

rgw - Documentation #24508

How to configure user- or bucket-specific data placement

06/12/2018 08:48 PM - Casey Bodley

Status:	Resolved	Start date:	06/12/2018
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Tags:		Affected Versions:	
Backport:	luminous mimic	Pull request ID:	26997
Reviewed:			

Description

There is some minimal documentation about the radosgw-admin 'zone/zonegroup placement' commands, and some documentation about the use of pools in radosgw. But there's nothing to describe how they fit together, how to add new placement targets, and how to configure specific users/buckets to use them.

This question was raised on the [reddit.com/r/ceph](https://www.reddit.com/r/ceph/), and I responded with these instructions:

The zonegroup contains a list of placement_targets, with default-placement as the initial/only target. The zone contains a list of placement_pools, which map placement targets from the zonegroup into specific pool names.

Assuming you leave the default-placement target as the safe/replicated pools, you'd want to add a new unsafe placement target to the zonegroup:

```
$ radosgw-admin zonegroup placement add --placement-id=unsafe
```

Then map that placement id to different pools in the zone config:

```
$ radosgw-admin zone placement add --placement-id=unsafe --index-pool=default.rgw.buckets.unsafe-index  
--data-pool=default.rgw.buckets.unsafe-data- --data-extra-pool=default.rgw.buckets.unsafe-non-ec
```

(note that the index-pool and data-extra-pool must not be erasure-coded, because they rely on omap)

You can then specify a default per-user placement target by setting the default_placement field of the user info. The radosgw-admin user modify command doesn't support setting this directly, so you'd have to resort to operating on the json format:

```
$ radosgw-admin metadata get user:<user-id> > user.json  
$ vi user.json  
$ radosgw-admin metadata put user:<user-id> < user.json
```

You can also specify a per-bucket placement target by adding a LocationConstraint to the PUT Bucket request. The format of this LocationConstraint would be api:placement-id, where api comes from the zonegroup's api_name field (and must be non-empty). This placement target would override any default_placement configured on the user.

Related issues:

Related to rgw - Documentation #38008: How to configure storage classes for l...	Resolved	01/22/2019
Copied to rgw - Backport #38796: luminous: How to configure user- or bucket-s...	Resolved	
Copied to rgw - Backport #38797: mimic: How to configure user- or bucket-spec...	Resolved	

History

#1 - 08/24/2018 08:09 AM - Konstantin Shalygin

As we can see comments on reddit this also need configured realm.

#2 - 01/22/2019 06:13 PM - Casey Bodley

- Related to Documentation #38008: How to configure storage classes for lifecycle transitions added

#3 - 03/18/2019 01:16 PM - Casey Bodley

- Status changed from New to Pending Backport
- Backport set to luminous mimic

#4 - 03/18/2019 01:17 PM - Casey Bodley

- Copied to Backport #38796: luminous: How to configure user- or bucket-specific data placement added

#5 - 03/18/2019 01:17 PM - Casey Bodley

- Copied to Backport #38797: mimic: How to configure user- or bucket-specific data placement added

#6 - 03/22/2019 01:18 PM - Nathan Cutler

- Pull request ID set to 26997

#7 - 03/22/2019 01:23 PM - Nathan Cutler

- Status changed from Pending Backport to Resolved