

RADOS - Bug #24687

Automatically set expected_num_objects for new pools with >=100 PGs per OSD

06/28/2018 02:49 PM - Douglas Fuller

Status:	Resolved	Start date:	06/28/2018
Priority:	Normal	Due date:	
Assignee:	Douglas Fuller	% Done:	0%
Category:	Administration/Usability	Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:	mimic,luminous	ceph-qa-suite:	
Regression:	No	Component(RADOS):	Monitor, OSD
Severity:	2 - major	Pull request ID:	
Description			
Field experience has demonstrated significant performance impact from filestore split and merge activity. The expected_num_objects hint to ceph osd pool create mitigates this impact. Assume that pools created with 50 or more PGs per OSD are expected to store a large number of objects and pre-split their collections automatically.			
Related issues:			
Blocked by RADOS - Bug #24686: change default filestore_merge_threshold to -10		Resolved	06/28/2018
Blocks RADOS - Cleanup #25124: Add message to consult pgcalc for expected_num...		New	07/26/2018
Copied to RADOS - Backport #25144: mimic: Automatically set expected_num_obje...		Resolved	
Copied to RADOS - Backport #25145: luminous: Automatically set expected_num_o...		Resolved	

History

#1 - 06/28/2018 02:50 PM - Douglas Fuller

- Blocked by Bug #24686: change default filestore_merge_threshold to -10 added

#2 - 06/28/2018 02:50 PM - Douglas Fuller

- Subject changed from Automatically set expected_num_objects for new pools with >50 PGs per OSD to Automatically set expected_num_objects for new pools with >=50 PGs per OSD

#3 - 06/29/2018 05:59 PM - Douglas Fuller

Also include >1024 PGs overall

#4 - 07/03/2018 10:45 AM - Nathan Cutler

- Backport changed from luminous to mimic,luminous

#5 - 07/16/2018 02:24 PM - Douglas Fuller

- Status changed from In Progress to Need Review

Because a value for expected_num_objects is too difficult to determine automatically, instead we print a suggestion that the user consult pgcalc.

#6 - 07/16/2018 02:24 PM - Douglas Fuller

<https://github.com/ceph/ceph/pull/23072>

#7 - 07/17/2018 05:39 PM - Neha Ojha

- Subject changed from Automatically set expected_num_objects for new pools with >=50 PGs per OSD to Automatically set expected_num_objects for new pools with >=100 PGs per OSD

- Status changed from Need Review to Testing

#8 - 07/26/2018 05:17 PM - Douglas Fuller

Removed pgcalc message while pgcalc updates are considered

#9 - 07/26/2018 05:20 PM - Douglas Fuller

- Blocks Cleanup #25124: Add message to consult pgcalc for expected_num_objects added

#10 - 07/26/2018 05:47 PM - Douglas Fuller

- Status changed from Testing to Pending Backport

#11 - 07/28/2018 09:11 AM - Nathan Cutler

- Copied to Backport #25144: mimic: Automatically set expected_num_objects for new pools with ≥ 100 PGs per OSD added

#12 - 07/28/2018 09:11 AM - Nathan Cutler

- Copied to Backport #25145: luminous: Automatically set expected_num_objects for new pools with ≥ 100 PGs per OSD added

#13 - 10/05/2018 02:34 AM - Nathan Cutler

- Status changed from Pending Backport to Resolved