

RADOS - Bug #22847

ceph osd force-create-pg cause all ceph-mon to crash and unable to come up again

02/01/2018 04:29 AM - Frank Li

Status:	Resolved	% Done:	0%
Priority:	High	Spent time:	0.00 hour
Assignee:			
Category:			
Target version:	v12.2.3		
Source:	Community (user)	Affected Versions:	v12.2.2
Tags:		ceph-qa-suite:	
Backport:	luminous	Component(RADOS):	
Regression:	No	Pull request ID:	
Severity:	2 - major	Crash signature (v1):	
Reviewed:		Crash signature (v2):	

Description

during the course of trouble-shooting an osd issue, I ran this command:

```
ceph osd force-create-pg 1.ace11d67
```

then all of the ceph-mon crashed with this error:

```
--- begin dump of recent events ---
```

```
0> 2018-01-31 22:47:22.959665 7fc64350e700 -1 ** Caught signal (Aborted) *
```

```
in thread 7fc64350e700 thread_name:cpu_tp
```

```
ceph version 12.2.2 (cf0bae00000000000000000000000000000000) luminous (stable)
```

```
1: (()+0x8eae11) [0x55f1113fae11]
```

```
2: (()+0xf5e0) [0x7fc64aafa5e0]
```

```
3: (gsignal()+0x37) [0x7fc647fca1f7]
```

```
4: (abort()+0x148) [0x7fc647fcb8e8]
```

```
5: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x284) [0x55f1110fa4a4]
```

```
6: (()+0x2ccc4e) [0x55f110ddcc4e]
```

```
7: (OSDMonitor::update_creating_pgs()+0x98b) [0x55f11102232b]
```

```
8: (C_UpdateCreatingPGs::finish(int)+0x79) [0x55f1110777b9]
```

```
9: (Context::complete(int)+0x9) [0x55f110ed30c9]
```

```
10: (ParallelPGMapper::WQ::_process(ParallelPGMapper::Item*, ThreadPool::TPHandle&)+0x7f) [0x55f111204e1f]
```

```
11: (ThreadPool::worker(ThreadPool::WorkThread*)+0xa8e) [0x55f111100f1e]
```

```
12: (ThreadPool::WorkThread::entry()+0x10) [0x55f111101e00]
```

```
13: (()+0x7e25) [0x7fc64aaf2e25]
```

```
14: (clone()+0x6d) [0x7fc64808d34d]
```

```
NOTE: a copy of the executable, or `objdump -rds &lt;executable>` is needed to interpret this.
```

here is the current ceph -s output:

```
[root@dl1-kaf101 fri]# ceph -s
```

```
cluster:
```

```
id: 021a1428-fea5-4697-bcd0-a45c1c2ca80b
```

```
health: HEALTH_WARN
```

```
nodown flag(s) set
```

```
32 osds down
```

```
4 hosts (32 osds) down
```

```
2 racks (32 osds) down
```

```
1 row (32 osds) down
```

```
Reduced data availability: 10240 pgs inactive, 5 pgs down, 3431 pgs peering, 3643 pgs stale
```

```
Degraded data redundancy: 215798/552339 objects degraded (39.070%), 10240 pgs unclean, 4666 pgs degraded, 4666 pgs
```

```
undersized
```

```
services:
```

```
mon: 5 daemons, quorum dl1-kaf101,dl1-kaf201,dl1-kaf301,dl1-kaf302,dl1-kaf401
mgr: dl1-kaf101(active), standbys: dl1-kaf201
osd: 64 osds: 19 up, 51 in
      flags nodown
```

data:

```
  pools: 3 pools, 10240 pgs
  objects: 179k objects, 712 GB
  usage: 2135 GB used, 461 TB / 463 TB avail
  pgs: 20.879% pgs unknown
      79.121% pgs not active
      215798/552339 objects degraded (39.070%)
      3381 stale+undersized+degraded+peered
      3169 peering
      2138 unknown
      1285 undersized+degraded+peered
      262 stale+peering
      5 down
```

Related issues:

Copied to RADOS - Backport #22942: luminous: ceph osd force-create-pg cause a...

Resolved

History

#1 - 02/02/2018 05:46 PM - Sage Weil

- Status changed from New to Need More Info
- Priority changed from Normal to High

Can you attach the entire osd log for the crashed osd? (In particular, we need to see what assertion failed.) Thanks!

#2 - 02/02/2018 05:46 PM - Sage Weil

- Project changed from Ceph to RADOS
- Category deleted (Monitor)

#3 - 02/02/2018 06:09 PM - Sage Weil

- Status changed from Need More Info to Fix Under Review

<https://github.com/ceph/ceph/pull/20267>

#4 - 02/02/2018 06:21 PM - Frank Li

```
4> 2018-01-31 22:47:22.942381 7fc641d0b700-1 10.102.52.37:6789/0 <== mon.0 10.102.52.37:6789/0 0 ==== log(1 entries from seq 1 at
2018-01-31 22:47:22.934106) v1 ===== 0+0+0 (0 0 0) 0x55f11b065f80 con 0x55f11b5a4000
-3> 2018-01-31 22:47:22.942394 7fc641d0b700 5 mon.dl1-kaf101@0(electing) e6_ms_dispatch setting monitor caps on this connection
-2> 2018-01-31 22:47:22.942405 7fc641d0b700 5 mon.dl1-kaf101@0(electing).paxos(paxos recovering c 28110997..28111530) is_readable = 0 -
now=2018-01-31 22:47:22.942405 lease_expire=0.000000 has v0 lc 28111530
-1> 2018-01-31 22:47:22.942422 7fc641d0b700 5 mon.dl1-kaf101@0(electing).paxos(paxos recovering c 28110997..28111530) is_readable = 0 -
now=2018-01-31 22:47:22.942422 lease_expire=0.000000 has v0 lc 28111530
0> 2018-01-31 22:47:22.955415 7fc64350e700 -1
/home/jenkins-build/build/workspace/ceph-build/ARCH/x86_64/AVAILABLE_ARCH/x86_64/AVAILABLE_DIST/centos7/DIST/centos7/MACHINE_SIZ
E/huge/release/12.2.2/rpm/el7/BUILD/ceph-12.2.2/src/osd/OSDMapMapping.h: In function 'void OSDMapMapping::get(pg_t, std::vector<int>*, int*,
std::vector<int>*, int*) const' thread 7fc64350e700 time 2018-01-31 22:47:22.952877
/home/jenkins-build/build/workspace/ceph-build/ARCH/x86_64/AVAILABLE_ARCH/x86_64/AVAILABLE_DIST/centos7/DIST/centos7/MACHINE_SIZ
E/huge/release/12.2.2/rpm/el7/BUILD/ceph-12.2.2/src/osd/OSDMapMapping.h: 288: FAILED assert(pgid.ps() < p->second.pg_num)
```

#5 - 02/02/2018 06:22 PM - Frank Li

Please note the Crash happend on the monitor, not the OSD, the OSDs all stayed up, but all the monitor crashed.

#6 - 02/02/2018 08:29 PM - Frank Li

Just for future operational references, is there anyway to revert the Monitor map to a previous state in the case of an operation that crashed all the monitor daemons ?

#7 - 02/02/2018 09:14 PM - Frank Li

I've updated all the ceph-mon with the RPMs from the patch repo, they came up fine, and I've restarted the OSDs, but now they seems to be stuck in the booting stage:

```
[root@dl1-cephosd101 ceph]# ceph --admin-daemon /var/run/ceph/dl1aprod-osd.7.asok status {
"cluster_fsid": "021a1428-fea5-4697-bcd0-a45c1c2ca80b",
"osd_fsid": "29f0db5d-bd64-4c4e-9ef7-04a8ac2b88e7",
"whoami": 7,
"state": "booting",
"oldest_map": 133275,
"newest_map": 133997,
"num_pgs": 483
}
```

here is the OSD debug log:

```
2018-02-02 21:13:03.959263 7f09fa6d9700 2 osd.0 133997 ms_handle_reset con 0x559fbdf82800 session 0
2018-02-02 21:13:03.959268 7f0a0bd3d700 5 -- 10.102.52.11:6800/38382 reap_dead delete 0x559fbdf82800
2018-02-02 21:13:03.959276 7f0a0bd3d700 5 -- 10.102.52.11:6800/38382 reap_dead delete 0x559fbdf1000
2018-02-02 21:13:03.959280 7f0a0bd3d700 5 -- 10.102.52.11:6800/38382 reap_dead delete 0x559fbdf2800
2018-02-02 21:13:03.959283 7f0a0bd3d700 5 -- 10.102.52.11:6800/38382 reap_dead delete 0x559fbdf5800
2018-02-02 21:13:03.959286 7f0a0bd3d700 5 -- 10.102.52.11:6800/38382 reap_dead delete 0x559fbdf8800
2018-02-02 21:13:04.285164 7f0a0ad3b700 1 -- 10.102.52.11:6800/38382 >> - conn(0x559fbdf8800 :6800 s=STATE_ACCEPTING pgs=0 cs=0 l=0)._process_connection sd=166
2018-02-02 21:13:04.285329 7f0a0ad3b700 1 -- 10.102.52.11:6800/38382 >> - conn(0x559fbdf8800 :6800 s=STATE_ACCEPTING_WAIT_BANNER_ADDR pgs=0 cs=0 l=0).read_bulk peer close file descriptor 166
2018-02-02 21:13:04.285347 7f0a0ad3b700 1 -- 10.102.52.11:6800/38382 >> - conn(0x559fbdf8800 :6800 s=STATE_ACCEPTING_WAIT_BANNER_ADDR pgs=0 cs=0 l=0).read_until read failed
2018-02-02 21:13:04.285351 7f0a0ad3b700 1 -- 10.102.52.11:6800/38382 >> - conn(0x559fbdf8800 :6800 s=STATE_ACCEPTING_WAIT_BANNER_ADDR pgs=0 cs=0 l=0)._process_connection read peer banner and addr failed
2018-02-02 21:13:04.285376 7f0a0ad3b700 2 -- 10.102.52.11:6800/38382 >> - conn(0x559fbdf8800 :6800 s=STATE_ACCEPTING_WAIT_BANNER_ADDR pgs=0 cs=0 l=0)._stop
2018-02-02 21:13:04.285427 7f09fa6d9700 2 osd.0 133997 ms_handle_reset con 0x559fbdf8800 session 0
2018-02-02 21:13:04.695847 7f0a07d77700 10 osd.0 133997 tick
2018-02-02 21:13:04.695864 7f0a07d77700 10 osd.0 133997 do_waiters -- start
2018-02-02 21:13:04.695867 7f0a07d77700 10 osd.0 133997 do_waiters -- finish
2018-02-02 21:13:04.866376 7f0a07576700 10 osd.0 133997 tick_without_osd_lock
2018-02-02 21:13:04.899985 7f09e9eb8700 20 osd.0 133997 update_osd_stat osd_stat(46278 MB used, 9266 GB avail, 9312 GB total, peers [] op hist [])
2018-02-02 21:13:04.900005 7f09e9eb8700 20 osd.0 133997 check_full_status cur ratio 0.00485325. nearfull_ratio 0.85. backfillfull_ratio 0.9, full_ratio 0.95, failsafe_ratio 0.97, new state none
2018-02-02 21:13:04.900022 7f09e9eb8700 5 osd.0 133997 heartbeat: osd_stat(46278 MB used, 9266 GB avail, 9312 GB total, peers [] op hist [])
```

#8 - 02/02/2018 09:15 PM - Frank Li

Frank Li wrote:

I've updated all the ceph-mon with the RPMs from the patch repo, they came up fine, and I've restarted the OSDs, but now they seems to be stuck in the booting stage:

Should I be updating the ceph-osd to the same patched version ??

```
[root@dl1-cephosd101 ceph]# ceph --admin-daemon /var/run/ceph/dl1approd-osd.7.asok status {
"cluster_fsid": "021a1428-fea5-4697-bcd0-a45c1c2ca80b",
"osd_fsid": "29f0db5d-bd64-4c4e-9ef7-04a8ac2b88e7",
"whoami": 7,
"state": "booting",
"oldest_map": 133275,
"newest_map": 133997,
"num_pgs": 483
}
```

here is the OSD debug log:

```
2018-02-02 21:13:03.959263 7f09fa6d9700 2 osd.0 133997 ms_handle_reset con 0x559fbdf82800 session 0
2018-02-02 21:13:03.959268 7f0a0bd3d700 5 -- 10.102.52.11:6800/38382 reap_dead delete 0x559fbdf82800
2018-02-02 21:13:03.959276 7f0a0bd3d700 5 -- 10.102.52.11:6800/38382 reap_dead delete 0x559fbdf1000
2018-02-02 21:13:03.959280 7f0a0bd3d700 5 -- 10.102.52.11:6800/38382 reap_dead delete 0x559fbdf2800
2018-02-02 21:13:03.959283 7f0a0bd3d700 5 -- 10.102.52.11:6800/38382 reap_dead delete 0x559fbdf5800
2018-02-02 21:13:03.959286 7f0a0bd3d700 5 -- 10.102.52.11:6800/38382 reap_dead delete 0x559fbdf8800
2018-02-02 21:13:04.285164 7f0a0ad3b700 1 -- 10.102.52.11:6800/38382 >> - conn(0x559fbdf8800 :6800 s=STATE_ACCEPTING pgs=0 cs=0 l=0)._process_connection sd=166
2018-02-02 21:13:04.285329 7f0a0ad3b700 1 -- 10.102.52.11:6800/38382 >> - conn(0x559fbdf8800 :6800
s=STATE_ACCEPTING_WAIT_BANNER_ADDR pgs=0 cs=0 l=0).read_bulk peer close file descriptor 166
2018-02-02 21:13:04.285347 7f0a0ad3b700 1 -- 10.102.52.11:6800/38382 >> - conn(0x559fbdf8800 :6800
s=STATE_ACCEPTING_WAIT_BANNER_ADDR pgs=0 cs=0 l=0).read_until read failed
2018-02-02 21:13:04.285351 7f0a0ad3b700 1 -- 10.102.52.11:6800/38382 >> - conn(0x559fbdf8800 :6800
s=STATE_ACCEPTING_WAIT_BANNER_ADDR pgs=0 cs=0 l=0)._process_connection read peer banner and addr failed
2018-02-02 21:13:04.285376 7f0a0ad3b700 2 -- 10.102.52.11:6800/38382 >> - conn(0x559fbdf8800 :6800
s=STATE_ACCEPTING_WAIT_BANNER_ADDR pgs=0 cs=0 l=0)._stop
2018-02-02 21:13:04.285427 7f09fa6d9700 2 osd.0 133997 ms_handle_reset con 0x559fbdf8800 session 0
2018-02-02 21:13:04.695847 7f0a07d77700 10 osd.0 133997 tick
2018-02-02 21:13:04.695864 7f0a07d77700 10 osd.0 133997 do_waiters -- start
2018-02-02 21:13:04.695867 7f0a07d77700 10 osd.0 133997 do_waiters -- finish
2018-02-02 21:13:04.866376 7f0a07576700 10 osd.0 133997 tick_without_osd_lock
2018-02-02 21:13:04.899985 7f09e9eb8700 20 osd.0 133997 update_osd_stat osd_stat(46278 MB used, 9266 GB avail, 9312 GB total, peers []
op hist [])
2018-02-02 21:13:04.900005 7f09e9eb8700 20 osd.0 133997 check_full_status cur ratio 0.00485325. nearfull_ratio 0.85. backfillfull_ratio 0.9,
full_ratio 0.95, failsafe_ratio 0.97, new state none
2018-02-02 21:13:04.900022 7f09e9eb8700 5 osd.0 133997 heartbeat: osd_stat(46278 MB used, 9266 GB avail, 9312 GB total, peers [] op hist
[])
```

#9 - 02/04/2018 07:29 AM - Frank Li

ALso, while the monitors came up and form a forum, but very soon, they would all stop responding again, and then I find three out of the five would not respond to the mon_status command from the asok socket, and two of the ones that responded are in probing mode and is considered out of the quorum.

#10 - 02/05/2018 03:52 AM - Kefu Chai

- Target version changed from v12.2.2 to v12.2.3

Should I be updating the ceph-osd to the same patched version ??

no need to update ceph-osd.

but very soon, they would all stop responding again

could you post the debug-mon=20 log of one of the mon nodes?

#11 - 02/06/2018 01:23 AM - Frank Li

- File dl1approd-mon.dl1-kaf101.log.prob.1.gz added

- File dl1approd-mon.dl1-kaf101.log.prob.2.gz added

please see attached logs for when the monitor was started, and then later got into the stuck mode.
I just replaced the IP address with xxxxx.

#12 - 02/06/2018 06:50 AM - Kefu Chai

- Status changed from Fix Under Review to Pending Backport

#13 - 02/06/2018 07:49 PM - Frank Li

so anything I can do to help recover the cluster ??

#14 - 02/07/2018 05:50 AM - Nathan Cutler

- Backport set to luminous

#15 - 02/07/2018 05:50 AM - Nathan Cutler

- Copied to Backport #22942: luminous: ceph osd force-create-pg cause all ceph-mon to crash and unable to come up again added

#16 - 02/07/2018 06:12 AM - Kefu Chai

```
2018-02-05 18:02:21.203957 7f9318dd5700 5 mon.dll-kaf101@0(electing).elector(1497) so far i have { mon.0: features 2305244844532236283 mon_feature_t([kraken, luminous]), mon.1: features 2305244844532236283 mon_feature_t([kraken, luminous]), mon.2: features 2305244844532236283 mon_feature_t([kraken, luminous]), mon.3: features 2305244844532236283 mon_feature_t([kraken, luminous]), mon.4: features 2305244844532236283 mon_feature_t([kraken, luminous]) }
2018-02-05 18:02:21.203974 7f9318dd5700 10 mon.dll-kaf101@0(electing).elector(1497) bump_epoch 1497 to 1498
..

2018-02-05 18:02:21.207901 7f9318dd5700 10 mon.dll-kaf101@0(electing) e6 win_election epoch 1498 quorum 0,1,2,3,4 features 2305244844532236283 mon_features mon_feature_t([kraken, luminous])
```

the cluster formed a quorum of [0,1,2,3,4] since 18:02:21. and it was not in probing mode after the quorum was formed.

and two of the ones that responded are in probing mode and is considered out of the quorum.

how do you come to this conclusion?

I find three out of the five would not respond to the mon_status command from the asok socket

did you sent mon_status to dl1approd-mon.dll-kaf101?

also i'd suggest you create another ticket to track this issue. as it's different from the one we've addressed in [#22847](#).

#17 - 02/07/2018 06:56 AM - Kefu Chai

```
2018-02-05 18:32:25.923127 7f931b5da700 0 log_channel(cluster) log [INF] : osd.0 marked down after no beacon for 900.076451 seconds
...
2018-02-05 18:32:25.925921 7f931b5da700 0 log_channel(cluster) log [INF] : osd.44 marked down after no beacon
```

```
for 900.076451 seconds
```

was any osd up when you were testing?

#18 - 02/07/2018 05:58 PM - Frank Li

Kefu Chai wrote:

[...]

was any osd up when you were testing?

Yes, but they were in Booting State, all of them, I'll file a different ticket for that.

#19 - 02/07/2018 06:01 PM - Frank Li

Kefu Chai wrote:

[...]

the cluster formed a quorum of [0,1,2,3,4] since 18:02:21. and it was not in probing mode after the quorum was formed.

and two of the ones that responded are in probing mode and is considered out of the quorum.

how do you come to this conclusion?

When the ceph -s stopped responding on kaf101, I tried to run the mon_status command on the asok socket, which also hang. When I tried the same mon_status command on kaf401, it responded and showed that it was probing and unable to form a quorum.

```
[root@dl1-kaf401 frli]# ceph --admin-daemon /var/run/ceph/dl1 approd-mon.*.asok mon_status {
"name": "dl1-kaf401",
"rank": 4,
"state": "electing",
"election_epoch": 35553,
"quorum": [],
"features": {
"required_con": "153140804152475648",
"required_mon": [
"kraken",
"luminous"
],
"quorum_con": "2305244844532236283",
"quorum_mon": [
"kraken",
"luminous"
]
},
```

I find three out of the five would not respond to the mon_status command from the asok socket

did you sent mon_status to dl1 approd-mon.dl1-kaf101?

also i'd suggest you create another ticket to track this issue. as it's different from the one we've addressed in [#22847](#).

#20 - 02/07/2018 06:02 PM - Frank Li

I'll open a separate ticket to track the monitor not responding issue. the fix for the force-create-pg issue is good.

#21 - 02/07/2018 06:06 PM - Frank Li

<https://tracker.ceph.com/issues/22952>

ticket opened for ceph-mon not responding issue.

#22 - 02/09/2018 10:08 PM - Frank Li

Just FYI, Using this new patch, the leader ceph-mon will hung once it is up, and any kind of OSD command is ran, like ceph osd unset nodown, or bring up/down an OSD will cause the leader to hang. I've tried it on another 12.2.2 cluster. see <https://tracker.ceph.com/issues/22952> for details. So while this patch fixed the original issue of crashing, I'm still no able to get my cluster into a healthy state and risk losing all the data, are there any other way to bring the cluster back to working mode ??

#23 - 02/12/2018 08:49 AM - Kefu Chai

- Related to Bug #22952: Monitor stopped responding after awhile added

#24 - 02/14/2018 06:51 AM - Frank Li

just curious, I saw this patch got merged to the master branch and has the target version of 12.2.3, does that mean it'll make it into the official 12.2.3 ?

#25 - 02/14/2018 09:10 AM - Nathan Cutler

Frank Li wrote:

just curious, I saw this patch got merged to the master branch and has the target version of 12.2.3, does that mean it'll make it into the official 12.2.3 ?

The target version field on this issue was presumably filled in by you when you opened the bug.

The luminous backport is <https://github.com/ceph/ceph/pull/20399> and it is still pending. It will not be in 12.2.3.

#26 - 02/14/2018 05:48 PM - Frank Li

ok, I'll wait for 12.2.4 or a 12.2.3 + the patch then.

#27 - 02/24/2018 04:31 AM - Kefu Chai

- Related to deleted (Bug #22952: Monitor stopped responding after awhile)

#28 - 03/29/2018 01:24 PM - Nathan Cutler

- Status changed from Pending Backport to Resolved

Files

ceph.versions	518 Bytes	02/01/2018	Frank Li
ceph.osd.tree.log	5.46 KB	02/01/2018	Frank Li
ceph_health_detail.log	13.3 KB	02/01/2018	Frank Li
dl1approd-mon.dl1-kaf101.log.prob.1.gz	960 KB	02/06/2018	Frank Li
dl1approd-mon.dl1-kaf101.log.prob.2.gz	902 KB	02/06/2018	Frank Li