

DragonFlyBSD - Bug #1069

rsync + hammer + nfs = trouble

07/13/2008 07:30 PM - aoiko

Status: Closed	Start date:
Priority: Normal	Due date:
Assignee:	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version:	
Description Tried using hammer for the first time today. Created a hammer volume on two partitions, exported it over nfs, mounted it on my fileserver and tried rsyncing some GB of data. After 20+ GB, the machine starts acting up: I can't start new processes (command line editing over ssh works though) and the rsync processes on the non-dfly box is stuck. Taking a look at the console, I can see three "cache_lock: blocked on 0x... (.filename.c.garbage)" messages. Some nfs-hammer interaction perhaps? Aggelos	

History

#1 - 07/13/2008 07:58 PM - dillon

:Tried using hammer for the first time today. Created a hammer volume on
:two partitions, exported it over nfs, mounted it on my fileserver and
:tried rsyncing some GB of data. After 20+ GB, the machine starts acting
:up: I can't start new processes (command line editing over ssh works though)
:and the rsync processes on the non-dfly box is stuck. Taking a look at the
:console, I can see three "cache_lock: blocked on 0x... (.filename.c.garbage)"
:messages. Some nfs-hammer interaction perhaps?
:
:Aggelos

It sounds like a buffer cache deadlock, which should be easy for me to fix. Were you reading from the NFS mount or writing to the NFS mount?

I'm assuming you are running the latest HEAD. If so, please drop into the debugger (e.g. ctl-alt-esc) and generate a crash dump, and get it onto leaf so I can take a look.

-Matt
Matthew Dillon
<dillon@backplane.com>

#2 - 07/13/2008 08:12 PM - aoiko

Only rsync (via nfs) was accessing the filesystem.

I am running HEAD from ~28h ago. Unfortunately, IIRC the machine doesn't have a dumpdev configured (yeah, I know). Is there any information I can get you via DDB?

Aggelos

#3 - 07/14/2008 12:21 AM - dillon

```
> :Aggelos
>
> It sounds like a buffer cache deadlock, which should be easy for me to
> fix. Were you reading from the NFS mount or writing to the NFS mount?
:
:Only rsync (via nfs) was accessing the filesystem.
:
> I'm assuming you are running the latest HEAD. If so, please drop
> into the debugger (e.g. ctl-alt-esc) and generate a crash dump, and
> get it onto leaf so I can take a look.
:
:I am running HEAD from ~28h ago. Unfortunately, IIRC the machine doesn't have
:a dumpdev configured (yeah, I know). Is there any information I can get you
:via DDB?
:
:Aggelos
```

Try it with latest head, to make sure the bug is still there.

I need to know whether rdist was reading HAMMER via NFS or writing to HAMMER via NFS, and what kinds of file structure was being copied? (big files, little files, hundreds of files, thousands of files.. what?).

-Matt
Matthew Dillon
<dillon@backplane.com>

#4 - 07/14/2008 12:49 AM - aoiko

```
[...]  
> Try it with latest head, to make sure the bug is still there.
```

Sure.

Just an rsync -av src dst, where dst is the nfs-mounted hammer volume, which was just created, so I don't think rsync could read much. It got stuck in the devel/ directory (in an old dragonfly sys/ checkout even). The file was

some version of ibcs2_stat.c. Before getting there it had copied over a few hundred thousands of files, some small, some average-sized. I'm going to make sure I can get a dump and then try it with latest head.

Aggelos

#5 - 07/14/2008 03:27 AM - dillon

:Just an rsync -av src dst, where dst is the nfs-mounted hammer volume, which
:was just created, so I don't think rsync could read much. It got stuck in
:the devel/ directory (in an old dragonfly sys/ checkout even). The file was
:some version of ibcs2_stat.c. Before getting there it had copied over a few
:hundred thousands of files, some small, some average-sized. I'm going to
:make sure I can get a dump and then try it with latest head.
:
:Aggelos

I'm doing an rdist test with pkgbox's /archive (about 56GB and 2 million inodes), but so far it is running smoothly.

-Matt
Matthew Dillon
<dillon@backplane.com>

#6 - 07/14/2008 12:46 PM - aoiko

On Monday 14 July 2008, Aggelos Economopoulos wrote:

[...]

> Just an rsync -av src dst, where dst is the nfs-mounted hammer volume, which
> was just created, so I don't think rsync could read much. It got stuck in
> the devel/ directory (in an old dragonfly sys/ checkout even). The file was
> some version of ibcs2_stat.c. Before getting there it had copied over a few
> hundred thousands of files, some small, some average-sized. I'm going to
> make sure I can get a dump and then try it with latest head.

Hit an unrelated hammer bug (latest HEAD) while trying to reproduce. Should be easy to fix; uploading dump to my ~/crash directory on leaf.

Aggelos

#7 - 07/14/2008 05:42 PM - dillon

> Just an rsync -av src dst, where dst is the nfs-mounted hammer volume, which
> was just created, so I don't think rsync could read much. It got stuck in

> the devel/ directory (in an old dragonfly sys/ checkout even). The file was
> some version of ibcs2_stat.c. Before getting there it had copied over a few
> hundred thousands of files, some small, some average-sized. I'm going to
> make sure I can get a dump and then try it with latest head.
:
:Hit an unrelated hammer bug (latest HEAD) while trying to reproduce. Should
:be easy to fix; uploading dump to my ~/crash directory on leaf.
:
:Aggelos

(oops, forgot to CC to the list)

I'm looking at the core file but I don't see anything wrong. There are
no processes stuck on anything as far as I can tell.

Could it be an issue with your network? Can you access the machine via
its console and see if the network still works?

-Matt
Matthew Dillon
<dillon@backplane.com>

#8 - 07/14/2008 05:51 PM - aoiko

This is not the same issue as I reported previously. AFAICT it's a null
pointer dereference:

```
(kgdb) bt
#0  dumpsys () at ./machine/thread.h:83
#1  0xc0256b2d in boot (howto=260)
at /home/aggelos/imports/vcs/dfgit/cz/sys/kern/kern_shutdown.c:375
#2  0xc0256df0 in panic (fmt=0xc04437c5 "from debugger")
at /home/aggelos/imports/vcs/dfgit/cz/sys/kern/kern_shutdown.c:800
#3  0xc016133d in db_panic (addr=-1070072130, have_addr=0, count=-1, modif=0xd45b25e4 "")
at /home/aggelos/imports/vcs/dfgit/cz/sys/ddb/db_command.c:447
#4  0xc01619a8 in db_command_loop ()
at /home/aggelos/imports/vcs/dfgit/cz/sys/ddb/db_command.c:343
#5  0xc0163f54 in db_trap (type=12, code=2)
at /home/aggelos/imports/vcs/dfgit/cz/sys/ddb/db_trap.c:71
#6  0xc03ed786 in kdb_trap (type=12, code=2, regs=0xd45b272c)
at /home/aggelos/imports/vcs/dfgit/cz/sys/platform/pc32/i386/db_interface.c:148
#7  0xc0403877 in trap_fatal (frame=0xd45b272c, eva=<value optimized out>)
at /home/aggelos/imports/vcs/dfgit/cz/sys/platform/pc32/i386/trap.c:1133
#8  0xc0403a07 in trap_pfault (frame=0xd45b272c, usermode=0, eva=0)
at /home/aggelos/imports/vcs/dfgit/cz/sys/platform/pc32/i386/trap.c:1039
#9  0xc0404006 in trap (frame=0xd45b272c)
at /home/aggelos/imports/vcs/dfgit/cz/sys/platform/pc32/i386/trap.c:719
#10 0xc03ee487 in calltrap ()
at /home/aggelos/imports/vcs/dfgit/cz/sys/platform/pc32/i386/exception.s:785
#11 0xc037febe in hammer_alloc_objid (hmp=0xd4827000, dip=0xd17c46d0)
at /home/aggelos/imports/vcs/dfgit/cz/sys/vfs/hammer/hammer_transaction.c:205
#12 0xc0373280 in hammer_create_inode (trans=0xd45b27dc, vap=0xd45b2a70, cred=0xc237aa30,
dip=0xd17c46d0, pfsm=0x0, ipp=0xd45b27d8)
at /home/aggelos/imports/vcs/dfgit/cz/sys/vfs/hammer/hammer_inode.c:524
#13 0xc0382865 in hammer_vop_nmkdir (ap=0xd45b2818)
```

```
at /home/aggelos/imports/vcs/dfgit/cz/sys/vfs/hammer/hammer_vnops.c:1050
#14 0xc02aa2c3 in vop_nmkdir (ops=0xc23add70, nch=0xd45b28e0, dvp=0xd47ed7e8, vpp=0xd45b28c8,
cred=0xc237aa30, vap=0xd45b2a70)
at /home/aggelos/imports/vcs/dfgit/cz/sys/kern/vfs_vopops.c:1029
---Type <return> to continue, or q <return> to quit---
---Type <return> to continue, or q <return> to quit---#15 0xc03075b2 in nfsrv_mkdir (nfsd=0xc237a9b0, slp=0xc23aca50, td=0xc23799c0,
mrq=0xd45b2ba8)
at /home/aggelos/imports/vcs/dfgit/cz/sys/vfs/nfs/nfs_serv.c:2583
#16 0xc0318f2d in sys_nfssvc (uap=0xd45b2cf0)
at /home/aggelos/imports/vcs/dfgit/cz/sys/vfs/nfs/nfs_syscalls.c:590
#17 0xc040469b in syscall2 (frame=0xd45b2d40)
at /home/aggelos/imports/vcs/dfgit/cz/sys/platform/pc32/i386/trap.c:1384
#18 0xc03ee536 in Xint0x80_syscall ()
at /home/aggelos/imports/vcs/dfgit/cz/sys/platform/pc32/i386/exception.s:876
#19 0x080498c0 in ?? ()
Backtrace stopped: previous frame inner to this frame (corrupt stack?)
(kgdb)
```

In the previous bug (hit the new bug while trying to reproduce it) the network was up; as I said, I could edit the command line over ssh just fine. In this one, the machine crashed, so obviously no network access was possible.

Aggelos

#9 - 08/28/2009 08:04 PM - tuxillo

There have been many fixes on HAMMER and NFS since this bug was reported. Is it possible that it could be reproduced with HEAD?

#10 - 03/26/2010 06:04 PM - sjg

User did not provide an update on request, assuming fixed. Please re-open if not.