

Convergent Technologies

RELEASE NOTICE

for

6.2 S/Series CTIX

(SAC-20Bx, SAC-2014x, SAC-20Cx, SAC-20Fx)

Revised July 18, 1990

Trademark Notice

Convergent Technologies, NGEN, MegaFrame and MightyFrame are registered trademarks of Convergent Technologies, Inc.

WGS/Office, WGS/Desktop Manager, WGS/Word Processor, WGS/Spreadsheet, WGS/Mail, WGS/Calendar, Workgroup Solutions, PC Exchange, S/640, S/480, S/320, S/280, S/222, S/221, S/220, S/120, S/80, MiniFrame, AWS, IWS, Server PC, PT, GT, CTIX, CTIX/386 and CTOS are trademarks of Convergent Technologies, Inc.

CTIX and CTIX/386 are derived from UNIX System V software, under license from AT&T. UNIX is a trademark of AT&T.

١٠٠

١٠١

١٠٢

TABLE OF CONTENTS

SECTION	TITLE	PAGE
1.	General Description of the Release	1
2.	Changes from 6.10 CTIX Release	3
2.1	Origins of the Release	3
2.2	Major Differences from Previous Releases	3
2.3	Other Differences	7
2.4	Other New Features	8
2.4.1	Enhanced Cluster Driver	9
2.4.2	More Open Files Per Process	10
2.4.3	Kernel Support for the SCSI/RS-232 Expansion Board	10
2.4.4	Support for the SCSI Half-Inch Tape Drive	12
2.4.5	S/80 SCSI/Ethernet Expansion Board Support	13
2.4.6	Concurrent Support for RFS and NFS	13
2.5	Bug Fixes	14
2.5.1	Bug Fixes in 6.10 CTIX	14
2.5.2	Bug Fixes from 6.10 CTIX Addendums	14
2.5.3	Bugs Fixed in this Release	15
2.5.4	Detailed Bug Fixes in System V Release 3.2 Kernel	15
2.5.5	Bugs Fixed by the ncc Compiler	17
2.5.6	C Compiler Bugs Fixed in 6.10 CTIX	19
3.	Contents of the Distribution Media	22
3.1	Raw/Update Installation Tape	22
3.2	Maintenance Mode	26
4.	Installation Requirements	27
4.1	Time Required for a Full Installation	27

4.2	Files Required for an Update	
	Installation	27
4.3	Space Required for an Update	
	Installation	27
4.4	Space Required for each GROUP	27
5.	Installation Procedures	30
5.1	Raw Installation	31
	5.1.1 Overview	31
	5.1.2 Preparation for Raw Installation on SMD Disks	33
	5.1.3 Raw Installation Process	34
5.2	Update Installation	40
5.3	Update Silent Mode Installation	46
5.4	Non-CORE Update Installation	47
6.	Using Maintenance Mode	49
7.	Deleting a Group of Files from the System	52
8.	Reconfiguring the System for More or Fewer Users	54
9.	The SCSI Subsystem	56
9.1	S/Series SCSI Hardware	56
9.2	S/Series SCSI Software	56
10.	Notes on System Operation	58
10.1	A Note on System Tuning	58
10.2	A Note on SCSI Half-Inch Tape Drives	58
10.3	Clist Usage	59
10.4	Tty Device Nodes	59
10.5	Process Sizes and Swap Space	59
10.6	Tuning the System for STREAMS Buffers	60
10.7	Configuring the System for the 4K Block File System	61
10.8	Accessing Additional Disk Space on Larger Drives	61
10.9	Re-Partitioning Additional Disks	61
10.10	Some Notes About Cross Development	61
11.	Known Errors, Restrictions, and Warnings	63

11.1	Known Errors	63
11.2	Restrictions	67
11.3	Compatibility Issues	68
11.4	Warnings	70
12.	Documentation Updates	74
12.1	Available Documentation	74
12.2	CTIX Documentation Update	74

•
C

C

C

1. General Description of the Release

Throughout this Release Notice the term **S/Series** stands for the current members of the family: **S/80, S/120, S/221, S/222, S/280, S/320, S/480, and S/640.**

The 6.2 S/Series CTIX ordered by (**SAC-20Bx, SAC-2014x, SAC-20Cx, SAC-20Fx**) contains documents:

<i>Title</i>	<i>Part Number</i>
6.2 S/Series CTIX Release Notice	B-09-01975-01-D
1.13 S/Series Remote I/O Processor Release Notice	B-09-02078-01-B
Programmer's Notes for C Cross-Compilation & Flexnames	B-09-00689-01-B

The 6.2 S/Series CTIX software release comprises 1 Quarter-Inch Tape:

<i>Name</i>	<i>Part Number</i>
Raw/Update Installation Tape	71-03352-01

The 6.2 S/Series package is documented by the following manuals:

<i>Title</i>	<i>Part Number</i>
CTIX Operating System, Ver C, Second Edition, Volumes 1-4	DAC-200
S/Series CTIX Administrator's Guide	DAC-220
CTIX Administration Tools Manual	DAC-210

Additional CTIX documentation that can be ordered:

<i>Title</i>	<i>Part Number</i>
Programmer's Guide: CTIX Supplement	DAC-180
S/Series Diagnostics Manual	DHC-120
S/80 and S/280 Diagnostics Manual	DHC-560
CTIX Network Administrator's Guide	DNC-180

Writing MIGHTYFRAME Device

Note that this manual was written for 5.XX CTIX; however, it is the only reference on the Kernel Debugger, on the use of the Interactive Loader, on integration of customized device drivers into CTIX, and on the Kernel Subroutines available to the writer of device drivers.

AT&T UNIX System V Release 3.2 Documentation:

<i>Title</i>	<i>Part Number</i>
Documenters Workbench Technical Reference	DAU-450
Documenters Workbench User's Guide	DAU-460
Network Programmer's Guide	DAU-430
Streams Programmer's Guide	DAU-410
Streams Primer	DAU-420
Assist User's Guide	DAU-440
User's Guide	DAU-300

2. Changes from 6.10 CTIX Release

2.1 Origins of the Release

This release of CTIX corresponds to AT&T UNIX™ System V Release 3.2, with the following exceptions:

1. The new AT&T 2K filesystem is replaced by a CTIX 4K filesystem.
2. New AT&T functionality for lp(1), access to printed forms and easier administration of lp interface filters, is not provided.
3. The new AT&T master/slave stream pipe driver is not included in 6.2 CTIX. The comments in the source code indicated that this driver will not be present in future releases.

2.2 Major Differences from Previous Releases

This section briefly states the major differences between the 6.10 and 6.2 S/Series CTIX release.

1. Single binary tape:
 - This means all commands "*do the right thing*" on all S/Series machines.
 - There are two kernels, one for S/640¹ and one for S/320² and S/80.³ There is a loader on the tape that can pick which kernel to bring off the tape.
 - A single install tape requires a single loader that knows about all machine architectures. The size of this loader will be much larger than the current maximum

-
1. For the purposes of this document, there are only four different machine architectures, and S/640 will imply S/480 and S/640, as these have the exact same software interface.
 2. Likewise, S/320 will imply S/120, S/22x, and S/320, as these have the exact same software interface.
 3. In most contexts S/80 will imply S/80 and S/280, as these have nearly the same software interface. In cases where there are differences, the specific machine name will be used.

size allowed on the disk (30KB). To work around this limit, there will be three different loaders:

- i. The "knows all" (normal) loader for the install tape, and machines with room in partition 0 for a larger loader. This loader has all presently known boot device drivers built in. Its size is 38K.
- ii. A small version of the loader (without tape drivers), for the boot disks with only 30KB of space for the loader. Its size is 30K.
- iii. The interactive version of the loader, which will be another "knows all" loader. Its size is 40K.

System disks with less than 30 KB of space for a loader must be raw installed. The partition 0 size must increase by 98 KB to accommodate the new disk description file in /usr/lib/iv. The disk description files in /usr/lib/iv have been changed to specify 128KB for the loader area.

When raw installing, the normal "knows all" loader is used, and a larger area for the loader is reserved in partition 0. When updating, if there is space in partition 0 (maybe by stealing blocks from the dump area if present), the normal loader is installed. If space cannot be found, then the small loader must be used.

- As a side-effect of the single binary tape, the stock kernel provided upon installation includes drivers for devices that might not be present on a given machine. For example, the scsitape driver and the QIC-2 driver are both present, but the QIC-2 hardware is present only on the S/320. Also the V/SMD3200 SMD driver is included. Optimum use of memory would require using the KFIG Group to build a new kernel with an appropriate dfile for that machine.
2. Support for S/280: *hinv(1)* and *syslocal(2)* are enhanced to allow user programs to configure the S/280.

3. A STREAMS-based tty driver is provided, and is now the default tty driver. Initial tests show that for certain situations, the STREAMS-based tty driver is faster than the non-STREAMS driver.

The corresponding STREAMS-based sxt and iop drivers are also provided. **Note that the sxt driver (shell layers) does not work on RIOP ports.**

4. The old wxt driver, provided with CTAM, does not work, except on Convergent cluster ports, including RIOP ports, until a new version is released which supports STREAMS.
5. The old and new STREAMS-based SLIP drivers are provided. The RIOP tty driver is not yet converted to STREAMS. *slattach(1M)* can detect which tty driver is in use and call the appropriate SLIP driver.

You must save your old */etc/netcf* customizations and add them to the new version of the file **or SLIP may fail.**

6. Entries for the Convergent Technologies DC2400, a Hayes Smart Modem 2400 compatible, are now included in */usr/spool/uucp/Dialers*.
7. *adman(1)* now allows configuration of any one Ethernet controller; the controller number is not limited to **en0**. You must configure manually for multiple Ethernet controllers.
8. The installation of the CROSS Group now works correctly. It will overwrite the SOFTWARE Floating Point libraries in */usr/lib* if you request it. The installation will also remove */cross* if you request that be done.
9. Selected AT&T System V Release 3.2 features as follows:
 - New filesystem size: Instead of AT&T's 2K filesystem, 6.2 S/Series CTIX incorporates a 4K filesystem. **The root filesystem cannot be a 4K filesystem**, because the default kernel has no 4K buffers configured for it. Note that to enable 4K filesystem mounts, the kernel must either be relinked or reconfigured using *uconf(1M)*. *mkfs(1M)* has a new

option for creating 4K file systems. Refer to the new man page.

- An enhanced *curses(3X)* library is included. It provides improvements to support color terminals. Support for the CT local features for line drawing has been dropped. The non-standard keywords **ldul**, **ldur**, **ldht**, **ldbr**, **ldbl**, and **ldatt**, have been deleted. Equivalent features are now supported by *curses*, using the standard construct **acsc**.

The *terminfo* description file keywords **goto_func_line** and **leave_func_line**, added to support Convergent PT and GT terminals, have been deleted. Programmers who need to address the PT/GT status line should use the AT&T standard keywords **to_status_line** and **from_status_line**.

- Enhanced security features are included. The shadow password file, a program *pwconv(1M)* to convert to using the shadow file, and a program *pwunconv(1M)* to revert to not using the shadow file are provided. The *rlogin(1)* and *telnet(1)* commands are also enhanced to handle the shadow password file.
- The new command *passmgmt(1M)* is provided which helps super-user shell scripts to make controlled changes to the password files **/etc/passwd** and **/etc/shadow**.
- The uucp group has been added to the file **/etc/group**. It is the group ID of most uucp files and commands.

10. New features in the AT&T System V Release 3.2 kernel:

- If you are removing a directory in a **S51K** file system and the sticky bit is set on the parent directory, then this directory is deleted only when:
 1. The deletion is performed as superuser.
 2. The deletion is performed as owner of the directory.
 3. The deletion is performed as owner of the parent directory.

4. The deletion is performed when the directory is writable.
- When the user executes *chmod(1)* to turn on the sticky bit in a **S51K** file system and is not **root**, *chmod(1) + t* is allowed on directories only.

2.3 Other Differences

Other changes from 6.10 CTIX are listed in this section.

1. New features from 6.10 CTIX Addendums:
 - Changes to kernel file system services to support NFS (without Yellow Pages).
 - For the S/320 and S/640, support for the SCSI/RS-232 Expansion Board.
 - For the S/320, there is a new version of the QIC-2 Tape driver.
 - Changes to *scsimap(1M)* and *adman(1M)* to recognize SCSI disks and the SCSI/RS-232 Expansion Board.
 - Fixes for the VME-based Half-Inch Tape driver.
 - Support for the Cipher SCSI Half-Inch Tape Drives, F880S and M990S.
 - S/80 SCSI/Ethernet Expansion board support, with a new scsi driver and system utilities.
 - Changes to *adman(1)* for recognition of the S/80 RS-232 Expansion board and for correct modifications of */etc/inittab* entries.
2. Other new features:
 - Support for the MC68882 Floating Point Co-processor on the S/280.

The MC68882 is equivalent to the MC68881. If the MC68882 is present on an S/280, *hinv(1M)* reports that an MC68881 is present. If you use the output of *hinv(1M)* in a script, take this into account.
 - Throughput improvements to IOP that increase total character throughput on transmit from 20K chars/sec

to about 24K chars/sec. Throughput on receive has yet to be measured, but should be similar.

- The new command *serstat(1M)* now supports IOP controlled ports.
- *fuser(1M)* and *fusage(1M)* are enhanced to support NFS.
- *cu(1)* and *uucp(1)* now run over TLI connections.
- The size of the callout table is now configurable by *uconf(1M)*. The variable name is **v_call** for the callout table size.
- There is a new *adman* capability, *createdev(1M)* in **/usr/local/bin** that will create device nodes given the type of device, *major* device major number, and starting *minor* device number. Refer to the new man page *createdev(1M)* for more information.

2.4 Other New Features

What follows are brief descriptions of other new or different features. Most were available in 6.10 CTIX and are included here for those upgrading from a 5.XX CTIX release.

1. A new program, *tio(1)* has been added, which uses several buffers in order to keep the tape device streaming. When *tio(1)* is used with a Half-Inch Tape drive, the Interphase V/Tape 3209 VME controller board must be at the currently shipped Revision Level for best performance.
2. If your system supports the VME bus and is using the Interphase V/SMD 3200 disk controller, the default DMA burst count has been changed from 8 to 2 words. This parameter, **gdvms32dmaburst**, is also now configurable in **/etc/system** using *uconf(1M)*. This change was made to minimize a problem with the CMC Ethernet VME Controller.
3. The gap sizes for SMD disks on the Interphase V/SMD 3200 controller have been changed from Gap1 = 16 and Gap2 = 16 to Gap1 = 12 and Gap2 = 12 in the disk description files in **/usr/lib/iv**. It is advisable, especially for the Northern Telecom 750MB drive, to reformat your SMD disks using the new gap sizes. (Be sure to backup

your disks before doing the format.)

2.4.1 Enhanced Cluster Driver

There are some enhancements to the Terminal Cluster Driver.

The first change is in the way `/dev/tty n n` numbers are assigned to RS-422 lines. In `/etc/inittab` in 6.10 and 5.XX CTIX, there are eight `/dev/tty???` device nodes per RS-422 line.

line	0	1	2	3
<code>/etc/inittab</code>	256-263	264-271	272-279	280-287
<code>cl_defdrops=8</code>	256-263	264-271	272-279	280-287

Prior to this release, if you changed `cl_defdrops` (drops per line) in `/etc/system`, the `/dev/tty n n` numbers were not correct. If you had changed the drops in `/etc/system` to two, the `/dev/tty n n` numbers would have been assigned as follows:

line	0	1	2	3
<code>cl_defdrops=2</code>	256-257	258-259	260-261	262-263

The new assignments are as follows:

line	0	1	2	3
<code>cl_defdrops=1-8</code>	256-263	264-271	272-279	280-287
<code>cl_defdrops=9-15</code>	256-263 288-294	264-271 296-302	272-279 304-310	280-287 312-318

Note:

1. ttys numbers 295, 303, 311, and 319 cannot be assigned or opened.
2. the administrator will have to add entries to `/etc/inittab` when `cl_defdrops` is greater than 8.

The second change is for support of RS-422 high speed devices such as PC Exchange/Vines, RIOP, and Telecluster. These run at 1.8M baud versus the PT/GT rate of 307K baud.

Note:

1. Each line on the RS-422 Board will auto-configure to the highest speed requested for any one connected device.
Do not mix high and low speed devices.

2.4.2 More Open Files Per Process

The number of open files per process has become configurable in the range of 20 to 100. The default is 20. To generate a kernel with a larger number of open files specify **nofile** in your `/usr/sys/cf/dfile`, and do a `make(1)`; refer to the *S/Series CTIX Administrator's Guide*.

Note: Any program compiled in 5.XX CTIX that uses fields of the user structure that occur after the open file table (such as `u_arg[]`, `u_signal[]`, or `u_ar0`) must be recompiled, since the offsets of the fields in the user struct have changed from those of 5.XX CTIX and earlier.

Warning: The **stdio** library still has a limit of 20 open files: although you may have 100 file descriptors for open files, at most 20 of them can be acquired by using `fopen(3)`. Attempting to open the twenty-first file, `fopen(3)` will return the value 24 in **errno**, which is **EMFILE**, *Too many open files*.

2.4.3 Kernel Support for the SCSI/RS-232 Expansion Board

There are changes to the SCSI driver for SCSI/RS-232 Expansion Board support.

In order to configure the system for SCSI/RS-232 board(s), entries in the **!SCSIMAP** section of `/etc/system` must be added to use the SCSI devices.

To configure your system for the SCSI/RS-232 upgrade, take the following steps:

1. Shut down the system and install the SCSI/RS-232 board into the first open I/O slot.
2. Power up and reboot the system. Login as **root**.
3. Execute: **hinv -p**.

If the additional board is not seen, check the number of controllers using the DMA channels. Note: The bus

number to use for SCSI entries is equal to the reported slot number.

4. Edit the **!SCSIMAP** section of **/etc/system**.
(Note: ALL disk entries are "disk-c0dN". ALL tape entries are "tape-dN".)
5. Create device node entries in **/dev/rdisk**, **/dev/dsk**, and **/dev/rmt** using **mknod(1M)**. See the Subsection on "Creation of SCSI Disk Device Nodes" for creation of disk device nodes. If your system is an S/320, see the later Section on "Support for the SCSI Half-Inch Tape Drive" for creation of tape device nodes.
6. Execute: **scsimap -u**
Execute: **hinu -p**
(Note: All SCSI disk units should be displayed. No SCSI tape units are shown by **hinu(1M)**.)

See the *S/Series CTIX Administrator's Guide* for details on configuring SCSI disks.

If your system is an S/320, note that there are only 2 DMA interrupt channels on the system. This means that the SCSI/RS-232 Board cannot be supported if a VME Interface Board and an IOP/RS-422 Board are already installed. If there is no VME Interface Board, the SCSI/RS-232 Board chooses the VME DMA channel, otherwise it takes the IOP/RS-422 channel. If either a VME Interface Board or an IOP Board is already installed, then the system cannot support a mix of COMBO (SCSI/RS-232 or SCSI/Ethernet) boards. The first type of COMBO board is recognized and the second is ignored.

If your system is an S/640, there are 3 DMA channels on the system. The SCSI/RS-232 Board always takes one of the non-VME DMA channels.

The default entries for onboard SCSI-based systems in **/etc/system** look like:

```
!SCSIMAP
disk-c0d0    bus=0 target=6 lun=0 parity reselect
tape-d0      bus=0 target=1 lun=0 parity reselect
```

The S/320 systems do not have an onboard SCSI bus, so the first bus number that can be used is 1. The default entries for S/320 systems in `/etc/system` look like:

```
!SCSIMAP
disk-c0d0   bus=1 target=6 lun=0 parity reselect
tape-d0     bus=1 target=1 lun=0 parity reselect
```

Warning: for **bus 0** targets **4** and **5** must be disks, not tapes, otherwise the system will not boot.

2.4.4 Support for the SCSI Half-Inch Tape Drive

The Cipher SCSI Half-Inch Tape drives are supported.

In order to configure the system for SCSI Half-Inch Tape drives, entries in the SCSSIMAP section of `/etc/system` must be added.

To configure your system for additional drives, do the following:

1. Shut down the system and power it off.
2. Add the new drives (disk or tape).
3. Reboot the system and login as **root**.
4. Execute: **hinv -p**
All SCSI disk units should be displayed. No SCSI tape units are shown by `hinv(1M)`. You should see the disks reported.
5. Edit the SCSSIMAP section of `/etc/system`.
(Note: ALL disk entries are "disk-c0dN". ALL tape entries are "tape-dN".)
6. Create device node entries in `/dev/rmt` using `mknod(1M)`. See the later Section on "Creation of SCSI Half-Inch Tape Device Nodes" for creation of tape device nodes.
7. Execute: **scsimap -u**
Execute: **hinv -p**
8. Add entries for the Half-Inch Tape drive, using the default `/etc/system !SCSSIMAP` section lines as templates. Be sure the target number is different for each

drive on the same bus. *hinv(1M)* reports which slot number the board is in, therefore which bus number to use. The default entries plus the Half-Inch tape drive look like:

```
!SCSIMAP
disk-c0d0 bus=0 target=6 lun=0 parity reselect
tape-d0 bus=0 target=1 lun=0 parity reselect
tape-d1 bus=0 target=2 lun=0 parity reselect
```

Note that an additional field, **halfinch**, is added to the end of the line for the entry for Half-Inch Tape drive, to distinguish Half-Inch from Quarter-Inch Tape drive.

9. 6.2 CTIX does not have device nodes provided for SCSI tape drives for the S/320. You must make nodes for the tape drives, which are on controller 2. Refer to Chapter 7, "*Configuring Device Drivers and the Kernel*", in the *S/Series CTIX Administrator's Guide* for more information.

2.4.5 S/80 SCSI/Ethernet Expansion Board Support

S/80 SCSI/Ethernet Expansion Board support is provided, with a new SCSI driver and system utilities. There is a new version of the enet driver, **mitienet.o**, and there are changes to the general disk, SCSI disk, and SCSI tape drivers. A new ioctl command, **SCSIENET**, is added to support SCSI/Ethernet. The corresponding modules in kernel libs in **/usr/sys** are changed. In **/usr/include/sys**, system header files **socket.h**, **scsi.h**, and **if_enp.h** are changed, and **se_cmd.h** is new. System command *hinv(1M)* is updated, and system initialization files **/etc/init.d/devices**, **/etc/rc2.d/S79devices**, and **/etc/enp/enpload** are also changed.

NOTE: be sure that **/etc/hosts** and **/etc/rcopts/NODE** are correct after the installation is finished.

Refer to Chapter 7, "*Configuring Device Drivers and the Kernel*", in the *S/Series CTIX Administrator's Guide* for more information.

2.4.6 Concurrent Support for RFS and NFS

6.2 CTIX now supports RFS and NFS running concurrently. To use RFS and NFS together, you must be running 1.1 RFS.

2.5 Bug Fixes

2.5.1 Bug Fixes in 6.10 CTIX

The following bugs have been closed since 5.XX CTIX. Corrections or the proper feature functionality are described.

1. **Getty** recognizes B19200 and B38400 in the `/etc/gettydefs` file.
2. The download image for the IOP Board has been modified so that the system does not hang under high character throughput levels.
3. **stty 0** hangs up the phone line immediately, as it should.
4. When shell commands are typed on a system with an IOP while output is coming to the terminal, the shell does not hang and become unkillable.
5. **tidc(1)** has been replaced by **infocmp(1)**, which correctly decompiles supported *terminfo* files.
6. Pressing the BREAK key while in *cu(1)* does not cause every other character to be read in.
7. The installation update script now asks whether the user wants to save `/etc/termcap`.
8. **cs**h scripts run in the background are no longer killed by hitting the DELETE key.
9. **getty** and **uugetty** now correctly recognize the BREAK character.

2.5.2 Bug Fixes from 6.10 CTIX Addendums

Bug fixes from 6.10 CTIX Addendums:

1. A fix to the kernel Direct I/O.
2. A new version of *iv(1M)* that formats SCSI disks in the manufacturer's defaults regardless of the disk description file.
3. Changes to SCSI tape driver error handling, so that when a filemark is encountered and any recoverable errors occur, the driver does not miss the filemark.

4. There is a new version of the kernel ST506 disk driver. It fixes a problem with Hitachi disk drives/by suppressing error reporting when the error W_ABC (Command Abort) is received. A retry then takes care of the actual disk transfer.

2.5.3 Bugs Fixed in this Release

The problems that have been fixed since 6.10 are listed below:

1. **SPR 15526:** *rmail(1C)* is fixed. If *rmail* did not see a From_ line, *uf* was uninitialized but was referenced. If there is no From_ line, *rmail* now gets the From_ line from \$LOGNAME. If LOGNAME is not set, *rmail* then tries *getlogin()*, then *getpwuid()*.
2. *mailx(1C)* now works for sending mail to local users.
3. The utility *shutdown(1M)* has been changed. It now tests for incorrect command name invocation at the start of the script and it checks whether there are at least 16 free blocks in / before continuing. Also upper case input is now handled correctly.

If *shutdown(1M)* exits with a warning about lack of free space in root, it is necessary to remove insignificant files before continuing. It is wise to have more than 100 blocks free, but 16 will suffice for shutdown to occur.

2.5.4 Detailed Bug Fixes in System V Release 3.2 Kernel

Detailed bug fixes in the System V Release 3.2 kernel:

1. On a mount of an **S51K** file system, the updating of the *u* struct occurs **after** all error condition tests have been passed.
2. The *stats* system call on a **S51K** file system which is not mounted now checks to see if the device number in the inode struct is in range before trying to read in the superblock.
3. When freeing storage from a specified inode in the **S51K** file system, the calculation of the index into the inode's block list of the last block is now correctly calculated.

4. When getting directory entries from a **S51K** file system, the inode access time is now allowed to be updated.
5. When obtaining a physical disk block number on a device given the inode and the logical block number in a file, there is a scan of the mount table to see if this block device is mounted. During this scan of the mount table, a check is now done to see if the size from the mount table (the source) is greater than the size in the inode struct (the destination). If it is not then ignore the mount table entry.
6. When doing a **clone open**, CTIX now protects against another **clone open** on a different inode.
7. When a STREAMS module gets popped off the stream and it's not an ordinary release and it's not a disconnection request, **don't** send a disconnection request back up to the application.
8. In the UNIX kernel STREAMS module read queue put procedure, if there is a message type of **T_ERROR_ACK** or **T_OK_ACK**, send an **M_ERROR** message back up to the application.
9. When doing raw I/O, RFS now checks to see if the byte count requested is greater than **NPGPT*NBPP**. If so, then loop doing blocks of **NPGPT*NBPP** bytes until the requested byte count is satisfied.
10. In the UNIX kernel idle loop, CTIX will now give higher priority to the STREAMS processing when there are no runnable processes [see **swtch()**].
11. The following STREAMS functions, are now protected from interrupts:

 freemsg(), **dupb()**, **dupmsg()**, **copymsg()**, **linkb()**,
 unlinkb(), **rmvb()**,
 flushq(), **insq()**, **strinit()**, **noenable()**, **enableok()**.
12. When doing the steams functions **pullupmsg()** and **adjmsg()**, CTIX will now protect against possible zero address references.

13. When doing a **streams open** call, the **streams** table is checked to make sure that there are not two streams to one device. This closes the timing window between opening and closing a stream.
14. When doing a **streams clean** call, if the inode byte count and the request byte count passed to the clean routine are equal to **one**, then CTIX decrements the **streams** table entry's reference count by one. This closes another timing window the between opening and closing of the last open of a stream.
15. When doing a **streams read** call, if the **FNDELAY** flag is set then CTIX does not cause the **streams read** to fail when there is nothing available from the read queue.
16. Now allow **M_COPYIN** and **M_COPYOUT** cases for **STREAMS** driver or module ioctl data within the UNIX kernel.
17. CTIX does not allow an RFS *link* system call to be handled by the RFS **server** process. Symbolic links are not supported by RFS.
18. The signal **SIGWINCH** has been integrated into signal processing.
19. CTIX allows a null string from a just received block of data from the network in the name server process.
20. A *typedef* bug in the network services library related to the connection ID was fixed.

2.5.5 Bugs Fixed by the ncc Compiler

In 5.25.2 CTIX and in 6.10 CTIX, a new C compiler, *ncc*, was introduced to fix a number of problems. This compiler was given a new Group, **NCC**. For the benefit of those who are updating from a 5.25.1 CTIX release, the compiler fixes are reiterated in this Release Notice.

- **SPR 2571:** When compiled with the **-g** option, the compiler would sometimes take the size of a function, causing an "internal compiler error" message.

- **SPR 12819/13543:** The compiler generated bad code for an expression which dereferenced a **double** array element. The generated code used a data register for two different purposes at the same time.
- **SPR 10714:** Compiler was generating a **bfin**s instruction with the source operand in an address register, which is illegal.
- **SPR 11638:** The compiler generated a "cmp.w" instruction for a comparison involving byte variables.
- **SPR 12134:** The compiler generated bad code when using variables of type **enum** as array indexes, causing incorrect offsets.
- **SPR 13039/12771:** The compiler did not generate correct code for the following:

```
char *foo(); return (foo()[i]);
```

- **SPR 11763:** The compiler would generate the following error message for some code segments:

```
"compiler error: cfix trouble"
```

- **SPR 12610:** An error message was incorrectly generated for a shift involving a variable of type "enum".
- **SPR 12611:** The following program generated the error "compiler error: bad bigsize 012" on the cast:

```
enum col { red, green, blue };
main() {
    char str[3];
    enum col e = red;
    str[0] = (char) e;
}
```

- **SPR 12648:** The following code caused the compiler error, "can't deal with op AUTO":

```
char x == (1 && 1) ? 1: 1;
```


- **SPR 13065:** The 'cc' command used the wrong libraries when profiling is used even though LIBROOT and CENVIRON were correctly set.

2.5.6 C Compiler Bugs Fixed in 6.10 CTIX

A number of problems with the 5.25.1 CTIX C compiler were fixed in 6.10 CTIX. These were fixes to libraries that could not be back-ported to 5.XX CTIX due to source licensing restrictions. For the benefit of those who are updating from a 5.XX CTIX release, the compiler fixes are listed in this Release Notice.

- The compiler could not deal with complex expressions such as the following:

```
int b, c, f; char *d;
b ? (c ? (*d == 0) : f) : e() ;
```

This expression caused an "internal compiler error" message.

- The compiler generated a warning when an **enum** was used in a **switch** expression.
- The optimizer could generate a jump subroutine with a data register as the address, which is illegal. This caused an assembler syntax error.
- The compiler front end did not pick up an illegal use of a structure in an **if** expression. This caused an "internal compiler error" message.
- The compiler front end did not pick up an illegal use of a structure in a conditional expression. This caused an "internal compiler error" message.
- The optimizer generated an incorrect short multiply instruction between an address and data register, causing an assembler syntax error.
- The compiler generated a cryptic error message for the following, incorrect code:

```
int f(a); f(b) { a=b; }
```

- The compiler incorrectly evaluated expressions of the form:

```
if ((signed == unsigned) < 0)
```

The compiler removed this code, under the false assumption that the expression could not be less than zero.

- The compiler generated bad code for bit field comparisons in which the high bit of the field is set.
- The compiler would fail to compile an expression of the following form:

```
double *a, *b, *c; *a = *b = *c;
```

- Some code segments exist which could make the optimizer go into an infinite loop.
- The compiler would generate bad code for the following:

```
short a = -1;  
unsigned long l;  
l = ((unsigned short) a);
```

The long variable contained -1, when it should have contained 0xffff.

- The compiler would generate incorrect code for float to unsigned conversions, double to unsigned conversions, unsigned to double conversions, and unsigned to float conversions, where the number being converted was beyond the range of a positive signed integer.
- Bit fields were sometimes improperly typed to a signed type rather than an unsigned type. Then, for a conditional test on a bitfield, code could be generated which would test the 'N' (negative) bit of the condition code register. This was incorrect, since an unsigned type can never be negative.
- The 5.2X CTIX version (based on AT&T System V Release 2.2) of the Section 3C library lacked the subroutines to manipulate the hardware Floating Point chip. 6.2 CTIX contains those routines: *fpgetmask(3C)* and *fpsetmask(3C)*. These routines should be used before relying on Floating

Point Exception detection, because they set the mask that controls output of the Exception from the Floating Point coprocessor.

- Correct code which contained bit fields which were 32 bits wide could cause the compiler to generate the following cryptic compiler message:

Compiler Error: line 97: compiler error: sconv:t 05.

- The compiler would not increment the register short pointer in the following example:

```
proc()
{
    register short *p = val;
    ...
    func(*p++);
    ...
    func2(p);    /* any reference to p */
}
```

This problem occurs only with register short pointers; any other combination of unsigned, or type length produced the proper code.

3. Contents of the Distribution Media

The 6.2 S/Series CTIX distribution comprises 1 Quarter-Inch Tape:

- Raw/Update Installation Tape (Part # 71-03352-01)

3.1 Raw/Update Installation Tape

The Raw/Update Installation Tape contains all the files necessary for either initial or update installation. Files numbered 5 and greater are *cpio(1)* archives in *-Qc* format.

<i>File #</i>	<i>Description</i>
0	VHB and loader
1	stand-alone UNIX kernel
2	stand-alone file system (reserved area 1)
3	minimum file system (reserved area 2)
4	download images
5	installation tools
6	installation control files for CORE Group
7	customizable files for CORE Group
8	required files for CORE Group
9	installation control files for ADMAN Group
10	required files for ADMAN Group
11	installation control files for LP Group
12	required files for LP Group
13	installation control files for UUCP Group
14	customizable files for UUCP Group
15	required files for UUCP Group
16	installation control files for RIOP Group
17	required files for RIOP Group
18	installation control files for LAN Group
19	required files for LAN Group
20	installation control files for SYSM Group
21	customizable files for SYSM Group
22	required files for SYSM Group
23	installation control files for ACCT Group
24	customizable files for ACCT Group
25	required files for ACCT Group
26	installation control files for BCE Group
27	required files for BCE Group
28	installation control files for ACE Group

29	required files for ACE Group
30	installation control files for DOC Group
31	required files for DOC Group
32	installation control files for KFIG Group
33	required files for KFIG Group
34	installation control files for GAMES Group
35	required files for GAMES Group
36	installation control files for TINFO Group
37	required files for TINFO Group
38	installation control files for HELP Group
39	required files for HELP Group
40	installation control files for ASSIST Group
41	required files for ASSIST Group
42	installation control files for CROSS Group
43	optional files for CROSS Group
44	required files for CROSS Group
45	installation control files for MISC Group
46	required files for MISC Group
47	installation control files for SMAIL Group
48	customizable files for SMAIL Group
49	required files for SMAIL Group
50	installation control files for NCC Group
51	required files for NCC Group

This release is structured to allow installation of one or more selected subsystems, or groups of commands. A brief description of each of the 20 groups follows.

<i>Group</i>	<i>Description</i>
CORE	required CTIX CORE commands. This group is mandatory and is treated specially in that the system is automatically rebooted upon installation completion.
ADMAN	Menu-oriented System Administration tools.
LP	Line Printer spooling system.
UUCP	UNIX system to UNIX system CoPy programs.
RIOP	Remote IOP system.
LAN	Local Area Network protocol independent programs.
SYSM	SYStem Maintenance commands including system profiling and file system repair.
ACCT	ACCounTing system which reports CPU usage,

user disk usage, command usage and more.

BCE Basic 'C' Environment, including the 'C' compiler, loader, assembler, libraries. This suite is configured to generate code for a 68020 CPU with software floating point.

ACE Advanced 'C' Environment: SCCS, lex, yacc, awk, profiling, lint, etc.

DOC Documentation and text formatting tools: nroff, spell, terminal filters, fonts, etc.

KFIG Kernel conFIGuration files and commands.

GAMES CTIX GAMES.

TINFO Terminal INfOrMation files in /usr/lib/terminfo.

HELP The HELP sub-system.

ASSIST Screen oriented command interpreter.

CROSS 68020-68881 CROSS-development environment used to generate code for the 68881 floating point hardware on your own machine, if you have it; or to generate code for machines other than your own.

MISC MISCellaneous CTIX commands.

SMAIL Sendmail and associated support files.

NCC C compiler with several bug fixes.

To obtain an actual list of the contents of the cpio archives on the Raw/Update Installation Tape, execute:

```
tsioctl -c rewind /dev/rmt0
tsioctl -c skip /dev/rmt4 5
while true
do
    /bin/cpio -icvQt < /dev/rmt4
    if [ $? = 2 ]
    then
        break
    fi
done
```

3.2 Maintenance Mode

The 6.2 S/Series CTIX Raw/Update Installation tape provides the maintenance tape function of previous releases. Maintenance mode is entered by choosing the *maintenance* option when the install tape is booted.

4. Installation Requirements

4.1 Time Required for a Full Installation

The installation of all Groups will take about thirty minutes.

4.2 Files Required for an Update Installation

There are no required files to perform an update installation. The installation tape is completely self contained.

Note that the installation renames each file in */etc/rcopts* to 'OLDfilename' before it loads new files into */etc/rcopts*. Save your customizations and apply them to the new files.

When updating from 5.25 or 6.10 CTIX, you will see an error message about *make(1)* not knowing how to make *instclobber*. Ignore the message.

4.3 Space Required for an Update Installation

Installation of the CORE Group requires approximately **2048** free blocks on the root file system at the time of installation script invocation. The command:

```
df /
```

reports the number of free blocks on the root file system. If there are not **2048** free blocks, clean up the file system to obtain them before executing the installation script. This space is required to save necessary files and to read in the customizable files.

4.4 Space Required for each GROUP

The entire 6.2 CTIX release uses **53600** 512-byte disk blocks; **20800** of which are in */*, and **32800** of which are in */usr*. The number of blocks and inodes each Group uses is tabulated here:

Group	root	rootinodes	usr	usrinodes	total
CORE	7110	1282	3660	180	10770
ADMAN	10	4	5224	161	5234
LP	0	0	738	75	738
UUCP	0	0	1492	86	1492
RIOP	376	20	0	0	376
LAN	814	22	0	0	814
SYSM	638	26	1530	47	2168
ACCT	74	11	384	40	458
BCE	3384	35	2542	314	5926
ACE	168	8	4518	108	4686
DOC	0	0	2554	167	2554
KFIG	114	3	2376	18	2490
GAMES	0	0	1194	53	1194
TINFO	0	0	2208	1227	2208
HELP	6	2	850	35	856
ASSIST	0	0	2376	223	2376
CROSS	7022	103	0	0	7022
MISC	118	5	768	27	886
SMAIL	34	4	284	14	318
NCC	954	4	0	0	954
TOTAL	20822	1529	32698	2775	53520

Installing all groups leaves very little free space in the **root** and **usr** (if separate) filesystems. An easy way to save some of the space in the root filesystem is to separately mount **/cross** before installing the CROSS Group.

The groups can be clustered together into several categories:

Basic: CORE ADMAN LP TINFO MISC
 necessary for minimal system operation

Network: UUCP LAN SMAIL (1)
 required for electronic mail

SGS: BCE ACE CROSS (2)
 required for software development

Sysadm: SYSM ACCT HELP KFIG (3)
 useful for system administration

Option: RIOP
required for RIOP support

Misc: GAMES ASSIST
miscellaneous software groups

There are some dependencies, indicated by the numbers:

1. SMAIL expects that UUCP has been installed.
2. BCE must be installed before ACE, CROSS, or NCC.
3. KFIG depends on BCE being installed.

5. Installation Procedures

If your system has 5.XX CTIX installed, be aware that system initialization has changed (new scripts, directories, etc). If you are doing an update instead of a raw install, learn the differences before proceeding. Refer to the man pages for *rc0(1M)*, *rc2(1M)*, and *rc3(1M)*, in the *CTIX Operating System Manual, Version C*. In order to get the new *ctinstall(1)*, use the 5.XX CTIX installation procedure:

```
tsioctl -c skip /dev/rmt4 5
cpio -icvQdum install/ctinstall < /dev/rmt0
/install/ctinstall
```

There are three modes of installation described here: Raw, Update, Update Silent Mode. These correspond to the *ctinstall(1)* options **install**, **update** and **update in silent mode**, respectively.

Do not attempt to update CORE using update from single user level.

- If your disk does not already have a running CTIX system on it, perform a **raw** install.
- If your system is already running CTIX, you may perform an **update** install.
- If your system is already running CTIX and it is acceptable that all existing customizable files (such as */etc/passwd*) are renamed to OLD <filename> and replaced with the release versions, then perform an update in **silent mode**, which does not ask you many questions.

WARNING: If your installation fails during the CORE Group, the file */etc/inittab* has been removed and the */usr* file system may not be mounted properly. To correct this situation, do the following:

1. Execute **umount /usr**
2. Create */etc/inittab* and */etc/mnttab*, by executing **> /etc/inittab ; > /etc/mnttab**

3. Execute **rm -rf /install/bin**
4. Reboot.

If your installation fails for some reason, **do not use the restart capability of *ctinstall(1)***. *restart* cannot handle all cases of install failure and your root file system might end up in a strange state. Instead, fix the problem that caused the failure and start the entire installation over again.

A list of installed files does not appear on the screen.

Remember that the new loader has size greater than 30K and that the new disk description files have been changed to reserve 128K for the loader in the system disk slice 0.

Changing to a single binary tape for all S/Series systems has required providing all files for each type of specific machine. For some machine types, not all provided files are needed and can be removed to gain some disk space. The following files are among this group:

1. For the S/80 and S/280, there is no RS-422 cluster line support, so the files **/usr/lib/iv/ws*** may be removed.
2. For the same reason, the **RIOP** Group should not be installed. Use the *ctinstall(1)* **omit** mode for installation of groups.

If you should respond to the prompt for Group selection with only **RETURN**'s or with incorrect responses, which causes *ctinstall(1)* to prompt again for selection, the list of Groups selected will end up having multiple entries for some or all Groups. *ctinstall(1)* goes through the list and discards multiple instances of each Group it installs.

5.1 Raw Installation

5.1.1 Overview

6.2 S/Series CTIX can be installed on ST506, SCSI, or SMD based systems. (Note: SMD disks are only supported on systems which have a VME bus.) Northern Telecom SMD disks are supported. Installation on SMD disks requires several preliminary steps.

Note: Raw installation initializes the system disk, and you will be able to change the disk configuration (number and size of partitions, separate /usr partition, etc). **The installation does not check the current disk partitioning.** You must use the default disk description file parameters supplied with the release or you must request modification. If necessary, make note of your current disk configuration using `iv -t`. The drive should be formatted and have a valid bad block table and Volume Home Block.

If you plan to install and use TCP/IP and RFS or NFS, you should not use the default sizes. Use the following sizes instead:

<i>Partition</i>	<i>Logical 1024-byte blocks</i>	<i>Physical 512-byte blocks</i>
slice 1	17600	35200
slice 2	8000	16000
slice 3	24000	48000

Otherwise you must omit some groups, in order to have enough space.

Note that to boot off an SMD disk requires a boot PROM with a revision level of 1.2 or above. Also note that for S/640 machines to boot off the SCSI disk (onboard SCSI bus only) requires a boot PROM with a revision level of 3.0 or above.

The loader attempts to boot off the following devices in order:

1. Quarter-inch tape
2. ST506 disk
3. SMD disk
4. SCSI disk

Please refer to the Section *"Initializing and Configuring Disks"* in Chapter 8, *"Disks and File Systems"*, of the *S/Series CTIX Administrator's Guide* if you want to add disks other than those shipped with the system.

5.1.2 Preparation for Raw Installation on SMD Disks

If your system does not have SMD disks or you are not installing on an SMD disk, you may skip this section. Before 6.2 S/Series CTIX can be installed for the very first time on an SMD disk, the following preparatory steps must be performed.

1. Go into maintenance mode after booting from the tape. Refer to Section 6, "Using Maintenance Mode" for a description of the Maintenance Mode.
2. Skip this step if the SMD controller is already at address **C1000200**. Edit the ramdisk file `/etc/system` if you want to place the SMD controller at an address other than **C1000200**. Remove the `loadvs32` field. (On the controller board itself, address **C1000200** corresponds to base address switch 1 off and all others on.)
3. Execute the `ldeeprom(1M)` command to load the electrically erasable programmable read-only memory on the VME interface card. This command might take up to 10 minutes to complete.
4. Execute the hardware inventory command

hinv -p

to verify that controller number 1 exists and has your drive on it.

5. Installation of 6.2 CTIX requires a valid Volume Home Block in slice 0 of the disk drive. Execute

iv -t /dev/rdisk/c1d0s0

to determine the validity of the Volume Home Block on SMD drive 0. If the drive does not contain a valid Volume Home Block, you will get an error message about not being able to access the drive. You must format the drive with a Diagnostics tape or contact your Convergent Technologies Field Service Representative.

6. Having verified that a valid Volume Home Block exists on the drive, reboot the tape and continue with the Raw Installation steps outlined below.

7. Optional: if your machine has an existing CTIX on an ST506 system disk and you are switching to an SMD system disk, you may want to save any important files from the file systems on the ST-506 disk then re-initialize it.

5.1.3 Raw Installation Process

To raw install 6.2 CTIX perform the following steps.

WARNING: THIS PROCEDURE DESTROYS ALL EXISTING DATA ON THE DISK.

(If you do not have an SMD controller, ignore any messages about its absence.)

1. The raw installation requires an RS-232 terminal at 9600 baud on channel 0 *or* an RS-422 terminal on cluster line 0 (with no RS-232 terminal connected to channel 0).
2. Insert the distribution Quarter-Inch Tape into the drive and close the latch.
3. If the system is not powered on then turn power on, otherwise open the access door on the S/Series enclosure and press the reset button.
4. **WARNING:** The sizes of disk partitions must be even multiples of **32**, in order to create partition starting block numbers that are on disk track boundaries. This is the simplest means of satisfying all disk type (ST506, SMD, and SCSI) requirements.
5. The following output appears on the screen. User supplied responses are shown in **boldface** type. A carriage return is implied after every user input. The example presented here is taken from an SMD disk installation on a system having both SMD and ST506 disks. The corresponding ST506 or SCSI messages appears for their respective installations.

40 MB Drives		
Partition 1	12000	/root
2	6000	swap
3	16000	/usr.
Drives above 40 MB		
Partition 1	16000	/root
2	8000	swap
3	20000	/usr.

CTIX is growing larger than 40MB disks can comfortably support. It is still possible to operate with such a disk, but you must choose carefully which groups to install. If you are unsure of what to choose for a 40MB disk install, consult with S/Series Technical Support.

If you answer 'no' to this question skip to the next section.

If you answer 'yes' the installation continues:

You have N megabytes of dump area
 Do you wish to change this number?
 Type 'yes' to confirm/continue: **yes**

Enter size in megabytes: **0**

NOTE: The dump area on the disk is used to contain what was in memory at the time of a system crash. In most cases you do not need this area. However if you are writing device drivers or other operating system functions, you might want one. If you do, you need a megabyte of dump area for every megabyte of physical memory that is in your system. This also may be set to '0' for no dump area.

NOTE: Slice 2 will be used as SWAP device!
 Default for slice one will be 2000 less than total available blocks to allow for minimum swap slice!

Enter new slice sizes

slice 1: (default=NNNNNN) NNNNNN

NOTE: The installation does not allow the first slice (slice 1) of the root disk to be under a size that can contain the CORE Group of the OS release. It also does not allow you to have a swap slice of less than 2000 blocks. If you want a separate /usr file system you need to configure for at least 3 slices with the third slice being used for the /usr file system. (You can configure for three or more slices and still choose not to have a separate /usr file system.) Just pressing <RETURN> ends the input and continues to the next question. Example:

slice 1: (default=90000) **18000**
slice 2: (default=74000) **8000**
slice 3: (default=66000) **24000**
slice 4: (default=42000) **<RETURN>**

Are these partition sizes satisfactory?
Type 'yes' to confirm/continue: **yes**

NOTE: You see this next question only if you configured for 3 or more slices!

Do you want /usr to be mountable on Partition 3?
Type 'yes' to confirm/continue: **yes**

6. The installation continues:

The Bad Block Table contains NN entries
SMD disk drive #0 is setup
>>>Adding a swap area

path	dev	swaplo	blocks	free
/dev/dsk/c1d0s2	1,2	0	NNNNN	NNNNN

>>>Root file system made."
>>>Usr file system made."
mount: warning: <> mounted as </mnt>
>>>Root file system mounted."

>>>Usr file system mounted."

7. The installation proceeds by loading the programs essential to the installation.

Loading in minimum installation file system
NNNNN blocks

Loading in installation tools
NNNNN blocks

8. The *ctinstall(1)* command is now executed from the raw install script:

ctinstall.sh 6.32

Positioning the Tape for Product Installation.

9. The *ctinstall* script asks for user-input at this point. The usual response in a raw install is **'all'**. (Refer to the section on "Some Notes about Cross Development" for information on the **CROSS** Group).

CORE ADMAN LP UUCP RIOP LAN SYSM ACCT BCE ACE DOC
KFIG GAMES TINFO HELP ASSIST CROSS MISC SMAIL NCC

Please enter your group choices for CTIX

If you'd like all of the groups, type **'all'**;
If you'd like none of the groups, type **'none'**
If you'd like to omit some groups, type **'omit'**
followed by groups to omit: **all**

10. Since you are installing CORE the installation continues with:

Starting to Install Group(s) CORE ADMAN LP UUCP RIOP
LAN SYSM ACCT BCE ACE DOC KFIG GAMES TINFO HELP
ASSIST CROSS MISC SMAIL NCC into /.
Installing Group CORE.

11. The space required for the new files is determined and the changes reported.

Calculating size required for group CORE.

NNN 512 byte blocks will be freed on /
NN additional inodes will be used on /
NNN 512 byte blocks will be used on /usr
NN additional inodes will be used on /usr.

Reading Customizable CORE files.
Installing Customizable CORE files.
NNNNN blocks

CORE Customizable files installed.

Installing required CORE files.
NNNNN blocks

Creating 0 length files.
Checking permissions and modes on new CORE commands.
Completed Installation of Group CORE.
CORE installation cleanup ...

[Other groups selected will now be installed.]

Rewinding tape.

Installation Complete.

12. When the installation has completed, the installation script prompts:
Group ACCT has been installed. Do you want Accounting enabled?
Type 'yes' to confirm: **answer**

Group LP has been installed. Do you want the LP spooler enabled?
Type 'yes' to confirm: **answer**

What do you want the NODE name of the system to be?
answer

Group SAR has been installed. Do you want system activity reporting enabled?

Type 'yes' to confirm: **answer**

Remove the tape and press return (the system will reboot):

13. After pressing RETURN, the system reboots. Installation is complete.

5.2 Update Installation

To update an existing CTIX system to 6.2 CTIX perform the following steps:

1. Log in as root. Provide a password if necessary. The system will respond with a # prompt.
2. **cd /** RETURN
3. Bring the system into single-user mode by entering:

/etc/shutdown RETURN

Before installing the product wait for the system to prompt with the message:

Ok To Stop Or Reset Processor

4. Insert the distribution Quarter-Inch Tape into the drive and close the latch.
5. Reboot the system with the following commands:

sync ; reboot RETURN

6. The following output appears on the screen. User supplied responses are shown in **boldface** type; a carriage return is implied after every user input. The example presented here is taken from an SMD disk installation on a system having both SMD and ST506 disks.

path	dev	swaplo	blocks	free
/dev/dsk/c1d0s2	1,2	0	16000	16000

running fsck on root file system

File System: Volume:

** Phase 1 - Check Blocks and Sizes
** Phase 2 - Check Pathnames
** Phase 3 - Check Connectivity
** Phase 4 - Check Reference Counts
** Phase 5 - Check Free List and Bitmap
FS STATE SET TO OK
NNNN files NNNNN blocks NNNN free

Mounting root file system

mount: warning: <> mounted as </mnt>

8. If **/usr** is a mounted file system then fsck is run and **/usr** mounted.

running fsck on /usr file system

File System:usr Volume:usr

** Phase 1 - Check Blocks and Sizes
** Phase 2 - Check Pathnames
** Phase 3 - Check Connectivity
** Phase 4 - Check Reference Counts
** Phase 5 - Check Free List and Bitmap
NNNN files NNNNN blocks NNNN free

Mounting /usr file system

[It is important to fix any problems before continuing.]

9. The installation proceeds by loading the programs essential to the installation.

Loading in minimum installation file system
NNNNNN blocks

10. Loading in installation tools NNNNNN blocks

11. The `ctinstall(1)` command is now executed from the raw install script:

```
ctinstall.sh 6.32
```

Positioning the Tape for Product Installation.

12. The `ctinstall` script asks for user-input at this point. The usual response in an update installation is 'all'; however, the only *required* choice is 'CORE'. If you have already installed CORE and want to update other groups, refer to section 5.4 of this release notice. Refer to Section 10.10 of this release notice for information on the **CROSS** Group.

```
CORE ADMAN LP UUCP RIOP LAN SYSM ACCT BCE ACE DOC
KFIG GAMES TINFO HELP ASSIST CROSS MISC SMAIL NCC
Please enter your group choices for CTIX
```

If you'd like all of the groups, type 'all';

If you'd like none of the groups, type 'none'

If you'd like to omit some groups, type 'omit'

followed by groups to omit: **CORE**

13. Since you are installing CORE the installation continues with:

Files in the following directory are being removed:

```
/usr/sys - Install the KFIG group for the new kernel
libraries and build tools
```

(Several files in `/usr/sys` are mentioned as being removed. If you are updating from 5.XX or 6.10 CTIX, ignore a message about "instclobber".)

Removing files that no longer exist or are being moved

`/etc/bcheckrc` is no longer customizable,

see `/etc/rcopts/README`

`/etc/bcheckrc` moved to `/etc/OLD/bcheckrc`

`/etc/profile` is no longer customizable,

see `/etc/rcopts/README`

/etc/profile moved to /etc/OLDprofile
/etc/cprofile is no longer customizable,
see /etc/rcopts/README
/etc/cprofile moved to /etc/OLDcprofile

14. The installation continues with:

This procedure will update the / file system with
6.2 CTIX group(s)

XXX

Type 'yes' to confirm/continue: **yes**

Starting to Install Group(s) XXX into /.
Installing Group XXX.

15. The space required for the new files is determined and the
changes are reported.

Calculating size required for group XXX.

NNN 512 byte blocks will be freed on /
NN additional inodes will be used on /
NNN 512 byte blocks will be used on /usr
NN additional inodes will be used on /usr.

Reading Customizable XXX files.

Installing Customizable XXX files.

16. You are then queried, for each customizable file, as to
whether you would like to:

[r] replace the old one with the new

[s] save the old one in *OLDfile* and replace with the new

[i] ignore the new one

[d] perform an sdiff(1) between the old one and the new

[rd] replace the old one with the previous diff.

If you do not want to destroy an old version, save the old
one and replace with the new. Then, after the installation

is complete, compare the two and update the new one with your changes.

17. The installation continues:

XXX Customizable files installed.

Installing required XXX files.

NNNNNN blocks

18. Creating 0 length files.
Checking permissions and modes on new XXX commands.
Completed Installation of Group XXX.
XXX installation cleanup

[Other groups selected are installed also.]

Rewinding tape.

Installation Complete.

19. When the installation has completed, the installation script prompts:

Remove the tape and press return (the system will reboot):

20. After pressing RETURN, the system reboots. Installation is complete.
21. If **CORE** is installed, the "CORE installation cleanup" runs *iv(IM)* to install a new loader and download files.
22. If **CROSS** is one of the groups chosen to be updated, decide which cross development libraries, if any, are to be installed. If your machine has floating point hardware, decide whether to install **2fp** in the root file system replacing the default **2sw** libraries. If **2fp** is installed in the root file system, it is removed later from **/cross**. The following messages appears on the screen:

Querying Optional CROSS files.

The cross development libraries available for installation are:

68020 with software floating point (2sw)
68020 with 68881 hardware floating point(2fp)

23. If the machine has software floating point, the installation script echos:

This 68020 machine has software floating point.

24. If the machine has hardware floating point, the script echos:

This 68020 machine has a 68881 chip (hardware floating point).

Type 'y' if you'd like to install 2fp into the ROOT file system: **y**

Changing CENVIRON in /etc/profile to
CENVIRON="CPU=68020,FPU=68881"

25. The installation continues:

Please enter your choices separated by blanks.
Possible choices are:

all none 2sw 2fp

Please enter 'all' for both; 'none' for none; or,
any combination of '2sw 2fp' separated by blanks: **all**

26. If **CORE** was one of the updated groups, upon completion the installation script prompts:

Remove the tape and press return (the system will reboot):

27. Installation is complete.

5.3 Update Silent Mode Installation

To non-interactively update your system, follow these instructions.

1. Perform steps 1 through 5 listed in Section 5.2, "Update Installation", above.

2. Step 6 listed in Section 5.2, "Update Installation", is essentially the same except that when selecting the installation option, specify silent instead of update.

Do you wish to install, update, or do maintenance?
Choices are ('install', 'update' or 'maintenance')? **update**

Do you want this to be a silent update?
Silent update uses fewer prompts and saves customizable files as 'OLDfilename'
Type 'yes' to confirm or 'no' to deny: **yes**

3. The installation proceeds very much as in Section 5.2, "Update Installation", except that no questions are asked by the installation script, with the exception of:
 1. Which groups do you want to install?
 2. Then at the end some questions regarding options.

Each old customizable file is saved in *OLDfile* and replaced with the new file. After the installation is complete you may wish to compare your old customizable files with the newly installed ones and update the new ones with your customizations.

5.4 Non-CORE Update Installation

The procedures outlined above for RAW installation, update, and update silent mode installation are those required to install the CORE Group, normally the additional groups that are required would be installed as part of that process. However if you want to install some groups later, after installing CORE (it is **essential** that CORE is the first group installed from a set of release tapes), the following procedure can be used:

1. Log in as root. Provide a password if necessary. The system responds with a # prompt.
2. `cd /` RETURN
3. Bring the system into single-user mode by entering:
`/etc/shutdown` RETURN

Before installing the product wait for the system to prompt with the message:

Ok To Stop Or Reset Processor

4. Re-mount /usr if it is a mounted file system:

`/etc/mount /usr` for 6.XX CTIX

5. Insert the distribution Quarter-Inch Tape into the drive and close the latch.
6. Rewind the tape with the following command:

`tsioctl -c rewind /dev/rmt0` RETURN

7. Execute the installation script:

`ctinstall` RETURN

8. The following output appears on the screen. User supplied responses are shown in **boldface** type; a carriage return is implied after every user input.

`ctinstall.sh 6.32`

Positioning the tape for Product Installation

9. The `ctinstall` script asks for user-input at this point. You should not select CORE, because `ctinstall` will not allow a CORE install without booting from the tape. (Refer to Section 10.10 of this release notice for information on the **CROSS** Group).

Please enter your group choices for CTIX separated by blanks.
Your choices are:

CORE ADMAN LP UUCP RIOP LAN SYSM ACCT BCE ACE DOC
KFIG GAMES TINFO HELP ASSIST CROSS MISC SMAIL NCC

If you'd like all of the groups, type 'all';
if you'd like none of the groups, type 'none': < **X** >

10. Further output is similar to that described in section 5.2, steps 14 through 25.

6. Using Maintenance Mode

The 6.2 S/Series CTIX Raw/Update Installation tape contains the maintenance mode capability of previous releases. The maintenance mode allows you to:

1. boot CTIX
2. use provided tools to determine the cause of the problem
3. fix the problem.

Below are some examples of problems:

1. When you are locked out of the system (for example, because you have forgotten **root's** password and you cannot modify **/etc/passwd** without being **root**),
2. Loading the VME Interface Board EEPROM.

A partial list of the files available in maintenance mode follows.

/bin/cat	/bin/chmod	/bin/cpio
/bin/cp	/bin/ed	/bin/ln
/bin/ls	/bin/mv	/bin/mkdir
/bin/rm	/bin/sh	/bin/stty
/bin/sync	/dev/console	/dev/vme/a16
/dev/vme/a24	/dev/vme/a32l	/dev/vme/a32h
/dev/vme/eprom	/dev/dsk/c0d0s[0-f]	/dev/dsk/c0d1s[0-f]
/dev/dsk/c0d2s[1-f]	/dev/dsk/c1d0s[1-f]	/dev/dsk/c1d1s[1-f]
/dev/dsk/c1d2s[1-f]	/dev/dsk/c2d0s[1-f]	/dev/dsk/c2d1s[1-f]
/dev/dsk/c2d2s[1-f]	/dev/rdisk/c0d0s[0-f]	/dev/rdisk/c0d1s[0-f]
/dev/rdisk/c0d2s[0-f]	/dev/rdisk/c1d0s[0-f]	/dev/rdisk/c1d1s[0-f]
/dev/rdisk/c1d2s[0-f]	/dev/rdisk/c2d0s[0-f]	/dev/rdisk/c2d1s[0-f]
/dev/rdisk/c2d2s[0-f]	/dev/kmem	/dev/mem
/dev/null	/dev/prf	/dev/rmt0
/dev/rmt4	/dev/tty000	/dev/tty256
/dev/tty	/etc/fsck	/etc/fsdb
/etc/icode	/etc/icode1	/etc/iv
/etc/mkfs	/etc/mnntab	/etc/mount
/etc/passwd	/etc/reboot	/etc/splocate
/etc/swap	/etc/umount	/etc/system
/etc/ldeprom	/etc/hinv	/etc/iv.desc.cpio
/shlib/libc2sw_s		

Note: Access to the maintenance mode should be restricted, since it poses an obvious security risk.

The number of files that can be included in the ramdisk file system is limited by the space needed for the kernel, a user area in which to run processes, and the files in the maintenance file system. The maximum space is limited by the minimum physical memory available on any one of the S/Series machines: 2MB. For this reason, many files, such as the loader `/usr/lib/iv/loader`, are not included.

If you need the loader, in order to initialize the system disk, you must use one of several means:

1. Remove some files from the ramdisk files system and use `cpio(1)` to bring in the loader from another 6.2 CTIX system to `/usr/lib/iv`.
2. Mount the system disk `/usr` file system (if possible) on `/mnt` and copy `/mnt/usr/lib/iv/loader` to `/usr/lib/iv`. Unmount the `/usr` file system before doing the `iv(1)`.
3. Perform an **install** of the CORE Group.

The file `/etc/iv.desc.cpio` is a `cpio(1)` archive containing minimally, the following files:

```
/usr/lib/iv/desc.40PS  /usr/lib/iv/desc.80PS
/usr/lib/iv/desc.85I  /usr/lib/iv/desc.85M
/usr/lib/iv/desc.85S  /usr/lib/iv/desc.80QS
/usr/lib/iv/desc.140M /usr/lib/iv/desc.140MS
/usr/lib/iv/desc.190M /usr/lib/iv/desc.300N
/usr/lib/iv/desc.350MS /usr/lib/iv/desc.650MS
/usr/lib/iv/desc.750N
/usr/lib/iv/desc.tape  /usr/lib/iv/desc.tdump
```

If you need to use one of these files, it must first be extracted from the archive. Note that the desc files do not have entries for the dump area and the download images. You must edit the desc files if you want to add them.

```
cpio -iv desc.XXXXXX < /etc/iv.desc.cpio
```


In maintenance mode the following table describes controller to controller-type configuration (again note that **S/80** means S/80 and S/280, **S/320** means S/120, S/22X, and S/320, and **S/640** means S/480 and S/640):

controller	S/80	S/320	S/640
c0	none	ST506	ST506
c1	none	SMD	SMD
c2	SCSI	none	SCSI

Please refer to "Using the Maintenance Tape" in Chapter 13 of the *S/Series CTIX Administrator's Guide* for complete instructions.

Warnings:

1. When accessing the Quarter-Inch Tape drive, use the device nodes `/dev/rmt0` and `/dev/rmt4`. These are linked at tape bootup time to the device nodes in `/dev/rmt` that correspond to your machine type. The device nodes `/dev/rmt/c0d0` and `/dev/rmt/c0d0n` are used for the QIC-2 drive of the S/320 systems and have major device numbers 18. The other device nodes in the directory `/dev/rmt` are used for the SCSI Quarter-Inch Tape drive of onboard SCSI-based systems and have major device numbers 65.
2. If you enter a `<CTRL>D` to the maintenance mode shell the stand-alone kernel will panic. The maintenance mode `init exec's` the shell. `<CTRL>D` kills the shell, and then there are no processes, a fatal condition.

7. Deleting a Group of Files from the System

WARNING: DO NOT REMOVE THE CORE GROUP.

Do not remove the BCE and KFIG groups if you need to reconfigure your kernel.

To delete a group of CTIX release files from your system, follow these directions. Groups that can be deleted are:

ADMAN LP UUCP RIOP LAN SYSM ACCT BCE ACE DOC KFIG
GAMES TINFO HELP ASSIST CROSS MISC SMAIL NCC

1. Insert the distribution Quarter-Inch Tape into the drive and close the latch.

2. Make sure the tape is rewound by typing

```
tsioctl -c rewind /dev/rmt0          RETURN
```

3. Skip to the installation control files for the group you want to delete. To do this, type

```
tsioctl -c skip /dev/rmt4 [arg]      RETURN
```

where **arg** is the *File #* listed in Section 3.1 of this release notice.

4. **cd /** RETURN

5. Extract the installation control files for the group.

```
cpio -ivcQdum < /dev/rmt0          RETURN
```

6. Do an *ls(1)* to see which file lists exist for this group.

```
ls GROUP*                            RETURN
```

7. Remove the files in this group.

```
cat [GROUP_file_lists] | xargs rm -f RETURN
```

where **GROUP_file_lists** can be one of the files:

GROUP, GROUP.cust, GROUP.noup, GROUP.noqu
GROUP.fopt

8. Remove the file-list files.

```
rm GROUP*
```

RETURN

For example, to delete the **GAMES** Group from the system, execute:

```
tsioctl -c rewind /dev/rmt0  
tsioctl -c skip /dev/rmt4 34  
cd /  
cpio -ivcQdum < /dev/rmt0  
cat GAMES | xargs rm -f  
rm GAMES*
```

RETURN
RETURN
RETURN
RETURN
RETURN
RETURN

Note that if you are removing the **SMAIL** Group, you must also perform the following:

```
mv /bin/OLDRmail /bin/rmail
```

8. Reconfiguring the System for More or Fewer Users

As distributed, 6.2 S/Series CTIX is a 16 user system. The size of the **proc table** and the in-core **inode table** have been fixed to support 16 users. You can reconfigure your system for a number of users different from the 6.2 CTIX release default. **If you intend to increase the number of users to the system limit, consult with S/Series Technical Support for the correct number of maximum users for your system configuration. Pushing your system above the limits might produce a configuration unsupportable by CTIX.** If you never have more than 8 users on your system at a time, you might reconfigure for 8 users. This would maximize available user memory. (Refer to Appendices B and C in the *S/Series CTIX Administrator's Guide* for information about which configuration parameters are affected by the number of users specified).

1. Login as **root**.
2. Change the entries in **/etc/inittab** for gettys required. Refer to Chapter 5 and Appendix B in the *S/Series CTIX Administrator's Guide* for instructions on how to configure **/etc/inittab**.
3. Change the kernel parameter **v_proc**, using *uconf(1M)*, or make the change in your version of **dfile** and rebuild the kernel.
4. Reboot the system.

The nominal upper bounds on the systems in the S/Series are:

<i>System</i>	<i>Upper bound on maximum users</i>
S/80	16
S/280	22
S/221	32
S/222	32
S/320	32
S/480	48
S/640	64

NOTES:

1. The utility *muser(1M)* has been discontinued.
2. The output from **hinv -u** or **hinv -p** will now show a value of 128 maxusers for all S/Series machines. **This value is meaningless to the user and should no longer be used in scripts that attempt to calculate the maximum users set for the system.** The -u option of *hinv(1M)* will be discontinued in future releases.

9. The SCSI Subsystem

9.1 S/Series SCSI Hardware

The S/Series machines provide the capability of attaching SCSI devices to the system:

1. Through a SCSI bus on the SCSI/RS-232 Expansion board for the S/320 systems.
2. Through either an onboard SCSI bus, or a SCSI bus on the SCSI/RS-232 Expansion board for the S/640 systems.
3. Through the onboard SCSI bus for the S/80 systems.

The Western Digital chip WD3393 is the SCSI bus interface controller. Together with the CPU, this handles all necessary bus interface and SCSI protocol communication with other SCSI devices. SCSI devices can be connected to the internal SCSI bus, or connected to the SCSI bus through an external SCSI adaptor. For the S/80 and S/280 this can also be done using the connector on the SCSI/Ethernet board.

Refer to your SCSI Hardware Reference manuals for controller and peripheral configuration details.

9.2 S/Series SCSI Software

There are three drivers in the system to provide the necessary software to interface to the SCSI hardware. When the kernel is configured, the corresponding entries have to be present in the system configuration file, `/usr/sys/cf/dfile`.

The low level driver, `scsi`, which uses the device node `/dev/scsi` provides the interface to the system hardware and handles the SCSI bus interface and protocol. All the high level drivers interface through this device to the low level SCSI drivers. `scsi` must be an entry in the `dfile`, after the disk driver entries.

The disk driver, `scsidisk`, provides the interface to SCSI disks. The current implementation only allows one copy of `scsidisk` in the system, supporting a maximum of 16 disks. Since there is only one copy of `scsidisk` in the system, only one line of `scsidisk` is needed in `/usr/sys/cf/dfile`, even if you have

multiple SCSI buses in the system.

The tape driver, **stape**, provides interface to SCSI tape drives, both SCSI Quarter-Inch Tape drives and SCSI Half-Inch Tape drives. The current implementation supports a maximum of 8 tape drives. **stape** may be an entry in the tape section of dfile, or the SCSI tape driver, `/etc/lldrv/stape.o`, may be dynamically loaded.

Refer to the *S/Series CTIX Administrator's Guide* for details on software configuration of the SCSI subsystem.

10. Notes on System Operation

10.1 A Note on System Tuning

It is extremely important to tune your system for the Ethernet driver. If not configured properly, the number of stream buffers can become a bottle neck for Ethernet performance. It is recommended that the system should have at least 32 of 2k buffers. To reconfigure your system, add the following line to `/etc/system`, in the `!TUNABLES` section, under the `'dsabldbg=1'` line:

```
v_nblk2048=32
```

Note that the symbol `'*'` in column 1 is a comment delimiter.

Then run `uconf` to update the changes:

```
uconf -w
```

The changes are effective next time CTIX is rebooted. **This note applies to both RFS and NFS network operation.** See the `uconf(1M)` manual page for details.

The system administrator should use `netstat -m` periodically to monitor stream buffer usage to tune the system. For example, if the failure rate in 1k stream buffer blocks is high, the number of 1K buffers should be increased. Refer to the *S/Series Network Administrator's Guide* for more information.

10.2 A Note on SCSI Half-Inch Tape Drives

The Cipher Half-Inch Tape drives come up in fixed block mode as the default, so the mode must be changed by using the `tapeset(1)` command. The command does a MODE SENSE to the tape drive, removes the fixed block length bit, and then does a MODE SELECT to the drive. This sequence turns the tape drive into variable-length block mode. The drives do not support a density command.

Some of the SCSI Half-Inch tape drives do not accept commands when off-line. Be sure your drive is on-line at boot time and at any time that a `tapeset` command is issued to the drive.

In the file `/etc/tapedrives`, the following lines have been added:

```
F880S s 131072 12 00000000000000000000000000000000 00007fffff00000000  
M990S s 131072 12 00000000000000000000000000000000 00007fffff00000000
```

Create the file `/etc/rcopts/KSCSI`, if necessary, and add the following lines:

```
tapeset -t F880S /dev/rmt/c0d1c  
tapeset -t M990S /dev/rmt/c0d1c
```

Note that jumper W7 on the SCSI adapter for the Cipher F880S and M990S should be removed, otherwise the drive will reset the SCSI bus every time power is cycled.

10.3 Clist Usage

The RS232 ports are now controlled by a STREAMS-based driver that does not make use of *clists*. The remaining serial devices that use *clists* are RS422 ports and the CT Window Manager driver, *wxt*. Check `/etc/inittab` to obtain the number of active RS422 ports with *getty* spawned. Obtain the estimated maximum number of RS422 ports that will be using the CT Window Manager. The number of active **clist** users is the number of active RS422 ports plus two times the estimated number of CT Window Manager uses. The number of *clists* should be changed to the following number:

$$((\text{number of active clist ports}) * 6) + 40$$

10.4 Tty Device Nodes

Note that cluster device nodes in `/dev` are not provided for `/dev/tty288` and above, because they are needed only in case of large configurations. If you need the nodes, you must make them. Refer to the new man page *createdev(1M)* for information on creation of nodes.

10.5 Process Sizes and Swap Space

The **sar -r** report can be used to track the available swap space; if necessary, additional space can be allocated via the *swap(1M)* command.

It is recommended that the swap area be at least one-half again

as much space as there is physical memory on the system. To obtain more swap space, use the *swap(1M)* command with the *-a* option. There is no need to reconfigure the kernel or to reconfigure the disk to enlarge the original swap slice (*/dev/rdisk/c0d0s2*) when this is done. Please refer to the *swap(1)* entry in the *CTIX Operating System Manual, Version C, Second Edition, Volume 1* and to *Adding Swap Space with the swap Command* in *Chapter 8* of the *S/Series CTIX Administrator's Guide*.

As an additional safety measure, CTIX restricts the maximum size that a process can be to 1/4 of the total available swap space plus 1MB. This limit dynamically changes when a **swap -a** or **swap -d** is issued. This default 'floating' max process size can be overridden by changing the **maxumem** variable in **dfile** and building a new kernel. A maximum value of 6144 will allow a 24 MB process to be created.

Note that the amount of virtual memory and physical memory is different for various members of the S/Series family. For the S/320 systems and the S/80, the maximum virtual memory is 32 MB and the maximum physical memory is 16 MB. For the S/640 systems, the maximum virtual memory is the full address space of the MC68020 (4 GB) and the maximum physical memory is 64 MB.

10.6 Tuning the System for STREAMS Buffers

The STREAMS buffer resources should be configured as shown below:

Name	Number per User			
	Buffer Size	1 User	16 User	32 User
nblk2048	2048	0	16	32
nblk1024	1024	1	48	96
nblk512	512	5	64	128
nblk256	256	5	256	512
nblk128	128	10	320	640
nblk64	64	20	640	1280
nblk16	16	20	512	1024
nblk4	4	10	512	1024

10.7 Configuring the System for the 4K Block File System

If you want to create 4K file systems, you must configure your system to optimize performance.

1. The system should have at least 3MB of physical memory.
2. The 4K file systems should be created on disks of size 85MB or more.
3. The number of 4K buffers should be about one-fourth the number of 1K buffers. The number of 4K buffers in the kernel should be in the range of 16 to 32. The number of 1K buffers is auto-configured at bootup. Check the number of 1K buffers, using the `-v` option of `uconf(1M)`, and configure the number of 4K buffers accordingly, using the `-w` option of `uconf(1M)`.

10.8 Accessing Additional Disk Space on Larger Drives

S/Series systems shipped with disk drives of size greater than 40MB have many additional megabytes of space in a fourth partition which was formatted at installation but was not initialized with a valid file system structure. Refer to Chapter 8, "Disks and File Systems", in the *S/Series CTIX Administrator's Guide* for instructions on how to create and mount file systems.

10.9 Re-Partitioning Additional Disks

S/Series systems shipped with more than one disk drive have only the first (system) disk initialized with more than one file system. The remaining disks are initialized with a Volume Home Block in partition 0 and a large file system in partition 1.

For instructions on how to re-partition these disks into multiple file systems, refer to Chapter 8, "Disks and File Systems", in the *S/Series CTIX Administrator's Guide*.

10.10 Some Notes About Cross Development

Installation of the **CROSS** Group provides the capability to do software cross development for the following processors:

- 68020 with software floating point (2sw)
- 68020 with 68881 hardware floating point (2fp)

The default libraries installed into your root file system under

`/lib` and under `/usr/lib` contain 68020 and software floating point objects. If your system has 68881 hardware floating point, and you plan to develop software *only* for the 68020 with 68881 hardware floating point, you might want to install the 2fp libraries as the default libraries. (This eliminates the need for setting the LIBROOT environment variable to `'/cross/2fp'` upon compiling and linking.)

- If the CROSS Group is not already installed, perform an update installation of the Group. (Refer to Section 5.2 of this release notice.)
- If the entire CROSS Group is already installed, execute the `/cross/crossins` script to install one set of cross development libraries into the root file system. This script also changes the definition of CENVIRON in `/etc/profile`, and allows the removal of extraneous cross development libraries from `/cross`.

Please refer to the `cc(1)` and `ld(1)` pages in the *CTIX Operating System Manual, Version C, Second Edition, Volume 1* and to the *Programmer's Notes for C Cross Compiler and Flexnames*

11. Known Errors, Restrictions, and Warnings

11.1 Known Errors

Known errors at the time of release of 6.2 CTIX are listed below.

1. When NFS runs out of inodes, you get the warning:

nfs_iget: rnode table overflow

Shortly after that the system panics with a

Page 0 access in kernel

To avoid this problem, increase the default value of **50** for the **nfs_rnum** entry in **/etc/system**.

2. 1.1 RFS will only allow you to execute 64 pairs of the commands *rfstart(1M)* and *rfstop(1M)*. Thereafter, you must reboot the system, in order to bring up RFS.
3. *crash(1M)* has not been completely converted to handle both STREAMS-based and clist-based serial ports. The following *crash(1M)* commands currently work correctly: **ser**, **vt**. The following commands do not work: **pt**, **gt**, **wxt**, **clist**, **cbk**. The command **tty** works in so far as it reports STREAMS-based ports. The command **ser -f** works correctly for non-IOP controlled ports but will report bogus values for the **device** structure (see **/usr/include/sys/serial.h**).
4. If you add SCSI disk drives to the onboard SCSI bus **0**, you must add them in decending sequence, in order to satisfy the built-in SCSI map. Use the example below as a guide.

```
disk-c0d0 bus=0 target=6 lun=0 parity reselect
disk-c0d1 bus=0 target=5 lun=0 parity reselect
disk-c0d2 bus=0 target=4 lun=0 parity reselect
tape-d0 bus=0 target=1 lun=0 parity reselect
```
5. *crash(1M)* might successfully read an S/280 disk dump area.
6. If *rcmd(1)* calls a remote command which calls another command in the background, the original *rcmd(1)* will not terminate until the background job terminates.

7. If you try to create a disk crash dump area of size 64 MB, *w(1M)* writes a 0 size dump area, but the installation script *RawInstall* reserves 64 MB of disk space. The work around is to create a dump area of less than 64 MB. This problem is of minor severity, since it is unlikely that you would want to give that much disk space to crash dumps.
8. **iv -iv** cannot format SCSI disk drives if there are trailing spaces on any line of the description file.
9. Systems with a SCSI bus might report accumulating SCSI data I/O errors, even though no data corruption or loss occurs. The errors cause retries that are eventually successful, but the errors are still being logged.
10. The new utility *serstat(1M)* might dump core when the mode (scan, auto-scan, or continuous) is changed.
11. You will get an inappropriate message when executing a command with *dd* piped to *tio* with different block sizes specified for writing and reading the tape:

reach end of medium

instead of

change tape ...

12. On the S/320 systems and the S/80, when you execute **tsioctl -c erase** with a write-protected tape, the command returns immediately with a status value of 0.
13. With the VME Half-Inch Tape driver, **tsioctl -c skip** does not identify two end-of-file (EOF) marks as End-of-tape (EOT). This can cause the unit to wind the tape off the supply reel.
14. Setting **dbconsole** in */etc/rcopts/LOCDRVLD* causes the system to never get to *login* state. Enter the command after the system has come to multi-user run level.
15. If the kernel debugger and symbols drivers are loaded but **dsabldbg** is not set to 0, then if the system panics symbols are not displayed.
16. When in the kernel debugger, you cannot select output to go to printer or to log file.

17. If the superuser executes

```
cc -o dirname somefile.c
```

the loader, *ld(1)*, will use the system call *unlink(2)* to remove the specified output file, **even though it is a directory!** A regular user will get the message

```
ld fatal: cannot create output file .
```

18. The C compiler link-loader, *ld(1)*, will hang forever if */usr/tmp* runs out of space.
19. If the user tries to set the Half-Inch Tape drive to low density by executing:

```
echo < /dev/rmt/c? d? l
```

the tape drive unloads the tape.

20. *usage(1)* in *print command list* mode causes the top line on the screen to be offset. Also when a command with a large usage description is printed, the first screenful is displayed with a *more* at the bottom, but nothing further can be displayed.
21. If you use **getline** on the last line of a file that does not have a final newline, *awk(1)* will coredump.
22. **SPR 12131:** The compiler does not correctly pad structures which contain fields of type **enum**. As an example, consider the following structure definition:

```
enum a { a,b,c,d,e };
struct str {
    enum a a1, a2, a3;
    char arr[6];
};
```

The size of this structure on an MC68010 based system is 18 bytes. On an MC68020 system it is 20. The MC68020 compiler pads the structure with two bytes after the 6 byte array. This is incorrect because it is incompatible with an earlier version of the compiler. Unfortunately, fixing this problem would cause significant compatibility problems for existing users of the compiler.

23. **Floating point:** There are a number of known floating point errors which are documented by internal SPR's. These errors are all related to incorrect return values from floating point conversions in overflow situations, such as when a NaN is involved in an expression. See the *IEEE Standard for Binary Floating-Point Arithmetic, P754*, for a detailed discussion of NaN.
24. If your system has a cluster of PT or GT terminals and one of them develops a hardware problem, the system might crash with a panic message:

```
pt_recv: too much input
```

The work-around is to pull the offending terminal out of the cluster chain.

25. If the CT Window Manager, *ctwm*, is used, the *passwd(1)* program cannot determine who you are. Use **passwd user_name**.
26. When illegal options are given to *uconf(1M)*, it might respond with bogus error messages. Also it does not check for out-of-range inputs, nor suggest proper values.
27. There are many minor bugs and omissions, too numerous to mention, in AT&T's ASSIST package. The package is provided but will not be upgraded to fix these deficiencies.
28. If too many printers are active in the spooler, error messages about "too many processes" are generated, because *lp(1)* as a user process, is limited to 25 processes. There are several solutions:
1. Change **/usr/sys/cf/dfile** to increase the parameter **maxproc** to 50 and then rebuild the kernel, link it to **/unix**, and reboot.
 2. Change the kernel parameter **v_maxup** with *uconf(1M)* and reboot the system.
 3. Set up the spooler to run with individual UID's per printer spool. (Keep the group ID the same.)

11.2 Restrictions

Restrictions known to apply at the time of release of 6.2 CTIX are listed below.

1. For ports controlled by an IOP, up to 4 ports are guaranteed to support SLIP lines.
2. The encryption/decryption functionality related to *crypt(1)* is not contained in this release. There is not a man page for *crypt(1)* in the current CTIX manual set. There is a separate *crypt(1)* product.
3. You must have at least 3 MB of memory in order to run TCP/IP, RFS or NFS, and applications.
4. The utility *uugetty(1)* should be used with the **-r** option. Note that the new version (to handle STREAMS-related changes) will take as much as five seconds to return a login prompt.
5. If you execute *shutdown(1M)* and attempt to return to multi-user level (init state 2 or 3) without rebooting the system, some system services will not be restarted and if your **/usr** is mountable, it will not be mounted.
6. *ftp(1)* does not support *ksh* (the new AT&T Korn shell). You will get the message "access denied" if you try to use *ftp(1)* from this shell.
7. 6.2 S/Series CTIX does not support the Convergent TO-300 Terminal PC compatible keyboard.
8. You must add at least 1 MB of Expansion Memory to an S/80 machine in order to run 6.2 CTIX without problems. The 6.2 CTIX kernel is much larger than the 6.10 CTIX version, because all drivers are now included in the kernel. The S/320 systems drivers (the SMD disk driver, QIC-2 tape driver, RS-422 cluster device driver, ST506 disk driver, VME interface driver, and the VME Half-Inch Tape driver) are not needed by the S/80 systems. To remove them, you must change **/usr/sys/cf/dfile**, rebuild and install the kernel, and reboot your system.
9. The Convergent Window Manager *ctwm* does not work in 6.2 CTIX, due to the STREAMS-based implementation of **tty**.

10. The *adman(1)* menu for adding a printer does not allow you to choose your model file. Also *adman(1)* does not provide a menu for deleting a printer model.
11. If *adman* is used to add a dial-in only modem, and then you try to delete the modem later, *adman* will claim that there are no modems configured. Since the modem was defined as dial-in, there is no entry for it in */usr/lib/uucp/Devices*.
12. *PCX* will not support RS-422 cluster devices in 6.2 CTIX.
13. *adman(1)* does not support administration of the TCP/IP daemon *named(1C)*. *adman(1)* should not be used for administration of TCP/IP.
14. If *adman(1)* is used on the system console (usually */dev/tty000*), there may be messages sent to the screen that *adman* cannot trap, especially if the kernel debugger is loaded and enabled. When *adman(1)* is using *hinv(1)*, the output will be sent directly to the console and it will overwrite the *adman* screen.
15. *fstyp(1M)* always reports a file system as **S51K**, whether it is 1K or 4K, because the size block is independent of the file system type.
16. The *lp(1)* spooler does not allow printer names longer than 13 characters.
17. The *-s* option for *tsioctl(1)* is not implemented for SCSI tape drives. You will get an error message if you attempt to use it.

11.3 Compatibility Issues

The following changes have been made:

1. The Motorola MC68010 cross compilation libraries are no longer provided in the CROSS Group. The 68010 compiler has not been removed.
2. The *openi(2)* system call is no longer supported.
3. The *libtermcap.a* and *libtermplib.a* libraries are now provided by linking to *libcurses.a*. This means that they are now based on *terminfo* and not on *termcap*. This

means that the **TERMCAP** environment variable is not recognized by programs linked with this library. The old *termcap* versions of these libraries are provided as **libtermcap.a** and **libtermplib.a**.

4. The capability to set a negative nice value (and thus fixing the priority of a process) has been changed so that it must be enabled in the system before it is operative. This change has been made so that the released system can be **SVID** compatible. The mechanism may be enabled by executing **/etc/rtpenable -e** (as **root**), or a program executing the **syslocal(SYSL_RTNICE,1)** system call (again as **root**). This capability can be automatically enabled at system boot by creating the zero length file **/etc/rcopts/ENABRTNICE**.
5. A change in System V.3 (and hence in 6.2 CTIX) is the new definition of the signal system call from "extern int(*signal())();" to "extern void(*signal())();".
6. The files controlling the boot process are substantially different in 6.2 CTIX from those in 5.XX CTIX. The files in **/etc/rcopts** control options in the initialization scripts. The files in **/etc/rcopts** are not provided on the CTIX Distribution tape; hence, customizations via this mechanism simplify reconfiguration on new installations. A lot of actions have been extracted from **/etc/rc**, which is no longer used and has been replaced by **/etc/rc0**, **/etc/rc2**, and separate scripts in **/etc/init.d**. These scripts are then linked to files in **/etc/rc0.d** and/or **/etc/rc2.d**. These files are not marked as customizable in the 6.2 S/Series CTIX distribution. It is intended that their actions are modified by the presence of files in **/etc/rcopts**. In many cases local customizations may be provided by creating scripts in **/etc/rcopts**. See **/etc/rcopts/README** for a brief description of the options.
7. A number of customizable files (**/etc/inittab**, **/etc/checklist**, **/etc/drvload**), contain changes that are important for the successful operation of the system. As such it is important that these files are not ignored during installation. It is also important that when applying customizations from an earlier release that the

customizations be added to the new versions. An example is `/etc/checklist`, the new version specifies the root file system as the blocked device and all other file systems as the raw disk. If the root file system is specified in `/etc/checklist` as the raw disk the system enters a loop of booting, entering ADMIN mode, you run `fsck(1M)`, reboot the system and enter ADMIN mode, etc. `/etc/drvload` manipulates a number of files that are involved with controlling the new capabilities provided by `/etc/lldrv/lldrv` in this release, hence it is important to use the new `drvload` or drivers will not load. `/etc/inittab` now invokes `/etc/rc2` to enter init state 2. The old `inittab` invoked `/etc/rc` (which is moved to `/etc/OLDrc` as part of the installation), hence the old `inittab` must not be used.

11.4 Warnings

1. The products listed below are known not to function correctly under 6.2 CTIX.

Revision	Product
2.20	S/Series CTAM
2.20	PCX

2. S/Series WGS installation overwrites some of the **ADMAN** Group files, so you must reinstall the **ADMAN** Group.
3. *Shell layers* does not work on an RS-422 line. The shell layers driver and the RS-422 cluster line driver both use multiplex line discipline and cannot be active concurrently.
4. **You must use the new 3.2 TCP/IP release with 6.2 CTIX.**

If you want the system to be a gateway, you must make some modifications. If the system is at multi-user level, execute a shutdown. When the system is at single-user level, do the following:

- To the file `/etc/system` on the gateway system, add the line:
`net_ipforwarding=1`

- To update your kernel, execute:
 uconf -w
 - Reboot your system.
5. **If you are using NFS, you must first install 6.2 CTIX, then 3.2 TCP/IP, and finally reinstall 1.0 NFS.** 6.2 CTIX will overwrite some *NFS* files.
 6. **If you are using RFS, you must update to 1.1 RFS.** You must first install 6.2 CTIX, then 3.2 TCP/IP, and finally 1.1 RFS.
 7. The *semop(2)* system call returns the value of the semaphore in 5.2X CTIX. In 6.00, 6.10, and 6.2 CTIX, it returns 0 if successful and -1 otherwise.
 8. *uconf(1M)* must be used carefully. It does not check validity of some inputs; for example, if **reccs** is increased beyond a reasonable value, the kernel will fail to reboot. (Always keep a copy of the old kernel until certain that a new reconfiguration won't fail.)
 9. Newer versions of the SCSI-based systems use a 150 MB SCSI Quarter-Inch Tape drive. Tapes written on this drive cannot be read by the older 60 MB SCSI Quarter-Inch Tape drives. Tapes written on the 60 MB drive can be read on the 150 MB drive. **Crash dump tapes labeled on the 60 MB drive cannot be used on a 150 MB drive. First erased the tape and then re-label it.**
 10. For the S/80 and S/640, the disk node **c0d0** on bus **0**, is the **rootdev**. Don't change the entry for this node in the **!SCSIMAP** section of the system configuration file, **/etc/system**.
 11. On an S/80, on bus 0, the target ID **0** is reserved for the SCSI/Ethernet controller. Don't use this ID number in the **SCSIMAP** section of the system configuration file, **/etc/system**.
 12. If your system is running NFS and the server runs out of disk resources, you will get messages about "WARNING: NFS getattr failed for server XXXX: TIMED OUT", and eventually your system will panic. The administrator of the server system must monitor working disk space. This

bug is fixed in UNIX System V Release 4.

13. The S/PC CTIX/386 RFS routine *getservaddr* swaps the first two bytes of the expanded internet address. This causes remote mounting failures between S/Series and S/PC systems. The fix is to edit the S/PC file */usr/nserve/rfmaster* to swap the bytes.
14. If you are converting network file system server from RFS to NFS or from NFS to RFS, you must convert remote mount entries in the file */etc/fstab*.
15. If a kernel is built using */usr/sys/cf/dfilerfs* or */usr/sys/cf/dfilefs* without the entry for **socket**, the drivers that require the **socket** driver will fail to load. *drvload(1M)* will not load the socket driver. It will complain about redefines and not being able to find entry points.
16. *Adman* will allow you to add a modem to a port that already has a *getty* for a terminal running on it.
17. If your system is configured for Ethernet, you cannot take the system to single-user run-level and then go from single-user run-level to multi-user run-level. The system almost comes all the way up, but never starts any *gettys*.
18. There is a swap area problem that can cause the system to panic. It occurs when an I/O error occurs during a read or write of the swap area.
19. An lp destination printer name of 14 character length is disabled by the scheduler, so make the names 13 characters or less.
20. If you *iv(1M)* a SCSI disk and mistakenly specify the size of the last slice to be greater than the actual size, and then make a file system on the slice using the */dev/dsk* device node, the system will crash with a SCSI panic message.
21. Be careful using the *iv(1)* command. *iv -i* of */dev/rdisk/c0d0s0* reinitializes the system disk, which implies destroying the existing CTIX system.
22. The kernel libraries in */usr/sys* and the dfile files in */usr/sys/cf* will be overwritten. If you are customizing

the kernel for a product, please review the product Release Notice.

23. The `-f` option for `as(1)` is not supported.
24. If you are using **pseudoterminals**, you may want to increase the number of **clists**. Consult the *S/Series CTIX Administrator's Guide*.
25. The file `/usr/lib/mailx/mailx.rc` executes a **source** on the file `/usr/lib/mailx/local.rc`, which as delivered is empty. `mailx.rc` is used by `mailx(1)`.

You must create **local.rc**, as in the following example:

```
# local.rc
set sendmail = /usr/lib/sendmail
# Management system
group   mgradmin      greg
alias   greg           adminsys!greg
alias   Greg_Holden   greg
# Inventory system
group   invadmin      sam greg
alias   sam            inventory!sam
alias   Sam_Wooster   sam
```

Refer to the *S/Series CTIX Administrator's Guide* for more information.

12. Documentation Updates

12.1 Available Documentation

The following documents accompany this software release:

1. CTIX Operating System Manual, Version C, Second Edition, in four volumes.
2. S/Series CTIX Administrator's Guide
3. CTIX Administration Tools Manual

12.2 CTIX Documentation Update

Some documentation deficiencies will be addressed in this section.

1. The man page for *mcs(1)* is incorrect. *mcs(1)* does not use the environmental variable **TMPDIR**. It explicitly uses **/usr/tmp**.
2. The man page for *masterupd(1M)* is included in the latest version of the CTIX Operating System Manual, but the utility is not included in the 6.2 CTIX release.
3. 6.2 S/Series CTIX *uucp(1C)* is simpler to configure for autodialers or smart modems. The **/usr/lib/uucp/modemcap** file no longer exists. It is replaced by the **/usr/lib/uucp/Dialers** file. Please refer Chapter 10, "UUCP", in the *S/Series CTIX Administrator's Guide*.
4. The *lp(1)* command has been added to the default commands executable from a remote system with *uux(1C)*. The other defaults are *rmail(1)* and *rnews(1)*.
5. See the *S/Series CTIX Administrator's Guide* for information on how to configure your system for Ethernet.
6. The SMAIL Group installs the *sendmail(1C)* package. Refer to Chapter 15 of the *S/Series CTIX Administrator's Guide* for advice on configuration and operation. The *sendmail* package is difficult to administer, so install it only if you want it.