Tru64 UNIX

Release Notes for Version 4.0F

Order Number: AA-QTLMF-TE

April 1999

Product Version: Tru64 UNIX Version 4.0F

This book contains notes on software and documentation restrictions for Compaq Tru64™ UNIX® (formerly DIGITAL UNIX) Version 4.0F and the bundled layered products that ship with the operating system.

This book also describes significant new and changed features in this version of the Tru64 UNIX operating system and lists features and interfaces scheduled for retirement in future releases.

Compaq maintains an online *Technical Update* that contains information about restrictions and problems that have been discovered since Version 4.0F began shipping. To view this document on the Web, go to the following URL:

```
http://www.unix.digital.com/faqs/publications
/updates/V40F/TITLE.HTM
```

Compaq recommends that you visit this site periodically to see if any new information has been added.

Compaq Computer Corporation
Houston, Texas

Contents

About This Manual

1   New and Changed Features

2   Installation Notes

3   Processor-Specific Notes

4   Base System Software Notes

5   Development Environment Notes

6   Window System Software Notes

7   Documentation Notes

8 Features and Interfaces Scheduled for Retirement

A   System Limits

B   Software Subset Information

Index

Tables

About This Manual

_____

This manual contains release notes for Compaq Tru64™ UNIX® (formerly DIGITAL UNIX) Version 4.0F.

This manual also describes significant new and changed features in this version of the Tru64 UNIX operating system and lists features and interfaces scheduled for retirement in future releases.

Audience

These release notes are for the person who installs the product and for anyone using the product following installation.

Organization

This manual is organized as follows:

Chapter 1	Contains an overview of new and changed features.
Chapter 2	Contains notes pertaining to installation.
Chapter 3	Contains processor-specific information.
Chapter 4	Contains information about the base operating system software.
Chapter 5	Contains information about the development environment.
Chapter 6	Contains information about the window system software.
Chapter 7	Contains information about the documentation.
Chapter 8	Contains information about features scheduled for removal in future versions of Tru64 UNIX.
Appendix A	Contains information about the system limits.
Appendix B	Contains information about the software subsets for the Tru64 UNIX software kit.

Related Documents

In addition to this manual, you should have the following documentation available during the installation of this release.

- The hardware documentation for your system

- The online or hardcopy reference pages
- The HTML book files that are provided on the CD–ROM
- *Read This First* letter

  This letter provides general information pertaining to the Tru64 UNIX Version 4.0F software.
- *Console Firmware Release Notes*

  This document includes the console firmware revision numbers and release notes.
- *Installation Guide*

  This document describes in detail how to install Tru64 UNIX.
- *Sharing Software on a Local Area Network*

  This document describes in detail how to set up and manage a Remote Installation Services (RIS) area for Tru64 UNIX systems. It also contains information on setting up Dataless Management Services (DMS).
- *System Administration*

  This document contains information on administering and maintaining your system.

The printed version of the Tru64 UNIX documentation set is color coded to help specific audiences quickly find the books that meet their needs. (You can order the printed documentation from Compaq.) This color coding is reinforced with the use of an icon on the spines of books. The following list describes this convention:

Audience	Icon	Color Code
General users	G	Blue
System and network administrators	S	Red
Programmers	P	Purple
Reference page users	R	Green
Device driver writers	D	Orange

Some books in the documentation set help meet the needs of several audiences. For example, the information in some system books is also used by programmers. Keep this in mind when searching for information on specific topics.

The *Documentation Overview* provides information on all of the books in the Tru64 UNIX documentation set.

Reader's Comments

Compaq welcomes any comments and suggestions you have on this and other Tru64 UNIX manuals.

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: UEG Publications, ZK03-3/Y32

- Internet electronic mail: `readers_comment@zk3.dec.com`

  A Reader's Comment form is located on your system in the following location:

  `/usr/doc/readers_comment.txt`

- Mail:

  Compaq Computer Corporation
  UEG Publications Manager
  ZK03-3/Y32
  110 Spit Brook Road
  Nashua, NH 03062-9987

  A Reader's Comment form is located in the back of each printed manual. The form is postage paid if you mail it in the United States.

Please include the following information along with your comments:

- The full title of the book and the order number. (The order number is printed on the title page of this book and on its back cover.)

- The section numbers and page numbers of the information on which you are commenting.

- The version of Tru64 UNIX that you are using.

- If known, the type of processor that is running the Tru64 UNIX software.

The Tru64 UNIX Publications group cannot respond to system problems or technical support inquiries. Please address technical questions to your local system vendor or to the appropriate Compaq technical support office. Information provided with the software media explains how to send problem reports to Compaq.

Conventions

The following conventions are used in this guide:

`%` `$`	A percent sign represents the C shell system prompt. A dollar sign represents the system prompt for the Bourne and Korn shells.
`#`	A number sign represents the superuser prompt.
`%` **`cat`**	Boldface type in interactive examples indicates typed user input.
`file`	Italic (slanted) type indicates variable values, placeholders, and function argument names.
[|] **{ | }**	In syntax definitions, brackets indicate items that are optional and braces indicate items that are required. Vertical bars separating items inside brackets or braces indicate that you choose one item from among those listed.
`. . .`	In syntax definitions, a horizontal ellipsis indicates that the preceding item can be repeated one or more times.
`cat`(1)	A cross-reference to a reference page includes the appropriate section number in parentheses. For example, `cat`(1) indicates that you can find information on the `cat` command in Section 1 of the reference pages.
Ctrl/*x*	This symbol indicates that you hold down the first named key while pressing the key or mouse button that follows the slash. In examples, this key combination is enclosed in a box (for example, Ctrl/C).

New and Changed Features 1

The Tru64 UNIX Version 4.0F software replaces earlier versions of the operating system software and provides new or additional support for the following:

- Maintenance for previous versions of the operating system
- Hardware-related support, including the following:
  - Support for new AlphaServer processors
  - Support for Fibre Channel based storage
  - Logical partitions for AlphaServer GS140 system
- New and enhanced features, including the following:
  - Upgraded Dynamic Host Configuration Protocol (DHCP)
  - Change to ATM PACING Virtual Circuits bandwidth accounting
  - Extensions to the `atom` command's programmable interfaces
  - Ability to examine the stack of a user program that caused a kernel crash dump
  - Introduction of a new Compaq Analyze utility
  - Secure version of the `syslog` facility
  - Changes to VFS and AdvFS tunable attributes
  - Addition of Visual Threads
  - Enhancement to the bootable tape facility
  - Introduction of Compaq Advanced Printing Software Version 1.0
  - A new version of Netscape Communicator
  - Introduction of Compaq Insight Manager
  - Support for Universal Serial Bus (USB) keyboards and mice
  - AltaVista search capability on documentation CD-ROM
  - Support for Digital Versatile Disc (DVD)
  - Availability of COM for Tru64 UNIX

For more information about some of these features, see the following sections. You should also refer to Chapter 8, which describes features and

interfaces scheduled to be retired in future versions of the operating system.

The following sections provide more information on the major new features included in Tru64 UNIX Version 4.0F.

1.1 New Hardware Support

Tru64 UNIX Version 4.0F adds support for the following systems:

- Compaq AlphaServer DS20
- Compaq AlphaServer ES40

1.2 Support for Fibre Channel Based Storage

Tru64 UNIX now contains support for Fibre Channel (FC) based storage. Initially, device support is limited to the KGPSA adapters and HSG80 disk arrays interconnected via FC switches. The FC devices appear to the system as normal SCSI devices. As such, Fibre Channel configurations are bounded by the SCSI-2 namespace exported via the standard SCSI disk (rz) and tape (tz) drivers. Many of the configuration restrictions will be removed in a future release.

As the Fibre Channel address space is much larger and more sparsely-populated than parallel SCSI, a persistent translation of SCSI target IDs to FC addresses is necessitated in the device driver for the FC adapter. Refer to the emx(7) and emx_data.c(4) reference pages for further information on how this is performed.

For the Version 4.0F release, support for Fibre Channel storage is limited to data devices only. A Fibre Channel disk device cannot be used for boot, swap, or dump. Support for these features will appear in a later release.

Cautions

The Version 4.0F installation software does not discern between Fibre Channel and Parallel SCSI disks. All are presented as potential install-to devices. If a Fibre Channel disk is selected, a portion of the installation process will succeed before an ultimate failure is encountered. This issue will be resolved in an upcoming release.

Fibre Channel disk devices can be added as secondary swap devices via fstab(4) or swapon(8). However, doing so usually places the FC device into the list of potential dump devices for the system. In the event of a system panic, the FC disk can be selected as a dump target. If so, the system will encounter a fatal error and will abort the dump process. This issue will be resolved in an upcoming release.

1.3 Logical Partitions for AlphaServer GS140 System

The logical partitioning feature allows a single AlphaServer GS140 system to be partitioned into as many as three logical partitions. Each partition runs a separate instance of the operating system. For a complete description of logical partitions, refer to the *System Adminstration* guide.

The minimum AlphaServer GS140 firmware revision for logical partitions support is Version 5.4.

At the time of Version 4.0F release, the AlphaServer GS140 Model 6-525 is the only system supported for logical partitioning. Refer to the *Systems and Options Catalog* for information on newly-supported systems.

1.4 Dynamic Host Configuration Protocol

This version of Tru64 UNIX includes an upgraded and improved Dynamic Host Configuration Protocol (DHCP) database. The new database supports a new database format. Therefore, if you are currently using DHCP from a prior version of Tru64 UNIX, you must convert the database to the new format.

The improvements to DHCP include:

- Force broadcast reply

- Ignore hardware type

- Return bootfile in header

- Server policy consistency check

- Allow the server canonical name

For more information on DHCP, see the *Network Administration* guide, the *Technical Overview*, and the `joind`(8) reference page.

1.5 ATM PACING VC Bandwidth Accounting Change

In this release, the bandwidth used by pacing virtual circuits (VCs) is not deducted from the total bandwidth in the constant bit rate (CBR) reserved resource pool. Therefore, ATM drivers can now overbook the driver's bandwidth. In addition, reserved resources can be applied to CBR VCs only.

In previous releases of the operating system, the bandwidth used by pacing VCs was deducted from the total bandwidth in the constant bit rate (CBR) reserved resource pool. This prevented ATM drivers that supported paced unspecified bit rate (UBR) traffic from allowing paced VCs to overbook the driver's bandwidth.

1.6 Extensions to atom's Programmable Interfaces

The `atom` command has been enhanced. It now provides a `-pthread` flag and RTL routines to help tool-writers develop thread safe program analysis tools, including support for deferred (but not asynchronous) thread cancellation. The cancellation support has some impact on the performance of thread-safe tools. See the `atom`(1), `Thread`(5), and `pthread_setcanceltype`(3) reference pages.

The `atom` command's `AddCallProto`(5) instrumentation RTL routine now provides the `FREGV` argument type, to let tool writers pass floating-point register values to analysis routines.

The `atom` command's `-gp` flag now retains all debug information in applications that were compiled with the `-g` option.

The `atom` command's new `-w2` switch suppresses warnings regarding analysis routines.

1.7 Ability to Examine the Stack of a User Program that Caused a Kernel Crash Dump

You can now view the user stack of a program that precipitated a kernel crash dump. For more information, see Section 5.5.

1.8 Compaq Analyze Utility

The Compaq Analyze utility enables you to analyze binary error log files (located in `/var/adm/binary.errlog`). Refer to the *System Administration* guide for information on using this new tool.

Compaq Analyze supports selected EV6-class processors in this release.

To function on a supported system, Compaq Analyze requires a modified Common Error Log Header, which is created automatically if `binlogd` detects that the value of `rpb_proctype` is equivalent to `ST_DEC_6600` (EV6-series processors). This modified header is supported by DECevent.

The AlphaServer 8400 (with EV6 processors) is not supported by Compaq Analyze. This system is supported by DECevent using the older error log format.

1.9 Secure Version of syslog

The security of the `syslog` facility has been enhanced in this release. Unless the domain host name of a remote host is entered in the local file, `/etc/syslog.auth`, the local system will not log any `syslog` messages from that remote host.

If you are installing the secure version of `syslogd` on a system, and you have configured or intend to configure other hosts to forward `syslog` messages to the system, refer to the `syslogd`(8) and `syslog.auth`(4) reference pages.

1.10 Changes to VFS and AdvFS Tunable Attributes

The default value of the Virtual File System (VFS) `vfs` subsystem attribute `name-cache-hash-size` has changed from 256 to a dynamically calculated value. The new value is `name-cache-size` divided by 8 and rounded up to the next power of 2, or 8192, whichever is less.

On systems where the `nchsz` configuration variable in `/sys/conf/param.c` has been modified by the customer, the customer-provided value will be used. Compaq recommends that any previous customization of `nchsz` in the `param.c` file be changed to:

```
int   nchsz;
```

Compaq also recommends that you use the default dynamic calculation of the size of the `namei` cache hash change table.

If a customer specified value is still desired, Compaq recommends modifying the value of `name-cache-hash-size` in `/etc/sysconfigtab`, instead of modifying `nchsz` in `/sys/conf/param.c`. See `sysconfigdb`(8) for more information.

The default value of the AdvFS subsystem attribute `AdvfsMaxDevQLen` has changed from 80 to 24 (number of I/O requests). The new default of 24 reduces the number of I/O requests queued to a device. A lower value favors response time by reducing the number of I/O requests that may be already queued ahead of synchronous read and write requests.

A new AdvFS subsystem attribute, `AdvfsCacheHashSize` is added. `AdvfsCacheHashSize` determines the size of the hash table for the buffers of the AdvFS buffer cache. The minimum value is 1024. The maximum value is 65536.

The default value is the size of the AdvFS buffer cache (`AdvfsCacheMaxPercent`) divided by 10 and rounded up to the next power of 2, or 8192, whichever is less. To save a small amount of space in the kernel at the expense of longer hash queue search time, the `AdvfsCacheHashSize` can be set to a smaller value.

For more information about VFS and AdvFS tuning, refer to the *System Configuration and Tuning* guide.

1.11 Visual Threads

Visual Threads is a new tool available on the *Tru64 UNIX Associated Products Volume 2 CD–ROM* that lets you analyze your multithreaded applications for potential logic and performance problems. It is licensed as part of the Developer's Toolkit for Tru64 UNIX.

You can use Visual Threads with DECthreads applications that use POSIX threads (Pthreads) and with Java applications.

Visual Threads features include the following:

- Detects violation conditions based on the application of particular rules in your application. Several predefined rules look for data protection errors, deadlock conditions, programming errors, and performance issues.

- Lets you use templates to define your own rules to specify criteria for violation conditions.

- Keeps track of the events associated with the violations.

- Records events to a trace file so that you can play back and analyze them later.

- Dynamically displays events as they occur, with controls for filtering.

- Suspends execution of the application when it detects violation conditions. You can choose from several options at this point, including invoking the debugger in the appropriate context for your application.

- Dynamically monitors multithreading objects (threads, mutexes, and so on) in real time.

- Provides easy access to object-level statistics and current state information, including use of resources (for example, the mutexes a particular thread holds).

1.12 Bootable Tape Enhancements

This release adds new functionality to the bootable tape facility. For more information, refer to the *System Administration* guide.

1.13 Compaq Advanced Printing Software

This release includes Compaq Advanced Printing Software Version 1.0 for Tru64 UNIX. The software subsets and documentation are located on the *Tru64 UNIX Associated Products Volume 2 CD–ROM*.

Advanced Printing Software, an implementation of Xerox Corporation's PrintXchange technology, is a flexible, distributed, client/server printing system based on ISO and POSIX standards.

This implementation supports a wide variety of popular printer devices and the common printer connection methods. The software subsets include server processes, command line utilities, and graphical user interface (GUI) tools that cater to the needs of end users and system administrators. The system can be configured as the primary print system, or it can work together with existing lpr/lpd based printing systems.

Documentation includes a *System Administration and Operation Guide*, a *Command Reference Guide*, a *User Guide*, an *Installation Guide*, CDE Help, and *Release Notes*.

1.14 Netscape Communicator Version 4.5

This release of Tru64 UNIX contains Version 4.5 of Netscape Communicator. For information about how to install, configure, and invoke Communicator, see the Tru64 UNIX *Installation Guide*.

1.15 Compaq Insight Manager

This release provides Web-based Enterprise Management (WBEM) capabilities by integrating the Compaq Insight Manager Agents into Tru64 UNIX. Insight Manager enables Web-based browsing and monitoring of local and remote system hardware. You can activate these Web browsing features from a dedicated HTTP port, or from the Compaq Insight Manager CIM32 or CIMXE Management Consoles running on Compaq Windows NT servers.

Refer to Section 4.3.9 for more information.

1.16 Support for Universal Serial Bus Keyboards and Mice

Universal Serial Bus (USB) is a communications architecture that gives a computer system the ability to interconnect a variety of devices using a simple four-wire cable. The USB is actually a two-wire serial communication (receive and transmit) link that runs at either 1.5 or 12 megabits per second (mbs). The remaining two wires are used for power and ground. USB protocols can configure devices at startup or when they are plugged in at run time. These devices are broken into various device classes. Each device class defines the common behavior and protocols for devices that serve similar functions.

Tru64 UNIX Version 4.0F provides limited USB support for keyboard and mouse devices only. For proper usage of USB keyboards and mice with Version 4.0F, the following guidelines and restrictions apply:

- USB keyboards and mice must be attached to the system before the system is powered on.

- A non-USB keyboard and a USB keyboard cannot be used together on the same system. In the event that both types are attached, the non-USB keyboard has priority. The same rule applies for mice.

- If a USB keyboard is used, a USB mouse must also be used. It is not advisable to mix a non-USB keyboard with a USB mouse, as the system will be inoperable. The same rule applies for USB keyboards and non-USB mice.

1.17   AltaVista Search on Documentation CD-ROM

The Tru64 UNIX Documentation CD-ROM now includes AltaVista CD–ROM Search software, with a complete index of the HTML documentation set. The AltaVista software runs on an x86-based PC with Windows 95, Windows 98, or Windows NT Version 4.0 or higher.

1.18   Support for Digital Versatile Disc (DVD)

Tru64 UNIX Version 4.0F includes preliminary support for the native DVD file system, called DVDFS, together with documentation. This preliminary implementation is not fully functional. See Section 4.3.14 for a description of the restrictions.

1.19   Availability of COM for Tru64 UNIX

COM for Tru64 UNIX is available for downloading from Compaq.

COM for Tru64 UNIX implements the Microsoft Common Object Model (COM) as well as required underlying Windows capabilities on the Compaq Tru64 UNIX platform.

Using COM for Tru64 UNIX, you can:

- Port COM server applications from 32-bit UNIX and Windows NT platforms to the Tru64 UNIX platform where they can take advantage of the Tru64 UNIX speed and storage capacity.

- Create wrappers for legacy Tru64 UNIX applications, providing COM access to these applications by Windows clients.

- Develop new distributed UNIX applications that take advantage of COM re-use, version independence, and language independence capabilities.

COM for Tru64 UNIX will be incorporated in the Tru64 UNIX operating system kit in a future release. It is licensed as an extension of the Tru64

UNIX product; no additional license is required to use the COM for Tru64 UNIX software on a system running the Tru64 UNIX operating system. It is available via download at no charge from the following URL:

```
http://www.unix.digital.com/com
```

Installation Notes 2

This chapter discuss the following topics:

- Update installation
- Installation problems using console device EWE0
- Join database migration

Do not attempt to install Tru64 UNIX Version 4.0F without first reading the notes appropriate to your processor in Chapter 3. Failure to read these notes can result in installation problems. Also, before you start your installation, be sure to review the hardware documentation that came with your system.

2.1 Update Installation

To update your operating system software to Version 4.0F, you must use the `installupdate` utility or the full installation procedure as described in the *Installation Guide*.

Version 4.0F supports update installations from Version 4.0D and Version 4.0E. See Appendix B for more information about disk space requirements before executing the `installupdate` utility.

Note that the `-i` flag for the `installupdate` command will be retired in a future version of Tru64 UNIX. See Section 8.28 for more information.

2.1.1 Recovering from Boot Problems After Update Installation

During an installation update from Version 4.0D or Version 4.0E to Version 4.0F, a number of system configuration files on the system are updated with new information. The changes to these configuration files are incorporated into the existing files on the installed system as part of the update processing. The changes are merged into the existing files after the system is rebooted and running the new Version 4.0F generic vmunix kernel.

Depending on the hardware configuration and the content of the existing system configuration files, the new Version 4.0F generic vmunix kernel may fail to boot and run the final merge. This can occur when there are entries in the existing files on the installed system that are incompatible with the new kernel.

If the new Version 4.0F generic vmunix kernel fails to boot and complete the configuration of the updated system software, retry the boot using the `c` flag. At the SRM console prompt, issue the following command:

```
>>> boot -flag c
```

Do not supply any other arguments. The `c` flag causes the booted kernel to use default values instead of the values contained in the old configuration files. After the new Version 4.0F generic vmunix kernel boots and completes merging the updates to all of the system files, the system should boot the `/genvmunix` kernel successfully without using the `c` flag.

2.2 Installation Problems Using Console Device EWE0

On a system with more than four Tulip-style network devices, a user is limited to installation from the first four.

The install kernel is built knowing only of devices `tu0` through `tu3`. With recent network cards such as the DE504 (four network devices per card), it will become increasingly likely that a user will attempt installation from a higher numbered device. The boot will progress to the point where it attempts to mount the root device (console `ewn0`, where `n` is `e` or further in the alphabet), and fail.

The workaround is to install from console devices `ewa0`, `ewb0`, `ewc0`, or `ewd0`.

2.3 Join Database Migration

If your system provides Dynamic Host Configuration Protocol (DHCP), Remote Installation Services (RIS), or Dataless Management Services (DMS) services to to other systems, you must update the database files for the `join` daemon after you complete the installation. See Section 4.4.11 for additional information.

Processor-Specific Notes 3

This chapter contains notes that apply to the following computers:

- General notes on processors
- Personal Workstation 433au, 500au, and 600au systems
- AlphaServer 1000 and 1000A systems
- AlphaServer 8200, 8400, GS60 and GS140 systems
- Alpha VME 4/224, 4/288, and 5/nnn single board computers
- DMCC EBM2n class and EBM4n class single board computers
- DEC 7000 system
- AlphaPC 164 systems

3.1 General Notes on Processors

The following sections apply to more than one processor type.

3.1.1 Upgrading Your Hardware

You can follow the instructions in the *Installation Guide* and those provided by your hardware and firmware documentation when you add new options or change your system hardware. However, if the new option is supported only in the newest version of Tru64 UNIX, you must perform the upgrade in the following sequence:

1. Update your operating system software to the version that supports the new hardware or option.
2. Upgrade your firmware.
3. Upgrade your hardware or install the new option.
4. Follow the instructions in the *Installation Guide* for rebuilding your system kernel.

3.1.2 PCI-Based Systems with ATI Mach64 Graphics Controllers

On PCI machines that include ATI Mach64 graphics controllers in the configuration, you must set the console environment `pci_parity` variable to `off`. Due to a hardware limitation, hardware machine checks will occur if this variable is not turned off. Use the following commands:

```
>>> set pci_parity off
>>> init
```

3.1.3 PCI Shared Interrupt Support

In Version 3.2D and subsequent releases of the operating system, the PCI-based system and I/O support code is capable of shared interrupts. The only exception is the 53C810 (onboard or option) SCSI controller.

You can use the following console firmware command to determine if your system has a 53C810 SCSI controller:

```
>>> show config
```

To set up your PCI-based system to share interrupts, you must ensure the following:

- The firmware revision for your system supports shared interrupts.
- You have placed the option cards in PCI slots that allow sharing.

Refer to your layered product documentation (for example, for MME or Open3D) for information about the shared interrupt capabilities of the software.

3.1.4 KZPSA Behind the PCI-to-PCI Bridge

On AlphaServer 1000A and 2100A class systems, updating the firmware on a KZPSA SCSI adapter is not supported when the adapter is behind the PCI-to-PCI bridge. See your hardware installation guide for further information. A later version of the console firmware will support this feature.

3.1.5 Qlogic ISP1040B CAM Errors

On systems with a Qlogic ISP1040B option, CAM errors like the following may occur during bootup:

```
pci2000 at pci0 slot 8
isp0 at pci2000 slot 0
isp0: QLOGIC ISP1020A
cam_logger: CAM_ERROR packet
cam_logger: bus 0
isp_probe
NVRAM parameters invalid, using driver Fast10 defaults
```

To alleviate the error, you must use the `eeromcfg` utility to program the NVRAM with the proper set of parameters. The `eeromcfg` utility is provided in the `/mnt-pnt/utility` directory of the *Alpha Systems Firmware Update* CD-ROM. Consult the `readme.txt` file in that same directory for information about how to use the utility.

3.1.6 DJ-ML200-xx PCI NVRAM Hardware Rev Must Be E01

The revision of the ML200-xx 2/4/8MB PCI NVRAM adapter must be revision E01.

3.1.7 RIS boot Options

To start a RIS installation on an AlphaServer 1200, 4000, 4100, 8200, or 8400, you must no longer specify the `n` option with the `boot` console firmware command, as in the following example:

```
>>> boot -fl n ewa0
```

If you specify the `n` option, the RIS installation kernel will panic. You can safely initiate a RIS installation by omitting the `n` option:

```
>>> boot ewa0
```

3.2 Personal Workstation 433au, 500au, and 600au Systems

The following notes are specific to Compaq Personal Workstation class systems.

3.2.1 EIDE (ATAPI) CD–ROM Support

The Compaq Personal Workstation class systems contain two EIDE (ATAPI) ports.

Version 4.0F supports the use of ATAPI CD–ROM devices attached to the internal EIDE adapter of the Compaq Personal Workstation class systems. The CD–ROM devices appear as SCSI rz devices. The device names for devices attached to the first EIDE port are rz0 and rz1. The device names for the devices attached to the second EIDE port are rz8 and rz9. The EIDE master devices are rz0 and rz8. The EIDE slave devices are rz1 and rz9.

The device names are reserved even if no devices are connected to those ports at the present time. Disks connected to the SCSI buses on the system are numbered rz16 and above.

Any Tru64 UNIX application designed to use a SCSI CD–ROM drive can use a supported EIDE CD–ROM drive through this software interface.

3.2.2 64-Bit PCI Option Cards

The 64-bit PCI slots, slots 4 and 5, are intended only for those cards listed in the *Systems and Options Catalog* as supported for slots 4 and 5. The console prevents system operation and displays the following error if an unsupported card is present in these slots (*n*):

```
Illegal device detected on primary bus in physical slot n
Power down the system and remove the unsupported device from
slot n
```

3.2.3 Incorrect Default Keyboard Mappings

If you use a PCXLA-NA keyboard on a Compaq Personal Workstation 433au, 500au, or 600au class system, the keys will not map properly unless you reconfigure the keyboard driver to use the correct keymaps.

You can do this by executing the following command:

```
# sysconfig -r gpc_input kbd_scancode=2
```

If you prefer, you can use the sysconfigdb command to add the following entry to the /etc/sysconfigtab file:

```
gpc_input:
kbd_scancode = 2
```

Note that if you execute the sysconfig command to reconfigure the driver, you must do this every time you reboot the system. Using the sysconfigdb utility to make the change preserves the information across reboots, and no other user intervention is required.

3.3 AlphaServer 1000 and 1000A Systems

The following notes are specific to AlphaServer 1000 and 1000A systems.

3.3.1 EISA Configuration Utility Diskette Version 1.10

This note applies to users who utilize the onboard Cirrus VGA graphics controller.

The default setting for the VGA graphic controller when running the EISA Configuration Utility (ECU) Version 1.10 diskette is Disabled. For previous versions the default is Enabled.

When you run the ECU Version 1.10 for the first time on a system that was previously configured with an earlier version of the ECU diskette, the setting for the onboard VGA graphic controller is automatically set to Disabled. While running the ECU, select Step 3: View and edit details and set the VGA graphic controller to Enabled before exiting. If you do not set the VGA graphic controller to Enabled prior to booting Tru64 UNIX, your X

server will not start and you will use generic console support when you boot Tru64 UNIX.

3.3.2 Graphics Resolution

The default graphics resolution for Compaq AlphaServer 1000A systems containing built-in Cirrus video with 1 MB of video ram is 1024x768. If the optional 512 KB of video ram is not present, Tru64 UNIX supports resolutions of 640x480 (by default) or 800x600 only.

The default resolution for Compaq AlphaServer 1000 systems that contain built-in Cirrus video with 512 KB of video ram is 640x480. This configuration also supports 800x600 resolution.

To use 800x600 resolution, edit the following line in the `/usr/lib/X11/xdm/Xservers` file:

```
:0 local /usr/bin/X11/X
```

Change the line to:

```
:0 local /usr/bin/X11/X -screen0 800
```

To use 800x600 resolution for the CDE session manager, edit the following line in the `/usr/dt/config/Xservers` and `Xservers.conf` files:

```
:0 Local local@console /usr/bin/X11/X :0
```

Change the line to:

```
:0 Local local@console /usr/bin/X11/X :0 -screen0 800
```

Before editing these files for XDM or CDE, be sure that your system's monitor supports 800x600 resolution.

3.4 AlphaServer 8200, 8400, GS60 and GS140 Systems

The following note applies to the AlphaServer 8200, 8400, GS60, and GS140 systems.

3.4.1 Minimum Firmware Requirements

Before installing the Tru64 UNIX software kit, ensure that your system has the correct firmware version. The minimum firmware version required Compaq AlphaServer 8200, 8400, GS60, and GS140 systems is Version 5.4. If you have an earlier firmware version, update your firmware before installing the Version 4.0F software. For information on how to update your firmware, refer to the AlphaServer hardware documentation.

To determine the version of firmware on your system, enter the show

`version` console firmware command at the prompt:

```
P##>>> show version
```

Note that `##` is the primary processor number (usually 00).

3.5 Alpha VME 4/224, 4/288, and 5/nnn Single-Board Computers

The Compaq Alpha VME 4/nnn and 5/nnn series are Compaq's latest VME-based single board computers (SBCs).

Support for the VME VIP/VIC64 VME adapter on the Compaq Alpha VME 4/nnn and 5/nnn series SBCs is the same as the support for this adapter on Compaq AXPvme SBCs and Compaq Alpha VME 2100 systems. The following notes are specific to Compaq Alpha VME 4/nnn and 5/nnn SBCs.

3.5.1 Verify CPU Version

The `sizer` utility identifies Compaq Alpha VME 4/nnn and 5/nnn SBCs as Compaq Alpha 4/nnn or 5/nnn family members.

The `sizer -c` command displays the following output for Compaq Alpha VME 4/224 and 4/288 SBCs:

```
sysname> sizer -c
cpu      "DECALPHAVME_224"
```

The `sizer -c` command displays the following output for Compaq Alpha VME 5/nnn SBCs:

```
sysname> sizer -c
cpu      "DECALPHAVME_320"
```

3.5.2 Firmware Requirements

Before installing the Tru64 UNIX software kit, make sure that your system has the correct firmware version. The minimum firmware versions required are Version 1.2 or higher for a Compaq Alpha VME 4/224 or 4/288 SBC, and Version 1.0 or higher for a Compaq Alpha VME 5/nnn SBC. If you have an earlier firmware version, update your firmware before installing the Version 4.0F software. For information on how to update your firmware, refer to the firmware documentation.

To determine the version of firmware on your system, enter the following

command at the console firmware prompt:

```
>>> show version
```

3.5.3 Restrictions and Known Problems

The following restrictions apply to Compaq Alpha VME 4/nnn and 5/nnn SBCs.

3.5.3.1 Byte Addressability Not Supported by mmap (5/nnn Only)

In addition to sparse and dense addressing modes, the Compaq Alpha VME 5/nnn SBC hardware supports byte addressability. This includes the ability to perform simple math on a byte address to derive and access other byte addresses.

Currently, byte addressability is not supported by the Tru64 UNIX mmap interface.

VMEbus device drivers running on Compaq Alpha VME 5/nnn SBCs can continue to use the same sparse and dense space macros as in previous Tru64 UNIX releases and in existing VMEbus device drivers. Note that using the macros to access memory can help maintain the driver's binary compatibility (portability) across VMEbus platforms.

Byte addressability for Compaq Alpha VME 5/nnn SBCs will be supported in a future release of Tru64 UNIX.

3.5.3.2 VME Autovectors

The Compaq Alpha VME 4/nnn and 5/nnn SBCs do not support VME autovectors.

3.5.3.3 Network Port Termination

A Compaq Alpha VME 4/nnn or 5/nnn SBC that has the network configured in an UP state must have its external network connection properly terminated. If the network connection is unplugged or not properly terminated, then the network software will periodically time out and perform a network reset. This is normal for an unterminated Compaq Alpha VME system. However, it will cause high system latencies during the reset period, resulting in delays of about 10 milliseconds that can affect the realtime performance of the system.

Note that a loopback connector is not sufficient to terminate the network connection.

3.5.4 Writing VMEbus Device Drivers

For information about writing VMEbus device drivers and configuring the VMEbus on Compaq Alpha VME systems, refer to the `vme_manual_setup`(7) reference page.

3.5.5 CPU Boards Cannot Be Installed Until Operating System Is Installed

If you are not currently running Version 4.0E, and are upgrading CPU boards, then you must install Version 4.0F before you upgrade your machine with the new AlphaServer 8200/8400 (Compaq AlphaServer GS60,GS140) 6/525 CPU boards.

3.5.6 Futurebus Backplane Is Not Supported

The Futurebus backplane is not supported when an AlphaServer 8200/8400 is upgraded to the Compaq AlphaServer GS60 and GS140 6/525. This hardware must be removed from the system.

3.6 DMCC EBM2n and EBM4n Class Single Board Computers

The Modular Computing Components (DMCC) EBM2n-xx and EBM4n-xx are Compaq's latest PCI/ISA-based single board computers (SBCs).

The EBM2n-xx is a PICMG-compliant processor card based on the Compaq Alpha 21164 CPU. The EBM4n-xx is a PICMG-compliant processor card based on the Compaq Alpha 21064A CPU.

The following notes are specific to DMCC EBM2n class and EBM4n class single board computers.

3.6.1 Verify CPU Version

The `sizer` utility identifies Compaq DMCC EBM2n class and EBM4n class SBCs as DMCC EBM2n or EBM4n family members. The `sizer -c` command displays the following output for Compaq DMCC EBM2n class SBCs:

```
sysname> sizer -c
cpu      "DECEV56_PBP"
```

The `sizer -c` command displays the following output for Compaq DMCC

EBM4n class SBCs:

```
sysname> sizer -c
cpu      "DECEV45_PBP"
```

3.6.2 Firmware Requirements

Before installing the Tru64 UNIX software kit, make sure that your system has the correct firmware version. The minimum firmware versions required for Compaq DMCC EBM2n and EBM4n class SBCs is Version 4.7 or higher. If you have an earlier firmware version, update your firmware before installing the Version 4.0F software. For information on how to update your firmware, refer to the firmware documentation.

To determine the version of firmware on your system, enter the following console firmware command at the prompt:

```
>>> show version
```

3.6.3 Restrictions and Known Problems

The following restrictions and known problems apply to Compaq DMCC EBM2n class and EBM4n class SBCs.

3.6.3.1 Option Card Restrictions

You can use the DMCC EBM2n class and EBM4n SBCs on DMCC backplanes in the ETMXB-xx and ETMAB-xx family. The list of supported DMCC backplanes currently includes the following:

Part Number	Description of Backplane
ETMXB-BA	5-slot PICMG (2 PCI, 1 PCI/ISA, 1 ISA, 1 CPU)
ETMXB-AA	7-slot PICMG (2 PCI, 1 PCI/ISA, 3 ISA, 1 CPU) with onboard SCSI controller
ETMXB-DA	7-slot PICMG (3 PCI, 1 PCI/ISA, 1 ISA, 2 CPU)
ETMXB-CA	9-slot PICMG (5 PCI, 1 PCI/ISA, 1 ISA, 2 CPU)
ETMXB-AB/ETMAB-AB	14-slot PICMG (7 PCI, 6 ISA, 1 CPU)
ETMXB-BB/ETMAB-BB	14-slot PICMG (10 PCI, 3 ISA, 1 CPU)
ETMXB-AC/ETMAB-AC	19-slot PICMG (10 PCI, 7 ISA, 2 CPU)
ETMXB-BC/ETMAB-BC	19-slot PICMG (13 PCI, 4 ISA, 2 CPU)

All ETMXB-xx/ETMAB-xx backplanes, except the 5-slot ETMXB-BA and the 7-slot ETMXB-AA and ETMXB-DA, use PCI-to-PCI bridge (PPB) technology to provide both primary (in front of the bridge) and secondary

(behind the PPB) slots.

All ETMAB-xx backplanes are PCI Version 2.1 compliant.

Only the ETMXB-AA contains an onboard SCSI controller. The other ETMXB-xx/ETMAB-xx backplanes require a SCSI option card.

The option cards shown in Table 3-1 work behind (as well as in front of) the bridge. You can plug these cards in to any available slot.

Table 3-1: Supported Options Behind the Bridge

Option Type	Part Number	Description
Graphics	SN-PBXGB-AA[a]	TGA2 2MB PowerStorm 3D30
SCSI	KZPAA-AA	PCI-SCSI host bus adapter
SCSI	SN-KZPBA-CA	PCI-SCSI Ultra Wide adapter
Network	DE450-CA	PCI NIC (TP, TW, AUI)
Network	DE500-BA	PCI NIC (TP)

Table note:

a. The SN-PBXGB-AA (TGA2 PowerStorm 3D30) video card will work behind a bridge in multiple configurations if the first card is within the primary bus. For restrictions on jumper settings and X Server DMA for the PowerStorm 3D30 card, see Section 3.6.3.2.

When used with DMCC EBM2n SBCs (only), the SN-KZPBA-CA (PCI-SCSI Ultra Wide adapter) requires the following console parameter to be set:

```
 >>> set pci_prefetch SMS
```

3.6.3.2   PBXGB-AA (TGA2 PowerStorm 3D30) Video Card Restrictions

The following restrictions apply to the PBXGB-AA (TGA2 PowerStorm 3D30) video card that is listed in Table 3-1:

• EV5 Alias jumper setting (EBM2n only)

For EBM2n SBCs only, you must set the EV5 Alias jumper on the PowerStorm 3D30 card to On.

• VGAEN jumper settings

Only one PowerStorm 3D30 card can have its VGAEN jumper set to On. This card must be positioned in a primary PCI slot for the SRM Console to be displayed. All other PowerStorm 3D30 cards must have their VGAEN jumpers set to Off but may be positioned in any PCI slot.

For more information about the jumpers, see the *PBXGB-AA/CA PCI Graphics Option Owner's Guide* provided with the card.

- X Server DMA write operations must be disabled for some configurations

  Some configurations of PowerStorm 3D30 cards on Compaq EBM2n and EBM4n SBCs require that you disable X Server direct memory access (DMA) write operations. Specifically, you must disable these operations if there are multiple PowerStorm 3D30 cards in the system, or if any PowerStorm 3D30 graphics card is installed behind the PCI-PCI bridge (PPB).

3.6.3.3   General Restrictions

The Tru64 UNIX operating system does not support the operator control panel or watchdog timer. These server management features are supported only in the hardware and firmware.

3.6.4   Hardware Configuration for DMCC Eight-Headed Graphics Systems

This section describes how to configure a Modular Computing Components (DMCC) system to run eight-headed graphics under the Tru64 UNIX operating system.

You can configure DMCC platforms that contain a EBM2n-AZ Alpha PICMG single-board computer (SBC) and multiple PowerStorm 3D30 graphics cards to run multiheaded graphics, controlling up to eight monitors at a time.

3.6.4.1   Hardware and Software Requirements

Running eight-headed graphics on a DMCC system requires the following:

- An EBM2n-AZ Alpha PICMG SBC and eight PowerStorm 3D30 graphics cards within a fully configured DMCC system.

- A DMCC backplane and enclosure with at least 10 PCI slots, 512 MB main memory, a DMCC-qualified Ethernet card, and all the other storage and I/O options normally required for such a system. See the *DMCC Order Configuration Guide*, which is available on the Compaq OEM Web site at:
  `http://www.digital.com/oem/`

- The SBC occupies a CPU slot and the graphics cards must occupy eight PCI slots.

- You must be running Tru64 UNIX Version 4.0E or higher, and you must be running the latest DMCC SRM code from Version 5.2 or higher of the

Firmware CD-ROM.

The following DMCC configuration has been qualified by Compaq for running eight-headed graphics under Tru64 UNIX Version 4.0E:

- Alpha 21164/366 MHz SBC with 2 MB cache and Tru64 UNIX license (EBM21-AZ)

- 512 MB main memory (2 x EBXMA-HC, for a total of four 128MB DIMMs)

- Desktop enclosure with 14-slot PICMG backplane: 10 PCI, 3 ISA, 1 SBC (ETM33-BD)

- Eight PowerStorm 3D30 graphics cards (8 x SN-PBXGB-AA)

- PCI Ethernet card (DE450-CA)

- PCI fast/narrow SCSI controller (KZPAA-AA)

- Mandatory or associated options such as floppy, hard, and CD–ROM drives, cable kit for PICMG enclosure, and power cord

- DMCC SRM code from the Version 5.2 Firmware CD–ROM

3.6.4.2 Hardware Setup

When you configure the DMCC 15-slot platform for eight-headed graphics, placement of the graphics cards is critical.

The configuration described in the Hardware and Software Requirements section, and qualified by Compaq, uses an ETM33-BD desktop enclosure with 14-slot backplane. Within that enclosure, the PCI option cards must be placed into PCI slots in top to bottom order as follows:

- Secondary 32-bit PCI bus connectors

    – KZPAA SCSI card

    – PowerStorm graphics card: SCREEN 2

    – PowerStorm graphics card: SCREEN 3

    – PowerStorm graphics card: SCREEN 4

    – DE450 Ethernet card

    – PowerStorm graphics card: SCREEN 5

    – PowerStorm graphics card: SCREEN 6

    – PowerStorm graphics card: SCREEN 7

- Primary 64-bit PCI bus connectors

    – PowerStorm graphics card: SCREEN 0 (VGA ENABLED)

–   PowerStorm graphics card: SCREEN 1

For reference, the power connector is situated above the PCI slots, and the CPU and ISA connectors below.

All PowerStorm cards must have the ALIAS jumper in and VGA EN jumper out, except the SCREEN 0 card, which must be VGA enabled.

3.6.4.3   Software Setup

After you complete the hardware configuration for the eight-headed system, you can set up the Tru64 UNIX software to operate the eight screens as one row of eight screens (8x1) or two rows of four screens (4x2).

By default in a multiheaded configuration, the screens are operated as 8x1. To set up the screens in a 4x2 combination, you must edit your system's X Window System server configuration file, `/usr/var/X11/Xserver.conf`.

Instructions for editing this file in order to customize the X server configuration are provided in the `Xserver`(1X) reference page.

To set up 4x2 operation, you need to specify `-edge_top`, `-edge_bottom`, `-edge_right`, and `-edge_left` command-line arguments that arrange and attach the screens as you want them. Each argument takes *scr1* and *scr2* values, which are the numbers of the screens you are attaching.

For example, you could arrange the eight screens as follows:

```
/--------------\
| 4 | 5 | 6 | 7 |
|--------------|
| 0 | 1 | 2 | 3 |
\--------------/
```
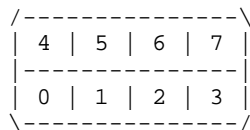
To achieve this combination, add the following information to the command-line arguments section of `Xserver.conf`:

```
! you specify command line arguments here
args <
  -pn
  -edge_top0 4    -edge_top1 5    -edge_top2 6     -edge_top3 7
  -edge_bottom4 0 -edge_bottom5 1 -edge_bottom6 2 -edge_bottom7 3
  -edge_right0 1  -edge_right1 2  -edge_right2 3
  -edge_right4 5  -edge_right5 6  -edge_right6 7
  -edge_left1 0   -edge_left2 1   -edge_left3 2
  -edge_left5 4   -edge_left6 5   -edge_left7 6
>
```

3.7 AXPvme Single Board Computers

For information on how to configure AXPvme single board computers, see the *Writing VMEbus Device Drivers* manual.

The following notes are specific to the AXPvme SBCs.

3.7.1 Firmware Upgrade Required

Compaq AXPvme SBCs must upgrade to Version 17.0 or higher of the AXPvme firmware to run Tru64 UNIX Version 4.0F.

3.7.2 Unpredictable Results When Performing Master Block Transfers

The following restriction applies to the VIP/VIC adapter used on Compaq AXPvme SBCs and Compaq Alpha VME 2100 systems.

Performing master block transfers with a data width of D64 can produce unpredictable results in the following cases:

- If D64 slave access is performed before memory has been mapped to the VMEbus.

- If memory access does not coincide with the appropriate access mode, such as attempting user access to memory specified as supervisory mode access.

- If the AXPvme SBC or Alpha VME 2100 system is a VME interrupter and is targeted for D64 slave access. The interrupt vector presented by the VME interrupter may not be the vector specified in the vba_post_irq function.

Memory must be mapped to the VMEbus prior to D64 slave access.

Access to memory must coincide with the appropriate access mode. If supervisory mode access is specified when memory is mapped, memory accesses must use supervisory mode. If user mode access is specified, both supervisory and user access are allowed.

See the *Writing VMEbus Device Drivers* manual for more information on slave and master block transfers.

3.8 DEC 7000 System

The following note is specific to the DEC 7000 system.

3.8.1 Misleading Error Message Displayed During Installation

During installation of Tru64 UNIX Version 4.0F software on a DEC 7000 system, the following message might appear on the console screen:

```
prom_saveenv: console returned illegal value
```

You should ignore this message and follow the instructions displayed on the screen by the installation process.

3.9 AlphaPC 164 Systems

During an initial installation of Tru64 UNIX Version 4.0F software on some AlphaPC 164 family systems, the installation procedure might hang while setting the console environment variables prior to rebooting the system with the new software. The problem does not occur during an update installation because console environment variables are not reset.

This problem is due to a limitation in the console firmware for the systems on which it is observed. The use of the `consvar` command, or attempts to modify parameters in the `bparm` subsystem through use of the `sysconfig -r` command also cause these systems to hang. To avoid this problem, change console environment variables directly at the SRM console prompt instead of through the privileged kernel interfaces.

If your AlphaPC 164 family system hangs during initial installation, you should power cycle the system and then set the console environment variables to reflect the correct boot device and boot mode. When you reboot the system, the installation procedure will continue with system configuration.

For example, if the system software is installed on the disk connected as unit 0 on SCSI bus 0 (that is, on rz0), then enter the following commands when the system powers up to the SRM console prompt (>>>):

```
>>> set bootdef_dev dka0
>>> set boot_osflags A
>>> set auto_action boot
```

The next time the system powers up, it will automatically boot from the freshly-installed disk and prepare to come up in multiuser mode. You can boot immediately from the console prompt by entering the `boot` command, or you can power cycle the system.

Base System Software Notes 4

_____

This chapter contains notes about issues and known problems with the base operating system and, whenever possible, provides solutions or workarounds to those problems.

The following topics are discussed:

- Commands and utilities
- SysMan System Management graphical user interface
- System administration
- Network and communications
- Local Area Transport
- File systems

4.1 Commands and Utilities

The following notes apply to commands and utilities.

4.1.1 Escaped Comment Symbols in a Makefile

The `make` command will not recognize escaped comment symbols as literal characters in a Makefile. Comment lines that begin with a # (number sign) and all text following this symbol up to the end of the line are considered part of a comment. This is also true even if the symbol is preceded with a backslash (\).

4.1.2 Security

The following notes describe problems that may occur when using commands and utilities under certain security settings.

4.1.2.1 Unexpected Command Behavior with ACLs

Programs cannot reliably inspect the permission bits in the `stat` structure and determine the access that will be granted to a particular user. On local file systems, read-only mounts and Access Control Lists (ACLs) can both modify the access that will be allowed. On remote file systems, in addition

to read-only mounts and ACLs, there may be additional controls that can alter the permitted access such as:

- ID mapping
- Mandatory access control
- Additional authentication requirements

Programs that copy files to update them, rather than updating them in place, often do not preserve ACLs. Some programs that have this problem are `gzip`, `compress`, and `emacs`.

The best solution for programs that need to make access decisions is for the program to use the `access()` call to determine what access will be granted. Note that even this may not work as the access protections of the file could be changed between the `access()` call and the `read`, `write`, or `execute` operation.

For programs that copy files, the following command will copy a file while preserving ACLs and any other extended attribute (property list):

```
# cp -p
```

See the `acl(4)` and `proplist(4)` reference pages for more information.

4.1.2.2 ACLs and Archive Tools

The `pax`, `tar`, and `cpio` archive tools might not restore ACLs on files in the manner you think that they should be restored. Always check the ACLs on your files after saving and restoring them with any of these tools.

4.1.3 Emacs Can Lose ACL File Settings

By default, the Emacs editor will rename the original file and save the new file as a copy under the original name. If the original file had an Access Control List (ACL) it will now apply to the backup file. If the directory had a default ACL, the new file (original file name) will now have the default ACL instead of the original ACL. If the directory did not have a default ACL, the new file will be protected only by the file permission bits.

The Emacs editor has some user-preference variables that you can set to control which file will retain the original ACL. The relevant Emacs variables are:

- `backup-by-copying`
- `backup-by-copying-when-mismatch`
- `backup-by-copying-when-linked`

4.1.4 Some Emacs Command Line Options Fail

Compaq ships the Emacs software as it is received from the source. The
following command line options do not work as documented in the
`emacs(1)` reference page: `-cr`, `-geometry`, `-i`, `-ib`, `-iconic`,
`-iconname`, `-in`, `-internal borderwidth`, `-mc`, `-T`, and `-title`.

In some cases, a workaround is available by using an appropriate X resource.

4.1.5 New Tunable Parameter for Running Oracle8

If Oracle8 running on Tru64 UNIX hangs, look for the following console
message:

```
malloc_wait:X : no space in map.
```

If this message is present add the following to `/etc/sysconfigtab`:

```
generic:
kmem-percent=XXX
```

In this example, `XXX` is the size of the kernel malloc map. By default, `XXX`
is set to 25 percent of physical memory. Increase this parameter to 50 and
reboot the system. If Oracle8 continues to hang, increase this parameter up to
100.

4.1.6 Change to Format of ipcs Command

The output format of the `ipcs` command has been modified. The `KEY` field
now presents the value in hexadecimal format instead of decimal. This
change is compatible with other UNIX implementations and conforms to the
upcoming X/Open UNIX 98 standard.

4.1.7 Sendmail

The following notes apply to `sendmail`.

4.1.7.1 New sendmail Version 8.8.8

The `sendmail` program has been upgraded to Version 8.8.8 as the default
version for Tru64 UNIX. The current version, Version 5.65, is still available
for use.

The new version, known as V8, has become the de facto standard in the
industry, especially for ISPs. Some of the features and enhancements
provided with this new version include masquerading, virtual domains, relay
control, and spam blocking. For more information, see the sendmail book by
O'Reilly and Associates, or the *sendmail Installation and Operations Guide,*
or the `sendmail(8)` and `sendmail.cf(4)` reference pages. The *sendmail
Installation and Operations Guide* is included in PDF format on the

Supplementary Documents bookshelf of the online Tru64 UNIX documentation.

You can configure `sendmail` using the `mailconfig` or `mailsetup` applications. During an update installation, the system will attempt to update your existing `sendmail.cf` file to comply with this new version. This updated `sendmail.cf` is compatible with both new and old version of `sendmail`.

To continue using the previous version of `sendmail`, invoke the following command as root:

```
/sbin/init.d/sendmail select old
```

This will change all links related to `sendmail` to point to the version previous to Version 8.0. Once you are ready to use the new version, you can switch back by invoking the following command:

```
/sbin/init.d/sendmail select v8.8.8
```

4.1.7.2 Initial sendmail Configuration

When the system is first booted, after a full installation, the `sendmail` software is automatically configured. This takes a noticable amount of time, and no status message is displayed until the following warning message is displayed as a result of starting `sendmail`:

```
warning: local host name(hostname) is not qualified;
fix $j in config file.
```

This means the system does not have a qualified name. This is because neither `bind` nor `mail` has been configured. The `sendmail` program will continue to operate.

Ignore this warning.

4.1.7.3 Warning Message When sendmail Updates the alias Database

The `sendmail` program gives the following warning message when it is started and its alias database is updated:

```
WARNING: writable directory /var/adm/sendmail
```

When `sendmail` initializes the alias database, it checks the modes on system directories and files to determine if they have been compromised.

Because the permission on the `/var` directory is 775, the `sendmail` program gives a warning message because it expects the permission on `/var`

to be set as follows:

```
chmod go-w /var
```

Ignore this warning.

4.1.7.4 Mail to user@localhost Fails

When a user sends mail to `<user>@localhost`, the `sendmail` utility
resolves it as `<user>@localhost.<domain name>`. As a result, if
you do not have a system `localhost` in your domain, your mail will
bounce back.

Instead of sending mail to `<user>@localhost`, send it to `<user>` only.

4.1.8 The tar Command h Option Does Not Work As Documented

The `h` option to the `tar` command, which expands symbolic links to real
files and directories in the archive, does not work as documented. Do not use
the `h` option. This problem will be fixed in a future release.

If you must use the `h` option, use the version of the `tar` command installed
in `/usr/opt/obsolete` from the `OSFOBSOLETE440` optional subset.

4.1.9 Netscape Communicator

The following notes apply to the Netscape Communicator.

4.1.9.1 Netscape Communicator Dumps Core Running in CDE

Netscape Communicator dumps core when the application posts a file
selection dialog (`XmFileSelectionBox`). Typically, this occurs when
you are running the application in the Common Desktop Environment (CDE)
and select the Save As option in the File pulldown menu of the Navigator
browser. It can also occur when you select a link to download a file or save
an attachment to a mail message in the Messenger Mailbox component.

To avoid this problem, invoke Netscape using the
`/usr/bin/X11/netscape` script.

As long as this script is used to start Netscape Communicator, the application
will display the file selection dialog within CDE without core dumping. Use
the `-xrm '*nsMotifFSBCdeMode: True'` command-line option if
you are starting Netscape Communicator using some other means.

For more information, see the Communicator on UNIX release notes at the

Netscape Web site:

```
http://home.netscape.com/eng/mozilla/4.0/relnotes/unix-4.0.html
```

4.1.9.2 Netscape Communicator Core Dumps If mailto: Link Is Selected In a Web Page

Clicking on a `mailto:` HTML link in a web page displayed with Netscape Communicator's Navigator web browser causes Communicator to core dump. This is a known problem scheduled to be fixed in a future release. The workaround is to start the Mailbox Messenger component from the Messenger option in the Communicator pulldown menu before selecting the `mailto:` link.

4.1.9.3 Deleting Multiple Mail Messages Causes Netscape Communicator to Core Dump

Deleting multiple mail messages in Netscape Communicator's Messenger Mailbox component sometimes causes Communicator Version 4.5 to core dump. Usually, it requires several multiple deletions of mail to make Communicator core dump. If Communicator does not core dump immediately, deleted messages might reappear in the mail folder from which they were deleted.

4.1.9.4 Netscape Communicator Dumps Core Intermittently

Netscape Communicator intermittently dumps core and returns the following error in the terminal window from which it is started:

```
Memory Fault - (core dumped)
```

This core dump occurs with different hardware and software configurations and under different circumstances. Sometimes it hangs for a time, taking most of the CPU time, then it crashes. At other times, its process has to be killed and the application restarted. Numerous problems of this nature have been reported. None are resolved at this time and no workaround is available. In all cases, the behavior cannot be reproduced consistently.

4.1.9.5 Cannot Delete Mail Messages From Inbox to Trash When Using IMAP Server

After upgrading to Netscape Communicator Version 4.5 from a previous version of Communicator, an IMAP mail user cannot move messages to the Trash folder in the Messenger component. All Delete options in the user interface are insensitive (greyed-out). Setting the Move it to trash folder option in the IMAP mail server preferences window does not work. This behavior is the result of a new feature in Netscape Communicator Version 4.5 that might require user customization after upgrading to this version.

Starting with Netscape Communicator Version 4.5, the Namespace extension to the standard IMAP protocol is used to locate the users folders on the IMAP mail server. This feature will not work if you are using an older IMAP server that does not support the Namespace extension to the protocol. Use the following procedure to customize Netscape Communicator to be able to locate a user's Trash folder on an old IMAP server:

1. Select the Preferences option in the Edit pulldown menu and choose the Mail and News Servers option in the Preferences window.

2. In the options displayed below this option, select the Mail Servers option.

3. Find the IMAP server in the scrolled list of servers. Select it and choose the Edit button to edit the server configuration.

4. In the pop-up dialog box that is displayed, choose the Advanced tab.

5. Ensure that the Namespace field in the tab reads (quotes and period included) as follows:

   ```
   Namespace: "INBOX."
   ```

6. Click on the OK button in the pop-up window and again in the Preferences window to save the settings.

7. Exit and restart Communicator.

You should be able to move messages to the Trash folder and all the Delete options should now be sensitive (dark letters). Because IMAP mail server configurations differ (including the location of the user's folders on the server), check with your IMAP mail server administrator if the above procedure fails to resolve the problem.

4.1.9.6   Communicator Returns sh: /usr/bin/X11/showps: not found

When you select a link to a PostScript file in the Navigator component of Communicator, it might return the following error:

```
sh: /usr/bin/X11/showps: not found
```

The `showps` helper application has been retired from Tru64 UNIX as a result of licensing changes to Adobe Display PostScript. The user may have customized the PostScript Document MIME type to use the showps helper application in `$HOME/.mailcap` and `$HOME/.mime.types` files.

To resolve this problem, you must obtain a new PostScript viewer and reconfigure the helper application for the PostScript Document MIME type in Communicator. Use the Edit option in the Edit-Preferences-Navigator-Applications pulldown menu of Communicator to edit your PostScript Document helper application and replace `/usr/bin/X11/showps` with the path to your new PostScript viewer.

4.1.9.7 Installing Open3D Causes Netscape Communicator to Core Dump

Netscape Communicator crashes when scrolling Web pages in Navigator or when editing bookmarks on a system with Open3D installed and a PowerStorm 4D5T graphics option after upgrading from Version 4.0D or later of the operating system. In general, Communicator and some other applications become unstable on this system configuration. The workaround is to set shells DISPLAY variable to local:0 before starting Netscape Communicator. Another option is to display Netscape Communicator remotely from another system without Open3D installed on it.

4.1.9.8 Netscape Communicator Version 4.5 Form Handling Problem in Asian Locales

Netscape Communicator Version 4.5 has a form handling problem when running under the Asian locales, such as Japanese or Chinese. The input data field is cleared after the data is submitted. Web applications that use the input field value might not work properly. A workaround is to start Communicator in C locale, for example:

```
% setenv LANG C
% netscape
```

4.1.10 Do Not Use Current Directory Character with setld

Do not specify the current directory "." character as the location argument for setld, as shown in in the following example:

```
% cd /mnt/ALPHA/BASE
% setld -l .
```

Instead, use the directory path as the location argument, for example:

```
% setld -l /mnt/ALPHA/BASE
```

4.2 SysMan System Management Graphical User Interface

The following sections apply to restrictions on using the SysMan applications.

4.2.1 Account Manager

The following notes apply to Account Manager, dxaccounts.

4.2.1.1 Copying Accounts and UID

When copying user accounts via cut and paste or drag and drop, the Allow Duplicate UIDs option in the General Preferences dialog box will be honored. For example, when making a copy of user account that has a UID of 200, if the Allow Duplicate UIDs check box is off (the default), the resulting copy will have a unique UID automatically generated. If the Allow Duplicate UIDs check box is on, then the copy will have an identical UID. The same rules apply to copying groups.

4.2.1.2 Account Manager Restrictions

The Account Manager has the following restrictions on both base security and enhanced security (C2) systems:

- Using mouse button 1 (MB1) to drag and drop users, groups, or templates does a copy operation, not a move operation. This is different from the CDE/Motif default where MB1 performs a drag and drop move operation and Shift-MB1 is required to perform a copy operation. For example, using MB1 to drag a user from the Local Users view and dropping it in the NIS Users view creates a copy of that user in NIS.

  Workaround: Delete the original icon after the copy has been completed.

- The Account Manager allows you to specify the minimum and maximum range for a UID or GID. However you cannot specify the starting value in the range.

  Workaround: Use the `usermod` or `groupmod` commands to set a starting value within the range:

  ```
  usermod -D -x next_uid=xxx
  usermod -D -x next_gid=xxx
  ```

  Suppose the minimum UID is 100 and the maximum UID is 10000. Then the following line would cause the Account Manager to start generating UIDs from 5000:

  ```
  usermod -D -x next_uid=5000
  ```

- If you change a user's UID with Account Manager, the ownership of the user's files and subdirectories do not change, and under certain circumstances, the home directory ownership may not change either. For instance, if you attempt to change the UID of user `johndoe` from 200 to 201, the files and subdirectories under his home directory will still belong to UID 200. Furthermore, if `johndoe` does not own his home directory, the ownership of that directory will not change either.

  Workaround: Use the `chown` command to change the directory and files, if applicable.

- You cannot drag and drop items across different instances of Account Manager. For example, if Account Manager A on system 1 and Account Manager B on system 2 are both being displayed on the same workstation, then you cannot drag and drop between Account Manager A and B.

  Workaround: Use the copy/paste feature to copy users, groups, or templates from Account Manager A to B.

- Restriction: Two system administrators should not run two different concurrent instances of Account Manager.

  Workarounds: Account manager correctly allows two or more system administrators to work on the same password files simultaneously. The proper file locking will occur and new accounts can be added or modified. However, the local groups file, `/etc/group`, and the NIS groups file, `/var/yp/src/group`, are written out after each group modification. Therefore, the last system administrator to make a change in a group's view window would overwrite any prior changes from a different system administrator. For this reason, running multiple, concurrent Account Manager instances is not recommended.

- When running Account Manager from a terminal window, occasionally the following message might be written to stdout:

  ```
  Warning: DtComboBoxWidget: Unable to find item to select
  ```

  Workaround: None. You can safely ignore these messages.

4.2.1.3 Account Manager Problems

Leading and trailing white space is not stripped from text entry areas. This could lead to confusion, for example, if a field on the Find dialog box contains a space character before the desired search string. The search string would not match because of the spurious space character.

4.2.1.4 Enhanced Security Account Manager Problems

The following problems apply to Account Manager when running on enhanced security systems:

- The Lock/Unlock Toolbar and Menu Options are inactive for the Template views.

  Workaround: Change the template lock setting on the Create/Modify Template dialog screen after selecting the template by double clicking on the template icon in the Template view icon box.

- The C1Crypt Encryption Type restricts the password length to between 6 and 8 characters even though the Password Controls' Minimum Length and Maximum Length fields imply otherwise.

Workaround: Set passwords through `/usr/tcb/bin/dxchpwd` or the `/usr/bin/passwd` command when the C1Crypt Encryption type is chosen.

- Do not set a template Encryption Type to C1Crypt as this will invalidate the template.

  Workaround: Set the C1Crypt Encryption type for the user from the Create/Modify User dialog.

- Account Manager does not enforce the minimum/maximum password length limitation when setting passwords.

  Workaround: Set passwords through `/usr/tcb/bin/dxchpwd` or the `/usr/bin/passwd` command if the minimum/maximum password length limitation is necessary.

- The Pointer Focus Prompt message in the status line of the Icon Box will display Delete instead of Retire when the mouse pointer is in the Retire toolbar icon.

- Error messages generated from the Create/Modify Template dialog box refer to the user name when they should refer to the template name.

- On enhanced security system, you typically retire users instead of deleting them. However, there are times when you might want to delete a user account. Account Manager supports retiring user accounts but not deleting them.

  Workaround: To delete a user account you must do the following:

  1. Manually edit the `/etc/passwd` and `/etc/group` files to remove references to the user.

  2. Use the following command to remove the user from the protected password database:

     # **/usr/tcb/bin/edauth -r <user name>**

- Renaming a user by changing the Username field of the Create/Modify User dialog box in Modify mode does not clean up the protected password database entry for the old name.

  Workaround: Use the following command to remove the dangling protected password database entry:

  # **/usr/tcb/bin/edauth -r <user name>**

- Do not rename a template by changing the Template name field of the Create/Modify Template dialog box in Modify mode. Account Manager actually creates a new template without removing the old template. However, the old template's icon is removed from the Icon Box.

  Workaround: Restart Account Manager to restore the former template icon. Use the Delete Toolbar icon or the Edit->Delete... option from the

Template view to delete the undesired template

- Accounts and templates inherit their settings either from locally defined values in their protected password database entry or from the templates that they may reference. All accounts and templates implicitly reference a default template. The default template is not served by NIS. This creates an inconsistency for Account Manager when displaying NIS user accounts and templates on an NIS master. The user and template values displayed may be the default template values of the NIS master. When an NIS user logs into a NIS client, the NIS client's default template might be different from the NIS master's default template. The client's default template is used to establish the user's account settings.

- Using drag and drop to copy a user account copies the user's template references by value when the user account is being dropped on a different view. This means that the template itself will no longer be referenced by the newly created account. Instead, the template's values will be contained directly in the new user's protected password database entry. For example, assume the local user Joe has an account based on the developers template. If you drag and drop Joe's account from the Developers view into the NIS Users view, the attributes from the developers template will be placed into the protected password database entry for Joe's account. This preserves Joe's developer attributes and overrides any corresponding attributes from the default template for NIS users.

    Workaround: Modify the copied user's account and change the template from the default to the desired template. Note that the template reference is maintained if the user is dropped within the same view.

- Dropping a template icon on a user changes the user's account to use that template. However, the template's lock attribute is not honored. For example, if the template developer has the lock field enabled, then dropping this template on a user should cause the account to be locked but it does not.

    Workaround: Only the drag and drop method of template assignment has this problem. You can use the Create/Modify dialog box to change a single user's template or use the Modify Selected dialog box to change templates for several selected users. Both methods correctly propagate the template's lock field.

- Dragging and dropping a template onto a user's account displays a confirmation message in the view's status line. This message incorrectly displays `template %2` instead of the template's name.

    Workaround: None.

- Deleting a newly created template (the template was created after starting Account Manager) will cause the application to crash.

Workaround: Restart Account Manager and then delete the template.

- After deleting a template, the NIS maps are not remade.

  Workaround: Manually remake the NIS maps or perform an Account Manager function (for example, Account Modification) that will trigger the maps to be remade. To manually remake the maps do the following:

  ```
  # cd /var/yp
  # make all
  ```

4.2.2   BIND Configuration

The following problems apply to configuring BIND servers with the BIND configuration graphical user interface.

- Once entered, a zone cannot be deleted by clicking on the delete button in the Zones window.

  Workaround: Edit the /etc/namedb/named.boot file to remove the zone entry that needs to be deleted.

- At least one Forwarder must be entered in the Forwarder window if any Zones are to be added.

  Workaround: If the Forwarder is not needed, you can edit the /etc/namedb/named.boot file to remove the Forwarder entry after you Commit to the setup in the BIND configuration interface.

- If your host is already configured as BIND server, and the BIND configuration interface is used to add, delete, or append forwarders, all of the existing Zones in /etc/namedb/named.boot file will be deleted if the Zones button is not clicked first.

  Workaround: Whenever you modify your BIND setup using the BIND configuration interface, click on the Zones button before you Commit to any changes.

- The Server Addr field of secondary domain in the Modify Zone window cannot be modified using the BIND configuration interface.

  Workaround: The address can be modified by editing the secondary domain entry in the /etc/namedb/named.boot file.

4.2.3   The dxshutdown Application Does Not Prohibit Logins

The dxshutdown application does not create the /etc/nologin file as described in the documentation. This means that users will be able to log in to a machine that is being shut down up until the actual time of the shut down.

Note that this behavior differs from that of the shutdown command that creates the /etc/nologin file at 5 minutes prior to the shutdown.

4.3 System Administration

The following notes apply to system administration.

4.3.1   License Management Facility Allows Year 2000 Testing

Certain software license PAKs include expiration dates that currently limit the ability to run software when the date is set well into the future (into the year 2000). While most customers do not have PAKs with expiration dates, those who do (such as CSLG, ASAP, or Partner PAK Program members) may benefit from the following enhancements.

The lmf utility and supporting code have been enhanced to allow customers with expiring software license PAKs to set the system date beyond the expiration dates within a specific time window to allow Year 2000 (Y2K) testing.

Also, a test PAK, 00Y2K-TESTING, is available for use with Version 4.0D and higher of the operating system. If you have expiring license PAKs, you can install this test PAK, allowing your PAKs to operate beyond their expiration dates between the dates of December 1, 1999 and March 2, 2000. The test PAK allows you to conduct Y2K testing within the specified time window.

For information about obtaining the 00Y2K-TESTING PAK and other Y2K issues, see the *Tru64 UNIX Year 2000 Readiness* white paper, available in HTML format on the *Tru64 UNIX 4.0F Documentation, Volume 1* CD–ROM. This document is also available on the Tru64 UNIX web page at the following URL:

```
http://www.unix.digital.com/unix/year2000/whitepaper.html
```

Once you obtain and install the 00Y2K-TESTING PAK and set the system date forward for Y2K testing, you must execute the following command from the root account the first time run level 3 (also referred to as "init 3" or "multiuser mode") is entered after each system boot:

```
# /usr/sbin/lmf reset
```

This will ensure that all software licenses are loaded while the system date is set within the Y2K testing window. You need to execute this command only once within run level 3 between each system boot. You can toggle the system between run level 1 and run level 3 after the first execution of the command in run level 3. You will only be required to execute the command again after rebooting the system.

4.3.2   EISA Configuration Utility Revision Requirements

For Tru64 UNIX Version 4.0F and its software supplements, the supported version of the EISA Configuration Utility (ECU) is Version 1.10 or higher. If your system is configured with an EISA bus, you should update the ECU to this supported version.

4.3.3   Open3D Support

Consult the *Open3D Software Product Description* (SPD) before installing Open3D to ensure that this layered product is supported on your system.

Installing Open3D on systems not supported by the Open3D layered product can leave your system in an unusable state.

4.3.4   Opening a Write-Protected Tape in Write Mode

The behavior of the `open` call to a tape device has changed. You can no longer use `write` mode to open a write protected tape. The attempt to open the tape will fail, returning the following message:

```
EACCES (permission denied).
```

If an application is written so that it attempts to open the tape device with `O_RDWR` when the intention is only to read the tape, the open attempt will fail. Applications should be changed to open the device with `O_RDONLY`. For applications that cannot be changed, use the following command to obtain the previous behaviour of the open call:

```
# sysconfig -r cam_tape open_behaviour=0
```

4.3.5   Restrictions for a Dataless Environment

The Tru64 UNIX Server Extensions includes Dataless Management Services (DMS) support for installing and operating systems in a dataless configuration. A server system maintains the `root`, `/usr`, and `/var` file systems for all client systems. The server maintains one copy of root for each client. The `/usr` file system is exported read-only and is shared by all clients registered to the environment. Each client has their own `/var` file system. Dataless clients access the file systems maintained on the server utilizing NFS.

A dataless environment should be considered in the following scenarios:

• There is a need to reduce the disk space requirements on the client systems. Local disk space is required only for swap and dump, the `root`, `usr` and `var` file systems are resident on the server.

• There is a need to have all software and data centrally located on the server. A dataless client is not designed to support local access to data.

You should consider the following limitations when deciding to implement a dataless environment:

- Hardware limitation of the network – the majority of file system usage performed by any of the clients is done via NFS. This will consume network bandwidth, which, depending on the type and configuration of the network, can become exhausted as more clients are added to the environment.

- Hardware limitation on the file systems contained on the server – all clients are configured to share a /usr file system and all client specific data is contained in a /clients file system on the server. As clients are added to the dataless environment, the demands on these file systems on the server increase and may surpass the capabilities of the server.

- DMS may not be the best way to address centralized system management requirements. The use of DMS introduces added system management tasks in the server, and requires system setup and management on each client system.

- Patch management and update installation are not fully supported in a dataless environment. An installed dataless environment cannot be modified after clients have been configured. All clients must be deconfigured before any patches are applied or other changes made. When a client is deconfigured, the client specific customizations that have been made in the client's writable file systems, such as system configuration data, will be lost.

4.3.6 Alternate Root Installation May Change Host File Dates

During an alternate root installation of base operating system subsets, such as is done using the dmu utility to set up a Dataless Management Services environment, the file access dates on some of the files in the host server's file system may be changed to correspond to those from the subset's file inventory. When the release installed into the alternate root is different from that installed on the host system, these changed dates appear invalid since they may be newer (or older) than the actual file dates from the host system's installation kit.

This occurs when the pax utility is invoked by the setld utility to copy symbolic links from the kit subsets, and the symbolic links target absolute paths that correspond to actual files in the host system's file system. The pax utility attempts to adjust the dates for the symbolic link, but the file system actually adjusts the dates for the target of the symbolic link.

The changed dates have no operational impact on the host system. The content of the affected files has not been changed. However, because the dates have changed, the behavior of utilities that examine file dates (such as the find command or archivers) may be affected.

4.3.7   Enhanced Security

The following notes apply to the use of enhanced security features.

4.3.7.1   Distribution of Enhanced Security Profiles via NIS

The following restrictions apply to distributing enhanced security profiles via NIS:

- Successful and unsuccessful login attempts for NIS-shared accounts require the completion of the following steps:

  1. The master system's `rpc.yppasswdd` daemon must respond and update the last successful and last unsuccessful login fields in the `prpasswd` NIS map.

  2. The NIS slave servers must answer to the `yppush` operation initiated from the `rpc.yppasswdd` daemon. (Most successful logins do not require a `yppush` operation, but login failures and password changes do.)

  The login process will not continue or terminate until both of these steps are completed.

  The more NIS slave servers that are present in a given NIS domain, the more time `rpc.yppasswdd` takes to complete these steps. Also, nearly-simultaneous login attempts are processed sequentially by the NIS master, each waiting on a possible `yppush` for the previous attempt to succeed. Therefore, if several simultaneous attempts arrive at once, some may timeout and require you to log in again. You can alleviate this problem to some extent by using the `-p` option of `yppush`. One way to do this is to modify the `/var/yp/Makefile` file and change the `YPPUSH=` line. The following example allows up to 6 simultaneous transfers to NIS slave servers (the default number is 4):

  ```
  YPPUSH=$(YPDIR)/yppush -p 6
  ```

- The time allowed for responses to RPC requests is only 25 seconds. The more profiles that are present in the `prpasswd` map, the more likely the time limit is to expire during a login attempt, causing that attempt to fail. Simultaneous or nearly-simultaneous login attempts will fail if the NIS master server does not respond quickly enough to the pending login processes. If the total time taken on the NIS master for the following commands exceeds 25 seconds, then there will be circumstances under which only one user will succeed in logging in at a time:

  ```
  # cd /var/yp
  # make passwd prpasswd PRPWDPUSHONLY=1 NOPUSH='"'
  ```

  You can decrease the time required for map transfers if you use the `btree` format to store the maps on all of your NIS servers.

  With successful logins, the `rpc.yppasswdd` daemon will defer pushing

the maps if the login notification comes from a Version 4.0D client. Therefore, the yppush operation is only completed when an older client initiates the operation or when it is necessary to clear a failed login count.

- Sites that cannot use NIS to share prpasswd information may be able to use NFS to share the /tcb/files and /var/tcb/files directories instead. This requires you to export the directories with root access to the participating nodes (with −root=0 or −root=client1:client2:client3 as appropriate). It also requires you to enable NFS locking to ensure that no database corruption occurs. For more information, see the exports(4) reference page.

In previous releases of the operating system, NIS slaves that were listed in the ypservers NIS map on the NIS master but that did not already have a copy of the prpasswd and prpasswd_nonsecure NIS maps may not have succeeded in transferring those maps during the yppush operation. This problem has been fixed for Version 4.0D and higher.

4.3.7.2 Disaster Recovery

Because the user profile and tty information is now stored in database files, the previous recovery method of editing the files while in single-user mode is no longer available. However, as long as the /usr (and, if separate, /var) file systems are mounted, you can use the edauth utility in single-user mode to edit extended profiles and ttys database entries.

If the /etc/passwd file is somehow lost, but the extended profiles are still available, then you can use a command sequence as in the following example to recover some of the missing data (the "\" characters indicate line continuation):

```
# bcheckrc
# /tcb/bin/convuser -dn | /usr/bin/xargs /tcb/bin/edauth -g | \
sed '/:u_id#/!d;s/.*:u_name=//;s/:u_id#/:*:/;s/:u_.*$/:/' \
>psw.missing
```

This will create a psw.missing file containing entries like the following:

```
root:*:0:
```

Primary group information, finger information, home directory, and login shell are not recorded in the extended profile. You must recover the data for those fields by other means.

4.3.7.3 The pw_id_map and gr_id_map Files Are Obsolete

The Enhanced Security routines pw_idtoname, pw_nametoid, gr_idtoname, and gr_nametoid (described in pw_mapping(3)) previously used the /etc/auth/system/pw_id_map and /etc/auth/system/gr_id_map files to find the required information for mapping names to numeric identifiers, and vice versa. The disk space

required by those files imposed a limit on how many accounts a system could support.

The Enhanced Security routines no longer use the `pw_id_map` and `gr_id_map` files. If you are running Tru64 UNIX Version 4.0D or later and still have those files, it is recommended that you remove them to recover the space occupied on the `root` parition.

4.3.7.4 Behavior of useradd, usermod, and userdel Commands

The `useradd` command correctly honors the default administrative lock value found in the `/.sysman/Account_defaults file`. If `Account_defaults` does not exist, the internal default for `useradd` is to create locked accounts. Use the `administrative_lock_applied` extended command line option to override the default. In the following example, `useradd` creates a locked account for `foo` regardless of the default value for administrative lock:

```
useradd -x administrative_lock_applied=1 foo
```

For base security, a locked account has the text `Nologin` in the password field in the `/etc/passwd` file. If an account is unlocked and has no password, that account has no value in the password field. The account is open and accessible to anyone. A warning is displayed if an unlocked account with no password is created.

For enhanced security, all accounts have an asterisk (`*`) in the password field in `/etc/passwd`, but the lock flag in the protected password database is correctly set to reflect the lock status. As with base security, an unlocked account with no password is accessible to anyone.

The `usermod` command correctly sets the lock flags for enhanced security when the `administrative_lock_applied` option is given on the command line. If `usermod` is used to unlock a locked account with no password, a warning is displayed.

The `userdel` command will retire, instead of remove, accounts on a system running enhanced security.

4.3.8 Enabling Extended UIDs in the Kernel

By default, extended UIDs are not enabled in the kernel. To enable this feature, use `sysconfig` or the `dxkerneltuner` interface to set the value of the variable `enable_extended_uids` to 1 (enabled).

If you choose to do this, please note the following:

*   While `enable_extended_uids` can be set dynamically, you cannot disable it dynamically. To disable `enable_extended_uids`, set the boot time value to 0 (disabled) and reboot your system.

- If you cannot easily reboot your system, note that enabling extended UIDs has very little impact on overall system performance.

4.3.9 Insight Manager Agents

Insight Manager Agents for Tru64 UNIX include SNMP-based subagents and Web Based Enterprise Management (WBEM) capabilities to present SNMP data in a format viewable by a Web browser. It provides a rich view of the data using HTML 2.0 and JavaScript in the form of Web pages.

4.3.9.1 Installation

Insight Manager is installed automatically when you install the operating system's mandatory subsets. They are OSFCLINET440, which includes the Compaq Mibs and SNMP subagent, and OSFIMXE440, which includes the Insight Manager subagents. For further information on these subsets, see the *Installation Guide*.

4.3.9.2 Starting and Stopping the Insight Manager Daemons

The Insight Manager daemons are started as part of the init 3 startup process. The Compaq SNMP subagent (/usr/sbin/cpq_mibs) is initialized and started or stopped as part of the snmpd startup script:

```
# /sbin/init.d/snmpd start
```

or

```
# /sbin/init.d/snmpd stop
```

The WBEM daemon is started and stopped from a new script called insightd:

```
# /sbin/init.d/insightd start
```

or

```
# /sbin/init.d/insightd stop
```

4.3.9.3 How to Monitor the Tru64 UNIX Managed System

The Compaq Insight Manager Agents provide access to device management data over the industry-standard HTTP protocol. This lets you view management data from a Web browser, either locally or remotely.

Use the following URL to view data locally:

```
http://127.0.0.1:2301/
```

Use this URL if localhost is configured:

```
http://localhost:2301/
```

Use this URL to view data remotely:

```
http://machine:2301/
```

In this example, `machine` is the IP address or host name under DNS.

The Insight Manager Device home page provides a single service icon entry called Compaq Insight Management Agents. Clicking on this icon displays the device pages of the target system. The home page also presents three hyperlinks:

- Refresh
- Options
- Devices

The Options link provides various Agent configuration options.

The Devices link provides auto-discovery of other Insight Manager Agents and Insight Manager XE Servers running on the local network, allowing direct access to their 2301 ports.

The Compaq Insight Manager Agents page consists of three frames, as follows.

The upper left frame displays a conditions legend and the following links:

Menu Item	Description
Agent Help	Use this link for online help.
Summary	Use this link to navigate back to the Summary Page.
Device Home	Use this link to return to the Device home page.
Options	Use this link to navigate to the Options Page and set options for Display Mode, Help Icons, and AutoRefresh intervals.

The lower left frame is a navigational frame that provides links to various system components grouped under Configuration, Mass Storage, NIC (Networking), Utilization, and Recovery. Clicking on the system component links displays component-specific data in the main pane window.

Not all the components are instrumented in this first release of the Insight Manager Agents; more will be made available in future releases. The following table shows data availability in Tru64 UNIX Version 4.0F.

```
Configuration
        System Info             Yes
        System Board            Yes
        Software Version Info   N/A
```

Mass Storage
        File System Space Usage   Yes
        Diskette Drives   Yes
        SCSI Controllers   Yes
                Device Information   Yes
                Device Statistics   N/A
                Bus Information   N/A

NIC
        Virtual NIC   N/A
        Network Controllers   N/A

Utilization
        File System Utilization   Yes

Recovery
        Reboot   N/A
        Power On Messages   N/A
        Environment   N/A
        Power Supply   N/A

Note

CPU Cache and Memory module information are available only in AlphaServers supporting FRU V4 configuration tables, specifically: AlphaServers 1000, 1200, 4100, and 8400.

4.3.9.4 Browser Requirements and Security

The minimum browser requirements include support for tables, frames, Java, JavaScript, and the Java Development Kit (JDK) 1.1. Additional browsers, or the browsers mentioned, used with different operating systems, might not work correctly, depending upon their specific implementations of the required browser technologies.

The requirements are TCP/IP and one of the following browsers:

Operating System	Browser Requirements
Windows 95/98 or NT Version 3.51/4.0	Microsoft Internet Explorer 4.0 Version 4.72.2106.8 or above
	Netscape Navigator Version 4.04 or above
Tru64 UNIX Version 4.0F and higher	Netscape Communicator 4.5 or higher

Operating System	Browser Requirements
DIGITAL UNIX Version 3.2C and higher	Netscape Communicator 4.06 or higher

Important: You must turn on the following options so that the Server Agents work properly:

```
Enable Java
Enable JavaScript
Accept all cookies or
Accept cookies originating from the same server as the page being
viewed
```

4.3.9.5 Security Access

In this release, you can access anonymously the device web page without logging in. If needed, the Insight Manager administrator can set user privileges.

There are only three predefined users (user, operator, and administrator) and passwords.

There are three types of data: Default (read only), Sets (read/write), and Reboot (read/write). The `WEBAGENT.INI` configuration file specifies the level of user that has access to data.

Refer to the Agent Help link in the online help for further information on setting security and on the default user, operator, and adminsitrator passwords.

To change the default passwords:

1. Navigate to the Device home page.

2. Click on the anonymous link in the text:

   ```
   Login Account: anonymous.
   ```

   The Account Login page is displayed.

3. Click on the `changed` link in the text:

   ```
   The Password for a login account may be changed at any time by
   an Account Administrator.
   ```

Another Account Login page is displayed. This page gives the appropriate instructions.

In this release, the Web-enabled Compaq Server Management Agents do not provide SNMP sets for system parameters, or reboot capability.

In Tru64 UNIX, the user access configuration file is located in `/usr/share/sysman/web/im/webagent/WEBAGENT.INI`. In a dataless environment, you will need to copy that file to the local directory `/var/im/webagent` if you need to change its configuration setting.

4.3.9.6 Known Problems

The following problems exist in current version of the Insight Manager:

- Contrary to the message presented on the initial web page (`http://machine:2301`), this version of Insight Manager does not generate traps. Thus, alerts are not generated.

- The Insight Manager Device Discovery web page (`http://machine:2301/cpqdev.htm`) may show inconsistent or incorrect data on some platforms, as active discovery is not fully functional.

- On some browsers, the login dialog box, which consists of text fields for Name and Password, opens with the initial focus on the Password text field.

- The information displayed in the Capacity field for disks described in the Mass Storage web pages is calculated as `blocksize` * 1,000,000. This will be changed to the defined calculation of `blocksize` * 1,048,576 in a future release.

- The Insight Manager AutoRefresh option, when set for less than 60 seconds, might stop refreshing web pages when run from a NetScape browser running on Tru64 UNIX.

- The Expansion Bus information displayed on the Insight Manager System Information web page may be incorrect or incomplete. For example, on systems where both ISA and PCI buses exist, only ISA is displayed.

- The Insight Manager is not supported on the DEC 7000 systems. The `cpq_mibs` daemon will fail to initialize on this platform. This problem will be fixed in the first available patch kit.

4.3.10 Startup Messages Lost in Large Configurations

On systems which display a large number of console messages at system initialization (typically, systems configured with a large number of devices), some of the messages may be missing from the `/var/adm/messages` file. This problem can be corrected by increasing the size of the kernel's message buffer.

Use either of the following procedures to change the buffer size. You must be `root` to make the change.

To change the buffer size using graphical administration tools, use the following steps:

1. Start the `dxkerneltuner` application.

2. Select the `generic` subsystem.

3. Set the Boot Time Value entry for the `msgbuf_size` attribute to the new value.

4. Apply the change before exiting.

To change the buffer size from the command line, use the following steps:

1. Create a temporary file, `/tmp/msgbufsize`, containing the following lines, but replacing the 32768 with the size appropriate to your system:

    ```
    generic:
    msgbuf_size = 32768
    ```

2. Enter the following command:

    ```
    % sysconfigdb -f /tmp/msgbufsize -m
    ```

    If a different entry was already present in the database, `sysconfigdb` will display a warning message to advise you of the change in size.

The increase will take effect at the next system reboot. After rebooting, you can verify the change by entering the following command:

```
% sysconfig -q generic | grep msgbuf_size
```

Note

The default size of the message buffer is 4 KB, and the example above sets it to 32 KB. Since the space used by the buffer is not returned for general use after initialization, you should set the size only high enough to correct the problem.

4.3.11 Incoming Remote Print Jobs Disable Queue If They Create spooldir Lock File

A print job to a Version 4.0F system (from any other version of Tru64 UNIX) might cause the print queue on the receiving system to become disabled.

This occurs when the lock file is not present in the spool directory for the queue when the remote host receives the job. This condition exists if the print queue was newly created, or if the contents of the spool directory were manually purged (which would delete the lock file).

When a remote job is received, the print daemon creates a lock file with a protection mode that signifies that the queue is disabled. The initial job

completes, but any subsequent jobs submitted to the same queue do not print because the queue is in the disabled state. When the queue is disabled by this condition, the protection modes on the spool directory lock file are displayed by the `ls -l` command as:

```
-rwxr-xr--
```

Normally the lock file should have these protection modes set:

```
-rw-r--r--
```

To restart the queue, enter the `lpc up queuename` command, to clear the disabled condition and restart the queue.

To avoid the problem, create the lock file by submitting any job on the local host to the queue, or by entering the following command:

```
# touch spooldirpathlock
```

4.3.12 The lpstat Command Fails

The `lpstat` command fails if the `PRINTER` environment variable is not set and the command is issued with the following syntax:

```
# lpstat -u username
```

The command fails with the following message:

```
lpstat: lp: unknown printer
```

If the `PRINTER` variable is set, the command returns queue information for the queue specified by `PRINTER`.

4.3.13 The Alias Name Field Is Empty in the Disk Configuration

The Disk Configuration application creates the correct `ls /dev/alias` files. It should also show the `alias` in the Alias Name: field in the Disk Configuration: Configure Partitions box after naming it in the previous session, but it does not. Use the `ls /dev/alias[a - h]` command to verify and use the disk alias.

4.3.14 Problems when Trying to Mount DVD–ROM

The preliminary support for the native DVD file system in Version 4.0F, called DVDFS, is not fully functional. Attempts to invoke the DVDFS file system result in an error message of `function not supported`.

When mounting DVD media that contains an ISO-9660 directory structure (which may be present on DVD–ROM disks formatted according to ECMA-167 that include the bridge format), any attempts to mount specifying `-t dvdfs` will use the CDFS file system (`-t cdfs`) automatically. When the on-disk format does not contain an ISO-9660 directory structure, this mount

will fail. Only DVD–ROM disks that include an ISO-9660 directory structure are supported in Version 4.0F.

For both CD–ROM media and DVD–ROM media that contains an ISO-9660 directory structure, Compaq recommends use of the `-t cdfs` option, as well as the `-o noversion` or `-o rrip` option where appropriate. Refer to mount(8) for a description of these options.

4.4 Network and Communications

The following notes apply to network and communications software.

4.4.1 IP Switching over ATM

The following restrictions apply when using IP switching over ATM:

- Do not use the `atmsig` command to start UNI signaling on a driver used for IP switching.

- Automatic startup of IP switching at boot time is not yet supported. For this release, you must configure IP switching interfaces in the `/etc/atm.conf` file, and not through the `atmsetup` utility.

4.4.2 Changes to ATM syslog Messages

In Version 4.0F, many syslog messages for routine ATM events are not displayed by default. The messages which show the duration of virtual connections (VC accounting messages) must be explicitly enabled.

The general verbosity of ATM messages can be controlled by setting the `global_message_level` parameter associated with the `atm` module in the `sysconfigtab` file. For example:

```
sysconfig -r atm global_message_level=2
```

Values for this parameter are:

- 0 – Normal level. Most routine or transient events are suppressed.

- 1 – Informational level. Adds messages for normal VC release and other state changes.

- 2 – Verbose level. Provides detailed tracing of most events and state changes.

Message levels can be set for individual ATM components to tailor the amount of information displayed. Refer to `sys_attrs_atm`(5) for more information.

4.4.3 Edit the /etc/hosts File After Running netsetup

In Tru64 UNIX Version 4.0F, once you configure your interfaces using
`netsetup`, view the `/etc/hosts` file and make sure that the hostname
and the IP address of your system has been added to this file. If not, then you
need to add this information in the `/etc/hosts` file. Edit the
`/etc/hosts` file as a root, and add the following line:

```
<system's ip address> <system's hostname> <any aliases for the system>
```

Example:

```
16.60.266.6    example.domain.com    example    example-alias2
```

where:

`16.60.266.6` = IP address of your system.

`example.domain.com` = fully qualified hostname of your system.

`example`, `example-alias2` = aliases for your system.

Refer to the `hosts`(4) reference page for further information regarding the
`/etc/hosts` file.

4.4.4 CDE's Static Dependency on the Network

The Common Desktop Environment (CDE) provides facilities and features
for applications to communicate in a networked environment. After the
network is configured and enabled, these features become available each time
a new desktop session is started. After a desktop session has started, the
current session has a static dependency on the state of the network
configuration. Network and system administrators should be very cautious
about dynamic changes to the network configuration while in a network-
aware desktop session.

Prior to making any dynamic network changes, such as changing the state of
your network adapter to off or changing your primary network address, add
the following entry to the `/.dtprofile` file:

```
export DTNONETWORK=true
```

The system administrator must then log out and back in as root for the
change to take effect. This change removes the dependency on the state of
the network. Failure to do this may result in a session hanging after clicking
on a CDE icon, such as the screen lock or Exit icons.

After all network changes are completed, remove the `export`
`DTNONETWORK=true` entry from the `/.dtprofile` file.

4.4.5   Support for NetRAIN Network Interface

Tru64 UNIX now provides support for detecting the physical loss of network connectivity, and subsequent automatic switchover to a working network interface. This feature is called Redundant Array of Independent Network adapters (NetRAIN).

NetRAIN uses two or more network interfaces to provide redundancy. Although only one interface at a time is actually used for communication, all interfaces are monitored to ensure that traffic is flowing on each. If the interface currently being used should suffer a loss of connectivity, NetRAIN will switch network traffic to the next working interface. All the context of the previous interface is maintained: hardware address, multicast addresses, and so on.

NetRAIN is configured through extensions to the `ifconfig` command. It also features a programming interface via `ioctl()` for manual program control.

NetRAIN supports the following network adapters:

- DE435 (PCI/Ethernet)
- DE500-XA (PCI/Fast Ethernet)
- DE500-AA (PCI/Fast
- DE504-AA (Four-port/Fast Ethernet)
- DE500-BA (PCI/Fast Ethernet)
- DEFEA (EISA/FDDI)
- DEFPA (PCI/FDDI)
- DE422 (EISA/Lance Ethernet)
- DE425 (EISA/Ethernet)
- DEMFA (XMI/FDDI)
- DEMNA (XMI/Ethernet)
- ATM Emulated LAN (ELAN) interface (any supported ATM adapter)

TruCluster Software users should consult the Software Product Descriptions (SPDs) for TruCluster Software for a list of network interfaces supported by that product.

For more details about NetRAIN, see the *Network Administration* guide and the `nr(7)` and `ifconfig(8)` reference pages.

4.4.6 NetRAIN Interface Configuration Change

The method for configuring a NetRAIN interface will be improved in a future release of the operating system. If you configure a NetRAIN set in Version 4.0F, or did so in a previous releases of the operating system, you will need migrate your set to the new configuration at that time.

4.4.7 Network Services Fail to Start with NetRAIN Interface

If you configure a NetRAIN set in the `inet.local` file as documented in the *Network Administration* guide, and the NetRAIN interface is the only network interface configured on your system, network services will fail to start automatically at boot time. To work around this problem, configure your NetRAIN interface in the `rc.config` file as follows:

1. Log in as root.

2. If you have already configured NetRAIN, remove the `ifconfig` line in the `inet.local` file that starts the NetRAIN interface. It should look similar to the following:

   ```
   ifconfig nr0 18.240.32.40 netmask 255.255.255.0 add fta0,fta1
   ```

3. Set the adapter name for the NetRAIN interface you are configuring. Using the parameters from the previous step, you would enter the following command:

   ```
   # rcmgr set NETDEV_0 nr0
   ```

4. Enter the following command to set the `ifconfig` parameters that are used to initialize the NetRAIN interface (the "\" character indicates line continuation):

   ```
   # rcmgr set IFCONFIG_0 nr0 18.240.32.40 netmask  \
   255.255.255.0 add fta0,fta1
   ```

5. Enter the following command to indicate to the system that you have configured a network interface:

   ```
   # rcmgr set NUM_NETCONFIG=1
   ```

   Note that network services will not start if NUM_NETCONFIG is set to 0.

6. Restart network services to effect the changes.

If you need to deconfigure the interface for any reason, use the `rcmgr` utility to delete the appropriate NETDEV_0 and IFCONFIG_0 entries and update the NUM_NETCONFIG variable to 0. Then reboot your system. See the `rcmgr`(8) reference page for more information.

If you plan to add a non-NetRAIN interface to your system in the future, you will need to deconfigure the NetRAIN set in the `rc.config` file and reconfigure it in the `inet.local` file as documented in the *Network*

Administration guide. If you do not make this change before adding the new interface, the Common Desktop Environment (CDE) Network Configuration application (netconfig) and the netsetup utility might yield unpredicable results.

This problem will be addressed in a future release of the operating system. If you configure a NetRAIN set, you will need to migrate your set to a new configuration at that time.

4.4.8 LAT and NetRAIN

By default, the latsetup utility attempts to configure LAT for operation over all network interfaces on a system.

If you have configured a NetRAIN virtual interface on your system, and the latsetup utility attempts to run LAT over any of the physical network adapters that comprise the NetRAIN set, the latstartup command will hang.

To prevent this problem, configure LAT to run over specific network interfaces on the system as documented in Section 5.7 of the *Network Administration* guide. Run LAT over an interface that is not part of the NetRAIN set.

4.4.9 NetRAIN and Microsoft Windows NT Clients

While using the default NetRAIN settings on a Tru64 UNIX server, Microsoft Windows NT clients reset their TCP/IP connections before the server is able to complete the failover to another interface. There are two possible solutions:

- For environments with a few Windows NT clients, you can modify the TcpTimedWaitDelay and TcpMaxDataRetransmissions NT TCP parameters as follows:

  ```
  TcpTimedWaitDelay= 30//The maximum retransmission timer.
  TcpMaxDataRetransmissions=20//The maximum number of retransmissions.
  ```

- For environments with many Windows NT clients, it is more easier to modify two NetRAIN kernel attributes on the server instead of modifying the TCP retry timing parameter on each client system. On the Tru64 UNIX server system, modify the nr_max_retries and netrain_timeout kernel attributes using the sysconfig -r command or dxkerneltuner. For example, you can set the attributes

as follows:

```
nr_max_retries=1     //The default is 4
netrain_timeout=100  //The default is 1000
```

Alternatively, you can set the `nrmaxretry` and `nrtmoisr` values, respectively, when you issue the `ifconfig` command to create the NetRAIN set.

See `sys_attrs_netrain`(5) and `ifconfig`(8) for information on NetRAIN parameters.

Regardless of the solution you choose, if the host system's network interface cards (NICs) are connected to a layer 2 switch or bridge that is running a Spanning Tree algorithm, the Spanning Tree configuration parameters on the switch/bridge might dictate the failover time. In this case, the actual behavior of the TCP retry might depend on the switch implementation and configuration. See the documentation accompanying the switch/bridge for information.

4.4.10   Sender Address Problem When Using DECnet Mail

Tru64 UNIX Version 4.0F includes a new version of the `sendmail` application, which includes the ability to specify an owner for a mail alias, which can cause the identity of the sender on the mailer command line to be different from the identity of the sender in the From: header line. The end result is that the recipient of the mail might see the wrong address identified as the sender of the message.

If this problem occurs, the corrective action is to create a file, `/var/dna/defaults/Dnetrc.defaults`, containing the following lines:

```
mail11.SenderFromHeader: on
mail11.FromLineToUse: From
```

The `Dnetrc.defaults` file is a configuration file read by the mail11 mailer. The specified options tell mail11 to take the sender id from the header lines, and specifically to get it from the From: line.

After you create the `Dnetrc.defaults` file, obtain the latest version of `mail11v3` and replace the current version in the `/usr/sbin` directory. The `mail11v3` mailer understands the option specified and uses the appropriate sender id.

4.4.11   DHCP Database Migration (joind and bootpd)

Starting with Tru64 UNIX Version 4.0F, DHCP database files are stored in an entirely new format that is incompatible with older formats. The operating system ships with an online document, provided by JOIN Systems, that explains the reasons behind this change, lists the files that are affected,

and provides instructions for converting the files to the new format. The document and conversion utility, README-DB237 and conv185-237, respectively, are located in the /etc/join directory.

4.4.12 TCP Timestamp and Selective Acknowledgement Support

Tru64 UNIX Version 4.0F includes support for the TCP Selective Acknowledgment Options (SACK) as well as the TCP Timestamp and PAWS (Protect Against Wrapped Sequence Numbers) options. These TCP options can help improve TCP performance over unreliable networks. RFC 2018 TCP Selective Acknowledgment Options describes the SACK options. RFC 1323 TCP Extensions for High Performance describes both the Timestamp and PAWS option.

4.4.13 Latent Support for Multiple Network Adapters in a Subnet

Tru64 UNIX Version 4.0F provides latent support for a single system to have multiple network adapters in the same subnet; for example, tu0 configured with IP address 192.24.156.20 and tu1 configured with IP address 192.24.156.21, both with the same netmask. Previous versions require each network adapter to be in a separate physical subnet.

On connection establishment, the kernel chooses the interface that has the fewest number of connections. This connection-balancing effect could lead to greater throughput than on a system with just one network adapter per subnet.

You should use this feature only if both of the following configuration situations are true:

*   The underlying network topology is switched, and it is of type Ethernet or FDDI.

*   You are not running a routing daemon.

The feature might affect the operation of network software or, for example:

*   Multicast transmission might not work properly.

*   Utilities such as traceroute might give varied output, since the interface used might change from packet to packet.

To enable this feature, you must use dbx on the kernel and set the routemode variable to 1. You must patch the kernel for this to persist

accross a reboot, for example:

```
# dbx -k /vmunix /dev/mem
(dbx) assign routemode=1        (to set)
(dbx) patch routemode=1         (to patch)
```

The dependence on this kernel variable will be removed, and the behavior of this feature may change in a future release.

For more details about dbx, see the *Kernel Debugging* guide and the *Programmer's Guide*.

For more details about adding and viewing routes, see the route(8) and netstat(1) reference pages, respectively.

4.5 Local Area Transport

The following notes apply to Local Area Transport (LAT).

4.5.1 Duplicate Minor Numbers and latsetup

The latsetup utility sometimes creates devices with duplicate minor numbers. If you manually create LAT BSD devices that do not match the valid BSD tty name space convention, latsetup can create devices with duplicate minor numbers. For example, creating device tty0 with a minor number 2 instead of 1 can cause this problem.

4.5.2 CTRL/A Causes LAT tty to Change the Case of Characters

When a CTRL/A character is typed during a LAT tty session, all lowercase characters are converted to uppercase. Another CTRL/A changes the mode back to normal.

4.5.3 Simultaneous llogin Connections

When doing a number of simultaneous llogin connections, you should use llogin with the -p option. To speed up an llogin connection, add the target host name as a reserved service.

4.5.4 LAT Interactions with NetRAIN

See Section 4.4.8 for information on LAT interactions with NetRAIN.

4.6 File Systems

The notes in this section apply to file systems.

4.6.1 The newfs Command No Longer Searches disktab File

Starting with Version 4.0D, the `newfs` command no longer searches the `/etc/disktab` file for hard disk geometry information. It now performs an `ioctl GETDEVGEOM` call to determine the characteristics of a disk.

4.6.2 Using ACLs over NFS

For an NFS client to make direct use of ACLs or extended attributes (property lists) over NFS, you must be enable the `proplistd` daemon on an NFS server. You also must use the `proplist` mount option when mounting on the client. Access checks are enforced by the server in any case, although NFSv2 client caching could sometimes cause inappropriate read access to be granted. Correctly implemented NFSv3 clients make the necessary access checks.

Start the `proplistd` daemon by selecting the number of `proplist` daemons to run when you use the `nfssetup` utility. You can also use the `proplistd` command to start the daemon manually:

```
# /usr/sbin/proplistd 4
```

On the client, the file system must be mounted with the `proplist` option by either of the following methods:

* Add `proplist` to the options field in the `/etc/fstab` file:

   sware1:/advfs /nfs_advfs nfs rw,proplist 0 0

* Add the option to the mount command as follows:

   # **mount -o proplist sware1:/advfs /nfs_advfs**

See the `acl(4)`, `fstab(4)`, `proplist(4)`, `mount(8)`, `nfssetup(8)`, and `proplistd(8)` reference pages for more information. Note that the `proplist` option is not documented in `mount(8)`.

4.6.3 ACL Size Limitations

On AdvFS file systems there is a hard limit of 1560 bytes for a property list entry. Since Access Control Lists (ACLs) are stored in property list entries, this equates to 62 ACL entries in addition to the three required ACL entries. The `EINVAL` error is returned if you attempt to exceed this limit.

To facilitate interoperation of the UFS and AdvFS ACLs, a configurable limit has been imposed on UFS ACLs. The default value of the UFS limit is 1548 bytes, equivalent to the 65-entry limit on AdvFS. The UFS configurable limit on ACLs has been added to the `sec` subsystem and has been given the attribute name `ufs-sec-proplist-max-entry`. You can use the `sysconfig` utility to dynamically configure the attribute or you can use `sysconfigdb` or `dxkerneltuner` to statically configure the attribute in

the `/etc/sysconfigtab` file.

A configurable property list element size for UFS has also been added to the
`sec` subsystem and has been given the attribute name `ufs-proplist-`
`max-entry`. The value of `ufs-proplist-max-entry` must be larger
than `ufs-sec-proplist-max-entry` by enough space to hold a
property list element header. The `sysconfig` utility adjusts the `ufs-`
`proplist-max-entry` attribute automatically. The default value of
`ufs-proplist-max-entry` is 8192 bytes.

See the `cfgmgr`(8), `seconfig`(8), `seconfigdb`(8), `sysconfig`(8), and
`sysconfigdb`(8) reference pages for more information.

4.6.4   Advanced File System (AdvFS)

The following notes discuss features, problems, and restrictions of the
Advanced File System (AdvFS).

4.6.4.1   AdvFS and fsync()

You can use the `fsync()` system call to synchronously write dirty file data
to disk. There are two ways a file can have dirty data in memory. One way
is via the `write()` system call. The other is from a memory write reference
after an `mmap()` system call. For AdvFS files, the `fsync()` system call
writes out dirty data only from the `write()` system call. If dirty data from
an `mmap()` also needs to be written then you must also use the `msync()`
system call.

4.6.4.2   Reusing AdvFS Partitions

You can reuse a partition that was previously part of an AdvFS domain.
However, before you reuse the partition, you must remove the domain on the
partition you want to reuse. Use the `rmfdmn` command to remove the entire
domain. After the unused domain is removed, you can create a new domain
on the partition.

4.6.4.3   Disk Usage Information

Under certain conditions, the disk usage information on an AdvFS file system
may become corrupted. To correct this, turn on quotas in the `/etc/fstab`
file for the affected file system, and then run the `quotacheck` command on
the file system. This should correct the disk usage information.

4.6.4.4 Cleanly Unmount File Systems Before Changing Operating System Versions

If a system crashes or goes down unexpectedly due to a loss of power or other similar circumstances, AdvFS will perform recovery the next time that the filesets that were mounted at the time of the crash are remounted after rebooting. This recovery keeps the AdvFS metadata consistent and makes use of the AdvFS log file.

Different versions of the operating system use different AdvFS log record types. Therefore, it is important that AdvFS recovery be done on the same version of the operating system that was running at the time of the crash. For example, if your system is running Version 4.0F and the system crashes, do not reboot using Version 3.2G since that version of AdvFS may not be able to work with the log records that the Version 4.0F system put into the log.

Therefore, if you want to reboot using a different version of the operating system, make sure that any mounted AdvFS filesets are unmounted cleanly before rebooting. In addition, if the system panicked or an AdvFS domain was domain-panicked, it is best to reboot using the original version of the operating system and run the /sbin/advfs/verify command to make sure that the domain is not corrupted. If it is not, it is then safe to reboot using a different version of the operating system and remount the filesets.

4.6.4.5 Running Verify on AdvFS File Systems

Use the AdvFS verify utility (/sbin/advfs/verify) to check the consistency of the on-disk metadata in an AdvFS domain. This utility has been enhanced to better detect certain potential problems. Compaq recommends that you run the verify utility during your regular maintenance schedule.

You should run the verify utility when an entire AdvFS domain is scheduled to be offline, as no fileset in the domain can be mounted when you run the verify utility. If you create a cron job, ensure that it executes the appropriate commands to dismount the file systems.

You should also run the verify utility prior to doing an update installation.

4.6.4.6 The vdump Utility Requires Two kill Signals

Testing of AdvFS on a multiprocessor system showed that occasionally (about 30 percent of the time) when a test suite was run, it would fail because vdump processes would not respond to kill -9 signals. The workaround is to send such a process a second kill -9 signal.

4.6.4.7 Verify Fails on Dual-Mounted Domain

While running `verify` on a dual-mounted domain, it is possible to receive the following type of message:

```
Corruption in frag file type list for frag type 0
Expected to find set ID 36f68cbd.0006ae8c.1.8001 but found set ID
36f68c94.000b1b83.1.8001
```

This message can be ignored as it does not reflect an actual problem in the frag file, but instead an artifact of the dual-mounted domain.

4.6.4.8 Very Small AdvFS Domains Created In lockmode 4

Creation of domains smaller than 20 MB while running in lockmode 4 can cause the system to panic when the first fileset is created within the small domain. Such domains may safely be created and used if the system is not running in lockmode 4. To see which lockmode the system is running in, enter the following command:

```
% sysconfig -q generic lockmode
```

4.6.5 Logical Storage Manager

The following notes describe problems and restrictions of the Logical Storage Manager (LSM).

4.6.5.1 LSM Sparse Plexes May Erroneously Mask an I/O Failure

Under certain hardware failure scenarios, an LSM volume configured with a sparse plex may erroneously return success to the file system or application when in fact the I/O failed. This does not occur when you do not configure volumes with sparse plexes.

4.6.5.2 Restrictions with LSM root and swap Volumes

Root, primary swap, and secondary swap volumes configured under LSM have the following restrictions:

- Root and primary swap must be on the same physical disk and you must configure them using the encapsulation tools.

- If a secondary swap volume is configured under LSM, you must configure the root and primary swap.

- You must configure the root, primary swap, and any secondary swap volumes into the disk group named `rootdg`.

4.6.5.3 Possible Problems Accessing Physical Block 0 with LSM

Physical block 0 on Tru64 UNIX disks is typically write protected by default.
If a disk is added to LSM by using the `voldiskadd` utility, physical block
0 is skipped. However, if a partition that includes physical block 0 is
encapsulated into LSM by using the `volencap`, `vollvmencap`, or
`voladvdomencap` utility, physical block 0 is not skipped. This is not a
problem because the file system software already skips block 0 and does not
write to it.

A problem can occur when an LSM volume that contains a write-protected
block 0 is dissolved and its disk space is reused for a new purpose. Neither
the new application nor LSM know about the write-protected physical disk
block 0 and a write failure can occur.

To fix this problem, use the following steps to remove the write-protected
physical disk block 0 from the LSM disk before it can be assigned to the new
volume:

1. Use the `voldg` and `voldisk` commands to remove the disk from LSM.

2. Use the `voldiskadd` command to add either a specific partition of the
    disk or the entire disk to LSM.

4.6.5.4 Using LSM with SWXCR-P(A/B) and SWXCR-E(A/B) RAID Controllers

When you create an LSM mirror using a disk that is configured as Just-a-
Bunch-of-Disks (JBOD) with either the SWXCR-P or SWXCR-E RAID
controllers, a disk failure requires that you reconfigure the disk on the
controller. The disk is in an unusable state once it is set off line by the
controller and cannot be used by LSM until it is reconfigured. Refer to the
*StorageWorks RAID Array 200 Subsystem Family Installation and
Configuration Guide*.

4.6.5.5 Enabling LSM After Installation Requires Rebuilding the Kernel

If you use the `setld` utility to install LSM after you originally install Tru64
UNIX, you must rebuild the system kernel to enable LSM.

To rebuild the kernel, run the `doconfig` utility with no command flags.

Be sure to select all the appropriate options for your system, and be ready to
edit the configuration file as needed to reimplement any changes required by
your system.

Note that the `doconfig` menu display does not include LSM. However, the
`doconfig` utility will build a kernel that includes LSM. Refer to the
Logical Storage Manager guide for more information.

4.6.5.6   The volrootmir Script Supports Only LUN 0 on HSZ

Only LUN 0 is supported as a boot device by the console. Hence, you can only mirror the LSM rootvol and swapvol volumes to LUN 0 in an HSZ. Therefore, when you use the volrootmir script to mirror rootvol and swapvol, use only LUN 0 on an HSZ as an argument to the volrootmir script.

4.6.5.7   Using LSM rootvol Requires sysconfigtab Parameters

If you use the LSM rootvol volume for the root file system and the swapvol volume is in use as a primary swap volume, LSM adds the following entries to the /etc/sysconfigtab file to enable rootability:

```
lsm:
lsm_rootvol_is_dev=1
lsm_swapvol_is_dev=1
```

If these entries are deleted or if the /etc/sysconfigtab file is deleted, the system will not boot. If this happens, you can boot the system interactively as follows (the "\" character indicates line continuation):

```
>>> boot -fl i
Enter kernel_name option_1 ... option_n: vmunix \
lsm_rootdev_is_volume=1
```

Use the sysconfigdb utility to add the LSM entries as shown above to the /etc/sysconfigtab file after the system boots. Then, reboot the system for the changes to take effect.

4.6.5.8   Error While Dissociating a Preferred LSM Plex

When using LSM, a preferred plex can not be dissociated from an active mirrored volume. Attempting to do so will cause the following error:

```
fsgen/volplex: Unexpected kernel error in configuration update
```

Before dissociating the plex, change the volume's read policy so the plex is no longer the preferred plex. For example:

```
# volume -g my_diskgroup rdpol round my_volume
# volplex -g my_diskgroup dis my_plex
```

4.6.5.9   Presto Reset and Flush Operations May Cause System Crash

Due to a problem in the presto driver, avoid using the presto reset and flush operations (using the command presto -R or presto -F) to flush presto data. There is a possibility that the system will panic if a presto reset/flush operation is attempted. A workaround is to perform an orderly shutdown to flush all data that can be flushed and use the console command prflush -z to purge any remaining data.

Note that reset/purge operations will destroy file system data and so should be performed only after all attempts to recover devices and data have failed.

4.6.6 Mounting Associated Products CD-ROMs on Tru64 UNIX Versions Other Than Version 4.0F

The Associated Products CD–ROMs (APCDs) can be mounted with the mount command on Tru64 UNIX systems running Version 4.0E or later, as follows:

```
# mount -r /dev/rz4c /mnt
```

To mount the APCDs on releases prior to Version 4.0E, you must mount the CD–ROMs with the following options:

```
# mount -r -t cdfs -o rrip /dev/rz4c /mnt
```

On versions prior to Version 4.0D you might receive the following error message indicating that CDFS support is not built in to the current running kernel:

```
# mount -r -t cdfs -o rrip /dev/rz4c /mnt
/dev/rz4c on /mnt: No valid filesystem exists on this partition
```

If you receive this error, you need to build your kernel with the following option:

```
ISO 9660 Compact Disc File System (CDFS)
```

Development Environment Notes   **5**

This chapter contains notes about issues and known problems with the development environment software and, whenever possible, provides solutions or workarounds to those problems. The following topics are discussed:

- General programming
- Realtime programming
- DECthreads (pthreads)
- Profiling
- Debugging with dbx
- Java Programming

5.1 General Programming

The following note applies to general programming.

5.1.1 The malloc Function Is Now Tunable for Better Multithreaded Performance

The C runtime library `malloc` function (and associated functions) have been modified to allow significantly better concurrency when used by multithreaded applications. Additionally, three new memory allocator tuning variables have been added to allow more control of allocator behavior:

- `__first_fit`
- `__max_cache`
- `__delayed_free`

As always when developing applications making significant use of dynamically allocated memory and requiring maximum speed of execution, you should carefully read the Tuning Memory Allocation section of the `malloc(3)` reference page.

5.1.2 New DEC C Default Tuning Could Impact Applications That Directly Map I/O Space

Applications that directly map and access I/O space with bytes or shorts may be impacted by the new DEC C compiler.

The default tuning for the DEC C compiler has advanced its focus from EV4-EV5 architectures to EV56-EV6 architectures. With this change in tuning, the compiler now generates amask-guarded byte and word instruction sequences for some loops. The amask guards assure that the byte and word instructions will not execute on processors that do not support them. Less efficient instructions will execute instead.

The net result of this change is that users who recompile their applications with the default tuning may see a slight increase in object code size, a very slight decrease in performance on EV4-EV5 processors, and a sizable increase in performance on EV56-EV6 machines.

This change may be disruptive for applications that use special device driver interfaces that directly map I/O space for devices that do not support 8-bit and 16-bit access granularity.

If those applications are compiled without −Wf, −static and are run on EV56-EV6 machines they may corrupt I/O memory. To avoid this possibility, those applications should be compiled with −tune ev5 which disables byte/word instruction generation.

5.2 Realtime Programming

The following notes apply to realtime programming.

5.2.1 SA_SIGINFO Not Visible Under Certain Namespace Conditions

The symbol SA_SIGINFO, defined in sys/signal.h, is not visible under certain namespace conditions when _POSIX_C_SOURCE is explicitly defined in the application or on the compile line.

The SA_SIGINFO symbol is visible if you do not explicitly define _POSIX_C_SOURCE. For most applications, unistd.h provides the standards definitions needed, including _POSIX_C_SOURCE. As a general rule, avoid explicitly defining standards macros in your application or on the compile line. If you do explicitly define _POSIX_C_SOURCE, then SA_SIGINFO is visible if you also explicitly define _OSF_SOURCE.

5.2.2 POSIX 1003.1b Synchronized I/O and File Truncation

POSIX 1003.1b synchronized I/O using file status flags does not apply to file truncation. When file status flags are used to control I/O synchronization, no synchronization occurs for file truncation operations.

You can use the `fsync()` or `fdatasync()` function to explicitly synchronize truncation operations.

5.2.3 The fcntl() Function and F_GETFL with O_DSYNC File Status

A problem occurs when `fcntl()` is called with the `F_GETFL` request, and the file operated on has the `O_DSYNC` file status flag set. The return mask incorrectly indicates `O_SYNC` instead of `O_DSYNC`.

5.3 DECthreads (pthreads)

The following notes apply to DECthreads. See Section 8.10 and Section 8.11 for information about DECthreads interfaces that will be retired in a future release. See Section 1.11 for information about Visual Threads, a new product that lets you analyze your multithreaded applications for potential logic and performance problems.

5.3.1 Static Libraries

Users who desire optimal performance from DECthreads, and who are willing to relink on future versions of Tru64 UNIX, might want to use the DECthreads static libraries that are located in the `CMPDEVENH440` subset. Once this subset is installed, you can find the libraries in the `/usr/opt/alt/usr/lib/threads` directory.

Before using these static libraries, you should read the README file in the same location.

5.3.2 Signal Handling

Signal handling in the POSIX 1003.1c (`pthread`) interface of DECthreads is substantially different from signal handling for the draft 4 POSIX and the CMA interfaces of DECthreads. When migrating your application from the draft 4 POSIX or CMA interfaces to the POSIX 1003.1c interface, please see the IEEE POSIX 1003.1c standard or the *Guide to DECthreads* for a discussion of signal handling in threaded applications.

5.3.3 Scheduling Behavior (Contention Scope)

In releases prior to Version 4.0, thread scheduling attributes were systemwide. In other words, threads had a system contention scope. Since Version 4.0, thread policies and priorities are, by default, local to the process. No artificial limit exists for thread priorities of these process contention scope threads, the full priority range is accessible by every thread.

Previously, there was no way to control the contention scope of a thread. Starting with Version 4.0D, applications coded to the POSIX 1003.1c pthreads interface can set the desired contention scope upon thread creation. For more information on setting and determining thread contention scope, see the descriptions of the following routines in the *Guide to DECthreads*:

```
pthread_attr_setscope()
```

```
pthread_attr_getscope()
```

The guide also describes a problem with inheritance of the contention scope scheduling attribute in Versions 4.0D and higher.

Process contention scope threads provide faster context switches between threads in the same process, and reduce the demand on system resources without reducing execution concurrency. The Tru64 UNIX "two level scheduling" implementation (the code that supports process contention scope scheduling) automatically replaces kernel execution entities when a process contention scope thread blocks in the kernel for any reason, and it provides time-slicing of compute-bound threads. Therefore, there is no need to worry that using process contention scope will reduce parallelism or allow the execution of some threads to prevent other threads from executing.

The only code that should require system contention scope is code that must run on a specific processor via binding and code that must be directly scheduled by the operating system kernel against threads in other processes, particularly threads running inside the kernel. While the scheduling policy and priority of process contention scope threads is virtual and affects scheduling only against other threads within the process, the scheduling policy and priority of system contention scope threads (when the process runs with root access) can allow the thread to preempt threads within the kernel. While this can sometimes be valuable and even essential, extreme care must be used in such programs to avoid locking up the system. It might be impossible to interrupt such a thread.

5.3.4 Problems Using of the stackaddr Thread Creation Attribute

Compaq does not recommend using the `stackaddr` thread creation attribute which allows you to allocate your own stack for a thread. The semantics of this attribute are poorly defined by POSIX and the Single UNIX Specification, Version 2. As a result, code using the attribute is unlikely to be portable between implementations. The attribute is difficult to use reliably,

since the developer must, by intimate knowledge of the machine architecture and implementation, know the correct address to specify relative to the allocated stack. The implementation cannot diagnose an incorrect value because the interface does not provide sufficient information. Using an incorrect value might result in program failure, possibly in obscure ways.

5.3.5 DECthreads Read-Write Locks

DECthreads now supports read-write locks. A read-write lock is a synchronization object for protecting a data object that can be accessed concurrently by more than one thread in the same program. Unlike a mutex, a read-write lock distinguishes between shared read and exclusive write operations on the shared data object. A read-write lock is most useful in protecting a shared data object that is read frequently and modified less frequently. The following routines provide access to the read-write lock capability:

- `pthread_rwlockattr_destroy(3)`
- `pthread_rwlockattr_init(3)`
- `pthread_rwlock_destroy(3)`
- `pthread_rwlock_init(3)`
- `pthread_rwlock-rdlock(3)`
- `pthread_rwlock_tryrdlock(3)`
- `pthread_rwlock_trywrlock(3)`
- `pthread_rwlock_unlock(3)`
- `pthread_rwlock_wrlock(3)`

For more information about read-write locks, see the reference pages for these routines.

5.3.6 DECthreads Object Naming

DECthreads now allows you to assign names, as C language strings, to thread objects including threads, mutexes, condition variables, and read-write locks (see Section 5.3.5). During debugging, you can use these names to help identify individual objects by function rather than by the numeric identifiers the thread library assigns. The Ladebug debugger and the Visual Threads analysis tool (see Section 1.11) include these names when displaying information about thread objects. Other debuggers and analysis tools can also use the names you have assigned.

Use the following routines to assign and retrieve object names:

- `pthread_attr_getname_np(3)`

- `pthread_attr_setname_np(3)`
- `pthread_cond_getname_np(3)`
- `pthread_cond_setname_np(3)`
- `pthread_getname_np(3)`
- `pthread_key_getname_np(3)`
- `pthread_key_setname_np(3)`
- `pthread_mutex_getname_np(3)`
- `pthread_mutex_setname_np(3)`
- `pthread_rwlock_getname_np(3)`
- `pthread_rwlock_setname_np(3)`
- `pthread_setname_np(3)`

For more information about object naming, see the reference pages for these routines.

5.3.7 DECthreads Metering Capability May Not Be Reliable in Some Situations

In this release, the metering capabilities of DECthreads may not be reliable in a process that forks.

5.3.8 Memory Alignment Issue

Although older Alpha processors (prior to the 21264 chip) can only access memory in units of at least a quadword (8 bytes), multiple variables, each of which is less than eight bytes, can occupy the same quadword in memory. In such cases, multithreaded programs might experience a problem if two or more threads read the same quadword, update different parts of it, then independently write their respective copies back to memory. The last thread to write the quadword overwrites any data previously written to other parts of the quadword. This can happen even though each thread protects its part of the quadword with its own mutex.

The Tru64 UNIX C compiler protects scalar variables against this problem by aligning them in memory on quadword (8-byte) boundaries. However, in composite data objects such as structures or arrays, the compiler aligns members on their natural boundaries. For example, a 2-byte member is aligned on a 2-byte boundary. Because of this, any adjacent members of the composite object that total eight bytes or less could occupy the same quadword in memory.

Inspect your multithreaded application code to determine if you have a composite data object in which adjacent members could share the same

quadword in memory. If you do and if your project allows, Compaq recommends that you force alignment of each such member variable to a quadword boundary by redefining the variable to be at least eight bytes, or by defining sufficient padding storage after the variable to total eight bytes.

Alternatively, you can create one mutex for each composite data object in which adjacent members can share the same quadword in memory. Then use this single mutex to protect all write accesses by all threads to the composite data object. This technique might be less desirable because of performance considerations.

5.3.9   DECthreads pthread_debug() and pthread_debug_cmd() Routines

In order to allow for the possibility of a more comprehensive and robust threads debugging environment, it has become necessary to remove the `pthread_debug()` and `pthread_debug_cmd()` routines. To prevent existing binaries from failing, the routines will continue to be recognized. However, a call to either routine now results in an immediate return to the calling program. The `pthread_debug_cmd()` routine returns a zero (0) indicating success. Debuggers such as Ladebug and TotalView provide functionality formerly provided by these routines.

5.3.10   DECthreads SIGEV_THREAD Notification Mechanism

The `SIGEV_THREAD` notification mechanism works correctly, starting in Version 4.0D. Using this notification mechanism, a user-defined function is called to perform notification of an asynchronous event. The function is run as though it were the start routine of a thread and can make full use of the DECthreads synchronization objects.

The `SIGEV_THREAD` notification mechanism and the function to be called are specified in the sigevent structure. This mechanism is useful for programming with the POSIX 1.b realtime signal interfaces such as timers and asynchronous I/O. For information and cautions concerning the use of signals in a multithreaded environment, see the *Guide to DECthreads.* For more information about using `SIGEV_THREAD`, see the IEEE POSIX 1003.1c-1996 standard and The Open Group Single UNIX Specification, Version 2.

5.4   Profiling

The following notes apply to the profiler tools.

5.4.1   Change to hiprof's Profiling of Threaded Programs

The `-cputime` option of the `hiprof`(5) profiler now provides an
instruction-count profile for threaded programs, the same as the `-calltime`
option, because the CPU cycles reported for kernel-threads by the RPCC
instruction can not be mapped to `pthread`(3) threads.

The only significant difference is that the profile is displayed as the number
of instructions executed instead of CPU seconds used. The `-cputime`
option still profiles CPU seconds for nonthreaded programs.

5.4.2   Change in Naming of Files by cc Profiling Option

The `cc` command's `-prof_gen` option (which causes the `pixie` profiler
to be run after the executable is linked) names files differently from the way
it did in releases prior to Version 4.0E.

The new naming scheme is necessary to support formal benchmarking, which
is the primary purpose of the `-prof_gen` option.

Before Version 4.0E, the uninstrumented executable produced by the `ld`
linker and provided as input to `pixie` was named `a.out` (or as indicated
with the `-o` option). The instrumented executable produced by `pixie` was
given the usual `.pixie` filename extension.

Starting with Version 4.0E, the instrumented executable is named `a.out` (or
as indicated with the `-o` option). The uninstrumented executable is given a
`.non_pixie` file name extension.

5.5   Debugging with dbx

The following note applies to debugging with `dbx`.

5.5.1   Examining the User Program Stack in a Kernel Crash Dump

When debugging a crash dump with `dbx`, you can examine the call stack of
the user program whose execution precipitated the kernel crash. To examine a
crash dump and also view the user program stack, you must invoke `dbx`
using the following command syntax:

```
# dbx -k vmunix.n vm[z]core.n path/user-program
```

The version number (n) is determined by the value contained in the bounds
file, which is located in the same directory as the dump files. The *user-
program* parameter specifies the user program executable.

The crash dump file must contain a full crash dump. For information on
setting system defaults for full or partial crash dumps, see the *System
Administration* guide. You can use the `assign` command in `dbx`, as shown
in the following example, to temporarily specify a full crash dump. This

setting stays in effect until the system is rebooted.

```
# dbx -k vmunix.3
dbx version 5.0
    .
    .
    .
(dbx) assign partial_dump=0
```

To specify a full crash dump permanently so that this setting remains in
effect after a reboot, use the `patch` command in `dbx`, as shown in the
following example:

```
(dbx) patch partial_dump=0
```

With either command, a `partial_dump` value of 1 specifies a partial
dump.

The following example shows how to examine the state of a user program
named test1 that purposely precipitated a kernel crash with a syscall after
several recursive calls:

```
# dbx -k vmunix.1 vmzcore.1 /usr/proj7/test1
dbx version 5.0
Type 'help' for help.

stopped at  [boot:1890 ,0xfffffc000041ebe8]    Source not available

warning: Files compiled -g3: parameter values probably wrong
(dbx) where     1
>  0 boot() ["../../../../src/kernel/arch/alpha/machdep.c":1890,
0xfffffc000041ebe8]
1 panic(0xfffffc000051e1e0, 0x8, 0x0, 0x0, 0xffffffff888c3a38)
["../../../../src/kernel/bsd/subr_prf.c":824, 0xfffffc0000281974]
2 syscall(0x2d, 0x1, 0xffffffff888c3ce0, 0x9aa1e00000000, 0x0)
["../../../../src/kernel/arch/alpha/syscall_trap.c":593, 0xfffffc0000423be4]
3 _Xsyscall(0x8, 0x3ff8010f9f8, 0x140008130, 0xaa, 0x3ffc0097b70)
["../../../../src/kernel/arch/alpha/locore.s":1409, 0xfffffc000041b0f4]
4 __syscall(0x0, 0x0, 0x0, 0x0, 0x0) [0x3ff8010f9f4]
5 justtryme(scall = 170, cpu = 0, levels = 25) ["test1.c":14,
0x120001310]
6 recurse(inbox = (...)) ["test1.c":28, 0x1200013c4]
7 recurse(inbox = (...)) ["test1.c":30, 0x120001400]
8 recurse(inbox = (...)) ["test1.c":30, 0x120001400]
9 recurse(inbox = (...)) ["test1.c":30, 0x120001400]
    .
    .
    .
30 recurse(inbox = (...)) ["test1.c":30, 0x120001400]
31 main(argc = 3, argv = 0x11ffffd08) ["test1.c":52, 0x120001518]
(dbx) up 8     2
recurse:  30     if (r.a[2] > 0) recurse(r);
(dbx) print r     3
struct {
a = {
[0] 170
[1] 0
[2] 2
[3] 0
```

```
       .
       .
       .
(dbx) print r.a[511]    4
25
(dbx)
```

1. The `where` command displays the kernel stack followed by the user
   program stack at the time of the crash. In this case, the kernel stack has 4
   activation levels; the user program stack starts with the fifth level and
   includes several recursive calls.

2. The `up 8` command moves the debugging context 8 activation levels up
   the stack to one of the recursive calls within the user program code.

3. The `print r` command displays the current value of the variable `r`,
   which is a structure of array elements. Full symbolization is available for
   the user program, assuming it was compiled with the `-g` option.

4. The `print r.a[511]` command displays the current value of array
   element 511 of structure `r`.

5.6 Java Programming

The following note applies to Java programming.

5.6.1 Name Space Conflict Between Java and SVE

A file system conflict exists between Java and the System V Environment
(SVE) on Version 4.0 and later systems.

The problem arises because both Java and SVE use the file system path name
string `/usr/bin/alpha` for different purposes. Java creates
`/usr/bin/alpha` as a directory. SVE (specifically, the optional
SVEBCP4** Base Compatibility Package subset) creates
`/usr/bin/alpha` as a symbolic link to the
`/usr/opt/svr4/usr/bin/alpha` directory. The order in which these
applications are installed determines if the customer will experience a
problem. Here are three ways to avoid the problem:

- If SVE is already installed and the `/usr/bin/alpha` link exists, it is
  safe to remove the link. The link is not used after the SVE installation
  and for all other situations SVE will look for the directory location,
  `/usr/opt/svr/usr/bin/alpha`. That directory will be found and
  does not cause a conflict.

- If SVE has not yet been installed, Compaq recommends that SVE be
  installed without the Inclusive View option. If Inclusive View is not
  selected during SVE installation, the `/usr/bin/alpha` link will not be
  needed during SVE installation.

- If SVE has not yet been installed and the Inclusive View option is required, follow these steps:

  1. Temporarily rename the `/usr/bin/alpha` directory that is used by Java

  2. Install SVE

  3. Remove the `/usr/bin/alpha` link created by the SVE installation

  4. Restore the Java `/usr/bin/alpha` directory

There will be no patch or other resolution mechanism for this problem other than the workaround provided here.

Window System Software Notes 6

This chapter contains notes about issues and known problems with the windowing software and, whenever possible, provides solutions or workarounds to those problems. The following topics are discussed in this chapter:

- Hardware notes and restrictions
- X servers
- X clients
- CDE clients
- Windows programming
- Internationalization

6.1 Hardware Notes and Restrictions

The following notes apply to graphics hardware restrictions.

6.1.1 Qvision Graphics Display Error

Different versions of Qvision graphics boards demonstrate `fillsolid` drawing problems, leaving a line at the bottom of the screen, which is evident when running CDE blank lock screen. The line varies in color and intensity depending on the version of the Qvision board.

6.2 X Servers

The following notes apply to X servers.

6.2.1 Limited Multiscreen Display Support with CDE

CDE provides limited support for X servers with more than one screen. While a multiscreen environment is possible, a number of inconsistencies are noticeable. For example, colors in secondary screens may not be correct, icons may not display properly, and applications may not appear on the screen where they are invoked.

6.2.2 Do Not Modify the keymaps.dir File

The final revision of the X Keyboard Extension, XKB Version 1.0, will be different from XKB Version 0.65, shipping with this release. The format of `/usr/lib/X11/xkb/keymaps.dir` will change. Do not modify this file as it will not be preserved with future updates of the operating system.

To force the server to use a specific XKB keymap, add the `-xkbmap` option to the server options line in `/usr/lib/X11/Xserver.conf`. Refer to the `Xdec`(1X) reference page on for more information.

6.3 X Clients

The following notes apply to X clients.

6.3.1 X Window Colormap Resources

Color rich applications, such as Netscape, exhaust a large number of colormap resources. This results in problems with other graphical applications. For example, you may notice that icons normally displayed by the CDE Application Manager are not displayed when a color rich application is currently running on the system. Graphics applications and online help volumes may also be affected. The icon editor, `dticon`, may not be able to open a pixmap that contains a large number of colors.

In most cases this is a visual problem, and it may not be necessary to take any corrective actions. You can use the CDE icon labels in the same way as the icon for user actions such as drag-and-drop, and single and double click.

The simplest solution is to exit the color rich application. There are several alternate actions:

- Use the CDE Style Manager's Color application to select the number of colors used by the desktop or an application. Alternatively, if supported:
  - Set an application resource to limit color usage.
  - Start the application with a flag that controls color map installation. (Refer to the application documentation for further details.)

- In the case of Netscape, two workarounds are available; the first is preferred:
  - Use Netscape's `maxImageColors` resource to limit the number of colors that Netscape uses. A suggested limit is 96. You can do this

by placing the following line in the `$HOME/.Xdefaults` file:

```
Netscape*maxImageColors:       96
```

– Start Netscape with the `-install` flag, which specifies that Netscape should install its own colormap. Although this is supported, there are side affects such as:

- Noticeable application and background color changes when focus is moved in and out of the Netscape window

- Inability to clearly view the contents of the Netscape window when focus is moved out of the Netscape window

6.3.2 Old X Environment

If the XDM graphical login environment is selected instead of the default CDE environment, you may need to install the `OSFOLDX11440` subset to regain the expected user environment. If this subset is missing, the default X session will consist of a single `xterm` window and the `twm` window manager instead of the XDM session manager.

6.4 CDE Clients

The following notes apply to CDE clients.

6.4.1 Inaccessible Dialog Buttons

When running CDE with 640 x 480 graphics resolution, the OK, Apply, Cancel, and Help buttons of some application dialogs may be inaccessible. If this happens, you can correct it by setting the `DXmfitToScreenPolicy` resource to `as_needed` in the application's defaults file or, for systemwide problems, in the `/usr/dt/config/$LANG/sys.resources` file.

6.4.2 Screen Savers Prevent Efficient Power Management

When the screen on a DPMS-capable monitor is switched to standby, suspend, or off mode, the X server continues to run the screen saver. In CDE, where there are a number of active screen savers, this may defeat the CPU slowdown features for power management on certain Energy Star-compliant platforms. To minimize power consumption, you should stop using active screen savers by doing any of the following steps:

- In the Screen Saver panel of the Screen dialog under the Style Manager, select "Blank Screen" and deselect any active screen savers that may be running.

- Select the "Off" button in the same dialog box.
- Execute `xset s off` from a terminal client window.

6.4.3   Remote Invocation of CDE File Manager dtfile

File Manager, Application Manager, and Trash Manager are different views supported by the `dtfile` application. Avoid invoking `dtfile` from a remote system with the `DISPLAY` environment variable set appropriately. This restriction is necessary because of the client-server model used by the `dtfile` application and its close interaction with the tooltalk messaging system.

In the event of an unexpected behaviour from any of these utilities, close down all windows associated with the File Manager, Application Manager, and Trash Manager. Then kill all processes associated with `dtfile`. You can get the `pid` for each process by using the following command:

```
# ps -aef | grep dtfile
```

6.4.4   The dtmail Application

The following notes apply to the `dtmail` application.

6.4.4.1   MAILLOCKING Environment Variable

The `dtmail` application does not honor the user configurable mail-locking environment variable, `MAILLOCKING`. If you are using NFS, you must have NFS locking enabled on both client and server systems.

6.4.4.2   Tooltalk Locking Disabled by Default

The `dtmail` application disables tooltalk locking by default. To enable it, select the following option from the Mail Options->Advance dialog box:

```
Use network aware mail file locking
```

Alternatively, set the following option in your `$HOME/.mailrc` file:

```
cdenotooltalklock='f'
```

If tooltalk locking is enabled, and the `rpc.ttdbserverd` daemon is not running, you will get the following message:

```
Mailer is unable to obtain exclusive access to this
mailbox because the system is not responding.

For this time only, you can choose to open this mailbox
read-only, or to open it read-write without exclusive
access (use only if no one else is using this mailbox).
```

You can either quit `dtmail` and then start the `rpc.ttdbserverd` daemon, or click on read-write, which allows you to continue without tooltalk locking.

If you are running the `automount` daemon, the `dtmail` application may not be able to access your new mail inbox and you will see a dialog box showing the following message:

```
Unable to access an object required to complete the operation
```

If you see this message, copy the contents of your current mail inbox to a temporary file as a backup and perform the following steps:

1. Check that the following line is in your `.mailrc` file:

   ```
   set cdenotooltalklock
   ```

2. If the line is not in your `.mailrc` file, edit the file to include it.

Alternatively, you can use the following method while in `dtmail`:

1. From the Mailbox pull down menu, select Mail Options

2. From the Mail Options dialog box, select the category Advanced.

3. From the Advanced options, click on the option named: Use network aware mail file locking. This action should remove the check mark.

4. Close all dialog boxes and restart the `dtmail` application. Check the contents of your current mailbox and the backup mailbox to ensure that no mail was lost during this process.

6.4.4.3 Row and Column Display

Changing the values for rows and columns in the Mail Options->Message View will not take effect immediately when you click on OK or APPLY. Exit `dtmail` and restart it again from the control panel, command line, or file manager.

6.4.4.4 Attachments Are Not Included

If you compose a mail message by selecting one of the following from the `dtmail` main window, attachments will not be included:

• Compose->New, Include All

• Compose->Reply to Sender, Include

• Compose->Reply to All, Include

Include the attachments manually by using the drag and drop feature.

6.4.5 CDE's Static Dependency on the Network

CDE has a static dependency on the state of the network configuration. For more information, see Section 4.4.4.

6.4.6 The xnlLanguage Resource Settings May Cause Problems with CDE

CDE users should remove any `xnlLanguage` resource settings from their `.Xdefaults` files. Those settings are typically left over from one of the user's earlier DECwindows sessions, where the user selected a language from the session manager's language menu and then saved that setting.

The `xnlLanguage` settings should be removed because they override whatever language you select from the language menu in the CDE login window.

6.4.7 Possible Failure in the XOpenDisplay Call

When logging in to the CDE desktop, some applications may not restart. The X server process may not be able to handle all of the requests for new open connections, causing some to fail in the `XOpenDisplay` call. Some applications, like `xterm`, log startup errors in the `dxconsole` window, such as the following error:

```
xterm error: can't open display :0
```

To avoid this problem, add the following resource to your `$HOME/.Xdefaults` file:

```
Dtsession*contManagement:    2
```

This resource enables a handshake protocol between the CDE session manager and window manager during the login phase to control the appearance of new windows. While it may marginally increase the time before the login completes, it better assures that all applications will be restarted.

This can be added to the `/usr/dt/app-defaults/C/Dtsession` file to make the change for all users automatically.

6.5 Windows Programming

The following note applies to windows programming.

6.5.1 X11 Font Problem

In the `Lucida-Typewriter-medium-R-normal` fonts, the glyphs for multiplication and division are reversed. The multiplication symbol is where the division symbol should be, and vice versa. If the representation of these glyphs is important to your application, use a different font.

6.6 Internationalization

The following notes apply to restrictions on use of internationalization features in the windowing environments.

6.6.1 Japanese Keyboard Support in Console Mode

When running in single-user or console mode, Tru64 UNIX now supports two new Japanese keyboard types (JIS and ANSI) on AlphaStation and AlphaServer systems. (Japanese keyboard support is not available on TURBOchannel-based machines.)

To use JIS-type Japanese keyboards, like the PCXAJ-AA and LK411-JJ, you must set the `language` console environment variable to 50, as in the following example:

```
>>> set language 50
```

To use ANSI-type Japanese keyboards, like the LK411-AJ, you must set the `language` console environment variable to 52, as in the following example:

```
>>> set language 52
```

6.6.2 System-Default Keyboard Setting May Prevent User Login

When any user logs on to a system, the system-default keyboard setting must be appropriate for the keys that the user presses when entering characters in the username and password fields. Otherwise, characters that are correct from the user perspective, given the keyboard being used, might be treated as invalid. In this case, the user cannot log on the the system. This situation most often arises when a keyboard is being used in one language and the system-default keyboard setting is another language. You can change the system-default keyboard setting at the console prompt or, if the required language is not available at the console level, by editing the `Xserver.conf` file to change the keymap used by the X Server. See `keyboard`(5) for more detailed information about changing keyboard settings.

6.6.3 Dxterm Does Not Support Latin-2, Latin-4, and Latin-Cyrillic Characters

The `dxterm` application is not able to display Latin-2, Latin-4 and Latin-Cyrillic Characters even when the locale is set correctly. Therefore, `dxterm` should not be used for displaying the following languages: Czech, Hungarian, Lithuanian, Polish, Russian, Slovak, and Slovene. The `dtterm` application should be used in this case.

Documentation Notes 7

_____

This chapter contains release notes that apply to Tru64 UNIX Version 4.0F documentation. It discusses the following topics:

- *Installation Guide*
- *System Administration*
- *Assembly Language Programmer's Guide*
- *Programming with ONC RPC*
- *Gate Daemon Configuration Guide*
- *System Configuration and Tuning Guide*
- *Guide to Preparing Product Kits*
- Online help volumes
- Device Driver Documentation

7.1 Installation Guide

The following notes apply to the *Installation Guide*.

7.1.1 Chapter 12 (Performing a Worldwide Update Installation)

In both Sections 12.7 and 12.8, Step 7 shows the following incorrect path for invoking the `wwinstallupdate` script:

```
% /mnt/ALPHA/WORLDWIDE/wwinstallupdate
```

The correct path for invoking the `installupdate` script is as follows:

```
% /mnt/Worldwide_Language_Support/kit/wwinstallupdate
```

7.1.2 Appendix E (Associated Product Descriptions)

The following section is an addition to Appendix E.

UniCensus

The `unicensus` directory contains the `UNICEN301` software subset.

UniCensus is a tool for Tru64 UNIX that you can use to collect and archive system configuration information. UniCensus uses `sys_check`(8) to

produce a HTML report showing system configuration information, revision levels, storage subsystem configuration and other information. UniCensus is typically configured to collect configuration information on a periodic basis for archiving and transport to Compaq Services. It can also be configured to run on system reboot.

7.2 System Administration Guide

The following notes describe additional information for the *System Administration* guide.

7.2.1 Section 10.5.2.1 (Print Symbols for Advanced Printing Software)

The description of the `printcap` entries for Advanced Printing Software is inaccurate.

When setting up Advanced Printing Software, you should set the following print symbols:

*   `rm` – specify `@dpa` to indicate that jobs submitted to this printer should be directed to the Advanced Printing Software inbound gateway. The inbound gateway submits the job to an Advanced Printing Software spooler.

*   `rp` – specify the name of the Advanced Printing Software logical printer

7.2.2 Section E.4.1 (Hardware Requirements)

The Logical Partitions feature is supported on the AlphaServer GS140 system. An AlphaServer 8400 (upgraded to a GS140 by replacing the processor modules) is also supported.

XMI hardware may be used with Logical Partitions. However, XMI controllers and devices must be configured into partition zero (0). This is a console firmware restriction.

The minimum requirement for a partition is one IOP module. A partition may include a second IOP module. The maximum number of IOP modules for the entire system (the sum of all partitions) is three.

The restriction of a graphics device to windowing software (cannot be the console device) applies only to secondary partitions. A supported graphics device can be the console for the primary partition (partition 0). To use a graphics console, set the `console` environment variable to `BOTH` before

initializing partitions. For example:

```
P00>>> set console BOTH
```

7.2.3 Section E.4.3.3 (Configuring Logical Partitions)

The instructions for setting the IO port mask (`lp_io_mask`) do not take in to account a partition with two IOP module. For this case, the `lp_io_mask` is formed by combining (logical or) the masks for individual IOP modules slots. For example, to assign IOP modules in slots 7 and 8 to partition 1, `lp_io_mask1` would be 180.

When assigning IOP modules to secondary partitions, it is important to remember that one of the IOPs assigned to the partition must be connected to a DWLPB with a KFE72 option installed. The KFE72 provides the console serial port for secondary partitions.

7.2.4 Section E.4.3.7 (Disabling Automatic Boot Reset)

The Tru64 UNIX *Installation Guide* recommends setting the `boot_reset` console environment variable to ON. This setting is not compatible with Logical Partitions. The `boot_reset` console EV must be set to OFF.

7.2.5 Section E.4.3.8 (Set Memory Interleave Mode)

While setting the interleave mode to none, the console may echo a series of Y characters to the console display screen (there could be several lines of Y characters). This is not an error and can be ignored.

7.2.6 Section E.4.4 (Initializing Partitions)

The normal sequence of commands is as follows:

```
P00>>> set lp_count n
```

(where n is the number of partitions)

```
P00>>> init
```

(initialize the primary partition)

```
P00>>> lpinit
```

(start secondary partitions)

```
P00>>> boot
```

(boot primary partition)

```
P##>>> boot
```

(boot secondary partitions)

Improper operation results if the `lpinit` command is omitted. The console firmware prevents this by automatically executing the `lpinit` command if the `lp_count` is nonzero and a `boot` command is issued on the primary partition's console terminal.

On startup, each secondary partition displays configuration information (similar to the message in Section E.4.3.1). It is possible for this message to be proceeded by a series of Y characters. This is not an error and can be ignored.

7.2.7   Section E.4.5 (Installing the Operating System)

The operating system can be installed from a CD–ROM or over the network from a Remote Installation Server (RIS). It may not always be practical to configure a CD–ROM drive on all partitions and/or a RIS server may not be available. One alternative (assuming a local network is available) is to install the operating system to one partition from a CD–ROM, then configure that partition as a RIS server for the other partitions. Refer to *Sharing Software on a Local Area Network* for instructions on setting up a Remote Installation Server.

7.3   Assembly Language Programmer's Guide

The *Assembly Language Programmer's Guide* (Chapter 5) should include a description of the `.rconst` directive. This directive instructs the assembler to add subsequent data into the `.rconst` section. This behavior is similar to the `.rdata` directive, except that the entries cannot be relocatable.

7.4   Programming with ONC RPC

Sections 2.1.5 and 2.1.6 of the *Programming with ONC RPC* guide are correct only for a system running Tru64 UNIX or earlier version of the operating system. The code may not be portable to other UNIX systems.

If the client machine is not running Tru64 UNIX or earlier version of the operating system, copy the source code to the client. Check that the header calls are appropriate for the programming conventions on the client operating system and then compile the code locally.

7.5 Gate Daemon Configuration Guide

References to the *Gate Daemon Configuration Guide* in the *Technical Overview* and possibly other books are obsolete because this document has been removed from the *Tru64 UNIX 4.0F Documentation* CD–ROM. Information from the *Gate Daemon Configuration Guide* has been incorporated into the following reference pages: `gated.conf(4)`, `gated.control(4)`, `gated.proto(4)`, `gated_intro(7)`, and `gated(8)`.

7.6 System Configuration and Tuning Guide

In prior releases, reference information about kernel subsystem attributes was found in Appendix B of the *System Configuration and Tuning Guide*. Starting with Version 4.0F, reference information about subsystem attributes resides in the reference pages. See `sys_attrs(5)` for more information.

Note that the current set of reference pages for subsystem attributes includes more subsystems than were formerly covered in the book appendix; however, there are still subsystems with attributes that do not have reference pages. This omission will be addressed in a future release of the product.

The following errors exist in the `sys_attrs_ufs(5)` reference page:

*   The default value for the `inode-hash-size` attribute is listed as 2048. The correct default value is 512.

*   The default value for the `ufs-blkpref-lookbehind` attribute is listed as 16. The correct default value is 8.

7.7 Guide to Preparing Product Kits

The following information should be included in the *Guide to Preparing Product Kits*.

If you are developing an Asynchronous Hardware kit for Version 4.0F, you must ship symbolic links from the `/sys/BINARY` directory to every kernel module in your kit. In Version 4.0F, the kernel build process searches only the `/sys/BINARY` directory. As a result, symbolic links from the `/sys/BINARY` directory to your kernel module must be shipped in the same subset.

The following corrections should be made to Appendix A of the *Guide to Preparing Product Kits*.

7.7.1 Section A.1.12 (How to Build a Consolidated CD–ROM)

Step 3 of the instructions lists the incorrect paths:

```
% /usr/bin/mkisofs -D -R -a -d -o \
/spare/consolidate_digital_unix.cdfs /cdimage/
```

The correct paths are as follows:

```
% /usr/sbin/mkisofs -D -R -a -d -o \
/spare/consolidated_digital_unix.cdfs /cdimage/
```

7.7.2 Section A.2.2 (Building a Consolidated CD–ROM)

Step 4 of the instructions lists the incorrect paths:

```
% /usr/bin/mkisofs -D -R -a -d -o \
/spare/consolidate_digital_unix.cdfs /cdimage/
```

The correct paths are as follows:

```
% /usr/sbin/mkisofs -D -R -a -d -o \
/spare/consolidated_digital_unix.cdfs /cdimage/
```

7.8 Online Help Volumes

The notes in this section refer to problems with the online help volumes.

The Help viewer has the following known problems:

- In some cases, the help viewer is not correctly initialized. As a result, it will sometimes exhibit the expected behavior the second time an action is taken, but not the first time.

  For example, the first time a quick help dialog box displays a reference page, the Backtrack button may be enabled even though there is no place to which to backtrack. If the dialog box is closed and then opened again, the Backtrack button is dimmed.

- Similarly, if a request for on-item help displays the correct help, but one line has scrolled off the top, the online help will typically be displayed in exactly the right position when the on-item help request is repeated.

7.8.1 General Problems

The following known problems occur in the online help:

- The Appearance menu is not consistently documented.

  In the Archiver, License Manager, and System Information applications, there is an Appearance item on one of the menus. The item should offer three options: Text Only, Large Icon, and Small Icon. In some help

volumes, not all of these are documented.

- Some links from one help volume to another are displayed in a new view.

  There are a few links from one help volume to another that appear in a new help viewer window. In some situations, a distracting proliferation of help viewer windows can result.

  You can get a new view when you want one using the New Window item on the File menu in the help viewer.

7.8.2 Integration

The SysMan configuration applications on-item help does not work on the items in the menu bar.

In all the SysMan applications, the keyboard method of getting on-item help does not work on the Help menu.

7.8.3 Update Path in Installation Online Help Is Incorrect

The update installation path shown in the installation online help is not correct. The Preinstall Review List in the installation online help should state:

```
Your system must be running operating system Version 4.0D or 4.0E
to update to Version 4.0F.
```

7.8.4 Help Volumes by Application

The following problems apply only to help in specific applications:

- Kernel Tuner

  The Kernel Tuner application records its changes immediately, so if a system failure occurs while the Kernel Tuner is running, any boot-time changes will take effect the next time the system boots. If the /etc/sysconfigtab file contains invalid values, you can enter the following command at the boot prompt to boot with default values:

  ```
  boot -fl c
  ```

- Display Window

  The help volume for Display Window has opening instructions that show how to access the application from the CDE Application Manager. These should show that the Display Window icon appears in two groups.

- Network Configuration

  The names of the gated, joind, routed, and rwhod daemons are misspelled in the online help volume.

In the Configuring Interfaces dialog box, the fields under To Obtain IP Address are relevant for all interfaces.

- Printer Configuration

  In the Local Printer Settings dialog box and the Remote Printer Settings dialog box, the aliases in the Printer Aliases field must be separated by vertical bars because spaces are allowed in an alias.

- Shutdown

  The `/usr/sbin/shutdown` command now sends the Shutdown Message each time a reminder is sent of the time remaining before the system is shut down.

  The `/etc/nologin` file is not created until immediately before the shutdown occurs.

  System shutdown messages are sent to all users that are locally or remotely logged in to the system being shut down. Additionally, if the "Broadcast to NFS Clients" option is selected, shutdown messages are broadcast to all hosts that are NFS clients of the system being shut down.

- Disk Configuration

  The following corrected definitions replace the definitions in the glossary for Disk Configuration.

  – Skew

    A deviation from a reference direction, either by design or in response to lateral forces.

  – Track skew

    On a disk, the sector skew per track. The skew is the angle that sector 0 of the track changes from an imaginary radius line, due to a nonuniform number of sectors per track.

- BIND Configuration

  The following corrected definitions replace the definitions in the glossary for BIND Configuration:

  – BIND client

    A system that queries a BIND server for host name and address information, interprets the responses, and passes the information to requesting applications.

  – BIND server

    An authoritative source for information about one or more zones. It either maintains the master copy of the hosts database for the zone or obtains the information required to serve the hosts database from another server.

– DCE

Distributed Computing Environment. The capabilities of DCE are
defined by the Open Software Foundation (OSF).

– DCE cell

A logical group of systems that share services offered by DCE.

– DCE server

The server in a DCE cell.

– service type

In BIND Configuration, the available service types are BIND client
and BIND server. The service type determines whether a system is
configured to be a BIND client or a BIND server.

7.9 Device Driver Documentation

There are no changes to the Device Driver documentation for Version 4.0F.

Compaq maintains an online technical update page that contains any new
information about the Device Driver Kit since the product was released. To
view this document, go to the following URL:

```
http://www.unix.digital.com/faqs/publications/
pub_page/ddk_update.html
```

Features and Interfaces Scheduled for Retirement 8

This chapter lists features of Tru64 UNIX scheduled to be removed from, or changed in, future functional releases. Users and developers should plan to migrate away from or accomodate changes in these features in the near future.

8.1   DECwindows Applications

The following DECwindows utilities and tools will be retired in a future release of Tru64 UNIX. These dx* tools and utilities, commonly known as DECwindows, have been replaced by the dt* tools in Common Desktop Environment (CDE) that were introduced in Version 4.0 of the operating system. The replacement applications are listed in Table 8-1. Not all of the dx* applications that will be retired have a replacement due to the limited use or capability of that specific tool/utility.

Table 8-1:   Retired DECwindows Applications

Retired Tools/Utilities	Replacement Option(s)
dxprint	dtlp
dxcalendar	dtcm
dxcalc	dtcalc,xcalc
dxclock	Front Panel,xclock
dxpaint	dticon/dtstyle,bitmap
dxnotepad	dtpad
dxbook	dthelpview,Netscape
dxcardfiler	N/A
dxsession	xdm,dtsession
CDA	N/A

You should migrate to the dt* tools/utilities or other options as soon as possible.

8.2 Adobe Display PostScript, Client Libraries, and X Server Extension

Adobe has retired their Display PostScript (DPS) product, which includes the client libraries, X Server extension, and various applications and examples. Therefore, Compaq has also retired the Adobe DPS product from Tru64 UNIX. No replacements are available. Customers who have used the Adobe DPS libraries to develop their own applications do not have a migration path.

The following Adobe applications have been retired and are no longer included with the Tru64 UNIX operating system.

- `dpsclock`
- `dpsexec`
- `draw`
- `fontview`
- `libdps.a` and `libdps.so`
- `libdpstk.a` and `libdpstk.so`
- `libpsres.a` and `libpsres.so`
- `lib_adobe_dps.so`
- `makepsres`
- `pswrap`
- `scratchpad`
- `showps`
- `texteroids`
- `wonderland`
- `xepsf`
- All Adobe fonts under `/usr/lib/X11/fonts/Type1Adobe`

The following applications, which use Adobe DPS, have been retired and are no longer included with the Tru64 UNIX operating system:

- `dxvdoc`
- `dxbook`
- `dxnotepad` (internationalized version in WorldWide Language Support software in the *Tru64 UNIX Associated Products CD–ROM*)

During an installupdate from Version 4.0D or Version 4.0E to Version 4.0F, the DPS commands, libraries, and documentation will be marked by the `installupdate` procedure as obsolete inventory. The files will remain on the system until they are removed during a subsequent `updadmin` step.

Because these obsolete components have been retired from the operating system, they are no longer distributed in new versions, beginning with this version, and they might not be compatible with a future release. They are not supported as part of the Tru64 UNIX product beginning with this release.

If you have applications or procedures that depend on the DPS software that is now obsolete, you might want to keep the obsolete DPS software on your system, yet remove the other obsolete software at this time.

Before you perform the `updadmin` procedure, you can edit the file `/var/adm/smlogs/upd_obsolete_files` so that the DPS software components are not listed.

If you remove the DPS components from your system by using the `updadmin` utility, you will have the opportunity to archive all of the obsolete components. If you prepare an archive copy, and later discover a dependency on the obsolete DPS components, you can restore them to the system from the archive copy.

If you remove the obsolete DPS components without making an archival copy, there is no way to restore them from the product distribution media for the current release.

See Section 4.1.9.6 for information on possible user dependencies on the `showps` utility as a PostScript document viewer for the Netscape browser. The freeware application `gv` might be a suitable replacement for `showps`. It has a different user interface than `showps` but has no dependency on the obsolete DPS software libraries.

Some of the freeware applications delivered with Version 4.0D had dependencies on the DPS client libraries. Those applications have been rebuilt using the GhostScript libraries available on the freeware distribution media.

8.3 Nonconforming Curses Library

Version 4.0 of the operating system included a new X/Open compliant Internationalized Curses library. This library was not binary compatible with previous versions of the Tru64 UNIX Curses library, so compatible binaries (named `libcurses.a` and `libcurses.so`) were shipped in separate directories for Version 4.0. In that same release, the intent to put the compatible binaries into an obsolete subset and subsequently remove them in future releases was announced.

These compatible binaries were moved into the `OSFOBSOLETE` subset for Version 4.0D and will be removed in a future release of Tru64 UNIX.

8.4 The -oldc C Compiler

The official C compiler for Tru64 UNIX is DEC C. The previous compiler, implemented with -oldc, is no longer distributed in the default subset and will not be installed during a default system installation.

The -oldc compiler is still available in this version by installing the CMPDEVENH subset. Any problems reported with -oldc will be addressed in DEC C for Tru64 UNIX rather than in the -oldc compiler.

8.5 The dbx Debugger

The dbx symbolic debugger will be retired in a future release of Tru64 UNIX. The dbx debugger will be replaced by the ladebug debugger, which is a superset of the dbx functionality. The ladebug debugger is command line compatible with dbx and also supports a graphical user interface.

Compaq recommends that you begin using the ladebug debugger now and report any problems. This will provide a higher quality replacement when dbx is finally retired.

8.6 DEC C Compiler Default Change from -std0 to -std

The default language mode for the Tru64 UNIX C compiler will change from -std0 to -std. This will occur in a future release.

When the change occurs, it will be possible to revert back to -std0 by specifying -std0 on the compilation command line, by adding -std0 to the /usr/ccs/lib/cmplrs/cc/comp.config file, or by using an environment variable.

8.7 Change in struct utmp, struct utmpx, and struct lastlog

To provide compliance with several UNIX and Internet standards, the struct utmp, struct utmpx, and struct lastlog structures will be changed in the next major release (Version 5.0) of Tru64 UNIX. These changes will affect the /usr/include/utmp.h, /usr/include/utmpx.h, and /usr/include/lastlog.h files:

- The time field in struct utmp will change from a time_t structure to a struct __ut_timeval structure (to be consistent with /usr/include/utmpx.h).

- The ut_host field size (in struct utmp and struct utmpx) will be increased to comply with relevant Internet RFCs.

- The `ll_line` and `ll_host` manifest constants in
  `/usr/include/lastlog.h` will change to allow their sizes to
  correspond to the `ut_line` and `ut_host` fields in `struct utmp` and
  `struct utmpx`.

These changes will also affect the format of the `/var/adm/utmp`,
`/var/adm/wtmp`, and `/var/adm/lastlog` files. The following
conversion programs will be supplied:

- `/usr/lbin/wtmpconvert`

- `/usr/lbin/llconvert`

The programs will enable you to convert your existing `/var/adm/wtmp`
and `/var/adm/lastlog` files to the new format or convert new format
files to the old format for use by existing programs. See the corresponding
reference pages for more information.

8.8   C Language long double Type Changing to 128 Bits

In a future release of Tru64 UNIX, the default size of the C language `long
double` type will change from 64 bits to 128 bits. This will allow
application writers to perform mathematical calculations on numbers much
larger in magnitude with more precision than possible with the current `long
double` type, which is treated identically to the `double` type. A compiler
option will be provided to allow existing source code that expects a 64-bit
`long double` type to continue to be compiled and executed.

The one binary incompatibility that an existing application (if linked using
the `-call_shared` switch) could experience with the new default is
related to the input and output of `long double` types. Currently, the
`printf` and `scanf` functions, and other associated functions, interpret the
format code `%Lf` (capital-L followed by f) as referring to a 64-bit `long
double` type. In a future release, the interpretation of this format code will
be changed to expect the new 128-bit data type. Programs that use this
format code will either need to be changed, or will need to be run with the
new compatibility library that will be provided. An extra step will be
necessary to cause the application to use this library, and will be documented
in the release in which the change actually takes affect.

8.9   C Library Functions and POSIX P1003.1C

As of Version 4.0 of the operating system, the following C library functions
exist in two versions due to conflicts between previous versions of Tru64
UNIX and the recent IEEE POSIX P1003.1C standard (these new interfaces
are in effect by default). The old interfaces are currently accessible by

defining the C preprocessor symbol _POSIX_C_SOURCE to 199309L.

```
asctime_r     getgrnam_r    getpwuid_r    localtime_r    readdir_r
ctime_r       getlogin_r    gmtime_r      rand_r         ttyname_r
getgrgid_r    getpwnam_r
```

Binary compatibility is maintained in Version 4.0F; however, these routines will be retired in a future release of the operating system. The obsolete versions should not be used in new designs. These routines formerly resided in `libc_r.a` and `libc_r.so`, but were merged into the standard C runtime library.

8.10 POSIX 1003.4a (draft 4) pthread Routines in DECthreads

The POSIX 1003.4a, Draft 4 interface of DECthreads will be retired in a future release of Tru64 UNIX. Applications that were written using the POSIX 1003.4a, Draft 4 API should be migrated to the IEEE Std 1003.1c-1995, POSIX System Application Program Interface provided by DECthreads. The POSIX 1003.1c standard interface is the most portable, efficient, and powerful programming interface offered by DECthreads. A compatibility mode for the Draft 4 POSIX 1003.4a API has been provided to help ease migration. This compatibility mode will be removed in a future release.

8.11 DECthreads CMA Interface

The CMA interface of DECthreads is obsolete beginning with this release of Tru64 UNIX. Obsolescence means that while the CMA API will continue to exist in Tru64 UNIX and will be supported, CMA routines will no longer be documented or enhanced. Compaq recommends that you port your CMA based application to the IEEE Std 1003-1c-1995, POSIX System Application Program Interface provided by DECthreads.

8.12 Asynchronous I/O Binary Compatibility

Data structures for asynchronous I/O, like `aio_read()` and `aio_write()`, changed between DIGITAL UNIX Version 3.2 and Version 4.0. The kernel currently handles these differences so that applications built under Version 3.2 continue to run when executed on Version 4.0x systems.

In the next major release of the operating system, support for applications built under Version 3.2x using asynchronous I/O will be discontinued. These applications will need to be recompiled and relinked in order to run properly under Tru64 UNIX.

8.13 Nemacs

Nemacs Version 3.3.2, a public domain Japanese implementation of `emacs`, will be retired in a future release of Tru64 UNIX. `Mule`, a public domain multilingual implementation of `emacs`, will be carried forward as the replacement functionality for Nemacs. The Nemacs subsets `IOSJPNEMACS440` and `IOSJPNEMACSSRC440` will be removed from the system. For more information on `Mule`, refer to the `mule`(1) reference page.

8.14 Berkeley Software Distribution TTY-NAME

The intent to retire the BSD `TTY-NAME` namespace was announced in DEC OSF/1 Version 3.0. This functionality will be retired in a future release.

8.15 SCSI Device Names

Support for `rz` SCSI device names will be retired in a future release of Tru64 UNIX. Any code that derives knowledge about a device from the ASCII name or minor number may be impacted.

All code that uses the current namespace will be compatible when the change occurs because a mechanism that ensures binary compatibility will be provided. Existing interfaces such as names and minor numbers will be fully supported.

8.16 The -x and -p Options in addvol and mkfdmn

The `-x` and `-p` options of the `addvol` and `mkfdmn` commands allow you to set the per-volume bitfile metadata table (BMT) when you create a new volume or file domain. Users typically set the BMT to prevent an extent exhaustion problem.

In Version 4.0D and later, the limitations in the operating system that caused the extent exhaustion problem were removed, hence the `-x` and `-p` options are no longer needed and will be retired in a future major release.

8.17 LSM Block Change Logging (BCL)

In the next major release of Tru64 UNIX, the Logical Storage Manager Block Change Logging (BCL) feature will be retired and replaced by a new logging method. This new logging method, called Dirty Region Logging (DRL), will log regions instead of blocks for writes to LSM mirrored volumes. For most environments, DRL will provide the same ability to quickly resynchronize mirrors after a failure as BCL, but with considerably less logging overhead.

The logging format and configuration for DRL will not be compatible with BCL; for example, DRL will require a log size of at least two blocks. The next major release of Tru64 UNIX will automatically reconfigure volumes for DRL when a BCL mirrored volume has a log of two or more blocks. Mirrored volume configurations with a BCL of one block will require manual reconfiguration to continue to take advantage of logging for faster mirror recovery.

Due to these changes, you can configure a BCL size of two or more blocks to simplify BCL to DRL migration in the future. For optimum DRL configurations, use the following guidelines when configuring volumes with BCL: use a BCL subdisk block size of one block plus an additional block for every 1 GB of volume storage, then round up the BCL size to the next even number. Table 8-2 shows some examples.

Table 8-2: BCL Configuration Examples

Volume Size (GB)	BCL Subdisk Size (Blocks)
<1	2
1	2
2	4
3	4
4	6
5	6
6	8
7	8
8	10
9	10
10	12

8.18 LSM volassist Command Syntax

In the next major release of Tru64 UNIX, the syntax of the `volassist` command will change. It will no longer support the following constructs:

- `[!]medianame,offset`

  You will still be able to use `[!]medianame,offset` construct but you will not be able to specify an offset. If you need to specify an offset, you will have to use the low-level commands to create subdisks, plexes, and volumes exactly as required.

- `alloc=size`

  A new construct, `alloc=storage-spec[,storage-spec,...]`, will replace `alloc=size`. However, the new construct will not allow you to specify sizes for all allocations. You will need to use the low-level commands to create subdisks, plexes, and volumes exactly as required.

- `align=size`

  Two new constructs, `diskalign` and `nodiskalign`, will replace `align=size`, allowing you to specify whether subdisks should be created on cylinder boundaries. If you require the ability to specify alignments for all allocations, you will need to use the low-level commands to create subdisks, plexes, and volumes exactly as required.

8.19 LSM volprint Command Format

In the next major release of Tru64 UNIX, the default output format of the `volprint` command will be changing.

Invoking `volprint` with no options currently displays all objects in `rootdg`, starting with the `rootdg` disk group record, followed by all of the disk-media records, subdisk records, plex records, and finally volume records. In a future release, invoking `volprint` with no options will display all records for all disk groups, with all objects arranged in a hierarchical fashion.

Invoking `volprint` with an object-type option (`v`, `p`, or `s`) currently displays all objects of that type in the specific disk group. In a future release, invoking `volprint` with any of these options will display all objects of that type, as well as all objects of all subsidiary types, in all disk groups.

By default, invoking `volprint` currently displays the record type, name, association, kernel state, length, and comment field. In a future release, it will display the record type, name, association, kernel state, length, plex offset, state, and the `tutil0` and `putil0` fields.

Users who use `volprint` in scripts should use the `-F` option to define the exact output format that they require.

8.20 LVM-to-LSM Migration Tool

The LVM-to-LSM Migration Tool was provided with Version 4.0 of the operating system to enable migration from the LVM interfaces that were retired in that release to Tru64 UNIX Logical Storage Manager volumes. This tool will become obsolete in later releases of Tru64 UNIX because most customers will have migrated to LSM volumes by that time; therefore, the LVM-to-LSM Migration Tool will be retired in a future release of the operating system.

There are no plans to retire UFS or the AdvFS Migration Tools at this time.

8.21 OSF/Motif Version 1.1.3

The Motif Version 1.1.3 libraries have been provided as run-time services for compatibility with applications that have not yet converted to Motif 1.2. Development support was retired in DEC OSF/1 Version 2.0.

In Version 4.0 of the operating system, the Motif 1.1.3 libraries were moved to an optional subset. Applications requiring the libraries will see an error from the loader and you must install the optional subset. This optional subset will be removed from the product in a future release.

8.22 XIE Version 3.0 X Client Extension

Tru64 UNIX Version 4.0F supports XIE Version 5.0. Support for XIE Version 3.0 server extensions was removed in Version 4.0, but client support will not be removed until a later release of Tru64 UNIX.

8.23 Microsoft Sound Board Driver

In Version 4.0 of the operating system, the device driver for the base audio on the DIGITAL AlphaStations and DIGITAL AlphaServers was removed from the base operating system. This device driver supported the Microsoft Sound Board (MSB), the AlphaStation Sound Card, and the built-in audio hardware shipped with certain AlphaStation systems.

The driver binaries are now available as part of the Multimedia Services for Tru64 UNIX kit on the *Tru64 UNIX 4.0F Associated Products, Volume 1* CD–ROM in the MMEDRVMSB subset.

The following files will be removed from the base operating system:

- /usr/sys/BINARY/msb.o

- /usr/sys/include/io/dec/eisa/msb.h

- /usr/sys/include/io/dec/eisa/msb_reg.h

You can also get support for this device from the Multimedia Services for Tru64 UNIX kit that is located on the Software Products Library CD-ROM. Support is also factory-installed on all AlphaStation Tru64 UNIX packaged systems. The license for this product is bundled with AlphaStations so that you can use it at no additional cost.

8.24 Graph Utility

The /usr/bin/graph utility will be removed in the next major release of
Tru64 UNIX.

8.25 The atmsetup Script

The atmsetup script is a new utility with Version 4.0D that allows you to
set up or change the current configuration of ATM on your system.

This script was designed to be an interim solution to simplify the setup
process for ATM. In future releases, it will be supplemented and finally
replaced by an application in the SysMan suite, with a full graphical user
interface.

For more information about how to use atmsetup, see the atmsetup(8)
reference page and the *Asynchronous Transfer Mode* guide.

8.26 Remote Prestoserve Support

In the next major release of Tru64 UNIX, PrestoServe support for remote
operations will be retired. This means that the use of the -h option of the
presto command, the dxpresto *hostname* variable, and the -n option
of the prestoctl_svc command will no longer be supported.

Users who require the ability to perform presto operations remotely can
still do so by using telnet, rlogin, or rsh to gain access to a shell on
the remote system and then to perform the operation locally.

8.27 Disk Size Requirement for Installation

The minimum disk size requirement for installing the Tru64 UNIX operating
system is changing in the next major release due to additional features and
services that are being added. The minimum disk size required to support a
single-disk operating system installation will be 1GB (such as an RZ26) for
both default and custom installation types.

8.28 Installupdate -i Option

The -i option to the /sbin/installupdate command will be retired in
a future release of the operating system.

The -i option currently allows you to interactively select kernel components
once the new software subsets have been installed. Starting with the next
major release, this flag will be unnecessary because you will be able to
interactively select optional kernel components at the beginning of the update
installation process, prior to software installation. These kernel components

will be built into the kernel automatically during the kernel build phase at the
end of the update installation; therefore, you need not be present at that time.

8.29 Different awk Versions

Tru64 UNIX contains three different versions of the awk command:

- oawk - Old awk is the original awk command.
- gawk - GNU awk is the Free Software Foundation's version.
- nawk - Is the new version of the awk command and is current linked to
  awk. This version of the awk command is the XPG4-compliant version
  of the awk command.

The oawk and gawk versions of the awk command will be removed in a
future version of the operating system. Also, the link between awk and nawk
will be removed, leaving awk and removing nawk. Users should ensure that
their scripts use /usr/bin/awk instead of any other version of the awk
command currently existing on the system.

8.30 The secsetup Script

The /usr/sbin/secsetup script, used to configure Enhanced Security
authentication support, will retire in a future release of Tru64 UNIX. It will
be replaced by sec_config , an application in the SysMan suite that
provides a full graphical user interface.

8.31 Open3D TURBOchannel and PVpci Device Support

Support for the following Open3D TURBOchannel devices has been retired
from this release of Tru64 UNIX (code is shipped as unsupported):

- HX+
- PV-L
- PV-M

These devices are also known as:

- ZLX-E1, E2, E3
- ZLX-L1, L2
- ZLX-M1, M2

Open3D V4.4 is the last version that supports these graphic adapters. The
DDX and device drivers will ship in future releases of Tru64 UNIX, but
support will be retired and Open3D will no longer respond to problem

reports. Tru64 UNIX 2D support will continue for future Tru64 UNIX V4.0n releases and the initial release of Tru64 UNIX Version 5.0 for the TX, HX, and HX+ graphic adapters.

The replacement functionality is PCI based systems and graphics devices supported in those systems.

In addition, the Open3D PVpci (ZLXp-L1 and ZLXp-L2) graphic adapter support ended in December 1998. The last version of Open3D to support these cards was Open3D Version 4.9. The last version of Tru64 UNIX to support these adapters is Tru64 UNIX Version 4.0F. The replacement functionality is the PowerStorm 4Dxx (Intergraph Cateyes) graphic devices.

8.32 System Management Utilities

The following setup scripts will be retired in a future release:

- `bindsetup`
- `mailsetup`
- `netsetup`
- `nfssetup`
- `ntpsetup`

These scripts have been replaced by tasks that you can perform by running the SysMan menu and navigating to the desired item or directly from the command line by using the appropriate menu accelerator. The following table maps the replacement functionality and accelerator.

Retired Script	SysMan Menu Item	SysMan Accelerator
`bindsetup`	Configure system as a DNS client	`dns_client`
	Configure system as a DNS server	`dns_server`
`mailsetup`	Configure the host's mail system	`mailsetup`
`netsetup`	Network Configuration Step-by-Step	`net_step`
`nfssetup`	Configure System as an NFS client	`nfs_client`
	Configure System as an NFS server	`nfs_server`
`ntpsetup`	Configure system as an NTP client	`ntp_config`

For more information, see the reference pages or the *Network Administration* manual.

8.33 Changed Protected Password Database Format

The enhanced security user authentication profiles, originally stored in the protected password files `/tcb/files/auth/<a..z>/<username>`, are now stored by default in a database format. The `convauth` program, run for an update installation, will convert existing files into the new format. Support for the original format will be removed in a future release of Tru64 UNIX.

8.34 Audit_setup Script

The `/usr/sbin/audit_setup` script, used to configure auditing support, will be retired in a future release of Tru64 UNIX. It will be replaced by `audit_config`, an application in the SysMan suite that provides a full graphical user interface.

8.35 VM Subsystem ubc_nfsloopback Variable Is Obsolete

The `/etc/sysconfigtab ubc_nfsloopback` Virtual Memory (VM) subsystem kernel configuration variable is unnecessary, and has been removed.

This configuration variable was needed in some TruCluster configurations with previous versions of Tru64 UNIX to avoid a recursion deadlock with NFS loopback-mounted file systems. In those configurations, setting the `ubc_nfsloopback` VM subsystem variable to 1 prevented the deadlock.

The file system code in Tru64 UNIX Version 4.0F has been modified to prevent the deadlock which removes the need for the `ubc-nfsloopback` variable.

If the variable is not removed from `/etc/sysconfigtab`, it will not cause a problem. However, after the kernel is loaded by the firmware, the console generates the following message to indicate that the variable is still present in `/etc/sysconfigtab`:

```
sysconfigtab:  attribute ubc_nfsloopback not in subsystem VM
```

8.36 libsecurity.a

`libsecurity.a` is the static archive form of the enhanced security library. This form of `libsecurity` is no longer delivered in Tru64 UNIX. Customers wishing to link against `libsecurity` should instead use the shareable form, `libsecurity.so`.

8.37 XPG3 ctab Locales

The XPG3 ctab locales in the obsolete OSFCTABLOC subset will be retired
in a future release. These locales were superseded by the XPG4 locales
distributed in OSFEURLOC and various subsets in the Worldwide Language
Support software on the Associated Products CD (APCD).

8.38 Security Interfaces

Several interfaces in libsecurity will be retired in a future release of
Tru64 UNIX. These interfaces are not capable of supporting new
functionality offered in Tru64 UNIX Version 4.0F.

Binary compatibility will be preserved until the interfaces are retired, which
will be no sooner than 12 months after the current release is shipped. Table
8-3 provides the complete list of documented interfaces. Some
undocumented interfaces might be retired at the same time.

Table 8-3: libsecurity Interfaces Scheduled for Retirement

```
getprtcent()      getprtcnam()       putprtcnam()
getprdfent()      getprdfnam()       putprdfnam()
getprfient()      getprfinam()       putprfinam()
getprlpent()      getprlpnam()       putprlpnam()
getdvagent()      getdvagnam()       putdvagnam()
getprpwent()      getprpwnam()       getprpwuid()
putprpwnam()
time_lock()       get_seed()         auth_for_terminal()
locked_out()      copydvent()
read_pw_fields()  store_pw_fields()
read_tc_fields()  store_tc_fields()
```

Table 8-4 provides the list of associated data structures that will be retired at
the same time as the interfaces.

**Table 8-4: libsecurity Data Structure Interfaces Scheduled for
              Retirement**

```
struct pr_field              struct pr_flag
struct t_field               struct t_flag
struct l_field               struct l_flag
struct dev_field             struct dev_flag
struct pr_passwd             struct pr_term
struct pr_file               struct pr_lp
```

Table 8-4: (continued)

```
struct pr_default              struct dev_asg
struct system_default_fields   struct system_default_flags
```

8.39 The ogated Routing Daemon

The `ogated` (the old version of the `gated` routing daemon) will be retired in a future release of Tru64 UNIX. If you use the `ogated` routing daemon, you should migrate to the `gated` routing daemon, which supports a superset of functionality in `ogated`.

8.40 PEX Extensions

Starting with Open3D Version 4.8, support for the PEX extension to the X Server has been retired. The extension libraries still ship on the Open3D product, but no support is available for them.

8.41 NL* and NL* libc Interfaces

A Worldwide Portability Interface (WPI) for internationalization based on the XPG4 standard was introduced in DIGITAL UNIX Version 2.0. Because the WPI interfaces supercede the `libc` OSF/1 proprietary interfaces, all `libc` interfaces that begin with the letters `NL` or `NC` will be removed in a future release of the operating system.

8.42 DEC Ada RTL

DEC Ada (UPI - 0HM) and DEC Ada PDO (UPI - 0VS) will be retired in a future release of Tru64 UNIX.

8.43 System V Environment

The System V Environment product (QL-0J6A*-**) will be reengineered and its functionality will be merged into the Tru64 UNIX base operating system in a future major release of the operating system. Today, Tru64 UNIX already delivers 80 percent of the System V Interface Definition (SVID) standard, as verified by the SVVS 3 and SVVS 4 test suites. As a result, Tru64 UNIX contains a substantial number of System V Release 4 (SVR4) features and delivers the highest composite SVR4 conformance of any implementation. SVR4 functionality will be further expanded in the base operating system when the System V Environment reengineering is complete,

eliminating the need for the layered product. A migration plan for upgrading to the appropriate version of the Tru64 UNIX base operating system has been developed to assist our customers who currently use the System V Environment layered product.

System Limits   A

_____

A.1   System Limits

This appendix lists the system limits for the major components of this release. For hardware information specific to your individual processor, see the Tru64 UNIX *Software Product Description* (SPD) and the *Systems & Options Catalog*.

You may be able to increase some of the system limits by changing kernel attribute values. Use the sysconfig command to display the current attribute values, and the maximum and minimum values. For information on how to modify attributes, see the *System Administration* manual and the *System Configuration and Tuning* manual, which also includes lists of tunable attributes.

The following sections describe the system limits.

A.1.1   Installation Requirements

A default installation of the Tru64 UNIX operating system requires at least 525 MB of disk space (for example, a RZ25 disk). This limit will change in a future release (see Section 8.27 for more information). The disk space that a custom installation requires depends on the optional subsets that you will install. See Appendix B for information about subset sizes.

A.1.2   Memory Limits

The memory limits are as follows:

- Physical memory requirements

  The minimum amount of physical memory that Tru64 UNIX requires is 64 MB.

  The system platform limits the maximum amount of physical memory. For more information on platform memory support, see the *Systems & Options Catalog*.

- Virtual address space limit

  The default amount of virtual address space available for each process is

1 GB. However, in many cases, available swap space may be exhausted before this limit is reached.

You can increase the available virtual address space to a maximum of 4 TB by modifying the value of the `vm-maxvas` attribute in the `vm` subsystem. For more information about setting this attribute, see the *System Configuration and Tuning* guide.

- Page size

  The physical page size is 8 KB and cannot be changed. The page size is hardware dependent and is set by the console at boot time.

A.1.3 Process Limits

The process limits for the system and for users are as follows:

- Per system

  The `maxusers` attribute in the `proc` subsystem specifies the number of simultaneous users that a system can support without straining system resources. System algorithms use the `maxusers` attribute to size various system data structures, and to determine the amount of space allocated to system tables, such as the system process table, which is used to determine how many active processes can be running at one time.

  The default value assigned to the `maxusers` attribute depends on the size of your system. The default value is 32 users on systems with at least 64 MB of physical memory. The minimum value is 8 and the maximum value is 4096.

  The maximum number of tasks that can run simultaneously on a system is determined by the value of the `task-max` attribute in the `proc` subsystem. The default value is 8213. The minimum value is 85, and the maximum value is 32768.

  The maximum number of kernel threads that can run simultaneously on a system is determined by the value of the `thread-max` attribute in the `proc` subsystem. The default value is 16424. The minimum value is 160, and the maximum value is 262136.

- Per user

  The maximum number of processes that a user can create is determined by the value of the `max-proc-per-user` attribute in the `proc` subsystem. The default value is 64. The minimum value is 0, and the maximum value is 32767.

  The maximum number of threads that a user can create is determined by the value of the `max-threads-per-user` attribute in the `proc` subsystem. The default value is 256. The minimum value is 0, and the maximum value is 18,446,744,073,709,551,615 ($2^{64} - 1$).

A.1.4 Device Addressing Limits

The limits for device addressing are as follows:

- Device access

  There are two types of disk device access: raw (character) and buffered (block).

  For raw or character access, the uio.uio_offset structure field describes the byte offset within the disk partition. In this release, the uio_offset is an unsigned 64-bit value, allowing an offset up to 2^{64} or 18 exabytes. This value is converted to a physical block/sector number, which is the data transfer start position. The physical block/sector number is limited by the buf.b_blkno structure field.

  For buffered or block access, the buf.b_blkno structure field describes the block/sector offset within the disk partition and is a signed 32-bit value. Because this release supports a fixed 512-byte block/sector size, as defined by DEV_BSIZE, the offset is limited to 1 TB.

- Major-minor numbers (dev_t)

  Devices are identified by a major-minor pair of numbers, where the major number specifies the device driver and the minor number identifies the device. In this release, this pairing is represented by a 32-bit value described by the type dev_t. The major number portion of dev_t consists of bits 20 to 31 (a total of 12 bits). Because each device driver requires 12 bits for its major number, you can configure only 4096 device drivers into the system.

  The minor number portion of dev_t consists of bits 0 to 19 (a total of 20 bits). The device driver is responsible for interpreting these bits. A device driver that utilizes all 20 bits for device addressing can address up to 1048576 devices for each major number. For device drivers that support disk devices, some of the bits in the minor device number represents the partition number. This release requires disk drivers to reserve the lower 6 bits for device attributes and partition numbers, and only supports eight partitions.

- SCSI/CAM addressing

  Common Access Method (CAM) is an ANSI-proposed standard for a common software interface to the Small Computer Systems Interface (SCSI). There are no restrictions or limitations within CAM for disk block addressing, because the address is an incoming value.

  For SCSI-2, the Command Descriptor Block (CDB) defines the starting disk block number for the transfer. In this release, the 10-byte CDB reserves 4 bytes for the disk block address. This is an unsigned 32-bit value that allows $2^{32} - 1$ or 4 gigasectors of addressing. Using a 512-byte block/sector size, this value corresponds to 2 TB.

In this release, the SCSI/CAM driver can address a maximum of 64 buses, with up to 7 device targets for each bus, and a maximum of eight logical unit numbers (LUNs) for each device target. According to these limitations, SCSI/CAM can address a maximum of 3584 devices.

- Redundant Array of Independent Disks (RAID)

  Refer to the following for complete information:

  – *Systems & Options Catalog*

  – Product description literature

  – Web pages:

    ```
    http://www.digital.com/products.html
    http://www.storage.digital.com
    ```

- Disklabel command

  The `disklabel` command specifies the partitions of a disk and their starting block/sector number. The starting block/sector number of a partition is defined by the `partition.p_offset` structure field, which is an unsigned 32-bit value that supports up to 2 TB of addressing, using a 512-byte block/sector size.

A.1.5 Device Limits

Tru64 UNIX supports a maximum of seven devices per SCSI bus (for example, one host bus adapter and seven disks). The maximum number of SCSI buses supported by a system depends on the platform, but the Tru64 UNIX operating system supports a maximum of 64 SCSI buses.

For all other device limits, see the *Systems & Options Catalog*.

A.1.6 Logical Storage Manager Limits

In this release, the Logical Storage Manager (LSM), supports a maximum of 768 disk groups and 256 disks, either in a disk group or across the system.

The LSM term **volume** refers to a virtual disk that represents an addressable range of disk blocks. File system data or raw I/O can be directed to a volume. This release supports a maximum of 1 TB of disk space in a disk group or on a system, and a maximum volume size of 1 TB. The maximum number of supported LSM volumes is 4093 for all disk groups in a system, which includes 4091 nonsystem volumes and 2 system (root or swap) volumes.

The LSM term **plex** refers to a physical disk or a set of disks that contain a complete copy of a volume's data. A mirrored volume consists of at least two plexes. In this release, the maximum number of supported plexes per volume is 8, and the maximum number of supported plexes per system is

4093 (or 4091, if root and swap volumes are not used).

The LSM term **subdisk** refers to a contiguous portion of a physical disk that can be striped or concatenated with other subdisks to form a plex. A maximum of 4096 subdisks can be associated with one plex, and Compaq supports 4096 subdisks in each disk group or on a system.

LSM object names (for example, volumes, plexes, subdisks, and disk groups), volume attribute names (such as user and group), and `dxlsm view` names are limited to 14 characters.

A.1.7   File System Limits

The file system limits are as follows:

- Advanced File System

  The Advanced File System (AdvFS) term **volume** refers to a single logical device, such as a disk, disk partition, or a logical volume. A **file domain** is a named set of one or more volumes. A **fileset** is a named and mountable logical file structure that is created in a domain. An **active fileset** is a fileset that is currently mounted.

  AdvFS supports a maximum fileset and file size of 16 TB minus 512 K, a maximum of 255 active file domains for each system, and a maximum of 256 volumes for each domain. However, because a single disk failure in a multivolume file domain can make the entire domain inaccessible, Compaq recommends that you use a maximum of eight volumes in a file domain.

  Although Tru64 UNIX supports an unlimited number of filesets per system, only 767 filesets can be mounted at one time. The maximum number of files in a fileset is 2^{31} and is limited by the tag that is used to uniquely identify a file in a fileset. Because of a sequence number limit, a tag can be used only 4096 times; therefore, the actual limit on files in a fileset decreases over time.

  Although AdvFS can support a page size larger than 13 bits, the maximum size of an AdvFS file and fileset is 16 TB - 512 K ($2^{13} * 2^{31}$), with a 13-bit page size and a 31-bit page number.

  AdvFS supports a maximum of 767 mounted filesets. However, each active file domain has a hidden mounted fileset that must be counted when determining the total number of mounted filesets. For example, if you have an active file domain with two visibly mounted filesets, the file domain actually has three mounted filesets.

  The AdvFS `verify` utility cannot validate domains with more than 767 filesets.

- UNIX File System

  In this release, the UNIX File System (UFS) file size is limited by the amount of space that can be addressed by the kernel `buf` structure. The `buf.b_blkno` structure field, defined as `daddr_t`, is a 32-bit signed value, and specifies the block/sector offset within a disk partition. The `DEV_BSIZE`, block or sector size, is 512 bytes. Theoretically, a UFS file system could be 1 TB ($2^{13} * 2^9$); however, Tru64 UNIX supports only 128 GB. The maximum LSM logical volume size also limits a UFS file system and file size.

  Tru64 UNIX supports up to 2,147,483,647 UNIX File System (UFS) and Memory File System (MFS) mounts. The `max-ufs-mounts` attribute controls the maximum number of UFS and MFS mounts. The default value is 1000.

- CD–ROM File System

  The size of ISO 9660 CD–ROM File System (CDFS) files and file systems is limited by the CD–ROM media on which they reside. Currently, the CD–ROM media supports approximately 600 MB. However, Tru64 UNIX is able to support larger ISO 9660 file systems on other media.

  This release supports a maximum of 512 CDFS mounts.

- Sparse files

  Tru64 UNIX supports **sparse files** on AdvFS and UFS; therefore, the logical size of a file (that is, the offset to the last data actually present) can exceed the size of the file system in which it resides. The maximum sizes for sparse files are as follows:

  – AdvFS – $2^{44} - 512K$

  – UFS – $2^{44} - 8$ K

- Network File System

  In this release, the theoretical maximum sizes of files that are accessible through Network File System (NFS) Version 2 and Version 3 are as follows:

  – NFS Version 2 – (2 GB – 1 byte)

  – NFS Version 3 – 18 exabytes ($2^{64} - 1$)

  However, Tru64 UNIX supports the following maximum file sizes:

  – NFS Version 2 – 2 GB

  – NFS Version 3 – 512 GB

  In addition, an NFS file system is always limited by the size of the local file system that is being exported.

An NFS Version 2 or Version 3 client can mount a maximum of 2048 files or directories.

- Memory mapped file limit

  The maximum supported size of a file that can be mapped into memory without segmenting the file depends on the virtual address space limits, as documented in the section on memory limits.

- Open files limit

  The `open-max-soft` and `open-max-hard` attributes control the maximum number of open file descriptors per process for all processes. The default value is 4096. When the `open-max-soft` limit is reached, a warning message is issued, and when the `open-max-hard` limit is reached, the process is stopped. The maximum number of open files per process is 65,536. If you increase the maximum number of open files per process, make sure that you adjust the value of the `max-vnodes` attribute.

  File descriptor entries in the per process file table are dynamically allocated after the initial 64 entries in the `utask` structure are used.

- File locking limits

  The Tru64 UNIX file record locking service allows applications to lock any number of bytes in a file in the range of 0 to $2^{63} - 1$, inclusive. File locking is supported by UFS, AdvFS, and both NFS Version 2 and Version 3. Because the NFS Version 2 protocol suite allows ranges to be specified only with 32-bit numbers, it supports a file locking range of 0 to $2^{31} - 1$, inclusive.

- Pathname limits

  AdvFS, UFS, CDFS, and NFS support a maximum pathname component of 255 bytes and a maximum file pathname of 1023 bytes.

A.1.8 Networking Limits

The networking limits are as follows:

- Pseudoterminals (`ptys`)

  The maximum number of supported `ptys` is 8192.

- LAT connections

  - The maximum number of incoming LAT connections is 4800.

  - The maximum number of outgoing LAT connections is 4000.

- IP alias addresses

  Tru64 UNIX allows the use of up to 5,000 IP alias addresses before

system performance begins to degrade.

- Packetfilter limits

  The packetfilter pseudo-driver can support up to 255 simultaneous open filters (each filter is usually mapped to one instance of an application program). The packetfilter can support a maximum of 255 devices. Use the `pfconfig` command to configure packet filters.

- Network transfer rates

  For information on network transfer rates, see the *Technical Overview*.

A.1.9 Limits for Backup Utilities

The limits for the backup utilities are as follows:

- `cpio`

Files per archive:	No limit
Files per file system:	No limit
File size:	4 GB
File name size:	256 bytes

- `dd`

Files per archive:	Not used
Files per file system:	Not used
File size:	4 GB
File name size:	Not used

- `dump`

Files per archive:	4,000,000,000
Files per file system:	4,000,000,000
File size:	4 GB
File name size:	No limit (part of the inode data)

- `vdump`

Files per archive:	maximum files per fileset
Files per file system:	file system dependent
File size:	file system dependent
File name size:	1024 bytes (with prefix)

- `tar`

Files per archive:	No limit
Files per file system:	No limit
File size:	8 GB (4 TB with extended option)
File name size:	1024 bytes (with prefix)

- `pax`

Files per archive:	No limit
Files per file system:	No limit
File size:	8 GB (4 TB with extended option)
File name size:	1024 bytes (with prefix)

Software Subset Information   B

_____

This appendix provides the sizes of all Tru64 UNIX software subsets for full, update, and RIS installations.

B.1 Disk Space Required for Software Subsets

Table B-1, Table B-2, and Table B-3 show disk space as the number of 512-byte blocks required in the `root`, `/usr`, and `/var` file systems to install each Tru64 UNIX software subset. The figures for each group of files within a subset have been rounded up to the next higher 512-byte increment; this means that the total space requirements listed are slightly greater than the space actually required.

To determine the subset size in megabytes (MB), divide the size in blocks by 2048.

For information on the contents of each subset, refer to the *Installation Guide*. If you want to add optional subsets after you install Tru64 UNIX Version 4.0F, use the `df` command to determine free disk space in blocks.

Table B-1: Tru64 UNIX Version 4.0F Operating System, Volume 1

Tru64 UNIX Version 4.0F Operating System

Subset	root	/usr	/var	Total
OSFACCT440	9.67	1077.52	91.03	1178.22
OSFADVFS440	6745.80	6982.87	—	13728.66
OSFADVFSBIN440	2181.95	3.06	—	2185.00
OSFADVFSBINOBJECT440	—	3441.70	—	3441.70
OSFAFM440	—	2160.73	—	2160.73
OSFATMBASE440	241.78	799.84	—	1041.61
OSFATMBIN440	4574.03	21.40	—	4595.44
OSFATMBINCOM440	—	423.57	—	423.57
OSFATMBINOBJECT440	—	16009.66	—	16009.66
OSFBASE440	24438.53	82419.21	848.09	107705.84
OSFBIN440	12819.52	1335.50	—	14155.02
OSFBINCOM440	45.47	28843.91	68.98	28958.36
OSFBINOBJECT440	—	19584.51	—	19584.51
OSFC2SEC440	861.78	1360.56	84.00	2306.34
OSFCDAPGMR440	—	3318.51	—	3318.51

Table B-1: (continued)

Tru64 UNIX Version 4.0F Operating System

Subset	root	/usr	/var	Total
OSFCDEAPPS440	—	11279.12	—	11279.12
OSFCDEDEV440	—	25497.76	—	25497.76
OSFCDEDT440	—	53143.40	—	53143.40
OSFCDEMAIL440	—	3370.60	—	3370.60
OSFCDEMANOP440	—	2055.95	—	2055.95
OSFCDEMANOS440	—	1425.27	—	1425.27
OSFCDEMIN440	—	13959.90	76.00	14035.90
OSFCLINET440	864.50	17363.30	41.00	18268.80
OSFCMPLRS440	—	25330.68	—	25330.68
OSFCTABLOC440	34.25	359.30	—	393.55
OSFDCMT440	—	1107.81	—	1107.81
OSFDCMTEXT440	—	4368.19	—	4368.19
OSFDECW440	—	1781.22	47.20	1828.42
OSFDMS440	—	87.26	73.00	160.26
OSFDOSTOOLS440	—	1005.00	—	1005.00
OSFEMACS440	—	39543.50	—	39543.50
OSFENVMON440	21.25	152.62	—	173.87
OSFEURLOC440	—	1811.22	—	1811.22
OSFEXAMPLES440	—	1995.58	—	1995.58
OSFEXER440	—	4107.00	—	4107.00
OSFFONT15440	—	3160.99	—	3160.99
OSFFONT440	—	2432.85	—	2432.85
OSFHWBASE440	23075.41	4266.22	34.40	27376.03
OSFHWBIN440	25718.38	2088.85	—	27807.22
OSFHWBINCOM440	—	3470.83	—	3470.83
OSFHWBINOBJECT440	—	22174.51	—	22174.51
OSFIMXE440	16.06	6685.81	279.68	6981.55
OSFINCLUDE440	—	5425.84	—	5425.84
OSFINET440	678.33	15188.52	538.73	16405.58
OSFJAVA440	—	21731.88	—	21731.88
OSFJAVADEV440	—	17695.53	—	17695.53
OSFJAVADOC440	—	33844.67	—	33844.67
OSFKBDLK201440	—	361.70	—	361.70
OSFKBDLK401440	—	248.44	—	248.44
OSFKBDLK411440	—	134.33	—	134.33
OSFKBDLK421440	—	16.42	—	16.42
OSFKBDLK444440	—	126.52	—	126.52
OSFKBDPCXAL440	—	273.12	—	273.12
OSFKTOOLS440	—	1486.21	7125.71	8611.92
OSFLAT440	516.59	923.71	7.82	1448.12
OSFLDBBASE440	—	13554.49	—	13554.49
OSFLDBDOC440	—	547.43	—	547.43
OSFLDBGUI440	—	7666.76	—	7666.76
OSFLDBSRV440	—	83.00	—	83.00

Table B-1: (continued)

Tru64 UNIX Version 4.0F Operating System

Subset	root	/usr	/var	Total
OSFLEARN440	—	3097.67	—	3097.67
OSFLIBA440	—	6555.07	—	6555.07
OSFLSMBASE440	5377.42	3038.92	48.27	8464.60
OSFLSMBIN440	542.96	3.05	—	546.01
OSFLSMBINCOM440	—	461.96	—	461.96
OSFLSMBINOBJECT440	—	555.91	—	555.91
OSFLSMX11440	—	2267.79	47.90	2315.69
OSFLVM440	—	2917.42	—	2917.42
OSFMANOP440	—	17037.07	—	17037.07
OSFMANOS440	—	14714.16	—	14714.16
OSFMANWOP440	—	8872.17	—	8872.17
OSFMANWOS440	—	1430.32	—	1430.32
OSFMH440	—	4416.17	—	4416.17
OSFMITFONT440	—	18773.90	104.01	18877.91
OSFMOTIF11440	—	12847.09	—	12847.09
OSFNETCONF440	—	2307.99	—	2307.99
OSFNETSCAPE440	—	49357.37	—	49357.37
OSFNFS440	50.21	1212.45	—	1262.67
OSFNFSCONF440	—	337.90	—	337.90
OSFOBSOLETE440	—	3375.14	—	3375.14
OSFOLDDECW440	—	14961.35	—	14961.35
OSFOLDX11440	—	1410.37	—	1410.37
OSFPGMR440	—	9258.06	—	9258.06
OSFPRINT440	112.77	8822.86	19.00	8954.63
OSFRCS440	—	1901.84	—	1901.84
OSFRIS440	—	189.63	143.00	332.63
OSFSCCS440	—	9976.66	—	9976.66
OSFSDE440	—	14540.53	—	14540.53
OSFSDECDE440	—	284.69	—	284.69
OSFSER3D440	—	12130.00	—	12130.00
OSFSER440	—	12641.99	28.32	12670.32
OSFSERPC440	—	3963.00	—	3963.00
OSFSERTC440	—	678.00	—	678.00
OSFSVID2440	31.59	615.95	—	647.54
OSFSYSMAN440	27.68	14214.09	—	14241.77
OSFTCLBASE440	—	3415.12	—	3415.12
OSFTERM440	—	3685.73	—	3685.73
OSFTKBASE440	—	5433.76	—	5433.76
OSFUUCP440	101.73	11768.01	266.00	12135.74
OSFX11440	33.40	43990.18	698.13	44721.71
OSFXADMIN440	—	7786.35	73.19	7859.54
OSFXC2SEC440	—	1272.02	—	1272.02
OSFXCDADEV440	—	523.15	—	523.15
OSFXDEMOS440	—	1854.78	—	1854.78

Table B-1: (continued)

Tru64 UNIX Version 4.0F Operating System

Subset	root	/usr	/var	Total
OSFXDEV440	—	2563.31	—	2563.31
OSFXEXAMPLES440	—	9228.08	—	9228.08
OSFXIEDOC440	—	1478.88	—	1478.88
OSFXINCLUDE440	—	7638.01	—	7638.01
OSFXLIBA440	—	15251.14	—	15251.14
OSFXMIT440	—	4177.32	—	4177.32
OSFXNEST440	—	390.00	10.11	400.11
OSFXOEM440	—	0.00	965.87	965.87
OSFXPRINT440	—	336.42	—	336.42
OSFXSYSMAN440	—	13565.21	152.74	13717.95
OSFXVFB440	—	246.00	10.11	256.11
Total	109121.05	934887.90	11882.30	1055891.24

Grand Totals	root	/usr	/var	Total
	109121.05	934887.90	11882.30	1055891.24

Table B-2: Tru64 UNIX Version 4.0F Associated Products, Volume 1

Alternative Development Environment Tools for Tru64 UNIX

Subset	root	/usr	/var	Total
CMPDEVALT440	14.29	14210.74	—	14225.03
CMPDEVENH440	—	12621.42	—	12621.42
Total	14.29	26832.16	—	26846.45

DEC Ada Runtime Library for Tru64 UNIX

Subset	root	/usr	/var	Total
ADALIB404	—	1102.86	—	1102.86
Total	—	1102.86	—	1102.86

DEC C++ Class Libraries Version 4.0 for Tru64 UNIX

Subset	root	/usr	/var	Total
CXLLIBA440	—	224.93	—	224.93
CXLSHRDA440	—	284.12	—	284.12
Total	—	509.04	—	509.04

Decimal RTL Version 2.4 for Tru64 UNIX

Subset	root	/usr	/var	Total
DCARTL250	—	2746.22	—	2746.22
O2ABASE250	—	2648.56	—	2648.56
Total	—	5394.79	—	5394.79

DECevent

Subset	root	/usr	/var	Total
DIABASE290	17.24	61339.61	25.17	61382.03
Total	17.24	61339.61	25.17	61382.03

DEC Fortran RTL

Subset	root	/usr	/var	Total
DFARTL369	—	4959.12	—	4959.12
Total	—	4959.12	—	4959.12

DEC Pascal RTL Version 5.4 for Tru64 UNIX Systems

Subset	root	/usr	/var	Total
DPORTL541	—	1890.15	—	1890.15
Total	—	1890.15	—	1890.15

Sort Runtime Library

Subset	root	/usr	/var	Total
SORLIB300	—	717.85	—	717.85
Total	—	717.85	—	717.85

INTERSOLV DataDirect

Subset	root	/usr	/var	Total
DAUDOC100	—	0.00	9546.46	9546.46
DAUJDBCODBC100	—	0.00	496.95	496.95
DAUODBCCON100	—	0.00	10217.61	10217.61
DAUSQLNKJAVA100	—	0.00	4738.95	4738.95
DAUSQLNKODBC100	33.32	—	5512.07	5545.39
DAUSQLNKSVR100	—	0.00	41441.10	41441.10
Total	33.32	—	71953.15	71986.47

Free Software Foundation GNU Source for Tru64 UNIX

Subset	root	/usr	/var	Total
FSFEMACSSRC440	—	55742.88	—	55742.88
FSFGAWKSRC440	—	1220.38	—	1220.38
FSFGZIPSRC440	—	1894.79	—	1894.79
FSFINDENTSRC440	—	1244.68	—	1244.68
FSFMKISOFSSRC440	—	975.94	—	975.94
FSFRCSSRC440	—	1907.47	—	1907.47
Total	—	62986.13	—	62986.13

Multimedia Services Version 2.4D for Tru64 UNIX

Subset	root	/usr	/var	Total
MMEDRVAV201244	—	900.80	—	900.80
MMEDRVAV300244	—	2240.32	—	2240.32
MMEDRVAV3X1244	—	1552.02	—	1552.02
MMEDRVBBA244	—	303.84	—	303.84
MMEDRVMMSESS244	—	635.46	—	635.46
MMEDRVMSB244	—	919.63	—	919.63
MMEMANRT244	—	323.40	—	323.40
MMERELNOTES244	—	1401.76	—	1401.76
MMERT244	20.97	21745.40	11.82	21778.19
MMERTCDE244	—	467.70	—	467.70
MMERTSMPLDAT244	—	12385.79	—	12385.79
Total	20.97	42876.13	11.82	42908.92

Netscape FastTrack Version 3.01 for Tru64 UNIX

Subset	root	/usr	/var	Total
WEBNETSCAPEFASTTRACK301	—	172955.13	—	172955.13
Total	—	172955.13	—	172955.13

Tru64 UNIX Version 4.0F X Window System PanoramiX Extension ADK

Subset	root	/usr	/var	Total
PRXADK200	—	16459.98	—	16459.98
Total	—	16459.98	—	16459.98

Porting Assistant Version 3.0 for Tru64 UNIX

Subset	root	/usr	/var	Total
PRTBASE300	—	50631.13	—	50631.13
PRTMAN300	—	18.45	—	18.45
Total	—	50649.58	—	50649.58

Graphical Program Analysis Tools Version 2.0 for Tru64 UNIX

Subset	root	/usr	/var	Total
GPABASE206	—	38958.36	—	38958.36
GPAMVIEW206	—	14481.00	—	14481.00
GPAPROFHEAP206	—	21251.91	—	21251.91
GPAPVIEW206	—	2226.05	—	2226.05
Total	—	76917.32	—	76917.32

Tru64 UNIX Worldwide Language Support Version 4.0F

Subset	root	/usr	/var	Total
IOSAACMENU440	—	0.00	—	—
IOSCACDEAPPS440	—	222.45	—	222.45
IOSCACDEDEV440	—	163.24	—	163.24
IOSCACDEDT440	—	1038.04	—	1038.04
IOSCACDEMAIL440	—	84.04	—	84.04
IOSCACDEMIN440	—	4008.92	—	4008.92
IOSCADECW440	—	177.65	—	177.65
IOSCAX11440	—	746.73	—	746.73
IOSCAXDEV440	—	94.68	—	94.68
IOSCSCDEAPPS440	—	215.37	—	215.37
IOSCSCDEDEV440	—	144.24	—	144.24
IOSCSCDEDT440	—	1067.61	—	1067.61
IOSCSCDEMAIL440	—	78.31	—	78.31
IOSCSCDEMIN440	—	495.43	—	495.43
IOSCSDECW440	—	750.62	—	750.62
IOSCSOLDDECW440	—	3542.02	—	3542.02
IOSCSOLDX11440	—	863.35	—	863.35
IOSCSUCSBASE440	—	116.43	—	116.43
IOSCSX11440	—	6559.27	—	6559.27
IOSCSXDEV440	—	94.57	—	94.57

Tru64 UNIX Worldwide Language Support Version 4.0F

Subset	root	/usr	/var	Total
IOSDECDEAPPS440	—	221.40	—	221.40
IOSDECDEDEV440	—	160.91	—	160.91
IOSDECDEDT440	—	1011.98	—	1011.98
IOSDECDEHLP440	—	18215.23	—	18215.23
IOSDECDEMAIL440	—	86.75	—	86.75
IOSDECDEMIN440	—	511.50	—	511.50
IOSDECHX11440	—	253.29	—	253.29
IOSDEDECW440	—	751.84	—	751.84
IOSDEOLDDECW440	—	3548.00	—	3548.00
IOSDEOLDX11440	—	877.02	—	877.02
IOSDEX11440	—	6385.16	—	6385.16
IOSDEXDEV440	—	94.67	—	94.67
IOSELFONT100M440	—	1111.38	—	1111.38
IOSELFONT100P440	—	1080.35	—	1080.35
IOSELFONT75M440	—	951.89	—	951.89
IOSELFONT75P440	—	926.76	—	926.76
IOSELOLDDECW440	—	545.88	—	545.88
IOSELOLFONT440	—	2014.36	37.23	2051.59
IOSELUCSBASE440	—	76.16	—	76.16
IOSELX11440	—	437.43	—	437.43
IOSESCDEAPPS440	—	217.63	—	217.63
IOSESCDEDEV440	—	159.09	—	159.09
IOSESCDEDT440	—	1022.86	—	1022.86
IOSESCDEHLP440	—	22216.91	—	22216.91
IOSESCDEMAIL440	—	85.09	—	85.09
IOSESCDEMIN440	—	523.28	—	523.28
IOSESDECW440	—	733.50	—	733.50
IOSESOLDDECW440	—	3597.96	—	3597.96
IOSESOLDX11440	—	886.37	—	886.37
IOSESX11440	—	6359.98	—	6359.98
IOSESXDEV440	—	94.73	—	94.73
IOSFRBEX11440	—	253.29	—	253.29
IOSFRCAX11440	—	247.47	—	247.47
IOSFRCDEAPPS440	—	212.69	—	212.69
IOSFRCDEDEV440	—	152.82	—	152.82
IOSFRCDEDT440	—	1020.63	—	1020.63
IOSFRCDEHLP440	—	18676.56	—	18676.56
IOSFRCDEMAIL440	—	88.34	—	88.34
IOSFRCDEMIN440	—	524.28	—	524.28
IOSFRCHX11440	—	253.30	—	253.30
IOSFRDECW440	—	725.20	—	725.20
IOSFROLDDECW440	—	3568.04	—	3568.04
IOSFROLDX11440	—	870.65	—	870.65
IOSFRX11440	—	6401.73	—	6401.73
IOSFRXDEV440	—	94.79	—	94.79
IOSHUCDEAPPS440	—	217.01	—	217.01

Tru64 UNIX Worldwide Language Support Version 4.0F

Subset	root	/usr	/var	Total
IOSHUCDEDEV440	—	160.93	—	160.93
IOSHUCDEDT440	—	1056.37	—	1056.37
IOSHUCDEMAIL440	—	80.75	—	80.75
IOSHUCDEMIN440	—	4029.62	—	4029.62
IOSHUDECW440	—	722.65	—	722.65
IOSHUOLDDECW440	—	3567.26	—	3567.26
IOSHUOLDX11440	—	855.98	—	855.98
IOSHUUCSBASE440	—	100.43	—	100.43
IOSHUX11440	—	6415.78	—	6415.78
IOSHUXDEV440	—	94.65	—	94.65
IOSITCDEAPPS440	—	213.17	—	213.17
IOSITCDEDEV440	—	158.90	—	158.90
IOSITCDEDT440	—	1009.87	—	1009.87
IOSITCDEHLP440	—	14581.31	—	14581.31
IOSITCDEMAIL440	—	88.63	—	88.63
IOSITCDEMIN440	—	529.38	—	529.38
IOSITDECW440	—	728.54	—	728.54
IOSITOLDDECW440	—	3579.18	—	3579.18
IOSITOLDX11440	—	891.73	—	891.73
IOSITX11440	—	6438.24	—	6438.24
IOSITXDEV440	—	94.73	—	94.73
IOSIWBASE440	—	186.07	—	186.07
IOSIWCDEDT440	—	242.71	—	242.71
IOSIWCDEMIN440	—	154.81	—	154.81
IOSIWDECW440	—	84.12	—	84.12
IOSIWFONT100M440	—	632.91	—	632.91
IOSIWFONT100P440	—	2157.05	—	2157.05
IOSIWFONT75M440	—	408.37	—	408.37
IOSIWFONT75P440	—	1901.38	—	1901.38
IOSIWOLDX11440	—	11.77	—	11.77
IOSIWOLFONT440	—	3041.34	62.44	3103.78
IOSIWUCSBASE440	—	158.39	—	158.39
IOSIWX11440	—	1704.15	—	1704.15
IOSIWXDEV440	—	930.67	—	930.67
IOSJPABASE440	—	3626.82	—	3626.82
IOSJPAMANOS440	—	35.76	—	35.76
IOSJPBASE440	550.85	17054.78	6.10	17611.72
IOSJPCDEAPPS440	—	475.68	—	475.68
IOSJPCDEDEV440	—	1167.70	—	1167.70
IOSJPCDEDT440	—	3742.31	—	3742.31
IOSJPCDEHLP440	—	32454.40	—	32454.40
IOSJPCDEHLPSJIS440	—	32817.55	—	32817.55
IOSJPCDEMAIL440	—	307.84	—	307.84
IOSJPCDEMIN440	—	1911.39	—	1911.39
IOSJPDECW440	—	857.52	—	857.52
IOSJPFONT100M440	—	12069.76	—	12069.76

Tru64 UNIX Worldwide Language Support Version 4.0F

Subset	root	/usr	/var	Total
IOSJPFONT100P440	—	11912.50	—	11912.50
IOSJPFONT75M440	—	8454.37	—	8454.37
IOSJPFONT75P440	—	8310.46	—	8310.46
IOSJPFONTM440	—	13553.31	—	13553.31
IOSJPLDBBASE440	—	1121.06	—	1121.06
IOSJPLDBGUI440	—	3568.32	—	3568.32
IOSJPMANOS440	—	6815.61	—	6815.61
IOSJPMANWOS440	—	184.39	—	184.39
IOSJPMSG440	—	1597.18	—	1597.18
IOSJPMSGSJIS440	—	532.35	—	532.35
IOSJPNEMACS440	—	23417.63	76.00	23493.63
IOSJPNEMACSSRC440	—	10272.62	—	10272.62
IOSJPNETSCAPE440	—	5967.04	—	5967.04
IOSJPOLDDECW440	—	5526.82	—	5526.82
IOSJPOLDX11440	69.39	956.94	—	1026.33
IOSJPPGMR440	—	2046.36	—	2046.36
IOSJPUCSBASE440	—	17159.06	—	17159.06
IOSJPWNN440	9.97	20149.03	95.00	20254.00
IOSJPWNNPGMR440	—	1213.74	—	1213.74
IOSJPWNNSRC440	—	10794.54	—	10794.54
IOSJPX11440	—	6576.92	—	6576.92
IOSJPXDEV440	—	126.00	—	126.00
IOSKOBASE440	—	1131.49	—	1131.49
IOSKOCDEAPPS440	—	185.78	—	185.78
IOSKOCDEDEV440	—	157.37	—	157.37
IOSKOCDEDT440	—	1339.99	—	1339.99
IOSKOCDEHLP440	—	11156.81	—	11156.81
IOSKOCDEMAIL440	—	130.05	—	130.05
IOSKOCDEMIN440	—	798.52	—	798.52
IOSKODECW440	—	181.13	—	181.13
IOSKOFONTM440	—	3826.33	—	3826.33
IOSKOFONTP440	—	9568.90	—	9568.90
IOSKOOLDDECW440	—	1952.09	—	1952.09
IOSKOOLDX11440	—	433.60	—	433.60
IOSKOOLFONT440	—	6159.31	3.07	6162.38
IOSKOPGMR440	—	168.60	—	168.60
IOSKOUCSBASE440	—	6971.20	—	6971.20
IOSKOX11440	—	4679.19	—	4679.19
IOSKOXDEV440	—	97.87	—	97.87
IOSLTX11440	—	198.84	—	198.84
IOSPLCDEAPPS440	—	199.49	—	199.49
IOSPLCDEDEV440	—	159.54	—	159.54
IOSPLCDEDT440	—	1053.98	—	1053.98
IOSPLCDEMAIL440	—	117.98	—	117.98
IOSPLCDEMIN440	—	4038.94	—	4038.94
IOSPLDECW440	—	740.93	—	740.93

Tru64 UNIX Worldwide Language Support Version 4.0F

Subset	root	/usr	/var	Total
IOSPLOLDDECW440	—	3548.81	—	3548.81
IOSPLOLDX11440	—	870.46	—	870.46
IOSPLUCSBASE440	—	116.43	—	116.43
IOSPLX11440	—	6654.71	—	6654.71
IOSPLXDEV440	—	94.69	—	94.69
IOSRUDECW440	—	733.34	—	733.34
IOSRUOLDDECW440	—	3717.49	—	3717.49
IOSRUOLDX11440	—	876.66	—	876.66
IOSRUUCSBASE440	—	79.20	—	79.20
IOSRUX11440	—	6501.25	—	6501.25
IOSRUXDEV440	—	94.85	—	94.85
IOSSKCDEAPPS440	—	207.22	—	207.22
IOSSKCDEDEV440	—	261.99	—	261.99
IOSSKCDEDT440	—	1045.82	—	1045.82
IOSSKCDEMAIL440	—	80.67	—	80.67
IOSSKCDEMIN440	—	4022.92	—	4022.92
IOSSKDECW440	—	710.03	—	710.03
IOSSKOLDDECW440	—	3523.65	—	3523.65
IOSSKOLDX11440	—	846.86	—	846.86
IOSSKUCSBASE440	—	116.43	—	116.43
IOSSKX11440	—	6045.19	—	6045.19
IOSSKXDEV440	—	94.58	—	94.58
IOSSLX11440	—	209.76	—	209.76
IOSSVCDEAPPS440	—	215.22	—	215.22
IOSSVCDEDEV440	—	154.46	—	154.46
IOSSVCDEDT440	—	979.67	—	979.67
IOSSVCDEHLP440	—	12757.42	—	12757.42
IOSSVCDEMAIL440	—	73.28	—	73.28
IOSSVCDEMIN440	—	471.67	—	471.67
IOSSVDECW440	—	709.32	—	709.32
IOSSVOLDDECW440	—	3361.60	—	3361.60
IOSSVOLDX11440	—	835.55	—	835.55
IOSSVX11440	—	6032.10	—	6032.10
IOSSVXDEV440	—	94.54	—	94.54
IOSTHBASE440	—	855.67	—	855.67
IOSTHBIN440	442.48	12.20	6.10	460.78
IOSTHCDEAPPS440	—	138.79	—	138.79
IOSTHCDEDEV440	—	255.62	—	255.62
IOSTHCDEDT440	—	1125.16	—	1125.16
IOSTHCDEMAIL440	—	67.13	—	67.13
IOSTHCDEMIN440	—	489.06	—	489.06
IOSTHDECW440	—	172.56	—	172.56
IOSTHFONTM440	—	134.48	—	134.48
IOSTHOLDDECW440	—	1821.06	—	1821.06
IOSTHOLDX11440	—	434.15	—	434.15
IOSTHOLFONT440	—	7385.15	123.70	7508.85

Tru64 UNIX Worldwide Language Support Version 4.0F

Subset	root	/usr	/var	Total
IOSTHPGMR440	—	124.15	—	124.15
IOSTHPRINT440	—	172.89	—	172.89
IOSTHX11440	—	2724.44	—	2724.44
IOSTHXDEV440	—	107.32	—	107.32
IOSTRFONT100M440	—	1115.60	—	1115.60
IOSTRFONT100P440	—	4113.84	—	4113.84
IOSTRFONT75M440	—	962.49	—	962.49
IOSTRFONT75P440	—	3468.61	—	3468.61
IOSTROLDDECW440	—	123.46	—	123.46
IOSTROLFONT440	—	5458.68	102.34	5561.02
IOSTRUCSBASE440	—	76.16	—	76.16
IOSTRX11440	—	502.37	—	502.37
IOSULUCSBASE440	—	73.08	—	73.08
IOSWWBASE440	18.42	1382.14	167.04	1567.61
IOSWWBIN440	911.24	129.38	12.20	1052.81
IOSWWBINCOM440	11.46	401.45	26.65	439.56
IOSWWCDEDT440	—	304.19	—	304.19
IOSWWFGC440	—	1492.74	—	1492.74
IOSWWFONTM440	—	702.40	—	702.40
IOSWWFONTP440	—	380.67	—	380.67
IOSWWLAT2FONT100M440	—	1195.91	—	1195.91
IOSWWLAT2FONT100P440	—	4363.31	—	4363.31
IOSWWLAT2FONT75M440	—	1022.86	—	1022.86
IOSWWLAT2FONT75P440	—	3677.32	—	3677.32
IOSWWLAT2OLFONT440	—	5576.47	102.34	5678.81
IOSWWLAT4FONT100M440	—	1183.15	—	1183.15
IOSWWLAT4FONT100P440	—	4386.89	—	4386.89
IOSWWLAT4FONT75M440	—	1014.75	—	1014.75
IOSWWLAT4FONT75P440	—	3699.16	—	3699.16
IOSWWLAT9FONT100M440	—	5291.42	—	5291.42
IOSWWLAT9FONT75M440	—	4470.69	—	4470.69
IOSWWLATCFONT100M440	—	1178.66	—	1178.66
IOSWWLATCFONT100P440	—	2424.95	—	2424.95
IOSWWLATCFONT75M440	—	1004.35	—	1004.35
IOSWWLATCFONT75P440	—	2030.64	—	2030.64
IOSWWLATCOLFONT440	—	3663.37	62.14	3725.52
IOSWWMOTIF11440	—	9837.15	—	9837.15
IOSWWMULE440	—	98245.13	—	98245.13
IOSWWMULESRC440	—	26977.74	—	26977.74
IOSWWPGMR440	—	320.78	—	320.78
IOSWWPHRASE440	389.12	853.19	6.10	1248.41
IOSWWPRINT440	38.18	1401.07	50.32	1489.57
IOSWWSVEDEV440	—	438.09	—	438.09
IOSWWSYSMAN440	—	388.63	4.78	393.41
IOSWWUCSBASE440	—	57217.66	—	57217.66
IOSWWUDCOS440	402.66	2277.19	6.10	2685.95

Tru64 UNIX Worldwide Language Support Version 4.0F

Subset	root	/usr	/var	Total
IOSWWUDCWOS440	—	122.60	—	122.60
IOSWWX11440	—	4741.92	—	4741.92
IOSWWXDEV440	—	2361.28	—	2361.28
IOSWWXFR440	13.81	1170.21	3.49	1187.52
IOSZHBASE440	—	440.94	—	440.94
IOSZHBIG5440	168.16	1712.37	3.05	1883.58
IOSZHCNBASE440	—	282.07	—	282.07
IOSZHCNCDEAPPS440	—	123.36	—	123.36
IOSZHCNCDEDEV440	—	153.52	—	153.52
IOSZHCNCDEDT440	—	1088.91	—	1088.91
IOSZHCNCDEHLP440	—	18580.19	—	18580.19
IOSZHCNCDEMAIL440	—	72.92	—	72.92
IOSZHCNCDEMIN440	—	481.51	—	481.51
IOSZHCNLOC440	—	759.15	—	759.15
IOSZHCNUCSBASE440	—	1860.69	—	1860.69
IOSZHCONV440	82.06	275.66	3.05	360.77
IOSZHEUCTW440	—	814.20	—	814.20
IOSZHHANYU440	—	683.05	—	683.05
IOSZHHANZI440	—	766.41	—	766.41
IOSZHHKBASE440	—	2014.28	—	2014.28
IOSZHHKUCSBASE440	—	1565.14	—	1565.14
IOSZHPGMR440	—	2452.41	—	2452.41
IOSZHSDECW440	—	173.38	—	173.38
IOSZHSFONTM440	—	3379.82	—	3379.82
IOSZHSFONTP440	—	27595.29	—	27595.29
IOSZHSOLDDECW440	—	1871.93	—	1871.93
IOSZHSOLDX11440	—	425.72	—	425.72
IOSZHSOLFONT440	—	15014.51	6.19	15020.70
IOSZHSX11440	—	3081.71	—	3081.71
IOSZHSXDEV440	—	198.22	—	198.22
IOSZHTDECW440	—	520.60	—	520.60
IOSZHTELEX440	151.03	2170.68	3.05	2324.76
IOSZHTFONTM440	—	9060.39	—	9060.39
IOSZHTFONTP440	—	25532.25	—	25532.25
IOSZHTOLDDECW440	—	5524.23	—	5524.23
IOSZHTOLDX11440	—	1284.83	—	1284.83
IOSZHTOLFONT440	—	28075.69	6.20	28081.89
IOSZHTWBASE440	—	1889.00	—	1889.00
IOSZHTWCDEAPPS440	—	367.62	—	367.62
IOSZHTWCDEDEV440	—	455.72	—	455.72
IOSZHTWCDEDT440	—	3289.17	—	3289.17
IOSZHTWCDEHLP440	—	11978.26	—	11978.26
IOSZHTWCDEMAIL440	—	196.19	—	196.19
IOSZHTWCDEMIN440	—	1400.21	—	1400.21
IOSZHTWLOC440	—	5487.71	—	5487.71
IOSZHTWUCSBASE440	—	4646.07	—	4646.07

Tru64 UNIX Worldwide Language Support Version 4.0F

Subset	root	/usr	/var	Total
IOSZHTX11440	—	7630.42	—	7630.42
IOSZHTXDEV440	—	942.60	—	942.60
IOSZHUCSBASE440	—	22580.40	—	22580.40
IOSZHX11440	—	6499.00	—	6499.00
Total	3258.83	1088946.54	974.67	1093180.05

Grand Totals	root	/usr	/var	Total
	3344.65	1614536.39	72964.81	1690845.87

Table B-3: Tru64 UNIX Version 4.0F Associated Products, Volume 2

Advanced File System Advanced Utilities

Subset	root	/usr	/var	Total
AFAADVANCED440	—	562.22	—	562.22
AFAADVDAEMON440	10.87	1709.69	57.00	1777.56
AFAADVGUI440	47.43	14164.12	19.00	14230.55
AFAADVMAN440	—	68.35	—	68.35
AFADEFRAGMENT440	—	149.77	—	149.77
AFAJPADVGUI440	—	1696.68	—	1696.68
AFAJPADVMAN440	—	93.04	—	93.04
Total	58.30	18443.88	76.00	18578.18

Advanced Printing Software

Subset	root	/usr	/var	Total
APXADMIN100	—	4431.92	—	4431.92
APXBASE100	—	7154.91	76.00	7230.91
APXGUI100	—	14387.37	—	14387.37
APXGW100	—	1078.09	—	1078.09
APXSVR100	—	2435.88	—	2435.88
Total	—	29488.16	76.00	29564.16

Advanced Server Version 4.0A for UNIX

Subset	root	/usr	/var	Total
ASUADM401	—	27601.94	—	27601.94
ASUADMJP401	—	26249.39	—	26249.39
ASUBASE401	10.80	36803.35	—	36814.16
ASUMANJP401	—	484.76	—	484.76
ASUMANPAGE401	—	461.20	—	461.20
ASUTRAN401	107.82	1782.24	1478.47	3368.54
Total	118.63	93382.90	1478.47	94980.00

COMPAQ Analyze Version 1.0

Subset	root	/usr	/var	Total
CABASE100	—	7243.55	—	7243.55
DESTA100	18.49	30480.42	—	30498.91
Total	18.49	37723.97	—	37742.46

NetWorker

Subset	root	/usr	/var	Total
BRXCLNT520	—	49804.47	—	49804.47
BRXMAN520	—	2059.74	—	2059.74
BRXNODE520	—	16840.00	—	16840.00
BRXSERV520	—	29705.32	—	29705.32
Total	—	98409.53	—	98409.53

DEC Open3D Version 4.93

Subset	root	/usr	/var	Total
O3DDWSBASE493	—	7350.06	—	7350.06
O3DDWSCMT493	—	323.61	—	323.61
O3DDWSCOMMON493	—	11091.24	6224.36	17315.61
O3DDWSCONFIG493	—	0.00	—	—
O3DDWSGLBASE493	—	4791.24	—	4791.24
O3DDWSGLEXAM493	—	12469.16	—	12469.16
O3DDWSGLMAN493	—	3187.46	—	3187.46
O3DDWSMITPEX493	—	1701.22	—	1701.22
O3DDWSMITPEXEXAM493	—	0.00	—	—
O3DDWSMITPEXMAN493	—	2943.81	—	2943.81
O3DDWSPCM493	—	821.79	—	821.79
O3DDWSPEX493	—	0.00	—	—
O3DDWSPHRJP493	—	2660.46	—	2660.46
O3DDWSPHRMAN493	—	36.28	—	36.28
O3DDWSPHRTO493	—	21450.31	—	21450.31

DEC Open3D Version 4.93

Subset	root	/usr	/var	Total
O3DDWSSTEREO493	—	180.08	—	180.08
O3DDWSSTEREOMAN493	—	12.92	—	12.92
O3DDWSZE3493	—	25987.26	—	25987.26
O3DDWSZLXE493	—	10833.82	—	10833.82
O3DDWSZLXL493	—	16400.99	—	16400.99
Total	—	122241.71	6224.36	128466.07

Performance Manager for Tru64 UNIX

Subset	root	/usr	/var	Total
PMGRAPP440	—	458.74	—	458.74
PMGRBASE440	—	3973.18	—	3973.18
PMGRCLUSTERS440	—	948.57	490.47	1439.04
PMGRGUI440	—	26016.07	—	26016.07
PMGRMAN440	—	73.94	—	73.94
PMGRUTIL440	—	19.09	1803.09	1822.18
Total	—	31489.59	2293.55	33783.14

SCSI CAM Layered Components Version 3.1D

Subset	root	/usr	/var	Total
CLCMC314	14.63	3837.39	—	3852.02
CLCOP314	16.94	1935.17	—	1952.11
Total	31.57	5772.56	—	5804.13

System V Environment

Subset	root	/usr	/var	Total
SVEADM425	434.71	64831.23	45.59	65311.53
SVEBCP425	44.16	22773.12	30.02	22847.29
SVEDEV425	11.82	33915.49	34.65	33961.95
SVEENV425	23.68	19.16	12.00	54.84
SVEMAN425	—	7476.22	—	7476.22
SVEPRINT425	14.02	54327.97	315.67	54657.66
Total	528.38	183343.19	437.93	184309.50

TruCluster ™ Software Version 1.6

Subset	root	/usr	/var	Total
TCRASE160	91.93	21904.34	3878.57	25874.84
TCRASEDT160	—	798.90	256.14	1055.04
TCRCMS160	—	10836.30	156.88	10993.18
TCRCOMMON160	58.04	1638.51	19.00	1715.54
TCRCONF160	—	50.26	—	50.26
TCRDSVC160	—	10648.21	2313.41	12961.62
TCRMAN160	—	1632.26	—	1632.26
TCRMCA160	—	4411.52	—	4411.52
Total	149.96	51920.29	6624.01	58694.26

UniCensus

Subset	root	/usr	/var	Total
UNICEN301	—	0.00	1318.32	1318.32
Total	—	0.00	1318.32	1318.32

Visual Threads Version 1.0 for Tru64 UNIX

Subset	root	/usr	/var	Total
DVTBASE100	—	27616.55	—	27616.55
Total	—	27616.55	—	27616.55

Grand Totals	root	/usr	/var	Total
	905.33	699832.33	18528.64	719266.30

B.2 Disk Space Required for Update Installations

For information about the disk space required to update the Tru64 UNIX operating system from Version 4.0D or Version 4.0E to Version 4.0F, see the Tru64 UNIX *Software Product Description*.

B.3 Disk Space Required for RIS Areas

The Remote Installation Services (RIS) area for Tru64 UNIX Version 4.0F requires approximately 2.4 GB of disk space. The space requirements are broken down as follows:

Table B-4: Disk Space Required for RIS Areas

Product Area	512-Byte Blocks
Tru64 UNIX Version 4.0F Operating System	
Alternative_Compiler	15940
DEC_Ada_RTL	556
DEC_C++_RTL	336
DEC_Cobol_RTL	2304
DEC_EVENT	16338
DEC_Fortran_RTL	2420
DEC_Pascal_RTL	818
DEC_Sort_RTL	376
Data_Direct	65560
GNUSRC	23878
Multimedia_Service	35020
Netscape_FastTrack_Server	79172
PanoramiX_ADK	9364
Porting_Assistant	20144
Program_Analyzers	24938
Worldwide_Language_Support	617330
AdvFS_Advanced_Utilities	9324
Advanced_Printing	16684
Advanced_Server	60642
NetWorker	98894
Open3d	69578
Performance_Manager	13750
SCSI_CAM_Layered_Components	3318
System_V_Environment	77246
TruCluster	26686
Visual_Threads	27614
Total	2417774

B.4 Disk Space Required for Documentation

The files for the Tru64 UNIX documentation set are contained on the Documentation CD–ROM. The files for the worldwide support documentation are contained on the Associated Products CD–ROM in the `/mnt-pnt/Worldwide_Language_Support/doc` directory. These files require the following amounts of disk space if moved from the CD–ROM onto a local disk:

- Tru64 UNIX operating system documentation – 107 MB
- Worldwide support documentation – 6.5 MB

Index

L

lastlog, 8–4

LAT

    and NetRAIN, 4–31

    tty, 4–34

latsetup, 4–34

libc, 8–16

libsecurity.a, 8–14

libsecurity.so, 8–14

llogin, 4–34

Local Area Transport

    See LAT

logical partitions, 1–3

    documentation updates, 7–2

Logical Storage Manager

    See LSM

long double type, 8–5

lost startup messages, 4–24

lpstat command, 4–26

LSM, 4–38, 8–9

    Block Change Logging (BCL), 8–7

    boot failure, 4–40

    dissociating LSM plex, 4–40

    kernel rebuild required, 4–39

    limits, A–4

    rootvol, 4–40

    using with SWXCR-E(A/B), 4–39

    using with SWXCR-P(A/B), 4–39

    volrootmir, 4–40

    writing to block 0, 4–39

Lucida-Typewriter-medium-R-normal font

    X11, 6–7

LVM-to-LSM migration, 8–9

M

MAILLOCKING, 6–4

mailsetup, 8–13

makepsres, 8–2

malloc, 5–1

memory limits, A–1

    physical memory, A–1

    virtual memory, A–1

memory requirements

    physical memory, A–1

Microsoft Sound Board driver, 8–10

Microsoft Windows NT clients

    and NetRAIN, 4–31

mkfdmn, 8–7

Modular Computing Components

    configuring eight-head support, 3–11

Motif

    Version 1.1.3, 8–10

mount, 4–35

mule, 8–7

multiscreen environment, 6–1

N

nemacs, 8–7

NetRAIN, 4–29

    and LAT, 4–31

    and Microsoft Windows NT clients, 4–31

    configuring, 4–30

    failure of network services, 4–30

    supported network adapters, 4–29

Netscape

    core dumps caused running from CDE, 4–5

    core dumps if mailto: link is selected, 4–6

    core dumps intermittently, 4–6

    core dumps when multiple mail messages are
        deleted, 4–6

How to Order Tru64 UNIX Documentation

You can order documentation for the Tru64 UNIX operating system and related products on the following Web site:

http://www.businesslink.digital.com/

If you need help deciding which documentation best meets your needs, see the Tru64 UNIX *Documentation Overview* or call **800-344-4825** in the United States and Canada. In Puerto Rico, call **787-781-0505**. In other countries, contact your local Compaq subsidiary.

If you have access to Compaq's intranet, you can place an order at the following Web site:

http://asmorder.nqo.dec.com/

The following table provides the order numbers for the Tru64 UNIX operating system documentation kits. For additional information about ordering this and related documentation, see the *Documentation Overview* or contact Compaq.

Name	Order Number
Tru64 UNIX Documentation CD-ROM	QA-MT4AA-G8
Tru64 UNIX Documentation Kit	QA-MT4AA-GZ
End user Documentation Kit	QA-MT4AB-GZ
Startup Documentation Kit	QA-MT4AC-GZ
General user Documentation Kit	QA-MT4AD-GZ
System and network Management Documentation Kit	QA-MT4AE-GZ
Developer's Documentation Kit	QA-MT5AA-GZ
General Programming Documentation Kit	QA-MT5AB-GZ
Windows Programming Documentation Kit	QA-MT5AC-GZ
Reference Pages Documentation Kit	QA-MT4AG-GZ
Device Driver Kit	QA-MT4AV-G8

Reader's Comments

DIGITAL welcomes your comments and suggestions on this manual. Your input will help us to write documentation that meets your needs. Please send your suggestions using one of the following methods:

• This postage-paid form

• Internet electronic mail: `readers_comment@zk3.dec.com`

• Fax: (603) 881-0120, Attn: UBPG Publications, ZKO3-3/Y32

If you are not using this form, please be sure you include the name of the document, the page number, and the product name and version.

Please rate this manual:	Excellent	Good	Fair	Poor
Accuracy (software works as manual says)	☐	☐	☐	☐
Completeness (enough information)	☐	☐	☐	☐
Clarity (easy to understand)	☐	☐	☐	☐
Organization (structure of subject matter)	☐	☐	☐	☐
Figures (useful)	☐	☐	☐	☐
Examples (useful)	☐	☐	☐	☐
Index (ability to find topic)	☐	☐	☐	☐
Usability (ability to access information quickly)	☐	☐	☐	☐

Please list errors you have found in this manual:

Page   Description

_____   _____
_____   _____
_____   _____
_____   _____
_____   _____

Additional comments or suggestions to improve this manual:

_____
_____
_____
_____
_____

What version of the software described by this manual are you using? _____

Name/Title _____ Dept. _____
Company _____ Date _____
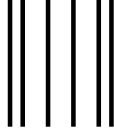Mailing Address _____
_____ Email _____ Phone _____

COMPAQ

BUSINESS REPLY MAIL
FIRST CLASS MAIL PERMIT NO. 33 MAYNARD MA
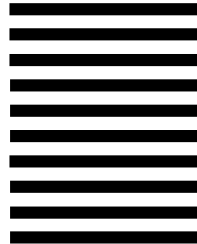
POSTAGE WILL BE PAID BY ADDRESSEE

COMPAQ COMPUTER CORPORATION
UBPG PUBLICATIONS MANAGER
ZKO3 3/Y32
110 SPIT BROOK RD
NASHUA NH 03062 9987