Multi Cluster and Protocol visibility with new Ceph Dashboard

Ankush Behl - IBM

Ankush Behl

ankush.behl@ibm.com

- Engineering Manager in IBM
- 10 year experience in storage Management tooling
- **3 years** with Ceph
- Experience with K8S storage, Openshift and Gluster storage

Agenda

- **Why Ceph Dashboard Matters?**
- **What's New in Tentacle?**
- **Deep Dive:**
  - Multi-cluster Dashboard
  - Management Gateway with OAuth2
  - RGW Automation & Sync Policies
  - Tiering, Rate Limiting & Account Control
  - SMB/NFS Protocol Support
  - Monitoring Dashboards (Application, NVMe, CephFS, SMB)
- **Road Ahead (Futures)**
- **Q&A**

Why Ceph Dashboard Matters?

🧩 Ceph is Growing in Complexity

- Multi-cluster, multi-protocol (NFS, SMB, NVMe)
- Trying to provide Ease-of-use

🎯 Operators Need a Unified View

- Make it unified workflow
- Making experience easy for administrators/users

💡 Dashboard = Monitoring Ceph cluster

- Visualize health, performance, and capacity
- Looking at Ceph Health status and Alerts that are being triggered

Ceph Dashboard

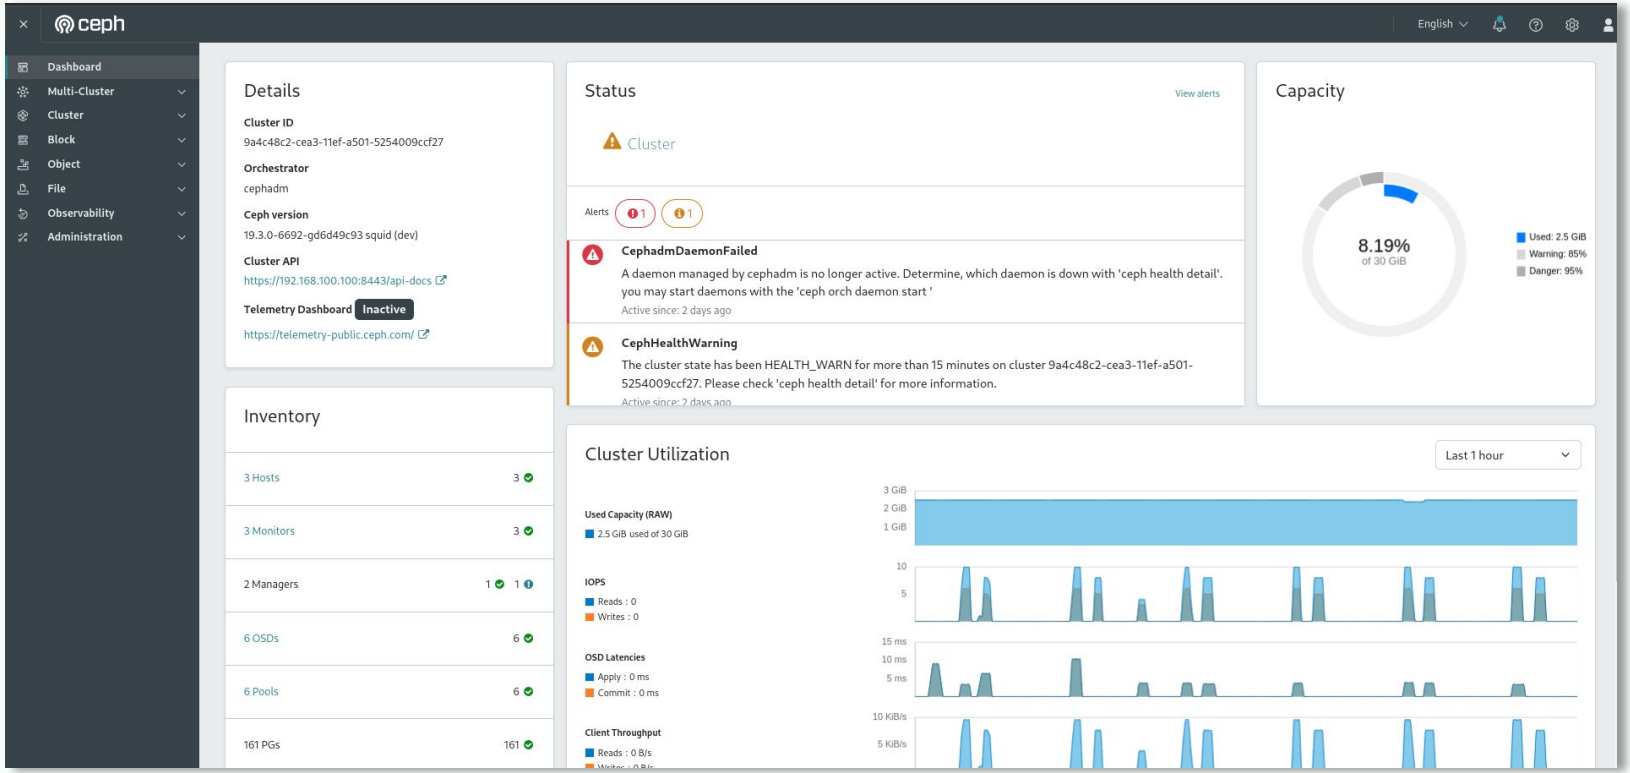

What's New in Tentacle?

Ceph Dashboard – Tentacle Highlights

🔁 Multi-Cluster Management & Monitoring

- Switch between clusters using a single dashboard view
- Multi-cluster overview from Grafana

🪣 Object (RGW) Enhancements

- Object replication automation
- Granular sync policy support(Per- bucket Replication)
- Storage class tiering
- S3 account management
- Per-user & bucket rate limiting

🔐 Management Gateway with OAuth2

- OAuth2-based SSO for centralized, secure access

What's New in Tentacle?

📁 File (CephFS) Improvements

- SMB management
- NFS UX update

📊 Monitoring Dashboards

- Multi-cluster Monitoring
- Application-level monitoring
- NVMe, and CephFS dashboards

Multi-Cluster Management & Monitoring

Multi-Cluster Management & Monitoring

Many organization which are using ceph have multiple Ceph clusters setup, each with its own Grafana, Prometheus, Ceph Dashboard, and Alertmanager. This makes it challenging for administrators to monitor and manage resources efficiently, as they must constantly switch between different dashboards.

Key Features

- Monitoring of multiple Clusters
- Alerting of multiple clusters
- Day 2 - Management of individual Ceph Cluster using single Ceph Dashboard UI
- Automated Replication setup for different storage protocol

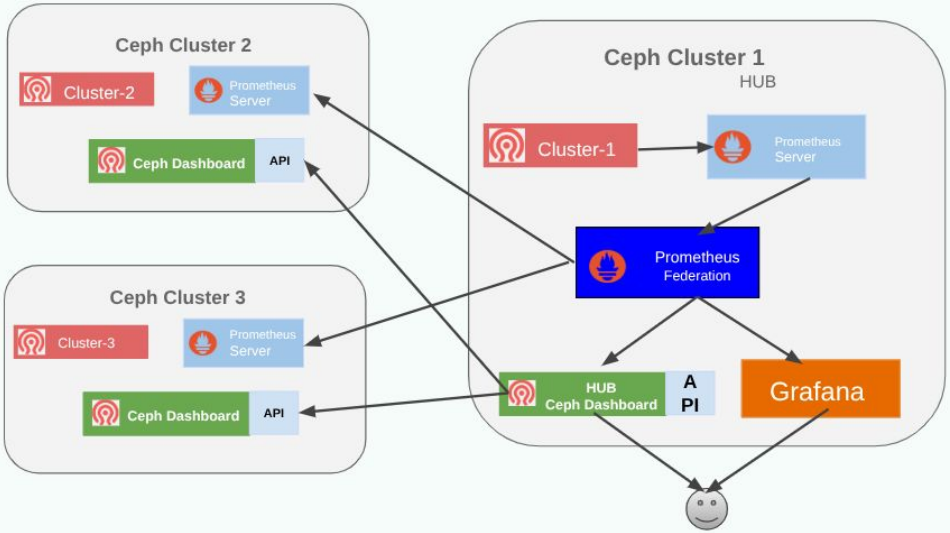

Multi-Cluster Management - Overview

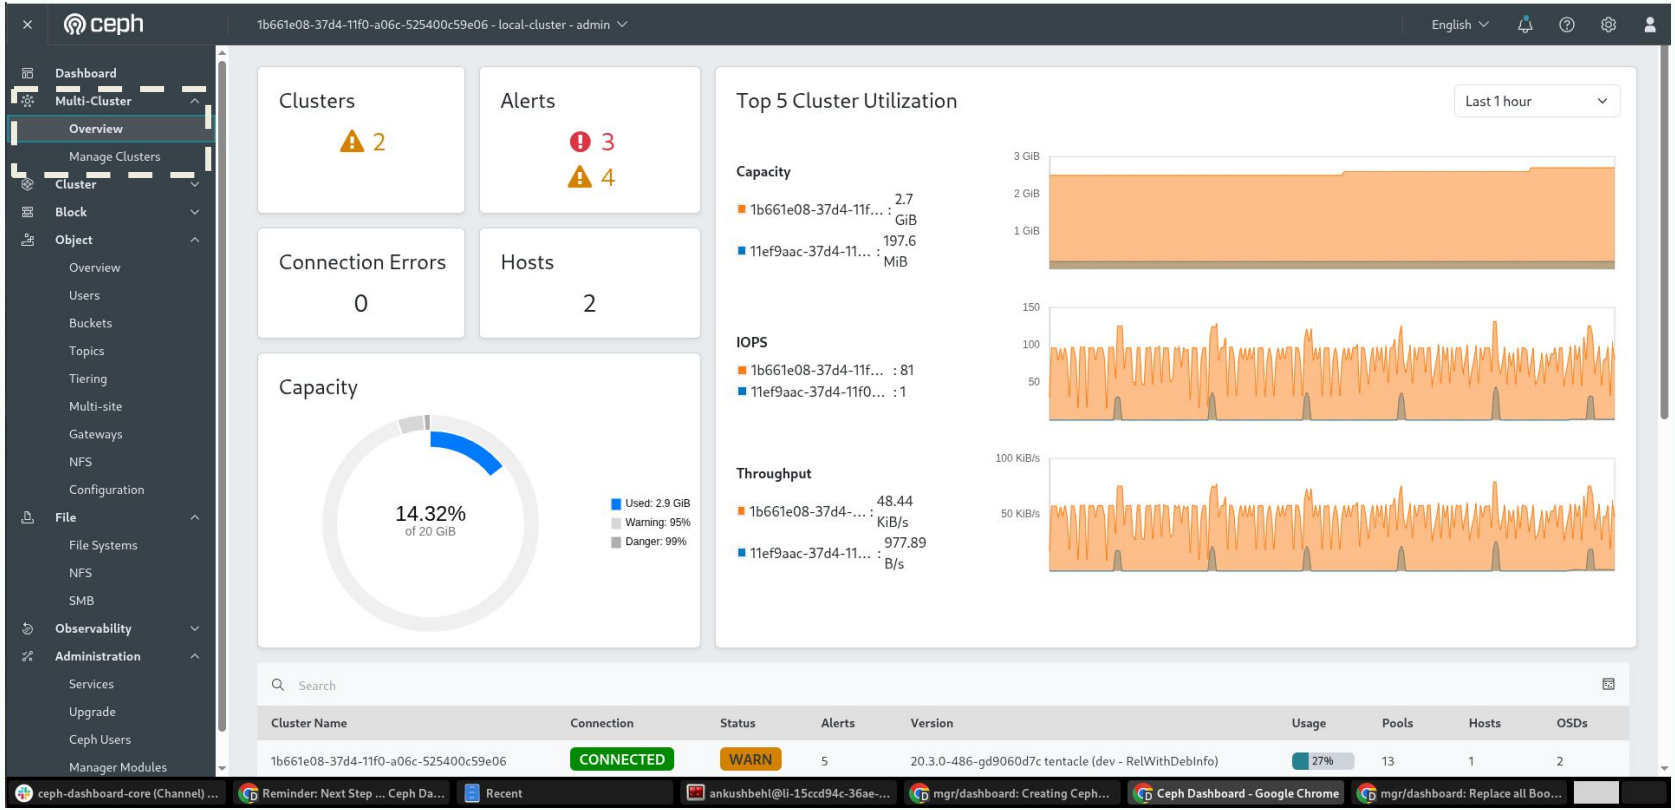

Multi-Cluster Management - Cluster List

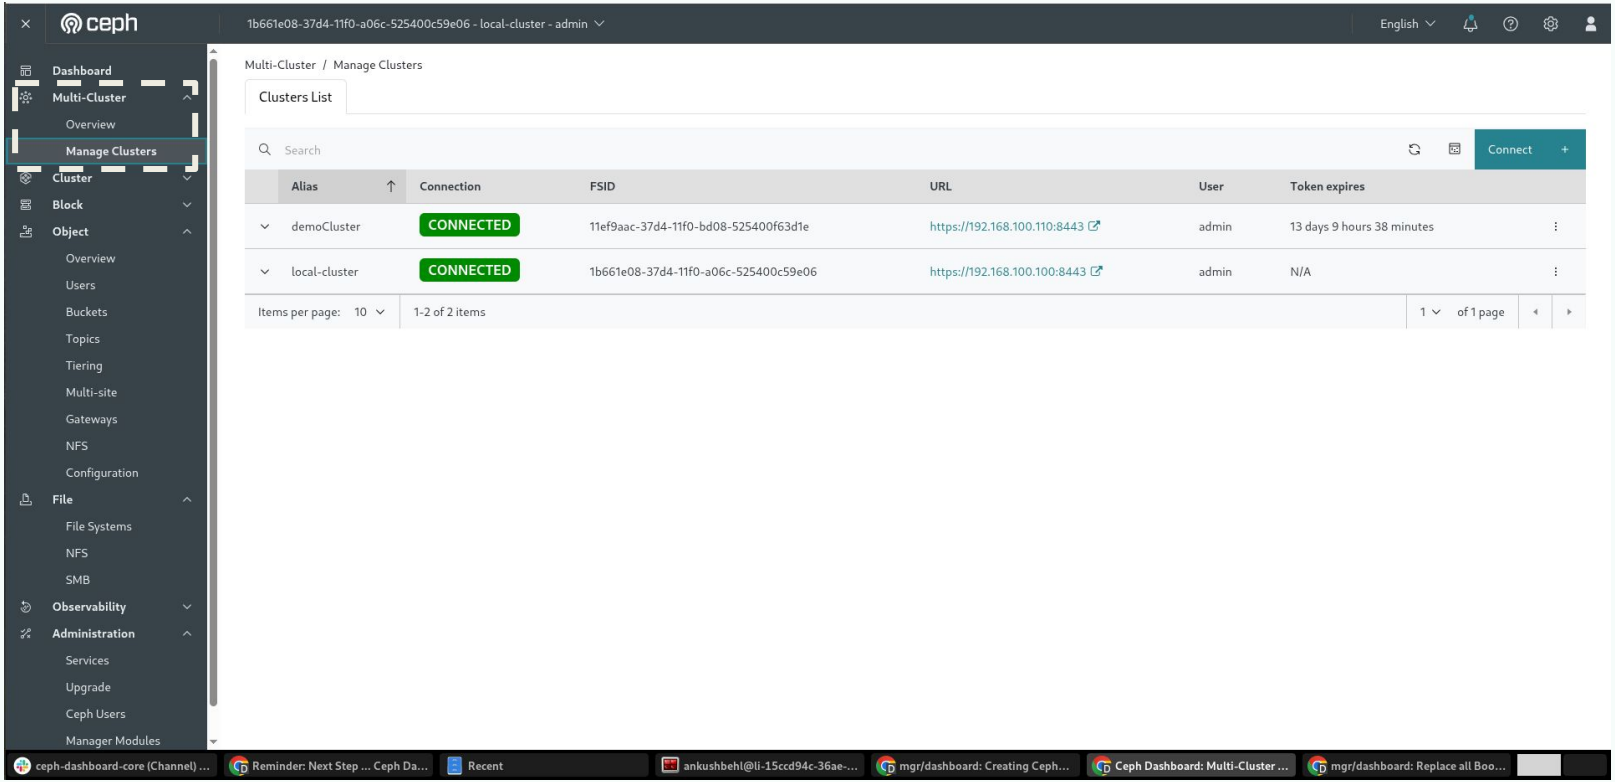

Multi-Cluster Management - Demo

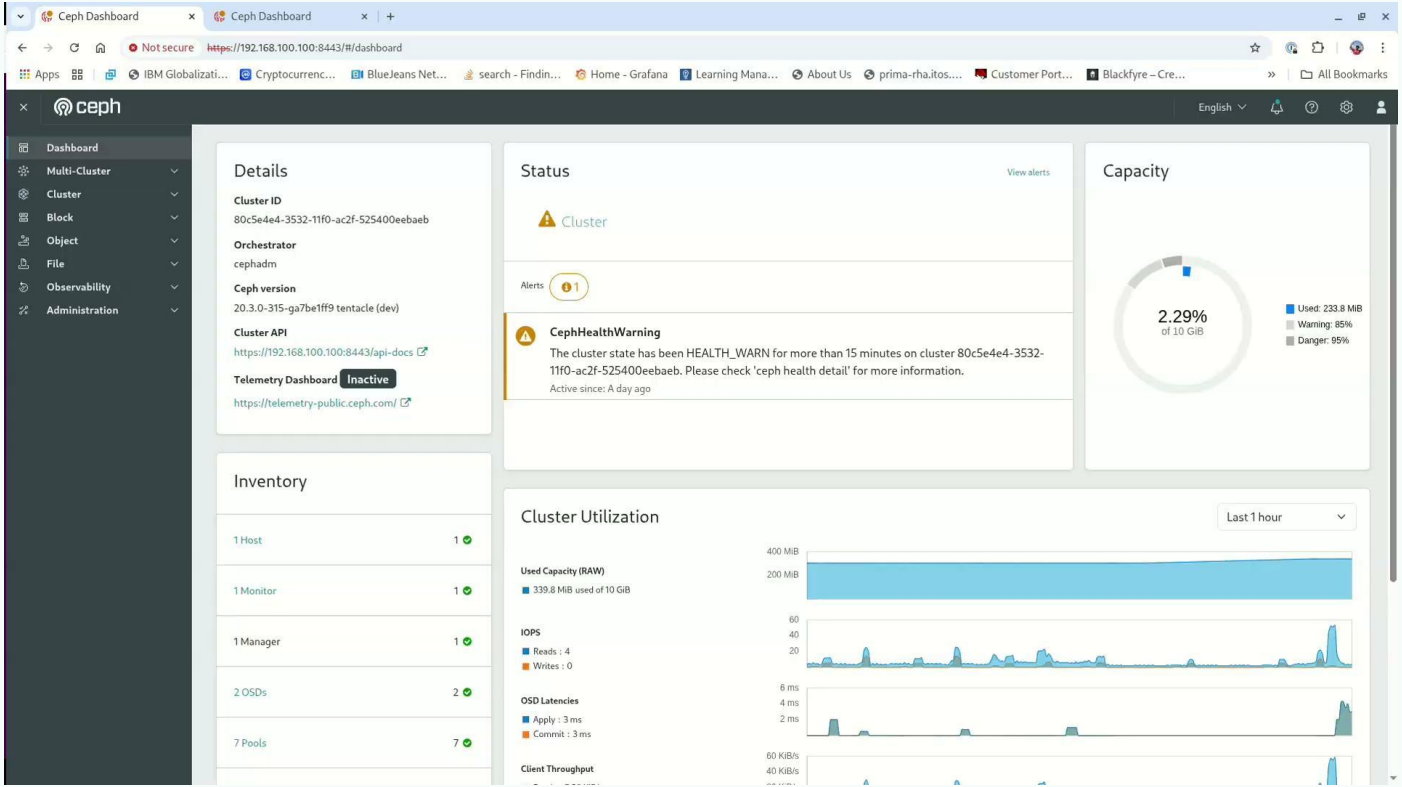

🪣 Object(RGW) Enhancements

Object - Replication Automation

Object Multi-site replication in Ceph can be complex to setup, involving realm, zones, zonegroups, and sync policies. We've simplified this with an intuitive, wizard-based UI workflow which reduces operational overhead, ensuring reliable replication, and making the process faster and less error-prone.

Key Features

- Existing setups can be converted into object replication.
- New multi-site replication configurations can be created from scratch.
- If clusters are connected in a multi-cluster setup, the entire process is automated for the user.
- In standalone setups, a token is provided that can be copied and pasted into the peer cluster for configuration.

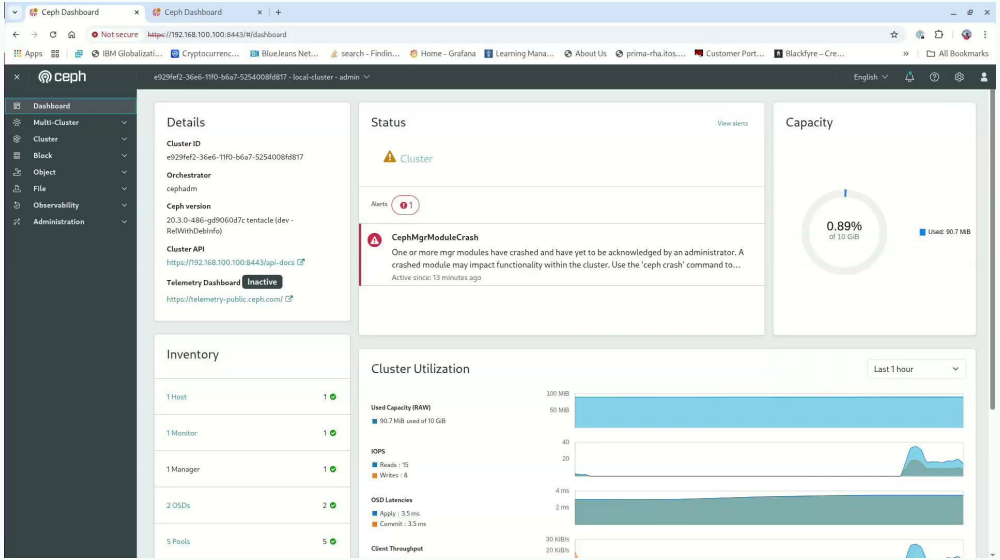

Blog: https://ceph.io/en/news/blog/2024/rgw-multisite-replication-wizard/

Object - Replication Automation - Multi-site

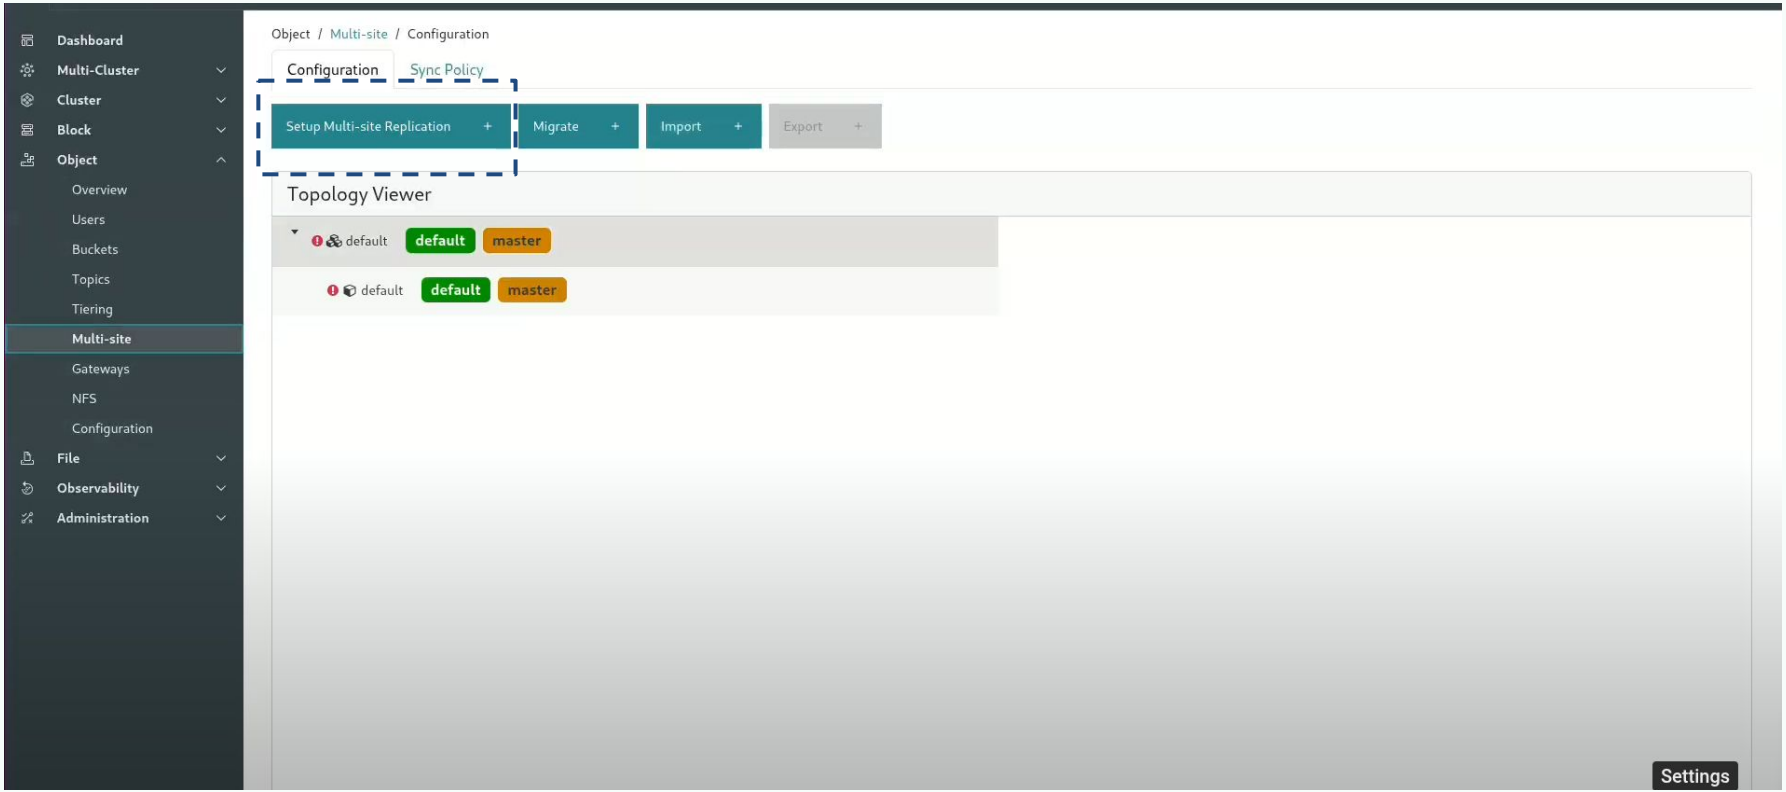

Object - Replication Automation - Wizard

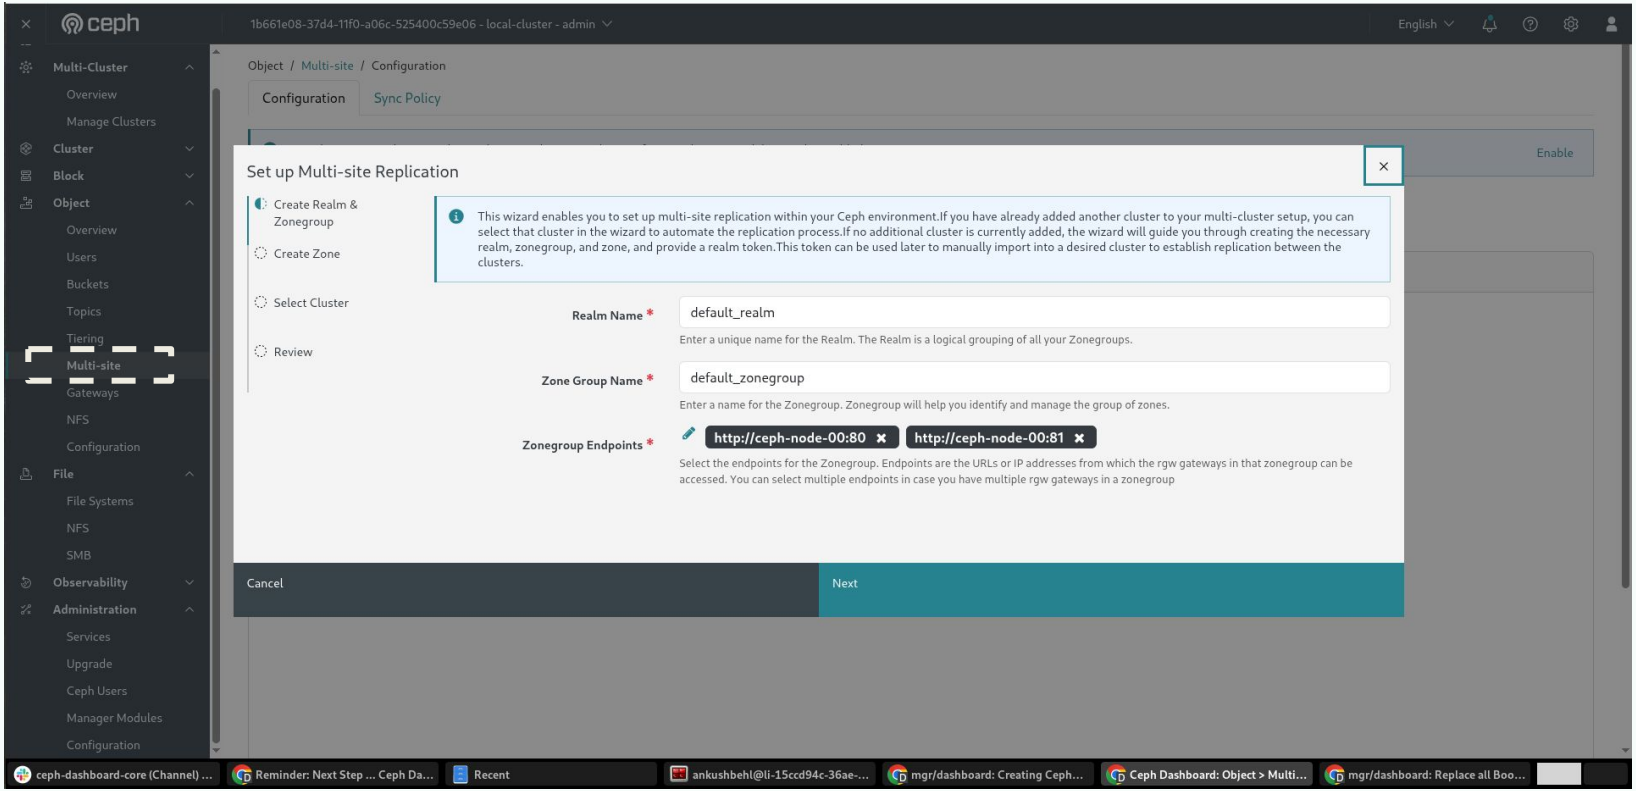

Object - Replication Automation - Step 2 & 3

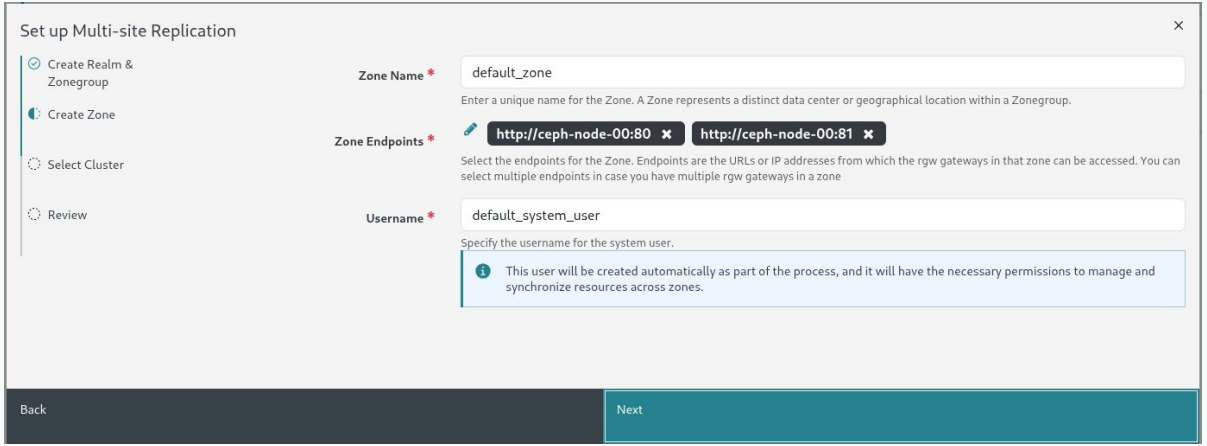

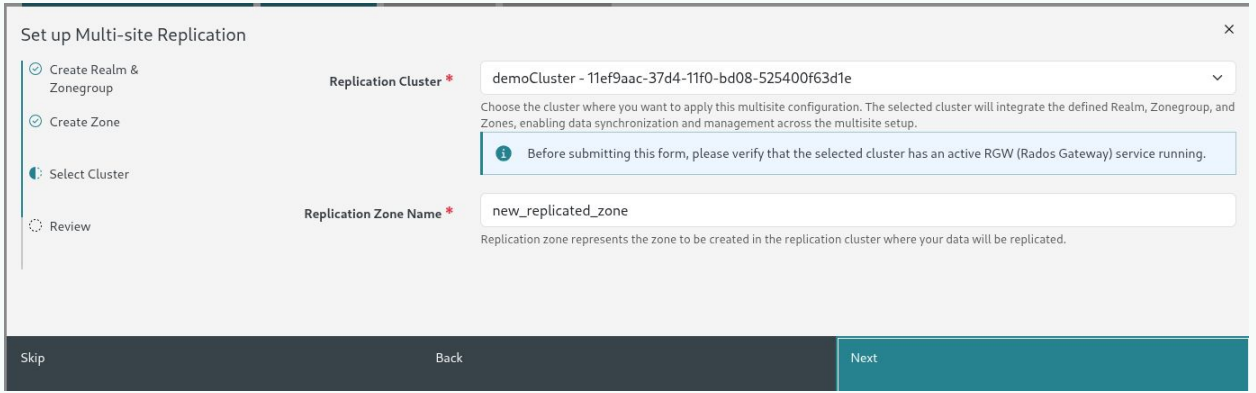

Object - Replication Automation - Step 4

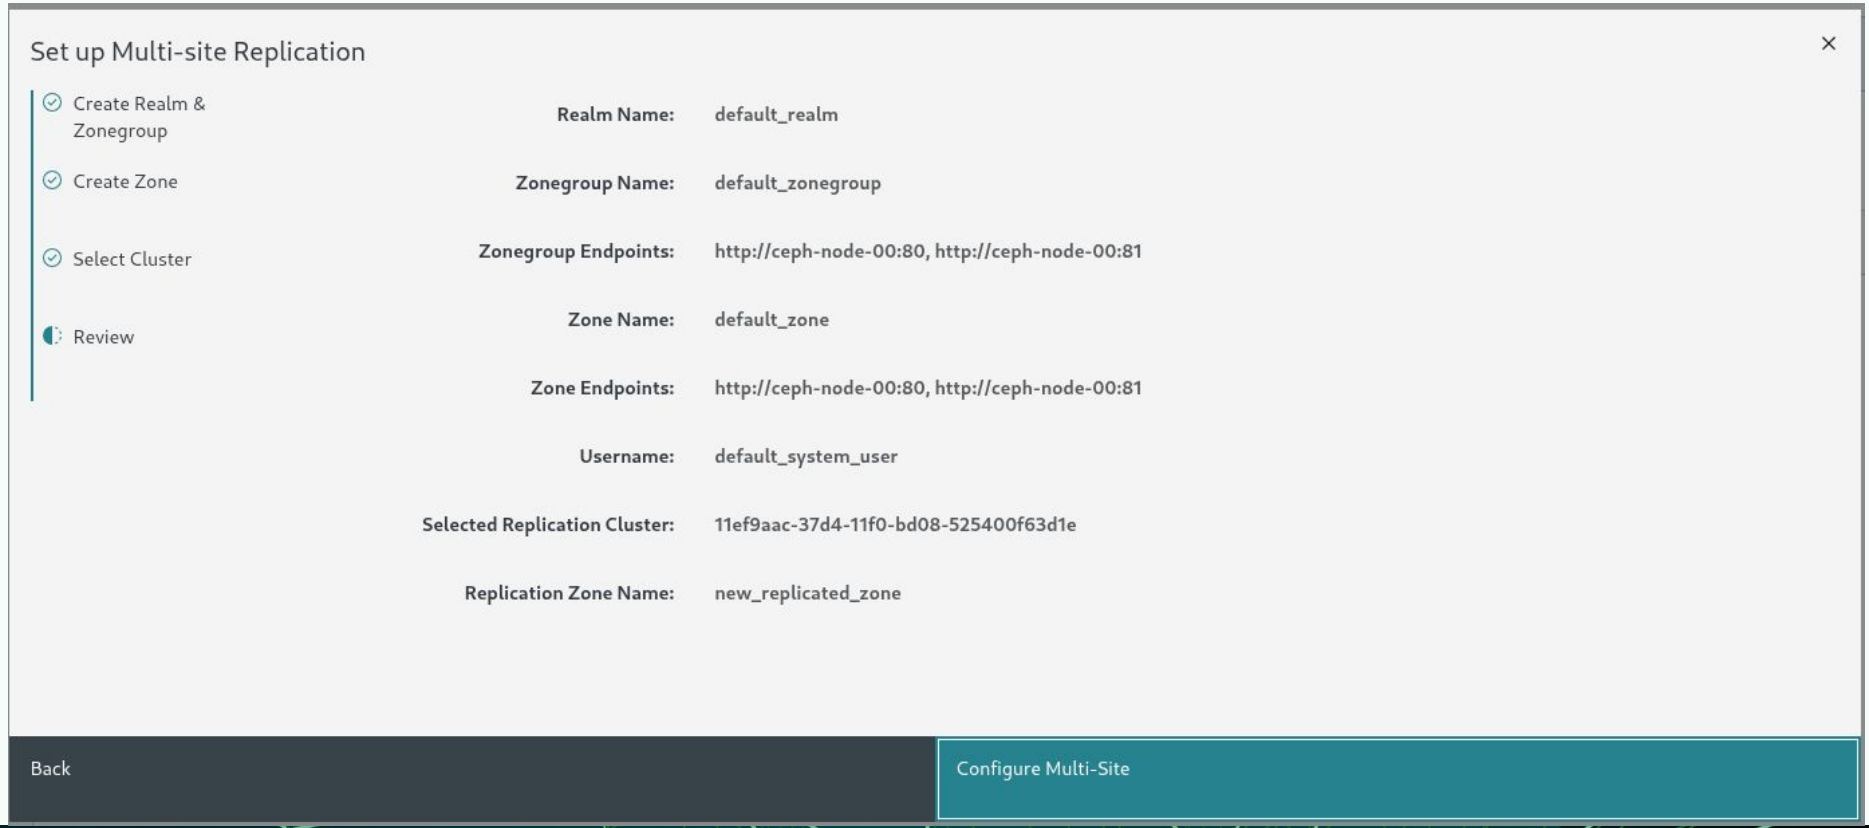

Object - Replication Automation - Progress

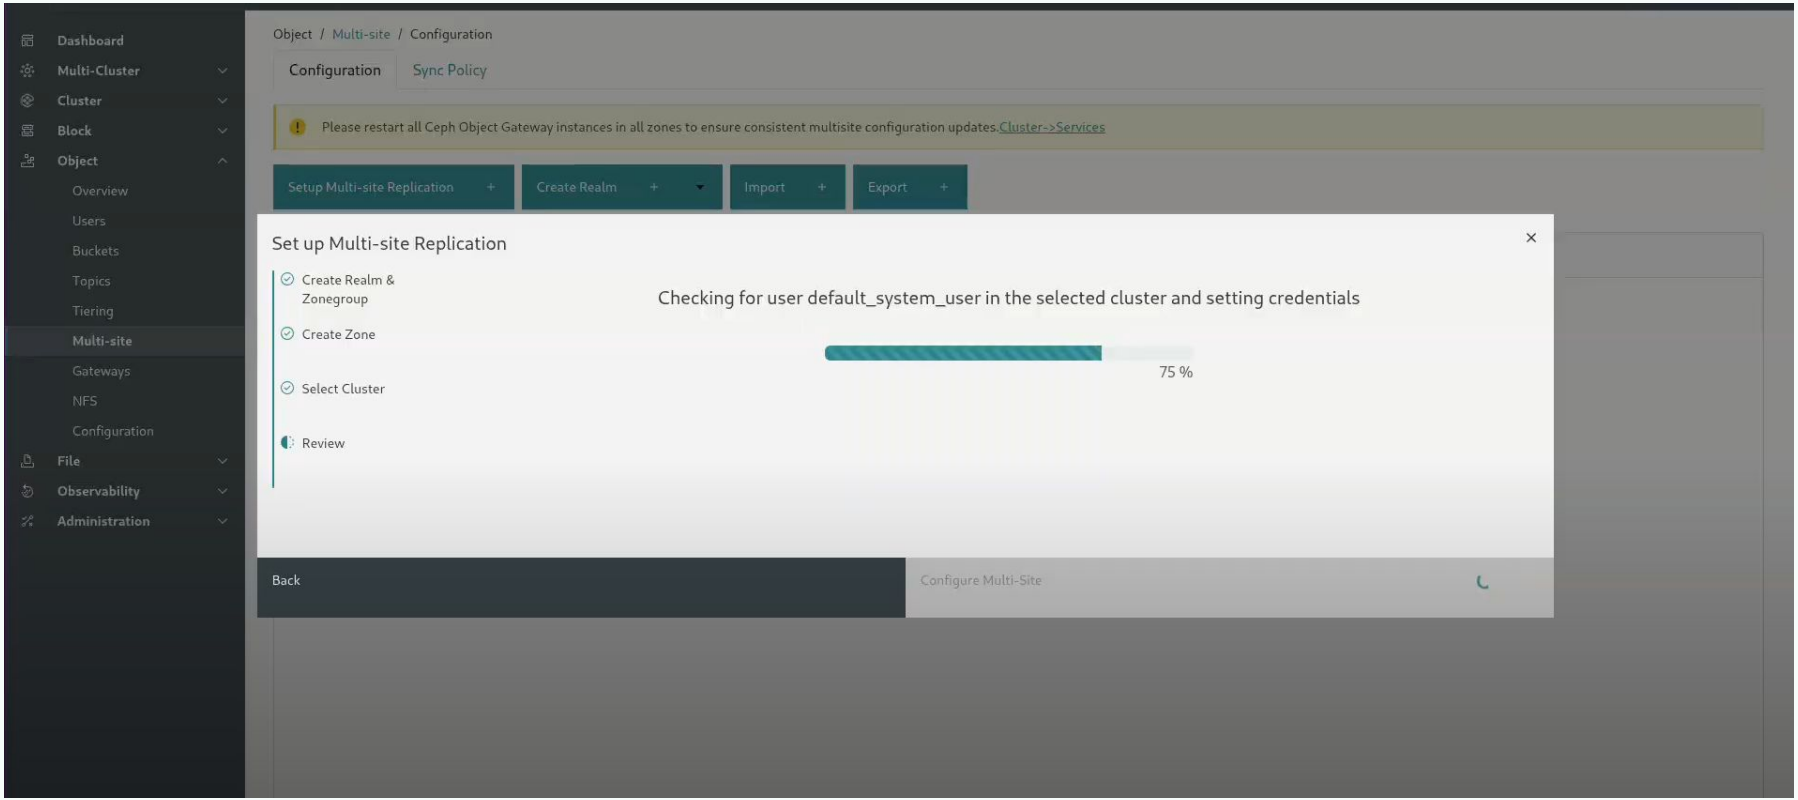

Object - Replication Automation - Final Result

Primary Site

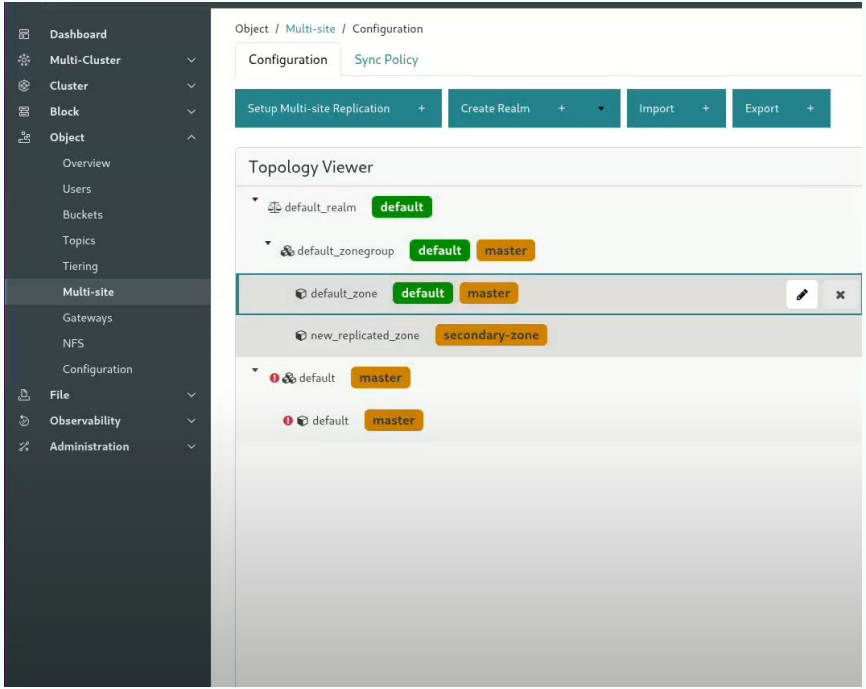

Secondary Site

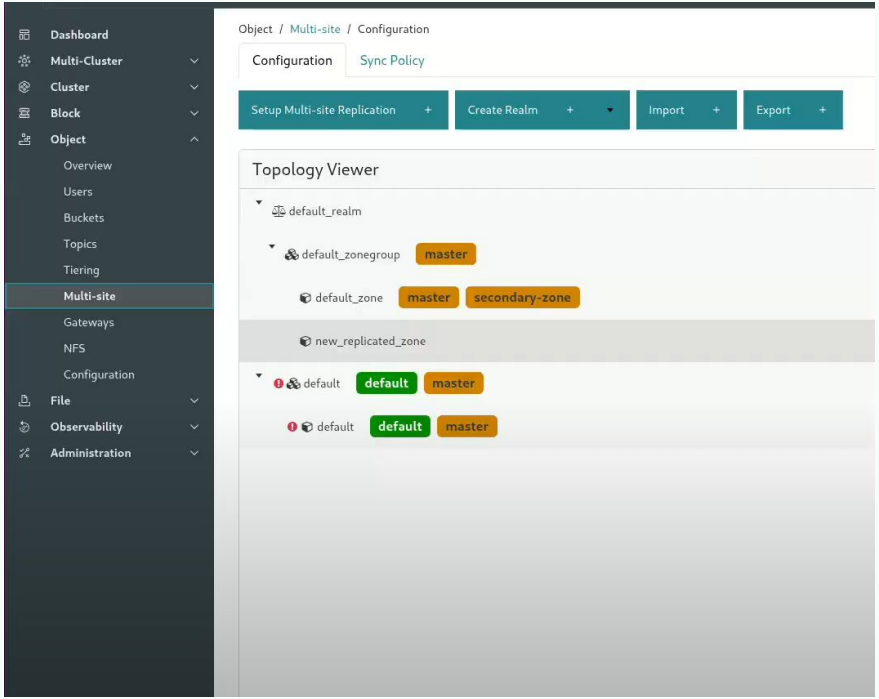

Object- Granular Sync

Replication Policy

Object - Granular Sync Policy Replication

Ceph's bucket-level granular sync policies offer more precise control—such as replicating specific buckets instead of entire zones—but have traditionally been complex and cumbersome to manage using CLI commands.

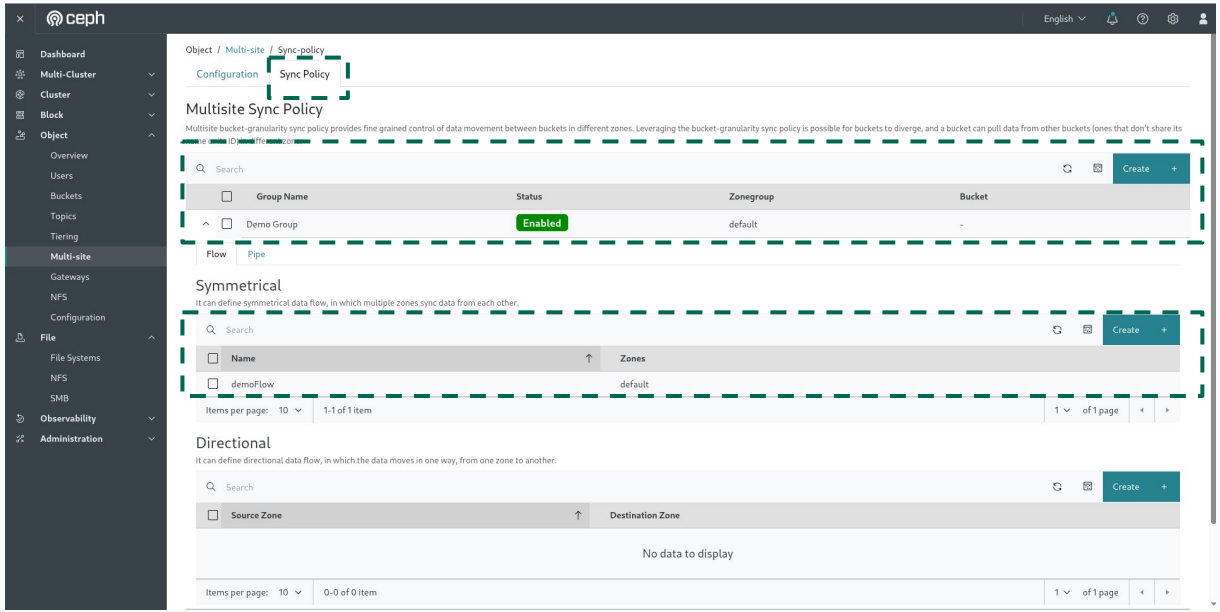

Groups: Hold data flow and pipe rules within a sync policy.

Data Flow: Can be **symmetrical** (bi-directional) or **directional** (one-way) between zones.

Pipes: Connect specific buckets and zones to data flows, defining what and how to sync.

Blog: https://ceph.io/en/news/blog/2025/rgw-multisite-sync-policy-dashboard/

Object - Granular Sync Policy Replication

Pipe

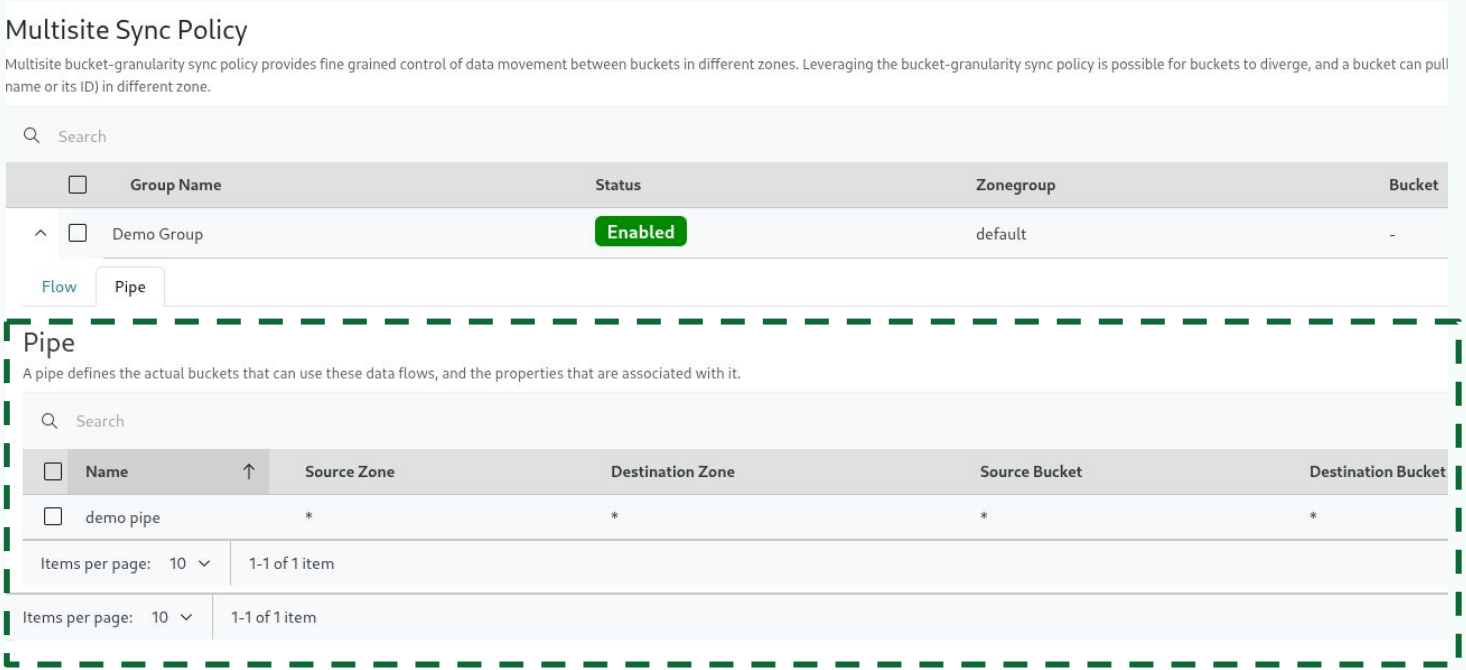

Object - Granular Sync Policy Replication

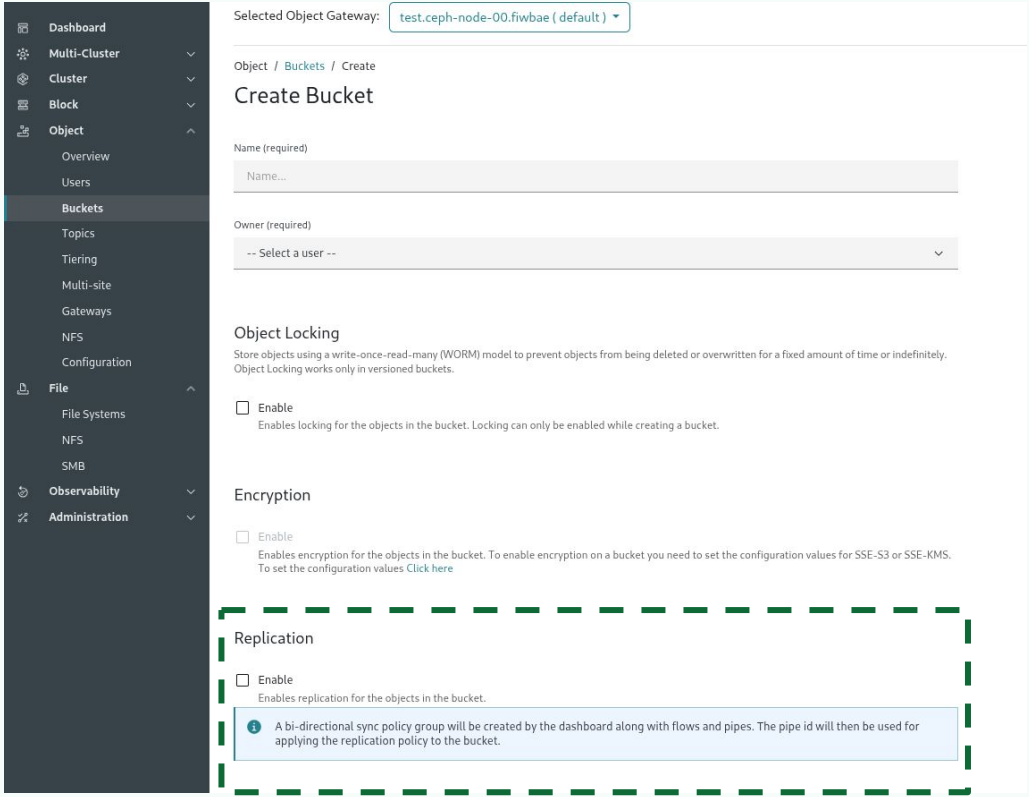

Bucket Level Replication

Object- Cloud Tiering

Object - Cloud Tiering

Ceph Object(RGW) supports object storage tiering, allowing seamlessly moves data between high-performance and archival tiers across on-prem and cloud environments.

Dashboard Features

- Create/List/Delete Cloud Tiering storage class
- Manage policy based archival for object in bucket
- Monitoring Lifecycle status

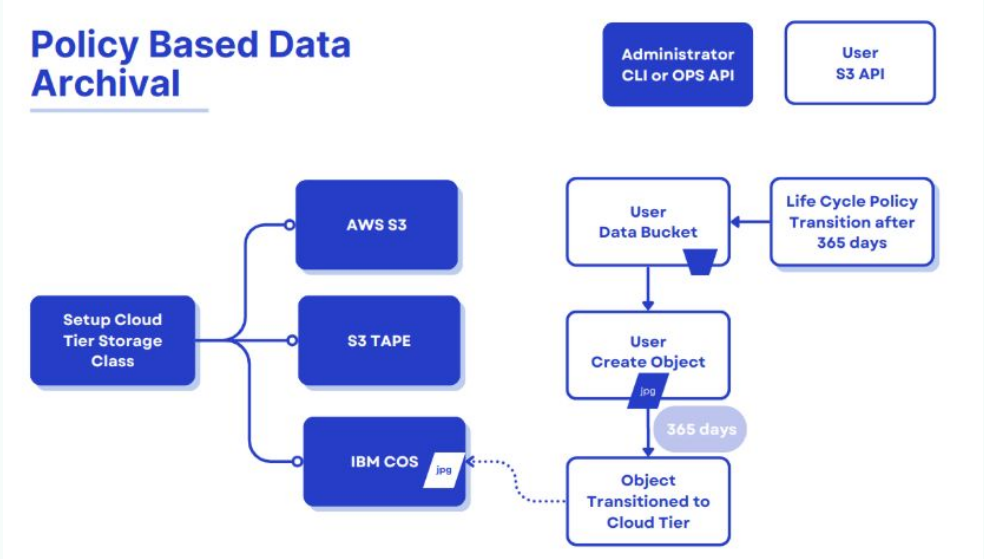

Object - Cloud Tiering - Create

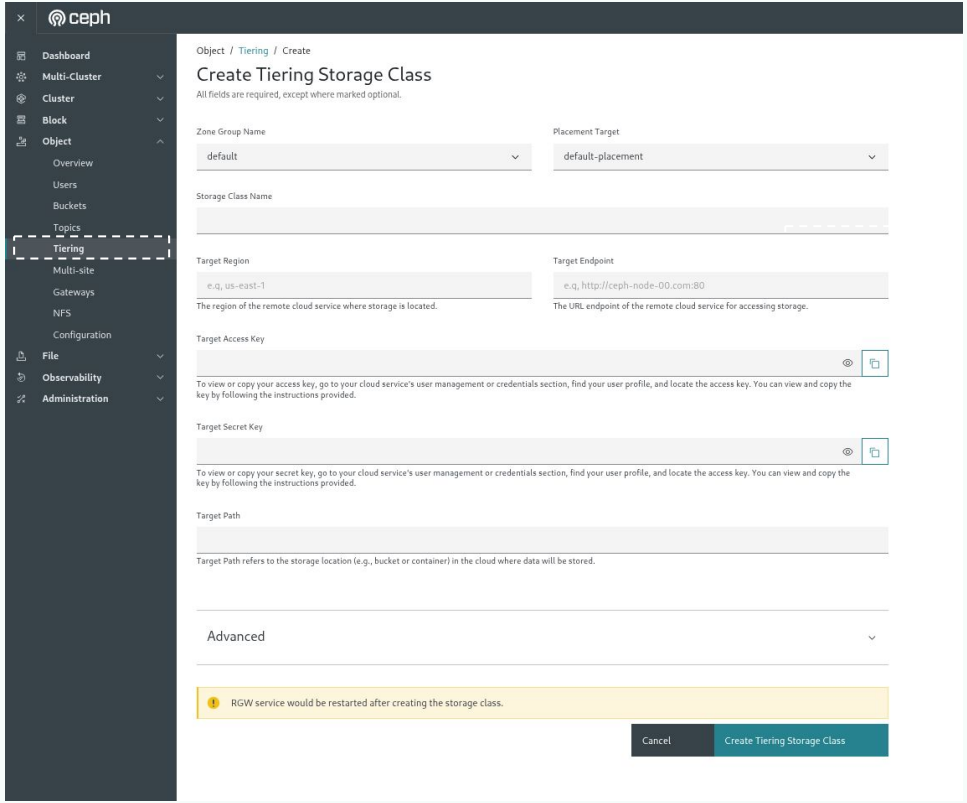

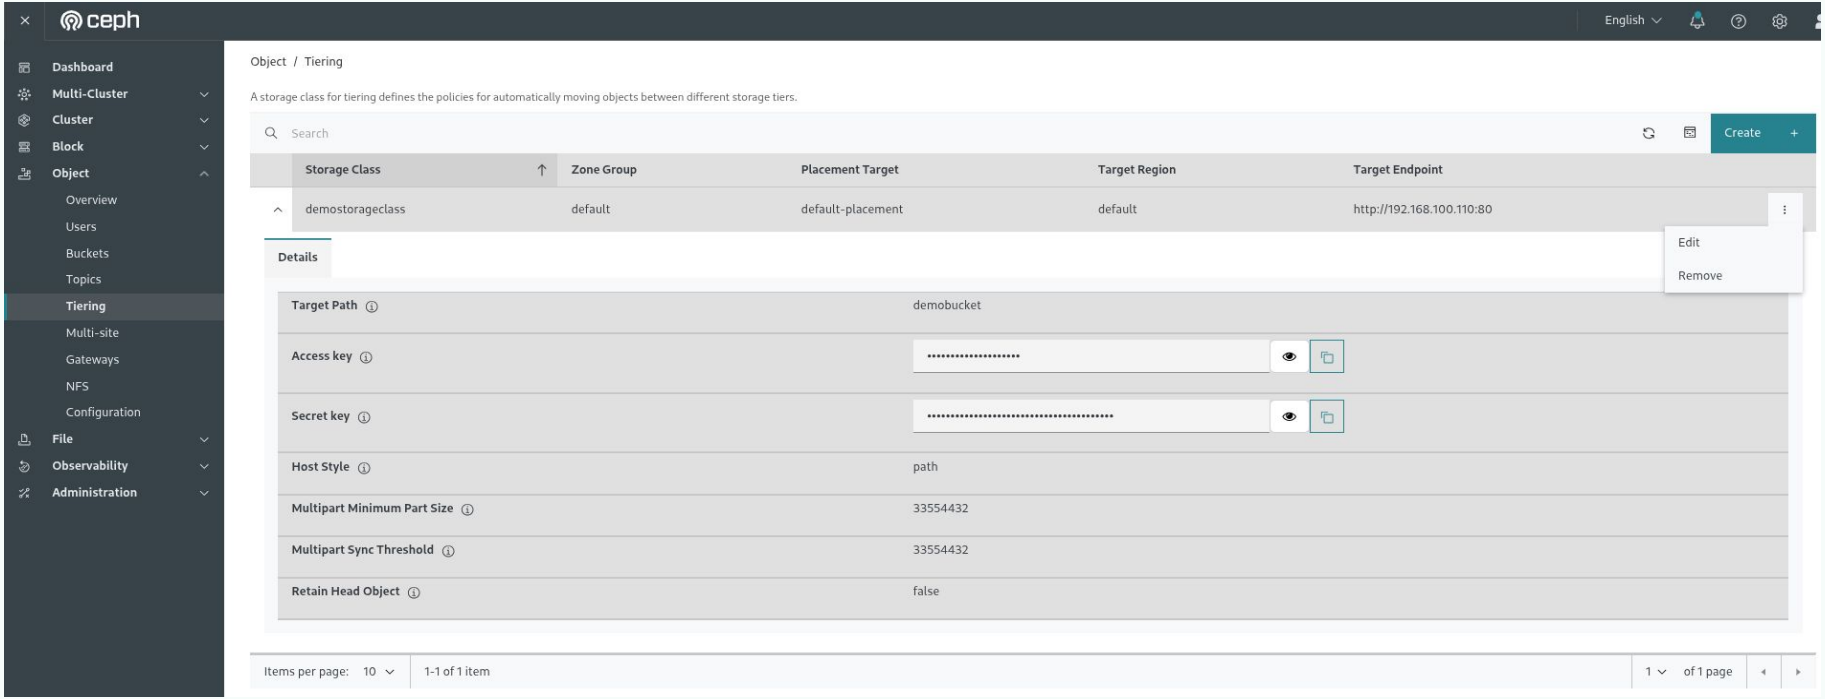
ceph
English
Dashboard
Multi-Cluster
Cluster
Block
Object
Overview
Users
Buckets
Topics
Tiering
Multi-site
Gateways
NFS
Configuration
File
Observability
Administration
Object / Tiering
A storage class for tiering defines the policies for automatically moving objects between different storage tiers.
Search
Create
Storage Class
Zone Group
Placement Target
Target Region
Target Endpoint
demostorageclass
default
default-placement
default
http://192.168.100.110:80
Edit
Remove
Details
Target Path
demobucket
Access key
Secret key
Host Style
path
Multipart Minimum Part Size
33554432
Multipart Sync Threshold
33554432
Retain Head Object
false
Items per page: 10
1-1 of 1 item
1
of 1 page

Object - Cloud Tiering - Screenshot

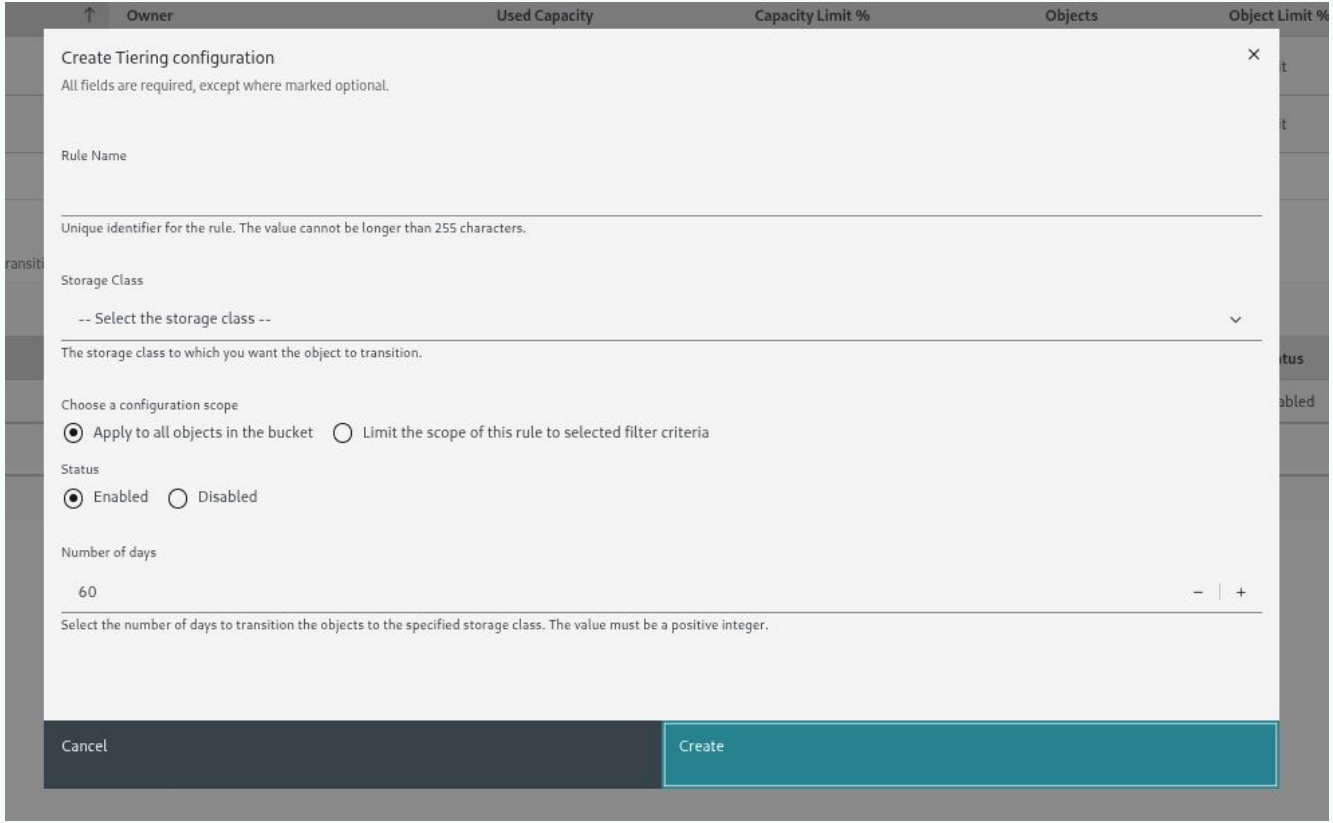

Object - Cloud Tiering - Screenshot

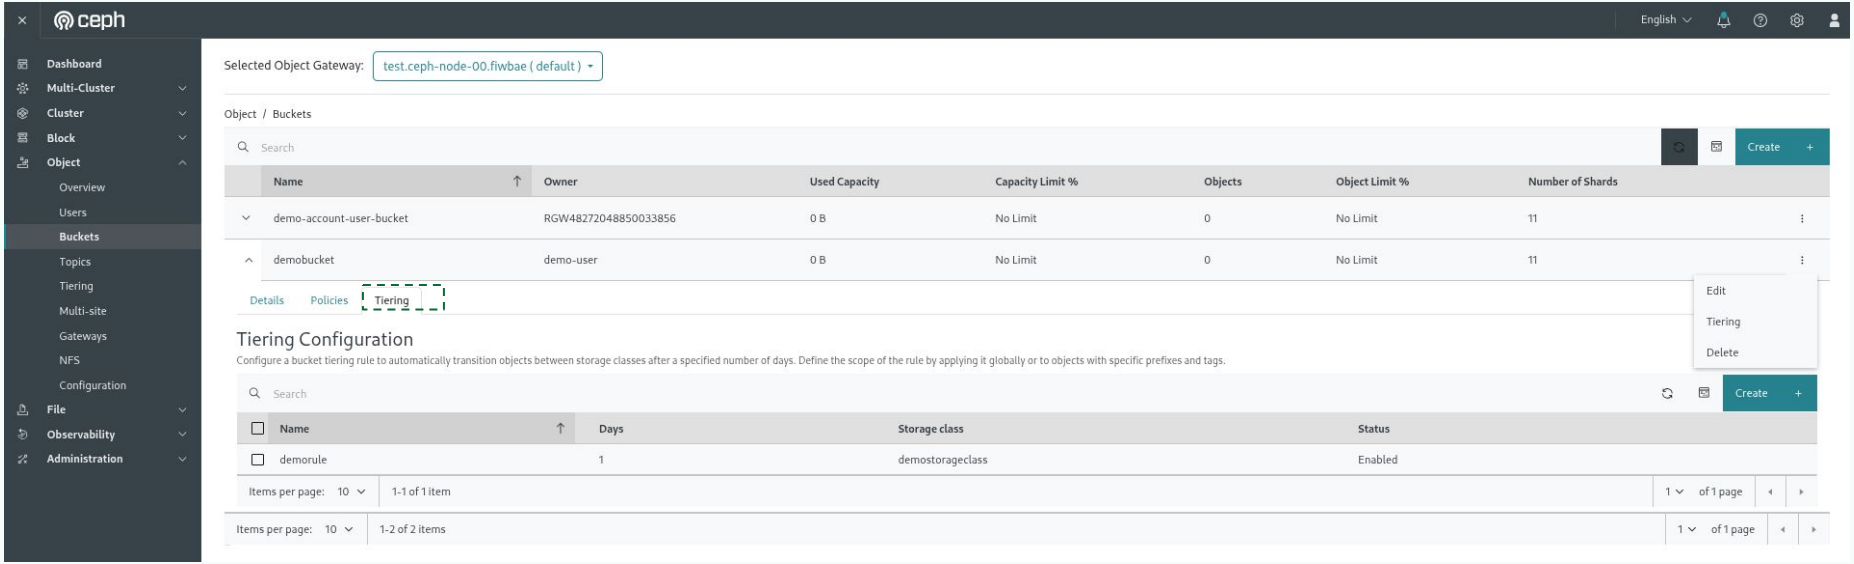

Object - Cloud Tiering - Screenshot

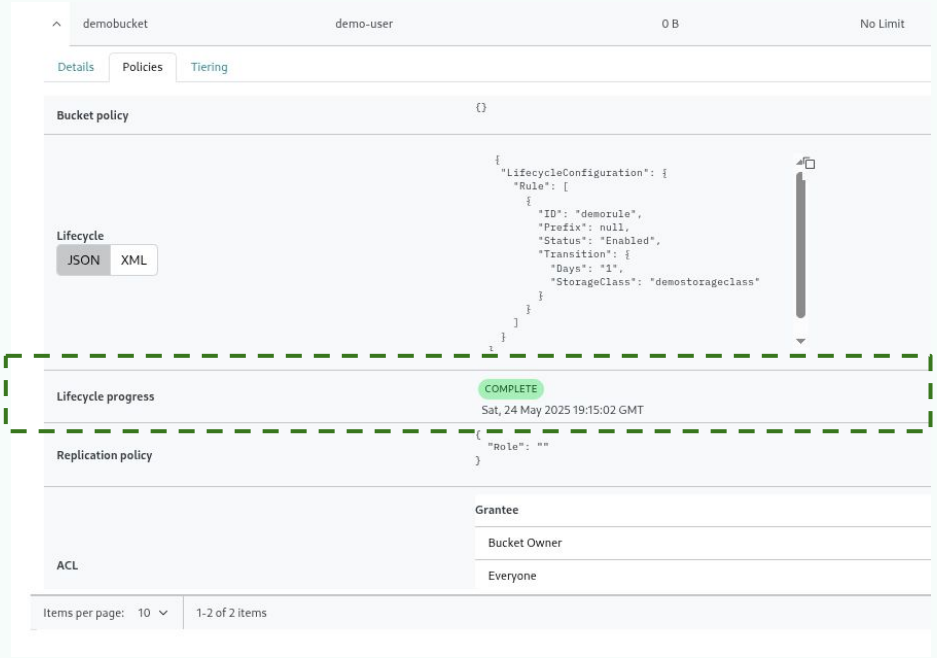

Object - Cloud Tiering - Demo

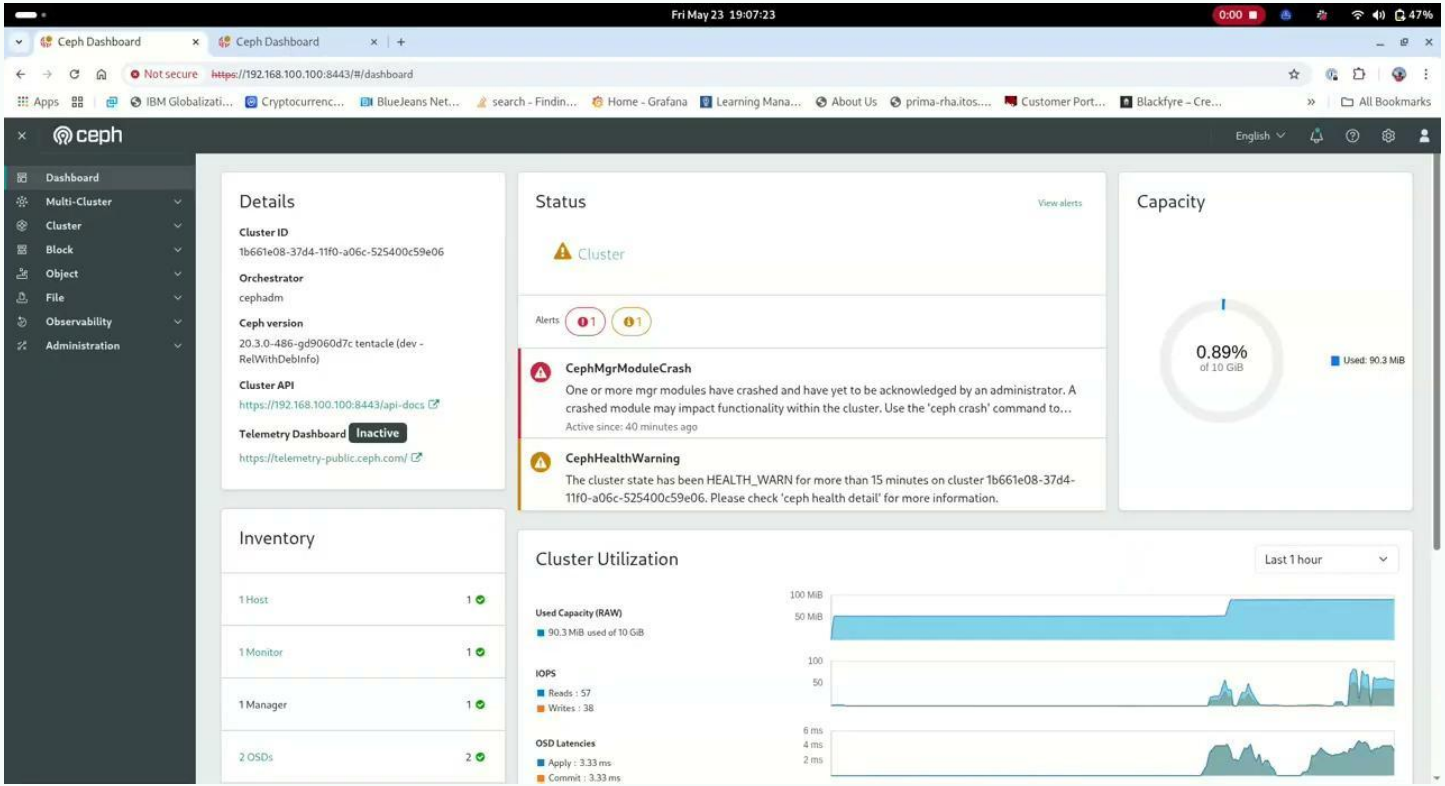

Object- S3 Account Management

Object - S3 Account Management

Ceph now supports AWS-compatible IAM accounts, enabling secure, self-service multitenancy for object storage. This feature empowers tenants to manage their own resources while reducing administrative effort.

Dashboard Features

- Create/List/Edit/Delete Account
- Add user to account
- Migrate an existing User into an Account
- Account Stats/Quota Management
- Creating a bucket into an account
- Migrating an existing bucket to an account

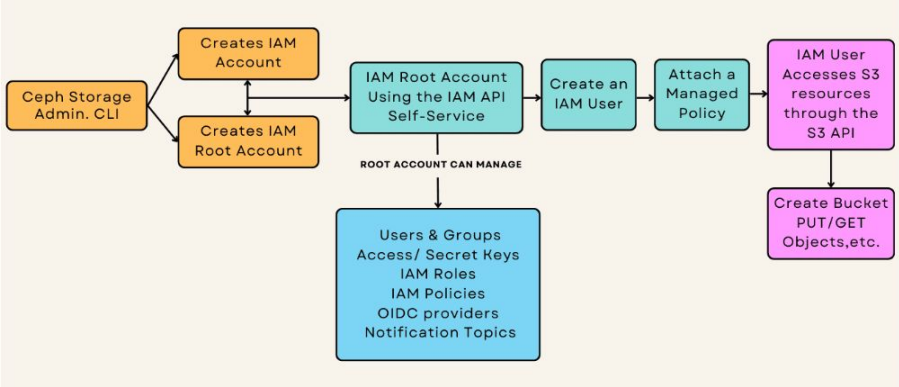

S3 Account Management - Account CRUD

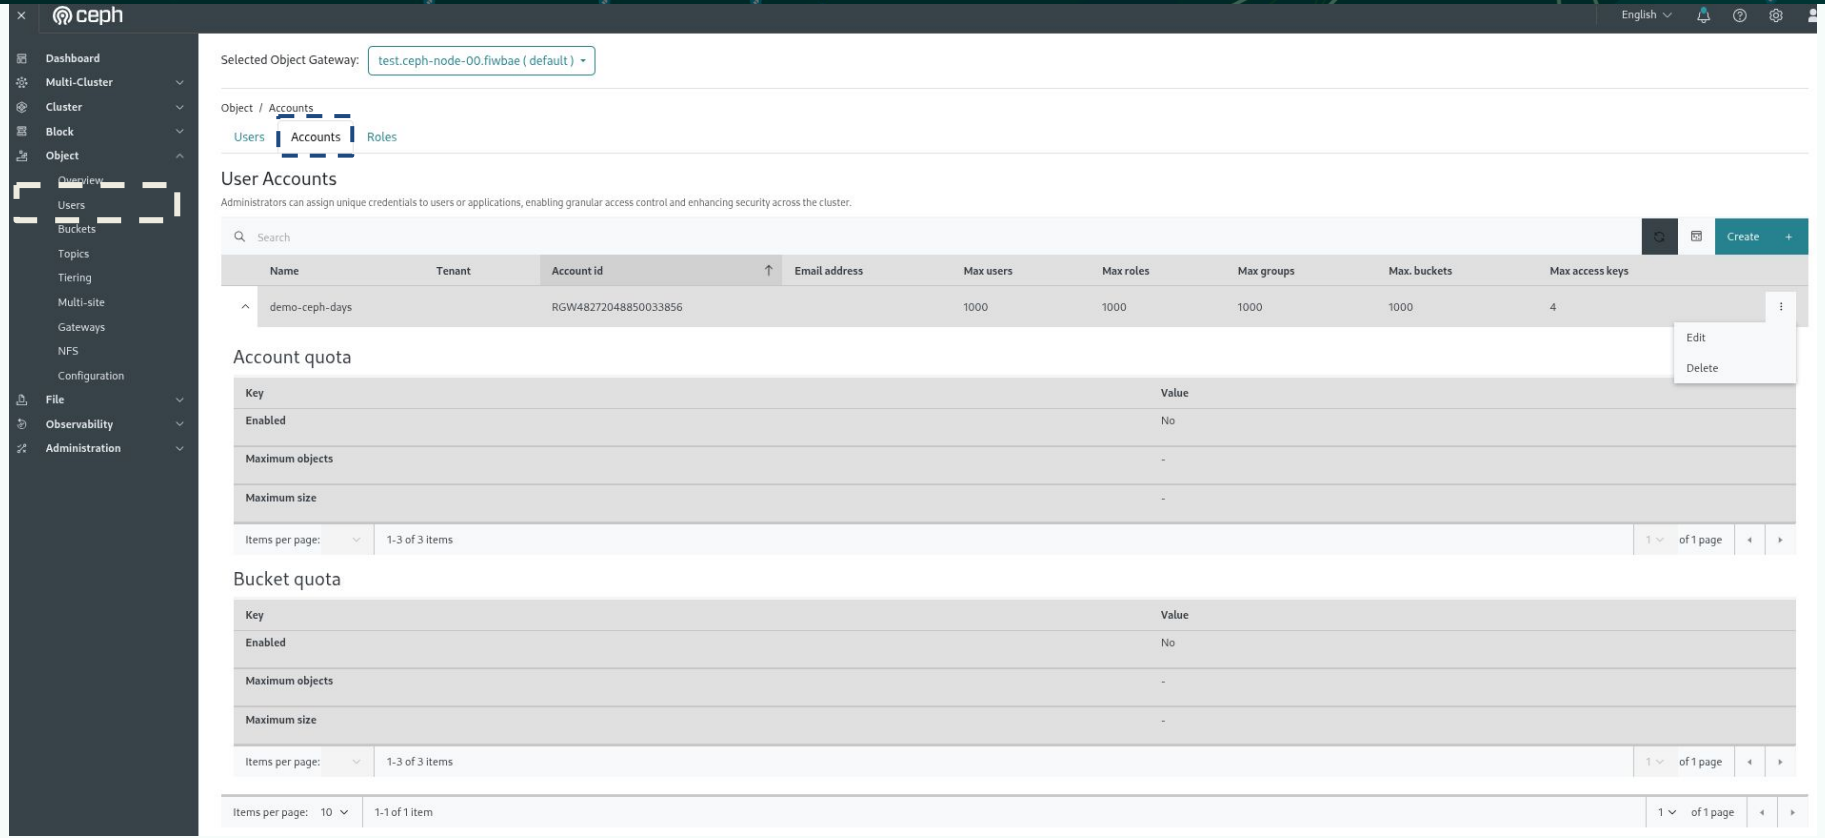

S3 Account Management - Account CRUD

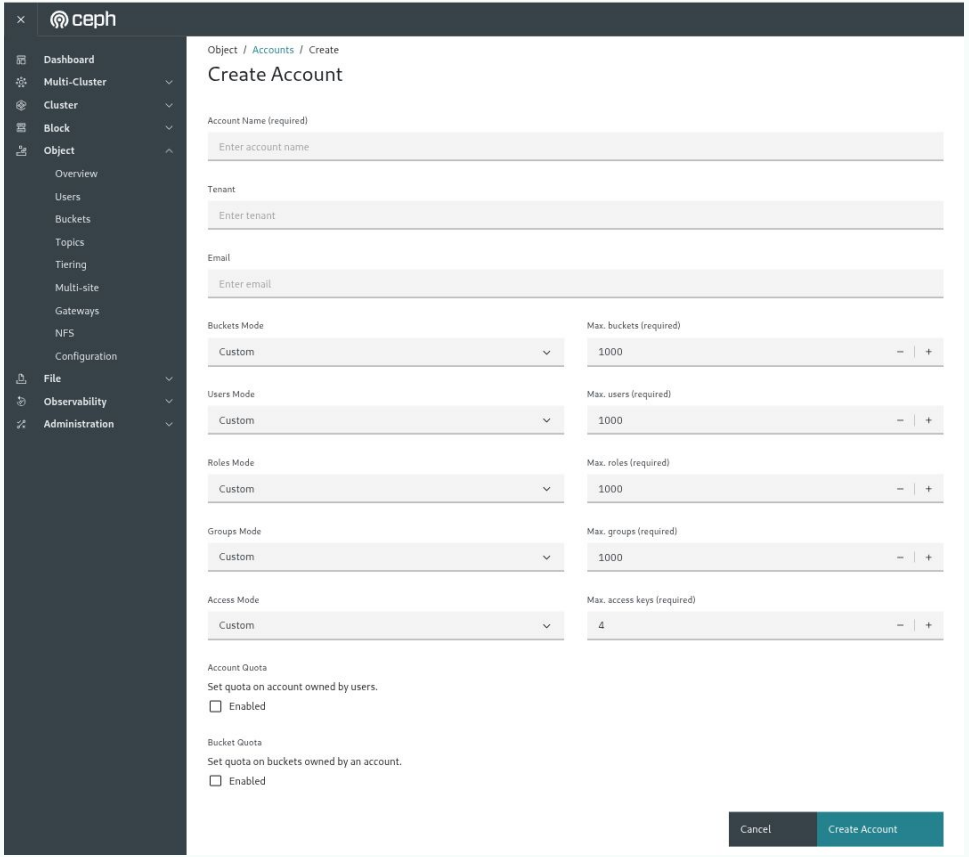

S3 Account Management - User CRUD

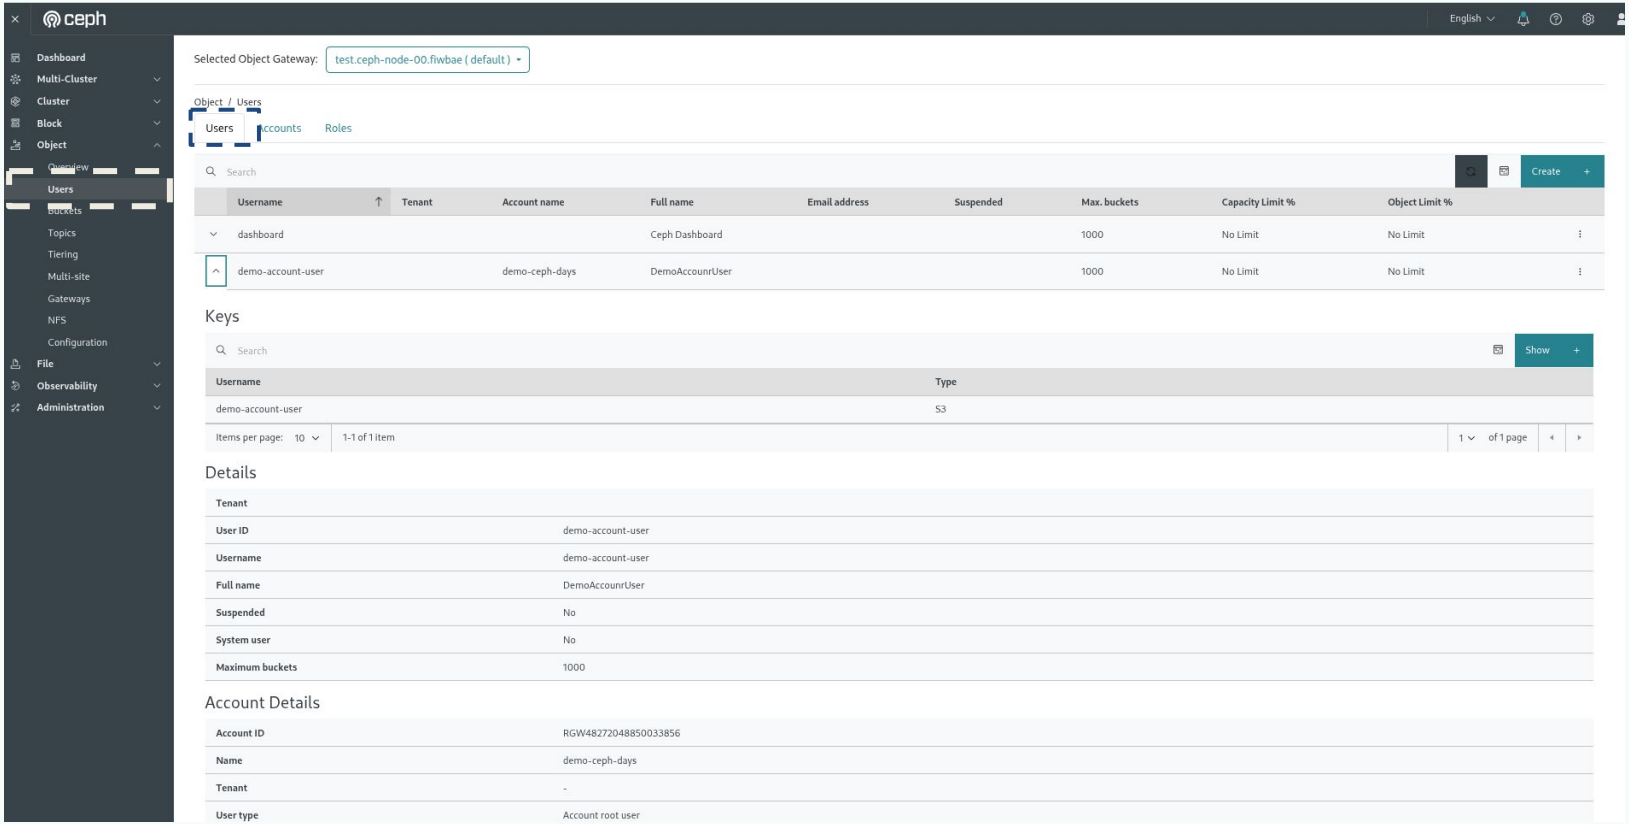

S3 Account Management - User CRUD

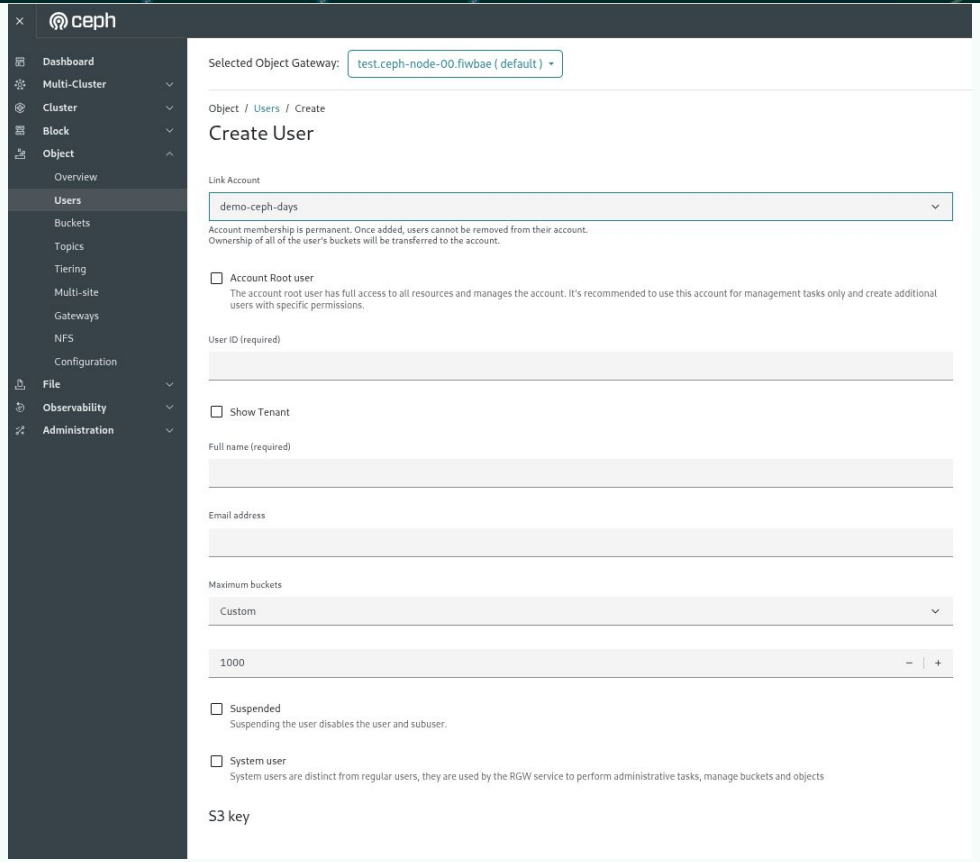

S3 Account Management - Bucket Linking

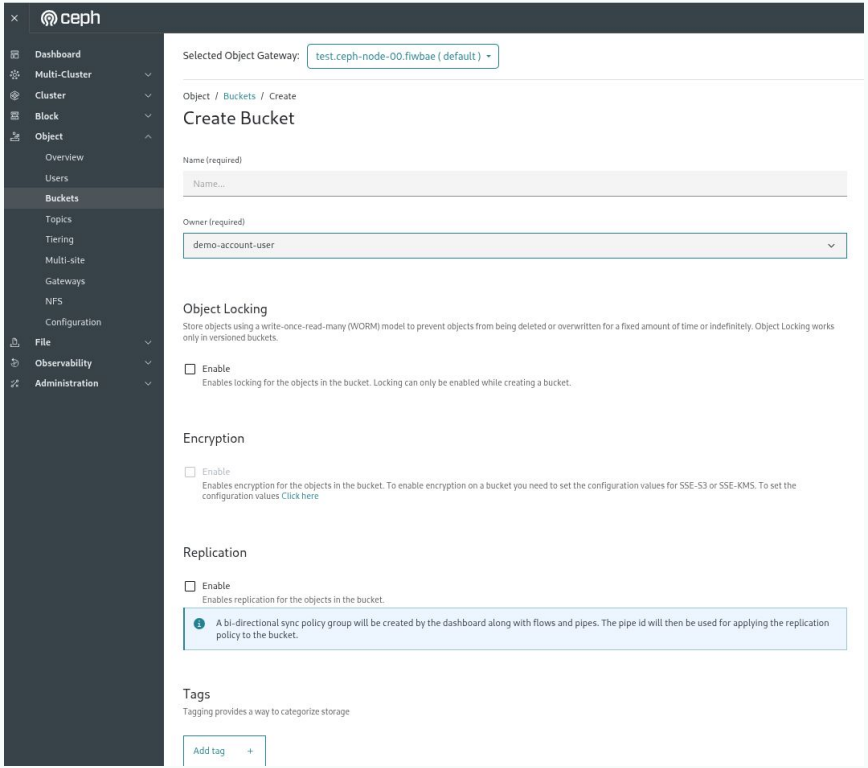

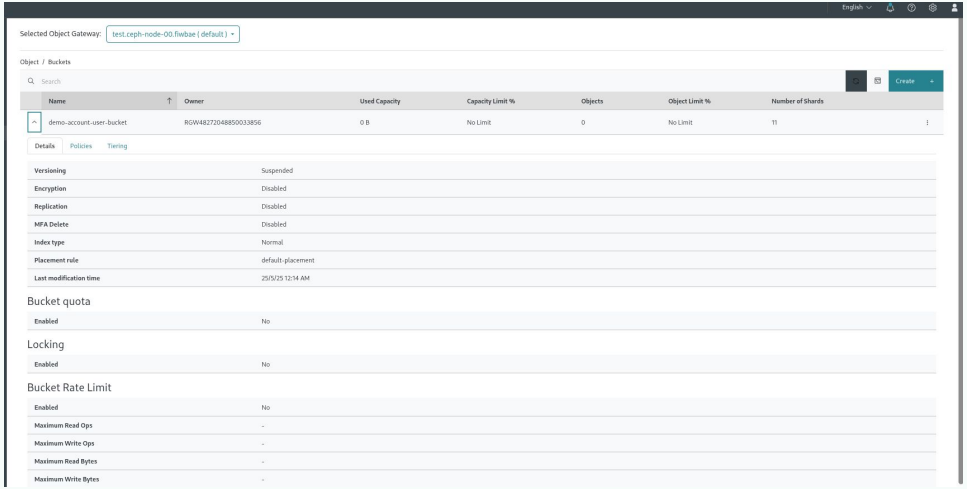

Object - S3 Account Management - Demo

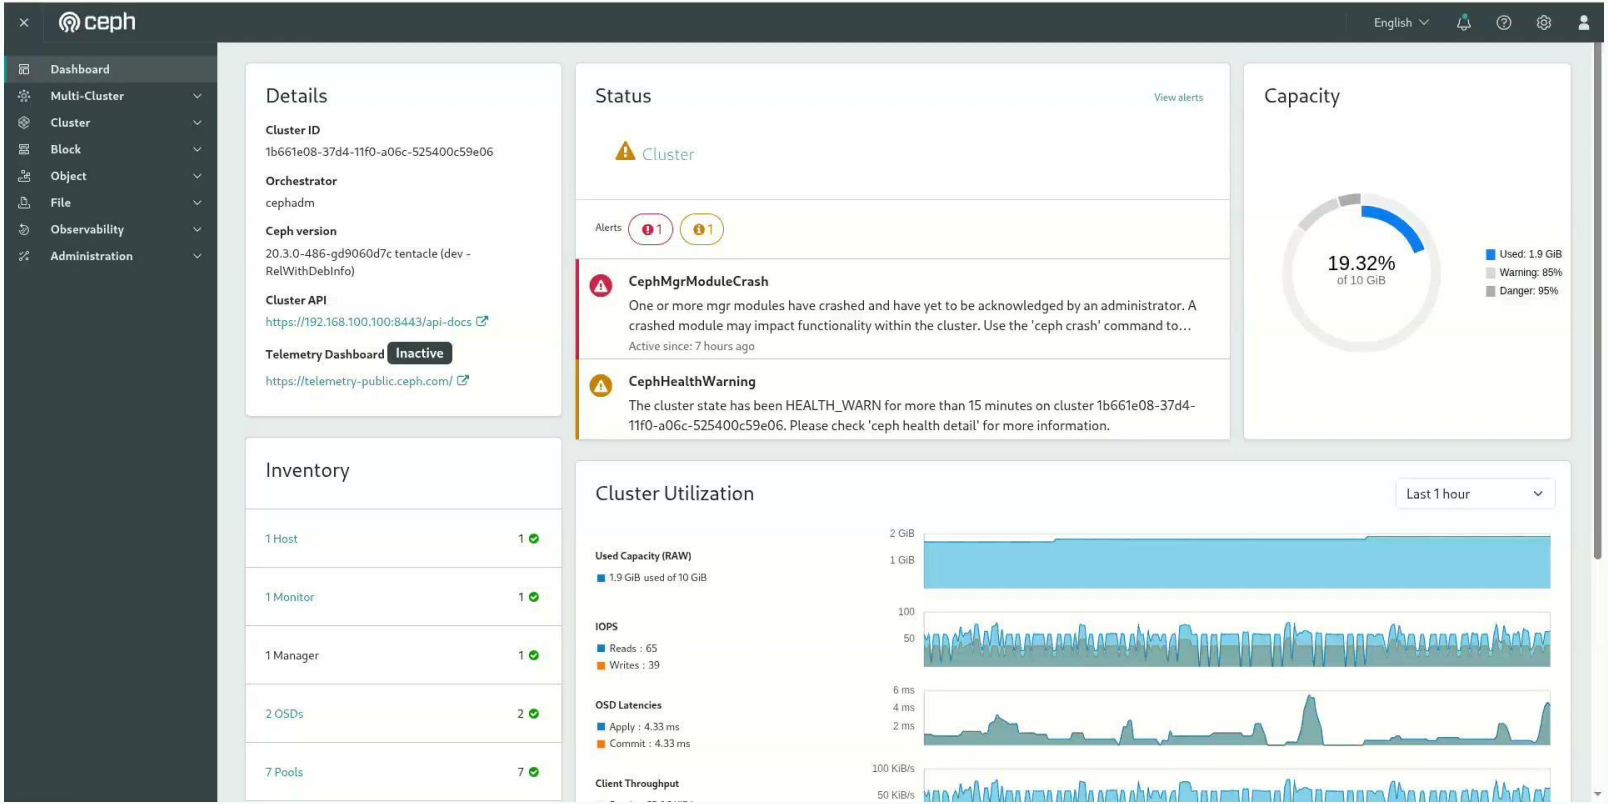

Object - Per User/Bucket Rate Limiting

Ceph Dashboard now supports RGW **rate limiting** at a **per-user** and **per-bucket** level, allowing administrators to control traffic and protect backend resources from overuse or abuse.

Dashboard Features:

- Individual users
- Specific buckets

Configure:

- **Request rate** (e.g., Read, Write)
- **Bandwidth limits** (Read, Write in bytes/sec)

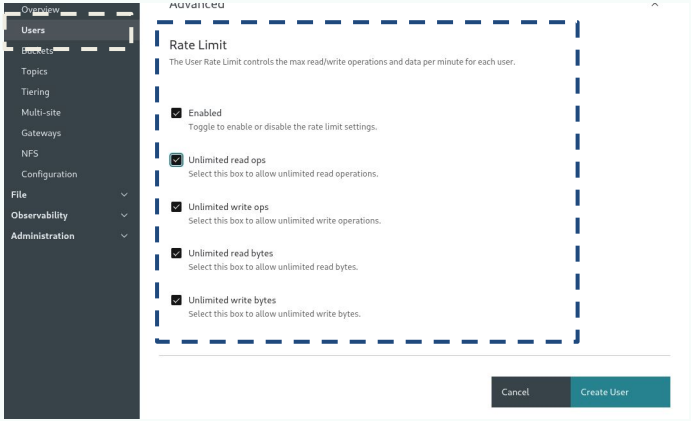

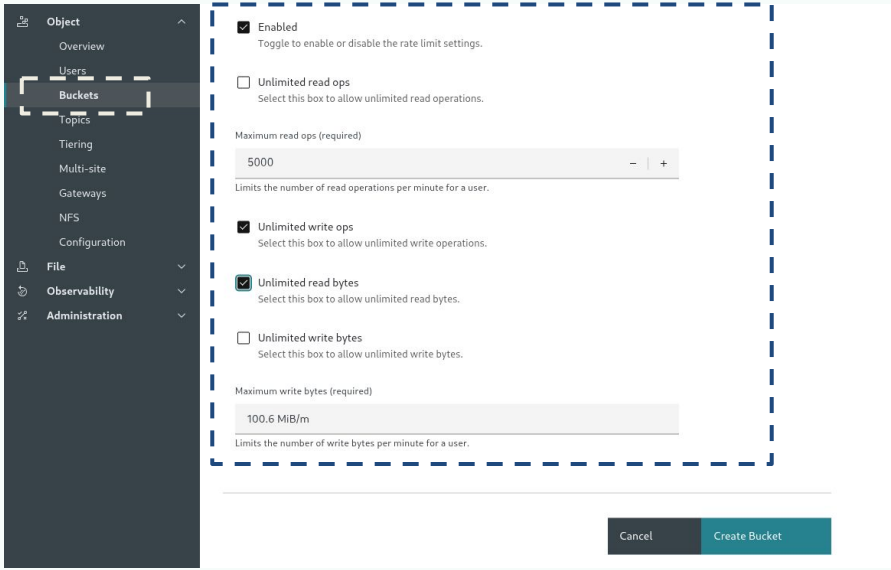

Management Gateway

With OAuth2

Management Gateway with OAuth2

Ceph introduces a new service called **Management Gateway** to unify access to management services, now enhanced with **OAuth 2.0** for secure, standards-based authentication.

Disadvantages of the legacy System

- Users has to know the ip:port for every monitoring service
- SSO support is reduced to SAML (and for Dashboard only)
- Only basic-auth and different user/password is needed for each service
- No Uniform authentication
- No high-availability support for monitoring

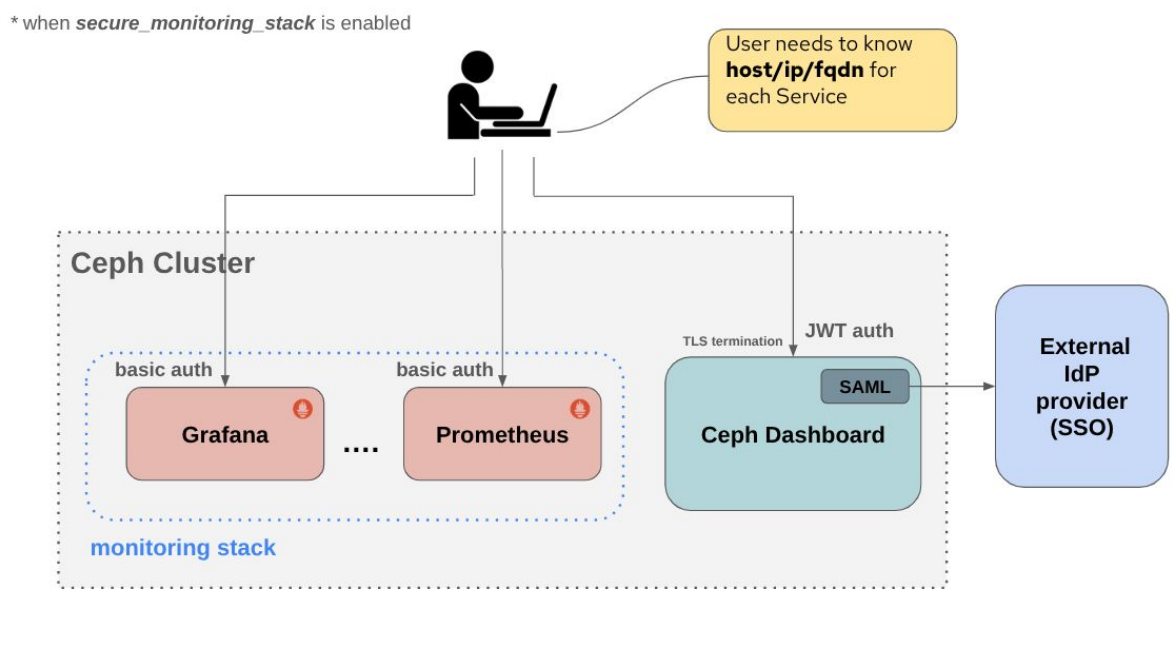

Management Gateway with OAuth2

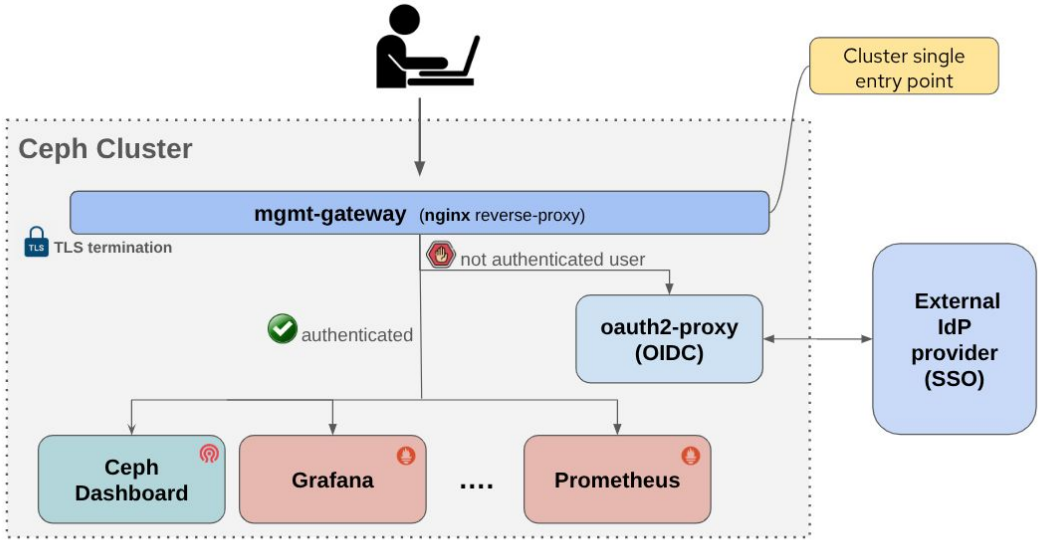

Advantages of the new architecture

- Single entry-point to all the cluster services
- Modern SSO support based on OIDC protocol (possibility of MFA, better security, etc)
- Uniform SSO authentication across all the cluster monitoring services
- Modern SSO infrastructure suited for modern hybrid environments
- Monitoring high-availability support

📁 File (CephFS) Improvements

File - SMB Management

Ceph now supports managing **SMB (Server Message Block)** shares directly through the Dashboard, simplifying file share setup

Dashboard Features:

- Creating SMB Cluster
- Managing users
  - Standalone user Management
  - Active Directory Management
- Create, Modify, Delete SMB Shares
- Overview Dashboard

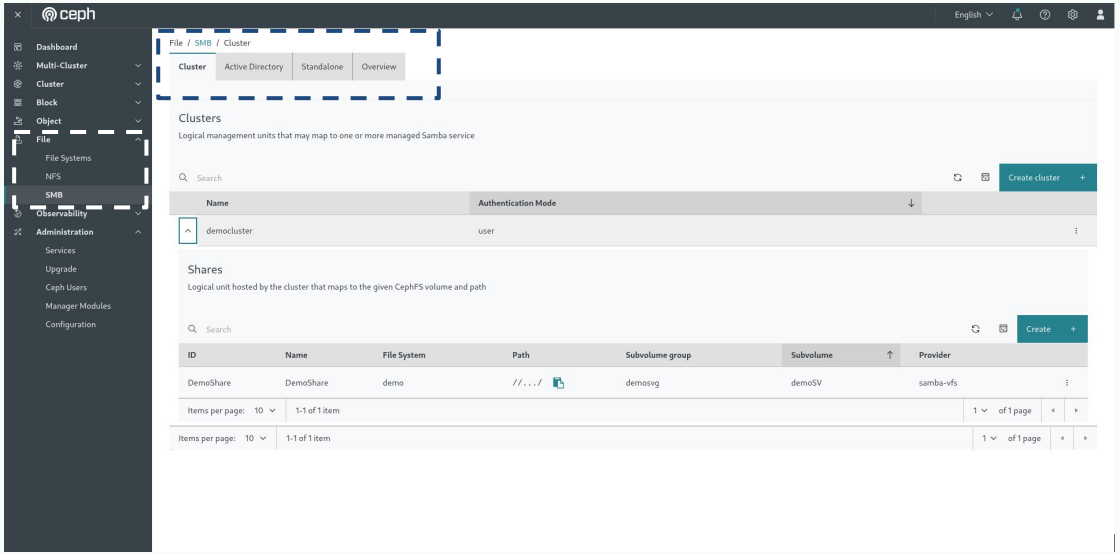

File - SMB Management - Create Cluster

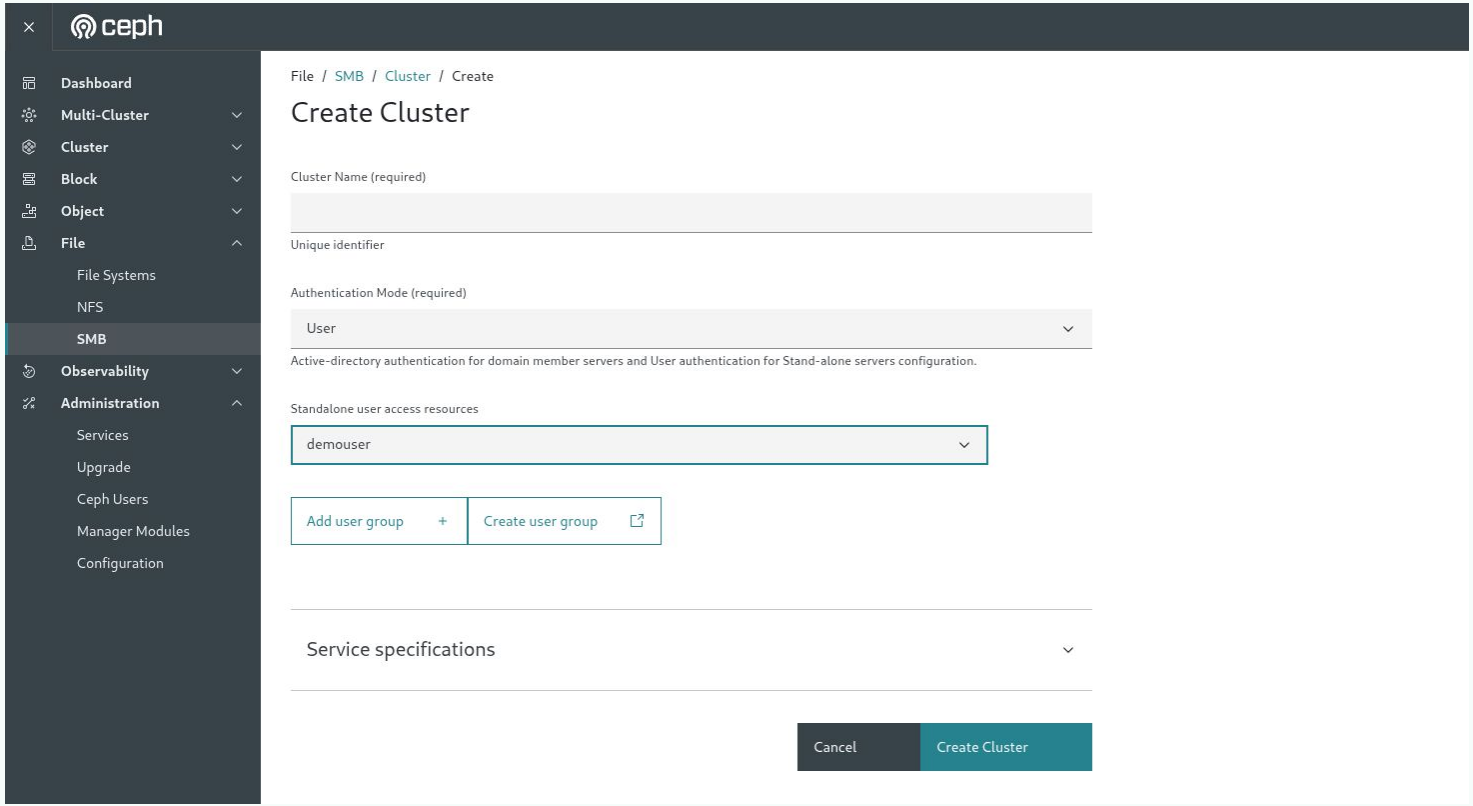

File - SMB Management - User Management

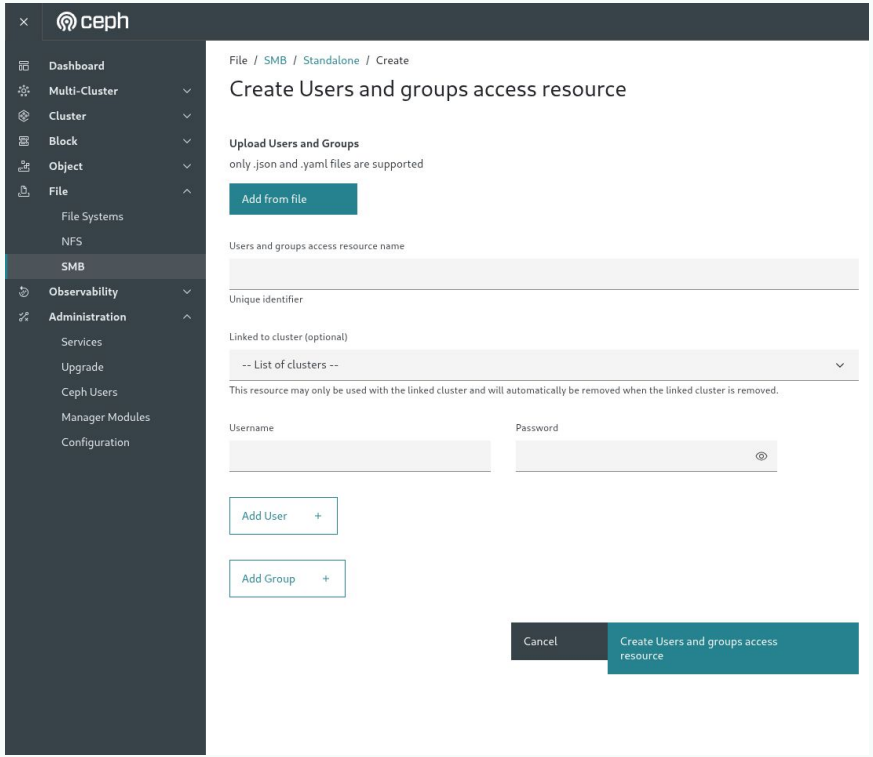

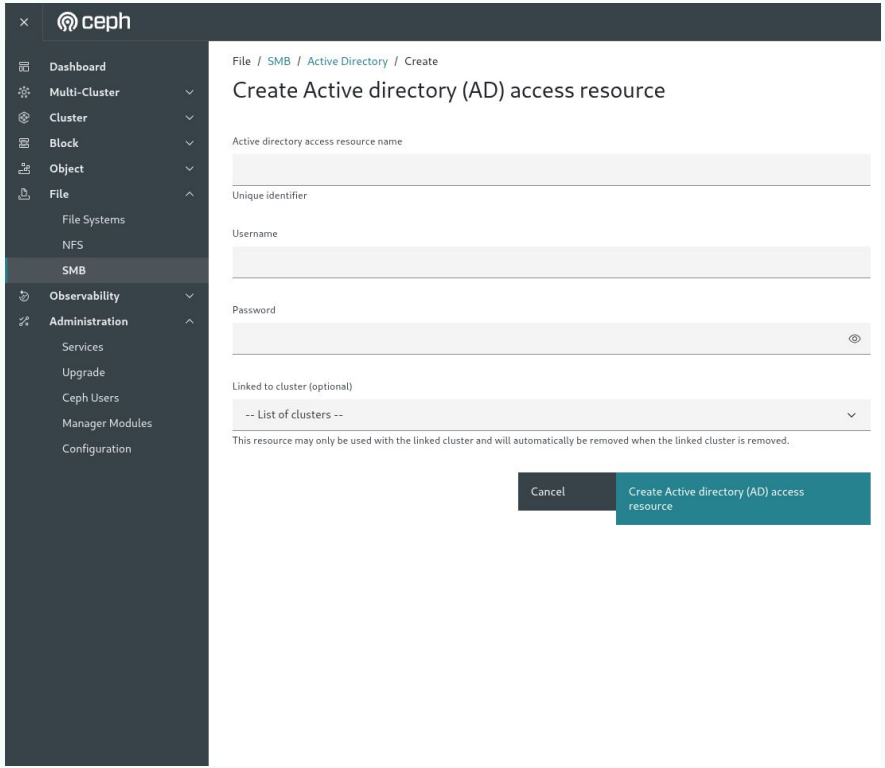

File - SMB Management - Monitoring

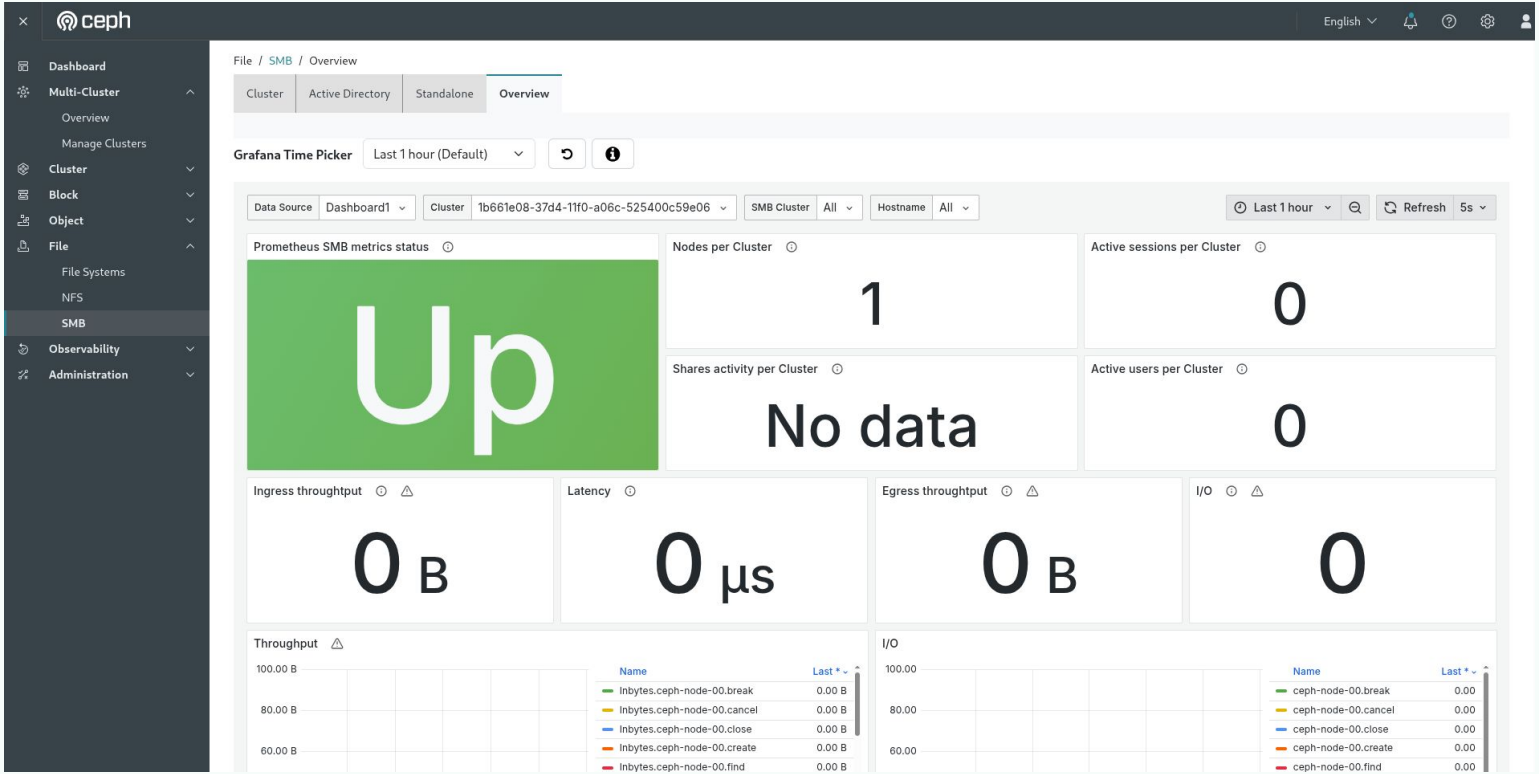

File - NFS UX update

Ceph Dashboard now offers a unified interface for managing both **SMB** and **NFS** file shares, streamlining file service provisioning for diverse client environments.

Dashboard Features

NFS management now mirrors SMB structure of Clusters -> Exports/Shares

GUI-driven, consistent setup for both protocols

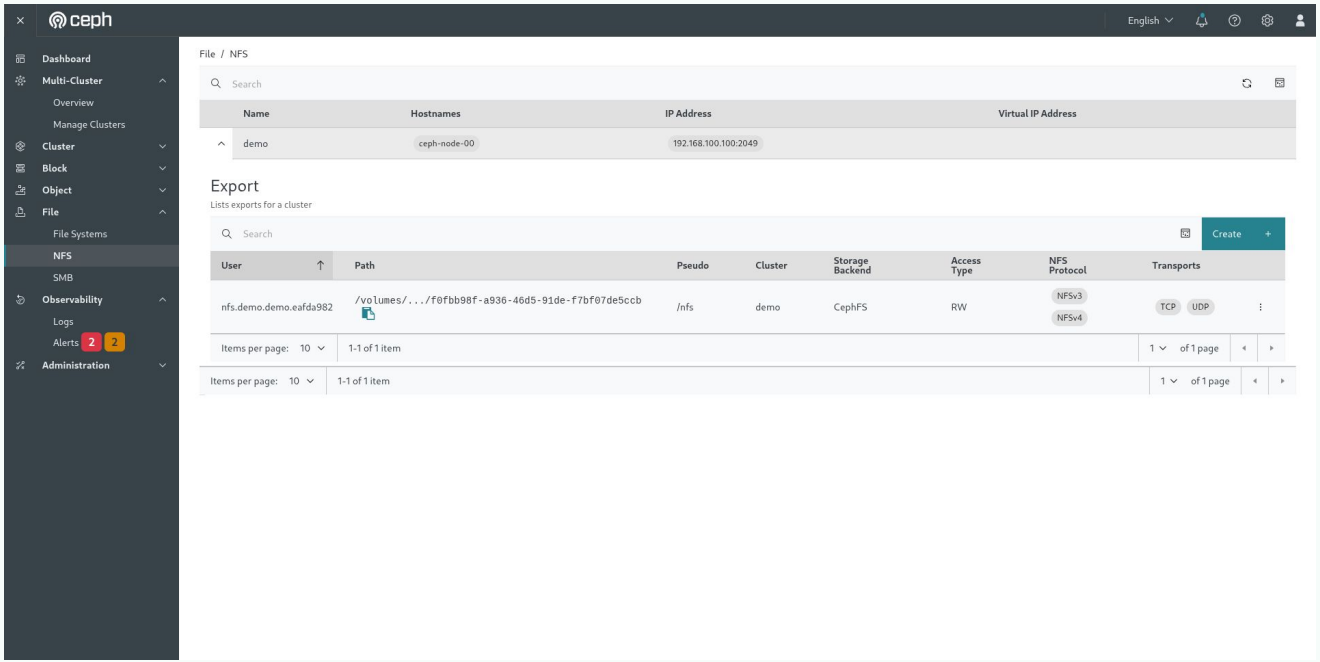

Monitoring Dashboards

Multi-Cluster Monitoring

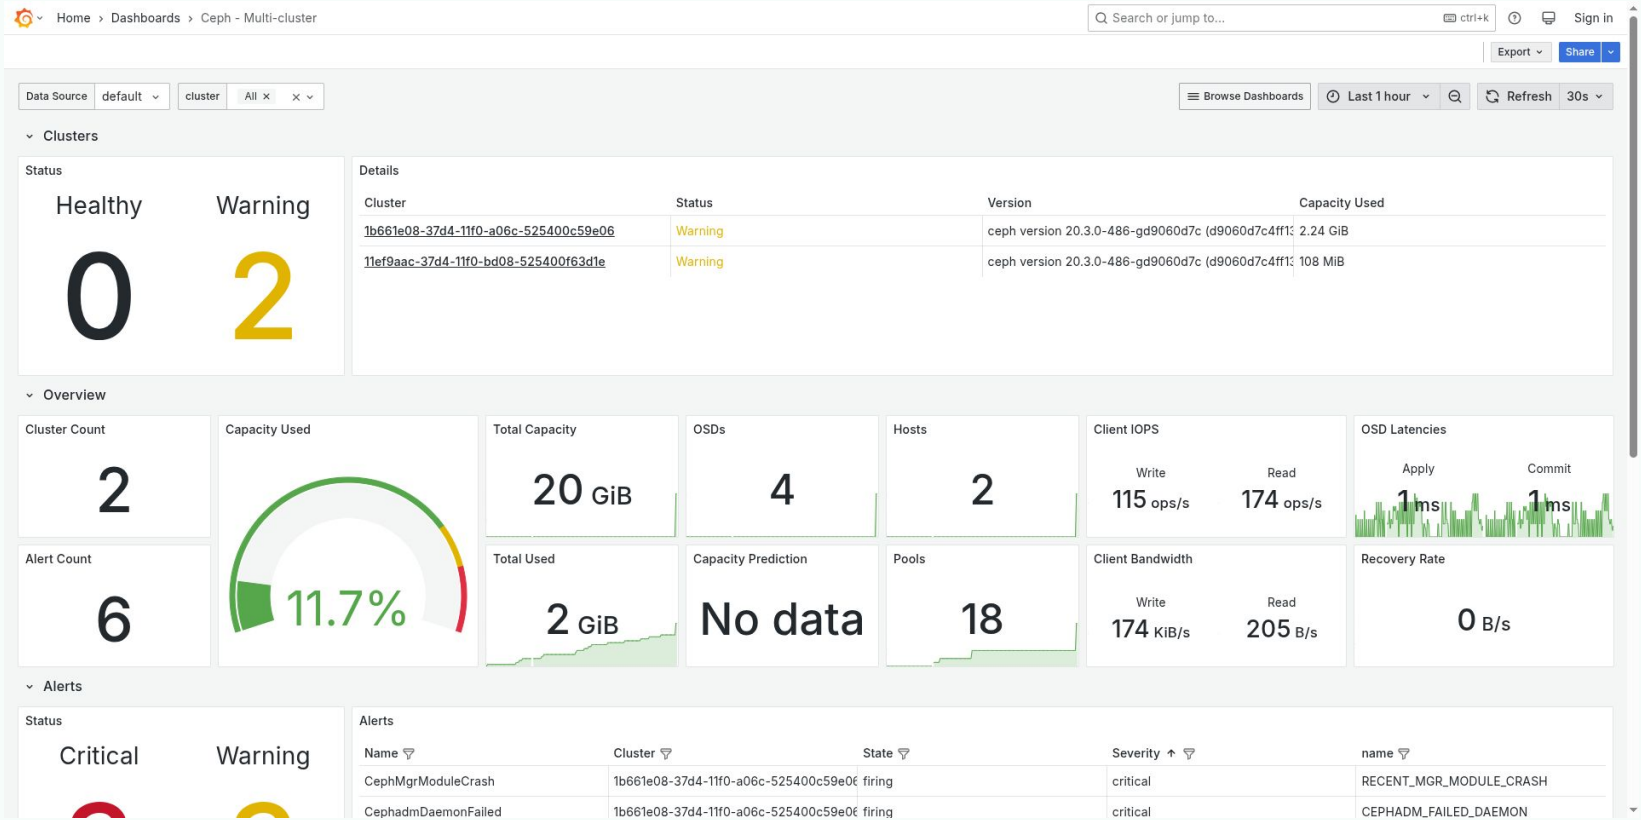

Application Overview

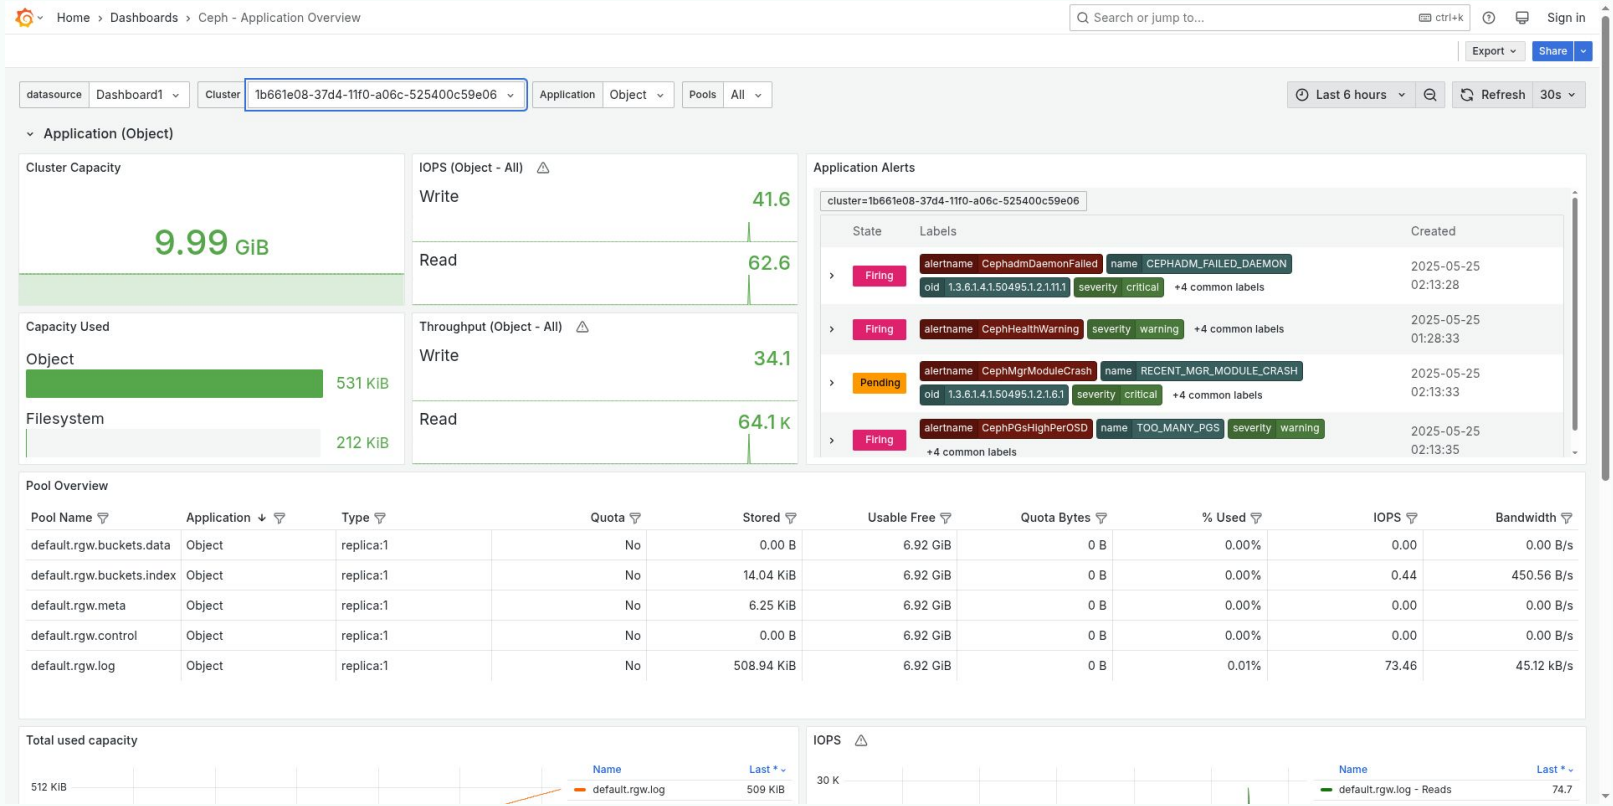

NVMe Overview

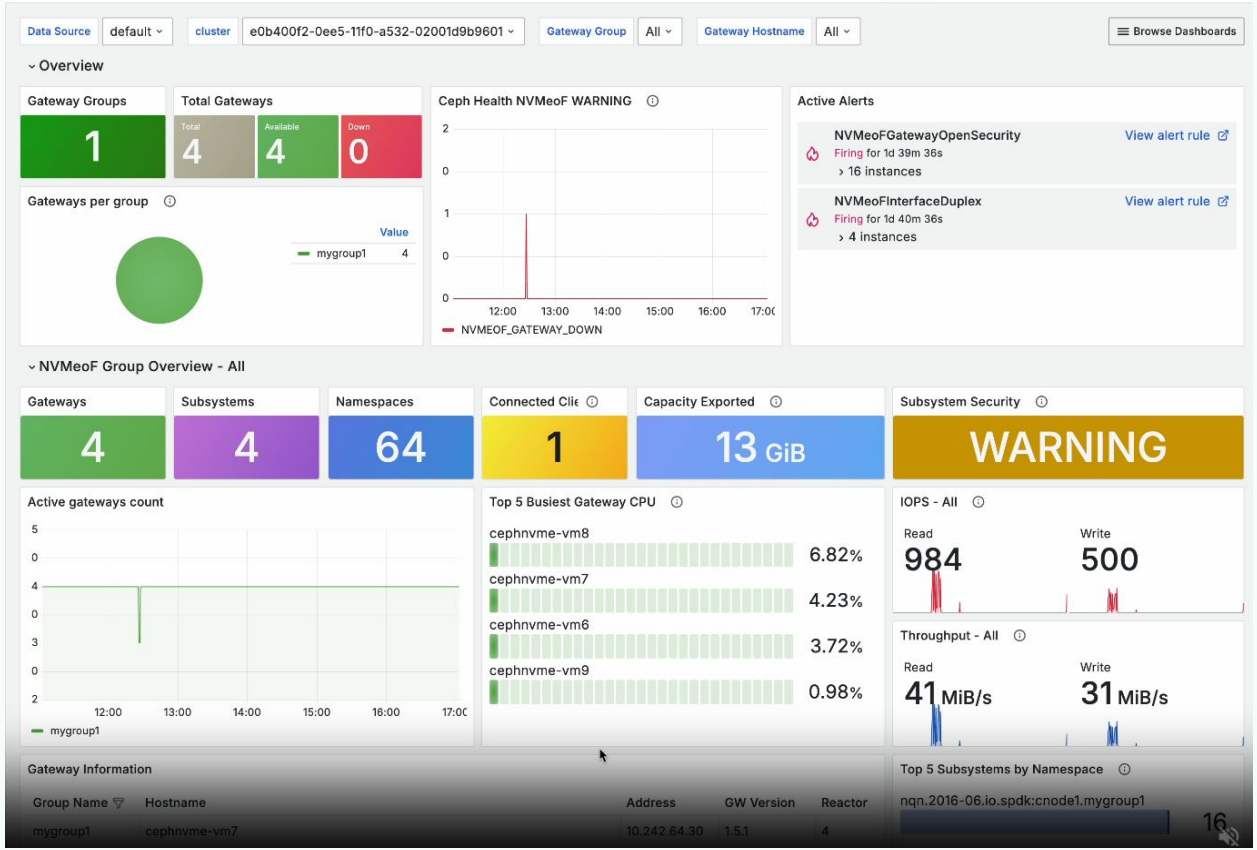

NVMe Performance

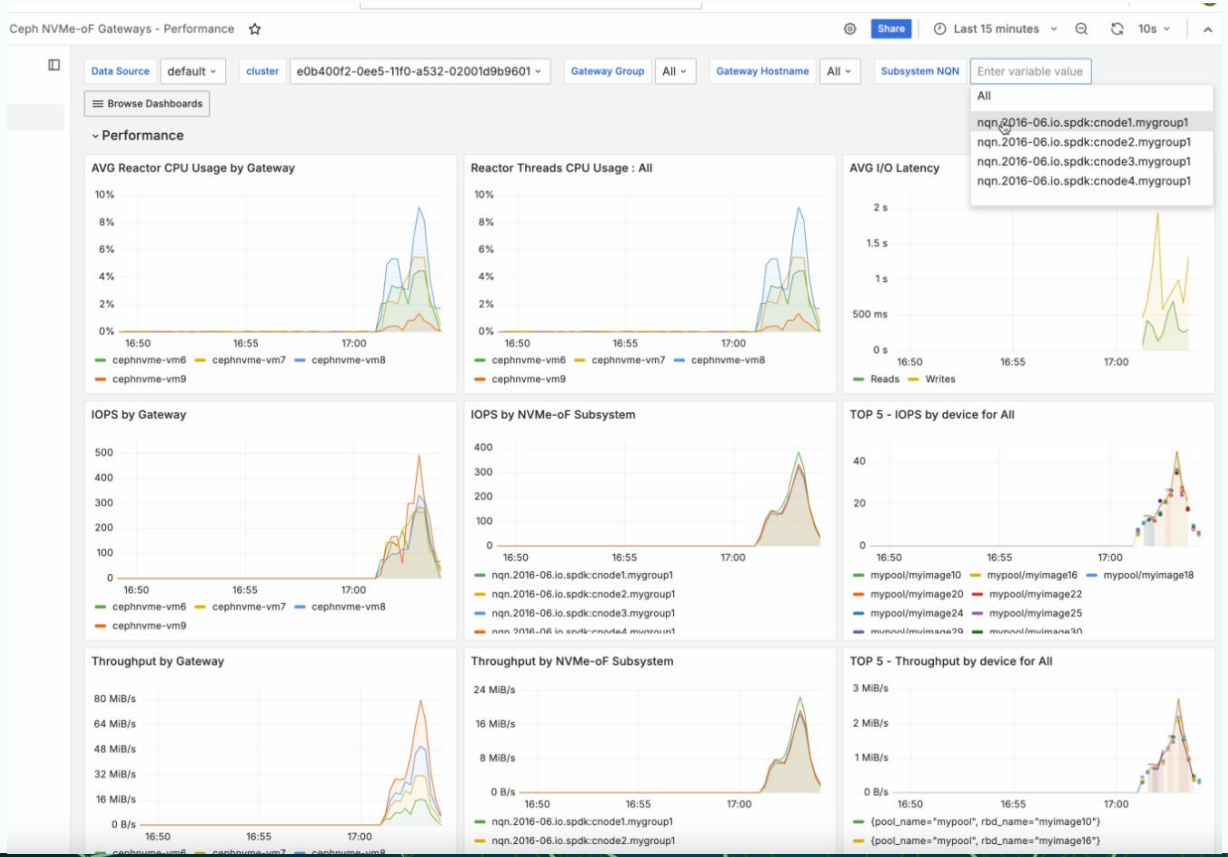

Ceph Filesystem Overview

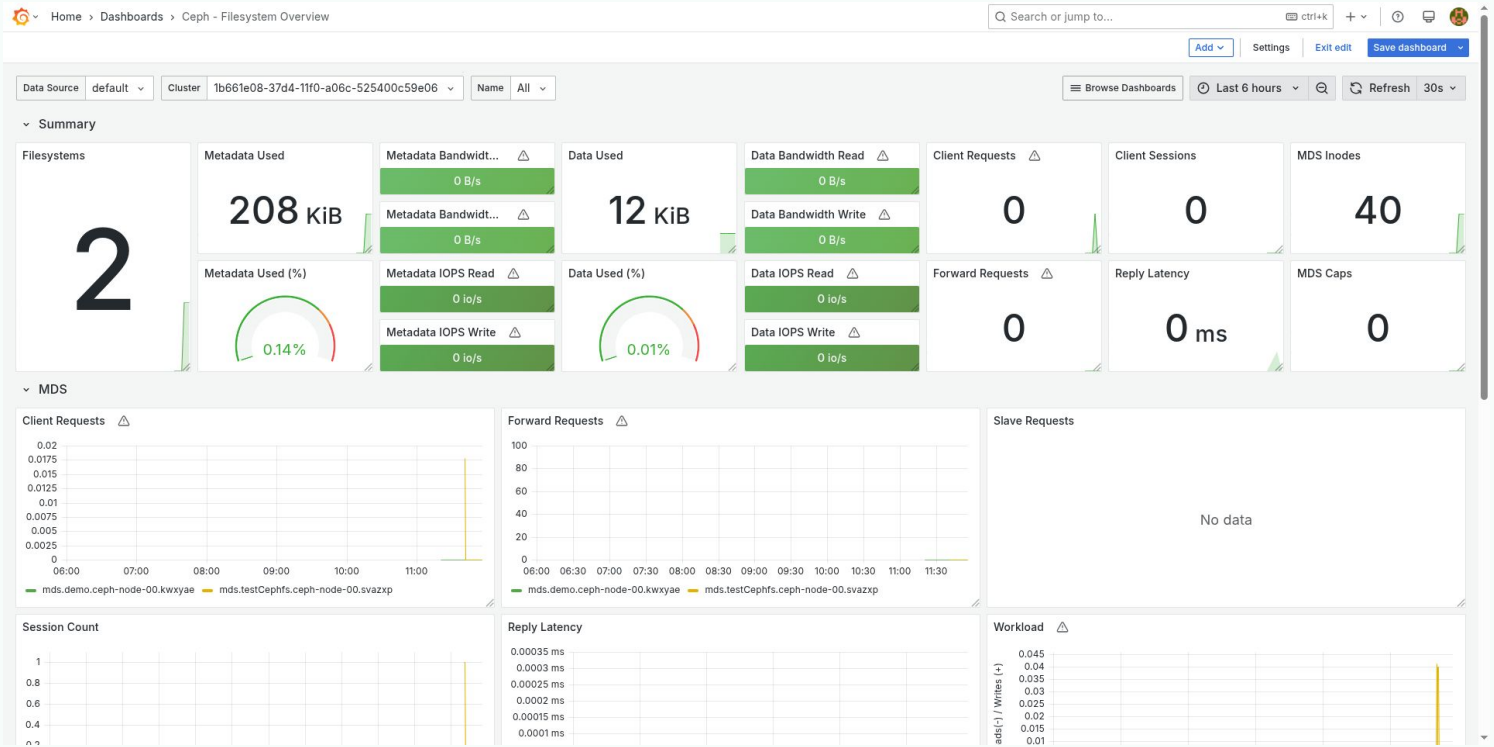

ceph

Future

>=20.x

Futures

Focus Area

- **Usability**
- **Overview**
- **Troubleshoot**
- **Scalability**
- **Monitoring & Alerting**

Usability

- Nvmeof workflow
- Multi- tenancy
- S3 Storage Browser
- Cluster level Visualization. Similar to https://tracker.ceph.com/issues/50980
- Install workflow - expand to be improved more
- Logging/progress

Overview

- What is using my data ? What application is using my data ?
- Users need fine tuning as well - filters
- Overview from all the perspectives - snapshot of everything, pool based, and specific resource
- Details, panels and graphs

Troubleshoot

- Notification management: Improvements in notifications area
- Same notifications keep coming again and again
- Need to silence, unsubscribe
- Have a dedicated page for notifications and tasks
- Keep track on ongoing/historic tasks
- A Centralized API for notification management.

Futures

Scalability

- Scalability at 10k and beyond - API level changes ??
- We can prioritize those endpoints and fix them up to stabilize for now... That should fix some of the issues. By the order of priority, it'd be
    1. Dashboard page
    2. Hosts page
    3. Alerts page
- https://tracker.ceph.com/issues/68418
- Making dashboard part of the core
    - Dashboard and API's outside the ceph container
    - Growing concern on the addition of more and more modules causing load on the host and mgr itself. (volumes, balancer etc)
- Host page
    - Even with server side pagination it breaks now for 000's of OSDs
    - Skeleton states (bug)Dashboard

Monitoring

- Cleanup for alerts
- RGW metrics
- Cephfs metrics
- Customizing Alerts from UI
- Runbooks for alerts
- Loki , promtail also occupy lots of space - needs design - being discussed.
- Log data retention
- Better Alert grouping

How to contribute

- Virtual **meetings**:
  - Ceph **Developer Monthly** (Wednesday)
  - Ceph **Users + Devs** Monthly
- As a **user**:
  - Getting help:
    - ceph-users@lists.ceph.com,
    - Slack #ceph-dashboard
  - Reporting issues or requesting features: tracker.ceph.com
- As a **developer**:
  - dev@ceph.io
  - github.com/ceph


ceph
Questions?


ceph
Thank you!