Digital UNIX (formerly DEC OSF/1)

Release Notes

Order Number: AA-QTLMA-TE

March 1996

Product Version: Digital UNIX Version 4.0

This book contains notes on software and documentation restrictions for Digital UNIX Version 4.0 and the bundled layered products that ship on the same CD–ROMs.

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

Digital Equipment Corporation
Maynard, Massachusetts

Contents

2 Installation Notes

3 Processor-Specific Notes

5 Development Environment Notes

6 Window System Software Notes

7 Documentation Notes

A Maximum System Limits

B Additional Notes

Index

Tables

About This Manual

This manual contains release notes for Digital UNIX® Version 4.0. The release notes are available in both hardcopy and an HTML version you can view with the Netscape Navigator. 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.

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

Audience

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

Organization

This manual is organized as follows:

Chapter 1	Contains an overview of new and changed features. This chapter is also part of the online (HTML) 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 on-line (HTML) book *New and Changed Features*.
Appendix A	Contains information about the maximum system limits.
Appendix B	Contains release notes that arrived after the base level was shipped. These notes are included in the hardcopy version but are not in the online version due to time constraints.

Related Documents

You should have the following documentation available during the installation of this release.

- The hardware documentation for your system
- The online or hardcopy reference pages
- The 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	Blue
System and network administrators	S	Red
Programmers	P	Purple
Device driver writers	D	Orange
Reference page users	R	Green

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

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

Reader's Comments

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

You can send your comments in the following ways:

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

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

`/usr/doc/readers_comment.txt`

- Mail:

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

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

Please include the following information along with your comments:

- The full title of the book and the order number. (The order number is printed on the title page of this book and on its back cover.)
- The section numbers and page numbers of the information on which you are commenting.
- The version of Digital UNIX that you are using.
- If known, the type of processor that is running the Digital UNIX software.

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

Conventions

The following conventions are used in this guide:

`%` `$`	A percent sign represents the C shell system prompt. A dollar sign represents the system prompt for the Bourne and Korn shells.
`#`	A number sign represents the superuser prompt.
`%` **`cat`**	Boldface type in interactive examples indicates typed user input.
`file`	Italic (slanted) type indicates variable values, placeholders, and function argument names.
**`[	]`** **`{	}`**	In syntax definitions, brackets indicate items that are optional and braces indicate items that are required. Vertical bars separating items inside brackets or braces indicate that you choose one item from among those listed.

. . .	In syntax definitions, a horizontal ellipsis indicates that the preceding item can be repeated one or more times.
`cat`(1)	A cross-reference to a reference page includes the appropriate section number in parentheses. For example, `cat`(1) indicates that you can find information on the `cat` command in Section 1 of the reference pages.
Ctrl/*x*	This symbol indicates that you hold down the first named key while pressing the key or mouse button that follows the slash. In examples, this key combination is enclosed in a box (for example, Ctrl/C).

New and Changed Features 1

This chapter provides brief descriptions of features that are new to the Digital UNIX system in this release or have changed significantly from previous releases.

1.1 Common Desktop Environment

The Common Desktop Environment (CDE) is the new default graphical user interface for Digital UNIX. 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 on Digital UNIX is based on the X Window System Release 6 (X11R6) and CDE/Motif 1.0 (OSF/Motif 1.2.4), and supplies the following desktop services and applications:

- Desktop Services:

Window Management	Workspace Management	SessionManagement
Help Management	File Management	Style Management
Data Typing/Actions	Tooltalk	Messaging System

- Desktop Applications:

Calendar	Calculator	MIME-capable Mail
Text Editor	Icon Editor	Terminal Emulator
Application Integrator	Print Queue Manager	Windowing ksh
Keyboard		

CDE is provided in seven software subsets that require a total of 57.81 MB of free disk space for installation. See the *Installation Guide* for information on the subset names, contents, and sizes.

The CDE kit contains the following migration tools:

- `mailcv` mail conversion

  This utility converts your `dxmail` folders to the conventional mail format used by CDE `dtmail`. If you plan to use the `mailcv` utility to convert your existing mail folders, back up the folders before converting them. Do not use the `-d` option with this version of the `mailcv` utility.

- `dxcaltodtcm` calendar conversion

  This utility converts a DECwindows Calendar, `dxcalendar`, database for use with CDE Calendar, `dtcm`.

1.1.1 The CDE Video Tour

A brief multimedia tutorial of CDE is located on the Digital UNIX V4.0 Associated Products Volume 1 CD–ROM. Once installed, the video tour can be accessed via the application manager in the "Information" folder. The user simply needs to double click the "CDE Video Tour" icon.

1.1.2 CDE Screen Savers

The CDE session manager supports X11R6 screen saver extensions and you can now select animated screen savers instead of a blank screen. This release also enables the automatic locking of screens after a specified idle time. Both features can be modified or disabled from the CDE Style Manager menu. Click on the Screen icon, and select the options you want.

1.2 X/Open-Compliant Curses

This release provides a new Curses implementation that incorporates the following sets of programming interfaces:

- X/Open Curses, Issue 4
- System V Multinational Language Supplement (MNLS)
- Minicurses
- BSD Curses

See Section 4.2 for information on usage and restrictions.

1.3 X11R6

This release of Digital UNIX supports Release 6 of the X Window System, Version 11 (X11R6) patchlevel 12. Prior versions of the operating system supported Release 5 (X11R5) patchlevel 26.

The Digital UNIX port of X11R6 supports all the features and functionality of previous releases of Digital UNIX. It also supports all X Consortium standard features of X11R6.

Included in new features are the following protocol extensions:

- BIG-REQUESTS

  Gives clients the ability to use requests that are arbitrarily large, rather than being limited to the size restriction of the core protocol. This can result in a significant performance improvement for applications that use large requests.

- DOUBLE-BUFFER

  Enables double buffering, using the new X Consortium standard.

- XIE (updated)

  Complete implementation of full XIE 5.0 protocol with a few exceptions.

- XKEYBOARD (Xkb)

  See the following section.

1.3.1 The X Keyboard Extension for X11R6 (XKB)

The XKB (X Keyboard) server extension is new for X11R6 and for Digital UNIX Version 4.0. XKB enhances control and customization of the keyboard under the X Window System by providing the following:

- Support for the ISO9996 standard for keyboard layouts.
- Compatibility with the core X keyboard handling (no client modifications are required).
- Standard methods for handling keyboard LEDs and locking modifiers such as CapsLock and NumLock.
- Support for keyboard geometry.

In addition, the X11R5 AccessX server extension for users with physical impairments has been incorporated into the XKB server extension. X11R5 applied to versions of Digital UNIX that preceded this release. These accessibility features include StickyKeys, SlowKeys, BounceKeys, MouseKeys, and ToggleKeys, and control over the autorepeat delay and rate.

Several applications that make use of XKB features are also new with Digital UNIX version 4.0. These applications include `Xdec`, `xkbcomp`, `xkbprint`, `xkbdfltmap`, `dxkbledpanel`, `dxkeyboard`, and `accessx`. Refer to the reference pages for more information.

Note that the final revision of the X Keyboard Extension, XKB Version 1.0, will be different from XKB Version 0.65, which is shipping with this release. Avoid creating code that directly references the XKB API and data structures. Any X clients created with direct references must be recompiled and relinked when XKB Version 1.0 is shipped in a future release. You may also have to modify your source code.

1.4 DXImageview

DXImageview version 1.0 is an OSF/Motif application for viewing a variety of image file types such as GIF and JPEG. DXImageview supports the Common Desktop Environment (CDE) recommended integration for applications. DXImageview is provided in the subset `OSFCDEAPPS400` and is installed as `/usr/bin/X11/dximageview.` The `-help` option displays information on using the application.

1.5 Commands and Utilities

The following new or changed commands and utilities are available in this release.

1.5.1 Changes to Mtools

`Mtools` software is included in the `OSFDOCTOOLS400` subset. In prior releases, the software was installed by an optional worldwide support subset.

1.5.2 The sendmail Utility Supports Configurable GECOS Fuzzy Matching

The `sendmail` utility now allows the user to configure the fuzzy logic for mail delivery. Previously, if the recipient's address did not precisely match any of the user names on the host, a best-match algorithm was applied against the GECOS field in the `passwd` file. If a unique best-match was found, the mail was delivered to this user. This behavior is now run-time configurable using the `-oG` option on the command-line. See `sendmail.cf`(4) for more information.

1.5.3 df Supports Large File Systems

The field width for the `Iused` and `Ifree` fields in the output of the `df` command has been increased to accommodate 12 digits when using the `-i` switch. This modification was made to support very large file systems where the number of inodes could exceed the field width that was previously set aside for these fields.

1.5.4 Compressed Reference Pages

To economize on disk space, reference pages are now shipped in compressed format. Compressed files were created with the `/usr/bin/gzip` utility. The `man` and `xman` utilities automatically uncompress the reference pages.

If required, you can uncompress reference page files manually with the `/usr/bin/gunzip` utility.

The catman command has also been enhanced to work with compressed catman files. All three commands, man, xman and catman, still provide support for uncompressed manpages. The CDE online help viewer also automatically uncompresses reference pages when they are accessed via a hyperlink in a help volume.

For more information, refer to the man(5) and the catman(8) reference pages.

1.5.5 Enhancements to terminfo

Terminal support has been enhanced to support non-Digital terminals. Entries have been added to the terminfo databases and the termcap file to enable this support. New tools have also been added to assist users in modifying or porting other termcap and terminfo entries to Digital UNIX. These include the following:

- captoinfo – Converts termcap files to terminfo entries.
- infocmp – Uncompiles and, if required, compares terminfo entries.

The tput and tic utilities have also been enhanced.

1.5.6 GNU emacs Version 19.28

GNU emacs has been updated to Version 19.28. This version is not upwardly compatible with GNU emacs Version 18.5, the previous version shipped with Digital UNIX. Refer to the appropriate GNU emacs documentation in /usr/lib/emacs/etc.

1.5.7 Netscape Navigator

This release contains Version 1.12I of the Netscape Navigator World Wide Web browsing program. Invoke the Netscape Navigator from a CDE desktop icon, located in the CDE Application Manager's Desktop_Apps group. The Netscape Navigator can be invoked directly from the command line by running /usr/bin/X11/netscape. You can access detailed help on the Netscape Navigator through the help menus.

A Digital UNIX home page can be found in /usr/doc/netscape in the the file named Digital_UNIX.html. The home page contains links to helpful documentation, including a local copy of the *Netscape Navigator User's Handbook* (consider adding this link to your list of Netscape bookmarks).

The Digital UNIX *Installation Guide* contains information on how to set up Netscape. See Chapter 6, which covers postinstallation setup tasks.

1.5.8 dxterm

The `dxterm` terminal utility has the following new features:

- Default resource file naming scheme
- A `loginShell` resource
- A `useWMHints` resource

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

1.5.9 Performance Manager

Performance Manager is a real-time performance monitor that allows users to detect and correct performance problems. Graphs and charts can show hundreds of different system values, including CPU performance, memory usage, disk transfers, file-system capacity, and network efficiency. Thresholds can be set to alert you to correct a problem when it occurs, commands can be run on multiple nodes from the graphical user interface, and archives of data can be kept for high-speed playback or long-term trend analysis.

1.5.10 utilupdate

A new utility, `utilupdate`, updates the `setld,` and optionally `ris` and `dmu` utilities on servers that are running a Version of Digital UNIX prior to Version 4.0. It is necessary to run this utility on the server prior to attempting to serve Digital UNIX Version 4.0 clients.

The `setld`, `ris`, and `dmu` utilities have been modified for Version 4.0 to allow them to work with the new format of the distribution media. For more information on the use of this utility, refer to the *Sharing Software on a Local Area Network* manual.

1.5.11 Bootable Tape

This release introduces the ability to create a standalone bootable tape of the operating system. You can boot from the bootable tape as easily as you can boot from CD-ROM or a RIS area, but without the overhead of selecting or installing subsets. When you restore your system from the bootable tape, you must reconfigure your system using the System Management applications. You will need to adjust system parameters, such as the host name or IP address,

The binaries and shell scripts needed to create and restore a bootable tape are installed with the base operating system. The files reside in `OSFBINCOM400` and no other subsets are needed. `OSFBINCOM400` is the Kernel Header and Common Files (Kernel Build Environment) subset. The files used by the bootable tape utility include:

- /usr/lib/sabt/etc/fstab
- /usr/lib/sabt/etc/inittab
- /usr/lib/sabt/etc/profile
- /usr/lib/sabt/sbin/finder
- /usr/lib/sabt/sbin/pickapart
- /usr/sys/bin/btcreate
- /usr/sys/bin/btextract
- /usr/sys/bin/fsmrg
- /usr/sys/bin/mksastape
- /usr/sys/bin/mktape
- /usr/sys/bin/pmerge
- /usr/sys/bin/sboot

You use the `btcreate` utility to create a standalone bootable tape. To extract and restore file systems from tape at the single-user level, you use the `btextract` utility.

For more information, see the `btcreate`(8) and `btextract`(8) reference pages. See Section 4.5.2 for information on supported systems and current restrictions on use.

1.5.12 Partition Overlap Checks Added to Disk Utilities

Partition overlap checks have been enhanced or added to the following commands:

```
newfs    ufs_fsck   mount
addvol   rmvol      rmfdmn
swapon   voldisk    voldisksetup
```

The checks ensure that partitions will not be overwritten if they are marked in use in the `fstype` field on the disk label. The overlap checks also ensure that no overlapping partition is marked in use.

If a partition or an overlapping partition has an in-use `fstype` field in the disk label, the following commands inquire interactively if a partition can be overwritten or not:

```
newfs    mkfdmn    addvol
swapon   voldisk   voldisksetup
```

Refer to the reference pages for more information.

Partition overlap checks have been generalized by creating two library functions: `check_usage` and `set_usage`. Two new `fstype` values have been added: `FS_RAW` and `FS_DB`. For example, the library function `set_usage` could be used by database applications to set the `fstype` field of a disk partition that is in use by the database. Similarly, `check_usage` can be used to determine the usage of a disk partition or any overlapping partition.

1.5.13 scsimgr Utility for Creating Device Special Files

The `scsimgr` utility creates device special files for newly attached disk and tape devices. This utility is automatically invoked at system boot time. You can execute the command to add device special files for all disk and tape devices attached to a specified SCSI bus at any time. See the `scsimgr`(8) reference page for further details.

1.6 SysMan System Management Applications

This release includes a new suite of graphical single system management applications called SysMan. These applications support installation, configuration, daily administration, and monitoring and tuning. Some of the SysMan applications require you to have superuser privileges or offer a restricted interface to non-privileged users.

At installation time, you can use the graphical Installation Setup application to select software and configure disks.

The graphical applications are integrated into the Common Desktop Environment (CDE). The SysMan applications are available from the CDE front panel by clicking on the Application Manager icon and double clicking on the System_Admin group icon.

There are five groups of system management utilities:

- Configuration

  The configuration applications modify system files that control various aspects of network connectivity and system configuration.

- Daily Administration

  The daily administration applications manage users, file systems, remote hosts, licenses, and backup; one application assists you in shutting down the system.

- Monitoring and tuning

  The monitoring and tuning applications monitor processes and tune kernel parameters.

- Storage management

  The storage management applications manage disk organization and performance.

- Tools

  The tools provide quick status reports.

The following sections list the system management utilities.

1.6.1 Installation

This release provides full installations for new systems, and update installations to update systems from Digital UNIX Version 3.2C, Version 3.2D-1, or Version 3.2D-2 to Digital UNIX 4.0.

1.6.1.1 New Installation User Interfaces

Two new user interfaces are available for the Digital UNIX full installation process:

- A text-based, menu-driven interface
- A graphical interface

The type of interface presented during the full installation is determined automatically based on your hardware. Systems with graphical consoles and 32MB or greater main memory will provide a graphical interface to the installation. Systems with consoles that do not have graphics capabilities or have insufficient memory to support graphical installations provide a text-based interface.

1.6.1.2 New Format of Distribution Media

As in previous releases, there are two types of media to install the Digital UNIX operating system onto your system:

- From a CD-ROM that contains the Digital UNIX base operating system Version 4.0
- Over a network connection to a remote installation services (RIS) server that is serving Digital UNIX Version 4.0

However, for Version 4.0, the CD-ROM contains file systems that are laid out just as the software would be installed on the system and contains directly accessible `root`, `/usr`, and `/var` areas. This format makes every operating system command and utility available to the installation process. This means that UNIX commands required for recovery procedures such as restoring corrupt file systems are readily available even if your operating system is not yet fully functional. The new format also eliminates the need for the environment previously known as the Standalone Environment, which was a primitive, limited operating system environment.

1.6.1.3 Installation Name Changes

In Digital UNIX Version 4.0, you can perform a default or custom full installation. In previous releases, these installation types were called basic and advanced, respectively.

The installation procedure still provides access to a UNIX shell that lets you recover from serious system problems such as root file system corruption or to perform general file system or disk maintenance tasks before or during the installation. In previous releases, this option was called the System Management option. In Version 4.0, it is known as the UNIX Shell.

1.6.1.4 Enhancements to the Installation Procedure

The following enhancements have been made to the installation procedure:

- During Custom installations, dependencies between optional software subsets are checked and the additional software is selected automatically.
- The amount of space remaining in the `root`, `/usr`, and `/var` areas is displayed as you select each optional software subset.
- If you are using the graphical interface, the custom installation provides access to the graphical disk partitioning utility for you to resize disk partitions.
- Online help is available for text-based and graphical installations.

1.6.1.5 Cloned Installations

The ability to perform a cloned installation is available in this release. A cloned installation enables you to duplicate the file system layout, file system type, and software subset selections from a similar system that has already been installed with Digital UNIX Version 4.0. Installation cloning can only be performed using RIS. If your system is registered to a RIS environment and has a configuration description file (CDF), the installation procedure retrieves the CDF and uses the system configuration information stored in the CDF to configure and install your system.

1.6.2 Configuration Applications

After a full installation, the SysMan Configuration Checklist will be displayed and can be used to configure your system. The order in which the applications are presented is a logical order in which to configure your system. You can skip any applications that are not needed for the current configuration.

The following graphical applications are also available in the Configuration group under System_Admin in the Application Manager group once the installation is complete:

- BIND configuration
- Disk configuration
- Mail configuration
- Network configuration
- NFS configuration
- Print configuration
- Internationalization configuration

In addition, the CDE Application Manager has a launchpoint for NIS Setup, which is a command-line interface for setting up Network Information Services (NIS), such as centralized password and user group files.

Each of the graphical configuration applications has a corresponding setup script that can be used through a character-based interface.

1.6.2.1 Bind Configuration

The bind configuration utility, `bindconfig`, provides streamlined management of BIND and DNS (Domain Name Services) services on a host. Bind configuration supports the creation and management of name service clients and servers and provides the complete range of configurations permitted by BIND.

Use this application to:

- Configure the host system as a master name server for one or more zones, as a caching-only name server, or as a name service client.
- Initialize and maintain BIND data files and generate resource records to manage areas such as the host-name-to-address and address-to-host-name mappings, host information, host name aliasing, name servers, mail exchange, and the cache file root server list.
- Manage the BIND boot files including creation of multiple zones of authority, designation of forwarders, indication of slave configurations, specify the BIND database directory, and specify the BIND cache file name.
- Maintain the resolver configuration file, specify queriable name servers, and specify an alternate domain search list.
- Modify the desired services for the hosts database service selection configuration file (`/etc/svc.conf`).

- Enable and disable the `named` daemon.
- Modify the host name of the system by modifying the BIND local domain name of the system.

1.6.2.2 Disk Configuration

The disk configuration utility, `diskconfig`, manages certain disk operations for the base system.

Use this application to:

- Modify disk configuration parameters
- Display attribute information for existing disks
- Create file system partitions on a disk
- Add disk alias names

1.6.2.3 Mail Configuration

The mail configuration utility, `mailconfig`, configures the `sendmail` utility.

Use this application to:

- Set up the routing and delivery of mail for any of the following types of workstations:
  - A standalone computer
  - A mail client (where mail is handled by another computer)
  - A mail server that addresses and routes mail to other systems
- Use NFS to share mailboxes across multiple hosts, even standalone systems.

1.6.2.4 Network Configuration

The network configuration utility, `netconfig`, manages network hardware and related services on a host. This application supports the configuration and management of networking adapters on a host, and provides maintenance of network services and related data files.

Use this application to:

- Configure network hardware interfaces
  - Ethernet
  - FDDI

- Token Ring

- Enable and disable the following daemons:
  - `gated`
  - `routed`
  - `rwhod`
- Configure the system as an IP router
- Add, modify and delete entries in the following files:
  - `/etc/routes`
  - `/etc/gateways`
  - `/etc/hosts`
  - `/etc/hosts.equiv`
  - `/etc/networks`

1.6.2.5 NFS Configuration

The NFS configuration utility `nfsconfig`, manages the Network File System (NFS) and related processes on a host.

Use this application to:

- Manage the configuration of the system as an NFS server, NFS client, or both. NFS client daemons are started only if the system is configured as a NFS client.
- Manage remote file systems by accessing the File Sharing application.
- Manage the configuration of the NFS daemons.

1.6.2.6 Print Configuration

The print configuration utility, `printconfig`, manages the `printcap` file, print daemon specific directories, and files.

Use this application to:

- View lists of configured printers and available printer type definitions
- Select a printer from the list of configured printers
- Modify basic and advanced settings for a selected printer
- Designate a selected printer as the system default printer
- Verify configuration changes and make modifications without leaving the application

- Configure local printers (that is, printers connected from an stty port) and printers served remotely (that is, attached to another Digital UNIX host)
- Configure printers connected to a LAT application port
- Create and configure the LAT device and application port

1.6.2.7 Internationalization Configuration

Internationalization configuration is available if the optional worldwide support subsets are installed. See Section 1.13 for more information.

1.6.3 Daily Administration Applications

There are eight applications for the daily administration of your system:

- Account Manager
- Archiver
- File Sharing
- Host Manager
- License Manager
- Shutdown
- System Information
- Audit Manager
- Display Window
- Power Management

1.6.3.1 Account Manager

The account manager utility, `dxaccounts`, is used to manage user accounts. It operates on both base security level systems and enhanced security (C2) level systems. You must have root privileges to modify account databases.

Use NIS to centrally manage user accounts in a network environment. NIS allows participating systems to share a common set of `passwd` and `group` files.

The account manager allows you manage both the local and NIS account databases. NIS uses a client-server model. To make changes to the NIS databases, you must run the account manager on the machine designated as the NIS server.

The account manager offers views of each of the system databases. By default, the account manager starts by displaying the Local User view showing the contents of the `/etc/passwd` file. The supported views are:

- Local Users (`/etc/passwd`)
- Local Groups (`/etc/group`)
- NIS Users (`/var/yp/src/passwd`)
- NIS Groups (`/var/yp/src/groups`)

Please refer to the account manager online help for more information.

1.6.3.2 Archiver

The archiver utility, `dxarchiver`, provides graphical access to the `tar`, `cpio`, `pax`, and `compress` commands.

Use this application to:

- Copy and store multiple files to a single, named archive file or output device such as a tape or floppy disk
- Uncompress incoming archive files and compress newly created files
- Retrieve stored files from an archive.

Archiver functions are classified as follows:

- Archiver input and output information

  The input and output areas provide icon containers and user text areas. You can drag icons that represent files and directories from the CDE File Manager, or icons representing tape and floppy drives from the system information application. Any icon that you can drag can also be dropped into the input icon container. Type the name of a file, directory, or device into the user input text area.

- Default Directory

  The Archiver requires you to choose whether to use absolute or relative pathnames. The use of absolute pathnames (those written out in full with an initial slash) and relative pathnames (those written without an initial slash) allows the archiver to control the placement of retrieved files. Files in an archive have either absolute or relative paths. Files with relative paths are restored to the current default directory. That is, they are restored to the directory that is displayed in the default directory area.

1.6.3.3 File Sharing

The file sharing utility, `dxfileshare`, allows the importing and exporting of NFS file systems. You can run it with or without superuser permissions. If you are not root, access to system files is restricted so that you can only view, but not modify, locally exported file systems and perform mounts to user directories.

Use this application to:

- Display mounted NFS file systems that are exported through NFS from the local machine.
- Display file systems available for mounting from a remote set of hosts.
- Perform mount and unmount operations with the ability to modify valid options.
- Add entries to the `/etc/fstab` file as permanent mount entries.
- Add, change, or delete NFS entries from `/etc/export` file.

1.6.3.4 Host Manager

The host manager utility, `dxhosts`, manages X clients. You can use this application to change X host permissions for hosts and to execute X applications remotely on a host. Other applications in this tool suite use it as a host name resource.

Use this application to:

- Drag and drop application icons from the CDE Application Manager onto displayed hosts to execute the application remotely and send the display back to the local screen.
- Double click on a host to execute a default command remotely (defaulted to `xterm`).
- Add or remove a host to or from the `xhost` access control list.
- Enable or disable X Window System access for all or selected hosts.

1.6.3.5 License Manager

The license manager utility, `dxlicenses`, provides access to Digital’s License Management Facility (LMF).

This application allows you to:

- Add, delete, issue, and list Digital UNIX licenses
- Load a license from a file

1.6.3.6 Shutdown Manager

The shutdown manager utility, `dxshutdown`, shuts down and optionally reboots your system.

The following options are available:

- Halt – halts the system and displays the console prompt

- Single-user mode – brings the system to single-user mode
- Reboot – reboots the system
- Message only – sends a message as if a shutdown was initiated, but a shutdown action does not occur.

When the main window expands, the following options are also available:

- Broadcast message – sends a broadcast to NFS Clients.
- Preshutdown script – enables you to specify a script to be run before the system is shut down.
- Fast – shuts down without a broadcast message and without performing file system checks when the system reboots.
- No disk sync – shuts down without flushing the system buffer cache.

1.6.3.7 System Information

The system information utility `dxsysinfo`, displays and monitors system resources on a Digital UNIX system.

Use this application to:

- Display the operating system version, amount of random access memory (RAM), and the number of CPUs on a system.
- Select the following views of system resources:
  - Activity displays the CPU percentage used.
  - Free Memory displays the current percentage of unused physical memory.
  - Available Swap shows the swap space available as a percentage of the total space on the swap device or devices.
  - File/Swap Warning Lights turns on the warning lights for file systems and swap space.
  - The File System Information area lists the mounted file systems. The user can select whether local, remote, or both local and remote file systems are displayed. For each file system, it shows the percentage of disk space used, the amount of disk space free, and an adjustable limit that triggers a warning light when a file system is full.
  - The Device Information area displays an icon for each disk drive, CD-ROM drive, tape drive, and floppy disk drive configured on the system.

1.6.3.8 Audit Manager

The Audit Manager, `dxaudit`, allows you to create audit reports. For more information see the *Security* manual and the `dxaudit`(8) reference page.

1.6.4 Monitoring and Tuning Applications

The applications described in this section help you to monitor and tune your system.

1.6.4.1 Kernel Tuner

The kernel tuner utility, `dxkerneltuner`, manages attributes of loadable kernel subsystems.

Use the kernel tuner to:

- Display a list of kernel subsystems
- Display boot time and current values of kernel subsystem attributes
- Load values for kernel subsystem attributes from a file
- Modify boot time values of kernel subsystem attributes
- Modify current values of updatable kernel subsystem attributes
- Save values of kernel subsystem attributes to a file

1.6.5 Storage Management

The CDE Application Manager provides launchpoints for Storage Management applications in the System Admin group.

The applications that are available depend on what software you have installed. A launchpoint for the Logical Storage Manager is available with the base operating system.

1.6.6 Tools

The CDE Application Manager provides launchpoints for Tools in the System Admin group.

The System Management suite includes four launchpoints in this area. Each runs an output-only command in a window that permits you to repeat the command on demand or at intervals.

The launchpoints and their commands are:

- I/O statistics: `iostat`
- Network statistics: `netstat`

- Virtual Memory statistics: `vmstat`
- Who>? `w`
- System Messages:

  `tail -f /var/adm/messages`

If you want to run another output-only command, you can run one of the commands in this group, then change the current command in the user interface.

1.7 Standards

This release complies with many new and changes standards. Refer to the `standards`(5) reference page for more information.

1.7.1 Realtime Is Compliant With the Final POSIX 1003.1b Standard Interfaces

As of this release, Digital UNIX completes the implementation of the POSIX 1003.1b standard interface as approved by the IEEE standards board in September 1993 (IEEE Std 1003.1b-1993, Realtime Extension). See the Guide to Realtime Programming for more information. The new features are described in sections Section 1.8.11, Section 1.8.12, and Section 1.8.13.

1.7.2 DECthreads is Compliant with the Final POSIX 1003.1c Standard Interfaces

As of this release, the DECthreads library `libpthread.so` implements the POSIX 1003.1c standard interface as approved by the IEEE standards board in June 1995 (IEEE Std 1003.1c-1995, POSIX System Application Program Interface). The new POSIX (pthread) interface supported with DECthreads is the most portable, efficient, and powerful programming interface for a multithreaded environment. These interfaces are defined by `pthread.h`. See the *Guide to DECthreads* for more information.

1.7.3 DEC C Compiler Complies With ISO C 94

In this release the DEC C compiler and `libc` implement the ISO C 94 standard. This behaviour is obtained by calling the DEC C compiler with the `-isoc94` flag. This feature is not available with the `-oldc` compiler option, which invokes the previous compiler.

1.8 Development Environment

This release includes the following enhancements to the development environment.

1.8.1 Tcl/Tk Availability

Tcl/Tk is now available as part of the base operating system. Tcl/Tk is a public domain unencumbered scripting language and graphical tool kit. In addition to Tcl/Tk, a popular extension package, TclX is also included. TclX provides many UNIX extensions to the Tcl command language. Tcl version 7.4, Tk version 4.0, and TclX version 7.4 are included in this release. See the *Installation Guide* for information on how to identify and install the appropriate software subsets.

The available programs are:

- `/usr/bin/tcl` – A tcl shell with TclX extensions.
- `/usr/bin/tclsh` – A hard link to `/usr/bin/tcl`.
- `/usr/bin/wishx` – A Tcl/Tk/tclX shell.
- `/usr/bin/wish` – A hard link to `/usr/bin/wishx`.
- `/usr/bin/tclhelp` – A graphical help browser for tcl help.

1.8.2 DEC C++

The following changes have beem implemented for DEC C++.

- DEC C++ version 5.3 Class Library is now threadsafe

  Refer to the *DEC C++ Class Library Reference Manual* for details on the threadsafe support, including a new Mutex Package.

- Complex division catches divide-by-zero errors

  The division routines within the Complex Library now catch divide-by-zero errors instead of signaling them.

- Iostream assignment operators

  For iostream assignment operators, there is no longer a memory leak when you use the `*_withassign` assignment operators to initialize an object for which you have called `xalloc()`. Previously, the memory allocated for the object by `xalloc()` was lost.

- String extraction operator

  The String extraction operator now takes care of dynamically allocating the String to accommodate the input.

- `ios::ate` mode

  When you open a file specifying `ios::ate` but not `ios::app` to the `filebuf open()` function, the file is no longer opened in `O_APPEND` mode. This incorrect behavior caused all data to be written to the end of the file, regardless of the current file position.

- Exception handling

  Various problems with exception handling have been fixed. Also, support for exception handling in DEC C++ version 5.3 has been added.

- Function `exp()` returns zero for underflow errors

  When the Complex Library `exp()` function detects an underflow error, the resulting value is now (0,0) instead of (+/- max-float, +/- max-float).

- Use of `clog()` and C++ Class Library iostream `clog`

  A single application is restricted from using both the math library function `clog()` and the iostream package's `clog` object. This restriction is due to the fact that `libm` and `libcxx` each contain a definition for the global symbol `clog` and those definitions are incompatible.

  Furthermore, applications which reference one of the `clog` symbols cannot include both `-lcxx` and `-lm` on their `ld` command line. An error will be generated by `ld` because `clog` is multiply defined.

- `catch(...)`clause

  The `catch(...)` clause now catches C structured exceptions.

- `fstream close()` clears the error state

  The `fstream`, `ifstream`, and `ofstream close()` member functions now clear the stream's error state when the close succeeds. Call the `clear()` member function after the call to `close`().

1.8.3 Replacement of the pixie Utility

The previous version of the `pixie` utility has been relocated to the file `/usr/opt/obsolete/usr/bin/pixie` and is provided as part of the `OSFOBSOLETE` subset. The previous version of `pixie` will be retired in a future release and is replaced with the Atom-based `pixie` utility. A driver for the Atom-based `pixie` is located in `/usr/bin/pixie`, for backward compatibility with the previous version. The driver accepts the same switches recognized by the previous version. Refer to the `pixie`(1) and `atom`(1) reference pages for more information on `pixie` and general information on the Atom tools.

1.8.4 Software Development Environment Repackaging

The Software Development Environment (SDE) has been repackaged to ease installation, simplify licensing, and create a product identity. The current SDE components have been repackaged into a single, new `OSFSDE` subset and all of the pieces outside the SDE have been moved into logical subsets, including:

- `OSFINCLUDE` for all include files
- `OSFLIBA` for all static libraries
- `OSFPGMR` for commands outside the scope of the SDE

Because the compiler is needed at installation time, some SDE components have remained in the mandatory `OSFCMPLRS` subset.

The Ladebug debugger subsets have been renamed to the `OSF*` subset name prefix and can now be installed during a custom installation of Digital UNIX. These changes have been made on the Digital UNIX Operating System Volume 1 CD–ROM. The FUSE Porting Assistant has been added to the Digital UNIX kit on the Digital UNIX Associated Products Volume 1 CD–ROM. This is a tool to help port code to Digital UNIX from a variety of platforms and operating systems.

The `OSFSDECDE` subset was also added to the Digital UNIX Operating System Volume 1 CD–ROM. It contains the files necessary to access DECladebug and the Porting Assistant from CDE.

1.8.5 DEC C Compiler

The C compilers, `cc` and `c89` on Digital UNIX have changed to Digital's DEC C compiler. The previous default C compiler is still a supported viable option and can be accessed by specifying `-oldc` on the `cc` and `c89` command lines.

DEC C uses Digital's back-end compiler (GEM) technology, which has been specifically developed and optimized for use with Alpha systems. Both compilers have full binary compatibility with each other.

1.8.6 Init Execution Order Modified for Static Executable Files

The execution order for `init` routines in static executable files has been modified to more closely match the execution order for `init` routines in dynamic executable files. The init routines loaded from an archive library will be executed prior to any init routines loaded from objects and archives occurring earlier on the linker command line. Prior to this change, `init` routines were executed in the order they were encountered in processing the `link` command from left to right. As a result, init order for static executable files was much different than the init order for equivalent shared executable files.

For existing applications that rely on the static `init`o rder used in prior releases of Digital UNIX, the new linker option: `-old_init_order` can be used to restore the strict left-to-right execution order for static executable files.

1.8.7 PC-Sample Mode of prof Command

The `prof` command's pc-sampling mode now supports profiling the shared libraries used by a program. Linking a call-shared program with the `cc` command's `-p` switch causes the resulting program to profile both the call-shared executable file and all the shared libraries. The following command displays a combined profile:

```
# prof -all
```

New `-all`, `-incobj`, `-excobj`, and `-stride` switches for the PROFFLAGS environment variable enable you to request per-procedure profiling of the shared libraries or to select particular libraries to profile.

Related enhancements are:

- Extended APIs to `monitor()`, `monstartup()`, and `profil()`
- Use of 32-bit pc-sampling counters instead of 16-bit for `cc -p` and `cc -pg` profiling (`gprof`), except for calls to the traditional `monitor()` API.
- Improved reliability in profiling multithreaded programs, and reference page guidelines for use of `monitor_signal()` with threads.
- `prof` and `gprof` checking.
- Profiling report formats are improved.

See the `prof(1)` and `monitor(3)` reference pages for further information.

1.8.8 atom and prof Commands and Threads

The following `atom` and `prof` commands now profile the shared libraries used by a program:

```
# atom -tool pixie -all
```

and

```
# prof -pixie -all
```

Also, the `threads` environment for the`atom`makes `pixie` tool thread-safe, though per-thread counts are not recorded.

See the `atom` (1), `prof` (1), and `pixie` (5) reference pages for further information.

1.8.9 Thread Independent Services (TIS) Interface

Digital UNIX Version 4.0 introduces the Thread Independent Services (TIS) application programming interface in the C run-time library `libc`. TIS provides services that assist in the development of thread-safe libraries.

Thread synchronization may involve significant run-time cost, which is undesirable in the absence of threads. TIS enables thread-safe libraries to be built that are both efficient in the non threaded environment, yet provide the necessary synchronization in the threaded environment.

When DECthreads (pthreads) are not active within the process, TIS executes only the minimum steps necessary. Code running in a non threaded environment does not encounter overhead incurred by the run-time synchronization that is necessary when the same code is run in a threaded environment. When DECthreads are active, the TIS functions provide the necessary thread-safe synchronization.

1.8.10 High-Resolution Clock

Digital UNIX Version 4.0 has an optional high-resolution clock. To enable this option, add the following line to the kernel configuration file and rebuild the kernel:

```
options MICRO_TIME
```

The system clock (`CLOCK_REALTIME`) resolution as returned by `clock_getres` will not change. Timer resolution remains the same. However, time as returned by the `clock_gettime` routine will now be extrapolated between the clock ticks. The granularity of the time returned will now be in microseconds. The time values returned are SMP safe, monotonically increasing, and have 1 microsecond as the apparent resolution.

1.8.11 POSIX 1003.1b Realtime Signals

Realtime signals have been implemented to conform to the POSIX 1003.1b standard. This new feature includes queued signals with optional data delivery, and 16 user-definable realtime signals.

The following functions to support realtime signals were implemented:

- `sigqueue`
- `sigtimedwait`
- `sigwaitinfo`
- `timer_getoverrun`

1.8.12 POSIX 1003.1b Synchronized I/O

Synchronized I/O (file synchronization) has been implemented to conform to the POSIX 1003.1b standard. New functions for synchronized I/O under the UFS and ADVFS file systems include:

- `aio_fsync`

  Asynchronously writes changes in a file to permanent storage

- `fdatasync`

  Writes data changes in a file to permanent storage

The `open` function now takes the following new flags for synchronized I/O:

- `O_DSYNC`

  Ensures synchronized I/O data integrity of the file accessed

- `O_RSYNC`

  Used for synchronized I/O read operations

1.8.13 POSIX 1003.1b _POSIX_C_SOURCE Symbol

For applications conforming to POSIX 1003.1b, the `_POSIX_4SOURCE` macro is supported for this release, but will be retired with the next release of Digital UNIX. The macro `_POSIX_4SOURCE` is part of an obsolete draft standard and is supported in this release for compatibility only. When possible, existing applications which use `_POSIX_4SOURCE` should be modified to use `_POSIX_C_SOURCE` instead.

The `_POSIX_C_SOURCE` macro is associated with a value, which allows an application to specify the namespace it requires. However, as a general rule, avoid explicitly defining standards macros when compiling your applications. For most applications, the header file `unistd.h` provides the standards definitions needed.

1.8.14 Digital Porting Assistant

The Digital Porting Assistant is a Motif-based tool to help you port your C, C++, and FORTRAN source code to Digital UNIX from other UNIX and proprietary platforms, including OpenVMS. The Porting Assistant provides the following features:

- uncovers 32-bit dependencies
- Checks your makefile commands and options
- Helps find functions that your application needs
- Helps develop code segments specific to Digital UNIX

- Provides additional information on porting your application

The Porting Assistant is licensed and provided to you with Digital UNIX Developers' Toolkit but requires separate installation.

To install Version 2.0 of the Porting Assistant, install subsets `PRTBASE200` and `PRTMAN200` (and their dependencies) from the *Digital UNIX Associated Products Volume 1* CD–ROM..

1.9 Networking

The following networking enhancements have been implemented.

1.9.1 New Version of the gated Daemon

This release includes a new version of the `gated` routing daemon. The update installation procedure will detect if your system is configured to run the `gated` routing daemon. If the Digital-supplied `gated` is detected then the `/etc/gated.conf` file is moved to `/etc/ogated.conf.` Otherwise, if a user supplied or customized `gated` is detected, then both the `/etc/gated.conf` and the `/usr/sbin/gated` are saved with the `.PreUPD` suffix.

When the system is installed, the new gated R3.5 is the default version in `/usr/sbin/gated`. The old `gated` version 1.9 is supplied in `/usr/sbin/ogated`. Also, corresponding older `gated` reference pages are saved with an `o` prefix.

1.9.2 Dynamic Host Configuration Protocol (DHCP)

This release contains both a client and a server Dynamic Host Configuration Protocol (DHCP) daemon. For DHCP client configuration, use the `netconfig` utility. For configuration of client parameters on the DHCP server, use the `/usr/bin/X11/xjoin` utility, which provides a graphical user interface to the `/etc/bootptab` file.

For more information on DHCP, refer to the `joinc`(8) and `joind`(8) reference pages.

1.9.3 Point-to-Point Protocol (PPP)

This release supports Point-to-Point Protocol PPP, including support for BSD-style compression of entire packets. This is a negotiated option. If a foreign peer cannot handle this, it should be gracefully rejected via the `Protocol-Reject` of LCP.

When using PPP with modems doing compression, it may be desirable to force no BSD-style compression. To do this, put `-bsdcomp` in either

/etc/ppp/options, or on the pppd command line.

PPP now has a configurable (at boot time) number of interfaces. The default is 1. To specify a higher value, add the following line to the /etc/sysconfigtab file and reboot the system:

```
ppp:nppp=x
```

PPP documentation is available in the pppd(8), pppstats(8), and chat(8), reference pages and in the *Network Administration* Guide.

1.9.4 NFS over TCP

This release contains Digital's first implementation of NFS support over the TCP transport. Previously, only the UDP transport was available. UDP is still available and may still be the preferred transport in local area networks. For NFS access over wide area, congested, or lossy networks, TCP may offer better performance.

Use the nfssetup script or the nfsconfig utility to enable TCP on the NFS server.

The nfsiod daemon has been changed to spawn kernel threads (instead of forking multiple processes as it did previously). Each nfsiod thread can handle UDP or TCP mounts, so the command nfsiod still accepts one argument.

1.9.5 Extensible SNMP

A new SNMP architecture is present in this release. The SNMP daemon, snmpd, is now an extensible master agent. End user programmers can develop subagent programs that communicate with snmpd to implement their MIBs on Digital UNIX systems.

The base operating system MIB support is implemented in a subagent program called os_mibs, which is started/stopped automatically with snmpd.

1.9.6 SNMP MIB support

This release supports the Host Resources MIB (RFC 1514). The MIB support daemon must query the system's devices to retrieve information required for this MIB. This query occurs when the daemon starts, and subsequently whenever a relevant SNMP request arrives.

This device querying is the default behavior, and may be configured off. See the snmpd(8) reference page for more information about configuring the SNMP agent.

1.9.7 telnet and mrouted

The following new features were implemented:

- A `NEW-ENVIRON` parameter was added to implement RFCs 1571 and 1572 for `telnet`
- Version 3.6 of `mrouted` was added, including the utilities:
  - `mtrace`
  - `map-mbone`
  - `mrinfo`

1.9.8 Network Time Protocol (NTP)

The `xntp` utility has been upgraded to Version 3.0, which is a complete implementation as defined in RFC 1305.

Note that the new `ntpdate` utility requires that the version of the server be specified with the `-o` flag. For example:

```
% ntpdate -o 2 server.oxo.fff.com
```

The default version is `3`.

The new `xntpd` daemon supports authentication. If you choose to enable this option, note that you only have to provide keys to the servers and peers that also support authentication.

You should provide names of three servers during setup with `ntpsetup`, choosing the mode `server` and the version option `V2` (In the case where you are not upgrading the servers).

1.10 Enhanced Security

This release provides the following new enhanced security features.

- support for per-user resource limits in user profiles, using `setrlimit`
- non-shadowed passwords are allowed, while using other extended profile features.
- The system administrator can control whether the `ttys` database is updated on logins.
- Wildcard support for `ttys` has been extended to X displays.
- User profiles and `ttys` information are stored in database files for faster access and update (resulting in faster logins).
- The new utilities `edauth` and `convuser` are available.

See the *Security* manual and the `setrlimit`(2), `edauth`(8), and `convuser`(8) reference pages.

1.11 File Systems

The following file system enhancements have been implemented in this release.

1.11.1 AdvFS

The following new features were introduced for AdvFS.

1.11.1.1 New Tuning Parameters for AdvFS

There is a new mechanism for limiting the amount of kernel memory that AdvFS uses for its access structures. This may be necessary only for systems with 64 Mb or less memory, and AdvFS as the default file systems. This is applicable to all hardware configurations.

There are two new kernel paramaters relevant to AdvFS that can be modified using the `sysconfig` or `sysconfigdb` commands. They are `AdvfsAccessMaxPercent` and `AdvfsAccessCleanupPercent`. There is a complete description of these parameters in the *Guide to File System Administration for the POLYCENTER Advanced File System and Utilities for Digital UNIX*.

1.11.1.2 AdvFS Now Supports Directory Truncation

Traditionally, AdvFS directories were never truncated, even though many of the files in the directory had been deleted. This created a problem if the directory file became very big. For example, if several hundred thousand files were added to a directory, then the directory file itself grew very large. Even though most of the files in that directory were subsequently deleted, operations that required scanning the directory remained inefficient because the entire directory file still needed to be read.

AdvFS now truncates directory files when all of the entries at the end of the directory have been deleted. This truncation is done on 8KB byte boundaries, so the size of a directory is always a multiple of 8192.

One ambiguity of directory truncation is that the truncation is done when files are created and not when they are deleted. This is done because of the efficiency of underlying algorithms, and is the same model used by UFS for directory truncation. For example, after most files in a given directory are deleted, the size of the directory file itself will not decrease until a new file is inserted into that directory.

1.11.2 File System Access Control Lists (ACLs)

ACLs (Access Control Lists) on files and directories are a new feature in this release. They are manipulated with the `getacl` and `setacl` commands. See the *Security* manual and the reference pages for more information.

1.11.3 Logical Storage Manager

Digital UNIX Version 4.0 provides the following new features for the Logical Storage Manager (LSM):

- Two new LSM commands, `volsave` and `volrestore`, provide an easy way to back up and restore the LSM configuration database. See the reference pages for these commands.
- The Basic Operations menu in LSM's graphical interface, `dxlsm`, now provides support for disk operations. For example, adding a disk to LSM.
- The LSM limits have increased as follows:
  - The maximum number of LSM volumes on a system has increased from 256 to 4093.
  - The maximum number of plexes on a system has increased from 256 to 4096.
  - The maximum number of subdisks in a plex has increased from 256 to 4096.
  - The maximum number of disks that can be added to LSM has increased from 128 to 256.
  - The maximum size of an LSM volume has increased from 128GB to 512GB.

The functionality and syntax of the LSM commands used for encapsulation, unencapsulation, and mirroring have changed in this release, as follows:

- The `volencap` command now supports the following features and functions. For details, refer to the `volencap`(8) reference page.
  - Allows the initialization of LSM and encapsulation of the system disk in one step. This requires the use of a free partition table entry.
  - Can be used to encapsulate all partitions on a disk. This requires the temporary use of a free partition table entry if the system disk is being encapsulated.
  - Can be used to encapsulate only the root and swap partitions.
  - Automatically creates a new disk group if specified.

- Subsumes the functionality of the `voladvdomencap` command.
- Takes multiple arguments.
- Uses a simple disk instead of a sliced disk for system disk encapsulation.
- For disk label characteristics, assumes that partition `c` maps the entire disk, and that an in-use partition has an *`fstype`* field other than `UNUSED`. (If a partition's *`fstype`* field is `UNUSED`, then `volencap` may allocate that partition table entry for its use.)

- The `volrootmir` command now supports the following features and functions. For details, refer to the `volrootmir`(8) reference page.
  - Can be used to mirror all volumes on the system disk by specifying the `-a` option. This option requires the target disk to be of the same type as the source disk.
  - Can be used to encapsulate only the root or swap partition by omitting the `-a` option. This procedure requires that the target root and swap partitions are large enough to hold `rootvol` and `swapvol`, but the target and source disks need not be of the same type.
- When used with the `-a` option, the `volunroot` command unencapsulates all LSM volumes on the system disk, not just `rootvol` and `swapvol`. The requirements for unencapsulation are:
  - The partition associated with the volume must have been initialized as a `nopriv` disk.
  - The volume must map directly to the partition (that is, the volume size must be equivalent to the partition size).
  - The volume must not be mirrored.

For details, refer to the `volunroot`(8) reference page.

1.11.4 Overlap Partition Checking

Two new functions, `check_usage` and `set_usage`3 are available for use by applications. These functions check whether a disk partition is marked for use and set the fstype of the partition in the disk label. See the reference pages for these functions for more information.

1.12 Bootstrap Linking and Kernel Modules

In Digital UNIX 4.0, a new method for configuring and building kernels is in place. It is designed to meet future needs for modularity, configurability, maintainability, scalability, and third-party device installation. Although the internal details of building kernels has changed significantly, the external interface is almost identical. The *System Administration* guide and *Writing Device Drivers: Tutorial* guide contain more detailed information on this process. For normal operations, a static `vmunix` is recommended. Bootstrap linking is intended to support foreign devices, where support for such devices is not delivered in Digital UNIX.

Previously, the Digital UNIX kernel was comprised of over 1000 individual `.o` files. The configuration procedure, `doconfig,` selected among those files, and compiled a static kernel image that was then installed and booted. The `.o` files are still available in the base kit and can be loaded separately if required.

In this release, the `.o` files are grouped into 150-200 modules, with the suffix `.mod`. Each module represents a functional piece of the kernel, such as a device driver or platform support module. These modules are the new form of the original dynamically loadable modules introduced in DEC OSF/1 Version 3.0.

Modules are selected by `doconfig` instead of the individual files previously used. This simplifies the Makefiles, saves considerable disk space, and reduces compilation time.

1.12.1 Defining Modules

There is a new syntax in the `files` file for making modules:

Module Definition	Syntax
`MODULE/[STATIC]/modname`	`optional opt1 ....`
`dir/file1.c`	`module modname`
`dir/file2.c`	`module modname`
`dir/file3.c`	`module modname`

This procedure defines a module: `modname.mod` (note that `.mod` is appended to `modname` automatically). Each module has three files: `file1.c`, `file2.c`, and `file3.c`. When each file has been compiled, the module is linked together, and is ready to build into a kernel.

Optionally, `STATIC` may be specified before the module name. If it is omitted, the object files are merely linked together into a module. If `STATIC` is used, the module is assumed to have a `modname_configure()` routine and `modname_attributes[]` table to make it a configurable and potentially loadable module. Linkages to these routines are made in the kernel's `static_subsys_list[]` table, and the configure routines will be called at system startup. See *System Administration* for more information on modules, making modules loadable, and the use of configure routines.

If an individual file in BINARY is rebuilt, the module that contains it must also be rebuilt to incorporate the change into the kernel. To help with this, an environment variable DO_MODULES is checked. Providing that the value of DO_MODULES is the string AUTO, a module will be automatically relinked whenever an individual file is rebuilt. Kernel developers are encouraged to put their files into modules, and to consider making them loadable if possible.

1.12.2 Bootstrap Linking

To enhance configurability, and help mainstream modules, a new technology called bootstrap linking has been created. This is a way to defer the actual linking of a kernel until boot time. It is provided as an alternative to booting a statically linked kernel.

When you run `doconfig` it builds a static kernel by default. To build a bootlink kernel, the command option `-b` must be explicitly specified.

A bootstrap linked kernel is actually a set of directives in a `sysconfigtab` file that tell the bootstrap program how to build the kernel in memory. This file is generated using the make `sctab` rule, which will be invoked by `doconfig`.

The updated file is placed into `/usr/sys/MACHINE/sysconfigtab`. It must be copied into the root file system, and the two machine-specific modules `MACHINE.mod` and `EXTRAS.mod` copied into the `/sys/BINARY` directory. The `sysconfigtab` file is then given to the bootstrap program instead of `vmunix` to build and boot the new kernel.

See also the `doconfig` (8) and `sizer` (8) Reference pages for more information.

There is one major restriction when using a bootstrap linked kernel. Because there is no `vmunix` executable file, programs that attempt to search the kernel's symbol table will not find one. This includes `dbx` and most`nm,` and `nlist` library call and pass the name of the booted file will work, even if that file is a `sysconfigtab` file. To avoid this restriction in Digital UNIX 4.0, execute:

```
# ld -o <filename> 'sizer -m'
```

This will generate an exact copy of the currently running kernel complete with symbol table in the file <*filename*>.

When building a bootstrap linkable kernel, `doconfig` first builds a static kernel. If the static kernel fails to build, do not attempt to bootstrap link the `sysconfigtab` file. Note that this static kernel will not be identical to the bootstrap-linked kernel as the kernels may be linked at different addresses. Use the `ld` command if required.

1.12.3 Related Changes

To effect full kernel modularization, several internal changes were made to the `config` program and the compile environment for kernel files. These changes are significant for anyone developing kernel drivers or layered products.

- New linker handling of globals.

  To make all modules the same, a `-dc` linker switch was added to build modules. This has the effect of forcing a full declaration and storage allocation for all commons (globals). Previously, if two files both declared a global `int foo;,` the linker would coalesce them into the same symbol when linking the kernel. This still happens within a module, but not across modules. Each module would have its own declaration *foo*, resulting in multiple definitions of the same global. Because of the multiple definition, kernel linking would fail. To avoid this, make all but one declaration external, using the `extern`. variable declaration

- Binary/NotBinary separation.

  The `BINARY` file is no longer configurable. All files that are listed as Binary are compiled, and placed into modules as defined. The `BINARY` configuration file no longer has any information about hardware configuration (bus, devices or controllers) or most options (such as CDFS or NTP_TIME). Neither are any of these header files built in `BINARY`. `BINARY` files that use the `#include` statement to include other files, and the substitution statement `#define`, will no longer compile. For example:

  ```
  #include <device.h>
  ```

  ```
  #define NDEVICE 4
  ```

  This includes files that use the following inclusion statement:

  ```
  #include <data/NotBinary_data.c>
  ```

See the `doconfig`(8) and `sizer`(8) reference pages for more information.

1.13 Internationalization and Language Support

The following new features have been implemented to support internationalization.

1.13.1 Internationalization (I18N) Configuration Utility for CDE

The I18N Configuration Tool, available through the CDE Application Manager, is one of the SysMan system administration configuration tools. The I18N Configuration Tool provides a graphical interface that enables you to configure internationalization-specific settings. It also provides a convenient way to see which countries, locales, fonts, and keymaps are currently supported on your system. Use this tool to remove unused fonts and unrequired country support.

1.13.2 Unicode Support

This release provides a new set of locales and codeset convertors that support the Unicode and ISO 10646 standards. The codeset convertor modules enable an application to convert between other supported codesets and UCS-4.

Digital UNIX also provides a function called `fold_string_w()` that maps one Unicode string to another, performing the specified Unicode character transformations. For more information on the `fold_string_w()` function, see the `fold_string_w`(3) reference page.

For more information on the Unicode support, see the `Unicode`(5) reference page.

1.13.3 The Worldwide Mail Handler No Longer Exists

Worldwide support subsets no longer install internationalized Mail Handler (MH) software in the `/usr/I18N/bin/mh` directory. In Digital UNIX Version 4.0, internationalization features have been merged into the default Mail Handler (MH) whose files are located in `/usr/bin/mh.` Check the value for the `mhpath` resource used to find the DECwindows Mail application. If necessary, change this value to be `/usr/bin/mh.`

1.13.4 Multilingual Emacs (mule)

The `mule` editor is a multilingual version of GNU Emacs and supports the following kinds of characters:

- ASCII (7-bit)
- ISO Latin-1 (8-bit)
- Japanese, Chinese, and Korean (16-bit) as specified by the ISO 2022 standard and its variants (EUC, Compound Text, and so on)
- Chinese in both GB and Big 5 encodings
- Thai as specified by the TIS 620 standard

The `IOSWWMULE400` subset installs Version 2.3 of the GNU mule editor and associated software. Corresponding sources are available in the `IOSWWMULESRC400` subset.

Digital UNIX does not include public domain fonts that you can use with mule. Refer to the `mule-2.3/README.Mule` file installed by the

IOSWWMULESRC400 subset to find out how you can obtain public domain fonts.

The Digital UNIX software is enhanced with lisp libraries that support the dechanzi codeset for Simplified Chinese and the dechanyu codeset and tsangchi input method for Traditional Chinese. These libraries are included in the IOSWWMULE400 subset and installed in the /usr/i18n/mule/lib/mule/site-lisp directory.

For more information about mule, see the mule(1) reference page.

1.13.5 Support for Catalan, Lithuanian, and Slovene

Digital UNIX Version 4.0 includes support for the the Catalan(5), Lithuanian(5), and Slovene(5) reference pages for information about associated codesets, locales, keyboards, and fonts.

1.13.6 man Command Supports Codeset Conversion

The man command can automatically invoke the iconv utility to perform codeset conversion of reference page files. This allows you to install one set of reference pages to support locales that have the same language and territory but different codesets, thereby reducing file redundancy on the system. For more information, refer to the man(1) reference page.

1.14 PCMCIA (PC Card) Support

PCMCIA (PC Card) support is a new feature in this release and is limited to the following capabilities:

- Support of a selected ISA to PCMCIA bridge adapter
- Support on the following platforms:
  - AlphaStation 255,
  - AlphaStation 200,
  - AlphaStation 400,
  - AlphaStation 600,
  - AlphaServer 1000
- 1 modem card, specifically Megahertz XJ2288 (28.8kpbs)
- hot swap capability of PC Cards

1.14.1 Restrictions

The following restrictions apply in this release.

- No support for loadable device drivers for PC Cards.
- If the system does not have enough available IRQ (interrupt) numbers to assign to the PCMCIA devices, PCMCIA devices cannot be configured. To support one PCMCIA adapter the system must have at least 3 unused IRQ numbers available. One IRQ is for the adapter and the other two are for each PCMCIA socket.
- Digital UNIX can support 2 PCMCIA adapters in a system provided that the necessary resources are available. In some systems, availability of interrupt lines will prohibit the use of multiple adapters. If you have sufficient resources and are going to support 2 adapters, the second adapter should be configured to use the I/O address 3E2.
- In order to use fax functions in a fax/modem PC Card, a commercial UNIX fax application software is required.
- The Megahertz XJ2288 is the only modem card fully qualified on Digital UNIX. However other modem cards of similar type (both 14.4kpbs and 28.8kpbs) may work. The following is the list of modem cards that are known to work:
  - card manufacturer: MEGAHERTZ product name: XJ2288
  - card manufacturer: MEGAHERTZ product name: XJ1144
  - card manufacturer: AT&T Paradyne product name: KeepInTouch Card
  - card manufacturer: Digital product name: PCMCIA V.32bis 14,400 Fax
- The selected ISA to PCMCIA bridge adapters are from SCM Microsystems, the SWAPBOX CLASSIC X2 ... Model MMCD-D2 which has the following features:
  - 3.5" Front Access
  - 2 slots (type II + type III) PC Card socket.
  - Standard PC-AT 16-bit ISA bus interface
  - PCMCIA Revision 2.X and ExCA compliant

  and the SWAPBOX PREMIUM COMBO ... Model MMCD-FC2:
  - 3.5" 1.44 Mbyte Floppy Drive Support.
  - One Type I, II or III front-access PC Card socket.
  - One Type I, II or III rear-access PC Card socket.
  - Standard PC-AT 16-bit ISA bus interface

– PCMCIA Revision 2.X and ExCA compliant

However, other ISA to PCMCIA bridge adapters using the Intel i82365SL or compatible chip may also work.

1.14.2 Configuring the PCMCIA Adapter Board from the Console

Before inserting the PCMCIA adapter board into your system, make sure to read the manual that came with the adapter from the adapter vendor and follow the instructions on how to connect the cables and install the board.

- Configuring on an ISA bus system:

1. If the system is an ISA bus system, the "isacfg" utility from the console must be used to configure the PCMCIA adapter.
2. After the PCMCIA adapter board is inserted to an ISA slot in the system, power-on the system.
3. To add an PCMCIA option to the platforms with an ISA bus, issue the following ISA option card configuration command at the console. The following example uses an AlphaStation 200 platform, but the commands should be the same in all 3 ISA bus platforms.

```
>>> isacfg -slot 1 -etyp 1 -dev 0 -mk -iobase0 3e0 -irq0
14 -enadev 1 -handle PCIC-PCMCIA
```

   NOTE that if the system is already using slot 1, use other slot number that is not used.
4. The IRQ (interrupt) number must also not conflict with those interrupt numbers that are already assigned to other default devices on the system. The system hardware manual usually indicates which default devices are on the system and which IRQ numbers those default devices use.
5. The recommended IRQ number to be used for the PCMCIA adapter is 14 (decimal).
6. If the number 14 is already used, number 10 should be set if it is not already used by other devices.
7. When you issue the above "isacfg" command, the console should print out the following line or something similar:

```
type >>>init to use these changes
```

8. After reinitializing the console, you can verify that you configured the

PCMCIA adapter correctly, by issuing the following command:

```
>>>isacfg -slot 1
```

9. You should see the following screen display:

```
==============================================================
handle: PCIC-PCMCIA
etyp: 1
slot: 1  dev: 0
enadev: 1
totdev: 1
iobase0: 3e0  membase0: 8000000000000000
iobase1: 8000000000000000  memlen0: 8000000000000000
iobase2: 8000000000000000  membase1: 8000000000000000
iobase3: 8000000000000000  memlen1: 8000000000000000
iobase4: 8000000000000000  membase2: 8000000000000000
iobase5: 8000000000000000  memlen2: 8000000000000000
rombase: 8000000000000000
romlen: 8000000000000000
dmamode0/chan0: 80000000 irq0: 14
dmamode1/chan1: 80000000 irq1: 80000000
dmamode2/chan2: 80000000 irq2: 80000000
dmamode3/chan3: 80000000 irq3: 80000000

==============================================================
>>>
```

- Configuring on an EISA bus system:
  1. If the system has an EISA bus, the ECU (EISA Configuration Utility) is used for configuring the PCMCIA adapter.
  2. If you are going to use PCMCIA adapter on EISA systems, run the EISA Configuration Utility (ECU), tell it that the PCMCIA adapter is there, and tell it to use the aisa3000.cfg config file.
  3. Refer to your system hardware documentation for complete instructions on how to run the ECU program.

1.14.3 Configuring and Using a PCMCIA Modem PC Card

Since a PC Card is a dynamic device (i.e. not a static device that is present all the time in the system hardware), and the serial-line device driver is a static device driver, when the system is installed initially, there will not be a corresponding `acex` entry created automatically by the `doconfig` of the target system. This is due to the fact that the system doesn't know when it is being installed that there will be a fax/modem card for PCMCIA since the card is not in the system yet.

If you want the system to automatically create the 'acex' entry for your PCMCIA fax/modem card, before you start installing the system, make sure that you have the PCMCIA adapter configured in the console and that the

PCMCIA fax/modem card is inserted into the slot. If you have a fax/modem card in the slot 0, for instance, when the system is installed and the target kernel is built, the system kernel configuration file built will have the following entry:

```
controller ace2   at pcmcia0  slot 0 vector aceintr
```

The installation will also create the device special file for this fax/modem card in the directory named `/dev`.

```
# ls -gl tty02

crw-rw-rw-   1 root     system    35,  2 Oct 16 13:22 tty02
```

If you didn't have the PCMCIA fax/modem card inserted in the slot when the system was installed, then you need to add the following line to your system kernel configuration file, `/sys/conf/HOSTNAME` where `HOSTNAME` is the name of your system. If you are just going to use one PCMCIA fax/modem card:

```
controller              ace2    at *    slot ? vector aceintr
```

If you plan to use 2 modem cards simultaneously, add the following 2 lines to your system configuration file:

```
controller              ace2    at *    slot ? vector aceintr
controller              ace3    at *    slot ? vector aceintr
```

Once the system configuration file is modified, use the following command to rebuild the new kernel and reboot the system.

```
# doconfig -c
```

1.14.4 Creating a Device Special File for the Modem Card

Normally the system installation will create the following two default `tty0x` device special files in the directory `/dev`, as the following command shows:

```
# ls -gl tty0*
```

This command produces output similar to the following:

```
crw-rw-rw-   1 root     system    35,  0 Oct 16 13:22 tty00
crw-rw-rw-   1 root     system    35,  1 Oct 16 13:22 tty01
```

This is because most systems have 2 imbedded serial lines. If the system has only one, only tty00 entry will be visible in the `/dev` directory.

Create additional device special files for the PCMCIA modem cards using the `MAKEDEV` utility in the `/dev` directory.

For example, in the /devdirectory:

```
# ./MAKEDEV ace2

MAKEDEV: special file(s) for ace2:
tty02
```

The generated special file should look like this:

```
crw-rw-rw-   1 root     system    35,  2 Oct 27 14:02 tty02
```

If you intend to have 2 PCMCIA modem cards working simultaneously, create device special files for the number of cards, in this case, 2. For example, in /dev directory:

```
# ./MAKEDEV ace2 ace3

MAKEDEV: special file(s) for ace2:
tty02
MAKEDEV: special file(s) for ace3:
tty03
```

The generated special file should look like this:

```
crw-rw-rw-   1 root     system    35,  2 Oct 27 14:02 tty02
crw-rw-rw-   1 root     system    35,  3 Oct 27 14:02 tty03
```

1.14.5 /etc/remote File

The /etc/remote file must be modified to add new access line definitions for the PCMCIA modem cards to be used. If you have a 28.8kpb modem card and will be using the full speed, the baud rate (br) in the /etc/remote file should be set to 38400.

For example, add the following line to /etc/remote file:

```
line2:dv=/dev/tty02:br#38400:pa=none:
```

Note that in the above line, "line2" can be any name you determine to be used with the "tip" command later on.

Once the PCMCIA modem card is inserted correctly and the system configures the card, the card can be used the same as any other modem devices.

1.14.6 Inserting a PCMCIA Modem Card

To use a PCMCIA modem card, insert the card to one of the PC Card slots in the PCMCIA adapter. Depending on the adapter type, there may be 2 front access card slots or one front access and one rear access card slot. When you insert the card into the slot 0, you should see the following message on the console terminal (or the Console Log window of the graphics head).

```
# PCMCIA socket 0: card manufacturer: MEGAHERTZ
product name: XJ2288
Configured: serial unit 2, type=16550A
ace2 at pcmcia0
```

This example used MegaHertz XJ2288 fax/modem card.

1.14.6.1 Possible Error Message

When a modem card is inserted, an error message such as the following example may appear on the Console Log window:

```
socket 0: card manufacturer: MEGAHERTZ, unknown modem card inserted

Using generic modem driver for this PC Card.

PCMCIA socket 0: card manufacturer: MEGAHERTZ

product name: XJ1144

socket 0: Couldn’t find usable config. for this card.
Please eject this PC Card.
```

The error message:

```
Couldn’t find usable config. for this card.
```

is generated if the card requires I/O resources that are already in use by other components in the system. If this error message is seen, the card should be ejected since it is not configured.

One possible way to correct this situation is to remove some other ISA/EISA devices in the system and reboot the system, thus freeing I/O resources that may be required by this card and trying the card insertion again.

1.14.7 Removing a PCMCIA Modem Card

Once you are finished using the modem card, push the button next to the card slot to eject the card previously inserted. You should see the following message on the console terminal or console Log window.

```
# stray interrupt on unit=2, intr_id=0
PCMCIA socket 0: PC Card removed
```

Note that you may or may not see the "stray interrupt..." message above when you eject the card. The message is generated by the serial-line driver if the PC Card generated an interrupt when the card got ejected.

1.15 Dynamic Device Recognition for SCSI Devices

Dynamic Device Recognition (DDR) is a framework for describing the operating parameters and characteristics of SCSI devices to the SCSI CAM I/O subsystem. You can use DDR to include new and changed SCSI devices into your environment without having to reboot the operating system. You do not disrupt user services and processes, as happens with static methods of device recognition.

Beginning with Digital UNIX Version 4.0, DDR is preferred over the current, static method for recognizing SCSI devices. The current, static method, as described in *System Administration*, is to edit SCSI device customizations into the `/sys/data/cam_data.c` data file, reconfigure the kernel, and shut down and reboot the operating system.

Note

> Support for the static method of recognizing SCSI devices will be retired in a future release of Digital UNIX .

Digital UNIX Version 4.0 supports both methods of recognizing SCSI devices. Both methods can be employed on the same system, with the restriction that the devices described by each method are exclusive to that method (nothing is doubly-defined).

The information DDR provides about SCSI devices is needed by SCSI drivers. You can supply this information using DDR when you add new SCSI devices to the system, or you can use the `/sys/data/cam_data.c` data file and static configuration methods. The information provided by DDR and the `cam_data.c` file have the same objectives. When compared to the static method of providing SCSI device information, DDR minimizes the amount of information that is supplied by the device driver or subsystem to the operating system and maximizes the amount of information that is supplied by the device itself or by defaults specified in the DDR databases.

You can also use DDR capabilities to convert customizations in the `cam_data.c` file to information in the DDR `/etc/ddr.dbase` text database.

For more information about DDR, see *System Administration*, `ddr_config`(8), and `ddr.dbase`(4).

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 space requirements

Note

Do not attempt to install Digital UNIX Version 4.0 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 process in general.

2.1.1 Kernel Object Files Not Installed

In this release, kernel object files are not needed for installation or kernel building. Only kernel modules are needed. See Section 1.12 for more information on these kernel modules.

The individual kernel object files are only needed for kernel development work. See Section 2.4.7 for instructions on when and how to load these objects if they are needed.

2.1.2 Installation Procedure may Fail to Detect Disks

The installation procedure automatically detects the disk(s) on a system, and asks you to specify the system (root) disk. If the disk label indicates that the `a` partition starts at a non-zero offset, the disk may not be detected by the installation procedure. This problem most commonly occurs with disks which have previously been used by LSM. LSM modifies the a partition to start at offset 16.

If all disks on the system cannot be detected because of this problem, the installation procedure will exit to a shell prompt #. If one or more disks are detected, while others are not, the undetected disks will not be presented to the user as devices where Digital UNIX can be installed.

You can correct this problem as follows. If the installation procedure does not automatically exit to the shell prompt, exit manually. In the graphical interface, click the UNIX Shell button, and click "OK" in the dialog box that appears. In the character-cell interface, press CTRL/C.

When you see the shell prompt, enter these commands:

```
# cd /dev
# ./MAKEDEV <device>
```

where <device> is replaced by the actual device name such as rz3. If more than one disk is not detected by the installation procedure, enter a MAKEDEV command for each undetected disk.

Enter the following commands to restart the installation procedure:

```
# cd /
# restart
```

2.1.3 Systems Without Graphics Devices

This note applies to all hardware configurations that do not have an attached graphics device. If the OSFXC2SEC350 subset is installed on a system that does not have a graphics controller installed on it, the following errors may be seen:

```
setld:
Error installing "Graphical System Administration Utilities"
(OSFXADMIN400)
This subset requires following subset(s) to operate correctly:

"Graphical Base System Management Utilities" (OSFXSYSMAN400)
setld:
Please install required subset(s) first.
```

These errors are caused by subset dependency ordering and if the errors occur, the following subsets will not be properly installed onto your system.

```
OSFXC2SEC400   OSFXSYSMAN400
OSFCDEMIN400   OSFXADMIN400
OSFTKBASE400
```

2.1.4 Worldwide Subsets

The following error message may be displayed during the installation of the Worldwide subsets:

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

The installation procedure will abort. This error may happen if the shell has too many environment variables and runs out of memory during installation.

This problem can be corrected by starting a new shell and removing the unused environment variables.

2.1.5 Invalid Password in Graphical Installation

After changing the Install Type field (default to custom, or custom to default) and entering an invalid password, an error dialog box will be displayed but may not be in focus. An invalid password is one which is less than six characters long or contains only lowercase characters. Pressing Return to dismiss the dialog box when it is not in focus may cause you to be temporarily locked out of all editable fields. To avoid this situation after entering an invalid password, ensure the error dialog is in focus and either press Return or click on OK with the mouse to dismiss it. This will correctly clear the error dialog box and allow re-entry of the password.

If you encounter this temporary lock-out situation, you can work around the problem by setting focus to a non-editable field in the Installation Setup window (i.e., clicking on a menu or button) and then setting focus back to the desired entry field. You can also set focus outside the Installation Setup window and back to it.

2.1.6 Minimum Version 2.10 for Qlogic Firmware

On booting, the Qlogic driver attempts to set the Data Overrun Recovery mode to that supported by Digital UNIX. If the operation fails, the driver will print a warning message on the system console:

```
isp<n>: WARNING - Data Overrun Recover Mode could NOT be set!
```

Where `<n>` is the number of the Qlogic controller encountered during system probe.

Digital recommends that the Qlogic ISP1020 or ISP1020A firmware be at a minimum revision level of 2.10. This is the minimum revision of Qlogic firmware required to implement reliable data overrun error detection and recovery.

2.1.7 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. Consult the hardware documentation for more information.

2.1.8 Incorrect Broken Pipe Messages

The installation process displays the following message during the installation of some subsets:

```
Broken Pipe
```

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

2.2 Layered Product Considerations

The versions of the layered products shown in Table 2-1 are compatible with, or were specifically developed to take advantage of new Digital UNIX features. They may also provide new hardware support or resolutions to problems. The table shows the correct layered product versions to use with this release.

Table 2-1: Layered Product Versions

Layered Product	Digital UNIX Version 3.2	Digital UNIX Version 3.2C	Digital UNIX Version 4.0
POLYCENTER AdvFS Utilities	Version 3.2	Version 3.2	Version 4.0
DECsafe Available Server	Version 1.2	Version 1.2A	N/A
Logical Storage Manager	Version 1.2	Version 1.2	Version 4.0
System V Environment	Version 3.2	Version 3.2	Version 4.0
DECnet	Version 3.0	Version 3.0	Version 4.0
SNA Peer	None	Version 1.2	Version 1.2 with ECO 1

2.2.1 DEC FUSE

DEC FUSE Version 2.1A and prior versions are not supported under Digital UNIX version 4.0. DEC FUSE Version 3.0 is supported under Digital UNIX version 4.0.

2.2.2 DECnet/OSI

Digital UNIX Digital UNIX Version 4.0 does not currently support DECnet/OSI. Any existing DECnet/OSI releases will not work with this release of Digital UNIX.

2.2.3 Subset Control Program Methods for Restricting a Product to Digital UNIX

Hardware releases based on Digital UNIX Version 4.0 may not install the `OSFBASE400.lk` file; however, they will install a `OSFBASE4??.lk` lock file.

If you want to tie your product with Digital UNIX Version 4.0 you should use wildcards to check for the `OSFBASE4??.lk` file instead of `OSFBASE400.lk` file.

If you hard code the test for the `OSFBASE400.lk` into your subset control programs (SCP) it will fail to find the lock file when run on a system with a Version 4.0-based hardware release installed.

This policy has always been in place; however, until now, not following it had not caused any problems.

If you wish to restrict your product from installing on Version 4.0-based hardware releases, you can hardcode the `OSFBASE400` string in your SCP.

2.3 Full Installation

A installation (default or custom) will overwrite any customization on your system.

2.3.1 Full Installation Overwrites /etc/fdmns

If your system is installed with the AdvFS file system, save a copy of the `/etc/fdmns` directory before doing a full installation. The installation process causes this directory to be overwritten, resulting in the potential loss of configuration data contained in the directory.

2.3.2 Full Installation 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.

2.3.3 Full Installation 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.

2.3.4 Full Installation Overwrites /etc/prestotab

If your system has Prestoserve configured on it, save a copy of the /etc/prestotab file 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.

2.3.5 The Installation Graphical Interface

See Section 4.4.1 for information on restrictions when using the Disk Configuration Manager.

2.4 Update Installation

The following notes apply to the update installation procedure.

2.4.1 Systems Without Graphics Devices

This note applies to all hardware configurations that do not have an attached graphics device. If the OSFXC2SEC350 subset is installed on a system that does not have a graphics controller installed on it, the following errors may be seen:

```
setld:
Error installing "Graphical System Administration Utilities"
(OSFXADMIN400)
This subset requires following subset(s) to operate correctly:

"Graphical Base System Management Utilities" (OSFXSYSMAN400)
setld:
Please install required subset(s) first.
```

These errors are caused by subset dependency ordering and if the errors occur, the following subsets will not be properly installed onto your system.

```
OSFXC2SEC400    OSFXSYSMAN400
OSFCDEMIN400    OSFXADMIN400
OSFTKBASE400
```

2.4.2 Disable Lightweight Wiring When Using Granularity Hints

When lightweight wiring is used with granularity hints, the system can experience the following panic message:

```
panic: lw_remove: light weight wiring(s) found
```

This panic can occur on any hardware configuration, but will generally only occur on large systems using large amounts of shared memory.

Before updating your system, remove the line in `/etc/sysconfigtab` containing `gh-chunks`. Lightweight wiring is enabled by default in Digital UNIX version 4.0. Removing `gh-chunks` from the `/etc/sysconfigtab` file will ensure that granularity hints is disabled after the update installation, and that both granularity hints and lightweight wiring are not enabled simutaneously.

If granularity hints is required, lightweight wiring must be disabled by setting `new-wire-method=0` in `/etc/sysconfigtab`. The system must then be rebooted for the change to take effect.

2.4.3 Disk Space Requirements for Update Installations

Before beginning your update installation from Version 3.2C, Version 3.2D-1, or Version 3.2D-2, you should refer to Appendix F.2 in the Digital UNIX Version 4.0 *Installation Guide* (supplied with your Version 4.0 software distribution kit). Ensure that your system meets the minimum additional disk capacity requirements to successfully update to Version 4.0. If it does not, you may need to remove some optional software subsets in addition to unnecessary core dump files, kernels, or other unrequired files, prior to starting the update procedure.

If you have already begun an update installation that exited due to insufficient additional disk space, the installation procedure will list which file systems require additional space, and how much space is required.

To identify base and layered product subsets that are currently installed, execute the following command as root:

```
# setld -i | grep installed
```

This command will list all potential candidates for deletion.

The left column lists the subset identifier, and the right column gives a brief text description of the subset.

Next, refer to Appendix E of the Version 3.2C Installation Guide. This appendix lists the disk space required per subset for the root, /usr, and /var file system. If you have Version 3.2D-1 or Version 3.2D-2 installed, you will need to add the additional amount of space required per file system that these releases add to Version 3.2C (this information is contained in the Version 3.2D-1 and Version 3.2D-2 release notes). Next, determine which subsets can be removed to satisfy the file system space requirements of the update installation.

For example, a Version 3.2C to Version 4.0 update installation exits, requiring an additional 1 MB of space in the root file system. The optional subsets `OSFATMBASE350`, `OSFCOMAGENT350`, and `OSFC2SEC350` are installed but are not being used. Referring to Appendix E in the Version 3.2C Installation Guide, the total of the root file system component (in 512-byte blocks) for these subsets is:

```
OSFATMBASE350            431.76
OSFC2SEC350             1513.77
OSFCOMAGENT350           182.79
--------------           ------
Total                    2128.32 (512-byte blocks)
Total                       1.04 (MB)
```

Therefore, removing these three optional subsets should reclaim enough space for the update procedure to complete:

```
# setld -d OSFATMBASE350 OSFC2SEC350 OSFCOMAGENT350
```

2.4.4 Update Install and Enhanced Security

The update installation process does not convert the extended user profiles and ttys information to the new database format. This can be done manually after the update installation completes with the following command:

```
# /tcb/bin/convauth
```

2.4.5 System Clock may Lose Time

During an update installation, the system clock (TOY, time of year) may lose a short amount of time possibly up to 3 minutes. This may happen when you update your system from any version preceding Version 4.0, to Version 4.0 or later.

During the boot process, a message will be printed to the console indicating the condition as follows:

```
over 5 years since last boot - CHECK AND RESET THE DATE
Setting time to last shutdown time stamp
```

the alternative message will be:

```
WARNING: preposterous time in TOY -- CHECK AND RESET THE DATE
```

The time will be set to the previous shutdown timestamp.

On some systems, similar behavior will occur everytime the systems transition between a version preceding Version 4.0 and version Version 4.0 or later. For example, consider a system which has one disk with Digital UNIX Version 3.2C installed and another disk with Digital UNIX Version 4.0 installed:

1. The time is set while either version is booted.
2. When the system is shut down and booted from the other disk, time will not be maintained by the clock. The time will be set to the previous shutdown time-stamp.

The problem will occur on the following systems:

```
DEC_3000 series   DEC_8000
DEC_4000          DEC_10000
DEC_7000          AXPVME64
```

Note that systems running `ntp` should not be affected by this problem, with the exception that a warning message may be printed to the console as the system boots. Open 3D is not supported for update installations in this release. Use the `setld` command to remove all Open 3D subsets before updating your system.

2.4.6 Message From swapon Utility

During an update installation, the following message from the system `swapon` utility may appear at the start of some update installations:

```
swapon: /dev/rz10b: already a swap device
```

The swap device specified in the message may be different than the example. You can ignore this message.

2.4.7 Kernel Object Files Are Deleted

After the Digital UNIX Version 4.0 update installation is complete you can remove kernel object files from your system to reclaim disk space. Kernel object files are only necessary if you are going to perform kernel development work on your system. If your system is not being used for kernel development work and you want to reclaim approximately 25 meg of disk space, use the following command to determine if any of the following subsets are loaded on your system:

```
# setld -i

OSFBINOBJECT400
OSFHWBINOBJECT400
OSFADVFSBINOBJECT400
OSFATMBINOBJECT400
OSFLSMBINOBJECT400
```

Use the following command to delete these subsets to reclaim additional disk space:

```
# setld -d
```

2.4.8 The installupdate Command May Select the Wrong Device

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

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

2.5 Server Extensions

This section provides information on installing Digital UNIX Server Extensions.

2.5.1 Error when installing Chinese subsets from RIS server

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

```
Common Chinese Unicode Support
Copying from hanax1 (inet)
setld: cannot access server mapping (rcp: ris3.alpha/rp_mapping:
No such file or directory)
setld: Load from hanax1 failed, subset IOSZHUCSBASE400
Verifying
setld:
There were verification errors for "Common Chinese Unicode Support
```

```
(IOSZHUCSBASE400)
```

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

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

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

```
Choose one of the following options:

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

Enter your choice:
```

2.6 Disk Space Requirements

The subset size tables have been moved to a permanent location in the *Installation Guide* for this release. If you want to add optional subsets after you install Digital UNIX Version 4.0, use the `df` command to determine free disk space in blocks.

Processor-Specific Notes 3

This chapter contains notes that apply to processors supported by Digital UNIX Version 4.0:

- General notes on processors
- Digital AlphaStation 200 and Digital AlphaStation 400 Series processors
- Digital AlphaServer 1000 processors
- Digital AlphaServer 2000, AlphaServer 2100 processors and Digital Alpha VME 2100 processors
- DEC 7000 processors
- Digital AXPvme single board computers
- Digital AXPpci 33 processors

3.1 General Notes on Processors

The notes in this section apply to more than one processor type.

3.1.1 PCI-based Systems with ATI Mach64 Graphics Controllers

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

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

3.1.2 PCI Shared Interrupt Support

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

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

```
>>> show config
```

In order to setup your PCI-based system to share interrupts, you must ensure the following:

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

Refer to your Digital UNIX Layered Product documentation (for example, for MME or Open3D) for information as to the shared interrupt capabilities of the software.

3.1.3 Support for RTS/CTS (Hardware) Flow Control

Digital UNIX Version 4.0, supports RTS/CTS (hardware) flow control in the serial drivers that are used on the DEC3000 series, DEC2000 series, AlphaStation, and AlphaServer systems. This support can be enabled or disabled using the `-crtscts` argument to the `stty` command. It can also be enabled on logins by specifying the `Fflag` option `CRTSCTS` on the appropriate `/etc/gettydefs` entry. By default, RTS/CTS flow control is disabled.

3.2 Digital AlphaStation 200 and Digital AlphaStation 400 Processors

The following notes apply to the Digital AlphaStation 200 and Digital AlphaStation 400 series processors.

3.2.1 Power Management Support for the Digital AlphaStation 255

This release supports power management for the Digital AlphaStation 255 system. The power management capability allows you to control:

- Disk spindown
- CPU slowdown
- Monitor powerdown (DPMS).

The user may control power management settings using a graphical user interface `/usr/bin/X11/dxpower` A command line interface `sysconfigdb` is also available. Specific monitor settings can also be modified using the `xset` command.

See the *System Administration* manual and the `xset`(1X) and `sysconfigdb`(8) reference pages for more information.

Note that a system (root) disk cannot spin down due to normal system activity. User disks will spin down unless applications are running and the applications require constant access to the user disks.

3.2.2 Monitor Power Management for the Digital AlphaStation 255

Monitors that do not support DPMS (Display Power Management Signaling) can be damaged by the activation of the DPMS feature. Check your monitor specifications.

Monitors that support DPMS and are put in a power savings state will vary in the time it takes to come out of power savings. Users will observe that the longer the monitor is in power-off state, the longer it takes for the display to return as a result of mouse or keyboard activity. This is the result of the monitor phosphor cooling down and the time required to heat it back up, and not a function of the operating system.

3.3 Digital AlphaServer 1000 4/200 Processors

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

3.3.1 EISA Configuration Utilty and RIS Installation

When upgrading a AlphaServer 1000 series processor from the firmware CD–ROM, you should boot the firmware file `AS1000_V5_1.EXE`, (version 5.1) and not `AS1000_V5_2.EXE` (version 5.2) as indicated in the *Alpha Systems Firmware Update FT1 Release Notes Overview*

If you use version 5.1 instead of version 5.2 you will not be able to perform an installation using RIS.

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

The notes in this section apply to the AlphaServer 2000, AlphaServer 2100 Server and Digital Alpha VME 2100 series processors.

3.4.1 SWXCR Raid Configuration Updates

Do not perform SWXCR RAID configuration updates with ARC Version 4.2-1 on the AlphaServer 2000 or AlphaServer 2100.

3.4.2 Digital Alpha VME 2100

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

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

3.4.2.1 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 3.6.2.

3.5 KDM70 Disk Controller on DEC 7000 Processors

The KDM70 disk controller is not supported when more than one gigabyte of memory is installed. After running the ECU over a serial line, the cursor will not display. You must change the terminal setup to redisplay the cursor.

3.6 Digital AXPvme Single Board Computers

The following notes apply to the Digital UNIX software running on the following single board computers (SBCs):

- Digital AXPvme 64 Single Board Computer
- Digital AXPvme 100 Single Board Computer
- Digital AXPvme 160 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 4.0 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. 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 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, 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. The following notes apply to the Digital AXPpci 33 series processors.

3.6.3 IRQs for PCI Options on Digital AXPpci 33 Systems

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.

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:

- Internationalization
- Commands and utilities
- SysMan System Management graphical user interface
- System administration
- Network and communications
- File systems

4.1 Internationalization

The following notes apply to restrictions on using functions that support internationalization or internationalized components.

4.1.1 Return Values of Functions iswdigit() and iswalnum() Under the Thai Locale

The return values of the functions `iswdigit()` and `iswalnum()` for Thai digits in the range 0xF0,...,0xF9 are false, even though these values are defined as digits in the *Wototo Standard, Version 2.0* relating to the Thai language.

4.1.2 Printing PostScript Format Files in Asian Languages

If two or more print jobs are sent to different queues of the same printer within a very short time, it is possible that some jobs will get blocked and cannot be printed. Restart the job using the `lpc` command.

4.1.3 Line Wrapping in vi

Depending on the composition of the line, the `vi` editor that supports Thai may wrap lines before the right boundary of the screen. For a normal 24x80 screen, a line wraps if more than 80 Thai or ASCII characters are entered, even when the display width of the line is fewer than 80 columns.

4.2 X/Open-compliant Curses

Digital UNIX version 4.0 contains a new curses implementation that incorporates the following sets of programming interfaces:

- X/Open Curses, Issue 4
- System V Multinational Language Supplement (MNLS)
- Minicurses
- BSD curses

These interfaces support the new curses standard and also maintain source compatibility with the curses interfaces found in the previous version of Digital UNIX. However, they do not maintain binary compatibility. Therefore, additional library files are also provided to maintain binary compatibility with existing programs

The new curses implementation can be found in these directories:

Directory	Contents
`/usr/include`	`curses.h, curshdr.h, term.h, unctrl.h`
`/usr/lib`	`libcurses.a`
`/usr/shlib`	`libcurses.so`

Libraries supporting the previous curses implementation will be found in the following directories:

Directory	Contents
`/usr/opt/lib`	`libcurses.a`
`/usr/shlib/osf.1`	`libcurses.so`

There are no header files provided for the old implementation. All programs that call curses directly will use the new header files when compiling and linking with the new libraries.

Existing binary files that were built with the old version of the curses shared library, `libcurses.so`, will continue to run without modification.

Through the technique called `versioning` old binary files will select the correct shared library.

When rebuilding applications that do not call curses directly, but which link with libraries that interface with the old curses implementation, set `LD_LIBRARY_PATH` to the proper directory as indicated in the preceding table, or else specify the complete pathname on the compile command.

4.2.1 Changes in Datatypes, Structures, and Macros

Many functions declared in `curses.h` and `term.h` are implemented both as routines and macros. In these cases, the default is to compile them as macros which execute more quickly. If you wish to execute the corresponding routines instead of using the macros, add the `_NOMACROS` compile-time switch.

The new curses library provides source compatibility with the old version with respect to function calls, data names, and macro names. However, the implementation of types, variables, structures and macros has changed considerably in some cases. If the existing user code takes advantage of the assumed implementation details of such elements, that program may need to be changed to work with the new curses. Elements whose implementation have changed include:

- datatype:
  - `chtype`
- structures:
  - `_win_st`
  - `WINDOW`
  - `SCREEN`
  - `term`
- macros:
  - `A_*`
  - `ACS_*`
  - `KEY_*`

In previous implementations, the ACS definitions in `curses.h` were statically defined constant values. In the new implementation, the definitions are dynamically defined at run time based on the current setting of the `TERM` environment variable. This implementation allows terminals with extra ACS capabilities to make them available to the user while providing a set of default ACS definitions for terminals with lesser capabilities.

The implementation changes may cause compile time failures for some programs that depend on the static definitions. For example, the following declaration will not compile when it occurs at the global level:

```
char A = ACS_ULCORNER;
```

The new implementation of the curses library defines the ACS definitions at run time and requires that all assignments be made after the `initscr()` function has been called.

4.2.2 Change in Behavior of cbreak() Function

The X/Open Curses standard requires that the `cbreak()` function disable the `ICRNL` input processing flag. In the previous Digital UNIX implementation, `cbreak()` did not disable this flag. In applications that relied on this default behavior to advance to new lines, subsequent output lines may now overwrite the last line addressed. Those applications should now set the `ICRNL` flag explicitly after the call to `cbreak()`. Here is a sample code fragment that sets the `ICRNL` flag:

```
#include <termios.h>

struct termios tty;

tcgetattr(0, &tty);
tty.c_iflag |= ICRNL;
tcsetattr(0, TCSANOW, &tty);
```

4.3 Commands and Utilities

The following notes apply to commands and utilities.

4.3.1 vdump and vrestore Compatability Problem

Note

This is an important note for users of `vdump` and `vrestore`

Backups made using `vdump` on Digital UNIX Version 4.0 cannot be restored using `vrestore` on earlier versions of Digital UNIX. Patches will be made available for earlier versions of `vrestore` to correct this problem.

Backups made using `vdump` on earlier versions of Digital UNIX can be restored using `vrestore` under Digital UNIX Version 4.0 without problems.

4.3.2 Security

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

4.3.2.1 Unexpected Command Behaviour with ACLs

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

- id mapping
- mandatory access control
- additional authentication requirements

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

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

For programs which copy files, the command:

```
# cp -p
```

will copy a file preserving ACLs and any other extended attribute (property list).

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

4.3.2.2 Archive Tools and Security

This note is to clarify the interactions between the archive tools `pax`, `tar`, and `cpio`, and files containing property lists or Access Control Lists (ACLs). or ACLs alone.

When you extract files with the above utilities without using the `-p` option, and the following conditions apply:

- you are extracting a tar or cpio archive into a directory that has default ACLs, or
- you are extracting a tar or cpio archive that contains directories with default ACLs, and

- not all of the files on the archive have associated ACLs.

Then, when files that do not have an associated ACL them are extracted, they will inherit the default ACL that is in force on the directory in which they are being created. This behavior was selected to allow file extractions to work as expected in as many cases as possible.

For times when the above behavior is not appropriate, an alternative behavior has been associated with the `-p` option. If you use the `-p` option, any file that is extracted from the archive will be given only the property list, ACL, and file permissions that were stored in the archive with that file. For `tar` archives you can use either the `-p` option to the `tar` command or the `-p p` option to the `pax` command. For `cpio` archives you can use the `-p p` option to the cpio command

See the `tar`(1), `pax`(1), `cpio`(1), and `tar`(4) reference pages.

4.3.2.3 emacs Can Lose ACL File Settings

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

The `emacs` utility has some user-preference variables which can be set to control which file will retain the original ACL. The relevent EMACS variables are:

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

4.3.3 Mail and mailx Command Behavior

Since the `mailx` and `Mail` command have become XPG4 compliant, their command behavior has changed. When using a carriage return with no arguments in command mode, you no longer see the next message. A carriage return with no arguments now behaves like the `print` command and not the `next` command. The current message is displayed and the message pointer stays on the current message.

4.3.4 Using gendisk on Diskette Devices

The `gendisk` utility is used to create product media. There is a problem in using it on FDI diskette devices, the diskette drive found on all non-Turbochannel bus Alpha platforms.

The solution involves making some hard links to the diskette device special files using the name of the device that `gendisk` will use.

1. Make the hard links as follows:

```
# cd /dev
# ln rfd0c rfl0c
# ln rfd0a rfl0a
# ln fd0a fl0a
# ln fd0c fl0c
```

2. Format the diskette disk as follows:

```
# fddisk -fmt /dev/rfd0c
```

   You will see the following messages:

```
NOTE:  Setting interleave factor to ‘‘-i2:4’’.
Use ‘‘-i<nnn>[:<ccc>]’’ option to override.
Disk type: 3.50 inch, HD  (1.44MB)
Number of sectors per track: 18
Number of surfaces:   2
Number of cylinders: 80
Sector size:  512
interleave factor:  2:4
Formatting disk...
Percentage complete: Format complete, checking...
Quick check of disk passes OK.
```

3. Use `disklabel` to label the diskette:

```
# disklabel -wr fd0 rx23
```

4. Run `gendisk` as follows:

   Note

   When running the `gendisk` utility on the diskette using these instructions, do not respond `yes` to the question asking to clean the disk.

   The following is an example of a `gendisk` command session:

```
# gendisk -d MYPRODUCT400 /dev/rfd0c
```

```
Generating MYPRODUCT400 Kit from <system address> on /dev/fl0c

WARNING: this will remove any information stored in /dev/fl0c.
Are you sure you want to do this? (y/n): y

Do you want to clean the entire disk first? Note: This will replace
your current disk label with a default one. (y/n) [n]: n

Preparing /dev/fl0c (floppy)
done.

Checking /dev/fl0c
/sbin/ufs_fsck /dev/rfl0c
** /dev/rfl0c
File system unmounted cleanly - no fsck needed

Mounting /dev/fl0c on /usr/tmp/cd_mnt8344

Writing Images (dd=/).

Image instctrl...done.
Image SVGASTATIC100...done.

Verifying Images (dd=/).

Image instctrl...done.
Image SVGASTATIC100...done.

Kit MYPRODUCT400 done.

Cleaning up working directories.
Unmounting /dev/fl0c
```

4.3.5 Some emacs Command Line Options Fail

Digital ships the emacs software as it is received from the source. The following command line options do not work as documented in the emacs reference page:

```
-geometry   -iconic
-ib         -T
-iconname   -in
-mc         -i
```

In some cases a solution is available using an appropriate X resource.

4.3.6 Directory Tree /usr/opt/sterling Renamed

The `/usr/opt/sterling` directory tree has been renamed to `/usr/opt/obsolete`.

All the files that were in the `/usr/opt/sterling` directory have been moved to the `/usr/opt/obsolete` directory tree. The following files have been moved within the /usr/opt file system.

Old location of commands	New location of commands
~sterling/sbin/cpio	~obsolete/sbin/cpio
~sterling/sbin/tar	~obsolete/sbin/tar
~sterling/usr/bin/cpio	~obsolete/usr/bin/cpio
~sterling/usr/bin/tar	~obsolete/usr/bin/tar

If you have scripts that rely on the old location, a symbolic link can be placed in /usr/opt to point to the new location as follows:

```
# cd /usr/opt
# ln -s sterling obsolete
```

4.3.7 POSIX Shell

To run the POSIX shell, the environment variable BIN_SH, must be set to xpg4. The POSIX shell is invoked when the user runs the command sh.

The POSIX shell is located in /usr/bin/posix/sh. If BIN_SH is not set to xpg4, the Bourne shell is invoked when the user runs the sh command. Relative or absolute paths are not determining factors; executing /usr/bin/sh gives the same result as sh. The determining factor is the environment variable BIN_SH.

4.3.8 Sendmail Configurable Maximum Hop Count

In sendmail, the maximum hop count is now configurable. If not specified, the hop count defaults to 17. Each time a message is forwarded through a host, the hop count is increased. When this count exceeds the maximum hop count value, the message is rejected, because it is automatically assumed that an endless loop has occurred.

The default value is acceptable in most installations but you may want to increase the value if too many messages are being lost.

4.3.9 cksum Reports Incorrect Values

The current values reported by the cksum command are incorrect according to the IEEE Std 100 3.2-1992. To conform with XPG4 requirements, new calculations have been made for the checksums.

For current compatibility, the default action of the cksum command is to report the present cksum values. To obtain the new checksums, set the environment variable CMD_ENV to the string xpg4. For example:

```
export CMD_ENV=xpg4
```

4.3.10 df and XPG4 Compliance

The default behavior for the df command is BSD SVR4 compliant. If XPG4 compliant behavior is desired, set the CMD_ENV environment variable to xpg4. The XPG4 compliant df command takes the following syntax:

```
df [-eiknPt] [-F fstype] [file | file_system ...]
```

See the df (1) reference page for more information.

4.3.11 Use of echo -n in Scripts

The default of the /usr/bin/echo command is compliant with the XPG4 standard. If the CMD_ENV environment variable is not set or is set to xpg4, the echo command will treat the option -n as a string. The echo command supports the -n option when the environment variable CMD_ENV is set to bsd.

4.3.12 awk Linked to nawk

In a future version of Digital UNIX, the link between awk and nawk will be removed, leaving an XPG4-compliant version of awk. You should ensure that your scripts use /usr/bin/awk in place of any other version of the command currently existing on the system.

The gawk command invokes GNUawk, the Free Software Foundation version of awk. This command has been moved out of the base subset to the Free Software Foundation subset. The oawk command has been removed.

4.3.13 Unsupported vmh Command

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

4.3.14 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.

4.4 SysMan System Management Graphical User Interface

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

4.4.1 Restrictions on Using the Disk Configuration Manager

The disk configuration manager is a new application that will allow the inspection and modification of disk attributes, partition information in particular. Type `/usr/sbin/diskconfig` at the command line to invoke the application and display the top level window. This note describes some restrictions on using the disk configuration manager in Digital UNIX, Version 4.0.

The `diskconfig` utility displays a list of disks attached to the system. This occurs implicitly upon invocation. The selection of a disk from this list presents the attributes of the disk as recorded on the disk label if a label is present and the drive is not defective. If no label is present and the drive is not defective, default information is presented. An error message is presented if the drive is defective. Much of the data displayed in the windows can be edited by the user to change the characteristics of a disk label. Known problems are as follows:

- The `diskconfig` utility will ignore an attached disk that has no associated special device files. Such a disk will not appear in the top level list.
- The `diskconfig` utility will create an AdvFS file system on a specified partition, only in the most primitive way. That is to say that the `diskconfig` utility does not take advantage of the powerful set of AdvFS file management capabilities. When you choose to create an AdvFS file system, a warning dialog will appear to state same.
- If you are creating AdvFS file systems on multiple partitions, you should commit them one at a time. That is, select AdvFS for the first partition, enter the domain and set names for that partition, and press commit. Then repeat the procedure for subsequent partitions. Failure to do so is not dangerous, but may be confusing as the `diskconfig` utility will not redisplay the domain and set names once another partition uses these areas of the display.
- The `diskconfig` utility does not currently support VGA 640x480 graphics. Consequently, the Partition Disks option on the installation GUI cannot be used on a system with VGA 640x480 graphics. The solution is to choose a UNIX shell from the installation GUI and use the shell to accomplish disk partitioning with disklabel. Continue with the installation when the disk is repartitioned.

- Disks moved out of a raid array and still bearing the label written while in the raid array, are not repairable by the `diskconfig` utility, only by setting the label to default values. Press `default` and `commit`. Then make your desired changes. Alternatively, you can use the `disklabel` command to change disk labels.
- The `none` option when selecting boot blocks means *write no boot blocks.* Boot blocks will not be removed; rather, they will remain as they were.

4.4.2 Account Manager

The following notes apply to the account manager, `dxaccounts`.

4.4.2.1 Usage Note

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

4.4.2.2 Account Manager Restrictions

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

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

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

- The account manager allows you to specifiy the minimum and maximim range for a UID or GID. However you cannot specify the starting value in the range.

  Workaround: Set a starting value within the range using the `usermod` or `groupmod` commands:

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

  For example, if the Minimum UID is 100 and the Maximum UID is

10000 then:

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

causes account manager to start generating UIDs from 5000.

- If you change the UID of a user, the ownership of user's home directory and files will not be changed.

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

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

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

- Restriction: Two system adminstrators should not run two different concurrent instances of account manager.

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

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

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

  Workaround: None. These messages can be safely ignored.

4.4.2.3 Account Manager Problems

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

4.4.2.4 Enhanced Security Account Manager Problems

The following problems apply to the account manager when running on enhanced security systems.

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

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

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

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

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

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

- Account manager does not enforce the Minimum/Maximum password length limitation when setting passwords.

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

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

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

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

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

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

  2. Remove the user from the protected password database using the

following command:

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

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

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

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

- Do not rename a template by changing the Template name field of the Create/Modify Template dialog box in Modify mode. Account manager will actually create a new template without removing the old template. However, the old template's icon will be removed from the icon box.

  Workaround: Restart the account manager to restore the former template icon. Delete the undesired template using the Delete Toolbar icon or the Edit->Delete... option from the Template view.

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

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

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

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

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

- Dragging and dropping a template on a user displayes and confirmation message in the view’s status line. This message incorrectly displays template %2 instead of the template’s name.

  Workaround: None.

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

  Workaround: Please restart account manager and then delete the template.

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

  Workaround: Manually remake the NIS maps or perform an account manager function (eg. Account Modification) that will trigger the maps to be remade. To manually remake the maps do the following:

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

4.4.3 Enhanced Security and Account Manager

- Only the root user can change a password using the Change Password graphical user interface, `/usr/tcb/bin/dxchpwd`. All non-root users must use the Change Password command line interface, `/bin/passwd` to change their passwords.

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

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

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

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

  Workaround: Set the C1Crypt Encryption type for the user from the

Create/Modify User dialog.

- The Account Manager does not enforce the Minimum/Maximum password length limitation when setting passwords.

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

- The Pointer Focus Help message in the Status Line of the Icon Box will display Delete instead of Retire when the mouse pointer is in the Retire toolbar icon.

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

- There is no mechanism to remove a user from the BSD and Protected databases.

  Workaround: The user must be removed from the BSD databases by manually editing the `/etc/passwd` and `/etc/group` files. The user can be removed from the Protected database using the command:

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

- Renaming a user by changing the User name field of the Create/Modify User dialog box in Modify mode does not clean up the Protected Password entry of the old name.

  Workaround: To remove a dangling protected entry, use the command:

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

- Do not rename a template by changing the Template name field of the Create/Modify Template dialog box in Modify mode. The account manager will actually create a new template without removing the old template. However, the old template's icon will be removed from the icon box.

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

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

- The Drag/Drop feature to copy a user copies template references by value when the user is being dragged to a different view (eg. Local to NIS). This means that the template is no longer referenced in the newly created account but the template values are contained int the new users profile. The template reference is maintained if the user is dragged within the same view.
- The Copy/Paste feature to copy a user always copies template references by value. This means that the template is no longer referenced in the newly created user account but the template values are contained in the new users profile.
- Template Assignment through the Drag/Drop feature does not update the user's lock field.

  Workaround: Manually update the Lock field by selecting the Lock toggle button or assign templates to users through the Template pull down list of the Create/Modify User dialog.
- Template Assignment through the Drag/Drop feature of a newly created template assigns the template by value. This means the user does not reference the template but the template values are contained in the users profile.

  Workaround: Manually assign the template to the user through the Template pull down list of the Create/Modify User dialog.
- Template Assignment through the Drag/Drop feature of a newly created template displays a confirmation message in the status area which references template %2 instead of the template name.
- Deleting a newly created template (the template was created after starting the Account Manager) will result in a core dump.

  Workaround: Restart the account manager and then delete the template.
- The NIS Maps are not remade after template deletion.

  Workaround: Manually remake the NIS maps or perform an Account Manager function (eg. Account Modification) that will trigger the maps to be remade.
- Leading and trailing white space is not stripped from text entry areas. This could lead to confusion when, for example, a field on the Find dialog contains a space before the entered text. The Find dialog includes the white space when attempting to determine the view members that match the given field.

4.4.4 Swap Warning in dxsysinfo

When using System Information, `dxsysinfo`, the swap warning Light will not illuminate if the available swap space falls below 10 percent free, unless the available swap meter is being displayed. Both of these options can be activated by selecting them from the View menu.

4.4.5 File System Percent Full Values May be Incorrect

The value of Percent Full in the file system area of `dxsysinfo` may be inaccurate.

For the correct value of Percent Full, use the `df` command and refer to the Capacity value.

4.4.6 The dxsysinfo Utility May Display /dev/prf as a Tape Device

The `dxsysinfo` application may display `/dev/prf` as a tape device depending on the subsets installed on the particular machine. Also, when Update Devices is selected, another `/dev/prf` icon may be added to the device area.

4.4.7 dxshutdown Does Not Prohibit User Logins

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

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

4.4.8 Print Configuration Manger

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

- Aliases that conflict with system-assigned names

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

- Altered attribute validation

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

- Trailing comments

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

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

4.5 System Administration

The following notes apply to system administration.

4.5.1 Enhanced Security

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

4.5.1.1 Distribution of Enhanced Security Profiles via NIS

This note covers problems that may occur when distributing of enhanced security profiles via NIS.

- The NIS master server must always be available, or `login` attempts will fail after a timeout.
- Logins and failed login attemps for NIS-shared accounts require updating the last unsuccessful and last successful login fields, which requires a response from the master's `rpc.yppasswdd` process. That process will block waiting for any previous updates to complete. Completion of those updates is dependent on the completion of the yppush operation. The `yppush` completion is dependent on having the NIS slave servers being reachable and responsive. The more NIS slave servers a given NIS domain has, the slower this update process will be. Also, simultaneous login attempts will be processed sequentially by the NIS master, each waiting on the `yppush` for the previous to succeed. Thus, with several simultaneous attempts, some may time out and need to be retried by the users. This can be alleviated to some extent by using the `-p` option to `yppush`. One way to do this is to modify the file `/var/yp/Makefile` and change the `YPPUSH=` line. For example:

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

  This example allows up to six simultaneous transfers to NIS slave servers.

- NIS slaves which are listed in the `ypservers` NIS map on the NIS master but which do not already have a copy of the `prpasswd` and prpasswd_nonsecure NIS maps will not succeed in transfering those maps during the yppush operation. However, that failure will involve a timeout in the yppush on the master, and will contribute to login failures while waiting since the timeout for the yppush is longer than for login attempts.

This can be addressed by ensuring that all the NIS slaves have these maps by a procedure like the following (to be executed on each slave server which does not yet have these maps):

```
# /var/yp/ypxfr -d <domainname> -h NISMASTER -c prpasswd#
/var/yp/ypxfr -d <domainname> -h NISMASTER -c
prpasswd_nonsecure
```

In the above, substitute the name of the local NIS master server for the NISMASTER token. This will transfer initial copies of those maps for those slave servers.

- The time allowed for responses to RPC requests is only 25 seconds. The more profiles which are in the prpasswd map, the more likely that time limit is to expire during a login attempt, causing that attempt to fail. In testing, 4000 accounts seemed to be close to the upper limit. Simultaneous or nearly-simultaneous login attempts will fail if the NIS master server does not respond quickly enough to the waiting login processes. If the total time taken on the NIS master for the command

  ```
  # make passwd prpasswd
  ```

  exceeds 25 seconds, then only one user will succeed in logging in at a time. Internal testing has demonstrated that 4000 profiles and infrequent logins (3 at a time) can work, but even fewer profiles can be accommodated if bursts of nearly-simultaneous logins are frequent.

- Login updates are checked by the login processes. This requires that the yppush operation for the prpasswd map must take place. Therefore that map (at least) must be pushed during the normal operation of the rpc.yppasswdd daemon. Setting the /var/yp/Makefile file variable NOPUSH is not recommended for such configurations.

- Sites which cannot use NIS to share prpasswd information may be able to use NFS to share the files /tcb/files and /var/tcb/files directories instead. This will require exporting the directories with root access to the participating nodes (with -root=client1:client2:client3 or -root=0 as appropriate, see the exports(4) reference page). It also requires that NFS locking is enabled to ensure that no database corruption will occur.

4.5.1.2 Enhanced Security and Disaster Recovery

Because the user profiles and ttys information are now stored in database files, the previous recovery method of editing the files while in single-user mode is no longer available. However, as long as the `/usr` (and, if separate, `/var`) filesystems are mounted, the `edauth`(8) utility can be used in single-user mode to edit extended profiles and ttys database entries.

If the `/etc/passwd` file is somehow lost, but the extended profiles are still available, then a command sequence like the following can be used to recover some of the missing data:

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

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

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

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

4.5.1.3 Username Length Restriction

Enhanced security will not allow usernames longer than the documented maximum of 8 characters.

4.5.2 Bootable Tape

For this release, the LSM product and AdvFS `addvol` utility are not supported. Also, not all platforms and tape drives support bootable tape. Supported processor platforms are:

- DEC 3000-500
- DEC 3000-400
- DEC 3000-600
- DEC 3000-300
- DEC 3000-300X
- DEC 3000-900
- DEC 2100 4/275
- AlphaStation 600 5/266
- AlphaStation 200 4/100
- AlphaStation 200 4/166

Supported tape devices are:

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

4.5.2.1 Disk Space Issues

You should be aware of the following disk space issues before you use the `btcreate` command:

- If you have installed all the subsets and selected all kernel options, you need 120,000 (512 blocks) in `/usr` before you start `btcreate`.
- If your system has all the subsets with all the kernel options installed then the `btcreate` utility cannot be run on an RZ26 or lower type of disk with the default partitions. This results in the `/usr` file system being full. In that case, you can remove some of the subsets or rebuild the kernel with only the necessary options. In other words, make sure that you have the additional 120,000 (512 blocks) needed in the `/usr` file system before you run the `btcreate` utility.

  If you want to keep all the subsets along with all the kernel options, do the following to make extra space. Note, the examples in the following procedure are for UFS. For AdvFS, use the `mkfdmn` and `mkfset` commands to create new partitions and mount them.

  1. Run the `newfs` command to recreate a new partition; for example:

     ```
     # newfs /dev/rz1d
     ```

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

     ```
     # cd /usr/sys
     ```

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

     ```
     # mkdir FLAWLESS.BOOTABLE
     ```

  4. Mount the new partition on the *`SYSTEM`*`.BOOTABLE`directory; for example

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

5. Create another new partition; for example:

   ```
   # newfs /dev/rz1b
   ```

6. Mount the partition; for example:

   ```
   # mount /dev/rz1b /mnt
   ```

7. Change the current working directory to the `/usr/sys/bin` directory.

8. Copy the contents of the `/usr/sys/bin` directory to the `/mnt` directory:

   ```
   # cp * /mnt
   ```

9. Unmount the `/mnt` directory:

   ```
   # umount /mnt
   ```

10. Mount the new partition on the `/usr/sys/bin` directory; for example:

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

After completing the these steps, start `btcreate`. If you are using AdvFS, the `/usr/sys/bin` file system must be dumped during `btcreate` in order to copy the entire contents of the `/usr` file system.

4.5.2.2 Setting Swap Space on a Restored Disk

After restoring your system from bootable tape, you must set the swap space at the bootable tape single-user mode as follows:

1. Mount the root file system:

   ```
   # mount -u /
   ```

2. Change the current working directory to the `/etc` directory:

   ```
   # cd /etc
   ```

3. Add the swap space to the `fstab` file:

   ```
   # echo "/dev/rz3b swap1 ufs sw 0 2" >> fstab
   ```

4. Change the disk label:

   ```
   # disklabel -s rz3b swap
   ```

After you complete these steps, shut down and reboot the system from the restored disk.

4.5.2.3 Setting an Existing File System on a Restored Disk

If the restored disk already contains a file system that has not been touched during the `btextract` process, do the following to see and use that partition:

1. Set the file system type (*`fstype`*) on the partition. In the following example, `rz1d` is the partition and `4.2BSD` is the file system type:

   ```
   # disklabel -s rz1d 4.2BSD
   ```

2. Mount the partition on the file system you want to restore; for example:

   ```
   # mount /dev/rz1d /yourfs
   ```

4.5.2.4 Tape Drive Restriction

To use a tape drive with any system, make sure that the kernel has been built with the tape drive attached to the system. Otherwise, you get dump errors and the system cannot boot from the tape.

4.5.3 Multimedia Services for Digital UNIX V2.0 and V2.0A

The initial release of Multimedia Services for Digital UNIX Version 2.0 required several modifications to fully support Version 4.0 and Version 4.0 components. As a result, Multimedia Services for Digital UNIX Version 2.0A is shipped on the *Digital UNIX Associated Products Volume 1* CD –ROM. The Multimedia Services Version 2.0A runtime should be installed instead of the Version 2.0 runtime. The Multimedia Services Version 2.0 development kit, not distributed with the Digital UNIX Version 4.0 distribution, can be used with the Multimedia Services Version 2.0A runtime kit. See the release notes for Multimedia Services for Digital UNIX Version 2.0A for more details. The release notes can be found on the CD –ROM as

`DOCUMENTATION/HTML/MME201_RELNOTES*`

4.5.4 settime After an Update Installation.

The `settime` utility is called twice at boot time if an update installation was performed.

To prevent this happening on subsequent reboots, remove the link:

```
# rm /sbin/rc3.d/S05settime
```

4.5.5 System Clock Loses Time

When switching between two different boot disks that are running different versions of Digital UNIX, the system clock may lose time.

See Section 2.4.5 for a full description of the problem.

4.5.6 Symbolically Linked Kernels

In this release, `osf_boot` supports the booting of symbolically linked kernels. For example, assume you have `/tmp/vmunix` symbolically linked to `../mdec/vmunix` as follows:

```
lrwxrwxrwx 1 anyuser system 14 Dec 6 22:41 /tmp/vmunix -> ../mdec/vmunix
```

In this case, `osf_boot` will detect the link and boot `/mdec/vmunix` as follows:

```
Digital UNIX boot - Wed Dec 6 17:02:04 EST 1995

symbolically linked kernel detected: Loading /mdec/vmunix ...
Loading at fffffc0000230000
Current PAL Revision <0x4000000010530>
Switching to OSF PALcode Succeeded
New PAL Revision <0x4000000020123>
```

4.5.7 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.

4.5.8 Additional SCSI Disk Informational Entries in Error Log

To enhance the ability to detect certain SCSI disk device errors, additional event logging now occurs. However, these events can also occur during normal operation of the system. The events are known as SCSI unit attention events. To distinguish between normal events and abnormal events, the context within the event as well as the surrounding events must be considered. These entries begin as follows:

```
----- CAM STRING -----

ERROR  TYPE               Soft Error Detected (recovered)
```

The logging of these events will be associated with the first access after the event has occurred. First examine the ASC and ASQ values of the packet:

```
----- ENT_SENSE_DATA -----

ERROR CODE                    x0070     CODE x70
SEGMENT              x00
SENSE KEY                     x0006     UNIT ATTN
INFO BYTE 3          x00
INFO BYTE 2          x00
INFO BYTE 1          x00
INFO BYTE 0          x00
ADDITION LEN         x0A
CMD SPECIFIC 3       x00
CMD SPECIFIC 2       x00
CMD SPECIFIC 1       x00
CMD SPECIFIC 0       x00
ASC                  x29               <<<<<<<<<<<<<<<<<<
ASQ                  x02               <<<<<<<<<<<<<<<<<<
FRU                  x00
SENSE SPECIFIC       x000000
```

If the ASC is not x29, this is a device event or error (the SCSI specification contains the details of the meaning of the error). For entries with the ASC of x29, there are several possibilities for newer SCSI devices (older devices do not report these types of events):

- ASQ = x01 – device was powered on (or a firmware rebooted occurred)
- ASQ = x02 – Bus reset occurred
- ASQ = x03 – Device reset occurred

Events of ASQ x02 or x03 should also contain other related events in the error log (such as, a bus reset, or a device reset) for the specified bus or device. These types of events are just informational and do not indicate any type of failure on the system.

Events of ASQ x01 are normal events when they occur as part of switching on the system. The first access after the system is switched on creates one of these events. If these events occur at times other than the first access after the system is switched on, then it may indicate a problem within the device that is causing it to restart its firmware.

Also note that not all SCSI disks are capable of reporting this information.

4.5.9 Multiple Queues

If a printer is connected to multiple queues through a LAT or a local tty port and different jobs are submitted to different queues within a short period, some of the jobs may be lost. If this happens, resubmit the print request.

4.6 Network and Communications

The following notes apply to network and communications software.

4.6.1 Using bind on an UNIX Domain (AF_UNIX) Socket

When using `bind()` on a UNIX domain (AF_UNIX) socket, the default modes of the socket have changed in Digital UNIX version 4.0. Previously, the mode of a newly created socket was always 0777, regardless of the value of the creating process's `umask`. In this release, this behavior has changed so the the mode of a newly created socket is as follows:

```
(0777 &~ umask)
```

The previous behavior (0777, regardless of umask) may be restored by setting the kernel configuration flag `insecure_bind` to a value of 1. This can be done by either or both of the following two methods:

1. inserting the following lines into the `/etc/sysconfigtab` file:

   ```
   generic:
   insecure_bind = 1
   ```

   Then you must reboot the kernel.

2. using the following command on the running system:

   ```
   # sysconfig -r generic insecure_bind=1
   ```

4.6.2 Changes to ATM (Asynchronous Transfer Mode)

The following changes and restrictions apply to ATM

4.6.2.1 Changes to ATM Startup Scripts

The command syntax for the `atmarp` command has changed due to Multiple LIS support. The `atmconfig` command has added many new options mainly due to CBR (Constant Bit Rate) support. Lastly a new command, `atmsig`, is now required in all ATM startup scripts. Therefore all ATM startup scripts including `/etc/atm.conf` will need to be modified. See the reference pages associated with `atmconfig`, `atmarp`, and `atmsig` for further details.

4.6.2.2 Digital Gigaswitch Interoperability

When the ATM end system is connected to a Digital Gigaswitch, that is running Version 1.3 (or less) of the Gigaswitch software, the following line specifying the `useesi` keyword is required:

```
atmconfig up driver=lta0 useesi=1-4 wait
```

4.6.2.3 Support of NFS Locking over ATM

The `lockd` daemon will need to be restarted after ATM is brought up. This is because the `lockd` daemon takes a one time look at the IP interface list. ATM interfaces such as `LIS0` are dynamically added.

4.6.2.4 Restriction of Using atmsig down

If UNI signalling is disabled with the `atmsig down` command, it will not be correctly restarted by an `atmsig up` command. Following a command such as:

```
atmsig down driver=ltaX
```

You must do the following to successfully restart signalling:

```
atmconfig down driver=ltaX
atmconfig up driver=ltaX
atmsig up driver=ltaX
```

4.6.3 Large User Stack Limits May Cause a System Panic

Your system may crash when the user stack limits are increased beyond 2 gigabytes in the `/etc/sysconfigtab` file and more than 2 gigabytes are accessed by an application. The crash happens at some time after the location beyond 2G is accessed, typically when the system is paging. The default stack limits are below 2G and therefore this is a problem only if the stack limits have been increased.

The solution is to reduce the stack limits in the `/etc/sysconfigtab` and reboot. The application will then receive a fatal error when accessing stack above the specified limit.

4.6.4 SVR4 Streams IFNET paradigm.

The IFNET paradigm allows the bridging of streams device drivers to sockets. This release supports SVR4 streams, but the IFNET paradigm is not fully supported. IFNET is only supported over the `ln` ethernet interface and the number of `ln` devices supported is limited to two devices.

4.6.5 Orderly Release in XTI

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

4.6.6 Restarting the Network When Using Interface Aliases

When you restart the network using the following command:

```
# /usr/sbin/rcinet restart
```

The `ifconfig` command is run by the `/usr/sbin/rcinet` script. This will clear and reset the primary network interface address.

Network interfaces with configured interface aliases use the alias address as a source address for outgoing packets. Resetting the primary network interface address can cause a problem for systems with a firewall or proxy-access configuration based on the primary address. Generally, alias addresses are not in the access control lists in such systems.

To avoid this problem, you can use one of the following solutions:

- Add alias addresses in the firewall or proxy-access list in addition to the primary address
- Delete interfaces that are configured with interface aliases before restarting the network. Use the following command so that the restart script will correctly use the primary address in the outgoing packets:

  ```
  # ifconfig <if_w_aliases> down delete
  ```

4.6.7 Incorrect Error Message Using netsetup and rcinet

When restarting the network using `netsetup`, an error message similar to the following will be displayed:

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

This problem also exists when running the following commands:

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

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

4.6.8 CDE's Static Dependency on the Network

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

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

```
export DTNONETWORK=true
```

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

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

4.7 Local Area Transport

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

4.7.1 Duplicate Minor Numbers and latsetup

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

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

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

4.7.3 Simultaneous llogin Connections

When doing a number of simultaneous `llogin` connections it is recommended to use `llogin` with the `-p` option. To speed up an `llogin` connection, it is also recommended to add the target host name as a reserved service.

4.7.4 LAT Kernel Module Is Dynamically Loadable

It is no longer necessary to build LAT into the kernel. LAT is not made a mandatory kernel option upon selecting the LAT subset and will not appear in the kernel configuration file. As LAT requires the Data Link Bridge (DLB), it is still necessary to build DLB into the kernel when using LAT.

The default behavior upon booting to multi-user mode is for LAT to be dynamically loaded into the running kernel. If LAT is not started at boot-time via the `/sbin/rc3.d/S58lat` script, the recommended method for starting and stopping LAT is to verify that `LATSETUP` is enabled in

/etc/rc.config and execute the /usr/sbin/init.d/lat program, using the start or stop options.

4.7.5 Group Codes Are Compatible with Version 2.0 LAT Driver.

In the slave-only Version 2.0 LAT implementation, service group codes are used in solicitation messages for host-initiated connections,

Many users rely on the ability to control access to server ports by changing the group codes of the locally offered services. Although it is contrary to the recommendations of the LAT protocol, this behavior is once again supported. Outgoing port group codes, normally used for this purpose, continue to be used in all other cases where they are required by the protocol.

4.7.6 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.

4.8 File Systems

The notes in this section apply to file systems.

4.8.1 Using ACLs over NFS

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

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

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

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

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

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

- Alternatively, add the option to the mount command as follows:

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

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

4.8.2 ACL Size Limitations

On AdvFS filesystems there is a hard limit of 1560 bytes for a property list entry. Since Access Control Lists (ACLs) are stored in property list entries, this equates to 62 ACL entries in addition to the 3 required ACL entries. The error `EINVAL` will be returned if you attempt to exceed this limit.

To facilitate interoperation of the UFS and AdvFS ACLs, a configurable limit has been imposed on UFS ACLs. The default value of the UFS limit is 1548 bytes, equivalent to the 65 entry limit on AdvFS. The UFS configurable limit on ACLs has been added to the `sec` subsystem and has been given the attribute name `ufs-sec-proplist-max-entry`. The attribute can be dynamically configured using the `sysconfig` utility or by setting the attribute in the file `sysconfigtab`.

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

See the `cfgmgr`(8), `seconfig`(8), `seconfigdb`(8), and `sysconfigtab`(4) reference pages for more information.

4.8.3 POLYCENTER Advanced File System

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

4.8.3.1 Backward Compatability Problem with vdump and vrestore

Note

This is an important note for users of `vdump` and `vrestore`

Backups made using `vdump` on Digital UNIX Version 4.0 cannot be restored using `vrestore` on earlier versions of Digital UNIX. Patches will be made available for earlier versions of `vrestore` to correct this problem.

Backups made using `vdump` on earlier versions of Digital UNIX can be restored using `vrestore` under Digital UNIX Version 4.0 without problems.

4.8.3.2 Log Half Full Error

Under some circumstances, AdvFS can panic with the following message:

```
log half full
```

This can occur when a big percentage of a very large file is truncated and the fileset containing the file has a clone fileset. Truncation occurs when an existing file is overlaid by another file, and explicitly by the `truncate` system call.

The same will happen if very large, very fragmented files are migrated. Migration occurs when the `balance`, `rmvol`, and `migrate` AdvFS utilities are run. Files with greater than 40000 extents are at risk, unless the transaction logsize is increased as follows:

- 40000 extents use a logsize of 768
- 60000 extents use a logsize of 1024
- 80000 extents use a logsize of 1280

The command

```
# showfile -x <filename> | grep extentCnt
```

will indicate how many extents a file is using. Backup and restore will help to defragment files as will copying (not moving with `mv`) a file to another name. However, if the fileset has a clone and if a large file is truncated as a result of the copy, the truncation panic could occur.

4.8.3.3 ACL Size Limitations

On AdvFS filesystems there is a hard limit of 1560 bytes for a property list entry. Since Access Control Lists (ACLs) are stored in property list entries, this equates to 62 ACL entries in addition to the 3 required ACL entries. The error `EINVAL` will be returned if you attempt to exceed this limit.

To facilitate interoperation of the UFS and AdvFS ACLs, a configurable limit has been imposed on UFS ACLs. The default value of the UFS limit is 1548 bytes, equivalent to the 65 entry limit on AdvFS. The UFS configurable limit on ACLs has been added to the `sec` subsystem and has been given the attribute name ufs-sec-proplist-max-entry. The attribute can be dynamically configured using the sysconfig utility or by setting the attribute in the file sysconfigtab.

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

See the cfgmgr(8), seconfig(8), seconfigdb(8), and sysconfigtab(4) reference pages for more information.

4.8.3.4 The vdump and vrestore Commands

The vdump and vrestore commands do not have the same functionality as dump and restore commands.

When a file is renamed between different level dumps, the file is not backed up on the later dump; the vdump command assumes it was backed up on the earlier dump.

When a file is deleted between different level dumps, the file is restored during the restore process.

When a directory entry changes type (for example, a file becomes a directory) between different level dumps, during the restore of the higher level dump, you get an error message saying it cannot restore the file. The work around on is to remove the file that changed types between the different levels of restore.

4.8.3.5 The verify -F Command Does Not Bypass Recovery

Running the verify commmand with the -F flag causes some recovery to be done on the domain before the attempt to mount it.

4.8.3.6 rmfset Command May Hang

Avoid using the rmfset utility on busy domains. If you attempt to remove a fileset using the rmfset command and the target domain is experiencing a lot of I/O, the rmfset operation may hang.

4.8.3.7 AdvFS Sparse File Save and Restore

The vdump and vrestore utilities correctly save and restore AdvFS sparse files. In previous versions, the holes in the sparse files were allocated disk space and filled with zeros. Note that sparse files that are striped are still handled as in previous versions.

4.8.3.8 AdvFS Verifies Domain Volume Sizes

AdvFS will now verify at mount time that all of the data in all of the volumes in a domain can be accessed. It does this by attempting to read the last block in each volume as specified by the disk label that was in use at the time that the volume was added to the domain. If it cannot read that block, it attempts to read the last block that AdvFS has marked as being currently used to hold data. If AdvFS cannot read the last in-use block for any volume in the domain, the mount will fail. If it can read the last in-use block but cannot read the last block as specified by the disk label, the mount will succeed but in read-only mode.

One reason that the last block may not be able to be read is that a disk may be mislabeled on a RAID array. The user should check the labels of the flagged volumes in the error message. If the disk label is incorrect, the user can repair the domain in one of the following ways.

Before attempting corrective action, you should back up all filesets in the domain. The corrective action depends on the state of the domain. If the domain consist of multiple volumes and has enough free space to hold two entire volumes, it is possible to remove the offending volumes one at a time, fix the disk label and add them back to the domain. Perform the following operation on each of the failed volumes:

1. Remove volume from domain.
2. Correct the disklabel of the volume.
3. Add the corrected volume back to the domain.
4. Balance the domain.

Step 4 is important. If the domain is not balanced after adding the corrected volume, the user runs the risk of filling up one of the incorrectly labeled volumes and inducing an I/O error.

If the domain's free space is less than two volumes, you should back up all the filesets in that domain, remove the domain, fix the disk labels of the volumes, and rebuild the domain. Then restore the filesets from the backups.

Another example of why a mount might fail in this way would be that an LSM volume upon which an AdvFS domain resides has been shrunk from its original size.

4.8.3.9 AdvFS Domain Panic

When log or metadata write errors occur (for example, due to a disk failure or media error) AdvFS initiates a domain panic rather than a system panic on any non root file domains. A domain panic prevents further access to the domain, but allows the filesets in the domain to be unmounted.

When a domain panic occurs, a message is displayed in the following format:

```
AdvFS Domain Panic; Domain <name> Id  <domain_Id>
```

For example:

```
AdvFS Domain Panic; Domain cybase_domain Id 2dad7c28.0000dfbb
```

After a domain panic, use the `mount` command to list all mounted filesets then use `umount` to dismount all filesets in the domain specified in the error message. You can then take the necessary steps to correct the hardware problem. After you have corrected the hardware problem, it is recommended that you run the `verify` command (the domain structure checker) on the domain before remounting it. This will determine if the write error compromised the domain.

4.8.3.10 AdvFS Supports UFS Quota Commands

The UFS and AdvFS user and group quota commands have been consolidated. The standard UFS quota commands can now be used to manage user and group quotas on AdvFS. The following list identifies the old and new AdvFS quota commands:

Old AdvFS Command	New Consolidated Command
vquotaon	quotaon
vquotaoff	quotaoff
vquotacheck	quotacheck
vquota	quota
vquot	quot
vrepquota	repquota
vedquota	edquota
vncheck	ncheck

AdvFS quota functions have not changed. Functional differences between UFS and AdvFS quotas exist and are described in the reference pages for the consolidated commands.

The `/sbin/init.d/quota` script now checks and enables quotas for both AdvFS and UFS. This script runs during system initialization to stop or start user and group quota enforcement.

Support for the existing AdvFS versions of the quota commands will continue for some time. Future versions of AdvFS will drop the unique quota commands. Until then, both versions of the quota commands will work.

4.8.3.11 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.

4.8.3.12 Possible Check and Repair Failure

If a disk has a partition erroneously labeled AdvFS that overlaps a UFS partition, a file system check and repair operation (`fsck` and `ufs_fsck`) will fail on a partition that overlaps the AdvFS partition. The solution is to relabel the AdvFS partition on the disk.

Conversely, if a disk partition that overlaps AdvFS is erroneously labeled UFS, an AdvFS file system check and repair operation, `verify`, will fail on a partition that overlaps the UFS partition. To correct the problem, relabel the UFS partition.

4.8.3.13 AdvFS Command and Quota Problems

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.

4.8.3.14 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. Remove the entire domain by using the `rmfdmn` command. After the unused domain is removed, you can create a new domain on the partition.

4.8.3.15 Data Element Size Limits for Extended Attributes

Support for extended attributes (vfs+) in AdvFS is limited to data elements of 2KB or less. Application programs attempting to set larger attributes will receive an error return value.

4.8.3.16 Increase Metadata Extent Size for Large Number of Files

On systems with domains that contain very large numbers of files (over 5000), the standard AdvFS metadata extent page allocation may be inadequate. As a result, an incorrect error message is displayed as follows:

```
out of disk space
```

To avoid this problem, AdvFS provides two ways to configure your file domain to handle large numbers of files. You can use the `mkfdmn` command with the `-x` or `-p` flags to create a file domain. Then, if the file domain is extended beyond one volume, use the `addvol` command with the same flags.

See the `mkfdmn` reference page for complete details on using these flags. A table is included in the reference page to indicate the number of extents required for the number of files in the file domain.

4.8.3.17 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 manually 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 UNIX file systems.

4.8.3.18 Disk Usage Information

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

4.8.3.19 AdvFS Does Not Support the Hierarchical Storage Manager

Any attempt to enable shelving for an AdvFS fileset using the hierarchical storage manager (HSM) `mkefs` command results in the following error message:

```
Can't get current fileset shelving info - ENOT_SUPPORTED (-1041)
```

Also, any attempt to mount an existing AdvFS fileset that already has shelving enabled results in the following error message:

```
AdvFS mount - shelving not supported
```

To access an existing AdvFS fileset that already has shelving enabled, restore the data into another fileset that does not have shelving enabled.

4.8.3.20 Unique File Identifiers

Both the `fstat`(2) reference page and the `/usr/include/sys/stat.h` file inaccurately state that the combination of the `st_dev` and the `st_ino` fields of the `stat` structure create a unique file identifier through time. Because the `st_ino` value can be reused, the only way to create a unique file identifier is the combination of the `st_ino`, `st_dev`, and `st_gen` fields. While this is true in UFS, it is even more important in AdvFS, which recycles the `st_ino` values rapidly.

4.8.3.21 NFS Server May Print Error Message on NFS Client Recovery

If an NFS server is exporting an AdvFS directory and the client to which it is exporting crashes, upon reboot of the client, the following error message may appear on the console of the server:

```
lockd : can't clear lock after crash of client client_name : invalid
argument.
```

This message does not disrupt NFS operations.

4.8.3.22 AdvFS Panic

When a large percentage of a very large file is truncated and the fileset containing the file has a clone fileset, the system will panic with a `log half full` error message. Truncation happens when an existing file is overlaid by another file or when the truncate system call is made.

The panic also occurs if very large, very fragmented files are migrated. Migration occurs when you run the AdvFS `balance`, `rmvol`, and `migrate` utilities. Files with greater than 40,000 extents are at risk. To determine how many extents a file is using, enter the following command:

```
# showfile -x filename | grep extentCnt
```

The `backup`, `restore`, and `copy` commands tend to defragment files. The `mv` command does not. However, if the fileset has a clone and a large file is truncated as a result of the copy, the `log half full` panic again could result.

4.8.4 Logical Volume Manager

Logical Volume Manager (LVM) support is being retired in this release of Digital UNIX and there will be no further support of LVM.

All volume management functions are provided by the Logical Storage Manager (LSM). LVM functions are disabled with the exception of the support necessary to encapsulate LVM volumes under LSM. You must encapsulate any LVM volumes under LSM to maintain access to any data in such volumes. In a future release of Digital UNIX, encapsulation support will be dropped and any data still under LVM control will be lost.

LVM volume groups can be encapsulated to the `rootdg` diskgroup. Attempting to encapsulate an LVM volume group to any other LSM disk group fails.

You can encapsulate an LVM volume group to a non-rootdg diskgroup by performing the following procedure:

1. Use the `vollvmencap` command to encapsulate the LVM group to an existing LSM diskgroup. For example, to encapsulate the `vg1` LVM group to the existing `lvmdg1` LSM diskgroup, enter the following command:

   ```
   # /usr/sbin/vollvmencap -g lvmdg1 /dev/vg1
   ```

   The `vollvmencap` command creates the LSM scripts in the `/etc/vol/reconfig.d/lvm.d` directory.

2. Change the current working directory to `/etc/vol/reconfig.d/lvm.d`:

   ```
   # cd /etc/vol/reconfig.d/lvm.d
   ```

3. Move the `dg` file in the `/etc/vol/reconfig.d/lvm.d` directory to a sub-directory whose name corresponds to the LVM volume group that is being encapsulated; for example:

   ```
   # mv dg vg1
   ```

   Then change the current working directory to root.

4. Execute the `/sbin/vol-lvm-reconfig` command to execute the LSM scripts created by the `vollvmencap` command.

   When you execute the `/sbin/vol-lvm-reconfig` command, an error message is displayed; ignore this message.

For more information, see the *Logical Storage Manager* manual.

4.8.5 Logical Storage Manager

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

4.8.5.1 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.

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.

4.8.5.2 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*.

4.8.5.3 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. Refer to the LSM Installation documentation for more information.

4.8.5.4 The volrootmir Script Supports Only LUN 0 on HSZ

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

4.8.5.5 Boot Failures May Occur When Using LSM Volume rootvol for the root File System

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

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

If these entries are deleted or if the `/etc/sysconfigtab` file is deleted, the system will not boot. If this happens, you can boot the system interactively as follows:

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

After the system boots, edit the `/etc/sysconfigtab` file and add the LSM entries as shown above. Reboot the system for the changes to take effect.

4.8.5.6 Block Change Logging Subdisk Size Recommendation

LSM Volumes enabled with Block Change Logging (BCL) requires two or more log subdisks that are at a minimum one sector long. If you intend to use BCL on any volumes that do not already have logging subdisks, Digital recommends that you allocate at least two sectors to each log subdisk.

Any volumes that currently use single sector logging subdisks will continue to work correctly. However, Digital recommends that you reconfigure as soon as convenient to avoid being forced to do so at a later date.

Implementing these recommendations now will make the transition to new requirements in future releases easier.

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 Some C Library Functions Have Two Versions

In this release of Digital UNIX, the following C library functions exist in two versions due to conflicts between previous versions of Digital UNIX (to maintain binary compatibility these interfaces remain the default in this version of Digital UNIX as well) and the recent X/Open Single UNIX Specification. The Single UNIX Specification interfaces will be used when the `c89` command is used or when the C preprocessor symbol `_XOPEN_SOURCE_EXTENDED` is defined.

```
a64l          dbm_forder    getsockname   readv         sigaction
accept        dbm_nextkey   l64a          recvfrom      signal
basename      dbm_store     nftw          recvmsg       sigpause
dbm_delete    dirname       pipe          sendmsg       sigset
dbm_fetch     ftw           putmsg        setpgrp       ttyslot
dbm_firstkey  getpeername   putpmsg       setrlimit
```

The Single UNIX Specification interfaces are recommended for new designs and implementations. They will ensure portability of code on all UNIX platforms. In most cases the changes in the function will not require changes to the application source code using the function, just a recompilation.

5.2 Missing Debugger Information

When using the `-g` flag with the C compiler, debug information is not produced for small integer types (8-bit and 16-bit) in old-style function definitions. Correct code is generated, it is only the debugger information that is incomplete. The following example illustrates the types of definitions that will result in the missing information:

```
"int sub1(a, b, c)"
"int a;"
"short b;"
"char c;"
"{ ... }"
```

In this example, both `b` and `c` are small integer types, and `sub1` has an old-style definition. The following example uses the new prototype-style definition, so there is no missing information:

```
"int sub1(int a, short b, char c)"

"{ ... }"
```

You will notice the missing information when using the `where` command in dbx as follows:

```
% cc -g dbx.c
% dbx a.out
(dbx) where
>  0 sub1(a = 1) ["dbx.c":1, 0x1200011c4]
1 main() ["dbx.c":11, 0x120001220]
```

The missing information can also cause a problem if using the dbx `call` command, as the debugger interprets `sub1` as only having a single parameter:

```
(dbx) call sub1(5, 6, 7)
too many parameters in call to sub1
```

It is possible to work around this problem by converting to the ANSI C prototype style function definitions or by avoiding the small integer types as parameters in old-style function definitions. Alternatively, use the `-oldc` option when compiling.

5.3 DEC Ada Is Not Supported

DEC Ada is not supported in Digital UNIX Version 4.0.

5.4 Init Execution Order Modified

The execution order for init routines in static executables has been modified to more closely match the execution order for init routines in dynamic executables. See Section 1.8.6 in Chapter 1 for more information.

5.5 Using pixie on Applications Built with -om

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

The Atom-based tools, including a new version of `pixie`, cannot currently process executables produced with the `-om` switch. This is a limitation with the `om` utility which will be corrected in a future release.

5.6 pixie Command Can Fail to Locate atom Command

The `/usr/bin/pixie` command fails to invoke the `atom` command it needs, if the atom command's file is in the last directory in the process environment variable `PATH`.

The `atom` command is usually found in the `/usr/bin` and `/bin` directories, so when you use the `pixie` command, make sure that one of these is included in the value of your process's `PATH` and that it is not the last one.

The default `.login` and `.profile` files in `/usr/skel` (used to create new user accounts by `adduser`) define values for `PATH` that allow use of the `pixie` command in various shells.

The following code is included in the default `.login` file:

```
set path=($HOME/bin /usr/bin .)
```

The following code is included in the default `.profile` file:

```
PATH=$HOME/bin:${PATH:-/usr/bin:.}
```

5.7 c89 Command Conforms to Spec1170

The `c89` command now conforms to the Spec1170 requirements for command option spelling. The old command syntax is also supported. See the updated `c89` reference page for details.

5.8 Alternate Compiler Suite

The `CMPDEVALT400` subset contains an alternate application compiler suite. The README file contained in the subset explains its use.

5.9 DEC C

The default C compiler on Digital UNIX has changed to Digital's DEC C compiler. The previous default C compiler, is still a supported viable option and can be accessed by specifying `-oldc` on the `cc` command line.

Much work has been done to minimize the differences between DEC C and the previous compiler. Both compilers have full binary compatibility with each other. This release note provides a list of differences that may be helpful for you to know as you begin to use DEC C in Digital UNIX version 4.0.

- DEC C offers many features that were not available with the previous compiler. See the `cc`(1) and `c89`(1) reference pages and the *Programmer's Guide* for details.

- Latent errors in your source code may be exposed with this new compiler. Several types of these that we have observed include:
  - When accessing unaligned short data types through pointers, you will get an unaligned access fault with DEC C and not with the previous compiler.
  - The stack frames generated by DEC C are packed much tighter than those of the previous compiler. Consequently, any code that overwrites stack data is much more likely to corrupt the next stack variable with DEC C than with the previous compiler.
  - The use of uninitialized variables that worked with the previous compiler are not guaranteed to work with DEC C.
  - If data structures are not properly aligned on their natural boundaries, runtime unaligned access faults will more likely result when compiled with DEC C.
  - Compiling undefined C programming constructs may produce different behaviour with DEC C. In C, some constructs are undefined by ANSI standards and their behaviour may differ when complied on DEC C instead of ACC.
- Parameter cleaning, which the previous compiler did by default in all modes, is not done by default in `-std1` mode with DEC C. For more information, see `-ansi_args` in the `cc`(1) reference page.
- DEC C assumes the ANSI C aliasing rules are adhered to in the `-std` and `-std1` modes, and as a result, allows the optimizer to perform more optimizations. The previous compiler did not perform these optimizations.

  If your program does access the same data through pointers of a different type (for example, by a pointer to int and a pointer to float), then you must not allow the compiler to assume ANSI C aliasing rules. Otherwise, incorrect code might be generated. See the `cc`(1) reference page for details.
- DEC C does not support the `-j` command line switch. Equivalent functions are provided by either the `-ifo` or the `-O<n>` switches, where <n> is a number.
- By default, when no optimization level is specified, DEC C does perform a few more optimizations than the previous compiler did. See the `cc`(1) reference page for details.
- If your link line did not require linking with `libc` in the past, it may be required now as DEC C utilizes `_Ots` support routines found in `libc.` Prior to this release, these support routines were in `libots.`
- DEC C defines the `__DECC` macro whenever it is used to compile C programs.

- The maximum length for external names allowed by DEC C is 1024 characters. The previous compiler allowed longer external names.

5.9.1 Remove the DEC C -migrate Switch

If you are currently using DEC C with the `-migrate` switch, Digital encourages you to remove this option and use the default `cc`.

Removal is recommended because:

- The optimization switches have different mappings. See the `cc`(1) reference page for more information.
- You no longer need to explicitly link with the `libots` library. The `libots` routines used by DEC C now reside in `libc`.

If you call `ld` directly in the command string, replace the switch `-lots` with `-lc`.

5.9.2 Object File Address Reference Limit

There is a restriction with the DEC C compiler in that a single object file cannot contain more than 4093 external address references. The compiler will issue the following error message when compiling a module that exceeds this limit:

```
GP-relative section overflows limit of 32752 bytes
```

To workaround this problem break your module into smaller compilation units or use the `-oldc` compiler option.

5.10 DEC C -ifo Switch

When using the `-ifo` switch in DEC C, symbol information for external variables may be incorrect. For example:

```
% cat simple.c

int global_array[100][4];

main()
{
int local_array[100][4];
}

% cc -O1 -g2 -ifo simple.c
% dbx a.out
(dbx) whatis global_array
int global_array;
(dbx) whatis local_array
local_array[100][4] of int ;
```

Note that the global array is reported as `int` and the local array is a two dimensional array of ints. Both should report the same thing.

If this program is executed without `-ifo`, the results are as expected. The problem is in the generation of the symbol table, not the debugger.

5.11 XTI and TLI Libraries Are Thread-Safe

The libraries `libxti.a`, `libtli.a`, `libxti.so`, and `libtli.so` now support multithreaded applications. New `xti.h` (for XTI) and `tiuser.h` (for TLI) are available to be included by applications. Binary backwards compatibility has been preserved and unithreaded applications built with previous include files will still run.

5.12 DECladebug Product Name Change (Ladebug)

Note that the product name DECladebug has changed to Ladebug.

The release notes for Ladebug are in the subset `OSFLDBDOC400` in the *Digital UNIX V4.0 Operating System Volume 1* CD–ROM and can only be accessed when Ladebug is installed. After installation the release notes can be found in:

```
usr/doc/ladebug/ladebug-relnotes.txt
```

5.13 Analysis Tool with Object Modification (ATOM)

The following notes apply to the Analysis Tool with Object Modification utility `atom`.

5.13.1 Profiling Signal Handlers and fork with Multithreaded hiprof

A multithreaded program instrumented with the default `hiprof` command, `atom` *`appl_prog`*`-tool hiprof -env threads` will sometimes deadlock (enter an infinite polling and sleeping loop). This can happen if the program calls the `signal`, `sigaction` or `fork` function or if it was compiled with a `cc -speculate ...` command (which introduces signal handlers).

You can avoid this problem by adding the `-toolargs=-calltime` or the `-toolargs='-calltime -threads'` option to the command line.

For programs that call `fork` function, instrumenting `libc.so` is always recommended. You should also add the `-all` or `-incobj libc.so` option. Instrumenting `libc.so` also minimizes the risk of deadlock with the default command, `hiprof`.

If the default `hiprof` command must be used for a multithreaded program that calls the `signal` or `sigaction` and if `libc.so` must be profiled, add `-exc __sigtramp -exc __Esigtramp` to the `-toolargs` option used to exclude the signal handling procedures and any other procedures they call.

Alternatively, or when the program is compiled with a `cc -speculate ...` command, do not instrument `libc.so` by omitting `-all` or adding `-excobj libc.so`.

Similarly, if the application's signal handlers call many other procedures directly or indirectly (for example, in `libc.so`), the whole library can be excluded instead of excluding a long list of individual procedures.

5.13.2 Third Degree and hiprof Errors with Multi-threaded Applications

The Atom `hiprof` tool may not work correctly with threaded applications. The problem may cause threaded applications instrumented with the `hiprof` tool to get a segmentation fault and dump core.

You can avoid this problem by passing the `-exc __sigemptyset` switch with the `atom hiprof` command; for example:

```
% atom -tool hiprof -env threads -toolargs="-exc __sigemptyset"
myapp
```

This switch causes the `sigemptyset` routine to be excluded from the profile, but avoids the segmentation fault.

The Atom `third` tool (Third Degree) also may not work correctly with threaded applications. The problem may cause threaded applications instrumented with Third Degree to get a segmentation fault and dump core.

You can avoid this problem by including the `ignore __sigemptyset` entry in your `.third` file. This entry causes the `sigemptyset` routine to be excluded from Third Degree's error checking, but avoids the segmentation fault.

5.13.3 Atom third Tool (Third Degree) May Not Work Correctly with C++

The Atom `third` tool may not work correctly with C++ applications that use the tasking library (`libtask.so` or `libtask.a`). The problem causes Atom to issue a number of the following warnings:

```
atom: Warning: Object 't' has invalid instruction at 0xhexaddress,
treating as data.
```

These are followed by the following Third Degree assertion failure:

```
Assertion failed: pc==InstPC(inst)
```

You avoid the problem by linking your application with the `-call_shared` switch and passing the `-excobj libtask.so` switch to Atom. For example:

```
% atom -tool third -excobj libtask.so myapp
```

This switch excludes the tasking library from Third Degree's error checking, but avoids the assertion failure.

5.13.4 Atom Error with Dynamically Linked Fortran Applications

If a dynamically linked Fortran application is instrumented with `-all` or `-incobj libc.so`, you may get a loader error similar to the following when running the instrumented application:

```
pid:myapp.pixie: /sbin/loader: Fatal Error: Multiple instances of
shared object libc.so loaded as both /usr/shlib/libc_r.so and
```

You can avoid this problem by doing one of the following:

- If you have the ability to rebuild your application, you can relink it using the `-non_shared` switch. This produces an application that does not use shared libraries, thereby avoiding the problem.
- If you must use shared libraries but do not need to instrument `libc.so`, you can avoid the problem by passing the `-excobj libc.so -excobj libc_r.so` switches to Atom. For example:

  ```
  % atom -tool pixie -all -excobj libc.so -excobj libc_r.so myapp
  ```

  Note, this workaround does not work if you use the Third Degree tool, because Third Degree needs to instrument `libc.so`.
- If you need to instrument `libc.so` (or if you are using the Third Degree tool), you can avoid the problem by passing the `-incobj libc_r.so` switch to Atom to ensure that both `libc.so` and `libc_r.so` are instrumented. After running Atom, remove the instrumented version of `libc_r.so`. Then, replace the instrumented `libc_r.so` library with a hard link to the instrumented version of `libc.so`. Then run the instrumented application. The following example demonstrates the

workaround using the `pixie` tool.

```
% atom -tool pixie -incobj libc.so -incobj libc_r.so myapp
% rm libc_r.so.myapp.pixie
% ln libc.so.myapp.pixie libc_r.so.myapp.pixie
% setenv LD_LIBRARY_PATH .
% ./myapp.pixie
```

5.14 Change to Static Libraries and Kernel Objects

To save disk space, all the static system libraries and kernel objects 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.

5.15 DEC C++ Class Library Fixes

The following notes apply to DEC C++.

5.15.1 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.

5.15.2 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.

5.15.3 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.

5.15.4 LANG Environment Variable No Longer Cleared

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

5.15.5 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.

5.15.6 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.

5.16 General Programming

The following notes apply to general programming.

5.16.1 New Interfaces to the C Runtime Library

Certain C runtime library functions have new interfaces when compiled with _POSIX_C_SOURCE set to a value greater than or equal to 199506L. This is the default setting in the relevant include files if you do not set _POSIX_C_SOURCE to any other value. The relevant library functions are:

```
asctime_r()    getpwuid_r()
ctime_r()      gmtime_r()
getgrgid_r()   localtime_r()
getgrnam_r()   rand_r()
getlogin_r()   readdir_r()
getpwnam_r()   ttyname_r()
```

To access the old version, define the preprocessor symbol _POSIX_C_SOURCE to the appropriate value. See Section 5.16.3.1. Existing binary code will continue to reference the old, binary compatible interfaces.

5.16.2 prof Command Information Messages

The prof command suppresses informational heading messages in cases where a full profiling report will not be generated.

For example, when the -merge switch is used to consolidate profiling data files, no report will be generated. The informational messages from prof will be suppressed. These informational messages give a time and date stamp, and display the prof command line.

5.16.3 Realtime

The following notes apply to realtime programming.

5.16.3.1 SA_SIGINFO Not Visible Under Certain Namespace Conditions

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

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

5.16.3.2 POSIX 1003.1b Synchronized I/O and File Truncation

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

The fsync() or fdatasync() function can be used to explicitly synchronize truncation operations.

5.16.3.3 fcntl() and F_GETFL with O_DSYNC File Status.

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

5.16.3.4 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 should provide an adequate balance between efficiency and responsiveness to critical device errors.

5.16.4 DECthreads (pthreads)

The following notes apply to DECthreads.

5.16.4.1 The cc Command Supports the -pthread Option

The `-pthread` option has been added to the `cc` command. This option will provide the IEEE Std 1003.1C-1995, POSIX System Application Program Interface environment for developers of multithreaded and threadsafe code.

For compatibility the `-threads` option remains and provides a P1003.4a Draft 4 environment (P1003.1C was previously named P1003.4a). The `-threads` option should not be used in applications utilizing the 1003.1c-1995 environment.

5.16.4.2 Libraries libc and libc_r are Merged

The threadsafe functions in `libc_r.so` and `libc_r.a` have been merged into `libc.so` and `libc.a`. It is no longer necessary to link with `libc_r`, and this library has now been replaced with a link to `libc` for binary and makefile compatibility. See Section 5.16.4.1 for information on how to compile and link multithreaded and threadsafe code.

To determine the thread safety of individual `libc` functions, consult the appropriate reference page. The number of thread safe functions has been considerably increased in Version 4.0. In many cases the original API of functions for which there is a `*_r` variant are now also threadsafe. In these cases the `*_r` API is also retained for compatibility, and documented as obsolete.

5.16.4.3 DECthreads Exception Handling

DECthreads Exception Handling now relies on the DEC C exception handling capabilities. Threaded applications must now link with `libexc`. Note that if you are compiling your threaded application as Digital recommends (using the `-thread` or `-pthreads` switch) this change will not affect the building of your application.

If you are using the DECthreads Exception Package with POSIX 1003.1c threads interface, you must include `pthread_exception.h` in your source file. This does not apply to applications which use the CMA or draft 4 POSIX threads interfaces.

5.16.4.4 DECThreads Signal Model Incompatibilities in Threaded Applications

Full support of the POSIX 1003.1c signal model may cause a binary compatability problem for applications using threads. This is not an API change but a change in runtime behavior. A thread that was never interrupted before may now be interrupted.

5.16.4.5 dbx and DECthreads

You cannot use `dbx` to debug threads (pthreads) in this release. Use the Ladebug debugger to debug threaded applications.

5.16.4.6 Threaded Application Coding Errors

This release contains substantial changes to threads that will likely expose programming errors in existing applications that use DECthreads. Such errors include, but are not limited to, the following:

- Attempting to unlock a mutex that is not locked
- Use of uninitialized variables
- Improper use of data structures (for example, using a `pthread_attr_t` instead of a `pthread_mutexattr_t` in a call to `pthread_mutexattr_create`)
- Improper data access synchronization
- Use of an undocumented or unsupported routine

5.16.4.7 Missing Thread Routines

The following routines:

- `pthread_bind_to_cpu_np`
- `cma_thread_bind_to_cpu`

are not provided in this release. If you use an application that requires one of these routines, it is recommended that you not migrate to Digital UNIX V4.0.

Note that applications using `pthread_bind_to_cpu_np` are using the draft 4 POSIX interface which is scheduled for retirement. These applications should be migrated as soon as possible to the POSIX 1003.1c compliant interface. (See the retirement notice in this volume for `pthreads` draft 4).

5.16.4.8 DECthreads Static Libraries

The following static libraries for DECthreads have not been provided in this release:

- libpthread.a
- libpthreads.a

The libraries will be provided in a future release.

5.16.4.9 DECthreads System Cancellation Points

Threaded applications that require system cancellation points must use the POSIX 1003.1c DECthreads interface. Cancellation points are not supported for the CMA or draft POSIX DECthreads interfaces. See the *Guide to DECthreads* Appendix A for the list of supported cancellation points.

5.16.4.10 DECthreads and Forking

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

5.16.4.11 DECthreads and Realtime

In previous releases of Digital UNIX, thread scheduling attributes were system wide. In this release, all thread policies and priorities are now local to the process. Thread priorities between processes cannot be controlled. No artificial limit now exists on thread priorities: the full priority range is now accessible by every thread. See the IEEE POSIX 1003.1c standard or Appendix A of the *Guide to DECthreads*
for further discussion of contention scopes.

5.16.4.12 DECthreads and Alpha Firmware Revisions

Use the current minimum firmware revisions for Digital UNIX 4.0 to get the best DECthreads performance. The correct firmware revisions for your processor are identified in the Digital UNIX *Installation Guide*.

5.16.4.13 DECthreads and Signal Handling

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

5.16.5 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 resolve 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
   .
   .
   .
```

5.16.6 Warnings from sysconfig About 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.

5.16.7 Device Driver Example Files

Previously, Digital has supplied the `none.c` example driver in the `/usr/examples/devdriver` area. In this release, the `none.c` and `none100` kit examples have been replaced by the `ed100` driver example now in `/usr/examples/devdriver`. The file `ed100` is a directory that contains the driver source and the associated files required to configure the driver in Digital UNIX Version 4.0. Please refer to the *Guide to Writing Device Drivers: Advanced Tutorial* for more information about writing device drivers.

5.17 Functions and System Calls

The following notes apply to functions and system calls.

5.17.1 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.

5.17.2 shmat() Fails With Invalid Argument (EINVAL)

Existing applications that use large shared memory segments (>= 8MB) may fail during a `shmat` system call. This is because, by default, the system allocates shared page tables when a segment size >= 8MB is specified in `shmget` and therefore requires that attach addresses are aligned at 8MB. If an application uses an unaligned address to attach to such a segment and does not specify `SHM_RND`, the attach operation will fail and `EINVAL` is returned.

This behaviour can be modified using the `sysconfig` utility on a running system as follows:

```
# sysconfig -q ipc ssm-threshold
ipc:
ssm-threshold = 8388608
# sysconfig -r ipc ssm-threshold=0
ssm-threshold: reconfiguredsysconfig -q ipc ssm-threshold
ipc:
ssm-threshold = 0
```

If this behaviour is to be preserved across a reboot then edit the `/etc/sysconfitab` file to add the following:

```
ipc:  ssm-threshold = 0
```

To make use of shared page tables the application must be modified to use 8MB aligned attach addresses.

Applications that use a large shared region (such as databases) are likely to use more than one segment to form a contiguous shared region. This could be due to either of the following conditions:

- the configured segment size is small
- the configured segment size is at the largest SystemV shared segment of 2GB (Gigabytes).

For this to work correctly the segments must be attached to contiguous virtual addresses. Typically this is done by attaching the next segment at the address formed by adding the size of the segment to the start address. If a segment size is not a multiple of 8MB then, the attach address will be unaligned and it will fail if SSM is enabled.

Application vendors typically use the `shm-max` value to determine the size of a segment and configure an aligned size for `shm-max`. This is all that needs to be done to continue using SSM. The recommended optimum size is 2GB - 8MB = 2139095040 (0x7f800000).

To reconfigure `shm-max`, edit the `/etc/sysconfigtab` file, and add the following

```
ipc:
    shm-max=0x7f800000
```

Then reboot your system.

Window System Software Notes 6

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

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

6.1 Hardware Notes and Restrictions

The following notes apply to graphics hardware restrictions.

6.1.1 Calcomp Tablet

This section provides information about how configure a CalComp DrawingBoard III tablet, an input device supported by the Xinput extension to the Xserver. Once the software for the tablet is installed on your system, you can configure it to emulate a system mouse.

6.1.1.1 Configuring the CalComp DrawingBoard III Tablet

If you intend to use the CalComp DrawingBoard III tablet software, you must edit `/usr/var/X11/Xserver.conf` to turn on support for the X Input extension. To do this, remove the comment characters surrounding the following lines. The backslash in this example indicates line continuation and is not in the actual display:

```
input <
< _dec_xi_db3  lib_dec_xi_db3.so  XiDb3Init /dev/tty00:1:12:12:16:\
1:8:1000:1:1 >
>
```

You should also review these lines to ensure that the options specified for the tablet are correct, especially the `tty` that is specified as the serial port where the tablet is connected to your system.

The last line of this file has the following syntax. The backslash in this example indicates line continuation and is not in the actual display:

```
device:mode:tabletWidth:tabletHeight:numbtns:corePointer:mouseScale:\
resolution:Xincrement:Yincrement
```

Table 6-1 can help you determine how to set up the entries for the tablet in the `/usr/var/X11/Xserver.conf` file.

Table 6-1: CalComp DrawingBoard III Tablet Configuration Options and Values

Option	Description
device	The port (`tty`) to which the device is connected.
	The default is tty00.
mode	This should be set to 1 for absolute motion.
tabletWidth	Width of the active tablet area in inches, not the physical size.
	The default is 12.
tabletHeight	Height of the active tablet area in inches, not the physical size.
	The default is 12.
numbtns	Number of buttons on the puck or pen.
	The maximum number is 16 and the default is 16.
corePointer	0 indicates a native tablet mode (no system mouse).
	1 indicates emulate core pointer (the mouse and tablet are both core pointer devices.
	The default is 1 (emulate core pointer).
mouseScale	1 to 50 scaling factor in relative mode. Determines the speed of the cursor; the higher the number, the slower the cursor moves.
	The default is 8.
resolution	1 to 2540 lines per inch (lpi).
	The default is 1000.

Table 6-1: (continued)

Option	Description
Xincrement	How much the X axis must be incremented to cause the tablet to send new coordinates to the Xserver. The range is 0 to 65536. The default is 1.
Yincrement	How much the Y axis must be incremented to cause the tablet to send new coordinates to the Xserver. The range is 0 to 65536. The default is 1.

The device option is required and specifies which `tty` device should be associated with the tablet. By default, the installation software assigns the CalComp DrawingBoard III tablet to tty00, which you may want to change if that `tty` is already allocated. For information on how to determine which serial port your tablet is connected to, see the hardware documentation that was shipped with your processor.

Note that when the stylus or puck is moved as far as the minimum Xincrement or Yincrement value, the value of the corresponding axis is updated. For example, if the Xincrement value is set to 10 and the tablet is moved 10 units along the X axis, the value of the Y axis will also be updated simultaneously with the X axis, even if the Yincrement value has not been reached. Keep this in mind when setting the Xincrement and Yincrement options.

After you have configured the `/usr/var/X11/Xserver.conf` file, you must follow these steps to turn on support for the tablet in the Xserver:

1. Plug the tablet into your system and turn it on.
2. Enter the following command to restart the Xserver so that the Xinput extension can recognize the tablet (the backslash in this example indicates line continuation and is not in the actual display):

   ```
   # /usr/sbin/shutdown -r +5 "Turning on support for the Calcomp \
   Drawingboard III tablet"
   ```

   When the system comes back up, the tablet will be configured into the Xserver and ready to use.

When the Xserver first accesses the tablet, it performs some hardware-specific initialization that can be saved in the on-board memory of the tablet. To save these settings, follow these steps:

1. Press the EXIT CONFIG button on the tablet's menu.

2. Under the SAVE button, press the DEFAULT button.
3. Press the EXIT CONFIG button to save the settings.

6.1.1.2 Notes and Restrictions

The following notes and restrictions apply to the CalComp DrawingBoard III tablet.

- If the puck or stylus is not used within a 5 minute period, the tablet will automatically shut off. To reactivate it, press any button on the puck or stylus while they are in close proximity to the tablet.
- If you configure the tablet as the system's core pointer, moving the puck and the system mouse simultaneously will cause the cursor to move in an unpredictable fashion.
- Use only one puck or stylus at a time. If you try to use both input devices simultaneously, you will encounter unpredictable behavior.
- Digital UNIX does not support manual configuration of the tablet via the tablet buttons. If you try to use these buttons to configure the tablet, the Xserver will malfunction and may even crash. If you need to reconfigure the tablet, edit the `/usr/var/X11/Xserver.conf` file and then reboot the Xserver.
- You can modify some parts of the tablet setup by programming the Xinput extension. For more information on how to do this, see the XInput specification provided by the X Consortium.

6.1.2 Qvision Graphics Display Error

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

6.2 X Servers

The following notes apply to X servers.

6.2.1 Window System Response and Paging Load

The windowing system on AlphaStation 600 systems may experience poor response times in a heavy paging environment. This problem occurs on all platforms, but is most pronounced on AlphaStation 600 series processors. Mouse movement is unaffected, but windowing operations such as exposing a window may be degraded. When this condition occurs you may not be able to use `telnet` or `rlogin`. The poor response period may be greater than 2 minutes or it may stop after a few seconds. The cause is usually a memory intensive application that depletes the number of available pages on the free list to an absolute minimum. When this occurs, only privileged tasks are able to access any free pages. All other tasks are blocked, most notably the window manager. To verify if your system is experiencing this problem, open a terminal with `xterm` while your memory-intensive application is running and enter the following command:

```
# vmstat 2
```

The output from this command will show the number of free pages available under the heading `free`. If this number falls below the reserved amount `vm-page-free-reserved`, your system is exhibiting the memory problem. The following is sample output on a machine experiencing this condition:

```
Virtual Memory Statistics: (pagesize = 8192)
procs    memory          pages                           intr       cpu
r  w  u  act  free wire fault cow zero react pin pout  in  sy  cs  us  sy  id
2 37 15  5832 4021 1108  20K 5975 4575  285 4761    0  62 196 266   1   1  98
2 37 15  5851 4002 1108   99   13   50    0    8    0   4  13 239   0   1  99
3 37 15  7039 2788 1134   26    0   26    0    0    0  17  51 255   1   9  90
3 37 15  9791   10 1140 7291   17 7258    0   13   48 288  12 422   1  68  31
3 37 15  9790    9 1149 1049    0 1048   39    1  146 472  15 527   0  32  68
3 36 15  9784    9 1116  743    0  741    0    2  106 446  12 466   0  22  78
3 36 15  9811    9 1096  460    0  459    0    1  107 434  12 459   0  20  79
```

In this example, `vm-page-free-reserved` is set to its default value of ten pages and the VM subsystem has been overwhelmed. The system will demonstrate degraded window performance and other performance degradations. If the output from `vmstat` on your machine does not exhibit similar results, do not tune your system as Digital recommends in this note.

An interim solution to this problem is to:

- Raise the priority of the `mwm` or `dtwm` window manager.
- Change some `vm` (virtual memory) tuning parameters.

  Make the changes to three tuning parameters in the `/etc/sysconfigtab` file as described in the next paragraph.

Digital recommends that you systematically double the parameter values until you reach an acceptable balance between user interactivity and application performance. The following parameter values provide reasonable interactivity with moderate performance degradation on a Alphatop Station 6000 with 128M of physical memory These changes are intended to keep a larger than normal number of pages on the free list as a buffer:

- `vm-page-free-min=2048`

  This value enables the system to start paging earlier than usual.

- `vm-page-free-target=4096`

  This value enables the system to stop paging later than usual.

- `vm-page-free-optimal=4096`

  This value enhances the swapping of tasks to force the memory intensive application to be swapped out occasionally.

You must reboot your system after making these changes.

You may find that your system performs acceptably with some manipulation of these parameters. However, these suggested tuning parameters were devised to provide some level of user interactivity in a heavy paging/swapping environment. They may not be the best solution for your particular needs, but may provide a reasonable starting place.

If your system does not have 128MB of memory, reduce the suggested numbers. The units for these values are pages (which are 8KB) resulting in a substantial 32MB of memory use. If you are uncertain, try doubling the default parameter values until you reach acceptable perfromance. You can check the default values using dbx on a running system as follows:

```
# dbx -k /vmunix
(dbx) p vm_page_free_target
2048
(dbx) p vm_page_free_optimal
1029
(dbx) p vm_page_free_min
20
(dbx) p vm_page_free_reserved
10
```

Refer to the *System Tuning and Performance Management* manual for information on these parameters and system tuning in general.

To raise the priority of the window manager, open a terminal window and find the process ID of the window manager found using the `ps` command. Stop the window manager process with the `kill` command.

Restart the window manager using `nice` as follows:

```
# nice -2 dtwm &
```

IN this example, replace `dtwm` with `mwm` or `twm` if you are using a different window manager. If your system is tuned correctly, there should be is a small latency when the memory intensive application starts, which will balance as the memory subsystem recovers. You should be able to expose windows and drag them around with a reasonable response time. If you expose a window that has not been used for some time, expect some latency. This is because you are now working in a heavy paging/swapping environment and those pages are not yet available.

To avoid system panics after modifying the `sysconfigtab` file, you must use this release of Digital UNIX. Keep in mind, you will need to kill and restart the window manager each time you terminate your xsession. Note that to use `nice`, you need to be a privileged (root) user.

6.2.2 Limited Multiscreen Display Support With CDE

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

6.2.3 Keymap Interoperability Problem

If you use the Keyboard Options desktop application `/usr/dt/bin/dxkeyboard`, to change the keymap, X11R5 clients will not be notified of the change. For example, a `dxterm` running on DEC OSF/1 Version 3.2 displaying to a Digital UNIX Version 4.0 system will not be notified that the keyboard mapping changed.

To avoid this problem, enter the following command after using `dxkeyboard` to change the keyboard mapping:

```
xmodmap -e "keycode 248 = NoSymbol"
```

6.2.4 Do Not Modify the keymaps.dir File

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

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

6.2.5 Interleaf Display Problem

Some versions of Interleaf running under Digital UNIX exhibit display anomalies when manipulating Interleaf PullDown or PopUp menus. This will appear as a section of the screen that is not correctly repainted when the menu is removed.

This is due to a problem with Interleaf in relying on the contents of a Save-Unders region without responding to expose events. The problem may be fixed in a later version of Interleaf but can also be eliminated by disabling Save-Unders on the X Server command line. Do this by editing the `/var/X11/Xserver.conf` file to add the `-su` argument to the `args` statement. Refer to the `Xdec`(1X) reference page for the exact syntax.

6.3 X Clients

The following notes apply to X clients.

6.3.1 X Window Colormap Resources

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

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

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

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

    ```
    Netscape*maxImageColors:        96
    ```

  - Start Netscape with the `-install` flag, which specifies that Netscape should install its own colormap. Although this is supported, there are side affects such as:
    - Noticeable application and background color changes when focus is moved in and out of the Netscape window.
    - Inability to clearly view the contents of the Netscape window when focus is moved out of the Netscape window.

6.3.2 Manual Merge of DXnotepad Resource Definition File

A new resource definition (`*DXmfitToScreenPolicy: AS_NEEDED`) has been added to the `/usr/lib/X11/app-defaults/DXnotepad` resource file for the C locale.

If you are doing an update installation, this new resource is not automatically merged into your existing `/usr/lib/X11/app-defaults/DXnotepad` file. Therefore, you must manually manually append the following text to your C-locale `DXnotepad` resource file:

```
*DXmfitToScreenPolicy:  AS_NEEDED
```

This definition tells the toolkit to add scroll bars to a dialog box if the dialog box does not fit on the screen. It is required to support new workstations with smaller monitor resolutions.

6.3.3 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.

6.3.4 Switching Locales in dxmail

When `dxmail` is run using a new locale setting, mail subjects in the mail folders are still encoded in the codeset used by the previous locale. Non-ASCII characters in these subjects will be displayed incorrectly.

To correct this, use the rescan mail folder function to convert all existing subjects lines to the codeset used by the new locale.

6.3.5 Clipboard Interoperability Problem

MH and A Motif clipboard interoperability problem prevents applications 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. `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.

6.3.6 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.

6.3.7 Motif Version 1.1.3 Clients

Motif applications built on a release of Digital UNIX prior to v2.0 rely on the Motif 1.1.3 libraries. These retired runtime libraries are no long installed by default. These libraries are to be found in subset `OSFMOTIF11400.`

If an application that relies on these retired libraries is executed, it will fail with an error message. For example:

```
# xmosaic
```

This command produces an error message similar to the following:

```
1166:/usr/local/bin/xmosaic: /sbin/loader: Fatal Error:
 object libXm.so from liblist in /usr/local/bin/xmosaic has
 version "", which does not match the found object:
 /usr/shlib/libXm.so (with version "motif1.2")
```

6.3.8 Old X Environment

If `/usr/sbin/xsetup` is used to switch login environments from the CDE to XDM, it may need to install the subset `OSFOLDX11400` to regain the expected user environment. If this subset is missing, the default X session will consist of a single `xterm` and the `twm` window manager instead of the more familiar Digital session manager.

6.4 CDE Clients

The following notes apply to CDE clients.

6.4.1 Remote Invocation of CDE File Manager dtfile

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

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

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

6.4.2 Autopause and Enhanced Security

When using enhanced security, Autopause does not allow you to resume a suspended CDE session.

If you are not logged in as root, the dtsession Autopause and Locking features will not unlock. The session must be unlocked either by typing the root password or remotely killing the dtsession process, terminating and restarting the session.

The dxpause program exhibits the same behavior and there is no workaround for this problem. To avoid the problem, turn off the Screen Lock toggle on the Style Manager menu, Screen option.

6.4.3 dtmail and the Environment Variable MAILLOCKING

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

6.4.4 dtmail Disables Tooltalk Locking By Default

DtMail will disable tooltalk locking by default. To enable it, select the following option from the Mail Options->Advance dialog box:

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

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

```
cdenotooltalklock='f'
```

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

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

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

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

6.4.5 dtmail Row and Column Display

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

6.4.6 dtmail Attachments are Not Included

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

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

Include the attachments manually using the drag and drop feature.

6.4.7 dtmail Start-up

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

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

```
"No memory available for operation"
```

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

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

6.4.8 CDE's Static Dependency on the Network

The Common Desktop Environment (CDE) has a static dependency on the state of the network configuration. For more information, see Section 4.6.8.

6.4.9 xnlLanguage Resource Can Cause Problems with CDE

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

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

6.4.10 Possible Failure in the XOpenDisplay Call

When logging in to the CDE desktop, it is possible that not all desired applications will be restarted. The X server process may not be able to handle all of the requests for new open connections causing some to fail in the `XOpenDisplay` call. Some applications, like `xterm`, will log startup errors, such as:

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

in the dxconsole window.
```

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

```
Dtsession*contManagement:   2
```

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

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

6.5 Windows Programming

The following note applies to windows programming.

6.5.1 X11 Font Problem

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

6.6 Internationalization

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

6.6.1 Using Taiwanese EUC Codeset with Chinese DECterm Software

Chinese DECterm software supports the use of DEC Hanyu, Big5 and DEC Hanzi as its terminal code. To process Taiwanese EUC codeset, users should do the following:

1. Start the Chinese DECterm software using the zh_TW.dechanyu locale,

```
% /usr/bin/X11/dxterm -xnl zh_TW.dechanyu
```

2. In the DECterm just started, issue the following command:

```
% setenv LANG zh_TW.eucTW
% /usr/i18n/bin/stty adec acode eucTW tcode dechanyu
```

   After that, users can use terminal applications to process Taiwanese EUC files.

6.6.2 Backspace Key Not Supported in Asian Input

On PC-style keyboards, the Hanyu, Hanzi, Hangul and Kanji input servers use the Delete key instead of the Backspace key to erase information.

6.6.3 Transport Layer In Input Servers

The Digital UNIX input servers for Hanyu, Hanzi, and Kanji support only X and local transports. Transports over DECnet and TCP/IP are not supported. DECnet and TCP/IP transports are available only for IM servers that have a single input-server atom name.

6.6.4 Asian Multiscreen Support

X11R6 input servers can connect to R4, R5 and R6 clients under the same locale. However, if the system is using multiscreen support, only R6 clients can be connected.

6.6.5 Hanyu IM Server Support of EDPC Under zh_TW.big5 Locale

When the Language setting is `zh_TW.big5`, EDPC support must be turned off. Otherwise, random characters will appear in the candidate list for unsupported 4-byte encodings. Select the Input Method Customization... option from the Options button of the input server window in the lower left corner of your screen. Then in the customization dialog box, click the EDPC support button to off.

6.6.6 String Conversion and the ja_JP.SJIS Locale

String conversion using `XStringStyle()` is not supported for interclient communication under the `ja_JP.SJIS` locale. This is because the X library assumes that the 7-bit string data has `ISO8859-1:GL` encoding, whereas the `ja_JP.SJIS` locale assumes the data has `JISX0201.1979-0:GL` encoding.

6.6.7 Using Bookreader to Print Chinese and Korean Characters

The new mdxbook command in the IOSWWOLDBKR400 subset is a variant of dxbook Version 3.0. Unlike dxbook, the mdxbook command supports printing files with Chinese and Korean characters.

The display quality of mdxbook is not as good as dxbook. Therefore, install mdxbook only if you need to print Chinese and Korean characters.

6.6.8 Hebrew Editing Style

The following notes apply to Hebrew support.

- Hebrew Input Method (HIM) supports the XIMPreeditNone editing style only (no preedit area is needed). HIM currently supports the following status styles
  - XIMStatusNothing (a common status area for all clients)
  - XIMStatusNone (no status display)
  - XIMStatusArea (currently w/o dynamic area change and negotiation).
- To have the XIMStatusNone status, you should also define the environment variable (the value does not matter):

  ```
  HIM_ ALLOW_STATUSNONE
  ```

  If XIMStatusNone is required but the variable is not defined, XIMStatusNothing will be used.
- In a Motif 1.2 application, Motif uses XIMStatusNone for Hebrew, but XIMStatusNothing will be applied unless the XIMStatusNone environment variable is defined.
- If XIMStatusNone is being used and the Keyboard Manager (KM) is running, keyboard status will be reported by HIM, as it is done for XIMStatusNothing.

6.6.8.1 Cut and Paste Operation

The environment variable ADECW_MOTIF12 is needed for the cut and paste operation of the DXmCSText to work properly in the Hebrew environment. If you encounter problems with the cut and paste operation of a mixed English/Hebrew/Numeric strings, use this workaround and define the variable in your working environment (the value is arbitrary).

6.6.9 Hebrew CDE Users Should Avoid Using Certain Applications

Hebrew CDE users should avoid using the following CDE applications:

- Text Editor – `dtpad`
- Terminal – `dtterm`
- File and Application Manager – `dtfile`

These applications can hang in the Hebrew locale because of problems in the Hebrew input method.

6.6.10 Notepad Window Display and Input Style

When the notepad application is started with off-the-spot input style, the input area attached at the bottom of the application disappears during split view and delete view.

6.6.11 Default Keyboard Mapping

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

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

For example:

```
-pn -xkbmap digital_japanese_lk411aj
```

The available keymap files are located in `/usr/lib/X11/xkb/keymap`, where there is one file for each locale. The individual keymaps for the locale are in the keymap file, and are specified by the keyword `xkb_keymap`. For example, the keymap file `/usr/lib/X11/xkb/keymap/digital_japanese`, contains an entry for the `lk411aj` keymap.

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

6.6.12 Lithuanian and Slovene Language Support

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

6.6.12.1 Mnemonics in Lithuanian and Slovene DX Applications

DX applications V4.0 do not support the use of mnemonics in Lithuanian and Slovene language variants.

6.6.12.2 DX Applications Integrated to CDE Desktop (Application Manager)

For the Visual Differences application, `/usr/bin/X11/dxdiff`, to display differences between two text files written in Lithuanian or Slovene language you must do the following. Edit the resource file `DxDiff` located on `/usr/lib/X11/app-defaults` directory and change the following line:

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

to:

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

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

6.6.12.3 DX applications not integrated into the CDE desktop (Application Manager)

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

Note that all these actions must be performed by superuser.

- Mail –/usr/bin/X11/dxmail

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

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

```
*Text.FontList: *-*-*-Medium-R-Normal--*-120-*-*-M-*-ISO8859-1
```

to:

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

- Notepad – `/usr/bin/X11/dxnotepad`

  To use Lithuanian and/or Slovene national characters in `dxnotepad`, edit the resource file `/usr/lib/X11/app-defaults/DXnotepad` and change the following line:

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

  to:

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

- Bookreader – `/usr/bin/X11/dxbook`

  In order to use Lithuanian and/or Slovene national characters in Bookreader bookshelf and book names, edit the resource file `/usr/lib/X11/app-defaults/DXBookreader` and change all occurences of string:

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

  to:

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

- Paint – `/usr/bin/X11/dxpaint`

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

- Clock – `/usr/bin/X11/dxclock`

  The `dxclock` application displays date and time using english day/month abbreviations and format.

- Calendar – `/usr/bin/X11/dxcalendar`

  The `dxcalendar` application displays date and time using english day/month abbreviations and format.

  To create calendar entries using Lithuanian and/or Slovene national specific characters, edit the resource file `/usr/lib/X11/app-`

`defaults/DXcalendar` and change the following two lines:

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

to:

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

6.6.12.4 DX Applications Invoked from DXsession

In addition to the new default Common Desktop Environment (CDE), Digital UNIX V4.0 still supports DECwindows interface as an optional user interface. However, if you want to use the old DECwindows user interface you need to log in after selecting Dxsession session from Options --> Session menu of Digital UNIX login screen.

6.6.12.5 Lithuanian Language Variant

Lithuanian language variant users that wish to operate in Dxsession environment must specify the default user interface font by editing the file `.Xdefaults` located on their home directory and adding the following line and log in again:

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

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

Users that wish to have the default user interface font to be used with the Motif window manager need to edit the resource file `/usr/lib/X11/app-defaults/Mwm` and change the following line:

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

to:

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

Note that this change must be performed by a superuser (root).

Documentation Notes 7

This chapter discusses notes that apply to Digital UNIX Version 4.0 documentation.

7.1 Release Notes

The on-line HTML and PostScript versions of the release notes are not complete due to last minute changes. Use only the printed version supplied with the software media.

7.2 Media CD–ROMs

The Digital UNIX distribution media is comprised of three CD-ROM discs.

- Digital UNIX V4.0 Operating System Volume 1

  Contains the following:

  – Digital UNIX operating system

  – The *Release Notes* and *Installation Guide*

    PostScript versions are located in the `mnt_point/DOCUMENTATION/POSTSCRIPT` directory. Text versions are located in the `mnt_point/DOCUMENTATION/TEXT` directory.

- Digital UNIX V4.0 Associated Products Volume 1

  Associated software. See the *Installation Guide* for a description.

- Digital UNIX V4.0 Documentation Volume 1

  Disc 3 Contains the Digital UNIX Bookreader and HTML files for the Digital UNIX documentation set. It also contains files for worldwide support documents and some documentation for Digital UNIX associated products such as DECevent. HTML files can be viewed with Netscape; see the *Digital UNIX Installation Guide* for more information.

7.3 General Information About the Documentation

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

7.3.1 New Manuals

A number of manuals have been added to the documentation set. Refer to the *Documentation Overview, Glossary and Master Index* for more information.

7.3.2 Revised Documents

The title page of each document has an updated operating system revision number if the book has been revised since Version 3.0

All other documents labeled for Version 3.0 or later of the DEC OSF/1 or Digital UNIX operating system are applicable to this release.

7.3.3 List of Available Patches

Lists of patches for known problems are provided on the Digital UNIX Operating System Volume 1 CD–ROM, in the `<mount`point>/DOCUMENTATION/TEXT directory. The files are named `Digital_UNIX_V3.2C_CLD_Fixes.txt`, `Digital_UNIX_V3.2D-1_CLD_Fixes.txt`, and `Digital_UNIX_V4.0_CLD_Fixes.txt`.

7.3.4 Compressed PostScript Files

Some PostScript format files on the CD–ROM may be stored in compressed PostScript. This compressed format saves disk storage space and it requires less time to copy the files to other media. To decompress the files, use the `uncompress`(1) or `zcat`(1) commands.

There is an instruction file on the CD–ROM named as follows:

```
/DOCUMENTATION/POSTSCRIPT/00-READ-ME-FIRST
```

7.4 Reference Pages

This section provides information about the Digital UNIX Version 4.0 reference pages.

7.4.1 [XPG4_UNIX] Used Incorrectly in Section 1 Reference Pages

In this release, reference pages for interfaces and commands that are included in the Single UNIX Specification include tags to differentiate standard and proprietary information. The `standards`(5) reference page explains how standard-conformant and proprietary information should be differentiated in reference pages.

The `[XPG4_UNIX]` tag should appear at the beginning of text paragraphs when the paragraph discusses a UNIX extension added by Issue 4, Version 2 of X/Open CAE specifications. Text should not be flagged when discussing features of a command or function that were included in Issue 4, Version 1 of X/Open CAE specifications. However, in many reference pages for user commands, the `[XPG4_UNIX]` tag appears before all descriptions of standard-conformant features rather than only before descriptions of the UNIX extensions.

7.4.2 webman Utility for Netscape

The Digital UNIX system's documentation CD-ROM includes a tool called `webman`, which allows you to view the reference pages with the Netscape Navigator World Wide Web browser.

However, because of security restrictions imposed by the browser, it is not possible to view the files by opening them locally on the user's own system. Making the reference pages viewable requires that you set up your system (or one of the systems on your network) as a World Wide Web server by installing and running a server daemon. With a server running, all of the systems on your network can view the reference pages. Server software is available from third-party vendors such as Netscape Communications Corporation.

When the documentation CD-ROM is mounted as instructed in the *Installation Guide*, instructions for installing `webman` are in the following file:

```
/usr/share/doclib/online/DOCUMENTATION/HTML/webman/Installing_webman.txt
```

7.4.3 Modes of UNIX domain in the The bind() Reference Page

The bind() manpage does not discuss the modes of UNIX domain sockets. In previous releases, the mode was always 777. In this release, Digital has introduced a tuneable kernel parameter `insecure_bind`, which can be set either in the `generic` stanza of `/etc/sysconfigtab`, or can be set using the `/sbin/sysconfig` utility. The new behavior for `bind()` is that if insecure_bind has a value of 0 as follows:

```
insecure_bind=0
```

Then new AF_UNIX sockets have the mode set as follows:

```
0777 &~ umask
```

If insecure_bind is not set to =0, new AF_UNIX sockets have mode 0777 The default setting is insecure_bind=0.

7.4.4 Nonexistant ACL Interfaces in the acl(4) Reference Page

The following ACL interfaces are mentioned in the acl(4) reference page, but do not exist. Where possible, alternatives are specified:

- acl_set_id – use acl_set_qualifier
- acl_set_name – use acl_set_qualifier
- acl_set_perm – use acl_set_permset
- acl_convert_posix – unsupported, there is no substitute
- acl_set_acl_fd – use acl_set_fd
- acl_set_acl_file –use acl_set_file
- acl_validate_and_sort – sorting is unsupported. For validation use acl_valid
- acl_make_dec – unsupported, there is no substitute

See the acl(4), acl_set_qualifier(3), acl_set_permset(3), acl_set_fd(3), acl_set_file(3), and acl_valid(3) reference pages.

7.4.5 Property List Reference Pages

The setproplist(3), getproplist(3), and delproplist(3) reference pages should be corrected as follows.

In the EXAMPLES section of the setproplist(3) reference page, the definition for struct proplistname_args is both incorrect (no variable) and unnecessary (only getproplist() and delproplist() use this structure.)

In the SYNOPSIS section of the getproplist(3) reference page, the int *buf argument to getproplist() should be int *min_buf_size. In the EXAMPLES section, the struct proplistname_args; definition should be struct proplistname_args getargs.

In the EXAMPLES section of the delproplist(3) reference page, the struct proplistname_args definition should be struct proplistname_args entry_names;.

7.4.6 mail_manual_setup Reference Page

In step 7 of the instructions for setting up `POP`, use the command `create=20` to create a file called `/usr/spool/pop/POP` instead of a directory.

7.4.7 quota and vquota Show Wrong Usage of Options

The `quota` and `vquota` reference pages show the wrong usage of options.

The correct usage for the `quota` reference page is as follows:

quota [**-qv**] **-u** *username ...*

quota [**-qv**] **-g** *groupname ...*

The correct usage for the `vquota` reference page is as follows:

vquota [**-guqv**]

vquota [**-qv**] **-u** *username ...*

vquota [**-qv-g** *groupname ...*

7.5 Device Driver Documentation

The following notes apply to device driver documentation.

7.5.1 Writing VMEbus Device Drivers

The following notes apply to writing VMEbus Device Drivers.

- Updating Digital UNIX Device Drivers to the Current Version of the Digital UNIX Operating System

  Although existing VMEbus device drivers are compatible with the current version of Digital UNIX, you should plan on migrating your device drivers to take advantage of the single binary module technology. Refer to Writing Device Drivers: Tutorial for more information on the single binary module technology, and for previous device driver mechanisms that are not being supported in this release, such as the `config.file`, `stanza.static,` and `stanza.loadable` files. The syntax in the `files` file fragment also has changed to accommodate the single binary module in this release of Digital UNIX.

- Device Driver Configuration

  VMEbus device drivers written according to guidelines for Digital UNIX V4.0 single binary module technology are supported by the Static configuration method only. The Dynamic configuration method will be supported in a future release of Digital UNIX.

- Defining Bus-Specific Name Constants

  Device drivers written according to Digital UNIX V4.0 single binary module technology must specify `vba` as the bus name when creating a controller structure. Device drivers cannot specify `*` for a VMEbus device driver.

- VMEbus Configuration Information

  In previous versions of Digital UNIX, VMEbus configuration information was passed to device drivers through the config.file file fragment. VMEbus csr, csr2, vector, and interrupt request priority level were specified using the configuration keywords

  – `csr`
  – `csr2`
  – `vector`
  – `priority`

  respectively. This mechanism has changed for Digital UNIX V4.0.

  As you read the Digital UNIX V4.0 Writing VMEbus Device Drivers book, you should refer to the following description of VBA_Option when the book discusses the configuration keywords `csr`, `csr2`, `vector`, and `priority`. The book will be updated in a future release of Digital UNIX to reflect these changes.

  New and existing device drivers need to supply VMEbus configuration information by specifying a `VBA_Option` entry in `/etc/sysconfigtab.`

  `VBA_Option` fields have the following definitions and restrictions:

Field	Definition and Restriction
`Manufact_Name`	Specify up to a 31-character manufacturer name for this field. Enclose the name in single quotes (').
`Product_Name`	Specify up to a 31-character product name describing the product that the device driver is supporting. Enclose the name in single quotes (').
`Bus_Instance`	

Field	Definition and Restriction
	Specify the `vba` bus number that this device driver is being installed on. In most cases the `Bus_Instance` will be specified as zero. If the system consists of multiple VME buses, specify the appropriate bus instance, i.e. 0, 1, 2, etc. The bus numbers are associated with the physical order in which `vba` buses are detected.
`Driver_Name`	Specify up to a 15-character device driver name. Enclose the name in single quotes ('). The Driver_Name must match the subsystem name used in the V4.0 single binary module interface `create_controller_struct` or the driver name specified in the configuration file fragment in previous versions of Digital UNIX.
`Driver_Instance`	Specify the instance of the hardware device for which this device driver is being installed. If this is the first instance specify 0. This value is used to match against the controller number assigned in the V4.0 single binary module when a controller structure is created, or the controller number specified in the configuration file in previous versions of Digital UNIX. The same device driver can be installed multiple times to control different VMEbus device options of the same type. Typically. the only differences are the VMEbus device register addresses, interrupt vector and possibly the interrupt request level.
`Csr1`	Specify a VMEbus device register or memory address to map into CPU I/O space. A nonzero value in this field and the device driver's driver structure elements `addr1_size` and `addr1_atype` will cause the autoconfiguration software to map the specified VMEbus address space into system I/O space. The controller structure element `addr` will receive the resultant `io_handle_t` address. This value will be passed to the device driver's probe interface.
`Csr2`	

Field	Definition and Restriction
	Specify an optional VMEbus device register or memory address to map into CPU I/O space. A nonzero value in this field and the device driver's driver structure elements `addr2_size` and `addr2_atype` will cause the autoconfiguration software to map the specified VMEbus address space into system I/O space. The controller structure element `addr2` will receive the resultant `io_handle_t` address. This value is not passed to the device driver's probe interface, but is available in the controller structure.
`Vector`	Specify a VMEbus interrupt vector at which interrupts from the device interrupt the processor. This value is passed to the device driver in the `ivnum` element of the controller structure. The device driver's probe interface uses this value plus the VMEbus interrupt request level obtained in the next entry to add and enable VMEbus interrupt handlers via `handler_add` and `handler_enable` interfaces.
`Bus_Priority`	Specify a VMEbus interrupt request level at which the device will present its hardware interrupt request. This value is passed to the device driver in the `bus_priority` element of the controller structure. This interrupt request level and vector specified above is used by device driver's probe interface to add and enable VMEbus interrupt handlers via `handler_add` and `handler_enable` interfaces.
`Type`	Specify the character `C` for this entry. This indicates the entry is for a controller.
`Adpt_Config`	Specify the character `N`. This entry will be used in a subsequent release of Digital UNIX.

The VBA_Option entry below is included in `/etc/sysconfigtab` to provide VMEbus configuration information for the `vb` device driver. The `vb` device driver is the Digital UNIX supported VMEbus Backplane Network Driver.

```
VBA_Option = Manufact_Name - 'Digital',
Product_Name - 'VME Backplane Network Driver',
Bus_Instance - 0, Driver_Name - vb, Driver_Instance - 0,
Csr1 - 0, Csr2 - 0, Vector - 0x1150, Bus_Priority - 7,
Type - C, Adpt_Config - N
```

VBA_Option information is supplied on a contiguous line with no "newline" characters inserted.

The data following the dash (-) is the data for the specified field. The comma (,) separates the fields within `VBA_Option.` ASCII data fields are enclosed in single quotes (’).

The above entry indicates that the manufacturer is Digital and that the product is the VME Backplane Network Driver. The driver name is `vb` and installed on the first instance of `vba` as `vba0`. The controller is for the first instance of `vb` as `vb0.` The device driver requires no device register or memory mapping by the autoconfiguration software. This is indicated by zero being specified for `Csr1` and `Csr2`. The interrupt information to the driver indicates that the driver needs to add and enable an interrupt handler to vector 0x1150 with a VMEbus interrupt request level of 7.

- VMEbus Backplane Network Driver, `vb`.

  The Digital UNIX book Writing VMEbus Device Drivers indicates that the `vb` driver interrupt vector may need to be changed. This depends upon the module switch interrupt being selected. This was done in previous releases of Digital UNIX by specifying the information on the controller line in the system configuration file. This is no longer valid. In order to change the interrupt vector now, the `/etc/sysconfigtab` file must be edited and the `Vector` field of the `VBA_Option` entry for the `vb` driver must be changed to the appropriate vector.

  The configuration line in section D.5.5 of Writing VMEbus Device Drivers was previously specified as follows:

  ```
  controller vb0 at vba0 vector vbintr 0x1150 priority 7
  ```

  The configuration line must now be specified as follows:

  ```
  controller vb0 at vba0
  ```

- Fictitious VMEbus `/dev/dmaex` device driver

  The `/dev/dmaex` device driver included in Appendix B of the Digital UNIX V4.0 book Writing VMEbus Device Drivers has not been updated to support the Digital UNIX V4.0 single binary module for Static configuration.

  The following sections provide an updated version of the `/dev/dmaex` device driver that supports the Digital UNIX V4.0 single binary module for Static configuration. The following sections also provides an example files file and `sysconfigtab` fragment for the `/dev/dmaex` device driver.

  Refer to the Digital UNIX V4.0 book Writing Device Drivers: Tutorial for examples, discussions, and interfaces that support the single binary module and differences between Static and Dynamic device driver configurations.

– VMEbus /dev/dmaex files file fragment

The following files file fragment assumes the /dev/dmaex 3rd party device driver kit has been installed at /usr/opt/DM100.

```
MODULE/STATIC/dmaex           standard Binary
/usr/opt/DM100/dmaex.c        module dmaex
```

– VMEbus /dev/dmaex sysconfigtab files file fragment

```
dmaex:
Subsystem_Description = DMAEX device driver
Module_Config_Name = dmaex
Method_Type = Static
Module_Type = Static
Device_Char_Major = Any
Device_Char_Minor = 0
Device_Char_Files = dmaex
Device_Major_Req = Same
VBA_Option = Manufact_Name - 'Digital',
Product_Name - 'DMAEX Driver',
Bus_Instance - 0,
Driver_Name - dmaex, Driver_Instance - 0,
Csr1 - 0x00100000, Csr2 - 0x0,
Vector - 0x24, Bus_Priority - 1,
Type - C, Adpt_Config - N
```

VBA_Option information is supplied on a contiguous line with no newline characters inserted. Line break characters have been added to the above text for clarity.

The data following the dash (-) is the data for the specified field. The comma (,) separates the fields within VBA_Option. ASCII data fields are enclosed in single quotes (').

The above entry indicates that the manufacturer is Digital and that the product is the DMAEX Driver. The driver name is dmaex and installed on the first instance of vba as vba0. The controller is for the first instance of dmaex as dmaex0.

The device driver requires the autoconfiguration software to map eight (8) bytes of VMEbus address at address 0x00100000 in A24 User Data Mode address space into CPU I/O space. This is done by specifying a value in Csr1 above and specifying values for addr1_size and addr1_atype in the /dev/dmaex driver structure. The mapped I/O space address, io_handle_t, is passed to dmaexprobe and stored in the controller structure's addr element. The /dev/dmaex device driver uses this I/O mapped address to access its device control and status registers. A value of zero (0) was specified for Csr2 above. Therefore, no VMEbus to CPU I/O space mapping will be performed by the autoconfiguration software for Csr2.

The interrupt information to the driver indicates that the driver needs to add and enable interrupt handlers to vector 0x24 with a VMEbus interrupt request level of 1. The `/dev/dmaex` device driver installs two interrupt handlers at 0x24 and 0x25. The vector and bus_priority is passed to the device driver in the controller structure elements `ivnum` and `bus_priority`.

- VMEbus `/dev/dmaex` `dmaexreg.h` header file

  See `~DOCUMENTATION/TEXT/dmaex_sample.txt` on the *Digital UNIX V4.0 Documentation Volume 1* CD–ROM for an example of a driver.

7.5.2 Device Driver Tutorial

Section 14.1.1 of the tutorial states that you should create a directory to contain your driver source in the form of:

```
# mkdir  /usr/sys/io/ESA100
```

When you create a new directory to replace ESA100 you must place it in the path `/usr/sys/io/` using your directory selection to replace ESA100.

7.6 Network Administration Manual

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

7.6.1 Changes to Section 4.2.2.2, "Verifying PPP Support in the Kernel"

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

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

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

1. Log in as root.
2. Save the `/vmunix` file.
3. Rebuild the kernel by running the `doconfig` program and selecting the Point-to-Point (PPP) option.
4. Copy the new `vmunix` file to `/vmunix`.

5. Edit the `/etc/sysconfigtab` file and add the following lines:

   ```
   ppp:
   nppp=2
   ```

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

6. Reboot the system.

7.6.2 Sections 4.2.3.1, 4.2.3.2, and 4.2.3.3 are obsolete

Use the instructions n the following sections instead.

7.6.2.1 Establishing a PPP Dial-Out Connection

After you have connected your modem to a serial port on your system, do the following:

1. Verify that you can communicate with the modem. Do the following:

   a. Edit the `/etc/remote` file and copy the `kdebug` entry.

   b. Modify the new entry, providing a system name for the entry, the correct Digital UNIX device (`tty00` or `tty01` depending on your system), the correct baud rate, and correct parity. See `remote`(4) for more information.

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

      ```
      % tip system_name
      ```

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

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

      AT Return

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

2. Contact the remote system administrator or your Internet Service Provider (ISP) and obtain the following information:

   – Your remote IP address and netmask, unless the remote system assigns the IP address dynamically

   – Characters that might need to be escaped

   – Instructions on how to log in and use the remote service

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

3. Create a file for commands that the `chat` program uses to direct the modem what number to dial and what to send the remote system in order

to start pppd. This file is called a **chat script**. Each entry in a chat script has the following format:

string_chat_expects string_chat_sends

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

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

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

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

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

4 chat expects the Password: string and sends the mypassword string. The eq prevents chat from logging it when you use the -v option.

5 chat expects the $ (the shell prompt) and sends pppd to start the pppd daemon on the remote machine. The eq cancels the effect of the previous eq.

See chat(8) for more information on chat and chat scripts.

Note

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

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

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

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

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

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

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

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

6 If the local system does not receive a response from the remote system after 5 LCP echo request frames, `pppd` considers the link dead and tears down the connection.

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

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

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

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

5. Edit the `/etc/syslog.conf` file and do the following:

   a. Add the `local2` facility (used by `pppd` and `chat`) to the line that specifies `/dev/console` as the message destination as follows:

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

   In this example, the `notice` level is specified.

   b. Add the following entry to the file to create a `ppp-log` file:

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

   c. Save the edits and close the file.

6. Stop and restart `syslogd` by entering the following commands:

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

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

script file.

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

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

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

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

7.6.2.2 Establishing a PPP Dial-In Connection

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

1. Set up your modem for dial-in access. See Section 4.3.2 in *Network Administration* for more information.
2. Edit the `/etc/passwd` file and create a dedicated entry for a PPP user. For the login shell field, specify `/usr/sbin/startppp`; for example:

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

3. Edit the `/etc/inittab` file and create an entry for each terminal device that is to run PPP. For example:

   ```
   modem:3:respawn:/usr/sbin/getty /dev/tty00 M38400 vt100
   ```

   See `inittab`(4) for more information.
4. Issue the `init q` command to start the `getty` process immediately.
5. If the dial-in system is going to be a gateway for the dial-out system to reach other systems on the LAN, the dial-in system must be configured as IP router and must also run `gated`. Edit the `/etc/gated.conf` file and delete the `nobroadcast` option (if specified) in the `rip` statement. See Chapter 2 of *Network Administration* for basic network setup information and `gated.conf`(4) for `gated` options.
6. Edit the `/etc/ppp/options` file and include the following `pppd` options required to support dial-in access for all remote users:

   ```
   netmask dd.dd.dd.dd  1
   proxyarp  2
   crtscts  3
   asyncmap 0  4
   :remote_ip_address  5
   debug  6
   ```

   1 Sets the interface netmask to the specified value.

   2 Adds an entry to the local system's Address Resolution Protocol (ARP) table containing the IP address of the remote system and the

Ethernet address of the local system. This is not necessary if `gated` is running.

3 Enables hardware flow control for the serial port.

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

5 Specifies an IP address for the remote system.

If you want to specify options for each individual serial port, create a `/etc/ppp/options.tty`*xx* file and include the remote IP address and any other options that apply to that specific serial port. See `pppd`(8) for a complete list of `pppd` options.

7. After an incoming call is received and a connection established, `startppp` runs in the background. The process ID is logged in the `/etc/ppp/ppp`*xx*`.pid` file.
8. Enables debugging. All log messages are sent to the file specified in the `/etc/syslog.conf` file. After your connection is working correctly, remove this entry from the PPP options file.

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

7.7 System Administration Manual Instructions for SysMan Applications

The *System Administration* omits the following information on starting SysMan applications.

7.7.1 Accessing the SysMan Tools

In this release of Digital UNIX, the Common Desktop Environment (CDE) is the default desktop. The SysMan suite of graphical system management applications is the preferred interface for system administration. The SysMan applications are also available in the DECwindows and base X Windows graphical environments.

In CDE, the SysMan applications are available in the Application Manager. You can access the Application Manager from the CDE Front Panel by clicking on its icon. The SysMan applications are organized into five groups within the System_Admin group. You can double click on the System_Admin group to access the SysMan Configuration Checklist, the Welcome to SysMan online help volume, and the five application groups.

You can access online help for the SysMan applications without running the applications. Click on the Help Manager icon on the CDE Front Panel to display the online help browser. The browser includes help families for CDE, the CDE Desktop, and Digital System Management.

In DECwindows, the SysMan applications are listed in the Session Manager's Options menu. You can use the Applications Definitions menu item to add SysMan applications that are used frequently to the Applications menu.

In other X Windows environments, the SysMan applications can be invoked from the command line. See the `sysman_intro`(8X) reference page for a list of the SysMan applications. This reference page also describes how to display the online help browser in graphical environments other than CDE.

To support nongraphical environments, some of the SysMan applications offer command line and question and answer interfaces.

The following applications have a command line interface. A single command starts the application, which then performs the actions specified by the command line arguments.

- Network Configuration
- BIND Configuration
- NFS Configuration
- Mail Configuration
- Account Manager

The following applications have a question and answer interface invoked using the following command-line argument: `-ui menu`. The application prompts the user interactively.

- Network Configuration
- BIND Configuration
- NFS Configuration
- Printer Configuration

The menu interface for Mail Configuration is called `mailsetup`.

Finally, there are a variety of scripts invoked by commands such as SysMan Configuration Checklist and `/usr/sbin/setup`. They are documented in other sections of the System Administration guide.

7.8 Technical Overview

Note that the information on Maximum System Limits has moved from the *Technical Overview* to the *Release Notes*.

7.9 Assembly Language Programmer's Guide

The Assembly Language Programmer's Guide (Chapter 5) should include a description of the `.rconst` directive. This derective instructs the assembler to add subsequent data into the `.rconst` section. This behavior is similar to the `.rdata` directive, except that the entries cannot be relocatable.

7.10 DEC C Language Reference Manual

In the *DEC C Language Reference Manual*, the conventions table states that the DEC C extensions to the ANSI C standard are shown in teal in the printed manual, and are shaded in the online manual. This is incorrect.

7.11 Guide to Preparing Product Kits

The following notes apply to the *Guide to Preparing Product Kits*. manual.

7.11.1 Guide to Preparing Product Kits is now available on CD-ROM

The *Guide to Preparing Product Kits* is a new manual that describes the procedures for creating, managing, and installing layered product kits. You can view this guide on line with the Netscape browser. The book is part of the Supplementary Documentation bookshelf.

A compressed PostScript version of the book is also available for printing. You can find it on the documentation CD-ROM in the following location:

```
/usr/share/doclib/online/DOCUMENTATION/POSTSCRIPT/AA-QTLVA-TE.ps.Z
```

If you cannot find the file at this location, contact your system administrator.

The file named `00-READ-ME-FIRST` in the `POSTSCRIPT` directory describes how to view and print compressed files. It also lists the other PostScript files in the directory.

7.11.2 Description of setld

The following information used to be in *Programming Support Tools*, and was left out of the *Guide to Preparing Product Kits*.

7.11.2.1 setld Functions

This release note describes the steps the `setld` utility performs when executed with each of its options.

Note

The `setld` command's action is divided into phases. Some phases have `PRE_`*`phase`* and `POST_`*`phase`* subphases. If a given subset's `PRE_`*`phase`* subphase fails during any applicable operation, `setld` displays a message indicating that the subset control program has declined the operation and does not proceed further with that subset. No attempt is made to run the *`phase`* or `POST_`*`phase`* code.

7.11.2.2 Loading Software

When you load software by using the `-l` option to `setld`, the utility performs the following steps:

1. Verifies access to *`location`*.
2. Copies product installation information from *`location`* into a temporary area. The information copied is in the `INSTCTRL` file for each product kit to be installed. If the `setld` command line included specific subset identifiers, only those subsets are considered; otherwise, all subsets in *`location`* are considered.
3. Determines which subsets to load by calling each subset's subset control program with the `ACT` environment variable set to `M`. Subsets whose subset control programs determine that their respective subsets are candidates for installation are divided into mandatory and optional groups according to the subset control flags contained in the *`subset-id`*`.ctrl` files.
4. Displays a list of the candidate subsets, listing the mandatory subsets (if any) and offering the optional subsets in a menu for selection by the user. If there are no optional subsets, no menu is displayed; instead, the mandatory subsets are listed and the user is asked for permission to proceed.
5. Performs the following operations for each subset to be installed:
   a. Verifies that the subset will fit on the system.
   b. Invokes the subset's subset control program to perform product-specific tasks that must be done before the subset is loaded (`ACT` set

to `PRE_L`). A nonzero return status from the subset control program causes `setld` to abort the load operation.

   c. Creates a subset corrupt lock file.

   d. Loads the subset, using the subset's control and inventory files; then verifies the subset and upgrades the lock file to indicate that the subset is correctly installed.

6. Performs the following steps for each subset after loading all of the selected subsets:

   a. Invokes the subset's subset control program to perform product-specific tasks that must be done after the subset is loaded (`ACT` set to `POST_L`). The subset control program's actions at this time usually include dependency locking.

   b. Invokes the subset control program to perform product-specific tasks that must be done after the subset is installed (`ACT` set to `C` and `$1` set to `INSTALL`). This step is bypassed if the `-D` command option was invoked.

The installation control files (*subset-id*`.ctrl`, *subset-id*`.inv`, *subset-id*`.scp`, and *subset-id*`.lk` or *subset-id*`.dw`) are stored in the `&./usr/.smdb.` directory. The kit's subset archives are not stored because their contents have been placed in the appropriate locations. If you specified an alternative root path, this directory path is created under the directory you specify.

7.11.2.3 Configuring a Subset

When you load a product, the next-to-last stage of the `setld` process is to invoke the subset's subset control program with the `ACT` environment variable set to `C` and the command argument (`$1`) set to `INSTALL`.

When you issue a command to reconfigure a subset (the `-c` option), `setld` first verifies that the specified subset exists. If it does, `setld` sets the `ACT` environment variable to `C` and calls the subset's subset control program with *message* as a command argument (`$1`). Usually, the only valid messages are `INSTALL` and `DELETE`. These two messages are reserved in their meaning. For special needs, a particular subset control program could be designed to accept other messages. The `setld` utility cannot pass other messages except in response to its `-c` option.

7.11.2.4 Verifying a Subset

When you load a product, the final stage of the `setld` process is to invoke the subset's subset control program with the `ACT` environment variable set to `V`. This action instructs the subset control program to run its verification test.

When you issue a command to verify a subset (the `-v` option), `setld` first verifies that the specified subset exists. If it does, `setld` runs the subset's Installation Verification Procedure (IVP), if there is one.

7.11.2.5 Removing Software

When you issue a `setld -d` command, the `setld` utility performs the following steps for each subset to be deleted:

1. Verifies that the subset is installed.
2. Verifies that the subset's sticky bit, originally specified in the product's key file, is not set. If the sticky bit is set, `setld` declines to remove the subset.
3. Checks dependencies. If the subset's lock file (*`subset-id`*`.lk`) names any subsets that depend on the one to be removed, `setld` displays their names and requests confirmation that the subset should be deleted.
4. Invokes the subset's subset control program to perform product-specific tasks that must be done before any deletions are made. (`ACT` set to `C` and `$1` set to `DELETE`.)
5. Invokes the subset's subset control program to perform product-specific tasks that must be done before the subset is deleted. (`ACT` set to `PRE_D`.) If the subset control program returns nonzero status, `setld` aborts the deletion operation.
6. Deletes all files contained in the subset.
7. Invokes the subset's subset control program to perform product-specific tasks that must be done after the subset is deleted. (`ACT` set to `POST_D`.)
8. Marks the subset as being uninstalled by deleting its lock file.

7.12 Online Help Volumes.

The notes in this section refer to problems with the online help volumes.

The Help viewer has the following known problem:

- In some cases, the help viewer is not correctly initialized. As a result, it will sometimes exhibit the expected behavior the second time an action is taken, but not the first time.

  For example, the first time a quick help dialog box displays a reference page, the Backtrack button may be enabled even though there is no place to which to backtrack. If the dialog box is closed and then opened again, the Backtrack button is dimmed.
- Similarly, if a request for on-item help displays the correct help, but one line has scrolled off the top, the online help will typically be displayed in

exactly the right position when the on-item help request is repeated.

For each of the SysMan applications online help is available from the Help menu or from the Help button in the main window.

The online help contains:

- An overview of the capabilities of the application
- A set of tasks illustrating typical uses of the application
- A reference section documenting every window and dialog box in the application.

The Using Help item on the Help menu displays a help volume supplied by CDE that explains how to use the online help.

7.12.1 General Problems

The following known problems occur in the online help:

- The Appearance menu is not consistently documented.

  In the Archiver, License Manager, and System Information applications, there is an Appearance item on one of the menus. The item should offer three options: Text Only, Large Icon, and Small Icon. In some help volumes, not all of these are documented.

- Some links from one help volume to another are displayed in a new view.

  There are a few links from one help volume to another that display in a new help viewer window. In some situations, a distracting proliferation of help viewer windows can result.

  You can get a new view when you want one using the New Window item on the File menu in the help viewer.

7.12.2 Integration

The SyMman configuration applications on-item help does not work on the items in the menu bar.

In all the SysMan applications, the keyboard method of getting on-item help does not work on the Help menu.

7.12.3 Help Volumes by Application

The following problems apply only to help in specific applications:

- Kernel Tuner

  The Kernel Tuner application records its changes immediately, so if a system failure occurs while the Kernel Tuner is running, any boot time changes will take effect the next time the system boots. If the file

/etc/sysconfigtab contains invalid values, you can enter the following command at the boot prompt to boot using default values:

```
boot -fl c
```

- Display Window

  The help volume for Display Window has opening instructions that show how to access the application from the CDE Application Manager. These should show that the Display Window icon appears in two groups.

- Network Configuration

  The names of the gated, joind, routed, and rwhod daemons are misspelled in the online help volume.

  In the Configuring Interfaces dialog box, the fields under To Obtain IP Address are relevant for all interfaces.

- Printer Configuration

  In the Local Printer Settings dialog box and the Remote Printer Settings dialog box, the aliases in the Printer Aliases field must be separated by vertical bars because spaces are allowed in an alias.

- Shutdown

  The /usr/sbin/shutdown command now sends the Shutdown Message each time a reminder is sent of the time remaining before the system is shut down.

  The /etc/nologin file is not created until immediately before the shutdown occurs.

  System shutdown messages are sent to all users that are locally or remotely logged-into the system being shutdown. Additionally, if the Broadcast to NFS Clients" option is selected, shutdown messages are broadcast to all hosts that are NFS clients of the system being shutdown.

- Disk Configuration

  The following corrected definitions replace the definitions in the glossary for Disk Configuration.

  – Skew

    A deviation from a reference direction, either by design or in response to lateral forces.

  – Track skew

    On a disk, the sector skew per track. The skew is the angle that sector 0 of the track changes from an imaginary radius line, due to a nonuniform number of sectors per track.

- BIND Configuration

The following corrected definitions replace the definitions in the glossary for BIND Configuration:

- BIND client

  A system that queries a BIND server for host name and address information, interprets the responses, and passes the information to requesting applications.

- BIND server

  An authoritative source for information about one or more zones. It either maintains the master copy of the hosts database for the zone or obtains the information required to serve the hosts database from another server.

- DCE

  Distributed Computing Environment. The capabilities of DCE are defined by the Open Software Foundation (OSF).

- DCE cell

  A logical group of systems that share services offered by DCE.

- DCE server

  The server in a DCE cell.

- service type

  In BIND Configuration, the available service types are BIND client and BIND server. The service type determines whether a system is configured to be a BIND client or a BIND server.

7.13 Online PostScript Format Documentation

The Digital UNIX documentation set includes documents that are available only in PostScript format. The following list describes these documents and their locations:

- Digital UNIX Software Product Description

  The Software Product Description (SPD) is a legal description of the Digital UNIX product. It describes the software and gives information about its capabilities and about the hardware it supports. This information is intended for anyone who needs a legal description of the Digital UNIX product.

  The SPD is provided on the operating system distribution media in both PostScript and text versions. To obtain a copy of the PostScript version, mount the Digital UNIX V4.0 Operating System, Volume 1 CD–ROM. Then change to the `*mnt_point*/DOCUMENTATION/POSTSCRIPT` directory and print the following files:

```
Digital_UNIX_Developers_Toolkit_SPD.ps
Digital_UNIX_Logical_Storage_Manager_SPD.ps
Digital_UNIX_Operating_System_SPD.ps
Digital_UNIX_Server_Extensions_SPD.ps
Prestoserve_for_Digital_UNIX_SPD.ps
```

To obtain a copy of the text version, change to the *mnt_point*/DOCUMENTATION/TEXT directory and print the following files:

```
Digital_UNIX_C_Developers_Extensions_SPD.txt
Digital_UNIX_Logical_Storage_Manager_SPD.txt
Digital_UNIX_Operating_System_SPD.txt
Digital_UNIX_Server_Extensions_SPD.txt
Prestoserve_for_Digital_UNIX_SPD.txt
```

- *GNU Emacs Manual*

  This manual (developed by the Free Software Foundation) provides information about how to use and customize the Emacs text editor. It is primarily a reference manual, but can also be used as a tutorial.

  This manual is intended for general users and anyone who uses Emacs. You can obtain a copy of this manual by printing the following file:

  `/usr/lib/emacs/doc/emacs.ps`

  This file is available only if the `FSFEMACSSRC300` subset is installed on your system. (To use Emacs, install the `OSFEMACS300` subset.)

- *GNU Emacs Lisp Reference Manual*

  This manual (developed by the Free Software Foundation) describes Emacs Lisp and presumes considerable familiarity with how to use the Emacs text editor. The earlier chapters describe Emacs features that have counterparts in many other programming languages. The later chapters describe features that are peculiar to Emacs Lisp or relate specifically to editing.

  This manual is for programmers. You can obtain a copy of this manual by printing the following file:

  `/usr/lib/emacs/doc/elisp.ps`

  This file is available only if the `FSFEMACSSRC300` subset is installed on your system. (To use Emacs Lisp, install the `OSFEMACS300` subset.)

- Display PostScript Documentation Supplements

  The Display PostScript system is described in the manual *PostScript Language Reference Manual* which is available in printed format from Digital. To update the Display PostScript documentation, Adobe System, Inc. provides a number of supplements describing new and changed features of the Display PostScript system.

The Display PostScript documentation supplements are provided in compressed format in the following directory:

```
/usr/share/doclib/dps
```

Before you can print a copy of one of the supplements, you must uncompress that supplement. For example, to uncompress the *Level 2 Changes for X,* issue the following command:

```
% uncompress /usr/share/doclib/dps/\
Developer-TechNote.1-001.ps.Z
```

You can then print the resulting, uncompressed file:

```
/usr/share/doclib/dps/\
Developer-TechNote.1-001.ps
```

You might want to remove the PostScript file once you have printed it to save space on your system.

The following list describes the Display PostScript documentation supplements:

– *Level 2 Changes for X*

  This supplement describes the changes made to the X implementations of the Display PostScript system for PostScript Level 2. This supplement is provided in the following file:

  `Developer-TechNote.1-001.ps.Z`

– *Type 2 Image Dictionary*

  This supplement describes the Type 2 image dictionary, an operand for the `image` operator in the Display PostScript system. The Type 2 image dictionary is an extension to the Type 1 dictionary. It allows the `image` operator to use pixel data from a pixmap, the current window, or another window as source when copying into the current window. This supplement is provided in the following file:

  `Developer-TechNote.1-002.ps.Z`

– *Multiple Master Fonts in the DPS Toolkit Font Panel*

  This supplement describes new support for multiple master fonts in the Display PostScript Toolkit for X font panel. The supplement also describes how the font panel supports nontypographic sorting and a value-changed callback. This supplement is provided in the following file:

  `Developer-TechNote.1-003.ps`.

– *Writing Applications That Use the Resource Location Library*

  This supplement describes how applications can use resource paths to become easier to use and customize. Resource paths are included in calls to the resource location library, which applications can use to find resources such as fonts. This supplement is provided in the following file:

`Developer-TechNote.1-004.ps`.

– *PostScript(tm) Language Reference Manual Supplement*

The version 2015 supplement to the *PostScript(tm) Language Reference Manual* which contains all updates that manual (232pp).

`2015supplement.ps.Z`.

– *Adobe ShowPS Quick Reference (Card)*

This Quick Reference card describes the command syntax and lists the available X resources, command options and keyboard commands.

`ShowPSReferenceCard.ps`

– *Adobe ShowPS User Guide*

General instructions for installing, starting and using the Adobe ShowPS PostScript previewer. Contains tutorial material and a troubleshooting guide (72pp).

`ShowPSUserGuide.ps.Z`.

– *Adobe ShowPS Brochure*

A sample brochure in overhead format that gives an overview of the Adobe ShowPS PostScript previewer (5pp).

`ShowPS_Brochure.ps.Z`.

- X Image Exension documentation

The X Image Extension (XIE) code (developed by the X Consortium) provides a powerful mechanism for the transfer and display of virtually any image on X-capable hardware. Documentation for XIE is provided in compressed format in the following directory:

`/usr/share/doclib/xie`

Before you can print a copy of one of the XIE manuals, you must uncompress that manual. For example, to uncompress the *X Image Extension Overview*, issue the following command:

```
% gzip -d /usr/share/doclib/xie/xie_overview.ps.Z
```

You can then print the resulting, uncompressed file:

`/usr/share/doclib/xie/xie_overview.ps`.

You might want to remove the PostScript file once you have printed it to save space on your system.

The following list describes the manuals that describe the XIE code. A README file is also available in `/usr/share/doclib/xie`.

– *X Image Extension Overview*

  This manual provides general information about the X Image Extension (XIE) code. Topics covered include: XIE design goals, XIE historical summary, XIE architecture, element definitions, and subsetting.

  This manual is provided in the following file:

  `/usr/share/doclib/xie/overview.ps.gz`

– *XIElib Specification*

  This manual contains reference information about the XIElib functions, XIElib events, and XIElib errors. The Functions section covers the following types of functions: startup, LUT, photomap, ROI, photoflo, client data, abort and await, photoflo element, technique, and free.

  This manual is provided in the following file:

  `/usr/share/doclib/xie/xielib.ps.gz`

– *XIE Sample Implementation Architecture*

  This manual provides an architecture overview of XIE, including chapters on the following topics: extension initialization, memory management, request dispatching, data representation, data structures, protocol requests, DIXIE photoflo management, DDXIE photoflo management, and photo elements.

  This manual is provided in the following file:

  `/usr/share/doclib/xie/xieSIarch.ps.gz`

– *X Image Extension Protocol Reference Manual, Version 5.0*

  This manual specifies the X wire protocol for XIE. It defines the syntax, structure, and semantics of the XIE protocol elements. Topics covered include syntax specification, parameter types, resources, pipelined processing, import elements, process elements, export elements, events and errors, techniques, service class, and protocol encodings.

  This manual is provided in the following file:

  `/usr/share/doclib/xie/XIEProto.ps.gz`

• OSF/Motif Release Documentation

  OSF/Motif release notes and problems have been provided by the Open Software Foundation (OSF). These notes are contained in the following PostScript files:

  – `/usr/doc/motif/Motif_rel_notes.ps`

- A list of known problems in Release 1.2.3:
  `/usr/doc/motif/OPENBUGS`

  These files are available only if you have installed the OSFXDEV200 subset on your system.

- *STREAMS DLPI Paper*

  This paper provides the specification for a STREAMS Data Link Provider Interface (DLPI) It complies with DIS 8886 and Logical Link Control (LLC) DIS 8802/2

  This paper is provided uncompressed in the following file:

  `/usr/share/doclib/dlpi/dlpi.ps`.

Depending on the optional software subsets or environments installed on your system, a number of other documents may also be available in the `/usr/doc`

7.13.1 PostScript Manuals Replacing Online or Printed Volumes

In this release, several X Window and DECwindows manuals are shipped in PostScript format only. To access these documents, mount the DigitalUNIX V4.0 Documentation Volume 1CD–ROM and read the following file:

`<mount point>/DOCUMENTATION/POSTSCRIPT/00-READ-ME-FIRST`

This file provides a list of documents and instructions for decompressing the files. Depending on the installation at your site, the files may already be mounted at the following location:

`/usr/share/doclib/online/DOCUMENTATION/POSTSCRIPT .`

Features and Interfaces Scheduled for Retirement 8

This chapter lists features of Digital UNIX Version 4.0 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 Retired in This Release

This section lists features of Digital UNIX that have been retired in this release. These changes were announced in previous releases of DEC OSF/1.

- Support for ULTRIX RIS to Digital UNIX client functionality
- The `oawk` version of the `awk` command
- Duplicate `libc` and `libm` routines
- The `-n` option from `/usr/bin/echo` and `/bin/echo`
- Ethernet trailer encapsulation
- Linkworks run-time library
- Logical Volume Manager
- Obsolete POSIX real-time interfaces
- XIE V3.0 interface, server support

  Runtime support will still be provided transparently through the client
- The POLYCENTER Common Agent (extensions to the SNMP V1.0 agent).

8.2 Nonconforming Curses Library

Digital UNIX 4.0 is now shipping a new X/Open-Compliant Internationalized Curses library. This is not binary compatible with previous versions of the Digital UNIX Curses library, so compatible binaries (named `libcurses.a` and `libcurses.so`) are also shipped in separate directories. These compatible binaries will be retired on this schedule:

- In the next major release (5.0) they will be moved to an obsolete subset.
- In the following major release (6.0) they will be removed.

8.3 ACC

The C compiler for Digital UNIX will be officially replaced by DEC C for Digital UNIX in a future release of Digital UNIX. Both compilers were available in Digital UNIX Version 3.2 and both continue to be available in Digital UNIX Version 4.0. The licensing terms and conditions remain the same as they were in Digital UNIX Version 3.2 and earlier products.

In Digital UNIX Version 4.0, the newer DEC C is the default compiler invoked using the `cc` command. In releases prior to Digital UNIX (Version 3.2) it was invoked using the following option:

```
# cc -migrate
```

DEC C offers additional language and compiler features while also offering better, smaller, and faster executable files.

The older C compiler had been the default compiler for Digital UNIX Version 3.2 and earlier. It remains available in Digital UNIX Version 4.0 and is invoked using the following option:

```
# cc -oldc
```

Refer to the `cc`(1) reference page or the *Developers' Toolkit for Digital UNIX, Version 4.0 Software Product Description 44.36.11* for additional information.

8.4 Kernel Build Procedures

Several parts of the kernel interface and kernel build procedures will be retired in a future release of Digital UNIX. These interfaces and procedures are not compatible with new features being implemented in the Digital Unix Kernel.

Each item being retired has a new interface that can be migrated to. Consult the *Writing Device Drivers Tutorial* manual for specific information.

The following items are scheduled for future retirement:

- The `if_dynamic` keyword in `conf/files` and `conf/alpha/files` is being replaced with the new `MODULE` syntax.
- The use of `NotBinary` files. As part of the single-binary module format and dynamic load/topology requirements, `NotBinary` files are no longer recommended for use.

- The following kernel interfaces:

```
bdevsw_add,     bdevsw_del,     cdevsw_add,     cdevsw_del,
dualdevsw_add,  dualdevsw_del,  disable_option, enable_option,
kalloc,         kfree,          kget,           zinit,
zalloc,         zchange,        zfree,          zget
ldbl_ctlr_configure,    ldbl_ctlr_unconfigure,  ldbl_stanza_resolver,
```

- The kernel data structures `bdevsw`, `cdevsw`, and `zone`.
- The following stanza fields:

```
Subsystem_Description,
Module_Type,           Module_Path,           Module_Config,
Device_Block_Open,     Device_Block_Close,    Device_Block_Strategy,
Device_Block_Dump,     Device_Block_Psize,    Device_Block_Flags,
Device_Block_Ioctl,    Device_Block_Funnel,   Device_Char_Flags
Device_Char_Open,      Device_Char_Close,     Device_Char_Read,
Device_Char_Write,     Device_Char_Ioctl,     Device_Char_Stop,
Device_Char_Reset,     Device_Char_Ttys,      Device_Char_Select,
Device_Char_Mmap,      Device_Char_Funnel,    Device_Char_Segmap,
```

8.5 TURBOchannel Denali Graphics

The TURBOchannel Denali graphics subsystem (formerly marketed by Kubota Graphics Corporation and by Digital Equipment Corporation) is not supported by Digital UNIX Version 4.0. The KWS_TD device driver support, which was present in the Version 3.0 and later releases of DEC OSF/1 and Digital UNIX, is removed.

The TURBOchannel Denali graphics subsystem is supported using the generic character-cell style support, but will not operate with the X Window System software. No device dependent X (DDX) support compatible with the X11R6 based X Window System software is available from Kubota Graphics Corporation or its business successors.

Due to the limited performance of the generic character-cell style support provided, customers installing Digital UNIX Version 4.0 on a DEC 3000 family system with a TURBOchannel Denali graphics subsystem should consider installing a supported graphics subsystem (such as the Digital ZLX 2D or 3D options), or using a character cell terminal on systems that support operation in server mode. Refer to your system's hardware documentation for information on using a character cell terminal as the system console.

8.6 dbx Debugger

The `dbx` symbolic debugger will be retired in a future release of Digital UNIX. The `dbx` debugger will be replaced by the Digital `ladebug` debugger, which is a superset of the `dbx` functionality. The Digital ladebug debugger is command line compatible with dbx and also supports a graphical user interface.

Digital recommends that you begin using the Digital ladebug debugger now and report any problems. This will enable Digital to provide a higher quality replacement when dbx is finally retired.

8.7 Nemacs

Nemacs V3.3.2, a public domain Japanese implementation of emacs, will be retired in a future release of Digital UNIX. Mule, a public domain multilingual implementation of emacs is shipping in Digital UNIX Version 4.0 and will be carried forward as the replacement functionality for Nemacs. The Nemacs subsets IOSJPNEMACS400 and IOSJPNEMACSSRC400 will be removed from the system. For more information on Mule, refer to the mule(1) reference page.

8.8 Security Interfaces

Several interfaces in libsecurity will be retired in a future release of Digital UNIX. These interfaces are not capable of supporting new functionality offered in Digital UNIX 4.0.

Binary Compatibility will be preserved until the interfaces are retired which will be no sooner than 12 months after the current release is shipped. The complete list of documented interfaces follows. There may be some undocumented interfaces that will be retired at the same time.

```
getprtcent(), getprtcnam(), putprtcnam(), getprdfent(),
getprdfnam(), putprdfnam(), getprfient(), getprfinam(),
putprfinam(), getprlpent(), getprlpnam(), putprlpnam(),
getdvagent(), putprfinam(), getprlpent(), getprlpnam(),
putprlpnam(), getdvagent(), getfvagnam(), putdvagnam(),
getprpwent(), getprpwuid(), getprpwnam(), putprpwnam(),
getprlpent(), getprlpnam(), putprlpnam()

read_pw_fields(),    store_pw_fields(), read_tc_fields(),
store_tc_fields(),   time_lock(),       get_seed(),
auth_for_terminal(), locked_out().
```

The following list of associated data structures will be retired at the same time as the interfaces:

```
struct pr_field,   struct pr_flag,  struct t_field,
struct t_flag,     struct l_field,  struct l_flag,
struct dev_field,  struct dev_flag, struct pr_passwd,
struct pr_term,    struct pr_file,  struct pr_lp,
struct pr_default, struct dev_asg,
struct system_default_fields,
struct system_default_flags.
```

8.9 Security Windows-Based Administration Utilities

The functions previously performed with the `XIsso` and `XSysAdmin` programs have been moved to other graphical user interfaces (`dxaccounts`, `dxaudit`, and `dxdevices`) The `XIsso` and `XSysAdmin` programs in this release are only interfaces to the other GUIs and support for `XIsso` and `XSysAdmin` will be discontinued in a future release.

8.10 System Management Utilities

The following system management utilities will be retired in a future release of Digital UNIX.

These applications are being replaced by new more robust implementations under the new system management framework.

Old Application	Replacement Application
lprsetup	Printer Configuration
netsetup	Network Configuration
nfssetup	NFS Configuration
mailsetup	Mail Configuration
bindsetup	Bind Configuration
adduser	User and Group Account Manager
removeuser	User and Group Account Manager
XISSO, Xsysadmin	dxaccounts, dxaudit and dxdevices

8.11 BSD TTY-NAME

The intent to retire the BSD `TTY-NAME` namespace was announced in DEC OSF/1 Version 3.0. The retirement will not be implemented in Digital UNIX Version 4.0 but will be deferred to a later release.

8.12 NC and NL Interfaces in the Standard C Library

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 interfaces will be removed in a future version of Digital UNIX:

```
NCchrlen
NCcollate
NCcoluniq
NCdec
NCdechr
NCdecode
NCdecstr
NCenc
NCencode
NCencstr
NCflatchr
NLchrlen
NLctime
NLflatstr
NLflattab
NLfprintf
NLgetamsg
NLisNLcp
NLxcol
```

8.13 Desktop Applications

The introduction of the Common Desktop Environment (CDE) in Digital UNIX 4.0 includes new graphical desktop tools for:

- Session management
- Calculator
- Calender
- Text editing

The Motif X applications corresponding to the new CDE applications are dxsession and `dxpause`, `dxcalc`, `dxcalender`, and `dxnotepad`. These existing applications will be retired in a future release of Digital UNIX.

8.14 SCSI Device Names

Support for `rz` SCSI device names will be retired in a future release of Digital UNIX. Retirement is not expected to be implemented for at least two major releases. Any code that derives knowledge about a device from the ASCII name or minor number may be impacted.

All code that uses the current namespace will be compatible in Digital UNIX Version 4.0 because a mechanism that ensures binary compatability is provided. Existing interfaces such as names and minor numbers are fully supported.

No code changes are required in this release, unless programmers want to take advantage of the new extended address capability. Conversion to the new interfaces is strongly recommended.

8.15 POSIX 1003.4a (draft 4) pthread Routines in DECthreads

The POSIX 1003.4a, Draft 4 interface of DECthreads is being retired and will not work in a future release. Applications that were written using the POSIX 1003.4a, Draft 4 API should be migrated to the new IEEE Std 1003.1c-1995, POSIX System Application Program Interface provided by DECthreads. The POSIX 1003.1c standard interface is the most portable, efficient, and powerful programming interface offered by DECthreads. A compatibility mode for the draft 4 POSIX 1003.4a API has been provided in this release to help ease migration. This compatibility mode will be removed in a future release.

8.16 DECthreads CMA Interface

The CMA interface of DECthreads will be made obsolete in a future release. Obsolescence means that this API will always exist in Digital UNIX and will be supported, but will no longer be documented or enhanced. It is recommended that you port your CMA based application to the IEEE Std 1003-1c-1995, POSIX System Application Program Interface provided by DECthreads.

8.17 Functions Retired from libc

The following functions existed in both the `libm` and `libc` libraries until version Digital UNIX Version 3.2:

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

The `libc` versions of these functions have been retired. In this version, and in future versions of Digital UNIX the functions will only be available in `libm`.

8.18 AdvFS Quota Commands

The UFS and AdvFS user and group quota commands have been consolidated for Digital UNIX Version 4.0. The standard UFS quota commands can now be used to manage user and group quotas on AdvFS. AdvFS quota functions have not changed; AdvFS-specific options are now incorporated into the UFS commands.

Support for existing AdvFS versions of the quota commands will continue until a future release. Until then, both the UFS and AdvFS versions of the quota commands will be supported.

8.19 LVM-to-LSM Migration Tools

With the retirement of LVM in this release, the LVM-to-LSM migration tools will become redundant in later releases of Digital UNIX.

These migration tools were provided to enable migration from the retired LVM interfaces to Digital UNIX Logical Storage Manager volumes.

The LVM-to-LSM Migration Tool will be retired in a future release of the operating system. The UFS and AdvFS Migration Tools are not planned to be retired. 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.

8.20 OSF/Motif Version 1.1.3

The Motif Version 1.1.3 libraries have been provided as run-time services for compatibility with applications that had not yet converted to Motif 1.2. Development support was retired in DEC OSF/1 Version 2.0.

In Digital UNIX Version 4.0 the Motif 1.1.3 libraries have been moved to an optional subset. Applications requiring the libraries will see an error from the loader and you must install the optional subset. This optional subset will be removed from the product in a future release.

8.21 XIE Version 3.0 X Client Extension

Digital UNIX Version 4.0 supports XIE Version 5.0. Support for XIE v3.0 Server extensions has been removed in Digital UNIX Version 4.0, but Client support will not be removed until a later release of Digital UNIX. Digital UNIX Version 4.0 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.

8.22 ULTRIX Remote Installation (RIS) Support

Support for ULTRIX Remote Installation Services (RIS) to Digital UNIX client functionality has been retired in Digital UNIX Version 4.0.

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.

8.23 Disk Space Requirement for Installation

Due to additional functionality being planned, the disk space requirement for a custom installation 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.

8.24 MSB Driver

In Digital UNIX V4.0, the device driver for the base audio on the Digital AlphaStations and Digital AlphaServers will not be part of the base operating system. This device driver supports the Microsoft Sound Board, the AlphaStation Sound Card, and the built-in audio hardware shipped with certain AlphaStation systems.

Instead, the driver binaries will be available as part of the Multimedia Services for Digital UNIX kit available on the System Integrated Product CD–ROM in the `MMEDRVMSB201` subset.

The following files will be removed from the base operating system:

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

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

8.25 C Library Functions and POSIX P1003.1C

As of Digital UNIX V4.0 the following C library functions exist in two versions due to conflicts between previous versions of Digital UNIX and the recent IEEE POSIX P1003.1C standard (these new interfaces are in affect by default). The old interfaces are currently accessible by defining the C preprocessor symbol `_POSIX_C_SOURCE` to `199309L`.

```
asctime_r     getgrnam_r    getpwuid_r    localtime_r   readdir_r
ctime_r       getlogin_r    gmtime_r      rand_r        ttyname_r
getgrgid_r    getpwnam_r
```

Binary compatibility is maintained in Digital UNIX V4.0, however these routines will be retired in a future release of Digital UNIX, no sooner than April 1997. The obsolete versions should not be used in new designs. These routines formerly resided in libc_r.a and libc_r.so, but were merged into the standard C runtime library (see the Development Chapter of this document for additional information).

Maximum System Limits A

A.1 Maximum System Limits

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

- Backup Utility Limits
  - `cpio`

Files per archive:	No limit
Files per file system:	No limit
File size:	4 GB
File name size:	256 bytes

  - `dd`

Files per archive:	Not used
Files per file system:	Not used
File size:	4 GB
File name size:	Not used

– `dump`

Files per archive:	4 GF
Files per file system:	4 GF
File size:	4 GB
File name size:	No limit (part of the inode data)

– `tar`

Files per archive:	No limit
Files per file system:	No limit
File size:	4 GB
File name size:	256 bytes (with prefix)

- Device Addressing Limits
  - Device Access

    In this release there are two types of disk device access: raw or character and block or buffered.

    For raw or character access, the structure field `uio.uio_offset` describes the byte offset within the disk partition. In this release the `uio_offset` is an unsigned 64-bit value, allowing an offset up to 2^{64} or 18 Exabytes. This value is converted to a physical block/sector number that is the data transfer start position. The physical block/sector number is limited by the structure field `buf.b_blkno`.

    For block or buffered access, the structure field `buf.b_blkno` describes the block/sector offset within the disk partition and is a signed 32-bit value. Since this release supports a fixed 512-byte block/sector size defined by `DEV_BSIZE`, the offset is limited to 1 TB.

– Major-Minor Numbers (`dev_t`)

Devices are described by a major-minor pair of numbers, where the major number describes the device driver and the minor number describes the device. In this release these pairings are represented by a 32-bit value described by the type `dev_t`. The major number portion of `dev_t` consists of bits 20 to 31 (12 bits). Since each device driver requires 12 bits for its major number, 4096 device drivers may be configured into the system.

The minor number portion of `dev_t` consists of bits 0 to 19 (20 bits). The content of these bits is left up to the device driver. A device driver utilizing all 20 bits for device addressing could address up to 1048576 devices per major number. For device drivers supporting disk devices, some number of bits in the minor device number will be reserved for the partition number. This release requires disk drivers to reserve the lower 6 bits for device attributes and partition numbers. However, This release restricts support to 8 partitions.

– SCSI/CAM Addressing

Common Access Method (CAM) is an ANSI-proposed standard for a common software interface to Small Computer Systems Interface (SCSI). There are no restrictions or limitations within CAM for disk block addressing; the address is an incoming value.

For SCSI-2, the Command Descriptor Block (CDB) defines the starting disk block number for the transfer. In this release the 10-byte CDB has 4 bytes reserved for the disk block address. This is an unsigned 32-bit value allowing $2^{32} - 1$ or 4 Gigasectors of addressing, which corresponds to 2 TB given the 512-byte block/sector size.

In this release the SCSI/CAM driver can address a maximum of 64 buses, with up to 7 device targets per bus, and a maximum of 8 LUNs per device target. As a result, in this release, SCSI/CAM can address a maximum of 3584 devices.

– Redundant Array of Independent Disks

This release supports three Redundant Array of Independent Disks (RAID) controllers: two for the SCSI bus (HSZ10 and HSZ40) and one for the EISA or PCI bus (SWXCR). Each RAID device is seen by the operating system as a single target device (that is, as a single disk) with up to 8 Logical Unit Numbers (LUNs) on the SCSI controllers and 8 Logical Units (LUs) on the EISA or PCI controller, regardless of the number of disks on each RAID device.

The HSZ10 SCSI RAID controller supports a maximum of 35 back-end disks; the HSZ40 SCSI RAID controller, a maximum of 42 back-end disks. As a result of hardware constraints, the maximum number of HSZ10 disks that can be concatenated into a logical volume is 5; the maximum number of HSZ40 disks that can be concatenated into a RAID 0 set for a logical volume is 14 with a total size limit of 32 GB.

The EISA and PCI RAID controllers supports a maximum of 8 Logical units with a maximum of 8 drive groups. A Logical Unit refers to an amount of storage space presented to the host operating system as a single storage device. A drive group consists of one to eight physical drives that operate (defined and addressed) as a single unit. Logical volume sizes are not fixed sizes as compared to other disk devices. The size of a logical volume is configurable based on needs with a total size limit of 32 GB. In addition, the SWXCR controller may have either a one or a three channel SCSI adapter which supports 7 or 21 back-end SCSI disks, respectively. In addition, the SWXCR controller may have either one or three SCSI channels which support 7 or 21 back-end SCSI disks, respectively.

Although RAID theoretically increases the number of addressable disks significantly, Digital recommends that the maximum number of devices for each system – even with RAID configured – should not exceed the numbers listed in the following section on device limits per processor.

– Disklabel

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

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

    Note that DSA supports a maximum of 96 RA/TA devices.
  - Processors
    * Alpha Station 200 series – 21 RZ/TZ Devices

      (1 Baseboard PCI) + (2 PCI SCSI Adapters) = 3 SCSI Buses * 7 devices/bus
    * Alpha Station/Server 400 series – 28 RZ/TZ Devices

      (1 Baseboard PCI) + (3 PCI SCSI Adapters) = 4 SCSI Buses * 7 devices/bus
    * Alpha Station 600 series – 35 RZ/TZ Devices

      (1 Baseboard PCI) + (4 PCI SCSI Adapters) = 5 SCSI Buses * 7 devices/bus
    * Alpha Server 1000 series – 28 RZ/TZ Devices

      (1 Baseboard PCI) + (3 PCI SCSI Adapters) = 4 SCSI Buses * 7 devices/bus
    * AlphaServer 8200/8400 – 32 SCSI Buses * 7 devices/bus
    * DEC 2000 series
      + Model 300 – 28 RZ/TZ Devices

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

        (4 Single SCSI Adapters) = 4 SCSI buses * 7 devices/bus
    * DEC 2100 series – 35 RZ/TZ (SCSI) Devices –42 RE (RAID) Devices

      (1 PCI Baseboard Single SCSI) + (3 Single SCSI PCI Adapters) + (1 EISA Baseboard Single SCSI) + (2 Tri-SCSI Adapters) = 6 SCSI buses * 7 devices/bus

* DEC 2100 4/233, 4/275, 5/250, 5/300

  (8 PCI slots, each supporting 1 SCSI adapter) + (3 EISA slots, 1 SCSI internal)

* DEC 3000 series

  + Model 800/800S – 98 RZ/TZ Devices

    (1 Baseboard Dual SCSI) + (6 Dual SCSI TURBOchannel Adapters) = 14 SCSI buses * 7 devices/bus

  + Model 600S – 56 RZ/TZ Devices

    (1 Baseboard Dual SCSI) + (3 Dual SCSI TURBOchannel Adapters) = 8 SCSI buses * 7 devices/bus

  + Model 600 – 42 RZ/TZ Devices

    (1 Baseboard Dual SCSI) + (2 Dual SCSI TURBOchannel Adapters) = 6 SCSI buses * 7 devices/bus

  + Model 500 – 98 RZ/TZ Devices

    (1 Baseboard Dual SCSI) + (6 Dual SCSI TURBOchannel Adapters) = 14 SCSI buses * 7 devices/bus

  + Model 500X – 84 RZ/TZ Devices

    (1 Baseboard Dual SCSI) + (5 Dual SCSI TURBOchannel Adapters) = 12 SCSI buses * 7 devices/bus

  + Model 400S – 56 RZ/TZ Devices

    (1 Baseboard Dual SCSI) + (3 Dual SCSI TURBOchannel Adapters) = 8 SCSI buses * 7 devices/bus

  + Model 400 – 42 RZ/TZ Devices

    (1 Baseboard Dual SCSI) + (2 Dual SCSI TURBOchannel Adapters) = 6 SCSI buses * 7 devices/bus

  + Model 300/300X – 35 RZ/TZ Devices

    (1 Baseboard Single SCSI) + (2 Dual SCSI TURBOchannel Adapters) = 5 SCSI buses * 7 devices/bus

  + Model 300L/300LX – 7 RZ/TZ Devices

    (1 Baseboard Single SCSI) = 1 SCSI bus * 7 devices/bus

* DEC 4000 series

  28 RZ/TZ Devices

  (4 Baseboard Single SCSI) = 4 SCSI buses * 7 devices/bus

* DEC 7000/10000 series

  96 RA/TA Devices (DSA)

112 RZ/TZ Devices (SCSI)

(8 Dual SCSI XMI Adapters) = 32 SCSI buses * 7 devices/bus

- CPU Limits Per Processor

AlphaStation 200 series:	1
AlphaStation/Server 400 series:	1
AlphaStation/Server 600 series:	1
AlphaServer 1000 series:	1
AlphaServer 2000:	2
AlphaServer 2100:	4
AlphaServer 2100A:	4
AlphaServer 8200:	6
AlphaServer 8400:	12
DEC 2000 series:	1
DEC 2100 series:	4
DEC 3000 series:	1
DEC 4000 series:	2
DEC 7000/10000 series:	6

- File System Limits
  - Maximum Size
    * Logical Storage Manager

      In this release the Logical Storage Manager (LSM) supports a maximum of 768 disk groups and 256 disks either in a disk group or across the system.

      In LSM, the term **volume** describes a virtual disk representing an addressable range of disk blocks used by applications such as file systems or databases. This release supports a maximum of 512 GB of disk space in a disk group or on a system, with a maximum supported volume of 512 GB. The maximum number of supported LSM volumes is 4093 for all disk groups in a 4093 for all disk groups in a system: 4091 non-system volumes and 2 system (root/swap) volumes.

      In LSM, the term **plex** describes the physical disk or disks that contain a complete copy of a volume's data. So, for example, a mirrored volume would be made up of at least two plexes. In this release the maximum number of supported plexes per volume is 8 and the maximum number of supported plexes per system is 4093 (or 4091 if root and swap volumes are not used).

In LSM, the term **subdisk** describes a contiguous portion of a physical disk which can be striped or concatenated together to form a plex. A maximum of 4096 subdisks can be associated with one plex, and Digital supports 4096 subdisks per disk group or per system.

LSM object names (such as volumes, plexes, subdisks, disk groups), volume attribute names (such as user and group), and `dxlsm view` names are limited to 14 characters.

* Advanced File System

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

  Although the architectural maximum limit of domains is 2048, in this release the AdvFS supports a maximum fileset and file size of 512 GB, up to 100 active file domains per system, and a maximum of 256 volumes per domain.

  Since a single disk failure in a domain can make the entire domain inaccessible, Digital also recommends that you create no more than 8 volumes per domain.

  Note that while Digital UNIX supports an unlimited number of filesets per system, only 512 filesets can be mounted at one time.

  The number of files per fileset is 2^{32}, limited by the tag used to uniquely identify a file in a fileset.

  Note that over time the actual limit of files per fileset decreases, since a tag can only be used 4096 times due to a sequence number limit.

  Although AdvFS can support page sizes larger than 13 bits, the theoretical maximum AdvFS file and fileset size is 16 TB (2^{13} ^*^ 2^{31}) with a 13-bit page size and 31-bit page number.

* UNIX File System

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

* CD-ROM File System

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

* Sparse Files

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

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

* Network File System

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

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

  Digital supports the following maximum file sizes:

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

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

– Memory Mapped File Limit

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

– Mounts

  * Advanced File System

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

* UNIX File System

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

* CD-ROM File System

  this release supports a maximum of 512 CD-ROM File System (CDFS) mounts.

* Network File System

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

– Open files

  The maximum number of files a process can open is set to 4096 by default in the `OPEN_MAX_SYSTEM` variable in the file `/usr/sys/include/sys/param.h`. This number can be adjusted between 64 and 4096, either in individual programs by using the `setrlimit`(2) system call or on a system-wide basis by editing the file `/usr/sys/conf/param.c`, changing the `open_max_soft` and `open_max_hard` variables, and then relinking or rebuilding the kernel. Note that file descriptor entries in the per process file table are dynamically allocated after the initial 64 entries in the `utask` structure are used.

– File Locking Limits

  The Digital UNIX file record locking service allows applications to lock any number of bytes in a file in the range of 0 to $2^{63} - 1$ inclusive. File locking is supported by UFS, AdvFS, and both NFS Version 2 and Version 3. Note that since the NFS Version 2 protocol suite only allows ranges to be specified with 32-bit numbers, it supports a file locking range of 0 to $2^{31} - 1$ inclusive.

– Pathname Limits

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

• Installation Limits

  This section lists the disk space required for `root`, `/usr`, and `/var` when performing a default or a custom installation of this release. These are approximate numbers, as the values will depend on the system configuration and are based on the UFS files system.

– Default Installation (mandatory subsets)

`root:`
39.7 MB

`/usr:`
168.4 MB

– Custom Installation (mandatory and all optional subsets)

`root:`
49.1 MB

`/usr:`
333.7 MB

`/var:`
6.1 MB

- Memory Limits
  - Physical Memory

    The maximum supported memory is different for each individual processor, although the DEC 7000/10000/8200/8400 series a total of 14 GB of physical memory. The operating system, however, supports a total of approximately 4 GB of physical memory.

    For more information on supported memory, see the *Systems and Options Catalog*. and the SPD.
  - Virtual Memory
    * Per process

      The default virtual memory per process is 2 GB, although available swap space may in many cases be exhausted before this limit is reached.

      This value can be increased to a maximum of 8 TB by defining the `MAXVAS` variable in the system configuration file and relinking or rebuilding the kernel.
    * Page size

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

– Pseudoterminals (`ptys`)

  The maximum number of supported `ptys` is 8192.

– LAT Connections

  • The maximum number of incoming LAT connections is 4800

  • The maximum number of outgoing LAT connections is 4000

– IP Alias Addresses

  Digital UNIX allows the use of up to 5,000 IP alias addresses before system performance begins to degrade.

– Packetfilter limits

  The packetfilter pseudo-driver can support up to 255 simultaneous open filters (each filter is usually mapped to one instance of an application program). The packetfilter can support a maximum of 255 devices, where each device can be displayed using `ifconfig -a`.

– Network Transfer Rates

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

• Process Limits

  – Per system

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

    You can increase this value by either changing the value of `MAXUSERS` in the system configuration file or by adding the `maxproc` variable to the system configuration file and relinking or rebuilding the kernel.

    You might increase the value of `MAXUSERS` to allow more users to log in to your system or to allow applications that run as `root` to fork more processes than `NPROC` allows by default. However, increasing the number of processes per system reserves more system memory, so the upper limit of `NPROC` is dependent upon the system's total memory, the number of actual users on the system, and the requirements of your applications.

– Per user

  The number of processes that each user can fork is set to 64 by default through the `CHILDMAX` variable in the file `/usr/sys/include/sys/syslimits.h`. The number of processes per user can be varied by adding the `maxuprc` variable to the system configuration file, setting its value to some number of processes, and then relinking or rebuilding the kernel.

  You might increase this value if you had an application that needed more processes than `CHILDMAX` is set to by default. However, increasing the value of `maxuprc` reserves more system memory, so the upper limit of `maxuprc` is dependent upon the system's total memory, the number of actual users on the system, and the requirements of your applications.

Additional Notes B

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.

B.1 Demonstration Program dxmandel

The source code as provided in `/usr/examples/motif/dxmandel` does not produce a working version of the demonstration. This can be corrected if you change the first function call of `XtToolkitThreadInitialize()` to read `XInitThreads()`. Then remove the second call to `XtToolkitThreadInitialize()`.

B.2 Additional Bootable Tape Restrictions

Additional restrictions apply to the use of bootable tape, described in Section 4.5.2 has changed. The first paragraph should read as follows:

For this release, the following features are not supported:

- The LSM product.
- The AdvFS `addvol` utility
- 32MB systems using AdvFS file systems

In the section titled Disk Space Issues, there are two references to 120,000 (512 blocks) of disk space required in `/usr`. The number of blocks should be 156,000 (512-blocks).

Step 4 of the instructions in this section refers to mounting device `/dev/rzid`. This step should mention that device `rz1d` should have 75,000 (512-blocks) available.

Step 10 of the instructions refers to mounting device `/dev/rzib`. This step should mention that device `rz1b` should have 156,000 (512-blocks) available.

The rest of Section 4.5.2 still applies.

B.3 Maximum System Limits

The release notes appendix on Maximum System Limits has been updated. please refer only to the information in this printed version of the release notes, or in the versions supplied on the Digital UNIX Version 4.0 Operating System Volume 1 CD–ROM. These files are located in the `/DOCUMENTATION/POSTSCRIPT` and `/DOCUMENTATION/TEXT` directories.

B.4 Remote Installation Services (RIS) Version Requirement

There are incompatibilities between utilities used by Digital UNIX Version 4.0 (formerly DEC OSF/1) and DEC OSF/1 V2.0. Therefore, it is not possible for a OSF/1 V2.0 server to serve Digital UNIX V4.0 clients. You must upgrade the server operating system software to a minimum revision of DEC OSF/1 V3.0 to be able to serve Digital UNIX V4.0 clients.

B.5 Upsolete Files After an Update Installation

If the following obsolete files exist on your system after a V4.0 update installation, they can be removed. These files will not appear in the `/var/adm/smlogs/upd_obsolete_files` log file and must be removed manually.

- `/isl/sas/.profile`
- `/isl/sifsync`

Index

A

E

F

G

H

I

M

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-QTLMA-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

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