

Ceph - Bug #12782

Straw2 bucket created even when "tunable straw_calc_version 1" is created

08/25/2015 06:09 PM - James P

Status:	Closed	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:			
Category:			
Target version:			
Source:	other	Affected Versions:	v0.94.2
Tags:		ceph-qa-suite:	
Backport:		Pull request ID:	
Regression:	No	Crash signature (v1):	
Severity:	2 - major	Crash signature (v2):	
Reviewed:			

Description

- Hammer cluster running on kernel 3.19

- Tunables set to "hammer" via "ceph osd crush tunables hammer"

- Tunables dump:

```
{
  "choose_local_tries": 0,
  "choose_local_fallback_tries": 0,
  "choose_total_tries": 50,
  "chooseleaf_descend_once": 1,
  "chooseleaf_vary_r": 1,
  "straw_calc_version": 1,
  "allowed_bucket_algs": 54,
  "profile": "hammer",
  "optimal_tunables": 0,
  "legacy_tunables": 0,
  "require_feature_tunables": 1,
  "require_feature_tunables2": 1,
  "require_feature_tunables3": 1,
  "has_v2_rules": 0,
  "has_v3_rules": 0,
  "has_v4_buckets": 0
}
```

- Create a new bucket via "ceph osd crush add-bucket ssdroot root"

- dump crush map:

---snip---

```
root ssdroot {
  id -9      # do not change unnecessarily

  1. weight 0.000
    alg straw2
    hash 0    # rjenkins1
}
```

---snip---

This prevents client nodes from mapping RBDs with the following error:

"libceph: mon2 172.16.1.173:6789 feature set mismatch, my 2b84a842a42 < server's 102b84a842a42, missing 1000000000000"

History

#1 - 08/25/2015 06:31 PM - James P

Apparently "allowed_bucket_algs" is the one which dictates that buckets should be STRAW2 and not "straw_calc_version".

So this could probably be just a documentation issue. At the very least, an entry that states the minimum kernel version for hammer/optimal tunables which, if I understand correctly, is 4.1 to support STRAW2.

#2 - 09/08/2015 02:44 PM - Sage Weil

- *Status changed from New to Need More Info*

Yeah. Can you point to where in the docs it is confusing? Thanks!

#3 - 04/12/2017 05:38 PM - Sage Weil

- *Status changed from Need More Info to Closed*