==== ==== ==== ==== ==== ==== ====

Note 2 C:/bin/.lpr makes unreadable links 1 response
iv 3:54 pm Nov 18, 1987

Description: There is a bug in /bin/.lpr (as we call it on the 6000) which has been around forever (at least since v7 days).

If you lpr a file on the same filesystem as /usr/spool/lpd, 'lpr' will make a link to it in /usr/spool/lpd. This is fine unless the file isn't world readable, in which case lpd can't read it (it setuid's itself to "daemon"). In this case, the print job just seems to disappear (you get a banner if BANNERS > 0, but nothing else).

Amazed-Look: I'm astounded that no one has ever reported this. If you attempt to lpr a file that isn't world-readable, it looks like the job is ignored, especially if BANNERS=0.

Solutions: "lpr -c file" will work, since the copy of file will end up readable by lpd. Various forms of "lpr < file" will work also.

The Fix: Patch /bin/.lpr as follows:

```
# patch /bin/.lpr
byte offset (<ENTER> to exit)? 29da
29da:  6c |l| > 63
29db:  00 |.| > q
```

With this patch, lpr will no longer link files. They will always be copied. While this is not the best of solutions, the "proper" fix would require source code additions. At least this way, the users will not lpr a file and then complain that XENIX is broken.

==== ==== ====
Response 1 to Note 2
gordon 12:13 pm May 13, 1988

If anyone cares, I have a fix for this.

==== ==== ==== ==== ==== ==== ====
Note 3 C:tsh: broken again (new for 3.2)
iv 11:39 am Dec 15, 1987

If you say something like "restore :0 -ds /foofile", tsh will obligingly restore the ENTIRE archive, not just "/foofile". Oops.

tsh 1.4(11), on XENIX 3.1.0, did not do this. It worked correctly.

You can get 1.4(12) to work correctly by adding an extra dummy argument before the file you wish to restore, thusly:

```
restore :0 -ds - /foofile
```

The extra "-" will make it work.

==== ==== ==== ==== ==== ==== ====
Note 4 C:SPR 6460-004 TTY lockups in 3.2 3 responses
uhclem 11:14 am Dec 23, 1987

First complaint for 3.2 from the field, soon to be arriving in the form of a SPR:

>I just wanted to drop you a note to brighten your day by letting you know that
>we're having as many getty lockups under 3.2 as we were getting under 3.1 ...
>
>We've gotten some new hardware, a 62Meg 23ms Priam drive that *really* makes
>things fly (like, it's a joy to use even with 7 people and two uucps going --
>LET'S SEE THE 4000 ... er, I digress) ... anyhow, the new hard drive made the
>biggest difference in lockups: faster-drive == less-lockups. I've generalized
>the problem down to "it's getting caught with it's pants down -- it's catching
>a second interrupt in the driver before it has time to deal with the first
>one." It only does this when the system gets quite busy.
>
>The good news is that the autoboot sort-of solves our problem. At least we can
>re-boot from remote and I'm writing a script that will catch lockups and do a
>re-boot.

In a follow-up letter, I received some additional information.

There are three different systems around the country that this guy takes care of that are seeing this problem. San Antonio, San Jose CA, and Houston. The machines are two Tandy 6000 and one Model II with an extra slot wire-wrapped on. Yes, the problems also happen on the normal systems.

The lockups are with user calls, not uucp. They have not seen a uucp as the last activity on a line. The lines are always draining. Wchan for 3.2 was 1320c getty. Kill -1 to 16 does not get rid of the job. User A has a normal session and logs out. Line gets hung. User B calls in and gets a dead line. So does C, D, E, ...

They are using Hayes Smartmodems at 1200 baud. They are running with two multiterminal boards, but they may not be as close to the CPU as they can be. I have not gotten confirmation on this. They can get additional processes to hang on the line if they start messing it, but thats normal.

They have just started increasing their buffer count to reduce Z80 activity. The systems average 7 users logged in at once except in the early morning.

They average 500 dialups on the multiterminal boards each day, and see one lockup an hour. (Because of the constant rebooting, the dialup count could be higher.) They moved one dialup to the SIO and decreased the lockup count to 1/20 the original rate, but this was only a single day sample. The tty in question was tty08. I have advised them that this sounds like the PCI is too far from the CPU and to recheck the configuration. No answer on this.

I have been able to reproduce the lockup, but it took quite a while. The Z80 was still handling some operation on the line but never acknowledged its completion. I have added some code to cover something mentioned in an Intel application note (not the spec sheet) and we have switched Hal6000 to use multiterminal lines for all the modem I/O. If it locks up, we have debug code standing by to give us another glimpse at the problem.

Iv, how about posting a summary of what the 68000 thinks is being done when that we have that wait channel.

This is somewhat embarassing as this area is one of the things that they thought 3.2 was going to fix. Merch may have thought this too. In fact, we did have problem reports about lockups, but they always involved uucp and we did all sorts of things to reduce or eliminate those. On checking with customer service, they assumed that this type was just a variation of what they had already reported, so they did not turn in any new reports.

Hal has had the single modem on a multiterminal board get lost from time to time, but it was like once a month, not once an hour.

More as this develops.

==== ==== ====
Response 1 to Note 4
iv 2:05 pm Dec 23, 1987

The 68000 is waiting for a transmit-finished interrupt (checks by looking for the TXINT bit). It will wait all eternity if necessary.

The bad thing about that condition is that we've coded it NOT to bother the chip with any status change type stuff until the channel is idle. So, it'll stay stuck.

The wait channel you always see is a "waiting-for-output-to-drain" state. It's what the kernel does just before completing a close.

==== ==== ====
Response 2   to Note 4
        uhclem   2:00 pm Jan 28, 1988

More information. By running a debug z80ctl on Hal and a few other systems, we have been able to catch the sucker in a hung-up state. Of course, by the time we notice it, usually another user/computer has attempted to call in on the line and has wiped out the status of the actual moment of lock-up.

There appear to be two causes, (1) some problem with the driver, and (2) a race condition in uucico which should be discussed in a different base note. The remainder of this note discusses condition 1.

When we look at the wreckage, the state indicates the 68k was trying to transmit. It is not always the closing. Here are some snapshots we have accumulated (these are using a stock 3.2 kernel)

Line=7 Chip=85 Mode=00 Count A=01 Wctl=07 Baud=39 Iopb_csr=05 Count B=02
Pend=00 Mask=02 Chars=0a (Wchan=13252)

Line=7 Chip=85 Mode=00 Count A=02 Wctl=07 Baud=37 Iopb_csr=05 Count B=02
Pend=00 Mask=02 Chars=0d =0a (Wchan=13252)

Line=7 Chip=B7 Mode=00 Count A=02 Wctl=07 Baud=39 Iopb_csr=05 Count B=02
Pend=00 Mask=02 Chars=0d =0a (Wchan=13252)

Line=4 Chip=85 Mode=00 Count A=01 Wctl=07 Baud=39 Iopb_csr=05 Count B=02
Pend=00 Mask=02 Chars=39 (Wchan=13180)

Line=5 Chip=85 Mode=00 Count A=02 Wctl=07 Baud=37 Iopb_csr=05 Count B=03
Pend=00 Mask=02 Chars=20 =08 (Wchan=131c6)

(IV, please research the Wait channels.)

We have experimented with three changes to the driver to see if it improves the situation. One prevents the status of the chip (baud rate, word size) from being changed until the last character has been completely transmitted. Previously the 68k would not issue a request to change status until the last character had been handed to the hardware, but at low baud rates, it is possible for a status change to overtake a character being transmitted.

Another change was to delay between each command (including parity error reset) and the next transmission of data and vice versa. The book

implies in one place that both commands and data pass through a common buffer, but elsewhere it implies that it means an electrical buffer and that no delay is needed. Tests with the emulator have shown that the chip does not care how close a data/command come to one another. Needless to say, the Intel documentation on this chip is really bad. The code for this has been removed since it resulted in no improvement.

One final change has been put into test which causes the status of the chip to be read twice after all interrupts appear to have been serviced. The problem is that the pin on the chip to generate an interrupt for a transmit-empty or received-character-available goes high up to 28 clocks before the status in the port is updated. Tests with the emulator with probes on the signals and repeated IN opcodes shows that as many as 1.5 IN instructions can be executed between the time the interrupt condition is valid and the time we can find out about it. This results in a race condition since we have to resolve all interrupt sources from the chip before exiting the interrupt service routine or we won't get anymore without resetting the chip. If the two reads of the status show differences, then more service is needed.

More as we gain knowledge.

==== ==== ====

Response 3 to Note 4

uhclem 5:09 pm Jul 25, 1988

A problem report (SPR) on this subject has arrived. SPR 26-6013. This arrived on 21-Jul-88. On 25-Jul-88, I was informed that the SPR number has been changed to SPR 26-6460-004.

==== ==== ==== ==== ==== ==== ====
Note 5 C:Race condition in floppy driver
uhclem 1:36 pm Dec 28, 1987

There is a race condition in the floppy driver that causes XENIX 3.2
to hang and all terminals to go dead. The only recovery is a reboot.

The problem occurs when a buffer is being copied to the 68000 (in a read)
while another is copied from the FDC to the Z80 memory. If the
FDC transfer completes and its interrupt is serviced before all of the
previous buffer is copied, AND the last sector of the transfer was
just read, the system can hang waiting for the copy routine to finish,
which will not happen until the path that is waiting releases control.

It is hard to get this to happen with 512 byte sectors, however I suspect
it would be easier to see if the 68000 was only running at 6MHz. 128 byte
sectors also make the problem easier to see. Adding printfs (z80 equivalent)
of debug information in the driver also causes the problem to show itself.

In Edit 125 and later, code corrects this problem.

==== ==== ==== ==== ==== ==== ====
Note 6 C:Techsup says field getting HDTime 2 responses
uhclem 9:36 am Jan 22, 1988

Larri in Tech support reports that some shop claims that systems that
were upgraded from 3.1.2 to 3.2.0 are experiencing HDTimeOut errors.
The hard disks in question do not fail when hooked up to different CPU's.
They say they have three of these systems. I has asked for additional
information, specifically, the diagnostic information that is displayed
with the Timeout error.

I put this in here so we can keep track of when these were mentioned/reported.

==== ==== ====
Response 1 to Note 6
meek 11:56 am Jul 26, 1988

Anyone ever heard anything from this report? Been six months now....

Thanks..Don

==== ==== ====
Response 2 to Note 6
uhclem 1:52 pm Jul 26, 1988

Well, Larri has left Tandy, but I'll ask around...

==== ==== ==== ==== ==== ==== ====
Note 7 C:shutdown goes up to MU in SU mode 4 responses
uhclem 10:40 am Feb 26, 1988

A problem from a customer:

If you run /etc/shutdown and are in single-user mode, strange things happen.
If you use -h or -r, the screen clears and you get:

```
Tandy XENIX 3.2 (foobar)
login:  * Normal System Shutdown *
[Z80 Control System Halted]
```

If you use the -s option, it goes into multi-user mode, sort of. Not all the deamons are started and the system remains "strange" until you haltsys.

We need to see if shutdown can handle being run in single-user mode a little better since there is no good mechanism for a user's scripts to use to tell if running shutdown is a bad idea. In any event, it should not try to go multi-user.

==== ==== ====
Response 1 to Note 7
gordon 5:58 pm Feb 26, 1988

Assuming I had a fix for this, which was not patchable, what should I do with it?

==== ==== ====
Response 2 to Note 7
iv 10:57 am Feb 29, 1988

It might be noted that this behaviour has been around since before the beginning. V7 did it. The old sh script on XENIX did it as well.

==== ==== ====
Response 3 to Note 7
gordon 5:25 pm Mar 14, 1988

I do have a moderately tested fix for this. What am I supposed to do with it?

==== ==== ====
Response 4 to Note 7
uhclem 4:35 pm May 9, 1988

Hang onto that fix. I see a 3.2.1 far, far in the distance.

==== ==== ==== ==== ==== ==== ====
Note 8 C:shutdown barfs on v7 mount output 7 responses
uhclem 10:46 am Feb 26, 1988

Another problem from the field:

/etc/shutdown does not work properly on systems that use PROFILE.
It cannot handle the modified output produced by /etc/mount when
the environment variable SYS3INSTALL (is that right?) is present.
(That variable makes mount output in the old V7 format, which does
not display the "/dev/" on the front of mounted devices.)

/etc/shutdown needs to either eliminate the variable before invoking
mount or edit the fields if the variable is present.

One other question: Is there a newer version of Profile that is native
to System III which would eliminate the need for this variable? If
so, and it does not leave the variable around after upgrading, then
this should be the procedure taken to correct this problem, unless
there are other packages that expect the same kludge.

==== ==== ====
Response 1 to Note 8
iv 12:14 pm Feb 26, 1988

Why is this environment variable set when you shutdown??

It should ONLY be set when you are installing. What ever the
applications do internally to the environment is their business,
but they should *NOT* require modification of the default behaviour.

==== ==== ====
Response 2 to Note 8
gordon 1:40 pm Feb 26, 1988

Unfortunately, Profile instructs users to have the variable set
when they RUN profile (because, apparently, profile has the inverse
problem of what shutdown does - it runs df, mount, or something, and
expects the *OLD* format). So it probably ends up in everyone's
.profile file. (It's named .profile, so obviously it's controlled by
Profile, right?)

I don't understand why they couldn't have put this in the profile front-end
shell script, (as I recall, it *IS* a shell script) and take it out of
.profile/.login, but they didn't.

Now, which is harder to do: release a core upgrade disk (3.3?), or
modify the instructions (documentation) for Profile and get users to
edit some scripts?

==== ==== ====
Response 3 to Note 8
iv 2:17 pm Feb 26, 1988

Shoot, if we're going to have users fixup the profile botches, why not have them do:

cp /bin/mkdir ~profile/bin/makedir

and get rid of a security hole in the process.

==== ==== ====
Response 4 to Note 8
uhclem 3:46 pm Feb 26, 1988

R1>Why is this environment variable set when you shutdown??

Gordon was quite correct in his response. Profile runs mount and examines the output to figure out what is mounted where. It has a lot of paths hardcoded. Remember when you tried to move some of Jeri's Profile stuff to a secondary drive and it went ape?

We should be able to either have shutdown prevent the environment variable from reaching /etc/mount, have shutdown read the mount table itself or handle both formats.

==== ==== ====
Response 5 to Note 8
uhclem 4:37 pm May 9, 1988

Please look into "hiding" the environment variable from the mount/umount program when invoked by shutdown or whatever is a clean method. No hurry, just fill in a crack of time with this.

==== ==== ====
Response 6 to Note 8
iv 4:55 pm May 9, 1988

>> Please look into "hiding" the environment variable from mount/umount
No problem, we should just be able to use unsetenv() to get rid of it.

==== ==== ====
Response 7 to Note 8
gordon 10:32 am May 12, 1988

I have a better fix. Construct an environment, consisting only of PATH (shutdown is going to have problems if root doesn't have /etc in $PATH, also, and this fixes that problem at the same time), and use it instead of the passed one. Nothing in the user environment can interfere. We do have to test that there is no problem with normally-present environment variables being missing doesn't bother mount, qps, etc., but there is only one combination to test.

This change has been made and tested briefly on wombat (which does have a secondary drive to unmount). It does fix the nasty messages that come out under the same circumstances when you use the old one.

Now it needs to wait around for a core upgrade release ...

==== ==== ==== ==== ==== ==== ====
Note 12 C:Field wants DTR dropped on crash 2 responses
uhclem 9:58 am May 2, 1988

Here is one I received over the net. It *is* different.

"Sigh. Here's another thing that "broke" with 3.2: When the 3.2 system crashes (and I've had a lot of that this week :-( ) the z80 subsystem does *not* drop DTR to any of the /dev/tty* devices. Since I have my modems configured to follow DTR (and not force it on), the modems will still answer even with the system down. Bad news if you are calling long distance: the modem answers (call is completed), but you won't get a login: prompt which the console saying [Z80 Control System Halted].

To the best of my knowledge, 3.1.2 and earlier versions of Xenix would drop DTR.

- paul" of devon!paul.

Gee, I didn't know my 3.1 drivers would restart the system long enough to reset 12 serial devices after a crash. Must have put that code in ROM or something. :-(Sorry, nothing different on this point on the Z80 side for 3.2 over 3.1.

On the 68000 side, does the kernel issue closes to the Z80 on a panic? Since commands are handled in the order received, a shutdown signal from the 68K would not "beat" any other command to the Z80.

==== ==== ====
Response 1 to Note 12
iv 11:57 am May 2, 1988

>> On the 68000 side, does the kernel issue closes to the Z80 on a panic?
>> Since commands are handled in the order received, a shutdown signal
>> from the 68K would not "beat" any other command to the Z80.

Nope. He just attempts to sync() out the superblocks, then issues a halt.

==== ==== ====
Response 2 to Note 12
gordon 12:24 pm May 2, 1988

Can this guy's problem be fixed by saying "turn on autoboot and reboot"? Does the boot rom reset everything (at least Serial A&B)?

If, for some reason, the system can't get to multi-user, but did get to the boot rom/boot track DTR will be left off (at least A&B). If it can come up multi-user, then it should answer the phone, because the system is

back up again. If it can't get to the boot rom then the system is in
real trouble anyway, and who knows what will happen.

Kernel 3.1(13) does NOT reset DTR on a panic (in this case, panic: blkdev).

==== ==== ==== ==== ==== ==== ====
Note 15 C:SPR 6460-003, DTR stays high
uhclem 2:09 pm Jun 29, 1988

I have received SPR 6460-003 on 29-Jun-88. It was written 22-Jun-88.
Very impressive! Here is the problem:

"XENIX keeps DTR high on its serial lines after a crash (panic or Bughlt),
or after a shutdn(S) call. However, when a shutdown is performed in an
orderly fashion using shutdown(C), DTR goes low. On systems with modems,
leaving DTR high is a problem.

If DTR remains high in a crash situation, any modems on the system may
continue to answer the phone (resulting in possible long distance connect
charges) even though XENIX is not running.

The system should make some attempt to drop DTR when a crash occurs."

The customer is Barry Capell @ General Electric.

==== ==== ==== ==== ==== ==== ====
Note 16 C:SPR 6460-005 Slow Printing in 3.2
uhclem 5:16 pm Jul 25, 1988

SPR 700-3039-004, received 21-Jul-88.
This problem report was renumbered by Sher on 25-Jul-88. It is now
SPR 26-6460-005.

The problem is that when XENIX 3.2 prints to a Daisy Wheel II or
a IDS PaperTiger (alien), the printing speed is very slow.

This was also observed on a DMP400 at a location in Fort Worth. The
site reported a job that too 6 hours to print with 3.1.x started
taking 2 days when 3.2.0 was installed.
A hardware problem (GND pin on buffer chip folded under) caused the
slow printing by not passing ACK signals back to the PIO.

The problem appears to revolve around the polling speed used when
interrupts are not present.

If this is the case, a Printer Selector I will make *any* printer
that is attached print slow, as this device does not pass ACK,
at least that is what its documentation says.

Using a Printer Selector II or a Printer Buffer creates its own ACK
signal which the PIO should act upon.

==== ==== ==== ==== ==== ==== ====
Note 17 C:Wasted user memory 1 response
gordon 6:11 pm Jul 25, 1988

The kernel wastes 4k of user memory by initially neglecting to free it.

This slows performance slightly and is also the cause of the problem reported in note 265 in the archive, "Swap Deadlock".

To reproduce: poke around the system with adb dumping _coremap, and use ps to locate processes. Observe that the 4k following the swapper's U page is allocated, but it is not allocated to anything in particular.

==== ==== ====
Response 1 to Note 17
iv 4:38 pm Jul 29, 1988

This is fixed in the next (ha ha hahahaha) release [3.3].

==== ==== ==== ==== ==== ==== ====
Note 18 C:Swapping large processes
gordon 6:13 pm Jul 25, 1988

Xenix's swapping strategy doesn't work well for large processes that grow slowly (that is, make many small requests for more memory).

To reproduce, run the following program, with maxmem set to the largest permissible (the user memory printed by the kernel less 4k):

```
main()
{
        char    *cp;
        extern char *sbrk();

        while ( (cp = sbrk(4096)) >= 0)
                printf("sbrk returned %lx\n", cp);
        return 0;
}
```

When you run this program, observe that the messages are printed quickly with an occasional pause until the program size reaches about half of maxmem. From then on, the program goes very slowly, swapping in and out between each message.

The swap strategy works like this: if a process needs to grow, try to allocate another chunk of memory of the new size (without releasing the old one), and copy the process into it. If the allocation fails, swap. Once the process grows to over half of user memory, the allocation will always fail, and the process will always swap, even if there is already memory immediately following the process available.

There are quite a few processes that can grow quite large with small requests, so the process will grow slowly and swap frequently:

        C compiler/optimizer

make
pathalias (we don't supply this one)

and it wouldn't be too surprising if these were affected:

Profile
Unify

==== ==== ==== ==== ==== ==== ====
Note 19 C:/etc/init mis-handles etc/inittab 1 response
uhclem 12:37 pm Jul 26, 1988

/etc/init is not clearing its internal buffers between lines when
reading /etc/inittab. For example, if /etc/inittab looks like this:

```
0:01::/etc/mygetty foobar
0:02::/etc/mygetty
```

When the processes startup, you get:

```
$ ps -ft 01 -t 02
    UID   PID  PPID  C    STIME TTY  TIME COMMAND
   root  8980     1  0 12:06:52 01  0:00 -3 foobar
   root  8981     1  0 12:06:52 02  0:00 -4 foobar
$
```

Only one of the processes should have the optional argument.

This problem was originally described by paul@devon.

==== ==== ====
Response 1 to Note 19
gordon 1:48 pm Jul 26, 1988

This problem fixed on grumpy's source tree. What should I do with
the fixed version?

==== ==== ==== ==== ==== ==== ====
Note 20 C:diskutil wont work w/ repl. drive
uhclem 1:43 pm Jul 26, 1988

I have received word that Tandy has stopped buying Micropolis 70 Meg
hard disks and is now buying Rodimes. The Rodime is about $200 cheaper
and is a half-height unit, so I can see the logic. It is possible that
someone with a 70 Meg primary or secondary drive attached to the 16/16B/6000
could receive a Rodime if the bubble ever croaks and it is replaced.

The problem is that the Micropolis is 1024 x 8. The Rodime is 1224 x 7.
(The Rodime is a tad larger, with 8568 tracks vs 8192.)
The 1010 controller used in most of our hard disk controllers only
allows 1024 cylinders and 8 heads. This would make about 12 meg unusable.

If the 1010 controller chip is replaced with a 2010 controller (the 1010
is no longer made, and the 2010 is being used as the replacement anyway),
the hardware will allow up to 2048 cylinders and 8 heads. (All of this does

not apply to 8X300 controllers, but will work on internal and TB1 controllers.)

XENIX 3.2.0 allows drives larger than 1024 cylinders, but diskutil does not.

A simple patch to diskutil can change the cylinder limit in the field, but there is no simple way to tell if a 2010 controller has been installed. If you attempt to access cylinders above 1023 on a 1010 controller, the controller wraps back to zero. Eventually your swap area (or other data) would be written to low tracks and do all sorts of damage.

The appropriate thing to do would be to detect a cylinder count above 1023 and below 2048, then warn the user that they had better have the 2010 controller installed. Cylinder counts below 1024 would behave as before. This will probably require a new release of diskutil.

==== ==== ==== ==== ==== ==== ====
Note 21 C:System wont let go of error cond.
uhclem 1:23 pm Jul 28, 1988

Managed to get the system into a state where it would not format a diskette because it thought it was write-protected and it wasn't. Trying different drives and different disks did not help. It aborted all formats with a write-protect error. Then I asked it to just read from a pre-formatted disk. It did this. Then I tried formatting and it worked. Weird.

Gordon reports seeing a possibly-related problem where the system would report that all floppy drives were not ready even when they had media in them. He rebooted to clear the problem up.

==== ==== ==== ==== ==== ==== ====
Note 22 C:xmail doesn't notice disk overflo 3 responses
iv 9:36 am Aug 5, 1988

Problem: xmail doesn't recognize when a disk drive becomes full
(Hey, all these writes are failing. Wonder why...)

This has been fixed, and can be included in the next release [3.3]
(ha ha ha ha hah hah ha ha)

==== ==== ====
Response 1 to Note 22
uhclem 12:06 pm Aug 5, 1988

The problem with the write failing is that you lose your new mail without getting a chance to read it. THAT is the problem.

==== ==== ====
Response 2   to Note 22
iv
2:38 pm Aug 5, 1988

>> The problem with the write failing is that you lose your new mail
>> without getting a chance to read it. THAT is the problem.

THAT is only *one* of the many problems that ignoring output errors caused.
I just decided not to bore you with the whole gory list.

==== ==== ====
Response 3   to Note 22
uhclem
6:14 pm Aug 5, 1988

R1> The problem with the write failing is that you lose your new mail
R1> without getting a chance to read it. THAT is the problem.

R2>THAT is only *one* of the many problems that ignoring output errors caused.
R2>I just decided not to bore you with the whole gory list.

True, but losing your mail sounds worse than ignoring write errors.

==== ==== ==== ==== ==== ==== ====
Note 23   C:shutdown umounts out of order
iv
9:56 am Aug 5, 1988

Shutdown will try to umount things out of order if /etc/mnttab has been
set up with 'setmnt'. Example:

```
# mount
/dev/root on / read/write on Thu Aug  4 17:46:18 1988
/dev/hd1 on /usr read/write on Thu Aug  4 17:46:18 1988
/dev/hd2 on /usr/spool read/write on Thu Aug  4 17:46:18 1988
```

This will cause a problem when you try to umount '/usr' before you
umount '/usr/spool'.

Normally, if you had used 'mount' and not reset the order, they would
have appeared like this:

```
# mount
/dev/root on / read/write on Thu Aug  4 17:46:18 1988
/dev/hd2 on /usr/spool read/write on Thu Aug  4 17:46:18 1988
/dev/hd1 on /usr read/write on Thu Aug  4 17:46:18 1988
```

and all would be well.

The fix is to have shutdown reverse-sort the mount table and then
umount the world.

This has already been fixed and can be included any future releases.

==== ==== ==== ==== ==== ==== ====
Note 24 C:"dump f file &" infinite loops 3 responses
iv 10:04 am Aug 5, 1988

Remember that wonderful change to dump I did so that you could re-do a volume which had an I/O error on?? Well, it has come back and bit us.

If you dump to a file from cron or in background, something like this:

```
dump 4ukf /mnt/dump.root /dev/rroot
```

and the filesystem overflows or something like that, dump will infinite loop retrying the volume.

Fix: Don't ask to retry if stdout isn't a tty. I've already fixed this in hal's copy of dump, but haven't given this to gordon yet.

Should this be included in a future release(tm)??

==== ==== ====
Response 1 to Note 24
gordon 10:18 am Aug 5, 1988

> Fix: Don't ask to retry if stdout isn't a tty. I've already fixed this
> in hal's copy of dump, but haven't given this to gordon yet.

Well, what DO you do if stdout isn't a tty, and you need to switch volumes? Abort? Keep going and write over the first volume of the dump? (Note: the first volume of the dump is necessary to do anything useful with the other volumes)

There exist cute hacks with named pipes to communicate with background processes without tying up a terminal. The proposed "fix" would break this.

==== ==== ====
Response 2 to Note 24
iv 10:44 am Aug 5, 1988

>> There exist cute hacks with named pipes to communicate with background
>> processes without tying up a terminal. The proposed "fix" would break
>> this.

Oops. I lied. It aborts the dump if the output device is not a device. It doesn't look at stdin or stdout at all. Sorry 'bout that.

==== ==== ====
Response 3 to Note 24
gordon 10:54 am Aug 5, 1988

> Oops. I lied. It aborts the dump if the output device is not a device.

Ok. That one I like a lot better.

2	L	C:/bin/.lpr makes unreadable links	iv
3	M	C:tsh: broken again (new for 3.2)	iv
4	M	C:SPR 6460-004 TTY lockups in 3.2	uhclem
5	L	C:Race condition in floppy driver	uhclem
6	?	C:Techsup says field getting HDTime	uhclem
7	H	C:shutdown goes up to MU in SU mode	uhclem
8	H	C:shutdown barfs on v7 mount output	uhclem
12		C:Field wants DTR dropped on crash	uhclem SAME AS 15
15	M	C:SPR 6460-003, DTR stays high	uhclem
16	H	C:SPR 6460-005 Slow Printing in 3.2	uhclem
17	H	C:Wasted user memory	gordon
18	M	C:Swapping large processes	gordon
19	M	C:/etc/init mis-handles etc/inittab	uhclem
20	L	C:diskutil wont work w/ repl. drive	uhclem
21	?	C:System wont let go of error cond.	uhclem
22	H	C:xmail doesn't notice disk overflo	iv
23	L	C:shutdown umounts out of order	iv
24	L	C:"dump f file &" infinite loops	iv

5 HIGH
5 MEDIUM
5 LOW
? UNKNOWN