

Cloudera Operational Database ..

## Cloudera Operational Database Top Tasks

Date published: 2022-09-27

Date modified: 2023-01-12

# CLOUDERA

<https://docs.cloudera.com/>

# 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

<b>Creating a database using Cloudera Operational Database.....</b>	<b>4</b>
Creating a database using CDP CLI.....	6
<b>Deploying applications on Cloudera Operational Database.....</b>	<b>6</b>
<b>Importing and restoring data into Cloudera Operational Database     database.....</b>	<b>7</b>
<b>Example: run an application for the Apache HBase client.....</b>	<b>9</b>
Example: run an application for the Apache Phoenix thick client.....	10
Example to run application using the Phoenix thin client.....	10
<b>Compile an application for your Cloudera Operational Database.....</b>	<b>11</b>

# Creating a database using Cloudera Operational Database

You can create an operational database in your registered Cloudera environment using the Cloudera Operational Database.

## About this task

Required role: You must be logged into the Cloudera Operational Database as an OAdmin.

Watch the video at <https://www.youtube.com/watch?v=Zyok2K23XCs>

## Before you begin

- Understand the Cloudera environment and user management. For more information, see *User Management* and *Cloudera Environments* topics.
- Set up an environment that gives you credential and cloud storage. For more information, see *Register your first environment*.
- Ensure that you are authorized to create a database.
- To view the available instance types for your database deployments, see *View Instance Types in Cloudera Operational Database*.

## Procedure

1. In the Cloudera Operational Database 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 Cloudera environment from the list in which you want to associate the database.
  - 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. Provision your database by defining a scale for your database using a predefined Data Lake template.

The template helps you to structure your database automatically thereby saving your time and cost. Cloudera Operational Database 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.

You can create a small database using the Micro Duty template, which consists of one Gateway node and one Worker node. In a Micro database, the Gateway node carries out the processes involved in the Master or Leader nodes. You can consider using a Micro cluster for your testing and development purposes.

#### 4. Configure your database by selecting the storage type as Cloud Storage with Caching, Cloud Storage, or HDFS.

- 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.

You can also use this storage to configure time-based priority caching, where data within a specified time range is given a higher priority. In contrast, older data are likely to get evicted. For more information on this storage type, see *HBase Time-based Data Tiering using Persistent BucketCache*

You must have the `COD_DATATIERING` entitlement to use time-based priority caching.

- The storage type Cloud Storage, which resembles block storage, is equivalent to using `--storage-type CLOUD` option on CDP CLI while creating an operational database.
- The storage type HDFS is equivalent to using `--storage-type HDFS` option on CDP CLI while creating an operational database.

By default, Cloud Storage with Caching is selected.

Choose a type of your database setup.

- Choose Default for pre-selected and recommended configurations.
- Choose Custom if you need to modify any of the default values.
  - The HDFS Volume Type option appears only if you select HDFS as the storage type.
  - If you disable the Autoscaling option, the Worker Nodes and Compute Nodes options are hidden. Instead, a Node Count option appears for Worker Nodes.

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.

#### 5. Review the details before creating the database.

Click View 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-7218 --database-name test --scale-type LIGHT --storage-type CLOUD_WITH_EPHEMERAL --auto-scaling-parameters '{"minWorkersForDatabase":5, "maxWorkersForDatabase":100}' --num-edge-nodes 0 --java-version 8
```

#### 6. Click Create Database.

### Results

An information page is displayed that shows the status of the database. Your new database is ready to be used once its status becomes Available.



**Note:** Your database starts with a fixed size; however, it scales up and down as the workload applied to the database changes. For more information, see the *Auto Scaling* topic.

### Related Information

[Cloudera Operational Database Edge Node Overview](#)

[Cloudera Operational Database User Management](#)

[Cloudera Environments](#)

[Register your first environment](#)

[Cloudera Operational Database Auto Scaling](#)

[Cloudera Operational Database CLI command reference GitHub repository](#)

[Cloudera CLI BETA command reference GitHub repository](#)

## Creating a database using CDP CLI

You can create an operational database in your registered environment using CDP CLI Beta.

### Before you begin

- You must download and install the latest CDP CLI beta version.
- Understand Cloudera environment and user management. For more information, see *User Management* and *Cloudera Environments* topics.
- Set up an environment that gives you credential and cloud storage.
- Ensure you are authorized to create a database.

### About this task

- Required role: You must be logged into the Cloudera Operational Database as an ODAdmin.
- You can create a Cloudera Operational Database in an AWS, GCP, or Azure environment.

### Procedure

1. Log in to the CDP CLI beta tool.
2. Enter the following command:

```
cdp opdb create-database --environment-name cod7215 --database-name test123
```

Pass the `--use-hdfs` flag to create a database for Cloudera Operational Database using HDFS storage.

If you pass `--no-use-hdfs` flag or do not pass any flag, the database for Cloudera Operational Database is created using cloud storage. For example, Amazon S3, GCS, or ABFS.

### Results

Your new database is ready for use.

### Related Information

[Cloudera Operational Database User Management](#)

[Cloudera Environments](#)

## Deploying applications on Cloudera Operational Database

The edge node is a dedicated Cloudera Data Hub cluster that enables you to communicate with your Cloudera Operational Database instance and your applications. You can deploy a cluster that works as an edge node to access your Cloudera Operational Database instance. Deploy the edge node cluster in the same environment as the Cloudera Operational Database instance to ensure that the security groups and data ingress rules that apply to the Cloudera Operational Database instance must also apply to the edge node cluster.

### Procedure

1. From the Cloudera Management Console, click Data Hub Clusters.
2. Click Create Data Hub.
3. In the Selected Environment with running Data Lake drop-down list, select the same environment used by your Cloudera Operational Database instance.
4. Select the Cluster Definition.

5. In the Cluster Definition drop-down list, select the [\*\*\**RUNTIME VERSION*\*\*\*] COD Edge Node for [\*\*\**CLOUD PROVIDER NAME*\*\*\*].




For example, select the 7.2.10 COD Edge Node for AWS cluster template.

Data Hubs / Provision Data Hub

Provision Data Hub

Provision on-demand workload clusters with the combination of applications for various business needs such as enterprise data warehouse management and data science operations.

Selected Environment with running Data Lake

☒ Cluster Definition ☐ Custom

Services

Select the Cluster Definition option to create your cluster quickly by using one of the prescriptive cluster definitions included by default or one of your previously created custom cluster definitions.

Cluster Definition\*

Please select a Cluster Definition

- 7.2.10 - Flow Management Light Duty for AWS
- 7.2.10 - Operational Database with SQL for AWS
- 7.2.10 - Real-time Data Mart for AWS
- 7.2.10 - Streaming Analytics Heavy Duty for AWS
- 7.2.10 - Streaming Analytics Light Duty for AWS
- 7.2.10 - Streams Messaging Heavy Duty for AWS
- 7.2.10 - Streams Messaging Light Duty for AWS
- 7.2.10 COD Edge Node for AWS

Auto Scaling

☐ Currently autoscale is disabled

Advanced Options ☐

Provision Cluster Save As New Definition Show CLI Command Show Generated Cluster Template

6. In the Cluster Name field, provide a cluster name that you can identify later as an edge node of a specific Cloudera Operational Database instance.
7. Click Provision Cluster.

### What to do next

After you deploy the edge node, you can run your applications on this edge node using the [Client connectivity information](#). See how to compile applications for Cloudera Operational Database in [Compile an application for your database](#).

## Importing and restoring data into Cloudera Operational Database database

You can import your data into your Cloudera Operational Database by restoring your HBase table into Cloudera Operational Database.

### Before you begin

- Enable HBase replication on your Cloudera Operational Database cluster. For more information, see *Cloudera Operational Database data replication*.

- Have a location in cloud storage (for example, S3 or ABFS) with an exported snapshot in it, and have the name of the snapshot.

If you do not already have an exported HBase snapshot, you can export your data to cloud storage using the following command:

```
hbase org.apache.hadoop.hbase.snapshot.ExportSnapshot -snaps
hot [***SNAPSHOT NAME***] -copy-to [***CLOUD STORAGE LOCATION***] -mappers
10
```

For example, the data-from-onprem snapshot can be exported into s3a://cod-external-bucket/hbase:

```
hbase org.apache.hadoop.hbase.snapshot.ExportSnapshot -snapshot data-from-
onprem -copy-to s3a://cod-external-bucket/hbase -mappers 10
```

- Have an edge node with a configured HBase client tarball and know how to launch the hbase shell from it. For more information, see *Launching HBase shell*.

## Procedure

1. Get your CLOUD STORAGE LOCATION for your Cloudera Operational Database using the Cloudera Operational Database web interface.

s3a://my\_cod\_bucket/cod-12345/hbase

Databases / doc-test

The screenshot shows the Cloudera Operational Database web interface for a database named 'doc-test'. The interface includes a status bar at the top indicating 'Available - Updated just now'. Below this, the database name 'doc-test' is displayed along with its ID. A table below the name lists various attributes: ENVIRONMENT (cod-727-newsubnets), REGION (us-west-2), DATA LAKE (cod-727-newsubnets), SQL EDITOR (Hue), and CLOUD STORAGE LOCATION (s3a://cod-727-mowdev/cod-9zouq3ua3qqz/hbase). The 'CLOUD STORAGE LOCATION' field is highlighted with an orange box.

2. Add your bucket to the IAM policy used by IDBroker.

For more information, see one of the following documentation:

- [AWS Environments: Minimal setup for cloud storage](#)
- [Azure Environment: Minimal setup for cloud storage](#)

3. Launch the hbase shell from the edge node.

For more information, see *Launching HBase shell*.

4. From the edge node, run the ExportSnapshot command. Use the external bucket as the source location and the Cloudera Operational Database cloud storage as the target.

For example:

```
$ cd $HBASE_HOME
$ ./bin/hbase org.apache.hadoop.hbase.snapshot.ExportSnapshot -snapshot
"data-from-onprem" --copy-from s3a://cod-external-bucket/hbase --copy-to
s3a://my_cod_bucket/cod-12345/hbase
```

5. Use the list\_snapshots command and verify that your snapshot is listed.

```
$ cd $HBASE_HOME
$ ./bin/hbase shell
$ hbase> list_snapshots
[ 'snapshot_name' ]
```



6. Use the `restore_snapshot` or the `clone_snapshot` command to reconstitute the table.



**Warning:** The `restore_snapshot` command overwrites an existing table.

- `restore_snapshot`: Overwrites current table state with that of the snapshot. It means that any data modification applied after the snapshot was taken would be lost. If the table does not exist in the given cluster, the command automatically creates it.
- `clone_snapshot`: Accepts a new table name for the table in which it restores the table schema and data.

7. Validate that all rows are present in the table using the `hbase rowcounter` or `count` command in the `hbase` shell.

### Related Information

[Launching hbase shell](#)

[Cloudera Operational Database data replication](#)

[Using the CldrCopyTable utility to copy data](#)

## Example: run an application for the Apache HBase client

Checking the example of how to run a Maven application for the HBase client gives you better understanding about how to run your own application for the HBase client.

### Before you begin

The required `Kerberos` option is set to `true` for the Apache HBase client. The application must be run for this client from a computer which:

- Has internal network access to the public cloud in which the database is deployed
- Can resolve the internal hostnames of the database
- Can obtain a Kerberos ticket from the database's Kerberos Key Distribution Center (KDC)

One way to run the Apache HBase client applications is to launch an edge node in your cloud provider which meets the above requirements.

Ensure that you download the directory containing the Apache HBase client configuration files from `clientConfigurationUrl` provided in `describe-client-connectivity` response. This is a protected endpoint, and you have to use your Cloudera workload credentials to access this endpoint. You can also get the Apache HBase client configuration as a client tarball from the Cloudera Operational Database web user interface `Connect` tab.

### Procedure

1. Use `clientConfigurationURL` from the `describe-client-connectivity` response to obtain the necessary configuration details to communicate with Apache HBase:

```
$ cdp opdb describe-client-connectivity --database-name [***YOUR DATABASE NAME***] --environment-name [***YOUR ENVIRONMENT NAME***] | jq '.connectors[] | select(.name == "hbase") | .configuration.clientConfigurationUrl'
$ curl -k -o clientConfig.zip -u '***USERNAME***':[***CLOUDERA WORKLOAD PASSWORD***] https://client_configuration_url/.../services/hbase/clientConfig
```

You can build the application in your local machine and copy the JAR files to the remote node using the following commands:

```
$ scp -r target ec2-user@my-ec2-bastion-host.us-west-2.compute.amazonaws.com:
```

```
$ scp clientConfig.zip ec2-user@my-ec2-bastion-host.us-west-2.compute.amazonaws.com:
$ ssh ec2-user@my-ec2-bastion-host.us-west-2.compute.amazonaws.com "sudo
yum install -y java-1.8.0-openjdk"
$ ssh ec2-user@my-ec2-bastion-host.us-west-2.compute.amazonaws.com "unzip
clientConfig.zip"
```

2. Ensure that you have a Kerberos ticket:

```
$ kinit [***USERNAME***] Password: [***PASSWORD***]
$ java -cp target/nosql-libs/*:target/nosql-exemplar-0.0.1-SNAPSHOT.jar
:hbase-conf com.cloudera.odx.nosql.Client
```

3. Run your application.

## Example: run an application for the Apache Phoenix thick client

Checking the example of how to run a maven application for the Phoenix thick client gives you better understanding about how to run your own application for the Phoenix thick client.

### Before you begin

The requiredKerberos option is set to true for the Phoenix thick client. This means that the application must be run for this client from a computer which:

- Has internal network access to the public cloud in which the database is deployed
- Can resolve the internal hostnames of the database
- Can obtain a Kerberos ticket from the database's Kerberos Key Distribution Center (KDC)

One way to run an Apache Phoenix thick client application is to launch an edge node in your cloud provider which meets the above requirements.

### Procedure

1. Use the JDBC URL from the describe-client-connectivity command to run the example.

One method is to build on your local machine and copy the JAR files to the remote node:

```
$ scp -r target ec2-user@my-ec2-bastion-host.us-west-2.compute.amazonaws.com:
$ scp clientConfig.zip ec2-user@my-ec2-bastion-host.us-west-2.compute.amazonaws.com:
$ ssh ec2-user@my-ec2-bastion-host.us-west-2.compute.amazonaws.com "sudo
yum install -y java-1.8.0-openjdk"
```

2. Ensure that you have a Kerberos ticket and run your application for the Phoenix thick client

```
kinit USERNAME
$ java -cp target/sql-libs/*:target/sql-exemplar-0.0.1-SNAPSHOT.jar:hbase-
conf com.cloudera.odx.sql.Client "[***PHOENIX THICK JDBC URL***]"
```

3. Run your application.

### Related Information

[COD edge node overview](#)

## Example to run application using the Phoenix thin client

Checking the example of how to run a maven application for the Phoenix thin client gives you better understanding about how to run your own application for the Phoenix thin client.

The required Kerberos options set to false for the Phoenix thin client which means that it can be used from virtually any node. In this example, the client runs from the local machine.

For the Apache Phoenix thin client the describe-client-connectivity call returns a base JDBC URL. You must append the following attributes to the URL which are specific to your identity:

- `avatica_user`: your Cloudera username (required)
- `avatica_password`: your Cloudera workload password (required)
- `truststore`: a truststore for your Cloudera Knox gateway (optional)
- `truststore_password`: the password for the truststore file (optional)

You can use Maven to ease launching this application, but a standalone Java program is similarly launched:

```
$ mvn exec:exec -Dexec.executable=java -Dexec.args='-cp target/sql-libs/*:target/sql-exemplar-0.0.1-SNAPSHOT.jar com.cloudera.odx.sql.ThinClient "jdbc:phoenix:thin:url=[***PHOENIX THIN JDBC URL;serialization=PROTOBUF;authentication=BASIC;avatica_user=[***USERNAME***];avatica_password=[***PASSWORD***];truststore=[***CDP-TRUSTSTORE.JKS***];truststore_password=[***TRUSTSTORE PASSWORD***]"'
```

Or, you can launch the application without the help of Maven:

```
$ java -cp target/sql-libs/*:target/sql-exemplar-0.0.1-SNAPSHOT.jar com.cloudera.odx.sql.ThinClient "jdbc:phoenix:thin:url=[***PHOENIX THIN JDBC URL;serialization=PROTOBUF;authentication=BASIC;avatica_user=[***USERNAME***];avatica_password=[***PASSWORD***];truststore=[***CDP-TRUSTSTORE.JKS***];truststore_password=[***TRUSTSTORE PASSWORD***]"
```

## Compile an application for your Cloudera Operational Database

Once you have created your application and a database using Cloudera Operational Database, you have to compile your application for your database.

### Before you begin

- Set your Cloudera workload password. For more information, see the *Setting the workload password* documentation in the related information section.
- Grant the ODUser role to the machine user using the Cloudera Management Console **Environments** *ENVIRONMENT\_NAME* **Actions** **Manage Access** page for their Cloudera environment. By setting the ODUser role you can grant a number of rights in the system that allows you to access the Cloudera Operational Database using the machine user's workload password.
- Add synchronized users from User Management Service in the Cloudera Control Plane into the environment in which your Cloudera Operational Database database is running.

## Procedure

### 1. Get the Maven repository location to fetch JAR files.

There are two ways to get the necessary information:

- In command line: Using the `cdp opdb describe-client-connectivity --database-name <YOUR_DATABASE> --environment-name <YOUR_ENVIRONMENT>` command.
- In the Cloudera Operational Database user interface: Clicking the Connect bar and selecting the applicable client.

You have to use the version and the mavenURL attributes in your Maven project and configuration.

The following is an example about how to fetch the required HBase information to build your application:

```
$ cdp opdb describe-client-connectivity --database-name <YOUR_DATABASE> --environment-name <YOUR_ENVIRONMENT> | jq '.connectors[] | select(.name == "hbase")'
{
  "name": "hbase",
  "version": "2.2.3.7.2.0.0-219",
  "kind": "LIBRARY",
  "dependencies": {
    "mavenUrl": "HTTPS://REPOSITORY.CLOUDERA.COM/ARTIFACTORY/CLOUDERA-REPOS"
  },
  "configuration": {
    "clientConfigurationUrl": "HTTP://CLIENT_CONFIGURATION_URL/.../SERVICES/HBASE/CLIENTCONFIG"
  },
  "requiresKerberos": true
}
```

The following example fetch the required Phoenix information to build your application:

```
cdp opdb describe-client-connectivity --database-name <YOUR_DATABASE> --environment-name <YOUR_ENVIRONMENT> | jq ".connectors[] | select(.name == \"phoenix-thick-jdbc\") | .dependencies.mavenUrl"
cdp opdb describe-client-connectivity --database-name <YOUR_DATABASE> --environment-name <YOUR_ENVIRONMENT> | jq ".connectors[] | select(.name == \"phoenix-thin-jdbc\") | .version"

Phoenix-thick
"HTTPS://REPOSITORY.CLOUDERA.COM/ARTIFACTORY/CLOUDERA-REPOS"
"5.0.0.7.2.0.0-128"
Phoenix-thin
"HTTPS://REPOSITORY.CLOUDERA.COM/ARTIFACTORY/CLOUDERA-REPOS"
"5.0.0.7.2.0.0-128"
```

### 2. Modify your application's settings.

Ensure your application's settings.xml file uses the correct URL and version for your Cloudera Operational Database database.

An example when using NoSQL client:

```
<project>
  <dependencies>
    <!-- NoSQL client for Cloudera Operational Database -->
    <dependency>
      <groupId>org.apache.hbase</groupId>
      <artifactId>hbase-shaded-client</artifactId>
      <version>2.2.3.7.2.0.0-219</version>
    </dependency>
  </dependencies>
  ...
</project>
```

```

<repositories>
  <!-- Define our Cloudera Operational Database repository; this would be
  given to us by Cloudera Operational Database itself -->
  <repository>
    <id>nosql-odx</id>
    <url>HTTPS://MAVEN_URL/CDP-PROXY/HBASE/JARS</url>
    <name>Cloudera NoSQL Cloudera Operational Database Repository</
name>
    <snapshots>
      <enabled>>false</enabled>
    </snapshots>
  </repository>
</repositories>
</project>

```

An example when using SQL client:

```

<project>
  <dependencies>
    <!-- SQL client for ODX -->
    <dependency>
      <groupId>org.apache.phoenix</groupId>
      <artifactId>phoenix-client</artifactId>
      <version>5.0.0.7.2.0.0-128</version>
    </dependency>
    <dependency>
      <groupId>org.apache.phoenix</groupId>
      <artifactId>phoenix-queryserver-client</artifactId>
      <version>5.0.0.7.2.0.0-128</version>
    </dependency>
  </dependencies>
  ...
  <repositories>
    <!-- Define our Cloudera Operational Database repository -->
    <repository>
      <id>sql-cod</id>
      <url>HTTPS://MAVEN_URL/CDP-PROXY-API/AVATICA/MAVEN</url>
      <name>Cloudera SQL Cloudera Operational Database Repository</name>
      <snapshots>
        <enabled>>false</enabled>
      </snapshots>
    </repository>
  </repositories>
</project>

```

3. Build your application.
4. Run your application for the applicable client.

### Related Information

[Setting the workload password](#)