

rgw - Bug #45911

Cannot address buckets in different tenant

06/05/2020 10:10 PM - Chris Durham

Status: Triaged	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee:	
Category:	
Target version: v15.2.3	
Source: Community (user)	Reviewed:
Tags:	Affected Versions: v15.0.0, v15.2.1, v15.2.2
Backport:	ceph-qa-suite:
Regression: No	Pull request ID:
Severity: 2 - major	Crash signature:
Description	
Using release 15.2.2 on CentOS 8.1	
At: https://docs.ceph.com/docs/mimic/radosgw/multitenancy/	
There is a suggestion on some code that to address s3 buckets in a different tenant (your AWS creds are for a user in tenant t1, but the bucket you want to address is in tenant t2) you should use a ':' to separate the two, i.e. tenant2:bucketname.	
This does not work. Using the AWS CLI (boto3), when doing	
<pre>aws s3 ls s3://tenant2:bucketname</pre>	
I get an error:	
Invalid bucket name "tenant2:bucketname": Bucket name must match the regex "[a-zA-Z0-9.\-_{1,255}\$]" or be an ARN matching the regex "arn:(aws).*:s3:[a-z\0-9]+:[0-9]{12}:accesspoint[/:][a-zA-Z0-9_{1,63}\$]"	
(Again this is CentOS 8.1):	
In /usr/share/awcli/python/site-packages/botocore/handlers.py, there is a line:	
<pre>VALID_BUCKET = re.compile(r'^[a-zA-Z0-9.\-_{1,255}\$')</pre>	
If I change this regex appropriately to include a ':', then I can address the bucket with both the awcli as well as with boto3 code directly.	
It should be noted that in the document I mentioned above, the code shown is boto code, NOT boto3. Apparently with boto as opposed to boto3 there is not a problem.	
Without patching the regex as I show, there does not appear to be a way to address buckets in other tenants using python boto3 code which includes the awcli.	
AWS and the Ceph team need to coordinate a proper fix	

History

#1 - 06/05/2020 11:25 PM - Matt Benjamin

- Status changed from New to Triaged

Hi Chris,

This is a known issue, but as it's entirely in boto, we're not sure right now how to proceed. We've filed issues against upstream boto.

Matt

#2 - 06/06/2020 05:31 PM - Pritha Srivastava

Hi Chris,

Bucket name validation can be suppressed in boto3 using the following also:

```
s3client.meta.events.unregister('before-parameter-build.s3', validate_bucket_name)
```

Thanks,
Pritha

#3 - 06/08/2020 03:11 PM - Chris Durham

Pritha Srivastava wrote:

Hi Chris,

Bucket name validation can be suppressed in boto3 using the following also:

```
s3client.meta.events.unregister('before-parameter-build.s3', validate_bucket_name)
```

Thanks,
Pritha

That's great for boto3 code, but as far as I know, there is no config or environment variable I can use to do that for the cli without making the change I showed above, correct?

#4 - 06/11/2020 02:16 PM - Casey Bodley

Using : to separate bucket and tenant is just broken by design since it doesn't work with dns names.