

## Ceph - Bug #563

### osd: btrfs, warning at inode.c ( btrfs\_orphan\_commit\_root )

11/08/2010 11:05 AM - Wido den Hollander

<b>Status:</b>	Closed	<b>Start date:</b>	11/08/2010
<b>Priority:</b>	Immediate	<b>Due date:</b>	
<b>Assignee:</b>	Sage Weil	<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	v0.24.3	<b>Spent time:</b>	1.00 hour
<b>Source:</b>		<b>Reviewed:</b>	
<b>Tags:</b>		<b>Affected Versions:</b>	
<b>Backport:</b>		<b>ceph-qa-suite:</b>	
<b>Regression:</b>	No	<b>Pull request ID:</b>	
<b>Severity:</b>	3 - minor		

#### Description

I'm running the unstable branch and I'm seeing in my dmesg:

```
[ 9786.475700] -----[ cut here ]-----
[ 9786.475718] WARNING: at /build/builddd/linux-lts-backport-natty-2.6.37/fs/btrfs/inode.c:2143 btr
fs_orphan_commit_root+0xb0/0xc0 [btrfs] ()
[ 9786.475720] Hardware name: PDSMU
[ 9786.475721] Modules linked in: bonding btrfs zlib_deflate crc32c libcrc32c radeon ttm led_class
serio_raw drm_kms_helper drm lp parport shpchp i2c_algo_bit i3000_edac edac_core arcmsr e1000e fl
oppy
[ 9786.475739] Pid: 10, comm: kworker/0:1 Not tainted 2.6.37-2-server #9~lucid2-Ubuntu
[ 9786.475741] Call Trace:
[ 9786.475748] [<ffffffff8106377f>] warn_slowpath_common+0x7f/0xc0
[ 9786.475752] [<ffffffff810637da>] warn_slowpath_null+0x1a/0x20
[ 9786.475761] [<ffffffffffa0249bd0>] btrfs_orphan_commit_root+0xb0/0xc0 [btrfs]
[ 9786.475771] [<ffffffffffa0245391>] commit_fs_roots+0xa1/0x140 [btrfs]
[ 9786.475781] [<ffffffffffa02467c0>] btrfs_commit_transaction+0x350/0x730 [btrfs]
[ 9786.475785] [<ffffffffff810851f0>] ? autoremove_wake_function+0x0/0x40
[ 9786.475789] [<ffffffffff815c109e>] ? schedule+0x44e/0xa70
[ 9786.475799] [<ffffffffffa0246e40>] ? do_async_commit+0x0/0x30 [btrfs]
[ 9786.475808] [<ffffffffffa0246e5f>] do_async_commit+0x1f/0x30 [btrfs]
[ 9786.475812] [<ffffffffff8107e025>] process_one_work+0x125/0x440
[ 9786.475815] [<ffffffffff810807e0>] worker_thread+0x170/0x410
[ 9786.475817] [<ffffffffff81080670>] ? worker_thread+0x0/0x410
[ 9786.475820] [<ffffffffff81084c96>] kthread+0x96/0xa0
[ 9786.475824] [<ffffffffff8100cf64>] kernel_thread_helper+0x4/0x10
[ 9786.475827] [<ffffffffff81084c00>] ? kthread+0x0/0xa0
[ 9786.475829] [<ffffffffff8100cf60>] ? kernel_thread_helper+0x0/0x10
[ 9786.475831] ---[ end trace 29ffdec9a7921eec ]---
```

I joined the #btrfs channel and they said:

```
20:00 < widodh> I've just upgraded to 2.6.37-rc1 to test Ceph, but I'm seeing this right now in my
dmesg:
           http://pastebin.com/3HiGiNgy ( inode.c / btrfs_orphan_commit_root )
20:01 < josef> hrm looks like sage screwed something up with async commit
20:02 < widodh> aha, file it over there then?
20:03 < josef> well tell sage about it, he's the one who wrote the code so he'll probably be able
to fix it quicker
20:03 < widodh> yes, sure. I thought it was a btrfs bug, that's why I came along here
20:03 < widodh> thanks!
20:05 < josef> well it is a bug in btrfs, but he will be able to fix it quicker than i will
```

I've seen this while I had one VM running (qemu-rbd) with bonnie++ in it.

## History

### #1 - 11/08/2010 04:18 PM - Sage Weil

- Priority changed from Normal to Immediate

### #2 - 11/09/2010 01:23 PM - Sage Weil

- Assignee set to Sage Weil

### #3 - 11/22/2010 09:28 AM - Wido den Hollander

Just to update the issue, Sage asked me to change something in FileStore.cc, tried that for some days, but that didn't make a difference.

### #4 - 11/22/2010 12:59 PM - Sage Weil

Is the stack trace you're getting now identical, or different? The FileStore.cc change *should* have avoided the async snap create, which means you shouldn't see

```
[ 9786.475808] [<ffffffffffa0246e5f>] do_async_commit+0x1f/0x30 [btrfs]
```

in the stack dump. Either my hack was wrong, or there is another caller I'm forgetting about.

### #5 - 11/23/2010 07:44 AM - Wido den Hollander

I'll have to rebuild, since I didn't look at the messages that closely.

### #6 - 12/15/2010 10:39 AM - Wido den Hollander

I've just hit the bug again, this while I was running a rsync to my Ceph cluster.

```
[ 8701.939180] -----[ cut here ]-----
[ 8701.939216] WARNING: at /build/builddd/linux-lts-backport-natty-2.6.37/fs/btrfs/inode.c:2143 btrfs_orphan_co
mmit_root+0xb0/0xc0 [btrfs]()
[ 8701.939220] Hardware name: X7DVL
[ 8701.939222] Modules linked in: cryptd aes_x86_64 aes_generic ceph libceph btrfs zlib_deflate libcrc32c rade
on ttm drm_kms_helper drm i5000_edac edac_core ioatdma shpchp ghes i2c_algo_bit ppdev i5k_amb hed parport_pc l
p dca serio_raw joydev parport usbhid hid ahci e1000e libahci
[ 8701.939260] Pid: 23797, comm: kworker/1:1 Tainted: G          W   2.6.37-9-server #22-lucid1-Ubuntu
[ 8701.939263] Call Trace:
[ 8701.939275]   [<ffffffffff8106434f>] warn_slowpath_common+0x7f/0xc0
[ 8701.939280]   [<ffffffffff810643aa>] warn_slowpath_null+0x1a/0x20
[ 8701.939294]   [<ffffffffffa027df00>] btrfs_orphan_commit_root+0xb0/0xc0 [btrfs]
[ 8701.939308]   [<ffffffffffa0279551>] commit_fs_roots+0xa1/0x140 [btrfs]
[ 8701.939321]   [<ffffffffffa027a9d0>] btrfs_commit_transaction+0x350/0x730 [btrfs]
[ 8701.939328]   [<ffffffffff81085e80>] ? autoremove_wake_function+0x0/0x40
[ 8701.939334]   [<ffffffffff815c85f4>] ? schedule+0x454/0xa50
[ 8701.939348]   [<ffffffffffa027b050>] ? do_async_commit+0x0/0x30 [btrfs]
[ 8701.939361]   [<ffffffffffa027b06f>] do_async_commit+0x1f/0x30 [btrfs]
[ 8701.939385]   [<ffffffffff8107ecb5>] process_one_work+0x125/0x440
[ 8701.939390]   [<ffffffffff81081470>] worker_thread+0x170/0x410
[ 8701.939394]   [<ffffffffff81081300>] ? worker_thread+0x0/0x410
[ 8701.939398]   [<ffffffffff81085926>] kthread+0x96/0xa0
[ 8701.939404]   [<ffffffffff8100cf24>] kernel_thread_helper+0x4/0x10
[ 8701.939413]   [<ffffffffff81085890>] ? kthread+0x0/0xa0
[ 8701.939417]   [<ffffffffff8100cf20>] ? kernel_thread_helper+0x0/0x10
[ 8701.939420] ---[ end trace d66645fbbece7ada ]---
```

This was with the latest RC branch ( 89d5c91e7d207d646651f8959ee37a15ea199d1b ).

After this, the load on my machine went up, it seemed that 'btrfs-submit-0' went stale.

```
root@noisy:~# ps aux|grep btrfs-submit
root      22421  0.2  0.0    0    0 ?        D   17:37   0:18 [btrfs-submit-0]
root      22501  0.2  0.0    0    0 ?        D   17:37   0:21 [btrfs-submit-0]
root      22581  0.2  0.0    0    0 ?        S   17:37   0:21 [btrfs-submit-0]
root      22661  0.2  0.0    0    0 ?        D   17:37   0:22 [btrfs-submit-0]
root@noisy:~#
```

But my rsync was still running.

### #7 - 12/15/2010 01:01 PM - Wido den Hollander

I took the for-linus branch, since both next-rc and master wouldn't compile against 2.6.37-rc5. Due to this error:

```
root@noisy:/usr/src/linux-headers-2.6.37-9-server# make SUBDIRS=/usr/src/btrfs-unstable/fs/btrfs/
CC [M] /usr/src/btrfs-unstable/fs/btrfs/super.o
CC [M] /usr/src/btrfs-unstable/fs/btrfs/ctree.o
CC [M] /usr/src/btrfs-unstable/fs/btrfs/extent-tree.o
/usr/src/btrfs-unstable/fs/btrfs/extent-tree.c: In function 'btrfs_issue_discard':
/usr/src/btrfs-unstable/fs/btrfs/extent-tree.c:1750: error: 'BLKDEV_IFL_WAIT' undeclared (first use in this function)
/usr/src/btrfs-unstable/fs/btrfs/extent-tree.c:1750: error: (Each undeclared identifier is reported only once
/usr/src/btrfs-unstable/fs/btrfs/extent-tree.c:1750: error: for each function it appears in.)
/usr/src/btrfs-unstable/fs/btrfs/extent-tree.c:1750: error: 'BLKDEV_IFL_BARRIER' undeclared (first use in this function)
make[1]: *** [/usr/src/btrfs-unstable/fs/btrfs/extent-tree.o] Error 1
make: *** [_module_/usr/src/btrfs-unstable/fs/btrfs] Error 2
root@noisy:/usr/src/linux-headers-2.6.37-9-server#
```

Not sure why, the BLKDEV\_IFL\_WAIT marco was indeed removed in September 2010, btrfs should be updated I think?

Anyway, with the for-linus branch (6 weeks old atm) the message is still there, I even got a new one:

```
[17427.887664] -----[ cut here ]-----
[17427.887753] kernel BUG at /build/build/linux-lts-backport-natty-2.6.37/fs/btrfs/inode.c:6404!
[17427.887880] invalid opcode: 0000 [#1] SMP
[17427.887955] last sysfs file: /sys/module/btrfs/initstate
[17427.888038] CPU 0
[17427.888076] Modules linked in: cryptd aes_x86_64 aes_generic ceph libceph btrfs zlib_deflate libcrc32c rade
on ttm drm_kms_helper drm i5000_edac edac_core ioatdma shpchp ghes i2c_algo_bit ppdev i5k_amb hed parport_pc l
p dca serio_raw joydev parport usbhid hid ahci e1000e libahci
[17427.888861]
[17427.888898] Pid: 26125, comm: cosd Tainted: G          W   2.6.37-9-server #22~lucid1-Ubuntu X7DVL/X7DVL
[17427.889034] RIP: 0010:[<ffffffffa028297b>] [<ffffffffa028297b>] btrfs_truncate+0x22b/0x240 [btrfs]
[17427.889210] RSP: 0018:ffff880111ca3d78  EFLAGS: 0010286
[17427.889292] RAX: 00000000fffffe4  RBX: ffff88019bc37100  RCX: ffff8801a0addc60
[17427.889392] RDX: 0000000000000000  RSI: fffffea0005b26058  RDI: 0000000000000206
[17427.889494] RBP: ffff880111ca3da8  R08: 0000000000000001  R09: ffff880111ca3b68
[17427.889594] R10: 0000160000000000  R11: 6db6db6db6db6db7  R12: ffff880108296000
[17427.889694] R13: ffff880070b36768  R14: ffff880070b365d0  R15: 0000000000000000
[17427.889797] FS: 00007f519ee78700(0000)  GS:ffff8800cfc00000(0000)  knlGS:0000000000000000
[17427.889920] CS: 0010  DS: 0000  ES: 0000  CR0: 000000008005003b
[17427.890000] CR2: 00007f3dd097c000  CR3: 0000000198a84000  CR4: 00000000000006f0
[17427.890100] DR0: 0000000000000000  DR1: 0000000000000000  DR2: 0000000000000000
[17427.890201] DR3: 0000000000000000  DR6: 00000000ffff0fff0  DR7: 0000000000000400
[17427.890302] Process cosd (pid: 26125, threadinfo ffff880111ca2000, task ffff8801a0698000)
[17427.890426] Stack:
[17427.890464] ffff880111ca3dc8 ffff880070b36768 0000000000000000 ffff88019bc37100
[17427.890612] ffff880111ca3ec8 0000000000000000 ffff880111ca3dd8 ffffffff81118066
```

```

[17427.890759] ffff880111ca3ec8 0000000000000000 ffff880070b36768 ffff880108296000
[17427.890906] Call Trace:
[17427.890953] [<ffffffff81118066>] vmtruncate+0x56/0x70
[17427.891042] [<fffffffffa02871ea>] btrfs_setattr_size+0xaa/0x230 [btrfs]
[17427.891154] [<fffffffffa0287413>] btrfs_setattr+0xa3/0xb0 [btrfs]
[17427.891254] [<ffffffff8117bc30>] notify_change+0x170/0x2e0
[17427.891347] [<ffffffff81161bd4>] do_truncate+0x64/0xa0
[17427.891428] [<ffffffff8116d113>] ? generic_permission+0x23/0xc0
[17427.891523] [<ffffffff8116cf75>] ? get_write_access+0x45/0x70
[17427.891609] [<ffffffff81161d59>] do_sys_truncate+0x149/0x150
[17427.891691] [<ffffffff81161d6e>] sys_truncate+0xe/0x10
[17427.891775] [<ffffffff8100c102>] system_call_fastpath+0x16/0x1b
[17427.891866] Code: 85 c0 74 b2 0f 0b eb fe 4c 89 ea 4c 89 e6 48 89 df e8 1a 52 01 00 e9 c7 fe ff ff 0f 0b eb
fe 0f 0b eb fe 0f 0b eb fe 0f 0b eb fe <0f> 0b eb fe 0f 0b eb fe 0f 0b eb fe 66 0f 1f 84 00 00 00 00
[17427.892652] RIP [<fffffffffa028297b>] btrfs_truncate+0x22b/0x240 [btrfs]
[17427.892782] RSP <ffff880111ca3d78>
[17427.917828] ---[ end trace d66645fbbec7add ]---

```

#### #8 - 12/16/2010 02:09 AM - Wido den Hollander

Compiled from the kernel GIT ( a4851d8f7d6351a395d36ae8fdcf41745a832d76 ) last night and then started a rsync this morning.

The original error/message didn't come back, it changed slightly, I got:

```

[ 1488.110825] -----[ cut here ]-----
[ 1488.110873] WARNING: at fs/btrfs/inode.c:2143 btrfs_orphan_commit_root+0xb0/0xc0 [btrfs] ()
[ 1488.110877] Hardware name: X7DVL
[ 1488.110879] Modules linked in: cryptd aes_x86_64 aes_generic ceph libceph btrfs zlib_deflate libcrc32c rade
on ttm drm_kms_helper drm i5000_edac ppdev edac_core ioatdma i5k_amb i2c_algo_bit parport_pc lp parport ghes s
hpchp serio_raw hed dca joydev usbhid hid ahci libahci e1000e
[ 1488.110920] Pid: 3698, comm: cosd Not tainted 2.6.37-rc5-00338-ga4851d8 #1
[ 1488.110923] Call Trace:
[ 1488.110938] [<ffffffff81061f3f>] warn_slowpath_common+0x7f/0xc0
[ 1488.110943] [<ffffffff81061f9a>] warn_slowpath_null+0x1a/0x20
[ 1488.110958] [<fffffffffa02aea20>] btrfs_orphan_commit_root+0xb0/0xc0 [btrfs]
[ 1488.110973] [<fffffffffa02aa0c1>] commit_fs_roots+0xa1/0x140 [btrfs]
[ 1488.110989] [<fffffffffa02ab540>] btrfs_commit_transaction+0x350/0x730 [btrfs]
[ 1488.110996] [<ffffffff81083760>] ? autoremove_wake_function+0x0/0x40
[ 1488.111011] [<fffffffffa02d9aa3>] btrfs_mksubvol+0x363/0x380 [btrfs]
[ 1488.111025] [<fffffffffa02d9bad>] btrfs_ioctl_snap_create_transid+0xed/0x140 [btrfs]
[ 1488.111039] [<fffffffffa02d9cff>] btrfs_ioctl_snap_create+0xff/0x150 [btrfs]
[ 1488.111053] [<fffffffffa02dbc1f>] btrfs_ioctl+0x60f/0xa10 [btrfs]
[ 1488.111060] [<ffffffff81192b4a>] ? fsnotify+0x1ea/0x320
[ 1488.111066] [<ffffffff8116c6ef>] do_vfs_ioctl+0x9f/0x540
[ 1488.111070] [<ffffffff8116cc11>] sys_ioctl+0x81/0xa0
[ 1488.111077] [<ffffffff8100c042>] system_call_fastpath+0x16/0x1b
[ 1488.111081] ---[ end trace 7bcd0685dd3dae52 ]---

```

This message showed up shortly after the rsync started. It has been running now for about 1.5 hours and no new messages appeared.

#9 - 12/16/2010 08:11 AM - Wido den Hollander

My rsync (196k files, 74GB data) finished succesfully, but the btrfs warning repeated itself twice.

#10 - 12/17/2010 12:21 AM - Stefan Majer

Hi,

we see similar problems with kernel 2.6.37-rc6 and ceph build from yesterday (29480f42be8551f47d79282b7376a10a24eb98f4)

Any hints to further nail this down are welcome.

```
[65013.439144] -----[ cut here ]-----
[65013.444287] kernel BUG at fs/btrfs/inode.c:6403!
[65013.449429] invalid opcode: 0000 [#1] SMP
[65013.454014] last sysfs file: /sys/devices/system/cpu/online
[65013.460221] CPU 6
[65013.462264] Modules linked in: btrfs zlib_deflate libcrc32c sit tunnel4 bonding ipv6 ixgbe mdio iomemory_vs
l hpsa igb dca joydev pcspkr serio_raw ghes hed iTCO_wdt iTCO_vendor_support i7core_edac edac_core shpchp squa
shfs usb_storage [last unloaded: scsi_wait_scan]
[65013.488907]
[65013.490565] Pid: 2275, comm: cosd Tainted: G          W   2.6.37-0.rc6.git0.1.el6.x86_64 #1 /ProLiant DL180 G
6
[65013.501736] RIP: 0010:[<ffffffffffa023e0e8>] [<ffffffffffa023e0e8>] btrfs_truncate+0x441/0x477 [btrfs]
[65013.511838] RSP: 0018:ffff8805dcc0bd28  EFLAGS: 00010286
[65013.517754] RAX: 00000000ffffffe4  RBX: ffff8800041d9d78  RCX: ffff8805ddd54000
[65013.525707] RDX: 000000000250000  RSI: ffff8805ddd54090  RDI: 0000000000000206
[65013.533659] RBP: ffff8805dcc0bdc8  R08: ffff88060b65f080  R09: 000000000000005a
[65013.541609] R10: 0000000000000046  R11: 0000000000000011  R12: ffff8805ddd56110
[65013.549559] R13: ffff8805dd379048  R14: ffff8800041d9d78  R15: 0000000000000000
[65013.557510] FS: 00007f9722ffd710(0000)  GS:ffff8800bd200000(0000)  knlGS:0000000000000000
[65013.566524] CS: 0010  DS: 0000  ES: 0000  CR0: 0000000080050033
[65013.572925] CR2: 00007f96b0152000  CR3: 00000005dda91000  CR4: 000000000000006e0
[65013.580874] DR0: 0000000000000000  DR1: 0000000000000000  DR2: 0000000000000000
[65013.588823] DR3: 0000000000000000  DR6: 00000000ffff0ff0  DR7: 00000000000000400
[65013.596772] Process cosd (pid: 2275, threadinfo ffff8805dcc0a000, task ffff8805dcc047a0)
[65013.605786] Stack:
[65013.608024] 0000000000000000 0000000000000000 0000000000000000 00001000ffffffffff
[65013.616303] ffff8800041d9980 0000000000000000 00000000000000ff ffff8800041d9fe0
[65013.624585] ffff8800041d9fe0 ffff8800041d99f0 ffff8805ddd56110 0000000000000000
[65013.632863] Call Trace:
[65013.635588] [<ffffffffff810f1541>] vmtruncate+0x44/0x4f
[65013.641315] [<ffffffffffa024086f>] btrfs_setattr+0x1e9/0x23a [btrfs]
[65013.648302] [<ffffffffff81146182>] ? notify_change+0x18a/0x29d
[65013.654704] [<ffffffffff8114619a>] notify_change+0x1a2/0x29d
[65013.660915] [<ffffffffff8112fea5>] do_truncate+0x6c/0x89
[65013.666729] [<ffffffffff81139b41>] ? get_write_access+0x46/0x4d
[65013.673221] [<ffffffffff8113008e>] sys_truncate+0x111/0x113
[65013.679334] [<ffffffffff8100ac42>] system_call_fastpath+0x16/0x1b
[65013.686025] Code: 83 7b 5c 00 74 13 48 89 de 4c 89 e7 e8 30 e3 ff ff 85 c0 74 04 0f 0b eb fe 48 89 da 4c 89
ee 4c 89 e7 e8 cb ab ff ff 85 c0 74 04 <0f> 0b eb fe 4c 89 ee 4c 89 e7 49 8b 5c 24 20 e8 0b 7a ff ff 85
[65013.707741] RIP [<ffffffffffa023e0e8>] btrfs_truncate+0x441/0x477 [btrfs]
[65013.715232] RSP [<ffff8805dcc0bd28>]
[65013.719692] ---[ end trace c1a53430667a810c ]---
```

**#11 - 12/17/2010 09:33 AM - Sage Weil**

Stefan Majer wrote:

Hi,

we see similar problems with kernel 2.6.37-rc6 and ceph build from yesterday (29480f42be8551f47d79282b7376a10a24eb98f4)

It looks like update\_inode returned -ENOSPC. Is your disk full?

**#12 - 12/17/2010 09:37 AM - Sage Weil**

- Target version set to v0.24.1

**#13 - 12/18/2010 10:09 AM - Wido den Hollander**

Sage, like you asked:

```
2010-12-17 10:14:25.018605 --- 1038 opened log /var/log/ceph/osd.0.log ---
ceph version 0.24~rc (commit:89d5c91e7d207d646651f8959ee37a15ea199d1b)
2010-12-17 10:14:25.021485 7f0b69911720 filestore(/var/lib/ceph/osd.0) mount FIEMAP ioctl is NOT supported
2010-12-17 10:14:25.021597 7f0b69911720 filestore(/var/lib/ceph/osd.0) mount detected btrfs
2010-12-17 10:14:25.021612 7f0b69911720 filestore(/var/lib/ceph/osd.0) mount btrfs CLONE_RANGE ioctl is supported
2010-12-17 10:14:25.059253 7f0b69911720 filestore(/var/lib/ceph/osd.0) mount btrfs SNAP_CREATE is supported
2010-12-17 10:14:25.146531 7f0b69911720 filestore(/var/lib/ceph/osd.0) mount btrfs SNAP_DESTROY is supported
2010-12-17 10:14:25.146851 7f0b69911720 filestore(/var/lib/ceph/osd.0) mount btrfs START_SYNC got 0 Success
2010-12-17 10:14:25.146865 7f0b69911720 filestore(/var/lib/ceph/osd.0) mount btrfs START_SYNC is supported (tr
ansid 3760)
2010-12-17 10:14:25.251022 7f0b69911720 filestore(/var/lib/ceph/osd.0) mount btrfs WAIT_SYNC is supported
2010-12-17 10:14:25.251097 7f0b69911720 filestore(/var/lib/ceph/osd.0) mount WARNING: btrfs snaps enabled, but
no SNAP_CREATE_ASYNC ioctl (from kernel 2.6.37+)
2010-12-17 10:14:25.251340 7f0b69911720 filestore(/var/lib/ceph/osd.0) mount found snaps <1,648991,649204,6493
93>
2010-12-17 10:14:25.454825 7f0b69911720 filestore(/var/lib/ceph/osd.0) mount: enabling WRITEAHEAD journal mode
: btrfs SNAP_CREATE_ASYNC ioctl not detected (v2.6.37+)
```

**#14 - 12/19/2010 02:03 AM - Stefan Majer**

Sage Weil wrote:

Stefan Majer wrote:

Hi,

we see similar problems with kernel 2.6.37-rc6 and ceph build from yesterday (29480f42be8551f47d79282b7376a10a24eb98f4)

It looks like update\_inode returned -ENOSPC. Is your disk full?

Hi,

no there was about 100G of 6TB used.

**#15 - 01/11/2011 10:14 AM - Sage Weil**

- Target version changed from v0.24.1 to v0.24.2

**#16 - 01/22/2011 10:40 AM - Christian Andreetta**

Maybe it's "just" a btrfs issue, not a ceph-related one.

Btrfs was known to lack some consistency in metadata allocation (don't know about the last 2-3 months, but last summer I remember there were some comments in the kernel mailing list about the population/deletion algorithms managing the B-tree nodes):

<https://lkml.org/lkml/2010/6/18/144> (Edward Shishkin, former ReiserFS developer, points some unclear design issues)  
<https://btrfs.wiki.kernel.org/index.php/FAQ> (section 1.3)

**#17 - 01/28/2011 10:50 AM - Sage Weil**

- Target version changed from v0.24.2 to v0.24.3

**#18 - 01/29/2011 09:25 PM - Sage Weil**

- Status changed from New to Closed