



Bucket Synchronization Policies

Ceph RGW Multi-Site Replication with QuantaStor

~15 min presentation + live demo
By: Seth Cagampang & Sheetal Garde

Agenda



Zone vs Bucket Replication



Bucket Sync Policies Overview



Initial Setup: What to Expect



Uni- vs Bidirectional Replication



Storage Classes & Impact



Live Demo with QuantaStor CLI

Zone vs Bucket Replication

Zone-Level Replication

"Symmetrical replication"

All buckets replicate between zones automatically

Simplest to manage — everything replicates

Best for full DR / active-active sites

Bucket-Level Replication

"Bucket sync policies"

Only selected buckets replicate — granular control

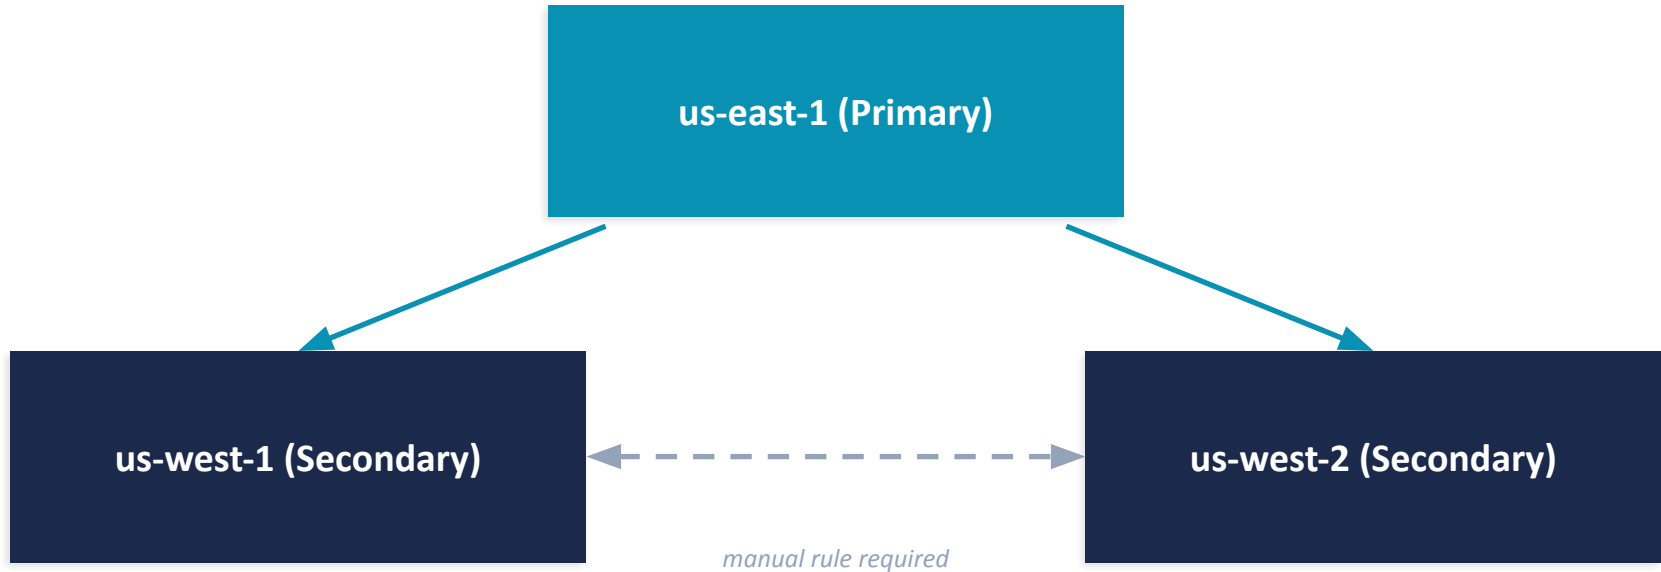
Default deny rules created between zones

Must add explicit enable rules per bucket

Best for selective replication or bandwidth-limited sites

Note: AWS supports intra-zone bucket replication, but Ceph RGW does not.

Bucket Sync Policies: How It Works



Key Behaviors:

Primary → Secondary: Default deny rules block replication until bucket-level enable rules are added

Secondary → Secondary: No default rules — must manually create default-deny (Allowed) rules

Initial Setup: What to Expect



Target cluster must be empty when initially linking — Ceph does not bidirectionally sync pre-existing data during initial setup.

1

Link Cluster A → B

Full metadata sync begins immediately. Zone group config, user accounts, bucket metadata all synchronize.

2

Metadata Sync Completes

All bucket listings, ACLs, and user info propagated. Data sync begins for each bucket's objects.


3

Object Data Sync

Objects sync based on timestamps. New objects written during sync are queued and replicated in order.


Unidirectional vs Bidirectional

Unidirectional



Writes only on Site A — Site B is read-only replica

Bidirectional



Both sites accept writes — conflicts resolved by timestamp



Conflict Resolution


Objects are timestamped at write time. When the same object is modified on both sites, the most recent timestamp wins. This is automatic — no manual intervention needed. Eventual consistency is guaranteed.

Storage Classes & Impact



Matching Storage Classes

us-east-1
STANDARD | STANDARD_IA




us-west-1
STANDARD | STANDARD_IA



Mismatched Storage Classes

us-east-1
STANDARD | STANDARD_IA



us-west-2
STANDARD only

Objects in STANDARD_IA on the source map to STANDARD on the destination. Storage class tiering is not preserved when the destination lacks a matching class.



Live Demo

QuantaStor CLI + Ceph RGW Multi-Site Replication

1. Show existing multi-site cluster configuration
2. Show buckets and sync policies
3. Upload objects and observe replication between clusters
4. Verify objects on Cluster A and Cluster B via CLI and S3 API

Demo: Key CLI Commands

Sync Policy Create

```
qs bucket-sync-policy-create --source-ceph-cluster=<primary>  
--destination-ceph-cluster=<secondary> --source-zone=<primary> --dest-zone=<secondary>  
--sync-policy-status=<enabled,forbidden> --direction=<symmetrical,directional>  
--bucket-sync-policy-name=<sync-name> --bucket-name=<target-bucket>
```

List Sync Policies

```
qs bucket-sync-policy-list
```

Upload Test Objects

```
qs-s3util create-objects --endpoint=http://<host>:7480 --accesskey=blah --secretkey=secret  
--object-count=5 --bucket-name=<target>
```

Verify Test Objects

```
qs-s3util list-objects --endpoint=http://<host>:7480 --accesskey=blah --secretkey=secret  
--bucket-name=<target>
```

Key Takeaways



Zone-level replication is all-or-nothing; bucket sync policies give granular control



Target cluster must be empty on initial link — plan data migration accordingly



Bidirectional replication uses timestamps for automatic conflict resolution



Mismatched storage classes fall back to STANDARD on the destination



Secondary-to-secondary sync requires explicit manual rules



QuantaStor simplifies Ceph RGW multi-site with a unified CLI and UI

Questions?

wiki.osnexus.com/Bucket_Synchronization_Policies

Thank you!