

Ceph - Bug #10353

the number of tries should be higher for the default erasure coded ruleset

12/17/2014 07:03 AM - Loïc Dachary

Status: Resolved	% Done: 80%
Priority: Normal	Spent time: 0.00 hour
Assignee: Loïc Dachary	
Category: OSD	
Target version:	
Source: other	Affected Versions:
Tags:	ceph-qa-suite:
Backport: giant,firefly	Pull request ID:
Regression: No	Crash signature (v1):
Severity: 4 - irritation	Crash signature (v2):
Reviewed:	

Description

It is common for people to try to map 9 OSDs out of a 9 OSDs total ceph cluster. The default tries (50) is often not enough and 200 should be enough. The ruleset created by default should have

```
step set_choose_tries 200
```

to address this case.

Associated revisions

Revision 2f87ac80 - 01/15/2015 08:26 PM - Loïc Dachary

crush: set_choose_tries = 100 for erasure code rulesets

It is common for people to try to map 9 OSDs out of a 9 OSDs total ceph cluster. The default tries (50) will frequently lead to bad mappings for this use case. Changing it to 100 makes no significant CPU performance difference, as tested manually by running crushtool on one million mappings.

<http://tracker.ceph.com/issues/10353> Fixes: #10353

Signed-off-by: Loïc Dachary <ldachary@redhat.com>

Revision 1e58bb49 - 02/27/2015 04:18 PM - Loïc Dachary

crush: set_choose_tries = 100 for erasure code rulesets

It is common for people to try to map 9 OSDs out of a 9 OSDs total ceph cluster. The default tries (50) will frequently lead to bad mappings for this use case. Changing it to 100 makes no significant CPU performance difference, as tested manually by running crushtool on one million mappings.

<http://tracker.ceph.com/issues/10353> Fixes: #10353

Signed-off-by: Loic Dachary <ldachary@redhat.com>
(cherry picked from commit 2f87ac807f3cc7ac55d9677d2051645bf5396a62)

Conflicts:
src/erasure-code/lrc/ErasureCodeLrc.cc
safely ignored because the file does not exist

Revision 5fca2320 - 03/17/2015 10:16 AM - Loic Dachary

crush: set_choose_tries = 100 for erasure code rulesets

It is common for people to try to map 9 OSDs out of a 9 OSDs total ceph cluster. The default tries (50) will frequently lead to bad mappings for this use case. Changing it to 100 makes no significant CPU performance difference, as tested manually by running crushtool on one million mappings.

<http://tracker.ceph.com/issues/10353> Fixes: #10353

Signed-off-by: Loic Dachary <ldachary@redhat.com>
(cherry picked from commit 2f87ac807f3cc7ac55d9677d2051645bf5396a62)

History

#1 - 12/17/2014 07:04 AM - Loic Dachary

- Backport set to *giant,firefly*

#2 - 12/17/2014 07:36 AM - Loic Dachary

- Status changed from 12 to *Fix Under Review*

- % Done changed from 0 to 80

<https://github.com/ceph/ceph/pull/3194>

#3 - 02/09/2015 02:31 PM - Loic Dachary

- Status changed from *Fix Under Review* to *Pending Backport*

#4 - 03/10/2015 08:22 AM - Loic Dachary

- Tracker changed from *Fix* to *Bug*

- Severity set to 4 - *irritation*

#5 - 03/17/2015 05:48 PM - Sage Weil

giant backport done, still waiting on firefly

#6 - 03/26/2015 03:36 PM - Loic Dachary

[5fca232](#) crush: set_choose_tries = 100 for erasure code rulesets (in *giant*), [1e58bb4](#) crush: set_choose_tries = 100 for erasure code rulesets (in *firefly*),

#7 - 03/26/2015 03:36 PM - Loïc Dachary

- Status changed from Pending Backport to Resolved