



This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

Cassandra Installation and Configuration Guide

Useful Tools

12/13/2025

Contents

- 1 Useful Tools
 - 1.1 cassandra-cli
 - 1.2 nodetool
 - 1.3 JConsole
 - 1.4 Resetting Cassandra to Initial State
 - 1.5 Routine Node Repair

Useful Tools

The Cassandra archive includes several helpful tools for viewing and configuring Cassandra. Most of which can be found in the %CASSANDRA_HOME%\bin directory.

cassandra-cli

NOTE: This tool is available only for Cassandra versions up to, and including, version 2.1.x. To use this tool, install a separate instance of Cassandra version 2.1.x. There is no need to configure or start this instance, this will be used only for `cassandra-cli` invocation. The defaults in the `yaml`, etc, can be left unchanged. For more information on the Cassandra-CLI utility, refer to the [Cassandra-CLI utility](#) documentation.

nodetool

The `nodetool` utility provides a couple of helpful features. The most important of which is the ability to view the status of a Cassandra cluster, for example:

```
C:\Cassandra\apache-cassandra-2.2.5\bin>nodetool.bat -h dwswin7 status
Starting NodeTool
Datacenter: datacenter1
=====
Status=Up/Down
|/ State=Normal/Leaving/Joining/Moving
-- Address          Load       Tokens     Owns    Host ID                               Rack
UN 172.21.82.83      267.79 KB  256 ?      c3af2e2e-ed72-4f96-8bce-f6f6f9c9eb98 rack1
UN 172.21.83.74      320.89 KB  256 ?      788429a3-c7f9-4263-aabd-21be37cc3e30 rack1
```

Refer to the following DataStax website for more information:

<https://docs.datastax.com/en/cassandra/2.1/cassandra/tools/toolsTOC.html>

JConsole

Available in the `jdk` bin directory. For example: `C:\Program Files\Java\jdk1.8.0_73\bin\jconsole.exe`.

After opening, select the remote process button and enter the `ip:jmxport` of the cassandra node. The `jmx` console will display heap memory usage, live threads, classes loaded, and CPU usage. Check out the `MBeans` tab for info regarding the Cassandra exposed information and tools. Multiple nodes may be monitored with one console - just select connection, new connection, and enter the Cassandra `ip:jmx port` for the new node.

Resetting Cassandra to Initial State

In some cases, due to errors in SXML scripts being exercised, there may be sessions which are started and never terminated. Depending on interaction volume, this may cause the Cassandra data to exceed the planned limits of the deployment. In this case, after removing issues with the SCXML scripts, which cause unterminated sessions:

- Stop all Orchestration nodes.
- Then stop all Cassandra nodes.
- Remove all entries in the yaml-defined paths for the data, saved_caches, and commitlog directories.
- Once this is complete for all Cassandra nodes in cluster, restart the Cassandra nodes, and reload the schema. Or allow Orchestration to load the schema (Orchestration versions 8.1.3 or later).

Routine Node Repair

On a production cluster, it is important to routinely run the Nodetool Repair command. Use this command to repair inconsistencies between nodes within a cluster. Stale data is updated by pulling the latest version from other nodes. Unless Cassandra applications perform no deletes at all, it is recommended to run scheduled repairs on all nodes at least once every `gc_grace_seconds`. If this procedure is not followed, deletes may be "forgotten" in the cluster following garbage collection.

Important

Nodetool Repair is a resource-intensive task and should be scheduled for low-usage hours. Avoid running nodetool repair on more than one node at a time. Only repair is required, compact is not recommended.

The following is an example for Windows.

```
nodetool.bat -h dwswin7 repair Orchestration
```

If Not Run within GCGraceSeconds

If Nodetool Repair is not run within GCGraceSeconds (default is 10 days), then you run the risk of forgotten deletes. This may lead to inconsistencies in the data returned by different nodes. Running Nodetool Repair will not correct the issue entirely. There are two recommended methods of dealing with this scenario:

1. Treat the node with inconsistent data as "failed" and replace it.
2. To minimize forgotten deletes, increase GCGraceSeconds for all Column Families via the CLI, perform a full repair on all nodes, and then change GCGraceSeconds back again.