Technical Support Bulletin

Issue 1984-1
13 January 1984

Technical Support Bulletin

Sun Microsystems, Inc.

Issue 1984-1
13 January 1984

1.1) Grounding of line terminator on D169 Fujitsu disk drive

Applies to: Only to systems with D169 (Fujitsu M2289, 169 MB) disk drives
Description:

The small line-termination board is used in socket OM1 at the back of the drive if there are no other disks on the same controller, or if the D169 is the last disk in a daisy-chain configuration. The line-termination board has a grounding wire, which the Fujitsu manual instructs you to fasten to the terminal "TRM1-2".

This wire should go to the screw labeled "0V(2)" (0v means "0 volts" = ground). Several customers mistakenly connected the grounding wire to the screw directly above the label "TRM1", which is NOT a proper ground. This causes intermittent disk errors, leading to random intermittent program and system crashes. This item also appears in "Read Me First" and (beginning with Release 1.1) in the System Manager's Manual.

See Also:
Fujitsu M2289 hardware manuals.

1.2) Miscellaneous f77 bugs

Applies to: Release 0.3, 0.4, 1.0
Description:

f77 -I2 flag problems: The -I2 command-line option to f77 (telling the compiler to generate 2-byte integers) can generate incorrect code for large subscripts and for do-loop indices. Fixed in 1.1.

Description:

f77 'rand' funtion generates incorrect values: The f77 'rand' function is supposed to generate random values uniformly distributed between 0 and 1.0. It actually generates values somewhere between 1 and 2, and not uniformly distributed. Fixed in 1.1.

Description:

f77 -O flag (optimize) can cause compile to fail: Compiling some f77 programs with the -O flag will produce a compiler error message, something to the effect of "expression too complicated". Work-around: removing the -O flag makes the problem go away. Fixed in 1.1.

1.3) Bad back-panel Ethernet connectors on some Model 150's

Applies to: Model 150's only, shipped in late 1983.
Description:

On some Model 150's, the back panel Ethernet connector fails to make proper contact with the Ethernet transceiver cable connector. This is caused by two related problems:

1) The washers under the screws which hold the connector to the back plate prevent the connector from seating properly.

2) The back plate is too thick.

The symptom is a dead Ethernet ("No server" or "connection timed out" messages). Handier customers can correct the problem in the field. First, confirm the problem by removing the back plate, unscrewing the connector from the back plate, and plugging the internal cable directly into the external cable. If the system now functions correctly, the fix is then made permanent by discarding the offending washers and filing down the inside of the back plate to make it thinner in the affected area before re-installing.

Less-handy customers should be referred to Field Support for replacement of the Ethernet connector backplate.

1.4) Ethernet debugging hints

Applies to: All systems with Ethernet.
Description:

A commonly-reported problem is a non-functioning Ethernet. The customer carefully hooks up his machines, only to get a message like "No server– giving up" (when trying to boot a diskless client) or "Connection timed out" (when trying to rlogin between two stand-alone systems).

Following are some suggested remedial actions. After each step, the net should be tested again.

1) Check /etc/hosts and the "ether" lines in /etc/nd.local to make sure that these files contain the correct Ethernet address for each host.
[NOTE: In 0.3 and 0.4 systems, there are no "ether" lines in /etc/nd.local.]
[NOTE: In 1.0 and later systems, the Internet addresses in /etc/hosts need not be the hardware Ethernet address, but the addresses must be consistent in the /etc/hosts files on all machines.]

2) In the server/client case, you can try specifying the hex address of the server in the boot command. The syntax is "b nd(*xxx*)vmunix" where *xxx* is the low-order three bytes of the hexadecimal ethernet address of the server. If this works (while the default boot doesn't) the server files are not configured correctly. Check /etc/hosts and/or /etc/nd.local.

3) If the affected workstation is a Model 150 that was received around the end of 1983, check for bad backpanel Ethernet connector [see item 1.3, "Bad back-panel Ethernet connectors on some Model 150's", above].

4) With its host system powered-on, each transceiver should feel slightly warm to the touch. If a transceiver is cold, it probably isn't receiving power. This could indicate a loose connection on either end of the transceiver cable, a loose connection of the internal Ethernet cable to the Ethernet board, or a faulty cable, transceiver (less likely), or

Ethernet board (even less likely). After checking for a bad backplate connector (step (3) above), try to narrow down the problem by checking and swapping suspected components; then refer the problem to Field Service if necessary. You could try the remaining hints below before giving up.

5) Check all Ethernet coax and transceiver cable connections to make sure they are solid. If boards were changed or moved in the workstation card cage, make sure that the internal connector (from back panel) is plugged into the Ethernet board.

6) On stand-alone host or server, try rlogin to yourself. Make sure that the network daemons are running on the machines that want to talk to each other. (In 0.3/0.4, rsh and rlogind; in 1.0 and above, inetd).

7) The command "netstat -i" (or other variations of netstat) tells you how many packets a machine thinks it's transmitting and receiving on each local network. (Look at the Ethernet network(s) named "ec"; the loopback-net "lo" doesn't actually go out on the wire). For example, on a server you may see the input packet count increasing each time a client tries to boot, while the output packet count remains steady. This suggests that the server is seeing the boot request packets from the client, but doesn't realize it's supposed to respond to them. This might be caused by an incorrect address in /etc/hosts or /etc/nd.local. On the other hand, if the input packet count is steady, the machine isn't seeing the packets at all, which suggests a different failure mode, possibly a hardware problem.

8) Is the network terminated on both ends? (It should be). Remove one of the terminators. You should measure about 50 ohms across the coaxial connector from which you unscrewed the terminator. This check is particularly indicated when the customer is using the "vampire tap" (clamp-on) type of transceiver, which tends to short-circuit the Ethernet coax cable. The threaded connectors on the type of transceivers supplied by Sun can also develop shorts at the connectors.

9) Remove one of the terminators, and try operating the network. You should get error messages on every machine, something to the effect of "Ethernet transmission error". If a machine just continues to give its previous error ("Connection timed out" or "no server") then it may not be connected correctly to its transceiver. As in step (4), the problem could be loose connections or a faulty connector, cable, transceiver, or Ethernet board.

10) Try swapping transceivers and/or transceiver cables. If spare Ethernet boards are available, try swapping boards. Even if there are only two machines on the net, exchanging transceivers, cables, or boards may be informative. When swapping Ethernet boards, remember to edit the Ethernet addresses in /etc/hosts and/or /etc/nd.local where applicable.

11) If you still haven't solved the problem, contact Field Service.

See Also:

System Manager's Manual
nd(4) manual page

1.5) Warning: backing up diskless clients (ndl partitions)

Applies to: All servers with diskless clients.
Description:

On a server, the /pub file system and diskless clients' file systems are allocated (in /etc/nd.local) as pieces of larger "label" partitions. On a machine with one disk, the public and client partitions are usually sub-pieces of partition g of disk 0 (for example, /dev/xy0g).

However, if you run "dump" on device /dev/xy0g on the server, you will only dump the /pub file system, NOT the client file systems. To dump the client file systems (or run fsck on them, or whatever) you must treat them as separate disk partitions, usually known on the server as /dev/ndl0, /dev/ndl1, and so on. The last digit on the "user" line in /etc/nd.local (if non-negative) is the "ndl" device number of the corresponding client partition. The lines ending in "-1" are usually swap areas, which are not file systems.

It is advisable to halt the client machine whenever accessing its ndl partition on the server.

See Also:

System Manager's Manual
nd(4) manual page

1.6) vp: unsupported ikon Versatec driver

Applies to: 0.3, 0.4, 1.0, 1.1
Description:

There is an unsupported driver for the Ikon 10071-5 Multibus Versatec parallel printer interface. This is included as part of the kernel in 0.3 and later releases. However, it is documented only in Release 1.0 and later (System Interface Manual, "vp(4)"). See 1.0/1.1 documentation for further details.

Starting with Release 1.1 there will be a *supported* driver for the Systech Centronics/Versatec parallel interface board (VPC-2200). The Ikon driver may be dropped from future software releases.

See Also:

vp(4) manual page

1.7) Why does df show less space than total size of disk?

Description:

Customers have inquired why the amount of free disk space shown by "df" seems less than the advertised size of the disk. Following is an answer that was sent to a customer, with reference to the D84 "84 Megabyte" Fujitsu disk (now advertised as "65 Megabytes" in our latest price list). The answer is analogous, of course, for other sizes of disk.

There is about 65MB of formatted space on the D84 drive. (There is about 6% less with the Xylogics 450 than with the Interphase 2180, since we use 32 sectors/track with the Xylogics and 34 sectors/track with the Interphase). Assuming we're talking about a typical network disk server, on the first disk we allocate space for the root file system (partition a), swap space (partition b), and the "network disk" partition (partition g) which includes the public file system plus possibly client partitions. The sizes of these should add up to the

total size of the disk, about 65MB. On the second disk, we usually use partition c (the whole disk) whose size should be also about 65MB. You can read the partition sizes in the diag program with the v command (verify label), or from UNIX using the /etc/dkinfo command (e.g., "/etc/dkinfo xy0" or "/etc/dkinfo xy1"). The sizes given are in 512 byte sectors.

The reason you don't see all the available space in your file systems is that 4.1c/4.2 requires that file systems not be filled more than 90% in order to get good efficiency from the new file system organization. So, when you do a "df" you'll see that a file system shows "100%" full when it's using only 90% of the available blocks. When it's really full, df shows "111%". It is advisable to stay within the nominal 100% whenever possible.

See Also:

nd(4) manual page
df(1) manual page

1.8) getty baud rate table for 0.4

Applies to: 0.3 and 0.4 only.
Description:

A Sun system can be made multi-user by attaching terminals, editing the file /etc/ttys, and inserting a "1" at the beginning of the line for each serial port on which logins are to be enabled. The super-user then types "kill -1 1" to make init read the ttys file. This is documented in the System Interface Manual under "ttys(5)". The second character on the line (following the "1") determines the baud rate for the corresponding line. However, the manual omits the table of supported baud rates. The omitted table follows:

Char	Baud rate
'0'-1-2-3	300,1200,150,110
'-'	Console tty 110
'1'	150
'2'	9600
'3'	1200,300
'4'	Console Decwriter
'5'	300,1200
'6'	2400
'7'	4800
'8'	9600,300
'9'	300,9600
'i'	Interdata Console
'l'	LSI Chess Terminal
'p'-'q'-'r'	9600-300-1200

Note that in 1.0 and later releases, the above scheme is superseded by a serial line description file named /etc/gettycap, documented in the System Interface Manual under "gettycap(5)".

See Also:

System Manager's Manual
ttys(5) manual page
ttytype(5) manual page

getty(8) manual page
init(8) manual page

1.9) Useful escape sequences for full-screen terminal emulation

Description:

In the full-screen terminal emulation your cursor sometimes disappears. It can be restored by echoing the escape sequence 'ESC[s' to the screen, where 'ESC' represents the ESCAPE character. NOTE: this does **not** apply to SunWindows. In a UNIX shell (or shell script) this is accomplished by typing

```
echo 'ESC[s'
```

where, again, ESC means you should type the ESC key. Some users put a line such as

```
alias fixcursor echo "'ESC[s'"
```

in their .cshrc file. Note that UNIX echoes the escape character as '^['.

Note also that there are other escape sequences to alter scrolling mode, change the display to white-on-black, etc.

See Also:

(Release 0.4 and earlier): *Programmer's Reference Manual for the Sun Workstation*
Section 5.1, "ANSI Terminal Emulation"
(Release 1.0 and later): *System Manager's Manual*
"Sun Workstation Monitor" Appendix, "ANSI Terminal Emulation" section

1.10) Disable daemons for better performance in 0.4

Applies to: 0.3 and 0.4 only
Description:

Users have noted very poor performance on workstations with 1 Megabyte of memory running UNIX Release 0.3 or 0.4, especially on diskless workstations. This is caused by excessive paging activity ("thrashing") due to an insufficient amount of available main memory. Besides the obvious remedy of adding a second megabyte of memory (which is quite effective), performance can usually be greatly improved by killing off unnecessary daemons. Daemons are processes that run constantly in the background to provide some service to other processes. These daemons are started up by the shell scripts /etc/rc and /etc/rc.local during booting, when the system goes from single-user to multi-user mode. For historical reasons, these daemons are started up by default in 0.3 and 0.4 even though they are often not used.

The daemons which have the most adverse effect on performance are the rwho daemon (/etc/rwhod) and the route daemon (/etc/routed). The rwho daemon is used only by the "rwho" and "ruptime" commands, which are not very important and (in 0.3 and 0.4) appear not to work correctly anyhow. The route daemon is ONLY needed on machines in networks containing one or more gateways. [A gateway is a Sun machine which interconnects between two physically distinct Ethernet networks]. The offending daemon is disabled by changing its name, for example:


```
cd /etc
mv rwhod rwhod.disabled
mv routed routed.disabled
```

The daemon does not die immediately, but it will not start up the next time the system is rebooted. You can kill it immediately (as root) by typing "ps ax" to find the daemon's process id (PID), then typing "kill *xxx*" where *xxx* is the process id number. Other daemons which may not be needed include /etc/ftpd, /etc/tftpd, and /etc/telnetd (ARPA-compatible file transfer and telnet programs); /etc/syslog (the network mailer's message logging program -- not much used); /etc/rexecd (remote procedure call -- doesn't even work properly!), and /usr/lib/sendmail. However, make sure a program is not used before disabling it.

In 1.0 and later releases, the above discussion no longer applies. One meg machines now perform reasonably, and the System Manager's Manual contains some updated suggestions for performance enhancement.

Besides disabling daemons, another way of freeing up memory is to reconfigure the kernel, making it smaller by removing unneeded device drivers. However, this is not well documented prior to Release 1.0 and should be undertaken on these releases only by users who know what they are doing -- and even they should think twice about it.

1.11) "Where are you?" message from biff command.

Description:

The command "biff y" tells the system to inform you immediately whenever you receive mail. However, if you put this command in your .cshrc file, then whenever you do an rsh or rcp from another machine to the machine with the biff, you will get the message "Where are you?". To avoid this, move the biff to .login instead.

1.12) Problems involving .cshrc

Applies to: 0.3 and 0.4 only.
Description:

There is a bug in csh which may cause it to hang when there is a comment (line beginning with #) in the .cshrc file. This tends to appear when doing an rsh, or when invoking a csh from inside of vi. Work-around: remove all comments from .cshrc files. Fixed in 1.0.

Applies to: 0.3, 0.4, 1.0
Description:

Another bug relating to .cshrc, involving the rcp command: rcp will not work if your .cshrc on the remote machine does ANY output to standard output (i.e., to your screen). For example, if you have the line

```
echo Hello, Fred.
```

in your .cshrc file on machine sun-a; then you login to sun-b and type an rcp command, for example:

```
rcp sun-a:/etc/passwd /tmp
```

you get the output "Hello, Fred.", but the file doesn't get copied. This appears to happen whenever .cshrc (or a program that it calls) outputs anything to your screen. Work-around: either don't put output-producing commands (like "biff"; see article 1.11 above) in .cshrc; or redirect their output to /dev/null. This bug has been reported, however it may not be fixed in 1.1.

1.13) Fixing bad disk blocks due to loose cables

Applies to: All systems with disks.
Description:

The console message "READ FAILED" or "WRITE FAILED" means that there is a bad block (sector) on the disk, which the disk controller is unable to read or write. The reason given in the message may be "ECC ERROR", "DATA OVERRUN", "SECTOR NOT FOUND", or a variety of other causes. These errors should *always* be taken seriously, since they indicate a disk hardware malfunction which may lead to total destruction of the user's file systems if not corrected promptly.

Sometimes upon checking the disk cables, the customer finds that a connector has been loosened or is partially unplugged. It is important to note that *even after the faulty cable is corrected*, the bad disk sectors may still be unreadable or unwriteable because the formatting information for the affected sector(s) may have been damaged. The hardware problem has been fixed, but a software problem remains on the disk. The only way to fix these sectors then is to use the "fix" command in diag to reformat the offending sectors. (With SCSI disks, the "fix" command is unavailable, but often the "write" command can be used to repair individual sector(s) in similar fashion).

After hardware problems have been corrected, you can attempt to salvage the data on the disk. The particular blocks at fault may be found by diag's "read" command, however the block number given by diag will not be accurate unless the "increment" and "number of sectors per transfer" (input to the "read" command) are both equal to 1. A quick way to find all the bad sectors is to do a "read" starting at sector 0 with number of sectors equal to the total size of the disk, increment and sectors/transfer both equal to (say) 32. Each time the read fails, write down the reported cylinder/head/sector address (for example, 38/3/12). When the read is completed, go back and run "read" again, starting somewhat before each reported bad sector (in the example above, 38/0/0), with number of sectors sufficient to include the bad sector (say, 1000), and with increment and sectors per track both equal to 1. This time the reported cylinder/head/sector addresses will be accurate, so you can run "fix" on the individual sectors.

Important: if the "fix" or "format" command responds with an error message ("FORMAT FAILED" or something of that nature) then you still have a hardware problem which must be corrected before proceeding.

Note that "fix" (and "write") destroy the data in the affected sector. After fixing all damaged blocks and certifying that the entire disk is readable, you can attempt to boot UNIX and run fsck on any file system that may have been damaged. If UNIX cannot be booted, you can attempt to restore the root file system from tape by following the installation procedure in the System Manager's Manual, beginning with loading the Standalone Copy program and ending with the running of xtr. *Do not* run setup, xtrusr, or xtrusrnd, since these will destroy the user or public file system. Instead, after running xtr, you can attempt to boot UNIX single-user and run fsck on the /usr (or /pub) partition and ndl partitions, if applicable. At the point you run fsck, it may remove some files or directories

whose blocks were overwritten during diag. You should write down any file names mentioned by fsck, and attempt to restore these files from tape, over the network if available, or from the "lost+found" directory if they were saved there by fsck.

If the damage to the file system was too extensive, or if you don't understand these instructions, there is always the recourse of re-installing the entire system from scratch. Follow the instructions in the *System Manager's Manual.*

See Also:

diag(8) manual page
fsck(8) manual page
System Manager's Manual for the Sun Workstation

1.14) Bad UNIX boot tapes causing "-:11 BUS ERROR"

Applies to: Some 0.4 systems.
Description:

Apparently a small number of bad 0.4 boot tapes were distributed with new systems around December of 1983. Upon trying to boot UNIX, the user gets the message "-:11 BUS ERROR" and the terminal mode of the console is found to be in a strange state. Then attempting to run programs such as **ed** or **stty** cause a crash and core dump. The solution is to obtain a good boot tape. Refer this problem to Field Support.

1.15) No ports C and D on 150's

Applies to: Model 150 and 150U.
Description:

Recently-shipped Model 150s have back panel connectors for four serial ports (labelled A, B, C, and D). A number of customers have called to ask us how to enable ports C and D. The sad answer is that there are no ports C and D. These connectors do not connect to anything.

1.16) Spurious ar error message

Applies to: 0.4 only.
Description:

Due to some slight system integration problems between Sun's quarter-inch tape controller and the Archive quarter-inch tape drive, a spurious error message may be generated upon the completion of a rewind command or read/write instruction which ends in a rewind command (because it specifies device /dev/rar*x* vs. /dev/nrar*x*). UNIX Release 0.4 deals properly with the problem but generates the following error message:

ar: interrupt without Rdy or Exc, state RdSt=a, Old Idle=1, bits=61080
<EnaExcep, Catch Ready, Online>

Since this message does not indicate an error in the functionality of the tape transfer, it should be ignored.

1.17) Miscellaneous information about "tar"

Description:

The **tar** command has a key letter **p** (useful only in conjunction with the **x** key letter) which was not documented in the 0.3/0.4 manual page. If "tar xp" is used, the named files are restored to their original ownership, creation date, and modes, ignoring the present user, date, and umask. This is useful when restoring backed-up files from tape, since you probably want to restore the file exactly as it was when saved. The **p** option also must be used when special permissions are involved, such as "set user ID on execution".

See Also:

chmod(1) manual page
tar(1) manual page
umask(2) manual page

Description:

When **tar**ing to tape, the **r** and **u** key letters do not work because the tape drives are unable to backspace. If you need to add additional tar(s) to a tape already containing one or more tar files, use the "no rewind" tape device (/dev/nrar0 or /dev/nrmt0 as opposed to /dev/rar0 or /dev/rmt0) in conjunction with the "mt" command to skip over existing files and write new files without rewinding. This is not quite the same thing as tar r or tar u, since a distinct physical file is written onto the tape, which must be positioned to using **mt** and read with a separate **tar** command. When writing more than one tar file onto the same tape, extreme care must be exercised since you may forget to specify the no-rewind device to tar, inadvertently rewind the tape, and write over the first file already on the tape. You may also forget that there is more than one file on the tape (tar can't tell you) so be sure to clearly label the tape reel.

See Also:

tar(1) manual page

1.18) When color board is added, must create /dev/cg0

Applies to: Systems with color upgrades.

Description:

In order to use the color graphics board, you must have a device /dev/cg0. This may be easily created by the super-user by typing:

```
cd /dev
MAKEDEV cg0
```

1.19) Tape exchange guidelines.

Applies to: Systems with half-inch (9-track) tape drives.

Description:

Our half-inch tape drive is advertised as "industry standard", which it is. However, it is much easier for a Sun to read a tape written on another vendor's machine if certain guidelines are observed.

Obviously, the tape must be 9-track half-inch tape. It must be 1600bpi (bits per inch). If the other system is a UNIX system, **tar** should be used to create the tape. Otherwise, the tape should be an unlabeled ANSI standard tape with a block size that is a multiple of 512 bytes (10240 works well). If possible, variable-length records should be used (although it's not too hard to strip trailing blanks off of fixed length records). If a large number of files are to be written on the tape (say, several dozen or more) then it is recommended that the first file be simply a list of the file names of the remaining files, one per line.

To read a tar tape, just use the **tar x** command. For an ANSI tape, use **dd**. dd has many options for converting EBCDIC to ASCII, unblocking fixed-length records, and so on. However it's easier to follow the guidelines above whenever possible. The list of file names recommended above can be used to drive a simple shell script which will automatically read all of the remaining files and assign the correct names to them. Writing this script is left as an exercise for the reader. (Call Tech Support if you need help).

One final note: tapes sent to the Technical Support group from the field offices *must* be in **tar format**; no other format will be accepted. It is the field office's (or customer's) responsibility to do any necessary conversion.

1.20) No 4014 emulation in Rev. L PROMs.

Applies to: Early Sun-2 upgrades (Models 100U and 150U only).
Description:

Many of the earliest Sun-2 upgrades shipped at the end of 1983 and early in 1984 have the Revision L monitor PROMs. These PROMs do *not* contain the 4014 terminal emulation which was present in earlier (Sun 1.5) PROM versions. The 4014 emulation is expected to reappear in forthcoming PROM releases. Customers who require the 4014 emulation will be able to exchange their PROMs for new ones when available. Details are yet to be worked out.