

Release Notice:
DG/UX™ System 5.4 for
AViiON® Computers

Release 3.00

January 1994

Part number 085-600365-01

This release notice applies to the following models:

P001
Q001

Restrictions and Trademarks

This software is made available solely pursuant to the terms of a DGC license agreement which governs its use.

Restricted Rights Legend: Use, duplications, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at [DFARS] 252.227-7013 (October 1988).

DATA GENERAL CORPORATION
4400 Computer Drive
Westboro, Massachusetts 01580

ADVANCED PICK is a U.S. registered trademark of Pick Systems, Inc.
AViiON is a U.S. registered trademark of Data General Corporation.
CLARiiON and **DG/UX** are trademarks of Data General Corporation.
Hayes Smart-modem is a trademark of Hayes Microcomputer Products.
INGRES is a trademark of INGRES Corporation.
LaserJet is a U.S. registered trademark of Hewlett-Packard Co.
Legato NetWorker is a trademark of Legato Systems, Inc.
LOOKING GLASS is a U.S. registered trademark of Visix Software, Inc.
NFS is a U.S. registered trademark of Sun Microsystems, Inc.
ONC and **SunOS** are trademarks of Sun Microsystems, Inc.
OpStar is a trademark of Data General Corporation.
OSF/Motif is a trademark of the Open Software Foundation, Inc.
Sniffer is a trademark of Network General Corporation.
UNIX is a U.S. registered trademark of UNIX System Laboratories (USL), Inc.
X Window System is a trademark of Massachusetts Institute of Technology.
X.desktop is a trademark of IXI Limited.

Release Notice:

DG/UX System 5.4 for AViiON® Computers Release 3.00

085-600365-01

Revision History:

Effective with:

Original Release – January 1994 DG/UX System 5.4 Release 3.00

Contents

1	Introduction	7
2	Product Description	7
2.1	Hardware Support	8
2.2	New Software Features	8
3	Environment	10
3.1	User Default License	10
3.2	Hardware	11
3.3	Software	11
3.4	Using Layered Software Products	11
3.4.1	OpStar	11
3.4.2	ADVANCED PICK	12
3.4.3	INGRES	12
3.4.4	LOOKING GLASS Replaced with X.desktop	12
4	Enhancements and Changes	12
4.1	New Hardware	12
4.1.1	AViiON 550	12
4.1.2	AViiON 5500	13
4.1.3	AViiON 8500	13
4.1.4	AViiON 9500	13
4.1.5	Multiple VME Chassis support	13
4.2	New Software Features	13
4.2.1	High Availability (HA)	13
4.2.1.1	Multi-Path LAN (MPL) I/O	14
4.2.1.2	IP Takeover	14
4.2.1.3	NFS Failover	15
4.2.1.4	Watchdog Timer (wdt)	15
4.2.1.5	On-Line Storage Management (OSM)	15
4.2.2	Kernel	19
4.2.2.1	Installer Kernels	20
4.2.2.2	Kernel Source Code Interface Changes	20
4.2.2.3	User scsi Device Driver	21
4.2.2.4	DOS and ISO-9660 CD-ROM File System Drivers	21
4.2.2.5	Enhanced dg_process_info	21
4.2.2.6	Real-time Extension Support	21
4.2.3	TCP/IP	23
4.2.3.1	Simple Network Management Protocol (SNMP)	23
4.2.3.2	gated	24
4.2.3.3	Diagnostics tools	25
4.2.3.4	Network Time Protocol (NTP)	27
4.2.3.5	Boot Protocol (bootp)	27
4.2.3.6	Serial Line IP (SLIP)	27
4.2.3.7	TCP/IP Protocol, Command, and Library Upgrades	27
4.2.3.8	Library Change for TCP/IP Programmers	29
4.2.3.9	sysadm Changes for TCP/IP	29
4.2.4	POSIX Threads (Pthreads)	30
4.2.4.1	Medium Term Scheduler (MTS)	30

4.2.4.2	Reentrant User Libraries	31
4.2.4.3	WARNING: Thread Cancellation	31
4.2.5	Debuggers	32
4.2.5.1	mxdb and dbx	32
4.2.5.2	sdb	32
4.2.5.3	gcore	32
4.2.5.4	Software Status Register (SWSR)	33
4.2.6	Compilers	33
4.2.6.1	Short Addressing	33
4.2.7	Libraries	33
4.2.7.1	WARNING: libnsl	33
4.2.8	Commands	33
4.2.8.1	syslogd	34
4.2.8.2	sysadm	34
4.2.8.3	New sar and ps commands	34
4.2.8.4	pread, pwrite	35
4.2.8.5	ls -L	35
4.2.8.6	login	35
4.2.8.7	man	36
4.2.8.8	Perl	36
4.2.8.9	ttymon	36
4.2.8.10	Archive Utilities	36
4.2.9	Standards	38
4.2.9.1	X/Open XPG3 Branding	38
4.2.10	Printer Services	38
4.2.10.1	Filter Enhancements	38
4.2.10.2	LP Enhancements	39
4.2.10.3	pcnfsd	39
4.3	Software Installation	40
4.3.1	Capacity Checker	40
4.3.2	Cleanup Procedures	41
4.3.3	Remote File System Loading	41
5	Notes and Warnings	41
5.1	Notes	41
5.1.1	Copying a CD-ROM Virtual Disk	41
5.1.2	Caching a CD-ROM Virtual Disk to a Magnetic Disk	41
5.1.2.1	On-Line Non-Persistent Caching	41
5.1.2.2	Off-Line Persistent Caching	42
5.1.3	Sizing Swap Area Needs for Debugging	43
5.1.4	sac	43
5.1.5	Sync Jumpering	43
5.1.5.1	AV530 and AV4600 Instructions	43
5.1.5.2	AV550 and AV5500 Instructions	44
5.1.6	Installing Packages in a Secondary Release Area	44
5.2	Warnings	45
5.2.1	Stand-Alone sysadm	45
5.2.2	sysadm Operations	45
5.2.3	Restrictions on vfork Child Process	46
5.2.4	Restrictions on Mounting /usr from a Server	46
5.2.5	Run newaliases	46
5.2.6	Streams Resource Configuration	47
6	Documentation	47

6.1	Printed Documentation	47
6.1.1	System Administration Manuals	47
6.1.2	Programming Manuals	48
6.1.3	New Manuals for DG/UX System 5.4 Release 3.00	48
6.2	On-line Documentation	48
6.2.1	New Manual Pages	48
6.2.2	Deleted Manual Pages	50
6.3	Manual Page Compression Format	50
6.4	Related Release Notices	50
7	Software Distribution	51
7.1	Media	51
7.2	Files	51
8	Installation Instructions	51
8.1	Keeping Your Previous Customizations when Upgrading	52
8.2	User Default License	52
9	Preparing a Software Trouble Report (STR)	52
9.1	Describing Problems	53
9.1.1	System Halts or Hangs	54
9.1.2	Incorrect Behavior by a DG/UX System Call or Device Driver	54
9.1.3	Incorrect Behavior by a DG/UX Command	54
9.1.4	Fast-Recovery File System	54
9.2	Preparing STR Tapes	55
9.2.1	The System Dump Dialog	55
9.2.1.1	Dump to Tape	56
9.2.1.2	Dump to Virtual Disk	56
9.2.1.3	Dump to a Network Device	56
9.3	STR Tape Examples	57
9.3.1	Copy a Panic Dump to Tape	57
9.3.2	Position to the End of a Dump Tape	57
9.3.3	Copy Your Tailored System File	58
9.3.4	Copy Additional Diagnostic Files	58
9.3.5	Preparing an STR Tape for a Diskless Workstation	58
9.3.6	Preparing an STR Tape Without a Memory Dump	59

1 Introduction

This release notice describes the DG/UX™ System 5.4 Release 3.00 operating system for Data General's family of AViiON® computers.

Data General OpStar™ 1.60 and previous OpStar releases numbered 1.XX will not run on DG/UX System 5.4 Release 3.00. If you are currently using any OpStar software release 1.XX, you will not be able to upgrade to Release 3.00 at this time. Please contact your Data General representative regarding OpStar software availability that supports DG/UX System 5.4 Release 3.00.

This release combines support for both M88100- and M88110-based systems. The memory requirement of the minimum equipment configuration (MEC) has increased from 12MB to 16MB to accommodate new features.

This notice also applies to the TCP/IP, ONC™, and NFS® packages, which are on the same release medium but are separately installable. There is a separate release notice for the GNU C compiler. See it for information about compilation in DG/UX System 5.4 Release 3.00.

This release is available in a base model and a workstation model. This release notice, which accompanies both, describes the base model. The workstation model includes the features described here, plus additional features appropriate to workstations and workstation servers. If you are using the workstation model, see also the release notices for the X Window System™ and X.desktop™ for DG/UX System 5.4 Release 3.00. These two notices describe graphics and windowing topics, including OSF/Motif™, for this release.

An on-line version of this release notice, suitable for line printers, is in the file `/usr/release/dgux_5.4R3.00.rm`. Should you find discrepancies between the typeset and on-line versions, the typeset version takes precedence.

2 Product Description

This is a full release of Data General's implementation of UNIX® System V Release 4.

For discussions of major DG/UX enhancements and important topics in open systems computing, obtain the following technical briefs from your Data General representative:

<i>The DG/UX File System</i>	(012-004054-01)
<i>A Look at High Availability Disk Systems</i>	(012-004035-01)
<i>Open Systems Terminology</i>	(012-004093-01)
<i>Operator Initiated Failover in the DG/UX 5.4.2 Operating System</i>	(012-004186-01)
<i>Taking Advantage of Symmetric Multiprocessor Systems</i>	(012-004301-00)
<i>Machine-Initiated Failover in the DG/UX 5.4.2 Operating System</i>	(012-004245-00)
<i>POSIX Real-time Extensions in the DG/UX 5.4R2.01 Operating System</i>	(012-004244-00)
<i>8-Way Symmetric Multiprocessor Systems</i>	(012-004246-00)
<i>Localization in the DG/UX 5.4R2.01 Operating System</i>	(012-004302-00)
<i>Support for Threads in the DG/UX 5.4 R3.00 Operating System</i>	(012-004405-00)
<i>Managing Virtual Disks with On-Line Storage Management (OSM)</i>	(012-004406-00)

New or expanded features in this release are described in the following two sections. Refer to "Enhancements and Changes" for more detailed information.

2.1 Hardware Support

Support has been added for the following:

- AViiON AV550 workstation
- AViiON AV5500 server
- AViiON AV8500 server (with one megabyte Secondary Cache)
- AViiON AV9500 server (with one megabyte Secondary Cache)
- Multiple VME Chassis for AViiON AV9500

2.2 New Software Features

New software features, as well as improvements to existing software features, have been made in the following areas:

- This release contains high-availability (HA) enhancements and changes as follows:
 - Multi-Path LAN (MPL) I/O, which is the ability to switch from one LAN interface to another in the event of an interface failure.
 - IP Takeover, which is the ability of two hosts in a dual-initiator configuration to coordinate the management of one or more IP addresses that will be used to provide high-availability IP services to their clients.
 - NFS Failover, which lets NFS clients of the failed system access remote mounted file systems through the backup system.
 - Watchdog Timer (wdt), which detects hardware and software system hangs and triggers a reset.
 - On-Line Storage Management (OSM), which lets you manipulate your physical disk partitions without having to take your data offline by unmounting file systems. Virtual Disk Management (VDM), which is part of OSM, gives you a consistent way to manage all disk resources.
- This release contains kernel enhancements and changes as follows:
 - Installer kernels were subsumed into the stand-alone **sysadm**. Some internal kernel source code interfaces are changed in this release.
 - The User scsi Device Driver was added, which lets users developing user space drivers that issue SCSI command descriptor blocks through an **ioctl** interface.

- DOS and ISO-9660 CD-ROM file system drivers were changed so that the pseudo devices **hfm** and **dfm** are no longer required to configure them.
- A new version (**DG_PROCESS_INFO_VERSION_3**) was added to the **dg_process_info** system call.
- Real-time extension support was updated from the previous DG/UX release.
- This release contains TCP/IP enhancements and changes as follows:
 - Simple Network Management Protocol (SNMP), based on code from SNMP Research Inc., was enhanced.
 - The following diagnostic tools were added: Berkeley Packet Filter (BPF) kernel interface, **tcpdump**, **nfc**, and **traceroute**.
 - Network Time Protocol (NTP), an Internet standard protocol for synchronizing time-of-day clocks on machines in a network, was added.
 - **bootp**, a commonly-used protocol for booting devices via the network, was added.
 - Serial Line IP (SLIP), which provides a simple long-distance solution for running TCP/IP between a pair of machines, was added.
 - TCP/IP protocols, commands, and libraries were upgraded.
 - TCP/IP **sysadm** interfaces were changed.
- This release contains support for POSIX Threads (Pthreads) as follows:
 - The configuration parameters, **MINBOUND** and **MAXBOUND**, were added.
 - Reentrant user libraries were added.
- This release contains support for debuggers as follows:
 - Support for debugging multi-threaded processes was added to **dbx** and the value-added product, **mxdb**; **sdb** does not have any special support for debugging multi-threaded processes.
 - The new **gcore** utility was added, which produces a coredump for a process without affecting the process's execution.
- This release contains compiler enhancements as follows:
 - **Short addressing**, an optimization that can reduce the number of MC88000 machine language instructions required to access a datum or to form its address, was added.
- The release contains the following new or enhanced commands:
 - **syslogd**, which controls the level of messages written to the */var/adm/message* file, was added.

- **sysadm** was changed to let you stop any port monitors after you have stopped any other port monitors.
- New **sar** and **ps** commands, which gather additional performance data, were added.
- **pread** and **pwrite**, which allow for an atomic seek and transfer operation to a file, were added.
- The behavior of **ls** under the **-L** option was changed to conform to the documentation.
- **login** was enhanced with a new configurable variable, **UPPEROLD**.
- **man** was enhanced to handle **.Z** files (files compressed by the **compress** program).
- **Perl**, a Practical Extraction and Report Language, is included in this release.
- **ttymon** was changed to once again log certain *Interrupted system call* messages.
- The archive utilities (**cpio**, **dd**, **dump2**, **tar**, **REELexchange**, and **systemtape**) were changed to report extended **errno** text information (if available) in addition to other information.
- This release contains the following printer services enhancements:
 - Filter enhancements were added.
 - LP enhancements were added.
- This release contains software installation enhancements and changes as follows:
 - Stand-alone **sysadm** replaced stand-alone **diskman** as the software installation utility in order to give the entire installation procedure a consistent interface.
 - A capacity checker was added to ensure sufficient space for a successful package load.
 - Cleanup procedures were added to assist the installer in removing obsolete files.
 - Support was added for installing or upgrading from a remote file system.

3 Environment

3.1 User Default License

All users must upgrade their user licenses with the new license token they receive with their DG/UX 5.4 Release 3.00 software, since the default user license is two users. See Section 8.2 of this notice for more information.

3.2 Hardware

This release can be installed from scratch on a new system or overloaded on the following versions of DG/UX:

- DG/UX Release 5.4.2
- DG/UX 5.4 Release 2.01
- DG/UX 5.4 Release 2.10

Note that the memory component of the recommended minimum equipment configuration (MEC) has increased from 12MB to 16MB to accommodate new features. **Do not attempt to install this release on an AViiON system that has less than 16 MB of memory.**

The disk component of the recommended minimum equipment configuration (MEC) remains the same (322 MB).

3.3 Software

Chapter 2 of *Installing the DG/UX™ System* (093-701087-05) contains detailed disk space storage requirements. Consult these tables before beginning installation.

This release includes fixes that were previously released as software patches. If any problem that was previously fixed by a patch reappears in this release, contact Data General. **Do not attempt to install any patches for previous releases on this release.**

This release also includes support for software and hardware that were previously released as supplements. **Do not attempt to install any supplements for previous releases on this release.**

DG/UX System 5.4 Release 3.00 is intended for use on all Data General AViiON computers.

3.4 Using Layered Software Products

Before attempting to use a layered software product with DG/UX System 5.4 Release 3.00, ask your Data General representative or the vendor about the product. Some layered products may require a particular DG/UX release.

3.4.1 OpStar

Data General OpStar 1.60 and previous OpStar releases numbered 1.XX will not run on DG/UX System 5.4 Release 3.00. If you are currently using any OpStar software release 1.XX, you will not be able to upgrade to Release 3.00 at this time. Please contact your Data General representative regarding OpStar software availability that supports DG/UX 5.4 Release 3.00.

3.4.2 **ADVANCED PICK**

To install **ADVANCED PICK®** on DG/UX System 5.4 Release 3.00, you must have **ADVANCED PICK** version 6.0.1 or higher. DG/UX 5.4 Release 3.00 does not support pre-6.0.1 versions of **ADVANCED PICK**.

3.4.3 **INGRES**

To install **INGRES™** on DG/UX System 5.4 Release 3.00, you must have **INGRES** version 6.4/04 or higher. DG/UX 5.4 Release 3.00 does not support pre-6.4/04 versions of **INGRES**.

3.4.4 **LOOKING GLASS Replaced with X.desktop**

LOOKING GLASS® was replaced with **X.desktop**. See the **X.desktop** and **X11** release notices for more information.

4 **Enhancements and Changes**

This section contains information about enhancements and changes made in this release.

4.1 **New Hardware**

This release provides support for the following new hardware discussed in this section:

- **AViiON AV550 workstation**
- **AViiON AV5500 server**
- **AViiON AV8500 server (with one megabyte Secondary Cache)**
- **AViiON AV9500 server (with one megabyte Secondary Cache)**
- **Multiple VME Chassis support**

This release combines support for both 88100 and 88110 systems.

4.1.1 **AViiON 550**

The **AViiON Model 550** is a midrange deskside workstation designed to meet the requirements of two-dimensional applications in areas such as GIS, Financial Services, and information management and retrieval. Based on the highly integrated Motorola 88110 RISC CPU, the **AV 550** offers 8-bit color graphics capabilities, with an option for 24-bit, in a compact, deskside package. Ethernet, SCSI-2, synchronous, and three asynchronous ports are all integrated on the system board. The **AV 550** also comes standard with one 3.5-inch disk and 16MB of memory together with a keyboard, mouse, and 17-inch monitor. A 19-inch monitor is available as an option. You can add up to 256MB of additional memory as well as up to three additional 3.5-inch

devices and a single 5.25-inch device (including QIC tape or CD-ROM for media distribution and load).

4.1.2 AViiON 5500

The AViiON 5500 server is a compact desktide package ideally suited for user populations from 20 to 100. The AV5500 is based on the Motorola 88110 RISC CPU and support single and dual CPU configurations. It provides 16KB of on-chip cache and 256KB secondary cache. The system board has three memory banks that take four SIMMs each, which allow for a fully-loaded system of 256MB.

4.1.3 AViiON 8500

The AViiON 8500 server is a high-end server also based on the Motorola 88110 RISC CPU. AV8500 systems provide dual and quad processors in a compact office package ideally suited to organizations with 75 to 300 users. Users can add large amounts of memory (2 GB) and disk (more than one terabyte) for application growth.

The AV8500 server allows automatic recovery from system hangs to maximize system availability.

4.1.4 AViiON 9500

The AV9500 rack-mount server offers very large applications even more processing power and system capacity than does the AV8500.

Both the AV8500 and AV9500 servers allow automatic recovery from system hangs to maximize your system availability.

4.1.5 Multiple VME Chassis support

Multiple external VME chassis support is possible on all rack-mount (but not desktide) implementations of the AV9500 platform. These new hardware systems offer two independent VME chassis. This release lets users plug controllers into either chassis.

4.2 New Software Features

This section describes new software features of DG/UX System 5.4 Release 3.00.

4.2.1 High Availability (HA)

This release adds to the DG/UX system's industry-leading commercial strengths by providing the following new high-availability features discussed in this section:

- Multi-Path LAN (MPL) I/O

- IP Takeover
- NFS Failover
- Watchdog Timer (wdt)
- On-line Storage Management (OSM)

For more information of these high-availability features, refer to *Managing the DG/UX System* (093-701088-04) and *Achieving High Availability on AViiON Systems* (093-701133-00).

4.2.1.1 Multi-Path LAN (MPL) I/O

Multi-Path LAN (MPL) I/O is the ability to switch from one LAN interface to another in the event of an interface failure. The primary motivation for this feature is increased availability. Ethernet and FDDI interfaces are supported. When moving from one interface to a backup, it is assumed that both interfaces are of the same type and are attached to the same Ethernet segment or FDDI ring. Both LAN interfaces must be in the same machine, although they may be in different VME chassis.

If a system is configured with two LAN controller cards connected to the same LAN cable, DG/UX automatically switches from the primary LAN controller to the backup LAN controller card if the primary LAN controller fails. The backup LAN controller must be an idle card and cannot be used for other operations while it is backing up the primary LAN controller.

Note that Multi-Path LAN I/O is used only within a single system, it cannot be used to fail from one system over to another system.

4.2.1.2 IP Takeover

IP Takeover is the ability of two hosts in a dual-initiator configuration to coordinate the management of one or more IP addresses that will be used to provide high-availability IP services to their clients. It is a useful way to capitalize on Machine-Initiated Failover (MIF). A primary-secondary host pair that shares a dual-ported disk and is properly configured for MIF allows the secondary to assume network duties, such as access to NFS-mounted file systems, for the primary in case the primary fails.

IP Takeover can assign Internet Protocol (IP) addresses to a set of disks instead of a server. One machine in a configuration has control of this "floating" IP address. The address can be transferred between systems. IP Takeover allows NFS mount points and remote logging to access disks regardless of which system controls the disk.

Note that IP Takeover is used only between two different systems on a network; it is not used within a single system.

The difference between Multi-Path LAN I/O and IP Takeover is that MPL works over controllers that reside in the same system, while IP Address Takeover works with more than one system. Use these two high-availability features together to provide redundancy both within a machine and between machines.

Note that IP Takeover is not supported for VTC controllers.

4.2.1.3 NFS Failover

A primary application of IP Takeover is NFS Failover. Once a backup system has taken over a failed server's floating IP address, the backup system can also failover the NFS services provided by that server. With the failover of NFS services, NFS clients of the failed system can access remote mounted file systems through the backup system. Note that these file systems must be on disks that also fail over to the backup system.

NFS services are unavailable only while the failover is taking place. Once NFS services have failed over, processing can continue with no further interruption. Remote applications using NFS services do not need to be restarted.

4.2.1.4 Watchdog Timer (wdt)

New features for the watchdog timer (wdt) were added to DG/UX System 5.4 Release 3.00. This release not only adds the AV550 to the list of systems with support for the hardware watchdog timer, but also adds for all AViiON platforms a watchdog timer feature that is supported by the **failovermon** monitor. See the **failovermon(1M)** manual page for more information.

The watchdog timer was first included in DG/UX 5.4 System Release 2.10 for the AV8500, AV9500 and AV5500 systems. When the watchdog timer detects either a hardware hang or a software system hang, it triggers a reset. When the system is reset on the AV8500 and AV9500 platforms, powerup diagnostics are run and any faulty field-replaceable hardware components are isolated and deconfigured when the system is rebooted. If the system has AV/Alert enabled, the faulty hardware is reported to the appropriate Data General Customer Support Center.

Panic code 53000060 is associated with a hardware watchdog timer reset. This panic code identifies a hardware watchdog timer reset only. Since the current state of the job processors is not saved with a system reset, a dump of the memory image is not taken with the hardware watchdog timer.

DG/UX System 5.4 Release 3.00 has added the ability to use **failovermon** monitors with the watchdog timer driver. This new feature lets AViiON systems with and without the hardware watchdog timer detect and recover from system hangs. For all AViiON systems, the watchdog timer can be configured through **Sysadm =>Device=>Configure**.

Once the wdt is configured, the kernel must be rebuilt and rebooted. Once the kernel is rebooted, the administrator can set up a **failovermon** monitor process to communicate with the watchdog timer driver and force a system panic if a system hang is detected. This new watchdog timer feature generates panic code 53000061. A 53000061 panic generates a memory image that can be dumped to tape or disk for analysis.

4.2.1.5 On-Line Storage Management (OSM)

This release introduces the new concept of On-line Storage Management (OSM) for managing disk partitions. The terminology for discussing disks has changed in this release. For example, a virtual disk is OSM's counterpart to the LDM (Logical disk manager) logical disk. For a complete

discussion of this concept see *Managing the DG/UX System* (093-701088-04) and the technical brief, *Managing Virtual Disks with On-line Storage Management* (012-004406-00).

Upgrading to DG/UX 5.4 Release 3.00

Make sure all logical disk caches are deleted prior to upgrading. See *Installing the DG/UX System* (093-701087) for detailed information on deleting logical disk caches.

During the upgrade installation, you will be prompted to convert your physical disks from logical disk format (LDM) to virtual disk format (VDM). The conversion step reads the system areas that are in LDM format and rewrites them into VDM format.

Note that only the system areas are converted, not user data. Thirty-three disk blocks are being read, converted, and rewritten to the virtual disk information tables (VDIT). Both LDM and VDM formats maintain redundant copies of the physical disk formats. If during the conversion process either copy of the VDM VDIT fails to write correctly, the physical disk remains in LDM format.

If your physical disks were organized using mirrors or stripes, those configurations will also be preserved during the conversion process. OSM provides greater extensibility of virtual disk partitions, aggregations, mirrors, stripes, and caches as detailed in the DG/UX documentation. However, during an upgrade to DG/UX System 5.4 Release 3.00, it is not necessary that you explicitly take advantage of any of these new features. Data General recommends that you convert all your physical disks to VDM format. After your system is back into production, you may wish to take advantage of the new features.

If you find it necessary to convert your system from DG/UX System 5.4 Release 3.00 to a previous release, you must convert the disk formats as well. Only those disks that were previously converted from LDM format to VDM format can be returned to LDM format. Any disks formatted under DG/UX 5.4 Release 3.00 will be in VDM format and cannot be converted to LDM format. Under revisions previous to this release, VDM-formatted disks cannot be accessed. For more information, refer to the `admpdisk(1M)` manual page and its discussion of the `-o convert` option.

CAUTION After you convert from LDM to VDM, do not boot a pre-DG/UX 5.4 Release 3.00 kernel or stand-alone `diskman`. If you accidentally boot one of them, the pre-DG/UX 5.4 Release 3.00 kernel will attempt to repair what it believes to be a damaged LDM-formatted physical disk. An attempted repair operation destroys the primary virtual information table (VDIT), which restricts the ways in which you can access data. For example, you can continue to read and write virtual disks but you can neither create new ones nor change the attributes of existing ones. To recover, shut down the system and boot a DG/UX System 5.4 Release 3.00 kernel from tape or disk. Execute the following `sysadm` operation:

Device -> Disk -> Physical -> Repair

which copies the secondary VDIT from the rear of the physical disk to the front, restoring the corrupted primary VDIT.

After the VDIT is restored, you may register the disk as follows:

Device -> Disk -> Physical -> Register

and access it as desired.

New Disk Management Features

This release provides the following new disk management features:

- This release supports the *virtual disk*, which is OSM's counterpart to the LDM (Logical disk manager) logical disk. A virtual disk is an amount of space that you reserve on a physical disk onto which data is stored. A virtual disk is further characterized by a type: partition, aggregation, mirror, cache, or combination of same virtual disk types. See *Managing the DG/UX System (093-701088-04)* and the technical brief, *Managing Virtual Disks with On-line Storage Management (012-004406-00)*, for more information.
- The on-line storage management (OSM) facility lets you manipulate your physical disk partitions without having to take your data off-line by unmounting file systems. OSM offers a flexible hierarchy that lets you dynamically restructure configurations. OSM enables on-line backups and on-line data storage failure recovery.
- This release supports new **rename** and **move** operations and improved **copy** and **expand** operations.
- This release supports booting from a multi-piece (aggregated) virtual disk as long as all the pieces reside on the same physical disk.
- In soft formatting, surface analysis is no longer recommended and installation of a bootstrap is optional. In addition, bad block remapping is not included with the creation of the virtual disk information table (VDIT) (formerly system areas).i
- This release supports a new operation to convert between logical and virtual disk formats.
- This release supports the use of either the NVRAM board or disk drives as cache devices. In addition, it supports front-end device sharing among multiple back ends.
- Because pre-DG/UX System 5.4 Release 3.00 **diskman** features were merged into **sysadm**, **sysadm** now provides all of the necessary features to manage disks. If a physical disk is prepared with the new virtual disk format (either formatted directly or converted from the logical disk format), it will contain virtual disk partitions and free space from which more partitions can be allocated. See Section 8, "Installation Instructions", for more information.
- Error messages generated by the file system will no longer use the FSID to identify the file system but will instead use the name of the mount point to help the system administrator diagnose problems.
- For details on allocating and managing disk space on virtual disk-formatted physical disks, consult *Installing the DG/UX System (093-701087)*, *Managing the DG/UX System (093-701088-04)*, and *Customizing the DG/UX System (093-701101-04)* for details. Also, consult the **admpdisk(1M)** and **admvdisk(1M)** manual pages for details on managing physical disks and virtual disks.
- The technical brief, *Managing Virtual Disks with On-line Storage Management (012-004406-00)*, ships with this release. This brief provides an overview of the new virtual disk formats and discusses the transition from LDM to OSM.

Virtual Disk Management (VDM) Cache

The Cache Device Manager (CDM) has been replaced by the VDM Cache. VDM Cache is completely integrated into `sysadm` and `admvdisk`. This integration gives you a consistent way to manage all disk resources plus all of the features of the VDM (such as cache creation) without having to take your data off-line.

The VDM Cache also supports both Non-Volatile RAM (NVRD) front-end devices and regular disk front-end devices. You can use a fast disk front-end device to cache a slower disk device, such as an optical disk.

Because of the additional features offered by the VDM Cache, CDM cache front-end devices cannot be converted to VDM cache format, but rather must be reformatted. To ensure that all your data is flushed to the back-end device before upgrading to DG/UX System 5.4 Release 3.00, you must unmount cached file systems (or shutdown cached databases) and then delete all CDM caches. After the upgrading to DG/UX System 5.4 Release 3.00 is complete, you can recreate the caches using the VDM Cache.

For a general discussion of the capabilities provided by the Virtual Disk Manager device driver, consult the `vdm(7)` manual page.

diskman

To provide a more consistent system management interface, the functions previously provided by the `diskman` utility are now part of `sysadm`. This extends to the stand-alone environment as well as the stand-alone environment.

To boot the stand-alone `sysadm` kernel, execute the following:

```
SCM> b device-specusr:/stand/sysadm ↵
```

For example:

```
SCM> b sd(incr(),0)usr:/stand/sysadm.↵
```

Bootstrap for Virtual Disks

You can now avoid using the `-q` boot command line switch to boot from virtual disks called something other than `root` and `swap`. To do this, set boot defaults on the physical disk from which you boot. Setting boot defaults lets you have unattended automatic reboots of systems using non-default names for `root` and `swap`. See the `admpdisk(1M)` manual page for details of the `set-defaults` option.

The bootstrap can now boot from multi-piece virtual disk aggregations as long as all of the pieces reside on the same physical disk. Therefore, you can now break the `root` and `usr` file systems up into multiple pieces and still boot from them.

Although you can boot only from partitions and aggregations, `root` and `usr` can exist as higher-level virtual disks as well (such as a mirror or a cache). However, you cannot boot from the cached or mirrored `root` and `usr` directly. Instead, you need to supply one of the children as the virtual disk from which to boot.

If **root** or **mirror** is a mirror, you need to select one of the images of the mirror on the boot line, as follows:

```
SCM> b sd(insc(),0)root.be:/dgux
```

If the **root** or **mirror** is a cache, you need to select the back-end device, as follows:

```
SCM> b sd(insc(),0)root_image1:/dgux
```

Keep in mind that the children of the mirror or cache need to be bootable virtual disks (a partition or multi-piece aggregation).

Quiescent Mirror Breaks

An image can be removed from a mirror (unlinked) while the mirror is in use (mounted as a file system, opened by an application, or swapped on). When the mirror is mounted as a file system, the image removed has all buffers flushed so that the image is mountable as a separate file system.

When a mirror is dismantled (unmirrored), if the unlinked image contains a valid mounted file system, that file system is synchronized to disk and marked as not needing **fsck**. The file system will be a consistent copy of the mounted file system at that time. This will make certain backup practices much easier to accomplish.

Because the **mount** command can tolerate some inconsistencies in the file system and does not aggressively attempt to repair inconsistencies in the way **fsck** does, an unlinked image from a mirror may not perform **fsck** cleanly, but can still be mounted.

VDM Mirror Performance

Mirror synchronization speed has improved by up to 35% because the internal size of the I/O request used by the sync read and write operations has increased from 32 blocks to 128 blocks.

4.2.2 Kernel

This release provides the following new and enhanced kernel changes discussed in this section:

- Installer Kernels
- Kernel Source Code Interface Changes
- User scsi Device Driver
- DOS and ISO-9660 CD-ROM File System Drivers
- Enhanced dg_process_info
- Real-time Extension Support

4.2.2.1 Installer Kernels

The roles of the former installer kernel (`/dgux.installer`), the diskless installer kernel (`/usr/stand/dgux.installer.diskless`), and the read-only installer kernel on the CD-ROM release disk have now been subsumed into stand-alone `sysadm`. To boot the stand-alone `sysadm` kernel in one of these other roles, you must specify one of the following additional options on the SCM boot command line:

- i to get the installer kernel.
- d to get the diskless installer kernel.
- r to get the read-only CD installer kernel.

The `/dgux.installer` kernel file is still delivered so that preloaded systems automatically boot up in installation mode when powered on, but the separate file for diskless systems (`/usr/stand/dgux.installer.diskless`) is no longer delivered.

4.2.2.2 Kernel Source Code Interface Changes

Some internal kernel source code interfaces are changed in this release. See *Programming in the DG/UX Kernel Environment* (093-701083-04) for more information. If you build any kernel components of your own (such as device drivers or STREAMS modules), you should check the source code for uses of the interfaces listed in this section and modify them as described below.

Errno and Status Code Macros

Two kernel macros that deal with error numbers and status codes are renamed. You must convert all uses of the `SC_NO_ERRNO` and `SC_ENCODE_STATUS` macros in your source code to `NO_ERRNO` and `ENCODE_STATUS`. To use the new macros, your code must `#include` the header file `ext/status_code.h` instead of the file needed for the old macros (`ii/i_sc.h`). The old macro names are still available in this release for backward compatibility, but they will be removed in the next major release of the DG/UX system.

In addition to the existing panic codes file in `/usr/release`, the status codes file is also now stored there. Any time a system console message mentions a status code (frequently this will be with the phrase `status = <number>`), you can look up the status code in the file in `/usr/release` for a more detailed explanation of the cause of the problem.

Device Address Reservation Functions

The interfaces to the functions `io_check_device_spec` and `io_forget_device_spec` have been changed to include a new argument, which precedes the `device_code` argument in both cases. The new argument is named `device_class`, and it identifies the device class in which the address and device code are being reserved. The argument's type is `uc_device_class_enum_type`. All calls to either of these two functions must be modified to include the appropriate device class as the new argument.

4.2.2.3 User scsi Device Driver

The user scsi device driver interface is a new feature of DG/UX System 5.4 Release 3.00. This driver lets users issue SCSI command descriptor blocks from user space through an **ioctl** interface. It also lets the user set various SCSI host bus adapter options for devices configured using this driver. The supporting header file is located in **/usr/include/sys/int_dg_user_scsi_ioctl.h**. Restrictions are described in this header file. For more information, see the **scsi(7)** manual page.

4.2.2.4 DOS and ISO-9660 CD-ROM File System Drivers

The pseudo devices **hfm** and **dfm** are no longer required to configure the DOS and ISO-9660 CD-ROM file system drivers. Using these in a system file will generate an error.

The ISO-9660 CD-ROM file system driver now supports Rock Ridge format, which is a superset of ISO-9660 with the addition of long filenames, extended DG/UX-like attributes and symlink capabilities. See the **hfm(4)** manual page for more details.

The DOS file system driver now supports partitioned hard disks. Each partition is specified by appending the appropriate letter in the range of **c** through **z** to the device that is to be mounted. See the **dfm(4M)** manual page for more details.

4.2.2.5 Enhanced dg_process_info

A new version (**DG_PROCESS_INFO_VERSION_3**) was added to the **dg_process_info** system call. The data structure for this version was expanded to provide additional information necessary for efficient performance monitoring and to provide useful, value-add information not available on competing systems. Refer to **/usr/include/sys/dg_process_info.h** for a detailed description of the process information available.

In addition to **dg_process_info**, the **dg_lwp_info** system call was added. **dg_lwp_info** works like **dg_process_info**, but provides detailed information about the lwps (lightweight processes) associated with a particular process. Refer to the **dg_lwp_info(2)** manual page and the **sys/dg_lwp_info.h** header file for more details.

4.2.2.6 Real-time Extension Support

This release contains a version of the P1003.4 real-time extensions updated from the previous DG/UX release. The real-time calls have been upgraded from Draft 12 of the Posix 1003.4 specification to Draft 14, as listed below.

This release contains all of P1003.4 Draft 14 Section 11, Synchronization. This was previously section 17 of P1003.4 Draft 12. There have been major changes in the synchronization implementation from P1003.4 Draft 12 to P1003.4 Draft 14. The following P1003.4 Draft 14 synchronization routines are either new or have changed their interface from P1003.4 Draft 12:

sem_init	interface change
sem_open	new routine
sem_close	new routine

sem_wait	new routine
sem_trywait	new routine
sem_post	new routine
sem_getvalue	new routine

The following P1003.4 Draft 12 synchronization routines are no longer supported in P1003.4 Draft 12:

sem_lock
sem_trylock
sem_unlock

This release contains all of P1003.4 Draft 14 Section 13, Execution Scheduling. This was previously section 21 of P1003.4 Draft 12. There are only minor changes in the scheduling implementation from P1003.4 Draft 12 to P1003.4 Draft 14. The following P1003.4 Draft 14 execution routine changed its interface from P1003.4 Draft 12:

sched_yield

This release contains a partial implementation of P1003.4 Draft 14 Section 14, Clock and Timer Functions. This was previously section 22 of P1003.4 Draft 12. There are only minor changes in the clock and timer implementation from P1003.4 Draft 12 to P1003.4 Draft 14. The following P1003.4 Draft 14 clock and timer routine changed its interface from P1003.4 Draft 12:

timer_create

The following P1003.4 Draft 14 routine was added to the partial implementation of clock and timers:

clock_gettime

Previously, the DG/UX system supported four per-process timers. DG System 5.4 Release 3.00 now supports all 32 per-process timers, as required by P1003.4 Draft 14.

Note that when P1003.4 is finalized, it is likely that the interfaces and/or functional descriptions of these routines will change. Data General's long-term goal for real-time extensions is to be P1003.4-compliant. As the specification becomes a standard, future DG/UX releases will be upgraded to the standard P1003.4 interfaces and will no longer support the draft versions. Therefore, Data General does not guarantee backward compatibility to these routines in future releases.

Because of the changes from P1003.4 Draft 12 to P1003.4 Draft 14, any code written to the P1003.4 Draft 12 standard must be modified, recompiled, and relinked. Existing programs that are P1003.4 Draft 12-compliant **will not function correctly** if executed on DG/UX System 5.4 Release 3.00. Please refer to individual manual pages for detailed programming information.

For additional information about the DG/UX real-time extensions, refer to specification *P1003.4 Draft 14, Draft Standard for Information Technology - Portable Operating System Interface (POSIX) - Part 1: System Application Program Interface (API) - Amendment 1: Real-time Extension [C Language]*. The specification is available by contacting the IEEE Standards Office, PO BOX 1331, 445 Hoes Lane, Piscataway, N.J. 08855-1331, phone number 909-562-3811.

4.2.3 TCP/IP

This release provides the following new and enhanced TCP/IP features discussed in this section:

- Simple Network Management Protocol (SNMP)
- **gated**
- Diagnostics tools
- Network Time Protocol (NTP)
- Boot Protocol (**bootp**)
- Serial Line IP (SLIP)
- TCP/IP Protocol, Command, and Library Upgrades
- **sysadm** Changes for TCP/IP

4.2.3.1 Simple Network Management Protocol (SNMP)

This release includes a new Emanate SNMP Agent and utilities based on code from SNMP Research Inc. The command line utilities **getone**, **getnext**, **getmany**, **getany**, **trap_send**, and **trap_recv** were renamed **snmpgetone**, **snmpgetnext**, **snmpgetmany**, **snmpgetany**, **snmptrapsend**, and **snmptraprecv**, respectively. The old names are supported in this release for backward compatibility but will be removed in the next major release of DG/UX. The output format of the new commands has changed as well. As with the command names, the old format will be removed in the next major release of DG/UX.

A complete list of the MIBs (Management Information Bases) supported by the new agent is given below. On-line versions of the MIBs are available in **/usr/etc/snmp** in compressed format. For more information about the SNMP agent, see *Managing TCP/IP on the DG/UX System* (093-701051-06).

- MIB II (RFC-1213) provides information about the system, network interfaces, and statistics for TCP/UDP/IP and SNMP.
- Interface Extensions MIB (RFC-1229) supplements the MIB II interface table and provides additional statistics, such as the number of broadcast and multicast packets, and a table of the addresses configured for each interface.
- Ethernet MIB (RFC-1398) provides Ethernet-specific statistics such as the number of collisions, framing errors, and CRC errors. The agent provides this information for all Ethernet devices on the system regardless of the network protocol using the device.
- FDDI MIB (RFC-1285) is based on the ANSI X3T9.5 SMT 6.2 specification and provides information about the SMT, MAC, PATH, PORT, and ATTACHMENT layers. The agent provides this information for all **pefn** devices on the system, regardless of the network protocol using the device.

- RS-232 MIB (RFC-1317) provides statistics about RS-232 ports on the system such as clock source, frame check errors, and the status of modem control signals. The Data General implementation supports synchronous devices on the system regardless of the application that may be using them.
- Hewlett Packard UNIX MIB provides general information about the UNIX Operating System, such as uptime and load, mounted file systems, and the process table.
- Data General MIB provides information specific to Data General Hardware and Operating Systems. This MIB allows management of the DG/UX printer subsystem and provides information about the hardware devices on your system.

Restrictions

The following restrictions apply for the MIBs provided with the SNMP Agent:

- Set operations are not supported for the TCP Connection Table, Interface Extensions MIB, FDDI MIB, and RS-232 MIB.
- RS-232 MIB supports only VSCi synchronous devices.
- Hewlett Packard MIB does not support all groups. Data General supports the computerSystem, fileSystem, processes, and ieee8023Mac groups.
 - Within the following supported groups, these objects are not applicable for the DG/UX system:
 - ComputerSystemNiceCPU
 - ProcessProcName
 - ProcessTime
 - ProcessesFss
 - The value 0 is always returned for the following:
 - ComputerSystemSwapConfig
 - ieee8023MacExcessiveCollisions
 - ieee8023MacControlFieldErrors
 - The Process Flag object returns the following:
 - 9 if bound and debugged.
 - 1 if bound.
 - 8 if debugged.
 - 0 if otherwise.

4.2.3.2 gated

The **gated** dynamic route management server (daemon), which can now replace the old **routed** program, does the following:

- Provides support for multiple routing protocols.

- Implements improved algorithms for route path fault detection.
- Provides greater configurability.
- Provides support for dynamic reconfiguration operations.
- Provides support for dynamic network interfaces.

Routing protocols currently implemented by **gated** include RIP, HELLO, and OSPF. OSPF and RIP version 2 can be configured to use IP multicast to simplify the **gated** configuration operation and to distribute routing information more efficiently. Support tools that monitor and manage **gated** include **ripquery**, **ospf_monitor**, and **gdc**. For more information, see the **ripquery(1M)**, **ospf_monitor(1M)**, and **gdc(1M)** manual pages.

The routing protocols comply with the following RFCs:

- RFC-891 HELLO Protocol
- RFC-1058 RIP Protocol
- RFC-1247 OSPF Version 2 Protocol

To initiate **gated**, refer to *Installing the DG/UX System (093-701087-05)*.

Note that existing **routed** users who are updating to DG/UX System 5.4 Release 3.00 will be automatically converted to use **gated**, provided the following:

- **routed** did not use the **/etc/gateways** file.
- **routed** was not invoked with any command switches.

Users who are not automatically converted to **gated** are strongly encouraged to perform the configuration operations required to convert to **gated**. See *Managing TCP/IP on the DG/UX System (093-701051-06)* and the **gated-config(4M)** manual page for details.

If the system time is changed after **gated** has started, you will need to stop and restart **gated**. See the **admtcpipdaemon(1M)** manual page for information on how to perform this operation.

4.2.3.3 Diagnostics tools

tcpdump

This release supports the Berkeley Packet Filter (BPF) kernel interface that allows raw data to be read from the network. The **tcpdump** command provides a powerful user interface to BPF. The **tcpdump** command lets any machine on the network function as a LAN analyzer. It puts a device driver into a promiscuous mode in which it can capture all packets on the network. **tcpdump** can decode a user-selectable subset of the packets in real-time or it can store them for later decoding.

Since **tcpdump** is a command, network traffic can be viewed simply by typing **tcpdump** at the shell prompt. Since your kernel already contains **bpf** by default, no rebooting is necessary and other processes running on the machine can continue to run. **tcpdump** can also efficiently filter out unwanted packets if given a boolean expression. This boolean expression can consist of a protocol name (e.g. ip, tcp, telnet, etc.), a command (e.g. host dawg), and/or an arbitrary comparison of the contents of each packet.

The following protocols can be displayed and processed by **tcpdump**:

ether
slip
llc
snap
ip
tcp
telnet
udp
arp
rarp
icmp
tftp
igmp
nfs
snmp
ntp
talk (apple talk)
bootp
dns
rip
ospf
hello

nfc

nfc (network file converter), a companion command, converts files containing raw network data from one format to another. File conversions currently supported are: from a **tcpdump** capture file to a Network General Sniffer™ save file, and from a Network General Sniffer file to a **tcpdump** file. This lets a Network General Sniffer display packets captured by **tcpdump** and vice versa.

traceroute

traceroute, another new command, prints the route a packet takes from the local host to a user-specified remote host. The time taken to reach each host is also printed. This command is intended for use in network testing, measurement, and management.

4.2.3.4 Network Time Protocol (NTP)

Network Time Protocol (NTP) is an Internet standard protocol for synchronizing time-of-day clocks on machines in a network. All machines that run the NTP program exchange time messages via the `xntpd` server program. Each machine performs appropriate adjustments to the designated local system clock so that all machines converge on a single time.

The NTP protocol complies with RFC-1305. See the `xntpd(1M)` manual page for more information.

4.2.3.5 Boot Protocol (bootp)

`bootp` is a commonly-used protocol for booting devices via the network. There is a client side (the device being booted) and a server side (the system that responds to client boot requests). The DG/UX system provides **only** the server side. Devices that use `bootp` as a download protocol include terminal servers, X-terminals and routers.

The device takes the information provided by the `bootpd` server and from a `tftp` request to get its second-stage bootstrap file. `bootpd` provides information such as the IP address of the device, hostname, subnet mask, gateway, boot server, various other servers, and the location and name of the second-stage bootstrap file.

The `bootp` protocol is defined by the following RFCs:

- RFC-951 Bootstrap Protocol
- RFC-1497 BOOTP Vendor Information Extensions

4.2.3.6 Serial Line IP (SLIP)

SLIP (Serial Line IP) provides a simple, low-cost, long-distance solution for running TCP/IP between a pair of machines via a custom async line or via async modems and phone lines. Standard TCP/IP utilities such as `telnet`, `ftp`, and NFS work over a SLIP connection, though performance is significantly less because async lines are much slower than Ethernet or other LAN media.

The SLIP protocol is defined by the following RFCs:

- RFC-1055 Transmission of IP Datagrams over Serial Lines
- RFC-1144 Compressing TCP/IP Headers for Low-Speed Serial Links

See the `slipd(1C)` manual page for additional information.

4.2.3.7 TCP/IP Protocol, Command, and Library Upgrades

The following changes were implemented in the TCP/IP protocol stack:

- The RAW socket interface was modified so that it does not require an IP header to be explicitly prepended to the data.
- The following global TCP/IP tunable parameters were added to allow for better kernel configuration:
 - IP Time-to-Live is the maximum hop count an IP packet can travel before being dropped.
 - IP Forwarding is a switch to prevent the forwarding of IP packets (except source-routed packets).
 - IP Non-Local Source Route Forwarding is a switch to prevent the forwarding of IP source-routed packets to a different net.
 - TCP Default Keep-Alive Idle Time is the time in seconds TCP will wait before sending Keep-Alive probes.
 - Base ARP Cache Time-out is the time in seconds used as the base value for determining ARP cache time-outs.
- The following per-connection TCP/IP tunable parameters were added:
 - TCP Keep-Alive Idle Time is the number of seconds TCP waits before sending Keep-Alive probes.
 - TCP Maximum Number of Retransmissions is the number of times TCP retransmits a packet before dropping the connection.
- Changes were made to the TCP/IP protocol stack to correct items of nonconformance with the Requirements for Internet Hosts RFC 1122.
- IP supports non-local subnet routing and route metrics.

The kernel now logs TCP/IP events that trigger error reporting. The following events are being logged: failure to make a new TCP connection, reset of an existing connection, and change of an interface state.

A programmable interface to FTP was provided in this release. The FTP Programmable Interface is a set of `libdgc` routines that implements a programmable interface to the File Transfer Protocol. The following functions that implement the programmable interface can be called from C or FORTRAN programs.

<code>ftp_xfer</code>	transfer file
<code>ftp_del</code>	delete file
<code>ftp_mkdir</code>	make directory
<code>ftp_rmdir</code>	remove directory
<code>ftp_rename</code>	rename directory
<code>ftp_dir</code>	display directory
<code>ftperror</code>	display error

For more information, refer to the **ftp_xfer(3N)** manual page and to *Programming with TCP/IP on the DG/UX System (093-701024-03)*.

The following commands and server program were changed:

- The **route** command now lets you set MTU (Maximum Transfer Unit) size and netmask for individual routes. **route** also provides other new commands and functions. See the **route(1M)** manual page for more information.
- The **ping** command was brought up to BSD4.4 functionality, and now offers more options and functions. See the **ping(1M)** manual page for more information.
- The **telnet** command now supports the **-e** switch to specify or disable the **telnet** escape character. See the **telnet(1C)** manual page for more information.
- The **ftpd** server program now supports enhanced logging, welcome messages and bannering, custom password prompts, restricted user accounts, selective access controls, and concurrent session limits. For more information, See the **ftpd(1M)** manual page.

4.2.3.8 Library Change for TCP/IP Programmers

/usr/include/arpa/inet.h has a revised prototype for **inet_addr(3N)**, which converts dot notation Internet addresses into internal form. The return type changed from a struct **in_addr** to an unsigned long to comply with prevailing practice and standards such as OCS Net. There is no object or binary impact because both return types are returned the same way by the compiler, as required by OCS.

The types of source code changes required include:

```
(struct in_addr) = inet_addr; becomes
(struct in_addr).s_addr = inet_addr;
```

and

```
(unsigned long) = inet_addr.s_addr; becomes
(unsigned long) = inet_addr;
```

4.2.3.9 sysadm Changes for TCP/IP

The following changes occurred in the **sysadm** interfaces for the TCP/IP product:

- The entire TCP/IP menu was redesigned to better manage server programs, databases, interfaces, parameters, and routes.
- The routing sub-menu was redesigned to better control static and dynamic route configuration.
- The interfaces submenu was changed to control each network interface individually via **sysadm**.

- A **sysadm** interface (**Logging -> Network**) provides an easy method by which to access the network error logs. The STREAMS logging mechanism was added for the following TCP/IP events: failure to make a new TCP connection, reset of an existing connection, and change of an interface state. A new **sysadm** interface provides access to this logging information.

4.2.4 POSIX Threads (Pthreads)

This release provides support for the following Posix Threads features discussed in this section:

- Medium Term Scheduler
- Reentrant User Libraries

This release includes support for Posix Threads (Pthreads). Because the Pthread standard (IEEE 1003.4a) is not yet finalized, DG/UX implements Draft 6 of Pthreads. DG/UX plans to support future drafts as well as the final standard. DG/UX plans to provide binary and source compatibility for at least one major revision beyond this release. This gives applications time to migrate to future drafts and eventually to the final standard. A similar migration path is planned for any future drafts.

DG/UX also provides several extensions to Draft 6, most notably Pthread Groups, which give an application more precise control over the execution locality of threads in a process without sacrificing portability to other vendors' Pthread implementations. Pthread Groups can be used in conjunction with the Hierarchical Processor Affinity that is also new in this release. Briefly, Pthread groups can be affined to sets of CPUs backed by the same shared cache, thus optimizing cache locality on the underlying SMP machine.

Pthread primitive support comes in the form of new headers (e.g. **usr/include/pthread.h**) and a shared library (**libpthread.so**). Because Pthreads are not yet standard and because DG/UX will provide support for future drafts, all DG/UX System 5.4 Release 3.00 thread programs must be compiled with the symbol **_POSIX4A_DRAFT6_SOURCE** defined. This can be done easily by defining this symbol on the compile line using the **-D** switch. Note that the Pthread primitives are available only in the Elf environment and can be linked-shared only.

For more details on programming with Pthreads, see the **threads(5)** manual page. This manual page gives an overview of standard Pthread calls and DG/UX Pthread extensions and details of programming with Pthreads in the DG/UX environment. Specific pthread manual pages are listed in the **threads(5)** manual page.

4.2.4.1 Medium Term Scheduler (MTS)

Two new configuration parameters were added to support the MTS. **MINBOUND** is the minimum number of user threads that can have resources to be dispatched. **MAXBOUND** is the maximum number of user threads that can have resources to be dispatched. If either of these is set to 0 (the default), the MTS calculates them based on system resources.

If the system becomes over-committed with too many threads competing for finite resources (memory and processor), the MTS begins restricting the number of threads that can be dispatched down to the minimum **MINBOUND**. This eases the demand for the finite resources. When the crunch time for resources ends, the MTS gradually increases the number of threads that can be

dispatched up to the maximum MAXBOUND. If MINBOUND equals MAXBOUND, the MTS algorithms for limiting dispatchable threads are turned off and the number of dispatchable threads is always MAXBOUND.

4.2.4.2 Reentrant User Libraries

The introduction of threads in DG/UX System 5.4 Release 3.00 lets you develop applications that simultaneously call C library functions from multiple threads. These functions must be made reentrant in order to handle simultaneous calls. The initial focus of this reentrancy work has been on the `libc` and `libm` functions specified in POSIX 1003.1.

In some cases, the standard interface to a function does not allow reentrancy (for example, functions that return a pointer to static data). POSIX 1003.4a Draft 6 provides a new function with the interface appropriately modified to allow for reentrancy. These new functions usually retain the name of the original function with an `_r` suffix.

When using library functions in threads applications, be aware that each library function consumes a varying amount of stack depending on the particular code paths taken. If the default stack size is inadequate for your combination of application code and library functions, you can select the size of the thread stack when the thread is created. Therefore, you need to recognize when a thread runs out of stack space. This usually shows up as a bus error or a segmentation violation.

Some features provided by library functions are global to your process (for example, locale state). If you change any aspect of the locale, such as timezone or message databases, that change affects all threads in the process.

See the `reentrant(3)` manual page for a detailed list of reentrant functions and general guidelines for creating threads applications. Individual manual pages for functions detail specific threads programming issues.

Note the following restrictions:

- 1003.4a/d6 priority ceilings and priority inheritance for mutexes are not supported in this release.
- 1003.4a/d7 is not yet implemented.
- Not all of `libc` is reentrant. See the `reentrant(3)` manual page for more information.

4.2.4.3 WARNING: Thread Cancellation

The POSIX standard for threads programming, P1003.4a, is in draft form. Thread cancellation is one area of P1003.4a that is still undergoing refinement. For example, Draft 6 requires only that vendors document which functions can recognize cancellation requests. Draft 8 requires that only certain functions be able to recognize cancellation requests. It is possible for application behavior to change or be broken completely whenever a function's cancellation behavior is changed. Data General therefore recommends that, if cancellation must be used, the application should make no assumption about the cancelability of any functions. A coding paradigm that insulates the application from changes in this area includes the following:

- If you are calling any DG/UX functions from the thread, always first push a cleanup handler if there is state that might need cleanup if any of the functions being called is a cancellation point. You can use the `pthread_setintr` call to explicitly turn cancelability off and on within a thread.
- Wherever a thread is required to recognize cancellation, use a call to `pthread_testintr`. Do not rely on any particular function to contain a cancellation point.

These steps should help insure that, as cancellation points are added or removed in future revisions of DG/UX, your application will continue to behave in the expected manner.

4.2.5 Debuggers

This release provides the following new and enhanced debugger features discussed in this section:

- **mxdb** and **dbx**
- **sdb**
- **gcore**

4.2.5.1 mxdb and dbx

Support for debugging multi-threaded processes has been added to **dbx** and the value-added product, **mxdb**, including:

- Focusing to different threads.
- Displaying the status of some or all threads.
- Displaying information about mutexes and condition-variables.
- Making stop events (such as breakpoints) thread-specific.

See the associated manual pages and/or on-line help for further details.

4.2.5.2 sdb

sdb does not have any special support for debugging multi-threaded processes. It can debug them but it cannot display any threads-related information.

4.2.5.3 gcore

The new **gcore** utility produces a core dump for a process without affecting the process's execution; for example, a core dump can be produced without terminating the process. This is useful in capturing the state of a process for examination with a debugger such as **sdb**, **dbx** or **mxdb**. See the `gcore(1)` manual page for more information.

4.2.5.4 Software Status Register (SWSR)

When requesting the machine state for a thread with **mxdb**, **dbx**, or **sdb**, users may notice that bits other than **MC_XREGS_DEFINED** are sometimes set in the software status register (SWSR). These bits are used by the DG/UX system to record implementation-specific thread-state information. They are not meaningful to the user and must not be modified.

The SWSR can contain the **MC_XREGS_DEFINED** bit when in fact the xregs (extended registers) were not modified by the thread. This inconsistency is caused by the kernel and will be removed in a later release.

4.2.6 Compilers

This release provides the following new and enhanced compiler features discussed in this section:

- **Short Addressing**

4.2.6.1 Short Addressing

Short addressing is an optimization that can reduce the number of MC88000 machine language instructions required to access a datum or to form its address. See *Porting and Developing Applications for the DG/UX System (069-701059-05)* for more information on the proper use of the new **-mshort-data-size** option to **gcc**.

4.2.7 Libraries

4.2.7.1 WARNING: libnsl

Support for the static library **libnsl.a** has been removed. Use the shared library **libnsl.so** instead.

4.2.8 Commands

This release provides the following new and enhanced commands discussed in this section:

- **syslogd**
- **sysadm**
- New **sar** and **ps** commands (**nps** and **nsar**)
- **pread** and **pwrite**
- **ls -L**
- **login**

- **man**
- **Perl**
- **ttymon**
- Archive Utilities (**cpio**, **dd**, **dump2**, **tar**, **REELexchange**, and **systemtape**)

4.2.8.1 **syslogd**

All messages sent to the system console by DG/UX software are also written to the `/var/adm/message` file by the `syslogd` server program. You can control the level of messages written to the file with the `syslogd` parameters. This is new for DG/UX 5.4 Release 3.00. Previously, many messages just went to the console; these messages could be easily lost or missed if they scrolled off the screen.

4.2.8.2 **sysadm**

A file system can be expanded while it is mounted and in use. This capability is provided through `sysadm` or `admfilesystem`.

`sysadm` now lets you stop any port monitors after you have stopped any other port monitors. Formerly, if you stopped a port monitor, `sysadm` did not offer as choices the monitors later than the one that you previously stopped.

For example, if you had three monitors (`ttymon1`, `ttymon2`, and `ttymon3`) and you stopped `ttymon2` and then re-entered the Stop Port Monitor menu, you could choose only `ttymon1`. `sysadm` now lets you choose `ttymon1` and `ttymon3`.

4.2.8.3 **New sar and ps commands**

There are new `sar` and `ps` commands that gather additional performance data. The commands, called `nsar` and `nps` respectively, will replace the existing `sar` and `ps` commands in the next major release of DG/UX when the existing commands will be named `osar` and `ops`, respectively.

nps

`nps` is a new utility that reports information about active processes and threads. For more information and options, refer to the `nps(1)` manual page.

`nps` provides information on threads with its new `-L` option. `nps` also provides access to many previously-unavailable forms of per-process information with the new `-o` option, while remaining compatible with the existing `ps` using the `ps` options. Among the new information are: process stack size, context switches, page faults, and input and output operations. The `-o` option also provides control over the output format. See the `nps(1)` manual page for more information.

For a quick start, type `nps -eWL` to get the status of all threads in the system.

WARNING: ps and nps May Report Larger Value in SZ Field

The `ps` and the `nps` commands may now report a larger value for the SZ field than in previous revisions of DG/UX because shared anonymous pages are now counted in the resident set size of a process. Consequently, a process that attaches a shared memory segment will now manifest a larger value in the SZ column of a `ps` listing than in the past.

nsar

`nsar` provides additional system performance information, such as paging and memory usage statistics, and information on virtual disks. `nsar` supports new standard `sar` options `-k`, `-g`, and `-p`. The new `-W Ostatlist` option provides access to over one hundred new statistics as well as more precise control over formatting of the standard statistics. The new `-W D=disktype` option allows control over the types of disks (virtual or physical) whose activity is monitored. A new data collector, `nsadc`, accompanies `nsar` and no longer has limitations on the number of disks it can sample. The new log file format is more compact than the existing `sadc` log format. See the `nsar(1M)` and `nsadc(1M)` manual pages for more information.

4.2.8.4 pread, pwrite

`pread` and `pwrite` allow for an atomic seek and transfer operation to a file. This lets multithreaded applications seek and transfer data on a file descriptor atomically without worrying about locking the file position during the I/O transfer.

4.2.8.5 ls -L

The behavior of `ls` under the `-L` option is changed to conform to the documentation. If the argument to the `-L` option is a symbolic link, `ls` displays the target filename and the statistics for the target file. If the final target file does not exist, `ls` displays no information.

If the argument to `-L` is the default argument of `.` (current directory), `ls` does not resolve links in that directory, since `.` is not a symbolic link.

4.2.8.6 login

Your username is now always echoed as you log on because of a new configurable variable, `UPPEROLD`, offered in this release. `login` assumes that you have an uppercase-only terminal and sets special terminal characteristics if you enter a username that contains uppercase characters. You can now disable this feature by setting `UPPEROLD` to `NO` in the file `/etc/default/login`.

Previously, if your `/etc/issue` file was larger than 500 bytes, it and the ensuing login prompt were not always displayed properly as you logged on. Now, `/etc/issue` files as large as an 80x24 character screen are displayed properly.

4.2.8.7 man

The **man** command was enhanced to handle **.Z** files (files compressed by the **compress** program). If you add local on-line manual pages to your system, you can save space by running **compress** on the manual page. This produces a smaller file than does **pack**, the program previously used to compress DG/UX manual pages.

4.2.8.8 Perl

Perl, a Practical Extraction and Report Language, is included in this release. **Perl** is an interpreted language optimized for scanning arbitrary text files, extracting information from those text files, and printing reports based on that information. It is also a good language for many system management tasks. The language combines some of the best features of **C**, **sed**, **awk**, and **sh**, so anyone familiar with those languages should have little difficulty with it.

Please read the **perl(1)** manual page for more information. You can also order *Programming Perl* (069-100513-00) from Data General.

4.2.8.9 ttymon

ttymon now supports several new configuration variables. The default values of these variables relieve problems that were seen primarily on VTC lines. See the **ttymon(1M)** manual page for additional information.

Also, in DG/UX System 5.4 Release 2.01, **ttymon** logged certain *Interrupted system call* messages. In previous revisions, these messages were not logged. The previous behavior has been restored.

4.2.8.10 Archive Utilities

cpio, **dd**, **dump2**, and **tar** now report extended **errno** text information (if available) in addition to other information. Now, for example, when you try to write to write-protected media, you get an easy-to-understand error message.

cpio

Formerly, **cpio** did not notice errors when trying to load files onto NFS-mounted file systems without appropriate privileges. The most common instance was when **root** on **systemA** tried to load a small file onto **systemB**'s disk, but **systemB** had not exported the file system with **-root=systemA**. **cpio** could remove any existing file and create a new file, but it could not write data to it and could not report an error. Now, **cpio** notices all such errors and preserves any such existing files.

The logic **cpio** uses to compute the block count as it finishes is now more accurate.

Formerly, **cpio -i** ignored any **-c** or **-H** options. When loading, it always decided what header format to use by examining the archive. Now, if you supply **-c** or **-H**, **cpio** considers only your requested header format. If you do not supply **-c** or **-H**, **cpio** considers only **asc**, **bin**, **crc**, and

odc header formats. This makes it much more likely that **cpio -ik** can work with corrupt archives, or when starting with the second or later volume of a multi-volume archive. You can now use **-Hbin** to explicitly request the default header format. You must now use **-H** to load **tar** or **ustar** formats.

Do not use a buffer size larger than 65536 bytes when dumping to tape with **cpio** because you may produce a multi-tape archive that **cpio** cannot load.

tar

The **tar** program has changed as follows:

- **tar**'s maximum blocking factor has increased from 32 (16384 bytes) to 128 (65536 bytes).
- **tar** now continues if it encounters an error while trying to open the second or later volume in a multi-volume archive.
- **tar** supports loading an archive that contains very long pathnames.
- Formerly, **tar** could either report errors and not set time stamps on directories, or not report errors and set time stamps on the wrong directories, when using the **-A** option to load an archive that was created without using the **-A** option. This problem has been fixed.
- If you use the **n** function modifier when loading from a diskette, **tar** stops at the end of the meaningful data, rather than reading the entire diskette.

REELexchange

REELexchange's **twrite** command now supports an **/etc/passwd** file larger than 3000 bytes when REELexchange's **taccess** command runs.

When using REELexchange to transfer binary files, use **tkey** to set **tchar=none**. If you do not, REELexchange displays an informative message when it goes too long without finding a record delimiter. Previously, it could dump core.

systemtape

The **systemtape** command was enhanced for this release. The command now produces tapes in a format that turns off the package load feature that copies root images to all OS clients. This eliminates the need on the load/restore operation for a root file system that is three times larger than the size of the root originally captured on the systemtape. See the **systemtape(1M)** manual page for more information.

4.2.9 Standards

4.2.9.1 X/Open XPG3 Branding

DG/UX System 5.4 Release 3.00 has not been XPG3-branded. XPG3 branding will be pursued in a future release of DG/UX.

4.2.10 Printer Services

This release provides the following new and enhanced printer features discussed in this section:

- Filter Enhancements
- LP Enhancements
- **pcnfsd**

4.2.10.1 Filter Enhancements

The LaserJet® line printer (**lp**) filter now supports the LaserJet 4 series of printers. This enhancement enables you to select many standard LaserJet 4 features such as font, print quality, and margins when printing with the **lp** command. The LaserJet 4 features can be accessed through the **-y mode** option of the **lp** command.

To configure a printer queue so that the LaserJet 4 features are available, choose **hplaserjet** as the input content type when adding or modifying a printer queue to your system.

In addition to the new LaserJet 4 filter support, a set of new terminfo entries is provided for the LaserJet printer. When configuring the LaserJet 4 printer queue, you need to specify one of the following terminfo entries as the printer type:

```
hplaserjet4
hplaserjet4-PS
hplaserjet4-PS-b
hplaserjet4-PCL
```

You can set the printer type and content type through the **sysadm** or the **lpadmin** command.

The **pageprint** and **postprint** printer filters now include new page formatting options, as follows:

- **lp -ypage_nums file**, which prints centered page numbers at the bottom of each page.
- **lp -ydate file**, which prints a left-justified date string on the first line of each page.
- **lp -ytitle=*my title file***, which prints a centered title string on the first page of the document.

These features are available with the **lp** command using the **-y mode** option.

4.2.10.2 LP Enhancements

The **lp** subsystem now supports the **abort** printer fault recovery option. This feature cancels a print job if a fault occurs during printing. This option is useful if you are running unattended print jobs. To configure a printer to abort a job, select the **abort** fault recovery option when adding or modifying a printer queue with **sysadm**.

This release provides terminfo entries that support using one printer in both PS and PCL emulations concurrently. The following new printer-specific terminfo entries initialize the printer emulation mode when a job is printed:

```

dg6640T-PS      hplaserjet3-PS
dg6640T-PS-b    hplaserjet3-PS-b
dg6640T-hp2     hplaserjet3-PCL
dg6646T-PS      hplaserjet3Si-PS
dg6646T-PS-b    hplaserjet3Si-PS-b
dg6646T-hp2     hplaserjet3Si-PCL
dg6771-PS       hplaserjet4-PS
dg6771-PS-3     hplaserjet4-PS-b
dg6771-hp3      hplaserjet4-PCL
dg6772-PS
dg6772-PS-b
dg6772-hp2

```

To use a printer in PS and PCL modes concurrently, you need to define two virtual printer queues. Set the printer type of one queue with a printer-specific PS terminfo entry (for example, **hplaserjet4-PS**, and set the printer (for example, **hplaserjet4-PCL**). Then, define both queues to use the same device (for example, **/dev/lp**). Use the **sysadm** or **lpadmin** to configure your printers.

4.2.10.3 pcnfsd

The **pcnfsd** V2.0 server program is an authentication server that verifies username and password information for all NFS mount requests to the server from DOS clients. The mounted file systems on the NFS servers appear to be DOS disk drives. These NFS drives operate in a manner similar to local and networked drives. **pcnfsd** is also responsible for printing to NFS-mounted redirected printer devices.

pcnfsd hardware requirements are as follows:

- A licensed AViiON server with any valid network controller.
- A client PC, which can be any PC or PC-compatible machine.

pcnfsd software requirements are as follows:

- A licensed AViiON server with DG/UX System 5.4 Release 3.00, TCP/IP, NFS, YP/NIS (optional).
- A client PC with MS DOS 3.2 or higher and LAN WorkPlace for DOS, including Novell's ODI drivers.

Printing from the pcnfsd V2.0 Server Program

Note the following when printing from **pcnfsd**:

- You must export the **var/spool/pcnfs** directory on your AViiON server so that any PC can access it for printing.
- Type the following on your PC to enable printing:

```
C:> net link prn: \\hostname\printer_name ↵
```

Where:

hostname: is the name of your AViiON server with the desired printer and has the file system **/var/spool/pcnfs** exported to the net group (everyone).

printer_name: is the name of the printer on the AViiON server.

For example:

```
C:> net link prn: \\slick\laser ↵
```

```
C:> print c:\autoexec.bat ↵
```

4.3 Software Installation

This release provides new software installation procedures that are performed through stand-alone **sysadm** instead of stand-alone **diskman**. In addition to the improved interface, the following new features were added:

- Capacity Checker
- Cleanup Procedures
- Remote File System Loading

For more information on these installation features, refer to *Installing the DG/UX System* (093-701087-05).

4.3.1 Capacity Checker

A capacity checker was added to the installation procedures to check the available disk space against the space requirements of packages that you have selected for loading. If the capacity checker detects that there is not enough space to load the software packages, you are given the option to expand virtual disks or clean up your system. Disk capacity is checked before the packages are loaded.

4.3.2 Cleanup Procedures

If the capacity checker detects that there is not enough space to load the packages you have selected, you have the opportunity to clean up first. The cleanup procedures automatically locate obsolete files on the system and present a list of files for you to remove.

4.3.3 Remote File System Loading

The installation procedures now provide a new loading option that allows you to load packages from a remote file system that contains a release area. Multiple systems can install or upgrade simultaneously from the same remote file system.

5 Notes and Warnings

This section contains notes and warnings about this release.

5.1 Notes

5.1.1 Copying a CD-ROM Virtual Disk

To copy a read-only virtual disk (such as one on a CD-ROM device), do not use the `sysadm` virtual disk copy operation. Should you attempt to copy such a virtual disk, you will receive an error. For more information, refer to *Managing the DG/UX System* (093-701088-04).

5.1.2 Caching a CD-ROM Virtual Disk to a Magnetic Disk

If you are using a CD-ROM device that contains virtual disks and are not satisfied with the performance of the CD-ROM, copy the virtual disk from the CD-ROM to a higher-performance magnetic disk. See the instructions elsewhere in this release notice on copying virtual disks that reside on a CD-ROM.

Should you decide to use a cache virtual disk in conjunction with a virtual disk that resides on a CD-ROM, you must follow one of the following two procedures.

5.1.2.1 On-Line Non-Persistent Caching

On-line non-persistent caching can be used while the CD-ROM virtual disk is in use (mounted), but the results are not persistent. In other words, the next time the machine is rebooted, the cache will not be there and you will have to follow this procedure again to recreate it if you want it again.

In both of these procedures, the virtual disk that resides on the CD-ROM is called `+root`, and the virtual disk on a magnetic disk that we are going to use as a front-end device for the cache is called `+root.fe`.

This procedure is as follows:

1. Put the CD-ROM virtual disk in the floating state by typing the following:

```
# admvdisk -o modify -pf +root ↵
```

2. Use **Sysadm Device->Disk->Virtual->Caches->Cache** to establish the cache.

5.1.2.2 Off-Line Persistent Caching

Off-line persistent caching results in a persistent cache that is automatically reestablished each time the system is booted, but this procedure cannot be used while the CD-ROM virtual disk is in use (mounted).

In both of these procedures, the virtual disk that resides on the CD-ROM is called **+root**, and the virtual disk on a magnetic disk that we are going to use as a front-end device for the cache is called **+root.fe**.

This procedure is as follows:

1. Make sure that **+root** is not mounted.
2. Create the cache, but make it temporary as it is created by typing the following:

```
# admvdisk -o create -t cache -pt -c +root +root.cache ↵
```

3. Link the front end to the cache as follows:

```
# admvdisk -o link -c +root.fe +root.cache ↵
```

4. Make the cache persistent as follows:

```
# admvdisk -o modify -pp +root.cache ↵
```

Making the cache persistent will generate an error message like the following:

```
Error: Could not make virtual disk "+root.cache" be persistent.
Error: Cannot store the attributes on some of the required disks.
The changes are incomplete.
```

In addition, a message like the following will appear on the console:

```
Jan 7 13:57:09 jagger dg/ux: Warning: Cannot store subdriver
entry for ID 9B80FCDC 'vdmcache' on disk
'sd(ncsc@13(FFFB0080,7),3,0)' (status = 4006006).
Jan 7 13:57:10 jagger dg/ux: Error: Cannot store the attributes
for virtual disk 'vdm(+root.cache,2D2DB06E,373A2264,0)' on disk
'sd(ncsc@13(FFFB0080,7),3,0)' (status = 4006006).
```

Despite these messages, the cache will be persistent on the magnetic disk that contains **+root.fe**, and the cache will be recreated when the system is rebooted.

Note that the cache is named **+root.cache**. Do not give it the same name as the cache back end (here, **+root**) because having two virtual disks with the same name can be confusing.

5.1.3 Sizing Swap Area Needs for Debugging

When debugging a process, the operating system currently allocates swap pages for the entire segment when any part of the segment is modified. For example, if a process has 100 pages of instructions in its text segment and the user creates a breakpoint for one instruction, then 100 pages of swap space are immediately allocated for that process. These pages remain allocated to the process until it terminates—even if the debugger detaches from it.

For shared ELF processes, the debugger automatically sets a breakpoint in **libc.so.1** to monitor shared-object open/close activity. This causes approximately 100 pages of additional swap space (the size of **libc**'s text segment) to be allocated for the process.

5.1.4 **sac**

Do not purge **sac**'s log file (**/var/saf/_log**) in a manner that changes or removes that file's inode while **sac** is running. See *Managing the DG/UX System* (093-701088-04) for details on how to clean up this log file without causing problems for **sac**.

5.1.5 Sync Jumpering

If you use the integrated sync controller on AV530, AV550, AV4600 or AV5500 computers and you are using external clocking, you **MUST** change the jumpers on the System Board for your sync applications to work correctly with DG/UX System 5.4 Release 3.00. Instructions on how to do this follow. Contact your customer support center if you do not have the additional jumpers you will need or if you have configuration questions that are not answered below.

5.1.5.1 AV530 and AV4600 Instructions

Be sure that the machine has been powered OFF before proceeding. Follow the procedures discussed in Chapter 10 of *Setting Up, Starting, Expanding, and Maintaining AViON 530 and 4600 Computers* (014-002091-01) to access the system board. Figure 10-11 of that manual will be helpful in performing this procedure. Be sure to use an ESD kit immediately after removing the cover of your machine. To access the system board jumpers you may need to remove an I/O option card. The jumpers that need to be changed are a row of ten jumpers located immediately below a ribbon cable on the system board near the back of the chassis.

OUT means no jumper connects the pins and at least one pin is fully exposed. IN means a jumper covers both pins.

The previous DG/UX jumper configuration (left to right) was as follows:

OUT IN OUT IN OUT OUT IN OUT OUT IN

If you are using external clocking on channel A (Com A) only, one additional jumper needs to be added (left to right), as follows:

IN IN OUT IN OUT OUT IN OUT OUT IN

If you are using external clocking on channel B (Com B) only, one additional jumper needs to be added (left to right), as follows:

OUT IN OUT IN OUT IN IN OUT OUT IN

If you are using external clocking on both channels, two additional jumpers need to be added (left to right), as follows:

IN IN OUT IN OUT IN IN OUT OUT IN

5.1.5.2 AV550 and AV5500 Instructions

Be sure that the machine has been powered OFF before proceeding. Follow the procedures discussed in *Installing, Expanding, and Maintaining AViiON 5500 and 550 Series Computers* (014-002250) to access the system board. Be sure to use an ESD kit immediately after removing the cover of your machine. To access the system board jumpers, you may need to remove an I/O option card. The jumpers that need to be changed are a row of five jumpers located near the rear of the system board approximately two-thirds of the way from the top.

OUT means no jumper connects the pins; at least one pin is fully exposed. IN means a jumper covers both pins.

The previous DG/UX jumper configuration (left to right) was as follows:

P10	P6	P5	P4	P3
OUT	IN	OUT	OUT	IN

If you are using external clocking, one additional jumper needs to be added (P5). On some machines, a “spare” jumper is available and is attached to only one of the pins at jumper location P5. The correct jumper configuration (left to right) for external clocking is as follows:

P10	P6	P5	P4	P3
OUT	IN	IN	OUT	IN

5.1.6 Installing Packages in a Secondary Release Area

Information is missing from Chapter 8, “Adding Secondary Operating System Releases,” of *Customizing the DG/UX System* (093-701101-04).

Page 8-6 shows how to modify `/etc/fstab` on the diskless client to grant read/write access to the `usr` file system. Add the following action at the end of Step 6.

Locate the following line in `/etc/fstab`:

```
server_dgux_543:/srv/release/dgux_54R300/usr/opt/X11  nfs ro x 0
```

Revise the line to read as follows:

```
server_dgux_543:/srv/release/dgux_54R300/usr/opt/X11  nfs rw x 0
```

5.2 Warnings

5.2.1 Stand-Alone sysadm

The stand-alone **sysadm** utility uses a memory-based **root** file system. The size of this **root** file system is limited to the amount of physical memory on your system. On small systems (such as 16MB), when you boot stand-alone **sysadm**, if you do not select a swap device right away, you can run out of paging space. This will result in panic number 57000145.

To avoid this problem, you should select a swap device as soon as you enter stand-alone **sysadm**. This can be accomplished by choosing the **Install Software** menu. Then run the first two install steps, as follows:

1. **Prepare physical disks**
2. **Prepare virtual disks**

If you are running on a system that has 16MB and includes a large number of disks (that is, more than six disks), you may still receive this panic when attempting to select the swap device. In this case, you should select only the disk that will contain the swap device, when preparing physical disks.

5.2.2 sysadm Operations

The following **sysadm** operations contain a subtle bug:

```
Device->Disk->Virtual->Create
Device->Disk->Virtual->Expand
Device->Disk->Virtual->Cache->Cache
Device->Disk->Virtual->Cache->Link
File System->Local Filesys->Expand
```

In each of these operations, you will be queried with one of the following prompts:

```
Select Space by:
```

or

```
Select Front End by:
```

At such a query, you might select the following:

Disk to partition and partition size

and then select a disk to partition and specify a number of blocks. If you then decide that you have made a mistake and you want to choose something other than **Disk to partition and partition size**, do not back up to the **Select by** query. In other words, neither type ^ in **asysadm** nor click on **Prev...** in **xsysadm**. Instead, cancel the operation either by typing **q** in **asysadm** or by clicking on **Cancel** in **xsysadm**, and then re-start the operation.

In addition, if you use **sysadm Device->Disk->Virtual->Create** to create a a striped virtual disk, **sysadm** will query you for each piece with the following prompt:

```
Select Space by: [Disk to partition and partition size]
```

Although **sysadm** appears to allow you to specify a different choice for each piece, it can fail if you do so. To reliably create a striped virtual disk, specify either **Disk to partition and partition size** for all pieces or **Name of an existing virtual disk** for all pieces.

5.2.3 Restrictions on vfork Child Process

If a program uses any signal functions in a **vfork(2)** child process, the program might halt. More specifically, when a process calls **vfork** before it calls any signal setting functions (namely **signal**, **sigvec**, **sigaction**, or **sigset**), and then calls the signal setting function within the **vfork** child process, the program will halt. Upon return to the parent process, any use of signals results in indeterminate results, including process termination because of addressing violations.

5.2.4 Restrictions on Mounting /usr from a Server

If your AViiON system mounts **/usr** from a server machine, you will not be able to invoke the following commands: **strerr**, **strclean**, and **strace**. This restriction affects Diskless AViiONs as well as any other AViiONs that mount the **/usr** file system from a server machine.

To execute **strerr**, **strclean**, and **strace**, you need to change their permissions from read and execute for owner to read and execute for owner, group, and other. To make this change, type the following on the server machine:

```
# chown 0555 /usr/sbin/strerr /usr/sbin/strclean /usr/sbin/strace ↵
```

5.2.5 Run newaliases

The **tcpip** setup scripts do not set up the **aliases** files properly. This means that **sendmail** does not have access to local aliases defined in the **/etc/aliases** file, which may result in incorrectly-delivered mail. Run the **newaliases** command to make local aliases available. For more information, see the **newaliases(1m)** and **aliases(4)** manual pages.

5.2.6 Streams Resource Configuration

When configuring a system for STREAMS resource usage, we strongly recommend that you configure the value of the kernel parameter `NQUEUE` to be the number of expected remote users multiplied by 40. As pointed out in *Analyzing DG/UX System Performance (093-701129-01)*, you should raise the value of `NQUEUE` from its default if a large number of network connections must be made to or from your system. There is no performance penalty for having a high value for `NQUEUE`.

When a system reaches the system limit of allocated STREAM queue pairs (the `NQUEUE` value), an error message is reported on the current operation. However, the system may also panic with the following panic code:

```
42000036 (SFM_PANIC_FINAL_CLOSE_NOT_SERIAL_OP)
```

Having the proper value for the `NQUEUE` kernel parameter avoids the problem.

Some network layered products change the value of `NQUEUE` in their delivered system configuration file. Please ensure that after installing such a product, the value for `NQUEUE` is appropriate for your system before you rebuild the kernel.

6 Documentation

6.1 Printed Documentation

The following manuals have been revised for DG/UX System 5.4 Release 3.00.

6.1.1 System Administration Manuals

- *Installation Roadmap -- A Guide to AViiON Systems Installation Documentation (069-701084-04)*
- *Guide to AViiON and DG/UX System Documentation (069-701085-04)*
- *Installing the DG/UX System (093-701087-05)*
- *Managing the DG/UX System (093-701088-04)*
- *Customizing the DG/UX System (093-701101-04)*
- *Managing TCP/IP on the DG/UX System (093-701051-06)*
- *Analyzing DG/UX System Performance (093-701129-01)*
- *Legato NetWorker Administrator's Guide (069-100495-01)*
- *Legato NetWorker User's Guide (069-100496-01)*

6.1.2 Programming Manuals

- *Programming with Perl* (069-100513-00)
- *Programming with TCP/IP on the DG/UX System* (093-701024-03)
- *POSIX.1 Conformance Document for the DG/UX System* (069-701078-03)
- *Programming in the DG/UX Kernel Environment* (093-701083-04)
- *Porting and Developing Applications for the DG/UX System* (069-701059-05)
- *Using the DG/UX Kernel Debugger* (093-701075-03)

6.1.3 New Manuals for DG/UX System 5.4 Release 3.00

- *Installing and Configuring Printers on the DG/UX System* (09-701132-00)
 This manual describes how to install and configure printers on the DG/UX system. It provides instructions for connecting cables, selecting the proper stty options and emulation modes, and troubleshooting printer problems.
- *Achieving High Availability on AViiON Systems* (093-701133-00)
 This manual introduces the concept of High Availability (HA) to DG/UX administrators. It describes the hardware and software components of HA, how they work together, and provides typical scenarios for achieving HA.
- *Managing Modems and UUCP on the DG/UX System* (069-000698-00)
 This manual describes installing and setting up a modem that is compatible with the Hayes Smartmodem™ in the DG/UX environment.

6.2 On-line Documentation

The following on-line manual pages were added, modified, or deleted.

6.2.1 New Manual Pages

Following are the DG/UX manual pages added for DG/UX System 5.4 Release 3.00:

a2p(1)	lmhostid(1M)	pthread_yield(3T)
adapter_driver(3K)	lmremove(1M)	rand_r(3T)
adapter_manager(3K)	lmreread(1M)	readdir_r(3T)
admfailoverip(1M)	lmstat(1M)	sc_panic(3K)
admiopath(1M)	lmver(1M)	sem_close(2)
admpf(1M)	lock_management(3K)	sem_getvalue(2)

admstrlog(1M)	memory_allocation(3K)	sem_open(2)
admsyslog(1M)	misc_format_line(3K)	sem_post(2)
admuserlicense(1M)	nfc(1M)	sem_trywait(2)
admvdisk(1M)	nsar(1M)	sem_wait(2)
audio8(7)	pm_is_super_user(3K)	server_messages(3K)
create_license(1M)	process_management(3K)	sigwait(2)
ctermid_r(3T)	pthread_attr_init(3T)	status_code_macros(3K)
ctime_r(3T)	pthread_attr_setscope(3T)	strtok_r(3T)
data_access_validation(3K)	pthread_cancel(3T)	su_str(3K)
device_driver(3K)	pthread_cleanup_push(3T)	sysconfig(1M)
dg_lmgrd(1M)	pthread_cond_init(3T)	system_clock(3K)
dg_lwp_info(2)	pthread_cond_signal(3T)	system_configuration(3K)
dgpdp(3C)	pthread_cond_wait(3T)	tcpdump(1M)
event_counters(3K)	pthread_condattr(3T)	threads(5)
flockfile(3T)	pthread_create(3T)	ttyname_r(3T)
fs_check_self_id(3K)	pthread_detach(3T)	userlicense.dat(4M)
getc_unlocked(3T)	pthread_equal(3T)	usermond(1M)
getgrgid_r(3T)	pthread_exit(3T)	vdm(7)
getlogin_r(3T)	pthread_getschedattr(3T)	vdmaggr(7)
interrupt_management(3K)	pthread_join(3T)	vdmcache(7)
io_buffer_vector(3K)	pthread_key_create(3T)	vdmdummy(7)
io_hex_str_to_int(3K)	pthread_kill(3T)	vdmirr(7)
io_nodvice(3K)	pthread_mutex_init(3T)	vdmphys(7)
io_select(3K)	pthread_mutex_lock(3T)	vdmremap(7)
iopath.params(4M)	pthread_mutexattr(3T)	vdmtest(7)
izscd(7)	pthread_once(3T)	wdt(7)
kmi(7)	pthread_self(3T)	nps(1)
lmdown(1M)	pthread_setintr(3T)	scsi(7)
lmgrd(1M)	pthread_setspecific(3T)	modem(5)
gcore(1)	perl(1)	
uucp(5)		

Following are the TCP/IP manual pages added for DG/UX System 5.4 Release 3.00:

admbootpclient(1M)	ripquery(1M)	trapsend(1M)
admntp(1M)	snmpgetmany(1M)	xntpd(1M)
admslipdial(1M)	snmpgetnext(1M)	xntpdc(1M)
admslipuser(1M)	snmpgetone(1M)	gated-config(4M)
bootpd(1M)	snmpgettab(1M)	slipdialinfo(4M)
gated(1M)	snmpsetany(1M)	slipusers(4M)
gdc(1M)	traceroute(1M)	ftp_xfer(3N)
ntpdate(1M)	traprecv(1M)	slipd(1C)
ntp(1M)	ospf_monitor(1M)	admtcpipparams(1M)
ifconf(1M)	tcpip.params(4M)	

The new kernel manual pages in `/usr/release/dgux.man_5.4R3.00.fi` are those whose names end in `.3k.Z`.

6.2.2 Deleted Manual Pages

Following are the DG/UX and TCP/IP manual pages deleted for DG/UX System 5.4 Release 3.00:

admindisk(1M)	sem_trylock(2)	getnext(1M)
admmdisk(1M)	sem_unlock(2)	getone(1M)
diskman(1M)	dg_telnetd(1M)	trap_rcv(1M)
osysadm(1M)	getmany(1M)	trap_send(1M)
sem_lock(2)		

For a complete list of DG/UX System 5.4 Release 3.00 manual pages, see the following files:

Package	Pathname
DG/UX	/usr/release/dgux.man_5.4R3.00.fl
gcc	/usr/release/gcc.man_2.4.5.6.fl
TCP/IP	/usr/release/tcpip.man_5.4R3.00.fl
ONC	/usr/release/onc.man_5.4R3.00.fl
NFS	/usr/release/nfs.man_5.4R3.00.fl
X11	/usr/opt/X11/release/X11.man_5.4R3.00.fl
X11 SDE	/usr/opt/X11/release/X11.sde.man_5.4R3.00.fl
NetWorker	/usr/opt/networker/release/networker.man_1.0.fl
X.desktop	/usr/opt/xdt/release/xdt_3.5.fl

6.3 Manual Page Compression Format

The DG/UX System 5.4 Release 3.00 manual pages are compressed using **compress**, which does a more efficient job than **pack**, which was used previously. Because of this change, any Release 3.00 manual page copied to an earlier system will not be recognized by the **man** command. The old **man** command recognized only uncompressed manual pages and those compressed with **pack**.

6.4 Related Release Notices

The following release notices are on the DG/UX tape (Model P001A) in other packages:

Package	Pathname
gcc	/usr/release/gcc_2.4.5.6.rn
X11	/usr/opt/X11/release/x11_5.4R3.00.rn
X.desktop	/usr/opt/xdt/release/xdt_5.4R3.00.rn
NetWorker	/usr/opt/networker/release/networker_1.0.rn

7 Software Distribution

7.1 Media

DG/UX System 5.4 Release 3.00 is released on QIC 150 tapes and compact disks (CD). This release is available via the following model numbers:

P001A The *DG/UX Operating System with X Window System* package. The part numbers are: 079-600353-01 for the tape and 068-600080-01 for CD.

Q001A The *DG/UX Operating System* package. The part numbers are: 079-600352-01 for the tape and 068-600079-01 for CD.

7.2 Files

A list of files that are loaded when the DG/UX system is loaded on your disk is in the file `/usr/release/dgux_5.4R3.00.fl`. The names of the file-list files in `/usr/release` for other packages documented in this release notice are as follows:

`tcpip_5.4R3.00.fl`
`onc_5.4R3.00.fl`
`nfs_5.4R3.00.fl`

See the related release notices described above for details on the `gcc` and NetWorker file lists. If you are installing the workstation model (P001A), see the DG/UX X Window System and X.desktop release notices for the names of other on-line files listing the contents of that model.

8 Installation Instructions

The primary instructions for installing DG/UX System 5.4 Release 3.00 from the release medium (tape or CD-ROM) and from remote tape or file system are contained in the manual *Installing the DG/UX System* (093-701087-05). Additional setup information is provided in *Customizing the DG/UX System* (093-701101-04). For full explanations of the `sysadm` installation and management program and other management topics, see *Managing the DG/UX System* (093-701088-04).

Before beginning installation, see Section 3, "Environment," for a list of the supported hardware environments and section 5.2.2 for information about swap area and debugging. The minimum memory requirement has increased from 12MB to 16MB.

Note that during the installation process, all existing logical disks will be converted into virtual disks.

8.1 Keeping Your Previous Customizations when Upgrading

When building a new kernel during the upgrade procedure, to keep your previous system customizations, make sure you supply the name of your old system configuration file. Verify your system file name when prompted:

```
System configuration file name: [moe] ↵
```

The default, usually your host's name (**moe**, for example), is appended to the prefix, **system.**, to yield **system.moe** as the system configuration file name upon which the kernel itself is based. Verify the system configuration file name before answering this prompt. The correct (previously-created) system file should be supplied by default. If you supply a new name to this prompt, you will create a new system configuration file and kernel, which will not include your previous customizations.

You can verify your selection through the confirmation:

```
[system.moe] Correct? [yes] ↵
```

If, by accident, you create (and boot) a kernel that is based on a new system configuration file, you can recover by using **sysadm** to repeat the kernel building and booting procedure, this time specifying the correct (old) system file. See **/usr/src/uts/aviion/Build** for a list of system configuration files.

Refer to the “Kernel” section of Section 4, “Enhancements and Changes” for related information.

8.2 User Default License

DG/UX System 5.4 Release 3.00 has a default license for two users. At installation time, you should upgrade this license to the appropriate number of users. Installation materials include a page containing the *license token* necessary to install this proper license and instructions on how to install the license. The **admuserlicense** command or the **System->License sysadm** menu can be used to perform the upgrade.

9 Preparing a Software Trouble Report (STR)

If you believe that you have found an error in the DG/UX system or its documentation or if you have a suggestion for enhancing or improving the product, do the following:

- Describe the error or suggestion using a Data General Software Trouble Report (STR) form. See “Describing Problems” below. For software problems, prepare a tape containing diagnostic information and problem illustrations. See “Preparing STR Tapes” below.
- To simply make a comment, edit a copy of the Reader Comment form in **usr/release/doc_comment_form**, and send it to the address contained on the form.

A copy of the standard STR form suitable for printing on a line printer is in the file `/usr/release/STR_form`. STR forms are also available from the nearest Data General office, from your representative, or from the Customer Support Center. Depending on your support contract, you can submit STR reports to the Customer Support Center or directly to your Data General representative.

9.1 Describing Problems

For each problem you report, fill in the three-page STR form completely. Where asked for the release, enter `5.4R3.00`. For the name and model number, enter one of the following selections:

- TCP/IP and C001A for TCP/IP problems.

For assistance in diagnosing, describing, and solving TCP/IP problems, see *Managing TCP/IP on the DG/UX System* (093-701051-06).

- NFS and C002A for NFS problems.
- DG/UX and Q501A for other DG/UX problems.

Also, indicate the type of STR you are writing:

- An enhancement request: Describe it clearly and explain why you want it.
- A documentation error: Identify the document—the manual page name, the document title, ordering number, the page number(s) and section(s) containing the error; and explain the error.
- A software problem: Clearly describe the problem, indicating how to reproduce it.

If you are reporting a problem that involves or results in a system panic or hang, your STR form *must* be accompanied by a tape containing:

- In file 0, a system memory dump
- In file 1, a `cpio` copy of your tailored system file (usually named `/dgux`)
- Optionally, in file 2, a `cpio` copy of other files (logs, program source and executable files, shell scripts, and so on) that cause, illustrate, or document the problem.

An accompanying tape is also essential for problems caused by programs not supplied by Data General. Include a copy of the executable program file *and* the source file.

If your STR includes a tape, indicate in the problem description the tape's organization and contents.

9.1.1 System Halts or Hangs

If the system hangs perform the following steps:

- If you are running on an M88110 platform, an “abort” button is recessed in the back panel of the system near the “Warm reset” button. Press the “abort” button. The system should generate a panic and prompt you for a dump. If it does not, use the “reset” button as described later in this section.
- If you are running on a M88100 platform, try the “hot key” sequence. While holding down **Ctrl**, alternately press [and] until the DG/UX System Panic, Panic Code 30000027 appears. The system prompts you for a dump.
- If neither of these procedures work, reset your machine using the “reset” switch or button. Once in the SCM, enter “**reset**” followed by “**start 1000**”. This invokes the dump dialog, described in “Preparing STR Tapes” below.

Note that on some systems, especially MC88100 based, the **start 1000** command from the PROM SCM prompt may take two or three minutes to respond with the prompt, **Do you want to take a dump?**

If the **start 1000** command seems hung, you should wait three minutes before attempting another reset, since it may proceed properly in time.

You should always perform the above actions in order before cycling the power on your system. If your system has the watchdog timer enabled, refer to Section 4.2.1.4 of this notice.

9.1.2 Incorrect Behavior by a DG/UX System Call or Device Driver

Write the *shortest possible* program that demonstrates the problem. This can be a shortened version of your actual program, without unnecessary variables and logic. Supply both source and executable copies of this program and a copy of the system image (*/dgux*), along with a detailed description of how to reproduce the problem.

9.1.3 Incorrect Behavior by a DG/UX Command

Write the *shortest possible* shell script that demonstrates the problem. Supply a copy of this script with a detailed description of how to reproduce the problem. Before submitting the problem, be sure that it can be duplicated under “sterile” conditions; for example, with a standard path, standard permissions on files, and so on. In other words, try to ensure that your environment is not the cause of the problem.

9.1.4 Fast-Recovery File System

If you are filing an STR about the fast-recovery file system, include */etc/log/fast_fsck.log* on the STR tape. Include also on the tape all output produced by the fast-recovery or normal run of **fsck** that reported problems. When **fsck** is run as part of an **init** script, that output is placed in */etc/log/fsck.log*. If you run **fsck** manually, get this information for the STR tape by redirecting

output to a file or piping it through `tee`.

9.2 Preparing STR Tapes

In most cases, a software error report should include both a filled-in STR form and an STR tape. While Data General wants to know of any problems you may encounter, those reported without accompanying diagnostic information are very difficult to solve.

There are two types of system error: those associated with system panics or hangs (type 1) and those having to do with failure of specific programs or components only (type 2). On a type 1 STR tape, you *must* include a system memory dump (tape file 0) and a copy of the system image (tape file 1); you may also include other requested files (tape file 2). A type 2 STR tape should include only the latter: program source, executable, and output files that demonstrate the problem, in tape file 0.

Please indicate on the problem description form whether or not you are providing a tape and, if so, how the tape is organized. Also, be sure to record the panic code.

The behavior of your system after a panic depends on the following two configurable system parameters:

- The default dump device, which may be the name of a tape drive, a virtual disk, or a network controller.
- The default behavior, which may be one of the following:
 - Prompt the operator to confirm whether memory should be dumped, and to which device. (This, the default, is the usual choice when the default dump device is a tape drive.)
 - Dump memory to the default dump device without asking. (This is a good choice when the dump device is a virtual disk.)
 - Do not dump memory.

Data General does not recommend that you dump memory. However, to force a memory dump after a panic when this option is set, press the reset button, then boot with the `-d` option to initiate the system dump dialog.

9.2.1 The System Dump Dialog

When the default behavior, prompt the operator, is set, a dialog ensues after a panic at the system console. The dialog differs slightly depending on whether the dump device is a tape drive, a virtual disk, or a network device. At each prompt, you can accept the displayed default or enter an overriding value.

To the prompt that asks whether to limit the dump to kernel memory, take the default (yes) unless you are told otherwise by a Data General representative. A full memory dump is seldom required.

9.2.1.1 Dump to Tape

If the default dump device is the local tape drive `st(insc(),4)`, a dialog like the following ensues:

```
Do you want to take a system dump? [Y] ↵
Limit the main memory dump to kernel memory only (no user memory)? [Y] ↵
Dump destination device? st(insc(),4) ↵
Mount tape. Type newline when tape is ready. ↵
```

If the system memory image is too large to fit on one tape, these prompts appear:

```
Tape volume 1 completed.
Mount tape. Type newline when tape is ready. ↵
```

If any tape volume write fails, you can restart the dump at the beginning of that volume:

```
Hard error on tape volume 1. Restarting volume from checkpoint.
Mount tape. Type newline when tape is ready. ↵
```

By default (you can change the default with `dg_sysctl`) the final volume rewinds upon completion of the dump:

```
System dump completed successfully.
```

Before submitting this tape with an STR, append a copy of your system file and other relevant material. See “STR Tape Examples” below for illustrations.

9.2.1.2 Dump to Virtual Disk

If the default dump device is a virtual disk named `/dev/rdisk/sys_dump` on the physical disk `sd(insc(),0)`, this dialog occurs:

```
Do you want to take a system dump? [Y] ↵
Limit the main memory dump to kernel memory only (no user memory)? [Y] ↵
Dump destination device? vdm_dump(sd(insc(),0),sys_dump) ↵
```

To submit an STR, transfer the memory dump to tape with the `-t` option of the `lsd` command, append a copy of your system file, and append other relevant material. See “STR Tape Examples” below for illustrations.

9.2.1.3 Dump to a Network Device

If you are at a diskless workstation, a dialog like the following occurs after a panic:

```
Do you want to take a system dump? [Y] ↵
Limit the main memory dump to kernel memory only (no user memory)? [Y] ↵
Dump destination device? inen() ↵
```


Only the workstation's communications controller name, **inen** or **dgen**, will work as the dump device. If any other value appears, override it.

The memory dump is written to a file on the server's disk specified in the server's **/etc/bootparams** file. Then the workstation boots from the boot path specified in **/etc/bootparams**.

If you encounter a problem, such as insufficient file space or incorrect permissions on the server, go to the server and fix the problem. Then come back to the workstation and boot, specifying the option that first initiates the dump dialog:

```
SCM> b inen() -d ↵
```

To submit an STR, transfer the memory dump to tape and append a copy of the system file. (Sample commands are given in the next section.)

9.3 STR Tape Examples

9.3.1 Copy a Panic Dump to Tape

If you wish, you can specify a virtual disk as the default system dump device. If you have selected this option, the first step in preparing an STR tape (for a problem connected with a system panic) is to transfer, from the virtual disk to tape, the system memory image that was dumped after the panic. The command below transfers a memory dump from a virtual disk named **/dev/rdisk/sys_dump** to the tape in drive **/dev/rmt/0n**:

```
# lsd -t /dev/rdisk/sys_dump /dev/rmt/0n ↵
```

The memory image is copied to file 0 of the tape. Note that the no-rewind option (**n**) is specified for the drive. This is so that the tape will be ready for writing file 1. You can skip the next section and go to "Copy Your Tailored System File."

9.3.2 Position to the End of a Dump Tape

Your system may be set up so that, in the event of a panic, it copies a memory dump to file 0 of a tape in a specified drive. If so, you prepare an STR tape by appending the system image to the tape containing the panic dump. But first, to prevent overwriting file 0 on the tape, make sure the write head is located at the end of file 0 and ready to write file 1 (next step). Assuming the tape drive is named **/dev/rmt/0n**, the command is:

```
# mt -f /dev/rmt/0n fsf ↵
```

This step may not be necessary, but it will do no harm. (By placing a semicolon after the command above, you can combine this step with the next step.)

9.3.3 Copy Your Tailored System File

In all cases, an STR tape that includes a memory dump image must also include a copy of your tailored system file. Without an image of the operating system that was running when the panic occurred, a memory dump is useless for diagnosis. The following command transfers the system file named `dgux` to tape drive `/dev/rmt/0n`:

```
# cd /; echo dgux | cpio -oBcv > /dev/rmt/0n ↵
```

The system file is copied to file 1. The tape, because it was not rewound, is ready for writing any additional files to file 2.

9.3.4 Copy Additional Diagnostic Files

For copying any additional files to the STR tape, use a command line like the following:

```
# ls filenames | cpio -oBcv > /dev/rmt/0 ↵
```

This example assumes `filenames` is a template for files in the current directory that you wish to copy. The tape rewinds after this command line. It is now ready for forwarding to Data General with the STR and problem description forms.

9.3.5 Preparing an STR Tape for a Diskless Workstation

The following sample commands make an STR tape for a diskless workstation:

```
# cd /srv/dump ↵
# ls diskless_client_name | cpio -oBcv >/dev/rmt/0n ↵
# cd /srv/release/PRIMARY/root/_Kernels ↵
# ls dgux.diskless | cpio -ocvB >/dev/rmt/0n ↵
```

These commands: (1) change to the server directory containing the memory dump; (2) transfer the memory dump to file 0, without rewind; (3) change to the server directory containing the tailored system file; (4) transfer the system file to file 1. The tape is set for copying additional files to tape file 2.

The actual pathname of the memory dump file for a workstation—`/srv/dump/diskless_client_name` in this example—is specified in the file `/etc/bootparams` on the workstation's server. Also, the name of the custom kernel may be something other than `dgux.diskless`.

The commands above assume you are logged into the server. If you are instead logged into the workstation, make sure `/srv/dump` is mounted, issue commands (1) and (2), and substitute the following commands for (3) and (4):

```
# cd / ↵
# ls dgux | cpio -ocvB >/dev/rmt/0 ↵
```

9.3.6 Preparing an STR Tape Without a Memory Dump

For problems that do not involve a system dump, put all files associated with the problem in tape file 0, in **cpio** format, using a command like the following:

```
# ls filenames | cpio -oBcv > /dev/rmt/0 ↵
```

Do not use absolute pathnames (pathnames starting with **/**) for the filenames being listed.

If you need to include a directory such as **/var/spool/lp** that also includes links, change to that directory and use the **tar** command. For example, **tar cvfL /dev/rmt/0 ***

10 Index

A

Archive utilities 10, 36
Availability, high 8, 13

B

Boot Protocol 27

C

C compiler 7
 man pages 50
 release notice 50
Caching, off-line persistent 42
Caching, on-line non-persistent 41
CD-ROM 41, 51
Commands 33
Compiler 9
Compiler, GNU C 7
 man pages 50
 release notice 50
Compilers 33

D

Debuggers 9, 32
Device address reservation functions 20
Device driver
 building 20
DG/UX System 5.4
 installing Release 3.00 51
 man pages 50
 Release 3.00 environment 10
 Reporting problems 52
dg_process_info 9
dgux.installer 20
dgux.installer.diskless 20
Diagnostics tools 25
Disk, magnetic 41
Disk, virtual 17, 18, 41, 51
diskman 18
Distribution of software
 available media 51
 compact disks 51

Distribution of software (*cont.*)

 QIC 150 tape 51
 Documentation, *see* Installation documentation;
 Kernel programming documentation
 list of new manual pages 48
 revised manuals 47
 DOS and ISO-9660 CD-ROM file system
 drivers 21
 Dump
 system 55

E

ENCODE_STATUS macro 20
Enhanced dg_process_info 21
Errno and status code macros 20
Errors in the DG/UX system or documentation,
 reporting 52

F

Failover 8, 15

G

GNU C compiler 7
 man pages 50
 release notice 50

H

High availability 8, 13

I

i_sc.h file, *see* status_code.h file
Installation, *see* DG/UX System 5.4, installing
 Release 3.00
Installation, software 40
Installer kernels 20
io_check_device_spec function 20
io_forget_device_spec function 20
IP Takeover 8, 14

K

Kernel 8, 19, 52
Kernel source code interface change 20

L

libc 43
libc.so.1 43

M

Magnetic Disk 41
Medium Term Scheduler 30
Memory dump
 from diskless workstation 56
 to disk 56
 to tape 56
Memory requirement 7
Motif 7
Multi-Path LAN I/O 8, 14
Multiple VME Chassis 8, 13

N

Network Time Protocol 9, 27
NetWorker
 man pages 50
 release notice 50
NFS 7
 man pages 50
 Preparing STR 53
NFS Failover 8, 15
NO_ERRNO macro 20

O

Off-line persistent caching 42
On-line non-persistent caching 41
On-line storage management 8, 15, 17
ONC 7
 man pages 50
OSF/Motif 7

P

Panic dump
 to disk 56
 to tape 56
POSIX Threads 9, 30

Preparing a Software Trouble Report (STR) 52
Printer service enhancements 10
Printer services 38
Pthreads 9, 30

R

Real-time extension support 9, 21
Reentrant user libraries 9, 31
Reporting errors in the DG/UX system or
 documentation 52

S

SC_ENCODE_STATUS macro, *see*
 ENCODE_STATUS macro
SC_NO_ERRNO macro, *see* NO_ERRNO
 macro
scsi 8, 21
Serial Line IP 9, 27
Shared object 43
Short addressing 9, 33
Simple Network Management Protocol 9, 23
Software distribution
 available media 51
 compact disks 51
 QIC 150 tape 51
Software Installation 10, 40
Software Status Register 33
Software Trouble Report (STR)
 re fast-recovery file system 54
Software Trouble Report (STR)
 dumping to network device 56
 dumping to tape 56
 dumping to virtual disk 56
 preparing tape 55
 preparing tape for diskless workstation 58
 preparing tape without memory dump 59
Software Trouble Report(STR)
 preparing 52
Standards 38
status_code.h file 20
STR, *see* Software Trouble Report (STR)
STREAMS 47
STREAMS module
 building 20
syadm 52
Sync jumpering 43
sysadm 8, 9, 10, 29, 40, 45, 51
System dump 55

T

Tape

- STR examples 57
- TCP/IP 7, 9, 23
 - man pages 50
 - Preparing STR 53
- TCPIP 29
- Thread cancellation 31
- Threads 9, 30

U

- uc_device_class_enum_type type 20
- User default license 10, 52
- Utilities, archive 10, 36

V

- VDM Cache 18
- VDM mirror performance 19
- vfork 46
- Virtual disk 17, 18, 41, 51

W

- Watchdog Timer 8, 15

X

- X Window System 7, 12, 51
 - man pages 50
 - release notice 50
- X.desktop System 7, 12
 - man pages 50
 - release notice 50
- X/Open XPG3 branding 38

End of DG/UX Release Notice