

Fast autoscaling in Cloudera operational database (Preview)

Date published: 2023-12-18

Date modified: 2024-10-01

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.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Contents

Legal Notice	2
Contents	3
Fast autoscaling in Cloudera operational database	4
Introduction to Compute Instance Group	4
Metrics categorization	5
Fast autoscaling mechanism	6
Enabling fast autoscaling in COD	6
Disabling fast autoscaling in COD	8
Runtime upgrade of the cluster	8

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Fast autoscaling in Cloudera operational database

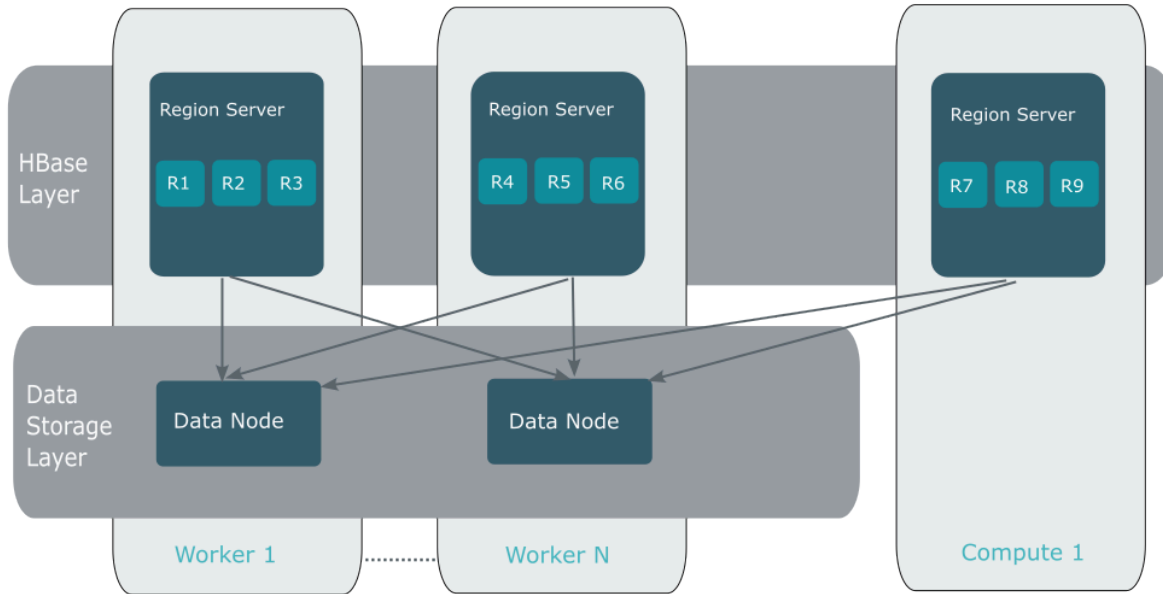
Cloudera Operational Database (COD) supports fast autoscaling for higher computing requirements. COD enables fast autoscaling whenever high CPU utilization or higher RPC latency is observed in the system.

The fast autoscaling is achieved using the following mechanisms.

Introduction to Compute Instance Group

A Compute node instance group in COD consists of nodes that host all the services, similar to those on Worker nodes, except the HDFS data services. The region servers on these nodes communicate with HDFS (data nodes) running on the Worker nodes.

The Compute nodes add to the system's higher compute capacity. During scale-up operations due to higher compute requirements, the Compute nodes are scaled up using the same storage as the existing Worker nodes. Hence, during the scale-down of these Compute nodes, the data remains in the Worker nodes and does not require data redistribution. The Worker nodes continue to host the data services and are scaled (up or down) according to the data requirements (such as Data storage and Regions per server).



Compute

<input type="checkbox"/>	Instance ID	Status	FQDN
>	<input type="checkbox"/> i-0c8a437905093ec9d	● Stopped	cod-zfo9yuhetvix-compute2.jrh19-co.xcu2-8y8x.dev.cldr.work
>	<input type="checkbox"/> i-0c7d468b3b29aa1e7	● Stopped	cod-zfo9yuhetvix-compute0.jrh19-co.xcu2-8y8x.dev.cldr.work
>	<input type="checkbox"/> i-093b3c2b99fb1f45d	● Running	cod-zfo9yuhetvix-compute1.jrh19-co.xcu2-8y8x.dev.cldr.work

Metrics categorization

The metrics used for autoscaling are categorized into the following groups: The data metrics scale the Worker instance group, while the compute metrics scale the Compute instance group.

- Worker Group (Data-related metrics)
 - a. HDFS usage
 - b. Region density (Regions per server)
 - c. Total Store File size across region servers
- Compute Group (Compute-related metrics)
 - a. CPU utilization
 - b. RPC latency

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cludera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Fast autoscaling mechanism

COD relies on suspending and resuming (start and stop) autoscaling, where the maximum allowed number of nodes in the system are instantiated but kept in a Stopped state. This mechanism is used for the scaling of the Compute instance group.

When a requirement for higher compute power arises in the cluster, COD starts the required number of Compute nodes. The scale-up time is drastically reduced because the Compute nodes are already initialized. Similarly, when the requirement for computing goes down, the Compute nodes are stopped but not dropped. They remain in the cluster in the Stopped state. The scale-up and scale-down time is reduced since the instance initialization is skipped during every scale-up or scale-down operation.

COD supports the suspending and resuming of fast autoscaling for the Compute nodes only. The Worker nodes continue to use the existing autoscaling mechanism.

For more information, see the following topics:

- [Autoscaling clusters](#).
- [Faster Auto-scaling for Higher Computing Requirements](#)

Enabling fast autoscaling in COD

Learn how to enable the fast autoscaling mechanism in COD.

Before you begin

- The cluster must have the COD setup in the environment.
- The node on which the steps are executed must have the CDP CLI installed.
- You must have the `COD_USE_COMPUTE_ONLY_NODES` entitlement.
- You must use COD 1.37.0 or a higher version.

Steps

1. Launch the CDP CLI tool.
2. Use the CDP CLI `create-database` command to create the compute instance group. Use the `--auto-scaling-parameters` option to specify the maximum and minimum number of compute nodes.

```
cdp opdb create-database -environment-name <env_name>
-database-name <db_name> --auto-scaling-parameters
```

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

```
'{"minComputeNodesForDatabase":<min_compute_nodes>,
"maxComputeNodesForDatabase": <max_compute_nodes>}'
```

Where,

- `<max_compute_nodes>` specifies the maximum number of Compute nodes required in the system.
- `<min_compute_nodes>` specifies the minimum number of Compute nodes in the system.

When the command is successfully executed, the newly created database contains the instantiated `<max_compute_nodes>` nodes. However, in the subsequent, scaling-executor cycle, the Compute nodes are stopped until the required number of nodes or the count `<min_compute_nodes>` is reached.

Alternatively, if the database is already created on 1.37.0, we can use the update database CDP CLI command to update the autoscaling parameters mentioned above.

3. Configure the autoscaling parameter to have an appropriate `maxCpuUtilization` which represents the average CPU utilization percent across Worker nodes and Compute nodes.

```
cdp opdb update-database --environment-name <env_name>
--database-name <db_name> \
--auto-scaling-parameters '{"maxCpuUtilization": 80. . . }'
```

The subsequent scaling cycles start or stop the Compute nodes as per the compute requirements.

4. Execute the update-database CDP CLI command to start all the Compute nodes in the system, if the cluster needs to be stopped or upgraded.

```
cdp opdb update-database --environment-name <env_name>
--database-name <db_name> \
--auto-scaling-parameters
'{"minComputeNodesForDatabase":<max_compute_nodes>}'
```

All the Compute nodes are started. The nodes are running after the cluster upgrade or restart.

5. Revert to the original `<min_compute_node>` after the restart or upgrade.

```
cdp opdb update-database --environment-name <env_name>
--database-name <db_name> \
--auto-scaling-parameters
'{"minComputeNodesForDatabase":<min_compute_nodes>}'
```

The subsequent scaling cycles stop the additional Compute nodes in the cluster.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

For more information, see [CDP CLI beta documentation](#).

Disabling fast autoscaling in COD

You can disable the fast autoscaling in your COD cluster. However, the existing autoscaling mechanism operates as usual in the cluster.

Before you begin

- The cluster must have the COD setup in the environment.
- The node on which the steps are executed must have the CDP CLI installed.
- Fast autoscaling must be enabled in your cluster.

Steps

1. Launch the CDP CLI tool.
2. Set the minimum and maximum Compute node counts to zero. This command deletes all the Compute nodes in the cluster.

```
cdp opdb create-database --environment-name <env_name>
--database-name <db_name> --auto-scaling-parameters
'{"minComputeNodesForDatabase":0, "maxComputeNodesForDatabase":
0}'
```

The subsequent autoscaling executions only impact the Worker instance group.

For more information, see [CDP CLI beta documentation](#).

Runtime upgrade of the cluster

Learn how to resume the stopped nodes in the cluster before performing a runtime upgrade.

The cluster upgrade might fail if the cluster contains nodes that are in the Stopped state. You must restart these nodes before executing the upgrade or drop the stopped nodes before the upgrade. They can be added later after the upgrade.

Ensure that you execute the following command to increase the minimum number of nodes in the cluster to the maximum number of allowed Compute nodes in the cluster. This triggers the restart of the stopped nodes. After that, the cluster can be upgraded.

```
./clients/cdpcli/cdp.sh opdb update-database --environment-name
jrh16-cod-7216 --database-name jrh13-7216 \
```

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

CLOUDERA TECHNICAL PREVIEW DOCUMENTATION

```
--auto-scaling-parameters  
'{"minComputeNodesForDatabase":<MAX_COMPUTE_NODES>}'
```

Revert the `minComputeNodesForDatabase` to the original value after the upgrade.

Alternatively, drop all the stopped Compute nodes in the cluster using the following CDP CLI command.

```
cdp opdb create-database -environment-name <env_name> -database-name  
<db_name> --auto-scaling-parameters '{"maxComputeNodesForDatabase":  
<number_of_running_compute_nodes_in_the_system>}'
```

For more information, see [CDP CLI beta documentation](#).