

rgw - Bug #41652

rgw lc does not delete objects that do not have the same tags

09/04/2019 01:51 PM - Or Friedmann

Status:	Resolved	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:	Or Friedmann		
Category:			
Target version:	v15.0.0		
Source:	Community (dev)	Reviewed:	
Tags:		Affected Versions:	v15.0.0
Backport:	nautilus	ceph-qa-suite:	rgw
Regression:	No	Pull request ID:	30151
Severity:	3 - minor	Crash signature:	

Description

It is possible that object will have multiple tags more than the rule that applied on.

For example:

Example 4: Tag-based Filtering and Resulting Conflicting Lifecycle Actions

Suppose that you have the following lifecycle policy that has two rules, each specifying a tag filter:

Rule 1 specifies a tag-based filter (tag1/value1). This rule directs Amazon S3 to transition objects to the GLACIER storage class 365 days after creation.

Rule 2 specifies a tag-based filter (tag2/value2). This rule directs Amazon S3 to expire objects 14 days after creation.

The lifecycle configuration is shown following:

```
<LifecycleConfiguration>
<Rule>
<ID>Rule 1</ID>
<Filter>
<Tag>
<Key>tag1</Key>
<Value>value1</Value>
</Tag>
</Filter>
<Status>Enabled</Status>
<Transition>
<StorageClass>GLACIER<StorageClass>
<Days>365</Days>
</Transition>
</Rule>
<Rule>
<ID>Rule 2</ID>
<Filter>
<Tag>
<Key>tag2</Key>
<Value>value1</Value>
</Tag>
</Filter>
<Status>Enabled</Status>
<Expiration>
<Days>14</Days>
</Expiration>
</Rule>
</LifecycleConfiguration>
```

The policy is fine, but if there is an object with both tags, then S3 has to decide what to do. That is, both rules apply to an object and in effect you are directing Amazon S3 to perform conflicting actions. In this case, Amazon S3 expires the object 14 days after creation. The object is removed, and therefore the transition action does not come into play.

Object is not being deleted if not all tags exactly the same as in the rule.

Related issues:

Copied to rgw - Backport #45502: nautilus: rgw lc does not delete objects tha...

Resolved

History

#1 - 09/19/2019 05:45 PM - Abhishek Lekshmanan

- Assignee set to Abhishek Lekshmanan

#2 - 09/19/2019 05:46 PM - Abhishek Lekshmanan

- Assignee changed from Abhishek Lekshmanan to Or Friedmann

#3 - 09/19/2019 05:46 PM - Abhishek Lekshmanan

- Status changed from New to 12

#4 - 12/05/2019 09:37 PM - Patrick Donnelly

- Status changed from 12 to New

#5 - 01/13/2020 08:25 AM - Or Friedmann

solved by this PR: <https://github.com/ceph/ceph/pull/30151>

#6 - 04/16/2020 10:53 AM - Or Friedmann

- Pull request ID set to 30151

#7 - 05/11/2020 02:00 PM - Or Friedmann

- Status changed from New to Pending Backport

- Backport set to nautilus

#8 - 05/11/2020 02:31 PM - Nathan Cutler

- Copied to Backport #45502: nautilus: rgw lc does not delete objects that do not have the same tags added

#9 - 05/18/2020 09:21 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".