

Cloudera Runtime 7.3.2

Ozone Performance Tuning

Date published: 2020-07-28

Date modified: 2026-03-31

CLOUDERA

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2026. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Configuring Ozone services.....	4
Calculating Ozone heap sizing.....	6
Performance tuning for Ozone.....	7
Tuning Ozone for faster deletion performance.....	8
Node maintenance.....	9

Configuring Ozone services

This section helps you to understand the purpose of the parameters you need to set and the disk type you require for services like DataNode, OM, Recon, and SCM.



Note:

- Cloudera recommends you to configure the correct values during the initial configuration of the Ozone service for the below properties. Changing this configuration later can cause service disruption.

For example, Cloudera does not recommend moving or changing the data directory once Ozone service is in use. Changing or moving the data directory later is disruptive and brings down the Ozone service.

- You must configure the Ozone Manager Metadata directory and Ozone Manager Data directory in the same filesystem. If you configure the Ozone Manager Metadata directory and Ozone Manager Data directory on different filesystems on the host server, the hard link will fail causing the snapshot installation job to fail.
- Multiple Ozone deployments are not supported in the Cloudera cluster. Ensure that you have only one Ozone deployment Service ID in the `ozone.om.service.ids` configuration of the Cloudera cluster. If you have multiple Ozone deployment Service IDs in the `ozone.om.service.ids` configuration, then remove all of them except one.

DataNode

Setting name	Default value	Purpose	Disktype	Mount point	Minimum size (GB)	Notes
<code>ozone.metadata.dirs</code>	<code>/var/lib/hadoop-ozone/datanode/ozone-metadata</code>	One or more directories used for storing Ozone metadata. OzoneManager, SCM, and Datanode will write the metadata to this path.	HDD	<code>/var/app/lib/</code>	100	
<code>dfs.container.ratis.data.dir</code>	<code>/var/app/lib/hadoop-ozone/datanode/ratis/data</code>	One or more directories used for storing Datanode Ratis metadata.	NVMe (Preferred)	<code>/var/app/lib/hadoop-ozone/datanode/ratis</code>	1024	RAID 1 NVMe or SSDNVMe/SSD is not required but strongly recommended.
<code>hdds.datanode.dir</code>	None	List of directories where user data is stored on each DataNode.	HDD	<code>/data/{1..N}</code>	Size and Number varies	The mount points are numeric values for each data drive 1 - n These disks must not be shared with HDFS or another storage system. If not configured this falls back to <code>dfs.datanode.data.dir</code> . However, Cloudera recommends you to configure this explicitly.

Ozone Manager

Setting name	Default value	Purpose	Disktype	Mount point	Minimum size (GB)	Notes
ozone.metadata.dirs	/var/lib/hadoop-ozone/om/ozone-metadata	One or more directories used for storing Ozone metadata. OzoneManager, SCM, and Datanode will write the metadata to this path.	NVMe (required)	/var/app/lib/hadoop-ozone/om/	4096	RAID 1 NVMe required
ozone.om.ratis.storage.dir	/var/lib/hadoop-ozone/om/ratis	This directory stores Ozone Manager's Ratis metadata like logs.	NVMe (required)	/var/app/lib/hadoop-ozone/om/		
ozone.om.db.dirs	/var/lib/hadoop-ozone/om/data	Directory where the Ozone Manager stores its metadata.	NVMe (required)	/var/app/lib/hadoop-ozone/om/		
ozone.om.snapshot.diff.db.dir	/var/lib/hadoop-ozone/om/data	Directory for Ozone snapshot diff scratch space.	HDD or SSD	/data/1	256 (One HDD with 256GB free space will be sufficient to store snapshot diffs.)	This is not available in Cloudera Manager. You must configure this through the safety valve. By default, if no value is specified here then the same locations as specified in ozone.metadata.dirs will be used.

Recon

Setting name	Default value	Purpose	Disktype	Mount point	Minimum size (GB)	Notes
ozone.metadata.dirs	/var/lib/hadoop-ozone/recon/ozone-metadata	One or more directories used for storing Ozone metadata. OzoneManager, SCM, and Datanode will write the metadata to this path.	HDD	/var/app/lib/	4096	These 3 settings can share the same NVMe dedicated to Recon
ozone.recon.db.dir	/var/lib/hadoop-ozone/recon/data	Directory where the Recon Server stores its metadata.	NVMe (required)	/var/app/lib/hadoop-ozone/recon		
ozone.recon.om.db.dir	/var/lib/hadoop-ozone/recon/om/data	Where recon keeps OM snapshot DB.	NVMe (required)	/var/app/lib/hadoop-ozone/recon		
ozone.recon.scm.db.dir	/var/lib/hadoop-ozone/recon/scm/data	Directory where the Recon Server stores StorageContainerManager's metadata	NVMe (required)	/var/app/lib/hadoop-ozone/recon		

Storage Container Manager

Setting name	Default value	Purpose	Disktype	Mount point	Minimum size (GB)	Notes
ozone.metadata.dirs	/var/lib/hadoop-ozone/scm/ozone-metadata	One or more directories used for storing Ozone metadata. OzoneManager, SCM, and Datanode will write the metadata to this path.	NVMe or SSD	/var/app/lib/hadoop-ozone/scm	4096	These 3 settings can share the same NVMe dedicated to the SCM. RAID 1 NVMe or SSD required.
ozone.scm.db.dirs	/var/lib/hadoop-ozone/scm/data	Directory where the Storage Container Manager stores its metadata	NVMe or SSD	/var/app/lib/hadoop-ozone/scm		
ozone.scm.ha.ratis.storage.dir	/var/lib/hadoop-ozone/scm/ratis	Storage directory used by SCM to write Ratis logs.	NVMe or SSD	/var/app/lib/hadoop-ozone/scm		



Note: If NVMe is not available, then a SAS SSD can be used instead.

Calculating Ozone heap sizing

Learn about the heuristic approach to configuring heap memory for Apache Ozone.

Unlike HDFS, which scales heap linearly with the number of blocks or files, Ozone utilizes RocksDB for metadata persistence and a container-based replication mechanism through the Storage Container Manager (SCM).

This architectural shift moves the bottleneck from raw data volume to workflow intensity, allowing the cluster to scale beyond 1 billion objects. The following guidelines provide calculation formulas for the Ozone Manager (OM), SCM, and DataNodes, combining a static base heap with dynamic multipliers based on peak operational metrics.

- Ozone Manager Heap

```
OM_heap = 8 GB + (write_ops_peak × 2 KB) + (delete_ops_peak × 2 KB) + (read_ops_peak × 200 B) + 2GB listing overhead
```

Where,

- write_ops_peak = om_metrics_num_key_allocate parameter value in the Prometheus metrics
- delete_ops_peak = om_metrics_num_key_deletes parameter value in the Prometheus metrics
- read_ops_peak = om_metrics_num_key_lookup parameter value in the Prometheus metrics
- Storage Container Manager

```
Heap SCM_heap = 6 GB + (number_of_containers × 1.4 KB) + (peak_delete_transactions × 500 B)
```

where,

- number_of_containers = Total number of containers in a cluster
- peak_delete_transactions = scm_block_deleting_service_num_processed_transactions parameter value in the Prometheus metrics

- DataNode Heap

```
DN_heap = 4 GB + (active_pipelines × 150 MB) + (containers_on_node × 200 B)
```

where,

- active_pipelines = Get the active_pipelines value using the `ozone admin pipeline list | grep --count $(hostname)` command
- containers_on_node = Container count from the Recon Web UI / Number of DataNodes in the cluster

Performance tuning for Ozone

Learn how to use the Ozone configuration properties to tune Ozone to work optimally on your cluster.

For performance optimization and to overcome known issues, Cloudera recommends the following additional configurations:

Configuration Property	Value	Remarks
Ozone Configuration		
Maximum Process File Descriptors	100,000	
Java Heap Size of Ozone Manager	31GB	
Java Heap Size of Storage Container Manager	31GB	
Java Heap Size of Recon	31GB	Can be increased to 64GB for higher load.
Java Heap Size of S3 Gateway	31GB	
Java Heap Size of DataNode	31GB	
Ozone Service Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml	<pre><property> <name>hdds.prometheus.endpoint.token</name> <value>*****</value> <description>Prometheus Token</description> </property></pre>	<p>Prometheus is an optional role and this configuration can be ignored if Prometheus is not being used.</p> <p>Disables Prometheus SPNEGO and uses Token Based Authentication.</p>
Ozone Configuration Ozone DataNode Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml		
ozone.container.cache.size	8192	
ozone.container.cache.lock.stripes	8192	
hdds.datanode.du.factory.classname	org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory	
dfs.container.ratis.leader.pending.bytes.limit	2GB	
Ozone Configuration Storage Container Manager Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml		
ozone.scm.datanode.pipeline.limit	10	If DataNodes have different number of disks, pick the highest number.
ozone.scm.ec.pipeline.per.volume.factor	4	Default value is 1.
ozone.scm.ec.pipeline.minimum	10	Default value is 5.
ozone.scm.pipeline.owner.container.count	10	
ozone.scm.pipeline.creation.auto.factor.one	FALSE	
ozone.scm.container.placement.impl	org.apache.hadoop.hdds.scm.container.placement.algorithms.SCMContainerPlacementCapacity	Flag on the SCM to optimize placement on capacity available instead of random allocation.

Configuration Property	Value	Remarks
Ozone Configuration Ozone Recon Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozzone-site.xml		
ozone.recon.task.pipelinesync.interval	120s	
ozone.recon.task.missingcontainer.interval	3600s	
Ozone Configuration Ozone Recon Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozzone-site.xml		
hdds.datanode.replication.work.dir	/tmp (is the default value)	The /tmp directory is used for staging transient files as part of data replication. Cloudera recommends a minimum of 50 GB size for the Datanode Replication Working Directory. Cloudera recommends you to use SSD for better performance during data replication. It can use the same SSD as datanode metadata.
hdds.datanode.replication.work.dir	</path/to/custom/replication/work/dir>	To use a non-default directory as Datanode Replication Working Directory
ozone.om.enable.filesystem.paths	true	
scm.container.client.max.size	<number of active pipelines in the cluster>	
scm.container.client.idle.threshold	120s	

Tuning Ozone for faster deletion performance

Learn how to use the Ozone configuration properties to tune Ozone for faster deletion performance under various workload conditions.

About this task

For deletion performance optimization and to overcome known issues, Cloudera recommends the following additional configurations based on different workload conditions:

Procedure

1. If you are experiencing slowness specifically when deleting deeply nested directory structures or a large number of directories, set the deletion configurations in the OzoneManager (OM) to the following values:

```
ozone.thread.number.dir.deletion = 30
ozone.directory.deleting.service.interval = 30s
```

2. When you observe general deletion slowness across the cluster or require faster overall delete throughput, update the following parameters in their respective services:

- For the OzoneManager (OM) configurations, use the following command:

```
ozone.block.deleting.service.interval = 30s
ozone.key.deleting.limit.per.task = 200000
ozone.directory.deleting.service.interval = 30s
```

- For the StorageContainerManager (SCM) configurations, use the following command:

```
hdds.scm.block.deletion.per-interval.max = 2000000
```

- For the DataNode (DN) configurations, use the following command:

```
hdds.datanode.block.deleting.limit.per.interval = 50000
ozone.block.deleting.service.workers = 30
hdds.datanode.block.deleting.service.interval = 30s
```

Related Information

[Configure optimized rename and recursive delete operations in Ranger Ozone plugin](#)

Node maintenance

You must provide the value for the `hdds.scm.replication.maintenance.remaining.redundancy` parameter to put the cluster's nodes under maintenance.

The `hdds.scm.replication.maintenance.remaining.redundancy` parameter defines the number of redundant containers in a group which must be available for a node to enter maintenance.

If you are putting a node into maintenance mode and reduces the redundancy below the value defined, the node will remain in the `ENTERING_MAINTENANCE` state until a new replica is created.

Parameter	Default value	Description
<code>hdds.scm.replication.maintenance.remaining.redundancy</code>	empty	The value is taken from the Ozone default configuration.
	1	The EC container has at least <code>dataNum + 1</code> online allowing the loss of 1 more replica before data becomes unavailable.
<code>hdds.scm.replication.maintenance.replica.minimum</code>	1	This is for Ratis containers only.



Note: For EC containers, if nodes are under maintenance, reconstruction reads might be required if some of the data replicas are offline. This is seamless to the client but affects the read performance.