring in the SLVM environment. Therefore, you must still use the `lvcreate` and the `lvextend` commands to create mirrored logical volumes. The only software code changes were made to the HP-UX kernel and do not affect any LVM manpages, or the MirrorDisk/UX version B.11.11 products, which are:

- B5403BA MirrorDisk/UX License for Workstations

- B2491BA MirrorDisk/UX License for Servers

To make use of the LVM mirroring capability, you may want to add extra disks to the volume group to mirror the data.

Compatibility Issues

There is no need to make any changes to scripts or makefiles to make use of the LVM mirroring capability in the SLVM environment.

NOTE SLVM mirroring is NOT supported for striped logical volumes and is ONLY supported in a two-node environment. SLVM mirroring does *not* support spared disks in a shared volume group. You should disable sparing by using the `pvchange -z n` *<path>* command on shared volume disks.

Process Resource Manager (PRM)

Process Resource Manager (PRM) enables system administrators to guarantee CPU, real memory, and disk bandwidth resources to users and applications on the system.

NOTE Process Resource Manager and HP-UX Workload Manager both make use of the PRM API. Consequently, *only one* of the products should be used at a time. (See also "HP-UX Workload Manager" on page 86.)

---

updated for September 2001 HP PRM version C.02.00 provides the following:

- new memory features

    — in-kernel memory management
    — dedicated memory for PRM groups
    — proportional overachievement for PRM groups
    — memory isolation - PRM uses the `prm2D` memory manager to optionally specify a group's memory resource to be restricted by use from other groups and processes on the system.

- support for processor sets

  — HP PRM introduces a new type of PRM group, called a PSET PRM group, that is based on processor sets. The PSET PRM's CPU entitlement is specified by assigning it to a subset of the system's processors.

- `prmanalyze` updated to perform new functionality

- new PRM toolset for SCM

NOTE   Processor Sets can be installed from the December 2001 Software Pack (Optional HP-UX 11i Core Enhancements) CD included in your HP-UX 11i December 2001 media kit. You can also download Processor Sets (or order a CD from the web) from the HP Software Depot at the following web site:

   `http://software.hp.com`

---

new at 11i original release   With version C.01.08.02, PRM provides the following features:

- Distribution of resources through shares, hierarchical PRM groups, in-kernel memory management, a Simple Network Management Protocol (SNMP) agent, remote management of PRM, and an improved GUI.

- Provides a shares model of distributing resources instead of static percentages. This model facilitates configuration changes. PRM groups can be nested, allowing for more convenient partitioning.

- Memory controlled in the kernel, through the `prm2d` daemon, rather than in user space, through the `prm0d` daemon.

- Syntax for memory records is essentially the same, with the only difference being that the optional SUPPRESS field is no longer needed and is ignored if present:

   `#!PRM_MEM:{PRMID|GROUP}:SHARES:[CAP]:::`

- CAP value is treated as a hard limit to the group's memory usage. Previously, it was a soft limit that could be crossed.

- `prm2d` in-kernel memory manager is the default for HP-UX 11i. If you prefer to use the previous manager (`prm0d`), follow the steps below:

  1. As root, go to the PRM install directory:

     `# cd /opt/prm/bin/`

  2. Make a backup of `prm2d`:

     `# mv prm2d prm2d.original`

  3. Create a symbolic link from `prm2d` to `prm0d`:

     `# ln -s prm0d prm2d`

- Provides an enhanced GUI, which enables PRM to be remotely managed from any system that has Java Runtime Environment 1.2.2 (or later). PRM has a SNMP agent that makes configuration and resource information available.

Installation Requirements

Process Resource Manager (PRM) requires a minimum of 9MB of disk space and 2MB of memory.

Compatibility Issues

PRM can be used with any 11.x version of GlancePlus.

Documentation

- The following documents are available on the HP-UX 11i Instant Information CD and on the Web at `http://docs.hp.com/`:

  — The *HP Process Resource Manager User's Guide*
  — *HP PRM version C.02.00 Release Notes for 11.0 and 11i*, and the *PRM Product Overview*

- Manpages:

  — *prm* (1) (revised)
  — *prmagt* (1) (new for 11i)
  — *prmanalyze* (1) (revised)
  — *prmavail* (1) (revised)
  — *prmconf* (4) (revised)
  — *prmconfig* (1) (revised)
  — *prmlist* (1) (revised)
  — *prmloadconf* (1) (revised)
  — *prmmonitor* (1) (revised)
  — *prmmove* (1) (revised)
  — *prmrecover* (1) (revised)
  — *prmrun* (1) (revised)
  — *xprm* (1) (revised)

HP-UX 11i Mission Critical Operating Environment

**new at 11i
original release**

The HP-UX 11i **Mission Critical** Operating Environment is a high-availability Operating Environment for HP 9000 servers. In addition to the features found in the two previously described environments, the Mission Critical OE includes:

- Enterprise Cluster Master (ECM) Toolkit (see page 85)
- HP-UX Workload Manager (see page 86)
- HP-UX Workload Manager Oracle , Database Toolkit (see page 87)
- MC/ServiceGuard (see page 88)
- MC/ServiceGuard NFS Toolkit (see page 89)

For the available selectable products for this OE, see Table 4-2 on page 66. See also "HP-UX 11i Enterprise Operating Environment" on page 80 and "HP-UX 11i Operating Environment" on page 67.

Enterprise Cluster Master (ECM) Toolkit

The Enterprise Cluster Master Toolkit is a set of templates and scripts that allow you to configure ServiceGuard packages for the HP Domain Internet servers as well as for several third-party database management systems. The toolkit also includes other specialized tools for monitoring your mission critical environment.

This release of the Enterprise Cluster Master Toolkit is for use with HP-UX 11.0 and HP-UX 11i, and has all the features found in earlier versions in addition to new features and defect repairs.

**updated for
December 2001**

With new version B.01.07, the following changes have or will be made:

- This will be the last release to include the HA Foundation Monitor Toolkit.

- This release contains support for IBM's DB2 version 7.1, Informix XPS version 8.31, and Progress version 9.1.

- This release contains support for Directory Server version 4.13.

**new at 11i
original release**

Version B.01.06 includes a new toolkit that supports the use of Oracle 8i's Oracle Standby Database in Continental Cluster configurations.

Impact

The Enterprise Cluster Master Toolkit does not provide Native Language Support. However, separate Japanese language versions of the Toolkit's documentation are available as a part of product B5139DA with option ABJ.

Installation Requirements

The disk space requirement is 1.2MB.

Documentation

**updated for
December 2001**

For further documentation, see *Enterprise Cluster Master Toolkit* (part no. B5139-90043), available in hardcopy and on the Web at `http://docs.hp.com/`.

The new toolkit for Oracle 8i Standby Database includes a `README` file,
`/opt/cmcluster/toolkit/SGOSB/README-CC`, which explains how to use the toolkit for
data replication in a Continental Cluster.

HP-UX Workload Manager

HP-UX Workload Manager (WLM) provides goal-based workload management, enabling
automatic resource allocation and application performance management.

HP-UX WLM uses a configuration file specifying workloads and their prioritized
service-level objectives (SLOs). HP-UX WLM automatically allocates CPU resources to
the workloads based on priorities and current performance.

NOTE	Process Resource Manager and HP-UX Workload Manager both make use of the PRM API. Consequently, *only one* of these products should be used at a time. (See also "Process Resource Manager (PRM)" on page 82.)

updated for September 2001

For version A.01.02, new features include the following:

- Example configuration files
- Ability to grant a certain amount of CPU per metric
- Integration with Oracle databases
- Integration with iCOD (instant Capacity On Demand)
- Integration with SCM (Servicecontrol Manager)

new at 11i original release

Besides resource shares, the bundling of HP-UX WLM version A.01.00.02 with PRM
provides:

- An SNMP agent for accessing PRM data

- Display of resource percentages to two decimal places (was formerly integer
  precision)

- Expanded field length for PRM group names in PRM's monitoring utilities
  (30-character field was formerly 14-character field)

- In-kernel memory management

This version of HP-UX WLM does not take advantage of PRM's hierarchical groups.

Installation Requirements

HP-UX WLM requires 9MB of disk space and 5MB of memory.

Compatibility Issues

HP-UX WLM can be used with any 11.x version of GlancePlus. However, GlancePlus
should be used only for monitoring, and not for changing, PRM entitlements.

Documentation

Documentation includes:

- The *HP-UX Workload Manager User's Guide* and the *HP-UX Workload Manager A.01.02 Release Notes for HP-UX 11.0 and HP-UX 11i* available on the HP-UX 11i Instant Information CD and on the Web at:

  `http://docs.hp.com/`

- Manpages:

  — *wlm* (5)

For more information on HP-UX Workload Manager, see the following web site:

`http://www.hp.com/go/wlm`

HP-UX Workload Manager Oracle ® Database Toolkit

new for September 2001

As described earlier ("HP-UX Workload Manager" on page 86), HP-UX WLM provides goal-based workload management. This management enables automatic resource allocation and application performance management through the use of prioritized service-level objectives (SLOs). It provides this functionality by building on HP Process Resource Manager (PRM) functionality.

WLM Oracle Database Toolkit (ODBTK) A.01.00, product number T1302AA, provides an easy mechanism for feeding Oracle database metrics into HP-UX WLM.

ODBTK offers the following features:

- Keeps response times for your transactions below a given level by setting response-time SLOs

- Increases an instance's available CPU when a particular user connects to the instance

- Increases an instance's available CPU when more than n users are connected

- Increases an instance's available CPU when a particular job is active

- Gives an instance n percent of the CPU for each process in the instance (This feature is available only with HP-UX WLM version A.01.02 and later.)

- Gives an instance n percent of the CPU for each user connection to the instance (This feature is available only with HP-UX WLM version A.01.02 and later.)

These features are provided by ODBTK's:

- Ability to gather and send Oracle database metrics (elapsed walltime and SQL values) to HP-UX WLM for use in service-level objectives

- Ability to average extreme values in database metrics before sending them to HP-UX WLM (This feature maintains a more smooth and predictable stream of metrics.)

Compatibility

- HP-UX WLM version A.01.02 on HP-UX 11.0 and HP-UX 11i
- HP-UX WLM version A.01.01 on HP-UX 11.0
- Oracle 8.0.x, Oracle 8.1.5, Oracle 8.1.6, and Oracle 8.1.7

Documentation

Documentation includes:

- The *HP-UX Workload Manager Oracle Database Toolkit User's Guide* and the *HP-UX*

Workload Manager Oracle Database Toolkit A.01.00 Release Notes for HP-UX 11.0 and HP-UX 11i, available on the HP-UX 11i Instant Information CD and on the Web at `http://docs.hp.com/`.

*   *wlmtk* (5) manpage

MC/ServiceGuard

Multi-Computer/ServiceGuard (MC/ServiceGuard) is a specialized facility for protecting mission critical applications from a wide variety of hardware and software failures.

NOTE   When you purchase the Mission Critical OE, it is assumed that you want MC/ServiceGuard, *not* ServiceGuard OPS Edition. MC/ServiceGuard and ServiceGuard OPS Edition cannot coexist on the same system. Users of ServiceGuard OPS Edition are encouraged to purchase and install the Enterprise OE and then install ServiceGuard OPS Edition. For detailed installation information, see the *ServiceGuard OPS Edition version A.11.14 Release Notes* (B5161-90043).

---

updated for March 2002

Version A.11.14 of MC/ServiceGuard contains all the features of earlier releases in addition to new features, defect repairs, and support for new hardware configurations.

New features include the following:

*   Support for the use of a Quorum Server system in place of a lock disk as a cluster lock
*   Support and documentation of ServiceGuard Manager administration feature
*   Package control scripts that now allow you to specify concurrent activation, mounting, and `fsck` operations for storage, thereby improving performance at package startup and shutdown

updated for December 2001

ServiceGuard Distributed Components CD

The ServiceGuard Manager version A.02.00 is being shipped on the ServiceGuard Distributed Components CD, which can automatically install ServiceGuard Manager in either English or Japanese. ServiceGuard Manager installs on three Operating System platforms: HP-UX, Linux, and Windows.

The latest upgrades of the ServiceGuard Manager GUI are also available at no charge from `http://software.hp.com` in the "high availability" area.

ServiceGuard Release Notes can be viewed or printed from `http://docs.hp.com` in the "high availability" area.

updated for September 2001

Version A.11.13 of MC/ServiceGuard includes all the features found in earlier versions in addition to new features, defect repairs, and support for new hardware configurations.

New features include the following:

*   Support for VxVM and CVM disk groups created with VERITAS 3.1 Volume Manager (available only on HP-UX 11i or later)
*   Support for up to 60 packages per cluster
*   Support for ServiceGuard Manager on HP-UX 11i
*   Support for Auto-Port Aggregation Auto Mode on Gigabit Ethernet cards

new at 11i original release

Version A.11.09 of MC/ServiceGuard is for use with HP-UX 11.0 and HP-UX 11i, and has all the features found in earlier versions in addition to new functionality, defect repairs, and support for new hardware configurations.

Auto-port aggregation is fully supported with 100BaseT network interface cards. The contents of MC/ServiceGuard releases A.11.01 through A.11.08 have been incorporated, including all A.11.08 patches.

New diagnostic error messages are written to the `syslog` file (`/var/adm/syslog/syslog.log`) when an attempt to obtain the cluster lock fails. Internal error codes are returned to facilitate troubleshooting.

MC/ServiceGuard supports the online replacement of network and I/O interface cards, which is allowed by the HP-UX 11i Operating Environments.

Impact

Disk space required for MC/ServiceGuard is 47MB. The memory required is 6MB, plus 70KB per package in the cluster, plus 300KB per Event Monitoring Service (EMS) resource in the cluster. This total amount is required on all cluster nodes, regardless of whether a given package or resource is on that node or not.

The EMS-CORE file set is no longer included as part of the MC/ServiceGuard product, even though EMS (B7609BA) is still a dependency and must be installed with MC/ServiceGuard.

NOTE

MC/ServiceGuard does not provide Native Language Support. However, separate native language versions of documentation are available as a part of product B3935DA and B3936EA with the following options:

- AB0 - Traditional Chinese
- AB1 - Korean
- AB2 - Simplified Chinese
- ABA - English
- ABJ - Japanese

Documentation

- *MC/ServiceGuard Version A.11.14 Release Notes* (B3935-90050)

- *Managing MC/ServiceGuard*, Ninth Edition (B3936-90065)

These documents are available on the HP-UX 11i Instant Information CD and on the Web at `http://docs.hp.com/hpux/ha/`.

MC/ServiceGuard NFS Toolkit

Multi-Computer/ServiceGuard Network File Server (MC/ServiceGuard NFS) Toolkit version A.11.11 uses MC/ServiceGuard to set up highly available NFS servers. NFS servers are hosts that "export" their local directories and make them available for client hosts for mounting. On the NFS client, these mounted directories look like part of the client's local file system.

MC/ServiceGuard NFS Toolkit is a set of configuration files and control scripts which can be customized to user-specific needs.

MC/ServiceGuard is used to create high-availability clusters of HP servers. Computer systems with high-availability clusters allow applications to continue in spite of a hardware or software failure. Such systems protect users from software failures as well as from failure of a system processing unit (SPU) or local area network (LAN) component. In the event that one component fails, the redundant component takes over and MC/ServiceGuard coordinates the transfer between components.

updated for December 2001

- MC/ServiceGuard NFS Toolkit versions A.11.00.05 and A.11.11.02 are major releases. Both of these versions have dependencies on MC/SG A.11.13 that provide the integrated solution for multiple MC/ServiceGuard partner software working together under the new MC/SG framework. MC/SG NFS Toolkit has been re-architected to fit under the new framework.

- NFS-related control functions and variables have been extracted from MC/ServiceGuard package control script to a separate NFS specific control script, `hanfs.sh`. When the MC/ServiceGuard package is started, it will check for this script under the package directory, and if it exists, will invoke NFS control script.

- An easy troubleshooting mechanism has been provided. Customers can easily identify if a problem resides in MC/SG NFS Toolkit from the MC/SG package log file.

- MC/ServiceGuard NFS Toolkit A.11.11.02 supports VxVM as an alternative disk storage management technology.

updated for September 2001

- MC/ServiceGuard NFS Toolkit can now monitor TCP services, when NFS over TCP is enabled.

- MC/ServiceGuard NFS Toolkit versions A.11.00.04 and A.11.11.02 can use the same NFS package control script, configuration file, and monitor script.

Installation Requirements

Installing MC/ServiceGuard NFS Toolkit requires only about 55KB of disk space. There are no other disk space and memory requirements.

Software Requirements

- Version A.11.00.04 and A.11.00.05 are supported only on HP-UX 11.0
- Version A.11.11.02 is supported only on HP-UX 11i

Documentation

Product documentation is available on the Web at `http://docs.hp.com/`.

Documentation related to the MC/ServiceGuard NFS products available include the following:

- *Managing MC/ServiceGuard NFS User Manual* (B5140-90011).

- *MC/ServiceGuard NFS Toolkit version A.11.00.02 Release Notes* (B5140-90006)

- *MC/ServiceGuard NFS Toolkit version A.11.00.05 and A.11.11.02 Release Notes* (B5140-90010)

HP-UX 11i Minimal Technical Operating Environment

**new for
June 2001**

New for June 2001, the **Minimal Technical** Operating Environment is the smallest and most fundamental OE that is defined specifically for HP workstations. It exists to offer an HP-UX 11i solution to the customer who is interested in a low-cost HP Workstation and a correspondingly basic Operating Environment. The Minimal Technical Operating Environment is directed to the Workstation OEM market and to those customers for whom the Technical Computing Operating Environment is not a suitable solution.

The Minimal Technical Operating Environment contains all the base functionality that is common to the other four OEs, including the base 32/64-bit HP-UX Operating System, network drivers, and other always-installed applications. However, compared to the Technical Computing Operating Environment, the set of additional applications is greatly reduced.

The Minimal Technical Operating Environment is positioned to deliver a lean set of high-demand applications that do not increase purchase cost, support cost, or license cost over the base Operating System.

The HP-UX 11i Minimal Technical Operating Environment includes these always-installed applications:

- Base VERITAS Volume Manager (VxVM) (see page 68)
- Event Monitoring Service (EMS) (see page 70)
- HP-UX Support Tools: STM, ODE, & EMS Hardware Monitors (see page 73)
- Netscape Communicator (see page 75)
- OpenGL 3D Graphics Developers Kit and Runtime Environment (see page 91)
- Perl Programming Language (see page 75)

For the available selectable products for this OE, see Table 4-2 on page 66.

OpenGL 3D Graphics Developers Kit and Runtime Environment

**updated for
September 2001**

The OpenGL 3D Graphics Developers Kit and Runtime Environment version 1.1 (B6268AA), provides the following 3D APIs: Starbase, PEX, Phigs and OGL.

Various defect fixes have been made to improve quality. New functionality has also been added, including:

- level 2 thread support for OpenGL
- 3DSLS and 3DSLS/d support under 11i on fx4 and fx6
- new graphics hardware fx5/fx10 support
- fxe support

Documentation

Both the Graphics Administration Guide and OpenGL Implementation Guide are available in new releases at the following site:
`http://www.hp.com/workstations/support/documentation/hpux_manuals.html`

For thread support, see the release notes for the appropriate versions of the Xserver, Xlib, and the kernel. The HP OpenGL Release Notes are located in the following directory after the product is installed: `/opt/graphics/OpenGL/11.00_Rel_Notes`.

HP-UX 11i Technical Computing Operating Environment

new at 11i original release

Like the Minimal Technical Operating Environment, the **Technical Computing** OE contains all the base functionality that is common to the other four OEs, including the base 32/64-bit HP-UX Operating System, network drivers, and other always-installed functionality. It is *not* a superset of the standard 11i OE; however, it *is* a superset of the Minimal Technical OE. Also, unlike the Minimal Technical OE, the Technical Computing OE is available on both technical servers and workstations.

The HP-UX 11i Technical Computing Operating Environment includes these applications:

- HP Apache-based Web Server for HP-UX (see page 70)
- CIFS/9000 Client and CIFS/9000 Server (see page 69)
- Event Monitoring Service (EMS) (see page 70)
- High Performance Math Libraries (HP MLIB) (see page 92)
- HP 3D Technology for the Java Platform (see page 94)
- HP Message-Passing Interface (MPI) (see page 94)
- HP Visualize Conference (see page 96)
- HP-UX Runtime Environment for the Java 2 Platform (see page 73)
- Netscape Directory Server (J4258BA) (see page 103)
- Netscape Communicator (see page 75)
- OpenGL 3D Graphics Developers Kit and Runtime Environment (see page 91)
- Perl Programming Language (see page 75)
- Pluggable Authentication Modules (PAM) Kerberos (see page 75)
- Plug-In for the Java 2 Platform for Netscape Communicator (see page 77)

For the available selectable products for this OE, see Table 4-2 on page 66.

FirstSpace VRML Viewer

The FirstSpace VRML (Virtual Reality Markup Language) Viewer version B.11.11 allows you to "drag-and-drop" a VRML model into the view space. FirstSpace has changed from revision 1a to 1.10.

obsoleted for March 2002

The FirstSpace VRML has been obsoleted and will not be installed by the current release of HP-UX 11i. The Update-UX utility will not remove existing instances of FirstSpace VRML at this time.

updated for December 2001

FirstSpace FRML Viewer has been deprecated (slated for future obsolescence) and will be removed from the operating environment in a future release.

High Performance Math Libraries (HP MLIB)

updated for September 2001

New features for version B.08.01 include the following:

- 64-bit integer interface for `VECLIB` and `LAPACK` libraries
- Sparse BLAS functionality
- Sparse solver library improvements

- `LAPACK` 3.0 compliance
- Archive and shared libraries

**new at 11i
original release**

The High Performance Math Libraries (HP MLIB) product contains both the Linear Algebra Package (`LAPACK`) and the Vector Library (`VECLIB`) subprograms, providing mathematical software and computational kernels for engineering and scientific applications involving linear equations, least squares, eigenvalue problems, and the singular value decomposition.

New features for version B.07.01 (described below) include the following:

- Shared libraries
- Basic Linear Algebra Subroutine (BLAS) Standard functionality
- LAPACK 3.0 tuned for HP PA-RISC 2.0 processors
- Simplified sparse solver interface with improved performance
- Improved performance of key routines
- Improved C and C++ usability

Impact

HP MLIB incorporates algorithmic improvements, and several tunable parameters are adjusted for good execution performance.

While you can use HP MLIB as archive or shared libraries, the performance of your applications is better when you use archive libraries. However, if you need to keep executable files small, you can use shared libraries on any PA-RISC 2.0 system running the HP-UX 11.0 or later operating system.

VECLIB is optimized by using a highly efficient implementation of the Basic Linear Algebra Subprograms (BLAS), Levels 1, 2, and 3, as well as a subset of the newly-defined BLAS Standard. Performance for key BLAS routines has been improved.

MLIB fully conforms to the public domain version 3.0 of LAPACK in all user-visible usage conventions. The internal workings of some subprograms have been tuned and optimized for Hewlett-Packard computers.

This version simplifies sparse solver interface use and improves its performance. Version 4.0 of the METIS reordering technology has been incorporated.

You can now use the C or C++ compiler to link applications built with MLIB. (Previous to this release, you were required to link using the Fortran compiler when using VECLIB or LAPACK.)

Documentation

For more detailed documentation and additional product information, see:

- The *HP MLIB User's Guide* (B6061-96015) and the *HP MLIB Release Notice* (B6061-96016) available at the following URLs:

    http://docs.hp.com/

    http://www.hp.com/rsn/mlib/mlibhome.html/

- Manpages:

    — BLAS Standard manpages (new for 11i)

    — LAPACK 3.0 manpages (revised)

— VECLIB manpages (revised)

HP 3D Technology for the Java Platform

The HP 3D Technology for the Java 2 Platform contains the classes for creating 3D applications. The HP 3D Technology for the Java Platform may be distributed with your Java applications as long as you adhere to the terms of the `LICENSE` file. Vendors also need to include an installer.

Documentation

For prerequisites, installation requirements, and information read the release notes included in the HP 3D software. Or for the most up-to-date information, go to the Web at `http://www.hp.com/go/java`.

HP Message-Passing Interface (MPI)

HP Message-Passing Interface (MPI) is a high-performance implementation of the Message-Passing Interface Standard. HP MPI complies fully with the 1.2 standard and partially with the 2.0 standard. HP MPI provides an application programming interface and software libraries to support parallel, message-passing applications that are efficient, portable, and flexible.

updated for March 2002

New feature of HP MPI version 1.7.2:

- With the use of `libmpirm.sl`, the PA-RISC version of Load Sharing Facility (LSF) can be run on the Itanium Processor Family (IPF) version of HP MPI.

updated for September 2001

New features of HP MPI version 1.7 include the following:

- New start up. The new HP MPI start-up requires that MPI be installed in the same directory on every execution host. The default is the location from which `mpirun` is executed. This can be overridden with the MPI_ROOT environment variable. We recommend setting the MPI_ROOT environment variable prior to starting `mpirun`.

  Previous versions of HP MPI allowed `mpirun` to exit prior to application termination by specifying the `-W` option. The option `-W` used with `mpirun` is no longer supported. To achieve similar functionality, place `mpirun` in the background.

- Support for shared libraries. When a library is shared, programs using it contain only references to library routines, as opposed to archive libraries, which must be linked into every program using them. The same copy of the shared library is referenced by each executable using it.

  An advantage of shared libraries is that when the library is updated (e.g. to fix a bug), all programs, which use the library immediately, enjoy the fix.

- Library names. Some of the libraries have been merged. Compilation wrappers have been provided for convenience. Wrappers can also be used as templates.

- Multi-thread mode. By default, the non thread-compliant library (`libmpi`) is used when running MPI jobs. Linking to the thread-compliant library (`libmtmpi`) is now required only for applications that have multiple threads making MPI calls simultaneously. In previous releases, linking to the thread-compliant library was required for multi-threaded applications even if only one thread was making a MPI call at a time.

- Additional MPI-2 support. HP MPI 1.7 expands MPI-2 support of one-sided communications to clusters. Refer to "Appendix C" in the *HP MPI User's Guide*, 6th edition, for a full list of MPI-2 support.

- New options for handling standard IO. HP MPI 1.7 supports several new options for handling standard IO streams.

  All standard input is routed through the `mpirun` process. Standard input to `mpirun` is selectively ignored (default behavior), replicated to all of the MPI processes, or directed to a single process. Input intended for one or all of the processes in an MPI application should therefore be directed to the standard input of `mpirun`.

  Since `mpirun` reads `stdin` on behalf of the processes, running an MPI application in the background will result in the application being suspended by most shells. For this reason, the default mode for `stdin` is off. Running applications in the background will not work with `stdin` turned on.

- Backtrace functionality. HP MPI 1.7 handles several common termination signals differently (on PA-RISC systems) than earlier versions of HP MPI by printing a stack trace prior to termination. The backtrace is helpful in determining where the signal was generated and the call stack at the time of the error.

- IMPI functionality. The Interoperable MPI protocol (IMPI) extends the power of MPI by allowing applications to run on heterogeneous clusters of machines with various architectures and operating systems, while allowing the program to use a different implementation of MPI on each machine.

- Fortran profiling interface. To facilitate improved Fortran performance, we no longer implement Fortran calls as wrappers to C calls. Consequently, profiling routines built for C calls will no longer cause the corresponding Fortran calls to be wrapped automatically. In order to profile Fortran routines, separate wrappers need to be written for the Fortran calls.

- Expanded support for collecting profiling information. You can collect profiling information for applications linked with the thread-compliant library in addition to those linked with the standard MPI library. Counter instrumentation (`MPI_INSTR`) is supported for the thread-compliant library regardless of thread level. Trace file generation (`XMPI`) is supported for all thread levels except `MPI_THREAD_MULTIPLE`.

- A new error checking flag (`-ck`) in the `mpirun` utility. The new error checking flag (`-ck`) allows you to check `appfile` set-up, host machine and program availability, and file permissions without creating MPI processes.

- The `mpirun` utility no longer makes assumptions about how long it will take before a process calls `MPI_Init`. Timeout errors before `MPI_Init` that may have been seen in older versions of the product do not occur in this version because `mpirun` no longer makes assumptions about time to `MPI_Init`.

new at 11i original release

New features of HP MPI version B.11.11 include the following:

- Support for the TotalView debugger.

- Support for collecting profiling information for applications linked with the thread-compliant library in addition to those linked with the standard MPI library.

- Expanded MPI-2 support.

- A new error-checking flag (`-ck`) in the `mpirun` utility.

- The `mpirun` utility no longer makes assumptions about how long it will take before a process calls `MPI_Init`.

Documentation

The *HP MPI User's Guide, Sixth Edition* and the *HP MPI Release Notes* are available on the HP-UX 11i Instant Information CD and at the following web sites:

- `www.hp.com/go/mpi`

- `www.docs.hp.com`

HP Visualize Conference

HP Visualize Conference version 1.4 (B.11.11.06) is a collaborative conferencing solution for HP Workstations that can interoperate with Microsoft's NetMeeting, Sun's SunForum, and SGI's SGImeeting products.

Functionality changes have not occurred in HP Visualize Conference since version 1.3 was released. Version 1.3 added support for NetMeeting 3 while maintaining backward compatibility to version 1.2 and NetMeeting 2 via two running modes:

- T.120 Compatibility mode (HP Visualize 1.2 and NetMeeting 2)

- NetMeeting 3 Compatibility mode

Impact

The HP-UX 11i release of HP Visualize Conference is functionally identical to HP Visualize Conference 1.4 for HP-UX 10.20. There is no impact on other system performance or other system components.

Compatibility Issues

HP Visualize Conference 1.4 (B.11.11.06) on HP-UX 11i is compatible with HP Visualize Conference on HP-UX 11.0 and 10.20, Microsoft's NetMeeting, Sun's SunForum, and SGI's SGImeeting.

Documentation

Online help was enhanced at HP Visualize Conference version 1.3 to cover how to utilize NetMeeting 3 functionality:

- TrueColor Application Sharing

- NetMeeting 3 Application Sharing and Control

- Microsoft variant of the T.126 Whiteboard protocol

Selectable Applications

The following applications are selectable and not automatically installed with the 11i Operating Environments. Table 4-2 on page 66 lists which applications are available for each Operating Environment.

- HP Intrusion Detection System/9000 (IDS/9000) (see page 97)
- Ignite-UX (IUX) (see page 98)
- IPFilter/9000 (see page 101)
- Netscape Directory Server (J4258BA) (see page 103)
- Java Out-of-Box (JAVAOOB) (see page 102)
- Selectable Network Drivers (see page 104)
- WebQoS Peak on HP-UX Packaged Edition (J4274AA) (see page 104)

HP Intrusion Detection System/9000 (IDS/9000)

**new for
December 2001**

The HP Intrusion Detection System/9000 (IDS/9000) is a host-based intrusion detection system for HP-UX 11.x. The IDS/9000, version 2.0 (J5083AA), includes the following features:

- **Integration with OpenView Operations** (OVO, formerly known as VPO or ITO) and the associated Smart Plug In (SPI)**:** Enables customers to monitor IDS/9000 alerts from the OVO management console. For configuration and control of IDS/9000 agent, the IDS/9000 administrative GUI is launched from within the OVO console.

- **New administrative GUI:** More task oriented and easier to use. Reduces the number of interactions for installing, configuring, monitoring and controlling IDS/9000 agents.

- **Web protection capability:** Documentation and a template script on how to customize IDS/9000 to provide near real-time protection against web server defacement.

- **Automatic restart of surveillance schedules after a reboot:** Users no longer need to manually activate the monitoring schedules after a system reboot.

- **Multiple response script capability:** The limit of having a single response script has been removed and users can now have multiple response scripts invoked when an alert is generated.

- **Rewritten and enhanced user's manual:** Task oriented documentation with new sections on agent configuration, response programs, error messages, and troubleshooting.

- **New certificate generation and management tools:** Accommodates a pure Java implementation of SSL for the administrative GUI.

Documentation

The following documents are available in the "Internet and Security Solutions" neighborhood of `docs.hp.com/hpux/internet/index.html`:

- *HP Intrusion Detection System/9000 Release 2.0 Release Notes*
  (product no. J5083-90006)

- *HP Intrusion Detection System/9000 Administrator's Guide*, version 2.0
  (product no. J5083-90007)

Ignite-UX (IUX)

Ignite-UX (IUX) is an HP-UX administration toolset to help you do the following:

- install HP-UX on multiple systems in your network

- create custom install configurations

- recover HP-UX systems remotely

- monitor system-installation status

The complete Ignite-UX product, with support for HP-UX 11i, 11.0 and 10.20, is available on the first CD of the 11i OE media kit.

For additional information about IUX, see "Ignite-UX Changed for 11i" on page 122.

updated for March 2002

Ignite-UX has been updated to version B.3.6 to support changes to the Operating Environments. Changes to B.3.6 include the following:

- An IUX feature has been enhanced so that newly installed core components and patches will be swmodify'ed with the correct file data to avoid swverify problems with those files.

- To avoid accidental tape rewinding, the print_manifest command will no longer call diskinfo (or provide a tape capacity estimate) for tape devices.

updated for December 2001

Ignite-UX version B.3.5 includes the following changes:

- Drivers such as X.25 (nioxb) now have their driver binding preserved during a recovery. This prevents them from becoming "unclaimed" after a recovery.

- The hw_instance_num config file keyword includes a new optional driver parameter. See the *instl_adm* (4) manpage for more information.

- The save_config command now sets the above driver parameter. A check was also added that ensures that the hardware path is of the format expected by the parser for HW_instance_num. This fixes the problem introduced by the sdm driver which has a hardware path of "es" for the virtbus class.

- Ignite-UX servers that have multiple IP addresses assigned to a single network card now work correctly—as far as doing network boots using instl_bootd is concerned. However, this does *not* fix the issue documented in the FAQ item #1.9. The FAQ (a text file) is installed with Ignite-UX in /opt/ignite/share/doc/FAQ.

- Ignite-UX now supports setting the block size for VxVS file systems. Previously this was only allowed to default. The user interface now permits this to be set to valid values and the save_config command saves it in order to restore it for a recovery. A sanity check now tests for valid values (also covering block sizes for HFS file systems) and applies this when laying out VxFS file systems.

- If an OS archive image contains volume group files (for example, the directory /dev/vg00 and the files underneath that), and the major and minor numbers match ones that have been created but have a different name (like if vg00 was renamed bpr00), then the /dev/vg00 directory and its contents will be removed and the action logged to install.log. This avoids problems with commands like swapinfo

which would report incorrect information since it would find `/dev/vg00` before `/dev/bpr00`.

- Additional sanity checks have been added to check for invalid mount points. This includes a missing mount point, a null string mount point, and a mount point that does not begin with a slash ("/") character.

- The subdirectory `/etc/opt/OV` (part of OpenView) is now permitted to be a mount point.

- The `make_bundles` command was incorrectly putting the entire product into bundles created when only a list of filesets on the command line were given. This ended up causing clients added for recovery from the ignite GUI to have the entire Ignite-UX product loaded instead of the subset needed for `make_net_recovery` as intended. If you don't need the entire Ignite-UX product on the clients (that is, if you don't use `make_tape_recovery`), you can manually remove the `Ignite-UX.FILE-SRV-*` filesets on the clients to free up space.

- The `make_medialif` command no longer prevents the usage of the `-a` and `-o` options to B.11.00 releases and later. There are clients which can run 10.20 and which require the WINSTALL kernel from the B-version of Ignite-UX in order to be installed. Now that it is possible to create a LIF header for such 10.20 clients which would include both WINSTAL and 10.20 SYSCMDS.

- An issue was found that will cause newer PA-RISC workstations running the B-version of IUX with the HP-UX 10.20 to be unable to boot from a tape created with either `make_tape_recovery` or `make_recovery`. This issue has been fixed with this version of IUX. There is no issue with these same systems using version A of IUX. (See section #1.20 in the FAQ for more information including a workaround. The FAQ, a text file, is installed with IUX in `/opt/ignite/share/doc/FAQ`.)

- If a mount point contains a symbolic link, it will now follow the same rules for archival as all other mount points. It either needs to be specifically included or marked as an essential item or it needs to be on a volume group or whole disk that already contains some item that has been included or marked as an essential item. If neither of these cases are met, then the mount point will not be included in the archive.

updated for September 2001

The following changes are included in Ignite-UX version B.3.4:

- If the DNS domain has been set, the line in the `/etc/hosts` file for the system being installed will be written out as:

    <IP address><tab><hostname>.<domain> <hostname>

For example:

    `15.16.17.18 myhost.corp.com myhost`

If the DNS domain has not been set, only the hostname without domain is used.

- The `add_release` command is not supported for the B.11.11 release and beyond. This command will be completely obsoleted in a future release. Instead of `add_release`, you may use the following commands:

`# make_depots -d <depot name> -s <source_device>`

`# make_config -c /var/opt/ignite/data/Rel_B.<XX.YY>/core_cfg -s <depot name>`

```
# manage_index -a -f /var/opt/ignite/data/Rel_B.<XX.YY>/core_cfg
```

- The B version of IUX now uses the ndd command to turn off subnet mask checking to allow for subnet masks such as 255.0.0.0. This is only done when installing HP-UX 11.00 or later.

- Removing volumes when recovering a system from a make_*recovery backup now correctly results in them being commented out of the /etc/fstab file. This also fixes a problem where if the /var volume was removed, the /var directory was left renamed to /var_iux after the system was recovered.

- A manual page for the fix_patches command has been included for the first time. See the *fix_patches* (1M) manpage for more information.

- The B version of Ignite-UX no longer supports installing or recovering disk devices connected to the older PCI Tachyon (A3740A) fibre channel cards. This has been the case since the B.3.0 release.

updated for June 2001

The following changes are included in Ignite-UX version B.3.3:

- A sanity check has been added that tests the relationships between memory size, allocated dump space, and free space in whatever logical volume contains /var/adm/crash. The check will produce messages if dump space is less than 50 percent of memory, and if the free space for /var/adm/crash is less than the dump space allocated. The check, at worst, produces warnings which can be ignored, and installations are allowed to proceed.

- File mtime, ctime, and atime behaviors have changed in the recovery commands. The commands make_tape_recovery, make_net_recovery, make_recovery, and make_sys_image no longer use the -t option to the pax command. In the past, this option was used so that the "atime" (access modification time) of the files being backed up was restored. However, a hidden side-effect of this -t option was that it caused the "ctime" (inode modification time) of files being backed up to change, which, in turn, triggered security tools to raise an alarm. Another side effect of the -t option also caused the "mtime" (modification time) of /etc/passwd and /etc/group to be changed due to the way make_recovery used -t in combination with the -s pax option.

- Device files for drivers that use dynamic major numbers are now modified to track any changes to the dynamic major assignments that may occur during the installation or recovery of a system. Prior to this change, some products (such as X.25 and EMS) may not work correctly after recovering a system, or when installing a system from an archive image. For this change to work, the OS archive image must contain the /stand/ioconfig file from the original system. Therefore, you will need to recreate old OS archive images with the new make_sys_image script that includes /stand/ioconfig in the archive.

- The make_sys_image command no longer retrieves variable values from the /tmp/install.vars or /var/opt/ignite/local/host.info files.

- A UID larger than 64K on an automount directory will no longer cause the contents of the automounted directory to be backed up as if it were a directory on the local disk or volume group. If a mount is detected as stale, the system will now issue a WARNING and continue the traversal of the volumes and disks that need to be recovered and put them in the archive image.

- When recovering systems that use Auto Port Aggregation Software (APA) installation, IUX no longer comments out the Internet Configuration parameters for LAN aggregates in the /etc/rc.config.d/netconf file. If the recovery is done

over the network, you will need to specify that the networking parameters are temporary for the system to come up on the LAN aggregate interface.

- A new option `-V` has been added to `bootsys`. One or more `-V` options may be supplied to set a value to a variable.

- All options of `make_bundles` other than `-f`, will generate content lines with product-level instead of fileset-level detail. This is a fix for duplicate entries in the bundle contents due to a product having both 32-bit and 64-bit filesets for 11.x and beyond. The `-f` option now lists the `fa=*` attribute in the bundle contents.

Documentation

The *Ignite-UX Administration Guide* has been updated for HP-UX 11i, and is available on the HP-UX Instant Information CD and on the `http://docs.hp.com/` web site. Another excellent source of information on Ignite-UX is the external web site:

`http://software.hp.com/products/IUX/`

IPFilter/9000

The security product, IPFilter/9000, provides system firewall capabilities by filtering IP packets to control traffic in and out of a system. IPFilter/9000 includes support for Static Linking, Gigabit Ethernet, Auto Port Aggregation (APA), and Virtual Local Area Network (VLAN).

updated for March 2002

IPFilter/9000 has been updated to version A.03.05.03 to incorporate defect fixes.

new for December 2001

IPFilter/9000, version A.03.05.02, is now available.

Impact

System performance may improve when IPFilter is installed but not running.

Documentation

For detailed, updated information, see the documents and web sites listed below:

- Documents (available at `www.docs.hp.com`):

  — *Installing and Administering IPFilter/9000* (B9901-90005)
  — *IPFilter/9000 Release Note*, Edition 5 (B9901-90006)

- Manpages:

ipf (4)	packet filtering kernel interface
ipf (5)	IP packet filter rule syntax
ipf (8)	alters packet filtering kernel's internal lists
ipl (4)	data structure for IP packet log device
ipmon (8)	monitors `/dev/ipl` for logged packets
ipstat (8)	reports on packet filter statistics and filter list
iptest (1)	test packet rules with arbitrary input

- Web Sites:

  — `http://techsolutions.hp.com/ipfilter/ipfilter_homex.html`
  — `http://software.hp.com`
  — `http://docs.hp.com`

Java Out-of-Box (JAVAOOB)

new for March 2002

Java Out-of-Box (JAVAOOB), version 1.0.00, is a small, installable product that modifies HP-UX kernel parameters and system tunables to allow large server-side Java applications to run.

While no claim is made that JAVAOOB will enable *all* Java applications to run (or that it will cause any particular Java application to exhibit optimum performance), the JAVAOOB parameter values were chosen to enable the widest range of Java applications to run successfully without significant impact on other (non-Java) processes. (The original default settings were too low to run the majority of Java applications.)

JAVAOOB modifies the following kernel parameters and system tunables:

```
maxusers              512
nproc                 2048
max_thread_proc       3000
nkthread              6000
nfile                 3000
maxfiles              2048
maxfiles_lim          2048
ncallout              6000
maxdsiz               2063835136
tcp_conn_request_max  2048
```

NOTE

JAVAOOB will not lower a parameter setting. If the current value of a parameter is higher than the value that JAVOOB would set, then JAVAOOB leaves the parameter value unchanged.

In addition to being provided through the OE media, JAVAOOB is available for download from `http://www.hp.com/go/java`.

Impact

Installing JAVAOOB will increase the amount of memory used by the kernel by approximately 8Mb above that used by the default settings. In theory, a system that is memory-bound, close to the edge, and running memory-intensive applications could experience additional paging activity, thereby lowering system or application performance.

Installation Requirements

Since it is intended for large, server-side Java applications, JAVAOOB is recommended only for systems with at least 512 Mb of physical memory (and not for systems used solely as desktops.)

Documentation

For further information, see the Java web pages at `http://www.hp.com/go/java` and the release notes that are part of the product.

For background information on tuning HP-UX for Java applications and for discussions of specific kernel parameters, see the technical paper "HP-UX Kernel Parameter Configuration for Java" at the developer portal (`http://www.hp.com/dspp`) under "topics" then "java."

Also at `http://www.hp.com/dspp` are other articles that may be of interest:

- "`hpjconfig` 2.0 for HP-UX 11.0 and 11i"

- "Java1 performance and tuning"

- "Java2 performance and tuning"

Netscape Directory Server (J4258BA)

The Netscape Directory Server Version is an industry-standard Lightweight Directory Access Protocol (LDAP) directory server. Netscape Directory Server for HP-UX is a selectable product for the HP-UX 11i OE.

updated for September 2001

Version B.04.13 includes features from earlier versions in addition to defect repairs.

Impact

You must purchase Extranet Client Access Licenses to use the Netscape Directory Server for HP-UX if the directory contains any entries for Extranet Users. An Extranet User is an entry in the Netscape Directory that represents a person that is neither an employee nor a full-time independent contractor of the company to which the Netscape Server is licensed. Contact your HP sales representative to purchase licenses. For contact information, see `http://eproducts.hp.com/buy2/index.html`.

Documentation

For more information, refer to the *Netscape Directory Server 4.13 for HP-UX Release Notes* (J4258-90003) available on the HP-UX 11i Instant Information CD and on the web at `http://docs.hp.com/`.

NOTE

Currently, Netscape Directory Server version 4.13 is included with the HP-UX 11i Operating Environments. If you wish to have the iPlanet 5.0 product, you may separately order the AR1201 CD for this product.

Selectable Network Drivers

updated for September 2001

The following table indicates which drivers are selectable during HP-UX 11i installation. Note that several new I/O adapters are fully supported as of June 2001 (A6684A, A6685A, A6748A, A6749A, and A6386A):

Table 4-4 **Selectable I/O Drivers on 11i OE Media**

I/O Driver	Description	I/O Adapter Product Number
100BaseT-00	EISA 100BaseT	A4308B
100BaseT-01	HP-PB 100BaseT	A3495A (on K and T-Class)
ATM-00	PCI ATM	A5483A, A5513A, A5515A, J3557A
ATM-01	HSC ATM	J2468A. J2469A, J2499A, J3420B, J3573A
FDDI-01	HSC FDDI HSC FDDI	A3722A (K and T) A3723A (B, C, D, J)
FDDI-02	HPPB FDDI	J2157B (K and T)
FDDI-03	EISA FDDI	A3659A, B5502BA
HyprFabrc-00	PCI HyperFabric HSC HyperFabric HSC HyperFabric PCI HyperFabric	A4919A (N and V) A4920A (K) A4921A (D and R) A6092A, A6386A (L, N, S, A4, A5)
TermIO-00	PCI MUX (8-port) PCI MUX (64-port)	J3592A, A6748A (L, N, A4, A5) J3593A, A6749A (S, V, L, N, A4, A5)
TermIO-01	EISA MUX	J2482A, A4930A
TokenRing-00	PCI Token Ring	A5783A, A4930A
TokenRing-01	HPPB Token Ring	J2166B (K and T)
TokenRing-03	EISA Token Ring	J2165B

new for June 2001

HSC HyperFabric2 PCI fiber adapter (for A400, A500, L-, N-, V-Class, and Superdome servers, B1000, C3000, J5000, J5600, J6000, and J7000 workstations) A6386A previously supported on HP-UX 11.0 is now supported on HP-UX 11i. See "HyperFabric2 PCI Fiber Adapter" on page 117.

WebQoS Peak on HP-UX Packaged Edition (J4274AA)

The HP Web Quality of Service (WebQoS) Peak Packaged Edition version B.01.02.06 is a Web-based solution that provides the quality of service needed to maintain your Web applications.

deprecated in December 2001

WebQos Peak has been deprecated (slated for future obsolescence) and will be removed from the operating environments in a future release.

NOTE   This product does *not* support the iPlanet Web Server.

Documentation

Other than the unchanged online help for this product, there is no hardcopy or other online documentation associated with it.

5   I/0 and Networking Cards

Online Addition and Replacement of I/O Adapters

**new at 11i
original release**

Online Addition and Replacement (OLAR) is a new HP-UX software feature that allows for adding and replacing PCI I/O cards (adapters) while a system is running, eliminating the need to reboot.

This feature enhances overall high-availability since the system can remain active while an I/O adapter is being added or replaced. When combined with other high-availability products, such as HP MC/ServiceGuard, system availability is significantly improved.

Systems Administration Manager (SAM) provides the system administration interface for OLAR. The first release of OLAR in HP-UX 11i provides support for L-Class, N-Class, and Superdome systems. Many future HP 9000 systems are being designed with this feature as well.

For more information about the OLAR feature, see:

- "Changes to System Administration Manager (SAM)" on page 133

- *Configuring HP-UX for Peripherals*, part no. B2355-90698

- *HP System Partitions Guide*, part no. B2355-90746

- *Managing Systems and Workgroups: A Guide for HP-UX System Administrators*, part no. B2355-90742

**updated
June 2001**

SAM patch PHCO_23004 for 11i (included in the HWEnable11i patch bundle and automatically installed) changes the behavior of the PCI card slot LED (Attention Indicator) to conform with the newly implemented PCI SHPC (Standard Hotplug Controller) specification.

NOTE

In contrast to their former behavior, PCI slot attention indicators now meet the PCI Hotplug specification by *blinking* during tasks involving PCI cards or card slots.

See "New Attention Indicator Behavior" on page 40 in Chapter 2 for more details.

Network Drivers

**new at 11i
original release**

New Network Driver btlan Pre-Installed

The networking driver for HP-UX 11i has been simplified and is now easier to install and upgrade.

The PCI and HSC-based Fast Ethernet network and I/O cards supported by drivers `btlan`, `btlan3`, `btlan4`, `btlan5` and `btlan6` have been combined into a single driver called `btlan`. This new driver is pre-installed as part of the kernel.

The result is to ease setup or upgrade of the networking and I/O products by eliminating driver installation and combining multiple drivers into one.

Impact

The `btlan` driver works seamlessly with existing HP LAN link administrative commands such as `lanadmin`, `lanscan`, `linkloop`, and `NetTL`.

The new `btlan` driver supports the same functions and features as the previous HP-UX 10.20 and 11.0-based drivers. In addition, it also supports the online addition and replacement of I/O cards on L-Class, N-Class and Superdome servers.

IMPORTANT   Although there you will see no impact in most cases, you will if you have scripts that refer specifically to `btlan3`, `btlan4`, `btlan5` or `btlan6`.

You will need to use the new driver name `btlan` with the following commands:

- `what` *string*
  For example: **what /stand/vmunix | grep btlan**

- `ioscan`
  For example: **ioscan -kfC lan | grep btlan**

You will also see the driver name `btlan` as the output in:

- the system file `/stand/system`
- `nettlgen.conf` and in the file `/var/admin/sw/nettl.LOG00`

Files Changed

The following files have changed to include the new `btlan` driver name (most involve just name changes):

- kernel library is now called `/usr/conf/libbtlan.a`

- `nettl` formatter/catalog files (no change except instead of `btlan3`, `btlan4`, `btlan5`, `btlan`, or `btlan6`, it will just refer to `btlan`)

- `debug`/q4

- `lanscan`/`lanadmin` support libraries/catalog files now have names to reflect `btlan` such as `libdsbtlan.a`, `dsbtlan.cat`, etc.

- master file

- `init scripts/conf` file

- The `init` script will be `hpbtlan` and the configuration file will be called `hpbtlanconf`.

- The configuration files under `/etc/rc.config.d/` will be replaced by `hpbtlanconf`. When a cold install is performed, this file will be installed for all `btlan` driver claimed cards. If, however, an upgrade is done, you can choose to merge the files using pre-update scripts. If you do not elect to merge during an upgrade, then the files will, by default, be saved as `.obsolete` files which can be later merged manually into the `hpbtlanconf` file.

Networking and I/O Card Drivers Pre-Installed

The drivers of all of the following networking, I/O, and mass storage cards are now pre-installed with (or built into) each of the HP-UX 11i Operating Environments:

updated for March 2002
- PCI 10/100Base-TX card A3738A will be discontinued starting February 1, 2002
- PCI 4-port 10/100Base-TX card A5506A has been replaced by A5506B

new for December 2001
- Core 1000Base-T/SCSI card A6794A for HP server rp7410
- PCI Tachyon XL2 Fibre Channel card A6795A

updated for June 2001
- PCI 1000Base-T (gigabit over copper) card A4929A
- PCI 1000Base-SX (gigabit over fiber) card A4926A
- PCI Combination Dual port 10/100Base-TX and Wide Ultra2 SCSI card A5838A
- PCI 4-port 10/100Base-TX cards A5506A and A5506B
- PCI 10/100Base-TX card A3738A
- PCI core 10/100Base-TX card for workstations and servers
- PCI 10/100Base-TX card A5230A for servers
- PCI 10/100Base-TX card B5509BA for workstations
- PCI Tachyon TL Fibre Channel card A5158A
- PCI FDDI card A3739B
- PCI RAID 4Si controller A5856A
- HSC 10/100Base-TX card J3514A opt #001 2-port for K-Class servers
- HSC 10/100Base-FX (fiber) card J3514A opt #002 2-port for K-Class servers
- HSC 10/100Base-TX card J3515A 1-port for workstations and D-Class servers
- HSC 10/100Base-TX card J3516A opt #001 2-port for workstations and D-Class servers
- HSC 10/100Base-FX (fiber) card J3516A opt #002 2-port for workstations and D-class servers
- HSC 10/100Base-TX card J3850A for T-Class server
- HSC 1000Base-SX (gigabit over fiber) cards A4924A, A4925A
- HSC Tachlite Fibre Channel cards A6684A, A6685A

Instructions for configuring built-in (or core) 100Base and 1000Base PCI cards can be found in Appendix C of the *HP-UX 11i Installation and Update Guide*.

Combining Pre-HP-UX 11i Configuration Files

In HP-UX 11i, the drivers for PCI and HSC-based Fast Ethernet networking are consolidated into one driver called `btlan`, which is pre-installed as part of the kernel. The configuration files used by these networking drivers in HP-UX 10.20 and

---

11.0—`hpbtlanconf` (`btlan`), `hpbase100conf` (`btlan3`), `hpgsc100conf` (`btlan4`), `hppci100conf` (`btlan5`) and `hpsppci100conf` (`btlan6`), which are in the `/etc/rc.config.d` directory—are merged into one file.

IMPORTANT   The configuration files used by the PCI and HSC-based Fast Ethernet networking drivers must be combined, either by use of a script or manually, into one configuration file *before* upgrading to HP-UX 11i.

Merging must be done before updating because the interface used to recognize the host might be PCI or HSC-based Fast Ethernet. One of these interfaces could have been configured in a pre-update configuration file, which would have to be maintained for the update to proceed (the link-up is needed during product configuration).

Merging with the `BTLAN.100` Script

When you run the pre-update script `BTLAN.100`, the configuration files used by the PCI and HSC-based Fast Ethernet networking drivers are merged into one file called `hpbtlanconf`.

Once the `BTLAN.100` script creates the merged configuration file, it is placed temporarily in the `/var/adm/sw/save_custom/UNIFIED_MER` directory, and the original configuration files are saved with the extension `.obsolete`. During the consolidated `btlan` driver installation, the merged configuration file is then moved to the `/etc/rc.config.d` directory as `hpbtlanconf`.

Note that the `BTLAN.100` script will only merge the driver-specific configuration files if the corresponding hardware is present on the system, and if the files have at least one LAN interface configured. (For example, if the system had the HSC-based Fast Ethernet networking driver [`btlan4`] and corresponding hardware [HSC cards], then the `hpgsc100conf` configuration file should have at least one card configured in it.)

Manual Merging

If the update process is abandoned because of preupdate script failures (other than syntax errors in configuration files or a duplicate LAN interface error), then the configuration files have to be merged manually. Once this has been done, you can then continue the update process. (Once the update process is restarted, the user should answer NO (N) at the prompt, `Do you want to proceed in merging the configuration files into one? [Y|N]`.)

Manual Merging Procedure

1. Find out which of the five configuration files listed below exists in the `/etc/rc.config.d` directory.

   - `hpbtlanconf` (`btlan`)

   - `hpbase100conf` (`btlan3`)

   - `hpgsc100conf` (`btlan4`)

   - `hppci100conf` (`btlan5`)

   - `hpsppci100conf` (`btlan6`)

2. For each of the files found in Step 1, use the command `ioscan –kfC lan` to find out if

---

corresponding hardware is present.

Driver Name: File Name

a. `btlan` -- `hpbtlanconf`
b. `btlan3` -- `hpbase100conf`
c. `btlan4` -- `hpgsc100conf`
d. `btlan5` -- `hppci100conf`
e. `btlan6` -- `hpsppci100conf`

If hardware is present, check to see if any LAN interface is configured. For example, the configuration for `hpbtlanconf` (`btlan`) would show as follows:

HP_BTLAN_INTERFACE_NAME[0]=lan1
HP_BTLAN_STATION_ADDRESS[0]=0x080009C4686E
HP_BTLAN_SPEED[0]=100HD

3. Create the file `hpbtlanconf.merge` in the directory `/etc/rc.config.d` as shown below:

```
#############################################################################
(#) hpbtlanconf
# hpbtlanconf: contains config values for HP PCI/HSC 100Base-T interfaces
#
# HP_BTLAN_INTERFACE_NAME    Name of interface (lan0, lan1, ...)
# HP_BTLAN_STATION ADDRESS   Station address of interface
# HP_BTLAN_SPEED             Speed and duplex mode
#
# The interface name, major number, card instance and ppa may be
# obtained from the lanscan (1m) command.
# The station address and speed are set through the landamin command.
#############################################################################

HP_BTLAN_INTERFACE_NAME[0]=
HP_BTLAN_STATION ADDRESS[0]=
HP_BTLAN_SPEED[0]=

#############################################################################
# The HP_BTLAN_INIT_ARGS are reserved by HP. they are NOT user changeable.
#############################################################################

HP_BTLAN_INIT_ARGS="HP_BTLAN_STATION ADDRESS HP_BTLAN_SPEED"

# End of hpbtlanconf configuration file
```

4. As described in steps a, b, and c below, merge every LAN interface (as shown to be configured in Step 2) into the file `hpbtlanconf` in the `/var/adm/sw/save_custom/UNIFIED_MER` directory.

a. First, note that a set of three parameters is required for each LAN interface:

- `HP_BTLAN_INTERFACE_NAME`,
- `HP_BTLAN_STATION_ADDRESS` **and**
- `HP_BTLAN_SPEED`.

b. Second, note that the index value used for the set of three parameters should be unique for each interface. For example:

```
HP_BTLAN_INTERFACE_NAME[0]=lan1
HP_BTLAN_STATION_ADDRESS[0]=0x080009C4686E
HP_BTLAN_SPEED[0]=100HD

HP_BTLAN_INTERFACE_NAME[1]=lan2
HP_BTLAN_STATION_ADDRESS[1]=0x080009C4B23C
HP_BTLAN_SPEED[1]=FULL
```

c. Finally, merge all LAN interfaces by creating a set of three parameters with a unique index for each:

- Copy the value of the interface name into the parameter `HP_BTLAN_INTERFACE_NAME`.

- Copy the value of the station address into the parameter `HP_BTLAN_STATION_ADDRESS`.

- If the `SPEED` parameter exists, and if its value (lowercase or uppercase) is 100FD or 100HD, 10FD or 10HD or auto_on, copy the value into the `HP_BTLAN_SPEED` parameter.

  However, if the `SPEED` parameter does not exist, and the `DUPLEX` parameter has been set and its value (lowercase or uppercase) is FULL or HALF, copy the value to `HP_BTLAN_SPEED` parameter.

  If neither the `SPEED` nor `DUPLEX` parameter values have not been set, do not put any value for the `HP_BTLAN_SPEED` parameter.

5. If the `hpbtlanconf` configuration file already exists (as `/etc/rc.config.d/hpbtlanconf`), copy it to `/etc/rc.config.d/hplanconf.obsolete`.

6. Move the file `hpbtlanconf.merge` to `hpbtlanconf` by using the command

   `mv -f /etc/rc.config.d/hpbtlanconf.merge /etc/rc.config.d/hpbtlanconf`

7. Restart the update process. Answer `NO` (`N`) at the prompt, `Do you want to proceed in merging the configuration file into one? [Y|N]`.

HSC and PCI 1000Base-SX/T (Gigabit Ethernet)

updated for March 2002

With the March 2002 release, the HP rp7410 server core 1000Base-T product is now supported. Customers should download the latest 11i release for the most recent version of Gigabit Ethernet.

updated for December 2001

The December 2001 release supports the core 1000Base-T/SCSI card (A6794A) for HP-UX 11i, to pre-enable future systems (64-bit only). (The driver, `igelan`, is available on software bundle `GigEther-01`.) The card is factory-installed, with a maximum of two cards. The card supports only MC/ServiceGuard in this release. Online files are located at `/opt/networkdocs`.

new at 11i original release

The PCI 1000Base-T card allows HP 9000s to connect to IEEE 802.3ab standard networks over Cat 5 or Cat 5E UTP copper cable. (The driver, `gelan`, is available on software bundle `GigEther-00`.) The card supports HP Auto-Port Aggregation, MC/ServiceGuard, and LAN Monitor. The card operates at 10, 100, or 1000 Mbit/s and

supports auto-negotiation and auto-sensing. There is no 1000Base-T card for HSC backplanes.

The HSC and PCI 1000Base-SX/9000 products provide the means for interfacing various types of HP 9000 computers to a 1000Base-SX multimode fiber network.

It is recommended that your system have at least 128 megabytes of memory when using this product.

The PCI Gigabit Ethernet cards are for use on:

- V2200, V2250, V2500, V2600, L1000, L2000, A180, N4000, and rp7400 servers running the HP-UX 11i operating system, and

- B-, C-, and J-Class workstations running the HP-UX 10.20, 11.0, and HP-UX 11i operating systems.

NOTE   For information on doing online addition and replacement of Gigabit Ethernet cards, refer to "Managing PCI Cards with OLAR" in the *Configuring HP-UX for Peripherals* manual.

The HSC card is for use on D-, K-, and R-Class servers running HP-UX 10.20, 11.0 and HP-UX 11i, with the following exceptions:

- The A4925A HSC 1000Base-SX card is *not* supported on D210, D220, D230, D310, D320, and D330 systems.

- The A4925A card *is* supported on the D250, D260, D270, D280, D350, D360, D370, D380, D390, R380, and R390 systems.

Please check with your HP representative for a definitive list of HSC card-supported systems.

Compatibility Issues

- The Gigabit Ethernet LAN software is for use with only the following protocols: TCP/IP, ARPA, and NFS

- The HSC and PCI 1000Base-SX/9000 cards operate only at 1000 Mbps. They do not interoperate with 100Base-FX cards.

- The PCI 1000Base-T card operates at 10, 100, or 1000 Mbps. Only full-duplex mode is supported at 1000 Mbps. The PCI 1000Base-T card and link partner (for example, a switch) must be set to autonegotiation to run at 1000Mbps.

- If using Jumbo Ethernet frames, ensure that all switches in the data path support the Jumbo frame size and that both ends of the connection support Jumbo frames. The 1000Base-T card supports Jumbo frames only at 1000 Mbps.

- On A180, B132, B132L, B160, B160L, B180, and B180L platforms running HP-UX 10.20, only one Gigabit Ethernet adapter is supported. (HP-UX 11.0 on these workstations supports multiple adapters.)

- MC/ServiceGuard is not supported on A-180 servers running HP-UX 10.20.

Documentation

For detailed information on advanced features and troubleshooting of the Gigabit Ethernet products, see *Using PCI 1000Base-T and HSC/PCI 1000Base-SX (Gigabit Ethernet)* available on the Instant Information CD or under "Networking and Communications" at `http://docs.hp.com`.

For information on installing and configuring the Gigabit Ethernet products, see *PCI 1000Base-T and HSC/PCI 1000Base-SX/9000 Quick Installation* available in the `/opt/networkdocs` directory and at `http://docs.hp.com`.

The following release notes are also located in the `/opt/networkdocs` directory:

- `gigether_relno.pdf`
- `gigether-01_relno.pdf`

Fibre Channel Tachlite Driver

new for December 2001

Now available is the new 2Gb, single port Fibre Channel Host Bus Adapter. The A6795A PCI Tachyon XL2 adapter operates at bus speeds of up to 66 MHz with 64 bits bus width and can transfer data at a 1Gb or 2Gb transfer rate. The adapter has auto-speed negotiation that enables it to negotiate for the highest speed (1Gb or 2Gb) that is common to the port and to the infrastructure connecting the port. Using the `fcmsutil` diagnostic tool, you can display link speed and Vital Product Data (VPD) information programmed on the adapter.

The A6795A adapter operates on the following platforms running HP-UX 11.00, 11i, or later Operating Environments:

- A400, A500, L-class, N-class, and Superdome servers
- J6000 workstations

The A6795A adapter will support all of the fabric topologies currently supported by the existing Tachlite adapter. Boot support for the first release will be only on the N- and L-class systems.

Summary of Changes

- As of December 2001, the HP-UX 11i Fibre Channel driver version is B.11.11.09.

  Although this driver is pre-installed with each of the HP-UX Operating Environments, you can also obtain it from the December 2001 application release CD. The Fibre Channel 11i driver is a common driver for the following Fibre Channel Host Bus Adapters:

  — A5158A PCI 1Gb Tachyon TL adapter

  — A6684A HSC Tachlite adapter

  — A6685A HSC Tachlite adapter

  — A6795A PCI 2Gb Tachyon XL2 adapter

- HP has created a new Fibre Channel Storage Network Industry Association (SNIA) Host Bus Adapter (HBA) Application Programming Interface (API). The SNIA HBA API is a C library interface that provides both a common HBA API library for all HP-UX HBA vendors and a vendor-specific API library for HP's Tachyon/Tachlite Fibre Channel adapters. The API software is part of the driver and is installed at the

same time. The intended users of this API are programmers who want to create applications for managing SANs networks.

- As of December 2001, you can use the HSC Tachlite adapters (A6684A and A6685A) on the following additional platforms:

— D320
— D330
— D220
— D230

The HSC Tachlite adapters (A6684A and A6685A) will also support boot on K260, K360, K460, D390, R390, K370, K570, and K580.

- Changes have also been made to the Fibre Channel online diagnostic tool, `fcmsutil`, to support the A6795A adapter. When you use the command, `fcmsutil/dev/tdx` (where x is the number assigned to the adapter), the following additional information will display:

— Link Speed (1Gb or 2Gb)
— Dino Present on Card
— Maximum Frame Size
— Driver Version

By using the command, `fcmsutil/dev/tdx vpd`, you can display Vital Product Data (VPD) information on the adapter. This information includes the following:

— Product description
— Part number
— Engineering date code
— Part serial number
— Miscellaneous information
— Manufacturing date
— Check sum
— EFI version
— Asset Tag

Documentation

Effective December 2001, the following documents (available at **http://www.docs.hp.com**) have changed or have been added:

- *Hewlett-Packard Fibre Channel Mass Storage Adapters Manual*
- *HP Fibre Channel Fabric Migration Guide*
- *HP FC SNIA HBA API Programmer's Guide* (new)
- *fcmsutil* (1M) manpage

updated for September 2001

The Tachlite driver enables the HSC Tachlite adapters (A6684A and A6685A), used in the K-, R- and D-Class servers, to have the same fabric capabilities as the PCI Tachlite adapter.

For a current list of driver versions, go to the ITRC web site, **http://itrc.hp.com**. Click *navigate knowledge trees and response center FAQs*. After logging in (if necessary), scroll to the Hardware area and navigate the links to *networks/mass storage connectivity/fibre channel/technical support*. Once you are on the Fibre Channel page, select the Tachlite driver table in the "at a glance" box.

Summary of Changes

- Boot support for the A6685A and A6685A HSC adapters on selected platforms will be available later in 2001. For the latest information on boot support for these adapters and other support information, see the web site mentioned above.

This version of the driver continues to support the following features:

- Fabric zoning and cascading

- Fabric boot and dump for supported platforms (PCI A5158A only)

- Supported configurations

  — Arbitrated private loop
  — Direct fabric attachment
  — QuickLoop (translative mode)
  — Public Loop

Documentation

For more information on supported Fabric topologies, see the *HP Fibre Channel Fabric Migration Guide* located at `http://docs.hp.com`.

The following documents have been revised:

- *HP Fibre Channel Fabric Migration Guide*
- *HP Fibre Channel Mass Storage Adapters Manual*
- *fcmsutil* (1M) man page

HP RAID 4Si Driver

HP RAID 4Si (A5856A) combines multiple small, inexpensive disks into an array that appears as a single logical unit or drive. The data is distributed across the disks in order to provide a method for data recovery or reconstruction in the event of a drive failure.

updated for March 2002

A new firmware version (U.01.06) for the RAID 4Si controller is available as a patch from the HP IT Resource Center (ITRC), at `http://itrc.hp.com`. This version supersedes the version (U.01.04) that is currently shipping with the HP RAID 4Si software.

As always we **strongly recommend** that you upgrade to the latest firmware at your convenience.

updated for June 2001

Summary of Changes

- Boot support (new feature)

  — Ability to boot over the RAID controller.

  — The customer can install the OS onto the logical drives configured on the RAID controller.

  — There is an option now in Ignite-UX called "Configure A5856A" which brings up the configuration tool called IRM that helps in configuring logical drives.

- OLA/R (new feature)

  — Ability to perform Online Addition/Replacement of the RAID controller.

— The customer can do an OLA/R of the RAID controller using SAM or the `rad` command.

- New firmware revision (U.01.04) of the RAID controller includes the following new features and fixes since earlier released version (U.01.01)

  — Provides support for disk firmware download.

  — Provides capability to handle NVRAM vs. Disk mismatch on replaced boot cards.

  — Provides fix for a problem which causes logical drives to become unavailable when a reconstruction is started from IRM (for example, when changing RAID levels).

  — Provides fix for I/O hang due to cache deadlock.

  — Provides fix for the problem of HOTSPARE sometimes not rebuilding when the disk that is part of the logical drive is pulled out.

  — The firmware image `a5856a_01_04.img` is present in the `/opt/raid4Si/firmware` directory on the system on which you install the RAID software. A help text file, `a5856a_u_01_04.txt`, for installing the image is available in the same directory.

- Diagnostic improvements

  — You can now make use of the new `-v` option of the `irdiag` command. Also, you can use the `irdfd` command to perform disk firmware downloads and the `irconcheck` command to perform consistency checks.

Documentation

For more information see the *HP A5856A RAID 4Si PCI 4-Channel Ultra2 SCSI Controller Installation and Administration Guide* at **http://docs.hp.com** and the following manpages:

- *irdfd* (1M)   (new)
- *irconcheck* (1M) (new)
- *irdisplay* (1M)   (revised)
- *irmd* (1M)   (revised)
- *irdiag* (1M)   (revised)
- *irm* (1M)
- *i2outil* (1M)

HyperFabric2 PCI Fiber Adapter

**new for
June 2001**

HyperFabric is a high-speed network link that runs on various HP 9000 servers and workstations. HyperFabric supports the IP network protocol stack, including TCP/IP, UDP/IP, and NFS.

Summary of Changes

- The A6386A HyperFabric2 PCI fiber adapter is now supported on HP-UX 11i on these HP 9000 systems:

  — A400, A500, L-, N-, V-Class, Superdome, and rp8400 systems

  — B1000, B2000, B2600, C3000, C3600, C3700, J5000, J5600, J6000, J6700, and J7000 workstations

- The `clic_ping` command is replaced by the `clic_probe` command

Documentation

See the *HP HyperFabric Release Note*s for Application Release 0601 (part number B6257-90029) for more information.

Performance Enhancements to the HSC FDDI Driver

HP-UX 11i includes code that enhances the HSC FDDI driver. These driver modifications increase performance of the FDDI link by up to 20%.

The performance enhancement was done by making the driver MBLK-based, whereas previously it was MBUF-based. Other code-path and function-call reductions have further improved performance and scalability. These changes are not directly visible to the user and have no effect on current documentation or support.

EISA Interface Cards Are Not 64-bit Compatible

HP-UX 11i 64-bit operation does not include support for EISA interface cards; however, they are supported on the 32-bit operating system. System configurations that include the following EISA interfaces cannot be configured to run HP-UX 11i 64-bit:

Table 5-1   EISA Interface Cards

Card	Description	Comments
J2482A	8-port MUX	
J2483A	64-port MUX	
A2679A	Single-Ended SCSI-2 interface	
A3658A	100BaseT adapter	Discontinued May 1999.
J2165A/B	802.5 Token Ring	J2165 A – discontinued. J2165B – active.
A3402A	Combined 10BaseT/100VG adapter	Discontinued October 1999.
J2815A	Dual-port X.25	
J2220A	SNA PlusLink adapter	Renamed as J2794A.
J2794A	X.25/ACC, SNAPlus/ACC adapter	
A3659A	FDDI (Single or Dual-Attach) adapter	Discontinued August 1999.
25560A	HP-IB Interface	Discontinued March 1998.
25567B	LAN/9000	Discontinued May 1998.
A4031A	Fibre Channel	Discontinued.
A4308A/B	100BT LAN	Discontinued.
B5502AA/BA	FDDI LAN	
J2159A	X.25 Adapter	Discontinued.
J2645AA	100 VG Any LAN	Discontinued.
J2802B	ATM	Discontinued.
J2720AA/BA	SNA Link	

Known Problems with HP-PB and EISA 10/100Base-TX/9000

The EISA and HP-PB 10/100Base-TX drivers do not implement the functionality to change SNMP MIB ifAdmin status.

6   Installation

Cold Install Changed

**new at 11i
original release**

The HP-UX operating system is now delivered as part of an operating environment; see
"The HP-UX 11i Operating Environments (OE)" on page 21 in Chapter 1 for details.

Although HP-UX 11i can be cold-installed with or without an Operating Environment
(OE), HP strongly recommends installing a complete OE. If you choose to install without
an OE, a minimum OS installation must include the following bundles: `HPUXBase32` or
`HPUXBase64`, `HPUXBaseAux`, and `OnlineDiag`.

CAUTION

Omitting the `OnlineDiag` bundle may prevent some of your peripheral devices from
working since they require the hardware monitors included with the Online Diagnostics.
Installing or removing individual products in the OE may lead to dependency issues.

The cold install program is used on all HP-UX systems to initialize a system from
scratch. The program is supplied as part of the bootable core/install CD-ROM media set
and is the first interaction in the install process.

The major changes to cold install are in response to the changing media structure and
the new Operating Environments. Key differences from pre-HP-UX 11i cold install
include the following:

- There is no longer a single core/install CD-ROM. Instead, there are two-CD-ROMs
  and with most installations, software will be loaded from both. The user is prompted
  when to insert the second CD-ROM.

- In addition to the base operating system, you can also select:

  — Operating Environments
  — some additional networking drivers previously available only on the Application
     Release media
  — other optional products previously available only on the Application Release
     media

- Other assorted software will always be loaded and will not be de-selectable.

Ignite-UX Changed for 11i

**new at 11i
original release**

Ignite-UX (IUX) is an HP-UX administration toolset to help you install and configure (or recover) HP-UX systems.

The complete Ignite-UX product, with support for HP-UX 11i, 11.0 and 10.20, is available on the first CD of the 11i OE media kit.

NOTE

Ignite-UX will not be installed by default when updating the entire operating system from 10.20 or 11.00 to 11i *even if* the system currently is an Ignite-UX server.

Updating your Ignite-UX server to 11i and Ignite -UX version B.3.4 (or newer) will ensure that you have the latest version of `make_net_recovery` and `make_tape_recovery`. The command `make_net_recovery` has added most of the command line options that used to be only in `make_recovery`. The command `make_tape_recovery` is the replacement for `make_recovery` and allows you to control tape creation on a client system from your Ignite-UX server. (For details on updating an Ignite-UX server, see the *HP-UX 11i Installation and Update Guide*.)

If you are cold-installing a system, you can select Ignite-UX from the software selection screen to be installed at the same time the rest of the OS is installed.

By default, running Ignite-UX will install the core HP-UX 11i files, networking drivers, and the OE. You can deselect the OE but you cannot install an OE subset using Ignite-UX. Also, some of the networking drivers are not deselectable.

NOTE

If you are using Ignite-UX version B.3.3 and experiencing any problems, upgrade to version B.3.5 or later.

Ignite-UX has been updated to version 3.5 in the December 2001 release and version 3.6 in the March 2002 release. For details, see "Ignite-UX (IUX)" on page 98 in Chapter 4.

Impact

Although Ignite-UX, version B, no longer supports the installation of 10.01 and 10.10 machines for 11i, it will continue to support 10.20 and later operating system revisions. However, version B will only operate on HP-UX machines that have at least 64MB of RAM. (Previously, HP-UX machines that had at least 32MB of RAM were supported. Version A is not affected by these changes.)

Table 6-1   Ignite-UX 11i Version B vs Version A Support

	Version B	Version A
IUX Revision Number	B.3.5	A.3.5
HP-UX Releases Supported	10.20, 11.0 or higher	10.01, 10.10 and 10.20

Table 6-1   **Ignite-UX 11i Version B vs Version A Support**

	Version B	Version A
HP-UX Releases Supporting IUX Servers & Systems	11.0 or higher	10.01, 10.10 and 10.20
Minimum Memory Size for IUX Servers & Systems	64MB	32MB

Since Ignite-UX server support for HP-UX B.10.01 and B.10.10 has been obsoleted, this functionality will be removed from your system upon updating Ignite-UX to the 11i version. If you wish to continue serving B.10.01 and B.10.10 clients, you should set up a separate Ignite-UX server with version B.2.5.136.

Additional information regarding Ignite-UX is described earlier under "Workstation Tuned Kernel Parameters" on page 59.

Documentation

The *Ignite-UX Administration Guide* has been updated for HP-UX 11i, and is available on the HP-UX Instant Information CD and on the `http://docs.hp.com/` web site. Another excellent source of information on Ignite-UX is the external web site:

`http://software.hp.com/products/IUX/`

Update-UX Changes

New Update-UX Product

**new for
March 2002**

The `update-ux` script, once a part of the `SW-DIST.SD-UPDATE` fileset, has been moved to the new Update-UX product in the `Update-UX.UUX` fileset. As a result, customers can now install the `update-ux` script and its accompanying `update-ux` manpage on a 10.20 or 11.00 system (using instructions in the *HP-UX 11i Installation and Update Guide*).

Documentation

For more information, see the *HP-UX 11i Installation and Update Guide* at **http://www.docs.hp.com**, and the manpage *update-ux* (1M).

Update-UX Command Replaces swgettools

**new at 11i
original release**

You update an existing HP-UX 10.20 or 11.0 system to 11i using the new `update-ux` command. Beginning at 11i, this command replaces `swgettools` to perform OS updates. With it, you can also add a new Operating Environment (OE), change an OE, or change the OS word-width from 32- to 64-bit on appropriate systems.

For more information on the `update-ux` command, see the *HP-UX 11i Installation and Update Guide*.

Software Distributor (SD-UX)

**new at 11i
original release**

Many changes have been implemented in Software Distributor since 11.0. The following sections detail the changes. (For additional Software Distributor information, see also Chapter 12.)

Multiple Target Management Capabilities Enabled

As part of the Servicecontrol Manager integration, capabilities previously only available through the OpenView Software Distributor version of SD-UX have been enabled. These include the ability to distribute software to multiple remote targets (individually or together), as well as job management capabilities for scheduling jobs and viewing (local or remote) agent logfiles.

POSIX Enhancements and Exceptions

Software Distributor has been enhanced to meet the IEEE Std 1387.2-1995 standard (also referred to as POSIX 7.2) This affects the behavior of the command line interface and the number of options. (See `/usr/lib/sw/sys.defaults` for a complete list of supported options, their descriptions, and default values.)

Exceptions to the POSIX 7.2 standard are as follows:

- Filesets are not allowed to span media.

- The command `swcopy` has not been modified to copy to tape (`swpackage` can be used for this instead).

- User interaction for tape changes is not handled in the command line of SD-UX.

- The only known exception to the distributed option of the POSIX 7.2 standard is that `swmodify` cannot be run against distributed systems.

Change in swlist Hides Superseded Patches by Default

In 11.0, `swlist` shows all installed patches including superseded ones. However, the 11i default behavior is *not* to show superseded patches. This can be overcome, returning to 11.0 standard behavior, by setting `-x show_superseded_patches=true` on the `swlist` command line or in the defaults files.

64-bit Capability Determined from System, Not /etc/.supported_bits

In 11.0, SD-UX reads the `/etc/.supported_bits` file to map model strings to either 32-, 32/64-, or 64-bit capability. From time to time, synchronization breakdowns between the model command and the contents of `/etc/.supported_bits` have created trouble on 64-bit systems.

To prevent synchronization breakdowns, SD-UX in 11i has changed to get the necessary information directly from the system, rather than using a look-up table.

CD Searched For Only When Requested

With the release of 11.0, SD-UX introduced the automatic discovery and mounting of a CD However, SD-UX always looked for the CD even if that was not what was wanted. This made the start-up of the GUI slower than necessary.

While the functionality is still available in the GUI, SD-UX now only performs this action when you push a new button in the Source Dialog called "Find Local CD."

GUI Streamlined

For 11i, the SD-UX GUI requires fewer confirmations. It has been streamlined to reduce the number of verification and informational popups.

Products Rather than Bundles Shown After Auto-Selection

In 11.0, GUI software selection using "Match What Target Has" or "Automatically select patches for software installed on target" could be confusing because the products in the bundles were not automatically marked for selection in the GUI.

The new 11i behavior provides a product-level view that shows which software has been matched. After inspecting the results of the automatic selection, you can then continue with the installation or change the view back to a bundle level.

See "SD-UX Changes to Patch Installation" on page 128 for details.

Software Groups Added to GUI

The `-f` option to `swinstall`, `swremove`, `swcopy` and `swlist`, which has allowed users to specify collections of software through a file, has been incorporated into the GUI. New actions have been added to allow SD-UX GUI users to save selected software in a "Software Group" (which creates a group) and to select that "Software Group" in subsequent sessions.

Layout Version No Longer Converted Automatically

In 11.0 SD-UX, commands automatically converted Installed Product Database (IPD) and depot catalogs to layout version 1.0 or the layout version specified via the `-x layout_version=...` option on the command line.

In 11i, no SD-UX command will automatically convert the layout version of an existing target, IPD, or depot catalog, even if the `-x layout_version=...` option is specified on the command line. To change the layout version of the IPD or depot, an explicit `swmodify` command is needed to make the conversion.

To convert a 0.8 depot or root to layout version 1.0, use the following:

```
swmodify -a layout_version=1.0 @ <depot_or_root>
```

To convert a 1.0 depot or root back to layout version 0.8 use the following:

```
swmodify -a layout_version=0.8 @ <depot_or_root>
```

As a result of this, the `-x layout_version` option should no longer be needed except when creating a depot that is to be in layout version 0.8 format. Then the `-x`

`layout_version=0.8` option is needed for the `swpackage` and/or `swcopy` commands used to initially create the depot. (By default, SD-UX commands that create depots will create them in layout version 1.0 format.)

Messages Improved or Eliminated

To eliminate unnecessary messages—and to make remaining messages more useful in diagnosing the problem or condition being recorded—many error, warning, and information messages have been removed from or have changed in the SD-UX log files.

Output of swlist Changed

The output of `swlist` has changed in the following ways:

- The `control_file` attribute is no longer displayed at the product- or fileset-levels by default when using the `-v` option, unless the `-l` file-level is also specified on the command line. Also, a new level, `control_file`, has been created to show just `control_file` attributes.

- Listing the product or fileset `control_file` attribute via `-a control_files` is unchanged. This provides performance improvement when listing products and filesets.

- The `source_path` attribute no longer exists in depots and is not displayed with the file-level attributes.

- The command `swlist -l bundle ...` (once used to list non-bundle products if there were no bundles in the source) has been changed so that it now listsnothing.

swpackage Produces Note Vs. Warning

The command `swpackage` no longer produces a warning when an unknown attribute name is encountered. Instead, it now produces a note stating that the attribute is being packaged as a "vendor defined attribute."

Newest Bundle Selected by Default

Previously, when you specified an unqualified bundle name for selection and the bundle name was ambiguous due to multiple revisions, SD-UX printed out an "ambiguous bundle" error message. Now, SD-UX selects the newest version of the bundle by default.

This change in behavior makes bundle selection consistent with what SD-UX does for products and filesets when multiple versions of these are available in the source.

control_utils File Improved

New functionality has been added and defect repair has been done to the `/usr/lbin/sw/control_utils` file. Documentation on the `control_utils` functions can be found at **http://software.hp.com/SD_AT_HP/info.html** under the link Control Script Guidelines (CSG). The `control_utils` library is a collection of shell functions which can help packagers produce better software packages.

New Environment Variable, SW_COMPATIBLE, Created

A new environment variable, `SW_COMPATIBLE`, has been created for use during the execution of a verify script that is called by the `swverify` command. If the software being considered is compatible with the system it is installed on, the variable will be set to `TRUE`. If it is incompatible, it will be set to `FALSE`. This new variable will help control-script writers determine if installed software is incompatible and should be removed from a system.

SD-UX Changes to Patch Installation

The SD-UX patch installation paradigm has changed for HP-UX 11i. To install patches on HP-UX 10.x systems, HP recommended that you use the `match_target` (Match What Target Has) option to match patches to the target. However, 10.x SD-UX could not identify specific software as patches.

With HP-UX 11i, SD-UX can recognize patches based on their "internal attributes." This provides more control over patch management than in previous releases.

patch_match_target

The `match_target` option still functions, but no longer matches patches to targets. With 11i, setting the `patch_match_target` option to `TRUE` automatically selects the latest patches that correspond to software on the target. The default setting is `patch_match_target=false`.

NOTE	Since the `patch_match_target` and `match_target` options cannot both be set to `TRUE` in the same `swinstall` command, you should use the `match_target` option to update from HP-UX 10.x, but use `patch_match_target` to install new patches on systems already running HP-UX 11i. (This option selects patches that apply to software already installed on an 11i system.

The 11i `autoselect_patches` option (`TRUE` by default) automatically selects patches to match software selected for installation. It lets you install patches at the same time you install base software. In addition to the base software selected by the `match_target` option, the `autoselect_patches` option provides the means for selecting appropriate patches during the update process.

Patch Filtering

With 11i, you can more interactively manage your patch process via Patch Filtering. By using the `category_tag` and `patch_filter` options plus various version qualifiers, you can select patches based on pre-defined criteria.

Category Tag Information

With 11i, SD-UX category tags are used to identify types of patches. These category tags can be used to select various patches for installation. The category tags include the following:

```
general_release
critical
hardware_enablement
defect_repair
```

```
corruption
enhancement
memory_leak
panic
halts_system
```

By specifying the category (c) tag in the SD-UX version specification, you can select all patches that contain that specific category tag. For example, using the SD-UX command line interface, you can select all patches in the depot that correspond to currently installed software (and that contain the category tag "critical") by entering the following:

```
swinstall -x autoreboot=true -x patch_match_target=true \
-x patch_filter="*.*,c=critical" -s depot_name
```

By using the pipe (|) function, you can combine category tags. For example, to install patches that are either critical OR hardware_enablement, enter the following:

```
swinstall -x autoreboot=true -x patch_match_target=true \
-x patch_filter="*.*,c=critical|hardware_enablement" \
-s depot_name
```

To preview the patches that are selected for a particular swinstall session, the -p (preview) option can be used. The -p option will cause SD-UX to analyze the installation, then exit (that is, the actual installation will not be performed). Look in the /var/adm/sw/swinstall.log file to determine which patches were selected.

To use category tags with the SD-UX Graphical User Interface, do the following:

Step 1. Under the Options menu, select Manage Patch Selection.

Step 2. Then select the box labeled "Automatically select patches for software installed on the target."

Step 3. In the "Filter..." text field, add the desired filters to the *.*. (For example, to select only the critical patches, the Filter... field would appear as *.*,c=critical. Likewise, to install all the patches that are critical OR hardware_enablement, the Filter...field should appear as *.*,c=critical|hardware_enablement.)

Step 4. Select OK.

Clicking the Filter... button will display a list of the predefined category tags already formatted for use in the Filter... field. Selecting the desired category tag from this list and then selecting OK will add that, and only that, category tag to the Filter... field. (Also shown under the Filter... field is the list of all category tags found in the source depot.)

The list of patches that were selected for install can then be viewed by double-clicking on the bundle in the main SD-UX window. You can then deselect any patches that you may not want to install. (Be careful not to break any documented patch dependencies.) Continue with the install (analysis) as with any other patch installation.

For more complete information on 11i Interactive Patch Management, refer to the manual *Software Distributor Administration Guide*, part no. B2355-90699.

Documentation

The *Software Distributor Administration Guide* has been extensively updated for accuracy and completeness (including many new examples) for HP-UX 11i and is

available on the HP-UX Instant Information CD and on the `http://docs.hp.com` web site. Another excellent source of information on SD-UX is the SD-UX web site:

`http://software.hp.com/SD_AT_HP/`

set_parms Enhanced

new at 11i original release

The `set_parms` program is a GUI/TUI interface that normally runs only the first time any HP-UX system is booted after installation if hostname/networking information has not been set up in advance.

For HP-UX 11i, `set_parms` has been enhanced to allow you to select which networking interface to set up. In prior releases, `set_parms` would pick the lowest numbered LAN interface to configure in the absence of any other information. This was often the wrong choice, especially when FDDI interfaces or other optional interfaces were present on the system, forcing users into extra steps to configure the system properly.

This change will allow you to pick the LAN interface to be configured in both the case of enabling DHCP (the user picks just after the decision to use DHCP) and in the normal mainline case of setting an IP address (the user picks the interface just before setting the IP address). With this change, no additional configuration steps are immediately needed to get the system operational.

This change does not fix any previous defects.

There is a new manpage, *set_parms* (1M), delivered at HP-UX 11i. However, the program itself is not new.

7 General System Administration and Performance Monitoring

Processor Sets Available on Software Pack

What are Processor Sets?

HP-UX Processor Sets offer a flexible mechanism for managing system processor resources among multiple workloads, users, and departments within an enterprise.

A processor set represents a set of processors grouped together for exclusive access to applications assigned to that processor set. Each application is assigned to a processor set and will run only on processors in the assigned processor set.

Processor sets allow partitioning of the system into multiple Scheduling Allocation Domains, such that workloads running in different processor sets do not contend with one another for system processor resources. This capability makes server consolidation on large systems more efficient.

Resource management based on processor sets is completely hardware platform independent and can be used on any HP-UX 11i multiprocessor system.

Processor sets are integrated with HP-UX Process Resource Manager. A PRM group can now be mapped to a processor set for processor resources rather than processor shares. Please see the *HP Process Resource Manager User's Guide* for further information.

For details on using processor sets on nPartition systems, refer to the *HP System Partitions Guide*.

NOTE   Any application that programatically uses the processor sets system call interfaces must be compiled with the header files installed with the Processor Sets product.

NOTE   In MeasureWare on HP-UX, currently the PRM_SYS group returns invalid data when multiple processor sets are enabled. There is no workaround at this time.

Where to Find Information

For detailed information about the Processor Sets feature see the following documentation:

- The Processor Sets Manual Pages:

  — *pset_create* (2)   — *pset_getattr* (2)
  — *pset_assign* (2)   — *pset_setattr* (2)
  — *pset_bind* (2)   — *psrset* (1M)
  — *pset_destroy* (2)   — *pstat_getpset* (2)
  — *pset_ctl* (2)   — *pthread_pset_bind_up* (3T)

- The Processor Sets documentation at the HP Documentation web site:

  `http://docs.hp.com/hpux/11i/index.html#Processor%20Sets`

- The *HP Process Resource Manager User's Guide* available on the September 2001 HP-UX 11i Instant Information CD and on the Web:

  — **http://docs.hp.com**

  — **http://www.hp.com/go/hpux**

The uname Command Outputs B.11.11

**new at 11i
original release**

The `uname` command for identifying the version of HP-UX that your system is running will return the following version name on an 11i system:

`B.11.11`

To guarantee compatibility on current and future platforms, you can use the new interfaces to `getconf` and `confstr` to retrieve unique machine identifiers. For more information, see the manpages and in Chapter 13, "Programming," in the section "Machine Identifier Changes to confstr" on page 227. See the *getconf* (1) and *confstr* (3C) manpages for more information on these commands.

You can also determine what versions of software and patches are installed on your system by using the following `swlist` command:

`$ /usr/sbin/swlist -l bundle`

New Option for top

A new `-h` option to the `top` command is provided to suppress the individual CPU state information for multiprocessor systems. If the `-h` option is specified, only the average of all CPU activities will be displayed.

The change enables `top` to display more processes on a standard (80x24) screen without the screen being dominated by state information of individual CPUs.

The *top* (1) manpage has been updated to include the new option.

Changes to System Administration Manager (SAM)

The `/usr/sbin/sam` command starts a menu-driven System Administration Manager (SAM) program that makes it easy to perform system administration tasks with only limited, specialized knowledge of the HP-UX operating system.

SAM has been enhanced to support new devices and features in the following areas within its interface: Disks and File Systems, Kernel Configuration, Network File Systems and Interface Cards, Peripheral Devices, Printers and Plotters, and Terminals and Modems. HP-UX and SAM discontinued support of NFS Diskless as of HP-UX 11.0.

Disks and File Systems Area

updated for June 2001

- Added support for the new HP SureStore Virtual Array 7100 (A6188A) disk array and will now identify Hitachi XP48 and Array LUNs.

new at 11i original release

- Added support for the HP SureStore Disk Array FC60. This allows you to manage larger volumes of mass storage. This support was also added in a patch to 10.20 and 11.0.

- Added support for the HP SureStore Disk Array XP256. This allows you to manage larger volumes of mass storage (up to 11 Terabytes). This support was also added in a patch to 10.20 and 11.0.

- Added recognition of disks being used by HP VERITAS Volume Manager (VxVM) and their associated Disk Groups and Volumes. You can launch the HP VERITAS Manager Storage Administrator (`vmsa`) tool from within the SAM Disks and File Systems area.

Kernel Configuration

- Added support for dynamic tunables. These changes allow the user to tune dynamic tunables using SAM and have the values take effect immediately. The user will see differences in modifying tunable parameters and in applying tuned parameter sets. When modifying a tunable, the value/formula selected by the user will take effect immediately if the tunable is dynamic. In applying tuned sets, the entire tuned set will take effect immediately if every tunable in the set is dynamic. Otherwise, the tuned set won't take effect until the kernel is rebuilt and the system is rebooted.

- Added support for kernel logging. Kernel logging is a new feature for 11i. Through SAM, system administrators will be allowed to modify options associated with this feature, such as turning logging on and off, for many of the kernel subsystems (such as VxFS File System, virtual memory, etc.) The subsystems are capable of recording information at various levels ranging from simple status messages to catastrophic failures. These logs can be read through SAM, either from kernel memory or from a file on disk.

Network File Systems

- Support for NFS over TCP/IP. NFS supports exporting a file system using the TCP/IP protocol. Accordingly, the Network File Systems area in SAM has been enhanced to support this new NFS feature. Now, the user can choose between TCP and UDP protocols to export file systems.

Network Interface Cards

- Added support for Gigabit Ethernet card. The system administrator can configure, modify the supported parameters, and also unconfigure Gigabit Ethernet cards. This new SAM support is provided to both 1000Base-T and 1000Base-SX Gigabit Ethernet cards. This support is also available on 10.20 and 11.0 through patches.

- Added support for 100-BaseT cards. The system administrator can configure the 100-BaseT cards, modify any of the supported parameters, and unconfigure the 100-BaseT cards. This support is also available on 10.20 and 11.0 through patches.

- Added support for Independent Hardware Vendor (IHV). The system administrator can configure network interface cards manufactured by third party vendors that conform to the exported SAM Networking interface. Users can configure, unconfigure, and modify any of the supported parameters of the card.

Peripheral Devices

updated for June 2001

- Added support for PCI Card Online Addition and Replacement (OLAR) on systems with OLAR-capable hardware, for Superdome systems, N-Class, and L-Class systems.

   This change allows you to add or replace PCI cards online without requiring a reboot. SAM shows the I/O slot number and the OLAR driver state (active, suspended, error, not OLAR-capable). Through the `Actions` menu in the `Peripheral Devices / Cards` screen, you can replace a card, turn on the I/O slot attention LED (so you can locate the slot more easily), bring a suspended card online, and view OLAR-specific information about the card and slot. In addition, different instructions are provided depending on whether or not the card is in a slot that is capable of OLAR.

NOTE   SAM patch PHCO_23004 for 11i (included in the HWEnable11i patch bundle and automatically installed) changes the behavior of the PCI card slot LED (Attention Indicator) to conform with the newly implemented PCI SHPC (Standard Hotplug Controller) specification.

Essentially, the meanings for OLAR "FLASHING" (BLINKING) and "ON" LED are reversed to meet the specification.

See "New Attention Indicator Behavior" on page 40 in Chapter 2 for more details.

- Modified cards screen layout to be hierarchical.

   You can open mass storage interface cards (SCSI, fibrechannel) to view the storage devices attached to the card. This enables you to quickly identify devices controlled through a particular card resource and therefore affected if the card were suspended. Because only the attached devices are displayed, the hierarchical view is more

convenient and easier to use then backtracking to the `Devices List` from the `Peripheral Devices` area.

- Added new `Actions` menu item, `Analyze Critical Resources`.

  This item displays a dialog listing all resources, including devices, file systems, device files, and processes that are affected if the selected card fails or is suspended from operation. It will work on systems and cards that are capable of OLAR, as well as those that are not. It determines if any of the resources are critical to HP-UX or SAM (and this would cause either to fail if the resource were lost). This menu item properly accounts for resources that have alternates, backups, or mirrors.

- Added support for DLT device tape densities (DLT_42500_24, DLT_42500_56, DLT_62500_64, DLT_81633_64).

- In the `Tape Devices` and `Backup and Recovery` areas, added support for the following new tape drives:

  — HP C7145/C7146 Autochangers
  — HP Surestore SDKT 220 tape drive

- Added device support for PCI multiplexer needed for `Terminals and Modems` area.

- Added device support for HP SureStore Disk Array FC60 and the HP SureStore Disk Array XP256 (as described in the section "Disks and File Systems Area" on page 133).

System Properties

**new for
June 2001**

- In the `System Properties` area, SAM recognizes new processors, if available on the system, and displays information.

Printers and Plotters

- Added support for the Super I/O Parallel Interface. This allows SAM to recognize this parallel interface and configure a printer for it.

- Updated the software products supported by SAM that allow the user to configure a network printer that has a JetDirect interface card.

- Removed obsolete JetDirect software (`/usr/lib/admin/hpnpcfg`) and added support for the new HP JetDirect Printer Installer (`/opt/hpnpl/admin/hppi`) software. You can now use SAM to configure a network printer using the new HP JetDirect Printer Installer. SAM still supports the HP JetAdmin Software (`/opt/hpnp/admin/jetadmin`).

Terminal and Modems

- Added support for PCI multiplexer cards.

- Added support for modems with hardware flow control.

Documentation Change

- The *sam* (1M) manpage and online help have been updated.

Possible Future Changes

Performance improvements are planned for Card OLAR code in the `Peripherals Devices` area.

In the future, SAM is planning on obsoleting the following:

- `Instruments` section under the `Peripheral Devices` area - HP-IB instruments will no longer be supported.

- Run SAM on `Remote Systems` area - The ServiceControl Manager product will be handling multi-system management.

- The `Process Management` area

- The `Routine Tasks` area

- `Backup and Recover` - The `fbackup` and `frecover` commands will still be available from the command line.

- `Distributed Print Service`

- `Kernel Configuration` - This will be replaced with a new tool

- `Uninterrupted Power Supply` section under the `Peripheral Devices` area

- In the `Performance Monitors` area, the following areas will be removed as SAM provides no added value above the existing commands (listed in parenthesis):

  — `Disk and Terminal Activity (iostat)`
  — `Inter-Process Communication Facility Status (ipcs)`
  — `Process with the Highest CPU Usage (top)`
  — `System Activity (sar)`
  — `Virtual Memory Activity (vmstat)`

Additional SAM Changes

The Guardian Service Processor was introduced for the N4000 mid-range servers with the May 1999 Extension Pack and subsequently on all new servers. (See "Service Processor (GSP or MP)" on page 51 in Chapter 3 for more information.) The new card has a port for the system console, as well as optional ports that can be used to connect terminals, modems, and Uninterruptable Power Supplies (UPS). SAM has been modified to aid in the configuration of these optional ports.

SAM allows you to launch the Partition Manager (`parmgr`), the new system administration tool that supports the initial and ongoing configuration of Superdome systems. See "Partition Manager (parmgr)" on page 35 in Chapter 2 for details.

In this release, SAM allows selecting workstation kernel parameter sets to tune your system for CAE/ME or EE applications. See "Workstation Tuned Kernel Parameters" on page 59 in Chapter 3 for more information.

syslog File Logging Changes for su and login

**new at 11i
original release**

This change will only affect you if you write or use programs or scripts that parse the syslog file.

The format of text messages logged in the syslog file by the `su` and `login` commands has changed slightly. Specifically, su events are now preceded by 'su:' and login events are now preceded by 'login:'. As a result, syslog output is now more consistent with the format of messages generated by other commands. It is also easier for programs to operate that parse syslog output.

Aside from affecting the text logged in the syslog file, this change may possibly impact any programs that parse the syslog file, such as certain security monitoring tools.

Programs that read syslog files looking for su and login events will need to take this change into account.

Process Resource Manager (PRM) Disk Bandwidth Control

This release allows Process Resource Manager (PRM) to report and control disk I/O bandwidth per LVM device. Based on user-configured priorities, PRM will reorder LVM disk queues to enforce the percentage of disk bandwidth a PRM group receives in an I/O-constrained environment. The functionality is documented in the PRM manpages.

HP Distributed Print Service Deprecated

**new at 11i
original release**   The HP Distributed Print Service (HPDPS) print environment is being deprecated at 11i and will be removed in a future release. HPDPS, however, will continue to be supported for HP-UX 11.0 and 10.20 until further notice.

Four options are available for print services in 11i:

1. the LP spooler, which is part of the HP-UX operating system

2. the HP DirectJet 4000 Printing Appliance

3. the HP Document Router

4. the HP/Dazel Output Server

Each option beyond the first is additive and increases in cost. Option 1 will provide base-level print services. Option 2 will spool print jobs from Windows clients and provide Web-accessible print queues, management, and configuration. Option 3 will automate the delivery of documents and provide a Web-based administrative interface for TCP-connected network printers, fax, and email. Option 4 will provide an enterprise-wide information delivery infrastructure to enable reliable document delivery and centralized print management across the entire enterprise.

Options 2 through 4 are available directly from HP; contact your local HP office for further information about functionality, platform support, and prices. Migration from HPDPS to Options 2 through 4 will require reconfiguration and setup of the print environment. Contact HP for assistance in performing this migration.

Option 3, the HP Document Router, most closely replicates the services that were provided by HPDPS although there is not a one-to-one mapping of commands.

HPDPS commands that are being deprecated are listed here for quick reference:

`pdclean`, `pdcreate`, `pddcesetup`, `pddelete`, `pddisable`, `pdenable`, `pdgwcfg`, `pdls`, `pdmod`, `pdmsg`, `pdmsghlp`, `pdpause`, `pdpr`, `pdpromote`, `pdps`, `pdq`, `pdresubmit`, `pdresume`, `pdrm`, `pdset`, `pdshutdown`, `pdstartclient`, `pdstartspl`, `pdstartsuv`, **and** `pdstopd`.

Diagnostics: EMS Hardware Monitors

The Event Monitoring System (EMS) Hardware Monitors allow you to monitor the operation of a wide variety of hardware products.

When monitors encounter failure or other unusual events, they generate messages with Description, Cause, and Action statements which can be used to prevent and reduce downtime caused by hardware failures.

EMS Hardware Monitors are installed automatically with the Support Tools Manager (STM). After installation, monitors must be enabled to begin operation. Most hardware monitors are supplied with a default configuration; additional configuration is optional. A few hardware monitors, such as the Fibre Channel Arbitrated Loop Hub Monitor, have special requirements. See the *EMS Hardware Monitors User's Guide* for details.

EMS Hardware Monitors require minimal maintenance once installed and enabled. Default notification definitions are provided so additional configuration is not necessary. New hardware resources added to the system are automatically included in the monitoring structure.

For more information, see `http://docs.hp.com/hpux/diag`.

Integration with Other Applications

The Event Monitoring System (EMS) Hardware Monitors can be integrated with applications responsible for maintaining system availability, such as MC/ServiceGuard. If configured to do so, they can provide event notification to system management applications such as HP OpenView IT/O and HP Network Node Manager.

The EMS Hardware Monitors use the same EMS framework as the EMS High Availability (HA) monitors, a separate set of monitors available at additional cost.

Some of the hardware monitors for fibre channel products write event information to text logs read by a new Predictive scanner, emsscan, which in turn may send events to the Response Center via On-line Predictive.

Improved ioscan Description Field for PCI Devices

The `ioscan` command displays I/O devices, memory modules, and CPUs in a tabular form for users.

Previously, PCI interface cards were listed in the `ioscan` output by cryptic values in the `ioscan` *description* field. These values have been replaced by PCI device header fields which provide a clearer description for most (common) devices. See the example provided below.

The changes are:

- Dropped "FRU" info. When PCI drivers update the description fields (as the SCSI interface driver `c720` currently does), this gets lost anyway. For more information on the SCSI interface driver, see "SCSI Drivers scsi3 and c720" on page 49 in Chapter 3.

- Added lookup of the class/subclass headers for most PCI specified devices (for example, PCI Ethernet). This provides useful and correct information when a device driver is **not** loaded or does not update the description.

- Dropped "PCI Bus Bridge" to shorten the CDIO description string. Epic CDIO would result in "EPIC Bridge." "PCItoPCI" is the name of the new CDIO for PCI-to-PCI Bridges (PPBs).

This is the old output:

```
H/W Path          Class            Description
8                 bc               Pseudo Bus Converter
8/0               ba               PCI Bus Bridge - GSCtoPCI
8/0/1/0           ba               PCI(10110024)
8/0/1/0/4/0       lan              PCI(10110019)
8/0/1/0/5/0       ext_bus          Ultra Wide SCSI
8/0/1/0/5/0.1     target
8/0/1/0/5/0.1.0   disk             HP C2247WD
```

Here's the new output:

```
H/W Path          Class            Description
8                 bc               Pseudo Bus Converter
8/0               ba               GSCtoPCI Bridge
8/0/1/0           ba               PCItoPCI Bridge
8/0/1/0/4/0       lan              PCI Ethernet (10110019)
8/0/1/0/5/0       ext_bus          Ultra Wide SCSI
8/0/1/0/5/0.1     target
8/0/1/0/5/0.1.0   disk             HP C2247WD
```

The `ioscan -F` option provides the same as well as additional information, separated by colons for parsing by scripts. This remains unchanged. Scripts can (and should) continue to use the `-F` option.

If scripts are parsing this output, the most significant "keys" remain the vendor/device ID (hex digits) and "PCI" string.

instant Capacity on Demand (iCOD)

iCOD and Pay Per Use Version B.04.00

updated for September 2001

Hewlett-Packard's instant Capacity on Demand (iCOD) and pay per use (PPU) software products provide the ability to increase or decrease processor capacity instantly on (specified) enterprise servers. iCOD and PPU software products are a part of the HP Utility Pricing Solutions program.

There are two products associated with HP's Utility Pricing Solutions program:

- **iCOD** (product number: B9073AA): You initially purchase a specified number of activated processors and pay a right-to-access fee for a specified number of inactive processors. After you activate an inactive processor, you pay an enablement fee.

- **PPU** (product number: T1322AA): You pay only for the average processor usage, based on a fixed/variable model. Pay per use is a new product and is available in September 2001.

Locating iCOD Information

For versions previous to iCOD 4.0, the Release Notes contain two types of information:

- what is new, changed, or obsolete in the most recent release of iCOD as compared with previous releases

- instructions on how to configure and manage iCOD systems

For iCOD version 4.0, the Release Notes contain only information on what is new, changed, or obsolete as compared to previous versions of iCOD. The instructions on how to configure and manage iCOD systems are contained in the *Instant Capacity on Demand (iCOD) and Pay Per Use (PPU) User's Guide for version B.04.00.*

- **iCOD Release Notes**

  You can find the *iCOD Release Notes for version B.04.00* in the following locations/media:

  — For the most recent information, visit the following HP documentation web site and search for "icod release notes":

    `http://docs.hp.com`

  — Printed (part number B9073-90014) - order from the following HP web site:

    `http:/www.hp.com/products1/unixservers/solutions/icod/index.html`

  — September 2001 Instant Information CD

  — In the iCOD 4.0 software product located in:
    `/usr/share/doc/icodRelNotes.pdf`. Note: this is an early version of this document. For the most recent information, see the documents listed above.

- **iCOD/PPU User's Guide**

  You can find the *Instant Capacity on Demand (iCOD) and Pay Per Use (PPU) User's Guide for version B.04.00* in the following locations:

— For the most current information, visit the following HP documentation web site and search for "icod users guide":

  `http://docs.hp.com/`

— Printed (part number 5185-6531) - order from the following HP Web site:

  `http://www.hp.com/products1/unixservers/solutions/icod/index.html`

— September 2001 Instant Information CD

— In the iCOD 4.0 software product located in:
  `/usr/share/doc/icodUserGuide.pdf`. Note: this is an early version of this document. For more current information, see the documents listed above.

- **Manpages**

  — *icod* (5)
  — *icod_modify* (1M)
  — *icod_notify* (1M)
  — *icod_stat* (1M)

For details on using iCOD and pay per use (PPU) software on nPartition systems, refer to the *HP System Partitions Guide*.

update for June 2001   The iCOD product was removed from the HP-UX Operating Environments at the HP-UX 11i June 2001 Release. You can find iCOD on the Support Plus media and at the following web site:

`http://software.hp.com.`

8   Process, Threads, Memory, and Kernel Parameters

HP-UX Gang Scheduling

**new at 11i
original release**

This release includes the ability to "gang schedule" MPI (Message Passing Interface) applications and multi-threaded processes. The gang scheduler permits a set of MPI processes, or multiple threads from a single process, to be scheduled concurrently as a group.

Only applications using the HP-UX 11.0 (or later) 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.

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.

The HP-UX gang scheduler is enabled or disabled by following environment variable:

```
MP_GANG [ON]  | [OFF]
```

Setting `MP_GANG` to `ON` enables gang scheduling; setting `MP_GANG` to `OFF` disables it. No action is taken if `MP_GANG` is not set, or if it is set to an undefined value. After the `MP_GANG` environment variable is set to `ON`, any MPI or pthread application executing and finding this variable will enable gang scheduling for that process.

You also can launch a program with gang scheduling enabled for it by using the `/usr/bin/mpsched -g...` command. (See "HP-UX SCA Process and Memory Management" on page 149 for details.)

Thread and process priorities for gangs are managed identically to timeshare policy. That is, the timeshare priority scheduler determines when to schedule a gang and adheres to the timeshare policies.

Although it is likely that scheduling a gang will preempt one or more higher priority timeshare threads, over the long run the gang scheduler policy is generally fair. All threads in a gang will have been highest priority by the time a gang is scheduled. Because all threads in a gang must execute concurrently, some threads do not execute when they are highest priority (the threads must wait until all other threads have also been selected, allowing other processes to run first).

Refer to the *mpsched* (1) or *gang_sched* (7) manpages for details about HP-UX gang scheduling.

Kernel Threads vs. CMA Threads

**new at 11i
original release**

The CMA threads (`libcma`) package, which is POSIX P1003.1a (Draft 4) compliant, is based on Concert Multi Thread Architecture (CMA). CMA is a user-level threads package in which thread scheduling and synchronization are handled within the user space without the kernel's assistance.

CMA threads have been deprecated (slated for future obsolescence) at HP-UX 11i. This development environment will not be shipped in a future HP-UX release. Also, there is no plan to release native IA-64 CMA threads on IA-64 platforms. Therefore, HP now strongly recommends that you use the currently supported kernel threads libraries and development tools. Thus, applications using CMA threads should start migrating to kernel threads.

Multi-threading is also supported in the HP-UX kernel at 11i and is known as kernel, POSIX or 1x1 threads. This kernel threads implementation, `libpthread`, is compliant with the approved POSIX 1003.1c (POSIX.1-1996 Draft 10) standard and will be replacing the CMA threads package.

The kernel threads implementation allows the application to take advantage of multiple processors in the system to parallelize execution of threads.

Compatibility Issues

It is expected that all existing CMA applications will continue to run on future releases. However, CMA applications may have to be ported to HP-UX POSIX threads in future releases, including those supporting IA-64, as there are differences in certain APIs between CMA threads and HP-UX POSIX threads.

As a POSIX standard, the kernel thread implementation facilitates better application portability on to POSIX-compliant vendor platforms.

Documentation

A white paper *Porting DCE Threads Programs to HP-UX 11.0 POSIX Threads* is available to help move from CMA threads to kernel threads (`http://devresource.hp.com/Docs/TechPapers/PortThreads.html`). The HP-UX Software Transition Kit (STK) for 11.0/11.x/IA-64 is also available (`http://devresource.hp.com/STK/`) for assistance. The STK contains documents that explain how to perform a source code or system transition. For more information see `http://devresource.hp.com/STK/toc_trans.html`

Large Private Data Space

**new at 11i
original release**

An additional 1 to 2GB of private address space is now supported for 32-bit programs (if enabled on a per process basis), at the expense of shared memory address space. This change increases the amount of private data space available for a process.

New Options

Two new options have been added to the `chatr` command that allow the user to control whether the 3rd quadrant (the 1GB of address space from `0x80000000-0xBFFFFFFF`) and the 4th quadrant (the 1GB of address space from `0xC0000000-0xFFFFFFFF`) of a process are part of the processes private address space or are shared with other running processes. Previously, the 3rd and 4th quadrants were dedicated for shared object usage. For example, System V shared memory and memory mapped files using a shared mapping (`MAP_SHARED`).

The new options are as follows:

- `+q3p` <**enable/disable**>

- `+q4p` <**enable/disable**>

See the *chatr* (1) manpage for more details.

In order to use this new feature, the `maxdsiz` kernel configurable variable will need to be increased appropriately. Also, the system will have to enable enough swap space to support processes with large private address spaces.

Compatibility Issues

Processes that enable a private 3rd quadrant (`q3p` processes) will reduce the amount of address space available for shared objects by 1GB. Also, `q3p` processes will not be able to share objects that were created by another, non-`q3p` process, even in the 4th quadrant, unless those objects were created by the non-`q3p` process using the `IPC_GLOBAL` flag (System V shared memory) or `MAP_GLOBAL` flag (`mmap`). If recompiling is not an option, it will be necessary to make *all* processes that share objects with the `q3p` process into `q3p` processes (**chatr +q3p enable <a.out>**).

Processes that enable a private 4th quadrant (`q4p` processes) will have no address space available for shared objects. This means that the process will not be able to use System V shared memory, shared mapped files, etc. Shared libraries will still work, although the kernel will map them as private. Note that a `q4p` process implies that the 3rd quadrant is private also. In other words, the kernel will not execute a process that only enables a private 4th quadrant.

Because the system call gateway page has to remain at address 0xC0000000 for binary compatibility reasons, the data segment cannot be extended past the beginning of the 4th quadrant. Therefore, the `brk()` and `sbrk()` system calls will only allow the data segment to be expanded up to that address.

To take advantage of private address space in the 4th quadrant, memory will need to be allocated using the `mmap()` system call with the `MAP_PRIVATE` option. The system call `malloc()` has been modified to do this automatically. No re-link will be necessary to take advantage of the new `malloc()` for a program that uses a shared version of the C

library. A program that was linked with a non-shared library version of the C library, however, will need to be re-linked.

These changes have no compatibility impacts if the feature is not enabled.

This feature can only be enabled for 32-bit programs running on the 64-bit version of HP-UX. The 32-bit version of HP-UX will silently ignore the request for a private 3rd or 4th quadrant.

Memory Windows

Running without memory windows, HP-UX has limitations for shared resources on 32-bit applications. All applications in the system are limited to a total of 1.75GB of shared memory (2.75GB if compiled as `SHMEM_MAGIC`). In a system with 16GB of physical memory, only 1.75 can be used for shared resources.

To address this limitation, a functional change has been made to allow 32-bit processes to create unique memory windows for shared objects like shared memory.

The memory window for default executables is 1GB.

This allows cooperating applications to create 1GB of shared resources without exhausting the system-wide resource. Part of the virtual address space remains globally visible to all processes, so that shared libraries are accessible no matter what memory window they are in.

Summary of Changes

The following customer-visible changes have been made for memory windows:

- A new kernel tunable, `max_mem_window`, allows you to configure the number of memory windows a system can have.

- A new set of commands and files (`setmemwindow`, `getmemwindow`, `/etc/services.window`) enables you to start applications in different memory windows.

- Three manpages for the new functionality have been created: *getmemwindow* (1M), *setmemwindow* (1M), and *services.window* (4).

See the *Memory Windows in HP-UX 11.0 White Paper* on `http://docs.hp.com` for details.

Compatibility Issues

Incorrect use of memory windows can lead to application failure. Although memory windows can be applied to any application, that does not mean the application is able to run in memory windows. Some interfaces may break when used under memory windows. (Since only the application owner or software provider can certify how and if an application can run under memory windows, HP does not consider this failure a compatibility failure.)

CAUTION   Errors due to incorrect usage may be subtle and normally not associated with memory windows.

In many cases software providers may have already certified their applications with memory windows. Contact HP to see if this is the case.

Configuration

By default, HP-UX ships with memory windows disabled.

To enable memory windows, set the kernel tunable parameter `max_mem_window` to the desired amount. Customers can change this value by placing the desired value in their kernel configuration file. The system must be rebooted for the new value to take effect.

As detailed below, `max_mem_window` represents the number of memory windows beyond the global default window:

- Setting `max_mem_window` to one (1) creates a single memory window to accompany the existing global memory window (or, a total of two memory windows: one default and one user-defined).

- Setting `max_mem_window` to two (2) produces a total of three memory windows: the default and the two user-defined.

- Setting `max_mem_window` to zero (0) leaves only one memory window: the default or global memory window.

What should the value be? That depends on the application requirements and the applications installed on the system. (HP recommends that each ISV/application should document how many windows it intends to use.)

Use of memory windows does not affect the performance of processes. There is no size requirement associated with memory windows. Any machine running HP-UX (32-bit or 64-bit) and any hardware supporting HP-UX release 11i can use memory windows.

HP-UX SCA Process and Memory Management

**new at 11i
original release**

For compatibility reasons, the HP-UX 11i release supports the Scalable Computing Architecture (SCA) programming, locality management, and memory management features that were introduced at HP-UX 11.10 for the HP V-Class SCA servers. However, these features do not provide any potential performance benefits and no previous HP-UX SCA features have changed.

NOTE

The HP V-Class SCA servers themselves are not supported by the HP-UX 11i release, and all 11i supported systems are non-SCA servers that consist of single "locality domains" that includes all of the systems' hardware resources. Consequently, any use of the HP-UX SCA features on HP-UX 11i systems will result in the default locality placement and memory allocation behaviors.

However, at the HP-UX 11i release you can use the `/usr/bin/mpsched` command to inquire about system processors; you can launch programs with "gang schedule" enabled; and you can bind or unbind processes to processors (CPUs). You also can use `mpsched` to inquire about process bindings.

Useful `mpsched` command options include:

- `-c` (bind command or PID to specified processor ID)
- `-g` (launch command with gang scheduling enabled)
- `-p` (specific PID: process ID)
- `-q` (inquire about process, requires -p)
- `-s` (list system status, including processor IDs)
- `-u` (unbind specified process, requires -p)

You should not use the `mpsched` command's locality placement policy features because they provide no benefits on supported HP-UX 11i systems.

See the *mpsched* (1) manpage for details.

Dynamic Tunables

**new at 11i
original release**

A new facility has been added which will allow the retrieving of all tunable values and the setting of a limited number of tunables. If a tunable is dynamic, a change will take place immediately, without the need to reboot the system. Such changes will persist across reboots.

Three parts of the system have been changed to allow retrieving and setting of dynamic tunable values:

- System Administration Manager (SAM). For details, see "Changes to System Administration Manager (SAM)" on page 133 in this document and the SAM Online Help Facility.

- `kmtune`. This command has been enhanced to allow the changing of dynamic tunables. See the *kmtune* (1M) manpage for further details.

- Three new system calls have been added: `gettune`, `settune`, and `tuneinfo`. For the function and use of each of these system calls, see their manpages: *gettune* (2), *settune* (2) and *tuneinfo* (2).

Currently, the following tunables are dynamic:

- `maxuprc` - maximum number of user processes
- `msgmax` - message maximum size in bytes
- `msgmnb` - maximum number of bytes on a message queue
- `shmmax` - maximum shared memory segment size in bytes
- `shmseg` - shared memory segments per process
- `maxtsiz` - maximum text segment size in bytes for 32-bit programs
- `maxtsiz_64bit` - maximum text segment size in bytes for 64-bit programs
- `maxfiles_lim` - hard file limit per process
- `core_addshmem_read` - include readable shared memory in process core dump
- `core_addshmem_write` - include writable shared memory in process core dump

For more information, see the white paper *Dynamically Tunable Kernel Parameters in HP-UX 11i* at **http://docs.hp.com**.

Asynchronous Disk Pseudo Driver (async) Compatibility

**new at 11i
original release**

The `async` driver is used mostly by databases for doing asynchronous I/O to the disk.

Applications that use the `async` driver must be owned by the superuser, or by a user who is a member of a group for which the privileges include MLOCK.

To check a group's privilege capabilities, issue this command:

`/usr/bin/getprivgrp` *group_name*

If the output of `getprivgrp` does not show that the group has the `MLOCK` privilege, set the group's privilege by issuing this command as `root`:

`/usr/bin/setprivgrp` *<group_name>* MLOCK

Impact

If the application accessing the `async` driver is not owned by superuser or by a user who is a member of a group that has `MLOCK` privilege, `ASYNC_CONFIG` and `ASYNC_ADDSEG` `ioctl()` will fail and `errno` will be set to `EPERM`.

Compatibility

An application running on HP-UX 11.0 with patch `PHKL_22126` (or any patch that supersedes it) installed will operate correctly when upgraded to 11i.

If the application using the `async` driver has been operating on a 11.0 system without `PHKL_22126` (or any patch that supersedes it), then, when migrating to HP-UX 11i, the group associated with that application must be modified to include the `MLOCK` privilege.

System-V IPC Message Queue Enhancement

**new at 11i
original release**

System-V IPC is the System-V InterProcess Communications package developed by AT&T and comprises mechanisms for arbitrary processes to send and receive data "messages," share virtual address space, and use semaphores to synchronize execution. This enhancement applies only to the message subsystem.

The System-V IPC kernel tunable MSGMNB, which sets the maximum number of bytes on a queue, has had its maximum upper limit increased from 64KB to 64MB. New or recompiled applications will automatically use new, larger fields in the `msqid_ds` structure which describes queue sizes. However, if queue sizes greater than 64KB are desired, a compilation feature macro, __BIGMSGQUEUE_ENABLED, must be defined. This may be done using the `-D` compiler option or the `#define` pre-processor directive prior to any `#include`. (This requirement is temporary and used to maintain compatibility during a transition period.)

As hardware system capacities (including CPU speed and memory) have increased, some customer and third-party applications have been placing a greater demand upon the System-V IPC message queues. By increasing the capacity of these queues, applications are able to transfer large messages in a much more efficient manner. You will be able to dedicate more system memory for this purpose. Also, this may ease porting of some applications from other vendors' platforms which use message queues.

This enhancement is available on all systems.

Impact

There is no impact on system resources unless you increase system memory dedicated to System-V IPC message queues. To do this, the size of message memory segments (MSGSSZ) kernel tunable and/or the number of these segments (MSGSEG) kernel tunable may need to be increased.

An individual message queue cannot exceed the maximum queue size (MSGMNB) kernel tunable. The size of an individual message cannot exceed the MSGMAX kernel tunable.

Compatibility Issues

This change allows execution of existing binary programs. However, as described in the *msgget* (2) and *msgctl* (2) manpages, if binaries built on pre-11i HP-UX are used, the queue should not be created in excess of 64KB. To ensure this, the programs which create the queue (that is, via the IPC_CREAT option to `msgget`) should not be recompiled with the symbol "__BIGMSGQUEUE_ENABLED" defined. Also, the IPC_SET command to `msgctl` should not specify a `msg_qbytes` value in excess of 64KB.

The reason for this is that pre-11i binaries use 16-bit fields in the `msqid_ds` structure for `msg_qbytes` and `msg_cbytes` queue size information. If the actual queue sizes exceed 64K, these fields are capped at 64K (that is, 65535 - the maximum value 16 bits can represent). It should be noted that binary programs which don't use these fields will operate properly even with larger queues. Even some of those programs which do use the fields may do so in such a manner that the inaccuracy does not adversly affect program behavior.

These concerns arise only for separately-built binaries which share common message queues. A group of binary programs which uses queues less than or equal to 64K are not affected by a separate group of programs which may be using other queues greater than 64K.

The special compile-time symbol, __BIGMSGQUEUE_ENABLED, selects the enhanced capabilities. It is anticipated that, at the major release to follow 11i, the default will be switched so that programs recompiled without this symbol will create big queues. Programs recompiled on 11i will be capable of handling the larger queue size fields, even if not compiled with __BIGMSGQUEUE_ENABLED. You should consider whether your applications should be recompiled on 11i to prepare for that future release.

Performance Issues

The purpose of increasing the size limits on System-V IPC message queues is to improve performance of applications which pass large messages between processes. Specifically, increased size limits eliminate the necessity to break messages into smaller pieces, as well as reduce the high rate of context switching associated with such techniques.

Documentation Changes

The *msgget* (2), *msgctl* (2), and *glossary* (9) manpages have been modified to reflect these changes.

System-V IPC SEMMSL Dynamic Kernel Tunable

**new at 11i
original release**

The System-V IPC kernel tunable configuration parameter SEMMSL, which sets the maximum number of semaphores per ID which can be grouped within a single System-V IPC semaphore set, has changed from a hard-coded value of 2048 in kernel code to a dynamic kernel tunable. Its minimum and default value is 2048 while its upper limit is 10240. For 11i, dynamic tune adjustments to SEMMSL may only be done using System Administration Manager (SAM).

With the increase in system sizes, applications have the ability to handle greater numbers of cooperating processes. Some applications synchronize operations of these processes by semaphores within a single set. Thus, the increase in the possible size of a semaphore set means that these applications may increase the number of processes they use.

This enhancement is available on all systems.

Impact

Increasing the SEMMSL kernel tunable to allow larger System-V IPC semaphore sets does not itself consume any additional kernel resources such as memory. However, in conjunction with increasing this tunable, you may need to increase the total number of semaphores in the system by increasing the SEMMNS kernel tunable. This will consume additional system memory.

Compatibility Issues

The change from a hard-coded SEMMSL to a dynamic tunable is transparent to applications.

In releases prior to 11i, a "SEMMSL" symbol was hardcoded to 2048 in `sys/sem.h`. This symbol will no longer always be an accurate representation of the maximum number of semaphores in a set. Uses of the symbol in programs should be removed. The `pstat` interface can return an accurate value for this kernel tunable.

Performance Issues

Some applications will be able to scale to utilize larger systems, where scaling depends upon the size of a semaphore set.

SCSI Queue Depth Management

**new at 11i
original release**

If you have multiple active paths to a SCSI device (LUN), you might need to manage your device queue depths to maximize the device's performance. This is particularly true with dynamic multi-pathing applications—such as EMC's PowerPath application—which allow all multiple paths to a LUN to be in use simultaneously. In such cases, you should check the queue depth specified on each path. If it is set to a value that is more appropriate for an environment where only one path is active at any point in time, you might need to lower the value.

Even in single-pathing or static multi-pathing environments, management of device queue depths can be important to maximize the performance and throughput of the storage device.

A single hard-coded default queue depth of 8 existed originally on 11.0 and could be changed only one device at a time via an `ioctl` to the device. But, it does not meet the needs of all devices and configurations. So, HP-UX 11i contains the following enhancements to the SCSI device queue depth management:

1. A dynamic tunable called `scsi_max_qdepth` has been added. This tunable allows you to set the default queue depth that will apply to devices that have not been individually set via the SIOC_SET_LUN_LIMITS `ioctl` or `scsictl` commands. This tunable is "dynamic," which means that it can be changed and will be applied without having to reboot the system.

2. On 11.0, the queue depth could be changed on a per-device basis via the SIOC_SET_LUN_LIMITS `ioctl` or the `scsictl` command. However, the settings were not persistent across device opens and closes. That is, on 11.0, the queue depth setting on a device would disappear on the last close of the device and would go back to the system default of 8 when the device was re-opened. On 11i, the per-device queue depth settings will persist across opens and closes. This allows you to set the queue depth only once during or after boot up to maintain a desired value.

The 11i `scsi_max_qdepth` tunable can be changed or read via the `kmtune` command. See the *kmtune* (1M) manpage for details. The only change in the behavior of the per-device queue depth settings is the persistence across device opens and closes, as described above. Otherwise, these can be set or read in the same way as they could on 11.0 via the SIOC_SET_LUN_LIMITS and SIOC_GET_LUN_LIMITS `ioctl` command or the `scsictl` command.

Changes to mpctl() System Call

new at 11i original release

Some extensions have made to the `mpctl()` system call interface to pre-enable processor set functionality in the HP-UX 11i release.

The `mpctl()` interface provides command requests to query system information, such as the total number of processors and locality domains in the system, and the IDs of all processors and locality domains. When processor set functionality is implemented in HP-UX, these command requests will return information about the processor set of the calling thread, and not the entire system.

A new set of `mpctl()` command options query for system-wide topology information, regardless of which processor set contains those resources. The following seven new command requests are provided in HP-UX 11i:

New Request	Its Equivalent in 11i
MPC_GETNUMSPUS_SYS	MPC_GETNUMSPUS
MPC_GETFIRSTSPU_SYS	MPC_GETFIRSTSPU
MPC_GETNEXTSPU_SYS	MPC_GETNEXTSPU
MPC_GETNUMLDOMS_SYS	MPC_GETNUMLDOMS
MPC_GETFIRSTLDOM_SYS	MPC_GETFIRSTLDOM
MPC_GETNEXTLDOM_SYS	MPC_GETNEXTLDOM
MPC_LDOMSPUS_SYS	MPC_LDOMSPUS

The new command requests are mapped to their current equivalent requests in the 11i release, so applications in the 11i release are not affected.

When the processor set functionality becomes available, applications that rely on `mpctl()` to return system level information will need to be changed to use new command requests. For these applications we recommend using the new commands in the 11i release to avoid any issues when processor sets are available.

All applications that use `mpctl()` to query the available processors and locality domains to scale and bind for optimal performance will not require any changes when the processor set functionality is made available. The `mpctl()`'s existing commands will return information about what processors are available for binding.

When processor set functionality becomes available, if the system is not partitioned into more than one processor sets, no applications using the `mpctl()` interface with current command requests will be affected.

9 New and Changed Disk and File Management Features

Additional Support for Striping and Mirroring

**new at 11i
original release**

Logical Volume Manager (LVM) now supports striping and mirroring for shared volume groups. Previously under HP-UX 11.0, a volume group could not be activated in shared mode if any of its logical volumes were striped or mirrored. This restriction has now been lifted.

Shared volume groups are provided to support ServiceGuard Oracle Parallel Server (OPS), so this change only affects customers using OPS. Those customers are now free to stripe their logical volumes (to improve throughput) or mirror them (for higher availability).

NOTE

Where performance is limited by I/O throughput, striping may help.

See also "MirrorDisk/UX" on page 81.

New Whitepaper on File and File System Sizes

**new for
June 2001**

There is a new whitepaper, *Supported File and File System Sizes for JFS and HFS*, (5971-2383) available at the HP Documentation web site:

`http://docs.hp.com`

New Version of Journaled File System (JFS)

**new at 11i
original release**

HP-UX 11i includes a new version of the Journaled File System (JFS): version 3.3, as opposed to the previous version 3.1. (JFS is also known as the VERITAS File System or VxFS).

New features in JFS 3.3 include the following:

- support for access control lists (ACLs), the only HFS feature unavailable in JFS 3.1 (see *aclv* (5), *getacl* (1), *setacl* (1), *acl* (2), and *aclsort* (3C)).

- a new disk layout (that is, version 4)

- a new command for tuning a VxFS file system: `vxtunefs` (see *vxtunefs* (1M) and *tunefstab* (4))

- a new command, `vxfsconvert`, for converting an HFS file system to a JFS file system. This command also converts HFS ACLs to JFS ACLs, with some limitations (see *vxfsconvert* (1M)).

- performance enhancements

- new packaging and licensing strategy for HP OnLineJFS 3.3 (see *vxlicense* (1M) and *vxenablef* (1M))

- a better solution for the file system shrink limitation when using the version 4 disk layout

With the HP-UX 11i release, JFS becomes a superset of the functionality available in HFS now that JFS includes support for ACLs. This enables all users to consider migration from HFS to JFS.

In HP-UX 11i, one kernel library contains the kernel functionality for both the JFS and the OnLineJFS products. When you install the JFS product, all the software for OnLineJFS is also installed, but its features are not enabled unless you also purchase it.

Having all the kernel functionality for both products in one library resolves many of the patching problems that existed in previous releases.

With the JFS version 4 disk layout in JFS 3.3, you are much less likely to encounter the file system shrink limitation that existed in earlier JFS versions in HP-UX 10.20 and 11.0. In other words, JFS previously could not shrink a file system if there were file extents residing in the area being reduced. Now, JFS 3.3 (with the version 4 disk layout) will attempt to move extents off the area of the file system being reduced. This provides a greater chance of success when shrinking JFS file systems. However, there may still be some occasions where JFS cannot move extents off the area of the file system being reduced, in which case a shrink will still fail.

Documentation

All VxFS manual pages are updated, as are the manual pages for generic HP-UX commands and functions which accommodate ACLs (`cp` and `find`, for instance). See *Managing Systems and Workgroups: A Guide for HP-UX System Administrators*, part no. B2355-90742, for a description of JFS ACLs and how to use them.

The *HP JFS 3.3 and HP OnLineJFS 3.3 VERITAS File System 3.3 System Administrator's Guide* is available on HP's documentation web site at

`http://docs.hp.com` and on the Instant Information CD, in both HTML and PDF formats.

Compatibility Issues

JFS ACLs use a different format from HFS ACLs. However, the new command, *vxfsconvert* (1M) will convert an HFS file system to a JFS file system, as well as HFS ACLs to JFS ACLs (with the limitation that HFS ACLs with no JFS ACL equivalents are not converted). See *Managing Systems and Workgroups: A Guide for HP-UX System Administrators*, part no. B2355-90742, for a description of the procedure for converting a file system.

Note that JFS ACLs require a file system with the new disk layout (version 4). To upgrade a file system from an older disk layout to version 4, you can use the `vxupgrade` command.

JFS 3.3 uses new header files. As far as the JFS module is concerned, a well-behaved application will not need to be recompiled. However, a kernel-intrusive application will need to be recompiled with the new header files, and possibly with some corresponding code changes. You should check with the application provider before upgrading.

Performance Issues

JFS 3.3 generally outperforms previous releases, as design changes have reduced the number of bottlenecks resulting from globally shared locks. JFS includes tunables and features to support improved performance in the OLTP, DSS, and technical computing markets. With appropriate tuning, JFS 3.3 also outperforms HFS in all categories.

JFS 3.3 includes a new command, `vxtunefs`, for tuning a VxFS file system. See *vxtunefs* (1M) and *tunefstab* (4). Also see the *HP JFS 3.3 and HP OnLineJFS 3.3 VERITAS File System 3.3 System Administrator's Guide* for information on tuning a JFS file system.

NOTE   The `volcopy` and `labelit` commands will be obsoleted in a future release. You should use `vxdump` and `vxrestore` for backup and restore, or you can use an application-specific utility. You can use `dd` to make a literal copy of the file system.

Network File System Support on TCP/IP

**new at 11i
original release**

With for versions 2 and 3, Network File System (NFS) is now supported over the connection-oriented protocol, TCP/IP, in addition to running over User Datagram Protocol (UDP).

As a result of this new functionality, NFS is now supported over wide-area networks (WANs). As long as TCP is supported on the WAN, then NFS is supported also. (TCP transport increases dependability on WANs. Generally, packets are successfully delivered more consistently because TCP provides congestion control and error recovery.)

The `mount_nfs` command now supports a `proto=` option on the command line where the value for `proto` can be either UDP or TCP. (In the past, this option was ignored.) This change allows administrators to specify which transport protocol they wish to use when mounting a remote file system.

If the `proto=` option is not specified, then NFS, by default, will attempt a TCP connection. If that fails, it will then try a UDP connection. Thus, by default, you will begin using TCP instead of UDP for NFS traffic when you begin using the 11i version of HP-UX. This should have little impact you. You do, however, have the option to specify either UDP or TCP connections.

If you specify a `proto=` option, only the specified protocol will be attempted. If the server does not support the specified protocol, the mount will fail.

The `nfsd` daemon now opens TCP transport endpoints to receive incoming TCP requests. For TCP, `nfsd` is multi-threaded. For UDP, `nfsd` is still multi-processed.

Kernel TCP threads execute under the process `nfskdtcp`. When counting the number of `nfsd` processes, keep in mind the following algorithm: An equal number of `nfsd`s that support UDP will be created per processor, but only one `nfsd` that supports TCP will be created. In the case of a four-way machine and NUM_NFSDS=14 (set in `/etc/rc.config.d/nfsconf`), 17 `nfsd`s will be created: 16 for UDP (4 per processor) and 1 for TCP.

The `nfsstat` command now reports TCP RPC statistics for both client and server. The TCP statistics are under the connection-oriented tag and the UDP statistics are under the connectionless-oriented tag.

AutoFS supports the `proto=` option in the Automounter maps and has the same behavior described above under the `mount_nfs` command. In the past, this was an invalid option.

However, Automounter will not support NFS over TCP.

Unlike the 11.0 patch release of NFS over TCP, there is no enablement flag in the 11i release for NFS over TCP. By default, NFS will attempt to use TCP.

The kernel RPC layer has been modified to support TCP connections over NFS. A new streams module, `rpcmod`, has been added to manage the TCP connections. These changes are internal to the NFS implementation and are not user accessible.

Documentation Changes

The following manpages have been modified for this new feature:

- *mount_nfs* (1M)
- *nfsd* (1M)
- *automount* (1M)
- *nfsstat* (1M)

Other NFS Changes

**new at 11i
original release**

Three additional features have been added to NFS:

- Loopback transport support has been added to transport-independent RPC.
- Automatic user-space thread generation has been enabled in the RPC library.
- NFS server-side performance enhancements have been added.

Loopback Transport Support

Loopback transport provider devices (`tlclts`, `tlcots`, and `tlcotsod`) have been added to the TI-RPC definition file, `/etc/netconfig`. Also, the system now has the following new loopback transport-specific directories and files:

`/etc/net/`*loopback_transport_name*`/hosts`

`/etc/net/`*loopback_transport_name*`/services`

The following ONC/NFS daemons support loopback transport requests:

- `/usr/sbin/rpcbind`
- `/usr/sbin/keyserv`
- `/usr/sbin/rpc.nisd`
- `/usr/sbin/nis_cachemgr`

The `netid` and address fields in the `rpcinfo` call (which queries `/usr/sbin/rpcbind` to determine what services have been registered) now give the loopback device name plus an address name, rather than just the `netbuf` address provided by `udp` and `tcp` transport.

Additionally, the `ticlts` loopback transport device has a randomly generated string address.

User-Space Thread Generation

To process incoming RPC requests, the `svc_run()` function call in the RPC library automatically generates a thread on behalf of the application. The threads are managed by the RPC library software. RPC threads may be created when calling the RPC library.

NFS Server-Side Performance Enhancements

The NFS server daemon, `/usr/sbin/nfsd`, has been modified to enhance performance. As a result, the user may see more `nfsd` daemon processes running than requested, depending on the number requested and the number of processors configured. This change is documented in the *nfsd* (1M) manpage.

The NFS client-side buffer cache management has been modified to improve server performance from a VxFS file system mounted on the client.

The performance enhancements included have given HP industry leading NFS SPECsfs benchmark values on our V-Class platforms.

Mounting and Unmounting NFS File Systems Automatically Using AutoFS

HP-UX 11i provides a daemon that mounts and unmounts NFS file systems automatically. This feature is known as AutoFS.

AutoFS coexists with automount and performs the same functions as automount, but has a new, more reliable design. Additionally, AutoFS supports the NFS PV3 protocol whereas the automounter does not. The automount command has been replaced with a shell script that will either invoke the old automount daemon or the new AutoFS `automount` command, depending on the variable AUTOFS in `/etc/rc.config.d/nfsconf`.

- `AUTOFS=1` causes `/sbin/init.d/nfs.client` to start the AutoFS daemon (`automountd`) and run the AutoFS automount command.

- `AUTOFS=0` starts the old automount daemon. This is the default on newly installed or updated systems.

The old automount executable is located at:

`/usr/lib/netsvc/fs/automount/automount`

The new AutoFS executables are located at:

`/usr/lib/netsvc/fs/autofs/automountd /usr/lib/netsvc/fs/autofs/automount`

When AutoFS is executed, a process used by its kernel code for kernel thread support is also started. The `autofs_proc` process cannot be killed, except by a shutdown of the system.

Impact

From an operational standpoint, AutoFS functions comparably to the old automounter and returns the same values.

From the system administrator's standpoint, however, AutoFS is started, stopped, and updated differently than its predecessor. The `nfs.client` start-script automatically starts and stops the correct daemons depending on the value of AUTOFS in `/etc/rc.config.d/nfsconf`.

NOTE   If you do not use this script, you need to remember which implementation of automatic NFS file mounting you are using. Starting both AutoFS and automounter can lead to problems accessing the remote file system. You must reboot to switch between AutoFS and the old automounter.

Other Operational Differences

Any user-written scripts that expect the automount command to remain running as a daemon will have to be updated either to not expect this behavior or to check explicitly that `automountd` is running. AutoFS can no longer be shut down by killing the automount process; instead, you must shut it down by executing the following command:

```
/sbin/init.d/nfs.client stop
```

This will unmount all mounted AutoFS filesystems and then kill the `automountd` process.

To stop AutoFS without using the `/sbin/init.d/nfs.client` script, you must enter the following:

```
/usr/sbin/umountall -F autofs
```

```
kill automounted_pid
```

The automount `-n`, `-M`, and `-tw` options are not supported in AutoFS. The `-m` and `-tm` options are also not supported, but their behavior can be configured in different ways:

- by modifying the `nsswitch.conf` file to get the `-m` behavior

- by modifying the automount map entries to specify the time-out for the `-tm` option. The `-tl` option is accessed using `-t`.

Another difference between automounter and AutoFS is that AutoFS no longer uses symbolic links to access the mount points. Applications that depend on this explicit behavior will no longer work as expected.

Additional Information

The existing 11.0 automounter can be re-enabled, if desired, by setting the `AUTOFS` variable to 0 or by removing the `AUTOFS` variable from `/etc/rc.config.d/nfsconf`. In this configuration, automounter will not mount file systems via the NFS version 3 protocol.

For more information on how to migrate to AutoFS, see Chapter 2 in the *Installing and Administering NFS Services* manual.

Configuration

To enable AutoFS, you must add or set the AUTOFS variable to 1 in `/etc/rc.config.d/nfsconf`. Here is an example of this change:

```
#autofs configuration. See autofs(1m)
#
#For the 11.0 Release line both AUTOFS and the old Automount
#are delivered. In order to invoke the AUTOFS instead of
#you must set the AUTOFS flag to 1.
#
#/usr/sbin/automount is now a script that sources in this file
#Depending on the variable AUTOFS, either AUTOFS or the old
#automount process will execute. The nfs.client start script
#will also use this variable to start the appropriate process
#during the boot sequence.
#AUTOFS= 0 - use the old automount process.
#       1 - use the new AutoFS.
#AUTOMOUNT_OPTIONS= - options to the AutoFS automount command
#AUTOMOUNTD_OPTIONS= - options to the AutoFS automountd daemon
#
#The AUTOMOUNT flag still needs to be set for either the old
#automount or new AutoFS to be started by the nfs.client script.
#
```

```
AUTOFS=1
AUTOMOUNT_OPTIONS=""
AUTOMOUNTD_OPTIONS=""
```

Documentation Change

A new manpage, *automountd* (1M), describes the AutoFS automount daemon. The *automountd* (1M) manpage has been modified to describe both the old automount daemon and the new AutoFS command.

Obsolescence

Although all 11.0 patch bundles contain both AutoFS and automounter, AutoFS will replace automounter in a future release of HP-UX.

HP Fibrechannel High Availability Disk and Closure

The HP Fibrechannel High Availability Disk and Closure, also referred to as the FC10, is a Mass Storage Subsystem disk enclosure. This is the design center for Fibre Channel-Arbitrated Loop and future SCSI Enclosures (SES). Some of the features of the FC10 are as follows:

- enclosure temperature, fan speed, and power supply monitoring
- fault tolerance through redundancy of disk paths, IO modules, fans, and PSs
- FC-AL interconnect redundancy to disks and host
- hot pluggable disk modules, fans, and power supplies
- support for proactive maintenance

Fibre Channel Mass Storage Diagnostic Message and Kernel Tunable

The Fibre Channel Mass Storage product will return the following diagnostic message if the disk device violates the Fibre Channel Standard: `ECB_FRAME_RECV_BEFORE_ADISC`

You will see the this error message in the kernel log file. If this message is received frequently and persistently, please contact your HP Customer Support Representative.

Fibre Channel Mass Storage has a new kernel tunable: `fcp_large_config`. In a Fibre Channel Mass Storage configuration, if this parameter is set to 1, it allows for large loops with up to 126 `nports`. For example:

```
fcp_large_config 1
```

If the parameter is set to 0, you are limited to less than 64 `nports`.

10 Internet and Networking Services

IPv6 Available on Software Pack

What is IPv6?

new for December 2001

Internet Protocol version 6 (IPv6) is a new generation of the Internet Protocol that is beginning to be adopted by the Internet community. IPv6 is also referred to as "IPng" (IP next generation). It provides the infrastructure for the next wave of Internet devices, such as PDAs, mobile phones, and appliances. It also provides increased connectivity for existing devices such as laptop computers.

The most visible difference between today's commonly used version of IP (IP version 4) and IPv6 is the larger address space supported by IPv6. IPv6 supports 128-bit internet addresses, compared to the 32-bit internet address supported by IP version 4. Additionally, IPv6 offers greater ease of configuration and manageability as well as increased security.

Once the HP-UX 11i IPv6 software product bundle is installed on the server and the IPv4 and IPv6 interface(s) are configured, the server is considered to be an IPv6/IPv4 "dual stack" implementation. This implies that IPv4 and IPv6 both run concurrently and independently. The server can communicate with both IPv4 nodes and IPv6 nodes, and can identify packets as being IPv4 or IPv6. A dual stack implementation supports both IPv4 and IPv6 applications. Programmers can write IPv6 applications that communicate with both IPv6 and IPv4 peers. Existing IPv4 applications will continue to work.

The following section lists HP-UX functionality that has been IPv6-enhanced. If an area is not included in this list, then this functionality has not been IPv6-enhanced and can only run on IPv4.

What's Included in HP-UX 11i IPv6?

This section provides only a brief summary list of what is included in the HP-UX 11i IPv6 software product bundle. For more detailed information, refer to the *HP-UX 11i IPv6 Release Notes* (T1306-90001).

New and Changed Features

- IPv6/IPv4 Dual Stack Support.
- Transition Mechanisms. These enable IPv6/IPv4 dual stack hosts and routers to connect with other IPv6/IPv4 dual stack hosts and routers over the existing IPv4 Internet. HP-UX 11i IPv6 supports the following transition mechanisms: configured tunneling, automatic tunneling and "6to4." With tunneling, IPv6 datagrams are encapsulated within IPv4 packets.
- IPv6 Stateless address autoconfiguration. A mechanism where a host can automatically assign an address to configure an interface.
- IPv6 Neighbor and Router Discovery and Duplicate Address Detection.
- TCP and UDP over IPv6, PMTUv6, ICMPv6, IPv6 MIBs, and Sockets APIs.
- New *netconf-ipv6* file. Used to store IPv6 settings (similar to IPv4's *netconf* file).

- Network-Interface Administration Utilities for both IPv4 and IPv6:

  ```
  ifconfig
  netstat
  ping
  route
  ndd
  traceroute
  ```
  `ndp` (a new IPv6-only utility for neighbor-discovery; ndp is similar to the arp utility used with IPv4)

- Support in */etc/hosts* for both IPv4 and IPv6 Addresses. Look-up policies for IPv6 are identical to those of IPv4.

- IPv6 Name/Address Resolution for Name Service Switch: new entry (`ipnodes`) in */etc/nsswitch.conf.*

- IPv6-enhanced Internet Services:

  — Services included with the HP-UX 11i IPv6 software product bundle:

     inetd, internet daemon
     telnet
     r* commands
     name and address resolution resolver routines
     inetd.sec over IPv6 is also supported

  — Services *not* included with the HP-UX 11i IPv6 software product bundle:

     WU-FTPD 2.6.1
     BIND 9.2
     Sendmail 8.11.1
     DHCPv6

     (Although these services are not part of the HP-UX 11i IPv6 software product bundle, they are available independently from HP's Software Depot, at `http://software.hp.com`.)

- The Nettl utilities (`nettl`, `nettladm`, `netfmt`) have been enhanced to trace and filter new IPv6 subsystems.

- Where needed, IPv6-enhancements have been made to some C2 Audit and HP-UX commands (for example - `lp`, `syslogd`, `rlpdaemon`). There are some known problems with the `who`, `last` and `finger` commands. These are documented in the *HP-UX IPv6 Release Notes* (T1306-90001).

- DCE Clients support.

- IPv6-enhanced libc.

Identifying IPv6 Systems

Systems with the HP-UX 11i IPv6 software product bundle installed can be identified by running:

```
swlist -l bundle IPv6NCF11i
```

where the following will be returned:

```
IPv6NCF11i        B.11.11.0109.5D IPv6 11i product bundle
```

Where to Find Information

The following customer documentation is available on the Web at `http://docs.hp.com`. Note that these documents were written for an independent Software Depot release. The information in these documents still applies; however the "Installation" details are written for a Web download only.

- *HP-UX 11i IPv6 Release Notes* (T1306-90001)

- *Installing and Administering HP-UX 11i IPv6* (T1306-90002)

- *IPv6 Porting Guide*

You can find all of the above documentation as well as the *HP-UX 11i IPv6 Product Note* (T1306-90003) for Software Pack on the Software Pack December 2001 media. The product note contains additional information on known problems and how to install IPv6 from the Software Pack CD. The Software Pack CD is included in the HP-UX 11i media kit.

Base HP-UX Internet Services

NOTE
Although HP-UX 11i includes Sendmail 8.9.3, BIND 8.1.2, and WU-FTPD 2.4, you can download the following new versions of these products from HP's software depot at `http://www.software.hp.com` under "internet & security": Sendmail 8.11.1, BIND 9.2, and WU-FTPD 2.6.1.

You can find information on these versions' features and functionality at `http://www.docs.hp.com` under "Networking and Communications."

Sendmail-8.9.3

new at 11i original release
A new version of `sendmail`, sendmail-8.9.3, is included with HP-UX 11i. This version provides additional features compared to the previous version. The sendmail-8.8.6 `sendmail.cf` file is compatible with the sendmail-8.9.3 binary. However, to take advantage of all the new features provided in this version, HP highly recommends that you use the default `sendmail.cf` file provided in the `/usr/newconfig/etc/mail` directory. Any site specific changes will need to be made as required.

New Features

New features in sendmail-8.9.3 include:

- Lightwight Directory Access Protocol (LDAP) support for address lookup
- New configuration file options:

`MaxHeadersLength`

Used to limit the maximum length of a mail header. The default maximum length is 32768.

`MaxRecipientsPerMessage`

Used to limit the number of recipients for a single mail message if the recipients have their mailboxes on the same mail server. The maximum value for this option is 100.

`DontBlameSendmail`

Used to enforce a security check on the mode files that sendmail reads and writes. The default value of this option is "safe."

`QueueSortOrder`

This option is NOT case sensitive.

`EightBitHeader`

Used to allow eight bit header when set to TRUE.

`PrivacyOptions=noetrn`

The `noetrn` flag will disable the `SMTP ETRN` command that forces `sendmail` to process its queue asynchronously.

```
PrivacyOptions=noverb
```

> The `noverb` flag will disable the `SMTP VERB` command that causes `sendmail` to enter the verbose mode and activate the deliver mode.

- Support for new mailer and map class:

Mailer: `discard`

> A special internal delivery agent named `discard` is now defined for use with `check_*` rulesets and header checking rulesets.

Map class: `regex`

> Sendmail-8.9.3 supports regular expressions using the new map class `regex`. The `regex` map can be used to see if an address matches a certain regular expression. By using such a map in a `check_*` ruleset, you can block a certain range of addresses that would otherwise be considered valid.

- Anti-spam configuration control. To enable some of the new anti-spamming rulesets, a shell script `gen_cf` is provided in the `/usr/newconfig/etc/mail/cf/cf` directory.
- New header checks. New syntax to do limited checking of header syntax is available.

Documentation

Refer to the *Installing and Administering Internet Services* manual available on the HP-UX 11i Instant Information CD and on the web at **http://docs.hp.com/hpux/11i** for detailed information on new features.

BIND 8.1.2

new at 11i original release   BIND 8.1.2, is shipped with HP-UX 11i. This version supports Dynamic updates via the Dynamic Domain Name Server (DDNS).

NOTE   The Dynamic updates, however, are NOT secure, and you are advised to put security mechanisms in place before using this feature.

The following lists the new features:

- DDNS Change Notification (RFC 1996)
- Completely new configuration syntax
- Flexible, categorized logging system
- IP-address-based access control for queries, zone transfers, and updates that may be specified on a zone-by-zone basis
- More efficient zone transfers
- Improved performance for servers with thousands of zones
- The server no longer forks for outbound zone transfers
- Many bug fixes

New Configuration File

The `BIND` configuration file is now `named.conf`, with many more configurable variables

than in previous releases of BIND. (The configuration file in previous versions of BIND was `named.boot`.)

There are now entirely new areas of configuration, such as, access control lists and categorized logging. Many options that previously applied to all zones can now be used selectively.

The configuration file can be obtained by following these steps:

1. Make sure that Perl is installed on the system.

2. Copy the `hosts_to_named` script to `/usr/sbin` and manually provide a link from `/usr/bin`.

3. To convert the existing `named.boot` file to `named.conf` file, use the Perl script `named-bootconf.pl` available in `/usr/bin`.

4. Create the new BIND configuration file `named.conf`. Do this in either of two ways:

   • If the configuration file `named.boot` already exists, create new config file as follows:

      `/usr/bin/named-bootconf.pl named.boot > named.conf`

   • If a BIND configuration file does *not* exist, execute `hosts_to_named` with appropriate options.

New Configurable Resolver Options

The timeout value is a function of the RES_RETRY and RES_RETRANS options of the resolver routines. It is currently hardcoded as 5000 milliseconds for RES_RETRANS and 4 attempts for RES_RETRY. This results in a timeout value of 75 seconds, which is obtained when you assume that there is one nameserver. When there are multiple nameservers, the timeout value will increase. Hence, to help achieve shorter timeout values, and better performance, the resolver options RES_RETRY and RES_RETRANS are now configurable.

These resolver options can be configured using any of the three methods shown below. They are listed in order of priority, from highest (first) to lowest (last).

1. Use environment variables.

2. Use resolver configuration file `/etc/resolv.conf`.

3. Use the new APIs defined in `set_resfield`.

The RES_RETRY and RES_RETRANS options can be set with any positive non-zero integer.

"PAM-ized" rexecd and remshd

The `rexecd` and `remshd` services on HP-UX 11i now use Pluggable Authentication Modules (PAM) for authentication.

You can take advantage of using an authentication mechanism of your choice like DCE Integrated Login, UNIX, or Kerberos by making a change in the `/etc/pam.conf` file. By default, if you do not edit the `/etc/pam.conf` file, the `rexec` and the `remsh` services will use the authentication mechanism specified by the OTHER directive in the `/etc/pam.conf` file.

The earlier version of `rexecd` and `remshd` allowed only those UNIX users who were included in `/etc/passwd` to use the `rexecd` and `remshd` services. This limitation has

been eliminated with the introduction of the "PAM-ized" modules. By PAM-izing `rexec` and `remsh` services, users belonging to other authenticating services like DCE Integrated Login can use the `remsh` and `rexec` services.

/etc/pam.conf File Changes

To use PAM-ized `rexec` and `remsh`, the following lines have to be added to the `/etc/pam.conf` file:

```
rcomds    auth required      /usr/lib/security/libpam_unix.1
rcomds    account required   /usr/lib/security/libpam_unix.1
```

Using PAM-ized remshd in Secure Internet Services (SIS) Environment

`rexecd` is not Kerber-ized and hence will not work in the SIS environment. However, `remshd` is Kerber-ized. To take advantage of the PAM-ized modules, add the following line to the `/etc/pam.conf` file:

```
rcomds    auth required      /usr/lib/security/libpam_dce.1
```

Also in the Kerberos environment, `remshd` has command line options for combining the UNIX method and the Kerberos method of authentication. These command line options can be set in the `/etc/inetd.conf` file for the `kremshd` service. Refer to the *kremshd* (1M) manpage for a more detailed description of the options available.

Changes for GateD

With HP-UX 11i, the `HELLO` protocol of GateD will be obsoleted and no longer supported.

However, the BGP protocol available with GateD-3.5.9 on HP-UX 11.0 is also available and supported on HP-UX 11i.

DHCP with Nonsecure DNS Updates

The Dynamic Host Control Protocol (DHCP) available on HP-UX 11i is capable of updating the Dynamic Domain Name Server (DDNS). This feature updates the DDNS with the name and IP address of the client. This means that for every client to which DHCP assigns a name and IP address, it also adds an "A" and "PTR" resource record (RR) of that client to the DDNS.

To assign a name for every IP address, a new, Boolean tag, `pcsn`, has been introduced. If set, the DHCP server gives priority to the name (if any) provided by the client. The name should be a fully qualified domain name (FQDN). If it is *not,* then the DHCP server will try to append the domain name (if set using the `dn` tag); otherwise, it appends a "." and updates the DDNS. If the `pcsn` tag is set, then the DHCP server will try to assign a name of its choice for every IP address.

To enable the DHCP server to perform updates to the DDNS, you need to add a new tag, `ddns-address`, specifying the address of the DDNS server, as well as the `pcsn` tag (within the same entry) to the `DHCP_POOL_GROUP` or `DHCP_DEVICE_GROUP` keywords.

The following is a sample of a `DHCP_DEVICE_GROUP` entry that includes the `ddns-address` tag and the `pcsn` tag:

```
DHCP_DEVICE_GROUP:\
ba:\
pcsn:\
class-name=SUBNET_128_XTERMINAL_GROUP:\ class-id="xterminal:"\
subnet-mask=255.255.255.0 :\
addr-pool-start-address= 15.14.128.1 :\
addr-pool-last-address= 15.14.128.254 :\
ddns-address=1.2.3.4:\
lease-time=604800 :\
lease-grace-period=5
```

Network Transport

**new at 11i
original release**

Network Transport includes `ifconfig`, `ndd`, `netstat`, `virtual IP address`, `setsockopt`, and `t_optmgmt`. All are discussed in more detail below.

ifconfig

The `ifconfig` subnet mask default now allows all 1's or all 0's in the masked part of the `subnet` field. (The `subnet` field is that portion of an IP address that identifies the subnet beyond the network portion of the address.) This provides up to twice as many IP addresses as before. Even though the default behavior now allows more IP address and `subnet` mask combinations, any addresses working before will continue working without alteration.

In the following example, a class A IP address used with the mask 255.192.0.0 (0xffc000000) has a two-bit `subnet` field which is the 5th and 6th bits:

11111111 11 000000 00000000 00000000

Now, `ifconfig` can assign the following IP address and `subnet` mask to an interface, even though the `subnet` field (subnet portion of the address) is all 1's:

IP address: 15.192.1.1

`Subnet` mask: 255.192.0.0 (0xffc00000)

In binary:

00001111 11 000000 00000001 00000001

11111111 11 000000 00000000 00000000

Also, `ifconfig` can now assign the following IP address and `subnet` mask to an interface, even though the `subnet` field is all 0's:

IP address: 15.1.1.1

`Subnet` mask: 255.192.0.0 (0xffc00000)

In binary:

00001111 00 000001 00000001 00000001

11111111 11 000000 00000000 00000000

To disallow `subnet` fields with all ones or all zeroes, set the `ndd` parameter `ip_check_subnet_addr` to 1 in the `nddconf` file.

ndd

The networking configuration tool, `ndd`, is used to customize the networking kernel. To make an 11i system more Internet friendly and easier to run "out of the box," some of the `ndd` tunable parameters defaults have changed, some formerly unsupportable tunable parameters are now supported, and some new tunable parameters have been added. In part, these modifications reflect changes to networking industry standards.

Specific Changes

1. All 1's or all 0's are now allowed in masked bits of `subnet` address: `ip_check_subnet_addr` shows whether or not that RFC1122 or RFC1878 enforces the network `subnet` mask. If it is a 0 (zero), then the RFC1122 behavior is seen. If it is a 1 (one), then RFC1812 is seen. The default is now RFC1812 behavior. (See `ifconfig` in the preceding section for more information.)

    This new behavior, an enhancement, makes available up to twice as many IP addresses than a similarly configured RFC1122 machine.

2. TCP-supported `tcp_sack_enable` now enables selective acknowledgement. This enhancement could improve performance in networks with large transmission windows by allowing TCP recipients to indicate lost segments within large transmission blocks. The TCP sender can then retransmit only the lost segments. Supported parameter values are as follows:

2	Local system enables SACK if remote system first sends SACK (Default).
1	Local system requests the SACK option during a `connect()` request.
0	Local system never uses SACK.

3. Send and receive buffers are now limited by the following tunable: Specifically, `udp_recv_hiwater_max` (default 2 GB) sets the maximum *receive* buffer size that `setsockopt` or `t_optmgmt` can set for a UDP socket. The TCP version of this is `tcp_recv_hiwater_max` (default 1 GB). On the other hand, `tcp_xmit_hiwater_max` (default 2 GB) sets the maximum *send* buffer size that `setsockopt` or `t_optmgmt` can set for a TCP socket. There is no UDP version of this value. These system-wide parameters prevent processes from keeping large amounts of data in send or receive buffers, and thereby consuming system resources.

CAUTION   Although the enhancement `tcp_fin_wait_2` may be used to set how long a connection will be in `FIN_WAIT_2`, you should use this cautiously. If the remote TCP entity is slow, but would terminate normally (is not hung nor will terminate abnormally), TCP may close the connection prematurely. This could result in the flushing of data in the remote connections receive buffer. If this happens unexpectedly, then the data could become corrupted.

For more information on using `tcp_fin_wait_2`, enter the following command: `ndd -h tcp_fin_wait_2`.

With a new field, `ip_udp_status` now reports how many times a given UDP socket has overflowed. Although this enhancement only works on sockets currently open, it can be a very handy troubleshooting tool used when `netstat -p udp` shows socket overflows.

4. With a TimeStamps option now supported, `tcp_ts_enable` allows RFC 1323 TimeStamp extensions to TCP Headers. The TimeStamps are used for two purposes:

    a. RTTM (Round Trip Time Measurement) of the interval between the time a TCP sends a segment and the time the return acknowledgement arrives.

    b. PAWS (Protect Against Wrapped Sequences) on high-speed networks.

Supported parameter values are as follows:

2	Use Timestamps option if initiated by the remote system
1	Always try to initiate the use of Timestamps option
0	Never use Timestamps option

5. With the enablement of socket caching, `tcp_conn_strategy` can now increase performance by setting how many cached socket structures the system keeps. The default value of 0 (zero) disables the feature. A value between 1 and 512 sets a minimum of 512. Any number above 512 sets `tcp_conn_strategy` to that value. Enabling socket structure caching can increase system performance if there are many short-lived connections on the system.

6. Using the following formula, `tcp_cwnd_init` now allows you to configure the sender's initial TCP congestion window size:

   Min(tcp_cwnd_init * MSS), max(2 * MSS, 4380),

   where MSS is the maximum segment size for the underlying link. Default 4: (TCP implements RFC 2414). Range: 1-4

7. To prevent a type of Denial-of-Service attack, `ip_pmtu_strategy` "2" is not supported for 11i. In particular, a local system can no longer send its expected Path Maximum Transmission Unit (PMTU) value within an ICMP_ECHO request to a remote system or router.

   Type `ndd -h` for an online description of tunable parameters and other documentation.

Compatibility Issues

For the following commands, `ndd` displays IP addresses using the IP version 6 (IPv6) format:

ndd -get /dev/ip ip_tcp_status

ndd -get /dev/ip ip_udp_status

ndd -get /dev/ip ip_raw_status

When `ndd` maps IP version 4 (IPv4) addresses to IPv6 addresses, `ndd` displays the IPv4 addresses with the prefix `::ffff: . However,` `ndd` displays the remainder of the IPv4 address in dotted-decimal notation. This could cause scripts that are looking for a given output to fail.

Performance

None of the features will degrade performance. Enabling socket caching using `tcp_conn_strategy` could potentially increase performance by 10 to 20%.

netstat

There are two changes to `netstat`, which displays the statistics and configuration of the networking kernel: one is to `netstat -r`; the other to `netstat -I`.

Since `netstat -r` no longer updates the "Use" field, `netstat -r` no longer displays it.

Beginning at HP-UX 11i, `netstat -I <interface>` displays statistics accumulated since the last system reboot. This matches `netstat -I` output for HP-UX releases 10.20 and earlier.

Compatibility Issues

There could be some compatibility problems with scripts where they look for the "Use" field.

Virtual IP (VIP) Address for the System

Using the loopback interface lo0:1, lo0:2, and so on, the system will respond to the IP address assigned to these interfaces using any physical interface. Thus, a system can now have a "systemIP" address that will be available as long as one interface stays usable.

In some configurations, a system needs to keep a "well known" IP address that will always be available even if an interface goes down. With the new VIP feature, a remote user can specify an IP address that will respond regardless of the local interface from which the packet arrived. This feature is an enhancement.

setsockopt()

If you determine that certain applications always ask for the largest socket buffer allowed, then you may want to set these variables and limit the amount of memory used by such applications. (When an application opens enough of these large sockets and the system does not contain a lot of memory, then the system may starve for memory if the application quits reading from the socket.)

The system-wide kernel parameters, `tcp_recv_hiwater_max` (for TCP sockets) and `udp_recv_hiwater_max` (for UDP sockets), now limit the maximum buffer sizes specified in the `SO_SNDBUF` or `SO_RCVBUF` `setsockopt()` parameters.

Applications that request sockets with send or receive buffers larger than high-water marks set by the administrator will fail. In other words, a `setsockopt()` call with a `SO_SNDBUF` or `SO_RCVBUF` option that exceeds the corresponding kernel parameter value will fail, returning the errno value `EINVAL`.

T_OPTMGMT

If you determine that certain applications always ask for the largest buffer or transport service data unit (`tsdu`) allowed, then you may want to set these variables and limit the amount of memory used by such applications. (When an application opens enough of these large sockets and the system does not contain a lot of memory, then the system may starve for memory if the application quits reading from the endpoint.)

The kernel parameters `tcp_recv_hiwater_max` (for TCP sockets, default 2GB) and `udp_recv_hiwater_max` (for UDP sockets, default 2GB) now limit the `XTI_RCVBUF` parameter maximum buffer size. The kernel parameter `tcp_xmit_hiwater_max` (default 2GB) now limits the `XTI_SNDBUF` parameter's maximum buffer size.

Applications that request sockets with buffers or `tsdu`s larger than high-water marks set by you will fail. In other words, a `t_optmgmt()` call with a `tdsu` or `etsdu` option that exceeds the corresponding kernel parameter value will fail with `TBADOPT`.

New Versions of FTPD

**new at 11i
original release**

This release contains a new version of FTPD, which replaces the legacy FTPD. In addition to supporting the FTP protocol defined in RFC 959, the following new features are provided:

- Logging of transfers.

- Logging of commands.

- On-the-fly compression and archiving.

- Classification of users by type and location.

- Per-directory upload permissions.

- Restricted guest accounts.

- System-wide and per-directory messages.

- Directory alias.

- CD path.

- Filename filter.

- Virtual host support.

- Per-class limits (the ability to define "classes" of users according to their source IP addresses and/or hostnames, and to limit access according to user class).

Existing installations do not have to modify their FTP configuration unless they want to use the new features.

The major differences between legacy FTPD and the new version of FTPD are as follows:

Table 10-1   **New FTP daemon options**

`-d`	Logs debugging information in `syslog`.
`-m number of tries`	Specifies the number of tries for a `bind()` socket call.
`-a`	Enables the use of the `ftpaccess` file, which is used to configure the operation of FTPD.
`-A`	Disables the use of the `ftpaccess` configuration file.
`-i`	Logs all the files received by the FTPD server to `xferlog`.
`-o`	Logs all files transmitted by FTPD in `xferlog`.
`-L`	Logs all commands sent to the FTPD server into `syslog`.

Table 10-2   **New Commands**

`/usr/bin/ftpcount`	Shows current number of users per class
`/usr/bin/ftpwho`	Shows current process information for each user.
`/usr/bin/ftpshut`	Creates shutdown message file.

Table 10-2   **New Commands**

`/usr/bin/ftprestart`	Removes the shutdown message file created by the `ftpshut` utility.

Table 10-3   **New Configuration Files**

`/etc/ftpd/ftpaccess`	The primary configuration file defining the operation of the new FTP daemon.
`/etc/ftpd/ftpconversions`	Defines options for compression/decompression and `tar`/un-`tar` operations.
`/etc/ftpd/ftphosts`	Lets you allow/deny FTP account access according to source IP addresses and hostnames.
`/etc/ftpd/ftpgroups`	The group password file for use with the SITE GROUP and SITE GPASS commands.

Table 10-4   **New Logging Information**

`/var/adm/syslog/xferlog`	This file contains logging information from the FTP server daemon.

Virtual FTP Support

If you wish to manage an `ftp` server for two separate domains on the same machine, the virtual `ftp` server feature can be used. This allows you to configure systems, so that a user `ftp`'ing to `ftp.domain1.com` gets one `ftp` banner and `ftp` directory, and a user `ftp`'ing to `ftp.domain2.com` gets another banner and directory even though they are on the same machine and use the same ports.

NOTE   Setting up a virtual ftp server requires IP address aliasing. This is supported in HP-UX 10.30 and later.

Table 10-5   **Support Tools**

`/usr/bin/ckconfig`	Verifies path names of all FTP configuration files.

Secure Version of FTPD

new at 11i original release

At 11i, a unified binary is available for the new version of FTPD that can operate as both a Kerberos and non-Kerberos service.

To have the new FTPD operate in a secure environment, you enable the secure environment with the following command:

`/usr/sbin/inetsvcs_sec enable`

This updates the system file `/etc/inetsvcs.conf` with an entry `kerberos true`. At run-time, the services obtain the type of authentication mechanism to use.

Changes to rwhod

**updated for
December 2001**

The `rwhod` daemon now accepts hostnames with supported characters as per RFC 952 only. The supported characters include letters, digits, and the hyphen (-) sign. Hostnames with invalid characters, such as underscores, are now ignored by the `rwhod` daemon.

For more information on RFC952, please see `http://ietf.org`.

STREAMS/UX

new at 11i original release

Several enhancements have been made to STREAMS/UX, including support for the `select()` system call, an I/O forwarding mechanism, and Function Registering:

• The `select()` system call for STREAMS/UX devices examines the files or devices associated with the file descriptors specified by the `bitmasks`, `readfds`, `writefds`, and `exceptfds`.

  The `select()` system call can detect out-of-band (OOB) data on TCP by calling an internal command, `hpstreams_select_int2()`, which contains a check in the exception case for `T_EXDATA_IND` messages.

• STREAMS/UX contains an I/O forwarding mechanism that preserves the order of messages and forwards those messages. This mechanism is particularly useful on multi-node systems where driver events can only be executed on the node where the NIC resides.

• Function Registering enables modules and drivers to work in a mixed mode system. It provides the modules and drivers within the kernel a mechanism for correctly translating data that is being sent to and from the application when STREAMS/UX determines that the application has been compiled for 32-bit execution, but is operating on a 64-bit architecture.

  Function registering defines dynamic data structures and stream head flags, which will indicate when and if a dynamically specified function is to be executed. These data structures and flags can be set dynamically or on the fly.

NOTE

UP Emulation will no longer be supported on HP-UX in a future release. Therefore, drivers and modules that are configured as UP emulation drivers and modules should be made MP scalable in preparation.

For more information about these changes, see the *STREAMS/UX for the HP 9000 Reference Manual*.

Low Bandwidth X Extension (LBX)

The Low Bandwidth X extension (LBX) uses several compression and local caching techniques to improve performance on wide-area networks and on slower speed connections. These techniques reduce the amount of protocol data transported over the network and reduce the number of client-to-server round trips required for common application startup operations.

LBX is implemented in two pieces: an X server extension and a proxy application. The X server extension provides the new optimized protocol. The proxy application, `lbxproxy`, translates a normal client X protocol stream into an LBX stream. This permits any existing application to gain the benefit of the optimized protocol with no changes. The proxy is especially useful when multiple applications are running on the same local area network separated from the X server by a slower network. In this case, the full benefit of the local cache is shared by each application using the same proxy process.

The `lbxproxy` binary has been added to the `/usr/bin/X11` directory. It must be started by an end user either directly or through the Proxy Manager (`proxymngr`) and Find Proxy (`xfindproxy`).

Performance Issues

When X clients are separated from the X server by a slow connection such as a modem, performance will be improved by going through `lbxproxy`. However, when the client and X server are separated by a fast connection such as a local area network, performance may be degraded by running through `lbxproxy`.

Proxy Manager (proxymngr)

The Proxy Management Protocol is an ICE-based protocol that provides a way for application servers to easily locate proxy services such as the LBX proxy. (LBX is currently the only supported proxy service.)

Typically, a service called a "proxy manager" is responsible for resolving requests for proxy services, starting new proxies when appropriate, and keeping track of the available proxy services. The proxy manager strives to re-use existing proxy processes whenever possible.

The `proxymngr` executable has been added to the `/usr/bin/X11` directory. It must be started directly by the user. The `proxymngr` executable can also be used in conjunction with `xfindproxy`, which is also in `/usr/bin/X11`.

Remote Execution (RX) Service

The remote execution (RX) service specifies a MIME format for invoking applications remotely (for example, via a Web browser). This RX format specifies a syntax for listing network services required by the application (for example, an X display server). The requesting Web browser must identify specific instances of the services in the request to invoke the application.

There are two methods to demonstrate this service:

1. `xrx` (the helper program)

The `xrx` helper program has been added to the `/usr/bin/X11` directory. End users must set up their Web browsers to use this program for files with the `rx` extension.

2. `libxrx.6.3` (the Netscape plug-in)

The Netscape plug-in, `libxrx.6.3`, has been added to the `/usr/lib/X11R6` directory. End users must copy this to their `$(HOME)/.netscape/plugins` directory (or the equivalent) so that files with the `rx` extension are interpreted correctly. In order to use the plug-in, do *not* set up Netscape to also use the helper program.

Security Extension

The security extension adds the X protocol needed to provide enhanced X server security. This extension adds the concepts of *trusted* and *untrusted* clients. The trust status of a client is determined by the authorization used at connection setup. All clients using host-based authorization are considered trusted. Clients using other authorization protocols may be either trusted or untrusted depending on the data included in the connection authorization phase.

When a connection identifying an untrusted client is accepted, the client is restricted from performing certain operations that would steal or modify data that is held by the server for trusted clients. An untrusted client performing a disallowed operation will receive protocol errors.

When a client is untrusted, the server will also limit the extensions that are available to the client. Each X protocol extension is responsible for defining what operations are permitted to untrusted clients; by default, the entire extension is hidden.

Application Group Extension (XC-APPGROUP)

The application group extension provides new protocol to implement Application Groups (`AppGroups`). The `AppGroup` facility allows other clients to share the `SubstructureRedirect` mechanism with the window manager. This allows another client (called the application group leader) such as a Web browser to intercept a `MapRequest` made by a third application and re-parent its window into the Web browser before the window manager takes control. The `AppGroup` leader may also limit the screens and visuals available to the applications in the group.

This extension, along with the Netscape remote execution plug-in, allows Netscape to run programs remotely over the Web with the output appearing in the Web browser display.

The only way for an application to become a member of an `AppGroup` is by using an authorization generated using the new security extension. Whenever an application connects to the server, the authorization that it used to connect is tested to see if it belongs to an `AppGroup`. This means that the authorization data must be transmitted to the remote host where the application will be run. In the case of X, HTTP is used to send the authorization. Sites that have concerns about sending un-encrypted authorization data such as `MIT-MAGIC-COOKIE-1` via HTTP should configure their Web servers and Web browsers to use SHTTP or SSL.

SLS/d - Distributed SLS (HP Visualize Center Support)

SLS/d is an extension of the SLS (Single Logical Screen) functionality provided by the X server that allows the X desktop to span graphics displays that reside on distributed systems. By distributing the display across several systems, a larger *logical* array of

graphics displays can be achieved than otherwise would be possible with a single system with multiple graphics cards. SLS/d provides the X Window system support for part of the 3-D Visualize Center products.

SLS/d involves a low-level change in the X server that unites several distributed graphics displays into a *logica*l X Window system. The only user-visible changes are related to system configuration. The X Window system API remains unchanged in the SLS/d system, and thus is completely transparent to 2-D X window applications. The motivation behind this new functionality is to increase the size of the *logical* screen beyond what is possible using a single system with multiple graphics cards.

A new driver and a new X server extension have been added to the X server in order to implement this change. The functionality is enabled by modifying the server's X* screens file. The full documentation for the SLS/d functionality can be found in the X server information file, `/usr/lib/X11/Xserver/info/screens/hp`, and in the *Graphics Administration Guide.*

An SLS daemon and a configuration tool are delivered to aid system configuration. The daemon is controlled via start and stop scripts that reside in the `/sbin/init.d`, `/sbin/rc1.d`, and `/sbin/rc2.d` directories. The SLS daemon is started when the system enters run-level 2 or greater, and stopped when the system enters run-level 1. See the X server documentation for more details.

The performance of SLS/d depends on the performance of the underlying network to which the SPUs in the system are connected. On a dedicated network with a 100 Base-T backbone, the 2-D X Windows performance approaches that of a single SPU SLS system.

SLS/d is transparent to applications in the same manner as SLS. Once the system has been configured, it behaves identically to a single screen X Window system, albeit with a much larger screen size. One requirement is that the underlying graphics cards in the system be homogeneous. Although not a strict requirement, it is also desirable that the systems participating in the SLS/d system be homogeneous as well.

11 Security

Generic Security Services for Developing Secure Applications

**new at 11i
original release**

The Generic Security Services Application Programming Interface (GSS API) is a newly introduced product for HP-UX 11i. It contains all the GSS APIs as per RFC 2743 and is implemented as C programming language interfaces as defined in the RFC 2744, "Generic Security Service API: C-bindings." It provides security services for applications independent of various underlying security mechanisms. GSS API is also independent of communication protocols. The GSS API is available as a separate shared library. The security services available to an application include authentication, integrity, and confidentiality services.

A set of GSS APIs is already available in `libdce` libraries, which are a part of the DCE Core product in this release, as well as in previous HP-UX releases. However, these GSS APIs are dependent on the DCE security mechanism and cannot be used as general purpose APIs.

Because of GSS API independence, an application developer writing secure applications need only write the code once and need not change it whenever the underlying security mechanism changes. This will prove to be quite advantageous during this period where security technology changes are rather frequent.

An application developer who uses the GSS API C-binding interfaces will need to include `/usr/include/gssapi.h` in the program and will need to link with `libgss.sl`. The underlying security mechanism and its library can be specified in a configuration file called `/etc/gss/mech`. The library will then dynamically load the corresponding mechanism-specific shared library (for example, `libgssapi_krb5.sl` in the case of Kerberos). The default mechanism configuration file is `/etc/gss/mech`, which can be altered with the environment variable called `GSSAPI_MECH_CONF`.

In addition to this configuration file, there are two other configuration files, namely `/etc/gss/qop` and `/etc/gss/gsscred.conf` for `libgss.sl`:

*   The `/etc/gss/qop` file contains information about the GSS API-based quality of protection (QOP) for each underlying security mechanism.

*   The `/etc/gss/gsscred.conf` is a configuration file that selects how the underlying mechanism stores the `gsscred` table. The `gsscred` table is used to store the mapping between a security principal and the UNIX uid. In this release, the supported `gsscred` backend mechanism is only flat files. Therefore, the entry "files" must be specified in `/etc/gss/gsscred.conf` for the successful operation of the library.

The 32-bit and 64-bit versions of `libgss.sl` library is available at the `/usr/lib` and `/usr/lib/pa20_64` directories respectively.

Symbol Clashes

Since the symbols of GSS APIs in the `libdce` libraries clash with the symbols of `libgss.sl`, application programmers who want to use GSS API and DCE together may need to resolve the symbol clashes by linking the `libgss.sl` library first and then the `libdce` library.

Size Requirements

A minimum of 32MB RAM and 1.5MB hard disk space will be required for installation and usage of the product on HP-UX 11i systems.

Compatibility

The `libgss.sl` library has been tested with the Kerberos V5 backend mechanism library (`/usr/lib/gss/libgssapi_krb5.sl`) and is fully compatible. This library is in the KRB5-Client Software. See the next section for more information.

Documentation Changes

There are new manpages for each of the APIs of the GSS API product under the `/usr/share/man` directory. These manpages are different from the manpages for DCE GSS API which is available under the `/opt/dce/share/man` directory. For general information about GSS API, refer to the *gssapi* (5) manpage and for information about `libgss.sl`, refer to the *libgss* (4) manpage.

There is also information about GSS API in *Network Security Features for HP-UX 11i* at:

`http://www.unixsolutions.hp.com/products/hpux/hpux11/whitepapers/netsecur.pdf`

Execute Protected Stacks

**new at 11i
original release**

System security can be improved by enabling a new feature that execute protected program stacks.

A common method of breaking into systems is by maliciously overflowing buffers on a program's stack. Malicious unprivileged users often use this method to trick a privileged program into starting a superuser shell for them, or similar unauthorized actions. Detailed information on this type of attack may be found by doing a web search for "Smashing the Stack for Fun and Profit."

HP-UX 11i provides new mechanisms to defend against this type of attack without sacrificing performance.

By setting the kernel tunable parameter `executable_stack` to zero, HP-UX systems can be configured to execute protect program stacks, providing significant protection from many common buffer overflow   attacks. In the vast majority of cases, enabling this feature will not affect compatibility of any legitimate applications.

Please refer to the new `+es`   option section of the *chatr* (1) manpage for additional information on how to configure this feature and how to quickly detect and resolve any (very rare) compatibility issues that may result from enabling it.

To implement this feature, changes were made to kernel `execve()` and   virtual memory code, and to the `chatr`, `elfdump`, and `ld` commands.

Impact

One of the primary goals of this feature is to significantly improve   system security with the minimum possible effect on performance or   compatibility. It consumes essentially no disk space or memory, and has no functional impact on the vast majority of legitimate applications,   other than making them less vulnerable to malicious attacks. There is no measurable performance impact from this code.

Compatibility

In the default configuration, HP-UX is unaffected by these changes.   Users who want to use this feature must explicitly enable it by setting the kernel tunable parameter `executable_stack` to 0. HP strongly encourages you to enable this feature. Refer to the `+es`   section of the *chatr* (1) manpage for details of the possible trade-offs between security and compatibility.

ELF-64 programs linked on previous releases of HP-UX will not benefit   from this security feature until they are re-linked on HP-UX 11i or later, but will still function normally. 32-bit applications do not need to be re-linked.

The output of `chatr` and `elfdump` have changed slightly. `chatr` now supports an `+es` option.

IMPORTANT

Warning to Java Users

Disabling stack execution will cause Java 1.2 programs to fail if using JDK/JRE 1.2.2 versions older than 1.2.2.06. To determine the Java version you are using, run `java`

-version. **To download the latest version of the JDK/JRE, see**
http://www.hp.com/go/java.

To allow pre-1.2.2.06 programs to run, the executable from stack **attribute of the**
program must be set to enable. **To do this, invoke** chatr +es enable *file*, **where** *file*
is the executable file. This attribute will need to be set to enable **for all executables**
contained in the JDK and JRE. This includes all files contained in the following
directories:

```
/opt/java1.2/bin/PA_RISC/native_threads
/opt/java1.2/bin/PA_RISC2.0/native_threads
/opt/java1.2/jre/bin/PA_RISC/native_threads
/opt/java1.2/jre/bin/PA_RISC2.0/native_threads
```

Java 1.1 versions will execute with no problem.

Auditing Commands/System Calls To Be Obsoleted

The auditing commands audevent, audisp, **etc. and the system calls** audwrite,
audswitch, **etc. will be obsoleted in a future release. An interface will be provided in the**
form of a device driver, /dev/idds, **with additional functionality.**

At that time, both /dev/idds **and the current 11i auditing process will be supported for**
ease of transition.

Configurable Security Features

**new at 11i
original release**

Administrators now have a new convenient way to customize security features. A new `/etc/default/security` file is defined. Editing this file provides a way to configure new security features or to modify the behavior of existing security features.

A `PASSWORD_HISTORY_DEPTH=<n>` parameter can be added to `/etc/default/security` to enable a new password history feature, which forces users to choose passwords that do not match their most recent <n> passwords.

A `MIN_PASSWORD_LENGTH=<n>` parameter can be added to `/etc/default/security` to force users to choose passwords which have at least <n> characters.

A `SU_ROOT_GROUP=<groupname>` parameter can be added to `/etc/default/security` to allow users to su to root only if they are a member of the `<groupname>` group.

See *security* (4) for additional parameters and details.

Password History Feature on Trusted Systems

Password history is a new trusted-system feature of the `passwd` command, used to discourage users from re-using previously used passwords.

The system administrator enables the system-wide password history feature by creating (or opening, if it already exists) a file called `/etc/default/security` and appending an entry:

`PASSWORD_HISTORY_DEPTH=number`

Depending on the value of number (decimal integer from 1 through 10), the system checks the user's new password against that number of previously used passwords and prevents their usage. (For example, if number=5, the system will not allow a user to use any of the last five passwords he or she has previously used.)

Structurally, the password history feature is accomplished by a shared library, called `libpam_unix.1`, which is dynamically loaded at run time by the command. This structural characteristic is totally transparent to users; the end-user interface of the command is unchanged.

For further information, consult the *passwd* (1) manpage.

Kerberos Client Software

**new at 11i
original release**

Kerberos is a network authentication protocol. Kerberos Client Software, now provided with HP-UX 11i, enables integrating HP-UX into a secure enterprise environment. It provides tools and libraries to perform authentication and secure communication.

The Kerberos protocol is designed to provide strong authentication for client/server applications by using secret-key cryptography. It uses strong cryptography so that a client can prove its identity to a server and vice versa across an insecure network connection. After the client and the server have established their identities, they can also encrypt all of their communications to assure privacy and data integrity.

Kerberos Client Software is based on MIT Kerberos V5 1.1.1. It consists of libraries, header files, manpages, and Kerberos utilities which help in performing command line or programmatic authentication. Data encryption APIs can be used to protect data transmitted over the Internet. Kerberos Client Software supports both 32- and 64-bit development. The 64-bit libraries are placed in the `/usr/lib/pa20_64` directory.

Libraries

The following libraries are included:

- `/usr/lib/libkrb5.sl`, `/usr/lib/pa20_64/libkrb5.sl`:

  All of the Kerberos APIs are implemented by this library. This library implements APIs for authentication, verifying tickets, creating authenticator, context management, etc. For more information see *libkrb5* (3).

- `/usr/lib/libcom_err.sl`, `/usr/lib/pa20_64/libcom_err.sl`:

  This library implements `com_err` APIs. The `com_err()` functions print appropriate error messages to the `stderr` based on the error code returned by Kerberos APIs. For more information see *libkrb5* (3).

- `/usr/libk5crypto.sl`, `/usr/lib/pa20_64/libk5crypto.sl`:

  This library provides APIs for encryption and decryption. Internally, it uses DES (Data Encryption Standard). Currently, it supports 56-bit DES and is used by the Kerberos APIs. For more information see *libkrb5* (3).

- `/usr/lib/gss/libgssapi_krb5.sl`, `/usr/lib/pa20_64/gss/libgssapi_krb5.sl`:

  This contains the Kerberos support for GSS API as per RFC 2743/2744. This library is used by `/usr/lib/libgss.sl`, which is part of the GSS API product. For more information, see *libgss* (4) and *gssapi* (5) and the previous section.

Header Files

- `/usr/include/krb5.h`
- `/usr/include/profile.h`
- `/usr/include/com_err.h`

Utilities

- `/usr/bin/kinit`: obtain and cache the Kerberos ticket-granting ticket. See *kinit* (1).
- `/usr/bin/klist`: list cached Kerberos tickets. See *klist* (1).
- `/usr/bin/kdestroy`: destroy Kerberos tickets. See *kdestroy* (1).
- `/usr/bin/kvno`: print key version numbers of Kerberos principals. See *kvno* (1).
- `/usr/bin/kpasswd`: change a user's Kerberos password. See *kpasswd* (1).
- `/usr/sbin/ktutil`: Kerberos `keytab` file maintenance utility. See *ktutil* (1).

Manpages

- Manpages in `/usr/share/man/man1.Z` directory: *kinit* (1), *klist* (1), *kdestroy* (1), *kvno* (1), kpasswd (1), and *ktutil* (1)
- Manpages in `/usr/share/man/man4.Z` directory: *krb5.conf* (4)
- Manpages in `/usr/share/man/man3.Z` directory: *libkrb5* (3)

Special Considerations

Developing Secure Applications

Though Kerberos APIs are made available, these are for supporting existing Kerberos Applications to HP-UX 11i. Application developers are strongly encouraged to use GSS API for developing secure applications. See *gssapi* (5) for details.

libsis.sl

Most of the KRB-Support (`libsis.sl`) functionality is now available with Kerberos Client Software. It is recommended that developers compile and link with these libraries.

Unsupported Features

- Kerberos Client Software does not support Triple DES due to U.S. export regulations.

- Kerberos Client libraries are not thread safe.

Size Requirement

Kerberos Client Software requires 5MB of disk space.

Compatibility Issues

- Kerberos V5 1.1.1 Client Software is compatible with earlier versions of the Kerberos product supporting RFC 1510.

- Kerberos Client Software only supports the Kerberos 5 protocol as per RFC 1510. The product does not support the Kerberos 4 protocol and Kerberos 4 to Kerberos 5 request conversions.

HP-UX Kerberos Server

**new for
September 2001**
HP announces a new security product available on the application CD distributed in September 2001. The HP-UX Kerberos Server (T1417AA) provides key distribution facilities to implement the Kerberos authentication protocol in network-distributed enterprises. It is designed to provide strong authentication for client/server applications by using secret-key cryptography. A client can prove its identity to a server (and vice versa) across an insecure network connection. After a client and server have used Kerberos to prove their identity, they can also encrypt all of their communications to assure privacy and data integrity as they go about their business.

Single sign-on:
Using the Kerberos protocol, users have the foundation for secure single sign-on to applications and resources. The server stores user profile data. Clients initially use a password that is converted into an authorization ticket by the server. This authorization then creates a service ticket, which is used in all applications and services that have been Kerberized to authenticate the user and provide access to applications. In this way, a single sign-on provides credentials to automatically access multiple applications and services wherever they reside on the network.

**Cross-realm
authentication:**
The server provides both an authentication service as well as acts as a key distribution center (KDC). An implementation of MIT Kerberos 5 version 1.2.2, the server supports cross-realm authentication. One use is to work with Windows clients who gain Windows 2000 Kerberos credentials. These are then used to authenticate the user to the HP-UX Kerberos server which, in turn, creates credentials for HP-UX applications and services, all with a single sign-on.

GSSAPI support:
For development, HP-UX Kerberos provides a Generic Security Services Application Programmer Interface (GSSAPI). The GSSAPI provides a standard programming interface that is authentication-mechanism independent and is supported on HP-UX 11.0 and 11i. This allows application developers the flexibility of using alternative authentication technologies, including Kerberos.

The implementation of the Kerberos protocol, Kerberos server, Kerberos client, PAM Kerberos, and Kerberized applications provides an infrastructure of DES encryption and single sign-on ease for users in a network-distributed computing environment.

12 Compatibility

Compatibility from HP-UX 11.0 to 11i

Hewlett-Packard has a long record of providing HP-UX compatibility. Because it protects your investment and allows you to upgrade easily, compatibility is an important feature that HP has always recognized and that HP customers have come to expect.

Compatibility requirements span across HP-UX products to third-party products as well. All third-party products (and those products they call) are equally important components in the complete customer environment. Customer solutions often have complex, multiple chains of dependent applications spanning the entire range of HP-UX products as well as third-party products. One broken link in the chain of dependencies may result in an application that no longer works. Support for the unbroken string of compatibility on HP-UX is one of the biggest and best benefits provided by HP.

HP-UX supports forward compatibility from 11.0 to 11i. This chapter will describe what this means for executable applications, object files, source files, data, and libraries. Compatibility exceptions will also be discussed. (For additional compatibility exceptions, see also "EISA Interface Cards Are Not 64-bit Compatible" on page 119 and Chapter 13, "Programming," on page 209.)

IMPORTANT   **Superdome administrators:**
Use the new Superdome Machine Identifier to guarantee compatibility.

Because the `uname -i` command on your Superdome systems may not return a unique value for each system, you should use the new interfaces to *getconf* (1) and *confstr* (3C) to retrieve unique machine identifiers (and thereby guarantee compatibility on current and future platforms).

These interfaces are documented in the manpages and in Chapter 13, "Programming," on page 209 of this document.

General Compatibility Concerns

The following types of compatibility are supported from 11.0 to 11i for well-behaved applications:

- Binary compatibility
- Source compatibility
- Data compatibility
- Relocatable object compatibility
- Upgrade compatibility

(For known exceptions to compatibility, see "Known Compatibility Exceptions from HP-UX 11.0 to 11i" on page 198.)

A well-behaved application adheres to the following characteristics:

- Uses only documented, public APIs

- Adheres to the required practices that are specifically documented

- Does not use documented features that are specifically described as having platform, architecture, or configuration limitations

- Does not decompose an HP-UX product and then reuse the results of the decomposition

---

NOTE For compatibility issues relevant to a particular component, see the corresponding section elsewhere in this document.

Binary Compatibility

An application that has run on HP-UX 11.0 will continue to run with the same behavior on 32-bit and 64-bit HP-UX 11i. This includes executables, binary files that have been processed by the HP link editor with `ld` or indirectly with the compiler, and can be run by the HP-UX loader (`exec`).

Source Compatibility

Software that has been compiled on an HP-UX 11.0 release can be recompiled without change on HP-UX 11i. The term *source* includes input source to compilers, scripts, and makefiles.

Data Compatibility

An application can continue to access persistent data files (such as system files, backup/recovery formats, and HP-documented data formats) via supported APIs in the same manner as the previous release. For example, applications should access the password file information via `getpwent()` rather than directly reading the file in order to maintain data compatibility.

Relocatable Object Compatibility

A relocatable object can be a file (`.o`), shared library (`.sl`), or an archive library (`.a`). Several types of object binary compatibility are below (note that some executables are *not* supported):

- **Release-to-release relocatable object binary compatibility:** If an executable is created by linking with forward-compatible, relocatable objects from different releases—or by using `shl_load()` and `dlopen()` to dynamically load shared libraries built on a different release—than the application is *only supported* from HP-UX 11.0 to 11i.

  However, linking pre-HP-UX 11.0 libraries and HP-UX 11.0 and 11i libraries in one relocatable object/executable is *not* supported.

CAUTION Even though the linker will permit the linking of pre-HP-UX 11.0 libraries and HP-UX 11.0 and 11i libraries in one relocatable object/executable (and will not exhibit any warning or error messages), the executable may exhibit incorrect behavior.

- **Archive and shared relocatable object compatibility:** An executable that is created by linking with a shared library that has dependencies on an archive library (a situation that typically occurs when linking with archive system libraries) is *not* supported.

- **Data model relocatable object compatibility:** An executable created by linking with a mixture of 32-bit and 64-bit objects is *not* supported. The loader will not permit this.

Upgrade Compatibility

Customized configurations and data from HP-UX 11.0 are preserved upon installation and upgrade to HP-UX 11i.

HP-UX 10.*x* Applications on HP-UX 11i

HP-UX 10.*x* applications that have been compiled and ran on 11.0 can be recompiled and run on HP-UX 11i without change.

Known Compatibility Exceptions from HP-UX 11.0 to 11i

In the following sections, a short description of an 11i change is followed by details of the compatibility exception with which it is associated. Please note that all of these compatibility exceptions are *rare corner cases* for well-behaved applications. (The exceptions have been arranged alphabetically in four groups: 1) Library-Related; 2) Miscellaneous; 3) Networking, Internet Services, and Security; and 4) Software Distributor.)

Library-Related

- **Customization of `wctype` Methods**

  To provide performance improvements, this change removes the ability to use customized locale methods for accessing `wctype`, wide-character classification APIs. If an application is built for locales with *localedef -m* and the method library includes custom functions for `iswalpha()`, `iswupper()`, `iswlower()`, `iswdigit()`, `iswxdigit()`, `iswalnum()`, `iswspace()`, `iswpunct()`, `iswprint()`, `iswgraph()`, `iswcntrl()`, `wctype()`, `iswctype()`, the application should now be linked with the method library and call the method functions directly.

- **`libc qsort()` Algorithm Change**

  This change improves performance by enhancing `qsort()` so that it sorts "tied" elements differently than the previous implementation. Well-behaved applications are not affected, since the manpage warns that the order in the output of two equal items is unpredictable.

- **`libc SYSTEM_ID callbraph` Change**

  The `callbraph` of `libc` has changed. As a consequence, applications that have been linked to the archival version of `libc` (as well as any shared libraries linked to that application) may fail.

NOTE Linking an application with a shared library that depends on an archive library is *not* a supported configuration. Applications linked in this way do not qualify as well-behaved because this configuration is *not* supported.

- **`libc atof()` Algorithm Change**

  This change fixes a defect in `atof()` to convert denormalized floating point numbers correctly. Applications which disregard the recommended coding practice of using floating point ranges (rather than relying on specific hard-coded floating point numbers) can be affected.

- **Linker Support for PBO of 64-bit Shared Libraries**

  This change enables 64-bit PBO to function with shared libraries. Only those who link `-noshared` instrumented applications and try to use `HP_LD_FDP_INIT` to specify an alternative version of `fdp_init.o` will be affected. If this is the case, you will have to use `HP_LD_FDP_INIT_NS` instead. If the `HP_LD_FDP_INIT_NS`

environment variable is not set and `fdp_init_ns.o` is in the default location, the link will fail with the `file not found` error message.

- **ONC+/NFS Security Correction**

  This change corrects a security problem in NIS+. However, applications that are linked to the archived version of the `libnsl` library may have a compatibility problem. (Applications linked to the *shared* version of `libnsl` will not exhibit these symptoms.) The symptoms include:

  — Daemon registration will fail when UDP/TCP is used instead of the local loopback transport device.

  — In the NIS+ environment, applications will not be able to authenticate themselves.

  — NIS+ performance degradation will occur due to not being able to contact the `nis_cachemgr`.

Miscellaneous

- **DNS Bind**

  This change is necessary to conform to the behavior found on other vendor platforms. Those who try to edit the file, `named.boot`, could find it missing, or if the file exists, they may try to edit the file, but will find that their changes have not taken effect. This is primarily a system administration change, but in the rare instance where scripts might be written to edit `named.boot`, the scripts would need to be modified to edit `named.conf` instead, both for the new file name and syntax.

- **ELF Undocumented Symbol Table Change**

  To make some tool development easier, the `ELF` symbol table type of some thirteen linker-defined symbols has changed from `STT_OBJECT` to `STT_NOTYPE`. Although the names of these symbols have been documented, their types and meaning have not. However, only applications that are not well-behaved and read 64-bit `ELF` executable files are affected.

- **Fortran 90 GETARG Intrinsic Function**

  This change causes the semantics of the `index` argument to the HP-supplied F90 intrinsic routine, `GETARG`, to be compatible with older HP F77 and other vendor implementations of this routine. Those affected will have to change and recompile their source code to use the industry-standard indexing scheme.

- **IOSCAN Usability Enhancement**

  This change improves the usefulness of the IOSCAN output for PCI interfaces. However, because the description field for PCI interface cards has been changed to be more descriptive, scripts that scan for hard-coded values may need modifications. (The description field for non-PCI devices has remained the same.) See "Improved ioscan Description Field for PCI Devices" on page 140 for more information.

- **MAX_PROCS Changed to Enable 128 CPU Support**

  This change enables support for 128 CPUs. The kernel macro `MAX_PROCS` has changed from 32 to 128 in the LP64 kernel and has changed the ABI for the undocumented system calls `ki_call()` and `ktest_ioctl()`.

  The `MAX_PROCS` change will cause an ABI incompatibility for kernel-intrusive applications or drivers which access internal kernel arrays sized by the `MAX_PROCS` macro.

- **MAXTID Removed; MAXPID Changed**

  This change raises the number of processes or threads to 8 million. `MAXPID`, which is undocumented, has been changed. An alternate mechanism to dynamically determine the value of `MAXPID` has been introduced.

- **Memory Windows**

  This change enables applications to access up to 1GB of shared memory that is not otherwise allocated against the system-wide limit. However, enabling Memory Windows alters the semantics of some memory APIs and some POSIX APIs. Although these APIs will function correctly for applications running within their own Memory Window, they will not function correctly for applications running in different Memory Windows.

- **Process-Private Memory: Increase in Memory Size Limit**

  This change increases the memory size limit for process-private memory. However, when the 3rd quadrant private feature is enabled for a process, it can only allocate shared objects in the 4th quadrant. If the 4th quadrant fills up, the application may fail, where it would not have failed if the 3rd quadrant were available for allocation of shared objects.

- **`pstat_getdynamic` Interface: Change in Maximum Number of Active Processors**

  This change corrects a defect in the `pstat_getdynamic` interface so it adheres to the documentation when it reports the number of processors that are active on a system. Well-behaved applications will not be affected by this change. However, ill-behaved applications may overestimate the number of processors that are active on the system. Ill-behaved applications can be corrected to reference the correct field with a simple code change. See "instant Capacity on Demand (iCOD)" on page 141 for more information.

- **`strftime()` Support for Week Number**

  This change fixes a defect in `strftime()`. Now, applications that use the %V option of `strftime()` to obtain the week number will find that the return value is 52 instead of 53 when:

  — December 31 falls on a Friday, in a non-leap year, when the date passed in is January 1 or January 2 of that week. Some years affected include 1999, 2004, and 2032.

  — December 31 falls on a Saturday, when the date passed in is January 1 of that week. Some years affected include 2005, 2011, 2016, 2022, 2033 and 2039.

- **Support Tools Manager User Interface, EMS Hardware Monitors**

  This change improves the usability for the STM User Interface and the EMS Hardware Monitors. As a result, any script that depends on the specific output of the EMS Hardware Monitors or specific commands or displays in the STM User Interface may have to be modified. See "HP-UX Support Tools: STM, ODE, & EMS Hardware Monitors" on page 73 for more information.

- **System V Message Queues Expanded Beyond 64KB Limit**

  This change alters the message queue data structures to support queues larger than 64KB. Consequently, if one application is built to use the larger queues, all related applications that use the same message queue must also be built to use the larger queues. See "System-V IPC Message Queue Enhancement" on page 152 for more information.

Networking, Internet Services, and Security

- **EISA Interface Cards**

  Although EISA interface cards are supported on 32-bit operating systems, they are not supported on 64-bit systems. For further details, see "EISA Interface Cards Are Not 64-bit Compatible" on page 119.

- **IPv6 IPsec**

  This change allows support for a contemporary standard. The values for the following `defines` were changed to support standards:

  — `IPPROTO_ENCAP`
  — `IPPROTO_IPIP`

- **Non-Executable Stack**

  This change improves system security. The majority of your programs should be unaffected by execute-protecting program stacks. Only those that execute instructions from their stack (typically interpreters, simulators and debuggers) are affected.

  When enabled, the new functionality causes the termination of any program attempting to execute code located on its stack. If this occurs, you will be given an error message pointing to relevant documentation that explains the reason for the process termination and how to remedy the situation.

  See "Execute Protected Stacks" on page 189 for more information.

- **NFS Mount Access Control**

  This change is necessary to conform to *de facto* industry behavior. The behavior of `access=` has been modified to conform to a common behavior. If you are using this undocumented feature to disallow the NFS mounts, it will now succeed.

NOTE   Applications that use undocumented features are not "well behaved."

- **Tighter Security for NFS Mounts**

  This change is necessary to improve the security of NFS mounts. Without this change, when you export a file system using the `root=` option of `exportfs`, NFS-clients on the `root= option` are allowed to mount the NFS file system even when they don't appear on the `rw= list` and/or `access= list`. The new behavior prevents NFS clients from mounting the file system unless they appear in either a `rw=` and/or `access=` list.

- **Export Filesystem (NFS) Security Defect Correction**

  This change corrects the NFS implementation so it conforms to industry practice when exporting a file system. Well-behaved applications will not be affected by this change. However, applications that assume that exporting a symbolic link to a file system will result in the symbolic link being exported, rather than the directory to which the symbolic link points, will fail that assumption. Shell scripts and administrative processes may have to be changed to correct the assumption.

Software Distributor (SD)

- **SD Bundle Algorithm Change**

  This change allows you to maintain a depot with multiple versions of a Software Distributor bundle and automatically get the latest version of the bundle without specifying a version qualifier. The install process no longer prints an error message when you do not qualify which version of the bundle is intended to be installed.

- **SD Automatic Generation of Depot or root Layout Version**

  This change allows you to maintain 10.20 and 11.x depots on an 11.x system. It modifies the SD commands so they do *not* change the layout version of a depot or root automatically. Any scripts or processes that rely on the automatic conversion to `layout_version=1.0` will be broken.

- **SD Log Message Simplification**

  This change makes it easier for an administrator to identify real problems when scanning the log files. The SD log files now contain less "noise" (error, warning or note messages that contain no useful information).

- **SD System Update Process**

  Although you will have to learn a new process, this change gives you a more robust and easier to use the HP-UX 11i update process.

- **SD Reduces Amount of Information from Program `swlist`**

  This change improves the performance of some `swlist` options. Extraneous data is no longer displayed and the listing of bundles in a depot shows only bundles. Applications that depend on the old format and behavior will have to be modified.

- **SD Changes Behavior with Unknown Attributes**

  This change allows packagers to use new attributes in their software packages without requiring SD to know the attribute. Specifically, the `swpackage` program will no longer print error messages when an unrecognized attribute is encountered.

However, you must be careful when naming attributes because typographical errors will no longer be reported.

Obsolescence and Deprecation of APIs

This section defines the obsolescence of core system libraries and relocatable objects. Obsolescence of other products are covered in separate sections.

Rationale and Objectives

HP's rationale and objectives for obsolescence and deprecation of APIs are as follows:

- provide common, standard APIs across UNIX vendors
- facilitate portability for our ISVs
- reduce confusion for the selection of similar APIs
- reduce the size of `libc`, thus increasing performance of shared `libc`
- reduce the continued application turbulence for future architecture changes
- remove compatibility problems for applications linked to shared libraries that have dependencies on archive system libraries
- reduce satisfaction issues with APIs that have specific defects (for example, compatibility issues)
- reduce support costs for APIs that are not moving in the strategic direction of standards, the industry, and our customers
- minimize adoption issues for new releases on PA or IA-64

The intent is that there will be *no* gratuitous changes. Obsolescence of APIs and libraries will be acceptable only when initiated to avoid application breakage or duplicate functionality.

Terms and Definitions

Deprecated: A *deprecated* interface can have the following characteristics:

- functionality is available on the system
- deprecation is a step towards obsolescence
- the specification is in flux
- has less value to users
- functionality no longer makes sense
- functionality has been replaced
- support/enhancement expectations have been lowered
- usage is discouraged
- warnings against usage/alternatives have been provided
- the provider continues to test functionality
- migration plan/tools have been provided

The reasons for marking an interface as *deprecated* may include the following:

- marked "to be withdrawn" by standards
- support is available via more standard means
- equivalent, enhanced, more reliable counterparts exist
- any or all reasons listed in the "Obsolete" section below

Obsolete: An *obsolete* interface may have the following characteristics:

- functionality is no longer available on the system
- runtime support is undefined
- cannot develop or build with this interface
- documentation is not provided or recommends against usage
- the final stage of the product life cycle has been reached

The reasons for marking an interface as *obsolete* may include the following:

- underlying infrastructure in either the software or hardware is obsolete or not available
- changes to the system have decreased reliability
- miscellaneous business decisions such as those listed below:

  — a third-party's solution exists
  — not strategic
  — support costs are too high
  — not enough ROI

Archive/Static Libraries

Most archive system libraries, such as `libc.a` (with the exception of `libc.a`, `libcres.a`, and `libsbin.a`), will be obsolete and not shipped on future releases of HP-UX, including those supporting IA-64. For the resulting benefits to you and to HP, refer to "Rationale and Objectives" on page 204.

NOTE	In most cases, your makefiles will continue to work without the need for modifications.

CMA Threads Obsolescence

Background

CMA threads (`libcma`) is a user-space implementation of POSIX P1003.1a (Draft 4), which was based on Concert Multi-Thread Architecture (CMA).

Starting at HP-UX 11.0, multi-threading was also supported in the HP-UX kernel and was known as kernel or POSIX threads (`libpthread`). The POSIX threads implementation supports the approved POSIX 1003.1c (POSIX.1-1996 Draft 10) standard, which facilitates application portability onto POSIX-compliant vendor platforms. POSIX threads also enable the application to parallelize the execution of threads on multiple processors in a multi-processor system.

CMA threads (`libcma`) have been deprecated (slated for future obsolescence) at 11i, and their development environment will no longer be shipped on future releases of HP-UX, including those supporting IA-64 (there is no plan to release native IA-64 CMA threads). Also see "Kernel Threads vs. CMA Threads" on page 144.

Options

Applications using CMA threads have the following options:

- `libcma` PA applications will continue to run on future releases of HP-UX, including those supporting IA-64, via compatibility mode.

- Applications using `libcma` should start migrating to POSIX threads (`libpthread`).

- Where the `libcma` development environment is still available, `libcma` applications can maintain their existing development environment on 11.0 to 11i in order to continue to make application defect repairs. The applications can then be deployed on future releases of HP-UX, including those supporting IA-64.

Customer Transition Aids

Transitioning from CMA threads to POSIX threads is not a trivial endeavor. To help you with the transition, the 11.*x*/IA-64 Software Transition Kit (STK) provides tools and documentation transition aids at:

http://devresource.hp.com/STK

Additional transition aids include the following:

- The *Porting DCE Threads Programs to HP-UX 11.0 POSIX Threads* white paper at:

    http://devrsrc.external.hp.com/devresource/Docs/TechPapers/PortThreads.html

- STK tools that can detect `libcma` usage in customer code/binaries, available at:

    http:/devresource.hp.com/STK

- The *Introduction to Kernel Threads* white paper at:

    http://devresource.hp.com/STK/partner/threads.html

List of APIs to be Deprecated/Obsoleted

The following table provides a summary of the APIs that have been deprecated and/or obsoleted:

Table 12-1   APIs to be Deprecated/Obsoleted

Library/API	Description	Release Deprecated	Native on IA-64	Comments
Entire Libraries				
`libc.a` `pa20_64/libc.a`	Archive/static `libc`	11i 11i	No No	
`libp/libc.a` `pa20_64/libp/ libc.a`	Archive profile `libc`	11i 11i	No	
`libpicc.a`	Build custom `libc`	11i	No	
`libPW.a`	ATT Programmer's Workbench	10.30	No	Comparable APIs are in `libc`.

Table 12-1   APIs to be Deprecated/Obsoleted

Library/API	Description	Release Deprecated	Native on IA-64	Comments
libBSD.a	BSD 4.2 library	10.30	No	Comparable APIs are in libc.
/usr/old/ libmalloc3x.a /usr/old/ malloc3c.o	Old malloc() relocatable objects for compatibility with pre-9.x	10.01 11i	No	Use libc malloc().
libcma.a libcma.1 libcma.2	CMA threads	11.0 & 11i	No No No	Use libpthread().
libc APIs				
memorymap()	Display the contents of the memory allocator. 32-bit only (no 64-bit available)	11i	No	Use mallinfo() instead.
blockmode() family blclose(), blget(), blopen(), blread(), blset(), <blmodeio.h>	HP proprietary terminal interfaces	10.30	No	Use libxcurses() instead.
File system descriptor file entry 4.2 BSD: endfsent(), getfsfile(), getfstype(), getfsent(), getfsspec(), setfsent()	File system APIs for compatibility with 4.2 BSD.	10.30	No	Use getmntent() APIs instead.
gettxt(), setcat()	SVID message catalog facility	11i	Yes	Use catopen(), catgets() instead.
sys_errlist(), sys_nerr()	Array of message strings and largest message number in the array.	11i	Yes	Use strerror() instead.
ptrace(), ptrace64()	Process trace	11i	No	

Table 12-1   APIs to be Deprecated/Obsoleted

Library/API	Description	Release Deprecated	Native on IA-64	Comments
`nl_tools_16()` APIs and Macros: `byte_status()`, `firstof2()`, `secoof2()`, `c_colwidth()`	Tools to process 16-bit characters.	10.0	No	
Derived Definitions for Header files				

`_INCLUDE_AES_SOURCE`
 	
`_XPG4_EXTENDED`
	Replaced by `_INCLUDE_XOPEN_SOURCE_EXTENDED`	11i	No	

`_SVID2`	No longer supported.	11i	No	
`_XPG2`	No longer supported.	11i	No	
`_XPG3`	No longer supported.	11i	No	
`_XPG4`	Replaced by `_XOPEN_SOURCE`	11i	No	
`_POSIX1_1988`	Replaced by `_INCLUDE_POSIX_SOURCE`.	11i	No	

`_CLASSIC_ANSI_TYPES`
 	
`_CLASSIC_POSIX_TYPES`
 	
`_CLASSIC_XOPEN_TYPES`
 	
`_CLASSIC_ID_TYPES`
	Supported in HP-UX 7.*x*, 8.*x* for HP-UX 6.*x* compatibility.
	
	 HP-UX compatibility is not required for 10.*x*.
		11i	No	

13 Programming

What's in This Chapter?

This chapter covers the topics listed below. (For other relevant information, also see Chapter 12, "Compatibility," on page 195.)

- HP-UX Software Transition Kit (STK) (see page 210)

- Libraries (see page 211)

  — aC++ Runtime (libCsup*, libstd*, libstream*, librwtool*) (see page 211)

  — Judy Libraries (see page 212)

  — Changes to libc (see page 213)

  — Overall libc Performance Tuning (see page 214)

  — Performance Improvements to libc's ftw() and nftw() (see page 216)

  — Performance Improvements to libc's malloc() (see page 217)

  — The libcres.a Library (see page 219)

  — Changes to libm (see page 220)

- Miscellaneous (see page 221)

  — Linker and Object File Tools (ld, crt0.o, dld.sl, libdld.sl, chatr, and odump) (see page 221)

  — Changes to the linker/dld Interface (see page 222)

  — Instrumented Code Using PBO or +O4 Optimization (see page 223)

  — HP DCE/9000 (see page 223)

  — Extensions to pstat() (see page 225)

  — Changes to sendfile (see page 227)

  — Machine Identifier Changes to confstr (see page 227)

HP-UX Software Transition Kit (STK)

The HP-UX Software Transition Kit (STK) aids in transitioning your software to either the 32-bit or the 64-bit version of HP-UX 11i. To do this, many tools are available to help you resolve issues such as those involving data models and API changes. API file scanners are provided in the HP-UX STK, while other tools are part of the HP-UX operating system, are included in HP-UX language products, or are supplied by third parties.

The HP-UX STK provides step-by-step instructions for performing transitions, a complete set of background and technical documents, and file scanners to help you identify and resolve any required API changes in your source files.

The HP-UX Software Transition Kit will scan the following types of source files:

- C and C++ programs
- FORTRAN programs
- COBOL programs
- scripts
- makefiles

The STK file scanners can help you locate and fix any of the following that have changed or become obsolete:

- functions
- commands
- path names
- macros
- structures and structure members
- language keywords
- libraries
- variables

One of the HP-UX STK file scanners, `scansummary`, can help you plan your transition by summarizing the number and type of API impacts in your source files. The other tool, `scandetail`, can help you resolve those impacts by identifying the file name and line number where each impact occurs. Both tools provide links to more detailed information about each impact. (The file scanners can also identify opportunities for using some of the enhanced features of HP-UX 11i.)

To use the HP-UX STK, you must install it. The HP-UX STK is available free of charge via the Web at **http://devresource.hp.com/STK**.

updated for March 2002

Check this site often for updated content. (Now, for instance, the HP-UX STK version 1.7 includes tools and documentation to help you successfully transition your software from PA to IPF platforms.)

Libraries

aC++ Runtime (libCsup*, libstd*, libstream*, librwtool*)

**new at 11i
original release**

The aC++ runtime provides the run-time environment necessary for deploying C++ based (aC++ compiled) applications on HP-UX 11i.

This release of the aC++ runtime includes a new ANSI-compliant Standard C++ library. (The previous version of the runtime included the "classic" C++ STL library that corresponded to the pre-standard [Sept. 1998] definition of the C++ language and library.) Although the updated C++ runtime included for HP-UX 11i retains the "classic" C++ library functionality, it also includes new components (`libstd_v2` and `libCsup_v2`) that introduce a standard-compliant set of C++ interfaces, as required by the ISO/IEC 14882 Standard for the C++ Programming Language.

The added components, `libstd_v2` and `libCsup_v2`, are new libraries with functionality that did not exist prior to this release of the C++ runtime. The details of the newly added libraries are covered in the following files (which are available after installation of the aC++ product, version A.03.25 or later):

- `file:/opt/aCC/html/libstd_v2/stdug/index.htm`

- `file:/opt/aCC/html/libstd_v2/stdref/index.htm`

Over time, with the acceptance of the new library, the old, "classic" library is expected to be deprecated and possibly removed from future operating system releases.

Impact

Overall (file) size of the C++ runtime will increase by about 44%, with 10 new libraries.

This release provides access to the standard compliant C++ library for application developers (and deployment of such applications). This is by far the most heavily requested enhancement by the users of the aC++ compiler. However, the performance of the new library (`iostreams`) may be slower.

Compatibility

C++ application (source and binary) forward compatibility with 11.x is fully maintained by preserving the classic C++ library in the new runtime. Therefore, source files, build systems, and object files or libraries produced under HP-UX 11.0 with the previous version of C++ runtime should continue to work under the new runtime.

However, the new libraries are binary incompatible with the classic C++ libraries. The option `-AA` must be used to enable the new libraries and headers.

To preserve backward source and runtime compatibility from HP-UX 11i to 11.0, application developers who develop C++ applications with the use of the new standard C++ library must ensure that the June 2000 Application Release-dependent C++ library patches (specified below) are applied to the 11.0 system.

The C++ library and Header File patches PHSS_21906, PHSS_21947, PHSS_21950, PHSS_21075, and PHSS_22217 are shown at the following web site, now a part of the Developer and Solution Partner Portal:
http://devresource.hp.com/devresource/Tools/lang.html

Documentation

Detailed manpages for the new library are included with the Independent Software Unit release. The new library is also discussed in aC++ Online Help.

Judy Libraries

The Judy product is a C language library that enables an unbounded array capability. Judy Libraries provide a state-of-the-art core technology that replaces many traditional data structions and algorithms (such as arrays, sparse arrays, hash tables, B-trees, binary trees, linear lists, skip lists, and counting functions).

updated for December 2001

New functionality has been added to the Judy libraries to search for the next available (empty) location in a Judy array. Functions include the following:

- FirstEmpty
- NextEmpty
- LastEmpty
- PrevEmpty

See the information library for details at: `http://www.hp.com/go/judy`.

updated for September 2001

As of September 2001, the Judy Libraries product was removed as a selectable application and is now installed with the OEs.

new for June 2001

The Judy product delivers:

- Very easy-to-use API

- Improved overall performance and better memory management

- Scalable arrays that can grow dynamically to very large populations while maintaining excellent performance

Judy offers three types of arrays:

Judy1 functions provide a way to store, retrieve, and locate Boolean values (bit maps) in a Judy array.

JudyL functions provide a way to store, retrieve, and locate long-word values in a Judy array.

JudySL functions provide a way to store, retrieve, and locate strings as indexes (similar to associative arrays in awk, Perl, and Java).

Compatibility

The following table shows the location of the libraries that are provided with the Judy technology on the HP-UX system:

Hardware Architecture	Type	Location on system (from root)	
		32-bit	**64-bit**
HP-PA 1.1 (32-bit only)	archive	/usr/lib/libJudy.a	N/A
	shared	/usr/lib/libJudy.sl	N/A

Hardware Architecture	Type	Location on system (from root)	
		32-bit	**64-bit**
HP-PA 2.0	archive	none	/usr/lib/pa20_64/libJudy.a
	shared	/usr/lib/pa20_32/libJudy.sl	/usr.lib/pa20_64/libJudy.sl

---

NOTE   The 32-bit HP-PA 1.1 shared library (/usr/lib/libJudy.sl) is provided for compatibility only. For best performance on 32-bit machines, use the HP-PA 2.0 shared library (/usr/lib/pa20_32/libJudy.sl).

Documentation

For more information about the Judy technology, refer to the *Judy* (3x) manpages (installed with the product) or visit the Judy web site at **http://www.hp.com/go/judy**.

Changes to libc

- **Large Files Support for C++ Applications**

  To support large files for C++ applications, libc has been modified. C++ applications can now access files greater than 2 GB. This is done by setting _FILE_OFFSET_BITS to 64 in 32-bit mode. More details can be found in the *HP-UX Large Files White Paper* Version 1.4 on **http://docs.hp.com**.

- **HP CxDL Development Tool Support**

  In both 64-bit and 32-bit libc, support for HP CxDL Development tool has been included in the setjmp() and longjmp() family of APIs.

- **libdbm and libndbm**

  A new patch for the dbm libraries, *libdbm* (1) and *libndbm* (2), has been created to increase performance of dbm_nextkey().

- **Header Files**

  Header files, ftw.h and stdio.h, were patched to enable C++ large files support. In addition, numerous defects were fixed.

- **New Environment Variables for malloc()**

  To make them thread-safe, libc uses a single lock in the malloc() routines. In a multi-threaded application, there could be contention on this single lock if multiple threads are calling malloc and free at the same time. This patch provides multiple arenas (from where malloc() can allocate space) and a lock for each arena. Threads are distributed among the arenas. Two new environment variables are introduced:

  ```
  _M_ARENA_OPTS
  _M_SBA_OPTS
  ```

  These can be used to tune the number of arenas and the arena expansion factor for

threaded applications. In general, the more threads in an application, the more arenas should be used for better performance. Expansion factors control the number of pages to expand each time and assumes the page size is 4096 bytes. The number of arenas can be from 4 to 64 for threaded applications.

For non-threaded applications, only one arena is used regardless of whether this environment variable is set or not. However, you still can use this environment variable to change the expansion factor for non-threaded applications.

If the environment variable is not set, or the number of arenas is set to be out of the range, the default number of 8 is used. The expansion factor is from 1 to 4096; the default value is 32. Again, if the factor is out of the range, the default value will be used. For example:

```
$ export _M_ARENA_OPTS=8:32
```

where the number of arenas is 8, and the expansion size is 32*4096 bytes. In general, the more arenas you use, the smaller the expansion factor should be, and vice versa.

`_M_SBA_OPTS` turns on the small block allocator, and sets up parameters for the small block allocator, namely, `maxfast`, `grain`, `num_smallblocks`. Refer to `mallopt()` for details about the small block allocator, and its parameters. Applications usually run faster with the small block allocator turned on rather than off.

Although a small block allocator can be turned on through `mallopt()`, it is not early enough for C++/Java applications. The environment variable turns it on before the application starts.

The `mallopt()` call can still be used the same way. If the environment variable is set, and no small block allocator has been used, the subsequent `mallopt()` calls can still overwrite whatever is set through `_M_SBA_OPTS`. If the environment variable is set, and a small block allocator has been used, then `mallopt()` will have no effect. For example:

```
$ export _M_SBA_OPTS=512:100:16
```

where the maxfast size is 512, the number of small blocks is 100, and the grain size is 16. You must supply all 3 values, and in that order. If not, the default values will be used instead.

The `_M_ARENA_OPTS` and `_M_SBA_OPTS` environment variables have the following impact:

— Performance is improved for multi-threaded applications.

— Threaded applications may experience increased heap storage usage (but you can adjust the heap usage through `_M_ARENA_OPTS`).

NOTE   Threaded applications that are linked with archive `libc` (or with other shared libraries that have dependencies on shared `libc`) may break.

Overall libc Performance Tuning

**new at 11i
original release**   To decrease calling overhead, several header files have been changed in the system library `libc`, (`/usr/lib/libc.sl`). In addition, a new archive library has been added to

allow linking the string and memory routines archived but an application as a whole can be linked shared.

There are now two different 32-bit system libraries. One is built for use on a PA1.1 machine and the other is built for use on a PA2.0 machine. The correct library is installed at installation time. Other changes to these libraries include a decreased calling overhead for the shared library. Also, the build process makes use of pragmas introduced in release 10.20 to decrease the calling overhead in shared libraries.

In addition to the changes to the library builds, changes have been made to selected header files to allow building applications that have decreased calling overhead. These changes apply to both 32-bit and 64-bit applications.

Two new libraries are added: `/usr/lib/libcres.a` and `/usr/lib/pa20_64/libcres.a`. These archive libraries include the common string and memory functions along with a improved performance `qsort` routine. A few other selected small routines are also included. The intent of this library is to allow an application to link to this library archived while linking the application as a whole shared. The use of this archived library is a supported link mode and will not introduce the problems normally associated with a shared/archive link.

The 32-bit system libraries now have selected APIs built with the pragmas `HP_DEFINED_EXTERNAL`, `HP_LONG_RETURN`, and `HP_NO_RELOCATION`. When these three pragmas are used in the building of `libc.sl`, it is referred to as a *fastcalled* library. The result of this is that the export stubs for the selected interfaces have been in-lined in the library code, thus reducing call overhead. Applications that have already been built will benefit from this without any effort other than the replacement of this library. The benefit a given application will gain is very dependent on the application's use of the `libc` APIs that have been fastcalled.

Along with the changes to the build process for `libc.2`, the following header files have been changed:

```
ctype.h
grp.h
mntent.h
pwd.h
stdio.h
stdlib.h
strings.h
string.h
time.h
```

These header files now contain the necessary fastcall pragmas to enable building a fastcalled application. To make use of the pragmas to build the application, `_HP_SHLIB_CALLS` needs to be defined for the application compile. With this define, the application will now have the import stubs inlined in the application code, further reducing the shared libary call overhead.

CAUTION   An application that has been built with the _HP_SHLIB_CALLS define can *only* be used with a fastcalled libc. If the application also has APIs that are fastcalled and are part of the applications shared libraries, then that library must also be built with the fastcall technology

---

Although the build process for `/usr/lib/pa20_64/libc.2` library has not changed, the runtime architecture for HPPA-2.0 can make use of a reduced call overhead technology similar to that which exists with the 32-bit library. There is no restriction on matching the correct `/usr/lib/pa20_64/libc.2` with the fastcalled application like there is with the 32-bit library.

Impact

There is little to no impact from these changes, although there is a slight (125KB) increase in amount of disk space required for `libcres.a`. The changes to the system libraries are transparent to current applications.

Any performance gains for an application are highly dependent on the application's use of `libc.sl` and what interfaces in this library are used.

The fastcall technology will be delivered with all systems. If there are compatibility concerns, the applications should not be built with this technology.

More API's in `libc` may make use of the fastcall technology in future releases. Appropriate changes to the header files will be delivered to track these changes.

Compatibility

An existing PA1.1 application will not have a compatibility issue with the new 32-bit fastcalled `/usr/lib/libc.sl`. However, if the fastcall technology is used to build an application, then that application can only be used with a fastcall technology library.

Neither an existing 64-bit application, nor a 64-bit application built with the fastcall technology, should have any compatibility issues with the existing `/usr/lib/pa20_64/libc.sl` libraries. However, to make use of the application fastcall and the `libcres.a` features, changes will need to be made to existing `make` files.

Documentation

The *libcres.a* (5) manpage describes use of the `libcres.a` library more thoroughly.

Performance Improvements to libc's ftw() and nftw()

The `libc` functions `ftw()` and `nftw()` have been rewritten to operate faster, avoid stack overflow conditions, reduce data space usage, and improve parallelism in multi-threaded applications.

Both `libc`, itself, and the commands that call `ftw()` and `nftw()` are affected.

- **ftw()**

  `ftw()` was rewritten to eliminate internal recursion, thus avoiding the possibility of a stack overflow on deep file trees. A single fixed-size data structure is allocated in the stack instead of using `malloc()` to separate buffers for each depth of the tree. Use of `strlen()` was eliminated, as well as trivial comparisons such as strcmp(buf,"."). The file descriptor re-use algorithm was changed from most-recently-opened to least-recently-opened which can show significant performance gains on very deep file trees.

  `ftw()` will typically show 8% reductions in elapsed time and 50% or more reduction in heap space used.

- **nftw()**

  nftw() was rewritten similarly to ftw() with the same benefits. nftw() now fully conforms with the UNIX95 definition, including the fact that when the FTW_PHYS is not set, files are reported only once.

  Threaded applications can obtain greater concurrency when specifying absolute path names for the starting path, and FTW_CHDIR is not set. In addition, an internal unbalanced binary tree was replaced with a much more efficient splay tree. The effect of this tree change becomes significant as the number of object inodes being tracked increases. Directory inodes are always tracked, and when executing in UNIX95 mode and the FTW_PHYS option is not set, all files and directories are tracked. When the number of tracked objects reaches about 20,000, the user CPU time with the splay tree is about half the user CPU time for the old nftw(). At 100,000 tracked inodes, the user CPU time is about 90% less for the splay tree.

  Another performance improvement to nftw() eliminated calls to access() by checking the mode bits in the stat() buffer. This decreased system CPU time by approximately 4%.

  Two defects were fixed in nftw():

  — When the FTW_CHDIR option is set, directories are considered unreadable unless they have both read and execute permissions. (The old nftw() would try to chdir() into a directory without execute permissions and then abort the walk with an error).
  — When the FTW_CHDIR option is set, a directory object is reported to the user function *before* it is chdir()'ed into.

  nftw() improvements vary depending on options provided, with the most significant improvements seen in UNIX95 standard mode with the FTW_PHYS option not set, or when a very large number of directories exist in the file tree being traversed.

Impact

The code size of ftw() and nftw() has increased by about 40%, but the heap requirements are reduced by 50% or more.

Minimally, you should find that ftw() operates about 6% faster and nftw() 4% faster. On very large file trees where the number of tracked inodes is in the tens of thousands or more, the performance gain of nftw() could be 30% to 40% or more.

If you relied on the FTW_CHDIR defects which were mentioned above, there may need to be an application change.

Documentation

The *ftw* (3C) and *nftw* (3C) manpages have been updated, particularly with respect to the two defect fixes and means of achieving best concurrency in threaded applications.

Performance Improvements to libc's malloc()

A new environment variable, _M_CACHE_OPTS, is available to help tune malloc() performance in kernel-threaded applications. This environment variable configures a thread-private cache for malloc'ed blocks. If cache is configured, malloc'ed blocks are placed into a thread's private cache when free() is called, and may thereafter be allocated from cache when malloc() is called. Having such a cache potentially improves

speed performance for some kernel-threaded applications, by reducing mutex contention among threads and by deferring coalescence of blocks.

The thread-private cache is only available for kernel-threaded applications, i.e. those linked with the pthread library. The installed shared pthread library version must be PHCO_19666 or later, or the application must be statically linked with an archive pthread library that is version PHCO_19666 or later, or else cache is not available.

By default cache is not active and must be activated by setting `_M_CACHE_OPTS` to a legal value. If `_M_CACHE_OPTS` is set to any out of range values, it is ignored and cache remains disabled.

There are two portions to the thread private cache: one for ordinary blocks and one for small blocks. Small blocks are blocks that are allocated by the small block allocator (`SBA`), which is configured with the environment variable `_M_SBA_OPTS` or by calls to mallopt(3C). The small block cache is automatically active whenever both the ordinary block cache and the `SBA` are active. The ordinary block cache is active only when it is configured by setting `_M_CACHE_OPTS`. There are no mallopt() options to configure the thread-private cache.

The following shows `_M_CACHE_OPTS`'s subparameters and their meaning:

`_M_CACHE_OPTS=<bucket_size>:<buckets>:<retirement_age>`

`<bucket_size>` is (roughly) the number of cached ordinary blocks per bucket that will be held in the ordinary block cache. The allowable values range from 0 through 8*4096 = 32768. If *`<bucket_size>`* is set to 0, cache is disabled.

`<buckets>` is the number of power of 2 buckets that will be maintained per thread. The allowable values range from 8 though 32. This value controls the size of the largest ordinary block that can be cached. For example, if *`<buckets>`* is 8, the largest ordinary block that can be cached will be 2^8 or 256 bytes. If *`<buckets>`* is 16, the largest ordinary block that can be cached will be 2^{20} or 65536 bytes, etc.

`<bucket_size>` * *`<buckets>`* is (exactly) the maximum number of ordinary blocks that will be cached per thread. There is no maximum number of small blocks that will be cached per thread if the small block cache is active.

`<retirement_age>` controls what happens to unused caches. It may happen that an application has more threads initially than it does later on. In that case, there will be unused caches, because caches are not automatically freed on thread exit -- by default they kept and assigned to newly-created threads. But for some applications, this could result in some caches being kept indefinitely and never reused. *`<retirement_age>`* sets the maximum amount of time in minutes that a cache may be unused by any thread before it is considered due for retirement. As threads are created and exit, caches due for retirement are freed back to their arena. The allowable values of *`<retirement_age>`* range from 0 to 1440 minutes (=24*60, i.e. one day). If *`<retirment_age>`* is 0, retirement is disabled and unused caches will be kept indefinitely. It is recommended that *`<retirement_age>`* be configured to 0 unless space efficiency is important and it is known that an application will stabilize to a smaller number of threads than its initial number.

In general, kernel threaded applications that benefit in performance from activating the small block allocator may also benefit further by activating a modest-sized ordinary cache, which also activates caching small blocks (from which most of the benefit is derived). For example, a setting that might be tried to begin with would be:

```
_M_SBA_OPTS=256:100:8
_M_CACHE_OPTS=100:20:0
```

The smallest ordinary cache that is legal and will activate small block caching (if the SBA is also configured) is

```
_M_CACHE_OPTS=1:8:0
```

It can happen that activating small block caching with this minimum level of ordinary cache gives all the performance benefit that can be gained from malloc cache, and increasing the ordinary block cache size further does not improve matters. Or, increasing cache size further may give some further improvement for a particular application.

The `malloc()` per-thread cache is a heuristic which may or may not benefit a given kernel-threaded application that makes intensive use of malloc. Only by trying different configurations can you determine whether any speed improvement can be obtained from per-thread cache for a given application, and what the optimal tuning is for that application.

Impact

Although there will be no impact on performance if cache is not configured (or if application is not kernel-threaded), significant speed performance improvements are possible for some kernel applications if cache is configured.

A small additional space cost (in process heap size) is associated with the cache machinery. While there is no per-block space cost for caching small blocks, there is a small space cost per ordinary block cached. Therefore, ISVs whose applications are very memory intensive may want to configure only a minimum-sized or very small ordinary cache when experimenting with this feature.

The `malloc()` thread-private cache does not change the function of `malloc()` for nonthreaded or CMA-threaded applications. It does maintain binary compatibility. However, because it is a change in allocation policy, it can cause different sequences of addresses to be emitted for the same sequence of requests than a previous version of malloc would have emitted. This level of compatibility is more stringent than ordinary binary compatibility and has never been guaranteed across releases of malloc.

The libcres.a Library

Provided with the original release of 11i, `libcres.a` is a small archive library that contains string, memory and other functions, to provide customers running performance-critical applications with the benefit of a static link.

Linking statically with `libc` is not a supported method of linking an application. Any performance improvement is highly dependent on the application's use of the included functions. The functions included in this library are as follows:

```
abs(), bsearch(), div(), ffs(), insque(), labs(), ldiv(), memchr(),
memcmp(), memcpy(), memmove(), memset(), strcat(), strchr(), strcmp(),
strcpy(), strcspn(), strlen(), strncat(), strncmp(), strcpy(), strrchr(),
strspn(), strstr(), swab()
```

To make use of this library, existing makefiles must be modified to include it on the link line. Existing applications must be re-linked to use this library.

The modules of this library are compiled with the HP optimizing compiler using a +O4 flag. As a result, the applications using this library can be linked only by using the HP optimizing compiler.

The functions in this library cannot be overwritten with a user-defined function of the

same name, as is the case today with `libc` names. If this library is used, user libraries cannot contain identically named functions or unexpected results may occur.

Impact

Performance of some applications may improve by using this library. The improvement is highly dependent on the application's use of the included functions.

Documentation

The *libcres.a* (5) manpage describes the `libcres.a` library's use more thoroughly.

Changes to libm

The `fesetround()` and `fehold()` functions in `fenv.h` have been upgraded to the latest ISO C9x specification. Previously the functions returned nonzero to indicate success and zero to indicate failure; now they return zero to indicate success and nonzero to indicate failure.

Any code that depended on the return value will need to change. For example:

```
if (!fesetround(FE_UPWARD))

{/* deal with failure to set rounding direction */}
```

could be changed to:

```
if(fesetrod(FE_UPWARD))

{/* deal with failure to set rounding direction */}
```

Previous code that depended on the return value is not compatible beginning with the 11.0 May 1999 Extension Pack.

Miscellaneous

Linker and Object File Tools (ld, crt0.o, dld.sl, libdld.sl, chatr, and odump)

new at 11i original release

The following list summarizes the changes to linker and object file tools. Details for several of the linker changes follow the lists.

Linker Changes

- Incremental linking supported in 64-bit `ld` and `elfdump`.
- Unix 98 (32-bit `dl()`* calls) supported in `libdld.sl` and `dld.sl`.
- 32-bit Filtered shared libraries supported in `ld`, `dld.sl` and in `odump`.
- GProf 32-bit shared library supported in `crt0.o` and `dld.sl`.
- `ld +filter` option creates filtered shared libraries.
- Support for `ldd32`, which lists dynamic dependencies of executable files or shared libraries in `dld.sl`.
- Plabel cache, caches `PLABELS` at run-time, supported in `ld` and `dld.sl`.
- `ld +dependdb` and `+dependdb_outputdir` options generate dependency database, `ldb` file.
- `ld +objdebugonly` (in both 32-bit and 64-bit) ignores debug information from non-`objdebug` objects or archives and then proceeds in `+objdebug` mode.
- Special support for OGL's `TLS` shared library in `dld` (both 32- and 64-bit).

Tools Enhancements

- `elfdump +ild` displays incremental linking information.
- `ar -x` option allows modules from `lib` to keep datestamp.
- `odump -tlssym` option displays the `TLS` (thread) symbols.
- `chatr +q3p enable/disable` and `q4p enable/disable` option supports marking 3rd/4th quadrant for private data space.
- `odump -verifyall` option suppresses stub warnings on executable.
- `odump -filtertable` to displays the filtered shared library's implementation libraries.

Details of Linker Changes

Incremental linking: Incremental linking provides significant linktime improvements for compile-link-debug development cycles by processing only those input files that are actually modified between cycles. Files that are not modified do not need to be reprocessed. For large application, incremental linking may provide up to 10x (and sometimes greater) improvements in linktime.

Unix 98: Support for the APIs `dlopen`, `dlsym`, `dlerror` and `dlclose` has been added for

32-bit programs.

Filtered Libraries: Filtered shared libraries divide up a large library into one filter and several implementation libraries. The user links against the filter library, but the real definitions of data and functions actually resides in the implementation libraries. At run time, only those implementation libraries that are actually used are loaded. Filtered libraries can be nested; an implementation library can itself be a filtered library containing other implementation libraries.

`GProf` **32-bit support:** GProf, an enhanced version of `prof`, produces call graph profile data (along with the data generated by `prof`). While the profiling of shared libraries was not supported in earlier releases, this release will now support profiling of shared libraries using the environmental variable `LD_PROFILE`. No recompilation is required for profiling shared libraries.

PLabel cache: `+plabel_cache` is added to 32-bit linker and `dld.sl` to control the global symbol hash mechanism.

Compatibility

Various serious and critical defects were repaired.

Forward and backward compatibility are maintained. However, use of new features in this release may break backward compatibility.

Invoking `chatr` on some binaries built with an older linker may emit the following message:

```
chatr(error): dl_header_ext.size != sizeof(dl_header_ext). Please update
your version of the linker/chatr.
```

This message should be regarded as a warning rather than an error. The operation of `chatr` will be successful despite the warning.

Changes to the linker/dld Interface

updated for March 2002

The `linker/dld` interface support for CXperf no longer applies, as the CXperf product has been discontinued. For further details, see **http://hp.com/go/cxperf**.

new at 11i original release

Patches to the `linker/dld` interface include the following enhancements:

- Added support for the CXperf performance measuring tool in both 32-bit and 64-bit versions of the `ld` command. Both versions now recognize the `+tools` option, which enables CXperf information to be propagated to an executable program.

- Added support for huge data (.bss > 4GB)

- A defect was repaired whereby `+Oprocelim` removed more than it should have causing a runtime error.

- Performance shows a definite improvement:

  — 32-bit ld: approximately 30% link time improvement

  — 64-bit ld: approximately 8% link time improvement

- Support OBJDEBUG architecture in both 32-bit and 64-bit linker.

- Added support for executable stack.

- Added global symbol table support.

- Added support for object code repository reuse.

Neither functionality nor compatibility are affected by the code changes. However, for 64-bit programs, mixing object files having non-weakorder sections with object files having weakorder sections might cause the ordering of text sections to change.

The new version of the linker requires 34112 blocks.

Instrumented Code Using PBO or +O4 Optimization

This note pertains to the compilers and linker for HP C, HP aC++, HP C++ (`cfront`), HP Fortran 77, and HP-UX Linker.

When you compile your source code with the compiler shipped on HP-UX 11i, without any changes to source code, options, or makefiles, you might create relocatable object files or executables that are no longer backward compatible to an original 11.0 system. This condition will occur if you recompile with PBO (`+I` compiler or linker option) or the `+O4` option. You might create instrumented objects (ISOM) that an 11.0 system does not recognize.

Under these circumstances, one of the following types of error messages will be issued if you attempt to link the objects created using the HP-UX 11i compiler on an original 11.0 system.

- If you compile with `+O3` or `+O4` you will receive the following message and a stack trace: `report error(13-12299-434) to your nearest HP service representative(8911).`

- If you compile with `+O2 +I` you will receive the following message and a stack trace: `Backend Assert ** Ucode versions earlier than v.4 no longer supported(5172).`

NOTE   This code is not backward-compatible with the 11.0 release. Instrumented object files cannot be moved backward.

HP DCE/9000

**new at 11i
original release**   HP Distributed Computing Environment (HP DCE/9000) Version 1.8 provides a high-quality, comprehensive, standards-based framework to develop, administer, and use distributed applications.

Kernel threads application development is now supported on HP-UX. The 32-bit version of the kernel threads DCE library (`libdcekt`) is now part of the HP-UX base operating system. The 64-bit version of `libdcekt` is also included.

NOTE   While the DCE library (`libdcekt`) has been ported to 64-bit, the binaries and daemons that are part of the DCE products are still 32-bit.

The distinction between the International and US/Canada version of DCE components has been removed. That is, the 56-bit Data Encryption Standard (DES) which was

earlier restricted to US/Canada is now available for all customers. This means there will only be one version of the DCE library and `dced` daemon, which is based on the 56-bit DES version.

NOTE   The number of LAN interfaces supported by DCED is limited to 32 and the LAN interfaces supported by CDS is 12. If there are more than 32 LAN interfaces, the environment variable `RPC_SUPPORTED_NETADDRS` can be used to identify the list of 32 LAN interfaces that are used by the Remote Procedure Call (RPC) application.

Listed below are new environment variables that have been added to support RPC operations and to enable better usability:

- `RPC_PREFERRED_PROTSEQ`: This variable is used to set the preferred protocol sequence.

- `RPC_SUPPORTED_PROTSEQS`: This variable helps in restricting the protocol sequence. For example, setting this variable to `ncacn_ip_tcp` will enable only connection-oriented communication.

- `RPC_DISABLE_PRIVATE`: This variable disables the "private" setting of a socket. Specifically, the datagram protocol opens up one socket for each network address family supported on a host. Once opened, these sockets are kept in a pool for use whenever the process needs to make another RPC over that particular address family. If concurrent calls are made over the same address family, the calls share a single socket from the pool. However, this is inefficient for those applications that don't require this degree of concurrency.

  To remedy this situation, 1 or 2 sockets are tagged as "private," along with the usual shared sockets in the socket pool. You can disable this setting by exporting `RPC_DISABLE_PRIVATE=1`. The default behavior is for private socket to be enabled.

- `RPC_DISABLE_LOCAL`: For a RPC server and client on same host, UNIX domain sockets are used by default to reduce the overhead. This can be disabled by exporting `RPC_DISABLE_LOCAL=1`.

- `HPDCE_CLIENT_DISC_TIME`: With this variable (provided in the DCE RPC runtime), the idle association termination time can be tuned to be a value less than the architecture-provided value of 5 minutes (namely, any value ranging from 1 to 300 seconds). Please note that this variable is applicable only for connection-oriented protocol.

- `SCTE_UNCACHE_TIME`: This variable is applicable for datagram only and is used to reduce the server connection table (SCT) elements to be uncached sooner than the default value. The default time is 300 seconds. This would allow more SCT entries to be added to the SCT without resulting in cache exhausting heap.

- `DMS_FORCEON`: DCE Measurement Service (DMS) provides performance instrumentation for DCE servers and for the server side of applications that use DCE RPCs. When DMS is enabled, it collects data about RPCs that execute in the target process. The collected data is actually displayed using HP GlancePlus. By default, DMS is disabled. DMS can be enabled exporting `DMS_FORCEON=1`.

Also, since CMA threads are being obsoleted, it is recommended that all applications using CMA threads start migrating to kernel threads and use `libdcekt`. (See "Kernel Threads vs. CMA Threads" on page 144 for additional information.)

Compatibility Issues

All applications using the 64-bit library `libdcekt` may need to include
`/usr/include/dce/dce64.h`.

DCE server products are not supported on workstations (Series 700 machines).

Documentation

The site `http://devresource.hp.com/STK` contains 64-bit porting concepts and 64-bit
compiler and linker changes needed to port the application to 64-bit.

Extensions to pstat()

This extension provides new functionality to the `pstat()` system call, enabling various
system management and measurement tools to eliminate their dependency on the
`/dev/kmem` pseudo-driver.

Today, many system management and measurement tools read kernel data structures
through unsupported interfaces, such as the `/dev/kmem` pseudo-driver, to get
information about open files, resource usage, process activity, and so on. Because kernel
data structures change from release to release, this access method is fragile, incurring a
high maintenance cost. To insulate these applications from the release-to-release
variability in private kernel data structures, HP-UX 11i provides the enhanced `pstat`
system call and a new set of wrappers.

The pstat interface is designed to allow future expansion of the interface, while
preserving source and binary compatibility of programs written using pstat wrappers.
The pstat interface is available in both 64-bit and 32-bit versions. Replacing the
`/dev/kmem` access with calls to `pstat` wrappers will eliminate the need to re-release
applications with each new HP-UX release.

Currently, the `pstat()` system call provides information about various system contexts,
such as static, dynamic, virtual memory, process, open files, etc. HP-UX provides a
number of `libc` wrappers (`pstat_get()`*) and corresponding structures (`struct
pst_()`*) to get information from the kernel using `pstat()`. As part of this
enhancement, new `pstat()` wrappers and corresponding structures are added and
some existing ones are extended.

Compatibility is significantly improved by introducing a well documented interface that
guarantees binary compatibility for kernel intrusive applications between releases.
There is no impact to legacy behavior of current `pstat()` services.

There is no impact to application performance as compared to obtaining the data from
`/dev/kmem`. No impact to system performance is expected from these `pstat` extensions.

NOTE   This release includes an enhanced version of `pstat()`. This version repairs some
existing defects by adding more fields in `pst_status` struct to return process children
usage information. The *pstat* (2) manpage reflects this added functionality. The
enhancement poses no problem for 11.0 executables running on 11.0 Extension Pack or
11i, nor for any executables running on 11.0 Extension Pack, as long as they do not rely
on the additional functionality.

Note, however, relocatable objects may incorrectly presume that the size of returned
information is the same pre- and post-patch. It is possible to determine the size of
information returned. `pstat()` users can use the size return value of the system call to

---

maintain relocatable object compatibility and portability across the proposed change. This is documented in the manpage.

`pstat()` is not part of an industry standard, but was designed to accommodate changes of this nature while maintaining compatibility with earlier versions.

New Modules

The following table shows new pstat modules and the purpose of each:

`pstat_getfile2()`	Provides information about open files of a process
`pstat_getfiledetails()`	Provides `stat` equivalent information
`pstat_getsocket()`	Provides detailed socket information
`pstat_getstream()`	Provides detailed stream information
`pstat_getpathname()`	Provides full pathname of an opened file (Reverse Pathname Lookup)
`pstat_getmpathname()`	Provides a copy of the DNLC entries for a given file system

NOTE   Use of the call `pstat_getmpathname()` is limited to `uid` equal to 0. Use of the calls `pstat_getfiledetails()`, `pstat_getsocket()`, `pstat_getstream()`, and `pstat_getpathanme()` is limited to uid equal to 0 or effective ID match. In the case of effective ID match, access will only be granted if the target process is not and has never run as a `setuid` or `setgid` process.

New Data Structures

The following are new data structures being added to the PSTAT module:

`pst_fileinfo2`	Describes per-file information. For the specified process, there is one instance of this context for each open file descriptor.
`pst_fid`	Used to efficiently re-access the opened files. This value is returned by `pstat_getfile2()`, `pstat_getproc()`, and `pstat_getprocvm()` calls. This ID is then passed to subsequent PSTAT calls such as `pstat_getsocket()` to efficiently re-access the opened files.
`pst_filedetails`	This data structure contains detailed information specific to a particular open file. For a specified file, there is only one instance of this structure. This information includes `stat` equivalent information.
`pst_socket`	The PSTAT socket structure contains detailed information pertaining to an opened socket, such as type, state, protocol, address family, and options of the socket. For a specified socket, there is only one instance of this structure.

pst_stream	The PSTAT stream structure contains detailed information pertaining to a stream entity. This includes information about the head, names of modules pushed, and the driver of the stream.
pst_mpathnode	This structure is returned by pstat_getmpathname() routine that provides a copy of the DNLC entries for a given file system. The information contained in this structure includes id of the current file or directory, parent of the current entry, and the name of the current entry. By traversing the DNLC entries in the reverse order, one can obtain the pathname for an opened file to the mount point.

In addition to the above data structures, several existing PSTAT data structures have been extended. These include: pst_dynamic, pst_vminfo, pst_vm_status, pst_status, pst_static, and pstun.

Documentation

The existing *pstat* (2) manpage has been extended to reflect the added functionality.

Changes to sendfile

The sendfile() system call is used to send a file directly over the network without having to perform many separate send() commands.

In previous releases, sendfile() did not work properly with large files, that is, when an application made a call to sendfile() and was compiled with the following compiler flags: LARGEFILE(64)_SOURCE and/or FILE_OFFSET_BITS=64. These flags allowed a 32-bit application to access large files that were over 2GB in size.

These large file applications should be recompiled on 11i. If the "nbytes" parameter is not set to zero and they are not recompiled, these applications will not execute on 11i. To work correctly, the large file applications need to be recoded with the new bsize_t and sbsize_t types. See the *sendfile* (2) and *sendfile64* (2) manpages for more information.

32-bit or 64-bit applications that use sendfile() and are *not* compiled with the LARGEFILE(64)_SOURCE or FILE_OFFSET_BITS=64 flags do *not* need to be changed or recompiled for HP-UX 11i.

Machine Identifier Changes to confstr

new at 11i original release

New machine identifier, partition identifier, and serial number parameters have been defined for the confstr() library function.

The new parameters for confstr() are defined as follows:

_CS_MACHINE_IDENT

Identifier for each physical machine. Returned as an opaque string of printable ascii characters.This string has the same value for all partitions in a physical machine. For hardware classes first released with HP-UX 11i or later, this ID is unique across all hardware classes. For earlier hardware classes, the ID number is unique only within the hardware class. A null string is returned if no ID number is available;

this is expected to be the case only for prototype machines or other systems improperly configured in manufacturing.

_CS_PARTITION_IDENT

Identifier for each partition existing on a machine. Returned as an opaque string of printable ascii characters. For any machine not supporting partitions this value will be same as _CS_MACHINE_IDENT.

_CS_MACHINE_SERIAL

Machine serial number as found labeled on the external machine chassis. The value will be a printable ascii string. This string is not available on all classes of machines; if unavailable, the string will be empty. This string is not a unique identifier of the machine, since machines of different classes can have the same serial number.

If a unique identifier is needed, use _CS_MACHINE_IDENT or _CS_PARTITION_IDENT.

The preferred method of calling these functions is defined in the *confstr* (3C) manpage as the following:

```
bufsize=confstr(_CS_MACHINE_IDENT,NULL,(size_t)0);
buffer=(char *)malloc(bufsize+1);
confstr(_CS_MACHINE_IDENT,buffer,bufsize+1);
```

The first line will return the length of the string to be returned, allocate memory based on this value, and then call confstr() again to get the actual value.

14 Licensing Products

Future Change for LicensePower/iFOR

HP plans to remove LicensePower/iFOR from the Core HP-UX software in a future release. This licensing product can be obtained directly from Isogon Corporation, the owner of the product.

To download LicensePower/iFOR, go to Isogon's Web site:

http://www.isogon.com/support/sptlpifor/download/download.htm

Impending LSSERV Software Obsolescence

**new at 11i
original release**

HP-UX 11i is the last release that will contain the LSSERV licensing product as a bundled part of the operating system.

You can obtain this product directly from its owner, the Isogon Corporation. You can also visit the Isogon Corporation web site for further information about LSSERV support at `http://www.isogon.com`.

15 New and Changed Internationalization Features

Unicode Character Set

HP-UX 11i provides system level support for the Unicode 2.1/ISO-10646 character set. Hewlett-Packard's support for Unicode provides a basis of enabling heterogeneous interoperability for all locales.

ISO-10646 is an industry standard for defining a single encoding which uniquely encodes all the world's characters. Unicode 2.1 is the companion specification to ISO-10646, Unicode support conforms with existing X/Open (OpenGroup), POSIX, ISO C and other relevant UNIX-based standards.

HP-UX 11i supports Unicode/ISO-10646 by utilizing the UTF-8 (Universal Transformation Format-8) representation for persistent storage. UTF-8 is an industry recognized 8-bit multibyte format representation for Unicode. This representation allows for successful data transmission over 8-bit networking protocols as well as for safe storage and retrieval within a historically byte-oriented operating system such as HP-UX.

For internal processing, HP-UX utilizes the four-octet (32-bit) canonical form specified in ISO-10646. This support allows parity with HP-UX's current `wchar_t` implementation which has been based on a 32-bit representation.

Full systems level support is provided for all locales provided in this release.

For more information on the Unicode features of Asian System Environment, see `/usr/share/doc/ASX-UTF8`.

A select subset of locale binaries have been provided for 32-bit application processing:

Table 15-1   **Base**

C.utf8	C UTF-8
univ.utf8	universal

Table 15-2   **European**

fr_CA.utf8	French Canadian
fr_FR.utf8	French
de_DE.utf8	German
it_IT.utf8	Italian
es_ES.utf8	Spanish
sv_SE.utf	Swedish

Table 15-3   **Asian**

ja_JP.utf8	Japanese
ko_KR.utf8	Korean

Table 15-3   **Asian**

zh_CN.utf8	Simplified Chinese
zh_HK.utf8	Traditional Chinese (Hong Kong)
zh_TW.utf8	Traditional Chinese

To enable Unicode support in applications, set the environment variable to a desired `utf8` locale.

Locales are installed based on the current language file sets already installed on the target system. For example, if the system uses the International.German the German Unicode locale (`de_DE.utf8`) is installed.

Source files for ALL supported locales (34 total) have also been supplied for 64- or 32-bit applications.

To build Unicode locales use the `localedef` command. Refer to the *localedef* (1M) manpage. Systems must have the kernel parameters MAXDSIZ, MAXTSIZ, and SHMMAX set to at least 100 MB to ensure adequate swap space allowance for successful `localedef` compilation of these locales.

Unicode Euro Enhancement

HP-UX 11i provides expanded Unicode support to align the character repertoire with the ISO 8859-15 locales that are being provided for Euro support. This will ensure full interoperability with the newly added support for the ISO 8859-15 codeset.

Specific enhancements are provided to allow Euro display and input capabilities though Xlib and new fonts.

A subset of existing European (and French Canadian) locales have been modified:

Table 15-4   **Modified European Locales**

Locale	Country
fr_CA.utf8	French Canadian
fr__FR.utf8	French
de_DE.utf8	German
it_ IT.utf8	Italian
es_ ES.utf8	Spanish
sv_SE.utf8	Swedish

Source files for all supported European locales have also been modified. To build these locales, refer to the *localedef* (1M) manpage.

Size Requirement

Unicode support requires the following additional disk space requirements:

Base Unicode offering (installed on all systems): Approximately 10MB.

Table 15-5   **Unicode European locales and localized files**

French & French Canadian	8.4 MB
German	4.2 MB
Italian	4.2 MB
Spanish	4.2 MB
Swedish	4.2 MB

Table 15-6   **Unicode Asian locales and localized files**

Japanese	3.4 MB
Korean	2.4 MB
Simplified Chinese	2.5 MB
Hong Kong	1.7 MB
Traditional Chinese	4.2 MB

Performance Issues

Applications using Unicode support should see comparable performance as observed with other multibyte codesets. For those applications moving from a single-byte codeset to Unicode, some performance impact will be observed for some types of character-based operations.

Streams PTY Driver

UTF-8 is supported on the Streams PTY driver's line discipline (LDTERM) module. The user does not interact with the Streams PTY driver directly; it runs underneath the `dtterm` window. The Streams PTY driver is responsible for providing a UTF-8 communication channel while `dtterm` is responsible for processing the UTF-8 code and displaying the characters on the screen.

Refer to *eucset* (1), *ldterm* (7) and the *lp* (1) model script for details.

Corrected Character Mappings to iconv(1) and iconv(3C)

This release contains defect fixes for incorrect character mappings. The corrections concern the Simplified Chinese, Traditional Chinese, Japanese, and Korean characters of HP-UX.

Corrected character converter mappings allow for improved interoperability when sending or receiving converted character data to/from Unicode-aware systems.

Correction for Simplified Chinese

A patch corrects an incorrect character mapping that occurs when converting between hp15CN and Unicode (UCS2)/UTF-8 for Simplified Chinese.

Specifically, the Simplified Chinese character "Double Vertical Line" mapped incorrectly when converting between hp15CN and UCS2/UTF-8. This character was being mapped to the "Parallel To" character, which is a different character.

The following table summarizes the change applied to `iconv` tables:

hp15CN	incorrect UCS2	correct UCS2	Character Name
0xA1CE	-	0x2225	Parallel To
0xA1AC	0x2225	0x2016	Double Vertical Line

The hp15CN=ucs2 and ucs2=hp15CN `iconv` converter tables are affected. These tables are shared by both UCS2 and UTF-8 conversions.

No compatibility problems are anticipated. However, if compatibility concerns arise with regard to persistent data stored either in Unicode (UCS2) or UTF-8 on an HP-UX system, it is possible to generate a simple conversion script to search for each occurrence of an incorrect value in either UCS2 or UTF-8 and convert it to the correct value, based on the following mapping:

Old UCS2	UCS2	Old UTF-8	UTF-8	Char Name
0x2225	0x2016	0xe288a5	0xe28096	Double Vertical Line

Correction for Traditional Chinese

A patch corrects several incorrect character mappings that occur when converting between Big-5/EUC and Unicode (UCS2)/UTF-8 for Traditional Chinese.

In the case of Big-5 to/from UCS2/UTF-8, the "Ideographic Space" character was absent in the Unicode conversion table mapping:

big5	incorrect UCS2	correct UCS2	Char Name
0xA140	-	0x3000	Ideographic Space

The following table summarizes the changes applied for conversions between eucTW and UCS2:

eucTW	incorrect UCS2	correct UCS2	Character Name
0xa1a6	0x30fb	0x2022	Bullet
0xa1b7	0x2014	0x2013	EN Dash
0xa1b9	0x2013	0x2014	EM Dash
0xa1b6	0xfe31	0xff5c	Full-width Vertical Line
0xa1b8	0xfe32	0xfe31	Presentation form Vertical EN Dash
0xa1ea	0x2032	0x2035	Reversed Prime
0xa1eb	0x2035	0x2032	Prime
0xa2b9	0x2264	0x2266	Less-than over equal to
0xa2ba	0x2265	0x2267	Greater-than over equal to
0xa2c2	0xfe66	0xfe65	Small Greater-Than
0xa2c3	0xfe65	0xfe66	Small Equals Sign
0xa2de	0xff5c	0x2223	Divides
0xa2e1	0xfe67	0xff0f	Full-width Solidus
0xa2e4	0xffe5	0x00a5	Yen Sign
0xa2e6	0xffe0	0x00a2	Cent Sign
0xa2e7	0xffe1	0x00a3	Pound Sign

`iconv` conversions between eucTW and UCS2 or UTF-8 may be affected.

Big-5 conversions with UCS2/UTF-8 are not directly impacted as only a missing table entry has been added.

eucTW=ucs2, ucs2=eucTW, big5=ucs2 and ucs2=big5 `iconv` converter tables are affected. These tables are shared by both UCS2 and UTF-8 conversions.

No compatibility problems are anticipated. However, if compatibility concerns arise with regard to persistent data stored either in Unicode (UCS2) or UTF-8 on an HP-UX system, it is possible to generate a simple conversion script to search for each occurrence

of an incorrect value in either UCS2 or UTF-8 and convert it to the correct value, based on the following mappings:

Old UCS2	UCS2	Old UTF-8	UTF-8	Char Name
0x30fb	0x2022	0xe383bb	0xe280a2	Bullet
0x2014	0x2013	0xe28094	0xe28093	EN Dash
0x2013	0x2014	0xe28093	0xe28094	EM Dash
0xfe31	0xff5c	0xefb8b1	0xefbd9c	Fullwidth Vertical Line
0xfe32	0xfe31	0xefb8b2	0xefb8b1	Presentation form Vertical EN Dash
0x2032	0x2035	0xe280b2	0xe280b5	Reversed Prime
0x2035	0x2032	0xe280b5	0xe280b2	Prime
0x2264	0x2266	0xe289a4	0xe289a6	Less-than over equal to
0x2265	0x2267	0xe289a5	0xe289a7	Greater-than over equal to
0xfe66	0xfe65	0xefb9a6	0xefb9a5	Small Greater-Than
0xfe65	0xfe66	0xefb9a5	0xefb9a6	Small Equals Sign
0xff5c	0x2223	0xefbd9c	0xe288a3	Divides
0xfe67	0xff0f	0xefb9a7	0xefbc8f	Full-width Solidus
0xffe5	0x00a5	0xefbfa5	0xc2a5	Yen Sign
0xffe0	0x00a2	0xefbfa0	0xc2a2	Cent Sign
0xffe1	0x00a3	0xefbfa1	0xc2a3	Pound Sign

Correction for Japanese

A patch corrects four incorrect Japanese character mappings that occur between Shift-JIS/EUC and Unicode (UCS2)/UTF-8.

The following table summarizes the changes applied:

sjis	eucJP	incorrect UCS2	correct UCS2	Character Name
0x8150	0xA1B1	0xFFE3	0x203E	Overline
0x815C	0xA1BD	0x2015	0x2014	Em Dash
0x818F	0xA1EF	0xFFE5	0x00A5	Yen Sign
n/a	0x8FA2B7	0x02DC	0xFF5E	Full-width Tilde

Affected `iconv` conversions are conversions between sjis and UCS2 or UTF-8 as well as conversions between eucJP and UCS2 or UTF-8.

sjis=ucs2, ucs2=sjis, eucJP=ucs2 and ucs2=eucJP are the affected `iconv` conversion

tables. These tables are shared by both UCS2 and UTF-8 conversions.

No compatibility problems are anticipated. However, if compatibility concerns arise with regard to persistent data stored either in Unicode (UCS2) or UTF-8 on an HP-UX system, it is possible to generate a simple conversion script to search for each occurrence of an incorrect value in either UCS2 or UTF-8 and convert it to the correct value, based on the following mappings:

Old UCS2	UCS2	Old UTF-8	UTF-8	Char Name
0xFFE3	0x203E	0xefbfa3	0xe280be	Overline
0x2015	0x2014	0xe28095	0xe28094	Em Dash
0xFFE5	0x00A5	0xefbfa5	0xc2a5	Yen Sign
0x02DC	0xFF5E	0xcb9c	0xefbd9e	Full-width Tilde

Correction for Korean

A patch provides a defect fix to address standards non-conformance for Korean Unicode (UCS2)/UTF-8 character mappings.

The currently supplied Korean `iconv` converter tables do not conform to the Unicode 2.1 and ISO-10646 (with 1997 amendments) standards in addition to the Korean national standard, KSC-5700. The current mappings are considered obsolete by all noted standards organizations.

The enhancement provides a set of standards-conformant `iconv` converter tables for converting between eucKR and Unicode/UTF-8. Specifically, the obsolete region of 0x3d2e - 0x4dff has been re-mapped to the 0xac00 - 0xd7ff region specified in Unicode 2.1 for Hangul.

Without this modification, it is impossible to share data with any other system which is standards-conformant in adhering to the Unicode 2.1/ISO-10646/KSC-5700 standards.

Affected `iconv` conversions are any conversions between eucKR and UCS2 or UTF-8.

The `iconv` conversion tables affected by this modification are eucKR=ucs2 and ucs2=eucKR. These tables are shared by both UCS2 and UTF-8 conversions.

No compatibility problems are anticipated. However, if compatibility concerns arise with regard to persistent data stored either in Unicode (UCS2) or UTF-8 on an HP-UX system, it is recommended that the previously installed ucs2=eucKR table be saved and renamed prior to installation of this fix. Persistent data can then be converted back to eucKR using this old table and then reconverted to the correct Unicode/UTF-8 representation.

EURO (ISO 8859-15 Locales)

Euro support is provided via locale support for the ISO 8859-15 character set. ISO 8859-15 is a newly ratified character set that differs from ISO 8859-1 in that it supports eight new characters. Specific enhancements are provided to allow Euro display, input, and processing capabilities.

Fourteen new locales have been created based on ISO 8859-15:

Locale	Language (Country)
C.iso885915	"C"
da_DK.iso885915@euro	Danish (Denmark)
de_DE.iso885915@euro	German (Germany)
en_GB.iso885915@euro	English (Great Britain)
es_ES.iso885915@euro	Spanish (Spain)
fi_FI.iso885915@euro	Finnish (Finland)
fr_CA.iso885915@euro	French (Canada)
fr_FR.iso885915@euro	French (France)
fr_IS.iso885915@euro	Icelandic (Iceland)
it_IT.iso885915@euro	Italian (Italy)
nl_NL.iso885915@euro	Dutch (The Netherlands)
no_NO.iso885915@euro	Norwegian (Norway)
pt_PT.iso885915@euro	Portuguese (Portugal)
sv_SE.iso885915@euro	Swedish (Sweden)

Source files for supported European locales are also being supplied.

Applications must elect to enable ISO 8859-15 support, by setting the LANG environment variable to the desired locale.

ISO 8859-15 support is part of HP-UX and is available to all platforms. ISO 8859-15 support is not automatically turned on for any application. No special configuration is required and there are no compatibility issues involved with the addition of this new feature.

Locales are installed, based on which current language file sets are already installed on a target system.

The `LC_MONETARY` environment variable will be set to the euro for all locales listed above except C.iso885915 and fr_CA.iso885915. Standard euro formatting rules will apply to ALL locales where this environment variable is set to the euro. As a result, users may encounter a change to the decimal and thousands separators for the currency, whereas decimal and thousands separators outside the monetary area stay the same as in previous locales.

For example, in the French locale, the thousands separator is a space and the decimal point is a comma. However, the international standard for the thousands separator for the euro currency is a period. So, a user that has the `LC_MONETARY` locale category set to "fr_FR.iso885915@euro" will see the following behavior:

- The number one thousand five hundred and fifty and a half, outside the monetary area will be displayed as 1 550,50

- One thousand five hundred and fifty euro and 50 cents will be displayed as EUR 1.550,50.

The `LC_MONETARY` value can be changed by users to their national currency unit.

ISO 8859-15 support is not automatically provided in any application. Applications which use the Euro symbol must elect to enable ISO 8859-15 support, by way of setting the `LANG` environment variable to the desired locale. Users enable ISO 8859-15 automatically in some locales when logging in through the CDE.

For more information, please see:

`http://software.hp.com/products/EURO/index.html`

CDE Support

New functionality was introduced in the CDE product to support input and display of the Euro symbol. (These changes are for both the workstation and the server.)

X Window Support

New functionality was added to Xlib to support input and display of the Euro symbol. This was done by adding internal support for the ISO8859-15 character set (as well as support of UTF8 on 11i). When an Xlib application is started, Xlib internals determine if the locale is set to an ISO8859-15 character set. If it is, Xlib will perform character lookups using the eight new symbols present in the ISO8859-15 character set. Currently, only applications linked with X11R6 (X Window version X11 Release 6) will support the ISO8859-15 character set. Older X11 versions are not currently supported.

Libraries

The `libc` and `xlib` libraries support the Euro symbol.

Codeset Converters

New `iconv` tables exist to support conversion from/to ISO 8859-15 and ISO 8859-1, ucs2, and utf8. The additional disk space in HP-UX 11i is 6.42MB. No additional memory is required.

LaserJet Printers

An important aspect of the euro support is printing the new symbol on LaserJet printers using existing standard lp(1) model files.

The ISO8859-15 font set is resident on the HP 4500 Color LaserJet Printer, which contains the Euro symbol at position A4 (hexadecimal). Your data file must contain this code to print the Euro symbol.

A new utility will be provided to download the fonts to the printer RAM. These fonts will

then reside in the printer's RAM until the next power cycle.

Use the `lp` option `-ocs9N` (or `-oscs9N`) to select the ISO 8859-15 character set as the primary (or secondary) character set. For example:

```
lp -dprinter_name -ocs9N -oother_ options print_filename
```

NOTE   The case is significant. Be sure to use an upper case "N".

Euro - ISO 10646/Unicode Support

**new at 11i
original release**
HP-UX 11i provides system level support for the Unicode 2.1/ISO-10646 character set. Hewlett-Packard's support for Unicode provides a basis of enabling heterogeneous interoperability for all geographic areas.

ISO-10646 is an industry standard for defining a single encoding which uniquely encodes all the characters of the modern world. Unicode 2.1 is the companion specification to ISO-10646. Unicode specification at revision 2.1 includes the Euro symbol at 0X20AC code point.

Euro support to input, store, retrieve, display and print the Euro symbol has been added for this release. In addition to the base functionalities, HP-UX 11i is providing the following new functionalities:

- Dual currency support using @euro modifier.

- UTF-8 (Universal Transformation Format - 8 Bit) performance tuning.

- Euro display and processing capabilities for Asian UTF-8 locales.

- Additional converter tables.

Specific enhancements are provided to `locales`, `localedef`, `libc`, `Xlib` and `iconv` converter tables to achieve those new functionalities.

A subset of existing European locales has been modified to support dual currency to meet euro standard monetary formatting.

The following table gives the list of euro locales being supplied which support dual currency:

Locale	Language/Country
de_DE.utf8	German (Germany)
es_ES.utf8	Spanish (Spain)
fr_FR.utf8	French (France)
it_IT.utf8	Italian (Italy)
sv_SE.utf8	Swedish (Sweden)

The following table gives the list of locale sources being supplied which include dual currency support:

Locale	Language/Country
da_DK.utf8	Danish (Denmark)
de_DE.utf8	German (Germany)
el_GR.utf8	Greek (Greece)
en_GB.utf8	English (Great Britain)
es_ES.utf8	Spanish (Spain)

Locale	Language/Country
i_FI.utf8	Finnish (Finland)
fr_FR.utf8	French (France)
is_IS.utf8	Icelandic (Iceland)
it_IT.utf8	Italian (Italy)
nl_NL.utf8	Dutch (The Netherlands)
no_NO.utf8	Norwegian (Norway)
pt_PT.utf8	Portuguese (Portugal)
sv_SE.utf8	Swedish (Sweden)

When the `LANG` and/or `LC_*` environment variables are set to a euro-supported locale, the national monetary formatting rules are used. The `LC_MONETARY` environment variable should be set to the euro-supported locale name with `@euro` modifier to use/access euro monetary formatting rules.

For example, to specify the Euro as the currency for French, the following should be set:

```
LANG =3D fr_FR.utf8
LC_MONETARY =3D fr_FR.utf8@euro
```

Similarly, to specify French francs the following should be set:

```
LANG=3Dfr_FR.utf8
```

To access the monetary unit and the related monetary formatting rules programmatically, the programmer needs to toggle between the alternate monetary units via *setlocale* (3C) calls:

```
/* Handle euro in strfmon(), ... */
setlocale(LC_MONETARY, "fr_FR.utf8@euro");


...
/* Handle French francs in strfmon(), ... */
setlocale(LC_MONETARY, "fr_FR.utf8");
```

When the `LC_MONETARY` environment variable is set to `euro`, the formatting in monetary category will use euro standard formatting rules whereas other categories will still use local conventions in formatting. As a result, users may encounter a change to the decimal and thousandths separators for the currency, whereas decimal and thousandths separators outside the monetary area, like in numeric numbers, remain as per local conventions.

For example, in the French locale the thousandths separator is a space and the decimal point is a comma. However, the international standard for the thousandths separator for the euro currency is a period. So, a user that has the `LC_MONETARY` locale category set to "`fr_FR.utf8@euro`" will see the following behavior:

- The number "One thousand five hundred and fifty and a half" outside the monetary area will be displayed as `1 550,50`.

- The monetary number "One thousand five hundred and fifty euro and 50 cents" will be displayed as `EUR 1.550,50`

Commands

The *localedef* (1M) command has been enhanced to handle @euro modifier in order to build dual currency locale(s).

The *lp* (1) model scripts for the dual currency locales have been enhanced to print euro character.

libc

Standard libc supports @euro dual currency.

Codeset Converters

New `iconv` converter tables exist to support conversion from/to utf8, ucs2 and iso885915, PC code pages and IBM's euro enabled code pages.

New `iconv` converter tables are available to support conversion from utf8, ucs2, and iso885915 to IBM's euro enabled code pages and PC code pages:

Table 15-7 **utf8 and IBM's code pages (EBCDIC)**

utf8 <-> cp1140	utf8 <-> cp1141	utf8 <-> cp1142	utf8 <-> cp1143
utf8 <-> cp1144	utf8 <-> cp1145	utf8 <-> cp1146	utf8 <-> cp1147
utf8 <-> cp1148	utf8 <-> cp1149		

Table 15-8 **ucs2 and IBM's code pages (EBCDIC)**

ucs2 <-> cp1140	ucs2 <-> cp1141	ucs2 <-> cp1142	ucs2 <-> cp1143
ucs2 <-> cp1144	ucs2 <-> cp1145	ucs2 <-> cp1146	ucs2 <-> cp1147
ucs2 <-> cp1148	ucs2 <-> cp1149		

Table 15-9 **iso885915 and IBM's code pages (EBCDIC)**

iso885915<->cp1140	iso885915<->cp1141	iso885915<->cp1142	iso885915<->cp1143
iso885915<->cp1144	iso885915<->cp1145	iso885915<->cp1146	iso885915<->cp1147
iso885915<->cp1148	iso885915<->cp1149		

Table 15-10 **utf8 and PC code pages (EBCDIC)**

utf8 <-> cp437	utf8 <-> cp737	utf8 <-> cp775	utf8 <-> cp850
utf8 <-> cp852	utf8 <-> cp855	utf8 <-> cp857	utf8 <-> cp1860
utf8 <-> cp861	utf8 <-> cp862	utf8 <-> cp863	utf8 <-> cp864
utf8 <-> cp865	utf8 <-> cp866	utf8 <-> cp869	utf8 <-> cp874
utf8 <-> cp1250	utf8 <-> cp1251	utf8 <-> cp1252	utf8 <-> cp1253
utf8 <-> cp1254	utf8 <-> cp1255	utf8 <-> cp1256	utf8 <-> cp1257
utf8 <-> cp1258			

Table 15-11   **ucs2 and PC code pages (EBCDIC)**

ucs2 <-> cp437	ucs2 <-> cp737	ucs2 <-> cp775	ucs2 <-> cp850
ucs2 <-> cp852	ucs2 <-> cp855	ucs2 <-> cp857	ucs2 <-> cp1860
ucs2 <-> cp861	ucs2 <-> cp862	ucs2 <-> cp863	ucs2 <-> cp864
ucs2 <-> cp865	ucs2 <-> cp866	ucs2 <-> cp869	ucs2 <-> cp874
ucs2 <-> cp1250	ucs2 <-> cp1251	ucs2 <-> cp1252	ucs2 <-> cp1253
ucs2 <-> cp1254	ucs2 <-> cp1255	ucs2 <-> cp1256	ucs2 <-> cp1257
ucs2 <-> cp1258			

Impact

To use euro monetary formatting rules, the `LC_MONETARY` environment variable must be set to the euro supported locale name with the @euro modifier appended to it.

The size requirement for locale sources and binaries is 20.1 MB, while the converter tables size requirement is 191 KB.

There are no compatibility issues involved with the addition of these features.

Applications using UTF-8 locales should see improved collation performance as compared with UTF-8 locales delivered in the previous releases.

Asian System Environment (ASE

**new at 11i
original release**

HP-UX provides Asian systems for the Asian countries of the Far East, consisting of the following products:

JSE	Japanese System Environment
KSE	Korean System Environment
SSE	Simplified-Chinese System Environment
TSE	Traditional-Chinese System Environment

HP-UX provides several Asian enhancements as server features, including some new Asian codesets, UDC (User Defined Characters, or Gaiji), printing, and codeset conversions with mainframe codesets.

The new, changed, deleted features as well as some troubleshooting information is described below. For further information, see the following documentation:

- JSE

  — *Japanese System Environment User's Guide* (B3782-90873)
  — *HP XJIM Japanese Input Method Guide* (B3782-90869)
  — *ATOK8 Japanese Input Method Guide* (B3782-90870)
  — *EGBridge Japanese Input Method Guide* (B3782-90871)
  — *VJE-gamma Japanese Input Method Guide* (B3782-90872)

- *KSE - Korean System Environment User's Guide* (5969-4454)

- *SSE - Simplified Chinese System Environment User's Guide* (5969-4455)

- *TSE - Traditional Chinese System Environment User's Guide* (5969-4453)

To get release information on earlier versions of ASE, see the following files:

- JSE: `/usr/share/doc/ASX-JPN`
- KSE: `/usr/share/doc/ASX-KOR`
- SSE: `/usr/share/doc/ASX-SCH`
- TSE: `/usr/share/doc/ASX-TCH`

New Features

- ASE Common

  ❏ New printer model

    New printer models are supported on both the LP Spooler and HPDPS. You can print plain text file on the following printers by configuring the printer using the PCL5.nloo(PCL5.asian) model file on the LP Spooler or PCL5.asx(2BPCL5.asx) printer model on HPDPS:

    HP LaserJet 4000(N)
    HP LaserJet 4050(N)
    HP LaserJet 4500(N)
    HP LaserJet 5000(N)
    HP LaserJet 8000(N)
    HP LaserJet 8100N

NOTE

By installing optional Font DIMM on these printers, you can print text with TrueType fonts. To use TrueType fonts, you have to configure a printer with PCL5.asian model file for the LP Spooler, or with 2BPCL5.asx printer model for HPDPS.

❏ HPDPS common printer model directory

For HPDPS, the common printer model directories PCL5.asx, 2BPCL5.asx and ESCP.asx are provided for future new printer support. The user can copy these sample printer model directories to a directory under `/var/opt/pd/lib/model` with an appropriate name and customize it to be suited for the printer being configured.

• JSE

❏ ATOK X for HP-UX Preview Edition

The new version of ATOK is now supported. As a Kana-Kanji conversion feature, the ATOK12 engine is incorporated enabling you to achieve a comfortable and effective Japanese input environment. As this release of ATOK X is a Preview Edition, some of the customization tools are not yet available. In the next release, a full featured ATOK X for HP-UX will be provided.

❏ Unicode

Japanese UTF-8 locale ja_JP.utf8 is supported. Using this locale, you can input, display and print UTF-8 characters. It supports characters defined in standards JIS X 0201 (1976), JIS X 0208 (1990), and JIS X 0212 (1990). UDC (User Defined Characters or GAIJI) and VDC (Vender Defined Characters) are not supported.

For details, see the document `/usr/share/doc/ASX-UTF8`.

❏ USB (Universal Serial Bus) Japanese 109 Keyboard support

This allows for inputting Japanese characters by Japanese input methods.

❏ NEC VDC symbols for display on X Window System

NEC special characters are included in Japanese fonts. NEC VDC has 83 characters which occupy following code areas:

JIS[Kuten]: 13/01 - 13/92

Shift-JIS: 0x8740 - 0x879C

Those characters can be shown on X Window System.

❏ New Ricoh TrueType font package

The new Ricoh TrueType font package "TrueTypeWorld ValueFontD2" is supported. The supported fonts are Windows 3.1 version of WABUN (Japanese) fonts.

❏ New printer model

New printer models are supported on both the LP Spooler and HPDPS. You can print Japanese plain text file on the following printers by configuring the printer

using the specified model file on the LP Spooler or printer model on HPDPS:

Table 15-12   New Printer Models for JSE

Printer	LP Spooler Model File	HPDPS Printer Model File
HP LaserJet 5si[a]	PCL5.nloo (PCL5.asian)	PCL5.asx (2BPCL5.asx)
HP HITPCPDA	ESCP	ESCP.asx
HP HITHTS4A	ESCP	ESCP.asx
HP HITKD20A	ESCP	ESCP.asx
HP HITKD45A	ESCP	ESCP.asx
Canon LBP-850	LIPS4	LIPS4.asx
Canon LBP-930EX	LIPS4	LIPS4.asx
Canon LBP-2030	LIPS4	LIPS4.asx
Canon LBP-2040	LIPS4	LIPS4.asx
Canon LBP-2160	LIPS4	LIPS4.asx
OKI Microline 9XXPSII[b]	PS2.nlio	PS2.asx
OKI Microline 9XXPSIII[b]	PS2.nlio	PS2.asx
OKI Microline 703N(3)[b]	PS2.nlio	PS2.asx
EPSON VP-1800	ESCP	ESCP.asx
OKI 533OS	ESCP	ESCP.asx
OKI 835OS	ESCP	ESCP.asx
OKI 858OS	ESCP	ESCP.asx
NEC LL-15 (NPDL2)	NPDLII	NPDLII
NEC LL-30 (NPDL2)	NPDLII	NPDLII
NEC LL-15 (ESC/P) [c]	ESCP	ESCP.asx
NEC LL-30 (ESC/P) (*3)	ESCP	ESCP.asx

a. By installing optional Japanese Font DIMM on these printers, you can print Japanese text with TrueType fonts. To use Japanese TrueType fonts, you have to configure a printer with PCL5.asian model file for the LP Spooler, or with 2BPCL5.asx printer model for HPDPS. To see whether your printer has Japanese TrueType Font installed, follow these steps:

1. Press **Menu** on the control panel of the printer until "INFORMATION MENU" appears.
2. Press **Item** until "PRINT PCL FONT LIST" appears.
3. Press **Select** to print the font list.
4. If your printer has Japanese TrueType font, you will see "MS Mincho" and "MS Gothic" in the printed list.

b. Printing text files on expanded A3 (called "A3-Nobi" in Japan) paper is not supported.

c. There are restrictions of page length setting on ESC/P mode. For detail, see manual of the printer and online document /usr/share/doc/PRINTER-JPN-S[E].

❏ HPDPS common printer model directory

For HPDPS, the common printer model directories `LIPS3.asx`, `LIPS4.asx` and `PS.asx` are provided for future new printer support. The user can copy these sample printer model directories to a directory under `/var/opt/pd/lib/model` with an appropriate name and customize it to be suited for the printer being configured.

❏ Mainframe code set conversion

The Mainframe code set conversions are provided to convert code sets between Mainframe code sets Hitachi KEIS, NEC JIPS, Fujitsu JEF, and IBM EBCDIC with existing code sets SJIS, eucJP, and ucs2. These code conversions are used from *iconv* (1) and *iconv* (3C).

The following code sets are supported:

— Hitachi KEIS

— keis7k: KEIS78 (Hitachi MF code set based on JIS C6226-1978) + EBCDIK
— keis8k: KEIS83 (Hitachi MF code set based on JIS X0208-1983) + EBCDIK
— keis7c: KEIS78 (Hitachi MF code set based on JIS C6226-1978) + EBCDIC
— keis8c: KEIS83 (Hitachi MF code set based on JIS X0208-1983) + EBCDIC

— NEC JIPS

— jipsj: JIPS (NEC Mainframe code set) JIS
— jipsec: JIPS (NEC Mainframe code set) EBCDIC
— jipsek: JIPS (NEC Mainframe code set) EBCDIK

— Fujitsu JEF

— jefc: JEF (Fujitsu Mainframe code set) + EBCDIC (lower alphabet)
— jefk: JEF (Fujitsu Mainframe code set) + EBCDIK (katakana)
— jefc9p: JEF + EBCDIC designating 9 point size in printing
— jefk9p: JEF + EBCDIK designating 9 point size in printing

The code set conversions are provided between the above Mainframe code sets and the following existing code sets:

 SJIS

eucJP

ucs2

❏ New UDC feature

A new UDC environment is provided for client/server or distributed environments. You can share UDC font on a single server machine and print UDC from client machines. As a UDC font, TrueType font is supported. You can use UDC TrueType font created on X Window or provided by some vendors. Two typefaces are supported as UDC fonts. ESC/P and PCL printers are supported.

• KSE

❏ Unicode

The Korean UTF-8 locale ko_KR.utf8 is supported. On this locale, you can input, display and print UTF-8 characters. There is support for characters defined in standards KSC 5636 (1989) and KSC 5601 (1987). UDC (User Defined Characters or GAIJI) and VDC (Vender Defined Characters) are not supported. For details, see the document /usr/share/doc/ASX-UTF8.

The full Hangul Syllables in KS X 1005-1 (old name is KS C 5700-1995) are supported on ko_KR.utf8 locale. You can input full Hangul characters by XKIM and display on X Window System. With Korean font DIMM and PCL5.asian model file, you can print full Hangul characters.

❏ Euro and registered trademark ® symbols

The printing of the Euro symbol in the ko_KR.eucKR locale is supported. The registered trademark symbol ® is also supported. PCL printers are supported to print these symbols with PCL5.asian model file. Two typefaces, Dotum and Batang, are supported. You can print Euro and ® symbols without any printing options.

❏ USB (Universal Serial Bus) Korean 106 Keyboard

USB Korean 106 Keyboard is supported for inputting Korean characters by Korean input method XKIM.

❏ X Print Server

KSE supports printing via the X Print Server to PCL printers.

• SSE

❏ Unicode

Simplified Chinese UTF-8 locale zh_CN.utf8 is supported. On this locale, you can input, display and print UTF-8 characters. There is support for characters defined in standards ISO 646 (1991) and GB 2312 (1980). UDC (User Defined Characters or GAIJI) and VDC (Vender Defined Characters) are not supported. For details, see the document /usr/share/doc/ASX-UTF8

❏ USB (Universal Serial Bus) Simplified Chinese 104 Keyboard

The USB Simplified Chinese 104 Keyboard is supported for inputting Simplified Chinese characters by the input method XSIM.

❏ X Print Server

SSE supports printing via the X Print Server to PCL printers.

• TSE

❏ Unicode

Traditional Chinese UTF-8 locales zh_TW.utf8 and zh_HK.utf8 are supported. On these locales, you can input, display and print UTF-8 characters. There is support for characters defined in standards ISO 646 (1991), CNS 11643 (1992) plane 1, 2, 3 and 4, except for some characters which are not supported by Unicode 2.0. UDC (User Defined Characters or GAIJI) and VDC (Vender Defined Characters) are not supported. For details, see the document /usr/share/doc/ASX-UTF8.

❏ USB (Universal Serial Bus) Traditional Chinese 104 Keyboard

USB Traditional Chinese 104 Keyboard is supported for inputting Traditional Chinese characters by the input method XTIM.

❏ X Print Server

TSE supports printing via the X Print Server to PCL printers.

❏ HongKong big5 Support (new)

Locale support is provided with the big5 codeset for HongKong.

HP provides support for the HongKong big5 locale, zh_HK.big5. HongKong big5 locale is similar to Traditional Chinese big5 locale. The difference between these two locales are in monetary and date/time properties which reflect local cultural conventions.

CDE has been enhanced to support this new locale by providing the required app-defaults files to CDE applications.

Impact

Applications must elect to enable big5 support by setting the `LANG` and/or `LC_*` environment variables to the HongKong big5 locale.

The size requirement for locale source and binaries is 1.7 MB

Applications using HongKong big5 locales should see the same performance as of Traditional Chinese big5.

Changed Feature

• JSE

❏ **EISUU** key mode change for 106/109 keyboard

In the previous version, the **EISUU** key, **Shift** + **EISUU** (Caps Lock mode) keys, and **Alt** + **EISUU** (KANJIBANGOU mode) keys all worked as **Caps Lock**. Now they work as original features of the key/keys.

Deleted Features

• ASE Common

❏ Printing to LaserJet III series is now obsoleted. If you are currently using LaserJet III series printers, you should use newer printer models.

• KSE

❏ XDevice is not included from this release.

NOTE   The Japanese input methods EGBridge and VJE-gamma will be obsoleted in an upcoming release.

Troubleshooting Information

• JSE

❏ XJIM

— On a low-resolution display, customize window is cut off by default. Specify 14-dot font with `-fn` option or XJim*fontList resource.

— If you use 'KANA' input (not 'ROMAJI' input) as the key input method at 'YOMI' input, and you input a 'KANA' character and 'HANDAKUTEN' or 'DAKUTEN' successively, the input method server does not compose 'KANA' with 'DAKUTEN' or 'HANDAKUTEN' as one character, but displays the 'KANA' character and 'DAKUTEN' or 'HANDAKUTEN' symbol. In this case, you should make the composite character using 'ZENKAKU-HIRAGANA' conversion (press **Shift** + **F5** key), or 'ZENKAKU-KATAKANA' conversion (press **F6** key).

— If you install XJIM after NIS configuration, you will find that you can not use XJIM Conversion Server. To resolve this problem, move the following line in the `/etc/services` file

```
nuekks         6897/tcp         # nuekks daemon
```

to the position above the line which begins with a "+" sign indicating the start of NIS mapping.

❏ EGBridge

Closing the EGBridge main window during Kana-Kanji conversion on hpterm may also close hpterm. You should finish conversion before closing the EGBridge main window.

❏ IMS common (XJIM/ATOK8/EGBridge/VJE-gamma)

— Window focus sometimes cannot be moved by **Meta**(**Alt**)-**Tab** key if applications use XIMStatusNothing and they overlap each other with KANJI-ON state. To avoid this problem, set `stackChange` resource to `False` as follows:

XJIM	`XJim*stackChange: False`
ATOK8	`Atok8*stackChange: False`
EGBridge	`EGIms*stackChange: False`
VJE	`Vje*stackChange: False`

See the "Resource" section in each Input Method manual for details.

— On Motif 1.2 and Motif 2.1 applications, the **F10** and **Shift**-**F10** keys cannot be used as the Japanese input function key because those keys are used to switch focus to the menu bar. To assign these keys to certain functions for IMS, set the following:

— for DIN keyboard: `$ xmodmap -e "keycode 25 = F10"`
— for ITF keyboard: `$ xmodmap -e "keycode 38 = F10"`

— Japanese IMS is not available with X11R4 (including Motif 1.1) applications using PS2-DIN-JIS keyboard if $LANG is "ja_JP.SJIS" or "ja_JP.eucJP". To avoid this problem, set $LANG "japanese" or "japanese.euc" when invoking X11R4 (Motif 1.1) applications.

— Even if you merge UDC in X font after running the input method server, the server cannot display UDC in the pre-edit and the candidate. You should merge UDC in X font server before running the input method server. Re-login makes sure that the input method server displays UDC on CDE.

❏ JIS keyboard

— Do not set the `KBD_LANG` shell variable or Motif 1.1 applications will not work with a JIS keyboard.

— The **Yen** key on JIS keyboard with X terminal does not work correctly. To use the **Yen** key, execute the command.

```
$ xmodmap -e "keysym yen = backslash bar prolongedsound"
```

❑ 106/109 Keyboard

— You cannot turn off EGBridge (although you can turn on). The solution is to change the key map file `$HOME/.egb/EGBMap` (for personal use) or `/etc/opt/egb/config/EGBMap` (for system use). You open the key map file with an editor and change the following entry:

old: `LKONOFF = XK_Henkan XK_Meta_L`
new: `LKONOFF = XK_Henkan XK_Meta_L XK_Alt_L`

Then save the updated key map file and restart EGBridge. You can turn EGBridge on/off with the left "**Alt**" key.

❑ `udcload`

— When UDCs are not arranged in the code order in the UDC file, `udcload` cannot load UDC. Therefore. you should arrange UDCs in the code order. UDCs generated by `xudced` have no problem because `xudced` generates UDSs arranged in code order.

• KSE

❑ `xk0input`

 Xkim is not available with X11R4 (including Motif 1.1) applications using PS2-DIN keyboard if `LANG` is `ko_KR.eucKR`. To avoid this problem, set `LANG` to `korean` when invoking X11R4 (Motif 1.1) applications.

• ASE Common

❑ `xudced` (UDC editor)

When you select "Search…" in the main menu "Edit", you cannot specify the character directly. Only the Index number can be specified to search a character.

Enhanced Print Capabilities in the Asian System Environment

HP-UX 11i contains enhancements to the printer capabilities of four Asian-country system environments (JSE, KSE, SSE, TSE), as itemized below.

Changes Common to All ASEs

- LP Model File: Supports new printers: The `PCL5.nloo` model file supports Asian text printing on following printers.

    — HP LaserJet 4000
    — LaserJet 5000
    — LaserJet 8000

- HPDPS: Provides common printer model directories: Provides new printer model directories, `PCL4.asx`, `PCL5.asx` and `ESCP.asx` for future printer support. Users can use these model directories as *model* or *sample* implementation of a printer-model. Users can copy these sample printer model directories to a directory under `/var/opt/pd/lib/model` with an appropriate name and customize it for the printer being configured.

    Supports new printers: Users can print Asian text on the following printers through HPDPS by configuring the printer with `PCL5.asx` printer-model.

    — HP LaserJet 4000
    — LaserJet 5000
    — LaserJet 8000

For more information, see the following files in `/usr/share/doc/`:

`ASX-JPN, ASX-JPN-S, ASX-JPN-E, ASX-KOR, ASX-SCH, ASX-TCH`

Japanese System Environment (JSE)

- LP Model File: Supports new printers. The `PS.nlio` model file supports Japanese text printing on these printers:

    — OKI ML703N
    — ML600PSII

    The `ESCP` model file supports Japanese text printing on these printers:

    — OKI 5330S
    — 8350S
    — 8580S
    — EPSON VP-1800

    The `PCL5.asian` model file supports Japanese text printing on:

    — HP LaserJet 5Si with 2Byte Font SIMM
    — LaserJet 4000 with 2Byte Font DIMM
    — LaserJet 5000 with 2Byte Font DIMM
    — LaserJet 8000 with 2Byte Font DIMM

- HPDPS: Provides common printer model directories: Provides new printer model

directories, `LIPS3.asx`, `LIPS4.asx`, `PS.asx` and `2BPCL5.asx` for future printer support. Users can use these model directories as *model* or *sample* implementation of a printer-model. Users may copy these sample printer model directories to a directory under `/var/opt/pd/lib/model` with an appropriate name and customize it to suit the printer being configured.

Supports new printers: Users can print Japanese text on the following printers through HPDPS, by configuring the printer with the `2BPCL5.asx` printer-model:

— HP LaserJet 5Si with 2Byte Font SIMM
— LaserJet 4000 with 2Byte Font DIMM
— LaserJet 5000 with 2Byte Font DIMM
— LaserJet 8000 with 2Byte Font DIMM

Users can print Japanese text on following printers through HPDPS by configuring the printer with the `PS.asx` printer-model:

— OKI ML703N
— ML600PSII

Users can print Japanese text on following printers through HPDPS by configuring the printer with the `ESCP.asx` printer-model:

— OKI 5330S
— 8350S
— 8580S
— EPSON VP-1800

For more information, see the following files in `usr/share/doc/`: `ASX-JPN`, `ASX-JPN-S`, `ASX-JPN-E`, `PRINTER-JPN-S`, `PRINTER-JPN-E`

Korean System Environment (KSE)

• X Print Server: KSE supports printing via X Print Server to PCL printers.

• LP and HPDPS: Supports new print options. Supports new printers.

• HPDPS: Provides a common template model directory for each print language.

For more information, see the file: `/usr/share/doc/ASX-KOR`.

Simplified-Chinese System Environment (SSE)

• X Print Server: SSE supports printing via X Print Server to PCL printers

• LP and HPDPS: Supports new print options. Support new printers.

• HPDPS: Provides common template model directory for each print language.

For more information, see the file: `/usr/share/doc/ASX-SCH`.

Traditional-Chinese System Environment (TSE)

• X Print Server: TSE supports printing via X Print Server to PCL printers

• LP and HPDPS: Supports new print options. Supports new printers.

• HPDPS: Provides common template model directory for each print language.

For more information, see the file: `/usr/share/doc/ASX-TCH`.

Multibyte Support Extension and Unix98 Support

**new at 11i
original release**

A new set of multibyte APIs have been added to `libc` following the C99 specification (ISO/IEC 9899:1999), and the Unix98 specification.

These APIs extend the already existing multibyte and wide character APIs in order to be able to:

- perform input and output of wide character, or multibyte character, or both

- perform general wide string manipulation

- provide extended capabilities for conversion between multibyte and wide character sequences

Several new design concepts have been introduced:

- Stream orientation

- Restartable APIs and the conversion state

Stream Orientation

A stream can be either wide-character or byte-oriented. The orientation of a stream is a concept based on an input/output model that assumes that characters are handled as wide characters within an application and stored as multibyte characters in files, and that all the wide-character input/output functions begin executing with the stream positioned at the boundary between two multibyte characters.

After a stream is associated with a file, but before any operations are performed on the stream, the stream is without orientation. If a wide-character input or output function is applied to a stream without orientation, the stream becomes wide-oriented implicitly. Likewise, if a byte input or output operation is applied to a stream without orientation, the stream becomes byte-oriented implicitly. Once the stream becomes oriented, the orientation is fixed and cannot be changed until the stream is closed.

Restartable APIs and the Conversion State

A new set of APIs have been introduced to facilitate the conversion between multibyte character representations to wide character representations. These APIs use a new object type, `mbstate_t`, that can hold the conversion state information necessary to convert between sequences of multibyte characters and wide characters. The conversion state determines the behavior of a conversion between multibyte and wide-character encoding. For conversion from multibyte characters to wide characters, the conversion state stores information, such as the position, within the current multibyte character (as a sequence of characters or a wide character accumulator). For conversions in either direction, the conversion state stores the current shift state, if any, and possibly, the encoding rule.

As these APIs store the partial character information, a multibyte sequence can be processed one byte at a time, and the processing can be interrupted and continued (i.e., restarted) at some other point in time, so the new multibyte/wide-conversion utilities are thus made restartable by using the information in the `mbstate_t` object.

How to Get MSE/Unix98 Behavior

In order to get MSE/Unix98 behavior, the programs have to be compiled with the `-D_XOPEN_SOURCE=500` macro definition and the variable has to be defined in the environment.

Under the Korn, Bourne, and POSIX shells, this is done with:

```
UNIX_STD=98
export UNIX_STD
```

Under the C shell this is done using

```
setenv UNIX_STD 98
```

A `cc` compiler equal to HP92453-01 A.11.01.20 HP C Compiler or newer is required to get this functionality.

Below is a summary list of new and modified APIs. For further details, please refer to the corresponding manpages.

New Interfaces

The following APIs are newly added to libc and will not affect existing code:

btowc()

`btowc()` returns the wide-character representation of a given single-byte character.

fwide()

`fwide()` sets the stream orientation.

fwprintf(), swprintf(), wprintf()

These APIs print formatted wide-character output.

fwscanf(), swscanf(), wscanf()

These APIs process formatted wide-character input.

mbrlen()

`mbrlen()` returns the number of bytes in a wide character. Note that the behavior of this function is affected by the `LC_CTYPE` category of the current locale.

mbrtowc()

`mbrtowc()` converts a stream of bytes to a wide-character code. Note that the behavior of this function is affected by the `LC_CTYPE` category of the current locale.

mbsinit()

`mbsinit()` determines whether the object pointed to by the first argument, which contains shift state information, describes an initial conversion state.

mbsrtowcs()

`mbsrtowcs()` converts a character string to a wide-character string. Note that the behavior of this function is affected by the `LC_CTYPE` category of the current locale.

towctrans()

towctrans() is provided for character transliteration. The current setting of the
LC_CTYPE category should be the same as during the call to wctrans()-.

vfwprintf(), vswprintf(), vwprintf()

These APIs are provided for printing wide-character formatted output of a stdarg
argument. They are similar to *fwprintf* (3C) except that instead of being called with a
variable number of arguments, they are called with an argument list as defined by
<stdarg.h>.

wcrtomb()

wcrtomb() converts a wide-character to a multibyte character. It determines the
number of bytes needed to represent the character corresponding to the wide-character
code whose value is specified by the second argument.

wcsrtombs()

wcsrtombs() converts a wide-character string to a character string. Note that the
behavior of this function is affected by the LC_CTYPE category of the current locale.

wcsstr()

wcsstr() finds a substring in a wide-character string. Note that the behavior of this
function is affected by the LC_CTYPE category of the current locale.

wctob()

wctob() converts wide-character to single-byte.

wctrans

wctrans() defines character mapping in the current locale. Note that the values
returned by wctrans() are valid until a call to setlocale() that modifies the category
LC_CTYPE.

wmemchr(), wmemcmp(), wmemcpy(), wmemmove(), wmemset()

These APIs operate with wide-character in memory areas:

— wmemchr() finds a wide-character in a memory array.
— wmemcmp() compares wide-characters in memory.
— wmemcpy() copies wide-chracter in memory.
— wmemmove() copies wide-characters in memory with overlapping areas.
— wmemset() sets wide-characters in memory.

Modified interfaces

The following APIs may have a change in behavior or a parameter type change that
could affect existing HP-UX code when the Unix98 support is selected:

fprintf(), printf(), snprintf(), sprintf(), fscanf(), scanf(), sscanf()

printf (3S), *scanf* (3S) and related functions support the new qualifier l (the letter) to
select wide character conversion in a given format string and set errno to [EILSEQ] if the
data obtained from the input stream does not form a valid wide character.

fputwc(), putwc(), putwchar()

The type of first argument is changed from `wint_t` to `wchar_t`.

freopen()

Regardless of the mode of underlying stream, after a successful call to the `freopen()` function, the orienting of the stream is cleared and the associated `mbstate_t` object is set to describe an initial conversion state.

wcschr(), wcsrchr()

The type of second argument is changed from `wint_t` to `wchar_t`.