

rgw - Bug #10365

radosgw-agent: buckets to retry logged by full sync won't be fully synced if they were uploaded before radosgw logging was enabled in the source zone

12/17/2014 07:46 PM - Josh Durgin

Status:	Closed	% Done:	0%
Priority:	High	Spent time:	0.00 hour
Assignee:	Yehuda Sadeh		
Category:			
Target version:			
Source:	Development	Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:		Pull request ID:	
Regression:		Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			
Description			
<p>When buckets to retry are handled by incremental sync, it uses the bucket index log to figure out which objects within the bucket need syncing. This does not include objects uploaded before logging was enabled in the zone.</p> <p>One fix would be to store items to retry as tuples "[full incremental], \$item", and assume existing items to retry are from full sync for maximum coverage.</p> <p>Suggested scheme:</p> <p>sync initialization</p> <p>-----</p> <p>for each shard:</p> <ul style="list-style-type: none">• read marker and buckets to retry for shard's replica log<ul style="list-style-type: none">- if no marker, check shard_id->buckets indexif that doesn't exist yet, list all buckets and create shard_id -> buckets index• for each bucket in current shard<ul style="list-style-type: none">- if bucket doesn't have entry in replica log: add to replica log, mark for full sync- if bucket exists in replica log, go to (incremental sync -> check type) <p>incremental sync:</p> <ul style="list-style-type: none">• look up data log shard in replica log to get buckets that need retrying, and current marker• check type of sync for any buckets that need to be retried, and generate the list of objects based on that<ul style="list-style-type: none">- full -> list all objects in bucket- incremental -> read data log as usual• for each bucket<ul style="list-style-type: none">- list objects in bucket that need to be synced based on sync type, from full list, or data log- read the bucket instance log for the bucket starting at marker from replica log- sync object using the same method as full sync- if syncing an object fails, add it to a list to retry- update bucket instance replica log with last marker read and list of objects to retry• once data log shard is done, update replica log for that shard with<ul style="list-style-type: none">- new marker- list of bucket instances to retry <p>May need to be careful about updating with an empty marker, if e.g. lots of objects were uploaded before data logging was enabled. Perhaps use ' ' as the marker in that case, since it's before all markers the gateway will generate.</p>			

History

#1 - 01/20/2015 08:59 AM - Alfredo Deza

- *Description updated*

#2 - 01/20/2015 09:01 AM - Alfredo Deza

- *Description updated*

#3 - 04/07/2015 06:09 PM - Yehuda Sadeh

- *Assignee set to Yehuda Sadeh*

#4 - 03/23/2016 06:34 PM - Yehuda Sadeh

- *Status changed from New to Closed*

Fixed in v2