

RE: Ufs dead-locks on freebsd 6.2

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2007-05/msg00050.html>

- *From:* "Andrew Edwards" <aedwards@xxxxxxxxxxxx>
 - *Date:* Tue, 22 May 2007 13:35:09 -0400
-

It's been a couple of days with no response, how do I know if anyone is looking into this problem?

-----Original Message-----
From: owner-freebsd-fs@xxxxxxxxxxxx

[<mailto:owner-freebsd-fs@xxxxxxxxxxxx>]

On Behalf Of Andrew Edwards
Sent: Saturday, May 19, 2007 12:34 AM
To: freebsd-fs@xxxxxxxxxxxx; freebsd-performance@xxxxxxxxxxxx
Subject: RE: Ufs dead-locks on freebsd 6.2

Fsck didn't help but below is a list of processes that were stuck in disk. Also, one potential problem I've hit is I have mrtg scripts

that

get launched from cron every min. MRTG is supposed to have a locking mechanism to prevent the same script from running at the same time but

I

suspect since the filesystem was unaccessible the cron jobs just kept piling up and piling up until the system would eventually crash. I caught it when the load avg. was at 620 and killed all the cron's I could. That brought the load avg. down to under 1 however system is still taking up 30% of the processor time and the disks are basically idle. I can still do an ls -l on the root of all my mounted ufs and

nfs

filesystems but on one it's taking a considerable amount longer than

the

rest. This particular rsync that I was running is copying into the

/d2

fs.

The system is still running and I can make tpc connections and somethings I have running from inetd work but ssh stops responding

right

away and I can't logon via the console. So, I've captured a core dump of the system and rebooted so that I could use it again. Are there

any

suggestion as to what to do next? I'm debating installing an adaptec raid and rebuilding the system to see if I get the same problem, my worry is that it's the intel raid drivers that are causing this

problem

and I have 4 other systems with the same card.

```
PID TT STAT TIME COMMAND
2 ?? DL 0:04.86 [g_event]
3 ?? DL 2:05.90 [g_up]
4 ?? DL 1:07.95 [g_down]
5 ?? DL 0:00.00 [xpt_thrd]
6 ?? DL 0:00.00 [kqueue taskq]
7 ?? DL 0:00.00 [thread taskq]
8 ?? DL 0:06.96 [pagedaemon]
9 ?? DL 0:00.00 [vmdaemon]
15 ?? DL 0:22.28 [yarrow]
24 ?? DL 0:00.01 [usb0]
25 ?? DL 0:00.00 [usbtask]
27 ?? DL 0:00.01 [usb1]
29 ?? DL 0:00.01 [usb2]
36 ?? DL 1:28.73 [pagezero]
37 ?? DL 0:08.76 [bufdaemon]
38 ?? DL 0:00.54 [vnlru]
39 ?? DL 1:08.12 [syncer]
40 ?? DL 0:04.00 [softdepflush]
41 ?? DL 0:11.05 [schedcpu]
27182 ?? Ds 0:05.75 /usr/sbin/syslogd -l /var/run/log -l
/var/named/var/run/log -b 127.0.0.1 -a 10.128.0.0/10
27471 ?? Is 0:01.10 /usr/local/bin/postmaster -D
/usr/local/pgsql/data (postgres)
27594 ?? Is 0:00.04 /usr/libexec/ftpd -m -D -l -l
27602 ?? DL 0:00.28 [smbiod1]
96581 ?? D 0:00.00 cron: running job (cron)
96582 ?? D 0:00.00 cron: running job (cron)
```

RE: Ufs dead-locks on freebsd 6.2

96583 ?? D 0:00.00 cron: running job (cron)
96585 ?? D 0:00.00 cron: running job (cron)
96586 ?? D 0:00.00 cron: running job (cron)
96587 ?? D 0:00.00 cron: running job (cron)
96588 ?? D 0:00.00 cron: running job (cron)
96589 ?? D 0:00.00 cron: running job (cron)
96590 ?? D 0:00.00 cron: running job (cron)
96591 ?? D 0:00.00 cron: running job (cron)
96592 ?? D 0:00.00 cron: running job (cron)
96593 ?? D 0:00.00 cron: running job (cron)
96594 ?? D 0:00.00 cron: running job (cron)
96607 ?? D 0:00.00 cron: running job (cron)
96608 ?? D 0:00.00 cron: running job (cron)
96609 ?? D 0:00.00 cron: running job (cron)
96610 ?? D 0:00.00 cron: running job (cron)
96611 ?? D 0:00.00 cron: running job (cron)
96612 ?? D 0:00.00 cron: running job (cron)
96613 ?? D 0:00.00 cron: running job (cron)
96614 ?? D 0:00.00 cron: running job (cron)
96615 ?? D 0:00.00 cron: running job (cron)
96616 ?? D 0:00.00 cron: running job (cron)
96617 ?? D 0:00.00 cron: running job (cron)
96631 ?? D 0:00.00 cron: running job (cron)
96632 ?? D 0:00.00 cron: running job (cron)
96633 ?? D 0:00.00 cron: running job (cron)
96634 ?? D 0:00.00 cron: running job (cron)
96635 ?? D 0:00.00 cron: running job (cron)
96636 ?? D 0:00.00 cron: running job (cron)
96637 ?? D 0:00.00 cron: running job (cron)
96638 ?? D 0:00.00 cron: running job (cron)
96639 ?? D 0:00.00 cron: running job (cron)
96642 ?? D 0:00.00 cron: running job (cron)
96650 ?? D 0:00.00 cron: running job (cron)
29393 p0 D+ 22:04.58 /usr/local/bin/rsync

real 0m0.012s
user 0m0.000s
sys 0m0.010s
/

real 0m0.019s
user 0m0.000s
sys 0m0.016s
/var

real 0m0.028s
user 0m0.008s
sys 0m0.018s
/diskless

real 0m0.017s

RE: Ufs dead-locks on freebsd 6.2

user 0m0.008s
sys 0m0.007s
/usr

real 0m0.016s
user 0m0.000s
sys 0m0.015s
/d2

real 0m0.024s
user 0m0.000s
sys 0m0.023s
/exports/home

real 0m2.559s
user 0m0.216s
sys 0m2.307s

-----Original Message-----

From: owner-freebsd-fs@xxxxxxxxxxxxx

[\[mailto:owner-freebsd-fs@xxxxxxxxxxxxx\]](mailto:owner-freebsd-fs@xxxxxxxxxxxxx)

On Behalf Of Andrew Edwards

Sent: Friday, May 18, 2007 6:44 PM

To: freebsd-fs@xxxxxxxxxxxxx; freebsd-performance@xxxxxxxxxxxxx

Subject: RE: Ufs dead-locks on freebsd 6.2

Okay, I let memtest run for a full day and there has been no memory errors. What do I do next? Just to be on the safe side I'll fsck all of my fs's and try to reproduce the problem again.

I also don't know what zonelimit is, I see this on similarly

configured

machines but running 5.4. I know it's related to network as I periodically get network connections to work i.e. ssh, ftp (both

server

and client side) but eventually the box will deadlock. Should I start

a

different thread on this? Happens about once every 30 days on two server although I haven't checked the exact timing.

-----Original Message-----

From: owner-freebsd-fs@xxxxxxxxxxxxx

RE: Ufs dead-locks on freebsd 6.2

[\[mailto:owner-freebsd-fs@xxxxxxxxxxx\]](mailto:owner-freebsd-fs@xxxxxxxxxxx)

On Behalf Of Eric Anderson
Sent: Friday, May 18, 2007 3:09 PM
To: Kris Kennaway
Cc: freebsd-fs@xxxxxxxxxxx
Subject: Re: Ufs dead-locks on freebsd 6.2

On 05/18/07 14:00, Kris Kennaway wrote:

On Thu, May 17, 2007 at 11:38:20PM -0500, Eric Anderson wrote:

On 05/17/07 12:47, Kostik Belousov wrote:

On Thu, May 17, 2007 at 01:03:37PM
-0400, Andrew Edwards wrote:

Here it is.

```
db> show vnode
0xccd47984
vnode 0xccd47984: tag ufs,
type VDIR
usecount 5135, writecount
0, refcount 5137
mountedhere 0
flags (VV_ROOT)
v_object 0xcd02518c ref 0
pages 1
#0 0xc0593f0d at
lockmgr+0x4ed
#1 0xc06b8e0e at
ffs_lock+0x76
#2 0xc0739787 at
VOP_LOCK_APV+0x87
#3 0xc0601c28 at
vn_lock+0xac
#4 0xc05ee832 at
lookup+0xde
#5 0xc05ee4b2 at
namei+0x39a
#6 0xc05e2ab0 at
unp_connect+0xf0
#7 0xc05e1a6a at
uipc_connect+0x66
#8 0xc05d9992 at
soconnect+0x4e
#9 0xc05dec60 at
kern_connect+0x74
#10 0xc05debd5 at
connect+0x2f
```

RE: Ufs dead-locks on freebsd 6.2

RE: Ufs dead-locks on freebsd 6.2

#11 0xc0723e2b at
syscall+0x25b
#12 0xc070ee0f at
Xint0x80_syscall+0x1f

ino 2, on dev amrd0s1a

It seems to be the sort of things that cannot happen. VOP_LOCK() returned 0, but vnode was not really locked.

Although claiming that kernel code cannot have such bug is too optimistic, I would first make sure that:

1. You checked the memory of the machine.
2. Your kernel is built from pristine sources.

This looks precisely like a lock I was seeing on one of my NFS

servers.

Only one of the filesystems would cause it, but it was the same

one

each time, not necessarily under any kind of load. Things like mountd would get wedged in state 'ufs', and other things would get stuck in one of the lock states (I can't recall).

...so you cannot conclude that it looks "precisely like" this case.

Please, don't confuse bug reports by this kind of claim unless you have made a detailed comparison of the debugging traces to yours.

Understood – my mistake.

Eric

freebsd-fs@xxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-fs>

To unsubscribe, send any mail to "freebsd-fs-unsubscribe@xxxxxxxxxxx"

RE: Ufs dead-locks on freebsd 6.2

RE: Ufs dead-locks on freebsd 6.2

freebsd-fs@xxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-fs>

To unsubscribe, send any mail to "freebsd-fs-unsubscribe@xxxxxxxxxxx"

freebsd-fs@xxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-fs>

To unsubscribe, send any mail to "freebsd-fs-unsubscribe@xxxxxxxxxxx"

freebsd-performance@xxxxxxxxxxx mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-performance>

To unsubscribe, send any mail to "freebsd-performance-unsubscribe@xxxxxxxxxxx"