

AWS Graviton instances in Cloudera Operational Database (Preview)

Date published: 2024-12-11

Date modified: 2025-03-01

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
Overview	4
Prerequisites for using AWS Graviton in COD	4
Considerations for using AWS Graviton in your cluster	4
Creating AWS Graviton-based COD clusters	5

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.

Overview

AWS Graviton is a general-purpose, ARM-based processor family that currently delivers the best price-to-performance ratio for cloud workloads running in AWS Elastic Compute Cloud (EC2). With AWS Graviton, you can optimize costs.

Prerequisites for using AWS Graviton in COD

- To use this feature, you must have the following entitlements:
 - COD_USE_GRAVITON
 - CDP_AWS_ARM_DATAHUB
 - COD_USE_CUSTOM_INSTANCE_TYPES
- You must be logged into the COD as an **ODAdmin**.
- Verify if the selected instance types are available in your preferred regions

Considerations for using AWS Graviton in your cluster

Important:

This feature is under technical preview. Technical previews are considered under development. Do not use these features in production environments.

- You can create the Graviton-based COD clusters only on the AWS environments.
- The supported storage type is S3 with ephemeral cache enabled.
- The supported scale type is HEAVY.
- AWS is still rolling out Graviton, this feature is only available in some regions.
- Deploying COD on Graviton might take longer based on the cluster shape and size.
- To use COD on Graviton, you need to install beta CDP CLI. For more information, see the [Installing beta CDP CLI documentation](#).

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.

Creating AWS Graviton-based COD clusters

You can create AWS Graviton-based COD clusters using the Cloudera Operational Database UI or CLI command.

Using the Cloudera Operational Database UI:

1. In the COD web interface, click **Create Database**.
2. Specify the location of the database where you want to store it.
 - a. Provide a name for the database in the **Database Name** field.
 - b. Select the CDP environment from the list in which you want to associate the database. The environment does not have to be Graviton-based; it has to use a CDP version that supports COD on Graviton.
 - c. Click **Next**.

If an environment does not exist, you can create one by clicking **Create New Environment**.

For more information, see [Register your first environment](#).

3. Select ARM64 as the **Architecture** from the drop-down list.
4. Commission your database by defining a scale for your database using a predefined Data Lake template.

Select **Heavy Duty**.

The template helps you to structure your database automatically thereby saving your time and cost. COD creates the predefined number of LITE or HEAVY gateway and master nodes, a set of worker nodes, and also adds additional functionalists into the new database. In case you need to modify the default number of nodes defined in the template, you can do so after the database creation.

The available templates are Micro Duty, Light Duty, and Heavy Duty. By default, Light Duty is selected.

5. Configure your database by selecting the storage type as **Cloud Storage with Caching**.

The storage type Cloud Storage with Caching is equivalent to using `--storage-type CLOUD_WITH_EPHEMERAL` option on CDP CLI while creating an operational database.

6. Check or update the settings for your database.
 - a. Check all the default settings for your database under the **Default** tab.
 - b. Go to the **Advanced** tab to modify any of the default values.
 - The HDFS Volume Type option appears under the Advanced tab only if you select HDFS as the storage type in the Configuration step.

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.

- If you disable the Autoscaling option using the Advanced tab, the Worker Nodes and Compute Nodes options are hidden. Instead, a Node Count option appears.
 - c. The minimum and maximum number of worker nodes vary for different storage types.
 - Micro duty: Minimum node count: 1. Maximum node count: 5.
 - Light duty: Minimum node count: 3. Maximum node count: 100.
 - Heavy duty: Minimum node count: 3. Maximum node count: 800.
7. Review the details before creating the database.
 Click **Show CLI Command** to get the complete command details corresponding to your settings. You can use it to create the database using CDP CLI.
 Alternatively, you can use the following sample command to create the database using CDP CLI.

```
cdp opdb create-database --environment-name cod-731
--database-name test --architecture ARM64 --scale-type HEAVY
--storage-type CLOUD_WITH_EPHEMERAL --auto-scaling-parameters
'{"minWorkersForDatabase":5, "maxWorkersForDatabase":100}'
--num-edge-nodes 0
```

8. Click **Create Database**.

Result

An information page shows the database's status. Once it is available, your new database is ready to use.

Using the beta CLI command:

1. Launch the beta CDP CLI tool.
2. Run the following command to create Graviton-based COD clusters.

```
cdp opdb create-database --environment-name <env_name>
--database-name <database_name> --architecture <string>
--storage-type CLOUD_WITH_EPHEMERAL --scale-type HEAVY
```

For example,

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

```
cdp opdb create-database --environment-name aws_test  
--database-name cod_7214 --architecture ARM64 --storage-type  
CLOUD_WITH_EPHEMERAL --scale-type HEAVY
```

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.