

Feature: Configure Multi-site

User Story

As an Administrator, I want to configure multi-site on the first cluster (add)

ceph

English

Dashboard Cluster Pools Block NFS Filesystems Object Gateway

Object Gateway >> Multi-site >> Add

1 Realm 2 Zonegroups 3 Zones 4 System User 5 Review

Add Realm

Add new realm

Name *

Default

Use existing realm

Name *

Notes

Add Multi-site Configuration workflow Step 1 - Add Realm

- Workflow for creating a multi-site configuration.
- Triggered by user selecting the "Add" action previously.
- A multi-site configuration requires at least 2 Ceph storage clusters, preferably given a distinct cluster name. At least 2 Ceph object gateway instances (rgw), one for each Ceph storage cluster.
- A multi-site configuration requires a master zone group and a master zone. Additionally, each zone group requires a master zone.
- Zone groups may have one or more secondary or non-master zones.

IMPORTANT NOTES

- The master zone within the master zone group of a realm is responsible for storing the master copy of the realm's metadata, including users, quotas and buckets (created by the radosgw-admin CLI).
- This metadata gets synchronized to secondary zones and secondary zone groups automatically.
- Metadata operations executed with the radosgw-admin CLI MUST be executed on a host within the master zone of the master zone group in order to ensure that they get synchronized to the secondary zone groups and zones.
- Currently, it is possible to execute metadata operations on secondary zones and zone groups, but it is NOT recommended because they WILL NOT be synchronized, leading to fragmented metadata.

ceph

Dashboard Cluster Pools Block NFS Filesystems Object Gateway

Object Gateway >> Multi-site >> Add

1 Realm 2 Zonegroups 3 Zones 4 System User 5 Review

Add Zonegroup

Name *

Default
 Master

Endpoints *

<< Edit Realm Add Another Zonegroup Add Zones >> Cancel

Notes

Add Multi-site Configuration workflow Step 2 - Add Zonegroup

- Add first zonegroup.

ceph

English

Dashboard Cluster Pools Block NFS Filesystems Object Gateway

Object Gateway >> Multi-site >> Add

1 Realm 2 Zonegroups 3 Zones 4 System User 5 Review

Add Zonegroup

Name *

Default
 Master

Endpoints *

<< Edit Realm Add Another Zonegroup **Add Zones >>** Cancel

Notes

Add Multi-site Configuration workflow Step 2 - Add Another Zonegroup

- Add second zonegroup (optional sub-step).

ceph

Dashboard Cluster Pools Block NFS Filesystems Object Gateway

Object Gateway >> Multi-site >> Add

1 2 3 4 5

Realm Zonegroups Zones System User Review

Add Zone

Zonegroup * ← 1

Name *

Default

Master

Zone Type * Default ← 2

Endpoints * ← 3

<< Edit Zonegroups Add Another Zone Add System User >> Cancel

Notes

Add Multi-site Configuration workflow Step 3 - Add Zone

- Add first zone.

1 Zonegroup should automatically default to the zonegroup that was just created.

2 Zone Type options (only for when additional zones get added) include:

- default (standard zone)
- archive (Archive sync tier type); Archive is a new feature introduced in Nautilus.

3 Endpoint for the zone is typically different from the endpoint of the zonegroup.

ceph

English

Dashboard Cluster Pools Block NFS Filesystems Object Gateway

Object Gateway >> Multi-site >> Add

1 2 3 4 5

Realm Zonegroups Zones System User Review

Add Zone

Zonegroup * -- Select a zonegroup --

Name * Name...

Default ← 1

Zone Type * -- Select a zone type -- ← 2

Endpoints * e.g., http://fqdn, http://ceph-node1.example.com:80

<< Edit Zonegroups Add Another Zone Add System User >> Cancel

Notes

Add Multi-site Configuration workflow Step 3 - Add Another Zone

- Add second zone (optional sub-step).

1 Secondary zones added will not have a master checkbox option.

2 Zone Type options (only for when additional zones get added) include:

- default (standard zone)
- archive (Archive sync tier type); Archive is a new feature introduced in Nautilus.

Additional Notes on the archive zone type

- By including an archive zone in a multizone configuration, it allows you to have the flexibility of an S3 object history in only one zone while saving the space that the replicas of the versioned S3 objects would consume in the rest of the zones.
- <http://docs.ceph.com/docs/master/ra-dosgw/archive-sync-module/>

FUTURE

- support for ElasticSearch
- support for Pubsearch

ceph

English

Dashboard Cluster Pools Block NFS Filesystems Object Gateway

Object Gateway >> Multi-site >> Add

1 Realm 2 Zonegroups 3 Zones 4 System User 5 Review

Add System User

Master zone Master_zone_name

Username *

<< Edit Zones **Review >>** Cancel

Notes

Add Multi-site Configuration workflow Step 4 - Add System User for the Master Zone

- Creates a system user that is associated to the master zone.
- System takes care of ensuring the keys for this user is used by each of the zones being configured.
- Note: from Orchestrator perspective, the system user has to be created after master zone is created, and then modifying the master zone to attach the system user's keys.

ceph

Dashboard Cluster Pools Block NFS Filesystems Object Gateway

Object Gateway >> Multi-site >> Add

1 2 3 4 5

Review

- realm_name (realm)
 - zonegroup_1 (zonegroup) master
 - zone_11 (zone) master
 - zone_12 (zone)
 - zonegroup_2 (zonegroup) master
 - zone_21 (zone) master
 - zone_22 (zone)

zone	zone_22
id	###
default	No
master	No
endpoints	http://ceph-node01.example.com:80
zone attrib	zone value
zone attrib	zone value

System User

Username	zone11_user
----------	-------------

<< Edit System User Add Multi-site >> Cancel

Notes

Add Multi-site Configuration workflow Step 5 - Review choices

1 Details are shown upon selection of an object on the tree (similar to crushmap viewer).