ConvexOS and Utilities V11.5.1 Release Notice

Document No. 710-003430-066

September 1996

Hewlett-Packard Company
Convex Technology Center
Richardson, Texas
United States of America

ConvexOS and Utilities V11.5.1 Release Notice

Document No. 710-003430-066

Notice

The information contained in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance or use of this material.

 This entire book is recyclable.

Printed in the United States of America

Contents

Overview

1

This release notice describes the V11.5.1 release of ConvexOS and the ConvexOS Utilities. It highlights new features and changes to existing features, and supplements the permanent documentation with information developed too late for inclusion. This release notice also lists fixes and workarounds that may save time if you encounter a known problem. Always refer to this document before reporting problems; your questions may be answered here.

Prerequisites

The V11.5.1 release of ConvexOS has the following prerequisites:

- The VME ethernet controller firmware must be at revision 6.4 or later. If it is not, reconfiguring the eth devices will not be supported and instability can result.
- Your system must have the versions of SPU software listed in Table 1.

Table 1 SPU software dependencies

CONVEX system	SPU OS	System Diagnostics or Processor Diagnostics	I/O Diagnostics
C200, C3200 Series	V6.1 or later	V5.2 or later	1.1
C3400 Series	V6.1	V2.1.1 or later	1.1
C3800 Series	V2.0	V4.1 or later	4.2
C4600 Series	V2.0	V1.0.3 or later	1.0.3

- You should read *ConvexOS and Utilities V11.5.1 Installation Procedures* before beginning the installation. This installation procedures document describes both initial and upgrade installations.

Notes

Remote installations, kernel upgrades from the SPU, and installation from ct-format cartridge tape are not supported in ConvexOS V11.5.1. If your system does not have a local tape drive, contact the Technical Assistance Center (TAC).

CONVEX C1 systems are not supported by the ConvexOS V11.5.1 release.

- The ConvexOS V11.5.1 installation procedure may require machine-specific activation keys. (The only optional product bundled with ConvexOS that requires an activation key is the CONVEX Share Scheduler.) Activation keys are shipped in an envelope attached to the ConvexOS V11.5.1 installation tape. Be careful not to discard the activation keys before you have completed the installation process.

Optional products

This section describes requirements for optional products.

Most recent versions of optional products

Table 2 lists, in alphabetical order, the most recent versions of optional products that are compatible with ConvexOS and Utilities V11.5.1. Best-effort support will be provided under ConvexOS V11.5.1 for optional products appearing in shaded entries.

Table 2 Most recent versions of optional products

Product	V11.5.1-compatible version C2 & C3	V11.5.1-compatible version C4
CONVEX Ada	V2.1.3	V2.1.3
CONVEX ALL	V3.0.2	V3.0.3
CONVEX Application Compiler	V2.1	V2.1
CONVEX AVS	V5.0	V5.0
CONVEX C	V5.0/5.0.2	V6.1
CONVEX C++	V1.1.2	V1.1.2
COVUEnet	V3.0	V3.0
COVUEbatch	V2.2	V2.2
COVUEbinary	V1.0.1	V1.0.1
COVUEedt	V1.2.1	V1.2.1
COVUElib	V2.0.1	V2.0.1
COVUEshell	V8.2.3	V8.2.3
CONVEX CXbatch	V3.0	V3.0
CONVEX CXdb	V2.0	V3.1
CONVEX CXmetrics	V1.0	V1.0
CONVEX CXpa	V2.0	V3.1.1
CONVEX CXwindows	V3.1.1	V3.1.1
CONVEX FDDI	V3.1	V3.1
CONVEX FORTRAN	V8.0/8.0.4	V9.4
CONVEX Internet Services	V11.5.1	V11.5.1
CONVEX MLIB	V9.0	V9.0
CONVEX NFS	V11.5	V11.5
CONVEX OSI WAN	V1.2	V1.2
CONVEX PVM	V3.3.10	V3.3.10

Table 2 Most recent versions of optional products (continued)

Product	V11.5.1-compatible version C2 & C3	V11.5.1-compatible version C4
CONVEX Share Scheduler	V11.5	V11.5
CONVEX Toolbox	V1.1	V1.1
CONVEX UltraNet	V2.1	V2.1

Optional products requiring reinstallation

Due to changes to the ConvexOS kernel, the following optional products must be reinstalled after installing ConvexOS and Utilities V11.5.1:

- CONVEX COVUEnet
- CONVEX FDDI Interface
- CONVEX NFS
- CONVEX OSI WAN Transport
- CONVEX UltraNet Interface

CXTS will not run

Note

CXTS, the Convex Xpert Troubleshooting System, is not supported and will not run under ConvexOS V11.5.1.

UniTree+ V1.7 not supported

Note

UniTree+ V1.7 is not supported under ConvexOS V11.5.1.

UniTree+ FTP patch required

Due to changes in the `bind(2)/connect(2)` system calls required to resolve security issues, UniTree+ patch `unitree/2.0.218` is required with ConvexOS V11.5.1 to provide an updated UniTree+ FTP server, `uftpd`.

Checkpoint/Restart

Processes that are checkpointed before the upgrade to ConvexOS V11.5.1 cannot be restarted under ConvexOS V11.5.1. All checkpointed jobs should be completed before the upgrade.

Binary compatibility with earlier versions of ConvexOS

Although every effort has been made to preserve binary compatibility between ConvexOS V11.5.1 and executables compiled and linked on previous releases, CONVEX recommends that any executables last built on ConvexOS V8.1 or earlier be recompiled and relinked. Networking applications built on earlier versions of ConvexOS are the most likely to have problems.

Associated documentation

The following document is new for ConvexOS V11.5.1:

- *ConvexOS V11.5: Technical Overview and Managing ConvexOS Addendum,* Second Edition, document number 710-001430-215.

The following documents have been previously published and are current for ConvexOS V11.5.1:

- *CONVEX 3800 Series SPU System Manager's Guide,* First Edition (DSW-023)
- *CONVEX Architecture Reference,* Seventh Edition (DHW-300)
- *CONVEX Assembly Language Reference,* First Edition (DHW-301)
- *C4600 Assembly Language Reference,* First Edition (DHW-350)
- *CONVEX Compiler Utilities User's Guide,* First Edition (DSW-096)
- *CONVEX Networking Concepts,* Second Edition (DSW-128)
- *CONVEX POSIX Conformance,* Third Edition (DSW-311)
- *CONVEX PRTSMail User's Guide,* Fifth Edition (DSW-600)
- *CONVEX SPU System Manager's Guide,* First Edition (DSW-022)
- *CONVEX adb Command Summary Quick Reference,* First Edition (DSW-601)
- *CONVEX adb Debugger User's Guide,* Seventh Edition (DSW-009)
- *CONVEX vi Quick Reference,* First Edition, Rev. 1 (DSW-019)
- *ConvexOS Extensions User's Guide,* Second Edition. (DSW-053)
- *ConvexOS Primer,* First Edition. (DSW-133)
- *ConvexOS Tutorial Papers,* Eighth Edition (DSW-002)
- *The C Programming Language,* Second Edition (DSW-046)
  By Brian Kernighan and Dennis Ritchie, published by Prentice-Hall
- *GNU Emacs Manual,* Sixth Edition (DSW-050)
  By Richard Stallman, published by the Free Software Foundation
- *Managing ConvexOS: Configuration Guide,* Fourth Edition (DSW-030)
- *Managing ConvexOS: Operations Guide,* Fourth Edition (DSW-031)
- *Programming perl,* First Edition (DSW-051)
  By Larry Wall and Randal L. Schwartz, published by O'Reilly and Associates, Inc.

Internet Services and Network File Systems (NFS)

The following documents have been previously published and are current for ConvexOS V11.5.1:

- *CONVEX Network Programming Guide,* First Edition (DSW-106)
- *Managing CONVEX Internet Services and NFS,* First Edition (DSW-108)
- *Using CONVEX Internet Services and NFS Quick Reference,* Second Edition (DSW-118).

Share Scheduler

The Share Scheduler product has not changed and the following document is current:

- *CONVEX Share Scheduler System Manager's Guide,* First Edition (DSW-268)

New features

2

This chapter describes the following new features in ConvexOS V11.5.1:

- Quad SCSI Channel (QSC) support
- VME SCSI tape support
- Contributed Utilites
- Kernel panic signature
- Tunables
- HIPPI/TCP support
- ConvexOS Realtime and Utilities support

Quad SCSI Channel (QSC) support

Support for the Quad SCSI CCU (QSC) has been incorporated into the ConvexOS V11.5.1 release. The QSC is a PBUS CCU providing four fast and wide differential SCSI busses. The QSC and related hardware are designed to be used with many types of SCSI peripherals such as the Elite 9 disk drive and various high performance tape drives.

By eliminating the concept of virtual memory address space on the CCU and using the CPU's ability to do virtual-to-physical address translation, the QSC realizes a significant improvement over previous channel controllers in PBUS efficiency.

Example `/ioconfig` entries for Quad SCSI Channel devices are shown in the following example:

```
qsc 37
  scsi 0 SCS-021
     unit 0 subunit 0 type DKD-402 sd0
     unit 1 subunit 0 type DKD-402 sd1
  scsi 1 SCS-021
     unit 0 subunit 0 type DKD-402 sd2
     unit 1 subunit 0 type DKD-402 sd3
qsc 39
  scsi 0 SCS-021
     unit 0 subunit 0 type MTD-503
     unit 1 subunit 0 type MTD-503
  scsi 1 SCS-021
     unit 0 subunit 0 type MTD-503
     unit 1 subunit 0 type MTD-503
```

Supported Quad SCSI Channel devices are listed in Table 3.

Table 3 Supported Quad SCSI Channel devices

/ioconfig designation	Description	Device file
SCS-021	QSC SCSI controller	-
DKD-402	Seagate Elite9 SCSI disk	sd*
MTD-500	HP/Convex C590 IBM Magstar (NTP)	ctp*
MTD-502	STK SD-3 (Redwood)	rdw*
MTD-503	STK 9490 (Timberline)	tc*

VME SCSI tape support

Enhanced support for VME SCSI tape devices is provided in the ConvexOS V11.5.1 release. In addition to improved robustness of the device driver, support for several new tape devices has been added.

Table 4 lists the officially supported VME SCSI tape drives in ConvexOS V11.5.1.

Table 4 Supported VME SCSI tape devices

/ioconfig designation	Description	Device file
MTC-202	VME SCSI controller	-
MTD-207	Fujitsu SCSI 3480	tc*
MTD-217	Fujitsu SCSI 3480 with stacker	
MTD-227	Fujitsu SCSI 3480 with compression	
MTD-237	Fujitsu SCSI 3480 with stacker & compression	
MTD-208	DAT (Python)	dat*
MTD-228	DAT (Diamond Back)	
MTD-209	Metrum RSP-2150	rsp*
MTD-20A	Exabyte 8200	exb*
MTD-20B	Exabyte 8500	
MTD-20B	Exabyte 8505	
MTD-20C	DLT2000/DLT4000[a]	dlt*
MTD-22C	DLT2000/DLT4000 with compression	
MTD-20D	Fujitsu SCSI 3490E	tc*
MTD-21D	Fujitsu SCSI 3490E with stacker	
MTD-22D	Fujitsu SCSI 3490E with compression	
MTD-23D	Fujitsu SCSI 3490E with stacker & compression	

a. For DLT4000 tape drives, the following line must be added to bootcmd.local:

```
tune viop st_dlt_synch = 0
```

The following enhancements have been introduced into the VME SCSI tape device driver:

- VIOP memory optimization: elimination of dead code and other performance enhancements have resulted in smaller VIOP memory requirements for the VME SCSI tape device driver.
- Adapter recovery: a `reset` command has been added to the `mt` utility for adapter recovery without the need to reboot the machine. A new device node has been added in `/dev/MAKEDEV` for each device with a `spt` suffix. Use of this option is restricted to the superuser. Due to the complexity of many configurations and potential loss of data with this command, please call the HP/Convex Technical Assistance Center (TAC) for detailed reset procedures.
- The VME SCSI tape device driver now allows SCSI id (unit) to be in the range 0 to 6. Each unit can still have up to four subunits (0 - 3).
- Separate channel trace: the internal trace information has been reworked to store separate information by target id, including the adapter. This causes no performance penalty, but provides support with more information where a drive hangs and other drives continue to work. These internal trace buffers can be disabled in situations where VIOP memory is extremely limited.
- Improved error messages with device identifier of the form: `scsi_tape(vme%d/slot%d/unit%d/su%d) error_msg`
- Numerous bug fixes.

Contributed Utilities

ConvexOS V11.5.1 provides several popular contributed utilities for the convenience of the user community. These unsupported products are provided with optional source. All contributed utilities are installed in `/usr/contrib`.

The contributed utilities include:

- **GNU Emacs 19.28.** This is an updated version of the popular Emacs editor with considerable new functionality over the GNU Emacs 18.57 which was shipped with ConvexOS V11.0.
- **GNU zip (gzip) 1.2.4.** This popular utility uses Lempel-Ziv coding (LZ77) to provide compression generally much better than that achieved by LZW (as used in `compress`), Huffman coding (as used in `pack`), or adaptive Huffman coding (`compact`).
- **Lynx 2.5FM.** This is a general purpose distributed information browser for the World Wide Web for users running cursor-addressable, character-cell display devices.
- **tcsh 6.04.** This is an updated version of this popular shell with C-like syntax.
- **GNU xargs 4.1.** The GNU version of `xargs` will build and execute commands lines from standard input.

The new version of GNU Emacs provided in the Contributed Utilities does not remove the previous version of GNU Emacs (located in `/usr/convex`) provided with ConvexOS V11.0.

Kernel panic signature

ConvexOS V11.5.1 provides a kernel panic signature in the event of a kernel error. The panic signature provides a kernel backtrace and a 32-bit hex signature value. In some cases, this may be enough information to allow a problem diagnosis and avoid a crashdump. An example kernel panic signature is shown below.

```
[CPU01@15:51:13] Panic backtrace:
[CPU01@15:51:13] _Panic(78cc2,117114,26cc,1199f3)
[CPU01@15:51:13] _blkfree+0x5f0(1845920,1a83,2000)
[CPU01@15:51:13] _realloccg+0x514(1845920,cd093,cd020,2000,4000)
[CPU01@15:51:13] _bmap+0x844(1845920,3,0,4000,0,bea9b28)
[CPU01@15:51:13] _rwip+0x416(1845920,bea9d3c,1,0)
[CPU01@15:51:13] _ufs_rdwr+0x12c(1845928,bea9d3c,1,0,1ef6274)
[CPU01@15:51:13] _rfs_write+0x4a4(fe3be80,fd53200,fdc6280,bea9e60)
[CPU01@15:51:13] _rfs_dispatch+0x476(bea9e60,fe21100)
[CPU01@15:51:13] _svc_getreq+0x30c(fe21100)
[CPU01@15:51:13] _svc_run+0x78(fe21100)
[CPU01@15:51:13] _nfs_svc+0x342(1ef61dc)
[CPU01@15:51:13] _syscall+0x512(9b,800236fc)
[CPU01@15:51:13] _C2_user_syscall+0x160(...)
[CPU01@15:51:13]        ***  PANIC SIGNATURE: d7fd64ab  ***
[CPU01@15:51:13]      ConvexOS: FATAL ERROR: (ufs,9932)free:freeing free frag
[CPU01@15:51:13]        sp:      0bea9988        a1:      0bea995c
[CPU01@15:51:14]        a2:      00000000        a3:      7f3600dd
[CPU01@15:51:14]        a4:      fffffff7        a5:      001199da
[CPU01@15:51:14]        ap:      0bea999c        fp:      0bea9988
[CPU01@15:51:14]        s0: 0000000000000000    s1: 0000000000000021
[CPU01@15:51:14]        s2: 000000000000df9d    s3: 00000000000000ff
[CPU01@15:51:14]        s4: 00000000000000ff    s5: 000000000000000a
[CPU01@15:51:14]        s6: 0000000000000001    s7: 0000000000000000
[CPU01@15:51:14]        int. mask: 000000ff
[CPU01@15:51:15] syncing disks...
```

Tunables

Several new tunables have been added in ConvexOS V11.5.1 to accomodate new features and performance enhancements. These tunables are listed in Table 5.

Default or maximum values for several ConvexOS V11.5.1 tunables have been changed from previous versions of ConvexOS to better meet the typical resource requirements of the system. These tunables are listed in Table 6.

Some tunables have been recognized as unnecessary in ConvexOS V11.5.1. These tunables have been removed from the operating system and are listed in Table 7.

Note

If the Enhanced MBS System is disabled (`mbs_enabled` set to 0), the FDDI tunables `fd_max_recv` and `fd_max_send` should be tuned to smaller values than the ConvexOS V11.5.1 defaults. Suggested values are the ConvexOS V11.0 defaults of `fd_max_recv = 28` and `fd_max_send = 28`.

Table 5 New ConvexOS V11.5.1 tunables.

Tunable	Description
`bufcache_size`	Size of the buffer cache address space in megabytes. The buffer cache is usually sparsely populated, so the maximum value is limited by the kernel to no more than four times available physical memory. Default = 1024, min = 64, max = 1024.

Table 5 New ConvexOS V11.5.1 tunables.

Tunable	Description
duc_enable	This tunable enables/disables delayed unwire caching. Default = 1, min = 0, max = 0.
elite9_max_out	Maximum number of active requests allowed for an Elite9 QSC disk. Default = 8, min = 1, max = 20.
explicit_table_size inode_table_size file_table_size ncache_table_size dquot_table_size	Allows explicit tuning of table sizes. If explicit_table_size is zero (default), these table sizes are automatically based on physical memory and/or maxusers. If explicit_table_size is nonzero, the tuned values override the computed values. Default = 0, min = 0, max = 1. Default = 1024, min = 512, max = 4096. Default = 1024, min = 512, max = 4096. Default = 1024, min = 512, max = 4096. Default = 1024, min = 512, max = 4096.
halt_on_panic	Halt immediately upon panic. Allows more of machine state to be saved when hitting a panic. Default = 0, min = 0, max = 1.
mbs_enabled	Enables/disables Enhanced MBS system; enabled by default. Default = 1, min = 0, max = 1.
viop_mbs_limit idc_mbs_limit qsc_mbs_limit hippi_mbs_limit	Specifies the number of pages to reserve for an alternate MBS memory pool. Each page represents 64 messages with a maximum of 16 pages or 1023 messages (1 unused). The original MBS pool, which is still used, contains 11 pages or 703 messages. Individual CCUs are disabled and use the original MBS pool when their value is set to zero. Default = 16, min = 0, max = 16. Default = 4, min = 0, max = 16. Default = 16, min = 0, max = 16. Default = 16, min = 0, max = 16.
org_mbs_drain viop_mbs_drain idc_mbs_drain qsc_mbs_drain hippi_mbs_drain	Specifies the number of messages to read from each MBS pool before sequencing to the next pool. This has the effect of increasing a CCUs priority when more messages are read on each pass. Setting the value to zero causes the pool to be drained before sequencing to the next pool. Default = 2, min = 0, max = 10. Default = 2, min = 0, max = 10. Default = 5, min = 0, max = 10. Default = 5, min = 0, max = 10. Default = 5, min = 0, max = 10.
n_async_proc	Specifies number of asiodaemon processes to be started at system boot. Default = 4, min = 0, max = 128.

Table 5 New ConvexOS V11.5.1 tunables.

Tunable	Description
nfsfastsend nfsfastrecv	Enable NFS STREAMS bypass; default to enabled. Default = 1, min = 0, max = 1. Default = 1, min = 0, max = 1.
nfstsize_loopback nfstsize_hippi nfstsize_fddi nfstsize_eth nfstsize_default	Preferred NFS transfer sizes for each interface. Default = 61440, min = 8192, max = 61440. Default = 61440, min = 8192, max = 61440. Default = 8192, min = 8192, max = 61440. Default = 8192, min = 8192, max = 61440. Default = 8192, min = 8192, max = 61440.
tcphippi_default_window fddi_default_window ethernet_default_window	Advertised window scaling for each interface. Default = 2097152, min = 64512, max = 2097152. Default = 64512, min = 31744, max = 262144. Default = 31744, min = 31744, max = 64512.
nfs_ret_mig_stat	Set sticky bit to tell NFS client that a file is migrated. Default = 0, min = 0, max = 1.
qsc_max_io_reqs	Maximum number of outstanding QSC I/O requests per channel. Default = 400, min = 16, max = 400.
sig_cannot_mask	Certain signals (like SIGBUS, SIGSEGV) should not be ignored/masked since they are usually generated by a hardware condition and would immediately be reasserted on return to user code. This value is a bitmask. Default = 0x600, min = 0x0, max = 0x7fffffff.
si_mbs_limit	Defines the maximum total number of MBS messages that the SCSI si driver will attempt to allocate. The default value of this limit is rarely reached, except on systems with multiple hundreds of disk drives. This limit should be lowered if the system begins to panic due to a lack of available MBS messages. Default = 1023, min = 10, max = 4092.
somaxconn	Backlog of listen(). Default = 5, min = 5, max = 20.
st_3480_synch	Use synchronous data transfers for VME SCSI 3480 drives. Default = 1, min = 0, max = 1.
st_3490_read_retry	Number of retries on read errors for VME SCSI 3490E drives. Default = 64, min = 0, max = 127.
st_3490_synch	Use synchronous data transfers for VME SCSI 3490E drives. Default = 1, min = 0, max = 1.
st_3490_write_retry	Number of retries on write errors for VME SCSI 3490E drives. Default = 64, min = 0, max = 127.

Table 5 New ConvexOS V11.5.1 tunables.

Tunable	Description
st_dlt_synch	Use synchronous data transfers for VME SCSI DLT drives. Default = 1, min = 0, max = 1.
st_exb8500_synch	Use synchronous data transfers for VME SCSI Exabyte 85xx drives. Default = 1, min = 0, max = 1.
st_exb8505_112m_bit	Selects autosizing of 112m tape, when tape is not detected as <=54m. This is for VME SCSI Exabyte 85xx drives only. Default = 1, min = 0, max = 1.
st_exb8505_density	Determines density selected for VME SCSI Exabyte 8500/8505 tape drives. The default (st_exb8505_density = 0) Drive Type — rexb0/rexb0n — rexb0i/rexb0ni EXB-8200 (MTD-20A) — 2 GByte — N/A EXB-8500 (MTD-20B) — 2 GByte — 5 GByte EXB-8505 (MTD-20B) — 2 GByte — 5 GByte New Mode (st_exb8505_density = 1) Drive Type — rexb0/rexb0n — rexb0i/rexb0ni EXB-8200 (MTD-20A) — 2 GByte — N/A EXB-8500 (MTD-20B) — 5 GByte — N/A EXB-8505 (MTD-20B) — 5 GByte — 10 GByte Default = 0, min = 0, max = 1.
st_ignore_attach_err	Ignore attach errors for VME SCSI drives, and attach drive anyway. Default = 1, min = 0, max = 1.
st_max_busy_retry	Maximum device busy retries on a VME SCSI device. Default = 30, min = 0, max = 6500.
st_rsp2150_synch	Use synchronous data transfers for VME SCSI Metrum drives. Default = 1, min = 0, max = 1.
st_show_recovered_err	Show recovered errors to the errlog for VME SCSI 3480/3490E drives. Default = 0, min = 0, max = 1.
st_spt_synch	Use synchronous data transfers for VME SCSI SPT devices. Default = 1, min = 0, max = 1.
st_trace_pg_cnt	VME SCSI trace page count in viop 4k pages. Default = 2, min = 0, max = 4.

Table 5 New ConvexOS V11.5.1 tunables.

Tunable	Description
ta_viop_prt_mask	Controls printing of error messages for VME nine-track tape drives. The following bit flags apply to ta_viop_prt_mask: 0x0001 /* prt cmd/esb/fsb on err */ 0x0002 /* prt cmd/csb/esb/fsb/mux1-decode*/ 0x0004 /* prt mux[0-3]-hex */ 0x0008 /* prt mux1 velocity errors */ 0x0010 /* prt chain overruns */ 0x0020 /* prt offline/notready messages */ Default = 0xffbf, min = 0x0, max = 0xffff

Table 6 Existing ConvexOS V11.5.1 tunables assigned new default or maximum values.

Tunable	Description
du_mbs_limit	Defines the maximum total number of MBS messages that the IDC driver will attempt to allocate. The default value of this limit is rarely reached, except on systems with multiple hundreds of disk drives. This limit should be lowered if the system begins to panic due to a lack of available MBS messages. Default value increased from 500 to 1023; maximum value increased from 700 to 4096. Default = 1023, min = 10, max = 4096.
fd_max_recv	Specifies the number of buffers used to hold input packets received by the FDDI driver and ready to be handed to the IP layer. Default value increased from 28 to 128. Default = 128, min = 2, max = 128.
fd_max_xmit	Specifies the number of buffers used to hold output packets handed from the IP layer and ready to be shipped out by the FDDI driver. When these buffers are used up, new output packets handed from the IP layer will be discarded by the FDDI driver. Default value increased from 28 to 64. Default = 64, min = 4, max = 64.
getnewbuf_goal	Defines the number of buffer headers that are freed at one time. Default value increased from 32 to 512. Default = 512, min = 32, max = 4096.
nstbuf	Number of stripe buffers. Maximum value increased from 8192 to 16384. Default = 512, min = 128, max = 16384.
str_ctl_sz	Maximum STREAMS control message size. Default value increased from 16384 bytes to 65536 bytes. Default = 65536, min = 16384, max = 65536.

Table 6 Existing ConvexOS V11.5.1 tunables assigned new default or maximum values.

Tunable	Description
str_dblk_2048 str_dblk_1024 str_dblk_512 str_dblk_128 str_dblk_64 str_dblk_16 str_dblk_4 str_dblk_0	Maximum number of various datablocks for STREAMS. Various new default values changed from 1024. Default = 1280, min = 0, max = 2048. Default = 512, min = 0, max = 2048. Default = 1536, min = 0, max = 2048. Default = 256, min = 0, max = 2048. Default = 1536, min = 0, max = 2048. Default = 128, min = 0, max = 2048. Default = 64, min = 0, max = 2048. Default = 100, min = 0, max = 2048.
str_lo_pct	Percentage at which a BPRI_LO `allocb` will fail. Default value increased from 60% to 80%. Default = 80, min = 0, max = 100.
str_med_pct	Percentage at which a BPRI_MED `allocb` will fail. Default value increased from 80% to 90%. Default = 90, min = 0, max = 100.
str_ctl_sz	Maximum STREAMS control message size (in bytes). Default value increased from 16384 to 65536. Default = 65536, min = 16384, max = 65536.
str_msg_sz	Maximum STREAMS data message size (in bytes). Default value increased from 16384 to 2097152. Default = 2097152, min = 16384, max = 10485760.
tcp_loopback_mtu	Specifies the maximum transmission unit TCP will use when transferring data loopback to the local host. Default value increased from 16384 to 64512. Default = 64512, min = 1000, max = 64512.
udpcksum	Enables check-summing of User Datagram Protocol (UDP) datagrams; now enabled by default. Default = 1, min = 0, max = 1.

Table 7 Deleted ConvexOS V11.5.1 tunables.

Tunable	Description
ifree_panic ifree_rebuild trace_off_on_ifree_rebuild	Tunables provided a workaround for a filesystem bug. These tunables are no longer necessary.
nfs_disable_wc	Force NFS write clustering to be disabled. Use `nfs_enable_wc` to control NFS write clustering.

CONVEX TCP/ HIPPI interface support

Support for the CONVEX TCP/HIPPI interface has been included in ConvexOS V11.5.1. This provides the latest version of the `hippitcp` ccu kernel.

ConvexOS Realtime and Utilities support

ConvexOS Realtime and Utilities support has been integrated into ConvexOS V11.5.1 for the convenience of existing ConvexOS Realtime customers. Users interested in the features of ConvexOS Realtime should refer to Appendix B, Realtime Features User Guide, in the *ConvexOS Realtime and Utilities V11.0 Release Notice*, document number 710-003430-507.

Several ConvexOS Realtime tunables have been added to ConvexOS V11.5.1. These tunables are shown in Table 8.

Table 8 ConvexOS V11.5.1 Realtime tunables.

Tunable	Description
`rt_enabled`	Controls warning messages regarding realtime-specific functions. Warning messages are issued upon entry to realtime-specific functions if this tunable is 0; realtime sites should set this tunable to 1; setting this tunable to 2 will force kernel panics when entering realtime-specific functions. Default = 0, min = 0, max = 2.
`fast_clock`	Specify hardclock frequency. Default = 0 (100 Hz), min = 0 (100 Hz), max = 1 (1000 Hz).
`roundrobin_freq`	Sets the number of hardclock ticks between each roundrobin reschedule. Default = 10, min = 1, max = 1000.
`realtime_grpid`	Specifies a special realtime group. Members of this group will be permitted to execute selected realtime system calls that are otherwise restricted to root access. Default = 0, min = 0, max = 32767.

Changes to utilities and procedures

3

This chapter describes changes to existing utilities and procedures. Changes discussed in this chapter pertain to the following areas:

- ConvexOS and Utilities packaging
- Enhancements to ConvexOS utilities
- Enhancements to Internet Services
- PRTLOG
- Rebuilding NIS maps

Repackaged ConvexOS Utilities

ConvexOS Utilities, previously distributed as root and usr Upgrades prior to ConvexOS V11.0, have been split into 13 separate subproducts that may be released separately. Splitting the Utilities enables CONVEX to quickly develop and release patch or incremental releases. This means there is no longer a single ConvexOS version number associated with the kernel and utilities, but there is a ConvexOS configuration denoted by the subproduct versions saved in a system's Generic Installation Procedures (GIP) database.

Table 9 shows the new subproduct sets.

Table 9 Subproduct sets

Product name	Part number	Description
ConvexOS Accounting System	710-016215-007	Programs and files for accounting administration.
ConvexOS Contributed Utilities	710-017215-006	Contributed utilities, such as GNU EMACS, gzip and tcsh.
ConvexOS Core Utilities	710-000315-011	The basic set of utilities required to install the system. This subproduct contains all the utilities included in the mini-root.
ConvexOS Extended Utilities	710-016115-008	ConvexOS utilities not in the Core Utilities subproduct except for those associated with another specific subproduct.
ConvexOS Programming Tools	710-016315-007	Header files, libraries, and programs to support a development environment.
ConvexOS Info System	710-016415-006	The `infosys`, `learn` and `csh-help` subsystems.
ConvexOS Line Printer System	710-016515-006	Programs and files to support the line printer subsystem.
ConvexOS Mail System	710-016615-008	Programs and files to support the mail subsystem.
ConvexOS Notes System	710-016715-006	Programs and files to support the notes subsystem.
ConvexOS Plotter Control System	710-016815-006	Programs and files explicitly for plotter control.

Table 9 Subproduct sets (continued)

Product name	Part number	Description
ConvexOS Terminal Control System	710-016915-006	Programs and files used explicitly for terminal control.
ConvexOS Text Processing System	710-017015-006	`nroff` and related binaries and files.
ConvexOS UUCP System	710-017115-006	Programs and files to support the UUCP subsystem.

Of the subproducts listed in Table 9, you should at least install:

- ConvexOS Core Utilities, Extended Utilities, and Accounting system.
- ConvexOS Programming Tools, if you have any of the CONVEX compilers.
- ConvexOS Text Processing Tools, if you wish to view the man pages associated with any product you install.
- ConvexOS Domestic Tools, if your site is located within the United States and Canada (sites in other countries will not receive Domestic Tools).

The other subproducts are optional.

Enhancements to ConvexOS utilities

Changes have been made to several ConvexOS utilities:

- `cron`
- `find`
- GNU Emacs
- `login`
- `mkfs`
- `perl`
- `pstat`
- `sendmail`
- `sysstat`

`cron`

The reliability of the `cron` daemon has been improved by eliminating certain circumstances that could cause it to skip jobs for arbitrary periods of time, or simply terminate. The scheduling algorithm has been fundamentally redesigned.

`find`

The `find` utility now supports an HP-UX compatible option, `-fsonly`, that forces `find` to completely ignore other filesystems types. Previously you needed to use a fancy combination of `-xdev` and other options to get the same behavior.

GNU Emacs

The Emacs product is now based on the freely copyable GNU Emacs 19.28, a version which has been proven to be quite stable. It offers considerable new functionality over the version shipped with ConvexOS V11.0, which was based on GNU Emacs 18.57.

The X-windows support was minimal in Emacs 18.57. Windowing capabilities have been greatly expanded in 19.28.

Several packages have been added in Emacs 19.28, including the following:

- GNUS — a UseNet news reader
- Hexl mode — for editing binary files in hex dump format
- Emerge — allows interactively merging two sets of modifications to a single file
- New editing modes — including modes for C++, `awk`, makefiles, `perl`, and SGML

Additionally, there is now a history for the mini-buffer, allowing previous commands to be recalled. Numerous other keybindings and editing capabilities have been added to improve the overall environment. One especially nice feature is the addition of the variable 'next-line-add-newlines', which can be used to disable the annoying addition of blank lines when attempting to move beyond the end of the buffer. There are a few changes which break keybindings which may be familiar to users of older versions of Emacs. If this is a concern on your system, you can continue to run the old version by simply not installing the ConvexOS V11.5.1 version of Emacs.

`login`

All occurrences of the `system` function have been removed from `/bin/login` to prevent environment-based attacks similar to those described in CERT advisory CA-95:14.

`mkfs`

The output from the `mkfs` utility is now appended to the file `/etc/sbtab`, if it exists. This provides a record of the superblocks on any filesystem created after the installation of ConvexOS V11.5.1. An example is shown below:

```
# Super-block log
# The mkfs utility will append its output to this file; this is
# useful when running fsck with an alternate super-block.
#

/dev/rst0: super-block backups (for fsck -b#) at:
 32, 40544, 81056,

/dev/rdu2e: super-block backups (for fsck -b#) at:
 32, 40544, 81056, 121568, 162080, 202592, 243104, 283616,
 322592, 363104, 403616, 444128, 484640, 525152,

/dev/rdu2d: super-block backups (for fsck -b#) at:
 64, 40576, 81088,
```

perl

/usr/bin/perl has been upgraded from version 4.010 to version 4.036, which we feel is the most stable version of Perl that has been released. (Perl 5 has not been included due to compatibility issues.) Perl 4.036 offers numerous bug fixes and compiler enhancements. For example, on a V11.0 system, the command

```
/usr/bin/perl -e 'eval "1 #comment" && print "foo\n"'
```

produces no output. It actually should print "foo", which is indeed what it does on a V11.5.1 system. The compiler in Perl 4.036 is stricter in some ways than the one in 4.010. Certain instances of poor syntax slipped past the older version but are not accepted by the newer one. To insure compatibility with existing scripts, the installation procedure will save a copy of the existing perl binary before installing the new one.

pstat

The pstat utility has been enhanced to include additional system information:

- TCP/IP connection information
- Buffer information
- Stack backtrace
- Statistics on MBLKs, DBLKs, and queues
- All network STREAMS information
- All STREAMS queues
- USER area information
- MBS information

sendmail

The utilities release will include an upgraded version of sendmail, based on Sendmail 8.7.5. This provides numerous improvements over the currently available version, including, but not limited to the following:

- Much improved MIME support
- Accepts SIGHUP to restart the daemon
- Enhanced sendmail.cf readability with descriptive option names, instead of single letters
- Automatic lookup of the addresses of all the network interfaces to discover "alternate" hostnames for the machine

Some security fixes are also included with Sendmail 8.7.5:

- Checks for buffer overflows for DNS/NIS lookups
- Protects against overflowing syslog buffer
- Supports a maximum message size to prevent denial of service attacks
- Tries to validate outgoing messages in the queue to prevent forgery attacks on queued messages.

sysstat

sysstat is a new utility for ConvexOS V11.5.1. sysstat provides system performance statistics (similar to syspic) to standard output.

Enhancements to Internet Services

Several changes have been made to utilities and protocols in CONVEX Internet Services.

bind(2)/connect(2) restriction

A restriction on the `bind(2)` and `connect(2)` system calls has been implemented to resolve a security issue. `connect(2)` must either be called by root or by the same UID that called `bind(2)`.

Denial of Service attacks

There is a denial of service attack that involves spoofing a udp packet from one machine's echo port to another machine's echo port, thus setting up a ping-pong volley of packets that could theoretically saturate the network link between the two machines. The most effective method of preventing the attack is to disable the various services like echo, chargen, date, etc in `inetd.conf`. However, the services can be very useful for legitimate uses, thus the "semi-secure" -s option. The -s option prevents `inetd` from servicing requests on the echo, discard, chargen, daytime and time ports (both udp and tcp) that come from source ports less than 1024. This solution is not fool-proof, but does increase protection against the attack. This option is NOT turned on by default.

Name resolution service configuration

The `/etc/resolv.conf` file now supports a new configuration option, "search" which allows the admin to specify multiple domains to search when attempting to resolve a short (non-fully-qualified) hostname. For example, if `/etc/resolv.conf` contains the lines:

```
domain foo.com
search bar.com
```

and a user wants to telnet to "baz" the resolver will first look up "baz.foo.com" and if that address does not exist, it will try "baz.bar.com."

Only ConvexOS V11.5.1 utilities will be able to take advantage of the 'search' directive. Others that do hostname lookups, such as many of the X-windows programs, will be unable to take advantage of this feature.

The maximum number of searchable domains is 6. Previously, programs that made independent use of the resolver libraries, were limited to a maximum of three domains. This last change should be of little to no consequence for system adminstrators.

Network device entry security

File access permissions on network device entries (/dev/hyper, /dev/eth, /dev/sl, /dev/un, /dev/fddi, /dev/hippi) have been modified to disallow world access. These files have group ownership `kmem` to allow normal users to run `ifconfig`.

SNMP daemon

The `snmpd` daemon now supports a configuration file that lets the administrator specify values for certain MIB variables. The configuration file is: `/usr/etc/snmpd.conf` and the configurable variables are:

```
system.sysContact
system.sysLocation
system.sysDescr
```

Additionally, the community strings that the daemon will honor can be specified in the same file. A sample, innocuous `snmpd.conf` file is shipped. The man page describes the syntax of the file completely.

strstat

The **strstat** utility has been enhanced to give more detailed information about networking buffers. It now displays statistics regarding the number of DBLK allocation requests promoted to the next larger size because they could not be satisfied using the size requested. This provides administrators with the information needed to tune DBLK distribution for maximum efficiency.

tcpdump

Updates have been made to the `tcpdump` utility, which provides the ability to monitor all ethernet network traffic. HiPPI can now be monitored with `tcpdump`, and the traffic filtering capabilities have been significantly improved. The kernel's support for the ethernet driver in promiscuous mode has been made more robust and secure. Promiscuous mode is automatically turned off whenever a process using this mode is terminated. Log entries are also made whenever promiscuous mode is turned on or off. These changes significantly improve the usability of `tcpdump` over ethernet.

PRTLOG

Numerous corrections and enhancements to PRTLOG have been made to improve reliability and provide additional infomation in the `/mnt/errlog` file. Some of the behavioral changes in the PRTLOG logging are listed below:

- A message is logged as to the known cause of PRTLOG termination.
- A large buffer is used to gather write data to minimize time spent doing disk access.
- Race conditions eliminated.
- Every 10th occurrence of a 'repeat' message causes the original error message to be reprinted.
- If a repeated message doesn't occur after 2 minutes, the repeat mode is cleared.
- A 'repeat' message is logged after 30 seconds of active repeats. This was added to make it obvious there is a continuous flow of the same error message.
- When PRTLOG exits, all partial messages and repeat counts will be flushed to the log file.
- If an incoming message isn't completed within 10 seconds, it will be flushed with an 'incomplete' message attached.
- Separate console logger (sniff) is used to help reduce memory requirements.
- When SIGHUP is received the log file is closed and reopened. If the size of the file is in excess of 250,000 bytes and a back up log file does not exist, the log file will be renamed and a new one started. The printf window is

scanned with each SIGHUP. SIGHUP is also sent to the `sniff' utility so sniff will close and open the file too.

- No more spurious INT10 error messages.

Rebuilding NIS maps

NIS maps that were built under ConvexOS V10.1 and contain blank lines must be rebuilt. If the source file for a NIS map (e.g. /etc/services for 'services.byname') has a blank line in it, it will cause problems with `ypserv` when performing look-ups on the NIS map. To avoid this problem, rebuild the NIS maps after installing ConvexOS V11.5.1 by entering the following commands:

```
su root
cd /usr/etc/yp
make clean
make
```

This will take about 2 minutes.

syslog

CERT advisory CA-95:13 documents a vulnerability in the `syslog` routine which potentially makes it possible for local and remote users to execute arbitrary commands. In particular, a remote user could fool `sendmail` into calling this routine with arbitrary data that would overflow the internal buffer. Though exploitation would require in-depth knowledge of the C-Series architecture, `syslog` has now been modified to insure that such an overflow cannot happen.

Changes to ConvexOS kernel 4

This chapter describes changes to the ConvexOS kernel. Changes discussed in this chapter pertain to the following areas:

- Enhanced Message Based System (MBS)
- Per process addressibility
- UIDs and GIDs
- Virtual memory.

Enhanced Message Based System (MBS)

The Message Based System (MBS) is a transport mechanism for passing messages between the CPU, SPU, and CCUs. MBS messages are 64 byte packets containing 16 bytes of header information and 48 bytes of data. In previous versions of ConvexOS, there are 703 MBS messages in a shared memory pool available for use by the CPU, SPU and all CCUs. Each Processor (CPU, CCU, SPU) has a queue header that points to a list of MBS messages on an input queue for that processor. Before an MBS message is placed on an input queue, a free message index must be retrieved from the Free Message Stack.

The Free Message Stack represents a bottleneck in that all processors must lock the Free Message Stack, retrieve a free message index, then unlock the free stack whenever a message is passed to another processor. In ConvexOS V11.5.1, MBS has been redesigned to eliminate the Free Message Stack by implementing queues for every CCU/CPU combination.

The CPU's input message queue represents another performance bottleneck in that every CCU that passes a message to the CPU must first lock the CPU's input message queue, append a message onto the queue, and unlock the queue. The ConvexOS V11.5.1 design addresses this bottleneck by creating a new MBS pool of messages for selected CCU/CPU interfaces. This will dedicate an MBS message pool of up to 512 MBS messages for a single VIOP, IDC, HiPPI, or QSC CCU.

The CPU receives I/O initialization commands from the SPU's `/ioconfig` file and sends them to the appropriate CCUs. When the CPU notices an I/O initialization command to a CCU of interest (VIOP, IDC, HiPPI, QSC), and that CCU's MBS tunable is set to a non-zero value, then the following special processing takes place:

- Memory is allocated on the CPU for an alternate MBS message pool. Alternate CPU and CCU input queue headers are created and initialized.
- A special command is sent to the CCU along with the necessary configuration data for allocating and mapping the CCU windows to the CPU memory.

A tunable is available to selectively enable and disable this MBS scheme on individual CCUs. The tunable specifies the number of pages to be allocated for MBS messages (up to 8 pages - 512 messages), or disables the new MBS system when set to zero.

Note

If the Enhanced MBS System is disabled (`mbs_enabled` set to 0), the FDDI tunables `fd_max_recv` and `fd_max_send` should be tuned to smaller values than the ConvexOS V11.5.1 defaults. Suggested values are the ConvexOS V11.0 defaults of `fd_max_recv = 28` and `fd_max_send = 28`.

The old MBS and new MBS schemes will coexist with one another so every CCU has access to two MBS message pools. The CCU will receive a unique interrupt depending on which message pool has the message. All communication to devices attached to the CCU using the new system will be via the new MBS scheme while SPU communication (e.g. `adbccu`) to the VIOP itself will use the original MBS scheme. The original MBS scheme is also used to send the I/O initialization message to the CPU to relay to CCUs.

Per process addressibility

The maximum size of thread virtual memory has been expanded from 64 MB per thread to 384 MB per thread. This raises the maximum possible thread memory for a process from 512 MB to 3 GB, allowing the creation of a 4.6 GB process.

UIDs and GIDs

The number of possible UIDs and GIDs has been effectively doubled, now allowing 65,000 of each. Both ConvexOS and its utilities have been modified to support this change.

Virtual memory

Virtual memory frequently-used routines have been re-engineered to reduce system time (i.e., system overhead). In addition, various other changes have been made to improve paging performance.

Process initialization

Virtual memory initialization has been sped up by using the system's ability to invalidate the entire Address Translation Unit (ATU) cache. The performance improvement is especially noticeable when initializing a process with a large address space. A 2 GB process, for example, has 500,000 pages, whose ATU cache entries were previously updated on a page-by-page basis. This process was fairly time consuming. Applications that allocate many large arrays (totaling approximately 512 MB or more) will benefit most from this enhancement.

Thread address space paging

The I/O subsystem can read and write in 64 K increments per operation. Therefore, paging speed can be significantly improved by reading or writing multiple pages at once. Page-ins and page-outs of thread-private memory are now clustered to avoid the need to perform repeated operations on the same disk block. This change can speed up paging of thread memory by as much as a factor of sixteen. The

procedure for paging multiple threads has also been fundamentally redesigned, so that it is done on a per-thread basis, rather than across all threads simultaneously.

Thread virtual space allocation

Note

Existing versions of ConvexOS do not accurately track usage of thread virtual memory, making it possible for the kernel to grant virtual memory to processes for which it was unable to actually provide swap space. This situation could lead to serious problems, including the inability to create new processes or a complete hang of the operating system. Due to more accurate tracking in ConvexOS V11.5.1, some applications could require more swap space than when running on an V11.0/V11.1 system.

Shared memory

In earlier versions of ConvexOS, shared memory that was mapped with the mapping attribute MAP_ANON was found to be inefficient when paged in or out. The problem was due to a limitation of the virtual memory subsystem which prevented clustering. In ConvexOS V11.5.1, this limitation has been eliminated. The performance improvement this provides is similar to that for thread memory; i.e. a speed-up of as much as a factor of sixteen. Note that for both these types of clustering, performance will already be non-optimal since swapping is taking place; however, these changes reduce the severity of the performance loss when swapping.

Bug information

5

Fixed bugs

This section lists fixed bugs and their resolutions. Some bugs do not have PR numbers; these are listed as "PR unknown."

Fixes for ConvexOS Utilities

accounting

PR-58669

The accounting system checks to see if it has already been run on any particular day, week, or month, but it doesn't check to see if it is already running, so if it takes longer than the cron interval to complete, then two copies will run at once.

Suggest a lock file with the current pid, that is checked to see if it is still running if it exists.

Resolution: Fixed in ConvexOS V11.5.1. An exclusive file lock on /usr/adm/lastacct is used to prevent running multiple instances of accounting.

acct scripts

PR-50708

The problem is with the weekly11.0 and monthly11.0 scripts and the addition of the periodic and job accounting runs...

They need to re-order the scripts so that the acctsum.awk script is run after sa is run with the amF[w,m] options.

Resolution: Fixed in ConvexOS V11.5.1.

accton.8

PR-41982

The accton(8) man page and accton command "usage" implies one can change both acctfile and period on the same command line, but the /etc/accton command will complain "you can't start accounting and periodic records at the same time".

Bug is reported against "usage" of accton. If this option combination is not possible then the Man page needs to indicate that only one of [acctfile] or [-p period] is allowed.

Resolution: Modify accton.c to show two uses of accton. One usage has optional acctfile. (If no acctfile, accounting is stopped.) The other usage specifies the period with a -p command-line option.

```
# accton -x
accton: illegal option -- x
Usage: accton [acctfile]
       accton -p period
```

Also, adjust the sa(8) manual page to give the following output:

```
SA(8)                    ConvexOS Man Pages                        SA(8)
NAME
     sa, accton - system accounting
SYNOPSIS
     /etc/sa [ -abcdDefFgijJkKlmnpPrstuv ] [ acctfile ]
     /etc/accton [ acctfile ]
     /etc/accton -p period
```

adb

PR-41083

/bin/adb incorrectly uses a mixture of read() and gets() to read its input. Specifically, the $g command in adb uses libc/gets to read the user's response. This is the only place that uses the stdio package. If you redirect adb's input from a non-tty source and your input contains a $g, gets() will read an entire "block" into the stdio buffer. When adb goes to read its next command via the "read", it incorrectly encounters EOF.

Resolution: Fixed in ConvexOS V11.5.1.

autocron

PR-62604

autocron consumes file descriptors. It opens the cron spool directory with opendir but fails to use closedir. This can occur repeatedly as autocron operates. Also, autocron can open a socket but fail to connect, but leave the socket opened.

Resolution: Fixed in ConvexOS V11.5.1.

avail

PR-27967

Running /usr/spool/convex/avail -m avail from the crontab file causes the following output:

```
spu file /mnt/usr/lib/softlog: No such file or directory
```

In the file /usr/spool/convex/changes is the following command /usr/convex/spu -r /mnt/usr/lib/softlog > /usr/spool/convex/tmp.softlog, causing the above error. With the upgrade of software on the spu, this file has been moved to /mnt/softlog.

Resolution: Updated /usr/spool/convex/changes for the following items:

1) add "c46" as a supported CPU type

2) correct handling of softlog information.

bc

PR-33793

The sign of the output of a logarithm of a number less than 1 is wrong. The sign is always positive.

Resolution: Fixed in ConvexOS V11.5.1. Problem caused by bug in dc.

bc

PR-40750

The result of negative sign argument of sin() is always positive.

```
bc -l
s(-1.)
.84147098480789650665
```

Resolution: Fixed in ConvexOS V11.5.1. Problem caused by bug in dc.

cbc_crypt

PR-53387

Although there is a man page for "des_crypt" which describes use of cbc_crypt and ebc_crypt, we were unable to find those function calls in an library.

cbc_crypt and ebc_crypt were left out of the DES.a library when it was built for the domestic utilities product for ConvexOS V11.0.

Resolution: There is no man page for these functions in ConvexOS V11.5.1. Sites should use crypt(), setkey(), and encrypt(), rather than ecb_crypt(), cbc_crypt(), and des_setparity().

chsh

PR unknown

"chsh" execs "mkpasswd -p" which changes the permission bits on "/etc/shadow.dir" and "/etc/shadow.pag" to "-rw-r--r--". This is considered a security risk and needs to be addressed.

Resolution: Fixed in ConvexOS V11.5.1.

chsh now sets permissions on /etc/shadow.{dir,pag} to 0600 rather than 0644 to resolve a security problem.

chsh now execs "mkpasswd -s" when rebuilding the shadow files rather than "mkpasswd -p".

compress

PR-44167

The compression algorithm reports erroneous compression factors when the "-v" option is used:

```
(convex) # compress -v u.w.9.26.93
u.w.9.26.93: Compression: 101.84% -- replaced with u.w.9.26.93.Z
                          ^^^^^
```

Resolution: Simplify reporting of compression ratio, to avoid reporting more than 100%.

contact

PR-38645

The cpu message generated when contact starts is confusing. The changes proposed are to change the initial message to:

```
Welcome to contact version 0.24 (93/07/19)
Generating contact report for host mikey (serial number 8401).
```

If the :ac: field is set it should continue with this question:

```
Would you like to specify a different machine (yes | no)?
```

Resolution: Changed print statements as requested. contact will always print the line "Generating contact report for host yourname (serial number 00000)." If the :ac: flag or if the serial number corresponds to a Convex internal system (determined by a hard-coded list of system serial numbers in the code), the user will be prompted "Would you like to specify a different machine (yes | no)?"

contactcap

PR-58578

Here is a valid network based contactcap that we should distribute:

```
c0|contact|contact site configuration:\
     :ph=(800) 952-0379:ed=/usr/ucb/vi:mb=contact@convex.com:\
     :fb=contactfile@convex.com:cc=%s:uu:ta:pa=:ml#50000:ul#50000:
```

Resolution: Fixed in ConvexOS V11.5.1.

cron

PR-54413

PR-54844

British summer time started on Sunday March 26th. Between 23:00 and 00:00 on Saturday 25th, cron looked to be finding it impossible to do time calculations:

```
Mar 25 23:45:56 neptune cron[1277]: (root) /usr/lib/atrun
Mar 25 23:46:40 neptune cron[1277]: time computation error, event '/
usr/lib/atrun' for 'root' not queued
```

atrun is cron scheduled for 15 minute intervals and didn't start running again until 02:15 BST.

Resolution: Completely rewrote next_time function.

csh

PR unknown

Commands executed in a subshell have extraneous file-descriptors left open.

Resolution: Fixed in ConvexOS V11.5.1. Some temporary file descriptors weren't being marked as close-on-exec.

csh

PR unknown

If a file is sourced interactively and the user suspends a process run by this file, the shell will never print out a prompt, hanging the session. For example, run the following commands under csh:

```
echo 'vi /tmp/foo' > /tmp/bar
source /tmp/bar
```

then suspend the vi session. A shell prompt will never be printed out, and there is no way to return to the vi session.

Resolution: Corrected the way we check whether there is a tty associated with us, so that we no longer hang in this situation.

csh, ksh

PR-38144

If csh or ksh is used to run a process which exceeds a memory limit, and the process takes the default action for SIGINFO (which is death) the shell prints "Cputime limit exceeded".

Resolution: Error message updated in the tcsh 6.04 port.

dc

PR-55822

First Example:

```
#!/bin/csh -f
echo ''
bc -l <<!
a(-1.35519928311407713317)
!
echo ''
bc -l <<!
a(-1.35519928311407713319)
!
echo ''
bc -l <<!
a(-1.35519928311407713318)
!
echo 'We need to type ^C to abort'
echo 'Or wrong results'
```

Second Example:

```
#!/bin/csh -f
echo ''
bc -l <<!
s(-2*a(1))
!
echo 'instead  -1 '
echo ''
dc <<!
0k
_1.00 1/p
!
echo 'instead -1 '
```

Resolution: Remove previous changes involving variable "neg" in source code. These changes cause the change in behavior relative to other vendors' bc/dc and cause the problem with dividing -1.00 by 1 at scale 0.

Problems with bc--such as s(-1), c(2), l(0.5) all being positive instead of negative--are resolved by this change to dc.

df

PR-44023

'df' fails on UniTree controlled NFS filesystems.

Resolution: Fixed in ConvexOS V11.5.1.

diskuse

PR-40762

diskuse utility needs to be supported with "large files". If we support Accounting, then we should support the utilities to generate the necessary data as documented in the ConvexOS Operations Guide. Please refer to posted Comments #1 for examples on how diskuse breaks with large files.

Resolution: Fixed in ConvexOS V11.5.1.

diskuse

PR-43505

diskuse terminates with exit code 1 in the following situation: testlink is a symlink to a file on a remote filesystem, which is not readable for root. If diskuse checks testlink, it reports:

```
diskuse: cannot statfs "..../testlink": Permission denied
```

and terminates with exit code 1.

Some of our users have such links since last week. Our disk accounting don't work anymore (and we need it).There are two possibilities for this bug:

1. diskuse uses stat but should use lstat (this would be totally wrong)

2. diskuse uses lstat but then also stats the file behind (but why?)

Resolution: Added check of the return code of statfs (not stat) for ENOACCES. Prints error message, but no longer terminates.

du

PR-45142

The du command seems to be limited to 1000 entries per directory and returns garbage a when this limit is exceeded. This limit is far too low in a file server environment. See the size of the ml table and the value (1000) of ML in du.c. This limitation should at least be documented. A good program would make use of dynamic memory here.

Resolution: Fixed in ConvexOS V11.5.1.

du was limited to a max of 1000 different files with multiple hard links. After 1000 different files, du starts counting each link as an individual file, sometimes with outrageous results.

Solution was to dynamically allocate space for counting hard links, now done in chunks of 1000 files.

dump

PR unknown

FileServ dump/restore should behave like ConvexOS dump/restore. If a write error occurs, dump should mount another tape and begin that portion of the dump again. (FileServ dump/restore will be based on ConvexOS dump/restore in ConvexOS V11.0).

Resolution: fsdump uses ConvexOS /etc/dump in ConvexOS V11.0 and V11.5.1 which provides the desired feature.

dump

PR-41304

When performing dumps of migration filesystems, prefer that an option be available to dump only inodes of files. Dumping data blocks for large filesystems can be very time consuming and resource intensive. As a sys admin, I configure the directories on such filesystems so that all files on the filesystem are subject to migration; hence, all files will eventually be written to tape within the fileserv environment. It is overkill to have yet another tape copy of a file. One alternative is to truncate.

Resolution: Fixed in ConvexOS V11.5.1.

dump

PR-41989

Attempted to dump a filesystem that contained files that were migrated to tape using FileServ v2.2. Dump says:

```
151.loki:/export/fs2.2/fileserv>dump  0GIEf /dev/rmt20 /daemon
  DUMP: /daemon filesystem is mounted
  DUMP: Dumps on mounted filesystems can be unreliable
  DUMP: Filesystem is under control of a migration dmon,
migrated files cannot be dumped.
  DUMP: NEEDS ATTENTION: Do you want to continue?: ("yes" or "no")
```

Appears that ConvexOS V10.1 patch FS2.1.128 was not rolled up into ConvexOS V11.0 release .

Resolution: Fixed in ConvexOS V11.5.1.

dump

PR-46486

Doing "dump 0GI " the 9track tape uses 1600bpi instead of default 6250bpi. Giving /dev/rmt16 with it uses 6250 bpi. The dump =0GI should work as in the man-pages mentioned.

dumpmain.c - when passing G option which sets tape=/dev/rmt16, we later regress to /dev/rmt8 if the device is not alloc'd in tpq of tpdaemon, because a -f /devicename wasn't passed. Changed code to keep tape=/dev/rmt16 when -G passed, unless they passed in a -f device.

xdump.c is busted, too.

Resolution: Fixed in ConvexOS V11.5.1.

emacs

PR-58642

We have been using emacs 19.X for a long time within Convex without problems. We should update the version shipped with the ConvexOS V11.5.1 release. I can supply Emacs 19.28 or we can get the latest.

It builds right out of the box for ConvexOS.

Resolution: The 'contrib' product in ConvexOS V11.5.1 includes Emacs 19.28.

exports(5)

PR-41539

```
root=client[:client]...
     Give root access only to the root
     users from a specified client.  The
     default is for no hosts to be
     granted root access.  A client can
     either be a hostname, or a netgroup
     (see netgroup(5)).  If a netgroup
     is used and a change is made to the
     netgroup, the directory must be
     exported again usin
```

Resolution: Corrected manual page to read:

```
root=client[:client]...
     Give root access only to the root
     users from a specified client.  The
     default is for no hosts to be
     granted root access.  A client can
     either be a hostname, or a netgroup
     (see netgroup(5)).  If a netgroup
     is used and a change is made to the
     netgroup, the directory must be
     exported again using exportfs(8) in
     order for the change to be effec-
     tive.  A maximum of 200 clients can
     be specified for the 'root' option.
```

find

PR-38913

If find(1) is executed with "/" given as the search path and the user doesn't have permission to a top-level subdirectory of "/", find exits with: find: bad directory tree". This is due to an error in the logic of find'a descend() function; it nulls out the search pathname when an attempt to descend into such a subdirectory fails. The following commands can be executed to reproduce the problem.

```
% su root
# mkdir /find-test
# chmod 700 /find-test
# ^D
% find / -name xyzzy -xdev -print
```

Resolution: Reproduced problem under ConvexOS V10.1. Problem is not present in ConvexOS V11.5.1. Problem was resolved in a previous OS release.

find

PR-46268

In the routine descend a test on the (NULL) return of rindex is missing. This causes find to core dump when the path is a symbolic link without a slash.

Resolution: Fixed in ConvexOS V11.5.1.

find

PR-46556

find with -fstype 4.2 option should only use 4.2 systems, but it goes into NFS systems too.

Resolution: Fixed in ConvexOS V11.5.1. Added -fsonly option which behaves the way the customer wants.

ftp

PR unknown

The ftp client buffers the printing of hash marks so that a couple of hundred Kbytes need to be transferred before a single hash mark is printed.

Resolution: This isn't because of buffering; it's an artifact of the way ftp is implemented. There isn't any real good way to fix it.

getsysinfo

PR-44045

The man page for getsysinfo (section 2) shows that getsysinfo should be a pointer. It should really be an actual system_information structure. In the file /usr/man/man2/getsysinfo.2, the line:

```
struct      system_information  *sysinfo;
```

should be changed to:

```
struct      system_information  sysinfo;
```

Resolution: The man page is correct.

getsysinfo

PR-50226

getsysinfo -l lists num_vregs as one of the features getsysinfo knows about:

```
saturn5 % getsysinfo -l
```

```
getsysinfo: Note that you're using -l.  All known features will be
listed
        without regard to the configuration of this machine.

ieee_supported
ieee_default
intrinsics
parallel
op_under_mask
scalar_acc
secure_ucode
ex_vregs
native_default
system_sn 46005
cpu_type c46
cpu_count 3
mem_interleave_factor 128
num_vregs 16
```

but then getsysinfo fails when I try to use it:

```
saturn5 % getsysinfo -f num_vregs
getsysinfo: "num_vregs" is an unknown feature
```

Resolution: '-f num_vregs' can now be specified; on non-C4's, this returns 8, on C4's, it returns 16.

ifconfig

PR unknown

With the 11.0.181 streams security patch, ifconfig doesn't report interface configuration to normal users anymore.

Ifconfig needs to be setgid kmem and modified to only print configuration information if the user is not root.

Resolution: Fixed in ConvexOS V11.5.1.

ifconfig

PR unknown

Bug in ifconfig that prints the wrong MAC address for fddi interfaces and prints bogus MAC addresses for hippi/hyper/slip/ultra interfaces. (There should be no MAC address for such interfaces)

Resolution: Fixed in ConvexOS V11.5.1.

inetd

PR unknown

Inetd is vulnerable to udp spoofs between two machine's echo/chargen/etc ports.

Resolution: Fixed in ConvexOS V11.5.1. Added the -s (for semi-secure) option to make inetd ignore connections from ports under 1024.

install

PR-27277

The installation of ConvexOS overwrites files in /usr/skel. The installation should be smarter about installing files in /usr/skel. Customers' changes to these files are lost during an upgrade.

Resolution: ConvexOS V11.5.1 installation of core utilities saves the following files:

- usr/skel/.cshrc
- usr/skel/.exrc
- usr/skel/.login
- usr/skel/.logout

ConvexOS V11.5.1 installation of extended utilities saves the following file:

- usr/skel/.project

This was also done in ConvexOS V11.0.

install

PR-27860

PR-28753

ConvexOS V10.1 Installation Procedures lend confusion for DAT tape upgrades.

Resolution: ConvexOS V11.0 and V11.5.1 do not support kernel upgrades from the SPU. ConvexOS V11.5.1 is only available on 9-track and DAT installsw format tapes.

install notes

PR-42812

PR-43777

PR-46490

The install notes do not mention the fact that the files: /usr/adm/savacct and /usr/adm/usracct (summary files) must be removedprior to running sa on the upgraded system or it will core dump.

Resolution: ConvexOS V11.5.1 installation notes request the customer to close out accounting before upgrading the system. The ConvexOS V11.5.1 installation procedure moves /usr/adm/savacct and /usr/adm/usracct to <file>10.1.sav to avoid this problem.

libcurses.a

PR-51418

The functions mvprintw() and mvwprintw() in mvprintw.c call va_start with the arguments in the wrong order. This causes the format string passed to mvprintw() and mvwprintw() to be overwritten. See comments for details.

Resolution: Reversed incorrect order of arguments to va_start().

login

PR-45512

If a password has expired and domainname is set, /bin/login runs yppasswd instead of passwd even if the user is not in the NIS passwd map. This is wrong because local users can have entries in /etc/passwd even though they are not in the NIS passwd map. /bin/login needs to do a ypmatch to determine if yppasswd should be run.

Resolution: /bin/login now does a ypmatch to determine if it should run /usr/bin/yppasswd or /bin/passwd.

lpc

PR-45167

If 'lpc' can no longer resolve a machine name for remote printer, one will get an un-informative "Bus error" message.

Resolution: In use_local, if we can't get the hostname for the remote machine, print an error message and return 0.

lpc(8)

PR-42626

According to the man page, the st directive in /etc/printcap indicates the name of the status file. What it does not tell you is the value of the st directive is concatenated with the path of the queue.

Resolution: "lo" and "st" directives in /etc/printcap specify filenames relative to the "sd" spool directory. See comments for source code.

Modify manual page for printcap(5) to indicate that both "lo" and "st" are relative to the spool directory.

lpq

PR-21333

Specifying the -l option (long display) on lpq more than once for remote printers results in misleading error messages.

Resolution: Fixed in ConvexOS V11.5.1.

magic

PR-59771

Since we see a lot of PARISC binaries these days, we should merge our /etc/magic with the SPP /etc/magic (or gnu's /etc/magic) to make it a little smarter.

Resolution: Added SPP-UX magic entries to end of ConvexOS magic file, with the following exceptions:

The following entry was adapted to work with the ConvexOS archive entry:

```
0       long    0x213c6172      archive file
>68     long    0x020b0619      -s800 relocatable library
```

The following entries duplicate or conflict with existing ConvexOS entries:

```
0       short   0x71c7          cpio archive
0       long    0x0000ff65      old archive
0       long    0x3c61723e      VAX 5.0 archive
0       string  %!PS-Adobe-     postscript file
>11     string  1.0             -version %s
>11     string  2.0             -version %s
>15     string  EPS             -type %s
>15     string  Query           -type %s
>15     string  ExitServer      -type %s
0       string  %!              postscript file
```

mail

PR-7909

PR-8033

PR-14550

PR-25792

PR-58891

/usr/ucb/mail does not correctly parse RFC-822 addresses. In particular, it thinks space is a delimiter on incoming addresses, which is incorrect.

Resolution: Fixed in ConvexOS V11.5.1.

mail

PR unknown

An option for mail to escape lines beginning with "From " by inserting a ">" character is desired.

Resolution: ConvexOS V11.5.1 includes an option "escfrom" which enables the desired behavior.

mail

PR unknown

/usr/lib/conf/sendmail/convex/bin/smrestart does not regenerate /usr/lib/sendmail.fc

Problem is cured by replacing

```
system("smfreeze");
```

with

```
system("$smfreeze");
```

Resolution: Fixed typo and generally cleaned up smrestart, also taught smfreeze to alert people that they don't need to use smfreeze any more.

MAKEDEV

PR unknown

Comments such as "rsp* MTC-209" at the top of /dev/MAKEDEV are incorrect.

Resolution: MAKEDEV comments updated.

man pages

PR unknown

Typo in the setuid(3) man page:

```
int   setuid(uit_t uid)
               ^
```

Resolution: Fixed in ConvexOS V11.5.1.

mkfs

PR-59760

On the SPP, when mkfs runs it appends its output to /etc/sbtab. This makes a running list of filesystems built and where their alternate superblocks are. This should be pretty easy to do for ConvexOS as well.

Resolution: Adapted SPP code for logging super-block backups to ConvexOS version. If /etc/sbtab exists, mkfs will append the listing of super-block backups for the filesystem. See the first comment for example output. (If /etc/sbtab does not exist; no output is logged.)

Added following item to manual page mkfs.8 under the "NOTES" heading:

```
If the file /etc/sbtab exists, mkfs will append the list of
super-block backups (for fsck -b#) for each file system
created.
```

mpa

PR-62540

mpa core dumps if the PATH environment variable is not set. This is because it passes the pointer returned by getenv() to strlen() without checking to see if getenv() returned NULL.

Resolution: Fixed in ConvexOS V11.5.1. Code checks for null PATH; if PATH is null, do not cycle through PATH elements looking for executable program.

mvst(8)

PR-50870

"mvst" stats raw disk device files of failed partitions when it runs to determine their associated device major/minor numbers. "mvst(8)" should document that these device files must exist for "mvst" to function properly.

Resolution: Fixed in ConvexOS V11.5.1.

nm

PR-63809

"nm" fails to close any of the files it processes, so it tends to run out of file descriptors if it is given a large number of files to process (such as via 'nm *.o').

Resolution: Add fclose statements in nm.c in routines namelist() and read_archive() to handle two types of problems.

od

PR-42774

The makefile was modified to use -ext instead of -pcc, but od.c still has ifdefs for "convex" instead of "__convex__", this caused it to dump core because there was missing code.

Resolution: Fixed in ConvexOS V11.5.1.

op

PR-23344

PR-64737

/etc/op "bus errors" when given a long path name in the arg list.

Resolution: Fixed in ConvexOS V11.5.1.

perl

PR-46159

When running the following Perl command shown in the Repeat-By field, the error message "Can't reswap uid and euid" is displayed and the script is terminated.

Resolution: This is an artifact of the compiler options that we are currently using in the compile and link phases of building Perl. We are currently relying on this behaviour to avoid a security hole that was documented in CERT advisory CA-96.12, so we have no intention of changing the way we are building it.

pstat

PR unknown

pstat can be used to access files that regular users would not otherwise have access to. For example (as root):

```
cp /vmunix /tmp/vmunix
chown root.kmem /tmp/vmunix
chmod 660 /tmp/vmunix
```

(as a regular user):

```
pstat -p /tmp/vmunix /dev/mem
```

Resolution: If any files are specified on the command line, pstat now gives up its membership in group kmem.

rc

PR-45367

/etc/rc has the following line:

```
/etc/umount -at 4.2 | 2>&1 /bin/grep -v "not mounted"
```

It appears that it should remove errors from the umount that have "not mounted" in them, but it won't because the "2>&1" redirect is on the wrong side of the pipe, so it's redirecting grep's output, not umount's. The line should read:

```
/etc/umount -at 4.2 2>&1 | /bin/grep -v "not mounted"
```

Resolution: Fixed incorrect redirection of umount command

rcp

PR-58604

Rcp opens /dev/ip to get a list of network interfaces. After doing the ioctl it nevers closes the file, so if rcp hangs, one of the available stream links to /dev/ip is lost.

Resolution: If a local read error occurs, break out of the read/write for loop.

rdump

PR-59100

rdump doesn't support the 'S' option which sets the dump to unattended and specifies a script to run when dump needs to advance to the next tape. From what I saw from the source code all the flags regarding 'S' are in there. But it doesn't work. I've tested it myself. The site need this fixed immediately to do remote dumps between two C3800s.

Resolution: Fixed in ConvexOS V11.5.1.

restore

PR-62674

PR-64125

A restore onto the C partition of a disk that was just newfs'd is giving some error messages. on a directory called 'sockets' it's giving:

```
cannot create special file, too many open files.
```

on some regular files it just says:

```
too many open files.
```

restore should close the socket descriptor that's created in the makesoc() function in tape.c

Resolution: Added close(s) after bind to socket (for both success and failure) to be certain that the file descriptor is released.

sa

PR unknown

"sa -mP" reports incorrect cpu times. The problem is due to the following code segment in printpermoney():

```
            x = ptr->ap.PROC.ac_nutime.tv_sec + ptr-
>ap.PROC.ac_nstime.tv_sec +
```

```
>>>             (double) ptr->ap.PROC.ac_nutime.tv_usec +
                (ptr->ap.PROC.ac_nstime.tv_usec)/1000000;
```

The indicated line should be:

```
                ((double) ptr->ap.PROC.ac_nutime.tv_usec)/1000000 +
```

Resolution: Fixed in ConvexOS V11.5.1.

sed

PR-59150

not all sed commands allow for a seperation character other than'/'. ie, you cannot say:

```
sed -e '@^foo /tmp/xyz@r /home/bar
```

you can only use:

```
sed -e '/^foo \/tmp\/xyz/r /home/bar
```

all other vendor's have fixed their sed's (and the workaround

for this PR is to use gnu sed) so we should fix ours as well.

Should we use Gnu sed instead?

Resolution: Modified *address() to match HP-UX behavior. Same behavior is seen on SGI and Sun. Copied documentation from HP-UX for this feature and added to manual page.

seekdir,telldir

PR-56563

PR-61398

seekdir and telldir have problems with resolution that cause utilities that use them (du, pax, tar, cpio) to loop forever on directories with more than ~6000 files.

Resolution: Fixed in ConvexOS V11.5.1.

sh

PR-44866

'test' function in bourne shell doesn't handle the -h file operator. It produces the following error message...

```
$ [ -h /tmp ]
test: argument expected
```

The external test program /bin/test supports the -h operator. Since the manpage for sh refers the reader to the test(1) manpage for information regarding the builtin test, the builtin should behave like the external program.

Resolution: Added '-h' as an alias for '-l'.

sh

PR-41260

When running the following test script shell gives the following error

```
test3: 3568 Memory fault - core dumped
```

This does not happen with the Sun or the HP. It happens on ConvexOS V10.1 shell and V11.0 shell. Sample input is in comments.

```
test script
#!/bin/sh -x
while read class group
do
        echo $group | sed "s/^/$class/"
        done  < dom.classes.new >> output
```

Resolution: Properly handle jobs which were killed due to signals greater than 19. Don't do anything when stakchk() is called, to avoid an invalid memory access.

sh

PR-45309

The Bourne shell does not properly handle jobs which die due to signals greater than 19. The shell dies silently with an exit value of 0.

For instance, if a process is killed due to SIGXCPU, the shell will exit with status 0, without properly reporting the exit status of the killed process.

Resolution: Properly handle jobs which were killed due to signals greater than 19. Don't do anything when stakchk() is called, to avoid an invalid memory access.

sh

PR-46489

Using the sh and doing a cd /sys you will see with pwd the dir of usr/sys. Then doing cd .. and pwd you get "/" but ll shows you the dir of /usr. This mismatch of paths could be dangerous for root doing e.g rm - functions in the wrong directory.

Resolution: Disabled built-in 'pwd' command.

sh

PR-53644

the following result is from ConvexOS V11.0

```
% /bin/sh
   $ foo || echo doesnt work
   foo: not found
   $
```

While every other architecture (BSDI, SGI, SunOS, HPUX)

does the right thing:

```
   % /bin/sh
   $ foo || echo this works
   foo: not found
   this works
   $
```

Resolution: Modify xec.c to not perform longjmp to reset shell if the command is not found. Perform same error notification, set exitval to ERROR, and allow command execution to complete "normally" through routine execute.

stat

PR-54273

PR-55346

PR-56724

/etc/stat (extended/src/stat) does not collect networking status. It also breaks the record format with an extra place holder write_byte(): write_byte(filep, "***** ERROR NETWORKING IS BROKEN *****", (long) (in/dt), 1);

Resolution: Fixed /etc/stat to report networking information. Also added seestat.pl script to allow user see the output from files in /usr/adm/stat/

syspic

PR-39309

When just typing syspic under a C38xx running ConvexOS V10.2 you'll not see anything displayed in the box called 'CCU Busy' that gives you the percentage of CCU activity.

Resolution: Fixed in ConvexOS V11.5.1.

syspic

PR-34774

When using syspic the I/O of an installed METRUM drive does not show up in the "Tape" window.

Resolution: Fixed in ConvexOS V11.5.1.

syspic

PR-53914

PR-54119

syspic disk page always shows 0% ccu utilization. xstatd always gets 0% ccu utilization from kernel.

Resolution: Fixed in ConvexOS V11.5.1.

syspic

PR-57867

syspic doesn't display info on quad scsi disks but it does show all other system activity.

Resolution: Fixed MAKEDEV so that -ctl device of sd disk group is owned and readable by kmem. Existing devices will have to have their permissions and ownership changed manually.

syspic

PR-60005

syspic should ignore not configured network interface instead of exit with error message.

Resolution: syspic now ignores not configured network interface.

syspic

PR-65099

Syspic doesn't sum up the stripe I/O rates and only shows a sum of all non-striped filesystems. This bug exists since ConvexOS V10.0.

Resolution: Cannot be correctly fixed since two stripes which share one disk can be displayed in two different windows. Change the code to add the stripes in the sum and also document it in man page.

Change disposition to RESTRICTION.

tar

PR-42098

PR-52765

The group on special files is not being restored correctly (eg tar up /dev/[k]mem). This was reported for general files in a ConvexOS V11.0 beta & appears to have been fixed only for regular files.

Resolution: Only strip off setgid bit if user is not root.

tar.1

PR-45596

The -C option on the tar man page includes the example:

```
For example, to archive files from /usr/include and from /etc, one
might use:
                        tar c -C /usr include -C /etc .
```

The resulting tarfile would have entries like:

include/errno.h

...

./inetd.conf

This is not very useful and does not clearly identify the use of the -C flag. Untarring the files into a useful place would be difficult. Changing the example to something like this makes more sense:

```
tar c -C /usr include -C / etc
```

Resolution: Fixed in ConvexOS V11.5.1

tcsh

PR-53279

Having just installed patch coreutil/11.0.132.2 I found that the man1 file csh.1 had an error which prevented "man csh" from working. The file contained .so tcsh.1 instead of .so man1/tcsh.1

I would also point out that the tcsh man page could be better worded as a joint csh/tcsh man page. Currently there is no indication that they are the same - indeed it says:

```
The NEW FEATURES section describes major enhancements of
tcsh over csh(1) and explains the `(+)' and `(u)' markings
```

Resolution: Fixed in ConvexOS V11.5.1.

tcsh

PR unknown

the csh delivered with patch coreutil 11.0.130.2 set the variable shell to /usr/local/bin/tcsh. This should be /bin/csh or /bin/tcsh.

We worked around this by setting a symbolic link from /usr/local/bin/tcsh to /bin/tcsh

Resolution: Fixed in ConvexOS V11.5.1.

tellcron

PR-42702

tellcron allows you to delete or replace another user's crontab entry (including root's).

Resolution: Fixed in ConvexOS V11.5.1.

touch

PR-39007

If you try to touch (create) a file with name 00010001 with touch 00010001, the file is not created. touch just returns, without an error message.

It looks like touch is first checlking for the number of args to be > 1 (if you count argv[0] as well) and then checks if the first arg (argv[1]) is a number. This number is then interpreted as date and then a for loop is run for all remaining args (but there are none)

This behaviour is wrong.

a) 00010001 is NOT a valid date.

b) if there is only one arg given,

touch 001002 no workie also ;-)

Resolution: Fixed in ConvexOS V11.5.1.

uname

PR-36646

I would like to see the serial number of the machine reported via 'uname'. HP-UX's uname reports a machine id, as does IBM's AIX. The information is in the "utsname" structure returned by uname(2), it shouldn't be a big deal to add a switch to the uname utility to report this value.

Resolution: In ConvexOS V11.5.1, "/usr/bin/uname -a" returns the serial number of the machine. "/usr/bin/uname -c" prints the serial number.

```
> /usr/bin/uname -a
ConvexOS mikey C220 11.5 convex 8401
> uname -c
8401
```

uname

PR unknown

The uname that comes with ConvexOS doesn't support the "-c" flag. The version that comes with the cluster tools does. We should probably just include that version with the OS. (-c reports the system id, or serial number)

Resolution: 'uname -c' or 'uname -i' will report serial number. 'uname -a' includes this as well.

vipw

PR unknown

vipw creates tmp files with mode 0644. This makes shadow passwords pretty useless if someone can read them when vipw is being run.

Resolution: Fixed in ConvexOS V11.5.1.

vmstat

PR-55272

The utility seems to use fixed field widths. This leads to missing field separators (blanks) if values greater than this fixed field width are to be displayed. Any utilities evaluating the ``vmstat'' output (perl, awk, ...) will run into problems.

Resolution: Fixed in ConvexOS V11.5.1. Numbers are always preceded by a space and a little more care has been taken in allocating column widths so that when common numbers get real big, they still won't shift the display out of alignment with the column headings. However, scripts that depend on the physical location of the columns may encounter problems, best to use whitespace delimitation instead.

xdump

PR unknown

Running XDUMP from cron to dump filesystems on Metrum the logfile (standard out and stderror) bring up the following warnings/errors.

```
XDUMP: Couldn't get a valid time from entry `/dev/rdu0a       1 Sat
Jun 17 16:45:02 1995' of dumpdates
```

The XDUMP runs with the following options

```
"/etc/xdump 0uGEf /dev/rvt0n / -b 256"
```

If you take a look into /etc/dumpdates, everything is ok.

```
/dev/rdu0a        1 Sat Jun 17 16:45:02 1995
/dev/rst0         1 Sat Jun 17 15:45:48 1995
/dev/rst3         1 Mon Oct 16 23:45:24 1995
/dev/rst1         1 Sat Jun 17 12:47:57 1995
/dev/rdu0a        0 Fri Jul 14 19:36:02 1995
/dev/rst0         0 Fri Jul 14 18:37:19 1995
/dev/rst2         0 Fri Jul 14 17:45:47 1995
/dev/rst3         0 Fri Jul 14 16:57:16 1995
/dev/rst1         0 Fri Jul 14 16:00:12 1995
/dev/rst2         1 Sat Jun 24 05:45:15 1995
```

This occurs when the spu tunables for timezone differ from the zic info on the JP side. The date conversion function was still using the spu tunable info rather than the zic stuff for parsing the dumpdates file.

Resolution: Fixed in ConvexOS V11.5.1.

xdump

PR-62597

Attempts to xdump a big filesystem can fail with...

```
NEEDIND: too many blocks
Up the MAXIND constant and re-compile xdump
```

Resolution: Fixed in ConvexOS V11.5.1.

zic

PR-50655

PR-63215

The European Community has now announced daylight savings rules for 1995 to 1997; I'm told this can be found in the Official Journal of the European Communities 30.6.94 L 164/1 as the Seventh Directive 94/21/EC of the European Parliament and of the Council.

In 1996 there will be blessed harmony, and DST will end at 01:00 GMT on the last Sunday in October in all EC countries.

In 1995 in UK and Eire DST will end on 22nd October; this doesn't conform to the rule used for 1981-1994, nor to the rule for 1996 onwards.

Thus for 1995-1997:

Year	Start	End
1995	26 March	24 September (22 October in UK and Eire)
1996	31 March	27 October
1997	30 March	26 October

So /etc/zoneinfo/rules/europe needs changing on both the OS and the various SPUs (the C2 SPU rules were wrong for 1994 as well as for future years).

Here is a patch; I am pretty confident about the UK and Eire parts, but less confident about the continental parts because I don't know which time zones will follow the EC ruling. I don't know what will happen in non-EC countries.

Resolution: Fixed in ConvexOS V11.5.1.

Fixes for ConvexOS Kernel

accounting

PR-23701

PR-27668

PR-30867

PR-31580

The raw accounting files reflect inaccurate accounting data where CPU time is being charged against the wrong UID.

Resolution: The credentials within the thread structure are now explicitly copied from the proc within exit(), in case the exit() call is internal (as a result of a fault/trap/signal/exit sequence) as opposed to the exit() system call.

arch

PR-25819

The include file pagsiz.h is not protected with #ifdef and #endif. Also, the symbol NBPG is defined there and in ...

```
/usr/include/interfaces/kernel_if/arch/ISA/ui_arch.h
```

This makes it hard to update the C++ include files.

Resolution: Fixed as part of the C++-ification of the kernel header files.

arch

PR-29848

The size of the kernel stack "death_stack" was reduced from 511 words to 1k bytes. It should be large enough to hold a C3800 context stack frame (3920 bytes) plus several standard stack frames, i.e. approximately 8k bytes.

Resolution: The size of the stack used to trap improperly nested interrupts is increased from 1K to 8K; the 1K value was too small for C3800.

arch

PR-66199

The subroutine in cinterface.s set_ter_reg() is missing a "plch()". This cache purge is needed everytime set_ter_reg() enables micro-traps to prevent stale cache entries frombeing used.

Once that code is added, there is one line of code in rt_trap_subr.c "rt_syscall_set_utraps()"that can be deleted:

```
asm("plch");
```

Resolution: Added plch to set_ter_reg(), removed it from rt_syscall_set_utraps.

cnvx

PR unknown

Files which include <sys/acct.h> and are compiled via "cc -std" do not compile due to syntax errors. It appears that the typedef for u_short is #ifdef'ed out when compiling with -std.

Resolution: <sys/acct.h> is not an include file specified by Posix 1003.1; it should therefore not be included in source to be compiled in -std mode.

crashdump

PR-27537

'crashdump' fails if pwd is not in /mnt/os.

Resolution: This is not a problem in ConvexOS V11.0.

crashdump

PR-29227

Please have the /mnt/os/crashdump script run 'cpureg' (appending the output to the end of /mnt/errlog), and also run 'osclean' automatically.

Resolution: The ConvexOS V11.5.1 "crashdump" script ensures that 'osclean' has been run. It was decided not to run "cpureg >> /mnt/errlog" from the script.

crashdump

PR-32718

When doing a 'crashdump', you are prompted to enter a comment then a ^C to continue on. If you use ^C you abort the crashdump, without being given a change to change your mind. crashdump should prompt to confirm the abort, if it receives ^C when it should be getting a cntrl d.

Resolution: A cntl-C will now cause the program to prompt the user if aborting the crashdump is what is really desired. User must respond with a 'Y' or program will continue.

crashdump.8

PR-35849

Man page indicates -SH is acceptable usage. crashdump complains about the H unless usage is crashdump -S -H. Customer indicated that -S with other additional options besides H behaves similarly (ie: not inconsistent with man page).

Resolution: modified man page to reflect the options must be separated.

execl

PR-31430

The symbol ___ap$execl is not available in -str mode. This is called from system().

Resolution: moved exel.s (which contains the __ap$execl code) to sysc.a, where it

can be found by devc_str.

getmsg.2

PR-27327

The second paragraph of the getmsg(2) man page mistakenly has the word "mush" instead of "much."

Resolution: Fixed spelling error.

getpeername.2

PR-26383

The SYNOPSIS section of the getpeername man page should have the references to the header files in boldface type.

Resolution: Fixed in ConvexOS V11.5.1.

hang

PR unknown

Network hang. All 2k mblks are allocated and sitting on the NFS server read queue.

Resolution: Modified the STREAMS queue initialization to properly set the queue fields so that flow control is properly exerted.

include

PR unknown

Both <io/lib/ui_drvrlib.h> and <io/tape/cpu/ui_mtio.h> define the structures dev_buf_desc and dev_characteristics. If a program includes both <sys/ioctl.h> and <sys/mtio.h> you get a "tag redefined" error. See test os/man2_sys/ioctl ioctl08.

Resolution: The structure definitions were removed from drvrlib.h.

init

PR-27142

If the system goes down while a 'mvst' is in progress, the 'vvmdaemon' will attempt to restart the mvst upon reboot (from .initrc). The problem is that after executing the commands in /etc/.initrc, init kills off all but the system processes, which kills off the recently restarted mvst(s).

Resolution: Fixed in ConvexOS V11.5.1.

init

PR-36120

PR-42197

PR-45798

The dquot table size is based on nproc and maxusers and isn't really related to the size of the file table. The formula for the size tends to leave it too small for systems with lots of storage and users. (i.e. Data management machines.) There needs to be a different way of sizing ndquot without needlessly increasing the sizes of other tables to excessive amounts.

The results of the small table is table full errors and users unable to open files.

Resolution: Added new tunables in ConvexOS V11.5.1 to explicitly size file system related tables.

intr_asm.d

PR-32370

PR-33259

The death_stack is not large enough. At least 8K should be provided for all architectures, not just C3800s. Some panics on C2s require much more than 511 words.

Resolution: Beginning with ConvexOS V11.0, the death_stack of all multi-processor C-series kernels is increased to 8k bytes; the previous smaller architecture-dependent value (e.g. 512 bytes) was penny-wise and pound-foolish when the time arrived to analyze a crashdump.

iosw

PR-55213

PR-58694

The prtlog process on the SPU occasionally dies without a core file.

Resolution: The prtlog in ConvexOS V11.5.1 has been completely rewritten.

jpd

PR-36489

The jpd kernel debugger quits with the message:

```
symbol_value(): symbol buffer too small
```

in the case that the desired symbol hashes to the end of the rivet'ed symbols and does not match an existing symbol.

Resolution: A boundary condition error in the kernel debugger's symbol table lookup has been fixed; this avoids occasional spurious error exits by the debugger.

kern

PR-29961

Performance problem when MMAP is used to allocate memory. Would like MMAP to run more efficiently.

Resolution: Fixed in ConvexOS V11.5.1.

kern

PR-31669

Since ConvexOS V10.0, the kernel has contained invalid calls to the trace utility due to conflicting assumptions about tracepoints numbering between the base and io portions of the kernel. If tracing is enabled to the extent of allocating a trace buffer in the bootcmd.local file, this can lead to undesired tracepoints, which cannot be disabled, being entered into the the buffer.

Resolution: Although there's still no way to enable/disable the i/o "reserved" tracepoints with the trace utility, the kernel now comprehends the reserved range by allocating sufficient memory. This implies the reserved range of tracepoints stays reliably disabled now.

kern

PR unknown

Input to RESC_CVX_INIT prior to setreslim of resource value L_V_KUNLIMITED, should produce output from getreslim of resource value L_V_KUNLIMITED and output from get_resource_limit of resource value L_V_UNLIMITED. Actual output from getreslim is L_V_UNLIMITED and from get_resource_limit is L_V_OVERFLOW.

Resolution: Fixed in ConvexOS V11.5.1.

kern

PR unknown

A situation exists whereby the vector state of one cpu could be restored on another cpu causing the memory interface tohang on a C3800.

Resolution: Fixed in ConvexOS V11.5.1.

kern

PR-46910

PR-46911

PR-50999

PR-53027

PR-53966

PR-57797

PR-58427

PR-61032

Since loading ConvexOS V11.0, intermittently lose date time stamps, in errlog. This problem is independent of SpuOS, or Diags. Started after upgrading to ConvexOS V11.0.

Resolution: prtlog has been redesigned in ConvexOS V11.5.1.

kern

PR unknown

System crashes with the following error message:

```
rt_isr Attempt to interrupt non real time cpu.
```

The problem is that some system calls can enable the realtime interrupt on a non-realtime CPU. In this case, the realtime interrupt was RT_INTR_A, which is the same as INT_CHAN_5. INT_CHAN_5, is reserved for profiling OR realtime (if profiling is disabled).

The solution has 2 parts:

1) if profiling is enabled, don't allow the user to use INT_CHAN_5 as a realtime interrupt.

2) if profiling is disabled, don't allow the errant system calls to enable INT_CHAN_5.

Note, this is only a problem on machines running the realtime extensions, which is a patch to ConvexOS V11.0.

Resolution: Fixed in ConvexOS V11.5.1.

kern

PR-34802

PR-54057

PR-54124

PR-54496

Numerous customers requesting uping the MAXUSERID to something higher than 32767

Resolution: Maximum UID/GID has been raised to 65000.

kern

PR-54635

ip_uderr inproperly handles a 0 protocol.

Resolution: test 0 != ip_uplinks[i].queue now.

kern

PR-56198

1) Any user may flush routing tables

2) Any user may open for read/write streams equiv of raw sockets

3) Any user may open any network interface for raw read/write

4) Any user may make incomplete entries in arptable (cosmetic)

Resolution: Fixed in ConvexOS V11.5.1.

kernel

PR-29872

Something is hurting performance badly when running xterm with jump scroll disabled.

The xterm seems to be spending essentially all of its real time in select pended on selwait.

Resolution: A bug in tcp header prediction code caused the connection to appear active even while idle. This caused a delay in the processing of output, making local tcp connections slow with bad response time.

kernel

PR-29922

On a 3800 (and possibly all others), if the root file system is dirty, causing the machine to reboot, it will try to reboot to multiuser.

Resolution: This will not happen if the root filesystem is fixed via preen on reboot since the root filesystem is now remounted and a reboot does not occur anymore.

kernel

PR unknown

The SIOCGIFCONF ioctl is not working. It succeeds, returning zero. The returned ifc->ifc_len is unchanged but only the first 4 bytes of the buffer got any data, and it's junk.

By itself, I'd say that's priority C. However, at write_ioctl, the buffer (2nd arg) points to user space and gets bcopy'ed in while holding the stream lock, which I believe leads to a possible deadlock.

Kernel stack trace available on request.

Resolution: SIOCGIFCONF now works properly on both sockets and streams.

kernel

PR unknown

Attempts to bind to (apparently) any port on the localhost interface fail with EADDRNOTAVAIL. Try the program in example on ConvexOS V11.0.0.12 vs. ConvexOS V10.1 to experience the frustration.

Resolution: Fixing an incorrect initialization of the textual name of the interface -from "lo0n" to a null-padded "lo0" - fixes the problem.

kernel

PR-36520

The following unexpected error messages were logged to the console. These did not appear when booting ConvexOS V10.2.

```
Loading vmunix
errintd(R1.21) started, options: -h -s
```

```
mm_sniff: sniff rate: 32.14 MB/day (7.90 days/pass)
vmunix: text: 1441792  data: 118784  tdata: 8192  tbss: 28672  bss:
1523712

[CPU00@13:12:14] ConvexOS -- Version 11.0.0.2.1 Tue Sep 28 21:41:39
CDT 1993
/diag/bin/logmsg:  Invalid event code(EC_OS_START)
+++>
<Wed Sep 29 13:12:14 1993>   /diag/bin/logmsg(3408):boot:0
<Unkn
```

Resolution: The message has been removed.

kernel

PR unknown

faulty read-ahead code could cause system to panic with "cmalloc: count too big"

Resolution: Fixed in ConvexOS V11.5.1.

kernel

PR unknown

The kernel Panic function has some bugs that make it unreliable, in the sense that data from crashdumps can be compromised. This is most likely if the system stays unattended for a while after Panic'ing and before crashdumping.

Resolution: Fixes have been made for cases in which a kernel panic would fail to write out the buffer cache as intended, and/or lose information useful in crashdump analysis.

kernel crash

PR-45632

PR-45674

system crashed:

```
 [CPU00@07:47:58]        ConvexOS: FATAL ERROR: (arch,7044)
C2_user_syscall: t_ttflag not zero
 [CPU00@07:47:58]        sp:     0bd70f6c        a1:     0bd70f6c
 [CPU00@07:47:58]        a2:     015005b8        a3:     00e056f0
 [CPU00@07:47:58]        a4:     0bd70ea4        a5:     00086670
 [CPU00@07:47:59]        ap:     0bd70f80        fp:     0bd70f6c
 [CPU00@07:47:59]       s0: 0000000000000f2    s1: ffffffff00000000
 [CPU00@07:47:59]        s2:
```

Resolution: Move t_ttflag to a spare spot to avoid the corruption problem.

kernel/vm

PR unknown

The new non-inheritance of user page zero leaves the 0xffffd000 mapping uninitialized if the 0x80000000 mapping is explictly remap'ed to MAP_INHERIT.

Resolution: Fixed in ConvexOS V11.5.1.

kio

PR unknown

The "plch" instructions in streconstruct and stparitycheck were not included in ConvexOS V11.1.

Resolution: Fixed in ConvexOS V11.5.1.

kio

PR unknown

ConvexOS appears to lock up for no reason. In the errlog on the SPU, you will see CTRL/A characters as '^A' in the last message logged.

Resolution: To unlock JP, issue a kill -HUP <prtlog pid>. The fix is implemented in the next release of prtlog in ConvexOS V11.5.1.

krpc

PR unknown

CVDM's migdmon uses the kernel event dmon facilities to listen for filesystem events such as remove and truncate. When there are 3 "rm -r" going on for 3 directories of migrated files, CVDM does not get the last 2 events (2 truncates, or 1 remove and 1 truncate). This causes 1 or 2 of the "rm -r" to hang. The 2 "missed" event still sit on the krpc queue but somehow the krpc event seqno and reqno got out of sync.

Resolution: Fixed in ConvexOS V11.5.1.

krpc

PR unknown

A message arriving after krpc_recv is interrupted from esema_await(), but before the channel lock is acquired, will result in the seqno being one larger than expected. This may cause an event daemon process to "hang".

Resolution: Changed krpc_recv() such that if a message arriving after interrupted from esame_await() but before the channel lock is acquired, will redo the message again. Added did_esema_send variable to control the code path.

lockmgr

PR unknown

The kernel component of the lock manager for file locking over NFS does not accept the SIGALRM signal. This is causing problems for notes and other applications as a blocking lock request cannot be timed-out.

Resolution: The problem has been resolved in the kernel lock manager and RPC subsystems.

lpd

PR-28964

'lpd' does not always print what is in the queue. In a case where a job is submitted to the printer while the printer is "offline" and the queue is "empty"; the job gets queued up but the job does not print when the printer is put back "online". If the job is then removed from the queue, the printer put "online" and a job submitted; the job will print.

Resolution:

(1) It was discovered that the real problem so far as this printer is concerned is a bug in the EPROM in the printer. New EPROM's have been ordered for all customers with the Fujitsu printer supplied by special systems.

(2) It was also discovered that one of the PAL's on the VME async/printer card (ACM-201) does not implement reset correctly, thereby exacerbating the problem, since no software action could be taken to try to recover from a hung printer.

mremap.2

PR-26305

The mremap.2 man page does not document EFAULT as a valid error return.

Resolution: The EFAULT description was added to the man page.

msync.2

PR-26381

Since the declaration of msync() uses caddr_t as the type of its first argument, the user needs to include <sys/types.h>. This information should be part of the manpage.

Resolution: The #include has been added.

open.2

PR-27313

PR-29694

In the man pages for open(2) and lseek(2), reference is made to the "ConvexOS Large Files User's Guide" which apparently does not exist.

Resolution: Changed the Large Files User's Guide reference to Large Files chapter in the Extensions User's Guide.

OS install

PR unknown

The OS install script can get hung in an endless loop by taking the defaults in automatic mode when there is insufficient disk space.

Resolution: When running in unattended mode, the os installation script will use no as the default response on questions regarding disk space. It will therefore abort.

os_inst

PR-29424

Page 30, "Copying a new system image from the SPU": there should be a second command listed: "/bin/chmod 644 /vmunix".

Resolution: Fixed in ConvexOS V11.5.1 Installation Guide.

os_inst

PR-29424

Page 10, "Halting ConvexOS": if this is truly an initial installation, then this is probably unnecessary and may be impossible.

Page 11, "Removing /mnt/old_os": same comment.

Page 12, "Step 5" (osclean): same comment. This command will not exist one a virgin installed SPU.

Resolution: Fixed in ConvexOS V11.5.1 Installation Guide.

os_inst

PR unknown

Type-o in "ConvexOS and Utilities V11.0 Installation Procedures" March 1994 -

Pg. 48 in step 1. Should read:

if ["`/etc/knetdctl -q`" = "knetd not configured"]; then

not "'/etc/knetdctl -q'"

Resolution: Fixed in ConvexOS V11.5.1.

passwd

PR-36988

PR-54473

Customer would like user id range to be incerased from a 32k possible values to 64k possible values.

Resolution: Maximum UID/GID raised to 65000

pclose

PR unknown

pclose can reap a child process that it should not touch. It is acting like the pcc mode pclose (ie, reaping too many children).

Resolution: The feature test macros have been changed to one of the ones used in the extended/ANSI C build mode. Waitpid will be used when these feature test macros (_POSIX_SOURCE && _POSIX_OBJECT) are defined.

pm

PR-17704

PR-20217

PR-20309

Request that valid range of uids be extended beyond 32767.

Resolution: UID/GID range has been extended to 65000

pm

PR-28052

There seems to be no easy way to set 'hard' resource limits (e.g. CPU time, memory usage limits, etc.) for interactive users in ConvexOS. Some shells contain some kind of 'limit' built-in which can be used to set resource limits. Because those limit commands are setting only the soft limit of the resource, they can be redefined very easily by the user using the 'limit' or 'unlimit' command. This kind of limit command is almost useless.

Resolution: In ConvexOS V11.0 it will be possible to set login job limits by using the the new limits(1) utility. /bin/login will execute the file /etc/jobs if it exists passing the following arguments: job id of new job, process id of new shell, and uid of user. This allows the local sysadmin to write their own site configurable limits on a per user or system wide basis using the limits(1) utility. Limits(1) allows you to set the soft, hard, and absolute limits on system resources such as cpu time, memory use etc.

pm

PR-50749

Members of the realtime group, specified by the realtime_groupid tunable, get a permissions error when they call the sytem calls that make up the realtime extensions.

Resolution: Fixed in ConvexOS V11.5.1.

pm

PR-61661

A situation exists whereby a trap is handled incorrectly leading to a system "hang".

Resolution: Fixed in ConvexOS V11.5.1.

pm

PR-52974

We added the line

```
tune cpu limits_traditional = 0
```

to bootcmd.local and got during booting from single to multi-user :

```
/usr/spool/convex/reboot_script : 115 Memory fault - core dumped
```

We booted twice with that error. After removing that line the error disappeared.

Resolution: Unable to reproduce this problem on systems running ConvexOS V11.5.1.

pm

PR-66201

The following 3 lines of code need to be added to the bottom of the trap() subroutine. The analogous call is already already at the bottom of the "syscall()" subroutine.

This code will have a user's process jump to the realtime AST processing if a realtime AST was delivered while the process was executing a system trap such as a debugger breakpoint or page fault, etc.

Resolution: Made suggested change, but only executes if rt_enabled is set to 1.

popen

PR-30212

'faillogpr' dumps core in popen(3). The arguments to popen(3) look good so the problem may be in the library routines.

Resolution: The problem was ultimately in the pipe() library routine. popen was calling the direct system call, instead of the sanitized library routine. This has been corrected (the direct system call has been removed, since it does not belong, and the sanitized library routine put in its place.)

pos_conf

PR-29096

The 1990 POSIX.1 standard requires that <limits.h> define POSIX_SSIZE_MAX, POSIX_STREAM_MAX, and POSIX_TZNAME_MAX.

Resolution: Fixed in ConvexOS V11.5.1.

pos_conf

PR-29420

The file <sys/unistd.h> should set _POSIX_VERSION to '199009'. It still references the 1988 standard.

Resolution: The version has been changed to match the purple book.

pos_conf

PR unknown

According to POSIX.1 Sec. 3.3.1.3 (3f), fcntl() should be reentrant with respect to signals (it may be invoked from signal-catching fctns).

In testing this, PCTS calls fcntl() with an argument of F_SETLKW, and then issues it a SIGINT. According to POSIX.1 Sec. 6.5.2.4 and the ConvexOS man page, fcntl() should return -1 and errno = EINTR. This is not happening in beta 2 on orion. This failure was not present in beta 1 on dragonne.

Resolution: Modified the kernel lock manager to correctly handle signals.

pos_conf

PR unknown

According to POSIX.1 Sec. 3.3.4.2, when the SA_NOCLDSTOP flag is set in the sa_flags argument to sigaction(), a parent process should not receive a SIGCHLD whenever any of its child processes stop.

Our implementation fails this test since a SIGCHLD is received by the parent when SA_NOCLDSTOP is set and a child process is stopped.

Resolution: In prior ConvexOS releases, a call to sigaction to change the disposition of any signal other than SIGCHLD cleared the bit representing SA_NOCLDSTOP in the process flags. Thus, only a sigaction(SIGCHLD,x,y) call could set SA_NOCLDSTOP, but a sigaction call for any other signal would clear it. This has been corrected, i.e. only changing SIGCHLD disposition can modify the state of SA_NOCLDSTOP.

posix (tty)

PR unknown

This is a POSIX conformance bug, the original test is from NIST-PCTS. From IEEE Std. 1003.1 (POSIX.1) P136, Ln 329, "If IGNBRK is set, a break condition detected on input is ignored, that is, not put on the input queue and therefore not read by any process." But on our system when c_iflag is set to IGNBRK, the break condition still affects the input.

Resolution: If a serial port on the ACM-201 VME asynch controller detects both a framing error and a parity error, the framing error now takes precedence. As has always been the case for this device, a framing error is handled as a break.

schd

PR-50271

When a process pends on a semaphore, its priority is changed from its run priority to PUSER (the priority of normal Unix processes). This can cause a low priority process to be scheduled before the high priority process after the semaphore is acquired.

Resolution: Fixed in ConvexOS V11.5.1.

schd

PR-33489

These comments are from the original problem report:

```
> [CPU02@14:45:14] /scr: bad dir ino 2 at offset 192 : mangled entry
> [CPU02@14:45:14] /scr: bad dir ino 2 at offset 512 : mangled entry
> [CPU02@14:45:14] /scr: bad dir ino 2 at offset 0 : mangled entry
```

When the customer upgraded to ConvexOS V11.0, the problem reappeared. The cause of the problem isn't the file system. The problem is that whena CPU begins realtime processing, it disables microtraps, thus ending cache coherency. So if a realtime CPU attempts to do file system operations, it may have incorrect file system data encached, corrupting the file system. The solution is to enable microtraps when making a system call from a realtime cpu. The solution is 2 files:

```
base/sched/ISA.cseries/Cmn.mp/rt_trap_subr.c
base/syscall/syscall.c
```

Resolution: Fixed in ConvexOS V11.5.1.

schd

PR unknown

In the release ConvexOS Realtime V11.0, the user is allowed to put each cpu in realtime mode. This is incorrect. The user must save one CPU for Unix processing. An attempt to allocate the last CPU for realtime should return an error.

The file that contains the fix is:

```
base/sched/ISA.cseries/Cmn.mp/swtch_rt.c
```

Resolution: Fixed in ConvexOS V11.5.1.

sched

PR-29905

There is a vector scheduling wrong answer bug in ConvexOS V11.0 on at least C2's, such that when using cpu_conf to disable all but 1 head, vector processes that are handling signals will get their vector registers trashed. This causes wrong answers. The problem may lie in the use of numcpus/nreqcpus throughout the vector scheduling code.

Resolution: A fix for this bug was checked in for ConvexOS V10.2 but somehow failed to be built. ConvexOS V11.0 picked this up when using ConvexOS V10.2 as a base line for their builds.

setrlimit

PR unknown

When the cpu hard and soft limits are set to something small, and a child process is forked off, the child does not get a SIGKILL. See example.

Resolution: Fixed in ConvexOS V11.5.1.

sgen

PR-20973

The system parameter, n_async_proc, is useful when performance tuning the system. It would be beneficial to have this in the tunables file on the SPU and that the impact of this parameter on system performance be documented.

Resolution: n_async_proc is now a tunable

sgen

PR-64200

String Table Overflow failure.

Resolution: Fixed in ConvexOS V11.5.1.

signals

PR-32735

System calls using signals should return an error for signal=32. Currently, there is no signal sent, but no error is returned.

Resolution: kill(pid,32) and killpg(pid,32) now fail, returning -1 with errno set to EINVAL. The previous behavior was to indicate success by returning 0, yet do nothing.

signals

PR-50929

PR-53851

This is a POSIX conformance bug, the original test is from NIST-PCTS. In the test case, a child process was stopped by SIGSTOP when it was waiting to set an exclusive lock on a file . The file has already been locked by a shared lock. After SIGCONT signal is sent to the child process, the child process should continue to wait until time_out; However, the child process set the exclusive lock successfully regardless the existing shared lock. Thus, the test failed.

Resolution: Fixed in ConvexOS V11.5.1.

sync

PR-24611

/usr/include/sync/sema.h includes sync/ui_sema_debug.h more than once.

Resolution: A duplicate #include was removed, was removed.

sysc

PR-29428

Sysconf() always returns _POSIX_CHILD_MAX when queried for _SC_CHILD_MAX. It should return -1 because the value is indeterminate.

Resolution: Changed SC_CHILD_MAX to -1 since the number of processes per real uid is unknown. NOTE: effective uid is max_user_processes, but real uid is unknown.

sysc

PR-53487

This problem is only a bug in the ConvexOS Realtime V11.0 release. The problem is with the realtime extensions added to ConvexOS. The bug is that when running in realtime mode, microtraps are disabled on Realtime CPUs, so when manipulating the file system, a realtime CPU will have incorrect data encached.

The solution is to change 2 files:

```
./base/syscall/syscall.c
./base/sched/ISA.cseries/Cmn.mp/rt_trap_subr.c
```

In syscall.c, check to see if we're in realtime mode, and if so, enable microtraps before making the syscall. (and then turn them back off when done w. syscall) In the file rt_trap_subr.c I've added the routine that checks to see if it's a realtime system call, and if so, enables/disables the microtraps.

Resolution: Fixed in ConvexOS V11.5.1.

tape_ug

PR-51772

The number of tape error messages choked the system. In a seismic environment it is very common to have bad tapes. This crash was completely unnecessary and could have been avoided if either the default print mask was closed or a tunable was provided and documented to control the amount of messages printed.

Resolution: The following should be added in the release notice:

```
The following bit flags apply to "ta_viop_prt_mask":
0x0001    /* prt cmd/esb/fsb on err          */
0x0002    /* prt cmd/csb/esb/fsb/mux1-decode*/
0x0004    /* prt mux[0-3]-hex                */
0x0008    /* prt mux1 velocity errors        */
0x0010    /* prt chain overruns              */
0x0020    /* prt offline/notready messages   */
*** Bits 0x0001 and 0x0002 are mutually exclusive. Bit 0x0002 takes
precedence.

*** Sample entry in the SPU:/mnt/os/tunables file:
ta_viop_prt_mask            0xffbf          0     0xffff     viop

*** Sample entry in the SPU:/mnt/os/bootcmd.local file:
# Turn on "cmd/csb/esb/fsb/mux1-decode" and "mux1 velocity errors"
tune viop ta_viop_prt_mask = 0x000a

The following lists provide an example of each of the above messages:

0x0001    /* prt cmd/esb/fsb on err          */
"ta%d: cmd: %s esb/fsb: 0x%x<%s>\n"

0x0002    /* prt cmd/csb/esb/fsb/mux1-decode*/
"ta%d: cmd: %s csb: 0x%x<%s> esb: 0x%x<%s> fsb: 0x%x<%s> mux1:
0x%x<%s>\n"

0x0004    /* prt mux[0-3]-hex                */
"ta%d: mux0: 0x%x mux1: 0x%x mux2: 0x%x mux3: 0x%x\n"

0x0008    /* prt mux1 velocity errors        */
```

```
"ta%d: transport reported velocity error (tape slippage?).\n",

0x0010    /* prt chain overruns              */
"ta%d: chain overrun.\n"

0x0020    /* prt offline/notready messages  */
"ta%d: not online.\n"
"ta%d: not ready.\n",
```

tty

PR unknown

Timed reads on tty ports may cause a deadlock in rare circumstances. ttytimedsleep has a lock ordering problem with tp->t_trlock.

Resolution: Fixed in ConvexOS V11.5.1.

ttyent.h

PR unknown

/usr/include/ttyent.h:

```
#   if defined(_cplusplus)
```

should be

```
#   if defined(__cplusplus)
```

Resolution: Fixed in ConvexOS V11.0.

ufs

PR-24993

When the event daemon is associated to a directory and a hard link is created such that the associated directory is the target directory and the file being linked to is not associated with the event daemon, no call-outs are received.

Resolution: Fixed in ConvexOS V11.5.1.

ufs

PR-26405

PR-35420

PR-45137

System crashed with a pte violation. Problem appears to be a corrupted inode free list.

Resolution: Fixed in ConvexOS V11.5.1.

ufs

PR-24410

PR-27167

PR-27588

The system hangs with no error messages. Problem appears to be a corrupted inode free list.

Resolution: Fixed in ConvexOS V11.5.1.

ufs

PR unknown

```
[CPU02@14:45:14] /scr: bad dir ino 2 at offset 192 : mangled entry
[CPU02@14:45:14] /scr: bad dir ino 2 at offset 512 : mangled entry
[CPU02@14:45:14] /scr: bad dir ino 2 at offset 0 : mangled entry
```

Resolution: Fixed in ConvexOS V11.5.1.

ufs

PR-36631

When writing an initial file, the _F_CACHEBUFS option works as expected. The data is on disk when the call to fsync completes. However, when overwriting a file which has had the _F_CACHEBUFS fcntl issued on the file descriptor, fsync may return before the data is really on disk.

Resolution: Fixed in ConvexOS V11.5.1.

ufs

PR-41120

An event daemon is not notified of writes to a file that have continued after being pended awaiting free disk space.

Resolution: Modified the retry loop in ufs_rdwr to notify a daemon that it was continuing a write originally blocked due to lack of disk space.

ufs

PR-45471

PR-58120

PR-58979

Unitree NFS local loopback mounts deadlock in fsync waiting on buffer headers.

Resolution: Fixed in ConvexOS V11.5.1

ufs

PR-51293

Applications deadlock trying to acquire an inode lock. The most obvious symptom is when the same process holds the inode and is attempting to reacquire the lock. This is often seen with an inode free list that has the inode pointed to by ifreeh having a null ifreet and ifreeb.

Resolution: Fixed in ConvexOS V11.5.1.

ufs

PR-44424

PR-51450

PR-51685

PR-51975

PR-52076

Processes on systems with greater than 2 cpus can hang with the symptom of the inode free list as seen from the headconsisting of a single inode with it's front and back pointers null.

Resolution: Fixed in ConvexOS V11.5.1.

ufs

PR-54130

Writing to a file (with callouts suspended) which is marked migrated by the use of cvxtruncate is considerably slower than writing to a non-migrated file.

Resolution: Fixed in ConvexOS V11.5.1.

ufs

PR-61008

open(2) drops setgid bit when root creates file when root is not explicitly listed as a member of the group determined by the parent directory.

Problem was noted on ConvexOS V10.1 system during patch installation via pax(1) where root had been given gid=0 rather than gid=10. SetGID bits were dropped from patch versions of netstat, ps, syspic, vmstat, etc.

Resolution: Only drop setgid bit if user is not root.

ufs

PR-59954

PR-60747

PR-60892

The dquot table in the kernel can be corrupted under heavy quota usage. The dquota table freelist lock is not acquired in the dqrele subroutine, thus leaving a hole through which a race condition can be experienced.

Resolution: Fixed in ConvexOS V11.5.1.

ui_reg.h

PR unknown

/usr/include/arch/ISA/ui_reg.h does not compile in -std mode.

Resolution: Include file mis-match between the Programming Tools and ALL.

vfs

PR-50929

There are several spots in the kernel and libraries where SVID style lock structures are improperly handled. Both lockf (libc/CxOS/lockf.c) and vno_lockrelease (vfs_io.c) allocate lockf structures on the stack and fail to zero the structure's memory. This leaves a field (l_xxx) that is documented as reserved, but actually used, in a random state. If it happens to match the valid value (0x000f) it will cause unlock calls for this lock structure to fail and return EINVAL.

close_one_ufd and close_wait_ufd (which call vno_lockrelease) do not check the return value and the lock is left on the file after the file is closed.

I've fixed the two calls described above and also added a couple of tunables that allow fcntl (kern_descrip.c) to clear the l_xxx field if it is set in an original lock call. Since most unlocks are called with the same structureoriginally used to lock a file and l_xxx is only used in unlocks, forcing it to zero on the lock call will help prevent the problem for existing apps.

Resolution: Fixed in ConvexOS V11.5.1.

vfs

PR-50929

In fcntl, the code below is responsible for copying out the (possibly modified) lock structure back to user space. Comments in the code and the man page (and the Sun NFS 4.1 reference source) indicate that only F_GETLK commands are expected to return anything, but our source thinks it returns something and the Sun source does return something, though it is of the wrong size and likely breaks things. The lockmgr also expects that an fcntl call can modify the lock on a F_SETLK call, so at least it's not unexpected.

Our source has an if nesting bug that causes it to not return what it thought it was and returns the wrong thing in one case.

Resolution: Fixed in ConvexOS V11.5.1.

vfs

PR-53822

Disk buffer cache buffers belonging to AFS may not be flushed in some instances. This could lead to a system hang. "bclean" does not call "bcleanmnt" for buffers of type MOUNT_AFS.

Resolution: Fixed in ConvexOS V11.5.1.

vm

PR-27304

When mmap and munmap calls are mixed to create "holes", subsequent calls to mmap() may fail when using MAP_PRIVATE | MAP_FILE.

Resolution: Fixed in ConvexOS V11.5.1.

vm

PR-41576

The 'rdcontrol' RamDisk creator causes a kernel panic when run.

Resolution: Fixed in ConvexOS V11.5.1.

vm

PR-34545

The efficiency of the routine "initptespte" must be improved. Repeated use of "pate" instructions degrades system performance.

Resolution: Fixed in ConvexOS V11.5.1.

vm

PR-51415

ConvexOS: FATAL ERROR: (pm,8516) trap: unresolved kernel pte violation crash occurs because vnodetoobject does not check the size of the structure returned from MALLOC.

Resolution: vnodetoobject() will now pend until more memory becomes available.

vm

PR-52526

A situation exists in which a deadlock on cmfree_lock may occur if clearbufptes calls a function that ultimately calls memall.

Resolution: clearbufptes() releases cmfree_lock if it needs to do the swapfree().

vm

PR-55769

"cmgetblk assumption" panics involving optimizations surrounding cm_flags field in coremaps.

Resolution: ConvexOS V11.5.1 uses gcc 2.5.8 to build kernel.

Fixes for ConvexOS I/O

ACS

PR-59955

After installing the new ACS software it has been noticed that occassionally a siloquery volume XXX command (as issued by dismount.stk) will never exit. The dismount script will get hung in a read from the siloquery. I can't tell why the siloquery is hung, but it is in a call to the subroutine acs_response.

The new dismount.stk that was distributed with unitree/2.0.136 uses a lock file to serialize requests to siloquery. When a siloquery hangs the lock functionality causes a stream of dismount.stk requests to build up as pvrsrvr times out on them and issues new ones, which all stack up waiting on the lock file.

Until we figure out why the siloquery is hanging, if pvrsrvr is going to give up on a dismount request, it should kill the process group of the dismount (and make sure that the dismount is put in it's own group) after the timeout.

Resolution: Fixed in ConvexOS V11.5.1.

ACS

PR-61030

Apparently issuing silo commands while the ACS server is down causes the silo command to hang even after the ACS server comes back up. The siloquery drive command was killed at 07:34 and returned an error status then.

We are implementing a work around to timeout on a failure to respond within a suitable time. We would like to see this functionality added internally to the silocommands.

Resolution: Fixed in ConvexOS V11.5.1.

QSC

PR unknown

ioctl_tur does not check return code of TUR operation. It merely checks that the TUR was issued and a response was received. It should check the status of the TUR operation.

Resolution: Fixed in ConvexOS V11.5.1.

TLI

PR-61612

At IDRIS, the customer changed its old STK 4480 drives by new STK 4490 drives. Since this time, some errors are seen, that were not happening before:

When the drive is accessed 1st time:

```
[CCU39@12:37:06] tc7: Unit Check, Recovery Code = 0x4d, Reserved
(invalid)
[CCU39@12:37:06] tc7: Sense  0-15: 0248004d 00000020 0000155f
00000000
[CCU39@12:37:06] tc7: Sense 16-31: 00000002 00000000 e70edd20
05131100
```

(same error for each 4490 drive)

And from time to time (2-3 times a week):

```
[CCU39@16:35:05] tc: Wait CU Ready timeout (60 seconds).
[CCU39@16:35:05] tc7: Wait CU Ready timeout (60 seconds).
[CCU39@16:42:49] tc: Wait CU Ready timeout (60 seconds).
[CCU39@16:42:49] tc8: Wait CU Ready timeout (60 seconds).
[CCU39@17:33:31] tc6: Unit Check, Recovery Code = 0x22, Path
Equipment Check
[CCU39@17:33:31] tc6: Sense  0-15: 10400022 000efb20 d600250b
22320000
```

```
[CCU39@17:33:31] tc6: Sense 16-31: 88e0803f 04090000 e70edd20
05130000
[CCU39@17:34:39] tc: Wait CU Ready timeout (60 seconds).
```

We currently do not handle the 0x4d Unit Check which is a resetting condition that will show on 4490's after reboot of the machine. Also need to extend the CU Ready timeout, or possibly add tunables for timeout values.

Resolution: Fixed in ConvexOS V11.5.1.

TLI

PR-61206

Crashdump taken.

```
[CCU37@21:23:25] pq_lock: Lock (00400f7c) failed after 200 attempts,
still trying.
[CCU37@21:23:25] Controller Error: Port 1, Device 5, CCS 2be, 08
(DPED Over/Underflow)
[CCU37@21:23:26] pq_lock: Lock ( acquired, resume processing.
```

Resolution: Fixed in ConvexOS V11.5.1. Increase value to 1000.

UNIX IO

PR-46160

Took another panic(disk,3252). I verified that spuio was still running and the printer had not been selected for echo. They are going to make sure that the printer is normally on-line and ready.

They asked if we would look at a crashdump since the obvious reasons did not pan out and the panic has occured twice within the last four months. I told them we would. (I'm hoping we can determine what the majority of the MBS buffers were allocated for.)

Resolution: The new MBS system now has more capacity plus the du device driver retries message allocation attempts which fail.

VME SCSI

PR unknown

Resetting the VME SCSI bus while commands are pending on the Ciprico board, may cause the viop to die and flood the errlog with Spurious Interrupt messages from the 68020. There is a small window of opportunity for this to happen.

To reproduce I used 3480 stacker with a 3590 tape in it and "stacker loadmag". This causes an infinite retry of test unit ready on the "wrong" tape. Then mt -f /dev/dat2spt reset of a drive on same bus. May have to repeat a few times, as timing is critical. The important thing is to see the command aborted in the unit options phase, which means it was an active command. After reset complete the Spurious Interrupt messages will loop out.

```
[CCU10@13:21:29] WARNING: SCSI adapter (vme1/slot4) resetting...
[CCU10@13:21:29] WARNING: SCSI adapter (vme1/slot4) stop list
issued.
[CCU10@13:21:31] WARNING: SCSI adapter (vme1/slot4) bus reset
asserted.
[CCU10@13:21:31] WARNING: SCSI adapter (vme1/slot4) board options
set.
```

```
[CCU10@13:21:31] WARNING: SCSI adapter (vme1/slot4) dma windows
flushed.
[CCU10@13:21:31] WARNING: SCSI device  (vme1/slot4/unit0) unit
options.
[CCU10@13:21:31] WARNING: SCSI device  (vme1/slot4/unit0/su0) init.
[CCU10@13:21:32] WARNING: SCSI device  (vme1/slot4/unit2) unit
options.
[CCU10@13:21:32] WARNING: SCSI device  (vme1/slot4/unit2/su0) init.
[CCU10@13:21:32] WARNING: SCSI device  (vme1/slot4/unit3) unit
options.
[CCU10@13:21:32] scsi_tape(vme1/slot4/unit3/su0/stacker load mag/
0xcd) ADAPTER ERROR(0x76) SCSI adapt
[CCU10@13:21:32] r reset, commands disabled until device open
[CCU10@13:21:32] WARNING: SCSI device  (vme1/slot4/unit3/su0) init.
[CCU10@13:21:33] WARNING: SCSI adapter (vme1/slot4) cmd queue
started.
[CCU10@13:21:33] WARNING: SCSI adapter (vme1/slot4) reset complete.
```

Resolution: Clear ISR register prior to returning to sleephw loop after reset.

VME SCSI driver

PR-64514

Problem communicating with Metrum tape drive after cleaning the drive.

Customer has experienced repeated problems. Suspected some difficulties withprocedure; stopping UniTree adequately, etc. Getting SCSI timeouts on Metrum tape drive:

```
scsi_tape(0/2/0/0/dev open/0x0) adapter error (0x1e scsi select
timeout)
```

Received 500 in past 15 minutes.

Tested cleaning procedure on Metrum drives attached to sequoia. Performed following operations:

```
    rbtstate off
    rbtman on
    kill <utaped-pid>
    <clean the drives>
    mt -f /dev/rrsp0 off   <---noticed that drive was online,ready
with no tape                                  in drive
    rbtman off
    rbtstate on
    $STORAGE/adm/bin/utaped
```

First problem noticed was state of top drive after cleaning---set it offline after cleaning. After inserting tape into the drive, unable to read it.

Appeared to have mechanical problems with the lower drive attached to sequoia after cleaning---it was eating tapes.

Resolution: Provided scsi bus reset capability to recover "lost" drives. Metrum drives will become lost if access is attempted while cleaning.

VME SCSI tape

PR-64523

Desire support for Fujitsu 3490E SCSI tape in generic SCSI driver.

Resolution: Fixed in ConvexOS V11.5.1.

cpu_ta

PR unknown

Stress test is stuck in close/process_event.

Resolution: Fixed in ConvexOS V11.5.1.

crashdump

PR-30414

PR-30606

PR-58393

The crashdump utility currently knows how to write to only a few different tape devices. It needs to be able write to 3480 drives which support compression and possibly other newer device types.

Resolution: Added VME SCSI supported devices, fixed crashdump for tli drives.

er90

PR-46764

If the ER90 fails an eject command, then the ITC kill the port on which the recorder is connected. This requires that the convex be rebooted to recover the use of the recorder.

Resolution: Fixed in ConvexOS V11.5.1.

exabyte

PR unknown

The Exabyte drive 8505 needs to be supported. It currently is not and there are customers that want the expanded capabilities.

Resolution: Fixed in ConvexOS V11.5.1.

exabyte

PR-44246

Two systems in my region have exabyte subsystems. Both subsystems hang after a period of time of writing to the tape. The process cannot be killed in both systems cases. This has become a big problem to both these Seismic contractors as the clients want DATA on 8mm exabyte tapes. I am investigating the possibility of putting the driver in a trace back mode to determine what is getting hung. Can anyone help!!!

Resolution: Fixed in ConvexOS V11.5.1.

exabyte

PR-52966

The customer needs to create some Exabyte 8200 tapes using an Exabyte 8500 device, for compatibility with other systems. The current ConvexOS V11.0 driver does not provide any (documented) means to do so (the old ConvexOS V10.2 Exabyte 3.4 driver was able to manage 8200 tapes on write). The ConvexOS V11.0 driver should be modified to fully support 8200 mode. Meanwhile, I need a workaround in order to 'tar' or 'dump' in 8200 mode on the Exabyte 8500.

Resolution: Fixed in ConvexOS V11.5.1.

exabyte

PR-53021

PR-56713

Problem with exabyte driver 11.0.3/TAC patch 11.0.171.2. When mounting tape read-only (the tape is read only) on a 8500 if the write format on the tape is 8200 the driver do not detect or do not get the file protect mode/status from the drive. Using a tape write in 8500 mode works fine.

Resolution: Fixed in ConvexOS V11.5.1.

fujitsu-3480

PR-38989

Pressing the reset button at the wrong time, will hang the drive and/or formatter.

Resolution: No longer occurs, or power/cycle and bus reset capability will allow recovery ifit does occur again.

netinet

PR-58972

System crashed after patch 11.0.170 was loaded. The patch has now been removed from this machine.

```
[SPU  @17:12:16] <Thu Aug 31 1995>
[CPU00@17:45:10]        ConvexOS: FATAL ERROR: (pm,8516) trap:
unresolved kernel pte violation
[CPU00@17:45:10]        sp:     0b9b9fe4        a1:     025c9ab0
[CPU00@17:45:10]        a2:     025c9ab0        a3:     025c9a64
[CPU00@17:45:10]        a4:     00000006        a5:     00000c06
[CPU00@17:45:10]        ap:     0b9b9ff8        fp:     0b9b9fe4
[CPU00@17:45:10]       s0: 0000000000000000    s1: 0000000000000000
[CPU00@17:45:10]       s2: 4800000900080000    s3: 0000000000000301
[CPU00@17:45:11]       s4: c13f58020000b9be    s5: 000000000b9be000
[CPU00@17:45:11]       s6: f311fce600000001    s7: 4800002900402898
[CPU00@17:45:11]        int. mask: 000000ff
[CPU00@17:45:11] syncing disks...
[CPU00@17:45:12] 42 4 0 1
[CPU00@17:45:13] SYNC TIMED OUT!
[CPU00@17:45:13] halting processor
```

Seems to be a rarely encountered bug in ip source routing.

Resolution: Fixed in ConvexOS V11.5.1. Added type casting during calculation of ip_nhops in ip_srcroute().

scsi driver

PR-52914

PR-53203

PR-53205

Seems we cannot do a SCSI reset when one or more of our Metrum drives goes into the weeds. Have to reboot to clear the SCSI card.

Resolution: Fixed in ConvexOS V11.5.1.

tape

PR-33291

What happens to the data that is stored in the buffered log in the IBM 3480 cartridge subsystem? Operating systems such as IBM's MVS and VM get this data and record it. As some of the data is concerned with errors corrected by the 3480 subsystem, it is important to collect and monitor this data, so as to be able to spot cartridges that may be about to fail. What is ConvexOS doing with this data?

Resolution: IO_GET_XSTAT returns REQUEST LOG data followed by REQUEST SENSE data for 3480 on VME SCSI.

tape

PR unknown

An ioconfig with many different tape controller types will run out of entries and fail to attach some controllers:

```
ta: probe MTC-201 failed, tape types exceed 9
```

Resolution: Expand tape controller entries.

tape driver

PR-51940

After replacing the CPU which was a C2 by a C38 , a scsi device won't respond at boot time any more. The problem occurs during probe and attach phase. The same configuration works fine on a C2 but fails on a C38.

Resolution: Fixed in ConvexOS V11.5.1.

tape driver

PR-52296

PR-54247

There is a problem with backup scripts (using the dump utility) that abort when it's time to load a new magazine into the 3480 stacker. tpdaemon/stkdaemon get a bad status from the stacker when it's time to change magazine and the dump aborts.

Resolution: Fixed in ConvexOS V11.5.1.

termcap

PR-61362

When running on a SunSparcstation 5 with Solaris 2.5, Convex is not using a correct termcap entry.

When decompiling tic entry for /usr/share/lib/terminfo/dtterm from Solaris 2.5 and producing a reusable termcap entry the terminal is not accepting full terminal edition and complains about a too long entry for termcap.

The termcap entry needs to be converted to syntax that ConvexOS will understand, as most termcap codes need to be converted to the older two character equivs.

Resolution: dtterm is a special emulator provided by some workstation vendors to provide DEC terminal emulation. It closely matches that of a DEC VT220, so the solution was to add dtterm alias on vt220 entry.

tli

PR-54958

PR-55801

tli 36 hung after a "waiting CU" error.

Resolution: Modified the accounting of hardware errors to not occur if there was no active command.

viop

PR-50434

To avoid this kind of problem we need a soft reset command for Formater/ Device driver: this provide us a way to restart the formater and the device driver that is hanged without rebooting the c38.

Two examples :

1) First a 3480 driver (one among 8) : we can't use it , no access is available , the system sees it as not on line . Then when ConvexOS is hung we can find error messages in the errlog file .

2) When a 3480 driver has got a write error , the driver is in an infinite loop, and ConvexOS crashes if we don't kill the job using this driver.

Resolution: Fixed in ConvexOS V11.5.1.

viop

PR-53144

A ConvexOS 11.0.3 scsi tape driver deficiency could lead to a "ConvexOS: FATAL ERROR: (arch,7036) Exception_handler: Not enough context stacks" panic if "scsi_cntl: ADAPTER ERROR(0x1f) SCSI disconnect timeout" occurs.

Resolution: Fixed in ConvexOS V11.5.1.

viop

PR-55199

What used to take 2 minutes to position the tape drive now takes 20 or more minutes since the upgrade of TAC Patch 11.0.171.3.

Problem is scsi timeout is not multiplied in the block position command, as it is in other commands (rew, fsf, etc...).

Resolution: Fixed in ConvexOS V11.5.1.

viop

PR-60200

When writing to "bad" tape on a 3480 scsi drive w/ formatter, after aborting operation with I/O error, drive is no longer accessable until power reset the formatter. The commands used were:

```
tpmount -a /dev/rtc0nu -s TC
tar cvf TC /usr/tftpboot(several times until I/O Error)
mt -f TC stat(fails w/ I/O Error)
tpunmount -s TC
```

Using tpunmount frees link but device doesn't eject tape. Able to issue tpmount command again, but any access results in I/O Error.

Checking the stacker status now results in:

```
stacker: Unable to open ACL device: /dev/tc0stk
open: No such device or address
```

/mnt/errlog shows:

```
[CCU01@12:51:38] scsi_tape(1/4/0/0/wrt phys/0xa) Recovered error
ASC=0x5b ASCQ=0x2
[CCU01@13:04:00] scsi_tape(1/4/0/0/wrt phys/0xa) Medium error
ASC=0x0 ASCQ=0x0
[CCU01@13:04:00] scsi_tape(1/4/0/0/wrt eof/0x10) Medium error
ASC=0x0 ASCQ=0x0
[CCU01@13:04:25] scsi_tape(1/4/0/0/dev open/0x15) Medium error
ASC=0x0 ASCQ=0x0
[CCU01@13:05:03] Previous message from CCU01 repeated 9 times in 29
seconds
[CCU01@13:23:39] scsi_tape(1/4/0/0/stacker mode select/0xce) Medium
error ASC=0x0 ASCQ=0x0
```

Console also shows "SPU trap interrupt 10" and "Errno 5: unable to get status w/ MTIOCGET" over and over.

Resolution: Fixed in ConvexOS V11.5.1.

viop

PR-58131

The DLT drive manual claims the default behaviour for BUS RESET is to not rewind, but testing and customer sites prove that it is. This can be fixed by sending the MODE SELECT for the Vendor Unique page (0x3E) with REWINDONRESET 0. At least that's what the manual claims. Without this fix, a bus reset may occur, rewinding the tape, then closing the device may write EOD at current position. We either need to handle bus reset better or make sure tapes don't rewind on BUS RESET.

Resolution: mt reset command will fail current & pending commands. Also, a device must be closed and reopened after a bus reset, before any modification commands are allowed on that device.

Fixes for ConvexOS TCP/IP

ftp

PR-34445

If you ftp a file with a long file name, hashing gets turned on automatically.

Resolution: Fixed in ConvexOS V11.5.1.

ftp

PR-46267

ftpd lacks the page initialisation code in fx_async_init, which is incorporated in the corresponding routine in ftp. The effect of this is that the first file transferred from the Convex ftpd (remote ftp executes a get), if smaller than 10-20Mb, is transferred with speed of 500kb/s (or below) over FDDI.

Resolution: ftpd now touches all its pages, like the client does, which eliminates this performance problem.

ftp

PR-51504

If a ConvexOS V11 ftp client is used to transmit a file to another system and some error causes the remote system to cease accepting data and close the data socket, the ConvexOS V11 ftp client terminates immediately with no error message of any kind. The correct behaviour is to print an error indicating the broken pipe/ socket and to print any other error message returned by the remote ftpd.

This situation can arise (for instance) if one sends a file to an account which is over its disk quota, or which goes over quota during the transfer. The ConvexOS V10 client prints the following messages in this case:

```
netout: Broken pipe
452 Error writing file: Disc quota exceeded.
ftp>
```

and similar things occur with any other respectable client.

Note: the problem in ConvexOS V11 doesn't just cause ftp to stop sending ; the process itself actually terminates.

Resolution: There were two problems causing this bug. The auxilary async I/O process now ignores SIGPIPE. It should terminate with a Broken Pipe I/O error now. The other problem was caused by the Convex ftp server not flushing its status message upon a command which failed to complete successfully. This would cause the ftp client to hang until the server idle timeout occurred. The act of the server exiting flushed all of its STDIO buffers and would unhang the client.

ftp

PR unknown

ftp prints out the "local: filename" and "remote: filename" messages only when NOT in verbose mode. It needs to print the messages only when IN verbose mode.

Resolution: 'local: filename remote: filename' printed only in verbose mode, instead of only when NOT in verbose mode.

ftp

PR-64002

in.ftpd does not filter "pass" commands properly when in debug mode. It only checks for "PASS" when masking password information when logging; however, it is case-insensitive when accepting commands.

Resolution: Fixed in ConvexOS V11.5.1.

io2knetcf

PR unknown

io2knetcf needs to understand LAN-205 for the Eagle Ethernet board.

Resolution: Fixed in ConvexOS V11.5.1.

lib

PR unknown

There is a bug in syslog(), in which it does not check to see if the message being logged will overflow its internal buffer. This is currently being used to break security on several machines across the Internet (see CERT Advisory CA-95:13), and could potentially cause security problems for C-series machines, as well.

Resolution: Use our own sprintf() function, so that we can set the _bufsiz that gets passed to _doprnt. By making our output buffer bigger than this size, we should avoid the possibility of overflow. Also, make sure we don't overflow when we do %m replacement on the format string.

man

PR-47131

The man page describes a "-g" option for traceroute, but traceroute reports a usage message that does not include "-g".

The man page should reflect reality.

Resolution: Removed -g from man page in ConvexOS V11.5.1.

netmgt

PR-42131

A non-privileged user can ifconfig an interface down.

Resolution: Fixed in ConvexOS V11.5.1.

rcp

PR unknown

rcp maintains an open file descriptor to /dev/ip for the duration of its execution. It needs to close it once it has obtained the list of system interfaces.

Resolution: rcp now closes its /dev/ip file descriptor when it is through with it.

rdist

PR-41543

Every user can become ``root'' in five minutes if ``rdist'' is available on the local host!

Resolution: Fixed in ConvexOS V11.5.1.

route

PR unknown

Under certain circumstances, when a route is deleted from the table, it can not be immediately readded.

Resolution: Once deleted, routes can now be readded.

route

PR-44129

PR-45554

multiple redirected route entries (through the same gateway) are kept in

the route table.

Resolution: Fixed in ConvexOS V11.5.1.

route

PR unknown

gated breaks routing due to gated add a route over hippitcp interface and the kernel cannot support I-field to the route for hippitcp.

Resolution: Fixed in ConvexOS V11.5.1.

rsh

PR-51401

The ConvexOS V11.0 rsh server returns truncated stdout when used with certain clients. It only happens when the -n parameter is used, and it does not occur when the server is running ConvexOS V10. When it does occur, the command

```
rsh io -n cat /usr/lib/libveclib.a >junk
```

produces a truncated output file with approximately 62k bytes in it, and no error message (the file can be any file, but it needs to be big).

Can be reproduced with clients running on

Sunos 4.1.3

Solaris 2.0

HP-UX

VMS (Multinet TCP/IP)

Unicos 8.0

but failed to reproduce it with clients running on IRIX 5.2 or Convex-OS. It happens on both Ethernet and FDDI connections.

Resolution: Some non-Convex rsh clients create a half-closed state using the shutdown() system call. This state allows one end of a TCP connection to terminate its output, while still receiving data from the other end. This state places the socket in the SOSTRM_ORDREL state. When the server sends a large volume of data to the client, flow control operations are initiated to throttle the data. This causes some code to execute that erroneously detects a disconnected socket and aborts with a SIGPIPE signal.

The streams code was modified so that it checks for this half-closed state and does not abort with SIGPIPE if that is the case.

select

PR unknown

A file is opened only for writing, but select() call returns that the file descriptor is ready for reading.

Resolution: This is how select() works under ConvexOS, HP-UX, SunOS, and Linux. Fixing this feature would probably be a bad idea, since it could cause some poorly written programs to hang indefinitely.

socket

PR-41139

Ingres fails during the connect system call with ECONNRESET.

Resolution: The system no longer returns duplicate local/remote port/address pairings when sockets are bound during connection establishment.

socket

PR-44716

PR-44862

PR-45306

PR-46741

Data is sometimes corrupted over loopback TCP connections.

Resolution: Loopback IP connections have been modified, so that TCP will no longer corrupt data.

socket

PR-50794

udcompact() is causing problem for utilities that uses SOCKET_DGRAM type socket. for example, syslog entry will be misformatted due to compacting messages in udsoputbq(). udsodgrrecv() will read in all the compacted data once and syslogd will consider all the data as one record message, thus print without newline.

Resolution: Does not do udcompact() if the socket is not SOCK_STREAM in ConvexOS V11.5.1.

streams

PR unknown

A process can suck up all of the system's mblocks, causing all streams activity to stop. Example program included.

Resolution: Fixed in ConvexOS V11.5.1.

tcpip

PR unknown

Loopback icmp packets are not set with the proper length.

Resolution: Loopback icmp packets are now set to the proper length.

tcpip

PR-42946

udpcksum appears to be broken. When udpcksum is enabled, all UDP RPC services fail to work, and the remote machine's 'netstat -p udp' command reports bad udp checksum errors.

Resolution: The problem has been resolved in the ConvexOS UDP/IP protocol.

tcpip

PR-43689

PR-43916

When heavy load causes a system to run out of stream resources (dblks), the data transfer can experience repeated data blocks and/or corrupted data.

Resolution: Modified tcp_send code so non-blocking writes return the number of bytes written when a send fails due to no buffer space.

tcpip

PR-45067

System has experienced streams problems several times. Related events are a dump of a filesystem to DAT and running a NASTRAN job. At each case, the 256 data blk allocation failed numerous times. ex.

```
dblks(256)        1024       616       408     34353       619   3396655
```

resulting in rlogin hung, pipe stops to function.....

Resolution: ConvexOS V11.5.1 implements additional flow control to reduce the datablks a connection can hold.

tcpip

PR-46222

PR-46968

ConvexOS: FATAL ERROR: (pm,8516) trap: unresolved kernel pte violation. Caused by the possibility that a null tp value passed in tcp_ndata_out().

Resolution: Fixed in ConvexOS V11.5.1.

tcpip

PR-44728

Cray ftp put using window scale > 2 hung the ftp session due to the fact that Convex receive window was set to 40000 while Rcv_Wind_Scale was at 0. Thus the 16 bit th_wnd sent to Cray was value 0. note: tcphippi_default_window setting is cuasing this problem.

Resolution: Fixed in ConvexOS V11.5.1.

tcpip

PR-50161

system crashes with "kernel pte violation". When tcp_ordrelreq() close down rcv channel it will set the queue->next to null. When a final close on the socket() is done, it will close all the queues in unlinkq() first. Since the pointer is null already, crash with PTE violation.

Resolution: Fixed in ConvexOS V11.5.1.

tcpip

PR-50652

PR-51287

PR-51800

PR-52356

PR-54351

PR-54624

PR-54628

PR-57108

system crash due to null mb_freelist. Yet there is plenty of mblks/dblks.

Resolution: Disabled nfs write clustering in ConvexOS V11.5.1.

tcpip

PR-45541

tcp loopback connections loop due to an insane q_minpsz.

Resolution: Fixed in ConvexOS V11.5.1.

tcpip

PR-51789

Application reports connection reset after upgrade to ConvexOS V11.0.

Resolution: Fixed in ConvexOS V11.5.1.

tcpip

PR-52892

IP Spoofing vulnerability

Resolution: tcp sequence numbers are incremented randomly to reduce the chance of IP Spoofing in ConvexOS V11.5.1.

tcpip

PR-55427

System panic if tune class of unused dblks to 0 due to dblk allocation rollover does not check for whether there are dblks configured or not in the higher dblk classes.

Resolution: Fixed in ConvexOS V11.5.1.

tcpip

PR-51842

PR-56155

PR-56156

System panic due to ip_p=0 which cause 0->0->4thfield, ie, the exception handler vector code location , to be passed to canputd() as a queue.

Resolution: Fixed in ConvexOS V11.5.1.

tcpip

PR-55838

PR-56523

System hung with streams holding kmsema and looping on one head. caused by off syn sequence numbers.

Resolution: Readjust read pointer back when there is not enought space to hold tcp_hdr in tcp_input.c.

tcpip

PR unknown

With 11.0.181 and 11.1.140 a normal user can't do "ifconfig eth0".

Resolution:ifconfig now runs setgid group kmem.

tcpip

PR unknown

Negative byte totals in tcp section of "netstat -s".

Resolution: Changed source netstat/inet.c to use %u format rather than %d format for unsigned longs in structure tcpstat. Other outputs already use %u.

tcpip

PR unknown

Negative byte totals in tcp section of "netstat -s".

Resolution: Fixed in ConvexOS V11.5.1.

tcpip

PR-62609

Connections hang around forever due to the write side queue being full, and the process calling shutdown is pending in strpend(). The remote side sent us a TH_RST, but the shutdown is pending on the write queue.

Resolution: flush write side queue when received TH_RST from remote.

timedc

PR unknown

The mtime command to timedc hangs on ConvexOS V11.0 systems.

Resolution: Raw socket is now supported in ConvexOS V11.5.1, so timedc is working.

uipc

PR-43653

Unix domain socket stream implementation breaks pipe, causing user application to deadlock. See comment please.

Resolution: UNIX domain sockets now compact data into fewer buffers when too many buffers are being used.

Fixes for Convex NFS

nfs

PR-46390

PR-52330

NFS lookup is very slow when the attribute is not cached. ex. ls -F /common on sherwood, where /common is NFS mounted from lovecraft on sherwood, Turn on and off "doingcache" via. adb to see the system time of the ls command.

MIT also experiences slow system time with doingcache turned on, seems the caching is not working and as longer the system is up, the problem gets worse.

Resolution: Added safe_makenfsnode to choose between slow and fast makenfsnode. Call cmincal() only when needed.

nfs

PR-51794

A deadlock may occur if routines called by nfs_lookup try to lock the same vnode that nfs_lookup has already locked.

Resolution: Fixed in ConvexOS V11.5.1.

nfs

PR-51793

open(fd, O_CREATE | O_WRONLY) passes on a r--r--r-- nfs mounted file; ex.

```
% ll data
-r--r--r--  1 cwang           0 Jan 20 09:52 data
% cp data.s data
% ll data
-r--r--r--  1 cwang          26 Jan 20 10:16 data
```

Resolution: Added a variable nfs_write_restrict.

nfs

PR-54710

A deadlock is possible due to nfs_cache_check calling bdinval while holding the rnode semaphore.

Resolution: RNODE_BREL(rp) before call bdincal() in nfs_purge_cache().

nfs

PR-51431

"ConvexOS: FATAL ERROR: (vm,6636) cmgetblk assumption" can be caused by nfs mounts of filesystems with blocksizes that are not a power of 2.

Resolution: Round down blocksize to power of 2.

rpc

PR-62044

PR-64017

There is a bug in rpc.statd (aka statd on some systems) by which someone can delete/create files on a remote host.

Resolution: Fixed in ConvexOS V11.5.1.

rpcgen

PR-54752

rpcgen generates bad .h file with enum{} definition. Repeat-By:

```
rpcgen -h nfs_prot.x | less.
```

Resolution: Fixed in ConvexOS V11.5.1.

Fixes for Convex TCP/HIPPI

hippi

PR unknown

tcp/ip hippi uses 15872 (31*1024/2) as its mss value though hippi MTU is at about 64k. This decreases hippi performance. The mss value is affected by tcp_sendspace and tcp_recvspace, which are set to 31k in kernel.

Resolution: tcp advised window size is set before mss is calculated.

reset

PR-51959

From time to time, Hippitcp hung and there is no way of clearing it unless a reboot of the whole Convex system. We need a way to reset ccu and reinitialize the interface.

Resolution: Fixed in ConvexOS V11.5.1. ifconfig up/down now actually does something.

Fixes for Convex AFS

kernel

PR-45004

PR-45018

PR-46262

unable to allocate 16 bytes from memory

Resolution: Fixed in ConvexOS V11.5.1. This was due to a bug in ConvexOS streams network, namely a bug in the strstr() function, which was causing a memory leak.

VIOP device driver memory usage

A

This appendix provides information about memory usage in the VIOP by ConvexOS device drivers. For ConvexOS to configure and boot properly, several resources must be available. This resource list includes the amount of memory used on the VIOP by device drivers. This document describes the resources used by various ConvexOS device drivers and shows you how to determine if your particular configuration will have problems configuring or booting.

Note

Prior to installing ConvexOS V11.5.1, you should read this document and determine if your system has enough VIOP memory for your particular configuration. If you do not have enough memory when you upgrade to ConvexOS V11.5.1, your system may not boot.

To help you determine whether a particular configuration will work, this appendix presents information in the following sections:

- VIOP memory system overview
- Device driver memory usage
- VIOP memory configuration worksheet
- VIOP system configuration
- Example worksheet

VIOP memory system usage

The standard VIOP contains 512 kbytes of local memory. This memory is divided into 128 four-kbyte pages and is a resource used by all ConvexOS VME device drivers. The RTIOP/VIOP-2, contains 1 Mb of local memory divided into 256 pages, has twice as much memory as the standard VIOP.

ConvexOS VME device drivers use memory for

- **Event Governed Operating System (EGOS)**—EGOS is a simple operating system that runs on a VIOP and determines what device driver code to execute. The EGOS code is stored in VIOP memory.
- **Device driver text/data segments**—Each device driver present in your viop image has an associated text and data segment. These segments are stored in VIOP memory.
- **Device driver allocation**—Most device drivers dynamically allocate VIOP memory at system boot time to properly handle I/O requests. This memory is needed for controller interfaces, buffers and state information. The amount needed is different for each device driver.

Device driver memory usage

Each device driver uses some amount of memory for its exclusive use. This section describes the number of pages of VIOP memory each device driver uses. Table 10 contains a list of controllers and how much VIOP memory the device driver for each controller uses:

Table 10 VIOP memory use per controller

Controller	VIOP Memory Usage (Pages)
ACM-201	6
DKC-203 (1 disk drive)	3
DKC-203 (2 disk drives)	4
DKC-203 (3 disk drives)	5
DKC-203 (4 disk drives)	6
DKC-204 (1 disk drive)	3
DKC-204 (2 disk drives)	4
LAN-007	9
MTC-201 (1 tape drive)	11
MTC-201 (2 tape drives)	16
MTC-201 (3 tape drives)	21
MTC-201 (4 tape drives)	26
MTC-202 (first controller)	8
MTC-202 (each additional)	5

VIOP memory configuration worksheet

Fill out the following worksheet to determine how much VIOP memory a particular configuration will use. With standard VIOP, the maximum amount of memory can not exceed 128 pages. For the RTIOP or VIOP-2, the maximum amount of memory can not exceed 256 pages.

Text/data size

Fill in one line of this section. If you are running the standard ConvexOS viop image, the text/data size for the device drivers is 48 pages. Sysgening in the FDDI driver increases the text/data size by 22 pages. Fill in the proper values for your system.

Device driver size

For each controller in your VIOP, your system will use some amount of VIOP memory. In the chart below, enter the number of controllers for each type in the column labeled "Number."

Once this is done, multiply the figure just entered by the number of pages each particular controller uses. Enter this value in the rightmost column.

Total memory usage

Once you have entered your configuration, you should total the column on the right-hand side. If this total is greater than 128 for standard VIOP or 256 for RTIOP/VIOP-2, your system will not configure properly. You should review the previous section to determine ways of reconfiguring your system to allow it to properly boot.

Table 11 Total memory usage

Overhead	Pages used	Number	Total pages
EGOS/Text/data	-	-	-
Standard viop	48	1	
with FDDI	+22		
Device driver	-	-	-
ACM-201	6		
DKC-203	-	-	-
1 disk drive	3		
2 disk drives	4		
3 disk drives	5		
4 disk drives	6		
DKC-204	-	-	-
1 disk drive	3		
2 disk drives	4		
LAN-007	9		
LAN-208	11-15		
MTC-201	-	-	-
1 tape drive	11		
2 tape drives	16		
3 tape drives	21		
4 tape drives	26		
MTC-202 (first controller)	8		
MTC-202 (each additional)	5		
Total	-	-	
Maximum available for standard VIOP	-	-	**128**
Maximum available for RTIOP/VIOP-2			**256**

VIOP system configuration

With ConvexOS V11.5.1, VIOP memory configurations that have worked properly in the past may no longer configure or boot correctly.

Performance enhancements have been made to the ConvexOS device drivers, especially the FDDI driver. To improve performance, extra VIOP memory has been allocated by these device drivers.

If you have a system that uses too much VIOP memory, you can distribute your VMEbus controllers across multiple VIOPs, taking into account how much memory the device driver for each controller uses. Use the worksheet in the previous section to compute if a particular configuration will configure properly.

If you still have problems with using too much VIOP memory, contact the Technical Assistance Center (TAC).

Example worksheet

This section provides a sample /ioconfig file and a completed worksheet as an example system configuration.

The following is a sample /ioconfig file entry for a VIOP:

```
viop 0
   vme 0
       ctlr MTC-201 csr 0x1000 int 7
         unit 0 type MTD-204
       ctlr DKC-203 csr 0x200 int 2
         unit 0 type DKD-284
       ctlr LAN-007 csr 0xfe00 int 5
         unit 0 type ex
      vme 1
       ctlr MTC-202 csr 0xee00 int 4
         unit 0 subunit 0 type MTD-207
       ctlr MTC-202 csr 0xff00 int 5
         unit 0 subunit 0 type MTD-208
```

With the above /ioconfig file, the worksheet should be filled-in as follows:

Overhead	Pages used	Number	Total pages
EGOS/Text/data	-	-	-
Standard viop	48	1	48
with FDDI	+22		
Device driver	-	-	-
ACM-201	6		
DKC-203	-	-	-
1 disk drive	3	1	3
2 disk drives	4		
3 disk drives	5		
4 disk drives	6		
DKC-204	-	-	-
1 disk drive	3		
2 disk drives	4		
LAN-007	9	1	9
LAN-208	11-15		
MTC-201	-	-	-
1 tape drive	11	1	11
2 tape drives	16		
3 tape drives	21		
4 tape drives	26		
MTC-202 (first controller)	8	1	8
MTC-202 (each additional)	5	1	5
Total	-	-	**84**
Maximum available for Standard VIOP	-	-	**128**
Maximum available for RTIOP/VIOP-2			**256**