

Troubleshooting for Cloudera Migration Assistant

Date published: 2026-06-18

Date modified: 2026-06-18



Legal Notice

© Cloudera Inc. 2026. All rights reserved.

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

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

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

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

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

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

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

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

Contents

Troubleshooting for Cloudera Migration Assistant.....	5
Log locations and key file paths.....	5
Scan failures.....	6
Hive scan fails with "Failed to run async statement" after HiveServer2 restart.....	6
Hive scan fails with "Unable to retrieve HS2 host information from ZooKeeper".....	6
Ozone scan fails with "PERMISSION_DENIED" on volume or bucket access.....	7
Scan stays in IN_PROGRESS indefinitely.....	8
Kerberos authentication error during scan.....	8
Scan completes but shows no results.....	9
Execution failures.....	9
Execution step stays in IN_PROGRESS.....	9
Cloudera Replication Manager policy creation fails.....	10
Canceling a running execution.....	10
CMA Agent issues.....	10
Agent not connecting to Master.....	10
Agent port conflict (port 9090 already in use).....	11
CMA_AGENT service type not available in Cloudera Manager.....	11
Agent starts but cluster is not displayed in CMA Master UI.....	12
Authentication issues.....	13
Default credentials.....	13
JWT token issuer mismatch when calling CMA API remotely.....	13
Unresponsive commands.....	14
Scan or migration commands remain in IN_PROGRESS state.....	14
CMA upgrade issues.....	15
Liquibase checksum failure after upgrade.....	15
CMA Master slow to start.....	16
Database and data cleanup.....	16
Vault troubleshooting.....	17
Vault token file accidentally deleted or empty.....	18
Troubleshooting HDFS to Ozone migration.....	18

Troubleshooting Airflow DAG migration to Git [Technical Preview].....	20
--	-----------

Troubleshooting for Cloudera Migration Assistant

Diagnose and resolve issues with CMA Master, CMA Agent, upgrades, authentication, Vault, and diagnostics.

Troubleshooting procedures cover CMA Master, CMA Agent, and related operations. Issues are grouped by category. Expand each group to find the symptom you observed. For log paths and deployment directories, see *Log locations and key file paths*.

Related Information

[Log locations and key file paths](#)

Log locations and key file paths

Use these default log paths, parcel directories, and embedded database locations when diagnosing Cloudera Migration Assistant problems.

Related topics

For database topology and moving off embedded H2, see *CMA database backends*.

Parcel deployment (Cloudera Manager)

Table 1: Log paths for Cloudera Migration Assistant services in parcel deployment mode

Service	Log path
CMA Master Server	/var/log/cma/cma-server.log
CMA Agent Server	/var/log/cma-agent/cma-agent-server.log
CMA Auth Server	/var/log/cma/cma-auth-server.log
CMA Gateway Server	/var/log/cma/cma-gateway-server.log

Key file paths (parcel deployment)

Table 2: File paths in parcel deployment mode

Path	Description
/opt/cloudera/parcels/cma/	CMA parcel installation directory
/opt/cloudera/csd/	CSD JAR location
/var/lib/cma/data/db/	CMA Master databases (H2)
/var/lib/cma-agent/data/db/	CMA Agent databases (H2)
/var/lib/cma/data/backups/	Automatic database backups from upgrades
/var/lib/cma/conf/	CMA Master configuration
/var/log/cma/	CMA Master logs
/var/log/cma-agent/	CMA Agent logs

Related Information

[database backends](#)

Scan failures

Resolve scan failures such as Hive JDBC pool issues, ZooKeeper TLS port mismatches, Kerberos errors, Ozone Ranger permission denials, empty scan results, and scans stuck in IN_PROGRESS.

These scenarios apply to cluster scans run through the CMA UI. For scan workflow steps, see *Scanning clusters*.

Hive scan fails with "Failed to run async statement" after HiveServer2 restart

Restart the CMA Agent to clear stale Hive JDBC connections after HiveServer2 restarts.

About this task

Symptom: Hive scan fails immediately with HiveScannerException: Exception while querying database names - Exception message: Failed to run async statement. The scan was working before a HiveServer2 restart.

Cause: Cloudera Migration Assistant pools Hive JDBC connections for reuse across scans. When HiveServer2 is restarted, pooled connections become invalid because the server-side Thrift session no longer exists, but the client-side socket still appears open.

Procedure

1. Go to the CMA Agent service in Cloudera Manager.
2. Click Actions > Restart.
3. Retry the Hive scan.

Hive scan fails with "Unable to retrieve HS2 host information from ZooKeeper"

Hive scans fail with Unable to retrieve HS2 host information from ZooKeeper when zookeeper_tls_enabled routes JDBC discovery to the secure ZooKeeper port; set the property to false in Cloudera Manager so Cloudera Migration Assistant uses the standard client port.

Symptom

Hive table scan fails with:

```
HiveScannerException: Connection to Hive using URL jdbc:hive2://<host>:2182/;...;serviceDiscoveryMode=zooKeeper;zooKeeperNamespace=hiveserver2 failed.  
- Exception message: org.apache.hive.jdbc.ZooKeeperHiveClientException: U  
nable to retrieve HS2 host information from ZooKeeper
```

Cause

When zookeeper_tls_enabled is set to true in the ZooKeeper service configuration, Cloudera Migration Assistant connects to the secure ZooKeeper client port (default 2182) for HiveServer2 service discovery. The Hive JDBC driver does not pass the ZooKeeper TLS client parameters needed to negotiate TLS on that port, which results in a connection failure.

This failure commonly occurs on clusters where ZooKeeper has both zookeeper_tls_enabled set to true and zookeeper_client_portUnification set to true. With port unification enabled, the standard client port (2181) accepts both TLS and non-TLS connections, so using port 2182 is unnecessary.

Remedy

Procedure

1. In Cloudera Manager, go to the ZooKeeper service and open Configuration.
2. Search for `zookeeper_tls_enabled` and set it to false.

No service restart is required. Cloudera Migration Assistant reads this configuration dynamically when building the Hive JDBC connection URL.

3. Retry the Hive scan.

Cloudera Migration Assistant uses the standard ZooKeeper client port (2181) instead of the secure port (2182). When `zookeeper_client_portUnification` is enabled, port 2181 still accepts TLS connections from other services.

4. Verify which ZooKeeper ports are configured:

```
curl -su admin:admin "https://<cm-host>:7183/api/v58/clusters/<cluster>/services/ZOOKEEPER-1/config" | \
  python3 -c "import sys,json; d=json.load(sys.stdin); [print(i['name'], '=', i.get('value', '')) for i in d['items'] if 'tls_enabled' in i['name'] or 'client_port' in i['name'] or 'portUnification' in i['name']]"
```

Ozone scan fails with "PERMISSION_DENIED" on volume or bucket access

Ozone scans return `PERMISSION_DENIED` on Apache Ranger-managed clusters when the CMA Agent principal lacks read and list access; add the agent user to the global Ranger Ozone policy.

Symptom

Ozone scan fails with:

```
OzoneScanException: Error occurred while scanning volume s3v
- Exception message: PERMISSION_DENIED
```

The CMA Agent log shows an ACL check failure for the `cma-agent` user (or the configured Kerberos principal) when listing buckets or reading volume metadata.

Cause

On clusters where Apache Ranger manages Ozone authorization, native Ozone ACLs (set through `ozone sh volume addacl`) are ineffective. Ranger intercepts all access checks before native ACLs are evaluated, so granting access through `ozone sh` has no effect.

The Cloudera Migration Assistant Ozone scanner requires read and list permissions on volumes, buckets, and keys to inventory the cluster's Ozone storage.

Remedy

Procedure

1. Open the Ranger Admin UI (typically `https://<ranger-host>:6182`).
2. Click the Ozone service (for example, `cm_ozone`).
3. Find the policy named `all - volume, bucket, key` and click Edit.
4. In the Allow Conditions section, add a new entry:
 - Select User: `cma-agent` (or the Kerberos principal used by the CMA Agent)

- Permissions: read, list

5. Click Save.

The policy change takes effect within 30 seconds (Ranger's default policy refresh interval for plugins). No service restart is required.

6. Retry the Ozone scan.

7. Verify the Ranger Ozone policies:

```
curl -sk --negotiate -u : "https://<ranger-host>:6182/service/plugins/policies/service/name/cm_ozone" | \
python3 -c "import sys,json; [print(f' {p[\"name\"]}') for p in json.load(sys.stdin).get('policies', [])]"
```



Note:

If you do not have access to the Ranger UI (for example, when only Kerberos SPNEGO authentication is available), you can update the policy through the Ranger REST API using a host-based Kerberos keytab (for example, the rangeradmin keytab on the Ranger host).

```
kinit -kt /path/to/ranger.keytab rangeradmin/<hostname>@REALM
# GET the policy, add cma-agent to policyItems, then PUT it back
curl -sk --negotiate -u : "https://<ranger-host>:6182/service/plugins/policies/<policy-id>" -X PUT \
-H "Content-Type: application/json" -d @updated-policy.json
```

Scan stays in IN_PROGRESS indefinitely

When a scan never completes, cancel it from command history and retry. If cancel fails, see [stuck-command troubleshooting](#).

About this task

Cause: The CMA Agent lost connection to the CMA Master during the scan, or the scan process on the agent was killed.

Procedure

1. Cancel the scan from command history (click the history icon on the cluster page).
2. Retry the scan.
3. If cancel does not work, see *Scan or migration commands stuck in IN_PROGRESS state*.

Kerberos authentication error during scan

When a scan fails immediately with Kerberos authentication errors, verify Kerberos is enabled on the CMA Agent and the keytab is valid.

About this task

Symptom: Scan fails immediately with authentication errors in the CMA Agent log.

Procedure

1. Verify that `kerberos.auth.enable` is set to `true` in the CMA Agent service configuration in Cloudera Manager.

2. Verify that the keytab file is valid.

```
klist -kt /path/to/keytab
```

3. Retry the scan.

Scan completes but shows no results

When a scan finishes with empty results, verify the target service exists, scan paths contain data, and the scan user has read permissions.

About this task

Possible causes:

- The service being scanned is not deployed on the cluster (for example, scanning for HBase when no HBase service exists).
- The scan path is empty (for HDFS or Ozone scans, verify the configured scan path contains data).
- The scan user lacks permissions to read the target resources.

Procedure

1. Confirm the scanned service is deployed on the cluster.
2. For HDFS or Ozone scans, verify the configured scan path contains data.
3. Verify the scan user has read permissions on the target resources.
4. Retry the scan after correcting configuration or permissions.

Execution failures

Resolve migration execution failures such as steps stuck in IN_PROGRESS, Cloudera Replication Manager policy errors, and cancellation issues.

These scenarios apply to plan executions in the CMA UI.

Execution step stays in IN_PROGRESS

When a migration execution step remains in IN_PROGRESS, check the CMA Agent log for the underlying operation status before canceling and re-running.

About this task

Cause: The CMA Agent executing the step lost connectivity to the CMA Master, or the underlying operation (for example DistCp or a Cloudera Replication Manager policy) is still running.

Procedure

1. Check the CMA Agent log for the current step status.
Default log path: `/var/log/cma-agent/cma-agent-server.log`
2. If the underlying operation completed but status was not reported back, cancel the execution and create a new one.

Cloudera Replication Manager policy creation fails

Resolve execution failures when Cloudera Replication Manager policy creation returns HTTP 400 or 409 errors.

About this task

Symptom: An execution step fails with Failed to create replication policy or HTTP 400/409 from Cloudera Replication Manager.

For Cloudera Replication Manager setup and integration, see *Cloudera Replication Manager Integration*.

Procedure

1. Verify that Cloudera Replication Manager is running on the target cluster Cloudera Manager host.
2. If the policy name already exists, delete the existing policy through the Cloudera Manager API and retry.
See *Database and Data Cleanup* for listing and deleting replication policies.
3. Ensure the source cluster is registered as a peer in the target cluster's Cloudera Replication Manager configuration.
4. Verify that source paths in the migration plan exist on the source cluster.

Canceling a running execution

Cancel an active migration execution from the plan Execution tab. Cancellation is best-effort when DistCp or replication jobs are in progress.

About this task



Note: Cancellation is best-effort. If a DistCp job is in progress, it stops but partially copied data may remain on the target.

Procedure

1. Navigate to the plan's Execution tab.
2. Click the active execution.
3. Click Cancel.

CMA Agent issues

Resolve common CMA Agent problems such as port conflicts, connectivity to CMA Master, Cloudera Manager service visibility, and cluster registration delays.

These scenarios cover common CMA Agent problems when deployed as a Cloudera Manager add-on service. For deployment steps, see *Deploying the CMA server with parcels*.

Related Information

[Deploying the server with parcels](#)

Agent not connecting to Master

When the CMA Agent stays OUT_OF_SERVICE or never completes registration to CMA Master, check cma_gateway_url, TLS, firewalls, and the gRPC listener port on the Master.

About this task

Symptom: Agent health displays OUT_OF_SERVICE. The agent log shows connection errors or hangs at the Initializing connection to master message without a subsequent Sending agent instance confirmation message.

For gateway discovery and cross-cluster agent configuration, see *Deploying the CMA server with parcels*.

Procedure

1. Confirm that `cma_gateway_url` points to the CMA Gateway (for example `https://<master-host>:8093`), not the Master gRPC port directly.

The agent discovers the Master host, gRPC port, and Auth Server URI from the gateway endpoint (GET `/api/agent/config`).

2. Ensure outbound connectivity from the agent host to the gateway URL and to the discovered gRPC port (default 8091 for plaintext or 8891 for TLS).

From Bash on GNU/Linux you can test TCP reachability using `/dev/tcp`.

```
timeout 5 bash -c 'cat < /dev/tcp/<master-host>/8091' && echo OPEN
```

3. If the CMA Master uses TLS, set `cma_gateway_url` to `https://` and verify the agent can reach the TLS gRPC port.
4. If the gRPC port was customized during deployment, check the configured listener properties on the Master host.

```
grep grpc /var/run/cloudera-scm-agent/process/*-cma_master-CMA_SERVER/config/cma-server.properties
```

Related Information

[Deploying the server with parcels](#)

Agent port conflict (port 9090 already in use)

When port 9090 is already in use (often by HBase ThriftServer), change the CMA Agent HTTP and HTTPS ports in Cloudera Manager and restart the agent.

About this task

Symptom: The agent fails to start. The log displays the Web server failed to start. Port 9090 was already in use.error message.

Cause: Port 9090 commonly conflicts with HBase ThriftServer on the same host.

For complete parcel deployment procedures, see *Deploying the CMA server with parcels*.

Procedure

1. Open the CMA Agent service configuration in Cloudera Manager.
2. Set `cma_agent_http_port` to an unused port, for example, 9190.
3. Set `cma_agent_https_port` to 9193. Adjust the value if your site reserves a different HTTPS port.
4. Restart the CMA Agent service.

Related Information

[Deploying the server with parcels](#)

CMA_AGENT service type not available in Cloudera Manager

When `CMA_AGENT` does not appear as a service type, install both Cloudera Migration Assistant CSD files on the Cloudera Manager Server host and restart Cloudera Manager Server.

About this task

Cause: One or both Cloudera Migration Assistant CSD files (CMA_MASTER and CMA_AGENT) are not installed, or Cloudera Manager was not restarted after CSD installation.

For end-to-end parcel installation, see *Deploying the CMA server with parcels*.

Procedure

1. Download both CSD JAR files to /opt/cloudera/csd/ on the Cloudera Manager Server host.

```
wget -P /opt/cloudera/csd/ https://archive.cloudera.com/cma/VERSION/csd/CMA_MASTER-VERSION-BUILD.jar
```

```
wget -P /opt/cloudera/csd/ https://archive.cloudera.com/cma/VERSION/csd/CMA_AGENT-VERSION-BUILD.jar
```

2. Set ownership.

```
chown cloudera-scm:cloudera-scm /opt/cloudera/csd/CMA_MASTER-VERSION-BUILD.jar /opt/cloudera/csd/CMA_AGENT-VERSION-BUILD.jar
```

3. Set permissions.

```
chmod 640 /opt/cloudera/csd/CMA_MASTER-VERSION-BUILD.jar /opt/cloudera/csd/CMA_AGENT-VERSION-BUILD.jar
```

4. Restart the Cloudera Manager Server.

```
systemctl restart cloudera-scm-server
```

5. Wait for Cloudera Manager to restart (approximately 90–120 seconds), then verify that the CMA service types are available.

Related Information

[Deploying the server with parcels](#)

Agent starts but cluster is not displayed in CMA Master UI

When the CMA Agent is running but the cluster is not displayed in the CMA Master UI, verify registration timing, agent logs, and Cloudera Manager credentials.

About this task

Symptom: The CMA Agent starts successfully, but the cluster does not display in the CMA Master UI.

For more information on cluster registration in Cloudera Migration Assistant, see *Registering clusters*.

Procedure

1. Wait briefly after agent startup.

Allow about 30 seconds for the registration to finish.

2. In the agent log, look for the Sending agent instance confirmation message.
3. If the confirmation message does not appear, check the Cloudera Manager credentials configured for the agent.

The agent uses those credentials to query Cloudera Manager for cluster metadata.

Related Information

[Registering clusters](#)

Authentication issues

Review default login credentials and JWT issuer constraints when calling Cloudera Migration Assistant APIs from remote hosts.

Default credentials

Manage default credentials for evaluation installs. For production clusters, rotate passwords and configure the database users.

About this task

Cloudera Migration Assistant uses default administrative credentials and an embedded H2 database intended only for evaluation and testing. For production deployments, you must change all default passwords and configure custom, secure database user accounts.

Override embedded database credentials using the CMA_DBA_USER and CMA_DBA_PASS environment variables. For more information, see *CMA database backends*.

Procedure

Record the default accounts before you change them.

Table 3: Default credentials

Service	Default credentials	Notes
CMA Web UI	admin / admin	Login through the gateway on port 8090.
CMA Auth Server	admin / admin	Direct access on port 9000.
CMA Agent API	admin / admin	HTTP Basic authentication.
H2 database (server roles)	admin / s3cret	Override using the CMA_DBA_USER and CMA_DBA_PASS variables.

Related Information

[database backends](#)

JWT token issuer mismatch when calling CMA API remotely

Remote API calls can fail JWT validation when the Auth Server issues tokens with a localhost issuer.

About this task

Symptom: API calls return the WWW-Authenticate: Bearer error="invalid_token", error_description="The iss claim is not valid" error.

Cause: The Auth Server sets iss: http://localhost:<port> in the JWT. When you call the API from a different host, CMA Master rejects the token because the issuer does not match.

Procedure

Run API calls on the CMA Master host itself using SSH so localhost matches the issuer encoded in the token.

Unresponsive commands

Clear migration or scan commands that are unresponsive and remain in the IN_PROGRESS state when a cancellation request fails to update state.

These procedures modify the embedded H2 metadata store used by default deployments. For external databases and credential rotation, see *CMA database backends*.

Related Information

[database backends](#)

Scan or migration commands remain in IN_PROGRESS state

When a scan or migration command remains in the IN_PROGRESS state and a cancellation request fails to update the state, update the command rows directly in the embedded H2 database on the Master host.

About this task

Symptom: A scan or migration command remains in IN_PROGRESS indefinitely. Cancelling through the UI or API returns a success message but the state does not change.

Cause: A known issue prevents the cancel operation from persisting the state change to the database.

For database locations and moving to an external database, see *CMA database backends*.

Procedure

1. Plan which commands to reset.



Caution:

Run database updates only when you understand which commands are affected. Narrow WHERE clauses before updating production metadata.

2. On an embedded H2 parcel deployment, SSH to the Master host and run the update.

```
ssh root@<master-host>

# On the Master host:
H2_JAR=$(ls /opt/cloudera/parcels/cma/lib/h2-*.jar)
DB_URL="jdbc:h2:/var/lib/cma/data/db/cma-server;IFEXISTS=TRUE;AUTO_SERVER=TRUE"
java -cp $H2_JAR org.h2.tools.Shell -url "$DB_URL" -user admin -password s3cret \
  -sql "UPDATE CLUSTER_COMMAND_EXECUTION SET STATE = 'FAILED' WHERE STATE = 'IN_PROGRESS';"
```

To update only specific commands, use WHERE ID IN (9, 10, 11) instead of the WHERE STATE = 'IN_PROGRESS' clause.

Related Information

[database backends](#)

CMA upgrade issues

Recover from Liquibase checksum failures, long startups after upgrade, and coordinated database cleanup across Master and Agent roles.

For parcel lifecycle steps and restart ordering across clusters, see *Deploying the CMA server with parcels*. For database topology and Liquibase expectations, see *CMA database backends*.

Related Information

[database backends](#)

[Deploying the server with parcels](#)

Liquibase checksum failure after upgrade

After a CMA parcel upgrade, service roles might report bad health when Liquibase detects checksum mismatches against a lower embedded schema version.

About this task

Symptom: After a Cloudera Migration Assistant parcel upgrade, one or more roles display healthSummary: BAD. The application log contains the following validation exception:

```
liquibase.exception.ValidationFailedException: Validation Failed:  
  N changesets check sum
```

Cause: The new version modified database schema files. The existing database stores checksums from the previous build that no longer match.

For more information on how Liquibase runs across embedded H2 and external databases, see *CMA database backends*.

Determine whether you need to preserve your existing metadata or wipe embedded databases.

Production clusters require data preservation, which you achieve by resetting checksum rows. Development or throwaway environments might tolerate full wipes.

Procedure

1. To preserve metadata, reset Liquibase checksum entries.



Note:

Use this procedure only when you must keep existing plans and metadata.

- a. Stop the failing service role.
- b. Replace the sample IDs with the failing changesets reported by Liquibase.

```
java -cp /opt/cloudera/parcels/cma/lib/h2-*.jar org.h2.tools.Shell \  
-url 'jdbc:h2:/var/lib/cma/data/db/cma-server;IFEXISTS=TRUE;AUTO_SERV  
ER=TRUE' \  
-user admin -password s3cret \  
-sql "UPDATE DATABASECHANGELOG SET MD5SUM = NULL WHERE ID IN ('<change  
set-id-1>', '<changeset-id-2>');"
```

- c. Restart the service role. On startup, Liquibase recalculates the checksums.
2. To wipe embedded databases (destructive), remove all local database files for CMA Master, CMA Auth Server, and CMA Agents.

**Warning:**

This action permanently deletes all Cloudera Migration Assistant data, including migration plans, execution history, cluster registrations, and scan results.

Use this procedure only if the embedded database records can be safely destroyed.

- a. Stop all Cloudera Migration Assistant services. You must stop CMA Agents first, followed by the CMA Master.
- b. On the CMA Master host delete the server and authentication database files.

```
sudo -u cma rm -f /var/lib/cma/data/db/cma-server.* /var/lib/cma/data/db/cma-auth.*
```

- c. On each CMA Agent host clear the local agent database files.

```
rm -f /var/lib/cma-agent/data/db/cma-agent.*
```

- d. Start CMA Master first, then start every CMA Agent service.

**Important:**

If you wipe the CMA Master database, you **must** also wipe all CMA Agent databases. CMA Agents store references to cluster and credential records that no longer exist in a fresh CMA Master database.

3. When upgrading multiple clusters, restart CMA Agents before restarting the CMA Master.

For deployment topology guidance, see *Deploying the CMA server with parcels*.

Related Information

[database backends](#)

[Deploying the server with parcels](#)

CMA Master slow to start

Resolve a slow long CMA Master startup immediately after upgrade by tracking Liquibase migration work against the metadata database.

About this task

Symptom: CMA Master takes a long time to start or appears unresponsive after startup.

Cause: Liquibase database migration runs on first startup after an upgrade, which can take about 60–90 seconds.

For database expectations and external JDBC backends, see *CMA database backends*.

Procedure

1. Wait for startup to finish instead of restarting repeatedly.
2. Monitor the CMA Master application log for progress.

Default path: `/var/log/cma/cma-server.log`. For other log paths, see *Log locations and key file paths*.

Related Information

[database backends](#)

[Log locations and key file paths](#)

Database and data cleanup

Use these H2 defaults and cleanup ordering when purging migration metadata or replication policies.

About this task



Note: Do not use this procedure for production clusters. Production deployments override embedded defaults using the CMA_DBA_USER and CMA_DBA_PASS environment variables to target external JDBC backends.

The embedded defaults apply unless you have configured external JDBC backends. For more information, see *CMA database backends*. For Cloudera Replication Manager integration context, see *Cloudera Replication Manager Integration*.

Table 4: Embedded H2 database default credentials

Role	Database path	Default credentials
CMA Master Server	/var/lib/cma/data/db/cma-server	admin / s3cret
CMA Auth Server	/var/lib/cma/data/db/cma-auth	root / root
CMA Agent Server	/var/lib/cma-agent/data/db/cma-agent	admin / s3cret

Procedure

1. List Cloudera Replication Manager replication policies on the target cluster using the Cloudera Manager API.

```
curl -s -u admin:admin \
  "http://<target-cm-host>:7180/api/v49/clusters/Cluster%201/services/HDFS-1/replications"
```

2. Delete each policy through the Cloudera Manager API.

```
curl -s -X DELETE -u admin:admin \
  "http://<target-cm-host>:7180/api/v49/clusters/Cluster%201/services/HDFS-1/replications/<id>"
```



Note:

Do not use the Cloudera Replication Manager API (/dmx/api/policies) for deletion, because it fails with HTTP 404 error due to pipe characters in policy IDs.

3. Clean the H2 database on the CMA Master host.

Disable referential integrity temporarily when untangling foreign-key chains.

4. Restart CMA Master.

The CMA Master repopulates essential plugin data from YAML configuration files at startup.

Related Information

[database backends](#)

[Replication Manager Integration](#)

Vault troubleshooting

Recover Vault token files and correlate logs when secrets management misconfiguration blocks startup.

For administrators configuring Vault with Cloudera Migration Assistant, see *Storing secrets in Vault*.

Related Information

[Storing secrets in Vault](#)

Vault token file accidentally deleted or empty

Restore the Vault token from the backup that the startup script creates after first successful initialization.

About this task

After the first successful initialization, the startup script creates a backup of the token file in your home folder named `.cma-server_vault`.

For enabling Vault with Cloudera Migration Assistant, see *Storing secrets in Vault*.

Procedure

1. Replace the active Vault token file with the backup on the CMA Server host.

```
cp ~/.cma-server_vault /var/lib/cma/conf/.vault
```

2. Restart the Vault server and all dependent Cloudera Migration Assistant roles according to your deployment checklist.

Related Information

[Storing secrets in Vault](#)

Troubleshooting HDFS to Ozone migration

Interpret common Cloudera Replication Manager failures, scan gaps, and authentication errors when migrating HDFS data into Apache Ozone with Cloudera Migration Assistant.

Symptom

No directories appear after HDFS scans.

Cause

Overly narrow path filter or aborted scan commands.

Remedy

Procedure

1. Review scan logs.
2. Rerun the scan with broader paths (often `/`) until inventories refresh.

Symptom

Ozone scan exposes zero volumes.

Cause

Ozone service offline or degraded.

Remedy

Verify Ozone availability in Cloudera Manager before resubmitting scans.

Symptom

Structure mapping rejects configuration.

Cause

Destination Ozone nodes missing from deployed layout.

Remedy

Create volumes, buckets, and folders described in *Managing Ozone storage*, then redeploy the layout and remap.

Symptom

Execution fails at Kerberos checkpoints.

Cause

Expired principals or mismatched CMA Agent keytabs.

Remedy

Refresh Kerberos configurations for each CMA Agent, then restart the agents.

Symptom

Cloudera Replication Manager errors mid-run.

Cause

Unreachable Cloudera Replication Manager endpoint or stale credentials stored in CMA Master.

Remedy

Validate Cloudera Replication Manager URLs and secrets on CMA Master, then inspect Master logs for stack traces preceding the failure.

Symptom

Partial datasets under Ozone.

Cause

Skipped files flagged by DistCp owing to POSIX permissions.

Remedy

Review Cloudera Replication Manager job logs for skip reasons, then widen source read permission for the migration role.

Related Information

[Managing Ozone storage](#)

[Migrating HDFS data to Ozone](#)

Troubleshooting Airflow DAG migration to Git [Technical Preview]

Interpret common failures, scan gaps, Git publishing failures, and authentication errors when migrating Cloudera Data Engineering Airflow DAG assets into Git-backed control planes with Cloudera Migration Assistant.

Before you begin



Note: This feature is in Technical Preview and is not ready for production deployments. Cloudera recommends trying this feature in test or development environments and encourages you to provide feedback on your experiences.

Each resolution builds on *Setting up to Cloudera Workflow Orchestrator Git migration* prerequisites when credentials or Git provider routes fail.

Symptom

Discovery is empty after scans.

Cause

Incomplete command state or restrictive collection filters masking rows.

Remedy

Procedure

1. Validate command transcripts.
2. Temporarily clear filters to confirm that the artifacts exist.

Symptom

Virtual clusters list zero DAGs.

Cause

Knox authentication halted mid-scan (SKIPPED log entries).

Remedy

Work with administrators to reaffirm Cloudera Manager credentials powering the CMA Agent.

Symptom

Duplicate dag_id badges appear.

Cause

Separate virtual clusters authored identical identifiers.

Remedy

Rename offending DAG IDs on the Workloads tab editor in your DAG migration plan.

Symptom

Execution phase 8.1 faults occur.

Cause

Stale local artifact caches on CMA Agents.

Remedy

Rescan Cloudera platform workspaces, then regenerate executions.

Symptom

Git step fails authentication.

Cause

Expired personal access tokens or underscoped PAT scope.

Remedy

Administrators rotate CMA_AGENT_CLUSTER_GIT_TOKEN in the advanced configuration snippet and recycle CMA Agent services.

Symptom

Git rejects base branch selection.

Cause

CMA_AGENT_CLUSTER_GIT_BRANCH was deleted upstream.

Remedy**Procedure**

1. Restore branch names or refresh configuration.
2. Restart CMA Agent services once secrets synchronize.

Symptom

Pull Request links are missing after the run.

Cause

Upstream provider API errors swallowed in CMA Master diagnostics.

Remedy

Inspect Master logs with Git webhook availability from CMA Agent networks.

Related Information

[Setting up to Git migration \[Technical Preview\]](#)

[Migrating Airflow DAGs from to a Git repository \[Technical Preview\]](#)