



Collibra Data Intelligence Cloud

Platform Configuration

Collibra Data Intelligence Cloud - Platform Configuration

Release date: January 2, 2024

Revision date: January 02, 2024

You can find the most up-to-date technical documentation on our Documentation Center at

<https://productresources.collibra.com/docs/collibra/latest/#cshid=platform-configuration>

Contents

Contents	1
DGC service configuration	1
Configure Collibra email settings	4
Edit hyperlink settings	9
Configure the recommenders and matchers	11
Customizing the search index	14
Edit import configuration	51
Edit Excel export configuration	52
Edit CSV export configuration	53
Configure the logging of Collibra DGC API calls	54
Security configuration	57
Configure Collibra Connect	94
Edit global data source registration settings	98
Add a Jobserver to the DGC service	101
Configure data profiling behavior	103
Configure Cloud Data Classification Platform	107
Enable or disable Catalog experience	111
Enable the registration of a data source via Edge	112
Enable profiling for Edge	113
Enable data quality synchronization via Edge	116
Enable Tableau metadata API	117
Anonymize data via Jobserver	118
Enable or disable the Settings landing page	122

Configure the Search service	123
Search service configuration options	124
Jobserver service configuration	125
Jobserver authentication levels	126
Mutual authentication between Jobserver and DGC service	127
General specifications for certificates and private keys	129
Edit the Jobserver service settings	130
Connection from an on-premises Jobserver to a Colibra Data Intelligence Cloud	132
Connection from Colibra Data Intelligence Cloud to an on-premises Jobserver	148
Jobserver best practices	152
Service infrastructure configurations	156
Edit the DGC service infrastructure settings	157
Edit the Search service infrastructure settings	162
Managing environments	164
Create a Colibra environment	165
Environment statuses	166
Start an environment	167
Stop an environment	168
Delete an environment	168
Start a service	169
Stop a service	169
Reindex Colibra Data Intelligence Cloud	170
Restore to factory defaults	170
Monitoring your Colibra Cloud Self-Hosted environment	172
System metrics	173
Importing logs in an observability tool	174

Environment health check	176
Node management	180
Add a node to your infrastructure	181
View node details	182
Edit the node server settings	183
Remove a node	188
Back up and restore	190
Create a backup of Collibra Data Intelligence Cloud 2024.01	192
Backup options	194
The Backups page	196
Download a backup	197
Upload a backup	198
Delete a backup	199
Create a backup schedule	199
Backup schedule options	202
Backup schedules overview	203
Edit a backup schedule	205
Delete a backup schedule	205
Restore a backup	206
Customize Collibra DGC with a backup restore	209
Back up and download with the REST API	210
REST API - List of backups	220
REST API - Delete a backup	221
Diagnostic files	222
Diagnostic files	223
Create a diagnostic file	225

Download a diagnostic file	227
Delete a diagnostic file	227
Edit the environment log settings	228
Logging	229
Contents of a diagnostic file	231
Collibra Console settings	234
Collibra Console users	235
Edit the Collibra Console settings	250
Edit the Collibra Console server settings	272
Open a Collibra Console log file	277
Troubleshooting	278
Finding resource IDs	278
I don't see a workflow in the global create	281
DGC service configuration settings	i
General settings	i
Email configuration	iv
Hyperlinking configuration	viii
Recommender configuration	ix
Search index configuration	xi
Upload configuration	xvii
Statistics configuration	xviii
Import configuration	xix
Excel export configuration	xxi
CSV export configuration	xxi
User interface configuration	xxii
API call logging	xxiii

System metrics	xxiv
API configuration	xxv
Security configuration	xxv
Workflow engine configuration	xlvi
Collibra Connect	xlvi
Register data source	xlvi
Jobserver	li
Data profiling	lii
Beta features	lv
[Design system] features	lvii
Throttling	lvii
Hibernate cache configuration	lix
Graph query	lxiii
Table	lxiv
Purge configuration	lxvi
Cloud Data Classification configuration	lxvii
Reporting	lxix
Catalog Experience	lxxi
Data Notebook features	lxxi
Diagrams	lxxi
Everywhere Desktop configuration	lxxiii
Everywhere Mobile configuration	lxxvi
Collibra Browser Extension configuration	lxxvi
Edge	lxxvii
Tableau Metadata API	lxxvii
Job Service (Activities)	lxxviii

Identity	lxxviii
Lineage on Edge	lxxviii
Collibra Protect	lxxix
License configuration	lxxix
Data Marketplace configuration	lxxx
Data Privacy	lxxx
Export Configuration	lxxxi
Slack Integration Configuration	lxxxi
Asset external link	lxxxii
Appendix B - Spring Cron syntax	lxxxiii
Special characters	lxxxiv
Quartz Cron syntax	lxxxvii
Special characters	lxxxix

DGC service configuration

In Collibra Console, you can find all the Collibra Data Intelligence Cloud environment settings. After the installation of Collibra Data Intelligence Cloud, a default configuration is applied. If you have changed some settings, you can always go back to the original configuration by restoring the factory defaults. It is highly recommended to review the default settings and change them if necessary.

Disable or enable view permissions

Collibra Data Intelligence Cloud has a **View permissions** option that enables you to set permissions on objects. This option is enabled by default. When you disable the option, all objects are visible for every user.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **General settings** section, set the **Enable view rights** to **False** (disabled) or **True** (enabled).
3. Click **Save all**.

Edit the Help menu

In Collibra Data Intelligence Cloud, you have a Help menu to access the product documentation and various Collibra sites.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.

2. In the **Help menu** section, do one of the following:

Action	Description												
Add a menu item	a. At the bottom of the section, click Add. » A new link section appears. b. Enter the required information: <table> <tr> <th>Setting</th><th>Description</th></tr> <tr> <td>Links</td><td>The list of links in the help menu.</td></tr> <tr> <td>Menu item name</td><td>The name of the menu item as it will appear in Collibra Data Intelligence Cloud's help menu.</td></tr> <tr> <td>Menu index</td><td>The position of the menu item in the help menu. The top position starts with the value 1.</td></tr> <tr> <td>Menu URL</td><td>The target URL of the menu item.</td></tr> <tr> <td>Show admin only</td><td> ■ ✓ True: The menu item is only visible to users with the Sysadmin role. ■ ✗ False: The menu item is visible to every user. </td></tr> </table>	Setting	Description	Links	The list of links in the help menu.	Menu item name	The name of the menu item as it will appear in Collibra Data Intelligence Cloud's help menu.	Menu index	The position of the menu item in the help menu. The top position starts with the value 1.	Menu URL	The target URL of the menu item.	Show admin only	■ ✓ True: The menu item is only visible to users with the Sysadmin role. ■ ✗ False: The menu item is visible to every user.
Setting	Description												
Links	The list of links in the help menu.												
Menu item name	The name of the menu item as it will appear in Collibra Data Intelligence Cloud's help menu.												
Menu index	The position of the menu item in the help menu. The top position starts with the value 1.												
Menu URL	The target URL of the menu item.												
Show admin only	■ ✓ True: The menu item is only visible to users with the Sysadmin role. ■ ✗ False: The menu item is visible to every user.												
Edit an existing menu item	Make the necessary changes in the relevant link section.												
Remove a menu item	In the upper-right corner of the menu section, click  .												

- Add a new menu item:
- Edit an existing menu item by making the necessary changes in the relevant link section.
- Remove a menu item by clicking  in the upper-right corner of the menu section.

3. Click **Save all**.

4. Refresh the Collibra page to see the changes.

Configure Collibra email settings

You can use notification emails to notify users of any changes made to assets. In the email settings, you can:

- Configure your email server settings.
- Specify which roles will receive emails on which days of the week.
- Enable monthly summary emails and specify which roles will receive them.
- Configure how emails are handled.

Note

- Collibra Console and Collibra Data Intelligence Cloud use different email services. This topic describes the email configuration of Collibra Data Intelligence Cloud. The configuration of the Collibra Console email settings is described in [Configure Collibra Console email settings](#).
- This section applies only to on-premises Collibra environments. In cloud environments, you can only request to change the **From address**. Contact [Collibra Support](#) to change the From address, see also [Email configuration](#).

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.

2. In the **Email configuration** section, make the necessary changes.

Setting	Description
Default schedule (Requires restart)	The Cron schedule to send emails only at specific times. With this, you can send emails in batches and avoid an overload of mails. Keep in mind that these emails are only workflow emails and have nothing to do with the notification schedule. If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.
Template map	The location of template emails.
Password This setting requires the SUPER role.	The password paired with your username to sign in to your SMTP server.
From address	The email address used as the sender of all outgoing emails. Contact Collibra Support to change the From address, see also Email configuration.
Port This setting requires the SUPER role.	The port to connect to your SMTP server. The default value is 25 .
Host This setting requires the SUPER role.	The hostname or URL of your SMTP server.
Start TLS This setting requires the SUPER role.	✓ True: Use TLS (Transport Layer Security) to connect to your SMTP server. ✗ False (default): Do not use TLS to connect to your SMTP server.

Setting	Description
Username This setting requires the SUPER role.	The username to sign in to your SMTP server.
Sending threads This setting requires the SUPER role.	The number of threads that are used to send emails. The default value is 3 .
Max retries This setting requires the SUPER role.	The maximum number of retries before the system aborts the sending of an email. The default value is 5 .
Email address change notification This setting is only available for the ADMIN role.	If you change the email address to which notifications are sent, notification of the change is sent to the old email address.

3. In the **Notifications** section, make the necessary changes.

Setting	Description
Notification days	The days of the week on which Colibra sends notifications. The days are represented by numbers from 1 to 7, where 1 represents Sunday. Per row you can add one day.
Daily roles	The roles that receive notifications on the days defined in Notification days .
Enable monthly notifications	✓ True: The users receive a monthly summary. ✗ False (default): The users do not receive a monthly summary.

Setting	Description
Roles for monthly notifications	The roles that receive monthly notification emails. This is only relevant if Enable monthly notifications is ✓ True.

4. In the **Handlers** section, make the necessary changes.

Setting	Description
Host This setting requires the SUPER role.	The hostname or URL of the incoming mail server.
Port This setting requires the SUPER role.	The port to connect to your incoming mail server.
Protocol This setting requires the SUPER role.	1. The protocol to connect to your incoming mail server, with or without SSL (<i>POP3</i>, <i>POP3S</i>, <i>IMAP</i>, <i>IMAPS</i>). Note The additional S at the end of the abbreviations stands for the secure version of the protocol using SSL. Using this requires the SSL certificates to be correctly configured.
Force domain This setting requires the SUPER role.	✓ True: Only handle emails from the same domain as the handler's email address. ✗ False (default): Handle emails from any domain.
Handler list This setting requires the SUPER role.	The configuration of email handlers, which can poll emails on an email server, process those emails and perform actions based on the contents.

Setting	Description
Enabled This setting requires the SUPER role.	✓ True: The handler is enabled. ✗ False (default): The handler is not enabled.
Name This setting requires the SUPER role.	The name of the mail handler. We recommend to use a meaningful name to easily identify what this handler is used for.
Username This setting requires the SUPER role.	The username to connect to the incoming mail server.
Password This setting requires the SUPER role.	The password to connect to the incoming mail server.
Email address This setting requires the SUPER role.	The email address to which workflow action mails are sent.
Polling interval This setting requires the SUPER role.	The time in milliseconds between two pollings of the mail server.
Delete This setting requires the SUPER role.	✓ True: Delete messages from the mail server once the mail is processed. ✗ False (default): Keep messages on the mail server after the mail is processed. This option is only relevant if Protocol is <i>IMAP</i> or <i>IMAPS</i>.

Setting	Description
Alias filter	✓ True (default): Retrieve only the emails of which the To field contains the email address of the handler. ✗ False: Do not filter on the To field.
This setting requires the SUPER role.	

5. Click **Save all**.

What's next?

Restart the environment to apply your changes. For more information, go to [Stop an environment](#) and [Start an environment](#).

Edit hyperlink settings

In Collibra Data Intelligence Cloud, each asset is a potential target of an automatically created hyperlink. If a text attribute contains the name of another asset, it is automatically converted to a link to that asset.

Note The text editor on asset pages and domain pages includes an "Exclude from autohyperlinking" button, to disable autohyperlinking for text attributes.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.

- d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Hyperlinking configuration** section, make the necessary changes.

Setting	Description
Enable hyper-linking	✓ True: Hyperlinks are created automatically. ✗ False (default): Hyperlinks are not created automatically. For more information about automatic hyperlinks, see Hyperlinking. Warning If you enable this setting, the performance of Collibra can decrease.
Enable case sensitivity	✓ True: Hyperlinks are case-sensitive. ✗ False (default): Hyperlinks are not case-sensitive. Note If you edit this setting, you have to reindex Collibra.
Excluded asset type IDs	The list of asset types that are ignored by automatic hyperlinking. You can use this setting to exclude particular asset types so that not all the assets are potential targets for automatic hyperlinking. Each time you enter an asset type ID in a field, a new field appears for you to add another ID. To exclude multiple asset types, use separate fields, instead of entering multiple IDs in a single field separated by commas. Excluding assets reduces the amount of hyperlinks, which improves performance. Tip We recommend that you exclude technical asset types such as Column, Field, Table, Code Value, and Code Set. Note If you edit this setting, you have to reindex Collibra.

3. Click **Save all**.
4. [Reindex](#) Collibra Data Intelligence Cloud.

Configure the recommenders and matchers

Collibra Data Intelligence Cloud contains [recommenders](#) and [matchers](#) that recommend data sets or business assets.

You can configure them to optimize the recommendations.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Recommender configuration** section, make the necessary changes.

Setting	impacts	Description
Catalog recommender enabled	All recommendations	✓ True (default): The "Data sets you might like" section is included on the Data Catalog Home page. This section shows data sets you might be interested in, as determined by the recommender, which takes into account your data sets and the data sets of similar users. ✗ False: The "Data sets you might like" section is not included on the Data Catalog Home page.

Setting	impacts	Description
Data set recommender execution time	Recommendations of data sets to users	The schedule (CRON job) by which the data set recommender looks for recommended data sets for a user. By default the data set recommender does this every night.
Asset recommender execution time	Recommendations of business assets to data assets	The schedule (CRON job) by which the asset recommender looks for suggested relations between business assets and data sets.
Data set matcher execution time	Data set matcher	The schedule (CRON job) by which the data set matcher looks for similar data sets.
Data set similarity threshold	Data set matcher	The amount of business assets that have to be related to two data sets before the data sets are considered to be similar. This percentage is expressed by a decimal where 1,00 equals 100%. Example If this value is 0.3 and at least 30% of the related business assets are related to both data sets, they are considered to be similar.
Duplicate schema threshold	Schema matcher	The amount of assets that have to be related to both schemas before the schemas are considered to be similar. This percentage is expressed by a decimal where 1,00 equals 100%.

Setting	impacts	Description
Fuzzy vs exact matching strategy for business assets	Recommendations of business assets to data sets and of business assets to column assets	The percentage that determines to what extent assets with a similar name become more important. The ranking in the search engine results always has an impact on the suggestion score. However, similarity between the asset names can also be taken into account. If you decrease this percentage, the ranking of the search results becomes more important for the suggestion score, while the similarity between the asset names becomes less important. If you increase the percentage, assets with similar names will receive a higher suggestion score. This percentage is expressed by a decimal where 1,00 equals 100%. You can enter a value greater than 1,00.
Recommendation weights for data sets	Recommendations of data sets to users	An ordered comma-separated list of values that define the importance of properties for recommendations. The order of the values reflects the importance of the value. This setting is only used for data set recommendations if your Colibra does not yet have enough data for relevant results from the active recommendations algorithms. Possible values: ◦ <i>CERTIFIED</i>: Data sets that are certified are considered more relevant. ◦ <i>POPULARITY</i>: The number of visits to the data set page.

Setting	impacts	Description
Active recommendation algorithms	Recommendations of data sets to users and of business assets to data sets	A comma-separated list of algorithms that calculate recommendations. By default, all available algorithms are listed. Possible values: ◦ <i>BASELINE</i> ◦ <i>USER_MEAN</i> ◦ <i>IICF</i> (Item-Item Collaborative Filtering) ◦ <i>SLOPE_ONE</i> ◦ <i>WEIGHTED_SLOPE_ONE</i>
Data set elements threshold	Recommendations of data sets to users	The maximum number of elements per data set that the recommender will use to train the model. The data set elements are taken randomly. Lowering this number can prevent out-of-memory issues but also impacts the accuracy of recommendations for large data sets.

Warning If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.

3. Click **Save all**.

Note Depending on the configuration that you have applied, it is possible that you do not notice the recommendation updates immediately, but only the next day, for example when you update a schedule.

Customizing the search index

Before you customize the Collibra Data Intelligence Cloud search feature, it is important to learn how the search functionality works.

All text content in Collibra is stored in a search index to allow fast text search. To populate the search index, the text is split into separate words. The split is done by a component called the tokenizer and the words are often called tokens.

Every logical entity is stored in an index document. This document contains information about how many times a specific token occurs in the text. Separate index documents are stored for:

- Asset names
- Community names
- Domain names
- Text attributes
- Comments

When you search for text, the text is also tokenized in the same way. Then, the different words are searched for in the entire search index and a score is calculated for each of the matched documents. The calculation of this score is driven by different factors:

- The number of times the searched words occur in the document
- The size of the match relative to the size of the document

Tip You can influence this score by changing the boost factor. For more information about the search functionality in Collibra, see [Searching in Collibra DGC](#).

Configure the general search behavior

You can customize the search index configuration.

Requirements and permissions

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.

- d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Search Index configuration** section, modify the information as needed.

Setting	Description
UI search appends wildcard	✓ True (default): A wildcard (asterisk) is automatically added to each search query. An asterisk is not added if: ■ The query contains a tilde (~). ■ The query ends with a quotation mark ("). Note This applies only to queries via the user interface. A wildcard is not added automatically for REST API queries. ✗ False: A wildcard is not added to the search query.
Maximum batch size	The amount of resources that are scanned in a single operation for the search query. - Default value: 5,000 - Maximum value: 30,000
Maximum batch size for relations	The maximum batch size for relations reindex. - Default value: 500 - Minimum value: 50 - Maximum value: 10,000
Stop words (Requires restart)	A list of stop words that are ignored as tokens for the index. The default list of English stop words includes the following: a, an, and, are, as, at, be, but, by, for, if, in, into, is, it, no, not, of, on, or, such, that, the, their, then, there, these, they, this, to, was, will, with If you do not create your own list of stop words, the default list is used. If you create your own list of stop words, you need to: a. Reindex Colibra Data Intelligence Cloud. b. Restart the environment to apply your changes. For more information, go to Stop an environment and Start an environment.

Setting	Description
Relation-based search	✓ True (default in new environments): The Data Marketplace search considers certain assets and relation types between assets. As a result, your search results not only include assets that directly match the search criteria, but also assets that match the criteria through specific relation types. Example A column named Order is included in a data set named Customer. If the relation-based search is enabled and you search for Order in Data Marketplace, then the data set Customer appears in the search results because the data set contains this column. Tip For more information about this feature and the default relation types, go to Filtering and searching based on relations in Data Marketplace. ✗ False: The Data Marketplace search results do not consider relations. After you enable this setting, you must reindex Data Marketplace relations or reindex Collibra completely. Note In new Collibra environments, this setting is enabled by default. In upgraded Collibra environments, the previous status of this setting is retained.

3. Click **Save all**.

Stop words

Stop words are specified words that the search engine ignores in a search text. The list typically contains articles, prepositions, pronouns, and auxiliary verbs.

The default list of English stop words includes: a, an, and, are, as, at, be, but, by, for, if, in, into, is, it, no, not, of, on, or, such, that, the, their, then, there, these, they, this, to, was, will and with.

Note You can [create](#) your own list of stop words. If you choose not to, the default list applies.

Create a list of stop words

You can create a list of [stop words](#), which tells the search engine which words in your search text to ignore.

Note If you choose not to create a list of stop words, the default list applies.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Search index configuration** section, click in the **Stop words** field.
3. Type the first stop word, for example "a".
 - » Another field appears below, in which you can type the next stop word.
4. Enter all of your stop words.

Warning To work correctly, you must ensure that each stop word field contains only one word.

5. Click **Save all**.

6. Reindex Collibra Data Intelligence Cloud.

For more information about reindexing, see [Reindex Collibra Data Intelligence Cloud](#).

What's next?

- Restart the environment to apply your changes. For more information, go to [Stop an environment and Start an environment](#).

Delete words from your list of stop words

You can delete some or all of the words in your list of [stop words](#).

Note If you delete all of the words in your list of stop words, the default list applies.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.
- You previously [created](#) a list of stop words.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Search index configuration** section, click in the **Stop words** field.
3. Click  next to each of the stop words that you want to delete.
4. Click **Save all**.

5. Reindex Collibra Data Intelligence Cloud.

For more information about reindexing, see [Reindex Collibra Data Intelligence Cloud](#).

What's next?

- Restart the environment to apply your changes. For more information, go to [Stop an environment and Start an environment](#).

Edit the tokenizer settings

The tokenizer settings determine how the search text entered in the **Search** field is split into individual terms—called tokens—when a search is performed.

Important After you edit the tokenizer settings, you need to [reindex](#) Collibra Data Intelligence Cloud.

Requirements and permissions

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Search index configuration** section, in the **Tokenizer** subsection, enter the required information.

Setting	Description
Type	The tokenizer that determines how the search text is split. The Type field must contain either Standard (default) or Character. ◦ Standard: The Standard tokenizer is a method of splitting the search text into individual terms and is based on a default set of characters. This tokenizer follows the word break rules from the Unicode Text Segmentation algorithm. ◦ Character: The Character tokenizer allows you to customize when you want a search text to be split into individual terms. For more information, go to When and how to use the Character tokenizer.
Parameter map	A list of characters that the Character tokenizer allows. If you entered <i>Character</i> in the Type field, you must add a parameter map. To add a parameter map: - Click Add. » The Add map option dialog box appears. - In the Field key field, enter the following value: <i>allowedCharacters</i> - In the Field value field, enter the set of characters that you want the tokenizer to allow. For more information, go to When and how to use the Character tokenizer.

5.1.a Type

Character

5.1.b Parameter map

allowedCharacters: 

Add

- Click **Save all**.
 - » Your changes are saved.

What's next?

- Restart the environment to apply your changes. For more information, go to [Stop an environment and Start an environment](#).
- [Reindex Colibra Data Intelligence Cloud](#).

When and how to use the Character tokenizer

The Character tokenizer splits the search text into tokens (individual terms) when it encounters a character that is not in a defined set of allowed characters. This defined set contains the characters that you add to the parameter map when [editing](#) the tokenizer settings, in addition to the following set of characters: **0-9a-zA-Z**

When [editing](#) the tokenizer settings, consider using the Character tokenizer if you know which characters should be allowed during a search, and add those characters to the **allowedCharacters** parameter. The characters that you add, together with the set **0-9a-zA-Z**, form the defined set. For example, if you added **-** to the **allowedCharacters** parameter, then the defined set is as follows: **0-9a-zA-Z-**

Note You do not need to add **0-9a-zA-Z** to the **allowedCharacters** parameter because it is part of the default set of allowed characters for the character tokenizer.

5.1.a Type
Character

5.1.b Parameter map
allowedCharacters: -' 

Add

Example

Suppose that the **allowedCharacters** parameter contains an asterisk (*). If you search for **salesforce task under*review**, then the Character tokenizer will treat * as an allowed character and generate the following tokens:

- **salesforce**
- **task**
- **under*review**

The Character tokenizer does not split the search text when it encounters * because it is an allowed character.

Example

Suppose that an asset named **salesforce_task** exists in Collibra.

- **Behavior with the Standard tokenizer:** If the Standard tokenizer is used, when you search for **salesforce task** (without the underscore), **salesforce_task** (with the underscore) will not be shown in the search results. The Standard tokenizer allows **_** because it is one of the allowed characters in the Standard tokenizer pattern (**a-zA-Z0-9_**). As a result, **salesforce_task** is treated as **salesforcetask**.
- **Behavior with the Character tokenizer:** If you want **salesforce_task** to be shown in the search results, then use the Character tokenizer and ensure that **_** is not added to the **allowedCharacters** parameter. If **_** is not added to the **allowedCharacters** parameter, then the tokenizer will split the search text when it encounters **_**. As a result, **salesforce_task** is treated as **salesforce** and **task**.

Search boost settings

Search boost factors allow you to influence the order of the [search](#) results. You can edit the search boost factors to increase or decrease the importance of a resource type, field, or asset type in the search ranking. You can also [edit](#) the boost settings.

Note Boosting increases the base score of the search results. It does not set the order of the search results.

Values

The default boost factor for all resource types and fields is 1, with the exception of Name, which has a default value of 3.

The higher the value, the greater the [relevance score](#). Boost factors can also have decimals greater than 0.

Boosting object types

The search boost factor of an object type determines whether the objects of that type are shown higher or lower in the search results if the search text was found in any of its fields.

Boosting applies to the following object types:

- Asset
- Community
- Domain
- User
- User group

Search boost factors determine only the relative importance of these 5 object types. They do not apply to the individual resources of a given object type. For example, you can boost the Community object type, but you cannot boost a specific community or any of the subcommunities, domains, or assets in a specific community. Similarly, you can boost the User group object type, but you cannot boost a specific user group or users in a specific user group.

You can use the facets on the [Search page](#) to navigate to specific resources.

Example

Suppose that Asset has a search boost factor of 1 and Domain has a search boost factor of 3. If you search for a text that appears in any field of an asset and a domain, the domain is shown higher in the search ranking.

Boosting fields

The search boost factor of a field determines its importance for the search results, regardless of the resource type. The following fields are available for boosting:

- Name
- Comment
- Tag
- <any attribute type>
- <any asset type>

Example

Suppose that the Name field has a search boost factor of 2 and the Tag field has a search boost factor of 0.5. If you search for a text that appears in the Name field and the Tag field of different assets, the asset that contains your search text in the Name field is ranked higher than the asset that contains your search text in the Tag field.

Boosting asset types

The search boost factor by asset type allows you to assign greater importance to specified asset types in the search results.

Example

Suppose that the Business Process asset type has a search boost factor of 5 and the Risk asset type has a search boost factor of 1. If you search for a text that appears in any field of a Business Process asset and a Risk asset, the Business Process asset is shown higher in the search ranking.

When searching for an asset, common asset types usually rank higher in the search results. If, however, you are looking for a less common but important asset type, you can improve its visibility and frequency in the search results by applying an [asset boost](#).

Edit search boost factors

You can edit [search boost factors](#), to increase or decrease its importance of resource types, fields and asset types, in the search ranking.

By default, all resources and attribute types have a boost factor of 1. Values greater than 1 increase the importance, values less than 1 decrease the importance. All values must be greater than 0.

Note Search performance can be affected by the number of edited boost factors and the range of the values. For example, assigning a boost factor of 2 will likely result in better Search performance than assigning a boost factor of 200. We recommend values no greater than 10. Keep in mind that you can use decimals, for example 0.5 and 6.3.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Search index configuration** section, edit the search boost factors of the following resource types and fields.

Setting	Description
Asset	The search boost factor of assets.
Class Match	The search boost factor of data classes that were accepted during Automatic Data Classification. Note This is applicable only to users that have the Catalog global permission.
Community	The search boost factor of communities. Note This only affects the community asset itself, not the subcommunities, domains or assets in the community.
Domain	The search boost factor of domains. Note This only affects the domain asset itself, not the assets in the domain.
User	The search boost factor of users.
User group	The search boost factor of user groups.
Name	The search boost factor of the Name field.

Setting	Description
Comment	The search boost factor of the Comment field.
Tag	The search boost factor of the Tag field.

3. Optionally, edit the boost factors of the existing additional attribute types in the list.
4. Optionally, add a boost factor of additional attribute types.
 - a. Under **Attribute boost map**, click **Add**.
 - » The **Add map option** dialog box appears.
 - b. Enter the required information.

Setting	Description
Field key	The resource ID of the attribute type. Tip You can find the resource ID of an attribute type in the Collibra settings.
Field value	The search boost factor of the attribute type.

- c. Click **Add map option**.
5. Optionally, edit the boost factors of asset types.
 - a. Under **Asset boost map**, click **Add**.
 - b. Enter the required information.

Setting	Description
Field key	The resource ID of the asset type.
Field value	The search boost factor of the asset type.

- c. Click **Add map option**.

Note Asset type hierarchy is not considered when boosting. This means that editing the boost factor of a parent asset type does not change the boost factor of any child asset types. You have to individually edit the boost factor of each relevant asset type.

6. Optionally, enable or disable [exact matching](#):
 - ✓ True (default): If the name of an asset is exactly the same as the search text, put it at the top of the search results regardless of boost factors.
 - ✗ False: Use the regular search order, taking into account boost factors.
7. Optionally, enable or disable [partial exact matching](#):
 - ✓ True (default): For multi-word search text, the search engine considers the exact match percentage with the resource name, when ordering the results.

Example You enter search text "scheduled maintenance". Two example assets are ordered as follows:

 - a. An asset named "daily **scheduled maintenance**", as two of the three words (66%) match exactly.
 - b. An asset named "daily **scheduled maintenance** revised", as two of the four words (50%) match exactly.

 - ✗ False: The exact match percentage is not taken into account in the score calculation.
8. Click **Save all**.
9. [Reindex](#) Collibra Data Intelligence Cloud.

Tip For examples of how boost factors affect search results, see [Example search queries and analysis](#).

Delete an attribute boost map

You can delete an attribute boost map. As a result, the [search boost factor](#) of the attribute type is reset to 1.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Search index configuration** section, click **Boosting**.
3. In the **Attribute boost map** section, next to the boost map that you want to delete, click .
4. Click **Save all**.
5. Reindex Collibra Data Intelligence Cloud.

For more information about reindexing, see [Reindex Collibra Data Intelligence Cloud](#).

Exact match features

The exact match features help to ensure more intuitive ordering of search results.

You can [enable and disable](#) the exact match features in the boost settings. Both features are enabled, by default.

Exact Match Boost

This feature ensures that exact search text matches in the Name attribute of a resource always appear first in the search results, regardless of boost factors.

Example

Prerequisites

- The following assets exist in your Collibra environment:
 - A Schema asset with the name "Payment"
 - A Business Term asset with the name "Payment Type".

- A Data Attribute asset with the name "Payment Type".
- A Policy asset with the name "Payments".
- Other assets, as shown in the following image.
- Edit asset type boost factors, in Collibra Console, as follows:
 - Business Term: 3
 - Data Attribute: 2
 - Policy: 1.5


Search text

Enter search text "payment".

Results

The resources that exactly match the search text in the Name attribute of the resource appear first. In this example, there is only one exact match, the Schema asset with the name "Payment".

After the exact matches, the search results are sorted in order of descending relevance.



Partial Exact Match

This feature is only relevant for multi-word search text. In this case, the search engine considers the exact match percentage with the asset name, when ordering the results.

Example You enter search text "scheduled maintenance". Two example assets are ordered as follows:

1. An asset named "daily **scheduled maintenance**", as two of the three words (66%) match exactly.
2. An asset named "daily **scheduled maintenance** revised", as two of the four words (50%) match exactly.

Asynchronous indexing

When you make changes in your Colibra Data Intelligence Cloud environment—for example, you import millions of assets—the changes are logged in a PostgreSQL database table. The changes then need to be reflected in the [search index](#). Instead of processing these millions of changes all at once, the Search service processes them in batches or processing cycles. This process of updating a search index asynchronously in batches or processing cycles is called asynchronous indexing.

Asynchronous indexing helps ensure optimal performance of your environment while Colibra is indexing changes.

Note

- Although asynchronous indexing allows you to continue using your Colibra environment during a significant change to your environment, it may take some time to process all of the changes. For example, if you are importing millions of assets, while the changes are being processed, search results targeting the imported assets might be incomplete until processing is complete.
- You can edit the asynchronous indexing settings.

Benefits

- Fewer occasions where reindexing is necessary.
- Allows you to continue to use Collibra even after significant changes to your environment (for example, when importing millions of assets).
- Reduced memory demand when importing.

Limitations

- Importing performance could be slightly reduced.
- When processing changes, search results targeting the imported assets might be incomplete until processing is complete.

Edit asynchronous indexing settings

[Asynchronous indexing](#) helps ensure optimal performance of your environment while Collibra is indexing changes.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.

1. In the **Search index configuration** section, make the necessary changes.

Setting	Description
Processing frequency	The amount of time between batch processing cycles. Select a value between 10 and 60,000 milliseconds. The default value is 5,000 milliseconds.
Maximum batch size	The maximum number of asset changes processed, per processing cycle. Select a value between 100 and 30,000 changes. The default value is 5,000.

2. Click **Save all**.

What's next?

If you edited the value in the **Processing frequency** field, you need to restart the environment to apply your changes. For more information, go to [Stop an environment and Start an environment](#).

Relevance in search results

By default, search results are sorted in order of descending relevance.

What is relevance in the context of search results?

Relevance is a calculation of the similarity, measured across several lines of comparison, between your search text and the content of the resources in your Collibra environment.

In a set of search results, the relevance of each resource is represented by a positive number or score. The higher the score, the more relevant the resource is to your search text.

How are relevance scores derived?

To derive relevance scores, the Collibra search engine uses a combination of query clauses and boost factors.

Query clauses

When you perform a search, the Collibra search engine queries the database, using various query clauses. Each query clause compares the similarity between your search text and your Collibra resources along a different line of comparison.

Some examples of the objectives of different query clauses are as follows:

- Calculate the similarity between the spelling of your search term and the term found in a field in the database.
- Calculate how frequently your search term appears in a field. The more often it appears, the greater the relevance. A field containing 5 occurrences of a given term is more likely to be relevant than a field containing 1 occurrence of the term.
- Calculate the occurrence percentage of a term among all words in a particular field. For example, if your search term occurs twice in the 10-word description of an asset, that asset will have a higher relevance score than an asset for which your search term occurs twice in its 20-word description.

Boosting

Search boost factors allow you to influence the order of the [search](#) results. You can edit search boost factors to increase or decrease the importance of a resource type, field, or asset type in the search ranking.

For more information, go to [Search boost settings](#).

Examples

This topic explains how certain configuration settings can affect the search results.

Note The information in this topic is intended to:

- Help Collibra Data Intelligence Cloud administrators understand how the search configuration settings affect the search results.
- Help other Collibra users understand why certain search queries might not provide the expected results.

For more information, go to the respective sections on customizing the search index, boosting, stop words, and the exact match features.

Searching for assets that contain stop words in their names

Prerequisite

An asset named **On The Go** exists in your Collibra environment.

Search text


Enter the following search text: *On The Go*

Result

The asset **On The Go** is found.

Interpretation

In the following image, the word **Go** is shown in green. This indicates that **Go** is the match that produced the result. The words **On The**, which are shown in black, did not produce a match.



Furthermore, if you enter the search text *On The*, the asset **On The Go** or any other asset is not found. This is because "on" and "the" are **stop words**, which are filtered from indexing and searches.

Tip The best way to ensure thorough and intuitive search results is to name your assets, domains, and communities as thoughtfully as possible.

Searching for assets that contain more than one word in their names

Search text

Enter the following search text: *marketing team summit*

Interpretation

The search engine interprets your search text as follows: "marketing" or "team" or "summit"

Result

Assets with the following names are found:

Note The order of the actual results may differ from the following order.

- marketing team summit
- marketing_campaign_august
- team123
- summit_planning
- marketing_team_summit_august
- marketingTeamSummit
- marketing team summit August

Searching for assets that contain camel case in their names

Search text

Enter the following search text: *MarketingTeam*

Interpretation

The search engine interprets your search text as follows: "marketing" or "team"

Result

Assets with the following names are found:

- **MarketingTeam**
- **marketing_team_summit**
- **marketing team summit**

Searching for assets that contain underscores in their names

Search text

Enter the following search text: *marketing_team_summit*

Interpretation

The search engine interprets your search text as follows: "marketing" or "team" or "summit"

Result

Assets with the following names are found:

- **marketing_team_summit**
- **marketingteamsummit**
- **marketing_team_summit_August**
- **marketing team summit**

Searching for assets that contain special characters in their names

Search text

Enter the following search text: *team's calendar*

Interpretation

The search engine interprets your search text as follows: "team" or "calendar"

Result

Assets with the following names are found:

- **Team's calendar**
- **team123**
- **weekly calendar**
- **team_calendar_weekly**
- **marketing team calendar**

Searching for assets that contain URLs in their names

Search text

Enter the following search text: *xyz/unified/solutions/data*

Interpretation

The search engine interprets your search text as follows: "xyz" or "unified" or "solutions" or "data"

Result

Assets with the following names are found:

- **xyz/unified/solutions/data**
- **xyz/unified/solutions/data/reports**
- **unified solutions data**

- solutions data
- unified solutions

How boosting specified resource types affects search results

Prerequisites


- A resource, user, or user group with the name **verylongname** exists in your Collibra environment.
- Edit the resource type boost factors in Collibra Settings as follows:
 - Asset: 2
 - Community: 4
 - Domain: 6
 - User: 8
 - User group: 10

Search text

Enter the following search text: *verylongname*

Result

The search results are ordered according to the boost factor values of the respective resource types. The user group resource type, with a boost factor of 10, is the most relevant of the results. Asset, with a boost factor of 2, is the least relevant resource type.



How boosting specified fields affects search results

Prerequisites


- The following assets exist in your Collibra environment:
 - An asset named **superfeature**
 - An asset named **asset1**, with the following tag: **superfeature**
 - An asset named **asset2**, with the following comment: **superfeature**
- Edit the property boost factors in Collibra Settings as follows:
 - Name: 1
 - Comment: 5
 - Tag: 10

Search text

Enter the following search text: *superfeature*

Result

The search results are ordered according to the boost factor values of the respective fields. The tag field, with a boost factor of 10, is the most relevant of the results. Name, with a boost factor of 1, is the least relevant field.



How boosting attributes affects search results

Prerequisites

- The following assets exist in your Collibra environment:
 - An asset with the description **terminator**
 - An asset with the definition **terminator**
 - An asset with the note **definition**
- Edit the attribute boost factors in Collibra Settings as follows:
 - Description: 1
 - Definition: 2
 - Note: 3

Search text

Enter the following search text: *terminator*

Result

The search results are ordered according to the boost factor values of the respective attributes. The note attribute, with a boost factor of 3, is the most relevant of the results. Description, with a boost factor of 1, is the least relevant attribute.

terminator
Results: 3 Sort by Relevance Search In All fields
Data Governance Council ▶ Data Quality Dimensions BA differentAsset Note terminator
Data Governance Council ▶ Data Quality Dimensions BA asset123 Definition terminator
Data Governance Council ▶ Data Quality Dimensions BA assetAbc Descriptive Example terminator

How boosting asset types affects search results

Prerequisites


- The following assets exist in your Collibra environment:
 - A Business Term asset with the name **Payment**
 - A Data Attribute asset with the name **Payment Type**
 - A Policy asset with the name **Payments**
- Edit the attribute boost factors in Collibra Settings as follows:
 - Business Term: 3
 - Data Attribute: 2
 - Policy: 1.5

Search text

Enter the following search text: *payment*

Result

The search results are ordered according to the boost factor values of the respective asset types. The Business Term asset type, with a boost factor of 3, is the most relevant of the results. Policy, with a boost factor of 1.5, is the least relevant asset type.



How the exact match boost feature affects search results, regardless of boost factors

Prerequisites

- The following assets exist in your Collibra environment:
 - A Schema asset with the name **Payment**
 - A Business Term asset with the name **Payment Type**
 - A Data Attribute asset with the name **Payment Type**
 - A Policy asset with the name **Payments**
 - Other assets, as shown in the following image

- Edit the asset type boost factors in Collibra Settings as follows:
 - Business Term: 3
 - Data Attribute: 2
 - Policy: 1.5

Search text

Enter the following search text: *payment*

Results

The resources that match the exact search text in the Name attribute of the resource are shown first. In this example, there is only one exact match—the Schema asset with the name **Payment**.

After the exact matches, the search results are sorted in order of descending relevance.

payment Search

Results: 15 | Sort by: **Relevance** | Search in: **All fields**

Business Analysts Community > Schemas > default	Payment
Business Analysts Community > Financial Services > Trading Glossary	BT Payment Type Definition the type of payment
Business Analysts Community > Supply Chain Performance Management > SCPM Glossary	BT Price Definition The amount of money expected, required, or given in payment for a certain product
Business Analysts Community > Schemas > dwh-stock_exchange-refined	payment_type Original Name dwh-stock_exchange-refined > dwh-stock_exchange-refined > stock_exchange_refined > payment_type Description type of the payment Full Name dwh-stock_exchange-refined > dwh-stock_exchange-refined > stock_exchange_refined > payment_type
Business Analysts Community > Financial Services > Trading Logical Data Catalog	DATT Payment Type

How term frequency contributes to relevance scores

Prerequisites


- Several assets, each with a variation of **jedi** in the name, exist in your Collibra environment.
- Default boost factors for all resource types, properties, and attributes.

Search text

Enter the following search text: *jedi*

Result

The search results are ranked as shown in the following image. A combination of the total number of matches in a name and the percentage of matches per total words in a name affect relevance scoring of the results.



Recommended file types for uploads

In general, you can allow any file type to be uploaded to Collibra Data Intelligence Cloud. However, you can limit the types of files that users can upload.

Below is a list with the minimum of allowed file types that you need for Collibra to work properly:

File type	MIME type
.xls	application/vnd.ms-excel
.xlsx	application/vnd.openxmlformats-officedocument.spreadsheetml.sheet
.xml	application/xml
.zip	application/zip Note Only administrators can work with backups .
.java, .log, .txt, ...	text/plain Note For a list of all available file types, see http://www.sitepoint.com/web-foundations/mime-types-complete-list .
.csv	text/csv
.csv	text/plain
.bpmn	application/xml Note Only administrators can work with workflows .

Below is a list with an enumeration of additional recommended allowed file types:

File type	MIME type
.doc	application/msword
.docx	application/vnd.openxmlformats-officedocument.wordprocessingml.document
.dotx	application/vnd.openxmlformats-officedocument.wordprocessingml.template

File type	MIME type
.xltx	application/vnd.openxmlformats-officedocument.spreadsheetml.template
.ppt	application/vnd.ms-powerpoint
.pptx	application/vnd.openxmlformats-officedocument.presentationml.presentation
.potx	application/vnd.openxmlformats-officedocument.presentationml.template
.ppsx	application/vnd.openxmlformats-officedocument.presentationml.slideshow
.pdf	application/pdf
.html	text/html

Below is a list with an enumeration of recommended file types for avatars:

File type	MIME type
.jpeg, .jpg	image/jpeg
.png	image/png
.gif	image/gif

Edit the buffer settings of the Colibra DGC statistics

You can edit the buffer settings of the Colibra Data Intelligence Cloud

Prerequisites

- You have the **ADMIN** or **SUPER** role in Colibra Console.
- You have the **SUPER** role in Colibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Statistics configuration** section, make the necessary changes.

Setting	Description														
Buffer size	The maximum amount of statistics entries that the buffer can contain before saving them in the database. The default value is <i>10</i>.														
Buffer flush time	The maximum amount of time in milliseconds to keep statistic entries in memory before saving them in the database. The default values is <i>10,000</i>.														
Cron map	List of statistics, listed by their Cron name, and a Cron interval. These are the default values: <table> <tr> <th>Field key</th><th>Field value</th></tr> <tr> <td>workflow-task</td><td>0 59 23 * * ?</td></tr> <tr> <td>active-users</td><td>0 0/15 * * * ?</td></tr> <tr> <td>term-count</td><td>0 59 23 * * ?</td></tr> <tr> <td>vocabulary-count</td><td>0 59 23 * * ?</td></tr> <tr> <td>page-hit</td><td>0 0 * * * ?</td></tr> <tr> <td>task-count</td><td>0 0 * * * ?</td></tr> </table> If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.	Field key	Field value	workflow-task	0 59 23 * * ?	active-users	0 0/15 * * * ?	term-count	0 59 23 * * ?	vocabulary-count	0 59 23 * * ?	page-hit	0 0 * * * ?	task-count	0 0 * * * ?
Field key	Field value														
workflow-task	0 59 23 * * ?														
active-users	0 0/15 * * * ?														
term-count	0 59 23 * * ?														
vocabulary-count	0 59 23 * * ?														
page-hit	0 0 * * * ?														
task-count	0 0 * * * ?														

3. Click **Save all**.

Edit the schedule of Collibra DGC statistics

Collibra Data Intelligence Cloud collects statistics at regular intervals. The statistics are shown on various pages, such as the **Metrics** pages of the Collibra applications.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.

2. In the **Statistics configuration** section, make the necessary changes.

Setting	Description														
Buffer size	The maximum amount of statistics entries that the buffer can contain before saving them in the database. The default value is <i>10</i>.														
Buffer flush time	The maximum amount of time in milliseconds to keep statistic entries in memory before saving them in the database. The default values is <i>10,000</i>.														
Cron map	List of statistics, listed by their Cron name, and a Cron interval. These are the default values: <table> <tr> <th>Field key</th><th>Field value</th></tr> <tr> <td>workflow-task</td><td>0 59 23 * * ?</td></tr> <tr> <td>active-users</td><td>0 0/15 * * * ?</td></tr> <tr> <td>term-count</td><td>0 59 23 * * ?</td></tr> <tr> <td>vocabulary-count</td><td>0 59 23 * * ?</td></tr> <tr> <td>page-hit</td><td>0 0 * * * ?</td></tr> <tr> <td>task-count</td><td>0 0 * * * ?</td></tr> </table> If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.	Field key	Field value	workflow-task	0 59 23 * * ?	active-users	0 0/15 * * * ?	term-count	0 59 23 * * ?	vocabulary-count	0 59 23 * * ?	page-hit	0 0 * * * ?	task-count	0 0 * * * ?
Field key	Field value														
workflow-task	0 59 23 * * ?														
active-users	0 0/15 * * * ?														
term-count	0 59 23 * * ?														
vocabulary-count	0 59 23 * * ?														
page-hit	0 0 * * * ?														
task-count	0 0 * * * ?														

3. Click **Save all**.

What's next?

Restart the environment to apply your changes. For more information, go to [Stop an environment](#) and [Start an environment](#).

Edit import configuration

In Collibra Data Intelligence Cloud, you can import data from comma-separated value (CSV, Excel) files.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.

2. In the **Imports configuration** section, make the necessary changes.

Setting	Description
Enable work-flows during import	✓ True: Allow starting workflows upon importing assets. ✗ False (default): Do not allow to start workflows upon importing assets.
Asset responsibilities support	✓ True: Enable importing responsibilities at asset level. ✗ False (default): Disable importing responsibilities at asset level. Warning Setting specific responsibilities on a large number of resources will affect the performance and stability of the system.
Number of failed commands before stopping import job	An import job with the option to continue on error enabled will stop after the specified number of commands have failed. Any valid command is still committed to the database until the moment the job stops, which can lead to some resources being imported. The default and maximum value is <i>100</i>.
Temporary data location	The location of the temporary files used by the import job. The default value is <i>FILE</i>.
Import UI v2 This setting requires the SUPER role.	✓ True (default): Use the new import interface for importing assets and complex relations, with improved usability and performance. ✗ False: Use the original import interface for importing assets and complex relations.

3. Click **Save all**.

Edit Excel export configuration

In Collibra Data Intelligence Cloud, you can export Collibra data to CSV or Excel files.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Excel export configuration** section, make the necessary changes.

Setting	Description
The default CSV separator character	The default separator character of the CSV fields for complex relations.
The default CSV quote character	The default quote character of the CSV fields for complex relations.
Number of rows per chunk of data	When exporting views, the database is called repeatedly, each time fetching a chunk of data to build the export file. This option defines how many rows each chunk of data can contain. Lower values reduce the burden on memory. Higher values require more memory, but may slightly increase the speed of the export. The default value is 5,000.

3. Click **Save all**.

Edit CSV export configuration

In Collibra Data Intelligence Cloud, you can export Collibra data to CSV or Excel files.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **CSV export configuration** section, make the necessary changes.

Setting	Description
Always use quotes	◦ ✓ True: Use quotes for every cell in the CSV. ◦ ✗ False (default): Only use quotes when necessary.
Number of rows per chunk of data	When exporting views, the database is called repeatedly, each time fetching a chunk of data to build the export file. This option defines how many rows each chunk of data can contain. Lower values reduce the burden on memory. Higher values require more memory, but may slightly increase the speed of the export. The default value is 5,000.

3. Click **Save all**.

Configure the logging of Collibra DGC API calls

You can use Collibra Data Intelligence Cloud in a web interface, but also through REST API calls. Every API call is logged in a database. You can configure the methods that you want to log and the maximum number of entries to store.

Steps

1. Open the DGC service settings for editing:
2. In the **API call logging**, make the necessary changes.

Setting	Description
Enabled	✓ True: API call logging is enabled. ✗ False (default): API call logging is disabled.
Maximum number of log entries (Requires restart)	The maximum number of API calls to store in the database. Once this number is reached, the oldest records are overwritten. The default value is 1,000,000.
Pattern duration list	The list of methods and a corresponding minimum duration time. The minimum duration time is the minimum time before the method is stored in the database.
Minimum duration	The time in milliseconds that an API call must last before it is logged.
Method pattern	The method that you want to log in the database. For each pattern that you want to log, you have to add a new pattern.

3. Click **Save all**.

What's next?

Restart the environment to apply your changes. For more information, go to [Stop an environment](#) and [Start an environment](#).

Create a pattern duration for API call logs

You can create a pattern duration for API call logs.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **API call logging** section under **Pattern duration list**, click **Add**.
3. In the new record, fill in the required fields.

Setting	Description
Enabled	◦ ☒ True: API call logging is enabled. ◦ ☐ False (default): API call logging is disabled.
Maximum number of log entries (Requires restart)	The maximum number of API calls to store in the database. Once this number is reached, the oldest records are overwritten. The default value is 1,000,000.
Pattern duration list	The list of methods and a corresponding minimum duration time. The minimum duration time is the minimum time before the method is stored in the database.
Minimum duration	The time in milliseconds that an API call must last before it is logged.
Method pattern	The method that you want to log in the database. For each pattern that you want to log, you have to add a new pattern.

4. Click **Save all**.

What's next?

Restart the environment to apply your changes. For more information, go to [Stop an environment](#) and [Start an environment](#).

Security configuration

You can configure the security of Collibra Data Intelligence Cloud in various aspects, going from SSL settings to password settings.

Configuring LDAP

You can integrate Collibra Data Intelligence Cloud with your organization's LDAP servers. User data as well as group-related information can be synchronized and used by Collibra. When LDAP is configured, authentication (credential checking) is done directly on the LDAP server(s).

In order for Collibra to synchronize with the LDAP server, the LDAP fields have to be mapped to the corresponding Collibra fields and the LDAP server has to be configured correctly.

Configure LDAP

To configure one or more LDAP servers, follow these steps:

1. Open the DGC service settings for editing:
2. Under **Servers** of the **LDAP** section, click **Add**.
3. Fill in the required fields.
For more information about the fields, go to the **LDAP** section of the [DGC service configuration page](#).
For a complete [tutorial](#), visit Collibra Support Portal.
4. Click **Save all**.
5. [Reindex](#) the environment.

What's next?

Restart the environment to apply your changes. For more information, go to [Stop an environment and Start an environment](#).

How to enable logging for troubleshooting LDAP issues

Before you start to troubleshoot issues with LDAP, make sure that you have [configured LDAP as documented here](#).

If you still encounter issues, enable logging for further investigation.

Warning If you have investigated the LDAP issue, don't forget to revert all the changes from this section.

Steps

1. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the tab pane, click the **Data Governance Center** service of the environment whose log settings you need.
3. Click **Logs**.
4. Above the table, to the right, click **⚙ Settings**.
5. Click **Add logger**.
6. In the **Add logger** dialog box, enter the necessary information.
 - **Logger name:** Enter *com.collibra.dgc.core.service.ldap.impl.LdapServiceImpl*
 - **Logger level:** Select *DEBUG*.
7. Click **Add logger**.

Define password strength

Every Collibra Data Intelligence Cloud user has a password to sign in to Collibra. You can define the password strength so that your system can resist guessing and brute-force attacks. You can set the password complexity, as well as limit the number of sign-in attempts or set an expiry date for the password.

Steps

1. Open the DGC service settings for editing:
2. In the **Security configuration** settings, click **Password**.
3. Set the necessary password policy parameters.
For more information about the parameters, go to the **Password** section of the [DGC service configuration page](#), that also contains the default values for each of the password parameters.
4. Click **Save all**.

What's next?

Restart the environment to apply your changes. For more information, go to [Stop an environment and Start an environment](#).

Configure REST security

To change the REST security settings, follow these steps:

1. Open the DGC service settings for editing:
2. In the **REST** section, make the necessary changes.

Setting	Description
Limited CSRF	This option offers limited security, so we recommend upgrading to the Enhanced CSRF. ◦ ☒ True: The validity of a request is checked with a CSRF token. ◦ ☐ False (default): The validity of a request is not checked with a CSRF token.
Enhanced CSRF	If enabled, Collibra will check the validity of the request using a Spring Security CSRF token. ◦ ☒ True: The validity of a request is checked with a CSRF token. ◦ ☐ False (default): The validity of a request is not checked with a CSRF token.

Setting	Description
Referrer enabled	✓ True: The HTTP referrer header is used to identify the origin of the request. ✗ False (default): The HTTP referrer header is not used to identify the origin of the request. It is recommended to leave this option disabled.
Referrer checking allow empty	✓ True (default): The HTTP referrer header can be empty. ✗ False: The HTTP referrer header cannot be empty.

3. Click **Save all**.

What's next?

Restart the environment to apply your changes. For more information, go to [Stop an environment and Start an environment](#).

Configure the SSL settings

For secure communication from Collibra Data Intelligence Cloud to other services such as an LDAP server, you can activate SSL.

Tip To configure SSL to securely access Collibra from your web browser, see [Configure SSL to access Collibra DGC](#).

Steps

1. Open the DGC service settings for editing:
2. In the **SSL** section, make the necessary changes.

Setting	Description
Key store name	The name of the keystore file. The file is expected to be in the <collibra_data>/dgc/security folder.
Key store password	The password of the keystore.

Setting	Description
Key store type	The type of the keystore file. For example, <i>JKS</i> or <i>PKCS12</i> .
Trust store name	The name of the truststore file. The file is expected to be in the <collibra_data>/dgc/security folder.
Trust store password	The password of the truststore.
Trust store type	The type of the truststore file. For example, <i>JKS</i> or <i>PKCS12</i> .

3. Click **Save all**.

What's next?

Restart the environment to apply your changes. For more information, go to [Stop an environment](#) and [Start an environment](#).

Configure SSL to access Collibra DGC

If you want to connect to Collibra Data Governance Center in a secure way with your web browser, you have to use SSL. This procedure explains how you can activate SSL access to Collibra DGC.

Tip For secure communication from Collibra DGC to other services, for example an LDAP server, see [Configure the SSL settings](#).

Prerequisites

- You have knowledge of the JSON syntax.
- You have created a Java KeyStore according the procedure described by [Oracle](#), for example **clientkeystore**.
- You have noted the following data while creating the Java KeyStore:
 - KeyStore file name: *clientkeystore* in the Oracle example.
 - KeyStore alias: *client* in the Oracle example.

- KeyStore password: The password that you entered after executing the command of the first step in the Oracle example.
- KeyStore alias password: The password that you entered as last step of step 2 in the Oracle example.
- You have stored the Java KeyStore on the server that hosts the DGC service in the `<collibra_data>/dgc/security` folder, for example `/opt/collibra_data/dgc/security`.

Steps

To configure Collibra Data Governance Center for access over SSL, follow these steps:

1. Open a terminal session on the DGC node.
2. Open the file `<collibra_data>/dgc/config/server.json` for editing.
3. Fill in the following parameters in the **httpsConnector** section:

Add string values between double quotes.

Parameter	Description
port	The port on which the HTTPS connector must bind. The value must be higher than 1024 to avoid root permissions. Note If you want to use the default SSL port 443, you have to use a reverse proxy.
keyAlias	The KeyStore alias.
keyPass	The KeyStore alias password.
keystorePass	The KeyStore password.
keystoreFile	The full path to the KeyStore file name, for example <code>/opt/collibra_data/dgc/security/clientkeystore</code> .

Parameter	Description
Example:	
<pre>"httpsConnector" : { "port": 5404, "keyAlias": "your-alias", "keyPass": "your-password", "keystorePass": "your-password", "keystoreFile": "/opt/collibra_data/dgc/security/collibradgc.jks"}, }</pre>	

4. Save and close the file.
5. Open the DGC service settings for editing:
6. Click the **General settings** section.
7. Update the **Base URL** parameter with *https* and the new port.
8. Restart the environment.

Connect to your Collibra DGC environment via the Base URL.

Extra

To prevent regular HTTP traffic to Collibra DGC, update the **address** parameter with the value *127.0.0.1* in `<collibra_data>/dgc/config/server.json` and restart the environment.

This will not prevent the administration tools, for example Collibra Console, from connecting to Collibra DGC without SSL.

For more information, see the knowledge base on the [Collibra Support Portal](#).

Single Sign-On (SSO)

Single Sign-On (SSO) enables users to access Collibra Data Intelligence Cloud using a web client, without having to explicitly type their login credentials (username and password).

Collibra provides support for two types of SSO. Each SSO type can be used with or without LDAP (Light-weight Directory Access Protocol), resulting in the following SSO modes:

- [SAML 2.0](#), options:
 - [SSO SAML with attributes sync](#)
 - [SSO SAML LDAP](#)
- [SSO through header information](#), options:
 - [SSO Header with Collibra Data Intelligence Cloud users](#)
 - [SSO Header with LDAP sync](#)

Tip If you want to use a custom certificate in the SSO configuration for Collibra access, see [this section](#).

This section explains how to:

- [Choose the most appropriate mode of SSO](#).
- Configure the Collibra for the SSO mode you prefer.

Note On Collibra SAML support:


- The only supported SAML protocol version in Collibra is SAML 2.0 (urn:oasis:names:tc:SAML:2.0:protocol).
- When configuring SAML, in order to obtain the SP metadata from Collibra DGC, go to <https://<yourinstance>/rest/latest/saml/metadata>.
- If you want a full SP metadata, go to <https://<yourinstance>/rest/latest/saml/metadata?complete=true>.
- Collibra only supports assertions to come in through the HTTP-POST binding (as defined in the SP metadata file).
- For more information about this subject, see the knowledge base on the [Collibra Support Portal](#).

Deciding which SSO mode to use

The following flow chart can help you decide which SSO mode to use. If you already know which mode is appropriate for your environment, just skip to the correct section in order to learn how to configure the relevant SSO mode.

Warning In case you are considering the SSO Header configuration, please make sure that you are aware of its limitations and security requirements. Please involve your security team while configuring SSO Header.

Note Collibra Data Intelligence Cloud Cloud only supports SAML with Attribute sync.



SSO header

SSO header with Collibra Data Intelligence Cloud users

Since there is no standard for the behavior of the SSO Header modes, Collibra makes certain assumptions.

Once configured with the appropriate header information (see configuration step), Collibra assumes that:

- You are running a proxy responsible for authenticating the users connecting to Collibra.
- This proxy fills in the value of the header to the username of the user that has been authenticated by the proxy. This is the case in every single request, including for resources like javascript, CSS and image files.
- Collibra uses this username to start a session in the application, assuming the above steps were performed properly.
- Collibra does not perform extra authentication.

To configure SSO header with Collibra users, consult the [SSO header: configuration options](#) section.

SSO Header with LDAP synchronization

Much like the SSO Header with Collibra users, there are no standards for proxy behavior.

Collibra assumes that:

- You are running an LDAP server and Collibra has been configured to sync with that LDAP. For more information about LDAP synchronization, see [Configuring LDAP](#).
- You are running a proxy server that is responsible for authenticating the users who want to go to Collibra DGC.
- The proxy server fills in the value of the header to contain either the LDAP Distinguished Name (DN) or the identifying LDAP attribute of the user (as configured, see configuration step).
- Collibra:
 - Uses the DN or identifying attribute to fetch the username from the LDAP server.
 - Verifies that the user is allowed to use Collibra.
 - Uses this user to start a session.

To configure SSO header with LDAP synchronization, consult the [SSO Header LDAP: configuration options](#) section.

Single Sign-On FAQ

Search this FAQ section to see if you find an answer to your question. If you cannot find the right answer, please contact our support team at support@collibra.com.

Help! SSO is broken! I can't get into Collibra!

Even though you seem to be locked out of Collibra Data Intelligence Cloud, you can always gain access by using a Collibra local account. You can change the address of the page that you are visiting to end with /signin (So: <https://<your instance>/signin>)

A generic username/password form is displayed. You can only use accounts that have been created in Collibra itself.

Once you have regained access to your instance, you can try to figure out what exactly went wrong. To resolve any problems, you can start by retrieving the log files, and analyzing them. If something went wrong with the SAML authentication flow, some error will show up in these log files.

For more information about logging, see [Logging](#).

I'm getting the "SAML authentication failed. Please contact your Administrator for more information." error, but nothing wrong shows up in the log

Most likely, this is because you are trying to sign in without global application permissions.

In most cases, you will want to use the **Everyone** group to assign application permissions to everybody. That is including new users and guest users. To do this, sign in manually, with the sign-in page, using an admin account. Then navigate to **Settings > Roles**, and assign the **Everyone** group to the appropriate global roles. For more details, see [Collibra Console roles](#) and [Responsibilities](#).

In the cloud environment, when do we require a restart of the instance if we changed something?

A restart of the environment can be requested when:

- Your initial configuration has been done, and the SAML metadata file has been provisioned by [uploading](#) a SAML metadata file.

In order to help in the decision process, the following things are being reloaded automatically:

- The SAML metadata file is reloaded every 2 hours. The next reload point will be mentioned in the log file.
- All **configuration.xml** changes will be reloaded on every SAML Response (meaning: Collibra always uses the current setting in the **configuration.xml**).

You can request a restart of your cloud instance by contacting our support team at support@collibra.com.

How do I upload my SAML metadata file to a cloud instance?

If your environment is using a cloud instance of Collibra, the only supported mode of SSO is SAML with attribute sync.

To upload the SAML metadata file to your environment, see [Upload a SAML metadata file](#).

How can I use SSO with SAML to assign groups/roles automatically to users?

In the [SAML documentation](#), you can read how to set the attributes that Collibra expects. One of these attributes is to allow the IDP to send group information.

If you set the **Group** attribute to *user.group*, Collibra then parses the SAML Response to look for the *user.group* attribute. This attribute has to be a comma-separated list of groups to which this user has to be added.

Groups that do not exist, are created just in time. Users that are no longer part of a group per this definition are removed from that group.

Any groups that are empty after this operation are also removed.

Why is my Collibra user account removed from all groups when I log in with SSO?

In the context of group management, even if the Group attribute is not set, Collibra still assumes that you are trying to manage the groups from your IDP.

This question implies that you wish to use group management in Collibra itself. Set the option **Groups DGC managed** in **Security configuration > SSO** to *True*. Collibra then always assumes that group management is performed in Collibra itself.

Why is a new profile created each time I log in to Collibra?

Profiles are created when the username in the assertion is not yet found in the database.

You will notice that there is no <username> attribute linked. This is because the value Collibra uses to create and check usernames is actually the value of the **nameid** field in the SAML Response.

This field is governed by SAML and can ensure that the value is not Personally Identifiable Information (PII). The default nameid format used by Collibra is:

```
urn:oasis:names:tc:SAML:2.0:nameid-format:persistent.
```

The problem with profiles being regenerated on every sign-on always has to do with this nameid. In some cases, the nameid is overridden by the IDP to be a random value, causing a user to be created many times. At other times, the IDP does not properly support our default nameid-format.

In order to fix this problem, first talk to your IDP team to figure out if the IDP is indeed supporting `urn:oasis:names:tc:SAML:2.0:nameid-format:persistent`. If it cannot support this in any case, Collibra can override the nameid-format used in the SAML Request.

Warning Overriding the `nameid-format` is an advanced feature. Choosing to use it, makes the consequences your own responsibility.

If you still want to override the nameid-format, add the following tag in the <saml> section of the configuration.xml:

```
<force-nameid>Your fully qualified nameid-format</force-nameid>
```

Collibra expects you to provide the full string that describes the nameid that you want to use.

Configure Single Sign-On access

After you have decided which [SSO mode](#) you are going to use, you can configure Single Sign-On (SSO) for Collibra Data Intelligence Cloud.

Note In your IDP configuration itself, make sure that the URL to your Collibra environment has a trailing slash, for example *https://dgc.yourcompany.com/*, without the trailing slash, SSO will not work in your environment.

Steps

1. Open the DGC service settings for editing:
2. In the **Security configuration** section, click **SSO**.
3. Enter the SSO parameters according to the selected SSO mode:
 - [SAML_ATTRIBUTES](#), for a complete tutorial, see the knowledge base on [Collibra Support Portal](#).
 - [SAML_LDAP](#), for a complete tutorial, see the knowledge base on [Collibra Support Portal](#).
 - [SSO_HEADER](#)
 - [SSO_HEADER_LDAP](#)
4. Click **Save all**.

Note

On Collibra's Support Portal, you can find complete tutorials for the following SSO providers with SAML attributes:

- [Okta](#)
- [OneLogin](#)
- [PingOne Cloud](#)
- [PingFederate](#)
- [Azure](#)
- [Active Directory Federation Service \(ADFS\)](#)
- [ADFS with SAML LDAP mode](#)

What's next?

Restart the environment to apply your changes. For more information, go to [Stop an environment](#) and [Start an environment](#).

SSO SAML with attributes: configuration options

To configure SSO as SAML with attributes sync, fill in the SSO fields as follows:

SSO configuration parameter	Value
Mode	<i>SAML_ATTRIBUTES</i>
Header	<leave empty>
DN	<leave empty>
Attribute	<leave empty>

SSO configuration parameter	Value
Metadata HTTP	URL to saml.xml file (http://url.to.your/saml.xml).
Entity ID	The entity ID as defined in the metadata file. It defines which specific entity (IDP or SP) should be used in a metadata file. The SAML metadata file enables you to define multiple entities in one metadata file. This can also prove useful in combination with Collibra, in cases where planned upgrades are going to occur. You can then upload a new metadata file that contains both entities. When the time comes to switch, you only need to change the configuration option for the Entity ID.
Groups DC managed	• False: groups are managed by the SAML IDP. • True: groups are managed by DGC.
Service Provider Entity ID	Leave empty, unless the Base URL in General settings does not match the Service Provider Entity ID to be used.
Sign authentication requests	Set to True to use the SAML keypair to sign authentication requests. Note A SAML keypair in x509 is generated and stored in the SAML metadata file when Collibra is started for the first time.
Force authn	• True (default): Make the SAML request ask for authentication every single time, even when the user is already known. • False: Do not make the SAML request ask for authentication.
Force passive	• <i>False</i> (default): The IDP is allowed to take visible control of the user interface/authentication. • True: The IDP is forced to not take visible control. Only for specific setups, see the SAML 2.0 specifications for more information.

SSO configuration parameter	Value
Name ID	The nameID to be sent in the SAML Request. nameID has to have the following format: <pre>urn:oasis:names:tc:SAML:2.0:nameid-format:persistent</pre> For other options: See the SAML 2.0 specifications. IDP has to understand the nameID in the SAML Request. It is recommended to set this to what the IDP expects.
Name ID allow create	• True (default): Allow the IDP to create a new nameID to satisfy the SP SAML Request. • False: Do not allow the IDP to create a new nameID.
Disable	• False (default): Send the authentication context as configured in this section. • True: Disable the authentication context. Nothing in this section applies anymore. (see also Configuring requested authentication context)
Comparison type	Defines the authentication strength that is to be used by the IDP compared to the SAML requested authentication context. This is advanced configuration, see the SAML 2.0 specifications for more information. Possible values: • <i>minimum</i> (default) • <i>maximum</i> • <i>better</i> • <i>exact</i> See also Configuring requested authentication context.
Reference list	Contains a list of allowed references. Default: <pre>urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport</pre> This is advanced configuration, see the SAML 2.0 specifications for more information.
Declaration list	Similar to the Reference list, but is empty by default. This is advanced configuration, see the SAML 2.0 specifications for more information.

SSO configuration parameter	Value
Ignore endpoint scheme differences	• False: The scheme is always taken into account for validating the SAML response. Any difference there invalidates the SAML response. • True: The scheme is not taken into account for validating the SAML response.

SSO SAML LDAP: configuration options

To configure SSO as SAML with LDAP user provisioning, fill in the SSO fields as follows:

SSO configuration parameter	Value
Mode	<i>SAML_LDAP</i>
Header	<leave empty>
DN	If the SSO mode is SSO_HEADER_LDAP or SAML_LDAP, this field determines whether the distinguished name (DN) or attribute is used: • True: the header has to contain the distinguished name (DN) of the user in the LDAP. • False: the header has to contain the value of Attribute. If the SSO mode is DISABLED, SSO_HEADER or SAML_ATTRIBUTES, this field is ignored.
Attribute	Set to the unique identifier, usually uid of the LDAP directory when linking to LDAP through an LDAP attribute. The nameID contains the value of the attribute set here to look for in the LDAP service. For example, if the value equals <i>sAMAccountName</i>, the SAML response should contain the value for this attribute of the user being signed in.

SSO configuration parameter	Value
Metadata HTTP	Enter the URL to the SAML metadata file. For example: http://url.to.your/metadata.xml. Note This parameter is ignored if you uploaded the SAML metadata file.
Entity ID	The entity ID as defined in the metadata file. It defines which specific entity (IDP or SP) has to be used in a metadata file. The SAML metadata file enables you to define multiple entities in one metadata file. This can also prove useful in combination with Collibra, in cases where planned upgrades are going to occur. You can then upload a new metadata file that contains both entities. When the time comes to switch, you only need to change the configuration option for the Entity ID.
Groups DC managed	• False: groups are managed by the SAML IDP • True: groups are managed by DGC
Service Provider Entity ID	Leave empty, unless the Base URL in General settings does not match the Service Provider Entity ID to be used.
Sign authentication requests	Set to True to use the SAML keypair to sign authentication requests. Note A SAML keypair in x509 is generated and stored in the SAML metadata file when Collibra is started for the first time.
Force authn	• True (default): Make the SAML request ask for authentication every single time, even when the user is already known. • False: Do not make the SAML request ask for authentication.
Force passive	• False (default): The IDP is allowed to take visible control of the user interface/authentication. • True: The IDP is forced to not take visible control. Only for specific setups, see the SAML 2.0 specifications for more information.

SSO configuration parameter	Value
Name ID	The nameID to be sent in the SAML Request. nameID has to have the following format: <pre>urn:oasis:names:tc:SAML:2.0:nameid-format:persistent</pre> For other options: See the SAML 2.0 specifications. IDP has to understand the nameID in the SAML Request. It is recommended to set this to what the IDP expects.
Name ID allow create	• True (default): Allow the IDP to create a new nameID to satisfy the SP SAML Request. • False: Do not allow the IDP to create a new nameID.
Disable	• False (default): Send the authentication context as configured in this section. • True: Disable the authentication context. Nothing in this section applies anymore. See also Configuring requested authentication context.
Comparison type	Defines the authentication strength that is to be used by the IDP compared to the SAML requested authentication context. This is advanced configuration, see the SAML 2.0 specifications for more information. Possible values: • <i>minimum</i> (default) • <i>maximum</i> • <i>better</i> • <i>exact</i> See also Configuring requested authentication context.
Reference list	Contains a list of allowed references. Default: <pre>urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport</pre> This is advanced configuration, see the SAML 2.0 specifications for more information.
Declaration list	Similar to the Reference list, but is empty by default. This is advanced configuration, see the SAML 2.0 specifications for more information.

SSO configuration parameter	Value
Ignore endpoint scheme differences	• False: The scheme is always taken into account for validating the SAML response. Any difference there invalidates the SAML response. • True: The scheme is not taken into account for validating the SAML response.

When both **Attribute** and **DN** are defined, **DN** takes priority and the attribute-based configuration is ignored.

SSO header: configuration options

To configure SSO in SSO header mode, fill in the SSO fields as follows:

SSO configuration parameter	Value
Mode	<i>SSO_HEADER</i>
Header	The header to be monitored for the username to be signed in.
DN	<does not apply>
Attribute	<does not apply>

SSO Header LDAP: configuration options

To configure SSO in SSO header mode with LDAP user provisioning, fill in the SSO fields as follows:

SSO configuration parameter	Value
Mode	<i>SSO_HEADER_LDAP</i>

SSO configuration parameter	Value
Header	The header to be monitored for the username to be signed in.
DN	• True: The nameID in the SAML response refers to the DN (Distinguished Name) of the user to be signed in. • False: Fill in the attribute field. Set to true when linking to LDAP through the Distinguished Name.
Attribute	Set to the unique identifier, usually uid of the LDAP directory when linking to LDAP through an LDAP attribute. The nameID contains the value of the attribute set here to look for in the LDAP service. For example, if the value equals <i>sAMAccountName</i> , the SAML response should contain the value for this attribute of the user being signed in.

Note When both **Attribute** and **DN** are defined, **DN** takes priority and the attribute-based configuration is ignored.

Configuring requested authentication context

You can configure in which context requested authentication happens on IDP side by changing the following settings:

- **Disable**
- **Comparison type**
- **Reference list**
- **Declaration list**

If you leave the above settings as they are and Collibra Data Intelligence Cloud is configured for SAML, Collibra sends a requested authentication context with:

- The **exact** comparison
- The **urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport** class-ref
- No class-decl

Enable SAML response encryption

To increase the security of the communication between Collibra Data Intelligence Cloud and an SSO provider, Collibra supports [encrypted SAML responses](#).

Prerequisites

- You can access the Collibra REST API.

Steps

Enable response decryption

- Open the DGC service settings for editing:
- In the **Security configuration** section, click **SSO**.
- Set the option **Response decryption mode** to *OPTIONAL* or *FORCED*.

An encryption key pair is generated and added to the SAML keystore. A self-signed encryption certificate is generated and works in most situations.




- Click **Save all**.

Provide encryption certificate to Identity Provider

The encryption certificate of Collibra has to be imported in the Identity Provider. You can retrieve this encryption certificate via the Collibra REST API.

- Retrieve the certificate via the endpoint at `/rest/2.0/security/saml` and copy the base64 representation of the encryption certificate.



2. Copy the content of the **ds:X509Certificate** element into a PEM file. A PEM file is a plain text file with the extension **pem**.

```
-----BEGIN CERTIFICATE-----  
MIICrDCCAZSgAwIBAgI...  
-----END CERTIFICATE-----
```

3. Provide this PEM file to an administrator of the Identity Provider who can load it into the IdP.

What's next?

Your Collibra Data Intelligence Cloud environment is configured to accept encrypted SAML responses.

If your Identity Provider does not accept self-signed certificates, contact Collibra Support.

Encrypted SAML response

Collibra Data Intelligence Cloud supports encrypted SAML responses. Collibra implements [XMLEnc](#), the industry standard to encrypt SAML responses.

Encryption

SAML assertions are expected to be encrypted in XMLEnc mixed mode, which means that:

1. The Identity Provider (IdP) generates a random symmetric key and uses it to encrypt the assertion.
2. The symmetric key is encrypted using a public key provided by the Service Provider (SP), in this case Collibra.
3. The encrypted symmetric key is embedded into the SAML response alongside both the public key used in its encryption and encrypted assertion.
4. When Collibra receives the response, it decrypts the symmetric key using its own private key and then uses that symmetric key to decrypt the assertion.

Supported cryptographic algorithms

Different IdPs may use different cryptographic schemes and algorithms. As a Service Provider, Collibra has to decrypt the assertions of many IdPs, so Collibra supports the algorithms that are marked as "required" in the XMLEnc specifications.

The supported algorithms for symmetric cryptography are:

- 3DES
- AES-128-CBC
- AES-256-CBC
- AES-128-GCM
- AES-256-GCM

The supported algorithm for asymmetric cryptography is RSA-OAEP, including MGF1 with SHA1.

Collibra recommends to integrate with an IdP that uses AES-256-GCM for symmetric encryption.

Configure a custom certificate for SSO in Collibra

If you configure single sign-on for accessing Collibra Data Intelligence Cloud, a default certificate is used. You can use this certificate for signing SAML authn requests or for SAML assertion encryption/decryption.

Instead of the default certificate, you can use your own certificate. However, keep in mind that you can only configure SSO with your own certificate via a REST API call.

Prerequisites

- The certificate must meet the following requirements:
 - The certificate must be in PEM format.
 - The PEM file must be unencrypted (no password).
 - The PEM file must contain the server certificate the private key of that certificate.

Tip To convert a key to a PEM key: `openssl rsa -in <pem-key>.key -out <rsa-key>.pem`

Example PEM file:

```
-----BEGIN RSA PRIVATE KEY-----
MIIEvgIBADA ... bml6YXRpb252YWxza
....
z3P668YfhUbKdRF6S42Cg6zn
-----END RSA PRIVATE KEY-----

# Your certificate
-----BEGIN CERTIFICATE-----
MIIFaDCCBFC ... bml6YXRpb252YWxza
...

1ffygD5IymCSuuDim4qB/9bh7oi37heJ4ObpBIzroPU0thbG4gv/5blW3Dc
=
-----END CERTIFICATE-----
```

- A base64 encoding hash of an API user.

Tip You can calculate the base64 hash of the user as follows: `echo '<username>:<password>' | base64`, for example `echo 'apiUser:apiUserpassword' | base64` results in `YXBpVXNlcjphcGlVc2VycGFzc3dvcmQK`

Use certificate for SAML assertion encryption

```
curl --location --request POST \
  'https://<your_dgc_environment_url>/rest/2.0/se-
  curity/saml/certificate/ENCRYPTION' \
  --header 'Authorization: Basic <base-64 encoding hash>' \
  --form 'file=@"/path/to/pem-file"'
```

Use certificate for signing SAML authn requests

```
curl --location --request POST \
  'https://<your_dgc_environment_url>/rest/2.0/se-
  curity/saml/certificate/SIGNING'\
  --header 'Authorization: Basic <base-64 encoding hash>' \
  --form 'file=@"/path/to/pem-file"'
```



Working with SAML

SAML or Security Assertion Markup Language is an open standard that allows security credentials to be shared by multiple computers across a network. It describes a framework that allows one computer to perform some security functions on behalf of one or more other computers.

For more information, see the website of the [OASIS Security Services \(SAML\) Technical Committee](#).

Terminology

Term	Explanation
SP	The Service Provider of the SAML integration. In this case, the SP is Collibra Data Intelligence Cloud.
IDP	The Identity Provider of the SAML integration. In this case , the IDP depends on your company. Examples of IDPs: Microsoft's Active Directory (if the federation services are installed), PingFederate, SiteMinder, Okta, OneSign...

Term	Explanation
SP-initiated sign-on	A type of sign-on that starts from the SP (Collibra). In this flow, the SP asks the IDP to authenticate the current user, using a SAML request. The IDP then returns a SAML response. This SAML response contains the assertion of the user currently visiting the SP. That assertion contains the properly configured attributes of the user that is requested. Note If the assertion's SubjectConfirmationData element contains the Address attribute, ensure that the value is a URL. This value indicates the expected network address of the client. You can enable or disable the validation of the address attribute in Console.  <pre> sequenceDiagram participant SP as Service Provider (Collibra DGC) participant Client participant IDP as Identity Provider Client->>SP: 1. User requests access SP->>IDP: 2. Generate SAML authentication request IDP->>Client: 3. Post to the IDP's SSO URL Client->>IDP: 4. Parse and verify SAML authentication request IDP->>Client: 5. Authenticate Client->>IDP: 6. Generate SAML assertion (authentication response) IDP->>SP: 7. Post to the SP's SSO URL SP->>Client: 8. Parse and verify SAML assertion Client->>SP: 9. User is logged on </pre>
IDP-initiated sign-on	A type of sign-on that starts from the IDP portal towards the SP. This flow skips the SAML request sent in the SP-initiated flow, but still sends the SAML Response containing an assertion containing the properly configured attributes of the user that is requested.  <pre> sequenceDiagram participant SP as Service Provider (Collibra DGC) participant Client participant IDP as Identity Provider Client->>IDP: 1. User requests access IDP->>Client: 2. Authenticate IDP->>Client: 3. Generate SAML assertion Client->>SP: 4. Post to the SP's SSO URL SP->>Client: 5. Parse and verify SAML assertion Client->>SP: 6. User is logged on </pre>

Supported SAML 2.0 profile

Although SAML 2.0 is a clear standard, it is too large to support in its entirety. The designers of the SAML standard also realized this and that's why they developed SAML Profiles. These profiles are defined in the [SAML 2.0 specification's section](#).

The purpose of SAML profiles is to allow both IDP and SP to decide to only support a subset in the SAML protocol that is applicable to them, in their situation.

Collibra supports the Web Browser SSO Profile.

How does Collibra implement SAML?

Even though the web browser profile is properly defined in the standard's profile section, there are a lot of optional features, but not all of them are implemented. This section explains what is implemented and what not.

The profile defines that metadata may be used to build trust between SP and IDP. Collibra has chosen to enforce this, so the only way to integrate your IDP with Collibra is for the IDP to provide the metadata to Collibra. This can be done by either uploading a SAML metadata file, or by providing a public URL that contains the metadata.

The profile explains that signing and/or encryption of SAML requests and responses can be used and that Collibra expects the IDP to perform the signing and/or encryption. Since the communication between SP and IDP is always supposed to happen in the context of an HTTPS connection, encryption is not performed. The transport layer takes care of that. By default, SAML requests are not signed, but you can enable signed requests. SAML responses always have to be signed because the response has to come from the correct IDP.

The validUntil property in the SAML metadata

The SAML metadata may contain a `validUntil` attribute. This optional attribute indicates the expiration time of the metadata contained in the element, as well as its child elements.

The "valid until" field is taken into account by Collibra. SAML will stop working either when the validity period has expired or when the content in the metadata file is no longer valid.

Steps for integrating SAML in Collibra

Step	Task	Description
1	Bring in the IDP's SAML metadata in Collibra Console	Collibra needs the metadata of your IDP. You can bring in the metadata in two ways: • Upload the IDP's SAML metadata file in Collibra Console. This is the easiest approach, but the metadata will not be updated automatically if there are changes in the IDP. • Enter the URL to the IDP's metadata in the Metadata HTTP field of the SAML configuration of the SSO section of the environment settings. This requires the IDP to have a publicly accessible URL containing the SAML metadata. If the IDP's metadata changes, Collibra automatically updates the metadata as well.
2	Configure DGC	Configure the environment settings for SSO SAML with attributes or SSO SAML with LDAP. We strongly recommend that you enable Sign authentication requests. If you do not enable Sign authentication requests, Collibra's outgoing requests are not signed. If your IDP demands that all requests are signed, the unsigned requests from Collibra will be ignored.
3	Restart Collibra	Stop and start the affected environment. This will generate a key pair, which is required to sign authentication requests.

Step	Task	Description
4	Bring in Collibra's SAML metadata in your IDP	Your IDP needs the metadata of Collibra. Your IDP team should help with this. You can obtain the metadata in two ways: • Sign in and download the metadata file from Collibra, and import it in your IDP. You can download the metadata file here: <code>https://<your dgc instance>/rest/1.0/saml/metadata</code> • Enter the URL to the Collibra's metadata in the IDP: <code>https://<your dgc instance>/rest/1.0/saml/metadata</code>. Your IDP needs to be connected to the internet to access that URL.
5	Configure your IDP	You IDP may require further configuration. Your IDP team should help with or perform the complete configuration.

Upload a SAML metadata file

If you want to use [SAML](#) for user authentication, you need to provide your IDP's SAML metadata file.

If there is already a SAML metadata file, uploading another file will overwrite the existing one.

Steps

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. Click the name of an environment to show its details.
3. Click the **SAML** tab.
4. Click **Upload**.
 - » The **Upload SAML metadata** file dialog box appears.
5. Perform one of the following steps:
 - Drag and drop a valid SAML metadata file in the **Upload a file...** field.
 - Click in the **Upload a file...** field, select the SAML metadata file and click **Open**.

What's next

Configure the environment settings for [SSO SAML with attributes](#) or [SSO SAML with LDAP](#).

Download a SAML metadata file

You can download the IDP's SAML metadata file. For example, this can be used when you want to create a [backup](#).

Prerequisites

- You have [uploaded](#) a SAML metadata file.

Steps

1. Open Collibra Console with a user profile that has the **SUPER** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. Click the name of an environment to show its details.
3. Click the **SAML** tab.
4. Click .
- » The SAML metadata file is downloaded.

Delete a SAML metadata file

You can delete a SAML metadata file from your configuration.

This is mandatory when you want to start using an online SAML metadata provider in the **Metadata HTTP** field of the SAML configuration. If you do not delete the SAML metadata file, the URL is ignored.

Prerequisites

- You have [uploaded](#) a SAML metadata file.

Steps

1. Open Collibra Console with a user profile that has the **SUPER** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. Click the name of an environment to show its details.
3. Click the **SAML** tab.
4. Click .
- » The **Delete SAML metadata** file dialog box appears.
5. Click **Delete**.

What's next

Enter the URL of the IDP's metadata in the **Metadata HTTP** field of the SAML configuration of the [SSO section](#) of the environment settings.

JSON Web Token authentication

You can use JSON Web Token (JWT) authentication for your applications that interact with the Collibra REST API. During this process, your application requests an access token from your Identity Provider (IdP). The IdP acts as the authentication server and returns a signed JWT access token. When your application makes API calls to Collibra, it provides the JWT access token as a **Bearer** token in the HTTP **Authorization** header.

Example

```
curl -H 'Authorization: Bearer <your_token>' ...
```

Benefits of using JWT token authentication include:

- Keeping the authentication details at the IdP and separate from Collibra.
- Simplifying your security management.
- Limiting the time a token is valid for.

Set up and use JWT

The following table summarizes acquiring and using JWT in Collibra REST API requests.

Process	Steps
Initial setup	1. Create a client credential account with a secret in your IdP. 2. Determine the JSON Web Key Set (JWKS) endpoint URL for your IdP. 3. Register the JWKS endpoint with Collibra. 4. Create a user in Collibra for your client application account.  Tip Provide a meaningful first and last name to identify that this is a service account.
When your application starts	1. Authenticate your client application with your IdP. 2. Save the returned access token for use in REST API calls.
When your application calls the Collibra REST APIs	1. Include the JWT token in the authorization HTTP header as a bearer token. 2. If the API call responds with unauthorized, the access token or JWKS credentials may have expired. Re-authenticate and retry the request.

JWT dependencies on your Identity Provider

The following details must be provided by your IdP in the JWT token for Collibra to accept the token.

Section	Field	Description
Header	alg	The encryption algorithm. Collibra supports all standard encryption algorithms.
	kid	The JWKS identifier of the public key used to sign the JWT token. An identity provider may have multiple public key certificates, for example multiple valid keys may exist while a key rotation is in progress.
	typ	Optional field that defines the JWT type. Tip Use the MIME type format, for example <i>application/jwt</i> or <i>application/secevent+jwt</i> with the <i>application/</i> prefix removed, for example <i>jwt</i> or <i>secevent+jwt</i>.
Payload	iss	The issuer of the token. This field is used to check the token comes from the expected IdP.
	exp	The token expiry time.
	iat	The time the token was issued.
	sub	The subject or the principal ID.
Signature		The digital signature of the header and payload. The signature verifies the message did not change along the way. In the case of tokens signed with a private key, the signature also verifies the authenticity of the JWT sender.

Configure JWT settings

To enable or change the JSON Web Token configuration:

1. Open the DGC service settings for editing:
2. In the **JWT** section, make the necessary changes.

Setting	Description
JSON Web Key Set URL	The URL to retrieve public key information needed to verify the authenticity of JSON Web Tokens (JWTs), issued by an authorization server. This setting is required to enable JWT authentication.
JWT Token Types	A case-insensitive comma-separated list of accepted JWT media types coming in the typ header parameter. Leave blank if the authorization server does not provide a media type parameter. The default values is at+jwt,jwt.
JWT Algorithms	A comma-separated list of accepted JWT algorithms coming in the alg header parameter. See https://tools.ietf.org/html/rfc7518#section-3.1 for details. Leave blank to accept all digital signature algorithms.
JWT Issuer	The accepted issuer coming in the iss JWT claim. Leave blank if the authorization server does not provide an issuer claim.
JWT Audience	A comma-separated list of accepted audience values for the aud claim. The value for this field is a configuration setting in your authorization server, which identifies your Colibra environment as the intended recipient of the JWT. Leave blank if the authorization server does not provide an audience claim.
JWT Principal ID Claim Name.	The name of the JWT claim containing the principal's identity. See https://tools.ietf.org/html/rfc7519#section-4.1.2 for details. Defaults to the standard subject claim, sub. Change this setting only if your authorization server has other means of identifying the principal, for example, a client_id claim. This setting is required if JWT authentication is enabled.

Setting	Description
JWT Maximum Clock Skew	The maximum acceptable difference in seconds between the clocks of the machines running the authorization server and Colibra. Differences smaller than the given amount are ignored when performing time comparisons for token validation. The default value is 60 seconds if left blank.

3. Click **Save all**.
4. Restart the environment to apply your changes. For more information, go to [Stop an environment and Start an environment](#).

JWT troubleshooting

The following table contains JWT authentication error codes that may appear in the body of the **401 Unauthorized** HTTP response. You can use the error codes to determine the appropriate course of action for your application.

Error code	Description	Possible action
malformedToken	The JWT token is incorrectly encoded or has an incorrect syntax.	Check your IdP configuration to make sure it is returning a token in JWT format.
expiredToken	The JWT token has expired.	In your application, request a new token from the IdP and retry.
invalidToken	The JWT token is invalid. Common causes are: • Bad signature. • Wrong audience. • Wrong issuer. • Wrong media type. • Disallowed signing algorithm.	Check your IdP and Colibra JWT configuration to make sure the settings are consistent.
unableToProcessToken	The JWT token could not be processed.	Check your IdP and Colibra JWT configuration and if the problem persists contact Colibra Support .

Configure Collibra Connect

We have made the decision to transition away from Collibra Connect to provide customers a wider range of integration options.

Our native Collibra integrations (connectors) will be easier to implement and maintain, provide a better return on investment, and allow you to grow with and derive greater value from Collibra:

- Collibra integrations and Spring Boot based frameworks will replace Collibra Connect as options to build integrations going forward.
- You can choose any ESB or integration method for your use case.
- Our intention is to enable Collibra connectors to support ingestion as well as use cases for data profiling, data classification and other cloud functionalities.
- If you have an enterprise MuleSoft license, you can easily switch to it.

Rest assured Connect templates are and will remain compatible with our product, please contact us for any Connect-related question. Only support or any upgrades on these products will be discontinued.

Note As of September 2022, you will need a MuleSoft Community Edition license or your own proprietary paid license to run Connect templates.

With Collibra Connect, you can connect to Collibra Data Intelligence Cloud with a tool of your own choice. Collibra Connect acts as the gateway between your tool and Collibra. For more details on Collibra Connect, consult the [Collibra Connect user guide](#).

In this section, you can learn how to set the credentials to access Collibra Connect with your own tool.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. Go to the **Data Classification** section.

3. Enter the required information:

Setting	Description
Machine Learning platform URL This setting requires the SUPER role.	The address of the machine learning platform that will classify your data.
Requester Name This setting requires the SUPER role.	The unique name to identify the client when using Machine Learning platform.
API key This setting requires the SUPER role.	The API Key to authorize the requester when connecting to the Machine Learning platform.
Enable Data Classification	✓ True: Enable Collibra's data classification technology. ✗ False (default): Do not use Collibra's data classification technology are not accepted.

4. If needed, configure the automatic classification acceptance and rejection.

Setting	Description
Enable automatic classification acceptance and rejection	✓ True: The automatic acceptance and rejection of data classification suggestions is active. ✗ False (default): Data classification suggestions are not automatically accepted or rejected. Tip Start using the Automatic Data Classification tool by manually accepting and rejecting the data classification suggestions. Only activate the automatic acceptance and rejection feature if you are comfortable with the results the tool provides.

Setting	Description
Automatic acceptance threshold	The percentage from which data classification suggestions must be accepted automatically. If you set this value to 75, then the classification suggestions with a confidence level of 75% or higher are automatically accepted. If multiple classification suggestions meet the threshold condition for a column, the classification suggestion with the highest confidence level percentage is accepted automatically if this classification suggestion is the only one to have that confidence level percentage. Example You set the automatic acceptance threshold to 85%. You classify a table with 2 columns. - For column A, three classification suggestions are possible, one with confidence level 93%, one with 92%, and one with 90%. - For column B, two classification suggestions are possible. Their confidence level is the same, 86%. The results of the automatic acceptance will be: - For column A, the classification suggestion with 93% will be accepted automatically. - For column B, nothing is done, both suggestions will be visible. The default acceptance threshold is 90.
Automatic rejection threshold	The percentage from which data classification suggestions must be rejected automatically. If you set this value to 49, then all data classification suggestions with a confidence level of 49% or lower are automatically rejected. The default rejection threshold is 10.

Note If the acceptance threshold and rejection threshold are set to the same value, and a data classification suggestion has this confidence level percentage, the classification suggestion will be rejected.

5. Click **Save all**.

Edit global data source registration settings

In Collibra Data Intelligence Cloud, you can create a data profile when you register a data source. You can configure global settings of data source registration.

More information: [Registering a data source](#) in the User Guide.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the section **Data Ingestion**, make the necessary changes.

Setting	Description
Table types to ignore	A comma separated list of table types that are not ingested. For example, <i>INDEX</i> and <i>SEQUENCE</i> .

Setting	Description
AWS regions restriction	A list of AWS regions Data Catalog is allowed to connect to. For example, <i>eu-west-3</i> and <i>us-east-2</i>. For a list of all AWS locations, see the AWS documentation. ◦ If you want to allow Collibra to make a connection to any AWS region, leave the field empty. ◦ If you remove a region from this list and the region was previously used for an S3 integration, you may want to delete the Glue database from the previously used region manually. By default, Collibra does not remove it. The Glue database has the following naming convention: <code>collibra_catalog_<Asset Id>_<Domain Id></code> For example: <code>collibra_catalog_d3174a88-5ffe-4d50-8fbe-7bf0832ec3af_5d198ce9-4e56-4d0e-a885-58204da50741</code> ◦ When using Edge, a warning is added to the logs if an invalid region is detected in the restricted regions list.
AWS API call rate	Allowed number of AWS API calls per second. Use this option to limit the number of API calls per second to prevent throttling errors from the AWS API.
Database registration via Edge	An option to enable database registration via Edge. ◦ ✓ True: Register a data source via Edge. ◦ ✗ False: Register a data source via Jobserver only. Note Enabling data source registration via Edge does not prevent you from registering a data source via Jobserver as well.
Collibra Data Quality & Observability Synchronization UI via DQ Connector on Edge	An option to enable the Data Quality extraction interface in Collibra ◦ ✓ True: The Quality extraction tab is available on the configuration page of a database asset ◦ ✗ False (default): The Quality extraction tab is not available and as such, it is not possible to extract and synchronize data quality information. You can only enable Collibra Data Quality & Observability synchronization if you also enabled Database registration via Edge.

Setting	Description
Google Cloud Storage synchronization via Edge	An option to enable Google Cloud Storage file system registration and synchronization via Edge. ✓ True: You can register and synchronize a Google Cloud Storage file system via Edge. ✗ False: You can't register a Google Cloud Storage file system via Edge.
Amazon S3 synchronization via Edge	An option to enable Amazon S3 file system registration and synchronization via Edge. ✓ True: You can register and synchronize an Amazon S3 file system via Edge. ✗ False: You can only register an Amazon S3 file system via Jobserver. Note Enabling the registration of an Amazon S3 file system via Edge does not prevent you from registering an Amazon S3 file system via Jobserver. For more information, see Working with Amazon S3.
Databricks Unity Catalog synchronization via Edge	An option to enable the integration of Databricks Unity Catalog via Edge. ✓ True: You can register and synchronize Databricks Unity Catalog via Edge. ✗ False: You cannot integrate Databricks Unity Catalog. Set this option to True.

Setting	Description
Source Tags Synchronization via Edge	An option to register and synchronize the tags on individual columns and tables in the data source during the data source registration via Edge. This means the tags become available and searchable in Data Catalog and can be used in business or policy processes and in workflows. Note Currently, you can synchronize source tags only from Snowflake. ✓ True: The Include Source Tags option becomes available when you define the rules for the synchronization of a data source via Edge. If you include the source tags, the tags defined on the assets in the data source are registered and available from the Schema, Table, Database View, and Column assets in the Source Tags attribute. ✗ False: Source tags are not registered in Data Catalog.

3. Click **Save all**.

Add a Jobserver to the DGC service

To register a data source and create a data profile in Collibra Data Intelligence Cloud, you need the Jobserver service.

If you don't have a Jobserver installed and [configured](#) in your environment, the **Register data source** action will be grayed out in the global create menu of Collibra Data Intelligence Cloud.

Tip Execute this procedure on Collibra Console of your cloud environment. In this configuration, the DGC service sends jobs to the on-premises Jobserver, however we highly recommend to revert this communication path so that the [Jobserver polls the DGC service for jobs](#).

Steps

1. Open the DGC service settings for editing:
2. In the **Jobserver** section, click **Add**.

3. Enter the necessary information:

Setting	Description
Jobserver list	The list of registered Jobserver instances.
Name	The name of the Jobserver as it will appear when you register a data source in Data Catalog. The name is a freely chosen name but it is recommended to only use alphanumeric characters and dashes, for example Jobserver-1. You will have to use this name as the ID of the gateway and in the address of this configuration.
Protocol	The protocol that is used for the communication between the Data Governance Center service and the Jobserver service. It is recommended to use HTTPS, especially if the services are hosted in different network segments.
Address	The address (IP address, URL, hostname) of the Jobserver.
Trusted server CA certificate	The certificate of the trusted CA needed to validate the server certificate. If blank, the default truststore will be used. The default truststore is defined in the SSL configuration section of the DGC service. The CA certificate of the server party (Jobserver).
Client certificate	The client certificate offered by the DGC service to the server. If blank, you cannot select mutual authentication as the Jobserver service authentication level.
Client private key	The private key of the DGC service's certificate.
Table profiling data size	The approximate maximum disk size of the data in MB that will be used to profile a table. The value cannot exceed 10,000.
Test connection timeout	This timeout is a time limit (in seconds) after which the connection test is stopped and a timeout error is shown. The default value is 60 seconds.

4. Click **Save all**.

Tip You can add as many [Jobserver services](#) as you want.

Multiple Jobserver services

You can add more than one Jobserver service to one Data Governance Center service. To get the best performance out of the Jobserver service, is to install the service as close to the data source as possible, preferably in the same network to reduce the latency as much as possible.

If you use multiple data sources in different locations, you can install a Jobserver service in every installation. During the [data source registration](#), you can then select the Jobserver that is closest to the selected data source.

Note

- A Jobserver can only run one ingestion job at a time. Multiple Jobservers can run multiple ingestion jobs in parallel.
- In a Collibra Data Intelligence Cloud environment with an on-premises Tableau server, you can use an on-premises Jobserver to ingest Tableau data. Keep in mind that this Jobserver is dedicated for the Tableau server and that it cannot be used for ingesting JDBC or Amazon S3 data sources.

Configure data profiling behavior

In Collibra Data Intelligence Cloud, you can create a data profile when you register a data source. You can configure the behavior of data profiling.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the section **Data Profiling**, make the necessary changes.

Setting	Description
Maximum value length	The maximum length of a value extracted during profiling or sampling. Additional characters are trimmed.
Default date pattern	The default format used to decode dates. It is the default pattern used for detecting dates when the Date Pattern and/or Time Pattern attribute is not specified in Column assets.
Default time pattern	The default format used to decode times. It is the default pattern used for detecting times when the Date Pattern and/or Time Pattern attribute is not specified in Column assets.
Default combined date and time pattern	The default format used to decode combined dates and times. It is the default pattern used for detecting combined dates and times when the Date Pattern and/or Time Pattern attribute is not specified in Column assets.
Empty values	A comma separated list of strings enclosed in double quotes. A value that matches one of those expressions is considered an empty value. Please note that a database null value is always considered an empty value, for example "", "na" and "none".
Data type detection threshold	The percentage of matching Column values to reach for an Advanced Data Type to be considered a possible Data Type for that Column. This is expressed as a value between 0.0 and 1.0).

Setting	Description
Anonymize data (Jobserver)	An option to anonymize sensitive data. ✓ True: Content in columns with data type Text or Geo is removed or replaced by a random hash value before the profiling results are sent to the cloud. ✗ False (default): No content is removed or replaced by a random hash value. Tip For anonymization via Edge, see setting "Anonymize Edge profiling results for all data types".
Database profiling via Edge	An option to enable profiling and classifying of synchronized metadata via Edge instead of Jobserver. ✓ True: Profiling and classification via Edge. ✗ False: Profile via Jobserver and classify via the Data Classification Platform. Note You can enable Database profiling via Edge only if you also enabled Database registration via Edge.
Maximum duration of a profiling Edge job	The maximum time duration, in minutes, that a profiling Edge job can run before Data Profiling stops the job. The default value is 20,160 minutes, 2 days. You can increase this limit to a maximum of 4 days.

Setting	Description
Parallel data-base profiling via Edge	The maximum number of schemas that Edge can profile at the same time. By default, the value of this setting is 4. This means Edge processes four profiling jobs at a time. This can have a huge positive impact on the performance of the profiling activity. You can increase this number to a maximum of 16. Note ◦ If you increase this number to more than four jobs, make sure that your Edge site resources are aligned with the extra requests it will receive. ◦ If you decrease this number and the running number of jobs exceeds the limit, no job will be canceled. Instead, there won't be any room to schedule a new job until at least one running job is completed. Example The parallel schema profiling via Edge setting is set to 4. ◦ For 1 database that contains 3 schemas, we will process all 3 schemas at the same time. ◦ For 2 databases that contain 4 schemas in total, we will process all 4 schemas at the same time. ◦ For 1 database that contains 8 schemas, we will start with 4 schemas and then proceed to the next ones as soon as a job is completed.
Anonymize Edge profiling results for all data types	Enable this option to anonymize all Edge profiling results stored in Collibra. ◦ ✓ True: Profiling results via Edge are anonymized for all columns. ◦ ✗ False (default): Profiling results via Edge are anonymized only for columns with the Text or Geo data type.
Calculate Data Similarity (Beta)	Enables the data similarity feature in your environment. ◦ ✓ True: Extra algorithms can run during profiling via Edge allowing the calculation of data similarity scores. The data similarity scores are currently used in Data Marketplace to show similar Table assets. ◦ ✗ False: (default): The beta feature is not enabled.

3. Click **Save all**.

Enable Tableau provisioning

You can enable Tableau provisioning to allow users to [create](#) Tableau provisioning files from Data Catalog data sets.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Beta Features** section, set the **Tableau provisioning enabled** field to **True**.
3. Click **Save all**.

Configure Cloud Data Classification Platform

When you want to use the Cloud Data Classification Platform in Data Catalog, you first have to configure it.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. Go to the **Data Classification** section.

3. Enter the required information:

Setting	Description
Machine Learning platform URL This setting requires the SUPER role.	The address of the machine learning platform that will classify your data.
Requester Name This setting requires the SUPER role.	The unique name to identify the client when using Machine Learning platform.
API key This setting requires the SUPER role.	The API Key to authorize the requester when connecting to the Machine Learning platform.
Enable Data Classification	✓ True: Enable Collibra's data classification technology. ✗ False (default): Do not use Collibra's data classification technology are not accepted.

4. If needed, configure the automatic classification acceptance and rejection.

Setting	Description
Enable automatic classification acceptance and rejection	✓ True: The automatic acceptance and rejection of data classification suggestions is active. ✗ False (default): Data classification suggestions are not automatically accepted or rejected. Tip Start using the Automatic Data Classification tool by manually accepting and rejecting the data classification suggestions. Only activate the automatic acceptance and rejection feature if you are comfortable with the results the tool provides.

Setting	Description
Automatic acceptance threshold	The percentage from which data classification suggestions must be accepted automatically. If you set this value to 75, then the classification suggestions with a confidence level of 75% or higher are automatically accepted. If multiple classification suggestions meet the threshold condition for a column, the classification suggestion with the highest confidence level percentage is accepted automatically if this classification suggestion is the only one to have that confidence level percentage. Example You set the automatic acceptance threshold to 85%. You classify a table with 2 columns. - For column A, three classification suggestions are possible, one with confidence level 93%, one with 92%, and one with 90%. - For column B, two classification suggestions are possible. Their confidence level is the same, 86%. The results of the automatic acceptance will be: - For column A, the classification suggestion with 93% will be accepted automatically. - For column B, nothing is done, both suggestions will be visible. The default acceptance threshold is 90.
Automatic rejection threshold	The percentage from which data classification suggestions must be rejected automatically. If you set this value to 49, then all data classification suggestions with a confidence level of 49% or lower are automatically rejected. The default rejection threshold is 10.

Note If the acceptance threshold and rejection threshold are set to the same value, and a data classification suggestion has this confidence level percentage, the classification suggestion will be rejected.

5. Click **Save all**.

Enable or disable Catalog experience

Catalog experience allows you to use the improved user experience in the [Data Catalog asset pages](#).

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Catalog Experience** section, enter the required information:

Setting	Description
Enable Catalog experience	◦ ✓ True: Catalog experience is enabled. This will improve the layout of Data Catalog's asset pages, such as those of Data Set, Schema, Table and Column assets. ◦ ✗ False: Catalog experience is disabled.
Catalog Experience Titlebar theme	The theme for the Catalog experience. You can choose between the LIGHT and DARK. This option is only applicable if the Enable Catalog experience option is enabled.

3. Click **Save all**.

Enable the registration of a data source via Edge

You can enable Edge to register a data source. When you [register a data source](#) via Edge, an Edge site ingests data into Data Catalog instead of Jobserver.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Register data source** section, enter the required information:

Setting	Description
Database registration via Edge	An option to enable database registration via Edge. ◦ ☒ True: Register a data source via Edge. ◦ ☐ False: Register a data source via Jobserver only. Note Enabling data source registration via Edge does not prevent you from registering a data source via Jobserver as well.

Setting	Description
Source Tags Synchronization via Edge	An option to register and synchronize the tags on individual columns and tables in the data source during the data source registration via Edge. This means the tags become available and searchable in Data Catalog and can be used in business or policy processes and in workflows. Note Currently, you can synchronize source tags only from Snowflake. ✓ True: The Include Source Tags option becomes available when you define the rules for the synchronization of a data source via Edge. If you include the source tags, the tags defined on the assets in the data source are registered and available from the Schema, Table, Database View, and Column assets in the Source Tags attribute. ✗ False: Source tags are not registered in Data Catalog.

3. Click **Save all**.
4. Restart the environment to apply your changes. See [Stop an environment](#) and [Start an environment](#).

Enable profiling for Edge

To enable Edge profiling and classification of synchronized metadata in Data Catalog, you need to run a command and enable multiple settings.

Before you begin

- You have [enabled Database registration via Edge](#).
- You have run the command to [enable classification on your Edge site](#).

Required permissions

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Data profiling** section, enter the required information:

Setting	Description
Database profiling via Edge	An option to enable profiling and classifying of synchronized metadata via Edge instead of Jobserver. ◦ ✓ True: Profiling and classification via Edge. ◦ ✗ False: Profile via Jobserver and classify via the Data Classification Platform. Note You can enable Database profiling via Edge only if you also enabled Database registration via Edge.
Maximum duration of a profiling Edge job	The maximum time duration, in minutes, that a profiling Edge job can run before Data Profiling stops the job. The default value is 20,160 minutes, 2 days. You can increase this limit to a maximum of 4 days.

Setting	Description
Parallel schema profiling via Edge	The maximum number of schemas that Edge can profile at the same time. By default, the value of this setting is 4. This means Edge processes four profiling jobs at a time. This can have a huge positive impact on the performance of the profiling activity. You can increase this number to a maximum of 16. Note ◦ If you increase this number to more than four jobs, make sure that your Edge site resources are aligned with the extra requests it will receive. ◦ If you decrease this number and the running number of jobs exceeds the limit, no job will be canceled. Instead, there won't be any room to schedule a new job until at least one running job is completed. Example The parallel schema profiling via Edge setting is set to 4. ◦ For 1 database that contains 3 schemas, we will process all 3 schemas at the same time. ◦ For 2 databases that contain 4 schemas in total, we will process all 4 schemas at the same time. ◦ For 1 database that contains 8 schemas, we will start with 4 schemas and then proceed to the next ones as soon as a job is completed.
Anonymize profiling data for all data types in Edge	Enable this option to anonymize all Edge profiling results stored in Collibra. ◦ ✓ True: Profiling results via Edge are anonymized for all columns. ◦ ✗ False (default): Profiling results via Edge are anonymized only for columns with the Text or Geo data type.

Note

- You don't need to enable the **Anonymize data (Jobserver)** setting because this setting is not relevant for Edge.
- You don't need to enable the **Enable Data Classification** setting in the **Data Classification configuration** section. This setting relates only to the [Data Classification Platform](#).
If this setting is set to `true`, the **Classify** button is available on Column and Table asset pages. This button allows you to classify data via the Data Classification Platform. However, when using [profiling and classification via Edge](#), you don't need the Data Classification Platform.

3. Click **Save all**.

What's next?

[Continue the configuration for profiling and classification.](#)

Enable data quality synchronization via Edge

You can enable the [Quality extraction](#) tab on a Database asset page in Data Catalog.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.

- d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Register data source** section, enter the required information:

Setting	Description
Collibra Data Quality & Observability Synchronization UI via DQ Connector on Edge Edge	An option to enable the Data Quality extraction interface in Collibra ✓ True: The Quality extraction tab is available on the configuration page of a database asset ✗ False (default): The Quality extraction tab is not available and as such, it is not possible to extract and synchronize data quality information. You can only enable Collibra Data Quality & Observability synchronization if you also enabled Database registration via Edge.

3. Click **Save all**.

Enable Tableau metadata API

You can enable the [Tableau metadata API](#) if you want to ingest Tableau 2020.2 or newer in Data Catalog.

Note Make sure that the Tableau metadata API is [enabled](#) in Tableau before you ingest or synchronize Tableau in Data Catalog.

Warning If you upgrade to Tableau version 2020.2 or newer, but previously synchronized an older Tableau version via the REST API and XML mapping, you have to prepare the [migration procedure](#) to prevent losing manually added relations, attributes, tags, comments and stitching results.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Tableau metadata API** section, enter the required information:

Setting	Description
Enable Tableau metadata API	✓ True: Tableau metadata API is enabled. This enables you to ingest Tableau 2020.2 or newer into Data Catalog. ✗ False: Tableau metadata API is disabled. If you ingest Tableau 2020.2 or newer, the ingestion will fail. This prevents data loss of manually added relations and attributes.

3. Click **Save all**.

Anonymize data via Jobserver

You can enable or disable the option to anonymize the content of columns with data type TEXT and GEO after the profiling process via Jobserver.

Tip For information about Edge, go to [Anonymization via Edge](#).

Warning Currently, if you enable the data anonymization process you can no longer use automatic data classification via the Data Classification platform. However, you can still classify and anonymize profiling results if you [use Edge](#).

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Data Profiling** section, enter the required information:

Setting	Description
Maximum number of samples	The maximum number of samples you want to collect for a data source. The default value is 100. The maximum value is 1,000. This setting is specific to sample data .
Maximum value length	The maximum length of a value extracted during profiling or sampling. Additional characters are trimmed.
Default date pattern	The default format used to decode dates. It is the default pattern used for detecting dates when the Date Pattern and/or Time Pattern attribute is not specified in Column assets.
Default time pattern	The default format used to decode times. It is the default pattern used for detecting times when the Date Pattern and/or Time Pattern attribute is not specified in Column assets.
Default combined date and time pattern	The default format used to decode combined dates and times. It is the default pattern used for detecting combined dates and times when the Date Pattern and/or Time Pattern attribute is not specified in Column assets.
Empty values	A comma separated list of strings enclosed in double quotes. A value that matches one of those expressions is considered an empty value. Please note that a database null value is always considered an empty value, for example "", "na" and "none".

Setting	Description
Data type detection threshold	The percentage of matching Column values to reach for an Advanced Data Type to be considered a possible Data Type for that Column. This is expressed as a value between 0.0 and 1.0).
Anonymize data (Jobserver)	An option to anonymize sensitive data. ✓ True: Content in columns with data type Text or Geo is removed or replaced by a random hash value before the profiling results are sent to the cloud. ✗ False (default): No content is removed or replaced by a random hash value. Tip For anonymization via Edge, see setting "Anonymize Edge profiling results for all data types".
Database profiling via Edge	An option to enable profiling and classifying of synchronized metadata via Edge instead of Jobserver. ✓ True: Profiling and classification via Edge. ✗ False: Profile via Jobserver and classify via the Data Classification Platform. Note You can enable Database profiling via Edge only if you also enabled Database registration via Edge.
Maximum duration of a profiling Edge job	The maximum time duration, in minutes, that a profiling Edge job can run before Data Profiling stops the job. The default value is 20,160 minutes, 2 days. You can increase this limit to a maximum of 4 days.

Setting	Description
Parallel data-base profiling via Edge	The maximum number of schemas that Edge can profile at the same time. By default, the value of this setting is 4. This means Edge processes four profiling jobs at a time. This can have a huge positive impact on the performance of the profiling activity. You can increase this number to a maximum of 16. Note ◦ If you increase this number to more than four jobs, make sure that your Edge site resources are aligned with the extra requests it will receive. ◦ If you decrease this number and the running number of jobs exceeds the limit, no job will be canceled. Instead, there won't be any room to schedule a new job until at least one running job is completed. Example The parallel schema profiling via Edge setting is set to 4. ◦ For 1 database that contains 3 schemas, we will process all 3 schemas at the same time. ◦ For 2 databases that contain 4 schemas in total, we will process all 4 schemas at the same time. ◦ For 1 database that contains 8 schemas, we will start with 4 schemas and then proceed to the next ones as soon as a job is completed.
Anonymize Edge profiling results for all data types	Enable this option to anonymize all Edge profiling results stored in Collibra. ◦ ✓ True: Profiling results via Edge are anonymized for all columns. ◦ ✗ False (default): Profiling results via Edge are anonymized only for columns with the Text or Geo data type.
Calculate Data Similarity (Beta)	Enables the data similarity feature in your environment. ◦ ✓ True: Extra algorithms can run during profiling via Edge allowing the calculation of data similarity scores. The data similarity scores are currently used in Data Marketplace to show similar Table assets. ◦ ✗ False: (default): The beta feature is not enabled.

3. Click **Save all**.

Enable or disable the Settings landing page

You can enable or disable the Settings landing page.

Prerequisites

- You have the **ADMIN** or **SUPER** role in Collibra Console.
- You have the **SUPER** role in Collibra Console.

Steps

1. Open the DGC service settings for editing:
 - a. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
 - b. In the tab pane, expand an environment to show its services.
 - c. In the tab pane, click the Data Governance Center service of that environment.
 - d. Click **Configuration**.
 - e. Click **Edit configuration**.
2. In the **Beta features** section, enter the required information:

Setting	Description
Settings landing enabled	◦ ✓ True (default): Use the new Settings landing page in the Collibra user interface. ◦ ✗ False : Do not use the Settings landing page in the Collibra user interface.

3. Click **Save all**.

Configure the Search service

In the Search service configuration, you can edit the TCP and Transport port of the service.

1. Open Collibra Console with a user profile that has the **SUPER** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the tab pane, click the **Search** service of a Collibra DGC environment.
 - » The details of the **Search** service are shown.
3. Click **Configuration**.
4. Click **Edit configuration**.
5. Edit the search service settings.

Setting	Description
HTTP Port (Requires restart)	The TCP port to access the Search service via REST API. The default port is 4421.
Transport Port (Requires restart)	The TCP port for the communication between the DGC and the Search service. The default port is 4422.

6. Click **Save all**.

Search service configuration options

To edit the Search service configuration options, you need the [SUPER](#) role.

Setting	Description
HTTP Port (Requires restart)	The TCP port to access the Search service via REST API. The default port is <i>4421</i> .
Transport Port (Requires restart)	The TCP port for the communication between the DGC and the Search service. The default port is <i>4422</i> .

Jobserver service configuration

If you want to ingest data and profile that data, you need a Jobserver service, often referred to as Jobserver. If both services are in the same trusted network segment, you can use the unsecure HTTP protocol with no authentication. However, it is highly recommended to configure a secure communication.

The configuration of the Jobserver service is to enable the communication with the Data Governance Center service.

In this section, you will learn more about configuring the Jobserver security and adding the Jobserver service to the DGC service.

You can add more than one Jobserver service to the DGC service.

In this chapter

Jobserver authentication levels	126
Mutual authentication between Jobserver and DGC service	127
General specifications for certificates and private keys	129
Edit the Jobserver service settings	130
Connection from an on-premises Jobserver to a Collibra Data Intelligence Cloud	132
Connection from Collibra Data Intelligence Cloud to an on-premises Jobserver	148
Jobserver best practices	152

Jobserver authentication levels

The Jobserver service handles customer data to, for example, perform profiling on it or ingest data into Collibra Data Intelligence Cloud. As this is sensitive data, we have to make sure that the communication with the Jobserver is as secure as possible. This section provides more information about the authentication levels available at the Jobserver configuration.

The Jobserver service supports three authentication levels.

Authentication level	Description
MUTUAL	Mutual authentication or two-way authentication is the authentication system where the two parties authenticate each other during their handshake. We only support certificate based mutual authentication, not username/password-based.
SERVER	Server authentication is the authentication system where the client authenticates the server identity but not the other way around. Warning For security reasons, do not use this level when you are using an on-premises Jobserver and a Collibra Data Intelligence Cloud environment. We only support certificate based server authentication, not username/password-based.
NONE	No authentication, no encryption. This is plain HTTP communication. Warning This is an insecure connection, your data can be exposed.

No authentication level

If you don't set an authentication level, you communicate with the Jobserver via an unsecure connection, using the HTTP protocol. This authentication level should only be used when the Jobserver service runs in the same trusted network segment as the Data

Governance Center service. But even when both services run in the same network segment, we recommend to avoid this authentication level as malicious software or network security breaches could still expose your data.

Server authentication

If you select server authentication, the client party (DGC service) authenticates the server party (Jobserver). If the authentication succeeds, the DGC service is certain that it communicates with the selected Jobserver service.

Mutual authentication

If you select mutual authentication, both parties (DGC and Jobserver services) authenticate each other at the same time. If the authentication succeeds, both parties are certain about each other's identity. This is the most secure authentication level and is therefore the recommended authentication level for the Jobserver service.


For more detailed information about the mutual authentication between the Jobserver and DGC service, see [Mutual authentication between Jobserver and DGC service](#).

Mutual authentication between Jobserver and DGC service

If you want to ingest data and profile that data, you can use a Jobserver service, often referred to as Jobserver. If the Data Governance Center service and Jobserver service are in the same network segment, you can use the insecure HTTP protocol with no authentication. However, it is highly recommended to configure a secure communication.

In the following schema, you can see how the communication is established for mutual authentication between the two services. This authentication level is the recommended level whenever both services are in a different network segment, for example when you have a cloud-based Collibra Data Intelligence Cloud environment and an on-premises Jobserver.

If there is a reverse proxy server between these services, see [Mutual authentication with a reverse proxy server](#).



Step	Description
Hello	DGC service initiates the communication to the Jobserver service by sending a Hello message.
JS Certificate Public Key	The Jobserver service sends its public key to the DGC service. The DGC service can then authenticate the Jobserver service.
Validate JS Certificate Public Key	The DGC service validates the received public key with the Jobserver's CA certificate. This means that the DGC service has the Jobserver's CA certificate in its truststore.
DGC Certificate Public Key	The DGC service sends its public key to the Jobserver service. The Jobserver service can then authenticate the DGC service.
Validate DGC Certificate Public Key	The Jobserver service validates the received public key with the DGC CA certificate. This means that the Jobserver service has the DGC CA certificate in its truststore.

If all these steps are completed successfully, then both services are 100% sure that they are communicating with the counterparty that they expect.

It is only the DGC service that initiates the communication to the Jobserver service. The communication from a Jobserver to the DGC service is only possible with a [reverse proxy](#).

General specifications for certificates and private keys

The certificates that you can use for the authentication between the Jobserver service and Data Governance Center service must meet the specifications as described in the following table.

Element	Specifications
Certificate	• The file name of the certificate can be freely chosen. • The certificate format must be in PEM format.
Private key	• The file name of the private key can be freely chosen. • The key format must be in PEM format. • The key must be packaged in PKCS#8 format. • The key is not protected by a passphrase.

If you already have a private key, but it does not meet the requirements, you can [Convert a private key format](#) the private key format.

Convert a private key format

If your company uses private keys, you can use them for server or mutual authentication between the services in a Collibra Data Intelligence Cloud environment, provided that they meet the specifications.

If they do not meet the specifications, you can convert the keys with the following command:

```
openssl pkcs8 -topk8 \
-inform PEM \
-outform PEM \
-in <path_to_original_key> \
-nocrypt \
-out <path_to_new_key>
```

Edit the Jobserver service settings

The Jobserver service infrastructure settings allow you to edit the behavior of the connections to this service and the Java Virtual Machine (JVM) parameters.

1. Open Colibra Console with a user profile that has the **SUPER** role.
 - » Colibra Console opens with the **Infrastructure** page.
2. In the tab pane, click **Jobserver** service of a Colibra DGC environment.
 - » The details of the **Jobserver** service are shown.
3. Click **Infrastructure Configuration**.
4. Click **Edit configuration**.
5. Edit the [Jobserver service infrastructure configuration](#).
6. Click **Save all**.

Tip

- You can navigate to a specific section by clicking it in the tab pane.
- When you edit certain fields, the **i** icon is displayed next to it. When you click it, it displays the default value for that field and a **Reset** button.
- If you have to restart Colibra or execute extra actions to apply the new settings, it is indicated in the user interface.

Jobserver infrastructure configuration options

To edit the Jobserver infrastructure configuration options, you need the [SUPER](#) role.

Application server configuration

Setting	Description
Jobserver Monitoring Port (Requires restart)	The port that is used by the Monitoring service to monitor the Jobserver service. The default port is 4424 .
Spark Monitoring Port (Requires restart)	The port that is used by the Monitoring service to monitor the Spark service. The default port is 4434 .
Jobserver memory (Requires restart)	The memory that is assigned to the Jobserver service. The Jobserver controls the jobs that are executed for Data Catalog.
Spark memory (Requires restart)	The memory that is assigned to the Spark application. This application is responsible for the actual data ingestion and profiling.

Security configuration

Setting	Description
Authentication level	The authentication level to communicate with the Jobserver. The client (DGC service) must be configured according the here selected authentication level.
Server certificate chain	The certificate or certificate chain with the public key that is offered by the Jobserver to the DGC service.
Server private key	The private key that is part of the Jobserver's certificate or certificate chain.
Trusted client CA certificate	The certificate of the trusted CA used to validate the client certificate (DGC service). To restrict authentication to this client, the CA should be exclusively used by this server.

JVM configuration

The Java Virtual Machine configuration parameters to run the Jobserver service.

Warning

Only apply changes to these settings if you are really experienced with JVM parameters. Changing parameters may cause serious performance issues.

To remove an individual JVM property you must use the delete icon () alongside the property, otherwise, the service will interpret it as a blank line and fail to start correctly.

Restart the service after editing the JVM parameters.

Context JVM configuration

The context configuration parameters to run ingestion and profiling jobs.

Warning

Only apply changes to these settings if you are really experienced with JVM parameters. Changing parameters may cause serious performance issues.

To remove an individual JVM property you must use the delete icon () alongside the property, otherwise, the service will interpret it as a blank line and fail to start correctly.

Restart the service after editing the JVM parameters.

Connection from an on-premises Jobserver to a Collibra Data Intelligence Cloud

In default installations, a Jobserver is installed on-premises and a Collibra Data Intelligence Cloud sends ingestion and profiling jobs to it. However, we highly recommend to [reverse this communication](#), so that the on-premises Jobserver polls for jobs. This is often required for security reasons.

Note If you want a Collibra Data Intelligence Cloud to send jobs to your on-premises Jobserver, contact your security officer and network administrator.

In this section, you get more information on how to set up the communication from an on-premises Jobserver to a Collibra Data Intelligence Cloud.


On-premises Jobserver to Colibra Data Intelligence Cloud communication

By default, the Data Governance Center service sends ingestion and profiling jobs to the Jobserver. This means that if you are using a Colibra Data Intelligence Cloud environment with an on-premises Jobserver, there is an inbound connection to the customer network, which is often not possible for security reasons. To allow such connections, you can use a [reverse proxy server](#).

But, instead of using a reverse proxy, you also have the possibility to reverse this communication, where the on-premises Jobserver initiates the communication to Colibra Data Intelligence Cloud.

Communication overview

The following schema shows the communication paths from an on-premises Jobserver to Colibra Data Intelligence Cloud or to an on-premises Tableau server:



In case of an on-premises Tableau server, it is possible that inbound connections to it are not allowed. To establish a communication between your Colibra Data Intelligence Cloud and Tableau server, you will need a Jobserver that is dedicated to ingest from the Tableau server. The configuration of the communication from the Jobserver to the Tableau server is similar to the one from a Jobserver to a Colibra Data Intelligence Cloud environment.

Each Jobserver has to be a dedicated Jobserver, you cannot use a Jobserver to ingest from both Tableau server and S3 or JDBC data sources.

Components

To enable communication from an on-premises Jobserver to Collibra Data Intelligence Cloud, there are two new components:

New component	Description
Reverse HTTP proxy servlet	The reverse HTTP proxy is part of the DGC service. It acts as a server for all other Collibra services, whether they are installed on-premises or together with the DGC service in the cloud.
Gateway	The gateway is part of the Jobserver service. It polls the DGC service's reverse proxy to fetch tasks and send them to the Jobserver.

You only need the gateway to communicate with an on-premises Tableau server.

Encrypt passwords for basic authentication

In the Jobserver service configuration, you have to enter an encrypted password. To encrypt the password, use the **reversehttp-gateway-standalone** utility.

Prerequisites

- You have downloaded the [reversehttp-gateway-standalone-7.0.0.jar](#) file.
- You have the password of the Collibra user that you use to connect to your Collibra Data Intelligence Cloud environment.

Steps

Note For security reasons, we have truncated the encrypted password in the example.

1. Open a terminal or command prompt session.
2. Go to the folder that contains the downloaded JAR file.

3. Execute the following command:

```
java -jar reversehttp-gateway-standalone-7.0.0.jar encrypt  
  
Collibra Reverse HTTP Gateway  
Enter value to encrypt: <password of Collibra user>  
Re-enter value to encrypt: <password of Collibra user>  
Encrypted value: encrypted:k7ScuJ3...
```

Note If the entered values in this command don't match, it will ask the values again.

What's next?

If you use an encrypted password in a [configuration](#), use the full string of the **Encrypted value** result. This includes the prefix "encrypted:".

Check connectivity between an on-premises Jobserver and Collibra Data Intelligence Cloud

Before you install your on-premises Jobserver, we highly recommend to do some basic connectivity tests between the node on which you are going to install the Jobserver service and your Collibra Data Intelligence Cloud environment. If the node cannot reach your Collibra Data Intelligence Cloud environment, then first fix the connectivity before you start the Jobserver installation.

1. Open the connection and port from your on-premises Jobserver to your Collibra Data Intelligence Cloud environment.
2. Whitelist the DNS name of your Collibra Data Intelligence Cloud environment.
3. Check if the Jobserver node can reach your Collibra Data Intelligence Cloud environment with the following command:

```
curl -x "" -i https://<your-environment>.collibra.com/reversehttp-poll/1
```

You should receive an HTTP 401 response. If you don't receive an HTTP 401 response, fix the connectivity before proceeding with the configuration.

Tip If there is a proxy server between your on-premises Jobserver and your Collibra Data Intelligence Cloud environment, then add the proxy address to the command:

```
curl -x http(s)://proxy:port -i https://<your-  
environment>.collibra.com/reversehttp-poll/1
```

What's next?

When you receive an HTTP 401 response, you can start the [installation of the Jobserver](#).

Configure the Jobserver to Collibra Data Intelligence Cloud communication

By default, Collibra Data Intelligence Cloud sends jobs to a Jobserver but you also have the possibility to have the Jobserver poll Collibra Data Intelligence Cloud for jobs.

In this section, we describe how to configure the Jobserver to poll Collibra for jobs. You will have to configure both the Data Governance Center service and the Jobserver service.

Prerequisites

- You have Collibra Data Intelligence Cloud 2020.10 or newer.
- You have [created a keystore in the PKCS#12 format](#) on the node that hosts the Jobserver service.

Steps

Configure the Jobserver service

Warning

Only apply changes to these settings if you are really experienced with JVM parameters. Changing parameters may cause serious performance issues.

To remove an individual JVM property you must use the delete icon (🗑) alongside the property, otherwise, the service will interpret it as a blank line and fail to start correctly.

Restart the service after editing the JVM parameters.

Execute the following steps in the Collibra Console instance that manages your Jobserver:

1. Open Collibra Console with a user profile that has the **SUPER** role.

» Collibra Console opens with the **Infrastructure** page.

Tip The default address to access Collibra Console is `<server hostname>:4402`, but you may have set another port during the installation of Collibra Console. Keep in mind that a firewall of your operating system can block the access to Collibra Console.

2. In the tab pane, click the Jobserver service of a Collibra environment.
3. Click **Infrastructure Configuration**.
4. Click **JVM configuration**.
5. Click **Edit configuration**.
6. Add the following JVM settings:

Setting	Description
reversehttp.gateway	The setting to enable the Jobserver's gateway. To enable the communication from the on-premises Jobserver to the Collibra Data Intelligence Cloud environment., this value must be <i>true</i>. Example <code>-Dreversehttp.gateway=true</code>

Setting	Description
proxy.url	The URL of the Colibra environment, followed by <i>reversehttp-poll/<gateway-id></i>. This "gateway-id" must be identical to the one used in the Name parameter when you add the Jobserver to the DGC service. The value of this setting is case-sensitive. Example -Dproxy.url=https://<your-environment-url>/reversehttp-poll/Jobserver-1
target.url	The URL of your the target system, either an on-premises Jobserver or a Tableau server. Example ◦ Jobserver: - Dtarget.url=http://localhost:4404 ◦ Tableau server: - Dtarget.url=https://tableau-sales.yourcompany.com
http.proxy.host (optional)	The hostname of the HTTP proxy server for outbound connections to your Colibra Data Intelligence Cloud environment. This option is used to enable outbound traffic monitoring. Example - Dhttp.proxy.host=proxy.yourcompany.com
http.proxy.port (optional)	The port of the HTTP proxy server for outbound connections to your Colibra Data Intelligence Cloud environment. This option is used to enable outbound traffic monitoring. Example -Dhttp.proxy.port=8080

Setting	Description
username	The username of any Collibra user for basic authentication. <pre>Example -Dusername=john.fisher</pre>
password	The corresponding password of the Collibra user for basic authentication. <pre>Example -Dpassword=ChangeMe</pre> You can encrypt this password if necessary. <pre>Example -Dpassword=enc_ 2:t2rklBY6699aWV0...</pre>
keystore.path	The full path to the PKCS12 keystore. This keystore should contain the private key to sign the basic authentication header. <pre>Example -Dkeystore.path=/opt/collibra_ data/spark- jobserver/security/jobserver-1- keystore.p12</pre>
keystore.alias	The alias of the private key in the keystore. Each alias must be unique in your configuration. If you used the <code>name</code> argument during the creation of the keystore, then use the value of this <code>name</code> argument. If only 1 keystore is created, the default alias is "1". <pre>Example -Dkeystore.alias=1 -Dkeystore.alias=MyJobserver</pre>

Setting	Description
keystore.password (optional)	The password to access the keystore. If the keystore is not password-protected, don't add it to the JVM settings. You can encrypt this password if necessary. <pre>Example -Dkeystore.password=ChangeMe</pre>
keystore.key.password (optional)	The password to use the private key, only applicable if you secured the private key with a password. If the key is not password-protected, don't add it to the JVM settings. You can encrypt this password if necessary. <pre>Example -Dkeystore.key.password=ChangeMe</pre>
polling.backoff	The time in milliseconds between a polling failure and a next polling attempt. We recommend to not define this parameter, it then uses the default value of 5,000 milliseconds. <pre>Example -Dpolling.backoff=10000</pre>
max.connections.route	The maximum number of HTTP connections per route. We recommend to not define this parameter, it then uses the default value of 20. <pre>Example -Dmax.connections.route=30</pre>
max.connections.total	The maximum number of all HTTP connections. We recommend to not define this parameter, it then uses the default value of 40. <pre>Example -Dmax.connections.total=60</pre>

Setting	Description
<code>idle.connection.timeout</code>	The time in milliseconds that an idle connection is kept in the connection pool. We recommend to not define this parameter, it then uses the default value of 5,000 milliseconds. <pre>Example -Didle.connection.timeout=3000</pre>
<code>connection.timeout</code>	The time in milliseconds that the reverse HTTP server waits for a response from your Collibra environment or from the value in target.url. If you don't set this parameter, the value is 60,000 milliseconds. <pre>Example -Dconnection.timeout=30000</pre>
<code>connection.soTimeout</code>	The time in milliseconds that the reverse HTTP server waits for a response from your Collibra environment or from the value in target.url on socket level. If you don't set this parameter, the value is 60,000 milliseconds. <pre>Example -Dconnection.soTimeout=30000</pre>
<code>polling.timeout</code>	The time in milliseconds that the reverse HTTP server waits for a poll request from your Collibra environment to be submitted to the target.url. If you don't set this parameter, the value is 300,000 milliseconds. <pre>Example -Dpolling.timeout=100000</pre>

Setting	Description
polling.period	The time in milliseconds that the reverse HTTP server waits in between poll request sessions. In other words, after having received a poll request or no request from your Collibra environment, the reverse HTTP server waits a certain amount of milliseconds before contacting the Collibra environment again. If you don't set this parameter, the value is 100 milliseconds. <pre>Example -Dpolling.period=200</pre>
health.check.enabled (optional)	Enables the health check mechanism between the Collibra environment and the reverse HTTP server. If you don't set this parameter, the value is false. <pre>Example -Dhealth.check.enabled=true</pre>
health.check.period (optional)	The time in milliseconds that the reverse HTTP server waits between health checks of its connection with Collibra. If you don't set this parameter, the value is 5,000 milliseconds. <pre>Example -Dhealth.check.period=10000</pre>
health.check.timeout (optional)	The time in milliseconds that the reverse HTTP server waits for a health check response from Collibra. If you don't set this parameter, the value is 5,000 milliseconds. <pre>Example -Dhealth.check.timeout=10000</pre>

Note You have to use separate Jobobservers for the ingestion of S3 or JDBC data sources and Tableau server data.

7. Click the green **Save all** button.
8. Click **Security configuration**.

9. Click **Edit configuration**.
10. Set the **Authentication level** to *NONE*.

Note This means that there is a one-way outbound communication over TLS from the Jobserver to the Collibra environment, note that there is no authentication at all.

11. Click the green **Save all** button.

Add the Jobserver to the DGC service

Execute the following steps in Collibra Console of your Collibra Data Intelligence Cloud environment.

1. Open the DGC service settings for editing:
2. Go to the **Jobserver** section of the configuration.
3. Enter the required information.

Setting	Description
Name	The name of the Jobserver as it will appear when you register a data source. The name is a freely chosen name but it is recommended to only use alphanumerical characters and dashes, for example Jobserver-1. You will have to use this name as the ID of the gateway and in the address of this configuration.
Protocol	The protocol for this configuration has to be <i>HTTP</i> and not the recommended <i>HTTPS</i>, this is because of the Collibra internal architecture.
Address	The loopback address of the DGC service, followed by <i>/reversehttp/<gateway-id></i>. The "gateway-id" must be identical to the one used in the Name parameter of this configuration. Do not use the scheme in the address. Example localhost:4400/reversehttp/Jobserver-1

Setting	Description
Trusted server CA certificate	The certificate in PEM format that contains the public key of the Jobserver to validate the signature of the basic authentication header. In the example to create a keystore, this is the content of the file cert.pem. <pre> Example -----BEGIN CERTIFICATE----- MIICqDCCAZACCQCcy3Oq51c5YzANBgkqhkiG9w0BAQsF ADAWMRQwEgYDVQQDDAtq b2JzZXJ2ZXIt... -----END CERTIFICATE----- </pre>
Client certificate	This field is not used in this configuration.
Client private key	This field is not used in this configuration. Note This field always shows dots, even if it is empty.
Table profiling data size	The approximate maximum disk size of the data in MB that will be used to profile a table. The value cannot exceed 10 000.

If you use a truststore, go to [Generate keys, certificates and keystores](#).

4. Click the green **Save all** button.

If all settings and communication paths are correctly configured, you will see a notice on the Jobserver:

```
INFO [I/O dispatcher 1] reversehttp.gateway.PollingController -
proxy -> no requests polled (204)
```

What's next?

When you have set up this communication, you may want to [monitor the outbound traffic](#). You can do so by enabling a man-in-the-middle proxy.

Monitor outbound traffic

If you set up the communication from your on-premises Jobserver to your Collibra Data Intelligence Cloud environment, you may want to monitor the outbound traffic. You can do so by setting up a man-in-the-middle proxy (MITM proxy).

1. Open Collibra Console with a user profile that has the **SUPER** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the tab pane, click the Jobserver service of a Collibra environment.
3. Click **Infrastructure Configuration**.
4. Click **JVM configuration**.
5. Click **Edit configuration**.
6. Add the following JVM settings:

Setting	Description
http.proxy.host	The hostname of the HTTP proxy server for outbound connections to your Collibra Data Intelligence Cloud environment. <pre>Example -Dhttp.proxy.host=proxy.yourcompany.com</pre>
http.proxy.port	The port of the HTTP proxy server for outbound connections to your Collibra Data Intelligence Cloud environment. <pre>Example -Dhttp.proxy.port=8080</pre>

7. Add the CA certificate of this MITM proxy in the Jobserver's truststore
`($ (COLLIBRA_DIR) /jre/lib/security/cacerts).`

Generate keys, certificates and keystores

For a [secure communication](#) between the Jobserver and Collibra Data Intelligence Cloud, you can use certificates.

In the current configuration, certificates are used as containers for public keys and the keystore is used to store private keys and certificates.

- On the node that hosts the Jobserver service, the keystore must be in PKCS#12 format.
- On the node that hosts the Data Governance Center service, you need a certificate, in PEM format, which includes the public key.

Steps

Note The commands used in this procedure are only examples, ask your Security officer for more information.

1. On the node on which you want to install the keystore, certificate and private key, open a terminal or command prompt session.
2. Go to or create a directory in which you want to create the keystore.
3. Create the private key and certificate:

```
openssl req -x509 -newkey rsa -keyout key.pem -out cert.pem
-days 365
```

```
Generating a 2048 bit RSA private key
```

```
.....+++
.....+++
```

```
writing new private key to 'key.pem'
```

```
Enter PEM pass phrase: <optional password>
```

```
Verifying - Enter PEM pass phrase: <repeat password>
```

```
-----
```

```
You are about to be asked to enter information that will be
incorporated
```

```
into your certificate request.
```

```
What you are about to enter is what is called a Distin-
guished Name or a DN.
```

```
There are quite a few fields but you can leave some blank
```

```
For some fields there will be a default value,
```

```
If you enter '.', the field will be left blank.
```

```
-----
```

```
Country Name (2 letter code) [AU]:US
```

```
State or Province Name (full name) [Some-State]:New York
```

```
Locality Name (eg, city) []:New York
```

```
Organization Name (eg, company) [Internet Widgits Pty
Ltd]:Collibra
```

```
Organizational Unit Name (eg, section) []:
```

```
Common Name (e.g. server FQDN or YOUR name) []:Francois
Lemaire
```

```
Email Address []:francois.lemaire@collibra.com
```

4. Create a PKCS#12 keystore including a private key and certificate.

```
openssl pkcs12 -export -inkey key.pem -in cert.pem -out key-  
store.p12 -name <meaningful name>
```

```
Enter pass phrase for key.pem:<if password added in pre-  
vious step>  
Enter Export Password:  
Verifying - Enter Export Password:
```

Important We recommend that you provide the `name` argument with a meaningful name. You then have to use this name as the keystore alias in the [JVM configuration](#) of the Jobserver service. If you don't use the `name` argument and there's only one keystore, then the keystore alias is 1.

5. Copy the p12 file to %collibra_data%/spark-jobserver/security/.

Working with a truststore

1. Find the `fast_context_jvm.conf` file.

You can use `find / -name fast_context_jvm.conf`.

This file can, for example, be located in `/opt/collibra/spark-jobserver/conf/fast_context_jvm.conf`.

2. Edit the `fast_context_jvm.conf` file.

For example: `vi /opt/collibra/spark-jobserver/conf/fast_context_jvm.conf`

3. Add the below truststore properties to `fast_context_jvm.conf`

- `-Djavax.net.ssl.trustStore=<fullpath where the cacerts file is located>`
- `-Djavax.net.ssl.trustStorePassword=XXXXXX`

4. Restart Jobserver or the complete environment to activate the changes.

You can also check the [troubleshooting article](#) on the Support portal.

Connection from Colibra Data Intelligence Cloud to an on-premises Jobserver


In a typical on-premises installation, you install all the services in the same network. To enable the communication between environment services that are installed in different networks, you can use a reverse proxy server between those networks, for example, when you use Colibra as a cloud service and an on-premises Jobserver.

Note If you are using a Colibra Data Intelligence Cloud environment with an on-premises Jobserver, the Jobserver version must be **compatible** with the cloud version. You can find the version of your Colibra Data Intelligence Cloud environment at the bottom of the sign-in window, for example 2024.01.0.

Mutual authentication with a reverse proxy server

Mutual authentication is a means to ensure secure communications between the Data Governance Center service and the Jobserver service. If there is a reverse proxy server between these two services, you can still use mutual authentication.

In the following schema, you can see how the communication is established for mutual authentication between the two services with a reverse proxy server in between.



Step	Description
1	The DGC service initiates the communication to the proxy server by sending a Hello message.
2	The proxy server sends its public key to the DGC service. The DGC service can then authenticate the proxy server.
3	The DGC service validates the received public key with the proxy server's CA certificate. This means that the DGC service has the proxy server's CA certificate in its truststore.
4	The DGC service sends its public key to the proxy server. The proxy server can then authenticate the DGC service.
5	The proxy server validates the received public key with the DGC CA certificate. This means that the proxy server has the DGC CA certificate in its truststore.
6	The proxy server initiates the communication to the Jobserver service by sending a Hello message.
7	The Jobserver service sends its public key to the proxy server. The proxy server can then authenticate the Jobserver service.
8	The proxy server validates the received public key with the Jobserver service's CA certificate. This means that the proxy server has the Jobserver service's CA certificate in its truststore.
9	The proxy server sends its public key to the Jobserver service. The Jobserver service can then authenticate the client side of the reverse proxy server.
10	The Jobserver service validates the received public key with the proxy server's CA certificate. This means that the Jobserver service has the proxy server CA certificate in its truststore.

Set up mutual authentication with a proxy server

If you use [mutual authentication with a proxy server](#) between the Data Governance Center service and the Jobserver service, the configuration of both services is slightly different, especially with the certificates.

Prerequisites

- Your certificates meet the necessary [specifications](#). If not, [convert](#) them to the right format.

Edit the Jobserver service settings

Execute the following steps in Collibra Console that manages the Jobserver service.

1. Open Collibra Console with a user profile that has the **SUPER** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the tab pane, click the **Jobserver** service of a Collibra environment.
The details of the **Jobserver** service are shown.
3. Click **Infrastructure Configuration**.
4. Enter the security configuration:

Setting	Description
Authentication level	The authentication level to communicate with the Jobserver. The client (reverse proxy server) must be configured according the here selected authentication level.
Server certificate chain	The certificate or certificate chain with the public key that is offered by the Jobserver to the reverse proxy server.
Server private key	The private key that is part of the Jobserver's certificate or certificate chain.
Trusted client CA certificate	The certificate of the trusted CA used to validate the client certificate (reverse proxy server). To restrict authentication to this client, the CA should be exclusively used by this server.

5. Click **Save all**.

Add a proxy server to the DGC service

Execute the following steps in Colibra Console that manages your DGC service.

1. Open the DGC service settings for editing:
2. In the **Jobserver** section, click **Add**.
3. Enter the necessary information:

Setting	Description
Jobserver list	The list of registered Jobserver instances.
Name	The name of the Jobserver as it will appear when you register a data source in Data Catalog.
Protocol	The protocol that is used for the communication between the Data Governance Center service and the reverse proxy server. It is recommended to use HTTPS, especially if the services are hosted in different network segments.
Address	The address (IP address, URL, hostname) of the reverse proxy server.
Trusted server CA certificate	The certificate of the trusted CA needed to validate the server certificate. If blank, the default truststore will be used. The default truststore is defined in the SSL configuration section of the DGC service. The CA certificate of the server party (reverse proxy server).
Client certificate	The client certificate offered by the DGC service to the server. If blank, you cannot select mutual authentication as the Jobserver service authentication level.
Client private key	The private key of the DGC service's certificate.
Table profiling data size	The approximate maximum disk size of the data in MB that will be used to profile a table. The value cannot exceed 10,000.
Test connection timeout	This timeout is a time limit (in seconds) after which the connection test is stopped and a timeout error is shown. The default value is 60 seconds.

4. Click **Save all**.

You can still add multiple Jobserver services but then you will need one reverse proxy server per Jobserver. In the unlikely event that there are multiple Jobservers behind one reverse proxy server, you have to configure the reverse proxy server in such a way that there is a unique port per Jobserver.

Configuring the proxy server

Consult the documentation of your reverse proxy server to configure the server side and client side.

Jobserver best practices

This article is specific for an on-premises Jobserver, and the corresponding Jobserver, that connect with Collibra Data Intelligence Cloud.

Installation

- Ensure that you use the latest available installer version that matches your Collibra environment. Collibra is deployed on a monthly basis while the on-premises installer versions are only available on a quarterly basis.

Example Collibra Data Intelligence Cloud is released on a monthly basis while the on-premises environments can only be upgraded on a quarterly basis. For example, Collibra 2020.11 has a corresponding on-premises version 5.7.7. This on-premises version will remain the latest available one for the next two monthly Collibra Data Intelligence Cloud releases.

- A user with the correct permissions to execute the installation:
 - Linux: Ideally, a root or sudo user should be used. The user must be able to execute the installation script on the mounted file system.
 - Windows: An administrator user must be used. This user must have full rights on the intended installation drive and directories.
- Ensure that all required ports are open/listening and not in use by other programs or processes on the new Jobserver.

For the list of default TCP ports, see Overview default ports in Collibra.

- Ensure that you can connect to Collibra Console that is used to manage the on-premises Jobserver.
- Ensure that the server that will run the Jobserver meets the [minimum system requirements](#).
- Ensure that the server that will run the Jobserver can communicate with Collibra Data Intelligence Cloud. For more information, see how you can check the [communication](#).

Configuration

- Ensure that all [JVM properties](#) are entered correctly into Collibra Console.
 - Ensure that there are no spaces at the beginning or end of each JVM property added. These empty spaces will cause the Jobserver to enter a Failed state and prevent the creation of the **spark-jobserver.log** file.
 - Ensure the name specified in the **proxy.url** JVM parameter matches the name of the Jobserver defined in Collibra Console under section 18.a and 18.c.

Note This value is case-sensitive, so the name should match between both Collibra Console for Collibra Data Intelligence Cloud entries and Collibra Console for on-premises Jobserver entries.

- On Linux, the Jobserver may fail to start if the **keystore.p12** file is not owned by the "Collibra" user and group. You can verify this by looking at the last entry in the **spark-jobserver.log** once the Jobserver is in error state.

If this is the issue, you can fix it as follows:

- a. On the server, go to the directory where the p12 file is stored.
- b. Run the following command:

```
chown collibra:collibra keystore.p12
```

- c. Restart the Jobserver.

Jobserver usage

- Ensure that all drivers are configured properly in accordance to the related documentation. This includes applying all of the required properties, based on the Data Source and Driver type:
 - Manage Collibra-provided JDBC drivers
 - Manage your own JDBC drivers
- Ensure that your on-premises Jobserver can communicate with the servers that host your data sources. Please work with your Network/Server teams to resolve any network blockages or restrictions that would prevent the Jobserver from successfully connecting to the intended data sources.

Use the Jobserver behind a reverse proxy

See [Connection from an on-premises Jobserver to a Collibra Data Intelligence Cloud](#).

If you have configured everything correctly, you will see a notice on the Jobserver:

```
INFO [I/O dispatcher 1] reversehttp.gateway.PollingController -
proxy -> no requests polled (204)
```

In the logs of the Data Governance Center service, you will see the following message stream:

```
[http-nio-0.0.0.0-4400-exec-12] INFO c.c.r.-
proxy.ReverseHttpProxyServlet - client -> eea3e5fc-3ed4-4095-
beb8-f365cd984973: GET /processes
[ajp-nio-0.0.0.0-8080-exec-7] INFO c.c.r.-
proxy.ReverseHttpProxyServlet - gw <- gw polled eea3e5fc-3ed4-
4095-beb8-f365cd984973
[ajp-nio-0.0.0.0-8080-exec-2] INFO c.c.r.-
proxy.ReverseHttpProxyServlet - gw -> respond eea3e5fc-3ed4-
4095-beb8-f365cd984973
[http-nio-0.0.0.0-4400-exec-12] INFO c.c.r.-
proxy.ReverseHttpProxyServlet - client <- eea3e5fc-3ed4-4095-
beb8-f365cd984973: 200
```

Optional Jobserver service configurations

See [Jobserver to Colibra communication](#).

- [Dpolling.backoff](#)
- [Dmax.connections.route](#)
- [Dmax.connections.total](#)
- [Didle.connection.timeout](#)
- [Dhttp.proxy.host](#)
- [Dhttp.proxy.port](#)

Service infrastructure configurations

When you install the Colibra Data Intelligence Cloud environment, you can define some basic settings such as TCP ports and passwords. All other settings of the services will have default values.

Most of these default settings however, may not comply with your company's IT infrastructure. In this section, you will find more information about configuring the different service's infrastructure settings.

In this chapter

Edit the DGC service infrastructure settings	157
Edit the Search service infrastructure settings	162

Edit the DGC service infrastructure settings

The Data Governance Center service infrastructure settings allow you to edit the behavior of the connections to this service and the Java Virtual Machine (JVM) parameters.

1. Open Collibra Console with a user profile that has the **SUPER** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the tab pane, click **Data Governance Center** service of a Collibra DGC environment.
 - » The details of the **Data Governance Center** service are shown.
3. Click **Infrastructure Configuration**.
4. Click **Edit configuration**.
5. Edit the [DGC service infrastructure configuration](#).
6. Click **Save all**.

Tip

- You can navigate to a specific section by clicking it in the tab pane.
- When you edit certain fields, the **i** icon is displayed next to it. When you click it, it displays the default value for that field and a **Reset** button.
- If you have to restart Collibra or execute extra actions to apply the new settings, it is indicated in the user interface.

DGC infrastructure configuration options

To edit the DGC service infrastructure configuration options, you need the [SUPER](#) role.

Application server configuration

HTTP connector

A connector supporting the HTTP/1.1 protocol.

Setting	Description
URI encoding (Requires restart) *	The character encoding used to decode the URI bytes, URL-decoding has been performed.
Acceptor thread count (Requires restart) *	The number of threads to be used to accept connections. The actual request processing is done by separate threads, so you would never really need more than two.
Min spare threads (Requires restart) *	The minimum number of threads always kept running.
Max threads (Requires restart) *	The maximum number of request processing threads to be created by this connector. This determines the maximum number of simultaneous requests that can be handled.
Accept count (Requires restart) *	The maximum queue length for incoming connection requests that can be assigned to a request processing thread. When the queue is full, requests will be refused.
Compression (Requires restart) *	The connector may use HTTP/1.1 GZIP compression. • off: No compression • on: Allow compression • force: Allow compress
Compression min size (Requires restart) *	The minimum amount of data before it is compressed. This setting only has effect if Compression is on .
Compressible MIME type (Requires restart)	A comma separated list of MIME types allowing HTTP compression.
Connection timeout (Requires restart) *	The number of milliseconds that this connector waits for the destination URI to be presented by the request, after accepting a connection.

Setting	Description
Port (Requires restart) *	The TCP port to access your Collibra environment via your web browser.

HTTPS connector

A connector supporting the HTTP/1.1 protocol with SSL support enabled.

Setting	Description
URI encoding (Requires restart) *	The character encoding used to decode the URI bytes, URL-decoding has been performed.
Acceptor thread count (Requires restart) *	The number of threads to be used to accept connections. The actual request processing is done by separate threads, so you would never really need more than two.
Min spare threads (Requires restart) *	The minimum number of threads always kept running.
Max threads (Requires restart) *	The maximum number of request processing threads to be created by this connector. This determines the maximum number of simultaneous requests that can be handled.
Accept count (Requires restart) *	The maximum queue length for incoming connection requests that can be assigned to a request processing thread. When the queue is full, requests will be refused.
Compression (Requires restart) *	The connector may use HTTP/1.1 GZIP compression. • off: No compression • on: Allow compression • force: Allow compress

Setting	Description
Compression min size (Requires restart) *	The minimum amount of data before it is compressed. This setting only has effect if Compression is on.
Compressible MIME type (Requires restart)	A comma separated list of MIME types allowing HTTP compression.
Connection timeout (Requires restart) *	The number of milliseconds that this connector waits for the destination URI to be presented by the request, after accepting a connection.
Port (Requires restart) *	The TCP port to access your Collibra environment via your web browser.

AJP connector

A connector able to communicate with another web connector via the AJP protocol. Mainly for transparent integration with another HTTP server, for example Apache, nginx....

Setting	Description
URI encoding (Requires restart) *	The character encoding used to decode the URI bytes, URL-decoding has been performed.
Acceptor thread count (Requires restart) *	The number of threads to be used to accept connections. The actual request processing is done by separate threads, so you would never really need more than two.
Min spare threads (Requires restart) *	The minimum number of threads always kept running.

Setting	Description
Max threads (Requires restart) *	The maximum number of request processing threads to be created by this connector. This determines the maximum number of simultaneous requests that can be handled.
Accept count (Requires restart) *	The maximum queue length for incoming connection requests that can be assigned to a request processing thread. When the queue is full, requests will be refused.
Compression (Requires restart) *	The connector may use HTTP/1.1 GZIP compression. • off: No compression • on: Allow compression • force: Allow compress
Compression min size (Requires restart) *	The minimum amount of data before it is compressed. This setting only has effect if Compression is on .
Compressible MIME type (Requires restart)	A comma separated list of MIME types allowing HTTP compression.
Connection timeout (Requires restart) *	The number of milliseconds that this connector waits for the destination URI to be presented by the request, after accepting a connection.
Port (Requires restart) *	The TCP port to access your Collibra environment via your web browser.

Static Resources

The static resources are the resources that are reserved for the service.

Setting	Description
Maximum cache size (Requires restart) *	The maximum size that can be assigned to the cache of the service.

JVM configuration

The Java Virtual Machine configuration parameters to run the Data Governance Center service.

Warning

Only apply changes to these settings if you are really experienced with JVM parameters. Changing parameters may cause serious performance issues.

To remove an individual JVM property you must use the delete icon (🗑) alongside the property, otherwise, the service will interpret it as a blank line and fail to start correctly.

Restart the service after editing the JVM parameters.

Edit the Search service infrastructure settings

The Search service infrastructure settings allow you to edit the Java Virtual Machine (JVM) parameters.

1. Open Collibra Console with a user profile that has the **SUPER** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the tab pane, click **Search** service of a Collibra DGC environment.
 - » The details of the **Search** service are shown.
3. Click **Infrastructure Configuration**.
4. Click **Edit configuration**.
5. Edit the [Search service infrastructure configuration](#).
6. Click **Save all**.

Search infrastructure configuration options

To edit the Search infrastructure configuration options, you need the [SUPER role](#).

JVM configuration

The Java Virtual Machine configuration parameters to run the Search service.

Warning

Only apply changes to these settings if you are really experienced with JVM parameters. Changing parameters may cause serious performance issues.

To remove an individual JVM property you must use the delete icon () alongside the property, otherwise, the service will interpret it as a blank line and fail to start correctly.

Restart the service after editing the JVM parameters.

Managing environments

A Collibra Data Intelligence Cloud environment is a collection of services that allows you to govern your data. With Collibra Console, you can manage one or more environments. You can start and stop an environment, configure the environment behavior and so on.

Each environment must have Data Governance Center service, a Repository service and a Search service. You can install each service on a dedicated node or combine services on one node.

- Data Governance Center service: This is the service that contains the business logic of your environment.
- Repository service: This is the service that stores the data in the DGC database.
- Search service: This is the service that allows you to search for data in the Collibra user interface.

For more information about the services, take a look at the [Product architecture](#) section.

In this chapter

Create a Collibra environment	165
Environment statuses	166
Start an environment	167
Stop an environment	168
Delete an environment	168
Start a service	169
Stop a service	169
Reindex Collibra Data Intelligence Cloud	170
Restore to factory defaults	170

Create a Colibra environment

A Colibra Data Intelligence Cloud environment contains all the services that makes the Colibra platform operational for end-users. You can create multiple environments for different purposes, for example a development and production environment, but keep in mind that a service is always dedicated to one environment.

Note With Colibra Console, you can manage many nodes, these are the services on which you have installed the Colibra services. These nodes must be installed with the same installer version as your Colibra Console.


Steps

1. Open Colibra Console with a user profile that has the **SUPER** role.
 - » Colibra Console opens with the **Infrastructure** page.

Tip

- The default address to access Colibra Console is `<server hostname>:4402`, but you may have set another port during the installation of Colibra Console. Keep in mind that a firewall of your operating system can block the access to Colibra Console.
- The default credentials to sign in to Colibra Console are *Admin / admin*. We highly recommend that you [edit](#) the Colibra Console administrator's password after signing in for the first time.

2. In the tab pane, click **Add / Create**.
 - » The **Add / Create** dialog box appears.
3. Click **Create environment**.
 - » The **Create Environment** dialog box appears.
4. Enter a name.
5. Click **Create Environment**.



6. Perform one of the following steps:

- Click **Close** to end the wizard.
- Click **Add services to this environment** to immediately [add services](#).


Note If the node that hosts the service you want to add is not yet available in Collibra Console, click **Add services from a new node** under the drop-down list and [add](#) the node details.

Environment statuses

An environment can have the following statuses:

- **starting** (transitional status)
- **running**
- **stopping** (transitional status)
- **stopped**
- **error**
- **restoring** (transitional status)

In the following diagram you can find the relation between the statuses:



Start an environment

Starting an environment will start all services of the environment. You can also start the services individually, for more information, see [Start a service](#).

Note When you use a network service account to start the agent and console services, the account must be available when the node starts. If the account is not available, the startup will fail.

Steps

To start an environment, follow these steps:

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. Click the name of an environment to show its details.
3. Click ► **Start**.

What's next?

The environment first starts the Repository service and optionally the Jobserver service, then the Collibra service. When both services have the status **running**, the environment status becomes **running**.

Stop an environment

Stopping an environment will stop all services of the environment.

Tip You can also stop individual services. More information: [Stop a service](#).

Steps

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. Click the environment that you want to stop.
3. Click ■ **Stop**.
 - » The **Stop environment** dialog box appears.
4. Click **Stop environment**.

What's next?

The environment first stops the Collibra service and then optionally the Jobserver service and finally the Repository service.

Delete an environment

To delete an environment, follow these steps:

1. Open Collibra Console with a user profile that has the **SUPER** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the list of environments, click  at the end of the line.

The **Delete Environment** dialog box appears
3. Click **Delete**.

Start a service

By **starting** an environment, you start all services at once. Instead of starting the environment, you can also start the services manually.

Steps

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. Click the environment whose service you want to start.
3. Click the service that you want to start.
4. Click ► **Start**.

Stop a service

Instead of **stopping** an environment, you can also stop a service manually. You can only stop services if no other services depend on it. For example, you cannot stop the Jobserver service if the Data Governance Center service is still running.

Steps

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the infrastructure tree, click the service that you want to stop.
 - » The service details appear.
3. Click ■ **Stop**.
 - » The **Stop service** dialog box appears.
4. Click **Stop service**.

Reindex Collibra Data Intelligence Cloud

After you edit any hyperlink or search settings, you need to reindex Collibra Data Intelligence Cloud so that the existing hyperlinks and search index are updated.

Reindexing Collibra rebuilds automatic hyperlinks and the search index. This topic describes how to reindex Collibra from Collibra Console.

Tip You can also [rebuild](#) the search index from the Collibra settings page.

Steps

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the tab pane, click the **Data Governance Center** service of the environment whose hyperlinks and search index you want to reindex.
3. Click **Rebuild search index and automatic hyperlinks**.

Restore to factory defaults

If you no longer need the data from your environment, you can reset your environment to factory defaults. This action completely resets your repository and removes all the customizations of the Data Governance Center service, but it will keep the configuration of your environment.

Warning This operation cannot be undone, it will delete permanently the content of the repository.

Steps

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. Click the name of an environment to show its details.
3. Click **More** → **Restore to factory defaults**.
4. Click **Restore** to confirm that you want to restore the environment.

5. Optionally click **Continue in background** to continue working in Collibra Console.
 - » The environment is reset. The DGC service will restart during this process.

Monitoring your Collibra Cloud Self-Hosted environment

If you are running Collibra Cloud Self-Hosted environments, we cannot monitor their health like we can with cloud environments. However, you can set up your own monitoring systems to ingest data, alert, and visualize the usage of your environment.

In this section, we describe what and how you can look up the monitoring data.

System metrics	173
Importing logs in an observability tool	174
Environment health check	176



System metrics

You can monitor the following system metrics in your environment:

- CPU usage
- Memory, used and available
- File system, used and available

Tip We recommend setting an alert when file system usage is above 80%.

If you are using services or daemons, we recommend monitoring the health of the `collibra-agent` and `collibra-console` services.

The following is an example configuration for [OpenTelemetry collector](#) to collect basic system metrics. This may be different for the observability tool that you are using.

```
receivers:
  # See more details: https://github.com/open-telemetry/opentelemetry-collector-contrib/tree/main/receiver/hostmetricsreceiver
  hostmetrics:
    scrapers:
      cpu:
      memory:
      disk:
      load:
      paging:
      processes:
      network:
      filesystem:

processors:
  # See more details: https://github.com/open-telemetry/opentelemetry-collector-contrib/tree/main/processor/batchprocessor
  batch:

  # See more details: https://github.com/open-telemetry/opentelemetry-collector-contrib/tree/main/processor/resourcedetectionprocessor
  resourcedetection:
    detectors:
      - env
```

```

- system

exporters:
  # See more details: https://github.com/open-telemetry/opentelemetry-collector/tree/main/exporter/debugexporter
  debug:

  # Add appropriate exporter(s) here
  # See all available ones: https://github.com/open-telemetry/opentelemetry-collector-contrib/tree/main/exporter

service:
  pipelines:
    metrics:
      receivers:
        - hostmetrics
      processors:
        - resourcedetection
        - batch
      exporters:
        - debug
        # Add appropriate exporter(s) here

```

Importing logs in an observability tool

You can [open log files](#) of all Collibra services in Collibra Console. Instead of analyzing log files in Collibra Console, you can also collect them in the observability tool of your choice. To know where to find the log files and what the content is, go to the [Platform configuration section](#).

The path to the log files is typically `/path/to/collibra_data/<service name>/logs`.

The following is an example configuration for [OpenTelemetry collector](#) to collect the logging of the Data Governance Center service. This may be different for the observability tool that you are using.

```

receivers:
  # See more details: https://github.com/open-telemetry/opentelemetry-collector-contrib/tree/main/receiver/filelogreceiver
  filelog/collibra_dgc:
    include:

```

```

    - ${env:COLLIBRA_DATA_PATH}/dgc/logs/dgc.log
  resource:
    service.name: collibra-dgc
  multiline:
    line_start_pattern: \d{4}-\d{2}-\d{2}\s\d{2}:\d{2}:\d
{2}.\d{3}

  # Add more filelog receivers for each log that you want to
  collect

processors:
  # See more details: https://github.com/open-tele-
  metry/opentelemetry-collector/tree/main/processor/batchprocessor
  batch:

  # See more details: https://github.com/open-tele-
  metry/opentelemetry-collector-contrib/tree/main/processor/resourcedetectionprocessor
  resourcedetection:
    detectors:
      - env
      - system

exporters:
  # See more details: https://github.com/open-tele-
  metry/opentelemetry-collector/tree/main/exporter/debugexporter
  debug:

  # Add appropriate exporter(s) here
  # See all available ones: https://github.com/open-tele-
  metry/opentelemetry-collector-contrib/tree/main/exporter

service:
  pipelines:
    logs:
      receivers:
        - filelog/collibra_dgc
        # Add more filelog receivers for each log that you want
to collect
      processors:
        - resourcedetection
        - batch
      exporters:
        - logging
        # Add appropriate exporter(s) here

```

Environment health check

You can see the health of your [environment infrastructure](#) in Collibra Console indicated by a [colored disc](#) in front of the infrastructure element.

We also provide a REST API to monitor all the infrastructure elements.

Infra-structure element	API documentation	Example response
Environment	<a href="https://<your-console-url>/docs/rest/index.html#/environment/findAll_1">https://<your-console-url>/docs/rest/index.html#/environment/findAll_1	Example - Environment details <a href="https://<your-console-url>/rest/environment/">https://<your-console-url>/rest/environment/ <pre>[{ "createdAtDate":1699578732.344000-000, "modifiedDate":1699578732.34400-0000, "id":"<uuid of environment>", "name":"Default environment", "serviceIdSet":["<uuid service 1>","<uuid service 2>","<uuid service 3>","<uuid service 4>"], "status":"RUNNING" }]</pre>

Infra-structure element	API documentation	Example response
Service	https://<your-console-url>/-docs/rest/index.html#managed-services/findAll_2	Example - Services details https://<your-console-url>/rest/service <pre>[{ "createdAtDate":1699578732.134000-000, "modifiedDate":1699578732.13400-0000, "id":"<uuid of service", "nodeId":"<uuid of node>", "status":"RUNNING", "errorMessage":"", "requiredDependencies": ["SEARCH","REPOSITORY"], "optionalDependencies": ["SPARK"], "type":"DGC" }, { "createdAtDate":1699578732.118000-000, "modifiedDate":1699578732.11800-0000, "id":"<uuid of dgc service>", "nodeId":"<uuid of node>", "status":"RUNNING", "errorMessage":"", "requiredDependencies":[], "optionalDependencies":[], "type":"SPARK" }, { "createdAtDate":1699578732.076000-000,</pre>

Infra-structure element	API documentation	Example response
		<pre> "mod- ifiedDate":1699578732.07600- 0000, "id":"<uuid of search service>", "nodeId":"<uuid of node>", "status":"RUNNING", "errorMessage":"", "requiredDependencies":[], "optionalDependencies":[], "type":"SEARCH" }, { "cre- atedDate":1699578732.098000- 000, "mod- ifiedDate":1699578732.09800- 0000, "id":"<uuid of repo service>", "nodeId":"<uuid of node>", "status":"RUNNING", "errorMessage":"", "requiredDependencies":[], "optionalDependencies":[], "type":"REPOSITORY"}] </pre>

Infra-structure element	API documentation	Example response
Node	• <a href="https://<your-console-url>/-doc-s/rest/index.html#/node/findAll_3">https://<your-console-url>/-doc-s/rest/index.html#/node/findAll_3 • <a href="https://<your-console-url>/-doc-s/rest/index.html#/node/status">https://<your-console-url>/-doc-s/rest/index.html#/node/status	• Example 1 - Node details <a href="https://<your-console-url>/rest/node">https://<your-console-url>/rest/node <pre data-bbox="863 510 1426 1305">[{"cre-atedDate":1699578732.031-000000,"mod-ifiedDate":1699578733.67-0000000,"id":"<uuid of node>","hostName":"localhost","port":4401,"name":"Default node","managedServiceIdSet":["<uuid of service 1 on this node>","<uuid of service 2 on this node>","<uuid of service 3 on this node>","<uuid of service 4 on this node>"]}]</pre> • Example 2 - Status of a node <a href="https://<your-console-url>/rest/node/<node uuid>/status">https://<your-console-url>/rest/node/<node uuid>/status <pre data-bbox="863 1476 1426 1599">"UP"</pre>

Node management

A node is a physical server or virtual machine on which you install at least one service of a Collibra Data Intelligence Cloud environment. You can install all services on one node, but keep in mind that Collibra does not support this kind of installation.

We recommend that you have 5 dedicated nodes: one for the Data Governance Center service and Search service, one for Collibra Console, one for the Repository service, one for the Monitoring service, and one for Jobserver if applicable. Make sure that you have a fast network between the nodes.

You can only manage nodes when Collibra is running on multiple nodes.

Important This section is only applicable for Collibra Cloud Self-Hosted 2023.11 or newer or the on-premises Collibra Data Governance Center 5.x.

In this chapter

Add a node to your infrastructure	181
View node details	182
Edit the node server settings	183
Remove a node	188



Add a node to your infrastructure

A node is a physical server that runs one or more services of a Collibra Data Intelligence Cloud environment.

Prerequisites

- The node that you want to add to your infrastructure must be up and running and reachable from the Console that you are using.
- The version of the node must match the version of Collibra Console.

Tip To add a node that was previously registered to another Collibra Console, go to the knowledge base on the [Collibra Support Portal](#).

Steps

1. Open Collibra Console with a user profile that has the **SUPER** role.
 - » Collibra Console opens with the **Infrastructure** page.


Tip

- The default address to access Collibra Console is `<server hostname>:4402`, but you may have set another port during the installation of Collibra Console. Keep in mind that a firewall of your operating system can block the access to Collibra Console.
- The default credentials to sign in to Collibra Console are *Admin / admin*. We highly recommend that you [edit](#) the Collibra Console administrator's password after signing in for the first time.

2. In the tab pane, click **Add / Create**.
 - » The **Add / Create** dialog box appears.
3. Click **Add node**.
 - » The **Add node** dialog box appears.

4. Enter the necessary information.

Field	Description
Node name	Enter a meaningful name for the node.
Hostname	Enter the hostname or IP address of the node, for example <i>192.168.1.100</i> or <i>repository-node-A</i>. If you use a hostname, make sure that the Collibra Console can resolve the hostname. Note Do not reuse hostnames, every hostname must be unique. If you reuse a hostname for a node that will be used in a repository cluster, the cluster won't synchronize.
Port	Enter the agent port. This is the port that you defined during the installation of the Agent on that node. The default value is 4401. The agent port is the port through which Collibra Console connects to the node.

5. Click **Add node**6. Click **Close**.

View node details

To view details of a node, follow these steps:

1. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.

2. Perform one of the following steps:

Strategy	Description
Open a node via an environment	a. Click an environment name to open its details. b. In the Environment Services section, in the Node column, click a node.
Open a node directly	a. In the tab pane, click a node.

» The node details appear.

Edit the node server settings

To edit node server settings, follow these steps:

1. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
2. Click the name of an environment to display its details.
3. Click **Configuration**.
4. Click **Edit configuration**.
5. Edit the settings.

For a complete list of the settings, see [Node server settings](#).
6. Click **Save all**.

Tip

- You can navigate to a specific section by clicking it in the tab pane.
- When you edit certain fields, the **i** icon is displayed next to it. When you click it, it displays the default value for that field and a **Reset** button.
- If you have to restart Collibra or execute extra actions to apply the new settings, it is indicated in the user interface.

Node server settings

Application server configuration

HTTP connector

A connector supporting the HTTP/1.1 protocol.

Setting	Description
URI encoding (Requires restart) *	The character encoding used to decode the URI bytes, URL-decoding has been performed.
Acceptor thread count (Requires restart) *	The number of threads to be used to accept connections. The actual request processing is done by separate threads, so you would never really need more than two.
Min spare threads (Requires restart) *	The minimum number of threads always kept running.
Max threads (Requires restart) *	The maximum number of request processing threads to be created by this connector. This determines the maximum number of simultaneous requests that can be handled.
Accept count (Requires restart) *	The maximum queue length for incoming connection requests that can be assigned to a request processing thread. When the queue is full, requests will be refused.
Compression (Requires restart) *	The connector may use HTTP/1.1 GZIP compression. • off: No compression • on: Allow compression • force: Allow compress
Compression min size (Requires restart) *	The minimum amount of data before it is compressed. This setting only has effect if Compression is on .

Setting	Description
Compressible MIME type (Requires restart)	A comma separated list of MIME types allowing HTTP compression.
Connection timeout (Requires restart) *	The number of milliseconds that this connector waits for the destination URI to be presented by the request, after accepting a connection.
Port (Requires restart) *	The TCP port to access your Collibra environment via your web browser.

HTTPS connector

A connector supporting the HTTP/1.1 protocol with SSL support enabled.

Setting	Description
URI encoding (Requires restart) *	The character encoding used to decode the URI bytes, URL-decoding has been performed.
Acceptor thread count (Requires restart) *	The number of threads to be used to accept connections. The actual request processing is done by separate threads, so you would never really need more than two.
Min spare threads (Requires restart) *	The minimum number of threads always kept running.
Max threads (Requires restart) *	The maximum number of request processing threads to be created by this connector. This determines the maximum number of simultaneous requests that can be handled.
Accept count (Requires restart) *	The maximum queue length for incoming connection requests that can be assigned to a request processing thread. When the queue is full, requests will be refused.

Setting	Description
Compression (Requires restart) *	The connector may use HTTP/1.1 GZIP compression. • off: No compression • on: Allow compression • force: Allow compress
Compression min size (Requires restart) *	The minimum amount of data before it is compressed. This setting only has effect if Compression is on .
Compressible MIME type (Requires restart)	A comma separated list of MIME types allowing HTTP compression.
Connection timeout (Requires restart) *	The number of milliseconds that this connector waits for the destination URI to be presented by the request, after accepting a connection.
Port (Requires restart) *	The TCP port to access your Collibra environment via your web browser.

AJP connector

A connector able to communicate with another web connector via the AJP protocol. Mainly for transparent integration with another HTTP server, for example Apache, nginx....

Setting	Description
URI encoding (Requires restart) *	The character encoding used to decode the URI bytes, URL-decoding has been performed.
Acceptor thread count (Requires restart) *	The number of threads to be used to accept connections. The actual request processing is done by separate threads, so you would never really need more than two.

Setting	Description
Min spare threads (Requires restart) *	The minimum number of threads always kept running.
Max threads (Requires restart) *	The maximum number of request processing threads to be created by this connector. This determines the maximum number of simultaneous requests that can be handled.
Accept count (Requires restart) *	The maximum queue length for incoming connection requests that can be assigned to a request processing thread. When the queue is full, requests will be refused.
Compression (Requires restart) *	The connector may use HTTP/1.1 GZIP compression. • off: No compression • on: Allow compression • force: Allow compress
Compression min size (Requires restart) *	The minimum amount of data before it is compressed. This setting only has effect if Compression is on .
Compressible MIME type (Requires restart)	A comma separated list of MIME types allowing HTTP compression.
Connection timeout (Requires restart) *	The number of milliseconds that this connector waits for the destination URI to be presented by the request, after accepting a connection.
Port (Requires restart) *	The TCP port to access your Collibra environment via your web browser.

Static Resources

The static resources are the resources that are reserved for the service.

Setting	Description
Maximum cache size (Requires restart) *	The maximum size that can be assigned to the cache of the service, expressed in kilobytes.

Remove a node

When you no longer need a node in an environment, you can remove it. Keep in mind that this action does not remove the software from the node, which means that you can reuse it.

You can remove a node from different locations:

- From the node details.
- From the node list.
- From the node list.

Remove a node from the node details

1. Open Collibra Console with a user profile that has the **SUPER** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the tab pane, click the node that you want to remove.
 - » The node details appear.
3. Click  **Delete**.
 - » The **Remove node** dialog box appears.
4. Click **Remove**.

Remove a single node from the node list

1. Open Collibra Console with a user profile that has the **SUPER** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the tab pane, click **NODES**.
 - » The node list appears.
3. Click  in the row of the node that you want to remove.
4. Click **Remove**.

Remove one or more nodes from the node list

1. Open Collibra Console with a user profile that has the **SUPER** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the tab pane, click **NODES**.
 - » The node list appears.
3. Select the checkboxes in front of the nodes you want to remove.
4. Click **Delete**.
 - » The **Remove node** dialog box appears.
5. Click **Remove**.

Back up and restore

A backup corresponds to a picture of a Collibra Data Intelligence Cloud environment at a certain point in time. This picture consists of all application data, the operating model, assets, comments, and so on. You can restore a backup at a later point in time.

You can specify which aspects of the application data to include in the backup and again during the restore.

In this chapter

Create a backup of Collibra Data Intelligence Cloud 2024.01	192
Backup options	194
The Backups page	196
Download a backup	197
Upload a backup	198
Delete a backup	199
Create a backup schedule	199
Backup schedule options	202
Backup schedules overview	203
Edit a backup schedule	205
Delete a backup schedule	205
Restore a backup	206
Customize Collibra DGC with a backup restore	209
Back up and download with the REST API	210
REST API - List of backups	220

REST API - Delete a backup	221

Create a backup of Collibra Data Intelligence Cloud 2024.01

If you want to take a full or partial [backup](#) of a Collibra Data Intelligence Cloud environment, follow these steps:

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the main menu, click **Backups**.
 - » The [backups page](#) appears.
3. In the upper right corner, click **Create backup**.

Note If you have one or more Edge sites in this Collibra environment, you need to create a [backup](#) for each Edge site.

4. Enter the required information.

Option	Description
Environment	The Collibra Data Intelligence Cloud environment that you want to back up.
Include	You can create: ◦ a full backup: Includes all backup options. ◦ a partial backup: Includes only a couple of the backup options.
Data	The core (database) data contained in the Collibra environment. This is all the data related to communities, domains, assets, users, operating model and so on.

Option	Description
History	The history of the data (activity and snapshot-related). This can result in a very large file, depending on the activity in Collibra. If you create a backup of the data and exclude the data history, you are taking a backup of a subset of the database content. This results in a much smaller backup file, but it still contains a complete picture of the current content. This is useful when attaching a backup to a technical support ticket, if support has to reproduce a problem. Warning If you restore a backup that does not include history and was made from a version older than 5.7.x, the data quality metrics will not work correctly. To solve this issue, after you restore the backup in your latest environment, upgrade the data, create a new backup and restore the new backup in the same environment.
Configuration	The configuration of the Collibra environment.
Customizations	The customizations such as custom modules, styling, page-definitions and email template files. This will add files that are located in the Collibra user home directory and files that are used to customize the Collibra environment (/opt/collibra_data/console/config). This concerns all the files located in the following directories: ◦ email-templates ◦ groovy-lib ◦ images ◦ page-definitions ◦ security ◦ styling ◦ translations ◦ modules (if available)

Option	Description
Name	The name of the backup. You can only use the following characters in the backup name: ◦ letters ◦ numbers ◦ hyphen (-) ◦ underscore (_) ◦ period (.)
Description	The description to provide more information on the backup.

5. Click **Create backup**.

What's next?

A dialog box with a progress bar is displayed. You can let the backup continue in the background. If you do so, you can see the progress of the backup on the [backups page](#).

Backup options

Option	Description
Environment	The Collibra Data Intelligence Cloud environment that you want to back up.
Include	You can create: • a full backup: Includes all backup options. • a partial backup: Includes only a couple of the backup options.
Data	The core (database) data contained in the Collibra environment. This is all the data related to communities, domains, assets, users, operating model and so on.

Option	Description
History	The history of the data (activity and snapshot-related). This can result in a very large file, depending on the activity in Collibra. If you create a backup of the data and exclude the data history, you are taking a backup of a subset of the database content. This results in a much smaller backup file, but it still contains a complete picture of the current content. This is useful when attaching a backup to a technical support ticket, if support has to reproduce a problem. Warning If you restore a backup that does not include history and was made from a version older than 5.7.x, the data quality metrics will not work correctly. To solve this issue, after you restore the backup in your latest environment, upgrade the data, create a new backup and restore the new backup in the same environment.
Configuration	The configuration of the Collibra environment.
Customizations	The customizations such as custom modules, styling, page-definitions and email template files. This will add files that are located in the Collibra user home directory and files that are used to customize the Collibra environment (<code>/opt/collibra_data/console/config</code>). This concerns all the files located in the following directories: • email-templates • groovy-lib • images • page-definitions • security • styling • translations • modules (if available)
Name	The name of the backup. You can only use the following characters in the backup name: • letters • numbers • hyphen (-) • underscore (_) • period (.)
Description	The description to provide more information on the backup.

The Backups page


The **Backups** page contains a table of available backups. In the first column, you can find the name and date of the backup and who created it. In the second column, you can find the content that is included in the backup.

Asset View Column	Description
Name	The name of the backup. You can edit the name of the backup in this table.
Description	The description that was entered when the backup was created. You can edit the description in this table.
Data	Indication whether the backup contains all current data (export of the underlying database schema).
History	Indication whether the data (if included) also holds the history (activity and snapshot data). Note Your backup can only include history data if it also includes the current data.
Configuration	Indication whether configuration files and license file are included in the backup.
Customizations	Indication whether customization files are included in the backup.
Created on	The date on which the backup was created.
Version	The version of Collibra Data Intelligence Cloud in the backup.
Size	The size of the backup
[empty]	Links to download (⬇), restore (↶) and delete (🗑) the backup.

Download a backup

If you want to restore a backup on another environment, you have to download it. When you download a backup, you can encrypt the backup for security reasons.

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the main menu, click **Backups**.
 - » The **backups page** appears.
3. In the row of the backup that you want to download, click **Download**.
 - » The **Download** dialog box appears.
4. Type a password to encrypt the backup or leave the **Password** field empty for an unencrypted backup.



Note The password is used to encrypt the backup file, not to password protect it. Encrypted or not, the backup file is a ZIP archive. You can only extract the unencrypted backup files with the usual archive managers. You can not extract an encrypted backup, it can only be used to upload it with Collibra Console.

5. Click **Download**.
 - » The backup is downloaded to your default download folder as a ZIP archive file.

What's next?

To restore the backup on another environment, you have to [upload](#) it to that environment.

Upload a backup

If you are planning to restore a backup that you stored outside of Collibra Data Intelligence Cloud, or if you want to promote content from one Collibra environment to another, follow these steps:

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the main menu, click **Backups**.
 - » The [backups page](#) appears.
3. In the upper-right corner, click the **Upload backup** button.
 - » The **Upload backup** dialog box appears.
4. Enter the required information:

Option	Description
Password protected	The option to choose when the backup was encrypted when downloading it.
Not password protected	The option to choose when the backup was not encrypted when downloading it.
Password	The password to unencrypt the backup when uploading it. This option is required when the Password protected option is selected.

5. Do one of the following:
 - Browse to the location of your backup file with a file explorer, grab the backup file with your mouse pointer and drag it to the dotted box in the **Upload backup** dialog box.
 - Click , browse to the location of your backup file in the window that opens, select the file and click **Open**.
 - » Your backup upload starts automatically.

What's next?


You can now [restore the backup](#) on your environment.

Delete a backup

If you no longer need a certain backup or backups, you can delete them. This will free up disk space.

Steps

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the main menu, click **Backups**.
 - » The [backups page](#) appears.
3. Do one of the following:

To delete only one backup at a time:	Click  at the end of the line of the backup that you want to delete.
To delete multiple backups at once:	Select the check boxes of the backups that you want to delete and click Delete at the top of the table. 

» The **Delete Backups** dialog box appears.

4. Click **Yes** to confirm.

Create a backup schedule

A backup schedule allows you to automatically create backups in a defined schedule. This reduces the risk of forgetting to create backups and possible data loss.

Note The backups that are taken with a backup schedule are full backups, not incremental backups.

Steps

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the main menu, click **Backups**.
 - » The [backups page](#) appears.
3. In the tab pane, click **Backup schedules**.
4. In the upper-right corner, click **Create backup schedule**.
5. Enter the required information:

Option	Description
Environment	The Collibra Data Intelligence Cloud environment that you want to back up.
Include	You can create: ◦ a full backup: Includes all backup options. ◦ a partial backup: Includes only a couple of the backup options.
Data	The core (database) data contained in the Collibra environment. This is all the data related to communities, domains, assets, users, operating model and so on.
History	The history of the data (activity and snapshot-related). This can result in a very large file, depending on the activity in Collibra. If you create a backup of the data and exclude the data history, you are taking a backup of a subset of the database content. This results in a much smaller backup file, but it still contains a complete picture of the current content. This is useful when attaching a backup to a technical support ticket, if support has to reproduce a problem. Warning If you restore a backup that does not include history and was made from a version older than 5.7.x, the data quality metrics will not work correctly. To solve this issue, after you restore the backup in your latest environment, upgrade the data, create a new backup and restore the new backup in the same environment.
Configuration	The configuration of the Collibra environment.

Option	Description
Customizations	The customizations such as custom modules, styling, page-definitions and email template files. This will add files that are located in the Collibra user home directory and files that are used to customize the Collibra environment (/opt/collibra_data/console/config). This concerns all the files located in the following directories: ◦ email-templates ◦ groovy-lib ◦ images ◦ page-definitions ◦ security ◦ styling ◦ translations ◦ modules (if available)
Name	The name of the backup. You can only use the following characters in the backup name: ◦ letters ◦ numbers ◦ hyphen (-) ◦ underscore (_) ◦ period (.)
Description	The description to provide more information on the backup.
Max retained	The number of backups that you want to retain. The default value is 30.
Cron expression	The schedule of the backups. For more information about the CRON syntax, see the CRON appendix. If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.

6. Click **Create backup schedule**.

» The schedule appears in the [schedule table](#).

Tip You can create multiple schedules for one environment. For example, one schedule for monthly backups that are kept for two years and another schedule for weekly schedules that are kept for only a couple of months.

Backup schedule options

Option	Description
Environment	The Collibra Data Intelligence Cloud environment that you want to back up.
Include	You can create: • a full backup: Includes all backup options. • a partial backup: Includes only a couple of the backup options.
Data	The core (database) data contained in the Collibra environment. This is all the data related to communities, domains, assets, users, operating model and so on.
History	The history of the data (activity and snapshot-related). This can result in a very large file, depending on the activity in Collibra. If you create a backup of the data and exclude the data history, you are taking a backup of a subset of the database content. This results in a much smaller backup file, but it still contains a complete picture of the current content. This is useful when attaching a backup to a technical support ticket, if support has to reproduce a problem. Warning If you restore a backup that does not include history and was made from a version older than 5.7.x, the data quality metrics will not work correctly. To solve this issue, after you restore the backup in your latest environment, upgrade the data, create a new backup and restore the new backup in the same environment.
Configuration	The configuration of the Collibra environment.

Option	Description
Customizations	The customizations such as custom modules, styling, page-definitions and email template files. This will add files that are located in the Collibra user home directory and files that are used to customize the Collibra environment (/opt/collibra_data/console/config). This concerns all the files located in the following directories: • email-templates • groovy-lib • images • page-definitions • security • styling • translations • modules (if available)
Name	The name of the backup. You can only use the following characters in the backup name: • letters • numbers • hyphen (-) • underscore (_) • period (.)
Description	The description to provide more information on the backup.
Max retained	The number of backups that you want to retain. The default value is 30.
Cron expression	The schedule of the backups. For more information about the CRON syntax, see the CRON appendix. If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.

Backup schedules overview

The **Backup schedules** page contains a table of the backup schedules.

Note The backups that are taken with a backup schedule are full backups, not incremental backups.

Backup schedules column	Description
Name	The name of the backup schedule. You can edit the name of the schedule in this table.
Description	The description of the backup schedule. You can edit the description in this cell.
Data	Indication whether the backup includes the Collibra environment data. This is all the data related to communities, domains, assets, users, operating model and so on.
History	Indication whether the data also holds the history. Note Your backup can only include the data history if Data is selected.
Configuration	Indication whether configuration files and license file are included in the backup.
Customizations	Indication whether customization files are included in the backup.
Max retained	The number of backups that are kept on disk. When reaching this number, a new backup will still be created, but the oldest backup will be deleted.
Environment	The name of the environment that will be backed up.
Cron expression	The schedule when the backups will be taken. For more information about the Cron syntax, see cron syntax . If your Cron expression could affect the performance of your environment, for example multiple times an hour, an exclamation mark will be shown. If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.
Next execution	The date and time when the next backup will be created.
[empty]	Button to delete (🗑) the backup.

Edit a backup schedule

When you have created a backup schedule, you can always change it afterwards, for example to change the name or schedule. For more information about the different columns, see [Backup schedules overview](#).

Steps

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the main menu, click **Backups**.
 - » The [backups page](#) appears.
3. In the tab pane, click **Backup schedules**.
4. In the schedule overview table, double-click any value in the table and edit it.

Note You cannot change the **Environment** or **Next execution** columns.

5. Click ✓ to save the change.

Tip If you edit the maximum number of retained backups to a number that is less than the number of currently retained backups, the first cleanup process will clean all necessary backups so that only the specified number of retained backups will remain.

Delete a backup schedule

When you no longer need one or more backup schedules, you can always delete them.

Steps

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the main menu, click **Backups**.
 - » The [backups page](#) appears.
3. In the tab pane, click **Backup schedules**.

4. In the schedule overview table, click  next to the schedule that you want to delete.
To delete multiple schedules at once, select them in the first column and click **Delete**.
» The **Delete backups** dialog box appears.
5. Click **Delete**.

Restore a backup

In this section, we describe how to restore a backup that was created in the same Collibra Data Intelligence Cloud version as your current version. You can restore backups to recover from data corruption or to transfer your data to another environment.

Warning

- If you restore a backup, you will lose all current Collibra data.
- You cannot restore a backup that is created from a newer Collibra version, on an older Collibra environment. For example, if you have a backup that is created on a 5.9 environment, you cannot restore this backup on 5.8 or older. You can, however, [restore](#) an older version in a newer environment, for example a 5.8 backup in an 5.9 environment.
- Backups from versions older than 5.5.x do not include certificates. You have to reinstall any necessary certificates, for example for secure communication with the Jobserver.
- If you restore a backup that does not include history and was made from a version older than 5.7.x, the [data quality](#) metrics will not work correctly. To solve this issue, after you restore the backup in your latest environment, upgrade the data, create a new backup and restore the new backup in the same environment.

Important If the Collibra environment you want to restore has one or more Edge sites, you need to reinstall all of them. To reinstall one or more Edge sites, go to [Reinstall an Edge site](#).

Steps

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
» Collibra Console opens with the **Infrastructure** page.

2. In the main menu, click **Backups**.
 - » The [backups page](#) appears.
3. Click  **Restore** in the row of the backup that you want to restore.
4. Click the name of an environment to show its details.
5. Click **More** → **Restore from backup**
6. Fill in the **Restore from backup** form:

Option	Description
Environment	Select the environment in which you want to restore the backup.
Backup	Select the backup that you want to restore.
Include	You can restore a backup fully or partially. Note The backup itself should include the items you select below. For example, if the backup does not contain the configuration and you include configuration for restoring the backup, it will be restored successfully, though without the configuration.
Data	Select to include the core (database) data from the backup. This is all the data related to communities, domains, assets, users, operating model and so on.
History	Select to include all historic data (activity and snapshot-related) from the backup.
Configuration	Select to include the configuration and license from the backup. Warning Do not select this option if you restore a backup from another environment, as this will overwrite the configuration of your current environment. It can prevent your environment from starting up again.

Option	Description
Customizations	Select to include custom modules, styling, page-definitions and email template files. This will add files to the Collibra user home directory and files that are used to customize the Collibra DGC instance (<code>/opt/collibra_data/console/config</code>). This concerns all the files in the following directories: ◦ email-templates ◦ groovy-lib ◦ images ◦ page-definitions ◦ security ◦ styling ◦ translations ◦ modules (if available)

7. Click **Restore backup**.

The backup is restored on the selected environment. The time to restore a backup depends heavily on the amount of data that the backup contains. During the restore time, your environment is unavailable.

The restore process is completed when all services of your environment have again the running status.

What's next?

If the configuration file is restored, the database connection information in the backup is ignored. In other words, the restore process makes sure that the entire backed-up configuration is applied to the current configuration without modifying the current database connection. This means that the existing database connection is used to restore the data. Also, after the restore, the Collibra instance continues to use the same database connection.

Note

- You cannot restore a backup from a Collibra environment that is using a database with a case-sensitive collation sequence into a Collibra environment that is using a database with a collation sequence that is not case-sensitive. If you want to use backup and restore to migrate Collibra content, ask your database administrator to determine whether the collation sequences of the databases used by your source and target Collibra environments are compatible.
- If you restore a backup without the history, Collibra does not restore the activity stream as it was. Instead, it creates the entries in the activity stream as if they were done by the System User on the restore date.

Customize Collibra DGC with a backup restore

If you have customized your Collibra Data Intelligence Cloud instance and you want to keep those customizations after an upgrade, you can create a backup and restore the backup to replace the files in the Collibra user home directory.

1. [Create](#) a backup that includes customizations.
2. [Download](#) the backup.
3. Unzip the file in a new directory.
4. Open the new directory.
5. Modify the files.
6. Select all the unzipped folders and files, including **backup.properties** and zip them again.

Note Do not zip the containing directory but the individual files and directories, so that you get the same structure as the original ZIP file.

7. [Upload](#) this backup to your Collibra instance.
8. [Restore](#) the backup, including customizations.

What's next?

Test all customizations.

Back up and download with the REST API

Tip Back up the Collibra Data Intelligence Cloud environments on a regular basis.

In Collibra Console you can [manually back up](#) your environments at any time or [set up a schedule](#) to automatically create backups at a set interval. You can achieve the same using the `/backup` and `/backupSchedule` endpoints of the [CollibraManagement Console REST API](#).

Important Do not include any cookies in your requests. This results in a 403 - Forbidden error.

Create a backup

Before you start, you must know the ID of the environment you want to back up. You can retrieve it through the API or get it from Collibra Console:

Get the ID via the API

Retrieve a list of all environments with a GET method of the `/environment` endpoint:

```
curl -L -X GET 'https://<your_collibra_console_url>/rest/environment' \
-H 'Authorization: Basic <Base64_username:password_encoded_values>'
```

Response example:

```
[
  {
    "createdDate": 1699926678.042000000,
    "modifiedDate": 1699926678.042000000,
    "id": "a8f3ff8a-9bbd-43a2-833d-d54e96d4a798",
    "name": "Default environment",
```

```

        "serviceIdSet": [
            "3e50ee0f-d5dd-4935-9355-ad7da43a3eff",
            "13144bac-a319-45f6-b459-4e856a998488",
            "79c43e61-e51e-4f13-94f4-e27b9e94500f",
            "163342f3-4a7a-4cc8-b0c8-2b15ecfa31fa"
        ],
        "status": "RUNNING"
    }
]

```

Get the ID via Collibra Console

In Collibra Console, go to the **Environments** tab and click the relevant environment to open its details. The environment ID appears in the address bar of your browser, after the colon sign.

Example <https://console.collibra.com/#!/environments/Environment:a8f3ff8a-9bbd-43a2-833d-d54e96d4a798>

To create a backup, use the POST method of the `/backup{environmentId}` endpoint and provide the backup information in a JSON body:

cURL

```

curl -L -X POST 'https://<your_collibra_console_
url>/rest/backup/<environment_id>' \
-H 'Content-Type: application/json' \
-H 'Authorization: Basic <Base64_username:password_encoded_
values>' \
-d '{
    "name": "<your_backup_name>",
    "description": "<your_backup_description>",
    "dgcBackupOptions": [
        "CUSTOMIZATIONS",
        "CONFIGURATION"
    ],
    "repoBackupOptions": [

```

```

        "DATA",
        "HISTORY"
    ]
}'

```

Python - http.client

```

import http.client
import json

conn = http.client.HTTPSConnection("<your_collibra_console_
url>")
payload = json.dumps({
    "name": "<your_backup_name>",
    "description": "<your_backup_description>",
    "dgcBackupOptions": [
        "CUSTOMIZATIONS",
        "CONFIGURATION"
    ],
    "repoBackupOptions": [
        "DATA",
        "HISTORY"
    ]
})
headers = {
    'Content-Type': 'application/json',
    'Authorization': 'Basic <Base64_username:password_encoded_
values>'
}
conn.request("POST", "/rest/backup/<environment_id>",
payload, headers)
res = conn.getresponse()
data = res.read()
print(data.decode("utf-8"))

```

Node - Native

```

var https = require('follow-redirects').https;
var fs = require('fs');

```

```

var options = {
  'method': 'POST',
  'hostname': '<your_collibra_console_url>',
  'path': '/rest/backup/<environment_id>',
  'headers': {
    'Content-Type': 'application/json',
    'Authorization': 'Basic <Base64_username:password_
encoded_values>'
  },
  'maxRedirects': 20
};

var req = https.request(options, function (res) {
  var chunks = [];

  res.on("data", function (chunk) {
    chunks.push(chunk);
  });

  res.on("end", function (chunk) {
    var body = Buffer.concat(chunks);
    console.log(body.toString());
  });

  res.on("error", function (error) {
    console.error(error);
  });
});

var postData = JSON.stringify({
  "name": "<your_backup_name>",
  "description": "<your_backup_description>",
  "dgcBackupOptions": [
    "CUSTOMIZATIONS",
    "CONFIGURATION"
  ],
  "repoBackupOptions": [
    "DATA",
    "HISTORY"
  ]
});

req.write(postData);

req.end();

```

Important

The **name** must not contain special characters.

description is an optional parameter.

The possible values for **dgcBackupOptions** are "CUSTOMIZATIONS" and "CONFIGURATION". If you select none, you must add at least one of the **repoBackupOptions**.

The possible values for **repoBackupOptions** are "DATA" and "HISTORY". If you selected "HISTORY", you must also add "DATA". If you select none, you must add at least one of the **dgcBackupOptions**.

Response example:

```
{
  "createdDate": 1699965108.223000000,
  "modifiedDate": 1699965108.223000000,
  "id": "d8c09773-f56e-49db-bfea-a2305ecbb219",
  "backupInformation": {
    "name": "<your_backup_name>",
    "description": "<your_backup_description>",
    "appVersion": "2023.11",
    "createdByEmail": "<email_address>",
    "date": 1699965108.198000000,
    "backupSpecificationId": "8eaac39e-2114-459b-9ca3-2b66b0ca5fdc",
    "environmentId": "a8f3ff8a-9bbd-43a2-833d-d54e96d4a798",
    "dgcBackupOptions": [
      "CUSTOMIZATIONS",
      "CONFIGURATION"
    ],
    "repoBackupOptions": [
      "HISTORY",
      "DATA"
    ]
  },
  "size": 0,
  "stepStateMap": {
    "REPOSITORY": {
      "status": "WAITING",
      "errorMessage": ""
    },
    "PRE_PROCESSING": {
      "status": "COMPLETED",
      "errorMessage": ""
    }
  }
}
```

```

    },
    "DGC": {
        "status": "WAITING",
        "errorMessage": ""
    },
    "POST_PROCESSING": {
        "status": "WAITING",
        "errorMessage": ""
    }
},
"inProgress": true
}

```

Download a backup

Before you start, you must know the ID of the backup you want to download. You can see the ID in the response you get after creating a backup or you can use the GET method of the `/backup` endpoint to retrieve a list of all the available backups:

cURL

```

curl -L -X GET 'https://<your_collibra_console_
url>/rest/backup' \
-H 'Authorization: Basic <Base64_username:password_encoded_
values>'

```

Python - http.client

```

import http.client

conn = http.client.HTTPSConnection("<your_collibra_console_
url>")
payload = ''
headers = {
    'Authorization': 'Basic <Base64_username:password_encoded_
values>'
}

```

```
conn.request("GET", "/rest/backup", payload, headers)
res = conn.getresponse()
data = res.read()
print(data.decode("utf-8"))
```

Node - Native

```
var https = require('follow-redirects').https;
var fs = require('fs');

var options = {
  'method': 'GET',
  'hostname': '<your_collibra_console_url>',
  'path': '/rest/backup',
  'headers': {
    'Authorization': 'Basic <Base64_username:password_
encoded_values>'
  },
  'maxRedirects': 20
};

var req = https.request(options, function (res) {
  var chunks = [];

  res.on("data", function (chunk) {
    chunks.push(chunk);
  });

  res.on("end", function (chunk) {
    var body = Buffer.concat(chunks);
    console.log(body.toString());
  });

  res.on("error", function (error) {
    console.error(error);
  });
});

req.end();
```

Response example:


```
[
  {
    "createdDate": 1699965108.223000000,
    "modifiedDate": 1699965133.074000000,
    "id": "d8c09773-f56e-49db-bfea-a2305ecbb219",
    "backupInformation": {
      "name": "<your_backup_name>",
      "description": "<your_backup_description>",
      "appVersion": "2023.11",
      "createdByEmail": "<email_address>",
      "date": 1699965108.198000000,
      "backupSpecificationId": "8eaac39e-2114-459b-9ca3-
2b66b0ca5fdc",
      "environmentId": "a8f3ff8a-9bbd-43a2-833d-
d54e96d4a798",
      "dgcBackupOptions": [
        "CUSTOMIZATIONS",
        "CONFIGURATION"
      ],
      "repoBackupOptions": [
        "HISTORY",
        "DATA"
      ]
    },
    "size": 126533796,
    "stepStateMap": {
      "REPOSITORY": {
        "status": "COMPLETED",
        "errorMessage": ""
      },
      "PRE_PROCESSING": {
        "status": "COMPLETED",
        "errorMessage": ""
      },
      "DGC": {
        "status": "COMPLETED",
        "errorMessage": ""
      },
      "POST_PROCESSING": {
        "status": "COMPLETED",
        "errorMessage": ""
      }
    },
    "inProgress": false
  }
]
```

To download a backup, use the POST method of the `/backup/file{id}` endpoint.

Important

- You must include a 'Content-Type: application/x-www-form-urlencoded' header.
- You must add a 'key=<your_password>' data option. This encrypts the download and you must provide the password when you upload the file. If you don't want to password protect the download, leave the password empty.

cURL

```
curl -L -X POST 'https://<your_collibra_console_
url>/rest/backup/file/<backup_id>' \
-H 'Content-Type: application/x-www-form-urlencoded' \
-H 'Authorization: Basic <Base64_username:password_encoded_
values>' \
-d 'key='
-o 'backup.zip'
```

Tip To save the file locally, provide the `-o` option and the name of the file. In this example, the file is saved as **backup.zip** in your current directory.

Python - http.client

```
import http.client

conn = http.client.HTTPSConnection("<your_collibra_console_
url>")
payload = 'key='
headers = {
    'Content-Type': 'application/x-www-form-urlencoded',
    'Authorization': 'Basic <Base64_username:password_encoded_
values>'
}
conn.request("POST", "/rest/backup/file/<backup_id>",
payload, headers)
res = conn.getresponse()
data = res.read()
print(data.decode("utf-8"))
```

Node - Native

```
var https = require('follow-redirects').https;
var fs = require('fs');

var qs = require('querystring');

var options = {
  'method': 'POST',
  'hostname': '<your_collibra_console_url>',
  'path': '/rest/backup/file/<backup_id>',
  'headers': {
    'Content-Type': 'application/x-www-form-urlencoded',
    'Authorization': 'Basic <Base64_username:password_
encoded_values>'
  },
  'maxRedirects': 20
};

var req = https.request(options, function (res) {
  var chunks = [];

  res.on("data", function (chunk) {
    chunks.push(chunk);
  });

  res.on("end", function (chunk) {
    var body = Buffer.concat(chunks);
    console.log(body.toString());
  });

  res.on("error", function (error) {
    console.error(error);
  });
});

var postData = qs.stringify({
  'key': ''
});

req.write(postData);

req.end();
```

REST API - List of backups

To retrieve the list of backups via the REST API:

```
GET https://<your_dgc_console_host>/rest/backup
```

```
[{"createdDate":1496985583.834000000,
  "modifiedDate":1496985583.846000000,
  "id":"c2a3e499-bb0f-4016-85aa-6efcbcd66df3",
  "backupInformation":{"
    "appVersion":"5.1",
    "createdByEmail":"admin@example.com",
    "date":1496843737.891000000,
    "dgcBackupOptionSet":["CUSTOMIZATIONS"],
    "repoBackupOptionSet":["DATA","CONFIGURATION","HISTORY"]},
    "size":86869667,
    "stepStateMap":{},
    "InProgress":false},
{"createdDate":1496906907.017000000,
  "modifiedDate":1496906910.141000000,
  "id":"3a749282-cbc7-4721-b598-67d785f416bc",
  "backupInformation":{"
    "appVersion":"5.1",
    "createdByEmail":"john.doe@collibra.com",
    "date":1496906907.016000000,
    "dgcBackupOptionSet":["CUSTOMIZATIONS"],
    "repoBackupOptionSet":["DATA","CONFIGURATION","HISTORY"]},
    "size":7155216,
    "stepStateMap":{"
      "REPOSITORY":{"
        "status":"COMPLETED",
        "errorMessage":""},
      "DGC":{"
        "status":"COMPLETED",
        "errorMessage":""},
      "POST_PROCESSING":{"
        "status":"COMPLETED",
        "errorMessage":""},
      "PRE_PROCESSING":{"
        "status":"COMPLETED",
        "errorMessage":""}}},
    "InProgress":false}]
```

REST API - Delete a backup

To delete a backup via the REST API:

```
DELETE https://<your_dgc_console_host>/rest/-  
backup/file/<backup-id>
```

If the operation succeeds, there is no response. If something went wrong, a response is shown in JSON format.

Example error:

```
{  
  "statusCode": 500,  
  "titleMessage": "Internal Server Error",  
  "helpMessage": "Please report a problem as this situation  
should not occur.",  
  "userMessage": "The server encountered an unexpected con-  
dition which prevented \  
it from fulfilling the request. The message  
was: null"  
}
```

Diagnostic files

A diagnostic file is a ZIP file containing files that can be used by support teams to help solve a support ticket. It can contain log files and some configuration files of the database. It does not contain the database content. For more information about logs, go to [Logging](#).

Diagnostic files	223
Create a diagnostic file	225
Download a diagnostic file	227
Delete a diagnostic file	227
Edit the environment log settings	228
Logging	229
Contents of a diagnostic file	231



Diagnostic files

A diagnostic file is a ZIP file containing files that can be used by support to help solve a support ticket.

For example, it contains log files and some configuration files of the database. It does not contain any actual database content. For more information about logs, see [Logging](#).

Content

Tip You can also find the log files in the **collibra_data** directory (for example, **/opt/collibra_data/agent/logs**).

Console

- **installation:** Installation log file and Collibra installation configuration.
- **logs:**
 - **console.log:** Registration of all actions that are performed in Collibra Console (for example, sign in and out, start and stop an environment, and create an environment).
 - **console_wrapper.log**
 - **database*.log:** All transactions of the Collibra Console database.

Agent

- **config:** All configuration files of the agent.
- **installation:** Installation log file and Collibra Data Intelligence Cloud installation configuration.
- **logs:**
 - **agent.log:** Registration of all actions that are performed by the Collibra agent (for example, starting jobs and database operations).
 - **agent_commands.log**
 - **agent_wrapper.log**

DGC

- **logs:**
 - **dgc_jobs.log:** Logging of the actual execution of a job.
 - **dgc_recommender.log:** Logging of the recommendation tool to create data sets for Data Catalog.
 - **dgc_validations.log**
 - **dgc.log:** Logging of the DGC service.

Repo

- **logs:** Logging of the PostgreSQL operations.
- **repository_statistics:** Performance statistics and actions in the repository.
- **blockedLocksInfo.json**
- **componentCalls.dmp**
- **customMetrics.txt:** Statistics of the repository, such as the number of assets, assets per column, active users, and number of jobs.
- **locksInfo.json**

Search

- **logs:** All log files of the Search service.
- JSON files with the Search service's cluster health, cluster settings, cluster state, cluster statistics, and node statistics.

Spark-jobserver

- **logs:** All log files of the Jobserver service.

Monitoring

Note This is not available in cloud environments and Collibra Cloud Self-Hosted 2023.11 and newer.

- **logs:** All log files of the Monitoring service.
- **snapshot:** Snapshot of the monitoring database that is created when you create the diagnostics file.

Other files

- **benchmark-result.txt:** Performance benchmark results if selected during the [creation](#) of the diagnostic file.
- **infraReport.txt:** Details of the infrastructure on which Collibra is installed.
- **installation.log:** Logging of the installation process.
- **installationConfig.json:** Settings that were selected during the installation process, such as TCP ports and installation directory.

Create a diagnostic file

Steps

1. Open Collibra Console.
» Collibra Console opens with the **Infrastructure** page.
2. Click the **Diagnostic files** tab.
3. Click **Create a diagnostic file**.
» The **Create a diagnostic file** dialog box appears.
4. Enter the required information.

Field	Description
Environment	The environment from which you want to create a diagnostic file.
Name	The name for the diagnostic file. There are no specific limitations for the file name.
Select files to include	The files that you want to include in the diagnostic file.

Field	Description
Console files	The logging of the actions in Colibra Console.
DGC files	The logging of the DGC service and a Colibra thread dump.
Repository files	The logging of the Repository service and the Repository service configuration and metrics.
Jobserver files	The logging of the Jobserver service.
Node files	The logging of the nodes. These are the most important operating system log files.
Performance benchmarks	The performance benchmarks of your environment. These benchmarks will perform certain predefined actions on your environment and create a report of the actions. Note This option will slow down your system and will take a few minutes for the diagnostics to be created.
Monitoring data	The monitoring data of your environment. This will create a snapshot of the monitoring database. You can select this option only if you have installed the Monitoring service. Note This option will significantly increase the diagnostic file size.

5. Click **Create a diagnostic file**.

» The diagnostic file is created and it appears in the **Diagnostic files** section of Colibra Console.

Note For information about the file, go to [Contents of a diagnostic file](#).

Download a diagnostic file

Steps

1. Open Collibra Console.
2. On the main menu, click **Diagnostic files**.
3. Do one of the following.

To download	Follow these steps
a single diagnostic file	Click  on the far right of the diagnostic file that you want to download.
one or more diagnostic files	- Select the checkbox next to one or more diagnostic files. Tip If you want all the checkboxes to be selected, select the checkbox in the Name column. - Click Download.

» Depending on your browser and browser settings, the files are downloaded to a default location, or a dialog box appears to prompt you to specify the location for the download.

Delete a diagnostic file

Steps

1. Open Collibra Console.
2. On the main menu, click **Diagnostic files**.
3. Do one of the following.

To delete	Follow these steps
a single diagnostic file	Click  on the far right of the diagnostic file that you want to delete.

To delete	Follow these steps
one or more diagnostic files	a. Select the checkbox next to one or more diagnostic files. Tip If you want all the checkboxes to be selected, select the checkbox in the Name column. b. Click Delete.

Edit the environment log settings

Warning It is highly recommended not to edit the log levels. By doing so, you can influence the performance negatively.

To edit the environment log settings, follow these steps:

1. Open Collibra Console.
» Collibra Console opens with the **Infrastructure** page.
2. In the tab pane, click the **Data Governance Center** service of the environment whose log settings you need.
3. Click **Logs**.
4. Above the table, to the right, click  **Settings**.
5. Edit the settings to suit your company's needs by hovering over the **Level** column and click .

For more information, go to [Environment log settings for Repository services](#).

Environment log settings for Repository services

Slow query logging: With this parameter, you can define the minimum amount of time a query has to last before it is logged in the Collibra Data Intelligence Cloud log files as a slow query. For example, if you set 10 milliseconds, then Collibra only logs queries that take longer than 10 milliseconds.

Logger level:

- **Repository messages:** The minimum severity level of the SQL messages that you want to log.
- **Repository statements:** The minimum severity level of the SQL statements that you want to log.

Repository log levels from low to high (the lower the log level, the more issues that are logged):

- DEBUG5
- DEBUG4
- DEBUG3
- DEBUG2
- DEBUG1
- INFO
- NOTICE
- WARNING
- ERROR
- LOG
- FATAL
- PANIC

Warning It is highly recommended not to change the log levels. By doing so, you can influence the performance negatively.

Logging

When you encounter an issue in Collibra Data Intelligence Cloud, the log file can provide a lot of interesting information about what is happening/has happened.

Whenever there seems to be something wrong with the product, it is good practice to check the log file first.

When you report a problem to Collibra support, it is recommended to immediately attach the log file(s) to your support ticket.

Depending on the configuration, the log file shows you:

- Errors that have occurred on the server side
- Basic information about tasks that are performed
- Detailed information about every action that has been performed
- The date and time an event occurred
- The thread that is executing an action
- The log level of the message

You can see [consult](#) the log files in Collibra Console, or you can download the log files in Collibra Console.

If you download the log files, they are compressed (zipped). You have to decompress them before you can read them. You can use any text editor, like Notepad to read the log files.

Tip If the log file is displayed as a single, long line in Notepad, download and use the (free) text editor Notepad++ from <http://notepad-plus-plus.org>

The log files can also be found on the file system of the server `/opt/collibra_data/<component>/logs`. For more information on the Collibra directory structure, see Product architecture.

Open a log file

Whenever there is something wrong with the product, it is good practice to check the log file first. You can open the log files of each service in Collibra Console.

Tip To open the log files of Collibra Console itself, see [Open a Collibra Console log file](#).

Steps

1. Open Collibra Console.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the tab pane, click the name of an environment.
3. In the **Environment Services** table, click the name of a service whose log files you want to consult.

4. On the details page of the service, click **Logs**.
5. Click the name of a log file to open it.

Contents of a diagnostic file

The [diagnostic file](#) contains the following files.

A diagnostic file is a ZIP file containing files that can be used by support to help solve a support ticket.

For example, it contains log files and some configuration files of the database. It does not contain any actual database content. For more information about logs, see [Logging](#).

Content

Tip You can also find the log files in the **collibra_data** directory (for example, `/opt/collibra_data/agent/logs`).

Console

- **installation:** Installation log file and Collibra installation configuration.
- **logs:**
 - **console.log:** Registration of all actions that are performed in Collibra Console (for example, sign in and out, start and stop an environment, and create an environment).
 - **console_wrapper.log**
 - **database*.log:** All transactions of the Collibra Console database.

Agent

- **config:** All configuration files of the agent.
- **installation:** Installation log file and Collibra Data Intelligence Cloud installation configuration.

- **logs:**
 - **agent.log:** Registration of all actions that are performed by the Collibra agent (for example, starting jobs and database operations).
 - **agent_commands.log**
 - **agent_wrapper.log**

DGC

- **logs:**
 - **dgc_jobs.log:** Logging of the actual execution of a job.
 - **dgc_recommender.log:** Logging of the recommendation tool to create data sets for Data Catalog.
 - **dgc_validations.log**
 - **dgc.log:** Logging of the DGC service.

Repo

- **logs:** Logging of the PostgreSQL operations.
- **repository_statistics:** Performance statistics and actions in the repository.
- **blockedLocksInfo.json**
- **componentCalls.dmp**
- **customMetrics.txt:** Statistics of the repository, such as the number of assets, assets per column, active users, and number of jobs.
- **locksInfo.json**

Search

- **logs:** All log files of the Search service.
- JSON files with the Search service's cluster health, cluster settings, cluster state, cluster statistics, and node statistics.

Spark-jobserver

- **logs:** All log files of the Jobserver service.

Monitoring

Note This is not available in cloud environments and Collibra Cloud Self-Hosted 2023.11 and newer.

- **logs:** All log files of the Monitoring service.
- **snapshot:** Snapshot of the monitoring database that is created when you create the diagnostics file.

Other files

- **benchmark-result.txt:** Performance benchmark results if selected during the [creation](#) of the diagnostic file.
- **infraReport.txt:** Details of the infrastructure on which Collibra is installed.
- **installation.log:** Logging of the installation process.
- **installationConfig.json:** Settings that were selected during the installation process, such as TCP ports and installation directory.

Collibra Console settings

The Collibra Console settings allow you to define how you can access Collibra Console, to manage the Collibra Console users, to configure its mail settings and so on.

In this chapter

Collibra Console users	235
Edit the Collibra Console settings	250
Edit the Collibra Console server settings	272
Open a Collibra Console log file	277



Collibra Console users

The Collibra Console user is different from the Collibra Data Governance Center user. The Collibra Console user can only use Collibra Console and cannot access Collibra DGC, even if the credentials are identical.

Create a Collibra Console user

To create a user to Collibra Console, follow these steps:

1. In the main menu, click **Console settings**.
2. In the tab pane, click **Users**.
3. Above the table, to the right, click **Add User**.
 - » The **Add User** dialog box appears.
4. Enter the required information.

Field	Description
Username	Provide a username for the new Collibra Console user.
Email	Provide the corresponding email address for this user.
Role	◦ Role: Select a role for the user. For more information about roles, see Collibra Console roles. ▪ READ ▪ ADMIN ▪ SUPER

5. Click **Add User**.

What's next?

The new Collibra Console user receives an email with the username and a link to set a password.

Change the role of a Collibra Console user

To change the role of a Collibra Console user, follow these steps:

1. In the main menu, click **Console settings**.
2. In the tab pane, click **Users**.
3. In the **Role** column, hover your pointer over the role next to the user whose role you want to change and click .
4. Click ▼ and click on the new role for the user.

Username	Role	Email
collibra	SUPER READ ADMIN SUPER	admin@collibra.com  

5. Click ✓ to save the change.

Reset the password of a Collibra Console user

If you have forgotten your Collibra Console password, a user with at least the ADMIN role can reset your password.

Note

- A user with the ADMIN role cannot reset the password of a user with the SUPER role.
- You can only reset a password of an other user. To edit your own password, see [Edit your Collibra Console user password](#).
- At any time, users can [edit](#) their own password.

To reset the password of a Collibra Console user, follow these steps:

1. In the main menu, click **Console settings**.
2. In the tab pane, click **Users**.
3. Hover over  next to the user whose password you want to reset and click **Reset password**.
 - » The **Reset password** dialog box appears.
4. In the **Reset password** dialog box, click **Send reset email**.

What's next?

The user receives an email with a link to reset password the password.

Edit your Collibra Console user password

At any time, Collibra Console users can change their own password.

To change your own password, follow these steps:

1. In the main menu, click **Console settings**.
2. In the tab pane, click **Users**.
3. Click  next to your own username.
 - » The **Change password** dialog box appears.
4. Enter the required information.

Field	Description
Old password	Type your current password.
New password	Type your new password.
Repeat password	Type your new password a second time.

5. Click **Change password**.

Delete a Collibra Console user

To delete a Collibra Console user, follow these steps:

1. In the main menu, click **Console settings**.
2. In the tab pane, click **Users**.
3. Click  next to the user that you want to delete.

Note Keep in mind that deleting a user is an irreversible action.

4. Click **Delete** to confirm your action.

What's next?

The user is deleted and can no longer access Collibra Console.

Collibra Console roles

A Collibra Console user can have one of the following roles:

- **READ:** This role gives you read-only access to Collibra Console.
- **ADMIN:** This role includes all the permissions of the READ role, in addition to other permissions, such as the ability to restore backups.
- **SUPER:** This role includes all permissions, such as the ability to stop an environment.

Note The SUPER role is not available in Collibra Console of your Collibra Data Intelligence Cloud environment. Certain actions mentioned in this section can be performed only by Collibra. To perform such actions, [create a support ticket](#). However, note that certain actions cannot be performed at all.

Note Items marked with (*) are available only in on-premises environments.

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
Side menu		• Visible	• Visible	• Visible • Add / Create (environment / node / cluster / add service)
ENVIRONMENTS		• View	• View • Delete environment	• View • Delete environment

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
Default environment		• Create a diagnostic file	• Start • Stop • Create a diagnostic file • More (Create backup; Restore, Delete) • Upgrade if available	• Start • Stop • Create a diagnostic file • More (Create backup; Restore, Delete) • Upgrade if available
	Overview	• View	• View	• View • Add services
	SAML	• View	• View • Upload • Download • Delete	• View • Upload • Download • Delete
	License	• View • Refresh	• View • Refresh	• View • Refresh • Upload new license
	Upgrade details (if available) (*)	• View page • Download log file • View log file • Refresh logs view / change number of lines	• View page • Download log file • View log file • Refresh logs view / change number of lines	• View page • Download log file • View log file • Refresh logs view / change number of lines

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
Data Governance Center			• Stop / Start Data Governance Center service • Rebuild search index and automatic hyperlinks	• Stop / Start Data Governance Center service • Rebuild search index and automatic hyperlinks
	Overview	• View	• View	• View • Delete service
	Monitoring (*)	• View • Change refresh interval	• View • Change refresh interval	• View • Change refresh interval
	Configuration	Page not available	• View • Search • Edit configuration	• View • Search • Edit configuration • Reset to default
	Infrastructure configuration	Page not available	Page not available	• View • Search • Edit configuration • Reset to default

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
	Logs	• View page • Download log file • View log file • Refresh logs view / change number of lines • Open settings page (view only)	• View page • Download log file • View log file • Refresh logs view / change number of lines • Open settings page: ◦ Start logging ◦ Change log level ◦ Reset to defaults ◦ Add logger	• View page • Download log file • View log file • Refresh logs view / change number of lines • Open settings page: ◦ Start logging ◦ Change log level ◦ Reset to defaults ◦ Add logger
Monitoring (*)			• Stop / Start	• Stop / Start
	Overview	• View	• View	• View • Delete service
	Configuration	Page not available	• View • Search • Edit configuration	• View • Search • Edit configuration • Reset to default
Repository			• Stop / Start (*)	• Stop / Start (*)
	Overview	• View	• View	• View • Delete service (*)

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
	Monitoring	• View • Change refresh interval	• View • Change refresh interval	• View • Change refresh interval
	Logs	• View page • Download log file • View log file • Refresh logs view / change number of lines • Open settings page (view only)	• View page • Download log file • View log file • Refresh logs view / change number of lines • Open settings page: ◦ Update period ◦ Change log level ◦ Reset to defaults	• View page • Download log file • View log file • Refresh logs view / change number of lines • Open settings page: ◦ Update period ◦ Change log level ◦ Reset to defaults
Search			• Stop / Start (*)	• Stop / Start (*)
	Overview	• View	• View	• View • Delete service (*)
	Configuration	Page not available	• View • Search • Edit configuration	• View • Search • Edit Configuration • Reset to default

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
	Infrastructure configuration	Page not available	Page not available	• View • Search • Edit configuration • Reset to default
	Logs	• View page • Download log file • View log file • Refresh logs view / change number of lines	• View page • Download log file • View log file • Refresh logs view / change number of lines	• View page • Download log file • View log file • Refresh logs view / change number of lines
Jobserver			• Stop / Start (*)	• Stop / Start (*)
	Overview	• View	• View	• View • Delete service (*)
	Infrastructure configuration	Page not available	Page not available	• View • Search • Edit Configuration • Reset to default

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
	Logs	• View page • Download log file • View log file • Refresh logs view / change number of lines	• View page • Download log file • View log file • Refresh logs view / change number of lines	• View page • Download log file • View log file • Refresh logs view / change number of lines
REPOSITORY CLUSTERS (*)		• View	• View	• View • Delete cluster
	ClusterTest	• View	• View • Start • Stop	• View • Start • Stop • Delete • Set master • Add slaves • Delete Master • Delete Slave
Repository (Master) (*)				• Start / Stop
	Overview	• View	• View	• View
	Monitoring	• View • Change refresh interval	• View • Change refresh interval	• View • Change refresh interval

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
	Logs	• View page • Download log file • View log file • Refresh logs view / change number of lines • Open settings page (view only)	• View page • Download log file • View log file • Refresh logs view / change number of lines • Open settings page: ◦ Update period ◦ Change log level ◦ Reset to defaults	• View page • Download log file • View log file • Refresh logs view / change number of lines • Open settings page: ◦ Update period ◦ Change log level ◦ Reset to defaults
Repository (Slave) (*)	Overview	• View	• View	• View
	Monitoring	• View • Change refresh interval	• View • Change refresh interval	• View • Change refresh interval

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
	Logs	• View page • Download log file • View log file • Refresh logs view / change number of lines • Open settings page (view only)	• View page • Download log file • View log file • Refresh logs view / change number of lines • Open settings page: ◦ Update period ◦ Change log level ◦ Reset to defaults	• View page • Download log file • View log file • Refresh logs view / change number of lines • Open settings page: ◦ Update period ◦ Change log level ◦ Reset to defaults
NODES		• View	• View	• View • Delete node
Default node				• Change node name • Delete node
	Overview	• View	• View	• View
	Monitoring	• View • Change refresh interval	• View • Change refresh interval	• View • Change refresh interval

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
	Logs	• View page • Download log file • View log file • Refresh logs view / change number of lines	• View page • Download log file • View log file • Refresh logs view / change number of lines	• View page • Download log file • View log file • Refresh logs view / change number of lines
	Configuration	Page not available	Page not available	• View • Search • Edit configuration • Reset to defaults
BACKUPS		• View	• View • Create backup • Upload backup • Download • Restore • Delete • Edit name and description	• View • Create backup • Upload backup • Download • Restore • Delete • Edit name and description

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
Backup schedules		• View	• View • Create backup schedule • Delete • Edit name, description, max retained, cron	• View • Create backup schedule • Delete • Edit name, description, max retained, cron
DIAGNOSTIC FILES		• View • Create diagnostic file • Download	• View • Create diagnostic file • Download • Delete	• View • Create diagnostic file • Download • Delete

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
CONSOLE SETTINGS	Users	• View • Change own password • Change own email address	• View • Change own password • Reset password users who have same or lower permission level • Change email address of users who have same or lower permission level • Delete users who have same or lower permission level • Add users who have same or lower permission level	• View • Change own password • Reset passwords • Change email address • Delete users • Add Users

Page / Role → ↓		READ role	ADMIN role	SUPER role (Not available in cloud environments)
	Configuration	• View LDAP configuration • View Backup configuration • View Session configuration • View CSRF and security configuration • Search	• Edit LDAP configuration • Edit Backup configuration • Edit Session configuration • Edit CSRF and security configuration • Search • Edit configuration	• View whole configuration • Search • Edit configuration • Reset to defaults
	Application server configuration	Page not available	Page not available	• View • Search • Edit configuration • Reset to defaults
	Logs	• View page • Download log file • View log file • Refresh logs view / change number of lines	• View page • Download log file • View log file • Refresh logs view / change number of lines	• View page • Download log file • View log file • Refresh logs view / change number of lines

Edit the Collibra Console settings

To edit the Collibra Console settings, follow these steps:

1. Open Collibra Console.
2. In the main menu, click **Console settings**.

3. In the left navigation pane, click **Configuration**.
4. Click **Edit configuration**.
5. Edit the [settings](#).
6. Click **Save all**.

Tip

- You can navigate to a specific section by clicking it in the tab pane.
- When you edit certain fields, the  icon is displayed next to it. When you click it, it displays the default value for that field and a **Reset** button.
- If you have to restart Collibra or execute extra actions to apply the new settings, it is indicated in the user interface.

Collibra Console configuration

Note

- Settings that are indicated with an asterisk (*) are mandatory settings.
- To edit the Console configuration, you need the [ADMIN role](#), except for settings or sections marked with (**), which require the [SUPER role](#). As the SUPER role does not exist for cloud environments, settings that require the SUPER role are not available in Collibra Data Intelligence Cloud environments. If you want to edit one of these settings, please [create a support ticket](#).

Configuration

Setting	Description
Base URL* (**)	The base URL to access Collibra Console.
Startup timeout for Jobserver	The duration in milliseconds before the Jobserver service startup is considered as a failure. The default value is 300,000 milliseconds or 5 minutes.
Startup timeout for the Repository	The duration in milliseconds before the Repository service startup is considered as a failure. The default value is 300,000 milliseconds or 5 minutes.

Setting	Description
X-Frame options	The option to specify the content of the HTTP-header X-Frame-Options. It is set on all the rendered pages and should be used to avoid a click-jacking attack. By default, we specify that only pages from the same origin can use the rendered pages in a frame. Note Changing this options requires a reboot of the server.

Email configuration (**)

The configuration of email notifications. More information: [Configure Colibra Console email settings](#).

Setting	Description
Username	The username used to sign in to the SMTP server.
Password	The password used to sign in to the SMTP server. A password is not required.
From address (*)	The email address used as the sender of all outgoing emails. See Email configuration for Colibra Data Intelligence Cloud and Colibra Console to update this setting in cloud environments.
Host (*)	The URL or IP address of the SMTP server.
Port (*)	The port used to access the SMTP server. The default port is 25.
Start TLS	✓ True: The insecure connections to the SMTP server will be upgraded to a secure connection using SSL or TLS. ✗ False: The connection to the SMTP server does not use SSL or TLS.

LDAP configuration

The configuration of LDAP to handle the authentication. For a tutorial about the LDAP configuration of Collibra Console, see [Collibra Support Portal](#).

Setting	Description
Enable LDAP integration (*)	Global setting to enable or disable LDAP integration.
Enable LDAP SUPER users (**)	Allow LDAP users to have the Collibra Console SUPER role. The default value is false, any LDAP users with SUPER role will have same access rights as ADMIN role users.
Time limit (*)	Maximum time in milliseconds for LDAP search operations. The default value is 120,000 milliseconds or two minutes. Set it to 0 if you do not want a time limit.
Servers	This section configures one or more LDAP servers.
LDAP server URL	When using SSL, use the LDAP's protocol and use the correct port. The SSL section in the configuration should also be configured for this.
Bind DN	The main LDAP Domain Name (DN) to be able to connect to the LDAP server. This DN should be able to access all the users and groups that you want to sync.
Bind password	The password for the main LDAP DN to connect to the LDAP server.
Base DN	General base path to which all DNs are relative. As a general rule, this should stay empty for most configurations.
User base	The base DN where the users are located. Subtree search is used, so all DNs located below the base are searched for matching users.
LDAP user authentication filter	User filter to which users authenticating in Collibra Data Intelligence Cloud should comply.

Setting	Description
LDAP user synchronization filter	User filter to which users imported by the synchronization job should comply. This should be a subset of or equal to the Authentication user LDAP filter. When empty, the same filter will be used. as specified in the Authentication user LDAP filter.
Authentication type	LDAP authentication mechanism; must be one of the following: • <i>none</i>: No authentication is performed. • <i>simple</i>: Simple authentication is performed, using the Bind DN and Bind password as credentials. The credentials are sent as plain text. • <i>DIGEST-MD5</i>: Simple authentication is performed using the Bind DN and Bind password as credentials. The Bind password is hashed with the MD5 algorithm. • <i>TLS-SIMPLE</i>: A temporary secured TLS connection is set up before the credentials are sent as plain text. SSL must be configured. • <i>TLS-EXTERNAL</i>: A temporary secured TLS connection with external SASL authentication using a client certificate. SSL must be configured. Note When using TLS, dont forget to configure the SSL security section.
Shutdown gracefully	In case of TLS usage, if set to true, Colibra Console tries to shutdown all TLS connections.
Referral setting	Specifies how referrals should be handled; must be one of: • <i>throw</i> • <i>ignore</i> • <i>follow</i>
Group base DN	The base DN where all groups are located.
Membership attribute	The attribute to define that users are member of a group, for example uniqueMember. This field is mandatory when mapping LDAP groups to roles.
Group LDAP filter	The LDAP filter to which each group should comply to be synced.
Map the DN of an LDAP group to a role	The mapping of the LDAP users that are members of the group with a specified DN to a role.

Setting	Description
User field mapping	Configuration mapping of all the user fields. Here you can configure to what LDAP field a given user field should be mapped. Leave empty if it should not be considered in the synchronization.
User Name *	Username mapping is mandatory for LDAP to be enabled.
Role*	Role mapping is mandatory for LDAP to be enabled. This is the LDAP directory field that defines which console role the user should have. If you do not intend on using this configuration mapping use the default mapping of <i>role</i> .
Email *	Email mapping is mandatory for LDAP to be enabled.
Default role	The default role for LDAP users. The default role is READ. Only a Collibra Console SUPER user can grant the SUPER role to an LDAP user.

UI configuration (**)

The configuration of user interface features.

Setting	Description
Optimize CSS This setting requires the SUPER role.	✓ True (default): The CSS files are optimized to improve performance of the user interface. ✗ False: The CSS files are not optimized.
Optimize JavaScript This setting requires the SUPER role.	✓ True (default): The JavaScript code is optimized to improve the performance of the user interface. ✗ False: The JavaScript code is not optimized
Concatenate JavaScript This setting requires the SUPER role.	✓ True (default) ✗ False

Setting	Description
Velocity cache This setting requires the SUPER role.	✓ True: The velocity cache is enabled. ✗ False: The velocity cache is disabled. This allows you to reload velocity templates without restarting Collibra.
Modules JSON overrides This setting requires the SUPER role.	DISCLAIMER: If you choose to customize any aspects of Collibra, including CSS or other modules/page-definition customizations, these must be thoroughly tested between upgrades. Customizations are unsupported and can break between upgrades. We recommend your organization and the responsible parties maintain a list of customizations applied to Collibra and use that as a checklist for validating upgrades in a test or lower environment. If changes are needed to customizations, make appropriate preparation and testing plan to promote to your production instance.
Modules properties overrides This setting requires the SUPER role.	DISCLAIMER: If you choose to customize any aspects of Collibra, including CSS or other modules/page-definition customizations, these must be thoroughly tested between upgrades. Customizations are unsupported and can break between upgrades. We recommend your organization and the responsible parties maintain a list of customizations applied to Collibra and use that as a checklist for validating upgrades in a test or lower environment. If changes are needed to customizations, make appropriate preparation and testing plan to promote to your production instance.
Page definition overrides This setting requires the SUPER role.	DISCLAIMER: If you choose to customize any aspects of Collibra, including CSS or other modules/page-definition customizations, these must be thoroughly tested between upgrades. Customizations are unsupported and can break between upgrades. We recommend your organization and the responsible parties maintain a list of customizations applied to Collibra and use that as a checklist for validating upgrades in a test or lower environment. If changes are needed to customizations, make appropriate preparation and testing plan to promote to your production instance .

Backup configuration

Setting	Description
Backup process timeout (*)	The duration in milliseconds before the creation of a backup is considered as a failure. The default value is 43,200,000 milliseconds or 12 hours. During an upgrade, this value is also taken into account.
Backup cleanup interval (Requires restart) (**)	Interval in minutes for the cleanup of backups. The default value is 720 minutes.

1.4.1 Backup schedule configuration

Setting	Description
Default max retained (*)	The default maximum number of retained backups related to a backup schedule. 0 signifies retain all.

1.4.2 Backup retry configuration (**)

Setting	Description
Backup Retry attempts (*)	The number of times an automated retry will be attempted after encountering the initial backup error. The default value is 1.
Backup retry delay (*)	Time in milliseconds between retry attempts. The default value is 5,000 milliseconds.
Backup retry agent check timeout (*)	The maximum amount of time in milliseconds a retry has in order to establish agent communication before the retry will timeout. The default value is 30,000 milliseconds

CSRF configuration

Setting	Description
Enable CSRF protection	If enabled, Collibra Console checks the validity of the request using a CSRF token.

Session configuration

Setting	Description
Idle Session timeout (Requires restart) (*)	The duration in seconds of idle time before a Collibra Console session times out. The minimum value is 15 minutes, the maximum is 24 hours. The default value is 1,800 seconds or 30 minutes.




Security Configuration (**)

Setting	Description
User locked out interval (Requires restart) (*)	The number of seconds during which a user cannot sign in to Collibra Console after the specified number of incorrect Login attempts .
Login attempts (Requires restart) (*)	The number of times that a user can try to sign in to Collibra Console before being locked out for the time defined in User locked out interval .
Interval for consecutive login attempts (Requires restart) (*)	The number of seconds during which the user has the amount of login attempts specified in the Login attempts field.

1.7.1 SSO

Setting	Description
Mode	The SSO mode of Collibra Console. Collibra Console only supports SAML_ATTRIBUTES.
Disable the Collibra Management console signin page	When SSO is enabled, a user can still navigate to the /signin page and try to log in via that page. However, you can disable that page. ✓ True: Users cannot access the Collibra Console signin page. ✗ False (default): Users can access the Collibra Console signin page
Default role	The default role for all SSO users. The default role is READ. SSO users with a SUPER role will be downgraded to users with the ADMIN role. Only a Collibra Console SUPER user can enable the SUPER role for SSO users.
Enable SSO SUPER users (**)	✓ True: SSO users can have the SUPER role. ✗ False (default): SSO users with the SUPER role will be downgraded to users with the ADMIN role. This option is only visible if you have the SUPER role.
SAML	The configuration of SAML.
Metadata HTTP	The URL of the SAML metadata file to be used. The URL always has to be reachable by the Collibra environment.
Entity ID	The entity ID inside the metadata to be referenced. Note A metadata file can describe multiple entity IDs, make sure to use in the entity ID from the correct metadata file.
Consumer service URL	By default, this URL is the same as the URL of your Collibra environment but if your IDP expects another value, you can fill it out here. Warning Make sure that the intended destination endpoint (The Destination attribute in the SAML response) matches the URL being used here. So this is only to be used in specific IDP circumstances When setting this, and getting the error "SAML message intended destination endpoint did not match recipient endpoint" check the Destination attribute in your SAML response and this parameter.

Setting	Description
Disable client address	✓ True: The validation of the client IP address in the assertion message is disabled. ✗ False (default): The validation of the client IP address in the assertion message is enabled.
Sign authentication requests (Requires restart)	✓ True: Authentication requests have to be signed. Use a Collibra generated self-signed certificate to sign requests. The request that is generated by Collibra Console will be appended with a signature in the redirect URL of the response. ✗ False (default): Authentication request don't have to be signed.
Force authn	✓ True (default): The SP authentication request forces re-authentication. ✗ False: The SP authentication request does not force re-authentication.
Force passive	Configure whether the SP authentication should set the authentication to go passive. If True, the IDP or browser MUST NOT take visibly control of the user interface. See the SAML 2.0 specification for more details. This is only relevant if Force authn is <i>True</i>.
Name ID	Configure the nameID used in the SP authentication. If set, the full content will be sent as a nameID. Use a fully qualified nameID. Default nameID="urn:oasis:names:tc:SAML:2.0:nameid-format:persistent". See SAML 2.0 specification for accepted nameID values.
Group to Role mapping	Map groups from the group field to a role in the Collibra Console. The Field key must be a name of a group as defined in your IDP, the Field value is one of the Collibra Console roles (READ, ADMIN, SUPER).
Attribute fields	The mappings of attributes in the SAML response. The values are used as keys to look for in the SAML response.

Setting	Description
Username	The mapping for the user's username. The value of this attribute must be the name of the SAML attribute as defined in your IDP. In the following example, the value in Colibra Console must be <i>email</i>. This is the used name in your IDP software, see the following screenshot.  This field is mandatory.
Role	The mapping for the user roles. The value of this attribute must be the name of the SAML attribute as defined in your IDP. In the following example, the value in Colibra Console must be <i>role</i>. This is the used name in your IDP software, see the following screenshot. 
Group	The mapping to define which attribute holds group information. The value of this attribute must be the name of the SAML attribute as defined in your IDP. In the following example, the value in Colibra Console must be <i>groups</i>. This is the used name in your IDP software, see the following screenshot.  If there is more than one group attribute statement, add them as comma-separated list.
SAML Requested authentication context	Settings for the SAML requested authentication context. The authentication context is the way in which the IDP authenticates the user. By default the authentication context will mandate user/password authentication over HTTPS.
Disable	Disable the SAML requested authentication context. Set to True if you wish to configure the IDP to use specific authentication contexts, without the need to send one in the request.

Setting	Description
Comparison type	Specifies the comparison method used to evaluate the requested authentication context. One of: "exact", "minimum", "maximum", "better". The industry default is "exact", other options are "minimum", "maximum" and "better". • Exact: The authn context in the assertion MUST exactly match the full expected context specified on the SP. • Minimum: The authn context in the assertion MUST be at least as strong as one of the contexts specified on the SP. • Better: The authn context in the assertion MUST be stronger than any of the contexts specified on the SP. • Maximum: The authn context in the assertion MUST be as strong as possible for all of the contexts specified on the SP, without exceeding the strength of at least one context. For more details, see the SAML 2.0 specification (Section 3.3.2.2.1).
Reference list	All SAML authentication classes to be sent in the SAML authentication request. Use to tune the authentication context on the IDP side. Default- t="urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport" (Which means user/password over HTTPS). For more details, see the SAML 2.0 specification (Section 3.3.2.2.1 & Section 2.7.2.2).
Declaration list	All SAML authentication declarative classes to be sent in the SAML authentication request. For details, see the SAML 2.0 specification (Section 3.3.2.2.1 & Section 2.7.2.2).
Signout	The sign-out redirection settings.
Override signout URL	Enable this option to redirect to another page than the default sign-out page. The default page when you sign out is the Collibra Console sign-in page.
Signout redirect URL	The URL to redirect to when you sign out.

1.7.2 HTTP headers

The configuration of the [HTTP response headers](#) of Collibra Console.

Field	Description
URL pattern This setting requires the SUPER role.	The pattern of the URLs to which the HTTP response header is applied. This field supports wildcards such as <code>**</code> , <code>*</code> , and <code>?</code> . Tip The following pattern matches all URLs: <code>/**</code>
HTTP headers This setting requires the SUPER role.	The HTTP response headers in a key-value format. You can add new HTTP response headers by clicking Add at the bottom of the section, and entering the HTTP response header name as the field key and the HTTP response header value as the field value.

Configure Collibra Console email settings

As for Collibra Data Intelligence Cloud, you can set email settings for Collibra Console to enable mails for Collibra Console user activation or password recovery links.

You have to configure the Collibra Console email settings via the Collibra Console user interface.

Note

- The Collibra Console email settings do not apply to the [Collibra email settings](#).
- This section applies only to Collibra Console 5.2 or newer.
- This section applies only to on-premises versions of Collibra Console. For cloud environments, you can only request to change the **From address**. Contact [Collibra Support](#) to change the From address, see also [Email configuration](#).

Prerequisites

You have the **SUPER** role in Collibra Console.

Steps

1. In the main menu, click **Console settings**.
2. In the tab pane, click  **Configuration**.

3. In the middle column, click **Mail configuration**.
4. Above the table, to the right, click **Edit configuration**.
5. Enter the relevant data:

Setting	Description
Username	The username used to sign in to the SMTP server.
Password	The password used to sign in to the SMTP server. A password is not required.
From address (*)	The email address used as the sender of all outgoing emails. See Email configuration for Collibra Data Intelligence Cloud and Collibra Console to update this setting in cloud environments.
Host (*)	The URL or IP address of the SMTP server.
Port (*)	The port used to access the SMTP server. The default port is 25.
Start TLS	◦ ☒ True: The insecure connections to the SMTP server will be upgraded to a secure connection using SSL or TLS. ◦ ☐ False: The connection to the SMTP server does not use SSL or TLS.

6. Click **Save all**.

Configure SSL to access Collibra Console

If you want to connect to Collibra Console in a secure way with your web browser, you have to use SSL. This procedure explains how you can activate SSL access to Collibra Console.

Prerequisites

- You have knowledge of the JSON syntax.
- You have created a Java KeyStore according the procedure described by [Oracle](#), for example `clientkeystore`.

- You have noted the following data while creating the Java KeyStore:
 - KeyStore file name: *clientkeystore* in the Oracle example.
 - KeyStore alias: *client* in the Oracle example.
 - KeyStore password: The password that you entered after executing the command of the first step in the Oracle example.
 - KeyStore alias password: The password that you entered as last step of step 2 in the Oracle example.
- You have stored the Java KeyStore on the node hosting Collibra Console, in the **<collibra_data>/console/security** folder, for example **/opt/collibra_data/-console/security**.

Steps

1. Open a terminal session on the node on which Collibra Console is installed.
2. Open the file **<collibra_data>/console/config/server.json** for editing.
3. Fill in the following parameters in the **httpsConnector** section, if the section does not exist, you have to create it manually:

Note Add string values between double quotes.

Parameter	Description
port	The port on which the HTTPS connector must bind. The value must be higher than 1024 to avoid root permissions. Note If you want to use the default SSL port 443, you have to use a reverse proxy.
keyAlias	The KeyStore alias.
keyPass	The KeyStore alias password.
keystorePass	The KeyStore password.
keystoreFile	The full path to the KeyStore file name, for example /opt/collibra_data/console/security/clientkeystore .

Parameter	Description
Example:	
<pre>"httpsConnector" : { "port": 5404, "keyAlias": "your-alias", "keyPass": "your-password", "keystorePass": "your-password", "keystoreFile": "/opt/collibra_data/console/security/console.jks"}, }</pre>	

4. Save and close the file.
5. Open the DGC service settings for editing:
6. Click the **General settings** section.
7. Update the **Base URL** parameter with *https* and the new port.
8. Restart the environment.

What's next?

Connect to your Collibra Console instance via the Base URL.

Tip To prevent regular HTTP traffic to Collibra Console, update the **address** parameter with the value *127.0.0.1* in `<collibra_data>/console/config/server.json` and restart the environment.

For more information, see the knowledge base on the [Collibra Support Portal](#).

Collibra Console SSO configuration

You can access Collibra Console with user accounts that are created within Collibra Console. Besides those user accounts, you can also configure single sign-on (SSO) access to enable integration with your SSO infrastructure.

Tip If you want to use a custom certificate in the SSO configuration for Collibra Console access, see [this section](#).

Prerequisites



- Your identity provider (IDP) supports SAML 2.0.


Steps

1. Open Collibra Console with a user profile that has at least the **ADMIN** role.
 - » Collibra Console opens with the **Infrastructure** page.
2. In the main menu, click **Console settings**.
3. In the tab pane, click **Configuration**.
4. In the middle column, click **Security Configuration**.
5. Click **Edit configuration**.
6. Enter the required information:

Setting	Description
Mode	The SSO mode of Collibra Console. Collibra Console only supports SAML_ATTRIBUTES.
Disable the Collibra Management console signin page	When SSO is enabled, a user can still navigate to the /signin page and try to log in via that page. However, you can disable that page. ◦ ☒ True: Users cannot access the Collibra Console signin page. ◦ ☐ False (default): Users can access the Collibra Console signin page
Default role	The default role for all SSO users. The default role is READ. SSO users with a SUPER role will be downgraded to users with the ADMIN role. Only a Collibra Console SUPER user can enable the SUPER role for SSO users.
Enable SSO SUPER users (**)	◦ ☒ True: SSO users can have the SUPER role. ◦ ☐ False (default): SSO users with the SUPER role will be downgraded to users with the ADMIN role. This option is only visible if you have the SUPER role.
SAML	The configuration of SAML.
Metadata HTTP	The URL of the SAML metadata file to be used. The URL always has to be reachable by the Collibra environment.

Setting	Description
Entity ID	The entity ID inside the metadata to be referenced. Note A metadata file can describe multiple entity IDs, make sure to use in the entity ID from the correct metadata file.
Consumer service URL	By default, this URL is the same as the URL of your Collibra environment but if your IDP expects another value, you can fill it out here. Warning Make sure that the intended destination endpoint (The Destination attribute in the SAML response) matches the URL being used here. So this is only to be used in specific IDP circumstances. When setting this, and getting the error "SAML message intended destination endpoint did not match recipient endpoint" check the Destination attribute in your SAML response and this parameter.
Disable client address	✓ True: The validation of the client IP address in the assertion message is disabled. ✗ False (default): The validation of the client IP address in the assertion message is enabled.
Sign authentication requests (Requires restart)	✓ True: Authentication requests have to be signed. Use a Collibra generated self-signed certificate to sign requests. The request that is generated by Collibra Console will be appended with a signature in the redirect URL of the response. ✗ False (default): Authentication request don't have to be signed.
Force authn	✓ True (default): The SP authentication request forces re-authentication. ✗ False: The SP authentication request does not force re-authentication.
Force passive	Configure whether the SP authentication should set the authentication to go passive. If True, the IDP or browser MUST NOT take visibly control of the user interface. See the SAML 2.0 specification for more details. This is only relevant if Force authn is <i>True</i>.

Setting	Description
Name ID	Configure the nameID used in the SP authentication. If set, the full content will be sent as a nameID. Use a fully qualified nameID. Default nameID="urn:oasis:names:tc:SAML:2.0:nameid-format:persistent". See SAML 2.0 specification for accepted nameID values.
Group to Role mapping	Map groups from the group field to a role in the Collibra Console. The Field key must be a name of a group as defined in your IDP, the Field value is one of the Collibra Console roles (READ, ADMIN, SUPER).
Attribute fields	The mappings of attributes in the SAML response. The values are used as keys to look for in the SAML response.
Username	The mapping for the user's username. The value of this attribute must be the name of the SAML attribute as defined in your IDP. In the following example, the value in Collibra Console must be <i>email</i>. This is the used name in your IDP software, see the following screenshot.  This field is mandatory.
Role	The mapping for the user roles. The value of this attribute must be the name of the SAML attribute as defined in your IDP. In the following example, the value in Collibra Console must be <i>role</i>. This is the used name in your IDP software, see the following screenshot. 

Setting	Description
Group	The mapping to define which attribute holds group information. The value of this attribute must be the name of the SAML attribute as defined in your IDP. In the following example, the value in Colibra Console must be <i>groups</i>. This is the used name in your IDP software, see the following screenshot.  If there is more than one group attribute statement, add them as comma-separated list.
SAML Requested authentication context	Settings for the SAML requested authentication context. The authentication context is the way in which the IDP authenticates the user. By default the authentication context will mandate user/password authentication over HTTPS.
Disable	Disable the SAML requested authentication context. Set to True if you wish to configure the IDP to use specific authentication contexts, without the need to send one in the request.
Comparison type	Specifies the comparison method used to evaluate the requested authentication context. One of: "exact", "minimum", "maximum", "better". The industry default is "exact", other options are "minimum", "maximum" and "better". ◦ Exact: The authn context in the assertion MUST exactly match the full expected context specified on the SP. ◦ Minimum: The authn context in the assertion MUST be at least as strong as one of the contexts specified on the SP. ◦ Better: The authn context in the assertion MUST be stronger than any of the contexts specified on the SP. ◦ Maximum: The authn context in the assertion MUST be as strong as possible for all of the contexts specified on the SP, without exceeding the strength of at least one context. For more details, see the SAML 2.0 specification (Section 3.3.2.2.1).

Setting	Description
Reference list	All SAML authentication classes to be sent in the SAML authentication request. Use to tune the authentication context on the IDP side. Default-t="urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport" (Which means user/password over HTTPS). For more details, see the SAML 2.0 specification (Section 3.3.2.2.1 & Section 2.7.2.2).
Declaration list	All SAML authentication declarative classes to be sent in the SAML authentication request. For details, see the SAML 2.0 specification (Section 3.3.2.2.1 & Section 2.7.2.2).
Signout	The sign-out redirection settings.
Override signout URL	Enable this option to redirect to another page than the default sign-out page. The default page when you sign out is the Collibra Console sign-in page.
Signout redirect URL	The URL to redirect to when you sign out.

7. Click **Save all**.

Configure a custom certificate for SSO in Collibra Console

If you configure single sign-on for accessing Collibra Console, a default certificate is used. You can use this certificate for signing SAML authn requests.

Instead of the default certificate, you can use your own certificate. However, keep in mind that you can only configure SSO with your own certificate via a REST API call.

Prerequisites

- The certificate must meet the following requirements:
 - The certificate must be in PEM format.
 - The PEM file must be unencrypted (no password).
 - The PEM file must contain the server certificate the private key of that certificate.

Tip To convert a key to a PEM key: `openssl rsa -in <pem-key>.key -out <rsa-key>.pem`

Example PEM file:

```
-----BEGIN PRIVATE KEY-----
MIIEvgIBADA ... bml6YXRpb252YWxz
....
z3P668YfhUbKdRF6S42Cg6zn
-----END PRIVATE KEY-----

# Your certificate
-----BEGIN CERTIFICATE-----
MIIFaDCCBFC ... bml6YXRpb252YWxz
...

lfyD5IymCSuuDim4qB/9bh7oi37heJ4ObpBIzroPUOthbG4gv/5blW3Dc
=
-----END CERTIFICATE-----
```

- A base64 encoding hash of an API user.

Tip You can calculate the base64 hash of the user as follows: `echo '<username>:<password>' | base64`, for example `echo 'apiUser:apiUserpassword' | base64` results in `YXBpVXNlcjphcGlVc2VycGFzc3dvcmQK`

Steps

```
curl --location --request POST \
  https://<your-collibra-console-url>/rest/sam-
lconsole/certificate \
  --header 'Authorization: Basic <base-64 encoding hash>'
  --form 'file=@"/path/to/pem-file"'
```

Edit the Collibra Console server settings

To edit the Collibra Console server settings, follow these steps:

1. Open Collibra Console.
2. In the main menu, click **Console settings**.
3. In the left navigation pane, click **Application server configuration**.
4. Click **Edit configuration**.
5. Edit the [Collibra Console - Application server configuration](#).
6. Click **Save all**.

Tip

- You can navigate to a specific section by clicking it in the tab pane.
- When you edit certain fields, the  icon is displayed next to it. When you click it, it displays the default value for that field and a **Reset** button.
- If you have to restart Collibra or execute extra actions to apply the new settings, it is indicated in the user interface.

Collibra Console - Application server configuration

Application server configuration

HTTP connector

A connector supporting the HTTP/1.1 protocol.

Setting	Description
URI encoding (Requires restart) *	The character encoding used to decode the URI bytes, URL-decoding has been performed.
Acceptor thread count (Requires restart) *	The number of threads to be used to accept connections. The actual request processing is done by separate threads, so you would never really need more than two.

Setting	Description
Min spare threads (Requires restart) *	The minimum number of threads always kept running.
Max threads (Requires restart) *	The maximum number of request processing threads to be created by this connector. This determines the maximum number of simultaneous requests that can be handled.
Accept count (Requires restart) *	The maximum queue length for incoming connection requests that can be assigned to a request processing thread. When the queue is full, requests will be refused.
Compression (Requires restart) *	The connector may use HTTP/1.1 GZIP compression. • off: No compression • on: Allow compression • force: Allow compress
Compression min size (Requires restart) *	The minimum amount of data before it is compressed. This setting only has effect if Compression is on .
Compressible MIME type (Requires restart)	A comma separated list of MIME types allowing HTTP compression.
Connection timeout (Requires restart) *	The number of milliseconds that this connector waits for the destination URI to be presented by the request, after accepting a connection.
Port (Requires restart) *	The TCP port to access your Collibra environment via your web browser.

HTTPS connector

A connector supporting the HTTP/1.1 protocol with SSL support enabled.

Setting	Description
URI encoding (Requires restart) *	The character encoding used to decode the URI bytes, URL-decoding has been performed.
Acceptor thread count (Requires restart) *	The number of threads to be used to accept connections. The actual request processing is done by separate threads, so you would never really need more than two.
Min spare threads (Requires restart) *	The minimum number of threads always kept running.
Max threads (Requires restart) *	The maximum number of request processing threads to be created by this connector. This determines the maximum number of simultaneous requests that can be handled.
Accept count (Requires restart) *	The maximum queue length for incoming connection requests that can be assigned to a request processing thread. When the queue is full, requests will be refused.
Compression (Requires restart) *	The connector may use HTTP/1.1 GZIP compression. • off: No compression • on: Allow compression • force: Allow compress
Compression min size (Requires restart) *	The minimum amount of data before it is compressed. This setting only has effect if Compression is on .
Compressible MIME type (Requires restart)	A comma separated list of MIME types allowing HTTP compression.
Connection timeout (Requires restart) *	The number of milliseconds that this connector waits for the destination URI to be presented by the request, after accepting a connection.

Setting	Description
Port (Requires restart) *	The TCP port to access your Collibra environment via your web browser.

AJP connector

A connector able to communicate with another web connector via the AJP protocol. Mainly for transparent integration with another HTTP server, for example Apache, nginx....

Setting	Description
URI encoding (Requires restart) *	The character encoding used to decode the URI bytes, URL-decoding has been performed.
Acceptor thread count (Requires restart) *	The number of threads to be used to accept connections. The actual request processing is done by separate threads, so you would never really need more than two.
Min spare threads (Requires restart) *	The minimum number of threads always kept running.
Max threads (Requires restart) *	The maximum number of request processing threads to be created by this connector. This determines the maximum number of simultaneous requests that can be handled.
Accept count (Requires restart) *	The maximum queue length for incoming connection requests that can be assigned to a request processing thread. When the queue is full, requests will be refused.
Compression (Requires restart) *	The connector may use HTTP/1.1 GZIP compression. • off: No compression • on: Allow compression • force: Allow compress

Setting	Description
Compression min size (Requires restart) *	The minimum amount of data before it is compressed. This setting only has effect if Compression is on .
Compressible MIME type (Requires restart)	A comma separated list of MIME types allowing HTTP compression.
Connection timeout (Requires restart) *	The number of milliseconds that this connector waits for the destination URI to be presented by the request, after accepting a connection.
Port (Requires restart) *	The TCP port to access your Collibra environment via your web browser.

Static Resources

The static resources are the resources that are reserved for the service.

Setting	Description
Maximum cache size (Requires restart) *	The maximum size that can be assigned to the cache of the service, expressed in kilobytes.

Open a Collibra Console log file

To open the log files of Collibra Console, follow these steps:

1. In the main menu, click **Console settings**.
2. In the tab pane, click **Logs**.
3. Click the name of a log file to open it.

Troubleshooting

Finding resource IDs	278
I don't see a workflow in the global create	281

Finding resource IDs

Every single resource in Collibra Data Intelligence Cloud has a unique ID. Sometimes it happens that you need this unique resource ID in procedures.

In this section, you can find some examples on how to do this.

Find the UUID of a metamodel element

It is possible that you need the UUID (Universally Unique Identifier) of an asset, domain or community, for example to use it in API calls.

You can find this UUID by clicking an asset, domain or community in Collibra Data Intelligence Cloud. The UUID of the selected resource is then shown in the address bar of your browser.

The URL looks like `https://<yourdgcinstance>/<resource type>/bc40c085-352c-4a8c-8ee7-494fe821308e`.

Tip UUIDs that start with 00000000 are packaged resources.

Examples

- Community: <https://<yourdgcinstance>/community/bc40c085-352c-4a8c-8ee7-494fe821308e>
- Domain: <https://<yourdgcinstance>/domain/00000000-0000-0000-0000-000000006019>
- Asset: <https://<yourdgcinstance>/asset/00000000-0000-0000-0000-000000008044>

Find the resource ID of an asset type

You can find the resource ID of an [asset type](#). This may be useful for certain operations, for example for API calls.

Prerequisites

You have a [global role](#) that has the **System administration** [global permission](#).

Steps

1. On the main menu, click , and then click  **Settings**.
 - » The [Collibra settings page](#) opens.
2. Click **Operating Model**.
 - » The [operating model settings](#) appear on the **Asset types** tab page.
3. On the content toolbar, click  → **Columns** → **Resource Id**.
 - » The **Resource Id** column appears.

Tip You can also find the resource ID by clicking the asset type. Then look in the URL of your browser to find the resource ID.

The URL looks like <https://<yourdgcinstance>/assettype/00000000-0000-0000-0000-0000000031101>.

The resource ID of the selected asset type is 00000000-0000-0000-0000-0000000031101, in this example *Business Asset*.

Find the resource ID of a relation type

You can find the resource ID of a relation type. This may be useful for certain operations.

Prerequisites

You have a [global role](#) that has the **System administration** [global permission](#).

Steps

1. On the main menu, click , and then click  **Settings**.
 - » The [Collibra settings page](#) opens.
2. Click **Operating Model**.
 - » The [operating model settings](#) appear on the **Asset types** tab page.
3. In the tab pane, click **Characteristics** → **Relations**.
4. If you don't see the **Resource Id** column, add it to the table.
 - a. On the content toolbar, click  → **Columns** → **Resource Id**.
 - » The **Resource Id** column appears.
 - » The resource ID of the relation types appear in the **Resource Id** column of the table.

Find the resource ID of a complex relation type

You can find the resource ID of a [complex relation type](#). This may be useful for certain operations.

Prerequisites

You have a [global role](#) that has the **System administration** [global permission](#).

Steps

1. On the main menu, click , and then click  **Settings**.
 - » The [Collibra settings page](#) opens.
2. Click **Operating Model**.
 - » The [operating model settings](#) appear on the **Asset types** tab page.
3. In the tab pane, click **Characteristics** → **Complex Relations**.
4. If you don't see the **Resource Id** column, add it to the table.
 - a. On the content toolbar, click  → **Columns** → **Resource Id**.
 - » The **Resource Id** column appears.
 - » The resource ID of the complex relation types appear in the **Resource Id** column of the table.

Find the resource ID of an attribute type

In some actions, you need the resource ID of an attribute type, for example, to edit search a boost factor.

Steps

1. On the main menu, click , and then click  **Settings**.
 - » The [Collibra settings page](#) opens.
2. Click **Operating Model**.
 - » The [operating model settings](#) appear on the **Asset types** tab page.
3. In the tab pane, click **Characteristics** → **Attributes**.
 - » The table with attribute types appears.
4. On the content toolbar, click  → **Columns** → **Resource Id**.
 - » The **Resource Id** column appears.

I don't see a workflow in the global create

If you do not find a workflow in the global create, it means that it has not been added to the global create.

To add a workflow to the global create, follow these steps:

1. On the main menu, click , and then click  **Settings**.
 - » The [Collibra settings page](#) opens.
2. Click **Workflows**.
 - » The [Workflows](#) settings page appears on the **Definitions** tab page.
3. Click the name of the workflow that you want to add.
4. If required, change the **Applies To** section to **Global**.
5. In the **Other** section, select **Show in global create**.

DGC service configuration settings

The settings in this topic are applicable for Collibra Data Intelligence Cloud 2024.01. For on-premises installations, check the [compatibility table](#) to know which version of Collibra Data Intelligence Cloud is equivalent to the version of Collibra Data Governance Center you are using.

Note

- You can edit the DGC service configuration in Collibra Console or [Collibra settings](#). To edit the DGC service configuration, you need the [ADMIN role](#), unless it is specified that a [SUPER role](#) is required.
- The SUPER role does not exist in cloud environments. Therefore, settings that require the SUPER role are not available in Collibra Data Intelligence Cloud environments. If you need to edit one of such settings, [create a support ticket](#).
- The numbers used to identify the following sections may differ from those shown on the **Services Configuration** page of the Collibra settings.

General settings

The general settings of Collibra Data Intelligence Cloud.

Setting	Description
Default locale (Requires restart)	The default locale for new users. It has to contain a language code and may contain a country code. • The language has to be an ISO language code. • The country has to be an ISO country code. Examples: <i>pl</i>, <i>en_US</i>, <i>nl_BE</i>.
Enable view rights	• ✓ True (default): The view permissions feature is enabled. • ✗ False: The view permissions feature is disabled.

Setting	Description
Base URL This setting requires the SUPER role.	The base URL for this Colibra instance, for example <code>http://dgc.example.com</code>. It is the consistent part of the URL to access Colibra Data Intelligence Cloud. This is used amongst others to: • construct hyperlinks to the system, for example in emails. • display your profile picture. • display Tableau report images. • ...
Enable auditing (Requires restart) This setting requires the SUPER role.	• ✓ True (default): Audit and history information is stored. • ✗ False: Audit and history information is not stored.
Google Analytics tracking ID This setting requires the SUPER role.	The Google Analytics 3 Tracking ID for capturing web analytics for your Colibra environment. The tracking ID is used to have the code snippet embedded on the Colibra page to capture default Google Analytics 3 events that occur, such as page visits and form submission.
Google Analytics 4 tracking ID This setting requires the SUPER role.	The Google Analytics 4 Tracking ID for capturing web analytics for your Colibra environment. The tracking ID is used to have the code snippet embedded on the Colibra page to capture default Google Analytics 4 events that occur, such as page visits and form submission.
Show target asset type above relation table	• ✓ True (default): Show the asset type of the target asset in the title of relation tables on an asset page. The target asset can be either the head or the tail of the relation, depending on which asset page you have open. • ✗ False: Hide the asset type of the target asset. The default value is <code>true</code>.

Setting	Description
Collect Application Usage Data This setting requires the SUPER role.	The usage data is used to understand how users interact with Collibra. The information can be used to provide reporting and recommendations. When you enable this setting in a cloud environment, the data collection starts immediately. When you enable the setting in an on-premises environment, you may need to approve the tracking script (pendo.io, app.pendo.io, and cdn.pendo.io). ✓ True (default): Gathering usage data is enabled and sent to Collibra. ✗ False: Gathering usage data is disabled.
Usage Data API key This setting requires the SUPER role.	The Pendo usage data API key that you want to use to collect the usage data.
Homepage	✓ True: Enables the Homepage. When you sign in to Collibra, the Homepage is shown. The Homepage replaces your default dashboard. ✗ False: Disables the Homepage. When you sign in to Collibra, your default dashboard is shown, instead of the Homepage.

Help Menu

The configuration of the Help menu in Collibra Data Intelligence Cloud.

Setting	Description
Links	The list of links in the help menu.
Menu item name	The name of the menu item as it will appear in Collibra Data Intelligence Cloud's help menu.
Menu index	The position of the menu item in the help menu. The top position starts with the value 1.
Menu URL	The target URL of the menu item.
Show admin only	✓ True: The menu item is only visible to users with the Sysadmin role. ✗ False: The menu item is visible to every user.

Email configuration

The configuration of email notifications.

Note In a Collibra Data Intelligence Cloud environment, you cannot update the email server settings, such as host and port. For more information, see Collibra Data Intelligence Cloud infrastructure.

Setting	Description
Default schedule (Requires restart)	The Cron schedule to send emails only at specific times. With this, you can send emails in batches and avoid an overload of mails. Keep in mind that these emails are only workflow emails and have nothing to do with the notification schedule. If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.
Template map	The location of template emails.
Password This setting requires the SUPER role.	The password paired with your username to sign in to your SMTP server.
From address	The email address used as the sender of all outgoing emails. Contact Collibra Support to change the From address, see also Email configuration.
Port This setting requires the SUPER role.	The port to connect to your SMTP server. The default value is 25 .
Host This setting requires the SUPER role.	The hostname or URL of your SMTP server.

Setting	Description
Start TLS This setting requires the SUPER role.	✓ True: Use TLS (Transport Layer Security) to connect to your SMTP server. ✗ False (default): Do not use TLS to connect to your SMTP server.
Username This setting requires the SUPER role.	The username to sign in to your SMTP server.
Sending threads This setting requires the SUPER role.	The number of threads that are used to send emails. The default value is 3 .
Max retries This setting requires the SUPER role.	The maximum number of retries before the system aborts the sending of an email. The default value is 5 .
Email address change notification This setting is only available for the ADMIN role.	If you change the email address to which notifications are sent, notification of the change is sent to the old email address.

Notifications

The configuration of notification emails to users.

Note These settings can be overridden for every user in the preferences.xml file.

Setting	Description
Notification days	The days of the week on which Collibra sends notifications. The days are represented by numbers from 1 to 7, where 1 represents Sunday. Per row you can add one day.
Daily roles	The roles that receive notifications on the days defined in Notification days .
Enable monthly notifications	✓ True: The users receive a monthly summary. ✗ False (default): The users do not receive a monthly summary.
Roles for monthly notifications	The roles that receive monthly notification emails. This is only relevant if Enable monthly notifications is ✓ True.

Handlers

A mail handler can poll for emails on a mail server, process those emails and perform actions based on the contents.

Setting	Description
Host This setting requires the SUPER role.	The hostname or URL of the incoming mail server.
Port This setting requires the SUPER role.	The port to connect to your incoming mail server.
Protocol This setting requires the SUPER role.	The protocol to connect to your incoming mail server, with or without SSL (<i>POP3</i>, <i>POP3S</i>, <i>IMAP</i>, <i>IMAPS</i>). Note The additional S at the end of the abbreviations stands for the secure version of the protocol using SSL. Using this requires the SSL certificates to be correctly configured.

Setting	Description
Force domain This setting requires the SUPER role.	✓ True: Only handle emails from the same domain as the handler's email address. ✗ False (default): Handle emails from any domain.
Handler list This setting requires the SUPER role.	The configuration of email handlers, which can poll emails on an email server, process those emails and perform actions based on the contents.
Enabled This setting requires the SUPER role.	✓ True: The handler is enabled. ✗ False (default): The handler is not enabled.
Name This setting requires the SUPER role.	The name of the mail handler. We recommend to use a meaningful name to easily identify what this handler is used for.
Username This setting requires the SUPER role.	The username to connect to the incoming mail server.
Password This setting requires the SUPER role.	The password to connect to the incoming mail server.
Email address This setting requires the SUPER role.	The email address to which workflow action mails are sent.
Polling interval This setting requires the SUPER role.	The time in milliseconds between two pollings of the mail server.

Setting	Description
Delete This setting requires the SUPER role.	✓ True: Delete messages from the mail server once the mail is processed. ✗ False (default): Keep messages on the mail server after the mail is processed. This option is only relevant if Protocol is <i>IMAP</i> or <i>IMAPS</i>.
Alias filter This setting requires the SUPER role.	✓ True (default): Retrieve only the emails of which the To field contains the email address of the handler. ✗ False: Do not filter on the To field.

Hyperlinking configuration

The configuration of automatic hyperlinks. When you change a setting, you have to [rebuild the hyperlinks](#).

Setting	Description
Enable hyper-linking	✓ True: Hyperlinks are created automatically. ✗ False (default): Hyperlinks are not created automatically. For more information about automatic hyperlinks, see Hyperlinking. Warning If you enable this setting, the performance of Collibra can decrease.
Enable case sensitivity	✓ True: Hyperlinks are case-sensitive. ✗ False (default): Hyperlinks are not case-sensitive. Note If you edit this setting, you have to reindex Collibra.

Setting	Description
Excluded asset type IDs	The list of asset types that are ignored by automatic hyperlinking. You can use this setting to exclude particular asset types so that not all the assets are potential targets for automatic hyperlinking. Each time you enter an asset type ID in a field, a new field appears for you to add another ID. To exclude multiple asset types, use separate fields, instead of entering multiple IDs in a single field separated by commas. Excluding assets reduces the amount of hyperlinks, which improves performance. Tip We recommend that you exclude technical asset types such as Column, Field, Table, Code Value, and Code Set. Note If you edit this setting, you have to reindex Collibra.

Recommender configuration

The configuration of the recommender.

We recommend not to change the default values of this configuration.

Setting	impacts	Description
Catalog recommender enabled	All recommendations	✓ True (default): The "Data sets you might like" section is included on the Data Catalog Home page. This section shows data sets you might be interested in, as determined by the recommender, which takes into account your data sets and the data sets of similar users. ✗ False: The "Data sets you might like" section is not included on the Data Catalog Home page.
Data set recommender execution time	Recommendations of data sets to users	The schedule (CRON job) by which the data set recommender looks for recommended data sets for a user. By default the data set recommender does this every night.

Setting	impacts	Description
Asset recommender execution time	Recommendations of business assets to data assets	The schedule (CRON job) by which the asset recommender looks for suggested relations between business assets and data sets.
Data set matcher execution time	Data set matcher	The schedule (CRON job) by which the data set matcher looks for similar data sets.
Data set similarity threshold	Data set matcher	The amount of business assets that have to be related to two data sets before the data sets are considered to be similar. This percentage is expressed by a decimal where 1,00 equals 100%. Example If this value is 0.3 and at least 30% of the related business assets are related to both data sets, they are considered to be similar.
Duplicate schema threshold	Schema matcher	The amount of assets that have to be related to both schemas before the schemas are considered to be similar. This percentage is expressed by a decimal where 1,00 equals 100%.
Fuzzy vs exact matching strategy for business assets	Recommendations of business assets to data sets and of business assets to column assets	The percentage that determines to what extent assets with a similar name become more important. The ranking in the search engine results always has an impact on the suggestion score. However, similarity between the asset names can also be taken into account. If you decrease this percentage, the ranking of the search results becomes more important for the suggestion score, while the similarity between the asset names becomes less important. If you increase the percentage, assets with similar names will receive a higher suggestion score. This percentage is expressed by a decimal where 1,00 equals 100%. You can enter a value greater than 1,00.

Setting	impacts	Description
Recommendation weights for data sets	Recommendations of data sets to users	An ordered comma-separated list of values that define the importance of properties for recommendations. The order of the values reflects the importance of the value. This setting is only used for data set recommendations if your Colibra does not yet have enough data for relevant results from the active recommendations algorithms. Possible values: • <i>CERTIFIED</i>: Data sets that are certified are considered more relevant. • <i>POPULARITY</i>: The number of visits to the data set page.
Active recommendation algorithms	Recommendations of data sets to users and of business assets to data sets	A comma-separated list of algorithms that calculate recommendations. By default, all available algorithms are listed. Possible values: • <i>BASELINE</i> • <i>USER_MEAN</i> • <i>IICF</i> (Item-Item Collaborative Filtering) • <i>SLOPE_ONE</i> • <i>WEIGHTED_SLOPE_ONE</i>
Data set elements threshold	Recommendations of data sets to users	The maximum number of elements per data set that the recommender will use to train the model. The data set elements are taken randomly. Lowering this number can prevent out-of-memory issues but also impacts the accuracy of recommendations for large data sets.

Warning If you create an invalid Cron pattern, Colibra Data Intelligence Cloud stops responding.

Search index configuration

The configuration of the search index.

Setting	Description
UI search appends wildcard	✓ True (default): A wildcard (asterisk) is automatically added to each search query. An asterisk is not added if: - The query contains a tilde (~). - The query ends with a quotation mark ("). Note This applies only to queries via the user interface. A wildcard is not added automatically for REST API queries. ✗ False: A wildcard is not added to the search query.
Maximum batch size	The amount of resources that are scanned in a single operation for the search query. - Default value: 5,000 - Maximum value: 30,000
Maximum batch size for relations	The maximum batch size for relations reindex. - Default value: 500 - Minimum value: 50 - Maximum value: 10,000
Stop words (Requires restart)	A list of stop words that are ignored as tokens for the index. The default list of English stop words includes the following: a, an, and, are, as, at, be, but, by, for, if, in, into, is, it, no, not, of, on, or, such, that, the, their, then, there, these, they, this, to, was, will, with If you do not create your own list of stop words, the default list is used. If you create your own list of stop words, you need to: Reindex Colibra Data Intelligence Cloud. - Restart the environment to apply your changes. For more information, go to Stop an environment and Start an environment.

Setting	Description
Relation-based search	✓ True (default in new environments): The Data Marketplace search considers certain assets and relation types between assets. As a result, your search results not only include assets that directly match the search criteria, but also assets that match the criteria through specific relation types. Example A column named Order is included in a data set named Customer. If the relation-based search is enabled and you search for Order in Data Marketplace, then the data set Customer appears in the search results because the data set contains this column. Tip For more information about this feature and the default relation types, go to Filtering and searching based on relations in Data Marketplace. ✗ False: The Data Marketplace search results do not consider relations. After you enable this setting, you must reindex Data Marketplace relations or reindex Collibra completely. Note In new Collibra environments, this setting is enabled by default. In upgraded Collibra environments, the previous status of this setting is retained.

Tokenizer

The configuration of the tokenizer of the indexing mechanism. If you edit these settings, you need to restart and [reindex](#) your environment.

Setting	Description
Type	The tokenizer that determines how the search text is split. The Type field must contain either Standard (default) or Character. • Standard: The Standard tokenizer is a method of splitting the search text into individual terms and is based on a default set of characters. This tokenizer follows the word break rules from the Unicode Text Segmentation algorithm. • Character: The Character tokenizer allows you to customize when you want a search text to be split into individual terms. For more information, go to When and how to use the Character tokenizer.
Parameter map	A list of characters that the Character tokenizer allows. If you entered <i>Character</i> in the Type field, you must add a parameter map. To add a parameter map: - Click Add. » The Add map option dialog box appears. - In the Field key field, enter the following value: <i>allowedCharacters</i> - In the Field value field, enter the set of characters that you want the tokenizer to allow. For more information, go to When and how to use the Character tokenizer.

Boosting

The configuration of the [boosting](#) function.

Setting	Description
Asset	The boost factor of assets.
Class Match	The boost factor of data classes.
Community	The boost factor of communities.
Domain	The boost factor of domains.
User	The boost factor of users.

Setting	Description
User group	The boost factor of user groups.
Name	The boost factor of names.
Comment	The boost factor of comments.
Tag	The boost factor of tags.
Attribute boost map	The boost factor of attribute types. • Field key: The attribute type ID. • Field value: The boost factor of the attribute type.
Display exact match of name as first	• ✓ True (default): If the name of an asset is exactly the same as the search text, put it at the top of the search results regardless of boost factors. • ✗ False: Use the regular search order, taking into account boost factors.
Asset boost map	The boost factor of asset types. • Field key: The asset type ID. • Field value: The boost factor of the asset type.
Partial exact match enabled	Enables partial exact matching while searching for multi word phrases. • ✓ True (default): For multi-word search text, the search engine considers the exact match percentage with the resource name, when ordering the results. Example You enter search text "scheduled maintenance". Two example assets are ordered as follows: - An asset named "daily scheduled maintenance", as two of the three words (66%) match exactly. - An asset named "daily scheduled maintenance revised", as two of the four words (50%) match exactly. • ✗ False: The exact match percentage is not taken into account in the score calculation.

Slow logs configuration

The configuration of the slow logs function.

Setting	Description
Indexing threshold	The time limit, in milliseconds, after which an index query is logged in Elasticsearch. If the value is set to 0 (zero), all index queries are logged. Changes to this setting require a full reindex of your Collibra Data Intelligence Cloud environment.
Fetching threshold	The time limit, in milliseconds, after which a fetch query is logged in Elasticsearch. If the value is set to 0 (zero), all fetch queries are logged. Changes to this setting require a full reindex of your Collibra Data Intelligence Cloud environment.

5.4 Search Event Log configuration

The configuration of indexing.

Setting	Description
Automatic relation indexing	This setting keeps Data Marketplace up to date if relations between assets are created, updated, or removed. Example If the relation between asset A and asset B changes and this relation is used in relation-based filters or relation-based search, then the Data Marketplace search considers this change. ✓ True: Automatically index certain relation type changes between assets so that the relation information remains consistent between Collibra and Data Marketplace. The relation types that are considered are the relation paths used by relation-based search and filters. If such a relation type between assets changes, the change is reflected in the search index after some time. Tip For more information about this feature and the default relation types, go to Filtering and searching based on relations in Data Marketplace. Note Collibra does not automatically reindex relations between assets for relation paths that end with an attribute. You need to manually reindex the relations. You, for example, created a path that ends with an attribute: Table A contains Column B with attribute Privacy. If you have a column with attribute Privacy and value "sensitive data", a user searching for "sensitive data" can find Table A based on the relation path. However, changes to the attribute value will not be picked up automatically during automatic reindexing. ✗ False (default): Changes to relations are not automatically indexed. This can cause inconsistencies between Collibra and Data Marketplace. You can, however, manually reindex Data Marketplace relations.

Upload configuration

The configuration of the file upload service.

The file upload restrictions apply to the following actions in Collibra:

- [Importing and exporting assets and complex relations.](#)
- Uploading [attachments](#).
- Importing and exporting CMA files in the [Migration feature](#).

Setting	Description
Max file size	The maximum file size in bytes for uploads. • For cloud environments, the default value is 512 MB or 536,870,912 bytes. This value cannot be changed. • For on-premises environments, the default is 10 MB or 10,485,760 bytes.
Max per day	The maximum number of uploads per user per day. • For cloud environments, the default value is 1,235,465 uploads. This value cannot be changed. • For on-premises environments, the default is 150 uploads.
Accepted content types	The MIME type names of the files you want to allow for uploads. For example, type <i>application/pdf</i> for PDF files.
Restricted content types	Content types in MIME type format that cannot be uploaded. Restricted content types take precedence over an accepted content types: • application/vnd.ms-excel.addin.macroenabled.12 • application/vnd.ms-excel.sheet.binary.macroenabled.12 • application/vnd.ms-excel.sheet.macroenabled.12 • application/vnd.ms-excel.template.macroenabled.12 • application/vnd.ms-powerpoint.addin.macroenabled.12 • application/vnd.ms-powerpoint.presentation.macroenabled.12 • application/vnd.ms-powerpoint.slide.macroenabled.12 • application/vnd.ms-powerpoint.slideshow.macroenabled.12 • application/vnd.ms-powerpoint.template.macroenabled.12 • application/vnd.ms-word.document.macroenabled.12 • application/vnd.ms-word.template.macroenabled.12

Statistics configuration

The configuration of statistics.

Setting	Description														
Buffer size	The maximum amount of statistics entries that the buffer can contain before saving them in the database. The default value is <i>10</i>.														
Buffer flush time	The maximum amount of time in milliseconds to keep statistic entries in memory before saving them in the database. The default values is <i>10,000</i>.														
Cron map	List of statistics, listed by their Cron name, and a Cron interval. These are the default values: <table> <tr> <th>Field key</th><th>Field value</th></tr> <tr> <td>workflow-task</td><td>0 59 23 * * ?</td></tr> <tr> <td>active-users</td><td>0 0/15 * * * ?</td></tr> <tr> <td>term-count</td><td>0 59 23 * * ?</td></tr> <tr> <td>vocabulary-count</td><td>0 59 23 * * ?</td></tr> <tr> <td>page-hit</td><td>0 0 * * * ?</td></tr> <tr> <td>task-count</td><td>0 0 * * * ?</td></tr> </table> If you create an invalid Cron pattern, Colibra Data Intelligence Cloud stops responding.	Field key	Field value	workflow-task	0 59 23 * * ?	active-users	0 0/15 * * * ?	term-count	0 59 23 * * ?	vocabulary-count	0 59 23 * * ?	page-hit	0 0 * * * ?	task-count	0 0 * * * ?
Field key	Field value														
workflow-task	0 59 23 * * ?														
active-users	0 0/15 * * * ?														
term-count	0 59 23 * * ?														
vocabulary-count	0 59 23 * * ?														
page-hit	0 0 * * * ?														
task-count	0 0 * * * ?														

Import configuration

The configuration for imports.

Setting	Description
Enable workflows during import	✓ True: Allow starting workflows upon importing assets. ✗ False (default): Do not allow to start workflows upon importing assets.

Setting	Description
Asset responsibilities support	✓ True: Enable importing responsibilities at asset level. ✗ False (default): Disable importing responsibilities at asset level. Warning Setting specific responsibilities on a large number of resources will affect the performance and stability of the system.
Number of failed commands before stopping import job	An import job with the option to continue on error enabled will stop after the specified number of commands have failed. Any valid command is still committed to the database until the moment the job stops, which can lead to some resources being imported. The default and maximum value is <i>100</i>.
Temporary data location	The location of the temporary files used by the import job. The default value is <i>FILE</i>.
Import UI v2 This setting requires the SUPER role.	✓ True (default): Use the new import interface for importing assets and complex relations, with improved usability and performance. ✗ False: Use the original import interface for importing assets and complex relations.

Excel import configuration

The configuration of Excel import.

Setting	Description
The default CSV separator character	The default separator character of the CSV fields for complex relations.
The default CSV quote character	The default quote character of the CSV fields for complex relations.

Setting	Description
Number of rows per chunk of data	When importing views, the database is called repeatedly, each time importing a chunk of data from the import file. This option defines how many rows each chunk of data can contain. Lower values reduce the burden on memory. Higher values require more memory, but may slightly increase the speed of the export. The default value is 5,000.

Excel export configuration

The configuration of Excel export.

Setting	Description
The default CSV separator character	The default separator character of the CSV fields for complex relations.
The default CSV quote character	The default quote character of the CSV fields for complex relations.
Number of rows per chunk of data	When exporting views, the database is called repeatedly, each time fetching a chunk of data to build the export file. This option defines how many rows each chunk of data can contain. Lower values reduce the burden on memory. Higher values require more memory, but may slightly increase the speed of the export. The default value is 5,000.

CSV export configuration

The configuration of CSV export.

Setting	Description
Always use quotes	✓ True: Use quotes for every cell in the CSV. ✗ False (default): Only use quotes when necessary.
Number of rows per chunk of data	When exporting views, the database is called repeatedly, each time fetching a chunk of data to build the export file. This option defines how many rows each chunk of data can contain. Lower values reduce the burden on memory. Higher values require more memory, but may slightly increase the speed of the export. The default value is 5,000.

User interface configuration

The configuration of user interface features.

Setting	Description
Optimize CSS This setting requires the SUPER role.	✓ True (default): The CSS files are optimized to improve performance of the user interface. ✗ False: The CSS files are not optimized.
Optimize JavaScript This setting requires the SUPER role.	✓ True (default): The JavaScript code is optimized to improve the performance of the user interface. ✗ False: The JavaScript code is not optimized
Concatenate JavaScript This setting requires the SUPER role.	✓ True (default) ✗ False

Setting	Description
Velocity cache This setting requires the SUPER role.	✓ True: The velocity cache is enabled. ✗ False: The velocity cache is disabled. This allows you to reload velocity templates without restarting Collibra.
Modules JSON overrides This setting requires the SUPER role.	DISCLAIMER: If you choose to customize any aspects of Collibra, including CSS or other modules/page-definition customizations, these must be thoroughly tested between upgrades. Customizations are unsupported and can break between upgrades. We recommend your organization and the responsible parties maintain a list of customizations applied to Collibra and use that as a checklist for validating upgrades in a test or lower environment. If changes are needed to customizations, make appropriate preparation and testing plan to promote to your production instance.
Modules properties overrides This setting requires the SUPER role.	DISCLAIMER: If you choose to customize any aspects of Collibra, including CSS or other modules/page-definition customizations, these must be thoroughly tested between upgrades. Customizations are unsupported and can break between upgrades. We recommend your organization and the responsible parties maintain a list of customizations applied to Collibra and use that as a checklist for validating upgrades in a test or lower environment. If changes are needed to customizations, make appropriate preparation and testing plan to promote to your production instance.
Page definition overrides This setting requires the SUPER role.	DISCLAIMER: If you choose to customize any aspects of Collibra, including CSS or other modules/page-definition customizations, these must be thoroughly tested between upgrades. Customizations are unsupported and can break between upgrades. We recommend your organization and the responsible parties maintain a list of customizations applied to Collibra and use that as a checklist for validating upgrades in a test or lower environment. If changes are needed to customizations, make appropriate preparation and testing plan to promote to your production instance .

API call logging

The configuration of the API call logging.

Setting	Description
Enabled	✓ True: API call logging is enabled. ✗ False (default): API call logging is disabled.
Maximum number of log entries (Requires restart)	The maximum number of API calls to store in the database. Once this number is reached, the oldest records are overwritten. The default value is 1,000,000.
Pattern duration list	The list of methods and a corresponding minimum duration time. The minimum duration time is the minimum time before the method is stored in the database.
Minimum duration	The time in milliseconds that an API call must last before it is logged.
Method pattern	The method that you want to log in the database. For each pattern that you want to log, you have to add a new pattern.

System metrics

The configuration of metric collection.

Setting	Description
Enable (Requires restart)	✓ True (default): Metric collection is enabled. ✗ False: Metric collection is disabled.
Enable JVM metrics (Requires restart)	✓ True (default): JVM metric collection is enabled. ✗ False: JVM metric collection is disabled.
Enable advanced metrics (Requires restart)	✓ True: Advanced metrics collection is enabled. Enabling this option has a negative impact on the performance of your environment. ✗ False (default): Advanced metrics collection is disabled.
Enable minimal monitoring (Requires restart)	✓ True: Monitoring of the metrics is enabled. ✗ False (default): Monitoring of the metrics is disabled.

API configuration

The configuration of API settings.

Setting	Description
Enable maximum paging limit	✓ True (default for new environments): The maximum paging limit is set to 1,000 data elements per API call. Note Once the maximum paging limit has been enabled, it cannot be disabled. ✗ False : There is no maximum paging limit per API calls. We recommend you enable the maximum paging limit, as too many data elements per API call can cause your environment to crash.

Security configuration

The configuration of security.

Setting	Description
X-Frame options (Requires restart)	The content of the HTTP-header <code>X-Frame-Options</code> . This is set on all rendered pages and is used to avoid clickjacking attacks . By default, only pages with the same origin can use the rendered pages in a frame.
Limit user sessions	✓ True: A user can only open one session. ✗ False (default): A user can open multiple sessions.
Office research guest access	✓ True: The Office research integration is always allowed guest access via REST, regardless of the general Guest access setting. ✗ False (default): The general Guest access setting is kept. Note Currently, The Office research integration is only available when Collibra Data Intelligence Cloud is publicly available, which is why this override setting is necessary.

Setting	Description
Prevent advanced html features in text dashboard	Text widgets can contain full HTML. However, this means an attacker could potentially execute an XSS attack by injecting malicious HTML. For more information, see the Troubleshooting section. ✓ True: Potentially dangerous HTML elements are removed from text attributes when you save the text field. ✗ False (default): No HTML elements are removed from text attributes when you save the text field. Note If you enable this setting, the following HTML elements are deleted when you save: • script (including JavaScript) • svg • frame • frameset • iframe • any event handlers
Guest access This setting requires the SUPER role.	✓ True: Anyone that can access the URL, has viewing rights to the system. ✗ False (default): The user is asked to sign in before having access to any data.
Enable schema introspection	✓ True: Schema fields are shown during an introspection. ✗ False (default): Schema fields are hidden during an introspection.
Enable schema introspection for Public GraphQL APIs (Requires restart)	✓ True (default): Schema fields are shown when using introspection for public APIs. ✗ False: Schema fields are hidden when using introspection.
Enable customer validation functions	✓ True (default): Groovy scripts with custom validation functions can be loaded. ✗ False: Groovy scripts with custom validation functions cannot be loaded.

LDAP

The configuration of an LDAP server to handle the authentication.

Setting	Description
Enable LDAP integration (Requires restart)	✓ True: The LDAP integration is enabled. ✗ False (default): The LDAP integration is disabled.
Sync after restore	✓ True (default): LDAP data is synchronized with Collibra when an initial data set is bootstrapped. ✗ False: LDAP data is synchronized with Collibra only when the LDAP synchronization job is triggered.
User page size	The page size that is used when retrieving users during synchronization. The default value is 500. You can set it to 0 to disable paging. Note This is a global setting. If you are working with multiple LDAP servers, only the value for the main server is taken into account.
Group page size	The page size that is used when retrieving groups. You can set it to 0 to disable paging. Note This is a global setting. If you are working with multiple LDAP servers, only the value for the main server is taken into account.
Time limit	Specifies the time limit in milliseconds for all LDAP searches. The default value is 120,000. You can set it to 0 to disable the time limit. Tip - If you get Time limit Exceeded error messages, increase the default value or check why the LDAP search takes too long. - We recommend that you modify the User page size and Group page size settings before you modify this setting.

Setting	Description
Sync job enabled	✓ True (default): The synchronization job is enabled. ✗ False: The synchronization job is disabled.
Sync job cron	The schedule to perform an LDAP synchronization (CRON). The default value for this setting is daily at midnight. If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.
User field mapping	The configuration mapping of all the user fields. This determines which LDAP field is mapped to which user field. Empty fields are ignored during the synchronization.
Username	The unique user ID in the LDAP, typically UID. This is a mandatory field.
Email	The corresponding email field in the LDAP directory. This is a mandatory field.
First name	The first name field in the LDAP directory.
Last name	The last name field in the LDAP directory.
Middle name	The middle name field of the LDAP directory, this is usually givenName.
Enabled	Indication whether a user is active or inactive in LDAP.
Language	The language and locale of the user. It has to contain a language code and may contain a country code. - The language has to be an ISO language code. - The country has to be an ISO country code. Examples: <i>pl</i>, <i>en_US</i>, <i>nl_BE</i>.
Group	The LDAP property that defines to which groups the user belongs. If there is a group entry in the LDAP directory, use the Group field mapping settings.
Additional email list	An additional email list.
Instant messaging fields	The mapping for the user's IM locations.

Setting	Description
AIM	The mapping for the user's AOL IM account.
Google Talk	The mapping for the user's Google Talk IM account.
Icq	The mapping for the user's ICQ IM account.
Jabber	The mapping for the user's Jabber IM account.
Messenger	The mapping for the user's Live Messenger IM account.
Skype	The mapping for the user's Skype IM account.
Yahoo Messenger	The mapping for the user's Yahoo Messenger IM account.
Website map	Enter the field value and field key to map a social media website.
Phone	The mapping for the user's phone.
Fax	The mapping for the user's fax number.
Mobile	The mapping for the user's mobile number.
Pager	The mapping for the user's pager number.
Private	The mapping for the user's private number.
Work	The mapping for the user's work number.
Other	The mapping for any other phone number for this user.
Home address	The mapping for the user's home address.
Street	The mapping for the user's street.
Number	The mapping for the user's number.
City	The mapping for the user's city.
Post code	The mapping for the user's postal code.

Setting	Description
State	The mapping for the user's state.
Country	The mapping for the user's country.
Work address	The mapping for the user's work address.
Street	The mapping for the user's street.
Number	The mapping for the user's number.
City	The mapping for the user's city.
Post code	The mapping for the user's postal code.
State	The mapping for the user's state.
Country	The mapping for the user's country.
Gender	The mapping information for the user's gender.
Mapping	The attribute key for the gender value. If the content equals one of the male or female mappings, the user will be saved as male or female. Otherwise a default of <i>UNKNOWN</i> will be used.
Male value	The value for male users.
Female value	The value for female users.
Group field mapping	Groups can be defined as a separate structure or as a userField. The following section allows you to sync with a group structure that is unrelated to the user structure.
Group name field	The name of the group to use in the application.
Users field	The user DNs that are member of the group.
Servers	The Colibra parameters to map with your LDAP server parameters.

Setting	Description
LDAP server URL	The URL or IP address to the LDAP server, for example <i>ldap://ldap.yourcompany.com:389</i> or <i>ldaps://ldap.yourcompany.com:636</i> .
Bind DN	The DN of the administrator user that is used for authentication, for example <i>admin</i> .
Bind password	The password of the administrator user.
Base DN	The base DN for when you are working with relative DNs. This base DN is used for all LDAP look-ups.
User base	The base DN of where the LDAP users for Collibra are located. If a base has been specified, it is used as a prefix for this user base. Subtree search is used, so all DNs located below are searched for matching users.
Authentication user LDAP filter	The filter that specifies which users can authenticate in the application. By default, all the objects found in the user base are selected, including the root.
Synchronization user LDAP filter	The filter that specifies which users are imported by the synchronization job. The users have to be the same as, or a subset of, the Authentication user LDAP filter. If you provide no value for this setting, the same filter as specified for the Authentication user LDAP filter setting is used. That allows you to synchronize only the users that have to have access to the application, even if they have not logged in yet. Users in the Authentication user LDAP filter are synchronized each time they authenticate and are only available after the first sign-in to the application. This is the default setting.

Setting	Description												
Authentication type	The authentication mechanism for authenticating users on the LDAP servers.												
	<table><tr><th>Authentication type</th><th>Explanation</th></tr><tr><td>none</td><td>No authentication is performed.</td></tr><tr><td>simple</td><td>Simple authentication is performed, using the Bind DN and Bind password as credentials. The credentials are sent as plain text.</td></tr><tr><td>DIGEST-MD5</td><td>Simple authentication is performed, using the Bind DN and Bind Password as credentials. The Bind password is hashed with the MD5 algorithm.</td></tr><tr><td>TLS-SIMPLE</td><td>A temporary secured TLS connection is set up before the credentials are sent as plain text. SSL must be configured.</td></tr><tr><td>TLS-EXTERNAL</td><td>A temporary secured TLS connection with external SASL authentication using a client certificate. SSL must be configured.</td></tr></table>	Authentication type	Explanation	none	No authentication is performed.	simple	Simple authentication is performed, using the Bind DN and Bind password as credentials. The credentials are sent as plain text.	DIGEST-MD5	Simple authentication is performed, using the Bind DN and Bind Password as credentials. The Bind password is hashed with the MD5 algorithm.	TLS-SIMPLE	A temporary secured TLS connection is set up before the credentials are sent as plain text. SSL must be configured.	TLS-EXTERNAL	A temporary secured TLS connection with external SASL authentication using a client certificate. SSL must be configured.
	Authentication type	Explanation											
	none	No authentication is performed.											
	simple	Simple authentication is performed, using the Bind DN and Bind password as credentials. The credentials are sent as plain text.											
	DIGEST-MD5	Simple authentication is performed, using the Bind DN and Bind Password as credentials. The Bind password is hashed with the MD5 algorithm.											
	TLS-SIMPLE	A temporary secured TLS connection is set up before the credentials are sent as plain text. SSL must be configured.											
TLS-EXTERNAL	A temporary secured TLS connection with external SASL authentication using a client certificate. SSL must be configured.												
Shutdown gracefully	✔ True: The LDAP context is destroyed immediately. When using TLS, some servers require the connection to be shut down by the client before the LDAP context is destroyed. ✘ False (default): The LDAP context is not destroyed immediately.												

Setting	Description								
Referral Setting	Specifies what to do with referrals. Possible values: <table> <tr> <th>Referral setting</th><th>Explanation</th></tr> <tr> <td>throw</td><td>Throws an exception if a referral is encountered.</td></tr> <tr> <td>ignore (default)</td><td>All referrals are ignored.</td></tr> <tr> <td>follow</td><td>Follows the referral to the actual location of the entry on another server. This is recommended when using Microsoft Active Directory.</td></tr> </table> Note If you are experiencing slow searches on Microsoft Active Directory with the <i>follow</i> value for the Referral setting, try using the Global Catalog as Active Directory domain controller. The Global Catalog enables searching for Active Directory objects in any domain in the forest without the need for subordinate referrals. This can dramatically speed up searching. However, the Global Catalog only contains a subset of the attributes of an object. This solution is only viable if the attributes requested for the search results are stored in the global catalog. Note that the Global Catalog is accessible on port 3268/3269, not the standard 389/636 LDAP ports.	Referral setting	Explanation	throw	Throws an exception if a referral is encountered.	ignore (default)	All referrals are ignored.	follow	Follows the referral to the actual location of the entry on another server. This is recommended when using Microsoft Active Directory.
Referral setting	Explanation								
throw	Throws an exception if a referral is encountered.								
ignore (default)	All referrals are ignored.								
follow	Follows the referral to the actual location of the entry on another server. This is recommended when using Microsoft Active Directory.								
Group base DN	The base DN of where all the groups are located. If a base has been specified, that base is used as the prefix for this group base.								
Group LDAP filter	The LDAP filter to which each group has to comply to be synchronized.								
Batch synchronization	The synchronization of the users with the LDAP server happens in batches.								
Batch size	The number of users in each batch. If a batch fails, none of the users in that batch is updated and the user names are listed in the DGC service log. Other batches are processed as normal. After processing all batches, Collibra disables users that are no longer in LDAP, unless one or more batches failed. Set the value to 0 to disable batch processing.								

Password

The configuration of passwords.

Setting	Description
Minimum length (Requires restart)	The minimum length of passwords. The default minimum length is 12.
Maximum length (Requires restart)	The maximum length of passwords. The default maximum length is 1,024.
Digits required (Requires restart)	✓ True (default): Passwords have to contain one or more digits. ✗ False: Passwords do not have to contain digits.
Non alpha-numeric required (Requires restart)	✓ True (default): Passwords have to contain one or more non-alpha-numeric (special) characters. ✗ False: Passwords do not have to contain non-alphanumeric characters.
Uppercase required (Requires restart)	✓ True (default): Passwords have to contain one or more upper-case characters. ✗ False: Passwords do not have to contain upper-case characters.
Lowercase required (Requires restart)	✓ True (default): Passwords have to contain one or more lower-case characters. ✗ False: Passwords do not have to contain lower-case characters.
Username disallowed (Requires restart)	✓ True (default): Passwords cannot be the username. ✗ False: Passwords can be the username.
Expiration interval (months)	The number of months before users have to change their passwords. Set it to 0 if users never have to change their passwords. The default interval is 6 months.
Allowed login failures	The number of consecutive failed login attempts that are allowed before the user account is disabled. Set it to 0 for unlimited attempts. The default is 3 login failures.
No reuse count	The number of previous passwords users cannot reuse. The default is 1: the user cannot change his password to what it currently is. Set this to 0 to allow using the same password.

Setting	Description
Password reset link validity period	The number of minutes that a link to reset a password remains valid. Beyond this time, the user has to request a new password reset link. The default value is 60 minutes. The minimum value is 15 minutes, the maximum value is 1,440 minutes (24 hours).
Account lock-out duration	The number of minutes that a user cannot sign in after too many failed sign-in attempts. If the number of minutes is set to 0, a Colibra administrator must reset the password to unlock the account. This setting is only applicable if the "Allowed sign-in failures" setting is defined. A locked-out account does not mean that your account is disabled.

REST

The security configuration of the REST interface.

Setting	Description
Limited CSRF	This option offers limited security, so we recommend upgrading to the Enhanced CSRF. ✓ True: The validity of a request is checked with a CSRF token. ✗ False (default): The validity of a request is not checked with a CSRF token.
Enhanced CSRF	If enabled, Colibra will check the validity of the request using a Spring Security CSRF token. ✓ True: The validity of a request is checked with a CSRF token. ✗ False (default): The validity of a request is not checked with a CSRF token.
Referrer enabled	✓ True: The HTTP referrer header is used to identify the origin of the request. ✗ False (default): The HTTP referrer header is not used to identify the origin of the request. It is recommended to leave this option disabled.

Setting	Description
Referrer checking allow empty	✓ True (default): The HTTP referrer header can be empty. ✗ False: The HTTP referrer header cannot be empty.

SSL

The configuration of SSL.

Setting	Description
Key store name	The name of the keystore file. The file is expected to be in the <collibra_data>/dgc/security folder.
Key store password	The password of the keystore.
Key store type	The type of the keystore file. For example, <i>JKS</i> or <i>PKCS12</i> .
Trust store name	The name of the truststore file. The file is expected to be in the <collibra_data>/dgc/security folder.
Trust store password	The password of the truststore.
Trust store type	The type of the truststore file. For example, <i>JKS</i> or <i>PKCS12</i> .

SSO

The configuration of Single Sign-On (SSO) authentication.

Setting	Description
Mode	The SSO mode of Collibra. The possible values are: • SAML_ATTRIBUTES • SAML_LDAP • SSO_HEADER • SSO_HEADER_LDAP • DISABLED
Header	The name of the header to be checked. The contents of this header is used for the search query, which is <i>SSO_HEADER = username</i>. The value of the actual query depends on DN and possibly Attribute.
DN	If the SSO mode is SSO_HEADER_LDAP or SAML_LDAP, this field determines whether the distinguished name (DN) or attribute is used: • ✓ True: The header has to contain the distinguished name (DN) of the user in the LDAP. • ✗ False (default): The header has to contain the value of Attribute. If the SSO mode is DISABLED, SSO_HEADER or SAML_ATTRIBUTES, this field is ignored.
Attribute	This field is only used if the SSO mode is SSO_HEADER_LDAP or SAML_LDAP, and if DN is False. If the above criteria are met, the LDAP has to contain this value. Example If Attribute is <i>FirstName</i>, then the header should contain the FirstName of the user that was signed in.
Disable automatic user creation when signing in via SSO	If users try to sign in via SSO, they still need a user account in Collibra. You can either create the user accounts automatically when they sign in, or create the user accounts manually or via LDAP synchronization • ✓ True: User accounts are not created automatically. • ✗ False (default): User accounts are created automatically.

Setting	Description
Disable the Collibra signin page	When SSO is enabled, a user can still navigate to the /signin page and try to log in via that page. However, you can disable that page. ✓ True: Users cannot access the Collibra signin page. ✗ False (default): Users can access the Collibra signin page
SAML	The configuration of SAML.
Metadata HTTP	The URL of the SAML metadata file to be used. The URL always has to be reachable by the Collibra environment.
Entity Provider Entity ID	The entity ID inside the metadata to be referenced. Note A metadata file can describe multiple entity IDs, make sure to use in the entity ID from the correct metadata file.
Attribute fields	The mappings of attributes in the SAML response. The values are used as keys to look for in the SAML response. Examples of attribute fields are first name, last name, address information, phone numbers and so on.
First name	The mapping for the user's first name. This attribute is optional. The value can be empty.
Last name	The mapping for the user's last name. This attribute is optional. The value can be empty.
Email	The mapping for the user's email address. This attribute is optional for existing users, but mandatory for new users. Warning If the email address is invalid when you synchronize, the user is deactivated and the user information is not updated.
Enabled	The mapping that indicates whether the account of the incoming user is enabled.

Setting	Description
Group	The mapping (attribute) which indicates to which Collibra groups the user should be added. If the groups don't exist yet, they will be created. This attribute can have multiple values (groups) or the groups can be sent as a comma-separated list of groups. If passing groups in this attribute, you must set Groups DGC Managed to <i>False</i>.
Phone	The mapping for the user's phone.
Fax	The mapping for the user's fax number.
Mobile	The mapping for the user's mobile number.
Pager	The mapping for the user's pager number.
Private	The mapping for the user's private number.
Work	The mapping for the user's work number.
Other	The mapping for any other phone number for this user.
Home address	The mapping for the user's home address.
Street	The mapping for the user's street.
Number	The mapping for the user's number.
City	The mapping for the user's city.
Post code	The mapping for the user's postal code.
State	The mapping for the user's state.
Country	The mapping for the user's country.
Work address	The mapping for the user's work address.
Street	The mapping for the user's street.
Number	The mapping for the user's number.

Setting	Description
City	The mapping for the user's city.
Post code	The mapping for the user's postal code.
State	The mapping for the user's state.
Country	The mapping for the user's country.
Instant messaging	The mapping for the user's IM locations.
AIM	The mapping for the user's AOL IM account.
Google Talk	The mapping for the user's Google Talk IM account.
Icq	The mapping for the user's ICQ IM account.
Jabber	The mapping for the user's Jabber IM account.
Messenger	The mapping for the user's Live Messenger IM account.
Skype	The mapping for the user's Skype IM account.
Yahoo Messenger	The mapping for the user's Yahoo Messenger IM account.
Gender	The mapping information for the user's gender.
Mapping	The attribute key for the gender value. If the content equals one of the male or female mappings, the user will be saved as male or female. Otherwise a default of <i>UNKNOWN</i> will be used.
Male value	The value for male users.
Female value	The value for female users.

Setting	Description
Groups DGC managed	Option to configure that groups should be managed by Collibra, or that groups should be set by the SAML assertion (SAML+Attributes mode). This option is only relevant if Mode is <code>SAML_ATTRIBUTES</code>. ✓ True: The groups are fully managed by Collibra. In the UI the admin has the option to assign groups to users, without it being overwritten by SAML. ✗ False (default): The groups are managed by the SAML assertions. In this case the groups are managed by the SAML IDP. Be sure to configure the Group attribute in the Attribute Fields section.
Service Provider Entity ID	Field that determines the value of the <code>Entity ID</code> parameter in the service provider metadata returned by Collibra. The default value is empty, in which case Collibra uses the value of the <code>Base URL</code> field. Enter a custom value if the base URL does not match the <code>audience</code> configured in your SAML identity provider. Warning The value of the <code>audience</code> restriction in the SAML response has to be exactly the same as the value of this field. Note SSO does not work if the <code>Service Provider Entity ID</code> field contains the base URL with trailing forward slash (for example <code>www.collibra.com/</code>), and the <code>audience</code> of your IDP contains the base URL without a trailing forward slash (for example <code>www.collibra.com</code>). Both values need to be exactly the same. In this case, you can resolve the issue by changing the value in the configuration of your IDP, or the value of this field. It does not matter whether both have a trailing forward slash or not, as long as they contain the same value.
Sign authentication requests (Requires restart)	✓ True: Authentication requests have to be signed. ✗ False (default): Authentication request don't have to be signed.
Force authn	✓ True (default): The SP authentication request forces re-authentication. ✗ False: The SP authentication request does not force re-authentication.

Setting	Description
Force passive	✓ True: The reauthentication has to happen in the background. ✗ False (default): The reauthentication does not have to happen in the background. This is only relevant if Force authn is <i>True</i>.
Name ID	Name ID that is used in the SP authentication. The default value is <i>urn:oasis:names:tc:SAML:2.0:nameid-format:persistent</i>. The Name ID value is mandatory.
Name ID allow create	✓ True (default): The IDP can create a name ID to fulfill the SP authentication request. ✗ False: The IDP cannot create a name ID to fulfill the SP authentication request.
Disable client address	✓ True: The validation of the client IP address in the assertion message is disabled. ✗ False (default): The validation of the client IP address in the assertion message is enabled.
SAML Requested authentication context	Settings for the SAML requested authentication context. The IDP uses the authentication context to authenticate the user. By default, the authentication context mandates user/password authentication over HTTPS.
Disable	✓ True: The requested authentication context section is not sent in the SAML request. ✗ False (default): The requested authentication context section is sent in the SAML request.
Comparison type	The comparison type that is transmitted in the requested authentication context. Possible values: <i>minimum</i> <i>maximum</i> <i>better</i> <i>exact</i> (default value) For more information about the comparison type values, refer to the SAML specifications.

Setting	Description
Reference list	The list of class references in the requested authentication context. You can separate list items with the pipe character (). For more information about this list, refer to the SAML specifications.
Declaration list	The list of class declarations in the requested authentication context. You can separate list items with the pipe character (). For more information about this list, refer to the SAML specifications.
Response decryption mode	Enable the support for encrypted SAML responses. • DISABLED: Collibra only accepts plain-text SAML responses. • OPTIONAL: Collibra can handle both encrypted and plain-text SAML responses. • FORCED: Collibra only accepts encrypted SAML responses. Once OPTIONAL or FORCED is selected, the encryption key pair is generated and added to the Collibra SAML keystore. A self-signed certificate is generated and works in most situations. If your IdP rejects self-signed certificates, you will have to add a certificate that is signed by a trusted 3rd party.
Validity period of the SAML certificate	The SAML certificate expiry date in years. By default, the SAML certificate expires after 20 years.

Signout

The configuration of redirecting after signing out of Collibra.

Setting	Description
Override signout URL (Requires restart)	• ☒ True: Redirect the user to a specific website after signing out. • ☐ False (default): Redirect the user to the sign-in page after signing out.
Signout redirect URL (Requires restart)	The URL to be redirected to when signing out.

Session

The configuration related to sessions.

Setting	Description
Idle Session timeout (Requires restart) This setting requires the SUPER role.	The time after which you are signed out if you are inactive. • Minimum value: 5 minutes • Maximum value: 24 hours • Default value: 1,800 seconds

Import/Export

The configuration to avoid the Formula Injection vulnerability in Excel.

Setting	Description
Escape Excel formulas	The option to disable Formula Injection into Excel. When enabling this option, an escape character is added at the beginning of Excel formulas during the export and is removed when importing formulas. The escape character will be added to fields that start with one of the following characters: • equation: = • plus: + • minus: - • at-sign: @ This option is enabled by default.
Excel formulas escape character	The escape character for Excel formulas when exporting or importing data.

JWT

The JSON Web Token configuration.

Setting	Description
JSON Web Key Set URL	The URL to retrieve public key information needed to verify the authenticity of JSON Web Tokens (JWTs), issued by an authorization server. This setting is required to enable JWT authentication.
JWT Token Types	A case-insensitive comma-separated list of accepted JWT media types coming in the typ header parameter. Leave blank if the authorization server does not provide a media type parameter. The default values is at+jwt,jwt.
JWT Algorithms	A comma-separated list of accepted JWT algorithms coming in the alg header parameter. See https://tools.ietf.org/html/rfc7518#section-3.1 for details. Leave blank to accept all digital signature algorithms.
JWT Issuer	The accepted issuer coming in the iss JWT claim. Leave blank if the authorization server does not provide an issuer claim.
JWT Audience	A comma-separated list of accepted audience values for the aud claim. The value for this field is a configuration setting in your authorization server, which identifies your Collibra environment as the intended recipient of the JWT. Leave blank if the authorization server does not provide an audience claim.
JWT Principal ID Claim Name.	The name of the JWT claim containing the principal's identity. See https://tools.ietf.org/html/rfc7519#section-4.1.2 for details. Defaults to the standard subject claim, sub. Change this setting only if your authorization server has other means of identifying the principal, for example, a client_id claim. This setting is required if JWT authentication is enabled.
JWT Maximum Clock Skew	The maximum acceptable difference in seconds between the clocks of the machines running the authorization server and Collibra. Differences smaller than the given amount are ignored when performing time comparisons for token validation. The default value is 60 seconds if left blank.

HTTP headers

The configuration of the HTTP headers

Field	Description
URL pattern This setting requires the SUPER role.	The pattern of the URLs to which the HTTP response header is applied. This field supports wildcards such as <code>**</code>, <code>*</code>, and <code>?</code>. Tip The following pattern matches all URLs: <code>/*</code>
HTTP headers This setting requires the SUPER role.	The HTTP response headers in a key-value format. You can add new HTTP response headers by clicking Add at the bottom of the section, and entering the HTTP response header name as the field key and the HTTP response header value as the field value.

Whitelists

The configuration for whitelist placeholders that can be used in security headers.

Option	Description
connect-src whitelist	The 'connect-src' whitelist. To use this whitelist in a security header, use the '{connectSrcWI}' placeholder.
font-src whitelist	The 'font-src' whitelist. To use this whitelist in a security header, use the '{fontSrcWI}' placeholder.
frame-src whitelist	The 'frame-src' whitelist. To use this whitelist in a security header, use the '{frameSrcWI}' placeholder.
img-src whitelist	The 'img-src' whitelist. To use this whitelist in a security header, use the '{imgSrcWI}' placeholder.
script-src whitelist	The 'script-src' whitelist. To use this whitelist in a security header, use the '{scriptSrcWI}' placeholder.

Option	Description
style-src whitelist	The 'style-src' whitelist. To use this whitelist in a security header, use the '{styleSrcWI}' placeholder.
frame-ancestors whitelist	The 'frame-ancestors' whitelist. To use this whitelist in a security header, use the '{frameAncestorsWI}' placeholder.
Tableau frame-ancestors whitelist	The tableau 'frame-ancestors' whitelist. To use this whitelist in a security header, use the '{tableauFrameAncestorsWI}' placeholder.

Disclaimer

The configuration of a disclaimer upon signing in to Colibra.

Setting	Description
Disclaimer	✓ True: Upon signing in, show a disclaimer that you have to agree with before you can continue. ✗ False (default): Don't show a disclaimer.
Disclaimer message	The disclaimer message that is shown after signing in. If you leave this field empty, there is a default message. You can use basic html tags, such as headers, paragraphs, images and hyperlinks.

Workflow engine configuration

The configuration of the workflow engine.

Setting	Description
Activate default escalation This setting requires the SUPER role.	✓ True (default): Automatically add an escalation function to user tasks after a configurable period of time. ✗ False: Don't escalate task automatically. This option only works for tasks that do not yet have a configured escalation path configured in Colibra Data Intelligence Cloud.

Setting	Description
Default escalation timer duration This setting requires the SUPER role.	The duration before a task is escalated. For more information on the format, please refer to ISO 8601 .
Activate default task email notification This setting requires the SUPER role.	✓ True (default): Send email notifications to any candidate user of a user task in the workflow engine. ✗ False: Do not send email notifications to candidate users.

Collibra Connect

The configuration to communicate with Collibra Connect.

Setting	Description
Base URL	The URL to Collibra Connect.
Username	The username to connect to Collibra Connect.
Password	The password to connect to Collibra Connect.

Register data source

Global parameters that apply to Data Source Registration.

Setting	Description
Table types to ignore	A comma separated list of table types that are not ingested. For example, <i>INDEX</i> and <i>SEQUENCE</i> .

Setting	Description
AWS regions restriction	A list of AWS regions Data Catalog is allowed to connect to. For example, <i>eu-west-3</i> and <i>us-east-2</i>. For a list of all AWS locations, see the AWS documentation. • If you want to allow Collibra to make a connection to any AWS region, leave the field empty. • If you remove a region from this list and the region was previously used for an S3 integration, you may want to delete the Glue database from the previously used region manually. By default, Collibra does not remove it. The Glue database has the following naming convention: <code>collibra_catalog_<Asset Id>_<Domain Id></code> For example: <code>collibra_catalog_d3174a88-5ffe-4d50-8fbe-7bf0832ec3af_5d198ce9-4e56-4d0e-a885-58204da50741</code> • When using Edge, a warning is added to the logs if an invalid region is detected in the restricted regions list.
AWS API call rate	Allowed number of AWS API calls per second. Use this option to limit the number of API calls per second to prevent throttling errors from the AWS API.
Database registration via Edge	An option to enable database registration via Edge. • ✓ True: Register a data source via Edge. • ✗ False: Register a data source via Jobserver only. Note Enabling data source registration via Edge does not prevent you from registering a data source via Jobserver as well.
Collibra Data Quality & Observability Synchronization UI via DQ Connector on Edge	An option to enable the Data Quality extraction interface in Collibra • ✓ True: The Quality extraction tab is available on the configuration page of a database asset • ✗ False (default): The Quality extraction tab is not available and as such, it is not possible to extract and synchronize data quality information. You can only enable Collibra Data Quality & Observability synchronization if you also enabled Database registration via Edge.

Setting	Description
Google Cloud Storage synchronization via Edge	An option to enable Google Cloud Storage file system registration and synchronization via Edge. ✓ True: You can register and synchronize a Google Cloud Storage file system via Edge. ✗ False: You can't register a Google Cloud Storage file system via Edge.
Amazon S3 synchronization via Edge	An option to enable Amazon S3 file system registration and synchronization via Edge. ✓ True: You can register and synchronize an Amazon S3 file system via Edge. ✗ False: You can only register an Amazon S3 file system via Jobserver. Note Enabling the registration of an Amazon S3 file system via Edge does not prevent you from registering an Amazon S3 file system via Jobserver. For more information, see Working with Amazon S3.
Databricks Unity Catalog synchronization via Edge	An option to enable the integration of Databricks Unity Catalog via Edge. ✓ True: You can register and synchronize Databricks Unity Catalog via Edge. ✗ False: You cannot integrate Databricks Unity Catalog. Set this option to True.
Source Tags Synchronization via Edge	An option to register and synchronize the tags on individual columns and tables in the data source during the data source registration via Edge. This means the tags become available and searchable in Data Catalog and can be used in business or policy processes and in workflows. Note Currently, you can synchronize source tags only from Snowflake. ✓ True: The Include Source Tags option becomes available when you define the rules for the synchronization of a data source via Edge. If you include the source tags, the tags defined on the assets in the data source are registered and available from the Schema, Table, Database View, and Column assets in the Source Tags attribute. ✗ False: Source tags are not registered in Data Catalog.

Jobserver

The configuration of the Jobserver service.

Setting	Description
Jobserver list	The list of registered Jobserver instances.
Name	The name of the Jobserver as it will appear when you register a data source in Data Catalog. The name is a freely chosen name but it is recommended to only use alphanumeric characters and dashes, for example Jobserver-1. You will have to use this name as the ID of the gateway and in the address of this configuration.
Protocol	The protocol that is used for the communication between the Data Governance Center service and the Jobserver service. It is recommended to use HTTPS, especially if the services are hosted in different network segments.
Address	The address (IP address, URL, hostname) of the Jobserver.
Trusted server CA certificate	The certificate of the trusted CA needed to validate the server certificate. If blank, the default truststore will be used. The default truststore is defined in the SSL configuration section of the DGC service. The CA certificate of the server party (Jobserver).
Client certificate	The client certificate offered by the DGC service to the server. If blank, you cannot select mutual authentication as the Jobserver service authentication level.
Client private key	The private key of the DGC service's certificate.
Table profiling data size	The approximate maximum disk size of the data in MB that will be used to profile a table. The value cannot exceed 10,000.
Test connection timeout	This timeout is a time limit (in seconds) after which the connection test is stopped and a timeout error is shown. The default value is 60 seconds.

Data profiling

Profiling must be executed again after a change in this section.

Setting	Description
Maximum number of samples	The maximum number of samples you want to collect for a data source. The default value is 100. The maximum value is 1,000. This setting is specific to sample data .
Maximum value length	The maximum length of a value extracted during profiling or sampling. Additional characters are trimmed.
Default date pattern	The default format used to decode dates. It is the default pattern used for detecting dates when the Date Pattern and/or Time Pattern attribute is not specified in Column assets.
Default time pattern	The default format used to decode times. It is the default pattern used for detecting times when the Date Pattern and/or Time Pattern attribute is not specified in Column assets.
Default combined date and time pattern	The default format used to decode combined dates and times. It is the default pattern used for detecting combined dates and times when the Date Pattern and/or Time Pattern attribute is not specified in Column assets.
Empty values	A comma separated list of strings enclosed in double quotes. A value that matches one of those expressions is considered an empty value. Please note that a database null value is always considered an empty value, for example <code>""</code> , <code>"na"</code> and <code>"none"</code> .
Data type detection threshold	The percentage of matching Column values to reach for an Advanced Data Type to be considered a possible Data Type for that Column. This is expressed as a value between 0.0 and 1.0).

Setting	Description
Anonymize data (Jobserver)	An option to anonymize sensitive data. ✓ True: Content in columns with data type Text or Geo is removed or replaced by a random hash value before the profiling results are sent to the cloud. ✗ False (default): No content is removed or replaced by a random hash value. Tip For anonymization via Edge, see setting "Anonymize Edge profiling results for all data types".
Database profiling via Edge	An option to enable profiling and classifying of synchronized metadata via Edge instead of Jobserver. ✓ True: Profiling and classification via Edge. ✗ False: Profile via Jobserver and classify via the Data Classification Platform. Note You can enable Database profiling via Edge only if you also enabled Database registration via Edge.
Maximum duration of a profiling Edge job	The maximum time duration, in minutes, that a profiling Edge job can run before Data Profiling stops the job. The default value is 20,160 minutes, 2 days. You can increase this limit to a maximum of 4 days.

Setting	Description
Parallel database profiling via Edge	The maximum number of schemas that Edge can profile at the same time. By default, the value of this setting is 4. This means Edge processes four profiling jobs at a time. This can have a huge positive impact on the performance of the profiling activity. You can increase this number to a maximum of 16. Note • If you increase this number to more than four jobs, make sure that your Edge site resources are aligned with the extra requests it will receive. • If you decrease this number and the running number of jobs exceeds the limit, no job will be canceled. Instead, there won't be any room to schedule a new job until at least one running job is completed. Example The parallel schema profiling via Edge setting is set to 4. • For 1 database that contains 3 schemas, we will process all 3 schemas at the same time. • For 2 databases that contain 4 schemas in total, we will process all 4 schemas at the same time. • For 1 database that contains 8 schemas, we will start with 4 schemas and then proceed to the next ones as soon as a job is completed.
Anonymize Edge profiling results for all data types	Enable this option to anonymize all Edge profiling results stored in Collibra. • ✓ True: Profiling results via Edge are anonymized for all columns. • ✗ False (default): Profiling results via Edge are anonymized only for columns with the Text or Geo data type.
Calculate Data Similarity (Beta)	Enables the data similarity feature in your environment. • ✓ True: Extra algorithms can run during profiling via Edge allowing the calculation of data similarity scores. The data similarity scores are currently used in Data Marketplace to show similar Table assets. • ✗ False: (default): The beta feature is not enabled.

Beta features

The configuration of features in beta state.

Setting	Description
Tableau provisioning enabled	✓ True: Provisioning to Tableau is enabled. ✗ False (default): Provisioning to Tableau is disabled.
Max number of concurrent import jobs	The maximum number of import jobs that can be executed at the same time via the API. This is to avoid memory issues. Default value is 4, set to 0 if there is no limit.
Task sidebar	✓ True (default): Workflow tasks appear in the sidebar on both resource pages and the task management page. Task forms appear in the sidebar instead of dialog boxes. Users can seamlessly complete their tasks from the task management page and have a side-by-side view of the tasks and resource details on resource pages. ✗ False : Workflow tasks appear in the task bar on resource pages and in a sidebar on the task management page. Task forms appear in dialog boxes. The behavior is the same as with older versions of Collibra.
Settings landing enabled	✓ True (default): Show the new Settings landing page in your Collibra environment. ✗ False : Use the classic Settings page in your Collibra environment.
Allow access to the Workflow Designer	✓ True: Enable the Workflow Designer accessglobal permission which allows access to the Workflow Designer, a visual tool for creating process definitions. ✗ False (default): Disable the Workflow Designer access global permission.
Frontend enabled This setting requires the SUPER role.	✓ True: Enable the frontend (shell) in the application. This will be served under the /apps/ and select other endpoints if enabled. Enabling this option results in some new designs and may be required for new applications. ✗ False (default): Disable the frontend (shell) in the application.
Shell asset page enabled	✓ True: Enable the new asset page layout in the user interface. ✗ False (default): Disable the new asset page layout in the user interface.

Setting	Description
Unified Classification enabled	Enables the new Unified Data Classification method on Edge. ✓ True: The environment uses the new classification method, Unified Data Classification. This has an impact on the available data classes, the required capabilities, and the way you classify data. Note All existing data classes and classifications become unavailable. Tip A migration process will be available in a future release. This process will transfer all old data classes and classifications to the Unified Data Classification method. Old transferred data classes will need to be updated to include classification rules to work with the new automatic data classification method. ✗ False: (default): The beta feature is not enabled. If you deactivate the setting, you revert back to your previous data classification setup and the previously defined data classes and classifications become available again.
Collections	Enable this setting to activate the use of collections in Data Marketplace. ✓ True: Collections can be used in Data Marketplace. - Users can add an asset to a collection from an asset preview in Data Marketplace. They can also remove an asset from a collection from the asset preview. - Users can access and manage all their collections from an overview page via their avatar → Collections. ✗ False: (default): The beta feature is not enabled. - When disabled, collections are not accessible via the UI. - Existing collections are not removed.

Setting	Description
Uninterrupted Search	✓ True: The search function remains available and the current data is available for search, while the search index is rebuilding in the background. Search results, however, may not be up to date until the rebuild is complete. After the rebuild is complete, all new data and changes to the existing data become available for search. ✗ False (default): The search function becomes temporarily unavailable when the search index is rebuilding. Important If you enable or disable this setting, you need to rebuild the search index for the change to take effect.

[Design system] features

This section contains settings that control the appearance of certain products and features. These settings require the SUPER role and are therefore managed by Collibra.

Throttling

Throttling is a security mechanism where you can limit the number of requests per seconds to ensure security and performance of your environment.

Setting	Description
Collect metrics without throttling This setting requires the SUPER role.	✓ True (default): Apply throttle logic without actual throttling to collect metrics. ✗ False: No throttle logic is applied to collect metrics.

The throttling metrics are only used for evaluation by Collibra support engineers if you report significant performance loss of your environment. We only track the number of times the throttling limit is exceeded.

By default, the throttling limit is 100 API requests per second.

REST API version 1.0 throttling

Throttle configuration for REST API version 1.0.

Setting	Description
Throttling enabled This setting requires the SUPER role.	✓ True: REST API v1 throttling is enabled. ✗ False (default): REST API v1 throttling is disabled.
Number of requests This setting requires the SUPER role.	The number of allowed request for the configured number of seconds.
Number of seconds This setting requires the SUPER role.	The number of seconds during which the configured number of requests can be performed.

REST API version 2.0 throttling

Throttle configuration for REST API version 2.0.

Setting	Description
Throttling enabled This setting requires the SUPER role.	✓ True: REST API v2 throttling is enabled. ✗ False (default): REST API v2 throttling is disabled.
Number of requests This setting requires the SUPER role.	The number of allowed request for the configured number of seconds.
Number of seconds This setting requires the SUPER role.	The number of seconds during which the configured number of requests can be performed.

GraphQL throttling

Throttle configuration for GraphQL.

Setting	Description
Throttling enabled This setting requires the SUPER role.	✓ True: GraphQL throttling is enabled. ✗ False (default): GraphQL throttling is disabled.
Number of requests This setting requires the SUPER role.	The number of allowed request for the configured number of seconds.
Number of seconds This setting requires the SUPER role.	The number of seconds during which the configured number of requests can be performed.

Hibernate cache configuration

The configuration of the hibernate second level caching. Hibernate caching uses a buffer memory that lies between Colibra and your repository database. It stores recently used data in this buffer memory to reduce the number of requests to your database, thereby improving the performance of your environment.

Tip We recommend that you use the default values. If you choose to edit the values, contact the Colibra support department before doing so.

Setting	Description
Enabled This setting requires the SUPER role.	✓ True (default): Hibernate caching is enabled. ✗ False: Hibernate caching is disabled.

Setting	Description												
Configuration list This setting requires the SUPER role.	The list of associated cache configurations.												
<cache configuration> This setting requires the SUPER role.	The cache configurations. Click Add to create a new cache configuration. <table> <tr> <th>Setting</th><th>Description</th></tr> <tr> <td>Enabled</td><td> ✓ True: This cache configuration is enabled. ✗ False: This cache configuration is disabled. </td></tr> <tr> <td>Name</td><td> The code of the assets that you want to cache. The codes are displayed in the following table. Do not use other codes. </td></tr> <tr> <td>Max elements in memory</td><td>The maximum amount of elements in the memory.</td></tr> <tr> <td>Eternal</td><td> ✓ True: The cache configuration is eternal. ✗ False: The cache configuration is not eternal. This is the default value. We recommend that you do not change the default value. </td></tr> <tr> <td>Overflow to disk</td><td> ✓ True: Write data to disk if the cache is full. ✗ False: Do not write data to disk if the cache is full. This is the default value. We recommend that you do not change the default value. </td></tr> </table>	Setting	Description	Enabled	✓ True: This cache configuration is enabled. ✗ False: This cache configuration is disabled.	Name	The code of the assets that you want to cache. The codes are displayed in the following table. Do not use other codes.	Max elements in memory	The maximum amount of elements in the memory.	Eternal	✓ True: The cache configuration is eternal. ✗ False: The cache configuration is not eternal. This is the default value. We recommend that you do not change the default value.	Overflow to disk	✓ True: Write data to disk if the cache is full. ✗ False: Do not write data to disk if the cache is full. This is the default value. We recommend that you do not change the default value.
Setting	Description												
Enabled	✓ True: This cache configuration is enabled. ✗ False: This cache configuration is disabled.												
Name	The code of the assets that you want to cache. The codes are displayed in the following table. Do not use other codes.												
Max elements in memory	The maximum amount of elements in the memory.												
Eternal	✓ True: The cache configuration is eternal. ✗ False: The cache configuration is not eternal. This is the default value. We recommend that you do not change the default value.												
Overflow to disk	✓ True: Write data to disk if the cache is full. ✗ False: Do not write data to disk if the cache is full. This is the default value. We recommend that you do not change the default value.												

Name	Resource type	Recommendation for Max elements in memory
GR	Groups	Set a number that is a percentage of the total amount of groups.
CO	Communities	Set a number that is at least the total amount of communities if possible.

Name	Resource type	Recommendation for Max elements in memory
UR	Users	Set a number that is a percentage of the total amount of users.
VC	Domains	Use as many domains as possible in cache.
RP	Assets	Set a number that is a percentage of the total amount of assets.
CT	Asset Types	Set a number that is at least the total amount of asset types if possible.
VT	Domain Types	Set a number that is at least the total amount of domain types if possible.
TY	Attribute Types	Set a number that is at least the total amount of attribute types if possible.
BF	Relation Types	Set a number that is at least the total amount of relation types if possible.
ST	Status	Set a number that is at least the total amount of statuses if possible.

Default configuration

The default configuration for this cache.

Do not change any of the default values.

Setting	Description
Max elements in memory This setting requires the SUPER role.	Type the maximum amount of elements in the memory.
Eternal This setting requires the SUPER role.	✓ True: The cache configuration is eternal. ✗ False: The cache configuration is not eternal.

Setting	Description						
Time to idle (seconds) This setting requires the SUPER role.	Type the amount of seconds until the cache is idle.						
Time to live (seconds) This setting requires the SUPER role.	Type the amount of seconds until the cache starts.						
Overflow to disk This setting requires the SUPER role.	✓ True: The cache configuration overflows to the disk. When the cache has reached its maximum number of elements, the next elements will be stored on disk. ✗ False: The cache configuration does not overflow to the disk. This is the default value.						
Disk persistent This setting requires the SUPER role.	✓ True: The disk store is persistent between CacheManager instances. ✗ False: The disk store is not persistent between CacheManager instances. This is the default value. This option is only relevant if Overflow to disk is set to <code>True</code>.						
Disk expiry thread interval (seconds) This setting requires the SUPER role.	Enter the amount of seconds between runs of the disk expiry thread. This value is how often we check for expiry. This option is only relevant if Overflow to disk is set to <code>True</code>.						
Memorystore eviction policy This setting requires the SUPER role.	Enter the policy how items are deleted from disk: <table border="1"> <tr> <td>LRU</td><td>Least recently used items are removed as first from disk. This is the default value.</td></tr> <tr> <td>LFU</td><td>Least frequently used items are removed as first from disk.</td></tr> <tr> <td>FIFO</td><td>First in first out principle to remove items from disk.</td></tr> </table> This option is only relevant if Overflow to disk is set to <code>True</code>.	LRU	Least recently used items are removed as first from disk. This is the default value.	LFU	Least frequently used items are removed as first from disk.	FIFO	First in first out principle to remove items from disk.
LRU	Least recently used items are removed as first from disk. This is the default value.						
LFU	Least frequently used items are removed as first from disk.						
FIFO	First in first out principle to remove items from disk.						

Setting	Description
Statistics This setting requires the SUPER role.	This option is not used by a Collibra environment.

Graph query

The configuration of the Graph query engine which is used to retrieve data from the repository.

For the general Graph query settings in a cloud environment, you need the SUPER role. Contact [Collibra Support](#) if you want to edit these settings.. For on-premises environments, you can edit the settings yourself.

The Graph query settings are not available in on-premises environments.

Setting	Description
work_mem setting for output module queries This setting requires the SUPER role.	A custom amount of memory that is reserved for the output module SQL queries. This setting should only be used for diagnosing potential lack of memory in case of performance issues. Performance issues may arise when large sorting or large joins are needed. By default, this option is disabled.
Enable simple joins This setting requires the SUPER role.	✓ True: Enable an alternative logic in the query engine to execute join operations. The alternative logic improves the query engine performance if queries don't have "to-many" relations. ✗ False (default): Keep the default logic of the query engine.
Enable joins for view permission filtering This setting requires the SUPER role.	✓ True: Enable an alternative logic to calculate view permissions in output module queries to improve the performance. ✗ False (default): Keep the default logic of the output module.

Graph query limits

Setting	Description
Enables limiting of the number of root nodes in result	✓ True: Enable limiting the number of root elements as result of a Graph query. ✗ False (default): Disable limiting the number of root elements as result of a Graph query.
Maximum number of root nodes that can be requested with graph query API	The maximum number of root nodes that you can request in the view configuration of an API call (REST or workflow). If you exceed this value in the view configuration, an exception is shown. If no value is defined in the view configuration, then the default value is taken. The default value is 100,000. Note If the number of asset types or domain types exceeds the set number, the hierarchy will be incomplete. Make sure that the limit is always higher than the actual number of asset and domain types.
Maximum number of nodes that can be requested with the graph query API in a single page	The maximum number of both root and children nodes that can be requested through Output API in a single data page. If the value is outside of the allowed range an exception is thrown. The default value is 1 million.

Graph query timeouts

Setting	Description
Maximum number of minutes a graph query can run	The maximum number of minutes that the graph query runs before it will time out. The maximum is 1,440 minutes (1 day). The default value is 480.

Table

The configuration of tables.

Setting	Description
Time limit for loading data in tables in seconds	The time limit after which a table stops loading on a page. Example A value of <i>600</i> means that if a table hasn't loaded within 600 seconds, the task is canceled and a timeout error is shown. The default value is 60, the maximum value is <i>720</i> seconds.

Multi-column sort

The configuration of multi-column sorting.

Setting	Description
Multi-column sorting on tables	✓ True: Tables can be sorted on multiple columns. ✗ False (default): Tables can be sorted on one column.
Number of columns available for multi-sort	Type the maximum number of columns that can be used to simultaneously sort tables. The default value is 3, the minimum is 1, the maximum is 9. This setting is only relevant if Multi-column sorting on tables is <i>True</i>.

Inherited responsibilities

Setting	Description
Enable Inherited Responsibilities	✓ True: Shows both direct and inherited responsibilities when filtering or displaying assets in table and tile views. ✗ False (default): Shows only direct responsibilities when filtering or displaying assets in table and tile views. Note - This setting affects only asset views and tile sets. It does not affect the Responsibilities tab of asset pages. - If this setting is set to True, opening a table or tile view may take longer depending on factors such as asset count, selected columns, and applied filters.

Purge configuration

The configuration of the automatic purging of data from the repository database. Purging means to delete data of a specified age. This helps to keep your data relevant and keep the database from growing infinitely.

Setting	Description
Purge schedule (Requires restart)	A Cron expression specifying the timing and frequency of purge cycles. The default scheduled time is 002 * *, which equates to 02:00 every day. If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.
Maximum time for each purge cycle (Requires restart)	Maximum amount of time (in seconds) allowed for each purge cycle. The default value is 7,200, which is two hours. Any qualifying data that is not purged in the allowed time will be addressed in a subsequent purge cycle, which picks up where the previous cycle left off.

Setting	Description
List of data elements and age at which each will be purged (Requires restart)	The data elements that will be purged at the specified age, in a key-value format: • The data elements to be purged (Field key). Possible data elements: ◦ Statistics: Data that has been used to calculate data quality. ◦ Authentication events: Data about authentication in your environment. ◦ Validation results: Information about data validation. ◦ Jobs: Data about all jobs that are created in your environment. ◦ Workflows: Data about completed workflow instances and tasks. ◦ License usage: Data about how licenses are used in your environment. • The age (in months) at which each individual data element will be purged (Field value).
Enable removal of orphaned tags (Requires restart)	Option to enable the automatic deletion of tags that are not assigned to any assets. • ✓ True (default): Orphan tags are deleted according to the timing and frequency that you specify. • ✗ False: Orphan tags are not deleted.
Orphaned tags removal schedule (Requires restart)	A Cron expression specifying the timing and frequency of the deletion of orphan tags. The default scheduled time is 001* *, which equates to 01:00 every day. If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.

Cloud Data Classification configuration

With data classification you can automatically assign data classes to ingested data.

Note In a Collibra Data Intelligence Cloud environment, you have to create a support ticket to configure this feature.

Setting	Description
Machine Learning platform URL This setting requires the SUPER role.	The address of the machine learning platform that will classify your data.
Requester Name This setting requires the SUPER role.	The unique name to identify the client when using Machine Learning platform.
API key This setting requires the SUPER role.	The API Key to authorize the requester when connecting to the Machine Learning platform.
Enable Data Classification	✓ True: Enable Colibra's data classification technology. ✗ False (default): Do not use Colibra's data classification technology are not accepted.

Classification thresholds

Setting	Description
Enable automatic classification acceptance and rejection	✓ True: The automatic acceptance and rejection of data classification suggestions is active. ✗ False (default): Data classification suggestions are not automatically accepted or rejected. Tip Start using the Automatic Data Classification tool by manually accepting and rejecting the data classification suggestions. Only activate the automatic acceptance and rejection feature if you are comfortable with the results the tool provides.

Setting	Description
Automatic acceptance threshold	The percentage from which data classification suggestions must be accepted automatically. If you set this value to 75, then the classification suggestions with a confidence level of 75% or higher are automatically accepted. If multiple classification suggestions meet the threshold condition for a column, the classification suggestion with the highest confidence level percentage is accepted automatically if this classification suggestion is the only one to have that confidence level percentage. Example You set the automatic acceptance threshold to 85%. You classify a table with 2 columns. - For column A, three classification suggestions are possible, one with confidence level 93%, one with 92%, and one with 90%. - For column B, two classification suggestions are possible. Their confidence level is the same, 86%. The results of the automatic acceptance will be: - For column A, the classification suggestion with 93% will be accepted automatically. - For column B, nothing is done, both suggestions will be visible. The default acceptance threshold is 90.
Automatic rejection threshold	The percentage from which data classification suggestions must be rejected automatically. If you set this value to 49, then all data classification suggestions with a confidence level of 49% or lower are automatically rejected. The default rejection threshold is 10.

Note If the acceptance threshold and rejection threshold are set to the same value, and a data classification suggestion has this confidence level percentage, the classification suggestion will be rejected.

Reporting

For more information about these settings, go to [Insights Data Access](#).

Setting	Description
Cloud Provider	Cloud provider - AWS or GCP
Customer GUID	The GUID of your Collibra environment. Note This field is configured by Collibra Cloud Ops.
Insights download bucket name	The name of the AWS S3 bucket in which your reporting data is stored. Note This field is configured by Collibra Cloud Ops.
Insights AWS S3 Region	The AWS S3 region in which your data is processed. Note This field is configured by Collibra Cloud Ops.
Insights zip location pattern	A pattern with the format "/zip/insights_%s.zip", where "%s" is replaced by the Collibra Insights snapshot date. Note This field is configured by Collibra Cloud Ops.
Tableau report URL pattern	The Tableau URL pattern, which should contain {reportName}. Tip You can paste the URL from the Link field in Tableau, as described in Generate the dashboard reports you configured in Collibra Data Intelligence Cloud Settings.
Reports definitions	
Report view name	The report name, as you want it to appear on the report button in the Usage Analytics widget, for example "Data Maturity Dashboard".
Report name	The report name, as it appears in the URL of the Tableau report, for example "DataMaturityDashboard".

Catalog Experience

Data Catalog Experience improves the layout of Data Catalog's asset pages.

Setting	Description
Enable Catalog experience	✓ True: Catalog experience is enabled. This will improve the layout of Data Catalog's asset pages, such as those of Data Set, Schema, Table and Column assets. ✗ False: Catalog experience is disabled.
Catalog Experience Titlebar theme	The theme for the Catalog experience. You can choose between the LIGHT and DARK. This option is only applicable if the Enable Catalog experience option is enabled.

Data Notebook features

Setting	Description
Data Notebook enabled This setting requires the SUPER role.	✓ True: The Data Notebook feature is available in the shell. ✗ False (default): The Data Notebook feature is not available in the shell.

Diagrams

These settings determine dialog loading time and size limits.

Setting	Description
Maximum loading time for the backend	The time limit, in seconds, after which a diagram stops fetching data. The value must be a positive integer and cannot be greater than 3,600 (one hour). The default value is 300. Example A value of <i>300</i> means that if a diagram hasn't fetched all data within 300 seconds, the diagram stops fetching data and an empty diagram with a notification is shown.
Size limit for the backend	The maximum number of nodes plus edges that will be fetched by the backend, to build a diagram. The value must be a positive integer and cannot be greater than 100,000. The default value is 10,000. Example A value of <i>10,000</i> means that if the total number of nodes plus edges is greater than 10,000, the diagram does not load and a notification is shown.
Size limit for the frontend	The maximum number of visible nodes plus edges that can be shown on the page. The value must be a positive integer and cannot be greater than 10,000. The default value is 2,000. Example A value of <i>2,000</i> means that if the total number of visible nodes and edges is greater than 2,000, the diagram does not load and a notification is shown.
Maximum flow depth	The system-wide maximum number of flow relations between the start node and any other diagram node. The value must be an integer between 1 and 100. The default value is 50. Note • If the maximum flow depth is specified in the selected diagram view, that value supersedes the maximum you specify here. • You can also manually adjust the flow depth in the diagram.

Setting	Description
Diagrams Business Qualifier Filter (*)	✓ True: Users can filter diagrams by a specified Business Qualifier asset. ✗ False (default): Users are unable to filter diagrams by Business Qualifier.

Everywhere Desktop configuration

These settings determine some of the ways in which Collibra for Desktop interacts with Collibra Data Intelligence Cloud.

Note These settings are only applied in Collibra for Desktop if you have Collibra 2021.01 or newer in combination with Collibra for Desktop 1.2.1 and newer.

Setting	Description
Default search filter	The filter that is applied, by default, to search results. The value must be the UUID of the filter. To find the UUID, open the Collibra environment and click in the Search box. Click the name of a search filter. In the address bar you will see the UUID of the filter. Note Specifying a default search filter in the application will override the default filter that you specify here.
Custom Search box placeholder	Placeholder text that appears in the Search field before a user enters search text. The default text is "Search in Collibra".
Shortcut Search	Enable or disable the use of a keyboard shortcut to search for selected text in Collibra Data Intelligence Cloud from within your browser or another application. ✓ True (default): Users can use the keyboard shortcut you specify in the following setting to search in Collibra Everywhere. ✗ False: Keyboard shortcut is disabled.

Setting	Description
Custom Shortcut Search	The keyboard shortcut to search for selected text in Collibra Data Intelligence Cloud from within your browser or another application. Tip The keyboard shortcut has to be a combination of Control, Alt or Shift with one letter or number. On macOS you can also use the Command key. Note - To make available the keyboard shortcut, you have to enable the feature in the previous setting. Specifying a keyboard shortcut in the application will override the shortcut that you specify here.
Enable Auto Hyperlinking	Option to enable automatic hyperlinking within Collibra for Desktop. With this option enabled, the name of an asset automatically becomes a hyperlink when you fill out a text attribute. This option only works if the Enable hyperlinking option in Collibra Console is also enabled.
Enable Workflows	Option to enable workflows in Collibra Everywhere. This allows you to complete tasks or start a workflow in the app. The available workflows depend on the ones that you add to the Global workflows and Asset workflow configuration.
Recommender	The Recommender helps users by suggesting relevant business assets and data sets, based on certain relation types and the past actions of similar users. ✓ True: Recommender is enabled. ✗ False (default): Recommender is disabled. This feature only works if Analytics is enabled. You can enable Analytics in section 1 General of the DGC service configuration.

Setting	Description
Auto-updater	Option to automatically upgrade Collibra Everywhere when a new version is available. ✓ True (default): Collibra Everywhere is automatically upgraded when a new version is available. ✗ False: You need to manually upgrade Collibra Everywhere when a new version is available. Note If you enable automatic updates, you have to whitelist the S3 bucket collibra-otg-desktop-installers in the region eu-west-1.
Allow User Configuration	Option to allow users to edit personal settings in Collibra Everywhere. ✓ True (default): Users can edit personal settings. ✗ False: Users cannot edit personal settings.
Global workflows	The list of workflows that is available in the app's main menu. Enter the UUIDs of the workflows. An example workflow could be "Create issue".
Asset workflow	The list of workflows that is available on an asset page in the app. Enter the UUIDs of the workflows. An example workflow could be "Ask the expert".
No search result workflows	The workflows that are available if there are no search results found. Enter the UUIDs of the workflows. An example workflow could be "Propose new business term".
Enable autostart	Option to automatically start Collibra for Desktop when signing in to your operating system. ✓ True: The app starts automatically when signing in to your operating system. ✗ False (default): The app does not start automatically. If you have set this option in the Collibra for Desktop settings, this option is neglected.

Everywhere Mobile configuration

These settings determine some of the ways in which Collibra for Mobile interacts with Collibra Data Intelligence Cloud.

Setting	Description
Default search filter	The filter that is applied, by default, to your search results. The value must be the UUID of the filter. This filter overrules any search filter that is set in the app. To find the UUID, open the Collibra environment and click in the Search box. Click the name of a search filter. In the address bar you will see the UUID of the filter.
Custom Search box placeholder	The text that is shown in the search box of Collibra for Mobile before a user enters search text.
Enable Workflows	Option to enable workflows in Collibra for Mobile. This allows you to complete tasks or start a workflow in the app. The available workflows depend on the ones that you add to the Global workflows and Asset workflow configuration.
Global workflows	The list of workflows that is available in the app's main menu. Enter the UUIDs of the workflows. An example workflow could be "Create issue".
Asset workflow	The list of workflows that is available on an asset page in the app. Enter the UUIDs of the workflows. An example workflow could be "Ask the expert".
No search result workflows	The workflows that are available if there are no search results found. Enter the UUIDs of the workflows. An example workflow could be "Propose new business term".

Collibra Browser Extension configuration

The settings determine how and where you can use the Everywhere Chrome Extension.

Setting	Description
Domains	Add a web domain , for example of a web application such as Power BI or Tableau, on which the Browser Extension automatically appears as overlay.

Edge

Edge configuration options, when changing an option, you only have to refresh the page that runs your Collibra environment.

Setting	Description
Enable Edge jobs feature (beta)	✓ True: Enable the Edge jobs page, this page gives you a overview of all jobs and their status. ✗ False (default): Disable the Edge jobs page.

Tableau Metadata API

You need the [Tableau metadata API](#) to ingest Tableau 2020.2 and newer.

Warning If you upgrade to Tableau version 2020.2 or newer, but previously synchronized an older Tableau version via the REST API and XML mapping, you have to prepare the [migration procedure](#) to prevent losing manually added relations, attributes, tags, comments and stitching results.

Setting	Description
Enable Tableau metadata API	✓ True: Tableau metadata API is enabled. This enables you to ingest Tableau 2020.2 or newer into Data Catalog. ✗ False: Tableau metadata API is disabled. If you ingest Tableau 2020.2 or newer, the ingestion will fail. This prevents data loss of manually added relations and attributes.

Job Service (Activities)

Setting	Description
Number of executor threads for the Job Service	The maximum number of threads, or jobs, that the Job Service can run in parallel. Generally speaking, increasingly the number of jobs running in parallel reduces overall processing time. Conversely, it requires more system resources, which can negatively impact performance. It also increases the risk of job conflicts.

Identity

Setting	Description
Limit user information access	✓ True: Limits user access to information related to other users. ✗ False (default): Users have access to information related to other users. Note Users that have a role with the System or User administration permission have full access regardless of this setting.

Lineage on Edge

Setting	Description
DGC user name	The DGC user that is used to ingest technical lineage data into the environment via the technical lineage servers.
DGC user password	Password of the DGC user that is used to ingest technical lineage data into the environment via the technical lineage servers.
Collibra system name flag	Enable this option if Lineage uses a Collibra system name.
Power BI and Tableau synchronization enabled	✓ True: Enable ingestion and technical lineage for Power BI and Tableau via Edge. ✗ False (default): Disable ingestion and technical lineage for Power BI and Tableau via Edge.

Setting	Description
Technical lineage via JDBC connection enabled	✓ True: Enable technical lineage for data sources by using JDBC connections via Edge. For an overview of the supported data sources, go to the Technical lineage documentation. ✗ False (default): Disable technical lineage for data sources by using JDBC connections via Edge.
Technical lineage for SQL via folder, ETL Tools, and custom lineage on Edge enabled	✓ True: Enable custom technical lineage and technical lineage for ETL tools and SQL data sources by using a Shared Storage connection. Technical lineage ingestion by using the Shared Storage connection type on Edge is equivalent to ingestion by using the folder connection type when you use the lineage harvester. For an overview of the supported data sources, go to the Technical lineage documentation. ✗ False (default): Disable custom technical lineage and technical lineage for ETL tools and SQL data sources by using a Shared Storage connection.

Collibra Protect

Setting	Description
Protect scheduler fixed delay	The number of minutes in between synchronizations. The default value is 60 minutes.

License configuration

Setting	Description
User license view schedule This setting requires the SUPER role.	This sets how often the license numbers on the user table are refreshed.

Setting	Description
License usage snapshot cron job schedule	This sets how often the license usage snapshot is refreshed. This cannot run at an interval smaller than 60 minutes.

Data Marketplace configuration

The configuration of the Data Marketplace.

Setting	Description
Data Marketplace	✓ True (default): Data Marketplace is enabled. Anyone with the required permissions can use or configure the Data Marketplace application from the Applications icon . Note When Data Marketplace is enabled and you reindex Collibra completely, the relations are also reindexed automatically. You don't need to start it manually. However, reindexing the relations will not reindex Collibra completely. ✗ False: Data Marketplace is not enabled. After you enable this setting, reindex Data Marketplace relations or reindex Collibra completely. Note In new Collibra environments, this setting is enabled by default. In upgraded Collibra environments, the previous status of this setting is retained.

Data Privacy

The configuration of the Data Privacy landing page.

Setting	Description
Privacy landing page	✓ True: Enables the Privacy landing page. ✗ False: Disables the Privacy landing page.

Export Configuration

Setting	Description
Detailed Relation Export Format This setting requires the SUPER role.	✓ True (default): The exported relation column headers contain 3 elements of the relation: [Source Asset Type], Relation, and [Target Asset Type]. ✗ False: The exported relation column headers contain only 2 elements: Relation and [Target Asset Type].

Slack Integration Configuration

Setting	Description
Slack Integration Enabled This setting requires the SUPER role.	✓ True: The Slack integration feature is enabled. ✗ False (default): The Slack integration feature is disabled.
Slack Integration URL This setting requires the SUPER role.	The URL of the Slack integration service.

Asset external link

Setting	Description
URL patterns to external asset pages	A list of URL patterns to external asset pages: • External system id: The external system ID as defined by the <code>externalSystemId</code> field of the import file used with the Colibra REST Import API. • External system label: The name that the external system has on the asset page. • URL pattern to external asset page: The pattern to construct the URL to external asset page, for example <code>https://external-system.com/entity/{externalEntityId}</code>. The <code>externalEntityId</code> is the one defined by the same field of the import file used with the Colibra REST Import API.

Appendix B - Spring Cron syntax

Cron is a software utility that specifies commands to run on a given schedule. This schedule is defined by a Cron pattern, which has a specific syntax that will be described in this section.

Warning If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.

Note By default, Collibra Console uses Spring Cron expressions to schedule backups, while you use [Quartz Cron expressions](#), for example, to schedule your mail, LDAP synchronizations, Purge cycles, Tableau and S3 synchronizations or to create a statistics cron map.

The Cron pattern consists of six space-separated fields:

<second> <minute> <hour> <day of month> <month> <day of week>

Position	Field	Allowed values	Allowed special characters	Examples
1	second	0-59	, - * /	10: at the 10th second. */10: every 10 seconds.
2	minute	0-59	, - * /	30: at the 30th minute. */15: every 15 minutes. 5/10: every 10 minutes starting at the 5th minute after the hour
3	hour	0-23	, - * /	10: at 10 o'clock. 8-10: at 8,9 and 10 AM. 6,18: at 6 AM and at 6 PM.
4	day of the month	1-31	, - * ? / L W	3: on the 3rd day of the month. 1-4: every first four days of the month. 1,15: the first day of the month and the 15th day of the month.

Position	Field	Allowed values	Allowed special characters	Examples
5	month	1-12 or JAN-DEC	, - * /	12: in December. 1-3: every first three months of the year. <i>JUL,AUG</i>: every July and August. Tip The names of the months are not case-sensitive.
6	day of the week	0-7 or MON-SUN where 0 and 7 is Sunday.	, - * ? / L #	<i>TUE</i>: every Tuesday. 1-5: every weekday, Monday to Friday. <i>MON,WED,FRI</i>: every Monday, Wednesday and Friday. <i>L</i>: only in combination with a digit or short day notation indicating the last day of the month. For example, 7L or SUNL indicates the last Sunday of the month, 3L or WEDL indicates the last Wednesday of the month. <i>5#3</i>: on the 3rd Friday of the month. Tip The names of the days are not case-sensitive.

For more information, see the [Spring Cron documentation](#).

Special characters

Character	Description
*	Used to select all values within a field. Example * in the minute field corresponds with every minute.

Character	Description
?	Used to specify something in one of the two fields in which the character is allowed, but not the other, mainly used for days of the week. Example If you want your trigger to fire on a particular day of the month, for example the 10th, but don't care what day of the week that happens to be, you could put "10" in the day-of-month field, and "?" in the day of the week field.
-	Used to specify ranges. Example 10-12 in the hour field means "the hours 10, 11 and 12".
,	Used to specify additional values. Example MON, WED, FRI in the day-of-week field means "the days Monday, Wednesday, and Friday".
/	Used to specify increments. Example 0/15 in the seconds field means "the seconds 0, 15, 30, and 45". And 5/15 in the seconds field means "the seconds 5, 20, 35, and 50". You can also leave out the number before /, which is equivalent to having 0 before /. 1/3 in the day-of-month field means "fire every 3 days starting on the first day of the month".
L	The value L in the day-of-month field means "the last day of the month" - day 31 for January, day 28 for February on non-leap years. You can also specify an offset from the last day of the month, such as "L-3" which would mean the third-to-last day of the calendar month. You cannot use L in the day-of-week field by itself, you always have to prefix it by a digit (0-7) or the short day notation (MON-SUN). This combination means "the last xxx day of the month" - for example "6L" or "SATL" means "the last Saturday of the month".

Character	Description
W	Used to specify the weekday (Monday-Friday) nearest the given day. Example 15W in the value for the day-of-month field, means the nearest weekday to the 15th of the month: • If the 15th is a Saturday, the trigger will fire on Friday the 14th. • If the 15th is a Sunday, the trigger will fire on Monday the 16th. • If the 15th is a Tuesday, then it will fire on Tuesday the 15th. However if you specify 1W as the value for day-of-month, and the 1st is a Saturday, the trigger will fire on Monday the 3rd, as it will not 'jump' over the boundary of a month's days. The 'W' character can only be specified when the value in the day-of-month field specifies a single day, not a range or list of days. Tip The 'L' and 'W' characters can also be combined in the day-of-month field to yield 'LW', which translates to "last weekday of the month".
#	Used to specify "the nth" XXX day of the month. Example 6#3 in the day-of-week field means "the third Saturday of the month" (day 6 = Friday and "#3" = the 3rd one in the month). Other examples: 2#1 is the first Tuesday of the month and 4#5 is the fifth Thursday of the month. Note that if you specify #5 and there is not 5 of the given day-of-week in the month, then no firing will occur that month.

Example

- 0 0 * * * * = the top of every hour of every day.
- */10 * * * * = every ten seconds.
- 0 0 8-10 * * * = 8, 9 and 10 o'clock of every day.
- 0 0 6,19 * * * = 6:00 AM and 7:00 PM every day.
- 0 0/30 8-10 * * * = 8:00, 8:30, 9:00, 9:30, 10:00 and 10:30 every day.
- 0 0 9-17 * * MON-FRI = on the hour nine-to-five weekdays.
- 0 0 0 25 12 ? = every Christmas Day at midnight, no matter what weekday it is.

Quartz Cron syntax

Cron is a software utility that specifies commands to run on a given schedule. This schedule is defined by a Cron pattern, which has a specific syntax that will be described in this section.

For example, you can create a schedule for LDAP synchronizations, Purge cycles or to automatically send emails using cron patterns. You can also use it to create a Cron map for your statistics.

Note By default, you use [Spring Cron expressions](#) to schedule Collibra Console back-ups.

Warning If you create an invalid Cron pattern, Collibra Data Intelligence Cloud stops responding.

The Cron pattern consists of six or seven space-separated fields:

<second> <minute> <hour> <day of the month> <month> <day of the week> <year>

Position	Field	Mandatory	Allowed values	Allowed special characters	Examples
1	second	Yes	0-59	, - * /	10: at the 10th second. */10: every 10 seconds.
2	minute	Yes	0-59	, - * /	30: at the 30th minute. */15: every 15 minutes. 5/10: every 10 minutes starting at the 5th minute after the hour
3	hour	Yes	0-23	, - * /	10: at 10 o'clock. 8-10: at 8,9 and 10 AM. 6,18: at 6 AM and at 6 PM.

Position	Field	Mandatory	Allowed values	Allowed special characters	Examples
4	day of the month	Yes	1-31	, - * ? / L W	• 3: on the 3rd day of the month. • 1-4: every first four days of the month. • 1,15: the first day of the month and the 15th day of the month. • L: on the last day of the month. • L-3: on the third-to-last day of the month. • 15W: on the nearest weekday to the 15th of the month. If the 15th is a Saturday, then the trigger will be on the 14th, if the 15th is a Sunday, then the trigger will be on the 16th. Note If the 1st day of the month is a Saturday, then 1W corresponds to the 3rd day of the month, since the month is specified in the 5th value of the Cron expression. LW: on the last weekday of the month.
5	month	Yes	1-12 or JAN-DEC	, - * /	• 12: in December. • 1-3: every first three months of the year. • JUL,AUG: every July and August. Tip The names of the months are not case-sensitive.

Position	Field	Mandatory	Allowed values	Allowed special characters	Examples
6	day of the week	Yes	1-7 or SUN-SAT	, - * ? / L #	<i>TUE</i>: every Tuesday. <i>2-6</i>: every weekday, Monday to Friday. <i>MON,WED,FRI</i>: every Monday, Wednesday and Friday. <i>L</i>: on Saturday, the 7th day of the week. <i>2L</i>: at the last Monday of the month. <i>6#3</i>: on the 3rd Friday of the month. Tip The names of the days are not case-sensitive.
7	year	No	empty, 1970-2099	, - * /	<i><empty></i>: if your schedule doesn't require a year, you can leave this value empty. <i>2021</i>: in 2021. <i>2021-2025</i>: in the years 2021, 2022, 2023, 2024 and 2025. <i>2021,2022,2025</i>: in the years 2021, 2022 and 2025.

Special characters

Character	Description
*	Used to select all values within a field. Example * in the minute field corresponds with every minute.

Character	Description
?	Used to specify something in one of the two fields in which the character is allowed, but not the other, mainly used for days of the week. Example If you want your trigger to fire on a particular day of the month, for example the 10th, but don't care what day of the week that happens to be, you could put "10" in the day-of-month field, and "?" in the day of the week field.
-	Used to specify ranges. Example 10-12 in the hour field means "the hours 10, 11 and 12".
,	Used to specify additional values. Example MON, WED, FRI in the day-of-week field means "the days Monday, Wednesday, and Friday".
/	Used to specify increments. Example 0/15 in the seconds field means "the seconds 0, 15, 30, and 45". And 5/15 in the seconds field means "the seconds 5, 20, 35, and 50". You can also leave out the number before /, which is equivalent to having 0 before /. 1/3 in the day-of-month field means "fire every 3 days starting on the first day of the month".
L	Has different meaning in each of the two fields in which it is allowed. The value L in the day-of-month field means "the last day of the month" - day 31 for January, day 28 for February on non-leap years. You can also specify an offset from the last day of the month, such as "L-3" which would mean the third-to-last day of the calendar month. If you use L in the day-of-week field by itself, it means "7" or "SAT". But if used in the day-of-week field after another value, it means "the last xxx day of the month" - for example "6L" means "the last Friday of the month". When using the L option, it is important not to specify lists, or ranges of values, because you may get unexpected results.

Character	Description
W	Used to specify the weekday (Monday-Friday) nearest the given day. Example 15W in the value for the day-of-month field, means the nearest weekday to the 15th of the month: • If the 15th is a Saturday, the trigger will fire on Friday the 14th. • If the 15th is a Sunday, the trigger will fire on Monday the 16th. • If the 15th is a Tuesday, then it will fire on Tuesday the 15th. However if you specify 1W as the value for day-of-month, and the 1st is a Saturday, the trigger will fire on Monday the 3rd, as it will not 'jump' over the boundary of a month's days. The 'W' character can only be specified when the value in the day-of-month field specifies a single day, not a range or list of days. Tip The 'L' and 'W' characters can also be combined in the day-of-month field to yield 'LW', which translates to "last weekday of the month".
#	Used to specify "the nth" XXX day of the month. Example 6#3 in the day-of-week field means "the third Friday of the month" (day 6 = Friday and "#3" = the 3rd one in the month). Other examples: 2#1 is the first Monday of the month and 4#5 is the fifth Wednesday of the month. Note that if you specify #5 and there is not 5 of the given day-of-week in the month, then no firing will occur that month.

Example

- `0 0 * ? * * *` = the top of every hour of every day.
- `*/10 * * * * ?` = every ten seconds.
- `0 0 8-10 * * ? 2020` = 8, 9 and 10 o'clock of every day during the year 2020.
- `0 0 6,19 ? * *` = 6:00 AM and 7:00 PM every day.
- `0 0/30 8-10 ? * *` = 8:00, 8:30, 9:00, 9:30, 10:00 and 10:30 every day.
- `0 0 9-17 * * MON-FRI` = on the hour nine-to-five weekdays.
- `0 0 0 25 12 ?` = every Christmas Day at midnight, no matter what day of the week it is.
- `0 15 10 ? * 6L 2022-2025` = 10:15 AM on every Friday of every month during the years 2022, 2023, 2024 and 2025.
- `0 30 11 ? * 6#2` = 11:30 AM on the second Friday of every month.

Warning Quartz Cron only supports a value in either the 4th or the 6th position, but not in both. At the same time, both positions cannot be empty.