

rgw - Bug #19219

Data sync problem in multisite

03/08/2017 03:12 AM - zhang sw

Status:	Resolved	Start date:	03/08/2017
Priority:	High	Due date:	
Assignee:	Casey Bodley	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor		
Description			
<p>In zone A and zone B, if we upload a object obj to A. The data sync process is :</p> <ol style="list-style-type: none">1. B get the bi log from A;2. B fetch the obj from A according to the bi log, and then add a new bi log locally;3. A get the bi log from B.4. A fetch the obj from B according to the bi log, and B will return 304. <p>The process 4 isn't needed. Besides, some strange problems may occur because of that process. For example, users upload obj and delete it: In zone A, users have deleted obj, but the delete operation hasn't been synced; Then according to process 4, A will fetch the obj from B. Because A's obj has been deleted, B will return the obj's data. Then A will add a new bi log that obj is created and sync this bi log to B. Now B will have one delete operation followed by one put operation. If these two operations are in the squash map, The put will override delete. At last, the obj will not be deleted.</p>			

History

#1 - 04/19/2017 03:53 PM - Josh Durgin

- Project changed from Ceph to rgw

#2 - 04/20/2017 06:14 PM - Casey Bodley

- Status changed from New to Need Review

- Assignee set to Casey Bodley

<https://github.com/ceph/ceph/pull/13851>

#3 - 04/26/2017 05:47 PM - Yehuda Sadeh

- Priority changed from Normal to High

#4 - 06/06/2017 06:28 PM - Yehuda Sadeh

- Status changed from Need Review to Resolved