

Data Catalog Overview

Date published: 2023-10-10

Date modified: 2024-03-05

CLOUDERA

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

About Data Catalog.....	4
Data Catalog Terminology.....	5
Before you start.....	6
CDP Private Cloud Base cluster requirements.....	6
Prerequisite to access Data Catalog service.....	6
How to access Data Catalog service.....	8
Authorization for viewing Assets.....	9
Data Catalog Dashboard.....	10
Restricting access for certain users of Data Catalog.....	12
Understanding Datasets.....	15
Understanding Data Assets.....	16
About the Data Catalog Profiler.....	16
Understanding the Cluster Sensitivity Profiler.....	16
Understanding the Hive Column Profiler.....	18
Understanding the Ranger Audit Profiler.....	19

About Data Catalog

Data Catalog is a service within Cloudera Data Platform that enables you to understand, manage, secure, and govern data assets across the enterprise.

Data Catalog helps you understand data lying in your data lake (Private Cloud Base Cluster). You can search to locate relevant data of interest based on various parameters. Using Data Catalog, you can understand how data is interpreted for use, how it is created and modified, and how data access is secured and protected.



Data Catalog enables data stewards across the enterprise to work with data assets in the following ways:

- Organize and curate data globally
 - • Organize data based on business classifications, purpose, protections needed, etc.
 - • Promote responsible collaboration across enterprise data workers
- Understand where relevant data is located
 - • Catalog and search to locate relevant data of interest (sensitive data, commonly used, high risk data, etc.)
 - • Understand what types of sensitive personal data exists and where it is located
- Understand how data is interpreted for use
 - • View basic descriptions: schema, classifications (business cataloging), and encodings

- View statistical models and parameters
- View user annotations, wrangling scripts, view definitions etc.
- Understand how data is created and modified
 - Visualize upstream lineage and downstream impact
 - Understand how schema or data evolve
 - View and understand data supply chain (pipelines, versioning, and evolution)
- Understand how data access is secured, protected, and audited
 - Understand who can see which data and metadata (for example, based on business classifications) and under what conditions (security policies, data protection, anonymization)
 - View who has accessed what data from a forensic audit or compliance perspective
 - Visualize access patterns and identify anomalies

Related Information

[Data Catalog Terminology](#)

Data Catalog Terminology

An overview of terminology used in Data Catalog service.

Profiler

Enables the Data Catalog service to gather and view information about different relevant characteristics of data such as shape, distribution, quality, and sensitivity which are important to understand and use the data effectively. For example, view the distribution between males and females in column Gender, or min/max/mean/null values in a column named avg_income. Profiled data is generated on a periodic basis from the profilers, which run at regularly scheduled intervals. Works with data sourced from Apache Ranger Audit Logs, Apache Atlas Metadata Store, and Hive.

Data Lake

A trusted and governed data repository that stores, processes, and provides access to many kinds of enterprise data to support data discovery, data preparation, analytics, insights, and predictive analytics. In the context of Cloudera Data Platform, a Data Lake can be realized in practice with an Cloudera Manager enabled CDP cluster that runs Apache Atlas for metadata and governance services, and Apache Ranger for security services.

ECS

The Embedded Container Service (ECS) service enables you to run CDP Private Cloud Data Services by creating container-based clusters in your data center. In addition to the option to use OpenShift, which requires that you deploy and manage the Kubernetes infrastructure, you can also deploy a Embedded Container Service cluster, which creates and manages an embedded Kubernetes infrastructure for use with CDP Private Cloud Data Services.

OpenShift Container (OCP)

OpenShift is an enterprise platform for container orchestration.

Data Asset

A data asset is a physical asset located in the CDP ecosystem such as a Hive table which contains business or technical data. A data asset could include a specific instance of an Apache Hive database, table, or column. An asset can belong to multiple asset collections. Data assets are equivalent to “entities” in Apache Atlas.

Datasets

Datasets allow users of Data Catalog to manage and govern various kinds of data objects as a single unit through a unified interface. Asset collections help organize and curate information about many assets based on many facets including data content and metadata, such as size/schema/tags/

alterations, lineage, and impact on processes and downstream objects in addition to the display of security and governance policies.

Related Information

[About Data Catalog](#)

Before you start

Before you access Cloudera Data Catalog on Cloudera Data Platform (CDP) Private Cloud Data Services, you must deploy the Data Catalog service.

Before deploying Data Catalog, make sure you have reviewed and compiled with the requirements in the installation guide for your environment:

Data Catalog is supported on both **OpenShift** and **ECS** clusters. For more information, see:

- [Installing on OpenShift](#)
- [The Embedded Container Service \(ECS\)](#)

CDP Private Cloud Base cluster requirements

The CDP Private Cloud Base cluster is used for the Cloudera Data Catalog service.

As Atlas and Ranger services are installed in the Base cluster, Data Catalog communicates with these services to fetch metadata and security policies. Data Catalog profilers connect with Hive and HDFS services running in the Base cluster for data profiling operations.

For more information, check the following links:

- [Supported base clusters for Data Services and supported Cloudera Manager versions](#)
- [Supported Private Cloud Data Services versions](#)

Prerequisite to access Data Catalog service

To access Data Catalog service, you must have the required credentials.

Follow these instructions to provide the required access to the Data Catalog users.

The PowerUser must provide the requisite access to subscribers who plan to use Data Catalog as per the requirement.



Note: To launch and delete profilers you must be a PowerUser.

There are a couple of roles based on which you can use Data Catalog:

- DataCatalogCspRuleManager
- DataCatalogCspRuleViewer

Using DataCatalogCspRuleManager role, you can create, deploy new Custom Sensitivity Profiler rules, create new Regex, and run validations on newly created rules.

Using DataCatalogCspRuleViewer role, you can list and view the existing Custom Sensitivity Profiler rules.

Additionally, using Cloudera Manager, you must configure Atlas and Ranger services. Use the following instructions to complete the process.

1. Cloudera Manager > Clusters > Atlas > Configuration .

2. Enable the Kerberos Authentication for the Atlas service.

Enable Kerberos Authentication

atlas.authentication.method.kerberos

[kerberos.auth.enable](#)

☒ **ATLAS-1 (Service-Wide)**

3. In the search bar, look up the following property: `conf/atlas-application.properties_role_safety_valve` and enter the values for the Atlas service:

atlas.proxyuser.dpprofiler.hosts=*

atlas.proxyuser.dpprofiler.users=*

atlas.proxyuser.dpprofiler.groups=*

4. For the Ranger configuration updates, go to Cloudera Manager > Clusters > Ranger > Configuration .
5. In the search bar, look up the following property: `conf/ranger-admin-site.xml` and enter the values for the Ranger service:

Name: `ranger.proxyuser.dpprofiler.hosts`

Value: *

Name: `ranger.proxyuser.dpprofiler.users`


Value: *

Name: `ranger.proxyuser.dpprofiler.groups`

Value: *



Note: You can add new rows using the + icon.



Later, restart Atlas and Ranger services respectively.

How to access Data Catalog service

Before you commence using Data Catalog service note that when the Data Services installation is complete, the CDP Private Cloud Base cluster registration is automatically enabled and Data Catalog as a service is installed on the ECS cluster.

You can navigate to the Data Catalog service by clicking on the CDP homepage on your CDP console.



Data Catalog is supported on both **OpenShift** and **ECS** clusters. For more information, see:

[Installing on OpenShift](#)

[The Embedded Container Service \(ECS\)](#)

[About CDP Private Cloud Data Services](#)

Authorization for viewing Assets

Data Catalog users must have appropriate authorization set-up in Ranger to view assets.

Hive Ranger Policy - You must set-up Hive Ranger policies as per your requirement to work with Hive assets in Data Catalog.

For example, the following diagram provides a sample Hive Ranger policy.

Service Manager

Hadoop SQL Policies

Create Policy

Create Policy

Policy Details :

Policy Type

Access

Add Validity Period

Policy Name *

Hive_Ranger

enabled

normal

Policy Label

Hive

database

DB

include

none

Description

Hive Ranger Policy

Audit Logging

YES

Allow Conditions :

hide

Select Role	Select Group	Select User	Permissions	Delegate Admin	
	Select Groups		AllAlterCreateDropIndexLockReadRefreshReplAdminselectService AdminTemporary UDF AdminupdateWrite	☐	

Atlas Ranger Policy- You must set-up Ranger policies for Atlas in order to work with asset search and Tag flow management.

For example, the following diagram provides a sample Atlas Ranger policy.

Service Manager

cm_atlas Policies

Create Policy

Create Policy

Policy Details :

Policy Type

Access

Add Validity Period

Policy Name *

Atlas_Policy

enabled

normal

Policy Label

Atlas

Ranger

type-category *

include

Type Name *

Atlas

include

Description

Atlas Ranger Policy

Audit Logging

YES

Allow Conditions :

hide


Select Role	Select Group	Select User	Permissions	Delegate Admin	
Select Roles	Select Groups		Create TypeDelete TypeUpdateType	☐	

Related Information
[Restricting access for certain users of Data Catalog](#)
[Ranger Policies Overview](#)


Data Catalog Dashboard

The Dashboard provides quick access to vital service information at a glance, in the form of visual, actionable navigation for multiple operations. The user-friendly navigation enables viewing, filtering, and acting upon data quickly and in a simple manner.


Data Stewards can view the homepage at a glance, and also focus on the most important tasks, enabling faster decision making as well as immediate action. The application lets you perform multiple actions for different types of content that helps in visualizing information with ease.



The displayed sections (panes) are fully interactive, with clickable areas for easy navigation to relevant parts of applications. Users can access individual sections and narrow down the information displayed,



The overview page contains information pertaining to the Data Lakes, the total number of assets that are profiled, along with the assets that are scanned for data.




Additionally, you can manage the datasets, check the status of profilers, and search or discover assets.

Restricting access for certain users of Data Catalog

To have a fine-grained access to the same user from accessing the assets in Data Catalog, you can perform some additional changes. For example, if you want to restrict some users from accessing specific table information, you must set-up a Ranger policy such that these users will not have access to the asset details in Data Catalog.

To create the Ranger policy to restrict users from accessing asset details, refer to the following images:



Edit Policy

Last Response Time
10/31/2023 05:00:08 PM

[Service Manager](#) > [cm_atlas Policies](#) > [Edit Policy](#)

Policy Details

Policy Type

Access

[Add Validity Period](#)

Policy ID*

81

Policy Name*

rs_test

Enabled

Normal

Policy Label

Select...

Entity Type

hive_table

hbase_table

Include

Entity Classification *

rs_test

Include

Entity ID *

x

Include

None

Select...

Description

Audit Logging*

Yes

Allow Conditions:

hide

Select Roles	Select Groups	Select Users	Permissions	Delegate Admin	
Select...	Select...	Select...	Add Permissions +	☐	x

The following image displays the **Deny Conditions** set for the specific user.

Deny Conditions :

hide

Select Role

Select Groups

Select User

Permissions

Delegate Admin

Select Roles

Select Groups

Read Entity

☐

x

+

Exclude from Deny Conditions :

hide

Select Role

Select Groups

Select Users

Permissions

Delegate Admin

Select Roles

Select Groups

Select Users

Add Permissions

☐

x

+

Save

Cancel

Delete

Edit Policy

Last Response Time
10/31/2023 05:00:08 PM

[Service Manager](#) > [cm_atlas Policies](#) > [Edit Policy](#)

Exclude from Allow Conditions:

Select Roles	Select Groups	Select Users	Permissions	Delegate Admin	
Select...	Select...	Select...	Add Permissions		

Deny All Other Accesses: *

False

Deny Conditions:

hide

Select Roles	Select Groups	Select Users	Permissions	Delegate Admin	
Select...	Select...	hrt_sze_admin X	Read Entity Create Entity Update Entity Delete Entity		

Exclude from Deny Conditions:

Select Roles	Select Groups	Select Users	Permissions	Delegate Admin	
Select...	Select...		Add Permissions		

The result is depicted in the following image, where the user has no permissions to access the specified dataset (us_c customers).

dc-pro-ibmj9

[Ranger](#)[Atlas](#)[Create Dataset](#)

Search

Filters

TYPE

Hive Table

HBase Table

HBase Namespace

ML Model Build

Add New Value

OWNER

Add New Value

ww_customers

Hive Table

/hortoniabank

Created: Wed Jun 17 2020

Owner: hive

Source: hive

eu_countries

Hive Table

/hortoniabank

Created: Wed Jun 17 2020

Owner: hive

Source: hive

us_customers *

Hive Table

/-NA-


Created: -NA-

Owner: -NA-

Source: hive

* Some information might not be available to unauthorised users

14



Additionally, when you plan to restrict data access, please note the following:

- Audit summarisation for the asset evolves from the Ranger audit profiler and Metrics service.
- Various Hive Column Statistical metrics for columns of the asset evolves from Atlas as part of the profile_data of a column.

To ensure that the data related to audit summary and Hive Column Statistics are not visible to the subscribers, you must make sure to turn off the audit profiler and the Hive Column profiler respectively.

Related Information
[Authorization for viewing Assets](#)

Understanding Datasets

A Dataset is a group of assets that fit search criteria so that you can manage and administer them collectively.

Asset collections enable you to perform the following tasks when working with your data:

- Organize
Group data assets into Datasets based on business classifications, purpose, protections, relevance, etc.
- Search

Find tags or assets in your data lake using Hive assets, attribute facets, or free text.

Advanced asset search uses facets of technical and business metadata about the assets, such as those captured in Apache Atlas, to help users define and build collections of interest. Advanced search conditions are a subset of attributes for the Apache Atlas type hive_table.

- Understand
Audit data asset security and use for anomaly detection, forensic audit and compliance, and proper control mechanisms.

You can edit Datasets after you create them and the assets contained within the collection will be updated. CRUD (Create, Read, Update, Delete) is supported for Datasets.

 **Note:** Datasets must have less than 130 assets.

Related Information

[About the Data Catalog Profiler](#)

Understanding Data Assets

A data asset is a specific instance of a data type, including the related attributes and metadata. A data asset is a physical asset located in the data lake, such as a Hive table, that contains business or technical data.

Data assets are also known as *entities* in Apache Atlas.

About the Data Catalog Profiler

The Data Catalog profiler employs Kubernetes enabled job scheduling and runs profilers jobs on-demand.

These profilers create metadata annotations that summarize the content and shape characteristics of the data assets.



Note: You must be a PowerUser to launch the Profiler.

Profiler Name	Description
Cluster Sensitivity Profiler	A sensitive data profiler- PII, PCI, HIPAA and others.
Ranger Audit Profiler	A Ranger audit log summarizer.
Hive Column Profiler	Provides summary statistics like Maximum, Minimum, Mean, Unique, and Null values at the Hive column level.

For example, data profilers can create summarized information about contents of an asset and also provide annotations that indicate its shape (such as distribution of values in a box plot or histogram).

Related Information

[Understanding Datasets](#)

Understanding the Cluster Sensitivity Profiler

The Cluster Sensitivity Profiler automatically performs context and content inspection to detect various types of sensitive data and suggest suitable classifications or tags based on the type of sensitive content detected or discovered.

Auto-detected data types

Type of data

- Bank account
- Credit card
- Driver number (UK)
- Email

- IBAN number
 - Austria (AUT)
 - Belgium (BEL)
 - Bulgaria (BGR)
 - Switzerland (CHE)
 - Cyprus (CYP)
 - Czech Republic (CZE)
 - Germany (DEU)
 - Denmark (DNK)
 - Spain (ESP)
 - Estonia (EST)
 - Finland (FIN)
 - France (FRA)
 - United Kingdom (GBR)
 - Greece (GRC)
 - Croatia (HRV)
 - Hungary (HUN)
 - Ireland (IRL)
 - Iceland (ISL)
 - Italy (ITA)
 - Liechtenstein (LIE)
 - Lithuania (LTU)
 - Latvia (LVA)
 - Luxembourg (LUX)
 - Malta (MLT)
 - Netherlands (NLD)
 - Norway (NOR)
 - Poland (POL)
 - Portugal (PRT)
 - Romania (ROU)
 - Slovakia (SVK)
 - Slovenia (SVN)
 - Sweden (SWE)
- IP address
- NPI
- Name


- National ID number
 - Bulgaria (BGR)
 - Switzerland (CHE)
 - Czech Republic (CZE)
 - Denmark (DNK)
 - Spain (ESP)
 - Estonia (EST)
 - Finland (FIN)
 - Greece (GRC)
 - Ireland (IRL)
 - Iceland (ISL)
 - Italy (ITA)
 - Lithuania (LTU)
 - Latvia (LVA)
 - Norway (NOR)
 - Poland (POL)
 - Portugal (PRT)
 - Romania (ROU)
 - Slovakia (SVK)
 - Sweden (SWE)
- National insurance number (UK)
- Passport number
 - Austria (AUT)
 - Belgium (BEL)
 - Switzerland (CHE)
 - Germany (DEU)
 - Spain (ESP)
 - Finland (FIN)
 - France (FRA)
 - Greece (GRC)
 - Ireland (IRL)
 - Italy (ITA)
 - Poland (POL)
 - United Kingdom (UK)
- Bank Routing Number
- US Social Security Number
- Society for Worldwide Interbank Financial Telecommunication (SWIFT)
- Telephone

Understanding the Hive Column Profiler


You can view the shape or distribution characteristics of the columnar data within a Hive table based on the Hive Column Profiler.

There are different charts available to help visualize the shape and distribution of the data within the column as well as summary statistics (such as means, null count, and cardinality of the data) for a column. The profiler computes column univariate statistics that are displayed using an appropriate chart in the Schema tab.

Pie charts are presented for categorical data with limited number of categories or classes. Examples include data such as eye colors that only have a fixed list of values (categories or labels).



When the data within columns is numeric, a histogram of the distribution of values organized into 10 groups (decile frequency histogram) and a box plot with a five-number summary (mean, median, quartiles, maximum, and minimum values) are shown for the column.




Related Information
[Understanding the Ranger Audit Profiler](#)

Understanding the Ranger Audit Profiler

You can view who has accessed which data from a forensic audit or compliance perspective, visualize access patterns, and identify anomalies in access patterns using the Ranger Audit Profiler.

The audit profiler uses the Apache Ranger audit logs to show the most recent raw audit event data as well as summarized views of audits by type of access and access outcomes (allowed/denied). Such summarized views are obtained by profiling audit records in the data lake with the audit profiler.



Related Information

[Understanding the Hive Column Profiler](#)