

DragonFlyBSD - Bug #1990

/mnt too large to mount

02/15/2011 06:26 PM - peur.neu

Status:	New	Start date:	
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
Description			
Hi!			
beyond dma stuff i got a crash after several data parallel updates			
after reboot /dev/ad1s1d seems to be crashed but it's not.			
i did fsck /dev/ad1s1a booted ok.			
hammer still faulty!!!			
then i did hammer -f /dev/ad1s1d checkmap			
the filesystem is on place!!!			
then i did mount_hammer -o ro /dev/ad1s1d /mnt			
mounts fine!!!			
BUT doing mount_hammer -o rw /dev/ad1s1d /mnt			
UNDO_REDO etc			
/mnt too large			
error.			
HELP			
thank			

History

#1 - 02/16/2011 02:39 AM - dillon

```
:BUT doing mount_hammer -o rw /dev/ad1s1d /mnt
:
:UNDO_REDO etc
:/mnt too large
:error.
:
:HELP
:
:thank
```

I can't find that '/mnt too large' error in the code, what is the precise error message you are getting?

-Matt
Matthew Dillon
<dillon@backplane.com>

#2 - 02/16/2011 06:31 PM - peur.neu

Hi Matt

```
-----
Mounting root from HAMMER:da13s1d
tryroot da13s1d
-----
```

```
# mount_hammer -o rw /dev/ad2s1d /tmp/fs
<dmesg>
-----
HAMMER(ROOT) recovery check seqno=0006ceb2
HAMMER(ROOT) recovery range 300000000399da90-300000000399da90
HAMMER(ROOT) recovery nexto 300000000399da90 endseqno=0006ced3
-----
</dmesg>
that's HAMMER working nice on (RO) mode
then again.
```

```
# mount -o rw /dev/ad2s1d /tmp/fs
mount_hammer: mount on /tmp/fs: Result too large
```

```
<dmesg>
-----
HAMMER(ROOT) recovery check seqno=0370dbbf
HAMMER(ROOT) recovery range 300000000582bd0-3000000001f65c78
HAMMER(ROOT) recovery nexto 300000000af64ab0 endseqno=0371bd5b
HAMMER(ROOT) recovery UNDO 300000000582bd0-3000000001f65c78 (27144360
Bytes) (RW)
HAMMER(ROOT) Found REDO_SYNC 30000000055e528
HAMMER(ROOT) Ignoring extra REDO_SYNC records in UNDO/REDO FIFO
HAMMER(ROOT) Ignoring extra REDO_SYNC records in UNDO/REDO FIFO
HAMMER(ROOT) Ignoring extra REDO_SYNC records in UNDO/REDO FIFO
HAMMER(ROOT) Recovery complete
HAMMER(ROOT) Recovery redo 300000000582bd0-3000000001f65c78 (27144360
Bytes) (RW)
HAMMER(ROOT) Find extended redo 30000000055e528, 149160 extbytes
HAMMER(ROOT) Find extended redo failed 34, unable to run REDO
HAMMER(ROOT) End redo recovery
-----
</dmesg>
```

if i don't succeed i'll be patching the vfs this next week. that will succeed to mount (RW) and flush.

thanks

"Matthew Dillon" <dillon@apollo.backplane.com> escribió en el mensaje news:201102160238.p1G2cx4Z039002@apollo.backplane.com...

```
> :BUT doing mount_hammer -o rw /dev/ad1s1d /mnt
> :
> :UNDO_REDO etc
> :/mnt too large
> :error.
> :
> :HELP
> :
> :thank
>
> I can't find that '/mnt too large' error in the code, what is the
> precise error message you are getting?
>
> -Matt
> Matthew Dillon
> <dillon@backplane.com>
```

#3 - 02/16/2011 11:24 PM - thomas.nikolajsen

>Re: [issue1984] hammer mount fails after crash - HAMMER: FIFO record badhead signature ..

>From: "Diegu" <peur.neu@xxxxxxxxx>

>Date: Tue, 15 Feb 2011 17:35:39 +0100

>

>looks quite similar to my stuff.

Yes both issues seems to be jammed UNDO/REDO FIFO;

what is your system like:

HAMMER file system size, RAM amount, DragonFly version, i386 or x86_64?

>have you tried to flush the PFS from another fresh install ?

What do you mean by 'flush the PFS'?

I tried to mount HAMMER file system from another DragonFly install (as affected FS was root FS); result is described in issue1984.

>theres no flushing without a master/slave PFS then you need a fresh install
>to prove the point right?
I'm still not sure what you mean by 'flushing'.

>i've tried to prune over a RO but it seems not possible.
>by the way. HAMMER recovery works for me but it won't save the prev install.
Well, if FS is mounted RO it makes a lot of sense that it can't be modified,
like pruned.
If data can be read from RO mount, you can recover by just copy data to
another disk (or other partition on same disk); this won't save history.
Using 'hammer recover' could be used to the same effect.

Lets see which diagnose Matt does; FS will hopefully be revived.

-thomas