Cloudera Runtime 7.3.2

Cloudera Storage Optimizer

Date published: 2020-07-28
Date modified: 2026-03-31

https://docs.cloudera.com/

Legal Notice

Contents

Cloudera Storage Optimizer overview

Learn about Cloudera Storage Optimizer and it's benefits.

Cloudera Storage Optimizer is an intelligent data lifecycle management feature that automatically reduces storage usage by converting infrequently accessed data from replicated storage (RATIS 3×) to space-efficient erasure coding (EC). It analyzes access patterns and applies configurable policies to identify and convert cold data, reducing the storage footprint by 45% to 60% while maintaining data durability and availability. The benefit is additional usable capacity for the same licensed capacity.

Important:

- By default, Cloudera Storage Optimizer runs daily. However, this is configurable.
- Initial conversion may take several hours depending on the data volume.
- The converted data remains fully accessible with minimal latency impact.
- Monitor the Cloudera Storage Optimizer dashboard for conversion statistics and savings metrics.

Advantages

The Cloudera Storage Optimizer continuously monitors your data access patterns and automatically transitions cold data to more efficient storage formats allowing you to optimize the following:

- **Storage reduction**: Converts 3x replicated data to EC format with only 40 to 50 percent of overhead.

  Example,

  For 100 TB of data where 70% is cold (Data Under Management unchanged):

  - Before (3× replication): 100 TB × 3 = 300 TB storage used
  - After (3× replication):
    - Hot data of 30 TB at 3x replication: 30*3 = 90 TB
    - Cold data of 70 TB at EC of approximately 1.5x: 70*1.5 = 105 TB
    - Total storage used after conversion: 90 + 105 = 195 TB
    - Total storage freed after conversion: 300 - 195 = 105 TB (approximately 35 percent of storage usage saved)

  Important: Licensing is based on **Data Under Management** which is still 100 TB here. The benefit is extra usable storage capacity of 105 TB freed on the same infrastructure, but not a change in license.

- **Cost savings**: Reduces upto 45 to 60 percent of storage costs for cold data.
- **Automated management**: No manual intervention is required once configured.
- **Compliance ready**: Maintains full data durability and audit trails.
- **Performance optimized**: Hot data remains in high-performance RATIS format.

Limitations

Cloudera Storage Optimizer supports only one-way conversion from RATIS to EC for your keys. Once they are converted to EC keys, Cloudera Storage Optimizer does not support coverting back from EC to RATIS.

How does Cloudera Storage Optimizer works?

The Cloudera Storage Optimizer operates through a sophisticated five-stage pipeline that runs daily to optimize your storage automatically. This builds upon Ozone's existing replication pipeline to provide intelligent data management.

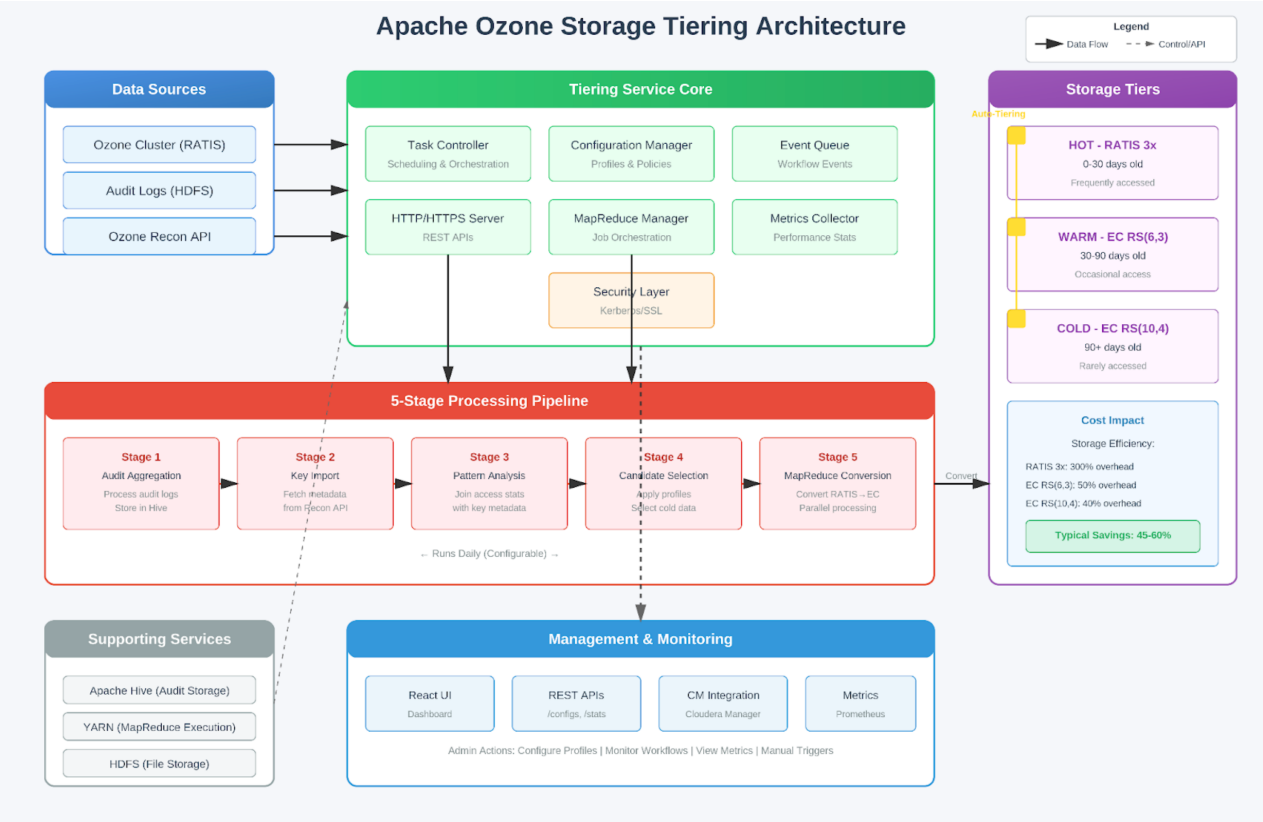

The five-stage processing pipeline works as follows:

- **Audit log aggregation**: The system collects and processes Ozone audit logs to understand data access patterns. Every read, write, and delete operation is tracked and aggregated into structured data using Cloudera Hive-on-Tez.
- **Metadata import**: The Cloudera Storage Optimizer fetches comprehensive metadata about all storage objects from Ozone Recon API, including file sizes, creation dates, and current replication types.
- **Access pattern analysis**: By joining metadata with access statistics, the system creates a complete view of your data usage patterns, identifying which files are truly **cold**. This analysis leverages Cloudera's data engineering capabilities.
- **Candidate selection**: Configurable policies determine which files should be converted. For example, files larger than 32 MB that are not accessed in 30 days become conversion candidates.
- **Distributed conversion**: MapReduce jobs perform the actual conversion from RATIS to Erasure Coding, processing data in parallel for maximum efficiency using YARN resources. All of this happens in the background without impacting your applications or users.

Installing Cloudera Storage Optimizer

Learn how to install Cloudera Storage Optimizer.

Prerequisites

Understand the necessary components and conditions required for installing Cloudera Storage Optimizer.

- Cloudera Storage Optimizer requires the following Cloudera components to be operational during the workflow execution:
  - Cloudera Ranger service: Used for streaming audit logs

  - Cloudera Hive-on-Tez service: Used for executing Hive queries against Storage Optimizer metadata tables
  - Cloudera YARN service: Required for running MapReduce jobs during data conversion
  - Cloudera ZooKeeper: Required for discovering Hive service used by Cloudera Storage Optimizer
  - (Optional) Cloudera Knox: Required for accessing secure clusters
- Make sure that your Cloudera Base on premises version is 7.3.2.0 and Cloudera Manager version is 7.13.2.0.

Adding Cloudera Storage Optimizer to an existing Ozone cluster after upgrade

Learn how to integrate Cloudera Storage Optimizer into an existing Ozone cluster after upgrading Cloudera Base on premises to 7.3.2.0 and Cloudera Manager to 7.13.2.0.

Before you begin

Before adding Cloudera Storage Optimizer to an existing Ozone cluster, do the following steps:

1. Sign in to Cloudera Manager.
2. In the left navigation, click Clusters and select the Ozone cluster.
3. Click Actions drop-down button and click Create Ranger Ozone Plugin Audit Directory.

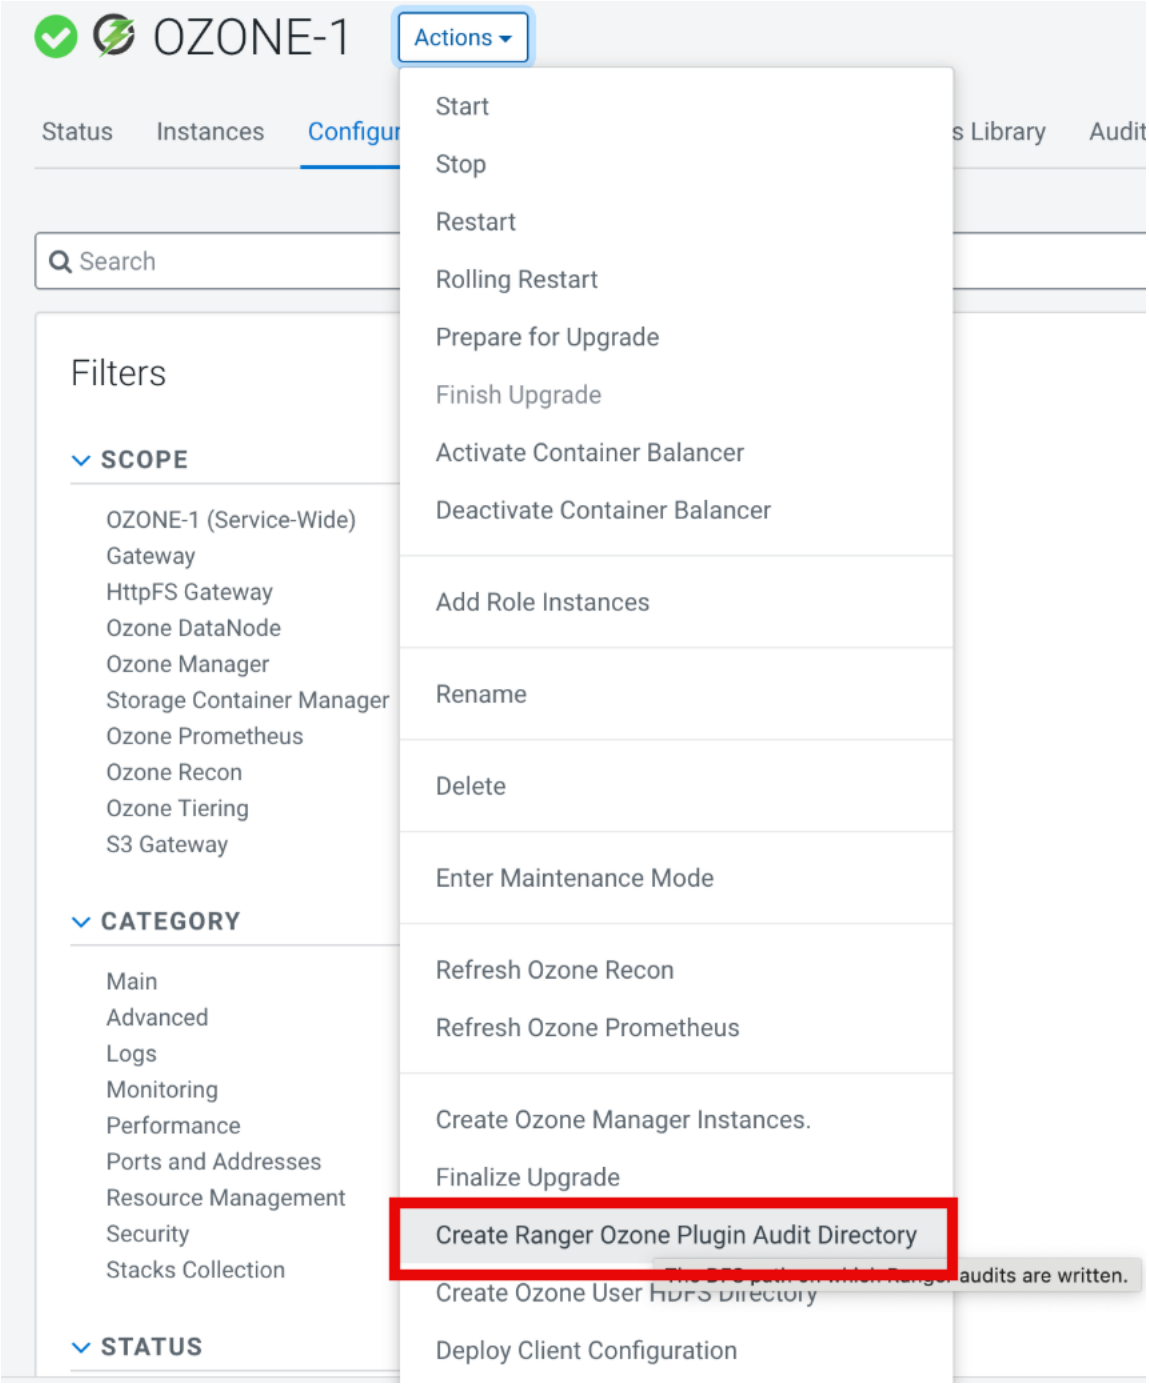

4. Click Actions drop-down button and click Create Ozone User HDFS Directory.

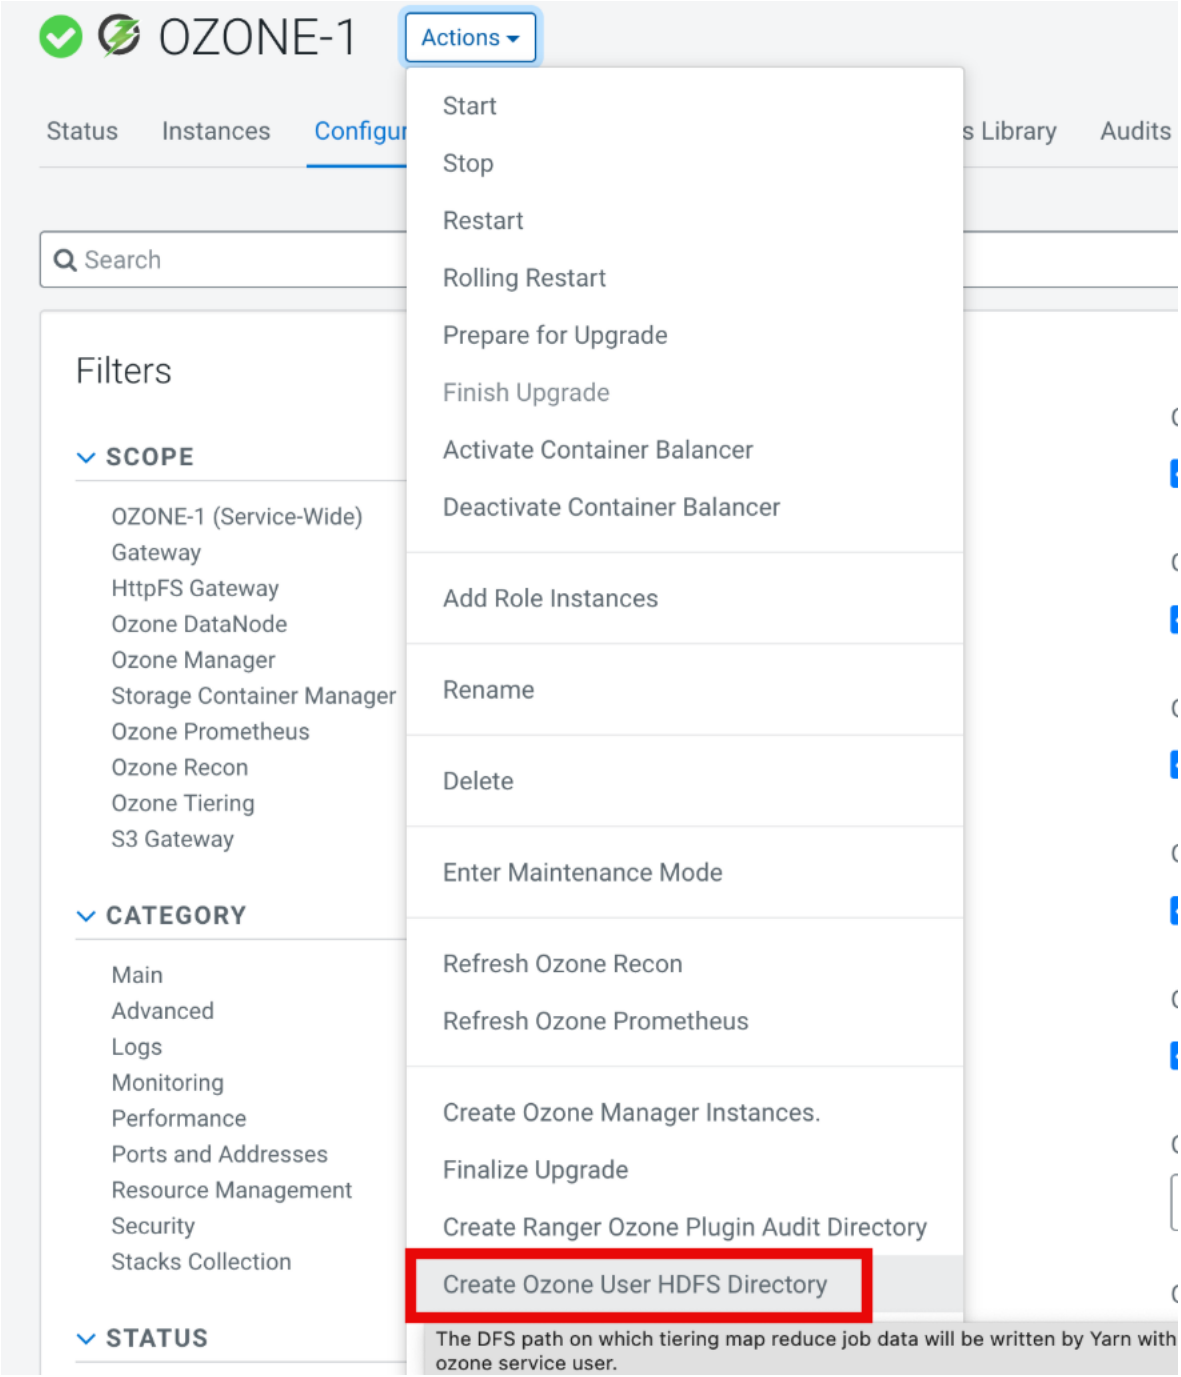

5. Perform the steps mentioned in Configuring Ranger policies.

About this task

Perform the following steps to add Cloudera Storage Optimizer to an existing Ozone cluster after upgrade:

Procedure

1. Sign in to Cloudera Manager.
2. In the left navigation, click Clusters and select the Ozone cluster.
3. Click Instances tab.
4. Click Add Role Instances button. Add Role Instances to <***CLUSTER-NAME***> page opens.

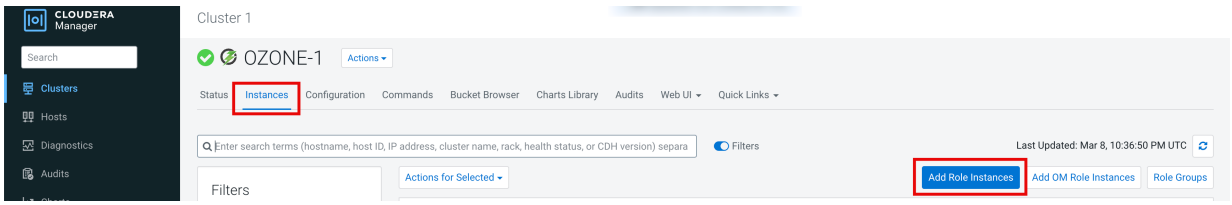

5. Click Ozone Tiering field. A popup window appears displaying all the hosts.

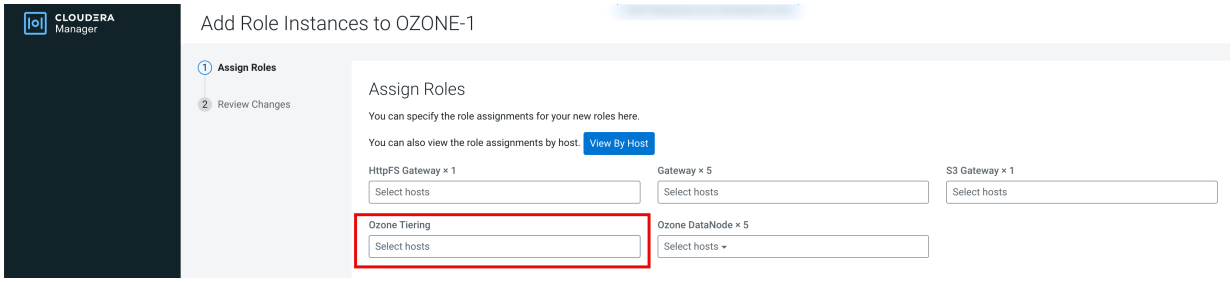

6. Select the hostname onto which you want to install the sevice and click OK.

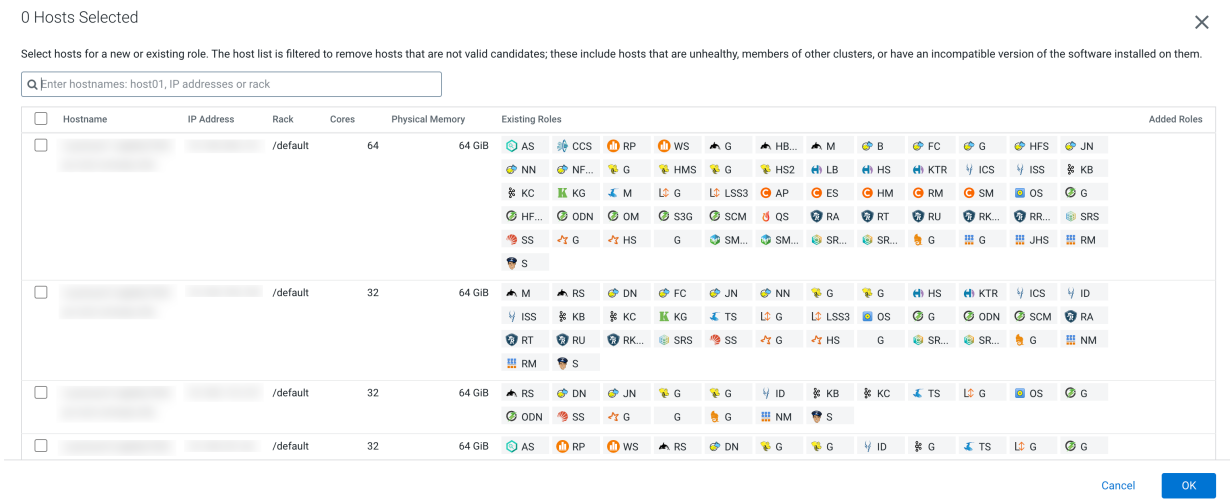

7. Click Continue Finish .
8. Once Cloudera Storage Optimizer is added, click Instances tab and click on Ozone Tiering in the Role Type column. The Cloudera Storage Optimizer UI page opens.
9. Click Actions drop-down button and click Start this Ozone Tiering.
10. Wait for service to display Active status and Good health.

Configuring Ranger policies

Learn how to configure Ranger policies to manage Cloudera Storage Optimizer.

Important: Configuring Ranger policies manually is needed only when upgrading Cloudera Base on premises version to 7.3.2 and Cloudera Manager version to 7.13.2. Fresh installations automatically configure these policies.

Policy 1: Hive Access Policy (Hadoop SQL)

1. Sign in to Ranger Admin UI and navigate to Resource Policies Hadoop SQL .
2. Click Add New Policy.
3. Enter the following configurations:

- Under Policy Details, enter the following configurations:
  - Policy Name: OZONE_TIERING_HIVE_POLICY
  - Description: Ozone tiering hive policy to grant access to om user on the URL for the Ozone storage
- Under Resources URL , enter ofs://${<***OZONE_SERVICE_ID***>}/tiering-volume/tiering-bucket/tables.
- Under Allow Conditions, add the following configurations:
  - Users: hdfs, om
  - Permissions: All

4. Click Save.

Policy 2: Volume Read Policy (cm_ozone)

1. Sign in to Ranger Admin UI and navigate to cm_ozone policy page.
2. Click Add New Policy.
3. Enter the following configurations:
   - Under Policy Details, enter the following configurations:
     - Policy Name: OZONE_TIERING_VOLUME_POLICY
     - Description: Policy for OZONE_TIERING_VOLUME_POLICY to grant hive user read access to tiering-volume
     - Audit Logging: Enable it to Yes
   - Under Resources Ozone Volume , enter tiering-volume.
   - Under Allow Conditions, add the following configurations:
     - Users: hive.
     - Permissions: Read
4. Click Save.

Policy 3: Bucket Access Policy (cm_ozone)

1. Sign in to Ranger Admin UI and navigate to cm_ozone policy page.
2. Click Add New Policy.
3. Enter the following configurations:
   - Under Policy Details, enter the following configurations:
     - Policy Name: OZONE_TIERING_POLICY
     - Description: Policy for OZONE_TIERING_POLICY to grant hive user all access to tiering-volume/tiering-bucket
     - Audit Logging: Enable it to Yes
   - Under Resources, add the following configurations:
     - Ozone Volume: Enter tiering-volume and enable Include toggle button on.
     - Ozone Bucket: Enter tiering-bucket and enable Include toggle button on.
     - Ozone Key: *
   - Under Allow Conditions, add the following configurations:
     - Users: hive.
     - Permissions: All, Create, Write, Read, List, Delete
4. Click Save.

Configuring Cloudera Storage Optimizer

Learn how to configure Cloudera Storage Optimizer for efficient data management and enhanced performance.

Configuring profiles through Cloudera Storage Optimizer UI

Cloudera Storage Optimizer offers a user-friendly web interface for configuring data conversion policies through the Cloudera Storage Optimizer UI. Learn how to configure profiles using the Cloudera Storage Optimizer UI.

About this task

Various types of profiles can be created and used, but only one will be default and activated at a time. Following are some examples which can be used as default profiles:

- Standard profile (balanced):
  - Keys older than 30 days
  - Size above 32 MB
  - Not accessed in the time window

  It converts files larger than 32 MB that are not accessed in 30 days to EC format.
- Aggressive profile (maximum savings):
  - Keys older than 7 days
  - Size above 1 MB
  - Ideal for logs and temporary data

  It aggressively converts log files to EC after just 7 days of inactivity.
- Conservative profile (minimal risk):
  - Keys older than 90 days
  - Size above 100 MB
  - For critical archived data

  It is useful for conservative approach for important data, and for only converting large files that are rarely accessed over 3 months.

Procedure

1. Sign in to Cloudera Manager.
2. In the left navigation, click Clusters and select the Ozone cluster.
3. Click Instances tab and click on Ozone Tiering in the Role Type column. The Cloudera Storage Optimizer UI page opens.
4. Click Configuration tab.
5. Click Profile Configurations to create and manage storage optimization profiles.

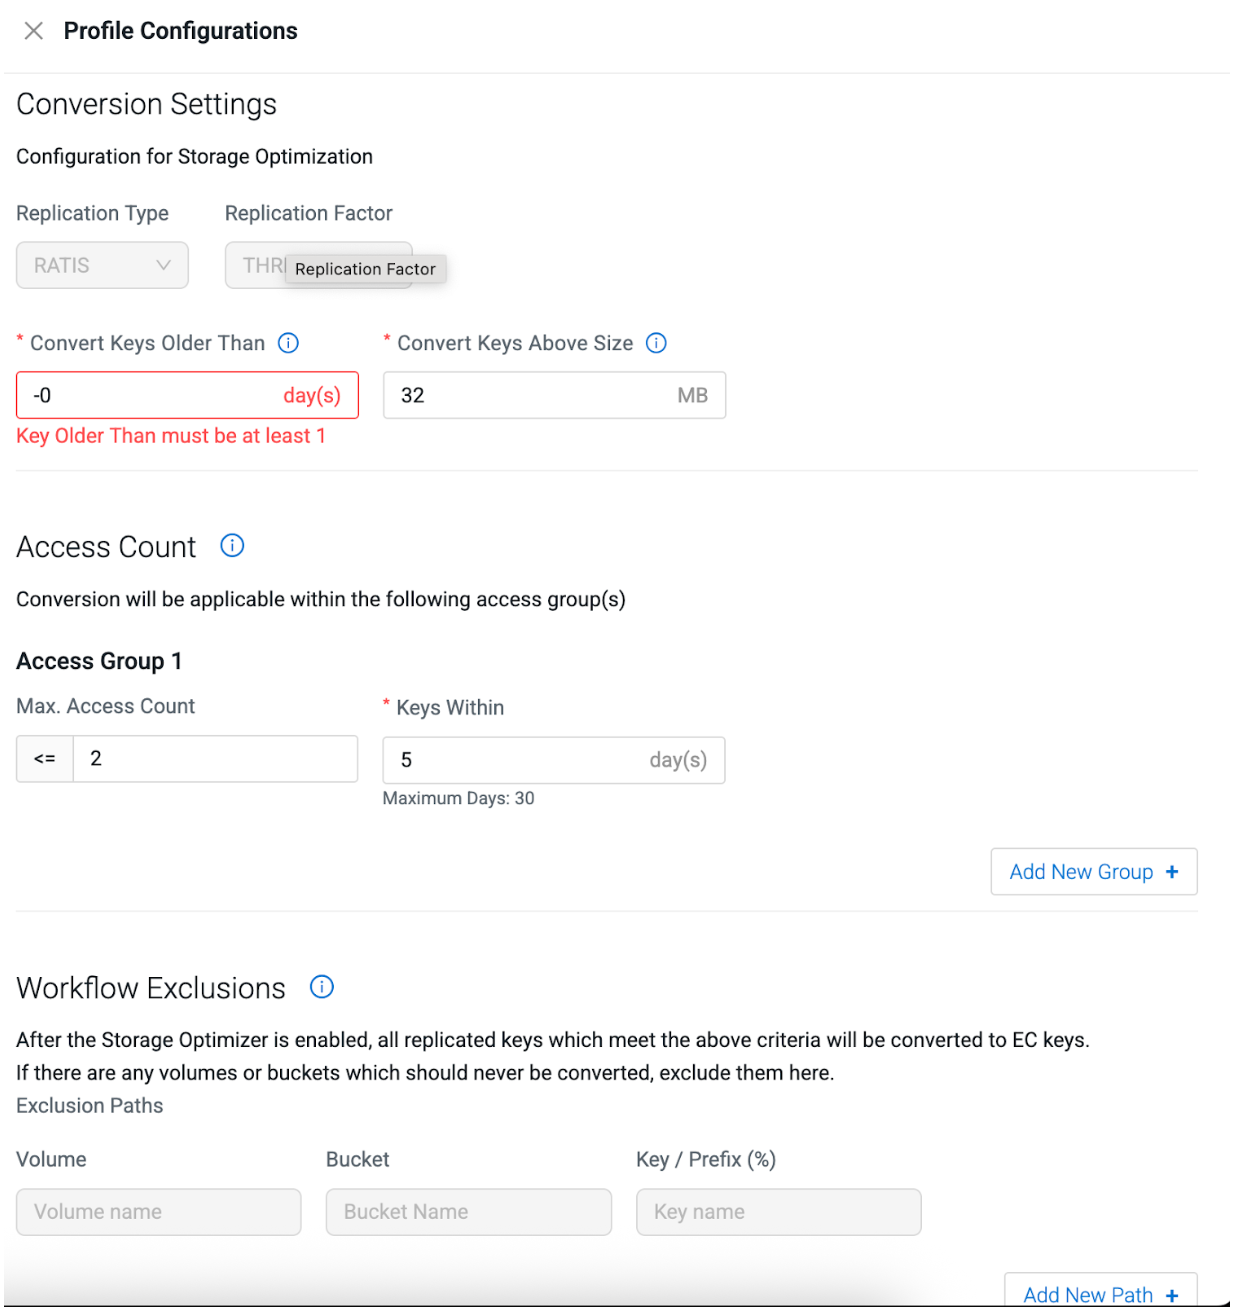

6. Update the following profile configurations:
   a. Conversion Settings: Configuration for Storage Optimization
      1. Replication Type: Select RATIS. By default, only RATIS replicated data can be converted to erasure coding (EC).
      2. Replication Factor: Enter THREE for standard 3x replication or your cluster's replication setting.
      3. Convert Keys Older Than: Enter the minimum number of days after which the files are eligible for conversion. Minimum value is 1 day. Default value is 30 day(s).

         Example, setting it to 30 days means that only files that are not accessed for 30 days or more will be converted to EC.
      4. Convert Keys Above Size: Enter the minimum size of files in MB that are eligible for conversion. Default value is 32 MB. However, Cloudera recommends to use 3 MB to 100 MB depending on your use case.

Note: Files that are small in size do not significantly benefit from EC conversion.

b. Access Count Configuration: Defines access patterns to identify truly cold data. You can create multiple access groups for granular control. Under Access Group, update the following:

1. Max. Access Count: Maximum number of times a file can be accessed (use "<=" for less than or equal).
2. Keys Within: Time interval in days for counting accesses. Maximum value is 30 days.

Example, if Max Access Countis set as <= 2 and Keys Within is set as 5 days, it means Cloudera Storage Optimizer convert files that are accessed 2 times or less in the past 5 days.

Click Add New Group + to define multiple access criteria. Files must meet all the Access Group conditions to be converted.

c. Workflow Exclusions: Protect critical data by defining paths that should never be converted, regardless of other criteria. Under Exclusion Paths, update the following:

1. Volume: Enter the volume name that you want to exclude.

   Example, critical-volume
2. Bucket: Specify the bucket name that you want to exclude.

   Example, real-time-bucket
3. Key / Prefix (%): Define the key patterns using % as wildcard.

   Example, /important/% excludes all keys under /important/

   Click Add New Path + to add multiple exclusion rules.

Important: After the Cloudera Storage Optimizer is enabled, all replicated keys which meet the criteria that is set will be converted to EC keys. Carefully review exclusions to protect critical data.

Configuring Prometheus endpoint

Learn how to configure Prometheus endpoint configuration ozone.tiering.prometheus.http.endpoint in your Cloudera Storage Optimizer UI.

Procedure

1. Identify the Prometheus endpoint URL.
   a) Sign in to Cloudera Manager.
   b) In the left navigation, click Clusters and select the Ozone cluster.
   c) Click Instances tab.
   d) Under Filters, click ROLE TYPE dropdown menu.
   e) Click Ozone Prometheus.
   f) Under Role Type column, click the relevant Ozone Prometheus instance. The Ozone Prometheus UI opens.
   g) Click the Prometheus Web UI tab.
   h) Copy the HTTP Hostname of the Ozone Prometheus Web UI and paste it in a text editor.
2. Sign in to Cloudera Manager.
3. In the left navigation, click Clusters and select the Ozone cluster.
4. Click Instances tab and click on Ozone Tiering in the Role Type column. The Cloudera Storage Optimizer UI page opens.
5. Click Configuration tab.
6. Search for Ozone Tiering Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml and add the following configuration:

   a. Name: ozone.tiering.prometheus.http.endpoint
   b. Value: <***OZONE-PROMETHEUS-ENDPOINT-URL***>: <***HTTP-PORT***>

Where,

- <***OZONE-PROMETHEUS-ENDPOINT-URL***> is the Prometheus Hostname that you have copied earlier in 1.h on page 13.
- <***HTTP-PORT***>: The default http port number is 9096.

Important: Make sure the Prometheus endpoint Hostname that you have copied and the port number are HTTP only. Do not use HTTPS endpoint Hostname and port number.

7. Click Save Changes.

Configuring YARN resources for MapReduce conversion jobs

Cloudera Storage Optimizer uses MapReduce jobs for data conversion. Learn how to configure sufficient YARN resources for MapReduce conversion jobs.

About this task

Table 1: Recommended YARN resources by data volume

Daily conversion volume	Concurrent mappers	YARN memory	YARN vCores	Estimated time for conversion
Less than 10 TB	10	10 GB	10	2 to 4 hours
10 to 50 TB	20	20 GB	20	3 to 6 hours
50 to 100 TB	30	30 GB	30	4 to 8 hours
Greater than 100 TB	50	50 GB	50	6 to 12 hours

Before you begin

Minimum requirements are as follows:

- Container Memory: 1 GB per mapper
- Container vCores: 1 vCore per mapper
- Concurrent Mappers: Default value is10 (configurable through UI setting Key Conversion Concurrent Mappers)
- Total Minimum: 10 GB RAM and 10 vCores available in YARN

Procedure

1. Sign in to Cloudera Manager.
2. In the left navigation, click Clusters and select the Ozone cluster.
3. Click Instances tab and click on Ozone Tiering in the Role Type column. The Cloudera Storage Optimizer UI page opens.
4. Click Configuration tab.
5. Click Ozone Key Settings.
6. Update the following configurations:
   - Key Conversion Batches Per Mapper: Specify the number of batches that can be converted per mapper to control the split size. Default value is 10.
   - Key Conversion Concurrent Mappers: Specify the number of key conversion mappers are needed concurrently. Increase the value for faster conversion rate. Cloudera recommends to use 20 to 50 mappers for large buckets. Conversion runs as standard priority jobs in the default queue.

Troubleshooting

Learn how to identify, analyze, and resolve technical issues related to Cloudera Storage Optimizer efficiently.

Custom Hive connection fails

If your custom Hive setup or the default Hive connection fails, then the Cloudera Storage Optimizer might fail to start and initialize. To resolve this issue, add the following configurations in the Ozone Tiering Advanced Configuration Snippet (Safety Valve) in Ozone configurations.

Procedure

1. Sign in to Cloudera Manager.
2. In the left navigation, click Clusters and select the Ozone cluster.
3. Click Configuration tab.
4. Search for Ozone Tiering Advanced Configuration Snippet (Safety Valve) and add the following configurations:

```
<property>
  <name>ozone.tiering.hive.connection.url</name>
  <value>[Custom Hive URL]</value>
</property>
```

Important: For correct Hive URL value, contact Cloudera Support as it requires environment-specific configuration.

Ozone cluster does not have exec permission on /tmp path

If the Ozone cluster does not have exec permission on the /tmp path, Cloudera Storage Optimizer might fail. To resolve this issue, add the following configurations in the Ozone configurations.

Procedure

1. Open Terminal in your device.
2. Create a temporary directory with exec permissions by running the following commands:

```
klist
kdestroy
kinit -kt <***KEYTABS_PATH***>/om.keytab om
mkdir -p /opt/tiering/tmp
chown om:om /opt/tiering/tmp
chmod 1777 /opt/tiering/tmp
```

3. Sign in to Cloudera Manager.
4. In the left navigation, click Clusters and select the Ozone cluster.
5. Click Configuration tab.
6. Search for ozone_java_opts and add the following configuration:

```
-Djava.io.tmpdir=/opt/tiering/tmp
```

7. Search for ozone_tiering_java_opts and add the following configuration:

```
-Dparquet.compression.codec.snappy.tmpdir=/opt/tiering/tmp
```

8. Click Save Changes.
9. Click Actions dropdown menu and click Restart to apply the above configurations.

Ranger service errors

Cloudera Storage Optimizer relies on Ranger audit files being streamed to the Ozone HDFS location for access pattern analysis.

Potential issues:

- Ranger service is down or unavailable.
- Audit file streaming is interrupted due to external issues.
- Network connectivity problems between Ranger and Ozone.

Error reporting:

- If audit files are not available at the expected Ozone HDFS location, Cloudera Storage Optimizer will detect the issue during workflow execution.
- Error messages will be updated at the appropriate task steps of the workflow run.
- These errors will be visible to users through the Storage Optimizer UI workflow status

Recommendation: Ensure all required services such as Ranger, Hive-on-Tez, and YARN are running and healthy before initiating Cloudera Storage Optimizer workflows. Monitor service health through Cloudera Manager and check Cloudera Storage Optimizer UI for any workflow errors.