

rgw - Bug #19830

radosgw-admin validate tier type while setting

05/03/2017 12:27 PM - Abhishek Lekshmanan

Status:	Need Review	Start date:	05/03/2017
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	v12.0.0
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor		

Description

One can get into difficult situations with multisite and incorrect tier-types
an eg would be:
on a secondary zone do

```
$ radosgw-admin zone modify --rgw-zone=us-west --tier-type=es --tier-config=endpoint=http://localhost:9200 ..  
$ radosgw-admin period update --commit # this succeeds
```

Now the radosgw's in this zone wouldn't start anymore and further period updates from rgw-admin will error out as sync module init fails

The only possible way to salvage this situation would be to go back to the master, modify zonegroup via infile to set the right tier type and go to the secondary and do a realm pull with keys etc, maybe input validation while doing a tier creation would be the easiest way out?

History

#1 - 05/04/2017 06:34 PM - Yehuda Sadeh

The command here in itself is not invalid. You could and should be able to modify the tier type of the zone. What happens if you try to revert the tier type (e.g., set it to "rgw" instead). Does the command fail to execute then?

#2 - 05/05/2017 08:55 AM - Abhishek Lekshmanan

Yehuda Sadeh wrote:

The command here in itself is not invalid. You could and should be able to modify the tier type of the zone. What happens if you try to revert the tier type (e.g., set it to "rgw" instead). Does the command fail to execute then?

Well the zone modify itself would succeed, the subsequent period commit fails with

```
$ ../src/mrun dc2 radosgw-admin period commit  
2017-05-05 10:54:31.662401 7f9d2dfaa6c0 -1 ERROR: failed to init sync module instance, ret=-2  
couldn't init storage provider
```

#3 - 05/05/2017 09:49 AM - Abhishek Lekshmanan

basically `init_complete()` fails since it will fail at `RGWSyncModulesManager::create_instance()` and since period commit or any other rgw admin comments that isn't a raw storage op would fail, so even though zone modify succeeds, it can't be committed in the period and hence subsequent changes are not possible from this zonegroup anymore. Maybe we could do a check in zone modify on the `tier_type` by asking the sync modules manager?

#4 - 05/09/2017 08:06 AM - Abhishek Lekshmanan

- *Status changed from New to Need Review*

master pr: <https://github.com/ceph/ceph/pull/14994>