

Storage Management Options

Date published:

Date modified:

CLOUDERA

Legal Notice

© Cloudera Inc. 2024. 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

Apache Ozone - Next-Gen Storage.....	4
Apache Hadoop HDFS.....	4
Migrating your data from HDFS to Ozone.....	4
Ozone Parcel.....	4
Introduction to Ozone Parcels.....	4
Ozone parcel block diagram.....	5
OZONE 718.2.5.....	5
Requirements.....	6
Downloading the Ozone parcels.....	6
Installing or Upgrading Ozone 718.2.5 parcel.....	6
Performance tuning for Ozone.....	7
Ozone HttpFS gateway.....	8
Configure S3 for system log using CDP CLI.....	14
Release notes.....	14
OZONE 718.1.0.....	32
Requirements.....	33
Downloading the Ozone parcels.....	33
Installing Ozone 718.1.0 parcel.....	33
Release notes.....	34

Apache Ozone - Next-Gen Storage

Apache Ozone is a scalable, redundant, and distributed object store optimized for big data workloads. Apart from scaling to billions of objects of varying sizes, applications that use frameworks like Apache Spark, Apache YARN and Apache Hive work natively on Ozone without any modifications. Ozone natively supports the S3 API and provides a Hadoop-compatible file system interface. Ozone is typically available in a CDP Private Cloud Base deployment.

Ozone consists of three important storage elements: volumes, buckets, and keys. Each key is part of a bucket, which, in turn, belongs to a volume. Only an administrator can create volumes. Depending on their requirements, regular users can create buckets in volumes. Ozone stores data as keys inside these buckets.

When a client writes a key, Ozone stores the associated data on DataNodes in chunks called blocks. Therefore, each key is associated with one or more blocks. Within a DataNode, multiple unrelated blocks can reside in a storage container.

To understand the concepts of Ozone, you can refer to the [Ozone overview and introduction](#) documentation.

For storing and configuring Ozone, refer to the [Ozone How to](#) documentation.

Apache Hadoop HDFS

Hadoop Distributed File System (HDFS) is a Java-based file system that provides scalable and reliable data storage. An HDFS cluster contains a NameNode to manage the cluster namespace and DataNodes to store data.

To understand the concepts of HDFS, you can refer to the [HDFS overview and introduction](#) documentation.

For managing, configuring, and accessing HDFS, refer to the [HDFS How to](#) documentation.

Migrating your data from HDFS to Ozone

You can migrate data from your existing HDFS deployment to Ozone for benefits such as scaling the metadata to billions of files; supporting denser, more cost-effective storage nodes as a part of a strategy to segregate storage and compute; or cold storage and archival storage needs.

This document assumes that you have an existing HDFS cluster from which you want a part or whole data to be moved to Ozone. The Ozone service can exist in the same cluster as HDFS or in a different cluster.

For more information, see the [Migration from HDFS to Ozone](#) documentation.

Ozone Parcel

This document helps you to install and manage Ozone parcels on top of a CDP Private Cloud Base version.

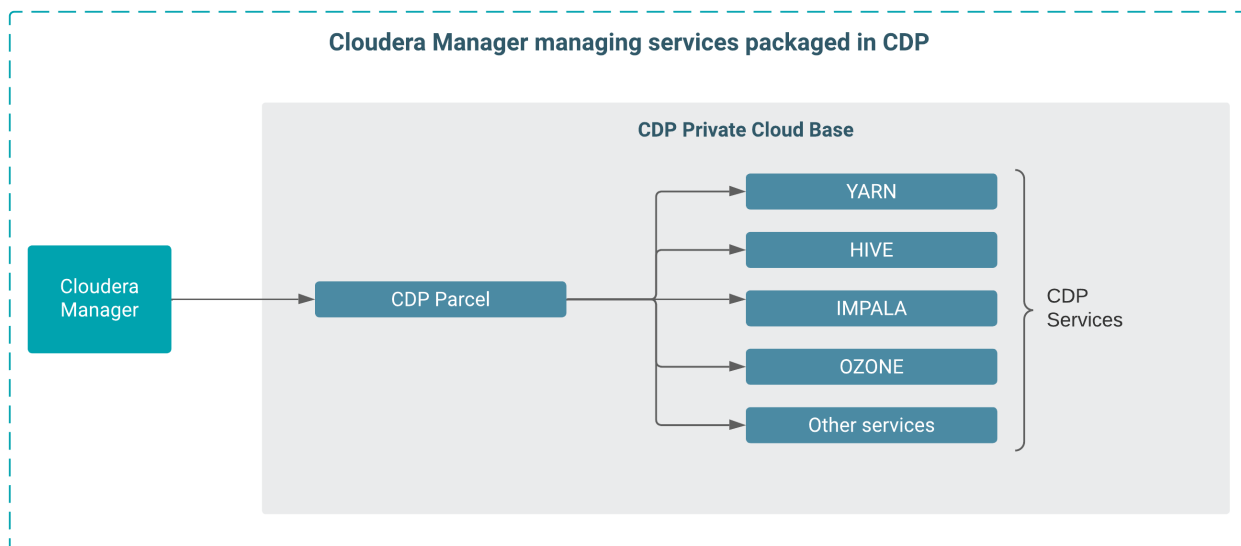
Introduction to Ozone Parcels

This document helps you to install and manage Ozone parcels on top of the CDP Private Cloud Base version.

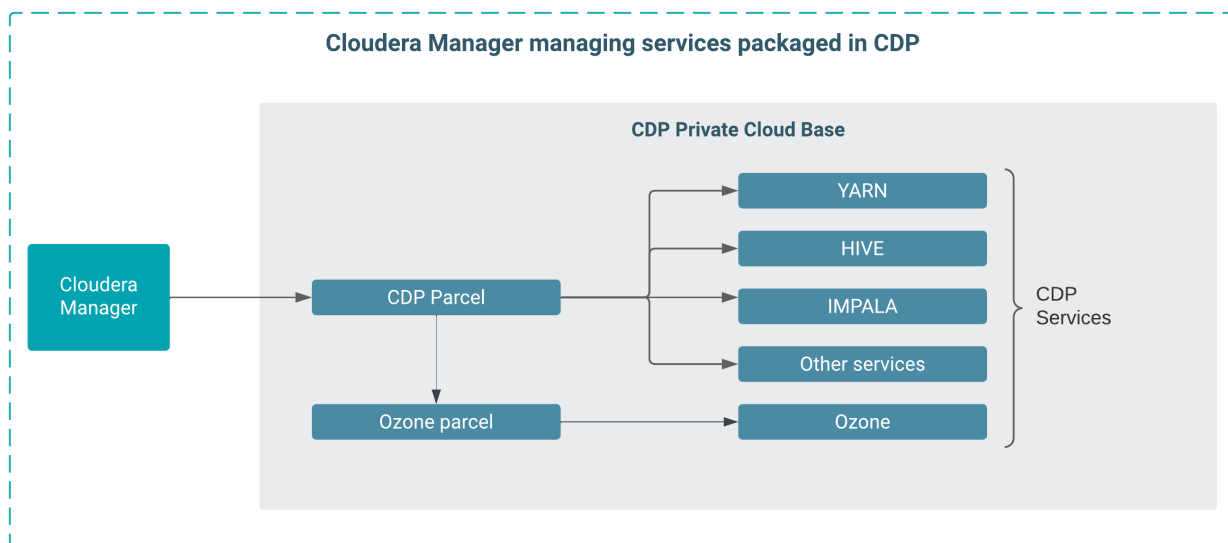
Ozone parcel releases enable you to upgrade the Ozone software at a cadence higher than LTS and feature releases of CDP. This enables delivering the upstream innovation in the open source version of Ozone sooner to customers.

Ozone parcel block diagram

This is a block diagram of CDP Private Cloud Base with Ozone and Ozone parcels as storage. You can upgrade Ozone parcels without upgrading the Ozone service or CDP Private Cloud Base cluster.



When the Ozone parcel is activated, only the Ozone software bits are picked up from the Ozone Parcel.



OZONE 718.2.5

Ozone 718.2.5 parcel must be installed on the CDP Private Cloud Base 7.1.8 cluster. You must be aware of the additional functionalities and improvements to features of components in Ozone 718.2.5 parcel. Learn how the new features and improvements benefit you.

- Refer to [Fixed issues](#) section for the list of issues fixed.

- Cloudera recommends you to upgrade to the latest Ozone 718.2.5 parcel.



Attention:

- If you are upgrading your cluster from Cloudera Runtime 7.1.8 with Ozone 718.2.x parcel to Cloudera Runtime 7.1.9, you must deactivate and undistribute the Ozone 718.2.x parcel and then proceed with the upgrade to 7.1.9.
- Using Cloudera Manager 7.7.1, if you are upgrading from [CDP Private Cloud Base 7.1.8 with Ozone parcels](#) to CDP Private Cloud Base 7.1.9 using Cloudera Manager 7.11.3, then you must
 - If you have [Ozone 718.1.0 parcel](#) on the CDP Private Cloud Base 7.1.8 cluster, you must deactivate and undistribute Ozone 718.1.0 parcel. If you have [Ozone 718.2.x parcel](#) on the CDP Private Cloud Base 7.1.8 cluster, you must only deactivate Ozone 718.2.x parcel.
 - Upgrade Cloudera Manager from 7.7.1 to 7.11.3. Do not restart the CDP cluster.
 - Upgrade from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9.

Requirements

Review the list of requirements before proceeding with the installation or upgrade.

- As a prevalidation for installation of the Ozone parcels, ensure that /var/lib/hadoop-hdfs is present on all the nodes and owned by the user hdfs.
- CDP 7.1.8 cluster with CHF 4 or greater installed and corresponding Cloudera Manager release.
- Ozone 718.2.1 parcel is certified to work with CHF4. However, parcels are expected to be not sensitive to CHF releases.
- Ozone Parcels are certified to work with Cloudera Manager version 7.10.1.
- The table below explains the product support matrix for Private Cloud Data Services, Cloudera Runtime, and Ozone parcel:

Product combination	Supportability
Private Cloud Data Services 1.5.1 with Cloudera Runtime 7.1.8 with CHF 4 or greater	Supported
Ozone 718.2.1 parcel on Cloudera Runtime 7.1.8	Supported
Private Cloud Data Services 1.5.1 with Ozone 718.2.1 parcel	Supported



Note: If you upgraded the cluster from CDP 7.1.7, then you must finalize the upgrade. Ozone parcel releases are not certified for rollback to CDP 7.1.7.

Downloading the Ozone parcels

This section helps you to download the Ozone parcel.

To download the Ozone parcel, use this link: <https://archive.cloudera.com/p/ozone/718.2.5/parcels/ozone/>



Note: Cloudera recommends you to upgrade to the latest Ozone 718.2.5 parcel.

Installing or Upgrading Ozone 718.2.5 parcel

To install or upgrade the Ozone 718.2.5 parcel, perform the following procedure:

Before you begin

- As a prevalidation for installation of the Ozone parcels, ensure that /var/lib/hadoop-hdfs is present on all the nodes and owned by the user hdfs.
- You must check the cluster state and run the required sanity checks before finalizing upgrade.

Procedure

1. Log in to Cloudera Manager
2. On the left navigation pane, click Parcels
3. Click Parcel Repositories & Network Settings
4. In the Remote Parcel Repository URLs section, click the "+" icon and add the *Ozone 718.2.5 URL* for your Parcel repository. For the Ozone parcel URL, refer to the [Downloading the Ozone parcel](#) section.
5. Click Save & Verify Configuration. A message with the status of the verification appears above the Remote Parcel Repository URLs section. If the URL is not valid, check the URL and enter the correct URL.
6. After the URL is verified, click Close.
7. Locate the row in the table that contains the new Ozone parcel and click the Download button.
8. After the download of the new Ozone parcel is complete, click the Distribute button.
Wait for the parcel to be distributed and unpacked before continuing. Cloudera Manager displays the status of the Cloudera Runtime parcel distribution. Click on the status display to view detailed status for each host.
9. Prepare the Ozone service for an upgrade before upgrading the Ozone software. Click Prepare for Upgrade under the Action drop-down.
10. Click Stop under the Action drop-down to stop the Ozone service
11. Click Activate on the Parcels page. If Ozone 718.1.0 parcel is already installed, this step will deactivate Ozone 718.1.0 parcel.
12. Click Finish Upgrade under the Action drop-down. Continue with next step.
13. Update YARN to use the new Ozone Client Jars. Ozone parcel activation or install should handle redeployment of YARN jars and clean-up cache. Perform the following steps:
 - a) Log in to Cloudera Manager UI.
 - b) Navigate to Clusters
 - c) Select the YARN service
 - d) Click Actions
 - e) Click Install YARN Service Dependencies
 - f) Click Install YARN MapReduce Framework JARs
14. Update Tez Configuration: (If this is not done as part of a previous Ozone parcel installation.)
 - a) For `tez.cluster.additional.classpath.prefix`, the value is `/var/lib/hadoop-hdfs/*` (Tez Additional Classpath)
 - b) For `tez.user.classpath.first`, the value is `true` (Tez Client Advanced Configuration Snippet (Safety Valve) for `tez-conf/tez-site.xml`)
15. Update Hive On Tez configuration: (If this is not done as part of a previous Ozone parcel installation.)
 - a) For `tez.cluster.additional.classpath.prefix` the value is `/var/lib/hadoop-hdfs/*` (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)
 - b) For `tez.user.classpath.first`, the value is `true` (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)
16. Restart the CDP 7.1.8 cluster: Click the Actions menu and click Restart.
17. Click Finalize Upgrade under the Action drop-down.

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	

Configuration Property	Value	Remarks
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	
hdds.container.report.interval	10m	
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.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	EdgortinSSCMContainerPlacementCapacity on capacity available instead of random allocation.
Ozone Configuration Ozone Recon Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-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/ozone-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	

Ozone HttpFS gateway

Ozone now supports HttpFS gateway. This allows Ozone to integrate with other tools through REST APIs. HttpFS service is a Jetty based web application that uses the Hadoop FileSystem API to talk to the cluster. HttpFS service is an optional role available in an Ozone cluster.

You can use HttpFS as a gateway to access data on an Ozone cluster behind a firewall. HttpFS is the only system that is allowed to cross the firewall into the Ozone cluster. You can use HttpFS to access data in Ozone using:

- HTTP utilities such as curl and wget
- HTTP libraries Perl from languages other than Java

HttpFS service provides REST HTTP gateway supporting File System operations such as read and write. HttpFS service is interoperable with the webhdfs REST HTTP API. You can use the webhdfs client FileSystem implementation to access HttpFS using:

- Ozone filesystem CLI (ozone fs)
- Java applications using the Hadoop FileSystem Java API.

HttpFS has built-in security that supports the following:

- Hadoop pseudo authentication
- Kerberos SPNEGO
- Other pluggable authentication mechanisms
- Hadoop proxy user support

Supported operations

Review the list of WebHDFS REST APIs and their state of support in Ozone.

File and directory operations

Operation	Support
Create and Write to a File	Supported
Append to a File	Not implemented in Ozone
Concat Files	Not implemented in Ozone
Open and Read a File	Supported
Make a Directory	Supported
Create a Symbolic Link	Not implemented in Ozone
Rename a File or Directory	Supported with limitations
Delete a File or Directory	Supported
Truncate a File	Not implemented in Ozone
Status of a File or Directory	Supported
List a Directory	Supported
List a File	Supported
Iteratively List a Directory	Unsupported

Other file system operations

Operation	Support
Get Content Summary of a Directory	Supported
Get Quota Usage of a Directory	Supported
Set Quota	Not implemented in Ozone FileSystem API

Operation	Support
Set Quota By Storage Type	Not implemented in Ozone
Get File Checksum	Unsupported
Get Home Directory	Unsupported
Get Trash Root	Unsupported
Set Permission	Not implemented in Ozone FileSystem API
Set Owner	Not implemented in Ozone FileSystem API
Set Replication Factor	Not implemented in Ozone FileSystem API
Set Access or Modification Time	Not implemented in Ozone FileSystem API
Modify ACL Entries	Not implemented in Ozone FileSystem API
Remove ACL Entries	Not implemented in Ozone FileSystem API
Remove Default ACL	Not implemented in Ozone FileSystem API
Remove ACL	Not implemented in Ozone FileSystem API
Set ACL	Not implemented in Ozone FileSystem API
Get ACL Status	Not implemented in Ozone FileSystem API
Check access	Not implemented in Ozone FileSystem API

Adding Ozone HttpFS gateway

Ozone now supports HttpFS gateway. This allows Ozone to integrate with other tools through REST APIs. HttpFS service is a Jetty based web application that uses the Hadoop FileSystem API to talk to the cluster. HttpFS service is an optional role available in an Ozone cluster.

Before you begin

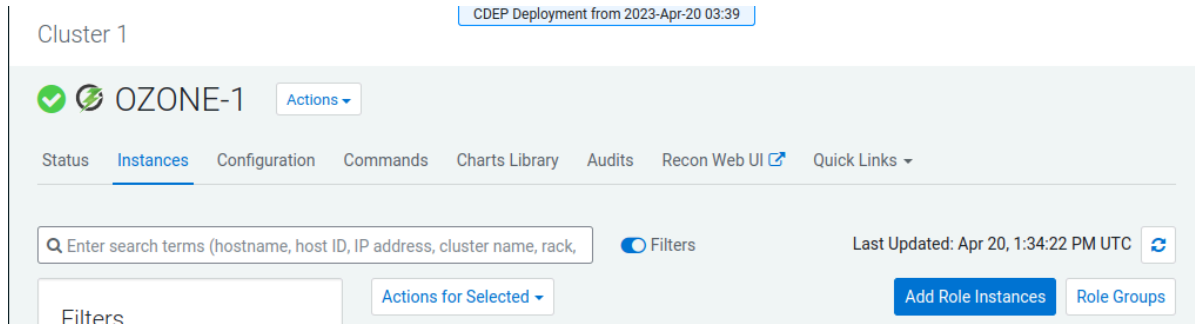
- You must have Ozone 718.2.x on the Ozone cluster.
- Ensure the Cloudera Manager version is 7.10.1.

Procedure

1. Restart the CDP Private Cloud Base cluster to ensure dependent components use the correct version of Ozone.

2. Assign the HttpFS Gateway role to one of the instances.

- a) Log in to Cloudera Manager
- b) Go to Clusters Ozone-1 Instances
- c) Click Add Role Instances



- d) Select one of the hosts for the HttpFS Gateway role and Save
You should see the following roles:



OZONE-1

Actions ▾

[Status](#)[Instances](#)[Configuration](#)[Commands](#)[Charts Library](#)[Alerts](#)

Health Tests

[Show 3 Good](#)[Ozone Basic Canary](#)

Test disabled by user: This health test checks if Ozone basic volume/bucket/key is working.

Status Summary

Gateway



3 None

HttpFS Gateway



1 Good Health

Ozone DataNode



3 Good Health

Ozone Manager



3 Good Health

Ozone Recon



1 Good Health

S3 Gateway



1 Good Health

Storage Container Manager



3 Good Health

[Hosts](#)

3 Good Health

e) Start the HttpFS Gateway role.

<input type="checkbox"/>	Status	Role Type	Tags	State
<input type="checkbox"/>	✓	Storage Container Manager		Started
<input type="checkbox"/>	●	Gateway		N/A
<input type="checkbox"/>	●	Gateway		N/A
<input type="checkbox"/>	●	Gateway		N/A
<input type="checkbox"/>	⬢	HttpFS Gateway		Stopped
<input type="checkbox"/>	✓	Ozone DataNode		Started
<input type="checkbox"/>	✓	Ozone DataNode		Started
<input type="checkbox"/>	✓	Ozone DataNode		Started
<input type="checkbox"/>	✓	Ozone Manager		Started
<input type="checkbox"/>	✓	Ozone Manager		Started
<input type="checkbox"/>	✓	Ozone Manager		Started
<input type="checkbox"/>	✓	Ozone Recon		Started
<input type="checkbox"/>	✓	S3 Gateway		Started
<input type="checkbox"/>	✓	Storage Container Manager		Started
<input type="checkbox"/>	✓	Storage Container Manager		Started

Configuring TLS/SSL encryption manually for Ozone HttpFS gateway

You can configure TLS/SSL encryption manually for Ozone HttpFS gateway using Cloudera Manager.

Before you begin


- The keystores containing certificates that are bound to the proper domain names are accessible on all the hosts on which at least one Ozone role instance is running.
- The hdfs user has read permissions to the keystore files for Ozone.
- You must specify absolute paths to the keystore file. These settings apply to all hosts on which the various Ozone role instances run. Therefore, the paths that you specify must be valid on all the hosts. In addition, the keystore file names for Ozone must be the same on all hosts.
- The Ozone HttpFS gateway is only available from the Ozone 718.2.0 parcel and from the Cloudera Manager version 7.10.1. For more information on other supported gateways, see [Configuring TLS/SSL encryption manually for Ozone](#).

Consider an example where you have separate certificates for the Ozone role instances on hosts node1.example.com and node2.example.com, and you have chosen to store the certificates in files with names ozone-node1.jks and ozone-node2.jks respectively. When deploying these keystores, you must provide them both with the same name on the target host, for example, ozone.jks.

Procedure

1. In Cloudera Manager, go to Ozone Configuration .
2. Search for *tls/ssl*.
3. Enter the TLS/SSL properties for the various Ozone roles.

Provide the values for the properties corresponding to the following fields on the Ozone configuration page:

- Enable TLS/SSL for HttpFS Gateway
- HttpFS Gateway TLS/SSL Server JKS Keystore File Location
- HttpFS Gateway TLS/SSL Server JKS Keystore File Password
- HttpFS Gateway TLS/SSL Server JKS Keystore Key Password
-  **Note:** Ensure that you configure the `ozone.prometheus.ca.file` property when configuring TLS/SSL encryption for Ozone. If the location of the root CA certificate is not specified, the location of the auto-TLS CA certificate is considered as the default value for this property.

4. Click Save Changes and restart the Ozone service.

Configure S3 for system log using CDP CLI

To configure S3 using CDP CLI, follow the below procedure.

Before you begin

- This operation requires CDP CLI version 0.9.84 or above.
- This procedure is applicable under following conditions:
 - If you did not have Ozone when you created the base cluster.
 - Later you decide to use CDE on the existing base cluster and add Ozone to the existing base cluster.

Procedure

1. Create an S3 bucket using the following steps:
 - a) Make use of the Cloudera Manager API call `/clusters/{clusterName}/getOzoneS3GatewayInfo`.
 - b) The `clusterName` is the base cluster that has the Ozone service.
 - c) You must provide a bucket name by following the guidelines mentioned in the [Bucket naming rules](#) documentation. You must note down the bucket name as this is different from the return bucket field mentioned below.
The API call is in progress and you must wait for the API call to complete.
 - d) After the API call is complete, save the returned json output. This looks similar to `{ "awsAccessKey": ... "awsSecret": ... "restUrl": ... "bucket": ... }`

2. Run the following `cdpcl` command using the values returned above:

```
cdp environments --endpoint-url <console-url> --no-verify-tls --no-
parameter-expansion --set-environment-setting --environment-name
<environment-name> --settings '{ "logType": "OZONE", "ozoneS3Key":
"<awsAccessKey>", "ozoneS3Secret": "<awsSecret>", "ozoneS3RestUrl":
"<restUrl>", "ozoneLogsBucket": "<bucket-name>", "ozoneLogsPath":
"<bucket>" }'
```

3. Verify that the environment settings have been correctly configured by running `cdp environments`

```
--no-verify-tls get-environment-setting --environment-name <environment-
name> --attrs logType ozoneS3Key ozoneS3Secret ozoneS3RestUrl
ozoneLogsBucket ozoneLogsPath
```

Release notes

This document helps you to install and manage Ozone 718.2.5 on top of CDP Private Cloud Base 7.1.8.

Fixed issues in Ozone parcel 718.2.5

You can review the list of reported issues and their fixes in Ozone parcel 718.2.5. Fixed issues represent selected issues that were previously logged through Cloudera Support, but are now addressed in the current Ozone parcel release. These issues may have been reported in previous versions of Runtime as a known issue; meaning they were reported by customers or identified by Cloudera Quality Engineering teams.

Ozone

- CDPD-69558: Disable ozone snapshot feature by default in Ozone parcel 7.2.x
- CDPD-69099: OM state machine should add response for every write request to the double-buffer
- CDPD-69192: Generate linear keys for freon ockrw
- CDPD-68197: EC Reconstruction fails with IOException: None of the block data have checksum after upgrade
- CDPD-68924: Do not fail read of EC block if the last chunk is empty
- CDPD-68723: EC Reconstruction creates empty chunks at the end of blocks with partial stripes
- CDPD-68724: EC Reconstruction does not issue put block to data index if it is unused
- CDPD-69780: Backport HDDS-9003 to OZONE 718.2.x
- CDPD-70830: OM DB cache metrics creation should be idempotent
- CDPD-66996: Few acceptance tests fail with Docker Compose V2
- CDPD-59986: Ozone freon tool extension for benchmarking listKeys operation
- CDPD-61466: Default reads/lists to 0 (freon ockrw)
- CDPD-52472: [FSO] S3G folder support fso layout filestatus s3A compatibility
- CDPD-56904: Clean up checked functional interfaces
- CDPD-63802: ChunkInputStream should use new token after pipeline refresh
- CDPD-71328: Ensure version is always set in ContainerCommandRequestProto
- CDPD-70061: EC Key read corruption when the replica index of container in DN mismatches
- CDPD-58159: Random EC read pipeline ID causes XceiverClient cache churn
- CDPD-71993: Revert client version bump done has part of HDDS-10983
- CDPD-72527: Fix hadoop2.7 test
- CDPD-52707: Support namespace summaries (du, dist & counts) for OBJECT_STORE buckets
- CDPD-72020: Some containers affected by HDDS-8129 may still be in the DELETING state incorrectly

Ratis or Ratis-thirdparty

- CDPD-69481: Backport RATIS-1876.

Known issues in Ozone parcel 718.2.5

You must be aware of the known issues and limitations, the areas of impact, and workaround in Ozone parcel.

OPSAPS-69859: Cloudera Manager does not generate correct configurations in a non-HA cluster.

This causes the Hadoop DistCp operation to fail. You must use the workaround if you are on Cloudera Manager 7.11.3 Cumulative hotfix 6 and below.



Note: This issue is fixed in Cloudera Manager 7.11.3 Cumulative hotfix 7 and higher. For more information, see the [Cloudera Manager 7.11.3 Cumulative hotfix 7](#) release notes.

1. Log in to Cloudera Manager.
2. Navigate to Clusters Ozone .
3. Click the Configuration tab.
4. In the search bar, search for the Ozone Client Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml property.

5. Add the following key-value pairs:
 - a. ozone.om.service.ids = <Ozone Service ID>
 - b. ozone.om.address = <OM primordial node with port number >
 - c. ozone.om.nodes.<ozone-service-id> = <valve>
 - d. ozone.om.address.<ozone-service-id>.<om-service-id> = <OM primordial node with port number>
6. Click Save Changes.
7. Restart the stale configuration.

The list of known issues for the Ozone parcel 718.2.5 includes the known issues from Ozone parcel 718.2.4. For more information, see [Known issues in Ozone parcel 718.2.4](#).

Fixed issues in Ozone parcel 718.2.4

You can review the list of reported issues and their fixes in Ozone parcel 718.2.4. Fixed issues represent selected issues that were previously logged through Cloudera Support, but are now addressed in the current Ozone parcel release. These issues may have been reported in previous versions of Runtime as a known issue; meaning they were reported by customers or identified by Cloudera Quality Engineering teams.

Ozone

- CDPD-63635: [AUTOSYNC] ContainerStateMachine.applyTransaction(..) should not validate token again
- CDPD-52478: Support Ozone OM Metadata Generator in `Ozone freon`
- CDPD-57631: [AUTOSYNC] Orphan blocks can leave empty container stuck deleting forever
- CDPD-52600: Do not use block count to determine if containers are empty
- CDPD-64419: OM nodes went down due to OOM, possible memory leak
- CDPD-63606: Datanodes do not Retry Pipeline Close Commands for SCM
- CDPD-63605: Verbose log messages regarding pipeline creation failure
- CDPD-63588: Do not show empty containers as missing in Recon UI
- CDPD-63281: [AUTOSYNC] Backport HDDS-9550 to legacy RM (missing containers which are empty)
- CDPD-62957: [AUTOSYNC] Container report shows missing containers when they actually appear empty
- CDPD-59221: Ozone Manager terminated in CreateKey request due to bucket not found
- CDPD-57621: [AUTOSYNC] Clean up datanode memory in case of an ozone stream write error
- CDPD-52398: Cannot delete empty OBS bucket if prefix of other, non-empty bucket
- CDPD-64709: HDDS-8993. Pin maven-antrun-plugin version to 3.1.0
- CDPD-64308: Legacy Replication Manager should consider that UNHEALTHY replicas might be decommissioning
- CDPD-64307: HDDS-9652. Legacy RM will not replicate all unhealthy containers when some are decommissioning

Ratis or Ratis-thirdparty

- Ratis-1928: Join the LogAppenders when closing the server.
- Ratis-1761: If LeaderStateImpl is not running, it should not restart a LogAppender.

Known issues in Ozone parcel 718.2.4

You must be aware of the known issues and limitations, the areas of impact, and workaround in Ozone parcel.

SSL Handshake fails between Ozone DataNodes if the two DataNodes have their certificate signed by different Ozone Storage Container Managers.

Ozone DataNode certificates are signed by the leader Storage Container Manager. Due to an issue in creating a TrustStore for DataNode to DataNode connections, the trust cannot be established between the two DataNodes if a different Storage Container Manager signs their certificate. These connections fail to establish and display an SSL Handshake Exception. This affects Pipeline creation and container replication (also EC container reconstruction). The symptoms vary,

depending on the number of the nodes that have different signer certificates, either these DataNodes do not participate in any Ratis-3 Pipeline, or have Pipelines exclusively in between groups with the same signer. Over time this can lead to an imbalance in DataNode usage, and it might cause decommission of a DataNode stuck if the data has to be replicated to a node with a certificate that has a different signer.

This problem affects all the 7.1.8 Ozone Parcel releases.

To identify if the problem is present on a cluster, the output of `ozone admin cert list` command must be examined. Ensure you define a sufficient number of certificates to be returned with the `-c` option to see all the certificates issued in the system.

If there are different Issuers for the latest DataNode certificates, this indicates the cluster is affected.

It is possible to avoid the problem by checking the `ozone admin scm roles` output and see if the Primordial node is the actual leader node before adding a new DataNode and starting it for the first time. If the leader SCM node is a different node, then calling `ozone admin scm transfer` to make the Primordial node the leader can put the cluster into the desired state before adding the new DataNode.

Ensure that all the certificates in the cluster are signed by the same Storage Container Manager node.

Perform the procedure below if the cluster is affected and has DataNode certificates signed by a different issuer:

1. Identify the current leader Storage Container Manager by running the command `ozone admin scm roles`.
2. If the leader did not sign the majority of DataNode certificates, change the leader that has signed the majority of DN certificates by running the `ozone admin scm transfer` command with the proper Ozone SCM Service ID (set in Cloudera Manager>Ozone>Configuration) and the UUID of the desired leader node.
3. Stop the minority of DataNodes that have a different sign than the majority.
4. Locate the Datanode Metadata Directory (set in Cloudera Manager>Ozone>Configuration) on the hosts of the stopped DataNode and move the directory to a backup location.
5. Start the previously stopped DataNodes.
6. After regenerating the certificates, check if the DataNodes are joining the Ratis-3 Pipeline.

To avoid service disruptions, you can stop the DataNodes one by one instead of all together. This can cause a data outage for the downtime of the node if there are files with Ratis-1 replication on the cluster and the single replica of these files reside on the node being restarted. Changing the leader node on an SCM HA-enabled cluster should not disrupt operations.

CDPD-65369: Service Host connectivity error: Target host is not specified

Run again with `confiq hive.server2.transport.mode`.

CDPD-56073: Impala-shell client throws DisconnectError when trying to connect.

None.

CDPD-65477: Ozone freon ommg does not work.

None.

CDPD-65491: Recon displaying DELETED container as missing.

None.

CDPD-64844: Hadoop put operation on File System Optimised is broken.

None.

CDPD-56486: [Spark] Ozone delete key failed error during Spark job completion

None.

The list of known issues for the Ozone parcel 718.2.4 includes the known issues from Ozone parcel 718.2.2/718.2.3 also. For more information, see the [Known issues in Ozone parcel 718.2.2](#) documentation.

Fixed issues in Ozone parcel 718.2.3

You can review the list of reported issues and their fixes in Ozone parcel 718.2.3. Fixed issues represent selected issues that were previously logged through Cloudera Support, but are now addressed in the current Ozone parcel release. These issues may have been reported in previous versions of Runtime as a known issue; meaning they were reported by customers or identified by Cloudera Quality Engineering teams.

Ozone

- HDDS-9489. LegacyReplicationManager: Do not count unique origin nodes as over-replicated.
- HDDS-9239. Ozone cli command to get container info should deal with empty values for --json.
- CDPD-63163: Make decom logging limit default to 1000 and be configurable.
- HDDS-9555. Decommission should not wait on deleting containers.
- HDDS-9354. LegacyReplicationManager: Unhealthy replicas of a sufficiently replicated container can block decommissioning.
- HDDS-9321. LegacyReplicationManager: Save unhealthy replicas with highest BCSID for a QUASI_CLOSED container.
- HDDS-9352. LegacyReplicationManager: Ignore any Datanodes that are not in-service and healthy when finding unique origins.
- HDDS-9327. LegacyReplicationManager: Handle all unhealthy replicas of a CLOSING container.
- HDDS-9254. Legacy replication manager uses mismatched replicas as replication sources.
- HDDS-7907. [hsync] KeyOutputStream is not thread safe.
- HDDS-9275. LegacyReplicationManager: Delete excess unhealthy with force=true.
- HDDS-7882. Delete empty containers that are stuck in closing state.
- HDDS-9257. LegacyReplicationManager: Unhealthy replicas could block under replication handling.
- HDDS-8536. ReplicationManager: Unhealthy replicas could block Ratis containers being replicated.
- HDDS-9075. QUASI_CLOSED Replica with incorrect sequenceID should be deleted by SCM.
- HDDS-8616. Underreplication not fixed if all replicas start decommissioning.
- HDDS-8459. Ratis under replication handling in a rack aware environment does not work.
- HDDS-9243. Add assertions to BlockOutputStream.
- HDDS-9254: Only choose matching replicas when available + unit test.

Ratis or Ratis-thirdparty

None

Known issues in Ozone parcel 718.2.3

You must be aware of the known issues and limitations, the areas of impact, and workaround in Ozone parcel.

The list of known issues for the Ozone parcel 718.2.3 remains the same as Ozone parcel 718.2.2. For more information, see the [Known issues in Ozone parcel 718.2.2](#) documentation.

Fixed issues in Ozone parcel 718.2.2

You can review the list of reported issues and their fixes in Ozone parcel 718.2.2. Fixed issues represent selected issues that were previously logged through Cloudera Support, but are now addressed in the current Ozone parcel release. These issues may have been reported in previous versions of Runtime as a known issue; meaning they were reported by customers or identified by Cloudera Quality Engineering teams.

Ozone

- HDDS-8417: Cap on queue of commands at datanode
- Revert CDPD-54746: Backport HDDS-8324 to Ozone Parcel 2
- HDDS-9081: Handling Netty back pressure when streaming containers for replication

- HDDS-9134: GRPC based replication can get stuck forever if the receiver is not available
- HDDS-9070: Expose Netty memory used metrics
- HDDS-9148: OM to reject hsync if ozone.fs.hsync.enabled is false

Ratis

- RATIS-1836: Frequent retry by leader on failure path for appendEntry
- RATIS-1868: Handling Netty back pressure when streaming ratis log

Technical Service Bulletins

TSB 2023-697: Potential Ozone data loss in Ozone Parcel 2 caused by hsync/hflush

Knowledge article: [TSB 2023-697: Potential Ozone data loss in Ozone Parcel 2 caused by hsync/hflush](#)

Known issues in Ozone parcel 718.2.2

You must be aware of the known issues and limitations, the areas of impact, and workaround in Ozone parcel.

Tez Configuration Changes

The following configuration changes have to be made to pick up the latest Ozone FS jar from the Ozone parcel (when installed):

CDPD-48540

For `tez.cluster.additional.classpath.prefix`, the value is `/var/lib/hadoop-hdfs/*` (Tez Additional Classpath)

For `tez.user.classpath.first`, the value is `true` (Tez Client Advanced Configuration Snippet (Safety Valve) for `tez-conf/tez-site.xml`)

CDPD-47605

For `tez.cluster.additional.classpath.prefix` the value is `/var/lib/hadoop-hdfs/*` (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

For `tez.user.classpath.first`, the value is `true` (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

Restart the TEZ and HIVE_ON_TEZ service as prompted by Cloudera Manager.

Update Yarn to use updated Ozone FS Jar

CDPD-48500: Ozone parcel activation or install should handle redeployment of YARN jars and clean-up cache.

Perform the following steps:

1. Log in to Cloudera Manager UI
2. Navigate to Clusters
3. Select the YARN service
4. Click Actions
5. Click Install YARN Service Dependencies
6. Click YARN MapReduce Framework JARs
7. Restart the CDP 7.1.8 cluster

Other issues

SSL Handshake fails between Ozone DataNodes if the two DataNodes have their certificate signed by different Ozone Storage Container Managers.

Ozone DataNode certificates are signed by the leader Storage Container Manager. Due to an issue in creating a TrustStore for DataNode to DataNode connections, the trust cannot be established

between the two DataNodes if a different Storage Container Manager signs their certificate. These connections fail to establish and display an SSL Handshake Exception. This affects Pipeline creation and container replication (also EC container reconstruction). The symptoms vary, depending on the number of the nodes that have different singer certificates, either these DataNodes do not participate in any Ratis-3 Pipeline, or have Pipelines exclusively in between groups with the same signer. Over time this can lead to an imbalance in DataNode usage, and it might cause decommission of a DataNode stuck if the data has to be replicated to a node with a certificate that has a different signer.

This problem affects all the 7.1.8 Ozone Parcel releases.

To identify if the problem is present on a cluster, the output of `ozone admin cert list` command must be examined. Ensure you define a sufficient number of certificates to be returned with the `-c` option to see all the certificates issued in the system.

If there are different Issuers for the latest DataNode certificates, this indicates the cluster is affected.

It is possible to avoid the problem by checking the `ozone admin scm roles` output and see if the Primordial node is the actual leader node before adding a new DataNode and starting it for the first time. If the leader SCM node is a different node, then calling `ozone admin scm transfer` to make the Primordial node the leader can put the cluster into the desired state before adding the new DataNode.

Ensure that all the certificates in the cluster are signed by the same Storage Container Manager node.

Perform the procedure below if the cluster is affected and has DataNode certificates signed by a different issuer:

1. Identify the current leader Storage Container Manager by running the command `ozone admin scm roles`.
2. If the leader did not sign the majority of DataNode certificates, change the leader that has signed the majority of DN certificates by running the `ozone admin scm transfer` command with the proper Ozone SCM Service ID (set in Cloudera Manager>Ozone>Configuration) and the UUID of the desired leader node.
3. Stop the minority of DataNodes that have a different sign than the majority.
4. Locate the Datanode Metadata Directory (set in Cloudera Manager>Ozone>Configuration) on the hosts of the stopped DataNode and move the directory to a backup location.
5. Start the previously stopped DataNodes.
6. After regenerating the certificates, check if the DataNodes are joining the Ratis-3 Pipeline.

To avoid service disruptions, you can stop the DataNodes one by one instead of all together. This can cause a data outage for the downtime of the node if there are files with Ratis-1 replication on the cluster and the single replica of these files reside on the node being restarted. Changing the leader node on an SCM HA-enabled cluster should not disrupt operations.

CDPD-56006: On providing an incorrect hostname/service ID in the Ozone URI, the filesystem client instead of failing, retries till exhaustion and the default retry is too high.

Configure `ozone.client.failover.max.attempts` to a lower value to avoid long endless retries.

CDPD-49137: Sometimes OM's kerberos token is not updated and it stops being able to communicate with SCM. When this occurs, writes will start to fail.

Restarting OM or setting the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true` will fix the issue.

CDPD-49808: Spark jobs against Ozone intermittently fail with ERROR spark.SparkContext: [main]: Error initializing SparkContext.java.lang.IllegalStateException: No filter named.

This is an intermittent failure which can be retried.

CDPD-50678: Deleting containers which have one or more replicas which are not empty on the Datanode can cause the container to be stuck in a deleting state indefinitely. Containers in this state can also block decommission or maintenance operations completing.

None.

CDPD-52571/HDDS-8178: CertificateClient and KeyStoresFactory support multiple Sub-CA certificates in the trust chain

None.

CDPD-35141: Error: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive.ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2> (state=08S01,code=40000) java.sql.SQLException: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive.ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2>.

We may see the above issue if the source and target buckets are different in Hive queries. For now, copying across the same bucket is only supported.

Avoid different buckets in source and target path.

CDPD-60578: [EC] Re-replication failing due to NullPointerException.

None.

Fixed issues in Ozone parcel 718.2.1

You can review the list of reported issues and their fixes in Ozone parcel 718.2.1. Fixed issues represent selected issues that were previously logged through Cloudera Support, but are now addressed in the current Ozone parcel release. These issues may have been reported in previous versions of Runtime as a known issue; meaning they were reported by customers or identified by Cloudera Quality Engineering teams.

RATIS-1783: MAX_OP_SIZE is not configurable on raft log read (ratis)

This issue is resolved.

RATIS-1760: Bump Netty to 4.1.86.Final to fix CVE-2022-41881 (ratis-thirdparty)

This issue is resolved.

Known issues in Ozone parcel 718.2.1

You must be aware of the known issues and limitations, the areas of impact, and workaround in Ozone parcel.

Tez Configuration Changes

The following configuration changes have to be made to pick up the latest Ozone FS jar from the Ozone parcel (when installed):

CDPD-48540

For `tez.cluster.additional.classpath.prefix`, the value is `/var/lib/hadoop-hdfs/*` (Tez Additional Classpath)

For `tez.user.classpath.first`, the value is `true` (Tez Client Advanced Configuration Snippet (Safety Valve) for `tez-conf/tez-site.xml`)

CDPD-47605

For `tez.cluster.additional.classpath.prefix` the value is `/var/lib/hadoop-hdfs/*` (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

For `tez.user.classpath.first`, the value is `true` (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

Restart the TEZ and HIVE_ON_TEZ service as prompted by Cloudera Manager.

Update Yarn to use updated Ozone FS Jar

CDPD-48500: Ozone parcel activation or install should handle redeployment of YARN jars and clean-up cache.

Perform the following steps:

1. Log in to Cloudera Manager UI
2. Navigate to Clusters
3. Select the YARN service
4. Click Actions
5. Click Install YARN Service Dependencies
6. Click YARN MapReduce Framework JARs
7. Restart the CDP 7.1.8 cluster

Other issues

SSL Handshake fails between Ozone DataNodes if the two DataNodes have their certificate signed by different Ozone Storage Container Managers.

Ozone DataNode certificates are signed by the leader Storage Container Manager. Due to an issue in creating a TrustStore for DataNode to DataNode connections, the trust cannot be established between the two DataNodes if a different Storage Container Manager signs their certificate. These connections fail to establish and display an SSL Handshake Exception. This affects Pipeline creation and container replication (also EC container reconstruction). The symptoms vary, depending on the number of the nodes that have different singer certificates, either these DataNodes do not participate in any Ratis-3 Pipeline, or have Pipelines exclusively in between groups with the same signer. Over time this can lead to an imbalance in DataNode usage, and it might cause decommission of a DataNode stuck if the data has to be replicated to a node with a certificate that has a different signer.

This problem affects all the 7.1.8 Ozone Parcel releases.

To identify if the problem is present on a cluster, the output of `ozone admin cert list` command must be examined. Ensure you define a sufficient number of certificates to be returned with the `-c` option to see all the certificates issued in the system.

If there are different Issuers for the latest DataNode certificates, this indicates the cluster is affected.

It is possible to avoid the problem by checking the `ozone admin scm roles` output and see if the Primordial node is the actual leader node before adding a new DataNode and starting it for the first time. If the leader SCM node is a different node, then calling `ozone admin scm transfer` to make the Primordial node the leader can put the cluster into the desired state before adding the new DataNode.

Ensure that all the certificates in the cluster are signed by the same Storage Container Manager node.

Perform the procedure below if the cluster is affected and has DataNode certificates signed by a different issuer:

1. Identify the current leader Storage Container Manager by running the command `ozone admin scm roles`.
2. If the leader did not sign the majority of DataNode certificates, change the leader that has signed the majority of DN certificates by running the `ozone admin scm transfer` command with the proper Ozone SCM Service ID (set in Cloudera Manager>Ozone>Configuration) and the UUID of the desired leader node.
3. Stop the minority of DataNodes that have a different sign than the majority.
4. Locate the Datanode Metadata Directory (set in Cloudera Manager>Ozone>Configuration) on the hosts of the stopped DataNode and move the directory to a backup location.
5. Start the previously stopped DataNodes.
6. After regenerating the certificates, check if the DataNodes are joining the Ratis-3 Pipeline.

To avoid service disruptions, you can stop the DataNodes one by one instead of all together. This can cause a data outage for the downtime of the node if there are files with Ratis-1 replication on the

cluster and the single replica of these files reside on the node being restarted. Changing the leader node on an SCM HA-enabled cluster should not disrupt operations.

CDPD-56006: On providing an incorrect hostname/service ID in the Ozone URI, the filesystem client instead of failing, retries till exhaustion and the default retry is too high.

Configure `ozone.client.failover.max.attempts` to a lower value to avoid long endless retries.

CDPD-49137: Sometimes OM's kerberos token is not updated and it stops being able to communicate with SCM. When this occurs, writes will start to fail.

Restarting OM or setting the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true` will fix the issue.

CDPD-49808: Spark jobs against Ozone intermittently fail with ERROR spark.SparkContext: [main]: Error initializing SparkContext.java.lang.IllegalStateException: No filter named.

This is an intermittent failure which can be retried.

CDPD-50678: Deleting containers which have one or more replicas which are not empty on the Datanode can cause the container to be stuck in a deleting state indefinitely. Containers in this state can also block decommission or maintenance operations completing.

None.

Technical Service Bulletins

TSB 2023-697: Potential Ozone data loss in Ozone Parcel 2 caused by hsync/hflush

A recently identified regression caused by an issue in the implementation of `hsync/hflush` Application Programming Interfaces (APIs) can result in data loss in Apache Ozone (Ozone) if the client application calls `hsync/hflush` before closing the file.

Once the file is closed, the blocks of the file will be removed by the Storage Container Manager (SCM) after the container with these blocks gets closed. Reading the file will fail with the error message similar to the following: `org.apache.hadoop.hdds.scm.container.common.helpers.StorageContainerException: Unable to find the block. BlockID : conID: 4 locID: 111677748019201149 bc sId: 382`

To avoid potential data loss caused by the issue described above, Users on the affected parcel versions must upgrade to OZONE-718.2.2 immediately.

Upstream Jira: [HDDS-9146\(fix\)](#)

Knowledge article: [TSB 2023-697: Potential Ozone data loss in Ozone Parcel 2 caused by hsync/hflush](#)

Fixed issues in Ozone parcel 718.2.0

You can review the list of reported issues and their fixes in Ozone parcel 718.2.0. Fixed issues represent selected issues that were previously logged through Cloudera Support, but are now addressed in the current Ozone parcel release. These issues may have been reported in previous versions of Runtime as a known issue; meaning they were reported by customers or identified by Cloudera Quality Engineering teams.

CDPD-56006: On providing an incorrect hostname/service ID in the Ozone URI, the filesystem client instead of failing, retries till exhaustion and the default retry is too high.

This issue is resolved.

Bug Fixes:

- HDDS-8495. Fix permissions and path handling for ozone sh token get command
- HDDS-8320. Fix ranger jackson version conflict
- HDDS-8173. Fix to remove entries from RocksDB after container gets deleted.
- HDDS-8068. Fix Exception: JMXServlet, getting attribute RatisRoles of Hadoop:service=OzoneManager.
- HDDS-8054. Fix NPE in metrics for failed volume
- HDDS-8221. Fix misplaced OzoneKey javadoc comment
- HDDS-8217. Fix `OmBucketInfo#equals` for comparing `defaultReplicationConfig`
- HDDS-8183. Fix edge case where delimiter is empty string in `BucketEndpoint#get`

- HDDS-8091. [addendum] Generate list of config tags from ConfigTag enum - Hadoop 3.1 compatibility fix
- HDDS-7721. Make OM Ratis roles available in /prom endpoint (fix compile error)
- HDDS-7596. Fix OM crash due to a corner case for FSO-enabled bucket
- HDDS-7871. Fix false positive in KeyManagerImpl#createFakeDirIfShould()
- HDDS-7372. Fix missing jars in classpath by specifying jar versions
- HDDS-7792. Fix package name typos in o.a.h.hdds.security.x509
- HDDS-7705. Fix OM Bootstrap request
- HDDS-7576. Prometheus metrics do not remove stale metrics until restart
- HDDS-7644. S3 multipart upload does not update quota namespace for missing parents
- HDDS-7641. Namespace quota validation is not present in multiple places
- HDDS-7696. MisReplicationHandler does not consider QUASI_CLOSED replicas as sources
- HDDS-7633. Compile error with Java 11: package com.sun.jmx.mbeanserver is not visible
- HDDS-7652. Volume Quota not enforced during write when bucket quota is not set
- HDDS-7801. Bucket not found when calling getKeyInfo with tenant context
- HDDS-7859. ICR processing does not remove container reference in NodeManager for a deleted replica
- HDDS-7838. gRPC channel created block input/output stream not shutdown properly
- HDDS-7969. CacheValue should not store value as an Optional.
- HDDS-7991. Do not return fake parent dir for deleted keys
- HDDS-8009. OM HA metrics should be unregistered if leader is not known
- HDDS-8029. [hsync] OutputStream in encrypted buckets do not return the correct stream capabilities.
- HDDS-8070. DBCheckpointMetrics is not unregistered during OM stop.
- HDDS-6176. Ozone service WebUI is not accessible with 404 error.
- HDDS-8108. gRPC channel created for replication client not shutdown properly
- HDDS-7930. [addendum] input stream does not refresh expired block token.
- HDDS-7930. input stream does not refresh expired block token.
- HDDS-8139. Datanodes should not drop block delete transactions based on transaction ID
- HDDS-8275. ManagedWriteBatch is not closed properly in SCMHADBTransactionBuffer
- HDDS-8258. RocksIterator not closed properly in KeyValueHandler.logBlocksIfNonZero
- HDDS-8257. RocksIterator not closed properly in LegacyBucketHandler
- HDDS-8242. Rename operation not working with FSO bucket destination
- HDDS-8496. S3 to return not found for object head/set when keyinfo indicates a directory.
- HDDS-8324. DN data cache gets removed randomly asking for data from disk.
- HDDS-7842. java.lang.NullPointerException:
hadoop.ozone.container.keyvalue.statemachine.background.StaleRecoveringContainerScrubbingService.
- HDDS-8045. Dependency convergence error for zookeeper.
- HDDS-7754. Download of container is failing with SSL/TLS error during re-replication
- HDDS-7455. ClassCastException: OzoneTokenIdentifier cannot be cast to String
- HDDS-8440. Ozone Manager crashed with ClassCastException when deleting FSO bucket.
- HDDS-8195. RDBStore.getUpdatesSince() throws RocksDBException: Requested array size exceeds VM limit.
- HDDS-8220. [Ozone-Streaming] Trigger volume check on IOException in StreamDataChannelBase
- HDDS-8057. Handle unchecked exception in KeyValueHandler more gracefully.
- HDDS-8019. IllegalStateException: call already closed at GrpcXceiverService.onCompleted.
- HDDS-7926. [hsync] Recon throws ClassCastException.
- HDDS-7714. Docker cluster ozone-om-ha fails during docker-compose up.

Upgrades:

- HDDS-8512. Bump Spring Framework to 5.3.27
- HDDS-8461. Bump jetty to 9.4.51.v20230217
- HDDS-8394. Upgrade Spring Framework to 5.3.26 due to CVE-2023-20861
- HDDS-8288. Upgrade moment.js to 2.29.4

- HDDS-7839. Upgrade Weld to 3.1.9
- HDDS-8263. [JDK17] Bump guice to 5.1.0, maven-shade-plugin to 3.4.1, remove guice-multibindings
- HDDS-8086. Bump snakeyaml from 1.33 to 2.0
- HDDS-7850. Bump maven-enforcer-plugin version to 3.2.1
- HDDS-7718. Bump Netty to 4.1.86 and gRPC to 1.51.1
- HDDS-8022. Add okio as dependency to the root pom
- HDDS-8022. Add okio as a dependency to the root pom.

Container Balancer:

- HDDS-8167. Inject MoveManager into ContainerBalancer.
- HDDS-8153. Integrate ContainerBalancer with MoveManager
- HDDS-8169. Delay Starting ContainerBalancer after SCM failover.

Erasur Coding improvements:

- HDDS-7122. Add validation for EC chunk size
- HDDS-7825. Warn when EC write exception occurs
- HDDS-7787. GetChecksum for EC files can fail intermittently with IndexOutOfBoundsException exception
- HDDS-7918. EC: ECBlockReconstructedStripeInputStream should check for spare replicas before failing an index
- HDDS-8216. EC: OzoneClientConfig is overwritten in ECKeyOutputStream
- HDDS-6572. EC: ReplicationManager - add move manager for container move
- HDDS-7931. EC: ManagedChannelImpl not cleaned up properly
- HDDS-7919. EC: ECPipelineProvider.createForRead should filter out dead replicas and sort replicas
- HDDS-7928. EC: Change ContainerReplicaPendingOps to store deadline rather than scheduled time
- HDDS-7923. [EC] Reconstruction is failing with IndexOutOfBoundsException
- HDDS-7917. EC: ECBlockInputStream should try spare replicas on error
- HDDS-7844. EC: Add normal and low priority to replication supervisor and commands
- HDDS-7841. EC: Remove ECReconstructionSupervisor and send reconstruction commands to ReplicationSupervisor
- HDDS-7833. EC: Refactor ReplicationSupervisor to allow Replication and ECReconstruction tasks
- HDDS-7761. EC: ReplicationManager - Use placementPolicy.replicasToRemoveToFixOverreplication in EC Over replication handler
- HDDS-7775. EC: Exception encountered while deleting UNHEALTHY replica in Datanode
- HDDS-7695. EC metrics related to replication commands don't add up
- HDDS-7729. EC: ECContainerReplicaCount should handle pending delete of unhealthy replicas
- HDDS-7727. EC: SCM unregistered event handler for DatanodeCommandCountUpdated
- HDDS-7666. EC: Unrecoverable EC containers with some remaining replicas may block decommissioning
- HDDS-7683. EC: ReplicationManager - UnderRep maintenance handler should not request nodes if none needed
- HDDS-7654. EC: ReplicationManager - merge mis-rep queue into under replicated queue
- HDDS-8287. Add toString() method to ECContainerReplicaCount
- HDDS-8172. ECUnderReplicationHandler should consider commands already sent when processing the container
- HDDS-8112. ECReconstructionCoordinator is not closed
- HDDS-8075. ECReconstructionCoordinatorTask.runTask should catch Exception
- HDDS-7649. S3 multipart upload EC release space quota wrong for old version

Test improvements:

- HDDS-8374. Fix flaky TestContainerStateCounts in Recon
- HDDS-8274. Intermittent timeout in acceptance-MR test setup
- HDDS-8219. Run HA secure tests from s3g container
- HDDS-7874. Disable flaky unit test: TestHddsSecureDatanodeInit.testCertificateRotationRecoverableFailure
- HDDS-7975. Rebalance acceptance tests

- HDDS-8152. Reduce S3 acceptance test setup time
- HDDS-8205. Reorder OM nodes in HA acceptance tests
- HDDS-7958. Ozone client not closed in integration tests
- HDDS-8144. TestDefaultCertificateClient#testTimeBeforeExpiryGracePeriod fails as we approach DST.
- HDDS-8150. RpcClientTest and ConfigurationSourceTest not run due to naming convention
- HDDS-8087. Intermittent crash in TestHddsDatanodeService
- HDDS-8024. Mark TestHSync#testOfsHSync as flaky
- HDDS-8035. Mark TestHSync#testO3fsHSync as flaky
- HDDS-8035. Mark TestOzoneManagerHAWithData#testOMHAMetrics as flaky
- HDDS-8012. Clean up after link loop test
- HDDS-7983. Intermittent OutOfMemoryError in TestOzoneRpcClientWithRatis#testUploadWithStreamAndMemoryMappedBuffer.
- HDDS-7588. Intermittent failure in TestObjectStoreWithLegacyFS#testFlatKeyStructureWithOBS
- Revert "HDDS-7588. Intermittent failure in TestObjectStoreWithLegacyFS#testFlatKeyStructureWithOBS"
- HDDS-7988. Run S3 tests with HA Proxy
- HDDS-7617. Remove flaky tag from TestECContainerRecovery
- HDDS-7617. Intermittent timeout in TestECContainerRecovery
- HDDS-7921. Migrate TestKeyManagerImpl to JUnit5
- HDDS-7617. Mark TestECContainerRecovery as flaky
- HDDS-3265. Intermittent timeout in TestRatisPipelineLeader
- HDDS-7864. Add integration test for replication
- HDDS-7881. Intermittent timeout in basic acceptance test
- HDDS-7886. Mark parts of TestRatisPipelineCreateAndDestroy as flaky
- HDDS-7868. Intermittent failure in TestOmAcls
- HDDS-7879. Intermittent BindException in HA integration tests
- HDDS-5626. Mark parts of TestAddRemoveOzoneManager, TestBlockOutputStreamFlushDelay and TestFailureHandlingByClient as flaky
- HDDS-7874. Disable flaky unit test: TestHddsSecureDatanodeInit.testCertificateRotation
- HDDS-7856. Fix timeout in TestPushReplicator
- HDDS-7856. Disable flaky TestPushReplicator until fixed
- HDDS-7806. Add unit tests for push replication
- HDDS-7808. Intermittent failure in TestReplicationManager#testUnderReplicationQueuePopulated
- HDDS-7719. [HTTPFSGW] Fix secure integration tests for the HttpFS module
- HDDS-7628. Intermittent failure in TestOzoneContainerWithTLS
- HDDS-7588. Intermittent failure in TestObjectStoreWithLegacyFS#testFlatKeyStructureWithOBS
- HDDS-5698 [HTTPFSGW] Port HTTPFS node and robot tests to ozone-ha, and ozonesecure(-ha)
- HDDS-5615 Add a simple test suite for HTTPFS GW.
- HDDS-8244. Selective checks: handle change in junit.sh

Logging improvements:

- HDDS-8148. Improve log for Pipeline creation failure
- HDDS-8245. Info log for keyDeletingService when nonzero number of keys are deleted.
- HDDS-7867. Clean up replication logs
- HDDS-8131. Add Configuration for OM Ratis Log Purge Tuning Parameters.
- HDDS-7869. Log configuration on component startup.
- HDDS-7959. Improve log in ECBlockInputStream and ECBlockReconstructedStripeInputStream
- HDDS-8037. Improve logging in EC Reconstruction putBlock precondition check
- HDDS-7857. Apply custom RocksDB log configuration for all schema versions
- HDDS-7082. Delete out of date audit logs
- HDDS-7097. Container scanner log output lacks useful information

- HDDS-7753. Simplify DatanodeDetails#toString to improve log messages
- HDDS-7726. EC: Enhance datanode reconstruction log message
- HDDS-7739. EC: Increase the information in the RM sending command log message
- HDDS-7716. Log read requests rejected with permission denied in OM audit
- HDDS-7631. Log format error on quotas when exceeding the space quota

HTTPFS protocol support:

- HDDS-7874. Disable Httpfs path check & mark corresponding APIs unsupported
- Revert "HDDS-5447. HttpFS support in Ozone"
- HDDS-5447. HttpFS support in Ozone
- HDDS-8088. [HTTPFSGW] Conform with new checkstyle rule in master #4354
- HDDS-8023. [HTTPFSGW] Remove support for satisfyStoragePolicy operation
- Change ozone-filessystem dependency scope to runtime in the HttpFS module
- HDDS-7737. [HTTPFSGW] Clean up dependencies
- HDDS-5828 [HTTPFSGW] Add proper handling for unsupported operations
- HDDS-6027 [HTTPFSGW] Fix dependency issues after master merge
- Revert "Quick fix for HTTPFS gateway dependencies after merge... to be revisited."
- Quick fix for HTTPFS gateway dependencies after merge... to be revisited.
- HDDS-5827 [HTTPFSGW] Remove non-server side related code from Ozone.
- HDDS-5829 [HTTPFSGW] Move to org.apache.ozone package from org.apache.hadoop.
- HDDS-5722 [HTTPFSGW] junit.jar and json-simple in jar report
- HDDS-5826 [HTTPFSGW] Remove or replace Hadoop shaded guava dependencies.
- Revert "[HTTPFSGW] HDDS-5695 Review ZK and Curator dependencies, and get rid of them."
- [HTTPFSGW] HDDS-5695 Review ZK and Curator dependencies, and get rid of them.
- HDDS-5520 Add lifecycle management to HttpFS server
- HDDS-5519 Remove unnecessary hadoop dependencies from httpfs module
- HDDS-5448 Copy over HttpFS module from Hadoop 3.3.1 to Ozone

Recon:

- HDDS-8127. Exclude deleted containers from Recon container count.
- HDDS-5331. Recon: Trigger PipelineSyncTask when DN becomes stale and ContainerHealthTask when DN becomes dead.
- HDDS-6056. Recon /containers endpoint should return SCM container data instead of OM container data.
- HDDS-7505. Recon: Rename OM DB Sync to DB Sync on UI.
- HDDS-7634. Recon: Show Datanode UUID on Pipeline page
- HDDS-7770. Recon namespace summary endpoint to carry basic entity information as well
- HDDS-7854. Recon NPE when trace enabled
- HDDS-4539. Container Health Task should not run until Recon has reached steady state.
- HDDS-8272. DBStore not closed properly in ReconStorageContainerManagerFacade
- HDDS-3486. Recon cannot track missing containers that were created and went missing while it is down.

Documentation:

- HDDS-8016. Updated the ozone doc for linked bucket and deletion async limitation.
- HDDS-7540. Document separate scheduled CI.
- HDDS-7409 [doc] Update documents for better presentation.
- HDDS-7684. Embed Matomo Web Analytics tracking code in docs.
- HDDS-7667. [Admin Doc] Observability: Add Grafana integration
- HDDS-7747. Document jq filtering examples for CLI responses.
- HDDS-7774. Update outdated Trash documentation
- HDDS-7556. Translate EC doc into Chinese
- HDDS-5966. [HTTPFSGW] Update module doc, and place it in Ozone project docs

CLI commands:

- HDDS-8042. Display certificate issuer in cert list command.
- HDDS-8168. Make deadlines inside MoveManager for move commands configurable
- HDDS-8233. ReplicationManager: Throttle delete container commands from over-replication handlers
- HDDS-8187. ReplicationManager: Datanode commands should be sent to nodeManager directly
- HDDS-6802. Implement JSON output for pipeline list CLI.
- HDDS-8241. Transfer leader command doesn't work in secure cluster
- HDDS-8158. Replication Manager: Make all handlers send commands immediately instead of returning commands.
- HDDS-8133. Create ozone sh key checksum command.
- HDDS-534. Remove unsupported jmxget subcommand.
- HDDS-7783. Add compile platform info in 'ozone version' command output
- HDDS-7137. Add CLI for Getting the failed deleted block txn.
- HDDS-3591. ozone.administrators support dynamic changes through the cli.
- HDDS-8039. Allow container inspector to run from ozone debug.
- HDDS-7489. Create a Freon tool to simulate datanodes

Security:

- HDDS-5043: Consistently use user's short name in secure mode
- HDDS-7697: Restrict change of bucket properties to owner and admins in NativeACL
- HDDS-7379. Use certificate bundles instead of the sole certificate
- HDDS-7399. Enable specifying external root ca
- HDDS-7398. Tool to remove old certs from the SCM db
- HDDS-8549. Restore client-side validation of bucket names
- HDDS-7573. Use keyManager and trustManager provided by keyStoreFactory in Ratis group.
- HDDS-8204. Add testuser principals for all Ozone containers.
- HDDS-8041. Let Ozone Client fail faster with wrong OM address in URI.
- HDDS-8145. ReadReplicas should close client.
- HDDS-8163. Use try-with-resources to ensure close rockdb connection in SstFilteringService.
- HDDS-8151. Support fine-grained lifetime for root CA certificate.
- HDDS-8030. Cleanup unused/unnecessary code related to CertificateClient.
- HDDS-8095. Unbuffer not supported in TDE buckets.
- HDDS-7590. Use keyManager and trustManager provided by keyStoreFactory in om grpc services.
- HDDS-8020. File checksum helper leaking client
- HDDS-7525. Migrate key digest from MD5 to SHA256 in Ozone shell.
- HDDS-7708. No check for certificate duration config scenarios.
- HDDS-7339. Implement Certificate renewal task for services.

Performance:

- HDDS-8289: Speed up FSO ListKeys, skip skipToFirst
- HDDS-8147: Introduce latency metrics for S3 Gateway operations
- HDDS-8270: Measure checkAccess latency for Ozone objects
- HDDS-8329: Performance Impact during container close when still Client is writing
- HDDS-8154. Perf: Reuse Mac instances in S3 token validation
- HDDS-8367. setquota should have a check on usedNamespace [NameSpacequota]
- HDDS-8285. Eliminate leftover Guava Optional from CacheValue.
- HDDS-8128. Deduplicate the ops in RDBBatchOperation.
- HDDS-8255. Remove unnecessary sleep at secure container startup
- HDDS-8076. Use container cache in Key listing API.
- HDDS-8074. Improve synchronization around command queue updates in Node Manager.

- HDDS-8129. ContainerStateMachine allows two different tasks with the same container id running in parallel.
- HDDS-8079. RocksDB included in client libs doubles shaded jar size.
- HDDS-7524. Compaction DAG node pruning.
- HDDS-7623. Do not compress container re-replication traffic by default.
- HDDS-7635. Update failure metrics when allocate block fails in preExecute.
- HDDS-7778. Add metrics for push replication

Stability/Reliability:

- HDDS-8424. Restore legacy bucket behavior for directory objects.
- HDDS-8292. Inconsistent key name handling for FSO bucket files.
- HDDS-8383. Misreplication cannot be resolved with single rack
- HDDS-7994. Expose information via the ClusterState endpoint on keys and Directories marked for deletion.
- HDDS-8253. Default Metadir in Current Working Directory.
- HDDS-8230. Let ReplicationManager decide the timeout for commands in Datanodes
- HDDS-8223. SCM delete block service should run wait for safemode to exit.
- HDDS-8155. SCM Block deleting service add transaction blocks for stale/dead/decommissioning DNs
- HDDS-7853. Add support for RemoveSCM in SCMRatisServer.
- HDDS-8090. When getBlock from a datanode fails, retry other datanodes.
- HDDS-6449. Failed container delete can leave artifacts on disk.
- HDDS-8171. Replicate commands could be sent to dead maintenance nodes if the same index is being decommissioned.
- HDDS-8111. ReplicationManager: Add RatisMisReplicationHandler into rm.processUnderReplicatedContainer.
- HDDS-8142. Check if no entries in Block DB for a container on container delete.
- HDDS-8119. Remove loosely related AutoCloseable from SendContainerOutputStream.
- HDDS-8118. Fail container delete on non-empty chunks dir.
- HDDS-8028. JNI for RocksDB SST Dump tool.
- HDDS-8024. When readChunk from a datanode fails, retry other datanodes.
- HDDS-7183. Expose RocksDB critical properties.
- HDDS-8036. Unprotected flush in SCMHADBTransactionBuffer.
- HDDS-6241. Follower SCM node repeatedly sending requests to Ratis server.
- HDDS-7156. Reset pending delete block count
- HDDS-7989. UnhealthyReplicationProcessor retries failure without delay.
- HDDS-7998. Synchronize on containerInfo in ReplicationManager and MoveManager.
- HDDS-8008. Move pendingOps into ContainerStateManagerImpl to ensure consistent state.
- HDDS-7782. OM lease recovery for hsync'ed files.
- HDDS-7463. SCM Pipeline scrubber never able to clean up allocated pipeline.
- HDDS-7621. Update SCM term in datanode from heartbeat without any commands.
- HDDS-7661. Ratis Misreplication Handler.
- HDDS-7584. Addition of new OM node expels itself from the Ratis ring after restart.
- HDDS-7473. Ratis integration for support of remove registration
- HDDS-7847. Handle Replication of Unhealthy Replicas in RM
- HDDS-7924. Remove non-conflicting transitive-only dependency versions
- HDDS-7915. Force close QUASI_CLOSED replicas only when the container is CLOSED in Legacy RM
- HDDS-7560. Placement Policy Interface changes to handle Overreplication
- HDDS-7738. SCM terminates when adding container to a closed pipeline

Usability:

- HDDS-8484. Allow Ozone volume name to have special characters
- HDDS-7586. Allow user to create bucket with non-s3-naming-convention
- HDDS-7495. Create OBS buckets by default from S3 API
- HDDS-8302. Add a flag to disable hsync by default

- HDDS-6064. Print proper JSON from ozone debug ldb scan.
- HDDS-8281. Add IntelliJ code style scheme import recommendation in CONTRIBUTING.md.
- HDDS-8222. EndpointBase#getBucket should handle BUCKET_NOT_FOUND
- HDDS-8098. Make Schema V3 as default in ozone debug ldb for container db.
- HDDS-8240. Disable JaCoCo for PRs and in forks
- HDDS-8065. Provide GNU long options.
- HDDS-8149. Refactor the way to notify keyStoreFactory about certificate renewed.
- HDDS-8099. Make unlimited length the default in ozone debug ldb.
- HDDS-8110. ReplicationManager: Introduce basic limits on ReplicateContainer commands.
- HDDS-8091. Generate list of config tags from ConfigTag enum.
- HDDS-7816. Add DataNode list to the SCM WebUI.
- HDDS-8073. Replace Usages of LegacyReplicationManager.MoveResult with MoveManager.MoveResult.
- HDDS-8078. Github CI: Allow PR title starting with Revert.
- HDDS-8046. Eliminate embedded newlines from config properties.
- HDDS-7210. Missing open containers show up as "Closing" on the container report.
- HDDS-7934. NPE in RandomKeyGenerator's shutdown hook.
- HDDS-8044. Ensure GrpcOutputStream is closed.
- HDDS-8032. SCM support reconfigurable dynamically.
- HDDS-8077. Enforce new checkstyle: NewlineAtEndOfFile
- HDDS-7301. Cleanup OmUtils.
- HDDS-7971. Support SCM transfer Ratis leadership.
- HDDS-7064. S3 get-object response emits tracing spans outside ObjectEndpoint#get.
- HDDS-7710. Support AWS s3 ListObjects API's encodingType request parameter.
- HDDS-7721. Make OM Ratis roles available in /prom endpoint.
- HDDS-8026. Replace import from shaded Guava.
- HDDS-7890. Refactor ContainerDeletionChoosingPolicy implementations.
- HDDS-7982. NPE in BlockInputStream due to null pipeline after refresh.
- HDDS-8025. ReplicationManager: Count a container once for missing, under, mis, or over replicated.
- HDDS-6743. Specify leader node for OM failover
- HDDS-7973. Let RatisMisReplicationHandler use the new RatisContainerReplicaCount constructor
- HDDS-7979. Replace import from shaded Guava
- HDDS-6650. S3MultipartUpload support update bucket usedNamespace.
- HDDS-7686. Cherry-pick proto.lock files change from ozone-1.3 release branch to master.

Known issues in Ozone parcel 718.2.0

You must be aware of the known issues and limitations, the areas of impact, and workaround in Ozone parcel.

Tez Configuration Changes

The following configuration changes have to be made to pick up the latest Ozone FS jar from the Ozone parcel (when installed):

CDPD-48540

For `tez.cluster.additional.classpath.prefix`, the value is `/var/lib/hadoop-hdfs/*` (Tez Additional Classpath)

For `tez.user.classpath.first`, the value is `true` (Tez Client Advanced Configuration Snippet (Safety Valve) for `tez-conf/tez-site.xml`)

CDPD-47605

For `tez.cluster.additional.classpath.prefix` the value is `/var/lib/hadoop-hdfs/*` (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

For `tez.user.classpath.first`, the value is `true` (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

Restart the TEZ and HIVE_ON_TEZ service as prompted by Cloudera Manager.

Update Yarn to use updated Ozone FS Jar

CDPD-48500: Ozone parcel activation or install should handle redeployment of YARN jars and clean-up cache.

Perform the following steps:

1. Log in to Cloudera Manager UI
2. Navigate to Clusters
3. Select the YARN service
4. Click Actions
5. Click Install YARN Service Dependencies
6. Click YARN MapReduce Framework JARs
7. Restart the CDP 7.1.8 cluster

Other issues

SSL Handshake fails between Ozone DataNodes if the two DataNodes have their certificate signed by different Ozone Storage Container Managers.

Ozone DataNode certificates are signed by the leader Storage Container Manager. Due to an issue in creating a TrustStore for DataNode to DataNode connections, the trust cannot be established between the two DataNodes if a different Storage Container Manager signs their certificate. These connections fail to establish and display an SSL Handshake Exception. This affects Pipeline creation and container replication (also EC container reconstruction). The symptoms vary, depending on the number of the nodes that have different signer certificates, either these DataNodes do not participate in any Ratis-3 Pipeline, or have Pipelines exclusively in between groups with the same signer. Over time this can lead to an imbalance in DataNode usage, and it might cause decommission of a DataNode stuck if the data has to be replicated to a node with a certificate that has a different signer.

This problem affects all the 7.1.8 Ozone Parcel releases.

To identify if the problem is present on a cluster, the output of `ozone admin cert list` command must be examined. Ensure you define a sufficient number of certificates to be returned with the `-c` option to see all the certificates issued in the system.

If there are different Issuers for the latest DataNode certificates, this indicates the cluster is affected.

It is possible to avoid the problem by checking the `ozone admin scm roles` output and see if the Primordial node is the actual leader node before adding a new DataNode and starting it for the first time. If the leader SCM node is a different node, then calling `ozone admin scm transfer` to make the Primordial node the leader can put the cluster into the desired state before adding the new DataNode.

Ensure that all the certificates in the cluster are signed by the same Storage Container Manager node.

Perform the procedure below if the cluster is affected and has DataNode certificates signed by a different issuer:

1. Identify the current leader Storage Container Manager by running the command `ozone admin scm roles`.
2. If the leader did not sign the majority of DataNode certificates, change the leader that has signed the majority of DN certificates by running the `ozone admin scm transfer` command with the proper Ozone SCM Service ID (set in Cloudera Manager>Ozone>Configuration) and the UUID of the desired leader node.

3. Stop the minority of DataNodes that have a different sign than the majority.
4. Locate the Datanode Metadata Directory (set in Cloudera Manager>Ozone>Configuration) on the hosts of the stopped DataNode and move the directory to a backup location.
5. Start the previously stopped DataNodes.
6. After regenerating the certificates, check if the DataNodes are joining the Ratis-3 Pipeline.

To avoid service disruptions, you can stop the DataNodes one by one instead of all together. This can cause a data outage for the downtime of the node if there are files with Ratis-1 replication on the cluster and the single replica of these files reside on the node being restarted. Changing the leader node on an SCM HA-enabled cluster should not disrupt operations.

CDPD-56006: On providing an incorrect hostname/service ID in the Ozone URI, the filesystem client instead of failing, retries till exhaustion and the default retry is too high.

Configure `ozone.client.failover.max.attempts` to a lower the value to avoid long endless retries.

CDPD-49137: Sometimes OM's kerberos token is not updated and it stops being able to communicate with SCM. When this occurs, writes will start to fail.

Restarting OM or setting the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true` will fix the issue.

CDPD-49808: Spark jobs against Ozone intermittently fail with ERROR spark.SparkContext: [main]: Error initializing SparkContext.java.lang.IllegalStateException: No filter named.

This is an intermittent failure which can be retried.

CDPD-50678: Deleting containers which have one or more replicas which are not empty on the Datanode can cause the container to be stuck in a deleting state indefinitely. Containers in this state can also block decommission or maintenance operations completing.

None.

Technical Service Bulletins

TSB 2023-697: Potential Ozone data loss in Ozone Parcel 2 caused by hsync/hflush

A recently identified regression caused by an issue in the implementation of hsync/hflush Application Programming Interfaces (APIs) can result in data loss in Apache Ozone (Ozone) if the client application calls hsync/hflush before closing the file.

Once the file is closed, the blocks of the file will be removed by the Storage Container Manager (SCM) after the container with these blocks gets closed. Reading the file will fail with the error message similar to the following: `org.apache.hadoop.hdds.scm.container.common.helpers.StorageContainerException: Unable to find the block. BlockID : conID: 4 locID: 111677748019201149 bc sId: 382`

To avoid potential data loss caused by the issue described above, Users on the affected parcel versions must upgrade to OZONE-718.2.2 immediately.

Upstream Jira: [HDDS-9146](#)(fix)

Knowledge article: [TSB 2023-697: Potential Ozone data loss in Ozone Parcel 2 caused by hsync/hflush](#)

OZONE 718.1.0

Ozone 718.1.0 parcel must be installed on the CDP 7.1.8 cluster. You must be aware of the additional functionalities and improvements to features of components in Ozone 718.1.0 parcel. Learn how the new features and improvements benefit you.

- Erasure Coding Improvements
- Replication Manager Optimizations
- Decommissioning Improvements
- Refer to [Fixed issues](#) section for the list of issues fixed.

**Attention:**

1. If you are upgrading your cluster from Cloudera Runtime 7.1.8 with Ozone 718.1.0 parcel to Cloudera Runtime 7.1.9, you must deactivate and undistribute the Ozone 718.1.0 parcel and then proceed with the upgrade to 7.1.9
2. Using Cloudera Manager 7.7.1, if you are upgrading from [CDP Private Cloud Base 7.1.8 with Ozone parcels](#) to CDP Private Cloud Base 7.1.9 using Cloudera Manager 7.11.3, then you must
 - a. If you have [Ozone 718.1.0 parcel](#) on the CDP Private Cloud Base 7.1.8 cluster, you must deactivate and undistribute Ozone 718.1.0 parcel.
 - b. Upgrade Cloudera Manager from 7.7.1 to 7.11.3. Do not restart the CDP cluster.
 - c. Upgrade from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9.

Requirements

Review the list of requirements before proceeding with installation.

- As a prevalidation for installation of the Ozone parcels, ensure that /var/lib/hadoop-hdfs is present on all the nodes and owned by the user hdfs.
- CDP 7.1.8 cluster with CHF 3 or greater installed and corresponding Cloudera Manager release.
- Ozone 718.1.0 parcel is certified to work with CHF3. However, parcels are expected to be not sensitive to CHF releases.
- Ozone Parcels are certified to work with Cloudera Manager version 7.7.1.
- The table below explains the product support matrix for Private Cloud Data Services, Cloudera Runtime, and Ozone parcel:

Product combination	Supportability
Private Cloud Data Services 1.5.0 with Cloudera Runtime 7.1.8	Supported
Ozone 718.1.0 parcel on Cloudera Runtime 7.1.8	Supported
Private Cloud Data Services 1.5.0 with Ozone 718.1.0 parcel	Not supported



Note: If you upgraded the cluster from CDP 7.1.7, then you must finalize the upgrade. Ozone parcel releases are not certified for rollback to CDP 7.1.7.

Downloading the Ozone parcels

This section helps you to download the Ozone parcel.

To download the Ozone parcel, use this link: <https://archive.cloudera.com/p/ozone/718.1.0/>

Installing Ozone 718.1.0 parcel

To install the Ozone 718.1.0 parcel, perform the following procedure:

Before you begin

As a prevalidation for installation of the Ozone parcels, ensure that /var/lib/hadoop-hdfs is present on all the nodes and owned by the user hdfs.

Procedure

1. Log in to Cloudera Manager
2. On the left navigation pane, click Parcels
3. Click Parcel Repositories & Network Settings

4. In the Remote Parcel Repository URLs section, click the "+" icon and add the *Ozone 718.1.0 URL* for your Parcel repository. For the Ozone parcel URL, refer to the [Downloading the Ozone parcel](#) section.
5. Click Save & Verify Configuration. A message with the status of the verification appears above the Remote Parcel Repository URLs section. If the URL is not valid, check the URL and enter the correct URL.
6. After the URL is verified, click Close.
7. Locate the row in the table that contains the new Ozone parcel and click the Download button.
8. After the download of the new Ozone parcel is complete, click the Distribute button.
9. Wait for the parcel to be distributed and unpacked before continuing. Cloudera Manager displays the status of the Cloudera Runtime parcel distribution. Click on the status display to view detailed status for each host.
10. Prepare the Ozone service for an upgrade before upgrading the Ozone software. Click Prepare for Upgrade under the Action drop-down.
11. Click Stop under the Action drop-down to stop the Ozone service
12. Click Activate on the Parcels page.
13. Click Finish Upgrade under the Action drop-down. Continue with next step.
14. Update Tez Configuration:
 - a) For `tez.cluster.additional.classpath.prefix`, the value is `/var/lib/hadoop-hdfs/*` (Tez Additional Classpath)
 - b) For `tez.user.classpath.first`, the value is `true` (Tez Client Advanced Configuration Snippet (Safety Valve) for `tez-conf/tez-site.xml`)
15. Update YARN to use the new Ozone Client Jars. Ozone parcel activation or install should handle redeployment of YARN jars and clean-up cache. Perform the following steps:
 - a) Log in to Cloudera Manager UI.
 - b) Navigate to Clusters
 - c) Select the YARN service
 - d) Click Actions
 - e) Click Install YARN Service Dependencies
 - f) Click Install YARN MapReduce Framework JARs
16. Update Hive On Tez configuration:
 - a) For `tez.cluster.additional.classpath.prefix` the value is `/var/lib/hadoop-hdfs/*` (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)
 - b) For `tez.user.classpath.first`, the value is `true` (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)
17. Restart the CDP 7.1.8 cluster: Click the Actions menu and click Restart.
18. Click Finalize Upgrade under the Action drop-down.

Release notes

This document helps you to install and manage Ozone 718.1.0 parcel on top of CDP Private Cloud Base 7.1.8.

Fixed issues in Ozone 718.1.0 parcel

You can review the list of reported issues and their fixes in Ozone 718.1.0 parcel. Fixed issues represent selected issues that were previously logged through Cloudera Support, but are now addressed in the current Ozone parcel release. These issues may have been reported in previous versions of Runtime as a known issue; meaning they were reported by customers or identified by Cloudera Quality Engineering teams.

- TSB (2024-748) - Potential Ozone data loss in CDP Private Cloud Base 7.1.8 caused by container delete
- CDPD-50524: HDDS-8142: Check if no entries in Block DB for a container on container delete
- CDPD-50524: HDDS-8118: Fail container delete on non empty chunks dir
- HDDS-8108: gRPC channel created for replication client not shutdown properly
- HDDS-8129: ContainerStateMachine allows two different tasks with the same container id running in parallel
- HDDS-8019: IllegalStateException: call already closed at GrpcXceiverService.onCompleted
- CDPD-49916: HDDS-7156: Reset pending delete block count

- CDPD-48932: HDDS-8044: Ensure GrpcOutputStream is closed
- CDPD-49916: HDDS-7091: Filter containers with pending deletion blocks before checking policy
- CDPD-49916: HDDS-7126: Clean deletedBlock records of deleted containers
- CDPD-49916: HDDS-7306: DeleteBlock retry count increased at incorrect place
- CDPD-49918: HDDS-8037: Improve logging in EC Reconstruction putBlock precondition check
- HDDS-8020: File checksum helper leaking client
- HDDS-7931: EC: ManagedChannelImpl not cleaned up properly
- HDDS-7829: SCM to reject adding container to closed pipeline
- HDDS-7738: SCM terminates when adding container to a closed pipeline
- HDDS-7920: User cannot list their own volumes without access to root volume
- HDDS-7915: Force close QUASI_CLOSED replicas only when the container is CLOSED in Legacy RM
- HDDS-7923: [EC] Reconstruction is failing with IndexOutOfBoundsException
- HDDS-7838: gRPC channel created block input/output stream not shutdown properly
- HDDS-7455: ClassCastException: OzoneTokenIdentifier cannot be cast to String
- CDPD-49042: Cherry pick HDDS-7859
- CDPD-47991: Ozone - Upgrade Netty to 4.1.86.Final due to CVE-2022-41881, CVE-2022-41915
- HDDS-7801: Bucket not found when calling getKeyInfo with tenant context
- HDDS-7755: Ensure that acquired locks are always released
- HDDS-7787: GetChecksum for EC files can fail intermittently with IndexOutOfBoundsException
- HDDS-7761: EC: ReplicationManager - Use placementPolicy.replicasToRemoveToFixOverreplication in EC Over replication handler
- HDDS-7775: EC: Exception encountered while deleting UNHEALTHY replica in Datanode
- HDDS-6447: Refine SCM handling of unhealthy container replicas
- HDDS-7753: Simplify DatanodeDetails#toString to improve log messages
- HDDS-7560: Placement Policy Interface changes to handle Overreplication
- CDPD-48225: HDDS-7754: Download of container is failing with SSL/TLS error during re-replication
- HDDS-7726: EC: Enhance datanode reconstruction log message
- HDDS-7739: EC: Increase the information in the RM sending command log message
- HDDS-7695: EC metrics related to replication commands don't add up
- HDDS-7729: EC: ECContainerReplicaCount should handle pending delete of unhealthy replicas
- HDDS-7727: EC: SCM unregistered event handler for DatanodeCommandCountUpdated
- HDDS-7666: EC: Unrecoverable EC containers with some remaining replicas may block decommissioning
- HDDS-7696: MisReplicationHandler does not consider QUASI_CLOSED replicas as sources
- HDDS-7661: Ratis Misreplication Handler
- HDDS-7692: Remove usage of MonotonicClock and use default system clock instead
- HDDS-7463: SCM Pipeline scrubber never able to cleanup allocated pipeline
- HDDS-7683: EC: ReplicationManager - UnderRep maintenance handler should not request nodes if none needed
- HDDS-7635: Update failure metrics when allocate block fails in preExecute
- HDDS-7654 EC: ReplicationManager - merge mis-rep queue into under replicated queue
- HDDS-7621: Update SCM term in datanode from heartbeat without any commands
- HDDS-7605: Improve logging in Container Balancer
- HDDS-7616: EC: Refactor Unhealthy Replicated Processor
- HDDS-7426: Add a new acceptance test for Streaming Pipeline
- HDDS-7478: [Ozone-Streaming] NPE in when creating a file with o3fs
- HDDS-7425: Add documentation for the new Streaming Pipeline feature
- HDDS-7438: [Ozone-Streaming] Add a createStreamKey method to OzoneBucket
- HDDS-7431: [Ozone-Streaming] Disable data steam by default
- HDDS-6955: [Ozone-streaming] Add explicit stream flag in ozone shell
- HDDS-6867: [Ozone-Streaming] PutKeyHandler should not use streaming to put EC key
- HDDS-6842: [Ozone-Streaming] Reduce the number of watch requests in StreamCommitWatcher

- HDDS-6592: [Ozone-Streaming] Fix ContainerStateMachine#applyTransaction assert error
- HDDS-5666: Add option to createKey via streaming api in Freon
- HDDS-6500: [Ozone-Streaming] Buffer the PutBlockRequest at the end of the stream
- HDDS-6137: [Ozone-Streaming] Refactor KeyDataStreamOutput
- HDDS-5798: [Ozone-Streaming] Setup TlsConf parameters
- HDDS-6388: [Ozone-Streaming] Streaming write support both pipeline model and star model
- HDDS-6355: [Ozone-Streaming] Fix CheckStyle problem
- HDDS-6229: [Ozone-Streaming] Data Channel abstraction on datanode
- HDDS-6282: Fix BlockDataStreamOutput#doFlushIfNeeded NPE
- HDDS-5487: [Ozone-Streaming] BlockDataStreamOutput support FlushDelay
- HDDS-6298: Add XceiverServerRatis stream config
- HDDS-6138: [Ozone-Streaming] Define a limit on the size of the retry bufferList
- HDDS-6178: [Ozone-Streaming] Fix NPE in HDDS-6139
- HDDS-6139: [Ozone-Streaming] Fix incorrect computation of totalAckDataLength
- HDDS-6130: [Ozone-Streaming] When releaseBuffers will get "Couldn't find the required future"
- HDDS-6039: Define a minimum packet size during streaming writes
- HDDS-5851: [Ozone-Streaming] Define a PutBlock/maxBuffer fixed boundary for streaming writes
- HDDS-5743: [Ozone-Streaming] Add option to write files via streaming api in ofs and o3fs
- HDDS-5879: [Ozone-Streaming] OzoneBucket add the createMultipartStreamKey method
- HDDS-5961: [Ozone-Streaming] update the usage space of Containers in the stream write
- HDDS-5987: [Ozone-Streaming] Add XceiverClientRatis stream config
- HDDS-5763: Provide an Executor for each LocalStream in ContainerStateMachine
- HDDS-5895: [Ozone-Streaming] Make raft.server.data-stream.client.pool.size configurable
- HDDS-5674: [Ozone-Streaming] Handle client retries on exception
- HDDS-5849: [Ozone-Streaming] Write exceptions occur after checksum is enabled
- HDDS-5486: [Ozone-Streaming] Streaming supports writing in Pipeline mode
- HDDS-5742: Avoid unnecessary ByteBuffer conversions
- HDDS-5705: [Ozone-Streaming] Change ByteBufStreamOutput to ByteBufferStreamOutput
- HDDS-5599: [Ozone-Streaming] drop BufferPool and ChunkBuffer to avoid buffer copying
- HDDS-5488: [Ozone-Streaming] Add a new BlockOutputStream/KeyOutputStream to support streaming api
- HDDS-5480: [Ozone-Streaming] Client and server should support stream setup
- HDDS-5481: Fix stream() and link() method in ContainerStateMachine
- HDDS-5452: Add link method to ContainerStateMachine for Ratis streaming
- HDDS-5366: [Ozone-Streaming] Implement stream method to ContainerStateMachine
- HDDS-7371: Create properties for all dependency versions
- HDDS-7620: Check term for async commands before processing
- HDDS-7521: EC: Misreplication Handler changes for Placement Policy interface changes
- HDDS-7636: Remove hadoop security dependency in org.apache.hadoop.HDDS.security.ssl package
- HDDS-7628: Disable flaky TestOzoneContainerWithTLS#testCreateOzoneContainer
- HDDS-7642: EC: Add debug logging to the Replication Manager check handlers
- HDDS-7592: Recursive delete of directory multiple time impacts some directory not removed
- HDDS-7627: EC: Bug fix for calculating Misreplication Count
- HDDS-7643: rocksdb iterator is not closed properly
- HDDS-7503: CLI admin namespace -quota should return whole file system disk capacity consumed by all replicas
- HDDS-7582: ECUnderReplicationHandler does not consider pending adds when finding targets
- HDDS-7640: EC: UNHEALTHY replicas not replaced by healthy replicas from a CLOSED container by RM
- HDDS-7008: Mark TestOzoneHARatisLogParser as flaky
- HDDS-7619: Update SCM term in datanode when command is received
- HDDS-7591: Add a handler for over replicated Ratis containers to RM
- HDDS-7628: Disable flaky TestOzoneContainerWithTLS#testContainerDownload

- HDDS-7618: Replication Commands should timeout if not processed on datanodes in time
- HDDS-7595: Add OmId in OM's UI
- HDDS-7589: Unify multipart stream reading
- HDDS-7614: Add subscription mechanism to ContainerReplicaPendingOps
- HDDS-7580: Add option to show key count in DBScanner
- HDDS-7572: Use keyManager and trustManager provided by keyStoreFactory in datanode grpc services
- HDDS-7617: Mark testECContainerRecoveryWithTimedOutRecovery as flaky
- HDDS-7587: Update NOTICE and LICENSE in license/bin/
- HDDS-7603: Enable Erasure coding and encryption flag on FileStatus for ofs
- HDDS-7265: ScatterRackPolicy for Ratis pipeline provider container placement
- HDDS-7502: Recon: Show number of open containers on the Overview page
- HDDS-7615: Import Guava classes directly, not from hadoop-thirdparty
- HDDS-7492: Extend Placement Policy Interface to select mis-replicated replicas to copy
- HDDS-7597: Clarify SCM admin access denied message
- HDDS-7464: Container Report at SCM is not coming separately for ICR and FCR in prometheus endpoint
- HDDS-7550: Recon: UI sorting not functioning as expected
- HDDS-7599: Bump commons-net from 3.6 to 3.9.0
- HDDS-7526: Avoid overwriting replication config on existing bucket when quota is set
- HDDS-7472: EC: Fix NSSummaryEndpoint#getDiskUsage for EC keys
- HDDS-5704: Ozone URI syntax description in help content needs to mention about ozone service id
- HDDS-7555: Upgrade Ratis to 2.4.2-8b8bdda-SNAPSHOT
- HDDS-7541: FSO recursive delete directory with hierarchy takes much time for cleanup
- HDDS-7581: Fix update-jar-report for snapshot
- HDDS-7253: Fix exception when '/' in key name
- HDDS-7579: Use Netty 4.1.77 for consistency
- HDDS-7562: Suppress warning about long filenames in tar
- HDDS-7563: Add a handler for under replicated Ratis containers in RM
- HDDS-7497: Fix mkdir does not update bucket's usedNamespace
- HDDS-7567: Invalid entries in LICENSE
- HDDS-7575: Correct showing of RATIS-THREE icon in Recon UI
- HDDS-7540: Let reusable workflow inherit secrets
- HDDS-7568: Bump copyright year in NOTICE
- HDDS-7394: OM RPC FairCallQueue decay decision metrics list caller username in the metric
- HDDS-7510: Recon: Return number of open containers in `/clusterState` endpoint
- HDDS-7561: Improve setquota, clrquota CLI usage
- HDDS-6615: EC: Improve write performance by pipelining encode and flush
- HDDS-7554: Recon UI should show DORMANT in pipeline status filter
- HDDS-7540: Separate scheduled CI from push/PR workflows
- HDDS-7500: Ozone recon server Memory Optimization
- HDDS-7536: Remove unused dependency declarations
- HDDS-7486: Support KeyStoreFactory which supports keyManager and trustManager reload
- HDDS-7289: Bump protobuf-java from 3.19.2 to 3.19.6
- HDDS-7535: Eliminate duplicated config in LegacyReplicationManager
- HDDS-7449: Avoid overwriting bucket encryption properties when quota/replication config are set on an existing bucket
- HDDS-7538: Remove unused OmUtils#getMD5Digest
- HDDS-7534: Intermittent failure in TestKeyDeletingService#checkDeletionForEmptyKey
- HDDS-7353: Cleanup of unused code in SCM
- HDDS-7481: Set quota with reduced size than actual used do not fail
- HDDS-7187: EC: Retry failed writes before rewrite to a new block group

- HDDS-7532: EC: ReplicationManager - remove calls to ECHealthCheck from under and over replication processing
- HDDS-7528: EC: ReplicationManager - refactor logic to send datanode commands into a central place
- HDDS-7092: EC: Offline Recovery with simultaneous Over Replication and Under Replication
- HDDS-7462: EC: Fix Reconstruction Issue with StaleRecoveringContainerScrubbingService
- HDDS-7518: Fix intermittent failure in ozonesecure replication test
- HDDS-5526: ContainerBalancer#checkConditionsForBalancing pre-emptively checks iteration limits
- HDDS-7280: New test addition to TestOMBucketCreateRequestWithFSO
- HDDS-7476: SCM root CA stored under HDDS.datanode.dir if ozone.metadata.dirs is not specified
- HDDS-4565: [FSO]Delete : quotaReleased on directory deletion should consider all sub paths
- HDDS-7498: Add permission check when `--user` is specified in `ozone sh volume list`
- HDDS-7346: Cannot set bucket args when the bucket has quota set
- HDDS-7470: [Quota] bucket is created crossing quota makes OM crash and unable to start
- HDDS-7260: dfs -du should return correct disk space consumed by all replicas
- HDDS-7474: Recon: Add count to containers page tab label
- HDDS-7245: Recon: Add button to trigger OM DB Sync (Frontend)
- HDDS-4402: Recon dashboard page does not load until missing containers API returns data
- HDDS-7488: EC: ReplicationManager: Move Mis-Replicated into a separate unhealthy state
- HDDS-7487: Support fine grained certificate lifetime for efficient test
- HDDS-7095: allow on demand scanning for containers
- HDDS-6971: EC: Add EC block checksum computer
- HDDS-7101: EC: ReplicationManager - handle UNHEALTHY replicas
- HDDS-7255: Add metrics for container reports events
- HDDS-7490: fix metaTable error
- HDDS-7468: Recon: clusterState API Changes for adding missingContainerCount
- HDDS-7485: Bump commons-codec to 1.15
- HDDS-7475: Remove outdated Hadoop dependencies
- HDDS-7471: EC: Notify ReplicationManager when a heartbeat updates datanode command counts
- HDDS-7446: Cleanup possibly onDisk unreferenced blocks
- HDDS-7469: Recon: Add more information to Open Containers column on the Datanodes page
- HDDS-4385: Search function using container ID on Recon Missing Container
- HDDS-7408: Recon: Refactor container status page
- HDDS-7456: Recon: Refreshed at not working on Pipeline and DataNode page
- HDDS-7323: Recon: Auto refresh toggle is switched back when visiting new site
- HDDS-7360: Investigate TestRootedOzoneFileSystem test cases
- HDDS-7314: Recon graceful shutdown while closing RocksDB handle
- HDDS-7440: Delete package-info class for org.apache.hadoop.io.erasurecode.rawcoder to avoid conflict with Hadoop
- HDDS-7453: Check certificate expiration at service startup, renew if necessary
- HDDS-2642: Expose decommission / maintenance metrics via JMX
- HDDS-5112: Recon Insights page does not list buckets when only one volume is present
- HDDS-6896: EC: ReplicationManager - Add relevant metrics to the various ReplicationManager classes
- HDDS-7326: Intermittent timeout in TestECContainerRecovery.testContainerRecoveryOverReplicationProcessing
- HDDS-7451: Add UT of SyncDataFromOM for SequenceNumberNotFoundException
- Upgrade RocksDB to 7.7.3
- HDDS-7460: Bump snakeyaml from 1.32 to 1.33
- HDDS-3419: Throw exception with correct code when available data nodes are not sufficient
- HDDS-7362: Bring back the mechanism to share LRU cache between RocksDB in datanodes
- HDDS-7419: Integrate the GetKeyInfo API to OFS

- HDDS-7324: S3G: Avoid calling getS3VolumeContext by using GetGeyInfo API
- HDDS-5866: Discrepancy in Trash directory in ofs vs o3fs
- HDDS-7423: DB Scanner should allow iteration from specific key
- HDDS-6757: Test "size.leaving.source.max" limit in ContainerBalancer
- HDDS-7090: EC: delete empty closed EC container
- HDDS-7025: Add cache metrics in OM
- HDDS-7435: ContainerOperationClient cleanup
- HDDS-6926: Add support for shaded protobufs used by hadoop-client/spark
- HDDS-7447: Leak of TableIterator in DirectoryDeletingService
- HDDS-7320: Enable Container SchemaV3 by default
- HDDS-7429: Remove Tag from SCMMetadataStoreMetrics
- HDDS-7363: Changes to Ozone CLI to validate if the jars in classpath files are present on an install
- HDDS-7290: provide a config to increase the list batch size in OzoneFileSystem
- HDDS-7244: Fix multiple reports queued up from same DN and using up heap
- HDDS-7146: Recon: Add button to trigger OM DB Sync (Backend)
- HDDS-7231: Integrate the GetKeyInfo API to key read flows
- HDDS-7384: EC: ReplicationManager - implement deleting container handler
- HDDS-7361: Add general metrics for queues in Datanode
- HDDS-7432: Move command summary into Commands object in CommandQueue
- HDDS-7420: Bump Spring framework from 5.2.20 to 5.3.23
- HDDS-7421: Respect OZONE_LOGLEVEL and OZONE_ROOT_LOGGER for CLI commands
- HDDS-7316: Print stacktrace to identify the location of RocksObject leaks
- HDDS-1157: TestOzoneContainerWithTLS is failing with SSLHandshakeException
- HDDS-7406: Remove unused the exception and improve debug log in KeyDeletingService
- HDDS-7328: Improve Deletion of FSO Paths
- HDDS-7370: Add pending commands in SCM to Datanode command count
- HDDS-7220: SCM should use sub-ca certificate for token signature without HA enabled
- HDDS-7424: Bump jetty to 9.4.49.v20220914
- HDDS-7407: EC: Block allocation should not be stripped across the EC group
- HDDS-7402: Adapt CommandQueue to track the count of each queued command type
- HDDS-7422: Bump woodstox-core from 5.0.3 to 5.4.0
- HDDS-7349: Flaky integration test have memory leak for RatisDropwizardExports
- HDDS-7342: Move encryption-related code from MultipartCryptoKeyInputStream to OzoneCryptoInputStream
- HDDS-7413: Fix logging while marking container state unhealthy
- Revert HDDS-7253: Fix exception when '/' in key name
- HDDS-7396: Force close non-RATIS containers in ReplicationManager
- HDDS-7121: Support namespace summaries (du, dist and counts) for legacy FS buckets
- HDDS-7258: Cleanup the allocated but uncommitted blocks
- HDDS-7381: Cleanup of VolumeManagerImpl
- HDDS-7253: Fix exception when '/' in key name
- HDDS-7182: Add property to control RocksDB max open files
- HDDS-7284: JVM crash for rocksdb for read/write after close
- HDDS-7368: [Multi-Tenant] Add Volume Existence check in preExecute for OMTenantCreateRequest
- HDDS-7403: README Security Improvement
- HDDS-7199: Implement new mix workload Read/Write Freon command
- HDDS-7248: Recon: Expand the container status page to show all unhealthy container states
- HDDS-7141: Recon: Improve Disk Usage Page
- HDDS-7369: Fix wrong order of command arguments in Nonrolling-Upgrade.md
- HDDS-6210: EC: Add EC metrics

- HDDS-7355: non-primordial scm fail to get signed cert from primordial SCM when converting an unsecure cluster to secure
- HDDS-7356: Update SCM-HA.zh.md to match the English version
- HDDS-6930: SCM,OM,RECON should not print ERROR and exit with code 1 on successful shutdown
- Revert HDDS-7199: Implement new mix workload Read/Write Freon command which meets specific test requirements
- HDDS-7354: SchemaV3 blockData not deleted in table
- HDDS-7199: Implement new mix workload Read/Write Freon command which meets specific test requirements
- HDDS-7351: Use jackson-bom to ensure consistent Jackson version
- HDDS-7305: Fix Hadoop imports
- HDDS-7341: EC: Close pipelines with unregistered nodes
- HDDS-7058: EC: ReplicationManager - Implement ratis container replication check handler
- HDDS-7352: OM log flooded by AWSV4AuthValidator
- HDDS-7149: Update ratis version to 2.4.0 and thirdparty version to 1.0.2.
- HDDS-7254: Document that moving SCM from non-HA to HA is currently unsupported
- HDDS-6893: EC: ReplicationManager - move the empty container handling into RM from Legacy
- HDDS-7340: Bump jackson-databind to 2.13.4.2
- HDDS-7214: Continuous start and stop can have hanging threads in stopping
- HDDS-7343: Do not log exception on file not found in getFileStatus()
- HDDS-7337: Replace set-output in Github Actions workflow
- HDDS-7261: Add container location cache metrics
- HDDS-7330: Cleanup of BucketManagerImpl
- HDDS-7303: EC: ECBlockReconstructedStripeInputStream should set initialized only at the end of init()
- HDDS-7269: Remove fair lock configuration for container state management flows
- HDDS-7300: Race condition between full data scan and block deletion
- HDDS-7322: Remove unused dependencies hsqldb and jettison
- HDDS-7308: Fix config example syntax highlight in OM HA
- HDDS-7270: Fix bug in checking healthy replica placement policy check in EC underReplication handler
- HDDS-7291: Fixing exception handling in case of non positive replica index
- HDDS-7308: Removed cluster2 from serviceIds
- HDDS-7315: SCM order of close of dbstore and other services
- HDDS-7271: Ozone Integration test shows memory leak (graceful shutdown cleanup)
- HDDS-6946: Bump Hadoop to 3.3.4
- HDDS-7313: Update github actions for Node16
- HDDS-7088: OM incorrectly detects SCM Ratis Group ID when OM and SCM are colocated with same Ratis storage directory
- HDDS-6440: Handle custom metadata (x-amz-meta) during put-object through S3 API
- HDDS-7236: Enable Recon SCM DB bootstrap by default
- Revert HDDS-7271: Ozone Integration test shows memory leak (graceful shutdown cleanup)
- HDDS-7271: Ozone Integration test shows memory leak (graceful shutdown cleanup)
- HDDS-7304: EC: EC Decode can fail when byteBuffer from elastic pool is larger than chunksize
- HDDS-7230: Implement GetKeyInfo API
- HDDS-7299: Migrate simple tests in HDDS-container-service to JUnit5
- HDDS-6664: Implements getUri in TrashOzoneFileSystem
- HDDS-7287: Send deleteBlocksRequest with correct retry count
- HDDS-7302: Fix inaccurate container block count caused by missing block
- HDDS-7294: Move ozone manager background services to a separate package
- HDDS-7286: Clean up ContainerManager
- HDDS-7293: Bump jackson2 to 2.13.4
- HDDS-7259: Fix uncounted blocksDeleted in BlockDeletingService

- HDDS-7208: Erasure coding and encryption are not flagged on FileStatus
- HDDS-7217: OM logs wrong bucket layout
- HDDS-7211: Remove dead code from KeyManagerImpl
- HDDS-7282: Add EstimatedKeyCount metrics for SCM DB
- HDDS-7250: Add HttpServer metrics
- HDDS-7219: Remove not applicable info from ozone bucket list for links
- HDDS-7285: Fix concurrency issue in DeleteBlocksCommandHandler
- HDDS-7235: EOFException occurs when executing TPC-DS using o3fs
- HDDS-7283: Set coreSize equal to maxSize for threadPoolExecutor
- HDDS-7268: EC: Fix tests for HealthCheck handlers of RM that use Replica Indexes for Ratis Containers
- HDDS-7229: Introduce container location cache in ScmClient
- HDDS-7206: Change the placement policy interface to allow existing nodes to be specified for Rack Scatter Policy
- HDDS-7263: Add a handler for Quasi Closed containers to RM
- HDDS-7232: Introduce container-level lock while handling container report events by SCM and Recon
- HDDS-7222: Recon Inactive pipelines not shown
- HDDS-6863: Add Group-Id and Ratis-Roles Information for OM UI
- HDDS-7257: Bump snakeyaml to 1.32
- HDDS-7160: Recon: Make Container and Bucket stats clickable in Overviews page
- HDDS-6492: Add metric for failed container moves
- HDDS-7145: Recon: Show last OM DB sync time on Overview page
- HDDS-7039: EC: Handle the placement policy check in ECUnderReplicationHandler
- HDDS-7075: Implement FileSystem listStatusIterator to support file list paging
- HDDS-7251: Replace Log4j 1.x with Reload4j
- HDDS-7240: List all volume operation should go through ACL check as well in order to trigger audit logging
- HDDS-7239: EC: Add a Handler for CLOSING containers in Replication Manager
- HDDS-7203: LookupKey Latency breakdown
- HDDS-7185: Encode asterisk for parsing signature
- HDDS-7221: EC: ReplicationManager - Encapsulate the under and over rep queues into a queue object
- HDDS-7207: Fixed typo "layload" in om-echo command description as payload
- HDDS-7225: Add Ratis tests for HealthCheck handlers of Replication Manager
- HDDS-7144: Recon: Make only hostname fixed in Datanodes page
- HDDS-7193: S3gateway Grpc Ranger Authorizer NPE Unhandled Exception
- HDDS-7114: CHAINED and STAND_ALONE replication-type should be removed from help message
- HDDS-7192: EC: ReplicationManager - create handlers to perform various container checks
- HDDS-7209: Bump Jersey2 to 2.34
- HDDS-7204: EC: Schedule UnderReplicatedProcessor and OverReplicatedProcessor threads in RM instead of StorageContainerManager
- HDDS-7174: Migrate SCM integration tests to JUnit 5
- HDDS-2476: Share more code between metadata and data scanners
- HDDS-7086: DataNode UI to show failed volumes too
- HDDS-7191: Separate prop for s3 admin
- HDDS-7143: Recon: Fix left to right scrolling on Pipelines page
- HDDS-7201: testContainerIsReplicatedWhenAllNodesGotoMaintenance is failing frequently
- HDDS-7197: Schema V3 RocksDB instance statistics register as metrics collision
- HDDS-7200: Fixed inaccurate get request status code
- HDDS-7171: Add SupportedSourceVersion in annotation processors
- HDDS-7125: Inaccurate numBlocksDeletion when getTransactions
- HDDS-7175: Timed out tests missing from bundle
- HDDS-7186: Add support for Rocksdb user info LOG configure

- HDDS-6975: EC: Define the value of Maintenance Redundancy for EC containers
- HDDS-7163: Ozone debug container CLI supports container export
- HDDS-7176: Sort the container replica according to the replica index
- HDDS-6312: Use KeyPrefixContainer table to accelerate the process of DELETE/UPDATE events
- HDDS-6621: Include annotation processor module in regular project hierarchy
- HDDS-7158: ldb cli command supports to scan container V3
- HDDS-2954: Support admin groups
- HDDS-7165: Remove AbstractContractUnbufferTest
- HDDS-6493: Improve TestOmKeyInfoCodec error handling
- HDDS-6967: EC: Handle maintenance replicas in ECUnderReplicationHandler
- HDDS-7132: GetFileStatus returns NULL for some paths
- HDDS-7172: Intermittent failure in Close pipeline smoketest
- HDDS-7173: Fix TestReconOmMetadataManagerImpl#testUpdateOmDB
- HDDS-7166: Memory leak in Recon when replacing DB for new checkpoint
- HDDS-7161: Make Checksum.int2ByteString() zero-copy
- HDDS-6243: sourceToTargetMap in ContainerBalancer doesn't support multiple entries with same source
- HDDS-7162: Add Double in ConfigType
- HDDS-7104: Remove WARNINGS while running ozone fs/CLI commands
- HDDS-6456: Update RocksDB Version to 7.x
- HDDS-7164: Skip installing awscli if already present
- HDDS-7153: User defined override of ozone shell hardcoded values for log files
- HDDS-7151: Avoid using GeneratedMessage in non-generated code
- HDDS-7159: Use Java 8 for integration tests
- HDDS-7157: Log the error msg explicitly in BlockDeletingService
- HDDS-7159: Upgrade ubuntu version from 18.04 to 20.04
- HDDS-7108: Make ReplicationConfig validation error message type-specific
- HDDS-7150: Freon fail-at-end option caused not shutdown
- HDDS-7154: Fixed code issues in org.apache.hadoop.HDDS.client reported by sonar
- HDDS-5779: Create s3g bucket creation performance test
- HDDS-7135: ofs file input stream should support StreamCapabilities interface
- Multiple RPC clients per test run
- HDDS-7108: Include supported EC data-parity combinations in error message
- HDDS-7139: INVALID_TOKEN is handled as internal error in S3G
- HDDS-7136: Memory leak due to ChunkInputStream.close() not releasing buffer
- HDDS-7147: DirectoryDeletingService interval unit mismatch
- HDDS-7113: Support overriding response header values
- HDDS-7134: NPE when Ranger client throws RangerServiceException without Status
- HDDS-6458: Add ReplicationConfig options for bucketArgs in MiniOzoneChaosCluster
- HDDS-7076: Log container file path when container cannot be written
- HDDS-7045: Election info is out of date in Recon
- HDDS-7120: Prometheus displays only one volume_io_stats per datanode
- HDDS-7074: DN EndpointStateMachineMBean to add getType()
- HDDS-7127: Fix Transparent Data Encryption documentation
- HDDS-7033: Include committed space in log for disk out of space
- HDDS-7102: Exclude unnecessary RangerClient dependencies
- HDDS-6726: Close RocksObject in Recon and tools
- HDDS-6127: file checksum to support both CRC32 and CRC32C
- HDDS-5504: Refactor namespace summaries (du, dist and counts) for upcoming legacy FS buckets support
- HDDS-7116: Avoid leaking RocksObject from DBProfile
- HDDS-7022: EC: Open EC container are not closed when SCM container was already closed

- HDDS-6581: Introduce KEY_PATH_LOCK in OMKeyCreateRequest class
- HDDS-7043: Disallow ellipsis in PR title
- HDDS-7053: Add client-side pipelines distribution metrics
- HDDS-7109: SCM should close LayoutVersionManager in stopping stage
- HDDS-7047: Add Recon to ozone-topology sample compose environment
- HDDS-6997: Add support for "echo" lookup to measure raw performance in OM
- HDDS-7066: Include number of tests requested in Freon output
- HDDS-7087: Manage RocksObjects to detect leaks
- HDDS-7085: Update gRPC to 1.48.1 to address OOM bug
- HDDS-7023: Add CLI for resetting the expired deleted block txn
- HDDS-4970: Significant overhead when DataNode is over-subscribed
- HDDS-7063: Increment createContainer metric when container is created while writing chunk
- HDDS-6662: Provide basic hash code generator for OBSKeyPathLockStrategy
- HDDS-7051: Fix offset Condition in ECKeyOutputStream
- HDDS-7081: EC: ReplicationManager - UnderRep handler should handle duplicate indexes
- HDDS-7077: EC: Fix block deletion not allowed due to missing pipelineID
- HDDS-7044: Ignore pr_title_check for selective checks
- HDDS-6951: Replace bucket.listKeys() with bucket.listStatus() in OmBucketReadWriteKeyOps
- HDDS-5564: feature/PrefixFSO.md translation
- HDDS-7073: Incomplete audit log for KeyBlockToDelete
- HDDS-7069: EC: ReplicationManager - Track nodes already used when handing under replication
- HDDS-7048: EC: Add debug logging with exception info when stripe write failed
- HDDS-7068: EC: Prematurely re-threwed the exception in reconstruction cleanup loop
- HDDS-6533: support balancing EC container
- HDDS-7059: S3G should not call lookup twice for Get Object
- HDDS-7050: Recon "/missing" endpoint return limit too high
- HDDS-6899: EC: remove warnings and errors from console during online reconstruction of data
- HDDS-7055: NPE in ec.reconstruction.TokenHelper
- HDDS-7052: Topology Dynamic List document issue
- HDDS-7056: EC: Ensure replica index is maintained when replicating a container
- HDDS-7057: EC: ReplicationManager - Over replication handler should set repIndex on delete cmds
- HDDS-6728: SCM UI not showing correct HA roles
- HDDS-6956: Initialize ozoneLockProvider earlier in OzoneManager
- HDDS-6980: support block composite-crc checksum
- HDDS-7042: Rebuilding tenant cache omits empty tenants
- HDDS-7046: Leader count missing in Recon
- HDDS-7041: Avoid ConcurrentModificationException in RepeatedOmKeyInfo
- HDDS-6767: Add network location in Recon's datanode page
- HDDS-7021: [Multi-Tenant] Fix warning message when OM BG Sync is unable to get Ozone service policy version from Ranger
- HDDS-7036: Enable RocksDB stats in monitoring sample config
- HDDS-6871: Bump netty to 4.1.79
- HDDS-7034: File handle leak on datanode in PutBlock
- HDDS-7024: Add queue size metric for OM double buffer
- HDDS-5251: Refactor ContainerTestHelper factory methods
- HDDS-7032: RDBStore#getUpdatesSince should throw SequenceNumberNotFou...
- HDDS-7015: Ensure excluded nodes have network location when used in placement policies
- HDDS-6995: Update ranger-intg to v2.3.0
- HDDS-6807: Remove tcnative.version property
- HDDS-6976: OGB data moved by container balancer after successful iteration

- HDDS-7016: EC: Implement the Over replication Processor
- HDDS-7009: Avoid calling getBucket in object APIs
- HDDS-7020: EC: ReplicationManager - skip processing open containers
- HDDS-6904: Cleanup customer related events in S3G logs
- HDDS-6895: EC: ReplicationManager - Logic to process the under replicated queue and assign work to DN
- HDDS-7007: Missing spaces in error messages
- HDDS-7012: Bump aws-java-sdk-s3 to 1.12.261
- HDDS-7006: Incorrect usedNamespace caused by exceeding bytes quota
- HDDS-6987: EC: Implement RECOVERING Container Scrubber
- HDDS-7005: Recursive delete through bucket link deletes linked data
- Add missing recon address to docker-config file
- HDDS-6994: Update Contributing guide
- HDDS-6998: Ignore keys with higher updateID in OpenKeyDeleteRequest
- HDDS-6433: Refactor OMFailoverProxyProvider to provide a base for OM and GrpcOM FailoverProxyProviders
- HDDS-6909: [Multi-Tenant] Use RangerClient for Ranger operations
- HDDS-6940: EC: Skip the EC container for balancer
- HDDS-6990: AuthorizerLockImpl.java#tryWriteLockInOMRequest: move sanity args check to after tryWriteLockThrowOnTimeout()
- HDDS-6999: Correct volume parameter description for freon test tools
- HDDS-6989: EC: Add Test for RECOVERING container cleanup when failure
- HDDS-6981: Add metrics for BlockDeletingService on Datanode
- HDDS-6889: EC: put key command with EC replication can use ReplicationConfig validator
- HDDS-6761: [SCM HA finalization] Handle restarts, crashes, and leader changes
- HDDS-6992: ListKeys : SubPaths with a mixture of files and dirs breaks sorted order
- HDDS-6993: [Multi-Tenant] Set QuotaInBytes and QuotaInNamespace during Tenant Create
- HDDS-6954: Migrate simple tests in ozone-recon to JUnit5
- HDDS-6725: Close Rocks objects properly in StorageContainerManager
- HDDS-6991: Throw java.lang.NullPointerException when running setOwner
- HDDS-6901: Configure HDDS volume reserved as percentage of the volume space
- HDDS-6978: EC: Cleanup RECOVERING container on DN restarts
- HDDS-6982: EC: Attempt to cleanup the RECOVERING container when reconstruction failed at coordinator
- HDDS-6968: Addendum: [Multi-Tenant] Fix USER_MISMATCH error even on correct user
- HDDS-6794: EC: Analyze and add putBlock even on non writing node in the case of partial single stripe
- HDDS-6900: Propagate TimeoutException for all SCM HA Ratis calls
- HDDS-6938: handle NPE when removing prefixAcl
- HDDS-6960: EC: Implement the Over-replication Handler
- HDDS-6979: Remove unused plexus dependency declaration
- HDDS-6957: EC: ReplicationManager - prioritise under replicated containers
- HDDS-6723: Close Rocks objects properly in OzoneManager
- HDDS-6942: Ozone Buckets/Objects created via S3 should not allow group access
- HDDS-6965: Increase timeout for basic check
- HDDS-6969: Add link to compose directory in smoketest README
- HDDS-6970: EC: Ensure DatanodeAdminMonitor can handle EC containers during decommission
- HDDS-6977: EC: Remove references to ContainerReplicaPendingOps in TestECContainerReplicaCount
- HDDS-6217: Cleanup XceiverClientGrpc TODOs, and document how the client works and should be used
- HDDS-6773: Cleanup TestRDBTableStore - fix checkstyle
- HDDS-6773: Cleanup TestRDBTableStore
- HDDS-6676: KeyValueContainerData#getProtoBufMessage() should set block count
- HDDS-6959: Fix move timeout in latest iteration metric
- HDDS-6879: Bucket layout as null on default layout type

- HDDS-6974: Container report processing in Recon is single threaded
- HDDS-6973: Refactor ObjectEndpoint
- HDDS-6894: EC: ReplicationManager - Add class to handle under-replication and form a command for a datanode
- HDDS-6968: [Multi-Tenant] Fix USER_MISMATCH error even on correct user
- HDDS-6928: ozone container balancer CLI went in hung state due to deadlock
- HDDS-6948: MultiTenantAccessAuthorizerRangerPlugin#deletePolicyByName should log and exit if getAccessPolicyByName returns null
- HDDS-6934: OM crashed with OzoneManagerDoubleBuffer
- HDDS-6945: EC: EC Reconstruction Command count queues should be included in DN heartbeat
- HDDS-6949: Audit bucket ACL changes in S3 Gateway
- HDDS-6947: Bump gson version to 2.9.0
- HDDS-6699: EC: ReplicationManager - collect under and over replicated containers
- HDDS-6943: Reject property change requests on link buckets
- HDDS-6944: EC: Handle reconstructECContainersCommand in heartbeat
- HDDS-6937: Close Rocks objects properly in TrashOzoneFilesystem
- HDDS-6906: Fix kerberos config for secure smoke tests
- HDDS-6941: Setting Bucket Property can corrupt bucket layout
- HDDS-6835: EC: Provide correct example for EC in ozone.server.default.replication
- HDDS-6812: Fix INTERNAL_ERROR message on failed write
- HDDS-5916: Datanodes stuck in leader election in Kubernetes
- HDDS-6873: EC: Add block token support for ECReconstructionCoordinator
- HDDS-6936: Mark TestContainerStateMachineFailures#testApplyTransactionFailure as flaky
- HDDS-6891: Add CapacityVolumeChoosingPolicy
- HDDS-6876: Migrate test with rules in HDDS-common to JUnit5
- HDDS-6932: [Multi-Tenant] Move Ranger plugin version to a separate tag
- HDDS-6868: Add S3Auth information to thread local
- HDDS-6877: Keep replication port unchanged when restarting datanode in MiniOzoneCluster
- HDDS-6907: OFS should create buckets with FILE_SYSTEM_OPTIMIZED layout
- HDDS-6875: Migrate parameterized tests in HDDS-common to JUnit5
- HDDS-6924: OBJECT_STORE isn't flat namespaced
- HDDS-6899: [EC] Remove warnings and errors from console during online reconstruction of data
- HDDS-6695: Enable SCM Ratis by default for new clusters only
- HDDS-4123: Integrate OM Open Key Cleanup Service Into Existing Code
- HDDS-6882: Correct exit code for invalid arguments passed to command-line tools
- HDDS-6890: EC: Fix potential wrong replica read with over-replicated container
- HDDS-6902: Duplicate mockito-core entries in pom.xml
- HDDS-6752: Migrate tests with rules in HDDS-server-scm to JUnit5
- HDDS-6806: EC: Implement the EC Reconstruction coordinator
- HDDS-6829: Limit the no of inflight replication tasks in SCM
- HDDS-6898: [SCM HA finalization] Modify acceptance test configuration to speed up test finalization
- HDDS-6577: Configurations to reserve HDDS volume space
- HDDS-6870: Clean up isTenantAdmin to use UGI
- HDDS-6872: TestAuthorizationV4QueryParser should pass offline
- HDDS-6840: Add MetaData volume information to the SCM and OM - UI
- HDDS-6697: EC: ReplicationManager - create class to detect EC container health issues
- HDDS-6841: EC: Validate the server default configuration on Ozone manager startup
- HDDS-6830: EC: SCMContainerPlacementRackScatter may choose fewer nodes than required
- HDDS-6260: EC: Standalone containers should not move to quasi-closed
- HDDS-6799: Refactor OMLockMetrics code in OzoneManagerLock

- HDDS-6878: Only export span if Ozone tracer init
- HDDS-6874: Migrate simple tests in HDDS-common to JUnit5
- HDDS-6869: Use UGI when checking Upgrade priv
- HDDS-6837: Avoid duplicate headers in HTTP response
- HDDS-6845: Add CI check for pull request title
- HDDS-6846: HeadOp ignored for link bucket
- HDDS-6839: Migrate tests in HDDS-client to JUnit5

Known issues in Ozone 718.1.0 parcel

You must be aware of the known issues and limitations, the areas of impact, and workaround in Ozone parcel.

Tez Configuration Changes

The following configuration changes have to be made to pick up the latest Ozone FS jar from the Ozone parcel (when installed):

CDPD-48540

For `tez.cluster.additional.classpath.prefix`, the value is `/var/lib/hadoop-hdfs/*` (Tez Additional Classpath)

For `tez.user.classpath.first`, the value is `true` (Tez Client Advanced Configuration Snippet (Safety Valve) for `tez-conf/tez-site.xml`)

CDPD-47605

For `tez.cluster.additional.classpath.prefix` the value is `/var/lib/hadoop-hdfs/*` (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

For `tez.user.classpath.first`, the value is `true` (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

Restart the TEZ and HIVE_ON_TEZ service as prompted by Cloudera Manager.

Update Yarn to use updated Ozone FS Jar

CDPD-48500: Ozone parcel activation or install should handle redeployment of YARN jars and clean-up cache.

Perform the following steps:

1. Log in to Cloudera Manager UI
2. Navigate to Clusters
3. Select the YARN service
4. Click Actions
5. Click Install YARN Service Dependencies
6. Click YARN MapReduce Framework JARs
7. Restart the CDP 7.1.8 cluster

Other issues

SSL Handshake fails between Ozone DataNodes if the two DataNodes have their certificate signed by different Ozone Storage Container Managers.

Ozone DataNode certificates are signed by the leader Storage Container Manager. Due to an issue in creating a TrustStore for DataNode to DataNode connections, the trust cannot be established between the two DataNodes if a different Storage Container Manager signs their certificate. These connections fail to establish and display an SSL Handshake Exception. This affects Pipeline creation and container replication (also EC container reconstruction). The symptoms vary, depending on the number of the nodes that have different singer certificates, either these DataNodes do not participate in any Ratis-3 Pipeline, or have Pipelines exclusively in between groups with the same signer. Over time this can lead to an imbalance in DataNode usage, and it might cause

decommission of a DataNode stuck if the data has to be replicated to a node with a certificate that has a different signer.

This problem affects all the 7.1.8 Ozone Parcel releases.

To identify if the problem is present on a cluster, the output of `ozone admin cert list` command must be examined. Ensure you define a sufficient number of certificates to be returned with the `-c` option to see all the certificates issued in the system.

If there are different Issuers for the latest DataNode certificates, this indicates the cluster is affected.

It is possible to avoid the problem by checking the `ozone admin scm roles` output and see if the Primordial node is the actual leader node before adding a new DataNode and starting it for the first time. If the leader SCM node is a different node, then calling `ozone admin scm transfer` to make the Primordial node the leader can put the cluster into the desired state before adding the new DataNode.

Ensure that all the certificates in the cluster are signed by the same Storage Container Manager node.

Perform the procedure below if the cluster is affected and has DataNode certificates signed by a different issuer:

1. Identify the current leader Storage Container Manager by running the command `ozone admin scm roles`.
2. If the leader did not sign the majority of DataNode certificates, change the leader that has signed the majority of DN certificates by running the `ozone admin scm transfer` command with the proper Ozone SCM Service ID (set in Cloudera Manager>Ozone>Configuration) and the UUID of the desired leader node.
3. Stop the minority of DataNodes that have a different sign than the majority.
4. Locate the Datanode Metadata Directory (set in Cloudera Manager>Ozone>Configuration) on the hosts of the stopped DataNode and move the directory to a backup location.
5. Start the previously stopped DataNodes.
6. After regenerating the certificates, check if the DataNodes are joining the Ratis-3 Pipeline.

To avoid service disruptions, you can stop the DataNodes one by one instead of all together. This can cause a data outage for the downtime of the node if there are files with Ratis-1 replication on the cluster and the single replica of these files reside on the node being restarted. Changing the leader node on an SCM HA-enabled cluster should not disrupt operations.

CDPD-56006: On providing an incorrect hostname/service ID in the Ozone URI, the filesystem client instead of failing, retries till exhaustion and the default retry is too high.

Configure `ozone.client.failover.max.attempts` to a lower the value to avoid long endless retries.

CDPD-49137: Sometimes OM's kerberos token is not updated and it stops being able to communicate with SCM. When this occurs, writes will start to fail.

Restarting OM or setting the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true` will fix the issue.

CDPD-49808: Spark jobs against Ozone intermittently fail with ERROR spark.SparkContext: [main]: Error initializing SparkContext.java.lang.IllegalStateException: No filter named.

This is an intermittent failure which can be retried.

CDPD-49918: ECContainerReconstructionThread hits a precondition failure for checksum validations. This results in Reconstruction thread failing.

None.

CDPD-50678: Deleting containers which have one or more replicas which are not empty on the Datanode can cause the container to be stuck in a deleting state indefinitely. Containers in this state can also block decommission or maintenance operations completing.

None.

CDPD-50665: Due to HDDS-8171, if a datanode is IN_MAINTENANCE and stopped, and another node is decommissioned which is hosting an EC replica of the same index as the one on the IN_MAINTENANCE node, decommission can get blocked.

Start the datanode process on the IN_MAINTENANCE node so its replicas are available again.

CDPD-48932: Datanode needs off heap memory due to direct buffers and rocks DB usage. This can lead to the memory footprint to be larger than the configured physical memory on the node.

Remove other memory intensive services from the node, make sure the physical memory is sufficient for the planned load and scale.