

RADOS - Feature #22448

Visibility for snap trim queue length

12/15/2017 08:11 AM - Piotr Dalek

Status: Resolved	% Done: 0%
Priority: High	Spent time: 0.00 hour
Assignee:	
Category: Snapshots	
Target version:	
Source:	Affected Versions:
Tags:	Component(RADOS):
Backport: jewel, luminous	Pull request ID:
Reviewed:	
Description We observed unexplained, constant disk space usage increase on a few of our prod clusters. At first we thought that it's because of customers abusing them, but that wasn't it. Then we thought that images are constantly filled with data, but space usage reported by Ceph wasn't consistent with filesystem. After further digging, we realized that snap trim queues for some of PGs are in 250k elements territory... We increased the snap trimmer frequency and number of parallel snap trim ops and disk space usage finally started to drop. Ceph needs a features to efficiently and conveniently access snap trim queue lengths so it can be used with monitoring, and a features to warn Ceph cluster admins when snap trim queues are long enough to be requiring some attention. https://github.com/ceph/ceph/pull/19520	
Related issues:	
Copied to RADOS - Backport #22449: jewel: Visibility for snap trim queue length	Resolved
Copied to RADOS - Backport #22450: luminous: Visibility for snap trim queue l...	Resolved

History

#1 - 12/15/2017 08:13 AM - Piotr Dalek

- Copied to Backport #22449: jewel: Visibility for snap trim queue length added

#2 - 12/15/2017 08:16 AM - Piotr Dalek

- Copied to Backport #22450: luminous: Visibility for snap trim queue length added

#3 - 12/15/2017 12:06 PM - Nathan Cutler

@Piotr: Please wait until the master PR is merged before starting the backporting process. Thanks.

#4 - 12/15/2017 12:07 PM - Nathan Cutler

- Status changed from New to Fix Under Review

master PR is <https://github.com/ceph/ceph/pull/19520>

#5 - 12/15/2017 12:08 PM - Nathan Cutler

- Backport set to jewel, luminous

@Piotr: It's OK to add e.g. "jewel, luminous" to the "Backport" field right from the beginning, though.

When the master PR is merged, the status of the ticket is changed to "Pending Backport" and then an automated script automatically creates the backport tickets from the value of the "Backport" field.

#6 - 12/15/2017 12:10 PM - Piotr Dalek

@Nathan: yeah, sorry, I thought this process is more manual.

#7 - 02/01/2018 11:40 AM - Nathan Cutler

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

#8 - 05/15/2018 12:29 PM - Piotr Dalek

- Status changed from *Pending Backport* to *Resolved*

Already merged to master, luminous and jewel.