

Ceph - Bug #1032

osd: Marked down and become zombies after killing

04/29/2011 01:08 AM - Wido den Hollander

Status:	Can't reproduce	Start date:	04/29/2011
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:	OSD	Estimated time:	0.00 hour
Target version:	v0.37	Spent time:	13.00 hours
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	

Description

I've discussed this a bit with Tv on IRC, but it's still unclear to me what is really happening.

At first I saw OSD's which were eating 100% CPU and after killing them they would go into Zombie state:

```
root@atom4:~# ps aux|grep cosd
root      9928  99.9  0.0      0      0 ?        Zsl  Apr22  9279:33 [cosd] <defunct>
root     12061   0.0  0.0    7672   824 pts/0    S+   09:51   0:00 grep --color=auto cosd
root     26958  18.6  25.8 4014048 1048648 ?        Ssl  Apr22  1846:04 /usr/bin/cosd -i 16 -c /etc/ceph
/ceph.conf
root     27066  95.1  0.0      0      0 ?        Zsl  Apr22  9410:03 [cosd] <defunct>
root     27260  40.1  39.4 3528796 1599692 ?        Ssl  Apr22  3973:45 /usr/bin/cosd -i 19 -c /etc/ceph
/ceph.conf
root@atom4:~#
```

This is an example of such OSD's, those are still eating a lot of CPU, but Zombies are not supposed to do so, are they?

I've been trying to reproduce this, but I haven't been able to do so yet. The question was if those OSD's are still logging, but because my logging was too low, I had to restart all my machines to get everything up and running again.

Right now I'm seeing this:

```
osd e36982: 40 osds: 2 up, 3 in
```

But, I have 35 of my 40 OSD's 'up and running', but they are being marked down by the monitor.

Tailing the log (debug osd = 20) from a 'marked down' tells me:

```
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.c4( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] clear_recovery_state
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.c4( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] up [36,34,28] -> [36,28], acting [36,34,28] -> [36,28],
role -1 -> -1
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.c4( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] on_change
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.c5( v 26'2 (0'0,26'2) n=2 ec=2 les
=32821 36964/36964/36964) [28,19] r=-1 lcod 0'0 stray] unaffected with [28,19]/[28,19] up/acting
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[2.c1( v 24'12 (24'10,24'12) n=12 ec=
2 les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] noting past interval(36964-36964
[36,34]/[36,34])
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[2.c1( v 24'12 (24'10,24'12) n=12 ec=
```

```

2 les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] cancel_recovery
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[2.c1( v 24'12 (24'10,24'12) n=12 ec=
2 les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] clear_recovery_state
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[2.c1( v 24'12 (24'10,24'12) n=12 ec=
2 les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] up [36,34] -> [36], acting [36,34
] -> [36], role -1 -> -1
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[2.c1( v 24'12 (24'10,24'12) n=12 ec=
2 les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] on_change
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.c3( v 1298'329 (1205'327,1298'329)
n=307 ec=2 les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] noting past interval(36
964-36964 [36,34]/[36,34])
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.c3( v 1298'329 (1205'327,1298'329)
n=307 ec=2 les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] cancel_recovery
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.c3( v 1298'329 (1205'327,1298'329)
n=307 ec=2 les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] clear_recovery_state
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.c3( v 1298'329 (1205'327,1298'329)
n=307 ec=2 les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] up [36,34] -> [36], act
ing [36,34] -> [36], role -1 -> -1
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.c3( v 1298'329 (1205'327,1298'329)
n=307 ec=2 les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] on_change
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.c2( v 26'2 (0'0,26'2) n=2 ec=2 les
=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] noting past interval(36964-36964 [36,3
4]/[36,34])
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.c2( v 26'2 (0'0,26'2) n=2 ec=2 les
=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] cancel_recovery
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.c2( v 26'2 (0'0,26'2) n=2 ec=2 les
=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] clear_recovery_state
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.c2( v 26'2 (0'0,26'2) n=2 ec=2 les
=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] up [36,34] -> [36], acting [36,34] ->
[36], role -1 -> -1
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.c2( v 26'2 (0'0,26'2) n=2 ec=2 les
=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] on_change
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[3.c0( v 25'27 (25'25,25'27) n=7 ec=2
les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] noting past interval(36964-36964 [
36,34]/[36,34])
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[3.c0( v 25'27 (25'25,25'27) n=7 ec=2
les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] cancel_recovery
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[3.c0( v 25'27 (25'25,25'27) n=7 ec=2
les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] clear_recovery_state
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[3.c0( v 25'27 (25'25,25'27) n=7 ec=2
les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] up [36,34] -> [36], acting [36,34]
-> [36], role -1 -> -1
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[3.c0( v 25'27 (25'25,25'27) n=7 ec=2
les=22727 36965/36965/36964) [36] r=-1 lcod 0'0 crashed+stray] on_change
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.c3( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] noting past interval(36960-36964 [36,34,28]/[36,34,28])
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.c3( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] cancel_recovery
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.c3( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] clear_recovery_state
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.c3( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] up [36,34,28] -> [36,28], acting [36,34,28] -> [36,28],
role -1 -> -1
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.c3( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] on_change
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[3.c1( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] noting past interval(36960-36964 [36,34,28]/[36,34,28])
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[3.c1( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] cancel_recovery
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[3.c1( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] clear_recovery_state
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[3.c1( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] up [36,34,28] -> [36,28], acting [36,34,28] -> [36,28],
role -1 -> -1
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[3.c1( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] on_change
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[2.c2( empty n=0 ec=2 les=21620 36965

```

```

/36965/36960) [36,28] r=-1 crashed+stray] noting past interval(36960-36964 [36,34,28]/[36,34,28])
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[2.c2( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] cancel_recovery
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[2.c2( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] clear_recovery_state
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[2.c2( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] up [36,34,28] -> [36,28], acting [36,34,28] -> [36,28],
role -1 -> -1
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[2.c2( empty n=0 ec=2 les=21620 36965
/36965/36960) [36,28] r=-1 crashed+stray] on_change
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[3.c3( v 25'33 (25'29,25'33)+backlog
n=6 ec=2 les=32823 36964/36964/36964) [28,19] r=-1 lcod 0'0 stray] unaffected with [28,19]/[28,19]
up/acting
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.be( v 26'1 (0'0,26'1) n=1 ec=2 les
=35081 36964/36964/36964) [28,36] r=-1 lcod 0'0 crashed+stray] unaffected with [28,36]/[28,36] up/
acting
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[3.bc( v 25'49 (25'47,25'49) n=10 ec=
2 les=35081 36964/36964/36964) [28,36] r=-1 lcod 0'0 crashed+stray] unaffected with [28,36]/[28,36
] up/acting
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[2.bd( v 24'8 (24'6,24'8) n=8 ec=2 le
s=35081 36964/36964/36964) [28,36] r=-1 lcod 0'0 crashed+stray] unaffected with [28,36]/[28,36] up
/acting
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.bf( v 1205'286 (1205'284,1205'286)
n=277 ec=2 les=35081 36964/36964/36964) [28,36] r=-1 lcod 0'0 crashed+stray] unaffected with [28,
36]/[28,36] up/acting
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.bd( empty n=0 ec=2 les=21660 36965
/36965/36965) [19] r=-1 crashed+stray] noting past interval(36964-36964 [34,19]/[34,19])
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.bd( empty n=0 ec=2 les=21660 36965
/36965/36965) [19] r=-1 crashed+stray] cancel_recovery
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.bd( empty n=0 ec=2 les=21660 36965
/36965/36965) [19] r=-1 crashed+stray] clear_recovery_state
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.bd( empty n=0 ec=2 les=21660 36965
/36965/36965) [19] r=-1 crashed+stray] up [34,19] -> [19], acting [34,19] -> [19], role -1 -> -1
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.bd( empty n=0 ec=2 les=21660 36965
/36965/36965) [19] r=-1 crashed+stray] on_change
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[0.bd( empty n=0 ec=2 les=21660 36965
/36965/36965) [19] r=-1 crashed+stray] [34,19] -> [19], acting primary 34 -> 19
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.bc( empty n=0 ec=2 les=21660 36965
/36965/36965) [19] r=-1 crashed+stray] noting past interval(36964-36964 [34,19]/[34,19])
Apr 29 09:54:15 atom2 osd.9[2278]: 7f696c4e8700 osd9 36965 pg[1.bc( empty n=0 ec=2 les=21660 36965
/36965/36965) [19] r=-1 crashed+stray] cancel_recovery

```

Another OSD:

```

Apr 29 09:55:03 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 get_map 19448 - loading and decoding 0
x19d10c00
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a47a64700 osd8 36971 update_osd_stat osd_stat(642 GB used,
1217 GB avail, 1863 GB total, peers []/[19])
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a47a64700 osd8 36971 heartbeat: stat(2011-04-29 09:55:04.22
9086 oprate=0 qlen=0 recent_qlen=0 rdlat=0 / 0 fshedin=0)
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a47a64700 osd8 36971 heartbeat: osd_stat(642 GB used, 1217
GB avail, 1863 GB total, peers []/[19])
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a47a64700 -- [2a00:f10:113:1:225:90ff:fe33:49f2]:6808/21687
--> osd19 [2a00:f10:113:1:225:90ff:fe33:49cc]:6805/27259 -- osd_ping(e0 as_of 36971) v1 -- ?+0 0x
19c15000
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 add_map 19448 0x19d10c00
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior one of [9,5,
31] possibly crashed, marking pg crashed
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior interval(1915
6-19230 [9,31]/[9,31] maybe_went_rw)
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd9 u
p_from 36815 and last clean interval 36379-36800 does not include us

```

```
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd9 i
s down
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd31
up_from 36382 and last clean interval 35075-35182 does not include us
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd31
is down
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior 2 osds possi
bly went active+rw, no survivors, including
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 get_map 19230 - loading and decoding 0
x19d10900
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a47a64700 osd8 36971 update_osd_stat osd_stat(642 GB used,
1217 GB avail, 1863 GB total, peers []/[19])
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a47a64700 osd8 36971 heartbeat: stat(2011-04-29 09:55:04.93
0670 oprate=0 qlen=0 recent_qlen=0 rdlat=0 / 0 fshed=0)
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a47a64700 osd8 36971 heartbeat: osd_stat(642 GB used, 1217
GB avail, 1863 GB total, peers []/[19])
Apr 29 09:55:04 atom2 osd.8[21690]: 7f0a47a64700 -- [2a00:f10:113:1:225:90ff:fe33:49f2]:6808/21687
--> osd19 [2a00:f10:113:1:225:90ff:fe33:49cc]:6805/27259 -- osd_ping(e0 as_of 36971) v1 -- ?+0 0x
19c15540
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 add_map 19230 0x19d10900
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior one of [9,31
] possibly crashed, marking pg crashed
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior interval(1906
6-19155 [9,6,31]/[9,6,31] maybe_went_rw)
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd9 u
p_from 36815 and last clean interval 36379-36800 does not include us
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd9 i
s down
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd6 u
p_from 36813 and last clean interval 36376-36801 does not include us
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd6 i
s down
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd31
up_from 36382 and last clean interval 35075-35182 does not include us
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd31
is down
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior 3 osds possi
bly went active+rw, no survivors, including
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 get_map 19155 - loading and decoding 0
x19d10600
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 add_map 19155 0x19d10600
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior one of [9,6,
31] possibly crashed, marking pg crashed
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior interval(1899
0-19065 [9,31]/[9,31] maybe_went_rw)
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd9 u
p_from 36815 and last clean interval 36379-36800 does not include us
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd9 i
s down
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
```

```
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd31
up_from 36382 and last clean interval 35075-35182 does not include us
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior prior osd31
is down
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 pg[1.994( empty n=0 ec=2 les=16774 369
71/36971/36971) [8,36] r=0 mlcod 0'0 !hml crashed+down+degraded+peering] build_prior 2 osds possi
bly went active+rw, no survivors, including
Apr 29 09:55:05 atom2 osd.8[21690]: 7f0a4ba6c700 osd8 36971 get_map 19065 - loading and decoding 0
x19d10300
```

But both OSD's (8 and 9) are marked down by the monitor.

It could be that the Atom CPU's are not fast enough, but right now I'm seeing a average idle of 50% on all machines, to me it seems the CPU is currently not the bottleneck.

My monitor is also running with high debugging, in this logs I'm seeing:

```
Apr 29 10:00:27 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
<== osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 232 ===== auth(proto 2 2 bytes) v1 ===== 28
+0+0 (1679083706 0 0) 0x1852a00 con 0x1c4b780
Apr 29 10:00:27 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
<== osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 233 ===== mon_subscribe({monmap=2+,osdmap=
36972}) v2 ===== 42+0+0 (2521356415 0 0) 0x1f74380 con 0x1c4b780
Apr 29 10:00:27 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
--> osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 -- osd_map(36972,36985) v1 -- ?+0 0x1852a
00
Apr 29 10:00:27 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
--> osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 -- mon_subscribe_ack(300s) v1 -- ?+0 0x34
b3480
Apr 29 10:00:30 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
<== osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 234 ===== mon_subscribe({monmap=2+,osdmap=
36972}) v2 ===== 42+0+0 (2521356415 0 0) 0x1f74700 con 0x1c4b780
Apr 29 10:00:30 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
--> osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 -- osd_map(36972,36985) v1 -- ?+0 0x30bb4
00
Apr 29 10:00:30 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
--> osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 -- mon_subscribe_ack(300s) v1 -- ?+0 0x28
76900
Apr 29 10:00:37 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
<== osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 235 ===== auth(proto 2 2 bytes) v1 ===== 28
+0+0 (1679083706 0 0) 0x18fd000 con 0x1c4b780
Apr 29 10:00:37 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
<== osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 236 ===== mon_subscribe({monmap=2+,osdmap=
36972}) v2 ===== 42+0+0 (2521356415 0 0) 0x2053e00 con 0x1c4b780
Apr 29 10:00:37 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
--> osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 -- osd_map(36972,36985) v1 -- ?+0 0x18fd0
00
Apr 29 10:00:37 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
--> osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 -- mon_subscribe_ack(300s) v1 -- ?+0 0x1c
df000
Apr 29 10:00:48 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
<== osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 237 ===== auth(proto 2 2 bytes) v1 ===== 28
+0+0 (1679083706 0 0) 0x18fd800 con 0x1c4b780
Apr 29 10:00:48 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
<== osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 238 ===== mon_subscribe({monmap=2+,osdmap=
36972}) v2 ===== 42+0+0 (2521356415 0 0) 0x2053380 con 0x1c4b780
Apr 29 10:00:48 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
--> osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 -- osd_map(36972,36985) v1 -- ?+0 0x18fd8
00
Apr 29 10:00:48 monitor mon.dga[30655]: 7f9ab0aa3700 -- [2a00:f10:113:1:230:48ff:fed3:b086]:6789/0
--> osd8 [2a00:f10:113:1:225:90ff:fe33:49f2]:6806/21687 -- mon_subscribe_ack(300s) v1 -- ?+0 0x19
10000
```

The OSD still seems to work and tries to subscribe with the monitor.

This is what I'm seeing with all my OSD's in the cluster, although they are running and seem to be alive, they are marked down and out.

Sometimes I get a message that a OSD is going into "boot state", although that OSD has been running for a long time, about 24 hours atm.

My osdmap shows:

```
epoch 36987
fsid 109c5e61-13e5-2ad2-cd6c-79f2919901e9
created 2011-03-31 10:39:37.629575
modified 2011-04-29 10:02:59.567400
flags

pg_pool 0 'data' pg_pool(rep pg_size 3 crush_ruleset 0 object_hash rjenkins pg_num 2560 pgp_num 25
60 lpg_num 2 lpgp_num 2 last_change 14 owner 0)
pg_pool 1 'metadata' pg_pool(rep pg_size 3 crush_ruleset 1 object_hash rjenkins pg_num 2560 pgp_nu
m 2560 lpg_num 2 lpgp_num 2 last_change 11 owner 0)
pg_pool 2 'casdata' pg_pool(rep pg_size 3 crush_ruleset 2 object_hash rjenkins pg_num 2560 pgp_num
2560 lpg_num 2 lpgp_num 2 last_change 22 owner 0)
pg_pool 3 'rbd' pg_pool(rep pg_size 3 crush_ruleset 3 object_hash rjenkins pg_num 2560 pgp_num 256
0 lpg_num 2 lpgp_num 2 last_change 18 owner 0)
pg_pool 4 '.rgw' pg_pool(rep pg_size 3 crush_ruleset 0 object_hash rjenkins pg_num 8 pgp_num 8 lpg
_num 0 lpgp_num 0 last_change 10 owner 0)
pg_pool 5 '.users' pg_pool(rep pg_size 3 crush_ruleset 0 object_hash rjenkins pg_num 8 pgp_num 8 l
pg_num 0 lpgp_num 0 last_change 20 owner 0)
pg_pool 6 '.users.email' pg_pool(rep pg_size 3 crush_ruleset 0 object_hash rjenkins pg_num 8 pgp_n
um 8 lpg_num 0 lpgp_num 0 last_change 21 owner 0)
pg_pool 7 '.log' pg_pool(rep pg_size 3 crush_ruleset 0 object_hash rjenkins pg_num 8 pgp_num 8 lpg
_num 0 lpgp_num 0 last_change 25 owner 0)
pg_pool 10 'southpark' pg_pool(rep pg_size 3 crush_ruleset 0 object_hash rjenkins pg_num 8 pgp_num
8 lpg_num 0 lpgp_num 0 last_change 28 owner 0)
pg_pool 11 'thesimpsons' pg_pool(rep pg_size 3 crush_ruleset 0 object_hash rjenkins pg_num 8 pgp_n
um 8 lpg_num 0 lpgp_num 0 last_change 31 owner 0)

max_osd 40
osd0 down out up_from 36383 up_thru 35925 down_at 36487 last_clean_interval 34283-35925
osd1 down out up_from 36942 up_thru 35466 down_at 36948 last_clean_interval 36320-36854
osd2 down out up_from 36802 up_thru 35457 down_at 36823 last_clean_interval 36322-36749
osd3 down out up_from 36892 up_thru 35280 down_at 36899 last_clean_interval 36315-36879
osd4 down out up_from 36392 up_thru 35256 down_at 36510 last_clean_interval 34291-35671
osd5 down out up_from 36961 up_thru 35303 down_at 36964 last_clean_interval 36320-36604
osd6 down out up_from 36813 up_thru 34554 down_at 36828 last_clean_interval 36376-36801
osd7 down out up_from 36957 up_thru 23861 down_at 36964 last_clean_interval 36320-36936
osd8 down out up_from 36971 up_thru 36885 down_at 36972 last_clean_interval 36333-36885
osd9 down out up_from 36981 up_thru 35455 down_at 36982 last_clean_interval 36379-36956
osd10 down out up_from 36873 up_thru 36849 down_at 36875 last_clean_interval 36380-36837
osd11 down out up_from 36392 up_thru 24147 down_at 36421 last_clean_interval 23809-24148
osd12 down out up_from 36848 up_thru 35204 down_at 36881 last_clean_interval 36378-36824
osd13 down out up_from 36392 up_thru 33189 down_at 36424 last_clean_interval 23811-33494
osd14 down out up_from 36902 up_thru 35317 down_at 36903 last_clean_interval 36329-36819
osd15 down out up_from 36921 up_thru 36454 down_at 36922 last_clean_interval 36339-36915
osd16 down out up_from 35460 up_thru 36912 down_at 36918 last_clean_interval 23796-35452
osd17 down out up_from 24387 up_thru 29291 down_at 29453 last_clean_interval 23817-24236
osd18 down out up_from 29318 up_thru 24143 down_at 29319 last_clean_interval 23802-24145
osd19 up in weight 1 up_from 36955 up_thru 36983 down_at 36951 last_clean_interval 23828-36952
[2a00:f10:113:1:225:90ff:fe33:49cc]:6809/27259 [2a00:f10:113:1:225:90ff:fe33:49cc]:6804/27259 [2a0
0:f10:113:1:225:90ff:fe33:49cc]:6805/27259
osd20 down out up_from 36958 up_thru 36039 down_at 36964 last_clean_interval 36377-36956
osd21 up in weight 1 up_from 36987 up_thru 36030 down_at 36966 last_clean_interval 36376-36982
[2a00:f10:113:1:225:90ff:fe33:497c]:6803/1804 [2a00:f10:113:1:225:90ff:fe33:497c]:6807/1804 [2a00:
f10:113:1:225:90ff:fe33:497c]:6808/1804
osd22 down out up_from 36849 up_thru 35456 down_at 36852 last_clean_interval 36324-36697
osd23 down out up_from 36384 up_thru 33408 down_at 36396 last_clean_interval 29317-33753
osd24 down out up_from 36956 up_thru 35662 down_at 36964 last_clean_interval 36378-36827
```

```
osd25 down out up_from 36809 up_thru 36812 down_at 36825 last_clean_interval 36378-36781
osd26 down out up_from 36772 up_thru 35454 down_at 36775 last_clean_interval 36327-36728
osd27 down out up_from 36787 up_thru 36032 down_at 36792 last_clean_interval 36341-36776
osd28 down out up_from 36954 up_thru 23310 down_at 36967 last_clean_interval 36350-36829
osd29 down out up_from 36830 up_thru 35249 down_at 36860 last_clean_interval 36345-36813
osd30 down out up_from 36577 up_thru 23356 down_at 36598 last_clean_interval 36336-36513
osd31 down out up_from 36382 up_thru 23323 down_at 36400 last_clean_interval 35075-35182
osd32 down out up_from 36931 up_thru 36717 down_at 36932 last_clean_interval 36377-36857
osd33 down out up_from 36924 up_thru 36864 down_at 36925 last_clean_interval 36377-36915
osd34 down out up_from 36953 up_thru 29565 down_at 36965 last_clean_interval 36348-36838
osd35 down out up_from 5615 up_thru 6169 down_at 6239 last_clean_interval 4764-4910
osd36 down out up_from 36960 up_thru 36737 down_at 36978 last_clean_interval 36351-36889
osd37 down out up_from 36974 up_thru 36911 down_at 36983 last_clean_interval 36378-36961
osd38 down out up_from 5580 up_thru 6396 down_at 6441 last_clean_interval 4809-5562
osd39 down out up_from 36378 up_thru 36409 down_at 36598 last_clean_interval 29335-36299
```

```
pg_temp 0.29b [16]
pg_temp 0.461 [16]
pg_temp 0.75c [16]
pg_temp 0.95b [16]
pg_temp 1.29a [16]
pg_temp 1.460 [16]
pg_temp 1.75b [16]
pg_temp 2.299 [16]
pg_temp 2.45f [16]
pg_temp 2.75a [16]
pg_temp 2.959 [16]
pg_temp 3.298 [16]
pg_temp 3.45e [16]
pg_temp 3.759 [16]
pg_temp 3.958 [16]
```

And my OSD processes are running:

```
root@monitor:~# dsh -g osd-mdb "pidof cosd|wc -w"
4
3
4
4
4
4
4
4
4
4
3
3
root@monitor:~#
```

History

#1 - 06/03/2011 10:23 AM - Sage Weil

- Target version set to v0.30

Wanted to check in on this one. Are you still seeing this problem? When the processes are zombies, are there any btrfs messages in dmesg? (Usually it's caused by a syscall that's blocked in the kernel and the process can't be cleaned up until it returns...)

#2 - 06/06/2011 04:35 AM - Wido den Hollander

I've been away for some time, so haven't got a chance to hunt this one further down. I'm back now, so I'll see if I can reproduce it.

As far as I remember, there were some btrfs messages in the dmesg, not sure which one though.

#3 - 06/08/2011 07:25 AM - Wido den Hollander

I'm just upgrading my cluster, I see a Zombie from May 26th, running v0.28.1 and 2.6.38.2.

I'll see if I can reproduce it with the latest master.

#4 - 06/20/2011 12:06 PM - Sage Weil

- Target version changed from v0.30 to v0.31

#5 - 06/21/2011 01:01 PM - Sage Weil

- Status changed from New to Closed

I'm going to close this one. if you see a zombie with a btrfs error in dmesg, it's a btrfs problem. if you don't, let's reopen!

#6 - 06/21/2011 08:40 PM - Sage Weil

- Target version changed from v0.31 to v0.32

#7 - 06/29/2011 04:05 AM - Wido den Hollander

I'm seeing it again. Trying to upgrade from v0.29.1 to v0.30.

My cluster was bouncing all over again, "wrongly marked me down" was everywhere, so I thought I'd give v0.30 a try.

The upgrade is pretty straight forward:

```
dsh -g osd-mdb "service ceph stop osd; sleep 10; dpkg -i /tmp/*.deb; service ceph start osd"
```

On for example atom0 this caused:

```
root@atom0:~# ps aux|grep cosd
root    13346 13.6  0.0    0     0 ?        Zsl   Jun23 1185:08 [cosd] <defunct>
root    15570 11.5  0.0    0     0 ?        Zsl   Jun23 1007:46 [cosd] <defunct>
root    26095  0.0  0.0   7676  824 pts/0    S+   13:02   0:00 grep --color=auto cosd
root@atom0:~#
```

Top shows them eating about a 100% CPU.

No btrfs messages, but what I do see is that two "btrfs-cleaner" cleaner processes are pretty busy, could it be that the OSD is still waiting for some stuff from btrfs? I guess so?

#8 - 06/29/2011 09:09 AM - Sage Weil

- Status changed from *Closed* to *Feedback*
- Target version changed from *v0.32* to *v0.31*

If you can reproduce this, the next step is to figure out what the kernel threads are doing. Do several rounds of

```
echo t > /proc/sysrq-trigger
```

every few seconds and then capture the kernel log output and attach here.

'w' would dump only threads in uninterruptible state (which I think should capture the cosd threads), but 't' will dump everything and tell us what btrfs-cleaner is doing that is blocking the cosd threads.

Thanks!

#9 - 07/05/2011 09:36 AM - Sage Weil

- Target version changed from *v0.31* to *v0.32*

#10 - 07/23/2011 08:36 PM - Sage Weil

- Target version changed from *v0.32* to *v0.33*

#11 - 08/08/2011 09:18 AM - Sage Weil

- Target version changed from *v0.33* to *v0.34*

#12 - 08/08/2011 11:45 AM - Sage Weil

- translation missing: *en.field_position* set to 1
- translation missing: *en.field_position* changed from 1 to 808

#13 - 08/20/2011 11:08 AM - Sage Weil

- Target version changed from *v0.34* to *v0.35*

#14 - 09/06/2011 09:49 PM - Sage Weil

- Target version changed from *v0.35* to *v0.36*

#15 - 09/07/2011 12:07 PM - Sage Weil

- translation missing: *en.field_position* deleted (860)
- translation missing: *en.field_position* set to 905

#16 - 09/25/2011 02:14 PM - Sage Weil

- Target version changed from *v0.36* to *v0.37*

#17 - 10/09/2011 08:39 PM - Sage Weil

- Target version changed from *v0.37* to *v0.38*

#18 - 10/09/2011 08:42 PM - Sage Weil

- Status changed from *Feedback* to *Can't reproduce*

#19 - 10/09/2011 08:42 PM - Sage Weil

- Target version changed from *v0.38* to *v0.37*

- translation missing: *en.field_position* deleted (938)

- translation missing: *en.field_position* set to 1

- translation missing: *en.field_position* changed from 1 to 943