

RE: kernel deadlock

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2003-07/0469.html>

From: Don Bowman (don_at_sandvine.com)

Date: 07/30/03

To: Don Bowman <don@sandvine.com>, 'Robert Watson' <rwatson@freebsd.org>, Dave Dolson <ddolson@sa>
Date: Tue, 29 Jul 2003 19:23:36 -0400

> *From: Robert Watson [mailto:rwatson@freebsd.org]*
> > *On Tue, 29 Jul 2003, Dave Dolson wrote:*
> >
> > > *To follow up, I've discovered that the system has*
> *exhausted its "FFS*
> > > *node" malloc type.*
> ...
> >
> > *Some problems with this have turned up in -CURRENT on large-memory*
> > *machines where some of the scaling factors have been off. In*
>
> *We currently have kern.maxvnodes=70354 set (automatically*
> *scaled). This*
> *is a 1GB box.*
>
> *I will try re-running the test with less.*
>
> *when it hits kern.maxvnodes, what will it do?*

So I dropped kern.maxvnodes in half (to 35000). This has a 1GB of physical memory in a 2x xeon (w/ HTT enabled, so 4 procs).

when it hit the limit, the system stopped switching amongst processes. my vmstat blocked in 'vlruwk'.

I merged kern/52425 (kern/vfs_subr.c 1.249.2.30, kern/vfs_syscalls.c 1.151.2.18, sys/mount.h 1.89.2.7) which is supposed to address this, but it didn't. [we're running 4.7].

after a long time, my ^C to the vmstat came through, and my shell prompt came back, but then bash stopped in 'inode'.

In this case i'm not short of memory.

the test is doing this:

```
for (i = 0; i < 100000; i++)  
  mkdir dir.$i  
  cd dir.$i
```

and while it was running i had:

```
while true  
do  
  vmstat -m | grep FFS  
  sleep 1  
done
```

running to watch it.

So it seems the problem may not be running out of memory
in the malloc pool, but in the vnode reclamation?

--don

freebsd-stable@freebsd.org mailing list

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

To unsubscribe, send any mail to "freebsd-stable-unsubscribe@freebsd.org"