

## Patch Level Notice:

# DG/UX for AViiON Systems Patch Level op-sys-x\_5.4.2.931203

January 1994

Part Number 017-600019-04

This Patch Level Notice applies to Models:

P001

Q001

### 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 at [DFARS] 252.227-7013 (October 1988).

DATA GENERAL CORPORATION  
4400 Computer Drive  
Westboro, Massachusetts 01580

Unpublished — all rights reserved under the copyright laws of the United States.

Copyright © Data General Corporation 1989, 1990, 1991, 1992, 1993.  
All Rights Reserved.

Licensed Material — Property of Data General Corporation.  
This software is made available solely pursuant to the terms of a DGC license agreement which governs its use.

---

DG/UX is a trademark of Data General Corporation.

AViiON is a registered trademark of Data General Corporation.

The X Window System is a trademark of Massachusetts Institute of Technology.

OSF/Motif is a trademark of Open System Foundation.

Ethernet is a registered trademark of Xerox, Inc.

SunOS is a trademark of Sun Microsystems, Inc.

ONC/NFS is a registered trademark of Sun Microsystems, Inc.

UNIX is a registered trademark of UNIX Systems Laboratories

## 1. Introduction

This Patch Level Notice describes Patch Level op-sys-x\_5.4.2.931203 for 5.4.2 DG/UX™ Systems running on Data General's family of AViiON™ computers. In addition, this notice also includes information not currently available in the product manuals (e.g., information developed after the current manuals were printed, or corrections to current manuals).

This printed notice always accompanies the software. You may print additional copies of this notice after you have installed the product. A copy suitable for line printers can be found in the file /usr/release/op-sys-x\_5.4.2.931203.pn. In the event of differences between the printed copy of the notice and the copy on the distribution medium, the printed copy takes precedence.

This Patch Level consists of the following parts:

Part Description	Part Number
DG/UX for AViiON Systems Patch Level op-sys-x_5.4.2.931203 Notice	017-600019-04
DG/UX for AViiON Systems Patch Level op-sys-x_5.4.2.931203 Media	079-600366-00

## 2. Product Description

### 2.1 Individual Patches

For DG/UX Revision 5.4.2, Data General releases patches named and numbered on a per product basis. The products included in this convention are, but are not limited to, dgux, tcpip, nfs, X11, X11.lg, aview, and gcc. As needed, Data General will produce and deliver individual patches in response to Software Trouble Reports. Using the product DG/UX Revision 5.4.2 the naming convention is:

dgux\_5.4.2.pmm

where mm refers to the individual patch number for the product. Per product patch numbers are unique and always increase numerically.

During the support period for a product, it may be necessary to release multiple patches that replace the same file. In this case, a later patch will contain all fixes released in previous patches; therefore when you have two patches to the same file, always load the patch with the higher number. The patch with the lower patch number is superseded and will not appear in the next product patch level. In general, however, patches which do not replace previously modified files are separate entities (ie. not cumulative).

### 2.2 Product Patch Levels

When a sufficient number of individual patches has been generated, they are collected into a patch level on a per product basis. A product patch level is assembled as a single sysadm loadable package. For example, all the current applicable patches for DG/UX Revision 5.4.2, are assembled in one package using the following naming convention:

dgux\_5.4.2.nn

where nn is the patch level number, corresponding to the greatest individual patch number for dgux available in this patch level. Note that the "p" in the individual patch naming convention is dropped in this case. Product patch levels are cumulative, thus you should load only the product patch level file with the greatest number.

### 2.3 Patch Level op-sys-x

A collection of product patch levels are made available as an op-sys-x patch level for easier distribution. The op-sys-x patch levels are named using a date in the naming convention instead of a revision number. This should better ensure you have the latest patch level. The name is of the form:

op-sys-x\_r.r.r.yymmdd

where r.r.r is the release number and yy, mm, and dd correspond to the year, month and day, respectively. This op-sys-x Patch Level is named:

op-sys-x\_5.4.2.931203

Op-sys-x patch levels are cumulative, thus you should load only the most recent op-sys-x patch level. This op-sys-x patch level is a sysadm loadable tape which contains one product patch level per sysadm loadable package. The packages on this op-sys-x\_5.4.2.931203 Patch Level are:

Package 1: dgux\_5.4.2.106  
 Package 2: tcpip\_5.4.2.14  
 Package 3: nfs\_5.4.2.03  
 Package 4: X11\_5.4.2.08

For each product you have loaded on your system, you are directed to load the corresponding product patch level (package). See the section titled "Installation Instructions" later in this document for specific loading instructions.

Every user is directed to load package number 1, the DG/UX product patch level. After this package has been loaded, the /etc/issue file reads as follows:

DG/UX Operating System Level 5.4.2.106

## 3. Environment

### 3.1 Hardware

Patch Level op-sys-x\_5.4.2.931203 of the DG/UX System will run on Data General AViiON series machines with revision E.2 or later of the MC88100 processor and a minimum of 12 MB of main memory, at least 322MB of disk storage, a system console (or graphics monitor for workstations), and a tape drive (for software distribution). On workstations, 16MB of main memory is recommended.

Patch Level op-sys-x\_5.4.2.931203 of the DG/UX System may also be run on a "diskless" AViiON workstation. This requires a console or workstation monitor, but does not require any disk or tape drive units.

### 3.2 Software

Your system needs to be running DG/UX 5.4.2 System software. Apply this patch ONLY on DG/UX Revision 5.4.2.

## 4. Patches

The subsection headings below correspond to the individual patch numbers that are included in this op-sys-x patch level. Where it appears that a patch has been omitted because the numbering is not consecutive, consider that patch to be superseded and the fix included in a succeeding patch.

This op-sys-x patch level supersedes the following previously released items against DG/UX Revision 5.4.2:

DG/UX for AViiON System Patch Level op-sys-x\_5.4.2.930514  
 DG/UX for AViiON System Patch Level op-sys-x\_5.4.2.930122  
 DG/UX for AViiON System Patch Level op-sys-x\_5.4.2.921016  
 Software Prerequisite dgux\_5.4.2.s03  
 Software Prerequisite dgux\_5.4.2.s04  
 Layered Software Prerequisite op-sys-x\_5.4.2.s05  
 AViiON Hardware Supplement for NVRam Board dgux\_5.4.2.s06  
 DG/UX for AViiON Systems Patch dgux\_5.4.2.p45

For those systems that have loaded individual DG/UX Revision 5.4.2 patches or extensions, follow this outline to determine if you need to reload any previously applied single patches.

If previously applied patches, patch levels, or extensions are included in this op-sys-x patch level, then they will be properly overwritten.

If previously applied patches are not included in this op-sys-x patch level AND have an individual patch number greater than the patch level number, then they must be reloaded. For example if you have installed the patch dgux\_5.4.2.p110, then it must be reloaded, because its patch number, p110 is greater than the patch level number 106, dgux\_5.4.2.106, for this patch level.

If previously applied patches are not included in this op-sys-x patch level AND have a patch number less than or equal to the patch level number, then these patches have been either superseded by another patch in this op-sys-x patch level or obsoleted and do NOT need to be reloaded. For example, the patch dgux\_5.4.2.p01 is not listed below as being included in this op-sys-x patch level and its number, 01 is less than or equal to the patch level number 106, dgux\_5.4.2.106, therefore it does not need to be reloaded.

If previously applied patches or extensions are not included in this op-sys-x patch level and appear to conflict with this op-sys-x patch level, call the Customer Support Center for assistance.

#### **4.1 Package 1 - Product Patch Level dgux\_5.4.2.106**

DG/UX patch level dgux\_5.4.2.106 contains the previously released DG/UX patch levels dgux\_5.4.2.25, dgux\_5.4.2.47 and dgux\_5.4.2.84 in their entirety. Both the old and the new contents are described below.

##### **4.1.1 dgux\_5.4.2.p08**

Patch dgux\_5.4.2.p08 modifies shutdown(1M) so that the /sbin/rc.init script runs only one time when the system run level is changed.

##### **4.1.2 dgux\_5.4.2.p09**

Patch dgux\_5.4.2.p09 corrects some problems with commands used by admterminal(1M). The result of this is that multiple copies of admterminal(1M) may run concurrently. Previously, when multiple copies of admterminal(1M) were run, users would see the message "cannot contact SAC." Additionally, the \_pmtab associated with the controlling tty was corrupted or deleted and the terminals being modified were left in an indeterminate state.

**4.1.3 dgux\_5.4.2.p17**

Patch dgux\_5.4.2.p17 corrects a problem with the handling of inheritance of debugging events after a vfork() call. When a process calls vfork(), it creates a child process that shares its address space. If the parent process is being debugged, and the child does not inherit debugging, the kernel removes all watchpoints from the child's address space, which, in this case, is also the parent's address space. This affects the parent's debugging. This patch will cause the kernel to leave the watchpoints in the shared address space after a vfork() call, even if the child does not inherit debugging. The result of this is that the watchpoints will remain in the child's address space, and if a watchpoint is tripped by the child, the child will receive a SIGTRAP and die. This is consistent with the way breakpoint inheritance is handled.

**4.1.4 dgux\_5.4.2.p19**

Patch dgux\_5.4.2.p19 corrects a problem in the bufcall processing in the FDDI driver. In the case that the system ran out of streams memory, the driver would never replenish the board's receive buffers, as well as corrupting the list of allocated buffers. This would lead to a network hang, followed by a system panic when the network was restarted.

**4.1.5 dgux\_5.4.2.p22**

Patch dgux\_5.4.2.p22 corrects three problems in libc.

This patch corrects a problem in the ELF environment with INGRES 6.4 and yellow pages. When the INGRES user name is defined only via yellow pages, the INGRES servers dump core on a call to getpwuid(). This was caused by routines in libc.so.1 calling DG's version of malloc instead of the INGRES malloc.

This patch also corrects a problem in the COFF and ELF environments where the user is unable to use the TZ variable to change when daylight savings time begins and ends.

Finally, this patch corrects a problem in the COFF and ELF environments where bcmp(3C) compares one more byte than you tell it to compare. The application could crash because bcmp(3C) attempted to read memory that had not been allocated.

If your application was linked static, then you must relink with this updated libc.a. If your application was linked shared, then it will automatically pick up the update once the patch is installed.

**4.1.6 dgux\_5.4.2.p23**

Patch dgux\_5.4.2.p23 prevents a bus error from being generated when using the SCCS get(1) command in a SCCS scriptfile. This patch also corrects a similar problem with the SCCS commands cdc(1), comb(1), delta(1), and rmdel(1).

**4.1.7 dgux\_5.4.2.p24**

This patch modifies login(1) so that the delay after an incorrect username/password pair is incremental from 0 to 4 seconds. The first time a user enters an invalid username/password pair, the user will wait 0 (zero) seconds before being prompted again. On a second incorrect login attempt, a wait of 1 second is enforced, and so on for a maximum of 4 seconds.

Logic has also been added to login to implement the -d switch. When login runs with -d, and is passed the login device name as an argument, login will not make a ttyname() call to get the terminal's name. Using this enhancement on systems with many device nodes in /dev (and its subdirectories), will enhance login performance. In extreme cases, system performance is also enhanced.

To take advantage of the -d switch, on directly connected lines controlled by ttymon(1M), the login command for an individual line should include the -d switch followed by the device pathname on which login(1) is started. For

example, for tty08, the service started should be `"/usr/bin/login -d %d."`

The `admterminal(1M)` command can be used to modify the service a terminal. For example:

```
# admterminal -o modify -p ttymon1 \  
-s "/usr/bin/login -d %d" tty03
```

modifies the service started on tty03 to `"/usr/bin/login -d %d."`

Alternatively, you may use `sysadm`'s Device -> Port -> Port Services -> Modify option to change the service.

If `login` determines that the `-d` argument is not the correct login device, it will ignore the `-d` switch and use the `ttyname()` routine to determine the correct device.

#### 4.1.8 dgux\_5.4.2.p25

Patch `dgux_5.4.2.p25` corrects the following problems:

1. `Dtl` no longer aborts when executed with an object file that does not contain debugging information. It issues a warning message instead.
2. `Dtl` now issues warnings for several cases that previously resulted in errors.
3. `Dtl` takes up less disk space and executes faster than previous versions.
4. The `-compress` option documented in the `legend(5)` man page works. Previously it had no effect.
5. Support for the GCC "long long" type variants has been added.
6. Use of the GCC "-g1" option no longer causes a `dtl` abort but no debugging information is output either. A warning message is produced in this case. Support for this option is planned for a later revision of `dtl`.
7. Some versions of `dtl` contained a bug that would cause the ELF debuggers to get the following error when attempting global name resolution:

```
Error: Expected to find global variable "{name}" in legend "{legend-name}", but couldn't
```

When this problem occurs the debugging information should be regenerated with this version of `dtl` to repair the problem.

#### 4.1.9 dgux\_5.4.2.p26

Patch `dgux_5.4.2.p26` releases an updated `make` command which removes the 14 character filename restriction previously imposed. Previously `make` would incorrectly update the target when its dependents had not been changed if either the target or dependents had a 14 or greater character filename.

#### 4.1.10 dgux\_5.4.2.p30

Patch `dgux_5.4.2.p30` corrects a problem with the commands `mail(1)` and `rmail(1)`. If the intended recipient did not exist on the destination machine, the `mail(1)` and `rmail(1)` processes would continually try to redeliver the mail and thus become `cpu` bound. This sometimes occurs when `sendmail.cf` is not available due to not loading the `TCP/IP` package.

According to "Managing the DG/UX System", a login name can be up to 32 alphabetic or numeric characters, but if you want to be compatible with traditional systems, the login name should be no more than eight characters long. `mail(1)` is one of those commands that conforms to the traditional system specification. This patch removes the restriction of the 8 character login name (14 when ".lock" is appended to the login name to create the lock file for reading and writing the mailbox). However, the 14 character limit still exists over `NFS` (which is also based on

traditional systems). Thus, if the mail reader is not local to the system that mail is being executed on, the restriction of 14 characters still remains.

#### 4.1.11 dgux\_5.4.2.p32

Patch dgux\_5.4.2.p32 addresses a problem where the group ID of a file accessed over NFS will be set to 65535 when only the UID was to be changed. If a user that issues a chown command or chown(2) system call on a file over NFS, the NFS client software on the AViiON sends a short int for the group ID. In this case of a chown, the group ID is set to -1 if only the owner is going to be changed. Some NFS server implementations read the short int -1 and set the group ID to 65535. The server was expecting the -1 to be a long int. This patch changes the group ID in this case to a long -1 before sending the data packet. This problem does not occur between AViiON NFS clients and servers, only some foreign NFS servers.

#### 4.1.12 dgux\_5.4.2.p34

Patch dgux\_5.4.2.p34 corrects a problem whereby a DLPI application accessing the ethernet interface on a AV/530 or AV/4600 system may not start due to a hardware failure, and cause other protocol stacks that are accessing the ethernet interface to hang.

If a user application uses multi-cast addresses, the opening and closing of the dgen ethernet interface requires that the interface be restarted. If the restart fails, the ethernet device is stopped and data traffic, inbound and outbound, stops. The user application will be returned an error 14 from the DLPI.

This patch allows the failure to be detected and a restart message to be queued to the device and a subsequent restart will be tried.

#### 4.1.13 dgux\_5.4.2.p36

Patch dgux\_5.4.2.p36 corrects the ps(1) command to return 0 on successful return rather than 164.

#### 4.1.14 dgux\_5.4.2.p39

Patch dgux\_5.4.2.p39 corrects a problem where the file placed in /usr/spool/cron/atjobs, as a result of "at -f <script>", contains an extra blank line between each command line.

#### 4.1.15 dgux\_5.4.2.p43

Patch dgux\_5.4.2.p43 corrects a problem where core files may be incomplete when the program that is dumping core uses mmap() to map a file into its address space. The problem will only occur if the file is smaller than the length specified in the mmap call.

In addition, this patch corrects a few problems in the standard DG/UX product which are manifested in the Trusted DG/UX product although they have some implications in the standard DG/UX operating system. In particular, this patch corrects mkdir when applied to an NFS mounted file system to use umask. Also, it allows a system to NFS mount a remote file system on a remote file system mount point. Additionally, this patch corrects a problem with the creation of core files that allow the appropriate permissions to be applied. This fix mostly has implications in the Trusted DG/UX product.

#### 4.1.16 dgux\_5.4.2.p47

Patch dgux\_5.4.2.p47 corrects a problem with the wanadm and syncadm commands. The wanadm command was modified to allow large max frame sizes to be passed to the ssid driver. Previously, the request was NACKED (Negative Acknowledged) by the ssid driver resulting in the subnet not initializing. The syncadm command was modified to allow the /usr filesystem to be mounted Read-Only on the server. Previously, the syncadm command

opened the controller code Read/Write resulting in an error when trying to download the controller.

#### 4.1.17 dgux\_5.4.2.p51

Patch dgux\_5.4.2.p51 corrects a problem with the iscd (Integrated Synchronous Controller) on the AV/4600 series computers. The crc (Cyclic Redundancy Check) routine would fail to detect a zero length message and process it incorrectly. The driver would dereference the default NULL pointer resulting in a panic. The same packet would be processed correctly by the VSC/3 device.

This patch also corrects a problem in the iscd device driver on the AV/53X and AV/46XX series machines. Under heavy stress, the iscd may generate a memory scribble if it is receiving large frames that have a shared flag set. This will result in various DG/UX panics being generated.

#### 4.1.18 dgux\_5.4.2.p53

Patch dgux\_5.4.2.p53 corrects a problem with upsd(1M). When the power supply monitor is attached to a serial port on a model 20 cluster, upsd(1M) does not initialize the port correctly. A symptom of this problem is the repeated messages "Warning line power dead, on battery power" and "Line power restored; no longer running on battery" when there is no loss of power. On some ports these messages appear on each poll of the port. Other ports may not display a message on each poll, but the messages appear frequently.

#### 4.1.19 dgux\_5.4.2.p54

Patch dgux\_5.4.2.p54 corrects a problem with the sys/types.h header file not defining certain data types in the \_SYSV4\_SOURCE environment. The types.h header file has been modified so that those data types are visible when the \_SYSV4\_SOURCE environment is used.

This patch also corrects a problem with the \_INLINE definition in the include file, /usr/include/sys/\_int\_features.h. Currently, C++ programs that include /usr/include/sys/signal.h will not compile because \_INLINE is defined incorrectly. This patch modifies the \_INLINE definition in /usr/include/sys/\_int\_features.h so that it will be set to "inline" (a C++ keyword) when \_\_cplusplus is defined.

#### 4.1.20 dgux\_5.4.2.p58

For all disk devices found in /dev/pdsk, the logical disk manager (LDM) checks the physical device at boot time to see if it is LDM formatted. If it is, then the device is registered. Patch dgux\_5.4.2.p58 corrects a problem whereby if the device is not LDM formatted, then an inadvertent update is written to block 1 of the disk device. This may cause corruption if any data is stored in that block.

Situations in which this problem will be seen include:

1. A writable DOS formatted floppy is in a floppy drive and accessible when the system is being booted. The floppy will not be able to be mounted following the inadvertent update.
2. A NVRAM board on the system is being used in its raw form (not LDM formatted into partitions.)

This problem occurred only in the 5.4.2 release of DG/UX.

#### 4.1.21 dgux\_5.4.2.p61

Patch dgux\_5.4.2.p61 corrects a few problems in the standard DG/UX product which are manifested in the Trusted DG/UX product although they have some implications in the standard DG/UX operating system. In particular, this patch corrects mkdir when applied to an NFS mounted file system to use umask. Also, it allows a system to NFS mount a remote file system on a remote file system mount point. Additionally, this patch corrects a problem with



the creation of core files that allow the appropriate permissions to be applied. This fix mostly has implications in the Trusted DG/UX product.

#### 4.1.22 dgux\_5.4.2.p63

Patch dgux\_5.4.1.p63 fixes a number of problems with ttymon(1M). Several problems with the use of modems have been fixed. In particular, ttymon(1M) now correctly handles modem hangup. Previously ttymon could get hung in a portion of code trying to enforce a delay after modem hangup. The result of this was that other lines controlled by that parent ttymon could no longer be used for login. This patch also corrects a problem where the text "Message not found!!!" was placed in /var/saf/<prmtag>/log instead of a message indicating that the terminal controller was still downloading. Finally in the case of a terminal controller which is in the process of downloading code to a cluster, ttymon(1M) now correctly identifies this state and the text indicating that a terminal controller was downloading a cluster, and that ttymon would "Try again later in 15 seconds" will only appear when this is actually the case.

This patch also fixes a problem where it appeared that two different users were "live" on the same tty line. This was a result of ttymon making an incorrect entry in /etc/utmp after receiving a signal indicating that a child process on a line had died, i.e. the user logged out. Upon receiving this signal, it was possible for ttymon to mark the wrong entry in /etc/utmp as dead. This happened more frequently on systems which were up for a long period of time, or systems where login/logouts were frequent. In short, this happened on systems where pids wrapped. With the application of this patch, it is no longer possible for the ttymon parent process to mark the wrong entry as dead in /etc/utmp.

#### 4.1.23 dgux\_5.4.2.p65

Patch dgux\_5.4.2.p65 corrects a problem whereby attempting simultaneous multiple binds of hardware mirrors on a HADA-I causes a field to not be initialized. This caused all but the first mirror to be built with an incorrect size.

#### 4.1.24 dgux\_5.4.2.p66

Patch dgux\_5.4.2.p66 corrects a problem whereby the registration of any disk on a system will cause all open software mirrors on the system to be resynced.

#### 4.1.25 dgux\_5.4.2.p71

Patch dgux\_5.4.2.p71 corrects a problem with ksh(1) arrays when assigned with the "set -A <array> <value>" syntax. If a one element array is assigned as described, then "print \${#<array>[\*]}" returns the number of characters in the array value instead of the number of values in the array. For example, if the array "foo" is set to the one element value of "bar", "print \${#foo[\*]}" returns 3 instead of 1.

#### 4.1.26 dgux\_5.4.2.p72

Patch dgux\_5.4.2.p72 corrects a problem which causes a 2000110 panic to occur when a zero length message is sent to a UNIX domain socket. This has been corrected by not allocating or freeing a memory buffer for this message if the length of the message is zero.

This patch also addresses a problem in the UNIX Domain socket subsystem whereby a bind(2) call would fail if the size of the argument passed was the size of the socket structure. Without this patch, the bind call must be made with the size being the sum of the size of the protocol family plus the string length of the path. This patch changes the bind(2) system call to accept the size of the socket structure for the namelen argument.

#### 4.1.27 dgux\_5.4.2.p73

Patch dgux\_5.4.2.p73 fixes a problem with the taccess(1) and twrite(1) commands used by Reexchange. If the /etc/passwd file is close to or larger than approximately 3000 bytes, a portion of Reexchange's memory is

corrupted by `taccess(1)`. When `twrite(1)` accesses this portion of memory it displays the message "Error: user label nnnnnnn - unknown Label Identifier" where nnnnnnn is a large number. It then dumps core. The `taccess(1)` command is modified so that memory is not corrupted, causing `twrite(1)` to dump core. The `twrite(1)` command is modified to report Label Identifier errors correctly.

This patch also corrects a problem where the Reelexchange utilities in DG/UX 5.4 were filling out the VSN to 6 characters with leading zeros. With the installation of this patch, it is now possible to read tapes whose labels have fewer than 6 characters.

#### 4.1.28 dgux\_5.4.2.p77

This patch corrects an incompatibility between patch dgux\_5.4.2.p45 and any of the following AV/52xx or AV/62xx system boards:

Part number 005-36290 Revision 56 or greater  
 Part number 005-36291 Revision 55 or greater  
 Part number 005-36292 Revision 55 or greater

This incompatibility causes the system to hang while booting a kernel which includes patch dgux\_5.4.2.p45 on a machine which is configured with one of these system boards.

In addition to correcting this incompatibility, this patch includes the code fixes introduced in patch dgux\_5.4.2.p45. These code changes correct a problem that may, in the unlikely event of a multi-bit hardware memory failure, result in improper program execution or data corruption on AViiON 5200, 5225, 6200, 6225, 6200-20, 6225-20, 5240, 6240, 6240-20, 6280, 7000, and 8000 systems. Without this patch, the systems listed above will compensate for single bit memory errors. With this patch, the systems will continue to compensate for single bit errors, and in addition, DG/UX will:

1. Use `syslogd` to report and log memory errors. The default configuration for `syslogd` enables the reporting of memory errors to the system console.
2. Write ECC corrected data back into main memory. This helps prevent single bit errors from degrading into multi-bit errors.
3. Use improved processor synchronization to avoid deadlock in multiprocessor systems when writing ECC corrected data back into the main memory.
4. Panic when the system encounters a multi-bit hardware memory failure. The panic prevents data corruption and minimizes data loss that may result from this hardware memory failure.

This patch also corrects a hardware irregularity that can cause corruption of a system control register during data access exception processing. This patch contains a kernel software workaround for this problem. The symptoms of this problem are varied and depend upon the state of the process at the time the problem occurs. Possible symptoms are improper program execution or data corruption.

In addition, this patch corrects two HADA-I disk array deconfiguration problems. The first problem may be seen as a panic code 2000075 or 4000467 during an unbind or while running `probedev`. The second problem corrected is exhibited as a panic code 3000032. This problem occurs when the `dgsb` (HADA-I) driver is deconfigured as the result of unbinding all disk units. The driver code was not properly checking to confirm that no processes still held the `dgsb` open before returning memory to the free list. A process which still had the `dgsb` open and tried to access this memory was likely to find that memory already in use by another process, thereby causing the panic to occur. This problem was corrected by modifying the `dgsb` driver code to properly check for processes holding the driver open before freeing the memory.

Finally, this patch also includes three corrections to the `dev.a` kernel library.

The first correction included is the elimination of an error message which gets returned from momkfs when trying to create system areas on a magneto-optical disk at the same time that sar -d is running. The error message returned is "Cannot create the System Bootstrap Area". This error message was returned because sar and diskman both tried to open the disk device simultaneously. This was corrected by allowing diskman to open the magneto-optical device even if sar already has it open.

The second correction included is a workaround for a hardware problem with the Sony multifunction optical disk drives. The problem occurred when a hardware error was detected while writing several sectors to a WORM platter. Attempted retries to write data that was already written caused the logical disk to seal. The workaround will prevent retry attempts and will now leave error recovery up to the OD (Optical Disk) driver. The OD driver will only attempt to write those sectors that were not written by the SCSI adapter, thus avoiding this problem.

The third correction included corrects a problem which causes DG/UX to panic (most commonly with panic code 2000075) if the /dev/lp device is not configured with any streams modules and the /dev/lp device is accessed. The default configuration for the /dev/lp (set via lp\_ARG in /etc/dgux.params) device pushes the three streams modules ptem, ldterm, and ttcompat onto the device. The panic will only occur when this configuration has been modified so that the /dev/lp device is configured without any streams modules.

#### 4.1.28.1 Soft Error Reporting

A single bit soft memory error indicates that a single bit in a memory word is not correct, but that the system has been able to compensate by reconstructing the correct memory data. There are many environmental factors that contribute to an occasional soft, single bit error. Single bit soft memory errors are expected on most modern computer systems.

Although soft memory error logging has previously been enabled on some AViiON models, patches dgux\_5.4.2.p45 (now obsoleted by dgux\_5.4.2.p77) and dgux\_5.4.1.p76 (now obsoleted by dgux\_5.4.1.p87) will enable the logging for AViiON models 5200, 5225, 6200, 6225, 6200-20, 6225-20, 5240, 6240, 6240-20, 6280, 7000, and 8000. These log reports will be stored in the system's error log file (/var/adm/messages) and, by default, will be displayed on the system console. Although some other vendors may not display soft memory errors, Data General feels that reporting these errors may provide customers with some useful information.

The AViiON servers in the list above all have an error correcting (ECC) feature that automatically compensates for single bit memory errors. Therefore, most soft memory error reports should cause no concern and should only serve as a reminder that your AViiON hardware is operating properly. However, if soft memory error reports become very frequent, then the reports may indicate that field service should be contacted. Accordingly, we ask that you contact field service only if the frequency of soft memory errors exceeds the following limits:

Board Size	Soft Errors per Month
192 MB	40
128 MB	27
64 MB	13
48 MB	40
32 MB	27
16 MB	13

#### 4.1.28.2 Additional Corrections:

1. Corrects a 2000075 panic which occurs on a dual-initiator configuration. The panic is caused by a timing problem which occurs when one initiator selects a device for I/O while the other initiator is in the process of configuring.

2. Eliminates parity errors which occur while data is being transferred out to a HADA-II adapter. This patch corrects the parity errors by resetting the SCSI bus and retrying all outstanding requests including the request that failed due to the parity errors.
3. Eliminates a 53000023 panic which occurs on dual-initiator configurations when an entry is inadvertently deleted from the disconnected request queue for `ncsc()` or `dgsc()` devices. This problem can occur during one of two scenarios: (1) when one initiator is booting, in which case the other initiator may experience the panic; (2) when a SCSI bus reset occurs, in which case the host experiencing the reset may panic.
4. Corrects a problem with the HADA-I driver. During a `gridman(1m)` bind operation a global system lock is held. If a `ps -f` command is issued for the `gridman` process during this time, `ps` will hang and all further attempts to fork or exit will also hang until the bind has completed. The problem is corrected by changing the implementation of the bind operation so that the global lock is not necessary.
5. Corrects a problem with NFS-mounted Opstar Magneto Optical file systems whereby the OD driver requires each buffer descriptor to be a particular size; NFS requests which do not meet this requirement fail. To correct this, the driver now checks the entire request size rather than that of the individual descriptors, which eliminates the NFS problem.
6. Changes the `nvrld` device driver to enable self-ID checking of file system metadata. This enables the file system to detect corrupt data and restrict access to a file system to avoid a system panic.
7. Corrects a problem with the SCSI disk driver handling of Model #6796 disk drives. These drives support command queueing whereby disk requests can be queued on the drive instead of in the driver. The `sd` driver is queueing more commands concurrently to the disk than the disk supports. The disk reports an error when an attempt is made to enqueue a command when the disk queue is full. If this error is returned when the system is trying to page something to/from the disk, panics 57000073 or 57000116 may occur. This patch corrects the problem by limiting the number of commands queued to the disk drive to the maximum supported by the disk.
8. Eliminates a race condition which can cause system failure. Use of the DIOSTART asynchronous disk I/O facility can provoke this problem. The panic code associated with the failure may vary from one occurrence to the next. While system failure due to this problem is rare, it may occur on any system; therefore, it is recommended that this patch be applied to all DG/UX 5.4.2 systems.
9. Prevents panic 04000110 from occurring on AV/43xx systems as a result of stray interrupts at the `duart` port before the `duart` is configured. This panic occurs while the system is booting, and the following error message is displayed: "Hardware failure. Too many spurious interrupts from integrated device 0x4. System Panic. Panic code 04000110." This patch masks off all `duart` interrupts at the chip level if a stray interrupt is detected. These interrupts are then unmasked by the `duart` driver during configuration of the `duart`.
10. Prevents panic 04000110 from occurring as a result of stray `ncsc()` interrupts. The panic occurs as a result of the stray interrupt handler not correctly clearing the interrupt mask and thus not clearing all the pending stray `ncsc()` interrupts. This problem has only been seen in dual-SCSI-initiator configurations when an integrated `ncsc()` adapter is used. This patch fixes the problem by correctly clearing the interrupt mask.
11. Corrects a problem in the SCSI disk driver. The current driver checks to see if the drive properly responds to the driver's attempt to set it into SCSI-2 mode. If it does, then it assumes the drive does command queueing and sets up to use command tags. Some drives respond to the SCSI-2 request but do not handle command tagging. This behavior is technically allowed under the SCSI-2 spec. This patch will cause the driver to check the command queueing bit in the identify data and only set up to use command tags if a drive is able to handle them.
12. Corrects a hardware irregularity which can cause corruption of a system control register during data access exception processing. This patch contains a kernel software workaround for this problem. The symptoms of this problem are varied and depend upon the state of the process at the time the problem occurs. Possible symptoms are improper program execution or data corruption.
13. Addresses a problem where some tape drives will not write the filemarks on the first try. This patch will force the retry of the write filemark command with the vendor specific bits cleared after an initial failure. HP 4mm

DAT drives were addressed with this patch.

14. Corrects a HADA-I disk array deconfiguration problem. This problem may be seen as a panic code 2000075 or 4000467 during an unbind or a probedev.
15. Corrects a problem with ucontext information which could result in the misalignment handler incorrectly trapping with a SIGSEGV. When a signal is delivered to a program with ELF-style siginfo, the value for the uc\_mcontext.gregs[R\_R0] portion of the struct ucontext should be zero. In DG/UX 5.4.2, this is not always true unless this patch is applied. If your application expects this value to be zero when handling signals, you should apply this patch. Also, if your application uses the misalignment handling facility, you should apply this patch.
16. Corrects an occasional problem with booting machines dual-ported to a HADA-I. The problem was corrected by enlarging a timeout value for the HADA-I (dgsb) device driver.
17. Works around a problem with the AV/5240, AV/6240, AV/7000 and AV/8000 parallel printer port hardware which causes the lp device driver to wait for the hardware to deliver a interrupt with the FIFO flag set to empty, which the hardware fails to deliver. This patch causes the lp driver to timeout every 5 seconds to check the state of the FIFO flag.
18. Enhances the line printer device driver to prevent printer hangs. A timeout has been added to prevent the driver from waiting indefinitely for an interrupt that indicates the parallel port has drained its FIFO. If the driver times out, it will check to see if the parallel port's FIFO can accept more data. If the port is not ready for data, the driver will continue to this process of setting the timeout, waiting for an interrupt, and checking the FIFO status.

#### 4.1.29 dgux\_5.4.2.p78

Patch dgux\_5.4.2.p78 provides new support for remote printers connected to BSD servers, the MiLAN 3100 Network Print Server DG Model 40263, and consolidates previous patches for postprint and termprinter support.

##### 4.1.29.1 Remote BSD Support

Remote printers served on BSD and Novell systems may now be accessed using lpNet (System V printing). Prior to this patch on DG/UX 5.4.2, remote printers to BSD and Novell systems were not supported. With this patch it is no longer necessary that you run both the lpd/lpr and lp schedulers to access your remote BSD servers.

To support printing to or from a remote BSD server, you must define that system as a BSD server. You may use the command:

```
lpsystem -t bsd system-name
```

or in sysadm, Device -> Printer -> Systems -> Add. See lpsystem(1M).

If you had been running both lp and lpd on the same system, it was necessary for you to delete the listener port task. This must be restored and the BSD lp daemon, lpd, terminated. The listenBSD task that is defined in /etc/saf/tcp/\_pmtab.proto as:

```
lpd::root:reserved:reserved:reserved:p: /var/spool/lp/fifos/listenBSD#Berkley lp service
```

must be added to /etc/saf/tcp/\_pmtab. Then you must stop and restart the tcp listen port monitor using sysadm.

The scheduler has been improved to allow faster scheduling. Under heavy load to a BSD server this improvement is over 10 times. As a result, jobs get sent to a remote or local server faster, and local commands are responded to more quickly. Before this change, under heavy load (50-100 jobs in a queue), jobs would remain long periods of time in the print queue waiting to be scheduled.

#### 4.1.29.2 Other lp Bug Fixes

The lpsched command has been modified so that /var/spool can be set up as a separate file system. In previous releases of DG/UX a /var/spool file system could not be unmounted due to a bug in lpsched.

A bug was fixed in lpsched, so that lpstat reports the correct destination for each request when printing to a printer class that is being modified.

The scheduler has been improved to fix the unexpected message errors, received when using lpstat or lp.

The scheduler is more robust due to a selective redesign which includes testing for memory use errors.

The scheduler no longer creates unnecessary files in the spool directory. This previously required an administrator to periodically clean out these files.

The scheduler names filtered files with the request id and the system name to provide unique names. Previously two filtered files could have the same name causing some files not to be printed.

The scheduler properly removes a cancelled request to a server that has been down. While the server is down the job is marked cancelled when doing an lpstat.

#### 4.1.29.3 MiLAN 3100 Network Print Server Support

The MiLAN Network Print Server DG Model 40263 connects remote printers to TCP/IP Ethernet networks. The FastPort server is connected directly to the network using 10BaseT, 10Base2, or 10Base5 Ethernet media and the device includes a high speed Centronics parallel port and an RS232 serial port. This patch provides two files for this support, fpfilter and fastport.

Limited sysadm support is provided in this patch for the MiLAN Network Print Server. Full and seamless sysadm support is provided in DG/UX 5.4 Release 2.01. We recommend for convenience of use, you upgrade to DG/UX 5.4 Release 2.01 for the MiLAN support. For those sites which are not in a position to upgrade, we provide these steps.

First refer to the manual that came with your MiLAN Print Server, FastPort 3000 and 3100 Network Print Server User's Guide Version 2.5 for hardware installation instructions. Use this patch notice for the software installation instructions. The software installation instructions in the FastPort User's Guide are for DG/UX 5.4 Release 2.01 Systems.

Using sysadm define your MiLAN printer as a local printer:

```
Printer name: <must be the same as the hostname for the
              MiLAN server>
Connection: local
Device: /dev/null
Interface: fastport
```

The default delivery in the patch provides support for the parallel port. For DG/UX Systems prior to DG/UX 5.4 Release 2.01 the printer queue name must match that of the hostname for the MiLAN Network Printer as defined in /etc/hosts. In the interface script, /etc/lp/interfaces/<printer-name>, you may also modify 'basename \$0' in the statement \${DG\_NETWORK\_HOST:= 'basename \$0' to any hostname desired. The default port is 2001 which is the parallel port. To select the serial port you must modify 2001 to 2000 in the statement \${PORT\_OPTION:=2001}. For additional information refer to the man pages, fastport(1) and fpfilter(1).

Data General recommends the "Use-ACK" setting for the parallel port when printing to impact (line or dotmatrix) printers. You should also manually set the IP address.

If you connect the FastPort print server to a multiport transceiver, and there are pauses of several minutes during printing, the LAN may be operating at utilization levels greater than the IEEE standard recommends. To eliminate the pauses, either (1) set the print server to ignore alignment errors using the E command (see section 11.3 of the MiLAN FastPort User's Manual); or (2) put th print server on a separate single port transceiver.

You should cycle the print server power (that is, turn it off, and then on again) if you change any of the following:

- \* Printers on the print server port
- \* Print server switch positions
- \* The setup of a printer (such as switching from SW flow control to HW flow control)

#### 4.1.29.4 Postprint(1)

This patch corrects a problem with the postprint(1) command. This command has been modified to accept `^D%!PS` (0x04%!PS) as a valid PostScript signature. When the postprint(1) filter encounters one of these PostScript signatures at the beginning of a file or piped file, the data will be written directly to stdout and NOT converted to PostScript.

```
^D%!PS
%!PS
%!PS-Adobe-
```

In a networked PC environment, MS-Windows was found to include a `^D` as a part of the PostScript signature. Since the print-ready file from MS-Windows is already in PostScript format, it is not necessary to filter the file using postprint(1). However postprint(1) did not recognize the `^D` and incorrectly filtered a PostScript print-ready file.

The `^D` is a end-of-file/end-of-job marker for PostScript files.

#### 4.1.29.5 TermServer and Datability Support

This patch corrects a problem with the termprinter interface script and the lp.termprinter program which starts a printer session with a terminal server. This patch as delivered supports the termserver model of a single ip address per printer port.

This patch notice also provides specific details for setting up the DG supported termserver:

Model	Description
4817	TermServer 2000 Series
4819	TermServer 2100 Series
40600-40611	Datability VCP300/1000 TCP Series

##### 4.1.29.5.1 lp.termprinter

This command was updated to support the following additional command arguments.

- n Suppress printing 8 KB of null characters at the end of each connection session to flush the data from the termserver buffer to the printer. The default is always print the null characters. Some termserver do not send the data in the termserver buffers to the printer when the connection from the host is closed. As a result these nulls are necessary to force the valid data from the termserver out to the printer. The

Datability VCP300/1000 does not require nulls. The TermServer 2000/2100 running software release 4.2 does not require nulls.

- l secs Uses the SO\_LINGER socket option when the session is established. The -l option takes an argument indicating in seconds the maximum of time to wait for the close. See socket(2) and /usr/include/sys/socket.h for additional information. By default, the termprinter interface script invokes lp.termprinter with -l 32000.
- p port Use the port number or port alias name when connecting to the TermServer with the telnet protocol. If no port option is given, then the default port number is the telnet port 23.
- f Translate 0xff bytes to 0xff0xff. This option is needed to correct a problem with printing graphics to a 2000/2100 TermServer printer.

#### 4.1.29.5.2 termprinter

This interface script has been restructured to send the banner page, the content file, and any interrupt/cancellation message all in a single data stream. Prior to this release, each were sent using individual sessions.

The termprinter interface script by default, includes the -l option for SO\_LINGER. It also by default prints nulls to flush TermServer buffer. If using release 4.2 of TermServer Software TCP/IP w/SNMP for the TermServer 2000/2100 series, you should also include the -n option to suppress writing the nulls. The line would be modified as follows:

```
TERMPRINTER="${TERMPRINTER} -h ${DG_NETWORK_HOST} ${PORT_OPTION} -l 32000 -f -n"
```

This interface script is located in /usr/lib/lp/model/termprinter. Each time a termprinter is created or modified, this script is copied into the directory /etc/lp/interfaces taking the name of the printer. Making the change to /usr/lib/lp/model/termprinter will save you modifying the script in /etc/lp/interface each time you use sysadm to change your termprinter definition.

#### 4.1.29.5.3 datability

This patch deliveries a new interface script named datability. It is a duplicate of the termprinter script with the following changes to the invocation line of the lp.termprinter program.

The 8 kb of nulls are not necessary for Datability, so the "-n" switch is used. Also the "-f" switch is removed because translating 0xFFs are not necessary for the Datability.

#### 4.1.29.5.4 Defining a TermServer Printer Under DG/UX 5.4

A printer under DG/UX 5.4 Systems, is defined as follows. You may use the sysadm interface or command line interface. For example purposes we have included the command line interface.

Place the printer name and its internet address in /etc/hosts or in the YP database if running Yellow Pages.

```
# sysadm addhost
at the following prompts:
Host Name? <printer_name>
Host Address? <printer_address>
```

Configure a local printer:



```

Connection: local
device: /dev/null
Interface: /usr/lib/lp/model/<termprinter or datability>

```

```
# lpadmin -p <printer_name> -mtermprinter -v/dev/null
```

If using the Datability termserver use the datability interface script:

```
# lpadmin -p <printer_name> -mdatability -v/dev/null
```

For ascii line printers:

```
Content type: simple
```

```
Printer type: printer
```

```
# lpadmin -p <printer_name> -Tprinter -Isimple
```

For PostScript Printers:

```
Content type: PS
```

```
Printer type: PS-b
```

```
# lpadmin -p <printer_name> -TPS-b -IPS
```

Next enable the printer and accept print requests:

```
# enable <printer_name>
```

```
# accept <printer_name>
```

When printing binary ( 8 bit) data, make sure the printer defined on DG/UX is using the stty setting of -opost. See stty(1) for additional information on -opost. The command to set this option follows:

```
# lpadmin -p <printer_name> -ostty=-opost
```

#### 4.1.29.5.5 Setup for Datability

In order to facilitate access to a specific port on a Datability termserver, you need to create a unique IP address for it. You also need to set the port characteristics appropriately. The following DATABILITY commands will create the address, associate that address with a specific port, and set the port's characteristics. You must have the Datability "Local>" prompt and have the super privilege turned on by issuing the command "SET PRIV". The default password is "SYSTEM".

The following example defines a printer port named P1 on termserver port 8, with an IP address of 192.41.217.209 and a speed of 9600 baud.

```
Local> CHANGE ADDRESS 192.41.217.209
```

```
Local> CHANGE PORT 8 AUTOBAUD DISABLE ACCESS REMOTE SPEED 9600
```

```
Local> CHANGE PORT 8 IO_FLUSH DISABLE
```

```
Local> CHANGE SERVICE P1 PORT 8 IP 192.41.217.209 TCP PORT 23
```

Local> CHANGE SERVICE P1 PORT 8 TELNET DISABLE LPR ENABLE

Do this for each printer port on the termserver. Note that the Change Address command just adds new addresses - it does not change the address of the unit. You can define up to 63 additional addresses in this way.

The following are recommended parameters for a printer connected to a Datability termserver. To display the port parameters, enter the command "SHOW PORT 8".

Character Size: 8  
 Flowcontrol: <printer's-protocol>  
 Parity: none  
 Stop Bits: 1  
 Access: Remote  
 Input Speed: <printer's-baud-rate>  
 Output Speed: <printer's-baud-rate>  
 Modem Control: Disabled  
 Xon Char: <printer dependent usually ^Q>  
 Xoff Char: <printer dependent usually ^S>  
 Type: Soft

Enabled Characteristics:  
 Input Flow Control, Line Mode, Loss Notification,  
 Output Flow Control, Verification

Once the printer hostname is defined, verify it is accessible via the local routing table on your DG/UX host. From your DG/UX host use the command: "ping <hostname>" to verify accessibility.

#### 4.1.29.5.6 Setup for TermServer 2000/2100

This text is repeated here from the TermServer Software TCP/IP w/SNMP Release Notice for Model 30152.

Use the following recommended parameters for a printer connected to a TermServer. Set other parameters as necessary. Set the Baud, FlowControlTo, and FlowControlFrom parameters to match the printer (use "setd <parameter> = ?" for permissible values). If the DTR line is used for flow control, make a special cable so the DTR signal is used as the CTS signal, and set the UseDTRin parameter to Ignore.

DeVice = Host, Glass, DeFault  
 Baud = <printer's-baud-rate>  
 DataBits = 8  
 DUplex = Full  
 LineProtocol = Asynchronous  
 PARity = None  
 StopBits = 1  
 UseDCDout = AlwaysAssert.NoToggle  
 UseDTRin = <printer-dependent>  
 FlowControlTo = <printer's-protocol>  
 FlowControlFrom = <printer's-protocol>  
 MOfde = Transparent

If using the Data General Model 664x and 677x series laser printers and software flow control, you should set the parameter, FlowControlFrom, to None. This is required when sending 8 bit data (PCL format) to these laser printers

in HP II mode.

#### 4.1.29.5.7 Additional Changes

In addition to specific support of the Datability Termserver, the following corrections have been included in this patch.

When requesting multiple copies for printing using the -n switch with the lp command, the correct number of copies will be printed.

The banner page is printed correctly when the printer is set for binary mode printing.

The following problems with the lp scheduler are also corrected:

1. If two print queues (A & B) are defined for the same device, and one queue (B) is disabled, all print jobs in the other queue (A) which were submitted after the job at the top of the disabled queue (B) will be suspended. To clear this, the scheduler must be cycled by issuing `sysadm stoplp` followed by `sysadm startlp`.
2. When specifying a printer class with `lp -d` the print request failed and the user would receive the error that a filter was unable to convert the given file.
3. When printing to a remote system V release 4 printer (lpNet), if multiple requests were made immediately after each other, the jobs would remain in the local print queue and not transmitted to the remote host. The scheduler would go into a state that indicated the remote printer is busy, when in fact it is not.
4. There are rare situations, when a user would receive an error message from the scheduler stating the user does not have access to specific file, when in fact the file has world read permission.

Using lpNet to print to a remote host when the local host, `host_1`, has multiple queues defined to use the same queue on a remote host, `host_2`, is a configuration that is not supported. Remote lpNet print queues are restricted to a one-to-one relationship. For example, `host_2` has a local printer named `lp_h2_1`. Next, `host_1` defines two remote printers:

```
lp: host_2 as printer lp_h2_1
lp2: host_2 as printer lp_h2_1
```

This configuration for remote printers is not supported, and thus restricted. Remote queues must have a one to one correspondence with a local queue for a local-remote host pair. It is a valid configuration for a third host, `host_3`, to define a queue to use the remote printer, `lp_h2_1`, on `host2`.

Note: it is a valid configuration to have multiple local queues define to use the same local printer. For example:

```
hp_binary: device=/dev/lp stty=-opost
hp_text: device=/dev/lp
```

The restriction is only with multiple remote queues from a single host printing to the same remote host and printer.

The problem will manifest itself by the lpNet log file growing in size because of the repeated "Connect" messages, however the print jobs are not printed and not removed from the local print queue. The lpNet log files from each host will indicate connects with the remote host, but the requests logfile will not indicate the requested file was printed.

**4.1.30 dgux\_5.4.2.p79**

Patch dgux\_5.4.2.p79 corrects a problem where sar is outputting "dgux restarts" instead of meaningful sar output.

**4.1.31 dgux\_5.4.2.p80**

Patch dgux\_5.4.2.p80 corrects 2000107 panics which occur when the system has run out of unwired memory. These panics most commonly occur in two situations: (1) when mounting a very large number of file systems on a DG/UX system; or (2) when remapping large sections of VME addresses on a system with a small amount of physical memory. This problem was corrected by doubling the default size of the unwired memory pool to two times the size of main memory. In addition, the UPOOL\_MIN configuration variable was added to allow even larger unwired memory pools to be created.

The UPOOL\_MIN configuration variable allows the administrator to specify a minimum unwired memory pool size, in megabytes. At boot time, this value is rounded up to the next 4 Mb. boundary and used to create the unwired memory pool. There are two constraints to this value:

- (1) The value cannot be smaller than the default unwired memory pool size of twice physical memory. Such requests will be silently ignored.
- (2) There is an upper limit on the size of the unwired memory pool. This limit is determined at boot time and is based on the kernel address space layout. If a value larger than this upper limit is specified, the value will be scaled down to the maximum supportable unwired memory pool size.

This patch also eliminates two race conditions in the virtual memory subsystem which can cause system failure. The panic codes associated with the failures may vary from one occurrence to the next. While the failures due to these problems are rare, they may occur on any system; therefore, it is recommended that this patch be applied to all DG/UX 5.4.2 systems.

In addition, this patch corrects a system hang condition which occurs when one process is hung waiting on I/O to a SCSI tape drive and the command "ps -ef" is executed. The hang occurs because ps holds the process existence lock and tries to get a write lock for the map entry of the hung I/O process. That write lock is never released (unless the system is rebooted) because the process holding it is hung. The rest of the system processes then hang pending the process existence lock held by ps. This patch corrects the problem by having ps obtain only a read lock on each process's map entry.

**4.1.32 dgux\_5.4.2.p81**

Patch dgux\_5.4.2.p81 corrects a problem in the faam.a kernel library which causes a panic code 3000011. This panic occurs because a routine in this library does not properly handle FIFO objects when clearing signal requests.

**4.1.33 dgux\_5.4.2.p83**

Patch dgux\_5.4.2.p83 is a new release of the firmware for the VLCi Ethernet LAN Controller (CMC-130). It corrects a problem where the firmware would not process a frame correctly if it was larger than the maximum size allowed on the ethernet. Although it may be legal to transfer a frame larger than 1514 bytes on a LAN, it will be discarded by TCP/IP. Sometimes a network failure may cause a frame to be corrupted/extended and this may also appear to be a ethernet frame larger than 1514 bytes to the controller.

Without this patch installed, the VLCi controller will appear to hang if a ethernet frame it received (via broadcast or direct addressing) and can only be recovered by rebooting the kernel. Stopping and starting TCP/IP will not recover the controller.

TCP/IP must be stopped and restarted for this patch to take affect.

**4.1.34 dgux\_5.4.2.p86**

Patch dgux\_5.4.2.p86 corrects LDU downgrades to readonly mode and possible data corruption which may occur as the result of a design flaw in the cdm.a kernel library. This patch corrects the problem by requiring a buffer lock to be held while testing the current state of the buffer. This will prevent a cache device buffer that is being modified from being cleaned by another thread of control.

This patch also resolves several issues with the use of NVRAM boards including enhancements associated with the Cache Device Manager (CDM) that relate to the use of a fast recovery file system on a NVRAM board and enhanced data protection by flushing all data on the front end device on shutdown.

This patch corrects a problem where configuring a cached device to use the pass through policy may result in a panic 2000075. This is caused by a reference to memory which has already been released.

Using a cached device with a back end device which contains a file system operating in fast recovery mode may result in system hangs due to incorrect management of asynchronous I/O operations by the cache device management code.

When using a fast recovery file system on a cache device there's a timing problem which can cause the I/O daemon to hang waiting on a lock, which thereby causes the entire system to hang.

This patch level also contains /usr/release/093701088\_2 which includes additional information for using cached devices.

**4.1.35 dgux\_5.4.2.p87**

Patch dgux\_5.4.2.p87 corrects a bug in the sc.a kernel library which causes the system to enter an infinite loop, thereby resulting in a soft system hang.

**4.1.36 dgux\_5.4.2.p90**

Patch dgux\_5.4.2.p90 corrects several problems in the failover operations.

First, this patch addresses some performance problems. Three operations provided by the -o option are noticeably slow when takeaway and/or giveaway databases are large. These operations are sync, give, and take. This patch to admfailoverdisk(1M) corrects a logic error and speeds up these operations.

There have been some reports where the processing of a takeaway database of approximately 200 entries has taken on the order of 20 minutes after a machine failover. These fixes to admfailoverdisk(1M) reduce the time to a matter of seconds, rather than minutes.

Further corrections have been made to the "admfailoverdisk -o take -h host" so that it is no longer 4 to 6 times slower than the analogous "admfailoverdisk -o give -h host" operation.

Second, this patch lifts an artificial limitation placed on the number of physical disks to be failed over. A symptom of this problem is that after a "give" operation to a remote machine, the giveaway and takeaway databases on the remote machine are not correct. The giveaway database contains only a subset of those disks which were failed over and the takeaway database contains the remainder. With this correction the giveaway database now contains all the disks failed over.

A third set of changes address a case where a host is set up to fail over some disks locally as well as support remote failover. These problems occur in a configuration where two hosts, LOCAL and REMOTE are dual ported to a CLARiiON. The host LOCAL normally uses all disks on the CLARiiON. LOCAL's failover databases are set up to do two items: (1) failover CLARiiON based disks to REMOTE in the case where LOCAL is down; and (2) failover

CLARiiON based disks to an alternate path from LOCAL in the case of an IOP failure on the CLARiiON or a SCSI failure on LOCAL. This patch corrects a problem where the giveaway database on LOCAL was incorrectly updated after disks were added to support (2) above, and a sync was performed. This patch also corrects a problem where the giveaway database on LOCAL was incorrectly updated after a trespass operation for case (2) above was performed.

Note that with the application of this patch, the integrity of the giveaway database will be maintained for the LOCAL failover relationship following a failover of the disks on LOCAL to LOCAL. However the entries in the giveaway database for the REMOTE relationship will not reflect the current names of the disks as they are known on LOCAL. Therefore, in order to re-establish the LOCAL/REMOTE relationship for the disks that were failed over locally, the giveaway database must be updated with the local disk names as they are now known on LOCAL and resynchronized with REMOTE. This can be done with an application script.

#### **4.1.37 dgux\_5.4.2.p91**

Patch dgux\_5.4.2.p91 eliminates a panic from occurring when an NFS server receives a packet from a user that has more than 16 supplementary group IDS. The panic code will most often be 3000032 or 2000075 but other panic codes can occur. Like DG/UX, most UNIX operating systems limit the number of supplementary group IDs to 16. This patch corrects the panic occurring when communicating with an operating system supporting more than 16 supplementary group IDs.

#### **4.1.38 dgux\_5.4.2.p92**

Patch dgux\_5.4.2.p92 corrects 2000075 panics which occur as the result of referencing an unwired routine while performing a system shutdown. This problem is corrected by calling the routine in question earlier in the shutdown process.

#### **4.1.39 dgux\_5.4.2.p93**

Patch dgux\_5.4.2.p93 corrects 41000050 panics which occur as the result of a bug in the su.a kernel library. The panics are due to the dupb() streams routine not checking the streams message datablock reference count prior to incrementing that count.

#### **4.1.40 dgux\_5.4.2.p98**

Patch dgux\_5.4.2.p98 corrects two problems seen when using 3rd party SCSI-2 disks.

3rd party Micropolis disks in SCSI-2 mode could not be registered. The drive indicated it performed SCSI-2 Command Tagged Queuing, but when DG/UX attempted to send a tagged command, the drive rejected it and the operation failed. Currently, DG/UX checks the inquiry block to determine if Tagged Queue Commands are supported. DG/UX now verifies that this feature is actually enabled in the Control Mode Page instead of assuming that Tagged Queuing is enabled. If DG/UX detects that Tagged Queuing is off but the drive indicates that it supports this feature, an attempt will be made to turn it on. If that attempt fails, tagged commands will not be sent to the drive.

3rd party Seagate and IBM disks in SCSI-2 mode would see the error 'unable to flush buffer' when data was written to them. These errors would eventually cause the filesystem to be downgraded to read-only mode. This problem was occurred because these drives have shallower queues in them than DG drives. When the queue filled up with commands, subsequent commands were rejected with the QueueFull status which DG/UX didn't expected. This status was incorrectly interpreted as the "EBUSY" status so the command was actually retried. Now, when DG/UX detects the QueueFull status, it will adjust the drive's command queue information so the disk's queue will not be flooded.

**4.1.41 dgux\_5.4.2.p99**

Patch dgux\_5.4.2.p99 corrects a problem in the ELF environment where the linker miscalculates the size of the dynamic string table which results in linker errors. Programs that have object files that reference `_etext`, `_edata`, or `_end` may run into this problem.

**4.1.42 dgux\_5.4.2.p100**

Patch dgux\_5.4.2.p100 closes a small window where a biod server could perform a read ahead on a closed file at the same time the file is being reopened. This reopen race results in a 2000075 nfm panic.

**4.1.43 dgux\_5.4.2.p101**

Patch dgux\_5.4.2.p101 increases the number of blocks that the kernel resident fsck can process from 600,000 to 4,000,000 blocks. Prior to this enhancement, the user would receive an error from fsck indicating the file system was too large.

This patch and its required patch level supersedes previous patches to fsck.a and ffm.a includes includes the following fixes:

- corrects a problem which can lead to the occurrence of a panic 13000057 when using fast recovery file systems. This problem occurred when a data element was not cleared as it should have been during the handling of a file truncation. This patch includes a code fix which was inadvertently left out of patch dgux\_5.4.2.p50 which was originally issued to correct this problem.
- corrects an occurrence of a panic 13000057. This problem occurred because DG/UX failed to convert a fast recovery file system to regular recovery mode if it was downgraded to readonly mode for some reason.

**4.1.44 dgux\_5.4.2.p102**

Patch dgux\_5.4.2.p102 corrects a problem in the att\_kbd streams module which allows all of available streams memory to be used and not released. This may eventually cause a system hang.

This patch provides additional support for the VTC device.

The VTC/256 controller is an ethernet controller supporting up to 256 network connections. The VTC/256 supports a general board Internet address and optional tty specific Internet addresses. The VTC/256 supports telnet, rlogin, and tcp connections. The VTC/256 also supports call out connections. See the `vtc_ttyaddr(1M)` man page for details.

The VTC/256 controller also has an SNMP agent. The VTC's SNMP communities, hosts, traps, and site specific objects can be specified by modifying the VTC SNMP files located in `/etc/tcload`. See the `vtc_snmp_files(4M)` man page for detailed information. The VTC's default SNMP community is public with read access only.

The `/etc/tcload/vtc.addr`s file may need to be edited to conform to the new format. To achieve equivalent functionality from the previous VTC software revision, no changes are necessary.

The changes that are needed relate to tty specific entries. If your `vtc.addr`s file has no tty specific entries then you do not need to make any changes, although you may want to make a change to your device entry if you wish to take advantage of the network rejection message (i.e. obtain specific network rejection messages from the VTC whenever it rejects a connection).

The new tty specific entries are as follows:

```
<Tty path> [<Inet Addr>][@<Port>] [<Protocol> [<Binary flag> [<Input flag>]]]
```

The old tty specific entries are as follows:

```
<Tty path> <Inet Addr> [<Binary flag> [<Input flag>]]
```

If none of your tty specific entries used the <Binary flag> or <Input flag> fields and you do not want to change the behavior of your tty specific entries then no changes are needed. In other words, if you don't use the binary and input flags and you want the tty lines to remain as "telnetin" lines at port 23, then you don't need to make any changes.

Otherwise, you must make changes to your tty specific entries. Specifically, you must add the <Protocol> field if you wish to change the protocol or do not want to use the default binary/input flag values. See the vtc.addr(4m) for complete details on tty specific entries.

The system administrator should install the vtc.snmp.config, vtc.snmp.hosts, vtc.snmp.communities and vtc.snmp.traps files from the corresponding vtc\_snmp\_files(4m) man page for more information.

The three new SYAC VTC related IOCTLs are SYAC\_SET\_PORTSERV, SYAC\_GET\_PORTSERV, and SYAC\_SET\_CHADDR. /usr/include/sys/\_int\_dg\_syac\_ioctl.h must be included to use these new IOCTLs. A description of each IOCTL and the related structure is defined below. Also see /usr/include/sys/\_int\_dg\_syac\_ioctl.h for all related definitions.

```
int ioctl (fildes, <new ioctls>, syac_network_param_type *npp)
```

Where:

fildes is a valid, active file descriptor.

```
typedef struct
{
  long    np_addr1; /* Internet address */
  union
  {
    long   un_addr; /* Gateway address */
    short  un_ports[2]; /* Protocol Ports */
    long   un_entry; /* Entry in VTC's port */
    struct /* service table */
    {
      short port;
      short channel;
    } un_pch; /* Port and channel */
  } np_un;
  short   np_flags; /* Various flags */
  short   np_service; /* Service type ID */
  char    np_enet[6]; /* Ethernet Address */
  short   np_protocol; /* Protocol ID */
}
/*<-----*/
syac_network_param_type;
/*>-----*/
```

### SYAC\_SET\_PORTSERV

The np\_un.un\_ports, np\_service, np\_protocol fields must be defined. This IOCTL sets the service for a range of ports (i.e. np\_un.unport[1] to np\_un.unport[2])



**SYAC\_GET\_PORTSERV**

The `np_un.un_entry` field must be defined. The VTC maintains a table of the ports and their corresponding protocol services. To obtain a complete copy of this table, the user must issue this IOCTL for all entries. This table begins with entry number zero. The lowest entry number that causes this IOCTL to return an error marks the end of the VTC's port service table.

**SYAC\_SET\_CHADDR**

The `np_addr1`, `np_un.un_pch.port`, `np_un.un_pch.channel`, and `np_service` fields must be defined. For this IOCTL, the `np_service` is used to define the type of channel and not the port's service. For example, the `np_service` set to `SERV_TELNET` means that this channel will call-out with the telnet service.

Also included in this patch are the following patches for ts.a:

A problem has been corrected whereby a process running on the duart port of the AV/46XX series computer may hang. This patch works around a problem with the duart chip to prevent the hang.

This patch will change the way in which DG/UX will handle an error returned from a syac IOCB. Currently, DG/UX will panic with a 3000017 (`TS_PANIC_SYAC_IOCB_RETURN_ERROR`). With this patch installed, Terminal Services will reset the syac that generated the IOCB error.

This patch corrects two separate VTC RAM dumping problems. These two problems are only encountered while attempting to dump a VTC's RAM after a system panic. The first problem results in the system hanging while attempting to dump the VTC's RAM. The second problem results in the system reporting that it is dumping VTC RAM with a size of zero. This patch prevents these problems from occurring by avoiding a VTC firmware bug and a VTC download software bug, respectively.

This patch corrects a problem with the `_INLINE` definition in the include file, `/usr/include/sys/_int_features.h`. Currently, C++ programs that include `/usr/include/sys/signal.h` will not compile because `_INLINE` is defined incorrectly. This patch modifies the `_INLINE` definition in `/usr/include/sys/_int_features.h` so that it will be set to "inline" (a C++ keyword) when `__cplusplus` is defined.

This patch also corrects a problem with the `sys/types.h` header file not defining certain data types in the `_SYSV4_SOURCE` environment. The `types.h` header file has been modified so that those data types are visible when the `_SYSV4_SOURCE` environment is used.

This patch can also be installed on DG/UX 5.4.2 Trusted (C2/B1) systems.

**4.1.45 dgux\_5.4.2.p103**

Patch `dgux_5.4.2.p103` changes UUCP so that the system node name is truncated to 7 characters when creating execute files. Prior to this change, the name of the execute file would be 15 characters in length, and on some remote systems, the execute file would not be processed.

**4.1.46 dgux\_5.4.2.p104**

Patch `dgux_5.4.2.p104` corrects a problem where file access was not being checked before writing a file in a directory or before truncating a file on a DG/UX NFS filesystem when mounted on another vendor's operating system. The problem is not apparent when mounting a DG/UX NFS file system on a DG/UX system.

Also included in this patch is previously released patch `dgux_5.4.2.p94`. This patch corrects a problem in the `lfm.a`

kernel library whereby the library routine `lfm_read_file` was not properly bracketing a call to `ffm_read_file` call with the `fsls_start` and `fsls_end` transaction calls. This problem causes the system to panic with panic code 54000005.

In addition, this patch includes a correction to a bug in the `lfm.a` kernel library which caused file record locks to be (incorrectly) released when a file was unlinked or a directory removed. This caused improper program execution.

#### 4.1.47 dgux\_5.4.2.p105

Patch `dgux_5.4.2.p105` corrects a problem with the `select(2)` system call. The system call, `write(2)`, when writing to a pipe with a size greater than `PIPE_BUF` (default 8192 characters) may have failed to satisfy a read `select(2)` event on the destination end of the pipe. This is corrected with this patch.

This problem is fixed in DG/UX 5.4 Release 3.00 and in patches against, DG/UX 5.4 Release 2.01 and DG/UX 5.4 Release 2.10.

This patch includes all the updates from previous patches to this set of subsystems and is documented based on the patch number which the fix was introduced.

#### 4.1.47.1 dgux\_5.4.2.p96

Patch `dgux_5.4.2.p96` addressed a problem in the `sfm.a` kernel library whereby DG/UX will panic if the running kernel allocates all of the stream queue pairs. The typical panic will be a 2000075 panic code. The number of queue pairs in the system is set by the kernel configuration parameter `NQUEUE`.

#### 4.1.47.2 dgux\_5.4.2.p89

Patch `dgux_5.4.2.p89` corrected a problem in the `sfm.a` kernel library which causes 2000075 panics. These panics occur because of incorrect handling of `M_FLUSH` messages in the streams code.

#### 4.1.47.3 dgux\_5.4.2.p84

Patch `dgux_5.4.2.p84` corrected two problems in the DG/UX kernel:

The first problem corrected by `dgux_5.4.2.p84` is seen as either a 2000075 panic or `lpsched` hangs. These occur as the result of two processes obtaining locks on one end of a pipe, followed by one of the processes releasing its lock and deleting the end of the pipe it held.

The second problem corrected by `dgux_5.4.2.p84` is also seen as a 2000075 panic, when a streams module is removed from a FIFO stream. This occurs due to a design flaw in the `sfm.a` kernel library whereby the `I_PUSH` and `I_POP` streams `ioctl` calls behave incorrectly when inserting and removing streams modules.

#### 4.1.47.4 dgux\_5.4.2.p82

Patch `dgux_5.4.2.p82` corrected a problem in the `sfm.a` kernel library which could cause a 3000032 panic when certain `ioctl()` operations are performed on pipes with one end closed.

#### 4.1.47.5 dgux\_5.4.2.p69

Patch `dgux_5.4.2.p69` corrected a problem in the signal handling code. This problem caused a process to not be allowed to terminate. This process consumes a large amount of CPU time and must be killed by hand. This patch corrects the problem by enabling the loss of the controlling tty to be detected and subsequently returning a `HANGUP` signal. This patch also fixes an atomic `opt` depth problem. It corrects panic 3000013 from occurring when `SFM` does not properly release a stream lock. In addition, the patch corrects the lock up of terminal services when a process attempts to get a `syac` lock which has not been properly released.

**4.1.47.6 dgux\_5.4.2.p62**

Patch dgux\_5.4.2.p62 corrected a race condition in the STREAMS open path which could cause panic code 2000107 to occur. The race condition incorrectly allowed more than one stream to be created with a single device number. This problem has been corrected by setting the device number while still holding a lock rather than doing it after the lock is released.

**4.1.47.7 dgux\_5.4.2.p40**

Patch dgux\_5.4.2.p40 corrected 2000075 panics which could occur if a process was closing a FIFO at the same time another process is trying to open the same FIFO. Depending upon the timing of these two operations, the opening process may encounter the stream structures in a partially disassembled state due to the close, resulting in the reference of an invalid pointer which leads to the panic.

**4.1.47.8 dgux\_5.4.2.p05**

Patch dgux\_5.4.2.p05 which corrected a problem where under heavy use of the streams logging utility strlog() while the streams log daemon [strerr user command] and/or the streams tracer [strace user command], the system may panic due to memory corruption. The typical panic code reported is 41000053, however, many other panic codes are possible. The sfm.a library has been modified to correct the problem with the locking associated with the strlog driver's queues.

**4.1.48 dgux\_5.4.2.p106**

Patch dgux\_5.4.2.p106 corrects a problem where DG/UX could miss an interrupt from the NCR53c710 chip on the "dgsc" controller and experience a SCSI timeout.

This patch fixes a problem that could sometimes cause 53000023 or 53000053 panics during SCSI reset/recovery.

Also, this patch fixes a problem where 9-track tape drives would be incorrectly configured as 8mm tape drives.

Finally, this patch fixes a SCSI polling problem that could possibly cause hangs on systems with CLARiiON tape arrays, CLARiiON disk arrays, 8mm tapes, or 4mm tapes.

Also included in this patch are the following changes from an earlier dev.a patch.

3rd party Micropolis disks in SCSI-2 mode could not be registered. The drive indicated it performed SCSI-2 Command Tagged Queuing, but when DG/UX attempted to send a tagged command, the drive rejected it and the operation failed. Currently, DG/UX checks the inquiry block to determine if Tagged Queue Commands are supported. DG/UX now verifies that this feature is actually enabled in the Control Mode Page instead of assuming that Tagged Queuing is enabled. If DG/UX detects that Tagged Queuing is off but the drive indicates that it supports this feature, an attempt will be made to turn it on. If that attempt fails, tagged commands will not be sent to the drive.

3rd party Seagate and IBM disks in SCSI-2 mode would see the error 'unable to flush buffer' when data was written to them. These errors would eventually cause the filesystem to be downgraded to read-only mode. This problem was occurred because these drives have shallower queues in them than DG drives. When the queue filled up with commands, subsequent commands were rejected with the QueueFull status which DG/UX didn't expect. This status was incorrectly interpreted as the "EBUSY" status so the command was actually retried. Now, when DG/UX detects the QueueFull status, it will adjust the drive's command queue information so the disk's queue will not be flooded.

**4.2 Package 2 - Product Patch Level tcpip\_5.4.2.14**

TCP/IP Patch Level tcpip\_5.4.2.14 contains the previously release TCP/IP Patch Level tcpip\_5.4.2.04 and tcpip\_5.4.2.14 in its entirety. Both the old and the new contents are described below.

#### 4.2.1 tcpip\_5.4.2.p02

Patch tcpip\_5.4.2.p02 corrects two problems. First is a problem with IP receiving a packet and decoding it incorrectly as a subnet broadcast. If a network is configured using subnetting and the subnets do not use the same number of bits for the subnet identifier, then if a client machine tries to reach a host on another subnet, and the AViiON is the router for this request, it is possible for the AViiON to think that the destination IP address of a packet is a subnet broadcast for one of its interfaces and reply to the request (typically a ping request).

The problem is that the AViiON determines the network ID based upon the class of the IP address, then for each interface examines the IP address for a match when masked with the subnet mask and the network ID of the interface. If there is a match of either, then the comparison is made against the last bits for a host ID and if they are all 1's, the AViiON thinks it was a broadcast.

The fix is to also check the subnet ID for a match before returning success for a subnet broadcast.

The second problem addressed is if the arp ioctl request is incorrectly formed, an error would not be returned to the user, and no arp entry would be added to the arp tables. Now, an error will be returned to the user if the ioctl is not in the correct format.

#### 4.2.2 tcpip\_5.4.2.p03

Patch tcpip\_5.4.2.p03 corrects a problem in sendmail whereby, regardless of the network configuration or sendmail configuration, sendmail will try to use DNS to fully qualify a hostname. This results in long delays sending mail in a NON-DNS environment if the user sends to a high user count. Each request to resolve a hostname takes about 40 seconds. For large mailing lists, the mail can be delayed for a long time (For example, for a mailing list with 60 users it takes 40 minutes to deliver the mail.)

This new sendmail fixes a bug where the simple hostname returned by NIS (or resolved via /etc/hosts), would invoke the resolver code erroneously.

This patch also addresses a problem whereby sendmail will not deliver mail if the To: line on the message header contains more than 80 usernames. Sendmail has been modified to use a larger buffer to handle long To: fields. This problem does not occur if a mailing list is used and it contains more than 80 usernames because sendmail expands the mailing list internally to send to all users, but the To: line only contains the mailing list name.

#### 4.2.3 tcpip\_5.4.2.p07

Patch tcpip\_5.4.2.p07 corrects a problem whereby TCP/IP may reference a memory buffer that has already been freed. The TCP set buffer size macros were modified to check the endpoint queue pointer as well as the the SS\_NOFDREF flag before updating the high and low water marks. The panic occurs when a connection is lingering on close and a connection request packet is received on the endpoint.

TCP, IP and UDP are modified to avoid deadlock trying to obtain a lock for which the subsystem already holds.

Also a minor leak in the TCP TPI engine is corrected. Leak occurs for embryonic connections which are aborted before reaching the connected state. The Released the unused T\_DISCON\_IND message buffer. There would be a 92 byte buffer that would not be freed for each connection that entered this state.

Also corrected in this patch is the IP statistics routine. It failed to check flowcontrol before queuing the DL\_GET\_STATISTICS\_REQ message for a blocked interface. As this is done periodically, the system could run out of STREAMS memory resulting in a system hang.

This patch also corrects a problem whereby TCP/IP may not release an endpoint lock, eventually causing all processes to hang as they execute streams code that requires the same locked resource. The processes do not have to

be network processes, any process accessing streams can be affected.

In addition, this patch address a problem where a socket endpoint timer could be started before the user closed the endpoint. TCP/IP has been changed to not start the timer unless the socket has been closed by the user. This problem can be observed by starting the "on" command in the background. After a period of inactivity (several hours), the command will terminate.

#### 4.2.4 tcpip\_5.4.2.p09

Patch tcpip\_5.4.2.p09 adds the `initgroup()` call to `ftpd`. This will allow for the initialization of a ftp user's supplementary group list. See the `initgroups(3C)` man page for more details.

This patch also modifies the check of `autologin`'s username/password in the `.netrc` file. Now, the lookup of `autologin`'s username/password will be based on the specified destination rather than converting system alias names back to the primary system name.

#### 4.2.5 tcpip\_5.4.2.p11

Patch tcpip\_5.4.2.p11 adds the `-e` option to `telnet` providing the capability of specifying an escape character string on the command line. The escape character string must follow the `-e` option as a separate argument and must be enclosed in quotes. A white space must appear between the `-e` option and argument string. For example, if one does not want an escape sequence, specify a null string as the argument to the `-e` option, i.e. `telnet -e ""`.

Note when specifying escape character strings some key sequences will not be recognizable when `editread` is enabled. Also, when specifying escape character strings make sure that there are no conflicts with any `stty` settings.

#### 4.2.6 tcpip\_5.4.2.p12

Patch tcpip\_5.4.2.p12 corrects a problem with the `select()` satisfy on a connection confirmation. TCP connect confirmations incorrectly satisfy `select(s)` with read intent. This could lead to a `recv()` incorrectly receiving an `EAGAIN` error. This patch contains previously released patch tcpip\_5.4.2.p08 correcting a problem where a `select` operation would incorrectly return a status of zero for a socket when multiple processes are using the same socket. The patch modifies the `select` code in the library `inso.a` to return the intent that satisfied the `select`, rather than the current status. This problem occurs with `DG/UX` revision 5.4.2 and occurs more often on multi-processor systems.

#### 4.2.7 tcpip\_5.4.2.p13

Patch tcpip\_5.4.2.p13 adds support for a kernel enhancement that will allow the keep alive timer and maximum segment size for TCP/IP to be configurable kernel parameters. The new kernel parameters are `TCPKEEPIDLE` and `TCPMSS`. To use, add these variables to the kernel configuration file, rebuild the kernel and reboot.

The current default for the keep alive timer is 7200 seconds (2 hours). The value of the keep alive timer represents the number of seconds an endpoint is idle before the first keepalive probe is sent. With this patch, the value of the keep alive timer can be adjusted. When configuring a new kernel with `"sysadm newdgux"`, specify the new TCP/IP kernel variable `TCPKEEPIDLE` and give it the desired timeout value, in seconds. The minimum value for `TCPKEEPIDLE` is 75 seconds (a probe interval) and the maximum value is 65535 seconds. These limits are silently enforced and if exceeded, the limits will be used.

This patch also adds a kernel enhancement that will allow the maximum segment size (MSS) for TCP/IP to be a configurable kernel parameter. The value of the maximum segment size represents the amount of data TCP will

send in one packet. The current method for determining the default maximum segment size depends upon the location of the peer. If the peer is on a directly connected network, the MTU from the network interface is used to calculate the TCP MSS. If the peer is non-local, the TCP MSS is calculated using IP\_MSS(576) as the MTU (i.e. TCP MSS will be 536). With this patch, the TCP MSS value can be specified directly. When configuring a new kernel with "sysadm newdgux", specify the new TCP/IP kernel variable TCPMSS and give it the desired maximum segment size value in bytes. The minimum value for TCPMSS is 0 which means that TCP will use the current default method of determining the maximum segment size. Its maximum value will be 65536 bytes

**WARNING!!!** Setting the TCPKEEPIDLE and TCPMSS will be supported by a different mechanism in DG/UX 5.4 Release 3.00. There are no plans to support the mechanism implemented in this patch in any later revision of DG/UX.

#### 4.2.8 tcpip\_5.4.2.p14

Patch tcpip\_5.4.2.p14 eliminates a vulnerability in sendmail. The CERT Coordination Center has issued a warning (Advisory Number CA-93:16) regarding security holes in the sendmail program of most vendors.

The vulnerability allows unauthorized remote or local users to execute programs as any system user other than root. This vulnerability effects the final destination sendmail host and can be exploited through an intermediate mail machine. Therefore, all sendmail recipient machines within a domain are potentially vulnerable. This patch corrects this problem.

**WARNING!!!** Setting the TCPKEEPIDLE and TCPMSS will be supported by a different mechanism in DG/UX 5.4 Release 3.00. There are no plans to support the mechanism implemented in this patch in any later revision of DG/UX.

#### 4.3 Package 3 - Product Patch Level nfs\_5.4.2.03

NFS Patch Level nfs\_5.4.2.03 contains NFS Patch Level nfs\_5.4.2.02 and is described below:

##### 4.3.1 nfs\_5.4.2.p01

Patch nfs\_5.4.1.p01 changes the way in which rpc.lockd accesses the TCP/IP transport system to send lock messages to clients and servers. Previously, rpc.lockd would allocate and try to bind to a privileged UPD port to generate the lock message for each new client. This operation may take several seconds if for some reason the system currently running has all of the privileged ports in use (ports numbered between 600 and 1023 are privileged). Netstat can be used to list all of the ports currently in use.

Rpc.lockd has been changed to allocate one privileged port and use it for all messages thereby avoiding this condition.

This problem may show up as a "spinning or run-away lockd".

##### 4.3.2 nfs\_5.4.2.p02

Patch nfs\_5.4.2.p02 corrects a problem where the rusersd daemon fails to return a list users when the login terminal name is of type tty[pqrs] and the environment variable DISPLAY has been set to ":0.0" The rusersd daemon now determines logins in the same manner as the who() command does.

#### 4.3.3 nfs\_5.4.2.p03

Patch nfs\_5.4.2.p03 corrects a problem where PCNFSD would incorrectly hold /dev/console open. On a B1 trusted system this prevented users from logging onto /dev/console more than once.

Note that this patch delivers revision 2.00 of PCNFSD. PCNFSD has been qualified with NFS Client for Lan Workplace. The uncompiled lwpnfs daemon included in the NFS client for Lan Workplace is not supported. In addition, BSD printing is no longer supported in this revision.

Also note that in order to connect to the PCNFSD Daemon, a user's UID must be between 101 and 60002. User's who UID is not within this range will receive an authentication error when trying to execute a net link command.

#### 4.4 Package 4 - Product Patch Level X11\_5.4.2.08

X11 Patch Level X11\_5.4.2.08 contains X11 Patch Level X11\_5.4.2.05 and is described below:

##### 4.4.1 X11\_5.4.2.p01

Patch X11\_5.4.2.p01 corrects four problems in the mterm client.

The mterm color menu option was greyed out on 24-bit systems, even after following the mterm man page instructions of setting the display visual for the Xserver to pseudo-color (-cc 3). This patch allows the mterm color menu to be used on a 24-bit system where the Xserver's display visual has been set to pseudo-color.

Changes mterm logging. The log file now mirrors what is displayed on the screen unless -debug was used on the command line. When -debug is used the log file format is in the 'debug' mode and prints all information as it has in previous revs.

Corrects a problem where the mterm D410 margins were not being set properly in compressed mode. The user is now able to set the margins.

The final problem corrected was reported by STR NASC-12548. Corrects a problem where mterm would override a user specified icon bitmap, if it was specified with the Mwm resource Mwm\*iconImage.

##### 4.4.2 X11\_5.4.2.p02

Patch X11\_5.4.2.p02 corrects two problems with the Xserver.

1. Lines being drawn by FrameMaker were not being clipped correctly. This caused the Xserver to enter an infinite loop. This patch corrects the problem.
2. Under heavy system load, a gif file displayed via the DGShm:0 shared memory transport did not display. This patch corrects the problem.

##### 4.4.3 X11\_5.4.2.p05

Patch X11\_5.4.2.p05 correct/supplies the following:

1. When certain GC algebra functions were used and a pixmap was used to redisplay an exposed area of the screen, the foreground and background had the same color.
2. Oracle CASE\*Designer would crash after choosing an option on the main menu bar.

3. All MIT patches up to and including patch # 2830.

#### 4.4.4 X11\_5.4.2.p07

Patch X11\_5.4.2.p07 Fixes a problem in which text areas do not display properly and some text areas are unreadable when running SoftPC with windows on a monochrome workstation.

#### 4.4.5 X11\_5.4.2.p08

Patch X11\_5.4.2.p08 corrects a problem in which applications generated with UIM/X revision 2.1 would not link. This patch supersedes X11\_5.4.2.p04 and includes libXm.a for static linking.

This patch also corrects a problem in libXm.so.2. If a program built with libXm.so.2 on DG/UX 5.4.1 is executed on a DG/UX 5.4.2 system, the XmListItemExists function returns false and the XmListItemPos function always returns 0. This behavior causes the SoftSwitch application to fail on a DG/UX 5.4.2 system.

If the same program is built and executed on 5.4.1 or built and executed on 5.4.2, the problem does not occur. It only occurs when it is built on 5.4.1 and executed on 5.4.2.

This patch also corrects a problem in which mwm would hang under certain conditions upon button release on certain widgets (e.g. - the scrollbar).

## 5. Files

A list of files for each product package included in patch level op-sys-x\_5.4.2.931203 is contained in a file in the /usr/release directory. To be able to back out of this patch level, these files should be saved prior to loading this patch level.

Package 1 - Product Patch Level dgux\_5.4.2.106

```

/dgux.installer
/dgux.starter
/etc/issue
/etc/tcload/vtc.addrs.proto
/etc/tcload/vtc.snmp.config/proto
/etc/tcload/vtc.snmp.hosts.proto
/etc/tcload/vtc.snmp.traps.proto
/etc/tcload/vtc/snmp/communities.proto
/sbin/fsck
/sbin/shutdown
/usr/bin/at
/usr/bin/cdc
/usr/bin/comb
/usr/bin/ct
/usr/bin/cu
/usr/bin/delta
/usr/bin/dtl
/usr/bin/get
/usr/bin/ksh
/usr/bin/login
/usr/bin/lpstat
/usr/bin/mail
/usr/bin/mailalias

```



/usr/bin/make  
/usr/bin/ps  
/usr/bin/rmail  
/usr/bin/rmdel  
/usr/bin/sar  
/usr/bin/uucp  
/usr/bin/uudecode  
/usr/bin/uuencode  
/usr/bin/uuglist  
/usr/bin/uuname  
/usr/bin/uustat  
/usr/bin/uux  
/usr/bin/xc  
/usr/catman/a\_man/man1/syac\_routes.1m.z  
/usr/catman/a\_man/man1/syac\_ttyaddrs.1m.z  
/usr/catman/a\_man/man1/tcload.1m.z  
/usr/catman/a\_man/man1/vtc\_routes.1m.z  
/usr/catman/a\_man/man1/vtc\_snmp\_start.1m.z  
/usr/catman/a\_man/man1/vtc\_snmp\_stop.1m.z  
/usr/catman/a\_man/man1/vtc\_ttyaddrs.1m.z  
/usr/catman/a\_man/man4/vtc.addrs.4m.z  
/usr/catman/a\_man/man4/vtc\_snmp\_files.4m.z  
/usr/catman/a\_man/man7/syac.7.z  
/usr/catman/u\_man/man1/datability.1.z  
/usr/catman/u\_man/man1/fastport.1.z  
/usr/catman/u\_man/man1/fpfilter.1.z  
/usr/catman/u\_man/man1/lp.termprinter.1.z  
/usr/catman/u\_man/man1/termprinter.1.z  
/usr/dglib/libc.so.1  
/usr/include/sys/\_int\_features.h  
/usr/include/sys/types.h  
/usr/lib/lan/cmc130.bin  
/usr/lib/libc.so  
/usr/lib/lp/bin/fpfilter  
/usr/lib/lp/bin/lp.termprinter  
/usr/lib/lp/bin/proprinter  
/usr/lib/lp/lpNet  
/usr/lib/lp/lpsched  
/usr/lib/lp/model/datability  
/usr/lib/lp/model/fastport  
/usr/lib/lp/model/standard  
/usr/lib/lp/model/termprinter  
/usr/lib/lp/postscript/postprint  
/usr/lib/mail/mail\_pipe  
/usr/lib/saf/ttymon  
/usr/lib/tcload/models/syac/LOAD  
/usr/lib/tcload/models/syac/VTC256  
/usr/lib/tcload/models/syac/vtclload  
/usr/lib/uucp/bnuconvert  
/usr/lib/uucp/remote.unknown  
/usr/lib/uucp/uucheck  
/usr/lib/uucp/uucico  
/usr/lib/uucp/uucleanup  
/usr/lib/uucp/uusched

/usr/lib/uucp/uuxqt  
/usr/options/dgux.name  
/usr/release/093701088\_2  
/usr/sbin/gridman  
/usr/sbin/lpsystem  
/usr/sbin/pmadm  
/usr/sbin/sacadm  
/usr/sbin/syncadm  
/usr/sbin/upsd  
/usr/sbin/vtc\_routes  
/usr/sbin/vtc\_snmp\_start  
/usr/sbin/vtc\_snmp\_stop  
/usr/sbin/vtc\_ttyaddrs  
/usr/sbin/wanadm  
/usr/sde/m88kbc/usr/lib/crt0.o  
/usr/sde/m88kbc/usr/lib/libc.a  
/usr/sde/m88kbc/usr/lib/libp/libc.a  
/usr/sde/m88kbc/usr/lib/mcrt0.o  
/usr/sde/m88kdguxcoff/usr/lib/crt0.o  
/usr/sde/m88kdguxcoff/usr/lib/libc.a  
/usr/sde/m88kdguxcoff/usr/lib/libp/libc.a  
/usr/sde/m88kdguxcoff/usr/lib/mcrt0.o  
/usr/sde/m88kdguxelf/usr/lib/libc.a  
/usr/sde/m88kdguxelf/usr/lib/libp/libc.a  
/usr/src/uts/aviion/lb/cdm.a  
/usr/src/uts/aviion/lb/dev.a  
/usr/src/uts/aviion/lb/dgen.a  
/usr/src/uts/aviion/lb/faam.a  
/usr/src/uts/aviion/lb/ffm.a  
/usr/src/uts/aviion/lb/fs.a  
/usr/src/uts/aviion/lb/fsck.a  
/usr/src/uts/aviion/lb/io.a  
/usr/src/uts/aviion/lb/iscd.a  
/usr/src/uts/aviion/lb/kfs.a  
/usr/src/uts/aviion/lb/lbm.a  
/usr/src/uts/aviion/lb/lfm.a  
/usr/src/uts/aviion/lb/nfm.a  
/usr/src/uts/aviion/lb/nmf.a  
/usr/src/uts/aviion/lb/pdep.a  
/usr/src/uts/aviion/lb/pfn.a  
/usr/src/uts/aviion/lb/pm.a  
/usr/src/uts/aviion/lb/rpc.a  
/usr/src/uts/aviion/lb/sam.a  
/usr/src/uts/aviion/lb/sc.a  
/usr/src/uts/aviion/lb/sfm.a  
/usr/src/uts/aviion/lb/so.a  
/usr/src/uts/aviion/lb/sys.a  
/usr/src/uts/aviion/lb/tpkg.a  
/usr/src/uts/aviion/lb/ts.a  
/usr/src/uts/aviion/lb/uc.a  
/usr/src/uts/aviion/lb/vm.a  
/usr/src/uts/aviion/lb/vmc.a  
/usr/stand/boot.aviion  
/usr/stand/diskman

**Package 2 - Product Patch Level tcpip\_5.4.2.14**

```

/usr/bin/ftp
/usr/bin/ftpd
/usr/bin/mailq
/usr/bin/newaliases
/usr/bin/sendmail
/usr/bin/smtp
/usr/bin/telnet
/usr/bin/telnetd
/usr/etc/master.d/tcpip
/usr/options/tcpip.name
/usr/src/uts/aviion/lb/conf.a
/usr/src/uts/aviion/lb/inet.a
/usr/src/uts/aviion/lb/inso.a
/usr/src/uts/aviion/lb/ip.a
/usr/src/uts/aviion/lb/tcp.a
/usr/src/uts/aviion/lb/tpi.a
/usr/src/uts/aviion/lb/udp.a

```

**Package 3 - Product Patch Level nfs\_5.4.2.03**

```

/usr/options/nfs.name
/usr/sbin/dg_lock_mgr
/usr/sbin/pcnfsd
/usr/sbin/rpc.lockd
/usr/sbin/rusersd

```

**Package 4 - Product Patch Level X11\_5.4.2.08**

```

/usr/opt/X11/bin/Xdg
/usr/opt/X11/bin/mterm
/usr/opt/X11/bin/mwm
/usr/opt/X11/catman/M_man/man1/mterm.1.z
/usr/opt/X11/lib/libXm.a
/usr/opt/X11/lib/libXm.so.2
/usr/opt/X11/options/X11.name
/usr/opt/X11/xhelp/C/Mterm/helpfile
/usr/opt/X11/xhelp/C/Mterm/helpfile.inx
/usr/opt/X11/xhelp/C/Mterm/helpfile.toc
usr/opt/X11/lib/app-defaults/Mterm
usr/opt/X11/xhelp/C/Mterm/helpfile
usr/opt/X11/xhelp/C/Mterm/helpfile.inx
usr/opt/X11/xhelp/C/Mterm/helpfile.toc

```

**6. Installation Instructions****6.1 Loading the Patch Level**

This op-sys-x patch level is in sysadm loadpackage format and consists of those packages listed in the section titled "Patch Level op-sys-x" containing the files as listed by package in the section titled "FILES" Section of this document. **IMPORTANT, LOAD ONLY THE PACKAGES THAT APPLY TO YOUR SYSTEM.**

To back out of this op-sys-x patch level, files listed in the section titled "Files" should be saved to another name or

backed up before proceeding. YOU SHOULD BE IN RUN LEVEL 1 WHEN LOADING THIS op-sys-x PATCH LEVEL TO AVOID OVERWRITING A PROGRAM THAT IS CURRENTLY RUNNING.

This op-sys-x patch level will load on any DG/UX 5.4.2 system.

### 6.2 Loading the op-sys-x Patch Level from Tape

To load the op-sys-x patch level, perform the following steps:

```
#sysadm loadpackage
Release Medium: [/dev/rmt/0] /dev/rmt/<tape_device>
Is /dev/rmt/<tape_device> ready? [yes]
Package Name(s): [all]
List file names while loading? [no]
OK to perform operation? [yes]
*Positioning the tape to load: <package name> .....
*Loading package: <package name> .....
*Package "<package name>" has been loaded.
Updating proto root (/srv/release/PRIMARY/root/root.proto).
Updating MY_HOST root (/srv/release/PRIMARY/root/MY_HOST).
Package load is finished.
The selected packages have been loaded.
```

\* These lines are repeated for each package selected for loading.

### 6.3 Loading the op-sys-x Patch Level from a Directory

If the op-sys-x Patch Level is delivered as a tar format file which consists of images that are sysadm loadable, you will need to extract the images into an empty directory. To extract these images from a provided file, perform the following steps:

```
# mkdir -p /var/tmp/patch_level
# cd /var/tmp/patch_level
# tar xvf <path_name>/op-sys-x_5.4.2.931203
```

To load the patch level from this directory structure, perform the following steps:

```
#sysadm loadpackage
Release Medium: [/dev/rmt/0] /var/tmp/patch_level
Package Name(s): [all]
List file names while loading? [no]
OK to perform operation? [yes]
*Loading package: <package name> .....
*Package "<package name>" has been loaded.
Updating proto root (/srv/release/PRIMARY/root/root.proto).
Updating MY_HOST root (/srv/release/PRIMARY/root/MY_HOST).
Package load is finished.
The selected packages have been loaded.
```

\* These lines are repeated for each package selected for loading.

To create a sysadm loadable patch tape which can be applied on any AViiON system with a tape drive, follow the

steps below:

```
# cd /var/tmp/patch_level
# ./Docs/MakeTape /dev/rmt/<tape_device>
```

#### 6.4 Installing the Patch Level

See Chapter 4 of the *Managing the DG/UX System* manual for instructions on reconfiguring the system. Use `sysadm newdgux` to build a new kernel. If you are a diskless server, your client kernels must be rebuilt or the new `dgux.diskless` in `/usr/stand` must be hard linked to the appropriate client's root area.

##### 6.4.1 Systems with Modems

If a `ttymon` on your system controls a tty line with a modem, then you must check and possibly modify the settings on this line. You should do this before rebooting the system, as described in a later section.

This patch requires that "Hangup" be set to "Yes" for all modem lines. This ensures that the modem will be hung up. If you do not set "Hangup" to yes, the modem will not hangup. To check the current status of the Hangup flag on your modem line use "admportservice -o list service\_tag" where "service\_tag" is usually the tty line associated with the modem. An example follows which shows how to check the settings on tty14.

```
# admportservice -o list tty14
PMTAG PMTYPE SVCTAG FLGS ID <PMSPECIFIC>
ttymon1 ttymon tty14 u root /dev/tty14 bh \
- /usr/bin/login 10 M2400 - login: ==tty14 disabled== \
#tty14--bi-directional uucp line
```

The presence of the "h" under <PMSPECIFIC> indicates that Hangup is set to No. Absence of an "h" under <PMSPECIFIC> indicates the line is configured correctly. As an "h" is present in this example, this port service must be modified.

To modify the port service use `sysadm` and select Device -> Port -> Port Service -> Modify.

#### Port Service Menu

- |               |                         |
|---------------|-------------------------|
| 1 Add ...     | Add port service        |
| 2 Delete ...  | Delete port service(s)  |
| 3 Modify ...  | Modify port service     |
| 4 List        | List port services      |
|               |                         |
| 5 Enable ...  | Enable port service(s)  |
| 6 Disable ... | Disable port service(s) |

Enter a number, a name, ? or <number>? for help, <NL> to redisplay menu,

^ to return to previous menu, or q to quit: 3

Controlling port monitor for service: [tcp (listen)] ttymon1

Service to Modify: [tty01] tty14

Port service tag: [tty14]

Service Userid: [root]

Create utmp entry? [yes]

```

File name of configuration script:
Comment: [tty14--bi-directional uucp line ]
Path name of terminal device: [/dev/tty14]
TTY Definition Label: [M2400]
Service command: [/usr/bin/login]
Hangup? [no] y
Connect on Carrier? [no]
Bidirectional? [yes]
Wait-read value: [none]
Timeout: (0-600) [10]
Prompt message: [login: ]
Modules to be pushed:
Disabled response message: [==== tty14 disabled ====]
OK to perform operation? [yes]
Service tty14 for port monitor ttymon1 has been modified.

```

Once you have modified the hangup setting on all your modem lines you can move onto the next installation step.

#### 6.4.2 Installing bootstrap

In order to install the bootstrap delivered with this patch level, you must bring your system down to the SCM and boot standalone diskman. To bring your system down to the SCM, perform the following steps:

```

# cd /
# shutdown -g0 -y
# halt -q

```

Boot standalone diskman. For example:

```
SCM> b sd(inc(0),0,0)usr:/stand/diskman
```

Follow the instructions below to install the bootstrap:

1. From the Diskman Main Menu select the "Physical Disk Management Menu".
2. From the Physical Disk Management Menu select the "Format a Physical Disk" option.
3. From the Format a Physical Disk Menu select the "Install a Bootstrap on a Physical Disk" option.
4. The system will prompt you for the disk specification of the disk where the bootstrap is to be installed:

Enter the physical disk specification in DG/UX common format:

Type in the physical disk specification, for example sd(inc(0),0,0).

5. The system will then ask if you want to install the bootstrap:

Physical disk sd(inc(),0) does not have an up-to-date bootstrap.  
Do you want to install one? [yes]

Answer "yes" to this question.

6. The system will then ask for the pathname of the bootstrap image:

Enter the pathname of the bootstrap image file: [/usr/stand/boot.aviion]  
Accepting this default will install the bootstrap on the selected physical disk. This procedure can be repeated for other bootable disks in the system by returning to step 4.

Once the bootstrap is installed, exit diskman and reboot your kernel.

--- End of Patch Level Notice ---

