

Ceph Dashboard

Tableau de bord du stockage Ceph à des fins de surveillance et de gestion.

Description

Le tableau de bord est accessible à partir d'un navigateur Web et permet de configurer un cluster, ainsi que de visualiser les informations et les statistiques de performance. Le tableau de bord utilise un serveur web hébergé par le démon ceph-mgr.

Initialement développé par openATTIC, le Dashboard a été redéveloppé à partir de zéro avec une architecture plus flexible et extensible et a été intégré dans le projet Ceph à partir de la version Luminous

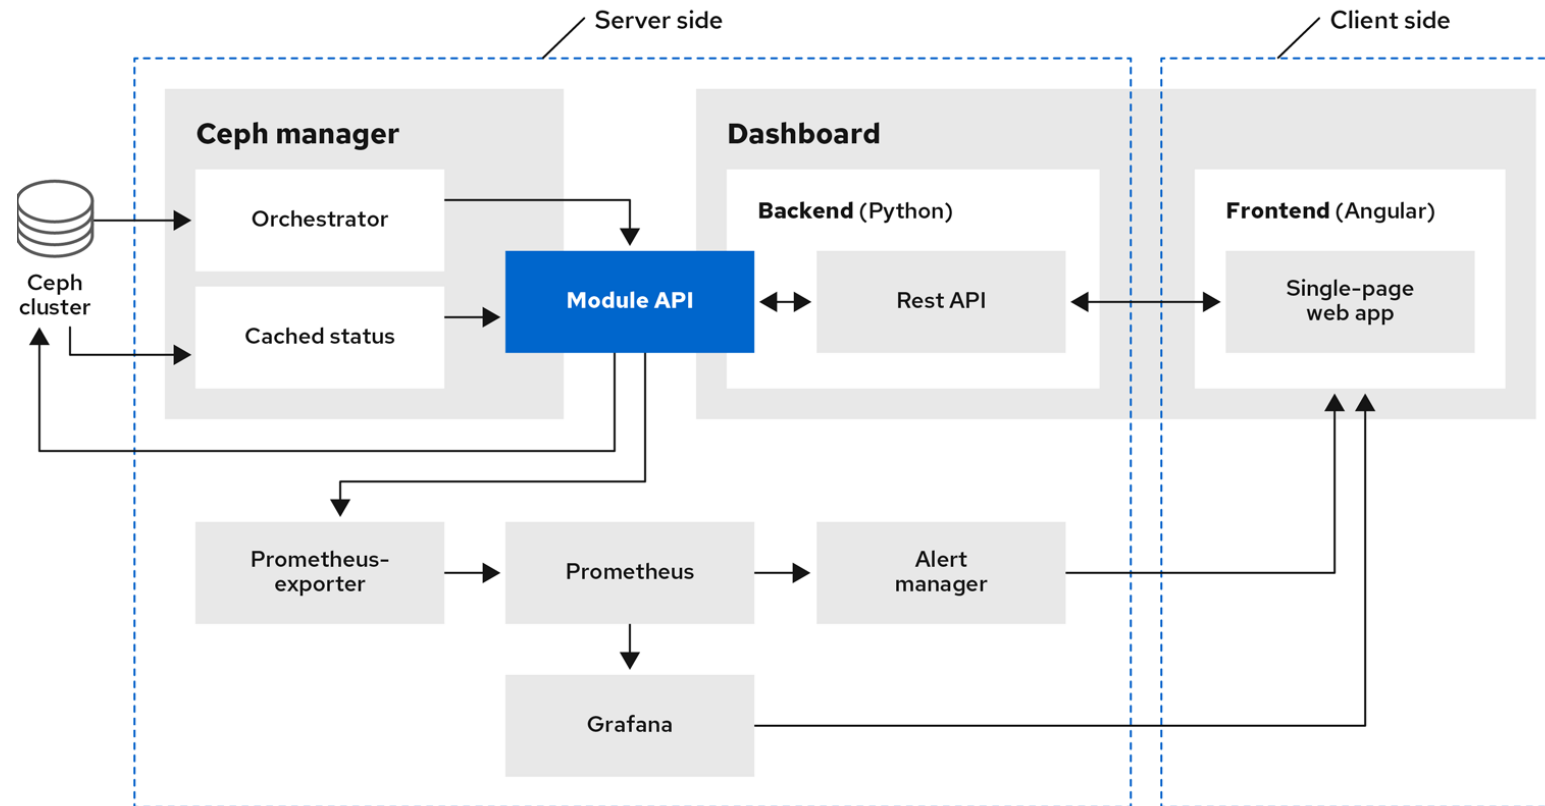
Le tableau de bord Ceph a gagné beaucoup de nouvelles fonctionnalités :

- Prise en charge de plusieurs utilisateurs / rôles
- SSO (SAMLv2) pour l'authentification des utilisateurs
- Nouvelle page d'accueil, affichant plus de mesures et d'informations sur l'état du cluster

Les nouvelles fonctionnalités de gestion

- Gestion de l'OSD (ajout, suppression, modifier les paramètres OSD, les profils de récupération)
- Éditeur de paramètres de configuration de cluster
- Gestion Ceph Pool (créer, modifier, supprimer)
- Configuration RBD (créer, modifier, supprimer, snapshot, QoS)
- Configuration de mise en miroir RBD
- Configuration CephFS et RadowsGW
- Tableaux de bord Grafana intégrés (dérivés de Ceph Metrics)
- Visionneuse de carte CRUSH
- Gestion NFS Ganesha
- Gestion des cibles iSCSI (via Ceph iSCSI Gateway)
- Gestion du module Ceph Manager (ceph-mgr)
- Gestion des alertes Prometheus

Dashboard architecture



https://access.redhat.com/documentation/en-us/red_hat_ceph_storage/5/html/dashboard_guide/ceph-dashboard-overview

Déterminer l'url du dashboard

```
[ceph: root@ceph1 /]# ceph mgr services
```

```
{  
  "dashboard": "https://172.16.7.242:8443/",  
  "prometheus": "http://172.16.7.242:9283/"  
}
```

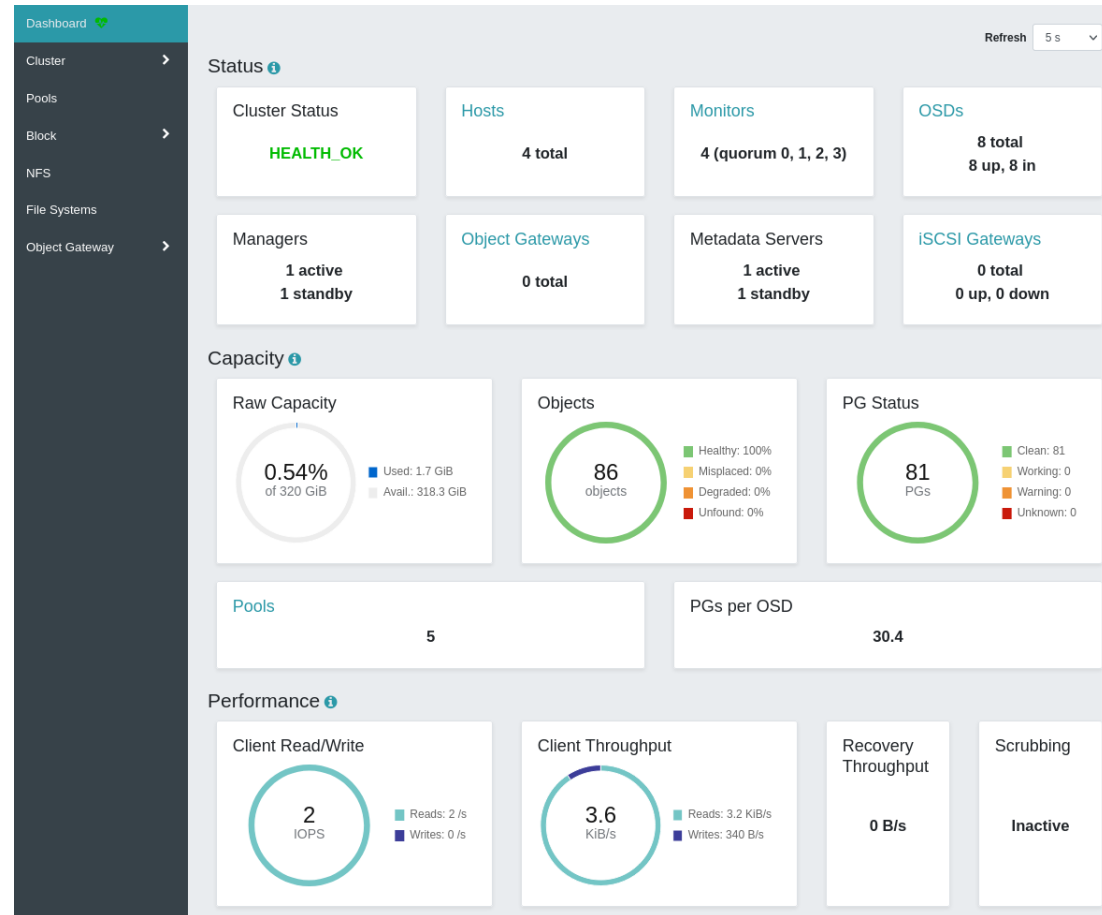
```
[ceph: root@ceph1 /]# ceph orch ls mgr
```

```
NAME PORTS RUNNING REFRESHED AGE PLACEMENT  
mgr      2/2 63s ago 12d count:2
```

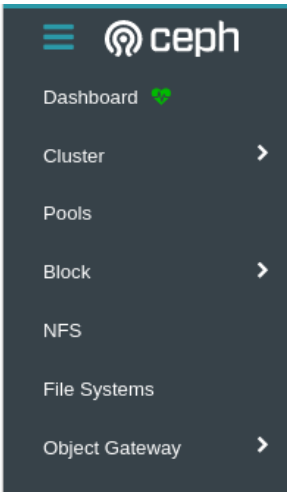
Page d'accueil

Informations sur :

- L'état du cluster et les différents services
- La capacité
- Les performances



Menu



- Dashboard: information sur l'état du cluster et retour page d'accueil
- Cluster : informations détaillées sur les hôtes, l'inventaire, les instances Ceph
- Monitor, les services, les OSD Ceph, la configuration du cluster, la carte CRUSH, les modules Ceph Manager, les journaux et la surveillance.
- Pools : Information et gestions des pools.
- Bloc : Information et gestion des images RBD, la mise en miroir, et iSCSI
- NFS : information et gestion des déploiement NFS Ganasha
- FileSystemes : Affichage et gestion des CephFS, des clients, des quotas, des snapshots
- Object Gateway : Affichage et gestion des RadosGw, des utilisateurs S3, des options par Buckets

Assistant de configuration

- Lors de la première connexion, un assistant de configuration vous accompagne pour
 - l'ajout des nodes supplémentaires
 - l'ajout d'osds supplémentaires
 - l'ajout de services supplémentaires
- Il est possible de retrouver ces assistants via le menu "Cluster"

Expand Cluster - add Hosts

The screenshot shows the Ceph dashboard interface for expanding a cluster. The left sidebar contains navigation options: Dashboard, Cluster, Pools, Block, NFS, File Systems, and Object Gateway. The main content area is titled 'Expand Cluster' and includes a progress indicator with four steps: 1. Add Hosts, 2. Create OSDs, 3. Create Services, and 4. Review. The 'Add Hosts' step is active, displaying a table of hosts to be added.

Hostname	Service Instances	Labels	Status	Model	CPUs	Cores	Total Memory	Raw Capacity	HDs	Flash	NICs
cna1	mds: 1 mgr: 1 mon: 1 osd: 2	admin		Virtual Machine (OpenStack Compute)	2	1	3.8 GiB	100 GiB	3	0	1
cna2	mds: 1 mgr: 1 mon: 1 osd: 2			Virtual Machine (OpenStack Compute)	2	1	3.8 GiB	100 GiB	3	0	1
cna3	cephfs-mirror: 1 mon: 1 osd: 2			Virtual Machine (OpenStack Compute)	2	1	3.8 GiB	100 GiB	3	0	1
cna4	mon: 1 osd: 2			Virtual Machine (OpenStack Compute)	2	1	3.8 GiB	100 GiB	3	0	1

0 selected / 4 total

Buttons: Cancel, Next

Expand Cluster - Create OSDs

The screenshot displays the Ceph web interface for expanding a cluster. The left sidebar contains navigation options: Dashboard, Cluster, Pools, Block, NFS, File Systems, and Object Gateway. The main content area is titled 'Expand Cluster' and shows a progress indicator with four steps: 1. Add Hosts, 2. Create OSDs (current step), 3. Create Services, and 4. Review. The 'Create OSDs' section includes three categories of devices, each with a '+ Add' button: Primary devices, Shared devices, and WAL devices. Below these, there are sections for 'DB devices' and 'Configuration'. The 'Configuration' section has a 'Features' subsection with an unchecked checkbox for 'Encryption'. At the bottom right of the main content area, there are 'Back' and 'Next' buttons.

Expand Cluster – create Services

The screenshot shows the Ceph dashboard interface. On the left is a dark sidebar with navigation options: Dashboard, Cluster, Pools, Block, NFS, File Systems, and Object Gateway. The main content area is titled 'Expand Cluster' and contains a progress indicator with four steps: 1. Add Hosts, 2. Create OSDs, 3. Create Services (highlighted in blue), and 4. Review. The 'Create Services' step is active, displaying a table of services to be created. The table has columns for 'Service' and 'Placement'. Above the table are controls for creating services, including a '+ Create' button, a refresh icon, a grid icon, a page size dropdown set to 10, and a search input field. At the bottom right of the main content area are 'Back' and 'Next' buttons.

Service	Placement
alertmanager	count:1
cephfs-mirror	cna3
grafana	count:1
mds.cephfs	count:2
node-exporter	*
osd.all-available-devices	*
prometheus	count:1

0 selected / 7 total

Dashboard - Settings

- Menu settings est accessible en haut à droite.
- Gestion des utilisateurs :
 - Permet de définir différents comptes d'accès au Dashboard.
 - Chaque compte peut avoir différents rôles : admin, read-only, block, cephfs, radosgw, ...
 - La politique des mots de passe est disponible via la cli
- Télémétrie configuration
 - Le module de télémétrie renvoie des données anonymes sur le cluster aux développeurs de Ceph pour aider à comprendre comment Ceph est utilisé et quels problèmes les utilisateurs peuvent rencontrer.
 - Suivi des évolutions : <https://telemetry-public.ceph.com/>
 - Il est conseillé d'activer la télémétrie

Cluster - Hosts

- Liste des nodes et des services
- Démarrage, arrêt, redéploiement d'un service
- Mode maintenance
- Performance des différents nodes
- Inventaire des périphériques
- Performance d'un node

Cluster » Hosts

Hosts List Overall Performance

Hostname	Service Instances	Labels	Status	Model	CPUs	Core	Total Memory	Raw Capacity	HDDs	Flash	NICs
cnb1	mds: 1 mgr: 1 mon: 1 osd: 2	admin		Virtual Machine (OpenStack Compute)	2	1	3.8 GiB	100 GiB	3	0	1
cnb2	mds: 1 mgr: 1 mon: 1 osd: 2			Virtual Machine (OpenStack Compute)	2	1	3.8 GiB	100 GiB	3	0	1
cnb3	mon: 1 osd: 2 rbd-mirror: 1			Virtual Machine (OpenStack Compute)	2	1	3.8 GiB	100 GiB	3	0	1
cnb4	mon: 1 osd: 2		maintenance	Virtual Machine (OpenStack Compute)	2	1	3.8 GiB	100 GiB	3	0	1

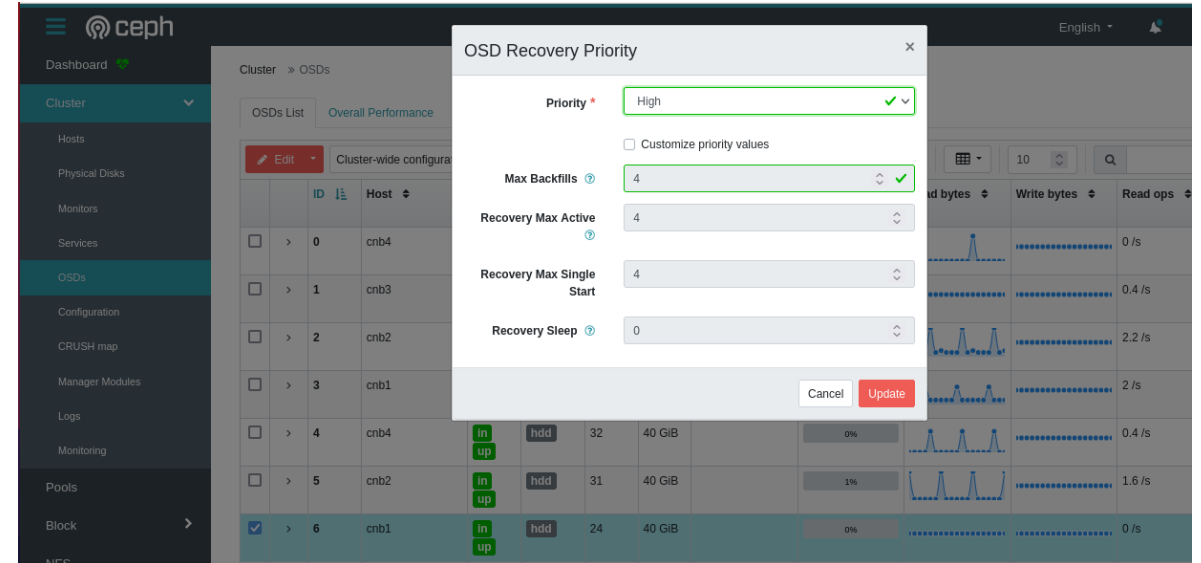
Devices Physical Disks Daemons Performance Details Device health

Daemon name	Version	Status	Last Refreshed	CPU Usage	Memory Usage	Daemon Events
crash.cnb4	17.2.1	stopped	33 minutes ago	5%	8.6 MiB	No data available
mon.cnb4	17.2.1	stopped	33 minutes ago	0%	510.5 MiB	No data available
node-exporter.cnb4		stopped	33 minutes ago	7%	23.9 MiB	No data available
osd.0	17.2.1	stopped	33 minutes ago	9%	480.9 MiB	No data available
osd.4	17.2.1	stopped	33 minutes ago	9%	534.2 MiB	No data available

0 selected / 5 total

Cluster - OSDs

- Gestion des OSDs : ajout, modification, suppression.
- OSD Flags: Gestion de « flags » noout, nodown, noscrub
- OSD Recovery priority : Permet d'activer un profil prédéfini ou de définir les priorités de récupération. un niveau bas ne va pas impacter la performance du cluster, mais le temps de récupération sera plus long.
- PG scrub : Définit les périodes où sont permis le scrub ou le deep scrub (heure et jour de la semaine)



Cluster - configuration

Permet de rechercher les options qui ont été modifiées.

Chaque option est accompagnée d'une explication et des valeurs supportées

Vous pouvez filtrer une configuration spécifique à l'aide des filtres suivants :

- Niveau : Basique, avancé ou dev
- Service : Tout, mon, mgr, osd, mds, common, mds_client, rgw
- Modifié : oui ou non



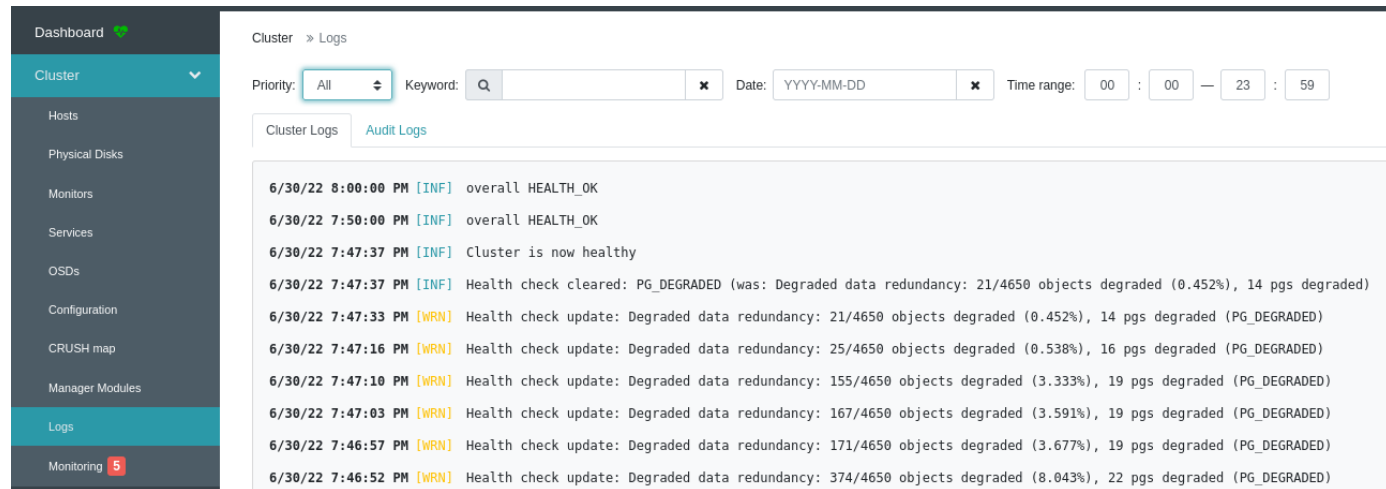
The screenshot shows the Ceph configuration interface. The left sidebar contains navigation options: Dashboard, Cluster (selected), Hosts, Physical Disks, Monitors, Services, OSDs, and Configuration. The main content area is titled 'Cluster > Configuration'. At the top, there is an 'Edit' button, a refresh icon, a table view icon, a page size dropdown set to '10', a search bar containing 'scrub_during', and a filter dropdown set to 'Level: advanced'. Below this is a table with the following columns: Name, Description, Current value, Default, and Editable. One row is visible for 'osd_scrub_during_recovery' with a description 'Allow scrubbing when PGs on the OSD are undergoing recovery', a current value of 'osd: true', a default of 'false', and an 'Editable' checkbox that is checked. At the bottom of the table, it says '0 selected / 1 found / 1788 total'. The footer of the slide contains the date '03/10/2022', the event name 'ANF CEPH 2022', the speaker's name 'Sébastien Geiger', and the page number '15'.

Name	Description	Current value	Default	Editable
> osd_scrub_during_recovery	Allow scrubbing when PGs on the OSD are undergoing recovery	osd: true	false	<input checked="" type="checkbox"/>

Cluster - Logs

Vous pouvez afficher et filtrer les journaux du cluster Ceph en fonction de plusieurs critères. Les critères incluent la priorité, le mot-clé, la date et l'intervalle de temps.

Remarques : l'affichage ne supporte que les 30 derniers événements de haut niveau. Les événements sont stockés en mémoire par le moniteur. Les entrées disparaissent après le redémarrage du moniteur. Si vous avez besoin d'examiner des journaux détaillés ou plus anciens, consultez les journaux basés sur les fichiers.



Cluster » Logs

Priority: All Keyword: Q Date: YYYY-MM-DD Time range: 00 : 00 — 23 : 59

Cluster Logs Audit Logs

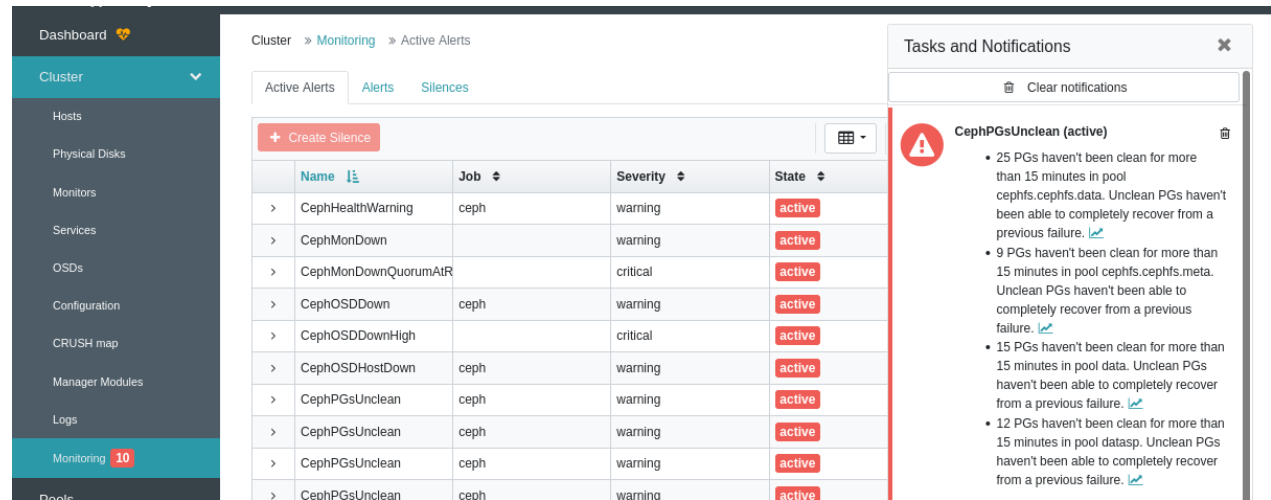
6/30/22 8:00:00 PM	[INF]	overall HEALTH_OK
6/30/22 7:50:00 PM	[INF]	overall HEALTH_OK
6/30/22 7:47:37 PM	[INF]	Cluster is now healthy
6/30/22 7:47:37 PM	[INF]	Health check cleared: PG_DEGRADED (was: Degraded data redundancy: 21/4650 objects degraded (0.452%), 14 pgs degraded)
6/30/22 7:47:33 PM	[WRN]	Health check update: Degraded data redundancy: 21/4650 objects degraded (0.452%), 14 pgs degraded (PG_DEGRADED)
6/30/22 7:47:16 PM	[WRN]	Health check update: Degraded data redundancy: 25/4650 objects degraded (0.538%), 16 pgs degraded (PG_DEGRADED)
6/30/22 7:47:10 PM	[WRN]	Health check update: Degraded data redundancy: 155/4650 objects degraded (3.333%), 19 pgs degraded (PG_DEGRADED)
6/30/22 7:47:03 PM	[WRN]	Health check update: Degraded data redundancy: 167/4650 objects degraded (3.591%), 19 pgs degraded (PG_DEGRADED)
6/30/22 7:46:57 PM	[WRN]	Health check update: Degraded data redundancy: 171/4650 objects degraded (3.677%), 19 pgs degraded (PG_DEGRADED)
6/30/22 7:46:52 PM	[WRN]	Health check update: Degraded data redundancy: 374/4650 objects degraded (8.043%), 22 pgs degraded (PG_DEGRADED)

Cluster – Monitoring – Active Alerts

Affichage des règles prometheus et de leurs explications

Vous pouvez afficher les différentes notifications depuis le menu cloche en haut à droite.

Vous pouvez également surveiller les alertes à l'aide du protocole SNMP (Simple Network Management Protocol).



The screenshot displays the Ceph Monitoring interface. On the left is a navigation sidebar with options like Dashboard, Cluster, Hosts, Physical Disks, Monitors, Services, OSDs, Configuration, CRUSH map, Manager Modules, Logs, and Pools. The main area is titled 'Cluster > Monitoring > Active Alerts' and contains a table of active alerts. A 'Tasks and Notifications' panel on the right shows a notification for 'CephPGsUnclean (active)' with details about pool recovery issues.

Name	Job	Severity	State
CephHealthWarning	ceph	warning	active
CephMonDown		warning	active
CephMonDownQuorumAIR		critical	active
CephOSDDown	ceph	warning	active
CephOSDDownHigh		critical	active
CephOSDHostDown	ceph	warning	active
CephPGsUnclean	ceph	warning	active
CephPGsUnclean	ceph	warning	active
CephPGsUnclean	ceph	warning	active
CephPGsUnclean	ceph	warning	active

Tasks and Notifications

CephPGsUnclean (active)

- 25 PGs haven't been clean for more than 15 minutes in pool cephfs.cephfs.data. Unclean PGs haven't been able to completely recover from a previous failure.
- 9 PGs haven't been clean for more than 15 minutes in pool cephfs.cephfs.meta. Unclean PGs haven't been able to completely recover from a previous failure.
- 15 PGs haven't been clean for more than 15 minutes in pool data. Unclean PGs haven't been able to completely recover from a previous failure.
- 12 PGs haven't been clean for more than 15 minutes in pool datap. Unclean PGs haven't been able to completely recover from a previous failure.

Cluster - Block

- Création , modification, suppression des images RBD
- Gestion namespace, snapshot, trash
- Configuration, modification de la géo réplication
- Configuration des iSCSI

Block » Mirroring

Site Name: silea Edit Site Name

Daemons

Instance	ID	Hostname	Version	Health
34615	cna3.ohukie	cna3	17.2.1	OK
1 total				

Pools

Name	Mode	Leader	# Local	# Remote	Health
data	pool	34615	2	2	OK
datasp	image	34615	2	2	OK
1 selected / 2 total					

Images

Issues Syncing Ready

Pool	Image	Description	State
data	imagea1	local image is primary	Primary
data	imagea2	local image is primary	Primary
datasp	imagesp1	local image is primary	Primary
datasp	imagesp2	replaying, ["bytes_per_second":0.0,"bytes_per_snapshot":0.0,"local_snapshot_timestamp":165	Replaying
4 total			

Cluster - FileSystem

- Configuration des quotas et des snapshots
- Visualisation de la configuration CephFS, des data pools, des clients, des répertoires, et de la performance
- La création des CephFS se fait toujours en cli.



File Systems

Name	Created	Enabled
moncfs	8/18/22 7:43:32 PM	✓

Details Clients 2 Directories Performance Details

▼ /

- ▶ rep1
- ▼ rep2
- nfs2
- repec
- ▶ volumes

/rep2

Quotas

Set

Name	Value	Origin
Max files		
Max size	3.7 GiB	/rep2

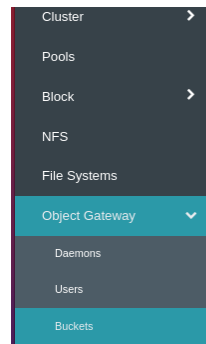
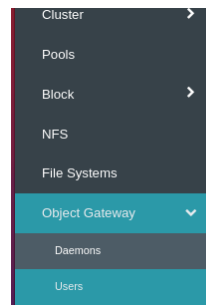
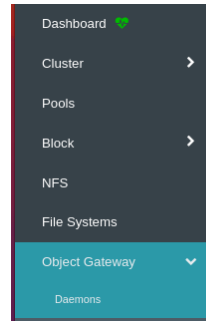
Snapshots

Create

Name	Created
<input type="checkbox"/> scheduled-2022-08-21-08_00_00	8/21/22 8:00:01 AM
<input type="checkbox"/> scheduled-2022-08-21-09_00_00	8/21/22 9:00:00 AM
<input type="checkbox"/> scheduled-2022-08-21-10_00_00	8/21/22 10:00:00 AM
<input type="checkbox"/> scheduled-2022-08-21-11_00_00	8/21/22 11:00:00 AM

Cluster - RadosGW

- Visualisation des services rgw par zone, de leur configuration et des performances
- Configuration des utilisateurs et des buckets



Selected Object Gateway: euzone.ceph3.umdnoa (euzone)

Object Gateway » Daemons

Daemons List Overall Performance Sync Performance

ID	Hostname	Zone	Zone Group	Realm	Version
> euzone.ceph3.umdnoa	ceph3	fr-east-1	euzone	demodom	16.2.10
> euzone.ceph4.gdkeiz	ceph4	fr-east-1	euzone	demodom	16.2.10

2 total

Object Gateway » Users

Username	Tenant	Full name	Email address	Suspended	Max. buckets	Capacity Limit %	Object Limit %
> dashboard		Ceph Dashboard			1000	No Limit	No Limit
> johndoe		John Doe			1000	No Limit	No Limit
> zone.user		Zone user			1000	No Limit	No Limit

0 selected / 3 total

Object Gateway » Buckets

Name	Owner	Used Capacity	Capacity Limit %	Objects	Object Limit %
> bquota	johndoe	500 MiB	47%	1	No Limit
> bworm3	johndoe	8 KiB	No Limit	2	No Limit
> mybucket	johndoe	8 KiB	No Limit	2	No Limit

0 selected / 3 total