

# Ceph Dashboard: Proposed RGW Features

[Ju Lim](#)

Red Hat

25 April 2019

# Proposed RGW Features for Ceph Dashboard

## Admin Workflows

- multi-site setup
- CRUD zone
  - Includes pool placement and storage class
- improved multi-site status
- get/set admin caps (delegate admin privs)
- additional status/stats
  - lifecycle
  - gc
- CRUD Object Gateway Configuration or Authentication (if we're not configuring anything else)
  - config/mgmt of external auth mechs (openid::connect, ldap, &c)
- CRUD user (this exists but is skeletal)
  - get/set policy for user

## Tenant / Consumer Workflows

- CRUD bucket (this exists but is skeletal)
  - create / update bucket
    - placement targets
    - get/set acls (on a bucket)
    - get/set policy for bucket
    - enable/disable versioning (on a bucket)
    - get/set lifecycle rules (on a bucket)
    - get/set CORS rules (on a bucket)
    - Mfa-delete
    - x-amz-storage-class
  - delete bucket (already exists)

Tracker: <http://tracker.ceph.com/issues/39478> (mgr/dashboard: new RGW workflows & RGW enhancements)

# RGW Multi-site User Stories

- As an Administrator, I want to delete the default zonegroup and default zone (per documented best practice).
- As an Administrator, I want to setup multi-site between 2 different Ceph clusters (add).
  - As an Administrator, I want to configure multi-site on the first cluster (add).
  - As an Administrator, I want to quickly configure the multi-site on the second cluster by leveraging the first cluster's multi-site configuration (import/export).
- As an Administrator, I want to be able to specify placement targets and storage classes on my local zone (edit).
- As an Administrator, I want to be able to rename the realm, zonegroup, or zone (rename).
- As an Administrator, I want to be able to view my multi-site topology to understand how the replication is configured (multi-site topology view)
- As an Administrator, I want to be able to quickly see my status and monitor the health and performance of my replication, as well as be alerted when there are problems (overview).

# Assumptions

- Assume RGW is configured, i.e.  $\geq 1$  RGW daemon
- Each initial RGW deployment in a Ceph cluster includes a default zonegroup and a default zone, but no default realm.
  - It is a best practice to not use the default zonegroup and default zone as it can cause unexpected replication behavior.
  - Specifically, users are expected to delete the default zonegroup and default zone before creating any new realm, zonegroup, zone.
- Most Ceph clusters typically have 1 realm, 1 zonegroup with 2 zones.
- Assumes the same access and secret keys are already configured on each gateway instance.
- Assumes Ceph Dashboard is running on each cluster used in a 2 cluster multi-site configuration.

# Ceph Dashboard: Current State

| RGW Epic                    | RGW Feature                | Description                                    | Implemented?                        | ceph-mgr orch impact?               |
|-----------------------------|----------------------------|------------------------------------------------|-------------------------------------|-------------------------------------|
| Configure RGW cluster       | Single Zone                | pool placement and storage class configuration | <input type="checkbox"/>            | <input checked="" type="checkbox"/> |
|                             | Muti-site                  |                                                | <input type="checkbox"/>            | <input checked="" type="checkbox"/> |
| Bucket Management           | Add/edit/delete            |                                                | <input checked="" type="checkbox"/> | <input type="checkbox"/>            |
|                             | Bucket Policies            |                                                | <input type="checkbox"/>            | <input type="checkbox"/>            |
|                             | Scaling                    |                                                | <input type="checkbox"/>            | <input type="checkbox"/>            |
| Users                       | Add/edit/delete            |                                                | <input checked="" type="checkbox"/> | <input type="checkbox"/>            |
|                             | User Policies              |                                                | <input type="checkbox"/>            | <input type="checkbox"/>            |
| Zone Configuration          | (Part of cluster creation) | pool placement and storage class configuration | <input type="checkbox"/>            | <input type="checkbox"/>            |
| Logging and Troubleshooting |                            |                                                | <input type="checkbox"/>            | <input type="checkbox"/>            |
| Performance Monitoring      | Buckets                    |                                                | <input type="checkbox"/>            | <input type="checkbox"/>            |
|                             | Users                      |                                                | <input type="checkbox"/>            | <input type="checkbox"/>            |
|                             | Daemons                    |                                                | <input checked="" type="checkbox"/> | <input type="checkbox"/>            |
|                             | Ips                        |                                                | <input type="checkbox"/>            | <input type="checkbox"/>            |
| Replication                 |                            |                                                | <input type="checkbox"/>            | <input type="checkbox"/>            |
| Authen/Authorz              | Cluster level              | LDAP, Keystone, Keycloak                       | <input type="checkbox"/>            | <input type="checkbox"/>            |
|                             | Bucket level               | OpenID connect, STS                            | <input type="checkbox"/>            | <input type="checkbox"/>            |

# RGW Designs: Table of Content

- RGW overview (landing page)
- Topology view for initial default configuration
- Delete initial default configuration
- Add multi-site
- Topology view for multi-site configuration (after multi-site configured)
- Edit multi-site
- Export and import realm
- Rename realm/zonegroup/zone
- Enable/disable versioning for a bucket
- RGW Metrics and Alerts Asks

# Feature: Improved multi-site status

## User Stories

As an Administrator, I want to be able to quickly see my status and monitor the health and performance of my replication, as well as be alerted when there are problems (overview).

ceph

Dashboard Cluster Pools Block NFS Filesystems Object Gateway

Object Gateway >> Overview

- Overview
- Daemons
- Multi-site
- Buckets
- Users
- Authentication

Daemons

2 total

Zoning

2 Zonegroups  
4 Zones

Buckets

40.3K buckets  
14.6TB  
800M objects

Users

200 total

default (zonegroup): default (zone) ← 1

| Metadata Sync   |     | Data Sync   |     |
|-----------------|-----|-------------|-----|
| no sync         |     |             |     |
| Metadata sync   | N/A | Data sync   | N/A |
| Metadata shards | N/A | Data shards | N/A |

Zonegroup1 (zonegroup): zone11 (master zone) <--> zone12 (zone) in realm\_name (realm) ← 2

| Metadata Sync   |                                                                                     | Data Sync                                                                            |                                                                                           |
|-----------------|-------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| sync failed     |                                                                                     | data is behind on 10 shards<br>change not applied since 2019-04-18 15:00:17.0.330225 |                                                                                           |
| Metadata sync   |  | Data sync                                                                            |  ← 3 |
| Metadata shards |  | Data shards                                                                          |  ← 4 |

**Notes**

**RGW overview page (dashboard)**

- For every zonegroup, a card will be presented.
- Note: synchronization status may take a little time to update, and we may need to factor some kind of progress/loading indicator when there is no data to show.

1 Example of the default zonegroup + default zone.

2 Example of a multi-site configuration.

3 Metadata / Data sync sparkline tracks delta of current vs. the oldest incremental change timestamp plotted over time.

- Goal is to see if the delta is expanding (not good) or decreasing (good) over time.

4 Metadata / Data shards sparkline tracks shards caught up with master plotted over time.

- Goal is to see if the shards are increasing/reducing over time (as it's an indicator of replication).

**Performance Asks**

- Top N Buckets by IOs (Bucket Name, IOs, Objects, Total Size, Shards, Owner/User)
- Top N Users by IOs (Users, IOs, Buckets Accessed, Objects, Total Size)
- Top N IP by IOs (IP, IOs, Buckets Accessed, Objects, Total Size)

# Feature: Delete initial default configuration

## User Stories

As an Administrator, I want to delete the default zonegroup and default zone (per documented best practice).

ceph

English

Dashboard Cluster Pools Block NFS Filesystems Object Gateway

Object Gateway >> Multi-site

+ Add

Expand All Collapse All

### Multi-site Topology Viewer

▼ default (zonegroup)

- default (zone) ← 1

2

Details Pools Performance Details Logs

|             |            |
|-------------|------------|
| zone        | default    |
| id          | ###        |
| default     | Yes        |
| master      | No         |
| endpoints   |            |
| zone attrib | zone value |
| zone attrib | zone value |

## Notes

### Multi-site default page

- shows the default zonegroup and default zone (initial first-time experience).
- A default zonegroup and default zone exist when RGW is enabled. The default zonegroup and default zone are created automatically first time radosgw runs without any multisite configuration.
- When the default zone is first created, associated pool names starting with default.rgw are also automatically created.
- Typically, users will delete the default zonegroup and default one to create a new multi-site configuration.

**1** Only single-selection supported on the tree.

- No multi-selection permitted.

**2** Tabs in the details area (below) appears when an item in the tree (above) is selected.

### Additional Notes

- Realm - namespace
- ZoneGroup - group of replicating zones
- Zone - represent a geographical location, and it cannot cross clusters
  - Period - current realm configuration. Updates are local and are only applied when committed.

ceph

English

Dashboard Cluster Pools Block NFS Filesystems Object Gateway

Object Gateway >> Multi-site

+ Add Edit Rename Export Import Delete

Topology Viewer

(zone)

Expand All Collapse All

Details Pools Performance Details

|                  |                 |
|------------------|-----------------|
| zonegroup        | default         |
| id               | ###             |
| default          | Yes             |
| master           | No              |
| endpoints        |                 |
| zonegroup attrib | zonegroup value |
| zonegroup attrib | zonegroup value |

## Notes

### Delete zonegroup and its zones

- Per Ceph docs, the best practice is to remove the default zonegroup and default zone due to it being buggy in previous releases. Note: More recently, one can use the default zonegroup and default zone, but the recommendation is to rename and reconfigure them.

#### 1 Actions (for realm, zonegroup, and zone)

- Add
- Edit
- Rename
- Export ("radosgw admin <realm | zonegroup | zone> get")
- Import ("radosgw admin <realm | zonegroup | zone> set")
- Delete
- Note: For Add and Import, no object selection on the topology viewer is needed. However, for all other actions, User has to click on an object (i.e. realm, zonegroup, or zone) before clicking on action (i.e. Edit, Rename, Export, Delete).

#### 2 User selects the default zonegroup for deletion.

ceph

Dashboard Cluster Pools Block NFS Filesystems Object Gateway

Object Gateway >> Multi-site

- Edit

### Multi-site Topology Viewer

- default (zonegroup)
  - default (zone)

**Delete Zonegroup**

Are you sure that you want to delete the selected zonegroup and its zones?  
This will delete the following pools and any data stored in these pools:

- .rgw.root
- .rgw.control
- .rgw.meta
- .rgw.log
- .rgw.buckets.index
- .rgw.buckets.data

Yes, I am sure.

Delete zonegroup Cancel

Expand All Collapse All

---

Details Pools Performance Details

|                  |                 |
|------------------|-----------------|
| zonegroup        | default         |
| id               | ###             |
| default          | Yes             |
| master           | No              |
| endpoints        |                 |
| system user      | xxx             |
| zonegroup attrib | zonegroup value |
| zonegroup attrib | zonegroup value |

**Notes**

**Delete Zonegroup Confirmation Modal**

1 Only the pools for the zone(s) would be listed.

Blank Slate Image Placeholder

Multi-site not yet configured. We'll put some useful information about what user should do or what the feature is about and let user know they can either add or import the multi-site configuration.

Learn more about this at [Ceph Documentation](#).

[Add Multi-site Configuration](#) [Import Multi-site Configuration](#) ← 1

Notes

**Blank Slate if no multi-site configuration exists**

- If there are no realms, zonegroups, zones to show, a blank slate is shown.

1 Add Multi-site Configuration - allows user to add a multi-site configuration.

2 Import Multi-site Configuration - allows user to add a multi-site configuration through importing a configuration from another cluster (which has an existing multi-site configuration.)

- Note: user would have to have done an export from the other cluster first.