

## rgw - Bug #27217

### rgw: Swift interface: server side copy fails if object name contains `?`

08/24/2018 09:41 AM - Nick Craig-Wood

<b>Status:</b>	Resolved	<b>% Done:</b>	0%
<b>Priority:</b>	Normal	<b>Spent time:</b>	0.00 hour
<b>Assignee:</b>	Casey Bodley		
<b>Category:</b>			
<b>Target version:</b>			
<b>Source:</b>	Community (user)	<b>Affected Versions:</b>	v12.2.7
<b>Tags:</b>		<b>ceph-qa-suite:</b>	rgw
<b>Backport:</b>	luminous mimic nautilus	<b>Pull request ID:</b>	25498
<b>Regression:</b>	Yes	<b>Crash signature (v1):</b>	
<b>Severity:</b>	3 - minor	<b>Crash signature (v2):</b>	
<b>Reviewed:</b>			

#### Description

Here is a demonstration of the problem with rclone. This sequence used to work with previous versions of CEPH and the same code works with openstack swift.

Create the object with a ? in hello?hello.txt

```
$ echo "hello" | rclone rcat 'swift:container1/hello?hello.txt'
```

Check it is OK

```
$ rclone ls swift:container1
6 hello?hello.txt
```

Do a server side copy of hello?hello.txt to hello2?hello2.txt

```
rclone copyto -vv --dump bodies 'swift:container1/hello?hello.txt' 'swift:container1/hello2?hello2.txt'
[snip - see below for HTTP transaction]
2018/08/23 16:11:46 ERROR : hello?hello.txt: Failed to copy: object not found
```

See what happened in the container

```
$ rclone ls swift:container1
6 hello2
6 hello?hello.txt
```

So ceph/rgw truncated the file name of hello2?hello2.txt at the ? to get hello2.

Here is the COPY http transaction

```
2018/08/23 16:11:46 DEBUG : HTTP REQUEST (req 0xc000198f00)
```

```
2018/08/23 16:11:46 DEBUG : COPY /swift/v1/659259d53fc74389b9d71602f8a59bf8/container1/hello%3Fhello.txt HTTP/1.1
Host: storage.reading-a.openstack.memset.com:8080
User-Agent: rclone/v1.42-126-g577fcd31-fix-2464-swift-auth
Destination: container1/hello2?hello2.txt
X-Auth-Token: XXXX
Accept-Encoding: gzip
```

```
2018/08/23 16:11:46 DEBUG : HTTP RESPONSE (req 0xc000198f00)
2018/08/23 16:11:46 DEBUG : HTTP/1.1 201 Created
Access-Control-Allow-Headers: Origin,Content-Type,Cache-Control,Content-Language,Expires,Last-Modified,Pragma,X-Requested-With,User-Agent,If-Modified-Since,Accept,X-Auth-Token
Access-Control-Max-Age: 10800
Bucket: 659259d53fc74389b9d71602f8a59bf8%2Fcontainer1
Content-Type: text/plain; charset=utf-8
Date: Thu, 23 Aug 2018 15:11:46 GMT
Etag: b1946ac92492d2347c6235b4d2611184
Last-Modified: Thu, 23 Aug 2018 15:11:46 GMT
X-Copied-From: container1/hello%3Fhello.txt
X-Copied-From-Account: 659259d53fc74389b9d71602f8a59bf8
X-Copied-From-Last-Modified: Thu, 23 Aug 2018 15:10:19 GMT
X-Object-Meta-Mtime: 1535037019.524649738
X-Openstack-Request-Id: tx000000000000000007c89-005b7eceb2-236cc2a-reading-a
X-Trans-Id: tx000000000000000007c89-005b7eceb2-236cc2a-reading-a
Content-Length: 0
```

Rclone writes the destination header as Destination: container1/hello2?hello2.txt

In <https://developer.openstack.org/api-ref/object-store/> it states under the COPY method

Destination header string The container and object name of the destination object in the form of /container/object. You must UTF-8-encode and then URL-encode the names of the destination container and object before you include them in this header.

rclone is not URL encoding the ? in the Destination which is out of spec. However this used to work with previous versions of CEPH and works with OpenStack swift also.

If I change rclone to URL encode the Destination: header I get exactly the same results though:

```
2018/08/24 10:06:47 DEBUG : HTTP REQUEST (req 0xc0000f6300)
2018/08/24 10:06:47 DEBUG : COPY /swift/v1/659259d53fc74389b9d71602f8a59bf8/container1/hello%3Fhello.txt HTTP/1.1
Host: storage.reading-a.openstack.memset.com:8080
User-Agent: rclone/v1.42-DEV
Destination: container1/hello2%3Fhello2.txt
X-Auth-Token: XXXX
Accept-Encoding: gzip
```

```
2018/08/24 10:06:47 DEBUG : HTTP RESPONSE (req 0xc0000f6300)
2018/08/24 10:06:47 DEBUG : HTTP/1.1 201 Created
Access-Control-Allow-Headers: Origin,Content-Type,Cache-Control,Content-Language,Expires,Last-Modified,Pragma,X-Requested-With,User-Agent,If-Modified-Since,Accept,X-Auth-Token
Access-Control-Max-Age: 10800
Bucket: 659259d53fc74389b9d71602f8a59bf8%2Fcontainer1
Content-Type: text/plain; charset=utf-8
Date: Fri, 24 Aug 2018 09:06:47 GMT
Etag: b1946ac92492d2347c6235b4d2611184
Last-Modified: Fri, 24 Aug 2018 09:06:47 GMT
X-Copied-From: container1/hello%3Fhello.txt
X-Copied-From-Account: 659259d53fc74389b9d71602f8a59bf8
X-Copied-From-Last-Modified: Thu, 23 Aug 2018 15:10:19 GMT
```



Content-Type: application/xml  
Date: Wed, 09 Jan 2019 10:44:00 GMT  
X-Amz-Request-Id: tx0000000000000000081eff9-005c35d070-3449ba4-dunsfold-a

```
<?xml version="1.0" encoding="UTF-8"?><Error><Code>NoSuchKey</Code><BucketName>container1</BucketName><RequestId>tx0000000000000000081eff9-005c35d070-3449ba4-dunsfold-a</RequestId><HostId>3449ba4-dunsfold-a-dunsfold-a</HostId></Error>
```

**#4 - 06/03/2019 01:00 PM - Casey Bodley**

- Status changed from Fix Under Review to Pending Backport
- Backport changed from luminous mimic to luminous mimic nautilus

**#5 - 06/04/2019 10:32 AM - Nathan Cutler**

- Copied to Backport #40127: luminous: rgw: Swift interface: server side copy fails if object name contains `?` added

**#6 - 06/04/2019 10:33 AM - Nathan Cutler**

- Copied to Backport #40128: mimic: rgw: Swift interface: server side copy fails if object name contains `?` added

**#7 - 06/04/2019 10:33 AM - Nathan Cutler**

- Copied to Backport #40129: nautilus: rgw: Swift interface: server side copy fails if object name contains `?` added

**#8 - 06/25/2019 10:58 AM - Nathan Cutler**

- Pull request ID set to 25498

**#9 - 11/20/2019 10:20 AM - Nathan Cutler**

- Status changed from Pending Backport to Resolved

While running with --resolve-parent, the script "backport-create-issue" noticed that all backports of this issue are in status "Resolved" or "Rejected".