

This is documentation for Apache Ozone **2.1.1**, which is no longer actively maintained.

For up-to-date documentation, see the [latest version](#) (2.2.0).

[Home](#) > [Core Concepts](#) > [High Availability](#) > [Ozone Manager High Availability](#)

Version: 2.1.1

Ozone Manager High Availability

Ozone has two metadata-manager nodes (*Ozone Manager* for key space management and *Storage Container Manager* for block space management) and multiple storage nodes (Datanode). Data is replicated between Datanodes with the help of RAFT consensus algorithm.


To avoid any single point of failure the metadata-manager nodes also should have a HA setup.

Both Ozone Manager and Storage Container Manager supports HA. In this mode the internal state is replicated via RAFT (with Apache Ratis).

This page explains the HA setup of Ozone Manager (OM). Please check the [SCM High Availability](#) page for SCM HA. While they can be setup for HA independently, a reliable, full HA setup requires enabling HA for both services.

How It Works

A single Ozone Manager uses [RocksDB](#) to persist metadata (volumes, buckets, keys) locally. HA version of Ozone Manager does exactly the same but all the data is replicated with the help of the RAFT consensus algorithm to follower Ozone Manager instances.



Client connects to the Leader Ozone Manager which processes the request and schedules the replication with RAFT. When the request is replicated to all the followers the leader can return with the response.

Learn More

- **Configuration:** For detailed configuration instructions on setting up OM HA, see the [OM HA Configuration](#) documentation.
- **Implementation Details:** For in-depth technical details about the OM HA implementation, including the double buffer approach and automatic snapshot installation