

Managing Apache ZooKeeper

Date published: 2019-08-21

Date modified: 2021-09-09



Legal Notice

© Cloudera Inc. 2025. 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

Add a ZooKeeper service.....	4
Use multiple ZooKeeper services.....	4
Replace a ZooKeeper disk.....	5
Replace a ZooKeeper role with ZooKeeper service downtime.....	5
Replace a ZooKeeper role without ZooKeeper service downtime.....	6
Replace a ZooKeeper role on an unmanaged cluster.....	7
Confirm the election status of a ZooKeeper service.....	7

Add a ZooKeeper service

If the Add Service wizard does not finish successfully, you can add the ZooKeeper service manually.

About this task

In a production environment, you should deploy ZooKeeper as an ensemble with an odd number of servers. As long as a majority of the servers in the ensemble are available, the ZooKeeper service will be available. The minimum recommended ensemble size is three ZooKeeper servers, and Cloudera recommends that each server run on a separate machine. In addition, the ZooKeeper server process should have its own dedicated disk storage if possible.



Note: If a customer has experienced fsync delays and other I/O related issues with ZooKeeper, ZooKeeper's dataDir and dataLogDir can be configured to use separate disks. It is hard to determine ahead of time whether this will be necessary; even a small cluster can result in heavy ZooKeeper activity.

When adding the ZooKeeper service, the Add Service wizard automatically initializes the data directories.

When you add ZooKeeper servers to an existing ensemble, a rolling restart of all ZooKeeper is required in order to allow all zooKeeper servers to have the same configurations.

If you quit the Add Service wizard or it does not finish successfully, you can initialize the directories outside the wizard by following these steps:

Procedure

1. In Cloudera Manager, select the ZooKeeper service.
2. Click Action.
3. Select Initialize.
4. Click Initialize again to confirm.

If the data directories are not initialized, the ZooKeeper servers cannot be started.

Use multiple ZooKeeper services

Only services that are not dependent can use different ZooKeeper services.

Cloudera Manager requires dependent services within Cloudera Runtime to use the same ZooKeeper service. If you configure dependent Cloudera Runtime services to use different ZooKeeper services, Cloudera Manager reports the following error:

```
com.cloudera.cmf.command.CmdExecException: java.lang.RuntimeException: java.l
ang.IllegalStateException: Assumption violated:
getAllDependencies returned multiple distinct services of the same type
at SeqFlowCmd.java line 120
in com.cloudera.cmf.command.flow.SeqFlowCmd run()
```

Because of a known limitation in Cloudera Manager, Cloudera Runtime services that are depending on each other can not use different ZooKeeper services. For example, Kafka does not depend on any services other than ZooKeeper, so Kafka can use one ZooKeeper cluster and the rest of the Cloudera Runtime services can use another one. At the same time, since Ranger also depends on ZooKeeper, if Kafka gets configured to use Ranger, then both Kafka and Ranger must use the same ZooKeeper cluster.

Replace a ZooKeeper disk

You can replace a ZooKeeper Disk but first you must stop the ZooKeeper role.

About this task

Procedure

1. In Cloudera Manager, select the ZooKeeper service.
2. Click the Configuration tab.
3. Search for data directory.
4. Find the Data Directory property and update as applicable.
5. Find the Transaction Log Directory property and update as applicable.
6. Click Save Changes.
7. Click the Instances tab.
8. Select one ZooKeeper role.
9. Click the Actions for Selected button.
10. Select Stop.
11. Move the contents to the new disk location and modify mounts as needed.
Ensure that the permissions and ownership are correct.
12. Select the same ZooKeeper role.
13. Click the Actions for Selected button.
14. Select Start.
15. Repeat step 8-14 for any remaining ZooKeeper roles.

Replace a ZooKeeper role with ZooKeeper service downtime

Replace a ZooKeeper role with ZooKeeper service downtime.

Procedure

1. In Cloudera Manager, select the ZooKeeper service.
2. Click the Instances tab.
3. Select the ZooKeeper role on the old host.
4. Click the Actions for Selected button.
5. Select Stop.
6. Click the Actions for Selected button.
7. Select Delete.
8. Click the Add Role Instances button.
9. Select the new host.
10. Click Continue.
11. Confirm the changes.
12. Click Continue.

13. Restart the old ZooKeeper servers that have outdated configuration.
14. Select the ZooKeeper service.
15. Click the Status tab.
16. Confirm the ZooKeeper service elected one of the restarted hosts as a leader.
17. Restart the newly added ZooKeeper server.
18. Restart any dependent services, such as HBase, HDFS, YARN or Hive, that are marked to have stale configuration.

Replace a ZooKeeper role without ZooKeeper service downtime

If server to server SASL authentication is not enabled, you can replace a ZooKeeper role without ZooKeeper service downtime.

About this task

This process is valid only if the SASL authentication is not enabled between the ZooKeeper servers. You can check this in Cloudera Manager, if you see the Enable Server to Server SASL Authentication under the Configuration of ZooKeeper.

Procedure

1. In Cloudera Manager, select the ZooKeeper service.
2. Click the Instances tab.
3. Select the ZooKeeper role on the old host.
4. Click the Actions for Selected button.
5. Select Stop.
6. Wait for 1 minute.
7. Click the Status tab.
8. Confirm that ZooKeeper service elected one of the remaining hosts as a leader.
9. Click the Instances tab.
10. Click the Add Role Instances button.
11. Select the new host.
12. Click Continue.
13. Confirm the changes.
14. Click Continue.
15. Click the newly added Server instance.
16. Click the Configuration tab.
17. Find the ZooKeeper Server ID property.
18. Change its value, so that it is the highest ZooKeeper Service ID set in the cluster.
19. Delete the old ZooKeeper role.
 - a) Go back to the ZooKeeper cluster.
 - b) Click Instances.
 - c) Select the old ZooKeeper role.
 - d) Click the Actions for Selected button.
 - e) Click Delete.
20. Click the Status tab.

21. Click Actions.
22. Select Start this Server.
The new ZooKeeper role is started.
23. Restart the ZooKeeper server that has an outdated configuration and is a follower.
24. Restart the leader ZooKeeper server that has an outdated configuration.
25. Confirm that a leader is elected after the restart, and the whole ZooKeeper service is in green state.
26. Restart any dependent services, such as HBase, HDFS, YARN or Hive, that are marked to have stale configuration.

Replace a ZooKeeper role on an unmanaged cluster

You can replace a ZooKeeper Role on an unmanaged cluster, using command line.

About this task

These instructions assume you are using ZooKeeper from the command line.

Procedure

1. Stop the ZooKeeper role on the old host.
2. Confirm the ZooKeeper Quorum elected a leader.
3. Add a new ZooKeeper role on the new server.
4. Identify the dataDir location from the zoo.cfg file.
This default to /var/lib/zookeeper.
5. Identify the ID number for the ZooKeeper Server from the myid file in the configuration: `cat/VAR/LIB/ZOOKEEPER/myid`.
6. On all the ZooKeeper hosts, edit the zoo.cfg file, so the server ID references the new server hostname.
For example:

```
server.1=zk1.example.org:3181:4181 server.2=zk2.example.org:3181:4181 se  
rver.4=zk4.example.org:3181:4181
```

7. Restart the ZooKeeper hosts.
8. Confirm the ZooKeeper Quorum elected a leader and the other hosts are followers.
9. Restart any dependent services, such as HBase, HDFS Failover Controllers with HDFS High Availability, or YARN with High Availability.
10. Perform a failover to make on HDFS NameNode active.

Confirm the election status of a ZooKeeper service

You can determine the election status of a ZooKeeper host if you have installed telnet or netcat and it runs from a host with network access to the ZooKeeper host.

The default ZooKeeper client port is 2181. Run the following command against each ZooKeeper host:

```
echo "stat" | nc server.example.org 2181 | grep Mode
```

For example, a follower host would return the message:

```
Mode: follower
```

You can use telnet, if you prefer:

```
$ telnet server.example.org 2181
```

Sample output would be similar to the following:

```
Trying 10.1.2.154...
Connected to server.example.org.
Escape character is '^]'.
stat
Zookeeper version: 3.4.5-cdh5.4.4--1, built on 07/06/2015 23:54 GMT
...

Latency min/avg/max: 0/1/40
Received: 631
Sent: 677
Connections: 7
Outstanding: 0
Zxid: 0x30000011a
Mode: follower <-----
Node count: 40
Connection closed by foreign host.
```