

## crimson - Bug #55550

### crimson: check\_past\_interval\_bounds() assert failure

05/04/2022 11:23 PM - Samuel Just

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

#### Description

Likely insufficient information to find the bug, but it's pretty reproducible. Killing and restarting an osd with IO running seems to result in this assert on startup during peering. At a guess, we're not recording the past\_intervals during activation correctly.

```
DEBUG 2022-05-04 23:20:57,583 [shard 0] osd - pg_epoch 17 pg[1.1( v 16'191 (0'0,16'191) local-lis /les=12/13 n=2 ec=9/9 lis/c=12/0 les/c/f=13/0/0 sis=12) [2,1] r=-1 lpr=17 pi=[12,17)/1 crt=16'191 lcod 0'0 mlcod 0'0 unknown noting past ([1 2,16] all_participants=0,1,2 intervals=([12,16] acting 0,1,2))
```

```
DEBUG 2022-05-04 23:20:57,583 [shard 0] osd - pg_epoch 17 pg[1.1( v 16'191 (0'0,16'191) local-lis /les=12/13 n=2 ec=9/9 lis/c=12/0 les/c/f=13/0/0 sis=17) [2,1] r=-1 lpr=17 pi=[12,17)/1 crt=16'191 lcod 0'0 mlcod 0'0 unknown on_new_interval  
DEBUG 2022-05-04 23:20:57,583 [shard 0] osd - pg_epoch 17 pg[1.1( v 16'191 (0'0,16'191) local-lis /les=12/13 n=2 ec=9/9 lis/c=12/0 les/c/f=13/0/0 sis=17) [2,1] r=-1 lpr=17 pi=[12,17)/1 crt=16'191 lcod 0'0 mlcod 0'0 unknown on_new_interval  
upacting_features 0x3f01cfbb7ffdffff from {2, 1}+{2, 1}
```

```
DEBUG 2022-05-04 23:20:57,583 [shard 0] osd - pg_epoch 17 pg[1.1( v 16'191 (0'0,16'191) local-lis /les=12/13 n=2 ec=9/9 lis/c=12/0 les/c/f=13/0/0 sis=17) [2,1] r=-1 lpr=17 pi=[12,17)/1 crt=16'191 lcod 0'0 mlcod 0'0 unknown on_new_interval  
checking missing set deletes flag. missing = missing(0 may_include_deletes = 1)
```

```
DEBUG 2022-05-04 23:20:57,583 [shard 0] osd - pg_epoch 17 pg[1.1( v 16'191 (0'0,16'191) local-lis /les=12/13 n=2 ec=9/9 lis/c=12/0 les/c/f=13/0/0 sis=17) [2,1] r=-1 lpr=17 pi=[12,17)/1 crt=16'191 lcod 0'0 mlcod 0'0 unknown init_hb_stamps n  
ow {}
```

```
DEBUG 2022-05-04 23:20:57,583 [shard 0] osd - pg_epoch 17 pg[1.1( v 16'191 (0'0,16'191) local-lis /les=12/13 n=2 ec=9/9 lis/c=12/0 les/c/f=13/0/0 sis=17) [2,1] r=-1 lpr=17 pi=[12,17)/1 crt=16'191 lcod 0'0 mlcod 0'0 unknown on_new_interval  
prior_readable_until_ub 0.000000000s (mnow 3.070197105s + 0.000000000s)
```

```
INFO 2022-05-04 23:20:57,583 [shard 0] osd - pg_epoch 17 pg[1.1( v 16'191 (0'0,16'191) local-lis /les=12/13 n=2 ec=9/9 lis/c=12/0 les/c/f=13/0/0 sis=17) [2,1] r=-1 lpr=17 pi=[12,17)/1 crt=16'191 lcod 0'0 mlcod 0'0 unknown start_peering_in  
terval up {2, 0, 1} -> {2, 1}, acting {2, 0, 1} -> {2, 1}, acting_primary 2 -> 2, up_primary 2 -> 2, role 1 -> -1, features acting 4540138303579357183 upacting 4540138303579357183
```

```

DEBUG 2022-05-04 23:20:57,583 [shard 0] osd - pg_epoch 17 pg[1.1( v 16'191 (0'0,16'191] local-lis
/les=12/13 n=2 ec=9/9 lis/c=12/0 les/c/f=13/0/0 sis=17) [2,1] r=-1 lpr=17 pi=[12,17)/1 crt=16'191
lcod 0'0 mlcod 0'0 unknown clear_primary_st
ate

DEBUG 2022-05-04 23:20:57,583 [shard 0] osd - pg_epoch 17 pg[1.1( v 16'191 (0'0,16'191] local-lis
/les=12/13 n=2 ec=9/9 lis/c=12/0 les/c/f=13/0/0 sis=17) [2,1] r=-1 lpr=17 pi=[12,17)/1 crt=16'191
lcod 0'0 mlcod 0'0 unknown on_change:
DEBUG 2022-05-04 23:20:57,583 [shard 0] osd - pg_epoch 17 pg[1.1( v 16'191 (0'0,16'191] local-lis
/les=12/13 n=2 ec=9/9 lis/c=12/0 les/c/f=13/0/0 sis=17) [2,1] r=-1 lpr=17 pi=[12,17)/1 crt=16'191
lcod 0'0 mlcod 0'0 unknown on_change: drop
ping requests

DEBUG 2022-05-04 23:20:57,583 [shard 0] osd - pg_epoch 17 pg[1.1( v 16'191 (0'0,16'191] local-lis
/les=12/13 n=2 ec=9/9 lis/c=12/0 les/c/f=13/0/0 sis=17) [2,1] r=-1 lpr=17 pi=[12,17)/1 crt=16'191
lcod 0'0 mlcod 0'0 unknown NOTIFY check_rec
overy_sources no source osds () went down

ERROR 2022-05-04 23:20:57,583 [shard 0] none - pg_epoch 17 pg[1.1( v 16'191 (0'0,16'191] local-li
s/les=12/13 n=2 ec=9/9 lis/c=12/0 les/c/f=13/0/0 sis=17) [2,1] r=-1 lpr=17 pi=[12,17)/1 crt=16'191
lcod 0'0 mlcod 0'0 unknown NOTIFY 1.1 past
_intervals [12,17) start interval does not contain the required bound [9,17) start

ERROR 2022-05-04 23:20:57,583 [shard 0] none - ../src/osd/PeeringState.cc:968 : In function 'void
PeeringState::check_past_interval_bounds() const', abort(%s)

past_interval start interval mismatch

Aborting on shard 0.

Backtrace:

Reactor stalled for 11600 ms on shard 0. Backtrace: 0x44700 0xda36731 0xd7e16ef 0xd7f9edb 0xd7fa37
e 0xd7fa60e 0xd7fa8d9 0x7ff142eedalf 0xcccd78 0x6e450d3 0x6e4704e 0x6e4c82b 0x6e4d49e 0x6e4db68 0x6
e41807 0x6e41cf3 0x6e42282 0x7ff142eedalf 0
x3d2a1 0x268a3 0xd4220b 0x3f20d5b 0x410834c 0x45b91cc 0x245ae64 0x420e7be 0x2126da6 0x328ac07 0x3
28b55a 0x1ed67e2 0x1f387ba 0x1f39421 0xd7bdd40 0xd81182c 0xd993d43 0xd995d54 0xd3f3a11 0xd3f6e53 0
x1917482 0x27b74 0x15f04bd
0# gsignal in /lib64/libc.so.6

1# abort in /lib64/libc.so.6

2# ceph::__ceph_abort(char const*, int, char const*, std::__cxx11::basic_string<char, std::char_t
raits<char>, std::allocator<char> > const&) at /home/sam/git-checkouts/ceph/build/../src/seastar/i
nclude/seastar/util/log.hh:106
3# PeeringState::check_past_interval_bounds() const at /usr/include/c++/11/bits/basic_string.h:67
2

4# PeeringState::Reset::react(PeeringState::AdvMap const&) at /home/sam/git-checkouts/ceph/build/
../src/osd/PeeringState.cc:4694

5# boost::statechart::simple_state<PeeringState::Reset, PeeringState::PeeringMachine, boost::mpl::
list<mpl_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na, mp
l_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na, mp
l_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na, mpl_::na>, (boost::statechart::history_mo
de)0>::react_Impl(boost::statechart::event_
base const&, void const*) at /home/sam/git-checkouts/ceph/build/boost/include/boost/statechart/res
```

ult.hpp:70

```
6# boost::statechart::state_machine<PeeringState::PeeringMachine, PeeringState::Initial, std::allocator<boost::statechart::none>, boost::statechart::null_exception_translator>::send_event(boost::statechart::event_base const&) at /home/sam  
/git-checkouts/ceph/build/boost/include/boost/statechart/state_machine.hpp:87
```

```
7# PeeringState::advance_map(boost::local_shared_ptr<OSDMap const>, boost::local_shared_ptr<OSDMap const>, std::vector<int, std::allocator<int> >&, int, std::vector<int, std::allocator<int> >&, int, PeeringCtx&) at /home/sam/git-checkouts  
/ceph/build/boost/include/boost/statechart/state_machine.hpp:275
```

```
8# crimson::osd::PG::handle_advance_map(boost::local_shared_ptr<OSDMap const>, PeeringCtx&) at /home/sam/git-checkouts/ceph/build/../src/crimson/osd/pg.cc:497
```

```
9# auto seastar::futurize_invoke<crimson::osd::PGAdvanceMap::start()::{lambda()#1}::operator()() const:{lambda(unsigned int)#1}::operator()(unsigned int) const:{lambda(boost::local_shared_ptr<OSDMap const>&&)#1}&, boost::local_shared_ptr<OSDMap const> >(crimson::osd::PGAdvanceMap::start()::{lambda()#1}::operator()() const:{lambda(unsigned int)#1}::operator()(unsigned int) const:{lambda(boost::local_shared_ptr<OSDMap const>&&)#1}&, boost::local_shared_ptr<OSDMap const>&  
&) at /home/sam/git-checkouts/ceph/build/../src/crimson/osd/osd_operations/pg_advance_map.cc:72
```

```
10# _ZN7seastar20noncopyable_functionIFNS_6futureIvEEON5boost16local_shared_ptrIK6OSDMapEEEE17dire  
ct_vtable_forIZNS1_IS7_E4thenIZZZN7crimson3osd12PGAdvanceMap5startEvENKUlvE_c1EvENKUljE_c1EjEUls8_  
E_S2_EET0_OT_EU1DpOT_E_E4callEPKSA_S8_ at /  
home/sam/git-checkouts/ceph/build/../src/seastar/include/seastar/util/noncopyable_function.hh:125
```

```
11# auto seastar::internal::future_invoke<seastar::noncopyable_function<seastar::future<void> (boost::local_shared_ptr<OSDMap const>&&)>&, boost::local_shared_ptr<OSDMap const> >(seastar::noncopyable_function<seastar::future<void> (boost::local_shared_ptr<OSDMap const>&&)>&, boost::local_shared_ptr<OSDMap const>&&) at /home/sam/git-checkouts/ceph/build/../src/seastar/include/seastar/core/future.hh:1213
```

```
12# void seastar::futurize<seastar::future<void> >::satisfy_with_result_of<seastar::future<boost::local_shared_ptr<OSDMap const> >::then_impl_nrvo<seastar::noncopyable_function<seastar::future<void> (boost::local_shared_ptr<OSDMap const>&&)>&, seastar::future<void> >(seastar::noncopyable_function<seastar::future<void> (boost::local_shared_ptr<OSDMap const>&&)>&&):&{lambda(seastar::internal::promise_base_with_type<void>&&, seastar::noncopyable_function<seastar::future<void> (boost::local_shared_ptr<OSDMap const>&&)>&, seastar::future_state<boost::local_shared_ptr<OSDMap const> >&&#1}::operator()(seastar::internal::promise_base_with_type<void>&&, seastar::noncopyable_function<seastar::future<void> (boost::local_shared_ptr<OSDMap const>&&)>&, seastar::future_state<boost::local_shared_ptr<OSDMap const> >&&) const:{lambda()#1}>(seastar::internal::promise_base_with_type<void>&&, seastar::noncopyable_function<seastar::future<void> (boost::local_shared_ptr<OSDMap const>&&)>&&) at /home/sam/git-checkouts/ceph/build/../src/seastar/include/seastar/core/future.hh:2120
```

```
13# seastar::continuation<seastar::internal::promise_base_with_type<void>, seastar::noncopyable_function<seastar::future<void> (boost::local_shared_ptr<OSDMap const>&&)>, seastar::future<boost::local_shared_ptr<OSDMap const> >::then_impl_n  
rvo<seastar::noncopyable_function<seastar::future<void> (boost::local_shared_ptr<OSDMap const>&&)>&, seastar::future<void> >(seastar::noncopyable_function<seastar::future<void> (boost::local_shared_ptr<OSDMap const>&&)>&&):&{lambda(seastar::internal::promise_base_with_type<void>&&, seastar::noncopyable_function<seastar::future<void> (boost::local_shared_ptr<OSDMap const>&&)>&, seastar::future_state<boost::local_shared_ptr<OSDMap const> >&&#1}, boost::local_shared_ptr<OSDMap const> >::run_and_dispose() at /home/sam/git-checkouts/ceph/build/../src/seastar/include/seastar/core/future.hh:1575
```

```
14# seastar::reactor::run_tasks(seastar::reactor::task_queue&) at /home/sam/git-checkouts/ceph/build/.../src/seastar/src/core/reactor.cc:2345
```

```
15# seastar::reactor::run_some_tasks() at /home/sam/git-checkouts/ceph/build/..../src/seastar/src/core/reactor.cc:2755

16# seastar::reactor::do_run() at /home/sam/git-checkouts/ceph/build/..../src/seastar/src/core/reactor.cc:2923

17# seastar::reactor::run() at /home/sam/git-checkouts/ceph/build/..../src/seastar/src/core/reactor.cc:2806

18# seastar::app_template::run_DEPRECATED(int, char**, std::function<void ()>&&) at /home/sam/git-checkouts/ceph/build/..../src/seastar/src/core/app-template.cc:265

19# seastar::app_template::run(int, char**, std::function<seastar::future<int> ()>&&) at /home/sam/git-checkouts/ceph/build/..../src/seastar/src/core/app-template.cc:156

20# main at /home/sam/git-checkouts/ceph/build/..../src/crimson/osd/main.cc:238

21# __libc_start_main in /lib64/libc.so.6
```

## Related issues:

Related to RADOS - Bug #49689: osd/PeeringState.cc: ceph\_abort\_msg("past\_inte...")

Resolved

## History

---

### #1 - 10/12/2022 05:09 PM - Samuel Just

- *Project changed from Ceph to crimson*

### #2 - 05/16/2023 04:57 PM - Matan Breizman

- *Related to Bug #49689: osd/PeeringState.cc: ceph\_abort\_msg("past\_interval start interval mismatch") start added*

### #3 - 05/17/2023 07:26 AM - Matan Breizman

This issue was fixed in Classic <https://github.com/ceph/ceph/pull/48706>.  
Similar changes should be applied in Crimson.

### #4 - 06/01/2023 09:13 AM - Matan Breizman

[https://pulpito.ceph.com/matan-2023-05-31\\_17:32:26-crimson-rados-wip-matanb-crimson-maybe-handle-osdmap-distro-crimson-smithi/7292139/](https://pulpito.ceph.com/matan-2023-05-31_17:32:26-crimson-rados-wip-matanb-crimson-maybe-handle-osdmap-distro-crimson-smithi/7292139/)

### #5 - 06/01/2023 09:14 AM - Matan Breizman

- *Assignee set to Matan Breizman*

- *Priority changed from Normal to High*

### #6 - 01/18/2024 10:42 AM - Matan Breizman

- *Status changed from New to Resolved*

Resolved by classical fix