

CDP Data Visualization Top Tasks

Date published: 2020-10-30

Date modified: 2024-12-18

CLOUDERA

Legal Notice

© Cloudera Inc. 2024. 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

Deploying Cloudera Data Visualization.....	4
Deploying Cloudera Data Visualization in Cloudera Data Warehouse.....	4
Deploying a Cloudera Data Visualization application in Cloudera AI.....	8
Creating a Cloudera AI project with Cloudera Data Visualization Runtime.....	8
Adding Cloudera Data Visualization Runtime to a Cloudera AI project.....	9
Creating a Cloudera Data Visualization application in Cloudera AI.....	9
Deploying a Cloudera Data Visualization application in Cloudera Data Science Workbench.....	10
Creating a Cloudera Data Science Workbench project with Cloudera Data Visualization Runtime.....	10
Creating a Cloudera Data Visualization application in Cloudera Data Science Workbench.....	11
Starting Cloudera Data Visualization in Cloudera Data Science Workbench.....	12
 Connecting to a data source in Cloudera Data Visualization.....	 13
Creating a Cloudera AI data connection to Impala.....	13
Creating a Cloudera AI data connection to a Hive data warehouse.....	20
Creating a Cloudera Data Warehouse data connection in Cloudera Data Visualization.....	26
Creating a Cloudera Data Science Workbench data connection to a data warehouse.....	29
 Creating a visual.....	 31

Deploying Cloudera Data Visualization

The first step to start using Cloudera Data Visualization is to deploy it on your chosen Cloudera platform. Cloudera Data Visualization provides consistent visualization capabilities across various Cloudera deployments and services. Once your Cloudera Data Visualization instance is up and running, you can start creating visuals, build customized dashboards, and develop analytical applications to enable users to explore data across the whole data lifecycle.

You can log into Cloudera Data Visualization using your Cloudera credentials and start analyzing data that you have access to. In Cloudera Public Cloud, Cloudera Data Visualization is included as a core capability in Cloudera Data Warehouse and Cloudera AI both for Azure and AWS. Cloudera Data Visualization is also available in Cloudera Private Cloud for Cloudera Data Warehouse and Cloudera AI, and it can be used in Cloudera Data Science Workbench, certified with CDH, HDP and Cloudera Private Cloud Base.

Deploying Cloudera Data Visualization in Cloudera Data Warehouse

Learn how to create and launch a Cloudera Data Visualization instance within the Cloudera Data Warehouse data service. You can use Cloudera Data Visualization with Cloudera Data Warehouse in both Cloudera Public Cloud and Cloudera Private Cloud to explore and visualize data stored in database catalogs providing data-driven insights throughout the entire data lifecycle.

About this task

If you want to create visuals based data stored in Cloudera Data Warehouse, you have to create a Cloudera Data Visualization instance and connect it to Hive or Impala Virtual Warehouse(s).

Cloudera Data Visualization is not tied to a particular Virtual Warehouse (VW). You can access data for your visualizations from multiple Data Catalogs using multiple Hive or Impala Virtual Warehouses in various environments. With multiple Cloudera Data Visualization instances attached to an environment, you can connect to different data sets, create dashboards for different groups and share your visualizations with different users.



Note: When you delete a Virtual Warehouse, your visual artifacts remain intact as long as the Cloudera Data Visualization instance is not deleted.

Before you begin

- You are logged into the Cloudera web interface and you have opened the Cloudera Data Warehouse service.
- You have DWAdmin role in Cloudera Data Warehouse.
- You have activated your environment.
- You have a Hive/Impala warehouse in running state.
- If you are using Cloudera Data Visualization with Cloudera Data Warehouse in Cloudera Private Cloud:

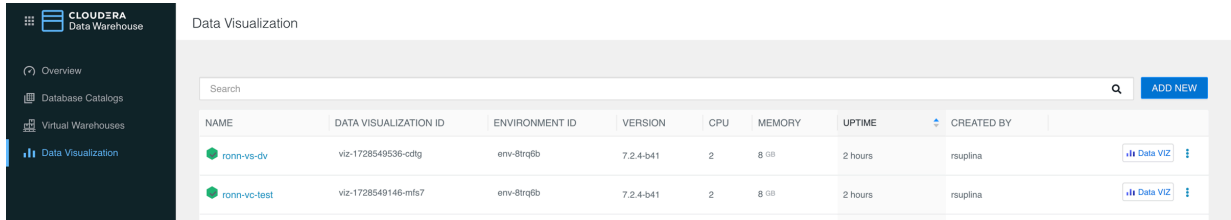
You have an admin group created in Management Console User Management . To log in using LDAP, select the Sync Groups on Login option from Management Console Administration Authentication , so that the associated groups can be imported when you log in to Cloudera Data Visualization.

Procedure

1. In Cloudera Data Warehouse, click Data Visualization in the left navigation panel.
A list of existing Cloudera Data Visualization instances appears, if there are any.

2. Click **ADD NEW** to create a new instance.

If you are creating the first Cloudera Data Visualization instance in Cloudera Data Warehouse, click **CREATE**.



The screenshot shows the Cloudera Data Warehouse interface. On the left is a dark sidebar with navigation options: Overview, Database Catalogs, Virtual Warehouses, and Data Visualization (highlighted). The main content area is titled "Data Visualization" and features a search bar and an "ADD NEW" button. Below these is a table listing existing instances.

NAME	DATA VISUALIZATION ID	ENVIRONMENT ID	VERSION	CPU	MEMORY	UPTIME	CREATED BY	
ronn-vs-dv	viz-1728549536-ctlg	env-8trq6b	7.2.4-541	2	8 GB	2 hours	rsuplina	Data VIZ
ronn-vc-test	viz-1728549146-mfs7	env-8trq6b	7.2.4-541	2	8 GB	2 hours	rsuplina	Data VIZ

3. Provide the following information in the **New Data Visualization** modal window:

Mandatory fields are marked with an asterisk.

- a) Name* – Specify a name for the instance.
- b) Environments* – Select the environment that you want Cloudera Data Visualization to be connected to.
- c) User Groups – Add user groups to allow user access to Cloudera Data Visualization for selected groups. If no group is added, all Cloudera users will have non-admin access.
- d) Admin Groups* – Add admin groups to allow configuration access to Cloudera Data Visualization for selected groups.

For more information on Cloudera Data Visualization permission granularity, see the *Security model*.

- e) Tagging – Enter keys and values to apply tags to your resources for organizing them into a taxonomy.



Note: This field is not available in Cloudera Data Warehouse on Cloudera Private Cloud.

- f) Resource Template – Select the Cloudera Data Warehouse resource template from the drop-down menu.
 - Default resources
 - Medium resources
 - Large resources

New Data Visualization
X

Name *

Environments *

se-sandboxx-aws
▼

User Groups ⓘ

Select Groups

You can select groups present in the CDP user management system. The group(s) must also exist in the external LDAP Identity Provider. Nested groups are not supported. Only users that are direct members of the group are allowed access.

Admin Groups * ⓘ

Select Admin Groups

You can select groups present in the CDP user management system. The group(s) must also exist in the external LDAP Identity Provider. Nested groups are not supported. Only users that are direct members of the group are allowed access.

Tagging ⓘ

Enter key

Enter value

+

Only alphanumeric and _-@: are allowed

Resource Template

Default resources
▼

Create

4. Click CREATE.

Instance creation starts. Wait until the Cloudera Data Visualization instance is in running state.

5. You can find the list of Cloudera Data Visualization instances and environments appears under the Data Visualization menu that you can open from the left navigation panel.

CLUSTER

Data Warehouse

Overview

Database Catalogs

Virtual Warehouses


Data Visualization

Data Visualization

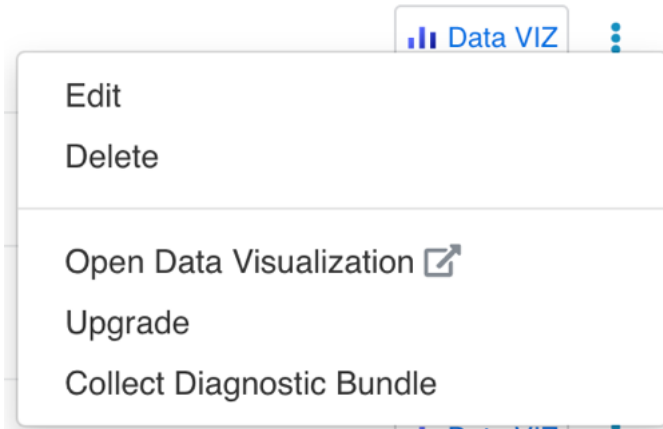
🔍
ADD NEW

NAME	DATA VISUALIZATION ID	ENVIRONMENT ID	VERSION	CPU	MEMORY	UPTIME	CREATED BY	
● Doc-demo-viz	viz-172855112-6wvm	env-wt479v	7.2.4-b41	2	8 GB	a minute	ifeher	Data VIZ ⋮
● ronni-vs-dv	viz-1728549536-cdtg	env-8trq6b	7.2.4-b41	2	8 GB	2 hours	rsuplina	Data VIZ ⋮
● ronni-vc-test	viz-1728549146-mfs7	env-8trq6b	7.2.4-b41	2	8 GB	2 hours	rsuplina	Data VIZ ⋮

6.

Select one from the list of running Cloudera Data Visualization instances and click  to start Cloudera Data Visualization.

Alternatively, you can click the launch/options menu (three dots) at the end of the row and click Open Data Visualization in the pop-up menu.



SSO authentication is enabled by default both in Cloudera Public Cloud and Cloudera Private Cloud. If you are logged in to the Control Plane, you will be logged into Cloudera Data Visualization automatically. If not, log into the Control Plane using your LDAP credentials. Cloudera Data Visualization opens in a new browser tab and you land on the Cloudera Data Visualization homepage, where you can explore sample dashboards and access the in-tool *Get Started guide* for help.

New users logging in to Cloudera Data Visualization are automatically added to the `viz_guest_group` group. You can assign the System Admin role to this group to grant new users the permission to create new connections. For more information, see *Assigning roles to users*.

Deploying a Cloudera Data Visualization application in Cloudera AI

Learn how to deploy Cloudera Data Visualization in Cloudera AI.

Creating a Cloudera AI project with Cloudera Data Visualization Runtime

Learn how to create a Cloudera AI project with Cloudera Data Visualization Runtime as the default runtime.

About this task

If you know that your project is going to be running Cloudera Data Visualization, you can add the Cloudera Data Visualization Runtime when setting up the project.

Procedure

1. Click Projects on the left sidebar of Cloudera AI data service.
2. Click New Project.
3. Enter a project name.
4. You can add a description for your project.
5. You can also select the visibility of the project.
6. Under Initial Setup, you can either create a blank project, or select a source for your project files.
7. Set a Cloudera Data Visualization Runtime for your project in the Runtime setup section.

You can use the Basic view to select the Cloudera Data Visualization Runtime and add all of its Standard Edition ML Runtimes for that kernel. In the Advanced view, you can add ML Runtimes based on more detailed Editor, Kernel, Edition, and Version criteria.

8. Click Create Project.

What to do next

After the project is created, you can start creating your application. If you added Cloudera Data Visualization Runtime as the only runtime during project set-up, it will be the only available runtime when creating your application. You can select which edition you want to use for the application.

Adding Cloudera Data Visualization Runtime to a Cloudera AI project

Learn how to add a Cloudera Data Visualization Runtime to an existing Cloudera AI project.

About this task

You need to add a Cloudera Data Visualization Runtime to your project if the workspace of your project is not set to use this Cloudera Data Visualization Runtime by default.

Procedure

1. Click Projects on the left sidebar of the Cloudera AI data service.
2. Click the project where you want to add Cloudera Data Visualization Runtime.
3. Open Project Settings from the left navigation bar and click the Runtime/Engine tab.
4. Click Add Runtime.

The Add new runtime to project modal window opens.

5. Select Cloudera Data Visualization from the Kernel drop-down list.
6. Select the edition of the Cloudera Data Visualization you want to add to your project.

The selected edition automatically defines the version of the runtime variant.



Note: New Cloudera Data Visualization Runtime versions are automatically added to the list of available editions, if internet connection is available.

7. Click Submit.

Results

You should now be able to select the Cloudera Data Visualization Runtime when creating a Cloudera Data Visualization application.

What to do next

Create a Cloudera Data Visualization application.

Creating a Cloudera Data Visualization application in Cloudera AI

Learn how to create a Cloudera Data Visualization application in Cloudera AI to help you visualize and interact with your data insights. This integration allows for seamless visualization of ML model outputs, data exploration, and reporting within the same platform.

About this task

The following steps will guide you on how to create a Cloudera Data Visualization application in a Cloudera AI project.



Important:

Each Cloudera AI project can host only one standalone Cloudera Data Visualization application. Since applications share the same metadata and logs, any changes made by one application can overwrite those made by another.

Before you begin

Ensure that a Cloudera Data Visualization Runtime is available in the Cloudera AI project where you want to create the Cloudera Data Visualization application.

For more information about ML Runtimes, see [Managing ML Runtimes](#) and [Using Runtime Catalog](#).

Procedure

1. Navigate to the Overview page of your Cloudera AI project.
2. On the left sidebar, click Applications.
3. Click New Application.
4. Provide the following details for your new application:
 - Name – Enter a name for the application.
 - Subdomain – Enter a subdomain that will be used to construct the URL for the web application. Ensure it only contains URL-friendly characters.
 - Description – Add a description for the application.
 - Script – Select the path to the startup script.



Note: Use the script located at: `/opt/vizapps/tools/arcviz/startup_app.py`

- Runtime – Configure the Runtime.
 - Editor – Workbench
 - Kernel – Select Cloudera Data Visualization for the kernel supported by the Runtime variant of the Cloudera AI project.
 - Edition – Select the edition of the Runtime variant you want to use for your application.



Note: The selected edition determines the version of the Runtime variant.

5. Click Create Application.

Results

After a few minutes, the application status will change from Starting to Running on the Applications page. Your Cloudera Data Visualization application is now fully operational.

You can restart, stop, or delete a Cloudera AI application from the supplemental menu. If you want to make changes to the application, navigate to [Application Details Settings](#).

What to do next

Start Cloudera Data Visualization.

Deploying a Cloudera Data Visualization application in Cloudera Data Science Workbench

Learn how to deploy Cloudera Data Visualization in Cloudera Data Science Workbench.

Creating a Cloudera Data Science Workbench project with Cloudera Data Visualization Runtime

Learn how to create a Cloudera Data Science Workbench project with a Cloudera Data Visualization Runtime as the default runtime.

Procedure

1. Click Projects on the left sidebar of Cloudera Data Science Workbench.
2. Click New Project on the Projects page.
3. Enter a project name.
4. Select the visibility of the project.
5. Under Initial Setup, you can either create a blank project, or select a source for your project files.
6. Click Create Project.

The created project uses Cloudera ML Runtimes that are managed in the Runtime Catalog.

What to do next

After the project is created, you can start creating an application. You can select which runtime edition you want to use for the application.

Creating a Cloudera Data Visualization application in Cloudera Data Science Workbench

Learn how to create a Cloudera Data Visualization application in Cloudera Data Science Workbench to help you visualize and interact with your data insights. This integration allows for seamless data visualization, exploration, and reporting within the same platform.

About this task

The following steps will guide you on how to create a Cloudera Data Visualization application within Cloudera Data Science Workbench.



Important:

Each Cloudera Data Visualization project can host only one standalone Cloudera Data Visualization application. Since applications share the same metadata and logs, any changes made by one application can overwrite those made by another.

Procedure

1. Navigate to the Overview page of your Cloudera Data Science Workbench project.
2. On the left sidebar, click Applications.
3. Click New Application.

4. Provide the following details for your new application:

- Name – Enter a name for the application.
- Subdomain – Enter a subdomain that will be used to construct the URL for the web application. Ensure it only contains URL-friendly characters.
- Description – Add a description for the application.
- Enable Unauthenticated Access – Mark the checkbox if you want to allow unauthenticated access to your application. You can also update this setting later on the Settings page of the application.



Note: To create public applications on an ML workspace, an Admin user needs to turn on the feature flag in **Admin Security** by selecting **Allow applications to be configured with unauthenticated access**. For more information on working with public applications, see *Application Limitations* in the Cloudera Data Science Workbench documentation.

- Script – Select the path to the startup script.



Note: Use the script located at: `/opt/vizapps/tools/arcviz/startup_app.py`

- Runtime
 - Editor – Workbench
 - Kernel – Select Cloudera Data Visualization for the kernel supported by the Runtime variant of the Cloudera Data Science Workbench project.
 - Edition – Select the edition of the Runtime variant you want to use for your application.



Note: The selected edition determines the version of the Runtime variant.

All available runtimes are listed in the Cloudera Data Science Workbench Runtime Catalog. For more information about ML Runtimes, see [Managing ML Runtimes](#) and [Using Runtime Catalog](#).



Note: New ML Runtime releases are automatically added to the deployment, if internet connection is available.

5. Click Create Application.

Results

After a few minutes, the application status will change from Starting to Running on the Applications page. Your CDV application is now fully operational.

You can restart, stop, or delete a Cloudera Data Science Workbench application from the supplemental menu of the application. If you want to make changes to the application, navigate to **Application Details Settings**.

What to do next

Start Cloudera Data Visualization in Cloudera Data Science Workbench.

Starting Cloudera Data Visualization in Cloudera Data Science Workbench

Learn how to start the Cloudera Data Visualization application you have created in Cloudera Data Science Workbench.

Procedure

1. On the Applications page, click the name of your Cloudera Data Visualization application to access the login interface.

2. Log in to Cloudera Data Visualization by entering your username and password. Use your workload credentials.

If you want to log in as an administrator, you can use the following default credentials:

- username: vizapps_admin
- password: vizapps_admin

When using the default credentials to log in to Cloudera Data Visualization, you are prompted to change your password at the first login.



Important: If you use the default credentials, security issues may arise. Cloudera recommends that you change the default username and password.

SSO authentication is disabled by default. See [Authentication](#) for information on how to permit user authentication with the Cloudera Data Science Workbench login credentials and log users in automatically.

After logging in, you land on the homepage view of Cloudera Data Visualization. Here you can explore some sample dashboards or access the in-tool Get Started guide for help.

Connecting to a data source in Cloudera Data Visualization

Cloudera Data Visualization allows you to create connections to many types of external data sources, enhancing your data analysis and visualization capabilities.

Cloudera Data Visualization currently supports the following connection types:

- Hive
- Impala
- MariaDB
- MySQL
- PostgreSQL
- Solr [Technical Preview]
- Spark SQL
- SQLite (not supported in CDW)
- Snowflake [Technical Preview]

This range of supported connection types provides flexibility and versatility in connecting to various data sources for comprehensive data analysis.

When using Cloudera Data Visualization with Cloudera Data Warehouse, the data connection is automatically set up, but you can connect to various other data sources as well to suit your specific requirements.

In Cloudera AI, you can connect to an Impala or a Hive data warehouse, or tie in data from predictive Cloudera AI models.

Creating a Cloudera AI data connection to Impala

Learn how to connect natively to data stored in Impala when using Cloudera Data Visualization in Cloudera AI.

About this task

Before you start using data modeling and visualization functions, you must connect to your data. The following steps show you how to create a new Cloudera AI data connection in Cloudera Data Visualization to an Impala data warehouse.

**Note:**

You must have the Manage data connections privilege or be an admin to be able to create and manage connections in Cloudera Data Visualization.

Setting user privileges requires administrator-level access. You can log in as an administrator using the default admin account with the following credentials:

- Username: vizapps_admin
- Password: vizapps_admin

When you create a connection, you automatically get privileges to create and manage the associated datasets. You can also build dashboards and visuals within these datasets.

- For more information on the Manage data connections privilege, see *RBAC permissions*.
- For instructions on how to define privileges for a specific role, see *Setting role privileges*.
- For instructions on how to assign the administrator role to a user, see *Promoting a user to administrator*.

Before you begin

If you are using a Cloudera Private Cloud Base cluster running Impala with Kerberos for authentication, make sure that Kerberos credentials are configured in Cloudera AI before creating a Cloudera AI data connection to the Impala data warehouse. This ensures seamless integration and authentication between Cloudera Data Visualization and the Impala cluster. If you add Kerberos credentials after launching the Cloudera Data Visualization app, you need to restart the app for the changes to take effect.

For more information on using Kerberos for authentication in Cloudera AI, see *Hadoop Authentication for AI Workspaces*.

Procedure

1. On the main navigation bar, click DATA.

The DATA interface opens, displaying the Datasets tab.

Title/Table	ID	Tags	Created	Last Updated	Modified By	# Dashboards
test-import-dashboard-dataset-association default.bbviz2077	546		Jun 20, 2024	a month ago	vizapps_admin	11
customers default.customers	527		May 31, 2024	2 months ago	vizapps_admin	0
Clone of Hr Capitalized default.hr_partitioned	514		May 22, 2024	2 months ago	vizapps_admin	0
tb_from_sql text.sql, default.bbviz2077, default.bbviz2077, ...	470		Apr 09, 2024	4 months ago	vizapps_admin	0
web_logs_test_joins default.web_logs, default.web_logs	469		Apr 09, 2024	4 months ago	vizapps_admin	0
Hr Capitalized default.hr_partitioned	441		Mar 12, 2024	4 months ago	ckoncz	1

2. On the side menu bar, click NEW CONNECTION.



Note: Only users with Manage data connections privilege or administrators can access the NEW CONNECTION button.

The screenshot shows the Cloudera Data Visualization interface. The top navigation bar includes 'HOME', 'SQL', 'VISUALS', and 'DATA'. A search bar is present on the right. The left sidebar contains a 'NEW CONNECTION' button (highlighted with an orange box) and a list of connections, including '2050_CloneSamples'. The main content area displays a table of datasets.

Title/Table	ID	Tags
Cereals main.cereals Data Connection: samples	11	Extract Source
cereal main.cereals, main.chicago_govt_pay Data Connection: samples	191	
cereal main.cereals, main.chicago_govt_pay Data Connection: 2050_CloneSamples	341	
csimport Solr.csaba_import_20240119 Data Connection: SolrTestConnection	417	

The Create New Data Connection modal window appears.

3. Choose Impala from the Connection type drop-down list and assign a name to your connection.

Create New Data Connection

Connection type: Impala

Connection name: doc-test

Basic | Advanced | Parameters | Data

Hostname or IP address: Enter IP address of the server where your data resides (example: prod_db.yourcompany.com or 10.0.1.20)

Port #: 443

Credentials

Username: [Empty]

Password: [Masked]

TEST | CONNECT

In this example, the Impala connection is made through Knox. Knox always uses TLS encryption and port 443 is the default HTTPS port.

4. Enter the hostname or IP address of the running coordinator.
You can retrieve this information from the JDBC URL of the Impala DW.
5. Add 443 in the Port # field.
6. Enter your workload username and password as credentials.

- Click the Advanced tab to configure additional details.

Create New Data Connection
✕

Connection type

Connection name

Impala

doc-test

Basic
Advanced
Parameters
Data

Connection mode

HTTP Path

Socket type

Authentication mode

Socket Timeout

Impersonation

Trusted Impersonation

Application Name

Binary
 HTTP

SQL path (default cliservice)

Normal
 SSL
 SSL with certificate

NoSasl
 Plain
 LDAP
 Kerberos

60

Enabled

Enabled

viz

TEST

CONNECT

- For HTTP connection mode, locate the Impala Endpoint for the Data Hub.

Name	URL
CM-API	https://jingalls-test-dm-gateway.euph-aw.a465-9q4k.cloudera.site/jingalls-test-dm/cdp-proxy-api/cm-api/
Impala	https://jingalls-test-dm-gateway.euph-aw.a465-9q4k.cloudera.site/jingalls-test-dm/cdp-proxy-api/impala/
Impala	jdbc:impala://jingalls-test-dm-gateway.euph-aw.a465-9q4k.cloudera.site:443/?ssl=1;transportMode=http;httpPath=jingalls-test-dm/cdp-proxy-api/impala;AuthMech=3;

- Copy and paste it into the HTTP Path field.

- c) Set any additional details as required.

- 8. Check the Parameters and Data tabs for more configuration options.

Create New Data Connection ✕

Connection type

Connection name

Basic Advanced **Parameters** Data

Parameter Name	Parameter Value	
Add new row		

Create New Data Connection ✕

Connection type

Connection name

Basic **Advanced** Parameters **Data**

Concurrency

Concurrency Per User

Query Timeout (Minutes)

Query Still Loading Warning (Seconds)

Row upload limit

Result Cache Enabled

Cache Retention Time (seconds)

9. Once you finish configuring the settings, click TEST to check the connection.

10. Click CONNECT to establish the connection.

Results

You have successfully set up a connection to a running Impala DW.

Related Information

[RBAC permissions](#)

[Setting role privileges](#)

[Promoting a user to administrator](#)

Creating a Cloudera AI data connection to a Hive data warehouse

Learn how to connect natively to data stored in Hive when using Cloudera Data Visualization in Cloudera AI.

About this task

Before you start using data modeling and visualization functions, you must connect to your data. The following steps show you how to create a new Cloudera AI data connection in Cloudera Data Visualization to a Hive data warehouse.



Note:

You must have the Manage data connections privilege or be an admin to be able to create and manage connections in Cloudera Data Visualization.

Setting user privileges requires administrator-level access. You can log in as an administrator, using the default admin account with the following credentials:

- Username: vizapps_admin
- Password: vizapps_admin

When you create a connection, you automatically get privileges to create and manage the associated datasets. You can also build dashboards and visuals within these datasets.

- For more information on the Manage data connections privilege, see *RBAC permissions*.
- For instructions on how to define privileges for a specific role, see *Setting role privileges*.
- For instructions on how to assign the administrator role to a user, see *Promoting a user to administrator*.

Before you begin

If you are using a Cloudera Private Cloud Base cluster running Hive with Kerberos for authentication, make sure that Kerberos credentials are configured in Cloudera AI before creating a Cloudera AI data connection to the Hive data warehouse. This ensures seamless integration and authentication between Cloudera Data Visualization and the Hive cluster. If you add Kerberos credentials after launching the Cloudera Data Visualization app, you need to restart the app for the changes to take effect.

Procedure

1. On the main navigation bar, click DATA.

The DATA opens, displaying the Datasets tab.

Title/Table	ID	Tags	Created	Last Updated	Modified By	# Dashboards
test-import-dashboard-dataset-association default.bbviz2077	546		Jun 20, 2024	a month ago	vizapps_admin	11
customers default.customers	527		May 31, 2024	2 months ago	vizapps_admin	0
Clone of Hr Capitalized default.hr_partitioned	514		May 22, 2024	2 months ago	vizapps_admin	0
tb_from_sql text.sql, default.bbviz2077, default.bbviz2077, ...	470		Apr 09, 2024	4 months ago	vizapps_admin	0
web_logs_test_joins default.web_logs, default.web_logs	469		Apr 09, 2024	4 months ago	vizapps_admin	0
Hr Capitalized default.hr_partitioned	441		Mar 12, 2024	4 months ago	ckoncz	1

- On the side menu bar, click NEW CONNECTION.



Note: The NEW CONNECTION button is only accessible to users assigned to roles with the Manage data connections privilege and to administrators.

The screenshot shows the Cloudera Data Visualization interface. On the left side menu, the 'NEW CONNECTION' button is highlighted with an orange box. The main content area displays a table of datasets with the following columns: Title/Table, ID, and Created.

Title/Table	ID	Created
Test Dataset main.census_pop	13	Dec 08, 2021
Food Stores Inspection in NYC main.retail_food_store_inspections_current_critical_vio...	12	Nov 22, 2021
Cereals main.cereals	11	Nov 22, 2021
Earthquake Data January 2019 main.earthquake_data2019	10	Nov 22, 2021
World Life Expectancy main.world_life_expectancy	9	Nov 22, 2021

The Create New Data Connection modal window appears.

- Choose Hive from the Connection type drop-down list and assign a name to your connection.

Create New Data Connection ✕

Connection type Hive

Connection name doc-test

Basic **Advanced** **Parameters** **Data**

Hostname or IP address Enter IP address of the server where your data resides
(example: prod_db.yourcompany.com or 10.0.1.20)

Port # 443

Credentials

Username

Password

TEST

CONNECT

4. Enter the hostname or IP address of the running coordinator.
You can get the coordinator hostname from the JDBC URL of the Hive DW.
5. Use port 443.
6. Enter your workload username and password as credentials.
7. Click the Advanced tab to configure the additional details.

Create New Data Connection ✕

Connection type

Connection name

Basic **Advanced** Parameters Data

Connection mode Binary HTTP

HTTP Path

Access Token

Socket type Normal SSL SSL with certificate

Authentication mode NoSasl Plain LDAP Kerberos

Socket Timeout

Impersonation ⓘ Enabled

Trusted ⓘ Impersonation Enabled

Application Name

- 8. Click the Parameters tab and set the hive.server2.async.exec.async.compile parameter to false.

Create New Data Connection ✕

Connection type

Connection name

Basic Advanced **Parameters** Data

Parameter Name	Parameter Value	
hive.server2.async.exec.async.compile	false	
Add new row		

9. Check the Data tab for more configuration options.

Create New Data Connection ✕

Connection type

Connection name

Basic Advanced Parameters **Data**

Concurrency

Concurrency Per User

Query Timeout (Minutes)

Query Still Loading Warning (Seconds)

Row upload limit

Result Cache Enabled

Cache Retention Time (seconds)

10. Once you finish configuring the settings, click TEST to test the connection.

11. Click CONNECT to establish the connection.

Results

You have successfully set up a connection to a running Hive DW.

Related Information

[RBAC permissions](#)

[Setting role privileges](#)

[Promoting a user to administrator](#)

Creating a Cloudera Data Warehouse data connection in Cloudera Data Visualization

Learn how to connect to data when using Cloudera Data Visualization in Cloudera Data Warehouse data service. You can connect Cloudera Data Visualization to a Virtual Warehouse to visualize your data. Similar to using a BI client, you can configure and connect to Virtual Warehouses from different clusters.

About this task

You must connect to your data prior to using the data modeling and visualization functions. You make the connection to the Virtual Warehouse when you select your warehouse in the steps below. The Cloudera Data Warehouse URL has the same compute instance ID as your Virtual Warehouse.



Note:

To create and manage connections in Cloudera Data Visualization, you must have the Manage data connections privilege or hold administrative privileges. In Cloudera Data Warehouse, these are the members of the Admin Groups associated with the Cloudera Data Visualization instance.

When you create a connection, you automatically gain privileges to create and manage datasets associated with this connection, and to build dashboards and visuals within these datasets.

- For more information on the Manage data connections privilege, see *RBAC permissions*.
- For instructions on how to define privileges for a specific role, see *Setting role privileges*.
- For instructions on how to assign the administrator role to a user, see *Promoting a user to administrator*.

When you are creating a Hive or Impala data connection within the same cluster, the connection is considered secure and trusted, and the connection details can be auto-populated with a default authentication user.

Procedure

1. Start Cloudera Data Visualization from the left navigation panel in Cloudera Data Warehouse.
2. On the main navigation bar, click DATA.

The DATA interface appears, open on the Datasets tab.

Title/Table	ID	Tags	Created	Last Updated	IF	Modified By	# Dashboards
test-import-dash-board-dataset-association default.bbviz2077	546		Jun 20, 2024	a month ago		vizapps_admin	11
customers default.customers	527		May 31, 2024	2 months ago		vizapps_admin	0
Clone of Hr Capitalized default.hr_partitioned	514		May 22, 2024	2 months ago		vizapps_admin	0
tbl_from_sql text.sql, default.bbviz2077, default.bbviz2077, ...	470		Apr 09, 2024	4 months ago		vizapps_admin	0
web_logs_test_joins default.web_logs, default.web_logs	469		Apr 09, 2024	4 months ago		vizapps_admin	0
Hr Capitalized default.hr_partitioned	441		Mar 12, 2024	4 months ago		ckoncz	1

3. In the side menu bar of DATA, click NEW CONNECTION.



Note: The NEW CONNECTION button is only accessible to users assigned to roles with the Manage data connections privilege and to administrators.

The screenshot shows the Cloudera Data Visualization interface. The top navigation bar includes 'HOME', 'VISUALS', and 'DATA'. The 'DATA' tab is active, and a search bar is present. In the left sidebar, the 'NEW CONNECTION' button is highlighted with an orange box. Below it, there are sections for 'All Connections' and 'samples'. The main content area shows a 'NEW DATASET' and 'ADD DATA' button, followed by a 'Datasets' section with a 'Connection Explorer' link. A table lists several datasets:

Title/Table	ID	Created
Test Dataset main.census_pop	13	Dec 08, 2021
Food Stores Inspection in NYC main.retail_food_store_inspections_current_critical_vio_...	12	Nov 22, 2021
Cereals main.cereals	11	Nov 22, 2021
Earthquake Data January 2019 main.earthquake_data2019	10	Nov 22, 2021
World Life Expectancy main.world_life_expectancy	9	Nov 22, 2021

The **Create New Data Connection** modal window appears.

4. In Connection type, select CDW Hive or CDW Impala, and provide a name for your connection.



Note: SQLite connection is not supported in Cloudera Data Warehouse, and the option is disabled in the connection list.

Create New Data Connection
✕

Connection type

Connection name

Basic
Advanced
Parameters
Data

Hostname or IP address

(example: prod_db.yourcompany.com or 10.0.1.20)

Port #

Credentials

Username

Password

TEST

CONNECT

5. Select a Cloudera Data Warehouse warehouse to connect to.

For Data connection within the same cluster

The following fields are auto-populated:

- Hostname or IP address
- Port #
- Username

For Data connection outside the cluster

Enter the following information:

- Hostname or IP address
- Port #
- Username
- Password

6. Click the Advanced tab and configure the additional details.



Important: Depending on the type of connection you are creating, there can be additional tabs in the Create New Data Connection modal window where you have to adjust further settings.

7. Click TEST.

If the connection is valid, the system returns a Connection Verified message.

8. Click CONNECT.

What to do next

You can create a data set, and then start creating visuals, dashboards, and applications. For more information, see *Creating datasets* and *Creating a visual*.

Related Information

[RBAC permissions](#)

[Setting role privileges](#)

[Promoting a user to administrator](#)

[Creating datasets](#)

[Creating a visual](#)

Creating a Cloudera Data Science Workbench data connection to a data warehouse

Learn how to connect natively to data stored in a data warehouse when using Cloudera Data Visualization in Cloudera Data Science Workbench.

About this task

You must connect to your data prior to using the data modeling and visualization functionalities. The following steps show you how to create a new Cloudera Data Science Workbench data connection to a running Impala system.



Note:

To create and manage connections in Cloudera Data Visualization, you must have the Manage data connections privilege or hold administrative privileges.

Setting user privileges requires administrator-level access. To log in as an administrator, you can use the default admin account with the following credentials:

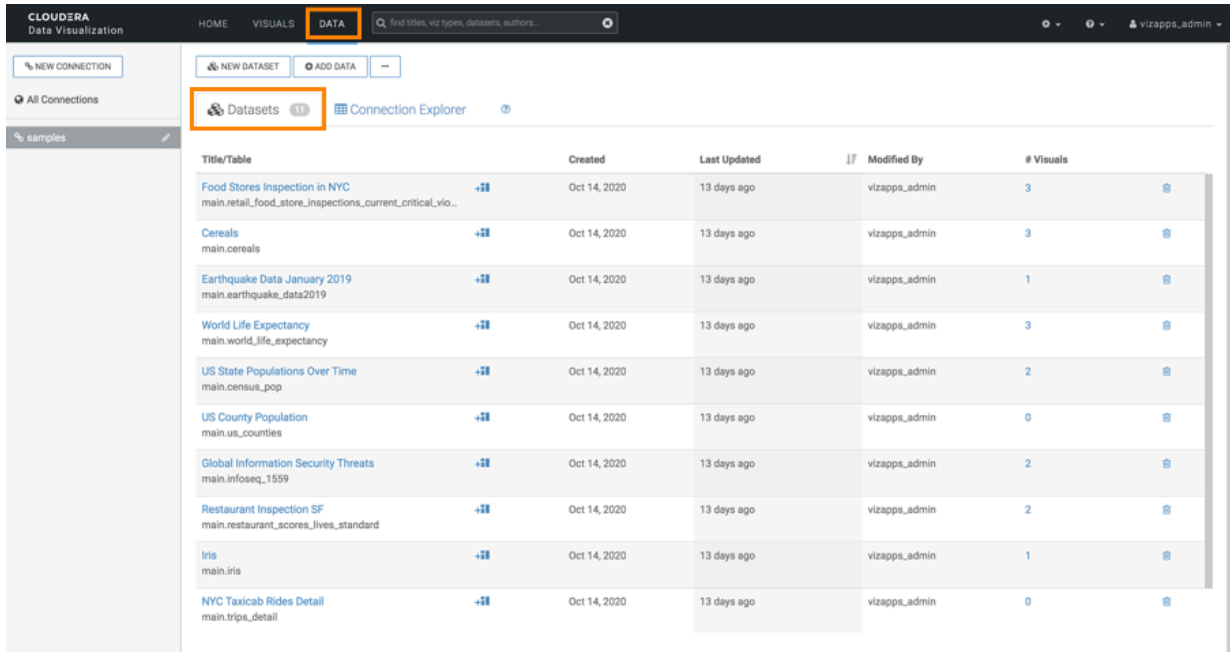
- Username: vizapps_admin
- Password: vizapps_admin

When you create a connection, you automatically gain privileges to create and manage datasets associated with this connection, and to build dashboards and visuals within these datasets.

- For more information on the Manage data connections privilege, see *RBAC permissions*.
- For instructions on how to define privileges for a specific role, see *Setting role privileges*.
- For instructions on how to assign the administrator role to a user, see *Promoting a user to administrator*.

Procedure

1. On the main navigation bar, click DATA.
The DATA interface appears, open on the Datasets tab.



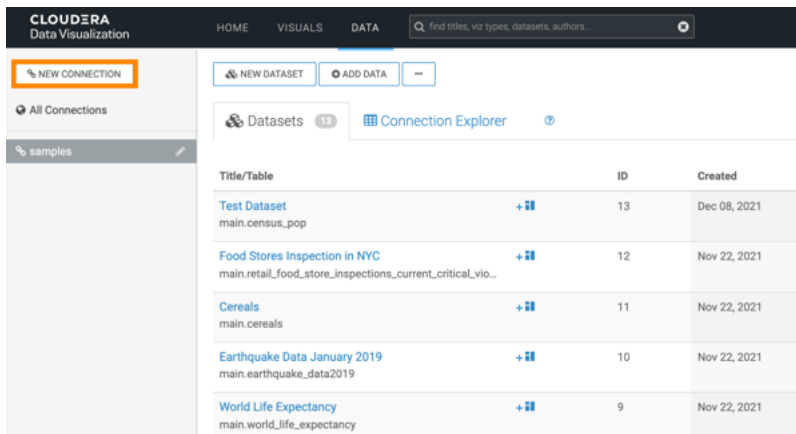
The screenshot shows the Cloudera Data Visualization interface. The top navigation bar includes 'HOME', 'VISUALS', and 'DATA' (highlighted with an orange box). Below the navigation bar, there are buttons for 'NEW CONNECTION', 'NEW DATASET', and 'ADD DATA'. The 'Datasets' tab is selected and highlighted with an orange box. The main content area displays a table of datasets with the following columns: Title/Table, Created, Last Updated, Modified By, and # Visuals.

Title/Table	Created	Last Updated	Modified By	# Visuals
Food Stores Inspection in NYC main.retail_food_store_inspections_current_critical_vio...	Oct 14, 2020	13 days ago	vizapps_admin	3
Cereals main.cereals	Oct 14, 2020	13 days ago	vizapps_admin	3
Earthquake Data January 2019 main.earthquake_data2019	Oct 14, 2020	13 days ago	vizapps_admin	1
World Life Expectancy main.world_life_expectancy	Oct 14, 2020	13 days ago	vizapps_admin	3
US State Populations Over Time main.census_pop	Oct 14, 2020	13 days ago	vizapps_admin	2
US County Population main.us_counties	Oct 14, 2020	13 days ago	vizapps_admin	0
Global Information Security Threats main.infoseq_1559	Oct 14, 2020	13 days ago	vizapps_admin	2
Restaurant Inspection SF main.restaurant_scores_lives_standard	Oct 14, 2020	13 days ago	vizapps_admin	2
Iris main.iris	Oct 14, 2020	13 days ago	vizapps_admin	1
NYC Taxicab Rides Detail main.trips_detail	Oct 14, 2020	13 days ago	vizapps_admin	0

2. In the Data side menu bar, click NEW CONNECTION.



Note: The NEW CONNECTION button is only accessible to users assigned to roles with the Manage data connections privilege and to administrators.



The screenshot shows the Cloudera Data Visualization interface with the 'NEW CONNECTION' button highlighted in the side menu. The main content area displays a table of datasets with the following columns: Title/Table, ID, and Created.

Title/Table	ID	Created
Test Dataset main.census_pop	13	Dec 08, 2021
Food Stores Inspection in NYC main.retail_food_store_inspections_current_critical_vio...	12	Nov 22, 2021
Cereals main.cereals	11	Nov 22, 2021
Earthquake Data January 2019 main.earthquake_data2019	10	Nov 22, 2021
World Life Expectancy main.world_life_expectancy	9	Nov 22, 2021

The Create New Data Connection modal window appears.

3. Select a Connection type from the drop-down list, and provide a name for the connection.
4. Enter the hostname or IP address of the running coordinator.
5. Under Port #, enter the port number.
6. Use your workload username and password as credentials.
7. Click the Advanced tab and make the appropriate selections.



Important: Depending on the type of connection you are creating, there can be additional tabs in the Create New Data Connection modal window where you have to adjust further settings.

8. Click TEST.

If the connection is valid, the system returns a Connection Verified message.

9. Click CONNECT.

Results

You have set up a connection to a running data warehouse.

Related Information

[RBAC permissions](#)

[Setting role privileges](#)

[Promoting a user to administrator](#)

Creating a visual


You can easily create new visual representations for your data in Cloudera Data Visualization.

About this task

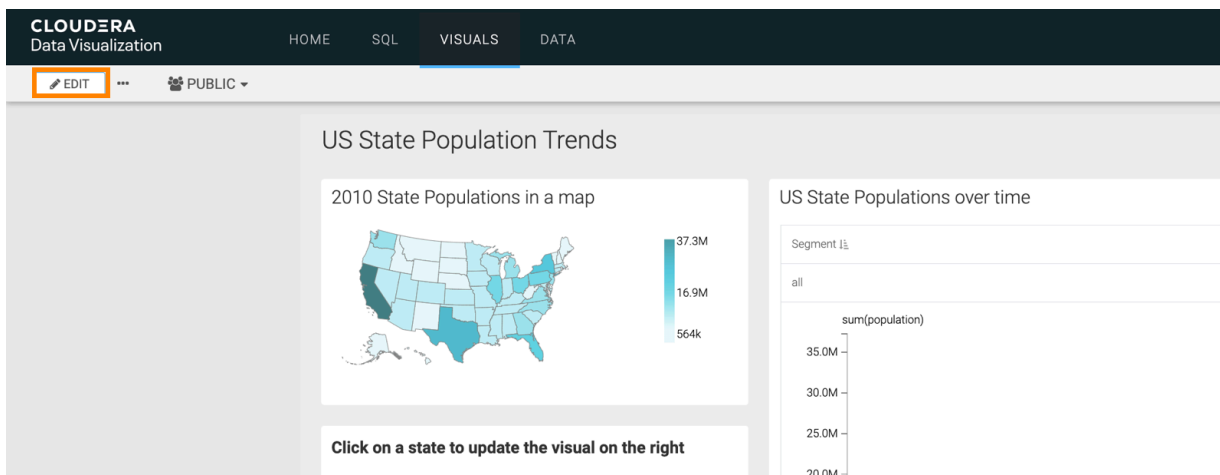
The following steps demonstrate how to create a new visual.

Procedure

1. On the main navigation bar, click VISUALS.
2. Click Create New Dashboard to open a new, blank dashboard.



If you already have an existing dashboard and you just want to add a new visual, open the dashboard and click EDIT to make changes to it.



- Choose the visual type you want to build.
In this example, the Bars visual has been selected.

Dashboard Designer

VISUALS

Bars

1234 LABEL

WORD CLOUD

ACTION

SQL

* X Axis

A content

* Y Axis

Record Count

Colors

drag fields to add here

Tooltips

drag fields to add here

Drill

drag fields to add here

Labels

drag fields to add here

Filters

drag fields to add here

REFRESH VISUAL

DATA

cm_logs_DL

Sample Mode: OFF

Search

Dimensions 3

cm_logs_goes_v2

A content

A timeoccurred

A severity

Measures 1

cm_logs_goes_v2

Record Count

DASH.

Visuals

Filters

Settings

Style

VISUAL

Build

Settings

Style

6. You can switch on Sample Mode for the visual you are building.

With sample mode, you create the visual on a subset of your data to test it. You can define this subset as a percentage of the original data.



Note: This feature is only available on Impala, Hive, and Spark connections.

Run in Sample Mode

Running in sample mode will allow you to run on a subset of your data which will make the chart load faster.

Sample Mode: % of data

7. Populate the shelves of the visual from the Dimensions and Measures fields:

In this example, the following fields have been added:

- state dimension on the X Axis shelf
- population measure on the Y Axis shelf



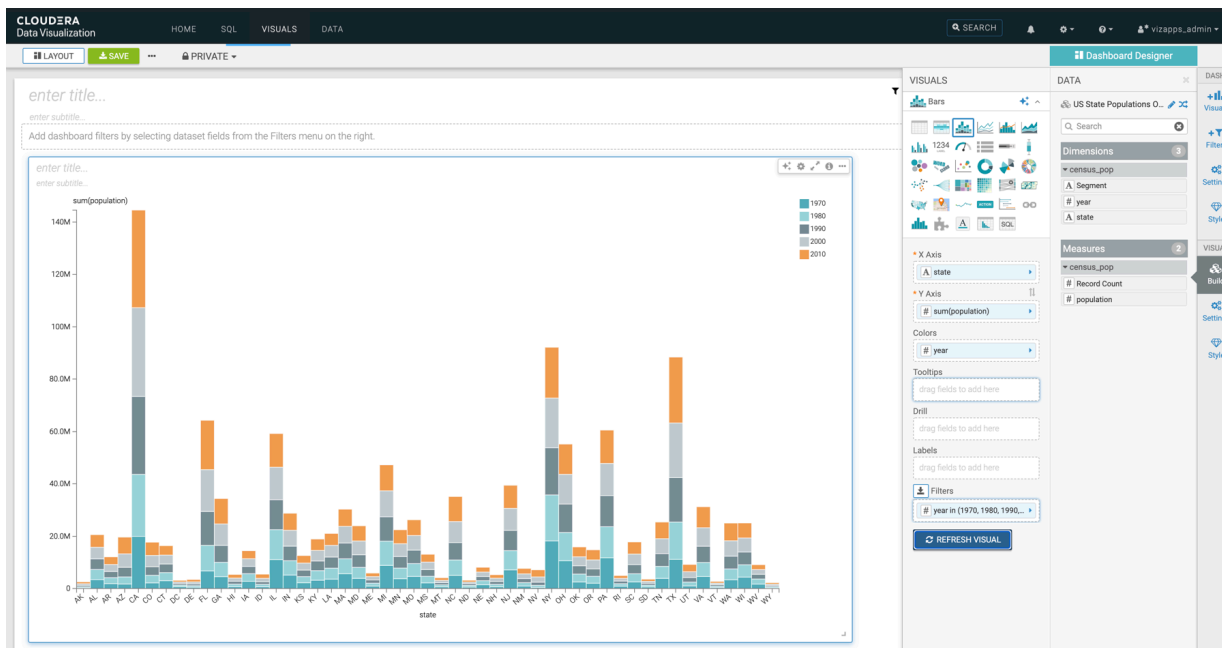
Note: The population measure appears with its default aggregation, as sum(population).

- year measure on the Colors shelf
- year measure on the Filters shelf



Note: The Filter for year module opens automatically. Years 1970, 1980, 1990, 2000, and 2010 have been selected.

8. Click REFRESH VISUAL.



9. Click the enter title... field and enter a name for the visual.

- In this example the title 'Recent US Census Results by State' has been added.
- You can also click the enter subtitle... field below the title of the visualization to add a brief description of the visual.

10. Click SAVE at the top left corner of the Dashboard Designer.