

Try out Ozone

- › 1 [Build from Source](#)
 - 1.1 [Build From Git Repo](#)
 - 1.2 [Build From a Source Release](#)
- › 2 [Download Binary Release](#)
- › 3 [Start Cluster Using Docker](#)
- › 4 [Single Node Development Cluster](#)
 - 4.1 [Configuration](#)
 - 4.2 [Start Services](#)
 - 4.3 [Run Ozone Commands](#)
 - 4.4 [Stop Services](#)
 - 4.5 [Clean up your Dev Environment \(Optional\)](#)
- › 5 [Multi-Node Ozone Cluster](#)
 - 5.1 [Pre-requisites](#)
 - 5.2 [Configuration](#)
 - 5.3 [Start Services](#)
 - 5.4 [Stop Services](#)

There are two ways to try out Ozone. Either you can build from source code or download a binary release.

Build from Source

Build From Git Repo

Get the Apache Ozone source code from the Apache Git repository. Then check out trunk and build it with the `hdds` Maven profile enabled.

```
git clone https://github.com/apache/ozone.git
cd hadoop-ozone
mvn clean install -DskipTests=true -Dmaven.javadoc.skip=true -Pdist -Dtar -DskipShade
```

Initial compilation may take over 30 minutes as Maven downloads dependencies. `-DskipShade` is optional - it makes compilation faster for development.

This will give you a tarball in your distribution directory. Here is an example of the tarball that will be generated.

`hadoop-ozone/dist/target/ozone-0.5.0-SNAPSHOT.tar.gz`

Build From a Source Release

Download and extract a source tarball from <https://ozone.apache.org/downloads/> E.g.

```
tar xf ozone-1.2.1-src.tar.gz
cd ozone-1.2.1-src/
mvn clean install -DskipTests=true -Dmaven.javadoc.skip=true -Pdist -Dtar -DskipShade
```

Download Binary Release

Download and extract a binary release from <https://ozone.apache.org/downloads/> E.g.

```
tar xf ozone-1.2.1.tar.gz
cd ozone-1.2.1/
```

Start Cluster Using Docker

If you downloaded and built a source release, to start the docker cluster using the package you built, please add `"-Ddocker.ozone-runner.version=dev"` to the maven build command line. Run the following commands to start an Ozone cluster in docker containers with 3 datanodes.

```
cd hadoop-ozone/dist/target/ozone-*-SNAPSHOT/compose/ozone
docker-compose up -d --scale datanode=3
```

If you downloaded a binary release, run the following instead.,

```
cd compose/ozone
docker-compose up -d --scale datanode=3
```

For more docker-compose commands, please check the end of the [Getting started with docker](#) guide

To Shutdown the cluster, please run the command **docker-compose down**

Single Node Development Cluster

This is the traditional way to start a development cluster from source code. Once the package is built, you can start Ozone services by going to the *hadoop-ozone/dist/target/ozone-*/* directory. Your Unix shell should expand the "*" wildcard to the correct Ozone version number.

Configuration

Save the minimal snippet to *hadoop-ozone/dist/target/ozone-*/etc/hadoop/ozone-site.xml* in the compiled distribution.

```
<configuration>
<properties>
<property><name>ozone.enabled</name><value>true</value></property>
<property><name>ozone.scm.datanode.id</name><value>/tmp/ozone/data/datanode.id</value></property>
<property><name>ozone.replication</name><value>1</value></property>
<property><name>ozone.metadata.dirs</name><value>/tmp/ozone/data/metadata</value></property>
<property><name>ozone.scm.names</name><value>localhost</value></property>
<property><name>ozone.scm.client.address</name><value>localhost</value></property>
<property><name>ozone.scm.block.client.address</name><value>localhost</value></property>
<property><name>ozone.om.address</name><value>localhost</value></property>
</properties>
</configuration>
```

Start Services

To start ozone, you need to start SCM, OzoneManager and DataNode. In pseudo-cluster mode, all services will be started on localhost.

```
bin/ozone scm --init
bin/ozone --daemon start scm
bin/ozone om --init
bin/ozone --daemon start om
bin/ozone --daemon start datanode
```

Run Ozone Commands

Once you have ozone running you can use these [Ozone shell](#) commands to create a volume, bucket and keys. E.g.

```
bin/ozone sh volume create /vol1
bin/ozone sh bucket create /vol1/bucket1
dd if=/dev/zero of=/tmp/myfile bs=1024 count=1
bin/ozone sh key put /vol1/bucket1/key1 /tmp/myfile
bin/ozone sh key list /vol1/bucket1
```

Stop Services

```
bin/ozone --daemon stop om
bin/ozone --daemon stop scm
bin/ozone --daemon stop datanode
```

Clean up your Dev Environment (Optional)

Remove the following directories to wipe the Ozone pseudo-cluster state. This will also delete all user data (volumes/buckets/keys) you added to the pseudo-cluster.

```
rm -fr /tmp/ozone
rm -fr /tmp/hadoop-${USER}*
```



Note: This will also wipe state for any running HDFS services.

Multi-Node Ozone Cluster

Pre-requisites

Ensure you have password-less ssh setup between your hosts.

Configuration

ozone-site.xml

Save the following snippet to *etc/hadoop/ozone-site.xml* in the compiled Ozone distribution.

```
<configuration>
<properties>
<property><name>ozone.scm.block.client.address</name><value>SCM-HOSTNAME</value></property>
<property><name>ozone.scm.names</name><value>SCM-HOSTNAME</value></property>
<property><name>ozone.scm.client.address</name><value>SCM-HOSTNAME</value></property>
<property><name>ozone.om.address</name><value>OM-HOSTNAME</value></property>
<property><name>ozone.handler.type</name><value>distributed</value></property>
<property><name>ozone.enabled</name><value>True</value></property>
<property><name>ozone.scm.datanode.id</name><value>/tmp/ozone/data/datanode.id</value></property>
<property><name>ozone.replication</name><value>1</value></property>
<property><name>ozone.metadata.dirs</name><value>/tmp/ozone/data/metadata</value></property>
</properties>
</configuration>
```

Replace SCM-HOSTNAME and OM-HOSTNAME with the names of the machines where you want to start the SCM and OM services respectively. It is okay to start these services on the same host. If you are unsure then just use any machine from your cluster.

ozone-env.sh

The only mandatory setting in ozone-env.sh is JAVA_HOME. E.g.

```
# The java implementation to use. By default, this environment
# variable is REQUIRED on ALL platforms except OS X!
export JAVA_HOME=/usr/java/latest
```

workers

The workers file should contain a list of hostnames in your cluster where DataNode service will be started. E.g.

```
n001.example.com  
n002.example.com  
n003.example.com  
n004.example.com
```

Start Services

Initialize the SCM

Run the following commands on the SCM host

```
bin/ozone scm --init  
bin/ozone --daemon start scm
```

Format the OM

Run the following commands on the OM host

```
bin/ozone om --init  
bin/ozone --daemon start om
```

Start DataNodes

Run the following command on any cluster host.

```
sbin/start-ozone.sh
```

Stop Services

Run the following command on any cluster host.

```
sbin/stop-ozone.sh
```