

## rgw - Bug #22121

### rgw: S3 interface: X-Amz-Copy-Source must be URL-encoded

11/14/2017 09:50 AM - Nick Craig-Wood

<b>Status:</b>	Resolved	<b>% Done:</b>	0%
<b>Priority:</b>	High	<b>Spent time:</b>	0.00 hour
<b>Assignee:</b>	Matt Benjamin		
<b>Category:</b>			
<b>Target version:</b>			
<b>Source:</b>	Community (user)	<b>Reviewed:</b>	
<b>Tags:</b>		<b>Affected Versions:</b>	v11.2.1
<b>Backport:</b>	luminous	<b>ceph-qa-suite:</b>	rgw
<b>Regression:</b>	No	<b>Pull request ID:</b>	
<b>Severity:</b>	3 - minor	<b>Crash signature:</b>	
<b>Description</b>			
<p>According to the AWS docs on COPY: <a href="http://docs.aws.amazon.com/AmazonS3/latest/API/RESTObjectCOPY.html">http://docs.aws.amazon.com/AmazonS3/latest/API/RESTObjectCOPY.html</a> it states about X-Amz-Copy-Source "This string must be URL-encoded".</p> <p>However URL encoding it does not seem to work with CEPH.</p> <p>Here are some examples generated with rclone trying to copy an object called "complex character £100.txt" from one bucket to another.</p> <pre>2017/11/14 09:22:18 DEBUG : HTTP REQUEST (req 0xc420352100) 2017/11/14 09:22:18 DEBUG : PUT /utf8-test2/complex%20character%20%C2%A3100.txt HTTP/1.1 Host: storage.dunsfold-a.openstack.memset.com:8080 User-Agent: rclone/v1.38-107-gc03f558d-vfs-rw Content-Length: 0 Authorization: XXXX X-Amz-Content-Sha256: e3b0c44298fc1c149afb4c8996fb92427ae41e4649b934ca495991b7852b855 X-Amz-Copy-Source: utf8-test%2Fcomplex+character+%C2%A3100.txt X-Amz-Date: 20171114T092218Z X-Amz-Metadata-Directive: COPY Accept-Encoding: gzip</pre> <pre>@2017/11/14 09:22:18 DEBUG : HTTP RESPONSE (req 0xc420352100) 2017/11/14 09:22:18 DEBUG : HTTP/1.1 404 Not Found Content-Length: 232 Accept-Ranges: bytes Bucket: 659259d53fc74389b9d71602f8a59bf8%2Futf8-test2 Content-Type: application/xml Date: Tue, 14 Nov 2017 09:22:18 GMT X-Amz-Request-Id: tx000000000000005bdf86-005a0ab5ca-205796c-dunsfold-a</pre> <pre>&lt;?xml version="1.0" encoding="UTF-8"?&gt;&lt;Error&gt;&lt;Code&gt;NoSuchKey&lt;/Code&gt;&lt;BucketName&gt;utf8-test2&lt;/BucketName&gt;&lt;RequestId&gt;tx0000000000000000 5bdf86-005a0ab5ca-205796c-dunsfold-a&lt;/RequestId&gt;&lt;HostId&gt;205796c-dunsfold-a-dunsfold-a&lt;/HostId&gt;&lt;/Error&gt;</pre> <p>@</p> <p>However if I don't URL encode X-Amz-Copy-Source it does work</p> <pre>2017/11/14 09:38:06 DEBUG : HTTP REQUEST (req 0xc420128900) 2017/11/14 09:38:06 DEBUG : PUT /utf8-test3/complex%20character%20%C2%A3100.txt HTTP/1.1 Host: storage.dunsfold-a.openstack.memset.com:8080 User-Agent: rclone/v1.38-DEV Content-Length: 0 Authorization: XXXX X-Amz-Content-Sha256: e3b0c44298fc1c149afb4c8996fb92427ae41e4649b934ca495991b7852b855</pre>			

X-Amz-Copy-Source: utf8-test/complex character £100.txt  
X-Amz-Date: 20171114T093806Z  
X-Amz-Metadata-Directive: COPY  
Accept-Encoding: gzip

@2017/11/14 09:38:06 DEBUG : HTTP RESPONSE (req 0xc420128900)  
2017/11/14 09:38:06 DEBUG : HTTP/1.1 200 OK  
Content-Length: 183  
Bucket: 659259d53fc74389b9d71602f8a59bf8%2Futf8-test3  
Content-Type: application/xml  
Date: Tue, 14 Nov 2017 09:38:06 GMT  
X-Amz-Request-Id: tx000000000000005bed5c-005a0ab97e-2027f9e-dunsfold-a

<CopyObjectResult  
xmlns="http://s3.amazonaws.com/doc/2006-03-01/"><LastModified>2017-11-14T09:38:06.634Z</LastModified><ETag>b1946ac92492d2347c6235b4d2611184</ETag></CopyObjectResult>  
@

However when trying with AWS the first of these works and the second doesn't. Here is the failing request from AWS

2017/11/14 09:40:53 DEBUG : HTTP REQUEST (req 0xc4202f4c00)  
2017/11/14 09:40:53 DEBUG : PUT /utf8-test4/complex%20character%20%C2%A3100.txt HTTP/1.1  
Host: s3.amazonaws.com  
User-Agent: rclone/v1.38-DEV  
Content-Length: 0  
Authorization: XXXX  
X-Amz-Content-Sha256: e3b0c44298fc1c149afb4c8996fb92427ae41e4649b934ca495991b7852b855  
X-Amz-Copy-Source: utf8-test/complex character £100.txt  
X-Amz-Date: 20171114T094053Z  
X-Amz-Metadata-Directive: COPY  
Accept-Encoding: gzip

@2017/11/14 09:40:53 DEBUG : HTTP RESPONSE (req 0xc4202f4c00)  
2017/11/14 09:40:53 DEBUG : HTTP/1.1 404 Not Found  
Transfer-Encoding: chunked  
Content-Type: application/xml  
Date: Tue, 14 Nov 2017 09:40:53 GMT  
Server: AmazonS3  
X-Amz-Id-2: 0JisWvRIBmOU7tZrndAcVTenG81tdAj2EUCrfeFEKXdUDVONLI5QRFIEgQLg6roGBHa8QTZhJUA=  
X-Amz-Request-Id: D57DF1E9FFE93CAC

103  
<Error><Code>NoSuchKey</Code><Message>The specified key does not exist.</Message><Key>complex character  
??100.txt</Key><RequestId>D57DF1E9FFE93CAC</RequestId><HostId>0JisWvRIBmOU7tZrndAcVTenG81tdAj2EUCrfeFEKXdUD  
VONLI5QRFIEgQLg6roGBHa8QTZhJUA=</HostId></Error>  
0  
@

This leads me to believe that this is a deviation from the S3 spec by CEPH and hence a bug.

This bug affects rclone users in two ways, firstly that certain files can't be server side copied and secondly that these files can't have their metadata changed so the users can't update their modification time (which is held as a metadata key).

Originally from: <https://github.com/ncw/rclone/issues/586>

#### Related issues:

Copied to rgw - Backport #22893: luminous: rgw: S3 interface: X-Amz-Copy-Sour...

**Resolved**

#### History

#1 - 11/16/2017 07:12 PM - Matt Benjamin

- Status changed from New to Triaged
- Assignee set to Matt Benjamin
- Backport set to Luminous

**#2 - 11/16/2017 07:13 PM - Matt Benjamin**

(we won't fix kraken, does this reproduce on Luminous or above?)

**#3 - 11/16/2017 07:15 PM - Matt Benjamin**

- Priority changed from Normal to High

**#4 - 11/16/2017 07:15 PM - Matt Benjamin**

- Status changed from Triaged to In Progress

**#5 - 11/24/2017 09:48 PM - Nathan Cutler**

- Backport changed from Luminous to luminous

**#6 - 01/12/2018 03:17 PM - Matt Benjamin**

- Status changed from In Progress to Fix Under Review

<https://github.com/ceph/ceph/pull/19936>

**#7 - 02/01/2018 07:07 PM - Matt Benjamin**

- Status changed from Fix Under Review to Pending Backport

**#8 - 02/01/2018 07:15 PM - Matt Benjamin**

(luminous backport staged at <https://github.com/ceph/ceph/pull/20236>)

**#9 - 02/01/2018 11:44 PM - Nathan Cutler**

- Copied to Backport #22893: luminous: rgw: S3 interface: X-Amz-Copy-Source must be URL-encoded added

**#10 - 04/29/2018 07:51 AM - Nathan Cutler**

- Status changed from Pending Backport to Resolved