DIGITAL UNIX

Release Notes for Version 4.0E

Order Number: AA-QTLMD-TE

November 1998

Product Version: DIGITAL UNIX Version 4.0E

This book contains notes on software and documentation restrictions for Compaq's DIGITAL UNIX Version 4.0E and the bundled layered products that ship with the operating system.

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

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

`http://www.unix.digital.com/faqs/publications/updates/V40E-update.html`

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

Compaq Computer Corporation
Houston, Texas

Contents

About This Manual

1 New and Changed Features

2 Installation Notes

3 Processor-Specific Notes

4 Base System Software Notes

5 Development Environment Notes

6 Window System Software Notes

7 Documentation Notes

8 Features and Interfaces Scheduled for Retirement

A Maximum System Limits

B Software Subset Information

C Enabling and Disabling Extended UID/GID Support

Index

Examples

Tables

About This Manual

This manual contains release notes for Compaq DIGITAL UNIX® Version 4.0E.

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

Audience

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

Organization

This manual is organized as follows:

Chapter 1	Contains an overview of new and changed features.
Chapter 2	Contains notes pertaining to installation.
Chapter 3	Contains processor-specific information.
Chapter 4	Contains information about the base operating system software.
Chapter 5	Contains information about the development environment.
Chapter 6	Contains information about the window system software.
Chapter 7	Contains information about the documentation.
Chapter 8	Contains information about features scheduled for removal in future versions of DIGITAL UNIX.
Appendix A	Contains information about the maximum system limits.
Appendix B	Contains information about the software subsets for the DIGITAL UNIX software kit.
Appendix C	Contains information about the new extended UID and GID support in this release.
Appendix D	Contains information about installing and configuring Netscape Communicator and Netscape FastTrack Server.
Appendix E	Contains information about how to increase the maximum number of open files per process.

Appendix F	Contains information about how to solve problems that you may encounter with the Advanced File System (AdvFS).
Appendix G	Contains information about how to enable the enhanced core file naming.
Appendix H	Contains information about early child process exit notification.
Appendix I	Contains information about Thread Local Storage (TLS) support in DEC C.
Appendix J	Contains information about creating a consolidated CD–ROM.

Related Documents

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

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

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

  This document includes the console firmware revision numbers and release notes.
- DIGITAL UNIX Version 4.0B *Installation Guide*

  This document describes in detail how to install DIGITAL UNIX Version 4.0B. Most of this information also applies to Version 4.0E.
- *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 DIGITAL UNIX systems. It also contains information on setting up Dataless Management Services (DMS).
- *System Administration*

  This document contains information on administering and maintaining your system.

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

describes this convention:

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

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

The *Documentation Overview, Glossary, and Master Index* provides information on all of the books in the DIGITAL UNIX documentation set.

Reader's Comments

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

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: UEG Publications, ZK03-3/Y32
- Internet electronic mail: `readers_comment@zk3.dec.com`

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

  `/usr/doc/readers_comment.txt`

- Mail:

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

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

Please include the following information along with your comments:

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

- The section numbers and page numbers of the information on which you are commenting.
- The version of DIGITAL UNIX that you are using.
- If known, the type of processor that is running the DIGITAL UNIX software.

The DIGITAL 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 DIGITAL technical support office. Information provided with the software media explains how to send problem reports to DIGITAL.

Conventions

The following conventions are used in this guide:

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

New and Changed Features 1

The Compaq DIGITAL UNIX Version 4.0E software is a replacement for the Version 4.0A, Version 4.0B, Version 4.0C and Version 4.0D operating system software and provides new or additional support for the following:

- Maintenance for previous versions of DIGITAL UNIX
- New and enhanced features, including the following:
  - AdvFS support for atomic write data logging
  - A method for preventing partial AdvFS data writes
  - Support for the Euro currency sign
  - Support for Redundant Array of Independent Network adapters (NetRAIN) in Available Server Environment (ASE) configurations
  - Support for the extensible SNMP agent API for asynchronous processing
  - A network enhancement to improve TCP connections for the TruCluster product
  - A new sys_check tool to collect system data
  - Enhanced security for TruCluster Version 1.5
  - Enhancement to the `disklabel` command
  - Support for ODBC and JDBC data access
  - Ability for users to create a customized installation CD–ROM.
- Hardware-related support, including the following:
  - Support for new AlphaServer processors
  - Support for the Univeral Serial Bus (USB) class driver
  - Support for the Intraserver UNIX driver
  - Support for Gigabit Ethernet
  - Support for KZPCM/LVD
- Software maintenance updates and new offerings, including the following:
  - Enhancement to the `sendmail` program
  - A new Prestoserve subsystem attribute

- Addition of arena malloc capabilities
- New versions of Netscape Navigator and Netscape FastTrack

For more information about some of these features, see the following sections. You should also refer to Chapter 8, which describes features and interfaces scheduled to be retired in future versions of the operating system.

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

1.1 New Processor Support

DIGITAL UNIX Version 4.0E adds support for the Compaq AlphaServer GS60 and GS140 processors, formerly known as AlphaServer 8200/8400.

1.2 USB Driver Support

DIGITAL UNIX Version 4.0E provides Universal Serial Bus (USB) class driver support for DIGITAL UNIX.

1.3 AdvFS Now Supports Atomic Write Data Logging

Writes of up to 8192 bytes in size may now be guaranteed to be atomic. That is, in the event of a system crash, either all of the `write()` buffer contents or none of it will be present in the file. It will not be possible that only part of the `write()` buffer contents will be present in the file. See Section 4.6.5.7 for more details.

1.4 Preventing Partial AdvFS Data Writes

To prevent partial AdvFS writes if a system crash occurs, you can use the `chfile -L on` command to enable atomic write data logging for a specified file. See Section 4.6.5.8 for more details.

1.5 Euro-character Support

DIGITAL UNIX Version 4.0E provides initial support for the new Euro currency sign. The Euro is a new pan-European currency that is scheduled to be introduced on January 1, 1999 and should completely replace local European currencies by the end of 2002.

The support for Euro character is provided by using the Unicode character set. A number of new European UTF-8 locales are introduced in the Worldwide Language Software subsets in this release. UTF-8 is one of the transformation formats for Unicode and is a multibyte character set in nature.

Applications that want to support the new Euro currency symbol have to be modified to support Unicode as process code and UTF-8 as file code.

See the `Euro`(5) and `Unicode`(5) reference pages for more information on Euro character support in this release.

1.6 Arena Malloc

DIGITAL UNIX Version 4.0E adds arena `malloc` capabilities to the DIGITAL UNIX `malloc` suite.

The following are new interfaces in `libc`:

- `acreate()`
- `adelete()`
- `amalloc()`
- `afree()`
- `arealloc()`
- `acalloc()`
- `arecalloc()`
- `amallinfo()`
- `amallopt()`

1.7 NetRAIN

The use of Redundant Array of Independent Network adapters (NetRAIN) eliminates the network adapter as a single point of failure on a DIGITAL UNIX system. NetRAIN functionality will be supported for client access for both ASE and Production Server and for cluster communications in an ASE configuration.

The NetRAIN virtual interface configures multiple interfaces on the same LAN segment into a single interface. One of the real interfaces is always active while the others remain idle. If the active interface fails, an idle interface is brought online.

For more information about the NetRAIN functionality, see Section 4.4.8.

1.8 Asynchronous Processing Support in the Extensible SNMP Agent API

DIGITAL UNIX Version 4.0E adds support in the extensible SNMP agent API (`libesnmp.so`) for asynchronous processing. This is accomplished by adding two new routines, `esnmp_poll_pdu()` and

`esnmp_pdu_response()`.

1.9 NFS Over TCP

DIGITAL UNIX Version 4.0E provides a required network enhancement to clean up the TCP connections for the TruCluster product. Without this functionality, certain error conditions are possible as a result of relocation of an NFS service.

1.10 Intraserver ITPS Device Driver

The Intraserver UNIX (ITPSA) driver is integrated in DIGITAL UNIX Version 4.0E, to support the new KZPAA-EB (SYMBIOS Low Voltage Differential (LVD)) UltraSCSI adapter board as well as the KZPCM-DA (ITI Combo adapter).

1.11 Gigabit Ethernet

DIGITAL UNIX Version 4.0E provides support for Gigabit Ethernet by supporting a driver for the Alteon Gigabit Ethernet adapter.

1.12 sys_check Tool

DIGITAL UNIX Version 4.0E provides a `sys_check` tool to collect system data.

An enhanced interface to `sys_check` is available. This interface provides a command line interface, a character cell interface, a graphical interface, and a Web-based interface to the `sys_check` command.

To run the sys_check interface, enter the following command:
`/usr/sbin/sysman escalation`

This can also be started by double-clicking the Escalation Report icon in the CDE Application Manager's System_Admin/Support category.

For information on the `sys_check` tool, see the online help of the graphical interface, or the `sys_check`(8) reference page.

1.13 Sendmail Enhancement

The `sendmail` program has been upgraded to Version 8.8.8 as the default version for DIGITAL UNIX. For more information, see Section 4.1.9 and the `sendmail`(8) reference page.

1.14 New Prestoserve Subsystem Attribute

Prestoserve now supports the `presto-buffer-hash-size` attribute. See Section 4.6.9 for details.

1.15 Enhanced Security for TruCluster Version 1.5

DIGITAL UNIX Version 4.0E supports a large (greater than 30K users) distributed authentication database with enhanced security for TruCluster Version 1.5. This solution has no single point of failure and supports failover of the NIS master in a TruCluster Version 1.5 environment.

1.16 Enhancement to disklabel Command

The `/sbin/disklabel` command supports operations on an image file. See Section 4.1.7 for more information.

1.17 Support for ODBC and JDBC Data Access

DIGITAL UNIX Version 4.0E includes the software to support Open Database Connectivity (ODBC) and Java Database Connectivity (JDBC) data access for UNIX and Windows NT applications. DIGITAL UNIX Version 4.0E includes INTERSOLV's DataDirect software on the Associated Products CD–ROM Volume 1.

DIGITAL UNIX customers have full use of this software on their applications and databases on DIGITAL UNIX systems. This software helps deliver the industry's best integration between UNIX and Windows NT for data access. This software is also available at:

`http://www.unix.digital.com/data-access`

1.18 Netscape Communicator Version 4.05

This release of DIGITAL UNIX contains Version 4.05 of Netscape Communicator, which integrates the Netscape Navigator World Wide Web browser with Netscape's Messenger Mailbox, Collabra Discussions, and Page Composer components. For information about how to install, configure, and invoke Communicator, see Section D.1 of this book and Chapter 6 of the DIGITAL UNIX *Installation Guide*.

1.19 Netscape FastTrack Server Version 3.01

This release of DIGITAL UNIX contains Version 3.01 of the Netscape FastTrack Server, an easy-to-use entry-level Web server designed to let you create and manage a Web site. For information about how to install and

configure FastTrack, see Section D.2.

1.20 Ability for Users to Create a Customized Installation CD–ROM

If you need to update many identical systems, this release of DIGITAL UNIX enables you to create a customized installation on a CD–ROM, using your own CD–ROM writer.

The customized CD–ROM will consolidate the latest firmware upgrade for a processor, and the operating system software into an ISO 9660-compliant CDFS image file. During installation, the `fwupgrade` utility is used to update the firmware.

Refer to Appendix J for instructions.

Installation Notes 2

The notes in this chapter discuss the following topics:

- General information about installation
- Update installation
- Server extensions

Do not attempt to install DIGITAL UNIX Version 4.0E 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 General Information About Installation

The following notes apply to the installation process in general.

2.1.1 Verifying the Installation

To verify that you have installed all the correct elements from the DIGITAL UNIX Version 4.0E software kit, check the DIGITAL UNIX revision level that appears in the `/etc/motd` file. The correct version string is:

```
DIGITAL UNIX 4.0E
```

2.1.2 Disk Space Requirements

For information about disk space requirements, see Appendix A, Appendix B, and Section 8.27. If you want to add optional subsets after you install DIGITAL UNIX Version 4.0E, use the `df` command to determine free disk space in blocks.

2.1.3 New Recommended Disk Partition Table

Starting in DIGITAL UNIX Version 4.0D, a recommended disk partition table is available for 1 GB (for example, RZ26) and larger disks during a full text-based or graphical installation. You have the option of applying the recommended disk partition table or using your existing disk partition table when you choose to install the operating system onto a single disk. If your

operating system is installed across multiple disks, this option will not be offered to you.

For single disk installations, you should choose the recommended disk partition table over the default partition table.

The recommended partition table creates the partitions listed in Table 2-1. Note that these partition sizes are the recommended sizes for Version 4.0E. They are subject to change in future releases as the size of DIGITAL UNIX continues to grow.

Table 2-1: Recommended Partitions

Partition	Description
a	The size of the a partition is 128 MB, regardless of the size of the disk. For example, a 1 GB disk and a 4 GB disk would both have a 128 MB a partition if you select the recommended disk partitions.
b	The size of the b partition is either its default value for a given disk or 128 MB, whichever is larger. For example, an RZ28 has a default b partition size of 196 MB. This default size will be maintained, instead of reducing it to 128 MB.
g	The size of the g partition is 700 MB, regardless of the size of the disk. For example, a 1 GB disk and a 4 GB disk would both have a 700 MB g partition if you select the recommended disk partitions. Note that if you will be installing many layered products, this value may be too small. Also, if sufficient disk space is available, you may want to consider making this value larger for future expansion.
h	The size of the h partition is whatever space is left over, unless it is less than 100 MB, in which case the space is added to the size of g. In the case of a 1 GB disk, approximately 45 MB would be left over, so it would be included in the size of g instead of h.

Partitions d, e, and f are evenly split between the size of g added to h, and they overlap g and h.

If the disk has an existing partition table with a, b, and g partitions each greater than 90% of their recommended sizes, then the existing partition table is accepted as the recommended partition table.

If you use `installupdate` to update an existing system, you will not be offered the new partition sizes because the procedure updates the system on your existing partitions. You may wish to adjust your disk partitions to meet

or exceed the recommendations in Table 2-1 before you begin the update. If you have a small system disk, you may want to migrate to a 1 GB or larger disk at this time.

2.1.3.1 New CDF Item Introduced

A new item, `Inst_disklabel`, has been introduced in the Configuration Description File (CDF) to support the recommended disk partition table. (You can use CDFs to perform installation cloning as described in Appendix C of the DIGITAL UNIX Version 4.0B *Installation Guide*.)

The `Inst_disklabel` item has several attributes associated with it. Table 2-2 lists these attributes.

Table 2-2: Inst_disklabel Attributes

Attribute	Description
`name`	A required attribute that specifies the software name of disk to which the recommended partition will be applied (for example, `rz0`)
`a_size`	The size of the `a` partition in 512-byte blocks
`a_offset`	The offset of the `a` partition (from block 0) in 512-byte blocks
`b_size`	The size of the `b` partition in 512-byte blocks
`b_offset`	The offset of the `b` partition (from block 0) in 512-byte blocks
`c_size`	The size of the `c` partition in 512-byte blocks
`c_offset`	The offset of the `c` partition (from block 0) in 512-byte blocks
`d_size`	The size of the `d` partition in 512-byte blocks
`d_offset`	The offset of the `d` partition (from block 0) in 512-byte blocks
`e_size`	The size of the `e` partition in 512-byte blocks
`e_offset`	The offset of the `e` partition (from block 0) in 512-byte blocks
`f_size`	The size of the `f` partition in 512-byte blocks
`f_offset`	The offset of the `f` partition (from block 0) in 512-byte blocks
`g_size`	The size of the `g` partition in 512-byte blocks
`g_offset`	The offset of the `g` partition (from block 0) in 512-byte blocks

Table 2-2: (continued)

Attribute	Description
h_size	The size of the h partition in 512-byte blocks
h_offset	The offset of the h partition (from block 0) in 512-byte blocks

When performing installation cloning, you can specify multiple Inst_disklabel items so that several disks can be repartitioned automatically during the cloning process, based on the values contained within the individual items.

The following example shows how an Inst_disklabel item might appear within a CDF file:

```
install:
        _item=Inst_disklabel
        name=rz1
        a_size=262144
        a_offset=0
        b_size=262144
        b_offset=262144
        g_size=1090979
        g_offset=524288
        h_size=435593
        h_offset=1615267
        _action=create
```

2.1.4 RIS boot -fl n Option

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
```

2.1.5 Worldwide Subsets

Sometimes the following error message is displayed during the installation of

worldwide subsets and the installation procedure aborts:

```
/usr/sbin/setld: /usr/lbin/depord: arg list too long
```

This problem occurs when too many subsets are chosen and the shell runs out of memory during installation. You can work around this problem by choosing fewer subsets to install at first, and then installing additional subsets at a later time.

2.1.6 I/O Error Message

During the installation process, you may encounter the following benign message:

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

The installation will complete successfully and the message can be ignored.

2.1.7 Broken Pipe Messages

The installation process may display the following messages during a RIS installation of some subsets:

```
Broken Pipe
```

The installation will complete successfully and the message can be ignored.

2.2 Update Installation

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

Version 4.0E supports update installations from DIGITAL UNIX Version 4.0B, Version 4.0C, and Version 4.0D. See the *DIGITAL UNIX Software Product Description* for more information about disk space requirements before executing `installupdate`.

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

The following sections apply to the update installation procedure.

2.2.1 Warning from Update of DDR Database

The DDR database format has been updated in Version 4.0D to include previously unrecognized SCSI device types and IDE/ATAPI entries.

If you are updating from DIGITAL UNIX Version 4.0B or Version 4.0C to Version 4.0E, during the update installation process, the DDR database is converted to the new format, but the running kernel is still based on the older DDR database format. Hence, when the kernel is notified of the DDR

update, it reports the apparent discrepancy in database formats by sending the following messages to the console:

```
DDR: Invalid Database Header
DDR: Invalid Database Header
```

You can ignore these messages. At the end of the update installation, a new kernel is built with the updated `ddr.mod` binary and the resulting kernel properly recognizes the database.

After the update installation, there should be no further warnings from the DDR subsystem, unless the system is booted with an older preupdate kernel. If this happens, the kernel sends the same messages to the console and indicates that it is using the defaults for each device.

This problem does not occur if you are updating from DIGITAL UNIX Version 4.0D to Version 4.0E.

2.2.2 Preparing AdvFS File Domains Before Upgrading

To protect your data on AdvFS file domains, perform the following procedure on them before you update your operating system to Version 4.0E. All steps in the procedure must be executed by a user with `root` privileges:

1. Use the `shutdown` command to bring the system to single-user mode.
2. Use the `umount` command to unmount all local file systems.
3. Run the `verify` utility on each local file system and if there are problems with the file systems, correct them before going further.
4. Use the `mount` command to mount all of the verified local file systems.
5. Use the `quotacheck` command to fix the quotas on the mounted local file systems. If there are problems running `quotacheck` on a local file system, you may have to edit the `/etc/fstab` file to fix the problem.

After successfully completing these steps, you can update your system. DIGITAL strongly recommends that you fix any problems before you update.

2.2.3 Installation of AdvFS Advanced Utilities in the DMS Environment

When installing the AdvFS Advanced Utilities Version 4.0E into a DMS software environment, subset AFAADVDAEMON435 may fail with the following error:

```
Loading 1 of 1 subset(s)....
awk: Cannot find or open file ./etc/services.

The source line number is 2.

Advanced File System Daemon, v4.0e, r435
```

```
Copying from /mnt/AFA435/kit (disk)
Verifying

1 of 1 subset(s) installed successfully.

Backing up ./etc/services to ./etc/services.advfsd.back....
cp: ./etc/services: No such file or directory
Backup failed.  Root filesystem is full.
 ./etc/services file not modified.

Aborting installation....
setld:
"Advanced File System Daemon, v4.0e, r435" (AFAADVDAEMON435) failed in
subset control program (POST_L).
```

If this occurs, select the DMU option `C` (CONFIGURE software environments) to configure the target software environment, and then reinstall the failed daemon subset.

2.2.4 Conflict with PATHWORKS in an Update Installation

The PATHWORKS product has been added to the list of layered products that will halt an update installation if detected during the analysis phase of the update. If the update installation detects the PATHWORKS product, it will output a message and exit without modifying the system. In order to complete the update installation, you must manually remove the PATHWORKS product. Refer to section 2.8.2.1 of the *Installation Guide* for details about layered products that halt the update installation and how to remove those products.

2.3 Server Extensions

The following sections provide information on installing DIGITAL UNIX Server Extensions.

2.3.1 Error When Installing Chinese Subsets from RIS Server

If you install the Chinese subsets from a RIS server, the following error may occur:

```
Common Chinese Unicode Support
Copying from server-name (inet)
setld: cannot access server mapping (rcp: ris3.alpha/rp_mapping:
No such file or directory)
setld: Load from server-name failed, subset IOSZHUCSBASE425
Verifying
setld:
There were verification errors for "Common Chinese Unicode Support
(IOSZHUCSBASE425)
```

This error causes the system to abort the installation process. The problem may happen in other Chinese subsets also.

The error occurs because of the large number of software subsets in the Worldwide kit.

To avoid this problem, create a RIS area with fewer subsets or Chinese subsets only. To do this, choose option 1 from the following menu when installing subsets to the RIS area:

```
Choose one of the following options:

1)  Extract software from [kit location]
2)  Create symbolic link to [kit location]

Enter your choice:
```

Do not choose to extract all software subsets; only extract those subsets that are necessary for installation.

Processor-Specific Notes 3

This chapter contains notes that apply to the following computers:

- General notes on processors
- DIGITAL Personal Workstation 433au, 500au, and 600au systems
- DIGITAL AlphaServer 1000 and 1000A systems
- DIGITAL Alpha VME 4/224, 4/288, and 5/nnn single board computers
- DIGITAL DMCC EBM2n class and EBM4n class single board computers
- DIGITAL AXPvme single board computers
- Compaq AlphaServer GS60 and GS140 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 DIGITAL 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 Chapter 2 of the DIGITAL UNIX *Installation Guide* for rebuilding your system kernel.

3.1.2 PCI-Based Systems with ATI Mach64 Graphics Controllers

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

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

3.1.3 PCI Shared Interrupt Support

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

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

```
>>> show config
```

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

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

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

3.1.4 KZPSA Behind the PCI-to-PCI Bridge

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

3.1.5 Qlogic ISP1040B CAM Errors

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

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

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

3.1.6 Monitor Power Management

Monitors that do not support the Display Power Management Signaling (DPMS) feature available on DIGITAL Personal Workstation 433AU, 500AU, 600AU, AlphaStation 255, and other systems can be damaged by the activation of monitor power management. Check your monitor specifications before using this feature.

The time it takes a DPMS-compatible monitor to come out of a power-saving state depends on the monitor. Users will observe that the longer the monitor is in the power-off state, the longer it takes for the display to return as a result of mouse or keyboard activity. This is the result of the monitor phosphor cooling down and the time required to heat it back up, and is not under the control of DIGITAL UNIX.

In addition, there are some monitors (for example, the VRC21-HX) that turn off the Power On/Off LED when they enter the DPMS_OFF state. Moving the mouse or typing at the keyboard will not bring the display back. Only by pressing the power switch off, then on again, will mouse and keyboard activity cause the display to return. Because of the varying behavior of monitors when in certain DPMS states, you should read your monitor specification to find out about the expected behavior and other visual features while in each power-savings state.

For more information about enabling power management, see the `dxpower`(8) reference page and the *System Administration* guide.

3.1.7 DJ-ML200-xx PCI Nvram Hardware Rev Must Be E01

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

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

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

3.2.1 EIDE (ATAPI) CD–ROM Support

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

Version 4.0E supports the use of ATAPI CD–ROM devices attached to the internal EIDE adapter of the DIGITAL 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. The SCSI buses on the system are numbered `rz16` and above.

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

3.2.2 64-Bit PCI Option Cards

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

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

3.2.3 Incorrect Default Keyboard Mappings

If you use a PCXLA-NA keyboard on a DIGITAL 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 from the command line:

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

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

```
gpc_input:  kbd_scancode = 2
```

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

3.3 DIGITAL AlphaServer 1000 and 1000A Systems

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

3.3.1 EISA Configuration Utility Diskette Version 1.10

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

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

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

3.3.2 Graphics Resolution

The default graphics resolution for DIGITAL 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, DIGITAL UNIX supports resolutions of 640x480 (by default) or 800x600 only.

The default resolution for DIGITAL 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.con files:

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

Change the line to:

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

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

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

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

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

3.4.1 Verify CPU Version

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

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

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

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

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

3.4.2 Firmware Requirements

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

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

command at the console firmware prompt:

```
>>> show version
```

3.4.3 Restrictions and Known Problems

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

3.4.3.1 EBVXG (TGA) Video Card Requires Disabling X Server DMA

3.4.3.2 Byte Addressability Not Supported by mmap (5/nnn only)

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

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

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

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

3.4.3.3 VME Autovectors

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

3.4.3.4 Network Port Termination

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

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

3.4.4 Writing VMEbus Device Drivers

For information about writing VMEbus device drivers and configuring the VMEbus on DIGITAL Alpha VME systems, refer to the DIGITAL UNIX manual *Writing VMEbus Device Drivers*. This manual is included in the DIGITAL UNIX *Device Driver Kit*.

3.5 DMCC EBM2n and EBM4n Class Single Board Computers

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

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

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

3.5.1 Verify CPU Version

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

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

The `sizer -c` command displays the following output for DIGITAL DMCC EBM4n class SBCs:

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

3.5.2 Firmware Requirements

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

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

console firmware command at the prompt:

```
>>> show version
```

3.5.3 Restrictions and Known Problems

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

3.5.3.1 Option Card Restrictions

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

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

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

All ETMAB-xx backplanes are PCI V2.1 compliant.

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

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

Table 3-1: Supported Options Behind the Bridge

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

Table note:

a. The SN-PBXGB-AA (TGA2 PowerStorm 3D30) video card will work behind a bridge in multiple configurations if the first card is within the primary bus. For restrictions on jumper settings and X Server DMA for the PowerStorm 3D30 card, see Section 3.5.3.2.

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

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

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

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

- EV5 Alias jumper setting (EBM2n only)

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

- VGAEN jumper settings

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

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

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

  Some configurations of PowerStorm 3D30 cards on DIGITAL EBM2n and EBM4n SBCs require that you disable X Server direct memory access (DMA) write operations. Specifically, you must disable these operations if there are multiple PowerStorm 3D30 cards in the system, or

if any PowerStorm 3D30 graphics card is installed behind the PCI-PCI bridge (PPB).

3.5.3.3 General Restrictions

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

3.6 DIGITAL AXPvme Single Board Computers

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

The following notes are specific to the DIGITAL AXPvme SBCs.

3.6.1 Firmware Upgrade Required

DIGITAL AXPvme SBCs must upgrade to Version 17.0 or higher of the AXPvme firmware to run DIGITAL UNIX Version 4.0E.

3.6.2 Unpredictable Results When Performing Master Block Transfers

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

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

- If D64 slave access is performed before memory has been mapped to the VMEbus.
- If memory access does not coincide with the appropriate access mode, such as attempting user access to memory specified as supervisory mode access.
- If the AXPvme SBC or Alpha VME 2100 system is a VME interrupter and is targeted for D64 slave access. The interrupt vector presented by the VME interrupter may not be the vector specified in the `vba_post_irq` function.

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

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

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

3.7 Compaq AlphaServer GS60 and GS140 Systems

The following notes are specific to the Compaq AlphaServer GS60 and GS140 Systems.

3.7.1 CPU Boards Cannot Be Installed Until DIGITAL UNIX Is Installed

DIGITAL UNIX Version 4.0E must be installed before you upgrade your machine with the new AlphaServer 8200/8400 (Compaq AlphaServer GS60,GS140) 6/525 CPU boards.

3.7.2 Futurebus Backplane is Not Supported

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

Base System Software Notes 4

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

The following topics are discussed:

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

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

`http://www.unix.digital.com/faqs/publications/updates/V40E-update.html`

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

4.1 Commands and Utilities

The following notes apply to commands and utilities.

4.1.1 Escaped Comment Symbols in a Makefile

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

4.1.2 Security

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

4.1.2.1 Unexpected Command Behavior with ACLs

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

- ID mapping
- Mandatory access control
- Additional authentication requirements

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

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

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

```
# cp -p
```

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

4.1.2.2 ACLs and Archive Tools

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

4.1.3 Emacs Can Lose ACL File Settings

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

ACL, the new file will be protected only by the file permission bits.

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

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

4.1.4 Some Emacs Command Line Options Fail

DIGITAL 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 Executing Program Images Cannot Be Overwritten

The `write()` system call may fail with an ETXTBSY error when an attempt is made to overwrite a running program or shared library. This prevents the image in memory from being overwritten accidentally, which can result in application crashes or hangs.

For example, using the `/usr/bin/cp` command to copy into an executing program will fail with the message `Text file busy` when the `write` system call is invoked:

```
% a.out &
% cp foo a.out
cp: a.out: Text file busy
```

A workaround is to use the `/usr/bin/mv` command:

```
% mv -f foo a.out
```

You may also see this error in a development or compilation environment where the `make` utility is used to build executables.

4.1.6 New Tunable Parameter for Running Oracle8

If Oracle8 running on DIGITAL UNIX hangs, look for console message

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

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

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

where 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.

4.1.7 disklabel Command Supports ISO 9660 CD-ROM Image File Operations

The /sbin/disklabel command supports operations on an image file, including the capability of writing a disk label and boot blocks. The default format is ISO 9660 Rock Ridge standard CD-ROM format known as the CD-ROM File System (CDFS).

With this feature, users can create a bootable CD-ROM, in CDFS format, from the image file.

4.1.8 Change to format of ipcs command

The output format of the ipcs command has been modified. The KEY field now presents the value in hexadecimal format instead of decimal. This change is compatible with other UNIX implementations and conforms to the upcoming X/Open Unix98 standard.

4.1.9 New sendmail Version 8.8.8

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

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

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

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

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

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

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

4.1.10 Sendmail Warning Message

When the system is first booted, after a full installation, the following warning message is displayed as a result of starting `sendmail`:

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

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

4.1.11 Warning Message when sendmail Updates the alias Database

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

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

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

Because the permission on the `/var` directory is 775, `sendmail` gives a warning message because it expects the permission on `/var` to be set as follows:

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

Ignore this warning.

4.1.12 Status Message From ifconfig Command

The `ifconfig` command will sometimes show "trustgrp unknown" in the status line for an interface. This is mainly restricted to ATM lis and elan interfaces, but might be seen on any interface configured after the system has booted.

For example, if you have an ATM lis0 interface:

```
# ifconfig lis0
lis0: flags=808e1<UP,NOTRAILERS,RUNNING,NOARP,SIMPLEX>
inet 212.180.32.49 netmask ffffff00  ipmtu 9180  trustgrp unknown
```

The message can be ignored and has no effect on your system. It will be fixed in a future release of DIGITAL UNIX.

4.1.13 tar Command 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.

4.1.14 Netscape Communicator

The following notes apply to the Netscape communicator.

4.1.14.1 Netscape Communicator Dumps Core Running in CDE

Netscape Communicator V4.05 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, edit the `/usr/bin/X11/netscape` script, as follows:

1. Locate the following line (92):

   ```
   exec $REAL_NETSCAPE -name netscape $*
   ```

2. Edit the line as to read as follows:

   ```
   exec $REAL_NETSCAPE -name netscape -xrm '*nsMotifFSBCdeMode: True' $*
   ```

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`

4.1.14.2 Netscape Communicator Hangs Starting Java on Some Processors

The Netscape Communicator V4.05 process hangs the first time a page containing a Java applet is loaded by the Navigator browser when running on a DEC 3000 or AlphaStation 200. The Netscape process also hangs if you select the Java Console option in the Communicator pulldown menu. How

long the process hangs depends on the processor type.

This hang only occurs the first time you start Java. Therefore, leave the Netscape process running in background until the hung process clears itself. After the hang has cleared once, you can continue to use Netscape Communicator normally.

4.1.14.3 Netscape Communicator Does Not Include Localization

Unlike the Netscape Navigator software in previous versions of the DIGITAL UNIX Operating System, the Netscape Communicator software bundled with the current version of DIGITAL UNIX does not currently come with any localization. Therefore, you cannot bring up a Japanese interface of the Communicator, for instance, even when your process is running in a Japanese locale.

Localizations for Netscape Communicator are planned. When available, these localizations will be provided on the DIGITAL UNIX web site, and in future releases of the operating system.

4.2 SysMan System Management Graphical User Interface

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

4.2.1 Account Manager

The following notes apply to Account Manager, `dxaccounts`.

4.2.1.1 Copying Accounts and UID

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

4.2.1.2 Account Manager Restrictions

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

- Using mouse button 1 (MB1) to drag and drop users, groups, or templates does a copy operation, not a move operation. This is different from the

CDE/Motif default where MB1 performs a drag and drop move operation and Shift-MB1 is required to perform a copy operation. For example, using MB1 to drag a user from the Local Users view and dropping it in the NIS Users view creates a copy of that user in NIS.

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

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

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

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

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

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

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

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

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

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

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

  Workarounds: Account manager correctly allows two or more system administrators to work on the same password files simultaneously. The proper file locking will occur and new accounts can be added or modified. However, the local groups file, `/etc/group`, and the NIS groups file, `/var/yp/src/group`, are written out after each group modification. Therefore, the last system administrator to make a change in a group's view window would overwrite any prior changes from a different system

administrator. For this reason, running multiple, concurrent Account Manager instances is not recommended.

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

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

  Workaround: None. You can safely ignore these messages.

4.2.1.3 Account Manager Problems

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

4.2.1.4 Enhanced Security Account Manager Problems

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

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

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

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

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

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

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

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

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

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

- Error messages generated from the Create/Modify Template dialog box refer to the user name when they should refer to the template name.
- On enhanced security system, you typically retire users instead of deleting them. However, there are times when you might want to delete a user account. Account Manager supports retiring user accounts but not deleting them.

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

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

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

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

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

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

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

  Workaround: Restart Account Manager to restore the former template icon. Use the Delete Toolbar icon or the Edit->Delete... option from the Template view to delete the undesired template
- Accounts and templates inherit their settings either from locally defined values in their protected password database entry or from the templates that they may reference. All accounts and templates implicitly reference a default template. The default template is not served by NIS. This creates an inconsistency for Account Manager when displaying NIS user accounts and templates on an NIS master. The user and template values displayed may be the default template values of the NIS master. When an NIS user logs into a NIS client, the NIS client's default template might be different from the NIS master's default template. The client's default template is used to establish the user's account settings.
- Using drag and drop to copy a user account copies the user's template references by value when the user account is being dropped on a different view. This means that the template itself will no longer be referenced by the newly created account. Instead, the template's values will be contained directly in the new user's protected password database entry.

For example, assume the local user Joe has an account based on the developers template. If you drag and drop Joe's account from the Developers view into the NIS Users view, the attributes from the developers template will be placed into the protected password database entry for Joe's account. This preserves Joe's developer attributes and overrides any corresponding attributes from the default template for NIS users.

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

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

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

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

  Workaround: None.

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

  Workaround: Restart Account Manager and then delete the template.

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

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

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

4.2.2 Print Configuration Manager

The Print Configuration Manager may have some problems with `/etc/printcap` files from DEC OSF/1 Version 3.2 or earlier, as follows:

- Aliases that conflict with system-assigned names

  Using `/etc/printcap` files in the current version of DIGITAL UNIX,

the system assigns printer names `lp[0-9]*`, `[0-9]*`, and for the default printer, `lp`. For example, the default printer may have a name field such as `lp0|0|lp|default|declaser3500:...`. Another printer may be named `lp7|7|some_alias|another alias:...`. Therefore, the system has difficulty with printers that have less than two names or that use these reserved names as aliases.

- Altered attribute validation

  Some of the attribute value checking is different between earlier versions and the current version. For example, some fields that were not required now are, and some attribute values that were legal no longer are.

- Trailing comments

  The Print Configuration Manager requires that all comments be associated with a printer. As a result, comments appearing after the last printer are truncated.

To avoid these problems, invoke the `printconfig` utility with the menu interface (`printconfig -ui menu`). This brings up the `lprsetup` utility, which is fully compatible with earlier `printcap` files.

4.2.3 BIND Configuration

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

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

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

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

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

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

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

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

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

4.2.4 The dxshutdown Application Does Not Prohibit Logins

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

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

4.3 System Administration

The following notes apply to system administration.

4.3.1 Change in mkpasswd Creates Potential Password Database Incompatibility

When the `/etc/passwd` file is very large, a performance degradation may occur.

When the number of `passwd` entries reaches the 30,000 to 80,000 range or greater, `mkpasswd` will sometimes fail to create a hashed (`ndbm`) database. Because the purpose of this database is to allow for efficient (fast) searches for `passwd` file information, failure to build it causes commands that rely on it to do a linear search of `/etc/passwd`. This results in a serious performance degradation for those commands.

For customers choosing to use the `mkpasswd -s` option to avoid this type of failure, a potential database or binary compatibility problem may arise. If a customer application that accesses the password database created by `mkpasswd` is built statically (nonshared), that application will be unable to read from or write to the password database correctly. This would cause the customer application to fail either by generating incorrect results or by possibly dumping core. Any statically linked application would be affected if it directly or indirectly calls any of the `libc ndbm` routines documented in the `ndbm`(3) reference page and then accesses the password database. To remedy this situation, you must re-link the application.

Customers who do not use the `mkpasswd -s` option will not see this compatibility problem.

4.3.2 License Management Facility Allows Year 2000 Testing

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

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

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

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

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

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

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

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

4.3.3 Disabling Compressed Crash Dumps

Compressed crash dumps have been enabled for this release of DIGITAL UNIX.

You may need to disable this feature if you have tools or scripts that do not work with compressed crash dumps. If necessary, use `dbx` to set the

compressed_dump variable to 0 in the running kernel, as follows:

```
(dbx) assign compressed_dump = 0
```

Note that this must be repeated each time the kernel is booted. Alternatively, you can use dbx to patch the value of compressed_dump to 0 in the kernel image file.

Chapter 4 of the *Kernel Debugging* guide provides more information about crash dump settings. Also, see the savecore(8), sysconfig(8), and dbx(1) reference pages.

4.3.4 Log Files in /var/adm/syslog.dated

The /var/adm/syslog.dated directory contains preserved copies of log files that are used for debugging. Normally, these files do not contain many entries. However, under certain error conditions, a DIGITAL UNIX subsystem might log an excessive amount of entries to a file and cause a problem.

You should either physically check the logs on a regular basis or use the cron utility to set up a regular job to clear the log files. The default root crontab file in the /var/spool/cron/crontabs directory contains the following sample line for clearing up the /var/adm/syslog.dated directory (the \ indicates line continuation):

```
40 4 * * * find /var/adm/syslog.dated -depth -type d -ctime +5 \
-exec rm -rf {} ;
```

If enabled, this cron job will be activated every morning at 4:40 a.m. and will delete any log file in /var/adm/syslog.dated that has not been updated for the last five days. You can edit the crontab file to uncomment and modify this line or add a similar line by using the following command:

```
# crontab -e
```

For more information, see the crontab(8) reference page.

4.3.5 Secure version of syslog

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

If you are installing the secure version of syslogd on a system, and you have configured or intend to configure other hosts to forward syslog messages to the system, complete the following steps:

1. Use su to become the superuser (root)
2. Create the file /etc/syslog.auth using a text editor. This file must be owned by root and have permissions set to 0600.

3. Add the names of any remote hosts that are allowed to forward `syslog` messages to the local system. Host names must meet the following criteria:
   - Each remote host name should appear in a separate line in `/etc/syslog.auth`. (A line started with the "#" character is considered as a comment and is ignored.)
   - A host name must be a complete domain name, for example:
     ```
     trout.fin.huk.com
     ```
   - If a domain host name is given, it must either appear in the local `/etc/hosts` file or the local system must resolve it through a name server (such as BIND).
   - A host name can have at most as many characters as defined by the `MAXHOSTNAMELEN` constant in `<sys/param.h>`, although each line in the `/etc/syslog.auth` file is limited to 512 characters.

4.3.6 Fixing Truncated Log Messages in /var/adm/messages

System configurations that are large, containing many adapters and devices, may exhibit incomplete message logging in the `/var/adm/messages` file.

If this happens, you should compensate for the large system configuration by increasing the value of the `msgbuf_size` attribute in the `generic` subsystem using `sysconfigdb` utility or the `dxkerneltuner` interface. The default value for `msgbuf_size` is 4096. Usually, setting it to 8192 is sufficient to resolve the problem. If you have a smaller configuration and you do not see this problem, you should not make the change.

Refer to the `sysconfigdb`(8), reference page and the *System Configuration and Tuning* guide for information about modifying system attributes.

4.3.7 EISA Configuration Utility Revision Requirements

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

4.3.8 Open3D Support

Consult the Open3D Software Product Description (SPD) before installing Open3D to ensure that this DIGITAL 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.

4.3.9 Bootable Tape

For this release, bootable tape will not work with the LSM product. Not all platforms and tape drives support bootable tape. The following processor platforms are supported:

- DEC 3000-500
- DEC 3000-400
- DEC 3000-600
- DEC 3000-300
- DEC 3000-300X
- DEC 3000-900
- DEC 2100
- AlphaStation 600
- AlphaStation 200
- AlphaServer 1000A
- AlphaServer 2100
- AlphaServer 4100

The following tape devices are supported:

- TLZ06, 4mm, 2.0GB/4.0GB
- TLZ07, 4mm, 4-8GB
- TZK10, QIC tape, 320MB-525 MB
- TZK11, QIC tape, 2.0GB
- TZ86, 5-1/4-inch cartridge

4.3.9.1 Disk Space Issues

To use the `btcreate` utility, your system must have at least 156,000 512-byte blocks of free space in the `/usr` directory.

You will not have enough space if your system uses an RZ26 or smaller disk with the default partitions and you have installed all of the subsets and kernel options.

To overcome this limitation, you can reclaim the required space by removing some subsets or by creating and mounting new partitions.

The following steps show you how to create and mount new partitions for a UNIX file system (UFS). If you prefer to use AdvFS, use the `mkfdmn` and `mkfset` commands.

1. Run the newfs command to recreate a new partition:

   # **newfs /dev/rz1d**

2. Change the current working directory to the /usr/sys directory:

   # **cd /usr/sys**

3. Make a *SYSTEM*.BOOTABLE directory under the /usr/sys directory, where *SYSTEM* is the system name:

   # **mkdir FLAWLESS.BOOTABLE**

4. Mount the new partition on the *SYSTEM*.BOOTABLE directory:

   # **mount /dev/rz1d /usr/sys/FLAWLESS.BOOTABLE**

   This device should have at least 75,000 512-blocks available.

5. Create another new partition:

   # **newfs /dev/rz1b**

6. Mount the partition:

   # **mount /dev/rz1b /mnt**

7. Change the current working directory to the /usr/sys/bin directory.
8. Copy the contents of the /usr/sys/bin directory to the /mnt directory:

   # **cp * /mnt**

9. Unmount the /mnt directory:

   # **umount /mnt**

10. Mount the new partition on the /usr/sys/bin directory:

    # **mount /dev/rz1b /usr/sys/bin**

After completing these steps, your system should have the necessary space to run btcreate.

If you are using AdvFS, the /usr/sys/bin file system must be copied during btcreate in order to copy the entire contents of the /usr file system.

4.3.9.2 Tape Drive Restriction

Ensure that the kernel has been built with the tape drive connected to your system. If the drive was not connected when the kernel was built, you will see dump errors and the system will not be able to boot from the tape drive.

4.3.9.3 Bootable Tape Notes for 32 MB Systems

Bootable tape will not function with the `-m mfs` option on systems with 32 MB memory configurations. After booting the kernel from tape, commands that use shared libraries will core dump.

Use the `-m ufs` option while creating the tape on systems with 32 MB memory configurations.

4.3.9.4 Bootable Tape Does Not Support Generic Kernels

Bootable tape does not support the bootable kernel built with the `/usr/sys/conf/GENERIC` kernel configuration file.

Be sure to use a system-specific custom kernel.

4.3.9.5 Cross-Platform Bootable Tapes Are Not Supported

Using a bootable tape on a platform other than the one on which it was created is not supported. For example, you cannot create a tape on a 4100 system and boot from it on a 1000A system.

4.3.9.6 QIC Tape Drives

When using QIC tape drives to create bootable tapes, you must use only high-density tapes of 320 or more megabytes. The QIC-24, QIC-120, and QIC-150 format tapes of fixed-512 blocks will not work. Tapes with a variable block size, such as the QIC-320 and QIC-525, will work with bootable tape.

Using an improperly configured QIC tape drive to create a bootable tape will result in an I/O error, a write error, or permission denied error. Therefore, you must take one of the following actions:

- Configure the drive at installation time.
- Rebuild the kernel if the drive was attached to the system after the installation.

4.3.9.7 Multiple Tapes

If creating a bootable tape with a UFS file system extends to multiple tapes, the `/sbin/dump` command displays a message indicating that the tape must be changed. If the tape is not changed promptly, warning messages repeat periodically until the tape is changed.

When you change the tape, the warning messages will stop.

4.3.9.8 Disk Overhead Requirements

When selecting disk partitions while restoring file systems from tape, add 5 percent to the needed file size displayed on the console.

4.3.9.9 Write Protection for Bootable Tape

A QIC tape created with the `btcreate` utility may fail with the following error when booted:

```
failed to send Read to mka...
```

Be sure that the tape is write protected before booting.

4.3.10 Opening a Write-Protected Tape in Write Mode

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

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

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

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

4.3.11 Restrictions for a Dataless Environment

The DIGITAL UNIX Server Extensions includes 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 f or each client. The `/usr` file system is exported read-only and is shared by all clients registered to the environment. Each client has their own `/var` file system. Dataless clients access the file systems maintained on the server utilizing NFS.

A dataless environment should be considered in the following scenarios:

- There is a need to reduce the disk space requirements on the client systems. Local disk space is required only for swap and dump, the `root`, `usr` and `var` file systems are resident on the server.
- There is a need to have all software and data centrally located on the server. A dataless client is not designed to support local access to data.

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

- Hardware limitation of the network – the majority of file system usage performed by any of the clients is done via NFS. This will consume network bandwidth, which, depending on the type and configuration of the network, can become exhausted as more clients are added to the environment.
- Hardware limitation on the file systems contained on the server – all clients are configured to share a `/usr` file system and all client specific data is contained in a `/clients` file system on the server. As clients are added to the dataless environment, the demands on these file systems on the server increase and may surpass the capabilities of the server.
- A dataless environment should not be considered to address central system management requirements.
- 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.

4.3.12 Enhanced Security

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

4.3.12.1 Distribution of Enhanced Security Profiles via NIS

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

- Successful and unsuccessful login attempts for NIS-shared accounts require the completion of the following steps:
  1. The master system's `rpc.yppasswdd` daemon must respond and update the last successful and last unsuccessful login fields in the `prpasswd` NIS map.
  2. The NIS slave servers must answer to the `yppush` operation initiated from the `rpc.yppasswdd` daemon. (Most successful logins do not require a `yppush` operation, but login failures and password changes do.)

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

  The more NIS slave servers that are present in a given NIS domain, the more time `rpc.yppasswdd` takes to complete these steps. Also, nearly-simultaneous login attempts are processed sequentially by the NIS master, each waiting on a possible `yppush` for the previous attempt to

succeed. Therefore, if several simultaneous attempts arrive at once, some may timeout and require you to log in again. You can alleviate this problem to some extent by using the `-p` option of `yppush`. One way to do this is to modify the `/var/yp/Makefile` file and change the `YPPUSH=` line. The following example allows up to 6 simultaneous transfers to NIS slave servers (the default number is 4):

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

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

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

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

  With successful logins, the `rpc.yppasswdd` daemon will defer pushing the maps if the login notification comes from a Version 4.0D client. Therefore, the `yppush` operation is only completed when an older client initiates the operation or when it is necessary to clear a failed login count.

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

In previous releases of DIGITAL UNIX, 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.

4.3.12.2 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.

4.3.12.3 pw_id_map and gr_id_map Files Are Obsolete

The Enhanced Security routines pw_idtoname, pw_nametoid, gr_idtoname, and gr_nametoid (described in pw_mapping(3)) previously used the /etc/auth/system/pw_id_map and /etc/auth/system/gr_id_map files to find the required information for mapping names to numeric identifiers, and vice versa. The disk space required by those files imposed a limit on how many accounts a system could support.

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

4.3.12.4 Restriction to 4000 NIS Accounts Removed

Logins with NIS-shared extended user profiles under Enhanced Security have been streamlined, thus lifting the former restriction of 4,000 accounts. However, depending on the method chosen for building the NIS maps (using nissetup or the /var/yp/Makefile file), the limits of the ndbm storage format may still impose a limit on the number of accounts that can be shared through NIS. If you are sharing more than 10,000 accounts with NIS, DIGITAL recommends that you use the btree storage format instead of ndbm (where practical). The limitations on NIS slave servers and NIS master availability for use of the prpasswd NIS map are unchanged.

4.3.12.5 Behavior of useradd, usermod, and userdel Commands

The useradd command correctly honors the default administrative lock value found in the /.sysman/Account_defaults file. If Account_defaults does not exist, the internal default for useradd is

to create locked accounts. You can use the `administrative_lock_applied` extended command line option to override the default. In the following example, `useradd` creates a locked account for `foo` regardless of the default value for administrative lock:

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

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

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

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

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

4.3.13 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 variable `enable_extended_uids` to. `1` (enabled).

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

- While `enable_extended_uids` can be set dynamically, you cannot disable it dynamically. To disable `enable_extended_uids`, set the boot time value to `0` (disabled) and reboot your system.
- If you cannot easily reboot your system, note that enabling extended UIDs has very little impact on overall system performance.

4.4 Network and Communications

The following notes apply to network and communications software.

4.4.1 Using netconfig

When using `netconfig` while CDE is running, avoid restarting network services after reconfiguring the primary network interface. This action can result in error dialog boxes and may even cause CDE to hang. The problems

may not be observed until you use `bindconfig` to set up BIND.

In particular, do not use the following `netconfig` features while running a CDE session:

- Do not answer `yes` to restart the network services from the `netsetup` or `netconfig` menu interfaces.
- Do not answer `yes` to the following prompt after reconfiguring the primary network interface from the `netconfig` graphical interface:

  ```
  restarting network services
  ```

- Do not use the Network pull-down menu in the `netconfig` graphical interface to start, stop, or restart the network.
- Do not exit from `netconfig` then use the `/usr/sbin/rcinet` stop, start, or restart options from the command line.

For the configuration changes to take effect, you must use `/sbin/reboot` or `/sbin/shutdown -r now` to reboot your machine from the command line.

4.4.2 IP Switching over ATM

The following restrictions apply when using IP switching over ATM:

- At most, one IP switching interface (`ips`) per host is supported.
- Other ATM protocols may not be used on an ATM driver that is being used for IP switching. Do not use the `atmsig` command to start UNI signaling on a driver used for IP switching.
- On an `ips` interface, `tcpdump` and `packetfilter` are not supported.
- 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.

4.4.3 Orderly Release in XTI

This release does not support Orderly Release in XPG4 XTI (default XTI interface). It is still available for users of XPG3 XTI. See the *Network Programmer's Guide* for information on using XPG3 XTI.

4.4.4 Incorrect Error Message When Using netsetup and rcinet

When you use `netsetup` to restart the network, an error message similar to

the following will be displayed:

```
kill: 204: no such process
```

This problem also appears when you execute the following commands:

```
# rcinet stop
# rcinet restart
```

The message is incorrect and has no effect on your system.

4.4.5 Edit the /etc/hosts File After Running netsetup

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

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

Example:

```
16.60.266.6    example.domain.com    example
```

where:

16.60.266.6 = ip address of your system.

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

example = aliases for your system.

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

4.4.6 CDE's Static Dependency on the Network

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

Prior to making any dynamic network changes, such as changing the state of your network adapter to off or changing your primary network address, add

the following entry to the `/.dtprofile` file:

```
export DTNONETWORK=true
```

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

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

4.4.7 Autosense Removed from Tulip Ethernet Driver

The autosense feature has been removed from the Tulip Ethernet and Fast Ethernet driver. This feature automatically determined whether your Ethernet connection was 10BaseT (UTP, Twisted Pair), 10Base2 (BNC, Thinwire), or 10Base5 (AUI, Thickwire) during the boot sequence. It also attempted to select between 10 Mbps and 100 Mbps operation if applicable, but not between half-duplex and full-duplex mode.

The Tulip driver used autosense in those systems where the Alpha SRM Console did not support or communicate (to the driver) the setting of the `EW*0_MODE` environment variable. Autosense was also used as the default mode for the EISA DE425 adapter.

In both cases, the default is now Twisted-Pair (half-duplex, 10 Mbps). If this new default is acceptable, then you do not need to do anything. Otherwise, you need to take one or more of the following actions:

- For the EISA DE425 adapter, run the EISA Configuration Utility (ECU) to select the desired media.
- For all other options, upgrade your Alpha SRM Console firmware to the most recent version available for your system. Then set the `EW*0_MODE` console environment variable for each `tu` or EW interface as desired.
- For all options (including the EISA DE425), use the `lan_config` command to select the desired mode of operation. This command overrides whatever was selected via the ECU or console (`EW*0_MODE` setting). You may use the `/etc/inet.local` configuration file to maintain `lan_config` settings across system restarts.

Refer to the `tu`(7), `lan_config`(8), and `inet.local`(8) reference pages for more information.

Note that the autosense feature is different from autonegotiation. The autosense feature uses a software algorithm to determine what media is currently present on the given device, and the autonegotiation feature uses specific hardware for determining the speed (10/100) and mode (full duplex/half duplex). The autonegotiation feature is still available in the Tulip driver and there are no plans to retire it.

4.4.8 Support for NetRAIN Network Interface

DIGITAL 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 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)

DECsafe ASE users should consult the DECsafe ASE software product description (SPD) for a list of network interfaces supported by that product.

When using NetRAIN over LANE it is recommended that you use UNI Version 3.1. With some ATM switches it is necessary to use UNI Version 3.1 to obtain acceptable failover times. This includes the Gigaswitch. If you use UNI Version 3.0, the failure over time may be long because of the T309 timer is default set to 90 seconds in some switches. If the T309 timer is adjustable on the switch, you may try to set the T309 timer to 10 seconds like UNI Version 3.1 for acceptable failover times.

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

4.4.8.1 LAT and NetRAIN Cannot Run Together

Local Area Transport (LAT) cannot be started on a system where a NetRAIN set has been configured.

4.4.8.2 NetRAIN and Layered Product Licensing

Licensing schemes that use a network adapter's Media Access Control (MAC) address to uniquely identify a machine can be affected by how NetRAIN changes the MAC address. All network drivers support the `SIOCRPHYSADDR` ioctl that fetches MAC addresses from the interface. This ioctl returns two addresses in an array:

- Default hardware address – this is the permanent address that is taken from the small PROM that each LAN adapter contains.
- Current physical address – this is the address that the network interface responds to on the wire.

Licensing schemes based on MAC addresses should use the default hardware address returned by the `SIOCRPHYSADDR` ioctl; do not use the current physical address as NetRAIN modifies this address for its own use. See the reference page for your network adapter (for example, `ln`(7) and `tu`(7)) for a sample program that uses the `SIOCRPHYSADDR` ioctl.

4.4.9 Sender Address Problem When Using DECnet Mail

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

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

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

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

Then obtain the latest version of mail11v3 and replace the current version in the `/usr/sbin` directory. The new mail11v3 will understand the above option lines and use the appropriate sender id.

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 Kernel Module Is Dynamically Loadable

You no longer need to build LAT into the kernel. LAT is not made a mandatory kernel option upon selecting the LAT subset and does not appear in the kernel configuration file. As LAT requires the Data Link Bridge (DLB), you must still build DLB into the kernel when using LAT.

The default behavior upon booting to multiuser mode is for LAT to be dynamically loaded into the running kernel. If LAT is not started at boot-time via the `/sbin/rc3.d/S58lat` script, the recommended method for starting and stopping LAT is to verify that `LATSETUP` is enabled in `/etc/rc.config` and execute the `/sbin/init.d/lat` program, using the `start` or `stop` options.

4.6 File Systems

The notes in this section apply to file systems.

4.6.1 UNIX File System Restriction

When using the UNIX file system (UFS), there is a problem when setting properties. Setting a property on a FAST symbolic link, a block special file, or a character special file causes `fsck` to erroneously detect contradictory block counts and produce inconsistent file system activity. There is no solution for this problem and it will be fixed in a future release.

4.6.2 The newfs Command No Longer Searches disktab File

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

4.6.3 Using ACLs over NFS

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

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

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

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

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

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

- Add the option to the mount command as follows:

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

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

4.6.4 ACL Size Limitations

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

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

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

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

4.6.5 Advanced File System (AdvFS)

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

For information about recovering from AdvFS domain panics and correcting an overlapping `frag` data corruption problem, see Appendix F.

4.6.5.1 AdvFS and fsync()

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

4.6.5.2 Reusing AdvFS Partitions

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

4.6.5.3 Disk Usage Information

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

4.6.5.4 AdvFS Split Mirror Dual Mount May Not Function Correctly

AdvFS ordinarily does not allow a domain to be mounted if another domain is already mounted with the same domain ID. However, in some cases, such as a split mirror dual-mount, you may want to bypass this restiction. You can do this by using the `-o` dual option to the `mount` command. When this option is used, AdvFS assigns a new domain ID to the domain being mounted.

4.6.5.5 Disabling Flushing to Disk AdvFS File Read Access Times

When a `read()` system call is made to a fileset's files, the default behavior is for AdvFS to update both the in-memory file access time and the on-disk stat structure, which contains most of the file information that is returned by the `stat()` system call.

You can improve AdvFS performance for proxy servers by specifying at mount time that AdvFS update only the in-memory file access time when a `read()` system call is made to a file. AdvFS will update the on-disk stat structure only if the file is modified.

To enable this feature, use the `mount` command `noatimes` option. See the `read`(2) and `mount`(8) reference pages for more information.

Updating only the in-memory file access time for reads can improve proxy server response time by decreasing the number of disk I/O operations. However, this behavior jeopardizes the integrity of read access time updates and violates POSIX standards. Do not use this feature if it will affect utilities that use read access times to perform tasks, such as migrating files to different devices.

4.6.5.6 Improving AdvFS Performance on Systems that Reuse Many Files

If your system opens and then reuses many files (for example, if you have a proxy server), you may be able to improve AdvFS performance by increasing the number of AdvFS access structures that the system places on the access structure free list at startup time.

AdvFS access structures are in-memory data structures that AdvFS uses to cache low-level information about files that are currently open and files that were opened but are now closed. Increasing the number of access structures

on the free list allows more open file information (metadata) to remain in the cache, which can improve AdvFS performance if the files are reused.

However, allocating more AdvFS access structures on the free list will decrease the memory that is available to processes and to the UBC.

Use the `AdvfsPreallocAccess` attribute to modify the number of AdvFS access structures that the system allocates at startup time. The default and minimum values are 128. The maximum value is either 65536 or the value of the `AdvfsAccessMaxPercent` attribute, whichever is the smallest value. The `AdvfsAccessMaxPercent` attribute specifies the maximum percentage of the malloc pool (pageable memory) that can be used for AdvFS access structures.

4.6.5.7 AdvFS Now Supports Atomic Write Data Logging

By using the `/usr/sbin/chfile` command, you can now activate atomic write data logging for individual files in an AdvFS domain. Atomic write data logging is a mode of performing writes to a file that differs from the normal, asynchronous mode that is typical of UNIX file systems, and from the forced synchronous write mode available through the `-l` switch to the `/usr/sbin/chfile` command.

When atomic write data logging is activated, writes to a file are done asynchronously. They are, however, also written to the AdvFS log file. This has the effect that if the system crashes any time during or after a write system call and the contents of the file are examined upon reboot, only entire write requests will be present in the file.

For example, if a write of a 8192-byte buffer were done to the file and either during the write system call or shortly thereafter the system crashed, upon reboot either the entire 8192 bytes of data would be found in the file or none of it would be found there. There is no possibility that only 1024, or 2048 bytes of the write, for example, would be in the file.

You can also activate and deactivate this feature using the `fcntl()` system call. In addition, both `/usr/sbin/chfile` and `fcntl()` can be used on an NFS client to activate or deactivate this feature on a file that resides on the NFS server.

Note that because atomic write data logging causes user data to be written to both the user's file and the AdvFS log file, writes to files with Atomic Write Data Logging enabled will be slower than normal, asynchronous writes.

Refer to the `fcntl`(2) and `chfile`(8) references pages for more information.

4.6.5.8 Preventing Partial AdvFS Data Writes

AdvFS writes data to disk in 8 KB chunks. By default and in accordance with POSIX standards, AdvFS does not guarantee that all or part of the data will actually be written to disk if a crash occurs during or immediately after the write. For example, if the system crashes during a write that consists of two 8 KB chunks of data, only a portion (anywhere from 0 to 16 KB) of the total write may have succeeded. This can result in partial data writes and inconsistent data.

To prevent partial writes if a system crash occurs, you can use the `chfile -L on` command to enable atomic write data logging for a specified file.

By default, each file domain has a transaction log file that tracks fileset activity and ensures that AdvFS can maintain a consistent view of the file system metadata if a crash occurs. If you enable atomic write data logging on a file, data from a write call will be written to the transaction log file before it is written to disk. If a system crash occurs during or immediately after the write call upon recovery, the data in the log file can be used to reconstruct the write. This guarantees that each 8 KB chunk of a write is either completely written to disk or is not written to disk.

For example, if atomic write data logging is enabled and a crash occurs during a write that consists of two 8 KB chunks of data, the write can have three possible states: none of the data is written, 8 KB of the data is written, or 16 KB of data is written.

Atomic write data logging may degrade AdvFS write performance because of the extra write to the transaction log file. In addition, a file that has atomic write data logging enabled cannot be memory mapped by using the `mmap()` system call.

A file cannot have both forced synchronous writes enabled and atomic write data logging enabled. However, you can enable atomic write data logging on a file and also open the file with an `O_SYNC` flag. This ensures that the write is synchronous, but also prevents partial writes if a crash occurs.

Use the `chfile` command with no flags in order to determine if forced synchronous writes or atomic write data logging is enabled. Use the `chfile -L off` command to disable atomic write data logging (the default).

To enable atomic write data logging on AdvFS files that are NFS mounted, the NFS property list daemon, `proplistd` , must be running on the NFS client and the fileset must be mounted on the client by using the mount command's `proplist` option.

If atomic write data logging is enabled and you are writing to a file that has been NFS mounted, the offset into the file must be on an 8 KB page boundary, because NFS performs I/O on 8 KB page boundaries.

4.6.5.9 Memory Mapping Will Fail on AdvFS Files Using Data Logging

Starting with DIGITAL UNIX Version 4.0E, attempts to memory map an AdvFS file using the `mmap()` system call will fail if the file has had atomic write data logging activated on it. Use the `chfile` command to determine if a file is using atomic write data logging. If the `chfile` command displays the following, attempts to memory-map the file using the `mmap()` system call will fail:

```
I/O mode = atomic write data logging
```

To deactivate the atomic write data logging, enter the following command:

```
# chfile -L off filename
```

For more information on atomic write data logging, see the `chfile`(8) references page.

4.6.5.10 Cleanly Unmount File Systems Before Changing Operating System Versions

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

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

Therefore, if you want to reboot using a different version of DIGITAL UNIX, 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 DIGITAL UNIX 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 DIGITAL UNIX and remount the filesets.

4.6.5.11 AdvFS Verify Reports an Error on Some Domains

The AdvFS tool `verify` will report the following error when it is run on a Version 4.0 domain which was created with the `mkfdmn` command using the

`-p` flag:

```
Checking mcell list ...
set_mcell_position: position field already set
disk: 1, mcell id (page.cell): 0.6
set tag: -2.0 (0xfffffffe.0x00000000)
tag: -6.0 (0xfffffffa.0x00000000)
```

This error message can be ignored.

4.6.6 Logical Storage Manager

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

4.6.6.1 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. DIGITAL recommends that you do not configure volumes with sparse plexes.

4.6.6.2 Restrictions with LSM root and swap Volumes

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

- Root and primary swap must be on the same physical disk and you must configure them using the encapsulation tools.
- If a secondary swap volume is configured under LSM, you must configure the root and primary swap.
- You must configure the root, primary swap, and any secondary swap volumes into the disk group named `rootdg`.

4.6.6.3 Possible Problems Accessing Physical Block 0 with LSM

Physical block 0 on DIGITAL 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 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.

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

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

4.6.6.5 Enabling LSM After Installation Requires Rebuilding the Kernel

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

To rebuild the kernel, run the `doconfig` utility with no command flags. Note that the `doconfig` menu display does not include LSM. However, the `doconfig` utility will build a kernel that includes LSM. Refer to the *Logical Storage Manager* guide for more information.

4.6.6.6 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.

4.6.6.7 Using LSM rootvol Requires sysconfigtab Parameters

If you use the LSM `rootvol` volume for the root file system and the `swapvol` volume is in use as a primary swap volume, LSM adds the

following entries to the /etc/sysconfigtab file to enable rootability:

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

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

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

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

4.6.7 Mounting Associated Products CD-ROMS on DIGITAL UNIX Versions Other Than 4.0E

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

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

On releases prior to Version 4.0E, the APCDs must be mounted with additional options, as follows:

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

On versions of DIGITAL UNIX prior to 4.0D you may receive the following error indicating that CDFS support is not built into the current running kernel:

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

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

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

4.6.8 Modifying the Maximum Number of Open File Descriptors

The proc subsystem's open-max-soft and open-max-hard attributes control the maximum number of open file descriptors for each process. When the open-max-soft limit is reached, a warning message is issued, and when the open-max-hard limit is reached, the process is stopped. These attributes prevent runaway allocations (for example, allocations within a loop that cannot be exited because of an error condition) from consuming

all the available file descriptors.

The `open-max-soft` and `open-max_hard` attributes both have default values of 4096 file descriptors (open files) per process, which is the maximum, systemwide value.

If an application requires many open files, you can increase the open file descriptor limit for that application above the maximum that is supported by the `proc` subsystem (4096). Increasing the maximum limit up to 65,536 provides more file descriptors to the process, but it increases the possibility of runaway allocations. In addition, if you increase the number of open files for a process, make sure that the `max_vnodes` attribute is set to an adequate value. See Appendix E for information about increasing the open file descriptor limit for an application.

Decreasing the open file descriptor limit decreases the number of file descriptors available to each process and prevents a process from consuming all the file descriptors. However,decreasing the limit may adversely affect the performance of processes that require many file descriptors.

4.6.9 New Prestoserve Subsystem Attribute

The Prestoserve subsystem attribute `presto-buffer-hash-size` controls the size of the Prestoserve buffer cache. The minimum value is 0, and the maximum value is 64 KB. The default value is 256 bytes. Under certain circumstances, Netscape Navigator may crash upon invocation when the current locale is `ja_JP.deckanji`. If this happens, a workaround for the problem is to add the following four lines to the `/usr/i18n/lib/X11/ja_JP.deckanji/app–defaults/Netscape` file:

```
netscape.xnlLanguage:   ja_JP.eucJP
netscape.XnlLanguage:   ja_JP.eucJP
Netscape.xnlLanguage:   ja_JP.eucJP
Netscape.XnlLanguage:   ja_JP.eucJP
```

This will force Navigator to run in the `ja_JP.eucJP` locale to avoid the crash.

Development Environment Notes 5

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

- Realtime programming
- DECthreads (pthreads)
- Analysis Tool with Object Modification (ATOM)
- Java Programming

5.1 Realtime Programming

The following notes apply to realtime programming.

5.1.1 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`.

5.1.2 POSIX 1003.1b Synchronized I/O and File Truncation

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

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

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

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

5.2 DECthreads (pthreads)

The following notes apply to DECthreads. See Section 8.10 and Section 8.11 for information about DECthreads interfaces that will be retired in a future release.

5.2.1 Static Libraries

Users who desire optimal performance from DECthreads, and who are willing to relink on future versions of DIGITAL UNIX, may want to use the DECthreads static libraries that are located in the `CMPDEVENH425` subset. Once this subset is installed, you can find the libraries in the `/usr/opt/alt/usr/lib/threads` directory.

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

5.2.2 Forking

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

5.2.3 Signal Handling

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

5.2.4 Scheduling Behavior (Contention Scope)

In releases of DIGITAL UNIX prior to Version 4.0, thread scheduling attributes were systemwide. In other words, threads had a system contention scope. In Version 4.0 and higher, 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 DIGITAL UNIX Version 4.0D.

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 DIGITAL 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 DIGITAL UNIX 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 may be impossible to interrupt such a thread.

5.2.5 Known DECthreads Problems

The following notes discuss problems with the current implementation of DECthreads.

5.2.5.1 Freeing User-Allocated Stacks

If you choose to use the new `stackaddr` thread creation attribute that allows you to allocate your own stack for a thread, beware that you cannot deallocate the stack immediately when a thread returns from `pthread_join` called with the thread ID of the thread using the stack. That is, in the following code segment, it is not safe to return the memory-

mapped stack after `pthread_join` returns:

```
pthread_t       thread;
pthread_attr_t  attr;
void            *stack;

stack = mmap (NULL, PTHREAD_STACK_MIN, PROT_WRITE|PROT_READ,
MAP_ANONYMOUS|MAP_PRIVATE, -1, 0);
pthread_attr_create (&attr);
pthread_attr_setstackaddr (&attr, stack);
pthread_create (&thread, &attr, routine, NULL);
pthread_join (thread, NULL);
/* The following call may cause an access fault in "thread",
because it may still be using the stack to complete
termination. */
munmap (stack, PTHREAD_STACK_MIN);
```

The problem is that, in technical contradiction to the POSIX standard, `pthread_join` may return before the target thread has completely terminated. This was not a problem before the implementation of the `stackaddr` attribute, because although the thread might have existed, it could not have accessed any program-visible resources. Now, because it may still be executing on its stack, the program cannot safely free that stack. There is currently no reliable way to determine when it is possible to free the stack.

This problem will be fixed in a patch to Version 4.0D and in future releases.

DIGITAL does not recommend using the `stackaddr` attribute. The semantics of this attribute are poorly defined by POSIX and the Single UNIX Specification, Version 2, and as a result, code using the attribute is unlikely to be portable between implementations. The attribute is difficult to use reliably, as the developer must, by intimate knowledge of the machine architecture and implementation, know the correct address to specify relative to the allocated stack. The implementation cannot diagnose an incorrect value because the interface does not provide sufficient information. Using an incorrect value may result in program failure, possibly in obscure ways.

5.3 Analysis Tool with Object Modification

The following notes apply to the Analysis Tool with Object Modification (ATOM) utility.

5.3.1 Using pixie on Applications Built with -om

If your application was built with the `-om` switch, you must use the previous version of `pixie` for performance analysis work on the application. This version is located in `/usr/opt/obsolete/usr/bin/pixie`.

The ATOM-based tools, including the latest version of `pixie`, cannot currently process executables produced with the `-om` switch. This is a

limitation with the `om` utility that will be corrected in a future release.

5.3.2 ATOM Routines Should Not Call Certain Libraries

Section 9.2.5 of the *Programmer's Guide*, "Writing Analysis Procedures," incorrectly states that an analysis routine can use any system call or library function. Some library routines cannot be called safely from analysis routines in all system configurations.

Standard C Library (`libc.a`) routines, including system calls, and Math Library (`libm.a`) routines can be called, except for:

- `unwind` other exception-handling routines
- `pthread_atfork`
- `tis` and the standard I/O routines

  They have certain differences in behavior, as described in Section 9.2.5.1 of the *Programmer's Guide*.
- Routines related to multi-threading or exception-handling (like `pthread`, `exc_*` and `libmach` routines)
- Routines that assume a particular environment (like X and Motif)

  They may not be useful or correct in an ATOM analysis environment.

5.4 Java Programming

The following note applies to Java programming.

5.4.1 Name Space Conflict Between Java and SVE

A name space conflict exists between Java and the System V Environment (SVE) on DIGITAL UNIX Version 4.0* systems.

The problem arises because both Java and SVE use the string `/usr/bin/alpha` for different purposes. Java creates `/usr/bin/alpha` as a directory. SVE (specifically, the optional SVEBCP4** Base Compatibility Package subset) creates `/usr/bin/alpha` as a symbolic link to the `/usr/opt/svr4/usr/bin/alpha` directory. The order in which these applications are installed determines if the customer will experience a problem. Here are three ways to avoid the name space 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 alternate link location, `/usr/opt/svr/usr/bin/alpha`. That link will always be there and does not cause a conflict.

- If SVE has not yet been installed, Compaq recommends that SVE be installed without the Inclusive View option. If Inclusive View is not selected during SVE installation, the `/usr/bin/alpha` link will not be created.
- If SVE has not yet been installed and the customer insists on having the Inclusive View option, then the customer must:
  1. Temporarily rename the `/usr/bin/alpha` directory that is used by Java
  2. Install SVE
  3. Remove the `/usr/bin/alpha` link created by the SVE installation
  4. Restore the Java `/usr/bin/alpha` directory

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

Window System Software Notes 6

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

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

6.1 Hardware Notes and Restrictions

The following notes apply to graphics hardware restrictions.

6.1.1 Qvision Graphics Display Error

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

6.1.2 Graphic Adapter Changes Needed to Support EV6 Linear Space Addressing

Graphics adapters that support sparse space addressing must be modified to support linear (byte) space addressing to run on EV6 class systems. This support has been provided for the S3Trio graphics adapter supplied in DIGITAL UNIX. Generic VGA support is also provided.

If you develop support for a graphics adapter that uses linear space addressing, you must modify the device driver and DDX library to support linear space addressing on the EV6 class of systems.

Note that the following information pertains only to graphics adapters that support both sparse and linear space addressing.

In addition to the specific changes required in the DDX library, you must compile routines that perform I/O with the `-arch ev56` switch to take advantage of the linear (byte) addressing capabilities of the system. Because this results in two versions of some routines, you must decide if you want to implement your solution as a single DDX library or as two libraries. The S3Trio and generic VGA DDX layers use the two library option.

6.1.2.1 Single and Two DDX Library Issues

The following lists a general issues and possible implementations for the single or two library approach. Note that this is not an exhaustive list.

- Single DDX Library
  - Two versions of addressing (sparse and linear) are required for all I/O routines and therefore these routines require two different names.
  - Two versions of I/O routines must coexist in the DDX library.
  - Linear versions of routines must be compiled with the `-arch ev56` switch.
  - Programs may contain conflicting definitions for variables for modules that are compiled twice and loaded into the same library.
  - The X server must be notified about which set of routines to use at run time; this is typically done during the DDX initialization phase.
  - You must set up a mechanism to inform the DDX layer which set of routines to load or run.
- Two DDX Libraries
  - All routines must be compiled twice and loaded into separate DDX libraries.
  - The linear version of library must be compiled with `-arch ev56` switch.
  - The X server has to be notified which library to load.
  - This implementation of the X server provides a mechanism to identify which version of the library to load.

6.1.2.2 Required Changes to Drivers

This section provides a general description of the driver and DDX library changes needed to support both linear and sparse space addressing for the same graphics adapter.

Within the X server environment, the driver is responsible for mapping memory. Before the memory is mapped, the driver must determine if the system supports sparse or linear space addressing. One method of doing this is to use the `get_info` kernel interface, as shown in Example 6-1.

Example 6-1: Determining the Type of Space Addressing

```
#define IOH_FLAGS       (HANDLE_BYTE | HANDLE_SPARSE_SPACE)
#define IOL_FLAGS       (HANDLE_LINEAR_SPACE)

int map_flags = 0;

request.next_function = NULL;
request.out_flags = 0;
request.in_flags = 0;
request.function = BYTEWORD_IO_CAPABLE;

if ((get_info(&request) == NOT_SUPPORTED) ||
(request.rtn_status == NOT_SUPPORTED))
map_flags = IOH_FLAGS;
else
map_flags = IOL_FLAGS;
```

Note that this example defaults to sparse space.

If you have chosen the single DDX library solution, you must provide a mechanism for the driver to inform the DDX layer which routines to load or run. Typically, this would be done during the DDX initialization phase.

If you have chosen the two DDX library solution, then you can use the mechanism provided with the X server to load the appropriate library during the initialization phase. To use this mechanism, you must notify the DDX layer about which type of mapping mapping (sparse or linear) to use.

The driver must set the `reserved1` field of the `ws_descriptor` structure to identify the mapping type.

Note that this value must be set before the X server makes the `GET_WORKSTATION_INFO` ioctl call. Placing this step in the `attach` routine will ensure this. A value of 1 instructs the X server to load the sparse space library while a value of 2 instructs it to load to load the linear space library. A value of 0 loads the current default.

Example 6-2 shows a sample attach routine to identify and set the mapping type.

Example 6-2: Informing the DDX Layer of the Mapping Type

```
#define SPARSE_SPACE    1
#define LINEAR_SPACE    2
#define IOH_FLAGS       (HANDLE_BYTE | HANDLE_SPARSE_SPACE)
#define IOL_FLAGS       (HANDLE_LINEAR_SPACE)
(map_flags defined above)
ws_info *wi = &ws_softc[0];

if (map_flags == IOL_FLAGS)
wi->ws.reserved1 = LINEAR_SPACE;
else if (map_flags == IOH_FLAGS)
wi->ws.reserved1 = SPARSE_SPACE;
else
wi->ws.reserved1 = 0;
```

Note

You can also include this code within Example 6-1 as long as you ensure that the `reserved1` field is set before the X server makes the `GET_WORKSTATION_INFO` ioctl call.

Once the the driver has determined the type of addressing to use it can map the memory, as shown in Example 6-3.

Example 6-3: Mapping Memory

```
io_handle_t handle;
caddr_t temp;
int nbytes;
(map_flags defined above)

handle = vp->mem_handle;
temp = (caddr_t) iohandle_to_phys(handle, map_flags);
temp = (caddr_t) PHYS_TO_KSEG(temp);

nbytes = iohandle_to_phys(handle + vp->mem_size, map_flags) -
iohandle_to_phys(handle + 0           , map_flags);

dp->pixmap = ws_map_region(temp, NULL, nbytes, 0600, (int *)NULL);

if ( dp->pixmap == (caddr_t) NULL )
return(ENOMEM);

handle = busphys_to_iohandle(IOREGS_BASE, BUS_IO, ctlr);
temp = (caddr_t) iohandle_to_phys(handle, map_flags);
temp = (caddr_t) PHYS_TO_KSEG(temp);

nbytes = iohandle_to_phys(handle + IOREGS_SIZE, map_flags) -
iohandle_to_phys(handle + 0,           map_flags);

dp->plane_mask = ws_map_region(temp, NULL, nbytes, 0600, (int *)NULL);
```

Example 6-3: (continued)

```
if (dp->plane_mask == (caddr_t)NULL)
return(ENOMEM);

handle = busphys_to_iohandle(ENHANCED_REG_BASE, BUS_MEMORY,ctlr);
temp = (caddr_t) iohandle_to_phys(handle, map_flags);
temp = (caddr_t) PHYS_TO_KSEG(temp);

nbytes = iohandle_to_phys(handle + ENHANCED_REG_SIZE, map_flags) -
iohandle_to_phys(handle + 0           , map_flags);

dp->physaddr = ws_map_region(temp, NULL, nbytes, 0600, (int *)NULL);

if (dp->physaddr == (caddr_t)NULL)
return(ENOMEM);

SET_VGA_MAPPED(vp);
```

6.1.2.3 Required Changes to the DDX Library

Regardless of which DDX layer solution you choose, you must modify all I/O routines to take advantage of the byte addressing capabilities with linear space addressing. In general, you must change I/O variables to bytes and compile your code with the `-arch ev56` switch.

If you have chosen the single DDX library solution, you must provide a mechanism for the driver to inform the DDX layer which routines to load. Typically, this is done during the DDX initialization phase. You can choose to load the appropriate I/O routine addresses into the DDX data structures during this phase or you can set an internal flag to select the appropriate I/O routine at run time. Only those DDX routines that perform I/O needs two versions: one for sparse space, and one for linear space that is compiled with the `-arch ev56` switch.

If you have chosen the two DDX library solution, you can use the mechanism provided with the X server to load the appropriate library. To use this mechanism, you must create two DDX libraries for the same graphics adapter.

The linear space version of the library must be compiled with the `-arch ev56` switch. The linear space libraries must have the same name as the sparse space version but with `_linear` appended to the library name. The X server attempts to load the linear version if the `reserved1` field is set to `LINEAR_SPACE` as previously described.

For the two library solution you must also add an entry for the second (linear space) library in the `Xserver.conf` file. The device name must be

devnameLINEAR with the library name as shown in the following example.

```
S3Trio DDX library names:
sparse space  - lib_dec_s3.so
linear space  - lib_dec_s3_linear.so
```

The devname must match the name used for the sparse space library in Xserver.conf.

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. DIGITAL is currently working with OSF to resolve issues related to multiscreen environments for future releases of CDE.

6.2.2 Do Not Modify the keymaps.dir File

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

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

6.2.3 CDE Does Not Work Properly With More than Seven Graphics Heads

If you are using more than seven graphics heads with the DMCC platform, you must use the XDM window manager. CDE does not operate properly on this configuration.

6.3 X Clients

The following notes apply to X clients.

6.3.1 The dxmail and mh Applications Require New MailScanFormat File

The `/usr/lib/X11/app-defaults/MailScanFormat` file, used by `dxmail` and other `mh`-based mailers, has been updated to support Year 2000 dates. Use of an old `MailScanFormat` file will result in date display problems and possible mail filename corruption. Therefore, you must replace any local, customized copies of the `MailScanFormat` file with the new file.

You should also update and review all local `dxmail` resource and `.mh_profile` files to verify that no other old versions of configuration files are being referenced.

6.3.2 X Window Colormap Resources

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

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

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

- Use the CDE Style Manager's Color application to select the number of colors used by the desktop or an application. Alternatively, if supported:
  - Set an application resource to limit color usage.
  - Start the application with a flag that controls color map installation. (Refer to the application documentation for further details.)
- In the case of Netscape, two workarounds are available; the first is preferred:
  - Use Netscape's `maxImageColors` resource to limit the number of colors that Netscape uses. A suggested limit is 96. You can do this by placing the following line in the `$HOME/.Xdefaults` file:

    ```
    Netscape*maxImageColors:        96
    ```

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

- Noticeable application and background color changes when focus is moved in and out of the Netscape window
- Inability to clearly view the contents of the Netscape window when focus is moved out of the Netscape window

6.3.3 Old X Environment

If the XDM graphical login environment is selected instead of the default CDE environment, you may need to install the `OSFOLDX11425` 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 more familiar DIGITAL session manager.

6.4 CDE Clients

The following notes apply to CDE clients.

6.4.1 Inaccessible Dialog Buttons

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

6.4.2 Screen Savers Prevent Efficient Power Management

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

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

6.4.3 Remote Invocation of CDE File Manager dtfile

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

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

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

6.4.4 The dtmail Application

The following notes apply to the `dtmail` application.

6.4.4.1 MAILLOCKING Environment Variable

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

6.4.4.2 Tooltalk Locking Disabled by Default

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

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

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

```
cdenotooltalklock='f'
```

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

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

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

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

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

showing the following message:

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

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

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

   ```
   set cdenotooltalklock
   ```

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

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

1. From the Mailbox pull down menu, select Mail Options
2. From the Mail Options dialog box, select the category Advanced.
3. From the Advanced options, click on the option named: Use network aware mail file locking. This action should remove the check mark.
4. Close all dialog boxes and restart the `dtmail` application. Check the contents of your current mailbox and the backup mailbox to ensure that no mail was lost during this process.

6.4.4.3 Row and Column Display

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

6.4.4.4 Attachments Are Not Included

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

- Compose->New, Include All
- Compose->Reply to Sender, Include
- Compose->Reply to All, Include

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

6.4.4.5 dtmail Startup

Incorrect default permissions on `/dev/zero` prevent `dtmail` from starting.

You may see the following message when starting `dtmail`:

```
No memory available for operation
```

If you see this message, set the permission mode on `/dev/zero` to 666, as follows:

```
# chmod 666 /dev/zero
```

6.4.5 CDE's Static Dependency on the Network

CDE has a static dependency on the state of the network configuration. For more information, see Section 4.4.6.

6.4.6 The xnlLanguage Resource Settings May Cause Problems with CDE

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

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

6.4.7 Possible Failure in the XOpenDisplay Call

When logging in to the CDE desktop, not all of the applications you want may 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.

For a multiuser system, this resource can be added to the `/usr/dt/app-defaults/C/Dtsession` file to make the change for all users automatically.

6.5 Windows Programming

The following note applies to windows programming.

6.5.1 X11 Font Problem

In the `Lucida-Typewriter-medium-R-normal` fonts, the glyphs for multiplication and division are reversed. The multiplication symbol is where the division symbol should be, and vice versa. If the representation of these glyphs are 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 Use of PC-style Keyboard with Asian Input Method

When using a PC-style keyboard, the Hanyu, Hanzi and Hangul input method servers may not recognize the Backspace key as the key to delete the previous character in the preedit area of the input method. In this case, you should use the Delete key to delete the previous character.

The default Alt-space key for activating and deactivating the input methods may not work under CDE. If that happens, go to the Input Method Customization dialog box and change the Start Input Method and End Input Method keys to other key sequence like Ctrl-space.

6.6.2 Japanese Keyboard Support

DIGITAL UNIX now supports two new Japanese keyboard types (JIS and ANSI) on AlphaStation and AlphaServer systems.

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

```
>>> set language 50
```

To use ANSI-type Japanese keyboards, like the LK411-AJ, you must set the `language` console environment variable to 52, as in the following example:

```
>>> set language 52
```

Japanese keyboard support is not available on TURBOchannel-based machines. See Section 6.6.3 for instructions on how to load Japanese keymaps when the `language` is not available at the system console.

6.6.3 Default Keyboard Mapping

The X server automatically chooses a keymap based on the language and keyboard settings of the system console. If your locale is not available on the system console, or you want the server to load a different keymap, you must set the system default keyboard map. The default keyboard map is specified by adding the `-xkbmap` option to the args list in the `/usr/var/X11/Xserver.conf` file. Add the `-xkbmap` option by using the following syntax:

```
! you specify command line arguments here
args <
-pn -xkbmap <keymap_file_name>_<keymap_name>
>
```

For example:

```
-pn -xkbmap digital_japanese_lk411aj
```

The available keymap files are located in `/usr/lib/X11/xkb/keymap`, where there is one file for each locale. The individual keymaps for the locale are in the keymap file, and are specified by the keyword `xkb_keymap`. For example, the `/usr/lib/X11/xkb/keymap/digital_japanese` keymap file contains an entry for the LK411-AJ keymap.

In addition, you can change keymaps after logging in by running the Keyboard Options desktop application `/usr/dt/bin/dxkeyboard` if you are using CDE. Use the keyboard setting option of the session manager if you are using `xdm`.

6.6.4 Lithuanian and Slovene Language Support

The following instructions are for modifying some system files to allow for some support of Lithuanian and Slovene.

6.6.4.1 Mnemonics in Lithuanian and Slovene DX Applications

DX applications in this release do not support the use of mnemonics in Lithuanian and Slovene language variants.

6.6.4.2 DX Applications Integrated to CDE Desktop (Application Manager)

To make the Visual Differences application `/usr/bin/X11/dxdiff` display differences between two text files written in Lithuanian or Slovene language, locate the following line in the `/usr/lib/X11/app-`

`defaults/DxDiff` resource file:

```
*dxdiff*textdisplay*fontList:           fixed
```

Change the line, as follows:

```
*dxdiff*textdisplay*fontList:   -*-terminal-medium-r-narrow--18-*-*-*-c-*-*-*
```

You need superuser (root) privileges to make this change.

6.6.4.3 DX Applications Not Integrated into the CDE Desktop (Application Manager)

Some DX applications that are not integrated as part of the CDE desktop, but can be invoked using the File Manager or command line, do not set all display fonts according to the locale in use. If users want these applications to provide full Lithuanian and/or Slovene language support, they must perform the following corrective actions on the system with the Lithuanian and/or Slovene language variant installed.

Note that the superuser must perform all of these actions.

- Mail – `/usr/bin/X11/dxmail`

  To enable the use of Lithuanian and/or Slovene national characters in `dxmail`, locate the following lines in the `/usr/lib/X11/app-defaults/DXMail` resource file:

  ```
  *outlineList*DXmfontListDefault: *-*-*-Medium-R-Normal--*-120-*-*-M-*-ISO8859-1
  *outlineList*DXmfontListLevel0: *-*-*-Medium-R-Normal--*-120-*-*-M-*-ISO8859-1
  *outlineList*DXmfontListLevel1: *-*-*-Medium-R-Normal--*-120-*-*-M-*-ISO8859-1
  *tocList*DXmfontListDefault: *-*-*-Medium-R-Normal--*-120-*-*-M-*-ISO8859-1
  *tocList*DXmfontListLevel0: *-*-*-Medium-R-Normal--*-120-*-*-M-*-ISO8859-1
  *Item.fontList: *-*-Helvetica-Bold-R-Normal--*-100-*-*-*-*-ISO8859-1
  *XmText.FontList: *-*-*-Medium-R-Normal--*-120-*-*-M-*-ISO8859-1
  *Text.FontList: *-*-*-Medium-R-Normal--*-120-*-*-M-*-ISO8859-1
  ```

  Change the lines, as follows:

  ```
  *outlineList*DXmfontListDefault: *-*-*-Medium-R-Normal--*-120-*-*-M-*-*-*
  *outlineList*DXmfontListLevel0: *-*-*-Medium-R-Normal--*-120-*-*-M-*-*-*
  *outlineList*DXmfontListLevel1: *-*-*-Medium-R-Normal--*-120-*-*-M-*-*-*
  *tocList*DXmfontListDefault: *-*-*-*-R-Normal--*-120-*-*-*-*-*-*
  *tocList*DXmfontListLevel0: *-*-*-*-R-Normal--*-120-*-*-*-*-*-*
  *Item.fontList: *-*-*-*-R-Normal--*-120-*-*-*-*-*-*
  *XmText.FontList: *-*-*-Medium-R-Normal--*-120-*-*-*-*-*-*
  *Text.FontList: *-*-*-Medium-R-Normal--*-120-*-*-*-*-*-*
  ```

- Notepad – `/usr/bin/X11/dxnotepad`

  To enable the use of Lithuanian and/or Slovene national characters in `dxnotepad`, locate the following line in the `/usr/lib/X11/app-`

defaults/DXnotepad resource file:

```
*textwindow.fontList: -*-Terminal-Medium-R-Narrow--*-140-*-*-C-*-ISO8859-1
```

Change the line, as follows:

```
*textwindow.fontList: -*-Terminal-Medium-R-Narrow--*-140-*-*-C-*-*-*
```

- Bookreader – /usr/bin/X11/dxbook

  To enable the use of Lithuanian and/or Slovene national characters in Bookreader bookshelf and book names, locate the following line in the /usr/lib/X11/app-defaults/DXBookreader resource file:

  ```
  -*-Menu-Medium-R-Normal--*-120-*-*-P-*-ISO8859-1
  ```

  Change the line, as follows:

  ```
  -*-*-*-R-Normal--*-120-*-*-P-*-*-*
  ```

- Paint – /usr/bin/X11/dxpaint

  The dxpaint Lithuanian and Slovene language variants do not support the insertion of Lithuanian and Slovene national characters.

- Clock – /usr/bin/X11/dxclock

  The dxclock application uses English day/month abbreviations and format to display the date and time.

- Calendar – /usr/bin/X11/dxcalendar

  The dxcalendar application uses English day/month abbreviations and format to display the date and time.

  To enable the use of Lithuanian and/or Slovene national characters to create calendar entries, locate the following lines in the /usr/lib/X11/app-defaults/DXcalendar resource file:

  ```
  *font_small_tb.fontList:     -*-Menu-Medium-R-Normal--*-100-*-*-P-*-ISO8859-1
  *font_medium_tb.fontList:    -*-Menu-Medium-R-Normal--*-120-*-*-P-*-ISO8859-1
  ```

  Change the lines, as follows:

  ```
  *font_small_tb.fontList:      -*-*-*-R-Normal--*-100-*-*-P-*-*-*
  *font_medium_tb.fontList:     -*-*-*-R-Normal--*-120-*-*-P-*-*-*
  ```

6.6.4.4 Lithuanian Language Variant

If you use the Lithuanian language variant and you want to operate in the DXsession environment, you must specify the default user interface font by adding the following line to the .Xdefaults file located in your home

directory and logging in again for the change to take effect:

```
*FontList:      -*-*-*-r-Normal--*-120-*-*-*-*-*-*
```

If you operate in a multilingual user interface environment, this corrective action will cause other language variants to display using a default font that is slightly different than the font family used before this action was performed.

If you want to use the default user interface font with the Motif window manager, locate the following line in the `/usr/lib/X11/app-defaults/Mwm` resource file:

```
Mwm*fontList: -*-Menu-Medium-R-Normal--*-120-*-*-P-*-ISO8859-1
```

Change the line, as follows:

```
Mwm*fontList: -*-*-*-R-Normal--*-120-*-*-P-*-*-*
```

You need superuser (root) privileges to make this change.

6.6.5 Dxterm Does Not Support Latin-2, Latin-4 and Latin-Cyrillic Characters

The `dxterm` application is not able to display Latin-2, Latin-4 and Latin-Cyrillic Characters even when the locale is set correctly. Therefore, `dxterm` should not be used for displaying the following languages: Czech, Hungarian, Lithuanian, Polish, Russian, Slovak, and Slovene. The `dtterm` application should be used in this case.

Documentation Notes 7

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

- *Documentation Overview*
- *Network Administration*
- *System Administration*
- *Assembly Language Programmer's Guide*
- *Programming with ONC RPC*
- *Developing Applications for the Display PostScript System*
- *Gate Daemon Configuration Guide*
- *System Configuration and Tuning Guide*
- Online help volumes

7.1 Documentation Overview

The online version of the *Documentation Overview, Glossary, and Master Index* has been changed in the following ways for this release:

- It has been divided into two separate books: *Documentation Overview* and *Glossary and Master Index.*
- Information in the *Documentation Overview* has been reorganized and expanded to provide a more comprehensive view of the DIGITAL UNIX documentation set and other information of importance to users of the DIGITAL UNIX operating system.

Because only the online version has changed, that version differs substantially from the printed version.

Also, the Master Index was not revised for this release; therefore, some of the index entries for books that have been revised since the Version 4.0 release may be inaccurate.

A new printed version of the *Documentation Overview* will be offered with the next major release of DIGITAL UNIX. At that time, the *Glossary and Master Index* will also be updated and provided as a separate book in both online and printed versions.

7.2 Network Administration Guide

The following sections describe changes to the *Network Administration* manual.

7.2.1 Section 4.2.2.2 (Verifying PPP Support in the Kernel) Is Obsolete

The following instructions replace the instructions in Section 4.2.2.2:

To verify that PPP is supported in the kernel, enter the following command:

```
# sysconfig -s | grep ppp
```

If it is not loaded and configured, do the following:

1. Log in as root.
2. Save the `/vmunix` file.
3. Rebuild the kernel by running the `doconfig` program and selecting the Point-to-Point Protocol (PPP) option.
4. Copy the new `vmunix` file to `/vmunix`.
5. Add the following entry to the `/etc/sysconfigtab` file using the `sysconfigdb` utility:

   ```
   ppp:
   nppp=2
   ```

   This provides for 2 PPP connections. If your system requires a greater number of PPP connections, increase the number.
6. Reboot the system.

7.2.2 Sections 4.2.3.1, 4.2.3.2, and 4.2.3.3 Are Obsolete

The following instructions replace the instructions in Sections 4.2.3.1, 4.2.3.2, and 4.2.3.3:

7.2.2.1 Establishing a PPP Dial-Out Connection

After you have connected your modem to a serial port on your system, or installed a supported modem card option, do the following:

1. Verify that you can communicate with the modem. Do the following:
   a. Edit the `/etc/remote` file and copy the `kdebug` entry.
   b. Modify the new entry, providing a system name for the entry, the correct DIGITAL UNIX device (`tty00` or `tty01` depending on your system), the correct baud rate, and correct parity. See the

remote(4) reference page for more information.

c. Use the tip command to access the modem as follows:

```
% tip system_name
```

system_name is the system name from the /etc/remote file.

d. If your modem is using the AT command language, enter the following command:

AT Return

If the modem is not in quiet mode, it responds with an OK message.

2. Contact the remote system administrator or your internet service provider (ISP) and obtain the following information:
   - Your remote IP address and netmask, unless the remote system assigns the IP address dynamically
   - Characters that might need to be escaped
   - Instructions on how to log in and use the remote service

   This information is used to create a chat script, which automates the dial-out process.

3. Create a file for commands that the chat program uses to direct the modem what number to dial and what to send the remote system in order to start pppd. This file is called a chat script. Each entry in a chat script has the following format:

   string_chat_expects string_chat_sends

   For example, the following file, named /etc/ppp/chat-script, contains the following information:

```
"" atdt2135476    1
CONNECT ""    2
login: myname      3
Password: "\qmypassword"    4
"$ " "\qpppd"       5
```

   1 The chat script expects nothing and sends a dial command to the modem.

   2 The chat script expects a CONNECT message and sends a carriage return (implied).

   3 The chat script expects the login: string and sends the myname string.

   4 The chat script expects the Password: string and sends the mypassword string. The \q escape sequence prevents chat from logging the password when you use the -v option.

5 The `chat` script expects the `$` (the shell prompt) and sends `pppd` to start the `pppd` daemon on the remote machine. The `\q` escape sequence cancels the effect of the previous `\q`.

See the `chat`(8) reference page for more information on `chat` and chat scripts.

Note

You might want to use the `tip` command to dial out and log in to the remote system and to write down the exact prompt, login sequence, and `pppd` start-up sequence.

4. Edit the `/etc/ppp/options` file and include the following `pppd` options as required by the remote system or ISP:

```
defaultroute   1
asyncmap 0  2
mru 296 3
netmask dd.dd.dd.dd  4
lcp-echo-interval 60  5
lcp-echo-failure 5   6
noipdefault  7
crtscts  8
debug  9
```

1 If your system is a standalone and you are connecting to the Internet through the remote system, add a default route via the remote host by specifying this option.

2 If the serial line is not completely 8-bit transparent, specify this option; `asyncmap 200a0000` is appropriate if the serial link includes a `telnet` link.

3 Reduces the maximum receive unit (MRU) on the local and remote systems to improve performance for multiple IP connections.

4 Sets the interface netmask to the specified value. Your ISP should provide this information.

5 Sends a Link Control Protocol (LCP) echo request frame to the remote system every 60 seconds. This determines whether the link to the remote system is still active.

6 If the local system does not receive a response from the remote system after five LCP echo request frames, `pppd` considers the link dead and tears down the connection.

7 Specifies that the remote system (ISP) is to provide an IP address to the local system, unless an IP address is specified explicitly on the command line or in an options file.

8 Enables hardware flow control on the serial device. If the modem does not support hardware flow control, do not add this entry. See your modem documentation to verify this information.

9 Enables debugging. All log messages are sent to the file specified in the `/etc/syslog.conf` file. After your connection is working correctly, remove this entry from the PPP options file.

See the `pppd`(8) reference page for a complete list of `pppd` options.

5. Edit the `/etc/syslog.conf` file and do the following:
   a. Add the `local2` facility (used by `pppd` and `chat`) to the line that specifies `/dev/console` as the message destination, as follows:

      ```
      kern.debug;local2.notice      /dev/console
      ```

      In this example, the `notice` level is specified.
   b. Add the following entry to the file to create a `ppp-log` file:

      ```
      local2.debug       /etc/ppp/ppp-log
      ```

   c. Save the edits and close the file.
6. Stop and restart `syslogd` by entering the following commands:

   ```
   # /sbin/init.d/syslog stop
   # /sbin/init.d/syslog start
   ```

7. Invoke `pppd` on the local system to connect to the remote system. For example, the following command starts a link on `tty01` and specifies the `connect` option to run the `chat` program and the specified chat script file.

```
% pppd /dev/tty01 38400 connect 'chat -v -f /etc/ppp/chat-script'
```

8. Issue the following command to monitor the `ppp-log` file and to determine whether the PPP connection is active:

   ```
   % tail -f /etc/ppp/ppp-log
   ```

If any problems occur while using PPP, see Chapter 13 in *Network Administration*.

7.2.2.2 Establishing a PPP Dial-In Connection

After you have connected your modem to a serial port on your system, to configure a dial-in system, complete the following steps:

1. Set up your modem for dial-in access. See Section 4.3.2 in *Network Administration* for more information.
2. Edit the `/etc/passwd` file and create a dedicated entry for a PPP user.

For the login shell field, specify /usr/sbin/startppp; for example:

```
ppp1:password:10:20:Remote PPP User:/usr/users/guest:/usr/sbin/startppp
```

3. Edit the /etc/inittab file and create an entry for each terminal device that is to run PPP. For example:

   ```
   modem:3:respawn:/usr/sbin/getty /dev/tty00 M38400 vt100
   ```

   See the inittab(4) reference page for more information.

4. Issue the init q command to start the getty process immediately.

5. If the dial-in system is going to be a gateway for the dial-out system to reach other systems on the LAN, the dial-in system must be configured as an IP router and must also run gated. Edit the /etc/gated.conf file and delete the nobroadcast option (if specified) in the rip statement. See Chapter 2 of *Network Administration* for basic network setup information and the gated.conf(4) reference page for gated options.

6. Edit the /etc/ppp/options file and include the following pppd options required to support dial-in access for all remote users:

   ```
   netmask dd.dd.dd.dd  1
   proxyarp  2
   crtscts  3
   asyncmap 0  4
   :remote_ip_address  5
   debug  6
   ```

   1 Sets the interface netmask to the specified value.

   2 Adds an entry to the local system's Address Resolution Protocol (ARP) table that contains the IP address of the remote system and the Ethernet address of the local system. This is not necessary if gated is running.

   3 Enables hardware flow control for the serial port.

   4 If the serial line is not completely 8-bit transparent, specify this option; asyncmap 200a0000 is appropriate if the serial link includes a telnet link.

   5 Specifies an IP address for the remote system.

   If you want to specify options for each individual serial port, create a /etc/ppp/options.tty*xx* file and include the remote IP address and any other options that apply to that specific serial port. See the pppd(8) reference page for a complete list of pppd options.

7. After an incoming call is received and a connection established, startppp runs in the background. The process ID is logged in the /etc/ppp/ppp*xx*.pid file.

8. Enables debugging. All log messages are sent to the file specified in the `/etc/syslog.conf` file. After your connection is working correctly, remove this entry from the PPP options file.

If any problems occur while using PPP, see Chapter 13 in *Network Administration*.

7.3 System Administration Guide

The following notes describe changes to the *System Administration* guide.

7.3.1 Hard Limit for File System Quotas

In the *System Administration* guide, file system hard and soft quota limits are discussed in:

- Section 6.7.3.1, Hard and Soft Quota Limits for the UFS File System
- Section 7.5, Managing File System and Fileset Quotas for the AdvFS File System
- Section 9.3.4, Setting File System Quotas for the UFS File System

None of these sections discuss the fact that a hard limit is one more unit (blocks, files, inodes) than will be allowed when the quota limit is active. The quota is up to, but not including the limit. For example, if you set a hard limit of 10,000 disk blocks for each user account in a file system, an account reaches the hard limit when 9,999 disk blocks have been allocated. If you want a maximum of 10,000 complete blocks for the user account, you must set the hard limit to 10,001.

7.3.2 Environmental Monitoring

In the Environmental Monitoring section, the *System Administration* guide incorrectly states:

When a fan failure is encountered, a message is broadcasted and an orderly shutdown ensues.

It should state:

When the cooling fan on an AlphaServer 1000A fails, the kernel logs the error, synchronizes the disks, then powers the system down. On all other fan failures, a hard shutdown ensues.

7.3.3 User Definable Messages in Environmental Monitoring

You can modify any messages broadcast or logged by the Environmental Monitoring utility. The messages are located in the following file:

`/usr/share/sysman/envmon/EnvMon_UserDefinable_Msg.tcl`

You must be root to edit this file. You can edit any message included in braces ({}). The instructions for editing this file are included in the comment (#) fields. Do not alter any other data in this file.

7.4 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.

7.5 Programming with ONC RPC

Sections 2.1.5 and 2.1.6 of the *Programming with ONC RPC* guide are correct only for a DIGITAL UNIX system. The code may not be portable to other UNIX systems.

If the client machine is not another DIGITAL UNIX system, copy the source code to the client. Check that the header calls are appropriate for the programming conventions on the client operating system and then compile the code locally.

7.6 Developing Applications for the Display PostScript System

The guide for *Developing Applications for the Display PostScript System* refers to a nonexistent example directory `/usr/examples/dps`. None of the examples specified in the documentation are provided.

7.7 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 *DIGITAL UNIX 4.0E Documentation, Volume 1* 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).

7.8 System Configuration and Tuning Guide

The DIGITAL UNIX Version 4.0D *System Configuration and Tuning Manual* contains an error in Appendix B.

The manual incorrectly states that the default value of the vm attribute `new-wire-method` is 0.

The correct default value of `new-wire-method` is 1.

This attribute is only for internal use. Do not modify the value of `new-wire-method.`

7.9 Online Help Volumes

The notes in this section refer to problems with the online help volumes.

For each of the SysMan applications, online help is available from the Help menu or from the Help button in the main window.

The online help contains:

- An overview of the capabilities of the application
- A set of tasks illustrating typical uses of the application
- A reference section that documents every window and dialog box in the application.

The Using Help item on the Help menu displays a help volume that explains how to use the online help.

The Help viewer has the following known problems:

- In some cases, the help viewer is not correctly initialized. As a result, it will sometimes exhibit the expected behavior the second time an action is taken, but not the first time.

  For example, the first time a quick help dialog box displays a reference page, the Backtrack button may be enabled even though there is no place to which to backtrack. If the dialog box is closed and then opened again, the Backtrack button is dimmed.
- Similarly, if a request for on-item help displays the correct help, but one line has scrolled off the top, the online help will typically be displayed in exactly the right position when the on-item help request is repeated.

7.9.1 General Problems

The following known problems occur in the online help:

- The Appearance menu is not consistently documented.

  In the Archiver, License Manager, and System Information applications,

there is an Appearance item on one of the menus. The item should offer three options: Text Only, Large Icon, and Small Icon. In some help volumes, not all of these are documented.

- Some links from one help volume to another are displayed in a new view.

  There are a few links from one help volume to another that appear in a new help viewer window. In some situations, a distracting proliferation of help viewer windows can result.

  You can get a new view when you want one using the New Window item on the File menu in the help viewer.

7.9.2 Integration

The SysMan configuration applications on-item help does not work on the items in the menu bar.

In all the SysMan applications, the keyboard method of getting on-item help does not work on the Help menu.

7.9.3 Update Path in Installation Online Help Is Incorrect

The update installation path shown in the installation online help is not correct. The Preinstall Review List in the installation online help should state:

```
Your system must be running DIGITAL UNIX operating system Version
4.0A, Version 4.0B, or Version 4.0C in order to update to
DIGITAL UNIX Version 4.0E.
```

7.9.4 Help Volumes by Application

The following problems apply only to help in specific applications:

- Kernel Tuner

  The Kernel Tuner application records its changes immediately, so if a system failure occurs while the Kernel Tuner is running, any boot-time changes will take effect the next time the system boots. If the `/etc/sysconfigtab` file contains invalid values, you can enter the following command at the boot prompt to boot with default values:

  ```
  boot -fl c
  ```

- Display Window

  The help volume for Display Window has opening instructions that show how to access the application from the CDE Application Manager. These should show that the Display Window icon appears in two groups.

- Network Configuration

  The names of the `gated`, `joind`, `routed`, and `rwhod` daemons are misspelled in the online help volume.

  In the Configuring Interfaces dialog box, the fields under To Obtain IP Address are relevant for all interfaces.

- Printer Configuration

  In the Local Printer Settings dialog box and the Remote Printer Settings dialog box, the aliases in the Printer Aliases field must be separated by vertical bars because spaces are allowed in an alias.

- Shutdown

  The `/usr/sbin/shutdown` command now sends the Shutdown Message each time a reminder is sent of the time remaining before the system is shut down.

  The `/etc/nologin` file is not created until immediately before the shutdown occurs.

  System shutdown messages are sent to all users that are locally or remotely logged in to the system being shut down. Additionally, if the "Broadcast to NFS Clients" option is selected, shutdown messages are broadcast to all hosts that are NFS clients of the system being shut down.

- Disk Configuration

  The following corrected definitions replace the definitions in the glossary for Disk Configuration.

  – Skew

    A deviation from a reference direction, either by design or in response to lateral forces.

  – Track skew

    On a disk, the sector skew per track. The skew is the angle that sector 0 of the track changes from an imaginary radius line, due to a nonuniform number of sectors per track.

- BIND Configuration

  The following corrected definitions replace the definitions in the glossary for BIND Configuration:

  – BIND client

    A system that queries a BIND server for host name and address information, interprets the responses, and passes the information to requesting applications.

- BIND server

  An authoritative source for information about one or more zones. It either maintains the master copy of the hosts database for the zone or obtains the information required to serve the hosts database from another server.

- DCE

  Distributed Computing Environment. The capabilities of DCE are defined by the Open Software Foundation (OSF).

- DCE cell

  A logical group of systems that share services offered by DCE.

- DCE server

  The server in a DCE cell.

- service type

  In BIND Configuration, the available service types are BIND client and BIND server. The service type determines whether a system is configured to be a BIND client or a BIND server.

Features and Interfaces Scheduled for Retirement 8

This chapter lists features of DIGITAL UNIX scheduled to be removed from, or changed in, future major functional releases. Users and developers should plan to migrate away from these features in the near future.

8.1 DECwindows Applications

The following DECwindows utilities and tools are scheduled for retirement in a future release of DIGITAL UNIX. These dx* tools and utilities, commonly known as DECwindows, have been replaced by the dt* tools in Common Desktop Environment (CDE) that were introduced in Version 4.0 of DIGITAL UNIX. The replacement applications are listed in Table 8-1. Not all of the dx* applications being considered for retirement have a replacement due to the limited use or capability of that specific tool/utility.

Table 8-1: DECwindows Applications to be Retired

Tools/Utilities to be Retired	Replacement Option(s)
dxprint	dtlp
dxmail	dtmail,xmh,exmh,Netscape mail
dxcalendar	dtcm
dxcalc	dtcalc,xcalc
dxclock	Front Panel,xclock
dxpaint	dticon/dtstyle,bitmap
dxnotepad	dtpad
dxbook	dthelpview,Netscape
dxcardfiler	N/A
dxsession	xdm,dtsession
dxpause	N/A
dxvdoc	ghostview, Adobe Acrobat
CDA	N/A

You should migrate to the dt* tools/utilities or other options as soon as possible.

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

Adobe has retired their Display PostScript (DPS) product, which includes the client libraries, X Server extension, and various applications and examples. Therefore, DIGITAL UNIX has no option but to also retire the Adobe DPS product. DIGITAL UNIX Version 4.0E will be the last version of DIGITAL UNIX to ship these components. 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 will be retired and will no longer be included with the DIGITAL UNIX operating system after DIGITAL UNIX Version 4.0E.

The following Adobe applications will be retired and will no longer be included with the DIGITAL UNIX operating system:

- `makepsres`
- `pswrap`
- `showps`
- `texteroids`

The following DIGITAL applications, which use Adobe DPS, will be retired and will no longer be included with the DIGITAL UNIX operating system after Version 4.0E:

- `dxvdoc`
- `dxmail`
- `dxbook`
- `dxnotepad` (Asian Language, with the exception of Japanese)

Compaq suggests that in future releases of DIGITAL UNIX customers use the GhostView product on the Freeware CD as a replacement viewer for PostScript documentation.

8.3 Nonconforming Curses Library

DIGITAL UNIX Version 4.0 included a new X/Open-Compliant Internationalized Curses library. This library was not binary compatible with previous versions of the DIGITAL UNIX Curses library, so compatible binaries (named `libcurses.a` and `libcurses.so`) were shipped in separate directories for Version 4.0. In that same release, DIGITAL announced the intent to put the compatible binaries into an obsolete subset and subsequently remove them in future releases.

These compatible binaries have been moved into the OSFOBSOLETE425 subset for Version 4.0D and will be removed in the next major release of DIGITAL UNIX.

8.4 Previous C Compiler

The `acc` compiler is being moved to the `CMPDEVENH` subset in DIGITAL UNIX Version 4.0E. This release note states that and informs the user that it is no longer distributed in the default subset.

The C compiler for DIGITAL UNIX has been officially replaced by DEC C for DIGITAL UNIX. The older compiler is no longer distributed in the default subset and will not be installed during a default system installation.

The older compiler was the default compiler for DIGITAL UNIX Version 3.2 and earlier. It is still available in this version by installing it from the CMPDEVENH subset and invoking it with the `cc -oldc` compiler command.

For more information, see the `cc-oldc`(8) and `cc`(1) reference pages or the *Developers' Toolkit for DIGITAL UNIX Software Product Description*.

8.5 The dbx Debugger

The `dbx` symbolic debugger will be retired in a future release of DIGITAL UNIX. The `dbx` debugger will be replaced by the DIGITAL `ladebug` debugger, which is a superset of the `dbx` functionality. The DIGITAL `ladebug` debugger is command line compatible with `dbx` and also supports a graphical user interface.

DIGITAL recommends that you begin using the DIGITAL `ladebug` debugger now and report any problems. This will enable DIGITAL to provide a higher quality replacement when `dbx` is finally retired.

8.6 DEC C Compiler Default Change from -std0 to -std

The default language mode for the DIGITAL UNIX C compiler will change from `-std0` to `-std`. This will occur in a future release.

When the change occurs, it will be possible to revert back to `-std0` by specifying `-std0` on the compilation command line, by adding `-std0` to the `/usr/ccs/lib/cmplrs/cc/comp.config` file, or by using an environment variable.

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

To provide compliance with several UNIX and Internet standards, the `struct utmp`, `struct utmpx`, and `struct lastlog` structures will be changed in the next major release (Version 5.0) of DIGITAL UNIX. These changes will affect the `/usr/include/utmp.h`, `/usr/include/utmpx.h`, and `/usr/include/lastlog.h` files :

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

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

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

The programs will enable you to convert your existing `/var/adm/wtmp` and `/var/adm/lastlog` files to the new format or convert new format files to the old format for use by existing programs. See the corresponding reference pages for more information.

8.8 C Language long double Type Changing to 128 bits

In a future release of DIGITAL UNIX, the default size of the C language `long double` type will change from 64 bits to 128 bits. This will allow application writers to perform mathematical calculations on numbers much larger in magnitude with more precision than possible with the current `long double` type, which is treated identically to the `double` type. A compiler option will be provided to allow existing source code that expects a 64-bit `long double` type to continue to be compiled and executed.

The one binary incompatability that an existing application (if linked using the `-call_shared` switch) could experience with the new default is related to the input and output of `long double` types. Currently, the `printf` and `scanf` functions, and other associated functions, interpret the format code `%Lf` (capital-L followed by f) as referring to a 64-bit `long`

`double` type. In a future release, the interpretation of this format code will be changed to expect the new 128-bit data type. Programs that use this format code will either need to be changed, or will need to be run with the new compatibility library that will be provided. An extra step will be necessary to cause the application to use this library, and will be documented in the release in which the change actually takes affect.

8.9 C Library Functions and POSIX P1003.1C

As of DIGITAL UNIX Version 4.0, the following C library functions exist in two versions due to conflicts between previous versions of DIGITAL 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 DIGITAL UNIX Version 4.0E; however, these routines will be retired in a future release of DIGITAL UNIX. The obsolete versions should not be used in new designs. These routines formerly resided in `libc_r.a` and `libc_r.so`, but were merged into the standard C runtime library.

8.10 POSIX 1003.4a (draft 4) pthread Routines in DECthreads

The POSIX 1003.4a, Draft 4 interface of DECthreads will be retired in a future release of DIGITAL UNIX. Applications that were written using the POSIX 1003.4a, Draft 4 API should be migrated to the IEEE Std 1003.1c-1995, POSIX System Application Program Interface provided by DECthreads. The POSIX 1003.1c standard interface is the most portable, efficient, and powerful programming interface offered by DECthreads. A compatibility mode for the draft 4 POSIX 1003.4a API has been provided to help ease migration. This compatibility mode will be removed in a future release.

8.11 DECthreads CMA Interface

The CMA interface of DECthreads is obsolete beginning with this release of DIGITAL UNIX. Obsolescence means that while the CMA API will continue to exist in DIGITAL UNIX and will be supported, CMA routines will no longer be documented or enhanced. DIGITAL recommends that you port your CMA based application to the IEEE Std 1003-1c-1995, POSIX System Application Program Interface provided by DECthreads.

8.12 Asynchronous I/O Binary Compatibility

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

In the next major release of the operating system, support for applications built under Version 3.2x using asynchronous I/O will be discontinued. These applications will need to be recompiled and relinked in order to run properly under DIGITAL UNIX.

8.13 Nemacs

Nemacs Version 3.3.2, a public domain Japanese implementation of `emacs`, will be retired in a future release of DIGITAL UNIX. `Mule`, a public domain multilingual implementation of `emacs`, will be carried forward as the replacement functionality for Nemacs. The Nemacs subsets `IOSJPNEMACS425` and `IOSJPNEMACSSRC425` will be removed from the system. For more information on `Mule`, refer to the `mule`(1) reference page.

8.14 Berkeley Software Distribution TTY-NAME

The intent to retire the BSD `TTY-NAME` namespace was announced in DEC OSF/1 Version 3.0. This functionality will be retired in a future release.

8.15 SCSI Device Names

Support for `rz` SCSI device names will be retired in a future release of DIGITAL UNIX. Any code that derives knowledge about a device from the ASCII name or minor number may be impacted.

All code that uses the current namespace will be compatible when the change occurs because a mechanism that ensures binary compatability will be provided. Existing interfaces such as names and minor numbers will be fully supported.

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

The `-x` and `-p` options of the `addvol` and `mkfdmn` commands allow you to set the per-volume bitfile metadata table (BMT) when you create a new volume or file domain. Users typically set the BMT to prevent an "extent exhaustion problem" that is common in the field.

In Version 4.0D, DIGITAL has removed limitations in the operating system that caused the "extent exhaustion problem," hence the `-x` and `-p` options

are no longer needed and will be retired in a future major release.

8.17 LSM Block Change Logging (BCL)

In the next major release of DIGITAL UNIX, the Logical Storage Manager Block Change Logging (BCL) feature will be retired and replaced by a new logging method. This new logging method, called Dirty Region Logging (DRL), will log regions instead of blocks for writes to LSM mirrored volumes. For most environments, DRL will provide the same ability to quickly resynchronize mirrors after a failure as BCL, but with considerably less logging overhead.

The logging format and configuration for DRL will not be compatible with BCL; for example, DRL will require a log size of at least two blocks. The next major release of DIGITAL UNIX will automatically reconfigure volumes for DRL when a BCL mirrored volume has a log of two or more blocks. Mirrored volume configurations with a BCL of one block will require manual reconfiguration to continue to take advantage of logging for faster mirror recovery.

Due to these changes, DIGITAL recommends configuring a BCL size of two or more blocks to simplify BCL to DRL migration in the future. For optimum DRL configurations, use the following guidelines when configuring volumes with BCL: use a BCL subdisk block size of one block plus an additional block for every 1 GB of volume storage, then round up the BCL size to the next even number. Table 8-2 shows some examples.

Table 8-2: BCL Configuration Examples

Volume Size (GB)	BCL Subdisk Size (Blocks)
<1	2
1	2
2	4
3	4
4	6
5	6
6	8
7	8
8	10
9	10
10	12

8.18 LSM volassist Command Syntax

In the next major release of DIGITAL UNIX, the syntax of the `volassist` command will be changing. It will no longer support the following constructs:

- `[!]medianame,offset`

  You will still be able to use `[!]medianame,offset`, but you will not be able to specify an offset. Users requiring the ability to specify an offset will need to use the low-level commands to create subdisks, plexes, and volumes exactly as required.

- `alloc=size`

  A new construct, `alloc=storage-spec[,storage-spec,...]`, will replace `alloc=size`. However, the new construct will not allow you to specify sizes for all allocations. You will need to use the low-level commands to create subdisks, plexes, and volumes exactly as required.

- `align=size`

  Two new constructs, `diskalign` and `nodiskalign`, will replace `align=size`, allowing you to specify whether subdisks should be created on cylinder boundaries. If you require the ability to specify alignments for all allocations, you will need to use the low-level commands to create subdisks, plexes, and volumes exactly as required.

8.19 LSM volprint Command Format

In the next major release of DIGITAL UNIX, the default output format of the `volprint` command will be changing.

Invoking `volprint` with no options currently displays all objects in `rootdg`, starting with the `rootdg` disk group record, followed by all of the disk-media records, subdisk records, plex records, and finally volume records. In a future release, invoking `volprint` with no options will display all records for all disk groups, with all objects arranged in a hierarchical fashion.

Invoking `volprint` with an object-type option (`v`, `p`, or `s`) currently displays all objects of that type in the specific disk group. In a future release, invoking `volprint` with any of these options will display all objects of that type, as well as all objects of all subsidiary types, in all disk groups.

By default, invoking `volprint` currently displays the record type, name, association, kernel state, length, and comment field. In a future release, it will display the record type, name, association, kernel state, length, plex offset, state, and the `tuti10` and `puti10` fields.

Users who use `volprint` in scripts should use the `-F` option to define the exact output format that they require.

8.20 LVM-to-LSM Migration Tools

The LVM-to-LSM Migration Tool was provided with DIGITAL UNIX Version 4.0 to enable migration from the LVM interfaces that were retired in that release to DIGITAL UNIX Logical Storage Manager volumes. This tool will become obsolete in later releases of DIGITAL UNIX because most customers will have migrated to LSM volumes by that time; therefore, DIGITAL plans to retire the LVM-to-LSM Migration Tool in a future release of the operating system.

There are no plans to retire UFS or the AdvFS Migration Tools at this time.

8.21 OSF/Motif Version 1.1.3

The Motif Version 1.1.3 libraries have been provided as run-time services for compatibility with applications that have not yet converted to Motif 1.2. Development support was retired in DEC OSF/1 Version 2.0.

In DIGITAL UNIX Version 4.0 the Motif 1.1.3 libraries were moved to an optional subset. Applications that require the libraries will see an error from the loader and you must install the optional subset. This optional subset will be removed from the product in a future release.

8.22 XIE Version 3.0 X Client Extension

DIGITAL UNIX Version 4.0E supports XIE Version 5.0. Support for XIE Version 3.0 Server extensions was removed in DIGITAL UNIX Version 4.0, but Client support will not be removed until a later release of DIGITAL UNIX.

8.23 Microsoft Sound Board Driver

In DIGITAL UNIX Version 4.0, the device driver for the base audio on the DIGITAL AlphaStations and DIGITAL AlphaServers was removed from the base operating system. This device driver supported the Microsoft Sound Board (MSB), the AlphaStation Sound Card, and the built-in audio hardware shipped with certain AlphaStation systems.

The driver binaries are now available as part of the Multimedia Services for DIGITAL UNIX kit on the *DIGITAL UNIX 4.0E Associated Products, Volume 1* CD–ROM in the `MMEDRVMSB241` subset.

The following files will be removed from the base operating system:

- `/usr/sys/BINARY/msb.o`
- `/usr/sys/include/io/dec/eisa/msb.h`
- `/usr/sys/include/io/dec/eisa/msb_reg.h`

You can also get support for this device from the Multimedia Services for DIGITAL UNIX kit that is located on the Software Products Library CD-ROM. Support is also factory-installed on all DIGITAL AlphaStation DIGITAL UNIX packaged systems. The license for this product is bundled with DIGITAL AlphaStations so that you can use it at no additional cost.

8.24 Graph Utility

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

8.25 The atmsetup Script

The `atmsetup` script is a new utility with DIGITAL UNIX Version 4.0D that allows you to set up or change the current configuration of ATM on your system.

This script was designed to be an interim solution to simplify the setup process for ATM. In future releases, it will be supplemented and finally replaced by an application in the SysMan suite, with a full graphical user interface.

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

8.26 Remote Prestoserve Support

In the next major release of DIGITAL UNIX, PrestoServe support for remote operations will be retired. This means that the use of the `-h` option of the `presto` command, the `dxpresto` *`hostname`* variable, and the `-n` option of the `prestoctl_svc` command will no longer be supported.

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

8.27 Disk Size Requirement for Installation

The minimum disk size requirement for installing the DIGITAL UNIX operating system is changing in the next major release due to additional features and services that are being added. The minimum disk size required to support a single-disk operating system installation will be 1GB (such as an

RZ26) for both default and custom installation types.

8.28 Installupdate -i Option

The `-i` option to the `/sbin/installupdate` command will be retired in a future release of the operating system.

The `-i` option currently allows you to interactively select kernel components once the new software subsets have been installed. Starting with the next major release, this flag will be unnecessary because you will be able to interactively select optional kernel components at the beginning of the update installation process, prior to software installation. These kernel components will be built into the kernel automatically during the kernel build phase at the end of the update installation; therefore, you need not be present at that time.

8.29 Different awk Versions

DIGITAL UNIX contains three different versions of the `awk` command:

- `oawk` - Old `awk` is the original `awk` command.
- `gawk` - GNU awk is the Free Software Foundation's version.
- `nawk` - Is the new version of the `awk` command and is current linked to `awk`. This version of the `awk` command is the XPG4-compliant version of the `awk` command.

The `oawk` and `gawk` versions of the `awk` command will be removed in a future version of DIGITAL UNIX. Also, the link between `awk` and `nawk` will be removed, leaving `awk` and removing `nawk`. Users should ensure that their scripts use `/usr/bin/awk` instead of any other version of the `awk` command currently existing on the system.

8.30 secsetup Script

The `/usr/sbin/secsetup` script, used to configure Enhanced Security authentication support, will retire in a future release of DIGITAL UNIX. It will be replaced by `sec_config` , an application in the SysMan suite that provides a full graphical user interface.

8.31 Open3D TURBOChannel and PVpci Device Support

Support for the following Open3D TURBOchannel devices will be retired in the nextrelease of DIGITAL UNIX:

- HX+

- PV-L
- PV-M

These devices are also known as:

- ZLX-E1, E2, E3
- ZLX-L1, L2
- ZLX-M1, M2

Digital Open3D V4.4 is the last version of Open 3D that supports these graphic adapters. The DDX and device drivers will ship in future releases of DIGITAL UNIX, but support will be retired and Open3D will no longer respond to problem reports. DIGITAL UNIX 2D support will continue for future DIGITAL UNIX V4.0n releases and the initial release of DIGITAL UNIX Version 5.0 for the TX, HX, and HX+ graphic adapters.

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

In addition, the Open3D PVpci (ZLXp-L1 and ZLXp-L2) graphic adapter support will end in December 1998. The last version of Open3D to support these cards is Open 3D Version 4.9. The last version of DIGITAL UNIX to support these adapters will be DIGITAL UNIX Version 4.0E. The replacement functionality is the Intergraph Cateyes (Powerstorm 4Dxx) graphic devices.

8.32 System Management Utilities

The following setup scripts will be retired in a future release:

- `bindsetup`
- `mailsetup`
- `netsetup`
- `nfssetup`
- `ntpsetup`

These scripts have been replaced by tasks that you can perform by running the SysMan menu and navigating to the desired item or directly from the command line by using the appropriate menu accelerator. The following table maps the replacement functionality and accelerator.

Retired Script	SysMan Menu Item	SysMan Accelerator
bindsetup	Configure system as a DNS client	dns_client
	Configure system as a DNS server	dns_server
mailsetup	Configure the host's mail system	mailsetup
netsetup	Network Configuration Step-by-Step	net_step
nfssetup	Configure System as an NFS client	nfs_client
	Configure System as an NFS server	nfs_server
ntpsetup	Configure system as an NTP client	ntp_config

For more information, see the reference pages or the *Network Administration* manual.

8.33 Changed Protected Password Database Format

The enhanced security user authentication profiles, originally stored in the protected password files /tcb/files/auth/<a..z>/<username>, are now stored by default in a database format. The convauth program, run for an update installation, will convert existing files into the new format. Support for the original format will be removed in a future release of DIGITAL UNIX.

8.34 Audit_setup Script

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

8.35 VM Subsystem ubc_nfsloopback Variable Is Obsolete

The /etc/sysconfigtab ubc_nfsloopback Virtual Memory (VM) subsystem kernel configuration variable is unnecessary, and has been removed.

This configuration variable was needed in some TruCluster configurations with previous versions of DIGITAL 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 DIGITAL UNIX Version 4.0E has been modified to prevent the deadlock which removes the need for the ubc-nfsloopback variable.

If the variable is not removed from `/etc/sysconfigtab`, it will not cause a problem. However, after the kernel is loaded by the firmware, the console generates the following message to indicate that the variable is still present in `/etc/sysconfigtab:`

```
sysconfigtab:  attribute ubc_nfsloopback not in subsystem VM
```

8.36 libsecurity.a

`libsecurity.a` is the archive form of the enhanced security library. This archive form of `libsecurity` will no longer ship. Customers wishing to link against `libsecurity` should instead use the shareable form, `libsecurity.so`.

Maximum System Limits A

A.1 System Limits

This appendix lists the system limits for the major components of this release. For hardware information specific to your individual processor, see the DIGITAL UNIX *Software Product Description* (SPD) and the *DIGITAL Systems & Options Catalog*.

You may be able to increase some of the system limits by changing kernel attribute values. Use the `sysconfig` command to display the current attribute values, and the maximum and minimum values. For information on how to modify attributes, see the *System Administration* manual and the *System Configuration and Tuning* manual, which also includes lists of tunable attributes.

The following sections describe the system limits.

A.1.1 Installation Requirements

A default installation of the DIGITAL UNIX operating system requires at least 525 MB of disk space (for example, a RZ25 disk). This limit will change in a future release (see Section 8.27 for more information). The disk space that a custom installation requires depends on the optional subsets that you will install. See Appendix B for information about subset sizes.

A.1.2 Memory Limits

The memory limits are as follows:

- Physical memory requirements

  The minimum amount of physical memory that DIGITAL UNIX requires is 24 MB, or 32 MB if you are using the Advanced File System (AdvFS). However, depending on the configuration and workload, you may need at least 64 MB for optimal performance.

  The system platform limits the maximum amount of physical memory. For more information on platform memory support, see the *DIGITAL Systems & Options Catalog*.

- Virtual address space limit

  The default amount of virtual address space available for each process is 1 GB. However, in many cases, available swap space may be exhausted before this limit is reached.

  You can increase the available virtual address space to a maximum of 4 TB by modifying the value of the `vm-maxvas` attribute in the `vm` subsystem. For more information about setting this attribute, see the *System Configuration and Tuning* guide.

- Page size

  The physical page size is 8 KB and cannot be changed. The page size is hardware dependent and is set by the console at boot time.

A.1.3 Process Limits

The process limits for the system and for users are as follows:

- Per system

  The `maxusers` attribute in the `proc` subsystem specifies the number of simultaneous users that a system can support without straining system resources. System algorithms use the `maxusers` attribute to size various system data structures, and to determine the amount of space allocated to system tables, such as the system process table, which is used to determine how many active processes can be running at one time.

  The default value assigned to the `maxusers` attribute depends on the size of your system. The default value is 32 on systems with at least 32 MB of physical memory. The minimum value is 8 and the maximum value is 4096.

  The maximum number of tasks that can run simultaneously on a system is determined by the value of the `task-max` attribute in the `proc` subsystem. The default value is 8213. The minimum value is 85, and the maximum value is 32768.

  The maximum number of kernel threads that can run simultaneously on a system is determined by the value of the `thread-max` attribute in the `proc` subsystem. The default value is 16424. The minimum value is 160, and the maximum value is 262136.

- Per user

  The maximum number of processes that a user can create is determined by the value of the `max-proc-per-user` attribute in the `proc` subsystem. The default value is 64. The minimum value is 0, and the maximum value is 32767.

  The maximum number of threads that a user can create is determined by

the value of the `max-threads-per-user` attribute in the `proc` subsystem. The default value is 256. The minimum value is 0, and the maximum value is 18,446,744,073,709,551,615 ($2^{64} - 1$).

A.1.4 Device Addressing Limits

The limits for device addressing are as follows:

- Device access

  There are two types of disk device access: raw (character) and buffered (block).

  For raw or character access, the `uio.uio_offset` structure field describes the byte offset within the disk partition. In this release, the `uio_offset` is an unsigned 64-bit value, allowing an offset up to 2^{64} or 18 exabytes. This value is converted to a physical block/sector number, which is the data transfer start position. The physical block/sector number is limited by the `buf.b_blkno` structure field.

  For buffered or block access, the `buf.b_blkno` structure field describes the block/sector offset within the disk partition and is a signed 32-bit value. Because this release supports a fixed 512-byte block/sector size, as defined by `DEV_BSIZE`, the offset is limited to 1 TB.

- Major-minor numbers (`dev_t`)

  Devices are identified by a major-minor pair of numbers, where the major number specifies the device driver and the minor number identifies the device. In this release, this pairing is represented by a 32-bit value described by the type `dev_t`. The major number portion of `dev_t` consists of bits 20 to 31 (a total of 12 bits). Because each device driver requires 12 bits for its major number, you can configure only 4096 device drivers into the system.

  The minor number portion of `dev_t` consists of bits 0 to 19 (a total of 20 bits). The device driver is responsible for interpreting these bits. A device driver that utilizes all 20 bits for device addressing can address up to 1048576 devices for each major number. For device drivers that support disk devices, some of the bits in the minor device number represents the partition number. This release requires disk drivers to reserve the lower 6 bits for device attributes and partition numbers, and only supports eight partitions.

- SCSI/CAM addressing

  Common Access Method (CAM) is an ANSI-proposed standard for a common software interface to the Small Computer Systems Interface (SCSI). There are no restrictions or limitations within CAM for disk block addressing, because the address is an incoming value.

For SCSI-2, the Command Descriptor Block (CDB) defines the starting disk block number for the transfer. In this release, the 10-byte CDB reserves 4 bytes for the disk block address. This is an unsigned 32-bit value that allows $2^{32} - 1$ or 4 gigasectors of addressing. Using a 512-byte block/sector size, this value corresponds to 2 TB.

In this release, the SCSI/CAM driver can address a maximum of 64 buses, with up to 7 device targets for each bus, and a maximum of eight logical unit numbers (LUNs) for each device target. According to these limitations, SCSI/CAM can address a maximum of 3584 devices.

- Redundant Array of Independent Disks (RAID)

  Refer to the following for complete information:

  – *DIGITAL Systems & Options Catalog*

  – Product description literature

  – Web pages:

  ```
  http://www.digital.com/products.html
  http://www.storage.digital.com
  ```

- Disklabel command

  The `disklabel` command specifies the partitions of a disk and their starting block/sector number. The starting block/sector number of a partition is defined by the `partition.p_offset` structure field, which is an unsigned 32-bit value that supports up to 2 TB of addressing, using a 512-byte block/sector size.

A.1.5 Device Limits

DIGITAL UNIX supports a maximum of seven devices per SCSI bus (for example, one host bus adapter and seven disks). The maximum number of SCSI buses supported by a system depends on the platform, but the DIGITAL UNIX operating system supports a maximum of 64 SCSI buses.

For all other device limits, see the *DIGITAL Systems & Options Catalog*.

A.1.6 Logical Storage Manager Limits

In this release, the Logical Storage Manager (LSM), supports a maximum of 768 disk groups and 256 disks, either in a disk group or across the system.

The LSM term **volume** refers to a virtual disk that represents an addressable range of disk blocks. File system data or raw I/O can be directed to a volume. This release supports a maximum of 1 TB of disk space in a disk group or on a system, and a maximum volume size of 1 TB. The maximum number of supported LSM volumes is 4093 for all disk groups in a system, which includes 4091 nonsystem volumes and 2 system (root or swap)

volumes.

The LSM term **plex** refers to a physical disk or a set of disks that contain a complete copy of a volume's data. A mirrored volume consists of at least two plexes. In this release, the maximum number of supported plexes per volume is 8, and the maximum number of supported plexes per system is 4093 (or 4091, if root and swap volumes are not used).

The LSM term **subdisk** refers to a contiguous portion of a physical disk that can be striped or concatenated with other subdisks to form a plex. A maximum of 4096 subdisks can be associated with one plex, and DIGITAL supports 4096 subdisks in each disk group or on a system.

LSM object names (for example, volumes, plexes, subdisks, and disk groups), volume attribute names (such as user and group), and `dxlsm view` names are limited to 14 characters.

A.1.7 File System Limits

The file system limits are as follows:

- Advanced File System

  The Advanced File System (AdvFS) term **volume** refers to a single logical device, such as a disk, disk partition, or a logical volume. A **file domain** is a named set of one or more volumes. A **fileset** is a named and mountable logical file structure that is created in a domain. An **active fileset** is a fileset that is currently mounted.

  AdvFS supports a maximum fileset and file size of 16 TB minus 512 K, a maximum of 100 active file domains for each system, and a maximum of 256 volumes for each domain. However, because a single disk failure in a multivolume file domain can make the entire domain inaccessible, DIGITAL recommends that you use a maximum of eight volumes in a file domain.

  Although DIGITAL UNIX supports an unlimited number of filesets per system, only 512 filesets can be mounted at one time. The maximum number of files in a fileset is 2^{32} and is limited by the tag that is used to uniquely identify a file in a fileset. Because of a sequence number limit, a tag can be used only 4096 times; therefore, the actual limit on files in a fileset decreases over time.

  Although AdvFS can support a page size larger than 13 bits, the maximum size of an AdvFS file and fileset is 16 TB - 512 K ($2^{13} * 2^{31}$), with a 13-bit page size and a 31-bit page number.

  AdvFS supports a maximum of 512 mounted filesets. However, each active file domain has a hidden mounted fileset that must be counted when determining the total number of mounted filesets. For example, if you have an active file domain with two mounted filesets, the file domain

actually has three mounted filesets.

- UNIX File System

  In this release, the UNIX File System (UFS) file size is limited by the amount of space that can be addressed by the kernel `buf` structure. The `buf.b_blkno` structure field, defined as `daddr_t`, is a 32-bit signed value, and specifies the block/sector offset within a disk partition. The `DEV_BSIZE`, block or sector size, is 512 bytes. Theoretically, a UFS file system could be 1 TB ($2^{13} * 2^{9}$); however, DIGITAL UNIX supports only 128 GB. The maximum LSM logical volume size also limits a UFS file system and file size.

  DIGITAL UNIX supports up to 2,147,483,647 UNIX File System (UFS) and Memory File System (MFS) mounts. The `max-ufs-mounts` attribute controls the maximum number of UFS and MFS mounts. The default value is 1000.

- CD–ROM File System

  The size of CD–ROM File System (CDFS) files and file systems is limited by the CD–ROM media in which they reside. Currently, the CD–ROM media supports approximately 600 MB. However, DIGITAL UNIX will be able to support larger CD–ROMs if they become available.

  This release supports a maximum of 512 CDFS mounts.

- Sparse files

  DIGITAL UNIX supports **sparse files** on AdvFS and UFS; therefore, the size of a file can exceed the size of the file system in which it resides. The maximum sizes for sparse files are as follows:

  – AdvFS – $2^{44} - 512K$

  – UFS – $2^{44} - 8$ K

- Network File System

  In this release, the theoretical maximum sizes of files that are accessible through Network File System (NFS) Version 2 and Version 3 are as follows:

  – NFS Version 2 – (2 GB – 1 byte)

  – NFS Version 3 – 18 exabytes ($2^{64} - 1$)

  However, DIGITAL UNIX supports the following maximum file sizes:

  – NFS Version 2 – 2 GB

  – NFS Version 3 – 512 GB

  In addition, an NFS file system is always limited by the size of the local file system that is being exported.

  An NFS Version 2 or Version 3 client can mount a maximum of 2048

files or directories.

- Memory mapped file limit

  The maximum supported size of a file that can be mapped into memory without segmenting the file depends on the virtual address space limits, as documented in the section on memory limits.

- Open files limit

  The `open-max-soft` and `open-max-hard` attributes control the maximum number of open file descriptors for all processes. The default values are 4096. When the `open-max-soft` limit is reached, a warning message is issued, and when the `open-max-hard` limit is reached, the process is stopped. The maximum number of open files per process is 65,536. If you increase the maximum number of open files per process, make sure that you adjust the value of the `max-vnodes` attribute.

  See Appendix E for information about increasing the open file descriptor limit. In addition, the maximum number of open files can be set, on a per-process basis, between 64 and 4096 by using the `setrlimit` function, or up to 65,536 by following the steps documented in Appendix E.

  File descriptor entries in the per process file table are dynamically allocated after the initial 64 entries in the `utask` structure are used.

- File locking limits

  The DIGITAL UNIX file record locking service allows applications to lock any number of bytes in a file in the range of 0 to $2^{63} - 1$, inclusive. File locking is supported by UFS, AdvFS, and both NFS Version 2 and Version 3. Because the NFS Version 2 protocol suite allows ranges to be specified only with 32-bit numbers, it supports a file locking range of 0 to $2^{31} - 1$, inclusive.

- Pathname limits

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

A.1.8 Networking Limits

The networking limits are as follows:

- Pseudoterminals (`ptys`)

  The maximum number of supported `ptys` is 8192.

- LAT connections

  - The maximum number of incoming LAT connections is 4800.
  - The maximum number of outgoing LAT connections is 4000.

- IP alias addresses

  DIGITAL UNIX allows the use of up to 5,000 IP alias addresses before system performance begins to degrade.

- Packetfilter limits

  The packetfilter pseudo-driver can support up to 255 simultaneous open filters (each filter is usually mapped to one instance of an application program). The packetfilter can support a maximum of 255 devices. Use the `pfconfig` command to configure packet filters.

- Network transfer rates

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

A.1.9 Limits for Backup Utilities

The limits for the backup utilities are as follows:

- `cpio`

Files per archive:	No limit
Files per file system:	No limit
File size:	4 GB
File name size:	256 bytes

- `dd`

Files per archive:	Not used
Files per file system:	Not used
File size:	4 GB
File name size:	Not used

- `dump`

Files per archive:	4,000,000,000
Files per file system:	4,000,000,000
File size:	4 GB
File name size:	No limit (part of the inode data)

- `tar`

Files per archive:	No limit
Files per file system:	No limit
File size:	8 GB (4 TB with extended option)
File name size:	1024 bytes (with prefix)

- `pax`

Files per archive:	No limit
Files per file system:	No limit
File size:	8 GB (4 TB with extended option)
File name size:	1024 bytes (with prefix)

Software Subset Information B

This appendix provides the sizes of all DIGITAL UNIX software subsets for full, update, and RIS installations.

B.1 Disk Space Required for Software Subsets

Table B-1, Table B-2, and Table B-3 show disk space as the number of 512-byte blocks required in the `root`, `/usr`, and `/var` file systems to install each DIGITAL UNIX software subset. The figures for each group of files within a subset have been rounded up to the next higher 512-byte increment; this means that the total space requirements listed are slightly greater than the space actually required.

To determine the subset size in megabytes (MB), divide the size in blocks by 2048.

For information on the contents of each subset, refer to the *Installation Guide*. If you want to add optional subsets after you install DIGITAL UNIX Version 4.0E, use the `df` command to determine free disk space in blocks.

Table B-1: DIGITAL UNIX V4.0E Operating System, Volume 1

DIGITAL UNIX V4.0E Operating System				
Subset	**root**	**/usr**	**/var**	**Total**
OSFACCT435	9.67	1077.52	88.03	1175.22
OSFADVFS435	4839.00	6421.66	—	11260.66
OSFADVFSBIN435	2145.93	3.06	—	2148.99
OSFADVFSBINOBJECT435	—	3393.74	—	3393.74
OSFAFM435	—	2160.73	—	2160.73
OSFATMBASE435	225.62	719.84	—	945.46
OSFATMBIN435	3289.40	15.28	—	3304.68
OSFATMBINCOM435	—	391.59	—	391.59
OSFATMBINOBJECT435	—	10912.33	—	10912.33
OSFBASE435	23413.95	79928.47	838.22	104180.64
OSFBIN435	12621.21	1508.99	—	14130.19
OSFBINCOM435	48.50	27872.14	67.88	27988.51
OSFBINOBJECT435	—	19230.06	—	19230.06
OSFC2SEC435	861.78	1354.16	84.00	2299.94
OSFCDAPGMR435	—	3301.75	—	3301.75

Table B-1: (continued)

DIGITAL UNIX V4.0E Operating System

Subset	root	/usr	/var	Total
OSFCDEAPPS435	—	11183.10	—	11183.10
OSFCDEDEV435	—	25559.38	—	25559.38
OSFCDEDT435	—	52907.88	—	52907.88
OSFCDEMAIL435	—	3002.43	—	3002.43
OSFCDEMANOP435	—	2055.95	—	2055.95
OSFCDEMANOS435	—	1575.70	—	1575.70
OSFCDEMIN435	—	13527.87	76.00	13603.87
OSFCLINET435	779.45	13838.39	41.00	14658.84
OSFCMPLRS435	—	24662.76	—	24662.76
OSFCTABLOC435	34.25	359.30	—	393.55
OSFDCMT435	—	1107.81	—	1107.81
OSFDCMTEXT435	—	4368.19	—	4368.19
OSFDECW435	—	1781.11	47.20	1828.31
OSFDMS435	—	87.26	73.00	160.26
OSFDOSTOOLS435	—	1005.00	—	1005.00
OSFDPSFONT435	—	5161.47	—	5161.47
OSFEMACS435	—	39367.50	—	39367.50
OSFENVMON435	21.25	152.62	—	173.87
OSFEURLOC435	—	1811.22	—	1811.22
OSFEXAMPLES435	—	1993.66	—	1993.66
OSFEXER435	—	3851.00	—	3851.00
OSFFONT15435	—	3160.99	—	3160.99
OSFFONT435	—	2432.85	—	2432.85
OSFHWBASE435	22230.49	3017.04	34.40	25281.93
OSFHWBIN435	21572.94	2104.07	—	23677.01
OSFHWBINCOM435	—	3273.34	—	3273.34
OSFHWBINOBJECT435	—	20420.69	—	20420.69
OSFINCLUDE435	—	5229.00	—	5229.00
OSFINET435	543.84	6773.52	538.73	7856.09
OSFJAVA435	—	21532.28	—	21532.28
OSFJAVADEV435	—	17340.34	—	17340.34
OSFJAVADOC435	—	24637.19	—	24637.19
OSFKBDLK201435	—	361.70	—	361.70
OSFKBDLK401435	—	248.44	—	248.44
OSFKBDLK411435	—	134.33	—	134.33
OSFKBDLK421435	—	16.42	—	16.42
OSFKBDLK444435	—	126.52	—	126.52
OSFKBDPCXAL435	—	273.12	—	273.12
OSFKTOOLS435	—	1468.75	7125.71	8594.46
OSFLAT435	513.07	907.71	7.82	1428.60
OSFLDBBASE435	—	17732.85	—	17732.85
OSFLDBDOC435	—	433.16	—	433.16
OSFLDBGUI435	—	10794.76	—	10794.76
OSFLDBSRV435	—	83.00	—	83.00

Table B-1: (continued)

DIGITAL UNIX V4.0E Operating System

Subset	root	/usr	/var	Total
OSFLEARN435	—	3097.64	—	3097.64
OSFLIBA435	—	6356.14	—	6356.14
OSFLSMBASE435	5217.42	3022.58	48.21	8288.21
OSFLSMBIN435	535.62	3.05	—	538.68
OSFLSMBINCOM435	—	461.96	—	461.96
OSFLSMBINOBJECT435	—	547.61	—	547.61
OSFLSMX11435	—	2251.79	47.90	2299.69
OSFLVM435	—	2917.42	—	2917.42
OSFMANOP435	—	16857.00	—	16857.00
OSFMANOS435	—	14042.53	—	14042.53
OSFMANWOP435	—	8883.66	—	8883.66
OSFMANWOS435	—	1510.91	—	1510.91
OSFMH435	—	4352.17	—	4352.17
OSFMITFONT435	—	18937.17	104.01	19041.18
OSFMOTIF11435	—	12847.09	—	12847.09
OSFNETCONF435	—	2307.91	—	2307.91
OSFNETSCAPE435	—	41196.22	—	41196.22
OSFNFS435	50.21	1180.45	—	1230.67
OSFNFSCONF435	—	337.90	—	337.90
OSFOBSOLETE435	—	3285.68	—	3285.68
OSFOLDDECW435	—	14945.35	—	14945.35
OSFOLDX11435	—	1394.37	—	1394.37
OSFPGMR435	—	8756.37	—	8756.37
OSFPRINT435	104.45	8371.66	19.00	8495.11
OSFRCS435	—	1853.84	—	1853.84
OSFRIS435	—	155.41	143.00	298.41
OSFSCCS435	—	9352.66	—	9352.66
OSFSDE435	—	14121.76	—	14121.76
OSFSDECDE435	—	284.69	—	284.69
OSFSER3D435	—	12130.00	—	12130.00
OSFSER435	—	19509.53	66.32	19575.85
OSFSERPC435	—	3947.00	—	3947.00
OSFSERTC435	—	678.00	—	678.00
OSFSVID2435	31.59	615.86	—	647.45
OSFSYSMAN435	8.56	13642.45	—	13651.01
OSFTCLBASE435	—	3415.12	—	3415.12
OSFTERM435	—	3685.73	—	3685.73
OSFTKBASE435	—	5433.77	—	5433.77
OSFUUCP435	101.73	11096.01	266.00	11463.74
OSFX11435	33.40	47039.72	698.13	47771.25
OSFXADMIN435	—	6778.57	73.19	6851.76
OSFXC2SEC435	—	1272.02	—	1272.02
OSFXCDADEV435	—	522.70	—	522.70
OSFXDEMOS435	—	2224.21	—	2224.21

Table B-1: (continued)

DIGITAL UNIX V4.0E Operating System				
Subset	**root**	**/usr**	**/var**	**Total**
OSFXDEV435	—	3902.30	—	3902.30
OSFXEXAMPLES435	—	9196.08	—	9196.08
OSFXIEDOC435	—	1478.88	—	1478.88
OSFXINCLUDE435	—	8342.31	—	8342.31
OSFXLIBA435	—	17145.47	—	17145.47
OSFXMAIL435	—	1186.55	—	1186.55
OSFXMIT435	—	9292.85	—	9292.85
OSFXNEST435	—	390.00	10.11	400.11
OSFXOEM435	—	0.00	965.87	965.87
OSFXPRINT435	—	336.39	—	336.39
OSFXSYSMAN435	—	13549.34	152.74	13702.08
OSFXVFB435	—	246.00	10.11	256.11
Total	99233.32	910842.81	11626.58	1021702.72

Grand Totals	**root**	**/usr**	**/var**	**Total**
	99233.32	910842.81	11626.58	1021702.72

Table B-2: DIGITAL UNIX V4.0E Associated Products, Volume 1

Alternative Development Environment Tools for DIGITAL UNIX				
Subset	**root**	**/usr**	**/var**	**Total**
CMPDEVALT435	14.29	13007.89	—	13022.18
CMPDEVENH435	—	12160.09	—	12160.09
Total	14.29	25167.98	—	25182.26

Digital UNIX V4.0 CDE Instructional Video				
Subset	**root**	**/usr**	**/var**	**Total**
CDEVIDEO100	—	156364.01	—	156364.01
CDEVIDPLR100	—	149.65	—	149.65
Total	—	156513.66	—	156513.66

DEC Ada runtime library for Digital UNIX V4.0D

Subset	root	/usr	/var	Total
ADALIB404	—	1102.86	—	1102.86
Total	—	1102.86	—	1102.86

DEC C++ Class Libraries Version 4.0 for DIGITAL UNIX

Subset	root	/usr	/var	Total
CXLLIBA435	—	224.93	—	224.93
CXLSHRDA435	—	284.12	—	284.12
Total	—	509.04	—	509.04

DIGITAL Decimal RTL V2.4-44 for DIGITAL UNIX

Subset	root	/usr	/var	Total
DCARTL250	—	2746.22	—	2746.22
O2ABASE250	—	2648.56	—	2648.56
Total	—	5394.79	—	5394.79

DECevent

Subset	root	/usr	/var	Total
DIABASE280	17.24	59848.72	28.37	59894.33
DIAKNL281	—	43291.76	—	43291.76
Total	17.24	103140.48	28.37	103186.09

DEC Fortran RTL

Subset	root	/usr	/var	Total
DFARTL369	—	4959.12	—	4959.12
Total	—	4959.12	—	4959.12

DEC Pascal RTL V5.4-18 for Digital UNIX Systems

Subset	root	/usr	/var	Total
DPORTL541	—	1890.15	—	1890.15
Total	—	1890.15	—	1890.15

Sort Runtime Library

Subset	root	/usr	/var	Total
SORLIB300	—	717.85	—	717.85
Total	—	717.85	—	717.85

DECtalk Runtime Kit V4.2A

Subset	root	/usr	/var	Total
DTKDFCRT420	—	18793.48	—	18793.48
DTKEMSRT420	—	1641.53	—	1641.53
DTKRT420	—	5845.04	—	5845.04
DTKRTCDE420	—	252.06	—	252.06
DTKRTDOC420	—	4661.71	—	4661.71
DTKRTRELNOT420	—	220.75	—	220.75
Total	—	31414.58	—	31414.58

INTERSOLV DataDirect

Subset	root	/usr	/var	Total
DAUDOC100	—	0.00	9546.46	9546.46
DAUJDBCODBC100	—	0.00	498.95	498.95
DAUODBCCON100	—	0.00	8303.69	8303.69
DAUSQLNKJAVA100	—	0.00	2539.09	2539.09
DAUSQLNKODBC100	33.32	—	17368.08	17401.40
DAUSQLNKSVR100	—	0.00	41441.10	41441.10
Total	33.32	—	79697.38	79730.70

Free Software Foundation GNU Source for DIGITAL UNIX

Subset	root	/usr	/var	Total
FSFEMACSSRC435	—	55742.88	—	55742.88
FSFGAWKSRC435	—	1220.38	—	1220.38
FSFGZIPSRC435	—	1894.79	—	1894.79
FSFINDENTSRC435	—	1244.68	—	1244.68
FSFMKISOFSSRC435	—	975.94	—	975.94
FSFRCSSRC435	—	1907.47	—	1907.47
Total	—	62986.13	—	62986.13

Multimedia Services V2.4B for DIGITAL UNIX

Subset	root	/usr	/var	Total
MMEDOC242	—	4898.63	—	4898.63
MMEDOCHW242	—	3685.72	—	3685.72
MMEDRVAV201242	75.57	791.66	—	867.23
MMEDRVAV300242	85.09	2121.56	—	2206.66
MMEDRVAV3X1242	70.99	1451.47	—	1522.46
MMEDRVBBA242	—	303.84	—	303.84
MMEDRVMMSESS242	—	532.12	—	532.12
MMEDRVMSB242	72.97	813.45	—	886.42
MMEMANRT242	—	323.40	—	323.40
MMERELNOTES242	—	1370.61	—	1370.61
MMERT242	20.97	21753.40	11.82	21786.19
MMERTCDE242	—	467.70	—	467.70
MMERTSMPLDAT242	—	12385.79	—	12385.79
Total	325.59	50899.35	11.82	51236.75

Netscape FastTrack V3.01 for DIGITAL UNIX

Subset	root	/usr	/var	Total
WEBNETSCAPEFASTTRACK301	—	157725.43	—	157725.43
Total	—	157725.43	—	157725.43

Digital UNIX 4.0 X Window System PanoramiX Extension ADK

Subset	root	/usr	/var	Total
PRXADK200	—	16459.98	—	16459.98
Total	—	16459.98	—	16459.98

Digital Porting Assistant V3.0-0 for Digital UNIX

Subset	root	/usr	/var	Total
PRTBASE300	—	50631.13	—	50631.13
PRTMAN300	—	18.45	—	18.45
Total	—	50649.58	—	50649.58

Graphical Program Analysis Tools V2.0-5 for DIGITAL UNIX

Subset	root	/usr	/var	Total
GPABASE205	—	38942.36	—	38942.36
GPAMVIEW205	—	19953.00	—	19953.00
GPAPROFHEAP205	—	39467.90	—	39467.90

Graphical Program Analysis Tools V2.0-5 for DIGITAL UNIX

Subset	root	/usr	/var	Total
GPAPVIEW205	—	4882.04	—	4882.04
Total	—	103245.31	—	103245.31

DIGITAL UNIX Worldwide Language Support V4.0E

Subset	root	/usr	/var	Total
IOSAACMENU435	—	0.00	—	—
IOSCACDEAPPS435	—	222.45	—	222.45
IOSCACDEDEV435	—	171.14	—	171.14
IOSCACDEDT435	—	1072.00	—	1072.00
IOSCACDEMAIL435	—	84.04	—	84.04
IOSCACDEMIN435	—	3987.55	—	3987.55
IOSCADECW435	—	177.65	—	177.65
IOSCAX11435	—	406.28	—	406.28
IOSCAXDEV435	—	94.94	—	94.94
IOSCSCDEAPPS435	—	215.37	—	215.37
IOSCSCDEDEV435	—	151.80	—	151.80
IOSCSCDEDT435	—	1092.53	—	1092.53
IOSCSCDEMAIL435	—	78.31	—	78.31
IOSCSCDEMIN435	—	474.06	—	474.06
IOSCSDECW435	—	750.62	—	750.62
IOSCSOLDDECW435	—	3949.19	—	3949.19
IOSCSOLDX11435	—	863.35	—	863.35
IOSCSUCSBASE435	—	119.50	—	119.50
IOSCSX11435	—	6723.10	—	6723.10
IOSCSXDEV435	—	94.77	—	94.77
IOSCSXMAIL435	—	353.78	—	353.78
IOSDECDEAPPS435	—	221.40	—	221.40
IOSDECDEDEV435	—	160.91	—	160.91
IOSDECDEDT435	—	1070.13	—	1070.13
IOSDECDEHLP435	—	18215.23	—	18215.23
IOSDECDEMAIL435	—	86.75	—	86.75
IOSDECDEMIN435	—	508.40	—	508.40
IOSDECHX11435	—	250.22	—	250.22
IOSDEDECW435	—	751.84	—	751.84
IOSDEOLDDECW435	—	3977.22	—	3977.22
IOSDEOLDX11435	—	877.02	—	877.02
IOSDEX11435	—	6632.24	—	6632.24
IOSDEXDEV435	—	94.92	—	94.92
IOSDEXMAIL435	—	410.34	—	410.34
IOSELFONT100M435	—	1111.38	—	1111.38
IOSELFONT100P435	—	1080.35	—	1080.35
IOSELFONT75M435	—	951.89	—	951.89
IOSELFONT75P435	—	926.76	—	926.76

DIGITAL UNIX Worldwide Language Support V4.0E

Subset	root	/usr	/var	Total
IOSELOLDDECW435	—	552.52	—	552.52
IOSELOLFONT435	—	2014.36	37.23	2051.59
IOSELUCSBASE435	—	76.16	—	76.16
IOSELX11435	—	431.61	—	431.61
IOSELXMAIL435	—	68.59	—	68.59
IOSESCDEAPPS435	—	217.63	—	217.63
IOSESCDEDEV435	—	159.09	—	159.09
IOSESCDEDT435	—	1083.51	—	1083.51
IOSESCDEHLP435	—	22216.91	—	22216.91
IOSESCDEMAIL435	—	85.09	—	85.09
IOSESCDEMIN435	—	520.19	—	520.19
IOSESDECW435	—	733.50	—	733.50
IOSESOLDDECW435	—	4004.10	—	4004.10
IOSESOLDX11435	—	886.37	—	886.37
IOSESX11435	—	6596.34	—	6596.34
IOSESXDEV435	—	94.99	—	94.99
IOSESXMAIL435	—	367.40	—	367.40
IOSFRBEX11435	—	250.23	—	250.23
IOSFRCAX11435	—	30.33	—	30.33
IOSFRCDEAPPS435	—	226.64	—	226.64
IOSFRCDEDEV435	—	152.82	—	152.82
IOSFRCDEDT435	—	1066.96	—	1066.96
IOSFRCDEHLP435	—	18676.56	—	18676.56
IOSFRCDEMAIL435	—	88.34	—	88.34
IOSFRCDEMIN435	—	521.19	—	521.19
IOSFRCHX11435	—	250.24	—	250.24
IOSFRDECW435	—	725.20	—	725.20
IOSFROLDDECW435	—	3963.56	—	3963.56
IOSFROLDX11435	—	870.65	—	870.65
IOSFRX11435	—	6639.24	—	6639.24
IOSFRXDEV435	—	95.04	—	95.04
IOSFRXMAIL435	—	384.21	—	384.21
IOSHUCDEAPPS435	—	217.01	—	217.01
IOSHUCDEDEV435	—	168.24	—	168.24
IOSHUCDEDT435	—	1081.57	—	1081.57
IOSHUCDEMAIL435	—	80.75	—	80.75
IOSHUCDEMIN435	—	4008.25	—	4008.25
IOSHUDECW435	—	722.65	—	722.65
IOSHUOLDDECW435	—	3954.59	—	3954.59
IOSHUOLDX11435	—	855.98	—	855.98
IOSHUUCSBASE435	—	103.50	—	103.50
IOSHUX11435	—	6576.63	—	6576.63
IOSHUXDEV435	—	94.86	—	94.86
IOSHUXMAIL435	—	361.43	—	361.43
IOSITCDEAPPS435	—	213.17	—	213.17
IOSITCDEDEV435	—	158.90	—	158.90

DIGITAL UNIX Worldwide Language Support V4.0E

Subset	root	/usr	/var	Total
IOSITCDEDT435	—	1068.05	—	1068.05
IOSITCDEHLP435	—	14581.31	—	14581.31
IOSITCDEMAIL435	—	88.63	—	88.63
IOSITCDEMIN435	—	526.28	—	526.28
IOSITDECW435	—	728.54	—	728.54
IOSITOLDDECW435	—	4006.67	—	4006.67
IOSITOLDX11435	—	891.73	—	891.73
IOSITX11435	—	6662.54	—	6662.54
IOSITXDEV435	—	94.99	—	94.99
IOSITXMAIL435	—	403.25	—	403.25
IOSIWBASE435	—	138.29	—	138.29
IOSIWCDEDT435	—	208.88	—	208.88
IOSIWCDEMIN435	—	106.00	—	106.00
IOSIWDECW435	—	84.12	—	84.12
IOSIWFONT100M435	—	632.91	—	632.91
IOSIWFONT100P435	—	2157.05	—	2157.05
IOSIWFONT75M435	—	408.37	—	408.37
IOSIWFONT75P435	—	1901.38	—	1901.38
IOSIWOLDX11435	—	11.77	—	11.77
IOSIWOLFONT435	—	3041.34	62.44	3103.78
IOSIWUCSBASE435	—	91.37	—	91.37
IOSIWX11435	—	1701.05	—	1701.05
IOSIWXDEV435	—	930.38	—	930.38
IOSJPABASE435	—	3546.82	—	3546.82
IOSJPAMANOS435	—	35.76	—	35.76
IOSJPBASE435	547.34	13901.91	6.10	14455.35
IOSJPCDEAPPS435	—	475.68	—	475.68
IOSJPCDEDEV435	—	1209.83	—	1209.83
IOSJPCDEDT435	—	3844.82	—	3844.82
IOSJPCDEHLP435	—	32523.57	—	32523.57
IOSJPCDEHLPSJIS435	—	32469.59	—	32469.59
IOSJPCDEMAIL435	—	307.84	—	307.84
IOSJPCDEMIN435	—	1899.02	—	1899.02
IOSJPDECW435	—	857.52	—	857.52
IOSJPFONT100M435	—	12069.76	—	12069.76
IOSJPFONT100P435	—	11912.50	—	11912.50
IOSJPFONT75M435	—	8454.37	—	8454.37
IOSJPFONT75P435	—	8310.46	—	8310.46
IOSJPFONTM435	—	13553.31	—	13553.31
IOSJPLDBBASE435	—	0.00	—	—
IOSJPLDBGUI435	—	0.00	—	—
IOSJPMANOS435	—	6855.71	—	6855.71
IOSJPMANWOS435	—	223.13	—	223.13
IOSJPMSG435	—	1596.55	—	1596.55
IOSJPMSGSJIS435	—	532.14	—	532.14
IOSJPNEMACS435	—	23417.63	76.00	23493.63

DIGITAL UNIX Worldwide Language Support V4.0E

Subset	root	/usr	/var	Total
IOSJPNEMACSSRC435	—	10272.62	—	10272.62
IOSJPOLDDECW435	—	6108.77	—	6108.77
IOSJPOLDX11435	69.39	956.94	—	1026.33
IOSJPPGMR435	—	2028.83	—	2028.83
IOSJPUCSBASE435	—	17159.06	—	17159.06
IOSJPWNN435	9.97	20149.03	95.00	20254.00
IOSJPWNNPGMR435	—	1216.16	—	1216.16
IOSJPWNNSRC435	—	10790.78	—	10790.78
IOSJPX11435	—	6993.54	—	6993.54
IOSJPXDEV435	—	126.53	—	126.53
IOSJPXMAIL435	—	791.29	—	791.29
IOSKOBASE435	—	1876.55	—	1876.55
IOSKOCDEAPPS435	—	185.78	—	185.78
IOSKOCDEDEV435	—	157.37	—	157.37
IOSKOCDEDT435	—	1386.45	—	1386.45
IOSKOCDEHLP435	—	21904.25	—	21904.25
IOSKOCDEMAIL435	—	130.05	—	130.05
IOSKOCDEMIN435	—	792.34	—	792.34
IOSKODECW435	—	181.13	—	181.13
IOSKOFONTM435	—	3826.33	—	3826.33
IOSKOFONTP435	—	9568.90	—	9568.90
IOSKOOLDDECW435	—	2111.14	—	2111.14
IOSKOOLDX11435	—	433.60	—	433.60
IOSKOOLFONT435	—	6159.31	3.07	6162.38
IOSKOPGMR435	—	168.11	—	168.11
IOSKOUCSBASE435	—	9122.83	—	9122.83
IOSKOX11435	—	4851.59	—	4851.59
IOSKOXDEV435	—	97.73	—	97.73
IOSKOXMAIL435	—	76.40	—	76.40
IOSLTX11435	—	43.87	—	43.87
IOSPLCDEAPPS435	—	199.49	—	199.49
IOSPLCDEDEV435	—	167.07	—	167.07
IOSPLCDEDT435	—	1078.94	—	1078.94
IOSPLCDEMAIL435	—	117.98	—	117.98
IOSPLCDEMIN435	—	4017.57	—	4017.57
IOSPLDECW435	—	740.93	—	740.93
IOSPLOLDDECW435	—	3945.02	—	3945.02
IOSPLOLDX11435	—	870.46	—	870.46
IOSPLUCSBASE435	—	119.50	—	119.50
IOSPLX11435	—	6813.58	—	6813.58
IOSPLXDEV435	—	94.89	—	94.89
IOSPLXMAIL435	—	372.84	—	372.84
IOSRUDECW435	—	733.34	—	733.34
IOSRUOLDDECW435	—	4123.13	—	4123.13
IOSRUOLDX11435	—	876.66	—	876.66
IOSRUUCSBASE435	—	82.27	—	82.27

DIGITAL UNIX Worldwide Language Support V4.0E

Subset	root	/usr	/var	Total
IOSRUX11435	—	6746.29	—	6746.29
IOSRUXDEV435	—	95.07	—	95.07
IOSRUXMAIL435	—	355.74	—	355.74
IOSSKCDEAPPS435	—	207.22	—	207.22
IOSSKCDEDEV435	—	275.49	—	275.49
IOSSKCDEDT435	—	1104.54	—	1104.54
IOSSKCDEMAIL435	—	80.67	—	80.67
IOSSKCDEMIN435	—	4013.70	—	4013.70
IOSSKDECW435	—	710.03	—	710.03
IOSSKOLDDECW435	—	3907.27	—	3907.27
IOSSKOLDX11435	—	846.86	—	846.86
IOSSKUCSBASE435	—	119.50	—	119.50
IOSSKX11435	—	6198.08	—	6198.08
IOSSKXDEV435	—	94.76	—	94.76
IOSSKXMAIL435	—	345.05	—	345.05
IOSSLX11435	—	54.98	—	54.98
IOSSVCDEAPPS435	—	215.22	—	215.22
IOSSVCDEDEV435	—	154.46	—	154.46
IOSSVCDEDT435	—	1023.96	—	1023.96
IOSSVCDEHLP435	—	12757.42	—	12757.42
IOSSVCDEMAIL435	—	73.28	—	73.28
IOSSVCDEMIN435	—	465.50	—	465.50
IOSSVDECW435	—	709.32	—	709.32
IOSSVOLDDECW435	—	3732.16	—	3732.16
IOSSVOLDX11435	—	835.55	—	835.55
IOSSVX11435	—	6262.17	—	6262.17
IOSSVXDEV435	—	94.80	—	94.80
IOSSVXMAIL435	—	345.22	—	345.22
IOSTHBASE435	—	805.45	—	805.45
IOSTHBIN435	439.39	12.20	6.10	457.68
IOSTHCDEAPPS435	—	138.79	—	138.79
IOSTHCDEDEV435	—	255.62	—	255.62
IOSTHCDEDT435	—	1187.13	—	1187.13
IOSTHCDEMAIL435	—	67.13	—	67.13
IOSTHCDEMIN435	—	485.97	—	485.97
IOSTHDECW435	—	172.56	—	172.56
IOSTHFONTM435	—	134.48	—	134.48
IOSTHOLDDECW435	—	1964.29	—	1964.29
IOSTHOLDX11435	—	434.15	—	434.15
IOSTHOLFONT435	—	7385.15	123.70	7508.85
IOSTHPGMR435	—	124.45	—	124.45
IOSTHPRINT435	—	172.89	—	172.89
IOSTHX11435	—	2877.22	—	2877.22
IOSTHXDEV435	—	107.04	—	107.04
IOSTHXMAIL435	—	74.33	—	74.33
IOSTRFONT100M435	—	1115.60	—	1115.60

DIGITAL UNIX Worldwide Language Support V4.0E

Subset	root	/usr	/var	Total
IOSTRFONT100P435	—	4113.84	—	4113.84
IOSTRFONT75M435	—	962.49	—	962.49
IOSTRFONT75P435	—	3468.61	—	3468.61
IOSTROLDDECW435	—	130.10	—	130.10
IOSTROLFONT435	—	5458.68	102.34	5561.02
IOSTRUCSBASE435	—	76.16	—	76.16
IOSTRX11435	—	502.37	—	502.37
IOSTRXMAIL435	—	68.58	—	68.58
IOSULUCSBASE435	—	73.08	—	73.08
IOSWWBASE435	18.42	1643.36	167.04	1828.82
IOSWWBIN435	901.82	129.38	12.20	1043.40
IOSWWBINCOM435	11.46	401.45	26.65	439.56
IOSWWCDEDT435	—	699.95	—	699.95
IOSWWFGC435	—	1492.74	—	1492.74
IOSWWFONTM435	—	702.40	—	702.40
IOSWWFONTP435	—	260.98	—	260.98
IOSWWLAT2FONT100M435	—	1195.91	—	1195.91
IOSWWLAT2FONT100P435	—	4363.31	—	4363.31
IOSWWLAT2FONT75M435	—	1022.86	—	1022.86
IOSWWLAT2FONT75P435	—	3677.32	—	3677.32
IOSWWLAT2OLFONT435	—	5576.47	102.34	5678.81
IOSWWLAT4FONT100M435	—	1183.15	—	1183.15
IOSWWLAT4FONT100P435	—	4386.89	—	4386.89
IOSWWLAT4FONT75M435	—	1014.75	—	1014.75
IOSWWLAT4FONT75P435	—	3699.16	—	3699.16
IOSWWLAT9FONT100M435	—	5291.42	—	5291.42
IOSWWLAT9FONT75M435	—	4470.69	—	4470.69
IOSWWLATCFONT100M435	—	1178.66	—	1178.66
IOSWWLATCFONT100P435	—	2424.95	—	2424.95
IOSWWLATCFONT75M435	—	1004.35	—	1004.35
IOSWWLATCFONT75P435	—	2030.64	—	2030.64
IOSWWLATCOLFONT435	—	3663.37	62.14	3725.52
IOSWWMOTIF11435	—	9837.15	—	9837.15
IOSWWMULE435	—	96197.59	—	96197.59
IOSWWMULESRC435	—	26982.03	—	26982.03
IOSWWOLDBKR435	—	5600.69	—	5600.69
IOSWWOLDDECW435	—	401.85	—	401.85
IOSWWPGMR435	—	320.78	—	320.78
IOSWWPHRASE435	386.22	837.19	6.10	1229.51
IOSWWPRINT435	38.18	1305.07	50.32	1393.57
IOSWWSVEDEV435	—	438.04	—	438.04
IOSWWSYSMAN435	—	388.63	4.78	393.41
IOSWWUCSBASE435	—	55821.29	—	55821.29
IOSWWUDCOS435	399.48	2261.19	6.10	2666.77
IOSWWUDCWOS435	—	122.60	—	122.60
IOSWWX11435	—	4733.72	—	4733.72

DIGITAL UNIX Worldwide Language Support V4.0E

Subset	root	/usr	/var	Total
IOSWWXDEV435	—	2357.28	—	2357.28
IOSWWXFR435	13.81	1170.21	3.49	1187.52
IOSZHBASE435	—	443.96	—	443.96
IOSZHBIG5435	168.07	1517.37	3.05	1688.48
IOSZHCNBASE435	—	419.92	—	419.92
IOSZHCNCDEAPPS435	—	123.36	—	123.36
IOSZHCNCDEDEV435	—	153.52	—	153.52
IOSZHCNCDEDT435	—	1139.06	—	1139.06
IOSZHCNCDEHLP435	—	14494.54	—	14494.54
IOSZHCNCDEMAIL435	—	72.92	—	72.92
IOSZHCNCDEMIN435	—	481.45	—	481.45
IOSZHCNLOC435	—	1239.15	—	1239.15
IOSZHCNUCSBASE435	—	1863.76	—	1863.76
IOSZHCONV435	82.05	259.66	3.05	344.76
IOSZHEUCTW435	—	814.20	—	814.20
IOSZHHANYU435	—	878.05	—	878.05
IOSZHHANZI435	—	750.41	—	750.41
IOSZHHKBASE435	—	3658.57	—	3658.57
IOSZHHKUCSBASE435	—	1565.14	—	1565.14
IOSZHPGMR435	—	2447.81	—	2447.81
IOSZHSDECW435	—	173.38	—	173.38
IOSZHSFONTM435	—	3379.82	—	3379.82
IOSZHSFONTP435	—	27595.29	—	27595.29
IOSZHSOLDDECW435	—	2023.17	—	2023.17
IOSZHSOLDX11435	—	425.72	—	425.72
IOSZHSOLFONT435	—	15014.51	6.19	15020.70
IOSZHSX11435	—	3236.29	—	3236.29
IOSZHSXDEV435	—	198.82	—	198.82
IOSZHSXMAIL435	—	75.07	—	75.07
IOSZHTDECW435	—	520.60	—	520.60
IOSZHTELEX435	150.94	2170.68	3.05	2324.67
IOSZHTFONTM435	—	9060.39	—	9060.39
IOSZHTFONTP435	—	25532.25	—	25532.25
IOSZHTOLDDECW435	—	5962.84	—	5962.84
IOSZHTOLDX11435	—	1284.83	—	1284.83
IOSZHTOLFONT435	—	28075.69	6.20	28081.89
IOSZHTWBASE435	—	3355.65	—	3355.65
IOSZHTWCDEAPPS435	—	367.62	—	367.62
IOSZHTWCDEDEV435	—	455.72	—	455.72
IOSZHTWCDEDT435	—	3447.43	—	3447.43
IOSZHTWCDEHLP435	—	18630.69	—	18630.69
IOSZHTWCDEMAIL435	—	196.19	—	196.19
IOSZHTWCDEMIN435	—	1371.94	—	1371.94
IOSZHTWLOC435	—	10095.71	—	10095.71
IOSZHTWUCSBASE435	—	4649.14	—	4649.14
IOSZHTX11435	—	8034.81	—	8034.81

DIGITAL UNIX Worldwide Language Support V4.0E

Subset	root	/usr	/var	Total
IOSZHTXDEV435	—	940.54	—	940.54
IOSZHTXMAIL435	—	233.12	—	233.12
IOSZHUCSBASE435	—	22580.40	—	22580.40
IOSZHX11435	—	6483.00	—	6483.00
Total	3236.54	1115916.56	974.67	1120127.77

Grand Totals	root	/usr	/var	Total
	3626.98	1888692.85	80712.24	1973032.05

Table B-3: DIGITAL UNIX V4.0E Associated Products, Volume 2

Advanced File System Advanced Utilities

Subset	root	/usr	/var	Total
AFAADVANCED435	146.36	562.22	—	708.58
AFAADVDAEMON435	10.87	1757.37	57.00	1825.24
AFAADVGUI435	47.43	14267.59	19.00	14334.03
AFAADVMAN435	11.45	68.35	—	79.80
AFAJPADVGUI435	—	1696.68	—	1696.68
AFAJPADVMAN435	—	92.86	—	92.86
Total	216.10	18445.07	76.00	18737.18

Advanced Server V4.0A for DIGITAL UNIX

Subset	root	/usr	/var	Total
ASUADM401	—	27601.94	—	27601.94
ASUADMJP401	—	26249.39	—	26249.39
ASUBASE401	10.80	36803.35	—	36814.16
ASUMANJP401	—	484.76	—	484.76
ASUMANPAGE401	—	461.20	—	461.20
ASUTRAN401	107.82	1782.24	1478.47	3368.54
Total	118.63	93382.90	1478.47	94980.00

Digital NetWorker

Subset	root	/usr	/var	Total
BRXCLNT520	—	49804.47	—	49804.47
BRXMAN520	—	2059.74	—	2059.74
BRXNODE520	—	16840.00	—	16840.00
BRXSERV520	—	29705.32	—	29705.32
Total	—	98409.53	—	98409.53

DEC Open3D V4.6

Subset	root	/usr	/var	Total
O3DDWSBASE460	—	10501.67	—	10501.67
O3DDWSCOMMON460	—	11815.88	6584.36	18400.25
O3DDWSCONFIG460	—	0.00	—	—
O3DDWSGLBASE460	—	7607.63	—	7607.63
O3DDWSGLEXAM460	—	12539.71	—	12539.71
O3DDWSGLMAN460	—	2424.31	—	2424.31
O3DDWSMITPEX460	—	1709.25	—	1709.25
O3DDWSMITPEXEXAM460	—	0.00	—	—
O3DDWSMITPEXMAN460	—	2943.81	—	2943.81
O3DDWSPCM460	—	788.88	—	788.88
O3DDWSPEX460	—	0.00	—	—
O3DDWSPHRJP460	—	2660.46	—	2660.46
O3DDWSPHRMAN460	—	36.28	—	36.28
O3DDWSPHRTO460	—	21450.31	—	21450.31
O3DDWSSTEREO460	—	196.38	—	196.38
O3DDWSSTEREOMAN460	—	12.92	—	12.92
O3DDWSZE3460	—	25579.35	—	25579.35
O3DDWSZLXE460	—	11249.82	—	11249.82
O3DDWSZLXL460	—	19808.99	—	19808.99
Total	—	131325.64	6584.36	137910.01

Performance Manager for Digital UNIX

Subset	root	/usr	/var	Total
PMGRAPP435	—	444.24	—	444.24
PMGRBASE435	—	4546.72	—	4546.72
PMGRCLUSTERS435	—	948.24	490.47	1438.71
PMGRGUI435	—	25957.58	—	25957.58
PMGRMAN435	—	74.98	—	74.98
PMGRUTIL435	—	19.09	1711.59	1730.68
Total	—	31990.85	2202.06	34192.91

SCSI CAM Layered Components V3.1D				
Subset	**root**	**/usr**	**/var**	**Total**
CLCMC314	14.63	3837.39	—	3852.02
CLCOP314	16.94	1935.17	—	1952.11
Total	31.57	5772.56	—	5804.13

System V Environment				
Subset	**root**	**/usr**	**/var**	**Total**
SVEADM425	434.71	64831.23	45.59	65311.53
SVEBCP425	44.16	22773.12	30.02	22847.29
SVEDEV425	11.82	33915.49	34.65	33961.95
SVEENV425	23.68	19.16	12.00	54.84
SVEMAN425	—	7476.22	—	7476.22
SVEPRINT425	14.02	54327.97	315.67	54657.66
Total	528.38	183343.19	437.93	184309.50

DIGITAL UNIX TruCluster(TM) Software V1.5				
Subset	**root**	**/usr**	**/var**	**Total**
TCRASE150	90.62	15755.23	1971.22	17817.07
TCRCMS150	—	10500.61	156.88	10657.48
TCRCOMMON150	54.42	1647.98	19.00	1721.40
TCRCONF150	—	58.40	—	58.40
TCRDSVC150	—	10136.21	2137.16	12273.37
TCRMAN150	—	1618.32	—	1618.32
TCRMCA150	—	3351.43	—	3351.43
Total	145.03	43068.19	4284.26	47497.48

Grand Totals	**root**	**/usr**	**/var**	**Total**
	1039.71	605737.93	15063.08	621840.74

B.2 Disk Space Required for Update Installations

For information about the disk space required to update the DIGITAL UNIX operating system from Version 4.0B, Version 4.0C, or Version 4.0D to Version 4.0E, see the DIGITAL UNIX *Software Product Description*.

B.3 Disk Space Required for RIS Areas

The Remote Installation Services (RIS) area for DIGITAL UNIX Version 4.0E requires approximately 1,113.7 MB of disk space. The space requirements are broken down as follows:

Table B-4: Disk Space Required for RIS Areas

Product Area	512-Byte Blocks
DIGITAL_UNIX_Operating_System	1073486
Alternative_Compiler	15268
CDE_Video	137454
DEC_Ada_RTL	556
DEC_C++_RTL	336
DEC_Cobol_RTL	2304
DEC_EVENT	25768
DEC_Fortran_RTL	2420
DEC_Pascal_RTL	818
DEC_Sort_RTL	376
DECtalk_Runtime	9236
GNUSRC	23878
Multimedia_Service	44476
Netscape_FastTrack_Server	71540
PanoramiX_ADK	9364
Porting_Assistant	20144
Program_Analyzers	33642
Worldwide_Language_Support	646974
AdvFS_Advanced_Utilities	9332
Advanced_Server	60642
NetWorker	98894
Open3d_460	73238
Performance_Manager	14016
SCSI_CAM_Layered_Components	3318
System_V_Environment	77246
TruCluster	21634
Total	2444860

B.4 Disk Space Required for Documentation

The files for the DIGITAL UNIX documentation set are contained on DIGITAL UNIX Documentation CD–ROM. The files for the worldwide support documentation are contained on the DIGITAL UNIX Associated

Products CD–ROM in the
/mnt–pnt/`Worldwide_Language_Support/doc` directory. These files require the following amounts of disk space if moved from the CD–ROM onto a local disk:

- DIGITAL UNIX operating system documentation – 100 MB
- Worldwide support documentation – 6.5 MB

Enabling and Disabling Extended UID/GID Support

C

With DIGITAL UNIX Version 4.0E, you can enable or disable extended UID and GID support, by following the instructions in this appendix.

C.1 Enabling Extended UID and GID Support

To enable extended UID and GID support, do the following:

1. Become the root user.
2. Use the `sysconfigdb` utility to add the following attribute to the `proc` subsystem of the `/etc/sysconfigtab` file:

   ```
   enable_extended_uids=1
   ```

   Alternatively, you can use the `dxkerneltuner` graphical user interface to set this attribute.
3. Reboot the system.

Extended UID and GID support is now enabled.

C.2 Disabling Extended UID and GID Support

To disable extended UID and GID support, do the following:

1. Become the root user.
2. Use `sysconfigdb` or `dxkerneltuner` to set the `enable_extended_uids` attribute to 0 in the `/etc/sysconfigtab` file:

   ```
   enable_extended_uids=0
   ```

3. Reboot the system.

Extended UID and GID support is now disabled.

Note

When extended UIDs and GIDs are disabled, files owned by a user with an extended UID or GID will be inaccessible to all users except root. Any user with an extended UID or GID will not have access to log in to the system or use the `su` command to access their accounts.

C.3 Checking for Extended UID and GID Support

The following sample program demonstrates how to check the maximum number of UIDs supported by a version of the DIGITAL UNIX operating system. The maximum number of GIDs will always be the same as the maximum number of UIDs:

```
#include <unistd.h>
#include <limits.h>
#include <errno.h>

#include <sys/table.h>

#ifndef TBL_UIDINFO
#define TBL_UIDINFO 56
#endif

main()
{
     uid_t uid_max;

     errno = 0;

     uid_max = table (TBL_UIDINFO, 0, (char *)0, 1, 0);

     if ((errno != 0) && ((int)uid_max < 0))
          uid_max = UID_MAX;

     printf(%d\n", uid_max);

}
```

C.4 Applications Affected by Extended UIDs and GIDs

The following programs are affected by extended UID and GID support.

C.4.1 Clusters

Extended UID and GID support can be enabled in a TruCluster Available Server Software or TruCluster Production Server Software configuration only after all member systems have installed (or upgraded to) Version 1.5 of the appropriate TruCluster software product. To enable extended UID or GID support, set the `enable_extended_uids` parameter on every system as directed in this appendix and reboot every system. Do not use extended UIDs and GIDs on any member system until you have rebooted the last member system.

Once you have enabled extended UID and GID support in a cluster, you must not disable it. Disabling this support will disrupt the operation of your

cluster.

C.4.2 Kerberos

Kerberos Version 4 does not support extended UIDs and GIDs. If you use Kerberos Version 4 and need extended UID and GID support, you should upgrade to Kerberos Version 5.

C.4.3 System V File System

The System V File System (S5FS) does not support extended UIDs and GIDs. File system `syscalls` that specify UIDs and GIDs greater than 65,535 will return an `EINVAL` error. Users assigned a UID or GID greater than 65,535 will not be able to create or own files on a System V File System. Consider using the UFS, MFS, or AdvFS for a workaround.

C.4.4 The ls Command

The `ls -l` command does not display the disk block usage on quota files or sparse files. This is not a result of the implementation of extended UIDs and GIDs, but rather a result of the behavior of the `ls -l` command. When extended UIDs and GIDs are enabled, quota files and sparse files may appear much larger than expected. To display the actual disk block usage for any file, use the `ls -s` command.

C.4.5 The cp Command

The `cp` command will incorrectly copy quota files or other sparse files. This is not a result of the implementation of extended UIDs and GIDs, but rather a result of the behavior of the `cp` command when it reads a file. When extended UIDs and GIDs are enabled, quota files and other sparse files may be copied to a new file that is much larger than expected. To correctly copy quota files or other sparse files, use the `dd` command with the `conv=sparse` parameter.

```
% dd conv=sparse if=inputfile of=outputfile
```

C.4.6 The vdump/vrestore Utilities and UFS File Systems

If a UFS file system that contains quota files or other sparse files is backed up using the `vdump` utility and restored using the `vrestore` utility, the quota files or other sparse files will be restored as follows:

- The first page of a file on disk will be restored as a fully populated page; that is, empty nonallocated disk blocks will be zero filled.

- Any additional pages on disk will be restored sparse.

C.4.7 The dxarchiver Utility

The `dxarchiver` utility does not support extended UIDs and GIDs. However, the `pax` and `tar` utilities do support extended UIDs and GIDs and can be used as alternatives. If you need to use the `dxarchiver` utility, you must not enable extended UID or GID support.

C.4.8 The cpio Utility

The `cpio` utility does not support extended UIDs and GIDs. However, the `pax` and `tar` utilities do support extended UID and GIDs and can be used as alternatives. If you need to use the `cpio` utility, you must not enable extended UID or GID support.

C.4.9 The pax Utility

The `pax` utility, used to extract, list, or write archive files has been modified to support long file names and extended UID/GID values

To take advantage of these enhancements, specify the `xtar` format with the `-x` option. The following text has been added to the description of the `-x` option on the `pax` reference page:

```
xtar  Extended tar interchange format.  The default blocking value
      for his format for character special archive files is 10240.
      Blocking values from 512 to 32,256 in increments of 512
      are supported.  This option lets the user archive long file
      names and extended UID/GID values.
```

On the reference page, the `xtar` format may be incorrectly identified as the `ustar` format. The correct syntax is `-x xtar`.

C.4.10 The tar Utility

The `tar` utility, used to extract, list, or write archive files has been modified to support long file names and extended UID/GID values

To take advantage of these enhancements, specify the `-E` option when using the `tar` utility.

C.4.11 PATHWORKS

PATHWORKS does not support extended UIDs and GIDs. If you use PATHWORKS and need extended UID or GID support, you should upgrade to Version 4.0 or higher of Advanced Server for DIGITAL UNIX (ASDU).

Netscape Communicator and Netscape FastTrack Server D

This appendix contains information about how to install and configure Netscape Communicator and Netscape FastTrack Server.

D.1 Netscape Communicator

This release of DIGITAL UNIX contains Version 4.05 of the Netscape Communicator. Netscape Communicator integrates the Netscape Navigator World Wide Web browser with Netscape's Messenger Mailbox, Collabra Discussions, and Page Composer components. The application can be invoked through its CDE desktop icon, which is located in the CDE Application Manager's Desktop_Apps folder. Netscape Communicator can also be invoked directly from the command line by running `/usr/bin/X11/netscape`. Detailed help on Netscape Communicator can be accessed through the application's help menus.

A sample resource defaults file for Netscape Communicator can be found at `/usr/doc/netscape/Netscape.ad`. Comments within this file indicate possible settings for each resource. It is not necessary to install this file. It is provided for informational purposes.

The file `/usr/bin/X11/netscape` is actually a shell script that performs the following actions:

- Creates a Communicator user configuration directory if `$HOME/.netscape` does not already exist. If you have a file or directory in `$HOME` called `.netscape`, it is highly recommended that you rename it prior to running this version. If you do not, there is a possibility that you could lose data contained in this directory (for example, old preferences, bookmarks, cookies, and so on).
- If a Communicator `preferences.js` file does not already exist at `$HOME/.netscape/preferences.js`, one is copied there from `/usr/doc/netscape/default-netscape-preferences`.
- The `MOZILLA_HOME` environment variable is set to point to the Communicator installation directory `/usr/lib/netscape`. The Communicator java class files, nethelp files, plugins, registry, and so on, are all installed under this directory. See the `/usr/lib/netscape/README` file for more information on `MOZILLA_HOME` and other environment variables.

- The actual Communicator binary at `/usr/bin/X11/real-netscape` is invoked.

The DIGITAL UNIX *Installation Guide* contains more information about how to set up Netscape. See Chapter 6, which covers postinstallation setup tasks.

D.2 Netscape FastTrack Server

This release of DIGITAL UNIX contains Version 3.01 of the Netscape FastTrack Server, an easy-to-use entry-level Web server designed to let you create and manage a Web site.

The Netscape FastTrack Server is provided in the DIGITAL UNIX `OSFNETSCAPEFASTTRACK301` subset on the *DIGITAL UNIX 4.0E Associated Products, Volume 1* CD–ROM. To install FastTrack on your system, perform the following steps:

1. Log in to the root account on your system.
2. Use the `/usr/sbin/setld` software subset management utility to install the `OSFNETSCAPEFASTTRACK301` subset. This will install the FastTrack server in the `/usr/opt/netscape/suitespot` directory.
3. Change your directory to the FastTrack kit directory:

   ```
   # cd /usr/opt/netscape/suitespot
   ```

4. Read the `readme.txt` file:

   ```
   # more readme.txt
   ```

   This file contains information from Netscape about the FastTrack kit. Because you have installed FastTrack from the `OSFNETSCAPEFASTTRACK301` subset and not from a CD–ROM that Netscape ships, you can ignore steps 1-5 of the installation instructions, which pertain to installing from the Netscape CD–ROM.
5. Begin the configuration by running the `ns-setup` program and following the instructions in the `readme.txt` file:

   ```
   # ./ns-setup
   ```

   The installation program will query you for several pieces of information and then configure the FastTrack server.

The installation of FastTrack Version 3.01 will not update FastTrack Version 2.01 servers installed on your system. If you have an existing FastTrack 2.01 server installed on your system and wish to upgrade it to version 3.01, use the following process:

1. Install FastTrack 3.01 as described in the previous steps.

2. Access the FastTrack 3.01 administration server as directed by the installation program.
3. Click on the Migrate option from the initial administration server screen. This option will allow you to migrate existing FastTrack 2.01 server instances to FastTrack 3.01.

Maximum Number of Open File Descriptors E

You can enable extended (64K) file descriptors for a specific application by following these steps:

1. Set the *SSI_FD_NEWMAX* option of the `setsysinfo()` system call to 1. When *SSI_FD_NEWMAX* is set to 1, the process `utask` bit is turned on and the process's hard file limit is raised to 64K. This setting is inherited by child processes.
2. Set the process's file descriptor soft limit to be greater than 4K by using the `setrlimit()` call:

   ```
   #include <sys/resource.h>

   struct rlimit *rlp;

   rlp->rlim_cur = 6000;
   rlp->rlim_max = 6000;
   setrlimit(RLIMIT_NOFILE, rlp);
   ```

   This setting is inherited by child processes. See the `setrlimit()` reference page for details.
3. This step is required only for applications that use `select()`, `fd_set`, `FD_CLR`, `FD_ISSET`, `FD_SET`, or `FD_ZERO`. Choose one of the following:
   - Override the default value (4k) for *FD_SETSIZE* in `<sys/select.h>` by specifying the new maximum (65536). You must do this before including the `<sys/time.h>` header file, which includes `<sys/select.h>`:

     ```
     #define FD_SETSIZE 65536
     #include <sys/time.h>
     ```

     This setting is not inherited by child processes; therefore, `FD_SETSIZE` must be set explicitly in the code of a child process that requires 64k file descriptor support.
   - Instead of using statically-defined `fd_set` structures, use `fd_set` pointers in conjunction with a `malloc` function to provide forward compatibility with any future changes to the maximum file descriptor

limit:

```
fd_set *fdp;

fdp = (fd_set *) malloc(
(fds_howmany(max_fds,FD_NFDBITS))*sizeof(fd_mask));
```

The `max_fds` would be the number of file descriptors to be manipulated. The soft file descriptor limit is a good choice. All other keywords are defined in the `select.h` header file:

```
#include <sys/time.h>
#include <sys/resource.h>


my_program()
{
fd_set *fdp;
struct rlimit rlim;
int max_fds;


getrlimit(RLIMIT_NOFILE, &rlim);
max_fds = rlim.rlim_cur;

fdp = (fd_set *) malloc(
(fds_howmany(max_fds,FD_NFDBITS))*sizeof(fd_mask));

FD_SET(2, fdp);

for (;;) {
switch(select(max_fds, (fd_set *)0, fdp, (fd_set
*)0,
struct timeval *)0)) {
}
```

System administrators should note that `max-vnodes` must be set high enough for the needs of any application that requires a high number of descriptors. The `max-vnodes` is the number of in-memory structures for files, and is usually set by default to 5 percent of system memory. A rough guide is 1 vnode per descriptor. For more information about this setting, see the *System Configuration and Tuning* guide.

To disable support for up to 64K file descriptors once the previous procedure has been carried out, set the *`SSI_FD_NEWMAX`* option of the `setsysinfo()` system call to 0. When *`SSI_FD_NEWMAX`* is set to 0, the process `utask` bit is turned off and the process's hard file limit returns to the default maximum of 4K, unless the process is using more than 4K file descriptors. In that case, the `setsysinfo()` call returns an error, `EINVAL`. Additionally, if a calling process's file descriptor hard or soft limit exceeds 4K, it is set to 4K. This setting can be inherited by child processes.

Advanced File System Issues F

This appendix provides information about how to solve problems you may encounter with the Advanced File System (AdvFS).

F.1 Recovering from AdvFS Domain Panics

Typically, an AdvFS domain panic signals a hardware problem. However, it also may indicate a software failure, particularly in a domain that contains critical system files, such as the root file system.

The benefit of an AdvFS domain panic is that it prevents further access by users to a single AdvFS file domain and allows the filesets in that domain to be unmounted and examined. In addition, all other AdvFS domains remain on line and in normal operation, unaffected by the domain panic.

To recover an AdvFS domain from a domain panic, you need to collect as much information as you can about the condition of AdvFS filesets and metadata, in case there are software problems you need to correct or document for a problem report. Then, you need to take the steps that will let you run the `verify` utility on the domain in an attempt to check the integrity of its filesets.

To recover from an AdvFS domain panic, perform the following steps. If you cannot successfully complete steps 1-6, go to step 8:

1. Use the `mount` command to obtain a list of all the filesets in the domain. Then, use the `umount` command to unmount all of the file sets in the domain. For example:

   ```
   # mount -t advfs
   # umount fileset_name
   ```

   Note that the filesets in a domain must be unmounted to run the `verify` utility that checks them. See the restrictions noted in the `verify`(8) reference page for more information.

2. Use the `ls -l` command on the `/etc/fdmns` directory for the domain in order to obtain a list of the AdvFS volumes in the domain. For

example:

```
# ls -l /etc/fdmns/staff_projects1
```

3. Use the `vfile` command to collect information about the metadata files for each volume in the domain, in case you cannot recover the domain with the `verify` operation. You need to record information about the bitfile metadata table (BMT), the storage bitmap (SBM), the root tag directory, and the transaction log file for each disk. See the `vfile`(8) reference page for additional information. For example:

```
# /sbin/advfs/vfile 0 0 rz3c > bmt_rz3c
# /sbin/advfs/vfile 0 1 rz3c > sbm_rz3c
# /sbin/advfs/vfile 0 2 rz3c > tag_rz3c
# /sbin/advfs/vfile 0 3 rz3c > log_rz3c
```

4. Use the `dia` utility to extract information about the domain panic from the binary error log, as documented in the `dia`(8) reference page.
5. If the problem is a hardware problem, fix it before going to step 6.
6. Run the `verify` utility on all of the filesets in the domain. For example:

```
# verify staff_projects1
```

7. If the `verify` utility exits successfully, mount all of the file sets you had unmounted in step 1. You can resume normal operations. If the `verify` utility indicates that there is a problem, go to step 8.
8. If there is a failure that prevents complete recovery, you must first re-create the domain and restore the domain's data from backup media. Then, mount all of the restored filesets in the domain and resume normal operations. You should also file a problem report with Compaq; include the information you collected during the procedure.

F.2 Correcting Overlapping Frag Data Corruption

In versions of DIGITAL UNIX prior to Version 4.0D, a bug existed in AdvFS that could result in data corruption of user files. Only sparse files, those files that contain offsets at which no data has been stored, were affected by this bug. The bug has been fixed in 4.0D, but files created using older versions of DIGITAL UNIX may still be corrupted. In these files, AdvFS may have stored two different versions of a particular page (an 8k segment). User intervention is necessary to correct this problem. The files must be recreated or fixed with the `verify` command.

The `verify` command, located in `/sbin/advfs`, can detect files that have been corrupted by this bug. You should run `verify`, as in the following example, to see if any files are corrupted. Then, optionally, execute the

command again with the `-f` flag to enable retrieval of the missing data.

```
# /sbin/advfs/verify test_domain
+++ Domain verification +++

Domain Id 32d3e638.000a46a0

Checking disks ...

Checking storage allocated on disk /dev/rz1a

Checking mcell list ...

Checking mcell position field ...

Checking tag directories ...

+++ Fileset verification +++

+++ Fileset test_fileset +++

Checking frag file headers ...

Checking frag file type lists ...

Scanning directories and files ...
Overlapping frag data corruption detected in:
File: <mount point>/50226.file.4
Page: 1
Run verify -f on this domain to enable recovery of this data.

Scanning tags ...

Searching for lost files ...

#
```

The `verify` utility has detected a corrupted file in the `test_fileset` fileset. The name of the file is `50226.file.4` and it is located in the uppermost directory of the fileset when it is mounted. The corrupted page is page 1. The `verify` utility also suggests running `verify` again, using the `-f` flag to enable recovery of the hidden data for page 1.

At this point, you have two choices:

- Delete the file and recreate it. The corruption problem has been fixed on the system. The newly created file will not exhibit the unwanted behavior.
- Execute `verify -f` to identify the corrupted data, as in the following example:

```
# /sbin/advfs/verify -f test_domain
+++ Domain verification +++

Domain Id 32d3e638.000a46a0

Checking disks ...

Checking storage allocated on disk /dev/rz1a

Checking mcell list ...

Checking mcell position field ...

Checking tag directories ...

+++ Fileset verification +++

+++ Fileset test_fileset +++

Checking frag file headers ...

Checking frag file type lists ...

Scanning directories and files ...
Overlapping frag data corruption detected in:
File: <mount point>/50226.file.4
Page: 1
Temporary files created representing the two versions of
page 1 of file <mount point>/50226.file.4
Refer to the Release Notes for a description
of how to use these temporary files to recover
from this overlapping frag corruption problem.
Scanning tags ...

Searching for lost files ...

#
```

The `verify` utility reports that it has created two temporary files in the same directory as the corrupted file. Mount the fileset to identify these two files:

```
# mount test_domain#test_fileset /test
# ls -l /test
total 169
drwx------   2 root     system      8192 Jan  8 13:23 .tags
-rw-r--r--   1 root     system     24576 Jan  9 12:27 50226.file.1
-rw-r--r--   1 root     system     40960 Jan  9 12:27 50226.file.2
-rw-r--r--   1 root     system     32768 Jan  9 12:27 50226.file.3
-rw-r--r--   1 root     system     24576 Jan  9 12:27 50226.file.4
-rw-------   1 root     system      8192 Jan 13 14:32 50226.file.4.page_1.ext
-rw-------   1 root     system      8192 Jan 13 14:32 50226.file.4.page_1.frag
-rw-r-----   1 root     operator    8192 Jan  8 13:23 quota.group
-rw-r-----   1 root     operator    8192 Jan  8 13:23 quota.user
#
```

The `.ext` and `.frag` files contain the following information from the

corrupted area:

- `50226.file.4`

  This is the original corrupted file.

- `50226.file.4.page_1.ext`

  This file contains the hidden version of page 1 of the corrupted file. A read() system call cannot retrieve this data.

- `50226.file.4.page_1.frag`

  This file contains the fragmented version of page 1 of the corrupted file. This is the same data that a read() of page 1 would return.

To fix the corrupted file:

1. View the `.ext` and `.frag` files to determine which to keep. Note that you may want to merge the two files. If the `50226.file.4.page_1.ext` file contains the data you want, enter:

   ```
   # ln -s 50226.file.4.page_1.ext desired_page_1
   ```

   If the `50226.file.4.page_1.frag` file contains the data you want, enter:

   ```
   # ln -s 50226.file.4.page_1.frag desired_page_1
   ```

   If you must merge the two files, do the merge and put the result into a new file called `desired_page_1`.

2. Next, use the corrupted file and the new file (`desired_page_1` in this example) to create a new fixed version of the corrupted file. Copy page 0 from the corrupted file into a new file:

   ```
   # dd if=50226.file.4 of=newfile bs=8192 count=1 > /dev/null 2>&1
   ```

3. Append the page 1 that you selected to the new file (the "\" character indicates line continuation):

   ```
   # dd if=desired_page_1 of=newfile bs=8192 count=1 seek=1 > = \
   /dev/null 2>&1
   ```

4. Append the remainder of the original file to the end of the new file (the "\" character indicates line continuation):

   ```
   # dd if=50226.file.4 of=newfile bs=8192 seek=2 skip=2 > = /dev/null \
   2>&1
   ```

   Run the `diff` command on the new and the original file to confirm that only page 1 has changed and that the difference is what you want.

5. Rename the new file and remove the temporary files:

```
# mv newfile 50226.file.4
# rm 50226.file.4.page_1.ext 50226.file.4.page_1.frag desired_page_1
```

   If you want, you can now run the `verify` command on the domain again to confirm that the data corruption problem is gone.

Enabling Enhanced Core File Naming G

With DIGITAL UNIX Version 4.0E, you can enable enhanced core file naming, by following the instructions in this appendix.

When enhanced core file naming is enabled, the system produces core files with names in the following format:

core.*prog-name*.*host-name*.**tag**

This file name can be broken down into four parts:

- `core`

  The literal string `core`.

- *`program_name`*

  Up to sixteen characters taken from the program name as shown by the `ps` command.

- *`host_name`*

  The first portion of the system's network host name, or up to 16 characters of the host name, taken from the part of the host name that precedes the first dot.

  For example, the fourth core file generated on host `buggy.net.ooze.com` by the program `dropsy` would be:

  ```
  core.dropsy.buggy.3
  ```

- *`numeric_tag`*

  This tag is assigned to the core file to make it unique among all the core files generated by a program on a host. The maximum value for this tag, and thus the maximum number of core files for this program and host, is set by a system configuration parameter.

  Note that the tag is not a literal version number. The system selects the first available unique tag for the core file. For example, if a program's core files have tags .0, .1, and .3, the system uses tag .2 for the next core file it creates for that program. If the system-configured limit for core file instances is reached, the system will not create any more core files for that program/host combination. By default, the system can create up to 16 versions of a core file.

Note

If you plan to save a number of uniquely named core files, be aware that core files can consume a lot of disk space. Allowing core files to be saved under different names in a file system with minimal free space can potentially fill your disk because the files are not overwritten when new core files are created. If you enable this feature, make sure to remove old core files when you are done examining them.

You can enable enhanced core file naming in one of two ways.

- You can enable it at the system level by setting the `enhanced-core-name` system configuration variable to 1 in the `proc` subsystem, as in the following example:

  ```
  proc:
          enhanced-core-name = 1
  ```

  The system manager can limit the number of unique core file versions that a program can create on a specific host system by setting the system configuration variable `enhanced-core-max-versions` to the desired value, as in the following example:

  ```
  proc:
          enhanced-core-name = 1
          enhanced-core-max-versions = 8
  ```

  The minimum value is 1, the maximum value is 99,999, and the default is 16.

- You can enable enhanced core file naming at the program level by calling the `uswitch` system call with the `USW_CORE` flag set, as in the following example:

  ```
  #include <signal.h>
  #include <sys/uswitch.h>

  /*
   * Request enhanced core file naming for
   * this process then create a core file.
   */
  main()
  {
          long uval = uswitch(USC_GET, 0);
          uval = uswitch(USC_SET, uval | USW_CORE);
          if (uval < 0) {
                  perror("uswitch");
                  exit(1);
          }
          raise(SIGQUIT);
  }
  ```

Early Child Process Exit Notification H

In DIGITAL UNIX Version 4.0E, the early exit for child process notification feature alerts a parent process when a child process is about to dump core and exit abnormally with a signal. Because writing a core file can take an indeterminate period of time, early notification allows the parent process to take immediate action to recover from the child process' impending termination.

A parent process requests early exit notification when establishing a signal handler for the SIGCHLD signal with the sigaction system call. The sigaction flag SA_CLDNOTIFY requests early notification; you should use this flag in combination with the SA_SIGINFO flag to provide complete information about the terminating child.

Setting the SA_CLDNOTIFY flag in combination with SA_SIGINFO causes two SIGCHLD signals to be sent to the parent for each abnormal termination of a child process:

1. The early-notification signal delivers a siginfo structure with the si_signo field set to SIGCHLD, the si_code field set to CLD_SIGEXITING, and the si_pid field set to the process ID of the child process that is about to write a core file and terminate abnormally. This signal tells the parent process that it can start its failover operation or take other appropriate action.
2. The final-notification signal delivers a siginfo structure with the si_signo field set to SIGCHLD, the si_code field set to CLD_DUMPED, and the si_pid field set to the process ID of the child that has terminated abnormally. This signal tells the parent process that the child process is now a zombie and should be cleaned with a call to wait.

Note that the use of SA_CLDNOTIFY in the recommended combination with SA_SIGINFO in effect turns the SIGCHLD signal into a "realtime" signal. Every instance of the SIGCHLD signal is queued first-in, first-out to the parent process. Every early child exit and every final child exit generates a SIGCHLD signal to the parent. Contrast this behavior with the default behavior of the SIGCHLD signal, in which multiple child exit signals can be compressed into a single signal when the SIGCHLD signal is already pending for the parent process.

The following program shows an example of using early notification of child process termination.

```
#include <stdio.h>
#include <sys/types.h>
#include <sys/signal.h>
#include <sys/siginfo.h>

volatile pid_t kid1, kid2;

/*
 * Handler for SIGCHLD signal. Note: printf() calls
 * are technically unsupported from signal handlers
 * and are shown for illustrative purposes only.
 */
void
sigchld_handler(int sig, siginfo_t *sip, void *extra)
{
        pid_t kid;
        kid = sip->si_pid;
        if (sip->si_code == CLD_SIGEXITING) {
                printf("SIGEXITING: Got signal %d, si_code %d"
                       " for kid %d\n",
                       sip->si_signo, sip->si_code, kid);
        } else if (sip->si_code == CLD_DUMPED) {

                printf("EXITED: Got signal %d, si_code %d"
                       " for kid %d\n",
                        sip->si_signo, sip->si_code, kid);
                kid = wait(0);
                printf("Parent got PID %d exiting\n", kid);
                if (kid == kid1)
                        kid1 = 0;
                else if (kid == kid2)
                        kid2 = 0;
        }
}

main()
{
        struct sigaction sa;
        int ret;

        /*
         * Set up SIGCHLD handler for early exit
         * notification.
         */
        sa.sa_sigaction = sigchild_handler;
        sigemptyset(&sa.sa_mask);
       sa.sa_flags = SA_CLDNOTIFY|SA_SIGINFO;

        ret = sigaction(SIGCHLD, &sa, 0);
        if (ret) {
                perror("sigaction");
                exit(1);
```

```
        }


        /*
         * Create 2 children to die with SIGABRT
         * and create core files.
         */
        kid1 = fork();
        if (kid1)
                printf("Parent forked %d\n", kid1);
        else {
                /*
                 * First child...
                 */
                sleep(1);
                printf("Kid %d will now die with core file...\n",
                        getpid());
                abort();
                /*NOTREACHED*/
        }
        kid2 = fork();
        if (kid2)
                printf("Parent forked %d\n", kid2);
        else {
                /*
                 * Second child...
                 */
                sleep(1);
                printf("Kid %d will now die with core file...\n",
                        getpid());
                abort();
                /*NOTREACHED*/
        }

        /*
         * Parent: keep busy until children exit.
         */
         while (kid1 || kid2)
               ;
               ;
        /*
         * Children have exited: verify by printing zeros.
         */
        printf("Parent: done -- kid1 %d, kid2 %d\n", kid1, kid2);
               ;
        /*
         * Children have exit: verify by printing zeros.
         */
        printf("Parent: done -- kid1 %d, kid2 %d\n", kid1, kid2);
}
```

Thread Local Storage

In DIGITAL UNIX Version 4.0D, Thread Local Storage (TLS) support has been added to the DEC C compiler.

I.1 Definition

Thread Local Storage is a name for data that has static extent (that is, not on the stack) for the lifetime of a thread in a multithreaded process, and whose allocation is specific to each thread.

In standard multithreaded programs, static extent data is shared among all threads of a given process, whereas Thread Local Storage is allocated on a per-thread basis such that each thread has its own copy of the data that can be modified by that thread without affecting the value seen by the other threads in the process. For a complete discussion of threads, see the *Guide to DECthreads* and the *Programmer's Guide*.

I.2 Background

The essential functionality of Thread Local Storage is and has been provided by explicit application program interfaces (APIs) such as POSIX (DECThreads) `pthread_key_create()`, `pthread_setspecific()`, `pthread_getspecific()`, and `pthread_key_delete()`.

Although these APIs are portable to POSIX-conforming platforms, using them can be cumbersome and error-prone. Significant engineering work is typically required to take existing single-threaded code and make it thread-safe by replacing all of the appropriate static and external variable declarations and their uses by calls to these thread-local APIs. Furthermore, for Win32 platforms there is a somewhat different set of APIs: `TlsAlloc()`, `TlsGetValue()`, `TlsSetValue()`, and `TlsFree()`, which have the same kinds of usability problems as the POSIX APIs.

By contrast, the Thread Local Storage language feature is much simpler to use than any of the APIs, and it is especially convenient to use when converting single-threaded code to be multithreaded. This is because the change to make a static or external variable have a thread-specific value involves only adding a storage-class qualifier to its declaration. The compiler, linker, program loader, and debugger effectively implement the complexity of the API calls automatically for variables declared with this

qualifier. Unlike coding to the APIs, you do not need to find and modify all uses of the variables, or to add explicit allocation and deallocation code. While the language feature is not generally portable under any formal programming standard, it is portable between DIGITAL UNIX and Win32 platforms.

I.3 The __thread Attribute

The C and C++ compilers for DIGITAL UNIX include the extended storage-class attribute, `__thread`.

You must use the `__thread` attribute with the `__declspec` keyword to declare a thread variable. For example, the following code declares an integer thread local variable and initializes it with a value:

```
__declspec( __thread ) int tls_i = 1;
```

I.4 Guidelines and Restrictions

You must observe the following guidelines and restrictions when declaring thread local objects and variables:

- You can apply the `__thread` storage-class attribute only to data declarations and definitions. You cannot use it on function declarations or definitions. For example, the following code generates a compiler error:

  ```
  #define Thread  __declspec( __thread )
  Thread void func();            // Error
  ```

- You can specify the `__thread` attribute only on data items with static storage duration. This includes global data objects (both static and external), local static objects, and static data members of C++ classes. You cannot declare automatic or register data objects with the `__thread` attribute. For example, the following code generates compiler errors:

  ```
  #define Thread  __declspec( __thread )
  void func1()
  {
  Thread int tls_i;            // Error
  }

  int func2( Thread int tls_i )    // Error
  {
  return tls_i;
  }
  ```

- You must use the `__thread` attribute for the declaration and the definition of a thread-local object, whether the declaration and definition

occur in the same file or separate files. For example, the following code generates an error:

```
#define Thread          __declspec( __thread )
extern int tls_i;      // This generates an error, because the
int Thread tls_i;      // declaration and the definition differ.
```

- You cannot use the `__thread` attribute as a type modifier. For example, the following code generates a compile-time error:

  ```
  char __declspec( __thread ) *ch;            // Error
  ```

- The address of a thread-local object is not considered a link-time constant, and any expression that involves such an address is not considered a constant expression. In standard C, the effect of this is to forbid the use of the address of a thread-local variable as an initializer for an object that has static or thread-local extent. For example, the following code is flagged as an error by the C compiler if it appears at file scope:

  ```
  #define Thread  __declspec( __thread )
  Thread int tls_i;
  int *p = &tls_i;            // ERROR
  ```

  Standard C permits initialization of an object or variable with an expression involving a reference to itself, but only for objects of nonstatic extent. Although C++ normally permits such dynamic initialization of an object with an expression involving a reference to itself, this type of initialization is not permitted with thread local objects. For example:

  ```
  #define Thread  __declspec( __thread )
  Thread int tls_i = tls_i;            // C and C++ error
  int j = j;                           // Okay in C++; C error
  Thread int tls_i = sizeof( tls_i )   // Okay in C and C++
  ```

  Note that a `sizeof` expression that includes the object being initialized does not constitute a reference to itself and is allowed in C and C++.

Creating a DIGITAL UNIX Consolidated CD-ROM J

J.1 Introduction

This appendix describes how to create a consolidated CD-ROM for DIGITAL UNIX Version 4.0E. A consolidated CD–ROM allows you to simultaneously upgrade the processor firmware while installing DIGITAL UNIX.

J.1.1 Description

This release includes the documentation (`disklabel`(8) and `mkisofs`(8) reference pages) and utilities that you need to construct a consolidated DIGITAL UNIX CD–ROM in ISO 9660-compliant Rock Ridge format.

Although support for booting a consolidated distribution ISO 9660-compliant Rock Ridge format CD–ROM was added in DIGITAL UNIX Version 4.0D, the tools to create the media were not supplied until Version 4.0E was released, making Version 4.0E the minimum revision for creating a consolidated CD–ROM.

J.1.2 Requirements

To build a DIGITAL UNIX consolidated CD–ROM, the following is required:

- The DIGITAL UNIX Version 4.0E Operating System CD–ROM in UFS format
- The DIGITAL UNIX Documentation CD-ROM in CDFS format
- The appropriate Alpha System firmware update CD-ROM

J.2 Build Instructions

The following sections explain how to prepare and build the DIGITAL UNIX consolidated CD–ROM.

J.2.1 Preparing for the Build

After you receive a new kit, perform the following steps to move the necessary files from the CD–ROM to working directories on the build machine.

Note

The examples in this document use the C shell.

1. Use the `disklabel` utility to set up a 635 Mb partition on a spare disk, starting at block 0, with a size of 1300480 512-byte blocks and a file system type of `unused`. Create a mount point for this partition (such as `/cdimage`) to use later. For example, enter the following commands to set partition `c` of device `rz6` to start at offset 0 with a size of 1300480, and create mount point `/cdimage`:

```
% disklabel -F -r -e rz6
write new label [y] y
% mkdir /cdimage
```

2. Mount the DIGITAL UNIX Version 4.0E Operating System CD–ROM to a temporary mount point (such as `/mnt`) and use the `tar` utility to copy the contents of the CD–ROM onto a suitably large directory on the system (at least 1.5Gb). Once this is done, unmount the CD–ROM. For example, using `/spare` as the target directory and `rz4` as the CD–ROM drive, enter the following commands

```
% mount -r /dev/rz4c /mnt
% cd /mnt
% tar cf /spare/digital_unix_v40e.tar .
% cd /
% umount /mnt
```

Note

This step may take as long as 60 minutes to complete.

J.2.2 Building the Consolidated CD–ROM

Once you have completed the steps in Section J.2, perform the following steps to consolidate the necessary data to a single CD–ROM in ISO 9660-compliant Rock Ridge format:

1. Use the `newfs` command to initialize a file system on the partition reserved in step 1 of Section J.2 , and mount it to the mount point `/cdimage`. If you are prompted for confirmation, enter `y`. Use the `tar` utility to copy the DIGITAL UNIX base operating system image created in step 2 of Section J.2

to /cdimage. For example, using /spare as the source and rz6c as the target partition, enter the following commands:

```
% newfs /dev/rz6c
% mount /dev/rz6c /cdimage
% cd /cdimage
% tar xf /spare/digital_unix_v40e.tar
% cd /
```

Note

This step may take as long as 60 minutes to complete.

2. Copy the firmware images to the target directory. This is a multipart operation:

   A. Mount the appropriate Alpha System firmware update CD–ROM to a temporary mount point such as /mnt. For example, using /dev/rz4c as the CD-ROM drive enter the following command:

   ```
   % mount -t cdfs -r /dev/rz4c /mnt
   ```

   B. Copy the SMM table from the appropriate Alpha System firmware update CD–ROM to the target directory.

   ```
   % cp /mnt/SMMTABLE.TXT;1 /cdimage/smmtable.txt
   ```

   Caution

   The target file name must be in lowercase with the ‘‘ ;1 ŕemoved from the end.

   C. Look in the SMM table to find the name and locations of the firmware images to be copied as follows:

   ```
   % more /mnt/SMMTABLE.TXT;1
   ```

   The entry for the EV5 AlphaServer 1000A platform is similar to this (the actual table entry is on one line):

   ```
   27      5       1270,1311,1558,1580-1581        \
   [ALPHA1000A]AS1000A_E5_V5_1.EXE;1       \
   6       5.1     !       AlphaServer 1000A 5/xxx
   ```

   In this case, the firmware file on the CD–ROM is AS1000A_E5_V5_1.EXE;1.

   D. Create the appropriate firmware directories in the target directory, and copy each of the platform firmware images that you want from the appropriate Alpha System firmware update CD–ROM.

Caution

The target file name must be in lower case with the ";1" removed from the end. Otherwise, the `fwupgrade` program cannot locate the firmware images. If the source file is `AS1000A_E5_V5.1.EXE;1`, the target file should be `as1000a_e5_v5_1.exe`.

For example, using the file names on the appropriate Alpha System firmware update CD–ROM:

```
% mkdir /cdimage/alpha800
% mkdir /cdimage/alpha1000a
% cp /mnt/ALPHA800/AS800_V5_1.EXE;1 \
/cdimage/alpha800/as800_v5_1.exe
% cp /mnt/ALPHA1000A/AS1000A_E5_V5_1.EXE;1 \
/cdimage/alpha1000a/as1000a_e5_v5_1.exe
```

E. Unmount and remove the Firmware CD–ROM.

```
% umount /mnt
```

3. Use the `mkisofs` program to build the target CDFS file image of the directory structure in `/cdimage.` For example, using `/spare` as the target location for the image, enter the following command:

```
% /usr/bin/mkisofs -D -R -a -d -o \
/spare/consolidated_digital_unix.cdfs /cdimage/
```

Refer to the `mkisofs`(8) reference page for information about using this program.

4. Use the `disklabel` program to insert a label into the file generated in step 3.

```
% disklabel -r -w -t cdfs -f \
/spare/consolidate_digital_unix.cdfs \
/mdec/rzboot.cdfs /mdec/bootrz.cdfs
```

Refer to the `disklabel`(8) reference page for information about using this program.

5. The CD image file `/spare/consolidate_digital_unix.cdfs` is ready to be loaded onto a CD.

J.3 Sample Build Session

The following assumptions are made for the examples in the following sections:

- The target partion is on `/dev/rz6c`.
- The `/spare` directory has at least 1.5 Gb of free space.

- The CD–ROM drive is `/dev/rz4`.

The examples in this document use the C shell.

J.3.1 Sample Preparation

This section shows an example of preparing for a build.

1. Log in as root:

```
% cd /
% disklabel -F -r -e rz6
write new label? [y] y
% mkdir /cdimage
```

2. Place the DIGITAL UNIX Version 4.0E Operating System CD–ROM in the CD–ROM drive:

```
% mount -r /dev/rz4c /mnt
% cd /mnt
% tar cf /spare/digital_unix_v40e.tar .
% cd /
% umount /mnt
```

3. Remove the DIGITAL UNIX Version 4.0E Operating System CD–ROM from the drive.

The preparatory steps are complete.

J.3.2 Sample CD–ROM Build

This section shows an example of building a CD–ROM.

1. Log in as root:

```
% cd /
% newfs /dev/rz6c
% mount /dev/rz6c /cdimage
% cd /cdimage
% tar xpf /spare/digital_unix_v40e.tar
% cd /
```

2. Place the appropriate Alpha System firmware update CD–ROM in the CD–ROM drive:

```
% mount -t cdfs -r /dev/rz4a /mnt
% cp /mnt/SMMTABLE.TXT;1 /cdimage/smmtable.txt
% more /cdimage/smmtable.txt
```

3. Determine the appropriate directory and file names for the firmware

images that you want:

```
% mkdir /cdimage/alpha800
% mkdir /cdimage/alpha1000a
% cp /mnt/ALPHA800/AS800_V5_1.EXE;1 \
/cdimage/alpha800/as800_v5_1.exe
% cp /mnt/ALPHA/AS1000A_E5_V5_1.EXE;1 \
/cdimage/alpha1000a/as1000a_e5_v5_1.exe
% umount /mnt
% /usr/bin/mkisofs -D -R -a -d -o \
/spare/consolidated_digital_unix.cdfs /cdimage/
Using OSFMANWO.000;1 for
/cdimage/ALPHA/BASE/instctrl/OSFMANWOS425.scp
(OSFMANWOP425.scp)
Using OSFMANWO.001;1 for
/cdimage/ALPHA/BASE/instctrl/OSFMANWOS425.inv
(OSFMANWOP425.inv) Using OSFMANWO.002;1 for
/cdimage/ALPHA/BASE/instctrl/OSFMANWOS425.ctrl
(OSFMANWOP425.ctrl)
   .
   .
   .
Using PROCFS_V.000;1 for
/cdimage/usr/sys/procfs/procfs_vnops_stubs.c
(procfs_vfsops_stubs.c)
3.92% done, estimate finish Thu Oct 23 15:36:59 1997
5.87% done, estimate finish Thu Oct 23 15:39:24 1997
   .
   .
   .
99.74% done, estimate finish Thu Oct 23 15:41:52 1997
Total extents actually written = 255673
Total translation table size: 0
Total rockridge attributes bytes: 2066594
Total directory bytes: 4239360
Path table size(bytes): 10130
Max brk space used b9ec60
255673 extents written (499 Mb)
% /disklabel -r -w -t cdfs -f \
/spare/consolidate_digital_unix.cdfs \
/mdec/rzboot.cdfs /mdec/bootrz.cdfs
```

The information is consolidated, and the file can be burned onto a CD–ROM.

Index

A

B

C

D

E

J

K

L

M

N

O

P

Q

R

S

T

U

V

X

Y

How to Order Additional Documentation

Technical Support

If you need help deciding which documentation best meets your needs, call 800-DIGITAL (800-344-4825) before placing your electronic, telephone, or direct mail order.

Electronic Orders

To place an order at the Electronic Store, dial 800-234-1998 using a 1200- or 2400-bps modem from anywhere in the USA, Canada, or Puerto Rico. If you need assistance using the Electronic Store, call 800-DIGITAL (800-344-4825).

Telephone and Direct Mail Orders

Your Location	Call	Contact
Continental USA, Alaska, or Hawaii	800-DIGITAL	Digital Equipment Corporation P.O. Box CS2008 Nashua, New Hampshire 03061
Puerto Rico	809-754-7575	Local Digital subsidiary
Canada	800-267-6215	Digital Equipment of Canada Attn: DECdirect Operations KAO2/2 P.O. Box 13000 100 Herzberg Road Kanata, Ontario, Canada K2K 2A6
International	————	Local Digital subsidiary or approved distributor
Internal[a]	————	SSB Order Processing – NQO/V19 *or* U. S. Software Supply Business Digital Equipment Corporation 10 Cotton Road Nashua, NH 03063-1260

[a] For internal orders, you must submit an Internal Software Order Form (EN-01740-07).

Reader's Comments

DIGITAL UNIX
Release Notes for Version 4.0E
AA-QTLMD-TE

Digital welcomes your comments and suggestions on this manual. Your input will help us to write documentation that meets your needs. Please send your suggestions using one of the following methods:

- This postage-paid form
- Internet electronic mail: `readers_comment@zk3.dec.com`
- Fax: (603) 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)	☐	☐	☐	☐
Completeness (enough information)	☐	☐	☐	☐
Clarity (easy to understand)	☐	☐	☐	☐
Organization (structure of subject matter)	☐	☐	☐	☐
Figures (useful)	☐	☐	☐	☐
Examples (useful)	☐	☐	☐	☐
Index (ability to find topic)	☐	☐	☐	☐
Usability (ability to access information quickly)	☐	☐	☐	☐

Please list errors you have found in this manual:

Page	Description
_____	_____
_____	_____
_____	_____
_____	_____
_____	_____

Additional comments or suggestions to improve this manual:

_____
_____
_____
_____
_____

What version of the software described by this manual are you using? _____

Name/Title _____ Dept. _____
Company _____ Date _____
Mailing Address _____
_____ Email _____ Phone _____

Do Not Cut or Tear − Fold Here and Tape

digital™

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

DIGITAL EQUIPMENT CORPORATION
UEG PUBLICATIONS MANAGER
ZKO3−3/Y32
110 SPIT BROOK RD
NASHUA NH 03062−9987

Do Not Cut or Tear − Fold Here

Cut on
Dotted
Line