Maintenance Update Notice:

DG/UX® System

Release 4.20MU05

078–600094–04

This update notice applies to the following model:
P001A

Notice

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

Restricted Rights: Use, duplication, 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 of Defense Federal Acquisition Regulation (DFAR) 252.227–7013 and in subparagraphs (a) through (d) of the Commercial Computer Software Restricted Rights clause of Federal Acquisition Regulation (FAR) 52.227–19, whichever may apply.

Data General Corporation
4400 Computer Drive
Westboro, MA 01580

AViiON, CLARiiON, and DG/UX are U.S. registered trademarks of Data General Corporation. UNIX is a U.S. registered trademark of Novell, Inc. X Window System is a trademark of the Massachusetts Institute of Technology.

All other trademarks mentioned herein are the property of their respective owners.

Maintenance Update Notice:
DG/UX® System

Release 4.20MU05

078–600094–04

Revision History: Effective with:

Original Release – December 1997   DG/UX® System Release 4.20MU01
Revision 01   – May 1998   DG/UX® System Release 4.20MU02
Revision 02   – September 1998   DG/UX® System Release 4.20MU03
Revision 03   – March 1999   DG/UX® System Release 4.20MU04
Revision 04   – August 1999   DG/UX® System Release 4.20MU05

Contents

7 Software distribution 55

8 Installation instructions 55

9 Preparing a software trouble report 58

1   Introduction

This update notice describes Maintenance Update 05 for DG/UX® System Release 4.20. A maintenance update contains product enhancements completed since the last release.

This update notice supplements information provided in the DG/UX® System documentation. It contains information from the R4.20 Release Notice to the current update notice.

An on–line version of this update notice suitable for line printers is in the file **/usr/release/dgux_R4.20MU05.un**. A PostScript® version is in the file **/usr/release/dgux_R4.20MU05.un.ps**.

Revision bars (vertical lines in the margin) in the typeset version identify changes that take precedence over the on–line version.

> **IMPORTANT**   To improve the ease of use of this notice, its content and organization differ from previously published revisions. Information about permanent product features was moved into the appropriate manuals. Some sections have fewer subsections since information from previous releases was consolidated into one section.

2   Product description

This product is a maintenance update (MU) for the DG/UX® System, one of three versions of the operating system for Data General's family of AViiON® computers:

*   DG/UX System (generic)

*   C2 Trusted DG/UX Option

*   DG/UX B2 Security Option

Each product has its own Release Notice and update notice; however, all relevant information in the generic DG/UX system notices is replicated in the B2 and C2 notices.

2.1 Releases provided on your CD–ROM

This maintenance update is provided on CD–ROM. The CD–ROM contains two different versions of the product that you can load on your system:

* DG/UX System R4.20 and R4.20MU05, which contains a full release of R4.20 and all maintenance updates through R4.20MU05.

* DG/UX R4.20MU05, which contains the maintenance update only.

During the installation process, you are asked to specify which release you want to load.

2.2 Application Capture Option (ACO)

The DG/UX system includes an Application Capture Option (ACO). The **dgux.aco** package enables certain UnixWare and SCO applications to run on the DG/UX system for Intel without being recompiled. You need to load the **sdk** package to use some features of ACO, and although it is on the release medium, you must specify the package suite name to install it. For more information, see the Application Capture Option update notice.

> **IMPORTANT** If you accidentally load ACO with the DG/UX system, the only clean way to remove it is to reload the DG/UX system without the **dgux.aco** package.

2.3 Software packages provided on your CD–ROM

The software packages provided on the compact disk are for the generic, C2, and B2 versions of the DG/UX system. This update notice describes those packages relevant to the generic DG/UX System.

dgux   The DG/UX operating system, containing system–level programs, facilities, and commands.

dgux.acl   The Access Control List package adds flexibility to the regular DG/UX file permissions. An ACL applied to an object can allow or deny read, write, and/or execute/search access either to specific user(s), specific group(s), or both. Permissions can be specified for up to 2000 different users and/or groups.

dgux.aco   The Application Capture Option package, that contains commands, files, and libraries, needed to run some UnixWare or SCO applications on Intel–based DG/UX systems.

networker The Legato Data Backup Utility (previously called Legato NetWorker product for the DG/UX system), a data backup utility. This utility lets you back up local disks to local tape drives or to another NetWorker server. If you install **networker**, you must install **tcpip** also.

nfs   The Network File System (NFS), a service that lets many users share file systems over a network. NFS is Data General's port of SunOS version 4.0 ONC(Trademark)/NFS to its family of AViiON computers. In addition to NFS, this service includes the Lock Manager and Status Manager. If you install **nfs**, you must install **tcpip** also.

onc   The ONC (Open Network Computing) programming platform, Remote Procedure Call (RPC), and External Data Representation (XDR), for writing heterogeneous distributed applications. This package also contains application–level services based on RPC/XDR. Network Information Service (NIS) is one such service. If you install **onc**, you must install **tcpip** also.

sdk   Software Development Kit (SDK), a necessary component for software development on the DG/UX system. It includes many important application development components, miscellaneous programming tools, and most system headers.

SecurID   The SecurID ACE/Client software is shipped with the DG/UX system; it allows a DG/UX system to become a client of SecurID ACE/Server software, which runs on a variety of operating systems as a single point of user authentication on a TCP/IP network. The SecurID software was developed as a solution for enterprise–wide security by Security Dynamics Technologies, Inc.

tcpip   A package of communications software that implements the TCP/IP family of network protocols on the DG/UX system.

X11   The X Window System, a networked window management system that allows your AViiON computer to operate in a networked environment with workstations from other vendors. It includes OSF/Motif.

X11.aco   The X11 portion of the Application Capture Option package. It contains libraries and links needed to run Graphical User Interface (GUI) UnixWare or SCO applications on the DG/UX system.

2.4 Manual page packages

Each of the software packages described in the preceding section has a corresponding manual (man) page package. Each manual page package has a **.man** suffix, for example, **dgux.man**. If a manual page package is absent in this maintenance update, it is because no manual pages in that package were added or changed.

If you load packages individually, be sure to load the corresponding manual page package. If you load packages by package suite, the corresponding manual page packages, if present, will be loaded accordingly.

2.5 Software package suites

The software packages described above are bundled into package suites for convenient loading. Your CD–ROM contains suites for the generic, C2, and B2 versions of the DG/UX system. The suites relevant to the generic version are as follows:

acl – Contains the DG/UX Access Control Lists packages: dgus.acl and dgux.acl.man.

aco– Contains only the Application Capture Option packages: **dgux.aco**, **dgux.aco.man**, **X11.aco**, and **X11.aco.man**. You can load this package by itself to add the ACO software to a system previously loaded without it.

development – Contains all of the DG/UX software packages listed in Section 2.5, plus all corresponding manual page packages that are present, except for the packages contained in the **acl suite**.

development_less_aco– Contains all of the **development** software packages except   **dgux.aco** and **dgux.aco.man**.

production – Contains all of the same basic DG/UX operating system software packages included in the **development** suite except the SDK software, **sdk**, **sdk.man**, **sdk.X11**, and **sdk.X11.man**.

production_less_aco– Contains all of the **production** software packages except **dgux.aco** and **dgux.aco.man**.

sdk_all– Contains only the Software Development Kit(SDK) packages: **sdk**, **sdk.man**, and **sdk.X11.man**.

2.6 NetWorker® Support Plus

NOTICE: TO ALL DG/UX AND NETWORKER CUSTOMERS

NetWorker Support Plus Service is required in order to receive continuous timely updates of licensed Legato NetWorker on the DG/UX System. This service is separate from and in addition to DG/UX Software Subscription Service(SSS).

The latest version of NetWorker products does not offer SSS. Support Plus or upgrade licenses are the only way to get the newest revision of the NetWorker products (v5.1).

This does not affect the Data Backup Utility (DBU), which is the unlicensed, restricted version of Legato NetWorker that is bundled with each new DG/UX System release. However, it does affect all licensed features of NetWorker.

Please contact your Data General sales representative for details about the NetWorker Support Plus Service.

3   Environment

This section describes the minimum hardware, firmware, and software required to run this maintenance update.

3.1 Hardware

This maintenance update requires one of the following Intel systems: AV 1600, AV 2000, AV 2100, AV 2600, AV 3000, AV 3600, AV 3600R, AV 3650, AV 3650R, AV 3700, AV 3700R, AV 4700, AV 4800, AV 4900, AV 5800, AV 5900, AV 6600, AV 20000, or AV 25000.

As with Release 4.20, the minimum equipment configuration is as follows:

*   32 MB of memory

*   CD–ROM drive

*   diskette drive

*   System console (or graphics monitor for workstations)

*   500 MB of hard disk space

This update supports only Intel–based AViiON computers. It does not support Motorola 88000–series processors. The 88K systems are supported by maintenance updates to DG/UX System R4.11.

3.2 Firmware

For proper operation, Release 4.20MU05 of the DG/UX system for Intel requires firmware revision numbers equal to or greater than those listed below.

AV 1600: BIOS 1.00.03, SCU 3.60 ovl 1.10
AV 2000: BIOS 1.00.07, SCU 3.08ovl 31
AV 2100: BIOS 1.00.04, ECU/SCU 3.60
AV 2600: BIOS 1.00.04, SCU 3.60
AV 3000: BIOS 1.00.03, SCU 3.11ovl39
AV 3600: BIOS 5.10.5, TSB 6.8.0, ECU 3.12
AV 3600R: BIOS 5.10.5, TSB 6.8.0, ECU 3.12
AV 3650: BIOS 5.10.8, TSB 8.9.0, ECU 3.17
AV 3650R: BIOS 5.10.8, TSB 8.9.0, ECU 3.17
AV 3700: BIOS 2, SSU 1.0
AV 3700R: BIOS 2, SSU 1.0
AV 4700: BIOS 5, SCM86 2.6.0
AV 4800: BIOS 4.04, SCM86 2.6.0
AV 4900: BIOS 1.00.05, TSB 5.3.0, SCU 350–140
AV 5800: BIOS 5, SCM86 2.6.0
AV 5900: BIOS 1.00.05, TSB 5.3.0, SCU 350–140
AV 6600: BIOS 5.10.6, TSB 7.10.0, ECU 3.17
AV 20000: BIOS 0.29, VCS 6.61, SMM 1.02, XDIAGS 10.7, M3D 1.02
AV 25000: BIOS Production 12, User Binary 12.00, VCS 7.28, SMM 3.30,
XDIAGS 12.12, M3W 3.30

3.3 Software

This section provides information about the software required to run this
maintenance update.

3.3.1   Configuration constraints

If you do not currently have the DG/UX System running on your system, you must
choose the product from your CD–ROM that contains both R4.20 and this
maintenance update (see Section 2.1). Do not try to install the MU–only release (see
Section 8.1) over any release or maintenance update prior to R4.20. If you are
running Cluster software, you must also upgrade to Cluster Software Release
1.25.The MU–only release cannot be installed on a new system or overloaded on any
R4.11 system.

Section 8 of this update notice, "Installation instructions," discusses package sizes
and disk space requirements. Before you install this update, please read that section.

3.3.2 DG/UX patches in this update

Data General produces and delivers individual patches in response to Software Trouble Reports (STRs). These patches are named and numbered by package. Patch numbers for a product are unique and increase sequentially.

During the support period for a product, it may be necessary to release multiple patches that replace the same file. If a file has more than one patch, the higher–numbered patch supersedes the lower–numbered patch and contains all fixes in it. Therefore when you have two patches for the same file, load the patch with the higher number.

This maintenance update contains the latest patches for the packages listed below:

Package	R4.20	MU01	MU02	MU03	MU04
dgux	p55	p101	*p77 p86	p96	*p54 p55
dgux.acl			p02	p02	
dgux.acl.man			p01		
dgux.aco	p02	p03	p02	p02	
dgux.man	p01	p03	p03	p02	
nfs	p01				p01
onc	p01	p02		p01	
sdk	p06	p08	p05	p07	p05
sdk.man	p02	p03			p02
tcpip	p07	p12	p12	p18	p08
tcpip.man				p01	p01
X11	p01	p03	p02	p02	p02

* indicates that a particular patch is not included in R4.20MU05. On this maintenance update do not load any patches released prior to R4.20MU05.

3.3.3 Using layered software

Some layered products may require a particular DG/UX release. For dependencies, please check your layered product documentation.

4   Enhancements and changes

This section lists the major enhancements and changes in Release 4.20 and subsequent updates. Enhancements are new features. A change may be a software problem that has been fixed or a feature that is modified for functional interoperability.

The section contains the following subsections:

*   New hardware features in R4.20MU05

*   Software enhancements and changes in R4.20MU05

*   Hardware features introduced in previous R4.20 MUs

*   Software enhancements and changes introduced in previous R4.20 MUs

4.1 New hardware features in R4.20MU05

4.1.1   Servers and Processors

4.1.1.1 AV 3700 processor enhancements

R4.20MU05 supports the 500MHz Pentium III Xeon processor on the AV 3700 server.

4.1.1.2 AV 25000 multi–block server enhancements

R4.20MU05 supports larger and faster AV 25000 systems. It supports AV 25000 configurations up to 8 blocks (32 CPUs); the previous maximum configuration was 6 blocks. This release also supports 500MHz and 550MHz Pentium III Xeon processors on AV 25000.

4.1.2   Controllers

4.1.2.1 Model 7487 and 7517 Fibre Channel Adapter Firmware Download

During device configuration, the **epfc** (7) driver checks the firmware and BIOS revisions on the system's Fibre Channel controllers, and automatically downloads the latest revisions if an older revision is detected.

The following revisions are downloaded by the **epfc** (7) driver:
Model 7487 (LP6000) firmware rev FF213A0 and BIOS rev FB120A08
Model 7517xxx (LP8000) firmware rev DD281x3 and BIOS rev DB120A08

4.1.2.2 Model 7495 ATI Graphics Card

R4.20MU05 supports the ATI 3D Pro Turbo+DVD PC2TV graphics card.

4.1.2.3 Model 7516 PCI Quad Channel Ethernet Adapter

R4.20MU05 supports the Model 7516 PCI Quad Channel Ethernet adapter. This model is a 10/100 Mbit adapter with four Ethernet ports. The device operates in both 32–bit and 64–bit PCI slots.

Each port may be independently configured to run using 10BASE–T or 100BASE–TX, either through explicit definition or through autoconfiguration, when connected to the appropriate 10BASE–T or 100BASE–TX Ethernet hub. The Model 7516 also supports ANSI/IEEE Standard 802.3 autonegotiation technology when connected to a device (hub, switch, or point–to–point connection) which also supports autonegotiation.

This device is supported in the AV 3700 and AV 25000 series of AViiONs. However, it is not supported in the AV 20000, AV 3600, AV 3650, AV 6600, AV 4900, or AV 5900 series of AViiON servers due to hardware incompatibility between the host PCI interface and the adapter. For more information, refer to the **apen** (7) on–line manual page.

4.1.2.4 Model 7517OPT and 7517COP Fibre Channel Adapters

R4.20MU05 supports two Model 7517 Fibre Channel Adapters, also known as the Emulex LP8000. Model 7517OPT includes an optical GBIC, and Model 7517COP includes a copper GBIC. These models offer improved performance over the Model 7487 (Emulex LP6000) adapter. Both adapters are supported by the **epfc** (7) driver.

4.1.2.5 Model 7522 Gigabit Ethernet Adapter

R4.20MU05 supports the Model 7522 Packet Engines G–NIC II Plus Gigabit Ethernet LAN adapter, which runs in a 32–bit or 64–bit PCI slot. The adapter supports IEEE 802.3z 1000base–SX connectivity using multi–mode fibre. For more information, see the **pgen**(7) on–line manual page.

4.1.2.6 Model 7523 Single Port Ethernet Adapter

R4.20MU05 supports the Model 7523 PCI Single Channel Ethernet (Intel EtherExpress Pro100+ Server Adapter) LAN adapter. The adapter can be configured to run using 10BASE–T or 100BASE–TX, either through explicit definition or through autoconfiguration, when connected to the appropriate 10BASE–T or 100BASE–TX Ethernet hub. The Model 7523 also supports ANSI/IEEE Standard 802.3 autonegotiation technology when connected to a device (hub, switch, or point–to–point connection) which also has an adapter capable of autonegotiation. For more information, refer to the **ipen**(7) on–line manual page.

4.1.3   Peripherals

4.1.3.1 Model 61040 40x CD ROM

R4.20MU05 supports the 40x CD ROM Model 61040.

4.1.3.2 Model 79589 and 79590 Brocade Fibre Channel Switches

R4.20MU05 supports the model 79589 8–port and model 79590 16–port fibre channel switches. from the Brocade Silkworm 2000 series. They support new features such as zoning, enhanced configurability, and on–line help in the telnet management interface.

Disks

4.1.3.3 New fibre channel and SCSI disks

The following new fibre channel disk is supported in R4.20MU05 :

79021–HFH 36 GB 7200 RPM CLARiiON® FC disk

The following new SCSI disk models are supported in R4.20MU05:

79013–HVH 9 GB 7200 RPM CLARiiON SCSI disk
79014–HVH 18 GB 7200 RPM CLARiiON SCSI disk
61038–SJE 36 GB 7200 RPM SCSI disk
61035A–SJL 9 GB 10K RPM SCSI disk for AV 3700
61035A–SJE 9 GB 10K RPM SCSI disk for CSS4
61036A–SJL 18 GB 10K RPM SCSI disk for AV 3700
61036A–SJE 18 GB 10K RPM SCSI disk for CSS4
61039–SJE 36 GB 10K RPM SCSI disk for CSS4

Tapes

4.1.3.4 Series 5800T DLT8000 Tape Drives

R4.20MU05 supports the Series 5800T DLT8000 tape drives and arrays. These new drives support high density 40 GB tapes; the maximum tape size supported on earlier DLT drives was 35 GB.

CLARiiON® Storage Systems　　　　　　　　　　　　　　　　　　　　　■

4.1.3.5 CLARiiON 5600 and 5700 Disk Arrays

R4.20MU05 supports the CLARiiON 5600 and 5700 full fibre disk arrays.

4.2 New Software enhancements and changes in R4.20MU05

4.2.1　Additional Library symbols in UnixWare version 7

Additional library symbols added for this release include new interfaces in the socket libraries (libsocket and libstreamsocket), as well as miscellaneous interfaces in the standard C and networking libraries (libc and libnsl). Certain third–party applications, such as Progress Version 9 and Sybase ASE 11.9.2, require these symbols in order to run on the DG/UX system.

4.2.2　Expanding the /usr file system during a rolling upgrade

If you are performing a rolling upgrade and the size of your **/usr** file system is less than 320,000 blocks, you must perform all of the additional steps below. Otherwise you may run out of space in the new **/usr** file system during the setup phase of the installation.

1. If you have not created your new release area, do so now. The following directions assume the release area is named "B".

2. You need to know the current size of your **/srv/release/B/usr** file system in order to know how much to expand it in the next step. If you do not know the size of the file system, use the following **sysadm** menu operation to list the virtual disk **usr.B** (or the **admvdisk**(1M) command):

   ```
   Device -> Disk -> Virtual -> List
   ```

3. Use the following **sysadm** menu operation [or the **admfilesystem**(1M) command] to increase the **/srv/release/B/usr** file system (virtual disk **usr.B**) to the new recommended size of 320,000 blocks or larger:

```
File System -> Local Filesys -> Expand
```

4. Perform this step before running any **admrelease –o check** operations on the new release area "B".

5. Run the following command to update the release area database:

   admrelease –o set –O FS_USR –m 320000:/usr B

   If you chose to expand the new **/usr** file system more than the minimum required in step 3, replace *320000* in the command above with the size (in blocks) of the file system after expansion.

6. Proceed with the normal rolling upgrade procedure as documented in *Installing the DG/UX® System.*

If you expand the new **usr** file system (step 3) without updating its size in the release area database (step 4), installation may succeed but **admrelease –o check** operations on the new release area will fail with messages warning about a wrong file system size.

4.2.3   gated command modified for Year 2000

The gateway routing server program has been modified to no longer support the HELLO routing protocol. This change was made to resolve a Year 2000 readiness issue. The HELLO protocol uses date in RT–11 format, which uses five bits to represent the year since 1972.

The prototype configuration file **/etc/gated.conf.proto** will no longer provide the option to turn on this protocol as it did in previous releases. If your system runs **gated**, prior to starting the system in multiuser mode, please complete the following tasks:

*   Save a copy of the current **/etc/gated.conf** file

*   Make a working copy of the **/etc/gated.conf.proto** file

*   Incorporate site–specific changes made to the existing file into the new file

*   Copy the configuration file and place it in **/etc/gated.conf** and make sure the access control attributes are correct

For information about formatting the **/etc/gated.conf** configuration file, configuring **gated**, or migrating off the HELLO protocol, see the manual *Managing TCP/IP on the DG/UX® System.*

4.2.4 Improved large file support for DBU

Data Backup Utility (DBU) support of save and recovery of very large files (i.e., those larger than 4 gigabytes) was introduced with DG/UX R4.20MU04. However, the recovery of large files by DBU 5.1 did not work properly. DBU revision 5.1.0 contains a fix which allows very large files supported by DG/UX to be correctly recovered from DBU tapes, beginning with R4.20MU04.

4.2.5 IP Multicast and IGMP support

R4.20MU05 includes a DG/UX System extension that supports IGMP protocol versions 1 and 2 and allows applications to send and receive IP Multicast messages.

Multicast is the concept of sending a packet to a set of hosts on a network. A multicast group address is a class D IP address. The addresses are usually written as dotted–decimal numbers which range from 224.0.0.0 to 239.255.255.255.

The Internet Group Management Protocol (IGMP) is used by multicasting hosts and routers. It lets the systems on a physical network know which hosts belong to a particular multicast group. For more information about IP Multicast and IGMP, see the **ip**(6P) and **netstat**(1M) on–line manual pages.

4.2.6 Kernel parameter changes for process scheduling

DG/UX R4.20MU05 introduces changed meanings and default values for the kernel parameters JPSL0DELAY, JPSL1DELAY, and JPSL2DELAY. These parameters influence operation of the Medium Term Scheduler (MTS), a kernel daemon that balances system load across processors and locales.

We recommend that you use the default values for these parameters even if you changed them in a previous DG/UX revisions. The default values allow the MTS to balance the system load correctly for the specific hardware platform.

For AV 20000 and AV 25000 systems, the default values are:
– JPSL0DELAY = 5 microseconds
– JPSL1DELAY = 0 milliseconds
– JPSL2DELAY = 10 milliseconds

For more information about process scheduling, MTS, and these kernel parameters, see the manual *Analyzing DG/UX® System Performance.*

4.2.7 Maximum file system size is two terabytes

DG/UX file systems may now be up to two terabytes in size. In keeping with this upper limit, the maximum size of a file in a DG/UX file system is two terabytes minus the amount of space required for inode storage for the file system.

Previously, expanding a file system to two terabytes produced an invalid superblock. Although the last DG/UX release allowed for the creation and expansion of file systems to two terabytes, doing so rendered the file system inaccessible.

4.2.8 Network Control

Network Control (**netctrl**) is available through the **sysadm** menu option `Networking -> TCP/IP -> Network Control`. It manages Network Address Translation (NAT), proxy, Address Pooling, and Proxy–ARP functions in TCP/IP networking.

Rules are used to specify filtering and other functions for each network interface using Network Control. A search of the ordered list of filters for each interface proceeds (in order) until a match is found. Matching is based on any of the following: Source IP address/port, Destination IP Address/port, in or outgoing interfaces, and/or session type (TCP, UDP, other). If a match is found, a corresponding action is taken. Possible actions include: reject or accept the packet, translate source or destination IP address using static or dynamic addressing, translate the source or destination port, or proxy to a local port.

Rules are stored in Network Control configuration files, by default **/etc/netctrl.config**. The `Add`, `Modify`, `Delete`, and `List` operations may be used to manage these files.

Changes are made to a work configuration file, which is specified in the Configuration menu. A set of rules in the work configuration file may be tested using the Kernel menu. A completed set of rules in a **work ... next page? (? for help) [yes]** configuration file may be installed as the boot–time default rules using the following Configuration menu: `... end pager? (? for help) [yes] [A]`.

For more information about Network Control and the Network Control Configuration files, see the following on–line manual pages: **netctrl**(1M), **netctrl.config**(4M), and **adm.netctrl**(1M).

4.2.9 New diagnostic messages in epfc driver

The **epfc** (7) fibre channel device driver provides better, quicker error reporting. New error messages and/or clearer explanatory text may be evident in some messages that were previously provided by the driver. This improvement provides better diagnosis of potential problems.

4.2.10   nps command improved

The **–WA** switch now shows an aggregation of central processing unit (CPU) affinity and other related process information. New options include `affinity`, which displays the CPU that a light–weight process (LWP) is allowed to run, and `lcpu`, which displays the last CPU the LWP ran on.

4.2.11   Removal of setuid and setgid bits

When an object with **setuid** or **setgid** bits is opened for modification, the **setuid** and **setgid** bits are removed unless the opening process has appropriate privilege. This prevents an unauthorized process from writing to the object.

4.2.12   sendmail command updated

The **sendmail** command was updated to version 8.9.3. This new version of **sendmail** has several new features and options which include the following:

* Relaying, which involves transmitting a message from outside the local domain to a destination outside of the local domain, is disabled by default.

* Sending and receiving mail is controllable on a per–user and per–host basis.

* Many security vulnerabilities have been addressed.

* Option and Macro names in the configuration file may now extend beyond a single character.

A prototype configuration file for the new **sendmail** command is delivered with this release as **/etc/sendmail.cf.proto**. As part of the software installation, the existing **/etc/sendmail.cf** file is saved (if there is one), and this file is replaced by the prototype file. This is necessary for the installation process to be successful. The previous files is saved under the name /etc/sendmail.cf.*<date>*:*<pid>*, where *<date>* is the date on which the system was upgraded (in *yyyymmdd* format) and *<pid>* is the process ID of the process that saved the file (to ensure that the name is unique).

Before booting the system in multiuser mode, you must customize the **/etc/sendmail.cf** file so mail can be sent and received. To do this, complete these steps:

* Review comments at the top of the **/etc/sendmail.cf** file. They describe the most common customizations and provide reference to a NOTE number so that the appropriate place in the file can be easily located.

* Make the necessary changes, if any, based o these comments. At the very least, you will need to supply the domain name (NOTE 2).

* Review any customizations made to the previous **/etc/sendmail.cf** file, and incorporate those changes into the new file, if they are still necessary. For detailed information about the format of the **sendmail.cf** file, see the manual *"Managing TCP/IP on the DG/UX System."*

  Note that the names and/or definitions of several macros were changed from the previous version of **sendmail**. In particular, the domain name is now "m", instead of "D"; the relay machine is "S", instead of "M"; and "W" is the short name of the current host, as opposed to the fully qualified name.

* Save the changes and quit the editor. Then test the changes using **sendmail**'s test mode:

  ```
  sendmail -bt
  ```

  If this produces any errors or warnings, check the changes you made to the configuration file. Also, this version of **sendmail** is very careful not to access world writable files, since they can very easily be modified by unauthorized users. If warnings about such files or directories appear, check the permissions on these files.

* Make sure the access control attributes (e.g., permission bits) of the **/etc/sendmail.cf** file allow read access (and only read access) to all users.

For further details about configuring **sendmail**, see the manual *"Managing TCP/IP on the DG/UX System."*

4.2.13   Two fields expanded in utmpx and wtmpx

Larger fields for username and hostname have been provided in the **/var/adm/utmpx** and **/var/adm/wtmpx** files. In both files, 32 characters are now available to record username information and 256 characters are available for hostname information. However, TCP/IP limits hostnames to 63 characters.

The following commands have been enhanced to read the longer username field:

* **logname**, **pppd**, and **slipd** now retrieve usernames which contain more than 8 characters from the **utmpx** file, as opposed to those which contain 8 characters or less.

* **passwd** and **yppasswd** may be used to change the passwords for users with long usernames, without requiring that the username be part of the command line.

* **finger** is now able to display information about users with long usernames, instead of simply reporting "in real life: ????".

Various commands and daemons which write to these files were enhanced to write the longer entries correctly. The following commands and daemons, which are used to log in, start processes, or perform related tasks, now write larger fields: **ftpd**, **init**, **login**, **mterm**, **pcnfsd**, **rexd**, **rlogind**, **sac**, **telnetd**, **ttymon**, and **xterm**.

Finally, a new option was added to the **fwtmp** command. The **–x** option allows users to convert **utmpx** and **wtmpx** files which use longer fields. For more information, see the **utmpx** and **wtmpx** on–line manual pages.

4.2.14   wdt () device included by default

The **wdt** () device driver provides a watchdog timer that helps increase system availability. In previous releases, the **wdt** () device was not included by default in the **system.device** file when building a kernel. With this release, **wdt** () is included by default. If the **wdt** () device is not desired, the "wdt()" line must be removed from the **system.device** file before building a kernel. See the **wdt**(7) on–line manual page for details.

All hardware configurations support a physical watchdog timer. The physical watchdog timer is a hardware device that must be periodically reset by software. In the case of a software hang, the software fails to reset the watchdog timer and the timer expires. At this point, the watchdog timer issues a non–mistakable interrupt to the system. If autodump is enabled, the system will attempt to take a memory dump. If autoboot is enabled, the system will reboot without user intervention. On some systems, a hardware reset is issued by the watchdog timer if the system does not respond to the non–maskable interrupt.

This chart shows the level of watchdog timer support on the AV models.

	Non–maskable interrupt	Reset
AV 1600	x	
AV 2000	x	
AV 2100	x	
AV 2600	x	
AV 2600R	x	
AV 3000	x	
AV 3600	x	
AV 3600R	x	
AV 3650	x	

AV 3700	x	
AV 3700R	x	
AV 4700	x	x
AV 4800	x	x
AV 4900	x	x
AV 5800	x	x
AV 5900	x	x
AV 6600	x	
AV 20000	x	x
AV 25000	x	x

4.3 Hardware features introduced in previous R4.20 MUs

4.3.1   Servers and Processors

AV 3650 and AV 3650R – In R4.20MU01 these products replaced the existing AV 3600 and 3600R. For information about these systems, see the documentation produced for the AV 3650 and AV 3650R systems.

AV 3700 and AV 3700R– R4.20MU04 introduced support for the AV 3700 and AV 3700R, which replaced the AV 3650 and AV 3650R. The new systems are available with up to four Intel Pentium II Xeon processors with either 512 KB or 1 MB integrated L2 cache. These systems support up to 4GB of memory and feature redundant cooling and hot–swap redundant power.

AV 6600 1 MB L2 cache – In R4.20MU01 support for the optional AV 6600 1 MB L2 cache was introduced. This option is available under the following model numbers: R7389: replace 200Mhz/512K CPUs with 200Mhz/1 MB and R7119: add–on tri–CPU board (200Mhz/1 MB).

AV 20000 server – R4.20 introduced this Intel–based ccNUMA system. It is a 4– to 32–processor, shared–memory system constructed of one to eight computing modules based on Intel quad–Pentium Pro processor motherboards. For more information, see the documentation produced for the AV 20000 server.

AV 20000 32–CPU server – In R4.20MU01 the AV 20000 was enhanced to support systems with up to 32 CPUs. The previous system introduced in R4.20 has a 16–CPU limit. For more information, see the documentation produced for the AV 20000 32–CPU server.

AV 20000 32–CPU server full support – R4.20MU02 contains software enhancements to fully support five– to eight–block (20– to 32–CPU) AV 20000 systems.

AV 25000 computer support – R4.20MU03 introduced support for single block AV 25000 systems. The single block incorporates four Intel Pentium II Xeon processors; an IO subsystem with ten PCI slots, including four 64–bit slots; up to 8 GB of memory; and redundant hot swappable power and cooling.

AV 25000 multi–block server– R4.20MU04 introduced support for multi–block AV 25000 systems (up to six blocks—24 CPUs). Prior to this release only single block (4 CPU) AV 25000 systems were supported.

Machine–initiated callout on systems without a TSB– In R4.20MU01, this feature, which was previously supported only on systems with an exTended Service Board (TSB), was enhanced to provide support on systems without a TSB. For information about this feature and the systems it is supported on, see the documentation produced for machine–initiated callout.

Pentium Pro processor– In R4.20 there was a shift in optimization away from Pentium towards Pentium Pro. The C compiler generates code that is optimized to run on the Pentium Pro systems, but will also run on the Pentium systems. For more information, see the documentation produced for the Pentium Pro processor.

4.3.2   Controllers

ATM adapter support– R4.20MU01 supports up to four ATM controllers per system. This is an increase over the previous limit of two per system in R4.20.

EPFC Emulex Fibre Channel driver automatic firmware download– R4.20MU03 prompts the EPFC Emulex Fibre Channel SCSI adapter driver to automatically download Emulex firmware and BIOS during device configuration. The download includes version FF2.10 firmware containing the Emulex SLI–2 interface and version FB1.10 BIOS. These items are downloaded to the EPFC Emulex Fibre Channel SCSI adapter. Downloading of this firmware will only occur when older firmware is detected by the EPFC driver.

Model 7463 PCI Ethernet LAN adapter – In R4.20MU01 the single–port DEC DE500–BA 10/100 Ethernet adapter, an updated version of the DEC DE500–XA adapter, was introduced. For more information, see the documentation provided for the Model 7463 PCI Ethernet LAN adapter.

Model 7466 Synchronous Controller– R4.20 added DG/UX System support for the Model 7466 Synchronous Controller. This controller allows two synchronous connections per card with any mix of media types (RS–232, RS–530, V.35 or X.21). It is supported on all Intel AViiON systems. For more information, see the **apsy**(7) on–line manual page.

Model 7467 PCI Ethernet LAN adapter– In R4.20MU01 the four–port DEC DE500–BA 10/100 Ethernet adapter, an updated version of the DEC DE500–XA adapter, was introduced. For more information, see the documentation provided for the Model 7467 PCI Ethernet LAN adapter.

Model 7468 Qlogic Differential SCSI Controller – R4.20 added DG/UX System support for the Model 7468 Qlogic Differential SCSI Controller. This controller supports differential SCSI devices in AV 20000 configurations. The **qpsc** adapter driver supports both fast and wide SCSI–2 and allows up to 15 SCSI devices per bus. The **qpsc** also supports multi–initiator SCSI environments. For more information, see the **qpsc**(7) on–line manual page.

Model 7469 SCSI adapter – In R4.20MU01 support was added for the Model 7469 Dual Channel Differential SCSI adapter. It is supported as an **apsc** controller. For more information, see the **apsc**(7) on–line manual page.

Model 7472 and 7473 PCI ATM adapters– In R4.20MU01 support was added for the model 7472 and 7473 Interphase 5575 PCI ATM adapters. These models run on all previously supported machines as well as AV 3600, AV 3650, and AV 6600 platforms. For more information, refer to the **ia**(7) on–line manual page.

Model 7487 Fibre Channel adapter firmware download– R4.20MU04 introduced automatic firmware download for the Emulex Model 7487 Fibre Channel adapter. Firmware and BIOS are automatically downloaded (if needed) by the **epfc** driver during device configuration. R4.20MU04 downloads firmware revision FF2. 13A0 and BIOS revision FB1.10A10 on the Emulex Fibre Channel adapter.

Model 7492 ATI graphics card– R4.20MU01 added support for the ATI 3D Pro Turbo PC2TV with the Rage II chip.

Model 7493 ATI graphics card– R4.20MU02 added support for the ATI 3D Pro Turbo PC2TV with the 3D Rage II+ accelerator chip.

Model 7497 PCI Ethernet LAN adapter– In R4.20MU02 this version of the four–port 10/100 Ethernet PCI adapter was made available. Model 7497 is an Adaptec ANA–6944B PCI 10/100 Quartet Ethernet adapter. For more information, see the documentation provided for the Model 7497 PCI Ethernet LAN adapter.

For information about a change in the hardware's port order, see "Change in Model 7497 External Connector port order" later in this update notice under "Hardware notes and warnings."

Model 7498 Packet Engines Gigabit Ethernet Adapter– R4.20MU04 supports the Model 7498 gigabit ethernet adapter, which runs in a 32–bit or 64–bit PCI slot. This adapter supports IEEE 802.3z 1000base–SX connectivity using multi–mode fibre. The following systems support model 7498: AV 3600, AV 3600R, AV 3650, AV 3650R, AV 3700, AV 3700R, AV 6600, AV 20000 and AV 25000. For more information, see the **pgen**(7) on–line manual page.

4.3.3   Peripherals

CLARiiON 5400 and 5500 Disk Arrays – R4.20MU03 added support for the CLARiiON 5400 and 5500 full fibre disk arrays.

Disk models supported in R4.20MU04 –

Part Number	Description
61034–SJL	4GB HS 7200RPM LVD disk for AV 3700
79017–HFL	9GB 7200 RPM Fibre Channel array disk
79018–HFL	18GB 7200 RPM Fibre Channel array disk
79019–HFL	9GB 10K RPM Fibre Channel array disk
79020–HFL	18GB 10K RPM Fibre Channel array disk

Model 61018 12X CD–ROM – R4.20 added support for this CD–ROM on all AViiON computers running R4.20 or later releases.

Model 7486 PCI based Redundant Array of Independent Disks – R4.20MU02 added support for the Mylex DAC960PDU RAID controller. This controller lets you configure external SCSI disks into RAID–0, RAID–1, and RAID–5 disk arrays. For more information, see the **mprc**(7) on–line manual page.

Model 79587 and Model 79588 Brocade Fibre Channel Switches – R4.20MU04 added support for these Brocade Silkworm fibre channel switches.

4.4 Software enhancements and changes introduced in previous R4.20 MUs

4.4.1   BSD 4.4 telnet ported to the DG/UX system (R4.20)

A new version of the **telnet** command (BSD revision 4.4) has been ported to the DG/UX system. The new version fixes bugs and provides changed and new features. For example, the new general commands **set**, **unset**, and **toggle** replace specific commands such as **crmod**, **debug**, and **escape** and add control previously not available. For more information, see the **telnet**(1M) on–line manual page.

4.4.2   ccNUMA technology (R4.20)

Data General ships the DG/UX Cache–Coherent Non–Uniform Memory Access (ccNUMA) technology on its AV20000 and 25000 series machines. Fundamentally an extension to the DG/UX Symmetric Multi–Processor(SMP) architecture, DG/UX ccNUMA offers significant performance improvements with no required modifications to existing SMP applications.

For detailed information about DG/UX ccNUMA system technology and how to modify configuration variables to improve ccNUMA system performance, see *Analyzing DG/UX® System Performance.*

4.4.3   Data Backup Utility (R4.20)

The DG/UX system provides a restricted version of Legato NetWorker called the Data Backup Utility(DBU). This utility enables you to back up your local file systems to a single local backup device. You can upgrade the DBU to NetWorker for DG/UX Network Edition or NetWorker for DG/UX Power Edition by purchasing the appropriate licenses. For more information, see *Managing the DG/UX® System.*

If NetWorker 4.2.5 or NetWorker 4.2.9 is installed on your computer, the DBU will not install. For more information, refer to the section entitled "Installing DBU on the DG/UX System leads to updated index file formats" later in this notice.

The multiclient version of Legato NetWorker is an optional, licensed product that can be purchased through Data General Corporation. It provides a complete backup solution for DG/UX AViiON systems and lets you back up more than one machine to more than one device. Legato NetWorker 5.1 for DG/UX R4.20MU04 also supports large file systems.

4.4.4   Date format changed in SCCS commands (R4.20)

The following SCCS commands were changed to have a new date/time format in order to allow users to specify a four–digit year from their respective command lines:

* **get**

* **prs**

For example, the **–c** <*cutoff*> option syntax can be specified in either a two– or four–digit year format.

* –c98/0228235959

* –c1998/0228235959

Note that the year must be separated from the date and time by a separator.

For more details, see the **get**(1), **prs**(1) and **sccsfile**(4) on–line manual pages.

4.4.5 Fewer dots printed during disk registration (R4.20)

On large systems that have hundreds of disk paths, like the AV 20000, fewer dots are printed during the disk registration message, making it easier to follow.

If there are more than 100 disk paths to register on the system, an asterisk (*) is printed for every 10 disks registered. This greatly reduces the number of indicators printed, generally down to one or two lines at most (two lines of * would mean 1400 disk paths registered). This would have been 17 lines of dots previously.

If there are fewer than 100 disk paths to register, the existing dot indicator is used, printing it once per disk path registered.

4.4.6 gcc and cc commands changed in R4.20

The default behavior of the **cc** and **gcc** commands for R4.20 is to produce code that is optimized for Pentium Pro systems, but that will run on both Pentium and Pentium Pro. The Pentium Pro system supports additional instructions that can improve performance, but code that uses these instructions does not run on a Pentium system. A command line option enables the use of Pentium Pro instructions:

* **gcc –march=pentiumpro**

* **cc –W0,–march=pentiumpro**

4.4.7   Internal data structures changed in R4.20

Starting with Release 4.20, the DG/UX system has changed the internal structure of **io_buffer_vector_type** and the size of **physical_address_type**. These changes accommodate the larger address spaces available on Intel Pentium Pro processors. These changes in the kernel do not affect application code. The **io_buffer_vector_type** is a kernel internal data structure used to communicate information about read and write requests between non–streams device drivers and the kernel proper. It is contained within an **io_operation_record_type**, which forms part of the interface between the kernel and device drivers. **physical_address_type** is used to hold a physical address; it is returned by **vm_get_physical_byte_address()** and is primarily used by devices (both streams and non–streams) that perform digital memory access (DMA).

Because of these changes, you need to recompile all device drivers written for the DG/UX system on Intel. For drivers that reference the internal fields of **io_buffer_vector_type**, some mechanical code changes will also be necessary. For more information, see the **io_buffer_vector_type**(3K) on–line manual page.

The kernel include files contain the structure definitions. Refer to the files below for the modified definitions.

*   The definition of `physical_address_type` is in the kernel include file **ext/os_generics.h**.

*   The definition of `io_buffer_vector_type` is in the kernel include file **ii/i_io.h**.

4.4.8   Modified admdevice operations (R4.20)

The following **admdevice** operations were modified for R4.20:

*   The **–o configure** operation automatically detects and configures all paths to a physical disk device. You can use the new **–p** option to override this behavior and force only the specified paths to be configured.

*   The **–o deconfigure** operation automatically detects and deconfigures all paths to a physical disk device. You can use the new **–p** option to override this behavior and force only the specified paths to be deconfigured.

4.4.9   Perl version 5 (R4.20)

The DG/UX system supports Perl version 5. Use the following line as the first line of any Perl script that you want to be executed using **perl5** instead of **perl4**, which is the default.

```
# /usr/bin/perl5
```

The directory **/usr/lib/perl5** contains **perl5** libraries and classes.

For more information, see the on–line manual pages.

4.4.10   Scalable network multiconnected LAN (R4.20)

Multiconnected LAN (MCL) provides multiple I/O paths to a single Ethernet network. Using MCL reduces the cache–coherency traffic on the AV 20000 SCI interconnect and enhances network availability. It also increases system bandwidth in switched networks.

MCL provides a limited form of fully connected LAN I/O. MCL allows multiple physical devices (one in each ccNUMA locale) to be joined into a single TCP/IP interface. All these devices should be connected to the same physical network. MCL directs outbound TCP/IP traffic to the device with the same locale as the memory buffer. This avoids cross locale direct memory access (DMA) and reduces the traffic on the SCI interconnect. MCL also balances the load of client access to the ccNUMA system by responding to address resolution protocol (ARP) requests in a round robin fashion. This form of load balancing is designed for large servers with many clients. In other environments, such as a server behind a router or a LAN with a few large clients, it can be more effective if each interface has its own IP address. Lastly, in some On–Line Transaction Environments (OLTP) environments it can be beneficial to limit network I/O to a single locale in the system, reserving other locales for disk I/O and data processing. Unlike with disk I/O, there is no single LAN solution for every networking application and topology.

MCL implementation involves using the **admipinterface** command to manage the TCP/IP interfaces built with MCL. The database of these interfaces is stored in the **/etc/tcpip.params** file. The **admipinterface** command has a new option, **–s** *<subdevice_list>*, which you use only when you specify the **–d mcl** option. Use commas or double quotes and spaces to separate the network devices specified with *<subdevice_list>*. All the MCL features are also controllable through **sysadm** choices in the `Networking -> TCP/IP -> Interfaces` menu.

The following example creates a TCP/IP interface "mcl0" using IP address "myhost" using Ethernet devices **dpen0** and **dpen1**. The Ethernet devices would be located one in each NUMA locale and attached to the same physical Ethernet segment.

```
admipinterface -o add -d mcl -s "dpen0,dpen1" -i mcl0 myhost
```

> **IMPORTANT**   MCL is designed for switched LANs; it should not be used on shared LANs.

4.4.11 admipinterface command (R4.20MU01)

R4.20MU01 added support for an **–s** option to the **admipinterface** command. The argument to this option specifies a list of one or more lower network devices to be linked under the multiplexer device. This option is valid only when the device is a multiplexing device, which includes Multiconnected LAN (MCL) and Virtual LAN (VLAN). The names in the list can be separated by a comma or separated by a space and enclosed in double quotation marks. For example, **"dpen1 dpen2"** is equivalent to **dpen1,dpen2**; and both are valid arguments to **–s**.

4.4.12 Application Capture Option (R4.20MU01)

The DG/UX system includes an Application Capture Option (ACO) package, **dgux.aco**, that enables certain UnixWare and SCO applications to run on the DG/UX system for Intel without being recompiled. This package and its corresponding manual page package (**dgux.aco.man**) are part of the **development** and **production** package suites because the software is required by a number of important third–party software products.

For information about deciding whether to load ACO or how to avoid loading it, see "Application Capture Option (R4.20MU01)" in the "Notes and Warnings" section of this notice.

4.4.13 Enhancements to MPL (R4.20MU01)

This release offers two enhancements to Multipath LAN I/O (MPL):

*   MPL has been enhanced to support multiple relationships per host in a machine–initiated failover configuration. Prior to this enhancement, only a single MPL relationship could be defined per host. Now you can configure any number of pairs of Ethernet or FDDI Links into MPL relationships.

*   Error logging and error recovery from LAN link start failures has been improved if the LAN link is the MPL primary or backup link.

4.4.14 Year 2000 readiness (R4.20MU01)

Data General has evaluated the DG/UX System to identify potential Year 2000 issues. The work identified during the evaluation was complete and delivered in R4.20MU01.

Based on the evaluation, the DG/UX operating system kernel is Year 2000 ready. The DG/UX kernel tracks time with an internal chronometer that counts "tics" from an arbitrary beginning date and converts the chronometer value to the country–specific format for display. File creation, access, and last–write times use the internal chronometer and therefore are unaffected by the Year 2000. No reformatting, conversion, or reloading of files is necessary.

During the evaluation, some DG/UX System layered commands and utilities, such as SCCS, were found to accept input with a two–digit year format that required minor changes for Year 2000 readiness. These identified changes were complete and delivered in R4.20MU01.

Note that software provided by customers and other third parties was not included in the evaluation. You should address questions about such software to the software supplier.

4.4.15 DG/UX features for AV 20000 computers (R4.20MU02)

This section describes new features related to booting and system parameter storage specific to AV 20000 machines running DG/UX R4.20MU02.

4.4.15.1 Boot command strings

On AV 20000 machines, boot command strings are stored in hardware, on the 3D Board, instead of on the disk as was done in DG/UX R4.20. The new software ignores the strings on disk. Each 3D Board supports three boot strings 255 characters. One must manually initialize the boot strings on the 3D Board using the **sysadm** `System -> NvRAM -> Set Parameter` menu sequence.

Each block of an AV 20000 computer contains a 3D Board. The 3D Board in the lowest numbered active block is the "master" 3D Board. When the system boots, it obtains the boot strings from the master 3D Board. You can write to the boot strings of any 3D board using **sysadm**. You should do this to prepare for a failure of all of the lower numbered blocks. The boot strings written to block one are used to prepare for a failure of block zero. These boot strings should not be dependent on any devices, such as disk controllers, from block zero.

4.4.15.2 Boot timeout

The default timeout when automatically booting is 15 seconds. The system displays the name of the kernel that will be booted and prompts for either the Escape or Enter key to be pressed to cancel or proceed with booting. You can add a **–t nnn** switch, where *nnn* is the timeout in seconds to each stored command line to override the default 15–second timeout. This feature was first included in DG/UX R4.20MU01.

4.4.15.3 /dev/scm(n)

For each block on an AV 20000 machine, the operating system creates a device entry called **scm(n)**, where *n* is the hexadecimal block number. The **/dev/scm(n)** device was not created with DG/UX R4.20. R4.20MU01 had only one **scm** device per system, not one per block. **admnvram** uses the **scm** device to read and write boot command strings and other parameters to the 3D Board. Since the **scm** device did not exist in DG/UX R4.20 or R4.20MU01, you should regenerate the device list when building the kernel to upgrade to DG/UX R4.20MU02.

4.4.15.4 halt command

To maintain compatibility with the DG/UX R4.20 AV 20000 release, when the **halt** command is issued, the system shuts down its current operation, resets, and autoboots. You can use **halt –r** on an AV 20000 machine, as on disk based systems, to stop the system without autobooting.

4.4.16   Support for Mylex DAC960 utilities (R4.20MU02)

DG/UX R4.20MU02 adds support for the Mylex DAC960PDU PCI based RAID controller (see the **mprc**(7) on–line manual page). This support includes the use of the Mylex Global Array Manager (GAM) client software, which runs on a networked PC running a Microsoft operating system. The GAM client lets the DG/UX system administrator monitor and control the operation of the Mylex DAC960PDU controller from a local desktop. For more information on this feature, see the **gamserv**(7) on–line manual page.

4.4.17   New file system type (R4.20MU03)

R4.20MU03 adds support for failover file systems. This new feature is intended for use with the DG/UX Cluster software. For more information, see the Release Notices that ship with the Cluster and NetWorker products.

4.4.18   POSIX ACLs (R4.20MU03)

R4.20MU03 adds support for POSIX access control lists (ACLs). This package is provided to support Oracle Express, though ACLs have other uses. The **acl** package suite consists of the packages **dgux.acl.** and **dgux.acl.man**. For more information about POSIX ACLs, see the **setacl** and **getacl** on–line manual pages.

4.4.19   Large File Support (R4.20MU04)

R4.20MU04 includes a DG/UX System extension that enables commands, functions, and applications to create and manipulate large files. A large file is one whose size is greater than two gigabytes.

WARNING If an exported DG/UX R4.20MU04 (or later release) file system
contains a large file, do not NFS mount the file system from a system
running an earlier DG/UX release. If you must mount the file system
using an earlier release, use the **–o vers=2 mount** option. Failure to
use this option may cause unexpected results.

For more information about large file support, such as how to write an application so
it handles large files correctly, see the *Porting and Developing Applications on the
DG/UX® System* manual, and the **largefile_intro**(5) on–line manual page.

4.4.20 LD_NO_LIBSTREAMSOCKET environment variable for runtime linker (R4.20MU04)

Support for a new environment variable has been added to the runtime dynamic
linker. The environment variable LD_NO_LIBSTREAMSOCKET is used by shared
third party applications (typically built on UnixWare) to specify how the dynamic
linker should load the socket services library (libsocket).

Third party applications that use streams operations on open socket descriptors
should set LD_NO_LIBSTREAMSOCKET to 0 (zero). This will cause the dynamic
linker to load **libstreamsocket.so.1** to resolve all socket system calls.

Third party applications that do not use streams operations on open socket
descriptors, and that want faster performance of kernel–based sockets, may set
LD_NO_LIBSTREAMSOCKET to 1 (one), which will cause the dynamic linker to load
libsocket.so.1.

If LD_NO_LIBSTREAMSOCKET is undefined, or defined as anything other than 0 or
1, the dynamic linker will print a warning message and behave as if
LD_NO_LIBSTREAMSOCKET is set to 0. If the third party application does not need
the socket services library, then LD_NO_LIBSTREAMSOCKET will have no effect.

4.4.21 libstreamsocket part of the base DG/UX System (R4.20MU04)

The file **libstreamsocket.so.1** is now delivered with the base DG/UX System (it was
previously available as part of the ACO option only). In addition the kernel modules
supporting this library (**sockmod**( ), **ticlts**( ), **ticots**( ), **ticotsor**( ), **pw**( ), **socksys**( ))
are now configured by default into the base DG/UX kernel.

4.4.22 make(1) correctly expands macros to completion (R4.20MU04)

The **make**(1) command has been modified to correctly expand macros in the form of
$ *MACRONAME:oldvalue=newvalue* to completion. For example, in the following
lines from a makefile, the **make**(1) command will replace instances of ".c" with ".o" in
the string "main.c a.c b.c".

```
A = main.c
B = a.c b.c
C = $(A) $(B)
target: $(C:.c=.o)
```

Previous versions of **make**(1) would not evaluate the $(A) and $(B) macros, but would try instead to replace any instances of ".c" with ".o" in the string "$(A) $(B)".

The **make**(1) command also supports use of the $@ internal macro in a dependency list. For more information, see the **make**(1) on–line manual page.

4.4.23   ntpq command modified for Year 2000 (R4.20MU04)

The **ntpq (1M)** command has been modified to treat "calculated" time values from RT–11 operating systems as indecipherable. For such queries this command will display a question mark (?). This change was made to resolve a Year 2000 readiness issue in which "calculated" time values from this operating system return two–digit year values.

4.4.24   passwd(1) (R4.20MU04)

A new option has been added to the **passwd**(1) command that causes it to operate on the files in the directory specified by SRC_DIR in **/etc/yp/Makefile**. The administrator on the NIS master will then be able to change a user's NIS password without having to know the old password. This option is not supported on either the DG/UX B2 Security Option or the C2 Trusted DG/UX Option. See the information on the **–Y** option in the **passwd**(1) on–line manual page for details.

5   Notes and warnings

This section contains the following subsections:

*   Hardware notes and warnings in R4.20MU05

*   Software notes and warnings in R4.20MU05

*   Hardware notes and warnings introduced in previous R4.20 MUs

*   Software notes and warnings introduced in previous R4.20 MUs

5.1 Hardware notes and warnings in R4.20MU05

5.1.1   Changes in /dev/fru after R4.20MU05

R4.20MU05 is the last DG/UX release that will provide CLARiiON FRU information about new CLARiiON storage system platforms through the **/dev/fru/interface**. The FRU_INQUIRY and FRU_PROBE ioctls defined in **/usr/include/sys/_int_dg_fru_ioctl.h** will continue to be supported for non–disk devices and legacy disk devices, but support for new CLARiiON storage systems will not be added starting with R4.20MU06.

R4.20MU05 is also the last release in which special hadall entries are created in **/dev/fru**. Entries such as

/dev/fru/hadall(aspsc(dppb(pci(1),D,0),4,0,6),0,0)

will no longer be created in R4.20MU06.

5.2 Software notes and warnings in R4.20MU05

5.2.1   pg and other commands that read from /dev/tty

Switching to another user (other than root) via **su** and then using **pg** or another command that accesses **/dev/tty,** causes an error. An example of the **pg** command error message for this type of operation is as follows:

```
UX:pg: ERROR: Cannot reopen stdout: Permission denied
```

Other commands which may exhibit this problem are as follows:

```
cpio -ir
ftp
login
passwd
xargs -p
```

The problem stems from the fact that the terminal is owned by the UID of the user who originally logged in, and the permissions allow only write access or no access to other users. Thus, when an unprivileged user attempts to access the terminal, access is denied.

In some cases, commands that function similar to the failing command may be used to replace it. For example, the **more** command may be used instead of **pg**. In other cases, alternate commands are not available.

R4.20MU05 provides a kernel configuration variable which may be set to disable this access check. To do this, rebuild the kernel and specify the `SKIPTTYACCESS` parameter. Then reboot the system using the new kernel. This only affects access attempts that use the **/dev/tty** name. Attempts to access the terminal directly via its name (e.g., **/dev/console** or **/dev/pts/5**) still results in an access failure.

5.2.2   Year 2000 testing

Data General recommends that you adhere to the following guidelines when testing applications software for Year 2000 readiness:

* Identify any data files which are time sensitive (accounting reports, system logs, etc.) and make appropriate backups of current data so that you can restore these files to their current states when Year 2000 testing is complete.

* Identify any user or role accounts that you will use for testing and set the password aging interval for these accounts to a time well past the planned test dates in the year 2000. This will ensure that the passwords will not expire or become extinct when you change the system date to a time appropriate for Year 2000 testing.

* Review any **cron** jobs that run on a periodic basis to determine whether there are any time–sensitive issues that might cause problems due to a sudden shift in system time and make any necessary changes.

* Depending on the type of hardware you have, you may need to change the date in both the BIOS and the DG/UX software. To change the date in the BIOS, take the system down to the SCM and reset the date to the desired test date. Then boot the system to **init** S and make a similar change for the operating system. Then reboot the system to multi–user mode. This will ensure that the hardware and operating system dates are consistent and that all system daemons start on the correct date.

* When application testing is completed, take the system back down to the SCM and reset the date in the BIOS. Boot the system to **init** S to change the software date, then restore any archived data files. Next, reset the aging parameters associated with test user accounts to the proper interval as set by your site security policy. Then return to multi–user mode.

5.3 Hardware notes and warnings introduced in previous R4.20 MUs

5.3.1   ATM adapter limitations and restrictions (R4.20)

During certain stress test conditions, Data General observes the potential for poor Network File System (NFS) performance and, in some cases, hung file transfers. This

can occur due to retransmission timeouts for packets that are discarded because of congestion. This occurrence is observed for ATM networks that use NFS remote–mounted file systems. You can change the **mount** command parameters for NFS remote–mounted file systems (timeo=20, rsize=8192, wsize=8192) to minimize NFS performance degradation due to NFS retransmissions caused by ATM traffic congestion. See the **mount**(1M) on–line manual page on how to change these mount parameters. You can use the **nfsstat** command to check the number of NFS retransmissions and timeouts. Large values reported in these statistics indicate that the **mount** command parameters need to be modified for the NFS remote–mounted file systems.

The following restrictions apply to ATM adapter models 7459, 7460, 7472, and 7473:

*   Simple Network Management Protocol (SNMP) does not support separate ATM Management Information Bases (MIBs).

*   ATM Forum Available Bit Rate (ABR) Service is not supported.

*   LAN Emulation Server (LES) and the Broadcast and Unknown frame Server (BUS) are not supported. The switch or another server in the ATM network must provide this support.

*   No Application Program Interfaces (API) is provided.

*   The adapters are not supported on the Data General Peripheral Component Interconnect (PCI) expansion chassis, models 7505, 7506, and 7507.

*   When running Classical Internet Protocol (IP), if the cable is unplugged for a short time (less than a minute), the ATM network connections do not recover. Unplug the cable for a longer period (about one minute) to restore Classical IP service; otherwise a system reboot is required for recovery.

*   Console error messages are displayed when the cable is pulled from an active ATM adapter regardless of which software driver is used. The following are samples of possible error messages that might occur running LAN emulation.

     If the system fails to connect to the LAN Emulation Configuration Server (LECS):

```
LEC0-0: Conn to LECS failed

network out of order
```

     If there is no response to polling:

```
Q.SAAL Status Indication SSCOP: vt(cc)>=Max cc

Q.SAAL Status Indication SSCOP vcc down
```

Similar messages are displayed for Classical IP over ATM. Other messages from the ATM signaling software may be displayed on the system console.

5.3.2   AV 6600 memory error reporting (R4.20)

On AV 6600 computers, Address Bit Permuting (ABP) is a factory–enabled feature that provides performance enhancement for most applications. When ABP is enabled, however, memory module error isolation is not correct. This may result in the incorrect identification of the failing dual in–line memory module (DIMM) or single in–line memory module (SIMM).

If you experience memory errors and are concerned that the wrong DIMM or SIMM is identified, contact your support center representative for information on how to address this problem. Give the support center representative the Field Alert Bulletin (FAB) number S1381 for reference.

5.3.3   DG/UX 3.5 diskette driver not supported on AV 20000 (R4.20)

The AV 20000 does not support a DG/UX driver for 3.5–inch diskette media. You can, however, use the boot diskette provided with the AV 20000 to boot the machine because the boot diskette does not require a driver to work.

5.3.4   PCI bus (R4.20)

AV 4700, AV 4800, and AV 5800 systems have two Peripheral Component Interconnect (PCI) buses. The first PCI bus is compliant with all PCI cards. The second bus does not allow I/O space access to registers on the PCI cards. This means that a card on the secondary bus must be fully operable without any I/O space access (that is, it requires only memory–mapped access). Since the DG Model 24025 PCI 10/100 Fast EtherLink adapter card is not fully operable memory–mapped, it does not work on the second bus of these systems. The model 7461 Computone asynchronous controller has the same restriction. For more information, refer to the **tcen**(7) on–line manual page.

5.3.5   SCSI device restriction (R4.20)

SCSI devices connected to an Adaptec PCI SCSI–2 Host Bus Adapter **apsc** cannot have a Logical Unit 15 (LUN 15; F in hex) on a SCSI ID of F. This is the only SCSI ID for which you cannot assign LUN 15. Any requests sent to such a LUN ID are refused by the **apsc** driver. Connection to LUNs 0 through 14 (E in hex) for SCSI ID F are still allowed.

5.3.6   Multiple DG/UX systems sharing a SCSI bus (R4.20MU01)

Running multiple non–clustered DG/UX systems on a shared SCSI bus is not supported, except for B2 DG/UX systems running Failover with single path active multipath I/O.

Sharing a SCSI bus was a common and acceptable configuration for the DG/UX Failover package. However, Data General now supports failover only on the DG/UX B2 Security Option. With the multipath active multipath I/O feature, the Virtual Disk Metafile (VDM) no longer places a SCSI reserve on the disk when the disk is registered. Multipath active multipath I/O lets both systems register the disk, which results in data corruption or loss. Even if you are careful to include in the kernel only disks that are used on the local system, booting **dgux.installer** or the installation software while the other system is up causes the problem. DG/UX Cluster Software must be used to provide application high availability.

5.3.7   Change in Model 7497 External Connector port order (R4.20MU02)

The external connector port order of the new Model 7497 four–port PCI 10/100 Ethernet LAN adapter is the reverse of the earlier Model 7458 and 7467 multichannel LAN adapters.

For the four ports on the 7458 and 7467 adapters, port 1 is the RJ45 connector closest to the adapter's host PCI Bus Slot connectors (bottom), and port 4 is the farthest (top). On these adapters, the link and activity light–emitting diodes (LEDs) for each port are immediately next to the RJ45 connector for that port within the end plate cut–out of that port. On the end plate, letters may be engraved: LNK next to port 1 (bottom) and ACT next to port 4 (top).

On the 7497 adapter, port 4 is the RJ45 connector closest to adapter's host PCI Bus Slot connectors (bottom), and port 1 is the farthest (top). The 7497 has two sets of four LEDs, one activity and one link LED for each of two ports. The two sets of LEDs are located between the card edge and the RJ45 connectors on the adapter end plate. The 7497 also has the letters 100TX engraved on the end plate between the bottom set of LEDs and the card edge. The RJ45 connector closest to the 100TX is port 4.

With all three adapters, the external connector mapping to PCI slot assignments is the same: Port 1 is in PCI slot 4 so that there is no way to distinguish the adapter port ordering from the PCI slot number information included as part of the DG/UX device node name. For example, all three adapters would create the same device names if installed in the same host chassis PCI slot:

Device name	7458/7467	7497
`dpen(dppb(pci(0),F,0),4,0) = port 1`	bottom	top
`dpen(dppb(pci(0),F,0),5,0) = port 2`		
`dpen(dppb(pci(0),F,0),6,0) = port 3`		
`dpen(dppb(pci(0),F,0),7,0) = port 4`	top	bottom

5.3.8   DG/UX 3.5 diskette driver not supported on AV 25000 (R4.20MU03)

The AV 25000 does not support a DG/UX driver for 3.5–inch diskette media. You can, however, use the boot diskette provided with the AV 25000 to boot the machine because the boot diskette does not require a driver to work.

5.3.9   Mylex RAID controller model 7486 physical disk copies (R4.20MU03)

Physical disk copies between Mylex based disks and non–Mylex disks are not supported due to differences in physical disk geometries.

5.3.10   AV 3700 may try to boot before CLARiiON completes powerup (R4.20MU04)

The AV 3700's POST testing can complete and the system attempt to boot from the hard disk before the CLARiiON disk is completely powered up. This is because the AV 3700 does not have an SCM prompt. If the CLARiiON is the boot disk, the system will fail to load the DG/UX bootstraps. The system will have to be reset or power–cycled to boot.

To configure the system for an unattended reboot:

* An internal hard disk is required. This disk must have the following: a DGUX_VDM PC partition; DG/UX bootstraps; the disk–based NvRam Database(NDB) installed; and it must be DG/UX soft–formatted.

* Disable BIOS on all controllers connected to the CLARiiON. This will prevent the CLARiiON from being seen as the "C:" drive from which the bootstraps are obtained.

* DG/UX can be installed either on the internal OR on the CLARiiON disk. Installing DG/UX on the CLARiiON offers RAID protection of the operating system.

* Set the default boot paths in the NDB on the internal disk. This should be the default location provided by **sysadm** and **admnvram**(1M) commands.

For best boot performance the boot paths should be setup as follows:

    boot_command_1="sd(npsc(pci(x),x),0,0) root –f /dgux –i 3"
    boot_command_2="sd(apsc(pci(x),x),0,0) root –f /dgux –i 3 –t 180"

The first boot path will try to boot the disk after 15 seconds, which is the
default boot timeout of the**–t boot** option. If the CLARiiON is not ready the boot
will fail and then try the second command, which will wait 3 minutes before
attempting to boot.

See the **boot**(8) on–line manual page for additional information on the **–t boot**
option.

5.3.11   Line speeds in half–duplex modes (R4.20MU04)

Half–duplex operation on all PCI synchronous controllers supports baud rates up to
56 KB. Higher line speeds are not supported in half–duplex mode.

5.3.12   Model 7487 Emulex Lightpulse Fibre Channel adapter may fail to configure (R4.20MU04)

On rare occasions, a Model 7487 Emulex Lightpulse Fibre Channel Adapter (**epfc**)
may fail to configure. This problem may correct itself when you configure a LUN on
the adapter, which triggers an attempt to configure the adapter itself. You can also
correct the problem manually by verifying that the cables are connected correctly and
retrying the configuration. **epfc** devices that have been successfully configured are
listed as entries in the **/dev/mpi/fc** directory.

5.3.13   Mylex 8GB setting not supported (R4.20MU04)

The Mylex RAID controller Model 7468 is configured to 2GB disk geometry. The 8GB
disk geometry setting described in Mylex documentation is not supported for booting
DG/UX. If you set your Mylex controller to 8GB using the BIOS configuration utility,
and you attempt to boot from the hard disk, the boot will fail and you will see the
following message: "BOOT2: no bootstrap in partition".

To fix this problem, run the Mylex BIOS configuration utility and set the disk
geometry to 2GB. Using the 2GB setting is appropriate for disks larger than 2GB.
The setting does not prevent the DG/UX System from accessing the entire disk.

> **IMPORTANT** Using the 2GB disk geometry may prevent you from
> performing physical disk copies from a non–Mylex to a Mylex disk. This is a
> known restriction.

5.4 Software notes and warnings introduced in previous R4.20 MUs

5.4.1   Denial of service attack (R4.20)

A User Datagram Protocol (UDP) denial–of–service attack targets vulnerable UDP services, such as UDP/Echo and UDP/Chargen, on a network host. The attacker sends a bogus request packet to a vulnerable service on one of two target hosts. The result is that the target hosts exchange UDP request/response packets at such a high rate as to render the LAN and/or host systems unusable.

The following symptoms may indicate that an attack is in progress:

*   **inetd** uses many CPU resources.

*   Sustained network utilization is much higher than usual.

To alleviate this situation, disable any unneeded UDP services on each host (in particular the **chargen** and **echo** services), and filter at the firewall or Internet gateway any packets that are received from outside the network and have a local source IP address. To filter packets, use the following **sysadm** menu path:
`Networking -> TCP/IP -> Netctrl.`

For more information, use a World Wide Web browser to read:

`http://www.cert.org/advisories/CA-96.01.UDP_service_denial.html`

If you do not have access to the World Wide Web, you can use FTP to access this information:

1. Enter the FTP command.

   ```
   $ ftp ftp.cert.org ↵
   ```

2. When the system prompts for a name and password, enter **ftp** and your e–mail address, respectively.

3. Change to the appropriate directory.

   ```
   $ cd pub/cert_advisories ↵
   ```

4. Get the files you need.

   ```
   $ mget CA-96.01.UDP_service_denial ↵
   ```

5.4.2   dpen driver (R4.20)

The **dpen** driver 10/100 autodetection mechanism may incorrectly determine that 10BASE–T is the highest performance mode of interoperation when the Model 7446 PCI LAN Controller/100 or Model 7458 or 7467 PCI multichannel/100 Ethernet adapter is connected to an IEEE 802.3u 10/100 switch or hub capable of autonegotiation.

You can set the adapter to 100BASE–TX mode through the media parameters of the **dpen** driver. For information on how to set device media parameters explicitly, see the **dpen**(7) on–line manual page.

5.4.3   Packet retransmissions on Fast Ethernet (R4.20)

During certain stress test conditions, Data General observes the potential for eventual dropped network connections. This can occur because retransmission timeouts for packets are discarded in PCI bus overrun conditions. This occurrence is observed for some Fast Ethernet (such as 100BASE–TX) networks that use ports on the Model 7458 PCI multichannel Ethernet adapter.

To correct the problem, you can reduce the load on the Model 7458 PCI multichannel Ethernet adapter by adding one of the following single–port PLC/100 Ethernet PCI adapters to your configuration:

*   Model 7446

*   Model 7463

Or you can replace the Model 7458 PCI multichannel adapter with the new Fast Ethernet Model 7467 PCI multichannel 10/100 adapter, which does not exhibit this problem.

5.4.4   sigwait changes (R4.20)

The **sigwait** system call provided by the DG/UX system behaves as specified in *IEEE Std 1003.4a Draft 6, Thread Extension for Portable Operating Systems*. The final version of that standard (renamed 1003.1c) specifies different behavior for the **sigwait** function. In a future revision of the DG/UX system, the behavior of the **sigwait** function will be changed to comply with the final version of the standard.

To prepare for this change, make sure that you include **<signal.h>** in each compilation that involves the **sigwait** function, after defining the macro **_POSIX4A_DRAFT6_SOURCE** (with **#define** or **–D**).

In addition, recompile object file and library references to **sigwait** with DG/UX R4.11MU01 (or later). This protects current invocations against the intended future change in **sigwait** behavior by making object–level references to **_d6_sigwait** instead of **sigwait**.

Object files and libraries built on DG/UX systems prior to R4.11MU01 make reference to **sigwait** instead of **__d6_sigwait**. Future binary compatibility concerns may require a change in the behavior of the object–level **sigwait** system call. Rebuild such object files and libraries as described in the previous paragraph. If you do not rebuild, your application may fail when relinked with future DG/UX revisions that provide the new behavior of the **sigwait** function, unless you direct the linker to search **libd6_sigwait.a** (as with **–ld6_sigwait**) before **libc.a**.

Existing executable files and shared objects that use the **sigwait** function experience no problem.

5.4.5   S_ISLNK macro defined (R4.20)

Beginning in Release 4.11, the DG/UX system defines the S_ISLNK file attribute macro in **<sys/stat.h>**, according to X/Open Common Applications Environment (CAE) specifications. You should not define a macro of the same name in any program that includes **<sys/stat.h>**. However, you can protect its definition with **#ifndef**, or you can remove the DG/UX macro with **#undef**, as needed.

5.4.6   Unattended reboot on AV 3600 computers (R4.20)

For AV 3600 BIOS revision 4.05.9, you receive the following message during a system boot if the keyboard is unplugged for console redirection:

```
Press SPACEBAR to resume
```

This characteristic of the BIOS is configurable in later revisions through the EISA Configuration Utility (ECU), but with revision 4.05.9, the requirement to press the spacebar means that unattended reboot does not work if the keyboard is unplugged.

To redirect the console and retain support for unattended reboot, you must leave the keyboard plugged in, and configure your default System Control Monitor (SCM) boot path with the –C switch. Note that if you use this method of console redirection and no monitor is connected to the VGA controller on your system, you may lose error messages reported by the Basic Input/Output System (BIOS) or DG/UX bootstraps early in the booting process.

5.4.7   Application Capture Option (R4.20MU01)

The **dgux.aco** package, mentioned earlier in this update notice, is required by a number of third–party software products. However, the ACO package consumes about 4.2 MB of disk and contains features that may interfere with the correct behavior of some existing DG/UX programs, though we are not aware of any specific examples.

If you have no dependencies on ACO or the specific features, you can avoid loading the ACO package or turn off some of its individual features. To avoid loading it, select the package suite less the **aco** package, i.e. **development_less_aco**, when prompted for packages to load during installation.

For more information, including how to turn off individual features, see the Application Capture Option Release Notice.

5.4.8   Binding logical units (LUNs) when disks resized causes system hang (R4.20MU01)

Do not bind LUNs while the system is rebooting. On some systems, it can take a while for LUNs to bind. If, for example, a host is down for maintenance and then is rebooted while the disk is still binding, then the host could hang during sizing of devices and not boot properly.

5.4.9   Installing bootstraps on all disks (R4.20MU02)

Data General recommends installing bootstraps on all disks. If any disks do not have bootstraps installed, you should consider disabling Basic Input/Output System (BIOS) on all adapters that can access those disks. Otherwise, a disk failure could prevent the machine from rebooting.

To learn how to disable BIOS, see the documentation for your specific adapter.

5.4.10   xsysadm bug fixed (R4.20MU02)

An **xsysadm** bug no longer exists. The bug had appeared when the user tried to return to a previous form dialog, and, under certain circumstances, resulted in a core dump.

5.4.11   rc.failover changes (R4.20MU03)

Before receiving Release 4.20MU03, you may have modified your **rc.failover** script to add the command: **admmpl –o start**. The **rc.failover** script in R4.20MU03 supports the start of Multipath LAN relationships at the time of reboot. This support relies upon information stored in the Multipath LAN database.

If you want the Multipath LAN to automatically begin when the system reboots, you should enter the command: **admmpl –o modify –S primary_board** Change your current "No" settings to "Yes" for each entry in the Multipath LAN database.

The system administrator's Mulipath LAN Menu can also be used to set the Multipath LAN Start On Reboot Option. See the **admmpl**(1M) on–line manual page for details on the use of the **–S** switch.

5.4.12   Address requirement for FDDI multipath LAN I/O (R4.20MU04)

When using PCI FDDI adapters in a multipath LAN I/O relationship, both the primary and backup adapters must be configured with an alternate address. An example of a device configuration using an alternate address follows:

```
dpfn(pci(0,1),D,0,08001ba69042)
```

The alternate addresses should be chosen carefully so that there are no duplicate addresses in the network broadcast domain. Specifying an alternate address on the adapter does not free its default address to be used elsewhere in the FDDI ring. See the **dpfn**(7) and **mpl**(7) on–line manual pages for more information.

5.4.13   Installing DBU leads to updated index file formats (R4.20MU04)

Both NetWorker 5.1.x and DBU 5.1.0 use an index file format that is incompatible with that used by NetWorker 4.2.5 and NetWorker 4.2.9. If you currently use either of these earlier versions of Legato NetWorker, your index files will be converted to the 5.1 file format when you first start the DBU after installation. Once the index files are updated to the new format, you cannot run NetWorker 4.2.5 or 4.2.9 unless you use the disaster recovery technique to reconstruct the bootstrap and index files from backups.

5.4.14   Large file support introduces revision lock (R4.20MU04)

The introduction of large file support in R4.20MU04 has created a revision lock. Any application built on R4.20MU04 (or later release) will not execute on DG/UX R4.20MU03 or earlier release.

5.4.15   Misuse of PORT command in the ftp protocol (R4.20MU04)

In this release, the **ftp** daemon is invoked with the **–p** option by default. This closes a security hole associated with the **ftp PORT** command. However, it may have compatibility implications. For instance, **bftp**(1C) will not work with this new default. For more detailed information, see the **ftpd** (1M) on–line manual page.

When this update is installed on a currently running system, the existing **/etc/inetd.conf** file is preserved. Therefore, you must manually add the **–p** option to the line which invokes **ftpd** and restart the **inetd** daemon.

To manually invoke the **–p** option, you can remove the **–p** from the **ftpd** entry in the **/etc/inetd.conf** file. You must then restart the **inetd** daemon, either by rebooting the system or by running the command **dg_kill –HUP inetd** as a user with appropriate privilege.

5.4.16   Navisphere utility blocks AV/Alert messages (R4.20MU04)

The Navisphere utility contains an option called ULOG which, when enabled can cause interaction with the DG/UX System resulting in the loss of AV/Alert messages about CLARiiON problems. To avoid this problem, set the EVENTLOG variable in **/etc/Navisphere/naviagent.config** to NO. For more information, see *Navisphere Agent / CLI for DG / UX Release Notice* (085–001587–01 or greater).

6   Documentation

The DG/UX System R4.20MU05 Documentation CD–ROM (068–600139) contains DG/UX manuals and software for viewing the manuals on–line on a PC running Windows or Windows NT. For instructions on setting up the software and viewing the on–line manuals, see the *Release Notice: DG / UX System Release 4.20MU05 Documentation CD–ROM* provided with the CD–ROM.

6.1 DG/UX Documentation

Two printed manuals ship with the DG/UX System: *Installing the DG / UX® System* and *Preventing and Fixing Problems on the DG / UX® System*. Below is a list of manuals provided on the documentation CD. An asterisk(*) indicates manuals that were revised for the current release.

Part number	Title
069–701142–03	*Analyzing DG / UX® System Performance*
069–701143–01*	*Installing and Managing Printers on the DG / UX® System*
069–701140–03*	*Installing the DG / UX® System*
069–701144–04*	*Managing Mass Storage Devices and DG / UX® File Systems*
069–000698–02	*Managing Modems and UUCP on the DG / UX® System*
069–701136–00	*Managing ONC™ / NFS® and Its Facilities on the DG / UX® System*
069–701137–01*	*Managing TCP / IP on the DG / UX® System*
069–701141–05*	*Managing the DG / UX® System*
069–701059–12	*Porting and Developing Applications on the DG / UX® System*
069–701145–02*	*Preventing and Fixing Problems on the DG / UX® System*
069–701139–00	*Programming in the DG / UX® Kernel Environment*
069–701135–00	*Programming with TCP / IP on the DG / UX® System*
069–701133–02*	*Setting Up and Managing Terminals on the DG / UX® System*

069–701134–00   *Using TCP/IP on the DG/UX® System*

014–002512–01   *Using the AViion® Diagnostics and the AV/Alert^{SM} Diagnostic Support System – DG/UX® Environment*

069–701138–00   *Using the DG/UX® Kernel Debugger*

069–701035–06*   *Using the DG/UX® System*

6.2 Correction to DG/UX driver installation documentation

Note the following correction to the *Addendum to 4515/5515 PCI ATM Adapter Users Guide*, Document No. UG05515–005, REVA (Release Date: October 17, 1996):

On page 9, the top line in Figure 1–1, Sample LECCONF File Contents should read: lecs: 1

The first bullet, beginning on page 9 and continuing on page 10 should read:

*   **lecs**: The **lecs** parameter can have two values: 0 and 1. If the value is 0, the client connects directly to the LES, bypassing the LECS. The address specified in the **atmprefix** and **atmesisel** parameters is the LES address. Otherwise, if the value is 1, the client first connects to the LECS. The address specified in **atmprefix** and **atmesisel** is the LECS address.

6.3 On–line manual pages added in R4.20MU05

The following generic DG/UX manual pages were added in R4.20MU05:

Manual Page	Description
apen(7)	Quad PCI Ethernet Network Adapter

6.4 On–line manual pages introduced in previous R4.20 MUs

6.4.1 Manual pages added in R4.20

The following generic DG/UX manual pages were added for R4.20:

Manual page	Description
adminterrupt(1M)	administer interrupts
adm_intro(1M)	introduction to DG/UX system administrative commands
perl5(1M)	Perl (Practical Extraction and Report Language) overview

Manual page	Description
perlapi(1M)	Perl 5 application programming interface for C extensions
perlbook(1M)	Perl book information
perlbot(1M)	Bag'o Object Tricks For Perl5 (the BOT)
perlcall(1M)	Perl calling conventions from C
perldata(1M)	Perl data structures
perldebug(1M)	Perl debugging
perldiag(1M)	various Perl diagnostics
perlform(1M)	Perl formats
perlfunc(1M)	Perl built–in functions
perlguts(1M)	Perl's Internal Functions
perlipc(1M)	Perl interprocess communication
perlmod(1M)	Perl modules (packages)
perlobj(1M)	Perl objects
perlop(1M)	Perl operators and precedence
perlovl(1M)	Perl overloading semantics
perlrun(1M)	how to execute the Perl interpreter
perlsec(1M)	Perl security
perlre(1M)	Perl regular expressions
perlref(1M)	Perl references and nested data structures
perlstyle(1M)	Perl style guide
perlsyn(1M)	Perl syntax
perltrap(1M)	Perl traps for the unwary
perlvar(1M)	Perl predefined variables
dg_get_label_region(2)	return the MAC region of the given MAC label
dg_mem_get_affinity(3)	get memory placement characteristics
dac(7)	Device Access Controller (Intel only)

6.4.2   Manual pages added in R4.20MU01

The following generic DG/UX manual pages were added in R4.20MU01:

Manual page	Description
admlpsched(1M)	Multiple lp schedulers
epfc(7)	Emulex PCI Fibre Channel Host Bus Adapter
fc–al(7)	Fibre Channel Arbitrated Loop adapters

fsd(7)	Fibre Channel Arbitrated Loop disk
mprc(7)	Mylex PCI RAID Controller
systemtape(1)	Build a bootable DG/UX system tape
pthread_atfork(3)	Register Fork Handlers
gpo(1)	Global Program Optimizer

6.4.3   Manual pages added in R4.20MU02

The following generic DG/UX manual pages were added in R4.20MU02:

Manual page	Description
gamevlog(1M)	Mylex Global Array Manager (GAM) event logger process
gamevent(1M)	Mylex Global Array Manager (GAM) event handler process
gamserv(1M)	Mylex Global Array Manager (GAM) server process
raidbld(1M)	Mylex Global Array Manager (GAM) array recovery process
mprc(7)	Mylex PCI RAID Controller

The following new R4.20MU02 manual pages are links to existing DG/UX manual pages:

New link	Existing manual page
atoll(3C)	**strtoll**(3C)
llabs(3C)	**abs**(3C)
lldiv(3C)	**div**(3C)
strtoll(3C)	**strtol**(3C)
strtoull(3C)	**strtol**(3C)
wcstoll(3C)	**wcstol**(3C)
wcstoul(3C)	**wcstol**(3C)
wcstoull(3C)	**wcstol**(3C)
gfsd(7)	**gscsi**(7)

6.4.4   Manual pages added in R4.20MU03

The following generic DG/UX manual pages were added in R4.20MU03:

Manual Page	Description
acl_library(3)	library of functions that support access control lists

fc–fabric(7)	Fibre Channel fabric adapters
getacl(1)	display access control lists (ACL)
mpl(7)	Multi–Path LAN Pseudo–Driver
pthread_atfork(3t)	register for handlers
setacl(1)	to change the acl of one or more objects
uptime(1)	shows how long the system has been up

6.4.5   Manual pages added in R4.20MU04

The following generic DG/UX manual pages were added in R4.20MU04:

Manual Page	Description
pgen	PCI Gigabit Ethernet Network Interface
largefile_intro(5)	DG/UX support for large files
lock_and_run(1)	Execute commands and arguments while file is locked
puser(1)	Use of file descriptors in report process
fseeko(3S)	Reposition a file pointer in a stream
ftello(3S)	Reposition a file pointer in a stream

7   Software distribution

The DG/UX System is released on compact disk. The model and part numbers are as follows:

Model number	Part number	Description
P001A–D	068–600216	Software CD–ROM
	090–601168	boot diskette

8   Installation instructions

To install the R4.20MU05 maintenance update, read this section and then follow the instructions in the manual *Installing the DG/UX® System*. This update notice provides information that supplements the *Installing* manual.

For the initial loading of the DG/UX release medium, you need to know whether your computer contains a System Control Monitor (SCM). The following table tells you whether or not your system has an SCM:

System model	SCM
AV 1600	No
AV 2000	No
AV 2100	No
AV 2600	No
AV 3000	No
AV 3600, 3600R	Yes
AV 3650, 3650R	Yes
AV 3700, 3700R	No
AV 4700	Yes
AV 4800	Yes
AV 4900	Yes
AV 5800	Yes
AV 5900	Yes
AV 6600	Yes
AV 20000	No
AV 25000	No

8.1 Local boot device and system disk names

The following table lists typical local boot device and system disk names for AViiON models.

AViiON system	CD–ROM	System disk
1600	sd(apsc(pci(0),d),4)	sd(apsc(pci(0),d),0)
2000	sd(apsc(pci(0),b),4)	sd(apsc(pci(0),b),0)
2100	sd(apsc(dppb(pci(0),14),5),5)	sd(apsc(dppb(pci(0),14),5),0)
2600	sd(apsc(dppb(pci(0),14),5),5)	sd(apsc(dppb(pci(0),14),5),0)
3000	sd(apsc(pci(1),d),4)	sd(apsc(pci(1),d),0)
3600, 3600R	sd(apsc(pci(1),c),5)	sd(apsc(pci(1),c),0)
3650, 3650R	sd(apsc(pci(1),b),5)	sd(apsc(pci(1),b),0)
3700, 3700R	sd(npsc(pci(0),8),5)	sd(npsc(pci(1),3),0)
4700, 4800	sd(npsc(cpci(0),2),4)	sd(npsc(cpci(0),3),0)
4900	sd(apsc(pci(1),b),4)	no default
5800	sd(apsc(pci(1),b),4)	sd(apsc(pci(1),c),0)

5900	**sd(apsc(pci(1),b),4)**	**no default**
6600	**sd(apsc(pci(1),b),4)**	**sd(apsc(pci(1),b),0)**
20000	**sd(apsc(dppb(pci(0,1),f),f),4)**	**sd(apsc(dppb(pci(0,1),f),f),0)**
25000	**sd(apsc(pci(0,0),A),6)**	**fsd(epfc(pci(0,2),4,0,7D),0,x,0)***

* This example is for a system that includes an FC arbitrated loop. If your system includes a FC switch, use the following example: fsd(epfc(pci(0,2),4,0,200013),200413,x,0).

8.2 Preparing virtual disks

The DG/UX system requires the **swap**, **root**, and **usr** virtual disks in order to operate. You must create these three virtual disks at the beginning of the installation process.

If you install optional software packages that ship with the DG/UX system, the installer program creates additional standard virtual disks. Create the additional virtual disks during the package load phase of installation.

We recommend that you create a virtual disk for a specific dump–to–disk area also. This dump–to–disk area may be created on any available SCSI system disk and should be called **dump**.

The following table gives the exact virtual disk names, mount point directory, and minimum size in blocks required by each of these disks.

Virtual disk name	Mount point	Minimum size	Required
swap	none	150,000	yes
root	/	70,000*	yes
usr	**/usr**	320,000	yes
dump (optional)	**none**		no
usr_opt_X11	**/usr/opt/X11**	100,000**	for X Window System
usr_opt_sdk	**/usr/opt/sdk**	60,000	for **sdk** and **sdk.X11** packages
usr_opt_networker	**/usr/opt/networker**	60,000	for Data Backup Utility package
var_opt_networker	**/var/opt/networker**	50,000	if host is NetWorker server

*The **root** virtual disk requires 70,000 blocks on one physical disk. If you are upgrading (overloading), the installer program tries to enlarge the **root** virtual disk to this size. If your disk does not have enough free blocks available to enlarge **root** sufficiently, you will need to move, delete, or shrink another file system on **root**'s disk, or move **root** to another physical disk with enough available space.

If you do not load the **sdk package, then **usr_opt_X11** is 80,000.

8.3 Ethernet controller device names

When your computer contains two or more Ethernet LAN controllers, you must specify a device that connects your computer to a TCP/IP LAN.

You identify controllers by the controller type and controller number, using the form *controller–type controller–num*. The following list provides Ethernet LAN controller device names.

Controller description	Name	Name example
DEC PCI 10/100 LAN controller for 10– or 100–Mbit Ethernet	**dpen**	**dpen(pci(0),d)**
Packed Engines Gigabit Ethernet Controller	**pgen**	**pgen(pci(0),9)**
Intel PCI 10/100 LAN controller for 10– or 100–Mbit Ethernet	**ipen**	**ipen(pci(0),9)**
DEC PCI Multi–channel LAN controller for 10– or 100–Mbit Ethernet	**dpen**	**dpen(dppb(pci (0),D,0)4,0)**
Adaptec 64 bit PCI Multi–channel LAN controller for 10– or 100–Mbit Ethernet	**ipen**	**ipen(pci(0),9)**

8.4 POSIX ACLs

POSIX ACL installation does not occur during the installation of the R4.20MU05 maintenance update. Thus, if you plan to use POSIX ACL, you must install this feature after installing the maintenance update. To do this, select the option to load the **acl** package suite, which includes the POSIX ACL packages.

9   Preparing a software trouble report

If you believe that you have found an error in the DG/UX® System or its documentation, consult the manual *Preventing and Fixing Problems on the DG/UX®*

System. This manual which is provided in both printed and CD–ROM versions with your system, explains the procedures for filing a software trouble report (STR).

A copy of the standard STR form suitable for printing on a line printer is in the file **/usr/release/STR_form**. For the name, model number, and release, enter DG/UX® System, P001A, and R4.20MU05.

To comment on the documentation, fill out the reader comment form in **/usr/release/doc_comment_form** and send it to the address on the form.

<div align="center">End of Notice</div>

Index

Symbols

−p option, 50

−t boot option, 45

Numbers

3D Board, boot command strings, 35

A

Access control list (ACL), 36
  dgux.acl package, 10

ACO, Application Capture Option, 10

Adapter
  Adaptec PCI, 42
  ATM, 27, 40
  Dual Channel Differential SCSI, 28
  Model 24025, 42
  Model 7446, 46, 47
  Model 7463, 47
  Model 7458, 46, 47
  Model 7467, 46
  Model 7459 restrictions, 41
  Model 7460 restrictions, 41
  Model 7472 restrictions, 41
  Model 7473 restrictions, 41
  Model 7467 avoids problem, 47
  Model 7469, 28
  Model 7472, 28
  Model 7473, 28
  Model 7487, Fibre Channel Firmware
    Download, 16
  Model 7487 fails to configure, 45
  Model 7497, 28
  Model 7516, PCI Quad Channel
    Ethernet, 17

Model 7517
  Fibre Channel, 17
  Fibre Channel Firmware Download, 16
Model 7522, 17
Model 7523, 17, 18
PCI ATM, 28
PCI Ethernet LAN, 27

Address Bit Permuting (ABP), 42

address space size, 32

admdevice command, 32

admipinterface command, 33, 34

admmpl −o modify −S primary_board
    command, 49

admmpl −o start command, 49

admnvram command, 36, 44

Application Capture Option (ACO), 10
  deinstalling, 10
  dgux.aco package, 10, 34, 48
  packages, 12
  X11.aco package, 11

ATI graphics card, 28

ATM adapter support, 27

ATM signaling software, 42

atmesisel parameter, 52

atmprefix parameter, 52

Attack, denial of service, 46

AV 20000 computer, 26
  32−CPU, 26, 27
  boot command strings, 35
  boot timeout, 35
  halt \\−r, 36
  new DG/UX features for, 35
  no DG/UX 3.5 diskette driver, 42, 44
  scm device, 36

AV 25000 computer, 27

W

Warnings, 38

wdt () device driver, provides watchdog
  timer, 24, 25

X

X Window System, 11

xsysadm program, bug fixed, 49

Y

Year 2000
  ntpq command, 38
  readiness, 34
  testing, 40