

rbd - Feature #21305

Just discard changed data since snapshot in "rbd rollback" command

09/08/2017 07:55 AM - Марк Коренберг

Status: Rejected	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee: Jason Dillaman	
Category:	
Target version:	
Source:	Reviewed:
Tags:	Affected Versions:
Backport:	Pull request ID:
Description	
Rolling back images is painfully slow. Yes, I know, "rbd clone", but this creates another image and automatic-snapshotting-backuppers get failed.	
Why not to bring image to the state of specified snapshot ?	

History

#1 - 12/05/2017 02:00 PM - Jason Dillaman

- Project changed from Ceph to rbd

- Subject changed from Just discard changed data since snapshot in "rbd roollback" command to Just discard changed data since snapshot in "rbd rollback" command

Only possible if you are rolling back to the most recent snapshot. If you are rolling back to an older snapshot, you would need to copy the older data to the HEAD revision instead of just deleting the HEAD revision to expose the most recent snapshot.

#2 - 12/05/2017 02:07 PM - Марк Коренберг

Yes, ability to do FAST rollback to latest snapshot will be huge improvement, IMHO.

#3 - 12/15/2017 01:32 AM - Jason Dillaman

- Status changed from New to In Progress

- Assignee set to Jason Dillaman

#4 - 12/15/2017 02:43 AM - Jason Dillaman

- Status changed from In Progress to Rejected

I quickly hacked this up and it turns out the OSDs don't actually allow you to do this. The low-level, internal transactions within the OSD will still clone the HEAD object and then "delete" it, which results in a zeroed object at HEAD.

#5 - 12/15/2017 06:16 AM - Марк Коренберг

I don't understand cause why this would not be implemented. Maybe open this issue until someone be able to fix it ?

#6 - 12/15/2017 02:08 PM - Jason Dillaman

There is nothing RBD can do -- you can perhaps open a new ticket against RADOS to optimize snap rollback directly.