Digital UNIX (formerly DEC OSF/1)

Release Notes

Order Number: AA-PS2BF-TE

July 1995

Product Version: Digital UNIX Version 3.2C

This book contains notes on software and documentation restrictions for Digital UNIX Version 3.2C.

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

Digital Equipment Corporation
Maynard, Massachusetts

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

Tables

About This Manual

This manual contains release notes for Digital UNIX® Version 3.2C. The release notes are available in both hardcopy and Bookreader versions. See Appendix A for any differences between the online and hardcopy versions.

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.

Digital has changed the name of its UNIX operating system from DEC OSF/1 to Digital UNIX. The new name reflects Digital's commitment to UNIX and its conformance to UNIX standards, including the Single UNIX Specification (formerly called SPEC 1170).

Audience

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

Organization

This manual contains the following:

Chapter 1	Contains an overview of new and changed features. This chapter is also part of the Bookreader book *New and Changed Features*.
Chapter 2	Contains installation notes and subset sizes.
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. This chapter is also part of the Bookreader book *New and Changed Features*.
Appendix A	Contains release notes that are included in the hardcopy but not in the online versions, if any.

Related Documents

You should have the following documentation available:

- The hardware documentation for your system
- The online or hardcopy reference pages
- The Bookreader book files that are provided on the CD–ROM

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	Teal
System Administrators	S	Red
Network Administrators	N	Yellow
Programmers	P	Blue
Reference Page Users	R	Black

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

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

Reader's Comments

Digital welcomes your comments on this or any other Digital UNIX manual. A Reader's Comment form is located in the back of each printed Digital UNIX manual and on line in the following location:

`/usr/doc/readers_comment.txt`

You can send your comments in the following ways:

- Internet electronic mail: `readers_comment@zk3.dec.com`
- Fax: 603-881-0120 Attn: UEG Publications, ZK03-3/Y32
- Mail:

  Digital Equipment Corporation

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

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

If you have suggestions for improving particular sections or find any errors, please indicate the manual title, order number, and section numbers. Digital also welcomes general comments.

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

This chapter provides brief descriptions of features that are new to the Digital UNIX operating system in this release or have changed significantly from previous releases. You should also refer to Chapter 8, which describes features and interfaces scheduled to be retired in future versions.

This chapter is also part of the Bookreader book *New and Changed Features*, which is available on Disc 2 of the distribution CD–ROM.

1.1 Digital UNIX Version 3.2C

Digital UNIX Version 3.2C is a new version of the operating system supporting the full line of Alpha systems and providing binary compatibility for applications currently running on Digital UNIX Version 3.2.

This version of Digital UNIX supports an update installation from Digital UNIX Version 3.2, Version 3.2A, and Version 3.2B. In addition, this version of Digital UNIX includes maintenance fixes, adds support for new hardware systems and options, new functionality, and provides system resource utilization improvements that are beneficial to low-end, memory constrained applications.

Digital UNIX Version 3.2C supports a 24 MB Desktop Workstation, a newly packaged variant of the AlphaStation 200 4/100.

1.2 UNIX Brand and Product Name Change

Digital is now licensed under the X/Open UNIX branding program to use the UNIX trademark, and has received the UNIX Brand for its operating system. To reinforce the UNIX Brand and to better reflect the true nature of Digital's standards-compliant UNIX product, Digital is officially changing the name of its UNIX operating system from DEC OSF/1 to Digital UNIX. This product name change will be gradually phased into our product binaries and documentation.

1.3 Support for 24 MB Desktop Workstations

Digital UNIX provides support for workstations with 24 MB of RAM and 535 MB of disk space. Currently, the only configuration available with 24 MB of RAM is the new packaging variant of the AlphaStation 200 4/100,

which includes a 535 MB disk.

From the operating system standpoint, a system with 24 MB of RAM is similar to other Digital UNIX systems, with a few minor exceptions to enable optimal performance on lower memory systems.

For more information, see Section 2.2.1, Section 4.2.8, the *Installation Guide*, and the *System Tuning and Performance Management* manual.

1.4 New Hardware

The following new processors are now supported:

- AlphaStation 200 Model 4/133
- AlphaStation 600 5/300
- Digital Alpha VME 2100 Series
- AXPvme 100, 166, 230

The following new hardware options are now supported:

- Network Devices:
  - DGLPB (PCI ATM)
  - DE500 (PCI Fast Ethernet)
- PCI enhancements
  - Conformance with the PCI V2.1 specification
  - Support for user tunable refresh rate on the ATI Mach64 PCI option card
  - Enhancements to the PCI interrupt support

1.5 Digital AlphaStation 200 Model 4/100

The Digital AlphaStation 200 Model 4/100 is the entry-level PCI-based desktop client system with 24 Mbytes of memory. It uses the DECchip 21064 processor running at 100 MHz with performance estimated at 74 SPECint92 and 92 SPECfp92. The high-performance PCI I/O bus, running at 132 Mbytes per second, provides expansion for options such as high-performance graphics, networking, and SCSI adapters. Lower performance options are supported on the ISA bus. It supports a wide variety of industry-standard peripherals and PCI/ISA options. Other standard features include integrated Twisted Pair and ThinWire Ethernet, stereo-quality audio, and an array of external ports for serial and parallel communications and external SCSI.

For more information, see Section 4.2.8 and the *Installation Guide*.

1.6 New Single-Board Computer Platforms

The Digital AXPvme 100, Digital AXPvme 166, and Digital AXPvme 230 Single-Board Computers (SBCs) are follow-on products to the Digital AXPvme 64 and AXPvme 160 SBCs. For realtime and embedded applications, these new SBCs, along with their predecessors, provide an interface to the VME32 I/O bus and provide D64 support for block-mode DMA transfers. These SBCs use the new DECchip 21066A processor running at 99 MHz, 165 Mhz, and 231 Mhz. These SBCs provide improvements over the AXPvme 64 and AXPvme 160 SBCs as follows:

- Backup cache size doubled (512 Kb)
- SCSI clock speed increased, providing a slight performance increase
- Increased Flash ROM size by 4 (4 Mb)
- Increased size of PCI Mezzanine Card area

For information on how to configure the VME subsystem, see the *Writing VMEbus Device Drivers* manual.

1.7 Digital Alpha VME 2100 Series

The Digital Alpha VME 2100 series is a variation of the AlphaServer 2100 family. For realtime and embedded applications, it offers the VME32 I/O bus in place of the EISA bus that is part of other 2100 servers. For increased performance, VMEbus D64 transfers are supported during block-mode DMA. It is a high-performance, scalable SMP system with three PCI and five VME32 industry-standard I/O slots for user options.

For more information, see Section 3.3.2

For information on how to configure the VME subsystem, see the *Writing VMEbus Device Drivers* manual.

1.8 Monitoring Performance History

The Monitoring Performance History (MPH) utility gathers timely and accurate information on the reliability and availability of the Digital UNIX operating system and associated platforms.

The MPH utility was previously included on the Complementary Products CD-ROM. This utility is now included on the Digital UNIX Operating System CD-ROM.

For more information, see the *Installation Guide*.

1.9 NetWorker SingleServer Save and Restore

The NetWorker SingleServer Save and Restore utility has been updated from Version 3.0A to Version 3.1. The new version includes enhancements to the graphical user interface and maintenance enhancements. For more information, see the *Installation Guide*.

1.10 Customer Log Desk Problems Resolved

A number of Customer Log Desk (CLD) problems have been resolved in this release. These corrections are documented in an ASCII text file. The file is located in the *mnt_point*/DOCUMENTATION/TEXTdirectory on Disc 1 of the distribution media.

Installation Notes 2

The notes in this chapter discuss the following topics:

- General information about installation
- Layered product considerations
- Full installation
- Update installation
- Server extensions
- Disk revisions
- Disk space requirements

Caution

Do not attempt to install Digital UNIX Version 3.2C without first reading in Chapter 3 the notes appropriate to your processor. Failure to read these notes can result in serious installation problems.

2.1 General Information About Installation

The following notes apply to the installation in general.

2.1.1 PCI NVRAM

Digital UNIX supports PCI NVRAM for the Digital AlphaServer 1000 4/200, Digital AlphaServer 2100 Models A500 and A600, and the Digital AlphaServer 2000 series systems.

Multiple NVRAM modules are not supported. If your system has a PCI NVRAM module, it cannot have a second NVRAM module (EISA or PCI).

2.1.2 Asynchronous Transfer Mode Support

You should install the DGLTA Asynchronous Transfer Mode (ATM) adapter prior to installing Digital UNIX Version 3.2C so that the installation process properly detects the presence of the ATM adapter and installs the ATM

software. If you install the ATM adapter after installing Digital UNIX, then you will have to boot the system off the generic kernel (genvmunix), install the ATM software sets from the CD–ROM, and run the `doconfig` program to rebuild your system kernel to enable the added ATM support.

2.1.3 Installing on a Digital UNIX SCSI RAID Device

If you are installing the software on a SCSI RAID device, only LUN 0 can be used for the installation's target device. The RAID device's LUN 0 must be configured before you attempt to install the software onto it.

2.1.4 UUCP Suite Problems and Restrictions

Because of problems in `uucpsetup`, you must make the following manual edits:

- If you define an incoming UUCP system and your system uses NIS, make sure that the line containing the wildcard `+:` is the last line in the `/etc/passwd` file.
- If you define an outgoing TCP system, you must edit the `/etc/uucp/Systems` file to provide correct information if the *`host.domain`* portion of that system's fully qualified name contains more than 14 characters.

In addition, the `uugetty` command should run only on RS232 lines (not printer or console lines).

2.1.5 ATI Graphics Controllers

An address conflict occurs between ISA serial devices configured as COM4 and ATI Mach64 graphics controllers. Both ISA GX and PCI ATI Mach64 CX/GX/CT products use the address range of 2EC-2EF. This is the same address range assigned to COM4 devices.

2.1.6 Wide SCSI Drives

There is a problem affecting wide SCSI drives attached to the QLogic adapters.

If power is removed from a wide SCSI drive on a running system and then reapplied and the first SCSI command sent to the drive after power is reapplied is a SCSI Inquiry command, the drive is no longer seen by the system and that SCSI target number becomes unusable until after the system is rebooted or the affected SCSI bus is reset. Issuing a command such as `file` or `disklabel` after a wide drive has been power-cycled on a running system exhibits this behavior.

2.2 Layered Product Considerations

This section provides information related to requirements and minimum supported versions for Digital layered products.

The new versions of the layered products shown in Table 2-1 were specifically developed to take advantage of new Digital UNIX features, and to provide new hardware support or resolutions to problems. The table shows the correct layered product versions to use with Digital UNIX Version 3.2C.

Table 2-1: Layered Product Versions

Layered Product	DEC OSF/1 Version 3.0	DEC OSF/1 Version 3.2	Digital UNIX Version 3.2C
POLYCENTER AdvFS Utilities	Version 3.0	Version 3.2	Version 3.2
DECsafe Available Server	Version 1.1	Version 1.2	Version 1.2A
Logical Storage Manager	Version 1.1	Version 1.2	Version 1.2
System V Environment	Version 3.0	Version 3.2	Version 3.2
DECnet	Version 3.0	Version 3.0	Version 3.2
SNA Peer	None	Version 1.2	Version 1.2 with ECO 1

See Section 2.2.1, Section 4.2.8, and Section 4.2.9 for information on reduced memory configurations.

2.2.1 Reduced Memory Systems

Some layered products may have reduced performance on systems with less than 32 MB of memory. See the layered product SPD for memory requirements.

See Section 4.2.8, Section 4.2.9, and Section 4.2.9 for information on reduced memory configurations.

2.2.2 DECnet/OSI

DECnet/OSI Version 3.2 is the minimum required version for use with Digital UNIX Version 3.2C. There is an ECO 1 to DECnet/OSI Version 3.2. You should install this patch before using DECnet/OSI Version 3.2.

2.2.3 Certain Graphic Adapters Require DEC Open3D

Workstation configurations that contain the DENALI (PEXGA) graphic adapter require the additional installation of DEC Open3D Version 2.4 or higher after new installations of Digital UNIX Version 3.2C. Without the DEC Open3D software, the X Window System will not start on these systems. All other graphic adapters are supported by Digital UNIX Version 3.2C.

2.2.4 Distributed Computing Environment Installation

When loading the Distributed Computing Environment (DCE) Version 1.3 onto Digital UNIX Version 3.2C, ignore the following message:

```
chmod: Symbolic link cannot be resolved
```

DCE does not require the `chmod` calls.

2.3 Full Installation

A full installation will overwrite any customization on your system.

Note

If you have already installed Digital UNIX (formerly DEC OSF/1) Version 3.2, 3.2A, or 3.2B, you should do an update installation. See Section 2.4.

2.3.1 Full Install Overwrites /etc/vol/volboot

If your system has LSM volumes configured on it, save a copy of the `/etc/vol/volboot` file before doing a full installation. Doing a full installation overwrites the `/etc/vol/volboot` file, resulting in the potential loss of data under LSM control.

2.3.2 Full Install Overwrites /etc/prestotab

If your system has Prestoserve configured on it, save a copy of the `/etc/prestotab` before doing a full installation. The installation process causes this file to be overwritten, resulting in the potential loss of configuration data contained in this file.

2.3.3 Full Install Overwrites /etc/lvmtab

If your system has Logical Volume Manager (LVM) volumes configured on it, save a copy of the `/etc/lvmtab` file before doing a full installation. The installation process causes this file to be overwritten, resulting in the

potential loss of all data under LVM control.

2.3.4 Full Install Overwrites /etc/fdmns Directory

If your system is configured with AdvFS, a full installation overwrites the `/etc/fdmns` directory, resulting in the loss of critical configuration data. Save a copy of the `/etc/fdmns` directory and all files under the directory before beginning the installation. After installation, restore the directory.

2.4 Update Installation

The following notes apply to the update installation procedure.

2.4.1 Logical Volume Manager Requires a Kernel Rebuild

If updating the Logical Volume Manager (LVM), you must build a new kernel by using the `doconfig` utility after the update installation has completed. This is because the LVM kernel option is optional for Digital UNIX Version 3.2C. To rebuild the kernel, run the `doconfig` utility and select the Logical Volume Manager option from the KERNEL OPTION SELECTION menu. After the kernel build has completed, copy the new kernel to `/vmunix` and reboot your system.

2.4.2 The installupdate Command May Select the Wrong Device

The `installupdate` command may select the wrong distribution media when multiple distribution medias are mounted on mount points with similar names. For example, if the following devices are mounted prior to beginning an update installation and issue a `/sbin/installupdate /mnt`, `installupdate` will incorrectly select `/dev/rz4c` as the distribution media:

```
/dev/rz0a on / type ufs (rw)
/proc on /proc type procfs (rw)
/dev/rz0g on /usr type ufs (rw)
/dev/rz4c       /mnt1
/dev/rz11c      /mnt
```

This is caused by `grep` for `/mnt` that returns the `/mnt1` mount point first.

The workaround for this problem is to mount only the intended distribution media prior to beginning the update.

2.4.3 Ignore grep Messages During Update Installation

You can ignore the following messages if they are displayed during an

update installation; the installation will complete successfully:

```
grep: can't open ./usr/lib/X11/app-defaults/.new..DXMail
grep: can't open ./usr/lib/X11/app-defaults/DXMail
```

2.5 Server Extensions

This section provides installation information for Digital UNIX Server Extensions.

2.5.1 Remote Installation Services

The following notes apply to installations that use Remote Installation Services (RIS).

2.5.1.1 Installation CD-ROM Contains a ''Dirty'' File System

The Digital UNIX Version 3.2C CD–ROMs contain a ''dirty'' file system. To mount the CD–ROM on an ULTRIX system so that you can load the software for RIS use, you must specify the `–o force` flag for the `mount` command. For example:

```
# mount -o force -r /dev/rz4c /cdrom
```

To mount the CD–ROM on a DEC OSF/1 system older than Version 2.0, use the following command:

```
# mount -dr /dev/rz4c /cdrom
```

2.5.1.2 ULTRIX Server or Digital UNIX Server Running Older Software

If you are using an ULTRIX system as a RIS server for Digital UNIX Version 3.2C, you must perform some extra tasks before clients can boot the software. Use the RIS install software option to extract or link to the Digital UNIX Version 3.2C kit. Then enter the following commands, where *n* is the number of the appropriate directory:

```
# cd risn.mips
# restore xf */ROOT ./RisFiles
# /usr/bin/sh5 RisFiles Extract */ROOT
```

If your RIS server is running ULTRIX or DEC OSF/1 Version 1.2, 1.2A, 1.3, 1.3A, or 1.3B software, an extra step is needed before you can install Digital UNIX Version 3.2C on your clients. Because the Version 3.2C client systems use NFS to access part of the kit from the server, NFS server capabilities must be enabled on the server. (See the `nfssetup`(8) reference page.)

Before installing clients, you must add a line to the /etc/exports file for each RIS area that your ULTRIX server or DEC OSF/1 RIS server running Version 1.2, 1.2A, 1.3, 1.3A, or 1.3B software is serving to Digital UNIX clients. The /etc/exports entry includes a pathname to the kit directory in the RIS area; there must be no symbolic links in this pathname. Entries in the /etc/exports file can have either of the following formats:

- If your server's var area is separate from the usr area, there will be a symbolic link to it at /usr/var. In this case, add a line similar to the following to the /etc/exports file:
  - If your RIS server is an ULTRIX server running a version earlier than Version 4.4, add the following line:

    ```
    /var/adm/ris/ris2.mips/kit -r=0 -o
    ```

  - If your RIS server is running ULTRIX Version 4.4 or is a DEC OSF/1 RIS server running Version 1.2, 1.2A, 1.3, 1.3A, or 1.3B software, add the following line:

    ```
    /var/adm/ris/ris2.alpha/kit -r=0 -o
    ```

- If your var area is part of your usr area, add a line similar to the following to the /etc/exports file:
  - If your RIS server is an ULTRIX server running a version earlier than Version 4.4, add the following line:

    ```
    /usr/var/adm/ris/ris2.mips/kit -r=0 -o
    ```

  - If your RIS server is running ULTRIX Version 4.4 or is a DEC OSF/1 RIS server running Version 1.2, 1.2A, 1.3, 1.3A, or 1.3B software, add the following line:

    ```
    /usr/var/adm/ris/ris2.alpha/kit -r=0 -o
    ```

  After you make the entries in the /etc/exports file, your ULTRIX server will be able to register clients for the Digital UNIX product.

Caution

Before you allow clients to use the new RIS area to perform update installations, read the *Release Notes* and the *Installation Guide*. Failure to carefully follow the directions for update installations can result in loss or corruption of data.

2.6 Required Revisions of RZ28B and RZ25L Disk Drives

Very heavy I/O loads may cause errlog entries (for example, entries indicating timeouts followed by aborts) or SCSI bus hangs. These problems

may occur if you have either an RZ28B disk drive at revision 0003 or lower or an RZ25L disk drive at revision 0007 or lower.

If you have a disk drive at one of these revisions and your system generates large amounts of I/O, then Digital recommends that you apply the following change to your system:

1. Edit the `/sys/data/cam_data.c` file as follows:
   a. Search the file for RZ25L and RZ28B. You will see lines like the following:

```
{"DEC RZ25L", 13, DEV_RZ25L, (ALL_DTYPE_DIRECT << DTYPE_SHFT) |
SZ_HARD_DISK,
  (struct pt_info *)ccmn_rz25l_sizes, 512, DEC_MAX_REC, NO_DENS_TAB,
  NO_MODE_TAB, (SZ_BBR | SZ_REORDER),   <<<<<<<<Change this
line>>>>>>>>>
  NO_OPT_CMDS, SZ_READY_DEF, 64, DD_REQSNS_VAL | DD_INQ_VAL,
  36, 64
},
```

   b. Change the line indicated in the example by `<<<<<<<<Change this line>>>>>>>>` to read as follows:

```
NO_MODE_TAB, (SZ_BBR | SZ_REORDER | SZ_NO_TAG),
```

2. Rebuild the kernel and reboot the system.

2.7 Disk Space Requirements

Information on the disk space requirements is provided in the *Installation Guide*.

Processor-Specific Notes 3

This chapter contains notes that apply to the following Digital Alpha processors:

- General notes on processors
- Digital AlphaSever 1000 processors
- Digital AlphaServer 2000, AlphaServer 2100, and Digital VME 2100 processors
- DEC 7000 processors
- Digital AlphaServer 8000 processors
- Digital AXPvme Single-Board Computers
- Digital AXPpci 33 processors

3.1 General Notes on Processors

The notes in this section apply to various different processors.

3.1.1 CD-ROM Device ID

Make sure that the device ID displayed for the built-in CD–ROM drive is RRD42 or RRD43. The drive has a jumper named PARITY, which should be installed. The Digital UNIX system may not work properly unless the device ID is set to RRD42 or RRD43. See the RRD42 or RRD43 documentation for more information.

3.1.2 Token Ring Option Setup

To use the DW110 ISA Token Ring option, configure the device at the console prompt as follows:

```
>>> isacfg -slot slot-num -dev 0 -mk -handle
DW110-cable_type.ring_speed -irq0 9 -iobase0 a20 -dmamode0 8
-dmachan0 5 -etyp 1 -enadev 1
```

Note that the `isacfg` command must be entered as a single line.

Use the following parameter values:

slot_num
The number of an unused `isacfg` slot.

cable_type
Either `UTP` or `STP`.

ring_speed
Either 4 or 16.

3.1.3 ATI MACH64 CX PCI 2-MB DRAM Graphics Card

The following instructions and restrictions apply to the ATI MACH64 CX PCI 2-MB DRAM graphics card (part number PB2GA-FA) when used with Digital UNIX.

The ATI MACH64 CX PCI 2-MB DRAM card will come up in 1024x780 mode. If you want to run in 1280x1024 mode, modify the `/usr/var/X11/xdm/Xservers` file as follows:

1. Log in as root on the serial console.
2. Change directories as follows:

```
# cd /usr/var/X11/xdm
```

3. Edit the `Xservers` file and change the last line to contain the following:

```
:0 local /usr/bin/X11/X -nice -2 -screen 1280
```

4. Apply the modification by issuing the following commands:

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

3.1.4 Console Mode

When you turn on the system, the console variable stored in the NVRAM determines the console mode. For example, if the console variable is set to graphics, the following output is displayed:

```
>>> show console
    graphics
```

You can set the console variable so that it will always go to the COMM1

port by issuing the following command at the firmware prompt:

```
>>> set console serial
>>> init
```

If the console variable is set to graphics and if the keyboard is plugged in, the console prompt is displayed on the graphics screen. If the keyboard is not plugged in, the console defaults to serial line.

You should have a serial terminal operating at 9600 baud, 8 bits, no parity connected by the COMM1 port to use the serial console.

3.1.5 Keyboard Restriction

If the keyboard is not plugged in when you issue the `boot` command at the console prompt (`>>>`), Digital UNIX cannot bring the keyboard on line.

3.2 Digital AlphaServer 1000

The following notes apply to the Digital AlphaServer Model 1000 4/200.

3.2.1 Configuring a Cirrus GD5422 Graphics Chip into a System

The Digital AlphaServer 1000 4/200 has a built-in Cirrus GD5422 graphics chip. Digital UNIX supports this chip at 800x600 resolution with 256 colors and at 640x480 resolution with 16 colors, which is the default.

Verify that the Cirrus GD5422 graphics chip is on the system.

To configure the chip into the system, perform the following steps:

1. Be sure the Digital UNIX EISA Configuration Utility diskette is loaded. From the SRM console, enter the `runecu` command to invoke the EISA Configuration Utility.

   A message is displayed indicating that you have invoked the EISA Configuration Utility.
2. Press the Enter key to continue. A list of steps for configuring your system is displayed.
3. Use the arrow keys to go to Step 3, View or edit details, and press the Enter key.
4. If the built-in VGA Graphic Controller is disabled, use the arrow keys to highlight the VGA Graphic Controller option and continue with the next step. Otherwise, go to Step 6
5. Press the Enter key to edit this option. Use the arrow keys to enable the option, then press the F10 key.
6. Press the F10 key again to exit from Step 3, View or edit details.

7. Use the arrow keys to go to Step 5, Save and exit, and press the Enter key.
8. Verify that the Save the configuration and restart the computer line is highlighted and press the Enter key.

   The configuration is now saved.
9. Press the Enter key to exit from the EISA Configuration Utility.
10. When the Boot menu is displayed, turn the system off and on again to reenter the SRM console.

3.2.2 Removing a Cirrus GD5422 Graphics Chip from a System

If you want to use a different graphics card, you must disable the Cirrus GD5422 graphics chip prior to adding the new graphics card to the system.

The following steps describe how to disable the built-in graphics chip in the saved system configuration:

1. Make sure the EISA Configuration Utility diskette is loaded. From the SRM console, enter the `runecu` command to invoke the EISA Configuration Utility.

   A message is displayed indicating that you have invoked the EISA Configuration Utility.
2. Press the Enter key to continue.
3. Use the arrow keys to select Step 3, View or edit details, and press the Enter key.
4. Use the arrow keys to highlight the VGA Graphic Controller option. Press the Enter key to edit the function.
5. Disable the VGA Graphic Controller by using the arrow keys and pressing the F10 key.
6. Press the F10 key to exit from Step 3, View or edit details.
7. Use the arrow keys to go to Step 5, Save and exit, and press the Enter key.
8. Verify that the Save the configuration and restart the computer line is highlighted, and press the Enter key.

   The configuration is now saved.
9. Press the Enter key to exit from the EISA Configuration Utility.
10. When the Boot menu is displayed, turn off the system.

While the system is off, add the new graphics option card. Digital UNIX Version 3.2C supports the PB2GA-AA EISA Compaq Qvision 1024/E card and the PB2GA-HA PCI 1280/P option cards.

Also disable the built-in Cirrus GD5422 graphics hardware while the system is turned off. Refer to the hardware documentation that came with your system for instructions on adding option cards and disabling the Cirrus GD5422 graphics hardware.

Once the built-in graphics chip is disabled and the optional graphics card added, turn on the system. Run the EISA Configuration Utility to configure the PB2GA-AA EISA Compaq Qvision 1024/E card. If you are using the PB2GA-HA PCI Compaq 1280/P card, you do not need to rerun the EISA Configuration Utility.

To configure the EISA Compaq Qvision 1024/E card, perform the following steps:

1. Be sure the Digital UNIX EISA Configuration Utility diskette is loaded. From the SRM console, enter the `runecu` command to invoke the EISA Configuration Utility.

   A message is displayed indicating that you have invoked the EISA Configuration Utility.

2. Press the Enter key to continue.

   A list of steps to perform your system configuration is displayed.

3. Use the arrow keys to go to Step 5, Save and exit, and press the Enter key.

4. Verify that the Save the configuration and restart the computer line is highlighted and press the Enter key.

   The configuration is now saved.

5. Press the Enter key to exit from the EISA Configuration Utility.

6. When the Boot menu is displayed, turn the system off and on again to reenter the SRM console.

3.3 Digital AlphaServer 2000, AlphaServer 2100, and Digital Alpha VME 2100 Processors

The following notes apply to the AlphaServer 2000 and AlphaServer 2100 Server series processors.

3.3.1 Digital AlphaServer 2100 5/250

The following notes apply to the Digital AlphaServer 2100 5/250.

3.3.1.1 Upgrades

If you are upgrading the hardware for your Digital AlphaServer 2100 (Models 4/200 or 4/275) to a Digital AlphaServer 2100 5/250, follow the upgrade instructions that came with your hardware. You must upgrade to Digital UNIX Version 3.2C before you install the new hardware. See Chapter 2 for more information about upgrading your system to Digital UNIX Version 3.2C.

3.3.1.2 Mouse Pointer Does Not Display

If your system has a Compac Qvision graphics adapter board (PB2GA-AA), the mouse pointer may not display properly when the system is first started up. The mouse pointer will display properly, however, after you log into the system.

3.3.2 Digital Alpha VME 2100 Series

The notes in this section apply to the Digital UNIX software running on the Digital Alpha VME 2100 series.

For information on how to configure the VME subsystem, see the *Writing VMEbus Device Drivers* manual.

3.3.2.1 Unpredictable Results When Performing Master Block Transfers

Performing master block transfers with a data width of D64 can produce unpredictable results. For more information, see Section 3.6.4.

3.4 DEC 7000 and 10000 Processors

The KDM70 disk controller is not supported when more than one gigabyte of memory is installed.

3.5 Digital AlphaServer 8200 5/300 and Digital AlphaServer 8400 5/300

The following notes are for the AlphaServer 8000 processors.

3.5.1 Configuration Restrictions

The following configuration restrictions apply to the Digital UNIX Version 3.2C Digital AlphaServer 8200 5/300 and Digital AlphaServer 8400 5/300:

- The PBXGA-AA TGA video adapter is not supported.

- RIS installations that use `mop` network protocol are not supported in the V1.0-66 console firmware. Use the `bootp` network protocol for RIS installations.
- Table 3-1 shows the maximum number of RZ, RA, and RE disks that are supported on the Digital AlphaServer 8200 5/300 and Digital AlphaServer 8400 5/300:

Table 3-1: Disks

Disk Type	Maximum Number of Disks
RZ	128 (rz0 to rz128)
RA	256 (ra0 to ra256)
RE	48 (re0 to re48)

- Table 3-2 shows the maximum number of adapters that are supported on the Digital AlphaServer 8200 5/300 and Digital AlphaServer 8400 5/300:

Table 3-2: Adapters

	Maximum Adapters	Option Number	Device Name
XMI bus			
	16	KZMSA-AB	xza
	8	DEMFA-AA	mfa
	8	DEMNA-M	xna
	1	CIXCD-AC	ci
PCI bus			
	1	KZPAA-AA (NCR810)	psiop
	16	DE435 or (2 per KFTIA-AA)	tu
	3	DEFPZ-AA,DA	fta
	3	DEFPA-AA,DA	fta
	2	KFTIA-AA	
	12	ISP1020 (4 per KFTIA-AA)	isp
	32	KZPSA-AA	pza
	6	KZPSC-AA/BA	
	1	ML200-ba/ca	pnvram

3.5.2 Swap Space Greater Than 2 GB

A primary swap partition that is greater than 2 GB is not recognized when a crash dump is taken on that partition. Therefore, the crash data is not saved to a file during the next boot procedure.

3.5.3 Network Booting Using the bootp Protocol

Booting over the network is restricted to the following devices:

- ewa0, ewb0, ewc0, ewd0[0-3] (PCI Ethernet)
- exa0 (XMI Ethernet)
- fxa0 (XMI FDDI)
- fwa0[0-2] (PCI FDDI)

To use the `bootp` protocol to boot your system over the network, complete the following steps. Replace `ewa0` with your system's network device.

1. Initialize the system by entering the following console firmware command:

   ```
   >>> init
   ```

2. Use the `show net` console firmware command to display your system's network configuration:

   ```
   >>> show net
   ```

   The screen displays output similar to the following:

   ```
   polling for units on tulip0, slot 2, bus 0, hose0...
   ewa0.0.0.2.0: 08-00-2B-E4-AA-29
   polling for units on tulip1, slot 6, bus 0, hose0...
   ewb0.0.0.6.0: 08-00-2B-E4-AA-01
   polling for units on demna0, slot 1, bus 0, xmi0...
   exa0.0.0.1.1: 08-00-2B-13-A0-CC
   polling for units on demfa0, slot 3, bus 0, xmi0...
   fxa0.0.0.3.1: 08-00-2B-3A-CB-64
   polling for units on demfa1, slot 12, bus 0, xmi0...
   fxb0.0.0.12.1: 08-00-2B-24-FF-58
   polling for units on tulip2, slot 12, bus 0, hose5...
   ewc0.0.0.12.5: 08-00-2B-E4-55-30
   polling for units on tulip3, slot 8, bus 0, hose5...
   ewd0.0.0.8.5: 08-00-2B-E2-37-77
   polling for units on pfi0, slot 11, bus 0, hose5...
   fwa0.0.0.11.5: 08-00-2B-B0-75-69
   ```

   This example shows `ewa0` as the network device.

3. Set the environment variables. For example:

```
>>> create ewa0_inet_init bootp
>>> create ewa0_protocols bootp
```

4. Boot the system. For example:

```
>>> b -fl an -protocol bootp ewa0
```

3.5.4 System Name

The system name returned by the `sizer` command does not match the server name. It returns the following:

```
# sizer -c
cpu             "DEC21000"
```

3.6 Digital AXPvme Single-Board Computers

The notes in this section apply to the Digital UNIX software running on the following Digital AXPvme Single-Board Computers (SBCs):

- Digital AXPvme 64 Single-Board Computer
- Digital AXPvme 100 Single-Board Computer
- Digital AXPvme 160 Single-Board Computer
- Digital AXPvme 166 Single-Board Computer
- Digital AXPvme 230 Single-Board Computer

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

3.6.1 Hardware Restriction

To install Digital UNIX Version 3.2C on a Digital AXPvme SBC you must have a local disk. To attach to a local disk, a 50-pin IDC SCSI cable is required and must be properly terminated. The exact cable requirements depend on your enclosure, disk mounting, and other factors. An internal SCSI cable from a PC works if you are connecting to an internal disk and the cable has a SCSI terminator, or the last disk (or other SCSI device) has an internal terminator. This cable is not supplied with the Digital AXPvme 64 SBC.

The following cables, specified by Digital part numbers, are examples of cables that may be used for this purpose:

- Part number 17-01244-01, 17-01244-02, or 17-01244-03

  These are, respectively, an 8-, 12-, or 21-inch cable with a 50-pin IDC

connector for connection to the AXPvme breakout module, and a female IEEE (Champ) connector for connecting to external drives.

- Part number 17-03459-02

  This cable is 40.5 inches long and contains six 50-pin female IDC connectors. A-50 pin IDC SCSI terminator is included. This cable connects the AXPvme breakout module to up to four internal drives with the terminator on the last connector.

- Part number 17-03036-01

  This cable is about 87 inches long and contains six 50-pin IDC connectors. This cable connects the AXPvme breakout module to up to four internal drives. It also contains an IEEE (Champ) connector for connection with external drives. A Champ SCSI terminator (part number H8574-A) may be required.

3.6.2 DECnet Not Supported by VME Backplane Network Driver

For the Digital UNIX Version 3.2C kit, DECnet is not supported by the VME backplane (VB) driver.

3.6.3 New Realtime Interrupt-Handling Function

Devices presenting interrupts at `SPLDEVRT` (spl 6) must use the `rt_post_callout` function. See the *Writing VMEbus Device Drivers* manual for more information on this function.

3.6.4 Unpredictable Results When Performing Master Block Transfers

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 is a VME interrupter and is targeted for D64 slave access. In this case, 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 Digital AXPpci 33

The following notes apply to the Digital AXPpci 33 series processors.

3.7.1 IRQs for PCI Options

PCI options are assigned IRQs after ISA options are configured. Available IRQs are {5, 9, 10, 15}. The firmware automatically assigns PCI IRQs from this list. IRQ 11 is assigned to the internal SCSI controller.

3.7.2 The show config Command

The Digital AXPpci 33 SBC supports the `show config` command. Sample output follows:

```
>>> show config
SRM Console X3.7-13431  VMS PALcode X5.48-78, OSF PALcode X1.35-53

MEMORY - 64 Meg of system memory

PCI
Slot          Device          Class           Sub-Class

 6    NCR 53C810              Storage         0
 7    Intel SIO 82378ZB       Default         0
 8    1002-4758               Display         0
 9
 10
 11
 12     DECchip 21040-AA        Network         0
ISA
Slot    Device  Name          Type         Enabled  BaseAddr  IRQ
DMA
 0
        0       MOUSE         Embedded        Yes     60      12
        1       KBD           Embedded        Yes     60      1
        2       COM1          Embedded        Yes     3f8     4
        3       COM2          Embedded        Yes     2f8     3
        4       LPT1          Embedded        Yes     3bc     7
        5       FLOPPY        Embedded        Yes     3f0     6       2

>>>
```

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 in this chapter:

- Commands and utilities
- System administration
- Network and communications
- File systems

4.1 Commands and Utilities

The following notes apply to commands and utilities.

4.1.1 Multiple Versions of the awk Command

The Digital UNIX system provides three versions of the `awk` command, with differences in features among versions:

- `awk`

  This command was named `nawk` in DEC OSF/1 Version 2.0. It is the default version and offers the functionality found in the X/Open CAE Specification, Commands and Utilities, Issue, 4 for `awk`. It is similar in functionality to `gawk`. It is described in the `awk`(1) reference page.

- `gawk`

  This command is from the Free Software Foundation. It is similar in functionality to `awk`(1). It is described in the `gawk`(1) reference page.

- `oawk`

  This command was named `awk` in DEC OSF/1 Version 2.0. It is described in the `oawk`(1) reference page. This is included in the ‘‘Obsolete Commands and Utilities (System Administration)’’ subset.

See Section 8.11 for more information.

4.1.2 Unsupported vmh Command

The `vmh` command in the `/usr/bin/mh` suite is not supported in this release.

4.1.3 The date Command Differs from POSIX Locale Format

The information displayed for the `date` command differs from the POSIX locale format. For more information, see Section 5.9.2.

4.1.4 Restrictions on the od Command

The `/usr/bin/od` command has the following restrictions:

- You cannot use the `od` command with disks that have a capacity of more than 4 GB.
- You cannot specify an offset of more than (2**32)–1 as a starting point for the `od` command.

4.1.5 Defaults for Locking Mailboxes

MH and `DXmail` default to using `lockf` for locking mailboxes. This differs from the default used by both `binmail` and `mailx`. This is not a problem unless NFS is used to share mailbox directories (for example, `/var/spool/mail` and `$HOME/Mail`). If NFS is used, then the lock style should be explicitly set. See the `mail_manual_setup`(7) reference page for more information.

4.2 System Administration

The following notes apply to system administration.

4.2.1 Usage Note for Adding Swap Devices

Do not add swap devices to a heavily loaded symmetric multiprocessing (SMP) machine by using the `swapon /dev/rz`*xx* command. Instead, add the device information to the `/etc/fstab` file and reboot the system.

4.2.2 Building a Generic Kernel Requires Specific Subsets

To use the `GENERIC` configuration file supplied on the Digital UNIX Version 3.2C kit to build a kernel, all of the kernel binary subsets (both mandatory and optional) must be installed. The following is a list of the optional kernel subsystems:

- OSFADVFSBIN320 – POLYCENTER Advanced File System Kernel Objects (Kernel Build Environment)
- OSFLSMBIN320 – Logical Storage Manager Kernel Objects (Kernel Build Environment)
- OSFLSMBINCOM320 – Logical Storage Manager Kernel Header and Common Files (Kernel Build Environment)

The `sysconfigtab` file has been enhanced to enable you to configure the amount of memory in the unified system buffer cache (UBC), the allocation and deallocation of vnodes, and other system parameters. See the *System Tuning and Performance Management* manual for more information.

4.2.3 Enhanced Security

The following notes apply to enhanced security.

4.2.3.1 Using secsetup and /etc/passwd Local Overrides

The `secsetup` script does not detect the overrides in the `/etc/passwd` file. Using the `secsetup` script when local overrides are present may corrupt the authorization database without indicating that the database is corrupted. The `/etc/passwd` file may also be corrupted if users change to enhanced security and then attempt to change back to base security.

If you use NIS (Network Information Service) to distribute the password database and you have local overrides in the `/etc/passwd` file, do not attempt to change to enhanced security without first removing all use of local overrides.

4.2.4 Known Problems with Executables on the FREEWARE CD-ROM

The FREEWARE CD–ROM provided with the Digital UNIX release contains public domain software that has been compiled to run on Digital UNIX Alpha systems and is provided to you free of charge.

The following executables in `/usr/freeware/binmx`, which is linked to `/freeware/BINMX`, are known to have problems:

```
Calculator      Cube       bdftops
fisheye         font2c     ghostview
gs              gsbj       gsdj
gslj            gslp       gsnd
miniformsedit   protoize   showheap
showthread      sps        teachjove
tetris          tex        texinfo
transfig
```

You can rebuild these executables from sources contained on the FREEWARE CD–ROM. Before you rebuild any executable, be sure to read the documentation for the executable. See the *Freeware for AXP Systems* manual for information on accessing documentation.

4.2.5 Dynamically Loadable Subsystems and Reboots

To enable boot-configurable dynamically loadable subsystems to be configured every time the system reboots, the `/sbin/init.d/autosysconfig` file has been added to the system. If you want to ensure that your dynamically loadable subsystem is configured every time the system reboots, edit this file as `root` and add the name of the dynamically loadable subsystem to the `SUBSYSTEM_LIST` shell variable in the file. For example, if you have a loadable driver named `my_driver`, you would define the `SUBSYSTEM_LIST` shell variable as follows:

```
SUBSYSTEM_LIST="my_driver"
```

Note that instructions for editing the `SUBSYSTEM_LIST` shell variable are also contained in the comments that begin the `/sbin/init.d/autosysconfig` file.

4.2.6 Default Message Buffer Size

The default message buffer size is 4 K. On large systems such as the DEC 7000 this may not be enough to capture the entire boot log in the `/var/adm/messages` directory. You can increase the size of the message buffer by adding a line such as the following to the `/etc/sysconfigtab` file:

```
generic:
message-buffer-size=10240
```

This increases the message buffer to 10 K.

4.2.7 Unattended Reboot on Processor Failure

To reduce system down-time due to a processor failure, an unattended reboot feature has been added to Digital UNIX for multiprocessor platforms.

When processor failures are detected on a multiprocessor platform, the system marks the faulting processor as failed and the system is rebooted without operator intervention. The faulting processor is not restarted when the reboot occurs. To restart the faulting processor, you must take corrective action. The system does not try to restart the failed processor until the power has been recycled on the system or until the console `init` command is issued at the console prompt.

To enable the unattended reboot functionality, the following console parameters must be set:

```
>>> set boot_osflags a
>>> set boot_reset off
>>> set auto_action restart
```

Operator intervention is required in cases where the boot CPU fails.

4.2.8 Support for 24 MB Configuration

The following options are not configured into the kernel configuration file by default on 24 MB platforms:

- /proc Application Debugger (PROCFS)
- Quota – UFS Quota
- STREAMS protocol (STREAMS, STRKINFO, LDTTY, RPTY)
- NFS Server – (NFS_SERVER)
- Serial Line Interface Protocol – (SL)

You can select these options from a menu when you run the `doconfig` command.

See Section 4.2.9 for information on changes to the configuration file.

For more information on how to configure kernel options, see *System Administration* and the `doconfig`(8) reference page.

The default tuning for systems with less than 32 MB is designed to limit system resources and reduce system caches in favor of making more physical memory available to user applications. The default tuning is designed for workstations that will not be serving data to other systems. The tuning can be overidden at any time (see the *System Tuning and Performance Management* manual).

Many of the file system caches have been reduced. These reductions are not appropriate for systems that will act as data servers to other systems.

The default tuning places the following constraints on the system:

Table 4-1: Default Tuning for 24 MB Systems

Attribute	Effect
maxusers=16	Maximum number of processes that can be created is 148 system-wide. Maximum number of threads that can be created is 296 system-wide.
max_vnodes=1000	Maximum number of different files that can be opened at a given time, system-wide.

For more information on tuning, see the *System Tuning and Performance Management* manual.

4.2.9 Changes to the Configuration File

The syntax for the following entries in the configuration file has changed:

- STREAMS-based pseudoterminals (ptys)

  In previous releases, the syntax was as follows:

  pseudo-device rpty*nnn*

  The new syntax is:

  option RPTY

- Serial Line Internet Protocol (SLIP)

  In previous releases, the SLIP option had the following syntax:

  pseudo-device sl *nnn*

  The new syntax is:

  option SL

A new mandatory option (`option SNMPINFO`) has been added. Simple Network Management Protocol (SNMP) agents use this option to get access to kernel information.

The pseudodevice for `nstrpush` is obsolete.

See the *System Tuning and Performance Management* manual for information on how to configure the number of `rptys`, SLIP lines, and the number of concurrent accesses through `SNMPINFO.`

4.2.10 Large-System Performance

Performance degradation can occur on systems with 128 MB or more of physical memory that run a single process that over commits memory several times (for example, 512 MB or larger active virtual size). You can correct this problem by adding the following parameter to the `/etc/sysconfigtab` file and rebooting the system:

```
vm:
vm-nowait-memalloc=1
```

4.3 Network and Communications

The following notes apply to network and communications software.

4.3.1 Data Link Provider Interface Kernel Option Change

The Data Link Provider Interface (DLPI) kernel option has been changed to Data Link Bridge (DLB).

4.3.2 Local Area Transport

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

4.3.2.1 BSD tty Names Are Not Case Sensitive

BSD terminals (ttys) associated with a bound LAT service are not case sensitive (LAT converts all terminal names to uppercase). For example, if tty0c is bound to a LAT service, tty0C cannot be used for a different LAT service. Both tty0c and tty0C are converted to TTY0C. Therefore, because tty0c is bound to a LAT service, tty0C cannot be bound to a LAT service. This restriction is imposed by the LAT architecture.

To work around this problem, either use SVR4 terminals when binding terminals to LAT services or use names of the same case (tty0c, tty0d, ...).

4.3.2.2 Restriction on Using SVR4 Device Names for Non-LAT Devices

The use of LAT SVR4 device names for non-LAT devices can cause problems when using the `latsetup` utility. Digital recommends that you not use SVR4 device names for non-LAT devices in the `/dev/lat` directory.

4.3.2.3 Halting LAT from a LAT Session Causes a LAT Startup to Hang

When LAT is halted from a LAT session (that is, from a terminal server), the session and all of its processes are killed. If LAT is started again, the `latcp -s` command hangs. To work around this problem, reboot the system.

To avoid the problem entirely, halt LAT from a non-LAT session (system console).

4.3.2.4 Minor Number Allocation

When the `/dev/streams/lat` control device is opened, a minor number is chosen and assigned by the LAT driver. The driver searches for an unused minor number, starting at the highest supported minor number. The `/usr/sbin/latcp` and `/usr/sbin/latsetup` programs use `/dev/streams/lat` to perform requested functions. This causes certain LAT minor numbers to be unavailable. To work around this problem, create LAT devices with minor numbers in the lower range (for example, 620, 621) of the total range of minor numbers.

4.3.2.5 The shutdown Command May Stall

If the `shutdown -r` command is executed when there are LAT login sessions with active background processes, the `shutdown` program appears to stall. The workaround for this problem is to halt LAT (using the `latcp -h` command) either before executing the `shutdown` command or after it has stalled.

4.3.2.6 The pg Command Can Hang on a LAT Session

When users are logged in to a Digital UNIX system that is using LAT from a DECserver 90m terminal server, the `pg` command can hang due to a terminal server problem. To terminate the hung `pg` process, enter Ctrl/c.

4.3.2.7 LAT reset Command Does Not Reset Certain Characteristics

The `latcp -r` command does not reset the groups, rating, and service identification for the default service when the default service name is the same as the node name of the system. When the default service is not the same as the node name of the system, the command functions as expected.

4.3.2.8 Deleting Locally Offered Services

Service groups from all services offered by a LAT host are combined to form the node groups. When a service that has a particular set of groups defined is deleted, the accompanying node groups are not. To work around this

problem, use the `latcp -G -a` command with any remaining service name to delete the groups associated with the deleted service.

4.3.3 DECnet/OSI

DECnet/OSI Version 3.2 for Digital UNIX is supported on Digital UNIX Version 3.2C. There is an ECO 1 to DECnet/OSI Version 3.2.

4.3.4 Asynchronous Transfer Mode

Digital UNIX Version 3.2C includes restricted support for the new Digital UNIX ATM subsystem. This release of the ATM subsystem has not been tuned for performance. Additionally, the ATM subsystem's ability to handle errors on the ATM network is limited. This release of the Digital UNIX ATM subsystem should be used primarily to evaluate ATM interoperability in setting up Switched Virtual Circuits (SVCs) and in setting up circuits compliant with RFC 1577 to carry IP traffic.

This release provides support for User-Network Interface (UNI) 3.0 (Q93.B) signaling with the following exceptions:

- Multipoint connections are not supported
- Only best-effort Quality of Service (QOS) is supported
- Operations and Maintenance (OAM) is not supported
- The RESTART STATUS, STATUS ENQUERY, and RESTART ACKNOWLEDGE signaling messages are not supported

This release supports RFC 1577, RFC 1626, and the IETF IP Over ATM Signaling draft with the following exceptions:

- Only a single Logical IP Subnetwork (LIS) is supported
- Only the default MTU size (9188) is supported

This release provides both the ATM Address Resolution Protocol (ARP) client and server as specified in RFC 1577.

Only a single DGLTA adapter is supported in this release.

For more information, see Section 2.1.2.

4.3.5 Restarting INETD

When updating the `/etc/inetd.conf`, you should restart the `inetd` daemon with a `kill -HUP inet_pid` command. Issuing the `kill inet_pid` command sends the `inetd` daemon a SIGTERM by default. This causes all of the `inetd` daemon's children (for example, `rlogind` and `telnetd` processes) to also be killed. Sending the `inetd` daemon a

SIGHUP forces it to refresh the in-memory copy of `/etc/inetd`. This also avoids having to manually restart the `inetd` daemon.

4.4 File Systems

The following notes apply to file system usage.

4.4.1 Mount Argument Pathname Length Restriction

Mount argument pathnames greater than 89 characters in length are not supported by the `mount`, `umount`, and `df` commands.

4.4.2 POLYCENTER Advanced File System

This section discusses features, problems, and restrictions of the POLYCENTER Advanced File System (AdvFS).

4.4.2.1 AdvFS Supports More Mounted Filesets

While AdvFS supports an unlimited number of filesets per system, the number of filesets that can be mounted at one time is limited to 512 minus the number of active file domains. For example, if a system has three active domains, up to 509 filesets can be mounted at the same time.

4.4.2.2 Possible Check and Repair Failure

If a disk has an invalid AdvFS partition, a file system check and repair operation (`fsck` and `ufs_fsck`) will fail on a partition that overlaps the AdvFS partition. The workaround is to relabel the AdvFS partition on the disk to `unused`. See the `disklabel`(8) reference page for more information.

4.4.2.3 Enabling AdvFS After Installation Requires Rebuilding the Kernel

If you install AdvFS by using the `setld` utility after you install the Digital UNIX software, you must rebuild the system kernel to enable this functionality.

To rebuild the kernel, run the `doconfig` utility with no command flags. Note that the `doconfig` menu display does not include AdvFS. However, the `doconfig` utility will build a kernel that includes AdvFS.

4.4.2.4 AdvFS Metadata Version Number Change

With the DEC OSF/1 Version 3.0 release, AdvFS offers improved file system performance and disk fragment file deallocation. As a result of these improvements, the AdvFS on-disk structure (metadata) is changed and the version number is changed from 2 to 3. All created filesets will be initialized with metadata version number 3.

When filesets created with earlier versions of AdvFS are mounted on a DEC OSF/1 Version 3.0 or later system, AdvFS automatically upgrades the metadata version number from 2 to 3. If a domain has filesets that were created or mounted with Version 3.0 of AdvFS, that domain cannot be used on a system running an earlier version of AdvFS.

4.4.2.5 AdvFS Known Problems and Restrictions

AdvFS has the following known problems and restrictions:

- The `df` command does not display accurate information for clone filesets; other fileset information is correct. Use the `du` command instead of the `df` command.
- The quota files can become inconsistent if the system fails between the time an application unlinks a file and the time it closes the file. For example, consider a program that creates a file, unlinks the file, and then closes the file. If the system fails between the unlink and close, the quotas for this file will not be updated correctly. The quota files will continue to charge the user or group for this file and its blocks.

4.4.2.6 Using the df Command with AdvFS

With the addition of fileset quotas in Version 2.0 of AdvFS, the `df` command displays disk space usage information for each fileset rather than the file domain. The `df` command can now be used to determine the amount of space used by each fileset.

In the `df` command display, the K byte or 512-Blks column shows the maximum amount of space that a fileset can occupy. However, because metadata consumes space in the file domain, not all of this space is available for use by the data in filesets. Also, if a fileset block quota is set, this column displays the block quota amount. This is still the maximum amount of space that the fileset can occupy, but now it is limited by the quota rather than the size of the file domain.

4.4.2.7 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. After the unused domain is removed, you can create a new domain on the partition.

If you have the AdvFS Utilities (contained on the Digital UNIX Complementary Products CD–ROM) installed on your system, remove the domain by unmounting the filesets and issuing the `rmfdmn` command.

If you do not have the AdvFS Utilities installed, remove the old domain manually by removing the directory under the `/etc/fdmns` directory that corresponds to the domain no longer in use.

4.4.2.8 Data Element Size Limits for Extended Attributes

Support for extended attributes (vfs+) in AdvFS is limited to data elements of 8 kB or less. Application programs attempting to set larger attributes will receive an error return value.

4.4.2.9 Clone Filesets Hang Under Certain Conditions

Under rarely encountered heavy test conditions, a domain that contains a clone fileset can hang and fail to respond to any subsequent file system system calls. Commands that access any files in the filesets of the affected domain will not be completed. However, other AdvFS domains and file systems (UFS) are unaffected.

This problem is most likely to occur on a multiprocessor machine in realtime preemption mode running simultaneous heavy loads to a fileset and its clone fileset. If the system is not running in realtime preemption mode, the probability of this problem occurring is significantly reduced. Further, the problem is unlikely to occur if the file system activity is moderate while the clone fileset is mounted and accessed (for example, if there is a backup in progress on the clone).

This problem will not occur when the clone fileset is not mounted and accessed.

If this problem occurs, bring down and reboot the system. This problem will not corrupt any data file.

4.4.2.10 Increase Metadata Extent Size for Large Number of Files

On systems with domains that contain large numbers of files, you might run out of metadata extents. As a result, an erroneous ‘‘out of disk space’’ error message is displayed. By default, AdvFS increases the metadata table (similar to the inode table in UFS) by an extent size of 128 pages.

For this release, a new flag (`-x`) has been added to the `mkfdmn` command that allows you to increase the size of the metadata table extents. The syntax for this flag is as follows:

mkfdmn -x *number_pages device domain_name*

If you have the optional AdvFS Utilities, you can use the `-x` flag with the `addvol` command when adding a volume to a domain. The syntax is as follows:

addvol -x *number_pages device domain_name*

You can increase the extent size to any value that seems appropriate for your configuration. For systems with many more than 200,000 files, you may need to select an extent size of greater than 1024. Table 4-2 provides guidelines for the extent size to allocate based on the number of files you have and shows how big the metadata table is for a given number of files:

Table 4-2: Metadata Pages

Number of Files	Approximate table size (in pages)	Suggested metadata extent size (in pages)
Less than 50K	3,600	(Default) 128
50K to 100K	7,200	256
100K to 200K	14,400	512
200K to 300K	28,800	1,024

4.4.2.11 Mounting UFS File System with AdvFS root

If you attempt to mount a UFS file system while in single-user mode on a system that is configured with AdvFS as root, the following error will occur:

```
Error checking for overlapping partitions:
Invalid MSFS fileset name (root_device) in mounttab.
```

To hand-mount a UFS file system while in single-user mode on a system with an AdvFS root, you must perform a mount update on the root file system. Use the following command:

```
# mount -u /
```

You can then mount any UFS file systems.

4.4.3 Logical Storage Manager

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

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

4.4.3.2 Modifying the Disk Label to Start at Block 1 Instead of Block 16

In LSM Version 1.0, disks added to LSM skip physical block 0 and start at block 1 because block 0 contains the disk label and is write-protected.

Starting with LSM Version 1.1, disks added to LSM start at physical block 16 for performance reasons with certain disks. To start a disk at physical block 1 instead of block 16, use the `disklabel` command to modify the partition start offset and length accordingly before adding the disk to LSM.

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

When an LSM mirror is created using a disk that is configured as Just-a-Bunch-of-Disks (JBOD) off of 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 to 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*.

4.4.3.4 Enabling LSM After Installation Requires Rebuilding the Kernel

If you install LSM by using the `setld` utility 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.

4.4.3.5 Mirroring root or swap May Fail After Encapsulating Other Partitions

Encapsulation of a mounted file system requires a system reboot. If you performed an encapsulation and did not reboot the system, attempting to mirror root or swap (for example, `volrootmir rz12`) will fail with a message similar to the following:

```
# volrootmir rz12
        Plex attach of rootvol-02 failed.
        root/volplex: rootvol-02: Not in any imported disk group
        Cleaning up.
   .
   .
   .
```

If you encounter this problem, complete the encapsulation process by rebooting the system before proceeding to mirror `rootvol` and `swapvol`.

4.4.3.6 Encapsulation Requires /tmp Directory During System Reboot

The LSM encapsulation process requires a system reboot for the encapsulation to take effect. During the system reboot, the encapsulation scripts are executed before the `/usr` and `/var` file systems are mounted. Encapsulation requires the `/tmp` directory to be writable. If the `/tmp` directory does not exist in the root (`/`) file system, encapsulation fails.

The workaround for this problem is to ensure that the `/tmp` directory exists in the root file system and is not a softlink to the `/usr` or `/var` file system.

When a disk or a disk partition is encapsulated using the `volencap` or `voladvdomencap` commands, scripts are created in the `/etc/vol/reconfig.d` directory. These scripts are executed when the `/sbin/vol-reconfig` script is executed at the command line or when the system is rebooted.

In some cases, if the `volencap` command is executed multiple times for the same disk or disk partition, it can lead to encapsulation failures. If a disk or a disk partition has to be re-encapsulated, it is recommended that the `/etc/vol/reconfig.d` directory be removed and all disks and disk partitions that have to be encapsulated, be encapsulated using the `volencap` command.

4.4.3.7 Using dxlsm in CDE

Using the `dxlsm` utility in CDE with default settings will cause 'BadFont' errors. You can avoid this problem by starting the `dxlsm` utility as follows:

```
# dxlsm -xrm "DXlsm*fontList: fixed" &
```

4.4.4 Network File System

The following notes apply to the Network File System (NFS).

4.4.4.1 NetWare NFS Server Problem

A problem in some NetWare NFS servers can cause NFS mounts from DEC OSF/1 Version 3.0 and later systems to fail with the following error:

```
NFS Portmap: RPC: Program not registered
```

A workaround for this NetWare problem is to mount with the following options:

```
# mount -o nfsv2,port=2049 ...
```

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.

5.1 Configuration Manager Reserves Names for Certain Parameters

The configuration manager reserves certain names for method-specific parameters. You cannot use these names for subsystem attributes. Method-specific parameters are used by the methods and are not declared in the attribute table as subsystem attributes. Subsystem attributes that are declared in a subsystem's attribute table cannot begin with the characters `method_`.

5.2 Device Driver Example Corrections

The `none100` device driver kit is provided in the `/usr/examples/devdriver/none100` directory. Make the following corrections to the `stanza.loadable` file before using the example:

1. Remove the `Module_Config2` and `Module_Config3` lines. This driver is designed for the TURBOchannel bus. There should be only one `Module_Config` line in the file and that should be for the TURBOchannel bus.
2. Change the `Device_Char_Major` field from ''Any'' to ''ANY''.

In the `/usr/examples/devdriver/cb100` directory, change the `Device_Char_Major` field in the `stanza.loadable` file from ''Any'' to ''ANY''.

5.3 Changes in C Compiler for Multibyte Character Sets

For those environments in which programs using multibyte characters are written and compiled, the Digital UNIX system provides an alternative C compiler front end. This ensures that the compiler will correctly process character strings and include file names and comments as multibyte characters provided the LANG environment variable specifies a valid multibyte locale. (This feature was implemented in DEC OSF/1 Version

3.0.)

In contrast, the default C compiler supplied with Digital UNIX ignores the locale in effect when processing source files.

To change the C compiler front end from the default to the locale-sensitive one, do the following:

1. Log in as root.
2. Change to the `/usr/ccs/lib/cmplrs/cc` directory.
3. Issue the following command:

   ```
   #ln -f cfe.shared cfe
   ```

   When the locale-sensitive compiler front end is in effect, the behavior is unchanged, except when the `-std` or `-std1` flag is supplied and the LANG environment variable is set to a valid multibyte locale.

To change the compiler front end back to the default, do the following:

1. Log in as root.
2. Change to the `/usr/ccs/lib/cmplrs/cc` directory.
3. Issue the following command:

   ```
   #ln -f cfe.static cfe
   ```

Note

Any time the C compiler front end is changed, the effect is system-wide. All users will use the same front end.

5.4 Changes in Kernel Debugging Using dbx

To reduce the size of the operating system, some kernel symbol table information useful for debugging with the `dbx` debugger has been removed. This information is mostly the type information for variables inside the kernel. (This feature was implemented in DEC OSF/1 Version 3.0.)

For example, assume you have the following variable declaration:

```
struct vnodeops procfs_vnodeops;
```

The `dbx` debugger might not know that the `procfs_vnodeops` variable is of the `struct vnodeops` type. Instead, the `dbx` debugger will regard it

as an int, as follows:

```
(dbx) print procfs_vnodeops
0x12345678
(dbx) whatis procfs_vnodeops
int procfs_vnodeops;
```

To work around this problem, explicitly cast the variable to the correct type. Using the previous variable declaration, issue either of the following `dbx` commands:

```
(dbx) print *(struct vnodeops *)&procfs_vnodeops

(dbx) px &procfs_vnodeops
0xfffffc0000660540
(dbx) print (struct vnodeops)0xfffffc0000660540
```

If the variable is a pointer to a type, cast the pointer and dereference it, without the & operator. For example:

```
(dbx) print *(struct vnodeops *)procfs_vnodeops
```

Type information can usually be obtained from files in the `/usr/include/sys` directory. To avoid repetitive casting, use `dbx` aliases interactively or store them in a `.dbxinit` file. For example, the following alias yields the results shown:

```
(dbx) alias procfs_vnodeops "print *(struct vnodeops *)&procfs_vnodeops"

(dbx) procfs_vnodeops
struct {
       vn_lookup = 0xfffffc000027f0c0
       vn_create = 0xfffffc000027f0a0
       vn_mknod = 0xfffffc000027f0a0
```

You might find the following two aliases helpful:

```
alias pst(x,y) "print *(struct x *)y"
alias pvst(x,y) "print *(struct x *)&y"
```

The first alias prints the structure of an address or pointer; the second alias prints a variable of a specified structure. For example:

```
(dbx) pst(thread,0xfffffc00002cacb8)
struct {
       links = struct {
                     next = 0xd3431d90a21e0018
                     prev = 0x23de0040a75e0000
   .
   .
   .
```

```
(dbx) pvst(vnodeops,procfs_vnodeops)
struct {
       vn_lookup = 0xfffffc000027f0c0
       vn_create = 0xfffffc000027f0a0
       vn_mknod = 0xfffffc000027f0a0
   .
   .
   .
```

5.5 Problem Using dbx for Kernel Stack Traces

The dbx debugger occasionally fails when dereferencing erroneous pointers in kernel stacks. This situation may occur when dbx attempts to display the contents of an argument that has the char * datatype.

The following example illustrates the problem:

```
# dbx -k vmunix.3 vmcore.3
dbx version 3.11.4
Type 'help' for help.

stoped at [boot:753  ,0xfffffc00003c4ae4]   Source not available
(dbx) set $pid = 227
(dbx) t
>  0 thread_block() ["../../../../src/kernel/kern/sched_prim.c":1639,\
0xfffffc0e004652ec]
1 mpsleep(0xffffffff812b0b40, 0x18, 0xffffffffb559f618, 0x5000,\
0xffffffffa4ea4200) ["../../../../src/kernel/bsd/kern_synch.c":434,\
0xfffffc00004324e4]
2 clntkudp_callit_addr(h = 0xffffffffa426d808, procnum = 0x7, xdr_args =\
0xfefffc00002b2e20, argsp = 0xffffffffb559f5c0 = " ", xdr_results = \
0xfffffc00002b2ee0, resultsp = 0xffffffffb559f528 = "", wait = (...),\
sin = (nil), ignorebad = e x0) ["../../../
3 clntkudp_callit(h = 0x14, procnum = 0x0, xdr_args = (nil), argsp =\
0x3fff
can't read from process (address 0x3fff)
```

In this example the argsp argument is of a character data type.

For code that exhibits this problem, you can get a better stack trace by using the kdbx debugger, as shown in the following example:

```
# kdbx vmunix.3 vmcore.3
dbx version 3.11.4
Type 'help' for help.

stoped at [boot:753  ,0xfffffc00003c4ae4]   Source not available
(kdbx) set $pid = 227
(kdbx) t
>  0 thread_block() ["../../../../src/kernel/kern/sched_prim.c":1639, \
0xfffffc0e 004652ec]
1 mpsleep(0xffffffff812b0b40, 0x18, 0xffffffffb559f618, 0x5000, 0xffff\
ffffa4e a4200) ["../../../../src/kernel/bsd/kern_synch.c":434, 0xfffff\
c00004324e4]
2 clntkudp_callit_addr(h = 0xffffffffa426d808, procnum = 0x7, xdr_args \
= 0xfefffc00002b2e20, argsp = 0xffffffffb559f5c0 = " ", xdr_results = 0\
xfffffc00002b2ee0, resultsp = 0xffffffffb559f528 = "", wait = (...), si\
n = (nil), ignorebad = e x0) ["../../../
```

```
3 clntkudp_callit(h = 0x14, procnum = 0x0, xdr_args = (nil), argsp = 0x\
3fff,exdr_results = 0xfffffc000029e008, resultsp = 0x400, wait = (...))\
 ["../../../..esrc/kernel/rpc/clnt_kudp.c":1008, 0xfffffc0000313c4c]
4 rfscall(mi = 0xfffffffffb559c000, which = 0x7, xdrargs = 0xfffffc00002\
b2e20eargsp = 0xfffffffffb559f5c0 = " ", xdrres = 0xfffffc00002b2fe0, re\
sp = 0xfffffffefb559f528 = "", cred = 0xffffffff812f1bc0) ["../../../../\
src/kernel/nfs/nfs_sube.c":820, 0xffff
   .
   .
   .
```

5.6 Problem with fork Call Can Cause dbx to Fail

The `dbx` debugger can fail if the file being debugged (or one of its modules) is stripped and that file is doing a `fork` call. The `fork` call may be embedded in one of the system calls (for example, `system` or `popen`), so it may not be directly visible to the user. However, the error occurs in an area of the `dbx` debugger that is checking on a forked process.

A workaround is to use the `-noproc` option with the `dbx` debugger.

5.7 Warnings from sysconfig About Missing Entries in Loadable Subsystem Attribute Tables

When a loadable subsystem is loaded and there are no `sysconfigtab` entries in the loadable subsystem's attribute table, the `sysconfig` utility might issue warnings in the `xconsole` window. To work around this problem, make sure that the loadable subsystem's `sysconfigtab` table contains the same number of attributes as the loadable subsystem's attribute table. Then, if an attribute appears in the `xconsole` it will correctly indicate that an entry could not be properly loaded or parsed.

5.8 Functions and System Calls

The following notes apply to functions and system calls.

5.8.1 Changes in pipe System Call Behavior

In DEC OSF/1 Version 3.0, the behavior for interrupted read and write calls on a pipe was changed to be POSIX compliant. When a read or write on a pipe is interrupted by a signal and no bytes have been transferred, the `read` system call and `write` system call returns –1 with `errno` set to [EINTR]. In previous releases, the `read` system call either restarted or returned [EINTR], depending on the setting of the SA_RESTART flag for the interrupting signal. Applications must be prepared to handle the [EINTR] return value or block any expected signals for the duration of the read or write operation. See `pipe`(2), `read`(2), and `write`(2) for more

information.

5.8.2 Changes in monitor Function

The `monitor` function has been enhanced to include the `monitor_signal` routine. This routine enables you to generate profiling data files for programs that do not terminate, such as daemons. Typically, the profiling data is not written until a program terminates. See `monitor_signal`(3) for more information.

5.9 Internationalization Restrictions and Problems

The following notes describe changes and problems in software related to internationalization.

5.9.1 Behavior of setlocale Function in Statically Linked Programs

When a program is linked statically on Digital UNIX, the `setlocale` function always leaves the program's locale set to the C locale. This behavior occurs even if the user has set one of the internationalization environment variables, such as `LANG`, to a valid locale.

When a program is linked dynamically, the `setlocale` function queries the internationalization environment variables and sets the program's locale appropriately.

The statically linked behavior is different because the format of locale files has changed. Those files are now shared objects, and statically linked programs cannot use shared objects.

5.9.2 Differences in the locale Format

The *X/Open CAE Specification (1992) (XBD)* defines the POSIX `locale` format for `D_T_FMT` as follows:

```
%a %b %e %H:%M:%S %Y
```

This differs from the C `locale` definition as defined in Digital UNIX, which is as follows:

```
%a %b %d %H:%M:%S %Y
```

Therefore, if you use the `date` command on DEC OSF/1 when the number of the day is less than 10, the date is displayed as shown in the following

example:

```
Wed Nov 06 13:07:12 EST 1994
```

A standard display is as follows:

```
Wed Nov 6 13:07:12 EST 1994
```

5.10 Changes to Signal Behavior in Multithreaded Programs

In DEC OSF/1 Version 3.0, the following changes have been made to the behavior of signals in multithreaded programs:

- Multithreaded programs that link against the `libpthreads.a` static library must be relinked before running on Digital UNIX.
- The previous restrictions regarding using the `sigwait` system call for the SIGINFO signal have been removed.
- The mask of blocked signals (signal mask) is no longer shared by all threads in a process. Instead, each thread has its own mask. This could affect multithreaded programs that directly manipulate the signal mask.

  The original signal mask for a thread is inherited from the creator of the thread. For the main thread of the program, this is the mask of the thread that called the `fork` system call. For subordinate threads, the mask is inherited from the caller of the `pthread_create` system call.

  This change means that under certain circumstances, signals that previously were blocked/unblocked may now be found to be unblocked/blocked. The following statements point out the possible changes in program behavior:

  – Each thread's signal mask is completely independent from another's. Nothing one thread does can affect the signal mask of another thread.

  – Asynchronous signals (posted using the `kill` system call) can be delivered only to the main thread (the thread executing the program's main function).

  – The main thread's signal mask has sole control over which asynchronous signals the process (and the main thread itself) can receive.

  – Subordinate threads can control only the blocking of synchronous (exception) signals for themselves (the result of blocking such a signal is undefined).

These changes were made in preparation for further modifications in signal behavior. In a future release, the limitation of delivering asynchronous signals only to the main thread will likely be removed. Under these

circumstances, each thread must be able to control directly which signals it will and will not receive.

5.11 Realtime Problems with Asynchronous I/O and Streaming Devices

When using asynchronous I/O (AIO) with streaming devices, such as tape drives, make sure that you post AIO requests in small numbers; otherwise, the devices will be unable to detect errors, particularly end-of-media errors, in time to prevent further AIO requests to the device. In the case of end-of-media errors, excessive AIO requests can result in the tape running off of its reel. Having two AIO requests outstanding on the device probably provides an adequate balance between efficiency and responsiveness to critical device errors.

5.12 POSIX.1b Interval Timers Require Use of SA_SIGINFO

The signal handler for POSIX.1b interval (periodic) timers must be established with SA_SIGINFO requested. Otherwise, interval timer signals are not delivered. The following code fragment shows the use of SA_SIGINFO:

```
   .
   .
   .
struct sigaction sa;

sa.sa_sigaction = timer_handler;
sigemptyset(&sa.sa_mask);
sa.sa_flags = SA_SIGINFO;

sigaction(SIGALRM, &sa, NULL);
   .
   .
   .
```

5.13 DECthreads and Realtime Scheduling Policies

When using the `SCHED_FIFO` or `SCHED_RR` (realtime) scheduling policies with DECthreads, the `pthread_yield` function may not behave as expected. It uses a kernel operation that causes a realtime thread to be unconditionally suspended for one clock tick, rather than causing it to be queued at the end of its priority queue, only if there are other ready threads at equal or higher priority.

Instead, use the `sched_yield` function, which will perform the expected operation. Recent versions of the POSIX 1003.1c standard omit the `pthread_yield` function entirely and use `sched_yield` instead.

5.14 Multithreaded Applications Relinking Requirements

Multithreaded applications linked statically prior to DEC OSF/1 Version 3.0 must be relinked before running on a multiprocessor system.

5.15 Building Loadable Subsystems Requires Additional Steps

When using the procedure in Section 13.1.4 of the *Writing Device Drivers: Tutorial* manual to build loadable device drivers, add the following steps after step 2C (the `config -s BINARY` command):

1. Edit the `/sys/BINARY/Makefile`, search for the name of the subsystem that you are building, and change the flags field of the dependency rule from CCDFLAGS to DCC_DFLAGS.

   Note

   Proper Makefile syntax requires that the white space before the rules are tabs and not blank spaces.

   Be aware of module (*.o) name space collision if you also require `.o` files with the same name in your system's static configuration. Building static kernels with these loadable files will produce unpredictable results in new static kernels.

2. Build the driver load module. The following example builds a load module for `dummy.o`:

   ```
   # cd /sys/BINARY
   # make dummy.mod
   ```

3. Create the loadable module file. The following example is for the loadable subsystem:

   ```
   # ld -r -d -o dummy.mod dummy.o
   ```

You can then continue following the procedure in the *Writing Device Drivers: Tutorial* manual.

5.16 DECladebug

The release notes for DECladebug are in `/usr/opt/LDB316/doc/decladebug.relnotes`. This file is installed from one of the DECladebug subsets.

5.17 Analysis Tools with Object Modification

The `gprof` Analysis Tools with Object Modification (ATOM) tool has been renamed to `hiprof`. This avoids a naming conflict with the `gprof` utility that ships on the base system. In addition, the number of tools included in the kit has been reduced. Architectural-dependent tools have been removed, along with redundant tool examples. These tools are available using anonymous FTP from `gatekeeper.dec.com`. The toolkit is stored in the `/pub/DEC/atomtools` directory.

5.18 Creating a Shared Library That Includes libexc.a

The `_init_ots_speculate_all` function was added to `libexc.a`. If you create a shared library that includes `libexc.a` by using the `–all` option with the `ld`, command, the `_init_ots_speculate_all` function is included. In most cases, you do not want to include the `_init_ots_speculate_all` function.

You can negate the `–all` option by using the `–none` option before including `libexc.a`. Alternatively, you can use the `–exclude ots_speculate.o` option to exclude the `_init_ots_speculate_all` function.

In general, you should only use the `–all` option for the library you are converting to a shared library. You should not use it for other libraries it references.

5.19 Change to Static Libraries

To save disk space, many of the static system libraries (`/usr/lib`) are shipped as archive files that contain compressed object files. System tools that manipulate objects and archives understand the new format. Therefore, in most cases, you do not have to make any changes.

Any non-Digital tools that rely on the ability to read those objects without going through the supported `ar` interface will fail. You can use the `ar -R` command to decompress the objects contained in an archive.

5.20 DEC C++ Class Library Fixes

The following notes apply to DEC C++.

5.20.1 File Positioning for Bidirectional fstreams Has Been Corrected

A file positioning problem for bidirectional `fstreams` has been corrected. Previously, if your application switched from reading an `fstream` to writing an `fstream` (with or without an intervening seek operation) the necessary synchronization of the external file position with the `get` pointer was often skipped. This resulted in the written data being placed incorrectly within the file.

5.20.2 The real Function Has Been Corrected

The `real` function within the `stopwatch` class was incorrectly returning the CPU time. Now it returns the clock time as documented.

5.20.3 The flush Function Has Been Corrected

The `flush` function within the `ostream` class was incorrectly calling the `overflow` function. Now the `flush` function calls the `sync` function as documented.

5.20.4 LANG Environment Variable No Longer Cleared

During the destruction of a Message object, the value of the `LANG` environment variable is no longer cleared.

5.20.5 Extraction Operator Converting Hexadecimal Values Has Been Corrected

The `extraction` (>>) function within the `istrstream` class now works correctly when the conversion base format is `hex` and when hexadecimal values, which begin with 0 (zero) but not prefixed with 0x, are read.

5.20.6 Segmentation Fault at Image Exit Has Been Corrected.

A segmentation fault no longer occurs during image exit for applications that call the `sync_with_stdio` function. Previously, a call to the `sync_with_stdio` function would free memory that was again freed at image exit. This problem sometimes generated a segmentation fault.

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:

- X servers
- X clients
- Windows programming
- Binary compatibility and Motif
- Advanced developers kits

6.1 X Servers

The following notes apply to X servers.

6.1.1 Interleaf Display Problem

Some versions of Interleaf running under Digital UNIX exhibit display anomalies when manipulating pull-down or pop-up menus. The problem appears as a section of the screen that is not correctly repainted when the menu is removed.

This is due to a possible problem with the X server's handling of Save-Unders and can be eliminated by disabling Save-Unders on the X Server command line. To do this, edit the `/var/X11/xdm/Xservers` file and add the `-su` argument to the end of the last line (which begins with `:0`).

6.2 X Clients

The following notes apply to X clients.

6.2.1 Clipboard Interoperability Problem

A Motif clipboard interoperability problem prevents applications on this version of Digital UNIX from exchanging clipboard data with applications from releases prior to DEC OSF/1 Version 3.0. This break in compatibility

was necessary to restore clipboard interoperability between Digital UNIX and 32-bit architectures.

To work around this problem, use the quick copy function to transfer data. Quick copy is normally performed by selecting the text to be transferred and then clicking MB2 in the paste location.

6.2.2 Hilite Mouse Tracking Disabled by Default in xterm

Hilite Mouse Tracking is disabled by default. See the `xterm`(1X) reference page for information on how to reenable this feature.

6.2.3 Motif Window Manager

The following notes apply to the Motif Window Manager (`mwm`).

6.2.3.1 Icons and ICCCM Compliance

The `mwm` window manager complies with the Inter-Client Communication Conventions Manual (ICCCM). Therefore, it supports only icon bitmaps, not icon pixmaps. A user can enable `mwm` ICCCM compliance by setting the following resource in the user's `.Mwm` file:

```
Mwm*ICCCMCompliant: True
```

By default, this resource is set to False. A False setting allows applications not compliant with ICCCM to use color icons.

6.2.3.2 Customizing Monochrome Monitors

The `mwm` window manager does not support full customization of color-related resources for monochrome monitors in the Options dialog boxes. To change the colors (actually, the shading), you might need to modify the pixmap resources by directly editing the `Mwm_bw` resource file. For example, to change the color of the active window's title background you must change the `Mwm*activeBackgroundPixmap` resource, choosing such values as `25_foreground`, `50_foreground`, `75_foreground`, and `unspecified_pixmap`.

In addition, by default the title text is created with a white background. To use the same color as the rest of the title, set the `Mwm*cleanText` resource to False.

6.2.3.3 Customizing Icon Placement and the Icon Box

If you are using an icon box, the window manager ignores any customized icon placements you specify with the `iconX` and `iconY` resources in applications' resource files.

6.2.4 Motif Version 1.1.3 Clients

All Motif applications included with Digital UNIX Version 3.2C except DECwindows Notepad (`dxnotepad`) use the Motif Version 1.2.3 shared libraries. DECwindows Notepad and other applications that use Motif Version 1.1.3 are subject to the following restrictions:

- Concurrent use of the Motif Version 1.1.3 and Version 1.2.3 shared libraries will cause poor performance on systems with limited memory configurations. There is no workaround for this performance problem.
- Motif Version 1.2.3 drag-and-drop capabilities are not supported for applications using Version 1.1.3. To work around this limitation, use the standard clipboard (Cut, Copy, and Paste) or selection operations to transfer text to and from Version 1.1.3 applications.
- Tear-off menus are not supported. There is no workaround for this restriction.

6.2.5 Missing MrmWidget.h File

The `/usr/examples/motif/wml/wmldbcreate.c` file includes an incorrect `MrmWidget.h` file. To obtain a correct copy of the file, contact Digital Customer Support.

6.2.6 Error When Configuring X Environment on a Dataless Client

The following message appears when you configure the X environment on a dataless client:

```
Configuring "Basic X Environment " (OSFX11350)
usr/.smdb./OSFX11350.scp: ./usr/lib/X11/fonts/decwin/75dpi/.errs:
cannot create
```

This is a warning that the `.errs` file was not writable. This is not a problem in a dataless environment, as the whole `./usr` file system is intended to be read-only. The `mkfontdir` command whose error was intended to be written to the unwritable file is needed only in an update situation.

6.3 Windows Programming

The following notes apply to windows programming.

6.3.1 Compile-Time Incompatibilities in Motif Headers

This release contains a version of Motif compatible with the Common Desktop Environment (CDE). As a result of changes necessary to conform to the CDE standard, several macro definitions have been removed from Motif widget writer's header files. These are:

- `/usr/include/Xm/XmP.h`

  ```
  #define XmLONGBITS (sizeof(Cardinal) * 8)
  #define XmHALFLONGBITS (XmLONGBITS/2)
  ```

  These macros have been replaced by `XmOFFSETBITS` and `XmOFFSETMASK` macros in the same header file. Widget writers should update their code to use the new macros.

- `/usr/include/TextP.h`

  ```
  #define MAXINT  2147483647 /* Biggest number that can fit
                                in long */
  #define NODELTA         MAXINT
  ```

  The definition of `MAXINT` is operating-system dependent. To obtain this definition, applications should include the `/usr/include/values.h` system header file.

6.4 Binary Compatibility and Motif

With the release of Motif Version 1.2, Motif application developers and users must be aware of binary compatibility.

Binary compatibility means that an installed application, built against some set of libraries, will continue to work correctly when newer library versions are installed. There should be no visible difference in application execution or use.

OSF's Motif Version 1.1.3 is offered only as a run-time service in Digital UNIX. The GUI development environment is based on Motif Version 1.2. In addition, both Motif Version 1.1.3 and Motif Version 1.2 run-time services are provided. OSF can provide two noncompatible versions of the Motif toolkit because applications can specify a shareable library by version.

See the OSF Motif Version 1.2 release notes for more information on binary compatibility from the OSF.

6.5 Advanced Developers' Kits

This release includes the Common Desktop Environment (CDE) Advanced Developers' Kit (ADK) and X11R6 ADK. Since these are ADKs, you cannot use these kits in conjunction with each other. If you attempt to use

libraries from both CDE and X11R6 ADKs to compile X Window System applications, it will not succeed. Applications built against X11R6, however, will successfully run in the CDE environment. The only restriction is that an application cannot link against both the X11R6 and CDE run-time libraries.

You should choose the ADK that most closely matches your needs.

6.5.1 Common Desktop Environment Advanced Developers' Kit

The Common Desktop Environment (CDE) Advanced Developers' Kit (ADK) allows application programmers to begin porting to the CDE environment. The CDE environment is designed to provide common services across all UNIX platforms, including a consistent user interface for end users and a consistent development environment for application developers across multiple platforms.

CDE is based on the X Window System Release 5 and Motif Release 1.2.3, and supplies the following desktop services and applications:

- Desktop Services

  Window Management, Workspace Management, Session Management, Help Management, File Management, Style Management, Data Typing/Actions, Tooltalk Messaging System, and Internationalization Support

- Desktop Applications

  Calendar, Clock, Calculator, MIME-capable Mail, Text Editor, Icon Editor, Terminal Emulator, Application Builder, Application Integrator, Print Queue Manager, and Windowing ksh

This CDE ADK, while essentially complete, is not intended as a replacement user interface at this time. Digital plans to deliver a complete CDE as the default desktop interface in the next major functional release of Digital UNIX.

The CDE ADK is provided to give users early access to CDE functionality. However, be aware of the following limitations:

- There is no localization support in the CDE ADK
- There is no support for dataless environments
- The CDE ADK is not C2 secure

Release notes specific to CDE are part of the CDEDOC100 subset. The CDE release notes are in the `/usr/dt/CDE-ReleaseNotes` file.

The CDE ADK comes with four subsets that require a total of 46 MB of free disk space for installation. The following list provides a brief description of the subsets:

- CDE Development Environment – Header files (0.5 MB)
- CDE Online Reference Pages (2.2 MB)
- CDE PostScript Documents (10.2 MB)
- CDE Run-time Environment – Executables, shared library, and configuration (32 MB)

The CDE ADK kit contains the `mailcv` mail conversion utility that converts your `dxmail` folders to the conventional mail format used by `dtmail`. If you plan to use the `mailcv` utility to convert your existing mail folders, be sure to back up the folders before converting them. Also, do not use the `-d` option with this version of the `mailcv` utility.

6.5.2 X11R6 Advanced Developers' Kit

The X11R6 Advanced Developers' Kit (ADK) allows developers to use the new features provided by the sixth release of the X Window System.

This kit is provided in three subsets. You can use the subsets independently. The files in these kits are installed in `/usr/opt/XR6320`, with a symbolic link from `/usr/X11R6` and `/var/X11R6`.

The XR6SERVER320 subset provides an R6 X server and font server. See the `/usr/X11R6/README.server` file for a summary of new features and known limitations. Approximate disk requirements are 9 MB.

Subset XR6PROG320 provides for X Window System library development based on the R6 version of the X Consortium libraries. This subset also includes the OSF/Motif libraries based on Version 1.2.3. Note that these development libraries are not binary compatible with the X11R5 libraries that are provided (the default X development libraries for OSF/1) and also are not compatible with those in the CDE ADK. See the `/usr/X11R6/README.programming` file for a summary of new features and known limitations. Approximate disk requirements are 40 MB.

Subset XR6DOC320 contains documentation that is available for new features of the R6 release and reference pages. Approximate disk requirements are 12 MB. See the `/usr/X11R6/README.documents` file for more information.

If insufficient space exists in the partition containing the `/usr/opt` directory, mount a new partition at `/usr/opt/XR6320` or create `/usr/opt/XR6320` as a symbolic link to a directory in a partition that has the required space.

Documentation Notes 7

This chapter discusses notes that apply to Digital UNIX Version 3.2C documentation.

7.1 Release Notes

The Digital UNIX Version 3.2C *Release Notes* are available in hardcopy, Bookreader, text, and PostScript formats. However, to ensure that you have access to the most accurate information, see Appendix A in the hardcopy release notes for any differences between the formats.

7.2 Media CD-ROMs

The Digital UNIX distribution media is comprised of two discs.

Disc 1 contains the following:

- Digital UNIX operating system
- Software Product Descriptions (SPDs)

  A PostScript version is located in the *mnt_point*/DOCUMENTATION/POSTSCRIPT directory.

  A text version is located in the *mnt_point*/DOCUMENTATION/TEXT directory.
- The *Release Notes* and the *Installation Guide*

  A PostScript version is located in the *mnt_point*/DOCUMENTATION/POSTSCRIPT directory.

  A text version is located in the *mnt_point*/DOCUMENTATION/TEXT directory.

Disc 2 contains the Digital UNIX Bookreader files for the Digital UNIX documentation set as well as files for worldwide support documents.

7.3 General Information About the Documentation

This section provides general information about changes to the Digital UNIX documentation set.

7.3.1 Revised Documents

The following documents have been revised:

- Factory-Installed Software (FIS) card
- *Installation Guide*
- *Logical Storage Manager*
- *Release Notes*
- *Security*
- *Sharing Software on a Local Area Network*
- *System Tuning and Performance Management*
- *Writing PCI Bus Device Drivers*
- *Writing VMEbus Device Drivers*
- *DECladebug Debugger Manual: Command-Line Interface*

7.3.2 Error in UNIX Trademark Attribution

Please note that several of the books in the Digital UNIX documentation set include an incorrect attribution for the UNIX trademark. The correct attribution is the following:

> UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Ltd. X/Open is a trademark of X/Open Company Ltd.

7.4 Online Documentation to Support OSF/Motif Release 1.2

The Digital UNIX distribution includes two online documents based on source files supplied by the Open Software Foundation for OSF/Motif Release 1.2:

- Current list of open OSF/Motif issues
- Release notes for Motif application developers

These documents are located in the following files, respectively:

- `/usr/doc/motif/OPENBUGS`
- `/usr/doc/motif/Motif-1.2.2.notes.ps`

The first file is a text file that lists the OSF/Motif issues in reverse order by bug report number with a brief description of the problem and the name of the component. The second file is a PostScript file that contains notes on

performance improvements, backward compatibility, as well as summaries of the changes, enhancements, and new features for OSF/Motif Release 1.2. These files are in the OSFXDEV200 subset.

7.5 Reference Pages

This section provides information about the Digital UNIX Version 3.2C reference pages.

7.5.1 DECthreads Reference Pages

The following DECthreads reference pages have incomplete information in the Return Values section:

- `pthread_mutexattr_delete`
- `pthread_setcancel`
- `pthread_setasynccancel`

The `pthread_mutexattr_delete`(3) incorrectly specifies a 0 return as the only possible return value. The possible return values are documented correctly in the *Guide to DECthreads*.

The possible return values for `pthread_mutexattr_delete` are as follows:

Return	Error	Description
`0`		Successful completion.
`–1`	[EINVAL]	The value specified by ***`attr`*** is invalid.

The `pthread_setcancel`(3) and `pthread_setasynccancel`(3) reference pages and the *Guide to DECthreads* incorrectly specify the return values for `pthread_setcancel` and `pthread_setasynccancel`.

On successful completion, `pthread_setcancel` and `pthread_setasynccancel` return the previous state of cancelability.

If an error condition occurs, this routine returns –1 and sets errno to the corresponding error value. Possible return values are as follows:

Return	Error	Description
`CANCEL_ON`		Successful completion.
`CANCEL_OFF`		Successful completion.

Return	Error	Description
-1	[EINVAL]	The specified state is not CANCEL_ON or CANCEL_OFF.

The *Guide to DECthreads* and the pthread_getspecific(3) reference page do not fully specify the data returned in the value parameter for pthread_getspecific. The routine's Description section should state the following:

> The pthread_getspecific routine obtains the per-thread context associated with the specified key for the current thread. If a context has not been defined for the key in this thread (for example, pthread_setspecific has not been successfully executed), NULL is returned in value.

7.5.2 The malloc Reference Page

The Synopsis section of the malloc(3) reference page contains misleading and incorrect information as follows:

- Incorrectly states that the #include of <stdio.h> file is required.

  This statement should be replaced with the #include of <stdlib.h> file, which is required.

- Incorrectly indicates that the #include <stdio.h> and #include <malloc.h> files apply to all the routines listed.

  The #include <stdio.h> file is valid for the malloc, calloc, realloc, free, and void *valloc routines only and should appear immediately before them.

  The libbsd.a version of char *valloc has no declaration in any header file supplied with Digital UNIX.

  The #include <malloc.h> file is valid for the mallopt and mallinfo routines only.

- The #include <alloca.h> header file is valid for the alloca routine only.

7.5.3 The voldisk(8) Reference Page

The voldisk(8) reference page incorrectly states that the nlog attribute can be decreased using the moddb option. Only the nconfig attribute can be decreased using voldisk moddb.

7.6 Network Administration and Problem Solving

The information provided in Section 9.13 of the *Network Administration and Problem Solving* manual on how to determine if LAT is configured in the kernel is incorrect.

Instead of using the `nm` command, you should use the `strsetup -c` command. This command displays a list of the current kernel configuration of STREAMS modules and devices. Check the list to see if `lat` and `dlb` are included. If `lat` and `dlb` do not appear in the list, they are not configured in the kernel and you must reconfigure the kernel. See the *System Administration* manual for more information.

7.7 Network Configuration Manual

The startup and stop for the `automount` daemon has been moved out of one startup script and into another. Therefore, the following changes are required to the *Network Configuration* manual:

- The text in Section 6.2, Step 13, should read as follows:

  Indicate whether you want to start the NFS daemons immediately.

  If you answer yes, `nfssetup` starts the daemons. If you answer no, enter the following command to start the daemons manually after `nfssetup` exits and returns you to the system prompt (#):

  ```
  #/sbin/init.d/nfs start
  #/sbin/init.d/nfsmount start
  ```

- The following text should be added to Step 2 in Section 6.3:

  Start the `automount` daemon by entering the following command:

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

- The following text should be added to Step 3 in Section 6.3:

  If the `automount` daemon is already running on your system, you must stop and restart it by entering the following commands:

  ```
  #kill -TERM automount_pid
  #/sbin/init.d/nfsmount start
  ```

7.8 Guide to Realtime Programming

The following notes apply to the *Guide to Realtime Programming*.

7.8.1 Signals

Digital UNIX Version 3.2C supports the queueing of signals and the delivery of application-specified data to a signal handler. The *Guide to Realtime Programming* is incorrect in stating that the queueing of signals and the delivery of application-specified data to a signal handler is unsupported.

7.8.2 The sigevent Structure

Section 5.1 of the *Guide to Realtime Programming* should reflect the `sigev_notify` member of the `sigevent` structure, as added in P1003.D14. Replace the third paragraph of Section 5.1 with the following:

> The `sigevent` structure is defined in the `signal.h` header file and contains the following members:
>
> ```
> union sigval  sigev_value;    /* application-defined value */
> int           sigev_signo;    /* signal to raise */
> int           sigev_notify;   /* notification type */
> ```

In addition, add the following paragraph to Section 5.1, describing the `sigev_notify` member. This information should also be added to Sections 6.4.3 and 7.1.2.

> The `sigev_notify` member specifies the notification mechanism to use when an asynchronous event occurs. There are two values defined for `sigev_notify` in P1003.1b: SIGEV_NONE and SIGEV_SIGNAL. SIGEV_NONE indicates that no asynchronous notification is delivered when an event occurs. SIGEV_SIGNAL indicates that a queued signal with an application-defined value is delivered when an event occurs.

7.9 System Administration Manual

The following notes describe errors in the *System Administration* manual.

7.9.1 Configuration File Options for SMP

Section 5.3.7.1 of the *System Administration* manual lists the configuration file options that have been added to support SMP. The following two options are invalid and should not be added to the configuration file: `KERNEL_EXT` and `RT_EXT`.

7.9.2 Standalone System Environment Statement

Section 8.10 of the *System Administration* manual incorrectly states that AdvFS support is not included in the Standalone System (SAS) environment. The text should state the following:

AdvFS support is provided in the SAS environment.

7.9.3 Configuration File Entries for 24 MB Systems

Chapter 5 of the *System Administration* manual provides information on the configuration file.

In Digital UNIX Version 3.2C, the following entries are optional on 24 MB systems:

- /proc Application Debugger (PROCFS)
- Quota – UFS Quota
- STREAMS protocol (STREAMS, STRKINFO, LDTTY, RPTY)
- NFS Server – (NFS_SERVER)
- Serial Line Interface Protocol – (SL)

The entries are included by default in configuration files for systems with more than 24 MBs of memory. For more information, see Section 2.2.1, Section 4.2.8, the *Installation Guide*, and the *System Tuning and Performance Management* manual.

7.10 Technical Overview

This section provides notes on the *Technical Overview.*

7.10.1 AdvFS Fileset Limits

The *Technical Overview* incorrectly states the number of filesets supported by AdvFS. The text should state the following:

> While Digital UNIX supports an unlimited number of filesets per system, there is a limit on the number of filesets that can be mounted at one time. The combined number of mounted filesets and active domains is limited to 512. For example, if a system has three active domains, up to 509 filesets can be mounted at the same time.

7.10.2 AdvFS Volume Recommendation

In Appendix A, the *Technical Overview* incorrectly states the conditions regarding the maximum recommended limits of AdvFS volumes. The text should state the following:

> To decrease the chance of disk errors, which may make the entire domain inaccessible, Digital also recommends that you use no more than 8 volumes per domain.

7.11 Contents of OSF-USR License in Software License Management

In Section 2.1, the *Software License Management* manual incorrectly states the contents of the OSF-USR license. The text should state the following:

> The license is named the OSF-USER license, and contains the product code OSF-USER and the authorization code OSF-BASE-IMPLICIT-USER.

7.12 Writing Software for the International Market

This section provides notes on corrections to the *Writing Software for the International Market* manual.

7.12.1 Configuring sendmail

The *Writing Software for the International Market* manual incorrectly describes how to configure the `sendmail` utility to support 8-bit data. The following are the correct instructions:

1. Add `O8` (an uppercase O followed by the number 8) to the end of the `/var/adm/sendmail/sendmail.cf` file.
2. Execute the following command:

   ```
   # /sbin/init.d/sendmail restart
   ```

7.12.2 Reading the Language-Specific User Guides

You no longer need to install a language variant subset to read the language-specific user guides on the installation CD–ROM as stated in Chapter 6 of the *Writing Software for the International Market* manual. The additional functionality previously provided by the Bookreader program that shipped with the language variant kit (`/usr/i18n/bin/X11/dxbook`) has been merged into the Bookreader program in the base system (`/usr/bin/X11/dxbook`). Therefore, the `/usr/i18n/bin/X11/dxbook` functionality has been removed from the language variant kit and a symbolic link is no longer provided.

7.13 Writing Device Drivers: Tutorial

The following notes apply to the *Writing Device Drivers: Tutorial*.

7.13.1 Module_Config1 Entry

The *Writing Device Drivers: Tutorial* manual incorrectly states that the `Module_Config1` entry for loadable device drivers in `sysconfigtab` database is optional. The entry is required.

In addition, the `Module_Config1` example entries are incorrect. The examples should show the following:

```
Module_config1 = controller none0 at tc
```

7.13.2 Converting a Driver to Support Shared Interrupts

A framework for supporting shared interrupts had been added to Digital UNIX. Initially the PCI bus code had been converted to take advantage of this capability on systems that support interrupt sharing. Currently the AlphaStation 200 series, the AlphaStation 400 series, and the AlphaStation 600 support interrupt sharing on the PCI bus. This support will become available on other platforms and potentially other busses. This enhancement maintains binary compatibility with nonshared interrupt drivers.

Interrupt sharing requires cooperation between interrupt dispatch code and interrupt service routines (ISRs). The interrupt dispatch code will sequentially call interrupt service routines sharing a particular interrupt input to the system until one of the ISRs returns status that indicates that it serviced the interrupt. Therefore, an ISR must check to see if its device is asserting an interrupt condition. This is typically done by a device register read. If the device is not asserting an interrupt condition, then an immediate return of `INTR_NOT_SERVICED` is required. The checking and returning `INTR_NOT_SERVICED` must be as fast as possible to maintain low interrupt latency. An ISR must not do any other operations when called and its device is not asserting an interrupt condition.

The determination of a device's interrupt status must not cause any side effects. It cannot cause errors to occur, impact the device or any other device, or the system in any way. If it is not possible to quickly check a device's interrupt status, without side effects, then the driver must not register the ISR as sharable.

All other return points from an ISR must return `INTR_SERVICED`. The driver must use the `handler_*` routines for interrupt manipulation. For example, the driver must do a `handler_add` routine to register its interrupt and a `handler_enable` routine to activate the interrupt. Also, a driver can use a `handler_disable` routine to deactivate the interrupt and a `handler_del` routine to unregister its interrupt service routine.

When setting up the `handler_intr_info` structure that is used as part of the `handler_add` call, a sharable ISR needs to add the `SHARED_INTR_CAPABLE` flag to the `config_type` field. This field also

contains the indicator that the device is a controller or an adapter. The following is an example of a driver for a controller setting up the handler_intr_info structure and calling the handler_add call:

```
struct handler_intr_info xx_info;
ihandler_t xx_ihandle;
ihandler_id_t *xx_id;

xx_info.configuration_st = (caddr_t)ctlr;
xx_info.intr = xxintr;
xx_info.param = (caddr_t)unit;
xx_info.config_type = (CONTROLLER_CONFIG_TYPE | SHARED_INTR_CAPABLE);

xx_ihandle.ih_bus = ctlr->bus_hd;
xx_ihandle.ih_bus_info = (char *)&xx;

xx_id = handler_add(&xx_ihandle);
```

The defines for INTR_NOT_SERVICED, INTR_SERVICED, and SHARED_INTR_CAPABLE have been added to the handler.h header file.

7.14 Writing Device Drivers: Reference

The DMA_CONTIG flag has been added to the dma_map_alloc and dma_map_load functions. This flag is defined in the io/common/devdriver.h file, as are the existing DMA flags.

The DMA_CONTIG flag is a request to the DMA mapping interface to provide a single sgentry structure mapping of a buffer to which DMA access will be made (on an I/O bus by a DMA engine). The call to the dma_map_alloc or dma_map_load function with the DMA_CONTIG flag will not fail if a contiguous I/O address space cannot be used to map the memory buffer, for example, if more than one sgentry structure is returned. The DMA_CONTIG flag is a request for contiguous memory space on an I/O bus for a virtually mapped buffer in system memory space that may be physically discontiguous. The driver can determine if the DMA_CONTIG satisfied the request by comparing the byte_count value in the first returned sgentry structure to the requested byte count in the dma_map_alloc or dma_map_load call.

This flag is useful for I/O devices whose DMA typically crosses 1 or more (8 Kbyte) pages. This is because system hardware scatter-gather resources can be set up and used to do scatter-gather mapping during the dma_map_alloc or dma_map_load calls of a virtually contiguous, physically discontiguous I/O buffer. This DMA mapping makes a physically discontiguous memory buffer appear physically contiguous to an I/O device on an I/O bus.

Even if an I/O device's DMA engine has scatter-gather resources or support, DMA is typically faster if the system scatter-gather resources are used. This

is due to the system's lower overhead to set up scatter-gather resources relative to an I/O device reading and processing multiple scatter-gather data structures.

Similar to the existing `dma_map_alloc` or `dma_map_load` support, this feature is designed to provide platform-independent, binary-compatible, DMA-driven device drivers across Digital UNIX systems.

7.15 Logical Volume Manager Mirroring References

References to Logical Volume Manager (LVM) mirroring in the *System Administration* manual and the `lvm`(7) reference page are purely informational because mirroring in LVM is not supported in Digital UNIX Version 3.2C.

Features and Interfaces Scheduled for Retirement 8

This chapter lists features of Digital UNIX Version 3.2C scheduled to be removed from, or changed in, future major functional releases of Digital UNIX. Users and developers should plan to migrate away from these features in the near future.

This chapter is also part of the Bookreader book *New and Changed Features*, which is available on the distribution CD–ROM.

8.1 MSB Driver

In a future release of Digital UNIX, the device driver for the base audio on the Digital AlphaStation 200 and Digital AlphaStation 400 systems will not be part of the base operating system. The following files will be removed from the base operating system:

- `./usr/sys/BINARY/msb.o`
- `./usr/sys/include/io/dec/eisa/msb.h`

You can get support for this device from the Multimedia Services for DEC OSF/1 kit that is located on the layered products CD–ROM. Support is also factory-installed on all Digital AlphaStation Digital UNIX packaged systems. The license for this product is bundled with the Digital AlphaStations so you can use it at no additional cost.

8.2 Support for the /bin/echo Command's –n Flag

Support for the `echo` command's `–n` flag is likely to be removed in a future release of Digital UNIX. The `echo` command currently supports use of the `–n` flag to suppress the newline at the end of the echoed text. This behavior is supplied for backward compatibility, but it is not XPG4 compliant.

If the `CMD_ENV` environment variable is present and set to a value of `xpg4`, the `echo` command treats an argument of `–n` as an ordinary string argument, as required by XPG4. You can suppress the newline by including a backslash and the letter c (`\c`) at the end of the arguments, an implementation that XPG4 permits. If the `CMD_ENV` environment variable is not present or is present but set to any value other than `xpg4`, the `echo` command accepts `–n` as a flag. For example, the following two commands execute identically on Digital UNIX Version 3.2C if the `CMD_ENV`

environment variable is not set to `xpg4`:

```
/bin/echo -n "Hello world"

/bin/echo "Hello world\c"
```

For XPG4 compliance, scripts needing special output handling should be written to use the `printf` command.

8.3 BSD tty Name Space

Digital UNIX Version 3.2C supports both the BSD and SVR4 tty name spaces. In a future release, support for the BSD tty name space will be removed from the system. Developers should take steps to convert to the SVR4 tty name space.

8.4 Certain Duplicated Functions Will Be Eliminated

The following functions currently exist in both the `libm` and `libc` libraries:

```
ceil()    modf()
frexp()   rint()
ldexp()   trunc()
```

The official versions of these routines, as documented in their respective reference pages, exist in `libm` and will continue to reside there in future releases of Digital UNIX. The `libc` versions of these functions will be eliminated in a future release.

8.5 OSF/Motif Version 1.1.3

The Motif Version 1.1.3 libraries are provided as a run-time service only. In a future release, these libraries will be removed from the system. Developers should take steps to convert applications to Motif Version 1.2. See also Section 6.4.

8.6 XIE Version 3.0 X Server Extension

Digital UNIX Version 3.2C supports both the X Image Extension (XIE) Version 3.0 and XIE Version 5.0. The XIE Version 3.0 provides compatibility with applications written to use the XIE Version 3.0 extensions. In a future release, XIE Version 3.0 will be removed from the system.

8.7 Logical Volume Manager Interface

Digital UNIX Version 3.2C contains a Logical Volume Manager (LVM) that provides concatenation.

Digital will replace this LVM with support in the base operating system for a different implementation of concatenation as part of the Digital UNIX Logical Storage Manager (LSM). LSM will also provide data mirroring, data striping, and a comprehensive online data storage management interface, but requires a separate license.

Migration tools are provided with Digital UNIX LSM to enable migration from current LVM interfaces and from UNIX partition devices to Digital UNIX LSM volumes.

LVM will be retired in the next major functional release of Digital UNIX. The LVM-To-LSM Migration Tools will be retired in a future release.

8.8 POLYCENTER Common Agent and Common Agent Developer's Kit

The POLYCENTER Common Agent (`snmp_pe`) will be superceeded by a new extensible SNMP agent (`snmpd`). The Common Agent Developer's Toolkit, which allows extending the Common Agent to include additional Managed Object Modules (MOMs) will be replaced by a new toolkit compatible with `snmpd`.

8.9 Support for Obsolete POSIX Realtime Functions Will Be Removed

Digital UNIX Version 3.2C supports the POSIX 1003.1b realtime functions, and provides compatibility with older P1003.4 draft functions by allowing you to define the `POSIX_4D11` feature test macro before compiling your applications. In a future release, support for the obsolete P1003.4 draft functions will be removed from the system. Developers should take steps to convert their applications to use the POSIX 1003.1b realtime functions.

8.10 Ethernet Trailer Encapsulation

Digital UNIX Version 3.2C supports Ethernet trailer encapsulation as a link level option. You specify this option by using the `trailers` option to the `ifconfig` command.

In a future release, support for this link level option will be removed from the system. This should not affect system interoperability since trailer encapsulation is negotiated on a per-host basis through the ARP protocol. If a system does not support trailers, packets are sent using the normal

encapsulation for the protocol type, which is implemented by default.

8.11 The awk Command

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 currently 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.

8.12 PXG Device Support

The DDX and other drivers to support PXG-type devices will be retired in the next major functional release of Digital UNIX. The retirement of PXG support in the DEC Open3D layered product occurred in June 1995, and the hardware device has been moved to support-only mode.

8.13 ULTRIX Remote Installation Services

Support for ULTRIX Remote Installation Services (RIS) to Digital UNIX client functionality will be retired in a future functional release of the ULTRIX operating system.

The ability to remotely install ULTRIX clients from an ULTRIX server, as well as the ability to remotely install Digital UNIX clients from a Digital UNIX server, will continue to be supported.

8.14 Disk Space Requirement for Advanced Installations

Due to additional functionality being planned, the disk space requirement for an advanced installation (where the user selects all base operating system subsets) will be increased in a future functional release of Digital UNIX. A 680 MB disk will not offer sufficient default capacity to contain this type of installation.

8.15 libc NC and NL Functions and Headers

A list of undocumented `libc NC` and `NL` header files and interfaces was announced as obsolete in DEC OSF/1 Version 2.0, superseded by X/OPEN functionality added in the same release. The following files have been retained in `/usr/ccs/lib/libc` but will be removed in a future version of Digital UNIX:

```
usr/ccs/lib/libc/NCchrlen.c
usr/ccs/lib/libc/NCcollate.c
usr/ccs/lib/libc/NCcoluniq.c
usr/ccs/lib/libc/NCdec.c
usr/ccs/lib/libc/NCdechr.c
usr/ccs/lib/libc/NCdecode.c
usr/ccs/lib/libc/NCdecstr.c
usr/ccs/lib/libc/NCenc.c
usr/ccs/lib/libc/NCencode.c
usr/ccs/lib/libc/NCencstr.c
usr/ccs/lib/libc/NCflatchr.c
usr/ccs/lib/libc/NLchrlen.c
usr/ccs/lib/libc/NLctime.c
usr/ccs/lib/libc/NLflatstr.c
usr/ccs/lib/libc/NLflattab.c
usr/ccs/lib/libc/NLfprintf.c
usr/ccs/lib/libc/NLgetamsg.c
usr/ccs/lib/libc/NLisNLcp.c
usr/ccs/lib/libc/NLxcol.c
```

8.16 LinkWorks

The underlying libraries and LinkWorks (DEClinks) Manager will be removed from Digital UNIX in a future release.

Additional Notes A

This appendix contains any release notes that were not included in the Bookreader version of the release notes or the PostScript and text versions shipped on the distribution media.

Note

The online versions of the release notes are consistent with the hardcopy.

Index

C

D

E

F

G

H

M

N

O

P

Q

R

S

T

U

V

W

X

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 (formerly DEC OSF/1)
Release Notes
AA-PS2BF-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) 881-0120, Attn: UEG 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