Ceph - Bug #13477

crush: crash if we see CRUSH_ITEM_NONE in early rule step

10/13/2015 01:43 PM - Sage Weil

Status:	Resolved	**% Done:**	0%
Priority:	Urgent	**Spent time:**	0.00 hour
Assignee:			
Category:			
Target version:			
Source:	Community (user)	**Affected Versions:**	
Tags:		**ceph-qa-suite:**	
Backport:	infernalis, hammer	**Pull request ID:**	
Regression:	No	**Crash signature (v1):**	
Severity:	3 - minor	**Crash signature (v2):**	
Reviewed:			

Description

see bugzilla https://bugzilla.redhat.com/show_bug.cgi?id=1231630

Description of problem:
Seeing Monitor Crash while creating erasure coded pool with wrong parameters

Version-Release number of selected component (if applicable):
ceph version 0.94.1 (e4bfad3a3c51054df7e537a724c8d0bf9be972ff)

How reproducible:
1/1

Steps to Reproduce:
1. Create a ec profile, with the below command.

ceph osd erasure-code-profile set myprofile plugin=lrc mapping=__DD__DD layers='[["_cDD_cDD", ""],["cDDD____", ""],["____cDDD", ""],]' ruleset-steps='[["choose", "datacenter", 3], ["chooseleaf", "osd", 0]]'

2. ceph osd pool create ecpool 12 12 erasure myprofile

Actual results: Monitoring is crashing.
BT:

```
ceph version 0.94.1 (e4bfad3a3c51054df7e537a724c8d0bf9be972ff)
 1: /usr/bin/ceph-mon() [0x901e52]
 2: (()+0xf130) [0x7f205a9ca130]
 3: (crush_do_rule()+0x291) [0x833501]
 4: (OSDMap::_pg_to_osds(pg_pool_t const&, pg_t, std::vector&lt;int, std::allocator&lt;int&gt; >
, int, unsigned int*) const+0xff) [0x77b76f]
 5: (OSDMap::_pg_to_up_acting_osds(pg_t const&, std::vector&lt;int, std::allocator&lt;int&gt; >
, int, std::vector&lt;int, std::allocator&lt;int&gt; >, int) const+0x104) [0x77be24]
 6: (PGMonitor::map_pg_creates()+0x268) [0x65b748]
 7: (PGMonitor::post_paxos_update()+0x25) [0x65bf35]
 8: (Monitor::refresh_from_paxos(bool*)+0x221) [0x575721]
 9: (Monitor::init_paxos()+0x95) [0x575ac5]
 10: (Monitor::preinit()+0x7f1) [0x57a881]
 11: (main()+0x24a1) [0x54d881]
 12: (__libc_start_main()+0xf5) [0x7f20593d0af5]
 13: /usr/bin/ceph-mon() [0x55d0f9]
 NOTE: a copy of the executable, or `objdump -rdS &lt;executable&gt;` is needed to interpret this.
```

Related issues:

Copied to Ceph - Backport #13654: crush: crash if we see CRUSH_ITEM_NONE in e...	**Resolved**

Copied to Ceph - Backport #13655: crush: crash if we see CRUSH_ITEM_NONE in e...	**Resolved**

Associated revisions

Revision 976a24a3 - 10/28/2015 12:55 AM - Sage Weil

crush/mapper: ensure bucket id is valid before indexing buckets array

We were indexing the buckets array without verifying the index was within
the [0,max_buckets) range. This could happen because a multistep rule
does not have enough buckets and has CRUSH_ITEM_NONE
for an intermediate result, which would feed in CRUSH_ITEM_NONE and
make us crash.

Fixes: #13477
Signed-off-by: Sage Weil <sage@redhat.com>

Revision 81d8aa14 - 10/30/2015 10:01 AM - Sage Weil

crush/mapper: ensure bucket id is valid before indexing buckets array

We were indexing the buckets array without verifying the index was within
the [0,max_buckets) range. This could happen because a multistep rule
does not have enough buckets and has CRUSH_ITEM_NONE
for an intermediate result, which would feed in CRUSH_ITEM_NONE and
make us crash.

Fixes: #13477
Signed-off-by: Sage Weil <sage@redhat.com>
(cherry picked from commit 976a24a326da8931e689ee22fce35feab5b67b76)

Revision ecb6aa23 - 11/18/2015 07:11 AM - Sage Weil

crush/mapper: ensure bucket id is valid before indexing buckets array

We were indexing the buckets array without verifying the index was within
the [0,max_buckets) range. This could happen because a multistep rule
does not have enough buckets and has CRUSH_ITEM_NONE
for an intermediate result, which would feed in CRUSH_ITEM_NONE and
make us crash.

Fixes: #13477
Signed-off-by: Sage Weil <sage@redhat.com>
(cherry picked from commit 976a24a326da8931e689ee22fce35feab5b67b76)

History

#1 - 10/13/2015 01:43 PM - Sage Weil

- Backport changed from hammer to infernalis, hammer

#2 - 10/13/2015 01:44 PM - Sage Weil

Good news is that latest code prevents the rule from being created in the first place (yay crushtool check). But crush shouldn't crash.

#3 - 10/13/2015 01:47 PM - Sage Weil

- File cm added

sample crush map attached

#4 - 10/13/2015 01:58 PM - Sage Weil

- Status changed from 12 to Fix Under Review

https://github.com/ceph/ceph/pull/6246

#5 - 10/27/2015 02:40 AM - Loïc Dachary

- Backport changed from infernalis, hammer to infernalis, hammer, firefly

#6 - 10/27/2015 02:40 AM - Loïc Dachary

- Assignee set to Sage Weil

#7 - 10/30/2015 06:45 AM - Sage Weil

- Status changed from Fix Under Review to Pending Backport

- Assignee deleted (Sage Weil)

#8 - 10/30/2015 08:54 AM - Nathan Cutler

- Copied to Backport #13653: crush: crash if we see CRUSH_ITEM_NONE in early rule step added

#9 - 10/30/2015 08:55 AM - Nathan Cutler

- Copied to Backport #13654: crush: crash if we see CRUSH_ITEM_NONE in early rule step added

#10 - 10/30/2015 08:55 AM - Nathan Cutler

- Copied to Backport #13655: crush: crash if we see CRUSH_ITEM_NONE in early rule step added

#11 - 01/28/2016 12:20 PM - Loïc Dachary

- Backport changed from infernalis, hammer, firefly to infernalis, hammer

#12 - 01/29/2016 03:54 AM - Loïc Dachary

- Copied to deleted (Backport #13653: crush: crash if we see CRUSH_ITEM_NONE in early rule step)

#13 - 02/08/2016 05:06 AM - Loïc Dachary

- Status changed from Pending Backport to Resolved

Files

cm	589 Bytes	10/13/2015	Sage Weil