

rgw - Bug #11563

RadosGW regression: COPY doesn't preserve Content-Type

05/07/2015 03:24 PM - Sylvain Munaut

Status:	Resolved	% Done:	0%
Priority:	Urgent	Spent time:	0.00 hour
Assignee:	Orit Wasserman		
Category:			
Target version:			
Source:	Community (user)	Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	hammer	Pull request ID:	
Regression:	Yes	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			

Description

- fix in master <https://github.com/ceph/ceph/pull/4995>

When creating an object with an explicit content-type then copying it to another key, the returned content-type is binary/octet-stream instead of what was set.

That definitely used to work fine in 0.87, but is now broken in 0.94. (Found it when running our own regression suite before updating prod)

Related issues:

Related to rgw - Bug #11639: Object copy bug	Resolved	05/15/2015
Related to rgw - Bug #12370: test_s3.test_object_copy_canned_acl ... FAIL	Resolved	07/16/2015
Related to rgw - Bug #13015: rgw: intra region copy does not preserve acl	Resolved	09/09/2015
Copied to rgw - Backport #12199: RadosGW regression: COPY doesn't preserve Co...	Resolved	05/07/2015

Associated revisions

Revision e41d97c8 - 06/17/2015 06:46 PM - Yehuda Sadeh

rgw: fix assignment of copy obj attributes

Fixes: #11563

Clarify the confusing usage of set_copy_attrs() by switching the source and destination params (attrs, src_attrs). Switch to use attrs instead of src_attrs afterwards. In one of the cases we originally used the wrong variable.

Signed-off-by: Yehuda Sadeh <yehuda@redhat.com>

Revision 82ea02ab - 09/09/2015 10:29 PM - Yehuda Sadeh

rgw: fix assignment of copy obj attributes

Fixes: #11563

Clarify the confusing usage of set_copy_attrs() by switching the source and destination params (attrs, src_attrs). Switch to use attrs instead of src_attrs afterwards. In one of the cases we originally used the wrong

variable.

Signed-off-by: Yehuda Sadeh <yehuda@redhat.com>
(cherry picked from commit e41d97c8e38bb60d7e09e9801c0179efe7af1734)

History

#1 - 05/19/2015 08:32 PM - Samuel Just

- Project changed from Ceph to rgw
- Category deleted (22)
- Affected Versions deleted (v0.94)

#2 - 06/03/2015 02:39 PM - Sylvain Munaut

Any news on this ?

Given that COPY is used when renaming objects (like if you do a streaming upload but you want the name of the object to be a sha256 of the name), it's pretty annoying to loose the metadata ...

#3 - 06/16/2015 06:16 PM - Sage Weil

- Priority changed from Normal to Urgent
- Backport set to hammer

#4 - 06/16/2015 06:21 PM - Sage Weil

- Source changed from other to Community (user)

Might be fixed by [#11639](#)

#5 - 06/16/2015 06:30 PM - Abhishek Lekshmanan

Is this in the swift or s3 API? In swift API there was a backport to hammer for COPY that *may* be the cause for this (<https://github.com/ceph/ceph/pull/4568>)

#6 - 06/16/2015 11:51 PM - Yehuda Sadeh

- Status changed from New to Can't reproduce

I cannot reproduce it on latest master.

#7 - 06/16/2015 11:51 PM - Yehuda Sadeh

also can't reproduce it on hammer

#8 - 06/17/2015 03:04 PM - Sylvain Munaut

- File `issue-11563.py` added

I just updated my build to the latest commit of the 'hammer' branch (3e8d60a80ce31860eac76a1f6489a35e1795a0c0) and I still have the issue.

See the attached python script to reproduce (using boto, edit the top part with credentials and host/port).

And this is all using the S3 API.

#9 - 06/17/2015 03:35 PM - Sylvain Munaut

Also tried pulling the commit for [#11639](#) into my branch, didn't change anything. Also that commit is only for same bucket copy, but the content-type (and other metadata) should be kept even when copying to other bucket. (unless new metadata are provided in the COPY call)

#10 - 06/17/2015 05:55 PM - Yehuda Sadeh

- Status changed from *Can't reproduce* to *12*

ok, managed to reproduce using your test. The key difference from my tests is that the object size also needs to be > 512k.

#11 - 06/17/2015 08:36 PM - Yehuda Sadeh

- Status changed from *12* to *Fix Under Review*

- Assignee set to *Orit Wasserman*

#12 - 06/18/2015 10:08 AM - Sylvain Munaut

I cherry picked `1a6c3e760b1e2c1f4698009e70fd6cd1515ec722` over the recent hammer branch (`3e8d60a80ce31860eac76a1f6489a35e1795a0c0`) but the behavior didn't change.

Do I need any other commits ?

#13 - 06/18/2015 05:55 PM - Yehuda Sadeh

I tested it again now and it works for me (tested commit:`e41d97c8e38bb60d7e09e9801c0179efe7af1734`). However, note that there's an issue with your test script, that when you check for the copy key, you're doing `b.new_key()` instead of `b.get_key()`.

#14 - 06/23/2015 12:28 PM - Sylvain Munaut

You're right my bad, the fix seems to work.

When running my full test suite, it now passes correctly with the fix. Seems I made a small mistake when writing the small .py test above and changing it to `get_key()` works.

#15 - 06/26/2015 04:48 PM - Yehuda Sadeh

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

#16 - 06/30/2015 05:56 PM - Yehuda Sadeh

- Assignee changed from *Orit Wasserman* to *Loïc Dachary*

#17 - 07/01/2015 10:26 PM - Loïc Dachary

- Description updated

#18 - 08/31/2015 02:10 PM - Loïc Dachary

- Assignee changed from *Loïc Dachary* to *Orit Wasserman*

Assigning Orit to record who was assigned the issue originally (I still care for the backport of course ;-))

#19 - 09/09/2015 09:43 AM - Sylvain Munaut

How is that not in 0.94.3 ?!?

Patch has been done for 3 month and it fixes a regression ...

#20 - 09/22/2015 08:10 PM - Yehuda Sadeh

Sorry for the time it takes. It fixed a regression but also introduced one, so we need to make sure everything works correctly now.

#21 - 10/12/2015 07:25 PM - Loïc Dachary

- *Status changed from Pending Backport to Resolved*

Files

issue-11563.py

1.44 KB

06/17/2015

Sylvain Munaut