

rgw - Bug #39551

multisite: RGWListBucketIndexesCR for data full sync needs pagination

04/30/2019 06:50 PM - Casey Bodley

Status: Resolved	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee: Shilpa MJ	
Category:	
Target version:	
Source:	Affected Versions:
Tags: multisite admin-api	ceph-qa-suite:
Backport: luminous mimic nautilus	Pull request ID: 28146
Regression: No	Crash signature (v1):
Severity: 3 - minor	Crash signature (v2):
Reviewed:	

Description

When a new peer zone is added, rgw starts in 'full sync'. This starts by listing all of that zone's bucket instances, in order to build a list of all of the bucket index shards we need to sync.

from RGWListBucketIndexesCR::operate():

```
/* FIXME: need a better scaling solution here, requires streaming output */
call(new RGWReadRESTResourceCR<list<string
>>(store->ctx(), sync_env->conn, sync_env->http_manager,
                                     entrypoint, NULL, &result));
```

This attempts to read the entire list of bucket instance ids from the remote zone's "/admin/metadata/bucket.instance" admin api. RGWOp_Metadata_List is the op for this admin api, and it will return a maximum of 1000 entries - so if there are more bucket instances than that, data sync won't see them. RGWOp_Metadata_List accepts a "marker" parameter that allows you to resume a listing from where the last request left off. RGWListBucketIndexesCR should use this marker in order to continue listing keys until the end (where truncated = false).

Related issues:

Copied to rgw - Backport #40352: nautilus: multisite: RGWListBucketIndexesCR ...	Resolved
Copied to rgw - Backport #40353: luminous: multisite: RGWListBucketIndexesCR ...	Rejected
Copied to rgw - Backport #40354: mimic: multisite: RGWListBucketIndexesCR for...	Resolved

History

#1 - 05/07/2019 05:35 PM - Casey Bodley

- Status changed from New to In Progress
- Assignee set to Shilpa MJ

#2 - 05/16/2019 02:45 PM - Casey Bodley

- Status changed from In Progress to Fix Under Review
- Backport set to luminous mimic nautilus
- Pull request ID set to 28146

#3 - 06/13/2019 01:39 PM - Casey Bodley

- Status changed from Fix Under Review to Pending Backport

#4 - 06/13/2019 09:02 PM - Nathan Cutler

- Copied to Backport #40352: nautilus: multisite: RGWListBucketIndexesCR for data full sync needs pagination added

#5 - 06/13/2019 09:02 PM - Nathan Cutler

- Copied to Backport #40353: luminous: multisite: RGWListBucketIndexesCR for data full sync needs pagination added

#6 - 06/13/2019 09:02 PM - Nathan Cutler

- Copied to Backport #40354: mimic: multisite: RGWListBucketIndexesCR for data full sync needs pagination added

#7 - 01/27/2021 07:11 PM - 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".