

Ceph - Feature #15910

Increase the default value of mon_osd_min_in_ratio

05/17/2016 05:59 PM - David Zafman

Status: Resolved	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee: David Zafman	
Category:	
Target version:	
Source: other	Reviewed:
Tags:	Affected Versions:
Backport:	Pull request ID:
Description	
<p>Why don't we default the mon_osd_min_in_ratio value to something larger than 30%? I would suggest somewhere around 70% - 85%.</p> <p>If there are networking problems, it is just going to wreak havoc when too many OSDs are marked out. Data movement and running out of space on the remaining OSDs are just some examples.</p> <p>Upon installation of a Ceph cluster the value for this config should be carefully considered. For example, a 3 rack installation may be able to tolerate 1 rack getting disconnected, but at that point further OSD failures should be ignored. This might mean that a value of 66% would be appropriate for this value.</p> <p>Another idea would be to have OSDMonitor::can_mark_out() consider fullness of the cluster as another criteria. If there are full OSDs or some other calculation based on losing a particular OSD, it could` return false even if above the mon_osd_min_in_ratio.</p> <p>The general problem of running out of space due to OSD failures has been around for a long time.</p>	
Related issues:	
Related to Ceph - Feature #2911: osd: Restrict recovery when the OSD full lis...	Duplicate 08/06/2012
Related to Ceph - Bug #15912: An OSD was seen getting ENOSPC even with osd_fa...	Resolved 05/17/2016

History

#1 - 05/17/2016 05:59 PM - David Zafman

- Related to Feature #2911: osd: Restrict recovery when the OSD full list is nonempty added

#2 - 05/17/2016 06:23 PM - David Zafman

- Assignee set to David Zafman

#3 - 03/10/2017 01:48 AM - David Zafman

- Related to Bug #15912: An OSD was seen getting ENOSPC even with osd_failsafe_full_ratio passed added

#4 - 03/10/2017 01:49 AM - David Zafman

- Status changed from New to Resolved

Fix included in fix for 15912

<https://github.com/ceph/ceph/pull/13425>