

Atlas Classifications

Date published: 2020-07-28

Date modified: 2024-12-10



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.

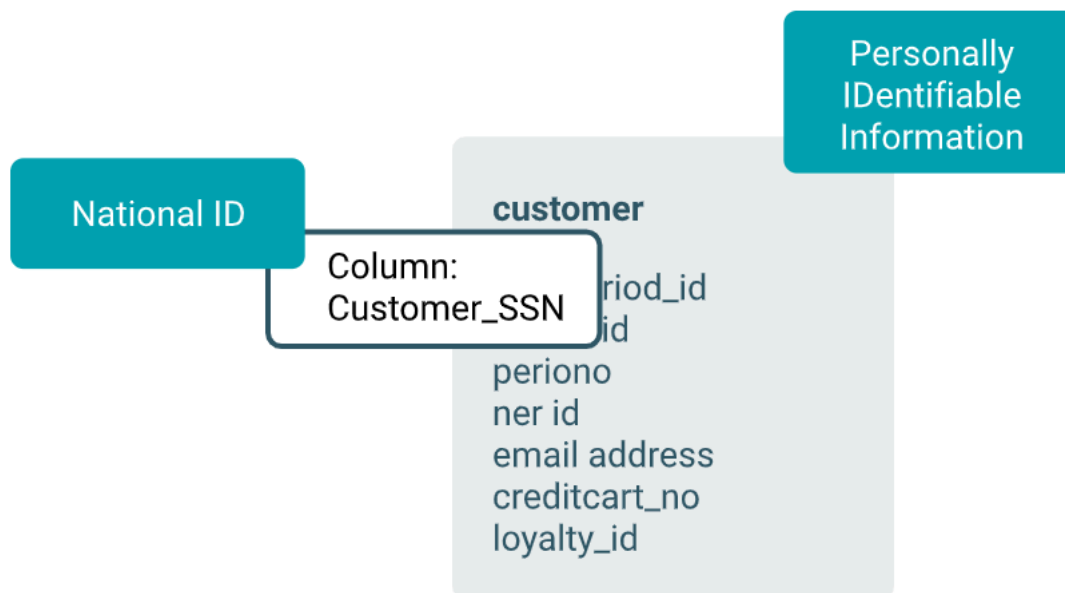
Contents

Working with Atlas classifications and labels.....	4
Text-editor for Atlas parameters.....	5
Creating classifications.....	7
Example for finding parent object for assigned classification or term.....	8
Creating labels.....	13
Adding attributes to classifications.....	14
Support for validating the AttributeName in parent and child TypeDef.....	16
Validations for parent types.....	16
Case for implementing backward compatibility.....	17
Associating classifications with entities.....	18
Propagating classifications through lineage.....	19
Searching for entities using classifications.....	21

Working with Atlas classifications and labels

Add metadata to Atlas entities using labels and classifications.

You can add metadata to Atlas entities to help your organization find, organize, and share your understanding of the data assets that drive business processes. Atlas provides two mechanisms for adding metadata to entities: labels and classifications. Both labels and classifications can be applied to entities to help describe the entity's content, status, or other business-driven value.



Labels

Labels are words or phrases (strings) that you can associate with an entity and reuse for other entities. They are a light-weight way to add information to an entity so you can find it easily and share your knowledge about the entity with others.

Anyone can create labels and associate labels with entities.

Classifications

Classifications are strings like labels, with added complexity and structure:

- Atlas includes precise search tools for finding entities using classifications.
- Classifications can automatically propagate to additional entities through lineage relationships.
- You can use classifications to drive access control policies in Ranger.
- You can enrich a classification with attributes in the form of key-value pairs and set the value to describe a particular entity.

Administrators can control who can create and assign classifications through Atlas policies in Ranger.

In short, use labels for annotating entities; use classifications to involve entities in processes inside and outside Atlas.

More about classifications: attributes

You can add key-value pairs, or *attributes* to the definition of a classification. A typical use for classification attributes would be to refine the meaning of a general category. Data assets identified with a classification of “PII” or Personally Identifiable Information can have classification attributes that indicate the nature of the information to drive data masking or expiration policies. Columns tagged with “PII” might be further separated into phone numbers, credit card numbers, and “other.” A Ranger policy based on the classification can use the attribute values to identify masks for the phone and credit card numbers and to block columns tagged as “PII” with attribute “Other.”



Planning for classifications and labels

Here are some questions to help you think about how you define classifications and labels in your system:

- Will you use the metadata to drive a workflow outside of Atlas? If so, use a classification. If not, a label may work fine.
- Do you need the text of the metadata to be exact values? You can limit who has the ability to create classifications: you can manage a single list for an organization. The ability to apply a classification to an entity can be controlled separately from defining classifications, which lets you potentially allow more users to applying the "official" classifications. Any users can create and add labels to entities, so it may be more difficult to standardize the content of labels.
- What types of entities do you expect the metadata to apply to? Labels can be added to any entity type; classifications can be defined to apply to a specific entity type. If you define a classification to only to apply to certain entity types, such as table columns, make sure that the name and description helps data stewards use the classification correctly.
- Are you adding metadata to make Atlas searches easier? Consider using classifications with attributes to allow refinement of search results. If your Atlas users are more likely to search using *terms*, you might consider connecting the classification to a glossary term so searches from either mechanism return the correct results. (Assigning the term to an entity automatically assigns the classification to the same entity.) For both labels and classifications, consider setting guidelines so everyone creating metadata uses the same conventions: underscores or hyphens? detail first or last? table-level metadata or column-level?
- Do you want a classification to follow lineage relationships and be assigned automatically as data is used in other entities? If not, you can use the classification description to help data stewards understand that they should turn off propagation when assigning this classification to entities or terms. Labels do not propagate.

Text-editor for Atlas parameters

In Atlas, while creating Classification, Glossary, and Business metadata, a new text editor is available, which can also be used for adding and customizing string attribute values of Business metadata on the Entity details page.

On the Atlas entities details page, when you add the business metadata and select the attribute type as "String", you can view a text editor with which you can customize content / description that is entered.



Important: While adding a Business metadata attribute, you **MUST** select the "Type" as "String". If you do not fulfill this criteria, you cannot add formatting to the value of the attribute using the text-editor on the Entity details page under the Business metadata tab.

Create Business Metadata

Name *

Description

B **I** **U** **S** **Link** **List** **Text** Plain Text

[+ Add Business Metadata attribute](#)

[Cancel](#) [Create](#)

[Back To Results](#) [adm](#)

Classifications: [Log Data](#) [+](#)

Terms: [Term2@G1](#) [+](#)

[Properties](#) [Lineage](#) [Relationships](#) [Classifications](#) [Audits](#) [Schema](#) [Tasks](#)

▼ **Technical properties** [Toggle](#)

createTime	1667278439138
description	log fact daily materialized view
lastAccessTime	1667278439138
name	log_fact_daily_mv
owner	Tim ETL
qualifiedName	log_fact_daily_mv@cl1
retention	1667278439138
tableType	Managed
temporary	false
typeName	Table

► **User-defined properties** [Add](#)

► **Labels** [Add](#)

▼ **Business Metadata** [Save](#) [Cancel](#)

[Add New Attribute](#)

Demo

stringType (string)

String Test

B **I** **U** **S** **Link** **List** **Text**

Note that while creating and editing Glossary, Term, and Category, the text editor is available only for the Long Description field. For creating and editing Business Metadata and Classification, text editor is available only for the Description field.

Create Glossary



Name*

G1

Short
Description

Short Description

Long
Description

1 **B** **I** **U** Plain Text

Cancel

Create

Create a new classification



Name*

Test

Description*

Test

1 **B** **I** **U** Plain Text

Select classification to inherit attributes(optional)

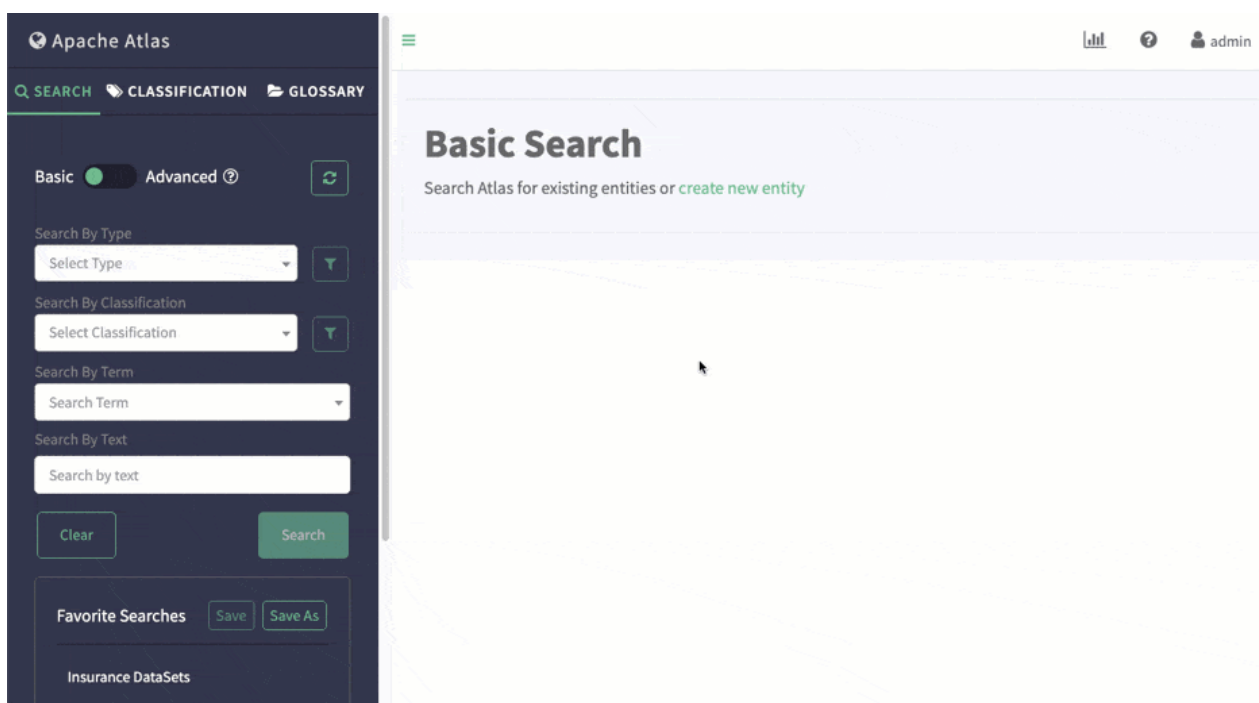
Attributes define additional properties for the classification

Search Classification

Creating classifications

Create classifications from the tab in the left navigation bar.

Create classifications from the Classifications tab in the left navigation pane.



Names can be 255 characters and contain ASCII characters and spaces. Descriptions can be 1024 characters and contain any characters.

When creating a classification, you can choose an existing classification to reuse its attribute definitions.

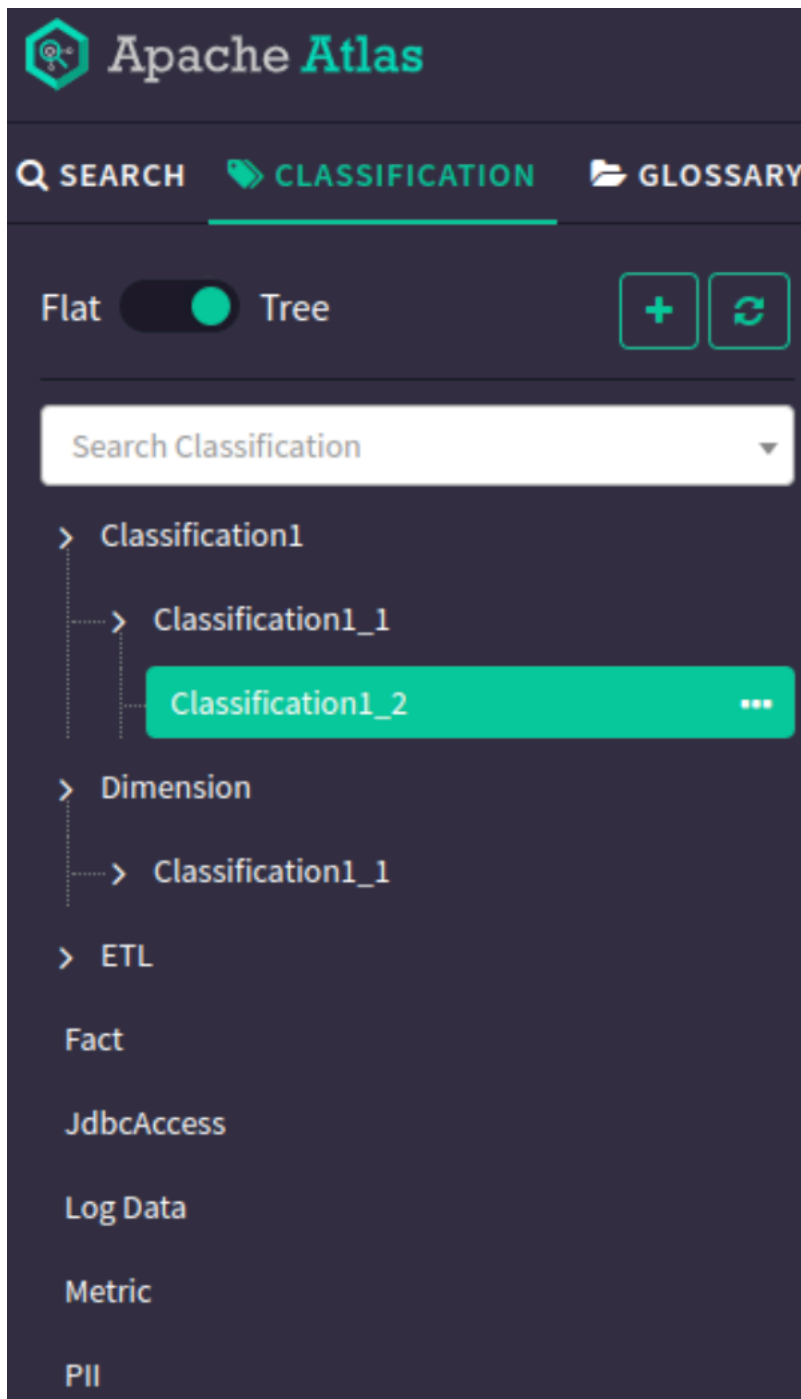
Example for finding parent object for assigned classification or term

You can view and search for Atlas entities by using the assigned name for the classification / term.

When you create an entity and assign a classification or term to the entity, the immediate parent of the assigned classification / term is displayed when you hover over the classification/term. The parent-child relationship lineage appears in the “assigned @parent1, parent2, parent3” format, depending upon the number of parent-child relationships. Click the assigned classification and term to view the details.

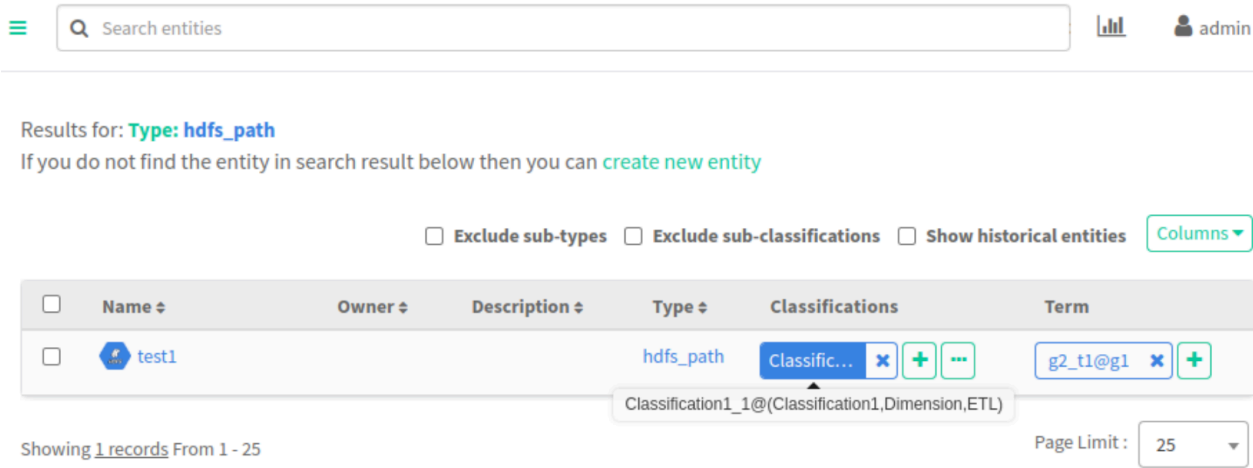
As an example, assume that there is an entity called `hdfs_path` for which a classification named “Classification1_1” is assigned. Under the “Classification” tab on the Atlas UI, you can search for the list of classifications > child classifications > child classifications > and so on.

The following image provides an example of how a classification is listed in Atlas.



The classification named Classification1_1 has parent classifications named Classification1 and Dimension respectively.

In the following image, the parent classifications are displayed on an entity for which Classification1_1 is assigned.




Note that hovering over the assigned classification displays the parent-child relationship. In this example: Classification1_1@(Classification1,Dimension,ETL)


Clicking on the assigned classification navigates to the classification details page.



The following image provides an example of how Glossary > Terms are listed in Atlas.



The term named g2_t1 has a parent term named g1, which is displayed on an entity for which g2_t1 is assigned. The parent term is visible for the assigned term in the “assigned term@parent Glossary” format. In this example: g2_t1@g1. Clicking on the assigned term navigates to the Term details page.



The respective entity pages for classification and term are as shown in the following images:

☰

[Back To Results](#)

Q

Search entities

admin

test1 (hdfs_path)

Classifications:

Classification1_1@(Classificat...

✕

+

Terms:

g2_t1@g1

✕

+

Propagated Classifications:

Dimension

Properties

Lineage

Relationships

Classifications

Audits

Tasks

✓ Technical properties

Edit

createTime

10/11/2022 12:00:00 AM (IST)

fileSize

0

isFile

false

isSymlink

false

modifiedTime

10/11/2022 12:00:00 AM (IST)

name

test1

> User-defined properties

Add

> Labels

Add

> Business Metadata

Add

☰

[Back To Results](#)

Q

Search entities

admin

test1 (hdfs_path)

Classifications:

Classification1_1@(Classificat...

✕

+

Terms:

Classification1_1@(Classification1,Dimension,ETL)

g2_t1@g1

✕

+

Propagated Classifications:

Dimension

Properties

Lineage

Relationships

Classifications

Audits

Tasks

✓ Technical properties

Edit

createTime

10/11/2022 12:00:00 AM (IST)

fileSize

0

isFile

false

isSymlink

false

modifiedTime

10/11/2022 12:00:00 AM (IST)

name

test1

> User-defined properties

Add

> Labels

Add

> Business Metadata

Add



Note: On the Term details page, you can assign terms to “Relation Types” on the “Related Terms” tab. The assigned term can be related to other terms. When you hover over the term, the assigned term displays in the “termname@glossaryname” format.

Apache Atlas

SEARCH CLASSIFICATION GLOSSARY

Terms Category

Search Glossary, Term

g1

g1_t2

g2_t1

test1

Search entities

Categories: +

Entities Properties Classifications Related Terms

Relation Types	Related Terms	Attributes
seeAlso	g1_t2@g1	
synonyms	g1_t2@g1 g1_t2@g1	
antonyms		
preferredTerms		
preferredToTerms		
replacementTerms		
replacedBy		
translationTerms		
translatedTerms		
isA		

Creating labels

Create and assign a label from the detail page for an entity.

From the detail page for an entity, the Properties tab includes a section for labels:

Results for: **Type: hive_table**
If you do not find the entity in search result below then you can [create new entity](#)

Showing 4 records From 1 - 25 ☐ Exclude sub-classifications ☐ Show historical entities [Columns ▾](#) ☐ Exclude sub-types

<input type="checkbox"/>	Name ↕	Owner ↕	Description ↕	Type ↕	Classifications	Term
<input type="checkbox"/>	custom...	admin		hive_table	+	+
<input type="checkbox"/>	web_logs	admin		hive_table	+	+
<input type="checkbox"/>	sample...	admin		hive_table	+	+
<input type="checkbox"/>	sample...	admin		hive_table	+	+

Labels can be 50 characters and contain ASCII characters, hyphens (-) and underscores (_). You can search for entities using the text from labels assigned to the entities.

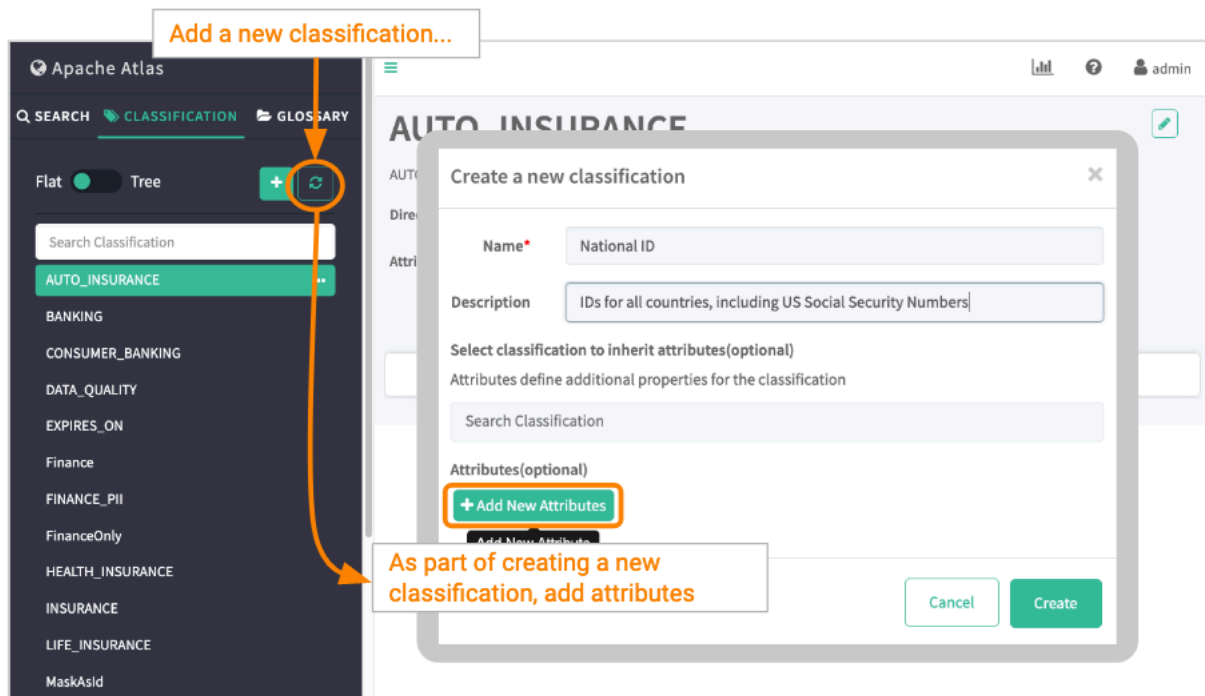
Related Information

[Using Free-text Search](#)

Adding attributes to classifications

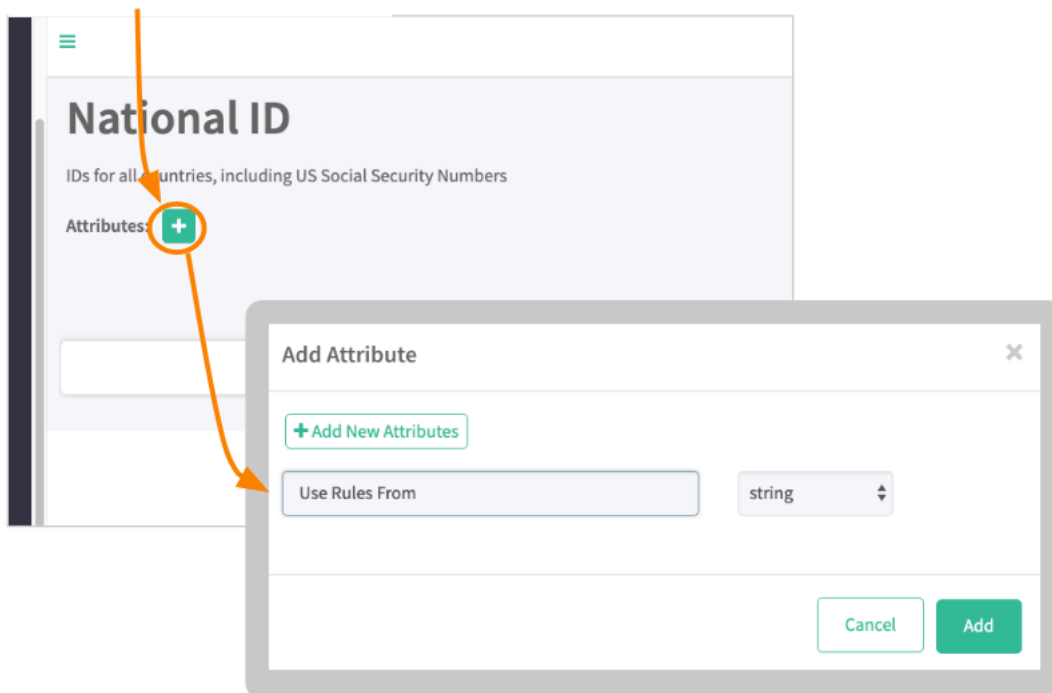
Attributes are key-value pairs defined for a classification that you can set to a value when you assign the classification to an entity

For a new classification, you can add attributes as part of the classification definition.



You can also add attributes to an existing classification. Select the classification in the list, then click the plus button next to Attributes.

From an existing classification...



Support for validating the AttributeName in parent and child TypeDef

The Apache Atlas service validates the attribute names of the entity types for those attributes having identical names as their parents' attributes.

Previously, until the Cloudera Runtime 7.1.6 release, Apache Atlas service did not validate the attributes of entity types which led to unpredictable situations, such as being unable to update attribute values or to assign valid values to those attributes in child types.

The checks or the validation of the attributes are performed for multiple use cases pertaining to

Entity and Classification type attribute names. Some of the common scenarios include:

1. Creating Entity or Classification type having attribute name which is same as one of the attributes from its Parent type
2. Updating the existing type to make it as a child of some other type, Atlas validates for the attribute names in both the types.
3. Updating the existing child type by adding an attribute, enables Atlas to validate if that attribute name is already present in its parent-type attributes.



Attention: These validations are also applicable for the Classification types.

Use case 1: If the user attempts to create an Entity or Classification type having an attribute name same as one of the attributes from its parent type, the Atlas service does not allow to create a child type and throws an exception.

As an example:

Consider there is a Type "ParentType" which has attributes "attributeA", "attributeB", and "attributeC". The user wants to create a child type of this "ParentType", say "ChildType". Note that the child type should not have any of the above attributes that are indicated. If any of these attributes are included at the time of "ChildType" creation, the Atlas service does not allow the creation of "ChildType".

Use case 2: If the user attempts to update the existing type and tries to change this type as a child to another type, Atlas service verifies the attribute names in both the types. If Atlas reports that the parent and child type have the same name attributes, Atlas does not allow updating the type and raises an exception.

As an example:

Consider there are 2 types "MyType01" having attribute "attributeA" and "MyType02" also having "attributeA". Atlas does not allow updating "MyType02" to become the child type of "MyType01" (and vice versa).

Use case 3: If a user attempts to update the existing child type and tries to add any attribute in child type, the Atlas service checks if that attribute name is already present in its parent-type attributes. If the attribute is present, Atlas service does not allow adding the attribute in child type and raises an exception.

As an example:

Suppose there are 2 types: "ParentType" having attribute "attributeA" and "ChildType" which is a child of "ParentType" and having attribute "attributeB". If a user tries to update the "ChildType" by adding an attribute "attributeA" in it, Atlas service does not allow adding the attribute in the "ChildType".

Validations for parent types

In Atlas, validations for scenarios where multiple parent types are involved for both entity types as well as Classifications.



Note: These validations are also applicable for the Classification types.

Use case 1: When there are two types having same name attributes, these two types cannot be assigned as parent-types to a child type. If a user tries to create such a child type, Atlas service does not allow it to create child type and raises an exception.

As an example:

Suppose two types "MyType01" and "MyType02" have a common attribute "attributeA".

A user tries to create a "ChildType" type by assigning both types "MyType01" and "MyType02" as its parents. Atlas service does not allow creation of "ChildType" because both parent types have the same name attribute.

Use case 2: A child type cannot be updated as a child of two parent types which have the same name attribute. If a user tries to update such a child type, Atlas service does not allow that and raises an exception.

As an example:

Suppose a type "ParentType01" has attribute "attributeA" and its child type is "ChildType". Also, another independent type "MyType02" which has the attribute "attributeA". The user wants to update the "ChildType" and tries to make it a child of "MyType02". Atlas service does not allow updating "ChildType", because both parent types have the same name attribute.

Case for implementing backward compatibility

When users of older versions want to upgrade their Apache Atlas versions to the latest one, they might face an issue because of certain validation checks in Atlas.

Existing types with attributes violating the attribute name rules (which were added in later Atlas versions) may have been added to the system.

All the existing types get populated at the time of Atlas startup and they all get validated.

If the users have Entity types or Classifications in the system that violate the naming rule, then Atlas raises an exception at the time of the service startup if the users have upgraded their environment to the latest version.

As a result of this exception that results at the time of Atlas startup, the typedef registry does not get populated. Even though users can log in to Atlas UI, they cannot view or use any types or Entities.

When users having data with such types try to upgrade their environment to the latest bits, they encounter issues as discussed earlier in this section.

Also, in their application.log they see an exception with a system output:

Error code: "ATLAS-400-00-09E" and Error message: "Invalid attribute name: <type name>.<attribute name>. Attribute already exists in another parent type: <parent type name>"

The backward compatibility feature is built to avoid these scenarios for such Entity types and Classifications.

How backward compatibility works

The backward compatibility feature is provided to skip exceptions coming from the attribute name rule.

From your Cloudera Manager instance > Configuration tab, a property must be added in your CM Safety Valve configuration.

atlas.skip.check.for.parent.child.attribute.name=true

Later, restart the Atlas service to apply this configuration.

When Atlas starts, no exceptions are raised and all the typedefs and Entities are displayed and available for use.

What Next:

After adding the backward compatibility configuration, all the validations and checks are disabled. You can create / update the types which have the same name attributes as their parent type attributes. The Atlas service operates in the same way as it was behaving prior to the Cloudera Runtime 7.1.7 version.



Note: When this feature is disabled, there will be corresponding Warning logs generated against those entity types, which may contain the information like this example.

Error code: "ATLAS-400-00-09E" and Error message: "Invalid attribute name: <type name>.<attribute name>. Attribute already exists in another parent type: <parent type name>"

If the user does not want any of their migrated existing types (which violate the name rule) to be present in the current environment, the user can delete these types using the Atlas REST APIs and later remove the configuration from the Cloudera Manager Safety Valve. Finally, restart the Atlas service.




Caution: Before deleting the typedefs, if there are any entities belonging to those types in the system, then the user must first delete and then purge those entities.

Associating classifications with entities


Assign classifications to entities in the entity detail page or in the search results.

There are two ways to associate a classification with an entity:

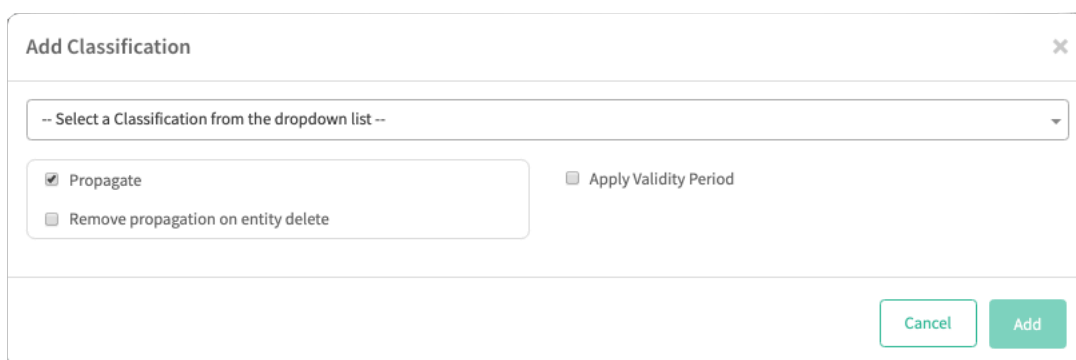
- Add the classification to the entity in the entity detail page.



- Add the classification to the entity in the search results.



When you associate a classification with an entity, you can set the classification's behavior with these controls:

A dialog box titled "Add Classification" with a close button (X) in the top right corner. It features a dropdown menu with the text "-- Select a Classification from the dropdown list --". Below the dropdown are two groups of checkboxes. The first group contains "Propagate" (checked) and "Remove propagation on entity delete" (unchecked). The second group contains "Apply Validity Period" (unchecked). At the bottom right are "Cancel" and "Add" buttons.

- **Propagate:** Any entities in the impact chain of this entity will automatically be tagged with this classification. The classification can't be removed from the "downstream" entities without being removed from this source entity. See more at [Propagate classifications to derived entities](#).
- **Apply validity period:** You can set a time period after which this classification is removed from the entity. Set the start time or end time or both. If you do not choose a time zone, the local time of the Atlas server is used.
- **Remove propagation on entity delete:** When an entity is deleted in the source and marked as deleted in Atlas, you can choose to maintain the propagation path, or to break propagation at the deleted entity.

You can change the behavior of classifications after they are assigned: go to the Classifications tab for an entity.

You can remove classifications from an entity using the X button on search results or entity detail page. Don't see an X on the classification name? That classification was assigned through propagation; to remove it, you have to turn off propagation for the source entity or term.

Related Information

[Propagating classifications through lineage](#)

Propagating classifications through lineage

Classifications on table columns can automatically apply to other tables where the same data is found when the data transformations are captured in lineage.

Viewing propagated classifications

When you look at the classifications assigned to an entity, you can see which classifications were originally assigned to this entity and which were propagated here through lineage. From the Classifications tab, you can:

- Show or hide propagated classifications (Show Propagated Classifications button)
- Navigate to the original entity where a classification was assigned (Propagated From button)

Navigate to the source of the classification

The screenshot shows the 'claims_view (hive_table)' interface. At the top, there's a 'Classifications:' section with a '+' button. Below it, a 'Term:' section shows 'Claim' with a dropdown and a '+' button. Underneath, 'Propagated Classifications:' lists 'DATA_QUALITY' and 'INSURANCE'. A tab bar at the bottom includes 'Properties', 'Lineage', 'Relationships', 'Classifications' (active), 'Audits', and 'Schema'. An orange arrow points to the 'Show Propagated Classifications' checkbox, which is checked. Below this, a table displays classification details:

Classification	Attributes	Action				
DATA_QUALITY	Propagated From	<table border="1"> <thead> <tr> <th>Name</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td>score</td> <td>0.51</td> </tr> </tbody> </table>	Name	Value	score	0.51
Name	Value					
score	0.51					
INSURANCE	Propagated From	NA				

Orange annotations include: 'Toggle the view of classifications associated with this entity by propagation' pointing to the checkbox, and 'Click to go to the entity that was first marked with this classification' pointing to the 'Propagated From' link for INSURANCE.

Setting propagation behavior

You can set propagation for a classification when you create the classification or in the lineage for an entity where the classification is assigned:

- Set propagation at the time when the classification is associated with an entity or term.

This screenshot shows the 'Add Classification' dialog box overlaid on the 'claims_view (hive_table)' interface. The dialog has a dropdown menu set to 'CONSUMER_BANKING'. It contains two checked checkboxes: 'Propagate' and 'Remove propagation on entity delete'. An 'Add' button is at the bottom right. Orange annotations include: 'Click to add a classification' pointing to the '+' button in the 'Classifications:' section, 'This classification is applied to entities in the lineage impact of this entity' pointing to the 'Propagate' checkbox, and 'When the current entity is deleted, this classification is removed from the entities where it propagated in the lineage impact.' pointing to the 'Remove propagation on entity delete' checkbox.

- Set propagation at a specific point in the lineage.

claim_savings (hive_table)

Classifications: DATA_QUALITY

Term: +

Properties Lineage Relationships Classifications Audits Schema

Current Entity → Lineage → Impact

create external t... claim_savings create view if no... claims_view

Click on a lineage impact line

Enable/Disable Propagation

Enable/Disable Propagation ☐ Select Classifications to Block Propagation

claim_savings → create view if not from cost_savings.claim_savings

hive_table → hive_process

none

Turn off propagation at this point in the lineage impact: no classifications assigned to the entity on the left will propagate to entities on the right.

Cancel Update

Enable/Disable Propagation

Enable/Disable Propagation ☒ Select Classifications to Block Propagation

Classification	Entity Name	Block Propagation
DATA_QUALITY	claim_savings (hive_table)	<input type="checkbox"/>

Choose the classifications to propagate to entities in the lineage impact.

Update

Related Information


[Associating classifications with entities](#)

[Controlling data access using tags](#)

Searching for entities using classifications

After you tag entities with classifications, it's very easy to search for the entities.

The Atlas dashboard makes it easy to find entities that you've tagged with classifications: in the left navigation pane, go to the Classification tab and select the classification you want. You can refine your search results by adding criteria on the Search tab (Basic search).



Related Information
[Using Basic Search](#)