

📢 Notice

This document is for a development version of Ceph.

[Report a Documentation Bug](#)

CEPH STORAGE CLUSTER

The [Ceph Storage Cluster](#) is the foundation for all Ceph deployments. Based upon [RADOS](#), Ceph Storage Clusters consist of several types of daemons:

1. a [Ceph OSD Daemon](#) (OSD) stores data as objects on a storage node
2. a [Ceph Monitor](#) (MON) maintains a master copy of the cluster map
3. a [Ceph Manager](#) daemon

A Ceph Storage Cluster might contain thousands of storage nodes. A minimal system has at least one Ceph Monitor and two Ceph OSD Daemons for data replication.

The Ceph File System, Ceph Object Storage and Ceph Block Devices read data from and write data to the Ceph Storage Cluster.

CONFIG AND DEPLOY

Ceph Storage Clusters have a few required settings, but most configuration settings have default values. A typical deployment uses a deployment tool to define a cluster and bootstrap a monitor. See [Cephadm](#) for details.

- [Configuration](#)
 - [Storage devices](#)

OPERATIONS

Once you have deployed a Ceph Storage Cluster, you may begin operating your cluster.

- [Operations](#)
 - [Operating a Cluster](#)
 - [Health Checks](#)
 - [Monitoring a Cluster](#)
 - [Monitoring OSDs and PGs](#)
 - [User Management](#)
 - [PG Calc](#)

APIS

Most Ceph deployments use [Ceph Block Devices](#), [Ceph Object Storage](#) and/or the [Ceph File System](#). You may also develop applications that talk directly to the Ceph Storage Cluster.

- [APIs](#)
 - [Introduction to librados](#)
 - [librados](#)
 - [librados \(C++\)](#)
 - [librados \(Python\)](#)

- BlueFS Spillover Cleaner
- Fast Crash Recovery for file-stored allocations
- ExtBlkDev, FCM plugin
- Configuring Ceph
- Common Settings
- Networks
- Temporary Directory
- Monitors
- Authentication
- OSDs
- Heartbeats
- Logs / Debugging
- Example ceph.conf
- Naming Clusters (deprecated)
- Network Settings
- Messenger v2 protocol
- Auth Settings
- Monitor Settings
- Looking up Monitors through DNS
- Heartbeat Settings
- OSD Settings
- DmClock Settings
- BlueStore Settings
- FileStore Settings
- Journal Settings
- Pool, PG & CRUSH Settings
- General Settings

- Pools
- Erasure code
- Cache Tiering
- Placement Groups
- Placement Group States
- Placement Group Concepts
- Using pg-upmap
- Operating the Read (Primary) Balancer
- Balancer Module
- Minimum Compatible Client
- Minimum OSD Release
- CRUSH Maps
- Manually editing the CRUSH Map
- Stretch Clusters
- Configuring Monitor Election Strategies
- Adding/Removing OSDs
- Adding/Removing Monitors
- Device Management
- BlueStore Migration
- Command Reference
- The Ceph Community
- Troubleshooting Monitors
- Troubleshooting OSDs
- Troubleshooting PGs
- Logging and Debugging
- CPU Profiling
- Memory Profiling

The Ceph Documentation is a community resource funded and hosted by the non-profit [Ceph Foundation](#). If you would like to support this and our other efforts, please consider [joining now](#).