

Ceph - Bug #1724

timeout during tiobench test

11/14/2011 04:03 PM - Anonymous

Status:	Resolved	% Done:	0%
Priority:	High	Spent time:	0.00 hour
Assignee:			
Category:			
Target version:	v0.39		
Source:		Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:		Pull request ID:	
Regression:	No	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			

Description

During the 11/11 nightlies, the tiotest task blocked multiple times. The first stack trace (from 1831/remote/ubuntu@sepia36.ceph.dreamhost.com/syslog/kern.log.gz) was

```
2011-11-11T02:34:33.345971-08:00 sepia36 kernel: [ 116.383388] libceph: client4101 fsid 39f48706-ca58-c4c1-6c90-929a3368fd5a
2011-11-11T02:34:33.345997-08:00 sepia36 kernel: [ 116.383515] libceph: mon1 10.3.14.131:6789 session established
2011-11-11T02:35:58.115977-08:00 sepia36 kernel: [ 201.150114] libceph: tid 409 timed out on osd0, will reset osd
2011-11-11T02:35:58.116009-08:00 sepia36 kernel: [ 201.150462] libceph: tid 699 timed out on osd1, will reset osd
2011-11-11T02:37:03.245972-08:00 sepia36 kernel: [ 266.280151] libceph: tid 4755 timed out on osd1, will reset osd
2011-11-11T02:37:03.246004-08:00 sepia36 kernel: [ 266.280259] libceph: tid 3785 timed out on osd0, will reset osd
2011-11-11T02:38:03.365968-08:00 sepia36 kernel: [ 326.400211] libceph: tid 3944 timed out on osd0, will reset osd
2011-11-11T02:38:03.365994-08:00 sepia36 kernel: [ 326.400373] libceph: tid 6799 timed out on osd1, will reset osd
2011-11-11T02:40:09.485971-08:00 sepia36 kernel: [ 452.520022] libceph: tid 7792 timed out on osd1, will reset osd
2011-11-11T02:40:09.486001-08:00 sepia36 kernel: [ 452.520325] libceph: tid 8015 timed out on osd0, will reset osd
2011-11-11T02:41:09.605962-08:00 sepia36 kernel: [ 512.640022] libceph: tid 8015 timed out on osd0, will reset osd
2011-11-11T02:41:14.615965-08:00 sepia36 kernel: [ 517.650023] libceph: tid 11879 timed out on osd1, will reset osd
2011-11-11T02:42:09.725961-08:00 sepia36 kernel: [ 572.760021] libceph: tid 11021 timed out on osd0, will reset osd
2011-11-11T02:42:14.735957-08:00 sepia36 kernel: [ 577.770022] libceph: tid 13292 timed out on osd1, will reset osd
2011-11-11T02:42:40.040498-08:00 sepia36 kernel: [ 603.070030] INFO: task tiotest:1277 blocked for more than 120 seconds.
2011-11-11T02:42:40.040664-08:00 sepia36 kernel: [ 603.076581] "echo 0 > /proc/sys/kernel/hung_task_timeout_secs" disables
this message.
2011-11-11T02:42:40.040724-08:00 sepia36 kernel: [ 603.084480] tiotest      D 0000000000000002  0 1277 1210 0x00000000
2011-11-11T02:42:40.040995-08:00 sepia36 kernel: [ 603.084552] ffff8800df91fc58 0000000000000046 0000000000000000
ffff880036e0e860
2011-11-11T02:42:40.041075-08:00 sepia36 kernel: [ 603.084575] ffff8800f2533f60 0000000000013b80 ffff8800df91ffd8
ffff8800df91e010
2011-11-11T02:42:40.041181-08:00 sepia36 kernel: [ 603.084580] ffff8800df91ffd8 0000000000013b80 ffff880036c7df10
ffff8800f2533f60
2011-11-11T02:42:40.041258-08:00 sepia36 kernel: [ 603.084586] Call Trace:
2011-11-11T02:42:40.041335-08:00 sepia36 kernel: [ 603.084595] [<fffffff81091cdd>] ? ktime_get_ts+0xad/0xe0
2011-11-11T02:42:40.041410-08:00 sepia36 kernel: [ 603.084601] [<fffffff8111b7e0>] ? __lock_page+0x70/0x70
2011-11-11T02:42:40.041486-08:00 sepia36 kernel: [ 603.084606] [<fffffff81600e0f>] schedule+0x3f/0x60
2011-11-11T02:42:40.041564-08:00 sepia36 kernel: [ 603.084610] [<fffffff81600ebc>] io_schedule+0x8c/0xd0
2011-11-11T02:42:40.041640-08:00 sepia36 kernel: [ 603.084613] [<fffffff8111b7ee>] sleep_on_page+0xe/0x20
2011-11-11T02:42:40.041716-08:00 sepia36 kernel: [ 603.084617] [<fffffff816015ff>] __wait_on_bit+0x5f/0x90
2011-11-11T02:42:40.041792-08:00 sepia36 kernel: [ 603.084621] [<fffffff8111ba33>] wait_on_page_bit+0x73/0x80
2011-11-11T02:42:40.041868-08:00 sepia36 kernel: [ 603.084626] [<fffffff810875c0>] ? autoremove_wake_function+0x40/0x40
2011-11-11T02:42:40.041944-08:00 sepia36 kernel: [ 603.084631] [<fffffff81127125>] ? pagevec_lookup_tag+0x25/0x40
2011-11-11T02:42:40.042021-08:00 sepia36 kernel: [ 603.084635] [<fffffff8111be93>] filemap_fdatawait_range+0x113/0x1a0
2011-11-11T02:42:40.042109-08:00 sepia36 kernel: [ 603.084640] [<fffffff8111c020>] filemap_write_and_wait_range+0x70/0x80
2011-11-11T02:42:40.042194-08:00 sepia36 kernel: [ 603.084665] [<ffffffa0209891>] ceph_fsycn+0x111/0x2f0 [ceph]
2011-11-11T02:42:40.042271-08:00 sepia36 kernel: [ 603.084672] [<fffffff811a2e7b>] vfs_fsycn_range+0x2b/0x40
2011-11-11T02:42:40.042343-08:00 sepia36 kernel: [ 603.084675] [<fffffff811a2eac>] vfs_fsycn+0x1c/0x20
```

```
2011-11-11T02:42:40.042416-08:00 sepia36 kernel: [ 603.084679] [<ffffff811a30ea>] do_fsync+0x3a/0x60
2011-11-11T02:42:40.042490-08:00 sepia36 kernel: [ 603.084683] [<ffffff811a3140>] sys_fsync+0x10/0x20
2011-11-11T02:42:40.042563-08:00 sepia36 kernel: [ 603.084687] [<ffffff8160bb42>] system_call_fastpath+0x16/0x1b
```

Note that although the trace is different from the one in bug 1827, this could (for all I know) simply be another instance of the same delay/loop merely caught at a different instant.

History

#1 - 11/14/2011 04:04 PM - Anonymous

(I should have said the other problem was filed as bug 1723)

#2 - 11/17/2011 10:49 AM - Sage Weil

- *Project changed from Linux kernel client to Ceph*

#3 - 11/17/2011 10:49 AM - Sage Weil

- *Target version set to v0.39*

#4 - 11/17/2011 10:52 AM - Sage Weil

- *Priority changed from Normal to High*

#5 - 11/17/2011 10:54 AM - Sage Weil

- *translation missing: en.field_position set to 8*

#6 - 11/17/2011 04:33 PM - Sage Weil

- *Status changed from New to Resolved*

this test ran commit:dfc3ddc8983fbc7c376394067335b360c68cd314, which did not include the root dentry fix in commit:774ac21da76f5c3018428725074e27a3fd40b128.