

rgw - Bug #17698

multisite: ECANCELED & 500 error on bucket delete

10/25/2016 04:28 PM - Abhishek Lekshmanan

Status:	Resolved	Start date:	10/25/2016
Priority:	High	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:	jewel	ceph-qa-suite:	
Regression:	No	Release:	
Severity:	3 - minor	Needs Doc:	No

Description

sometimes on bucket delete at the secondary, we're seeing secondary returning 500 on a bucket delete (caused by reraising ECANCELED from OSDs). However the bucket in question gets deleted actually and subsequent client requests fail with a 404

```
2016-10-25 13:59:21.999214 7f7a48417700 1 -- 127.0.0.1:0/1778501402 <== osd.0 127.0.0.1:6812/2741
2 2962 ==== osd_op_reply(6705 bucket26 [call,call,delete] v0'0 uv0 ondisk = -125 ((125) Operation
canceled)) v7 ==== 212+0+0 (1575458175 0 0) 0x7f7a20013ba0 con 0x564acad41190
2016-10-25 13:59:21.999244 7f7920ff1700 2 req 1192:0.312583:s3:DELETE /bucket26/:delete_bucket:co
mpleting
2016-10-25 13:59:21.999247 7f7920ff1700 0 WARNING: set_req_state_err err_no=125 resorting to 500
2016-10-25 13:59:21.999286 7f7920ff1700 2 req 1192:0.312625:s3:DELETE /bucket26/:delete_bucket:op
status=-125
2016-10-25 13:59:21.999289 7f7920ff1700 2 req 1192:0.312628:s3:DELETE /bucket26/:delete_bucket:ht
tp status=500
2016-10-25 13:59:21.999291 7f7920ff1700 1 ===== req done req=0x7f7920fee7e0 op status=-125 http_
status=500 =====
```

Related issues:

Copied to rgw - Backport #17886: jewel: multisite: ECANCELED & 500 error on b...

Resolved

History

#1 - 10/25/2016 04:35 PM - Abhishek Lekshmanan

- File *rgw-bucket26.log* added

Adding the rgw log of the relevant bucket, it seems we're issuing delete on the bucket object twice, once at

```
2016-10-25 13:59:21.931712 7f7a01ffb700 1 -- 127.0.0.1:0/1778501402 --> 127.0.0.1:6812/27412 -- osd_op(client.4111.0:6691 3.c8ef106b
bucket26 [call version.check_conds,call version.set,delete] snapc 0=[] ondisk+write+known_if_redirected e29) v7 -- ?+0 0x7f79e0018850 con
0x564acad41190
```

(which is probably mdlog/admin log?) and second when the secondary itself processes the delete.

#2 - 10/25/2016 04:59 PM - Yehuda Sadeh

Sounds like secondary is racing with the request forwarded from the master (which the secondary itself initiated). We should handle that ECANCELED gracefully.

#3 - 10/25/2016 05:04 PM - Casey Bodley

proposed fix at <https://github.com/ceph/ceph/pull/11648>

#4 - 11/01/2016 06:00 PM - Yehuda Sadeh

- Priority changed from Normal to High

#5 - 11/10/2016 06:46 PM - Casey Bodley

- Status changed from New to Pending Backport

#6 - 11/13/2016 08:47 AM - Nathan Cutler

- Source deleted (other)

- Backport set to jewel

#7 - 11/13/2016 12:26 PM - Nathan Cutler

- Copied to Backport #17886: jewel: multisite: ECANCELED & 500 error on bucket delete added

#8 - 06/19/2017 04:58 PM - Nathan Cutler

- Status changed from Pending Backport to Resolved

Files

rgw-bucket26.log

108 KB

10/25/2016

Abhishek Lekshmanan