

Ceph - Bug #17412

Applying ruleset halts monitor

09/27/2016 09:07 AM - Arvydas Opulskis

Status:	Duplicate	Start date:	09/27/2016
Priority:	Low	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:	other	Reviewed:	
Tags:		Affected Versions:	v10.2.2
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	

Description

Hi,

I noticed strange behavior of monitors during changing pool's crush_ruleset parameter. In some cases monitor halts and if command is not canceled (ctrl+c) other monitors goes down after some time too. After some investigation I've found what caused this:

1. Monitor halts only when applying last rule in crush map
2. This happens only if rules are not numbered in sequence

I had this layout in my crush map:

```
rule default {
  ruleset 0
  ...
}

rule xxx-sata {
  ruleset 2
  ...
}

rule sata {
  ruleset 3
  ...
}

rule ssd {
  ruleset 4
  ...
}

rule sas {
  ruleset 5
  ...
}
```

So, when applying "ceph osd pool set rbd crush_ruleset 5" I got monitor halt.

After changing numbering of rulesets in crush map from 0 to 4, there were no issues assigning ruleset 4 (sas) to rbd pool.

I also tried to set not existing ruleset (6) to pool, but then I got error about not known ruleset, which is ok.

I guess "ceph osd pool set" command checks if that rule exist, but when applying it, it ignores ruleset number from crushmap and

searches it by sequence.

I've tested it on Jewel 10.2.2 version.

Related issues:

Duplicates Ceph - Bug #16653: ceph mon Segmentation fault after set crush_rul...

Resolved

History**#1 - 09/27/2016 09:31 AM - Igor Podoski**

This could be duplicate of <http://tracker.ceph.com/issues/16653>

@Arvydas, please check your monitor log, if you're seeing something similar to <http://tracker.ceph.com/issues/16653#note-7>, it's for sure a duplicate.

I can see the same behavior on 10.2.2 at my cluster.

#2 - 05/24/2017 03:40 PM - Sage Weil

- *Status changed from New to Duplicate*

#3 - 05/24/2017 03:40 PM - Sage Weil

- *Duplicates Bug #16653: ceph mon Segmentation fault after set crush_ruleset ceph 10.2.2 added*