

RADOS - Bug #41064

OSD: assert(objiter->second->version > last_divergent_update) fails when there is only entry in "divergent entries"

08/05/2019 04:07 AM - Xuehan Xu

Status:	New	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:			
Category:	Correctness/Safety		
Target version:	v15.0.0		
Source:	Community (dev)	Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:		Component(RADOS):	
Regression:	No	Pull request ID:	29480
Severity:	3 - minor	Crash signature:	
Reviewed:			

Description

Recently, some OSDs in one of our cluster failed to start after a power outage

One OSD's log is as follows:

```
2019-08-02 13:33:53.740154 2b5541419700 10 merge_log divergent 155257'1755881 (143567'1708027) modify 3f1dfa35/rbd_data.dd1a7124d7da84.0000000000016a84/head//1 by client.101808178.0:72164 2019-08-01 17:02:31.918250
2019-08-02 13:33:53.742187 2b5541419700 10 _merge_object_divergent_entries: merging hoid 3f1dfa35/rbd_data.dd1a7124d7da84.0000000000016a84/head//1 entries: 155257'1755881 (143567'1708027) modify 3f1dfa35/rbd_data.dd1a7124d7da84.0000000000016a84/head//1 by client.101808178.0:72164 2019-08-01 17:02:31.918250
2019-08-02 13:33:53.748525 2b5541419700 10 _merge_object_divergent_entries: hoid 3f1dfa35/rbd_data.dd1a7124d7da84.0000000000016a84/head//1 prior_version: 143567'1708027 first_divergent_update: 155257'1755881 last_divergent_update: 155257'1755881
2019-08-02 13:33:53.773975 2b5541419700 -1 error_msg osd/PGLog.cc: In function 'static void PGLog::_merge_object_divergent_entries(const PGLog::IndexedLog&, const hobject_t&, const std::list<pg_log_entry_t>&, const pg_info_t&, eversion_t, pg_missing_t&, boost::optional<std::pair<eversion_t, hobject_t>>*, PGLog::LogEntryHandler*)' thread 2b5541419700 time 2019-08-02 13:33:53.748535 osd/PGLog.cc: 384: FAILED assert(objiter->second->version > last_divergent_update)
```

the pglog of that osd is as follows:

```
{
  "head": "155257'1755880",
  "tail": "154991'1752834",
  "log": [
    {
      "op": "modify ",
      "object": "9eb99a35\\rbd_data.83ce955b22648.0000000000024fca\\head\\/\\/1",
      "version": "154991'1752835",
      "prior_version": "154991'1752834",
      "reqid": "client.72368396.0:31936439",
      "extra_reqids": [],
      "mtime": "2019-08-01 04:08:38.039238",
      "mod_desc": {
        "object_mod_desc": {
          "can_local_rollback": false,
          "rollback_info_completed": false,
```

```
        "ops": []
      }
    },
  },
  .....
  {
    "op": "modify ",
    "object": "3f1dfa35\rbid_data.dd1a7124d7da84.0000000000016a84\head\//1",
    "version": "155257'1755881",
    "prior_version": "143567'1708027",
    "reqid": "client.101808178.0:72164",
    "extra_reqids": [],
    "mtime": "2019-08-01 17:02:31.918250",
    "mod_desc": {
      "object_mod_desc": {
        "can_local_rollback": false,
        "rollback_info_completed": false,
        "ops": []
      }
    }
  }
]
}
```

We believe that the failure is due to the assert "assert(objiter->second->version > last_divergent_update)" doesn't take into account the single-entry divergent entries case.

History

#1 - 08/05/2019 05:36 AM - Xuehan Xu

<https://github.com/ceph/ceph/pull/29480>