Tru64 UNIX

Release Notes for Version 4.0G

Part Number: AA-QTLMG-TE

May 2000

Product Version: Tru64 UNIX Version 4.0G

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

Compaq Computer Corporation
Houston, Texas

Contents

4 Base System Software Notes

5 Development Environment Notes

8 Features and Interfaces Scheduled for Retirement

About This Manual

This manual contains release notes for the Compaq Tru64™ UNIX® Version 4.0G operating system.

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 the 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 software subsets for the Tru64 UNIX software kit.
Appendix B	Contains information about the system limits.

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 Documentation CD–ROM
- *Read This First* letter

  This letter provides general information pertaining to the Tru64 UNIX Version 4.0G 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 the Tru64 UNIX operating system.
- *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.

Icons on Tru64 UNIX Printed Books

The printed version of the Tru64 UNIX documentation uses letter icons on the spines of the books to help specific audiences quickly find the books that meet their needs. (You can order the printed documentation from Compaq.) The following list describes this convention:

G	Books for general users
S	Books for system and network administrators
P	Books for programmers
D	Books for device driver writers
R	Books for reference page users

Some books in the documentation 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.

You can also view the Tru64 UNIX Version 4.0G *Technical Update* for any additional information not included in these notes. You can access the *Technical Update* from the following URL:

http://www-unix.zk3.dec.com:8083/faqs/publications/pub_page/update_list.html

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: UBPG Publications, ZKO3-3/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
  UBPG Publications Manager
  ZKO3-3/Y32
  110 Spit Brook Road
  Nashua, NH 03062-2698

  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 manual:

#	A number sign represents the superuser prompt.
`% cat`	Boldface type in interactive examples indicates typed user input.
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).
. . .	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.

1

New and Changed Features

The Compaq Tru64 UNIX Version 4.0G software replaces earlier versions of the operating system software. This chapter describes the follwoing new and changed features provided by this release of the operating system:

- Support for new hardware
- Netscape Communicator Version 4.7
- Support for the collect data collection tool
- Fibre Channel connectivity management tool
- Ability to monitor status of the 48-volt power supplies
- Support for an increased number of X clients

1.1 Support for New Hardware

The following new hardware is supported in this release of the operating system:

- Compaq AlphaServer GS160 systems
- AlphaServer DS320E systems
- 3X-DE600-AA: Single-port 10/100 UTP Ethernet Adapter
- 3X-DE602-AA: Dual-port 10/100 UTP Ethernet Adapter
- 3X-DE602-FA: Single-port, multi-mode fiber (MMF) "FX" add-on daughter card for 3X-DE602-AA
- DS-KZPCC-CC: 3-Channel, 64-bit PCI to Ultra2 LVD Backplane RAID controller with 64 MB cache

1.2 Netscape Communicator Version 4.7

This release of Tru64 UNIX contains Version 4.7 of Netscape Communicator. For information on using Netscape Communicator to display the online documentation, see Chapter 8 in the Tru64 UNIX *Installation Guide*.

To obtain the latest fixes to the Netscape Communicator problems described in Chapter 4 of these release notes, it is recommended that you download and install the latest version available of Netscape

Communicator for Tru64 UNIX from the Netscape Netcenter Download World Wide Web site. Enter the following URL in your web browser to obtain the latest version of Netscape Communicator:
http://home.netscape.com/download/index.html#clients

1.3 Collect Data Collection Tool Now Supported

Collect Version 2.0 is a Y2K-compliant tool that collects operating system and process data under all supported versions of Compaq Tru64 UNIX. Collect is designed for high reliability and low system resource overhead.

The Collect tool gathers and displays information for subsystems. You can set Collect to gather data for a single subsystem or any combination of the following:

Processes
LSM
Memory
Network
File systems
Disks
CPUs
Queues
Message
Tapes

Collect has two modes of operation: collection and playback. In collection mode, it gathers operating system and process data and writes it to standard output, to a binary file with a compressed format, or to both.

In playback mode, Collect reads from a previously written binary file and writes to standard output. The format of the data when written to standard output is identical, whether during realtime collection or playback mode.

See the `collect`(1) reference page for more information.

1.4 Fibre Channel Connectivity Management Utility

The `emxmgr` is a utility that can be used to maintain the Fibre Channel Worldwide name to Target ID mappings for emx devices. When you use the `emxmgr` utility to maintain mappings, you do not need to rebuild the kernel or reboot the system to complete the maintenance functions.

The `emxmgr` utility is used to update and maintain emx mappings when:

- A Fibre Channel device is removed from the system
- A Fibre Channel device is replaced (swapped)

- A host adapter is added or removed
- A symmetrical Fiber Channel configuration is created for a cluster

See the emxmgr(8) reference page for more information on using this utility.

1.5 New Networking Features

Networking enhancements include the following:

- Support for RFCs 1901-1908 (SNMPv2C) in the Extensible SNMP Agent, subagent developer's tools, and SNMP-related commands.
- Support for RFC 2089 (Mapping SNMPv2 to SNMPv1) in the bilingual extensible SNMP agent.
- Support for RFC 2257 (Agent Extensibility) in the Extensible SNMP Agent and subagent developer's tools.
- Improved scalability of the ATM LANE subsystem.

1.6 48-Volt Power Supply Monitoring on 8200, 8400, GS60, and GS40 Systems

AlphaServer 8200, 8400, GS60, and GS40 systems now have the ability to monitor the status of 48-volt power supplies and to report and log bad power supply status.

1.7 Maximum Number of X Clients

In this release of the operating system, the maximum number of X clients supported by the X Server has been increased from 128 to 256.

2

Installation Notes

This chapter discuss the following topics:

- RIS installation
- Update installation
- Join database migration
-
- Installation messages

Do not attempt to install Tru64 UNIX Version 4.0G 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 RIS Installation Notes

The follwoing notes describe restrictions and known problems with Remote Installation Services (RIS) installations.

2.1.1 Correcting RIS Gateway File Entries

The `/var/adm/ris/gateways` file contains information about the address of the gateway between the client system and the RIS server. When you register a new client, the `ris` utility queries this file to determine if a gateway is already specified for the client's network subnet. If not, you are prompted for the necessary information.

If the gateway address is entered incorrectly when it is first specified, you must edit the `/var/adm/ris/gateways` file manually to correct the entry.

Entries in this file have the following format:

```
subnet_addr:gateway_addr
```

The following example shows a typical `/var/adm/ris/gateways` file:

```
16.168.64:16.168.64.1
16.69.240:16.69.224.199
16.140.144:16.140.144.2
16.69.144:16.69.144.199
```

After you correct the entries in this file, you must use the `ris` utility to remove all clients using the bad gateway address and register them again.

2.1.2 Restriction for RIS Installation on IDE Disks

When using `ris` to install to an IDE disk, the system may not set up the `bootdef_dev` environment variable correctly. If the `bootdef_dev` variable is set incorrectly, the system will print the following when it attempts to reboot:

```
The installation software has successfully installed your system.

There are logfiles that contain a record of your installation.
These are:
/var/adm/smlogs/install.cdf     - configuration description file
/var/adm/smlogs/install.log     - general log file
/var/adm/smlogs/install.FS.log  - file system creation logs
/var/adm/smlogs/setld.log - log for the setld(8) utility
/var/adm/smlogs/fverify.log - verification log file
The above message is also recorded in /etc/motd for
your future reference.

Setting console environment variables
Current Firmware: 5.5-6
I/O error (errno 5) for block(0x40, 0x40) on device 255,1
syncing disks... done
rebooting.... (transferring to monitor)

CP - SAVE_TERM routine to be called
CP - SAVE_TERM exited with hlt_req = 1, r0 = 00000000.00000000
halted CPU 0

halt code = 5
HALT instruction executed
PC = fffffff002237f0

CPU 0 booting

Reseting I/O buses...
(boot dk*0.0.1.13.0 -flags A)
failed to open dk*0.0.1.13.0
Retrying, type ^C to abort...
```

At this point the system will retry booting from this device forever. The system has been successfully installed to disk, but it does not know how to access the disk that was just installed. To recover from this situation, you must set the `bootdef_dev` environment variable manually using the following procedure:

```
^c
>>>show dev
>>>set bootdef_dev  device_name
>>>boot
```

The system will now boot from the newly installed disk and complete the installation.

2.2 Update Installation

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

Version 4.0G supports update installations from Version 4.0E and Version 4.0F. 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 8.25 for more information.

2.2.1 Recovering from Boot Problems After Update Installation

During an installation update from Version 4.0E or Version 4.0F to Version 4.0G, 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 4.0G generic vmunix kernel.

Depending on the hardware configuration and the content of the existing system configuration files, the new Version 4.0G 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 4.0G 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 4.0G 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.

2.3 Join Database Migration

If your system provides Dynamic Host Configuration Protocol (DHCP), Remote Installation Services (RIS), or Dataless Management Services (DMS) services 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.

2.4 Installation Messages

The following sections describe error messages that may be reported during system installation.

2.4.1 I/O Error Message

After the installation process has completed installing all of the requested subsets, you may encounter the following benign error message:

```
I/O error (errno 5) for block ( xxx , xxx ) on device xxx , x
```

You can ignore this message. The installation will complete successfully.

2.4.2 Host Not Qualified Warning Message

After the installation process has completed installing all of the requested subsets and the kernel has been built, you may encounter the following warning message:

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

You can ignore this message. The system will continue to boot normally.

3

Processor-Specific Notes

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
- AlphaServer GS160 systems
- Alpha VME 4/224, 4/288, and 5/nnn 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 Floppy Drives May Fail to Configure

In some hardware configurations, floppy drives may fail to configure correctly during system boot. The following message is displayed when this failure occurs:

```
floppy driver: failed to allocate desired dma resources
fd internal driver error: FDI PROBE FAIL (A,1).
fdi0 not probed
```

After this message is displayed, the system will continue to boot normally. Becuase of this failure, you will not be able to use the floppy drive.

3.1.3 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 to disable PCI parity:

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

3.1.4 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.

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

3.1.5 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.

3.1.6 Qlogic ISP1040B CAM Errors

On systems with a Qlogic ISP1040B option, CAM errors like the following might occur when you boot the system:

```
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.

3.1.7 DJ-ML200-xx PCI NVRAM Hardware Revision Must Be E01

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

3.1.8 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
```

3.1.9 Shared SCSI Buses Report Command Timeouts

Shared SCSI buses with multiple hosts performing I/Os to the same logical unit number (LUN) may report command timeout errors. Some of the timeouts occur more than once in the same second for the same device.

These timeout errors are reported in the `binary.errlog` file when writes to the same storage device are performed from more than one node.

3.2 AlphaStation 500/400, AlphaStation 500/500, and AlphaStation 600 Systems

There is a memory restriction in the firmware for these systems that prohibits bootlinking the kernel on Version 4.0G of the operating system. An attempt to bootlink a kernel may result in a system hang. In some cases, system hangs may also be encountered when attempting to boot a customized kernel.

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

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

3.3.1 EIDE (ATAPI) CD–ROM Support

The Compaq Personal Workstation class systems contain two EIDE (ATAPI) ports. Version 4.0G of the operating system 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.

3.3.2 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
```

3.3.3 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.

3.4 AlphaServer 1000 and 1000A Systems

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

3.4.1 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.

3.4.2 Graphics Resolution

The default graphics resolution for 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 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.

3.5 AlphaServer 8200, 8400, GS60, and GS140 Systems

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

3.5.1 Minimum Firmware Requirements

Before installing the Tru64 UNIX software kit, ensure that your system has the correct firmware version. The minimum firmware version required for 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 4.0G 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).

3.6 AlphaServer GS160 Systems

The following notes apply to the AlphaServer GS160 systems.

3.6.1 GS160 Systems Do Not Support the svrSystem_mib Server MIB

The server MIB, `svrSystem_mib`, FRU parser does not support FRU Version 6.0. Information normally obtained from the FRU tables on GSxx systems is not provided by the MIB.

Because FRU Version 6.0 is not supported, errors similar to the following appear in the `/usr/var/adm/syslog.dated/date/daemon.log` log file:

```
Sep 24 11:00:44 wildthing [491]:  **ERROR svrsys_fru_parse.c line 2076:
FRU Table revision 6.0 is not supported
```

3.6.2 Hot Swap Features of GS160 Systems Are Not Supported

Version 4.0G of the operating system does not support use of the hot swap features of the GS160 systems.

3.7 Alpha VME and PCI/ISA (DMCC) Modular Single-Board Computers

For information about configuring the operating system on Alpha VME single-board computers (SBCs) and PCI/ISA modular SBCs, including the SMARTengine/Alpha 21264 SBC, see the *System Configuration Supplement: OEM Platforms*. (The PCI/ISA modular systems and components product family was formerly known as DIGITAL Modular Computing Components, or DMCC.)

Note

The April 2000 edition of the *System Configuration Supplement: OEM Platforms* that is provided with the operating system is labeled "Tru64 UNIX Version 5.0A or higher", but in fact its content applies equally to Tru64 UNIX Version 4.0G. This includes the manual's descriptions of four VMEbus backplane (`vb`) network driver parameters used to enhance data transfer throughput: `VB_Transfer_Type`, `VB_DMA_Threshold`, `VB_DMA_Dwidth`, and `VB_Maxmtu`. Although these parameters are flagged as requiring Version 5.0A or higher, they can be used in Version 4.0G as well.

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 4.0G 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 4.0G software on some Compaq 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.

4

Base System Software Notes

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 does 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 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 ToolTalk Security Enhancement

To prevent unauthorized access to your machine, a new security mechanism has been added to ToolTalk. This security mechanism, which was jointly developed by all companies shipping the Common Desktop Environment (CDE), requires that a ToolTalk message contain a valid cookie in order for the `ttsession` message server to deliver the message to its recipients. A

different cookie is generated by ttsession every time a user logs in using dtlogin.

The cookie resides in a new file called .TTauthority under your home directory. This permits you to send ToolTalk messages to the local ttsession message server. Other users who want to send a ToolTalk message to ttsession must place a copy of the cookie in their .TTauthority file. See the ttauth(1) reference page for instructions on how to share a cookie with other users.

For the special case of a root user sending messages to the local ttsession, ToolTalk looks for the cookie in the .TTauthority file of the user who owns the ttsession process. For messages being sent to a ttsession on a remote machine, ToolTalk looks for the cookie in root's .TTauthority file.

You can use the TTAUTHORITY environment variable to specify an alternate authority file.

Requiring all ToolTalk messages to contain a valid cookie might cause problems with some ToolTalk clients. Therefore, you now have the ability to start ttsession with relaxed security, full security, or no security.

Relaxed security is the default and requires a valid cookie only for the ToolTalk messages used to start an application on a remote machine. These messages contain a handler ptype or have an operation name that maps to a ptype in the ptype database. Other message types are always delivered. Relaxed security is ideal for situations in which notification messages are constantly being sent between ToolTalk clients.

To request full security, which requires that all messages contain a valid cookie, start the ttsession with the -F flag. You must use the -F flag in conjunction with the -a cookie flag (set by default).

To request no security, start the ttsession with the -a none flag. With no security, all messages are delivered without verification. This is not recommended, because it leaves your machine vulnerable to attack.

4.1.2.2 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 the following:

- 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 because 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.

4.1.2.3 ACLs and Archive Tools

The `pax`, `tar`, and `cpio` archive tools might not restore access control lists (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 backs up a file by renaming the original file and saving the new file as a copy under the original name. If the original file had an access control list (ACL), it now applies to the backup file. If the directory had a default ACL, the new file (original file name) now has 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. Following are the relevant Emacs variables:

- `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 format. This change is compatible with other UNIX implementations and conforms to the 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. 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, 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 Tru64 UNIX Documentation CD–ROM.

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 versions 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. After 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 following warning message is displayed as aresult 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 occurs because neither bind nor mail has been configured. Sendmail will continue to operate.

4.1.7.3 Warning Message When sendmail Updates the alias Database

When the sendmail program starts and its alias database is updated, it displays the following warning message:

```
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 were compromised.

Because the permission on the /var directory is 775, the sendmail program displays 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 @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 @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.

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 following script:

```
/usr/bin/X11/netscape
```

As long as you use this script 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 Deleting Multiple Mail Messages Causes Netscape Communicator to Dump Core

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

4.1.9.3 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.

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

After upgrading to Netscape Communicator 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 that might require user customization after upgrading.

The Namespace extension to the standard IMAP protocol is used to locate the user's 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.

4.1.9.5 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 message:

```
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 might 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.

4.1.10 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 the Account Manager, `dxaccounts`.

4.2.1.1 General Restrictions

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

- When copying user accounts via cut and paste or drag and drop, the Allow Duplicate UIDs option in the General Preferences dialog box is honored. For example, when making a copy of a user account that has a UID of 200, if the Allow Duplicate UIDs check box is off (the default), a unique UID is automatically generated for the resulting copy. If

the Allow Duplicate UIDs check box is on, then the copy will have an identical UID. The same rules apply to copying groups.

- Leading and trailing white space is not stripped from text entry areas. This can lead to confusion. For example, if a field in the Find dialog box contains a space character before the desired search string, the search string will not match because of the spurious space character.
- Using mouse button 1 (MB1) to drag and drop user accounts, groups, or templates does a copy operation, not a move operation. This is different from the CDE/Motif default, in which MB1 performs a drag-and-drop move operation and Shift/MB1 performs a copy operation. For example, if you use MB1 to drag a user account from the Local Users view and drop it in the NIS Users view, you create a copy of that user account in NIS. To avoid this problem, delete the original icon after the copy has been completed.
- If you change a user's UID with the Account Manager, the ownership of the user's files and subdirectories does not change and, under certain circumstances, the home directory ownership may not change, either. For example, if you change the UID of user `johndoe` from 200 to 201, the files and subdirectories under his home directory still belong to UID 200. Furthermore, if `johndoe` does not own his home directory, the ownership of that directory does not change. To avoid this problem, use the `chown` command to change the directory and files, if applicable.
- You cannot drag and drop items across different instances of the Account Manager. For example, if the Account Manager A on system 1 and the Account Manager B on system 2 are displayed on the same workstation, then you cannot drag and drop between Account Manager A and B. To avoid this problem, use the copy/paste feature to copy users, groups, or templates from Account Manager A to B. After paste operations, the Paste Errors dialog box might be displayed. You can ignore the error message and click OK to dismiss the dialog box.
- Although the Account Manager correctly allows two or more system administrators to work on the same password files simultaneously, only one system administrator can use the Account Manager at a time. If multiple instances of the Account Manger are run concurrently, the proper file locking occurs 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 modification of each group. Therefore, if more than one system administrator is working on the same file, the last one to change a group's view window overwrites any prior changes from a different system administrator. For this reason, running multiple, concurrent Account Manager instances is not recommended.

4.2.1.2 Account Manager and Enhanced Security

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

- The Lock/Unlock Toolbar and Menu Options are inactive for the Template views. To avoid this problem, change the template lock setting on the Add/Modify Template dialog box 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 Minimum Length and Maximum Length fields of the Password Controls imply otherwise. To avoid this problem, set the passwords by using the `/usr/tcb/bin/dxchpwd` or the `/usr/bin/passwd` command when using the C1Crypt Encryption Type.
- The Account Manager does not enforce the minimum and maximum password length limitations when setting passwords. To avoid this problem, set passwords by using the `/usr/tcb/bin/dxchpwd` or the `/usr/bin/passwd` command if the minimum and maximum password length limitation is necessary.
- On an enhanced security system, you typically retire user accounts instead of deleting them. However, at times you might want to delete a user account. Account Manager supports retiring user accounts but not deleting them. To delete a user account, do the following:

  1. Manually edit the `/etc/passwd` and `/etc/group` files to remove references to the user account.
  2. Use the following command to remove the user account from the protected password database:

     ```
     # /usr/tcb/bin/edauth -r user name
     ```

- When you rename a user account by changing the Username field of the Add/Modify User dialog box in Modify mode, the protected password database entry for the old name does not change. To avoid this problem, 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 Add/Modify Template dialog box in Modify mode. The Account Manager creates a new template without removing the old template, but renames the old template's icon from the Icon Box. To avoid this problem, restart the 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 reference. All accounts and templates implicitly reference a default template that is not served by the Network Information Service (NIS). This creates an inconsistency for the Account Manager when displaying NIS user accounts and templates on a NIS master. The user and template values displayed might be the default template values of the NIS master. When a NIS user logs in to an 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.
- When you use drag and drop to copy a user account on a different view, the user's template references are copied by value. This means that the template itself is no longer referenced by the new account. Instead, the template's values are 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 are 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. To avoid this problem, 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 account is dropped within the same view.
- After deleting a template, the NIS maps are not remade. Therefore, you will have to manually remake the NIS maps or perform an Account Manager function (for example, Account Modification) that will remake the maps. To manually remake the maps, do the following:

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

4.2.2 BIND Configuration

The following problem applies to configuring BIND servers with the BIND configuration graphical user interface:

- You cannot modify the Server Addr field of the secondary domain in the Modify Zone window using the BIND configuration interface. To avoid this problem, modify the address 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 are able to log in to a machine that is being shut down up until the actual time of the shutdown.

Note that this behavior differs from that of the `shutdown` command which 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 Known Problems with the Disk Configuration GUI

The following list describes known problems with the Disk Configuration graphical user interface (GUI):

- In the Disk Configuration list box, the Physical Location for an FC disk is not listed. The list box contains a question mark (?).
- The wrong help screen is accessed from the Disk Attributes screen. When you select Help in the Disk Configuration: Disk Attributes screen, the system displays the Configure Partitions dialog box.
- The megabyte conversion of disks is inaccurate.
- The start and end points of a size 0 (zero) partition are incorrect. The end is displayed as one less than the start.

4.3.2 Cleanup of the Binary Error Log File

The binary error log file as configured by `/etc/binlog.conf` (defaulting to`/usr/adm/binary.errlog`) is not automatically cleaned (older versions removed). As long as the `binlogd` daemon is running, the log file is kept open.

Normally, a well-managed system does not produce excessively large error log files, and you will want to maintain the history and continuity of error logs. To save and clean up `binary.errlog` files, do the following:

1. Save any existing log file that you want to keep.

Note

The saved log file will overwrite any previous files in /usr/var/adm/binlog.saved. If you want to keep previous log files, you must either rename them first, or move them to another location. It is a good idea to compress saved log files to save disk space.

2. After saving any existing logs that you want, run the following command to regain the disk space:

   #kill -USR1 'cat /var/run/binlogd.pid'

   This command renames the current log file to /usr/var/adm/binlog.saved/binary.errlog.saved and starts a new version of the log file.

If you prefer an automated cleanup approach, become the root user (superuser) and run the crontab command to uncomment the following crontab entry:

```
#0 2 1 * * kill -USR1 'cat /var/run/binlogd.pid'
```

See the crontab(1) reference page for more information. When implemented, this crontab entry will do the following:

1. Run at 2:00 AM on the first day of every month
2. Automatically rename the current log file to the following: /usr/var/adm/binlog.saved/binary.errlog.saved
3. Start a new copy of the log file, /usr/adm/binary.errlog

Note that the saved log file overwrites the previous version, so you should add a line to the crontab entry to copy the existing saved file to preserve it. Using this crontab entry means that the current and saved version combined will preserve up to two months of log records. Regular backups and a restore are required in order to reconstruct older binary error logs. You can also modify the crontab entry to adjust the frequency of the automatic cleanup.

4.3.3 EISA Configuration Utility Revision Requirements

For Tru64 UNIX 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, update the ECU to this supported version.

4.3.4 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.

4.3.5 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. You should change applications 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
```

4.3.6 Restrictions for a Dataless Environment

The Tru64 UNIX Server Extensions Software 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 reside 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.

4.3.7 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 might 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 because 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 is not changed. However, because the dates have changed, the behavior of utilities that examine file dates (such as the `find` command or archivers) might be affected.

4.3.8 Enhanced Security

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

4.3.8.1 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 time out 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.

4.3.8.2 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.

4.3.8.3 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 the Account_defaults file does not exist, the internal default for the useradd command is to create locked accounts. Use the administrative_lock_applied extended command-line option to override the default. In the following example, the useradd command 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 the `/etc/passwd` file, 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 you use the `usermod` command 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.

4.3.9 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, 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.

4.3.10 Compaq Insight Manager Agents

Compaq 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.

4.3.10.1 Installation

Compaq Insight Manager is installed automatically when you install the operating system's mandatory subsets. They are `OSFCLINET4xx`, which includes the Compaq MIBs and SNMP subagent, and `OSFIMXE4xx`, which

includes the Compaq Insight Manager subagents. For further information on these subsets, see the *Installation Guide.*

4.3.10.2 Starting and Stopping the Compaq Insight Manager Daemons

The Compaq 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
```

4.3.10.3 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 Compaq 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 Compaq Insight Manager Agents and Compaq 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 window.

Not all the components are implemented in this release of the Compaq Insight Manager Agents; more will be made available in future releases. The following table shows data availability in this version.

Configuration		
	System Information	Yes
	System Board	Yes
	Software Version Information	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.

4.3.10.4 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
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
```

4.3.10.5 Security Access

In this release, you can access anonymously the device web page without logging in. If needed, the Compaq 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 must copy that file to the local directory `/var/im/webagent` if you need to change its configuration setting.

4.3.10.6 Known Problems

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

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

- The Compaq 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 Compaq 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 Compaq Insight Manager may display incorrect cache or memory module sizes on AlphaServer platforms that are configured with EV6 CPUs and support FRU V5.2.
- The Compaq Insight Manager does not support partitioned memory on the AlphaServer GS140 platform. The "System Board" web page displays information about all hardware memory modules, instead of only those modules configured for the partition in which the Insight Manager is running.

4.3.11 Startup Messages Lost in Large Configurations

On systems that display a large number of console messages at system initialization (typically, systems configured with a large number of devices), some messages may be missing from the `/var/adm/messages` file. You can correct this problem 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 for your system:

```
generic:
msgbuf_size = 32768
```

2. Enter the following command:

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

If a different entry is present in the database, `sysconfigdb` displays a warning message to advise you of the change in size.

The increase takes 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. Because the space used by the buffer is not returned for general use after initialization, set the size only high enough to correct the problem.

See the *System Administration* guide for information on changing the buffer size.

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

A print job to a Version 4.0G 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 spooldirpath lock
```

4.3.13 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`.

4.3.14 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.

4.3.15 Problems When Trying to Mount DVD–ROM

The preliminary support for the native DVD file system in Version 4.0G, 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 4.0G.

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. See `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 4.0G, many `syslog` messages for routine ATM events are not displayed by default. The messages that show the duration of virtual connections (VC accounting messages) must be explicitly enabled.

You can control the verbosity of ATM messages 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. Supresses most routine or transient events.
- 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.

You can set message levels for individual ATM components to tailor the amount of information displayed. See `sys_attrs_atm`(5) for more information.

4.4.3 Edit the /etc/hosts File After Running netsetup

In Version 4.0G, after you configure your interfaces using `netsetup`, view the `/etc/hosts` file and make sure that the host name and the IP address of your system have 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 host name of your system.

example, example-alias2 = aliases for your system.

See 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
```

You 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)

Compaq 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 4.0G, or did so in a previous releases of the operating system, you will need to 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 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

The `sendmail` application, includes the ability to specify an owner for a mail alias, causing 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 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 Acknowledgment Support

This release of the operating system 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 options.

4.4.13 Latent Support for Multiple Network Adapters in a Subnet

Version 4.0G 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 because 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 the disktab File

Starting with Version 4.0D, the `newfs` command does not search 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 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 swarel:/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. Because 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 Use the quotacheck to Permanently Eliminate AdvFS Quota Underflow

This release includes a fix for a sporadic AdvFS problem that occasionally resulted in quota underflow messages on the system console. The fix for this problem will prevent quota summary information from becoming inconsistent in the future. However, existing AdvFS file systems may already contain inconsistent quota summary information.

For this reason, it is recommended that after booting a system with this release of the operating system, you should run quotacheck against all

filesets to insure that their quota information is correct. You can accomplish this by mounting all filesets and running the `quotacheck` command against them explicitly. The use of `quotacheck -a` will only run against those filesets that have quotas enabled in the `/etc/fstab` file. If you encounter the error message `"not properly enabled for quotas in /etc/fstab"`, then you must add the fileset to the `/etc/fstab` file with the `userquota` and `groupquota` options included in the fstab line, then repeat the `quotacheck` command.

After you have run `quotacheck` on all filesets, you should not receive further quota underflow messages on the system console.

4.6.4.2 The advfsstat Command Returns Incorrect Information for BMT Statistics

When gathering statistical information on read and write activity to the Bitfile Metadata Table (BMT), the `advfsstat` command returns unintelligible information and terminates unexpectedly. This problem occurs when using the `-B r` and the `-B w` options of the `advfsstat` command.

4.6.4.3 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()` system call also needs to be written, then you must also use the `msync()` system call.

4.6.4.4 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.

4.6.4.5 Unmount File Systems Cleanly 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 when 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 4.0G and the system crashes, do not reboot using Version 3.2G since that version of AdvFS might not work with the log records that the Version 4.0G 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.

4.6.4.6 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. It is recommended 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 off line, 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.

4.6.4.7 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.

4.6.4.8 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
```

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

4.6.4.9 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

In previous releases of the operating system the following restrictions applied to root, primary swap, and secondary swap volumes configured under LSM:

- 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.

These restricitions no longer apply. However, the following restriction still does apply:

- 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 after it is set off line by the controller and cannot be used by LSM until it is reconfigured. See the *StorageWorks RAID Array 200 Subystem Family Installation and Configuration Guide* for more information.

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. See 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 it to become root:

```
lsm:
lsm_rootdev_is_volume=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:

```
>>> 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 previously 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 cannot 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
```

4.6.6 Mounting Associated Products CD–ROMs on Tru64 UNIX Versions Released Before Version 4.0F

You can mount the Associated Products CD–ROMs (APCDs) 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 compact disc file system (CDFS) support is not built in to the kernel that is currently running:

```
# 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)
```

5

Development Environment Notes

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
- Java programming

5.1 General Programming

The following notes apply 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 that make 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 Compaq 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 Compaq C compiler.

The default tuning for the Compaq 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 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.

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 `CMPDEVENH` subset. After 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 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.

5.3.3 Problems Using the stackaddr Thread Creation Attribute

Using the `stackaddr` thread creation attribute, which allows you to allocate your own stack for a thread, is not recommend. 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, because you 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.

Alternatively, if you want to supply your own thread stacks, consider using the `pthread_attr_setstackaddr_np()` routine. Callers specify the thread stack using a base address and size, which avoids the worst problems with the standard interface.

5.3.4 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.

5.3.5 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 5.3.4). 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 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.

5.3.6 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.

5.3.7 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 8 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 8 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, It is recommended that you force alignment of each such member variable to a quadword boundary by redefining the variable to be at least 8 bytes, or by defining sufficient padding storage after the variable to total 8 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.

5.3.8 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 0 (zero) indicating success. Debuggers such as Ladebug and TotalView provide functionality formerly provided by these routines.

5.3.9 DECthreads SIGEV_THREAD Notification Mechanism

Using the `SIGEV_THREAD` notification mechanism, a user-defined function is called to perform notification of an asynchronous event. The function runs as though it were the start routine of a thread and can make full use of the POSIX Threads Library 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 the POSIX Threads Library*. For more information about using SIGEV_THREAD, see the *IEEE POSIX 1003.1-1996* 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 cannot 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 file name 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 Java Programming

The following note applies to Java programming.

5.5.1 Name Space Conflict Between Java and SVE

A file system conflict exists between Java and the System V Environment (SVE) on systems running Version 4.0 or higher of the operating system.

The problem arises because both Java and SVE use the file system pathname 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. Following are three ways to avoid this 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, it is recommended 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.

6

Window System Software Notes

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 note applies 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 the 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`. See 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 might notice that icons normally displayed by the CDE Application Manager do not display when a color-rich application is currently running on the system. Graphics applications and online help volumes might also be affected. The icon editor, `dticon`, might 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. Following 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 `OSFOLDX11xxxx` 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.

6.3.3 Information for Using DPS Libraries

To continue using the Display PostScript (DPS) libraries, you must edit your `/var/X11/Xserver.conf` file and remove the exclamation points commenting out the two lines that begin with: `< _adobe_dps`.

6.4 CDE Clients

The following notes apply to CDE clients.

6.4.1 Corrupted $HOME/.TTauthority File Disables CDE Tooltalk Messaging

The `$HOME/.TTauthority` file contains a key that ToolTalk clients read and send to the ttsession message server along with each message. The `ttsession` message server compares this key with the key it placed in the user's `$HOME/.TTauthority` file when the user logged into CDE. If the `$HOME/.TTauthority` is corrupt, the clients will no longer be able to send a valid key to the `ttsession` message server. As a result, CDE will no longer be able to function normally or may not start at all.

Use the `/usr/dt/bin/ttauth list` command to examine the contents of your `$HOME/.TTauthority` file. A corrupted file may contain null values which may cause `ttsession` to dump core while trying to read the file. The `$HOME/.dt/errorlog` file contains the following error message when `ttsession` cannot be started:

```
dtsession: Unable to start message server - exiting.
```

To correct this problem delete your $HOME/.TTauthority file, stop and then and restart CDE as the root user:

```
/sbin/init.d/xlogin stop
/sbin/init.d/xlogin start
```

When you log into CDE again, a new $HOME/.TTauthority containing valid keys is created. Users that share home directories, and thus the $HOME/.TTauthority file, must also restart CDE, and log out and in again.

6.4.2 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. To correct this problem, set 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.

6.4.3 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, which has 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:

- In the Screen Saver panel of the Screen dialog box, under the Style Manager, select Blank Screen and deselect any active screen savers that might be running.
- Click on the Off button in the same dialog box.
- Execute xset s off from a terminal client window.

6.4.4 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 unexpected behavior from any of these utilities, close all windows associated with the File Manager, Application Manager, and Trash Manager. Then kill all processes associated with dtfile. You can obtain the pid for each process by using the following command:

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

6.4.5 The dtmail Application

The following notes apply to the dtmail application.

6.4.5.1 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.

6.4.5.2 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 receive 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.

6.4.5.3 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.

6.4.5.4 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.

6.4.6 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.

6.4.7 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.

6.4.8 Possible Failure in the XOpenDisplay Call

When a user logs 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.

You can add this 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, 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, 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 user name 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 Language Restrictions

The `dxterm` terminal emulator does not work correctly when processing characters in the Lithunian, Slovene, and Cyrillic languages. It also does not support the Latin-9 or Unicode locales. Use the `dtterm` terminal emulator to process characters in these languages and locales.

7

Documentation Notes

This chapter contains release notes that apply to Tru64 UNIX Version 4.0G documentation. It discusses the following topics:

- *Installation Guide*
- *System Administration* guide
- *Assembly Language Programmer's Guide*
- *Programming with ONC RPC* guide
- *Gate Daemon Configuration Guide*
- *System Configuration and Tuning* guide
- *Guide to Preparing Product Kits*
- Online help volumes
- Device driver documentation
- SCSI CAM Layered Component Release Notes
- Prestoserve 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 an 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 note describes 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.

7.3 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 the Tru64 UNIX operating system. The code may not be portable to other UNIX systems.

If the client machine is not running Tru64 UNIX , 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 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 4.0G, you must ship symbolic links from the `/sys/BINARY` directory to every kernel module in your kit. In Version 4.0G, 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.

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.

.

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 4.0G.

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

7.10 SCSI CAM Layered Component Release Notes

The release notes document for the SCSI CAM Layered Component Version 3.1E product, shipped on the Associated Products CD–ROM Volume 2, incorrectly states that the product is only supported on of the operatin g system versions through Version 4.0F. The product has been tested on Tru64 UNIX Version 4.0G and is supported on this release.

7.11 Prestoserve Documentation

The following note applies to the *Guide to Prestoserve*.

7.11.1 Prestoserve File System Status

The *Guide to Prestoserve* describes the use of the `-L` option with the `presto` command. In the current version of the guide, the `-L` option is described as a method for displaying unusual presto states of file systems.

In Section 3.2.4, one possible unusual state is identified as `bounceio`. However, in this release, `bounceio` is the default state, and is the normal state of a file system when Prestoserve is running.

This book section will be updated with the correct information in a future release. The `presto`(8) reference page has been updated in this release.

8

Features and Interfaces Scheduled for Retirement

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 have been retired from 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. The replacement applications are listed in Table 8–1. Not all of the dx* applications that have been retired have a replacement due to the limited use or capability of that specific tool or utility.

Table 8–1: Retired DECwindows Applications

Retired Tools/Utilities	Replacement Option(s)
dxmail	dtmail
dxprint	dtlp
dxcalendar	dtcm
dxcalc	dtcalc, xcalc
dxclock	Front Panel, xclock
dxpaint	dticon/dtstyle, bitmap
dxnotepad	dtpad
dxbook	dthelpview, Netscape
dxcardfiler	None
dxsession	xdm, dtsession
dxvdoc	ghostview
libids	None

It is recommended that you migrate to the dt* tools and utilities or other options as soon as possible.

The CDA applications in Table 8–2 have also been retired.

Table 8–2: Retired CDA Applications

caspar	cdoc	ctod
ddifanls	ddifps	ddiftext
dtifanls	dtifddif	dtoc
textddif	vdoc	libcapsar
libcda	libcda_be	libcda_fe
libddif	libddif_be	libddif_fe
libimg	libdvs	libids_nox

8.2 Adobe Display PostScript, Client Libraries, and X Server Extension

-->

Adobe Systems Incorporated has retired its Display PostScript (DPS) product, which includes the client libraries, X server extension, and various applications and examples. Therefore, the Adobe DPS product has been retired from Tru64 UNIX. No replacements will be available. Customers who have used the Adobe DPS libraries to develop their own applications will 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 on the Associated Products CD-ROM)

During an update installation from Version 4.0D or Version 4.0F to Version 5.0, the DPS commands, libraries, and documentation noted above are marked by the installation update procedure as obsolete inventory and then removed from the system later in the procedure.

Because these obsolete components have been retired from the operating system, they are no longer distributed, and they may 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 has become obsolete, you may wish to have the obsolete DPS software remain on your system, yet remove the other obsolete software at this time.

To do this, you can edit the file /var/adm/smlogs/upd_obsolete_files so that the DPS software components are not listed, before you perform the updadmin procedure.

If you remove the DPS components from your system 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 archive copy, there is no way to restore them from the product distribution media for the current release.

The freeware application gv may 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 Curses library, so compatible binaries (named `libcurses.a` and `libcurses.so`) were shipped in separate directories for Version 4.0.

These compatible binaries have been retired and removed from this release of Tru64 UNIX.

8.4 The -oldc C Compiler

The C compiler for DIGITAL UNIX has been officially replaced by Compaq C for Tru64 UNIX. The older compiler is no longer distributed in this release of Tru64 UNIX. The `-oldc` switch, which was used to request the old C compiler from the `cc` command in previous versions, is no longer supported. The `cc` command now issues an error if the `-oldc` switch is specified.

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.

It is recommended 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 Tru64 C Compiler Default Change from -std0 to -std

The default language mode for the Tru64 UNIX C compiler has changed from `-std0` to `-std`.

You can revert back to the previous default language mode by adding the `-std0` flag to the `cc` command line, the `/usr/ccs/lib/cmplrs/cc/comp.config` file or the `$DECC_CC` or `$COMP_HOST_ROOT` environment variable definitions.

8.7 Change in struct utmp, struct utmpx, and struct lastlog

To bring them into compliance with several UNIX and Internet standards, the `struct utmp`, `struct utmpx`, and `struct lastlog` structures have been changed. These changes affect the `/usr/include/utmp.h`, `/usr/include/utmpx.h`, and `/usr/include/lastlog.h` files :

- The time field in the `struct utmp` structure has changed from a `time_t` structure to a `struct __ut_timeval` structure (to be consistent with the `/usr/include/utmpx.h` file).
- The `ut_host` field size (in the `struct utmp` and `struct utmpx` structures) has been increased to comply with relevant Internet RFCs.
- The `ll_line` and `ll_host` manifest constants in the`/usr/include/lastlog.h` file have changed to allow their sizes to correspond to the `ut_line` and `ut_host` fields in `struct utmp` and `struct utmpx` structures.

These changes also affect the format of the `/var/adm/utmp`, `/var/adm/wtmp`, and `/var/adm/lastlog` files. The following conversion programs are supplied:

- `/usr/lbin/wtmpconvert`
- `/usr/lbin/llconvert`

The programs 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

As of Tru64 UNIX Version 5.0 the default size of the C language long double type changed from 64 bits to 128 bits. This allows applications to perform mathematical calculations on larger numbers and with more precision than was possible with the previous long double type. Prior to Version 5.0, the long double type was the same size as the double type (64 bits).

If an application was linked using an earlier version than Tru64 UNIX Version 5.0, there is one binary incompatibility that may affect the application. If the application was linked using the `-call_shared` switch, it can experience an incompatibility with the new default related to the input and output of long double types. Previously, `printf`, `scanf`, and similar functions interpreted the format code `%Lf` (capital L followed by f) as a 64-bit long double type. Now, the format code `%Lf` specifies a 128-bit long double type.

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 4.0G; 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 POSIX Threads (formerly 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. 1003.1-1996, POSIX System Application Program Interface provided by POSIX Threads. The POSIX 1003.1c standard interface is the most portable, efficient, and powerful programming interface offered by POSIX Threads. A compatibility mode for the Draft 4 POSIX 1003.4a API has been provided to ease migration. This compatibility mode will be removed in a future release.

8.11 DECthreads CMA Interface

The CMA interface of POSIX Threads (formerly DECthreads) is obsolete beginning with this release. Obsolescence means that while the CMA API continues to exist and is supported in Tru64 UNIX, CMA routines are no longer documented or enhanced. It is recommended that you port your CMA-based application to the IEEE Std 1003.1-1996, POSIX System Application Program Interface provided by POSIX Threads Library.

8.12 Asynchronous I/O Binary Compatibility

Data structures for asynchronous I/O, like `aio_read()` and `aio_write()`, changed between Version 3.2 and Version 4.0 of the operating system. The kernel currently handles these differences so that applications built under Version 3.2 continue to run when executed on Version 4.0x.

In the next major release of the operating system, support for applications built under Version 3.2x using asynchronous I/O will be discontinued. You will need to recompile and relink these applications to run under Tru64 UNIX Version 4.0D or higher.

8.13 Nemacs

Nemacs Version 3.3.2, a public domain Japanese implementation of Emacs, has been removed from this release of Tru64 UNIX. `Mule`, a public domain multilingual implementation of Emacs, has been carried forward as the replacement functionality for Nemacs. The Nemacs subsets `IOSJPNEMACS500` and `IOSJPNEMACSSRC500` 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 and tz 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 until the change occurs because a mechanism that ensures binary compatibility has been provided.

8.16 The -x and -p Options in addvol and mkfdmn

The `-x` and `-p` options to 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 release.

8.17 LSM Block Change Logging (BCL)

The Logical Storage Manager Block Change Logging (BCL) feature has been retired. It has been replaced with Dirty Region Logging (DRL). DRL logs regions instead of blocks for writes to LSM mirrored volumes. For most environments, DRL provides the same ability as BCL to quickly

resynchronize mirrors after a failure, but with considerably less logging overhead.

The logging format and configuration for DRL is not compatible with BCL. Therefore, this release provides support to automatically migrate and reconfigure mirrored volumes from BCL to DRL where possible. However, some existing mirrored volume configurations with BCL may require manual reconfiguration to continue to exploit logging for faster mirror recovery.

8.18 LSM volassist Command Syntax

In Version 5.0 of Tru64 UNIX, the syntax of the `volassist` changed. It no longer supports the following constructs:

- `[!]medianame,offset`

  You will still be able to use `[!]medianame,offset` construct but you cannot 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,...]`, replaces `alloc=size`. However, the new construct does not allow you to specify sizes for all allocations. You need to use the low-level commands to create subdisks, plexes, and volumes exactly as required.

- `align=size`

  Two new constructs, `diskalign` and `nodiskalign`, 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 need to use the low-level commands to create subdisks, plexes, and volumes exactly as required.

8.19 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, requiring you to install the optional subset. This optional subset will be removed from the product in a future release.

8.20 XIE Version 3.0 X Client Extension

Tru64 UNIX Version 4.0G supports XIE Version 5.0. Support for XIE Version 3.0 server extensions was removed in Version 4.0. Client support will be removed in a future release.

8.21 Microsoft Sound Board Driver

In Version 4.0 of the operating system, the device driver for the base audio on the AlphaStations and 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 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 obtain 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.

8.22 Graph Utility

The `/usr/bin/graph` utility will be removed in the next major release of Tru64 UNIX.

8.23 The atmsetup Script

The `atmsetup` script introduced in Version 4.0D has been superceded by a new application. The new application is part of the SysMan suite, and provides a full graphical user interface. The `atmsetup` command now invokes this new SysMan application.

You can access the `atmsetup` script by including the `-old` flag with the `atmsetup` command.

The `atmsetup` script will be retired in a future release of Tru64 UNIX.

For more information on how to use the new `atmsetup`, application see the `atmsetup`(8) reference page and the *Asynchronous Transfer Mode* guide.

8.24 Remote Prestoserve Support

Prestoserve support for remote operations has been retired. This means that the `-h` option of the `presto` command is no longer supported.

Users who require the ability to perform `presto` operations remotely can do so by using `telnet`, `rlogin`, or `rsh` to gain access to a shell on the remote system and then performing the operation locally.

8.25 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 after 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.

8.26 Different awk Versions

The `gawk` and `oawk` versions of the `awk` command have been retired. Only the XPG4-compliant version of the `awk` command is supported in this release. Scripts that use variations of the `awk` command should access the `/usr/bin/awk` version.

8.27 The secsetup Script

The `/usr/sbin/secsetup` script, used to configure Enhanced Security authentication support, has been retired in this release of Tru64 UNIX. It has been moved to the `OSFOBSOLETE` subset. The script has been replaced by the `secconfig` application. The `secconfig` application is a graphical user interface and is integrated with the SysMan utilities.

8.28 Open3D TURBOchannel and PVpci Device Support

The following Open3D devices have been retired in this release of Tru64 UNIX:

- PV-L
- PV-M
- PV-P

These devices are also known as:

- ZLX-L1, L2
- ZLX-M1, M2
- ZLXp-L1, L2

Open3D Version 4.4 is the last version of Open3D that supports the PV-L and PV-M cards. Open3D Version 4.9 is the last version of Open3D to support the PV-P graphic adapters.

Additionally, 3D support for the HX+ and TGA has been retired. These are also known as the ZLX-E1, E2, E3; and ZLXp-E1, E2, E3. DIGITAL UNIX 2D support will continue for the initial release of Tru64 UNIX Version 5.0 for these cards. The initial release of Tru64 UNIX Version 5.0 will also be the last release to support the HX and TX graphics adapters.

The replacement functionality is PCI-based systems and graphics devices supported on those systems.

8.29 System Management Utilities

The following setup scripts have been retired in this release:

- `bindsetup`
- `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 Command
`bindsetup`	Domain Name Service (DNS(BIND))	`sysman dns`
`netsetup`	Network Setup Wizard	`sysman net_wizard`
`nfssetup`	Network File System (NFS)	`sysman nfs`
`ntpsetup`	Network Time Protocol (NTP)	`sysman ntp`

Do not use these scripts to configure systems in a TruCluster environment. For more information, see the reference pages for the retired scripts or the *Network Administration* guide.

8.30 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.

8.31 Audit_setup Script

The /usr/sbin/audit_setup script, used to configure auditing support, will be retired in a future release of the operating system. It will be replaced by the audit_config, which is a SysMan application that provides a full graphical user interface.

8.32 VM Subsystem ubc_nfsloopback Variable Is Obsolete

The /etc/sysconfigtabubc_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 4.0G 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
```

8.33 libsecurity.a

The archive form of the enhanced security library (libsecurity.a) has been retired. If you want to link against libsecurity use the shareable form, libsecurity.so.

8.34 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-ROM (APCD).

8.35 Security Interfaces

Table 8–3 provides a list of interfaces that have been retired from the libsecurity library. Other, undocumented interfaces, have been retired as well. These interfaces are not capable of supporting new functionality offered in Tru64 UNIX 5.0.

Table 8–3: Retired libsecurity Interfaces

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	

The associated data structures listed in Table 8–4 have been retired as well.

Table 8–4: Retired libsecurity Interfaces Associated Data Structures

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
struct pr_default	struct dev_asg
struct system_default_fields	struct system_default_flags

8.36 The ogated Routing Daemon

The `ogated` daemon (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 the `ogated` daemon.

8.37 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.

8.38 NL* and NL* libc Interfaces

A Worldwide Portability Interface (WPI) for internationalization based on the XPG4 standard was introduced in Version 2.0 of the operating system. Because the WPI interfaces supersede 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.

8.39 DEC Ada RTL

DEC Ada (UPI - 0HM) and DEC Ada PDO (UPI - 0VS) will be retired in a future release of Tru64 UNIX.

8.40 System V Environment

Tru64 UNIX provides 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 re-engineering 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 customers who currently use the System V Environment layered product. The System V Environment (SVE) product is not available as a separately licensed layered product with the Tru64 UNIX Version 5.0 family. Instead, many of its features are being re-engineered and will be merged into the Tru64 UNIX operating system in future releases.

A

System Limits

This section lists the maximum system limits for the major components of this release. For hardware information specific to your individual processor, see the *Software Product Description* (SPD) and the *Systems and Options Catalog*. For information on how to tune system parameters, see the *System Configuration and Tuning* and the *System Administration*.

A.1 Backup Utility Limits

The backup utility limits are as follows:

- `cpio`

Files per archive:	4 GB
Files per file system:	4 GB
File size:	32 GB
File name size:	256 bytes
`dev_t`:	64 KB
`dev_major`:	Not used
`dev_minor`:	Not used

- `dd`

Files per archive:	Not used
Files per file system:	Not used
File size:	32 GB
File name size:	Not used
`dev_t`:	Not used
`dev_major`:	Not used
`dev_minor`:	Not used

- `dump`

Files per archive:	4 GB
Files per file system:	4 GB
File size:	32 GB
File name size:	No limit
`dev_t`:	4 GB
`dev_major`:	Not used
`dev_minor`:	Not used

- `tar`

Files per archive:	No limit
Files per file system:	No limit
File size:	8 GB
File name size:	256 bytes (with prefix)
`dev_t`:	Not used
`dev_major`:	16 MB
`dev_minor`:	16 MB

A.2 Device Addressing Limits

The device addressing limits are as follows:

- Device Access

  There are two types of disk device access: raw or character and block or buffered. For raw or character access, the structure field `uio.uio_offset` 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} 18 Exabytes. This value is converted to a physical block/sector number that is the data transfer start position. The physical block/sector number is limited by the structure field `buf.b_blkno`. For block or buffered access, the structure field `buf.b_blkno` describes the block/sector offset within the disk partition and is a signed 32-bit value. Since this release supports a fixed 512-byte block/sector size defined by `DEV_BSIZE`, the offset is limited to 1 TB.

- Major-Minor Numbers (`dev_t`)

  Devices are described by a major-minor pair of numbers, where the major number describes the device driver and the minor number describes the device. In this release these pairings are 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 (12 bits). Since each device driver requires 12 bits for its major number, 4096 device drivers may be statically configured into the system. The minor number portion of `dev_t` consists of bits 0 to 19 (20 bits). The first 6 bits (0 to 5) of the minor number are device-specific, describing the type of device that is connected to the system. For disks, these bits are known as the partition number, and theoretically they can describe a maximum of 64 partitions. This release restricts support to 8 partitions, however. Bits 6 to 19 of the minor number (14 bits) represent the unit number, and therefore this release supports a theoretical maximum of 16384 units; however, the actual number of supported units is a function of how many buses each individual hardware platform supports. On a DECstation 3000 Model 500, for example, the number of supported units is 14.

- SCSI/CAM Addressing

  Common Access Method (CAM) is an ANSI-proposed standard for a common software interface to Small Computer Systems Interface (SCSI). There are no restrictions or limitations within CAM for disk block addressing; 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 has 4 bytes reserved for the disk block address. This is an unsigned 32-bit value allowing 2^{32} - 1 or 4 Gigasectors of addressing, which corresponds to 2 TB given the 512-byte block/sector size. In this release the SCSI/CAM driver can address a maximum of 64 buses, with up to 7 device targets per bus, and a maximum of 8 LUNs per device target. As a result, in this release, SCSI/CAM can address a maximum of 3584 devices.

- Redundant Array of Independent Disks

  This release supports three Redundant Array of Independent Disks (RAID) controllers: two for the SCSI bus (HSZ10 and HSZ40) and one for the EISA bus (SWXCR). Each RAID device is seen by the operating system as a single target device (that is, as a single disk) with up to 8 Logical Unit Numbers (LUNs) on the SCSI controllers and 8 Logical Units (LUs) on the EISA controller, regardless of the number of disks on each RAID device. The HSZ10 SCSI RAID controller supports a maximum of 35 back-end disks; the HSZ40 SCSI RAID controller, a maximum of 42 back-end disks. As a result of hardware constraints, the maximum number of HSZ10 disks that can be concatenated into a logical volume is 5; the maximum number of HSZ40 disks that can be concatenated into a RAID 0 set for a logical volume is 14 with a total size limit of 32 GB. The SWXCR EISA RAID controller supports 4 packs per controller (a pack is a collection of disks over which data is striped) and 8 RE disk drives per pack, for a total of 32 RE disks per controller. Logical volume sizes are not fixed sizes as compared to other disk devices. The

size of a logical volume is configurable based on needs with a total size limit of 32 GB. In addition, the SWXCR controller may have either a one or a three channel SCSI adapter which supports 7 or 21 back-end SCSI disks, respectively. Although RAID theoretically increases the number of addressable disks significantly, Compaq recommends that the maximum number of devices for each system – even with RAID configured – should not exceed the numbers listed in the following section on device limits per processor.

- Disklabel

  The disklabel defines the partitions of a disk and their starting block/sector number. The starting block/sector number of a partition is defined by the structure field `partition.p_offset`, which is an unsigned 32-bit value allowing 2 TB of addressing with a 512-byte block/sector size.

A.3 Device Limits Per Processor

- SCSI Bus
  - Buses/adapter – 1 or 2
  - RZ/TZ devices/Bus – 7
- Digital Storage Architecture (DSA)
  - CIXCD controller – 1 to 4 HSCs
  - HSC – Up to the limit supported by each HSC, with a combined maximum of 96 RA/TA devices
  - KDM controller – 8 RA/TA devices up to 6 controllers
  - CI Star – A maximum of 16 nodes

  Note that DSA supports a maximum of 96 RA/TA devices.
- Processors
  - AlphaStation 200 series – 21 RZ/TZ devices (1 baseboard PCI) + (2 PCI SCSI Adapters) = 3 SCSI buses * 7 devices/bus
  - AlphaStation/Server 400 series – 28 RZ/TZ devices (1 baseboard PCI) + (3 PCI SCSI Adapters) = 4 SCSI buses * 7 devices/bus
  - AlphaStation 600 series – 35 RZ/TZ devices (1 baseboard PCI) + (4 PCI SCSI Adapters) = 5 SCSI buses * 7 devices/bus
  - AlphaServer 1000 series – 28 RZ/TZ devices (1 baseboard PCI) + (3 PCI SCSI Adapters) = 4 SCSI buses * 7 devices/bus
  - AlphaServer 8000 series – RZ/TZ devices
  - DEC 2000 series

- Model 300 – 28 RZ/TZ devices (4 Single SCSI Adapters) = 4 SCSI buses * 7 devices/bus
- Model 500 – 28 RZ/TZ devices (4 Single SCSI Adapters) = 4 SCSI buses * 7 devices/bus

– DEC 2100 series – 77 RZ/TZ devices (1 PCI baseboard single SCSI) + (3 Single SCSI PCI Adapters) + (1 EISA baseboard single SCSI) + (2 Tri-SCSI Adapters) = 11 SCSI buses * 7 devices/bus

– DEC 3000 series

- Model 800/800S – 98 RZ/TZ devices (1 baseboard dual SCSI) + (6 dual SCSI TURBOchannel adapters) = 14 SCSI buses * 7 devices/bus
- Model 600S – 56 RZ/TZ devices (1 baseboard dual SCSI) + (3 dual SCSI TURBOchannel adapters) = 8 SCSI buses * 7 devices/bus
- Model 600 – 42 RZ/TZ devices (1 baseboard dual SCSI) + (2 dual SCSI TURBOchannel adapters) = 6 SCSI buses * 7 devices/bus
- Model 500 – 98 RZ/TZ devices (1 baseboard dual SCSI) + (6 dual SCSI TURBOchannel adapters) = 14 SCSI buses * 7 devices/bus
- Model 500X – 84 RZ/TZ devices (1 baseboard dual SCSI) + (5 dual SCSI TURBOchannel adapters) = 12 SCSI buses * 7 devices/bus
- Model 400S – 56 RZ/TZ devices (1 baseboard dual SCSI) + (3 dual SCSI TURBOchannel adapters) = 8 SCSI buses * 7 devices/bus
- Model 400 – 42 RZ/TZ devices (1 baseboard dual SCSI) + (2 dual SCSI TURBOchannel adapters) = 6 SCSI buses * 7 devices/bus
- Model 300/300X – 35 RZ/TZ devices (1 baseboard single SCSI) + (2 dual SCSI TURBOchannel adapters) = 5 SCSI buses * 7 devices/bus
- Model 300L/300LX – 7 RZ/TZ devices (1 baseboard single SCSI) = 1 SCSI bus * 7 devices/bus

– DEC 4000 series 35 RZ/TZ devices (5 baseboard single SCSI) = 5 SCSI buses * 7 devices/bus

– DEC 7000/10000 series 96 RA/TA devices (DSA) 112 RZ/TZ devices (SCSI) (8 dual SCSI XMI Adapters) = 16 SCSI buses * 7 devices/bus

A.4 CPU Limits Per Processor

AlphaStation 200 series:	1
AlphaStation/Server 400 series:	1
AlphaStation/Server 600 series:	1

AlphaServer 1000 series:	1
AlphaServer 8000 series:	12
DEC 2000 series:	1
DEC 2100 series:	4
DEC 3000 series:	1
DEC 4000 series:	2
DEC 7000/10000 series:	6
DEC 8200 series:	6
DEC 8400 series:	12

A.5 File System Limits

- Logical Storage Manager

  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. In LSM, the term **volume** describes a virtual disk representing an addressable range of disk blocks used by applications such as file systems or databases. This release supports a maximum of 512 GB of disk space in a disk group or on a system, with a maximum supported volume of 512 GB. The maximum number of supported LSM volumes is 4093 for all disk groups in a 4093 for all disk groups in a system: 4091 non-system volumes and 2 system (root/swap) volumes. In LSM, the term **plex** describes the physical disk or disks that contain a complete copy of a volume's data. So, for example, a mirrored volume would be made up 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). In LSM, the term **subdisk** describes a contiguous portion of a physical disk which can be striped or concatenated together to form a plex. A maximum of 4096 subdisks can be associated with one plex, and Compaq supports 4096 subdisks per disk group or per system. LSM object names (such as volumes, plexes, subdisks, disk groups), volume attribute names (such as user and group), and `dxlsm view` names are limited to 14 characters.

- Advanced File System

  In the Advanced File System (AdvFS), a **volume** is any single logical device which can be a partition on a physical disk or a logical volume. A **domain** is a named set of bound volumes on which filesets are placed. A **fileset** is a named collection of files that is bound to a single domain. An **active fileset** is a fileset that has been mounted, like a mounted UFS file system, for example. Although the architectural maximum

limit of domains is 2048, in this release the AdvFS supports a maximum fileset and file size of 512 GB, up to 100 active file domains per system, and a maximum of 256 volumes per domain. Since a single disk failure in a domain can make the entire domain inaccessible, Compaq also recommends that you create no more than 8 volumes per domain. Note that while Tru64 UNIX supports an unlimited number of filesets per system, only 512 filesets can be mounted at one time. The number of files per fileset is 2^{32}, limited by the tag used to uniquely identify a file in a fileset. Note that over time the actual limit of files per fileset decreases, since a tag can only be used 4096 times due to a sequence number limit. Although AdvFS can support page sizes larger than 13 bits, the theoretical maximum AdvFS file and fileset size is 16 TB (2^{13} * 2^{31}) with a 13-bit page size and 31-bit page number.

- UNIX File System

  In this release, the UNIX File System (UFS) file size is limited by the amount of space that the kernel `buf` structure can address. The structure field `buf.b_blkno`, defined as `daddr_t`, is the block/sector offset within a disk partition and is a 32-bit signed value. The block or sector size, `DEV_BSIZE`, is 512 bytes. As a result, the theoretical maximum file system size that this release supports is 1 TB (2^{31} * 2^{9}). Note that Compaq only supports 128 GB. The UFS file system and file size is limited by the maximum logical volume size supported by the Logical Volume Manager.› (the maximum logical volume size supported by the Logical Storage Manager).

- CD–ROM File System

  The size of the CD–ROM File System (CDFS) files and file systems is limited by the compact read-only optical disk (CD–ROM) media where they reside. Currently, the CD–ROM media supports approximately 0.60 GB. However, Tru64 UNIX is able to support larger CD–ROMs should they become available.

- Sparse Files

  Tru64 UNIX supports **sparse files** on AdvFS and UFS, which means that the size of a file may exceed the size of the file system where it resides. Tru64 UNIX supports the following maximum sizes for sparse files:

AdvFS	2^{43} +2
UFS	2^{44}-8K

  the operating system.

- Network File System

In this release, the theoretical maximum size of a file that is accessible through the Network File System (NFS) is as follows:

- NFS Version 2.0 – (2 GB - 1 byte)
- NFS Version 3.0 – 16 Exabytes (2^{64}-1)

Compaq supports the following maximum file sizes:

- NFS Version 2.0 – 2 GB
- NFS Version 3.0 – 512 GB

Note that an NFS server is always limited by the size of the underlying local file system.

- Memory Mapped File Limit

  The supported maximum size of a file that can be mapped into memory without segmenting the file is 1 GB.

- Mounts

  – Advanced File System

    The Advanced File System (AdvFS) supports a maximum of 512 mounted filesets. However, each active domain has an invisible mounted fileset associated with it which must be factored into the total number of mounted filesets. So, for example, if you have an active domain with two mounted filesets, the invisible fileset associated with the domain itself brings the total number of mounted filesets to three.

  – UNIX File System

    The UNIX File System (UFS) supports a total of 512 mounts, which are now allocated dynamically by the system rather than being dependent on statically configured mount tables as they were in previous releases of the operating system.

  – CD–ROM File System

    This release supports a maximum of 512 CD–ROM File System (CDFS) mounts.

  – Network File System

    A Network File System (NFS) Version 2.0 or Version 3.0 client can mount a maximum of 2048 files or directories. The `vnodes` necessary to support the NFS-mounts are now allocated dynamically rather than being dependent on a statically configured `vnode` table as they were in previous releases of

- Open files

  The maximum number of files a process can open is set to 4096 by default in the `OPEN_MAX_SYSTEM` variable in the file `/usr/sys/include/sys/param.h`. This number can be adjusted

between 64 and 4096, either in individual programs by using the `setrlimit`(2) system call or on a system-wide basis by editing the file `/usr/sys/conf/param.c`, changing the `open_max_soft` and `open_max_hard` variables, and then relinking or rebuilding the kernel. Note that file descriptor entries in the per process file table are dynamically allocated after the initial 64 entries in the `utask` structure are used.

A.6 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. Note that since the NFS Version 2 protocol suite only allows ranges to be specified with 32-bit numbers, it supports a file locking range of 0 to 2^{31}-1 inclusive.

A.7 Pathname Limits

AdvFS, UFS, CDFS, and NFS support a maximum pathname component of 255 bytes and a maximum file pathname of 1023 bytes.

A.8 Installation Limits

This section lists the disk space required for `root`, `/usr`, and `/var` when performing a basic or an advanced installation of this release.

- Basic Installation

```
root:30 MB
/usr:  148.5 MB
/var:  -
```

- Advanced Installation (taking all subsets)

```
root:  40 MB
/usr:  320 MB
/var:  6 MB
```

A.9 Memory Limits

- Physical Memory

  The maximum supported memory is different for each individual processor, although the DEC 7000/10000/8200/8400 series a total of 14 GB of physical memory. The operating system, however, supports a total

of approximately 4 GB of physical memory. For more information on supported memory, see the System and Options Catalog and the SPD.

- Virtual Memory
  - Per process

    The default virtual memory per process is 2 GB, although available swap space may in many cases be exhausted before this limit is reached. This value can be increased to a maximum of 8 TB by defining the `MAXVAS` variable in the system configuration file and relinking or rebuilding the kernel.
  - Page size

    The default page size is 8 KB and is not configurable. The page size is hardware dependent and is set up by the console at boot time.

A.10 Networking Limits

- Pseudoterminals (`ptys`)

  The maximum number of supported `ptys` is 1400.
- IP Alias Addresses

  Tru64 UNIX allows you to specify up to 1024 IP alias addresses before system performance begins to degrade.

A.11 Port Limits

Although there is no real limit to the number of connections that are possible to each port, there is a 32K limit to the number of network service ports per host. It is of course likely that your system will run out of memory before this 32K limit is reached.

- External Connections

  The `telnet` utility is allotted 1023 ports; however, owing to reserved ports for such protocols as mail and ftp, you can realistically `telnet` out of a Tru64 UNIX system approximately 800 or 900 times.
- Network Transfer Rates

  For information on network transfer rates, see the *Technical Overview*.

A.12 Process Limits

- Per system

  The number of processes per system depends on the value of `MAXUSERS`, which is configurable and set in the configuration file to 32 by default. With `MAXUSERS` at its default value, the number of processes per system is set to 276 in the `NPROC` variable in `/usr/sys/conf/param.c`. You

can increase this value by either changing the value of MAXUSERS in the system configuration file or by adding the maxproc variable to the system configuration file and relinking or rebuilding the kernel. You might increase the value of MAXUSERS to allow more users to log in to your system or to allow applications that run as root to fork more processes than NPROC allows by default. However, increasing the number of processes per system reserves more system memory, so the upper limit of NPROC is dependent upon the system's total memory, the number of actual users on the system, and the requirements of your applications.

- Per user

  The number of processes that each user can fork is set to 64 by default through the CHILDMAX variable in the file /usr/sys/include/sys/syslimits.h. The number of processes per user can be varied by adding the maxuprc variable to the system configuration file, setting its value to some number of processes, and then relinking or rebuilding the kernel. You might increase this value if you had an application that needed more processes than CHILDMAX is set to by default. However, increasing the value of maxuprc reserves more system memory, so the upper limit of maxuprc is dependent upon the system's total memory, the number of actual users on the system, and the requirements of your applications.

Index

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

How to Order Tru64 UNIX Documentation

You can order documentation for the Tru64 UNIX operating system and related products at 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-6ADAA-G8
Tru64 UNIX Documentation Kit	QA-6ADAA-GZ
End User Documentation Kit	QA-6ADAB-GZ
Startup Documentation Kit	QA-6ADAC-GZ
General User Documentation Kit	QA-6ADAD-GZ
System and Network Management Documentation Kit	QA-6ADAE-GZ
Developer's Documentation Kit	QA-6ADAG-GZ
Reference Pages Documentation Kit	QA-6ADAF-GZ

Reader's Comments

Tru64 UNIX
Release Notes for Version 4.0G
AA-QTLMG-TE

Compaq 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) 884-0120, 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)	☐	☐	☐	☐
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, department _____
Mailing address _____
Electronic mail _____
Telephone _____
Date _____

--------- Do Not Cut or Tear - Fold Here and Tape ----------

COMPAQ

NO POSTAGE
NECESSARY IF
MAILED IN THE
UNITED STATES

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 03062-2698

--------- Do Not Cut or Tear - Fold Here ----------

Cut on This Line