

Ceph - Bug #24524

Newly added OSDs do not start in Mimic

06/14/2018 10:09 AM - Vitaliy Filippov

Status: Duplicate	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee:	
Category: OSD	
Target version:	
Source:	Affected Versions: v13.2.0
Tags:	ceph-qa-suite:
Backport:	Pull request ID:
Regression: Yes	Crash signature (v1):
Severity: 2 - major	Crash signature (v2):
Reviewed:	

Description

Hi!

In my test cluster with Ceph Mimic installed (from scratch, not upgraded from luminous) newly added OSD fail to start with the following error:

```
-8> 2018-06-14 12:51:04.616 7f58cf8a0700 3 osd.3 0 handle_osd_map epochs [533,533], i have 0, src has [533,1182]
-7> 2018-06-14 12:51:04.616 7f58cf8a0700 -1 osd.3 0 failed to load OSD map for epoch 532, got 0 bytes
```

From reading the code here <https://github.com/ceph/ceph/blob/master/src/osd/OSD.cc#L7330> I suspect that when added OSDs try to load all osdmaps from oldest to newest (533-1182 in my case) for each Nth map it first loads (N-1)th map and compares them... But there is no (N-1)th map for the oldest one, so it dies with 'assertion failed'.

I'm trying to fix it like this:

```
diff --git a/src/osd/OSD.cc b/src/osd/OSD.cc
index a6cf188..b026401 100644
--- a/src/osd/OSD.cc
+++ b/src/osd/OSD.cc
@@ -7357,6 +7357,9 @@ void OSD::handle_osd_map(MOSDMap *m)
// check for deleted pools
OSDMapRef lastmap;
for (auto& i : added_maps) {
if (i.first <= first) {
+   continue;
+ }
if (!lastmap) {
lastmap = get_map(i.first - 1);
}
```

So please tell me if I'm correct and push this fix to your repository if yes :)

Related issues:

Duplicates RADOS - Bug #24423: failed to load OSD map for epoch X, got 0 bytes	Resolved	06/05/2018
--	-----------------	-------------------

History

#1 - 06/14/2018 10:11 AM - Vitaliy Filippov

```
diff --git a/src/osd/OSD.cc b/src/osd/OSD.cc
```

```
index a6cf188..b026401 100644
--- a/src/osd/OSD.cc
+++ b/src/osd/OSD.cc
@@ -7357,6 +7357,9 @@ void OSD::handle_osd_map(MOSDMap *m)
    // check for deleted pools
    OSDMapRef lastmap;
    for (auto& i : added_maps) {
+   if (i.first <= first) {
+     continue;
+   }
    if (!lastmap) {
      lastmap = get_map(i.first - 1);
    }
  }
```

#2 - 06/17/2018 08:39 AM - Sergey Malinin

More details here:

<http://tracker.ceph.com/issues/24423>

#3 - 06/18/2018 05:07 PM - Greg Farnum

- *Duplicates Bug #24423: failed to load OSD map for epoch X, got 0 bytes added*

#4 - 06/18/2018 05:07 PM - Greg Farnum

- *Status changed from New to Duplicate*

#5 - 06/19/2018 11:54 AM - Igor Fedotov

- *Duplicated by Bug #24450: OSD Caught signal (Aborted) added*

#6 - 06/19/2018 11:54 AM - Igor Fedotov

- *Duplicated by deleted (Bug #24450: OSD Caught signal (Aborted))*