

## Ceph - Bug #12652

osd/PGLog.cc: 732: FAILED assert(log.log.size() == log\_keys\_debug.size())

08/09/2015 02:45 PM - Sage Weil

<b>Status:</b>	Resolved	<b>Start date:</b>	08/09/2015
<b>Priority:</b>	Urgent	<b>Due date:</b>	
<b>Assignee:</b>		<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>		<b>Spent time:</b>	0.00 hour
<b>Source:</b>	Q/A	<b>Reviewed:</b>	
<b>Tags:</b>		<b>Affected Versions:</b>	
<b>Backport:</b>	hammer,firely	<b>ceph-qa-suite:</b>	
<b>Regression:</b>	No	<b>Pull request ID:</b>	
<b>Severity:</b>	3 - minor		

### Description

```
2015-08-08 12:39:07.732502 7f27206b4700 -1 osd/PGLog.cc: In function 'void PGLog::check()' thread
7f27206b4700 time 2015-08-08 12:39:07.583714
osd/PGLog.cc: 732: FAILED assert(log.log.size() == log_keys_debug.size())
```

```
ceph version 9.0.2-1176-gc2fc74d (c2fc74d18a5e65b4a5b30d760550935a174613e8)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x8b) [0xc3d32b]
2: (PGLog::check()+0x67d) [0x7a1bed]
3: (PGLog::write_log(ObjectStore::Transaction&, std::map<std::string, ceph::buffer::list, std::less<std::string>, std::allocator<std::pair<std::string const, ceph::buffer::list> > >*, coll_t const&, ghobject_t const&)+0x1de) [0x7a7fde]
4: (PG::write_if_dirty(ObjectStore::Transaction&)+0x7a) [0x80975a]
5: (OSD::process_peering_events(std::list<PG*, std::allocator<PG*> > const&, ThreadPool::TPHandle&)+0x2f0) [0x6df820]
6: (OSD::PeeringWQ::_process(std::list<PG*, std::allocator<PG*> > const&, ThreadPool::TPHandle&)+0x18) [0x7343f8]
7: (ThreadPool::worker(ThreadPool::WorkThread*)+0xa56) [0xc2df16]
8: (ThreadPool::WorkThread::entry()+0x10) [0xc2efc0]
9: (()+0x8182) [0x7f273b3b8182]
10: (clone()+0x6d) [0x7f27396fcfbfd]
NOTE: a copy of the executable, or `objdump -rdS <executable>` is needed to interpret this.
```

/a/sage-2015-08-08\_07:57:46-rados:thrash-wip-newstore-sort-distro-basic-multi/1006267

just prior,

```
2015-08-08 12:39:07.583516 7f27206b4700 5 write_log with: dirty_to: 956'202, dirty_from: 42949672
95'18446744073709551615, dirty_divergent_priors: false, divergent_priors: 0, writeout_from: 429496
7295'18446744073709551615, trimmed:
2015-08-08 12:39:07.583647 7f27206b4700 -1 log.log.size() != log_keys_debug.size()
2015-08-08 12:39:07.583649 7f27206b4700 -1 actual log:
2015-08-08 12:39:07.583650 7f27206b4700 -1 6'46 (0'0) modify 1/2ffc3b7c/mira03825864-195/hea
d by client.4122.0:976 2015-08-08 11:37:30.479123
2015-08-08 12:39:07.583658 7f27206b4700 -1 6'55 (6'46) modify 1/2ffc3b7c/mira03825864-195/he
ad by client.4122.0:977 2015-08-08 11:37:30.503925
2015-08-08 12:39:07.583662 7f27206b4700 -1 6'56 (6'55) modify 1/2ffc3b7c/mira03825864-195/he
ad by client.4122.0:978 2015-08-08 11:37:30.521613
2015-08-08 12:39:07.583666 7f27206b4700 -1 6'57 (6'56) modify 1/2ffc3b7c/mira03825864-195/he
ad by client.4122.0:979 2015-08-08 11:37:30.521663
2015-08-08 12:39:07.583670 7f27206b4700 -1 298'191 (6'57) modify 1/2ffc3b7c/mira03825864-195
```

```

/head by client.4122.0:5402 2015-08-08 11:57:21.552804
2015-08-08 12:39:07.583675 7f27206b4700 -1      298'192 (298'191) modify    1/2ffc3b7c/mira03825864-
195/head by client.4122.0:5403 2015-08-08 11:57:21.573882
2015-08-08 12:39:07.583679 7f27206b4700 -1      298'193 (298'192) modify    1/2ffc3b7c/mira03825864-
195/head by client.4122.0:5404 2015-08-08 11:57:21.573944
2015-08-08 12:39:07.583685 7f27206b4700 -1      956'199 (298'193) modify    1/2ffc3b7c/mira03825864-
195/head by client.4122.0:9105 2015-08-08 12:26:38.508262
2015-08-08 12:39:07.583689 7f27206b4700 -1      956'200 (956'199) modify    1/2ffc3b7c/mira03825864-
195/head by client.4122.0:9106 2015-08-08 12:26:38.535341
2015-08-08 12:39:07.583693 7f27206b4700 -1      956'201 (956'200) modify    1/2ffc3b7c/mira03825864-
195/head by client.4122.0:9107 2015-08-08 12:26:38.537499
2015-08-08 12:39:07.583698 7f27206b4700 -1      956'202 (956'201) modify    1/2ffc3b7c/mira03825864-
195/head by client.4122.0:9108 2015-08-08 12:26:38.537572
2015-08-08 12:39:07.583702 7f27206b4700 -1 log_keys_debug:
2015-08-08 12:39:07.583703 7f27206b4700 -1      0000000006.000000000000000000046
2015-08-08 12:39:07.583704 7f27206b4700 -1      0000000006.000000000000000000055
2015-08-08 12:39:07.583705 7f27206b4700 -1      0000000006.000000000000000000056
2015-08-08 12:39:07.583706 7f27206b4700 -1      0000000006.000000000000000000057
2015-08-08 12:39:07.583707 7f27206b4700 -1      0000000298.000000000000000000191
2015-08-08 12:39:07.583707 7f27206b4700 -1      0000000298.000000000000000000192
2015-08-08 12:39:07.583708 7f27206b4700 -1      0000000298.000000000000000000193
2015-08-08 12:39:07.583709 7f27206b4700 -1      0000000956.000000000000000000199
2015-08-08 12:39:07.583710 7f27206b4700 -1      0000000956.000000000000000000200
2015-08-08 12:39:07.583711 7f27206b4700 -1      0000000956.000000000000000000201

```

**Related issues:**

Related to Ceph - Bug #10718: osd/PGLog.h: 279: FAILED assert(log.log.size() ...	<b>Resolved</b>	<b>02/02/2015</b>
Copied to Ceph - Backport #12845: osd/PGLog.cc: 732: FAILED assert(log.log.si...	<b>Resolved</b>	<b>08/09/2015</b>
Copied to Ceph - Backport #12846: osd/PGLog.cc: 732: FAILED assert(log.log.si...	<b>Resolved</b>	<b>08/09/2015</b>

**Associated revisions**

**Revision f0ca14df - 08/19/2015 09:13 PM - Sage Weil**

osd/PGLog: dirty\_to is inclusive

There are only two callers of mark\_dirty\_to who do not pass max, and they are both in the merge\_log extending tail path. In that case, we want to include the last version specified in the log writeout. Fix the tail extending code to always specify the last entry added, inclusive.

Fixes: #12652  
Signed-off-by: Sage Weil <[sage@redhat.com](mailto:sage@redhat.com)>

**Revision 77624aff - 09/02/2015 07:26 PM - Sage Weil**

osd/PGLog: dirty\_to is inclusive

There are only two callers of mark\_dirty\_to who do not pass max, and they are both in the merge\_log extending tail path. In that case, we want to include the last version specified in the log writeout. Fix the tail extending code to always specify the last entry added, inclusive.

Fixes: #12652  
Signed-off-by: Sage Weil <[sage@redhat.com](mailto:sage@redhat.com)>  
(cherry picked from commit f0ca14df0641daa04eee39d98d8bd0faf46e4e6d)

osd/PGLog: dirty\_to is inclusive

There are only two callers of mark\_dirty\_to who do not pass max, and they are both in the merge\_log extending tail path. In that case, we want to include the last version specified in the log writeout. Fix the tail extending code to always specify the last entry added, inclusive.

Fixes: #12652

Signed-off-by: Sage Weil <sage@redhat.com>

(cherry picked from commit f0ca14df0641daa04eee39d98d8bd0faf46e4e6d)

## History

### #1 - 08/09/2015 02:46 PM - Sage Weil

merge\_log, just prior:

```

2015-08-08 12:39:07.582990 7f27206b4700 10 osd.5 pg_epoch: 1296 pg[1.17c( v 956'202 (956'202,956'202] local-les=1224 n=1 ec=5 les/c 1224/1224 1295/1295/1223) [5,4] r=0 lpr=1295 pi=1223-1294/1 crt=956'202 mlcod 0'0 peerin
g NIBBLEWISE] proc_master_log for osd.1: log((0'0,956'202], crt=956'202) missing(0)
2015-08-08 12:39:07.583002 7f27206b4700 10 merge_log log((0'0,956'202], crt=956'202) from osd.1 into log((956'
202,956'202], crt=956'202)
2015-08-08 12:39:07.583005 7f27206b4700 10 merge_log extending tail to 0'0
2015-08-08 12:39:07.583008 7f27206b4700 15 6'46 (0'0) modify 1/2ffc3b7c/mira03825864-195/head by client.4122
.0:976 2015-08-08 11:37:30.479123
2015-08-08 12:39:07.583016 7f27206b4700 15 6'55 (6'46) modify 1/2ffc3b7c/mira03825864-195/head by client.412
2.0:977 2015-08-08 11:37:30.503925
2015-08-08 12:39:07.583021 7f27206b4700 15 6'56 (6'55) modify 1/2ffc3b7c/mira03825864-195/head by client.412
2.0:978 2015-08-08 11:37:30.521613
2015-08-08 12:39:07.583026 7f27206b4700 15 6'57 (6'56) modify 1/2ffc3b7c/mira03825864-195/head by client.412
2.0:979 2015-08-08 11:37:30.521663
2015-08-08 12:39:07.583031 7f27206b4700 15 298'191 (6'57) modify 1/2ffc3b7c/mira03825864-195/head by client.
4122.0:5402 2015-08-08 11:57:21.552804
2015-08-08 12:39:07.583036 7f27206b4700 15 298'192 (298'191) modify 1/2ffc3b7c/mira03825864-195/head by clie
nt.4122.0:5403 2015-08-08 11:57:21.573882
2015-08-08 12:39:07.583041 7f27206b4700 15 298'193 (298'192) modify 1/2ffc3b7c/mira03825864-195/head by clie
nt.4122.0:5404 2015-08-08 11:57:21.573944
2015-08-08 12:39:07.583046 7f27206b4700 15 956'199 (298'193) modify 1/2ffc3b7c/mira03825864-195/head by clie
nt.4122.0:9105 2015-08-08 12:26:38.508262
2015-08-08 12:39:07.583051 7f27206b4700 15 956'200 (956'199) modify 1/2ffc3b7c/mira03825864-195/head by clie
nt.4122.0:9106 2015-08-08 12:26:38.535341
2015-08-08 12:39:07.583056 7f27206b4700 15 956'201 (956'200) modify 1/2ffc3b7c/mira03825864-195/head by clie
nt.4122.0:9107 2015-08-08 12:26:38.537499
2015-08-08 12:39:07.583061 7f27206b4700 15 956'202 (956'201) modify 1/2ffc3b7c/mira03825864-195/head by clie
nt.4122.0:9108 2015-08-08 12:26:38.537572
2015-08-08 12:39:07.583067 7f27206b4700 10 merge_log result log((0'0,956'202], crt=956'202) missing(0) changed
=1

```

it does mark\_dirty\_to(956'202)

but writeout does not include the 202 record.

**#2 - 08/09/2015 02:46 PM - Sage Weil**

- *Description updated*

**#3 - 08/09/2015 02:47 PM - Sage Weil**

- *Backport set to hammer,firely*

**#4 - 08/09/2015 02:49 PM - Sage Weil**

- *Status changed from New to Need Review*

<https://github.com/ceph/ceph/pull/5518>

**#5 - 08/20/2015 06:53 PM - Sage Weil**

- *Status changed from Need Review to Pending Backport*

**#6 - 10/20/2015 07:41 PM - Loic Dachary**

- *Status changed from Pending Backport to Resolved*