

Release Notice: DG/UX® System Release 4.20

July 1997

Part number 085-600492-00

This release notice applies to the following model:

P001A

Copyright © Data General Corporation 1997.

Unpublished—all rights reserved under the copyright laws of the United States and all other countries.

Printed in the United States of America.

Licensed material—property of Data General Corporation

Restrictions and Trademarks

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

Restricted Rights Legend

Use, 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 at Federal Acquisition Regulations (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. All other trademarks mentioned herein are the property of their respective owners.

Release Notice: DG/UX® System Release 4.20

085-600492-00

Revision History:

Effective with:

Original release – July 1997 DG/UX System Release 4.20

Contents

1	Introduction	7
2	Product description	7
2.1	ccNUMA technology	7
2.2	Application Capture Option (ACO)	9
2.3	Legato Data Backup Utility	9
3	Environment	10
3.1	Hardware	10
3.2	Software	10
3.2.1	Configuration constraints	10
3.2.2	Using layered software	11
3.3	Firmware required	11
4	Enhancements and changes	12
4.1	New hardware features	12
4.1.1	AV20000 server	12
4.1.1.1	3D Board device driver for AV 20000	12
4.1.1.2	Model 7468 Qlogic Differential SCSI Controller	13
4.1.1.3	Enhanced dump support	13
4.1.2	Pentium Pro processor	13
4.1.3	Model 7466 Synchronous controller	14
4.1.4	12X CDROM	14
4.2	New software features	14
4.2.1	New features in the Virtual Disk Manager (VDM)	14
4.2.1.1	Automation of multi-path disk I/O administration	15
4.2.1.2	Shorter disk names	18
4.2.1.3	Increased length limit on virtual disk names	19
4.2.1.4	Option to have the VDM zero disk blocks	20
4.2.1.5	New dynamic configuration variables	20
4.2.1.6	Halt dumps stored on virtual disks	20
4.2.2	Scalable network multi-connect (MCL)	21
4.2.3	Advanced 3-D system management software for the AV20000	22
4.2.4	Commands	22
4.2.4.1	Removal of admiopath(1M) and /etc/iopaths.params file	22
4.2.4.2	admpdisk and multi-path disk I/O administration	23

Contents

4.2.4.3	New admpdisk operations	23
4.2.4.4	admpdisk, admvdisk, and sysadm recognize short disk names	24
4.2.4.5	Modified admdevice operations	24
4.2.4.6	Modified admpdisk operations	24
4.2.4.7	admvdisk operations modified	25
4.2.4.8	New date format in SCCS commands	25
4.2.4.9	BSD 4.4 telnet ported to the DG/UX System	26
4.2.4.10	New VDM ioctl commands	26
4.2.4.11	gcc and cc command changes	26
4.2.4.12	New adminterrupt(1M) command	27
4.2.5	Fewer dots printed during disk registration on large systems	27
4.3	Perl version 5	27
5	Notes and warnings	28
5.1	Kernel builds take longer with AV20000	28
5.2	Floppy driver not supported on AV 20000	28
5.3	Line speeds in half-duplex modes	28
5.4	AV 6600 memory error reporting	28
5.5	Undesirable /etc/devlinktab results	29
5.6	SCSI device restriction	29
5.7	NetWorker documentation and Autochanger licenses	29
5.8	Unattended reboot on AV 3600 computers	29
5.9	PCI buses	30
5.10	dpen driver	30
5.11	Packet retransmissions on Fast Ethernet	31
5.12	ATM adapter limitations and restrictions	31
5.13	ATI Mach32 graphics accelerator	32
5.13.1	Set up your system for automatic dumps	33
5.13.2	Avoid taking hot key dumps	33
5.13.3	Terminating the X server	33
5.14	X server memory usage	34
5.15	/dev file system space	34
5.16	sigwait changes	34
5.17	Denial of service attack	35
5.18	S_ISLNK macro defined	36
5.19	Changes made to internal data structures	36
5.20	AV/Alert	37
5.21	Using QIC tapes for multi-tape dumps	37
6	Documentation	37
6.1	Correction to DG/UX driver installation documentation	38

6.2 New man pages 38
7 Software distribution 39
8 Installation instructions 40
9 Preparing a software trouble report 40

1 Introduction

This release notice provides you with specific information about the DG/UX® System Release 4.20 operating system that is not available in the standard DG/UX documentation set. An on-line version of this release notice, suitable for line printers, is in the file `/usr/release/dgux_R4.20.rn`. A post-script version is in the file `/usr/release/dgux_R4.20.rn.ps`

2 Product description

The DG/UX system is the operating system for Data General's family of AViiON® computers. This release of the DG/UX system includes support for all Intel-based AViiON computers.

This release does not support Motorola 88000 series processors. The 88K-based systems will continue to be supported by maintenance updates to DG/UX R4.11.

This release is provided on CD-ROM.

Important Note: DG/UX B2 Security Option R4.20 and C2 Trusted DG/UX Option R4.20 are shipped on their own media, separate from the DG/UX R4.20 system. For information on installing these security products, refer to *Release Notice: DG/UX B2 Security Option Release 4.20* or *Release Notice: C2 Trusted DG/UX Option R4.20*.

2.1 ccNUMA technology

ccNUMA is an acronym for Cache-Coherent Non-Uniform Memory Access. This phrase describes a class of computer systems in terms of the way memory is accessed on the system. Cache-Coherent means that the hardware keeps the contents of the various caches in the system coherent with each other and with primary memory; software need not concern itself with this. Non-Uniform Memory Access means that the time required to access different memory locations from any given processor or I/O controller may not be uniform.

As applications continue to grow, there is a continuing need to design computers that scale well without having to rewrite application code. The obvious solution of adding more processors does not work because of system bus congestion. While system buses can be made faster, the cost of doing so soon makes this solution impractical, especially for low-end systems.

The bus is the limiting component to SMP scalability. In a busy system, any attempt to add more processors to the bus will result in bus congestion and lower system performance.

ccNUMA technology connects multiple SMP locales together in such a way that there appears to be a single address space. This is NOT a cluster. From an application point of view, more processors and a larger single address space have been made available. Physical memory is divided across SMP locales. Any processor or I/O controller can access memory in its locale as easily and as quickly as before. In addition, any memory recently accessed from another locale is saved in a local far memory cache (FMC) for quicker access the next time.

Any processor can access memory located anywhere in the system (depending, of course, on process rights and privileges). Cache coherency is handled automatically by the memory system hardware. Access to memory in another locale, however, can take up to 10 times as long as access to local memory. This is the "non-uniform" part of ccNUMA. Obviously, this can be a bad thing if it happens too frequently. A certain part of this problem is solved by the hardware through the use of the Far Memory Caches (FMC). Additionally, the DG/UX operating system has been enhanced to work within the constraints imposed by the NUMA architecture.

To accommodate ccNUMA architectures, the operating system is called upon to assume a variety of new responsibilities:

- Inventory the system's hardware at boot time, and maintain a representation of it for use by the various kernel subsystems and application programs.
- Keep track of physical memory on a per-locale basis, and provide mechanisms for allocating and releasing memory for various purposes from the different locales, according to various strategies.
- Keep track of processors on a per-locale basis, and arrange to spread the load so that processing power is used effectively.
- Recognize redundancy in I/O facilities (e.g., multiple SCSI buses to a single CLARiiON disk unit) and manage these resources to provide optimum performance and availability.

- Distribute various internal data structures (including pools of special-purpose structures that are frequently allocated and released) across locales so as to minimize cross-locale memory accesses, and arrange to use such resource from the correct locale whenever it is needed.
- Provide a variety of interfaces to user space, so that NUMA-aware application programs can configure their own behavior to accommodate the ccNUMA architecture, and so that a system administrator may arrange for non-NUMA-aware applications to be executed in ways that optimize performance.
- Provide a plethora of system performance information, so that the system's behavior can be monitored.

Specific features and enhancements are discussed throughout this release notice.

2.2 Application Capture Option (ACO)

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

Note: If you accidentally load the ACO package with the DG/UX system, the only clean way to remove it is to reload the DG/UX system without the ACO package.

2.3 Legato Data Backup Utility

The DG/UX system provides a restricted version of NetWorker, called Legato Data Backup Utility. This utility enables you to back up your local file systems to a single local backup device. The Legato Data Backup Utility can be upgraded to NetWorker for DG/UX Server Edition or NetWorker for DG/UX Network Edition by purchasing the appropriate licenses. See the NetWorker release notice for more information.

Note: The Legato manuals are provided in hardcopy with this release. See Section 6, Documentation, for part numbers.

3 Environment

This section describes the minimum hardware, firmware, and software required to run DG/UX System Release 4.20.

3.1 Hardware

The minimum equipment configuration needed for running this release is listed below:

- 32 MB of memory
- 500 MB of hard disk space
- CD-ROM drive
- floppy disk

3.2 Software

This section lists the software necessary for this revision of the DG/UX system.

3.2.1 Configuration constraints

This release can be installed on a new system or overloaded on DG/UX System Release 4.11 or any maintenance update to R4.11.

The manual *Installing the DG/UX System* discusses package sizes and disk space requirements. Consult this manual before beginning installation.

This release includes fixes that were previously released as software patches. Do not attempt to install any patches, patch kits, or maintenance updates for previous releases on this release.

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

3.2.2 Using layered software

Some layered products may require a particular DG/UX release. Please check your layered product documentation for dependencies.

3.3 Firmware required

DG/UX for Intel requires firmware revision numbers equal to or greater than those listed below for proper operation.

AV3000: BIOS 1.00.03, SCU 3.11ovl39
AV3600: BIOS 5.10.5, TSB 6.8.0, ECU 3.12
AV3600R: BIOS 5.10.5, TSB 6.8.0, ECU 3.12
AV2600: BIOS 1.00.04, SCU 3.60
AV6600: BIOS 5.10.6, TSB 7.10.0, ECU 3.17
AV2000: BIOS 1.00.07, SCU 3.08ovl 31
AV2100: BIOS 1.00.04, ECU/SCU 3.60
AV1600: BIOS 1.00.03, SCU 3.60 ovl 1.10
AV4700: BIOS 5, SCM86 2.6.0
AV4800: BIOS 5, SCM86 2.6.0
AV5800: BIOS 5, SCM86 2.6.0
AV4900: BIOS 1.00.05, TSB 5.3.0, SCU 350-140
AV5900: BIOS 1.00.05, TSB 5.3.0, SCU 350-140
AV20000: BIOS 0.20, VCS 5.0 SMM 0.9

Note: The M3D firmware revision numbers are provided in the SMM release notice.

CLARiiON firmware revision 8.15 is required for use in multi-initiator SCSI configurations, including clusters.

4 Enhancements and changes

This section lists the major enhancements and changes specific to Release 4.20. Enhancements are new features or functionality. A change may be a software problem that has been fixed or a change required for consistency.

4.1 New hardware features

The following sections describe new hardware features for this release.

4.1.1 AV20000 server

The AV20000 server is supported by this release. This is DG's first Intel system based on ccNUMA technology. 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. The AV20000 brings new levels of scalability to DG's family of SMP products in terms of not only processor count, but also memory and I/O. It offers a growth path from a quad-processor entry-level server through a 32-processor server suitable for large data warehousing or OLTP applications. It provides balanced system performance and investment protection through a modular upgrade path.

4.1.1.1 3D Board device driver for AV 20000

A ccNUMA system is composed of separable functioning PCs, each with its own keyboard, VGA, mouse, and floppy interface. The 3D Board provides a method of multiplexing the output of these PCs to a single keyboard/VGA/floppy/mouse interface. The 3D Board routes all of the keyboard/VGA/mouse/floppy I/O to a management PC that virtualizes these devices.

The 3D Board contains a LAN controller that is used to communicate between DG/UX and a management PC. The device has read and write interfaces and is used for the service manager and other clients to send and receive configuration change and error information.

Note: The term "3D Board" is synonymous with the term "DAC."

4.1.1.2 Model 7468 Qlogic Differential SCSI Controller

This release adds DG/UX support for the Model 7468 Qlogic Differential SCSI controller. This controller supports differential SCSI devices in AV20000 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. Refer to the **qpsc** man page for more information.

4.1.1.3 Enhanced dump support

The AV20000 supports memory up to 32GB. Dump support has been enhanced to support this large memory size. The increase in memory support has been provided for the rare occasions when a full memory dump is necessary (kernel memory and user memory). These occasions should be limited to when someone at DG Support requests a full dump.

If DG Support requests a full dump, that is, user and kernel memory, the new hardware and dump support will support a large system dump. You should, however, consider the following:

- You must have the physical space to hold all the physical memory that results from a large system dump.
 - a. When dumping to a tape drive, you may need multiple cartridges.
 - b. When dumping to a disk, you need to have enough space on one disk, preferably (but not necessarily) contiguous space.
- Taking a large system dump is very time consuming. The actual amount of time depends on the transfer rates of your system and the total amount of physical memory involved, but even dumping to disk, which is usually quick, could take up to an hour or more for the largest memory configurations.

4.1.2 Pentium Pro processor

With this release, there is a shift in optimization focus 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.

Data General has overcome the memory limit of 32-bit systems by using the full 36-bit physical address capabilities of the Pentium Pro microprocessor on its DG/UX operating system. Applications running on DG/UX can access up to 32GB of physical memory (32GB of memory - 4GB per module -- is supported on AV20000 systems). This very large memory support can translate into transaction processing performance gains of up to 40 percent.

4.1.3 **Model 7466 Synchronous controller**

This release adds DG/UX 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). This controller is supported on all Intel AViiONs. Refer to the **apsy man** page for more information.

4.1.4 **12X CDROM**

The 12X CDROM, model 61018, is supported on all AViiONs with this release.

4.2 **New software features**

This section discusses support for new or enhanced software.

4.2.1 **New features in the Virtual Disk Manager (VDM)**

The Virtual Disk Manager (VDM) has been enhanced to support the ccNUMA technology. Enhancements to the VDM provide improved system performance, higher availability to disk data, and easier disk path location. The following sections detail the changes and enhancements to the Virtual Disk Manager for this release.

4.2.1.1 Automation of multi-path disk I/O administration

With this release, Data General introduces a fully connected multi-path input/output (mpio) model.

Multi-path disk I/O provides protection from single points of failure in disk controllers, I/O buses, cables, and CLARiiON(R) Storage Processors (SP). It does not provide protection from failures of the disk itself, however.

Multi-path disk I/O (mpio) for physical disks has been re-implemented to better support NUMA configurations. Most operations related to managing multi-path physical disks are now automated, making it much easier to use multi-path disks to provide high-availability to your data.

The previous implementation of mpio involved using the **admiopath(1M)** command to manage the multi-path I/O database stored in the **/etc/iopath.params** file. On installation of DG/UX R4.20, both the command **admiopath(1M)** and **/etc/iopath.params** file are removed from the system. Those physical disks with entries in the old database are converted to the newer implementation. Physical disk multi-path attributes are now stored on the physical disk itself and managed via the **admpdisk(1M)** command. The database is therefore no longer required.

In the case of the DG/UX cluster, physical disks configured for multi-path I/O using the old implementation will all be converted to the newer multi-path implementation as part of the rolling upgrade procedure for DG/UX clusters. For non-cluster systems, the conversion is done as part of the installation.

There are two new components of the automated mpio feature.

- Disk mpio mode selection
- Load-balancing mode selection

The options for each are described in following sections.

The mpio options for both disk mpio mode and load-balancing mode are controllable through **sysadm(1M)** choices in the Device -> Disk -> Physical -> Multi-Path Disks menu. All of the features are also controllable with the **admpdisk(1M)** command.

All of the mpio modes and load-balancing modes are fully supported on CLARiiON(R) disks. Both SPs in the CLARiiON(R) must be running microcode revision 8.00 or later in order to use the mpio modes. The mpio modes and load-balancing modes are only partially supported on non-CLARiiON(R) disks. In particular, if a standalone disk is in single path active mpio mode, you cannot register alternate paths to it. This is because non-CLARiiON(R) disks do not support the new disk signature mechanism that made these new mpio modes possible.

All of the mpio modes and load-balancing modes are fully supported on cluster disks. The same mpio mode and load-balancing modes will be used automatically on all nodes in the cluster. Any changes to the mpio modes will be propagated automatically to the other nodes in the cluster as well.

Note: If you have enabled any of the new mpio or load-balancing modes for a disk, the disk cannot be used on a pre-R4.20 system. Turn off the mpio and load-balancing modes for the disk before you move it to a pre-R4.20 system.

See the following sections and the **admpdisk(1M)** and **vdmpphys(7)** man pages for more details.

New disk mpio modes

All you have to do is put the disks into the mpio mode that you want to use, and the DG/UX system will take care of the rest. There are now three disk mpio modes to choose from:

- **No mpio:** There is only one path to the disk, and it is used for all I/O to the disk. No mpio protection is provided.
- **Single path active mpio:** This is the equivalent of the functionality provided by the pre-R4.20 mpio features, where there are several paths to the disk, but only one of the paths is used for I/O at any time. All the alternate paths are considered backup paths and are used only if the active path fails.

This is the default mode when you format new disks on single-locale systems (non-NUMA systems). This means that mpio is automatically enabled when you format new disks.

Existing mpio disks defined in `/etc/iopath.params` will be converted to this mode during installation when you upgrade from R4.11 to R4.20.

- **Multiple paths active mpio:** This mode allows I/O to the disk through multiple paths simultaneously. For example, if you had two disk controllers attached to the same Storage Processor (SP) on a CLARiON(R) disk array, I/O could be done through both of them at the same time.

This is the default mode when you format new disks on multiple locale systems (NUMA systems). This means that mpio is automatically enabled when you format new disks.

This is especially important on NUMA systems, allowing I/O from each locale in the system to be executed through the disk controllers within each locale. This provides the best system performance on NUMA systems.

All of the mpio modes can be used on all types of systems, both NUMA and non-NUMA. For best system performance on NUMA systems, we recommend using multiple paths active mode with load balancing enabled.

Warning: If you are running on a DG/UX system with information security, and you are using the Disk Failover Package to provide disk failover between hot swap systems, you **MUST NOT** put any failover disk into multiple paths active mpio mode. Doing so allows both failover hosts to access the disk simultaneously and leads to destruction of the data on the disk. Use only single-path active mpio mode if you are running on a DG/UX system with information security and using the Disk Failover Package.

The default `sysadm(1M)` and `admpdisk(1M)` listings now present only one path for each disk. This makes it much easier to identify the disks on the system, because you do not have to wade through all the alternate paths to the disks. You can still list the alternate paths with various listing options when you need them.

Each time the system boots, it finds all of the paths to the disks and sets them up automatically. Disks may now have an unlimited number of paths. Previously, they were limited to four paths each.

As you manipulate new disks you add to the system, the system finds the alternate paths to the disk automatically and handles them appropriately. For example, when you format a new disk, all of the alternate paths to the disk are registered automatically for you. When you deregister a disk, all of the alternate paths to it are deregistered automatically.

If you bring an already formatted disk onto the system, you need to put it into the `mpio` mode you want to use. From then on, the paths to the disk are handled automatically.

The `admiopath(1M)` command has been removed in R4.20, as has the `/etc/iopath.params` file. The `mpio` modes for a disk are now stored on the disk itself in the Virtual Disk Information Tables, which are managed by the DG/UX system.

New load-balancing feature for disk `mpio`

There is also a new load-balancing feature for disk `mpio`, with two configuration choices:

- **No load balancing:** With load balancing turned off, I/O for a given disk is executed using only one active controller in each locale. On non-NUMA systems, only one active controller in the system is used for each disk.

This is the equivalent of the functionality provided by the pre-R4.20 `mpio` features. It is also the default mode when new disks are formatted on single locale systems (non-NUMA systems).

- **Round robin load balancing within locale:** With this load-balancing mode turned on, I/O is executed on all active controllers in a locale for a given disk alternately, in a round robin fashion. This balances the load among the controllers and provides better system performance.

This is the default mode when new disks are formatted on multiple local systems (NUMA systems).

4.2.1.2 Shorter disk names

Disks now may have short disk names to make them easier to identify in `sysadm(1M)`, `admpdisk(1M)`, and `admvdisk(1M)` commands. Physical disk short names provide a way to identify registered physical disks. The shorter names are intended primarily to ease the administration of multi-path physical disks and may be assigned to any registered physical disk regardless of its multi-path mode.

By default, **sysadm(1M)** and **admpdisk(1M)** give a unique, small integer as a name to each new disk encountered. The first disk encountered on the system is named "disk(0)", the second is "disk(1)", and so forth. The utilities also present these short disk names by default in all listings because they are much easier to use than long disk path specifications, such as "sd(apsc(pci(0),B,0,3),0,0)".

You can use these short disk names in almost all **sysadm(1M)**, **admpdisk(1M)**, and **admvdisk(1M)** commands while you are manipulating physical and virtual disks. The names can be used whenever the disks are registered on the system.

The situations requiring long disk specifications are as follows:

- In boot command lines to the bootstrap. You must use a long disk specification when booting the system because the disks are not registered while the bootstrap is running.
- In halt dump virtual disk specifications. Because disks are no longer registered when the system is taking a halt dump, you must use a long disk specification when taking a halt dump.
- When configuring, formatting, or registering new disks into the system. You must use a long disk specification because the disks are not yet registered when you perform these operations.
- In Membership Manager Database entries on cluster systems. The database entries are used by the bootstraps when the system is booted, so you must use long disk specifications in the entries.

In all other cases you can use the short disk names instead of the long disk specifications. Using the short disk names makes creating, removing, and manipulating virtual disks much easier.

4.2.1.3 Increased length limit on virtual disk names

The length limit on virtual disk names has been raised from 31 to 63 characters in DG/UX System Release 4.20.

If you rename an existing virtual disk and give it a long name, or create a new virtual disk with a long name, you cannot use that disk on a pre-R4.20 system. Rename the virtual disk to a name 31 characters or less before you move the disk to a pre-R4.20 system.

In a cluster system, long virtual disk names cannot be used until a rolling upgrade of all nodes to R4.20 has been completed and you have changed the cluster revision to 1000.

4.2.1.4 Option to have the VDM zero disk blocks

You can now have the VDM zero the disk blocks in virtual disks when they are deleted. This prevents the data from being reused in a new virtual disk that overlaps some or all of the disk blocks that used to be in the deleted virtual disk. This is a security feature that ensures the data cannot be read again after it has been deleted.

There are several other situations where disks blocks will be zeroed if the feature is enabled. The feature is disabled by default on most systems. See the **vdm(7)** man page for more details.

4.2.1.5 New dynamic configuration variables

There are four new dynamic configuration variables that control various VDM operations:

- AFFINE_IO_TO_THREAD
- FAVORLOCALINACTIVEPATHS
- IOSTATS
- ZERODISKBLOCKS

They can be set in the system configuration file, or changed dynamically while the system is running. See the **vdm(7)** and **vdmphys(7)** man pages for more information on these variables.

4.2.1.6 Halt dumps stored on virtual disks

Virtual disks used to store halt dumps can now be unstriped aggregations of simple partitions, as long as all of the pieces are on the same physical disk. Up to eight pieces are supported. Previously they could be simple partitions only.

4.2.2 Scalable network multi-connect (MCL)

Multi-Connected LAN (MCL) provides multiple I/O paths to a single Ethernet network. Using MCL reduces the cache-coherency traffic on the AV20000 SCI interconnect and enhances network availability. It also increases system bandwidth in switched networks.

Multi-Connected LAN (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 DMA and reduces the traffic on the SCI interconnect. MCL also load balances client access to the ccNUMA system by responding to 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 may be more effective if each interface has its own IP address. Lastly, in some OLTP environments it may 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(1M)** 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(1M)** command has a new option, **-s subdevice_list**, which is used only when the **-d mcl** option is specified. 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(1M)** 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
```

Note: MCL is designed for switched LANs; it should not be used on shared LANs.

4.2.3 Advanced 3-D system management software for the AV20000

M3D (Monitor, Manage, and Maintain) software runs on a Windows NT PC. This is the console for AV20000. It allows for both multiple system management from a single workstation and remote system management. The management system software can be integrated with other DG management products, including Data General Enterprise Management (DGEM) products for the DG/UX operating system such as DG/UX CLARiiON Manager. The management system software is for the AV20000 system only.

4.2.4 Commands

This section describes commands and their operations that have been added, deleted, or modified. It also highlights new features associated with specific commands.

4.2.4.1 Removal of `admiopath(1M)` and `/etc/iopaths.params` file

The previous implementation of multi-path I/O involved using the `admiopath(1M)` command to manage the multi-path I/O database, stored in the `/etc/iopath.params` file. On installation of DG/UX R4.20, both the command `admiopath(1M)` and `/etc/iopaths.params` file are removed from the system. Physical disk multi-path attributes are now stored on the physical disk itself and managed via the `admpdisk(1M)` command.

The `admiopath(1M)` command was also used to manage multi-path LAN devices as well. That functionality has been moved to the new command `admmpl(1M)`. The database that holds the configurations for multi-path LAN devices are now in the file `/etc/mpl`. The conversion for the LAN configuration from the old database to the new is transparent as well. See the `admmpl(1M)` manual page for details on multi-path LAN devices.

In a cluster system, these changes are not done until you have changed the cluster revision to 1000.

4.2.4.2 **admpdisk and multi-path disk I/O administration**

Most operations related to managing multi-path physical disks are now automated. The **admpdisk(1M)** command controls the new features in the Virtual Disk Manager (VDM). See the **vdmpphys(7)** and **admpdisk(1M)** manual pages for details on multi-path physical disks and the section on "Automation of multi-path disk I/O administration" in this notice for more information.

4.2.4.3 **New admpdisk operations**

The following are new **admpdisk(1M)** operations.

- The **-o modify** operation can be used to manage most attributes of registered physical disks. These include the writable, cluster, and multi-path attributes. The **-o cluster** and **-o uncluster** operations are therefore obsolete and will not be supported in the next major release of the DG/UX system. The **admpdisk(1M) -o modify -c** operation should be used to manage the cluster attribute of physical disks.

The **-o modify** operation can also be used to activate, fail, and repair physical disk paths.

- The **-o rename** operation can be used to rename a physical disk. A disk short name can be assigned or removed.
- The **-o expand_vdit** operation can be used to expand the Virtual Disk Information Tables (VDITs) for physical disks. VDITs hold the virtual disk layout of a physical disk. VDITs are created when the physical disk is formatted. The default size depends on the disk size and is large enough for most physical disk uses. However, you can use the **-o expand_vdit** operation in cases where more VDIS space is required.
- The **-o switch_locale** operation can be used to manually control where physical disk I/O is issued on a per locale basis. The VDM manages where I/O is issued and performs automatic locale switches as needed, so this operation should be used rarely.

4.2.4.4 **admpdisk, admvdisk, and sysadm recognize short disk names**

Note that disk short names are recognized by management utilities such as **admpdisk(1M)**, **admvdisk(1M)**, and **sysadm(1M)** as long as the physical disk is registered. Unregistered physical disks must be addressed using a device specification. See "Shorter disk names" in the "New Features in the Virtual Disk Manager" section in this section for more information on this feature.

4.2.4.5 **Modified admdevice operations**

The following **admdvce(1M)** operations have been modified for R4.20.

- The **-o configure** operation automatically detects and configures all paths to a physical disk device. The new **-p** option can be used to override this behavior and force only the specified path(s) to be configured.
- The **-o deconfigure** operation automatically detects and deconfigures all paths to a physical disk device. The new **-p** option can be used to override this behavior and force only the specified path(s) to be deconfigured.

4.2.4.6 **Modified admpdisk operations**

The following **admpdisk(1M)** operations have been modified for R4.20.

- The **-o initialize** operation enables newly formatted physical disks for multi-path I/O by default. The actual mpio mode depends on the system configuration. For NUMA configurations with multiple locales, the disks are put into multiple paths active mode. For traditional UMA configurations or single locale systems, the disks are put into single path active mode.

The **-o initialize** operation also gives physical disks a simple integer disk short name like "disk(4)". The actual integer selected is based on the current pool of registered physical disks. The new **-N** option can be used to override this behavior and assigns the argument provided with the **-N** option as the name of the physical disk. The name provided can optionally be indexed by also including the **-i** option. The index selected is based on the

current pool of registered physical disks with the same prefix name.

- The **-o register** operation automatically registers all paths for multi-path physical disks. The new **-p** option can be used to override this behavior and force only the specified path to be registered. Additionally, if given the name of a disk that is already registered, the **-o register** operation will attempt to discover and register any unregistered paths to that disk.
- The **-o deregister** operation automatically deregisters all paths for multi-path physical disks. The new **-p option** can be used to override this behavior and force only the specified path to be deregistered.
- The **-o list** operation lists only one path per physical disk by default. This is useful for configurations that have many physical disks with multiple paths. The new **-A** option may be used to override this behavior and force a list of all paths. Additionally, the disk short name is listed for physical disks that have them (even on cluster systems) by default. The existing **-N** option can be used to override this behavior and force a list using the device specification for the physical disk.

The **-o list** operation now supports a new listing style via the new **-I** option. This operation provides a way to list any or all information about physical disks. The argument to the **-I** option specifies what information to list. Current arguments accepted by the **-I** option are: **basic, registered, mpio, paths, locales, size, or all.**

4.2.4.7 **admvdisk operations modified**

The following **admvdisk(1M)** operation has been modified for R4.20.

- The **-o list -l** (long listing) has been modified to list all registered physical disk paths and the related locale I/O direction information on systems with multiple locales.

4.2.4.8 **New date format in SCCS commands**

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 now be specified in either a two- or four-digit year format.

- `-c97/0228235959`
- `-c1997/0228235959`

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

See the **get(1)**, **prs(1)** and **sccsfile(4)** man pages for more details.

4.2.4.9 **BSD 4.4 telnet ported to the DG/UX System**

A new **telnet** command has been ported for the DG/UX operating system. BSD revision 4.4 of **telnet** is now available. This new command provides changed and new features that were not in the previous command. Please see the **telnet(1M)** man page for detailed information.

4.2.4.10 **New VDM ioctl commands**

There are several new VDM ioctl commands and several that have been modified with new versions. All VDM **ioctl** commands for the **vdmpio()** subdriver have been replaced with new VDM **ioctl** commands for the **vdmpphys()** subdriver. See the **vdm(7)** and **vdmpphys(7)** man pages for more details.

4.2.4.11 **gcc and cc command changes**

The default behavior of the **cc** and **gcc** commands for 4.20 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. There is a command line option to enable the use of Pentium Pro instructions.

- For the **gcc** command the option is: **-march=pentiumpro**
- For the **cc** command, the option is: **-W0,-march=pentiumpro**

4.2.4.12 New **adminterrupt(1M)** command

This is a new command that allows the system administrator to assign device interrupts to specific processors. This command is used for system performance tuning. See the **adminterrupt** man page for more information.

4.2.5 Fewer dots printed during disk registration on large systems

On large systems that have hundreds of disk paths, like the AV20000, 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, a '*' 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 indicator of '.' is used, printing it once per disk path registered.

4.3 Perl version 5

The DG/UX system now 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-specific libraries and classes.

See the man pages for more information.

5 Notes and warnings

This section lists the notes and warnings that are applicable to System Release R4.20.

5.1 Kernel builds take longer with AV20000

On AV20000 systems, configuring devices and building a kernel takes longer than it did on previous versions. Please be patient and allow your system several minutes longer.

5.2 Floppy driver not supported on AV 20000

The AViiON 20000 does not support a DG/UX driver for 3 1/4" diskette (floppy) media. You can, however, use the boot floppy provided with the AViiON 20000 to boot the machine because the boot floppy does not require a driver to work.

5.3 Line speeds in half-duplex modes

Half-duplex operation on all PCI sync controllers supports baud rates up to 19.2kb. Higher line speeds are not supported in half-duplex mode.

5.4 AV 6600 memory error reporting

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. So, when memory errors occur, this may result in the incorrect identification of the failing DIMM or SIMM.

If you experience memory errors and you are concerned that the wrong DIMM or SIMM is identified, then contact your local DG 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.

5.5 Undesirable /etc/devlinktab results

When the **admfailovervtc(1M)** command tries to sync the **devlinktab** files, it tries to determine an unused set of tty lines common to both hosts. In doing this, it may provide undesired results, especially if ttymons or printers are already configured.

To ensure your customizations are not lost, make backup copies of the **/etc/devlinktab** on both machines before running the **admfailovervtc -o sync** command. After running the sync, the backup copies can be restored and the **/usr/sbin/init.d/chk.devlink** program can be run to re-establish the original tty links.

5.6 SCSI device restriction

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.

5.7 NetWorker documentation and Autochanger licenses

The titles of the manuals in the NetWorker documentation set delivered with R4.20 indicate that they are for NetWorker 4.2. These documents are also applicable to the NetWorker 4.2.5 release.

A reference in *Managing the DG/UX System* lists incorrect NetWorker Autochanger licenses. Refer to *Release Notice: Legato NetWorker for AViiON Computers, NetWorker 4.2.5* for a correct description of these features.

5.8 Unattended reboot on AV 3600 computers

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 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 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, then you may lose any error messages reported by the BIOS or DG/UX bootstraps early in the booting process.

5.9 PCI buses

AV 4700, AV 4800, and AV 5800 systems have two 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). The DG Model 24025 PCI 10/100 Fast EtherLink adapter card is not fully operable memory-mapped, so it does not work on the second bus of these systems. For more information, refer to the **tcen(7)** man page.

5.10 dpen driver

The **dpen** driver 10/100 auto-detection mechanism may incorrectly determine that 10BASE-T is the highest performance mode of interoperation for the following connections: when the Model 7446 PCI LAN Controller/100 or Model 7458 or 7467 PCI Multi-Channel/100 Ethernet is connected to an IEEE 802.3u Auto-Negotiation capable 10/100 switch or hub.

You can set the adapter to 100BASE-TX mode through the media parameters of the **dpen** driver. See the **dpen(7)** man page for information on how to explicitly set device media parameters.

5.11 Packet retransmissions on Fast Ethernet

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 Multi-Channel Ethernet adapter.

To correct the problem, offload some of the load from the Model 7458 PCI Multi-Channel Ethernet adapter, such as also using one of the following single port PLC/100 Ethernet PCI adapters with the model 7458:

- Model 7446
- Model 7463

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

5.12 ATM adapter limitations and restrictions

During certain stress test conditions, Data General observes the potential for poor NFS performance and, in some cases, hung file transfers. This can occur because of retransmission timeouts for cell/packets that are discarded because of ATM congestion conditions. 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)` man page on how to change these mount parameters. The `nfsstat(1M)` command can be used to check on 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 adapters Models 7459 and 7460:

- SNMP does not support separate ATM MIBs.
- ATM Forum Available Bit Rate (ABR) Service is not supported.

- LAN Emulation Server (LES) and the Broadcast and Unknown frame Server (BUS) are unsupported. The switch or another server in the ATM network must provide this support.
- No API is provided.
- The adapters are not supported on the Data General PCI expansion chassis, Models 7505, 7506, and 7507.
- When running Classical IP, if the cable is unplugged for a short time (less than a minute), then 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:

From System:

```
LEC0-0: Conn to LECS failed
```

From System:

```
network out of order  
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.

5.13 ATI Mach32 graphics accelerator

The ATI Mach32 graphics accelerator chip does not switch correctly from graphics mode to text mode when the X server is terminated manually or by an application failure. Image and text information appears garbled on the system console, and you cannot interact with the kernel terminal emulator (which requires the accelerator to be in text mode).

The following procedures enable you to work around the problem.

5.13.1 Set up your system for automatic dumps

Set up your system to automatically dump if it halts because in graphics mode you are unable to read and respond to the dump questions. The `dg_sysctl` command sets the automatic dump state. You must have appropriate privileges to use this command.

If you plan to dump to tape, then enter a command similar to the following. (A writable tape must be in the drive, and only the kernel frames of memory are dumped.)

```
# dg_sysctl -r auto -d auto -f "st(npvc(pci()),E),4" -l kernel
```

If you plan to dump to disk, enter a command similar to the following, where `sys_dump` is a virtual disk to hold the dump:

```
# dg_sysctl -r auto -d auto -f
    "vdm_dump(sd(npvc(pci()),E),0),sys_dump" -l kernel
```

5.13.2 Avoid taking hot key dumps

Avoid taking S 1000 or hot key dumps. Such dumps require you to answer the dump questions, which, in graphics mode, are rendered unreadable.

5.13.3 Terminating the X server

If you must terminate the X server, use the normal X session exit methods (such as **logout** from the root menu). If this is not possible, then enter CTL-ALT-Backspace.

Avoid termination of the X server through use of the SIGKILL or SIGTERM signals. If you use these signals, then you must reboot the system to reset the graphics accelerator to text mode.

5.14 X server memory usage

Using a large number of images (such as in Web browsers, **xv**, and **xanim**) causes the X server to use an increasing amount of virtual memory with each image. This virtual memory increase can cause the system to run out of memory. The virtual memory usage is an aspect of **malloc** functionality, and is a tradeoff between performance and memory conservation. To avoid overconsumption of system memory, stop and restart your X server (log off and log back on again) every few days.

5.15 /dev file system space

On systems that support PCI terminal controllers, the **/dev** file system may run out of space if more than three controllers are configured. To solve this problem, set the **DEVFSSIZE** kernel configuration variable to 8192 blocks or higher to increase the size of the **/dev** file system.

5.16 sigwait changes

The **sigwait(2)** function 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 recompilation 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** routine. 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 revisions of DG/UX 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.

5.17 Denial of service attack

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 any packets at the firewall or Internet gateway that are received from outside the network and have a local source IP address. Filter the packets on a non-DG system.

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

```
ftp://cert.org/pub/cert_advisories/CA-95:01.IP.spoofing*  
ftp://cert.org/pub/cert_advisories/CA-96:01.UDP.service_denial
```

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 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-95.01* ↵  
# mget CA-96.01* ↵
```

5.18 S_ISLNK macro defined

Beginning in R4.11, 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 avoid definition of 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.

5.19 Changes made to internal data structures

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 are purely kernel internal and 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) which perform DMA.

Because of these changes, you need to recompile all device drivers written for DG/UX on Intel. For drivers that reference the internal fields of `io_buffer_vector_type`, some mechanical code changes will also be necessary.

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`.

5.20 AV/Alert

On AV20000 systems with fully-connected disks, AV/Alert shows a quantity of disks double what is actually attached in the System Hardware Configuration. Should an error occur on any disk, however, AV/Alert will report the error correctly, with correct disk and disk controller information. This will be fixed in R4.20MU01.

5.21 Using QIC tapes for multi-tape dumps

When you are using QIC tapes for multi-tape dumps, it is essential that each tape is of the same media type. Do not mix tapes of different cartridge types.

For example, if the first tape is a QIC-150, the second and subsequent tapes should also be QIC-150s.

6 Documentation

The DG/UX System Release 4.20 Documentation CD-ROM contains DG/UX System documentation and the WorldView Viewer software for the DG/UX and Microsoft Windows™ environments. With WorldView, you can view documentation on-line. The *Quick Start Guide for Viewing DG/UX Documentation* discusses how to use WorldView.

The following list describes the hard-copy manuals that are updated for System Release 4.20:

Part Number	Description
069-701140-00	Installing the DG/UX System
069-701135-00	Programming with TCP/IP on the DG/UX System
069-701035-05	Using the DG/UX System
069-701059-10	Porting and Developing Applications for the DG/UX System
069-701141-00	Managing the DG/UX System

069-701145-00	Preventing and Fixing Problems on the DG/UX System
069-701144-00	Managing Mass Storage Devices and DG/UX File Systems
069-701140-00	Installing the DG/UX System
069-701132-03	Managing a DG/UX Cluster
069-100496-03	Legato NetWorker User's Guide
069-100495-03	Legato NetWorker Administrator's Guide

6.1 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 bullit, 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.

6.2 New man pages

The following manual pages were added for this release:

Manual Page	Description
<code>perl5(1M)</code>	Perl (Practical Extraction and Report Language) overview
<code>perlapi(1M)</code>	Perl 5 application programming interface for C extensions
<code>perlbook(1M)</code>	Perl book information
<code>perlbot(1M)</code>	Bag'o Object Tricks For Perl5 (the BOT)
<code>perlcall(1M)</code>	Perl calling conventions from C
<code>perldata(1M)</code>	Perl data structures

perldebug(1M)	Perl debugging
perldiag(1M)	various Perl diagnostics
perlform(1M)	Perl formats
perlfunc(1M)	Perl builtin 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
admininterrupt(1M)	administer interrupts
dg_mem_get_affinity(3)	get memory placement characteristics
dac(7)	Device Access Controller (Intel only)
dg_get_label_region(2)	Return the MAC region of the given MAC label
adm_intro(1M)	Introduction to DG/UX system administrative commands

7 Software distribution

For installation on an Intel-based DG/UX System:

Model Number	Part Number	Description
P001A-D	068-600209-00	CD-ROM Intel DG/UX
P001A-D	090-601168-00	Intel Boot Floppy Diskette

8 Installation instructions

To install the R4.20 release, follow the instructions in the manual *Installing the DG/UX System*. The medium contains DG/UX System Release 4.20.

Install...	If you are running...
DG/UX System R4.20	DG/UX System R4.11MU04
	DG/UX System R4.11MU03
	DG/UX System R4.11MU02
	DG/UX System R4.11MU01
	DG/UX System R4.11
	no version of DG/UX

9 Preparing a software trouble report

Refer to the *Preventing and Fixing Problems on the DG/UX System* for instructions on how to file a software trouble report (STR).

Note: When you are using QIC tapes for multi-tape dumps (see Chapter 9), it is essential that each tape is of the same media type. Do not mix tapes of different cartridge types. For example, if the first tape is a QIC-150, the second and subsequent tapes should be QIC-150s also.

10 Index

3D Board

- Device driver 12
- LAN controller 12

A

- ACO, *see* Application Capture Option
- Address Bit Permuting (ABP) 28
- Application Capture Option 9
- AV/Alert 37

B

- Backing up your local file systems 9; *see also*
Legato Data Backup Utility

C

- C compiler 13
- Cache coherency 8
- Cache-Coherent Non-Uniform Memory
Access 7
- ccNUMA 7, 12, 14, 17, 21
 - and MCL 21
 - I/O mode 24
- Commands
 - admdevice 24
 - admfailovervtc 29
 - adminterrupt 27
 - admiopath 15, 18, 22
 - admipinterface 21
 - admimpl 22

Commands (*cont.*)

- admpdisk 15, 16, 17, 18, 19, 22, 23, 24
- admvdisk 18, 19, 24, 25
- cc 26
- dg_sysctl 33
- gcc 26
- get 25
- ioctl 26
- mount 31
- nfsstat 31
- prs 25
- sysadm 15, 17, 18, 19, 21, 24
- telnet 26

D

- Data General Enterprise Management
(DGEM) 22
- DG/UX CLARIION Manager 22
- Disk Failover Package 17
- Disk short names 24
- Dump support 13
- Dumping to a disk 13, 33
- Dumping to tape 33

F

- Far Memory Caches 8
- Four-digit year, specifying a 25

H

Halt dumps 20
Hung file transfers 31

I

I/O mode 24

K

Kernel builds 28

L

LAN controller 12
LAN I/O 21
Legato Data Backup Utility 9
Load balancing 21
Long disk names
 When to use 19

M

M3D system management software 22
MCL 21
Memory errors 28
mpio modes, *see* disk mpio modes
mpio options
 disk mpio modes 15, 16
 load-balancing modes 15, 18
Multi-connected LAN (MCL) 21
Multi-path disks
 Managing, *see* mpio options
Multi-path I/O model 15
Multi-path physical disk management 23, 24

Multi-path physical disks 24
 Administration of 18
 Automatic path deregistration 25
 Automatic path registration 25
 Disk registration 27
 Listing only one path 25
 Long listing 25
 Management 23
Multiple system management 22

N

Network connections dropped, Fast Ethernet
 31
NetWorker
 Documentation 29
 Legato Data Backup 9
NFS performance 31
Non-Uniform Memory Access (NUMA), *see*
 Cache-Coherent Non-Uniform Memory
 Access or ccNUMA

P

PCI buses 30
PCI terminal controllers 34
Pentium Pro processor 12, 13
 Improving performance 26
 Internal data structures changes to
 accommodate 36

Q

QLLogic differential SCSI controller 13

R

Remote system management 22

S

SCSI devices 13, 29

SDK Package 9

Security features

 Zero disk blocks 20

Security Option, B2 7

Security Option, C2 7

Short disk names 18, 24

sigwait 34

Synchronous controller 14

System performance tuning 27

T

TCP/IP

 and MCL 21

U

User Datagram Protocol (UDP) services 35

V

VDM, *see* Virtual Disk Manager

Virtual Disk Manager

 Disk block 20

 Enhancements 14

 Managing I/O 23

 Multi-path physical disk management 23

 New dynamic configuration variables 20

 New ioctl commands 26

Virtual disks

 Halt dumps stored on 20

 Increased name length 19

Virtual memory usage 34

W

WorldView Viewer 37

End of Release Notice

