RELEASE NOTE

Integrated Solutions UNIX† Release 4.0

490158 Rev. B

1.1 About This Release Note

This Release Note summarizes the major features of Integrated Solutions' current UNIX release, Release 4.0, Revision D. When you are finished with this Note, please insert it in the back of your *UNIX 4.3BSD System Administrator Guide* (SMM:1)†† for future reference.

NOTE TO EXISTING CUSTOMERS

If you are already running a version of ISI UNIX 4.2BSD, see the *Field Upgrade Procedure, UNIX 4.2BSD to UNIX 4.3BSD* for specific instructions on upgrading to the current release. If you have an earlier version of ISI UNIX 4.3BSD, refer to Section 1.4.1 for information on upgrading to this new release.

NOTE TO NEW CUSTOMERS

If you are a new customer and have just taken delivery of an Optimum system or workstation, please read your system *Installation Manual* followed by the *UNIX 4.3BSD System Administrator Guide* (SMM:1) for instructions on how to install, power up, and set up your system.

This note contains these major sections:

---

† **UNIX** is a registered trademark of AT&T in the USA and other countries. **4.2BSD and 4.3BSD** were developed by the Regents of the University of California (Berkeley), Electrical Engineering and Computer Sciences Departments.

†† References of the form (XXX:N) refer to a section of the seven-volume UNIX 4.3BSD Reference Set. **SMM:1**, for example, refers to the first section of the *UNIX System Manager's Manual* (SMM).

1.2 Release 4.0

Release 4.0 is the Integrated Solutions release of UNIX 4.3BSD. The 4.3BSD release of the UNIX operating system is fundamentally different from previous 4.2BSD releases.

4.3BSD is backward-compatible, meaning that all programs and binaries developed under 4.2BSD should run under 4.3BSD. However, you should not mix 4.2BSD and 4.3BSD systems for TRFS operations, as commands issued on a 4.3BSD machine may not execute properly in the 4.2BSD environment of another machine. See Section 1.7 for incompatibilities between 4.2BSD and 4.3BSD.

The directory */usr/old* contains some 4.2BSD binaries for execution on 4.2BSD systems. For example, if you want to talk with a user on a machine running 4.2BSD, you must run */usr/old/talk*.

See the *Field Upgrade Procedure, UNIX 4.2BSD to UNIX 4.3BSD* for complete instructions to upgrade an existing 4.2BSD system. See "Bug Fixes and Changes in 4.3BSD" (SMM:12) for specific details of changes in the UNIX operating system from 4.2BSD.

General changes from earlier ISI releases, beyond the changes involved in upgrading to UNIX 4.3BSD, include the following:

• */bin/[cc,pc,f77]*	new compilers replacing the Berkeley compilers (see Section 1.2.1)
• NFS†	the Network File System (NFS), now available as an option (see Section 1.2.2)
• Cluster/diskless	simplified boot procedure for cluster/diskless nodes (see Section 1.2.3)
• */sys/conf*	kernel configuration with **config** (8) (see Section 1.2.4)
• */INSTALL*	new script to automate UNIX installation (see Section 1.2.5)
• */etc/admin*	rewritten **admin** (8) system administration program (see Section 1.2.6)
• **csh** (1), **test** (1)	new *-l* option (see Section 1.2.7)
• */usr/lib/mail*	new directory for */usr/lib* mail files (see Section 1.2.8)
• Memory requirements	memory configurations for this release (see Section 1.2.9)
• Format of tape	format of files on release tape (see Section 1.2.10)
• **date**(1), **zic**(1)	changes in time zone handling (see Section 1.2.11)
• TRFS	ISI's implementation of multihop functionality (see Section 1.2.12)
• */UPDATE*	new script to update earlier Release 4.0 releases to Revision D (see Section 1.4.1)

† NFS (the Sun Network File System) is a product created and developed by Sun Microsystems, Inc. NFS is a trademark of Sun Microsystems, Inc.

1.2.1 New Compilers

Release 4.0 includes three new compilers for C, Pascal, and FORTRAN. These compilers use the same calling conventions for all subroutines, routines, procedures, and functions as the Berkeley UNIX compilers. ISI chose to use these compilers since they offer better optimizing capabilities and show significant improvements in speed over the standard Berkeley compilers. The compilers were included in Release 3.07 as an alternate compiler suite. Effective with Release 4.0, these compilers are the standard ISI compilers; hence, ISI no longer supports the Berkeley compilers.

Additionally, ISI Release 4.0 does not include support for **pi**(1) (Pascal Interpreter), **pix**(1) (Pascal Interpreter) or **fp**(1) (Functional Programming language compiler/interpreter). Support for **lisp**(1), **liszt**(1), and **lxref**(1) is available only by purchasing the Franz Extended Common Lisp software package. Contact your ISI salesperson for information on this package.

Please note that the extensions to Pascal (**pc**(1)) are not entirely compatible with extensions to Berkeley Pascal. See the man pages **cc**(1), **pc**(1), and **f77**(1) for compiler descriptions. Read the *UNIX Compiler Guide: C, Pascal, and FORTRAN 77* for more detailed information about the compilers.

1.2.2 The Network File System (NFS)

You can now purchase the Network File System (NFS) as an option. This network service facilitates communication between a variety of nodes on a local network. Only Revision B of NFS is compatible with this version of Release 4.0 (Revision D).

NFS allows nodes to share file systems across the network. Using layers of network services, nodes execute remote procedure calls to access data and programs on other nodes.

The *NFS Release Note* contains instructions for installing NFS. See the *Network File System (NFS) Guide* and corresponding man pages for a further description of NFS and for programming instructions. Hardcopies of the man pages for NFS-specific commands are contained in the *Network File System (NFS) Guide*; online copies are accessible only by customers who purchase the NFS option.

1.2.3 Booting Cluster and Diskless Nodes

NOTE

This section applies only to systems with CPU boards having PROM revision 4.0 or later.

The syntax for booting a diskless or cluster node has changed. Formerly, when booting a diskless or cluster node, the boot command was:

dev(0,0)*source*:/**vmunix HOST**=*remote* **SERVER**=*server*

where

dev is the network device (**vb, nw,** or **ex**)

source is the hostname of the node where the desired kernel resides

server is the hostname of the server node

remote is the hostname of the diskless or cluster node

To boot a cluster or diskless node now, simply enter the character "@" and press return. The node will boot to single-user mode automatically. The file */etc/bootd.conf* contains the strings that are used when "@"is entered. The bootd.conf file is delivered with sample lines so you can construct your own file if desired.

You can still use the older syntax. In fact, you must use the older syntax the first time that you boot a cluster or diskless node. Then you must set up */etc/bootd.conf* on the server. See the *UNIX 4.3BSD System Administrator Guide* (SMM:1) for further details and instructions. You must also use the older syntax for a diskless node if you want to specify one server among several possible servers, or if you want to use a kernel from one server while booting to another server.

1.2.4 Kernel Configuration with config (8)

The Release 4.0 version of **config** (8) was modified to produce kernels from the binaries in */usr/sys/OBJS.**. To correctly configure the system using these binaries, use the command

config -o *name*

in step 3 of Section 4.2.3 of the *UNIX 4.3BSD System Administrator Guide* (SMM:1). This is documented correctly in Rev. B of the manual, but was incorrect in Rev. A of the SMM:1.

Kernel configuration files now reside in the directory */sys/conf*. If you want to change the kernel, see these documents:

- *UNIX 4.3BSD System Administrator Guide* (SMM:1)
- *Building Berkeley UNIX Kernels with Config* (SMM:2)
- **config** (8)

If you do not want graphics, remove the GWS option and the vty and gp devices. Similarly, if you do not want TRFS support, remove the TRFS option.

Kernel Options	
Option	*Description*
DEBUGGER	Kernel debugger (enables kernel debugger)
GWS	Enables graphics support code in the kernel
QUOTA	Filesystem Quotas (enables quota system for limiting filesystem usage)
TRFS	Transparent Remote File System (enables all TRFS code in the kernel)
VBUS	VME bus
QBUS	Qbus
TB	Enables RS232 support code for graphics tablets
APPLETALK	AppleTalk protocols (enables kernel support for AppleTalk protocols via Ethernet)
ENETFILTER	Required for booting diskless nodes
ETHER	Enables kernel support for Ethernet access
SL	Serial Link (enables kernel support for TCP/IP over serial lines)

Psuedo Devices	
Device	*Description*
vty	Needed for ISI graphics systems
pty	Needed for rlogin/X windows
enet	Needed for diskless boot
gs	Needed for SCSI/U access
sl	Needed for TCP/IP across serial lines

VME Devices	
Device	*Description*
gd	SCSI/U disk device
gg	SCSI/U host adapter interface monitor
gp	ISI GIP/MWS device
ld	OSI laser disk device for old style SCSI host adapter
rd	Diskless swap device
sd	Old style SCSI disk device (requires Adaptec controller)
sio	CPU card serial device
sky	SKY floating point board device
sm	SMD disk device
sp	Spanned disk device
ts	1/4" and 1/2" tape device
ex	Excelan Ethernet device
nw	ISI Ethernet device
vb	Shared memory device (Cluster interface to host)

QBUS Devices	
Device	*Description*
dh	DH-11/DM-11 device
dl	DLV-11J device
dz	DZ-11/DZ-32 device
el	ISI extended rl disk device
hp	HP disk device
ib	IEEE488 device
lp	Line printer
ra	UDA50/RAxx disk device
rd	Diskless node swap device
rk	RK611/RK0 disk device
rl	RL01/02 disk device
rp	Rampage graphics device
rx	RX01/02 floppy disk device
sio	CPU card serial device
sky	SKY floating point board device
sp	Spanned disk device
tm	TM11 tape device
ts	TS11 tape device
il	Interlan Ethernet device
ex	Excelan Ethernet device

1.2.5 Simplified UNIX Installation

The Release 4.0 release tapes include an automatic UNIX installation script (*/INSTALL*). This script assists you in installing UNIX from release tapes.

When you follow the instructions of the *UNIX 4.3BSD System Administrator Guide* (SMM:1) to load UNIX from tape, the script will help you to configure your system, install the correct file systems, and load files from the tape.

1.2.6 Changes to admin (8)

The **admin** (8) system administration program has been rewritten to accommodate changes in network administration and other configuration procedures.

This is the **admin** top-level menu:

ADMIN TOP LEVEL MENU

(A) Display/modify admin parameters

(B) User account maintenance

(C) Diskless/cluster installation and deletion

(D) UUCP maintenance

(E) Printer maintenance

(?) Help

(Q or X) Exit from admin

The first execution of **admin** invokes an initialization routine that sets many basic system parameters (hostname, network address, and so forth). See the *UNIX 4.3BSD System Administrator Guide* (SMM:1) for further details and instructions.

1.2.7 Symbolic Link Test for csh (1) and test (1)

csh (1) and **test** (1) now include an additional test option: the *-l* option, which tests true when the named file is a symbolic link. The *-l* option works in the same manner as other **test** (1) options and **csh** (1) expressions.

1.2.8 Relocated */usr/lib* Mail Files

The following files, formerly in */usr/lib*, are now in */usr/lib/mail*:

aliases	sendmail.cf
aliases.dir	sendmail.fc
aliases.pag	sendmail.hf
	sendmail.st

Any references to */usr/lib/aliases*, for example, should now be to */usr/lib/mail/aliases*. Standard UNIX utilities now recognize the new location; you should alter any local programs that may refer to these files. This change is for easy administration of cluster and diskless configurations.

If you are implementing **uucp** at your site, you must alter the default */usr/lib/mail/sendmail.cf* file on the Release 4.0 release tape. To do this, use a UNIX editor (for example, **vi** (1)) to make the following changes to the */usr/lib/mail/sendmail.cf* file.

1. Locate this block of code near the top of the file:

```
# domain
DDXXX
CDLOCAL XXX
```

Change the string "XXX" to your local domain name. Do this for both occurrences. If you do not have a local domain name, choose one now. For more information on domains, see the *Sendmail Installation and Operation Guide* (SMM:7) and the following document:

Su, Zaw-Sing, and Postel, Jon
The Domain Naming Convention for Internet User Applications
RFC819
Network Information Center
SRI International
Menlo Park, CA
August 1982

2. If there is a host on the local net with access to the outside world via uucp, find the lines:

```
# host on LAN with UUCP (or other) connection
DRUNKNOWN
```

in **/usr/lib/mail/sendmail.cf**. On all machines on the local net, UNKNOWN should be changed to the hostname of the machine with the external connection. On the machine with the external connection, UNKNOWN should be changed to the hostname of the off-site node through which mail passes. Additionally, on the machine with the external connection, find the lines:

```
# forward other domains to a relay host ; change mailer to uucp to
# forward to outside domains
R$*<@$+.$+>$*                   $#ether$@$R$:$1<@$2.$3>$4user@host.DOMAIN
#R$*<@$+.$+>$*                  $#uucp$@$R$:$1<@$2.$3>$4user@host.DOMAIN
```

Comment out the line with *ether* and uncomment the line with *uucp*. This will allow mail to be forwarded from the local net to the outside world.

1.2.9 Memory Requirements

This release of the UNIX operating system requires these memory configurations:

- ≥2 Mbytes of board memory
- ≥6 Mbytes of disk for the root (/) file system
- ≥60 Mbytes of disk for the */usr* file system

1.2.10 Format of Release 4.0 Tape

Table 1-1 shows the format for the Release 4.0 tape. When running the */INSTALL* program, you will be prompted to insert the second tape.

Table 1-1. Format of the Release 4.0 Tapes

File			
1/4-Inch Tape	1/2-Inch Tape	Contents	Description
(tape 1)0-8	(tape 1)0-8	**miniroot**	the "miniroot" file system, with install software
9	9	**root**	**dump** (8) of root (/) file system
10	10	**kernels**	**tar** (1) of kernels for the root file system
11	11	**bin**	**tar** of */usr/bin*
12	12	**ucb**	**tar** of */usr/ucb*
13	13	**doc**	**tar** of */usr/doc*
14	14	**man**	**tar** of */usr/man*
15	15	**games**	**tar** of */usr/games*
16	16	**new**	**tar** of */usr/new*
17	17	**old**	**tar** of */usr/old*
(tape 2)0	(tape 2)0	**sys**	**tar** of */usr/sys*
1	1	**lib**	**tar** of */usr/lib*
2	2	**etc**	**tar** of */usr/etc*
3	3	**dict**	**tar** of */usr/dict*
4	4	**include**	**tar** of */usr/include*
5	5	**bench**	**tar** of */usr/bench*
6	6	**adm**	**tar** of */usr/adm*
7	7	**spool**	**tar** of */usr/spool*
8	8	**guest**	**tar** of */usr/guest*
9	9	**hosts**	**tar** of */usr/hosts*
10	10	**local**	**tar** of */usr/local*
11	11	**msgs**	**tar** of */usr/msgs*
12	12	**preserve**	**tar** of */usr/preserve*
13	13	**pub**	**tar** of */usr/pub*
14	14	**tmp**	**tar** of */usr/tmp*
15	15	**graph1**	**tar** of graphics files in */root*
16	16	**graph2**	**tar** of graphics files in */usr*
17	17	**X**	**tar** of X Window system files
18	18	**UPDATE**	**tar** of */UPDATE* script

1.2.11 Changes in Time Zone Handling

Significant changes have been made to the way time zones are handled by **date**(1). This command no longer uses the time zone information maintained in the kernel, although this kernel information can still be changed by **date**(1) as documented in the man page.

There are two ways to display the correct time zone.

1. Use the environment variable TZ. To do this, add the following line to all users' *.cshrc* files:

   setenv TZ *time zone*

   where *time zone* is the appropriate time zone abbreviation from */etc/timezone* (e.g., PST8PDT). This line should also be added to */usr/lib/admin/misc/.cshrc* so that all users added via the **admin**(1) program get this addition.

2. Create a time zone description file and compile it with **zic**(8) using the -l option. The man page, a copy of which is attached to this document, contains additional information. An example description file is in */etc/zoneinfo/example.zic*.

1.2.12 TRFS Multihop Functionality

Previous releases of ISI UNIX did not support pathnames that involved multiple symbolic links in a disjoint system. This release accommodates multiple hops between machines. For example, suppose you have a symbolic link from machine A to a file on machine B which is a symbolic link to a file on machine C. The kernel first evaluates and quantifies the link on machine B; it is interpreted as a link between machine A and machine C.

Loops in symbolic links are not detected by the kernel. A loop in a symbolic link will result in a "too many links" error message. The limitation on hops between machines is 20 links. Users with pre-Revision D software will occasionally see the following error message on the console:

TRFS: unknown request 52?

During the implementation of multihop functionality in Revision D, we discovered that ISI's interpretation of symbolic links over TRFS was incorrect. The current interpretation of */@*/tmp* is the */tmp* directory on system *. The correct interpretation of a symbolic link over TRFS is that any path beginning with "/" should be relative to the default root or a root set using the **chroot**(2) command. ISI has chosen not to change the current interpretation of */@*/tmp* because we feel that many symbolic links already in use over TRFS may be adversely affected. To allow **chroot**(2) to work as it was intended over the TRFS, we have added a new TRFS syntax. Symbolic links of the form /@/... will be interpreted relative to your current root **not** relative to your current location.

For example: */tmp* on system B is a symbolic link to */@/tmp*. If a user on system A executes **cd** */@B/tmp/f1*, *the* */tmp* referenced is located in */@A/tmp*. Likewise, if the same user does a **chroot** to */@B/* and then a **cd** to */@B/tmp*, the */tmp* on system B will be used.

NOTE

ISI has not changed the current interpretation of symbolic links over TRFS; therefore, existing links are not affected. This new capability is provided for users who need **chroot** to work over TRFS.

1.3 Documentation

Throughout this document and other UNIX documents, you will see references of the form (XXX:N). These references refer to sections of the seven-volume UNIX 4.3BSD Reference Set. **SMM:1**, for example, refers to the first section of the *UNIX System Manager's Manual* (SMM). The seven volumes are listed first in Table 1-2.

All manuals for this release contain a *Documentation Comments Form* as the last page. After using the manual, please take a minute to complete this form and return it to ISI. Your input will help the ISI Technical Publications Department identify and respond to your specific documentation needs.

If you have little or no experience with UNIX operating system, ISI recommends that you read one of the following books before tackling the ISI documentation. The first two books are available through ISI. All of the books are available at local bookstores.

Henry McGilton and Rachel Morgan, *Introducing the UNIX System*, New York: McGraw-Hill, 1983.

Gail Anderson and Paul Anderson, *The UNIX C Shell Field Guide*, New Jersey: Prentice-Hall, 1986.

Peter Birns, Patrick Brown, and John C. Muster, *UNIX for People*, New Jersey: Prentice-Hall, 1985.

The Waite Group, *UNIX Primer Plus*, Indianapolis, IN: Howard W. Sams and Co., 1983.

Sandra L. Emerson and Karen Paulsell, *troff Typesetting for UNIX Systems*, New Jersey: Prentice-Hall, 1987.

Table 1-2 shows a summary of the documentation that supports Release 4.0.

Table 1-2. Part Numbers for Documentation

Manual	Part Number
UNIX User's Reference Manual (URM)	490143
UNIX User's Supplementary Documents (USD)	490144
UNIX Programmer's Reference Manual (PRM)	490145
UNIX Programmer's Supplementary Documents, Volume 1 (PS1)	490146
UNIX Programmer's Supplementary Documents, Volume 2 (PS2)	490147
UNIX System Manager's Manual (SMM) *includes the System Administrator Guide* (SMM:1)	490148
User Contributed Software (UCS)	490149
Release Note for Release 4.0	490158
UNIX Compiler Guide: C, Pascal, and FORTRAN 77	490168
Programmer's Reference Manual for Graphics Software	490046
UNIX Programmer's Manual Supplement, Graphics System Software	490047
UNIX Documentation Roadmap	490174
Desktop Manager User's Guide	490048
Browse User's Guide	490107
Paint Program User's Guide	490059
Optimum V WorkStation X Window System Manual	490141

1.4 UNIX Upgrade for Existing Customers

There are two ways to upgrade an existing system to the current release. The method you choose depends on the version of software you are currently running.

- If you are running Release 4.0, pre-Revision D, do an incremental upgrade using the */UPDATE* script on the Release 4.0 release tape. This will add or modify files on your system to bring the system to Release 4.0, Revision D. This is documented in Section 1.4.1.
- If you are running Release 3.07 (or an earlier release of 4.2BSD), save your local (modified at your site) files, load the new release, and restore the local files. This puts a completely new release of UNIX on your disk. See Section 1.4.2 for instructions.

1.4.1 Incremental Upgrade for Existing Release 4.0 (4.3BSD) Customers

If you are currently running Release 4.0, Revision B (or an earlier revision), use this upgrade procedure to install the modified Release 4.0, Revision D files on your system.

Please note these cautions concerning the incremental upgrade procedure:

- The files in this upgrade will overwrite selected files on your system. Before overwriting anything that you may want to save, the upgrade script will save most local files. The files that need not be updated

will be restored from these saved files. However, you should do a level 0 dump of all file systems before upgrading, as a precaution.

- If your system has standard distribution files that are actually symbolic links to files on another machine (TRFS links), the upgrade script may not be able to install new versions of those files. This will generate permission error messages during the upgrade procedure. For example, if */usr/man* is a symbolic link to */usr/man* on another system *(/@systemname/usr/man)*, upgrades to */usr/man* files may fail.
- If you hold a source license, and */usr/sys* is a symbolic link to your source directory, you should answer "yes" when the upgrade script asks if you want to save your current */usr/sys* directory. This will prevent the extraction of the kernel binaries. Answering "no" may corrupt your source directories.
- The */root* for Release 4.0, Revision D is approximately 10% larger than other revisions of Release 4.0. Be sure your */root* partition has as few extra files as possible before running the */UPDATE* script.

Use this procedure to incrementally upgrade to Release 4.0, Revision D.

1. Boot the computer in single-user mode. See the *4.3BSD System Administrator Guide* for instructions on doing this.
2. Set the hostname by issuing the **hostname**(1) command.
3. Check that the */usr* file system is mounted. See **mount** (8) for details.
4. Load the second Release 4.0, Revision D release tape into the tape drive. Enter the command:

   mt rewind

   to ensure that the tape is rewound.
5. Enter the command

   mt fsf 18

   to position the tape at tape file #18 for the */UPDATE* script.
6. Enter the commands

   cd /
   tar xpbf 20 /dev/nrmt0

   to remove the */UPDATE* material from the tape.
7. Enter the command

   UPDATE

   to execute the */UPDATE* script. This script will prompt you for configuration information and step through the update procedure.
8. Reboot the system. If you are running NFS or System V, you must get NFS or System V releases from ISI that correspond with Release 4.0, Revision D.

1.4.2 UNIX Upgrade for Existing Release 3.07 (4.2BSD) Customers

The following instructions are intended for the user who is already running an ISI release of UNIX 4.2BSD. It is only a general description of the upgrade procedure. For detailed instructions, see the *Field Upgrade Procedure, UNIX 4.2BSD to UNIX 4.3BSD.*

If you are not running such a release already, refer to your *Installation Manual* and to the *UNIX 4.3BSD System Administrator Guide* (SMM:1) for instructions on the start-up procedures for a new system.

If you want to upgrade an existing system to the current release, you must follow this general procedure:

1. Using **tar**(1), archive all user (local) files on tape.
2. On a second tape, archive important system files that contain current information about your system.
3. Bring down the system, using the **shutdown**(8) command.
4. Follow the instructions in the *UNIX 4.3BSD System Administrator Guide* (SMM:1) for building the new release on disk.
5. If you have a WorkStation or if this system will be the server node for graphics cluster or diskless nodes, be sure to load graphics software from the release tape.
6. Restore the files that you archived with **tar**(1).
7. Since the */usr* file system will be quite large, you may need to delete files that will not be needed. See the *Field Upgrade Procedure* for a list of files for possible deletion.

1.5 Fixed Bugs in Release 4.0

This is a list of bugs which have been fixed since the last release.

1. **adb**(1)/**dbx**(1)
   - Displaying registers and global variables, the call function, and single stepping problems have been corrected in **dbx**(1).
   - High memory and user defined variables are displayed by **adb**(1).
2. **admin**(8)
   - User interface code for **admin**(8) has been completely rewritten for this release.
   - Problems with cluster/diskless node installations have been eliminated.
   - Mail file installation now operates correctly.
3. Spanned Disk
   - Spanned disk installation is supported during the installation phase.
4. Graphics
   - Using the text marking capability of the Desktop Manager no longer causes a system panic.
   - Problems with the kernel panicing for out of vt (Virtual Terminal) data space have been avoided by more carefully tracking data usage and issuing error messages when not enough space is available.
5. Utilities
   - Minor bugs have been fixed in create, **rcp**(1C), **more**(1), **refer**(1), **pstat**(8), **tip**(1C), **csh**(1), **killpg**(8), **rwhod**(8C), **swapon**(8),**strings**(1), **learn**(1), **quot**(8), **more**(1), and **man**(1).
6. Diskless/Cluster Swap Area Reclamation
   - This kernel removes all references to the swap files for diskless/cluster nodes such that they can be removed when diskless/cluster nodes are not in use.

- The kernel does not truncate the swap file so swap files can be fully allocated and remain that way. This allows you to preallocate swap files completely, thus avoiding the possibility of diskless/cluster nodes panicing for out of swap when the */usr* filesystem is full.

7. TRFS
   - A problem which allowed the root user of a diskless node to change files on a system other than its server has been corrected.
   - Symbolic links between two servers now work correctly.
8. NFS
   - Binaries can now be executed over NFS.
9. VB
   - A problem with the vb driver has been corrected that caused transfers of larger than 2048 bytes to fail.
10. cc(1)
    - This release does not implement any optimizations except the removing of unused code and data when the -O flag is not specified.
    - Subtracting 0 from an unsigned char produces the correct value.
    - A warning message appears when a program contains a call to alloca and the -X23 command line flag is not used.
    - The -g flag has been fixed so that it handles static variables correctly.
    - The assembler now generates the correct code for the movp instruction.
    - The -R flag is now passed to the assembler.
    - Profiling and gprof are supported in this release.
    - A problem that prevented the 4.2BSD C compiler from terminating when compiling a short program that used unsigned long variable types has been eliminated in 4.3BSD.
    - The C compiler no longer generates bad assembly language code for unsigned character operations. This was a problem with the 4.2BSD compiler.
    - Stack variables are long aligned where appropriate.
    - The sparse switch statements operate correctly in this release.
    - A problem with namelists not being generated for programs with no code in them has been eliminated.
    - The floating point arithmetic used by the compiler for constant evaluation is now accurate.
    - A problem with optimizations of static variables in recursive procedures has been fixed.
11. f77(1)
    - Units 5 and 6 can be used for I/O other than stdin/stdout.
    - A version of libbmf, compiled for the 68881 floating point coprocessor, is now included.
    - A bug in the cabs library routine has been fixed.
    - Lib.f77.a contains Z-abs.

- The -g flag can be used for compiling.
- A problem that caused programs compiled for the 68881 to terminate incorrectly on systems without the 68881 has been fixed by doing the appropriate testing before program execution begins.
- The IOSTAT command works correctly now.
- A linking problem caused by referencing the IOINIT subroutine has been fixed.
- Flmin, flmax, ffrac, dflmin, and dflmax no longer cause core dumps.
- The floating point arithmetic used by the compiler for constant evaluation is now accurate.

12. **pc**(1)
    - Two copies of the Pascal function "new" were linked into the Pascal library.

1.6 Known Problems with Release 4.0

1. User processes have been observed to go to sleep on heavily loaded systems and not wake up. This has happened after as little as 9 hours of continuous execution and as much as 28 hours.
2. UNIX 4.3BSD uses a different size for directories from UNIX 4.2BSD. When you install 4.3BSD, 4.2BSD directories may not have the correct 4.3BSD size (a multiple of 512). This will be noted and corrected when they are **fsck**'d the first time; disregard the associated error messages.
3. A write-protected tape may report an invalid error on rewind (there is no real error, only the message appears). Remove the write protection and the problem will go away. This problem will cause the installation scripts to report errors and abort if the release tape is write protected.
4. To log out from the graphics **demo** account or any account running */usr/isi/dg*, you must first kill the "Disk" window by placing the cursor in the window, pressing the middle button, and selecting "Exit."
5. Security issues have not been completely resolved in the current release of TRFS.

   One protection domain (i.e., one common */etc/passwd* file) should be established across the entire TRFS network in order to preserve the normal level of file security. When a file access request is then made via TRFS, the requester's user ID and group ID are shipped as part of the request, and the normal protection mechanisms apply.

   If a single protection domain is *not* established, then any user who does not have an account on a particular remote machine running TRFS can nonetheless access files there, using the user ID and group ID in the *passwd* file on the local machine. The user ID and group ID are subject to the normal protection mechanisms. As an exception, someone with user ID 0 on one machine running TRFS, will have the permission privileges associated with user ID -2 on any other machine on the TRFS network. This offers some protection against root access.

   If you desire a common protection domain across the network, you may wish to disable TRFS by building a new kernel without the TRFS option.
6. The **tar** options **–r** and **–u** do not work on quarter-inch cartridge tape drives. The QIC 02 industry-standard interface, the specification for the VME-QIC2 cartridge tape controller, causes this limitation. If these options are used on a quarter-inch drive, the results are unpredictable. This is a permanent limitation.
7. Optimum systems (based on the Q-bus) can only function as servers for diskless nodes when using an Excelan Ethernet controller. Interlan Ethernet controllers do not support TRFS.

8. The current release may not update graphics WorkStation windows correctly. A message to the console window, for instance, can interrupt a window movement on the screen and cause false window images to remain. You can erase these ghost images by moving an active window over them and then moving the window away; this will update that screen area. This problem is applicable to graphics WorkStations only.

9. Floating point exceptions are not generated for division by zero or other similar situations. Currently, the IEEE value for NOT A NUMBER is returned by the math library (both the software and the 68881).

10. The VME-ECX (Ethernet) board is incompatible with the older version (Revision A) of the VME-SMEM (cluster shared memory) board due to conflicts in 4.3BSD address space. This problem applies only to cluster systems. The solution to this problem is to use either a VME-EC (Excelan Ethernet) board or a newer version (Revision B or later) of the VME-SMEM board. These are the acceptable configurations:

    - VME-EC with any VME-SMEM
    - VME-ECX with new VME-SMEM (Revision B or later, shipped after December 1986) only

    See Figure 1-1 to identify the new and old versions of the VME-SMEM board. Two sets of jumpers on the Revision A board are missing on the Revision B board.

11. On systems with NFS, NFS does not support directory reads using the read system call. To access a directory, use the library calls provided for this purpose.

12. ISI does not support the SKY floating point boards in this release.

13. This release contains support for AppleTalk†. This feature is not delivered in the distribution kernel but it is possible to build kernels that have this feature enabled. If eight ISI systems are booted on the same network, all with this feature enabled, kernel panics ("mclfree") may occur.

14. There is a known problem in the F77 compiler regarding static variables in a recursive subroutine. This problem only occurs when the optimizer (-O) is turned on.

15. On Page 4-2 of the *UNIX 4.3BSD System Administrator Guide* (SMM:1), step 2 contains an incorrect filename. The correct filename is */sys/is68k/conf.c*.

16. The printcap file needs to be configured according to your environment.

17. Using *ps*, *ps -u*, or any variations of *ps* that list the processes of the current user display only the header when run on a diskless or cluster. Use *ps -aux | grep $user* as a workaround. The command *ps -a* works as expected by listing all processes.

18. The negative of an unsigned integer is treated as a signed integer.

19. Dflmin and dflmax return single precision values instead of double precision.

20. Unsigned integers are not cast to floats when expected.

21. The *UNIX User's Supplementary Documents (USD)* contains documentation on two programs that are not supported by ISI in Release 4.0. ISI does not support either Notesfiles (USD:11) or JOVE (USD:17).

22. The *UNIX User Contributed Software (UCS)* contains documentation on Courier. This remote procedure call program is not supported by ISI in Release 4.0.

† **AppleTalk** is a registered trademark of Apple Computer, Inc.

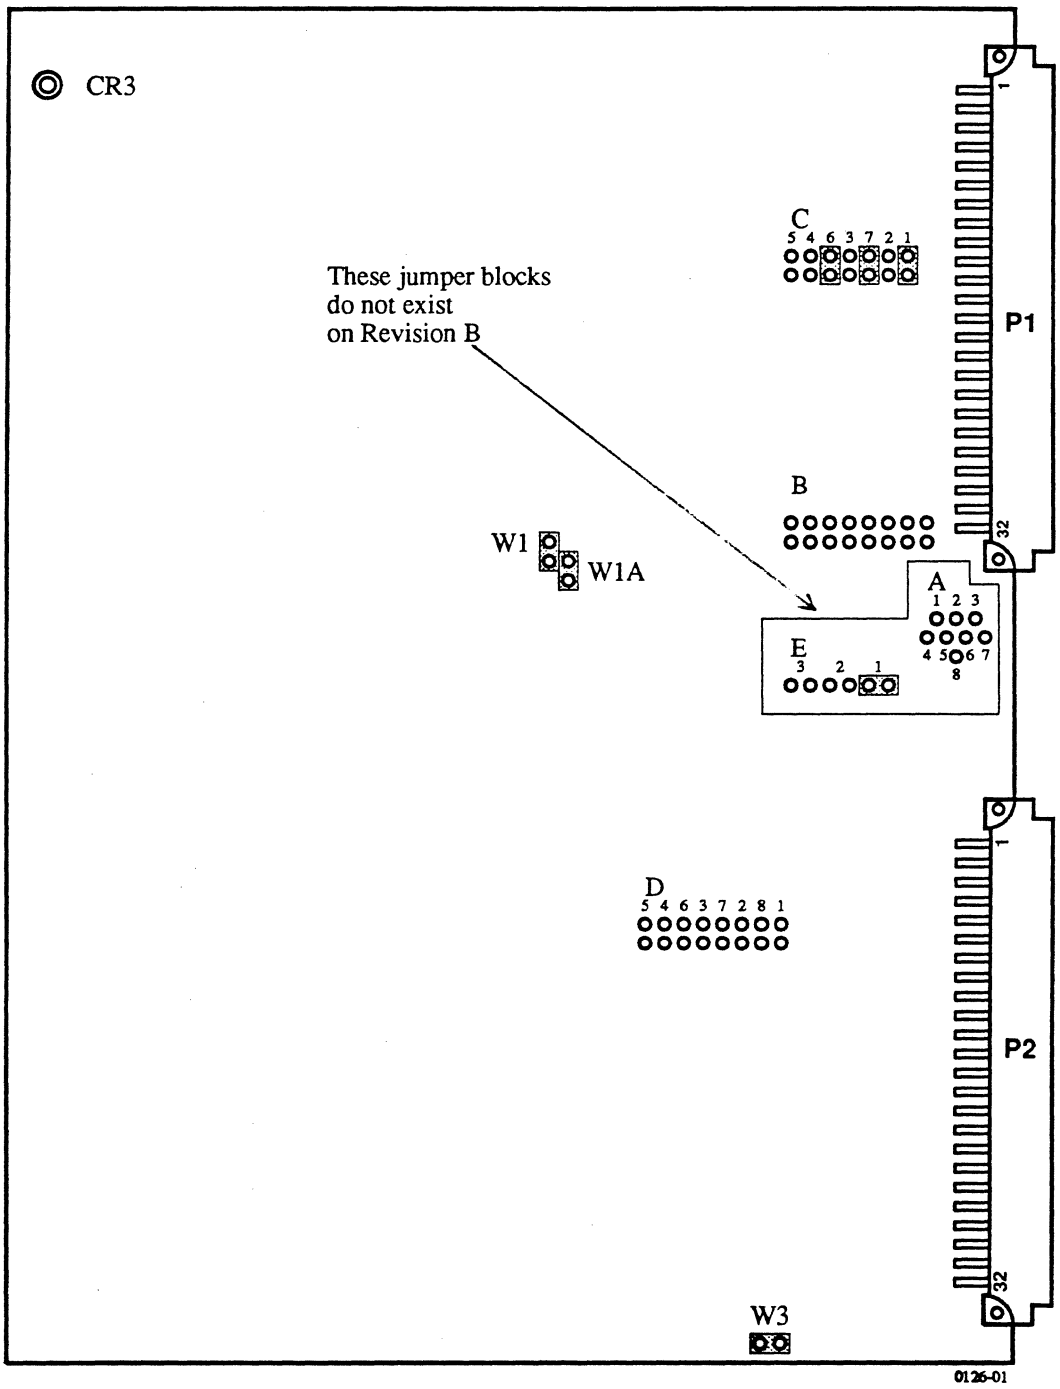

Figure 1-1. Identifying VME-SMEM Revision Level

1.7 Known Incompatibilities Between 4.2BSD and 4.3BSD

1. A 4.2BSD binary that uses the "getlogin" system call on a 4.3BSD system will always have a null string returned. This is not a bug. It is a difference between 4.2BSD and 4.3BSD which occurs because the format of the */etc/ttys* file differs between the two releases.
2. The Release 4.0 version of **ifconfig** (8C) does not distinguish between upper- and lower-case characters. This is because **gethostbyname** (2) maps all characters to lower case. Therefore, you should use only lower-case letters in hostnames to avoid confusion. This is a permanent feature of UNIX 4.3BSD.
3. Csh script files that are setuid must be executed with the -b command line flag.

1.8 UUCP Access to Technical Support Data Base

The Technical Support Department at Integrated Solutions maintains a data base, accessible over uucp. All customers are entitled to report bugs and to pick up the list of known bugs and bug fixes.

1.8.1 Data Base Organization

Table 1-3 shows the organization of the technical support data base.

1.8.2 Accessing the Data Base

Use this procedure to access the ISI Technical Support data base.

1. Ensure that uucp has been set up on your system.
2. Enter this line in your */usr/lib/uucp/L.sys* file:

   isi Any ACU 1200 4084340782 ogin-EOT-ogin-EOT-ogin: isiisi ssword: saturn

   Remove or add to the "408" in the phone number if your are in the local dialing area or if you require additional digits to make the call.
3. To send bug reports to ISI, use */usr/ucb/sendbug*. *sendbug* is a shell script that mails to ISI. If you include a line of the form:

   REPLY TO: *yourhostname yourloginname*

   ISI will queue a reply within five days indicating receipt of the bug report and stating primary disposition. These are the possible dispositions:

   — New bug

   — Known bug

   — Previously fixed bug

   — Unreproducible from information given

   — Not a bug

   In the last case, an explanation of why will be given. In the first three cases, a bug report number will be given, which can be used to track the status of the bug in the FIXED and BUGS files.

   It is the responsibility of the sender to poll ISI and pick up the reply. Do this with the command

   uupoll isi

4. To pick up a list of fixes (while in your home directory) enter:

   uucp isi\!~uucp/how/FIXED .
   or
   uucp isi\!~uucp/how/FIXED ~uucp

5. To pick up a list of changed files (while in your home directory) enter:

   uucp isi\!~uucp/how/NEWS .
   or
   uucp isi\!~uucp/how/NEWS ~uucp

Table 1-3. Technical Support Data Base Organization

Directory	Purpose
~uucp/how	Files with information about our file structure. This directory is readable by all customers. It includes README - describes the data base structure/access BUGS - list of known bugs. FIXED - list of fixes NEWS - list of files which have changed since the current release
~uucp/bugs	Files dealing with bugs. For example, if you have a program that demonstrates a bug, you could leave it here. This directory is accessible to all customers. It is the only directory writable by customers.
~uucp/root	This is a copy of the root of the most current root file system, with */usr* mounted on it, and any files that have been modified since the last release. All files are in the same place as on a newly installed system. All files are in compressed format.
~uucp/extra	Extra contains subdirectories containing additional software available from us. These directories are only accessible to registered owners of the software who have signed up for support.
~uucp/new	Special files for individual customers. Not generally accessible.

6. To pick up a list of bugs (while in your home directory) enter:

   uucp isi\!~uucp/how/BUGS .
   or
   uucp isi\!~uucp/how/BUGS ~uucp

7. To copy fixed or new files from the uucp root file system, *~uucp/root/*, specify complete pathnames. All files are located as in a distribution tape, and are in "compressed" format to reduce download time and cost.

   For example, to get a new copy of **more**, use

   uucp isi\!~uucp/root/usr/ucb/more ~uucp
   mv /usr/ucb/more /usr/ucb/more.old
   cat ~uucp/more | /usr/local/uncompress -c > /usr/ucb/more

   The commands above uncompress the file and install it, making sure to save the old copy of */usr/ucb/more*. Check that ownerships and permissions are correct for the new file.

   If you do not have a copy of */usr/local/compress* you can pick it up in uncompressed format from *~uucp/new* along with the man page, *compress.1l*. Install *compress* in */usr/local* and link it to

/usr/local/uncompress. The program will compress or uncompress depending on how it is invoked.

8. For those with graphics systems, updated graphics software is in the directory *~uucp/extra/graph*. This directory contains only files that have been modified since the last release. For example, a new **erode** could be transferred with

   uucp isi\!~uucp/extra/graph/.desktop/erode .

   This directory is not accessible unless your system name (set by **admin** (8)) has been registered in our data base of supported systems.

1.9 Revised Man Pages to Insert in UNIX Documentation

The attached man pages were revised after the documentation was printed. Please replace the existing man pages with the attached copies as follows:

Revised Man Pages		
Man Page	Document	Part Number
date(1)	UNIX User's Reference Manual	490143
ftp(1C)	UNIX User's Reference Manual	490143
grep(1)	UNIX User's Reference Manual	490143
rsh(1C)	UNIX User's Reference Manual	490143
tar(1)	UNIX User's Reference Manual	490143
mkpasswd(8)	UNIX System Manager's Manual	490148
zic(8)	UNIX System Manager's Manual	490148