DBeaver User Guide v.21.0

User Guide

Table of contents

General User Guide

- Installation
- Application Window Overview
- Views
  - Database Navigator
    - Filter Database Objects
    - Configure Filters
  - Projects View
  - Project Explorer
  - Query Manager
  - Background Tasks
- Database Object Editor
  - Properties Editor
- Data Editor
  - Navigation
  - Data View and Format
  - Data Filters
  - Data Refresh
  - Data Viewing and Editing
  - Panels
  - Managing Charts
  - Data Search
  - SQL Generation
  - Working with spatial/GIS data
  - Working with XML and JSON
  - Managing Data Formats
- SQL Editor
  - SQL Templates
  - SQL Assist and Auto-Complete
  - SQL Formatting
  - SQL Execution
  - Query Execution Plan

Visual Query Builder
Script Management
Client Side Commands
Debug
PostgreSQL Debugger
ER Diagrams
Database Structure Diagrams
Custom Diagrams
Search
File Search
DB Full-Text Search
DB Metadata Search
Schema compare
Data compare
MockData generation
Dashboards, DB monitoring
Projects
Project security
Team work (Git)
Bookmarks
Shortcuts

Database Management

Database Connections
Create Connection
Edit Connection
Connect to Database
Invalidate/Reconnect to Database
Disconnect from Database
Connection Types
Transactions
Auto and Manual Commit Modes
Transaction Log
Pending transactions

Database drivers
Tasks
Data export/import
Data migration
Data Import and Replace
Database backup/restore
Task management
Task scheduler
Composite tasks

Enterprise Databases support

MongoDB
Cassandra
InfluxDB
Redis
AWS DynamoDB
AWS DocumentDB
Google Bigtable
Couchbase
Apache Hive/Spark/Impala

Customizing DBeaver

Installing extensions - Themes, version control, etc

Troubleshooting

Command Line
Reset UI settings
Reset workspace
Posting issues
Log files
JDBC trace
Thread dump
Proxy configuration

Admin Guide

Managing connections
Managing drivers
Windows Silent Install

License management

License Administration
How to Import License
How to Reassign License

Installation

The installation process depends on the distribution type and your Operational System.

Windows / MacOS Installer

The installer distribution is the recommended way to install DBeaver on Windows and MacOS X. It contains all required dependencies. In addition, the installer automatically upgrades DBeaver to the new version if a previous version has already been installed. To install DBeaver, run the installer executable and follow the instructions on its screens.

NOTE:

- The installer does not change any system settings or the Java installation.
- The included JDK will be accessible only to DBeaver.

ZIP Archive

When installing DBeaver manually, without using an installer:

1. Extract the contents of the archive.
   NOTE: Do not unzip the archive over a previous DBeaver version. If you already have any version of DBeaver extracted in the same location - remove it before unzipping the new version.
   NOTE: All configurations, scripts and other necessary data are stored in a separate location (usually in the user`s home directory) so the program deinstallation does not affect them.
2. Run the **dbeaver** executable.

Debian Package

To install DBeaver using a Debian package:

1. Run `sudo dpkg -i dbeaver-<version>.deb`.
2. Execute `dbeaver &`.

RPM Package

To install DBeaver using RPM package:

1. Run `sudo rpm -ivh dbeaver-<version>.rpm`.
2. Execute `dbeaver &`.

NOTE: To upgrade DBeaver to the next version, use `sudo rpm -Uvh dbeaver-<version>.rpm` parameter.

Application Window Overview

The DBeaver window contains a menu bar, a toolbar, a shortcut bar, a workspace with one or more editors and views, and a status bar:

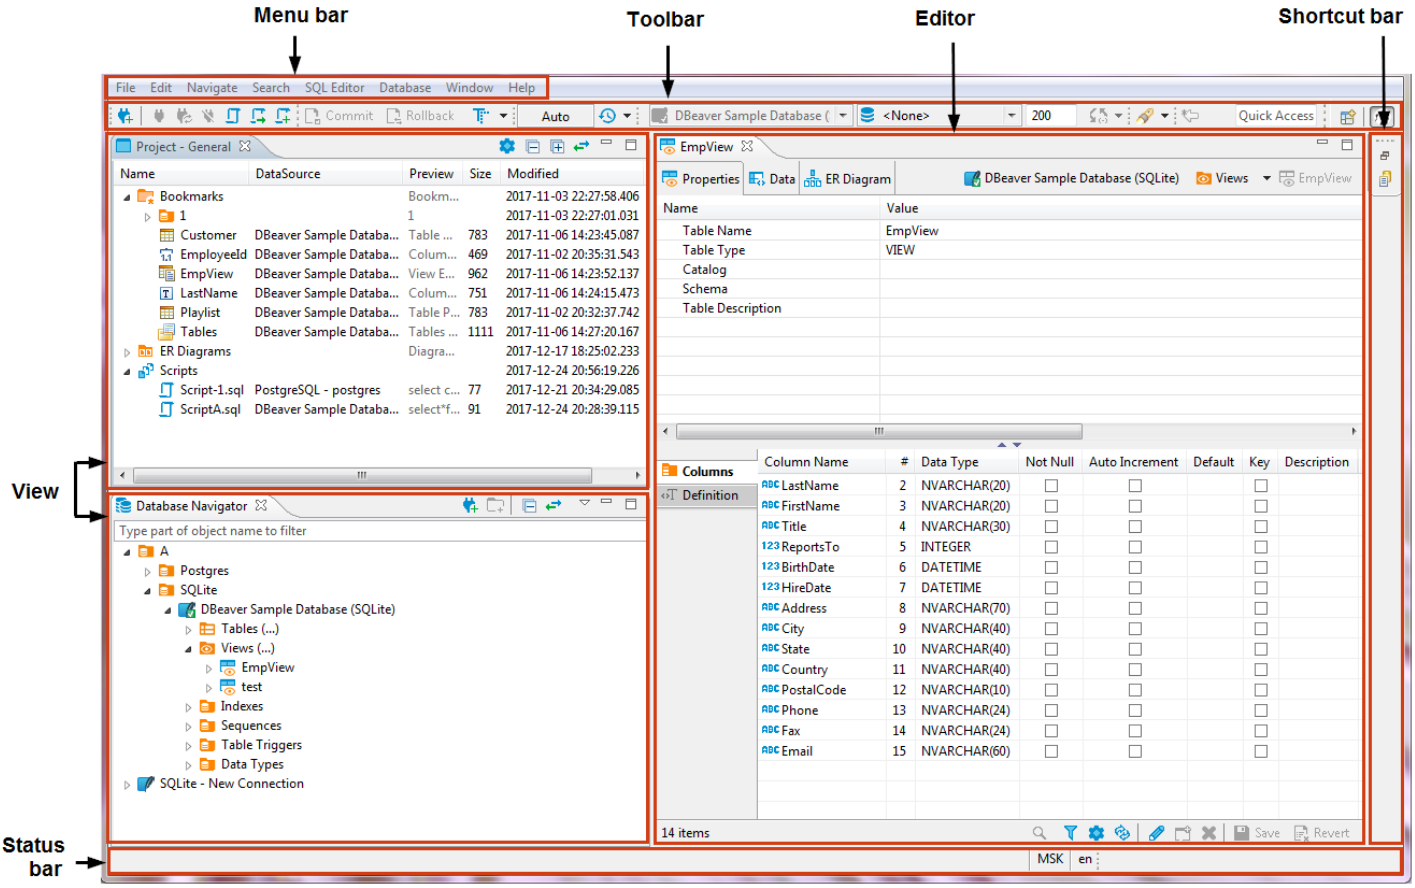

Menu Bar

By default, the menu bar contains the following menus:

- **File** menu contains menu items for the creation of files, folders, projects, database connections, database projects, and ER diagrams as well as Import and Export items.
- **Edit** menu contains global commands like Cut, Copy, Paste, and Delete targeted at the active element.
- **Navigate** menu allows navigation through scripts and database objects.
- **Search** menu provides options to search among files, database objects and across data.
- **SQL Editor** menu is for opening SQL Editor and managing its appearance.
- **Database** menu allows management of database drivers, connections and transactions, as well as reconnecting to and disconnecting from a database.
- **Window** menu includes items to manage the look of DBeaver window: show/hide and minimize/maximize views and editors, display bars, split editors, and manage other preferences.
- **Help** menu contains links to information and help resources, as well as menu items to check the version number and availability of updates.

You can customize the menu bar and the list of menu items to display, for this, go to **Window -> Customize Perspective -> Menu Visibility** tab.

Toolbar

The toolbar contains buttons for the most basic and frequently used commands:

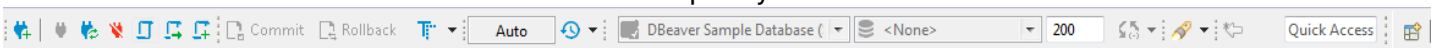

Some of the buttons are enabled (colored), others are disabled (grey). The sets of enabled and disabled buttons change depending on which editor is currently active in the workspace. Only enabled buttons are applicable to the active view or editor.

You can customize the toolbar, for this, go to **Window -> Customize Perspective -> Tool Bar Visibility** tab.

You can hide or show the toolbar in the application window. To do it, go to the Window menu, click **Appearance -> Hide Toolbar / Show Toolbar**.

Shortcut Bar

There are two shortcut bars - one on the left and one on the right side of the workspace zone. Shortcut bars host shortcuts of views and editors and appear if at least one view or editor is minimized, otherwise they are hidden.

Workspace: Views and Editors

Views are windows within the workspace that provide presentations and ways to navigate the information. For more information about particular views, see Views article.

Editors are windows in which you can interact with the content of files and databases. For more information about particular editors, see the Editors article.

Both views and editors can appear as separate windows or as tabs stacked with other views/editors in a tabbed window. The following image shows the title bar of a tabbed window. If tabs do not fit in the title bar of a tabbed window, they become hidden. To see the list of hidden tabs, click the Show List icon that also indicates their number:

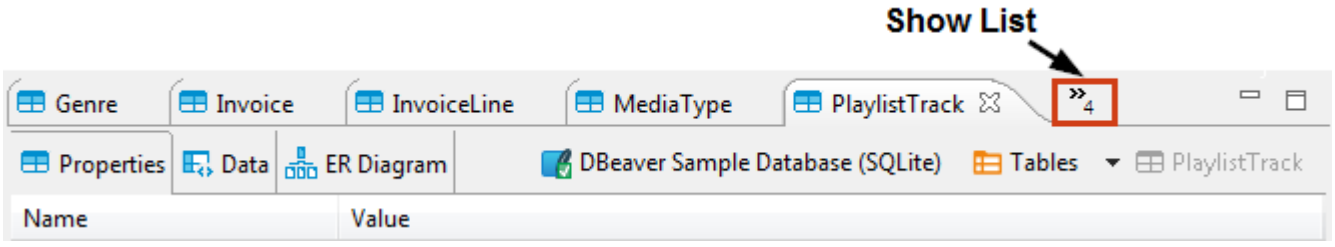

There can be several views and editors opened simultaneously in the workspace but only one of them can be active at a time.

You can change the layout of the workspace by opening and closing views, docking them in different positions in the workspace, collapsing them to the shortcut bar, or expanding them to occupy the whole workspace and restoring them to the latest docked position.

Changing Workspace Layout

You can move views and editors around the workspace and dock them in different positions:

- As a tab in a tabbed window
- As a separate window with a vertical or horizontal layout in any zone of the workspace

You can also swap locations of two views or editors.

To dock a view to a position in the workspace, press and hold the title bar of the view, then drag and drop it onto the desired position.

You can resize the view and editor windows. To resize, place the cursor to the border of the window until it changes to a double-ended arrow, then click and drag the border to the needed size.

To close a view or editor, click the Close button, or right-click the title bar of the view / editor, then click one of the options on the context menu (they change depending on the configuration of windows):

- **Close** - to close the active window or tab in a tabbed window
- **Close Others** (for editors and views that appear as tabs in tabbed windows) - to close all tabs of the current tabbed window except the active tab
- **Close Tabs to the Right / Left** (also for tabbed windows) - to close all tabs of the current tabbed window that are located to the right / left of the active tab
- **Close All** - to close all tabs of a tabbed window (close the window)

Maximizing, Minimizing and Restoring View and Editors

All views and editors have the Close, Minimize and Maximize buttons:

The Maximize button changes to the Restore button when a view or editor is maximized.

To maximize a view or editor to the size of the whole workspace, do one of the following:

- Click the Maximize button in the upper-right corner of the view.
- Double-click the title bar of the view or editor.
- On the Window menu, click **Appearance -> Maximize Active View or Editor**.

When one view is maximized, other views and editors appear as shortcuts on the shortcut bar.

To restore a maximized view or editor to its latest docked position, double-click its title bar or click the Restore button in its upper-right corner.

When you minimize a view, it wraps into a shortcut on the shortcut bar:

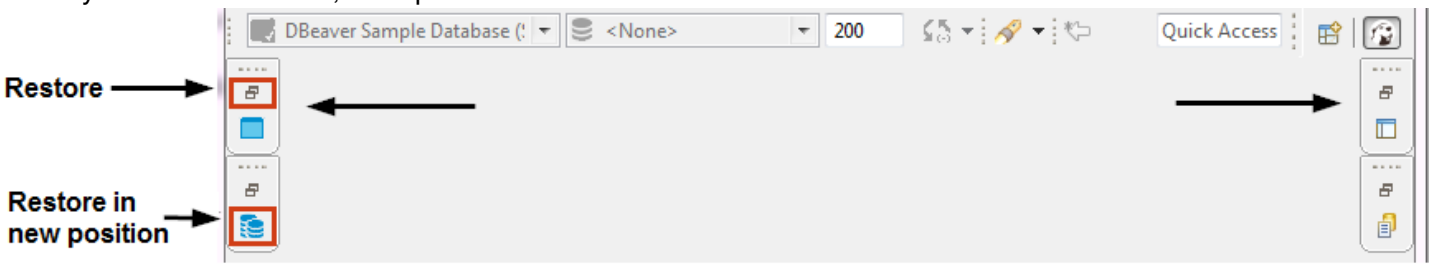

The shortcuts of views and editors may appear on the left or on the right shortcut bar depending on the latest docked position of the view or editor.

To minimize a view, do one of the following:

- Click the Minimize button in the upper-right corner of the view.
- On the Window menu, click **Appearance -> Minimize Active View or Editor**.

To restore a minimized view or editor to its previous position, click the Restore button on its shortcut in the shortcut bar. To restore a minimized view or editor to a new position, click the view / editor name button under the restore button.

Views

Views are windows within the workspace that provide presentations and ways to navigate the information. The main views in DBeaver are: Database Navigator, Projects and Project Explorer.

To open a view:

- On the Window menu, click **Show View** and then, on the submenu, click the name of the view. Click **Other** if the view is not visible on the submenu.
- For Database Navigator, Projects, and Project Explorer views, on the Window menu, just click the name of the view.

Some views open on demand, for example the Search view opens to show search results.

Views provide their own toolbar and menu:

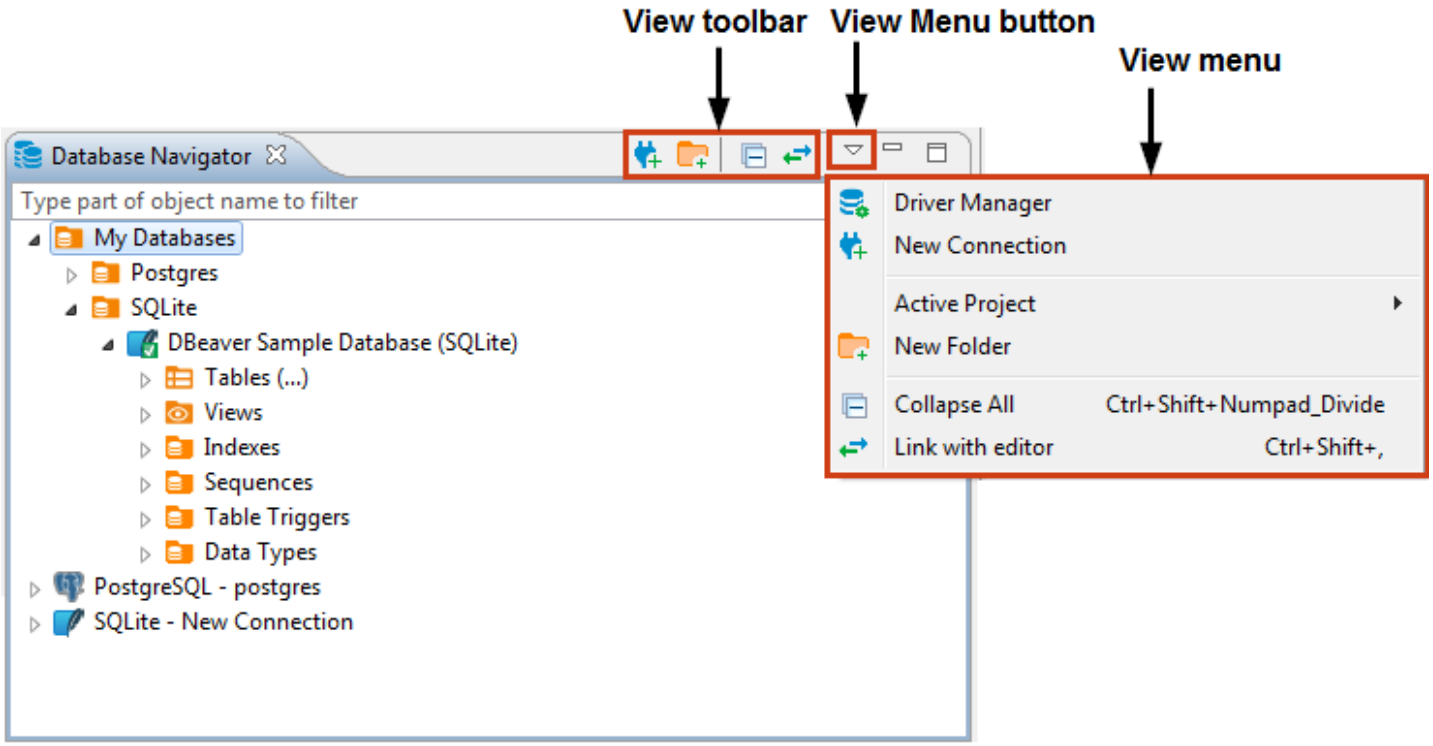

To open the view menu, click the View Menu button in the upper-right corner of the view`s title bar, next to the Minimize button.

The toolbar contains buttons applicable to the objects displayed in the view. The set of enabled and disabled buttons depends on the object in focus.

Views also provide context menus for objects they display. To open a context menu for an object, right-click the object.

Database Navigator

Database Navigator is the main view to work with the structure and content of databases. To open Database Navigator, on the Windows menu, click **Database Navigator**. For information on how to change the view layout, please see the Application Window Overview article.

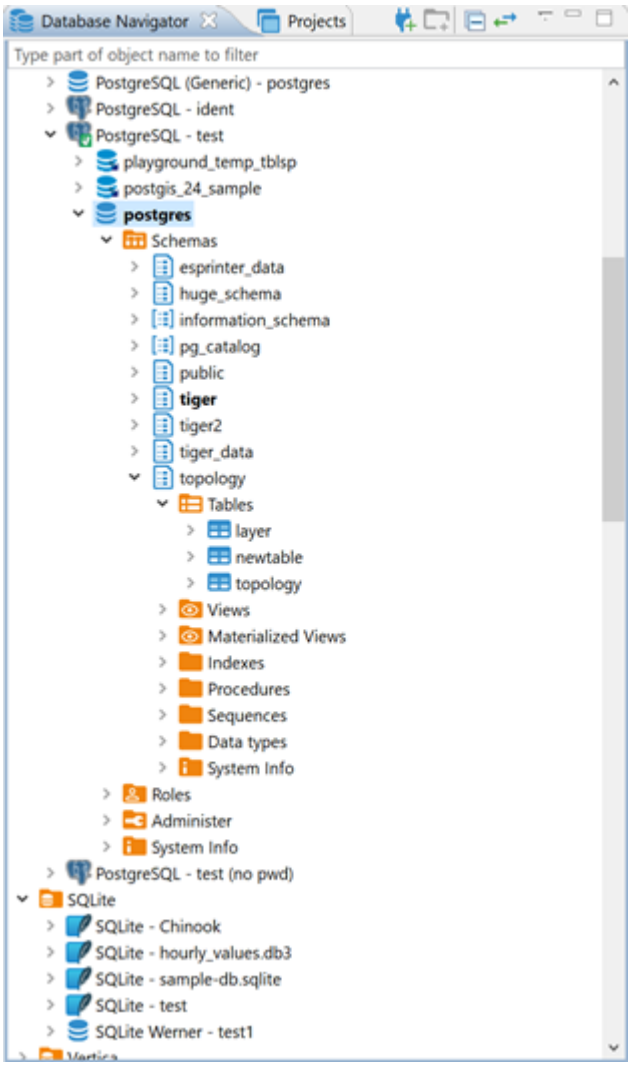

Database Navigator contains a tree of objects, a toolbar and View menu which contain generic items. Each object in the tree has its own context menu. The tree contains the following objects:

- Folders -
- Database connections - and other (icons differ depending on the database type)
- Database objects - various depending on the database type, such as Tables , Views , Columns , Indexes , etc.

To open the view menu of Database Navigator, click the View Menu button ( ) in the upper-right corner of the window. For more information on where to find the view toolbar and menu, please see the Views article.

The menu contains the following items:

Icon	Menu item	Description
	Driver Manager	Opens the Driver Manager window that allows you to create, edit and delete drivers for databases. See Database Drivers for information about managing database drivers.

	New Connection	Opens the Create new connection wizard. See Create Connection for information about creating connections.
(empty)	**Active Project**	Displays a submenu which allows you to choose a project. See Projects and Projects view for information about projects.
	New Folder	Opens a dialog box for creating a new folder
	Collapse All	Collapses the tree to the root level
	Link with editor	Synchronizes the active editor with the element in the database navigator

The toolbar is located in the title bar of the window. Its buttons duplicate the menu items, except for the **Active Project**.

To open the context menu for an object, right-click the object in the tree. The following table summarizes the context menu items for all types of objects that may appear in the tree. Note that the presence or absence of the context menu items for an object depends on the database and object types.

Menu item	Description
Open folder	Opens a folder in a separate view
Create new connections / Create New Connection	Opens the Create new connection wizard
New Folder	Opens a dialog box for creating a new folder
Copy	Copies an object to the clipboard
Paste	Inserts the copied object into a selected folder - most convenient for copy-pasting connections
Delete	Deletes an object **WARNING!** The Delete menu item removes the object not only from the tree but from the database itself or the file system, and this action is irreversible.
Rename	Opens the Rename [object] dialog box
Properties	Opens the Properties for [object] window which allows viewing and modifying the object`s properties
Refresh	Depending on the object, refreshes the object itself, or its parent, or its subnodes – mostly used for refreshing tables and schemes
Connect	Attempts to connect to the database
Invalidate/Reconnect	Checks the status of connection, if it is broken, attempts to reconnect
Disconnect	Disconnects from the database
SQL Editor	Opens a new SQL editor for the connection
Recent SQL Editor	Opens the most recently opened SQL editor
Edit Connection	Opens the Connection Configuration window that allows configuring connection settings
View [objects]	Opens the object in a separate viewer
Edit [object]	Opens the object in a separate editor
Create new [object]	Opens an editor in which you can specify properties and save the new object
	Opens a submenu of one or more filtering options (depending on the object):

Filter	- Hide [object] - Show only [object] - Configure [objects] filter - Toggle filter - Clear filter See Filters for information.
Copy Advanced Info	Copies the full name of an object
Read Data in SQL Console	Opens an SQL console displaying the object`s data
Compare	- Appears only if you select several objects of the same level - Opens the Compare objects wizard which guides you through the steps to generate a comparison report for the selected objects
Generate SQL	Opens a submenu on which you can select the type of SQL query to generate: - SELECT - INSERT - UPDATE - DELETE - MERGE - DDL Clicking one of the items (for example INSERT) generates a relevant query in a separate window.
Export Table Data	Opens the Data Transfer wizard that helps you select a format and export table data
Import Table Data	Opens a window with existing database connections in which you can select a table to import data from
Tools	Opens a submenu that provides tools for database backup and restore, vacuum, etc.

For information on how to filter database objects in Database Navigator, please see the Filter Database Objects article.

Filter Database Objects

In the Database Navigator and Database Object Editor you can filter database objects to include or exclude some of them from the view. You can filter schemas, tables, views and procedures. A dots sign `(...)` next to the node`s name indicates that a filter is applied to its sub-nodes: Tables (...)

There are several ways in which you can filter objects. One of the ways is to filter objects by the names of tables and views using the filter field above the tree of objects:

To filter objects by the name of a table and view, type the name in the field. The tree dynamically updates to show tables / views with that name. To reset the filter, click the Clear icon (✕) on the right end of the field.

Another way to filter objects is to use the **Filter** item on the context menu of a single object. To filter objects using the **Filter** menu, right-click the object, then click **Filter** on the context menu, and then click one of the items on the submenu:

Filter submenu item	Description
Hide '[object name]'	Hides the current object while displaying the other ones
Show only '[object name]'	Shows the current object while hiding the other ones
Toggle filter	Inverts the filtering – shows hidden objects and vice versa
Clear filter	Removes the filtering to display all objects
Configure [objects] filter	Appears only to the folder or parent nodes of database objects - like 'Tables', 'Indexes', etc. Allows the creation of a complex filter with multiple filtering criteria, see Configure Filters.

A third way of filtering is to use the **Filter** item on the context menu on several objects:

1. Select several objects of the same type using Ctrl or Shift keys.
2. Right-click the selection, then click **Filter**, and then choose one of the options on the submenu:

Filter submenu item	Description
Hide N objects	Hides the selected objects while displaying the rest
Show only selected objects	Shows the selected objects while hiding the rest

To reset such filters, right-click the parent (folder) node displaying the dots sign `(...)`, and then click **Filter -> Clear filter**.

Configure Filters

You can configure custom filters to filter database objects in the Database Navigator and Database Object Editor.

To configure a custom filter:

1. In the Database Navigator, right-click the object and on the context menu click **Filter -> Configure [objects] filter**. In the Database Object editor, in the toolbar of the **Properties** tab, click the Filter settings button ( ). The Filtering window opens.

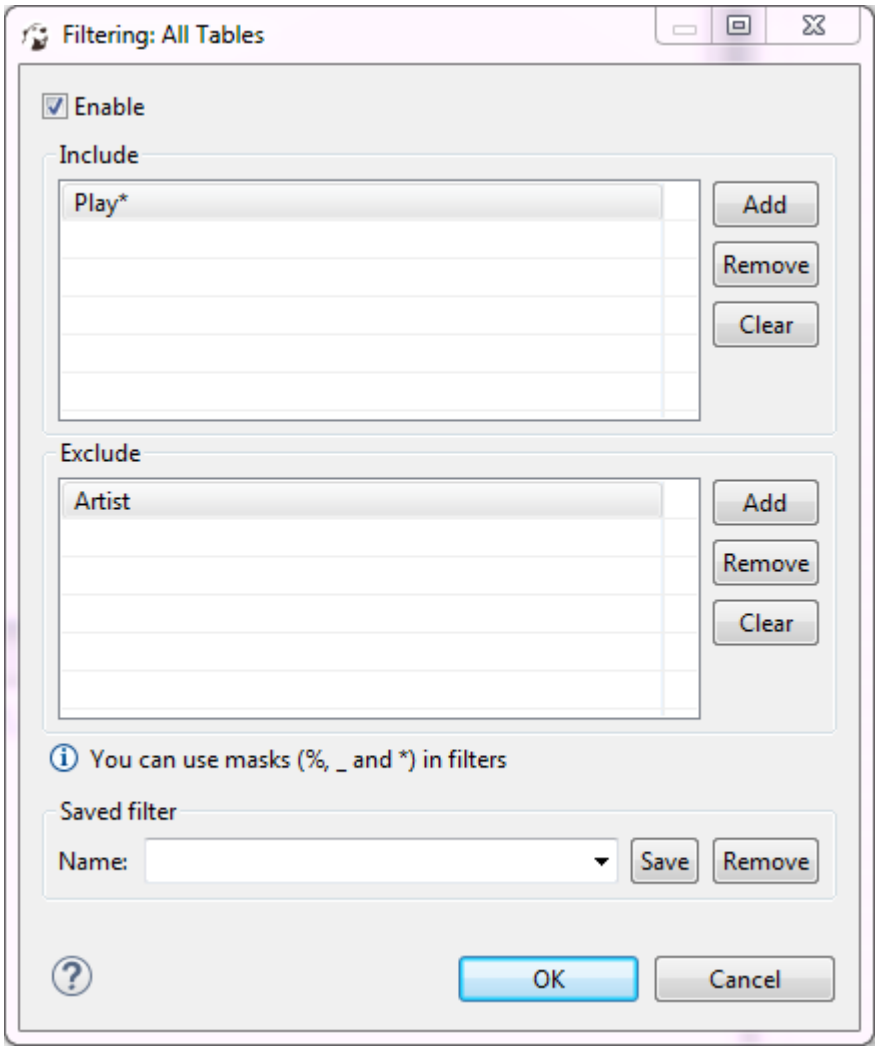

2. Select the **Enable** checkbox to activate the fields of the window.

3. If you want the filter to apply to all objects of a certain type, for example to all schemes, click **Show global filter** . Otherwise, the filter will apply only to the current object. NOTE: Once you apply the global filter, you cannot revert back to the local filer in the same window. To create a local filter, reopen the Filtering window, see Step 1.

4. For objects that you want to show, click **Add** next to the **Include** field and then, in the field itself, enter the name or combination of symbols to search. For objects that you want to hide, click **Add** next to the **Exclude** field and then, in the field itself, enter the name or combination of symbols to search. NOTE: You can use masks with `%` and `*` to replace one or more symbols and `_` to replace one symbol in the search combination.

5. To remove one filtering combination, click the combination in the field and then click **Remove**. To remove all combinations from either of the fields, click **Clear** next to the field.

6. Once you set all filtering criteria, you can save a filter to use for other objects. To save the filter, in the Saved filter area, in the **Name** field, enter the filter`s name and click **Save**.
7. You can also remove any of the saved filters. To remove a filter, in the **Name** drop-down list, click the filter name and then click **Remove**.
8. Click **OK** to apply the filtering criteria. Otherwise, click **Cancel**.

Projects View

You might need to classify and group database connections into projects. Projects store objects related not to a particular database but to all database connections. These are usually files stored on the file system.

The Projects view displays all projects created in the system and provides tools to manage them. To open the Projects view, on the **Window** menu, click **Projects** (or use `ALT+W+P` shortcut).

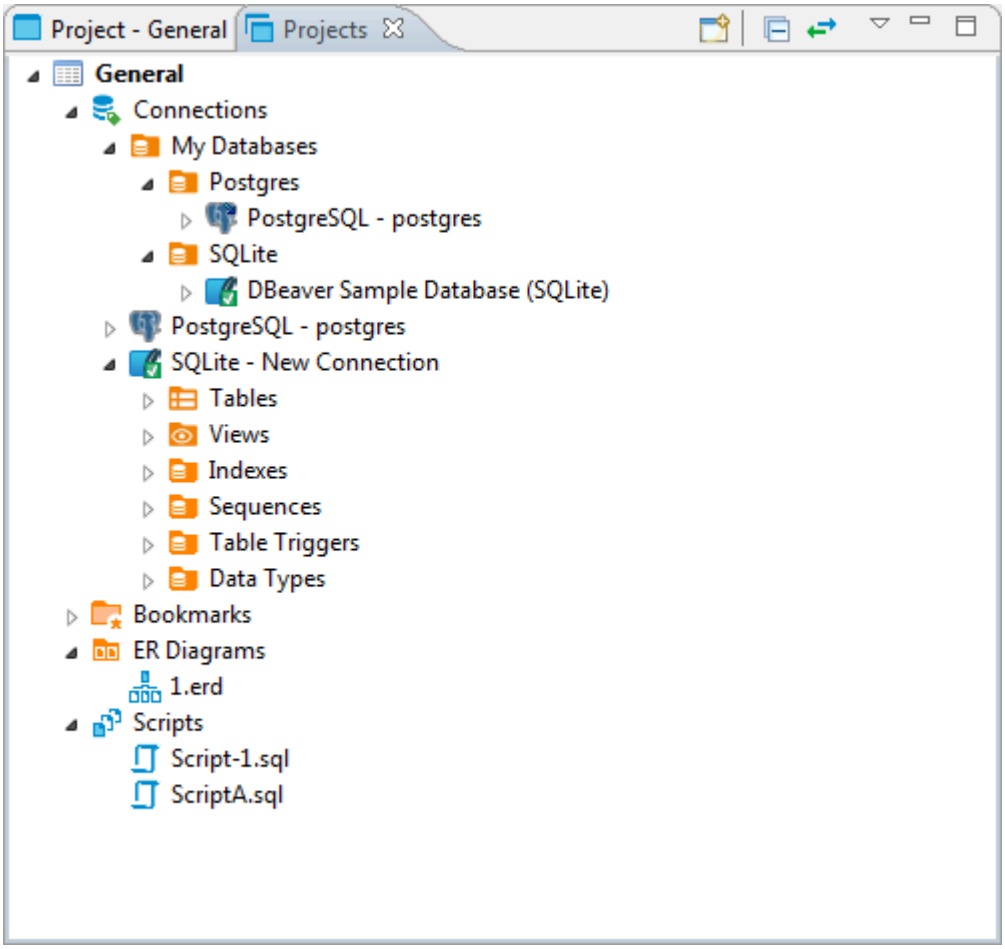

For information on how to change the view layout, please see the Application Window Overview article.

The projects are organized into a tree and all have the same high-level structure:

- **Connections** – repeat the content of the Database Navigator view for this project. You can perform the same actions over the objects of the databases as in the Database Navigator.
- **Bookmarks** – contains bookmarks – shortcuts to database objects, see …
- **ER Diagrams** - contains ER diagrams that you can drag-and-drop here from other folders
- **Scripts** – contains scripts that you can drag-and-drop here from other folders

The Projects view provides a toolbar and View menu which contain generic items. Each object in the tree has its own context menu.

To open the view menu of the Projects view, click the View Menu button ( ) in the upper-right corner of the window. The view menu contains the following items:

Icon	Item	Description

	Create Project	Opens the Create Project wizard
	Refresh Projects	Refreshes the projects tree to display changes caused by creating modifying or deleting projects
	Collapse All	Collapses the tree to the root level
	Link with editor	- Enabled when at least one editor is open, otherwise disabled - Highlights the object in the tree that has its editor open

The toolbar is located in the title bar of the window, its buttons duplicate the view menu items except for the **Refresh Projects** one.

To open the context menu for an object in the tree, right-click the object. For information about context menu items of all objects under the **Connections** node of the tree, please see Database Navigator. The context menus of other nodes in the tree contain some basic items for copy-pasting, renaming, deleting objects, managing their properties, creating folders, etc.

- The **Set Active Project** menu item (for a project root node) makes the project active, that is visible in the Database Navigator.
- The **Link File (SQL Script)** and **Link Folder** menu items allow creating links to files and folders in the file system.

For information about managing projects, please see Projects article.

Project Explorer

The Project Explorer view displays detailed contents of the currently active project. To open the Project Explorer, click **Window -> Project Explorer**.

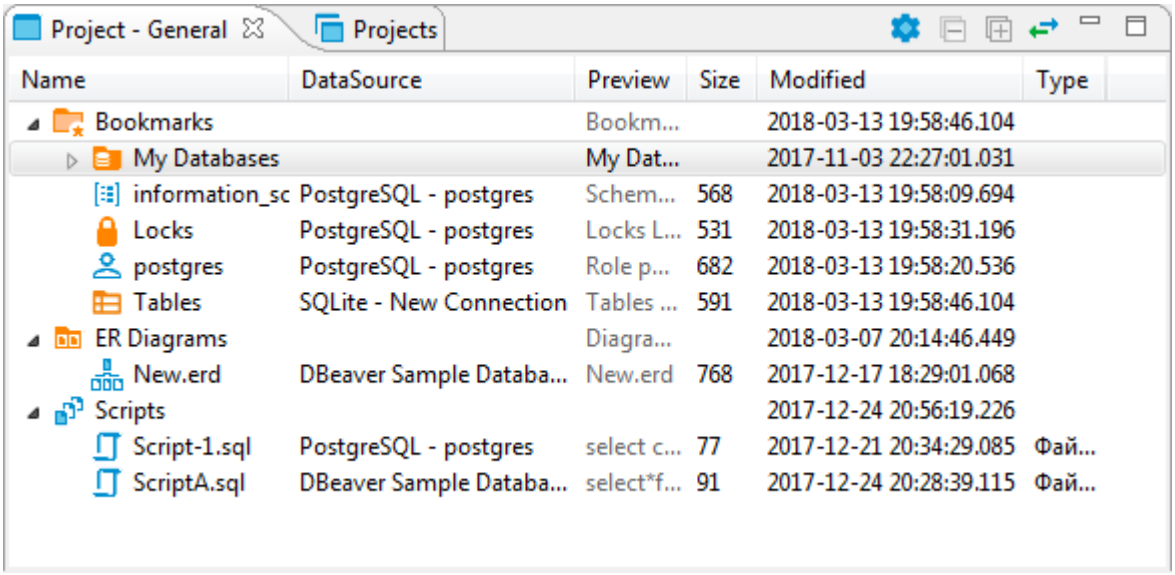

For information on how to change the view layout, please see the Application Window Overview article.

The title of the Project Explorer includes the name of the project: Project – [Project name]. **General** is a project that initially exists in the system by default.

The Project Explorer displays the content of a project with metadata. The content includes: **Bookmarks**, **ER Diagrams**, and **Scripts**. The metadata appears in columns which you can hide or show.

The Project Explorer view provides a toolbar that contains the following buttons:

Button	Name	Description
	Configure columns visibility	Opens a dialog box in which you can select columns to display in the view
	Collapse All	Collapses the tree to the root level
	Expand All	Expands the tree nodes
	Link with editor	- Enabled when at least one editor is open, otherwise disabled - Highlights the object in the tree that has its editor open

To sort the metadata in the table by a certain column, click the column header.

Query Manager

Query Manager is a view that shows the history of all SQL queries that DBeaver has executed during the current session.

NOTE: The DBeaver EE version persists all executed queries in the internal database so its execution history is available after the program restarts.

To open the Query Manager, do one of the following:

- Click the arrow next to the **Transaction Log** button in the toolbar and then click **Query Manager** on the dropdown menu:

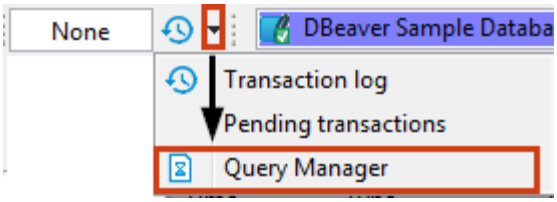

- On the **Window** menu, click **Show View -> Query Manager**:

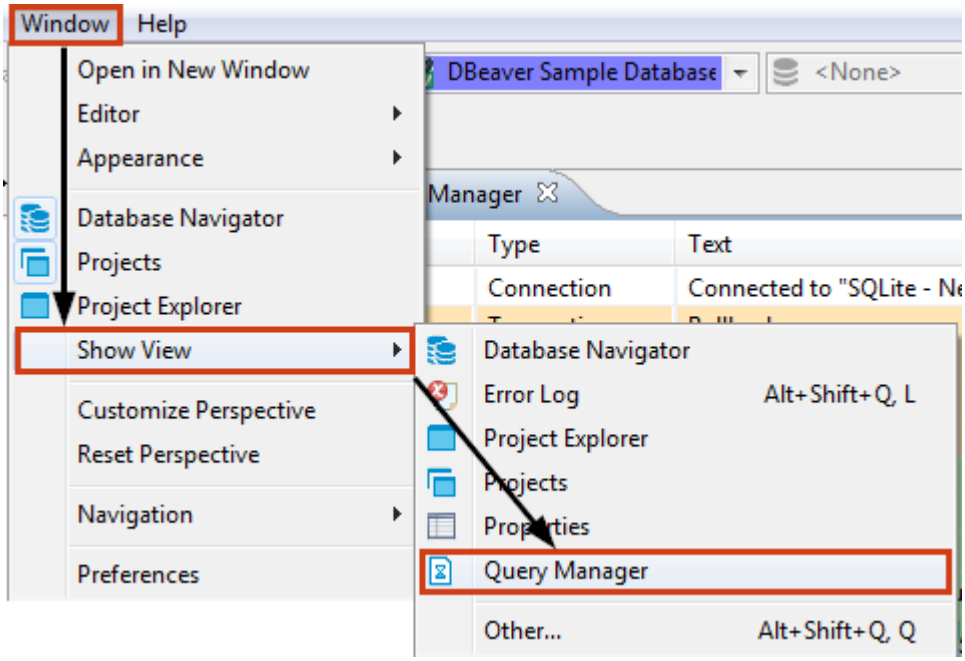

The Query Manager logs all queries together with their execution statistics (execution time, duration, number of fetched/updated rows, errors, etc.):

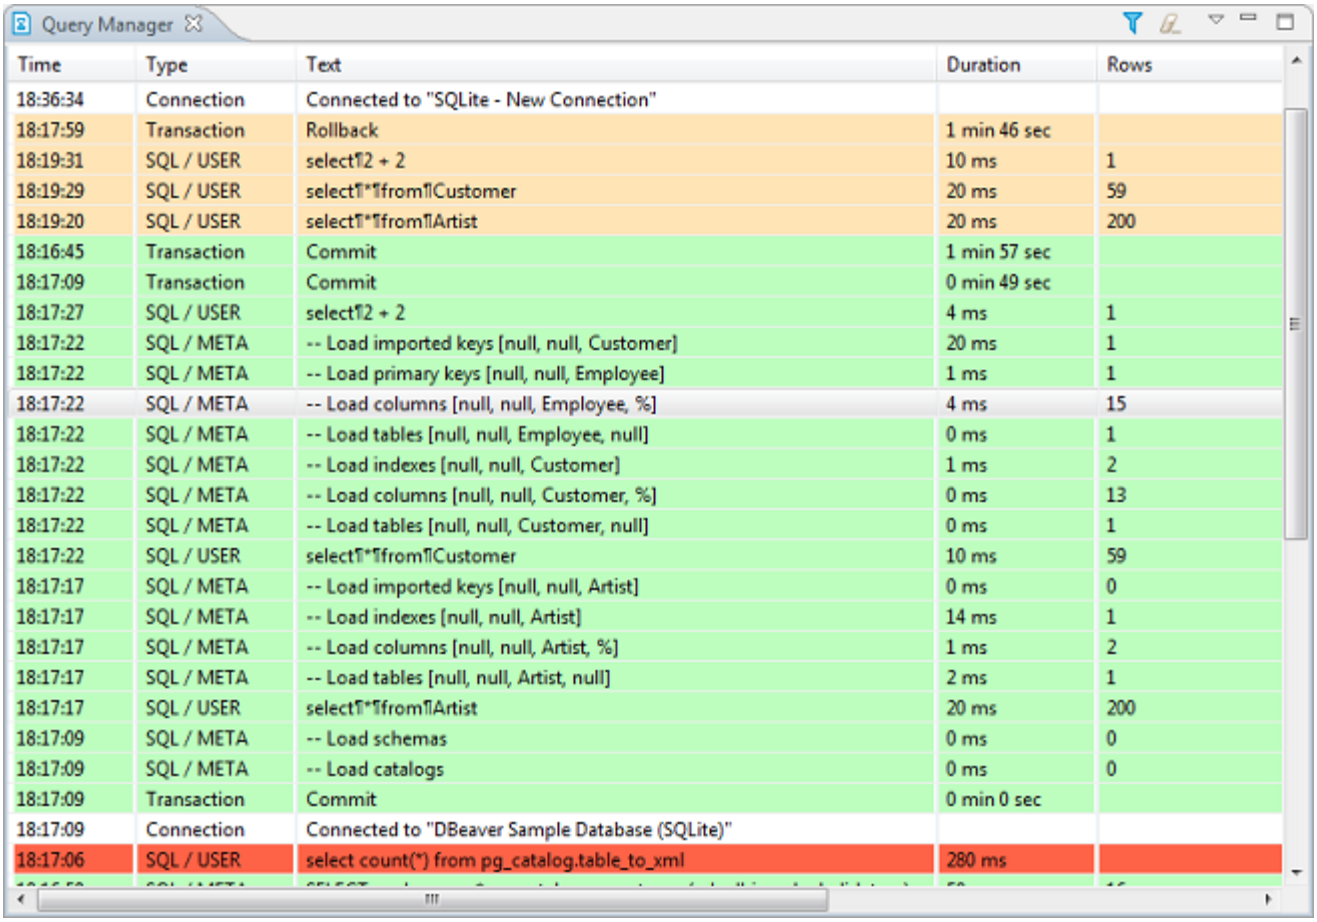

You can modify the look of the Query Manager by filtering queries and setting the number of entries displayed per page, as well as specifying some storage settings, see the 'Query Manager Properties' section below.

To erase all entries from the Query Manager, click the **Clear query manager log** button ( ) in the view`s toolbar.

Query Manager Properties

To manage the look of the Query Manager, filter the entries, and modify the storage settings, click the **Set query manager filter** button ( ) in the view`s toolbar. The Properties for Query Manager window opens:

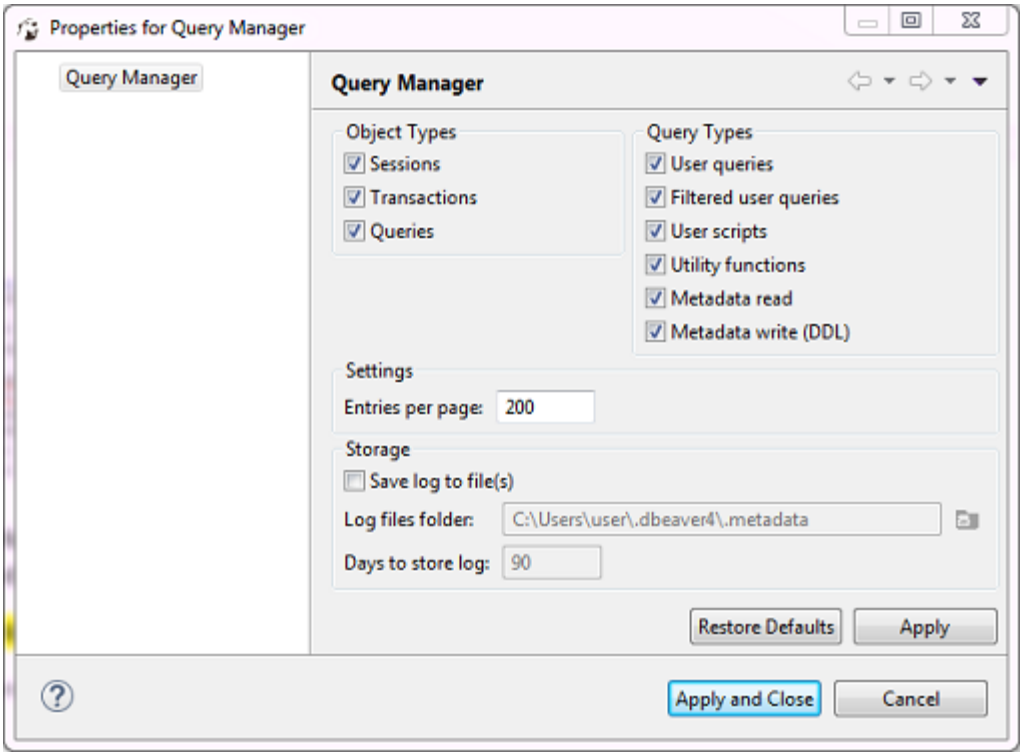

- To filter entries by object type, select or clear the checkboxes in the **Object Types** section. To filter entries by query type, select clear the checkboxes in the **Query Types** section.
- To change the number of entries displayed per page, enter the new number in the **Entries per page** field.
- To set DBeaver to save the query log in a file, select the **Save log to file(s)** checkbox and then specify the file location in the **Log files folder** field.

After you make all necessary changes to the settings, click **Apply** to apply the changes and keep the window open or click **Apply and Close** to apply the changes and close the window. To discard all changes and return to the previous state, click **Restore Defaults**.

Background Tasks

You can open the Background view from the main menu - click **Window -> Show View -> Other**, then in the Show View window, expand the **General** folder, click **Background Tasks** and then click **Open**:

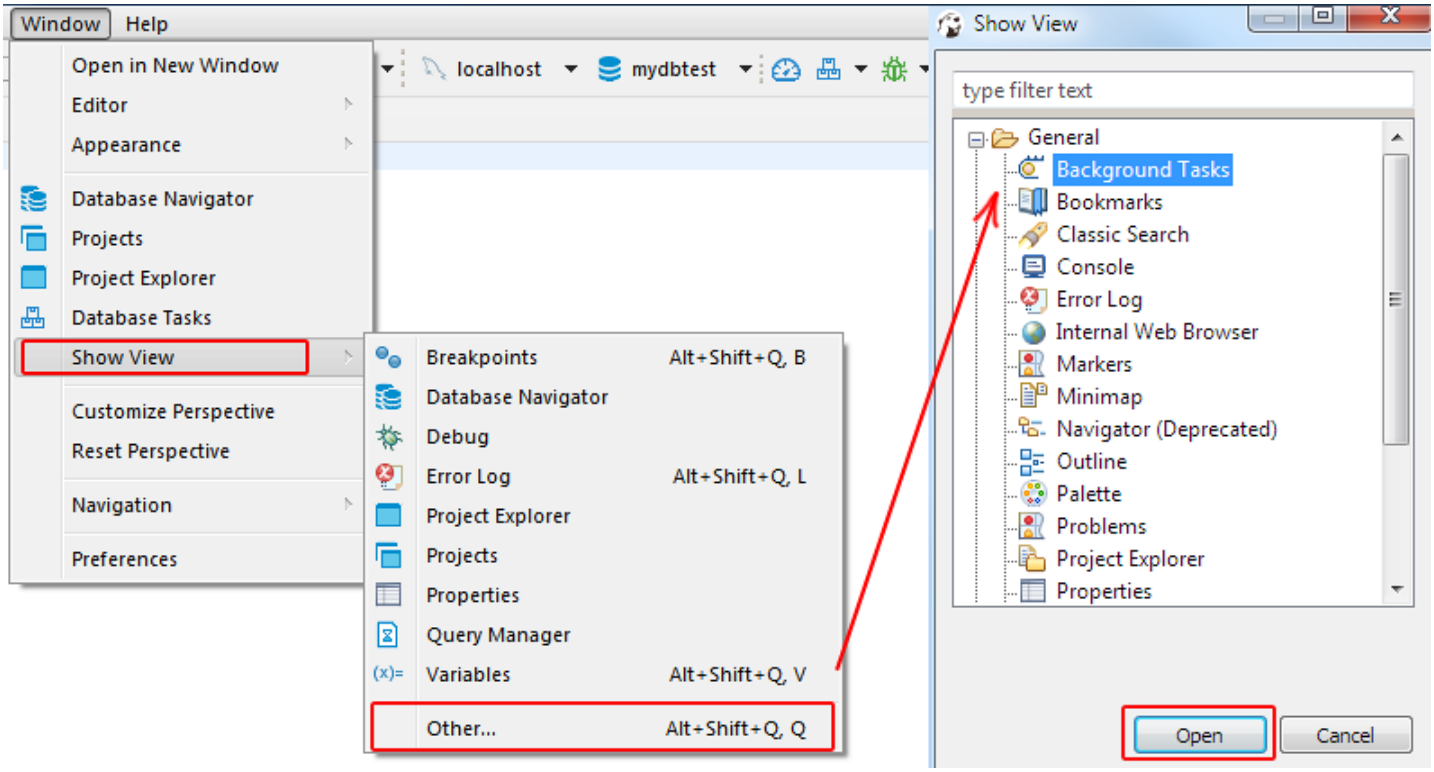

You can also open the Background Tasks view from some other views or editors by using a special button, for example from the Search view.

The Background Tasks view shows the progress of such background tasks as search, SQL query execution, etc. The view shows data when background tasks take some noticeable time and is useful when you want to track the progress of lengthy operations. If you open this view at a short task, the view will be empty.

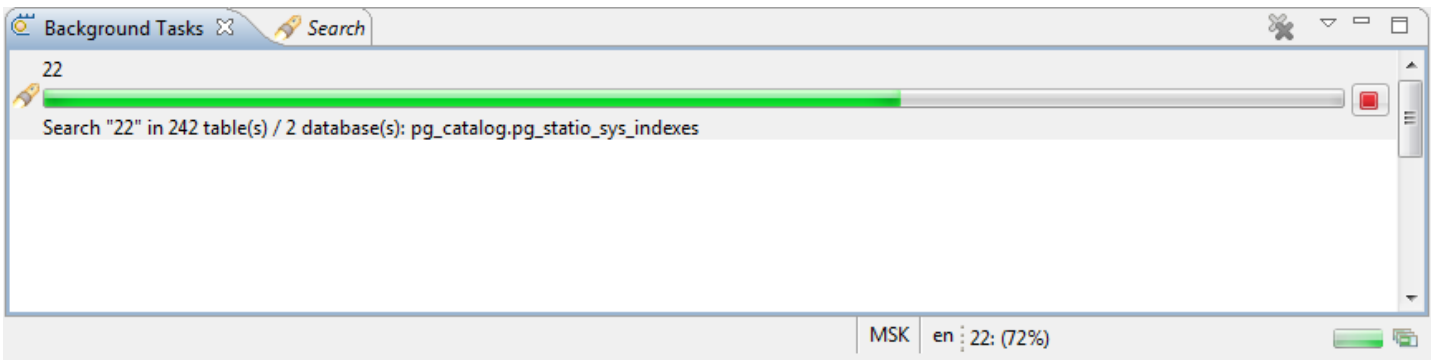

You can cancel the task in progress - click the **Cancel Operation** button ( ).

Database Object Editor

The Database object, or metadata, editor is available for multiple database objects such as tables, views and schemas. To open the metadata editor for an object, in the Database Navigator or in the Projects view:

- Double-click the database object
- Click the database object and press `Enter` or `F4`

The editor has three tabs:

- **Properties** tab appears for all objects, contains properties of the database object and its sub-entities, see further in this article
- **Data** tab appears for tables and views and represents the Data Editor
- **ER Diagram** tab appears for tables and schemas and displays ERD (Entity Relation Diagrams), see ER Diagrams and Database Structure Diagrams

The tabs have the following common parts:

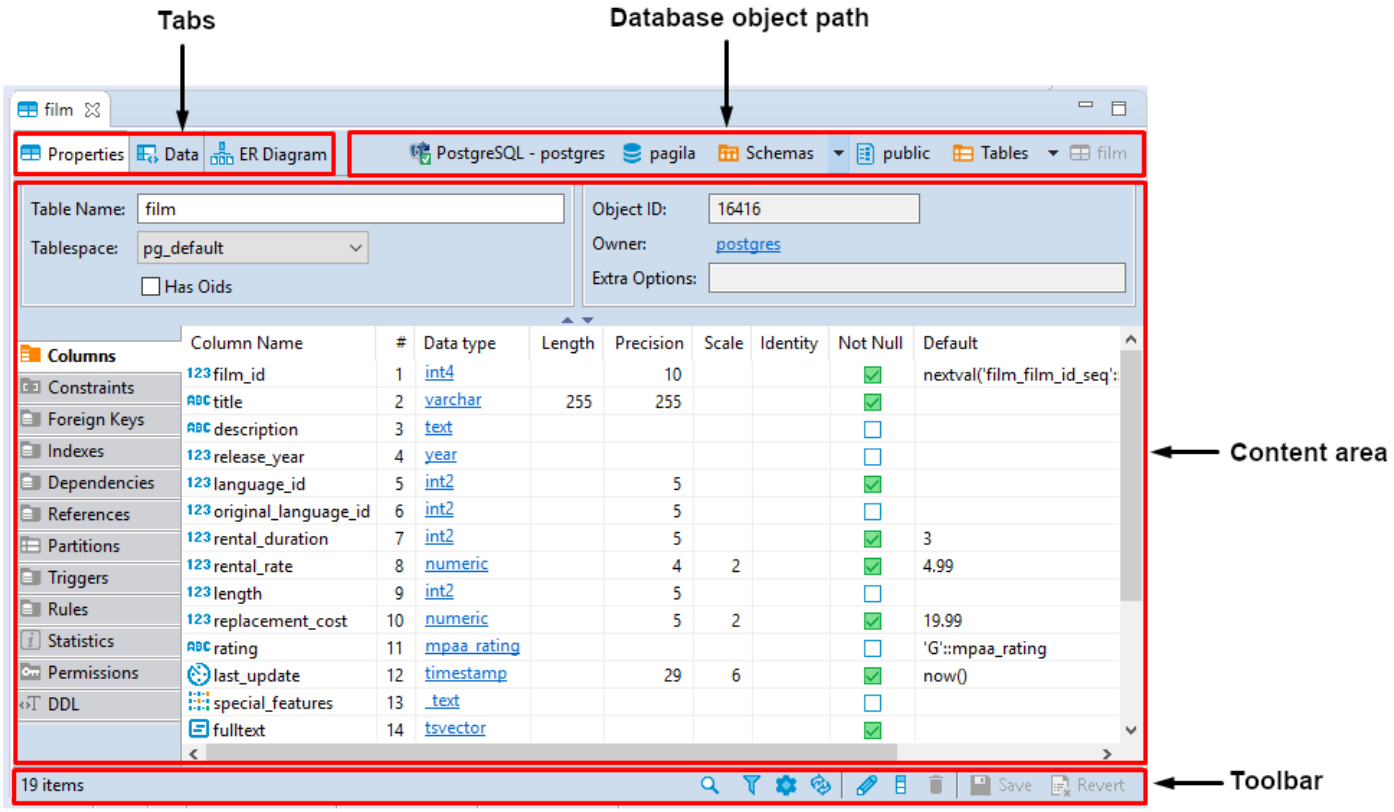

The object`s path shows the chain of all its parent entities. The entities are clickable: clicking an entity in the path, depending on its nature, either shows its children or opens an editor or a settings window.

The toolbar contains different tools on each of the three tabs.

An asterisk appears in the title of an editor if it contains unsaved changes:

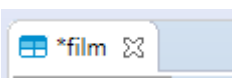

The Database Object editor supports the `Ctrl+Z` (undo) function.

Properties Editor

The Properties tab of the Database Object Editor provides you with the tools to view and edit the database object`s properties.

The content area of the Properties tab falls into two parts: the top pane displays properties of the current database object itself while the bottom pane contains properties of the object`s sub-entities or some complex properties like DDL (an SQL description of the current database object).

Properties of the sub-entities appear in the side tabbed editors – to open such an editor, click the tabs on the left side of the area:

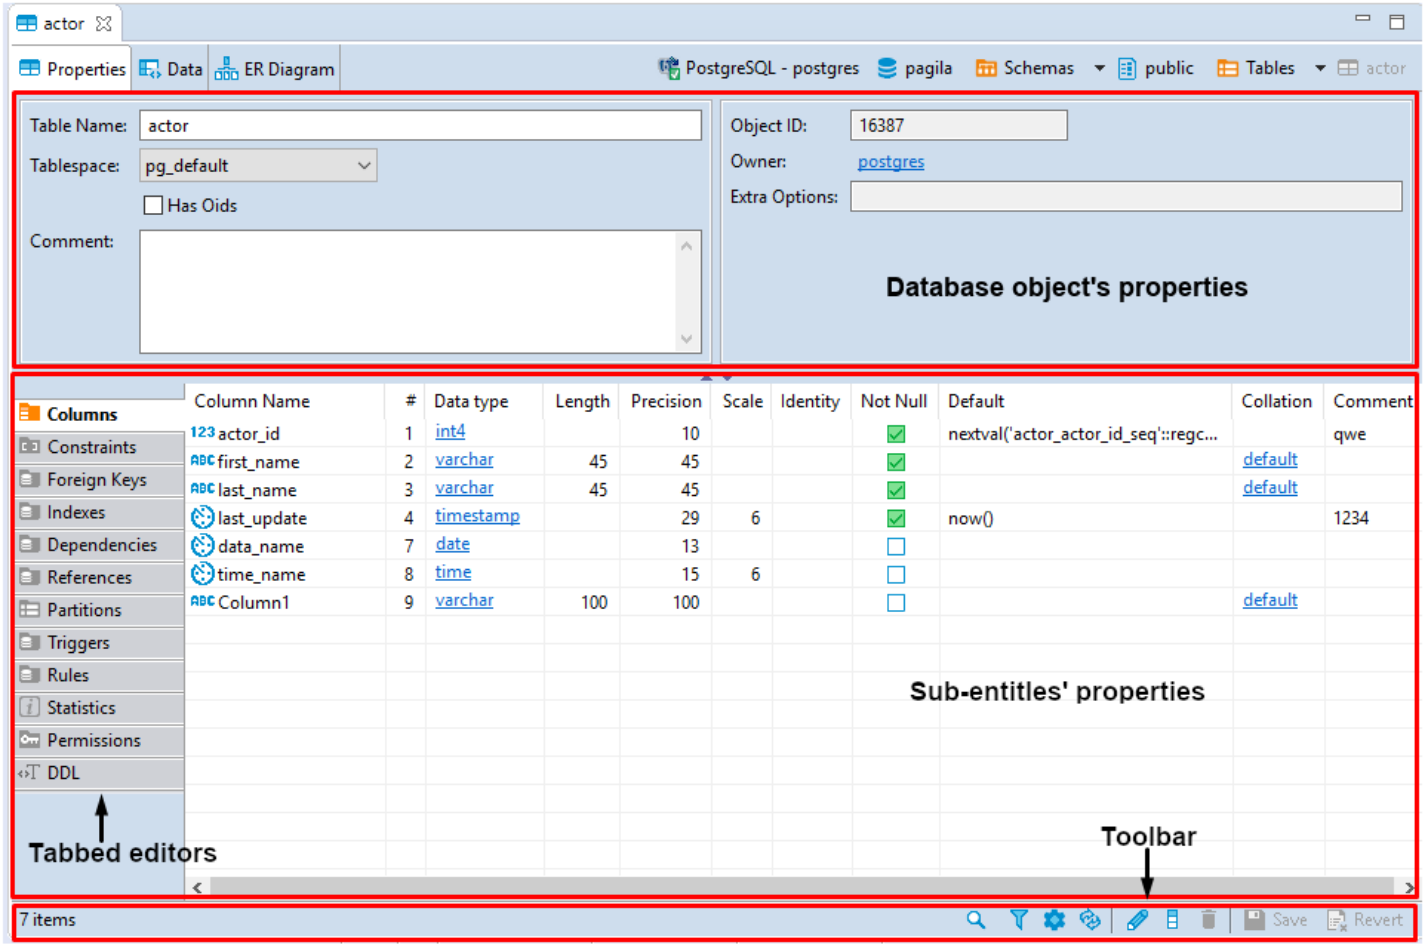

The toolbar at the bottom of the editor provides the following tools for the majority of sub-entities except for some specific ones like Permissions (in PostreSQL) or SQL based views (DDL and Source):

Button	Name	Description
	Search items	Displays a search field next to the button: - Type in the search combination - the content updates dynamically - To remove the filter, click the cross icon next to the search field
	Filter settings	Opens the Filtering window which allows setting a custom filter, see Configure Filters
	Configure columns	Opens the Configure columns dialog box in which you can select the columns to display or hide in the current view
	Refresh the selected items	Depending on the database type, refreshes either the current item or its parent or the whole database object – reloading data from the database
	View	Opens an editor/viewer for the item currently in focus

	Create new [items]	Creates a new item of the same type as currently displayed in the open view, for example, a column
	Delete database object	Deletes the item currently in focus
	Save the current contents	- Same as the **Save** button on the application main toolbar - Same as `Ctrl+S` - Opens the **Persist Changes** window that allows saving changes in the currently open sub-entity NOTE: DBeaver recommends saving work after each change.
	Revert to the last saved state	Reverts all changes made to the whole database object to the last saved state

Items in the tabbed editors have context menus which provide the same commands as those in the Database Navigator. To open a context menu for an item, right-click the item.

SQL Script Editors

SQL script editors (**DDL** and **Source**) of the Properties tab contain SQL script that you can either view or modify. The toolbar of the DDL and Source tabs provides the following tools:

Button	Name	Description
	Load form file	- Allows selecting a file from the file system - Disabled if the SQL code is read-only
	Save to file	Allows saving the current SQL code to a file
	Open in SQL console	Opens the SQL code in an SQL Editor

You can select parts of the SQL code and apply generic commands such as copy-paste or SQL-specific commands like formatting – using the context menu. To open the context menu, right-click the SQL code. See SQL Editor for information about SQL-specific commands.

NOTE: **SQL Assist**, **SQL Template**, and **SQL Context Information** menu items on the context menu are disabled if the SQL script is read-only.

Data Editor

The Data editor appears:

- As the **Data** tab of the Database Object Editor, which is only available for tables and views.
- As the **Results** tab when you run a custom SQL query in SQL Editor

The Data editor allows the viewing and data editing of a database table or view. The central part of the Data editor is the data table. The editor also provides two toolbars and a filter field:

To learn how many rows the data table contains, click the **Calculate total row count** button in the bottom toolbar. The number of rows appears in a status field next to the button: 8,715

To learn about ways to navigate data in the data table, see Navigation article.

The top toolbar contains the following buttons:

Button	Name	Description
	Apply filter criteria	Applies filter criteria entered in the filter field above the data table, see Data Filters article for more information
	Remove all filters /orderings	Removes all filters and orderings applied to the data
	Save filter settings for current object	Saves the current filter settings for the database object to apply next time when you reopen it in the editor, see details in the Data Filters article
	Custom Filters	Opens the Result Set Order/Filter Settings window, see article for more information
	Configure auto-refresh	Allows configuring data auto-refresh settings, see Data Refresh article for details
	Forward and backward - history navigation buttons	Navigate forward and backward in the Data Editor history, see *History* section of Navigation article for more information. The buttons are equivalent to pressing the key combinations: `Alt+Left` (backward) and `Alt+right` (forward).

The side bar contains the following tabs:

Button	Name	Description Chart_button
	Grid	Switches to grid view of data
	Text	Switches to plain text view of data
	Chart	Switches to chart view. For more details on charts, see the Managing Charts article.
	Record	- Same as pressing `Tab` - Switches the positions of rows and columns so that the columns appear as rows, and the rows hide in one **Value** column, see details in the *Table vs. Record Views* section of the article.

The bottom toolbar provides the following buttons:

Button	Name	Description
	Save	Saves all unsaved changes to the data such as adding, duplicating, deleting rows, inline editing of values, see the Data Viewing and Editing article for information
	Cancel	Discards all unsaved changes to the data
	Script	Opens the Preview Changes window in which you can see changes that you have made to the data, see details in the Data Viewing and Editing article
	Edit cell value in separate dialog /editor	Opens the cell in focus for editing in a separate editor or dialog box, see details in the *Cell Editor* section of the Data Viewing and Editing article
	Add new row	Adds a new empty row below the current row, see details in the *Adding, Copying and Deleting Rows* section of the Data Viewing and Editing article
	Duplicate current row	Copies the current rows and pastes the copy below the current row, see details in *Adding, Copying and Deleting Rows* section of Data Viewing and Editing article
	Delete current row	Colors the rows in focus in red to mark them for deletion, see details in the *Adding, Copying and Deleting Rows* section of the Data Viewing and Editing article
	Move to first row	Moves the focus (highlighting) from the current to the first row of the table
	Move to previous row	Moves the focus (highlighting) from the current to the previous row of the table
	Move to next row	Moves the focus (highlighting) from the current to the next row of the table
	Move to last row	Moves the focus (highlighting) from the current to the last row of the table
	Fetch next page of results	Fetches the next portion of data (next N rows) making it ready for display, see *Scrolling Results Page* section of Navigation article for more information
	Fetch all rows	Fetches the whole result set making it ready for display, see the *Scrolling Results Page* section of the Navigation article for more information
	Panels	Opens panels on the right side of the Data Editor, see the Panels for information
	Configure	Opens a dropdown menu with settings
	JSON	- Available in EE version only for MongoDB documents and JSON tables - Switches to JSON view of data
	XML	- Available in EE version only for XML tables - Switches to XML view of data
	Generate Mock Data	Available in EE version only. Opens the Mock Data Generator window

	Rows count details	Opens the Status details dialog box showing the timing details of fetching table rows
	Calculate total row count	Calculates the total number of rows in the table

Every cell in the data table has a context menu – right-click the cell to open the menu. The context menu provides the following items:

Menu Item	Description
Cut	Cuts the content of the current cell or column to the clipboard
Copy	Copies the content of the current cell or column to the clipboard
Advanced Copy	Opens advanced copy submenu that allows copying data with preset formatting parameters
Paste	Pastes the copied content to the cells in focus
Advanced Paste	Pastes several values delimited with a tabulation or line break
Delete	Deletes the row that has the cell in focus NOTE: In fact, when users click **Delete**, the system only highlights the red row while the actual deletion happens when users click **Save**.
Edit cell	- For CLOB/BLOB data format, opens the contents of the cell in a new tab - For all formats except CLOB/BLOB, opens a properties window for the cell
Inline edit	- Same as double-click on a cell - Makes the cell editable
Set to NULL	Sets the value of selected cells to NULL
Hide column	Hides the column currently in focus, see the Managing Display of Columns in Data Table section further in this article
Save to file…	- Appears only for columns with BLOB/CLOB data - Opens the standard Save As window that allows saving data contained in the cell to a file
Load from file…	- Appears only for columns with BLOB/CLOB data - Opens a standard window for opening files
Order /Filter	Displays a submenu that allows selecting filter criteria for the data. The submenu contains the most common filters that can be applied to the cell in focus – see details in Data Filters article. By default, DBeaver filters data by sending a request to the server (the Server-side results ordering checkbox selected). To filter data on the client side using DBeaver`s internal algorithm, clear the checkbox.
View /Format	Opens a submenu that provides tools for formatting and modifying the view of data, see Data View and Format
Navigate	Opens a submenu that helps users navigate throughout the data table, see Navigation
Layout	Changes the layout of data, see the *Table vs. Record Views* section of the Data View and Format article
Export Resultset	Opens the Data Transfer wizard that guides you through the steps to select a format and export data NOTE: The system exports the whole result set including records that are not visible in the screen and preserves all applied data filters and ordering.
Generate SQL	Opens a submenu on which you can select the type of SQL query to generate
Refresh	Refreshes the whole results set including all items that are not visible in the screen

For more information about using the Data Editor, please see the subsections of this article - open them via the contents tree on the right.

Navigation

Scrolling Results Page

If the result set has many rows, you can scroll the results page. To learn how many rows the data table contains, click the **Calculate total row count** button in the bottom toolbar. The number of rows appears in a status field next to the button:

. Alternatively, you can right-click a cell in the table and then click **Navigate -> Row Count** on the context menu.

By default, DBeaver limits the number of rows fetched to **200** (you can change this value in the main toolbar or in preferences). The maximum number of rows that DBeaver fetches to display in the Data tab is specified in the Maximum result-set size field in the main toolbar: ` 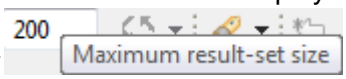

Once you scroll to the last row of the current result portion, DBeaver fetches the next portion (next N rows). You can disable this behavior in preferences. You can also manually fetch the next portion of data equal to the maximum result set size. To do so, click the **Fetch next page of results** button ( ) in the bottom toolbar or right-click the table and click **Navigate -> Fetch next page** on the context menu.

The number of rows fetched is visible in the status field under the data table:

To see the details, click the details button in the status field.

To fetch the whole result set, click the **Fetch all rows** button ( ) in the bottom toolbar or right-click the table and click **Navigate -> Fetch All Data** on the context menu.

NOTE: Be careful when fetching the whole result set. If it is huge, it might cause program hangup or out-of-memory errors.

You can navigate through the result set using standard shortcuts `Home`, `End`, `PgUp`, `PgDown`, `Ctrl+Home`, .

Data Rows

To jump to the first or last row or move one row forward or backward, use the navigation buttons in the bottom toolbar or on the context menu:

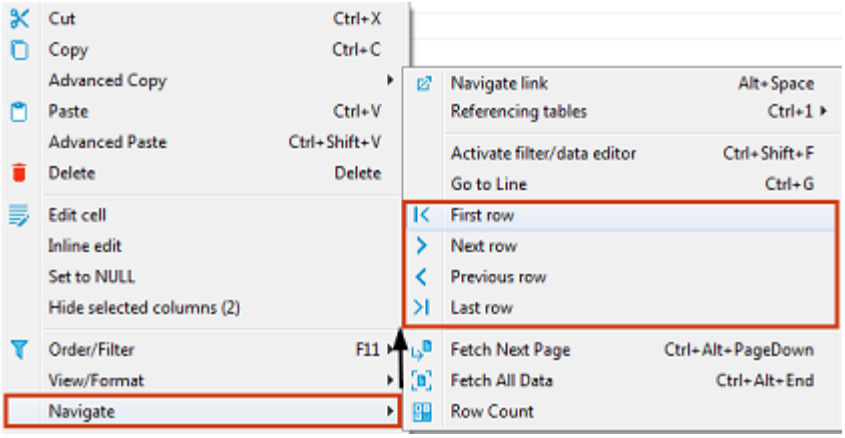

To jump to a specific line, right-click anywhere in the table and click **Navigate** -> **Go to Line** on the context menu. Then in the Go to Row dialog box, enter the row number and click **OK**.

History

DBeaver remembers the history of such actions as applying filters to data, opening reference tables or other tables via links. You can navigate among such tables and filtered views:

- Use the forward and backward buttons in the top toolbar:
- Click `Ctrl+Left` or `Ctrl+Right`

Hovering over these buttons displays the names of the tables or filtered views saved in the history.

Navigate Foreign Keys / Referencing Tables

You can navigate by foreign keys or to referencing tables – those that reference the current table. To open a referencing table, press `Ctrl+1` or right-click the cell and click **Navigate->Referencing tables->[table name]**:

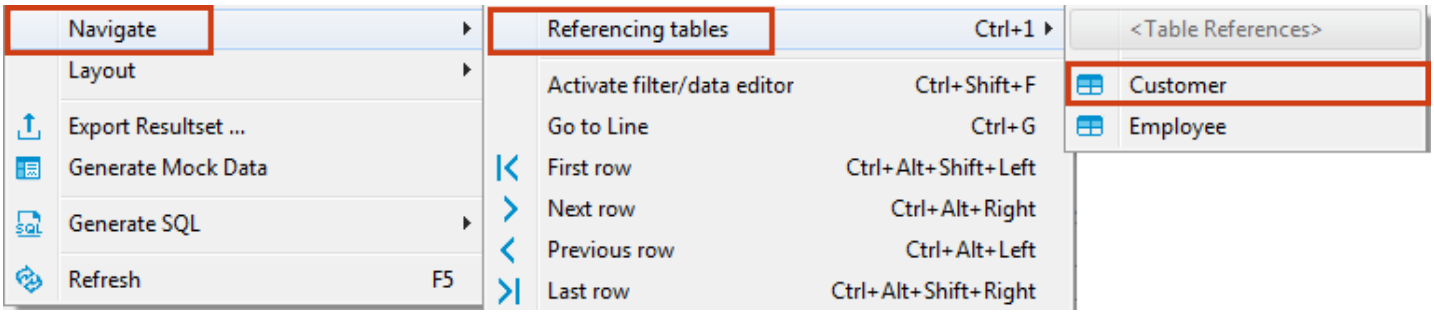

The referencing table opens in the same editor. To navigate back and forth between the initial and referencing tables, use the history navigation buttons ( ) in the top toolbar of the editor. In order to open referencing table in a new window use `Ctrl+Shift+1` shortcut to show menu.

Navigation Links

In the data editor, you can navigate to linked tables – the ones that the current table references. To open a linked table, click the Navigate link icon in a cell that contains it:

	AlbumId	Title	ArtistId
7	7	Facelift	5
8	8	Warner 25 Anos	6
9	9	Plays Metallica By Four Cel	7

Another way is to right-click such a cell and click **Navigate -> Navigate link** on the context menu. The linked table opens in the same editor, filtered by the cell value:

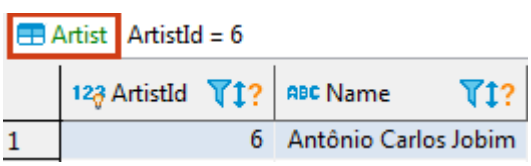

Artist ArtistId = 6

	ArtistId	Name
1	6	Antônio Carlos Jobim

NOTE: The table name in green above the table indicates which table is currently open in the editor.

To navigate back and forth between the initial and linked tables, use the history navigation buttons (⬅ ➡) in the top toolbar of the editor.
You can open a linked table in a separate editor. To do so, simultaneously hold the `Ctrl` key (or `# command` key on macOS) and click the Navigate link icon ( ) in the cell.

Data View and Format

The main tool for managing the appearance of the data table is the Result Set Order/Filter Settings window.

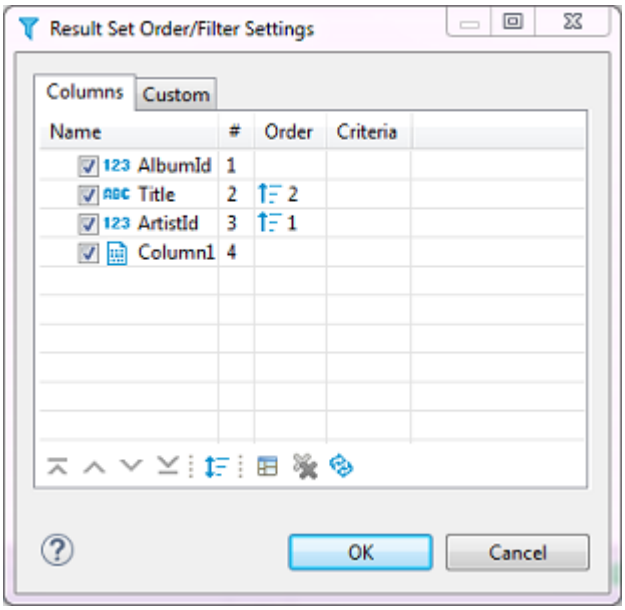

To open this window, click the Custom Filters button ( ) in the top toolbar of the editor or click the Configure button ( ) and then click **Order/Filter** on the dropdown menu.

The Result Set Order/Filter Settings window provides tools to:

- Order data inside columns
- Manage the display of columns in the table
- Manage the order of columns in the table
- Filter data in the table using an SQL expression, see … below

Another tool for managing data appearance are column headers. In the data table, every column header contains three elements each having its own function: Data (column) type icon, column name, filter icon, and ordering icon.

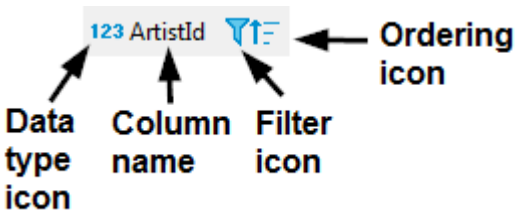

- Simply clicking the column name or column type icon highlights the whole column.
- You can click the column type icon and then drag and drop the column to a different position in the table.
- You can click the column name and then drag the cursor right or left to highlight multiple columns.
- Clicking the ordering icon allows ordering the data in the column in ascending or descending order - see 'Ordering Data in Columns' section further in this article
- Clicking the filter icon allows filtering the data by a cell value, see [TBA]

Ordering Data in Columns

You can order data in columns in one of the ways:

1. Click the ordering icon ( ) in the header of the column.

The icon has three states:

- Clicking once establishes ascending order ( )
- Clicking a second time changes the order to descending ( )
- Clicking a third time removes the ordering from the column ( )

To order data by several columns, go column by column, setting the order with the Ordering icon, starting from the column by which you want to order data first.

2. Click the Custom Filters button ( ) in the top toolbar of the editor to open the Result Set Order/Filter Settings window (see above):

a) Next to the column by which you want to order data in the first turn, set the ascending or descending order using the same three-state principle as described above.
b) Set the ordering in other columns by which you want to sort data in the second, third, etc. turn. The **Order** column indicates the order in which the sorting will happen.
NOTE: The number (**#**) column indicates the initial order of columns.

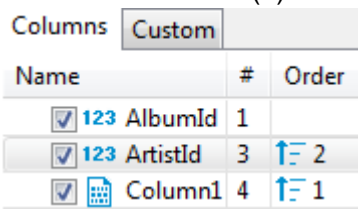

c) To easily move the ordering setting from column to column, you can use the Move up/down/to top/to bottom/ buttons:

To reset the data ordering to its initial state, click the Reset button ( ) in the same window.

Also, to remove all ordering settings, click the Remove All Filters/Orderings button ( ) in the top toolbar of the Data Editor.

Managing Display of Columns in Data Table

To hide a single column, right-click the column or any cell in it and click **View/Format -> Hide column** on the context menu. To unhide a hidden column, open the Result Set Order/Filter Settings window (see the image at the beginning of in this article) and select the checkbox next to the column name or click the Reset button ( ).

To display or hide columns in the data table, in the Result Set Order/Filter Settings window:

1. Select the checkboxes next to the columns that you want to see in the table and clear the checkboxes next to those that you want to hide.

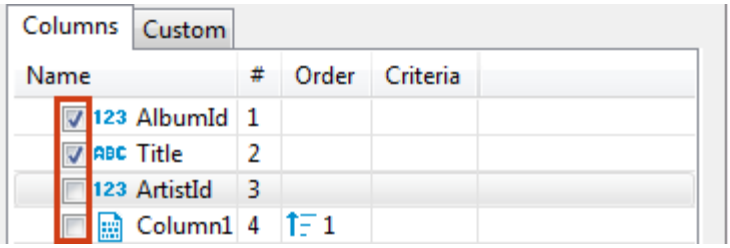

2. Use the Show All ( ) and Show None ( ) buttons at the bottom of the window.

Sorting Columns in Data Table

You can modify the order of columns in the data table in two ways:

1. Click the icon in the column header and drag-and-drop the column to a new position.
2. To sort column alphabetically, in the Result Set Order/Filter Settings window (open by clicking the Custom Filters button ( ) in the top toolbar of the editor), click the Sort button ( )
3. In the Result Set Order/Filter Settings window, click the column to set focus to it and then move it using the navigation buttons: ( )

Grid vs. Plain Text Views

You can switch between two data presentations in SE version and four presentations in EE version. Pressing `CTRL` switches available presentations in turn.

- To see data in a grid view, similar to Excel spreadsheet, click the **Grid** button (Grid) in the bottom toolbar of the editor.
- To switch to the plain text view, click **Text** (Text) in the bottom toolbar.
- To switch to JSON view (available in EE version only for MongoDB documents and JSON tables), click **JSON** in the toolbar.
- To switch to XML view (available in EE version only for XML tables), click **XML** in the toolbar.

Table vs. Record Views

The table view is a standard table (Excel-like) in which columns are vertical and rows are horizontal. This view is the default one. If you click the **Record** button in the bottom toolbar of the editor (Record), or press `Tab` , or right-click a cell and then click **Layout -> Record** on the context menu, the rows and columns switch positions – columns appear as rows, and rows hide in one **Value** column which now shows only one row of data, and column headers shift from the top of the table to its left side:

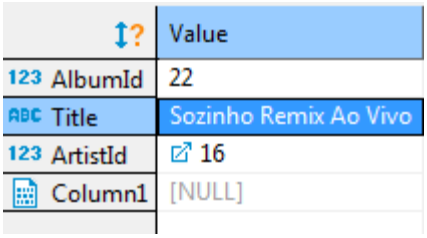

The Record view is useful if the table contains a big number of columns. To navigate from row to row of data, use the navigation buttons in the bottom toolbar of the editor:

To return back to the standard table view, click the **Record** button again.

Rows Coloring

In the data editor, you can color all rows that have the same value as a particular cell of a certain column. To do so, right-click the cell and click **View/Format -> Set row color for {column name = value}** on the context menu:

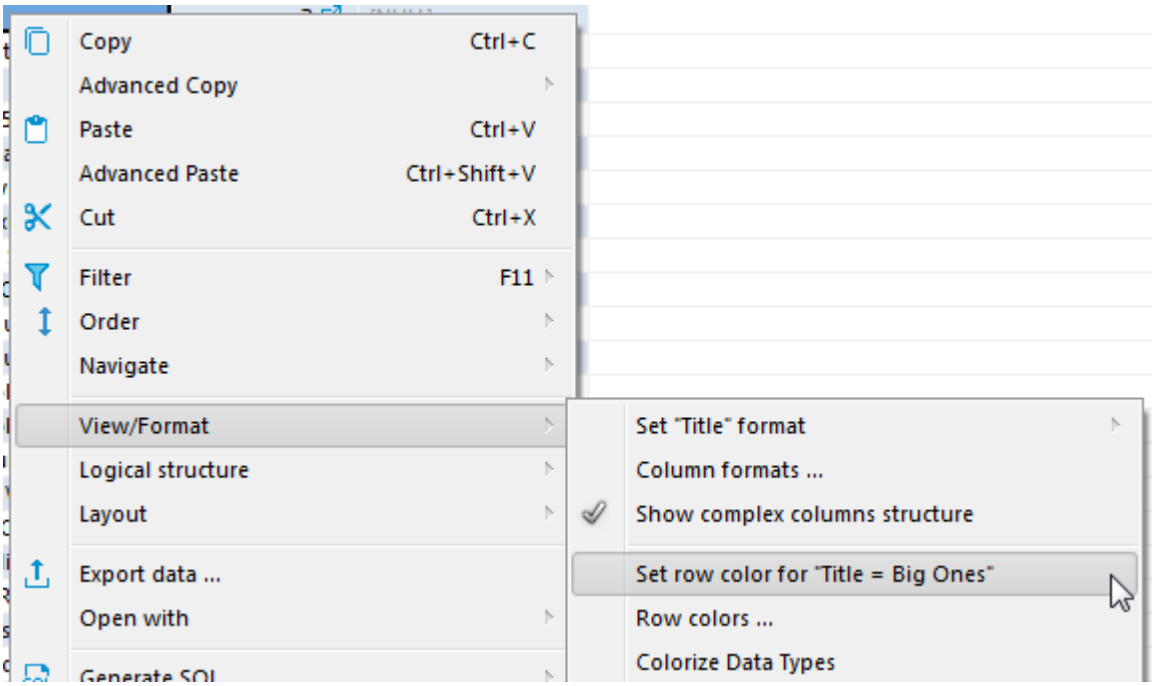

Then choose the color in the palette window that appears and click **OK**. The current row and all other rows that contain the same value change their color to the one you selected:

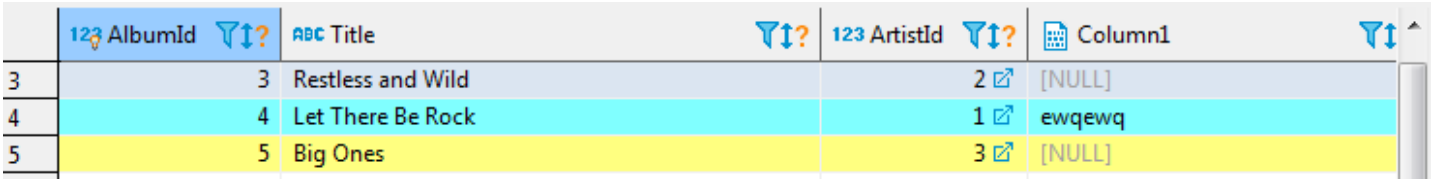

To remove the coloring by a particular column, right-click the cell again and click **View/Format -> Clear color for {column name = value}** on the context menu.

Coloring by Data Types

Besides coloring rows by a value, you can colorize values in columns by data types. To do so, right-click any cell in the table and, on the context menu, click **View/Format -> Colorize Data Types**. Values in cells are colored in different colors according to preferences currently set:

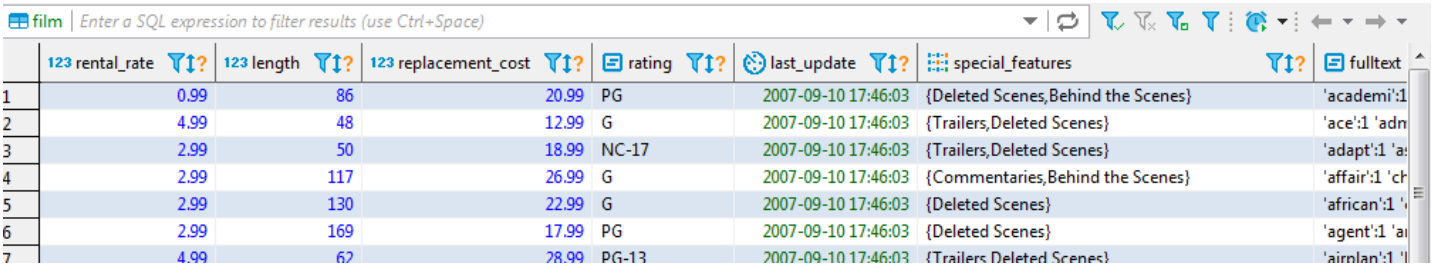

You can change the color preferences in the Preferences window: click **Window -> Preferences** on the main menu. Then, in the window, in the navigation pane on the left, expand **User Interface** and then **Appearance**, and then click **Colors and Fonts**:

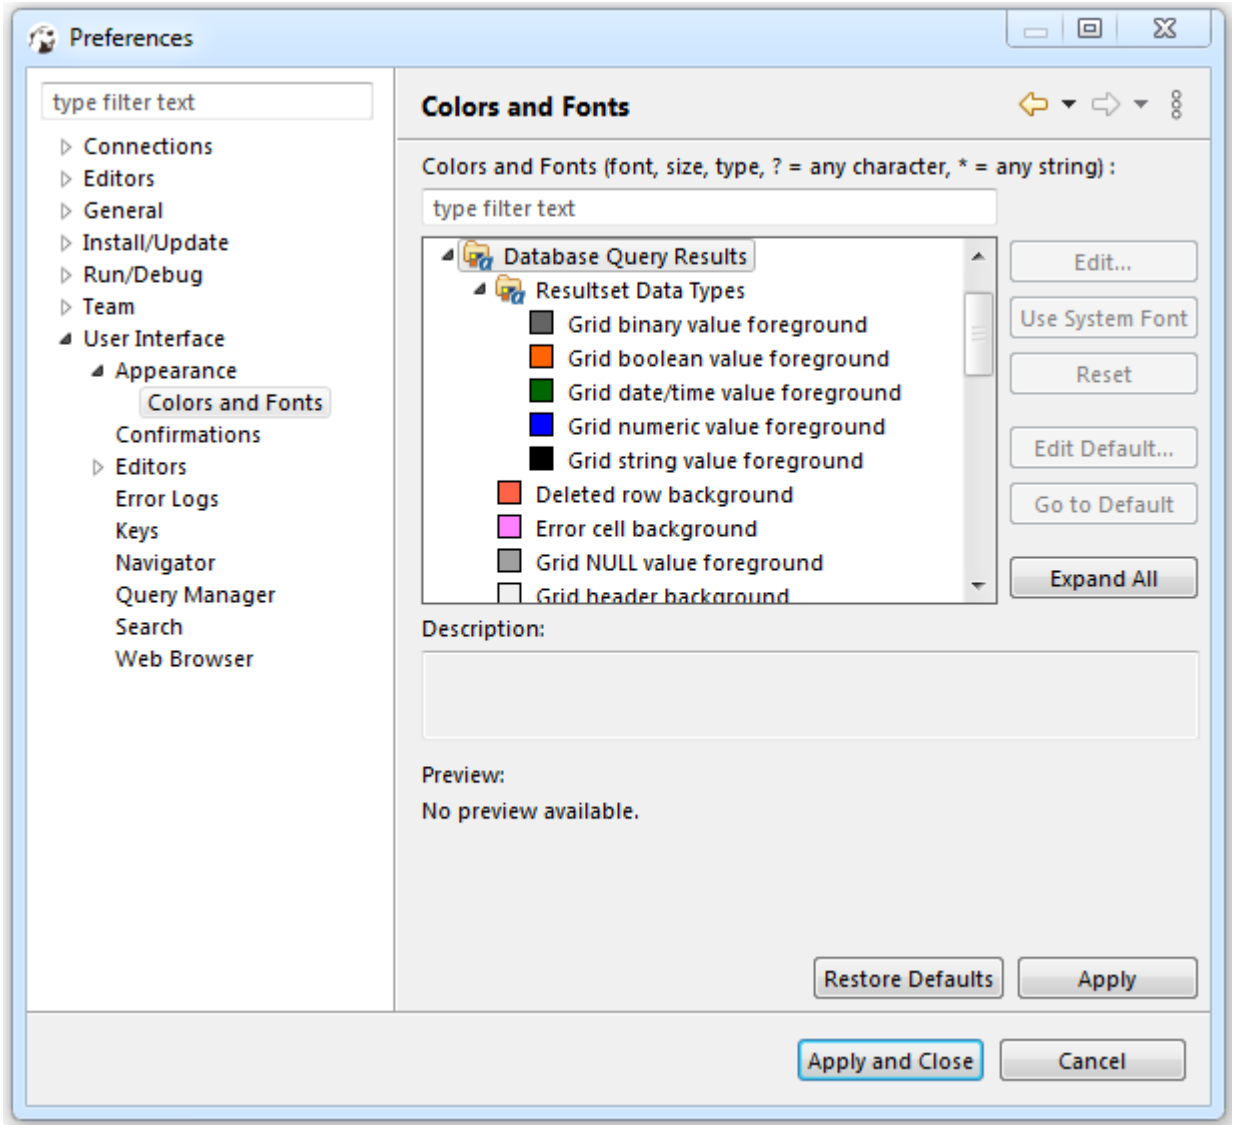

To remove coloring by data types, on the context menu, click **View/Format -> Colorize Data Types** again.

Transforming Data Presentation

For string and numeric data types, DBeaver provides tools to transform the data presentation into a number of formats, such as URL and Binary for strings and Epoch Time, Number Radix, etc. for numbers. To change the

data presentation in a certain column, right-click a cell in the column, then, on the context menu, click **View/Format -> Set {column name} format** and then click the presentation type name:

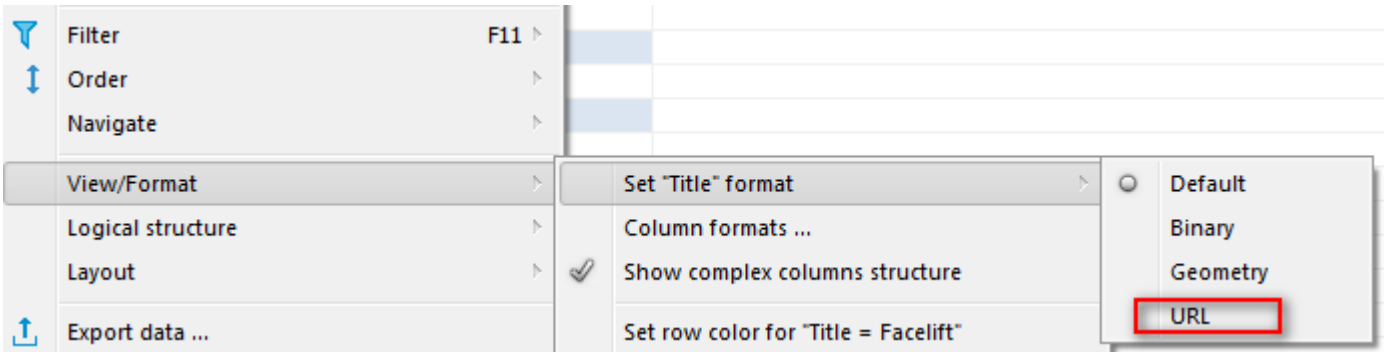

The Transformer settings window opens showing the value in the chosen format. Click **OK** to apply the change:

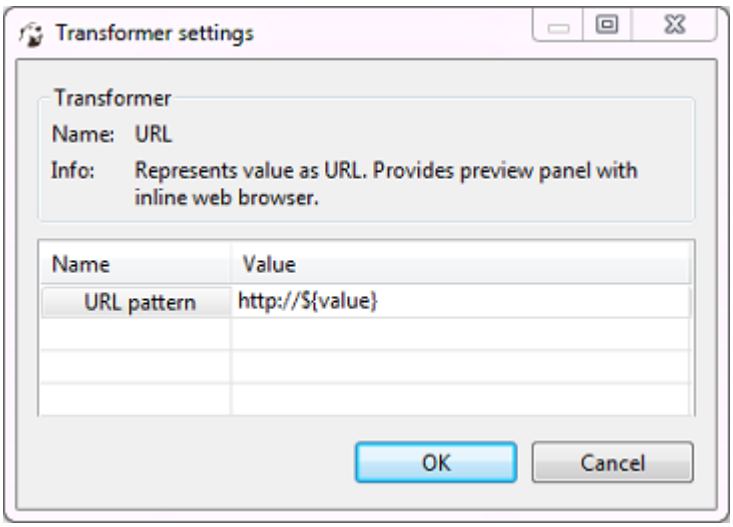

The values in the column appear in the new format.
NOTE: For URL format, the resulting cell provides a link to the URL in a browser window.

To roll back the changes to the default format, right-click any cell in the column, and on the context menu, click **View/Format -> View as -> Default**.

Structurizing Complex Data Types

For complex data types (that themselves represent a structure), such as objects, structures and arrays, DBeaver provides a tool for breaking them into columns:

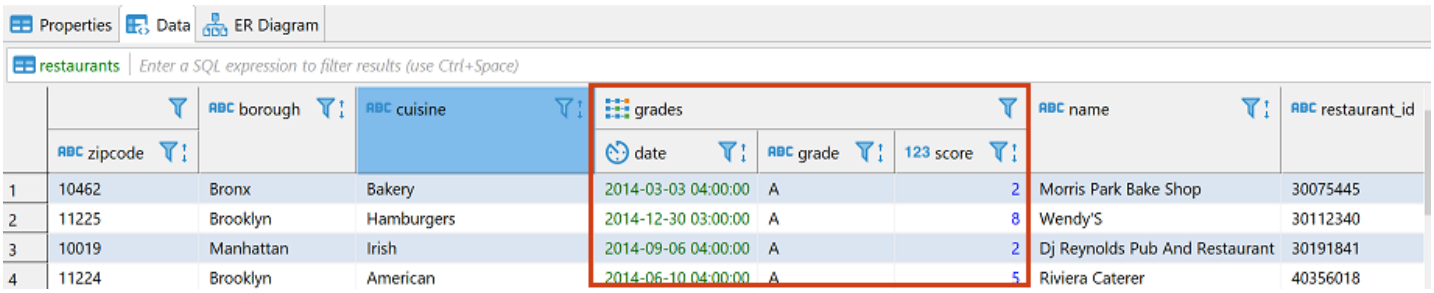

To do so, right-click a cell in the column and, on the context menu, click **View/Format -> Visualize complex columns**.

Configuring Numeric and Time Data Formats

You can specify the exact format of Time, Timestamp, Date, and Number data used in the currently open database or globally. To specify a format, right-click any cell in the table and, on the context menu, click **View /Format -> Data formats**. The Properties window opens displaying the **Data Formats** page:

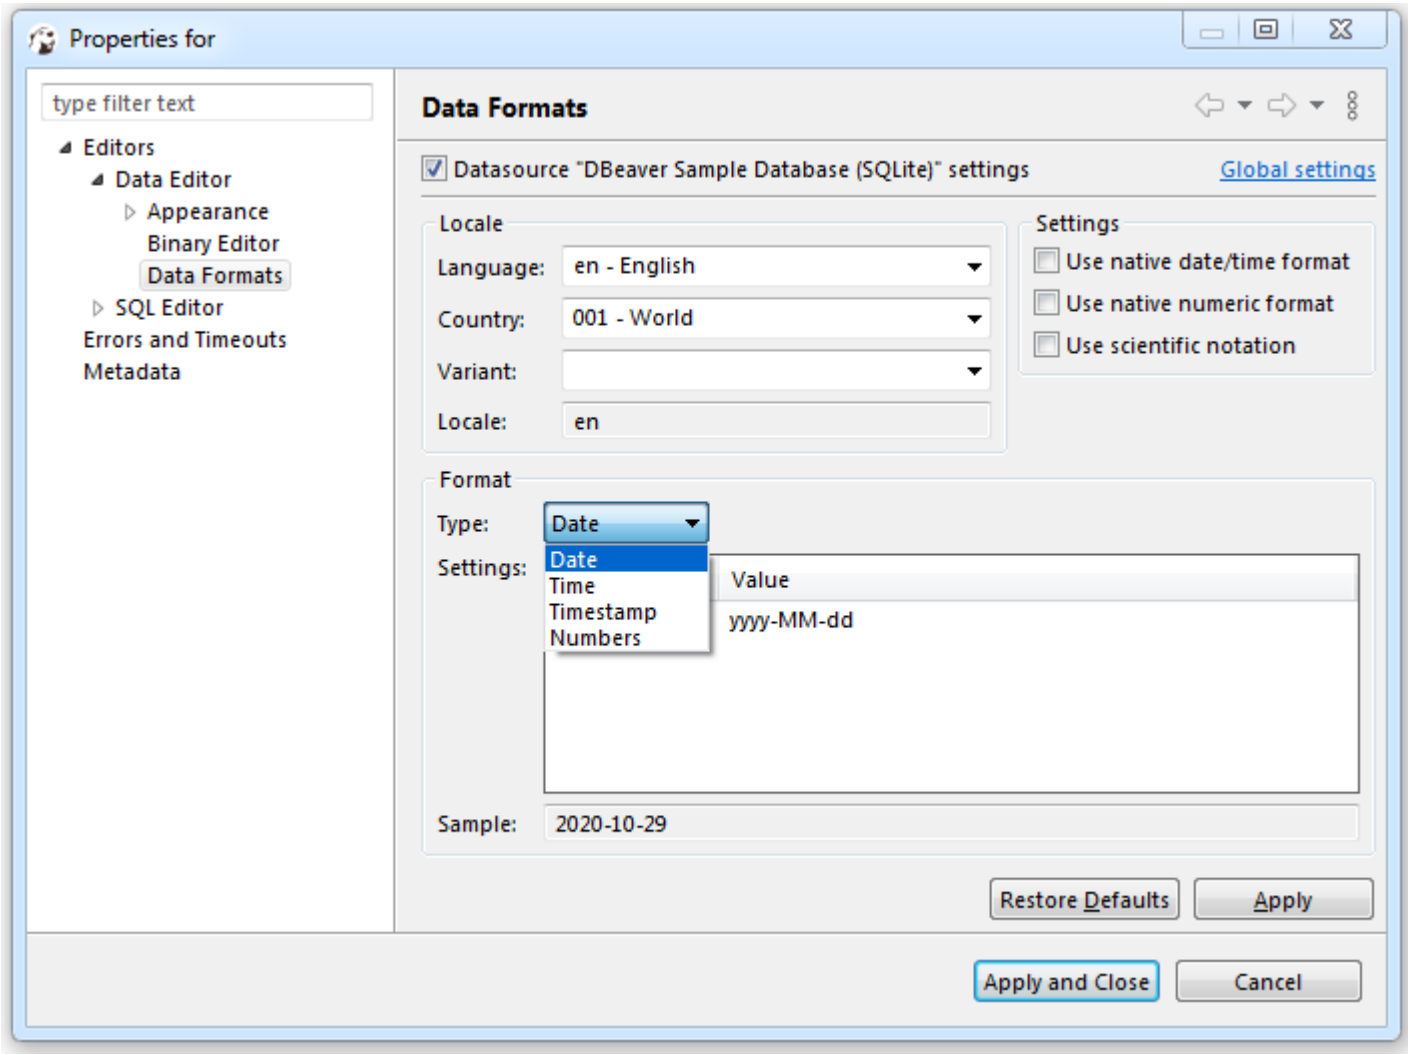

To configure the format for the current database only, select the **Datasource "[Connection name]" settings** checkbox. To configure the settings globally, to all databases that you have in DBeaver, click **Global settings**.
You can specify the locale for the data format in the **Locale** area, then, in the **Type** dropdown list, click the name of the data type and then, in the **Settings** table, click the required format.
To apply the changes and make them visible in the table, click **Apply and Close** and then refresh the window (F5).

Data Filters

You can apply custom filters to any table contents or query results. There are several ways in which you can filter data in the table.

One of the ways is to use the filter field above the table next to the top toolbar. To filter data, enter an SQL expression into the field and click the Apply filter criteria button ( ) next to the field or press `Enter`.

You can apply ready-to-use SQL expressions or SQL expression templates via the context menu. To select a ready SQL expression or a template, press `F11` or right-click the cell, then click **Order/Filter** on the context menu and then click one of the expressions.

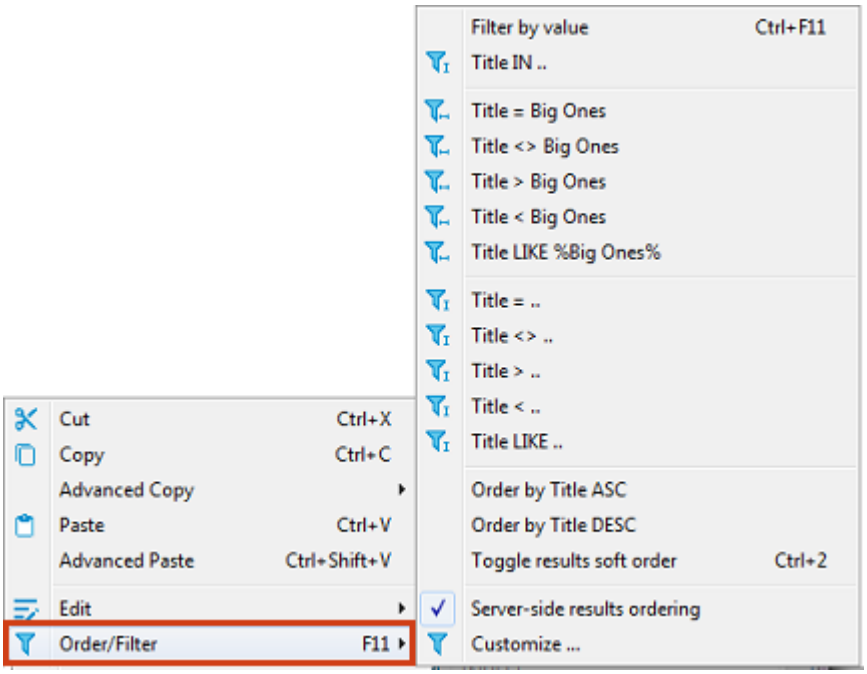

The third way is to filter data by a cell value using the filter icon in the column header. To filter data this way, click the filter icon in the column header and then double-click the cell value in the Filter by column value dialog box:

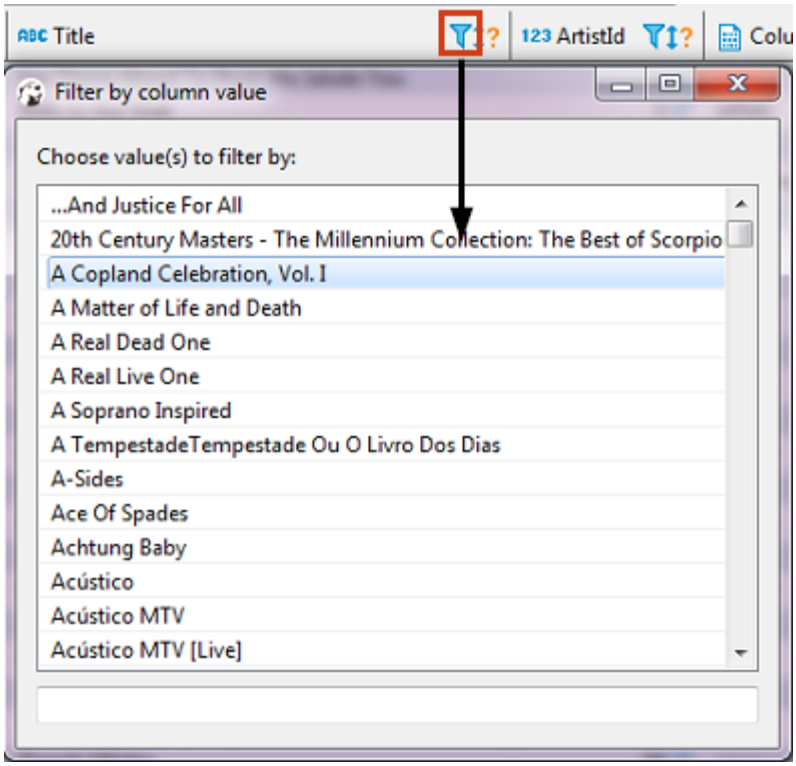

The data updates dynamically. To remove a filter, click the Remove All Filters/Orderings button ( ) in the top toolbar of the editor.

You can save the current filter settings for the database object to apply next time when you reopen it in the editor. To save the current filter settings, click the Save filter settings for current object button ( ) in the top toolbar.

Data Refresh

Refresh is necessary if the database contains changes made by other users working on it simultaneously with you and you want to see them in your DBeaver window. To refresh data manually, right-click anywhere in the data table and click **Refresh** on the context menu or press `F5`.

You can also schedule auto-refresh to happen on a regular basis. To auto-refresh the database on schedule:

1. Click the Configure auto-refresh button ( ) on the top toolbar of the editor. The Auto-refresh configuration dialog box opens:

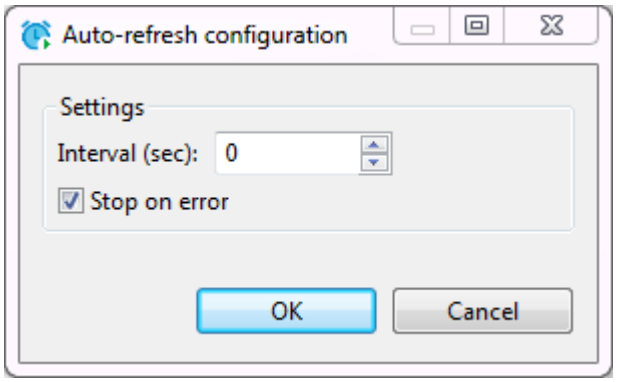

   a) Set **Interval** in seconds.
   b) Select the **Stop on error** checkbox if you want the refresh to stop when it encounters an error or clear it, if the refresh should ignore errors.
   c) Click **OK**.

2. Alternatively, click the arrow next to the Configure auto-refresh button ( ) to open the auto-refresh menu:

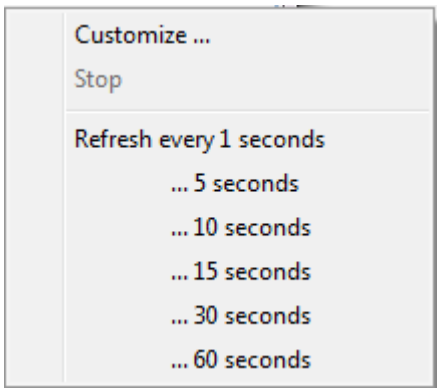

   On the menu, you can click one of the preset options or click **Customize** to open the Auto-refresh configuration dialog box, see option 1.

When you perform either of the two alternative options above, the system starts refreshing the data as scheduled and the Configure auto-refresh button changes to **Stop auto-refresh** button ( ). To stop the auto-refresh, click the **Stop auto-refresh** button or click the arrow next to it and click **Stop** on the auto-refresh menu.

Data Viewing and Editing

You can do inline editing (see the *Inline Editing* section below) as well as open the content of a cell in a separate editor (see the *Cell Editor* section below).

When you make any changes to the data and save them using steps described in this section, the changes apply to the database itself. Prior to saving the changes, you can preview the SQL script that the system sends to the database to apply the changes there. To see the SQL script, after making changes and before saving or discarding them, click the **Script** button (Script) in the bottom toolbar. The Preview Changes window opens, in which you can only view the SQL script and copy it, if necessary:

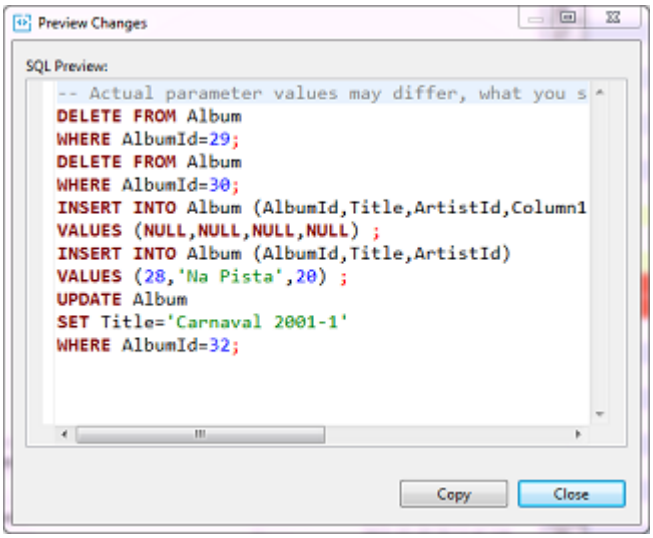

Inline Editing

Inline editing is when you modify the content right in the cell. To edit a cell inline, in the table do one of the following:

- Double-click the cell.
- Click the cell to set focus to it and press `Enter`.
- Right-click the cell and click **Inline edit** on the context menu.

The cell becomes editable, now you can change its value.

To set the cell value to NULL, right-click the cell and click **Set to NULL** on the context menu.

To save the changes, click the **Save** button (Save) in the bottom toolbar. To discard the changes, click the **Cancel** button (Cancel) in the bottom toolbar.

NOTE: Both the **Save** and **Cancel** buttons become editable only when you make changes in a cell and then jump to another cell.

Cell Editor

To edit data in a cell using a separate editor, do one of the following:

- Right-click the cell and click **Edit cell** on the context menu.
- Click the cell to set focus to it and press `Shift+Enter` or click the **Edit cell value in separate dialog/editor** button ( ) in the bottom toolbar.

For cells of CLOB/BLOB data format, this action opens the contents of the cell in a new tab. For all other formats except CLOB/BLOB, this action opens a properties window for the cell:

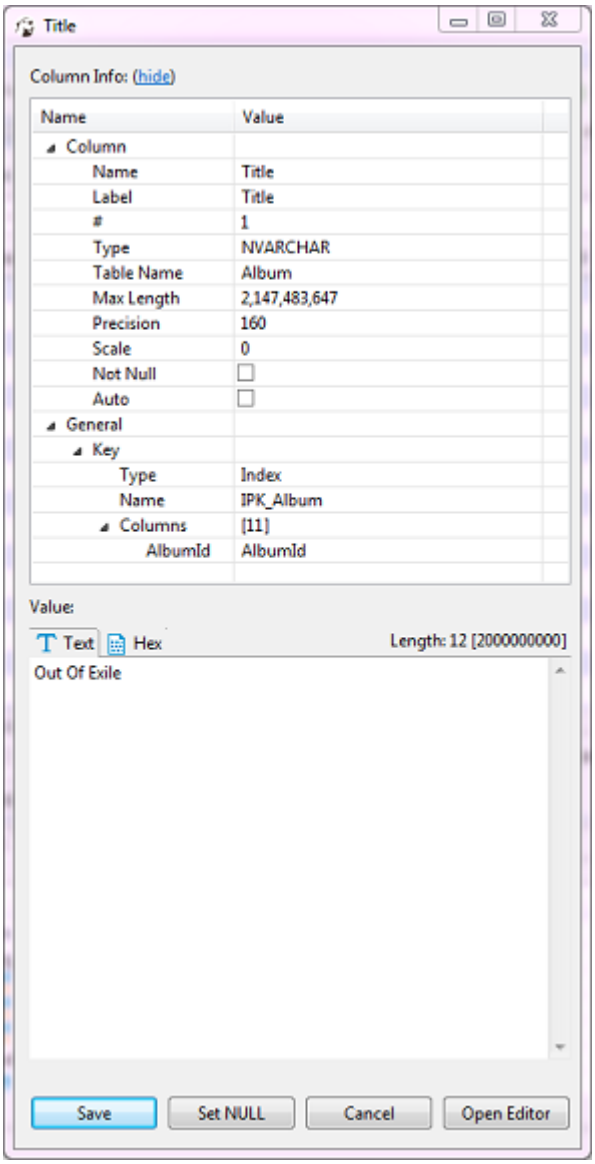

The window displays properties of the column in the **Column Info** section and provides the **Value** section where you can modify the value of the cell. Edit the value as required and click **Save**. To set the value to NULL, click **Set NULL**. To continue editing the cell in a separate editor (tab), click **Open Editor**.

NOTE: DBeaver has full support of CLOB/BLOB data types. You can view values, edit them and save back to the database. You can open CLOB/BLOB value in a separate editor (press `Shift+Enter` on a selected cell). You can save/load LOB value to/from regular files. DBeaver can recognize that some BLOB column keeps images (gif, png, jpeg, bmp). In this case DBeaver shows LOB contents as image. It is convenient to open value view panel (press `F7`) and browse images.

Adding, Copying and Deleting Rows

You can add an empty row below the row in focus. To add an empty row, click the **Add new row** button ( ) in the bottom toolbar. Use inline editing or open cell values in a separate editor to populate them with data (see the sections above).

You can copy any row or several rows currently in focus. To copy rows, highlight one or more rows and click the **Duplicate current row** button ( ) in the bottom toolbar. The duplicate rows appear below the rows in focus.

To delete a row or rows, set focus to the rows and click the **Delete current row** button ( ) in the bottom toolbar. The rows are colored red, which means that they are marked for deletion and will be deleted when you save the changes.

To save any of such changes, click the **Save** button (Save) in the bottom toolbar. To discard the changes, click the **Cancel** button (Cancel) in the bottom toolbar.

Copying/Pasting Cells

To copy the content of one or several cells to the clipboard in TAB-delimited format, press `Ctrl+C` or right-click the cell or cell selection and click **Copy** on the context menu. Then you can paste the copied selection into some spreadsheet editor (similar to Excel).

DBeaver provides the advanced copy option that allows configuring additional copy settings (copy with column names/row numbers, configure delimiter and choose value format). To copy cells with additional settings, press `Ct` or right click the cell(s) and click **Advanced Copy** on the context menu.

Pressing `Ctrl+V` on a cell pastes the copied content into the cell applying appropriate data type conversion. The **Advanced Paste** option on the context menu or pressing `Ctrl+Shift+V` pastes several cells.

Defining Virtual Keys

To be able to persist column value changes, a table must have some unique key (primary key or unique index). Some databases (Oracle, DB2, PostgreSQL) support a special virtual unique column that DBeaver can use to save changes. In other cases, you can define a virtual key – a set of columns that forms a unique combination of values. When you try to save changes in a table without a unique key, DBeaver displays the following error message:

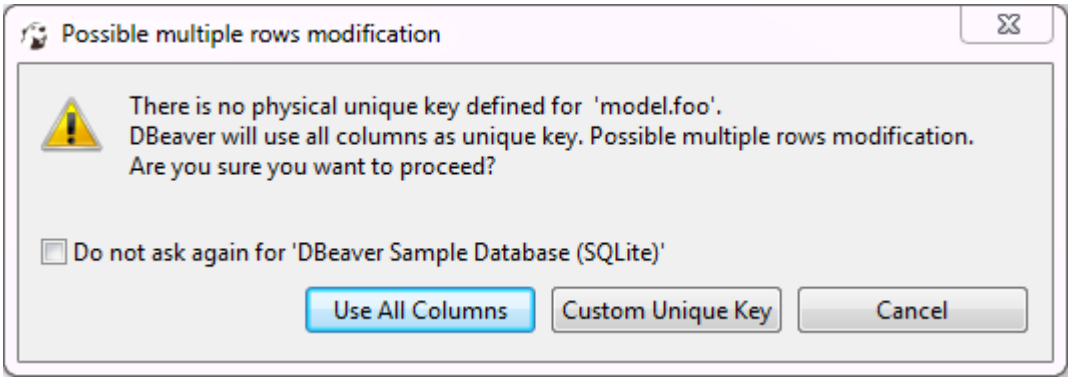

To use all columns as the virtual key, click **Use All Columns**. To create a custom key, click **Custom Unique Key**.

Alternatively, to create a custom unique key, you can click the **Configure** ( ) button in the bottom toolbar and then click **Define virtual unique key** on the Configure menu. The Define virtual unique identifier window opens:

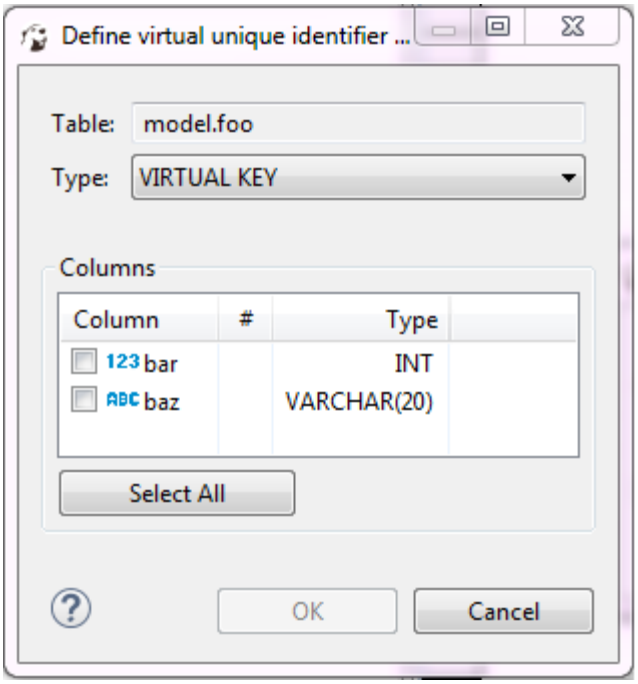

To define the key, select some of the columns or click **Select All** and then click **OK**. To remove a unique key from a table, click the **Configure** button in the bottom toolbar and then click **Clear virtual unique key**.

Panels

Panels provide additional space in the Data editor in which you can manipulate with data. The panels are handy if you work with complex types (structures, arrays), or long text data, or BLOBs. Panels appear as tabs in an additional pane in the right part of the Data tab:

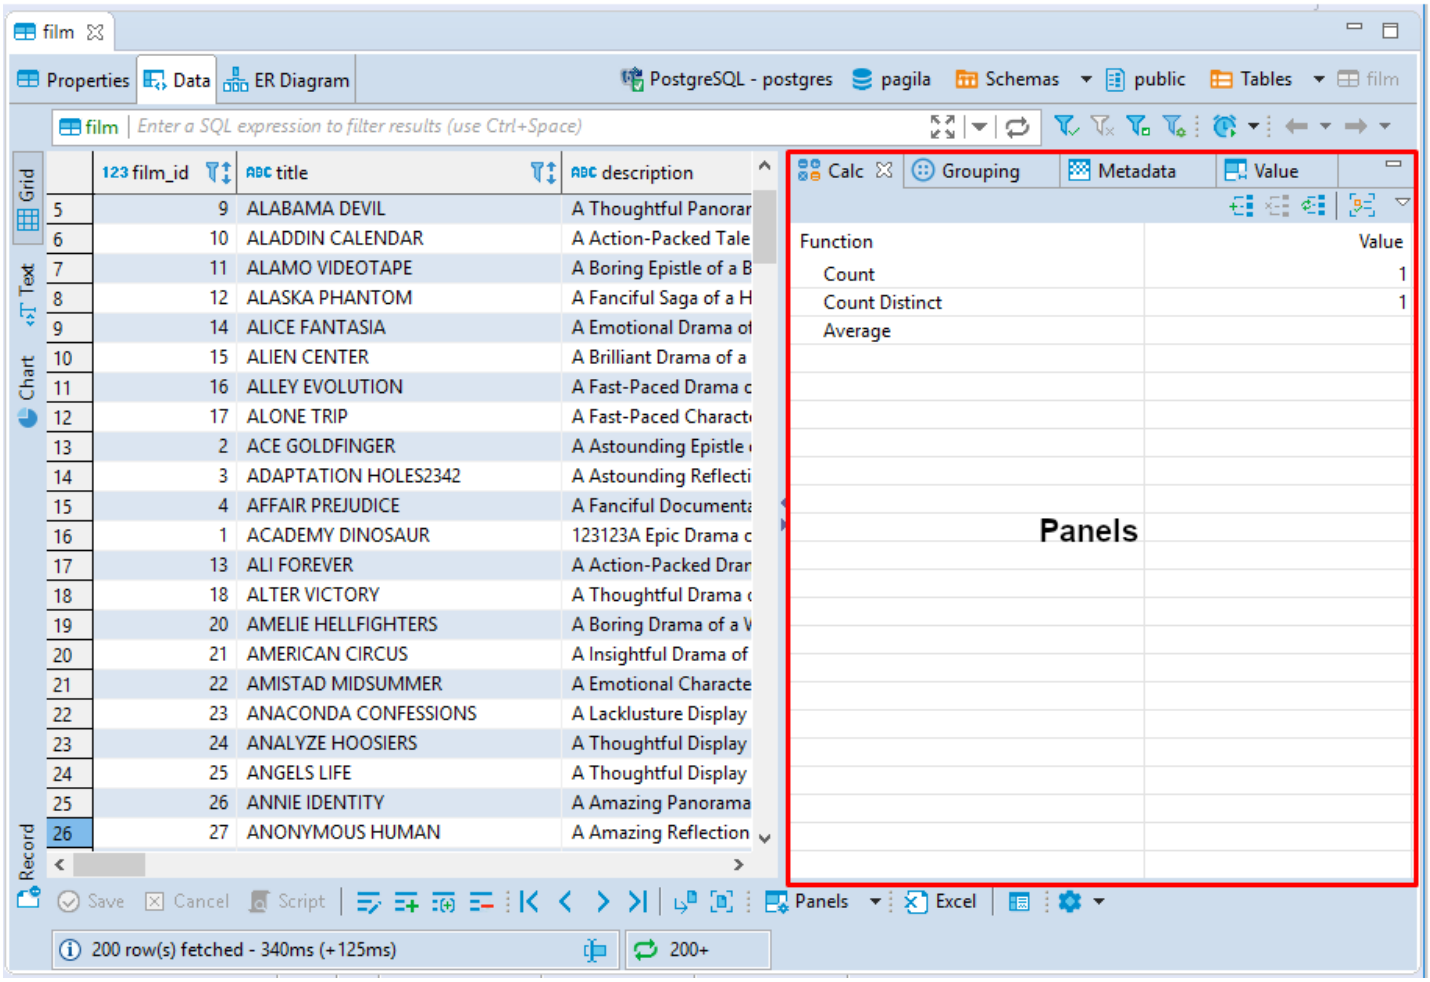

This additional pane appears only when you open one of the four panels:

- Calc
- Grouping
- Metadata
- Value viewer (default)

To open panels, click **Panels** in the bottom toolbar. By default, the Value viewer panel opens. Alternatively, you can open the Value panel by pressing F7 on a cell. To open the other panels, click the down arrow next to the **Panels** button and click the name of the panel on the menu:

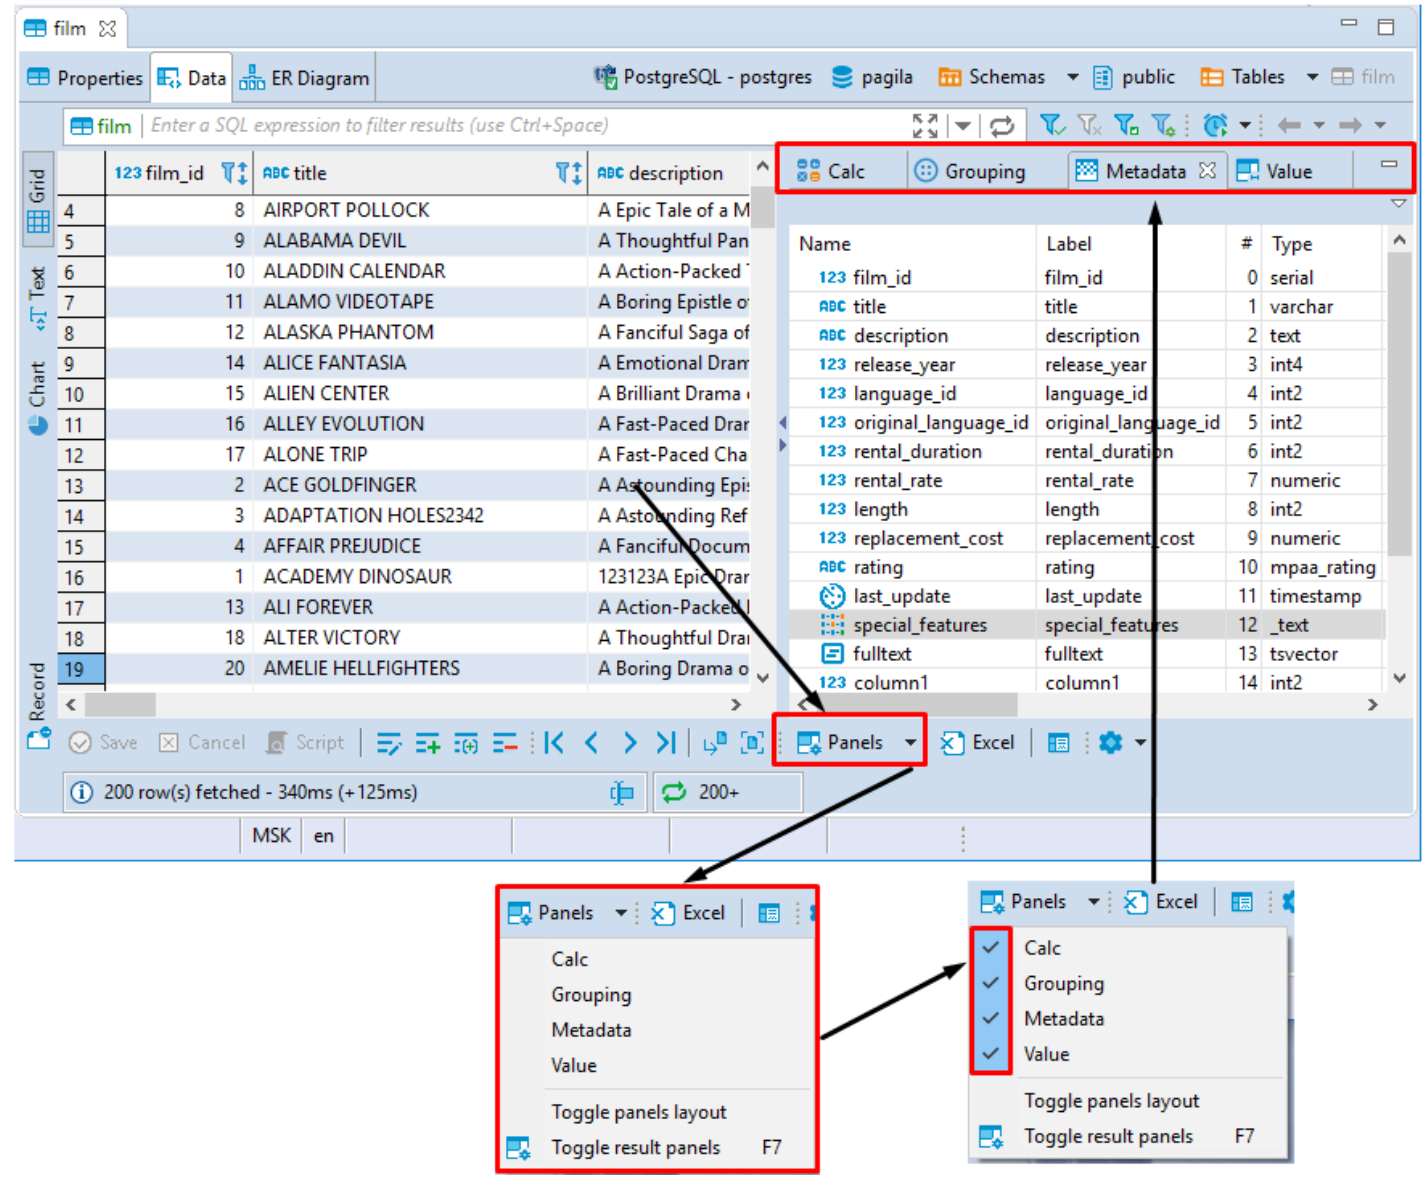

Panels also open if you try to inline-edit a cell with a complex data type.

To close panels, click the **Panels** button again or click the standard Close (cross) icon in the upper right corner of each panel.

You can also show and hide panels by clicking the **Configure** button ( ) in the bottom toolbar and then **Toggle result panels** on the Configure dropdown menu.

Value Viewer

The Value viewer panel displays just one value that is currently in focus and allows editing it.

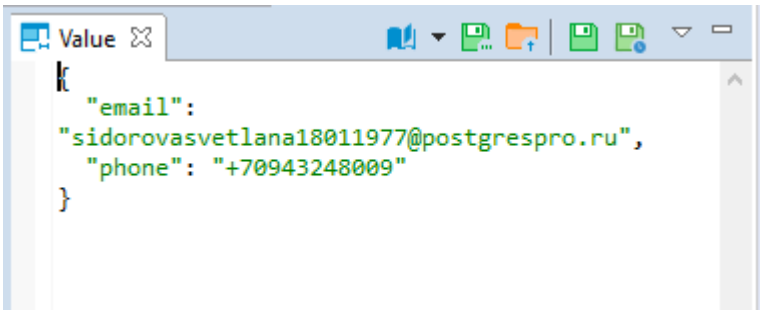

The toolbar of Value panel contains the following buttons:

Button	Name	Description
	Content viewer settings	Opens a menu with a set of options for content view change.
	Save to file...	Allows saving the content to a local file. **NOTE**: This button is only available for XML, JSON and Binary content.
	Load from file...	Allows uploading data from a local file. **NOTE**: This button is only available for XML,JSON and Binary content.
	Apply cell value	Displays in the data table the changes made in the Value viewer. **NOTE**: This button does not save changes made to the database. To save the changes in the database, you need to use the **Save** button in the bottom toolbar of the Data Editor..
	Auto-apply value	Enables automatic display of changes made in the Value viewer in the data table. When auto-saving is enabled, the changes appear in the data table at the same time when they are made in the Value viewer.**NOTE**: This button does not save changes made to the database. To save the changes in the database, you need to use the **Save** button in the bottom toolbar of the Data Editor.

Metadata Panel

The Metadata panel displays metadata for each cell in the row containing the cell currently in focus. You can just view the metadata.

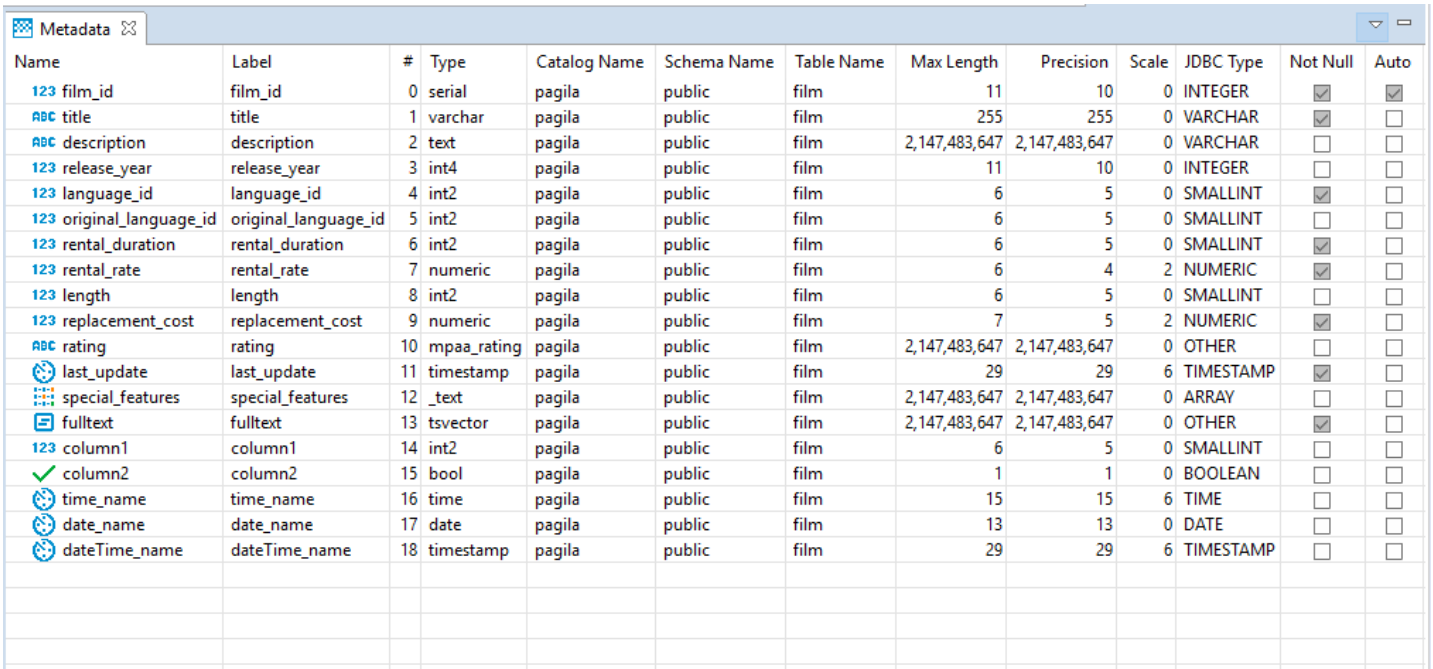

Name	Label	#	Type	Catalog Name	Schema Name	Table Name	Max Length	Precision	Scale	JDBC Type	Not Null	Auto
film_id	film_id	0	serial	pagila	public	film	11	10	0	INTEGER	☑	☑
title	title	1	varchar	pagila	public	film	255	255	0	VARCHAR	☑	☐
description	description	2	text	pagila	public	film	2,147,483,647	2,147,483,647	0	VARCHAR	☐	☐
release_year	release_year	3	int4	pagila	public	film	11	10	0	INTEGER	☐	☐
language_id	language_id	4	int2	pagila	public	film	6	5	0	SMALLINT	☑	☐
original_language_id	original_language_id	5	int2	pagila	public	film	6	5	0	SMALLINT	☐	☐
rental_duration	rental_duration	6	int2	pagila	public	film	6	5	0	SMALLINT	☑	☐
rental_rate	rental_rate	7	numeric	pagila	public	film	6	4	2	NUMERIC	☑	☐
length	length	8	int2	pagila	public	film	6	5	0	SMALLINT	☐	☐
replacement_cost	replacement_cost	9	numeric	pagila	public	film	7	5	2	NUMERIC	☑	☐
rating	rating	10	mpaa_rating	pagila	public	film	2,147,483,647	2,147,483,647	0	OTHER	☐	☐
last_update	last_update	11	timestamp	pagila	public	film	29	29	6	TIMESTAMP	☑	☐
special_features	special_features	12	_text	pagila	public	film	2,147,483,647	2,147,483,647	0	ARRAY	☐	☐
fulltext	fulltext	13	tsvector	pagila	public	film	2,147,483,647	2,147,483,647	0	OTHER	☑	☐
column1	column1	14	int2	pagila	public	film	6	5	0	SMALLINT	☐	☐
column2	column2	15	bool	pagila	public	film	1	1	0	BOOLEAN	☐	☐
time_name	time_name	16	time	pagila	public	film	15	15	6	TIME	☐	☐
date_name	date_name	17	date	pagila	public	film	13	13	0	DATE	☐	☐
dateTime_name	dateTime_name	18	timestamp	pagila	public	film	29	29	6	TIMESTAMP	☐	☐

Calc Panel

The Calc panel is useful for getting basic statistics across data in several columns and rows:

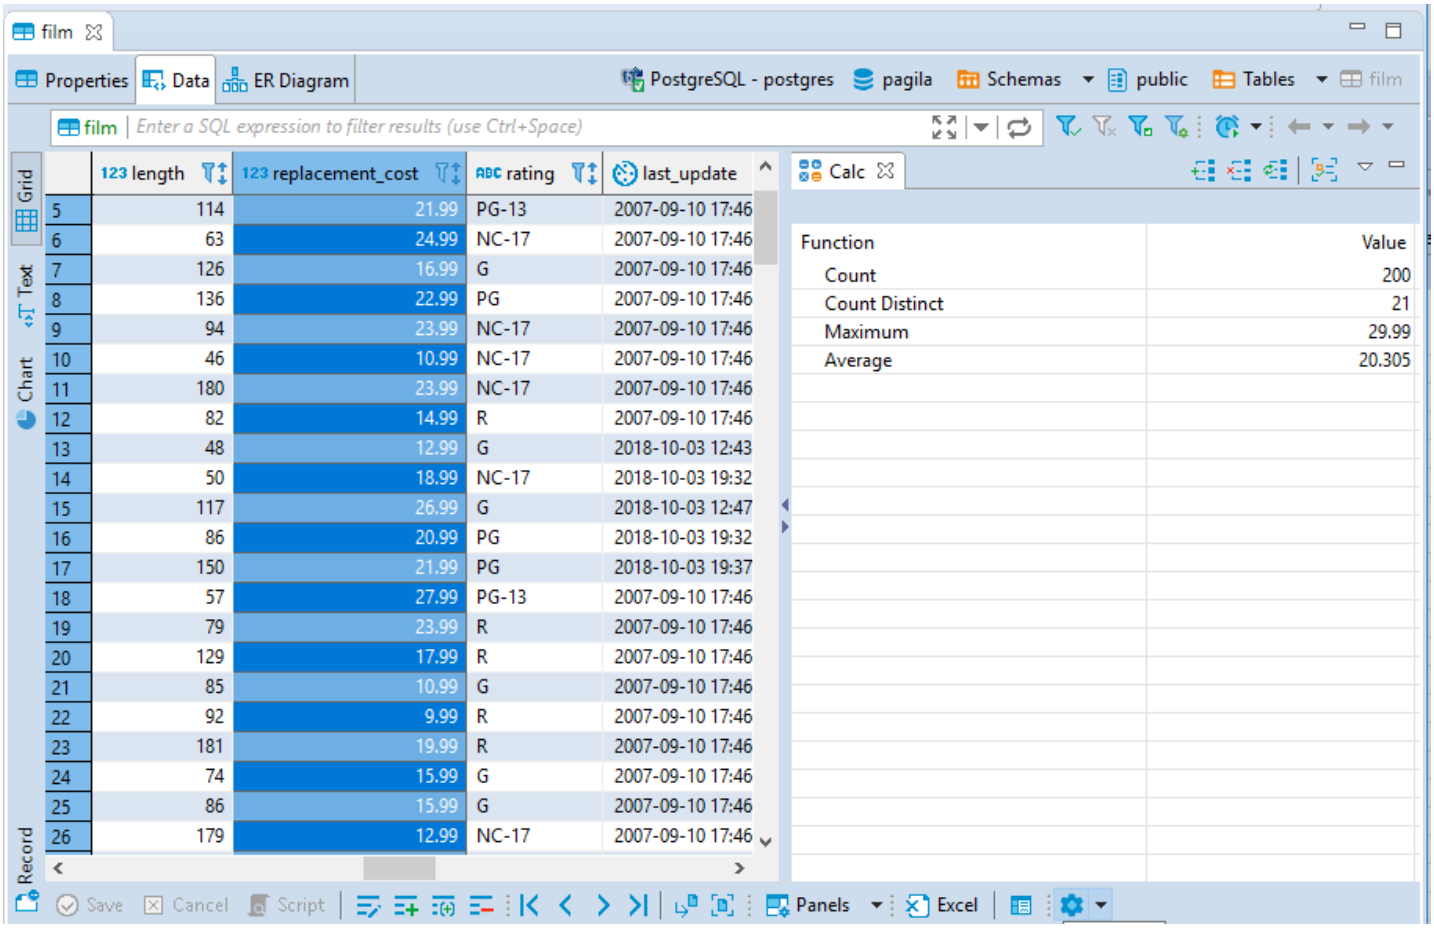

You can select several columns and rows in standard ways - by pressing and holding the left mouse button or by clicking cells while holding the `Ctrl` or `Shift` keys. The panel updates dynamically to shows statistics for the selected data.

To see data grouped by columns, click the Group by columns button ( ). To remove the grouping by columns and see summary values for all columns, click the same button again.

By default, the panel applies and displays results for two functions – **Count** and **Count Distinct**. To add other functions, click the **Add function** ( ) button in the toolbar of the panel or right-click one of the rows in the Aggregate panel and click **Add function** on the context menu and then click the name of the function. The following functions are available:

- Sum
- Average
- Minimum
- Maximum
- Median
- Mode

To remove an individual function, click the function and then click **Remove function** ( ) in the toolbar of the panel, or right-click the function and click **Remove function** on the context menu. To remove all functions, click **Reset** ( ) in the toolbar or on the context menu.

You can copy the value of a particular function to the clipboard - right-click the row and click **Copy Value** on the context menu.
You can also copy all functions with their values - right-click in the table and click **Copy All** on the context menu.

Grouping Panel

The Grouping panel provides tools to calculate statistics based on a table of a custom SQL query. It uses GROUP BY queries to extract unique values for COUNT (default), SUM, AVG, MIN, MAX and other analytics functions displaying the results in dedicated columns.
To obtain the grouping results for one or more columns of a data table, open the Grouping panel, then, in the results table, put the cursor onto the data type icon of the table header so that the cursor turns into a hand pointer ( ), and drag-n-drop the column(s) onto the panel:

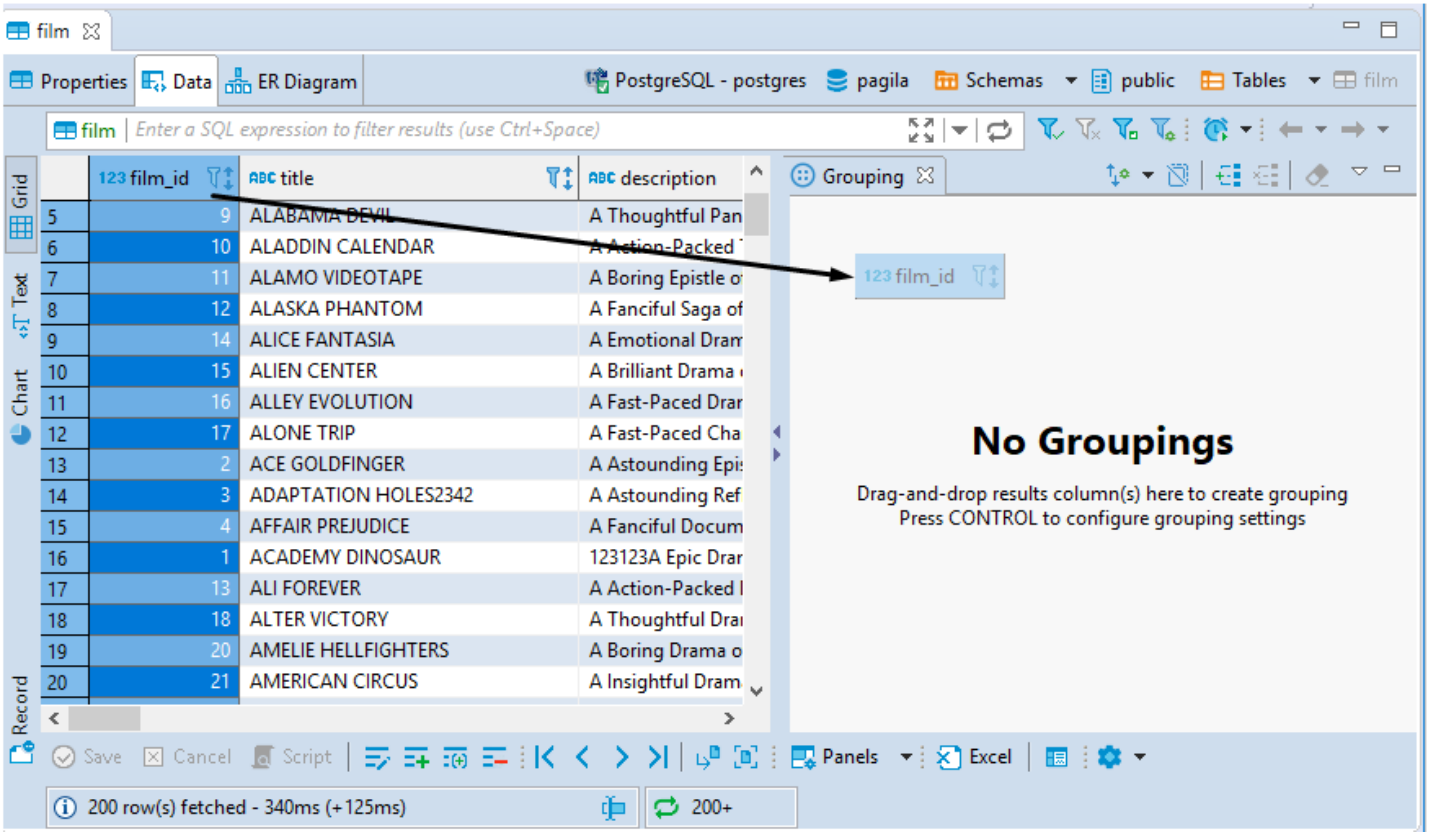

If you add several columns to the panel, DBeaver groups data in the order in which the columns go and calculates statistics based on the grouping.

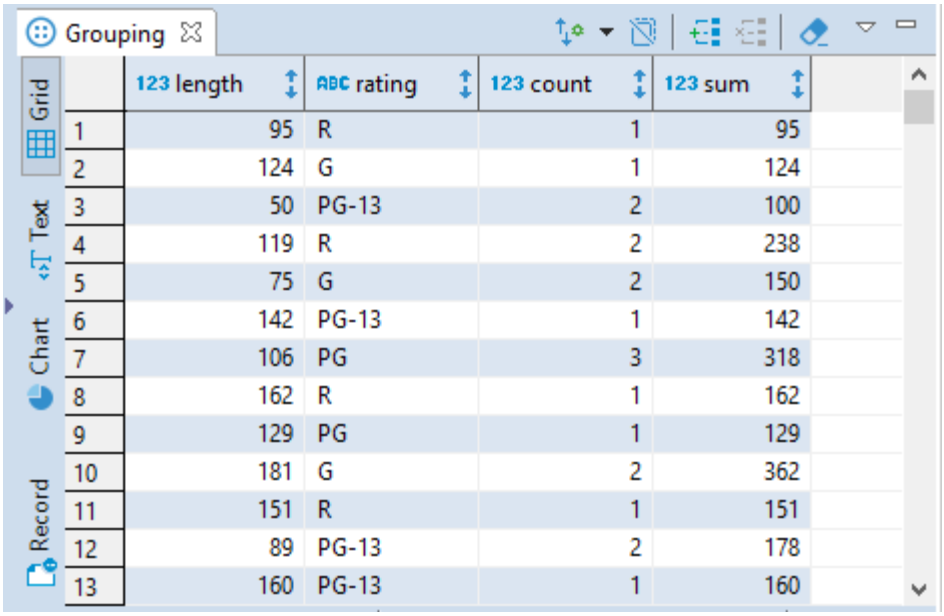

	123 length	ABC rating	123 count	123 sum
1	95	R	1	95
2	124	G	1	124
3	50	PG-13	2	100
4	119	R	2	238
5	75	G	2	150
6	142	PG-13	1	142
7	106	PG	3	318
8	162	R	1	162
9	129	PG	1	129
10	181	G	2	362
11	151	R	1	151
12	89	PG-13	2	178
13	160	PG-13	1	160

By default, the COUNT function is used. You can add other functions as well. To add a function:

1. Click the **Edit grouping columns** button in the panel`s toolbar.
2. In the Grouping Configuration window, in the **Functions** area, click **Add**, then type the function into the new row:
    - You can use auto-complete options DBeaver provides.
    - You need to indicate the column name in brackets. COUNT is the only function that supports `*` instead of column name.
3. Click **OK**:

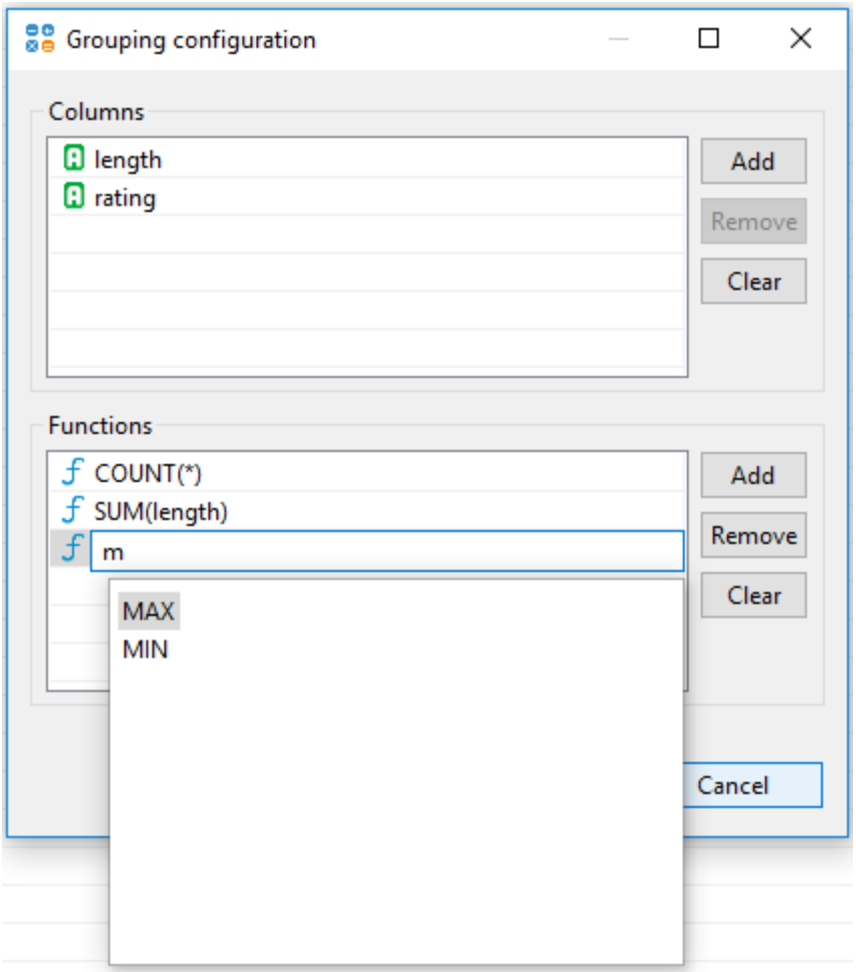

To remove a function, in the same Grouping Configuration window, click the function and click **Remove** and then **OK**. To remove all functions, click **Clear** and then **OK**.

You can also add or remove columns using the same Grouping Configuration window. To add a column:

1. Click the **Edit grouping columns** button in the panel`s toolbar.
2. In the Grouping Configuration window, in the **Columns** area, click **Add**, then type the column name into the new row (you can use auto-complete options DBeaver provides), and then click **OK**:

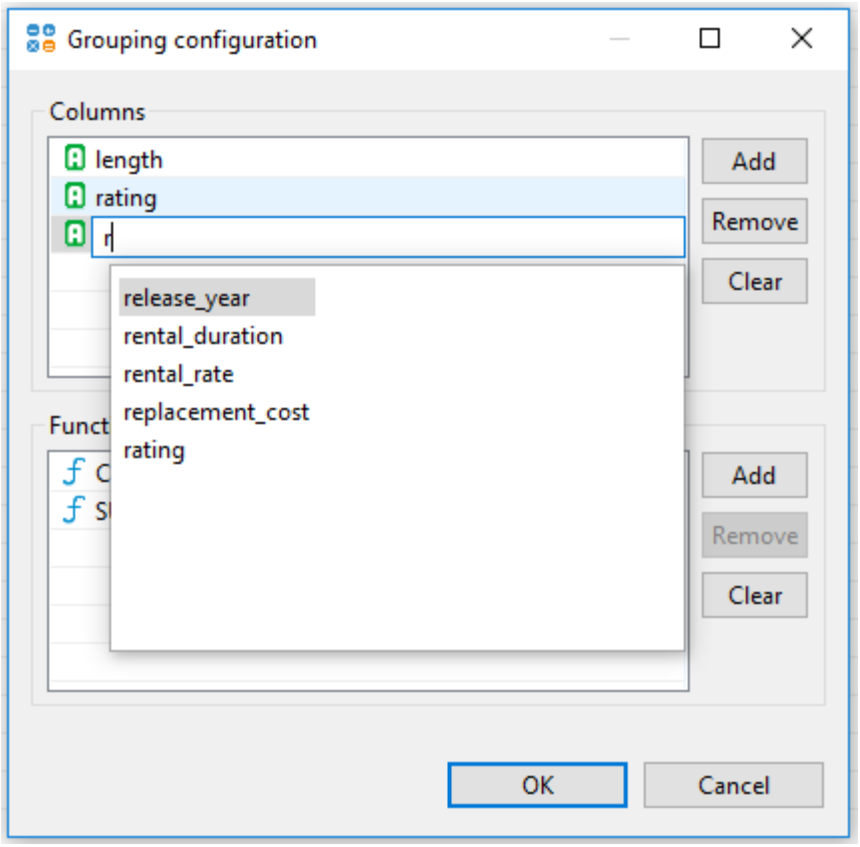

For MySQL/MariaDB databases you can also add a column with an expression - the expression will be calculated in the resulting column:

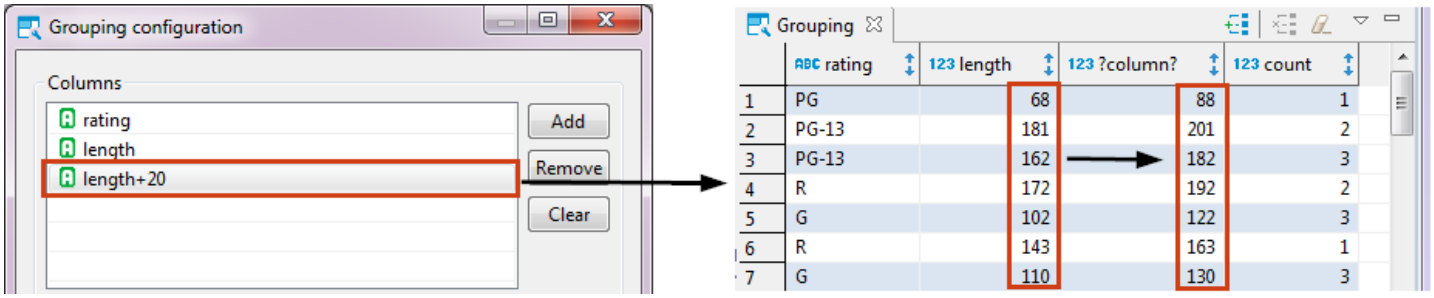

To remove a column, in the Grouping Configuration window, in the **Columns** area, click the column name, then **Remove** and **OK**. To remove all columns, click **Clear** and **OK**.

Another way to remove a column is to click the column and then the **Remove grouping column** button ( ) in the panel`s toolbar. Clicking the **Clear grouping** button ( ) removes all results from the Grouping panel.

Managing Charts

Note: This functionality is available only in Enterprise Edition.

The default grid view of query resulting data is not very much impressive, especially to business analysts and other end users. The **Charts** feature lets you quickly and easily turn your SELECT queries' output into a colorized bar chart.

Note: Analytical Charts present only in DBeaver Enterprise Edition

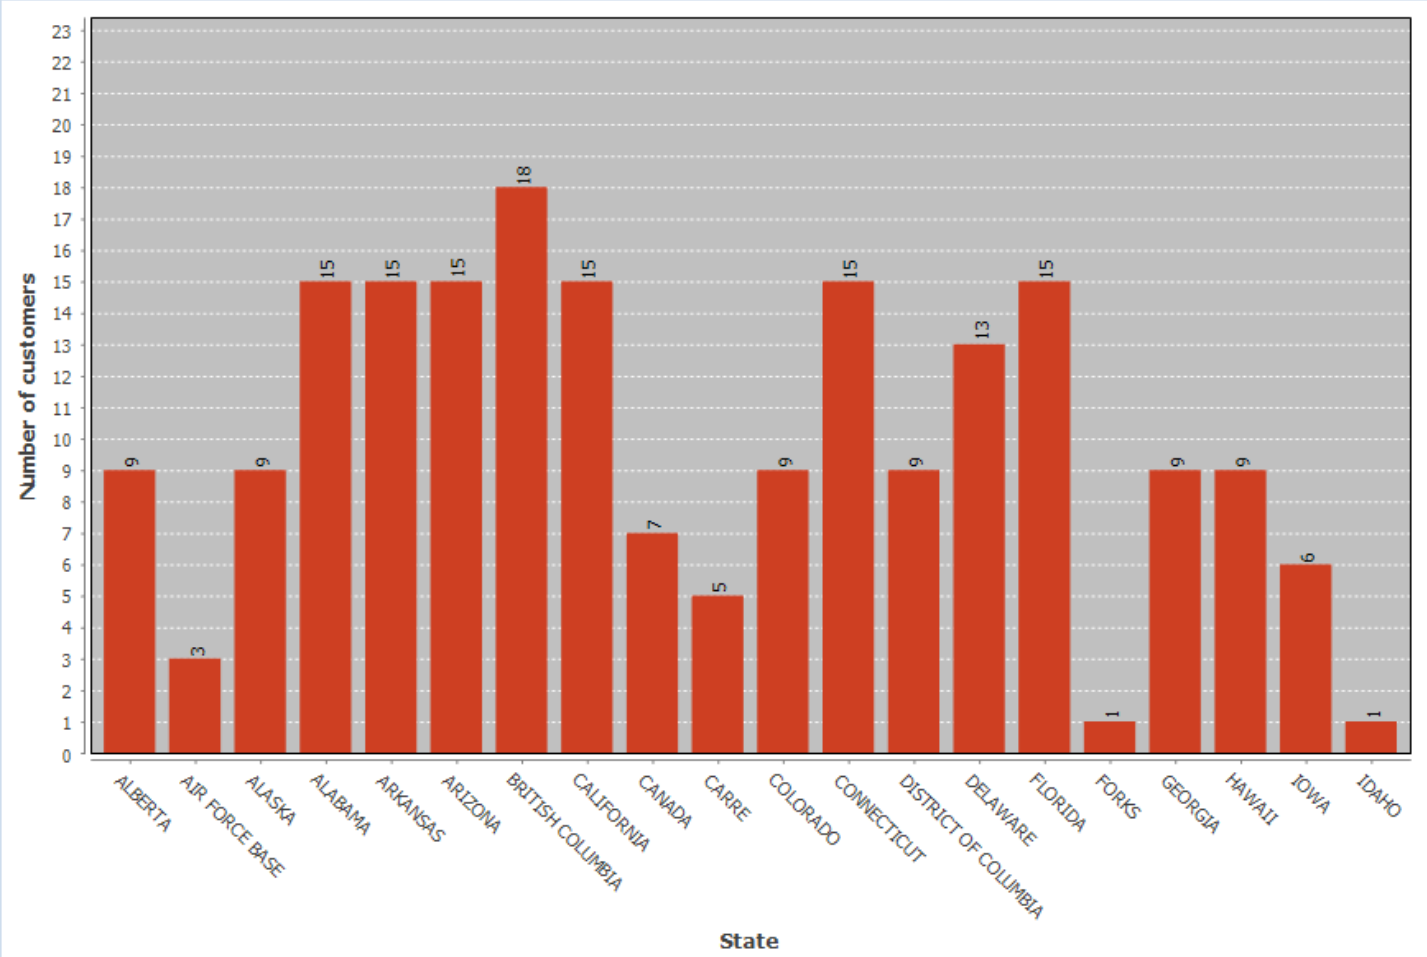

You can easily visualize your data by creating a chart bar both in **SQL Editor** and **Data Editor**.

Creating Charts In SQL Editor

Visual representation of vast data permits the analytical reasoning process to become faster and more focused. Charts make it easy for analysts to perceive salient aspects of their data quickly.

To build a bar chart in the **SQL Editor**, press the **Charts** button in the left vertical toolbar of the query results area.

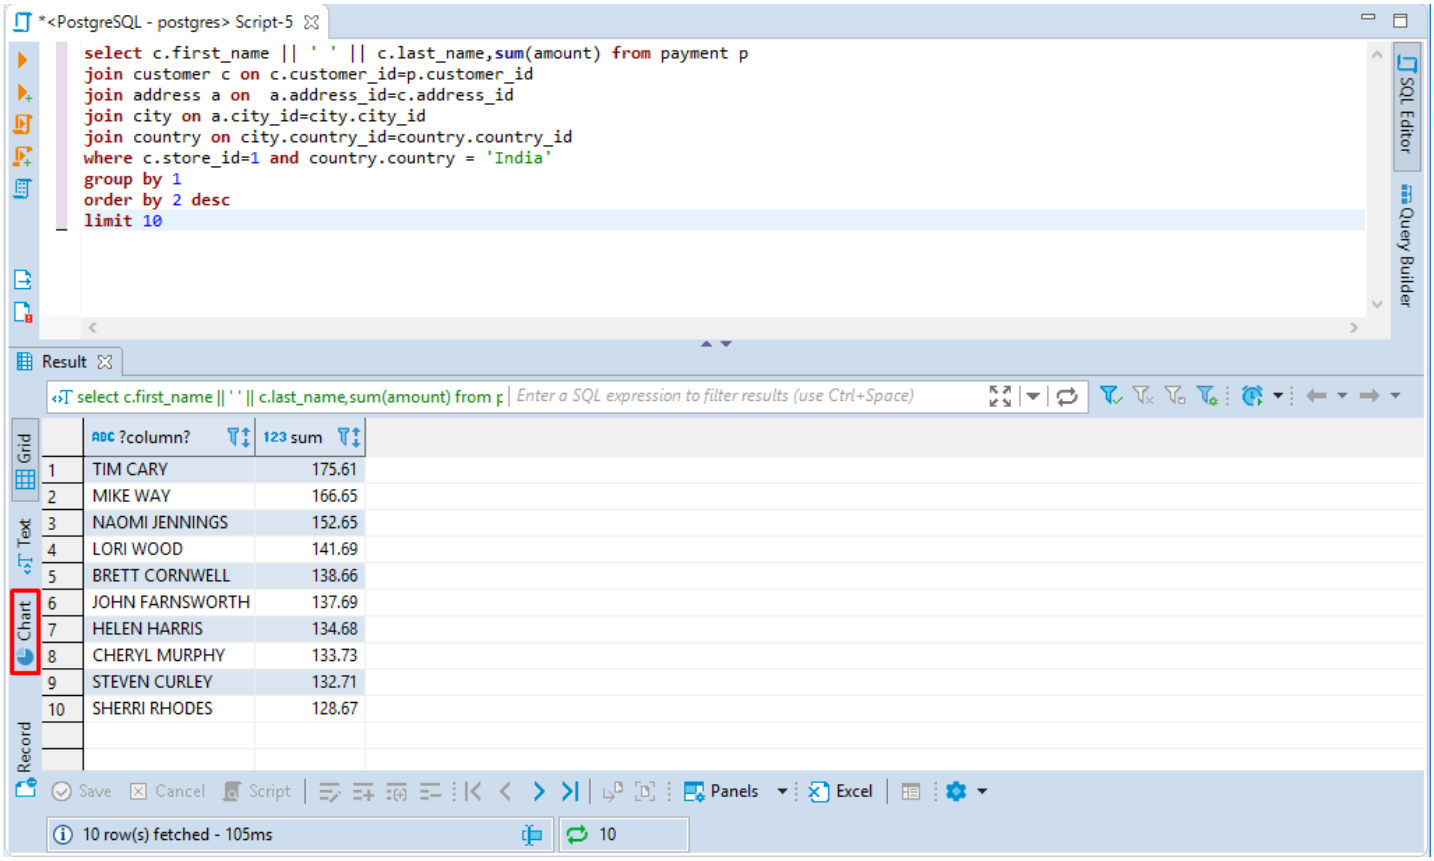

A bar chart will be created.

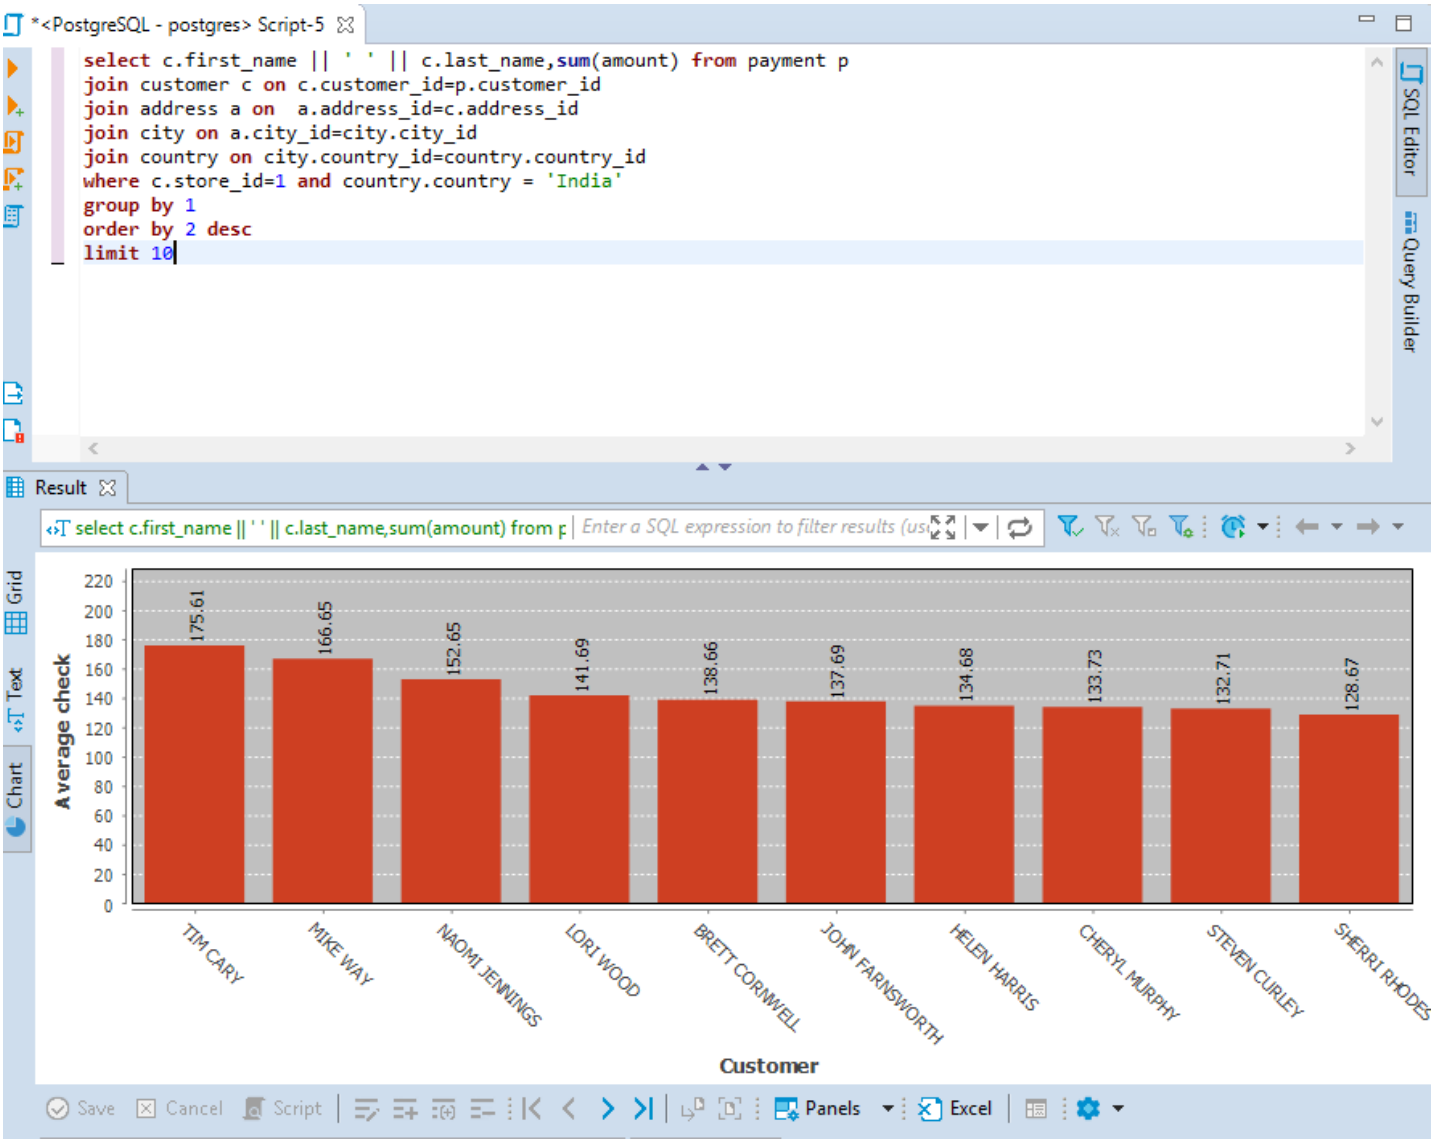

Creating Charts In Data Editor

Charts can be very helpful for visualizing structured analytical data stored as **Views**, for example.

You can also create a chart for any table but you have to structure its data by sorting and applying various filters to its columns first. All the structural changes you make will then affect the chart you build, this way you can adjust the chart representation to the desired one.

To build a chart in the **Data Editor**, press the **Charts** button in the left vertical toolbar.

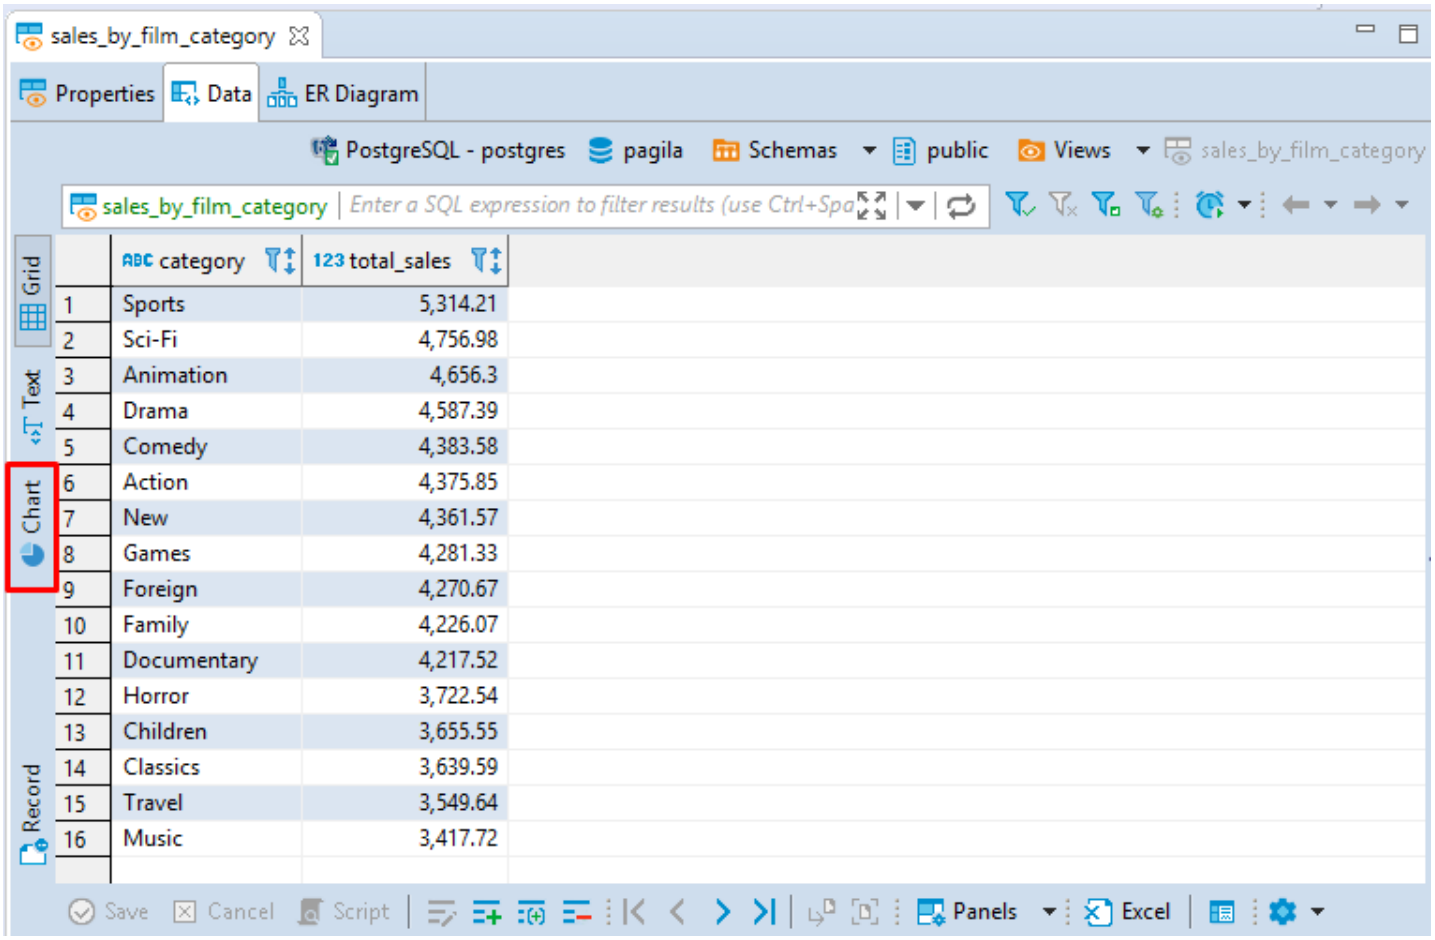

A bar chart will be created.

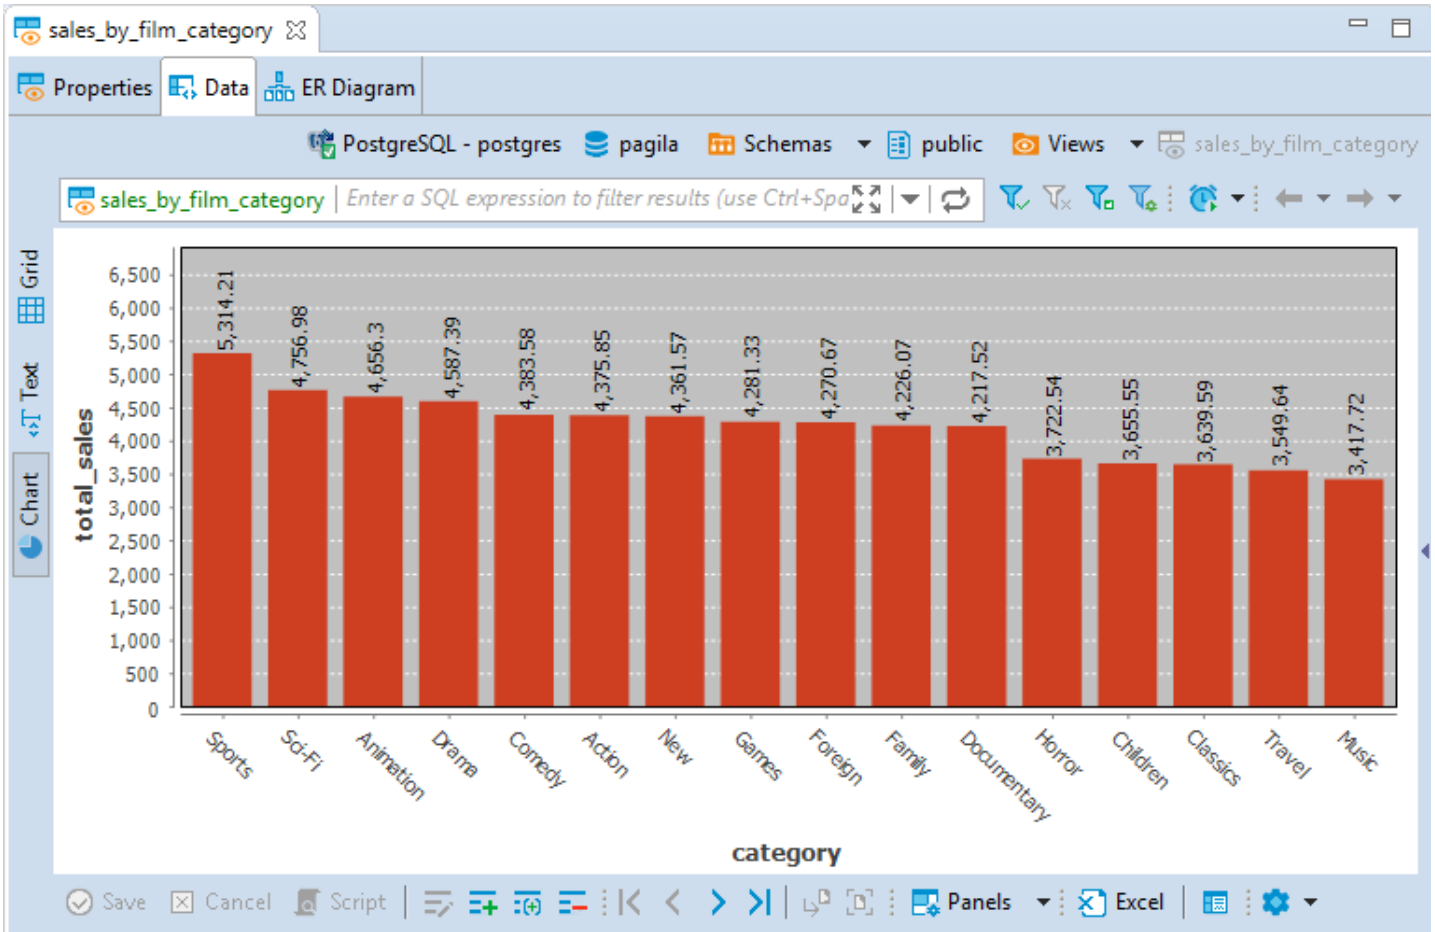

Note, that by default, the data for axis Y is taken from the first column of the table containing numeric values.

Creating Charts In Grouping Panel

Such analytical tool as the Grouping Panel also supports the **Charts** feature. In a chart built for a table containing the grouping results for one or more columns of a data table, you can easily change axes X and axes Y source data by switching the columns in the **Charts Editor**.

To build a chart in **Grouping Panel**, press the **Charts** button in the left vertical toolbar.

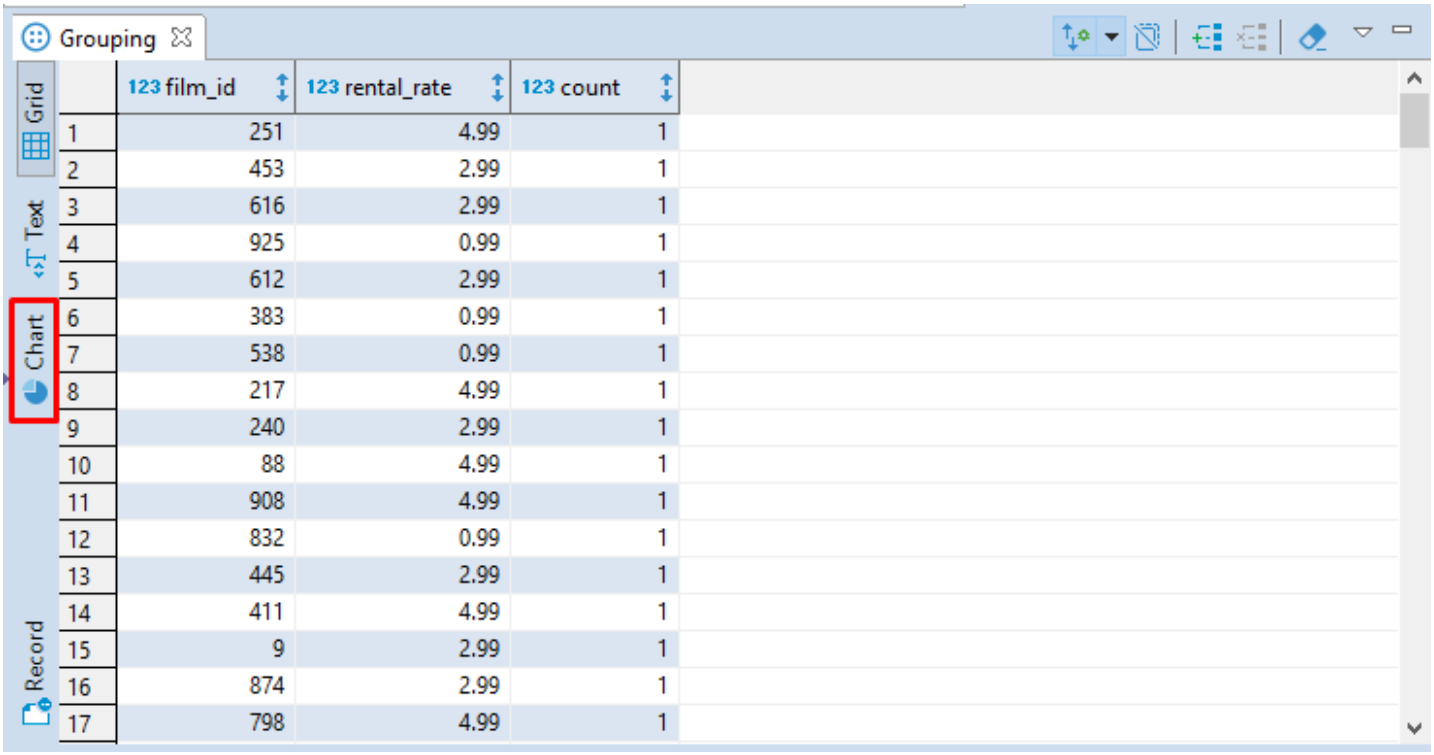

Editing Chart Settings

To edit chart settings select the **Charts...** option in the chart's context menu and the **Chart Editor** will appear.

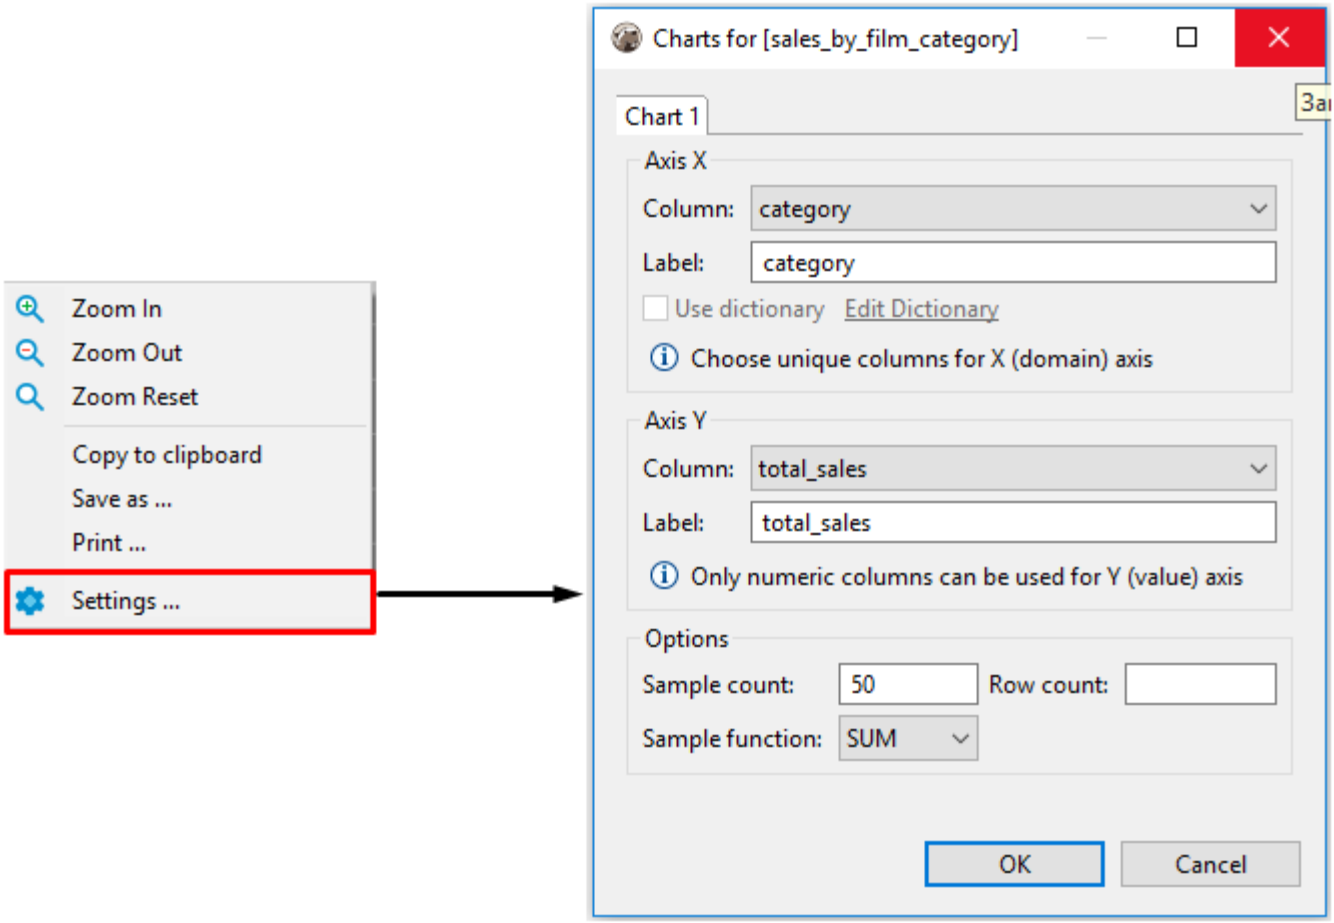

The following chart settings can be adjusted:

Setting Axis X

1. In the **Column** drop down list of available columns select a column whose data will be used on axis X of the bar chart. Make sure you choose unique columns for X axis.
2. Define a user-friendly axis name in the **Label** text field.

Setting Axis Y

1. In the **Column** drop down list of available columns select a column whose data will be used on axis Y of the bar chart. **Note**, that only columns containing numeric data can be used for axis Y.
2. Define a user-friendly axis name in the **Label** text field.

Setting Other Options

You can also set the following chart options:

- Sample count - maximum number of columns used for building a chart;
- Row count - maximum number of rows used for building a chart;

- Sample function - an aggregate function where the values of multiple rows are grouped together to form a single summary value displayed on axis Y.

The following sample functions are suppoerted:

Name	Description
AVG	Average value
SUM	The sum of all values
FIRST	The first value
LAST	The last value
COUNT	Total count of all values

Copying to clipboard

You can copy a chart to clipboard by selecting the **Copy to clipboard** option in the chart's context menu.

Exporting Charts

You can export a chart into PNG format by selecting the **Save as...** option in the chart's context menu.

Printing Charts

You can print a chart by selecting the **Print...** option in the chart's context menu.

Data Search

To search for data in the result set, press `CTRL+F` . The standard Find/Replace search dialog box opens:

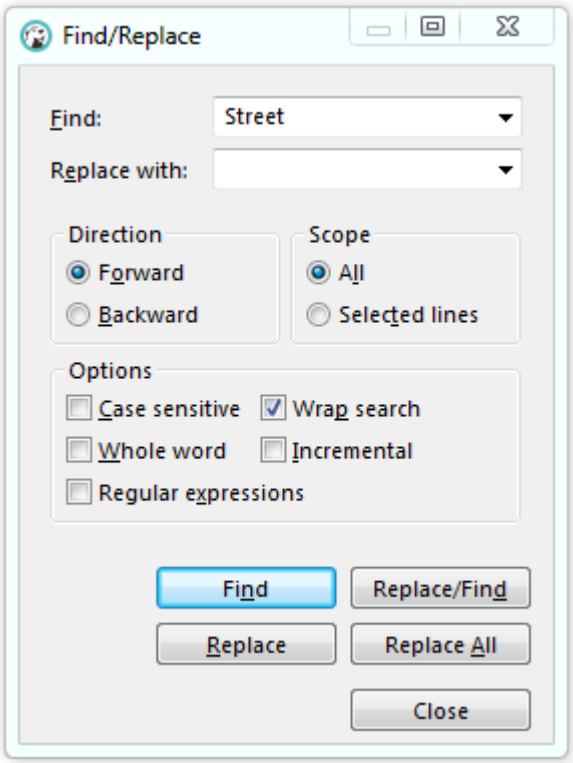

You can also use the Find and Replace feature.

NOTE: The system searches only in already fetched rows.

SQL Generation

You can generate SQL statements (SELECT/INSERT/UPDATE/DELETE) based on selected rows. To generate SQL, right-click the selected rows and click **Generate SQL** and then one of the SQL commands on the context menu:

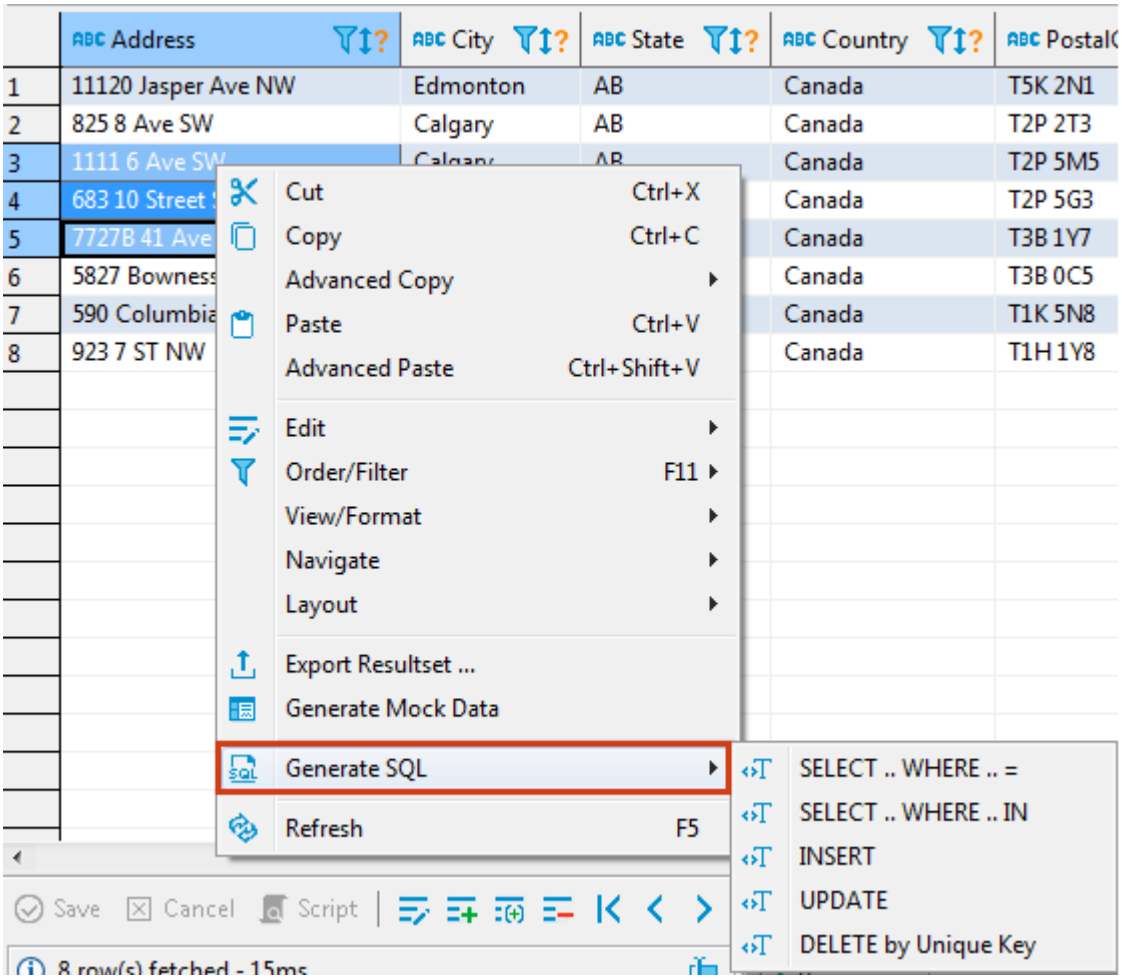

The SQL result opens in a separate window where you can view and copy it:

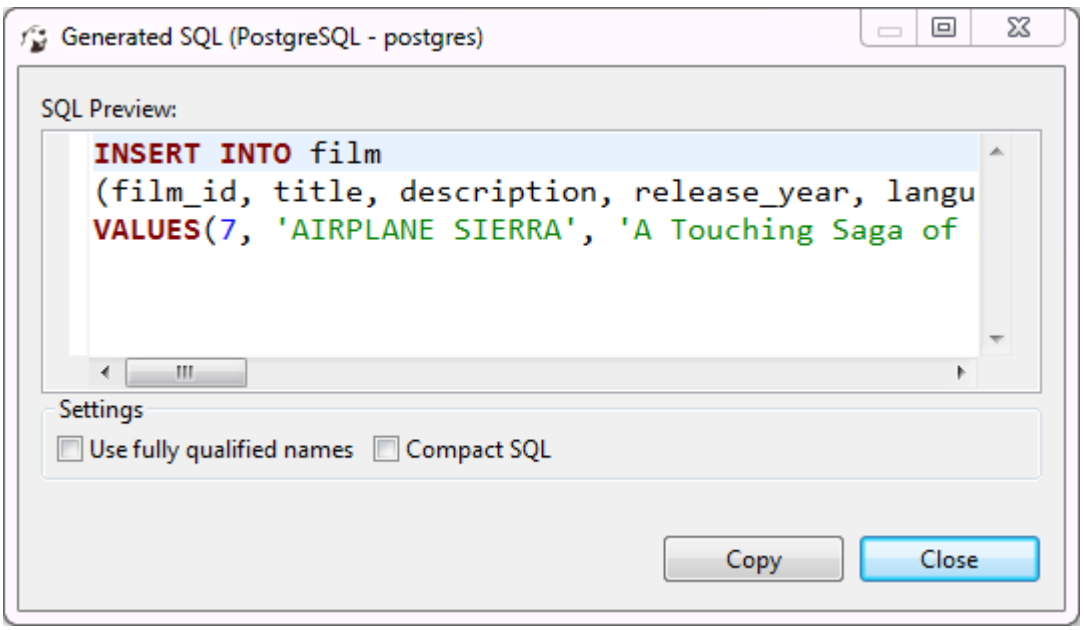

To use table names in the format '[schema name].[table name]', select the **Use fully qualified names** checkbox. To wrap the SQL query into one line, select the **Compact SQL** checkbox:

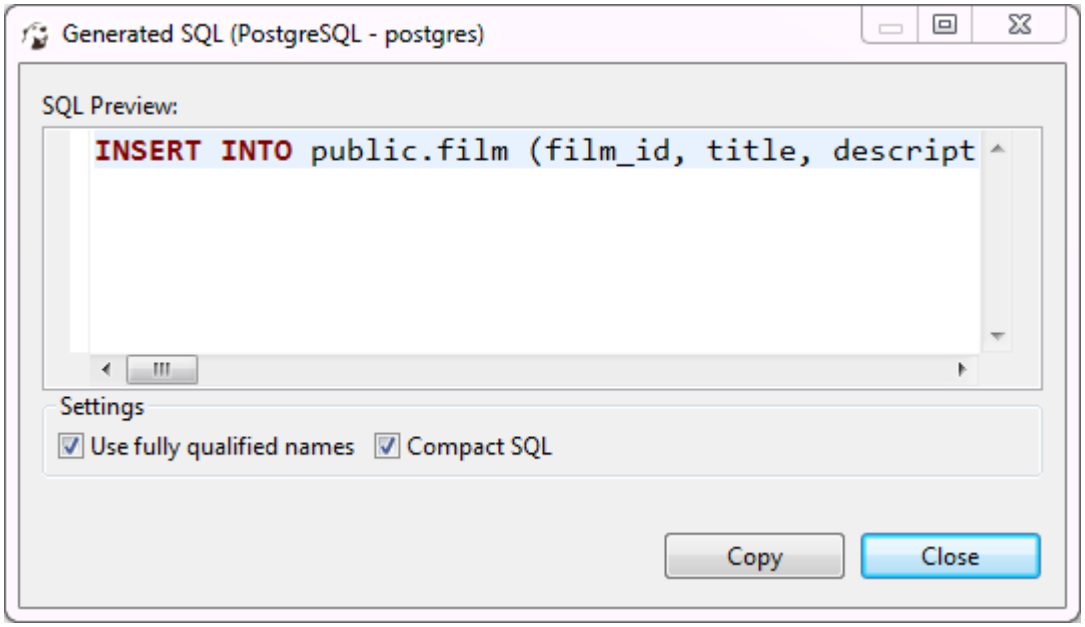

Working with spatial/GIS data

Spatial data is a geometry or geography value that can be represented on a map or as a graph. Geometry object consists of series of points. More details.

DBeaver's support of spatial data covers the following databases:

- PostgreSQL (PostGIS)
- MySQL
- SQLite (GeoPackage)
- H2GIS
- SAP HANA
- Oracle (EE)
- SQL Server (EE)

Spatial data viewer

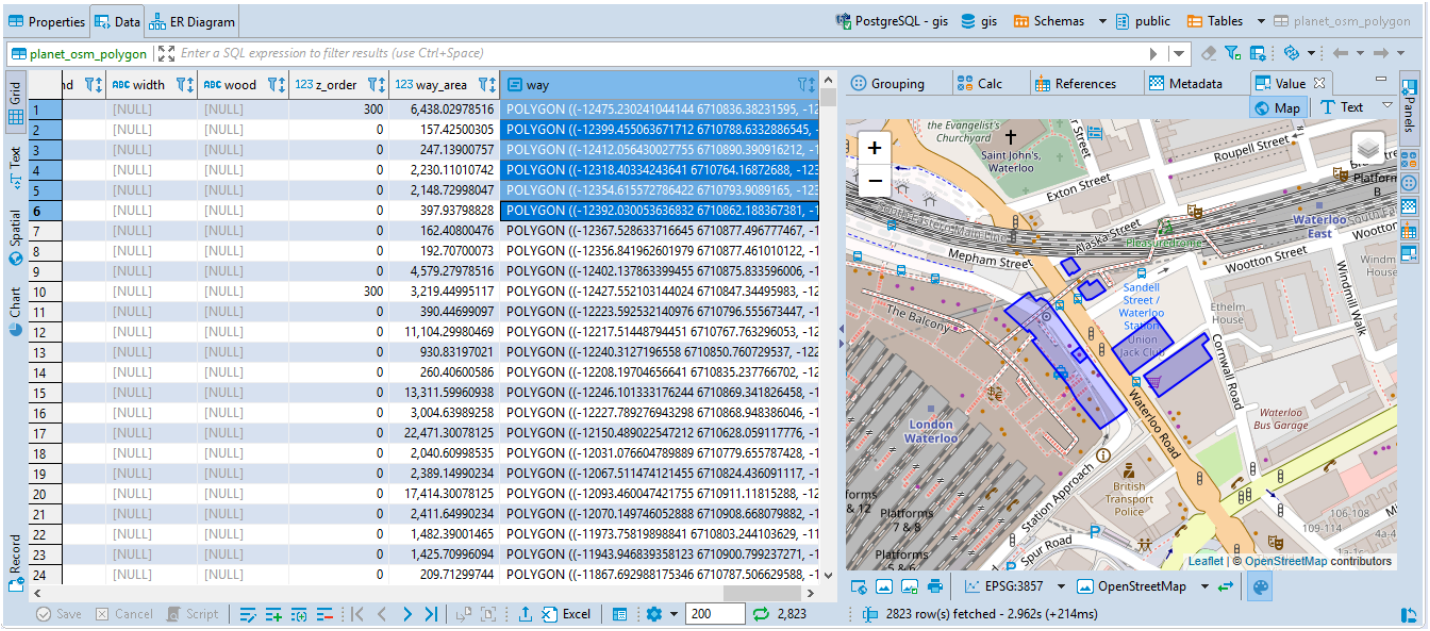

Tile layer management

DBeaver ships with several predefined map tiles. Tiles can be chosen with the combo below the viewer:

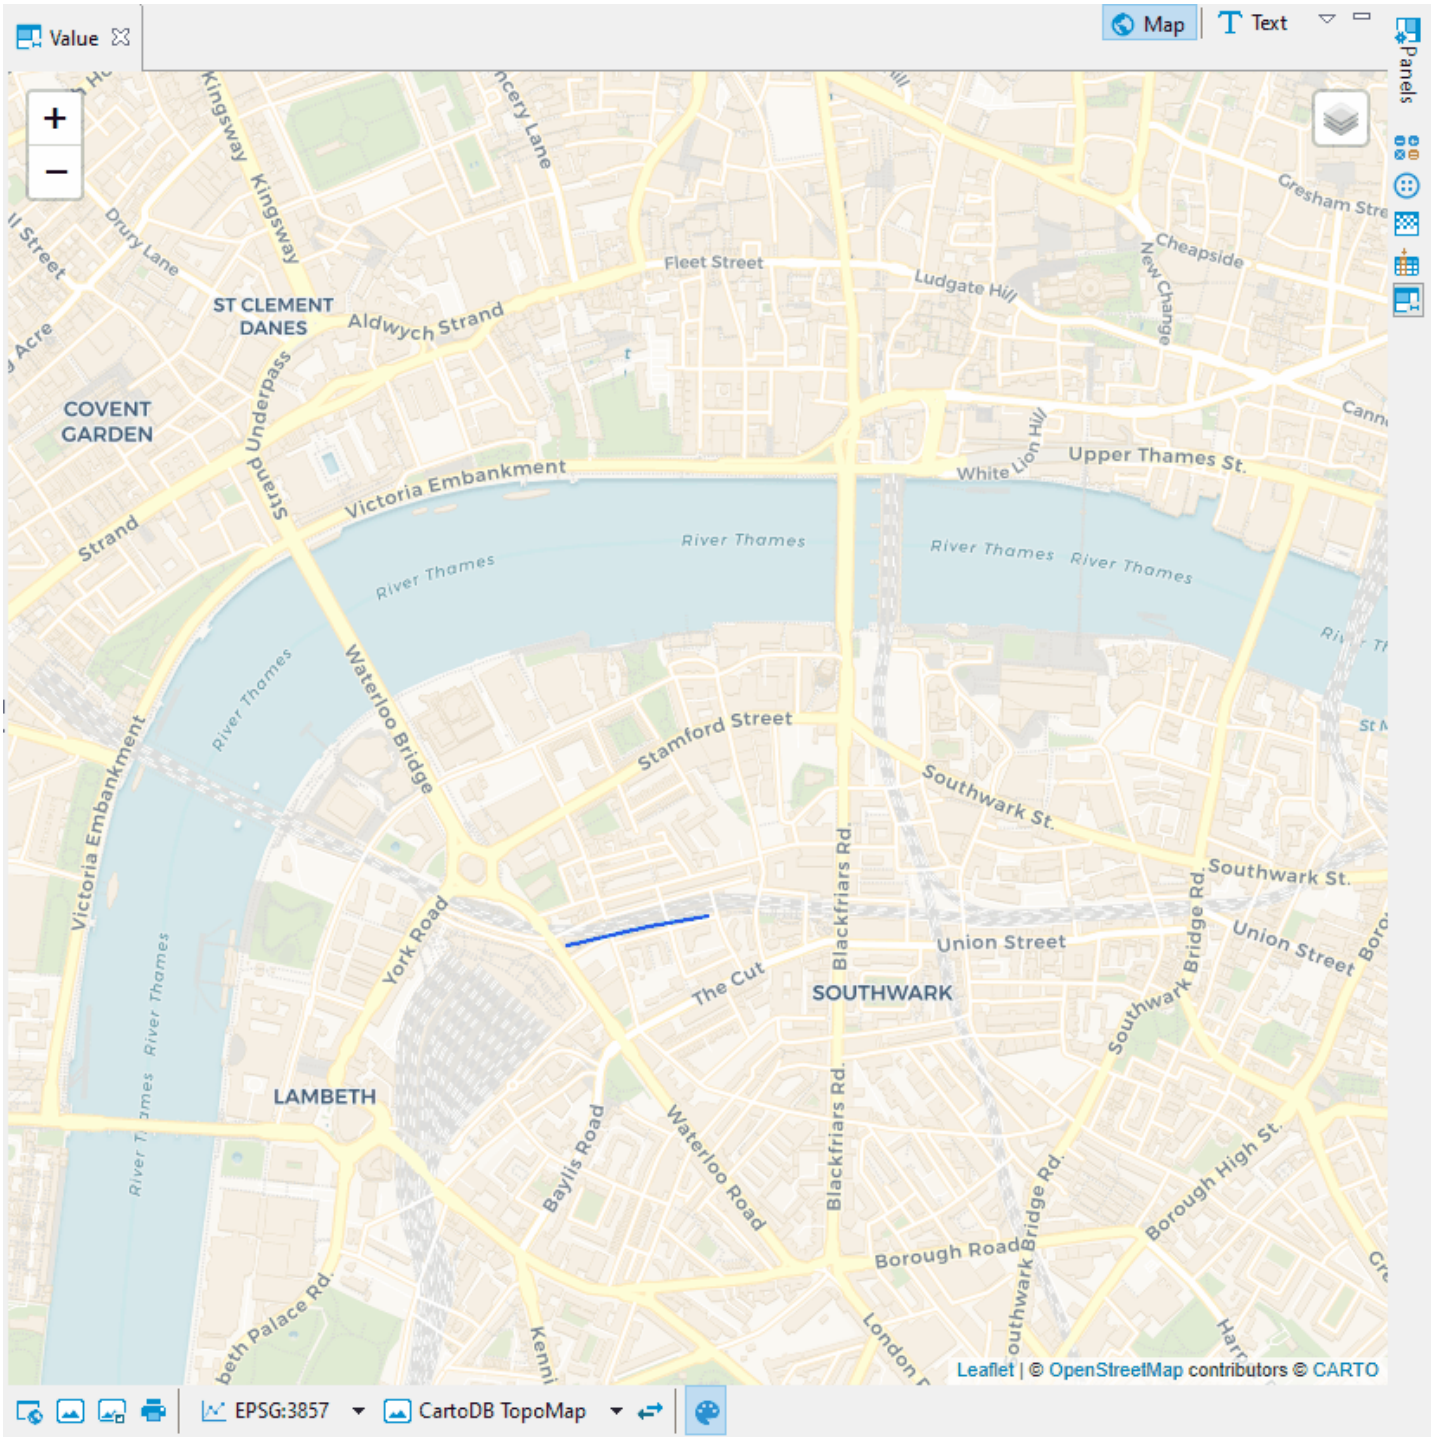

You can choose which tile layers you want to see in the combo in the *manage* dialogue:

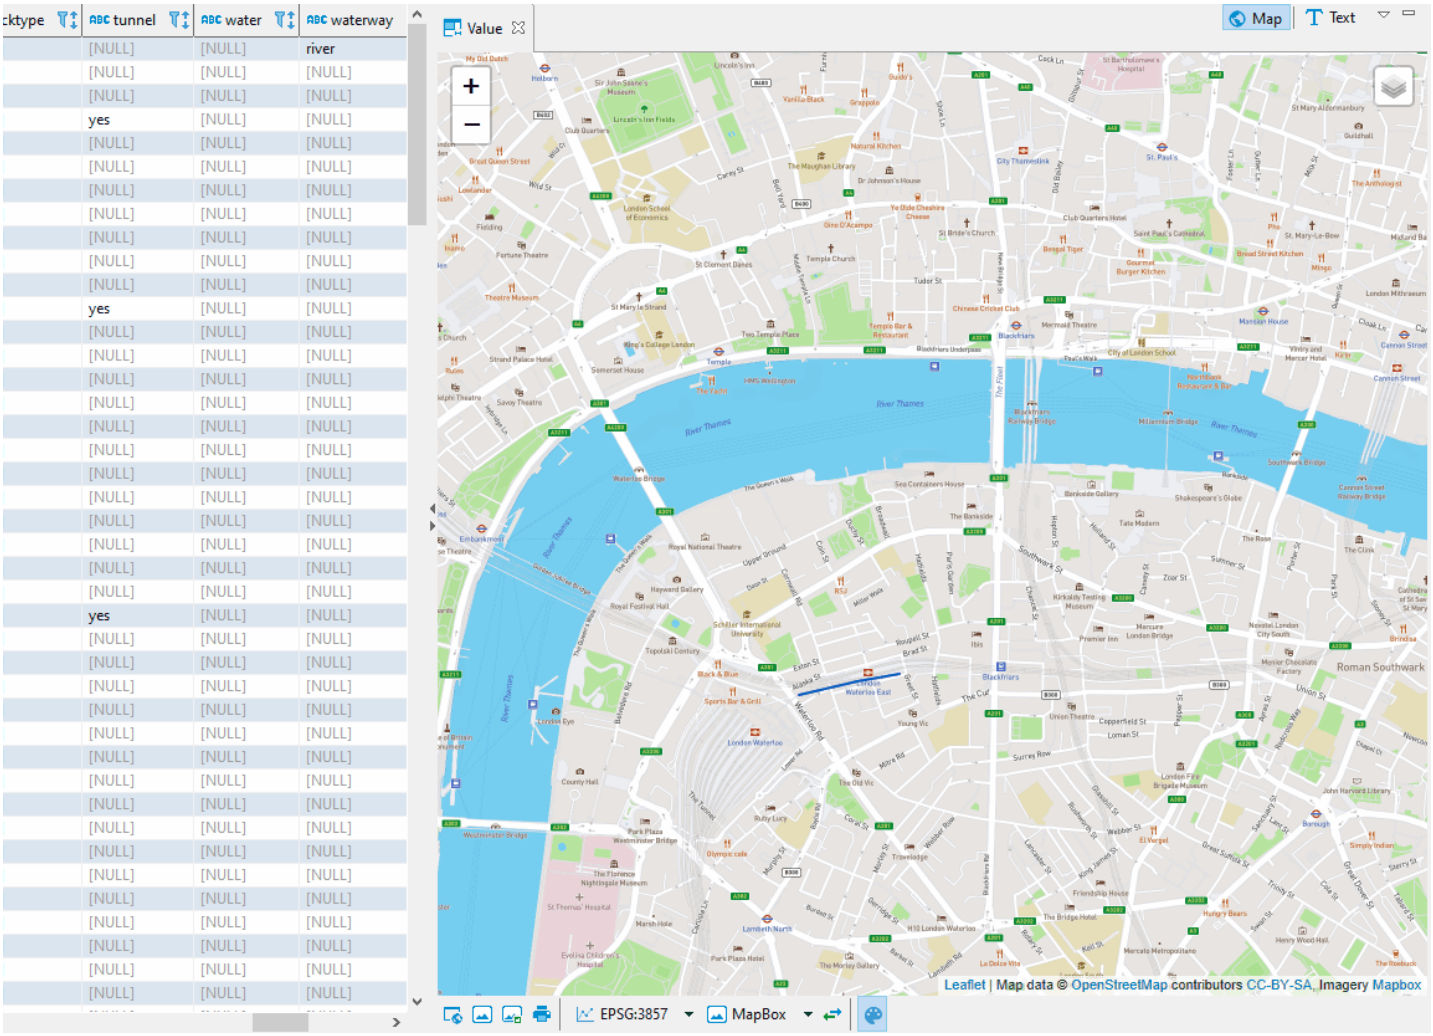

In the same manage dialogue, you can add new tile layers, edit layers you added previously, or delete them:

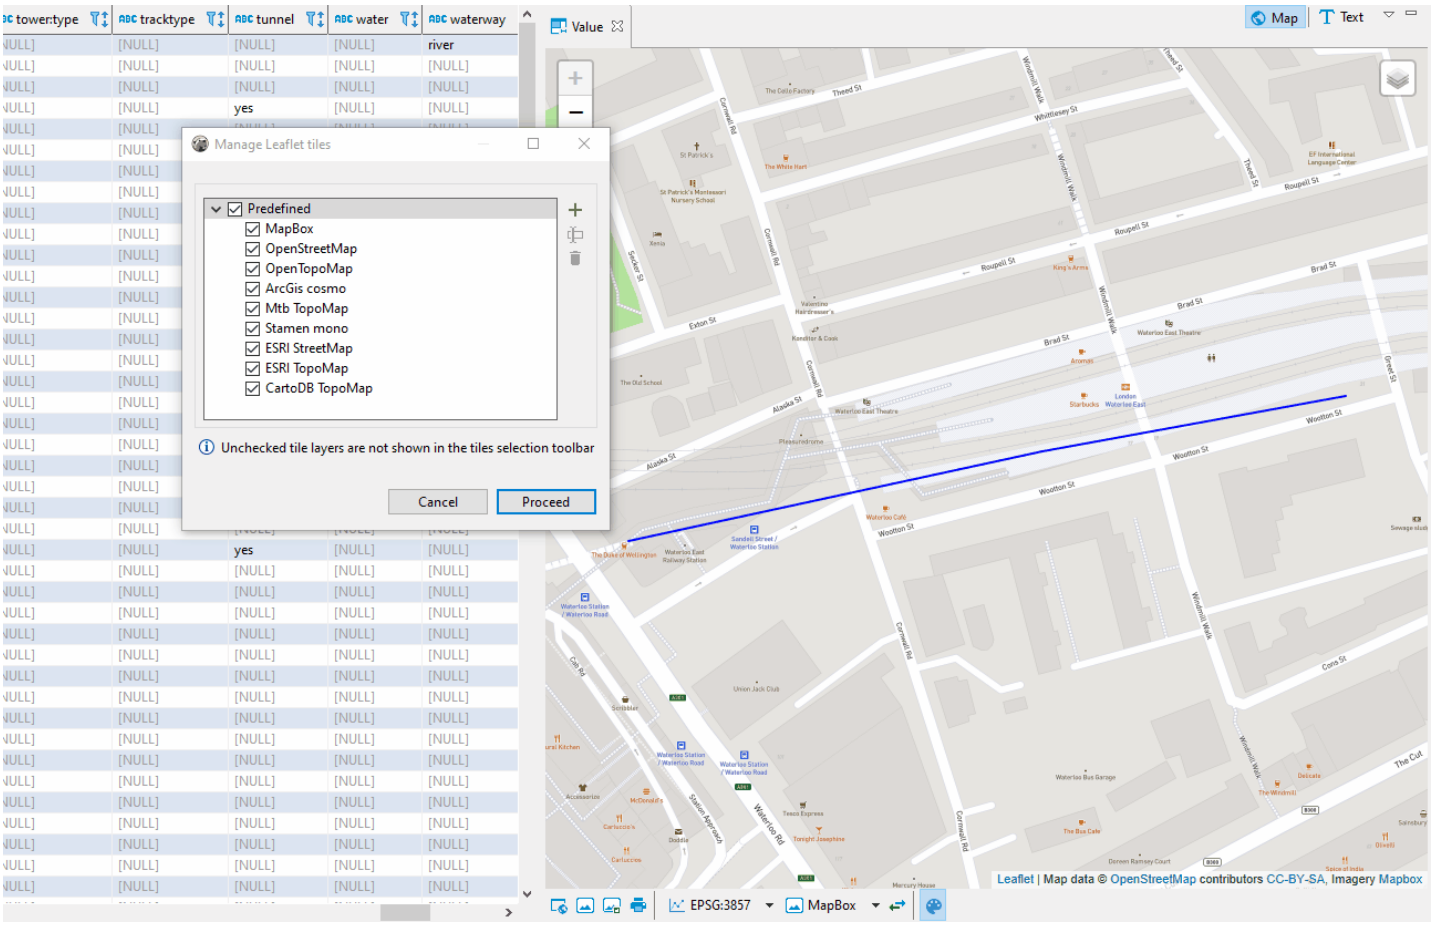

Defining custom tile layer

At this point, you may be wondering what to put in the Layers definition box. Here is a brief explanation.

DBeaver's spatial data viewer uses Leaflet (version 1.4.0 at the moment) under the hood. When providing Layers definition, you type the arguments for function L.tileLayer(), which installs a new tile layer. More on that function in the official Leaflet documentation. You can also see the definition of predefined tiles to help you get started.

Viewing string or binary data from any Database on a map

You can also see your geodata on the map if you select the data cell setting "View/Format", then "Set columnName format" and among the formats - Geometry. This works for both string and binary types of columns.

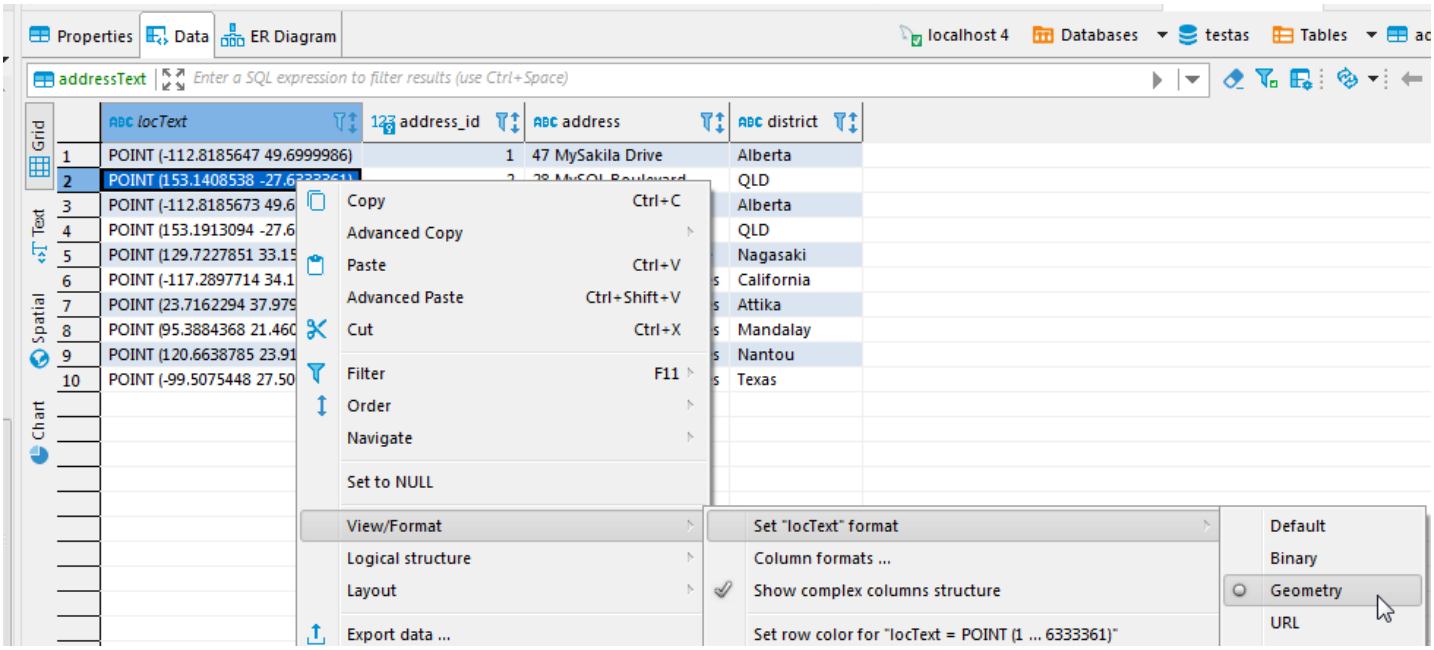

String column type to spatial.

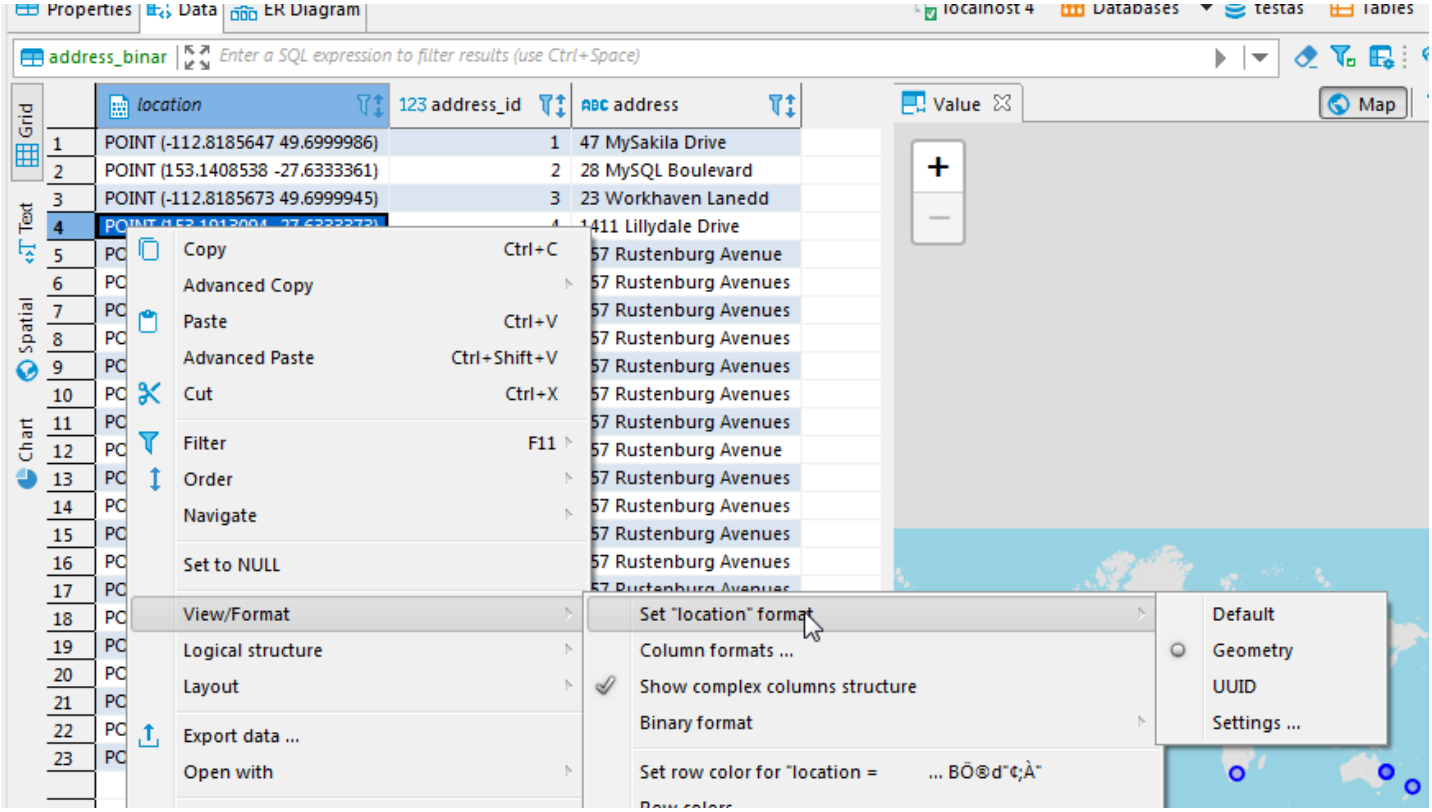

Binary column type to spatial.

Working with XML and JSON

DBeaver supports XML and JSON column types (in relational databases) by using standard JDBC interfaces. This feature was added in JDBC4 so you will need JDBC4 compliant driver for your database.

In the Data Editor, you can edit XML/JSON data right in the table cells, however, a big amount of data may require a larger editor, you may want to save XML/JSON scripts to a local file or upload this type of data from a local file.

To open the full-size XML/JSON editor click the cell containing data in XML/JSON format and press `Shift + Ente`
.

By default the editor opens on JSON tab, open XML tab to modify XML data.

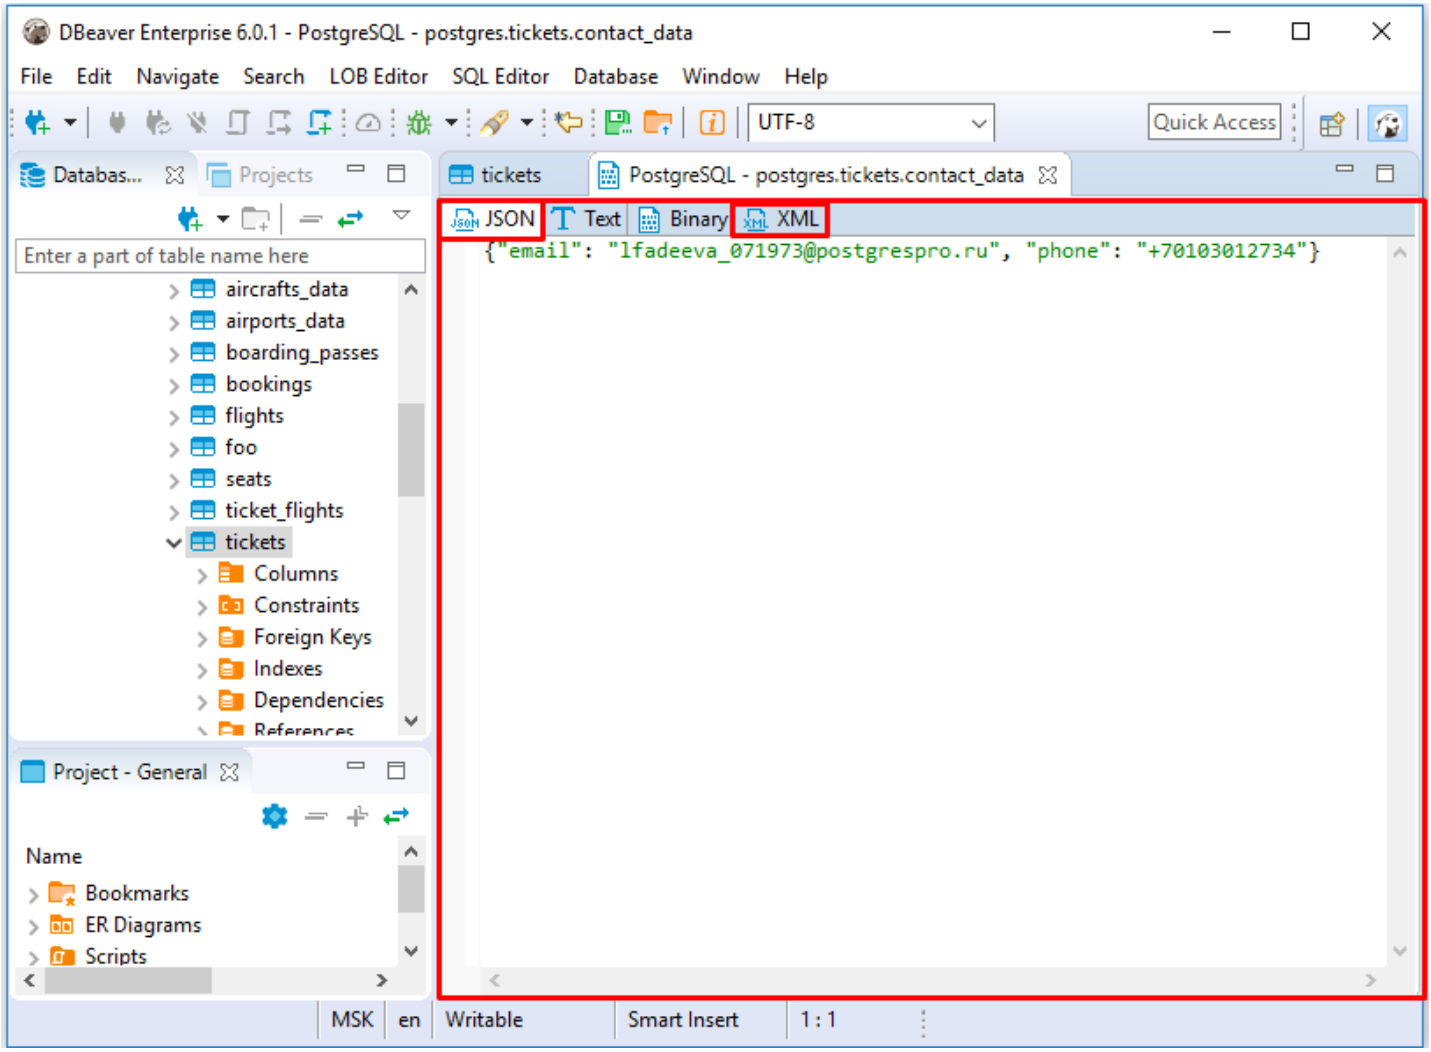

To auto-format XML/JSON script press `Ctrl + Shift + F` keyboard buttons.

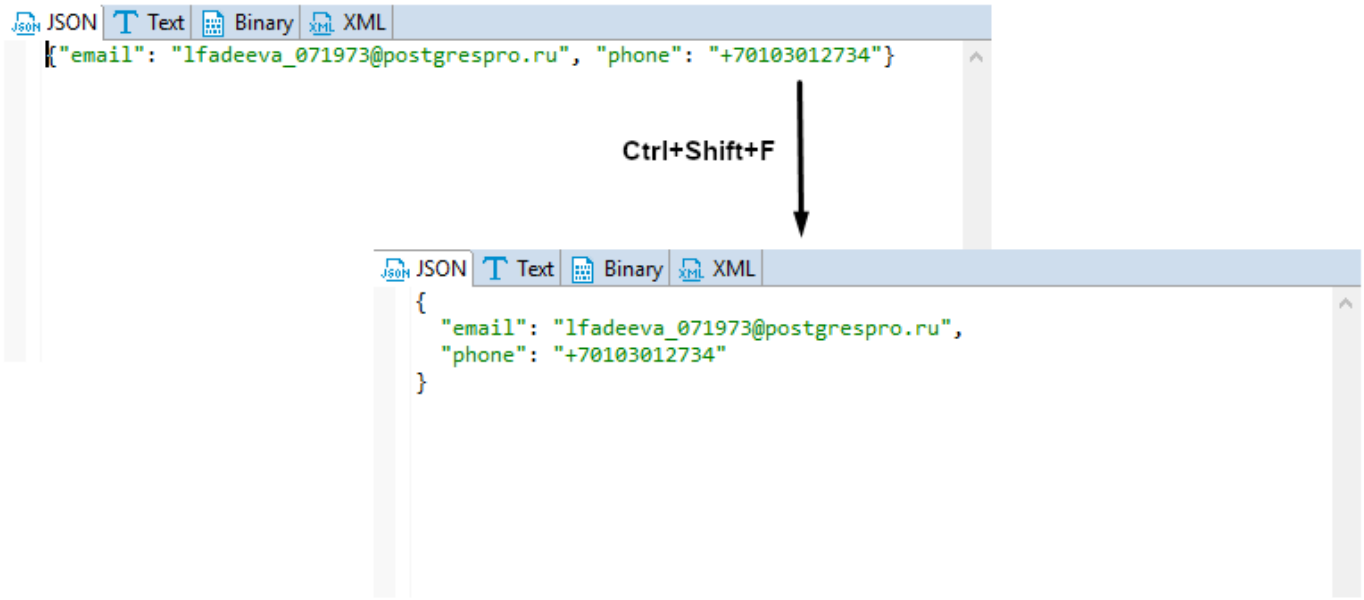

Use `Ctrl + S` keyboard shortcut to save the changes made.

You can also edit XML/JSON content, save it locally and upload it from a local file with the help of **Value** panel toolbar.

To upload data from a local file, press the **Load from file...** button .

To save the content to a local file, press the **Save to file...** button .

To switch between the formats, press the **Content viewer settings** button and select the format.

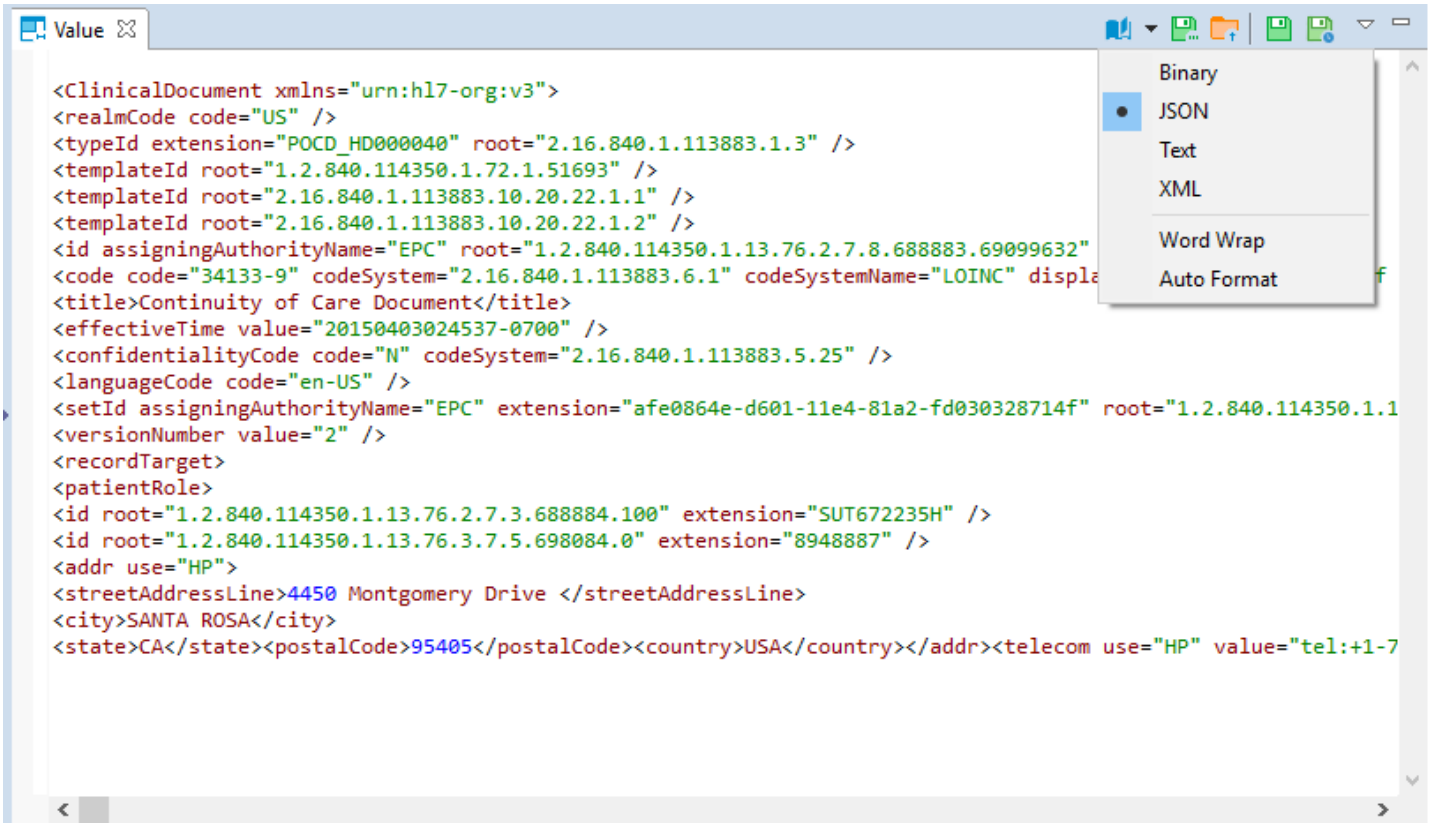

Use **Word Wrap** feature that wraps the text within a screen.

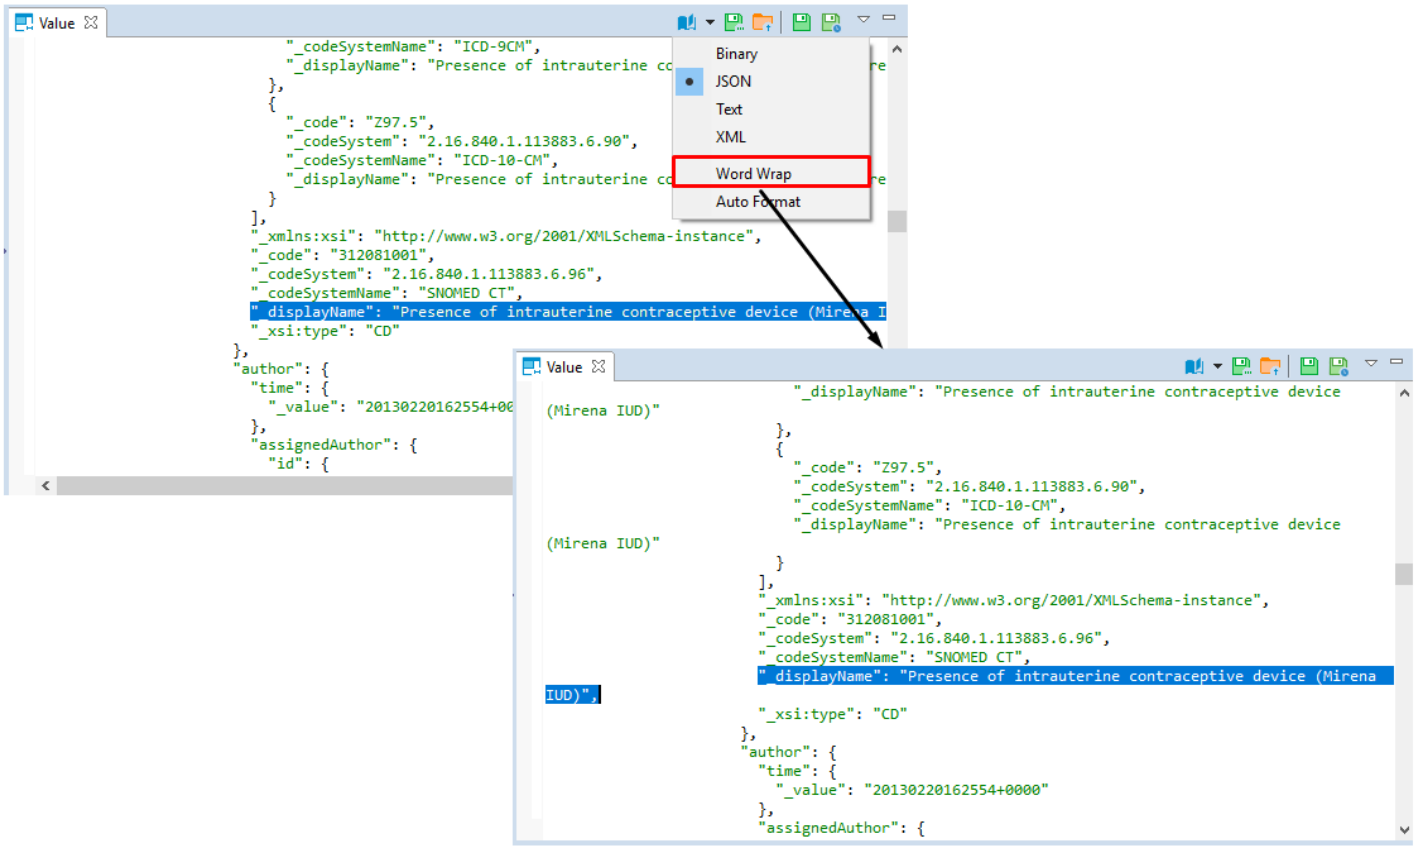

Use **Auto Format** feature to automatically change the appearance of XML/JSON script (fix spaces around operators / commas, fix indentation, etc) and make it more readable.

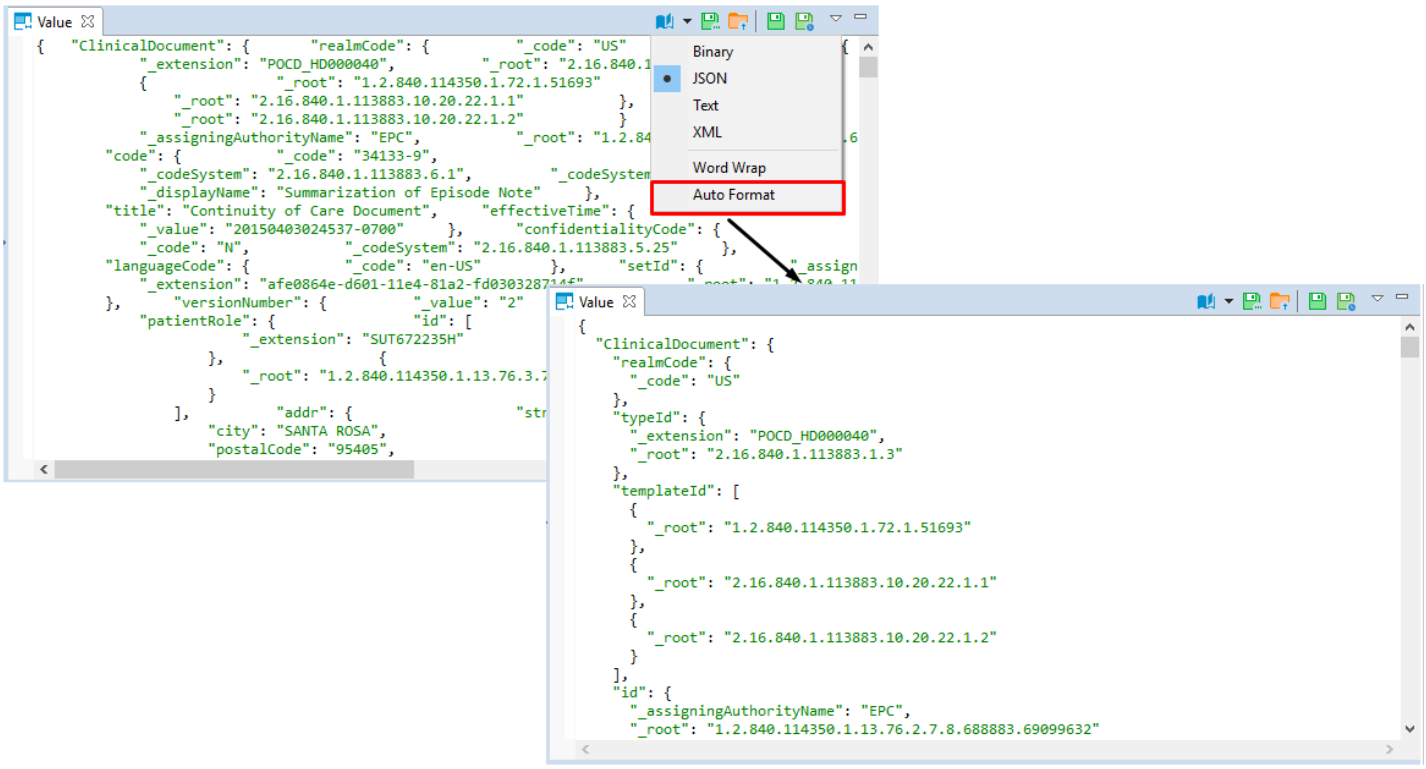

To learn more about **Value** panel, see Panels.

Managing Data Formats

The DBeaver formatting functions allow you to set up database locale and change datasource format settings. This feature can be very useful on database migration, for example.

To change data format settings use option Window -> Preferences in main menu.

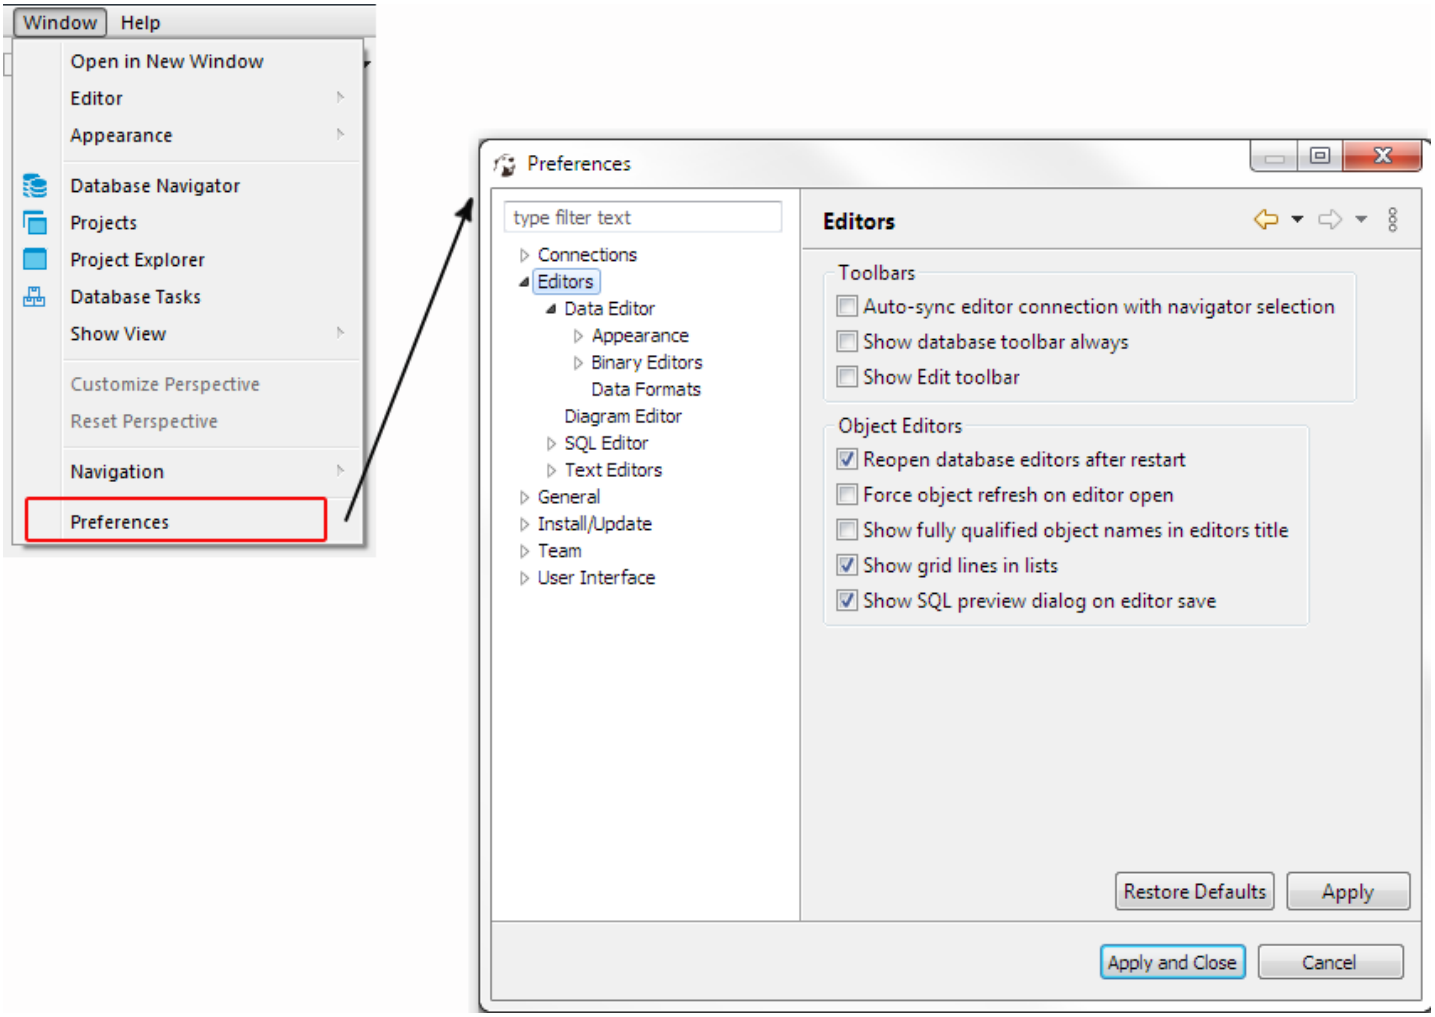

In the Preferences dialog box go to Editors -> Data Editor -> Data Formats.

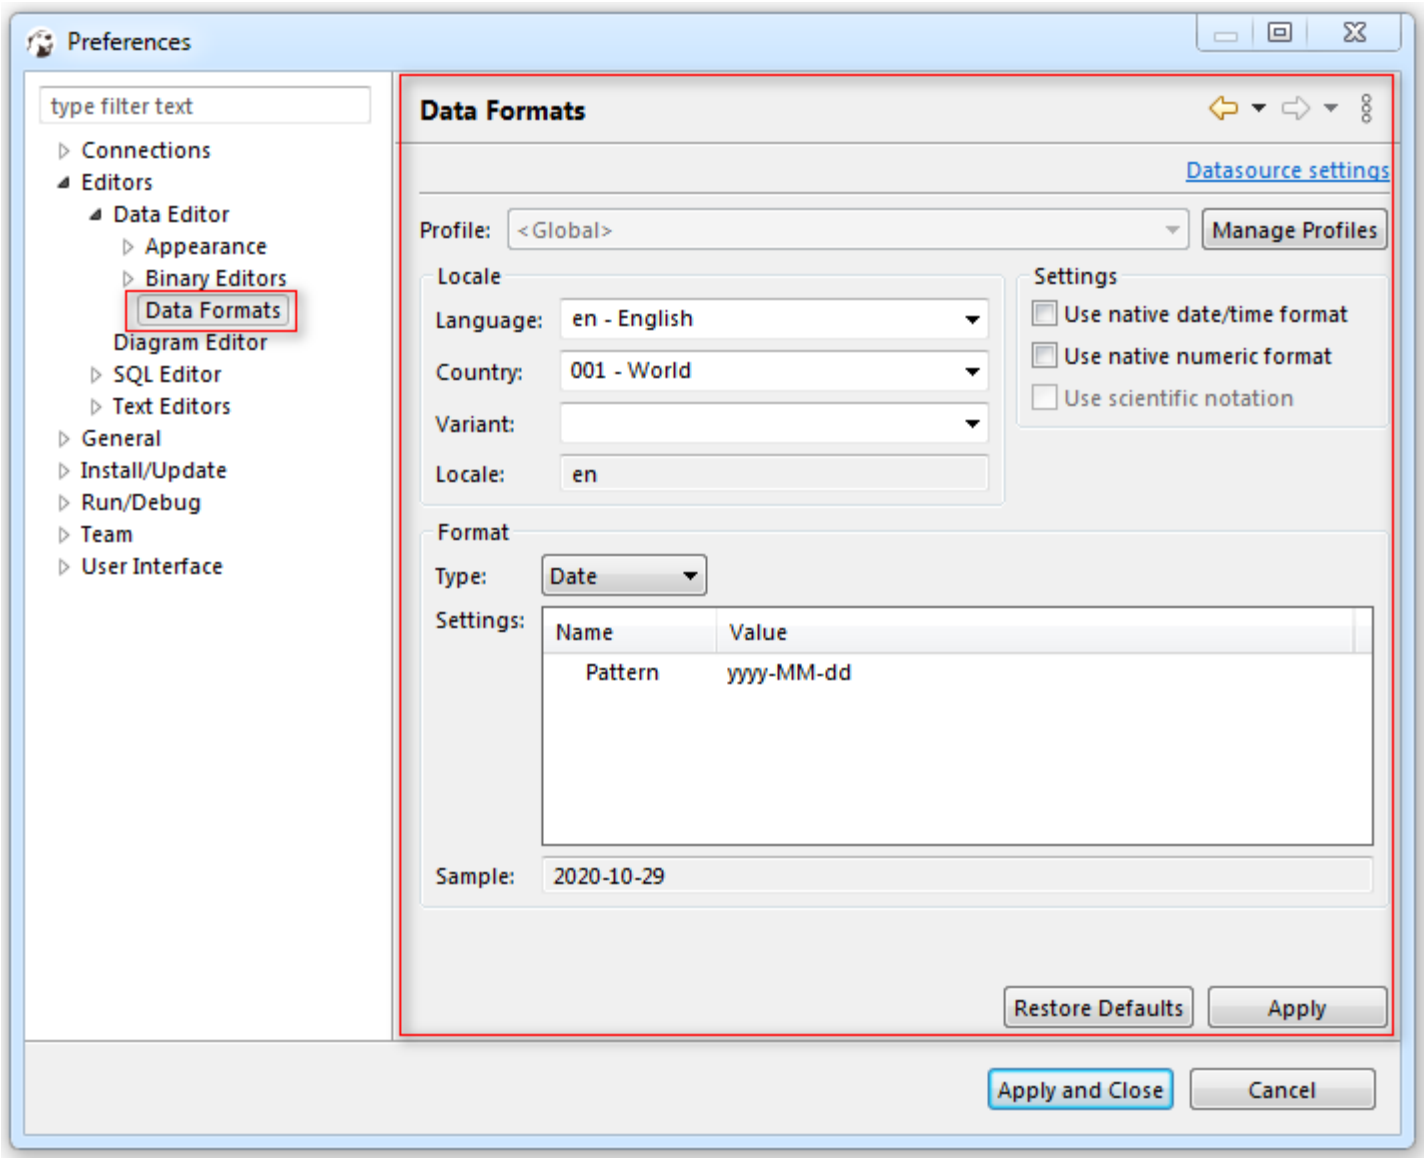

Or, in the Database Navigator right-click a connection and select **Edit Connection** menu option.

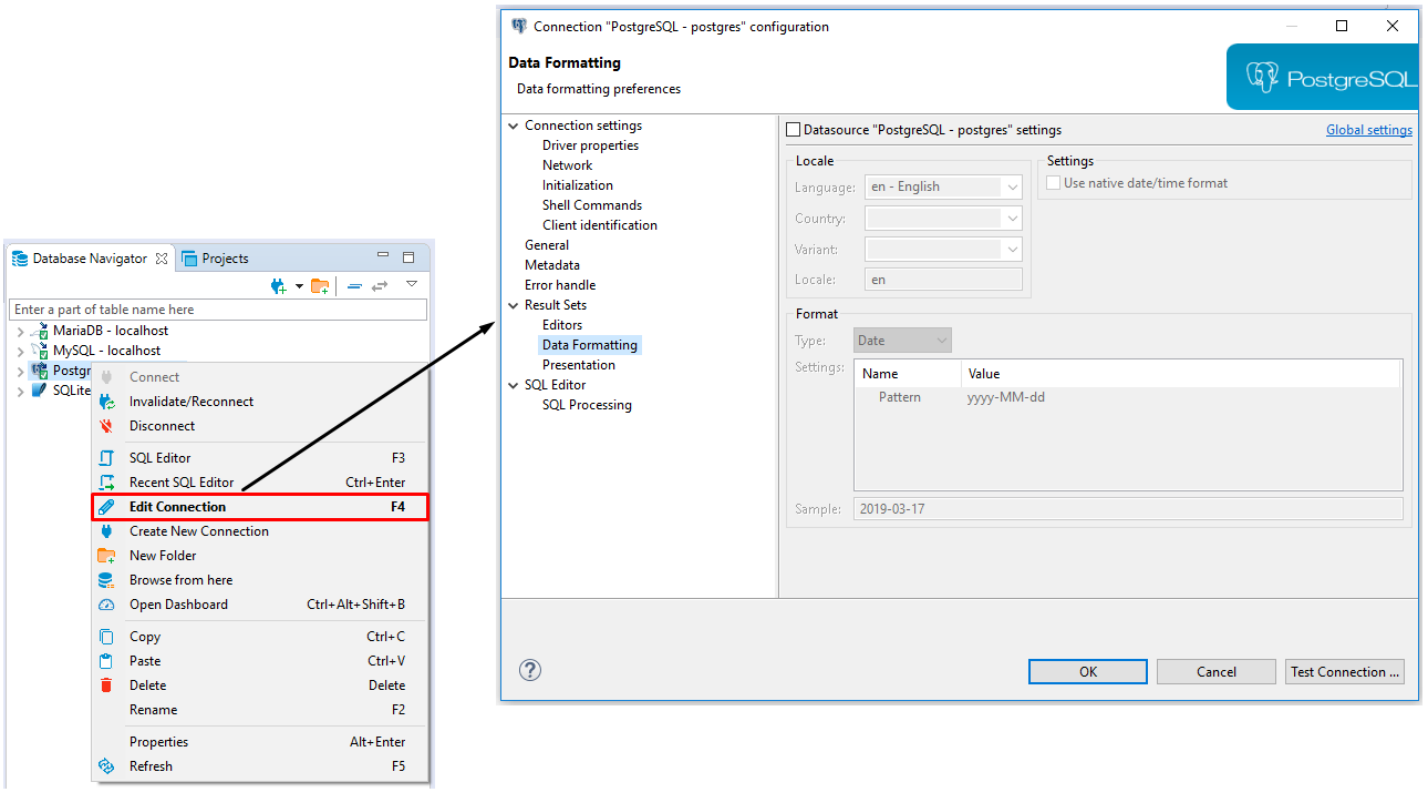

In the right area of the opened **Data formatting preferences** dialog window go to Data editor -> Data Formats and select the **Datasource settings** check box in the left area to customize the data format settings.

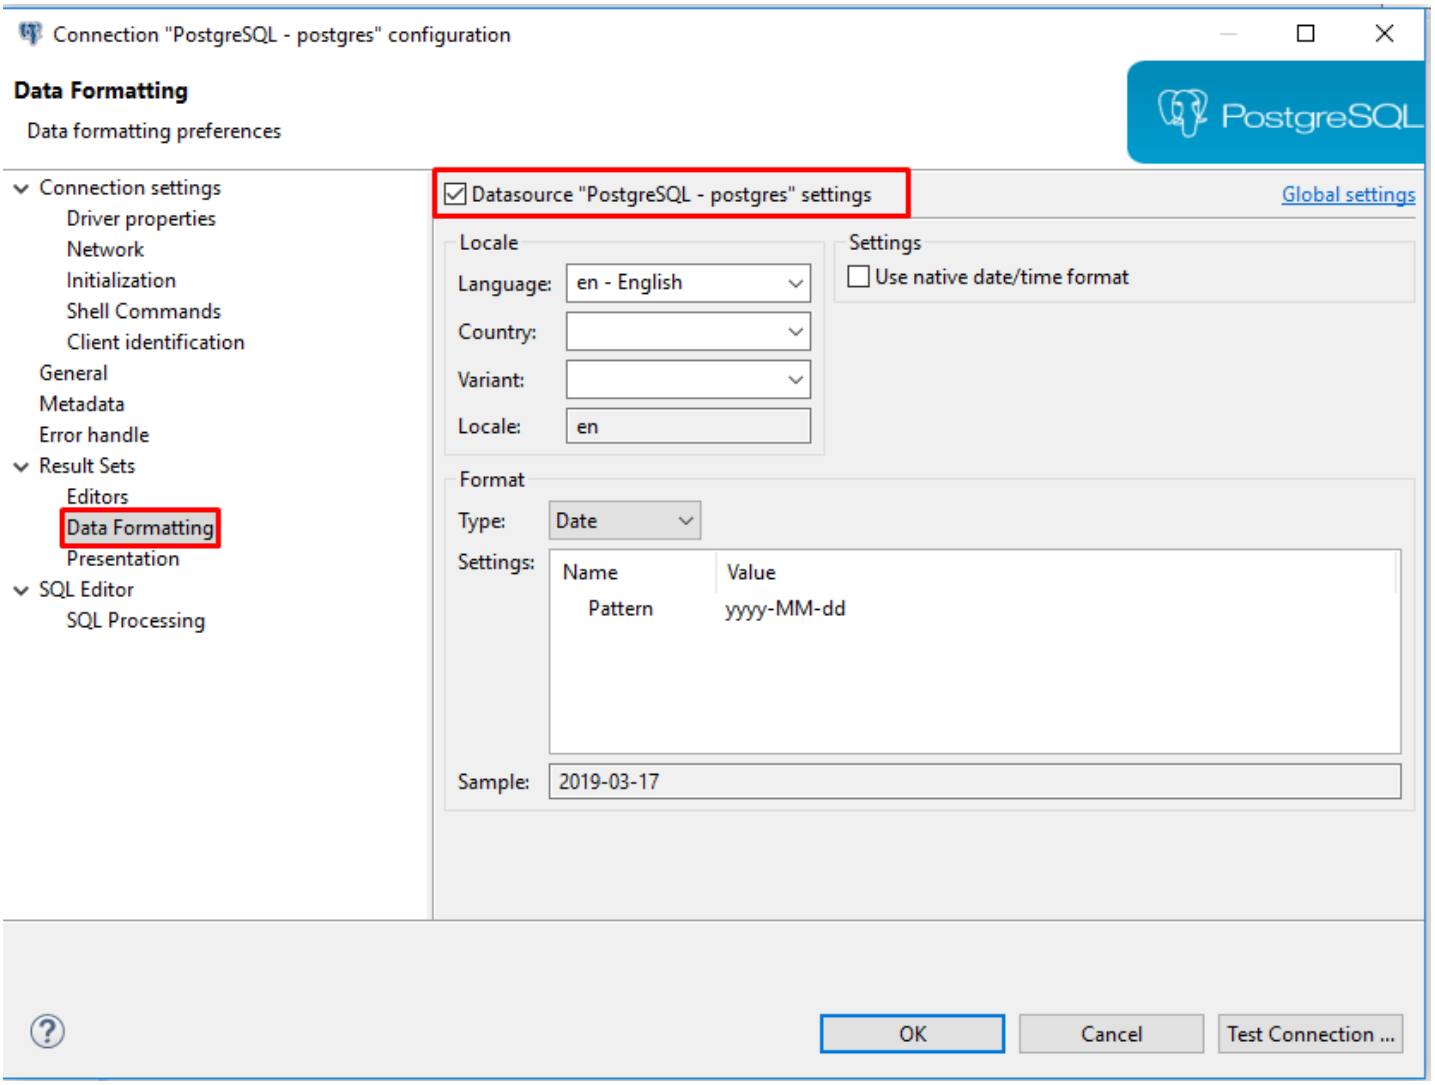

Data Format Profiles

Data format profiles allow you to apply a set of data format preferences to the whole current project by one click.

To create a data format profile press the **Manage Profiles** button. In the opened dialog window press button **New Profile**, define the name and press **Create**.

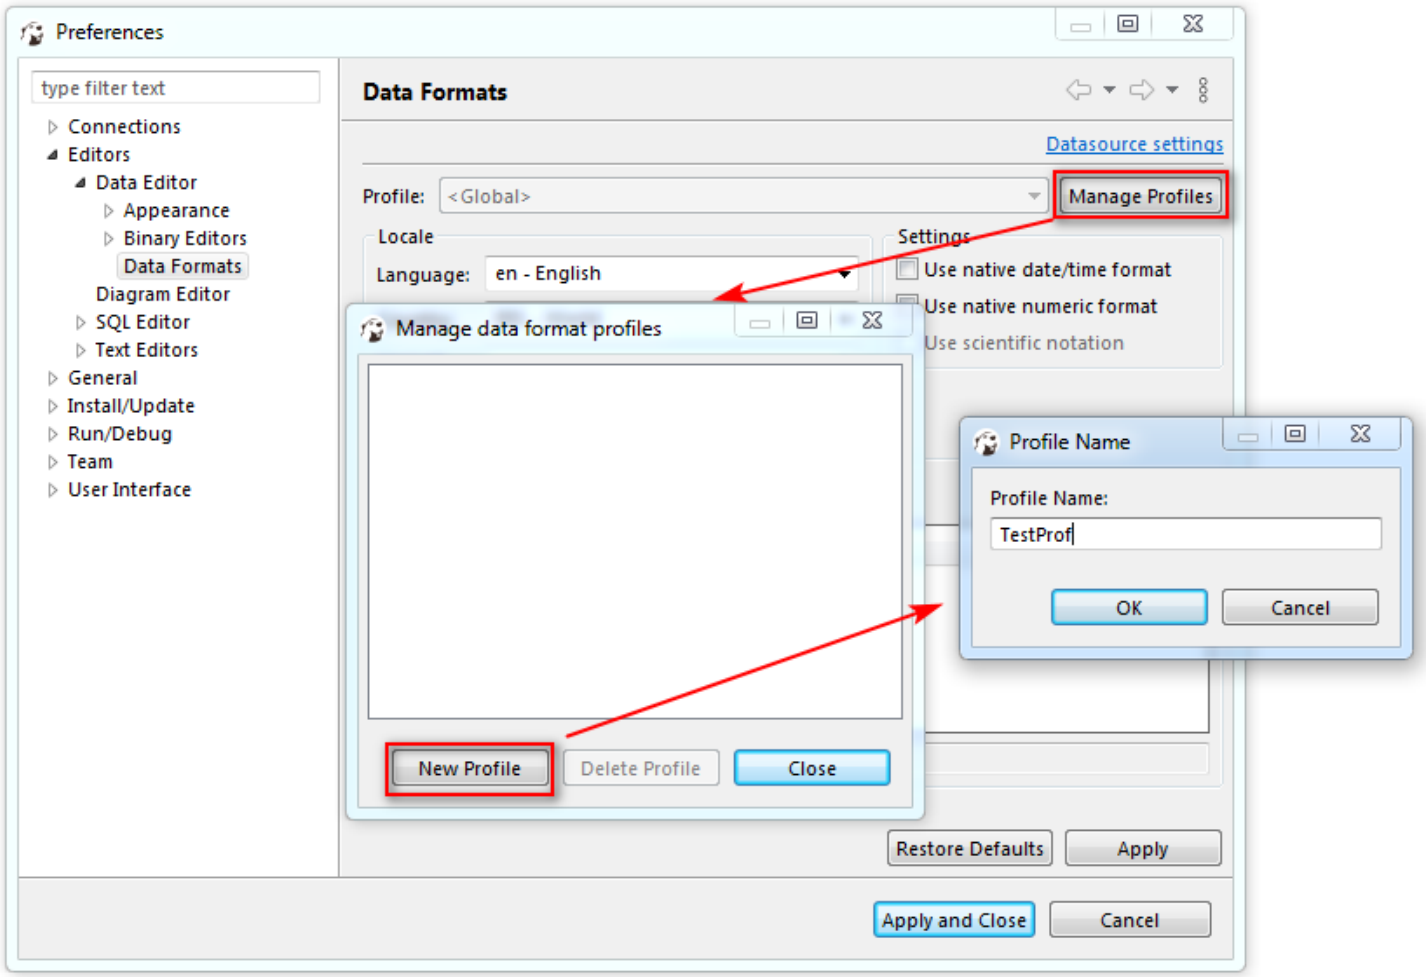

To delete a data format profile press the **Manage Profiles** button, then in the opened dialog window select the profile you want to delete and press the button **Delete Profile**.

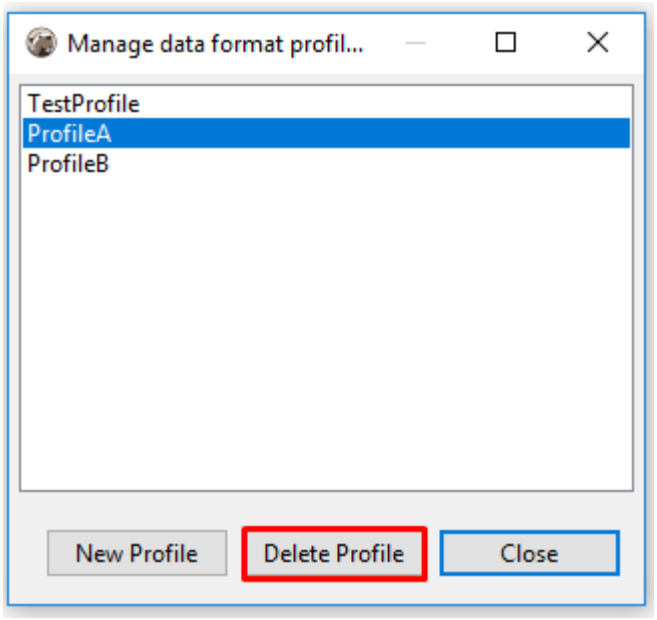

Changing Data Formats

The following groups of data format settings can be adjusted:

Locale

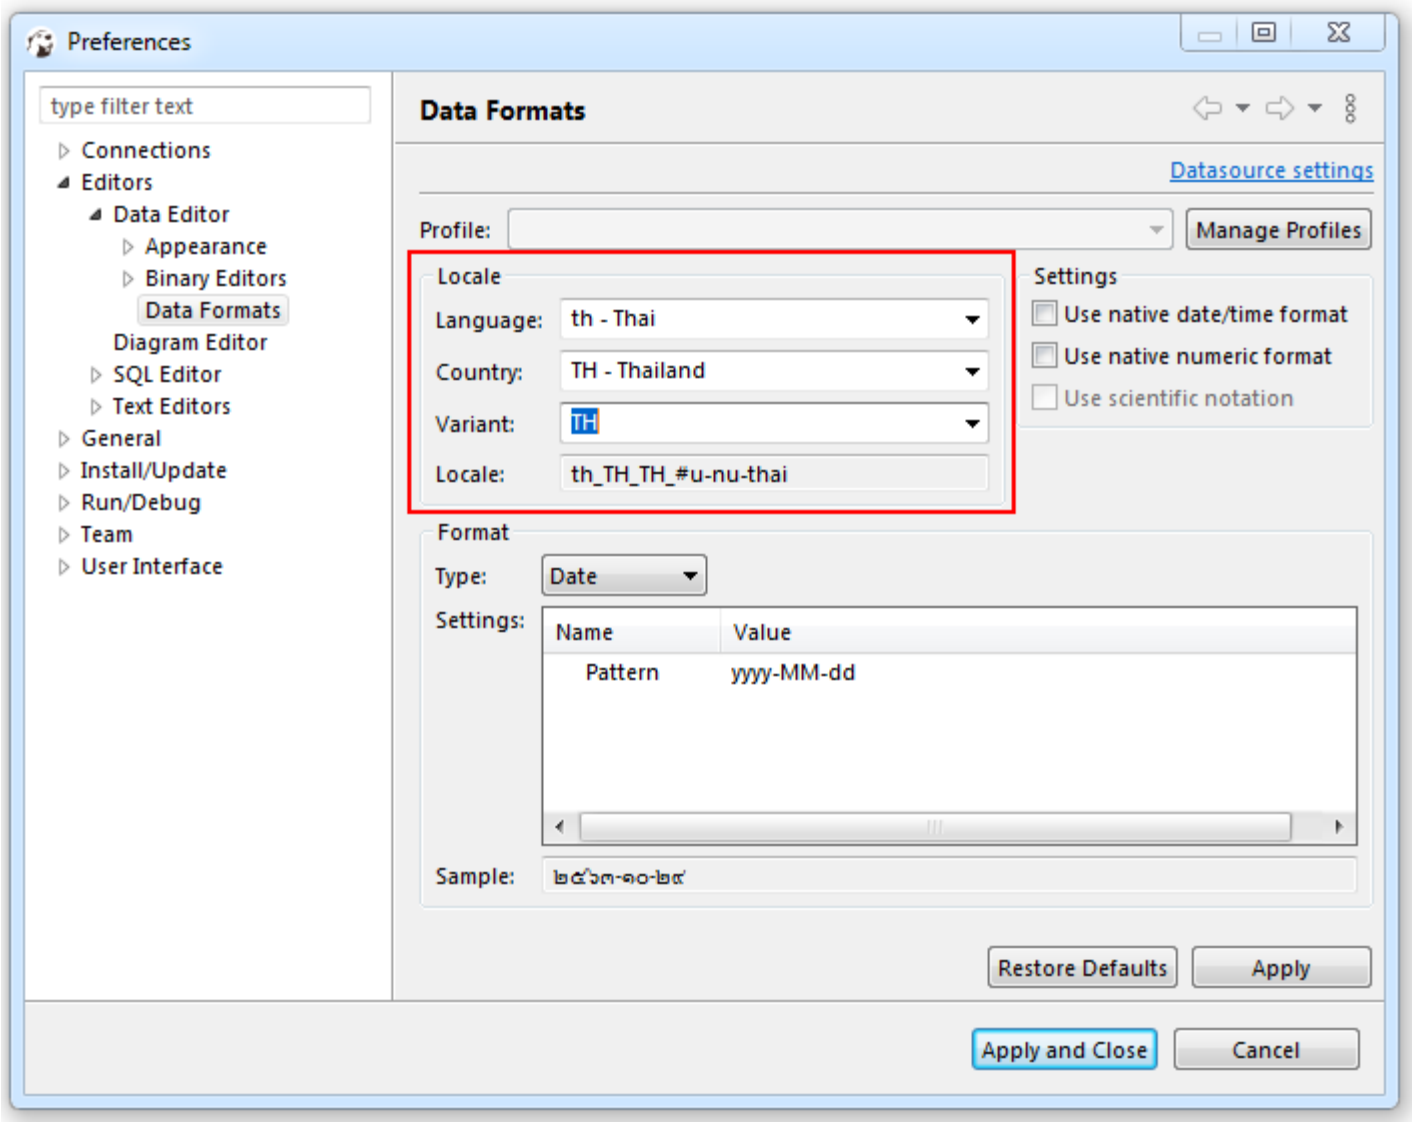

To define this setting select a language, country and variant if available.

Native Date/Time Mode

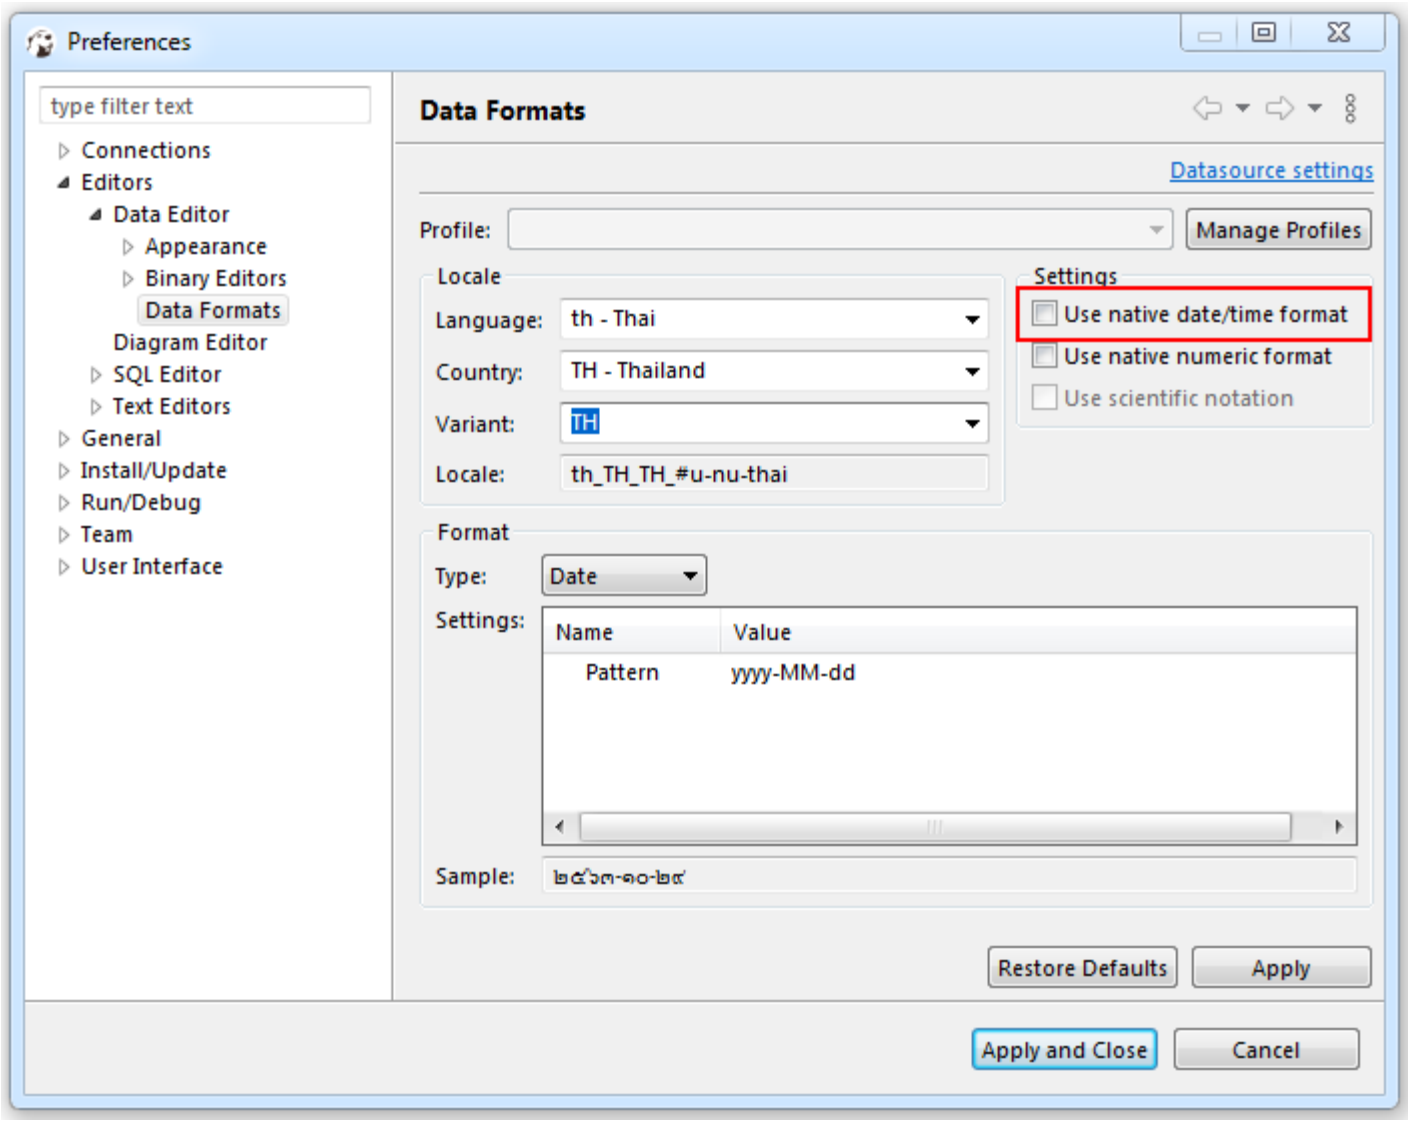

Select **Use native date/time format** check-box and the data format originally built-in to the datasource will be used.

You can change the format of the following data types:

Data Type Format

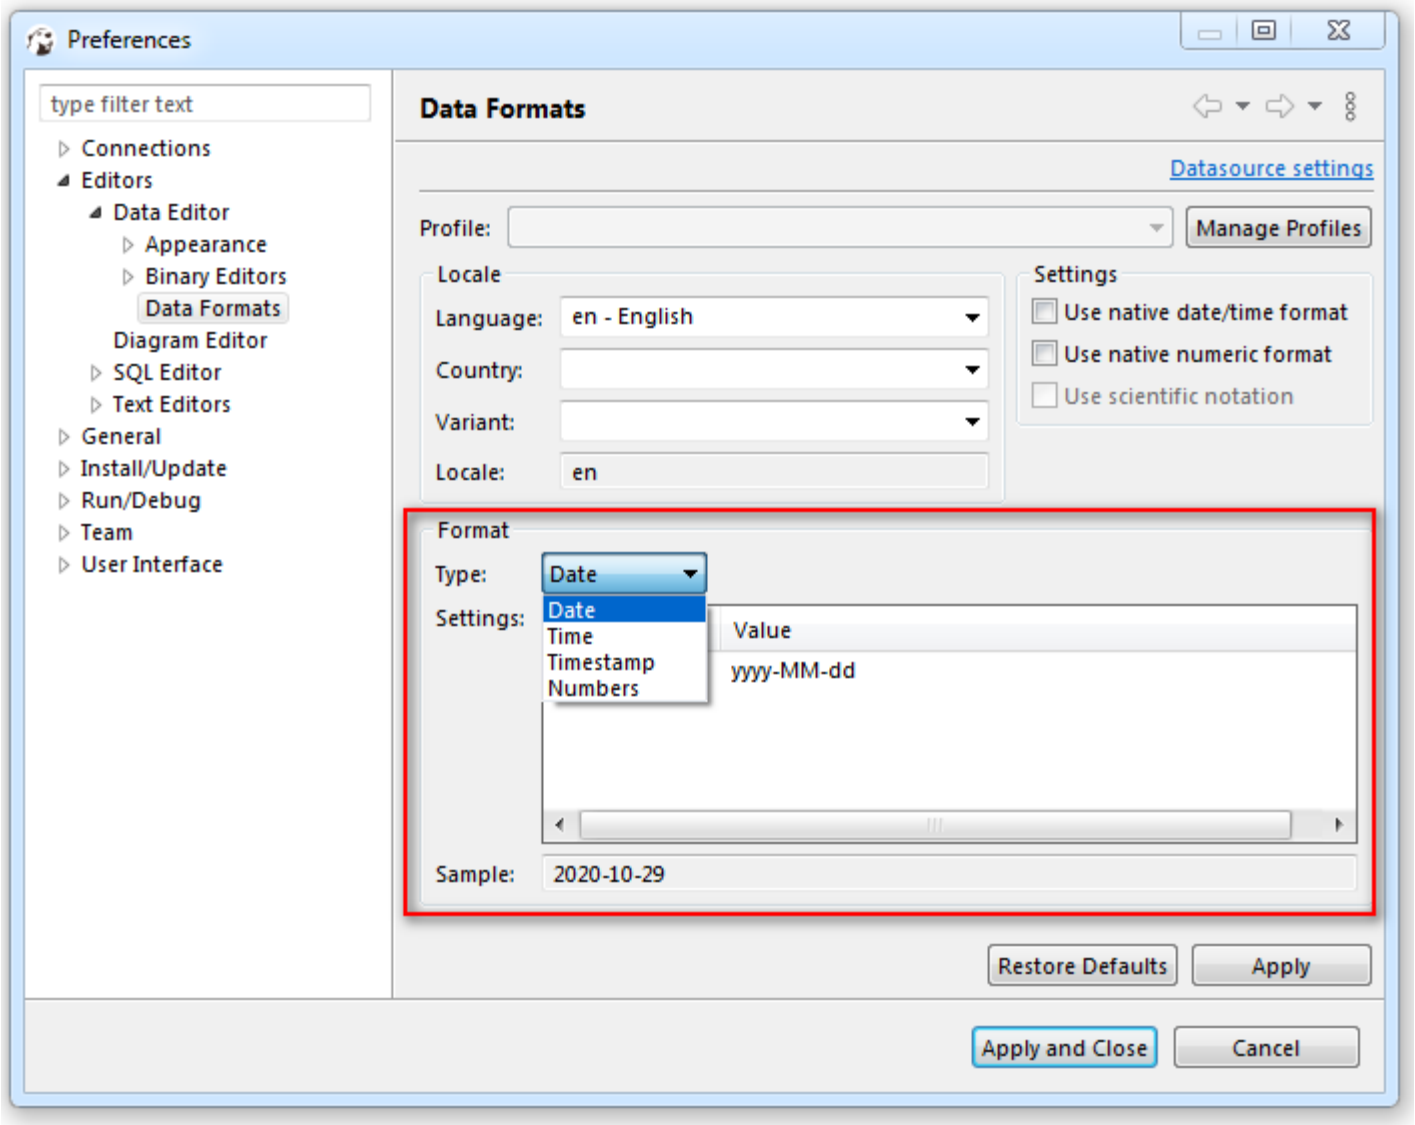

The format of the following data types can be customized:

- Date
- Time
- Timestamp
- Numbers

Date

The default value for this data type is *yyyy-MM-dd*.

Pattern	Description
yyyy or y	Year of era (4 digits)
yy	Year of era (2 last digits)
YYYY	Week year
M	Month in year without leading zeros
MM	Month in year

MMM	Short month name in year
MMMM	Month name in year
D	Day in year
d	Day in month without leading zeros
dd	Day in month
E	Day name in week
G	Era designator
z	General time zone
Z	RFC 822 time zone

Time

The default value for this data type is *HH:mm:ss*

Pattern	Description
H	Hour in day (0-23) without leading zeros
h	Hour in day in am/pm (1-12) without leading zeros
HH	Hour in day (0-23)
hh	Hour in day in am/pm (1-12)
a	Am/pm marker
m	Minute in hour without leading zeros
mm	Minute in hour
s	Second in minute without leading zeros
ss	Second in minute
S	Millisecond
ffffff	Microseconds

Timestamp

The default value for this data type is *yyyy-MM-dd HH:mm:ss*

Link on Java date pattern documentation

Numbers

The following parameters can be configured from this type of data:

- **Use Grouping** - Long numbers can be hard to read if they have too many digits. For example, the factorial of 30 is 33 digits long! Select this check-box to enable Grouping mode, in which digits are displayed in clumps of 3 or 4 (depending on the current radix) separated by commas.
- **Maximum integer digits** - Defines the maximum number of digits to the left of the decimal point.

- **Minimum integer digits** - Defines the minimum number of digits to the left of the decimal point.
- **Maximum fraction digits** - Defines the maximum number of digits to the right of the decimal point.
- **Minimum fraction digits** - Defines the minimum number of digits to the right of the decimal point.
- **Use data type scale for fraction digits** - Some numeric columns or parameters may have a predefined scale, that is the maximum number of digits to the right of the decimal point. Select this check-box if you want the predefined precision to be used.
- **Rounding mode** - Specifies a rounding behavior for numerical operations capable of discarding precision. Each rounding mode indicates how the least significant returned digit of a rounded result is to be calculated. To learn more, please refer to Oracle documentation.

To change the data type format, change the value displayed in the **Pattern** area,save the changes made by pressing the **Apply** button and observe the expected result in the **Sample** field.

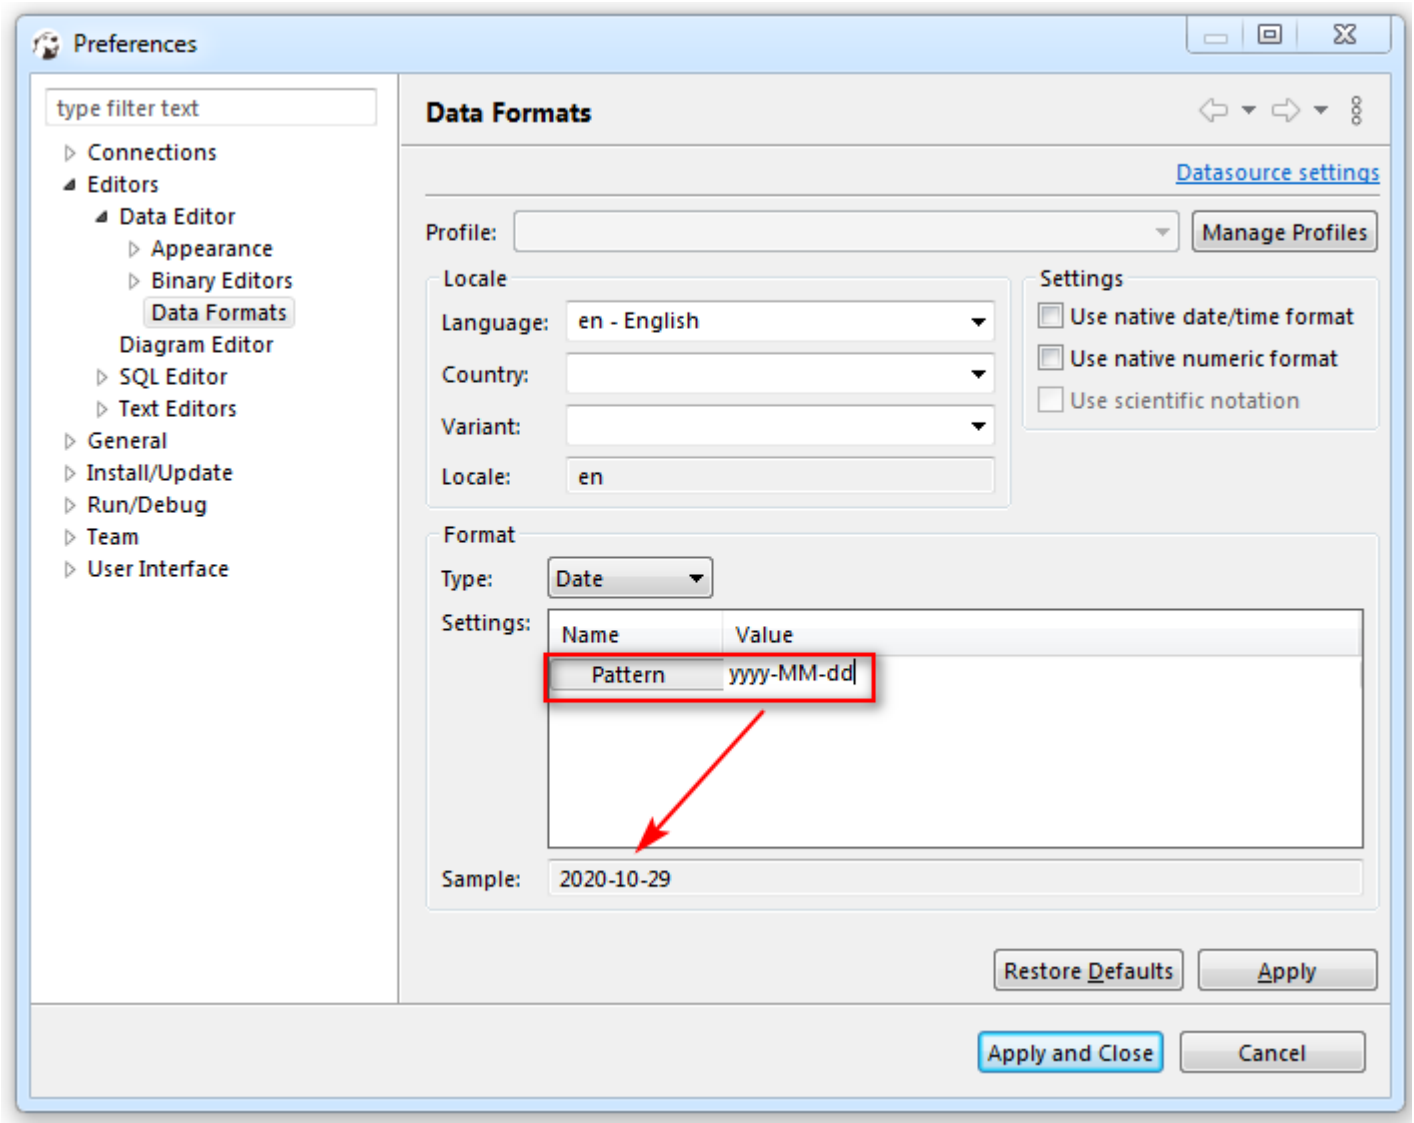

Restoring Default Data Formats

To restore the default data format settings, press the **Restore Defaults** button.

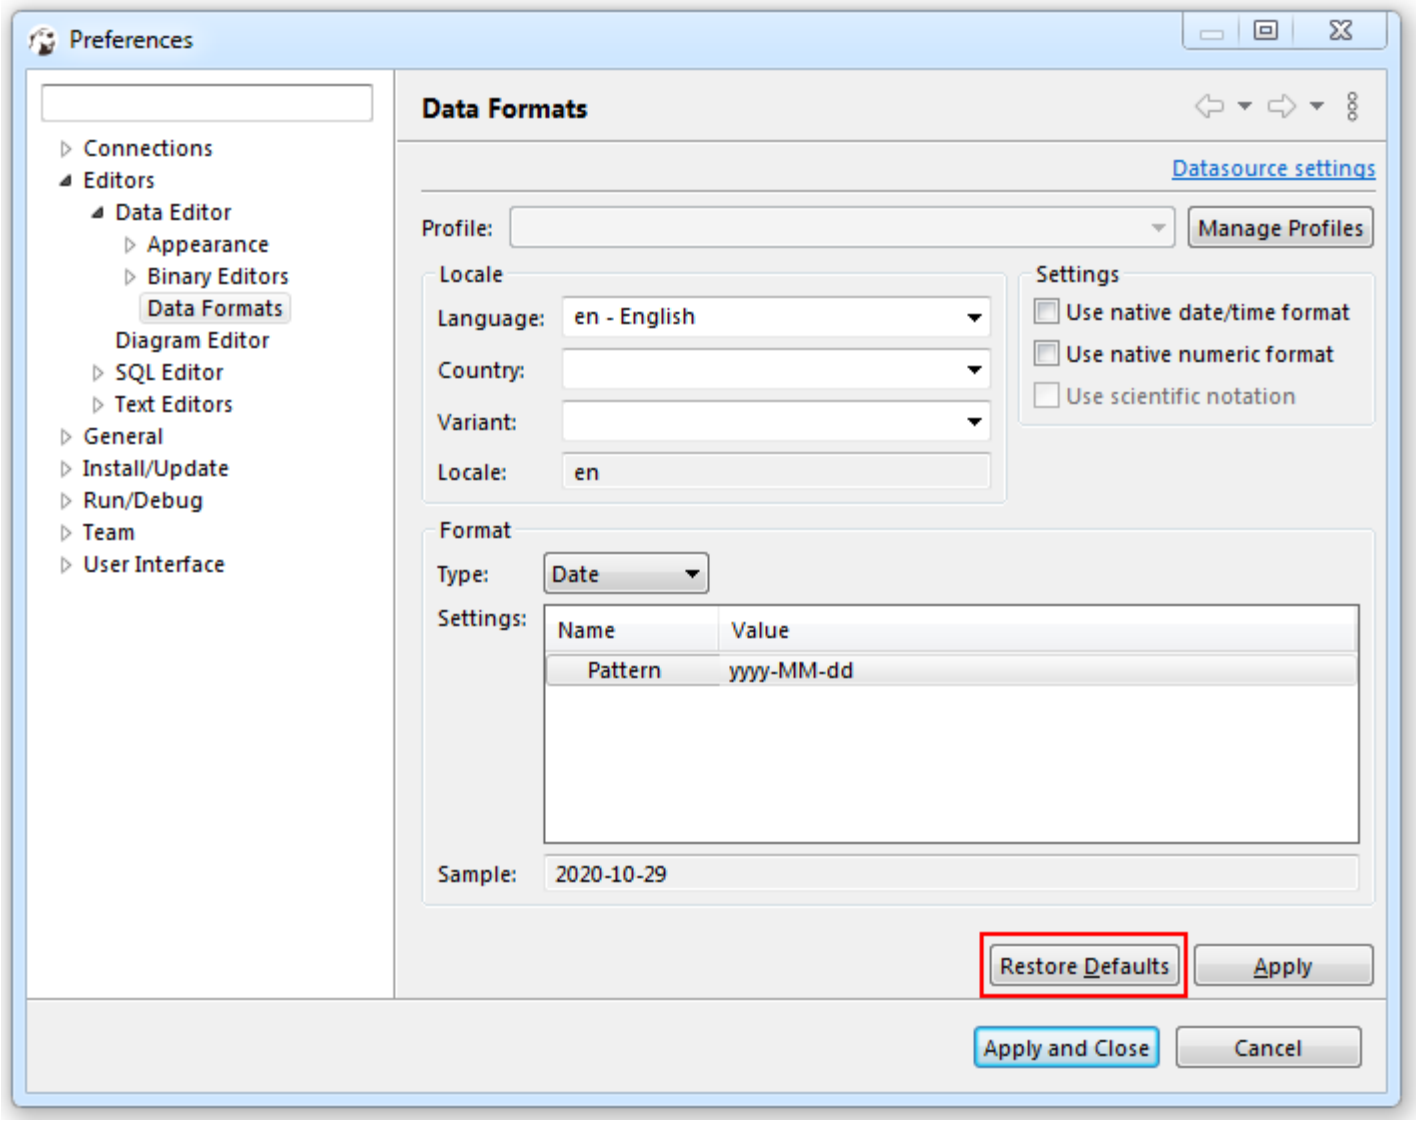

Datasource Settings

Press **Datasource settings** link to change data format settings for a particular datasource, then adjust the settings in the opened dialog box.

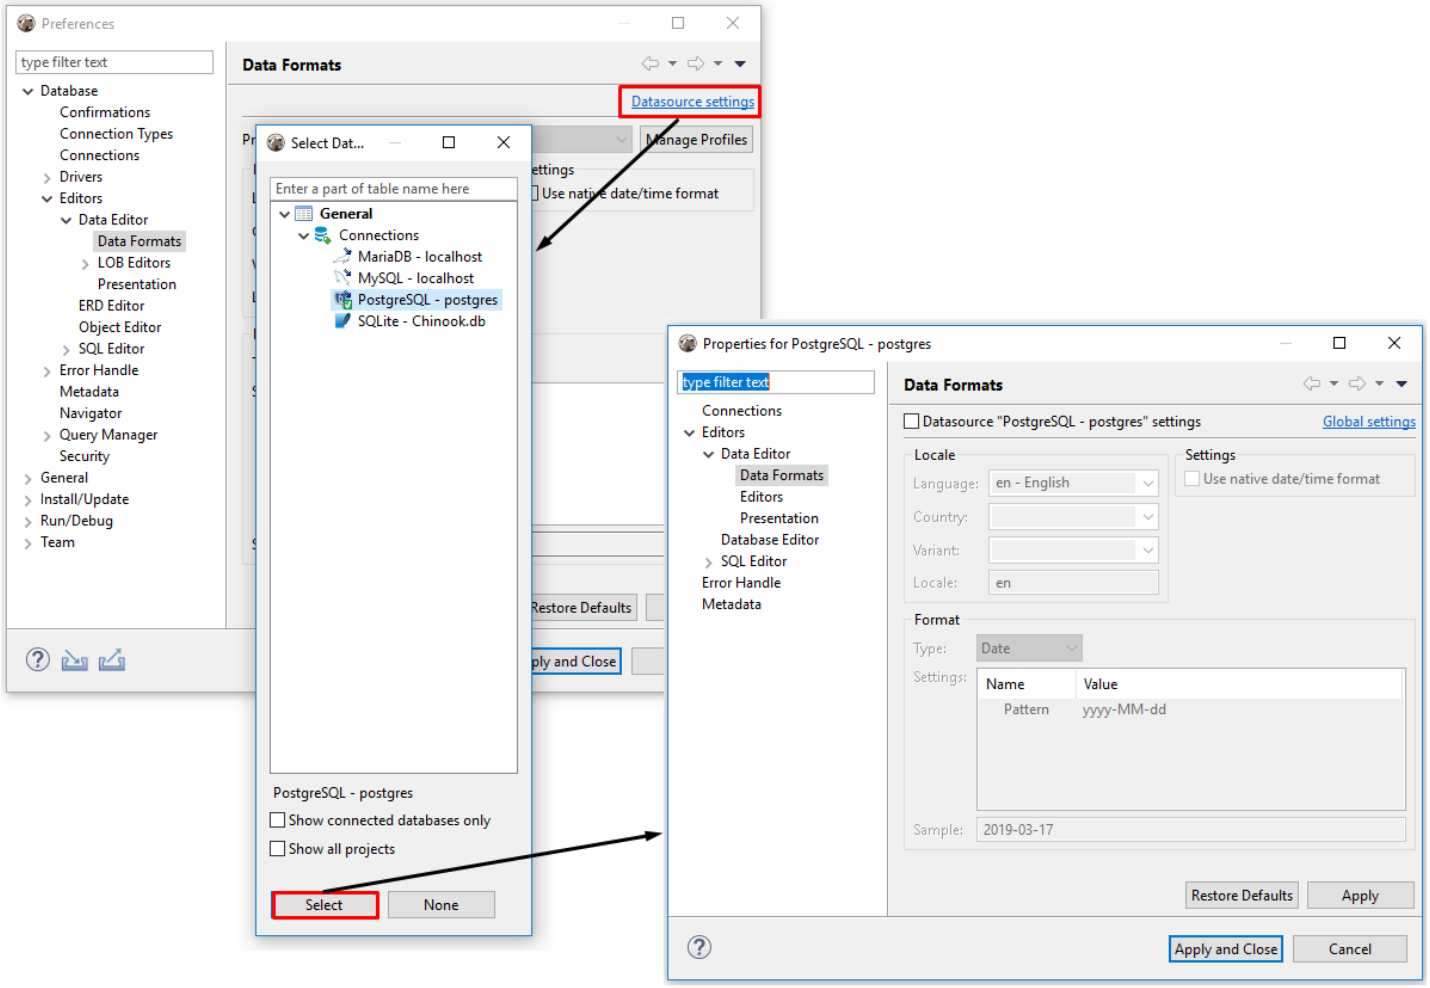

To save changes made press **Apply**.

SQL Editor

You can create multiple SQL scripts for a single connection. Every script opens in its own SQL editor. To open an SQL editor for some connection:

- Click this connection in the Database Navigator view and press `F3` or click **SQL Editor -> SQL Editor** on the main menu. Alternatively, you click **SQL Editor** on the context menu of this connection. DBeaver opens the Choose SQL script editor with saved SQL scripts linked to this connection. CLick the SQL script to open it in a separate tab.
- Click **Recent SQL Editor** on the context menu for this connection or on the main menu (**SQL Editor -> Recent SQL Editor**). This opens the latest used SQL editor. You can also open the most recent SQL editor using `Ctr` shortcut in the Database Navigator view.
- If you need to create a new SQL script, on the main menu, click **SQL Editor -> New SQL Editor** or press `F3` and then click **New Script** in the Choose SQL Script window.

DBeaver uses SQL syntax highlighting which depends on the database associated with the script. Different databases have different sets of reserved keywords and system functions.

NOTE: SQL Editor for a connection is different from SQL console for a table or view (right-click the table and click **Read data in SQL console**) in that, unlike the console, it can save scripts and changes to them.

You can see all your saved SQL scripts in the Project Explorer view in the **Scripts** folder.

The SQL editor includes the script panel at the top and results panel at the bottom:

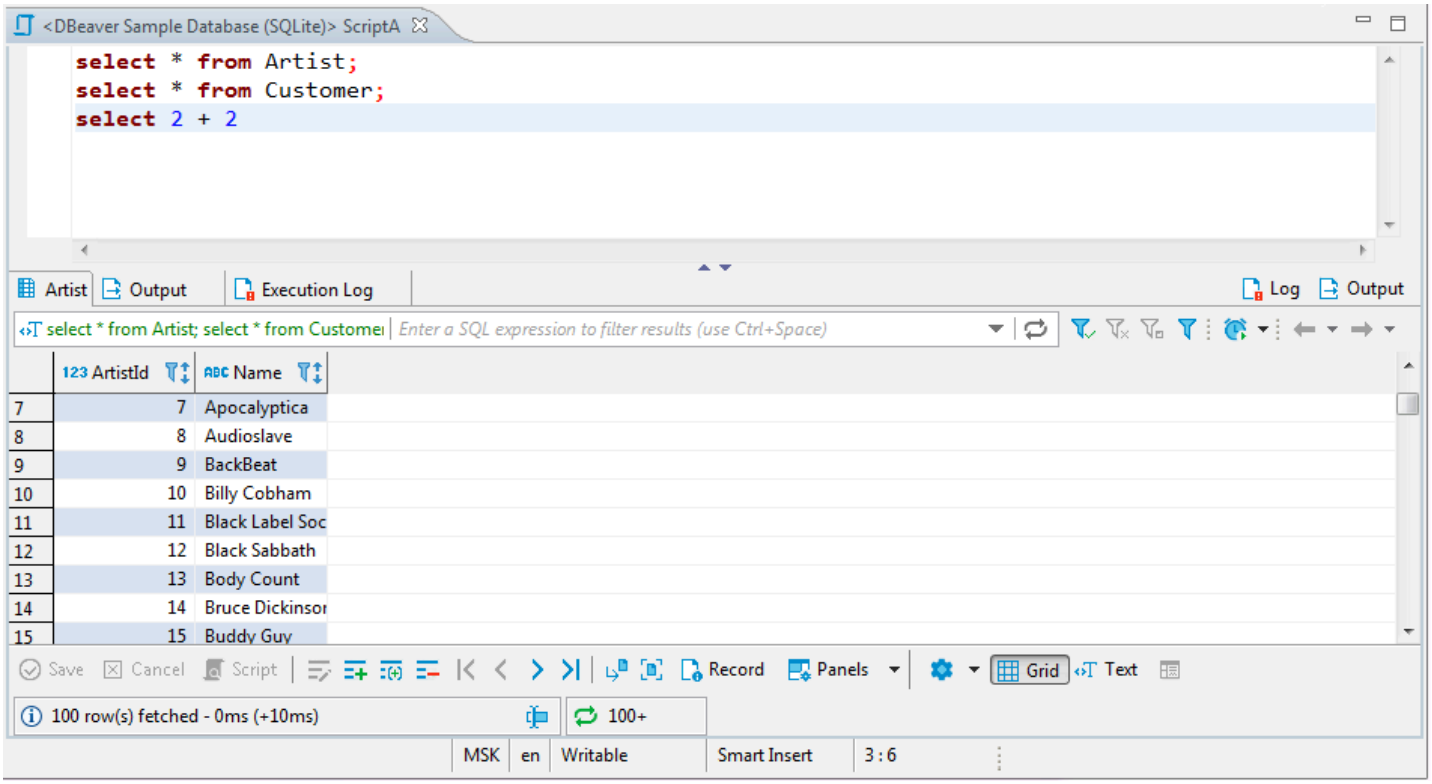

You can open SQL editor preferences by pressing `Alt+Enter`.

Results Panel

The results panel displays tabs with results in various formats. The tabs resulting from script execution represent instances of the Data Editor. You can create, edit and execute SQL scripts in the script panel and then see the results in the result tabs.

The results panel provides **Output** and **Log** views of results.

The execution **Log** tab contains all queries executed in the current SQL editor:

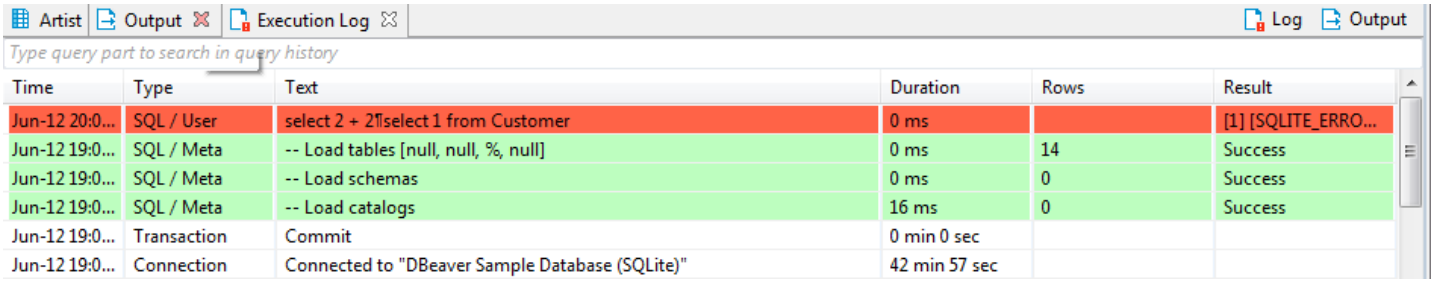

The **Output** tab contains all server-side database messages/warnings generated by a database when you execute queries. This feature is supported only by a few database engines (Oracle, SQL Server and some other ones).

Layout Adjustment

You can modify the layout of the SQL Editor by showing/hiding the results panel and changing the horizontal /vertical position of the panes.

- To toggle (hide/show) the results panel, press `CTRL+6` or right-click anywhere in the script pane and, on the context menu, click **Layout -> Toggle results panel**.
- To maximize the results panel, press `CTRL+Shift+6` , or double-click the results tab name, or right-click anywhere in the script panel and, on the context menu, click **Layout -> Maximize results panel**.
- To switch between the script panel and the results pane, press `Alt+6` or right-click anywhere in the script panel and, on the context menu, click **Layout -> Switch active panel**.

To position both panels horizontally, right-click anywhere in the script panel and, on the context menu, click **Layout -> Horizontal**.
To position both panels vertically, right-click anywhere in the script panel and, on the context menu, click **Layout -> Vertical**.

Hyperlinks

You can press and hold `Ctrl` and at the same time move the mouse over SQL text. If DBeaver recognizes some identifier as a table/view name, it presents it as a hyperlink. You can click the hyperlink to open this object`s editor:

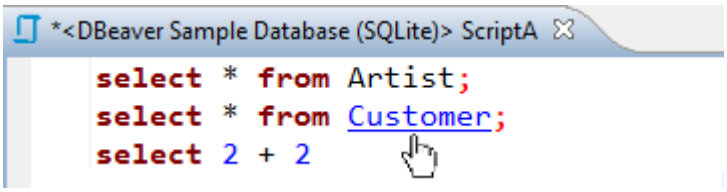

Active Database/Schema Selection

You can change the connection associated with the current SQL editor or change the active database/schema, at the same time retaining the SQL text.
To change the connection, press `Ctrl+9` or click the **Active datasource** box on DBeaver`s main toolbar:

The Select Data Source dialog box opens. In the tree of connections, click the required connection and then click **Select**. To disassociate the SQL Editor with any connection, click **None**:

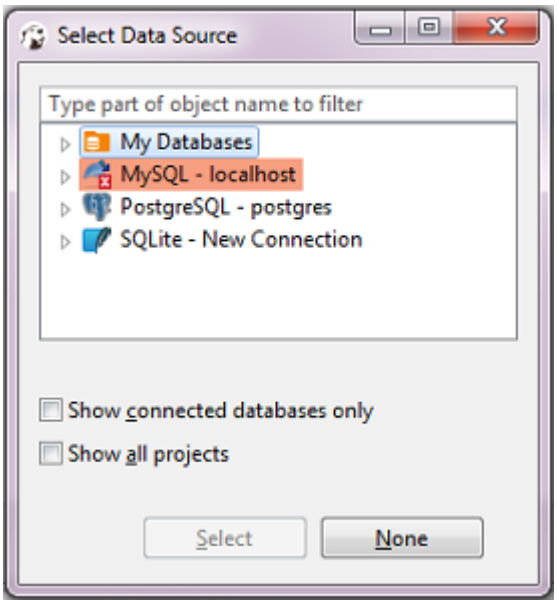

To change the active schema, press `Ctrl+0` or click the **Active Catalog/Schema** box in DBeaver`s main toolbar:

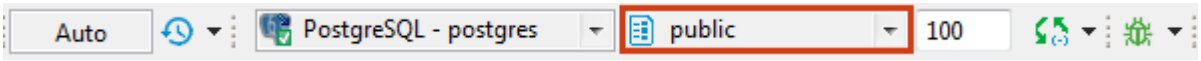

The Choose catalog/schema dialog box opens. In the list of schemas, double-click the required schema:

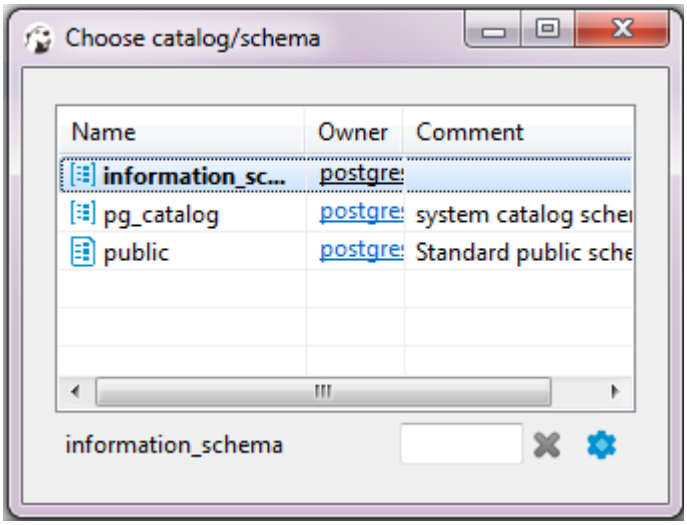

If there are many schemas and they do not fit in the dialog box use the search field to find the schema you need:

To configure the set of columns to be visible for each schema in the dialog box, click the **Configure columns** button ( ).

You can easily associate the SQL Editor with the connection that is currently in focus in the Database Navigator (the focus can be on any object of the connection - a table, a folder, etc.) - click the **Set connection from navigator** button in DBeaver`s main toolbar:

The reverse action is also possible: you can set the focus of the Database Navigator to the active connection of the SQL Editor - press `Ctrl+Shift+,` or click the arrow next to the **Set connection from navigator** button in DBeaver`s main toolbar and then click **Link with editor**:

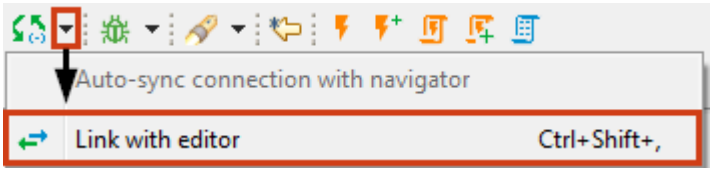

SQL Templates

Templates allow you to insert frequently used SQL statements into an SQL script.

To see available templates, press `Ctrl+Alt+SPACE` or right-click the line in the script pane and click **SQL Template** on the context menu. A box appears with a list of available templates:

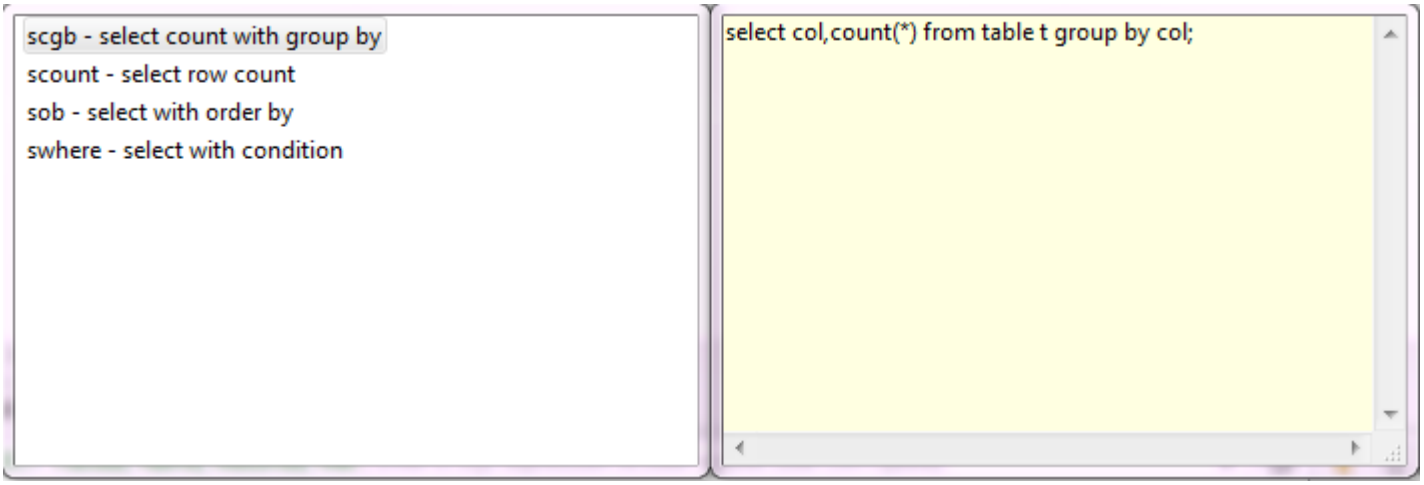

To apply a template, in the SQL Editor, in the script pane:

- Type the template name and press `Tab`
- Right-click the line where you want to insert a template expression, click **SQL Template** on the context menu, and then, in the list of templates, double-click the required template name.
  The template SQL statement appears in the script.

To edit/add/remove templates, click Configure () in the bottom toolbar, then click **Preferences -> SQL Editor -> Templates**. For more information about managing templates, please visit Eclipse Website.

Standard Eclipse templates:

Variable	Description
${cursor}	Specifies the cursor position when the template edit mode is left. This is useful when the cursor should jump to different place than to the end of the template upon leaving the template edit mode.
${year}	Takes the current year value
${date}	Takes the current date value
${time}	Takes the current time value
${dollar}	Takes the dollar sign $. Alternatively, two dollar signs can be used: $$.
${user}	Takes the user name
${word_selection}	Takes the content of the current text selection
${line_selection}	Takes content of all currently selected lines

DBeaver-specific templates:

Variable	Description

${schema}	Takes the current schema name
${catalog}	Takes the catalog name
${table}	Takes the current table name (from the active catalog/schema)
${column}	Takes the column name (from the current table)

SQL Assist and Auto-Complete

The SQL Assist feature provides auto-completion of database object names and SQL commands and other keywords in queries.

To perform some object name auto-complete, press `Ctrl+Space` or right-click the required place in the query and click **SQL Assist** on the context menu. DBeaver searches for potentially suitable objects in already loaded database metadata and in the database system tables.

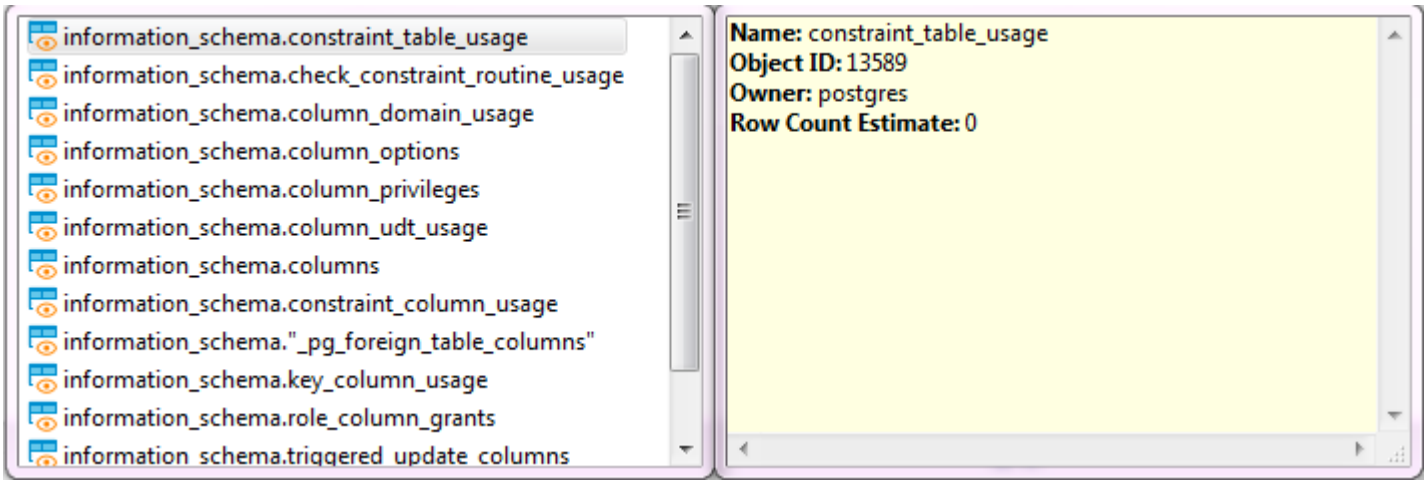

When you start typing an SQL keyword in a statement, DBeaver offers auto-complete options as well.
Another auto-complete function is search for completion only within already entered identifiers - press `Ctrl+Shif`
.

Also you can press `Ctrl+Space` after asterisk in query like "SELECT *FROM tableName" or like "INSERT INTO tableName ()*" (brackets are important) (you can use ()[]{} brackets) - the asterisk will be replaced with a list of all table columns.

SQL Formatting

To format SQL text, select it and press `Ctrl+Shift+F` or right-click the selected text and click **Format -> Format SQL** on the context menu.

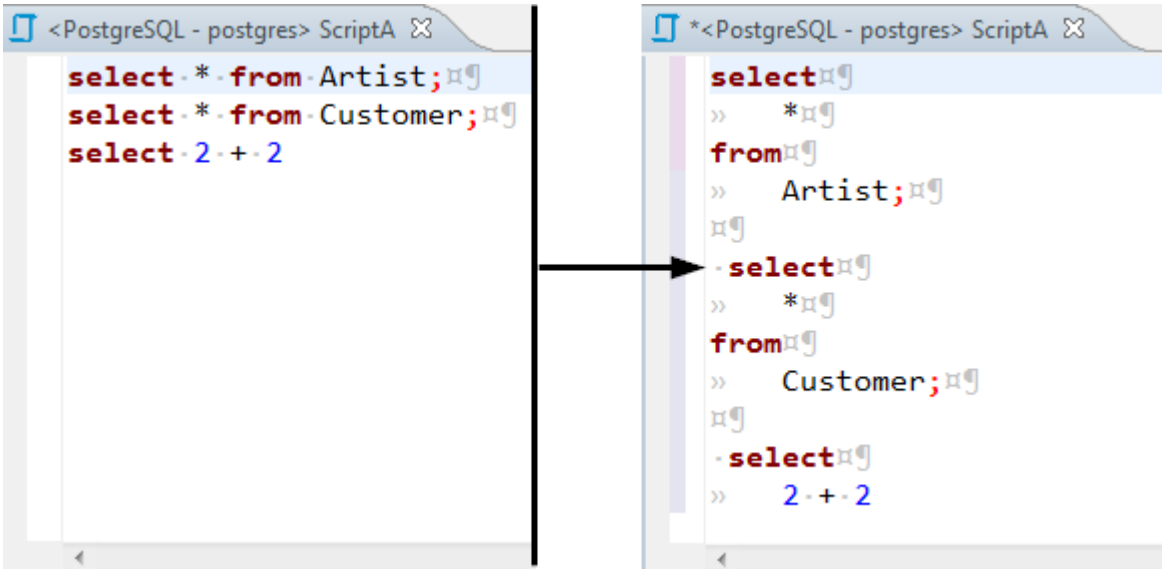

To format a script to upper or lower case, highlight the SQL text, then right-click it and click **Format -> To Upper Case / To Lower Case**, respectively, on the context menu.

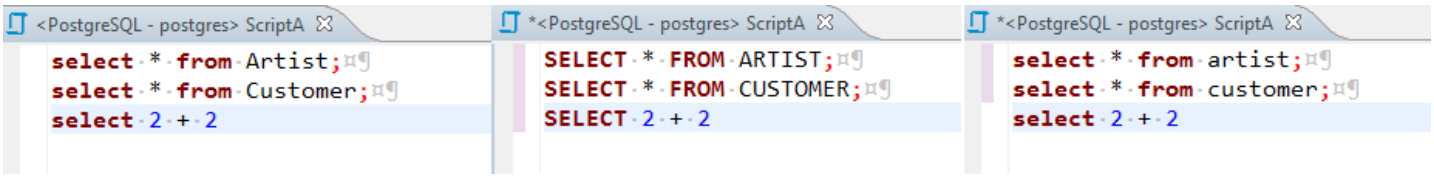

To comment out an SQL line, press `Ctrl+/` or right-click the line and click **Format -> Toggle Line Comment** on the context menu. To uncomment a commented line, manually remove the commenting syntax, or press the same button combination, or right-click the line and click the same item on the context menu.

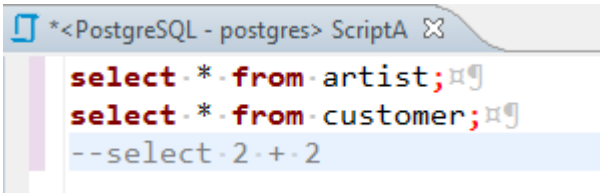

To comment out a block of text, select the text, then press `Ctrl+Shift+/` or right-click it and click **Format -> Toggle Block Comment** on the context menu. To uncomment a commented block of text, either manually remove the commenting syntax or select the same block of text, right-click it and click the same item on the context menu or press the same button combination.

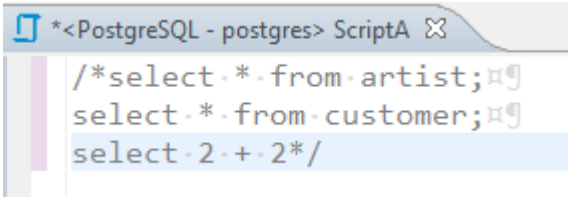

SQL Execution

You can execute one query, a highlighted portion of a script, or a whole script. You can execute them using:

- Shortcut key combinations (see details further in this article)
- Tools in the main toolbar:
- Context menu (right-click the query):

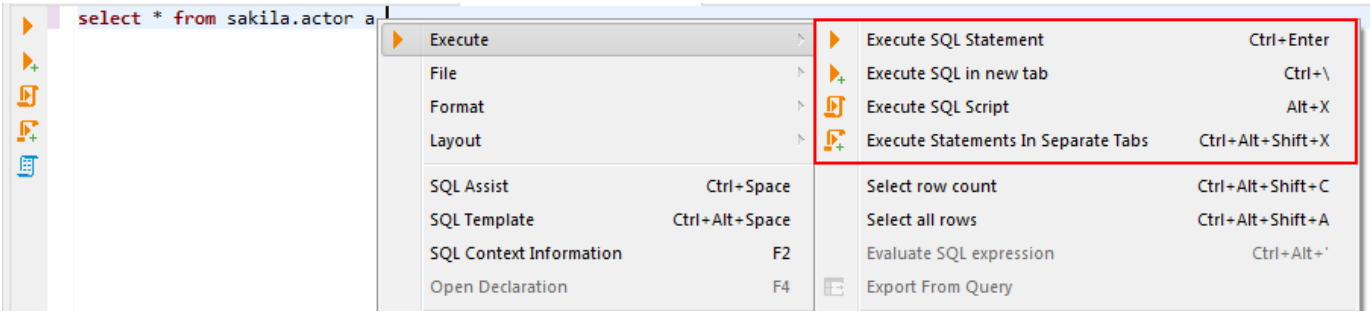

- DBeaver main menu:

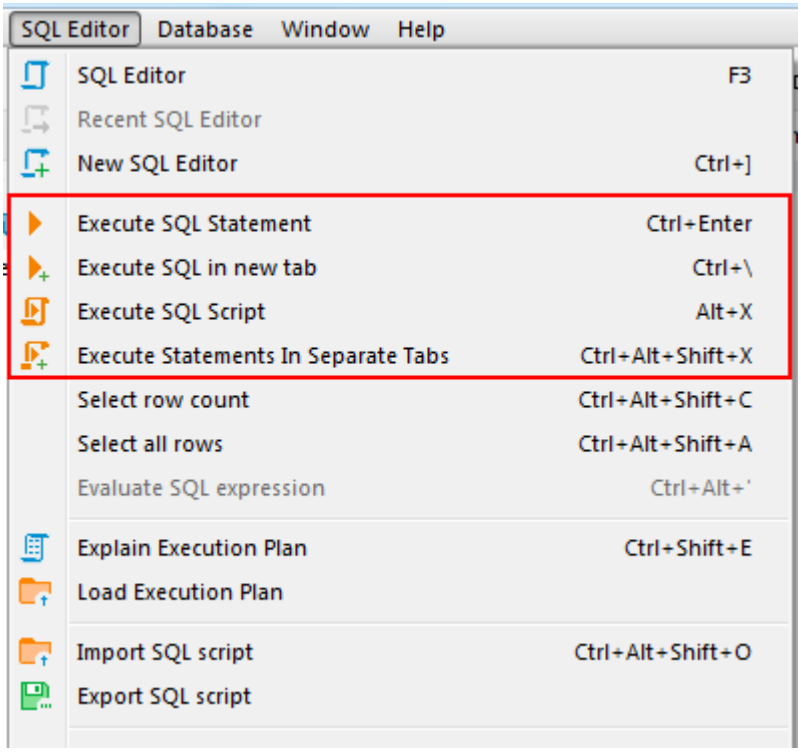

To execute a query under cursor or selected text, press `Ctrl+Enter` or right-click the query and click **Execute -> Execute SQL Statement** on the context menu. You can do the same using the main toolbar or main menu: **SQL Editor -> Execute SQL Statement**. This executes the SQL query under cursor or selected text and fills the results pane with the query results.

To execute a query under cursor in a separate tab, press `CTRL+\` or right-click the query and click **Execute -> Execute SQL in new tab** on the context menu. The same can be done using the main toolbar or the main menu:

SQL Editor -> Execute SQL in new tab. This executes the SQL query under cursor or selected text and creates a new results tab.

To execute the whole script, press `Alt+X` or click **Execute -> Execute SQL Script** on the context menu or **SQL Editor -> Execute SQL Script** on the main menu or in the main toolbar. This executes all queries in the current editor (or selected queries) as a script. DBeaver parses queries one by one using a statement delimiter (";" by default) and executes them consecutively. You can configure the script execution behavior in the SQL editor preferences (Right-click the script and click **Preferences** on the context menu).

To execute a script opening each query results in a separate tab, press `Ctrl+Alt+Shift+X` or click **Execute -> Execute Statements In Separate Tabs** on the context menu or **SQL Editor -> Execute Statements In Separate Tabs** on the main menu or in the main toolbar. The executes all queries in the script, but opens multiple result tabs. Each script query is executed in a separate thread (that is, all queries are executed simultaneously). NOTE: Be careful with this feature. If you execute a huge script with a large number of queries, it might cause unexpected problems.

Result tabs

Single query may generate several result sets represented by tabs. These tabs are linked to the query they are executed from.

- To close an individual tab, press `CTRL+Shift+\` or middle-click on a tab header.
- To close all tabs expect current, click **Close all result tabs except this** on the context menu of this tab.
- To close all tabs of desired query, click **Close all result tabs of same query** on the context menu of this tab.

Naming

Tabs are named in a form of `Results <A> (<B>)` , where:

- **A** is *an index of query*
- **B** is *an index of result set of this query*

Pinning

Tabs can be moved around by dragging them with mouse and pinned using **Pin tab** on the context menu of desired tab. Pinned tabs are stacked on the left, can be moved among other pinned tabs, but can't be mixed with unpinned tabs, also pinned tabs can't be closed without being unpinned first and can't be overwritten by executing a query in it (by making this tab active).

Dynamic Parameter Bindings

You can use dynamic parameters in your SQL queries. The parameter format is :name. When you execute a query which contains dynamic parameters, DBeaver displays a dialog box in which you can fill the parameter values:

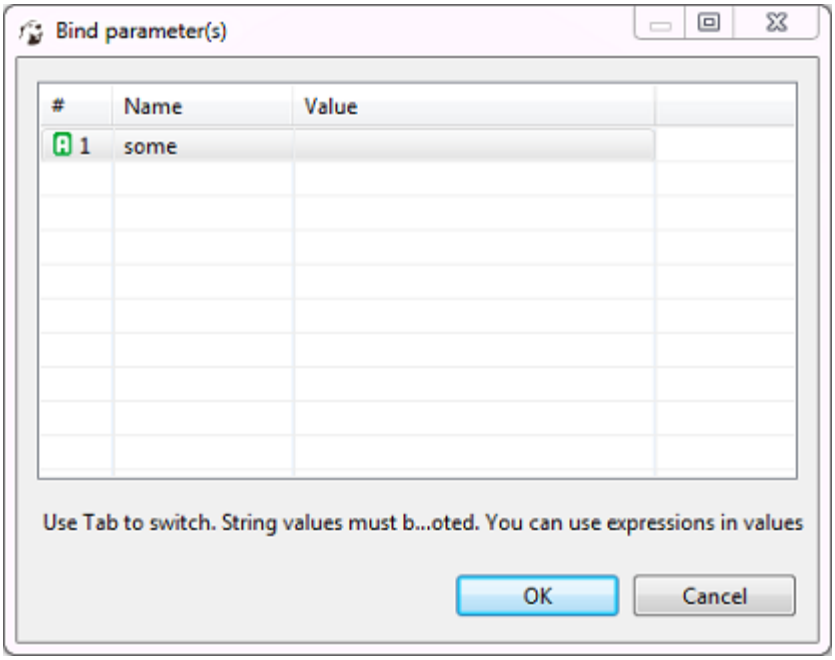

You can also use anonymous parameters (?), but you will need to enable them in SQL editor preferences:

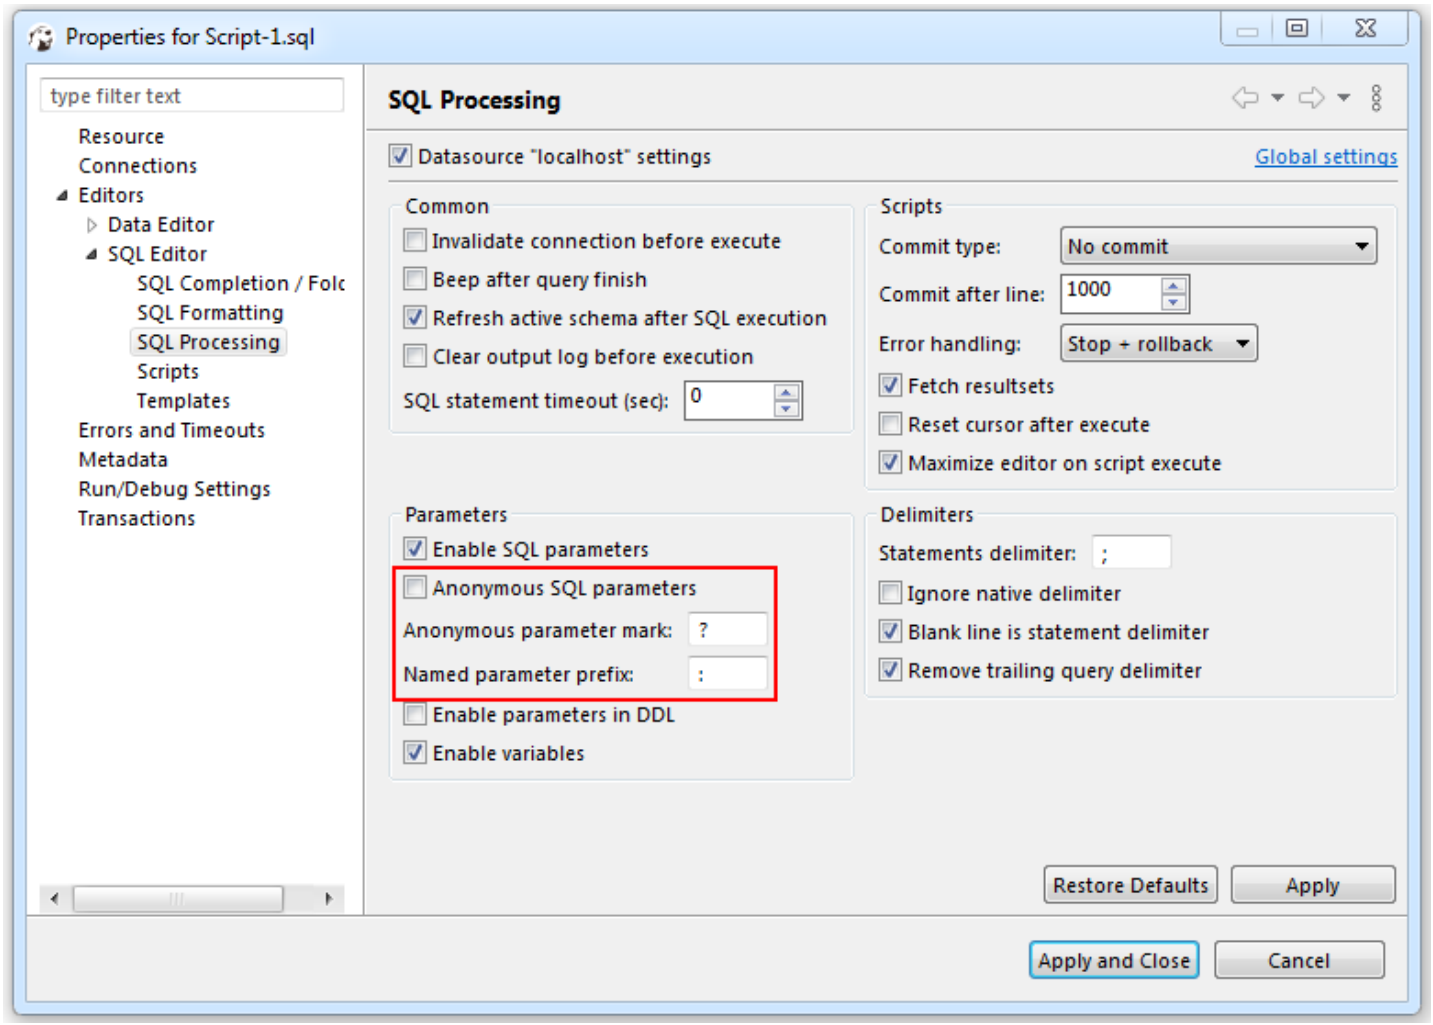

You can open SQL editor preferences by pressing `Alt+Enter`.

SQL Expression Evaluation

To evaluate an SQL expression, right-click the expression and click **Execute -> Evaluate SQL expression** on the context menu. This command basically performs a query of **SELECT [expression] FROM DUAL** type:

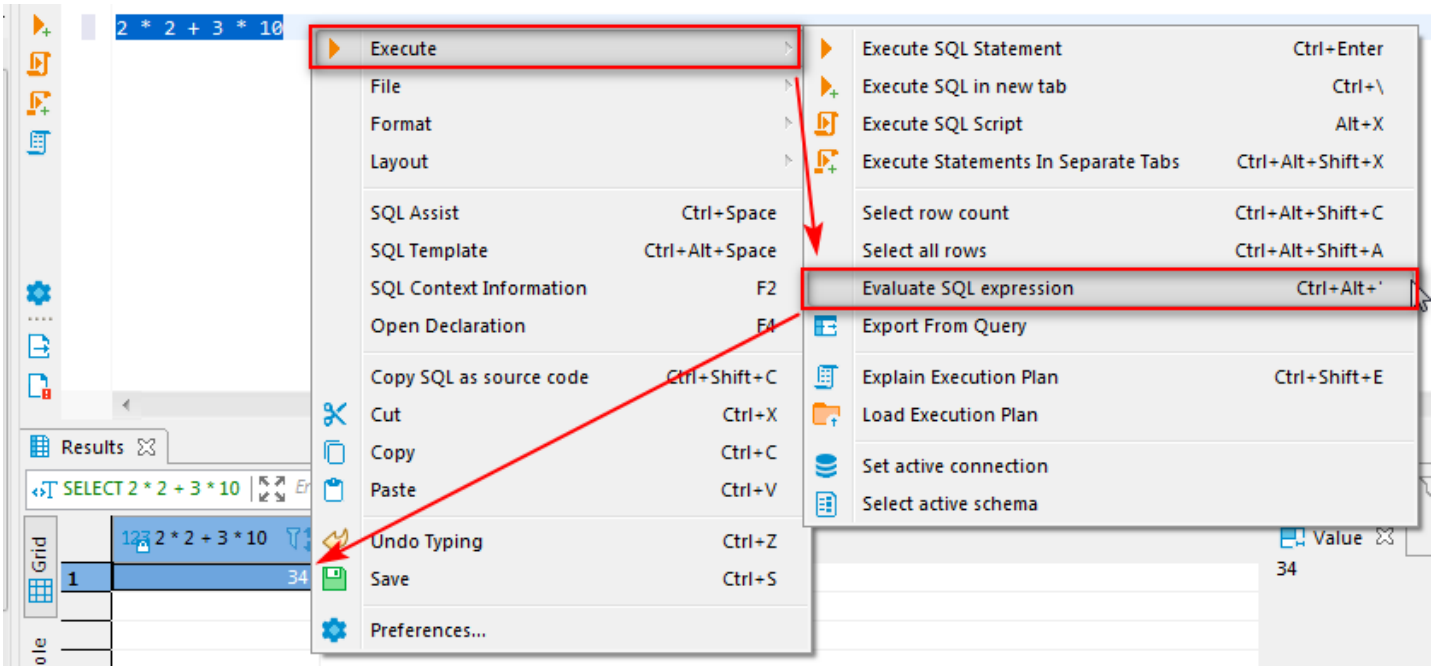

Row Count

If you want to know how many rows an SQL query will produce, you need to apply the Row Count feature – highlight and right-click the SQL text and then click **Execute -> Select row count** on the context menu:

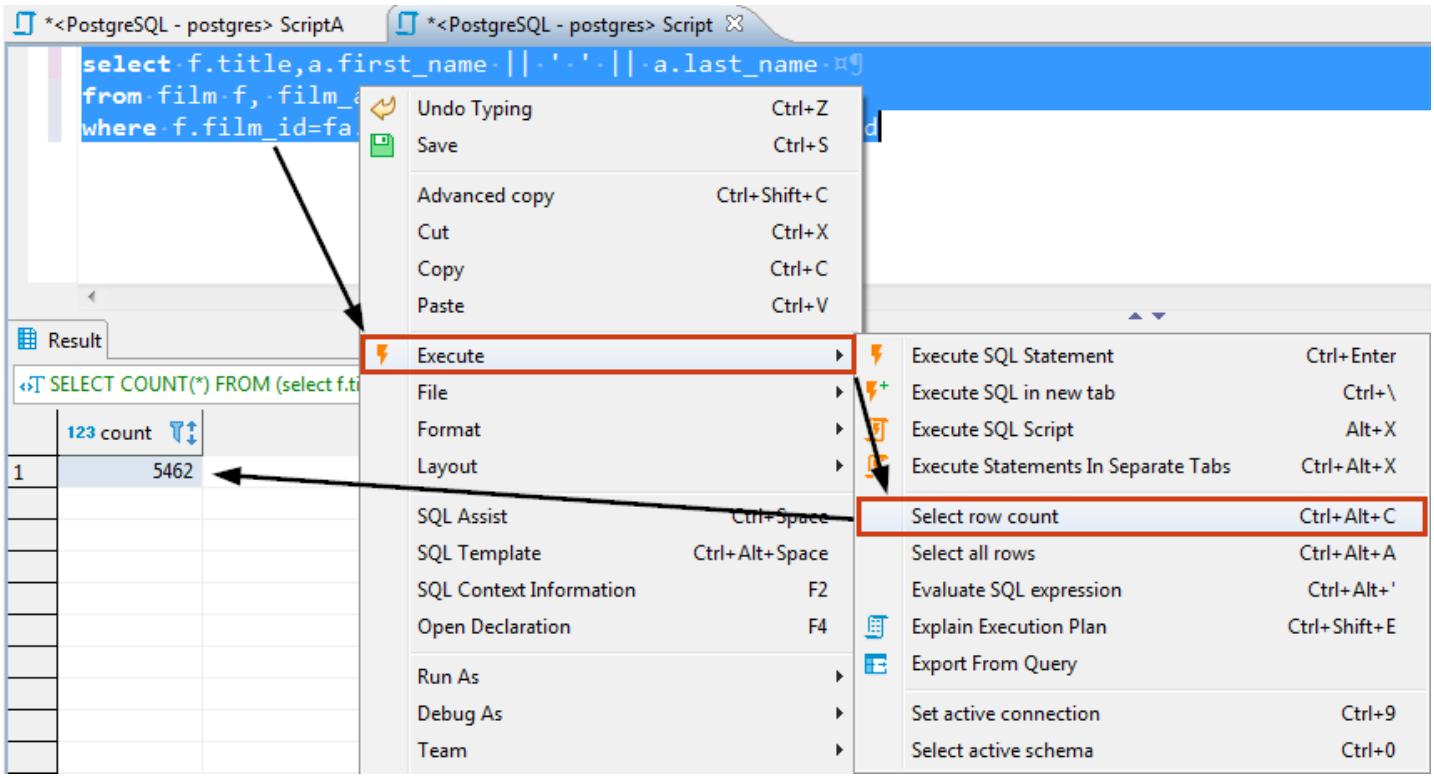

Query Export

It might be useful to export a query if you have a very long-running query and you do not need to see its results in the results panel. You can directly export the current query results to a file/table by right-clicking the query and then clicking **Execute -> Export From Query** on the context menu:

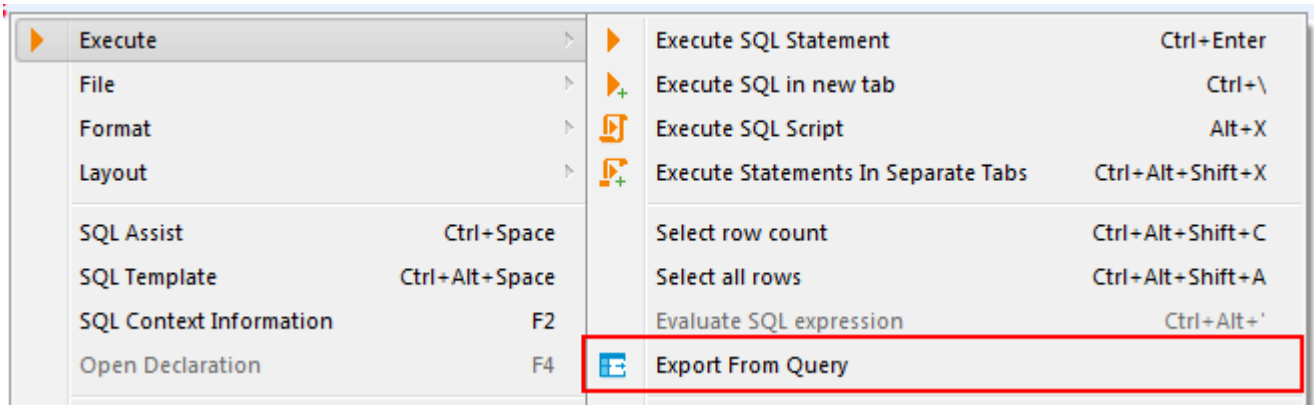

The Data transfer wizard opens. Go through its steps to complete the export of the query.

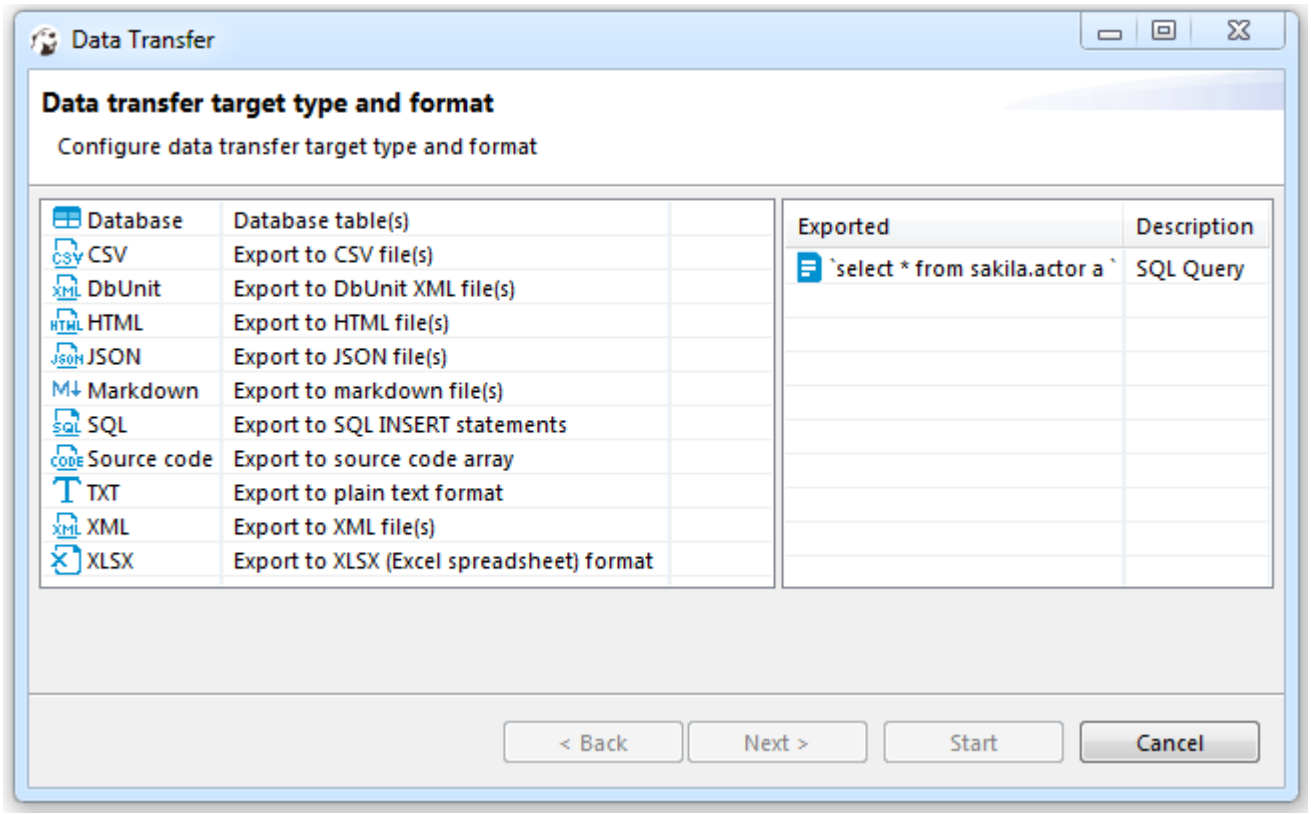

Client-side commands

You can use special commands in SQL scripts.
These commands are executed on DBeaver side, not on server-side.

Name	Description
@set var = value	Sets default value for SQL parameter
@echo text	Prints string into server output viewer
@include file	Includes script file from file system

Miscellaneous

- To select the current query row count, press `Ctrl+Alt+Shift+C`.
- To open the definition of the database object currently in focus (under cursor) in a viewer/editor, press `F4`.

Query Execution Plan

Execution Plan

Simple plan view

If a database driver supports execution plan visualization, you can see the execution plan of the current query (under cursor) by pressing `Ctrl+Shift+E` or clicking **Explain execution plan** on the context menu or in the main toolbar: The execution plan command generates a tree of query execution as one of the result tabs and is convenient in estimating if the query/script is quick/optimal enough:

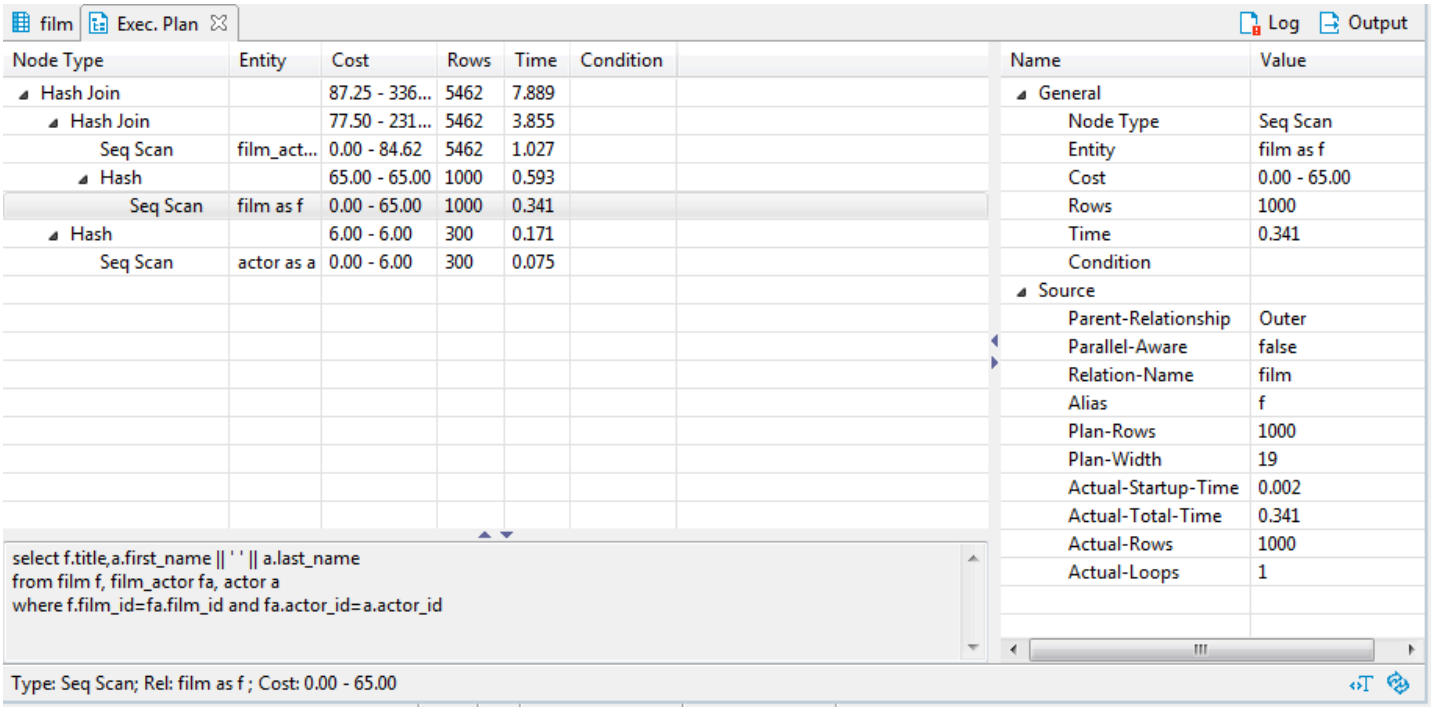

You can click the rows of the execution plan to see their details (statistics) in the panels below and to the right of the plan.

To reevaluate the plan, click the **Reevaluate** button ( ). To see the source script on which the plan is based, click the **View Source** button ( ).

Advanced plan view

In DBeaver Enterprise Edition you can use advanced (graph) execution plan visualization.
This visualization shows most expensive (cost-based) plan nodes. You can hide all irrelevant nodes, see node details, use horizontal or vertical pln layout, export it to image or save as json to pass plan information to a colleague.

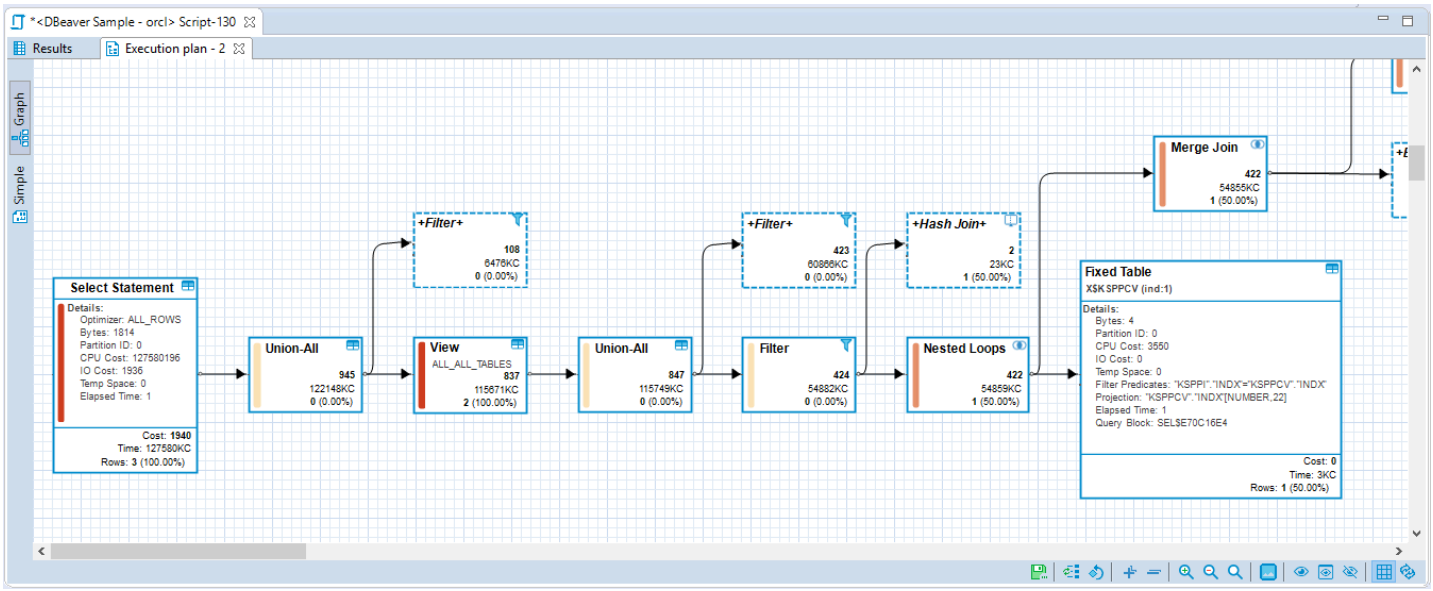
*<DBeaver Sample - orcl> Script-130
Results
Execution plan - 2
Graph
Simple
Select Statement
Details:
Optimizer: ALL_ROWS
Bytes: 1814
Partition ID: 0
CPU Cost: 127580196
IO Cost: 1936
Temp Space: 0
Elapsed Time: 1
Cost: 1940
Time: 127580KC
Rows: 3 (100.00%)
Union-All
945
122148KC
0 (0.00%)
+Filter+
108
6476KC
0 (0.00%)
View
ALL_ALL_TABLES
837
115671KC
2 (100.00%)
Union-All
847
115749KC
0 (0.00%)
+Filter+
423
60866KC
0 (0.00%)
Filter
424
54882KC
0 (0.00%)
+Hash Join+
2
23KC
1 (50.00%)
Nested Loops
422
54859KC
1 (50.00%)
Merge Join
422
54855KC
1 (50.00%)
Fixed Table
X$KSPPCV (ind:1)
Details:
Bytes: 4
Partition ID: 0
CPU Cost: 3550
IO Cost: 0
Temp Space: 0
Filter Predicates: "KSPPI"."INDX"="KSPPCV"."INDX"
Projection: "KSPPCV"."INDX"[NUMBER,22]
Elapsed Time: 1
Query Block: SEL$E70C16E4
Cost: 0
Time: 3KC
Rows: 1 (50.00%)

Visual Query Builder

Note: This functionality is available only in Enterprise Edition.

Query Builder is a user-friendly visualization tool that will help you make sense of your complex database designs. It can be useful when you need to understand the various relationships between different tables. Also, it can be helpful for those who is not much familiar with SQL scripting or if you don't want to insert script commands manually. The tool creates SQL scripts automatically based on visual schema you create.

*Note: Visual Query Builder presents only in DBeaver Enterprise Edition

Opening Visual Query Builder

To open **Visual Query Builder** click the **Open Query Builder** button in the SQL Editor tool bar.

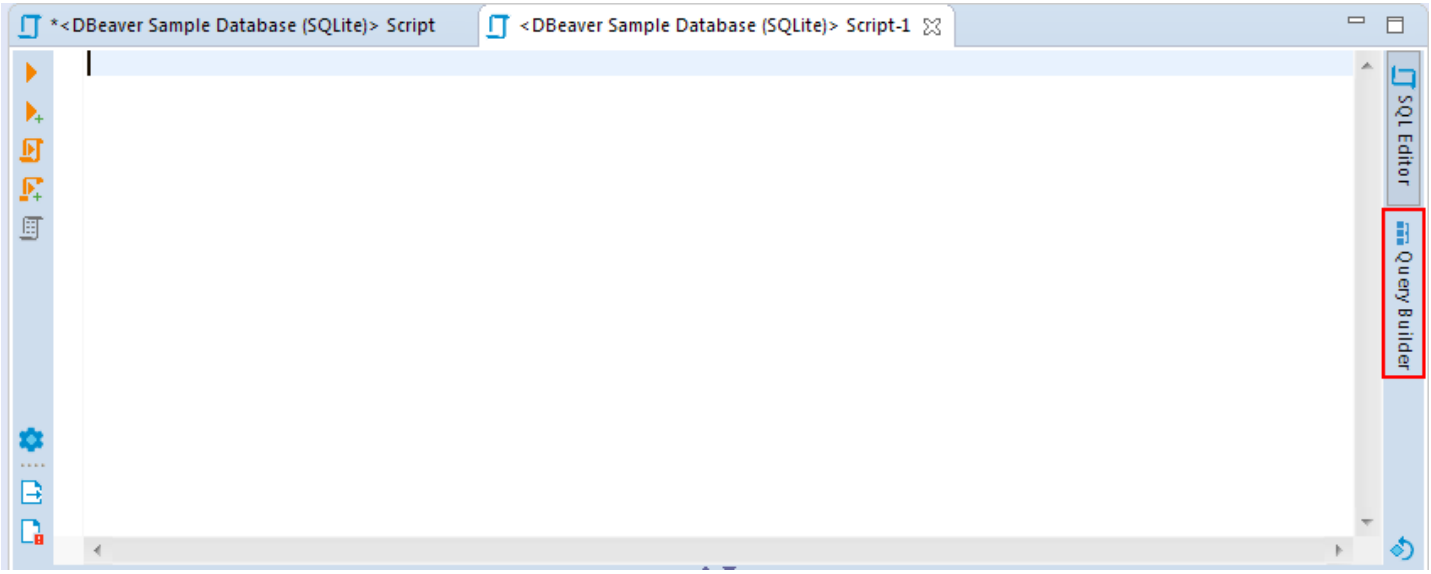

The **Visual Query Builder** will appear on the right.

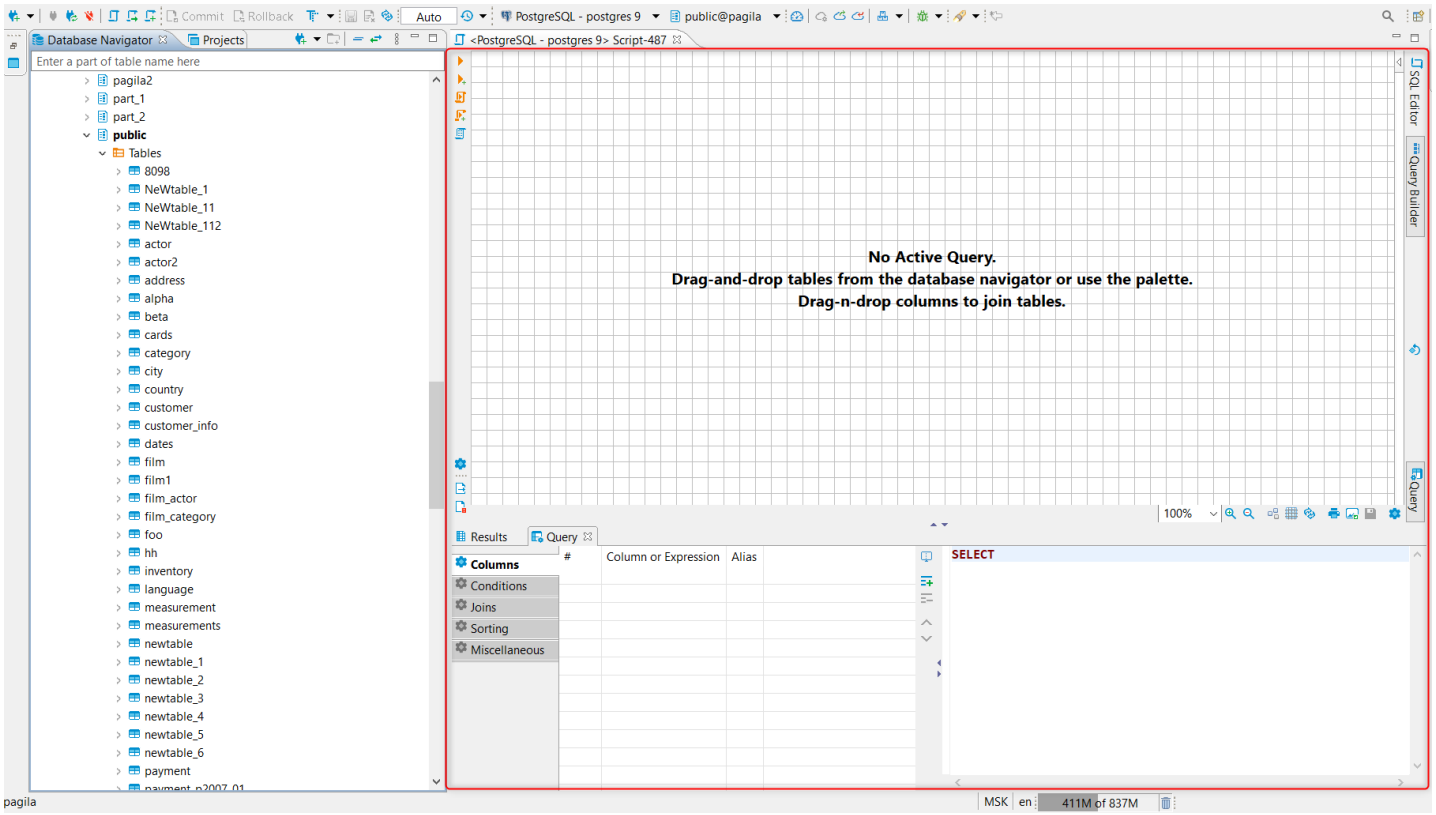

Creating Visual Query

Start creating a query from selecting query data source: drag-and-drop tables you want to work with from the **Database Navigator** pane into the **Visual Query Builder** area. All the connections existing between the tables will be shown automatically.

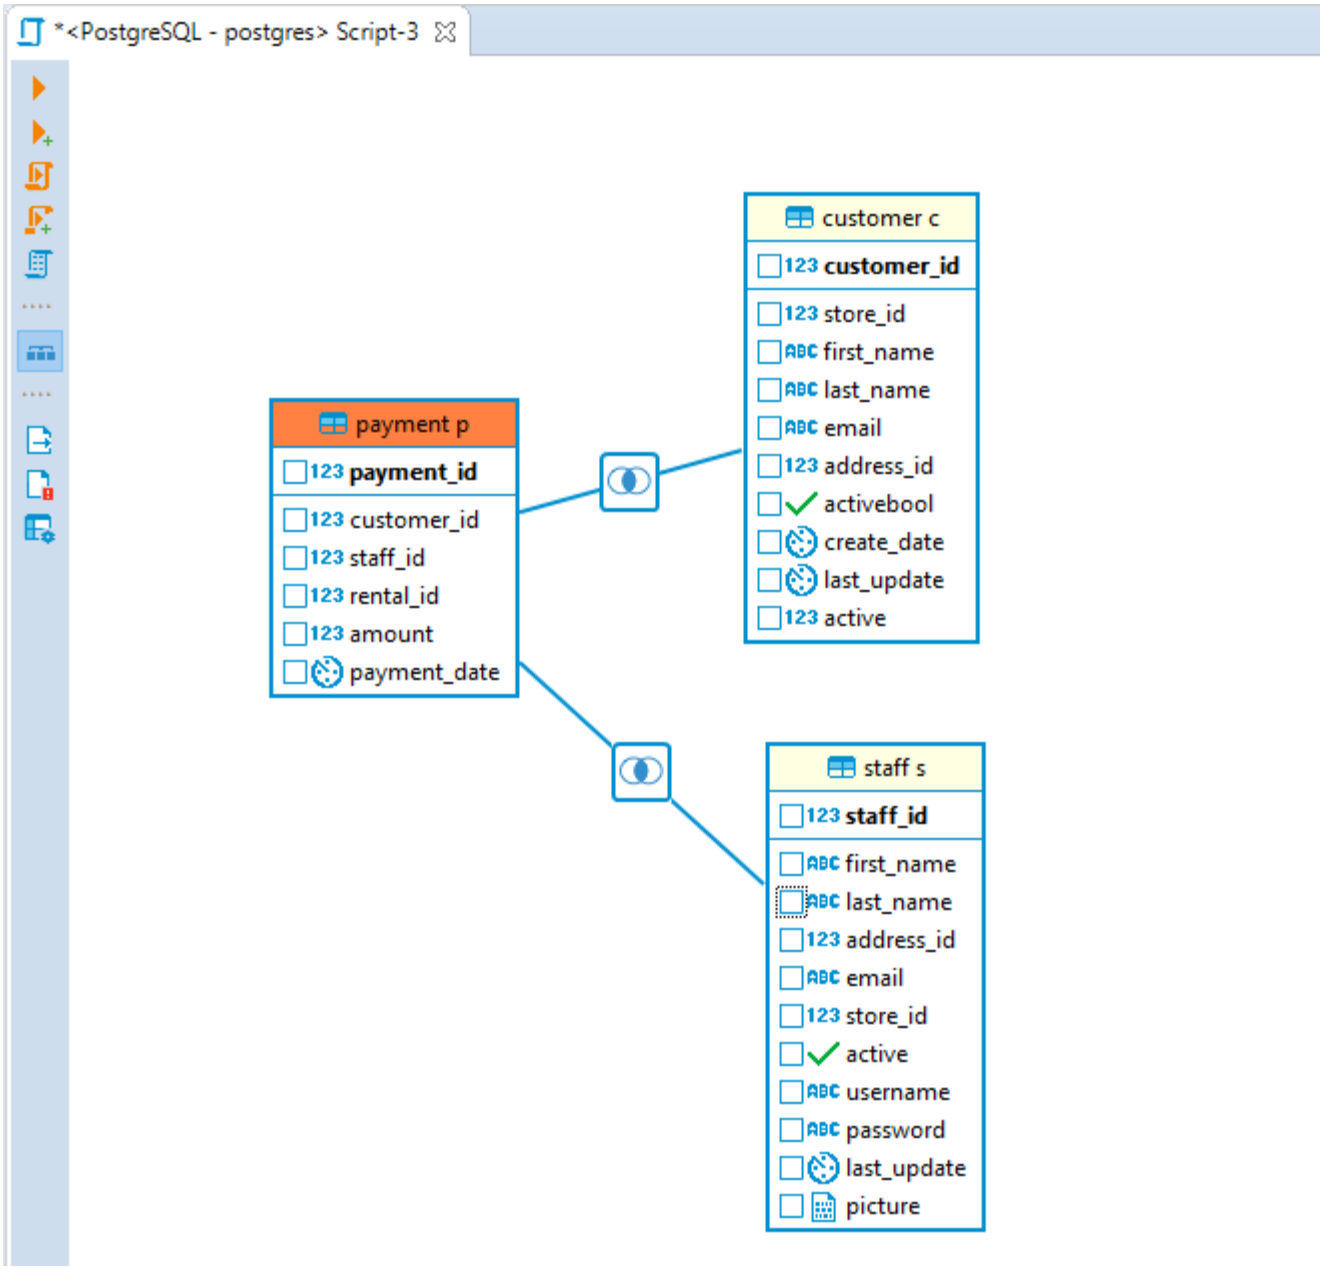

To create a new join between the tables press the left mouse button when the cursor is over the column of one table, holding the right mouse button pressed drag the cursor to the column of another table and release the right mouse button. The connection between the selected columns of the tables will be created visually and in the SQL script a new join will be added.

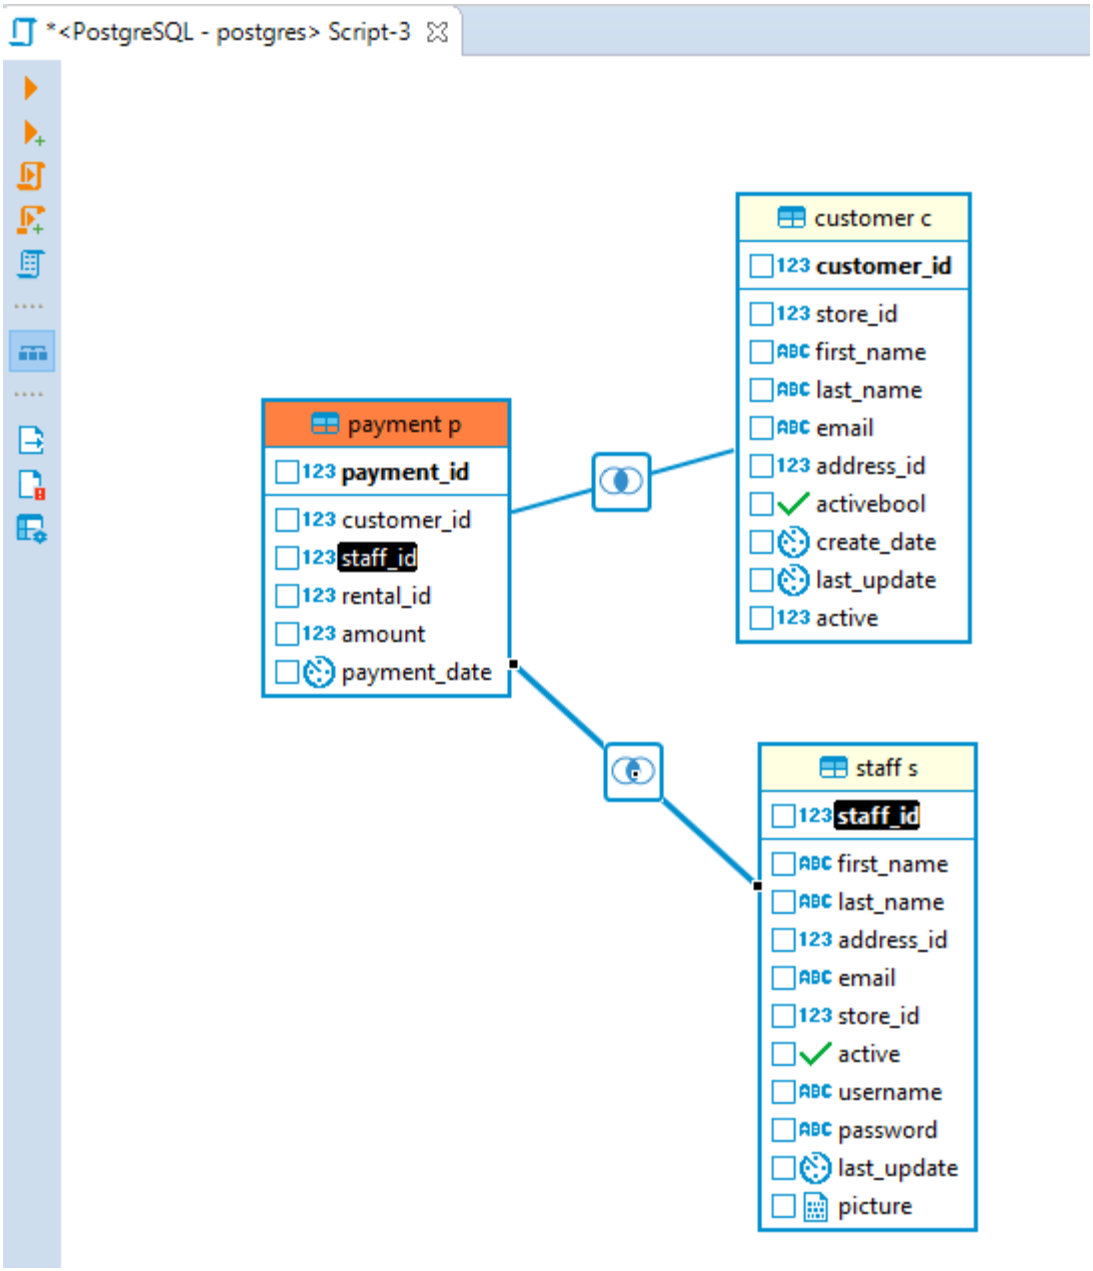

To remove a join between the tables click on it. The connection will be highlighted. Then, press `Delete` or use the **Delete** option in the context menu. The visual connection will be removed and the corresponding join will be automatically removed from the SQL script area.

To build a SELECT query you need to select columns in the tables you added. To select a column click the check-box next to its name - the column will be added to the **Columns** tab of the **Query Settings Editor** and SELECT query will be added to the SQL script area automatically.

Adjusting Query Settings

Visual Query Builder also allows setting query conditions and adjust representation of query results by means of **Query Settings Editor**.

To open **Query Settings Editor** use **Visual builder query settings** button in the vertical tool bar on the left.

Query Settings window contains five tabs described below.

Columns

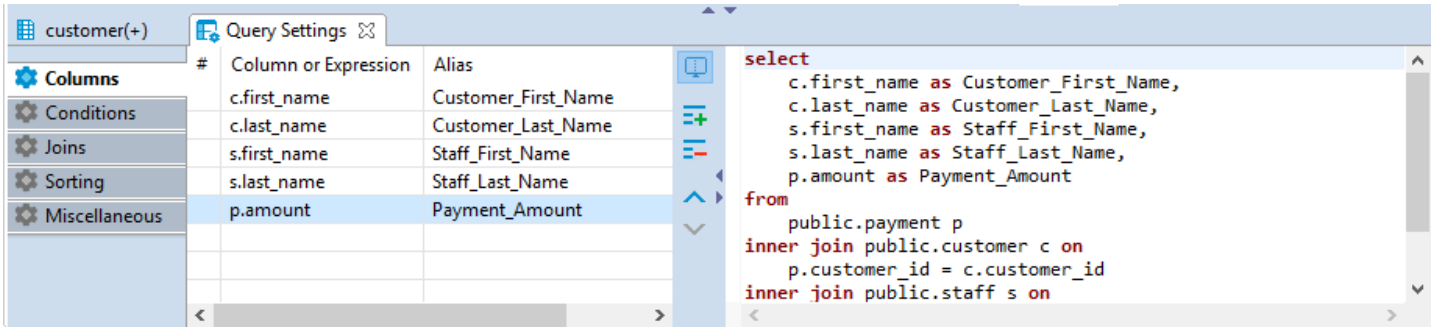

Columns tab of the **Query Settings Editor** contains all the columns you added by selecting column names in **Visual Builder** main window. In this tab you can add and remove columns using **Add** and **Remove** buttons correspondingly.

To add a column, press **Add** button and a new instance will be added to the table. Click on the first cell in **Column or Expression** column and select a column from the list of available columns displayed in the dropdown list appeared.

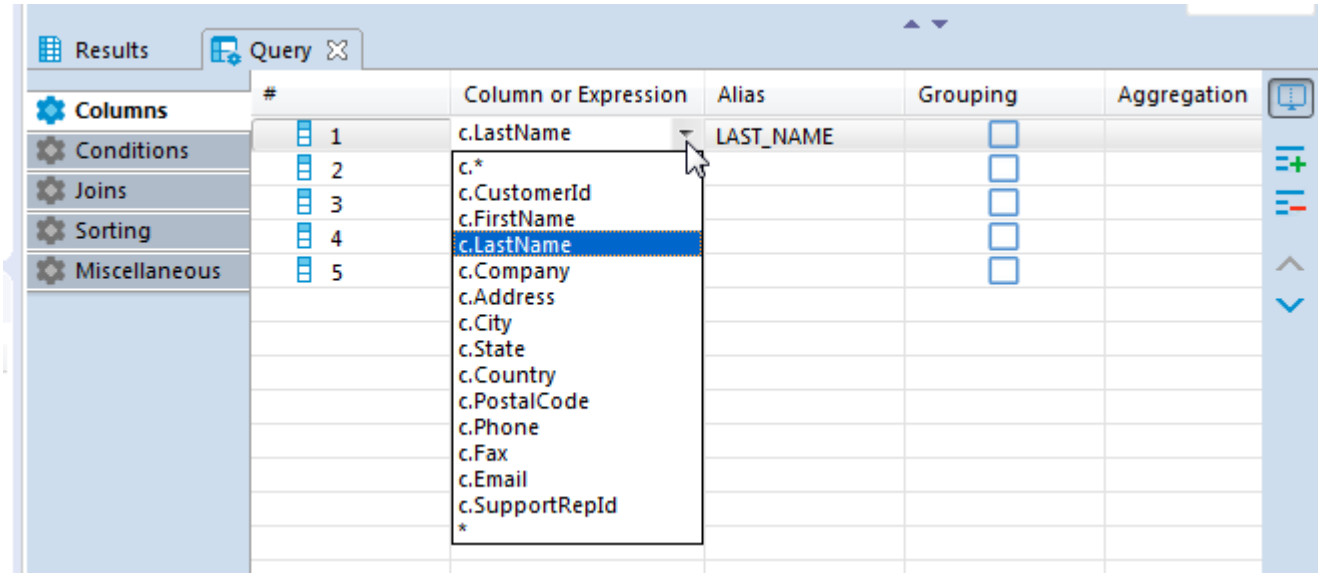

To remove a column, click on the row containing its name and press the **Remove** button on the right.

To change the display order of columns in the result table use **Move Up/Down** buttons .

You can also define a user-friendly name of the column to be displayed in the result table. To set a user-friendly name click on a cell in **Alias** column and insert the name. The change will be immediately displayed in the SQL script area.

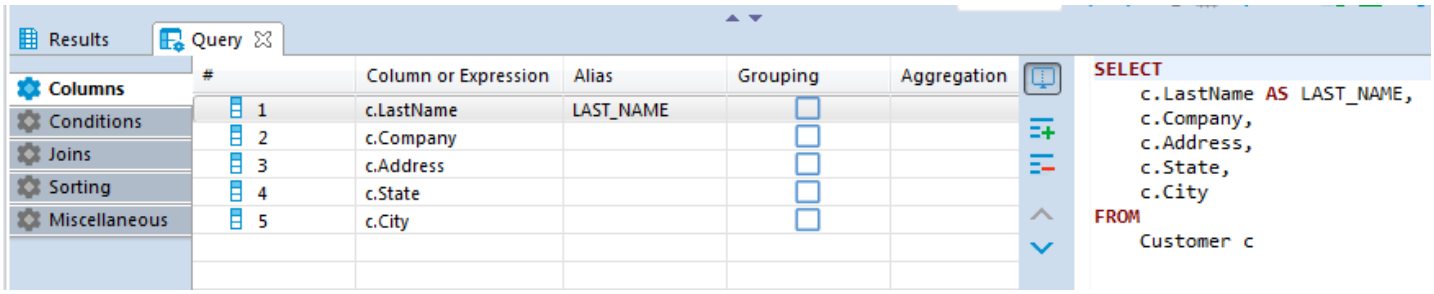

If you want a **grouping** condition to be added to your expression, then you can click on the checkbox in the column row. The expression will update automatically. The other previously selected columns will become aggregate. If there are no other columns, then the expression COUNT(*) will be automatically added.

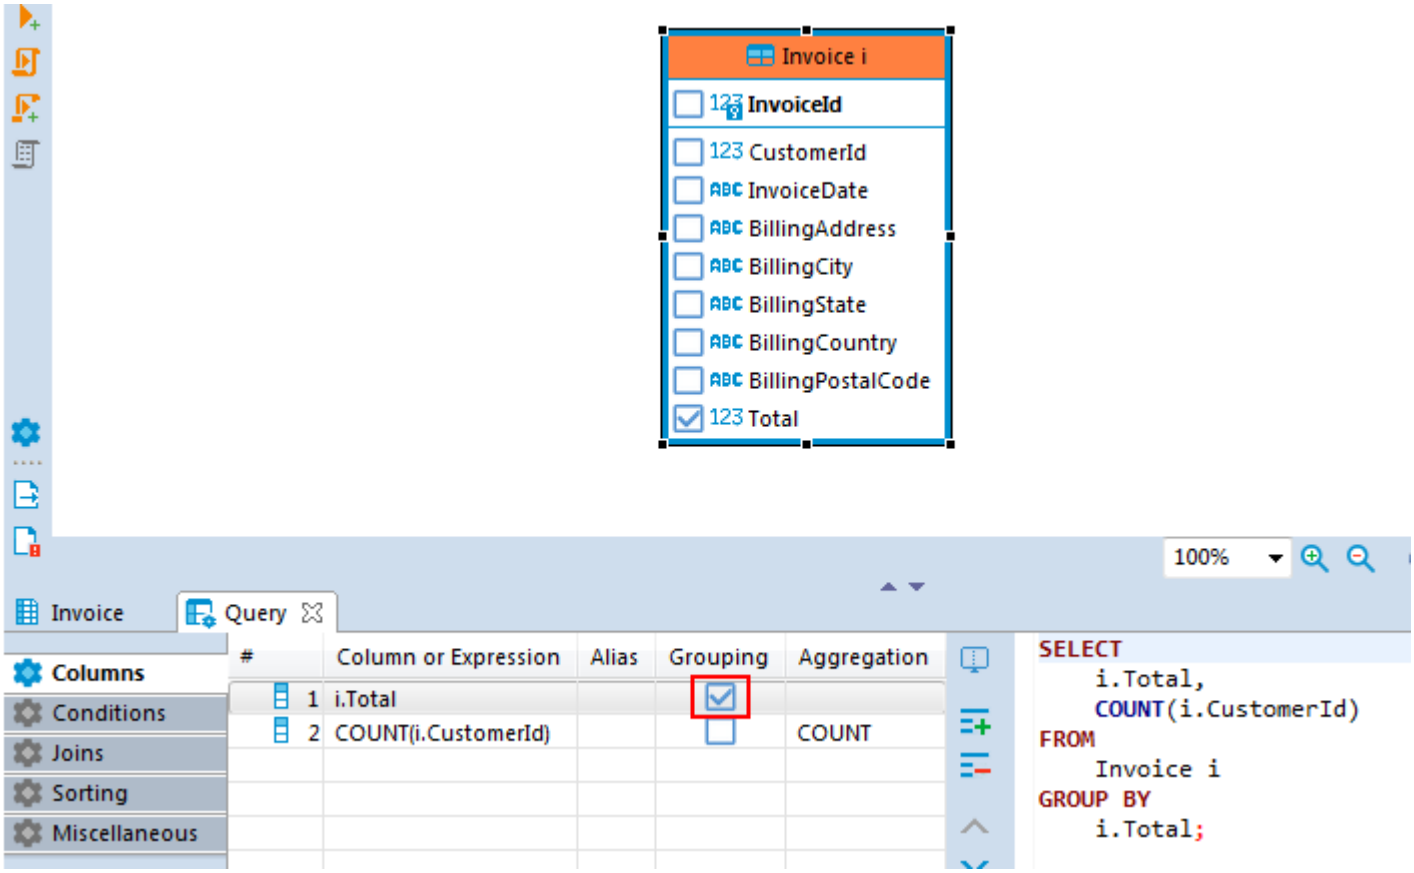

You can select other aggregation functions from the drop-down list. Or enter your own version in the cell.

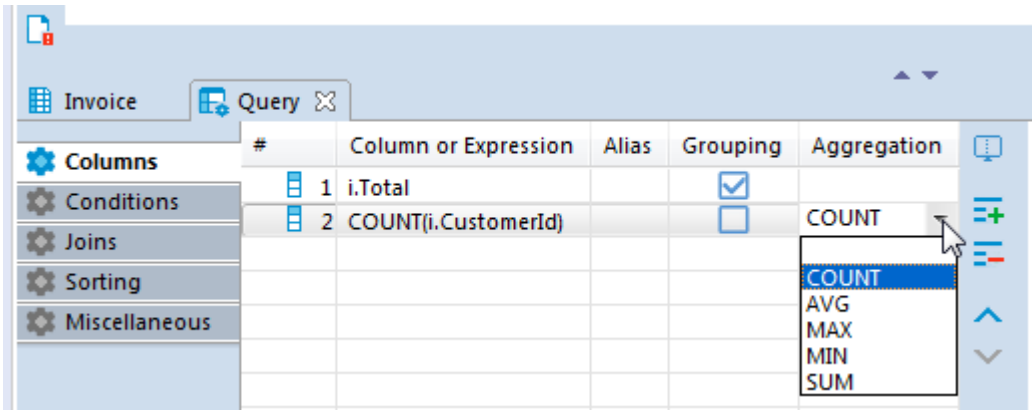

When removing columns from the list, they will be removed from the grouping expression as well. When adding new columns to the list, it is added to the grouping expression.

Conditions

Conditions tab is used for managing query conditional expressions.

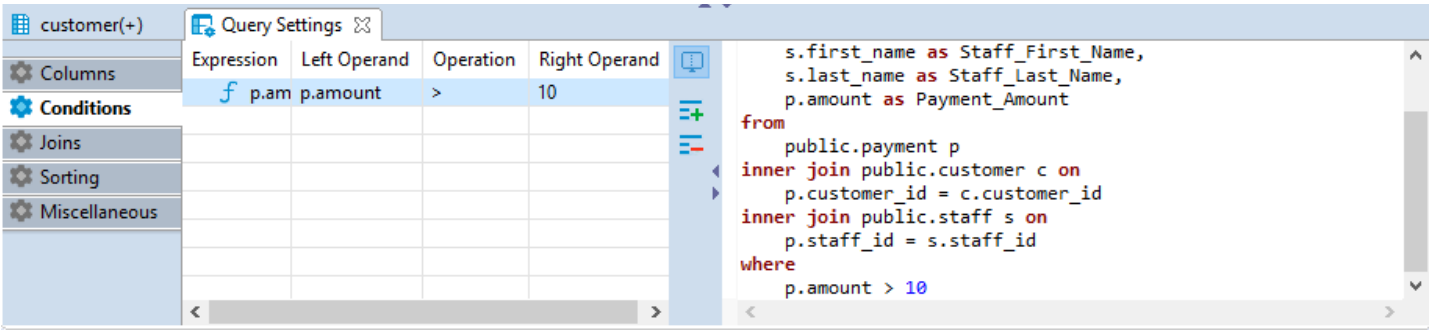

To add a new conditional expression use the **Add** button on the right - a new instance will be added and the default conditional expression WHERE will be added to the SQL script area automatically. This default conditional expression can be then adjusted to the one you need:

- **Left Operand** setting defines the left operand of the conditional expression. To set the left operand, click the cell in the **Left Operand** column and a drop down list of all available columns will be displayed. Select a column you want to use as the left operand in your conditional expression or insert a digit.

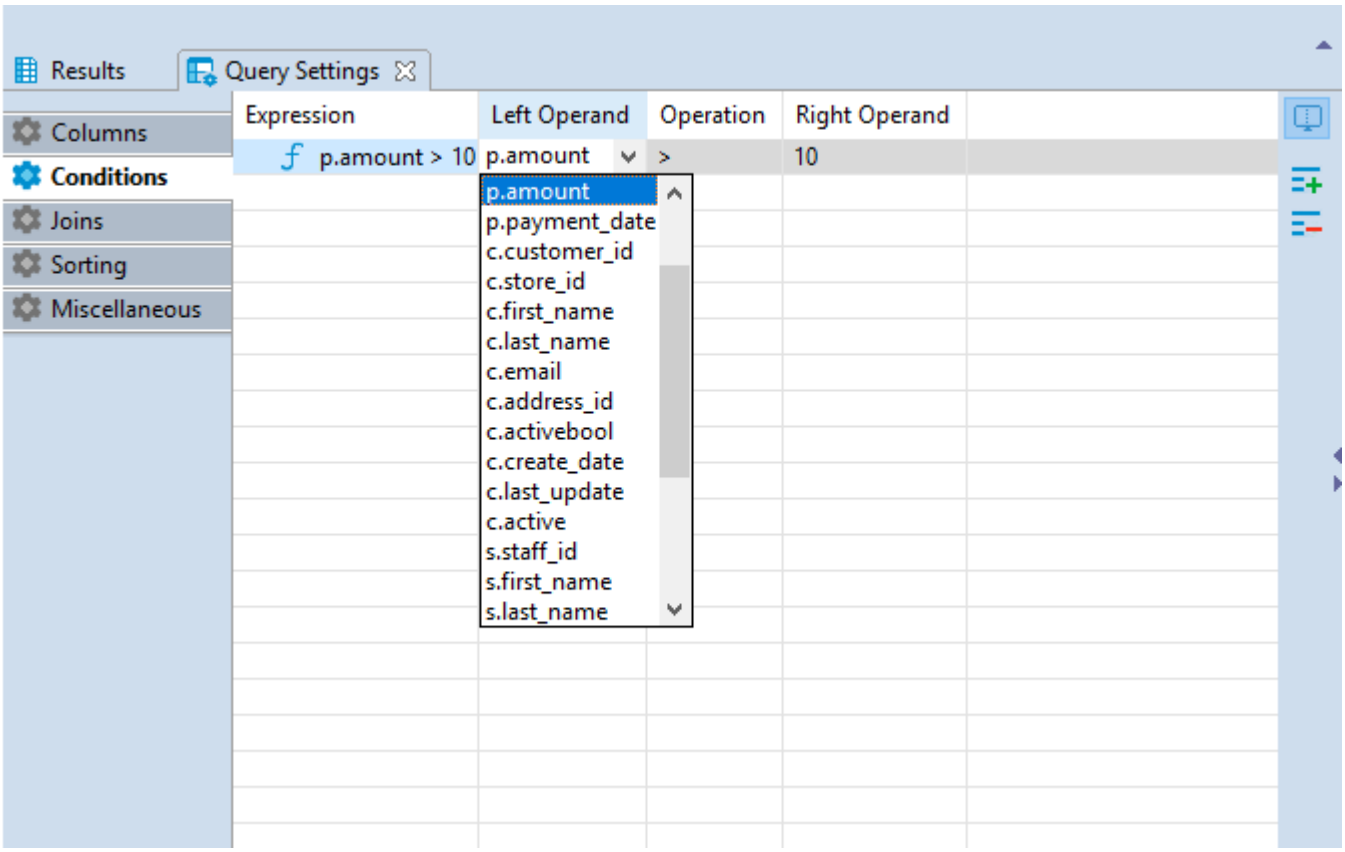

- **Operation** setting defines the comparison rule between the left and the right operands of the conditional expression. To set a comparison rule, click the cell in the **Operation** column and select the rule you need from the drop down list appeared.

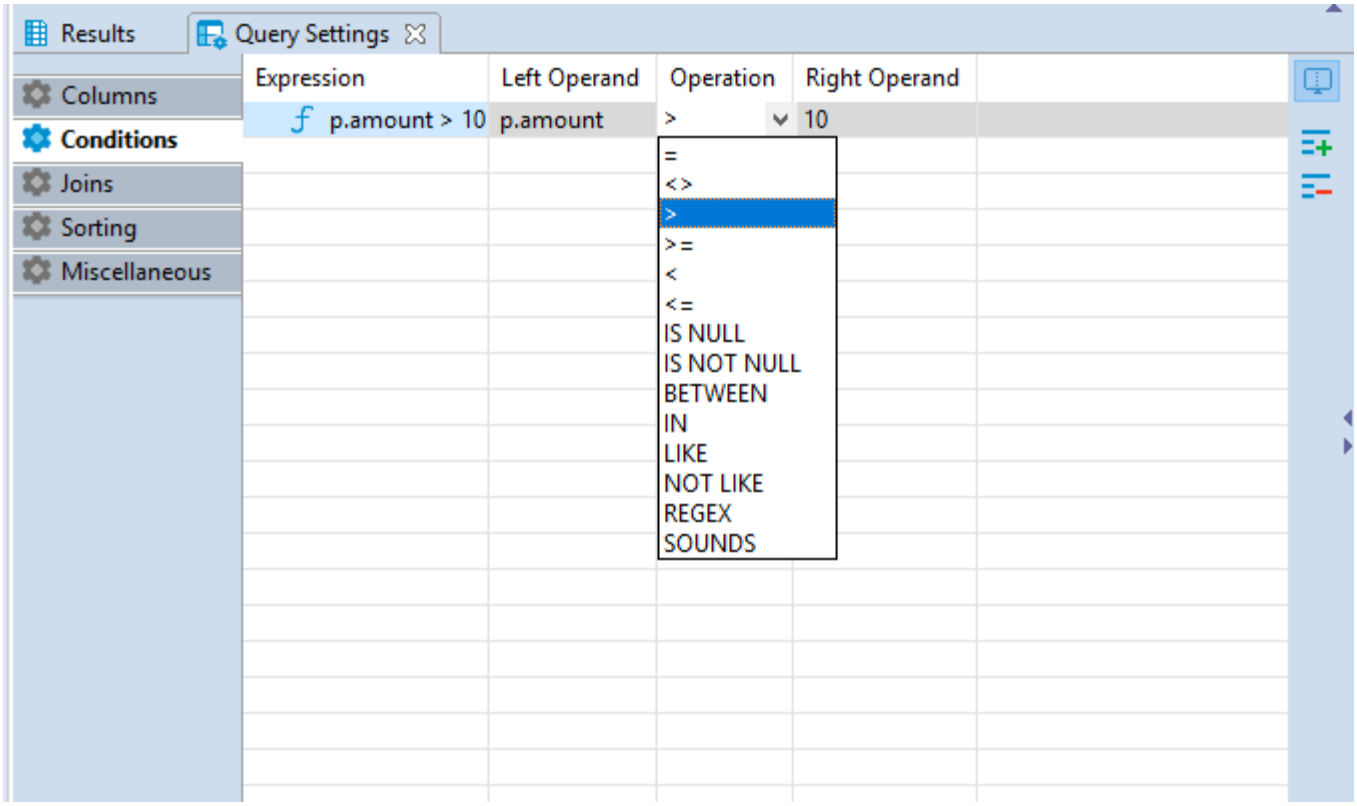

- **Right Operand** setting defines the right operand of the conditional expression. To set the right operand, click the cell in the **Right Operand** column and a drop down list of all available columns will be displayed. Select a column you want to use as the left operand in your conditional expression or insert a digit.

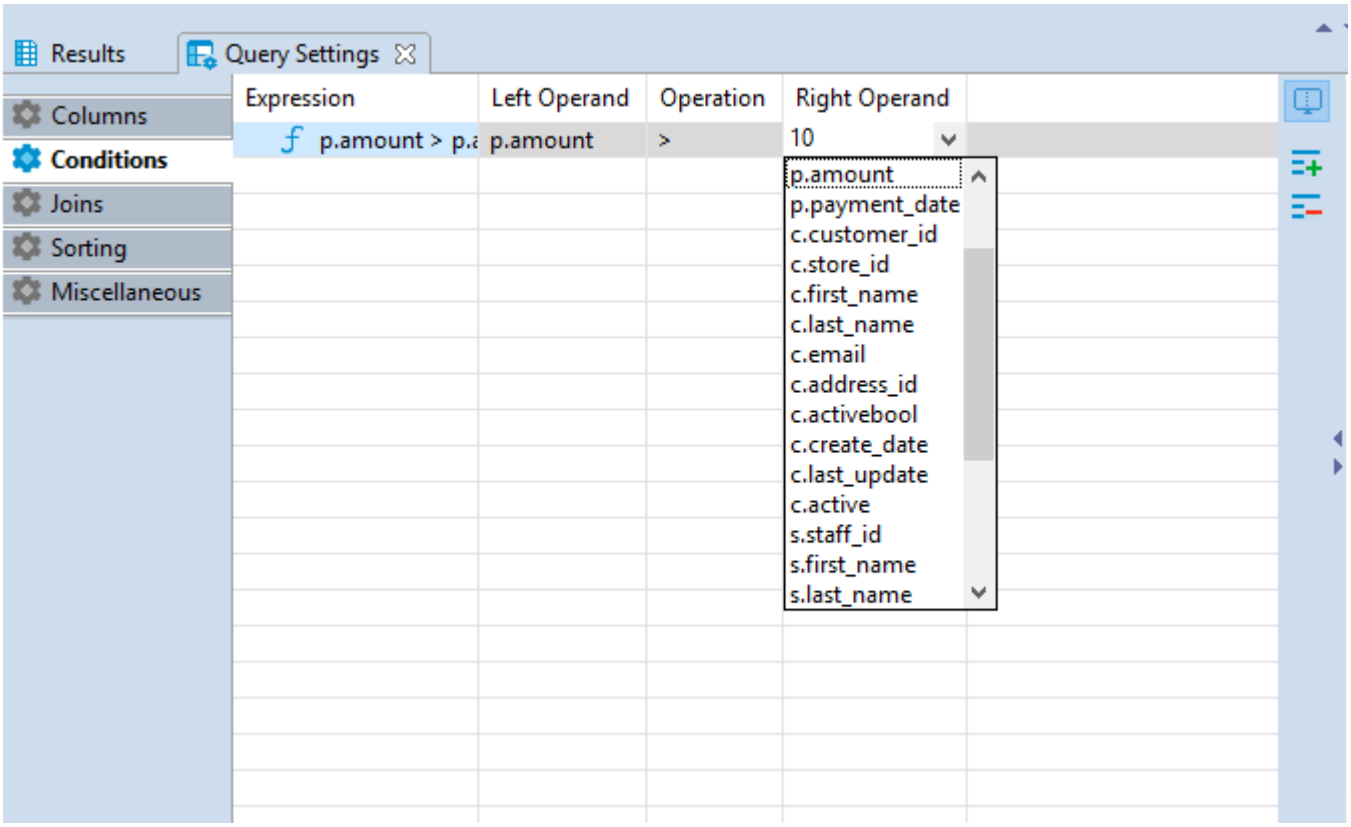

To remove a conditional expression, click on the row containing the expression and press the **Remove** button on the right.

Joins

All the joins existing between the tables in **Visual Query Builder** main window are displayed in the **Joins** tab of **Query Settings Editor**.

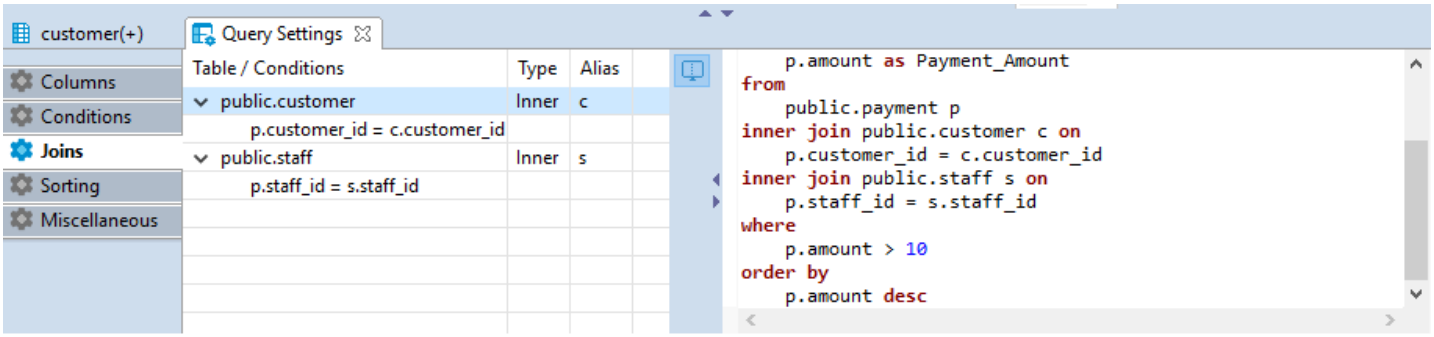

Joins cannot be added or removed by means of **Query Settings Editor**, however, the following join settings can be adjusted here:

- **Type** - defines the type of the join. Click the cell in the **Type** column - a drop down with available join types will be displayed. Select the required option from the list by clicking on it.

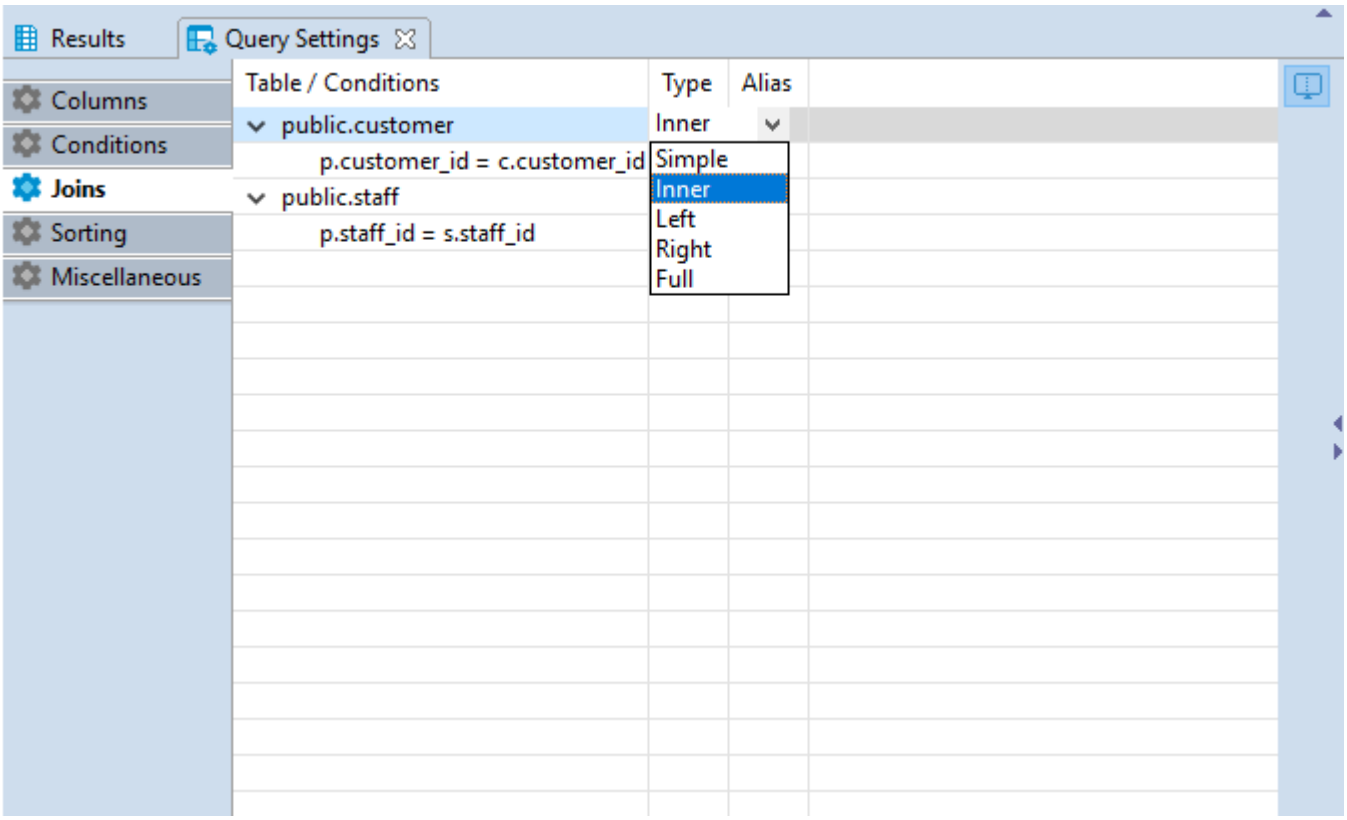

- **Alias** - defines a user friendly name of the join. To define this setting click on the cell in **Alias** column and input the name.

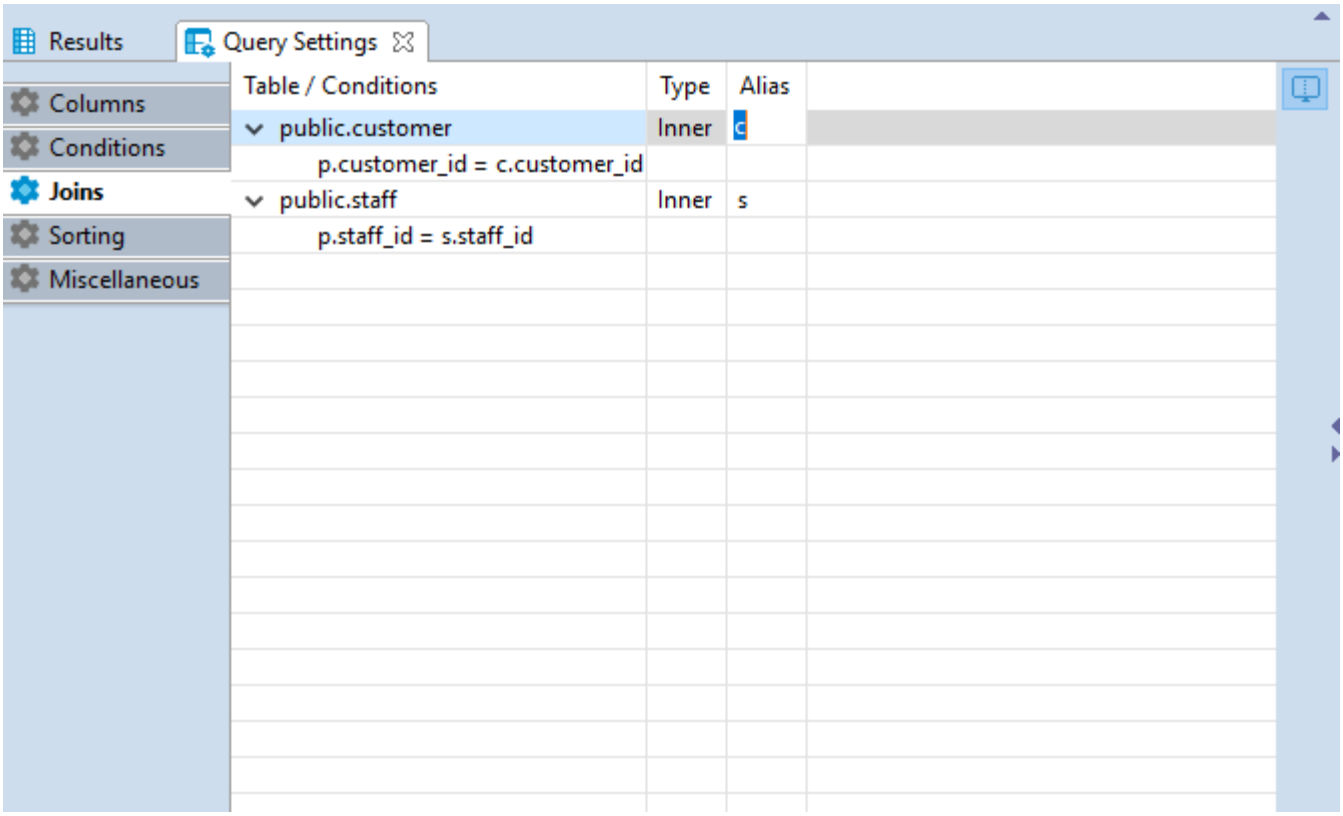

Sorting

In **Sorting** tab you can set the order of rows in the result table.

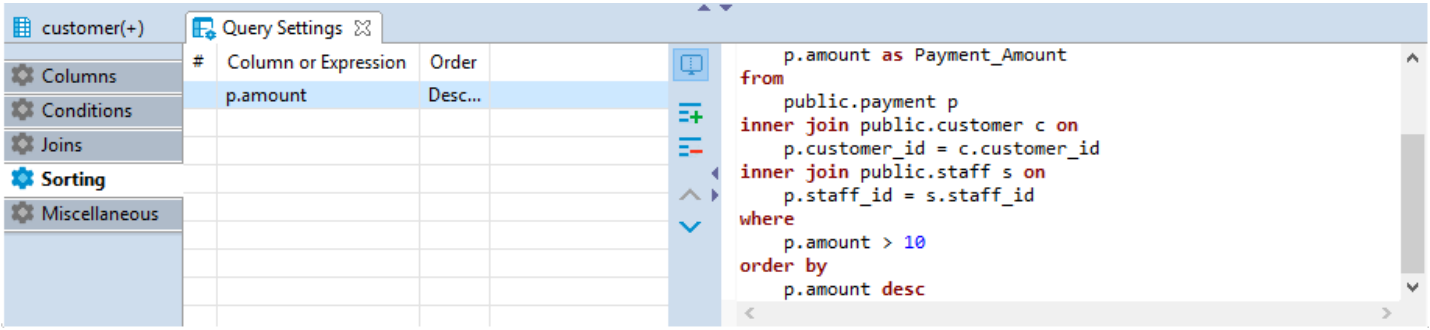

To add a new sorting condition press the **Add** button on the right and the default conditional expression ORDER BY will be added to the SQL script area automatically. This default conditional expression can be then adjusted to the one you need:

- Once a new condition is added, click the first cell in **Conditions or Expressions** column and a drop down list of all available columns will appear. Select the required column by clicking on its name.

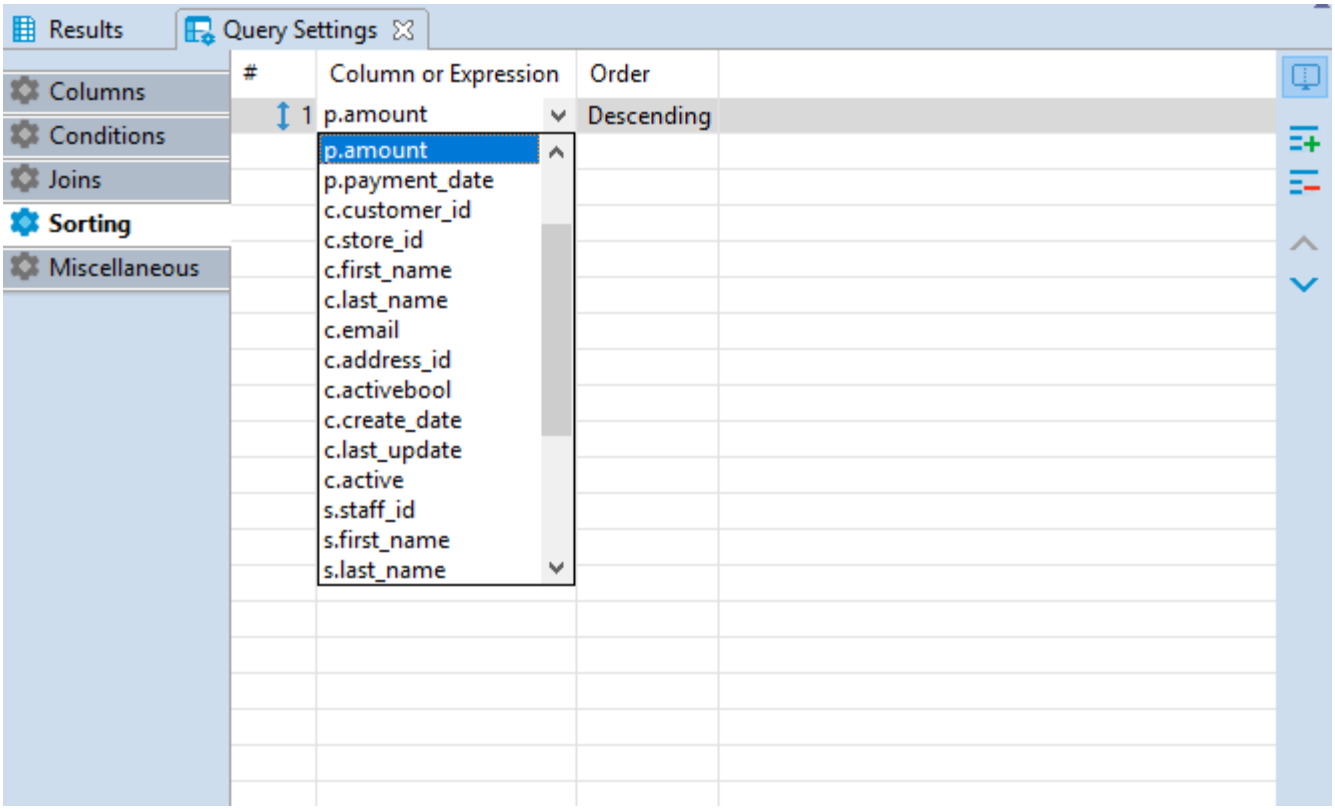

- In **Order** column you can define whether the rows of the selected column should be sorted in ascending or descending order in the result table. To set the order, click the cell in **Order** column and select the required option from.The order by command will be added to the script.

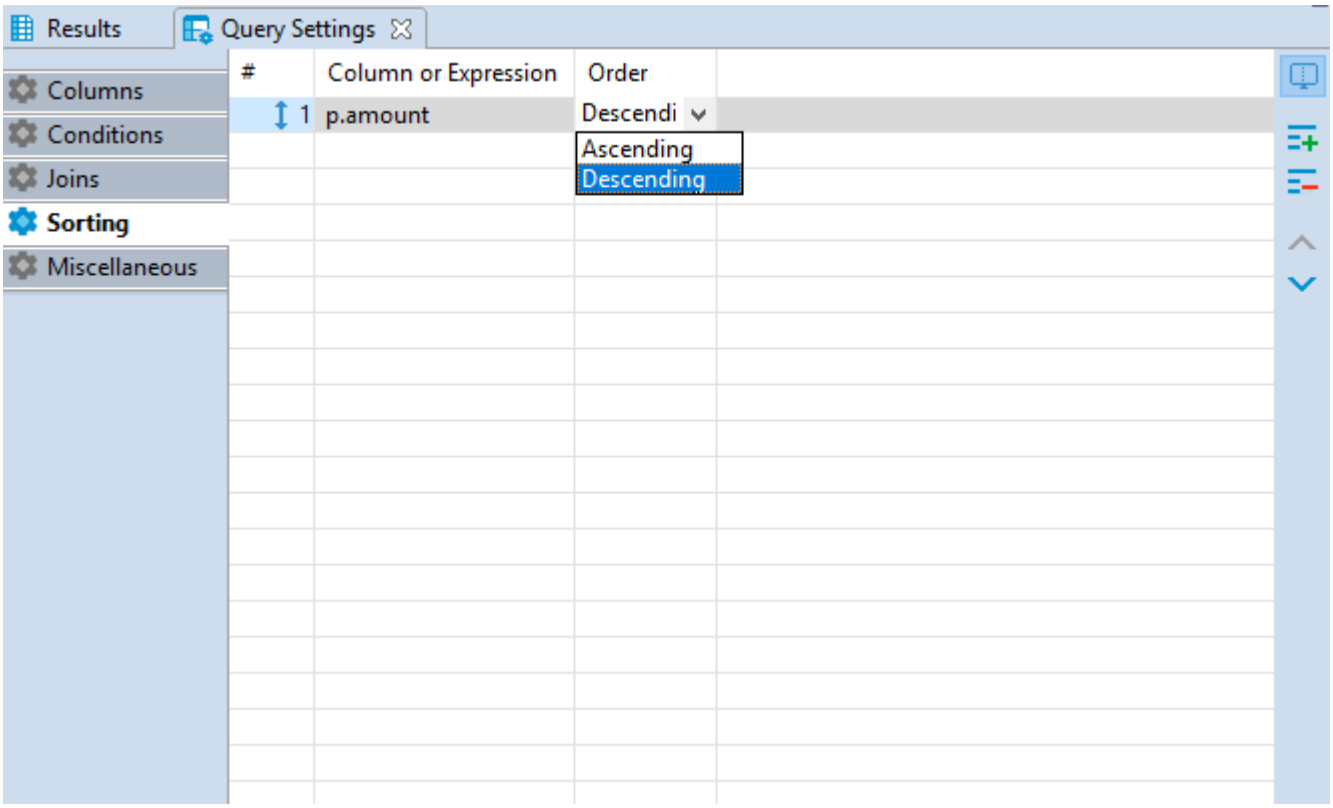

To remove a condition use the **Remove** button on the right.

Miscellaneous

In **Miscellaneous** tab it is possible to:

- Enable or disable the automatic generation of aliases for tables by selecting the **Add table aliases** check-box.
- Disable auto-completion for table names by selecting the **Use fully qualified table names** check-box.
- Enable Autosave on SQL-editor switch by selecting the **Autosave on SQL-editor switch** check-box.

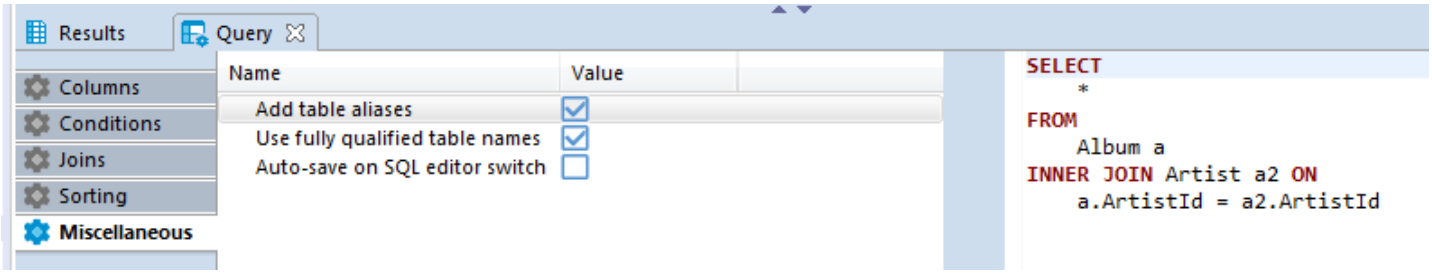

Executing Visual Query

To execute a query, use **Execute SQL statement** button to get the results in the same tab or **Execute SQL statement in new tab** button to get the results in a new tab. Both buttons are located in the **Visual Query Builder** vertical toolbar.

Script Management

Saving Scripts

You can save scripts to a predefined space in the currently active project or somewhere in the file system.

To save a script to the current project space, just press `Ctrl+S` or right-click the script and click **Save** on the context menu:

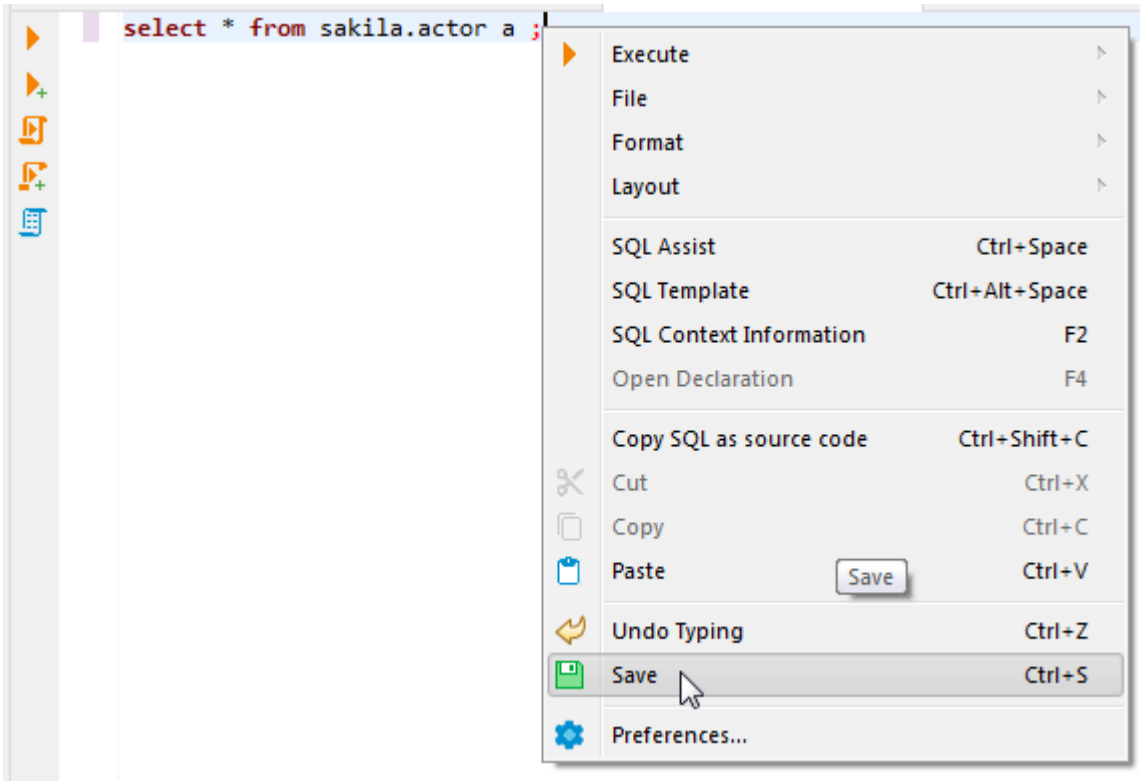

You can find the script saved this way in the Project Explorer view in the **Scripts** folder.

To save a script to the file system, right-click the script, click **File -> Export SQL script** on the context menu and then select the folder in the file system. You can also, click **SQL Editor -> Export SQL script** on the main menu:

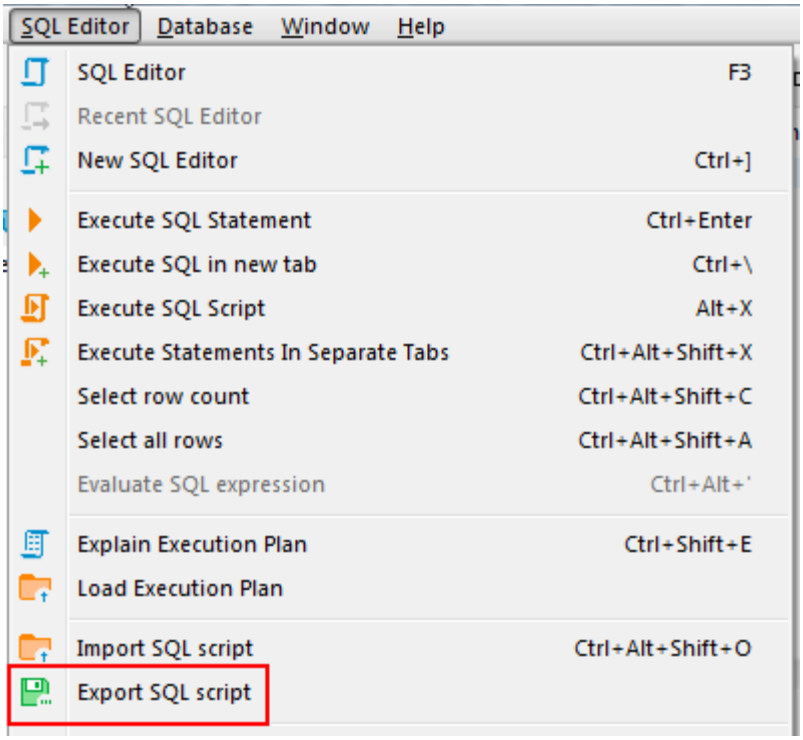

Unsaved data are highlighted in color on the left side of the editor, as well as an asterisk in the name of the script.

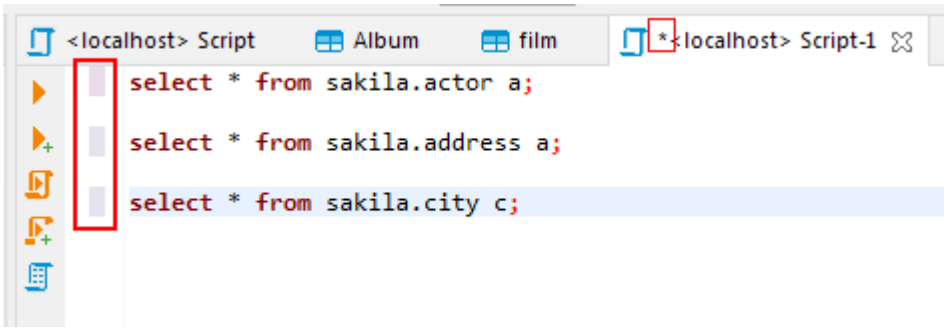

Loading Scripts

To load a script stored in the file system to the SQL Editor, press `CTRL+SHIFT+O`, or click **SQL Editor -> Import SQL script** on the main menu, or right-click the script panel and click **File -> Import SQL script** on the context menu:

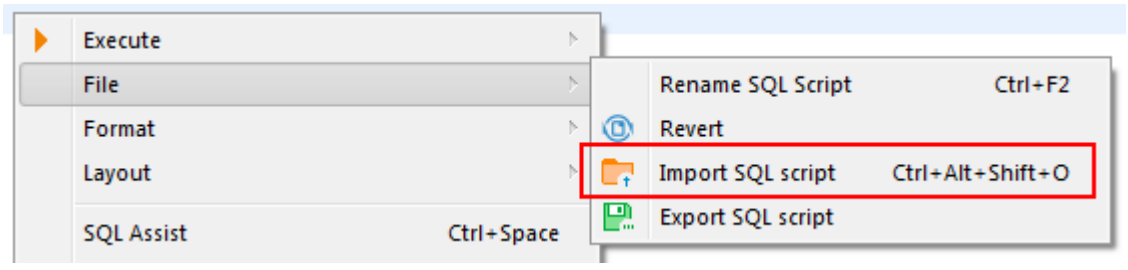

Renaming Scripts

To rename a script, right-click anywhere in the script panel, click **File -> Rename SQL Script** on the context menu or press `CTRL+F2`:

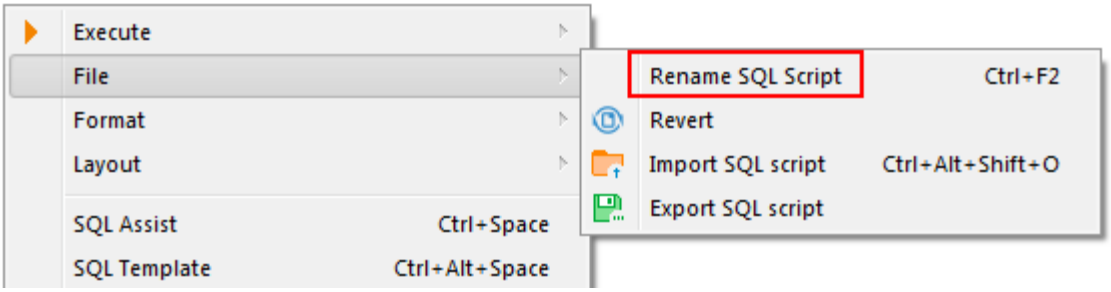

Then enter the new name in the Rename SQL script dialog box and click **OK**:

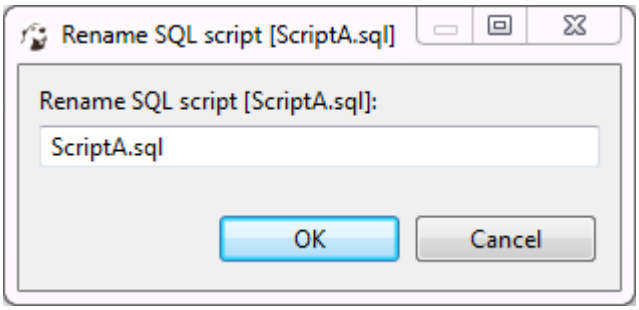

Reverting Changes

If you want to revert all changes made to the current SQL script and return it to its initial state (reload from disk), right-click anywhere in the script panel and click **File -> Revert** on the context menu.

SQL Console

In some cases you might want to execute a query and do not save it in a script. For example when you read table data using "Read data in SQL console" or open procedure/function source from DDL editor. SQL console doesn't has an associated .sql file. Its contents will be lost when you close it.

Client Side Commands

DBeaver supports the following commands:

Command	Database	Description
@set var = value	All	Sets a script variable. You can use expressions as a value. Variables can be used as SQL queries input parameters.
@unset var	All	Unsets a script variable.
@echo message	All	Prints message to output log. You can use a macro in message (for example ${var}).
@include fileName	All	- Executes a specified file name, - Can be used in scripts, - Opens a new SQL console with the specified file and processes SQL queries as in a regular SQL editor.
source fileName	MySQL	The same as @include but in MySQL CLI syntax
define var = value	Exasol	The same as @set but in Exasol EXAPlus syntax.

PostgreSQL Debugger

Prerequisites for Debugging

To enable interactive debugging of PL/SQL procedures on a Postgres server, you need to use _plugin*debugger*. The _plugin*debugger* is a typical interactive debugger delivered as an extension. It requires a shared library preload in Postgres to operate the shared_preload_libraries parameter in settings. The debugger is developed and maintained by EnterpriseDB. Its source code is available for examination and improvement.

The debugger provides the required server API for debugging PL/SQL procedures with:

- Breakpoint management;
- Step-by-step tracing;
- Variable acquisition and management.

Installation

PostgreSQL 12 on Ubuntu-based distros

If you happen to have a PostgreSQL 12 installed via *apt*, then the procedure is quite straightforward:

```
sudo apt install postgresql-12-pldebugger
sudo service postgresql restart
```

After that, run the following command in the database or databases that you wish to debug functions in:

```
CREATE EXTENSION pldbgapi;
```

Installation from source code

You can find the source code in this repository. Installation instructions are located in the README file.

Running debugger in DBeaver interface

Open the source code of the function you want to debug. To toggle breakpoints, place the caret on the line you want the function to be stopped at and use shortcut `Shift + Control + B`. Alternatively, you can toggle breakpoint with your mouse by clicking on a ruler, as demonstrated in the screenshot below:

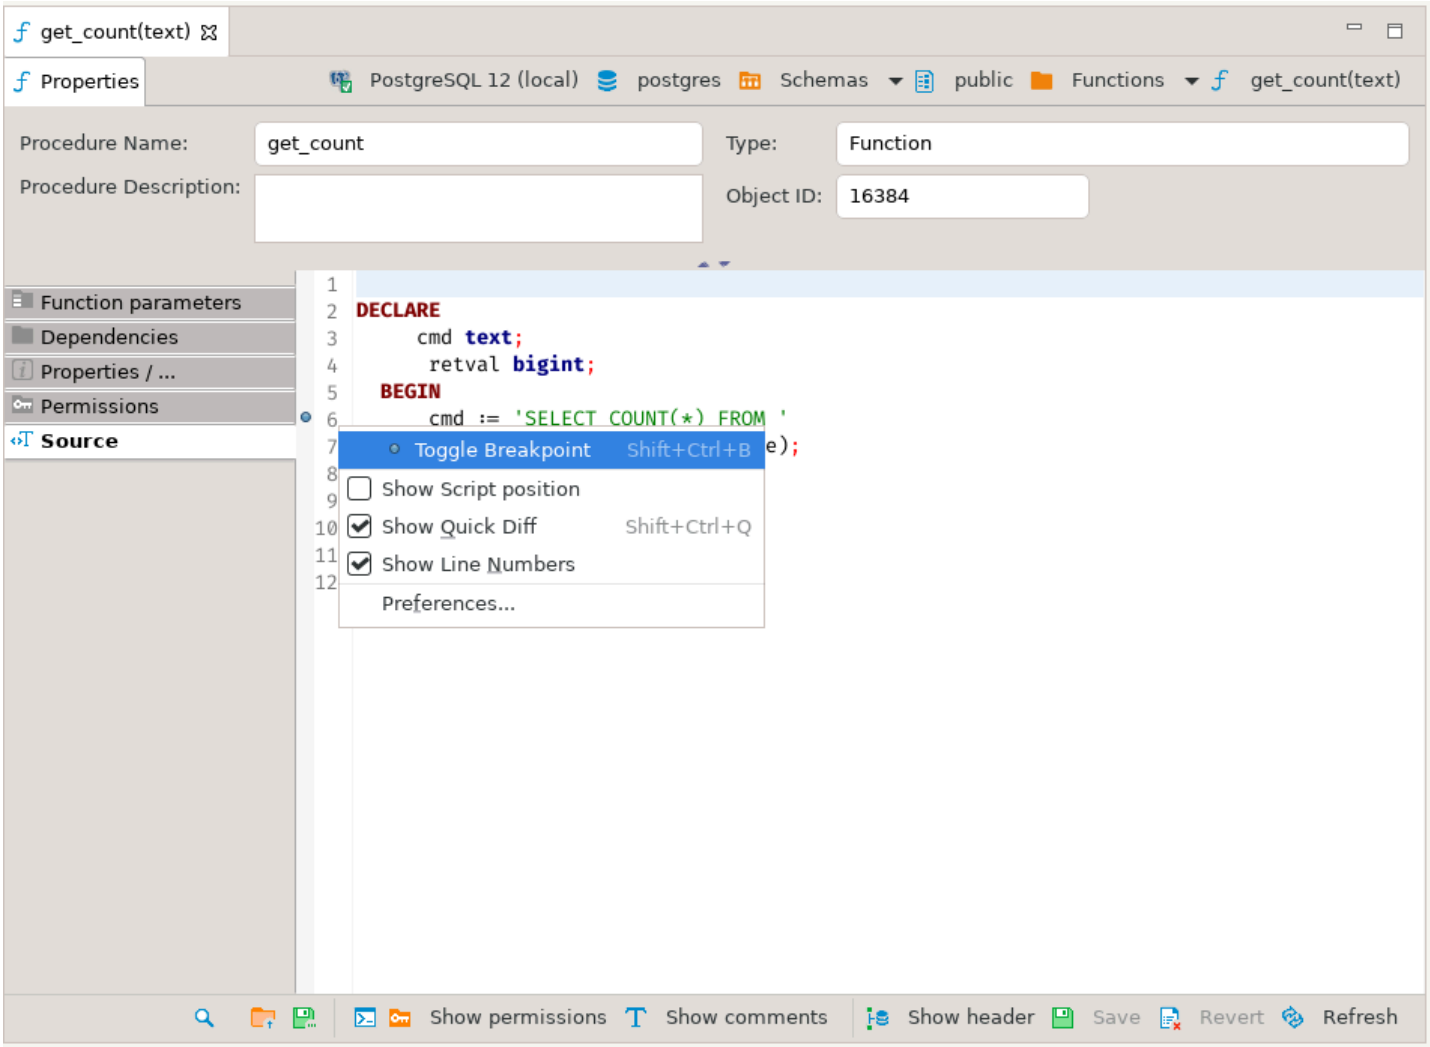

NB. You must toggle breakpoints only when *Show header* option is **not** on.

Then you need to set up a debugging configuration. Locate the downward-facing arrow right to the bug icon, click on it, then

Debug As -> Database Debug:

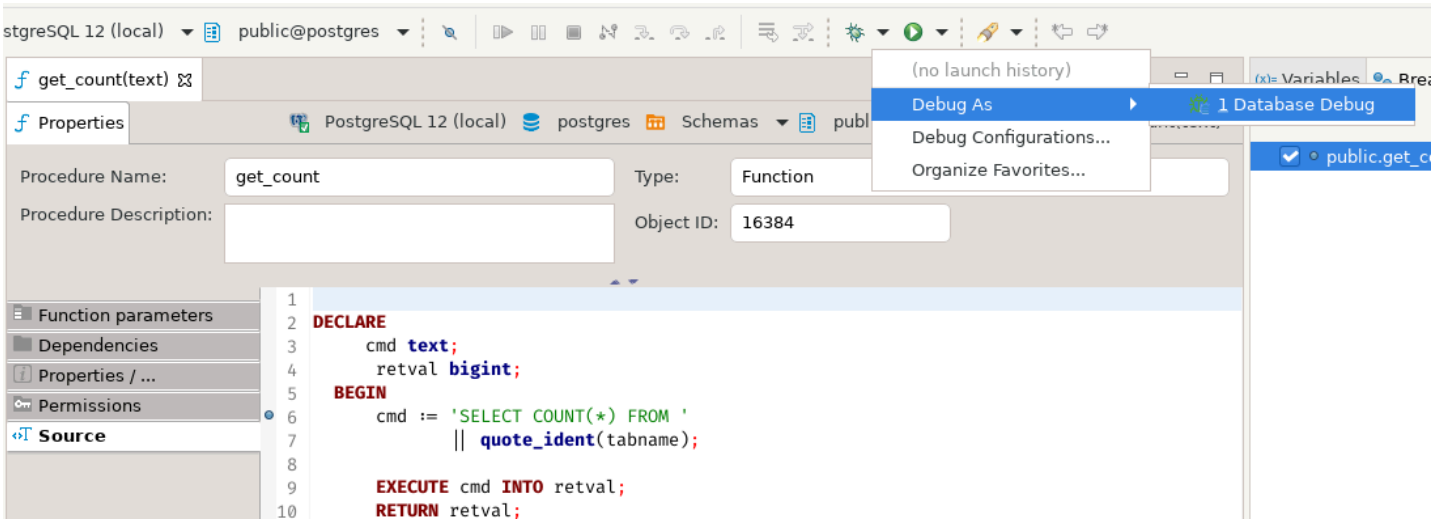

The *Edit Configuration* dialog opens. Set up input values in the table *Function parameters*.

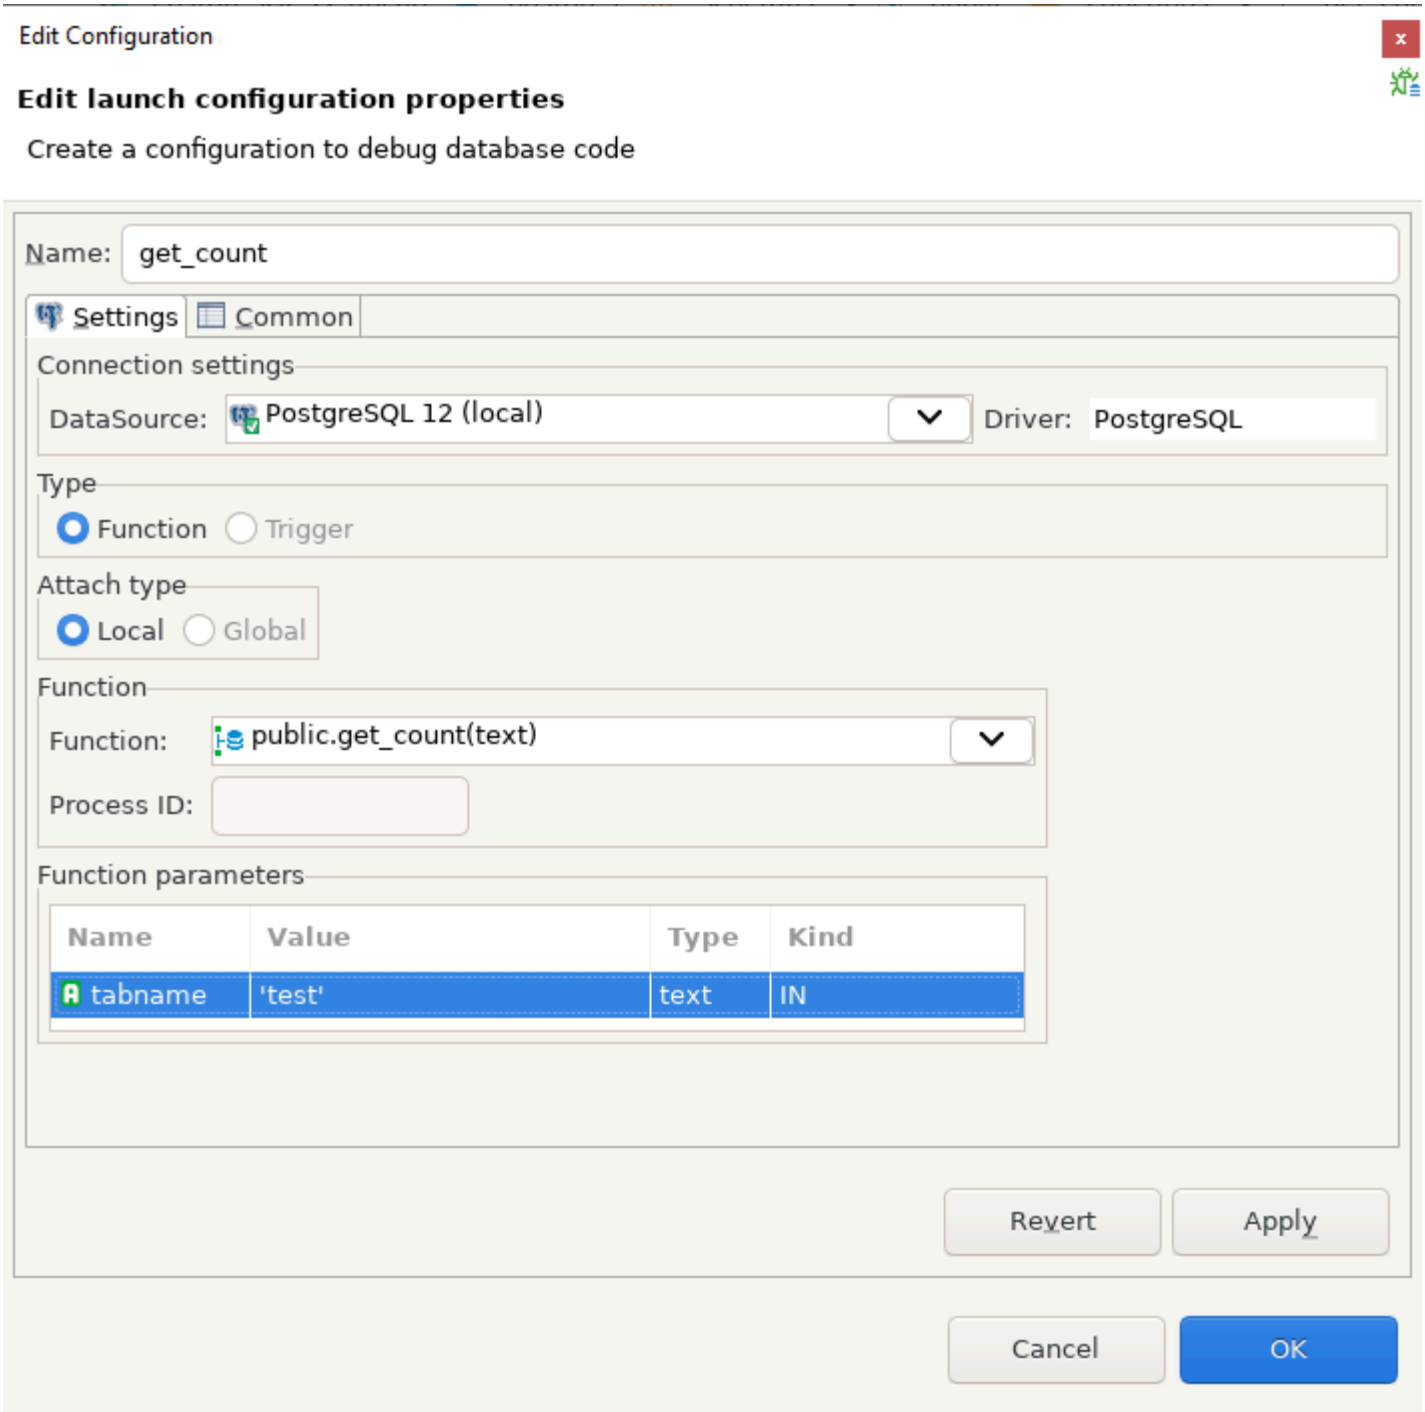

Click on *OK* button, and you are ready to go!

The usual buttons essential for debugging such as *Step Over* and *Continue* are located here:

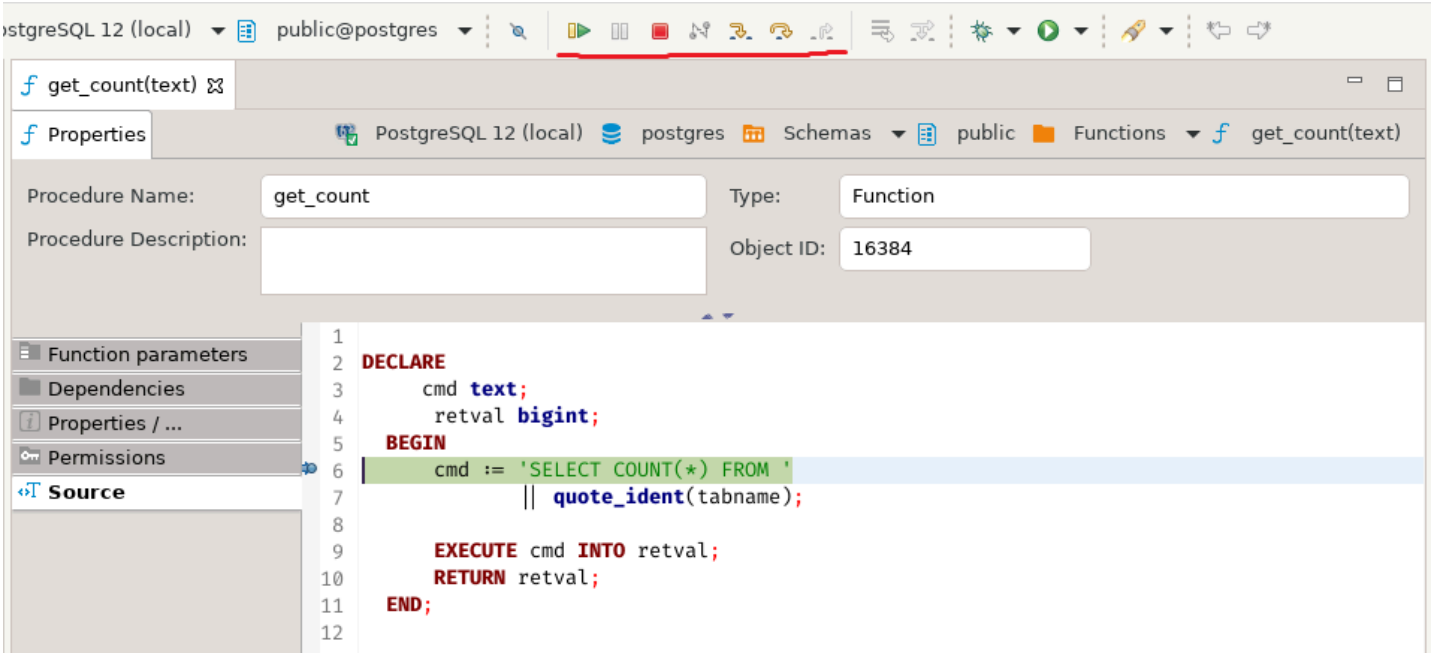
PostgreSQL 12 (local)
public@postgres
get_count(text)
Properties
PostgreSQL 12 (local)
postgres
Schemas
public
Functions
get_count(text)
Procedure Name:
get_count
Type:
Function
Procedure Description:
Object ID:
16384
Function parameters
Dependencies
Properties / ...
Permissions
Source
DECLARE
cmd text;
retval bigint;
BEGIN
cmd := 'SELECT COUNT(*) FROM '
|| quote_ident(tabname);
EXECUTE cmd INTO retval;
RETURN retval;
END;

ER Diagrams

ER diagrams appear on the rightmost tab of the Database Object Editor:

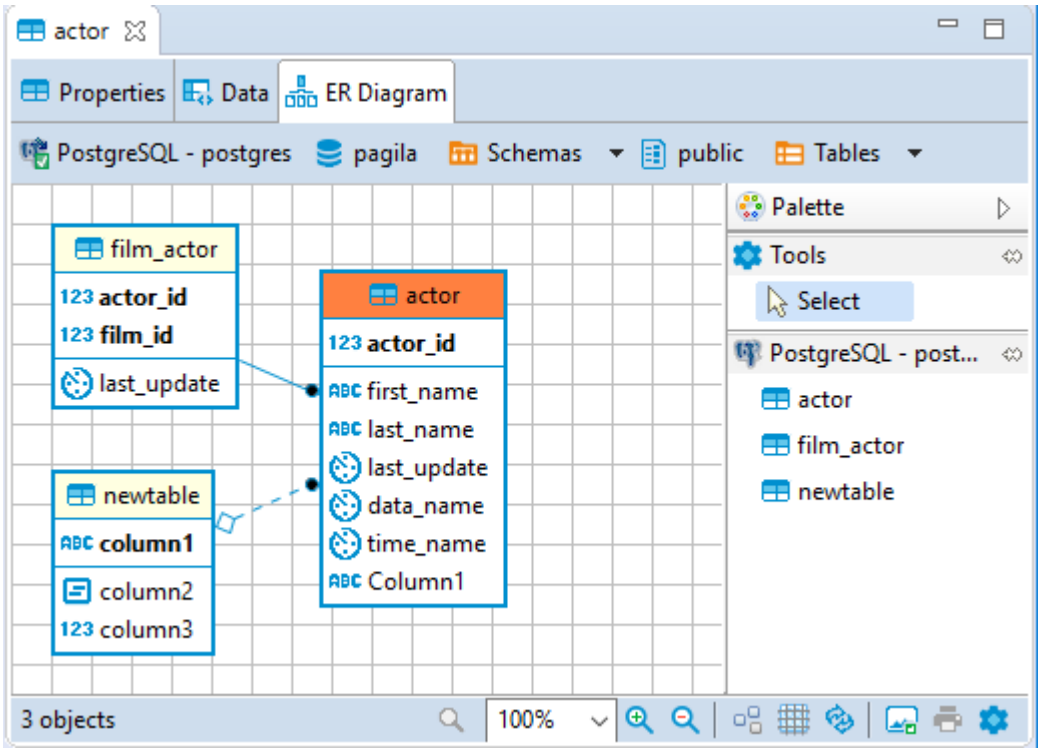

Entity Relation Diagrams (ERD) are graphic presentations of database entities and relations between them. DBeaver allows viewing diagrams of existing tables and whole database schemas, see Database Structure Diagrams, as well as create custom diagrams, see Custom Diagrams.
By default DBeaver uses IDEF1X notation.

Both types of diagrams provide the same tools to adjust their view and structure. They can be printed and exported to image file formats.

Selection of Elements in Diagrams

You can use one of the two tools to select elements in diagrams:

- Select – supports both, single and multi-select modes. To select a single element (table, connection, entity inside a table) in a diagram, just click that element. To select multiple elements, similar to using the Marquee tool, click outside the first element and draw until all elements you need are in focus:

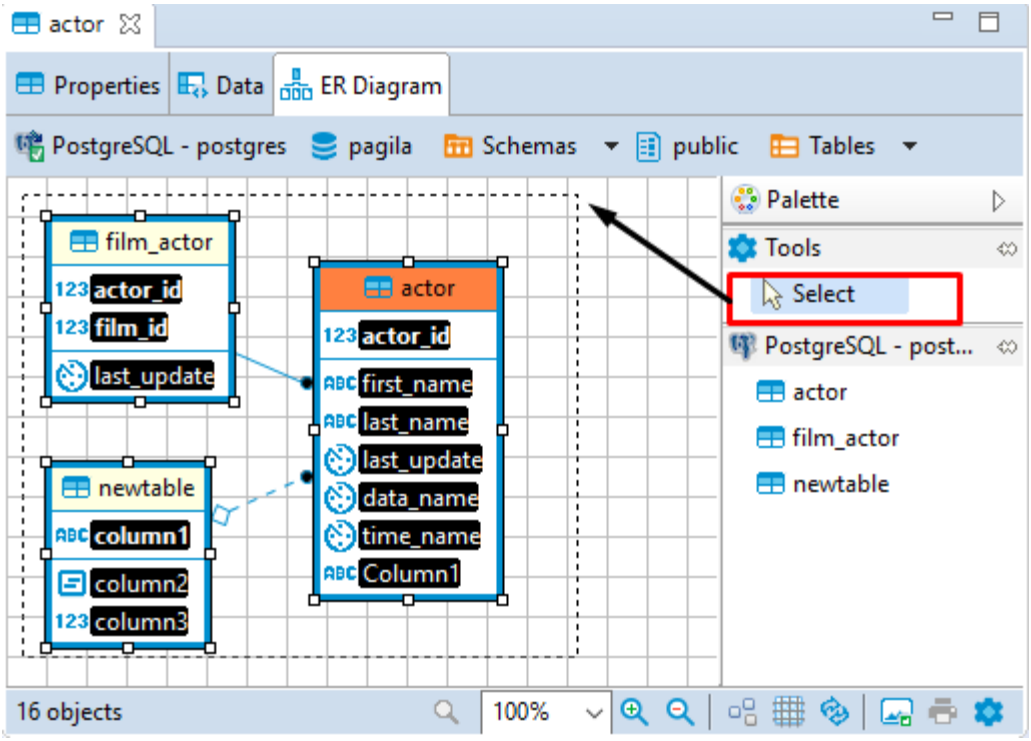

Structure Adjustment

NOTE: All changes to existing database schemas cannot be saved and are intended for exploration purposes only. You can do the following structure changes in diagrams.

- Add new tables to a diagram by drag-n-dropping them onto the diagram field from the Database Navigator.
- Rearrange tables in the diagram by dragging them all over the space. You can select several tables and drag them to a new location.
- Auto-arrange tables into a compact view after manual rearrangements: click the **Arrange Diagram** ( ) in the toolbar or on the context menu (right-click anywhere on the diagram tab).
- (Available for Custom Diagrams only) - connect tables with a connector: click the **Show Palette** button ( ) in the upper-left corner of the diagram tab and then, in the Palette panel, click **Connection**:

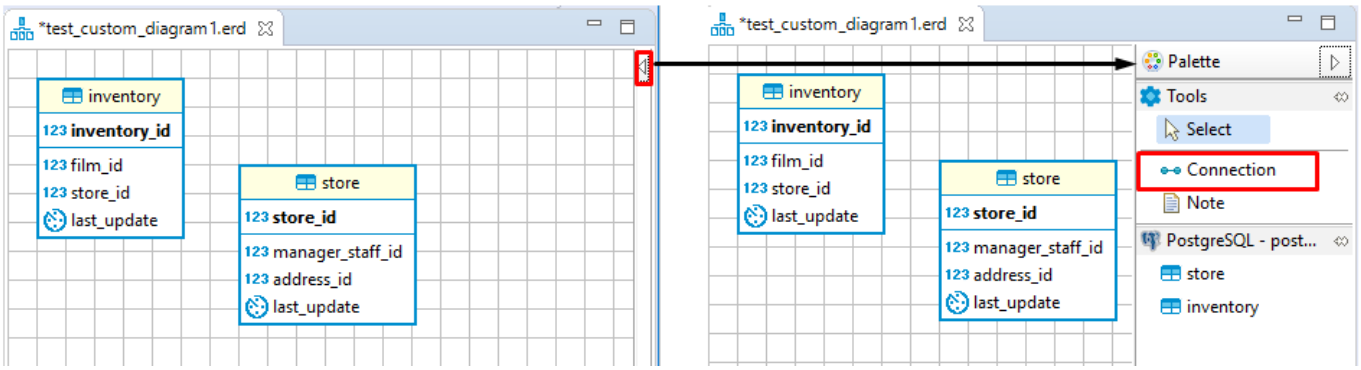

Now click the tables that you want to connect with each other in turn one by one, and then to stop the connection line double-clicking the last table.

- (Available for Custom Diagrams only) - remove tables and connections: right-click the table or conection and click **Delete** on the context menu or just click the table or connection and press `Delete`.

View Adjustment

You can adjust the view of any diagram in the following ways:

- Enable/disable the diagram grid: Click **Toggle Grid** ( ) in the toolbar.
- Modify attributes visibility: Right-click the diagram and, on the context menu, click **Show Attributes** and then one of the options:
  - **All** - all attributes
  - **Any keys** - primary and foreign keys
  - **Primary key** - only primary keys
  - **None** - no attributes
- Modify attributes presentation: Right-click the diagram and, on the context menu, click **View Styles** and then one of the options:
  - **Show Icons**
  - **Show Data Types**
  - **Show Nullability**
  - **Show Comments**
  - **Show Fully qualified names**
- Change the color of entities/notes: Right-click the header of the entity or comment and then click **Set color** on the context menu. Then you can select the color and click **OK**.
- For elements located in front of/behind others, bring an element to front or send to back: Right-click the element and then click **Bring to front** / **Send to back** on the context menu.
- Zoom the diagram in/out: Click the **Zoom In**/**Zoom Out** buttons or choose the scaling value in the dropdown list in the toolbar:

Refresh

To see changes made to the database schema by other users, you might need to refresh the diagram: click **Refresh Diagram** ( ) in the toolbar.

Notes

You can create notes only in Custom Diagrams. To create a note, click the **Show Palette** button ( ) in the upper-left corner of the diagram tab, then, in the Palette panel, click **Note** and then click anywhere in the diagram tab. Now you can double-click **Note** box to enter the note text:

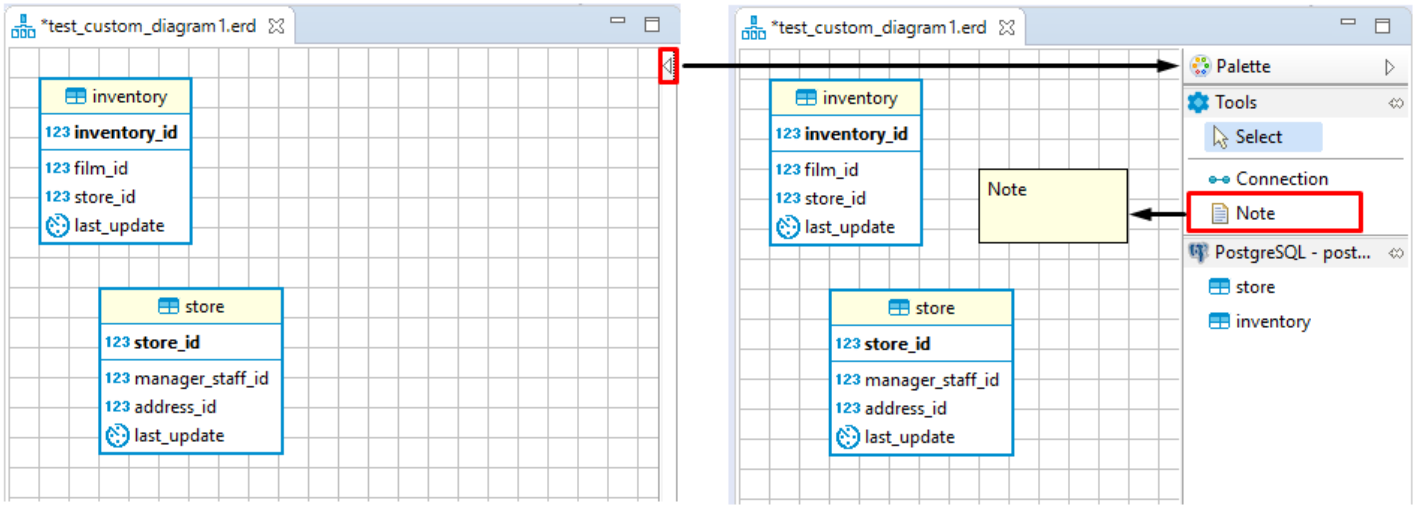

Search in Diagram Entities

To search among entities of a diagram, click the **Search items** button ( ) in the toolbar, then type in the search combination. The entities that contain the search combination are highlighted in the diagram. To remove the filter, click the cross icon next to the search field.

Diagram Export

You can export (save) a diagram as an image (PNG, GIF, BMP formats) or as a file in GraphML format. To export a diagram, click **Save diagram in external format** ( ) in the toolbar.

Diagram Printing

To print a diagram, press `CTRL+P` or click **Print Diagram** ( ) in the toolbar.

Settings

To modify the diagram settings, click **Configuration** ( ) in the toolbar.

Database Structure Diagrams

You can view a database structure in the standard ERD (Entity Relation Diagram) form. ER diagrams are available for all tables and schemas (databases).

ER diagram for a table shows the table itself and its relations with other tables inside the schema. To view the ER diagram for a table or view, double-click the table or view in the Database Navigator and then, in the Database Object Editor, switch to the **ER Diagram** tab:

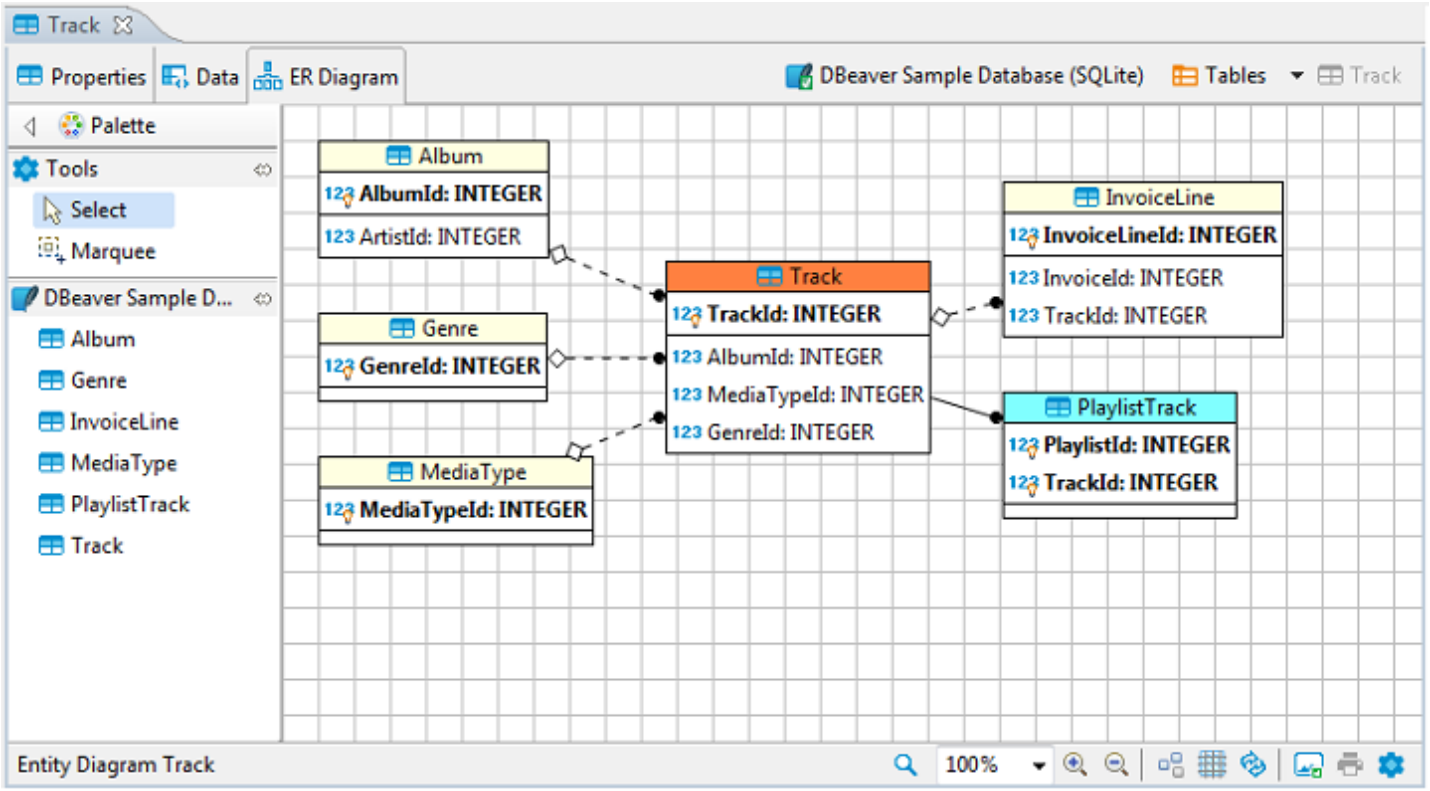

To view the ER diagram for a full database schema, double-click the schema name in the Database Navigator or the previous node in the path (usually - **Tables**):

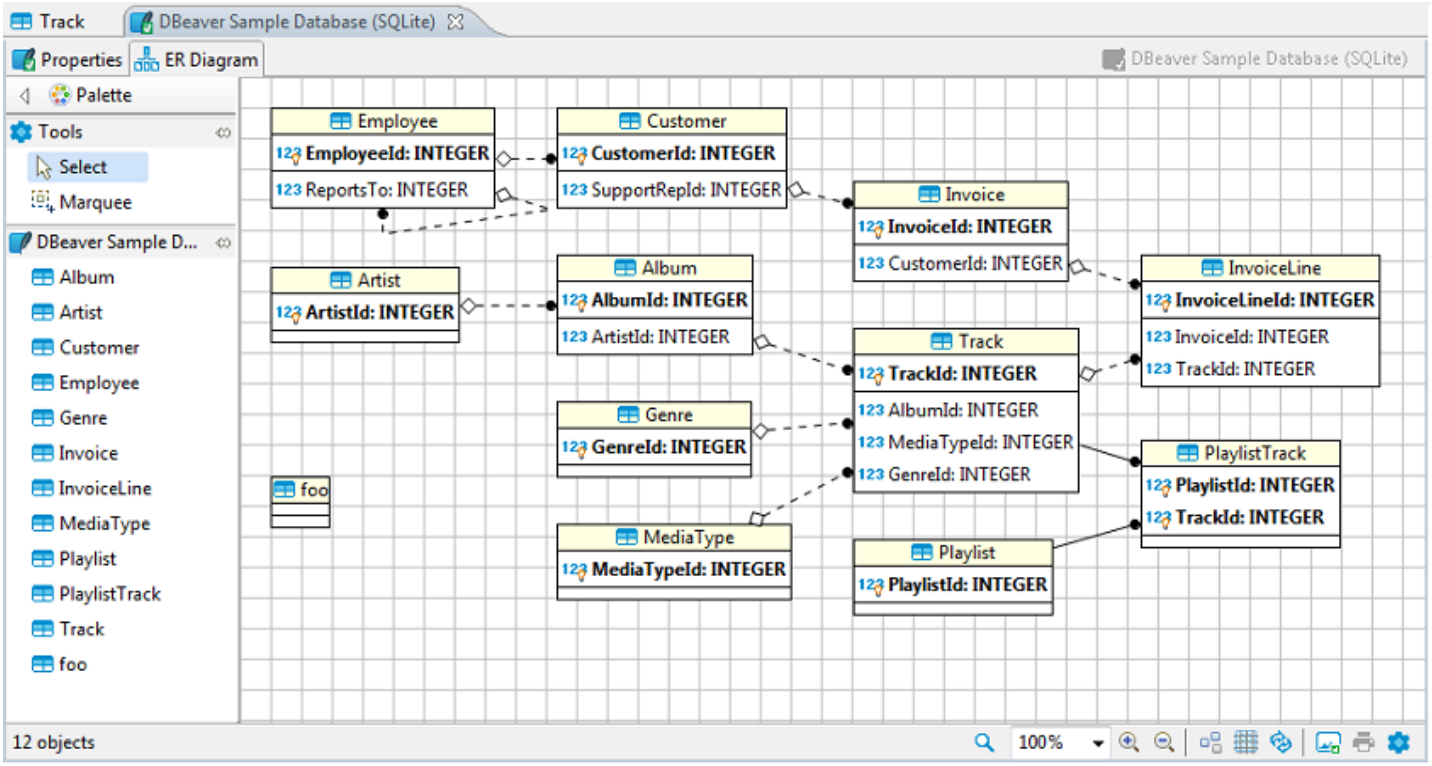

NOTE: Table and schema diagrams are read-only. You can rearrange the layout, drag-n-drop elements inside a diagram but you cannot save the changes state or delete/add anything. This is because the diagrams represent the actual state of databases.

Custom Diagrams

You can create custom ER diagrams that can contain any tables, relations and notes.
However, even custom diagrams may contain only real existing database entities (tables).

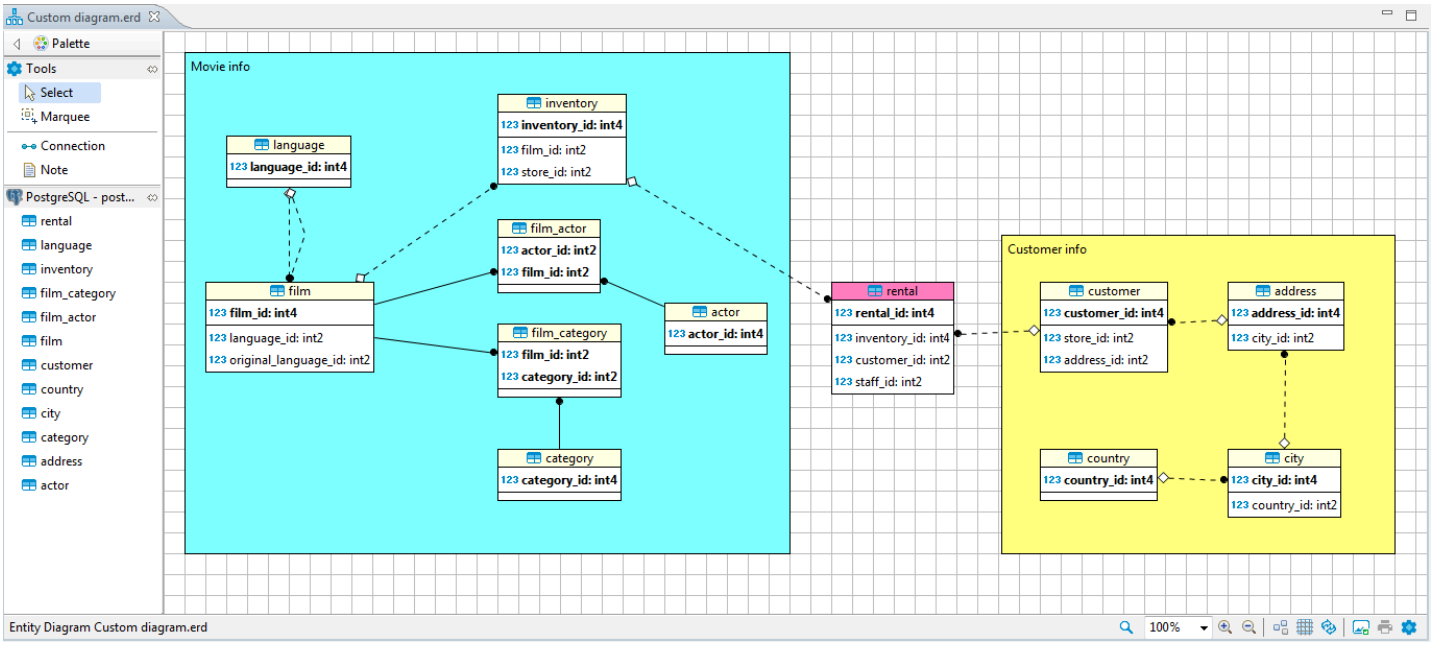

You can create a custom diagram in one of the ways:

1. On the DBeaver main menu, click **File -> New**. Then in the new diagram wizard, click **DBeaver -> ER Diagram** , and then **Next**:

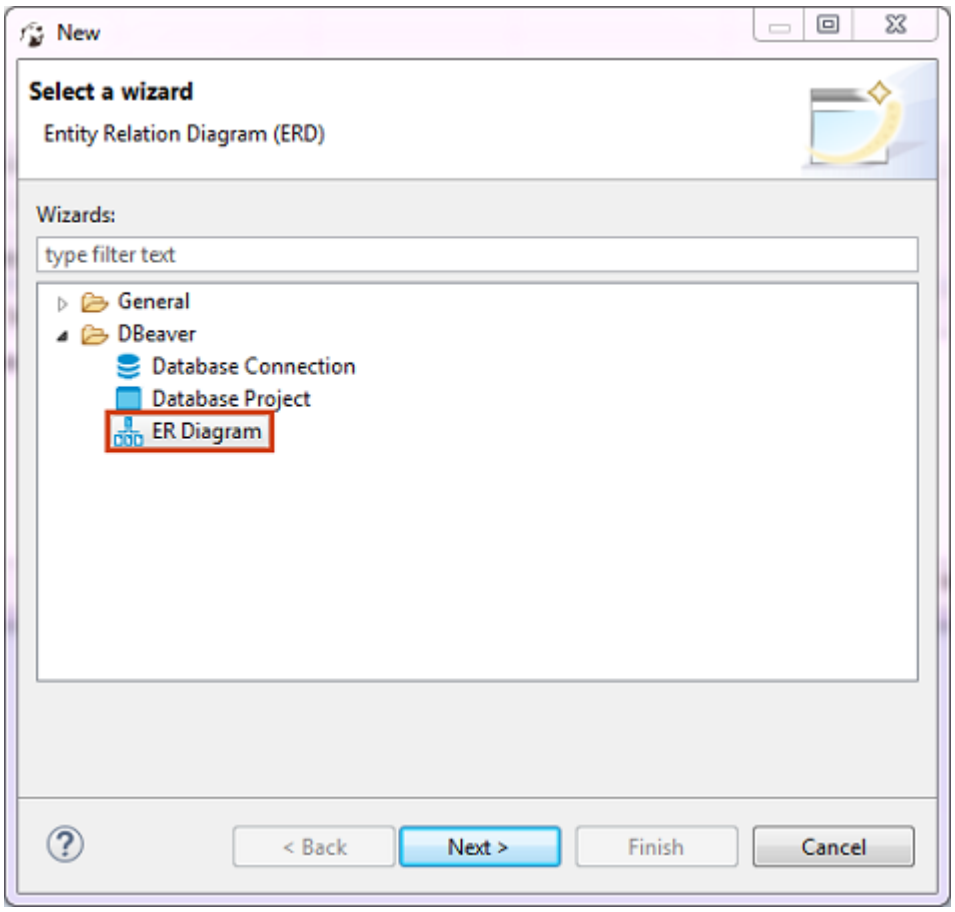

2. In the Project Explorer view, right-click the **ER Diagrams** node and then click **Create New ER Diagram** on the context menu.

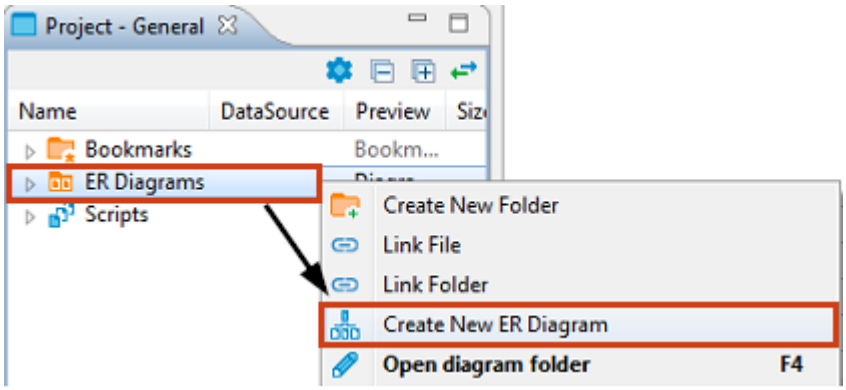

In both cases, in the Diagram Create Wizard, specify the diagram name and (optionally) choose initial diagram contents (set of tables):

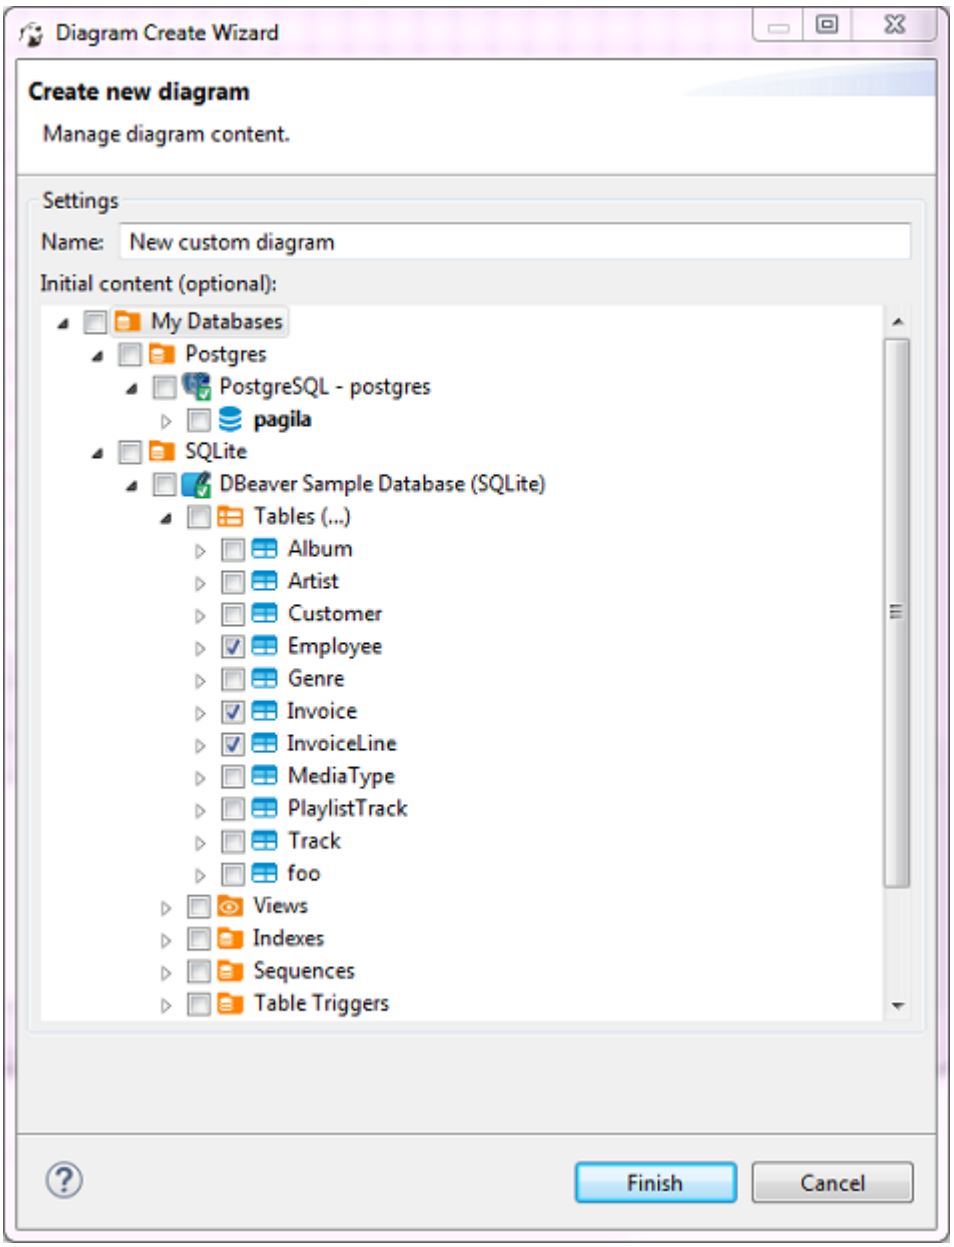

The new diagram appears in a separate editor. Now you can drag-n-drop any number of tables onto it. You can add tables from different connections and even different database type (for example, combine Oracle and MySQL tables in one and the same diagram).

You can also add notes and custom relations (associations) using the ERD palette on the left side of diagram tab, see details in ER Diagrams article. For example, to create a diagram similar to the one shown at the beginning of this article, you need to:

1. Add required tables and relationships between them and move them around to create a well-shaped structure (see *Structure Adjustment* section of the ER Diagrams article).
2. Add notes (see the *Notes* section of the ER Diagrams article).
3. Stretch the notes to cover the intended tables, then send the notes to back, and then set color to the tables and notes (see the *View Adjustment* section of the ER Diagrams article).

Undo/redo functions are fully supported in diagram editing.

Search

DBeaver provides:

- File search (search among file contents)
- Database full-text search
- Database metadata search

To use search, click the Search button in the main toolbar:

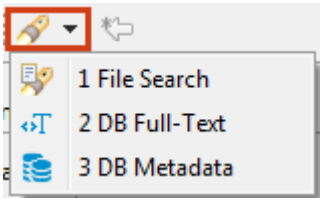

Please see dedicated articles for information about search of different types. This article describes common features of the three search types.

Search View

Search results for any of the search types appear in a separate Search view. The following image shows the Search view for the database full text search:

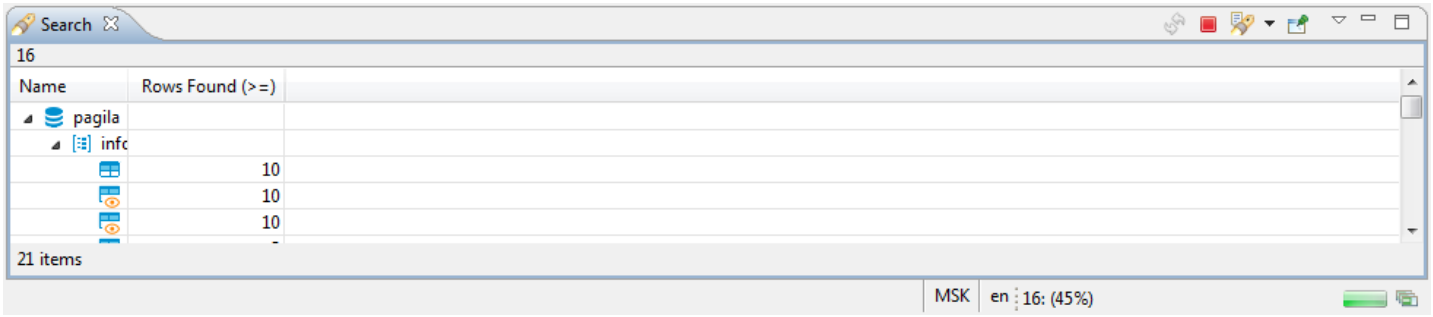

The view contains a toolbar that provides common tools for all types of search as well as specific tools for the File Search type. The following are common tools:

Button	Name	Description
	Run the Current Search Again	Repeats the search the results of which are displayed in the Search view
/	**Cancel Current Search**	- Active state (red) indicates that the search is still in progress and appears if the search takes some time to complete. Clicking the button in this state stops the current search. - Inactive state (grey) indicates that the search is complete. The button in this state is non-actionable.
	Show Previous Searches	- Clicking the button itself opens the Previous Searches window. - Clicking the arrow next to the button opens a dropdown menu. See the *Search History* section further in this article.
	Pin the Search	Ties the current search results to the Search view. If you click this button, the current results stay in the view while the results of the next search appear in a new Search view. Otherwise, every new

	View	search replaces the previous results with new results.

For information about specific tools of File search, see the File Search article.

If the search is short, the results appear almost instantly. But if it takes some time, the Search view indicates the progress in the following ways:

- The **Cancel Current Search** button in the toolbar has the Active state ( )
- The progress bar appears in the bottom-right corner of the view indicating the process:
- The button to show the search progress in a separate view ( ) appears in the bottom-right corner of the view next to the search progress bar. Clicking the button opens the Background Tasks view:

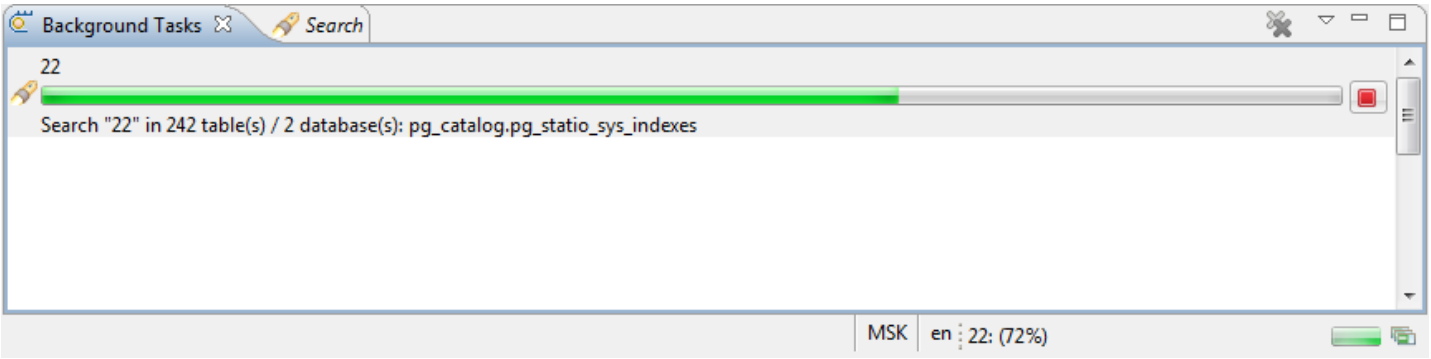

Search History

DBeaver stores the history of search queries made during the current session. You can reopen the Search view with results of a previous search query. You can also remove individual queries and clear the history. To manage the search history, use the **Show Previous Searches** button in the toolbar ( ).
To open the results of a previous search query, do one of the following:

- Click the arrow next to the **Show Previous Searches** button in the toolbar and then click the query in the dropdown list:

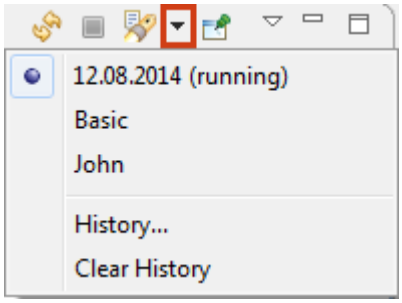

- Click the **Show Previous Searches** button itself or the arrow next to it and then **History** on the dropdown menu to open the Previous Searches window. Then, in the window, click the query and then either **Open** to open it in the active Search view or **Open in New** to open it in a new view:

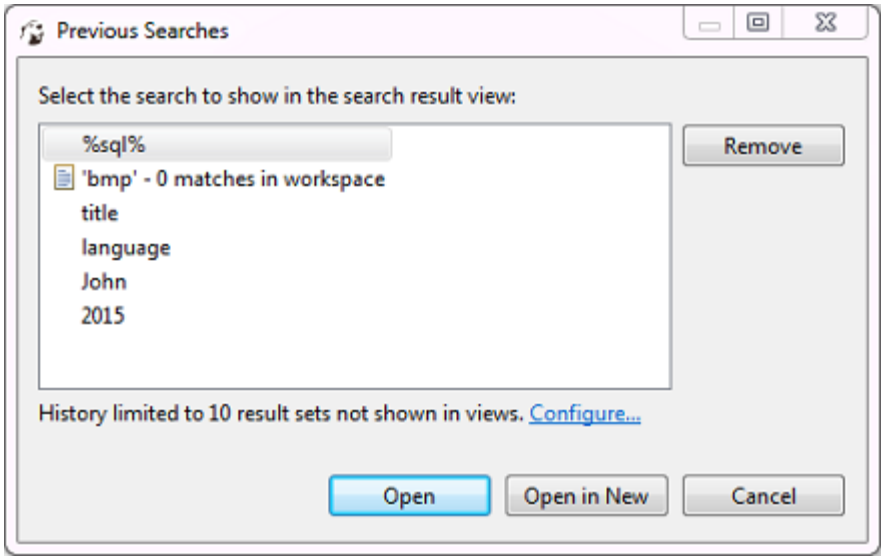

To remove one or more previous search queries:

1. Click the **Show Previous Searches** button in the toolbar or click the arrow next to it and then **History** on the dropdown menu. The Previous Searches window opens.
2. Click the query to remove or select several of them by clicking and simultaneously holding the `Ctrl` key.
3. Click **Remove**.

To clear the history by removing all previous queries, click the arrow next to the **Show Previous Searches** button in the toolbar and then click **Clear History** on the dropdown menu.

File Search

To search file contents for a string, click the Search button in the main toolbar or the arrow next to the Search button and then **File Search** on the dropdown menu:

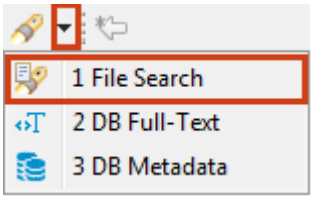

The Search window opens displaying the File Search tab:

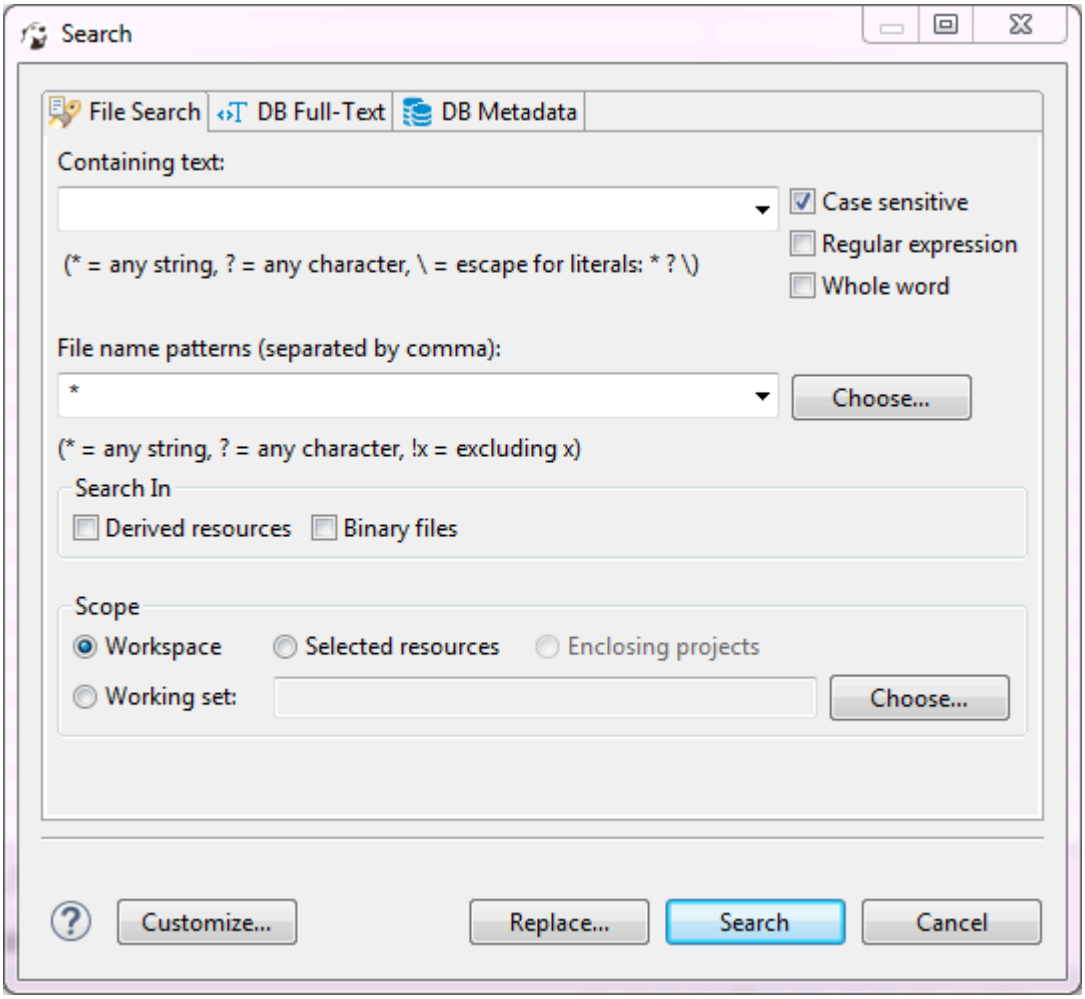

You can apply case sensitive search, search by regular expressions, search among particular file types (**File name patterns** field), and use the find and replace function.

After you click **Search**, the results appear in a Search view. The results represent a tree or list of files with the search combination highlighted:

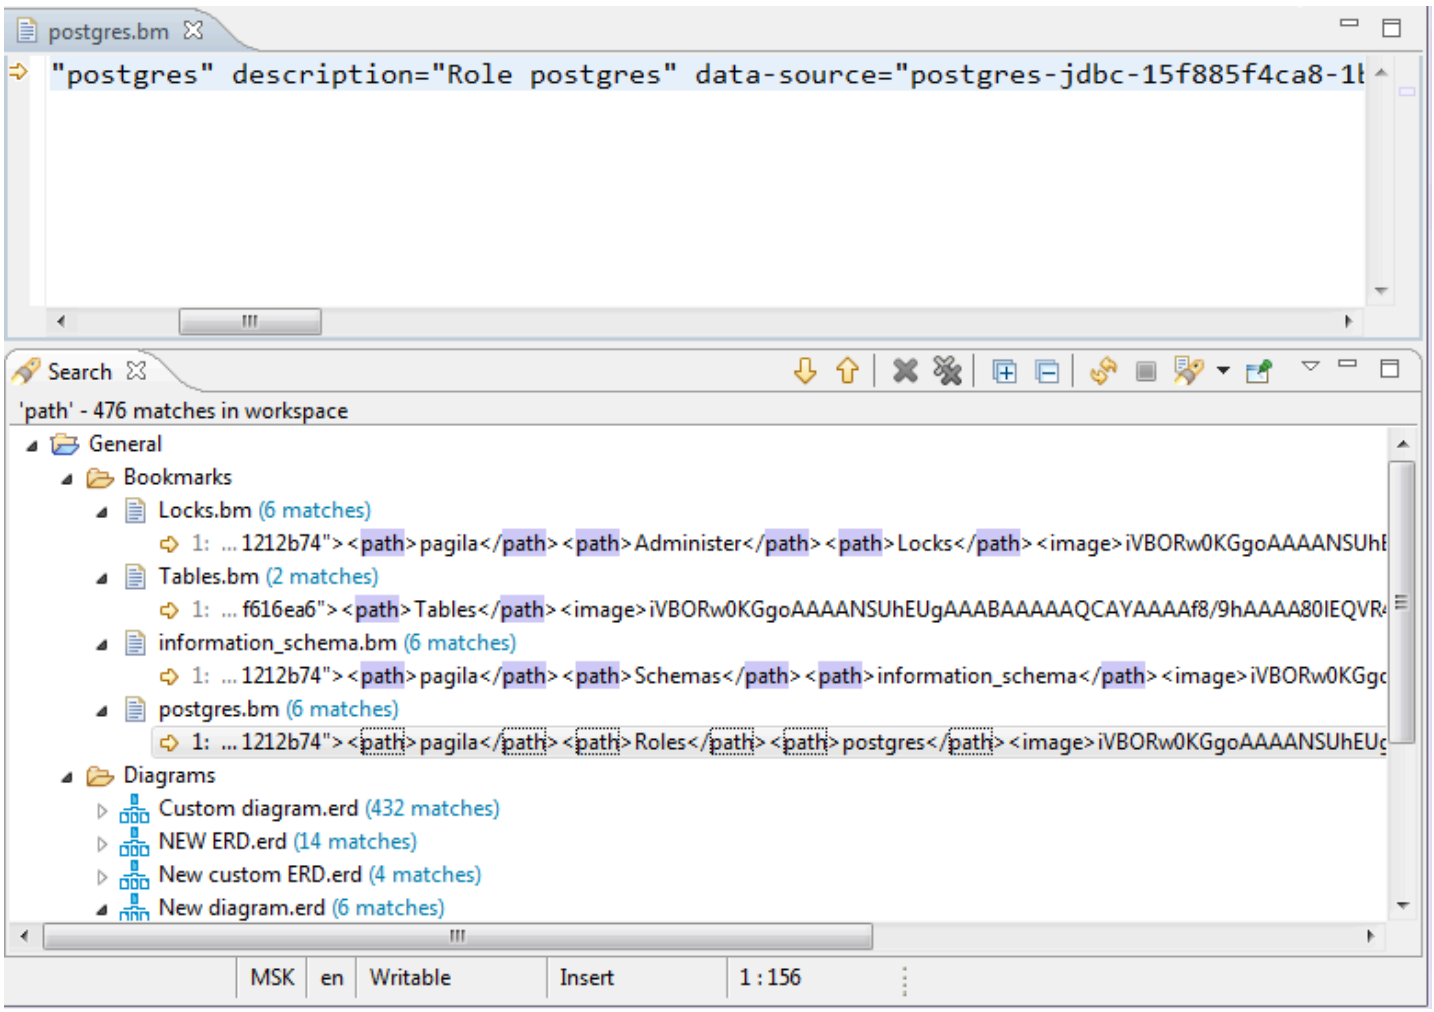

The toolbar of the Search view for File search provides more tools in addition to those available for all search types:

Button	Name	Description
	Show Next / Previous Match	Open the file in a separate viewer and move the highlight to the next/previous match, respectively
	Remove Selected Matches	Removes selected row (row in focus) of the results
	Remove All Matches	Removes all results in the view
	Expand/Collapse All	Expand/collapse the tree of results

The view also provides a view menu (click the **View Menu** button ( ) in the upper-right corner of the view) that contains the following items:

Menu item	Description
Show as List	Presents the results in the form of list
Show as Tree	Presents the results in the form of tree
Filters	Opens Search Filters dialog box
Preferences	Opens the Preferences window on the Search page

Double-clicking a results row opens it in a separate viewer.

DB Full-Text Search

To do full text search in the database contents, click the arrow next to the Search icon in the main toolbar and then click **File Search** on the dropdown menu:

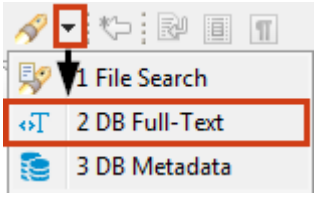

Alternatively, you can click the **Search** button on the main menu and then click the **DB Full-Text** tab in the Search window:

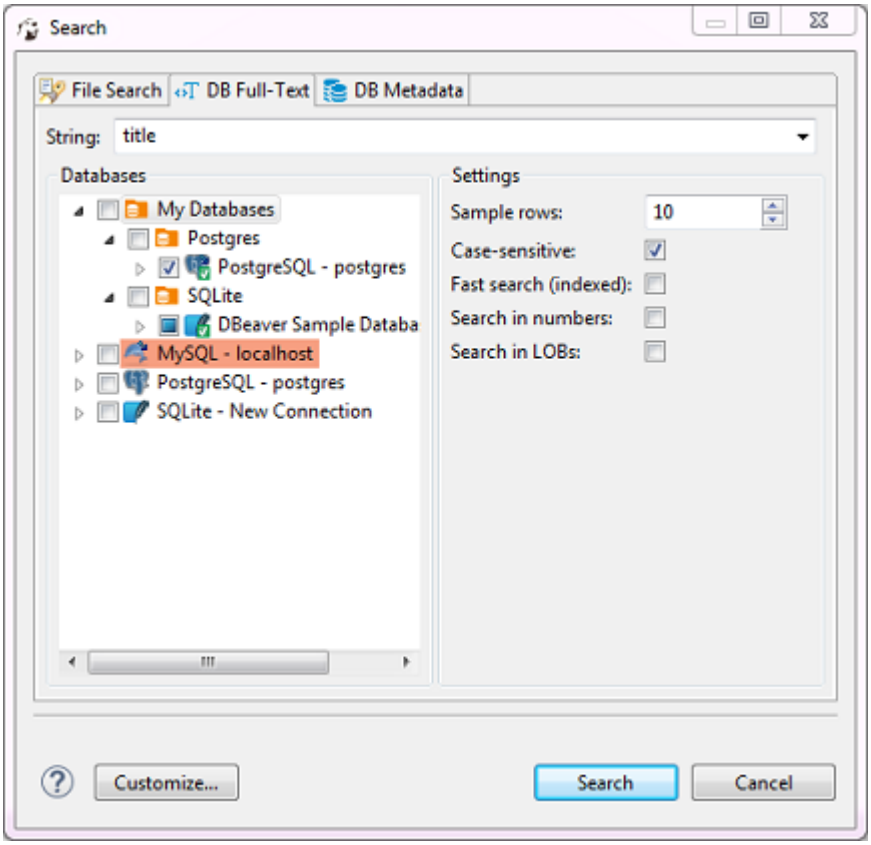

Now you need to choose the database connection or database objects against which to run the search – expand the tree in the **Databases** field to the database connections level or further down and select the checkboxes next to the required connections or database objects.
NOTE: The **Search** button is enabled only when you select the right level of checkboxes – database connections or lower nodes.

You can apply case-sensitive search, fast search and search in numbers and LOBs.

After you click **Search**, the results open in a Search view:

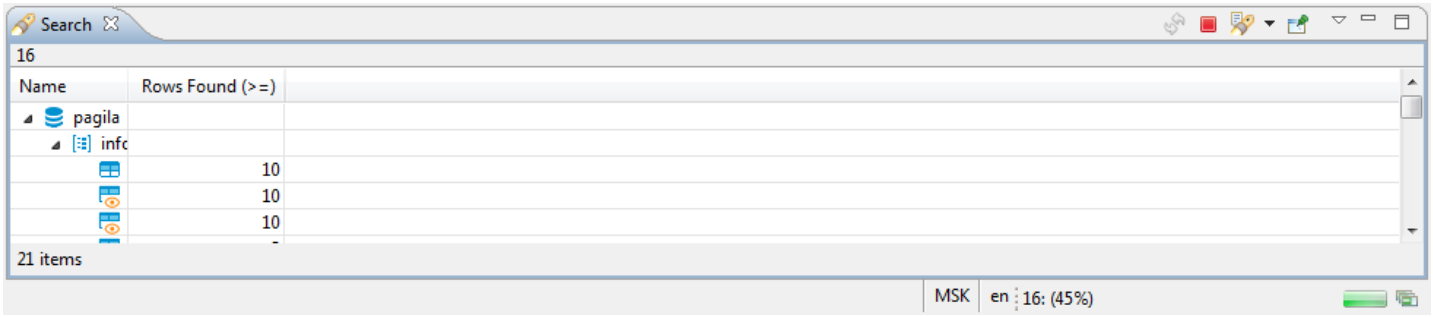

Double-clicking a row in the Search view opens the respective object in a dedicated Database Object editor.

DB Metadata Search

To search for database metadata, click the arrow next to the **Search** button in the main toolbar and then click **DB Metadata Search** on the dropdown menu:

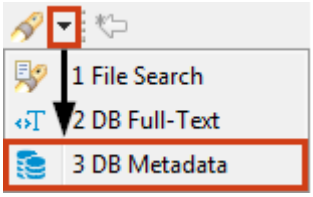

Alternatively, you can click the **Search** button on the main menu and then click the **DB Metadata** tab in the Search window:

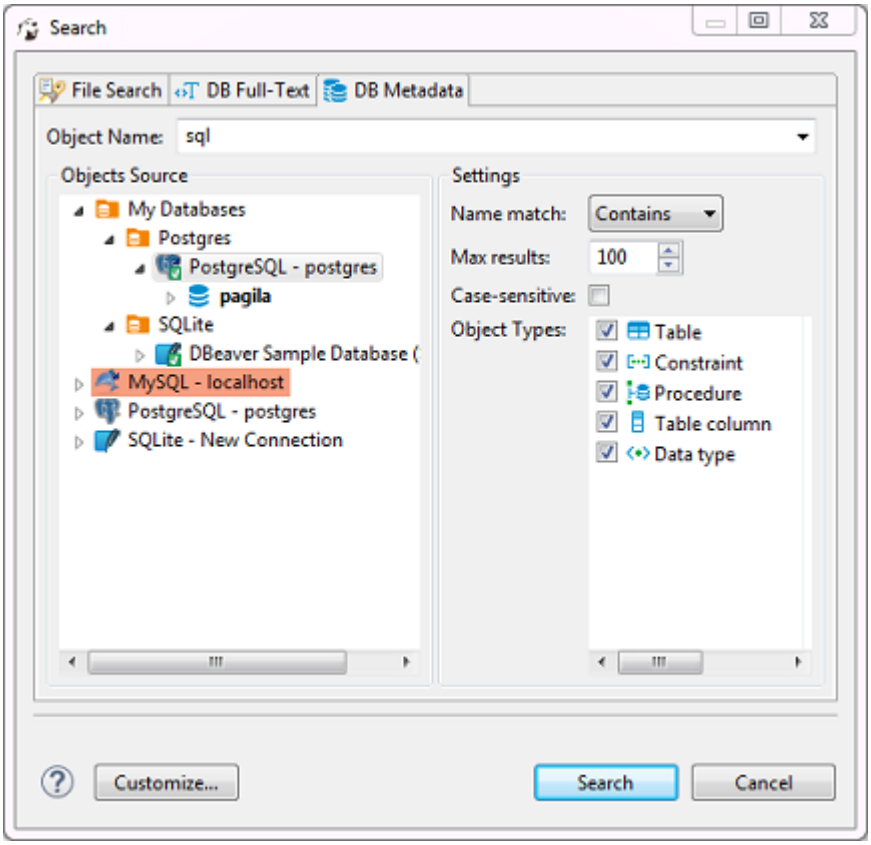

Now you need to choose the database connection against which to run the search – expand the tree in the **Objects Source** field to the database connections level and click the required database connection.
In the **Object Types** field, you can select the database objects among which DBeaver will run the metadata search – select or clear the checkboxes.
You can specify if the metadata should start with or contain or be similar to the search combination (**Name match** field). You can also set the maximum number of results to display (**Max results** field) and apply **Case-sensitive** search.

After you click **Search**, the results open in a Search view:

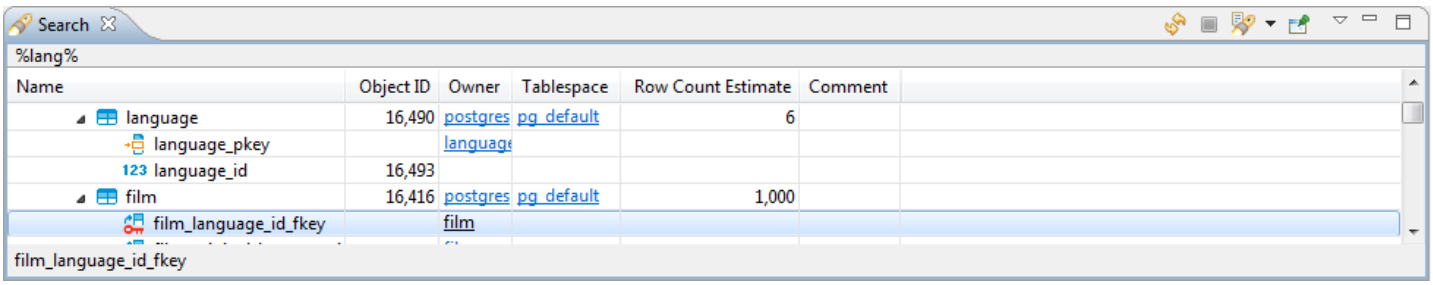

Double-clicking a row in the Search view opens the respective object in a dedicated Database Object editor.

Schema compare

Schema compare/migration

You can compare two schemas/databases structure and generate report as:

- DDL script (series of create/alter/drop statements)
- Diff diagram (sort of ER diagram)
- Liquibase change log
- Liquibase change report (json, yaml or plaintext)

Selecting objects to compare

- Select two objects (schemas, databases or tables) you want to compare
- Open context menu
- Open sub-menu `Compare/Migrate`
- Click on `Compare/Migrate Schema` element

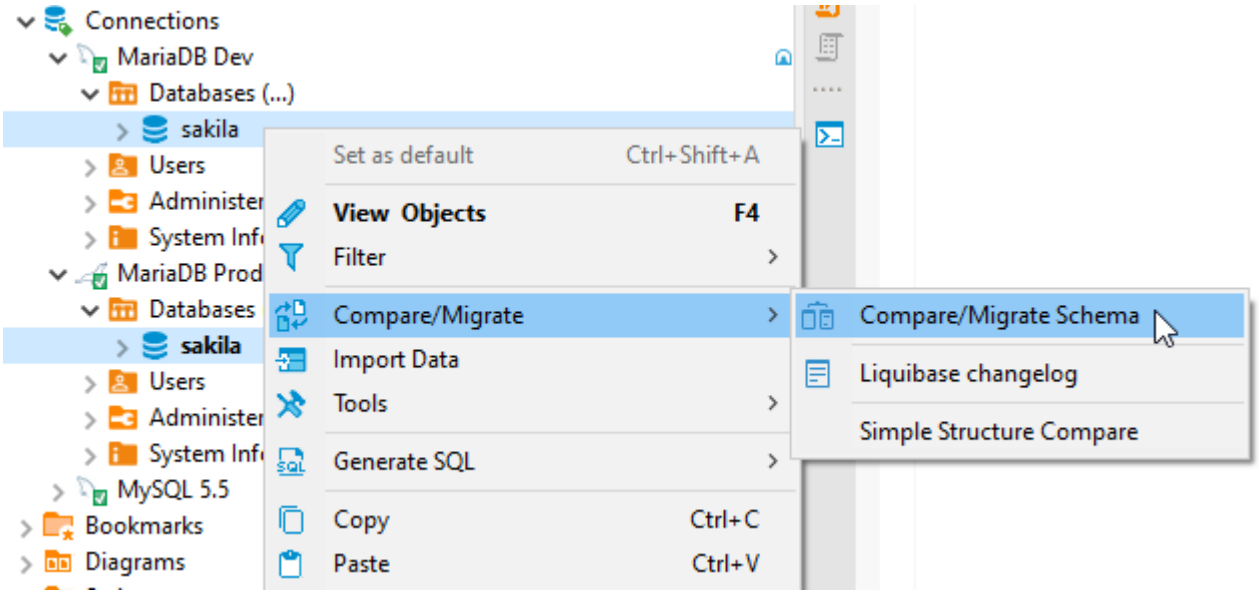

Compare settings

Re-validate that you chosen right objects to compare. You can also specify which changes to process: creates, drops, alters. By default all change types are enabled.

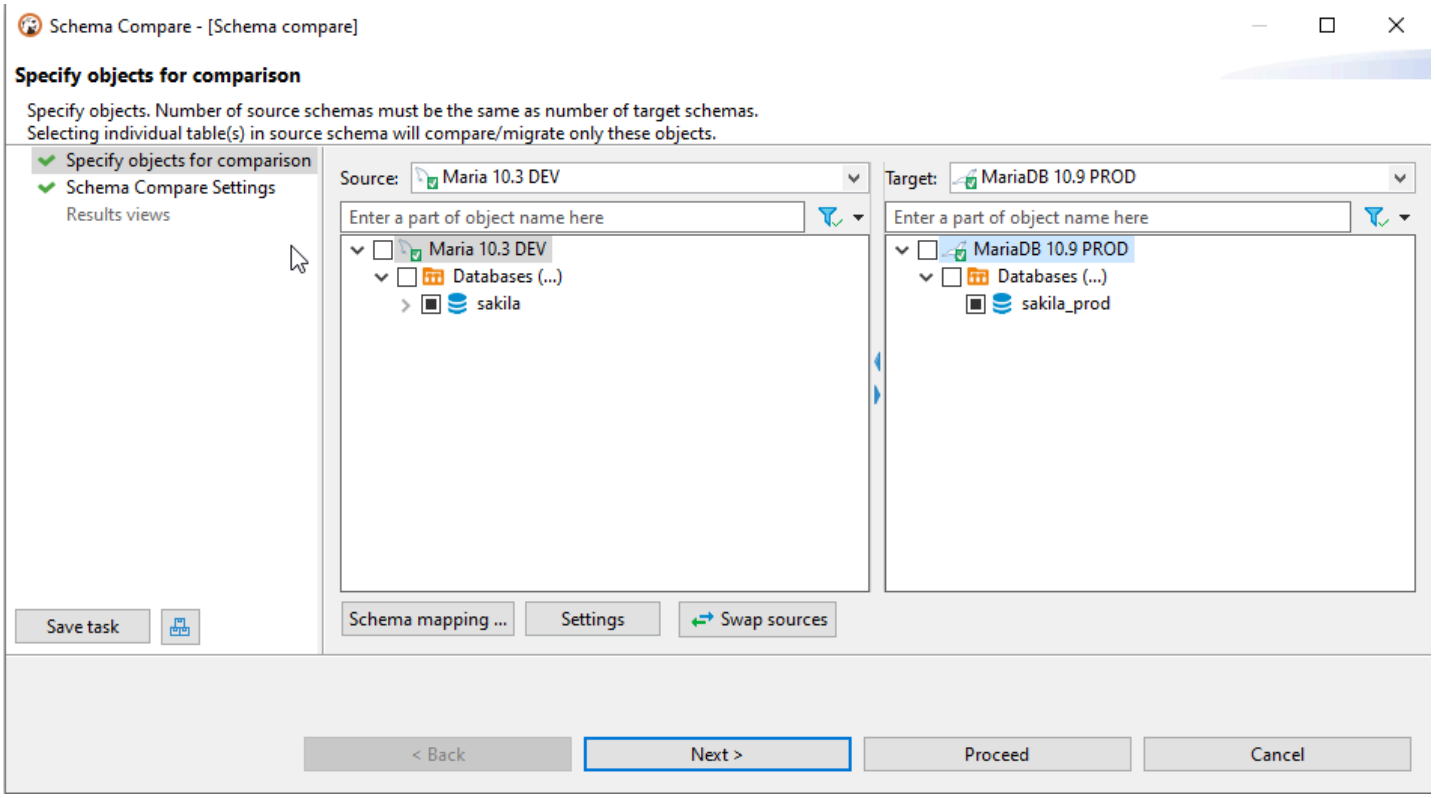

Also you can exclude the specific compared types of objects. For example, if you do not want to see in the final result comparison of sequences, views, external keys etc.

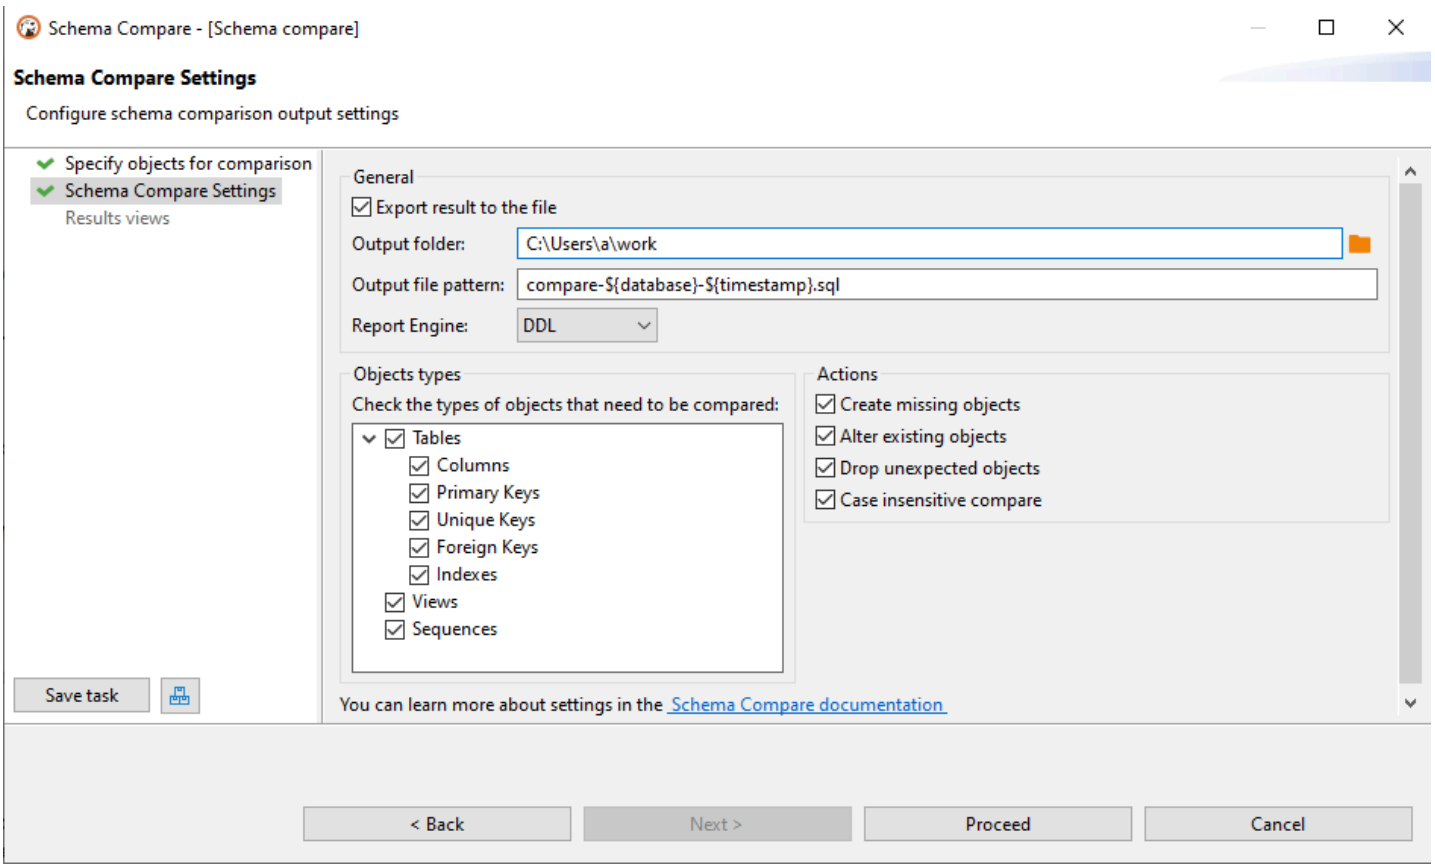

Compare results

Click on `Compare Schemas` to generate diff report.

By default DDL diff is generated. It contains series of create/alter/drop statements which will modify schema on the right side and thus will make it identical to the schema on the left side.

You can enable/disable particular changes in the tree on the left side of diff page:

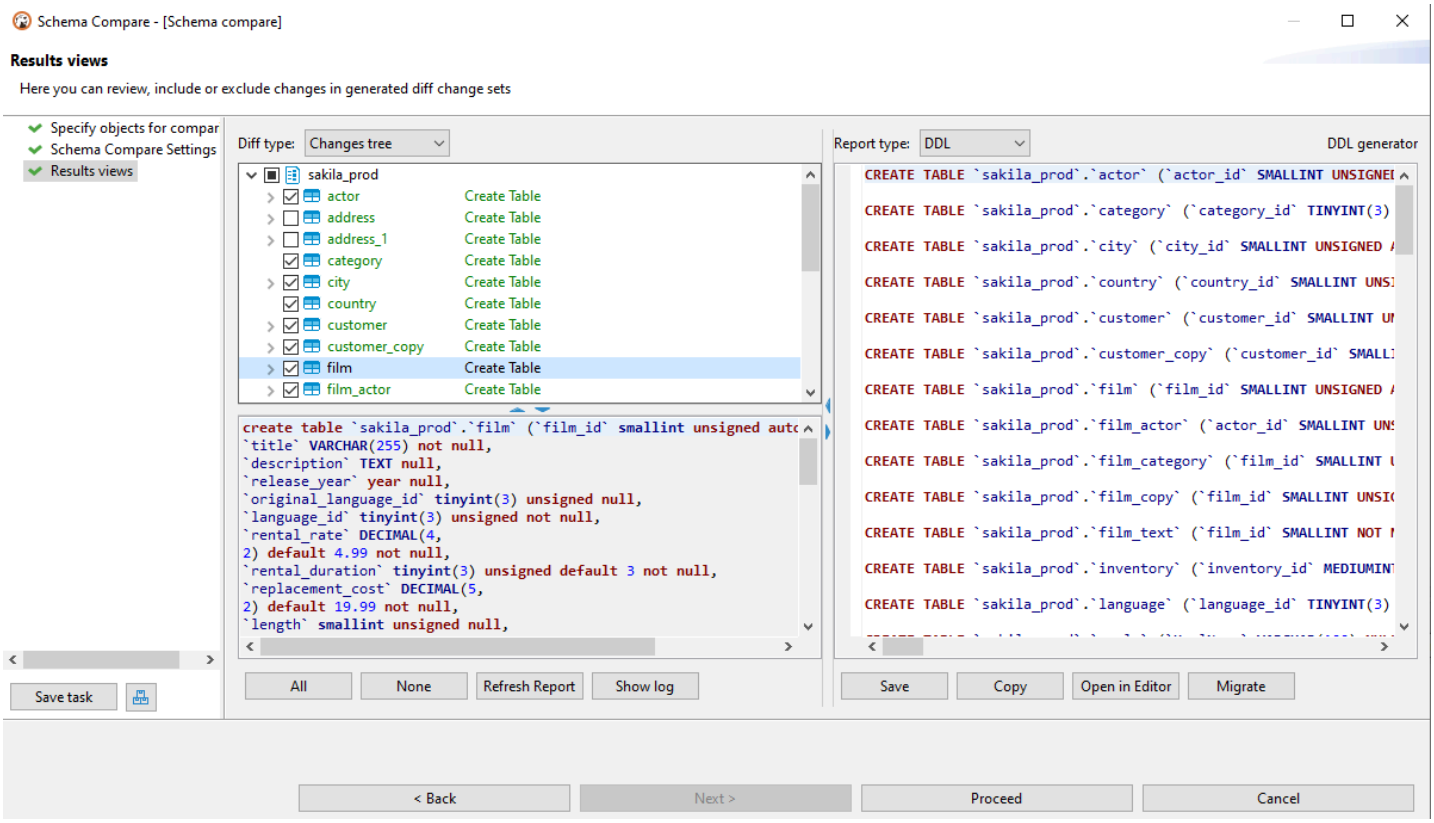

You can also switch to another diff report representation (diagram, json, yaml, plaintext).

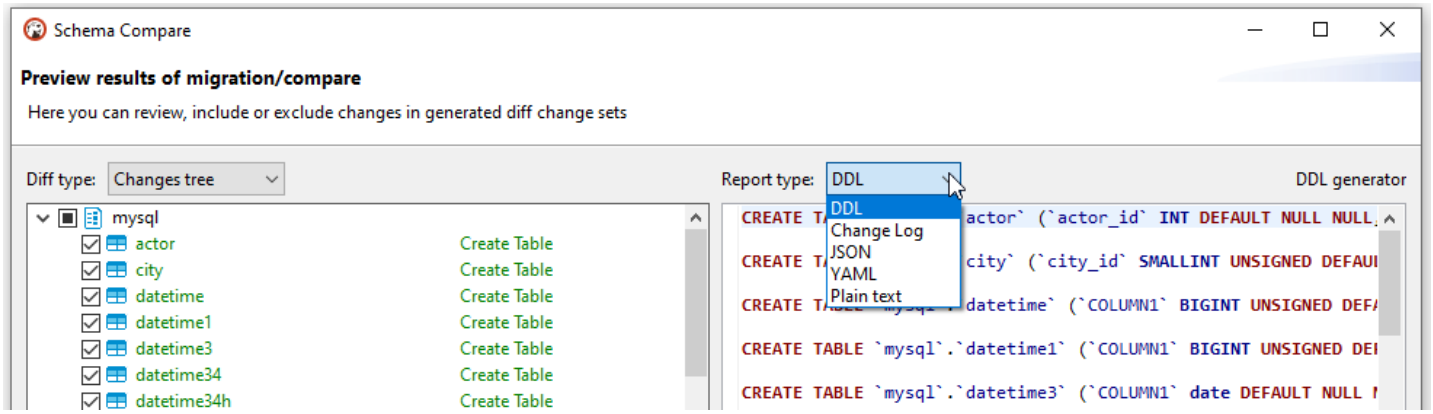

Compare schemaless bases

Some bases (like SQLite and Firebird) do not have catalogs and schemes that could be selected for comparison. In this case (and only for these databases), it is possible to compare the entire datasource entirely.

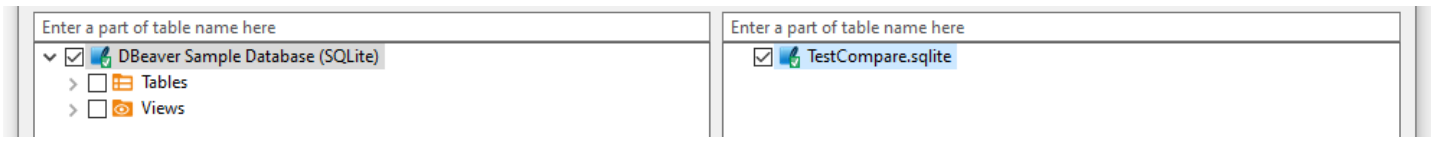

Data compare

Sometimes you need to compare data from two sources (tables) which have almost identical data with just few differences. There may be a plenty reasons to do so: quickly visualize all the differences and navigate through them by your sight, to copy differing rows or individual values or export them using Data Transfer.

Preparing the tool

1. You will need to choose one or two desired tables in the Database Navigator - it may be the tables from the same databases or from different databases or even from different RDBMS (e.g. PostgreSQL and MySQL).
2. Then choose `Database # Compare/Migrate # Data Compare` from the menu and the Data Compare Wizard will appear. Here you can preview selected tables or choose other ones. Then you can navigate to the second page.
3. After reaching the second page one may wonder about the purpose of it - don't worry, it's not that difficult: you should choose columns that will be used as a **unique key** during comparison (the amount of columns must be equal). If the *keys* chosen wrong, it may lead to invalid results. By default, if the table has unique key in it, it will be chosen automatically during the initial setup:

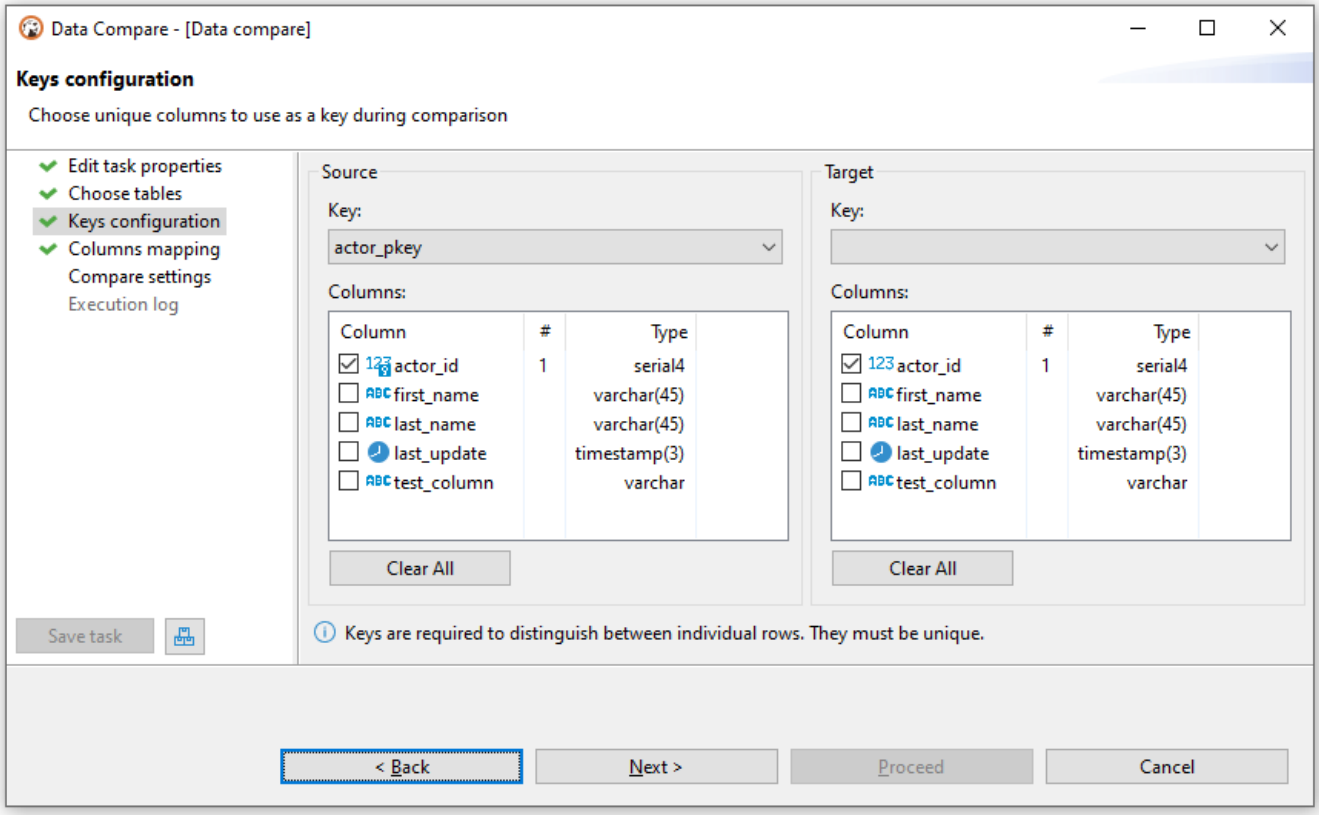

   Without unique keys the rows cannot be compared properly, since there's no way to distinguish between individual rows.
4. After reaching the last page, you can tweak limits (e.g. you only care about `N` first rows) and exclude categories of the resulting rows (e.g. you are not interested in modified rows). Also, you can press `Save tasl` button save configuration in a Task to use it later, or Schedule it, *although it is not very useful right now, but we have plans to support exporting results as an SQL script and more.*

Viewing the results

After finishing the wizard, you can press `Start` button to begin actual process of comparison - it may take some time depending on the databases you're comparing, the amount of data in them, and your network speed (*).

When the results are ready, DBeaver will play *beep* sound, and the editor will open:

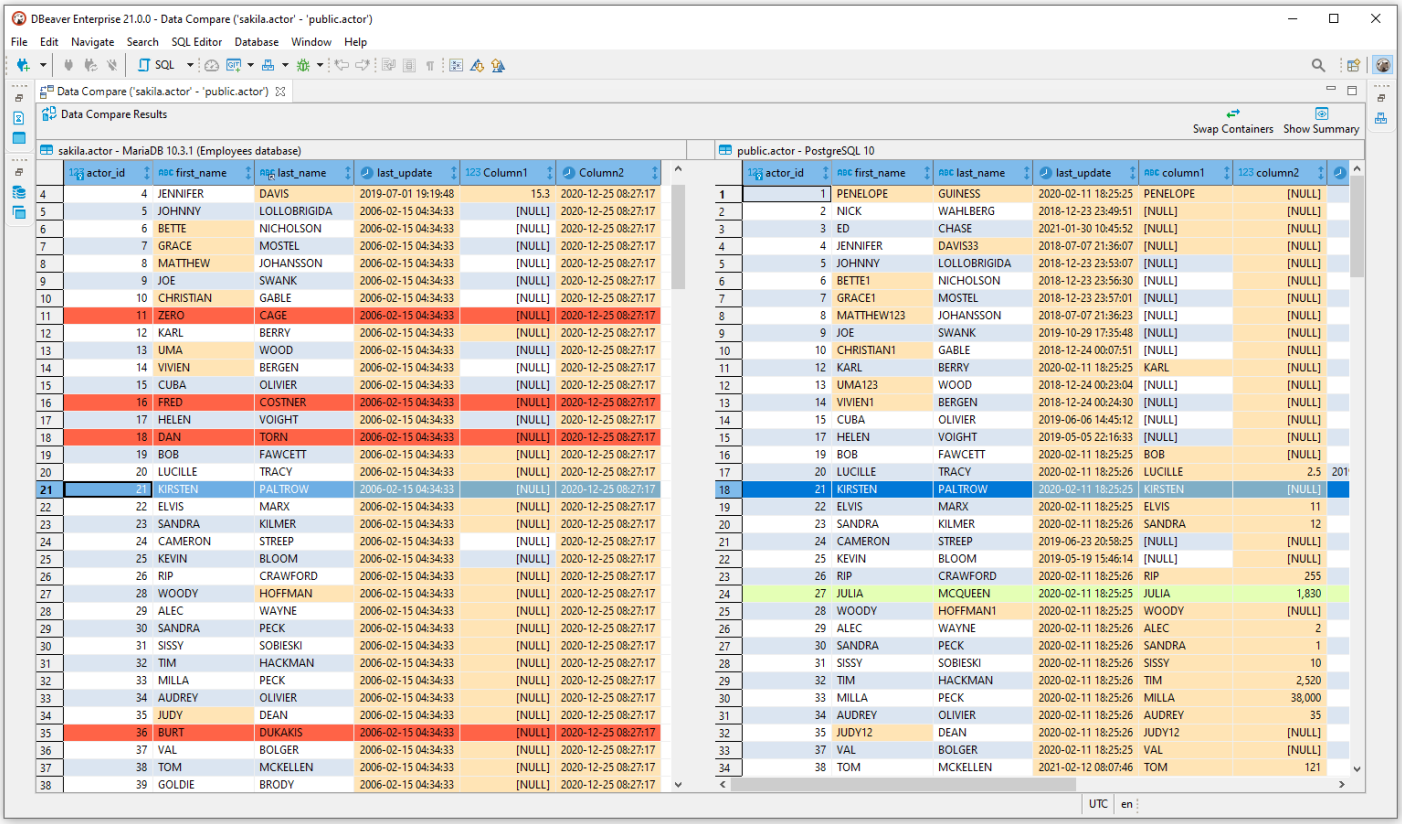

Here you can examine the differences, swap the panels using `Swap Containers` button in the toolbar, or preview summary including statistics by using `Show Summary` button in the toolbar too. Everything you can do in the regular data editor is it possible to perform here - you can copy data or transfer it to other database, except you cannot modify values.

** - Please note that the actual preview is only available when Data Compare is launched through the wizard - otherwise only statistics will be written into the Task Log.*

Under the hood

This section is under construction. It will include information about the engine's implementation details and more.

MockData generation

Note: since version 6.2 MockData generator extension is available only in Enterprise Edition.

Sometimes in software development we need to generate mock, but valid, data for testing. Populating a database manually is a time-consuming and exhausting process. It can be very complicated when you need to generate not just 5–10 users, but thousands of entities of different types. DBeaver Mock Data generator helps you generate test data much easier.

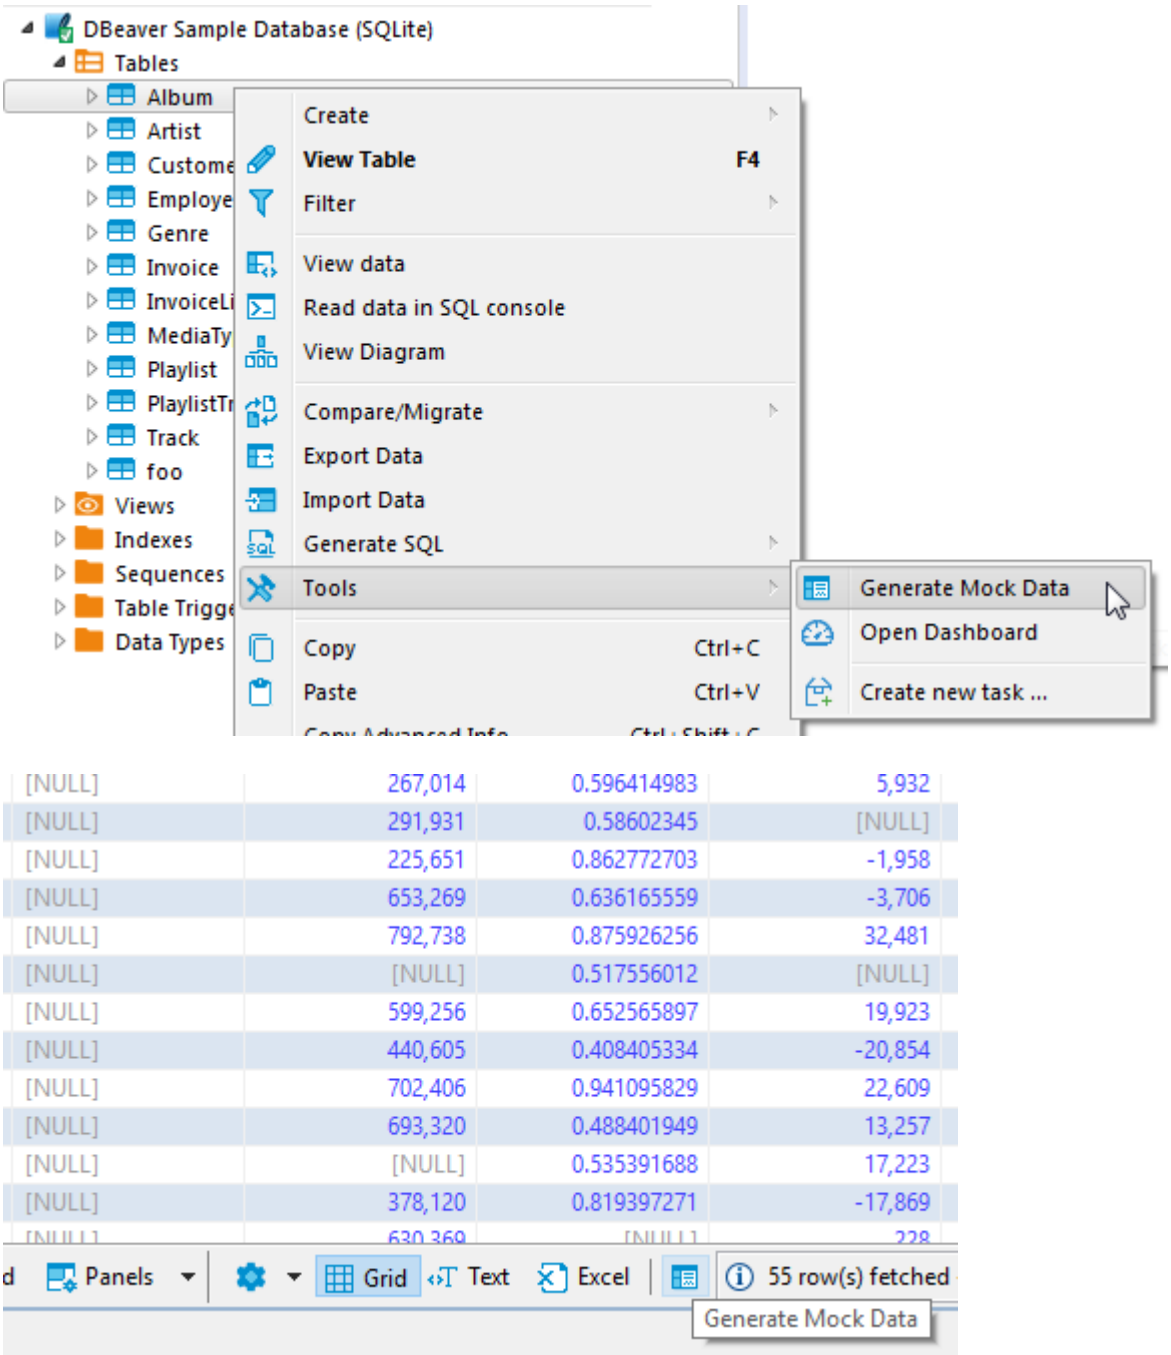

> *Disclaimer:* The idea behind Mock Data is to generate mock data in a table but it should **NOT TO BE USED IN PRODUCTION ENVIRONMENTS**. Please make sure you have a backup of your database before running the Mock Data generation process.

Th following are features of DBeaver Mock Data generator:

- Works for all the RDBMS that are supported by DBeaver (DB2, MS SQL Server, MySQL, Oracle, PostgreSQL, SQLite, etc.)

- Generates data that matches your database schema:
  - Generated data matches the database column types.
  - All base data types are supported.
  - Constraints (PK, FK, multi-column FK, unique) are supported.
- Supports over 20 configurable data generators (constants, randoms, sequences, names, domains, addresses, prices, regex based, etc.)
- Automatically associates a column with a generator based on the column characteristics
- Saves or overwrites old database data

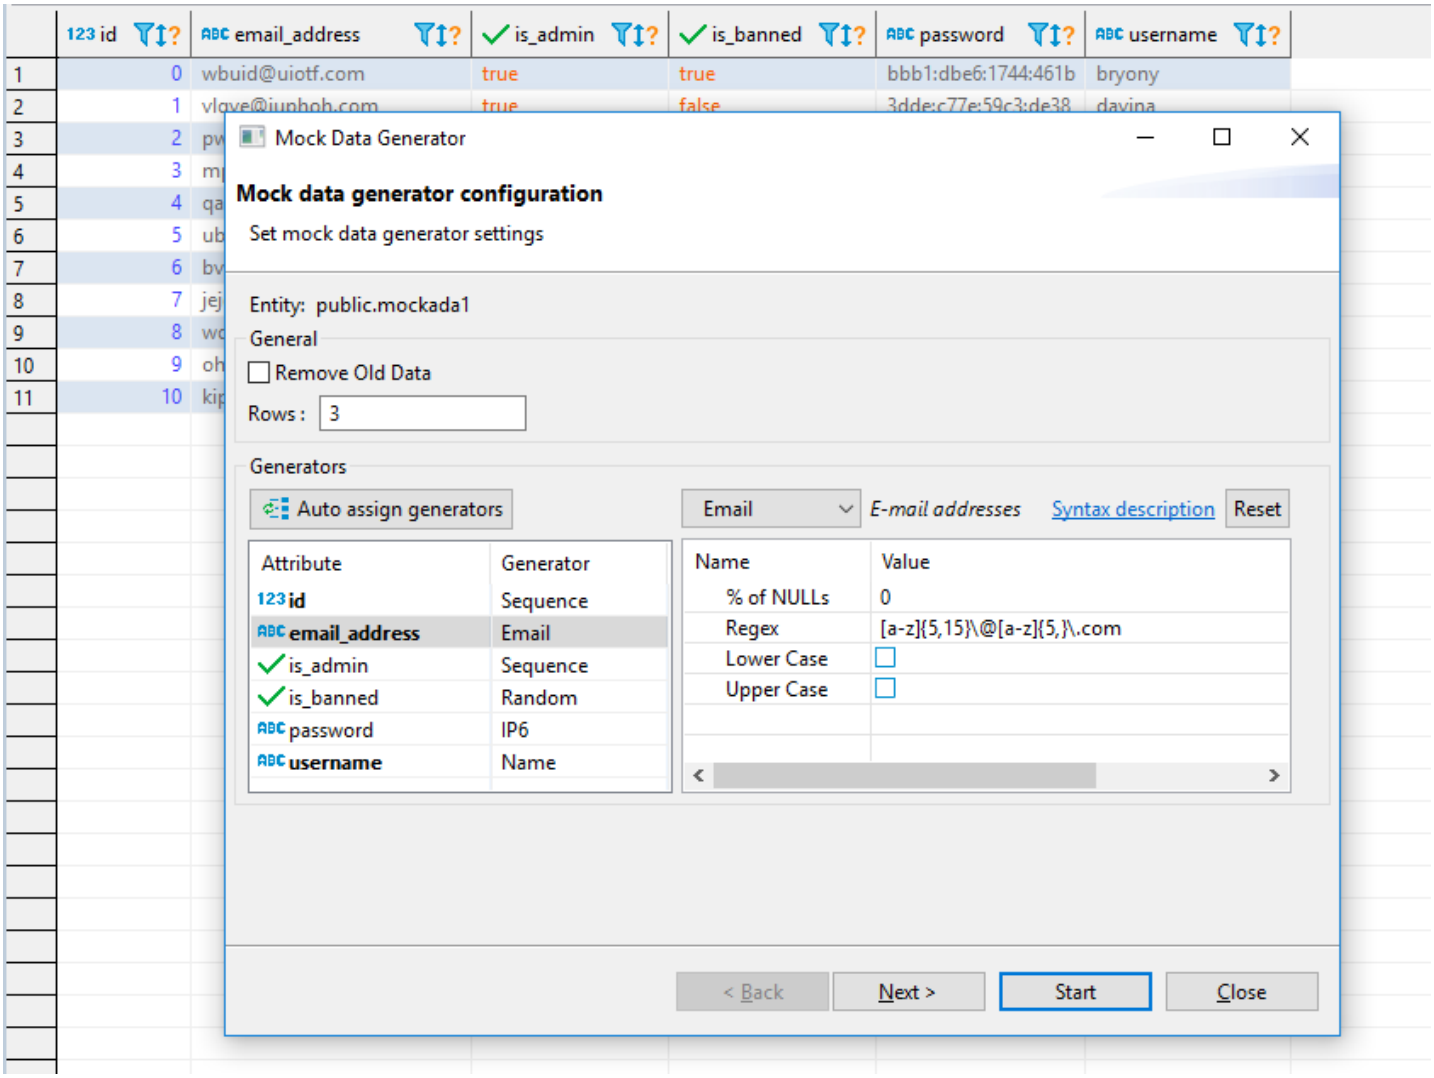

The following are mock data generators for data types with their configurable parameters:

- Boolean
  - Random
  - Sequence (initial, order)
- Date
  - Random (start, end)
  - Sequence (start, step, reverse)
- Numeric

- Random
- Sequence (start, step, reverse)
- Advanced (min, max, precision, scale) *
  - Price preset *
  - Coordinate preset *

- String
  - Text (template, min length, max length)
  - UUID
  - Address *
  - City *
  - Country *
  - Domain *
  - Email (gender, with surname, numeric suffix) *
  - Name (gender, with surname) *
  - Price (country, min, max) *
  - Regex based random (regex template) *
    - Credit Card preset *
    - Email preset *
    - Gender preset *
    - HEX Color preset *
    - IP4 address preset *
    - IP6 address preset *
    - Phone Number preset *
    - Postal Code preset *
    - Price preset *
  - Template with parametrized directives for other generators *:
    - address() - US postal address
    - city() - one of the world largest cities
    - country() - country

- domain() - one of the top Internet domains
- email(gender,surname) - e-mail address (gender is ALL|FEMALE|MALE, surname is true|false)
- name(gender,surname) - personal name (gender is ALL|FEMALE|MALE, surname is true|false)
- random(minimum,maximum) - random integer
- regex(pattern) - regex based value for the pattern
- sequence(start,step) - sequence of integers

- NULL values
- FK - data from the referenced table according to the constraint

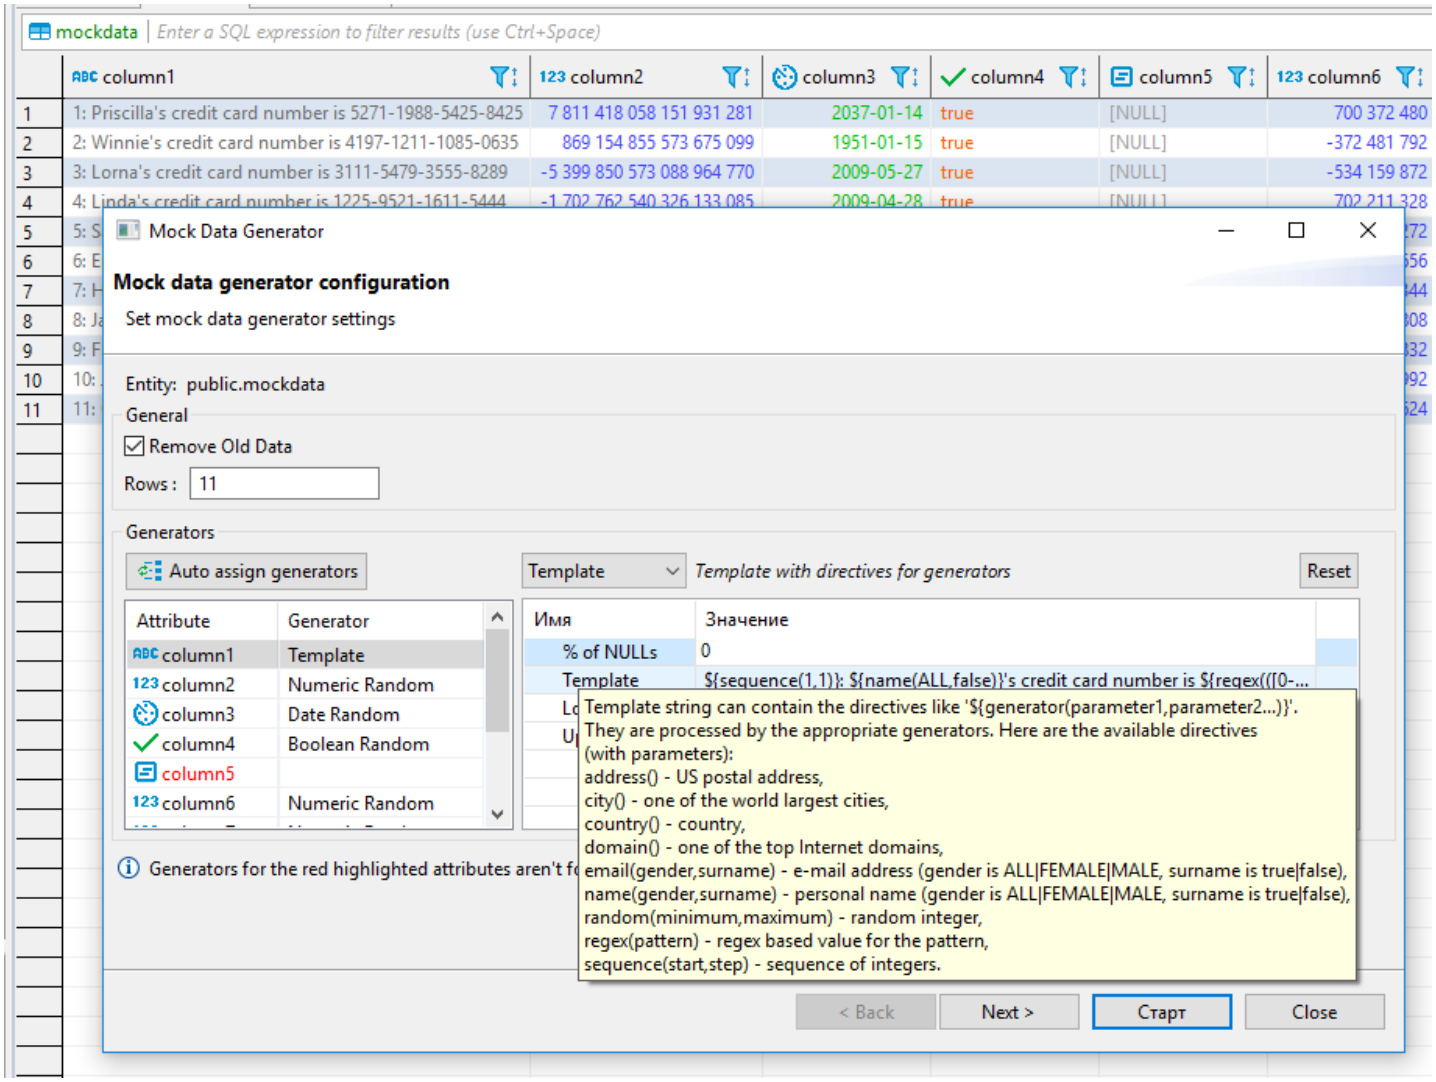

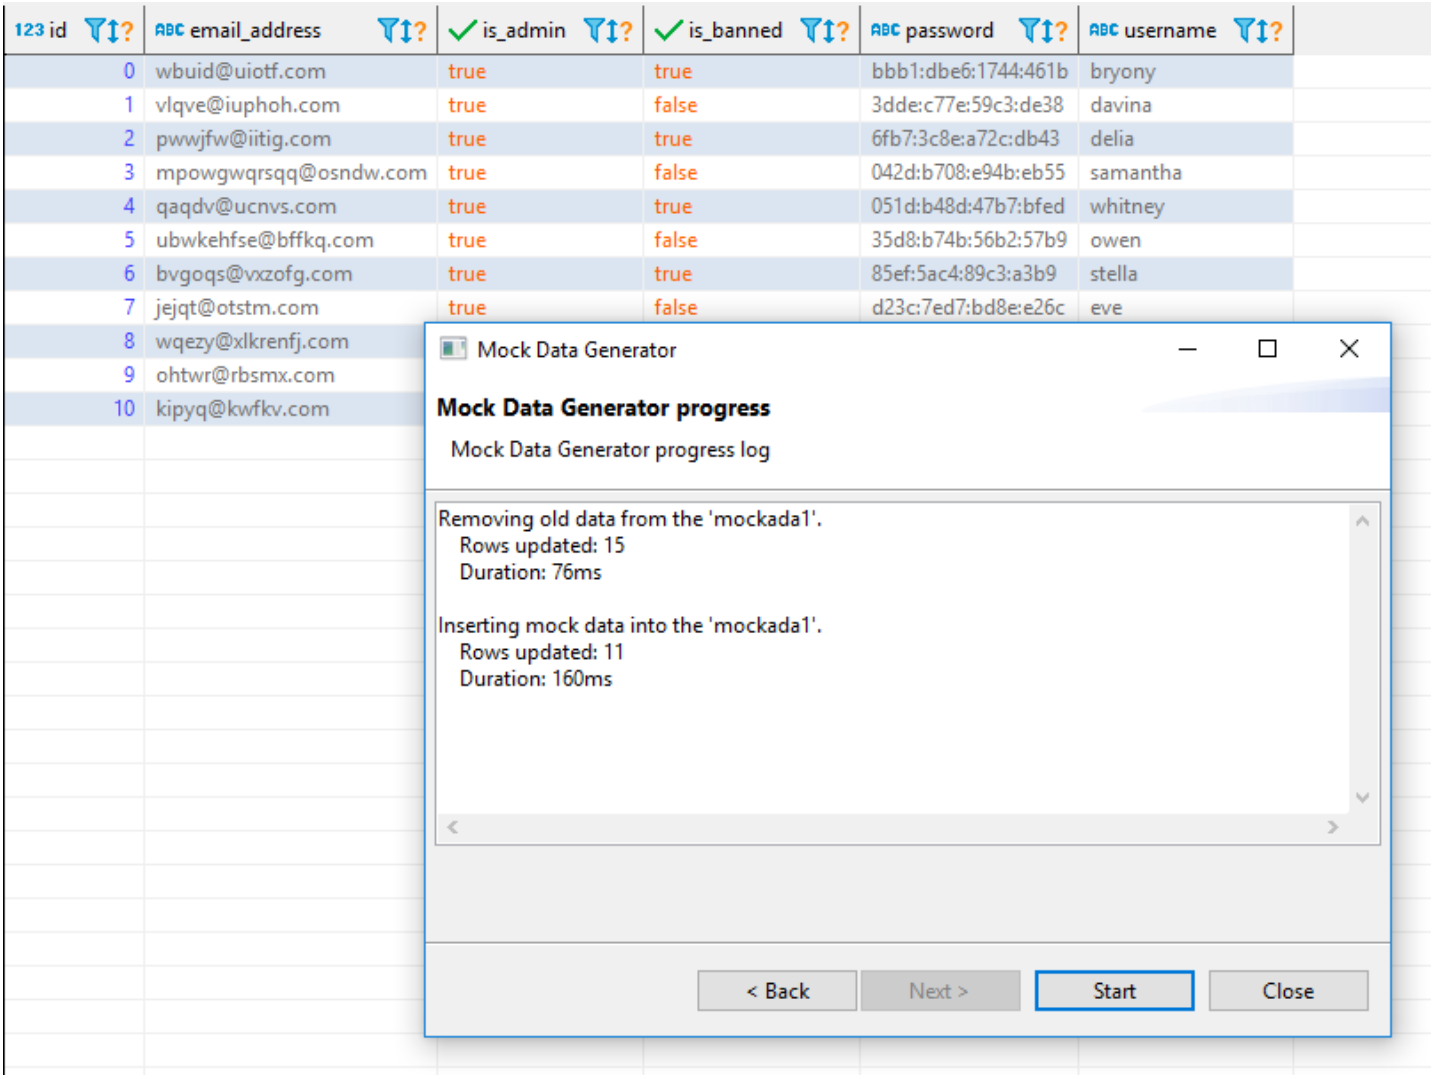

* These features are available in the DBeaver Enterprise Edition only.

Dashboards, DB monitoring

Dashboards tool allows DBAs and programmers to quickly identify performance, disk space issues, number of connections and other important KPIs associated with a single database connection. To learn more about database connections, see Database Connections.

By default, DBeaver is delivered with a number of predefined sets of dashboards for such data bases as PostgreSQL, MySQL, Oracle and Exasol. Custom dashboards are also supported. To learn more about custom dashboards, see Managing Dashboards section below.

Managing Dashboards Panel

Dashboards panel is a collection of real-time dashboards, that is dashboards that are continuously updated. Dashboards displayed on the dashboards panel are actually a combination of continiously run SQL SELECT queries and charts continiously built on the data fetched.

Opening Dashboard Panel

To open dashboards panel **Open Dashboard** button in the main toolbar. The default configuration of the dashboards panel for the current database connection will appear. To learn more about database connections, see Database Connections.

You can also right-click a connection name in the **Database Navigator** editor and select **Open Dashboard** menu option or use keyboard shortcut `Ctrl + Alt + Shift + B` and the dashboards panel will be opened.

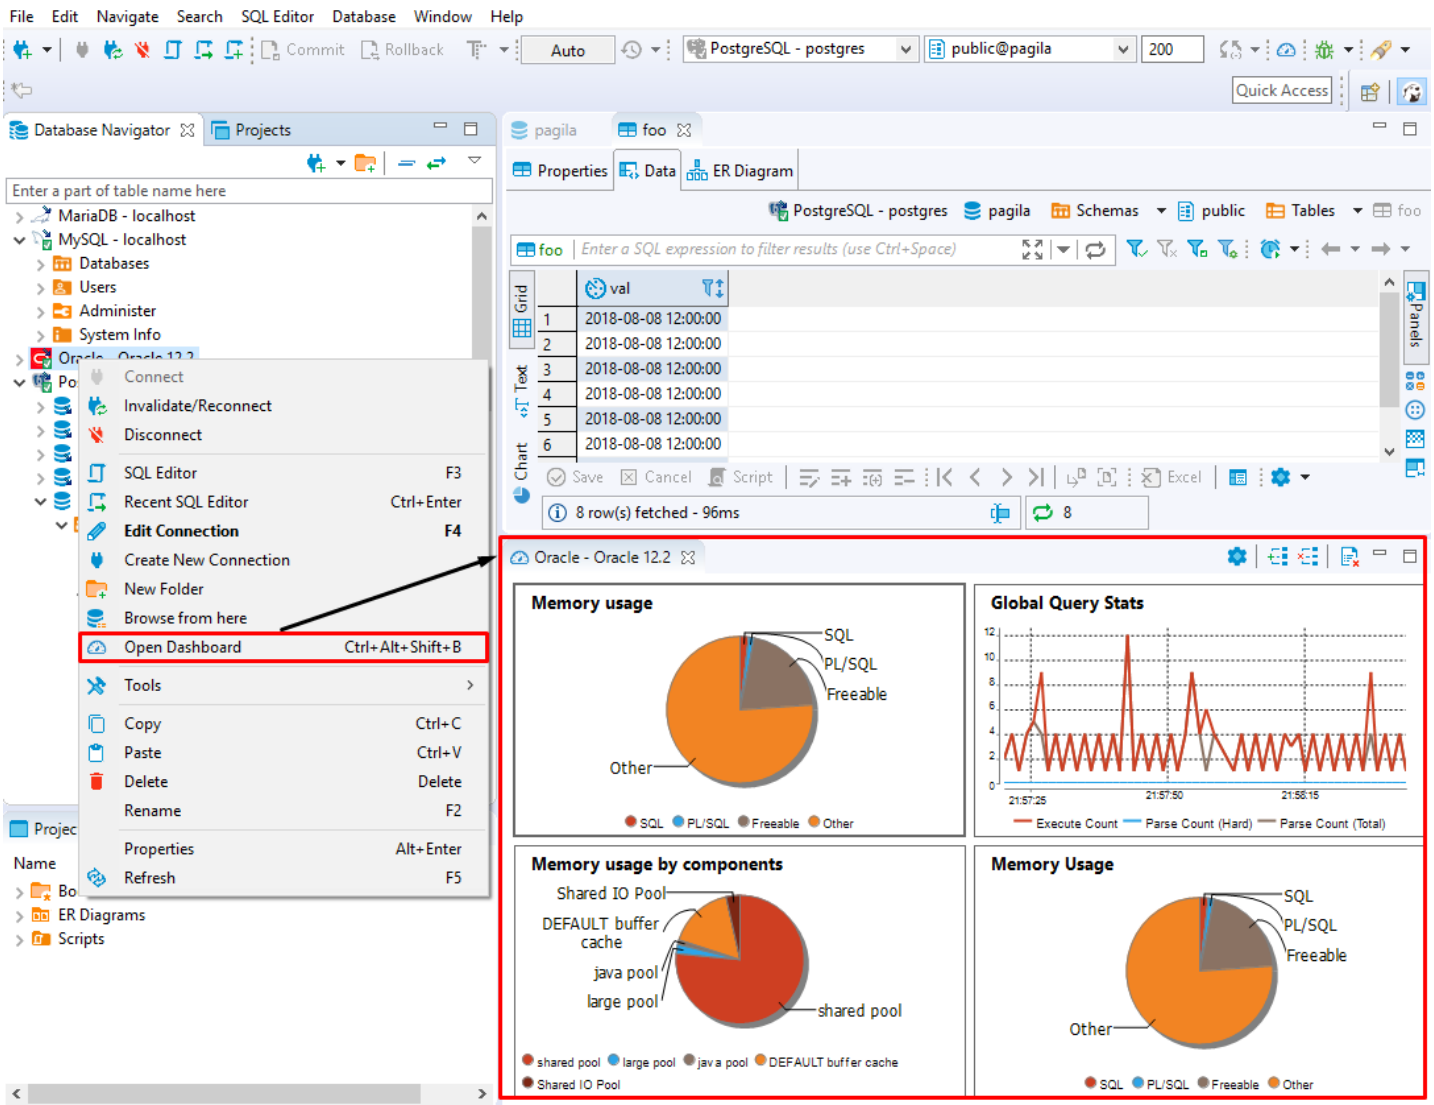

The following controls are available in the dashboards panel toolbar:

Icon	Name	Description
	Settings	Allows managing dashboards' configuration.
	Add dashboard	Allows to add dashboards to the dashboard panel.
	Remove dashboard	Allows to remove dashboards from the dashboard panel.
	Reset dashboards	Allows to restart dashboard calculation.

Adding Dashboards

To add a dashboard to the dashboards panel, press **Add dashboard** button in the dashboards panel's toolbar, choose one of the dashboards from the list of available dashboards and press **Add** button.

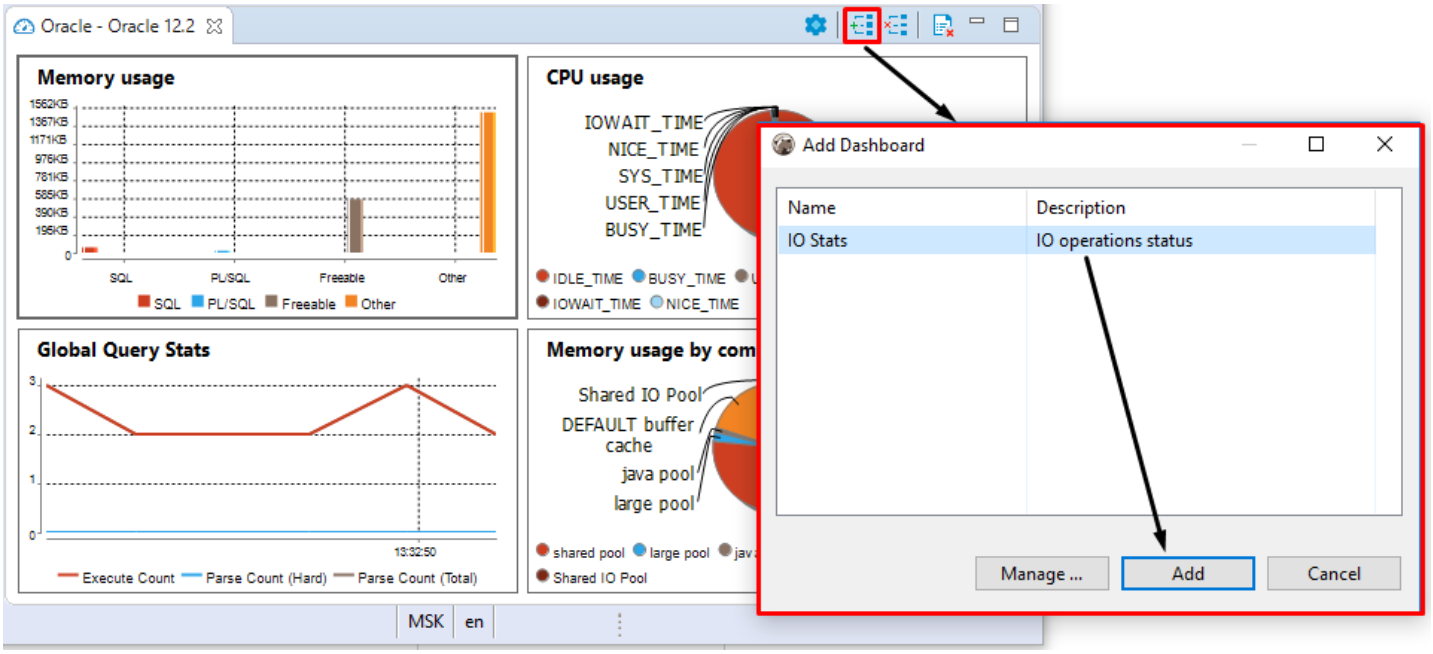

Note: Different databases have different sets of predefined dashboards. DBeaver is delivered with sets of predefined dashboards for such databases as Postgress SQL, MySQL, Oracle, and Exasol. It is also possible to create new custom dashboards, for more details see Managing Dashboards.

You can also add a dashboard by right-click in any place of the dashboards panel and then select the **Add dashboard** menu option.

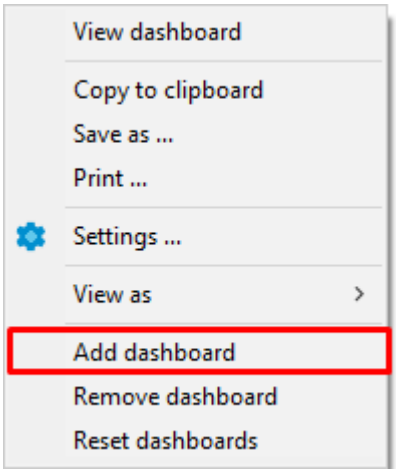

Removing Dashboards

To remove a dashboard from the dashboards panel, click on the dashboard you want to remove and press button **Remove dashboard** in the dashboards panel toolbar or select **Remove dashboard** option in the dashboard's context menu.

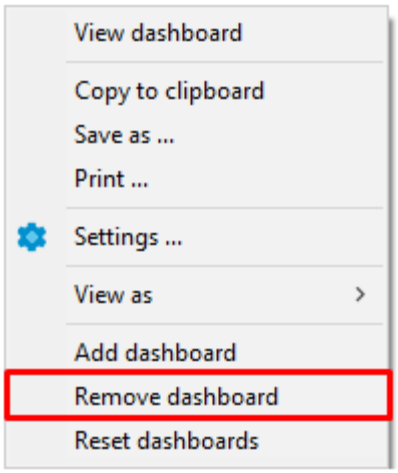

Resetting Dashboards

If you want to restart dashboard's calculation you can reset it.

You can reset all the dashboards displayed in the dashboards panel by a single click on **Reset dashboards** button in the dashboard panel's toolbar.

To reset a particular dashboard right-click on it and select **Reset dashboards** menu option or left click a dashboard and press **Reset dashboards** button in the dashboards panel's toolbar.

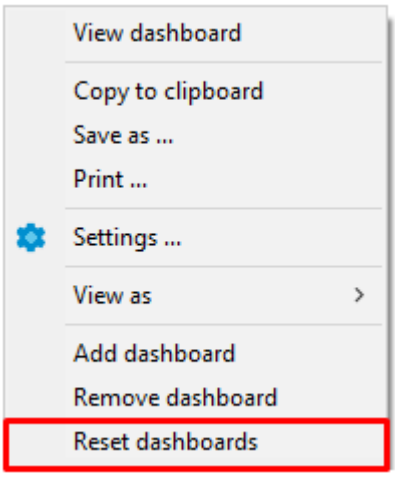

Changing Dashboard Representation

To adjust dashboard representation settings right click on a dashboard and select the **Settings** menu option, then, in the opened dialog change the parameters you want.

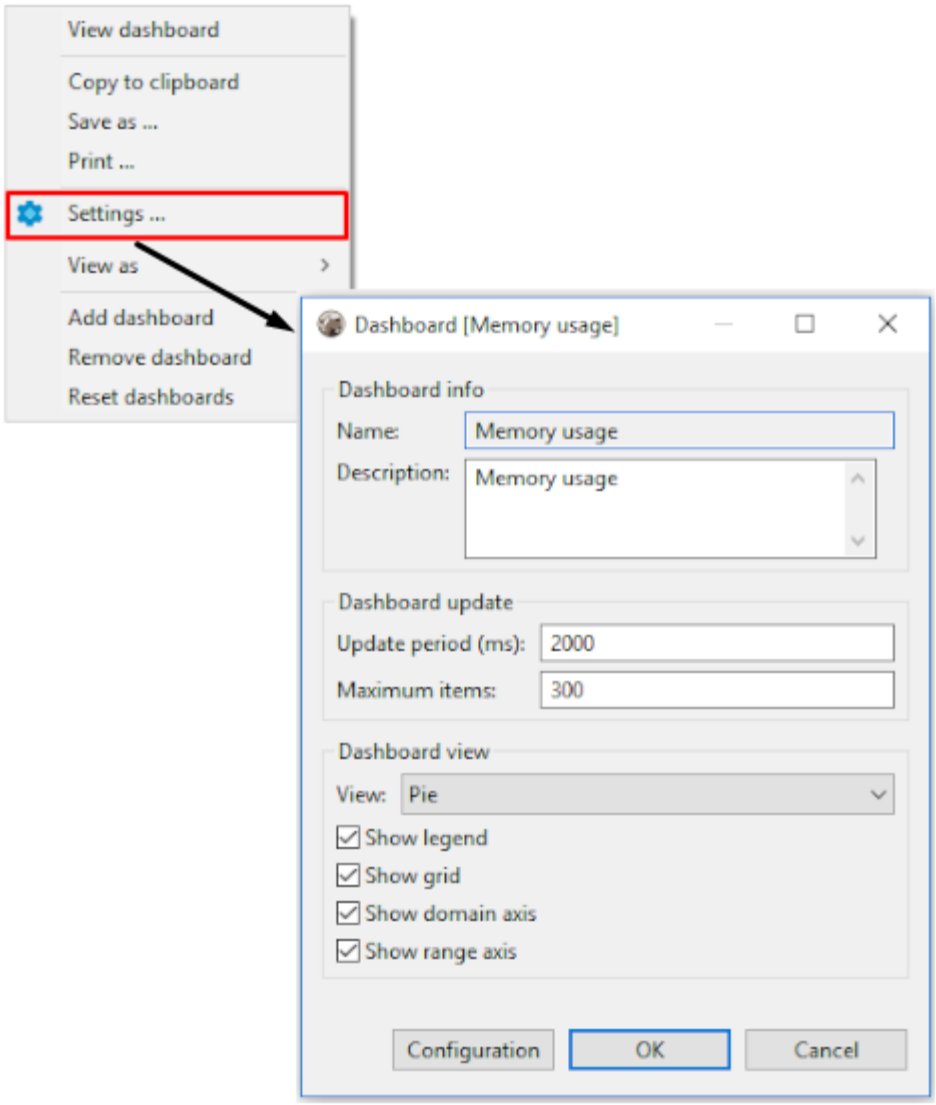

The following dashboard representation parameters can be adjusted:

Parameter	Description
Name	Defines a name of a dashboard.
Description	Defines dashboard's description. Use this field to make it easy to understand what kind of information the dashboard represents.
Update periods (ms)	Defines how often dashboard's rendering should be updated. The default value is 1000 ms.
Maximum items	Defines maximum number of fetched items. The default value is 300.
View	Defines visual representation of the dashboard. The following options are available: Bar, Pie, Time series.
Show legend	If this check-box is selected, the legend will be displayed on the dashboard.
Show grid	If this check-box is selected, the grid will be displayed on the dashboard.
Show domain axis	If this check-box is selected, the domain axis will be displayed on the dashboard.
Show range axis	If this check-box is selected, the range axis will be displayed on the dashboard.

Adjusting Dashboard Configuration

To adjust dashboard's configuration settings right-click on a dashboard, select the **Settings** menu option, then, in the opened dialog box press the **Configuration** menu option.

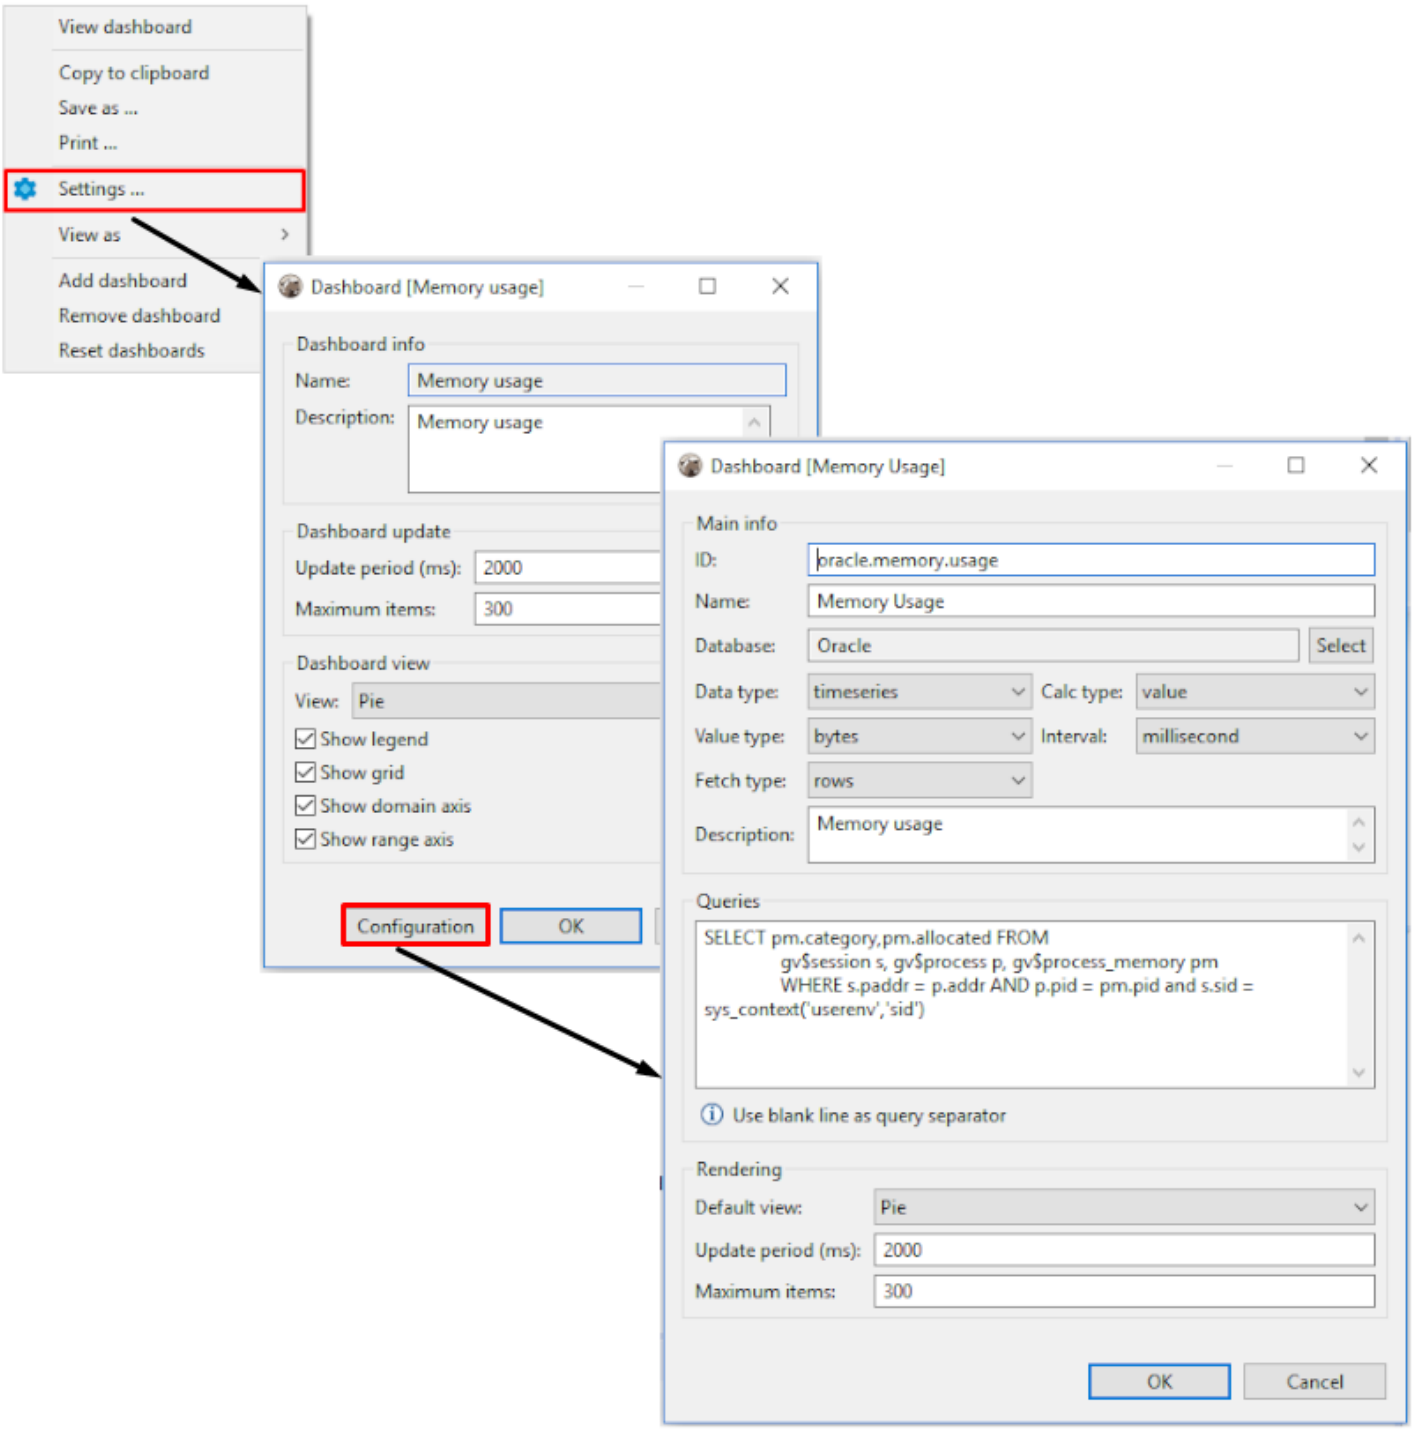

The following dashboard parameters can be configured:

Parameter	Description
ID	Defines dashboard's ID. Make sure that ID has numeric values in it.
Name	Defines dasboard's name.
Database	Defines the database driver. To learn moe about database drivers, see Database Drivers.
	Defines the data type. The following options are availabe: timerseries (the default option) and statistics. Select

Data type	timeseries type if you want to track the actual value returned by the server. Select statistics type if your dashboard will show historical data.
Calc type	Defines how the data should be calculated. The following options are available: value (the default option) and delta. Select value if you're interested in the current value. Select delta if you want to track the difference between the current value and the previous one. This may be very useful when you work with statistics data, for example.
Value type	Defines the value to be shown on the range domain. The following options are available: decimal (the default option), integer, percent, bytes. Choose the value type in accordance with your data, for example, memory usage is convinient to be tracked in KBytes.
Interval	Defines time interval to be shown on the domain axis. The following time intervals are available: millicecond(the default option), second, minute, hour, day, week, month, year.
Fetch type	Defines whether the query should fetch data from rows or columns.
Description	Defines the description of a dashboard. Use this field to make it easy to understand what kind of information the dashboard represents.
Queries	Defines an SQL query whose fetched data will be used to build the chart displayed on the dashboard.
Default view	Defines the default visual representation of a dashboard on the dashboard panel. The following options are available: Bar, Pie, Time series(the default option).
Update period(ms)	Defines how often the dashboard's rendering should be updated.
Maximum items	Defines maximum number of items to be fetched for the dashboard.

Note: Predefined dashboards are read-only and cannot be re-configured, but you can copy them and use as templates to create new dashboards with any query and other settings. To learn about creating new dashboards, see Managing Dashboards section.

Setting Connection Prefereces

By default, if there is no active connection to the database and you open its dashboards panel, all the dashboards on the panel will be empty.

You can force database connection on the dashboard panel's activation by pressing **Settings** button on the dashboards panel's toolbar and then selecting the **Connect on activation** check-box.

Detaching Dashboards

If you have several monitors and would like to place a dashboard into a separate screen, you can either detach the whole dashboards panel or a single dashboard and drag-and-drop them to any place you want.

To detach the whole dashboard panel right click on the dashboard's tab name and select the **Detach**menu option.

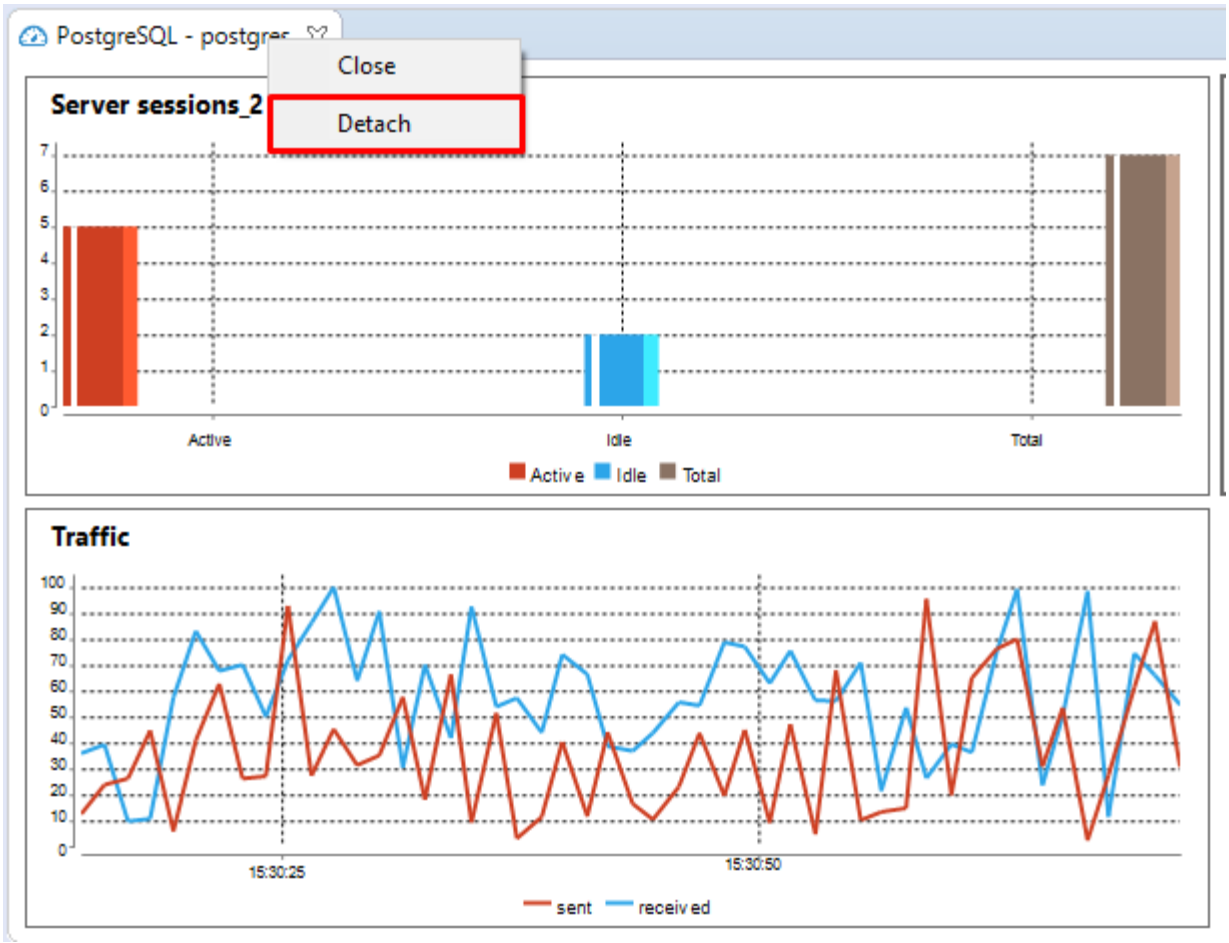

To detach a single dashboard make a double left click over it. You can also right click the dashboard and then, select the **View Dashboard** menu option, the dashboard will be detached from the panel and you will be able to move it to any place of your screen.

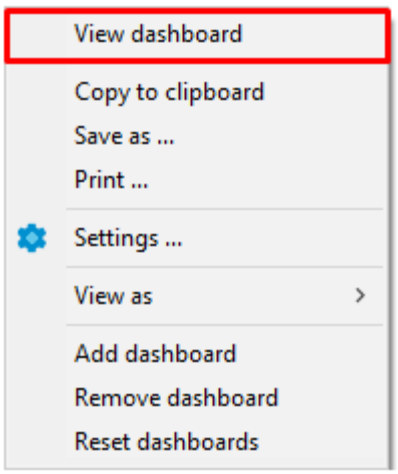

Changing Dashboard View

You can change the representation of a dashboard and view it as a Pie, Bar or Time series. To change the view of a dashboard, right click on it and select **View as** menu option.

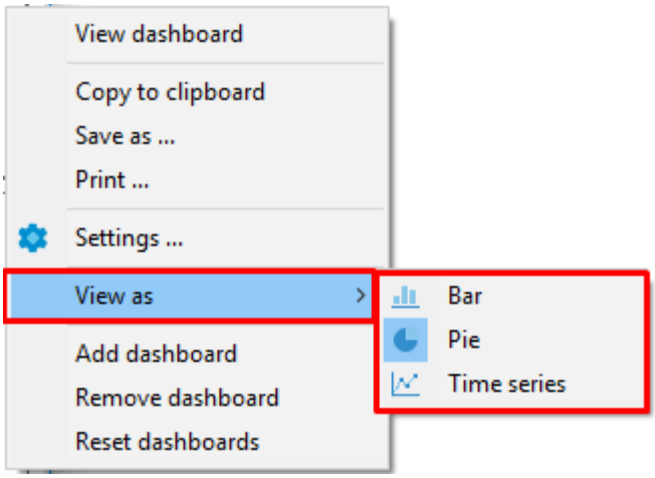

Copying Dashboards to Clipboard

To copy a dashboard into the clipboard, right click on the dashboard and use **Copy to Clipboard** menu option, the screenshot of the dashboard will be placed to the clipboard.

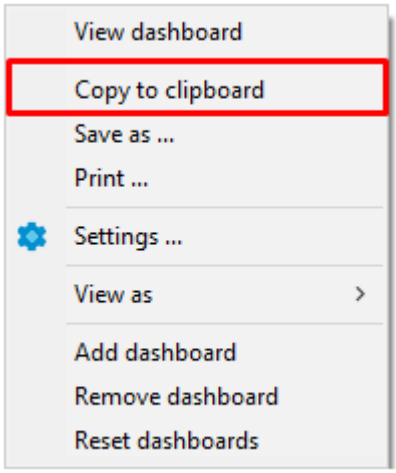

Saving Dashboards

If you want to save a screenshot of a dashboard locally in PNG format, right click on it and select the **Save as ...** option in the context menu displayed.

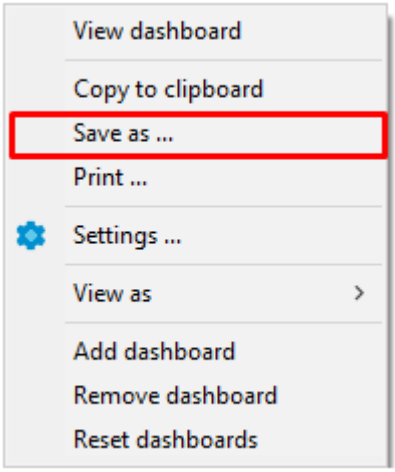

Printing Dashboards

If you want to print out a screenshot of a dashboard, right-click the dashboard to be printed and select the **Print…** option.

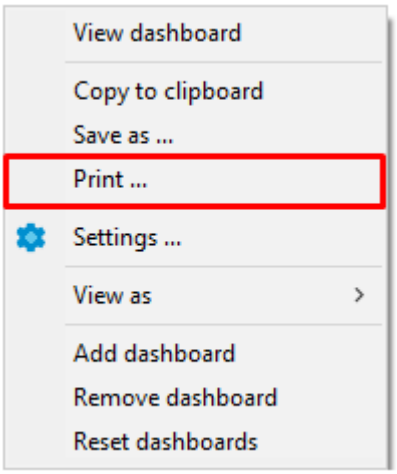

Zooming

For Time series and Bar dashboard representations the following zooming options are available in the dashboard's context menu:

- Zoom In
- Zoom Out
- Zoom Reset

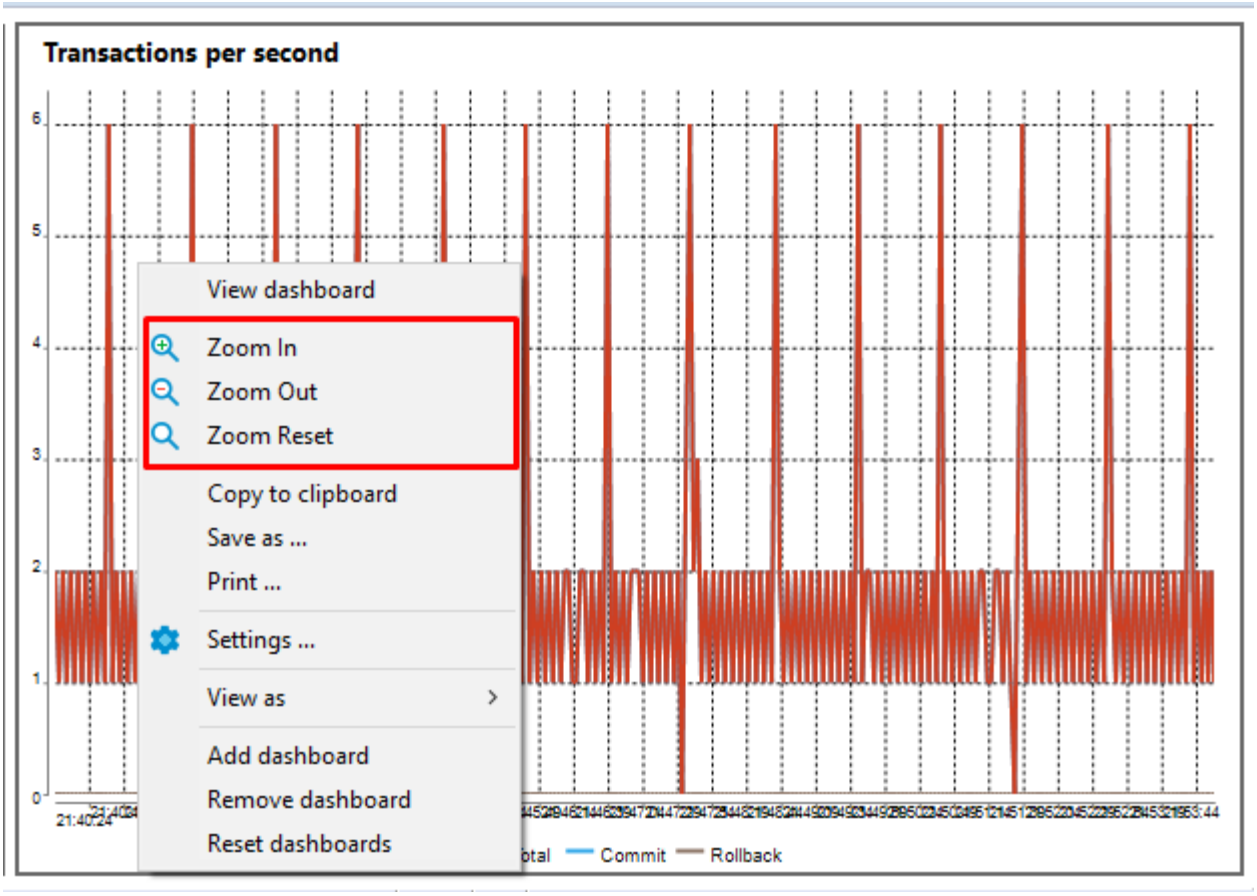

Managing Dashboards

You can extend the list of predefined default dashboards by creating your own custom dashboards. This section describes dashboards' list management.

Creating Dashboards

You can create a new custom dashboard either from scratch or from any of already existing dashboards.

To create a dashboard from scratch:

1. Press **Settings** button in the dashboards panel toolbar.
2. In the opened dialog box click **Manage...** button.
3. In the **Manage dashboards** window click **New dashboard...** button.
4. Set up all configurational parameters as required and press **OK**. To learn more about dashboard's configuration parameters, see Adjusting Dashboard Configuration.

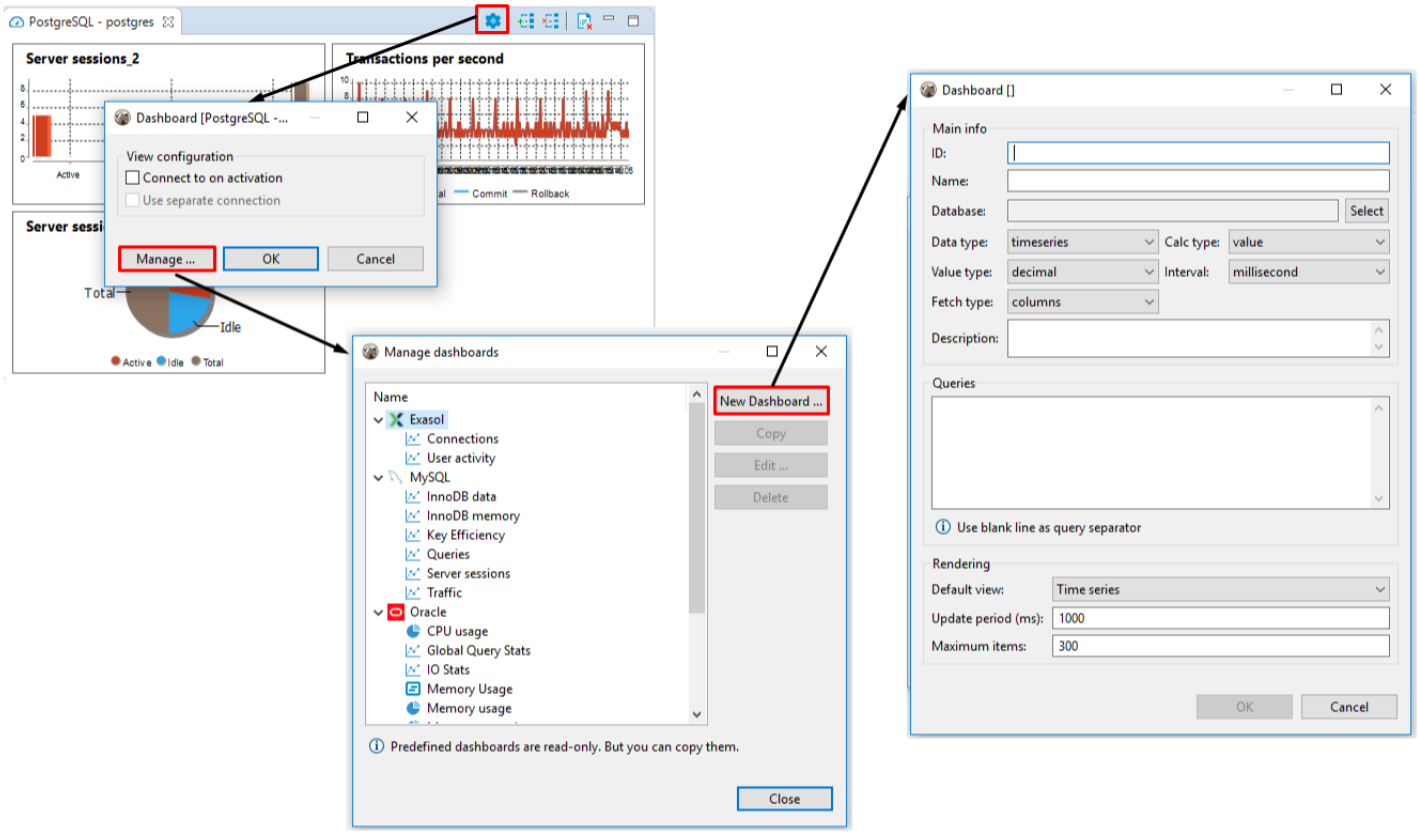

To create a dashboard from template:

1. Press **Settings** button in the dashboards panel toolbar.
2. In the opened dialog box click **Manage...** button.
3. In the **Manage dashboards** window select any of the existing dashboards from the list and click **Copy**.
4. Adjust all configurational parameters as required and press **OK**. To learn more about dashboard's configuration parameters, see Adjusting Dashboard Configuration.

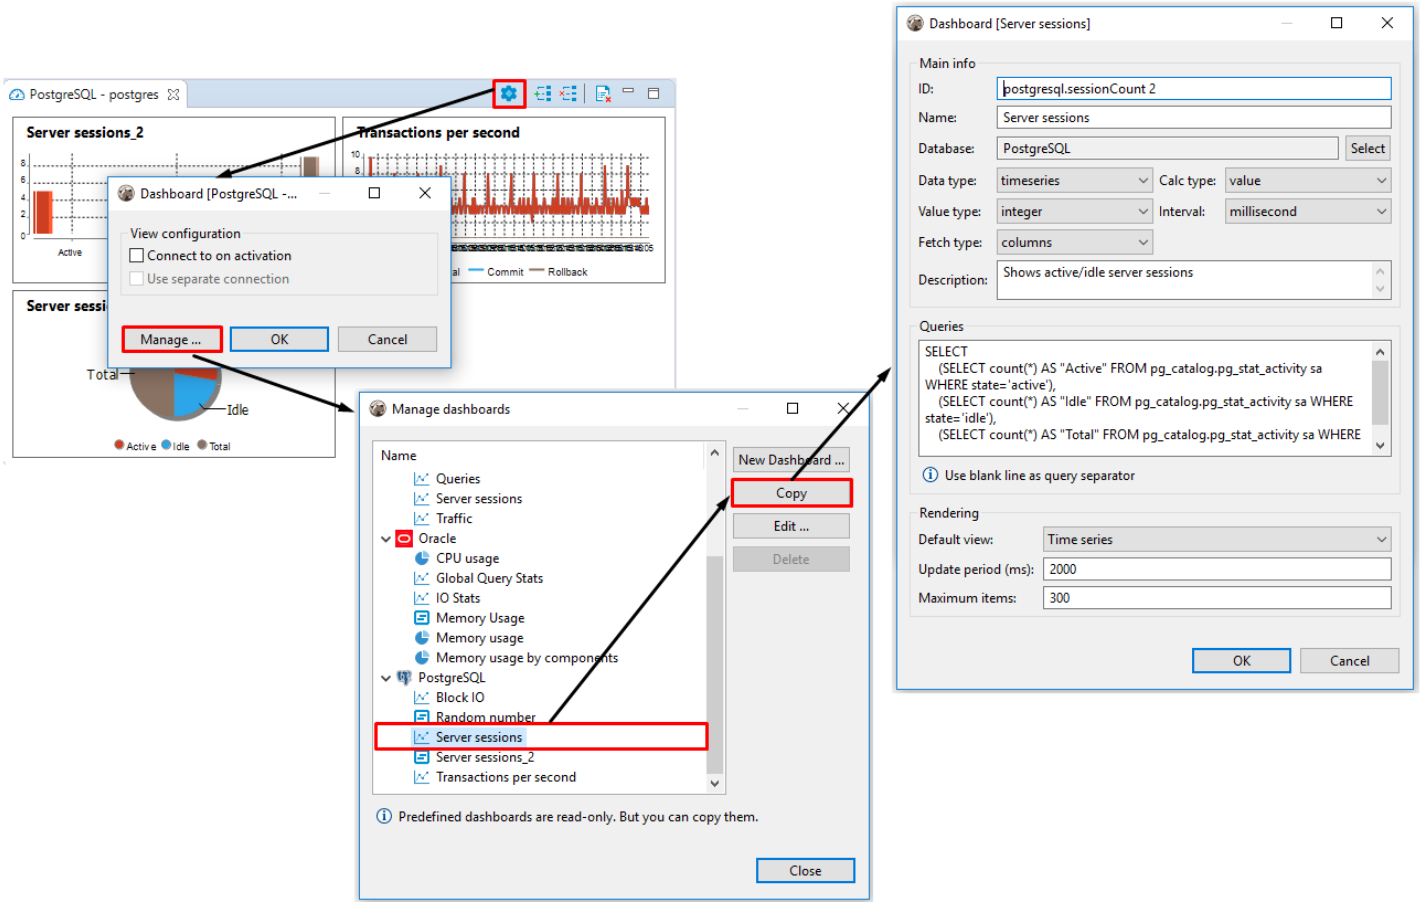

Editing Dashboards

If you need to change dashboard's name , ID or any other configurational setting, you can edit a dashboard.

Note: Only custom dashboards can be edited, predefined dashboards are read-only, but you can use them as templates and create a custom dashboard whose parameters will be editable. To learn how to create dashboards from templates, see Creating Dashboards.

To edit dashboard's configuration:

1. Press **Settings** button in the dashboards panel toolbar.
2. In the opened dialog box click **Manage...** button.
3. In the **Manage dashboards** window select any of the custom dashboards from the list and click **Edit...**.
4. Adjust all configurational parameters as required and press **OK**. To learn more about dashboard's configuration parameters, see Adjusting Dashboard Configuration.

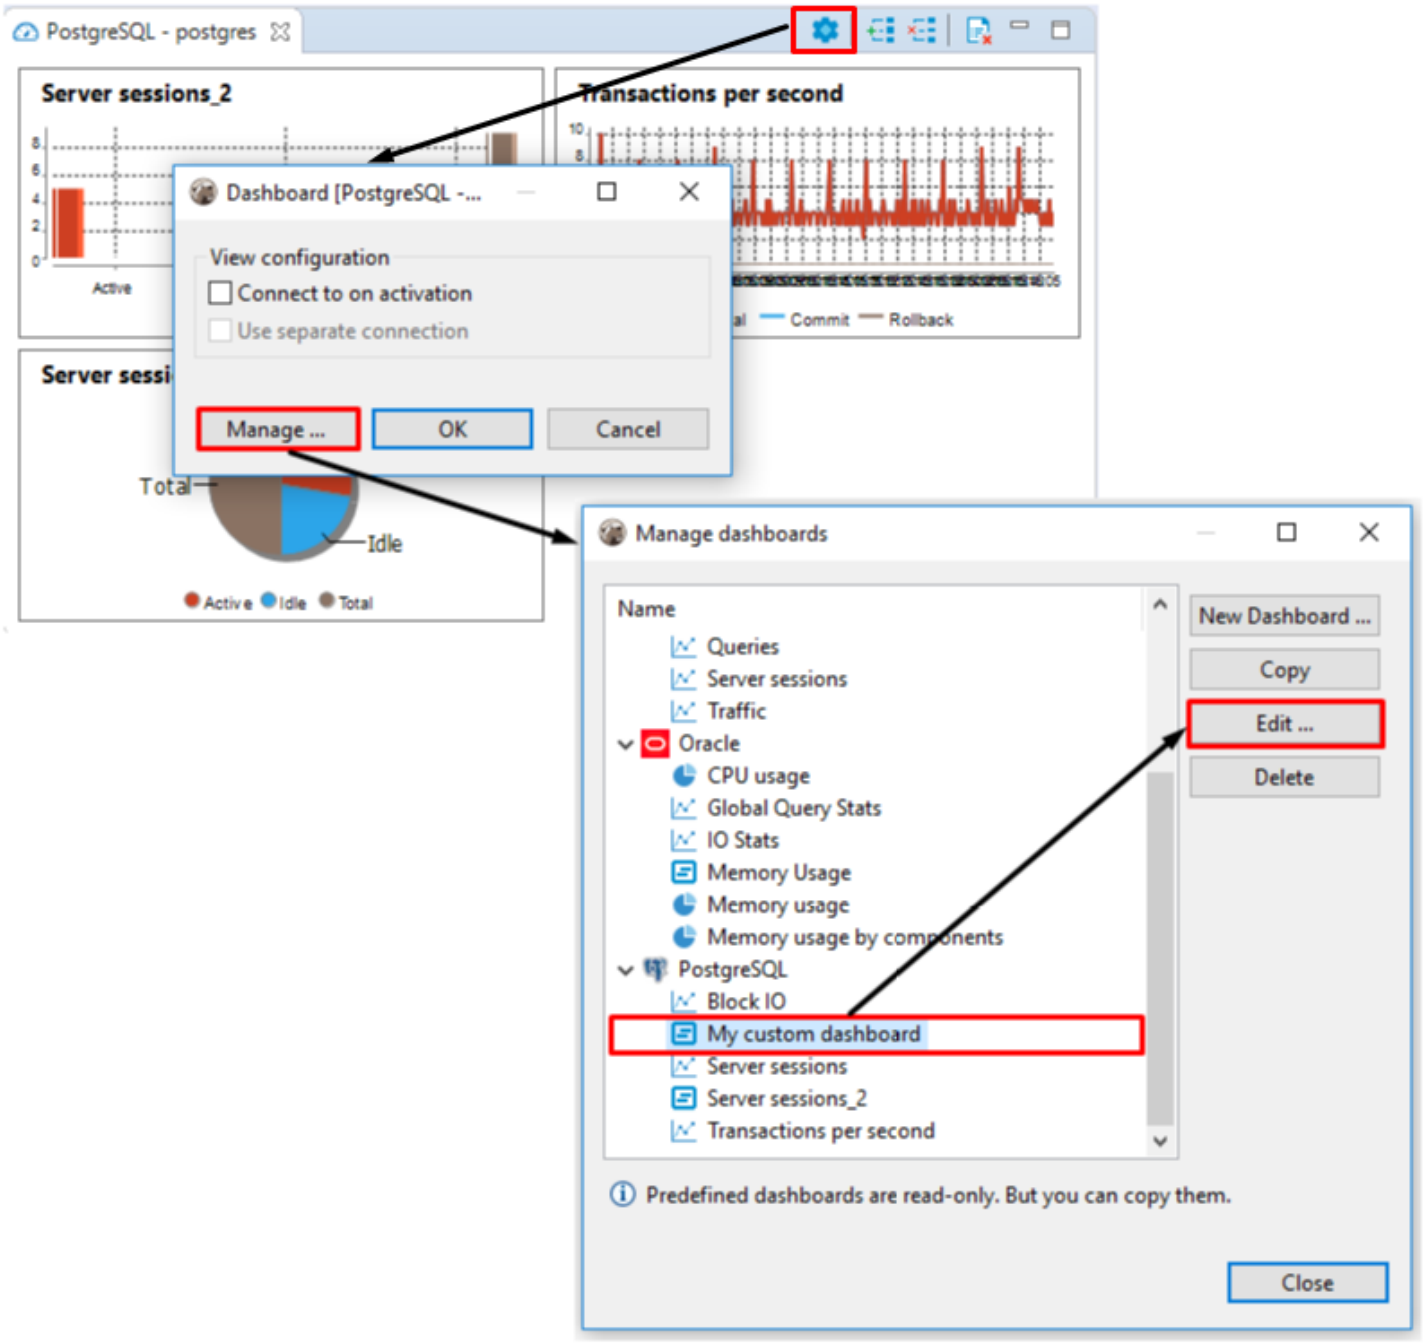

Deleting Dashboards

Note: Predefined dashboards cannot be deleted, but any of the custom dashboards can be deleted.

If you want to delete a dashboard, follow the steps described below.

To delete a dashboard:

1. Press **Settings** button in the dashboards panel toolbar.
2. In the opened dialog box click **Manage...** dashboards.
3. In the **Manage dashboards** window select any of the custom dashboards from the list and click **Delete**.

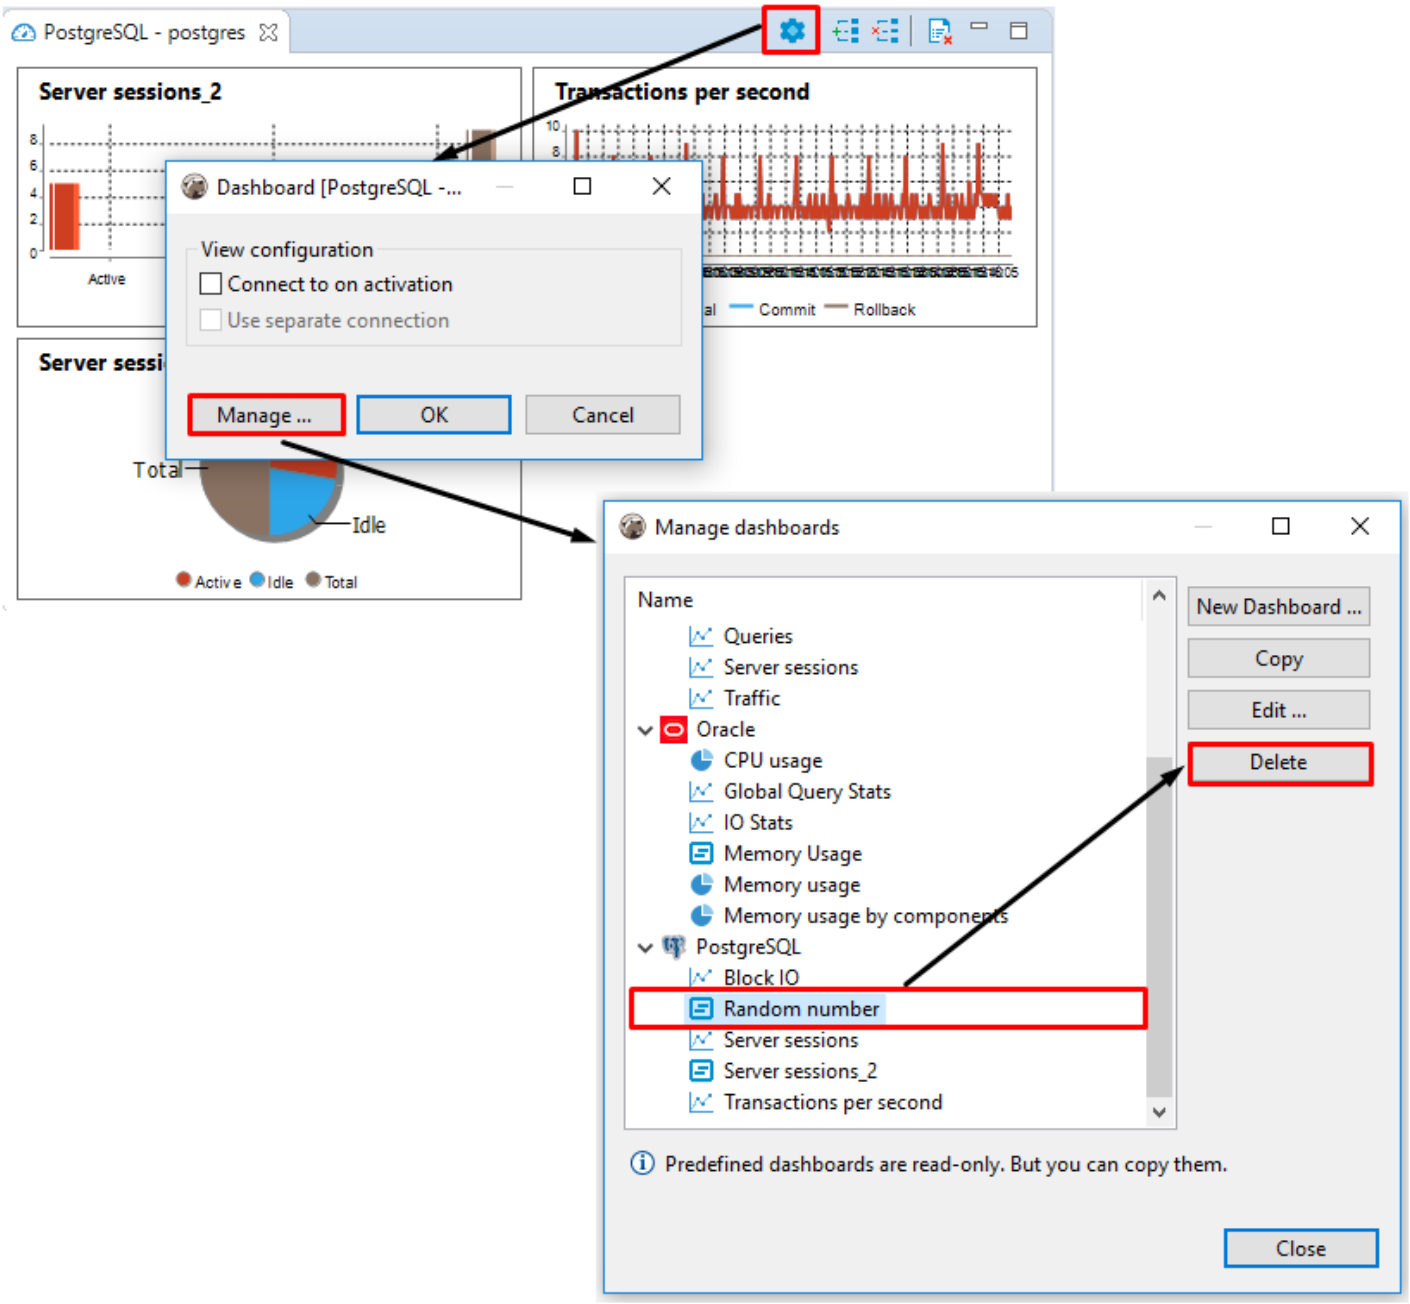
PostgreSQL - postgres
Server sessions_2
Active
Transactions per second
Commit
Rollback
Dashboard [PostgreSQL -...
View configuration
Connect to on activation
Use separate connection
Manage ...
OK
Cancel
Total
Idle
Active
Idle
Total
Manage dashboards
Name
Queries
Server sessions
Traffic
Oracle
CPU usage
Global Query Stats
IO Stats
Memory Usage
Memory usage
Memory usage by components
PostgreSQL
Block IO
Random number
Server sessions
Server sessions_2
Transactions per second
New Dashboard ...
Copy
Edit ...
Delete
Predefined dashboards are read-only. But you can copy them.
Close

Projects

The Projects view allows creating new projects as well as renaming and deleting projects that are not active.
NOTE: You cannot rename or delete a project that is set as active.

Creating Project

To create a project, in the Projects view, in the toolbar, click **Create Project** ( ). The Project Create Wizard opens.

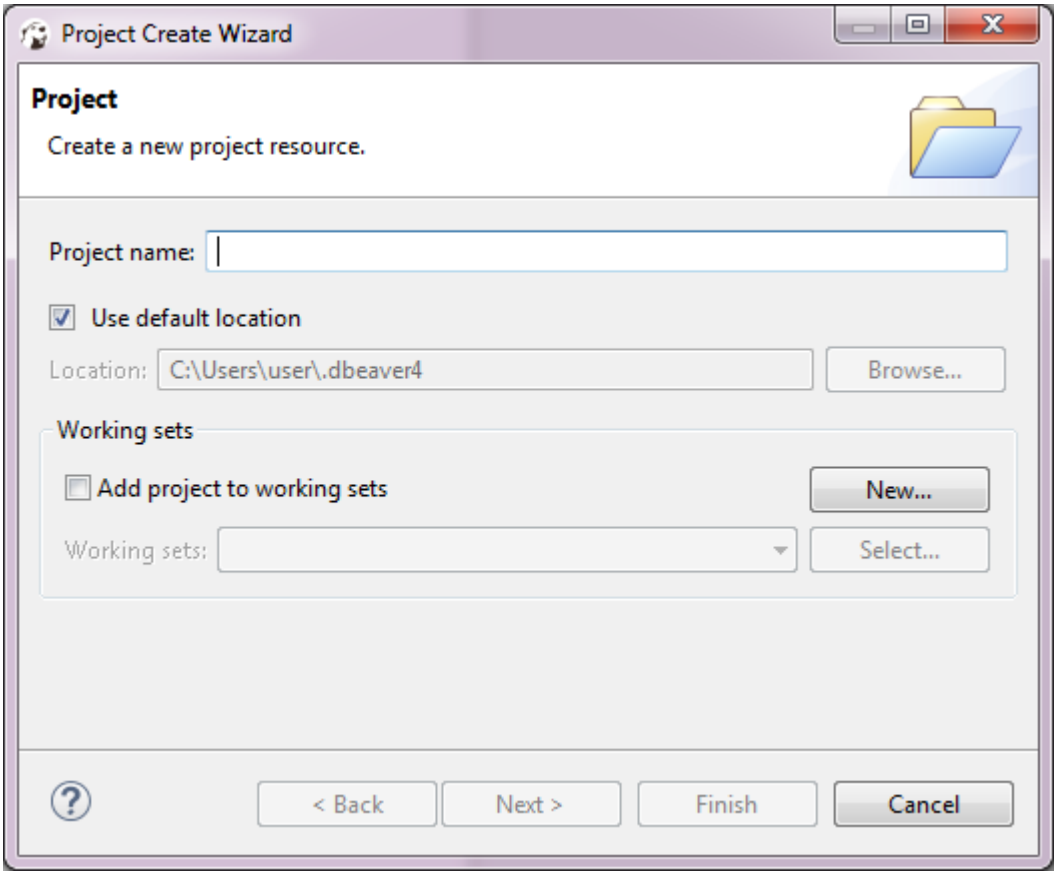

1. In the Project screen, in the **Project name** field, specify the name of the project.
2. To keep the default location to store the project, leave the **Use default location** checkbox selected. If you want to change the location, clear the checkbox and enter the name of the new directory into the **Location** field or click **Browse** and select the directory in the folder tree.
3. Click **Finish**. The new project appears in the projects tree.

Deleting Project

To delete a project, in the Projects view, right-click its name in the tree and click **Delete** on the context menu. Two confirmation dialog boxes appear one after another:

1. **Delete object** dialog box is to confirm the deletion of the project itself. Click **Yes** if you are sure you want to delete it. Otherwise, click **No**.

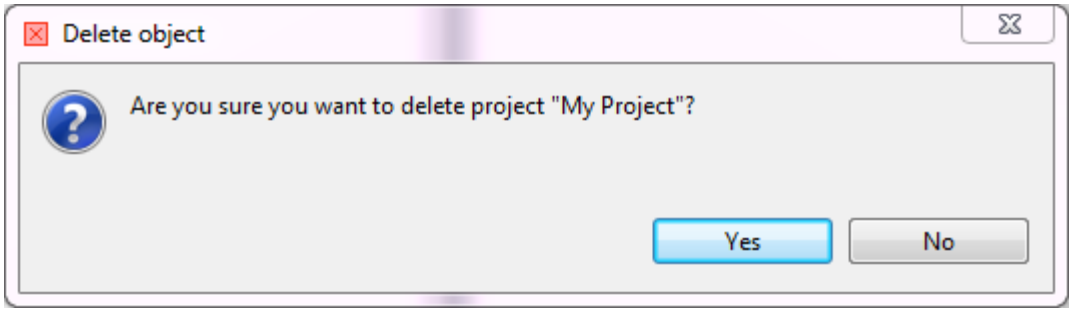

2. **Delete project** dialog box is to confirm the deletion of the project`s contents: these are the data stored in the file system, database connections are not affected. Click **Yes** if you want the contents to be deleted as well. To keep the contents, click **No**.

NOTE: If you have deleted a project and then re-create it with the same name, the new project picks up all the database connections of the deleted project.

Project security

Note: This functionality is available only in Enterprise-Edition.

DBeaver support local storage for connection secure data. It includes:

- Database server user credentials
- SSH tunnel user credentials
- Proxy user credentials

By default user names and passwords are stored in file `credentials-config.json`. This file is encrypted using AES key. However it is really secure as this key is not secure (can be found in DBeaver sources) and thus this file can be un-encrypted by 3rd party people using some 3rd party software.

In DBeaver Enterprise there is much more strong security support.

Master password for local configuration

It is possible to set master password for all projects in local workspace. Go to Preferences->Database->Security and enable option `Use secure passwords storage`. There are several password storage providers (you can see them on page General->Security->Secure Storage), `DBeaver Enterprise Password Provider` is the default one (in standalone DBeaver). It will ask you to specify master password. DBeaver doesn't store this password anywhere, it only encrypts user credentials in special local storage. It is not possible to decrypt these password without password (at least easily).

Side effect of this configuration - you cannot share your connections (with password) between different users. Because user credentials are stored in a completely separate location and they are protected by local user password.

Use Windows Integration password provider

You can disable default password provider and enable "Windows Integration" provider. This provider doesn't need master password but it uses randomly generated password stored in local user secure storage (in Windows). This is easier (as you don't need to remember master password) but less secure (anybody who have access to your Windows user account will have access to DBeaver stored credentials).

Project password

You may specify a password for a project. It will encrypt all project's configurations with this password. Also you will be able to share your project settings with other users (you will need to pass the project password as well).

In order to enable project password open project properties. You can do this by:

- Click on main menu File->Project security
- Click on "Configure" icon in project explorer view toolbar and switch to Project Security tab
- Press `ALT+Enter` on a project element in `Projects` view and switch to Project Security tab

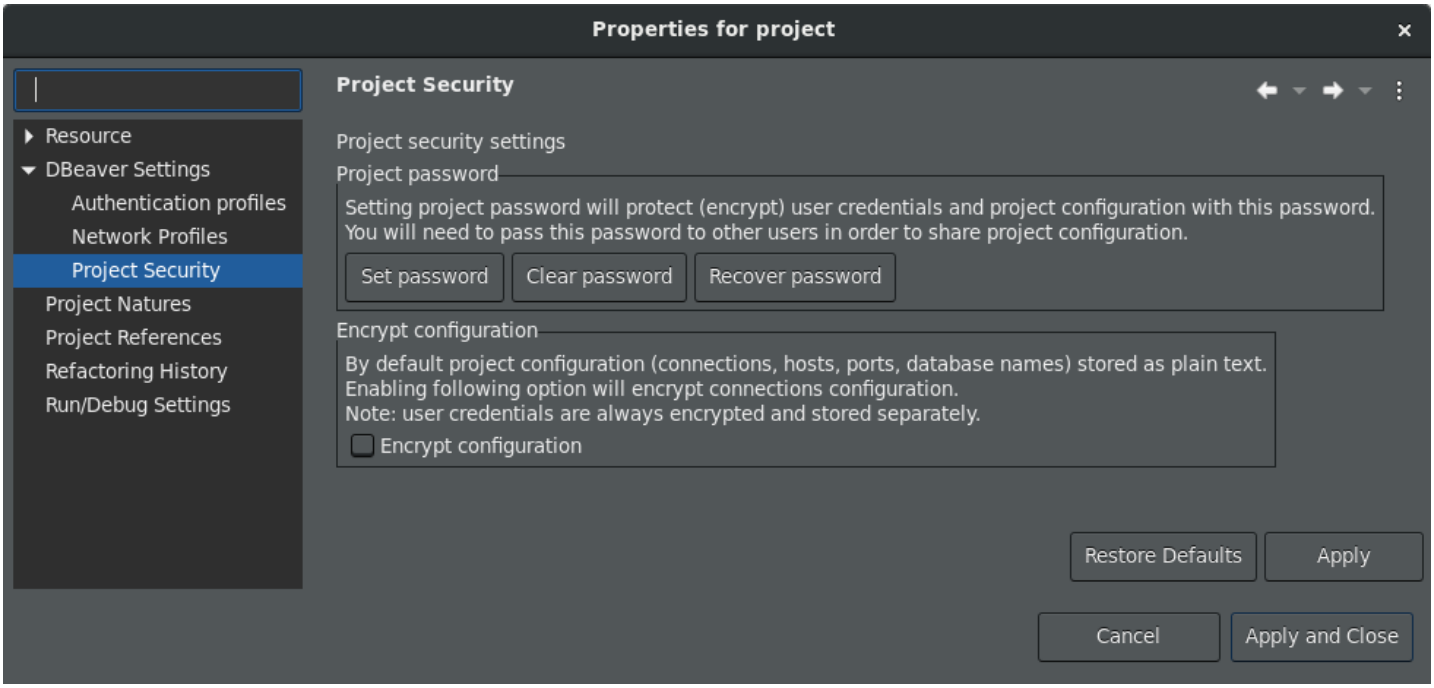

On project security page click on "Set Password" button to enable project password. Click on Clear to disable it (you will need to enter current project password to clear it).

"Encrypt configuration" option

Team work (Git)

DBeaver is fully integrated with Git team work/version control system. You can keep your project configuration, scripts, diagrams, bookmarks and other artifacts in a Git repository.

Managing connection configuration

Managing secured resources

Managing scripts and other resources

Bookmarks

Bookmarks are quick access links to objects of a database. They appear in the project tree inside the Projects or Project Explorer views.

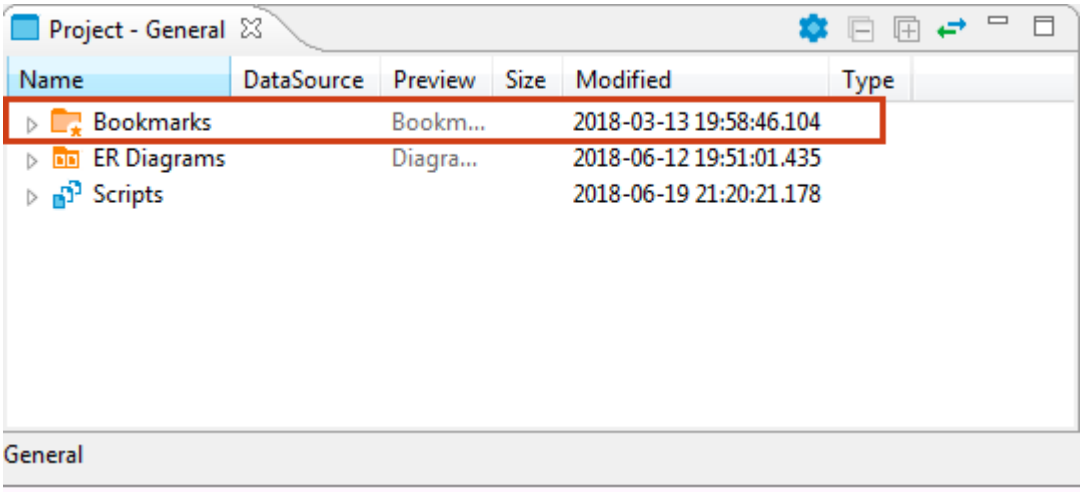

To create a bookmark:

1. In the Database Navigator or under **Connections** node of the Projects view, click the database object of interest to set focus on it.
2. Press `CTRL+d` . The **Bookmark Name** dialog box appears.
3. In the **Bookmark Name** field, enter the bookmark name, then in the **Bookmark folder** field, click the folder, and then click **OK**:

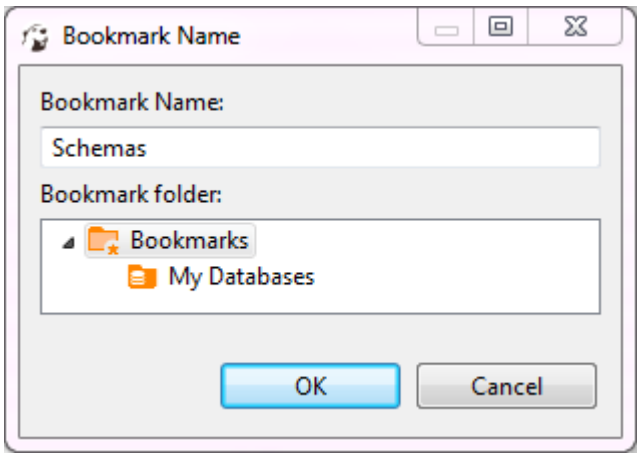

The bookmark appears in the selected folder of the related project.

To open an object using its bookmark, double-click the bookmark or right-click it and click **Open Bookmark** on the context menu. You can rename and delete bookmarks using the context menu as well.

Shortcuts

Brief list of the most important DBeaver shortcuts.
Of course you can redefine any (or almost any) of these shortcuts, here is the list of the default values.
Most of the following commands are accessible from DBeaver main menu, context menu or editor toolbar (or from all of them).
Use context menu wherever it is possible - it usually shows all actions accessible at this moment.

SQL Editor

Shortcut	Action
`CTRL+Enter`	Execute current query (*)
`CTRL+\`	Execute current query in a new tab
`ALT+X`	Execute current script (**)
`CTRL+ALT+'`	Execute selected SQL expression and print results
`CTRL+SHIFT+E`	Explain current query execution plan
`CTRL+ALT+SHIFT+X`	Execute queries of current script simultaneously, showing results in separate tabs
`CTRL+9`	Switch active connection (for SQL script)
`CTRL+Space` `Option+Space`	SQL completion proposals popup
`CTRL+ALT+Space`	SQL templates proposals popup
`CTRL+SHIFT+F`	Format current script (**) using current formatter
`CTRL+/` `CTRL+SHIFT+/`	Toggle single/multi line comment
`ALT+Up` `ALT+Down`	Jump to previous/next query
`CTRL+6` `CTRL+SHIFT+6` `ALT+6`	Toggle editor/results panels (maximize/minimize/switch)
`CTRL+SHIFT+X` `CTRL+SHIFT+Y`	Convert selected text into upper/lower case

Data viewer

Shortcut	Action
`TAB`	Switch to record/grid mode
`CTRL+~`	Switch presentation (grid, plain text, json ,etc)
`CTRL+1`	Foreign keys navigation menu
`ALT+Space`	Navigate to the link in active cell
`ALT+Left`	Navigate backward in history
`ALT+Right`	Navigate forward in history
`CTRL+2`	Toggle sorting by current column

F11	Current column filters menu
CTRL+F11	Current column filter dictionary panel
F7 CTRL+7	Toggle right panels on/off
F5	Refresh results (re-run query)

Data editor

Shortcut	Action
Enter	Activate inline editor
SHIFT+Enter	Open value editor dialog or separate value editor (for LOB values)
Delete ALT+Delete	Delete row
ALT+Insert	Add new row
CTRL+ALT+Insert	Copy current row
Escape	Cancel changes in current cell/row

Database Navigator

Shortcut	Action
F2	Rename current element (if supported)
F4	Open editor of selected element(s)
F5	Refresh selected element(s)
Delete	Delete selected element(s) (if supported)
CTRL+ALT+SHIFT+D	Add bookmark on selected element
Alt+Enter	Show properties of selected element
F3 CTRL+[	Open SQL editor for current connection (***). Shows script selector popup.
CTRL+F3 CTRL+]	Open new SQL editor for current connection (***). Always creates new script.
CTRL+Enter	Open recent SQL editor for current connection (***). Opens last modified script or creates a new script.

Other

Shortcut	Action
ALT+~	Shows database tools context menu
CTRL+0	Switch active schema/catalog (available if SQL/database editor is open)
CTRL+SHIFT+C	Advanced copy. Works in different contexts and performs "smart copy" operation (usually with parameters).
CTRL+SHIFT+V	Advanced paste. Same as "smart copy" but for "paste".

References

`*` - Current query is the query under cursor or the selected text. Query is separated from other script queries by delimiter (; by default) or by empty lines.
`**` - Current script is a set of all queries in the current SQL file. If there is a text selection then only queries in this selection are processed. Queries are separated from each other with a delimiter (; by default).
`**` - Current connection detected from active window and selection. If active (focused) window is SQL editor or database object editor then current connection is the same as in this editor. If active window is database navigator then active connection is "owner" connection of currently selected element. In other cases there is no current connection and DBeaver will ask you to choose connection explicitly.

Database Connections

To be able to manage your database in DBeaver, you need to create a connection to this database – see Creating Connections. A connection includes a driver and a number of configuration parameters including the location of the database and credentials to access it. You need to create a separate connection to every database you want to manage. Every database type requires its own set of connection parameters.

Connections reside in the Database Navigator and in the Projects views. In these views, you can:

- Edit connections, see Editing Connections
- Rename and delete connections - via corresponding context menu items, see Database Navigator
- Connect to and disconnect from databases using connections, see Connect to Database and Disconnect from Database.

Database connections might have the following states:

- not connected
- has network settings specified (such as SSH tunnel, etc.)
- connected
- connection error

Create Connection

DBeaver provides a wizard that guides you through the steps to create a connection. If you run DBeaver for the first time (standalone version), the new connection wizard appears automatically. In other cases, to create a connection, do one of the following:

- Click the **New Connection Wizard** button in the application toolbar or in the Database Navigator view toolbar:

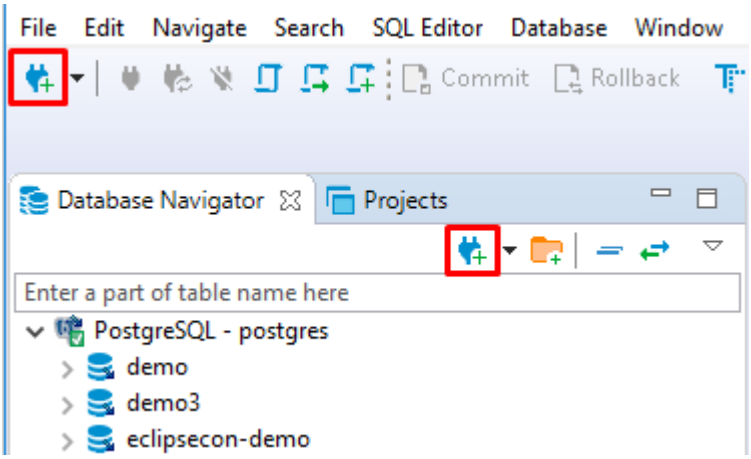

- Click **Database -> New Connection** in the menu bar:

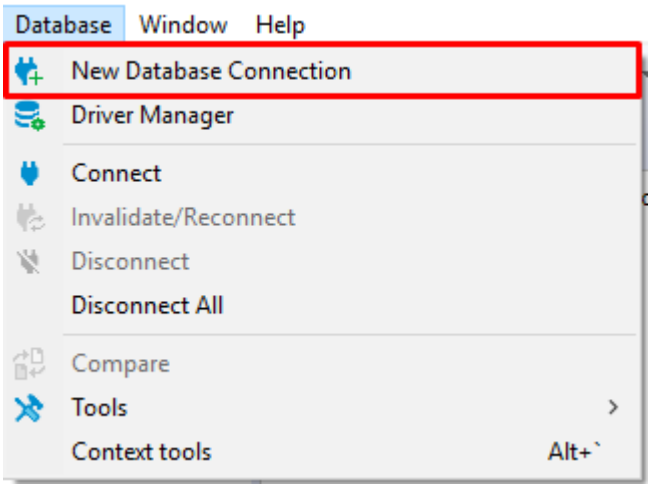

- Press `Ctrl+N` or click **File -> New** in the menu bar:

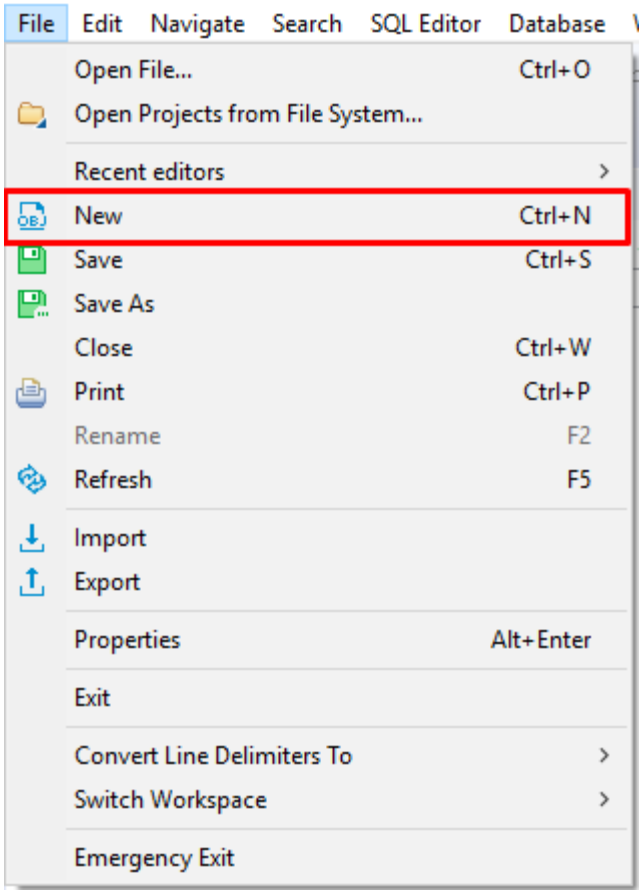

Then, in the wizard, click **Database connection** and then click **Next**:

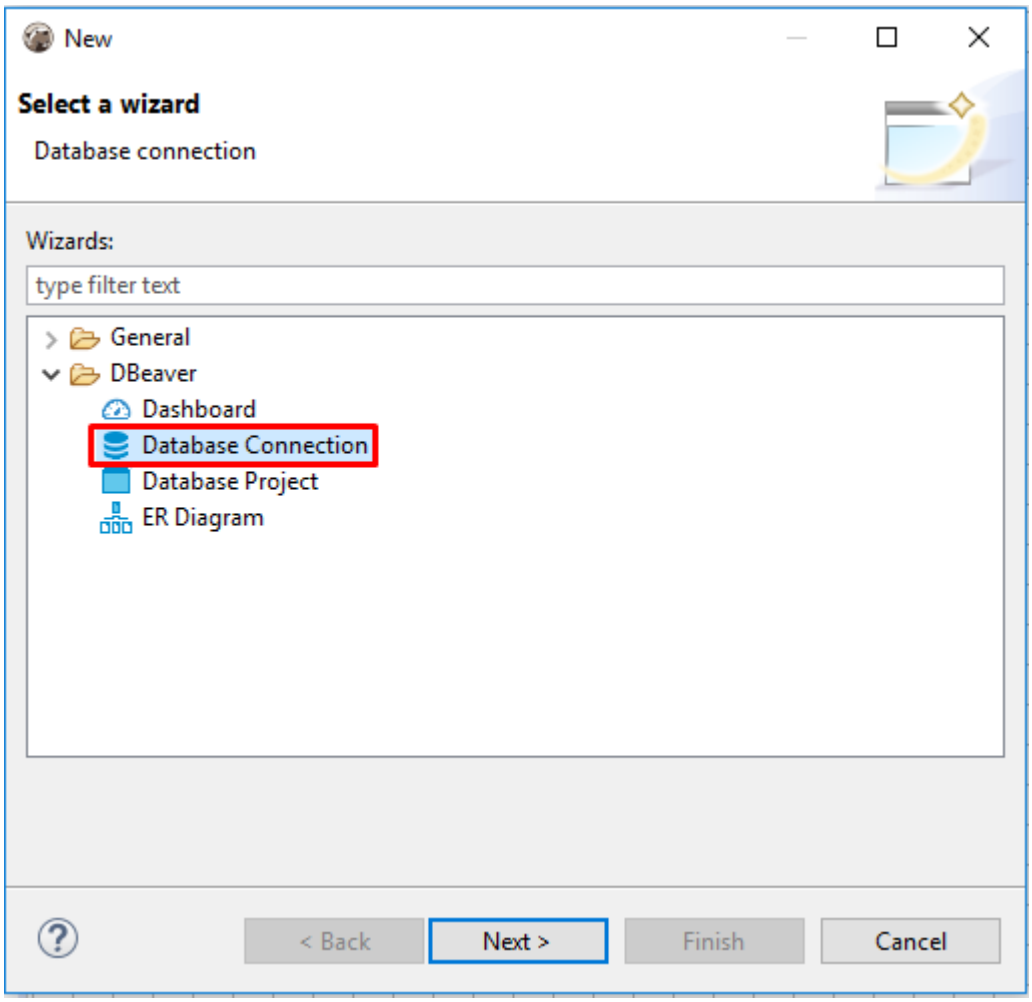

Then, in the **Create new connection** wizard:

1. Choose a driver for the new connection: click the name of the suitable database type in the gallery. Then click **Next**.

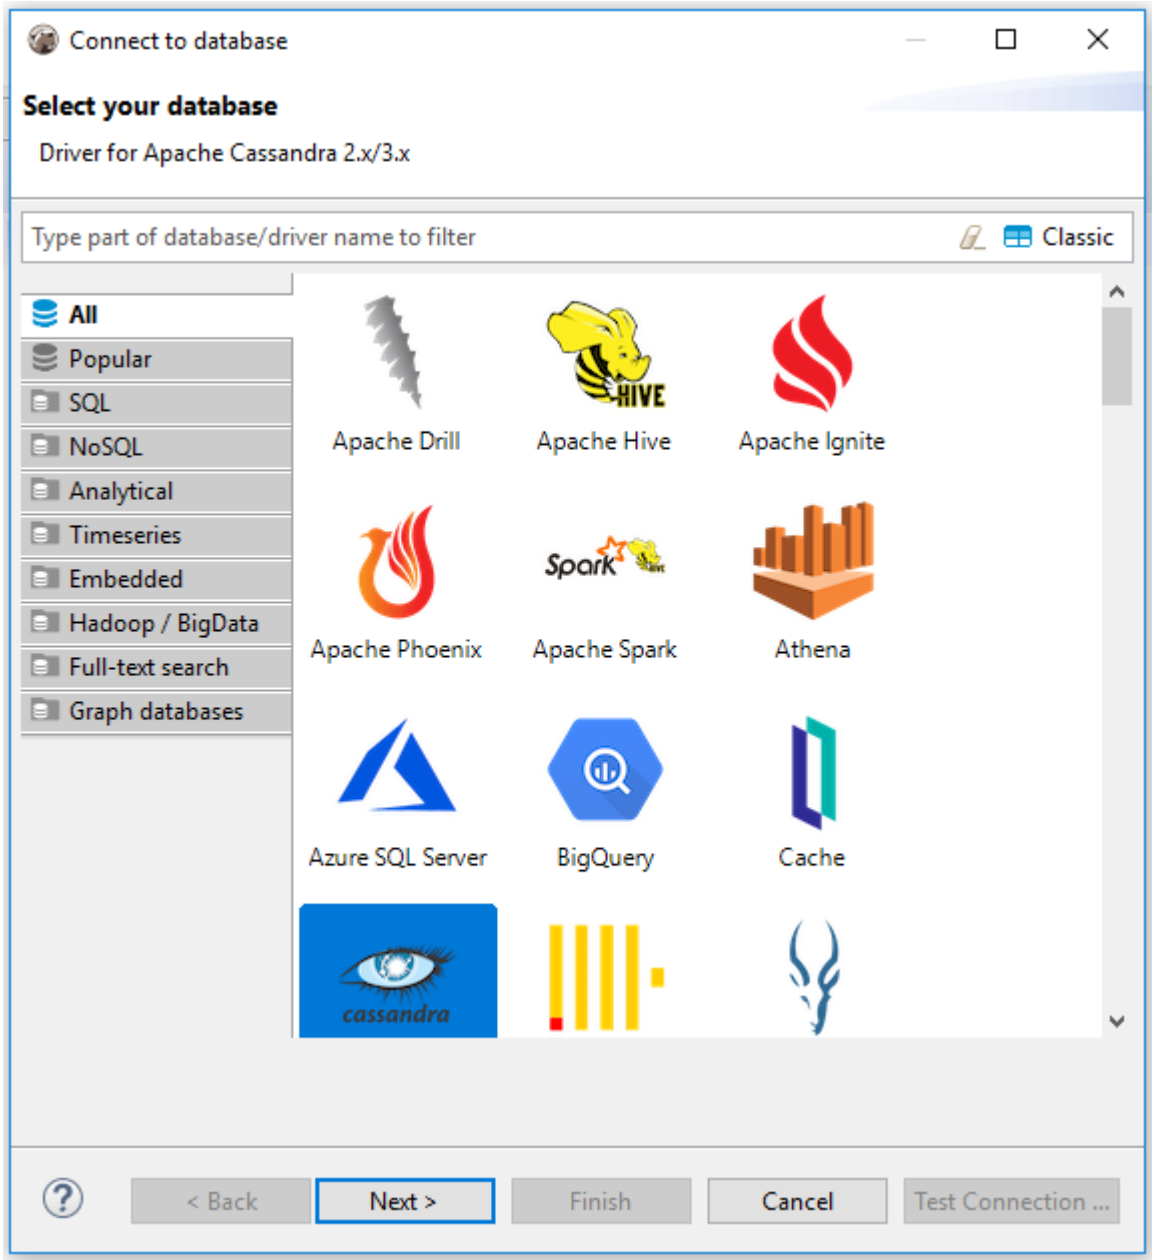

   To quickly find the needed driver, you can type a hint in the text field above the list of drivers.
   If you cannot find a driver for your database then probably there is no suitable driver and you need to create one. Please see Database Drivers article.

   NOTE: The list of database drivers diaplays the number of exising connections next to each driver. No number is displayed if there are no connections.

If you prefer the classic list view of the available drivers, use the **Classic** button.

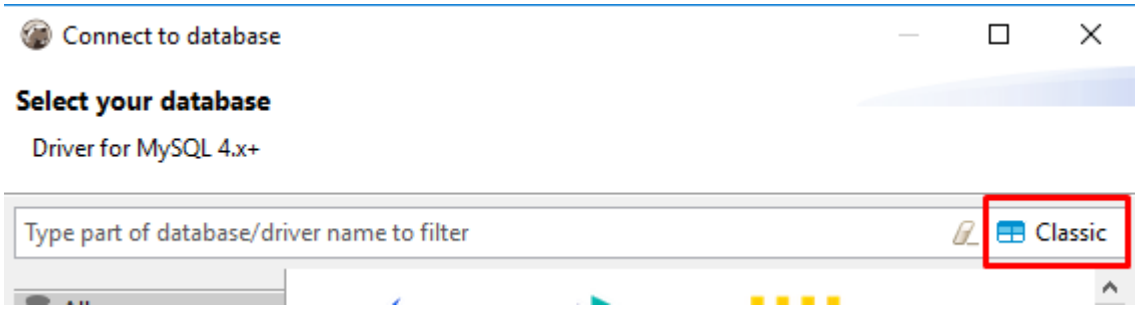

You can choose Simple mode on this step. Simple mode gives simplified access to the database. Basically with the ability to view data only in schemas and tables.

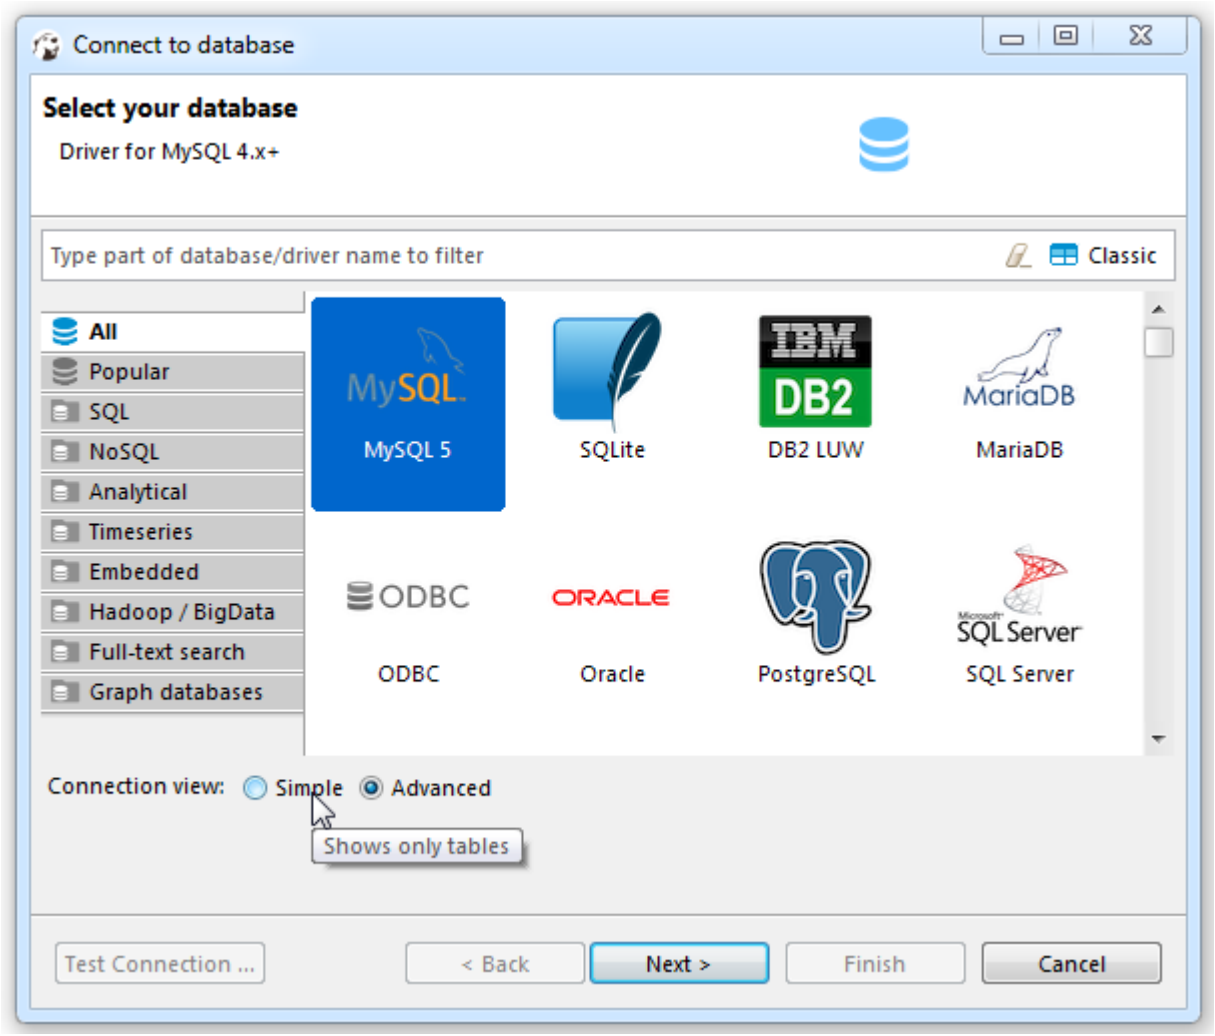

2. In the Connection Settings screen, on the General tab, set all primary connection settings:

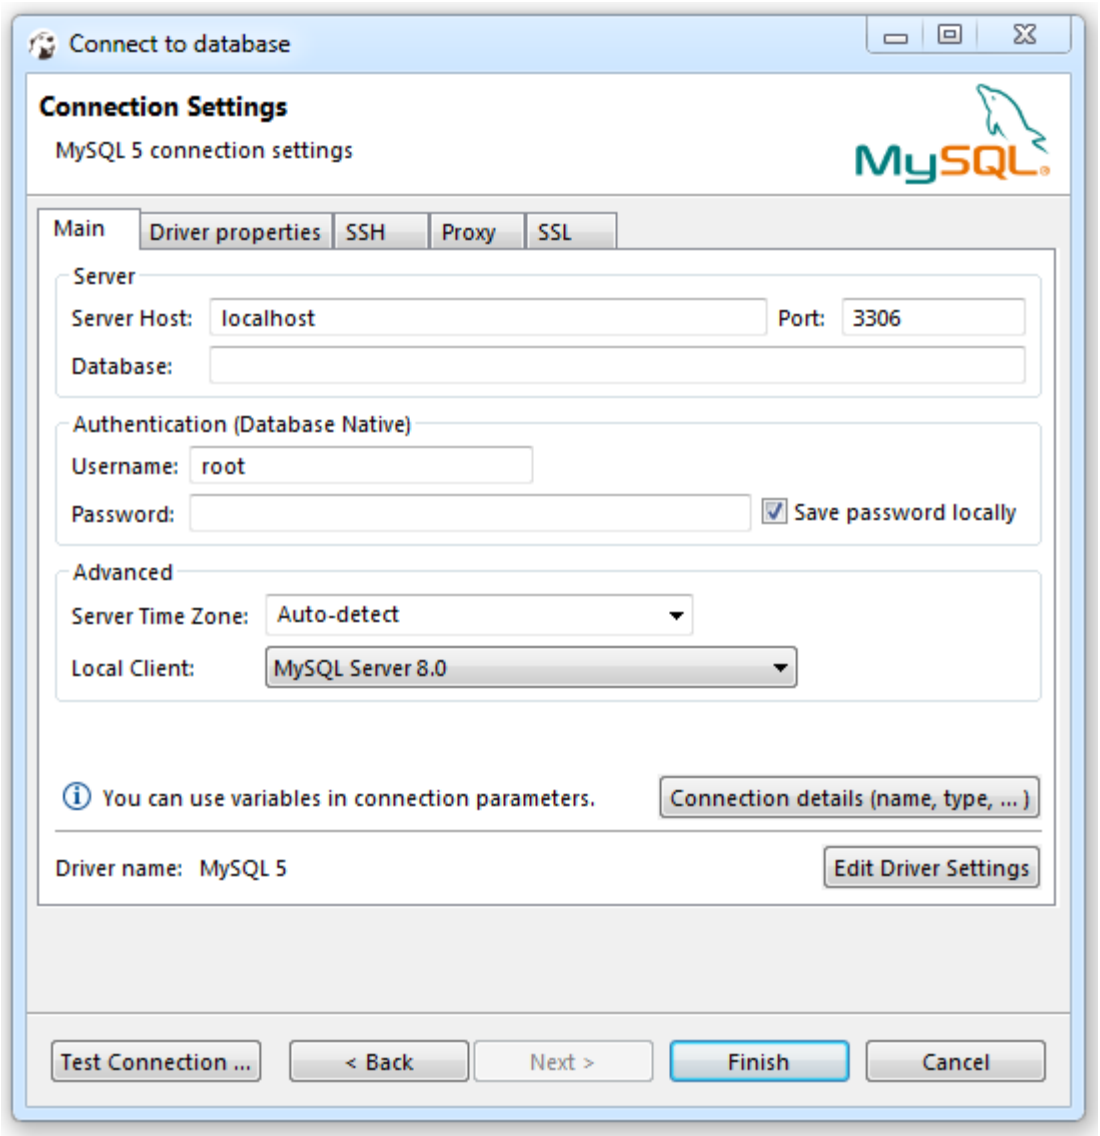

For most drivers required settings include:

- Host
- Port
- Database name
- User name and password

However, the number and type of connection properties are very dependent on the driver.
For example, embedded drivers (such as SQLite, Derby Embedded, HSQLDB, H2 Embedded), unlike remote ones, require only the path to the database.

3. If necessary, specify advanced settings, see **Advanced Settings** section below, and click **Next**.
4. To test if the connection works, click **Test Connection**.
5. Click **Finish**. The connection appears in the tree of connections in the Database Navigator and DBeaver actually connects to the database.

Advanced Settings

Network Settings (SSH, SOCKS, SSL)

If your database cannot be accessed directly, you can use SSH tunnel:

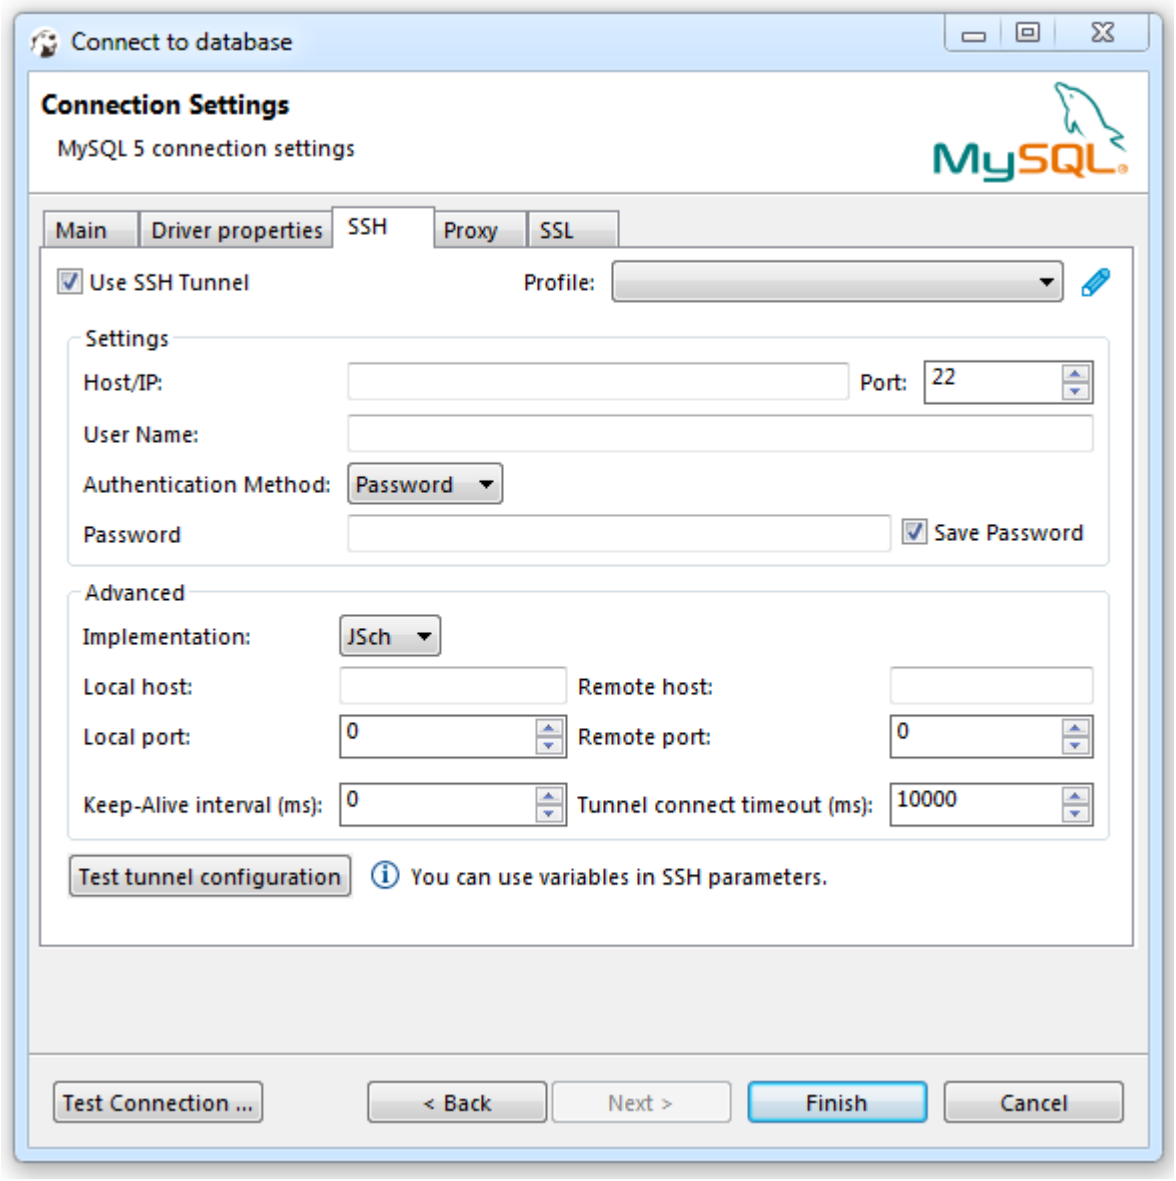

DBeaver supports following SSH authentication methods: user/password, public key authentication and agent authentication. Supported implementations for agent authentications are pageant and ssh-agent.

If a connection has network settings specified, such a connection appears in the application with a special 'arrow' icon such as this:

Connection Details (name, type, etc.)

You can also set the connection name, type and initial settings (such as bootstrap queries, transaction state, global filters, etc.).

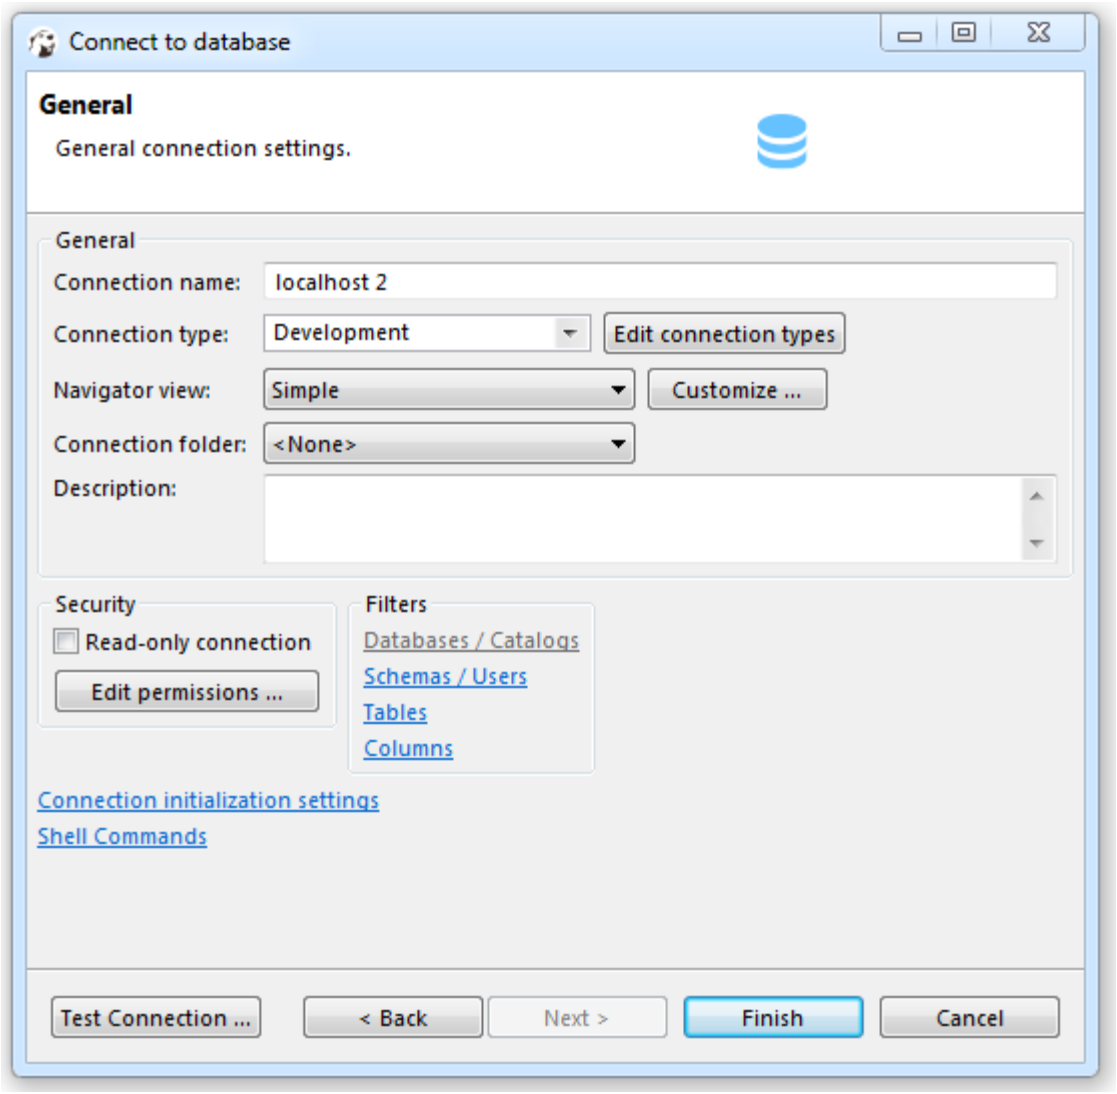

Driver Properties

Each driver has its own set of additional properties. Refer to the driver documentation to get information about available properties and their values.

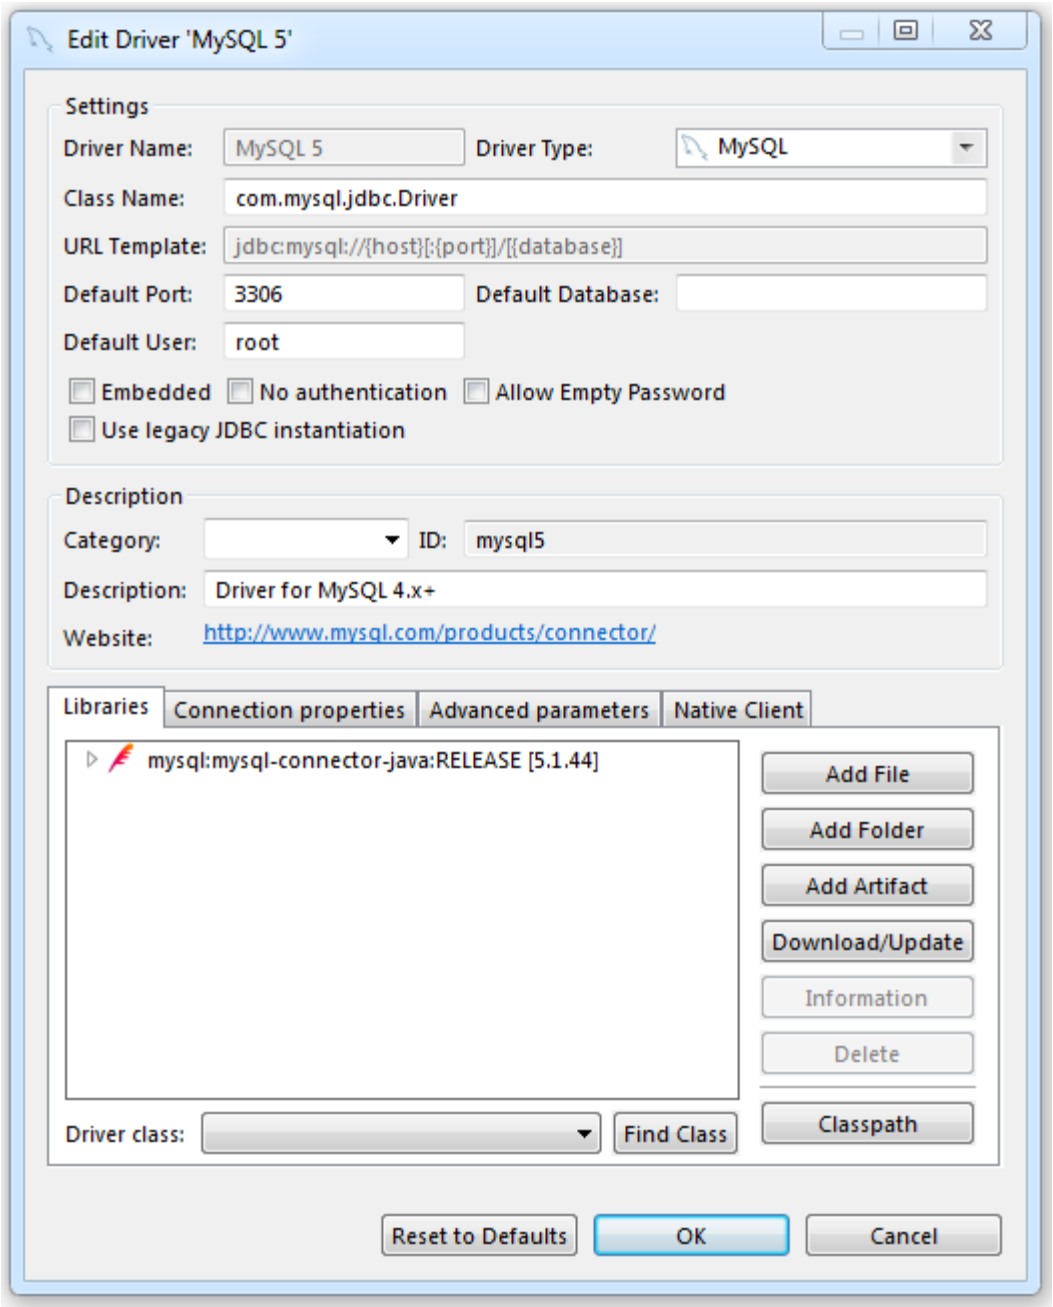

Variables in parameters

You can use variables in all connection parameters and in the driver properties. Variables are system environment variables or one of the following list:

Name	Value
${host}	Host name
${port}	Port number
${database}	Database name
${server}	Server name
${url}	Connection URL
${user}	User name

${password}	User password

Note: option `Use environment variables in connection parameters` must be turned on (see preferences).

Edit Connection

To edit configuration settings of a database connection, in the Database Navigator or in the Projects view, right-click the connection and click **Edit Connection** on the context menu. The Connection configuration window opens:

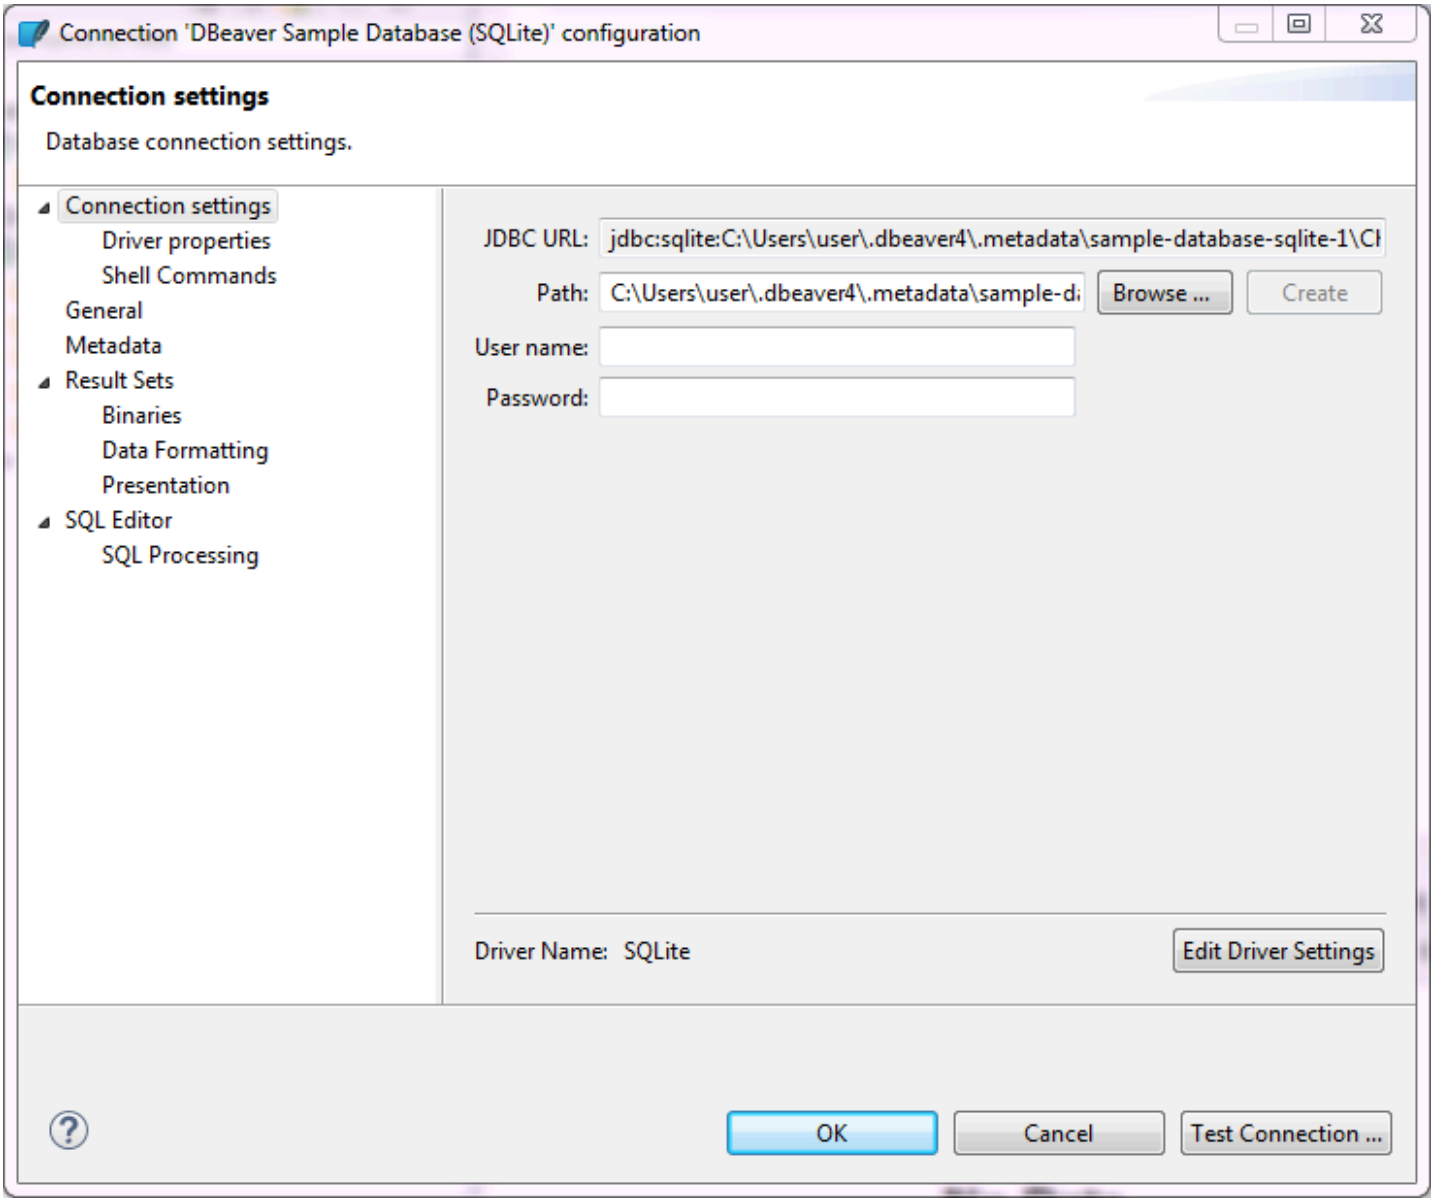

The navigation pane on the left displays configuration sections, most of which are the same as those in the Create new connection wizard, see Connect to Database. There are additional configuration sections as well, such as **Result Sets** and **SQL Editor**. Click the section name to open the configuration settings for editing.

You can test if your connection works with modified settings - click **Test Connection**. When you finish editing your connection, click **OK** to save the changes or **Cancel** to discard them.

Driver settings

- In development

JDBC Time Zones

Connect to Database

To be able to work with the content and structure of a database, you need to connect to it. When you create a new connection to a database, DBeaver automatically connects to the new database, see Create Connection.

To connect to a database using an existing connection, in the Database Navigator or Projects view, click the connection and then click the Connect button in the toolbar or click **Database -> Connect** on the main menu:

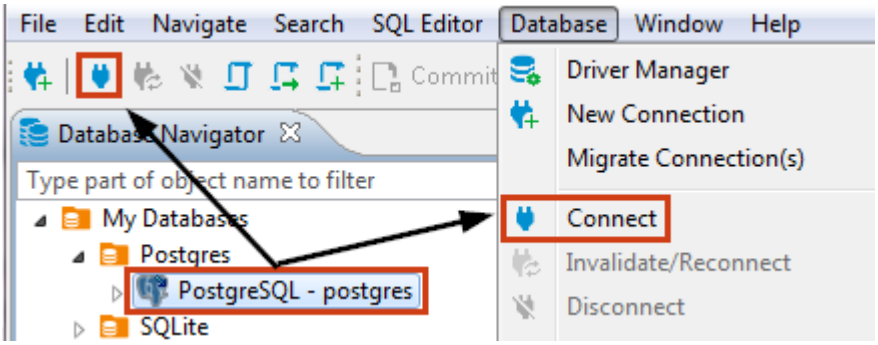

You can also right-click the connection and click **Connect** on the context menu.

If a database connection exists but DBeaver is not connected to the database, the connection appears with its original icon (for example, for PostgreSQL database). When DBeaver connects to the database, the icon changes to signal the connected status: .

If DBeaver cannot connect to a database, the connection appears with an error sign: . If you attempt to connect to such a database, DBeaver displays an error message describing the cause for the error.

Invalidate/Reconnect to Database

You might need to invalidate a database connection and then to reconnect to it again in such cases as connection to the server being lost, etc. To invalidate a database connection and then reconnect to the database, in the Database Navigator or Projects view, click the database connection and then click the **Invalidate/Reconnect** button in the toolbar or **Database -> Invalidate/Reconnect** on the main menu:

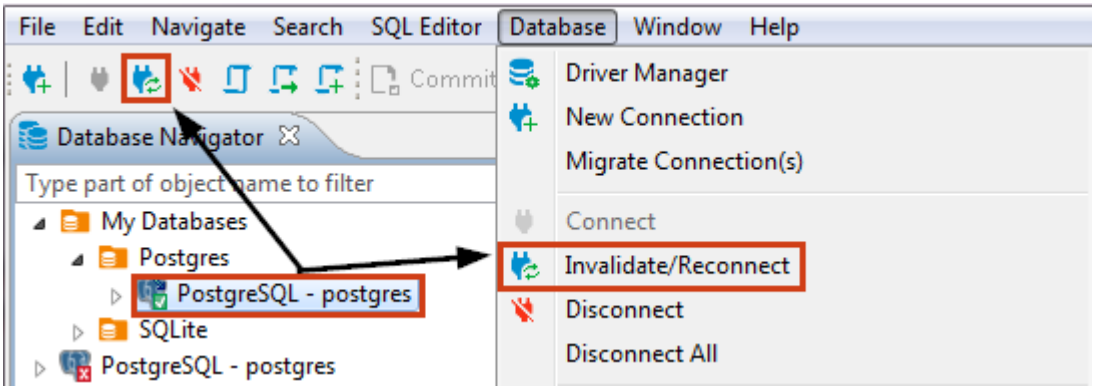

You can also right-click the connection and click **Invalidate/Reconnect** on the context menu.

Disconnect from Database

You might need to disconnect from a database to free up resources or close transactions. To disconnect from a database, in the Database Navigator or Projects view, click the connection and then click the Disconnect button in the toolbar or click **Database -> Disconnect** on the main menu:

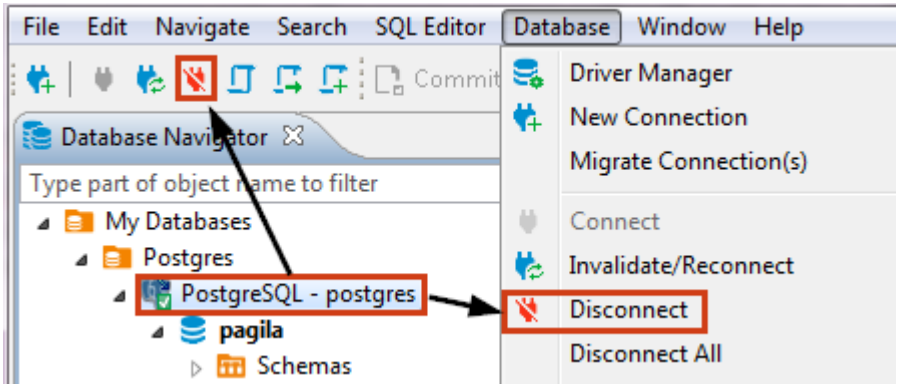

You can also right-click the connection and click **Disconnect** on the context menu.

NOTE: The Disconnect button and menu items are available only for those connections that are activated, that is, marked with the connected sign: .

When DBeaver disconnects from a database, its icon changes to its original state (not connected), for example, for PostgreSQL database.

To disconnect from all active connections, click **Database -> Disconnect All** on the main menu.

Connection Types

Connection types define how DBeaver behaves regarding:

- Default transactions commit mode - with or without automatic commit of changes to the database.

  NOTE: You can override the default commit behavior during your work with connections by changing the commit mode, see Auto and Manual Commit Modes.

- SQL statements execution (with or without user confirmation). If set to require user confirmation for SQL execution, DBeaver shows a confirmation message every time you attempt to execute a 'transaction' type of query (INSERT/DELETE/UPDATE, etc.):

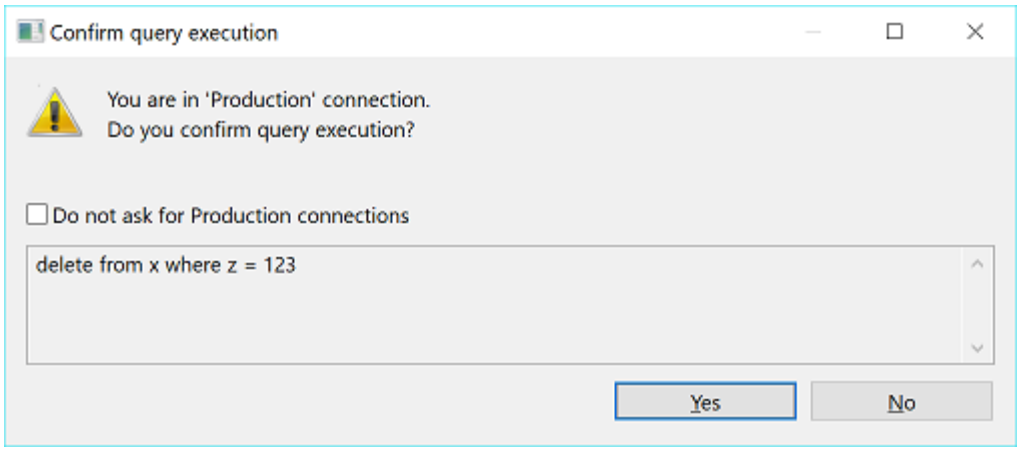

For your convenience, DBeaver supports color-coding of connection types so that you know at once which behavior to expect when you use a certain connection. Colored are database connections that use a certain connection type in the Database Navigator and Projects views as well as editors related to these connections:

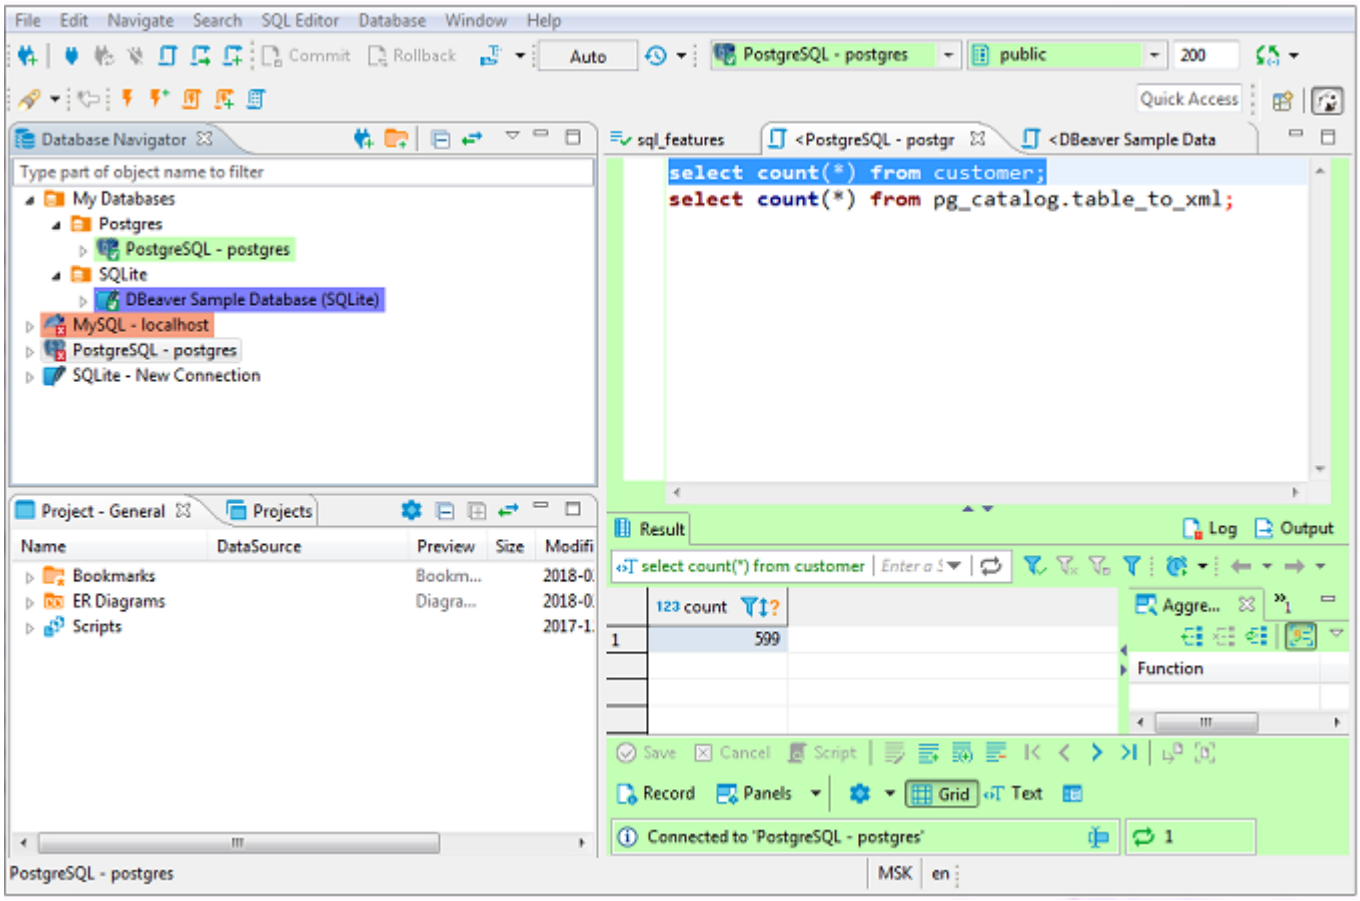

To manage connection types for a database connection, in the Database Navigator or Projects view, click the connection to set focus to it and then press F4 to open the connection properties window. Then, in the properties window, in the navigation pane on the left, click **General** to see the general settings. You can see **Connection Type** field among the settings:

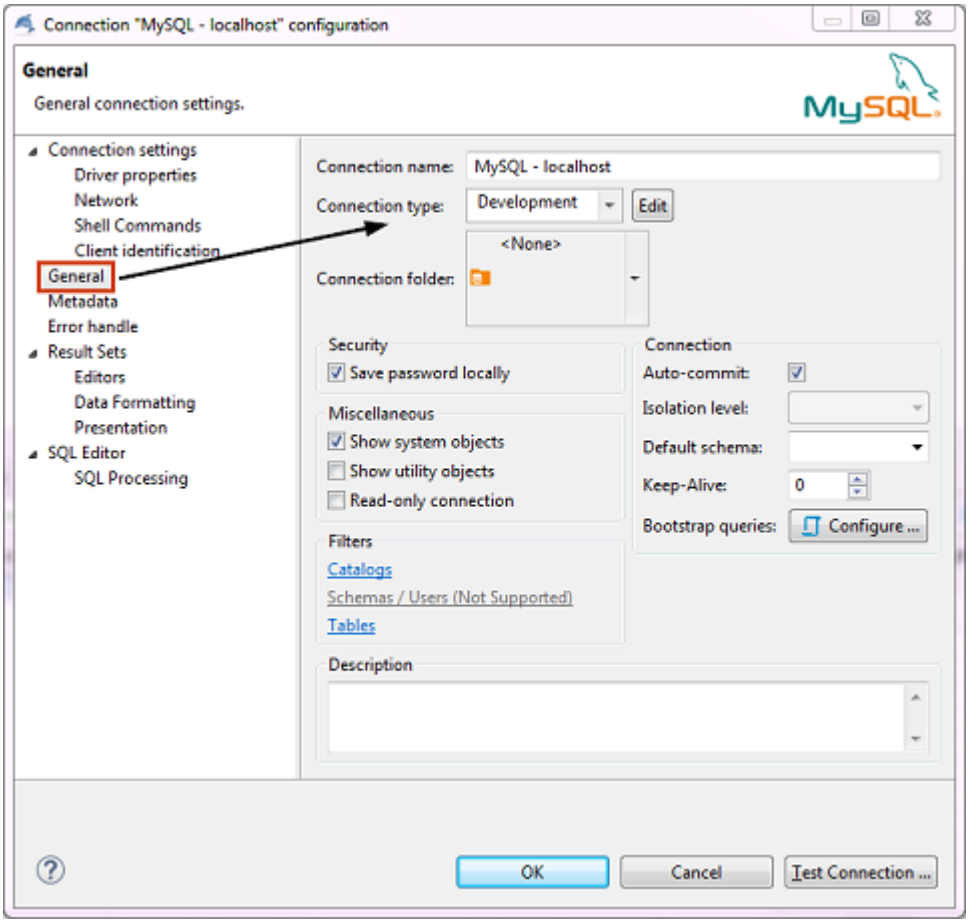

There are three default connection types – **Development**, **Test**, and **Production**. You can change the connection type for your database connection as well as you can create a new connection type, edit or delete an existing one.

Change Connection Type

By default, the **Development** connection is preset for all database connections. You can change the connection type to one of the default connection types or to a custom type, if there are any.

To change the connection type:

1. In the connection properties window, on the **General** page, click the **Connection type** field and then click the connection type in the dropdown list:

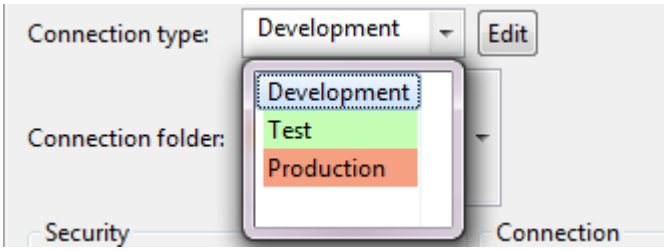

2. To test the connection, click **Test Connection**. To confirm the change, click **OK**.

Create Connection Type

To create a connection type:

1. In the connection properties window (F4 on a connection), on the **General** page, click **Edit** next to the **Connection type** field. The Properties for connection types window opens:

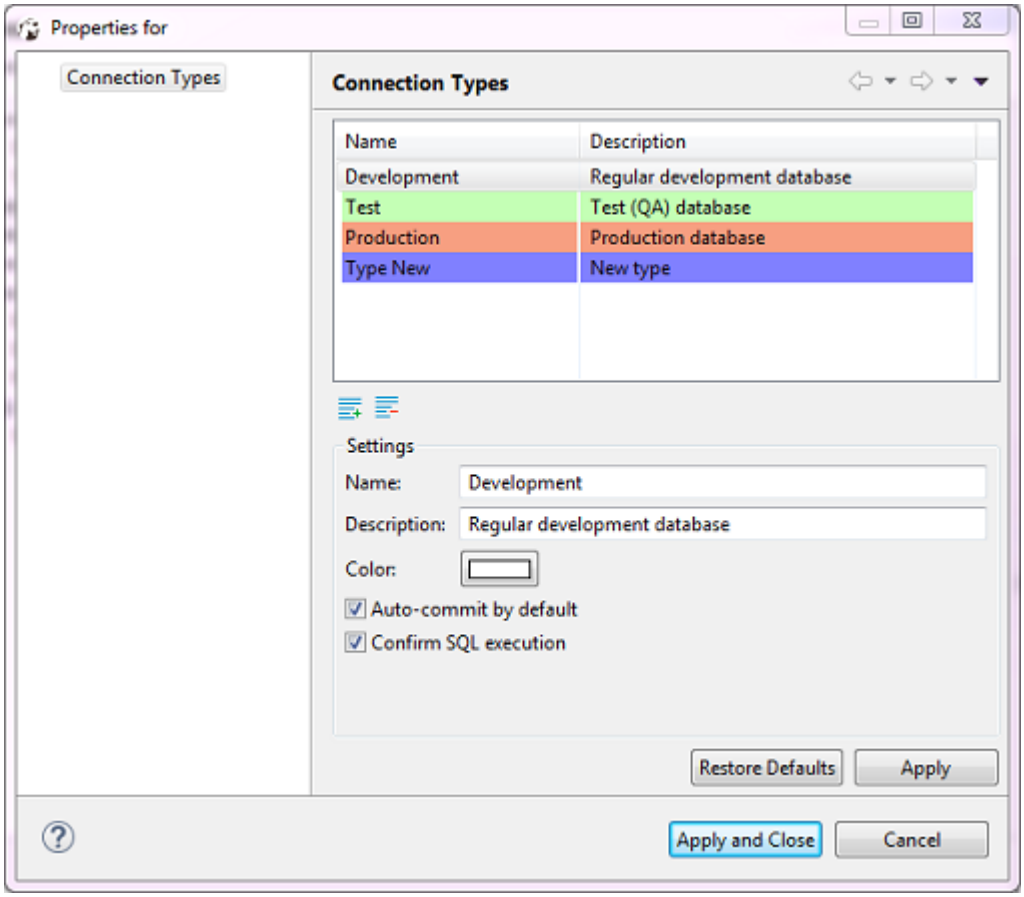

   The window displays existing connection types and their settings.

2. Click the new connection type button. A new connection type appears in the list:

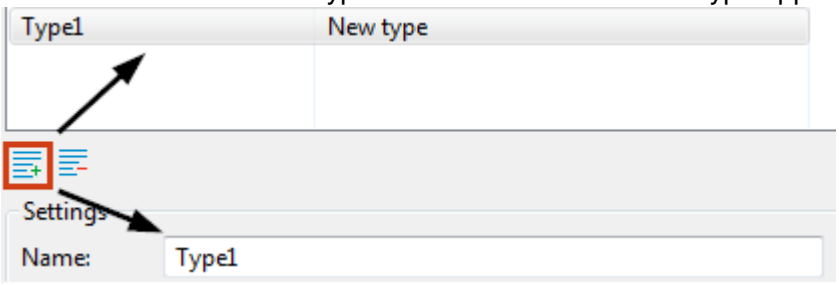

3. Now you can specify the settings for the new connection type:
   - Enter the connection type`s name into the **Name** field.
   - Enter a description into the **Description** field, if needed.
   - Click the **Color** box and select the color for the new connection type.
   - To set DBeaver to automatically commit changes to the database when connections use this connection type, select the **Auto-commit by default** checkbox. Otherwise, leave it empty.
   - To set DBeaver to ask for your confirmation at each execution of SQL statement of 'transaction' type, select the **Confirm SQL execution** checkbox. Otherwise, leave it empty.

4. Click **Apply** to apply the changes and keep the window open or click **Apply and Close** to apply the changes and close the window. To discard all changes and return to the previous state, click **Restore Defaults**.

Edit Connection Type

To edit a connection type:

1. In the connection properties window (`F4` on a connection), on the **General** page, click **Edit** next to the **Connection type** field.
2. Specify the settings for the new connection type the same way as when you create a connection type, see 'Create Connection Types' section above.
3. When you finish editing the connection types, click **Apply** to apply the changes and keep the window open or click **Apply and Close** to apply the changes and close the window. To discard all changes and return to the previous state, click **Restore Defaults**.

Delete Connection Type

To delete a connection type:

1. In the connection properties window (`F4` on a connection), on the **General** page, click **Edit** next to the **Connection type** field. The Properties for connection types window opens.
2. In the Properties window, in the list of connection types, click the connection type to set focus to it and then click the delete button under the list:
3. Click **Yes** in the confirmation dialog box to confirm the deletion. Otherwise, click **No**.
4. Click **Apply** to apply the changes and keep the window open or click **Apply and Close** to apply the changes and close the window.

Auto and Manual Commit Modes

DBeaver supports two modes for committing changes to the database:

- **Auto-commit** transfers all changes that you make immediately to the database.
- **Manual commit** requires your confirmation before committing a change to the database or rolling it back.

Though available in many cases, the two modes are actionable only in SQL Editor. See the next sections for details of using the modes.

To switch between the modes, use the mode selection button that appears in one of the two views: or .

Auto-Commit Mode

Auto-commit mode is the default one for the Development and Test connection types, see Connection Types. Auto-commit mode is on if you can see the auto-commit view of the mode selection button ( ) in the application toolbar. If you see the manual commit view ( ), then in order to switch to auto-commit mode, click the mode selection button – it changes to auto-commit. At the same time, this disables the two manual commit buttons in the toolbar: **Commit** and **Rollback** – these are available only in manual commit mode.

The statistics field next to the mode selection button always shows **Auto** in auto-commit mode:

.

Clicking the statistics field opens the Transaction Log.

Manual Commit Mode

Manual commit is intended to protect your database from inadvertent changes and that is why it is the default mode for Production connection type, see Connection Types.

Manual commit mode is on if you can see the Manual commit view of the mode selection button ( ) in the application toolbar. If you see the auto-commit view ( ), then in order to switch to manual commit mode, click the auto-commit button – it changes to manual commit. At the same time, this enables the two manual commit buttons in the toolbar: **Commit** (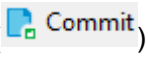

) and **Rollback** (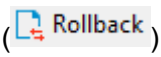

).

In manual commit mode, when you execute SQL statements (`Ctrl+Enter`), the number of database modifying statements that pend commitment to the database appears in the statistics field next to the mode selection button:

.

If you hover you mouse over the field, you can see statistics of your SQL statements:

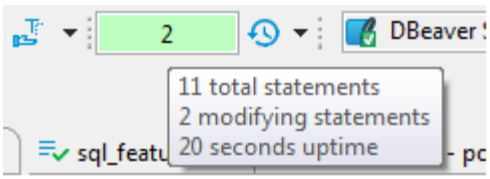

To commit statements to the database, click the **Commit** button in the toolbar. To discard them, click **Rollback**.

If no modifying statements have been made, the statistics field shows **None**: 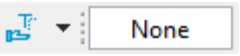

Clicking the statistics field opens the Transaction Log.

Smart Commit Mode

When smart commit is enabled and you are in auto-commit mode then DBeaver will monitor your activity. Once you will try to execute any data modifying query (UPDATE, INSERT, DELETE, UPSERT, MERGE, etc) DBeaver will switch to manual commit mode before executing your query. Also if you will edit table data and save your changes - DBeaver will also switch to manual mode before actual data modification.

If option "Return to auto-commit on transaction end" is on then DBeaver will switch back to auto-commit mode once you execute Commit or Rollback command (using main toolbar or main menu actions).

Smart commit mode is very useful if most of the time you are working in read-only mode. It doesn't lock tables when you perform SELECT queries. Transaction will be started only when you start to modify your data.

Transaction Isolation Level

For both, Auto and Manual commit modes, you can select the transaction isolation level. To do so, click the arrow next to the mode icon and then click the required option in the dropdown list:

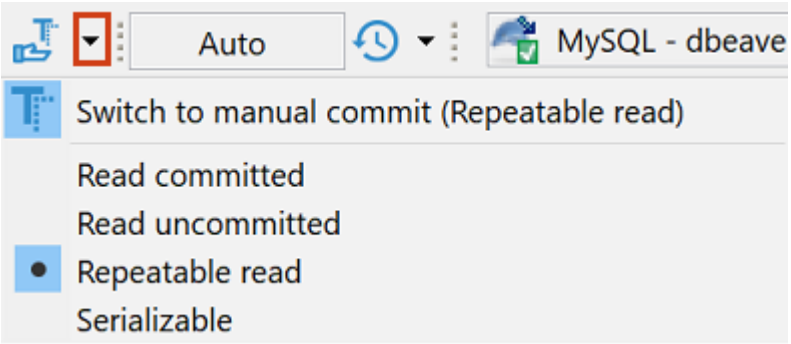

Transaction Log

Transaction Log shows all transactions (queries of 'transaction' type such as INSERT/DELETE/UPDATE and others) made during the current DBeaver session. To open the Transaction Log, click the **Transaction log** button ( ) in the toolbar or the statistics field to the left of it.

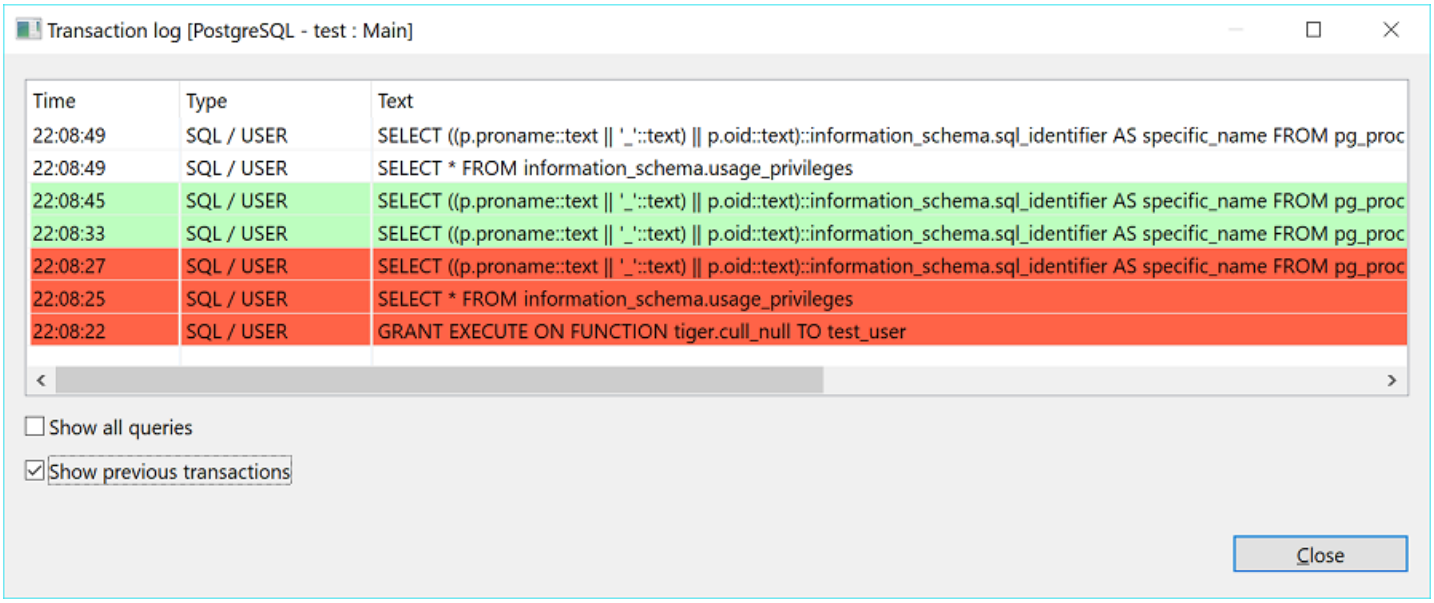

The Transaction Log window shows transactions that are:

- In progress or pending - shown without any special color
- Successfully committed – in green:
- Rolled back – in orange or red:

To see all previous transactions during the current session, select the **Show previous transactions** checkbox. To see all queries including non-transactional ones, select the **Show all queries** checkbox.

Pending transactions

It might be useful to check pending transactions because they might lock your database. To see pending transactions, click the arrow next to the **Transaction Log** button in the toolbar and then click **Pending Transactions** on the dropdown menu:

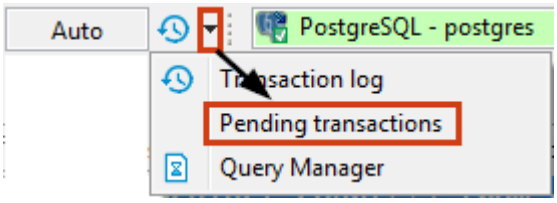

The upper table of the Pending transactions window shows currently active connections and the number of their transactions. The bottom table shows query details of the connection that is currently in focus in the upper table:

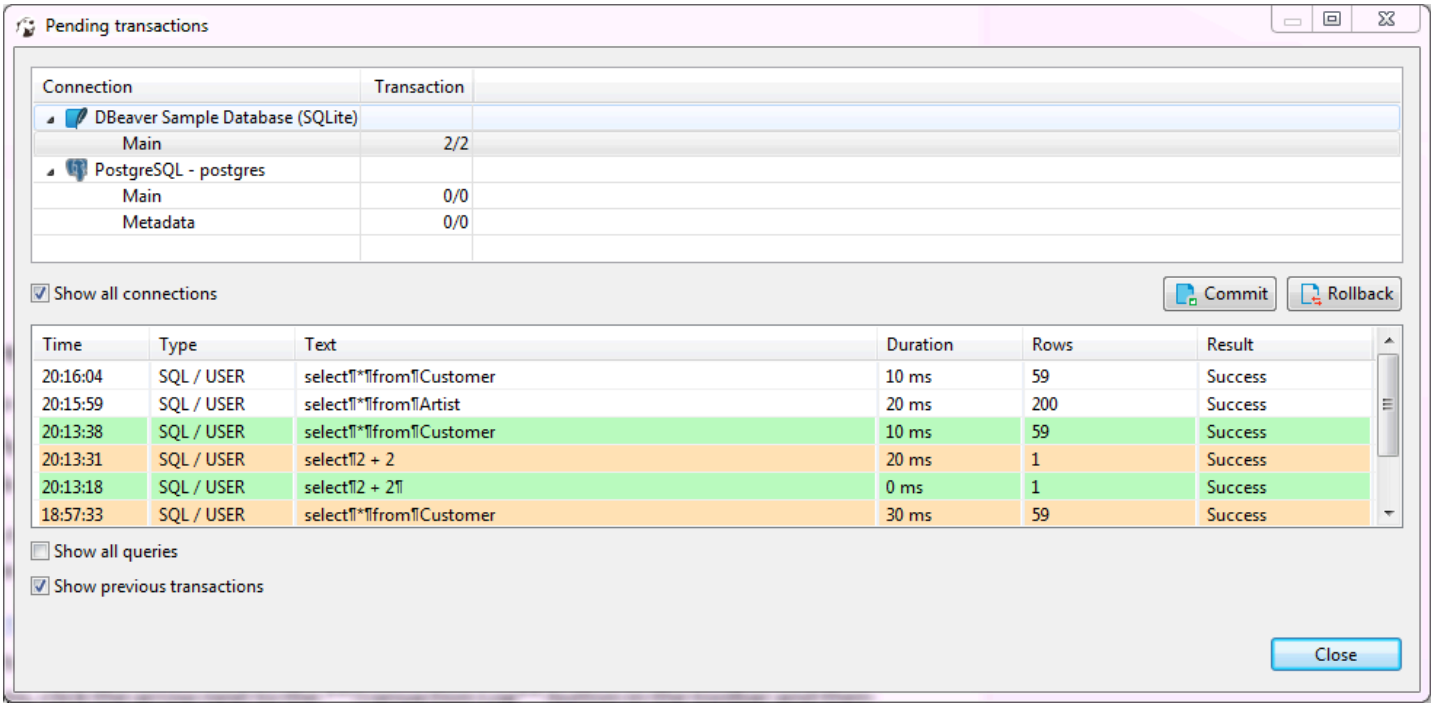

When the Pending transactions window opens, the upper table shows only those connections that have pending transactions. If no connections have pending transactions, the table is empty. To see all connections that are currently active (connected), select the **Show all connections** checkbox.

You can commit or roll back transactions right from the Pending transactions window: in the upper table, click the row with required uncommitted transactions and then click the **Commit** or **Rollback** button, depending on your purpose. If a transaction is committed/rolled back successfully, both buttons are disabled. If the operation is unsuccessful, the system displays an error message.

To see all previous transactions made during the current session, select the **Show previous transactions** checkbox. To see all queries including non-transactional ones, select the **Show all queries** checkbox. The green rows are committed transactions, orange (or red) ones are rolled back, rows without a special color are non-transactional or pending transactions.

Database drivers

You can use pre-configured database driver or create new driver.

DBeaver has a lot of pre-configured driver including SQL, NoSQL, key-value databases, graph databases, search engines, etc. But sometimes you need to connect to a database which was not configured in DBeaver yet.

All you need is JDBC driver of your database. The rest is easy.

Obtaining JDBC driver

JDBC driver is a program (in Java) which can connect and operate with some local or remote database server. It usually provides all needed functionality to cover 100% of database functionality. Usually, JDBC driver are provided by database vendors to let customers ability to work with their databases.

JDBC driver consists of one or multiple `jar` files. Jar file is a library which contains program code and some other files. You need to download driver's jar files before adding them in DBeaver. Sometimes jar files are included in database server distribution - in that case you need to refer your database documentation or ask your DBA.

Adding driver configuration in DBeaver

Open driver manager dialog

You can open driver manager from main menu:

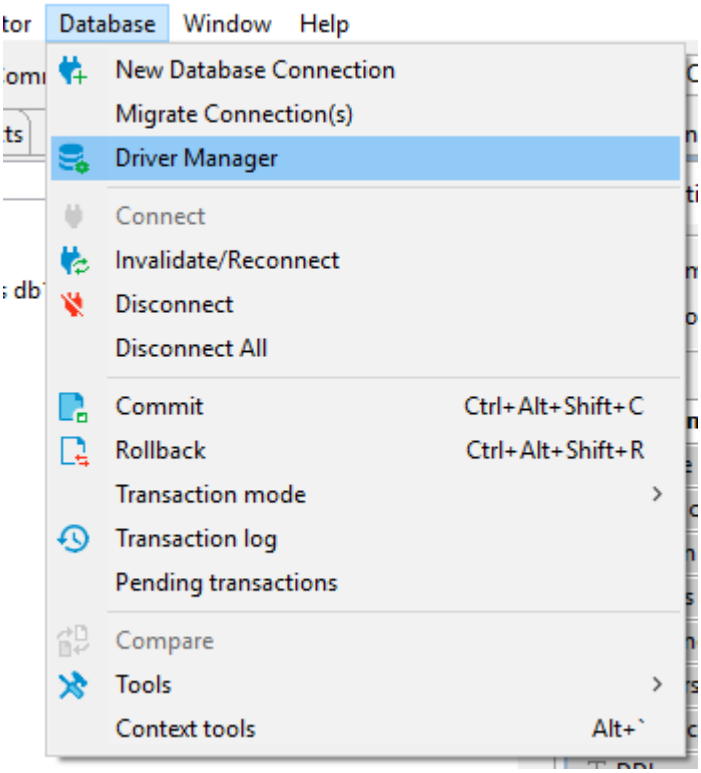

or from Database Navigator drop-down menu.

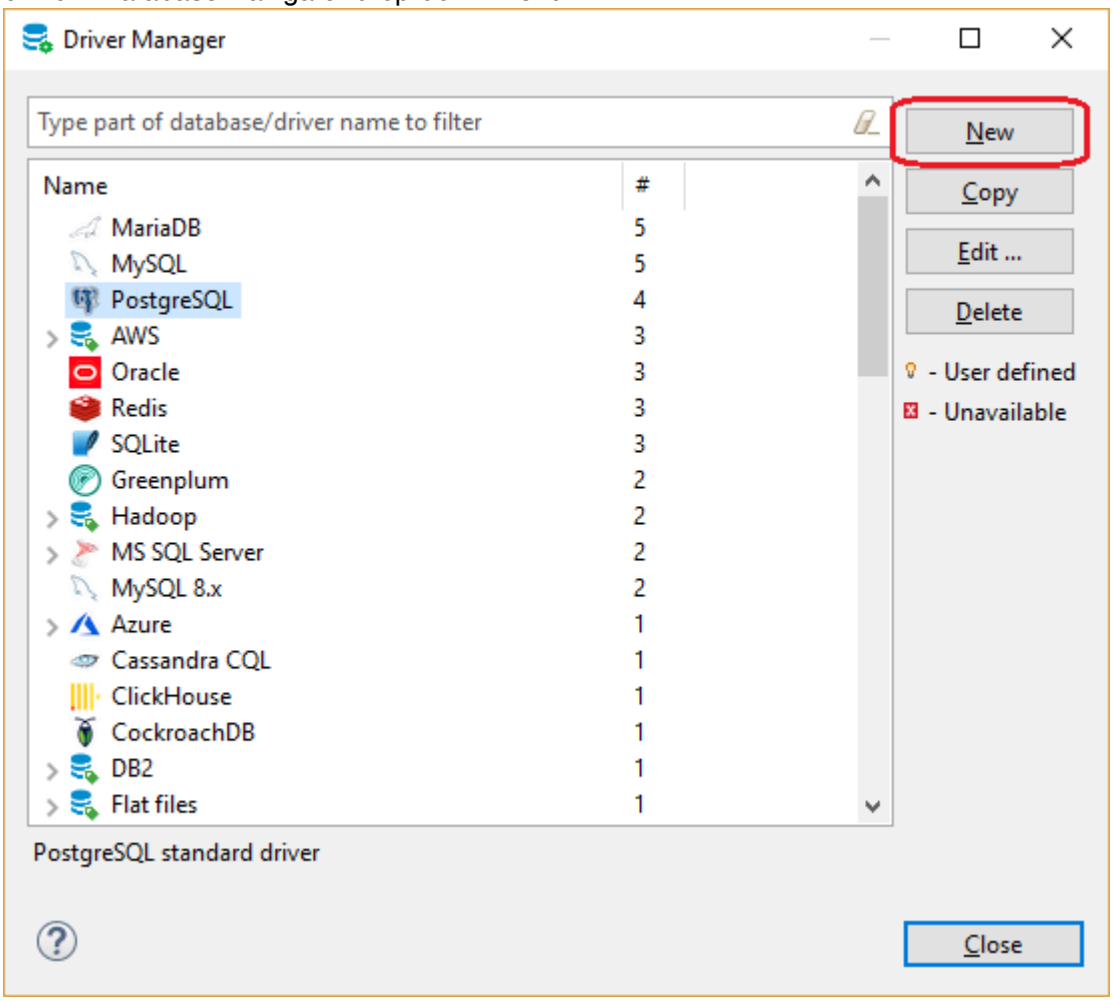

Add new driver

Just click the button New and create a new driver. On the driver edit dialog you need to enter all required information:

Main parameters

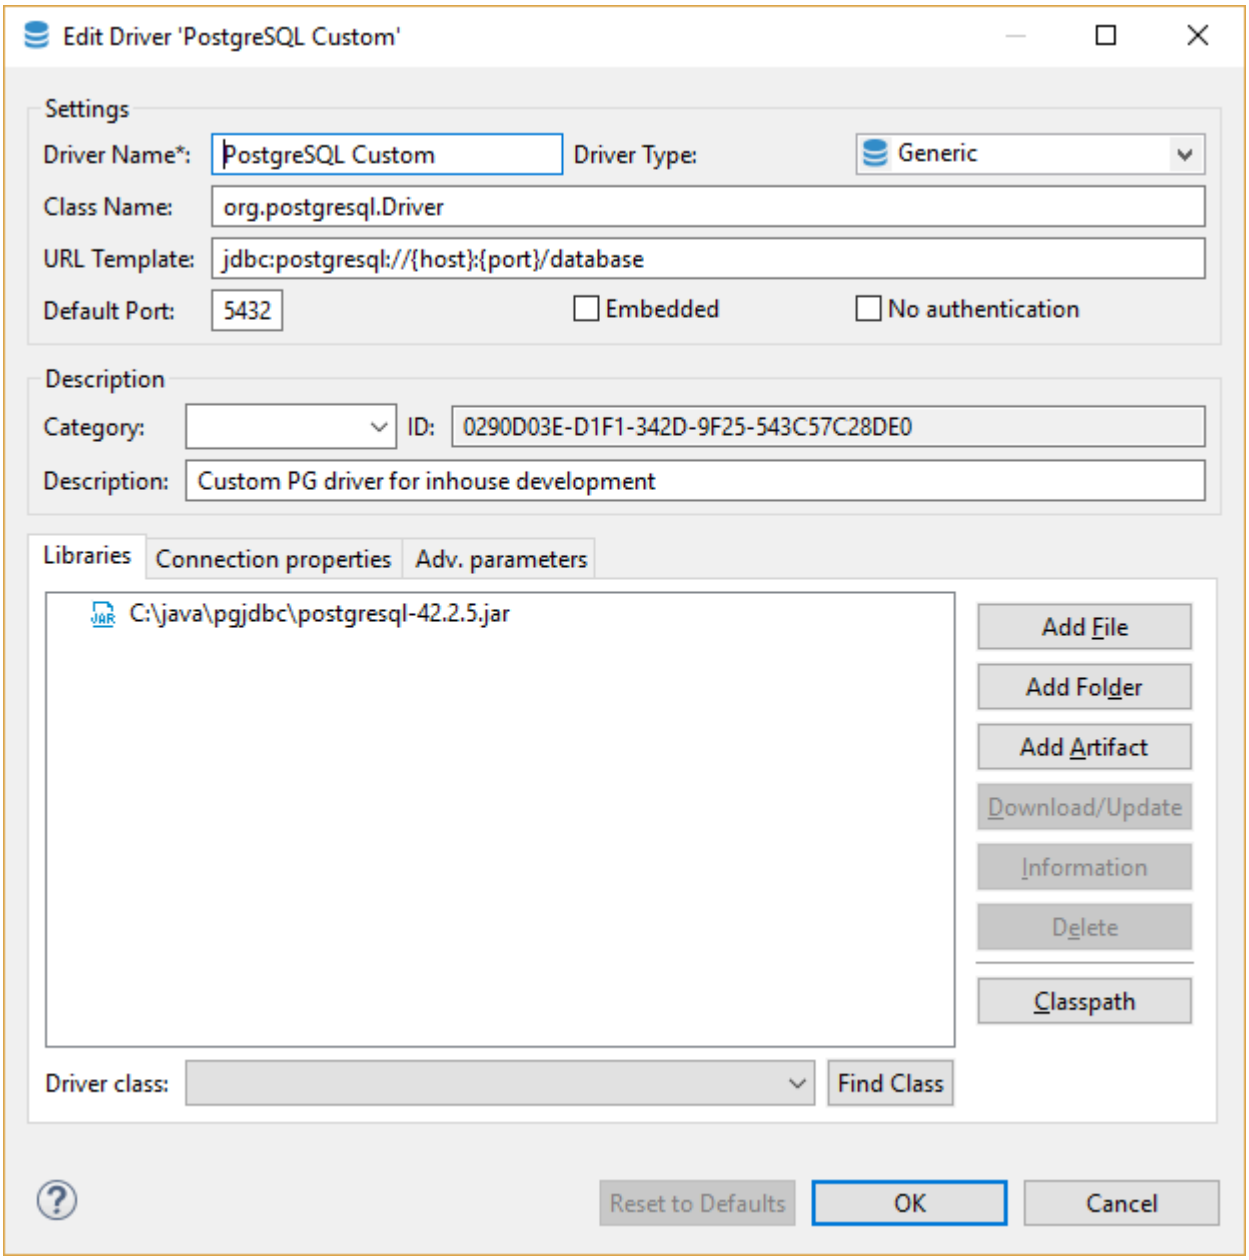

Parameter	Description
Driver Name	Name of your driver. It can be any name you like
Driver Type	Driver provider. In 99% cases you will need Generic driver (JDBC provider)
Class Name	JDBC driver class name. You can get it from the documentation or find it in jar files (see "Find Class" button description)
URL Template	Template of driver URL. You can leave it empty. But in this case you will be ready to set JDBC URL for each your connection. It is better to define a valid template, this will greatly simplify connections creation. See "URL Templates" for the detailed description
Default Port	Default database port. You can get it from documentation or leave it empty
Embedded	Enable it for server-less databases. This flag affects a few config options related to network/connections management

No Authentication	This means that driver doesn't require authentication (no user/password fields will be shown)
Category	Driver category, deprecated
ID	Driver unique ID, ignore it
Description	Driver description, it is shown on some dialogs/wizards as a hint

Libraries

This is the list of jar files, binary libraries (dll or so) and any other files required by driver. In most cases you need only jar files.
Click "Add File" to add single jar file, "Add Folder" to add folder with Java classes/resources and "Add Artifact" to add Maven artifact (see below).

After you add jar files you will be able to find all JDBC driver classes which present in these jars. Just click on the "Find Class" button and DBeaver will show all of them. In most cases there is just one driver class in the driver. If there are many of them then you need to refer to the driver's documentation.

Maven artifacts

DBeaver can download driver jars directly from Maven repository (it is a global public repository of Java libraries, usually open-source). If your database driver is published on some public repository you can use this feature. Maven artifacts are better than plain jar files because you can see all existing driver versions and can change driver version in runtime without any driver properties reconfiguration.

Saving driver, adding connection

After you finished configuring your driver just press Ok button.
Now you can create connection.

If you need to change some driver properties later you can access them directly from connection properties dialog:

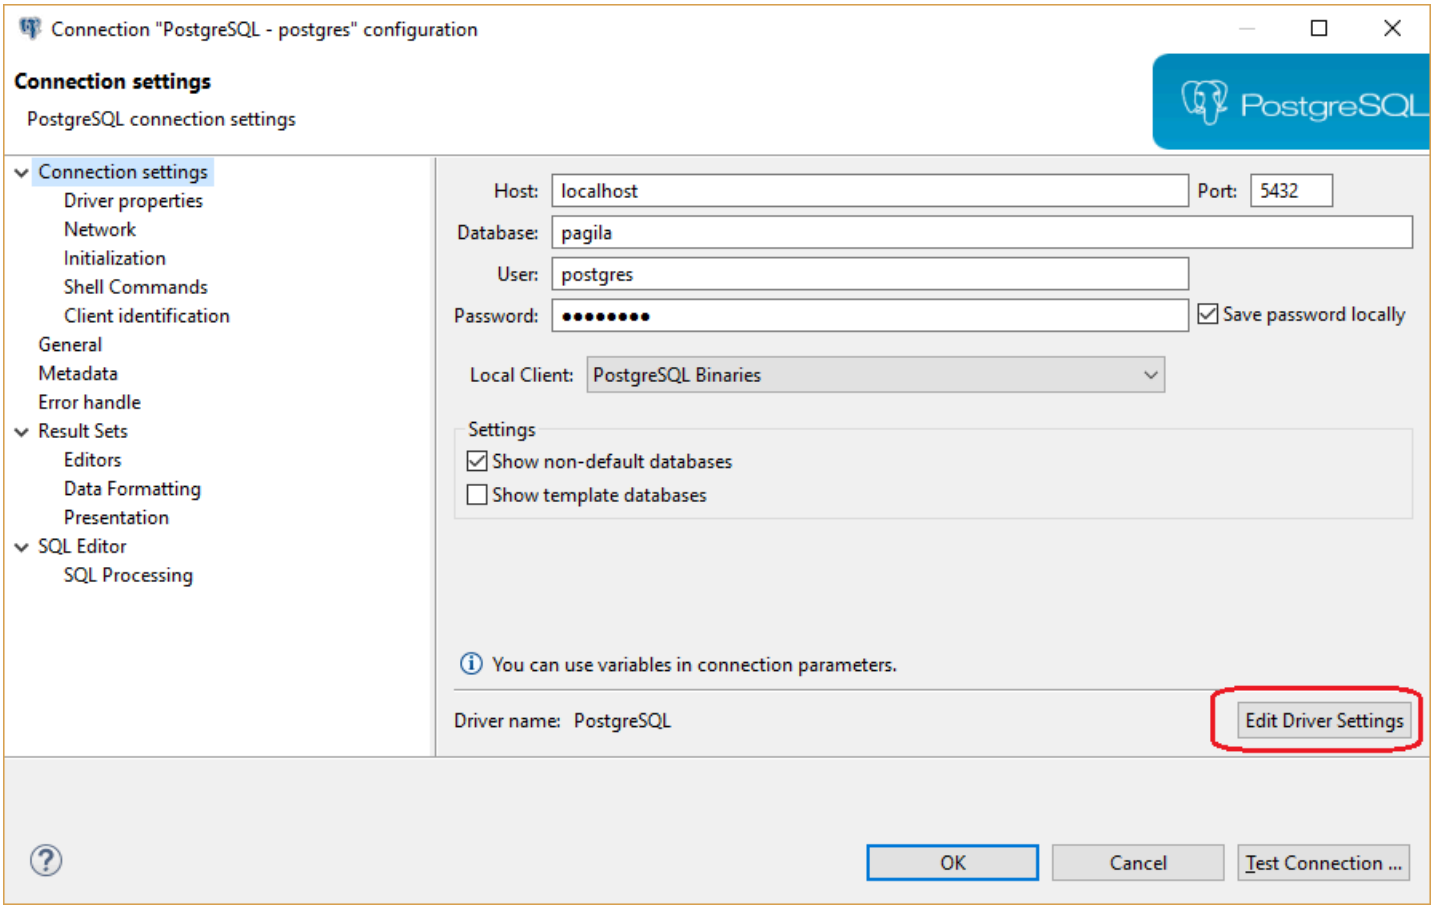

URL Templates

JDBC drivers use URLs to identify remote servers - strings similar to classic web URLs. Usually, URL has form `jdbc:vendor:host:port/database`, for example `jdbc:postgresql:localhost:5432/postgres'. It is not very convenient to edit such long and unobvious string. DBeaver can construct this URL from connection parameters (like host, port, etc).

For example above the URL template is: jdbc:postgresql://{host}:{port}/{database}
Host, port and database are parameters which you will need to enter on the connection configuration page.

Supported URL variables:

Parameter	Description
{host}	Database server host name
{port}	Database server port number
{database}	Target database name
{server}	Target server name (rarely used)
{folder}	Folder path (on the local file system). Used for embedded drivers
{file}	File path (on the local file system). Used for embedded drivers

Advanced settings

For most driver you don't need to change any advanced properties. But in some cases you can use this as a driver tuning, e.g. for better performance or for structure fixing.

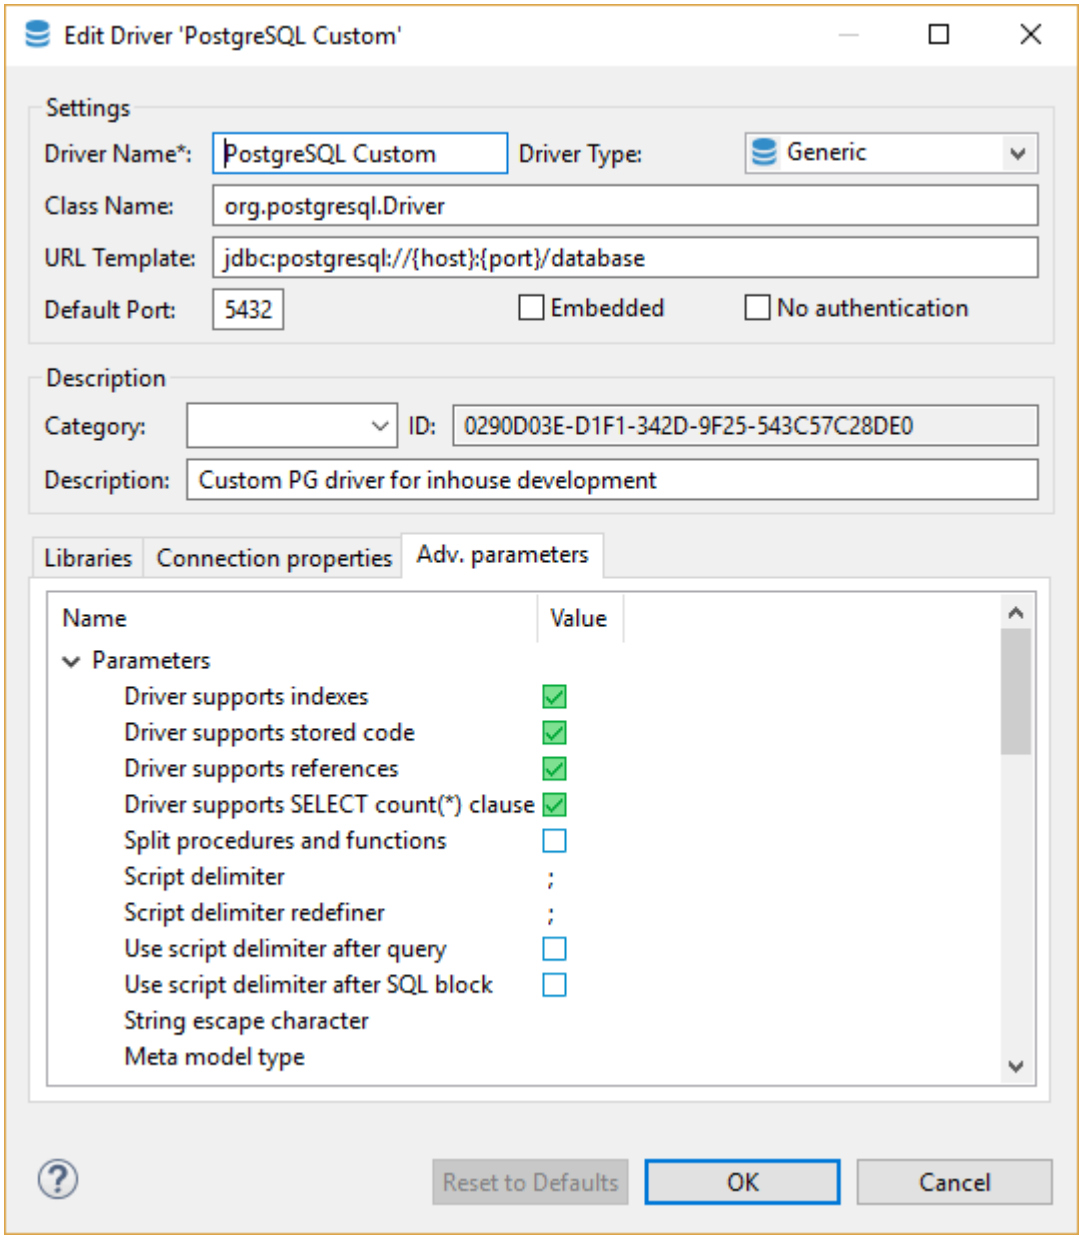

Main parameters

Parameter	Description
Driver supports indexes	Driver supports table indexes
Driver supports stored code	Whether this driver supports stored code (procedures, functions, packages, etc)
Driver supports references	Driver supports table references (foreign keys)
Driver supports SELECT count*() clause* / *Driver supports SELECT count*() clause	
Driver supports views	Driver supports table views
Split procedures and functions	Show procedures and functions in different folders
Script delimiter	Literal for SQL queries separation in scripts

Script delimiter redefiner	SQL clause which redefines script delimiter value
Use script delimiter after query	Keep SQL script delimiter after each SQL query
Use script delimiter after SQL block	Keep SQL script delimiter after SQL script blocks (BEGIN/END)
String escape character	Character used to escape special symbols in strings
Meta model type	Type of metadata reading model - standard or indexed
All Objects Pattern	SQL pattern for all metadata objects
Omit catalog(s)	Do not read and use catalog (aka database) information
Omit single catalog	Hide catalog (database) if there is only one catalog on server
Omit schema(s)	Do not read and use schemas information
Omit single schema	Hide schema if there is only one schema on server
Use schema filters	Use JDBC schema filters when database doesn't support catalogs. Otherwise just read all database schemas and filter on client-side
Omit type cache	Do not use data types provided by driver
Shutdown parameter	Database shutdown URL parameter
Create database parameter	Database create URL parameter
Driver supports multiple results	Driver supports multiple results for a single query
Driver supports result set limit	Driver supports multiple result set limit (max rows)
Driver supports structure cache	Driver supports structure cache reading. Enables schema columns, keys, etc
Driver supports TRUNCATE operation	Driver supports TRUNCATE command. It is much faster than DELETE without criteria

Queries (Custom driver queries)

Parameter	Description
Get active database	Query to obtain active database name
Set active database	Query to change active database
Shutdown database	Query to shutdown active database connection. Used for some embedded databases
PING query	Query to check connection state
Dual table name	Name of dummy 'DUAL' table which is used for expressions evaluation
Active object type	Type of selectable object (schema, catalog)
Driver supports results scrolling	Driver supports resultset scrolling
Quote reserved words	Quote columns/table names if they conflicts with reserved SQL keywords
Escape LIKE masks in search queries	Use to access JDBC metadata API. Enabled by default but should be disabled for some (broken) driverss

DDL (DDL generation options)

Parameter	Description
Drop column short syntax	Use 'ALTER TABLE DROP column-name' instead of standard syntax
Drop column - use brackets	Use 'ALTER TABLE DROP (column-name)' instead of standard syntax
Use legacy SQL dialect for DDL	Use legacy SQL dialect for DDL
Add COLUMN keyword in alter table query	Add COLUMN keyword after keyword ADD and before column name in alter table query

Formatting (SQL values formats)

Parameter	Description
Timestamp format	Format pattern for timestamp columns
Date format	Format pattern for date columns
Time format	Format pattern for time columns

Summary

If you have configured some driver, it works good and you think that it makes sense to have this driver configuration in standard DBeaver, please send your configuration to us. Just create a feature request issue on GitHub and copy/paste driver description to the ticket (in any suitable form).

Thank you :)

Data export/import

You can perform data export/import or migration for database table(s). We'll describe most typical use cases.

Exporting table data to CSV format

1. Select a table(s) you want to export. In the context menu choose "Export Data".
   (Note: you also can export data from custom SQL queries results. For that in results context menu choose "Export results").

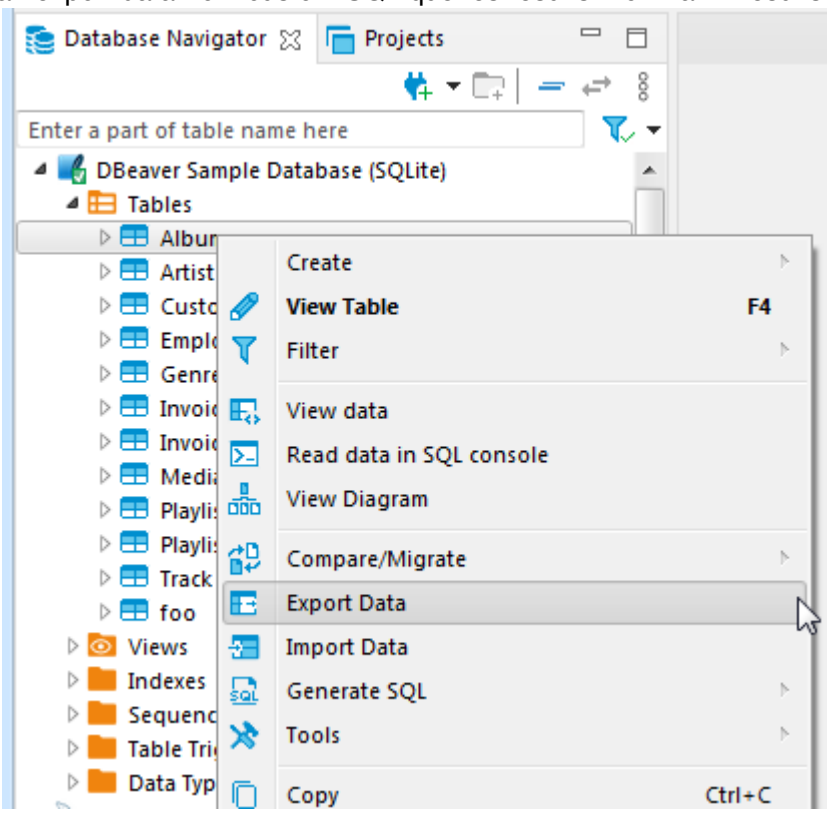

2. Choose export format. DBeaver support many different output formats including CSV, HTML, XLSX, etc:

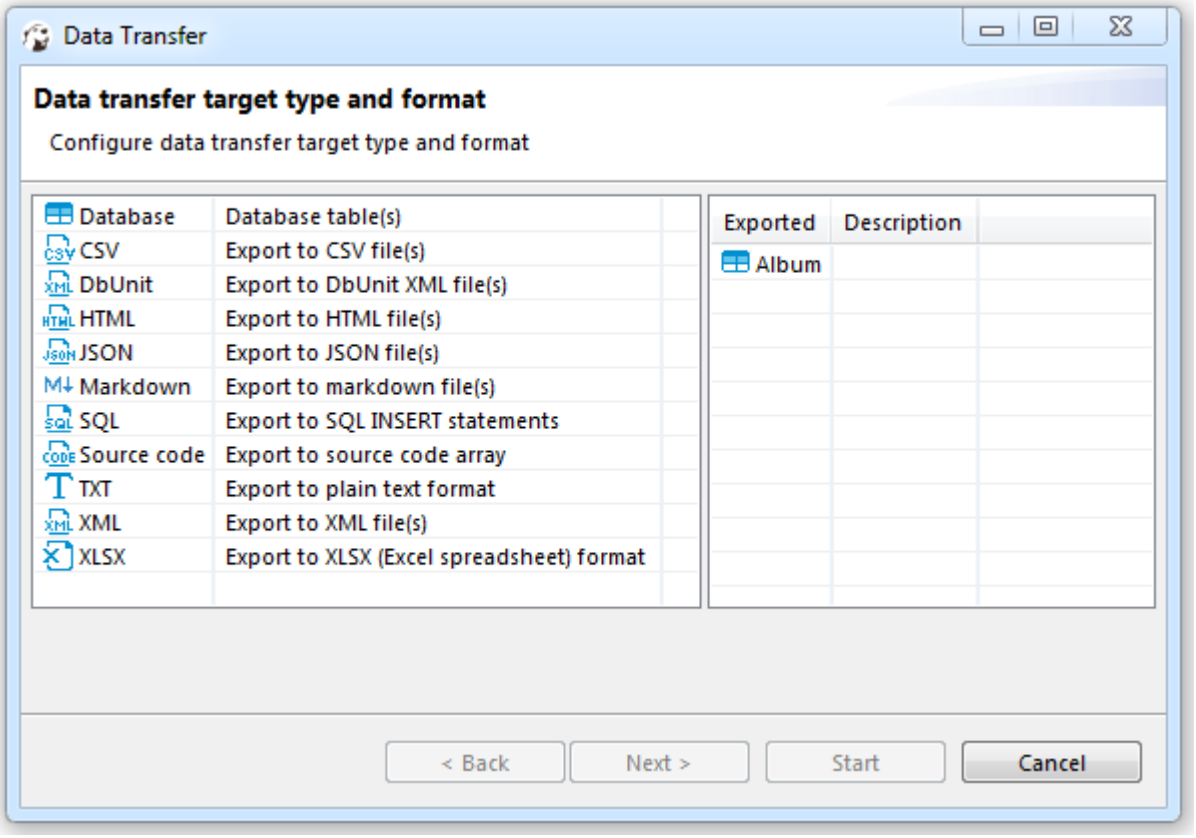

3. Set data extraction options (how data will be read from tables). This may affect extraction performance. And set export format option. They are specific to the data format you chose on step 2:

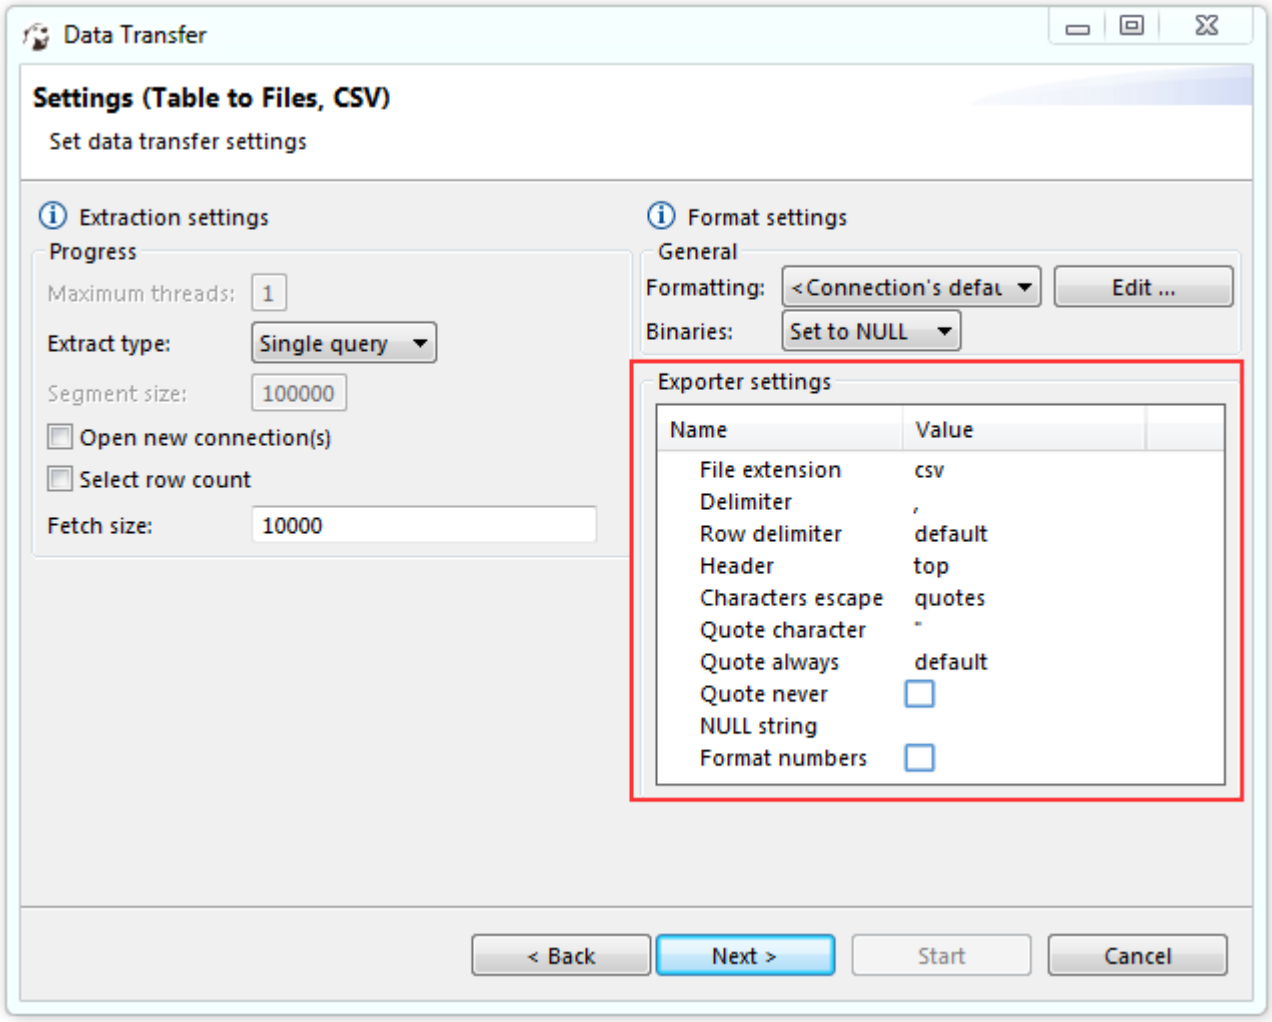

4. Set options for output files or clipboard. Note: Timestamp pattern here is used for target file name pattern:

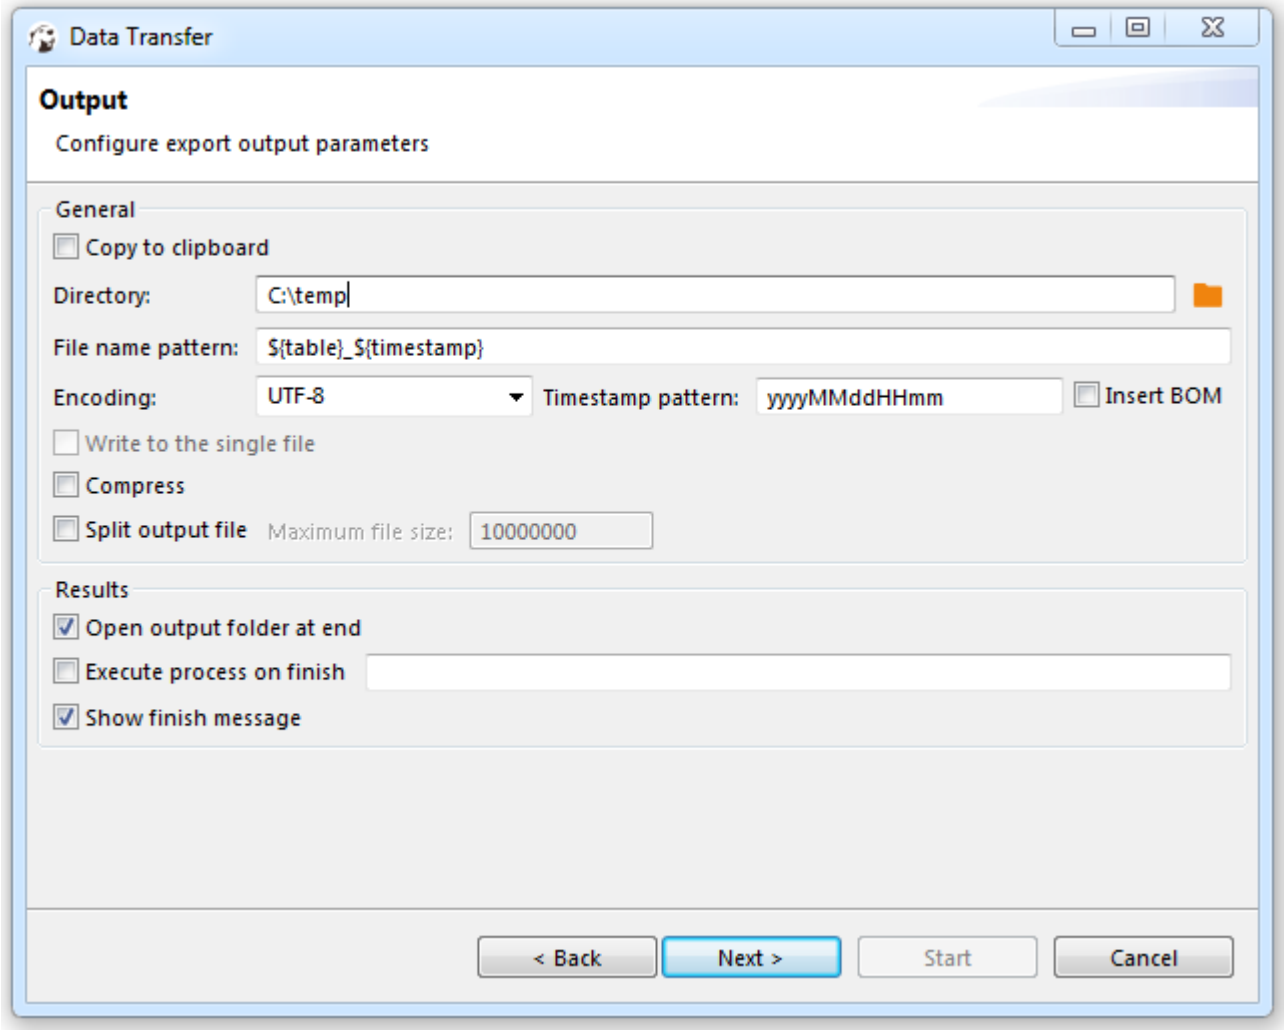

5. Review what and to what format you will export. You can also save all your settings as a task in this step or change task variables:

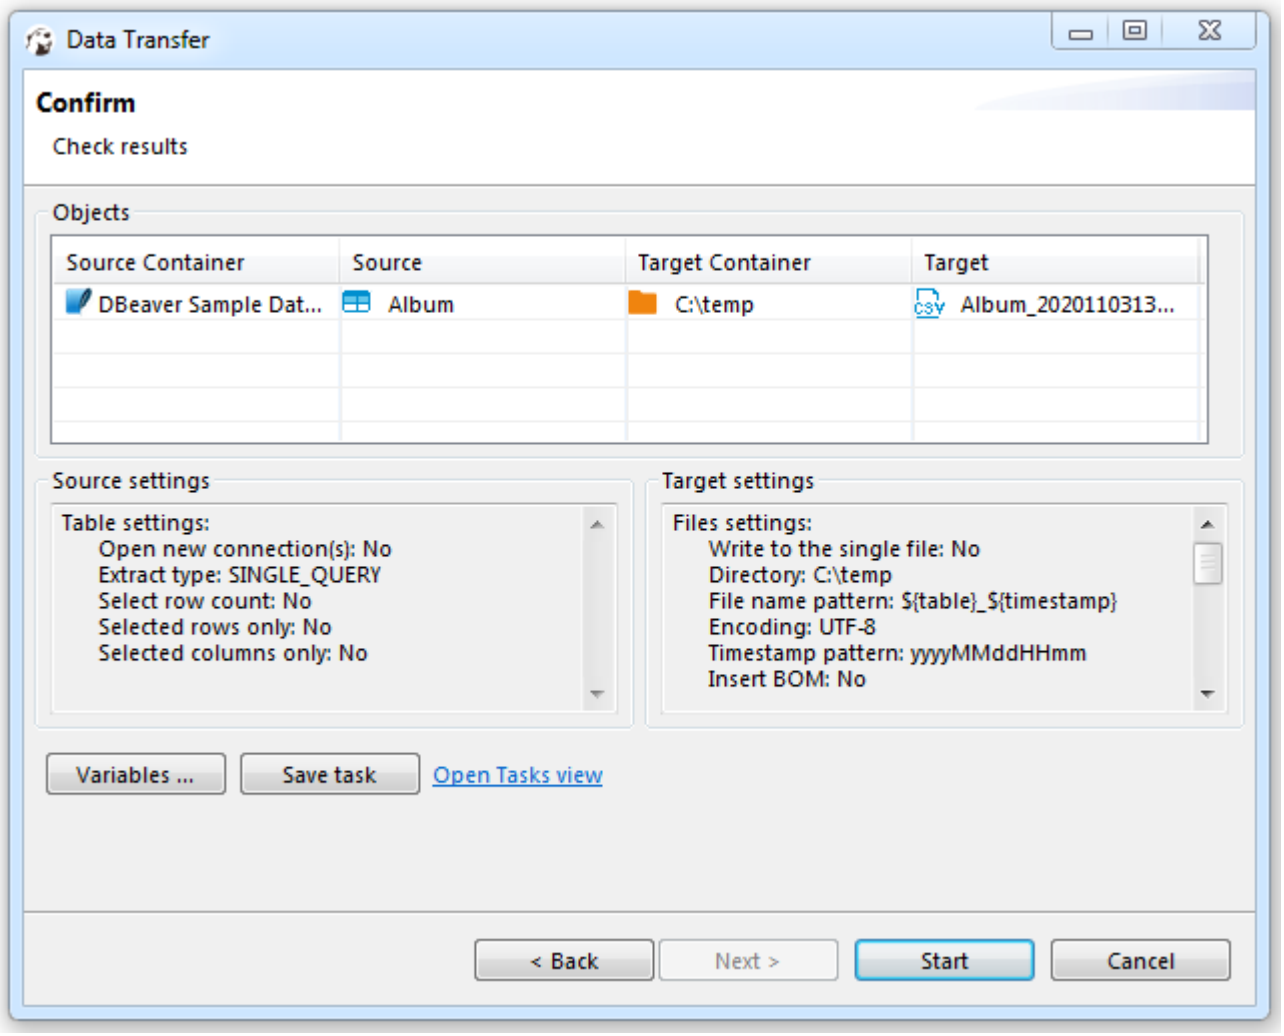

6. Press finish. See extraction progress. Actual data extraction will be performed in background, you can keep working with your database during export process. Note: avoid changing data in tables you selected for export while export is running. In the end you will see status message:

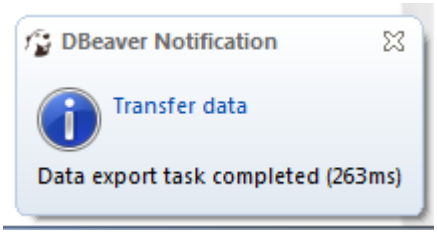

Importing data from CSV format

You can import data from CSV file(s) directly into your database table(s).

1. Select a table(s) to which you want to import data. In the context menu choose "Import Data":

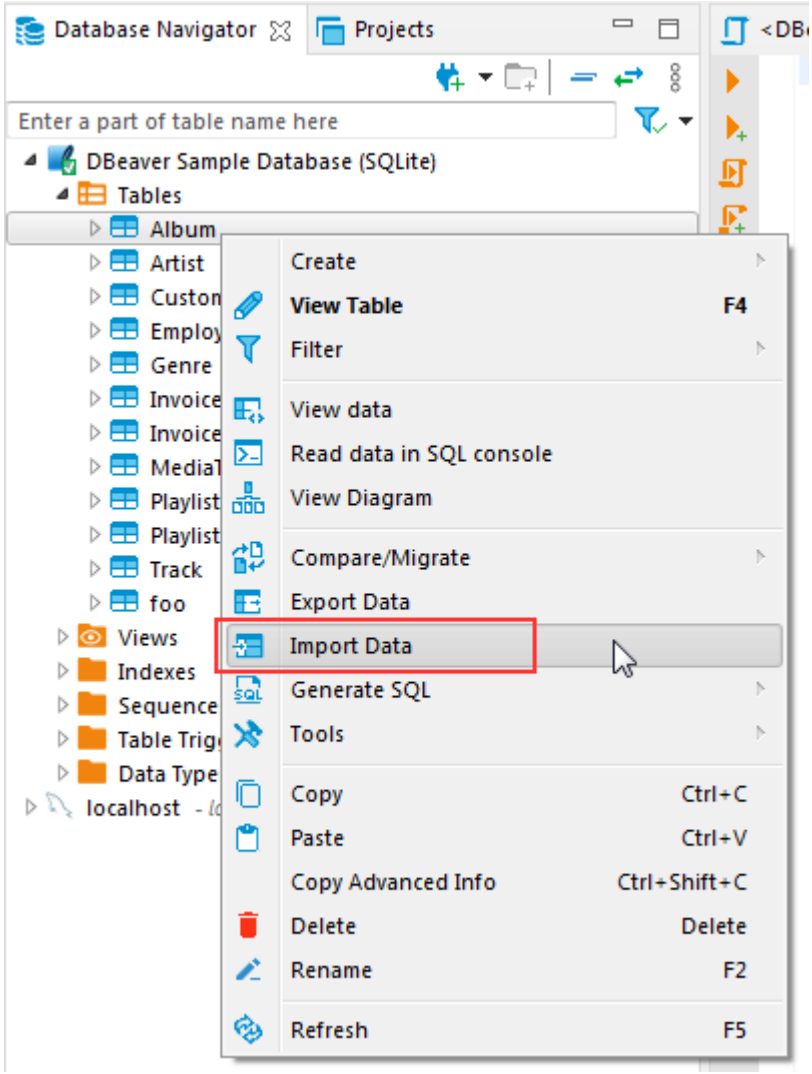

2. Choose import format (CSV):

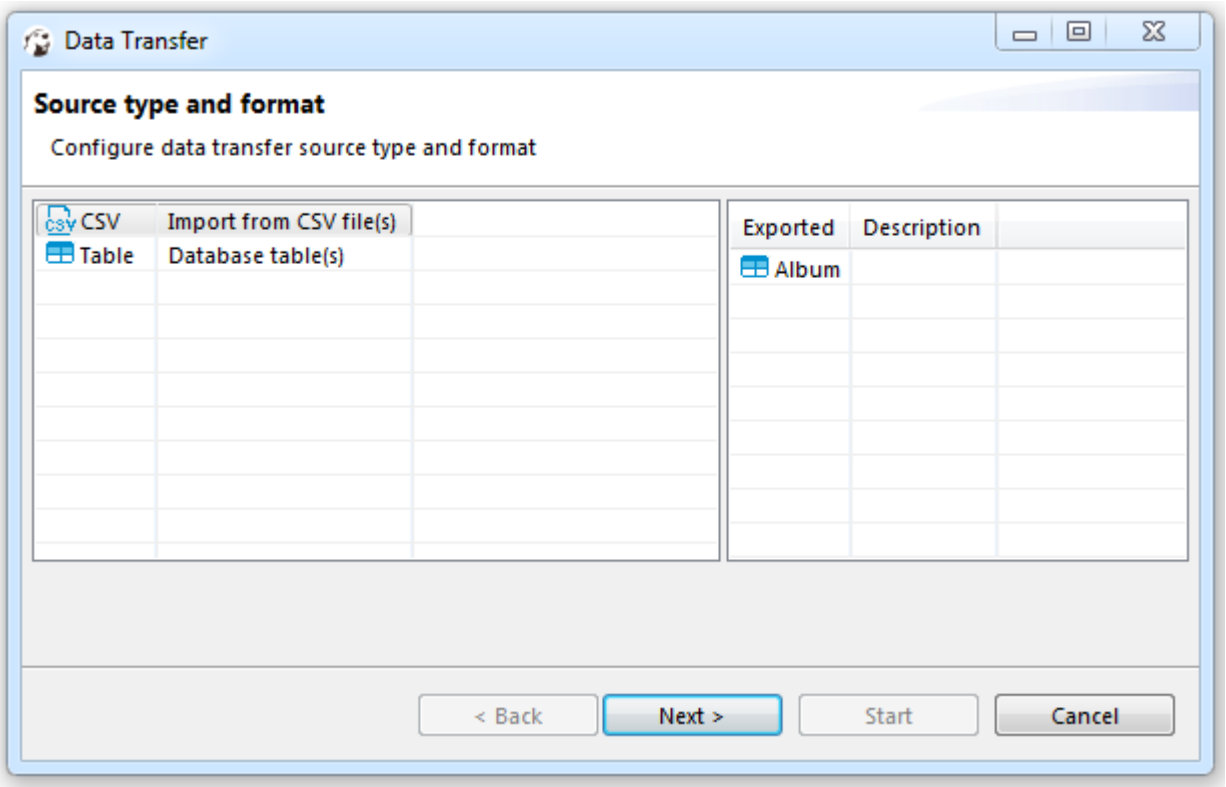

3. Select input CSV file for each table you want to import and you can change Importer settings (format specific) on this step:

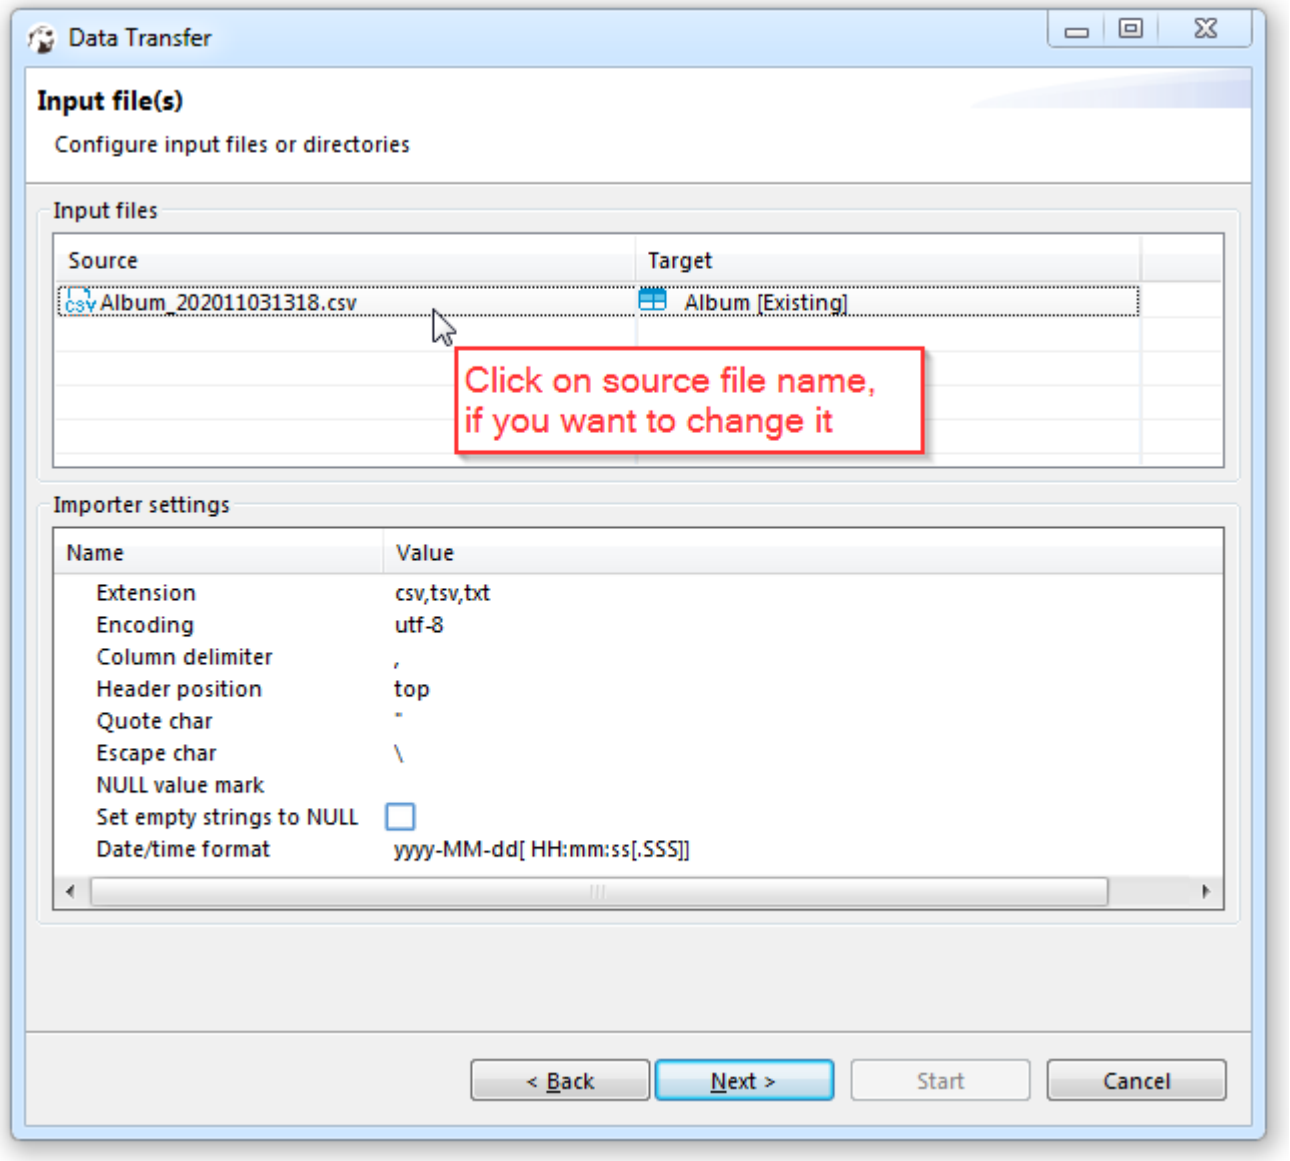

4. Set CSV-to-table mappings. You need to set some column in CSV file for each database table column. You can skip some column at all (in target table column value will be set to NULL). You can set some constant value for table column if there is no source column for it in CSV.

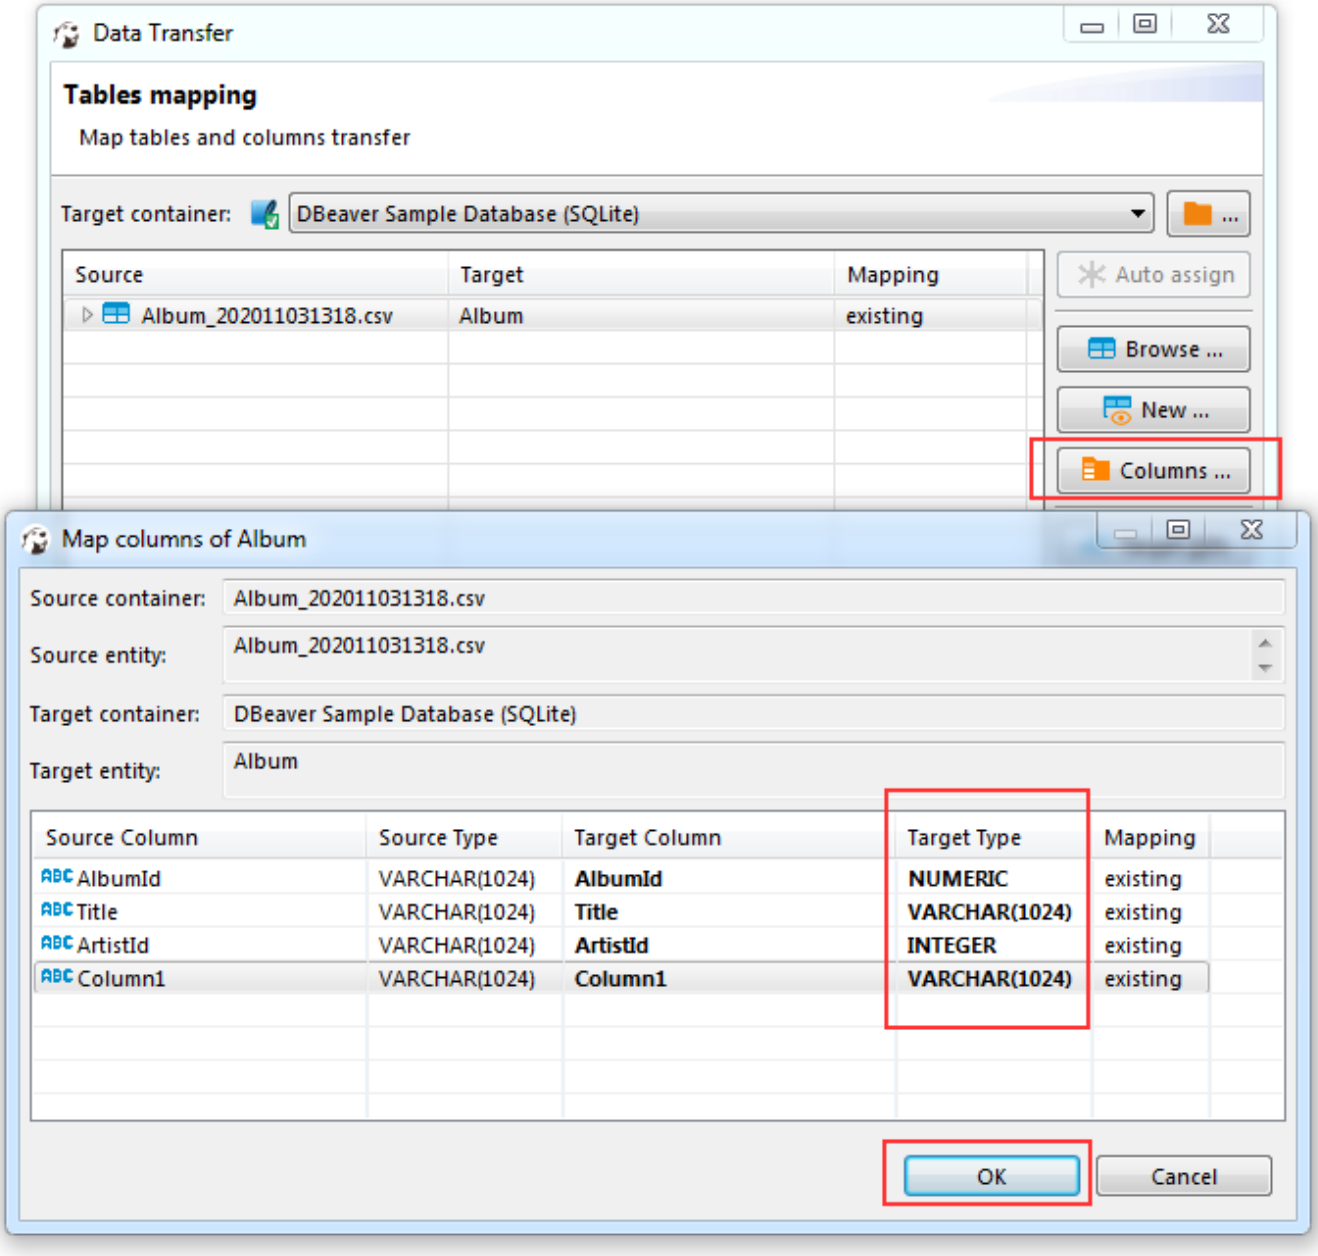

5. Set options for loading data in database. These options may affect performance:

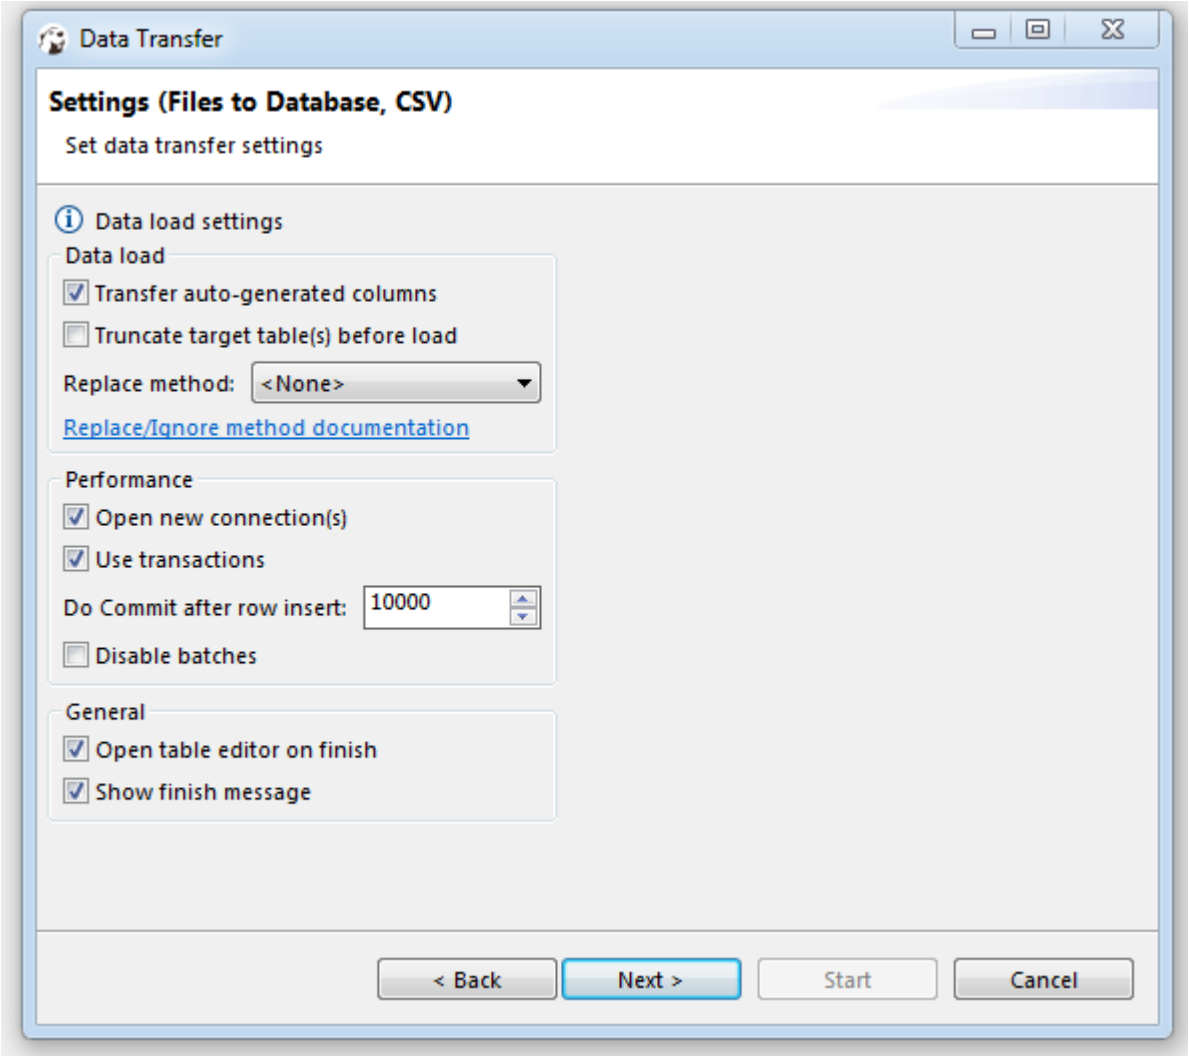

About replace method option you can read here.

6. Review what file(s) and to what table(s) you will import. You can also save all your settings as a task in this step:

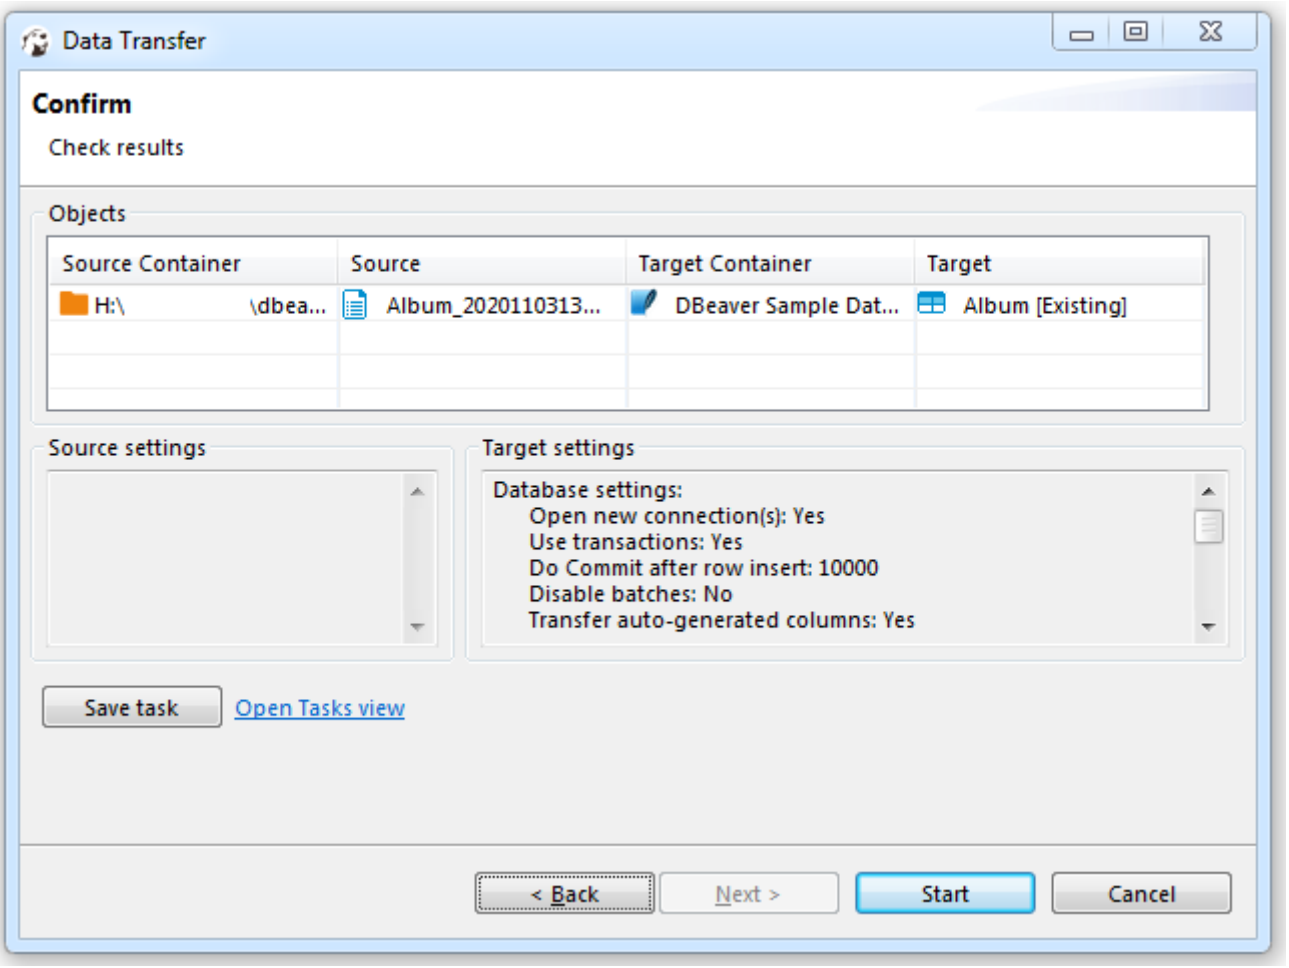

7. Press finish. See extraction progress. Actual data loading will be performed in background, you can keep working with your database during export process. Note: avoid changing data in tables you selected for import while import is running. In the end you will see status message:

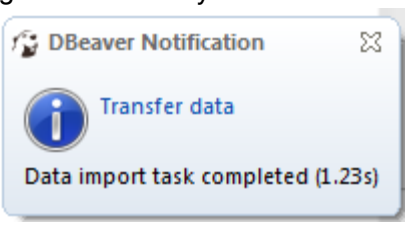

Related topic: Migrating table(s) data to another database table(s)

Data migration

DBeaver supports data migration from tables of one database to tables of another one.

To perform data transfer, please, follow the steps below.

Step 1: Define the data source

In the **Database Navigator** select one or more tables you want to export. In the context menu choose "Export Data". (Note: you also can export data from the custom SQL queries results. For that in the results context menu choose "Export results").

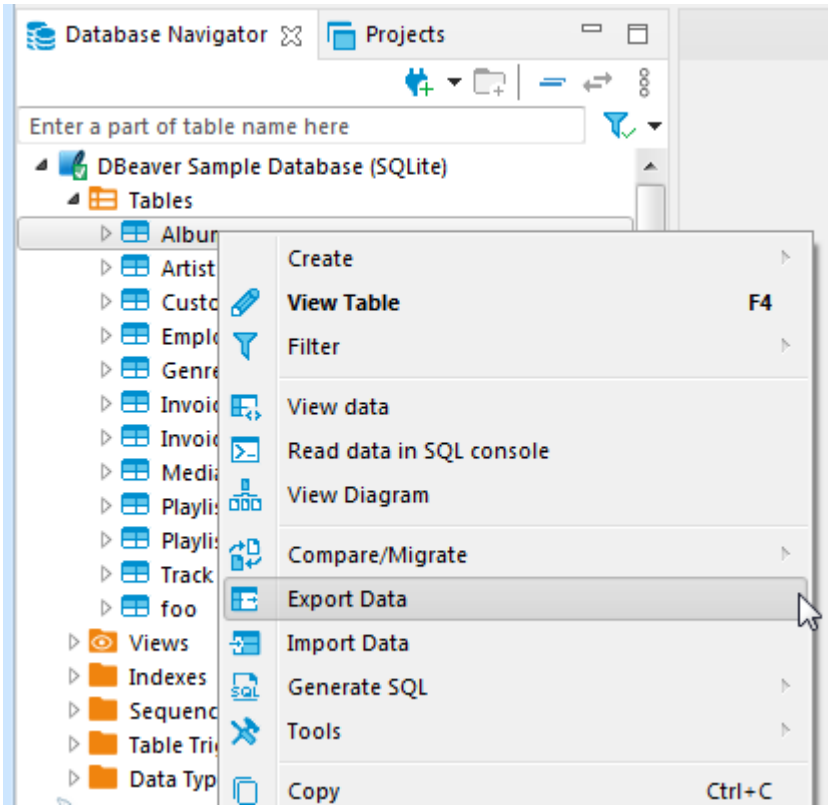

Step 2: Define data transfer target type

In the opened dialog box choose **Database** type as the data transfer target and press **Next**.

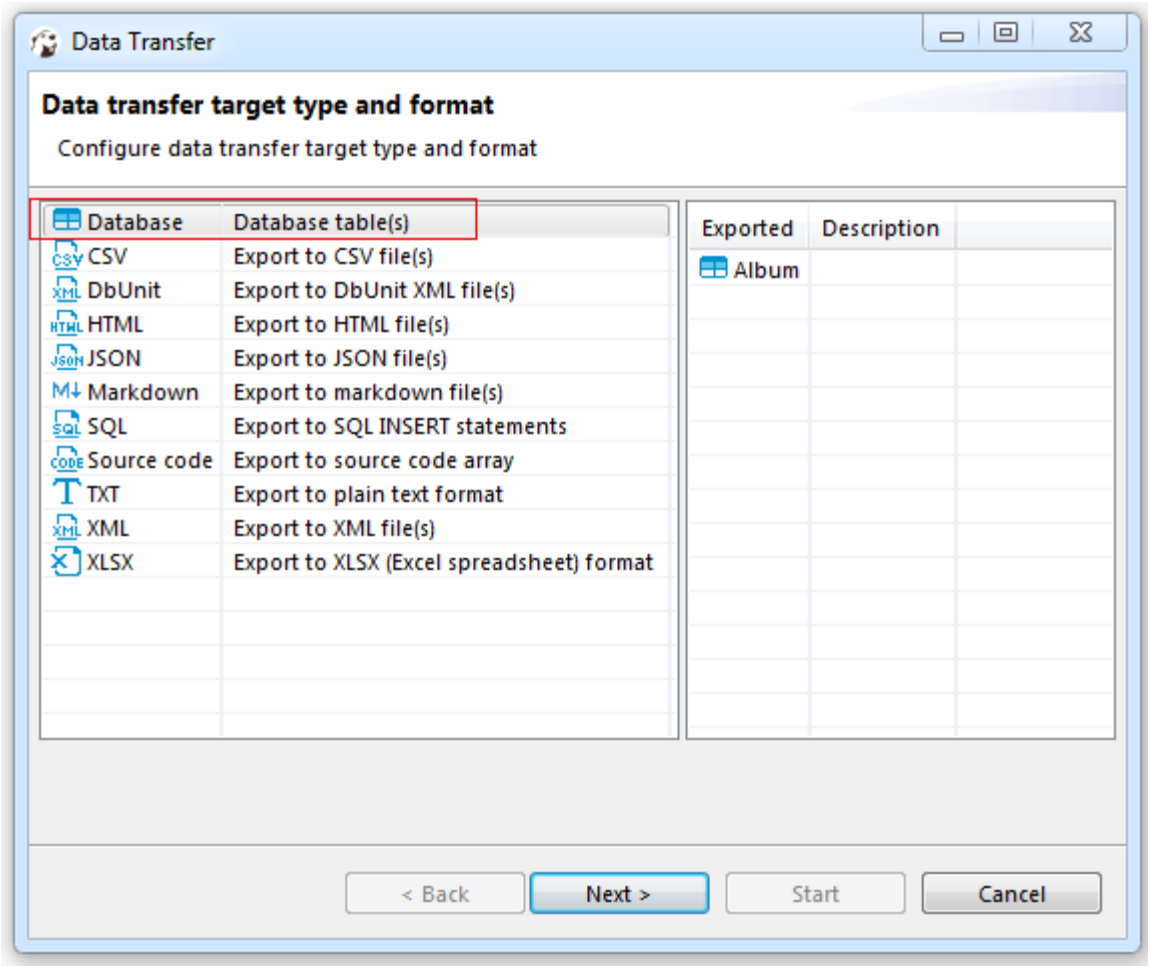

Step 3: Define data mapping

For proper table mapping the following options are available:

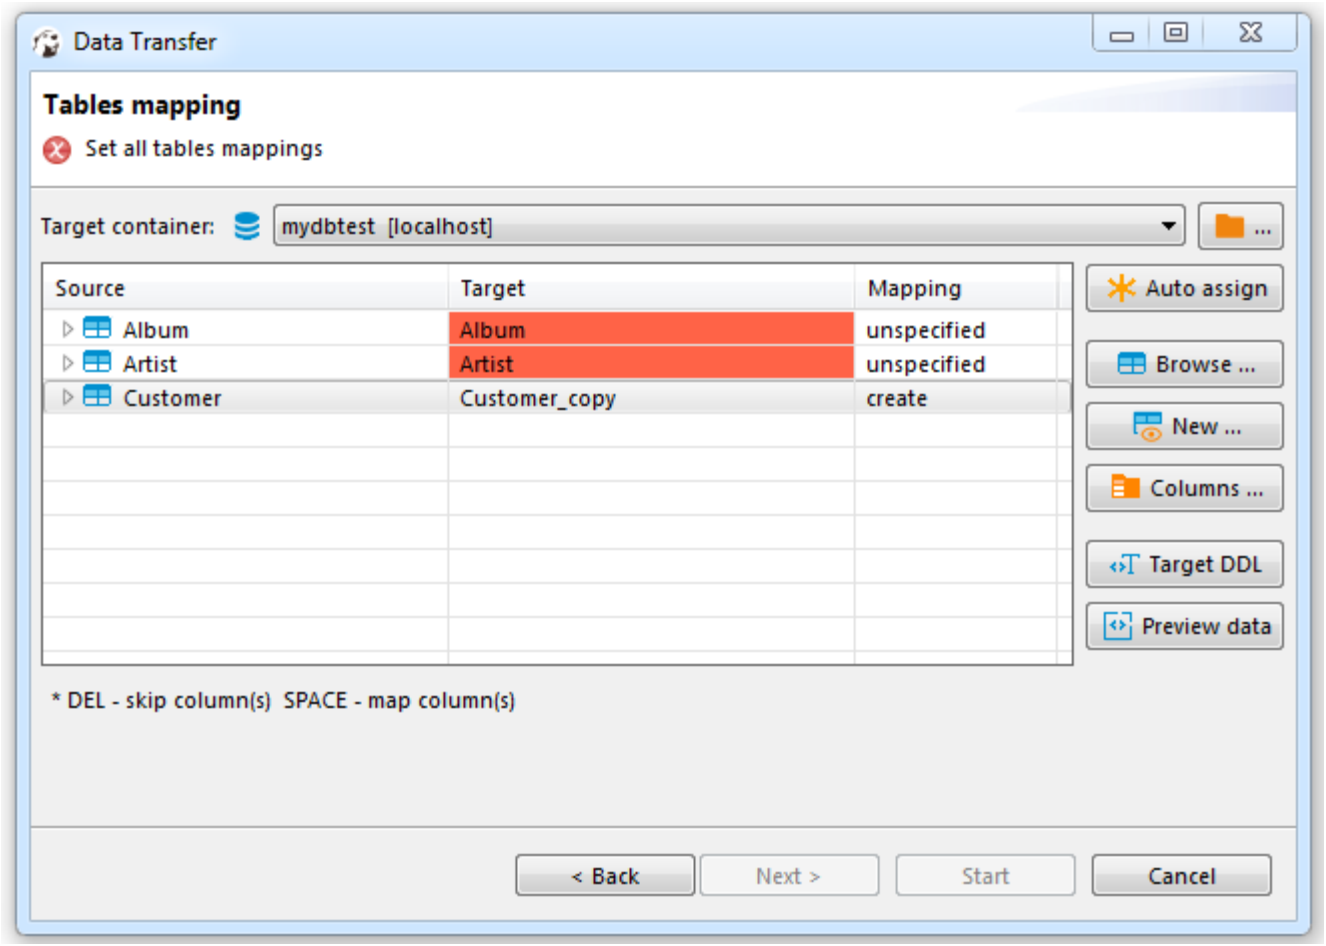

- **Target container** - defines a database or a scheme where the data will be transferred to. Press button and choose the container.

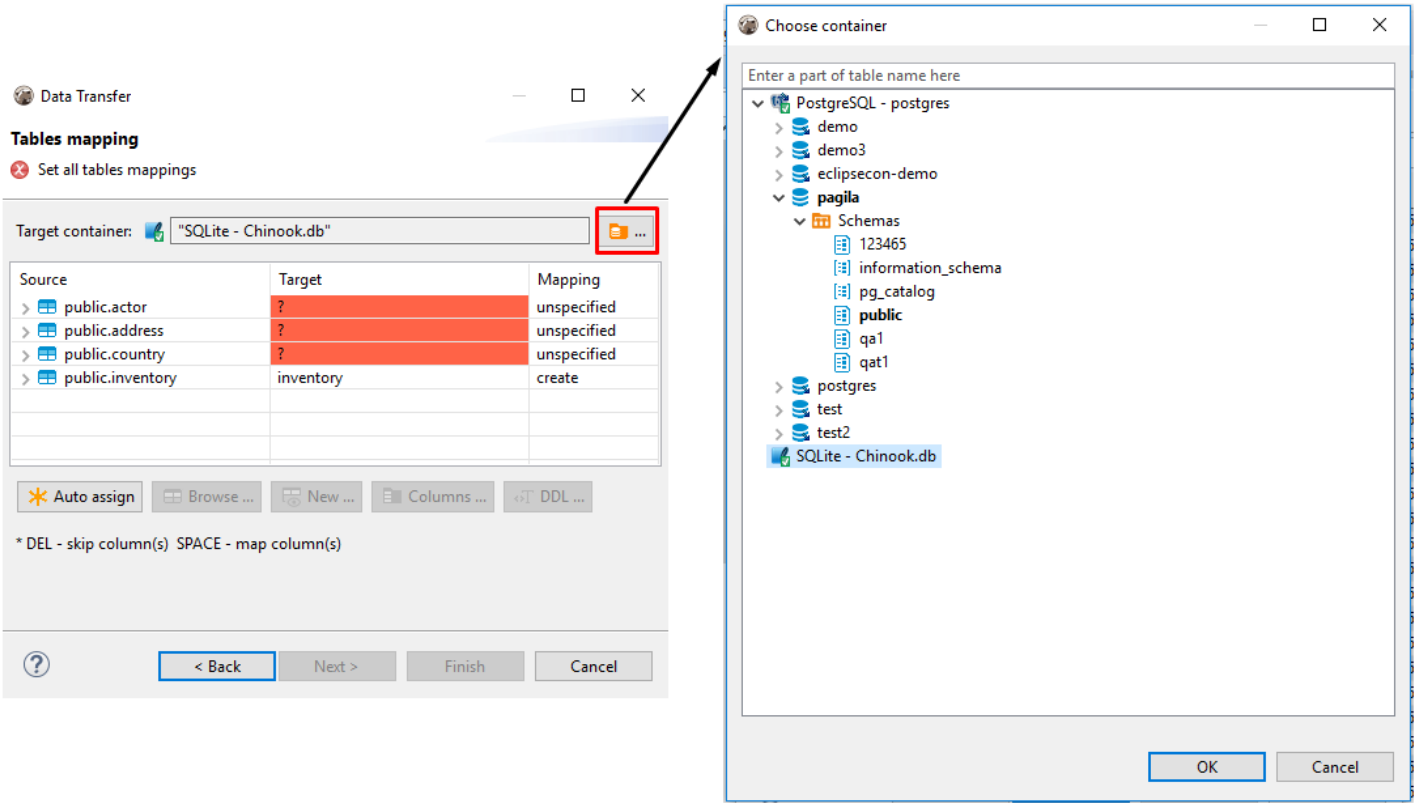

- **Source** - contains names of all the tables selected at step 1. You can also see the list of columns existing in the source table by pressing >
- **Target** - contains names of the tables where the data will be transferred to.
- **Mapping** - contains the list of actions to be applied to the source data on data transfer. The following options are available:

  Create- the source data will be populated into a newly created table or column of the target container.
  Skip - the source data will not be transferred to the target container.
  Existing - the source data will be transferred to the table that already exists in the target container.
  Unassigned - this value is set by default when there is no target defined.

If cells are marked as ? it means that in the target table there are no source tables with matching names, otherwise the names will be filled in automatically.

You can define a target table by clicking on a cell in the **Target** column and entering its name, or press the **New** button New ... and enter a new name in the opened dialog box.

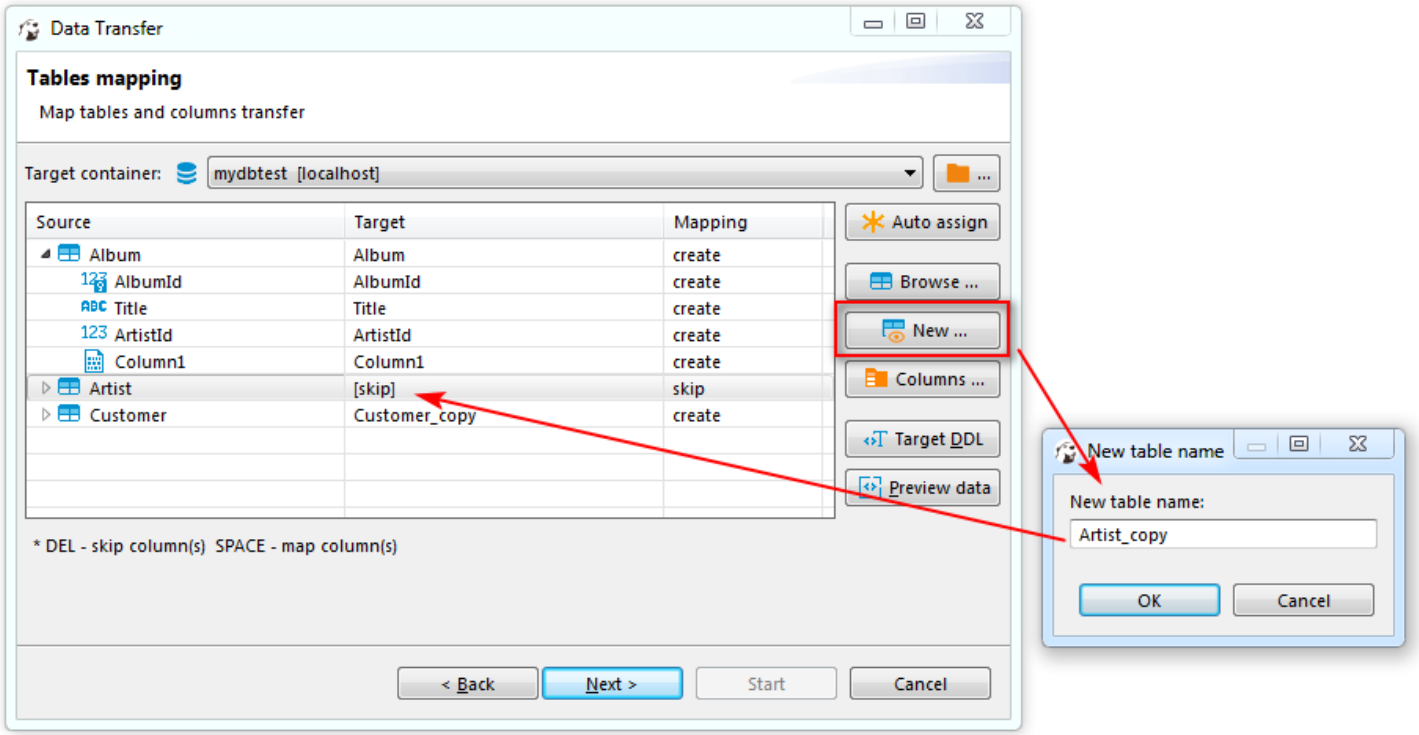

You can also choose a name for a target table from the drop-down list.

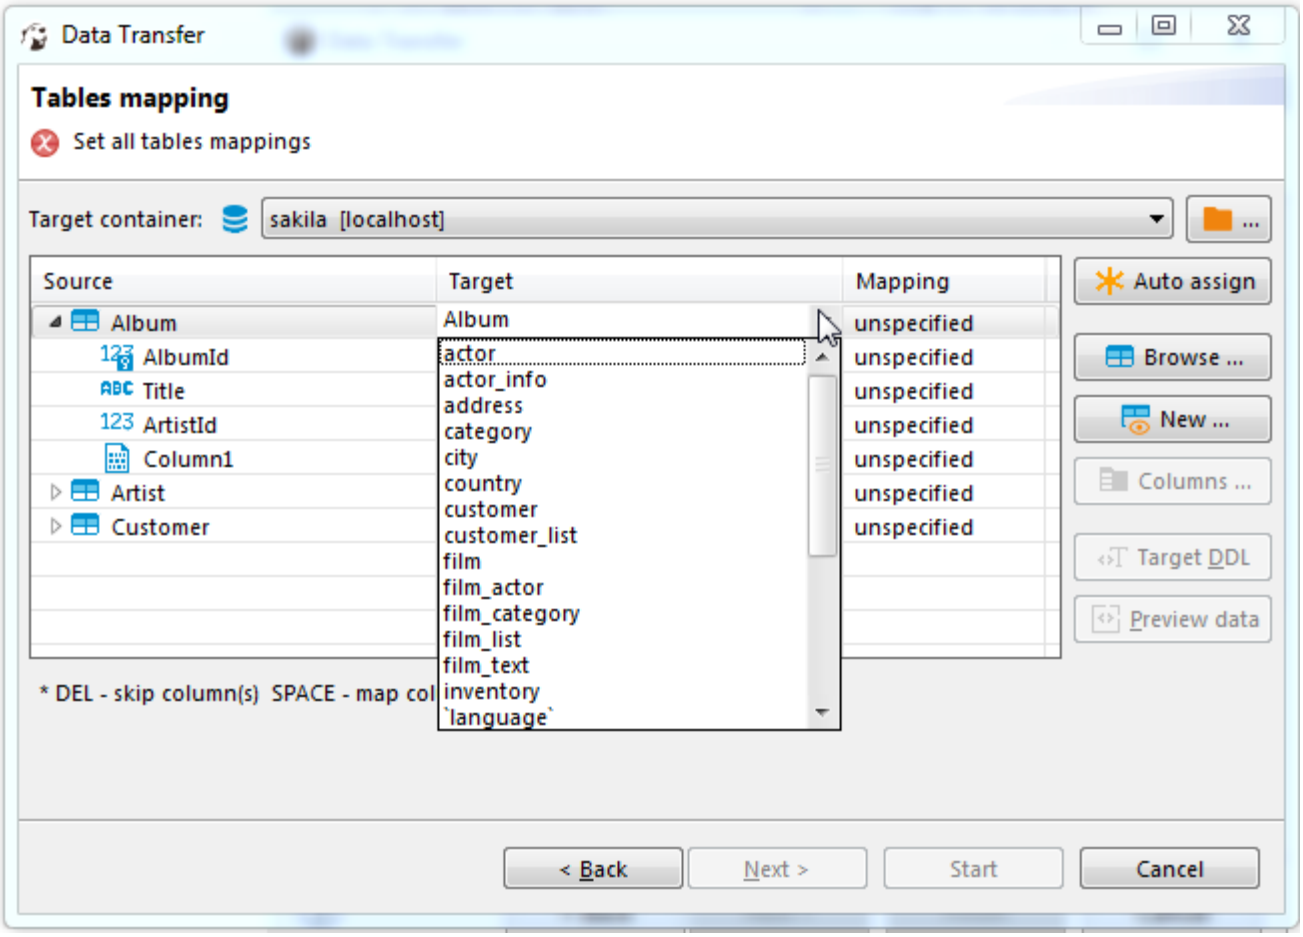

Or select from the list of tables already existing in the target container by pressing the **Browse** button 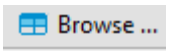

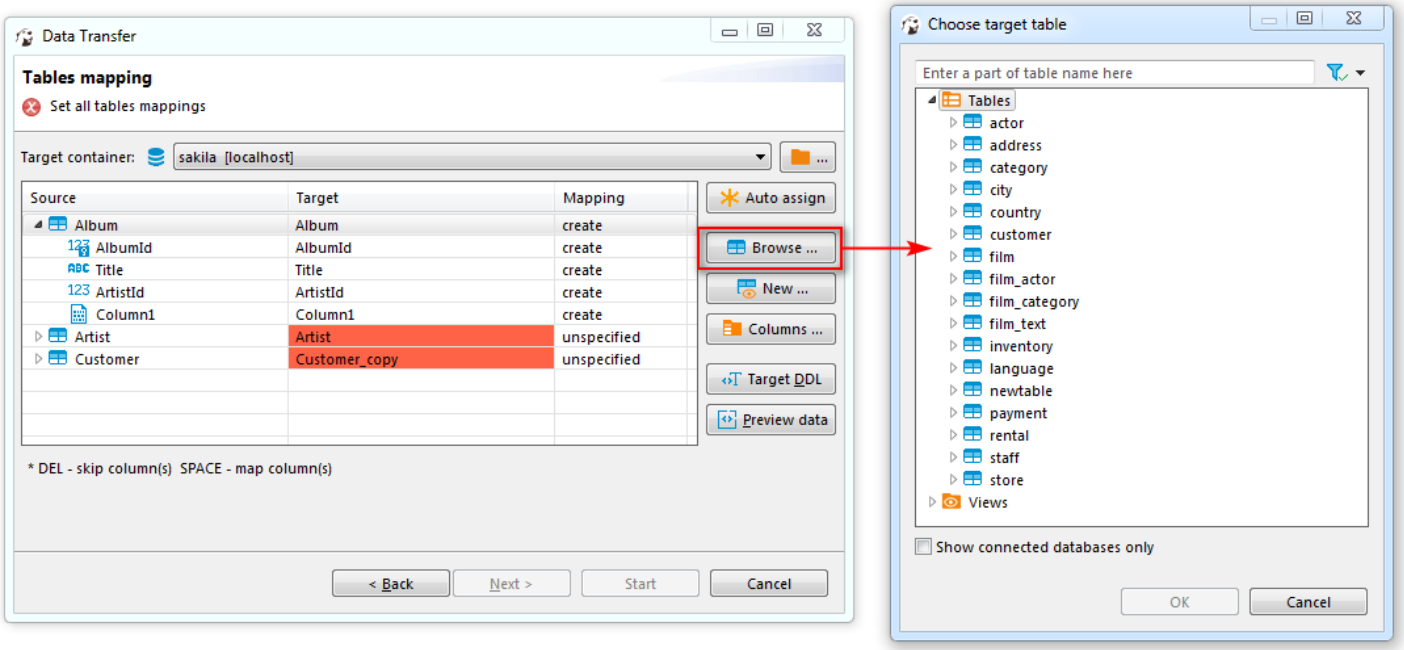

To define mapping setting for a column in a target table, release the list of source table columns by pressing next to the table's name in the **Source** column, then click the name of the target column and enter a new one or select one from the dropdown list. To collaps the list, press

If you want tables of the target container to be named like those of source, press the **Auto assign** button 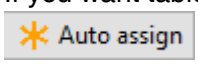

and the **Target** column will be automatically populated.

You can also define the names of target columns, as well as their data types, by clicking a row with a table name and pressing the **Columns** button

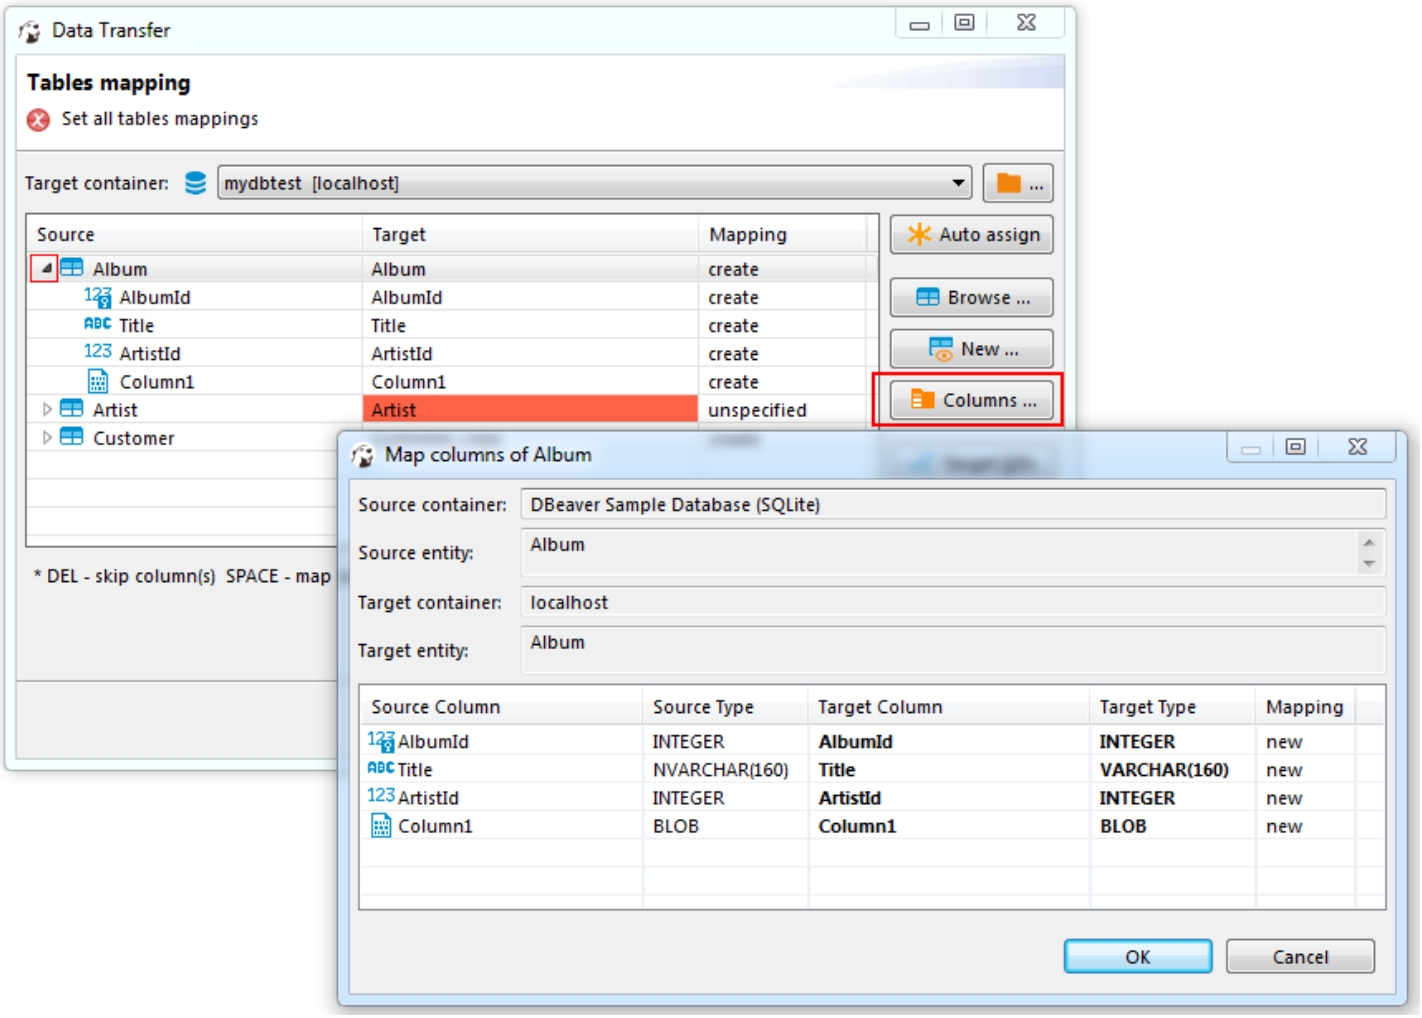

The following elements are available here:

- **Source column** -this column contains names of columns existing in the selected source table;
- **Source type** - this column contains the list of data types assigned to the columns in the selected source table;
- **Target column** - this column contains names of target table columns where the data from source column will be transferred to. To change the name, click the cell and enter a new name.
- **Target type** - this column contains the list of data types that will be assigned to the columns in the target table.

  IMPORTANT: Sometimes data types that are supported on the source database are not supported on the target or vice versa.
  To set a data type for a column in a target table, click the cell in the **Target Type** column and select one from the dropdown list of data types supported on the target.
- **Mapping** - this column contains the list of actions to be applied to the data on data transfer.

To change mapping type click a cell in **Mapping** column of **Table mapping dialog box** and select the required mapping type.

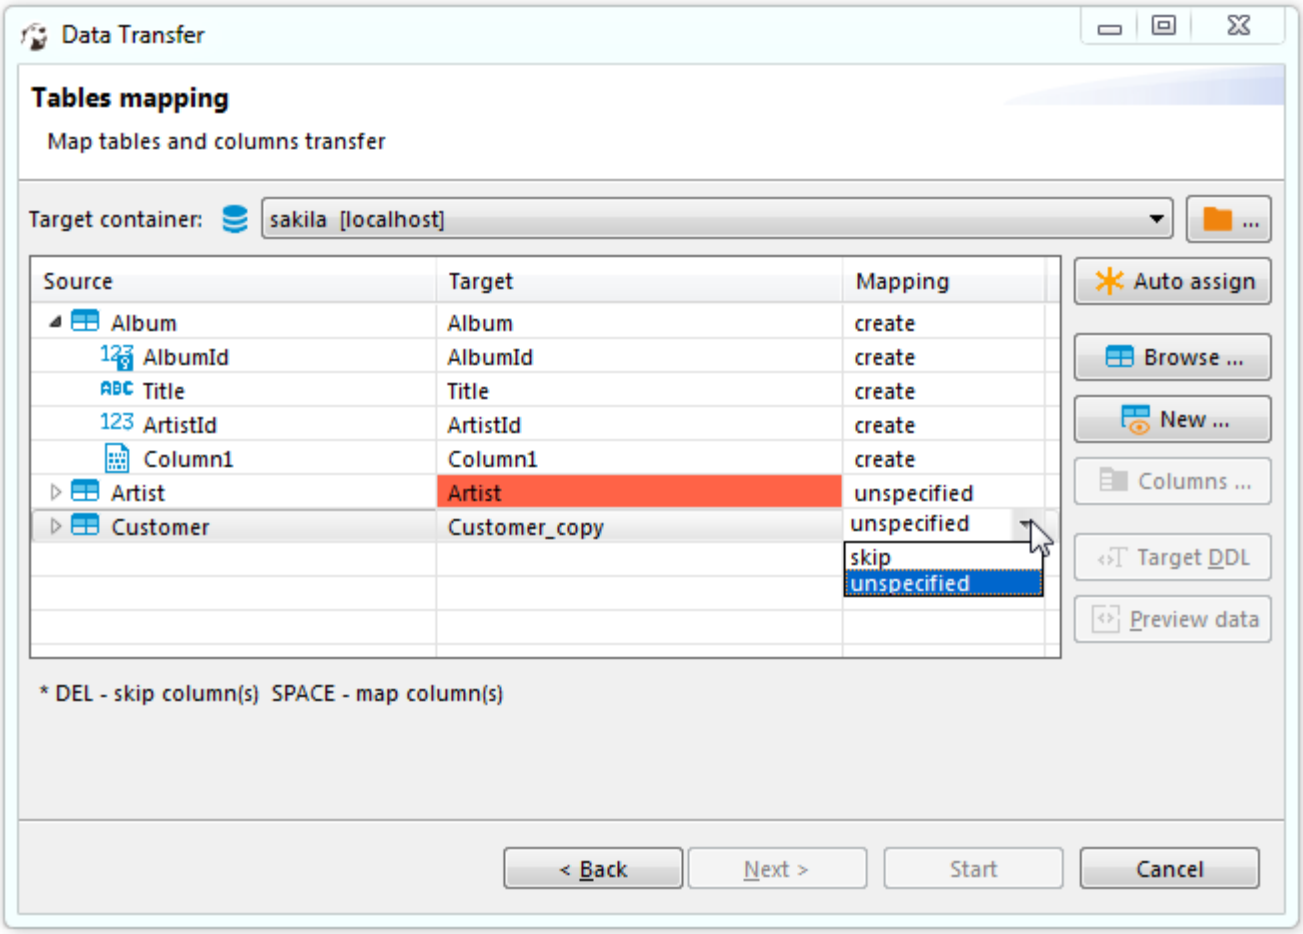

You can also view the SQL script that will be run on data transfer by pressing the **Target DDL** button

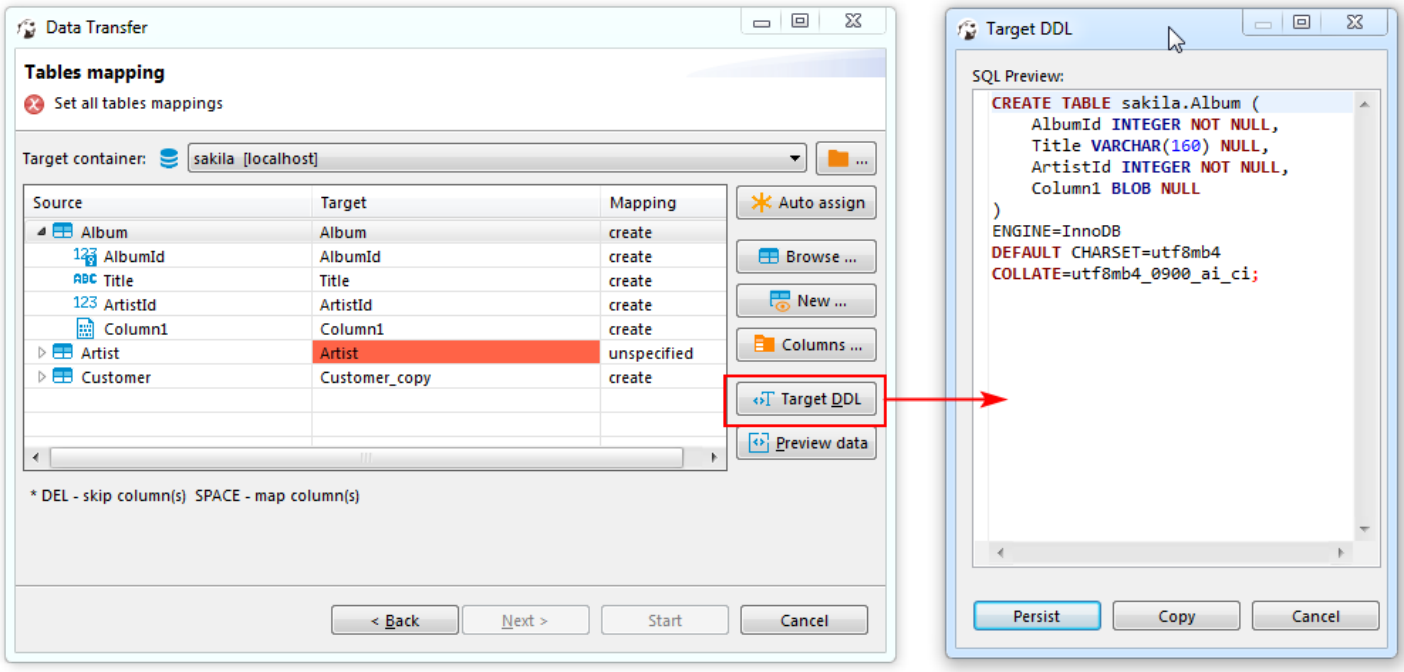

If you want to see in advance the preview of the imported data, you can select the **Preview data** button

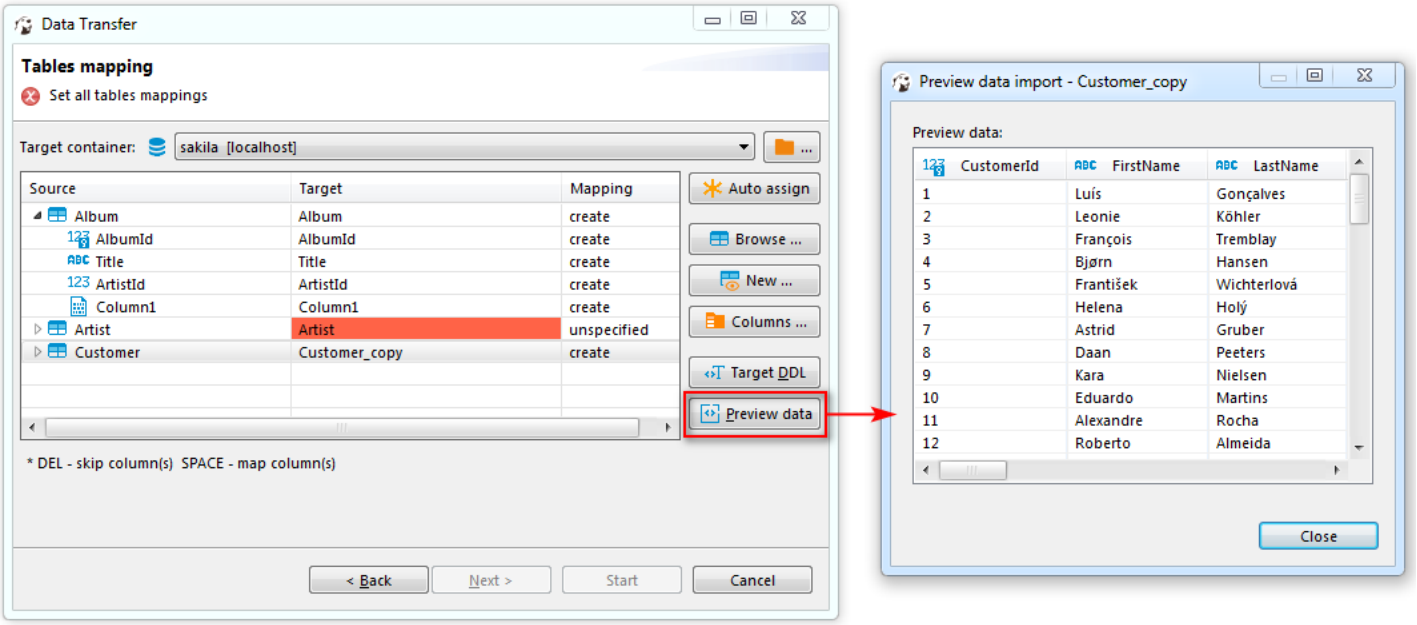

The following keyboard shortcuts for easy navigation within the mapping table area of **Table mapping** screen are supported:

Shortcut	Action
Up	Moves one row up.
Down	Moves one row down.
Right	Releases the list of source table columns.
Left	Swaps the list of source table columns.
Space	Auto-assigns the target.
Del	Sets mapping type to skip.

Configure data mapping and press **Next**.

Step 4: Define export settings

Data export settings are grouped into **Extraction settings** and **Data load settings**.

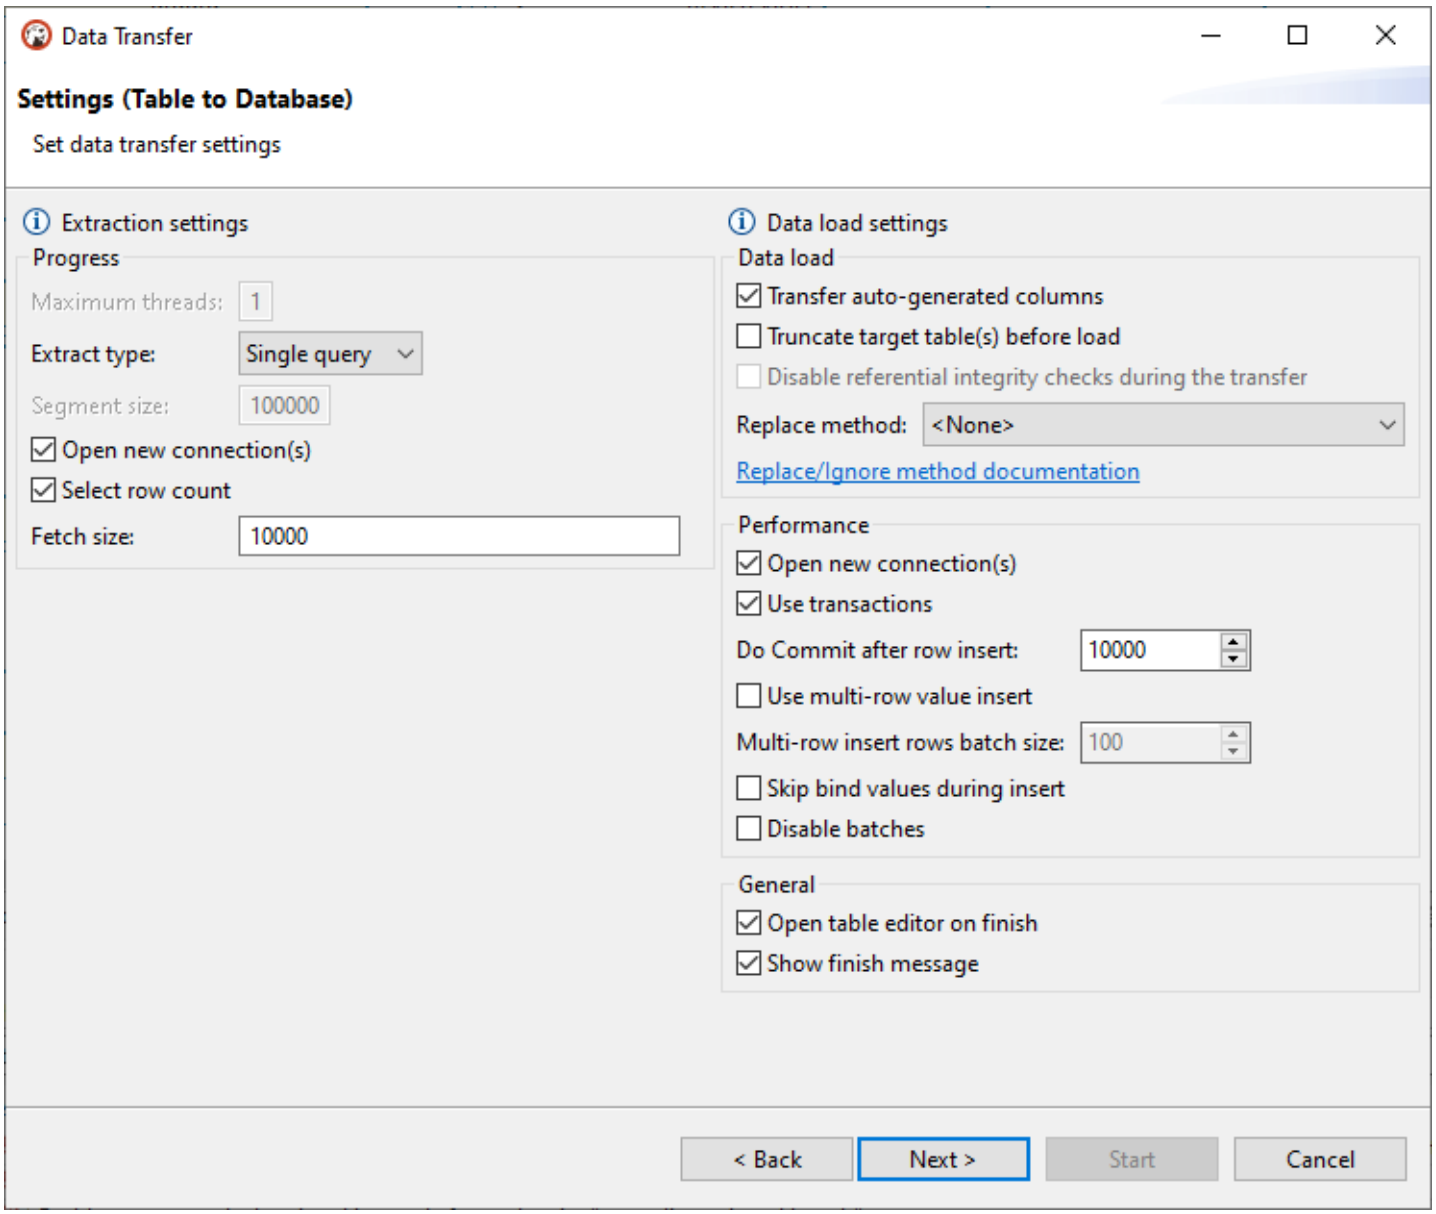

Extraction Settings

Extraction settings define how the data will be pulled from the source. The following options are available:

Option	Description
Maximum threads	Defines a number of threads to be used for data transfer.
Extract type	Select **Single query** option if your data load is not too big. Select **By segments** option if you need to migrate a solid amount of data. When this options is selected you can set the **Segment size** value, that is to define a number of rows to be transferred in each segment.
Open new connections	If selected, a new connection will be opened and data transfer will not interfere with other calls to the database whose data is being transferred.
Select row count	If selected, a progress bar displaying data migration process will be shown.

Data load settings

Data load settings define how the extracted data will be pushed to the target. The following options are available.

Option	Description
Truncate data load table before load	Select this check-box only if you want all the data be cleared from the target table. Be very careful with this option!
Replace method	Select this drop-down list if you want change import behavior in case of duplicate primary key value. Data Import and Replace page
Open new connections	Use this option to speed up data transfer. If selected, a new connection will be opened and data transfer will not interfere with other calls to the database where data is being transferred to.
Use transactions	This option allows to speed up data transfer and to define the number of rows for each transaction by setting **Commit after insert of** parameter.
Disable batches	Select this check-box if you want disable the use of batch imports. Import will be made row by row. Enabling this function will show all import errors, but make the import process slower.
Open table editor on finish	If selected, the table editor to be opened when data transfer is over.
Show finish message	If selected, a notification message will be shown when transfer is over.

Step 5: Confirm

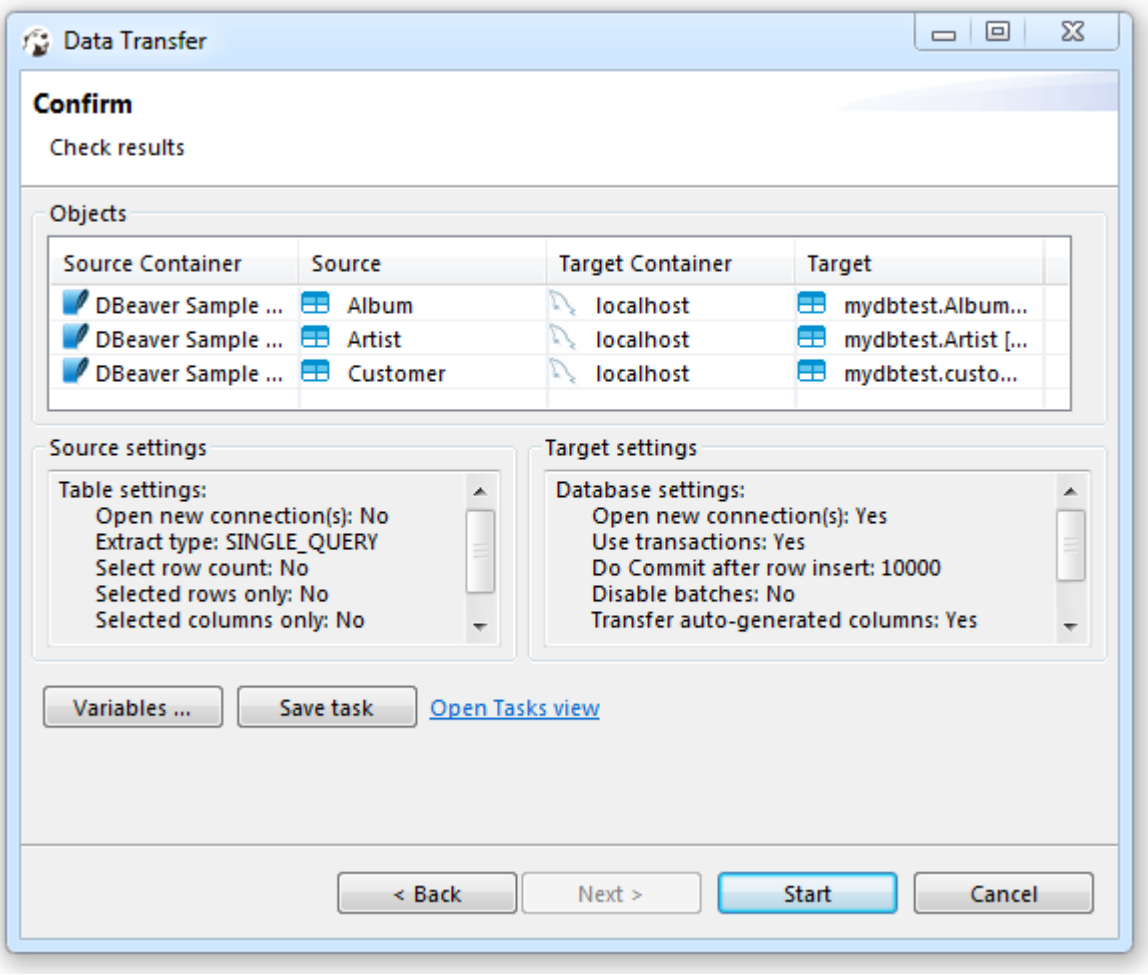

Check out data transfer settings and press **Start** or save as task.

Data Import and Replace

Data Import and Replace.

Sometimes there are situations when you want to ignore the current primary key value when importing into a table. Some databases have syntax constructs in addition to the INSERT INTO that may help.

The choice of the replacement method is in the import settings - in "Data load settings".

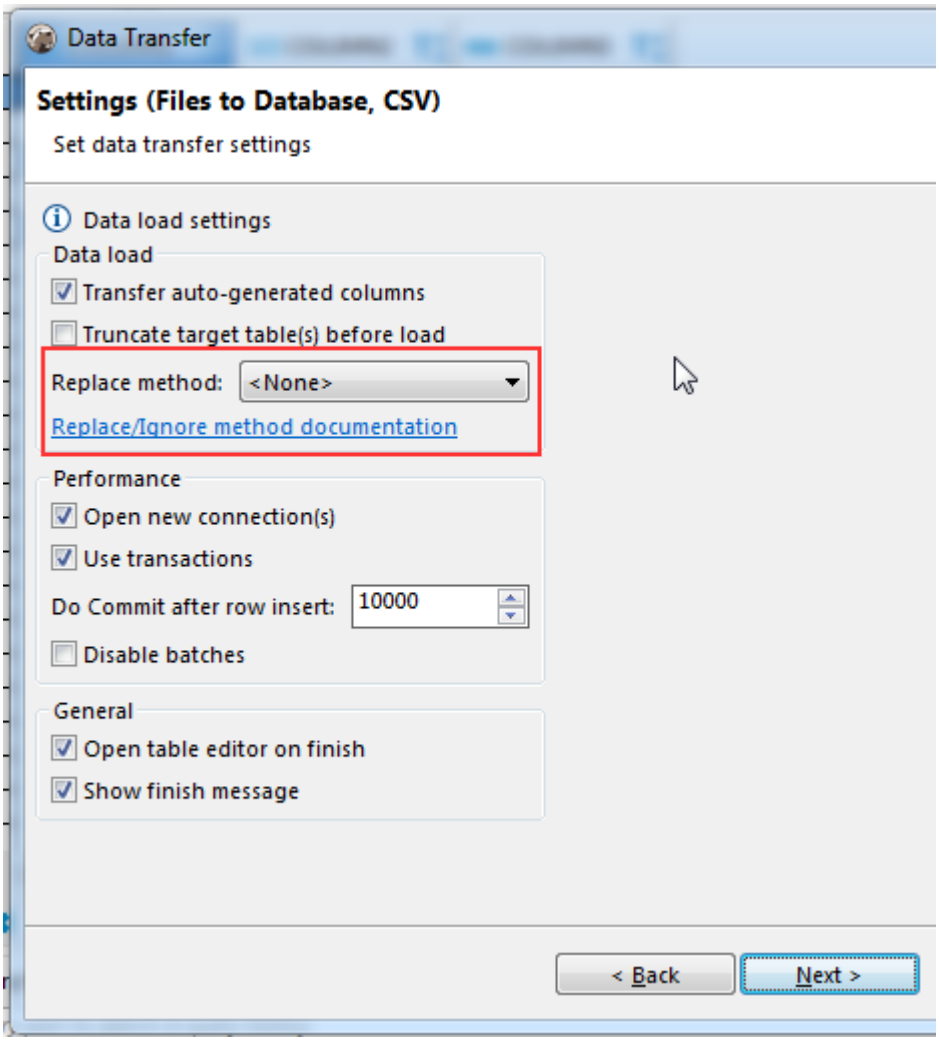

By default, the selection is `<None` >, you can select other options from the drop-down list. The options available depend on the target database you are importing to.

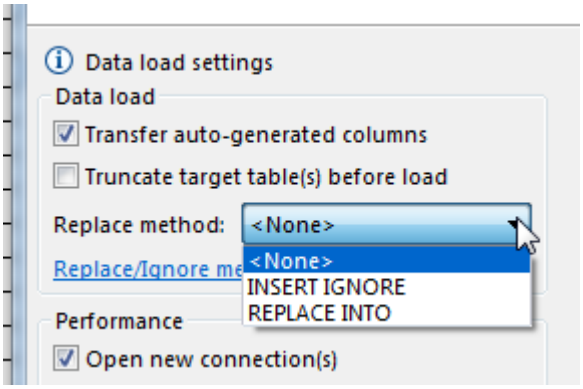

The database can only support the replace method or only the ignore method. In this case, the list of methods will consist of only one item except `<None>`. If the base does not support replacement methods or we have not added an implementation yet, then the combo with the list will be disabled.

Further you will find a list of databases supporting these methods and examples of syntax.

Let's take a look at an example of how this works. We use a simple small slightly modified Sakila (MySQL) table - sakila.language

```
CREATE TABLE language_insert (
language_id tinyint unsigned NOT NULL,
name char(20) NOT NULL,
last_update timestamp NOT NULL,
PRIMARY KEY (language_id)
);
```

```
INSERT INTO sakila.language_insert (language_id,name,last_update) VALUES
(1,'English','2006-02-15 05:02:19.0'),
(2,'Italian','2006-02-15 05:02:19.0'),
(3,'Japanese','2006-02-15 05:02:19.0');
```

If we try to execute this request twice, then the second time we get the following error: SQL Error [1062] [23000]: Duplicate entry '1' for key 'language_insert.PRIMARY' (This message may look different in other databases).

Let's take a new .csv file with the following content and try to use replace methods

```
"language_id","name","last_update"
1,Spanish,"2020-04-20 05:02:19.0"
2,Russian,"2020-05-20 05:02:19.0"
3,Belgian,"2020-06-20 05:02:19.0"
4,Mandarin,"2006-02-15 05:02:19.0"
5,French,"2006-02-15 05:02:19.0"
6,German,"2006-02-15 05:02:19.0"
```

If we set the "INSERT IGNORE" method in the settings, the result of the insert will look like this:

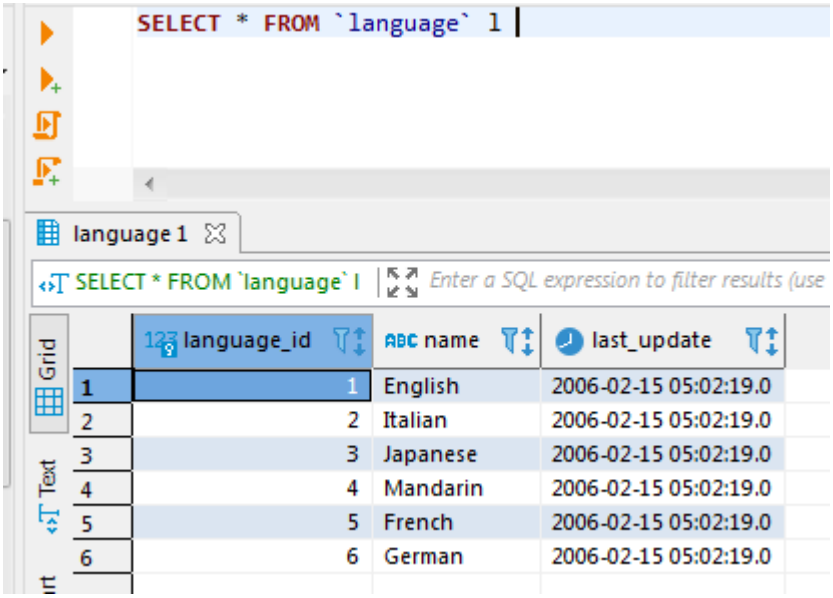

There will be no insertion errors, the first three lines will not change, and the fourth to sixth lines will be added to the table.

If we set the "REPLACE INTO" method in the settings, the result of the insert will look like this:

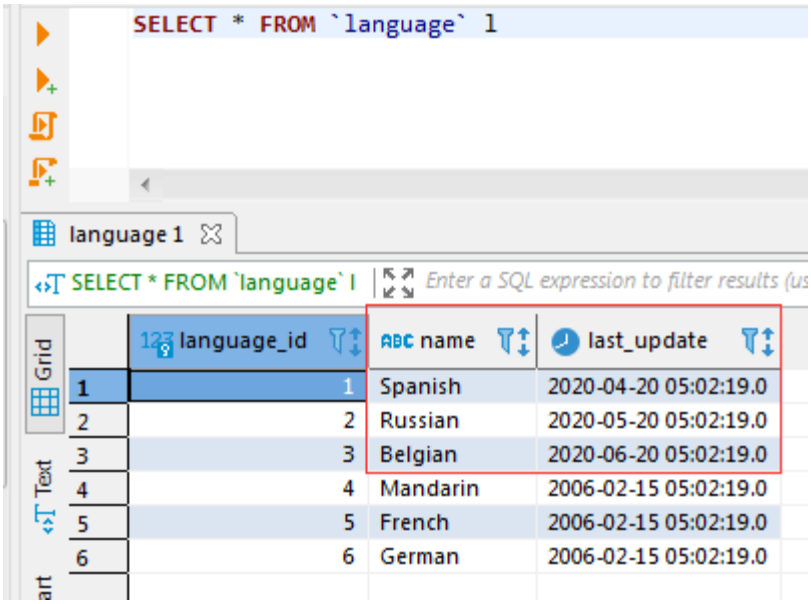

There will be no insertion errors, the first three lines will be replaced and the fourth to sixth lines will be added to the table.

Which databases support replace/insert methods?

MySQL

INSERT IGNORE and REPLACE INTO

Insert examples:

"INSERT IGNORE"

```
INSERT IGNORE INTO language_insert(language_id, name, last_update)
VALUES(1, 'English', '2006-02-15 05:02:19.0');
```

"REPLACE INTO"

```
REPLACE INTO language_insert(language_id, name, last_update)
VALUES(1, 'English', '2006-02-15 05:02:19.0');
```

SQLite

Documentation
"INSERT OR IGNORE" and "INSERT OR REPLACE"

Insert examples:

"INSERT OR IGNORE"

```
INSERT OR IGNORE INTO language_insert(language_id, name, last_update)
VALUES(1, 'English', '2006-02-15 05:02:19.0');
```

"INSERT OR REPLACE"

```
INSERT OR REPLACE INTO language_insert(language_id, name, last_update)
VALUES(1, 'English', '2006-02-15 05:02:19.0');
```

PostgreSQL

Available for PostgreSQL version 9.5.
"ON CONFLICT DO NOTHING" and "ON CONFLICT DO UPDATE SET"

Insert examples:

"ON CONFLICT DO NOTHING"

```
INSERT INTO language_insert(language_id, name, last_update)
VALUES(1, 'English', '2006-02-15 05:02:19.0') ON CONFLICT DO NOTHING;
```

"ON CONFLICT DO UPDATE SET"

```
INSERT INTO language_insert(language_id, name, last_update)
VALUES(1, 'English', '2006-02-15 05:02:19.0')
ON CONFLICT (language_id)
DO UPDATE SET (language_id, name, last_update) = (EXCLUDED.language_id, EXCLUDED.name, EXCLUDED.last_updat
```

FireBird

Available for FireBird version 2.1.
"UPDATE OR INSERT INTO"

Insert examples:

"UPDATE OR INSERT INTO"

```
UPDATE OR INSERT INTO language_insert(language_id, name, last_update)
VALUES(1, 'English', '2006-02-15 05:02:19.0');
```

Oracle

Available for Oracle version 11.2.
"INSERT IGNORE ROW INDEX"

Insert examples:

"INSERT IGNORE ROW INDEX"

```
INSERT /*+ IGNORE_ROW_ON_DUPKEY_INDEX(LANGUAGE_INSERT, LANGUAGE_INSERT_PK) */
INTO LANGUAGE_INSERT(LANGUAGE_ID, NAME, LAST_UPDATE) VALUES(1, 'English', TIMESTAMP '2006-02-15 05:02:19.0
```

Database backup/restore

Database Backup/restore

DBeaver supports native database backup/restore functions for following databases:

- PostgreSQL
- MySQL

Native backup restore differs from standard DBeaver data transfer feature. It uses database native dump formats and it may work much faster as it uses special utilities for direct high-performance database access.

These functions can be accessed from context menu Tools or main menu Database->Tools.

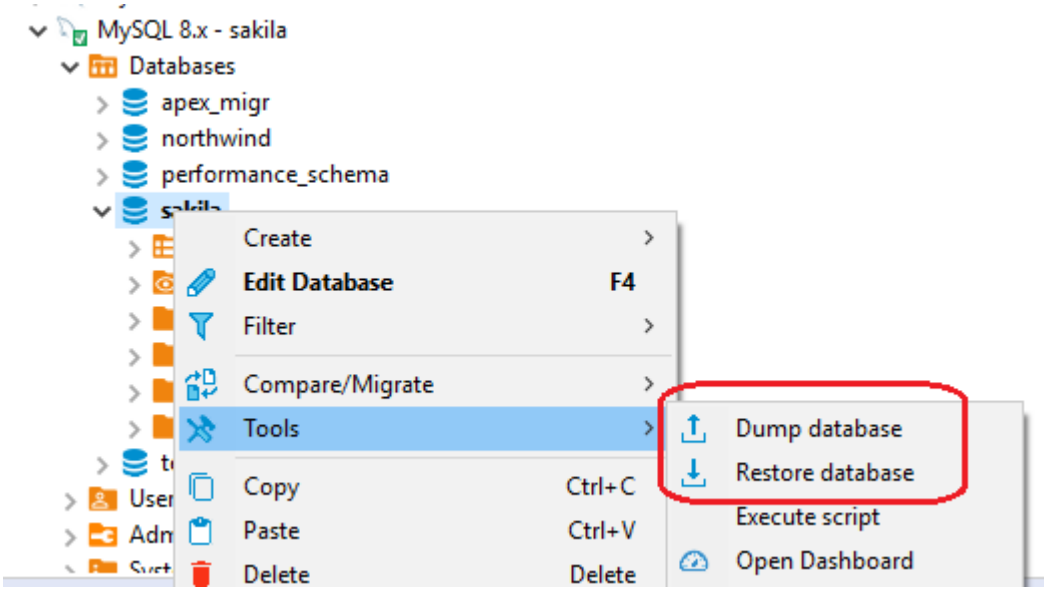

Native client configuration

In order to execute native backup/restore tools you need to configure database native client. Native client is a set of binaries (different for different OSes) which will be executed by DBeaver to process actual backup/restore. Native client configuration can be done in driver editor dialog or directly from backup/restore wizard. Just click on `Client ...` button in the button bar:

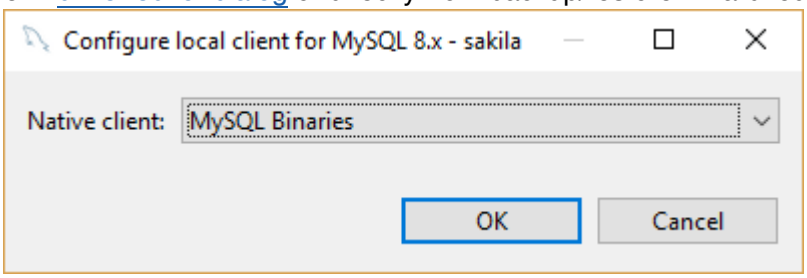

To configure new client location choose `Browse ...` item and ad new client in the following dialog:

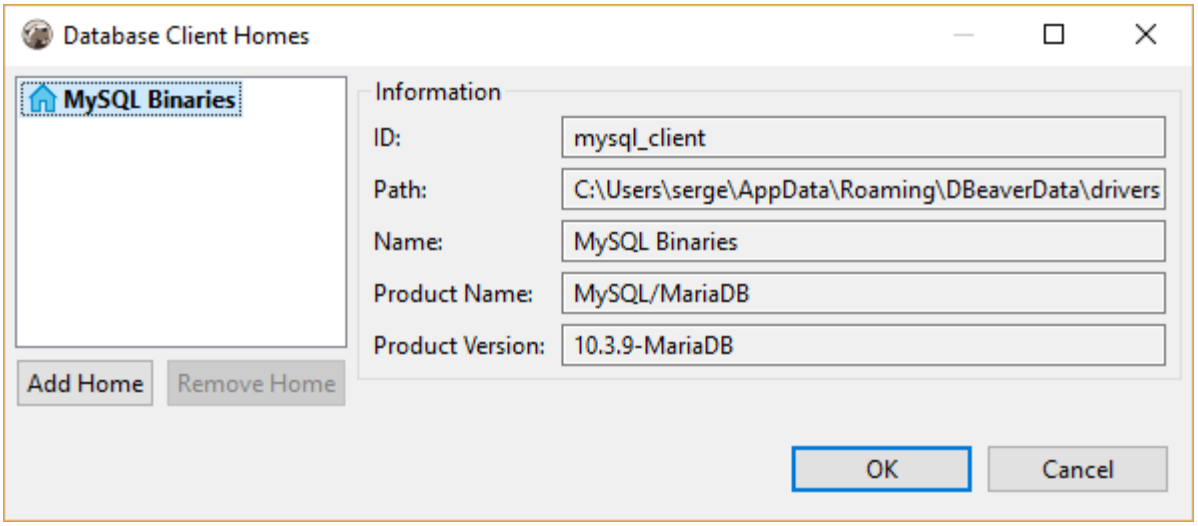

Database dump object selector

You can choose what schemas/tables you want to backup/dump:

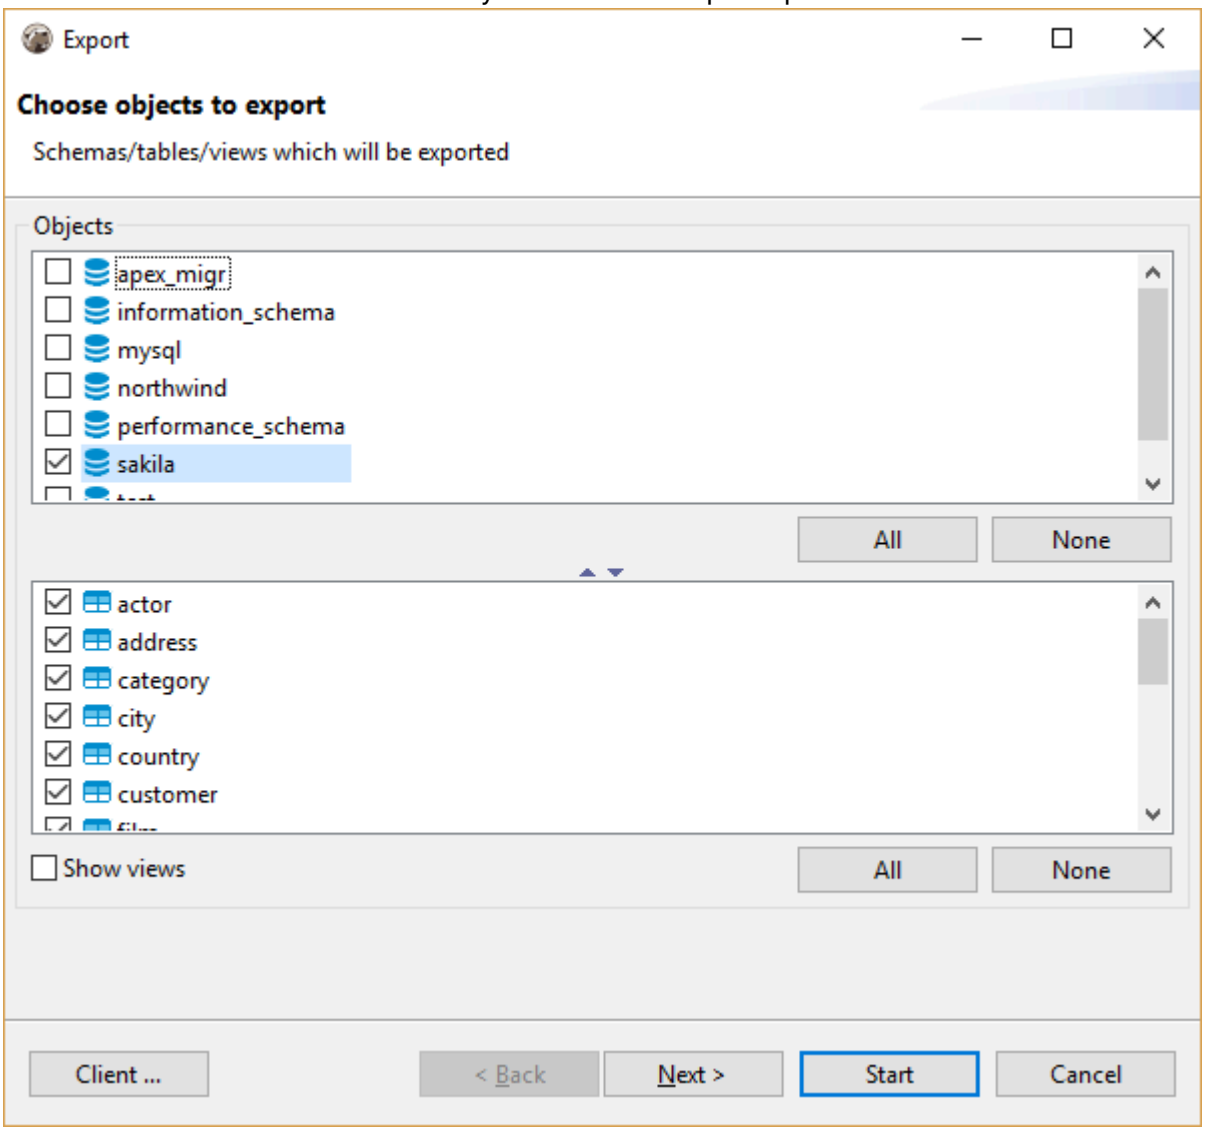

Database native tool configuration

You can pass a set of additional dump/restore parameters to the native tool. Particular set of configuration options depends on a database type.

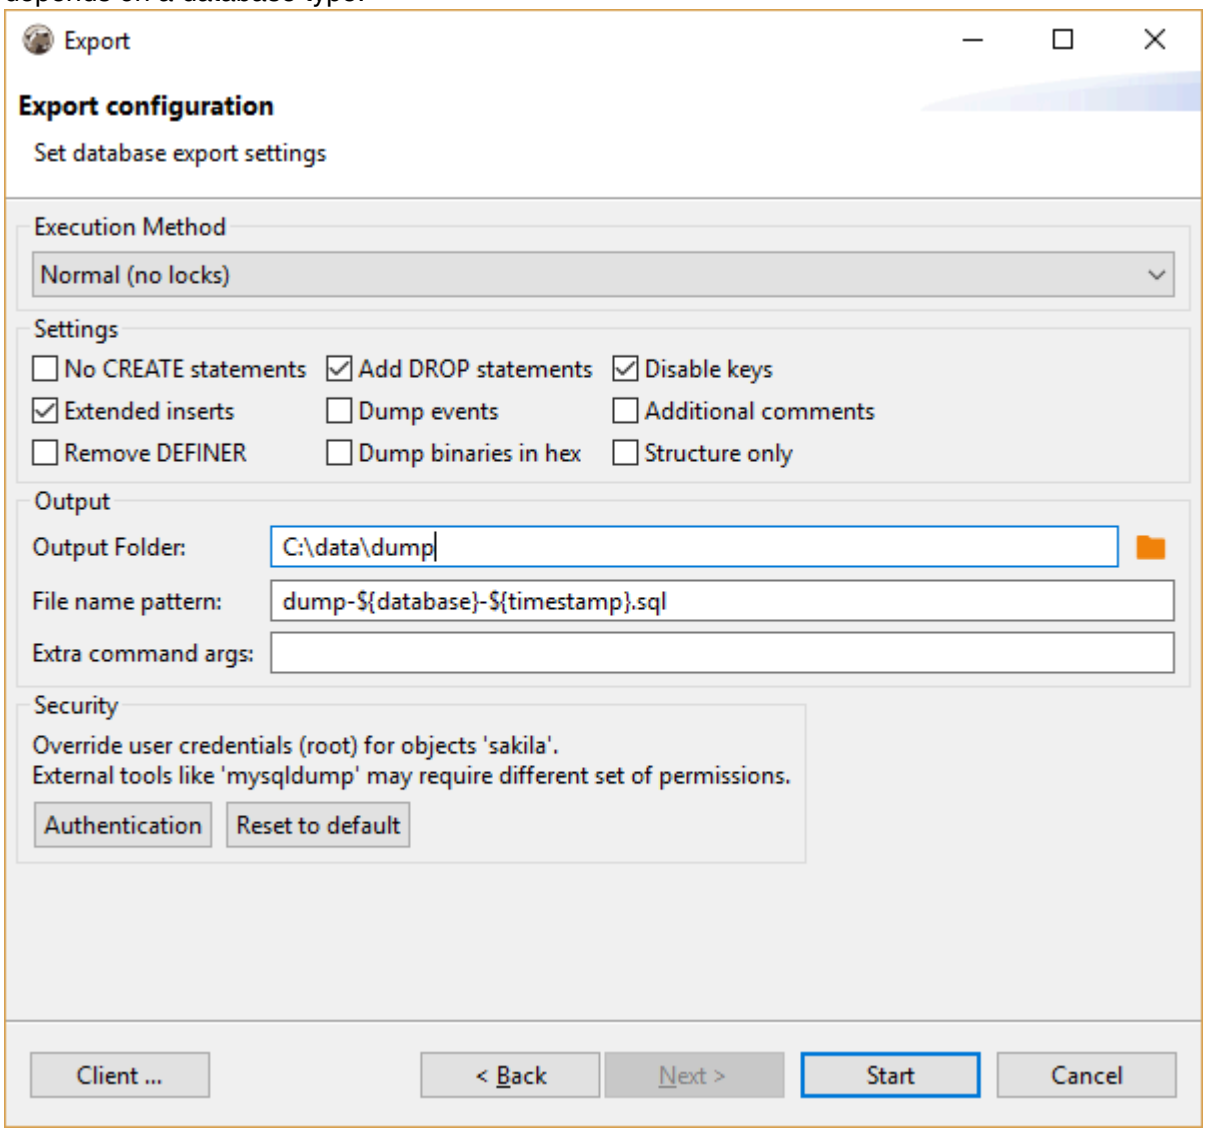

Task management

Creating tasks

Task is a saved configuration of some database tool. It can be started from task management view or from menu by a single click. You can create tasks for frequently used tools. Also tasks can be scheduled for regular execution.

Create task from tool configuration

You can save tool configuration into a task and run your task later with a single click. For example you can start Data Transfer wizard and configure data export from several tables in MySQL database into CSV files:

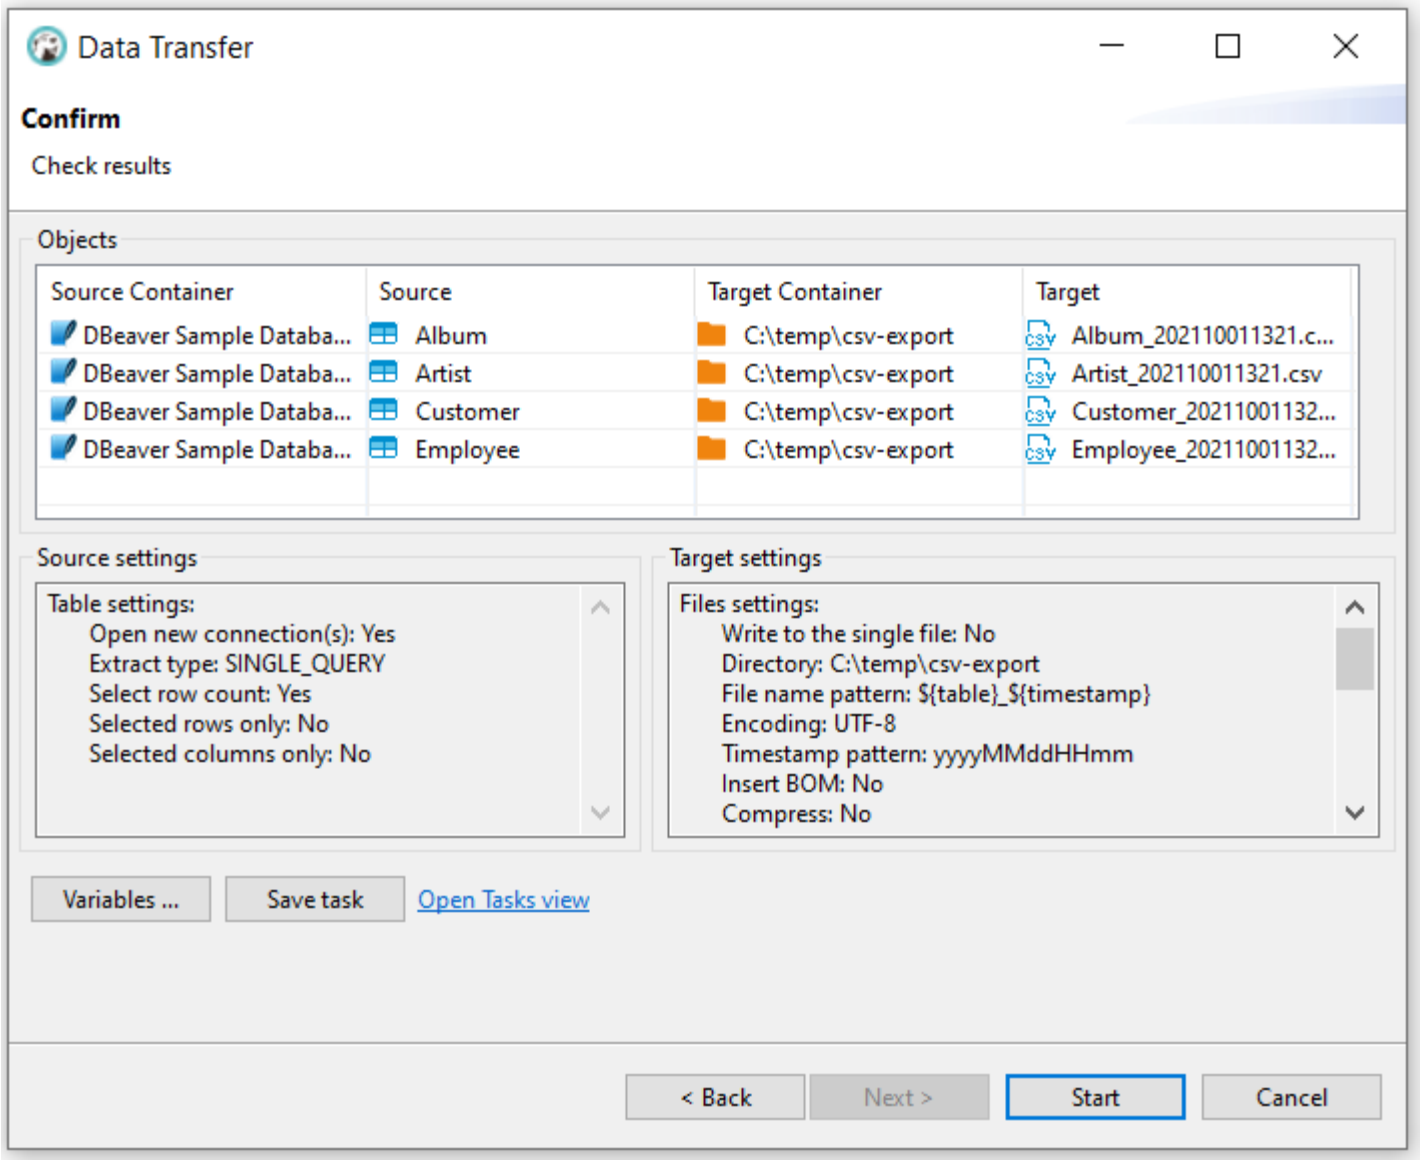

Click on `Save configuration as task` button and fill task properties:

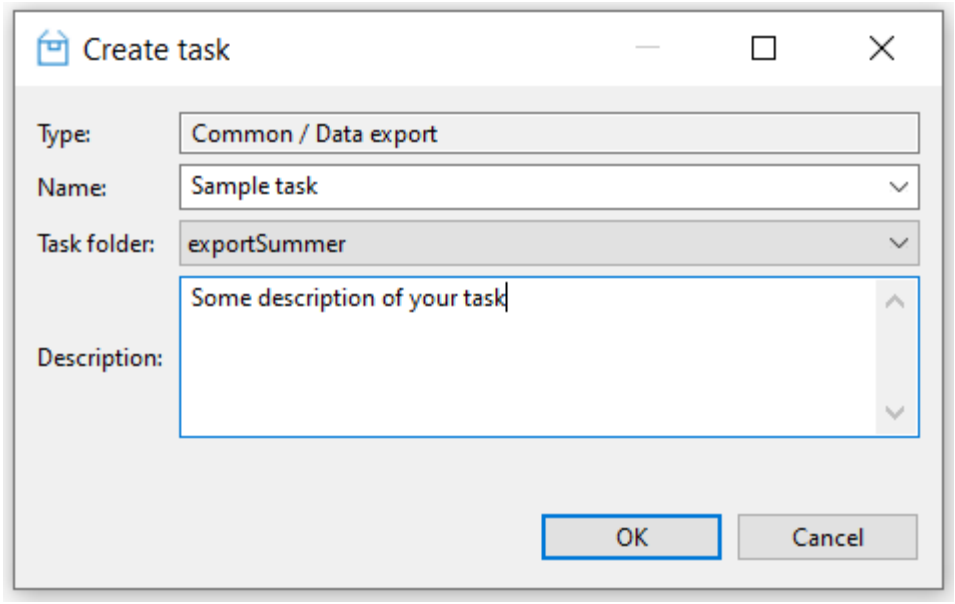

Now click on `Open Tasks view` link to open task list:

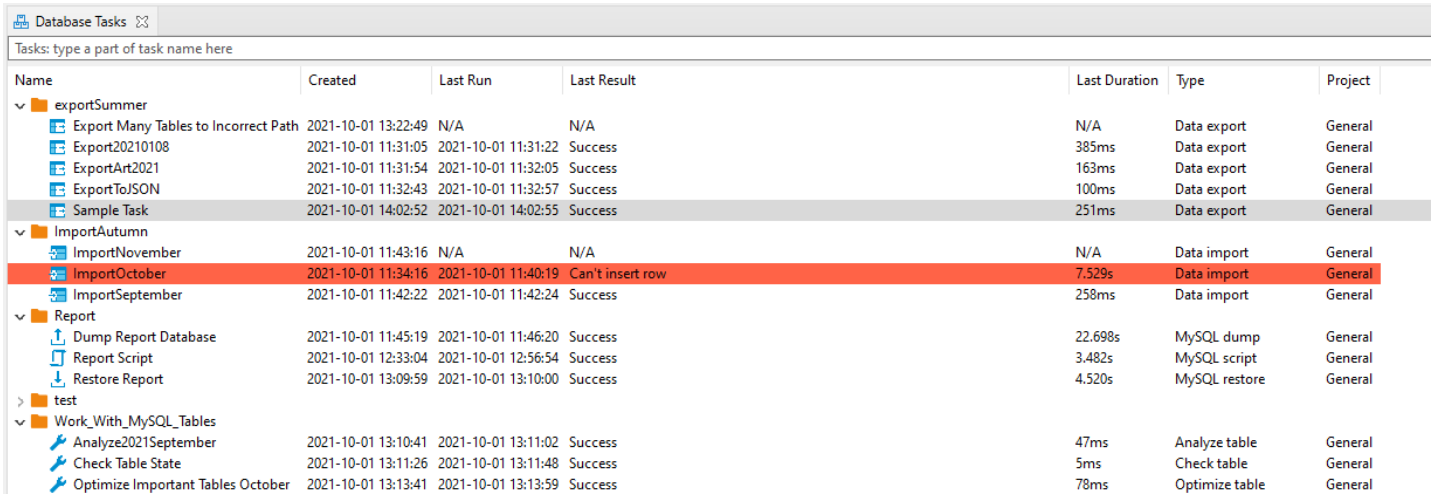

Editing/running tasks

From task view you can add, edit, remove and execute saved tasks. You can use context menu or view tools for that:

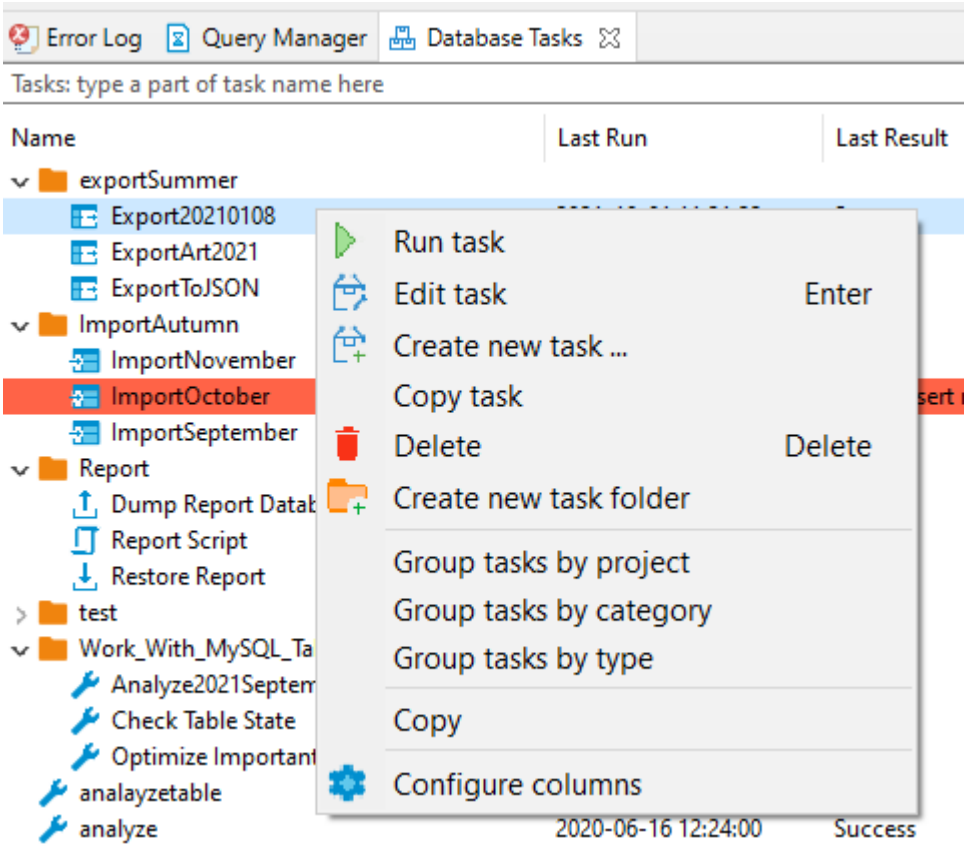

By clicking on `Edit` or by double-clicking on task you can open tasks edit wizard. In this wizard you can change task settings as well as actual tool configuration. You can change set of input objects for data transfer or any export configuration. After changing task settings click on `Update configuration in task` button (it is on the last page of task configuration wizard).

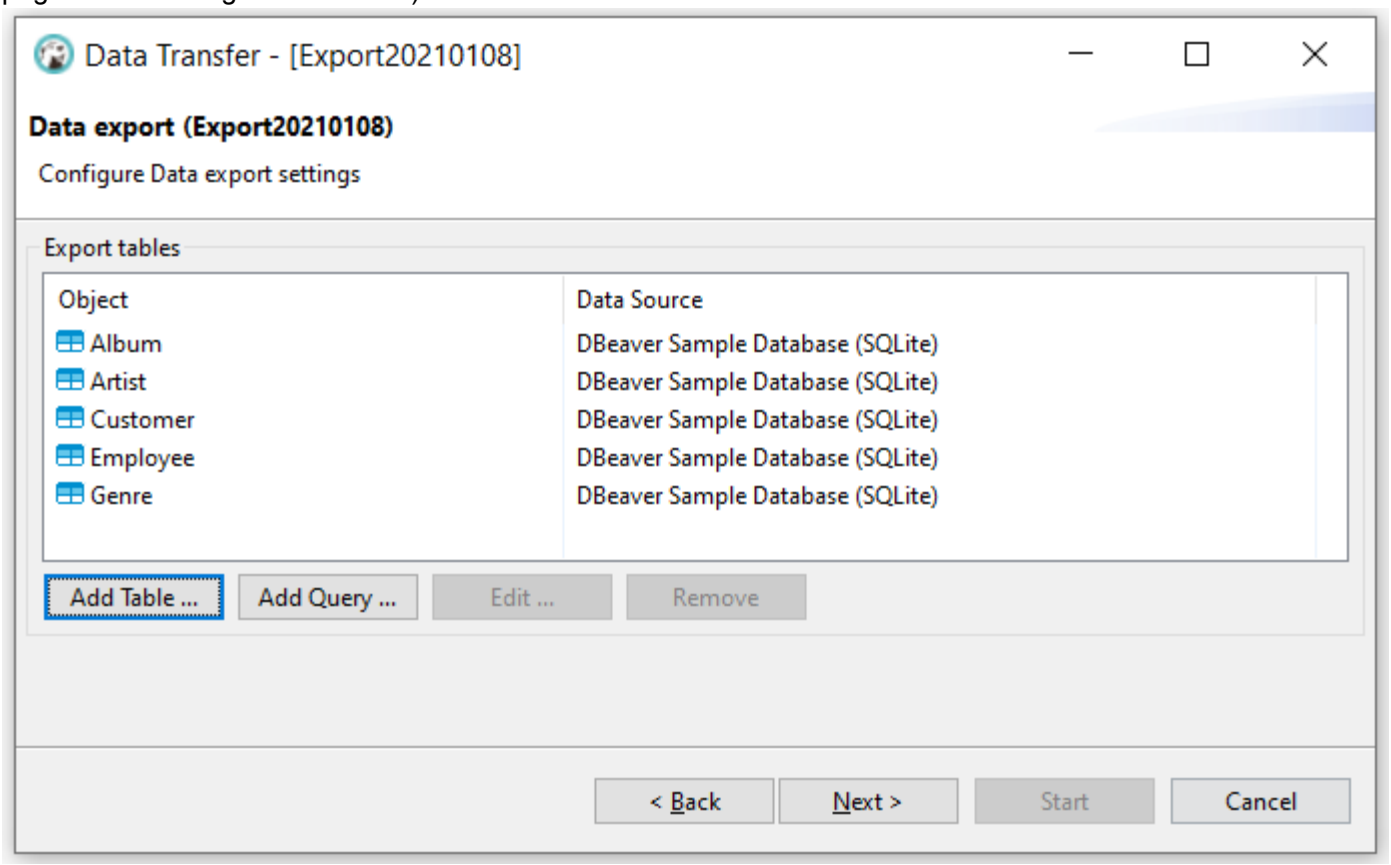

Create task from task management view

You can create task from scratch using tasks view. Open tasks view and `Create new task` button in the view toolbar or in the context menu. In task wizard you can choose task category, task type and name. On the next wizard pages actual tool configuration pages will be shown (they depend on chosen task type).

Scheduling tasks

You can schedule tasks for later/regular execution. See Task Scheduler article.

Task scheduler

Note: This functionality is available only in Enterprise Edition.

DBeaver can schedule task execution for regular executions. DBeaver supports `Windows Task Scheduler` on Windows and `cron` on macOS and GNU/Linux. On top of that as well, you can configure schedulers manually using command line.

Scheduling tasks from the Tasks view

Windows

You can open the tasks view from the main toolbar:

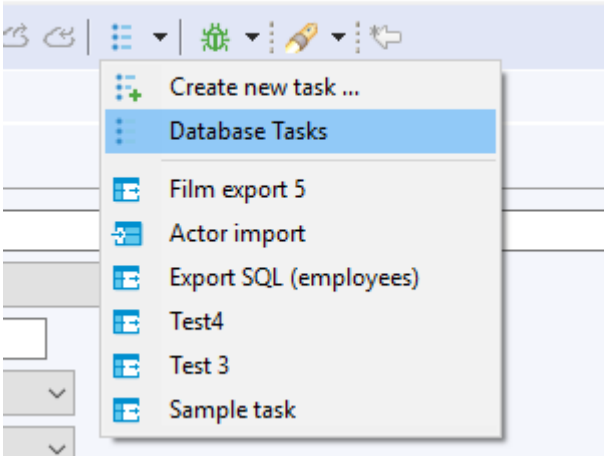

or from the main menu Window. Select a task that you want to schedule in the tasks view and open the context menu:

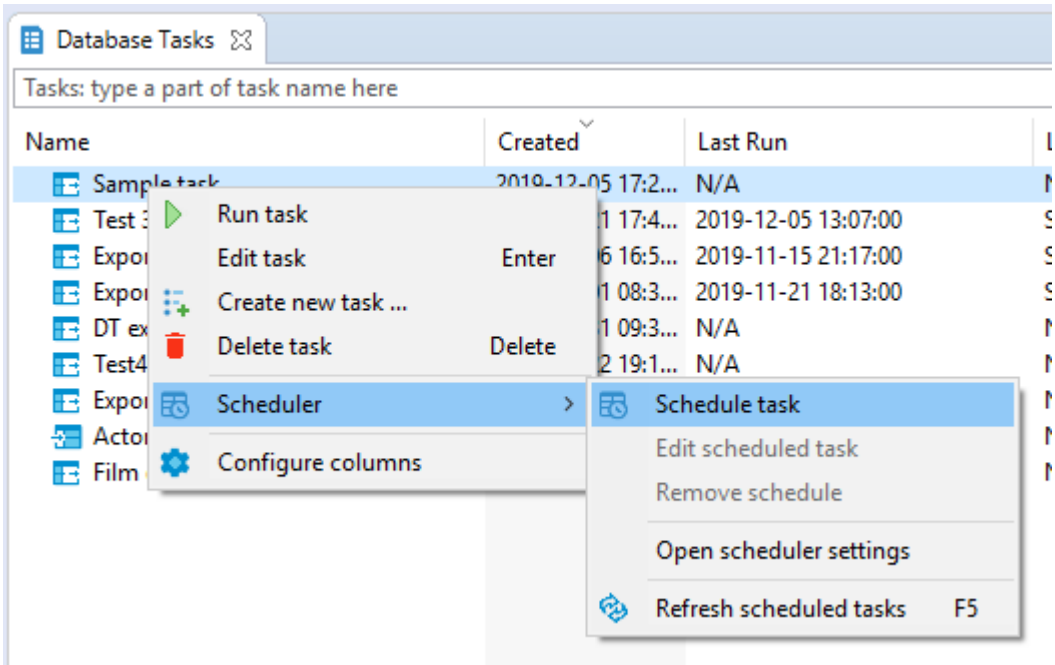

The scheduler configuration dialog will be opened. You can configure task frequency, recurrence period, and start time there:

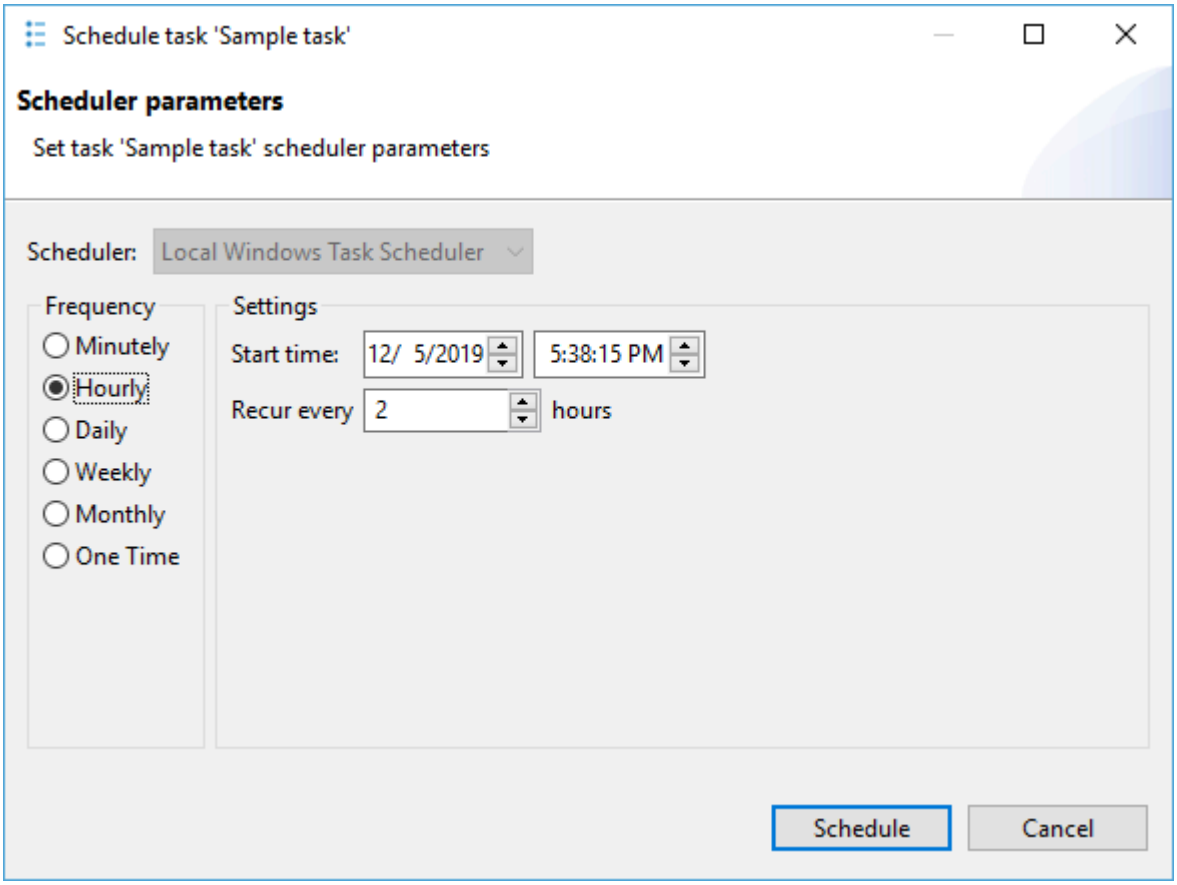

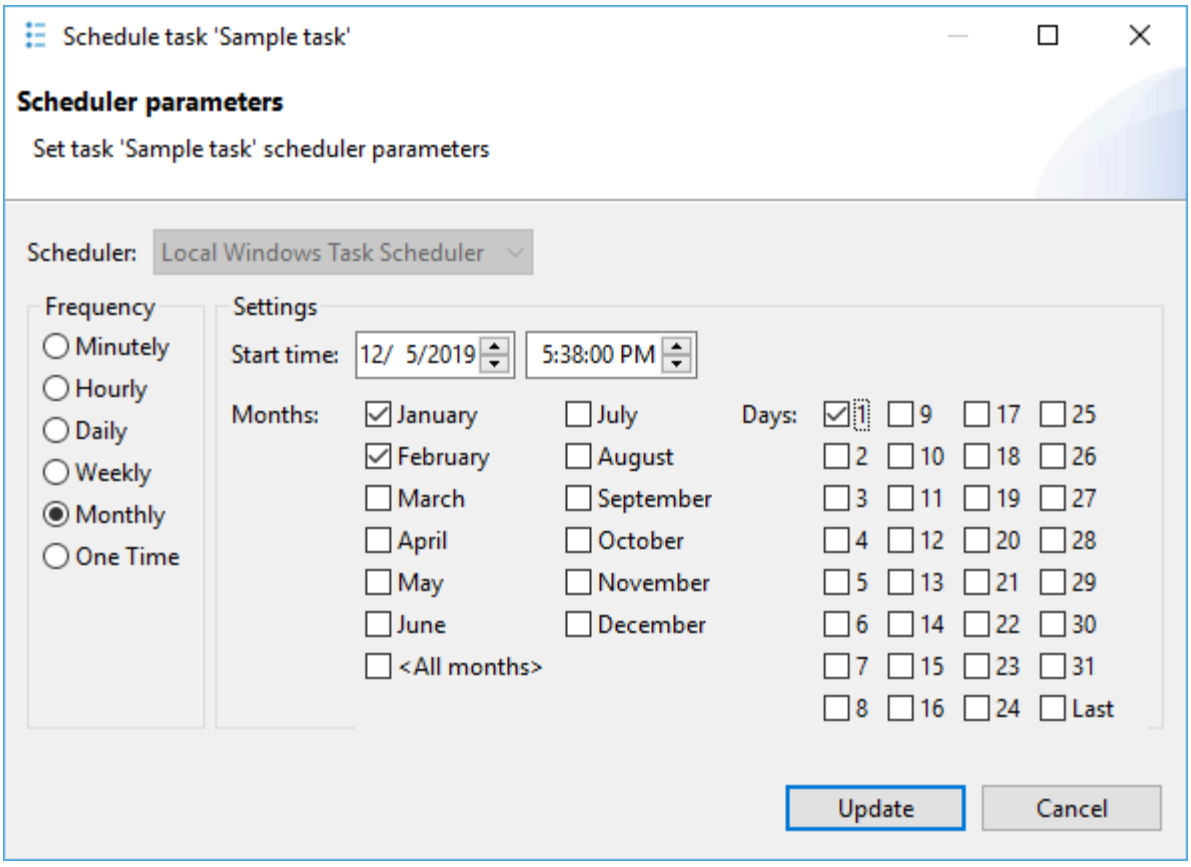

To schedule the task, click on the Schedule button. If everything is configured correctly you will see the confirmation dialog:

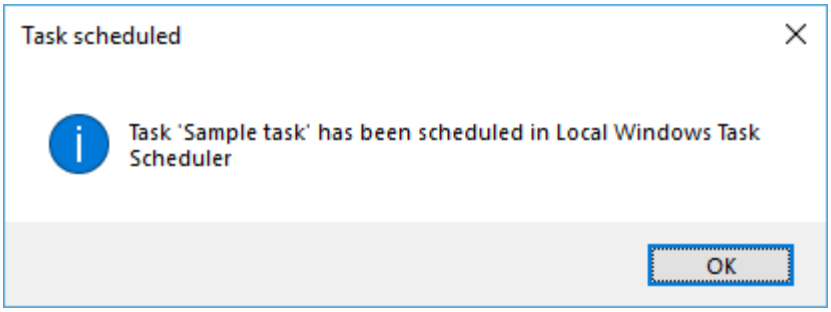

If anything goes wrong you will see an error message dialog. Error details can be viewed in the Error Log view.

You can change scheduler settings at any moment by choosing Edit scheduled task command from the context menu. You can also cancel the schedule by clicking on Remove schedule.

macOS or GNU/Linux

The first thing you need to open the tasks view. There are three ways to do that:

1. Database -> Tasks -> Database Tasks

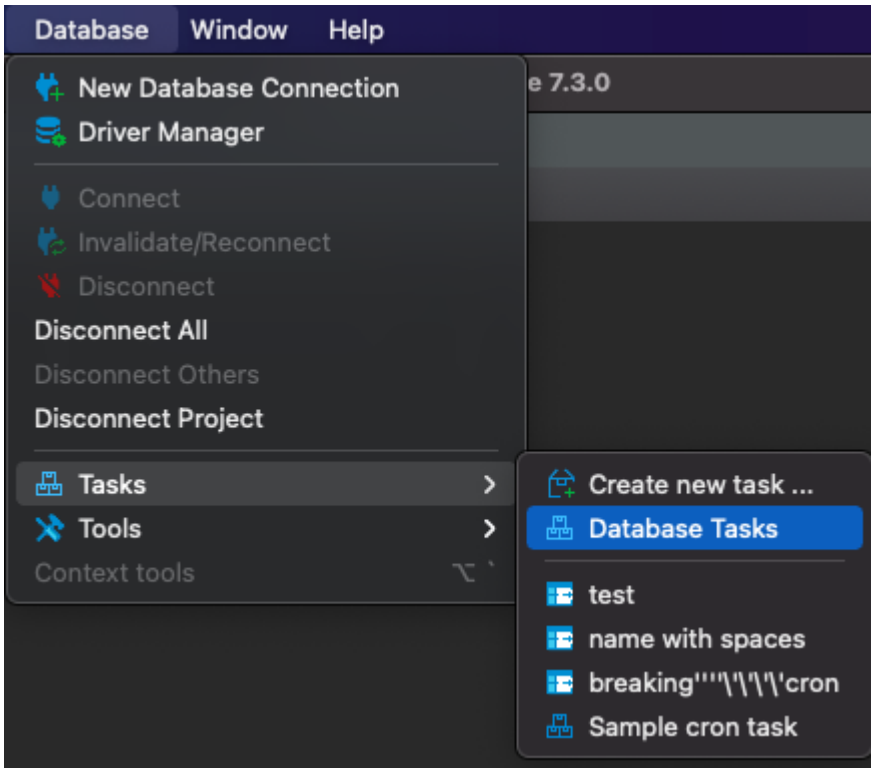

2. Window -> Database Tasks

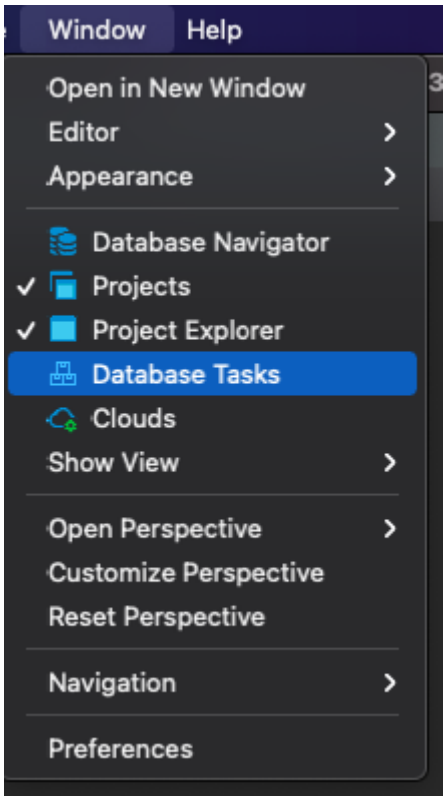

3. Click on 'Show View (Database Tasks)' icon

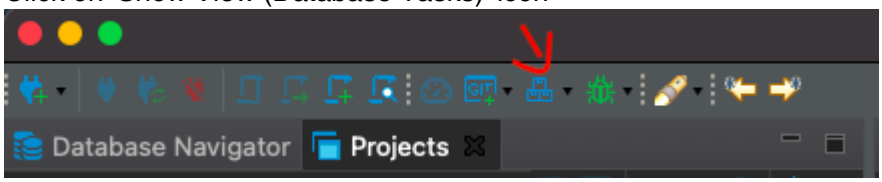

Select a task that you want to schedule in the tasks view. To open the scheduler dialog, either:

1. Open context menu with right-click -> Scheduler -> Schedule task

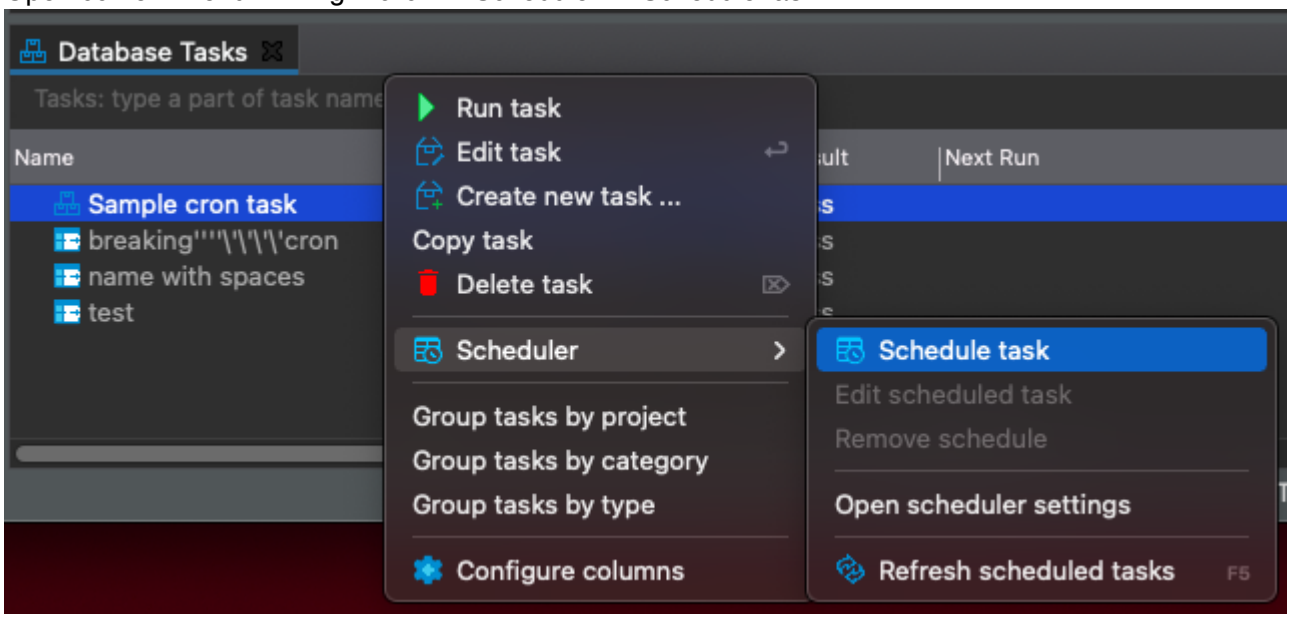

2. or click on 'Schedule task' icon

The scheduler dialog will be opened. It has a lot of similarities with the corresponding dialog on Windows, but unfortunately, there are a bit fewer settings on macOS and GNU/Linux due to the limitations of `cron`. For instance, when configuring an hourly task, you can only choose the minute at which the task executes. In the example below, the task executes at 1:42 PM, 2:42 PM, 3:42 PM, and so on:

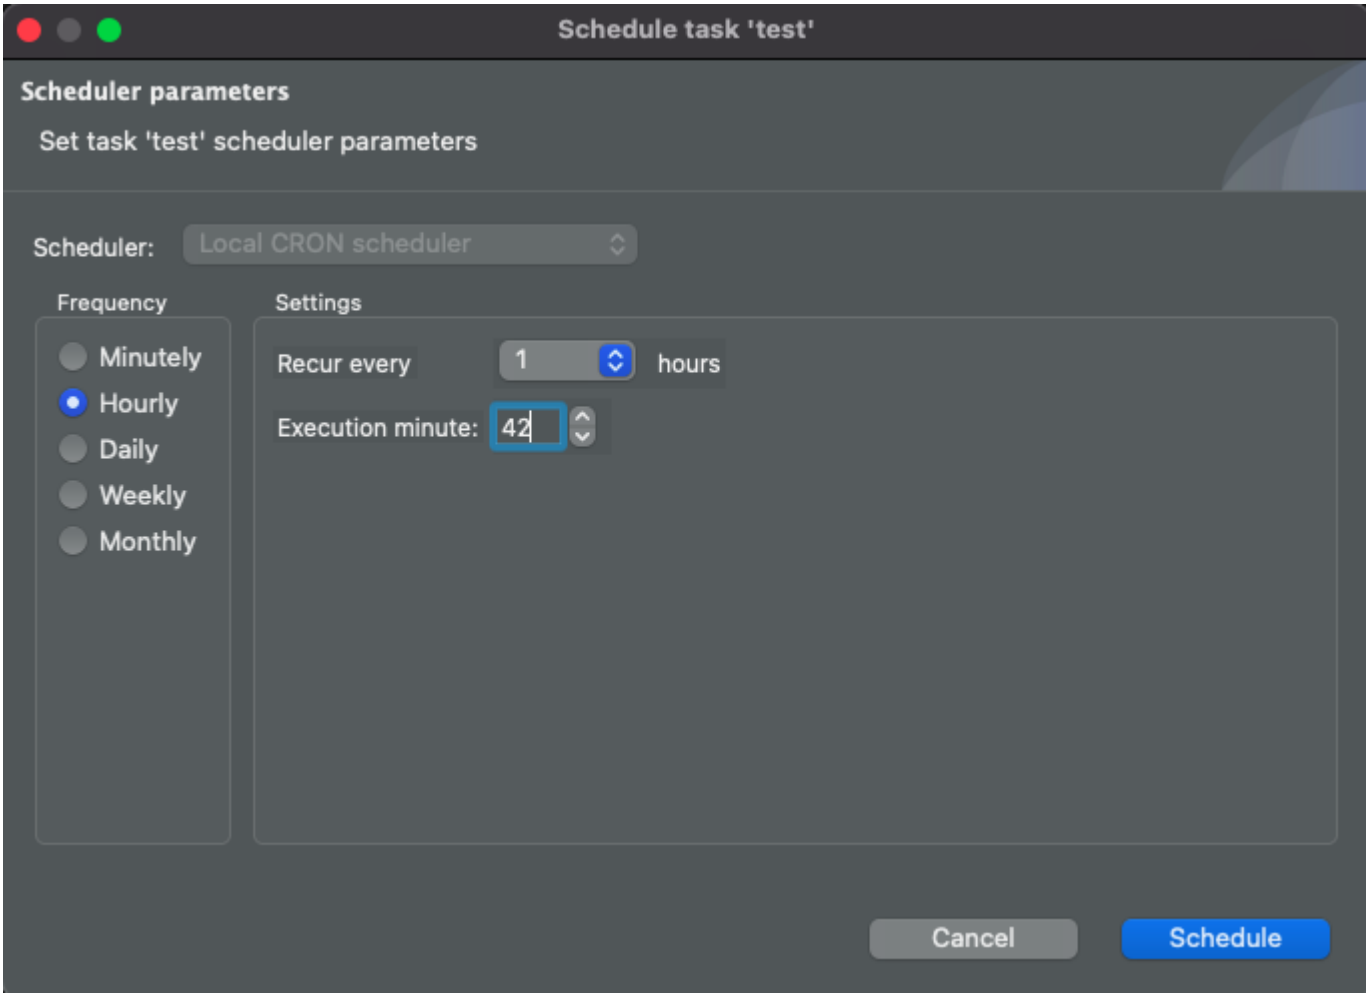

There is also no start date option, and, in case of minutely tasks, no start time either. The scheduler will execute the task at the specified recurrence, but there are no guarantees about when execution starts. It's also worth pointing out that even though you can specify seconds in the start time selector, they will be ignored. Most cron implementations do not support this type of granularity, and we are trying to be compliant with as many cron implementations as possible.

On macOS 10.15 or newer, when scheduling a task for the first time you will be prompted with something like this:

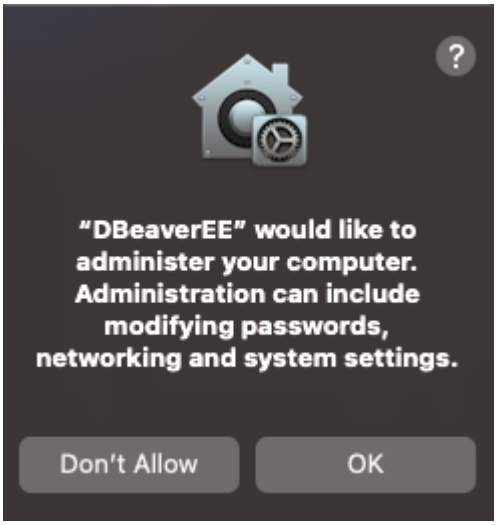

Click 'Yes' to proceed. The reason for that prompt is that `cron` settings (crontabs) are considered to be system settings by macOS, and DBeaver won't be able to change them without permission.

After that, if everything is configured correctly you will see the confirmation message. Just like on Windows, you can change scheduler settings at any moment by choosing 'Edit scheduled task' command from the context menu, or cancel schedule by clicking on 'Remove schedule'.

See schedule details

Windows

You can see and change scheduled task details in the Windows Task Scheduler. Click on the Open scheduler settings command in the task view context menu:

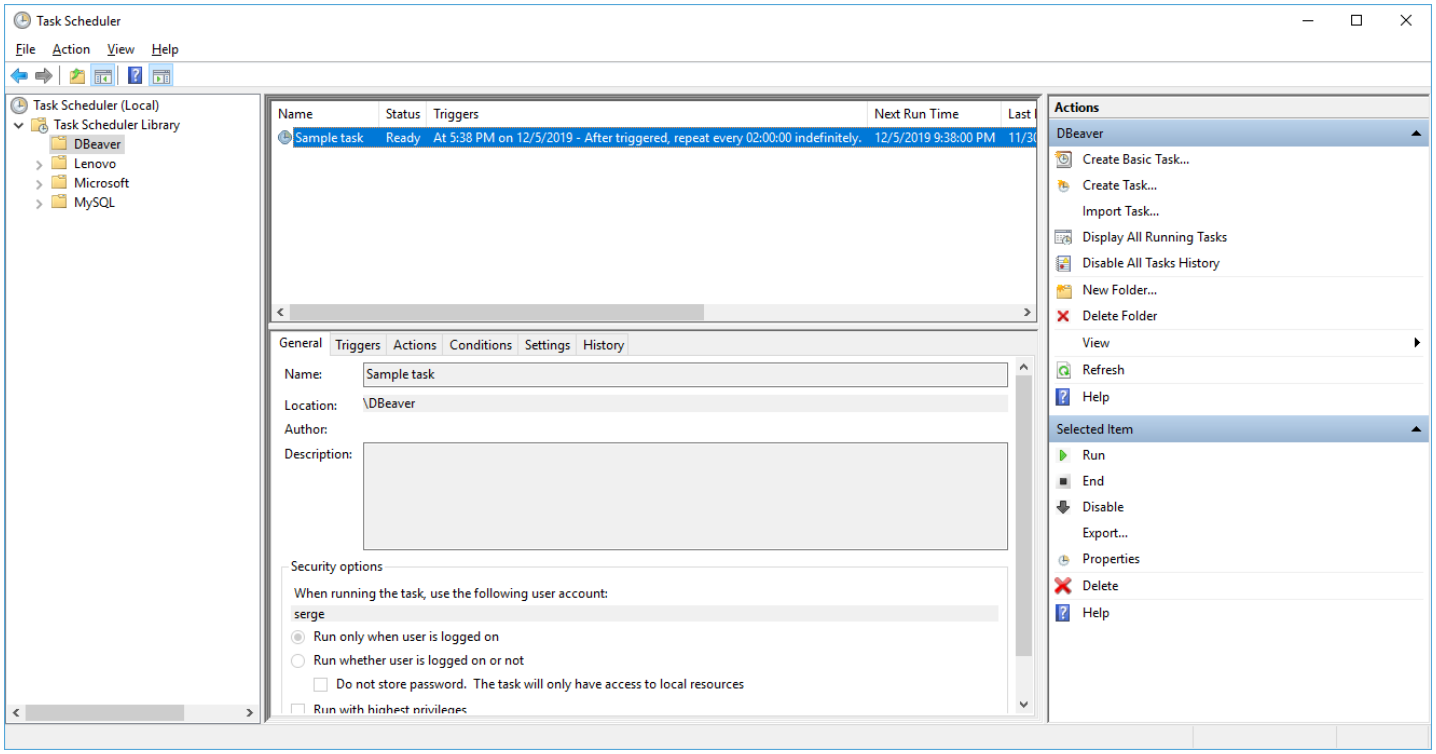

All DBeaver tasks are located in a folder DBeaver.

macOS or GNU/Linux

You can take a look at the crontab DBeaver uses to schedule tasks in `cron` by clicking the 'Open scheduler settings' command in the task view context menu. You can also do it in the terminal by using the command `crontab -l`. Although you can also edit the crontab by using `crontab -e`, we strongly do not recommend it.

Monitoring for task execution (any OS)

You can look through the task execution logs on the right side of the tasks view. By double-clicking on a task run item you can see the full log with all details, errors, and warnings:

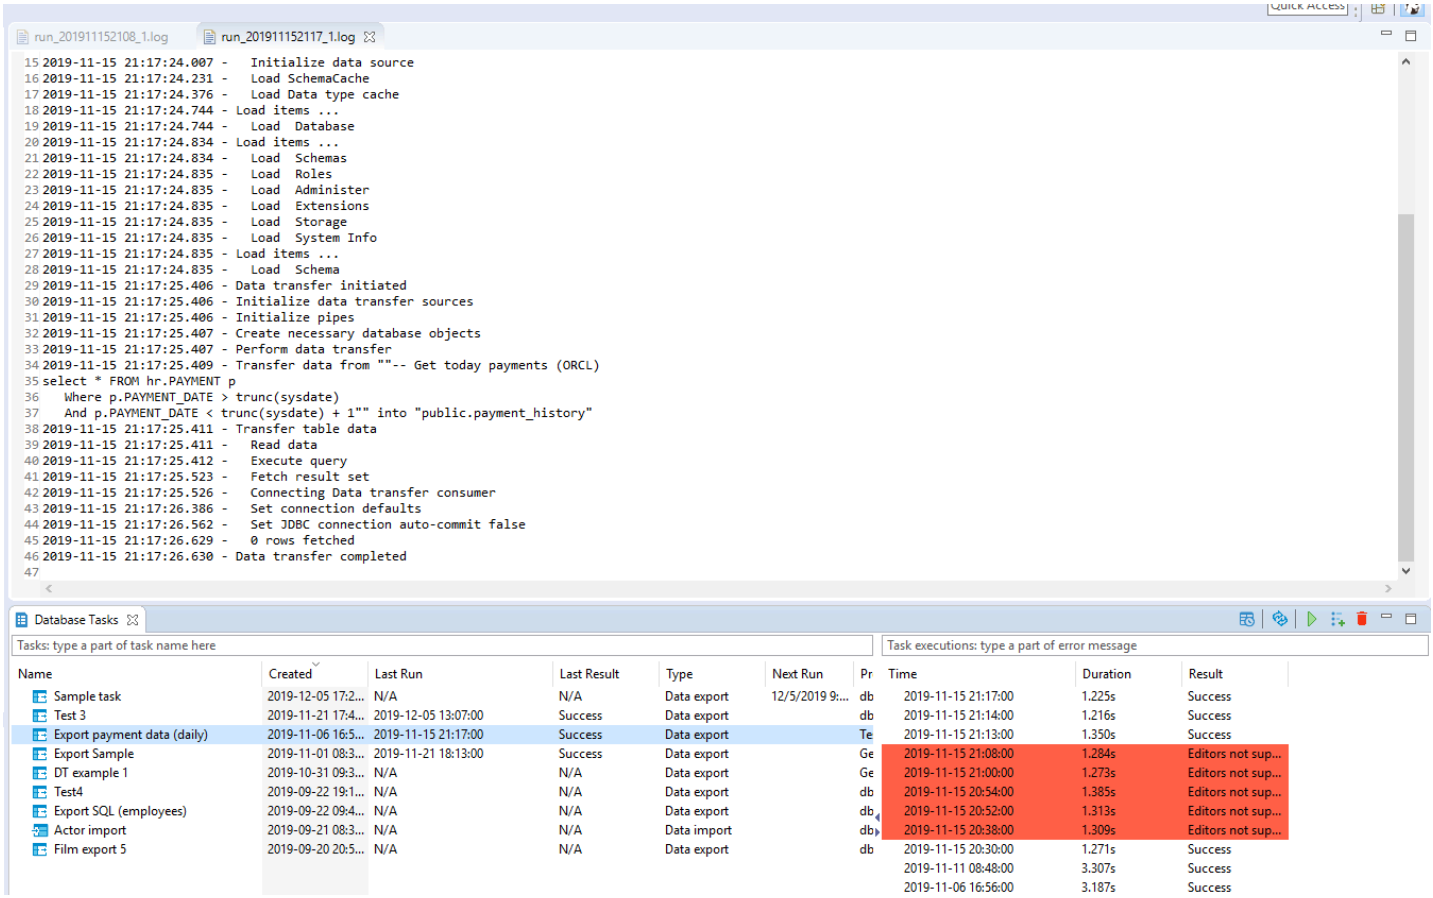

DBeaver keeps the task run logs in the workspace directory, subfolder .metadata/task-stats.

Running tasks from the command line (any OS)

The task scheduler uses the DBeaver command line interface to perform task execution. Command line parameter `-runTask TASK_ID` launches saved task execution (immediately). TASK_ID has form `@projectName:taskName`. You can omit the project name part if you have only one project in your workspace. On Windows, you can use `dbeaver-cli` executable to run tasks. Draw your attention that if you use `dbeaver` executable (by any reason) you will need to add command line parameter `-nosplash` to avoid splash screen appearance.

Troubleshooting

Windows Task Scheduler: incorrect date format

If you encounter an error on Windows which looks like this:
`ERROR: Invalid Start Date (Date should be in %some_format% format).`,

do the following:

1. Open file `dbeaver.ini` in the directory with your DBeaver installation
2. Place the line `-Ddbeaver.scheduler.windows.dateFormat=%some_format%` (where %some_format% is a format from the error message) below the `-vmargs` line.

This flag is available starting from the 7.3.4 EA version of DBeaverEnterprise and might be removed in the future.

Composite tasks

Note: This functionality is available only in Enterprise Edition.

As the name suggests, the *composite task* is a type of task that consists of other tasks. Just like the other type of tasks, the composite tasks can be scheduled via Task Scheduler. Let's take a look at what they can offer.

Creating a *composite task*

The first thing we need to open *Create a task* dialog. You can do it in multiple ways:

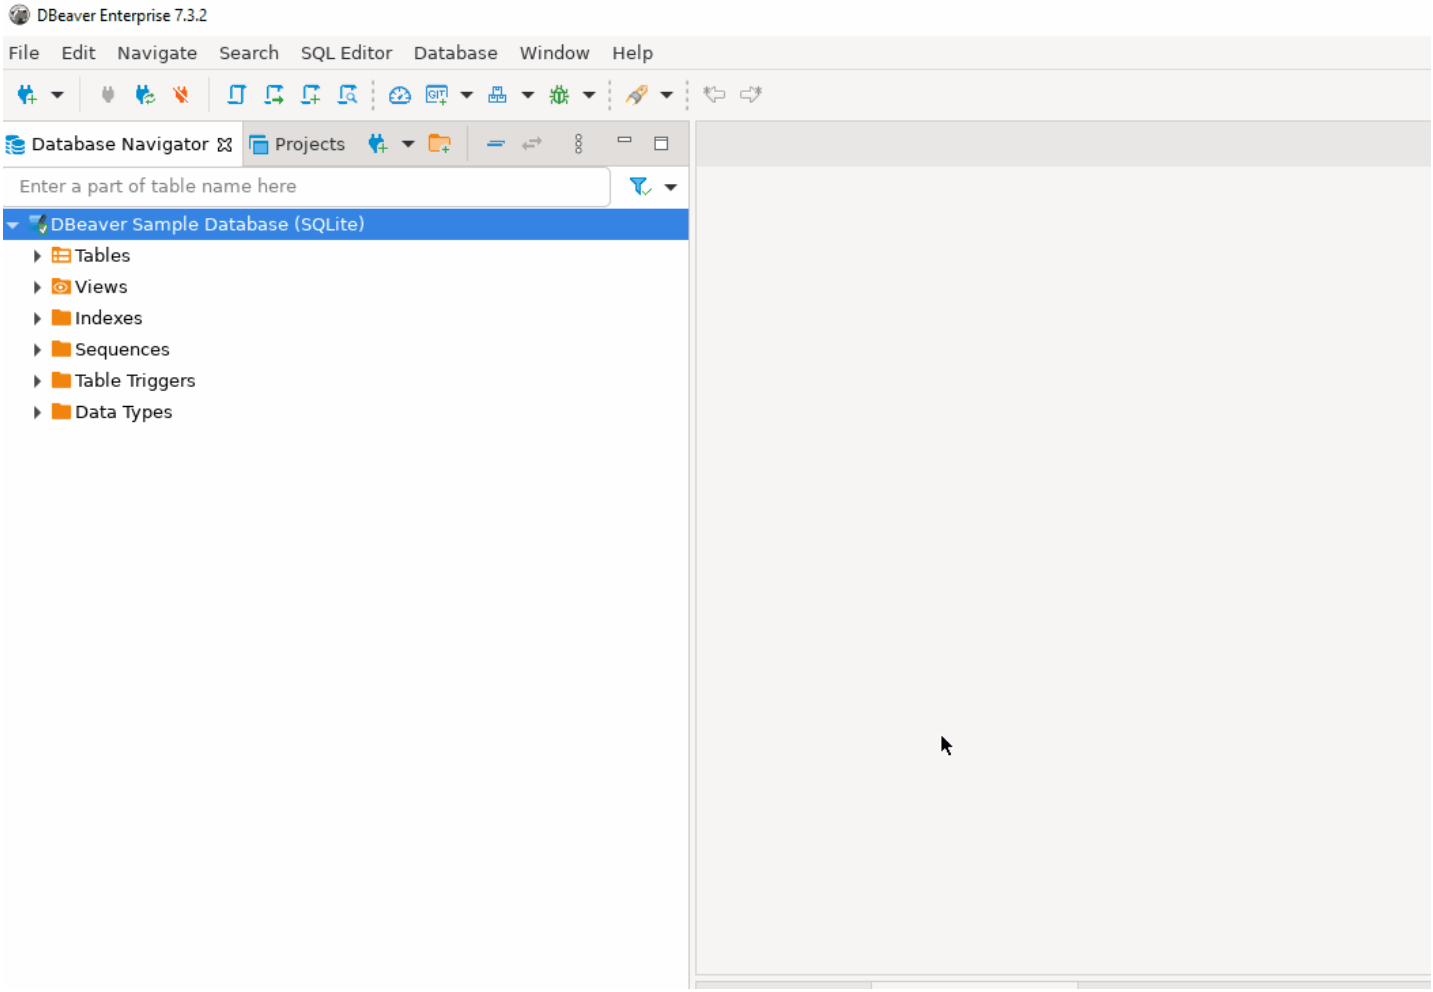
DBeaver Enterprise 7.3.2
File Edit Navigate Search SQL Editor Database Window Help
Database Navigator
Projects
Enter a part of table name here
DBeaver Sample Database (SQLite)
Tables
Views
Indexes
Sequences
Table Triggers
Data Types

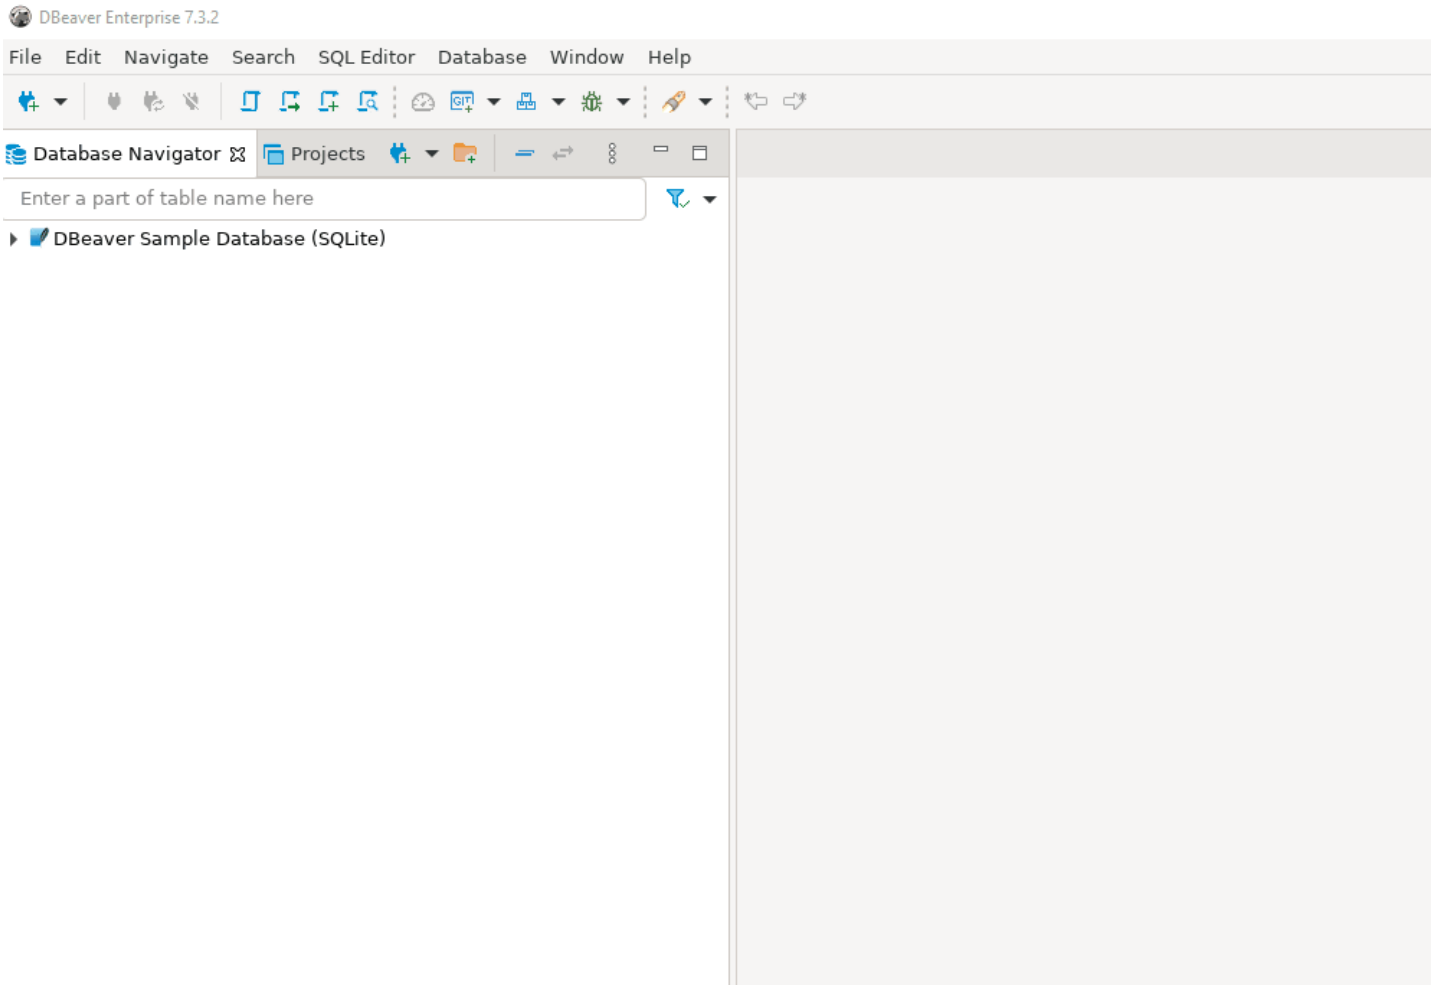
DBeaver Enterprise 7.3.2
File Edit Navigate Search SQL Editor Database Window Help
Database Navigator
Projects
Enter a part of table name here
DBeaver Sample Database (SQLite)

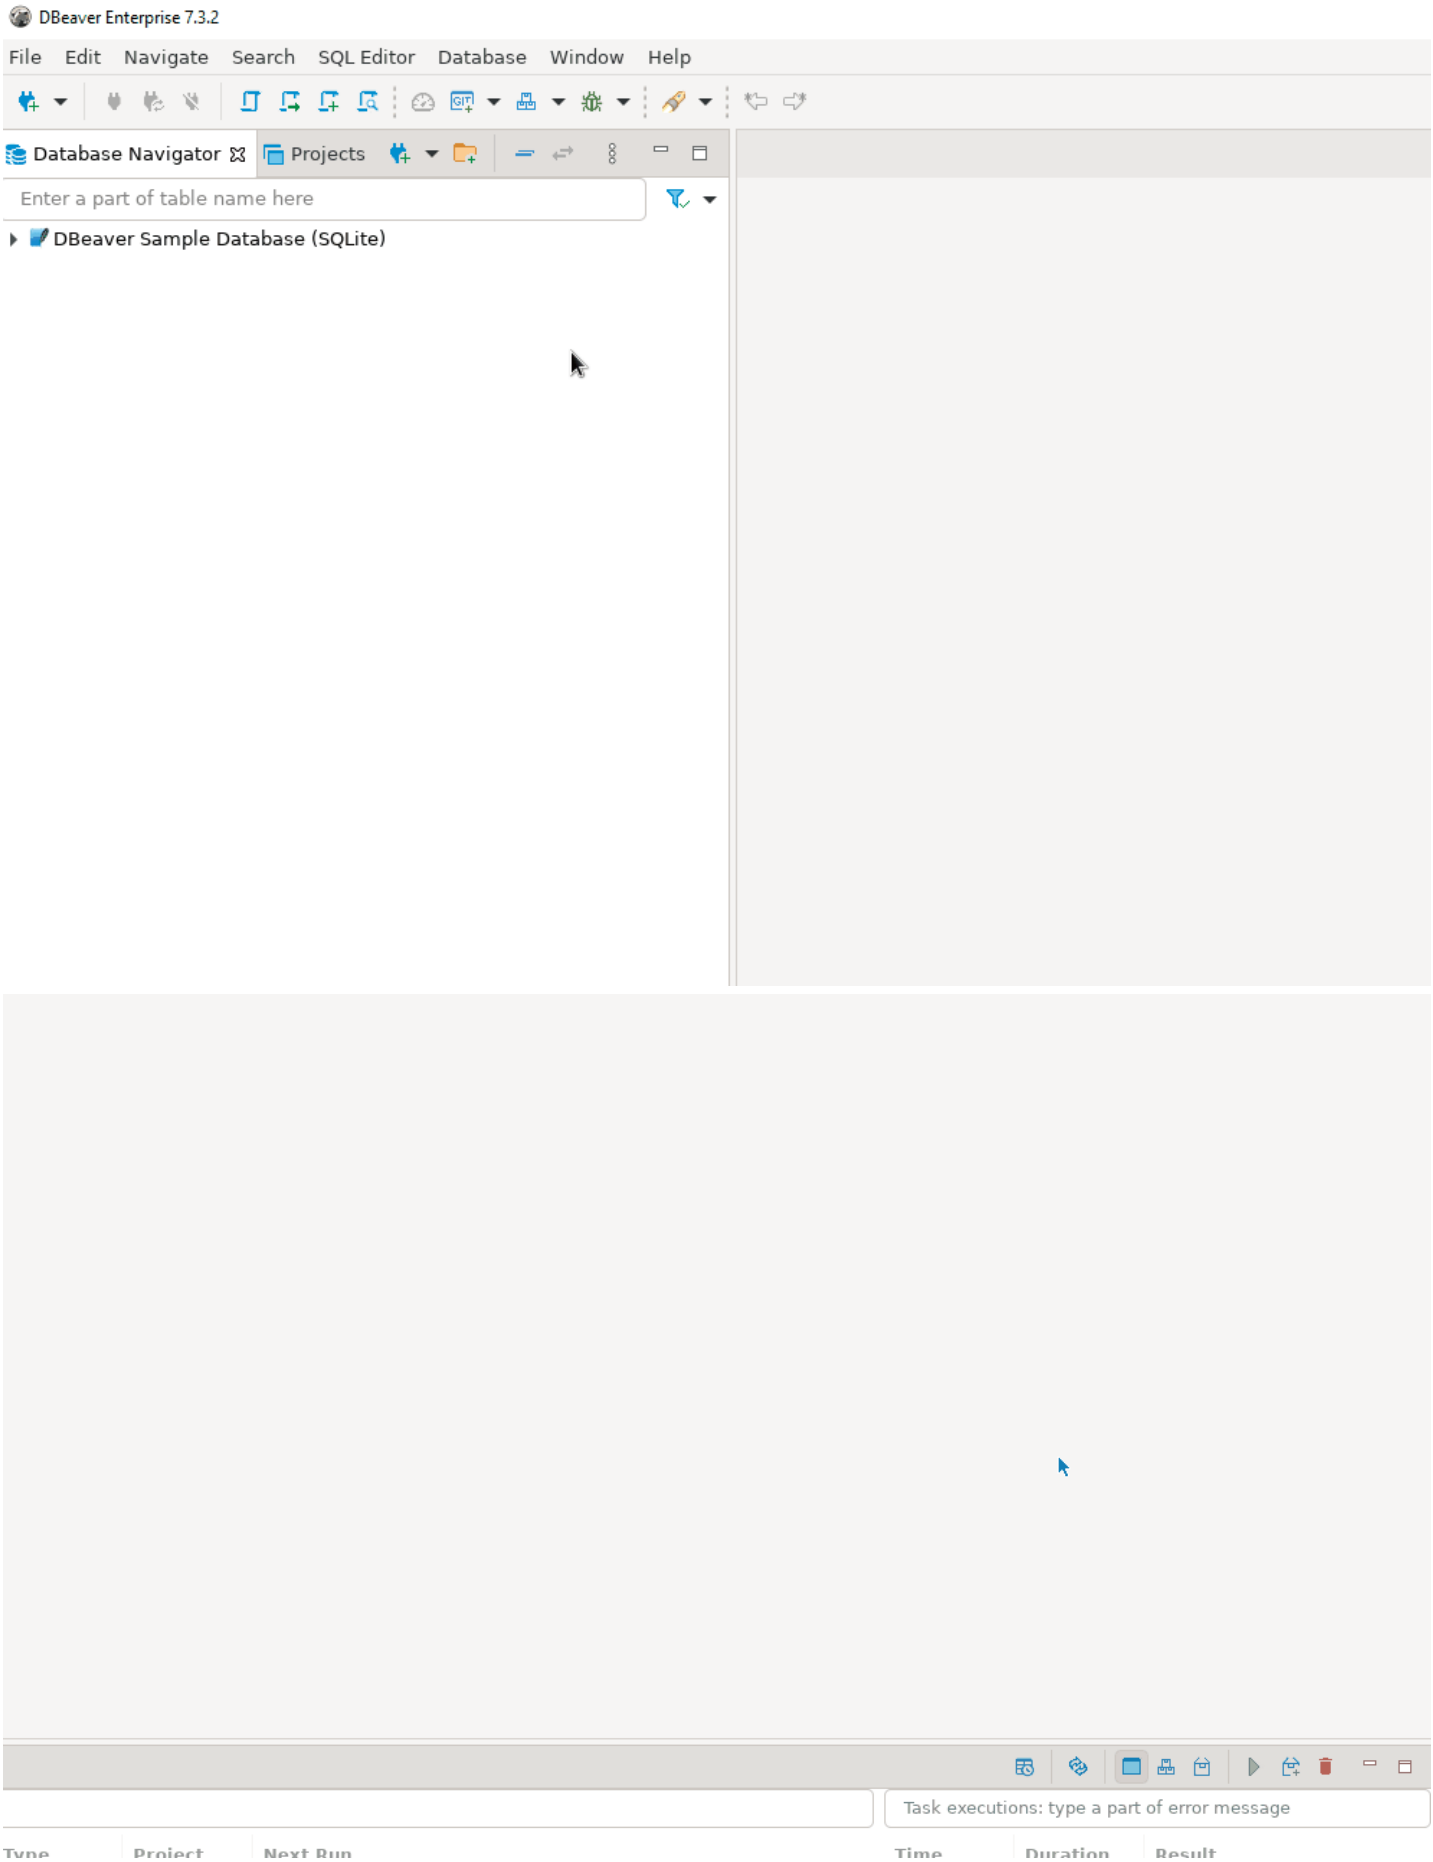
DBeaver Enterprise 7.3.2
File Edit Navigate Search SQL Editor Database Window Help
Database Navigator
Projects
Enter a part of table name here
DBeaver Sample Database (SQLite)
Task executions: type a part of error message
Type
Project
Next Run
Time
Duration
Result

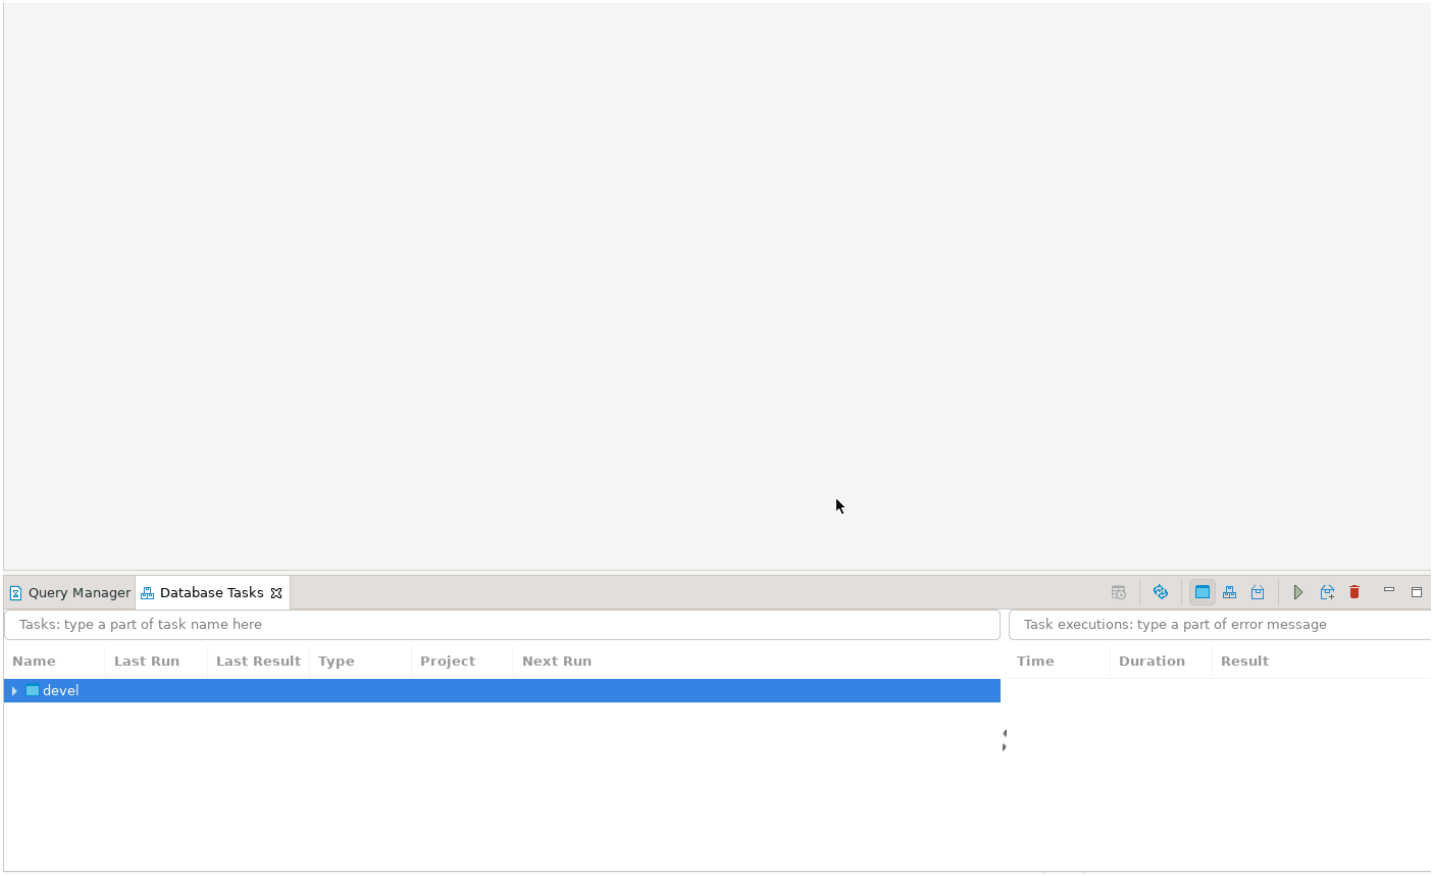

Choose *Composite task*, enter the task name, description (optional), and hit *Next*.

You will be presented with the following dialog:

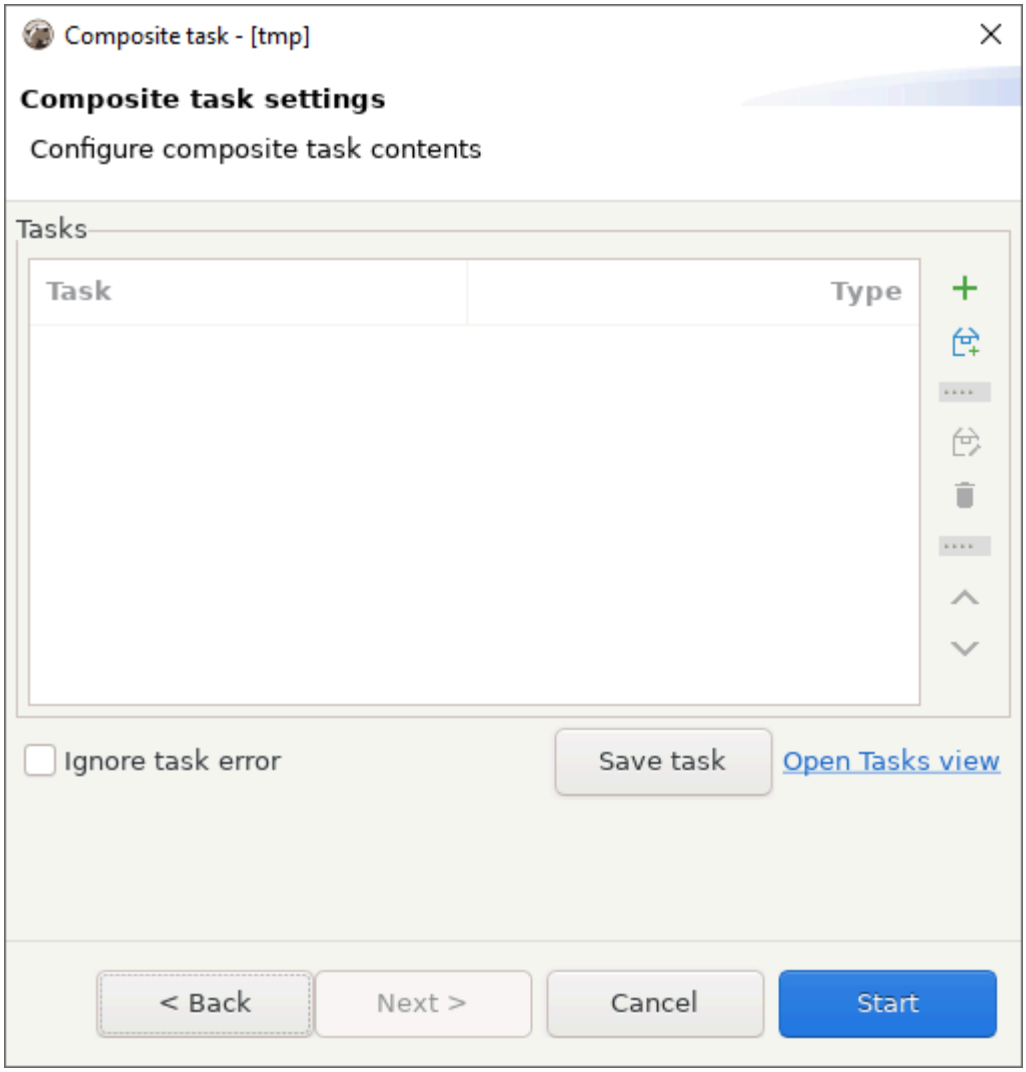

Setting up a *composite task*

When creating a composite task you need to specify which tasks the composite task consists of.

This can be done:

1. By adding an existed task:

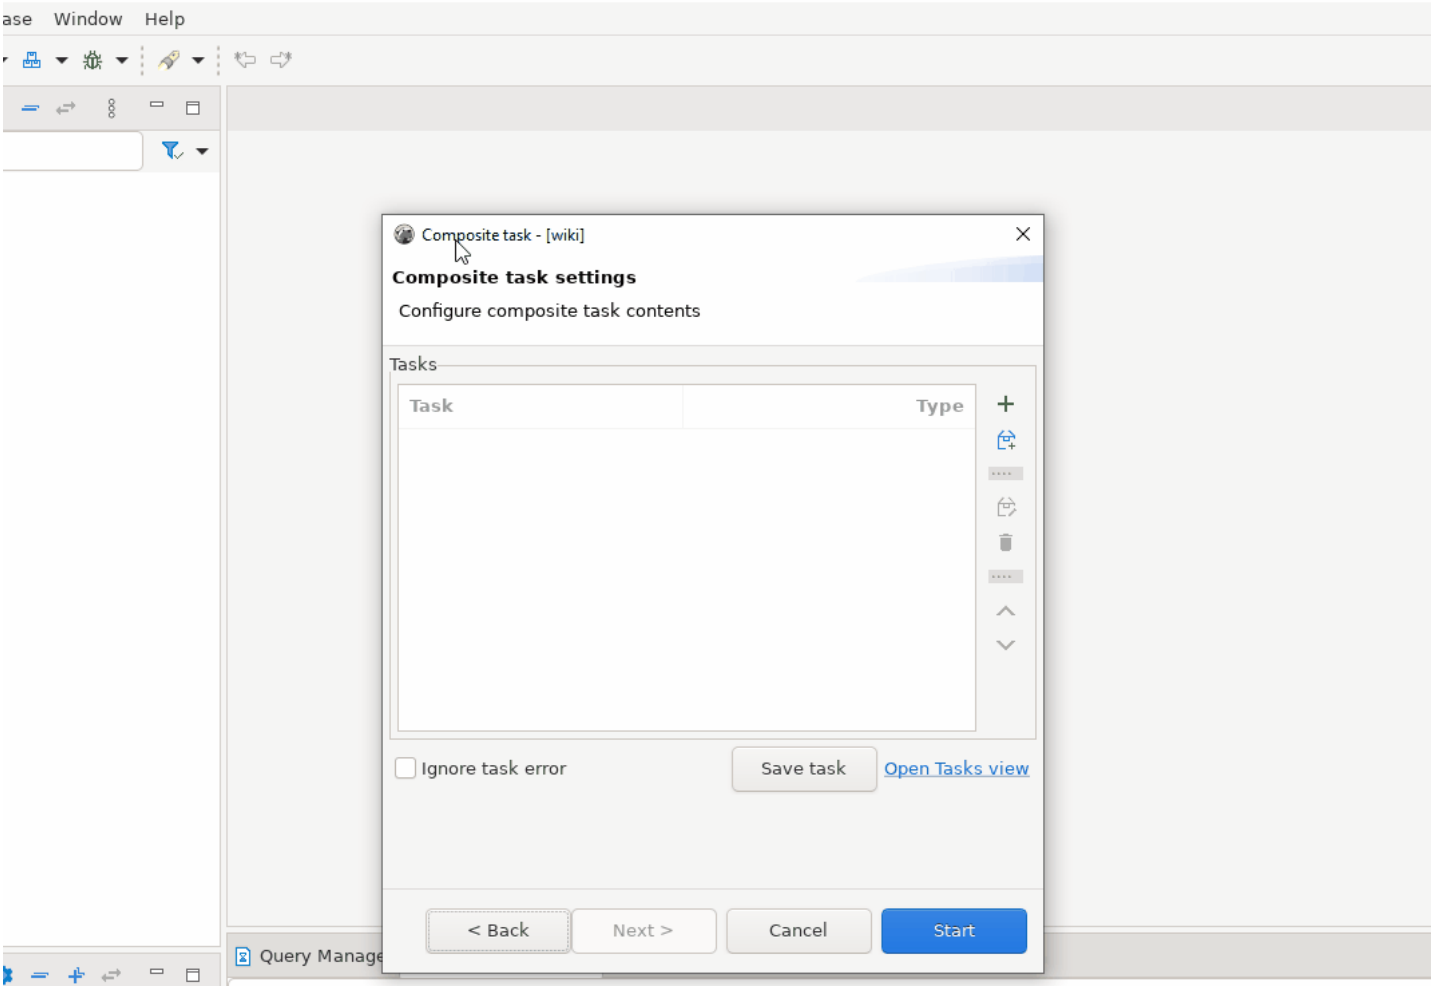

2. By creating a new task and adding it simultaneously:

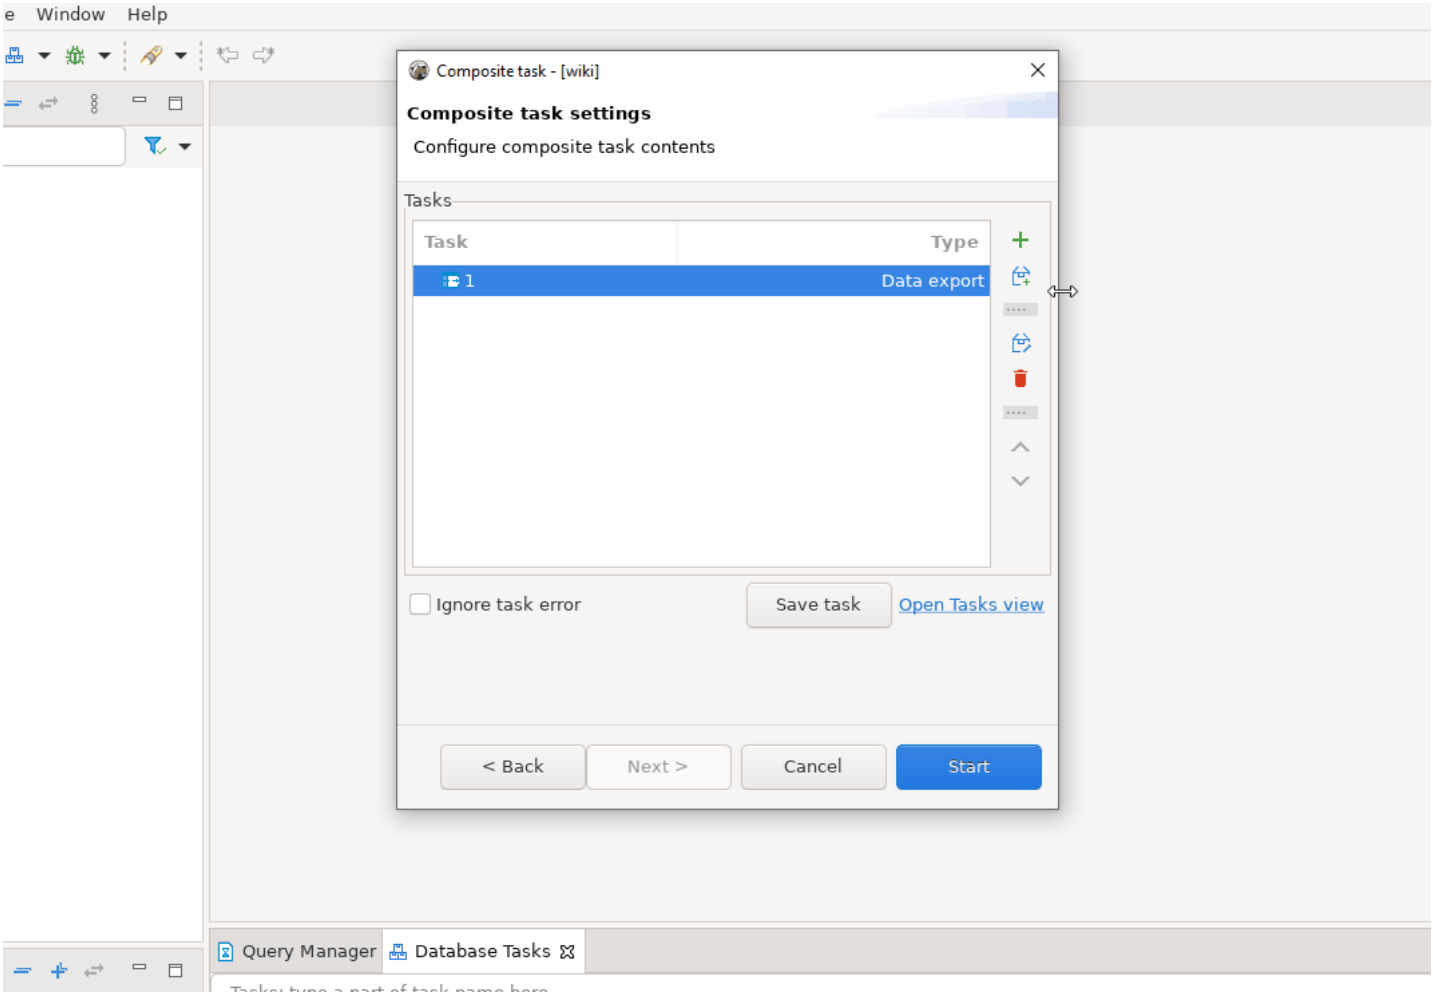

3. By drag-and-dropping a task from the *Database tasks panel*:

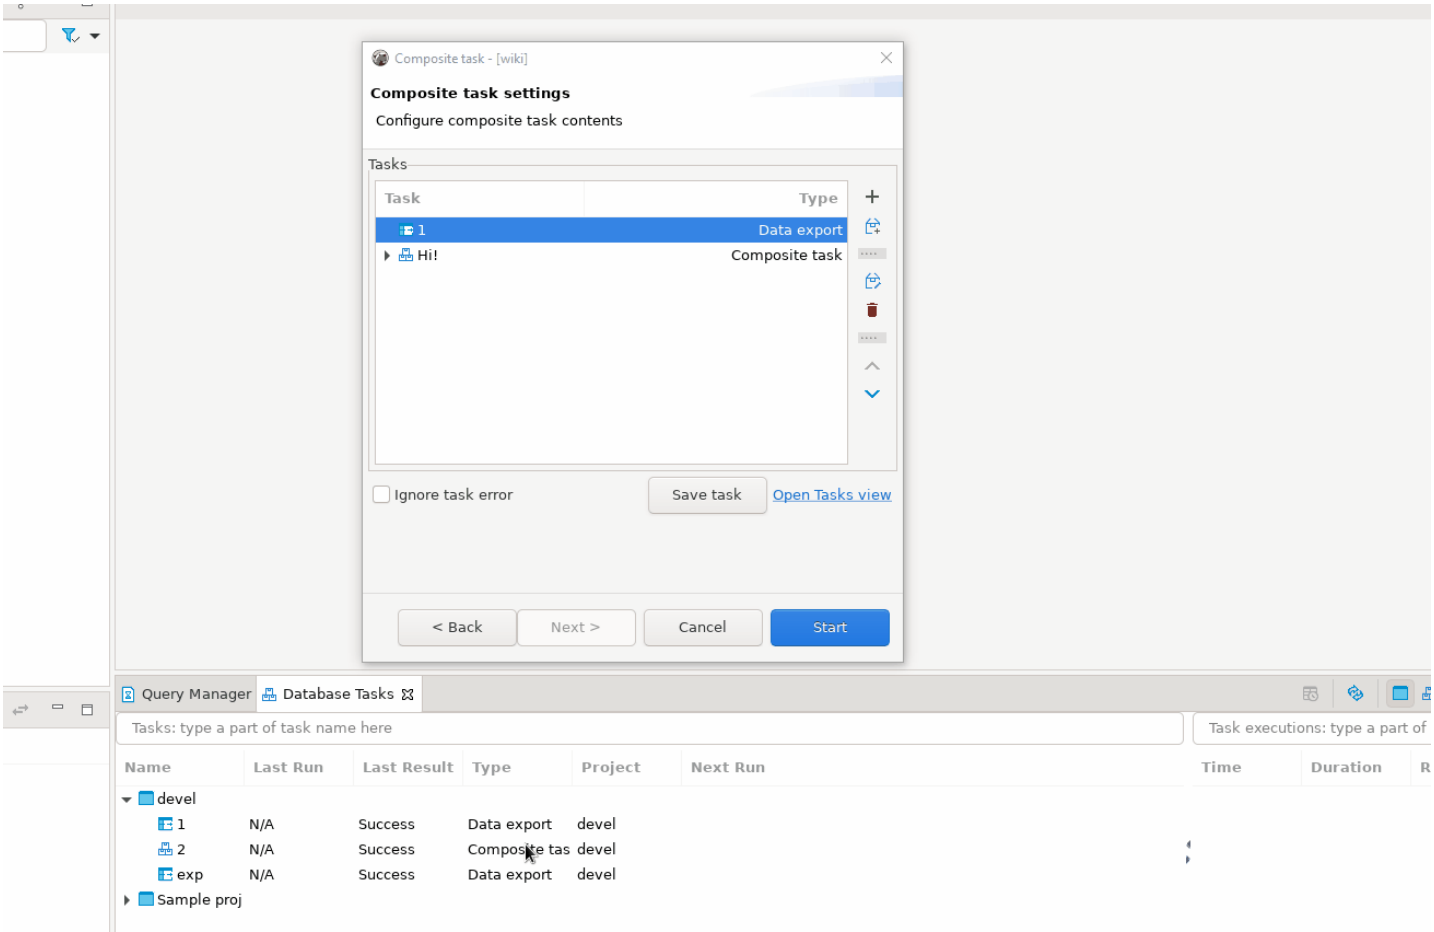

As a side note, you can add a composite task to your new composite task.

You can edit tasks in the same dialog, delete a task from a composite task, change the execution order:

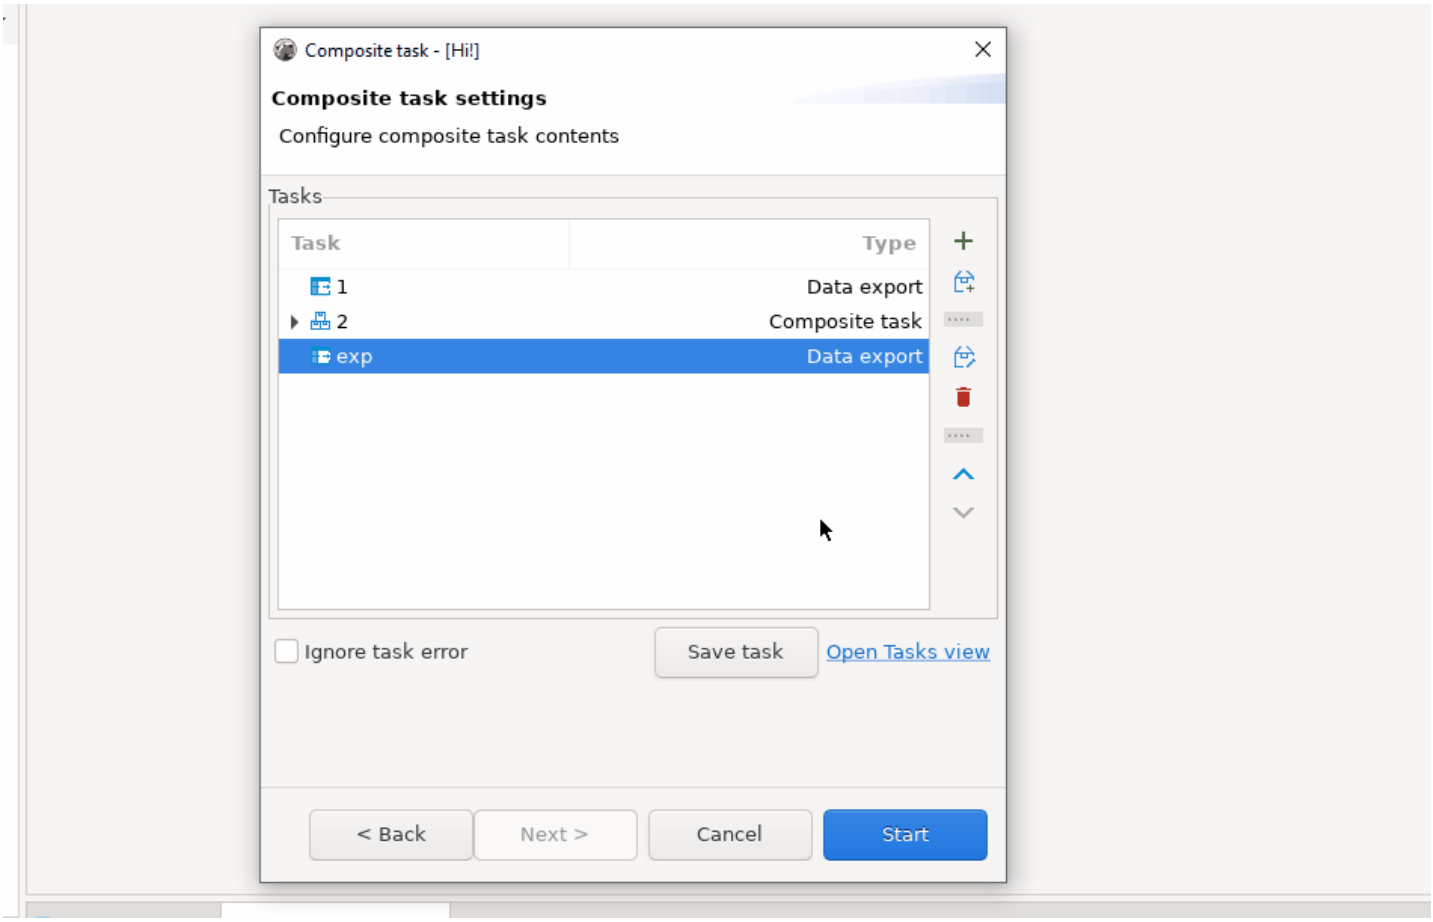

There is also a very important checkbox, *Ignore task error*. The tasks from the *composite task* are executed in the order they appear in the settings dialog. Executing a task from a *composite task* might produce an error that will block the next tasks from proceeding. The *Ignore task error* checkbox can be used to bypass this behavior.

MongoDB

Overview

DBeaver EE supports MongoDB schema browser, data viewer, SQL and JavaScript queries execution. Also it supports various administrative tools (like server sessions manager).
DBeaver uses MongoDB Java driver 3.8.0 to operate with server. It supports MongoDB servers from 2.x to 4.x.

Connecting to MongoDB Server

You can connect directly to a server or use SSH tunneling or SOCKS proxy.
You can specify server address as host/port/database configuration or you can enter target database URL with all necessary parameters:

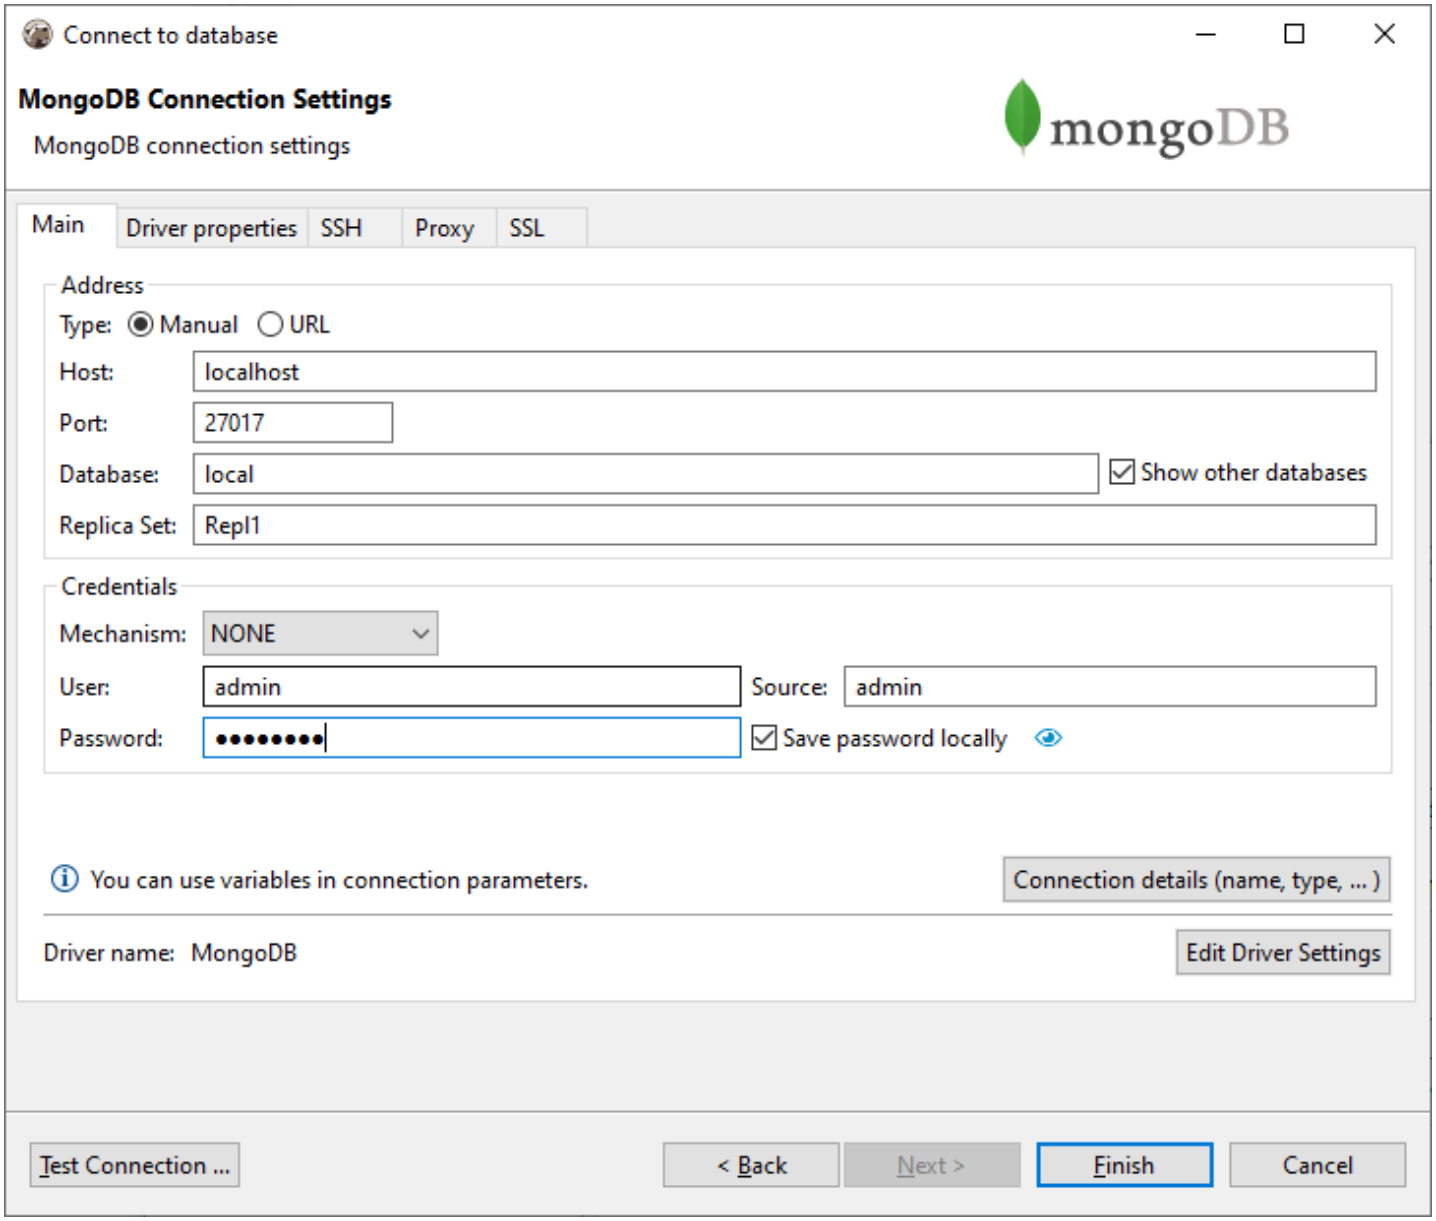
Connect to database
MongoDB Connection Settings
MongoDB connection settings
mongoDB
Main
Driver properties
SSH
Proxy
SSL
Address
Type: Manual URL
Host: localhost
Port: 27017
Database: local
Show other databases
Replica Set: Repl1
Credentials
Mechanism: NONE
User: admin
Source: admin
Password: ••••••••
Save password locally
You can use variables in connection parameters.
Connection details (name, type, ...)
Driver name: MongoDB
Edit Driver Settings
Test Connection ...
< Back
Next >
Finish
Cancel

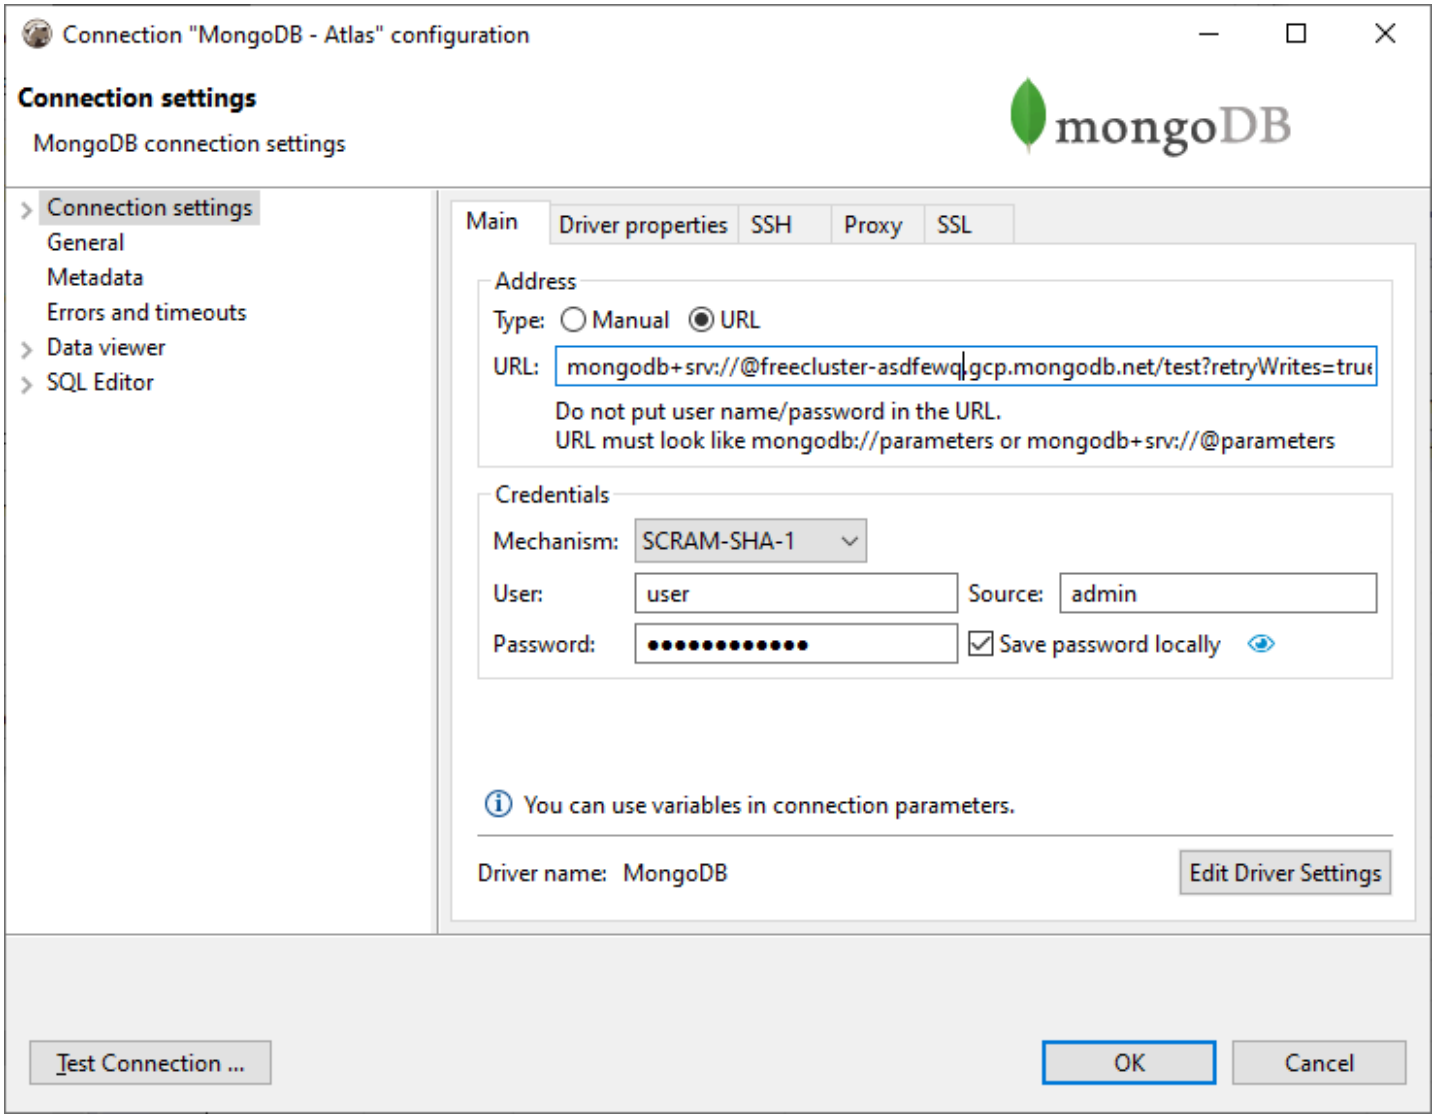
Connection "MongoDB - Atlas" configuration
Connection settings
MongoDB connection settings
mongoDB
Connection settings
General
Metadata
Errors and timeouts
Data viewer
SQL Editor
Main
Driver properties
SSH
Proxy
SSL
Address
Type: Manual URL
URL: mongodb+srv://@freecluster-asdfewq.gcp.mongodb.net/test?retryWrites=true
Do not put user name/password in the URL.
URL must look like mongodb://parameters or mongodb+srv://@parameters
Credentials
Mechanism: SCRAM-SHA-1
User: user
Source: admin
Password: ••••••••••••
Save password locally
You can use variables in connection parameters.
Driver name: MongoDB
Edit Driver Settings
Test Connection ...
OK
Cancel

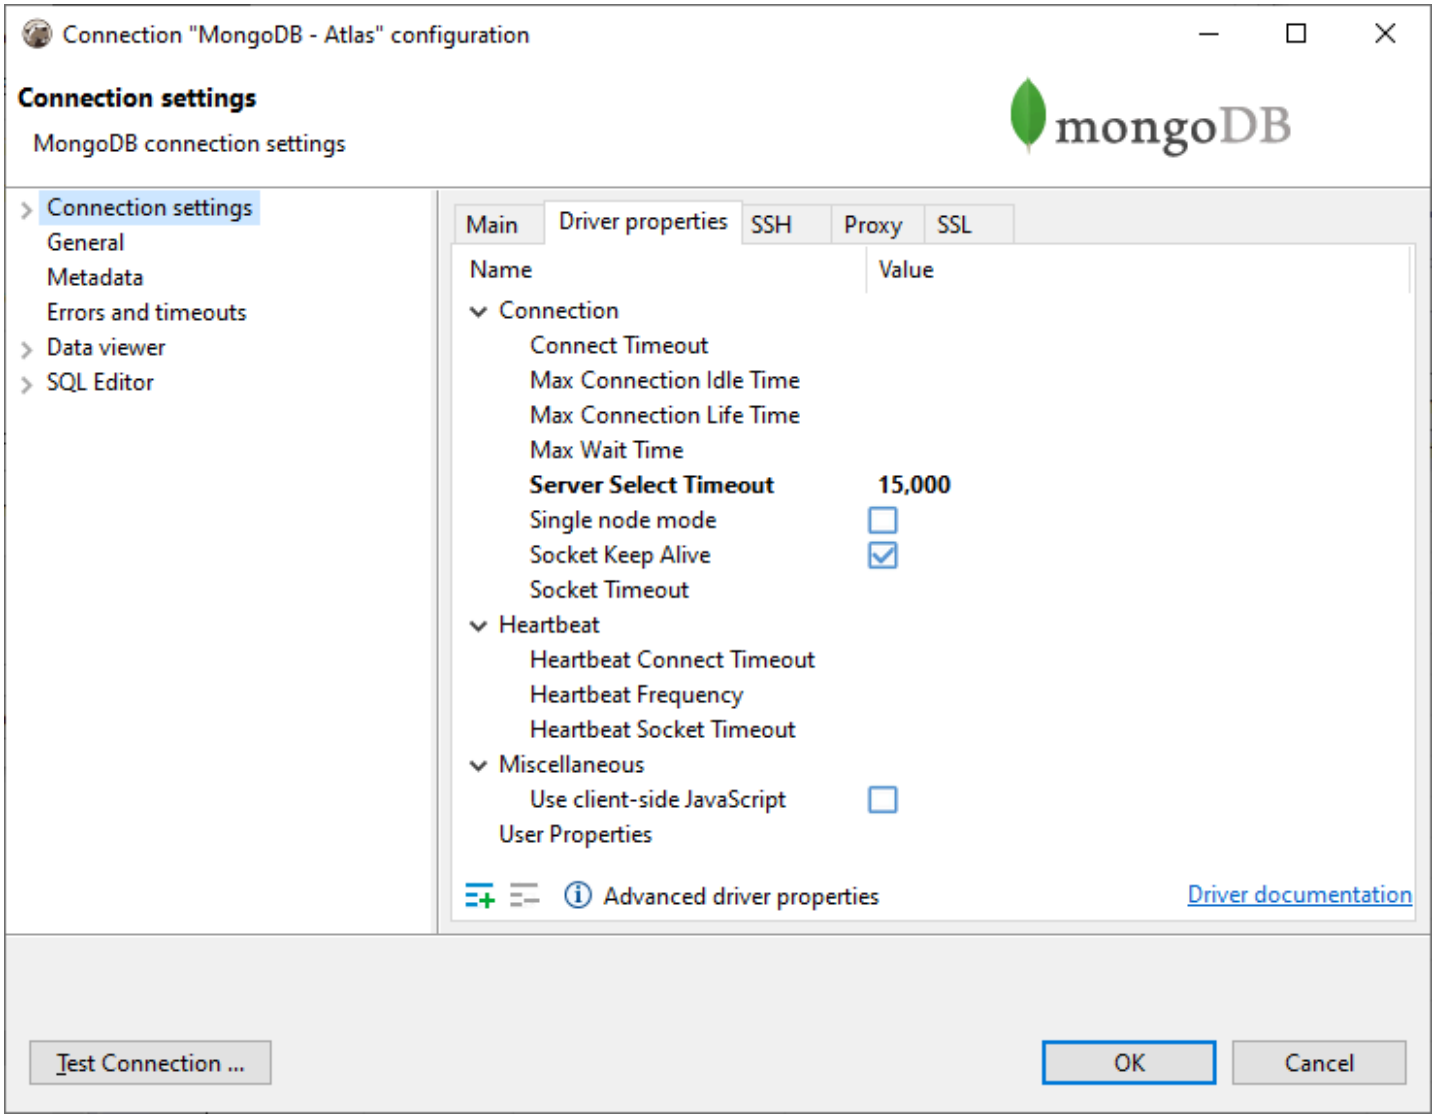
Connection "MongoDB - Atlas" configuration
Connection settings
MongoDB connection settings
mongoDB
Connection settings
General
Metadata
Errors and timeouts
Data viewer
SQL Editor
Main
Driver properties
SSH
Proxy
SSL
Name
Value
Connection
Connect Timeout
Max Connection Idle Time
Max Connection Life Time
Max Wait Time
Server Select Timeout
15,000
Single node mode
Socket Keep Alive
Socket Timeout
Heartbeat
Heartbeat Connect Timeout
Heartbeat Frequency
Heartbeat Socket Timeout
Miscellaneous
Use client-side JavaScript
User Properties
Advanced driver properties
Driver documentation
Test Connection ...
OK
Cancel

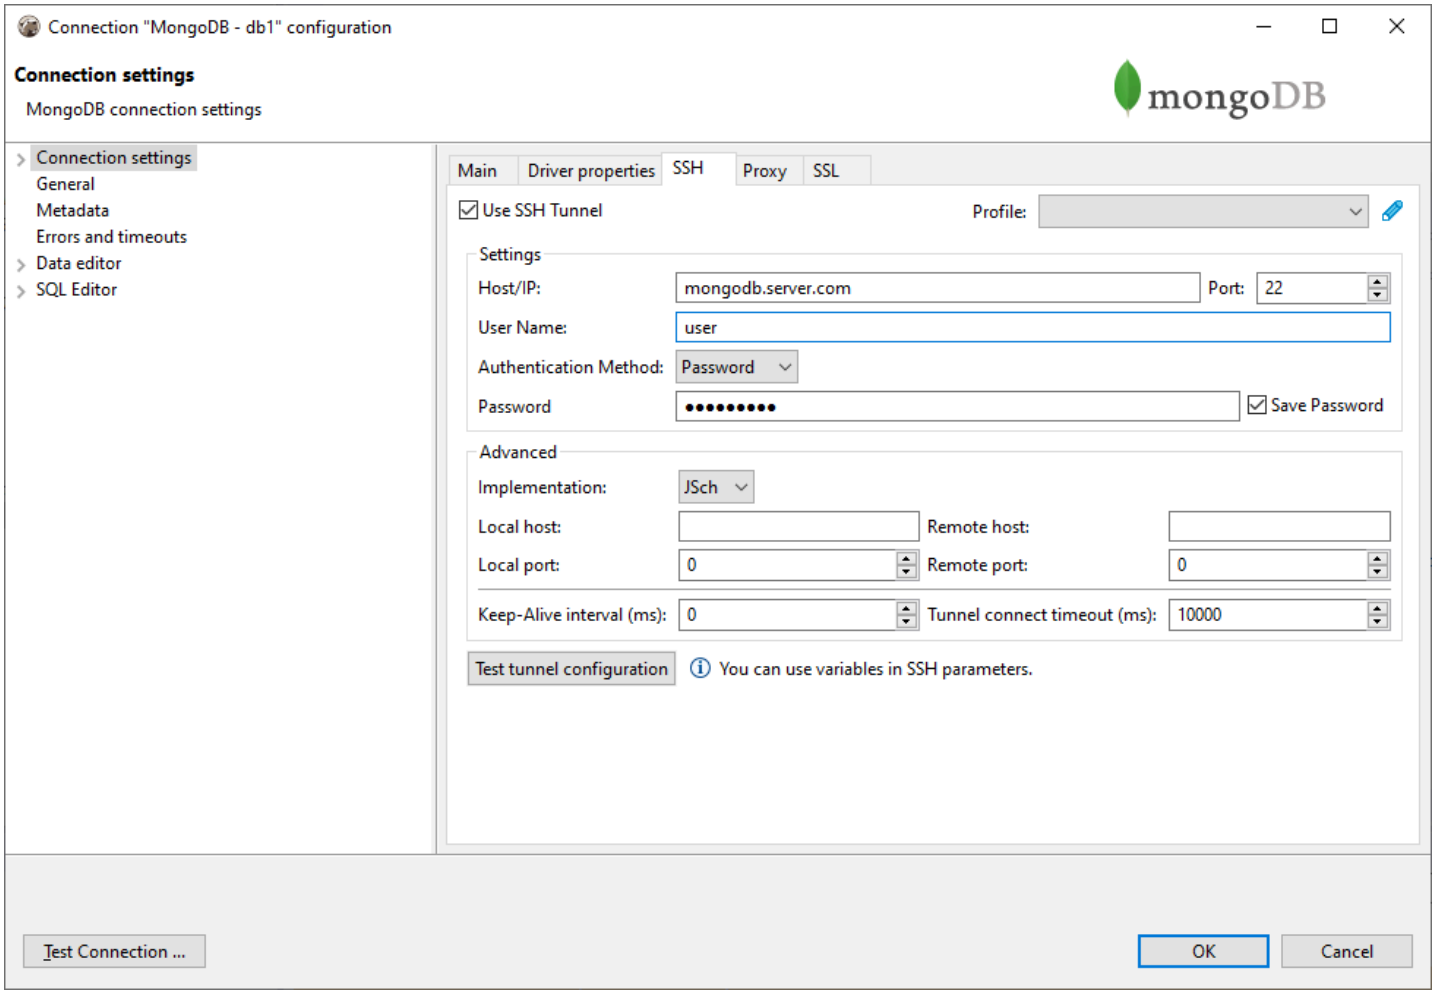

Browsing Mongo collections

You can view/edit MongoDB collections content as standard relational tables (grid/plain text presentations) or as JSON documents.
Presentation can be switched in the Results Viewer toolbar.
In grid DBeaver will try to unify all documents in some particular collection (as they have the same structure/the same set of properties).

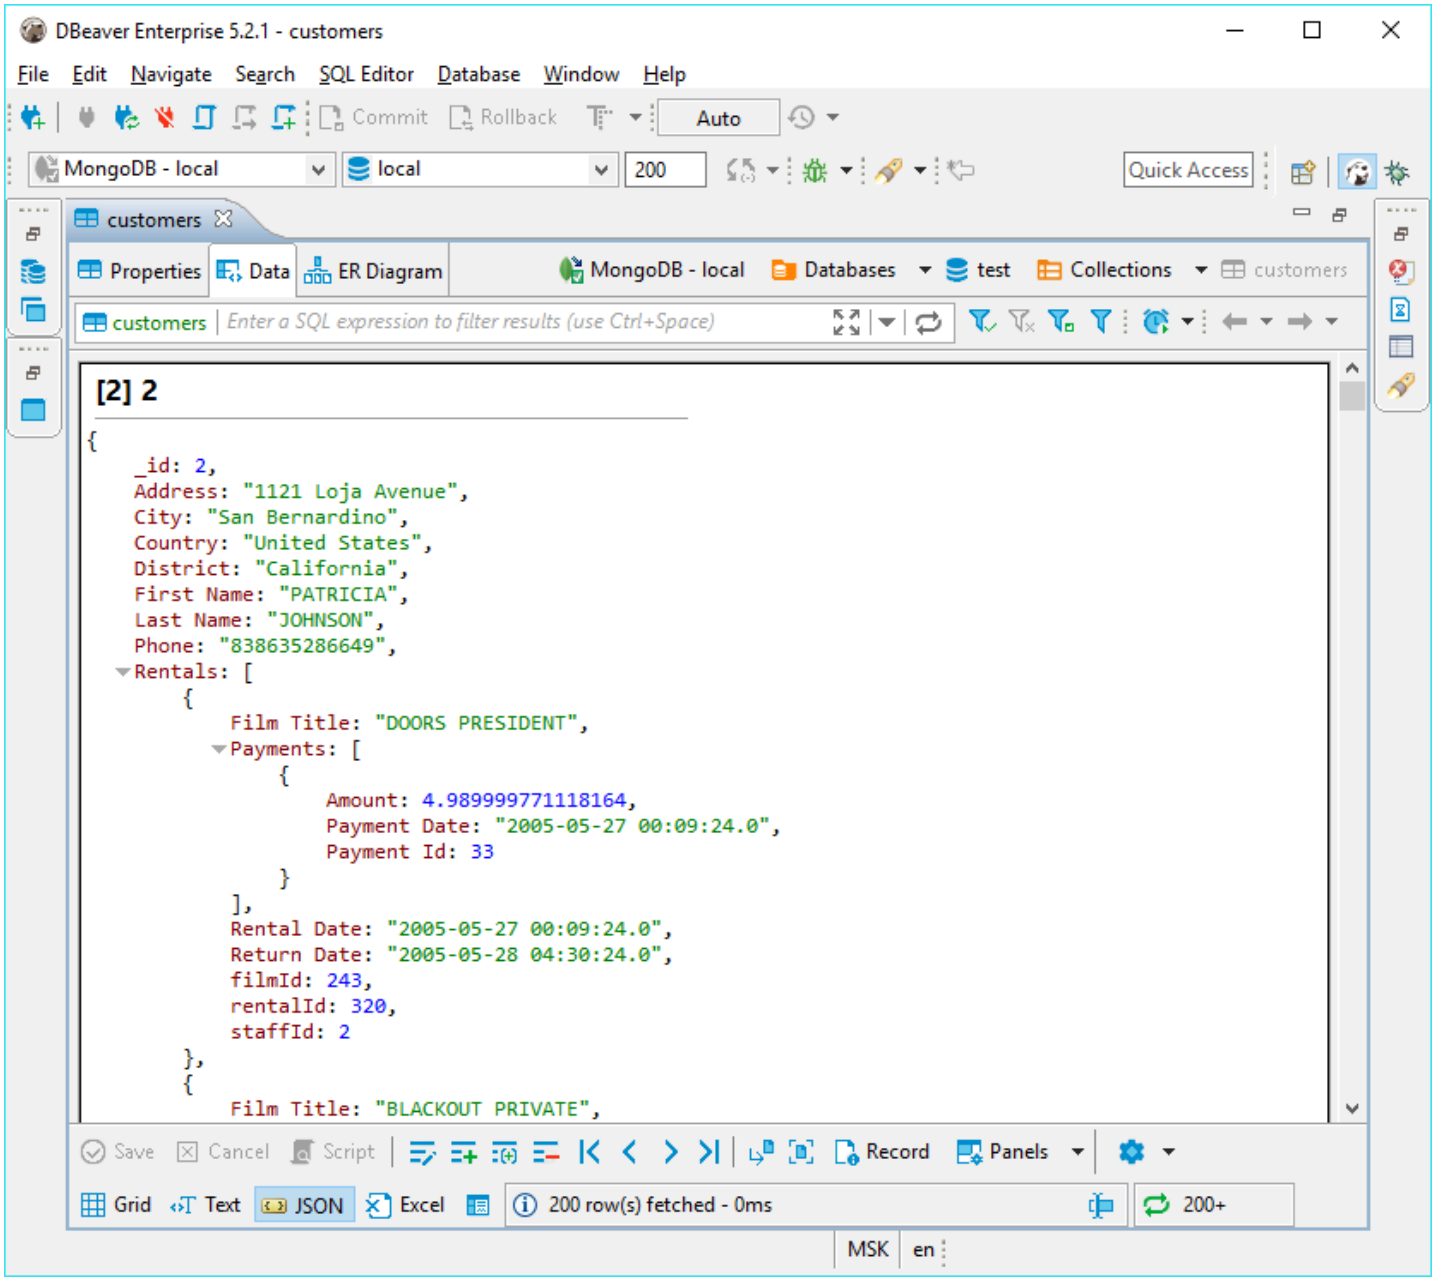
DBeaver Enterprise 5.2.1 - customers
File Edit Navigate Search SQL Editor Database Window Help
Commit Rollback Auto
MongoDB - local local 200 Quick Access
customers
Properties Data ER Diagram MongoDB - local Databases test Collections customers
customers Enter a SQL expression to filter results (use Ctrl+Space)
[2] 2
{
_id: 2,
Address: "1121 Loja Avenue",
City: "San Bernardino",
Country: "United States",
District: "California",
First Name: "PATRICIA",
Last Name: "JOHNSON",
Phone: "838635286649",
Rentals: [
{
Film Title: "DOORS PRESIDENT",
Payments: [
{
Amount: 4.989999771118164,
Payment Date: "2005-05-27 00:09:24.0",
Payment Id: 33
}
],
Rental Date: "2005-05-27 00:09:24.0",
Return Date: "2005-05-28 04:30:24.0",
filmId: 243,
rentalId: 320,
staffId: 2
},
{
Film Title: "BLACKOUT PRIVATE",
Save Cancel Script Record Panels
Grid Text JSON Excel 200 row(s) fetched - 0ms 200+
MSK en

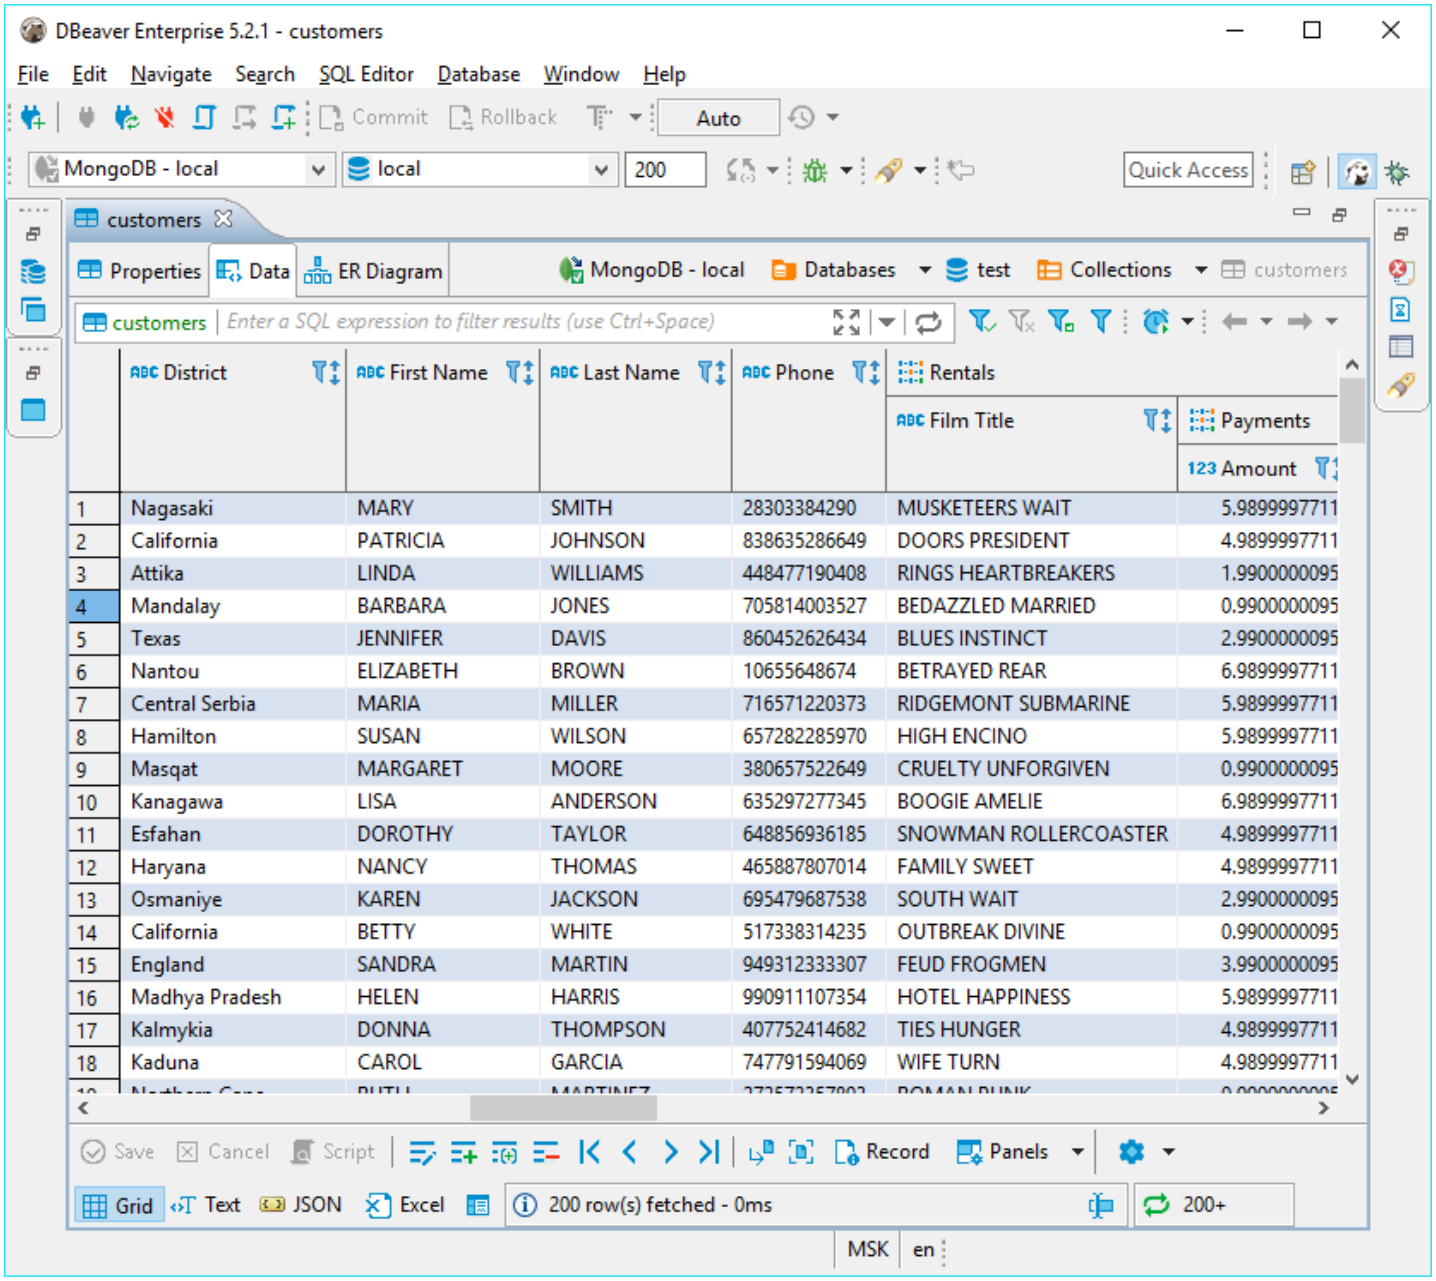
DBeaver Enterprise 5.2.1 - customers
File Edit Navigate Search SQL Editor Database Window Help
Commit Rollback Auto
MongoDB - local local 200 Quick Access
customers
Properties Data ER Diagram MongoDB - local Databases test Collections customers
customers Enter a SQL expression to filter results (use Ctrl+Space)
District First Name Last Name Phone Rentals
Film Title Payments
Amount
1 Nagasaki MARY SMITH 28303384290 MUSKETEERS WAIT 5.9899997711
2 California PATRICIA JOHNSON 838635286649 DOORS PRESIDENT 4.9899997711
3 Attika LINDA WILLIAMS 448477190408 RINGS HEARTBREAKERS 1.9900000095
4 Mandalay BARBARA JONES 705814003527 BEDAZZLED MARRIED 0.9900000095
5 Texas JENNIFER DAVIS 860452626434 BLUES INSTINCT 2.9900000095
6 Nantou ELIZABETH BROWN 10655648674 BETRAYED REAR 6.9899997711
7 Central Serbia MARIA MILLER 716571220373 RIDGEMONT SUBMARINE 5.9899997711
8 Hamilton SUSAN WILSON 657282285970 HIGH ENCINO 5.9899997711
9 Masqat MARGARET MOORE 380657522649 CRUELTY UNFORGIVEN 0.9900000095
10 Kanagawa LISA ANDERSON 635297277345 BOOGIE AMELIE 6.9899997711
11 Esfahan DOROTHY TAYLOR 648856936185 SNOWMAN ROLLERCOASTER 4.9899997711
12 Haryana NANCY THOMAS 465887807014 FAMILY SWEET 4.9899997711
13 Osmaniye KAREN JACKSON 695479687538 SOUTH WAIT 2.9900000095
14 California BETTY WHITE 517338314235 OUTBREAK DIVINE 0.9900000095
15 England SANDRA MARTIN 949312333307 FEUD FROGMEN 3.9900000095
16 Madhya Pradesh HELEN HARRIS 990911107354 HOTEL HAPPINESS 5.9899997711
17 Kalmykia DONNA THOMPSON 407752414682 TIES HUNGER 4.9899997711
18 Kaduna CAROL GARCIA 747791594069 WIFE TURN 4.9899997711
Save Cancel Script Record Panels
Grid Text JSON Excel 200 row(s) fetched - 0ms 200+
MSK en

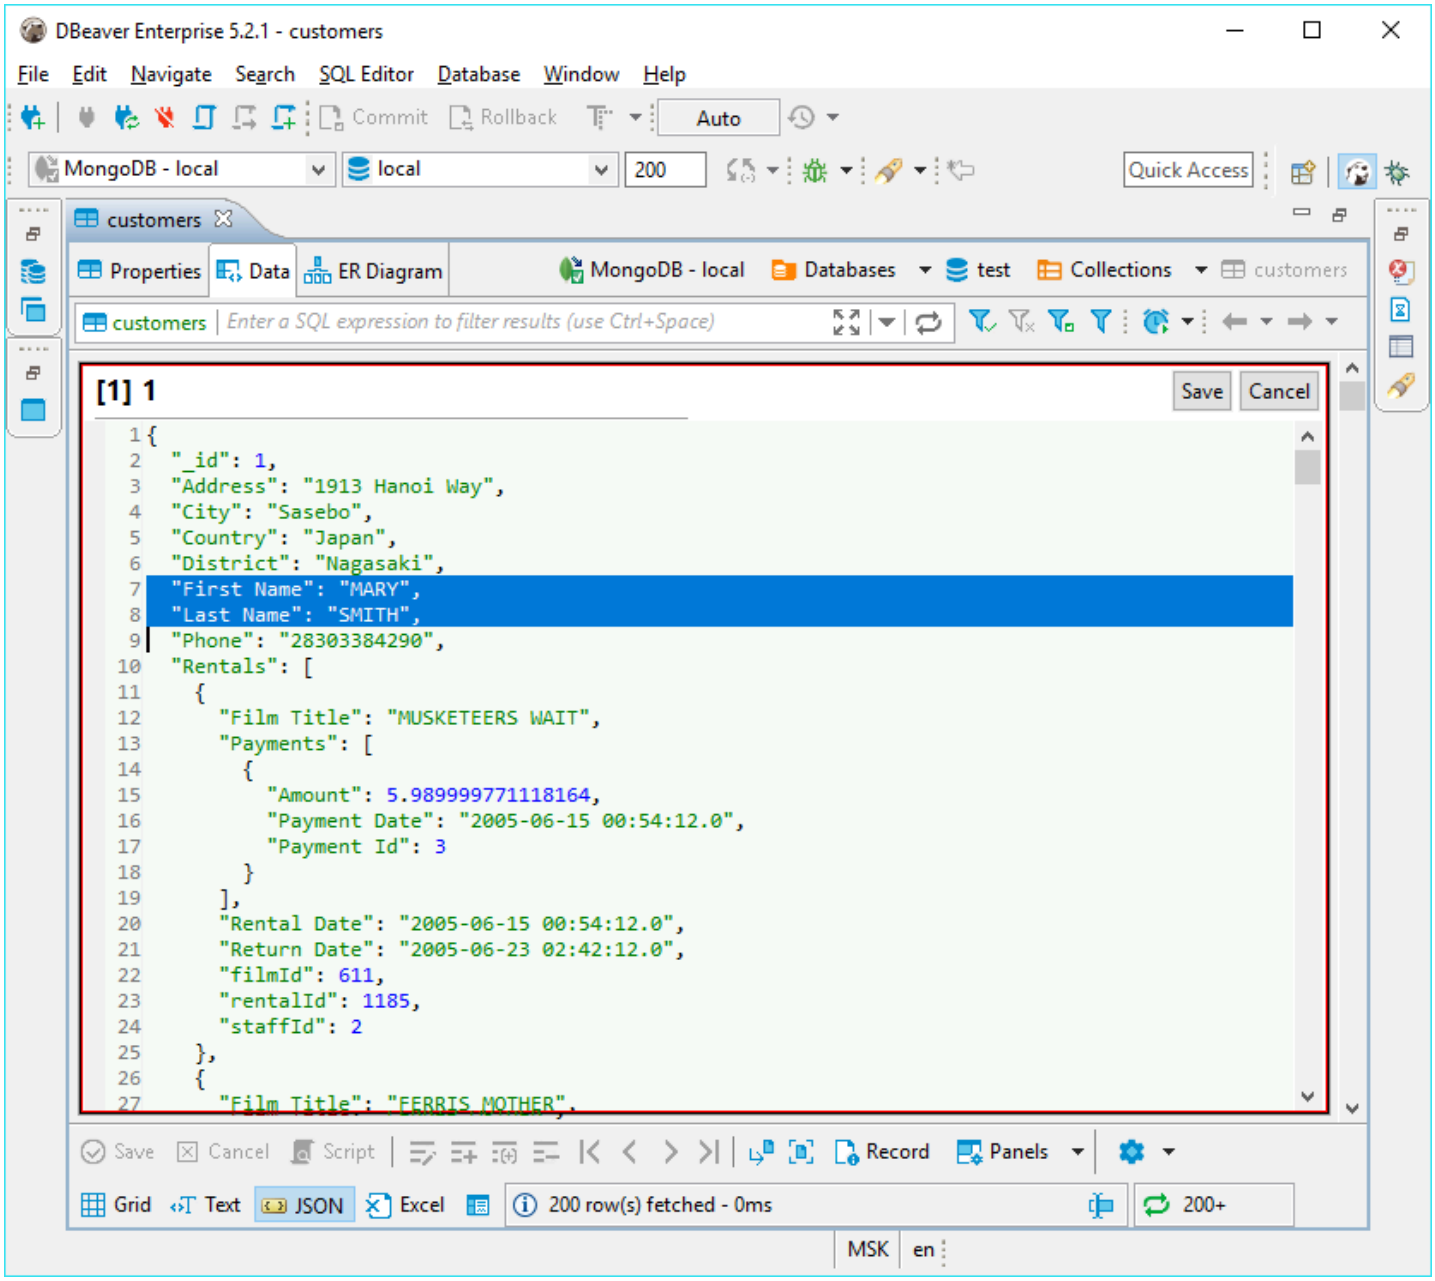

Executing JavaScript

JS statements can be executed in SQL editor as usual.

Mongo scripting reference

Following example creates a user in the current database.

```
db.createUser({
    user: 'testuser',
    pwd: 'test',
    roles: []
})
```

This example returns all documents in collection 'test_col':

```
db.test_col.find().toArray()
```

Note: script will be executed in the current database.
You can not set explicit database name in your query.
Current database can be changed in SQL Editor toolbar or in Database Navigator.

Executing SQL

You can use standard SQL statements (`SELECT` , `INSERT` , `UPDATE` , `DELETE`) to manipulate Mongo data.

```
SELECT * FROM test_col
WHERE propName.subProp='value'

UPDATE FROM test_col
SET propsName.val1=123
WHERE propName.subProp='value'
```

SELECT queries support `WHERE` , `ORDER BY` , `GROUP BY` , `JOIN` and `HAVING` clauses.

Nested JSON fields can be divided by dot. If your field contains any special characters (e.g. spaces, dashes, etc) you must enclose it with double quotes. For example:

```
SELECT title FROM movies WHERE info.'imdb-details'.rating > 6
```

Working with dates

If you need to operate with dates then you must specify them in ISO format. It is possible in both JavaScript and SQL dialect:

```
db.dates.insert([
    { value: new Date('2016-05-18T16:00:00Z') },
    { value: new Date('2017-05-18T16:00:00Z') },
    { value: new Date('2018-05-18T16:00:00Z') },
    { value: new Date('2019-05-18T16:00:00Z') },
    { value: new Date('2020-05-18T16:00:00Z') }
])
```

Querying data in JavaScript:

```
db.dates.find({
    value: { $gte: new Date('2018-05-18T16:00:00Z') }
}).toArray()
```

Querying data in SQL dialect (ISO and UNIX timestamp, in milliseconds):

```
SELECT value FROM dates
WHERE value > ISODate('2018-05-18T16:00:00.000Z')
ORDER BY value DESC

SELECT value FROM dates
WHERE value > ISODate(1526659200000)
ORDER BY value DESC
```

Working with object IDs

When you need to find document by ID you must use function `ObjectId` :

```
SELECT * FROM documents
WHERE _id = ObjectId('5f9c458018e3c69d0adc0fbd')
ORDER BY value DESC
```

Working with JOINs

Currently SQL dialect for MongoDB supports `LEFT JOIN` and `INNER JOIN` :

```
SELECT
    ar.Name as Artist,
    al.Title as Album,
    SUM(tr.Milliseconds) as Duration
FROM Track tr
INNER JOIN Album al ON tr.AlbumId = al.AlbumId
INNER JOIN Artist ar ON al.ArtistId = ar.ArtistId
GROUP BY Artist, Album
ORDER BY Duration DESC
```

The only limitation is that you have to specify aliases for both source and target tables in particular order:

```
SELECT *
FROM <source> <source-alias>
INNER JOIN <target> <target-alias> ON <source-alias>.column = <target-alias>.column
```

Note that executing following script will not result in merged document but separate documents for `Track` and `Album`:

```
SELECT *
FROM Track tr
INNER JOIN Album al ON tr.AlbumId = al.AlbumId
```

Cassandra

Overview

DBeaver EE supports Cassandra schema browser, data viewer and CQL queries execution. Also it supports various administrative tools.

Connecting to Cassandra cluster

You can connect directly to a server or use SSH tunneling or SOCKS proxy. DBeaver uses DataStax Java driver to operate with server. It supports Cassandra servers 2.x, 3.x or higher.

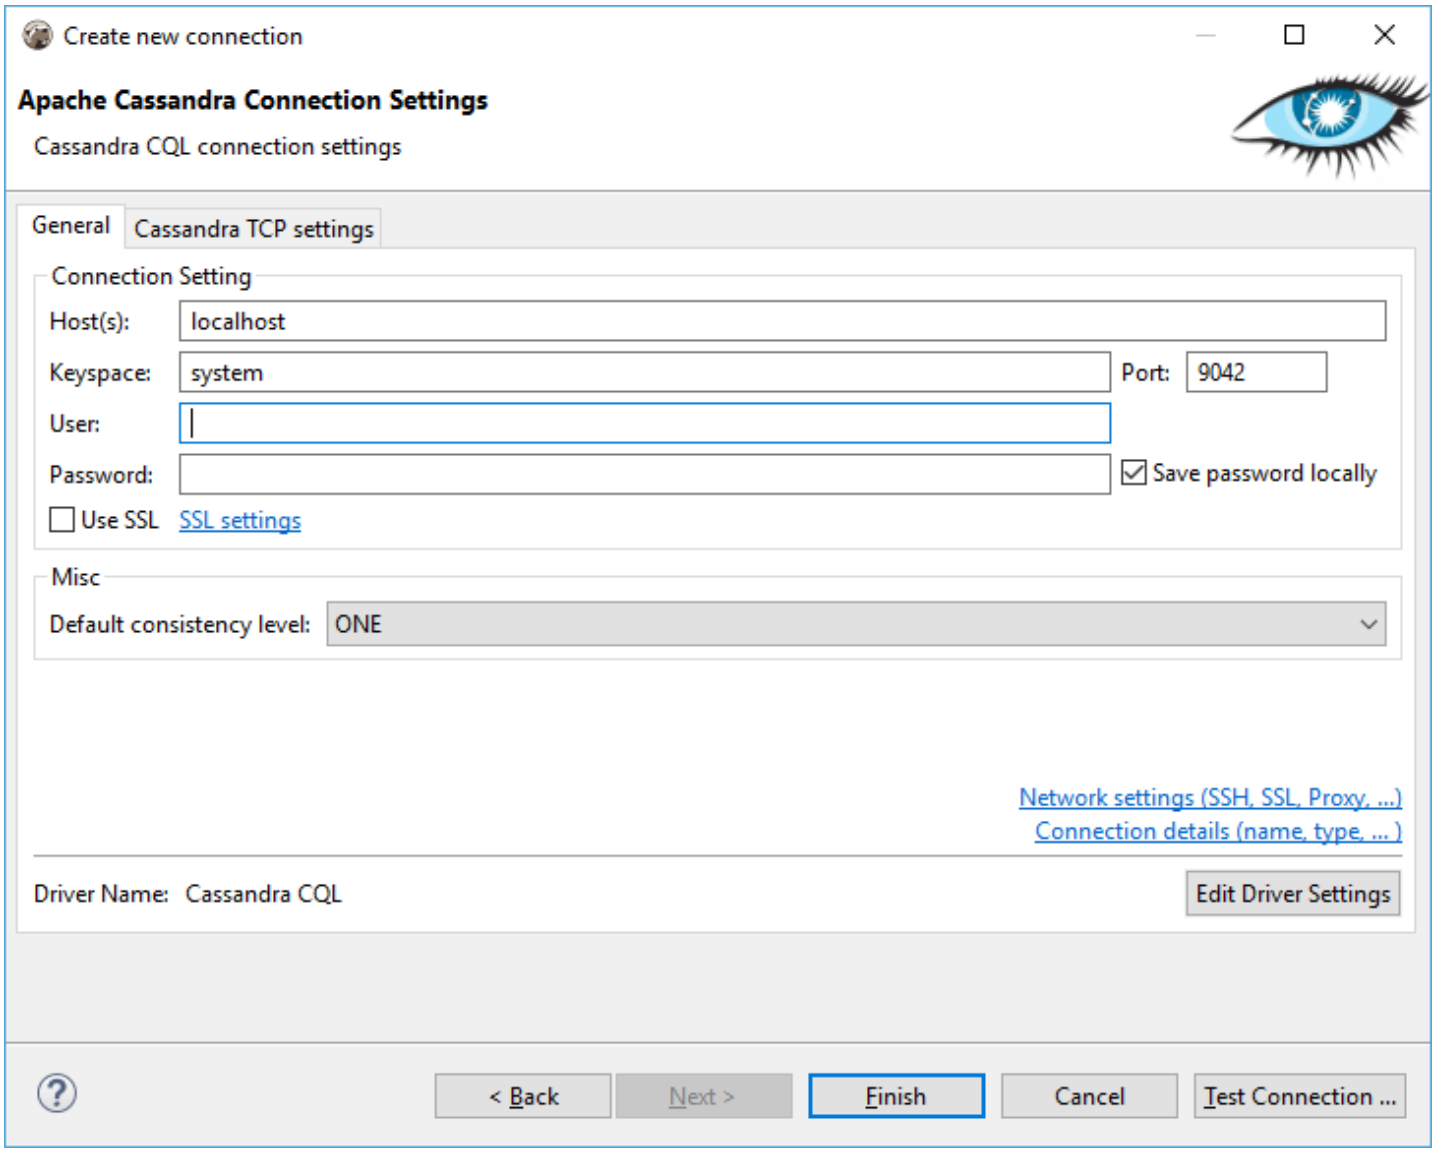
Create new connection
Apache Cassandra Connection Settings
Cassandra CQL connection settings
General
Cassandra TCP settings
Connection Setting
Host(s):
localhost
Keyspace:
system
Port:
9042
User:
Password:
Save password locally
Use SSL
SSL settings
Misc
Default consistency level:
ONE
Network settings (SSH, SSL, Proxy, ...)
Connection details (name, type, ...)
Driver Name: Cassandra CQL
Edit Driver Settings
< Back
Next >
Finish
Cancel
Test Connection ...

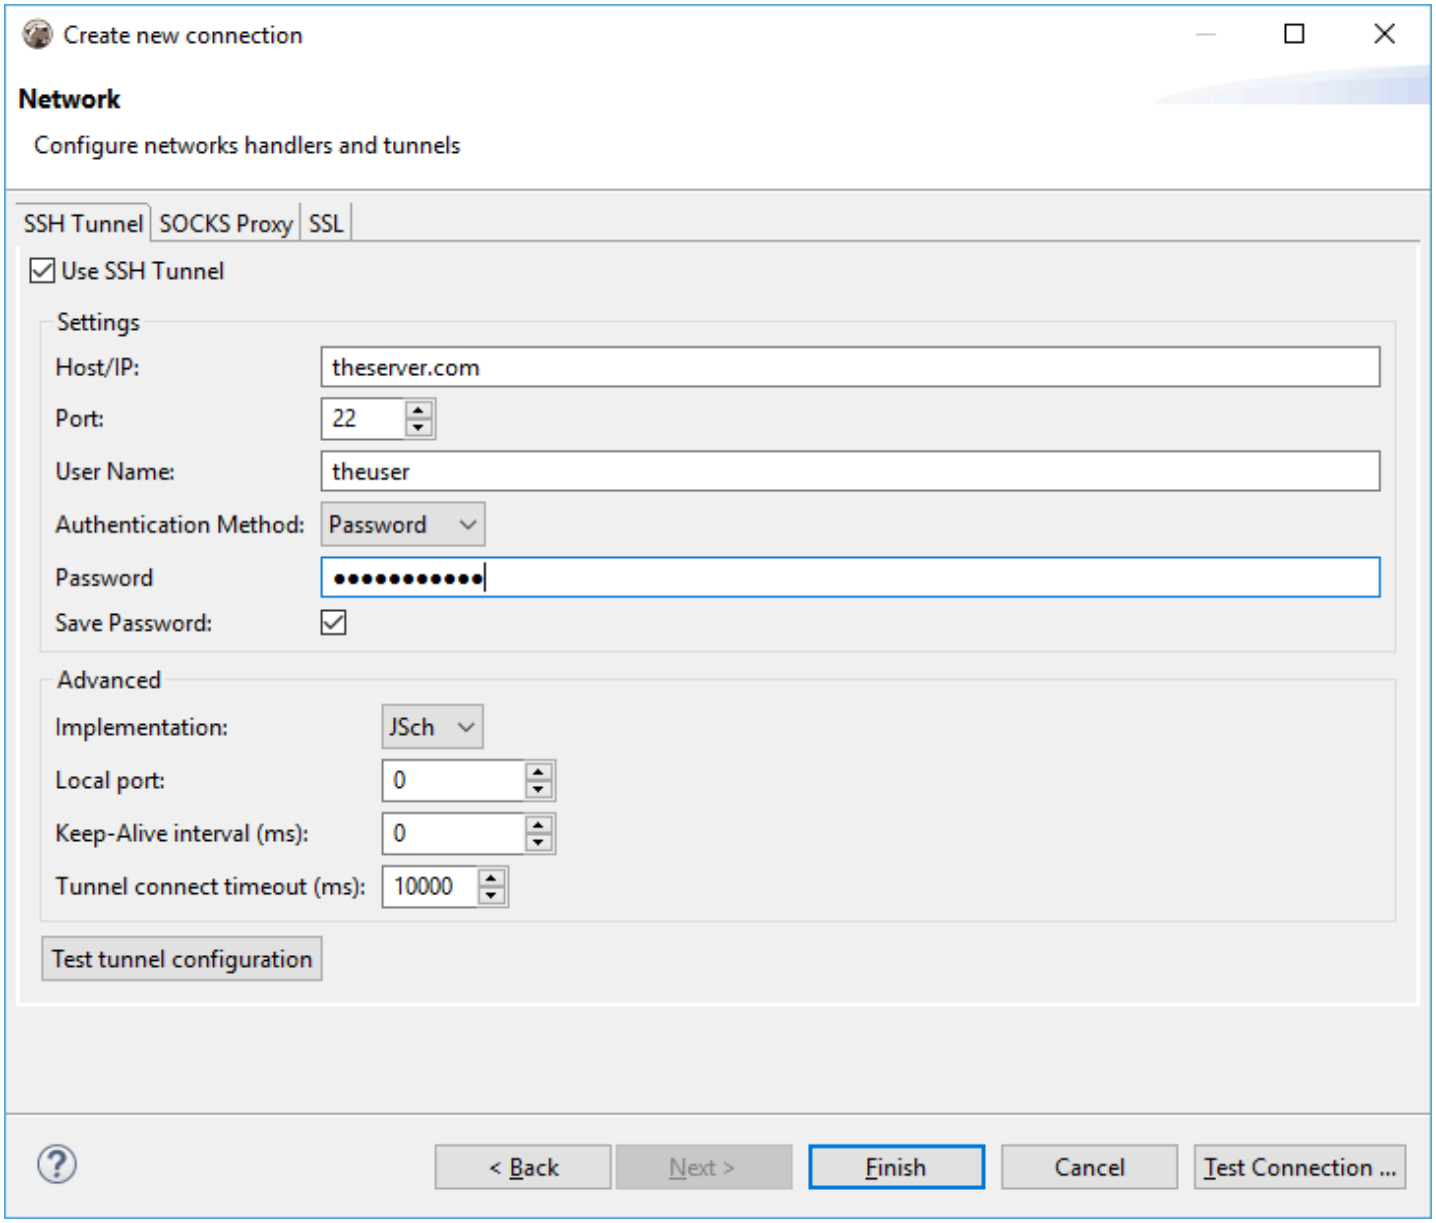
Create new connection
Network
Configure networks handlers and tunnels
SSH Tunnel
SOCKS Proxy
SSL
Use SSH Tunnel
Settings
Host/IP:
theserver.com
Port:
22
User Name:
theuser
Authentication Method:
Password
Password
Save Password:
Advanced
Implementation:
JSch
Local port:
0
Keep-Alive interval (ms):
0
Tunnel connect timeout (ms):
10000
Test tunnel configuration
< Back
Next >
Finish
Cancel
Test Connection ...

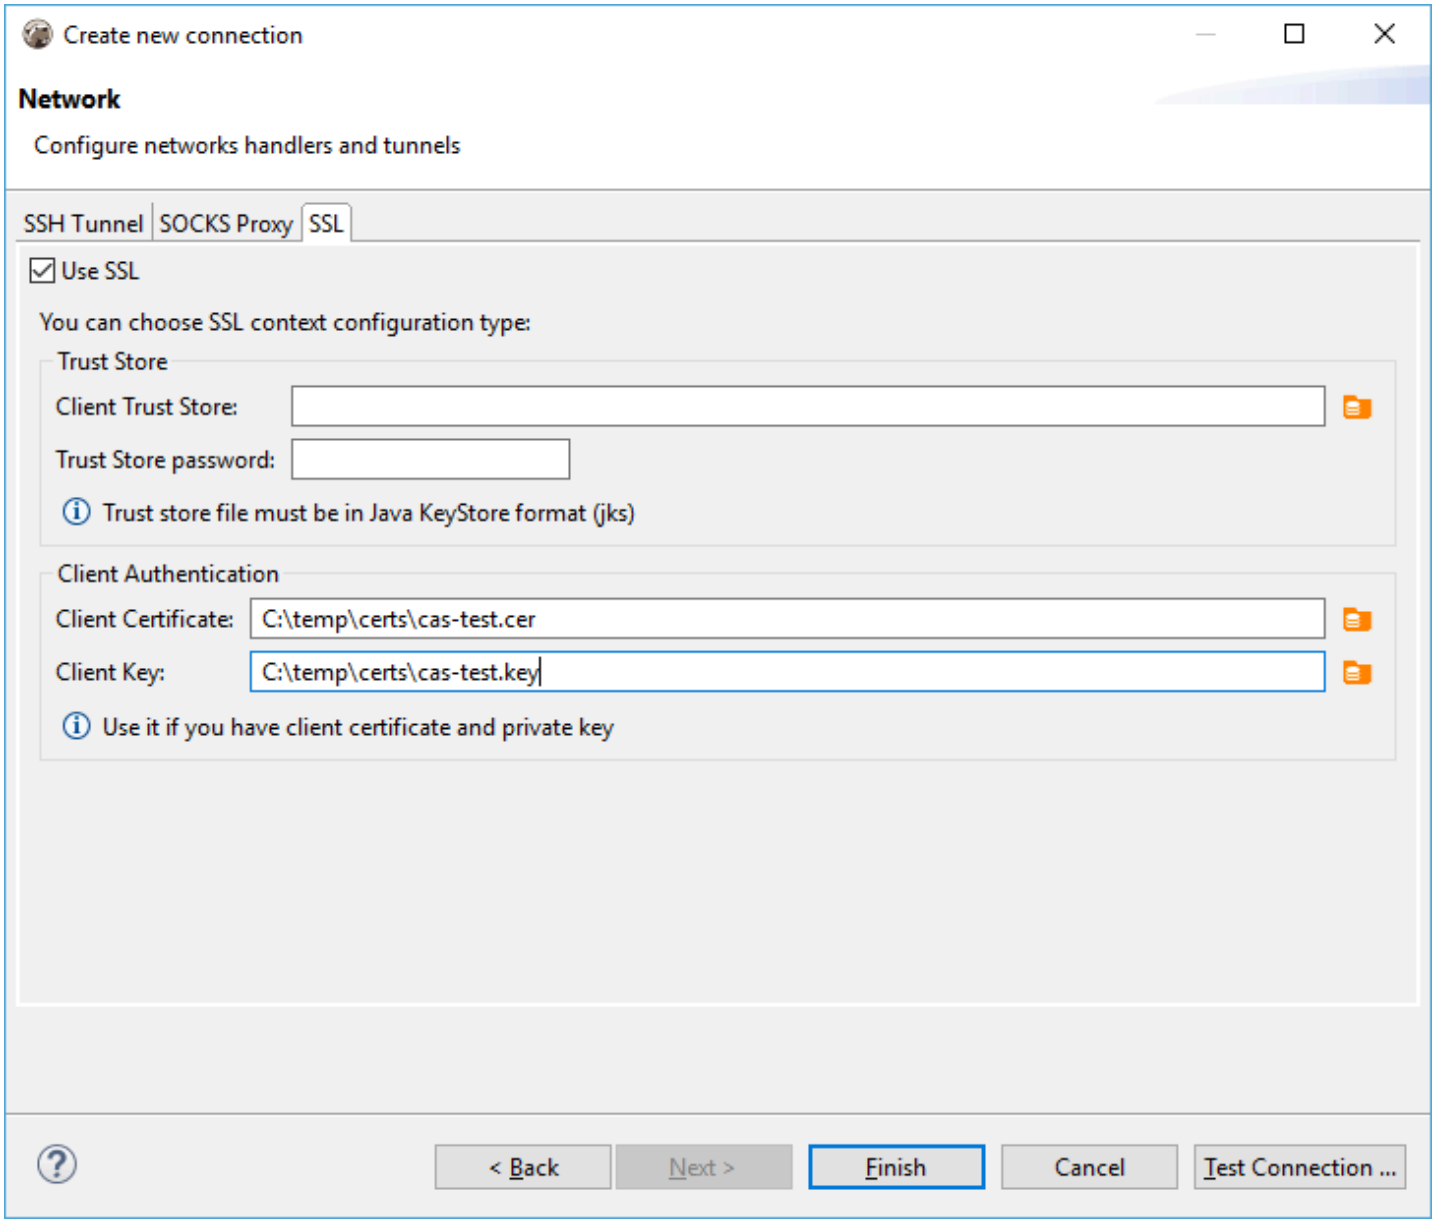
Create new connection
Network
Configure networks handlers and tunnels
SSH Tunnel
SOCKS Proxy
SSL
Use SSL
You can choose SSL context configuration type:
Trust Store
Client Trust Store:
Trust Store password:
Trust store file must be in Java KeyStore format (jks)
Client Authentication
Client Certificate:
C:\temp\certs\cas-test.cer
Client Key:
C:\temp\certs\cas-test.key
Use it if you have client certificate and private key
< Back
Next >
Finish
Cancel
Test Connection ...

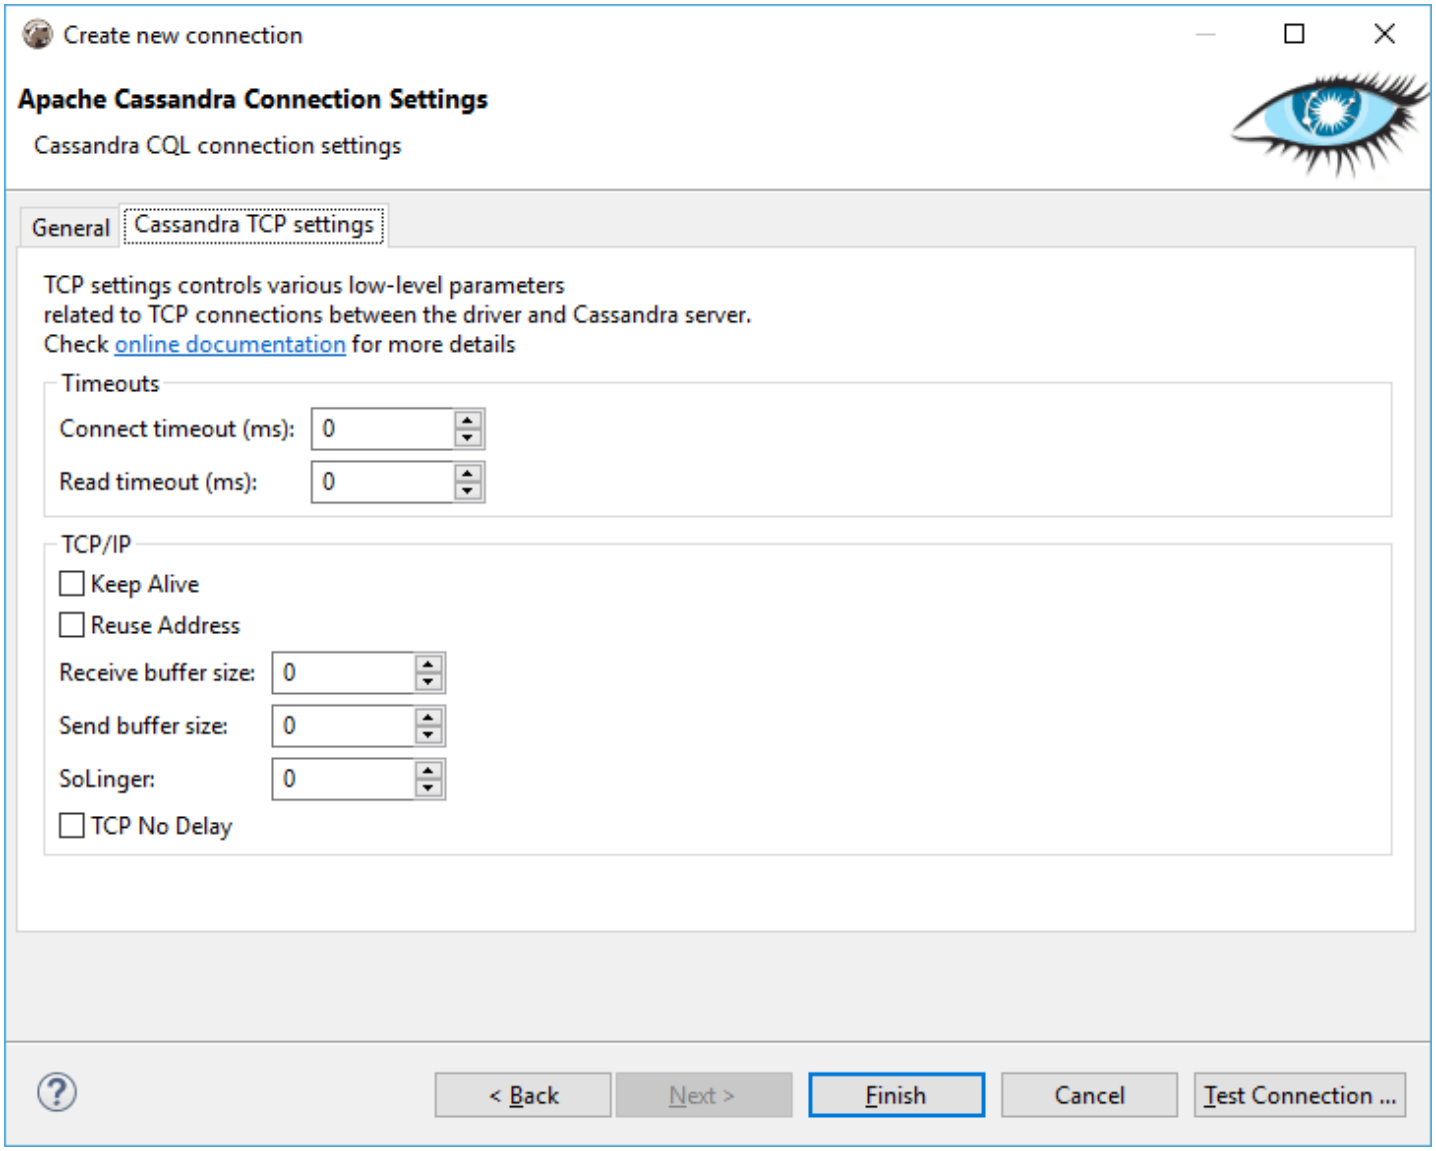

Browsing Cassandra tables

You can browse, view, edit and filter Cassandra tables the same way as with regular (relational) tables. However, being a distributed key-value database, Cassandra doesn't support any kind of referential integrity. There are no foreign keys, references, etc.
Note that Cassandra has very advanced (comparing to relational databases) data type system. Each column may be a collection, map or set of values (with very big number of values). In some cases this makes browsing data in the "Grid" mode inconvenient.

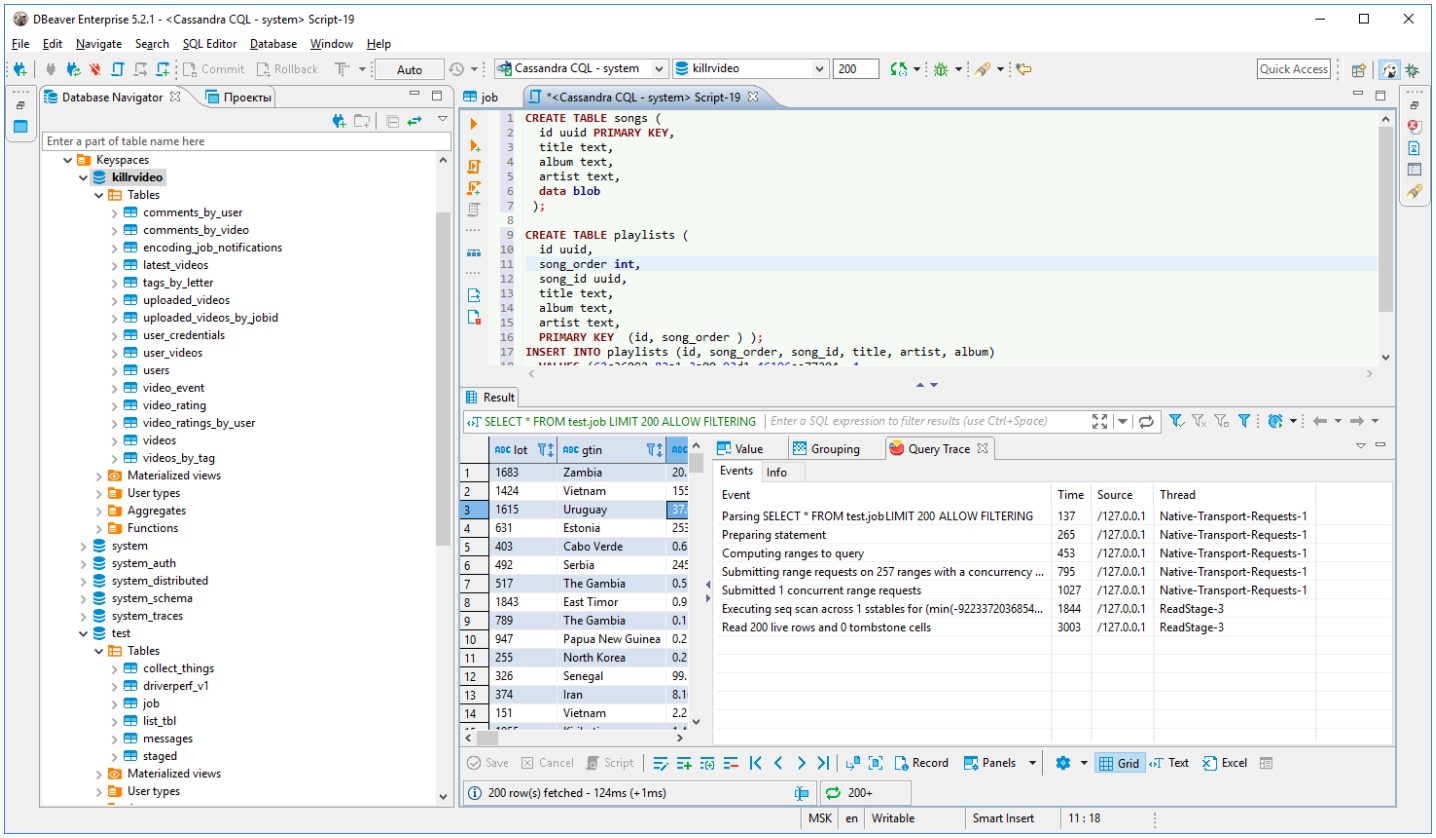
DBeaver Enterprise 5.2.1 - <Cassandra CQL - system> Script-19
File Edit Navigate Search SQL Editor Database Window Help
Database Navigator
Enter a part of table name here
Keyspaces
killrvideo
Tables
comments_by_user
comments_by_video
encoding_job_notifications
latest_videos
tags_by_letter
uploaded_videos
uploaded_videos_by_jobid
user_credentials
user_videos
users
video_event
video_rating
video_ratings_by_user
videos
videos_by_tag
Materialized views
User types
Aggregates
Functions
system
system_auth
system_distributed
system_schema
system_traces
test
Tables
collect_things
driverperf_v1
job
list_tbl
messages
staged
Materialized views
User types
job
*<Cassandra CQL - system> Script-19
CREATE TABLE songs (
id uuid PRIMARY KEY,
title text,
album text,
artist text,
data blob
);
CREATE TABLE playlists (
id uuid,
song_order int,
song_id uuid,
title text,
album text,
artist text,
PRIMARY KEY (id, song_order));
INSERT INTO playlists (id, song_order, song_id, title, artist, album)
Result
SELECT * FROM test.job LIMIT 200 ALLOW FILTERING
Enter a SQL expression to filter results (use Ctrl+Space)
lot
gtin
1 1683 Zambia
2 1424 Vietnam
3 1615 Uruguay
4 631 Estonia
5 403 Cabo Verde
6 492 Serbia
7 517 The Gambia
8 1843 East Timor
9 789 The Gambia
10 947 Papua New Guinea
11 255 North Korea
12 326 Senegal
13 374 Iran
14 151 Vietnam
Value
Grouping
Query Trace
Events
Info
Event
Time
Source
Thread
Parsing SELECT * FROM test.job LIMIT 200 ALLOW FILTERING 137 /127.0.0.1 Native-Transport-Requests-1
Preparing statement 265 /127.0.0.1 Native-Transport-Requests-1
Computing ranges to query 453 /127.0.0.1 Native-Transport-Requests-1
Submitting range requests on 257 ranges with a concurrency ... 795 /127.0.0.1 Native-Transport-Requests-1
Submitted 1 concurrent range requests 1027 /127.0.0.1 Native-Transport-Requests-1
Executing seq scan across 1 sstables for (min(-9223372036854... 1844 /127.0.0.1 ReadStage-3
Read 200 live rows and 0 tombstone cells 3003 /127.0.0.1 ReadStage-3
Save Cancel Script Record Panels Grid Text Excel
200 row(s) fetched - 124ms (+1ms)
MSK en Writable Smart Insert 11 : 18

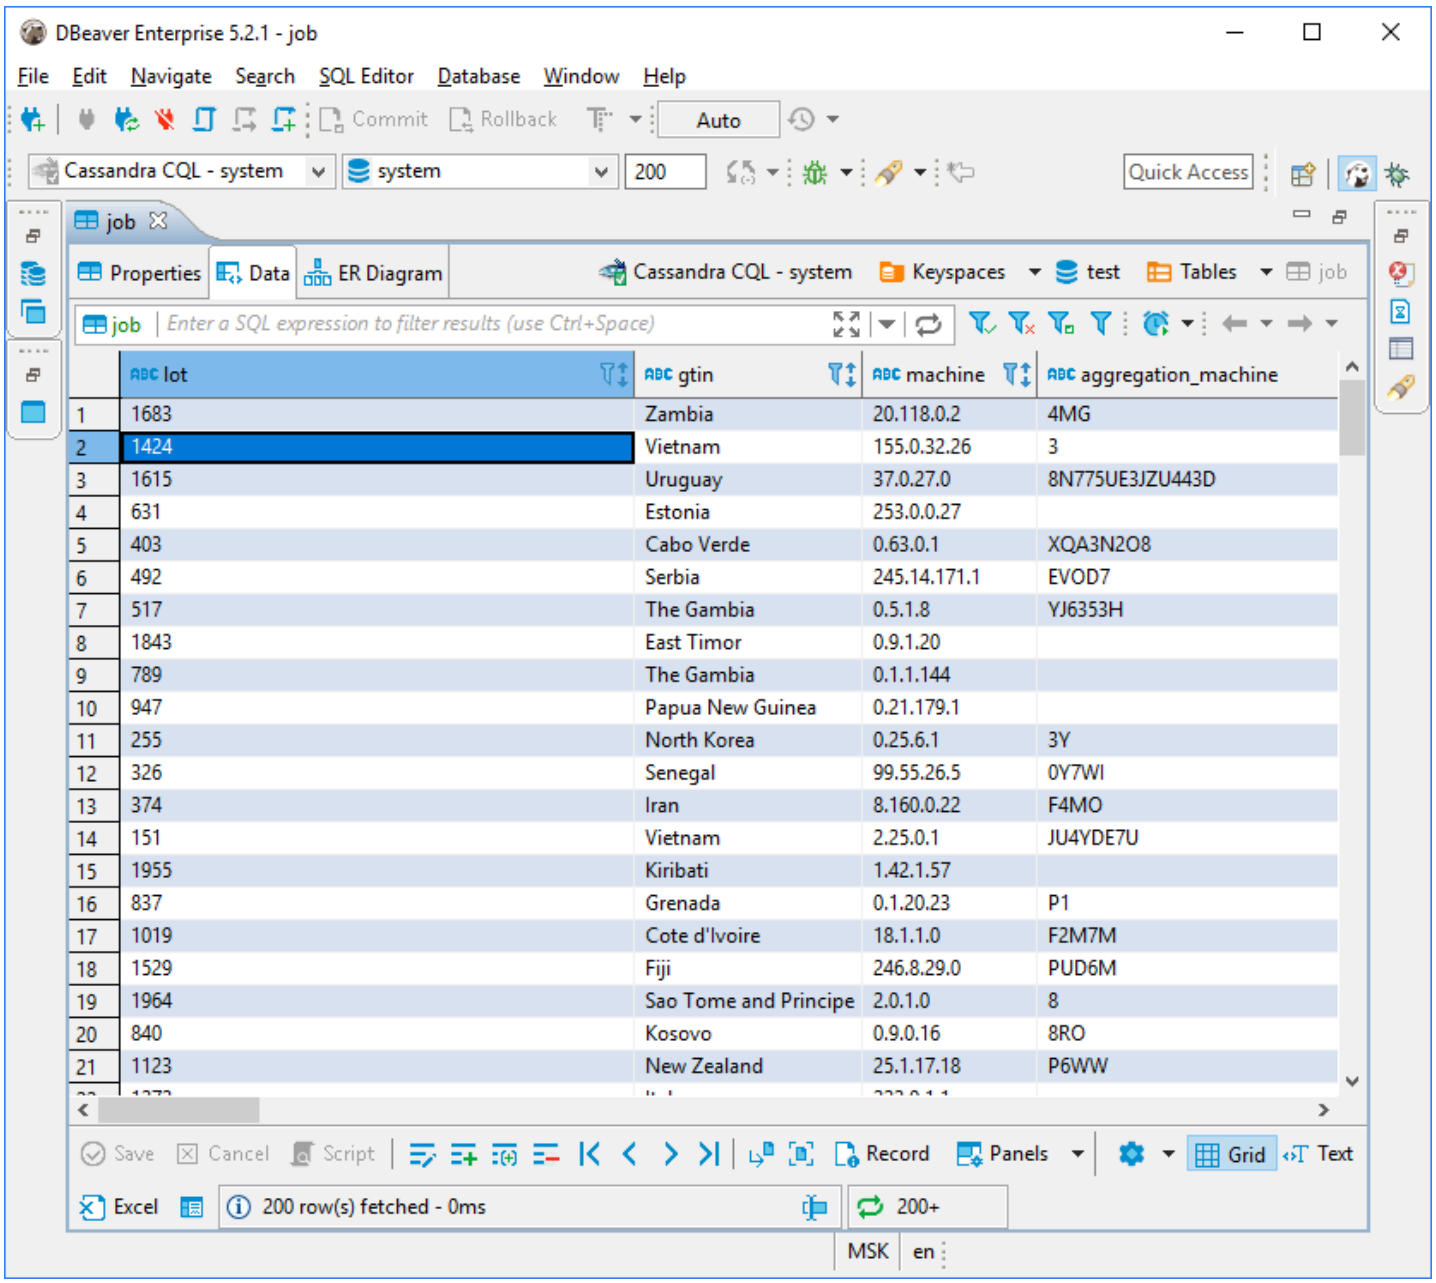

Executing CQL

CQL Cassandra Query Language is a kind of very simple SQL language dialect.
It supports simple SELECT queries, DDL statements (like CREATE TABLE) and some other.

You can use standard DBeaver SQL editor to execute CQL queries. DBeaver supports Cassandra query execution, results scrolling, data export/import, mock data generation and other features. Data viewer (of individual tables or custom CQL query results) query tracing is supported.

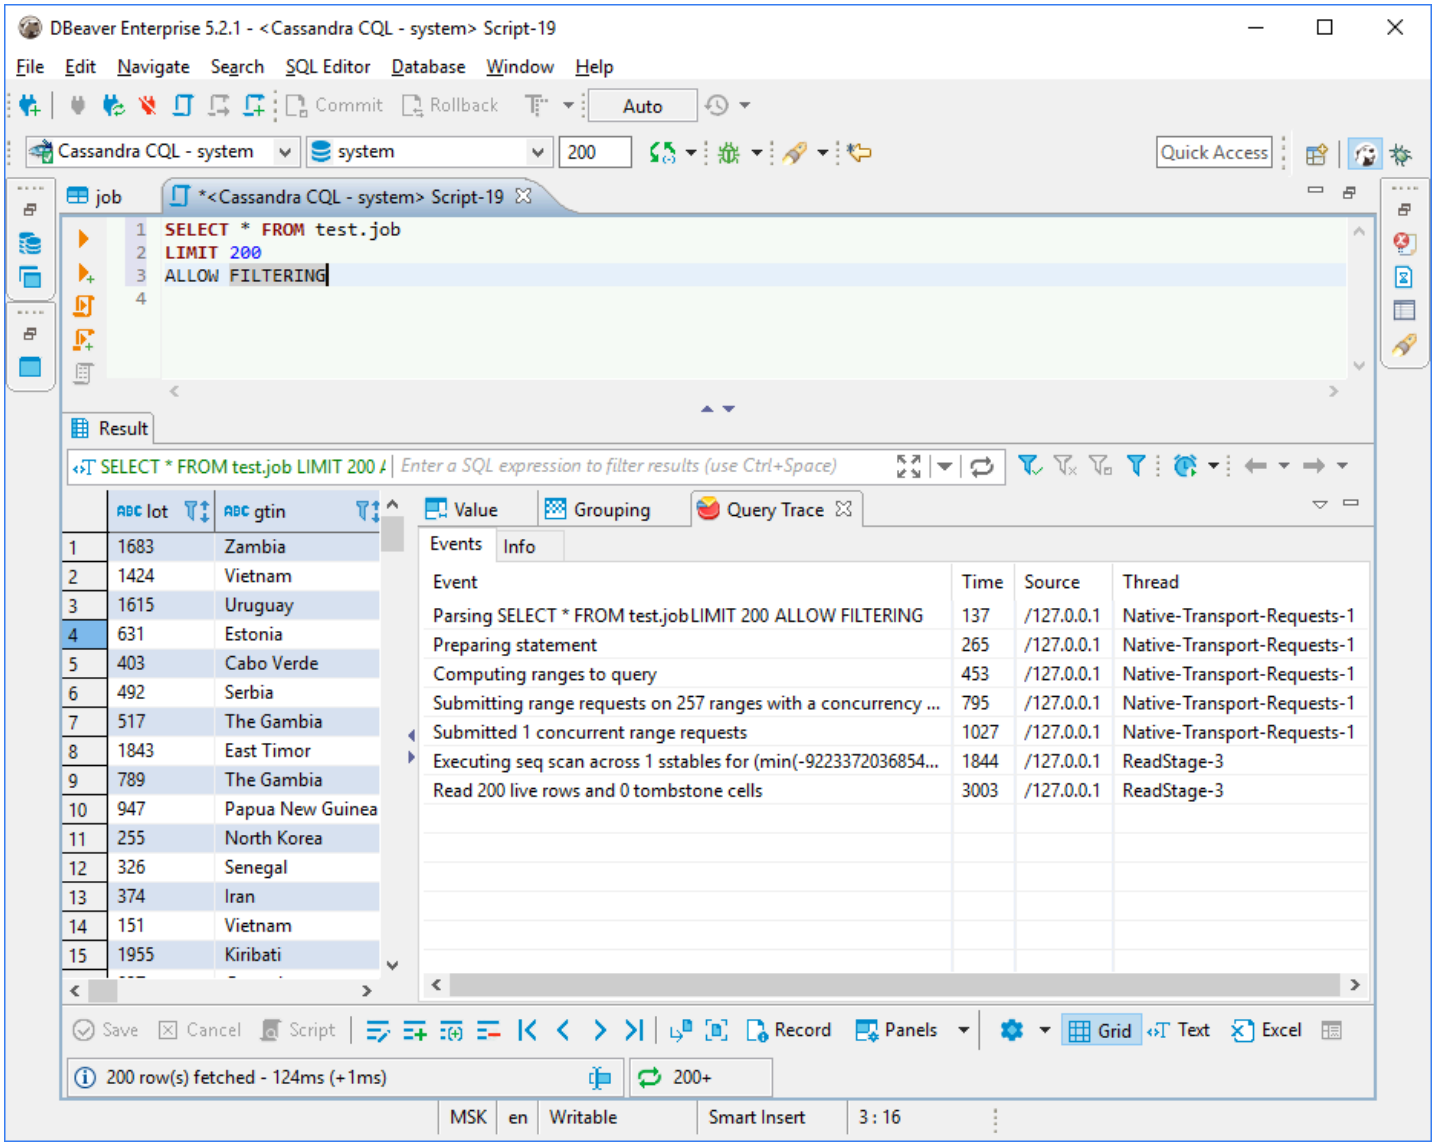

ERD

Physical ERD (Entity Relation Diagram) doesn't make much sense for Cassandra as there are no any foreign keys. However you can make you own custom ERD and connect actual Cassandra table with each other using logical associations.

InfluxDB

Overview

DBeaver EE supports InfluxDB schema browser, data viewer and InfluxQL queries execution.
DBeaver uses InfluxDB Java driver 2.12 to operate with the server over HTTP/HTTPS (standard InfluxDB protocol).
It supports InflixDB servers of any version (in the moment of writing).

Connecting to Influx Server

You can connect directly to a server or use SSH tunneling or SOCKS proxy.

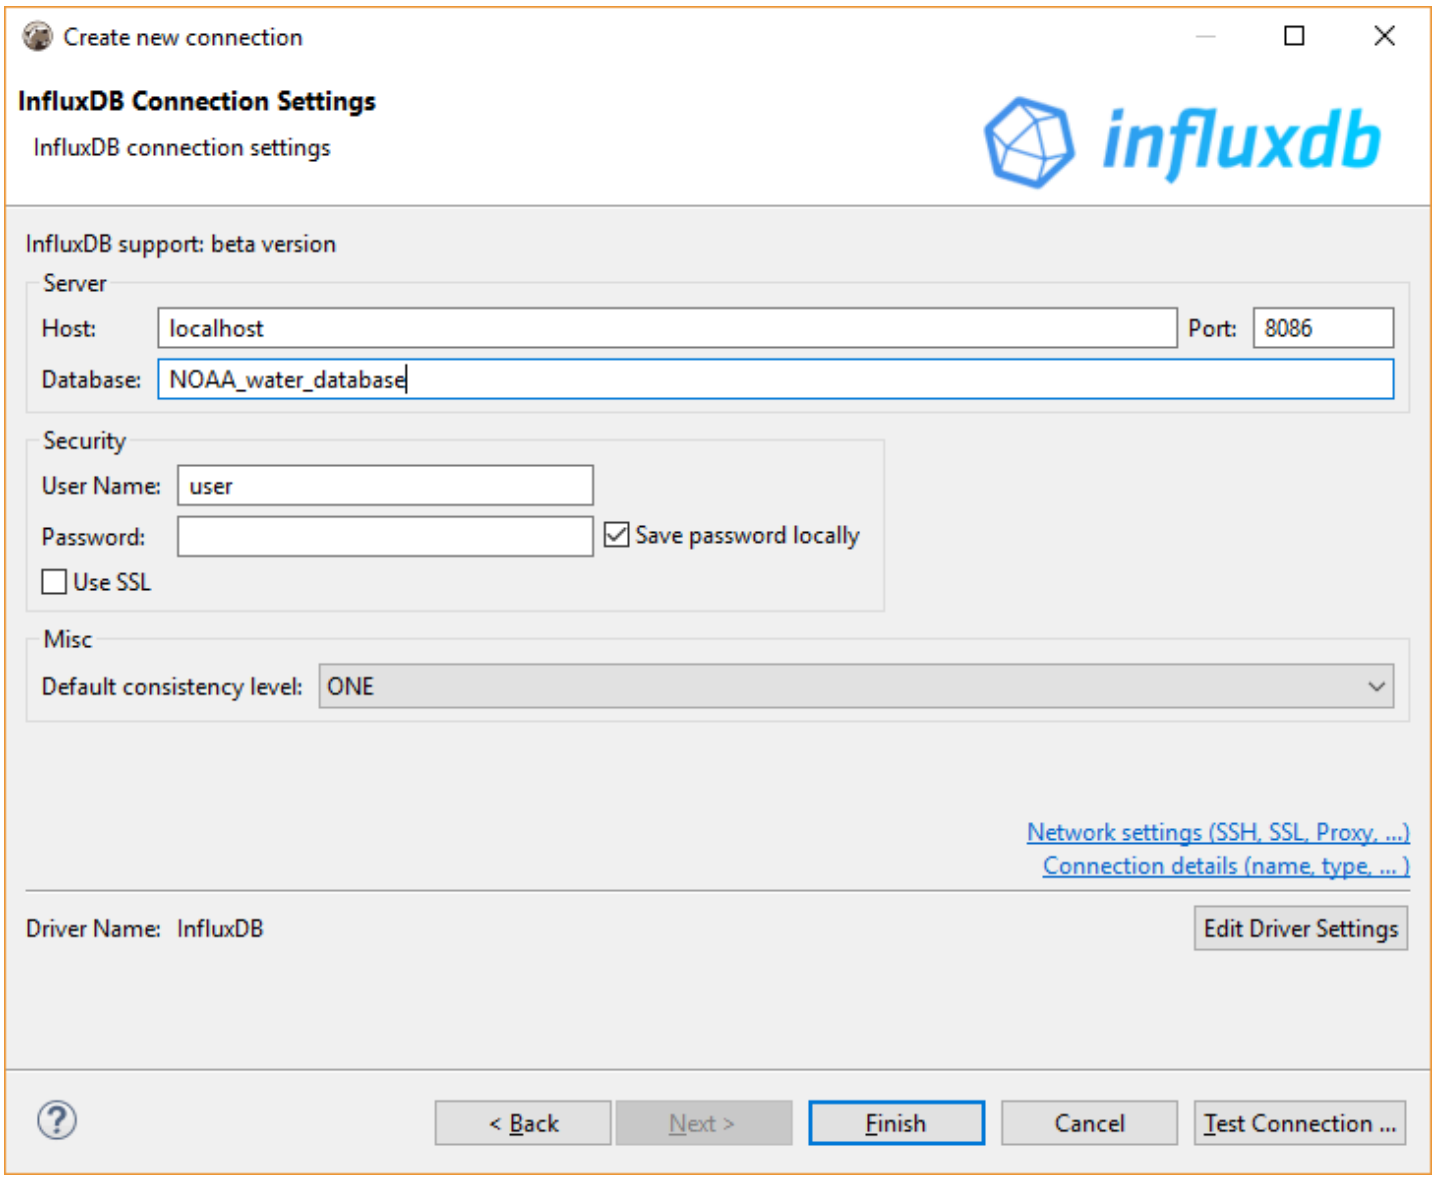

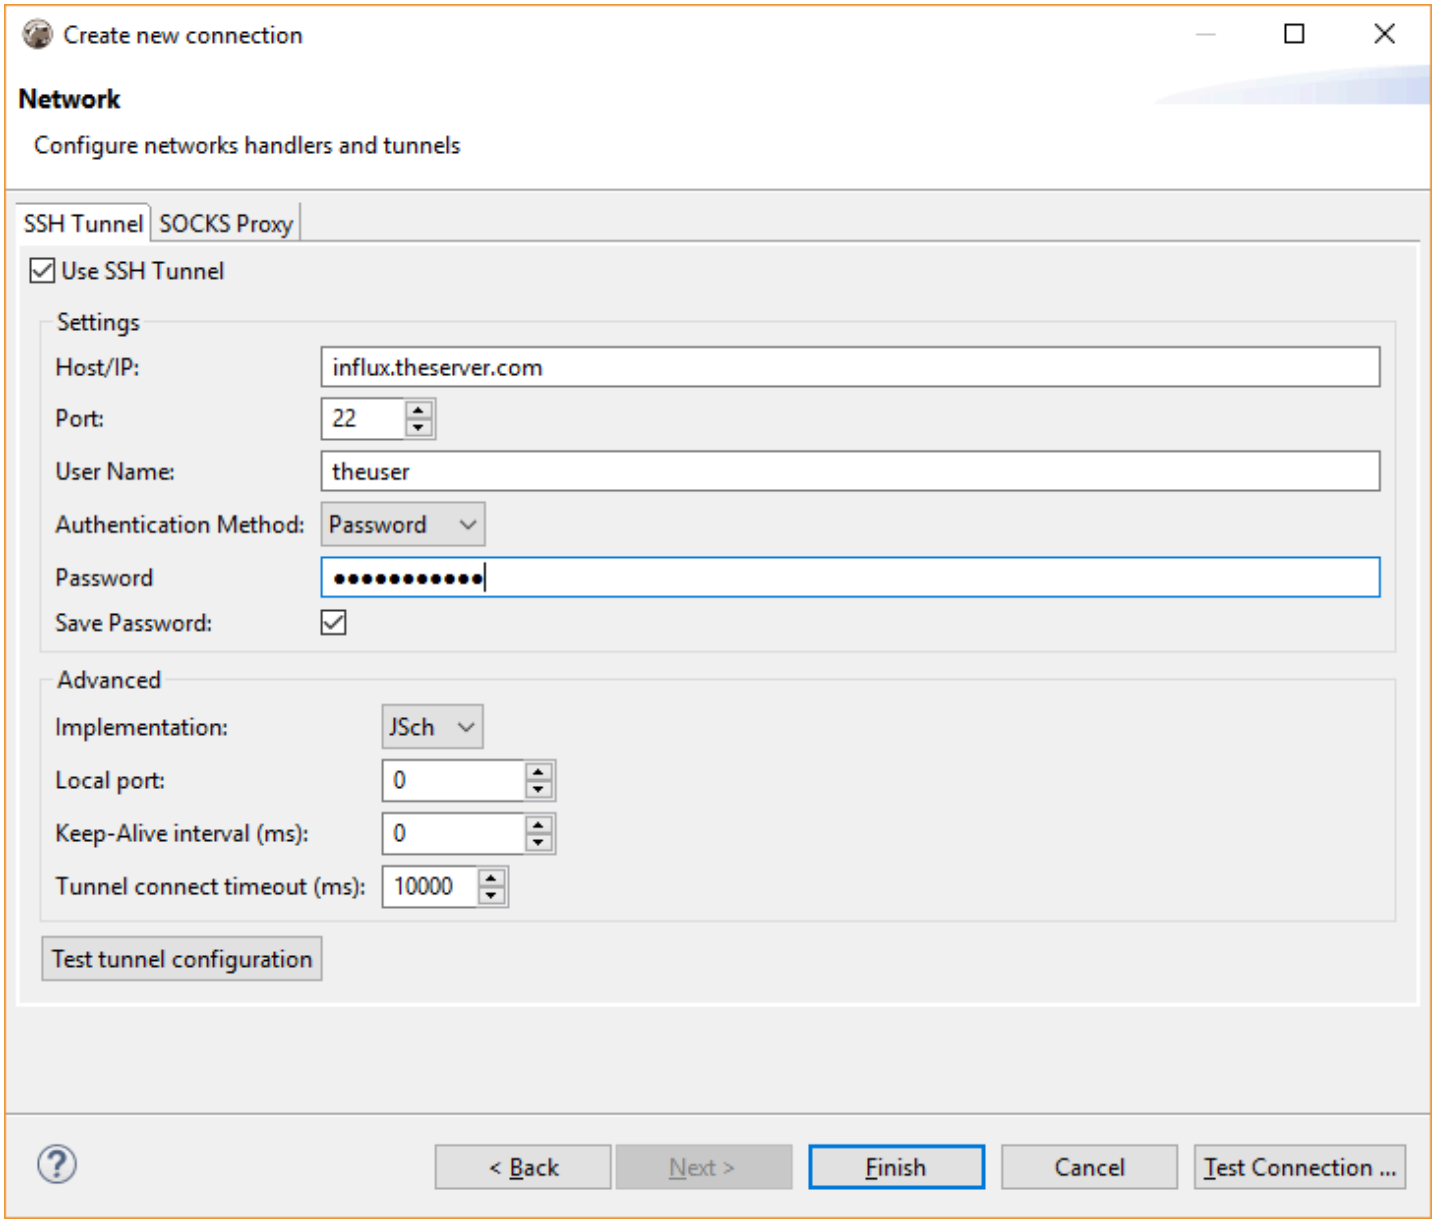

Browsing InfluxDB schema

InfluxDB is TimeSeries database, it doesn't support tables, foreign keys and other relational entities.

DBeaver doesn't support data insert/update in InfluxDB. Basically for DBeaver database is in read-only state. You can browse schema and view/analyse data.
While data itself is loaded by various sensors/data collectors in real time.
Instead of tables InfluxDB has measurements. Instead of columns it has fields and tags.

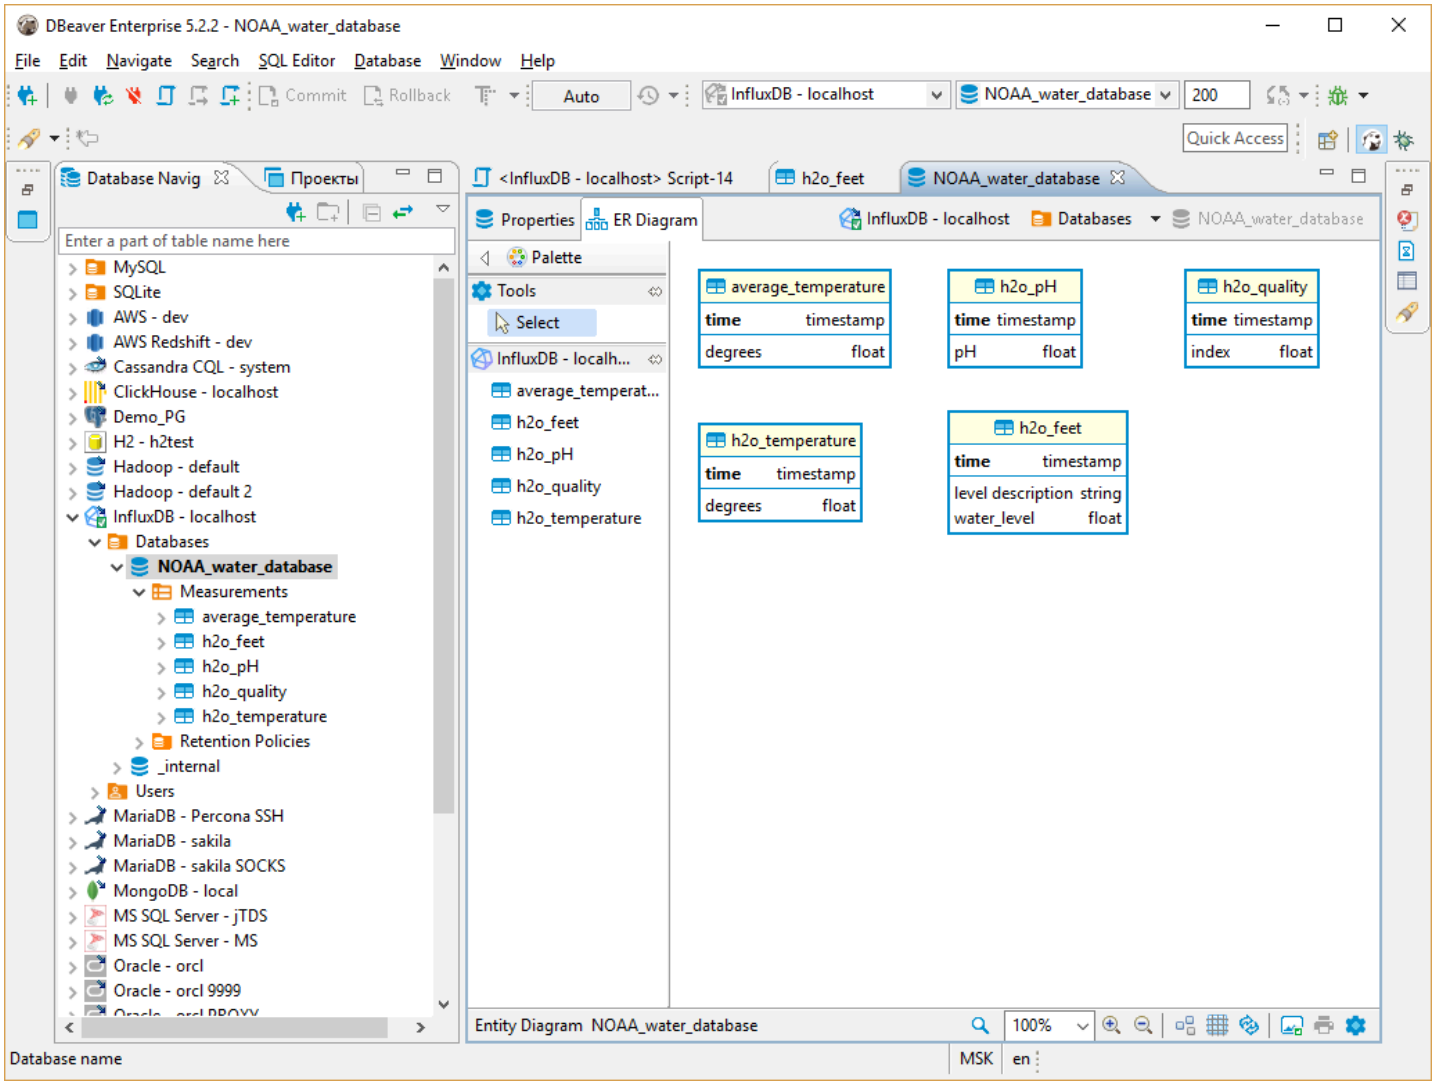

Executing InfluxQL

InfluxQL is a query language similar to SQL.
DBeaver fully supports all InfluxQL statements. Query results are presented as grid or as graphs:

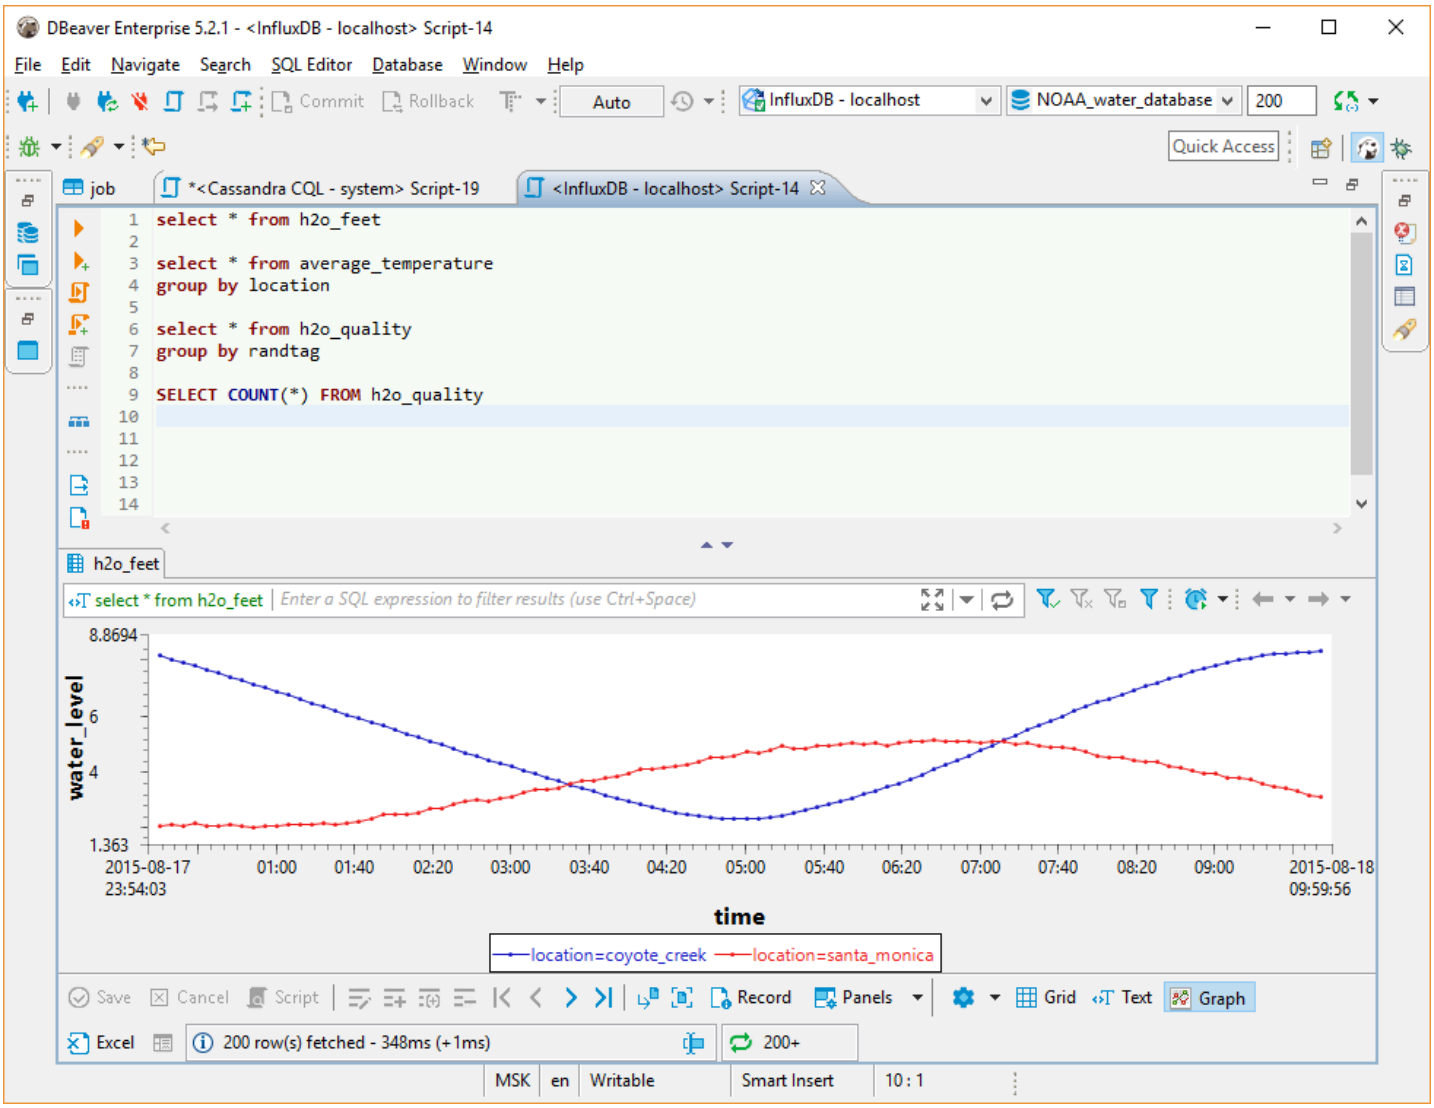
DBeaver Enterprise 5.2.1 - <InfluxDB - localhost> Script-14
File Edit Navigate Search SQL Editor Database Window Help
Commit Rollback Auto InfluxDB - localhost NOAA_water_database 200
Quick Access
job
*<Cassandra CQL - system> Script-19
<InfluxDB - localhost> Script-14
1 select * from h2o_feet
2
3 select * from average_temperature
4 group by location
5
6 select * from h2o_quality
7 group by randtag
8
9 SELECT COUNT(*) FROM h2o_quality
10
11
12
13
14
h2o_feet
select * from h2o_feet
Enter a SQL expression to filter results (use Ctrl+Space)
8.8694
6
4
1.363
water_level
2015-08-17 23:54:03
01:00
01:40
02:20
03:00
03:40
04:20
05:00
05:40
06:20
07:00
07:40
08:20
09:00
2015-08-18 09:59:56
time
location=coyote_creek
location=santa_monica
Save Cancel Script
Record
Panels
Grid
Text
Graph
Excel
200 row(s) fetched - 348ms (+1ms)
200+
MSK
en
Writable
Smart Insert
10 : 1

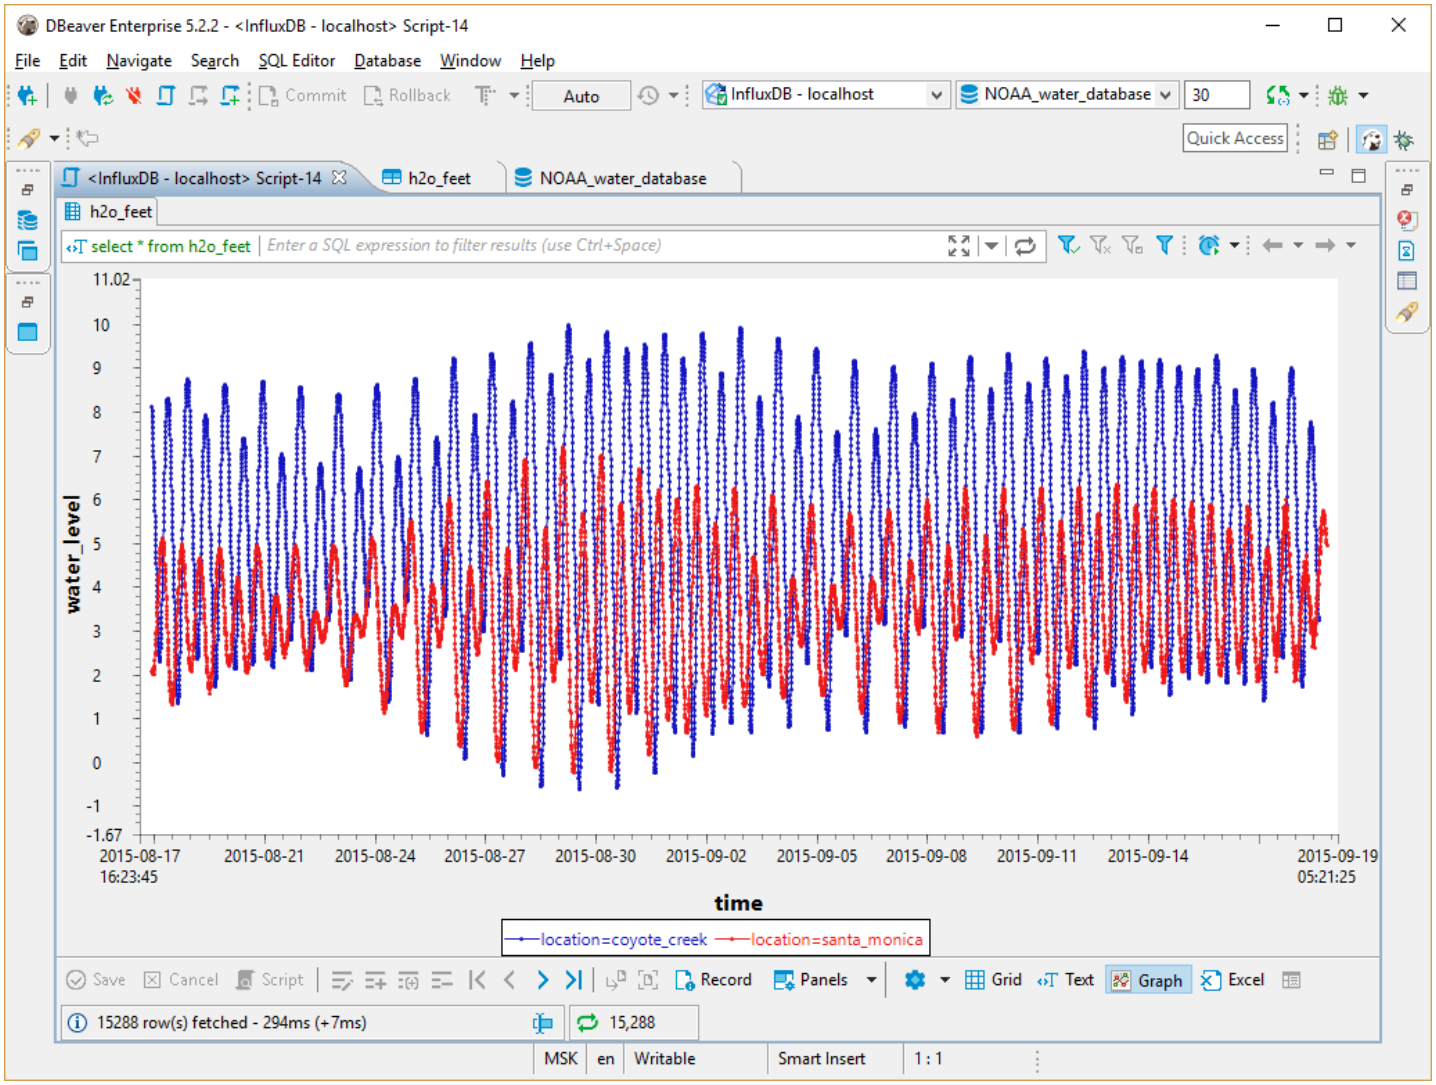
DBeaver Enterprise 5.2.2 - <InfluxDB - localhost> Script-14
File Edit Navigate Search SQL Editor Database Window Help
Commit Rollback Auto InfluxDB - localhost NOAA_water_database 30
Quick Access
<InfluxDB - localhost> Script-14
h2o_feet
NOAA_water_database
select * from h2o_feet
Enter a SQL expression to filter results (use Ctrl+Space)
11.02
10
9
8
7
6
5
4
3
2
1
0
-1
-1.67
water_level
2015-08-17 16:23:45
2015-08-21
2015-08-24
2015-08-27
2015-08-30
2015-09-02
2015-09-05
2015-09-08
2015-09-11
2015-09-14
2015-09-19 05:21:25
time
location=coyote_creek
location=santa_monica
Save Cancel Script Record Panels Grid Text Graph Excel
15288 row(s) fetched - 294ms (+7ms)
15,288
MSK en Writable Smart Insert 1 : 1

Redis

Overview

DBeaver EE supports Redis key browser, key value viewer and Redis commands shell.
DBeaver uses Jedis driver 2.9.0 to operate with Redis server. It supports Redis servers of any version.

Connecting to Redis Server

You can connect directly to a server or use SSH tunneling or SOCKS proxy.

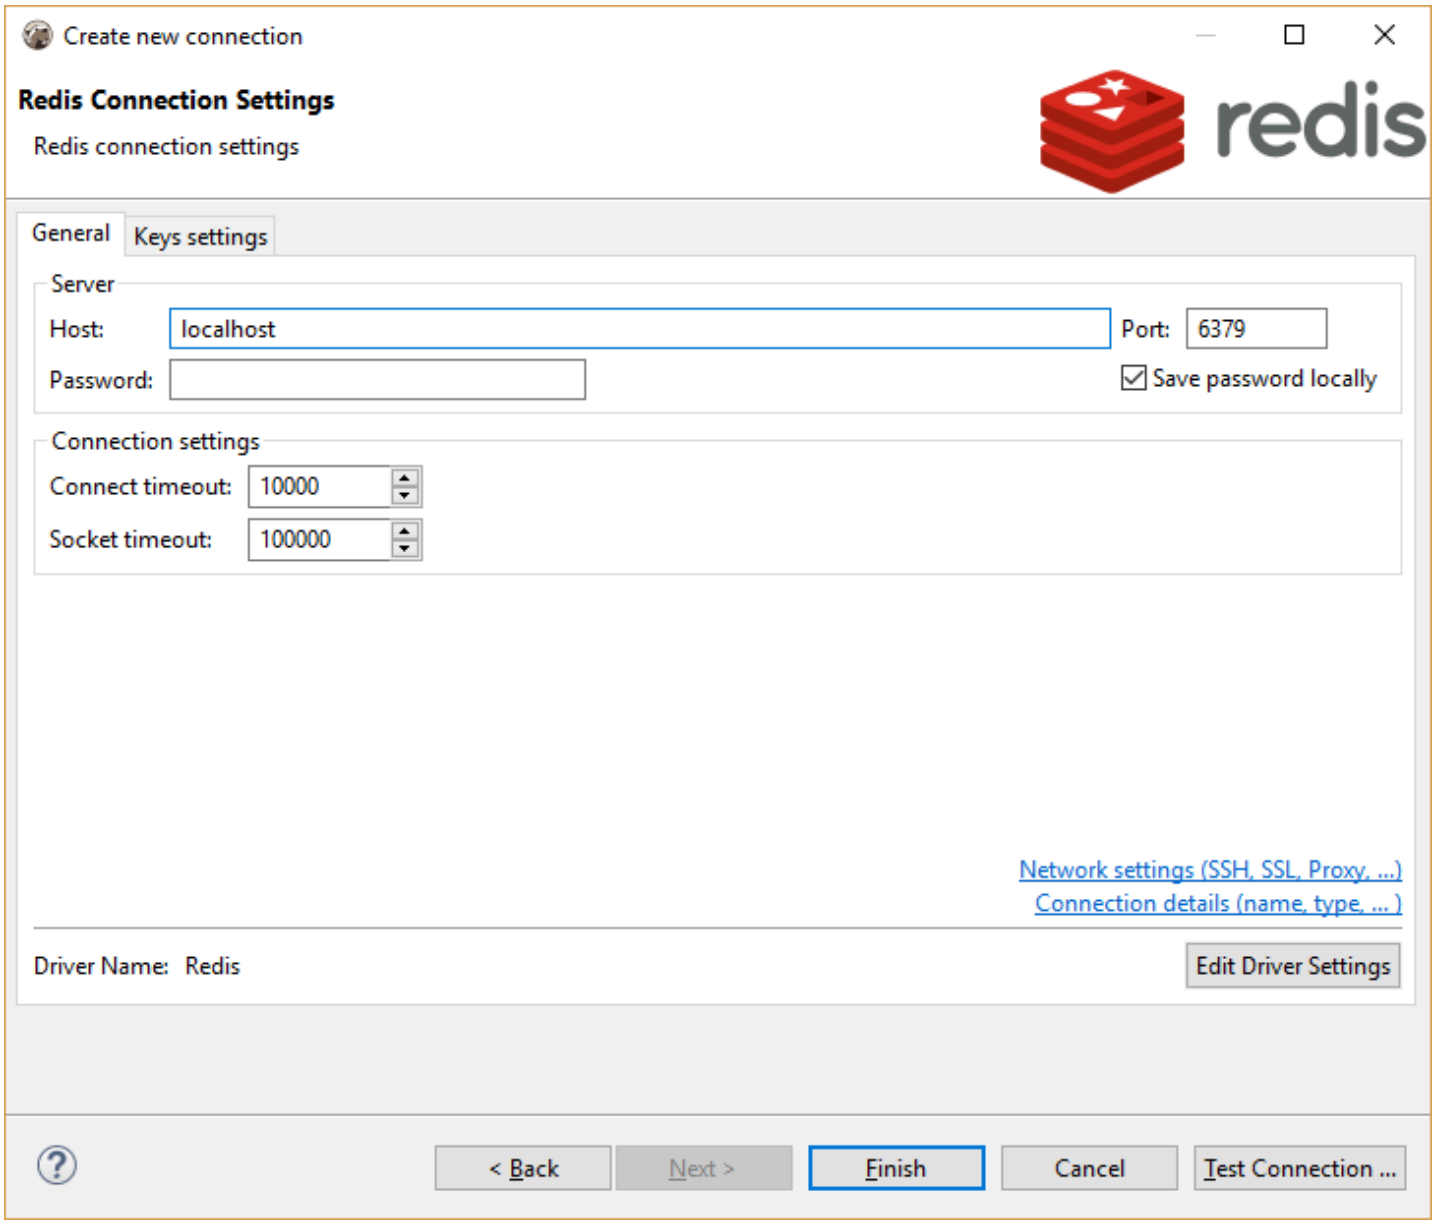
Create new connection
Redis Connection Settings
Redis connection settings
redis
General
Keys settings
Server
Host:
localhost
Port:
6379
Password:
Save password locally
Connection settings
Connect timeout:
10000
Socket timeout:
100000
Network settings (SSH, SSL, Proxy, ...)
Connection details (name, type, ...)
Driver Name: Redis
Edit Driver Settings
< Back
Next >
Finish
Cancel
Test Connection ...

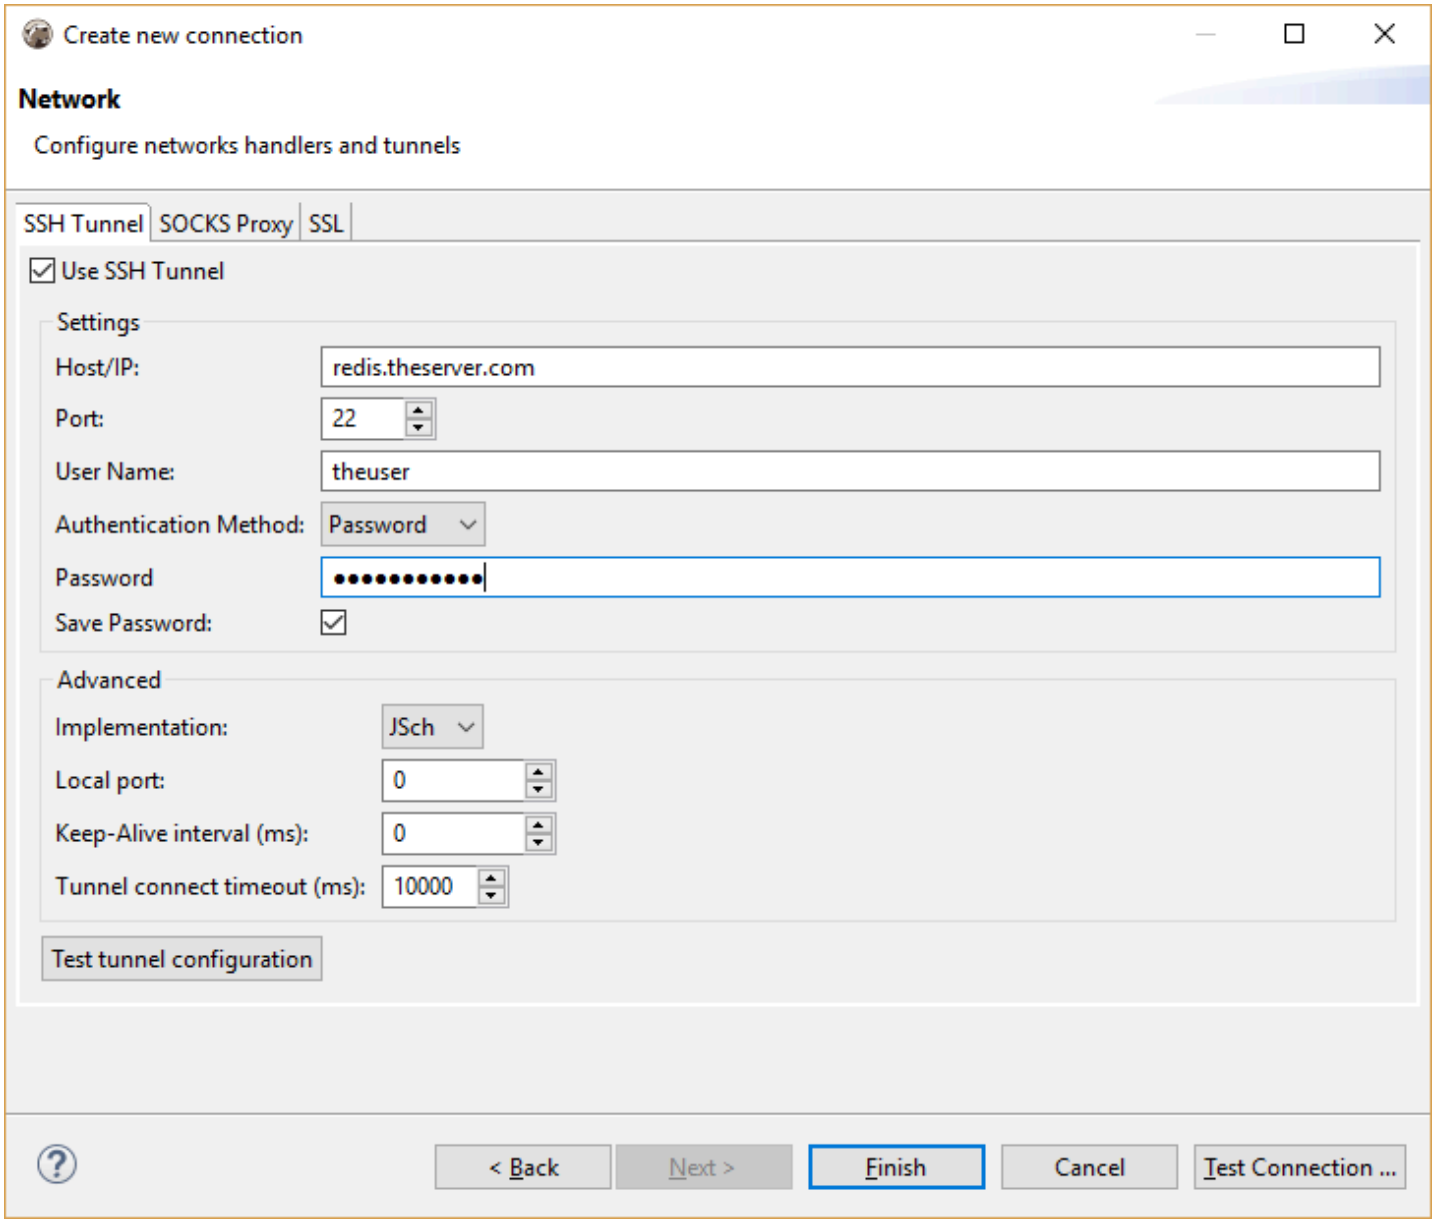
Create new connection
Network
Configure networks handlers and tunnels
SSH Tunnel
SOCKS Proxy
SSL
Use SSH Tunnel
Settings
Host/IP:
redis.theserver.com
Port:
22
User Name:
theuser
Authentication Method:
Password
Password
Save Password:
Advanced
Implementation:
JSch
Local port:
0
Keep-Alive interval (ms):
0
Tunnel connect timeout (ms):
10000
Test tunnel configuration
< Back
Next >
Finish
Cancel
Test Connection ...

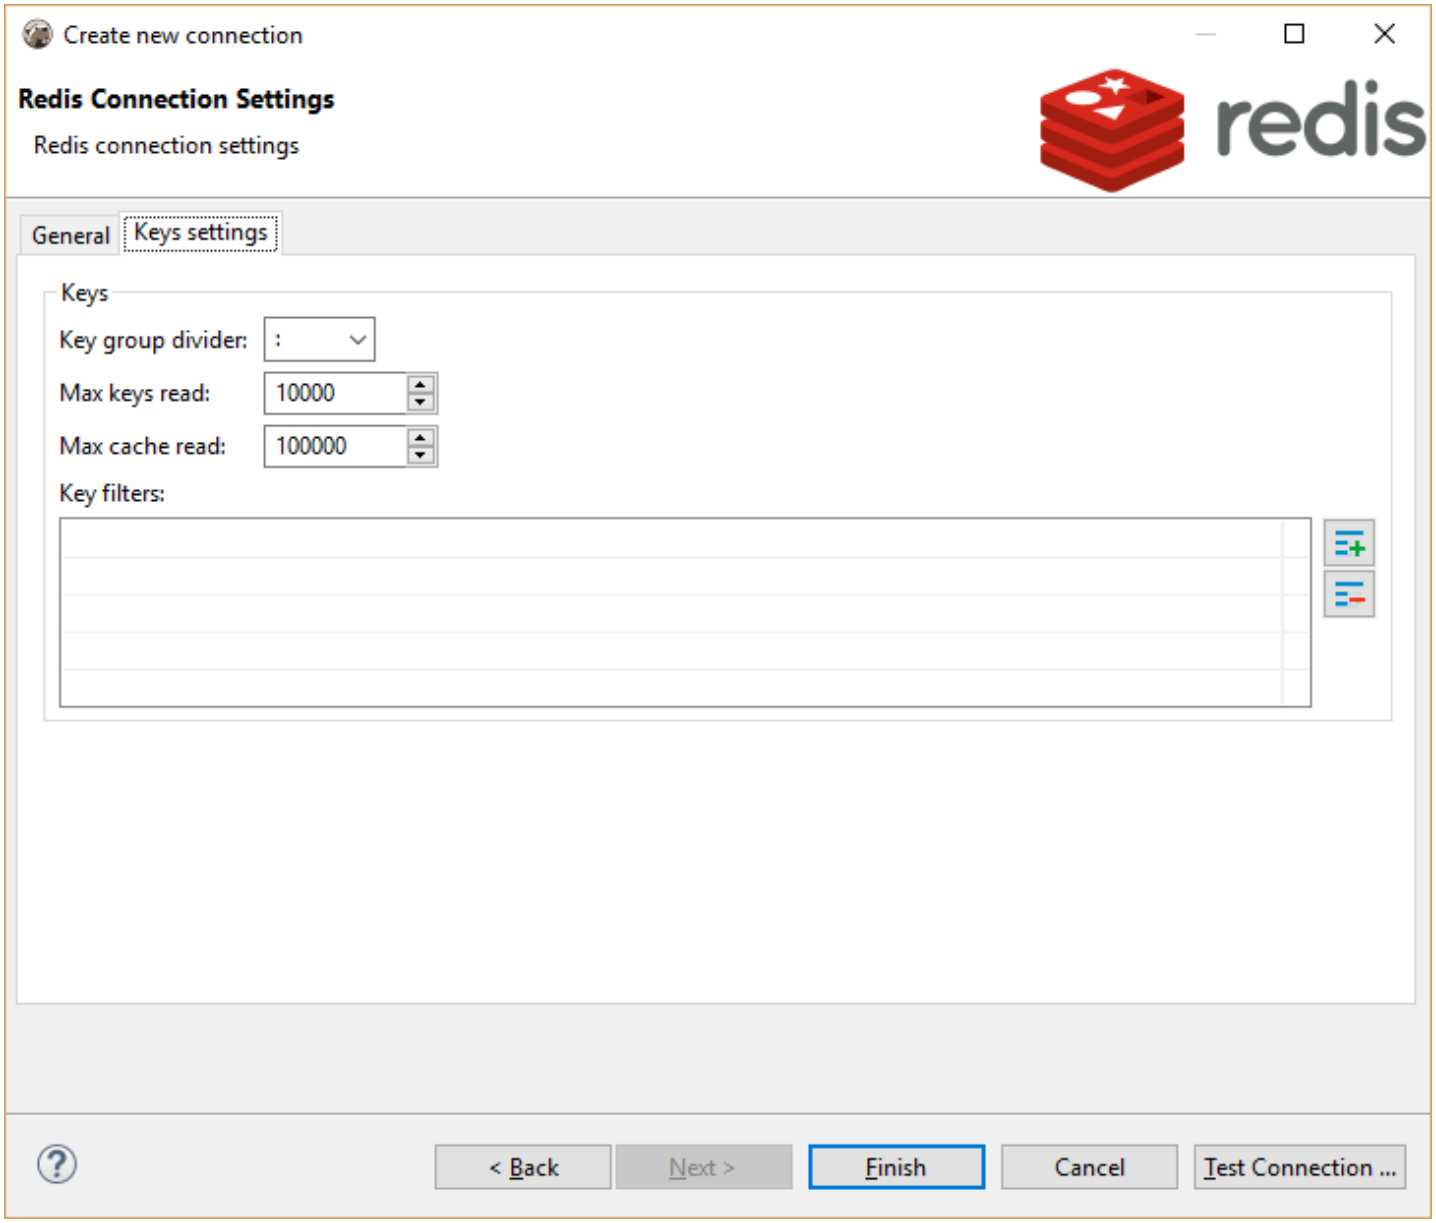

Browsing Redis keys

You can view/edit Redis keys a plain list. However Redis database usually contains a lot of keys (millions or even billions) and using list presentation is not convenient (or is not possible).
DBeaver supports hierarchy presentation of keys. Internally Redis doesn't support hierarchies but on application level key names may be divided on groups using some character (e.g. coma, dash or colon). DBeaver uses this pattern to show hierarchy. Group separator can be configured in connection properties.

Key browser may be convenient in some cases but in case of big databases it is very uneasy to find your key in navigator and SQL editor should be used instead. Redis commands are the most flexible way to operate with keys.

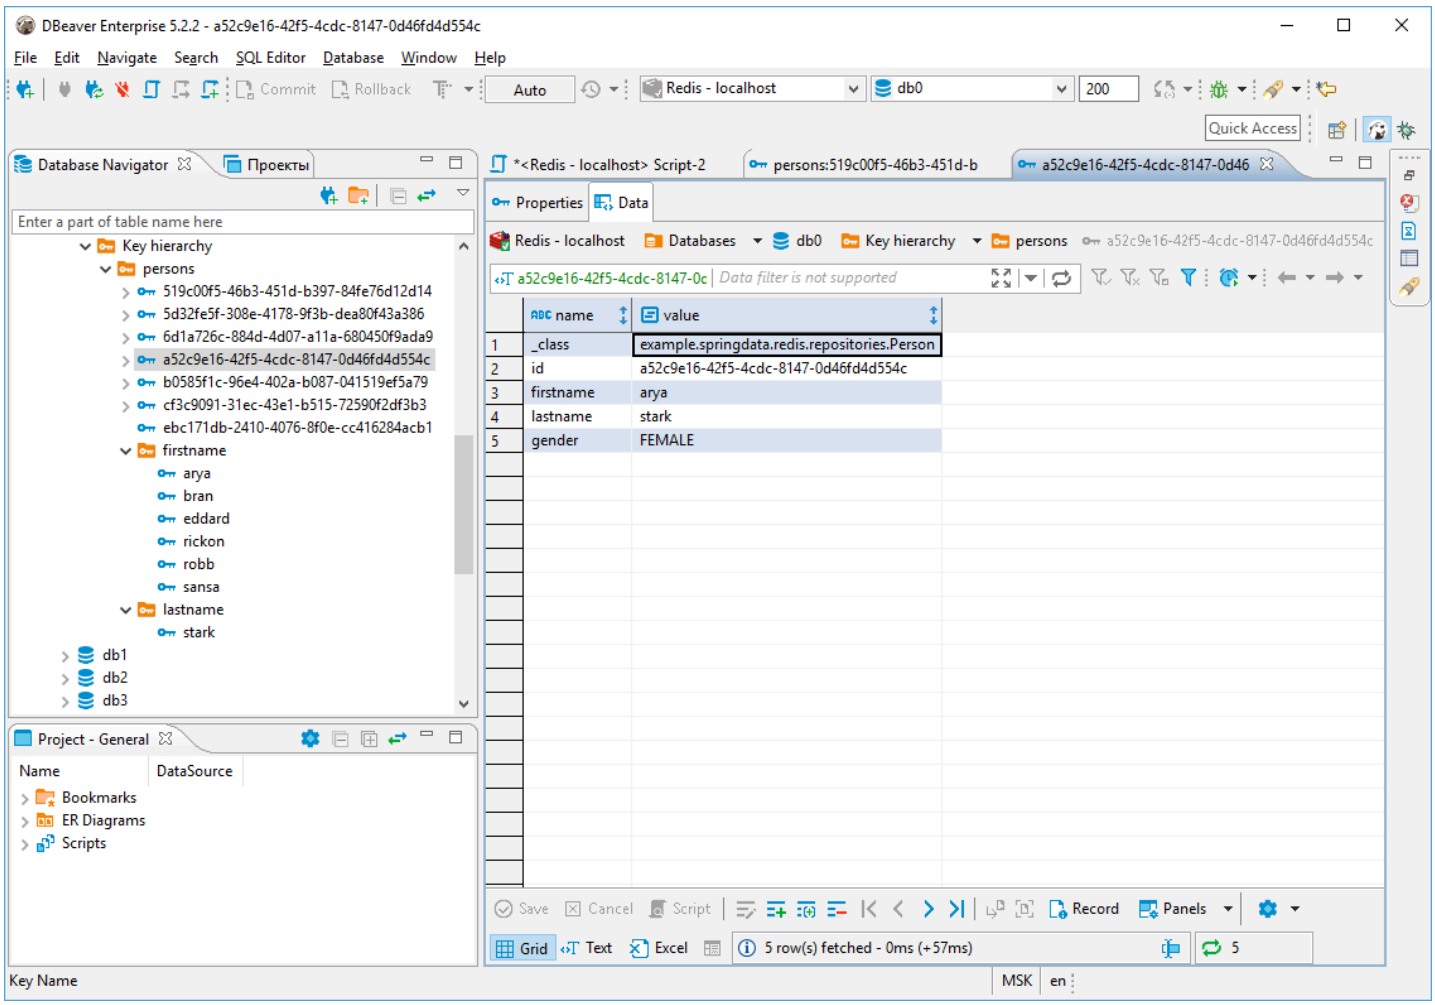
DBeaver Enterprise 5.2.2 - a52c9e16-42f5-4cdc-8147-0d46fd4d554c
File Edit Navigate Search SQL Editor Database Window Help
Commit Rollback Auto Redis - localhost db0 200
Quick Access
Database Navigator Проекты
Enter a part of table name here
Key hierarchy
persons
519c00f5-46b3-451d-b397-84fe76d12d14
5d32fe5f-308e-4178-9f3b-dea80f43a386
6d1a726c-884d-4d07-a11a-680450f9ada9
a52c9e16-42f5-4cdc-8147-0d46fd4d554c
b0585f1c-96e4-402a-b087-041519ef5a79
cf3c9091-31ec-43e1-b515-72590f2df3b3
ebc171db-2410-4076-8f0e-cc416284acb1
firstname
arya
bran
eddard
rickon
robb
sansa
lastname
stark
db1
db2
db3
Project - General
Name DataSource
Bookmarks
ER Diagrams
Scripts
*<Redis - localhost> Script-2
persons:519c00f5-46b3-451d-b
a52c9e16-42f5-4cdc-8147-0d46
Properties Data
Redis - localhost Databases db0 Key hierarchy persons a52c9e16-42f5-4cdc-8147-0d46fd4d554c
a52c9e16-42f5-4cdc-8147-0c Data filter is not supported
name value
1 _class example.springdata.redis.repositories.Person
2 id a52c9e16-42f5-4cdc-8147-0d46fd4d554c
3 firstname arya
4 lastname stark
5 gender FEMALE
Save Cancel Script Record Panels
Grid Text Excel 5 row(s) fetched - 0ms (+57ms) 5
Key Name
MSK en

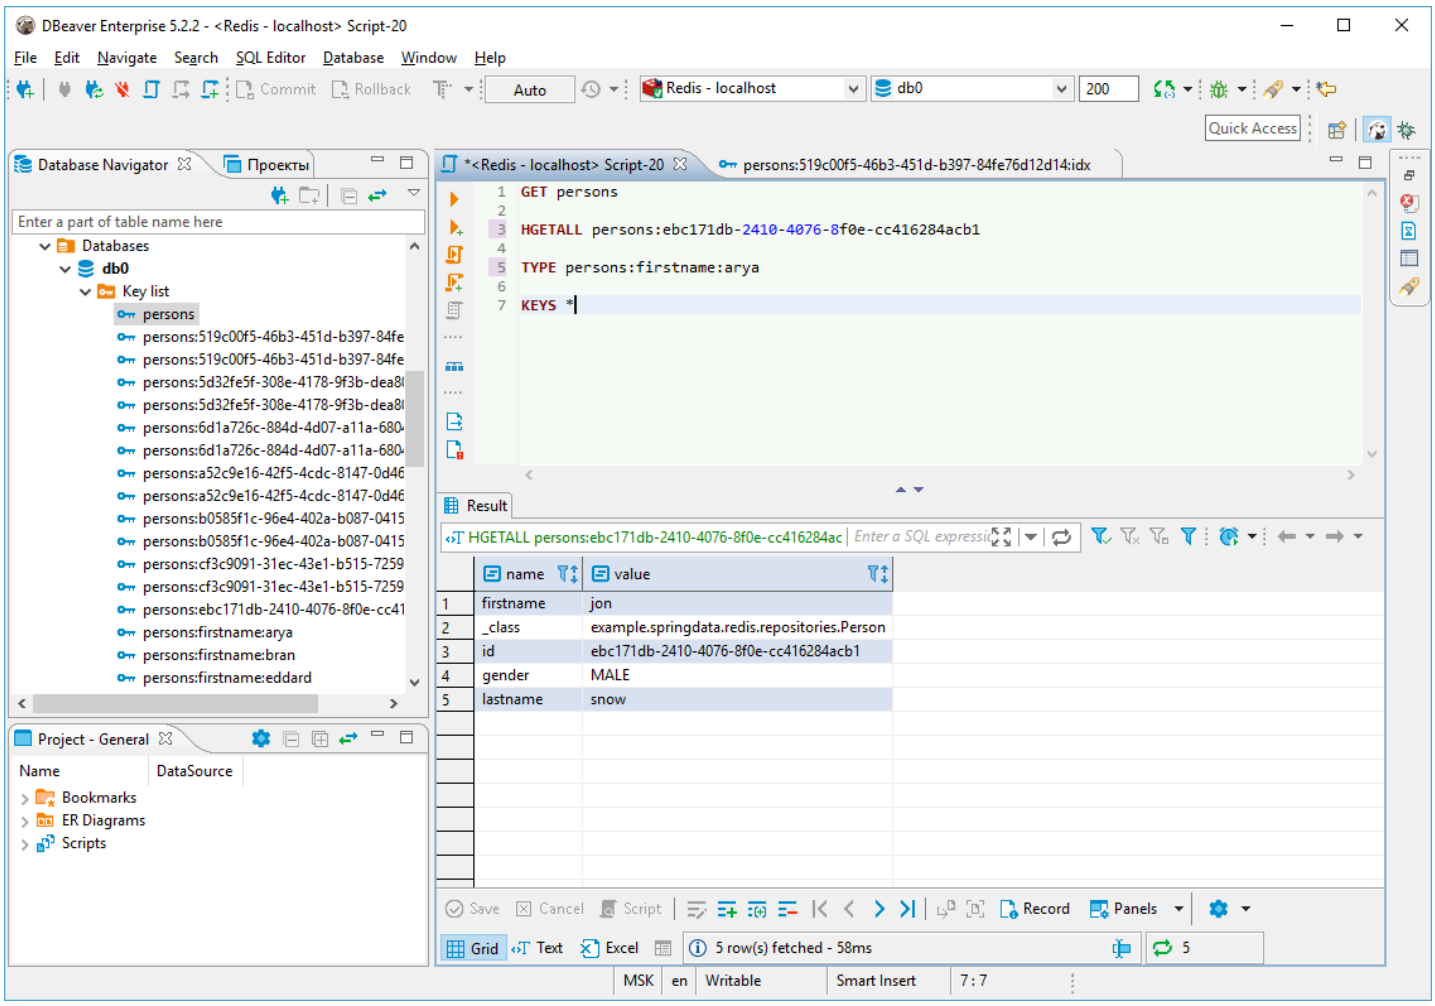
DBeaver Enterprise 5.2.2 - <Redis - localhost> Script-20
File Edit Navigate Search SQL Editor Database Window Help
Commit Rollback Auto Redis - localhost db0 200
Quick Access
Database Navigator Проекты
Enter a part of table name here
Databases
db0
Key list
persons
persons:519c00f5-46b3-451d-b397-84fe
persons:519c00f5-46b3-451d-b397-84fe
persons:5d32fe5f-308e-4178-9f3b-dea8
persons:5d32fe5f-308e-4178-9f3b-dea8
persons:6d1a726c-884d-4d07-a11a-680
persons:6d1a726c-884d-4d07-a11a-680
persons:a52c9e16-42f5-4cdc-8147-0d46
persons:a52c9e16-42f5-4cdc-8147-0d46
persons:b0585f1c-96e4-402a-b087-0415
persons:b0585f1c-96e4-402a-b087-0415
persons:cf3c9091-31ec-43e1-b515-7259
persons:cf3c9091-31ec-43e1-b515-7259
persons:ebc171db-2410-4076-8f0e-cc41
persons:firstname:arya
persons:firstname:bran
persons:firstname:eddard
Project - General
Name DataSource
Bookmarks
ER Diagrams
Scripts
*<Redis - localhost> Script-20
persons:519c00f5-46b3-451d-b397-84fe76d12d14:idx
1 GET persons
2
3 HGETALL persons:ebc171db-2410-4076-8f0e-cc416284acb1
4
5 TYPE persons:firstname:arya
6
7 KEYS *
Result
HGETALL persons:ebc171db-2410-4076-8f0e-cc416284ac Enter a SQL expressi
name value
1 firstname jon
2 _class example.springdata.redis.repositories.Person
3 id ebc171db-2410-4076-8f0e-cc416284acb1
4 gender MALE
5 lastname snow
Save Cancel Script Record Panels
Grid Text Excel 5 row(s) fetched - 58ms 5
MSK en Writable Smart Insert 7 : 7

Executing Redis commands

Redis doesn't support SQL or any other query language. Instead of that it supports build-in commands and LUA scripts.

Redis commands can be executed in the same way as in Redis comand line shell:

```
In order to execute a command run it using CTRL+Enter or ALT+X. All standard DBeaver SQL editor shortcuts

In order to execute LUA script surround it with {} brackets and run as single statement. If script contair
```lua
{
    return {1,2,{3,'Hello World!'}}
}
```
```

AWS DynamoDB

Supported features:

- Table data view
- Table data edit in document (json) mode
- Data filters
- SQL queries execution
- JSON queries execution
- Data export and import

DynamoDB connection

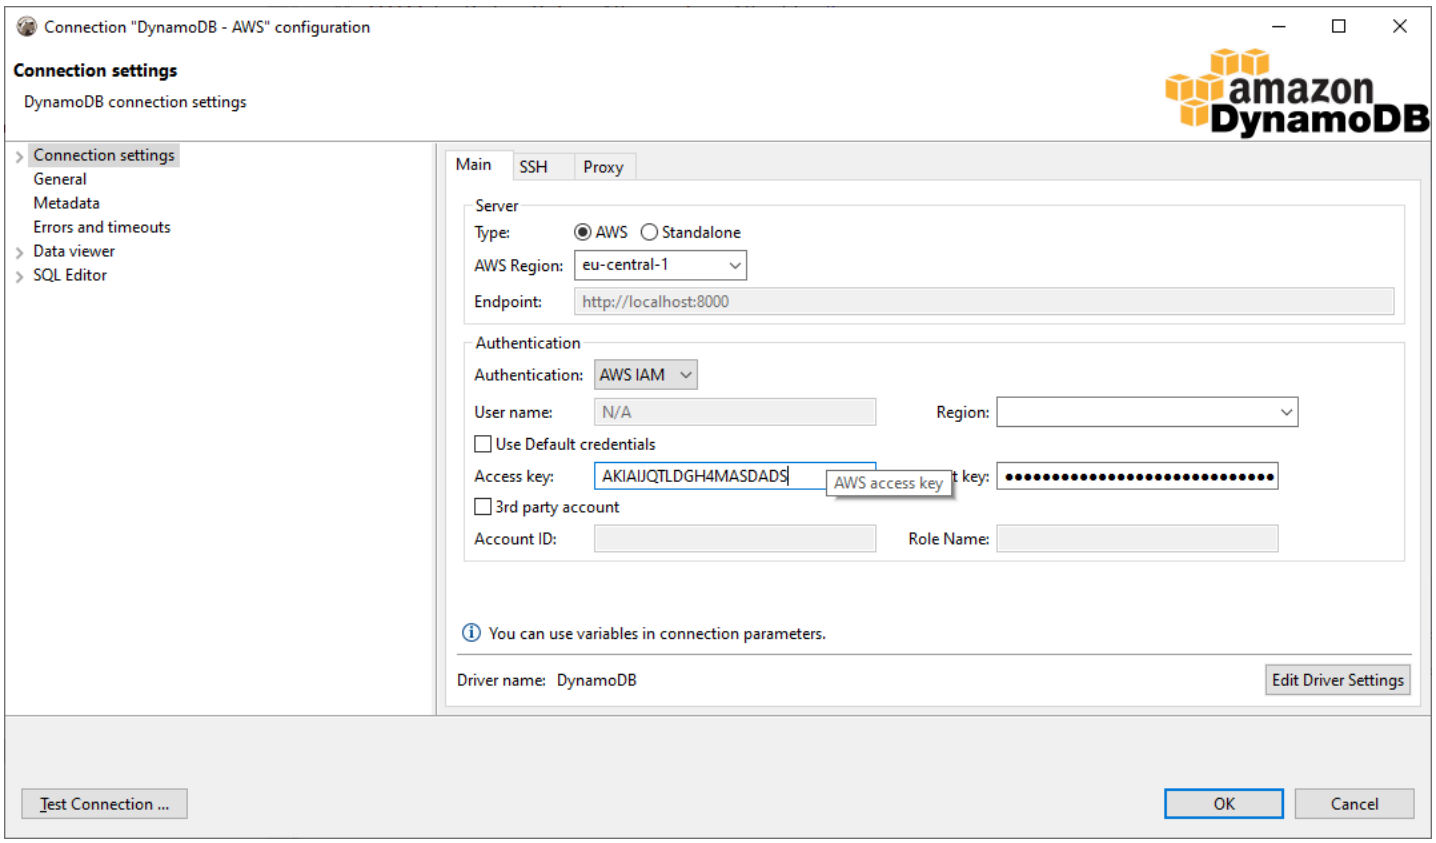

DBeaver supports AWS Cloud and Standalone versions of DynamoDB.
For standalone server you need to enter endpoint (http or https URL).
For cloud server you must enter AWS region. DynamoDB exists in all available regions in your AWS account but tables are different.

AWS Access Key and Secret Key are used for authentication.
For 3rd-party account access you must specify 3rd party account ID (12-digits number) and 3rd party role name. This role will be used for permission management. You account must be added to whitelist in the 3rd party account.

Press "Test Connection" to validate your connection settings.

Database navigation

DynamoDB has simple metadata structure. Basically you can access only Table and Global tables.
Table has primary attributes (a kind of primary key) and indexes.
DynamoDB is document-oriented database. Each table may have its own set of attributes and sub-attributes.

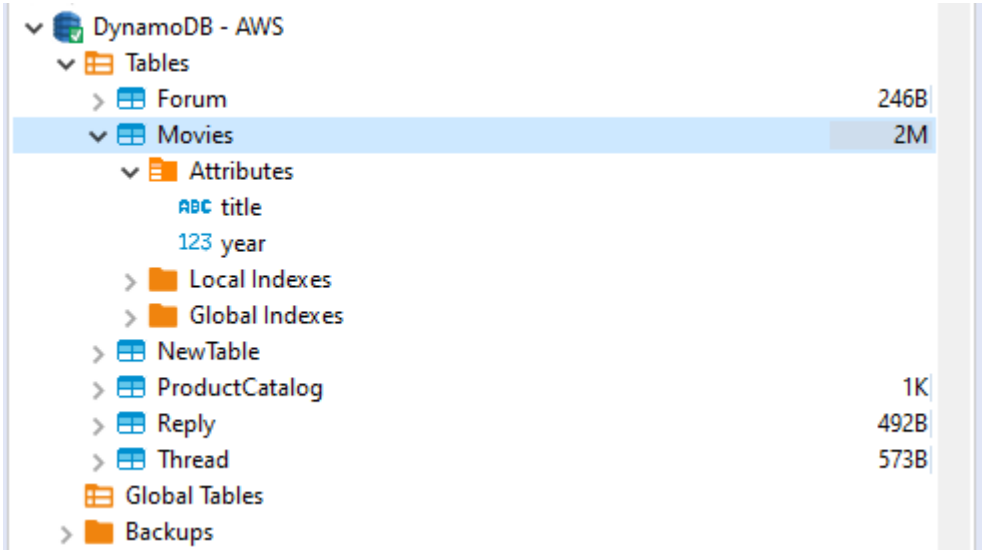

Viewing table data

You can open table editor and see table data.
You may need to switch to the "Data" tab. By default DBeaver converts DynamoDB documents into table format

but you can switch to another data representation.
You can use data filters in order to find documents.

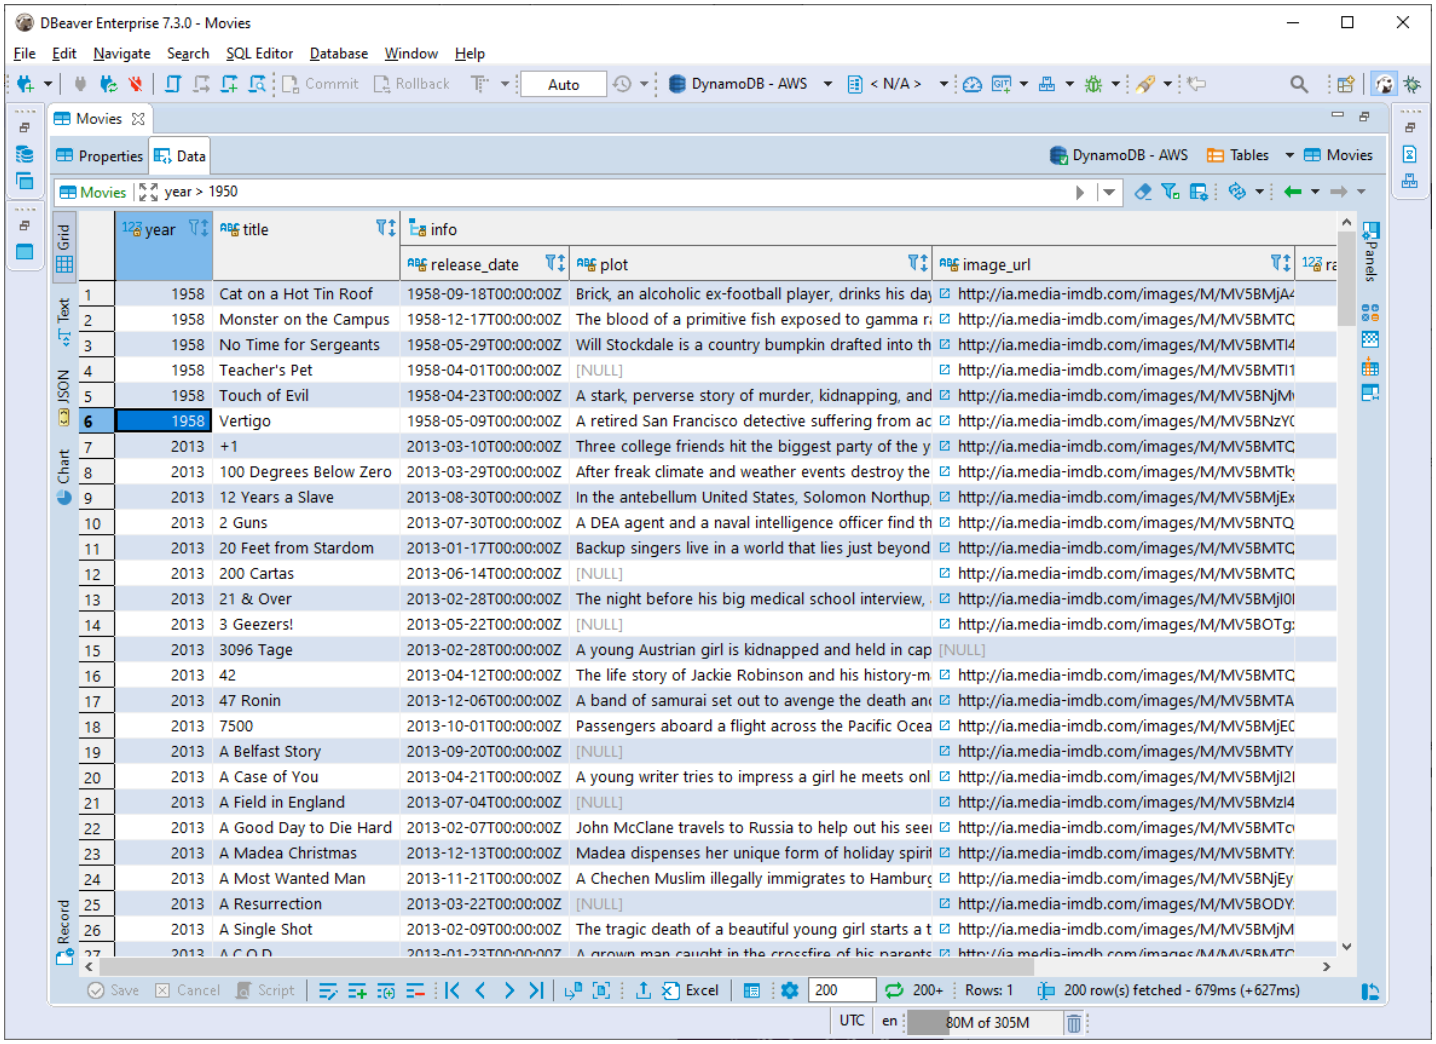

Viewing data in JSON document format

You view, search and edit JSON documents. Double-click on a document to activate editor.

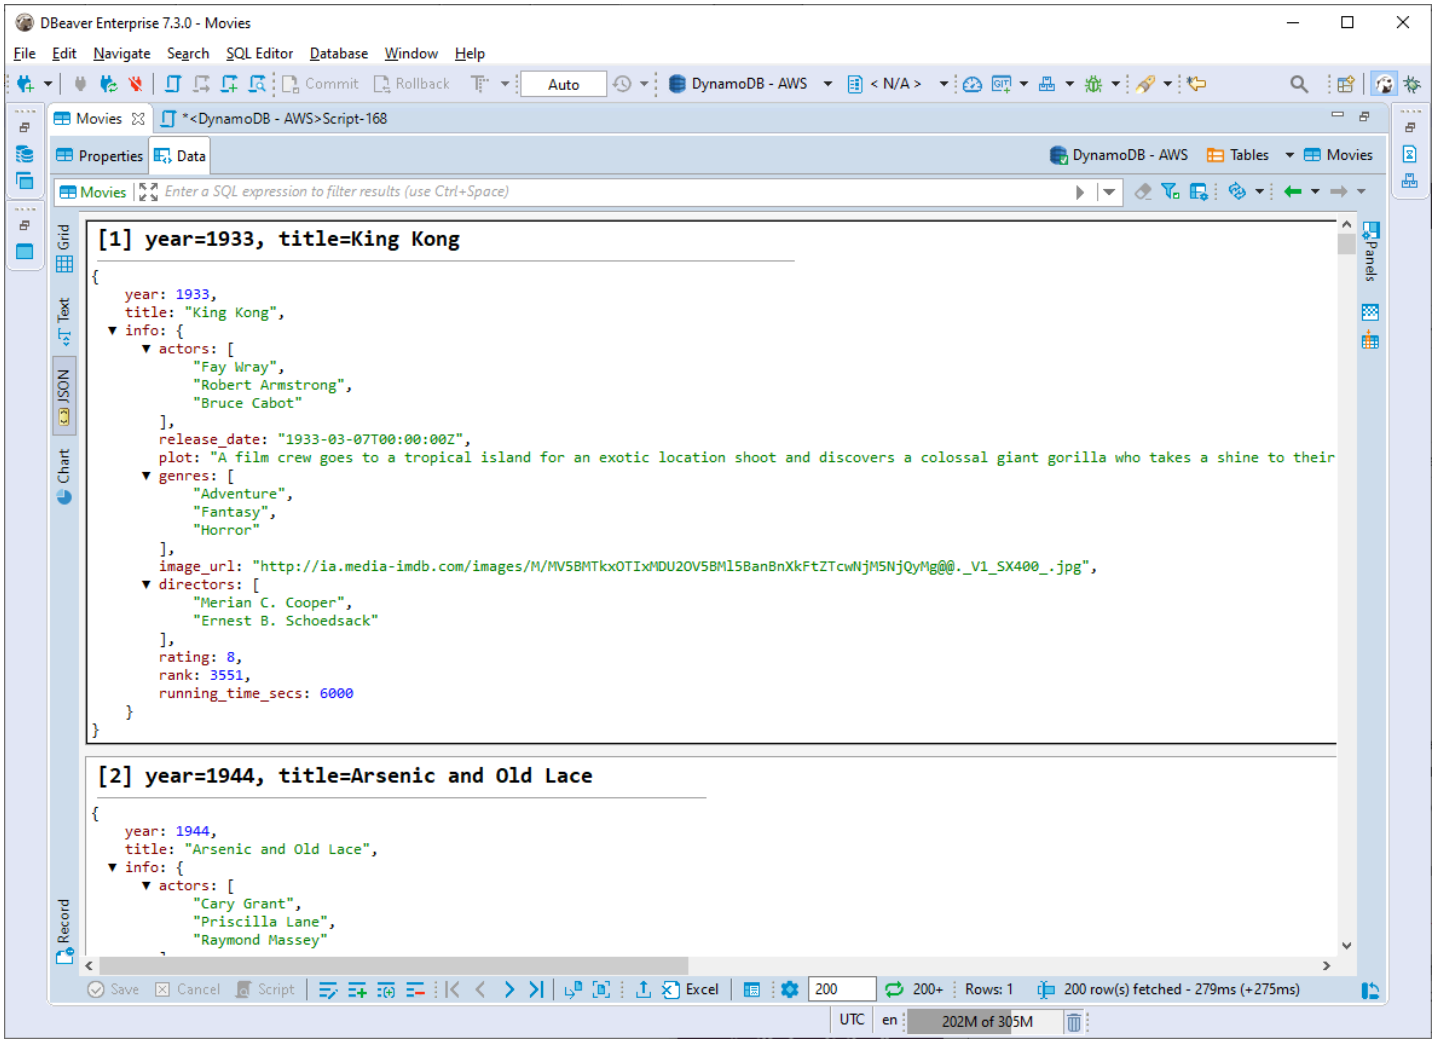

Executing queries

DBeaver support simple SQL dialect for DynamoDB.
You can use WHERE clause in the same fashion as in regular SQL in order to find or filter documents.

You can also use JSON requests syntax to query documents. See Amazon DynamoDB query reference.

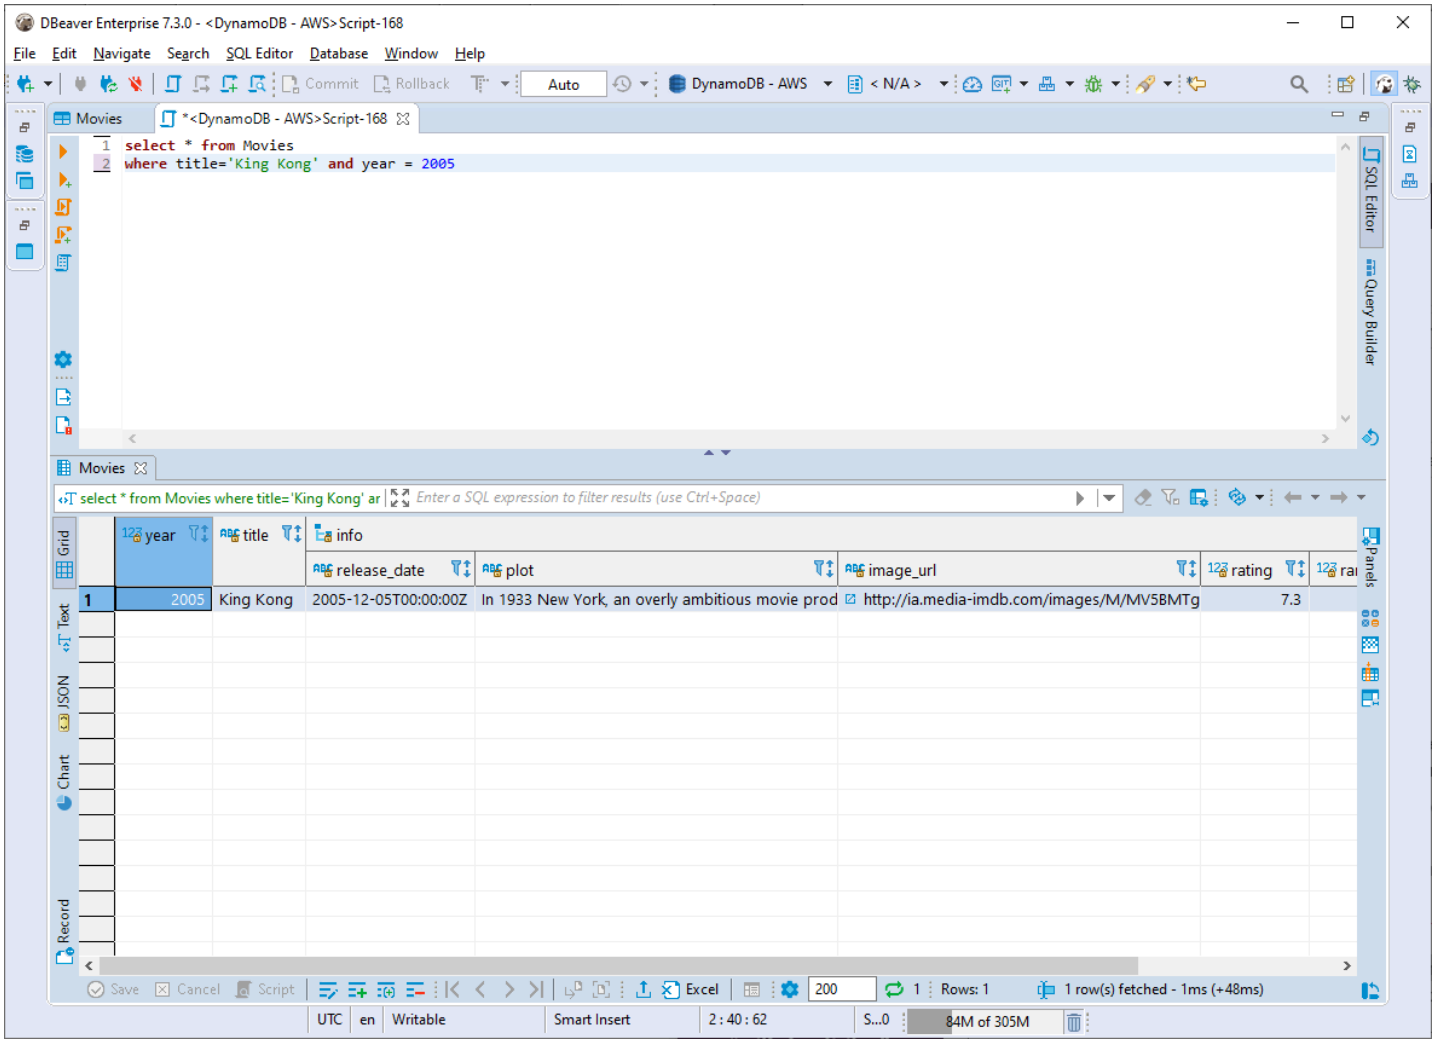

Exporting and importing data

You can export data from DynamoDB table in different file formats (CSV, XLSX, XML, JSON, etc) or export data directly to another table.

AWS DocumentDB

AWS DocumentDB is based on the MongoDB engine.
It has several minor differences in the query processing and network configuration.
However most features which work for MongoDB will work for DocumentDB as well. Please refer to MongoDB article.

Connections

AWS restricts direct access to DocumentDB clusters from outside of the cloud (region). So you can connect to it directly (using cluster host name) only when DBeaver is deployed on the EC2 instance.

In other cases you will need to use SSH tunnel thru a proxy machine to access DocumentDB instance. Please read AWS Documentation about proxy configuration: https://docs.aws.amazon.com/documentdb/latest/developerguide/connect-from-outside-a-vpc.html

In DBeaver you can use SSH tab on the connection settings page. Just enter proxy host, user name and specify private key file (it is provided by AWS as a keypair).

Queries

DBeaver processes DocDB SQL queries exactly like in MongoDB. It supports SELECT, UPDATE, INSERT and DELETE queries.
SELECT queries support WHERE, ORDER BY, GROUP BY and HAVING clauses.

DocumentDB restricts `eval` function so all JavaScript queries will be parsed on the client side and then evaluated at DocDB cluster one by one. Most of JS functions works exactly like in Mongo Shell.

Couchbase

Connecting to Couchbase server

Couchbase client uses multiple ports to connect to a cluster (8091-8096, 9140, etc). Some of these ports are dynamic (i.e. depend on server settings) and can't be overwritten. It makes SSH tunnelling impossible. Thus, if you work with remote Couchbase deployed behind a firewall you will need to setup VPN connection or SOCKS proxy.

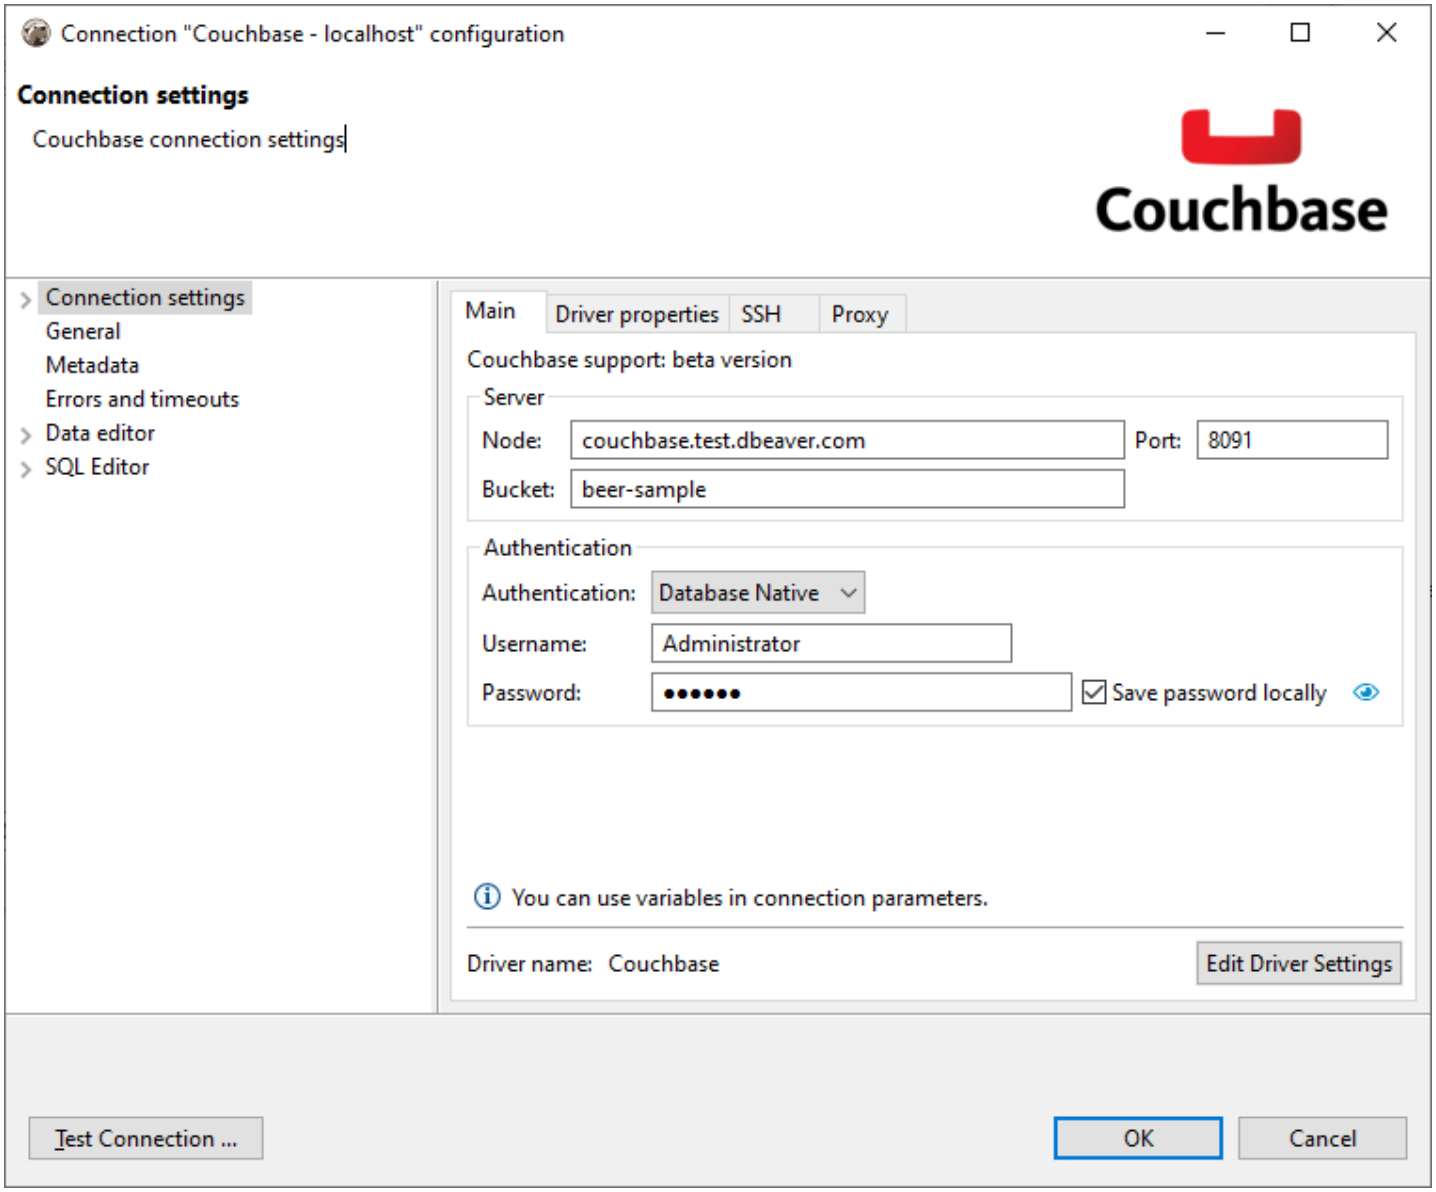
Connection "Couchbase - localhost" configuration
Connection settings
Couchbase connection settings
Couchbase
Connection settings
General
Metadata
Errors and timeouts
Data editor
SQL Editor
Main
Driver properties
SSH
Proxy
Couchbase support: beta version
Server
Node:
couchbase.test.dbeaver.com
Port:
8091
Bucket:
beer-sample
Authentication
Authentication:
Database Native
Username:
Administrator
Password:
Save password locally
You can use variables in connection parameters.
Driver name: Couchbase
Edit Driver Settings
Test Connection ...
OK
Cancel

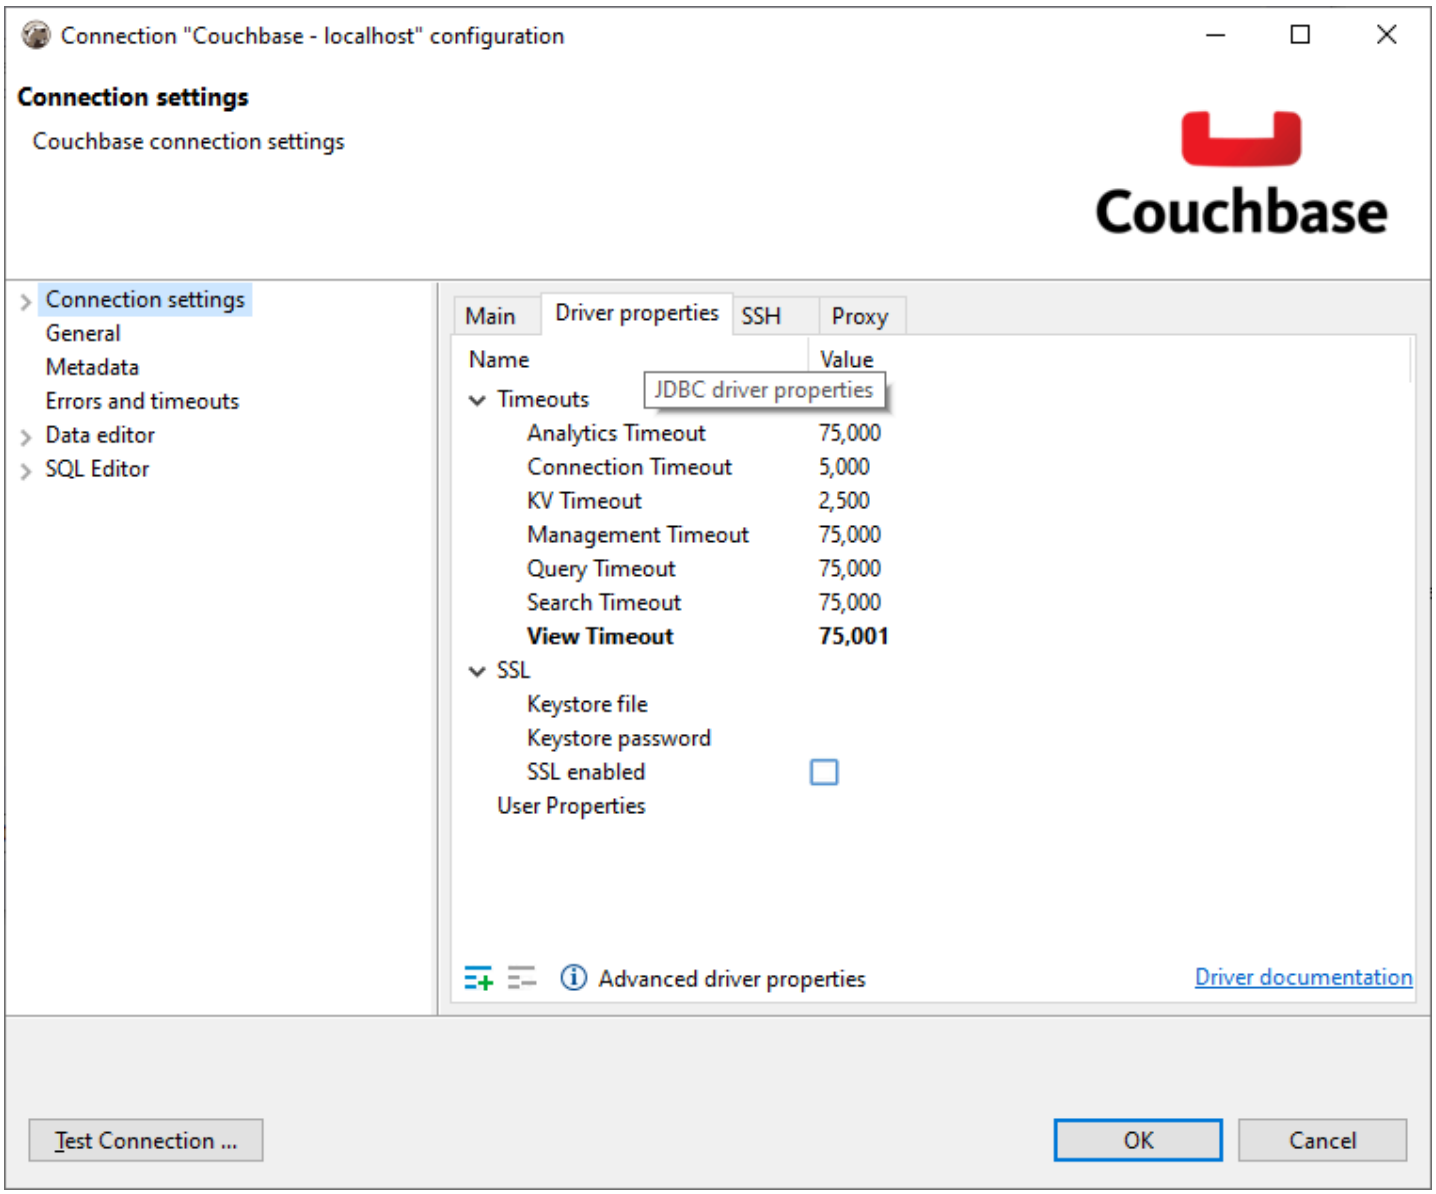

Viewing and editing Couchbase tables

Couchbase is a document-oriented database. It means that all documents may have different structure. You can view/edit buckets content as standard relational tables (grid/plain text presentations) or as JSON documents.

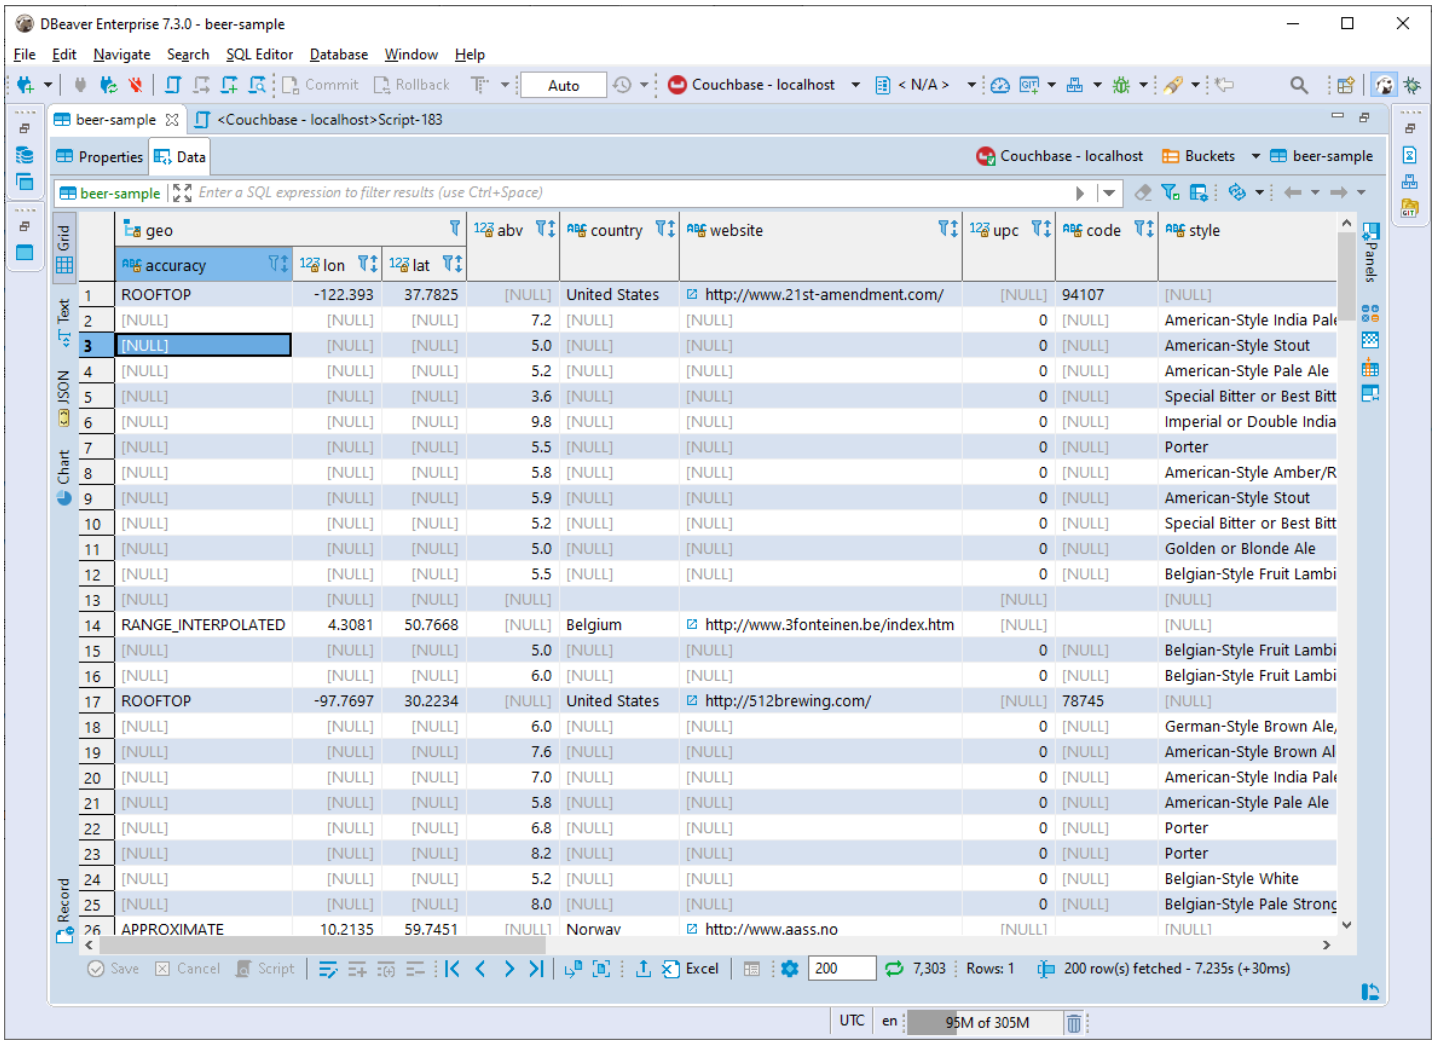

	geo			abv	country	website	upc	code	style
	accuracy	lon	lat						
1	ROOFTOP	-122.393	37.7825	[NULL]	United States	http://www.21st-amendment.com/	[NULL]	94107	[NULL]
2	[NULL]	[NULL]	[NULL]	7.2	[NULL]	[NULL]	0	[NULL]	American-Style India Pale
3	[NULL]	[NULL]	[NULL]	5.0	[NULL]	[NULL]	0	[NULL]	American-Style Stout
4	[NULL]	[NULL]	[NULL]	5.2	[NULL]	[NULL]	0	[NULL]	American-Style Pale Ale
5	[NULL]	[NULL]	[NULL]	3.6	[NULL]	[NULL]	0	[NULL]	Special Bitter or Best Bitt
6	[NULL]	[NULL]	[NULL]	9.8	[NULL]	[NULL]	0	[NULL]	Imperial or Double India
7	[NULL]	[NULL]	[NULL]	5.5	[NULL]	[NULL]	0	[NULL]	Porter
8	[NULL]	[NULL]	[NULL]	5.8	[NULL]	[NULL]	0	[NULL]	American-Style Amber/R
9	[NULL]	[NULL]	[NULL]	5.9	[NULL]	[NULL]	0	[NULL]	American-Style Stout
10	[NULL]	[NULL]	[NULL]	5.2	[NULL]	[NULL]	0	[NULL]	Special Bitter or Best Bitt
11	[NULL]	[NULL]	[NULL]	5.0	[NULL]	[NULL]	0	[NULL]	Golden or Blonde Ale
12	[NULL]	[NULL]	[NULL]	5.5	[NULL]	[NULL]	0	[NULL]	Belgian-Style Fruit Lambi
13	[NULL]	[NULL]	[NULL]	[NULL]			[NULL]		[NULL]
14	RANGE_INTERPOLATED	4.3081	50.7668	[NULL]	Belgium	http://www.3fonteinen.be/index.htm	[NULL]		[NULL]
15	[NULL]	[NULL]	[NULL]	5.0	[NULL]	[NULL]	0	[NULL]	Belgian-Style Fruit Lambi
16	[NULL]	[NULL]	[NULL]	6.0	[NULL]	[NULL]	0	[NULL]	Belgian-Style Fruit Lambi
17	ROOFTOP	-97.7697	30.2234	[NULL]	United States	http://512brewing.com/	[NULL]	78745	[NULL]
18	[NULL]	[NULL]	[NULL]	6.0	[NULL]	[NULL]	0	[NULL]	German-Style Brown Ale,
19	[NULL]	[NULL]	[NULL]	7.6	[NULL]	[NULL]	0	[NULL]	American-Style Brown Al
20	[NULL]	[NULL]	[NULL]	7.0	[NULL]	[NULL]	0	[NULL]	American-Style India Pale
21	[NULL]	[NULL]	[NULL]	5.8	[NULL]	[NULL]	0	[NULL]	American-Style Pale Ale
22	[NULL]	[NULL]	[NULL]	6.8	[NULL]	[NULL]	0	[NULL]	Porter
23	[NULL]	[NULL]	[NULL]	8.2	[NULL]	[NULL]	0	[NULL]	Porter
24	[NULL]	[NULL]	[NULL]	5.2	[NULL]	[NULL]	0	[NULL]	Belgian-Style White
25	[NULL]	[NULL]	[NULL]	8.0	[NULL]	[NULL]	0	[NULL]	Belgian-Style Pale Strong
26	APPROXIMATE	10.2135	59.7451	[NULL]	Norway	http://www.aass.no	[NULL]		[NULL]

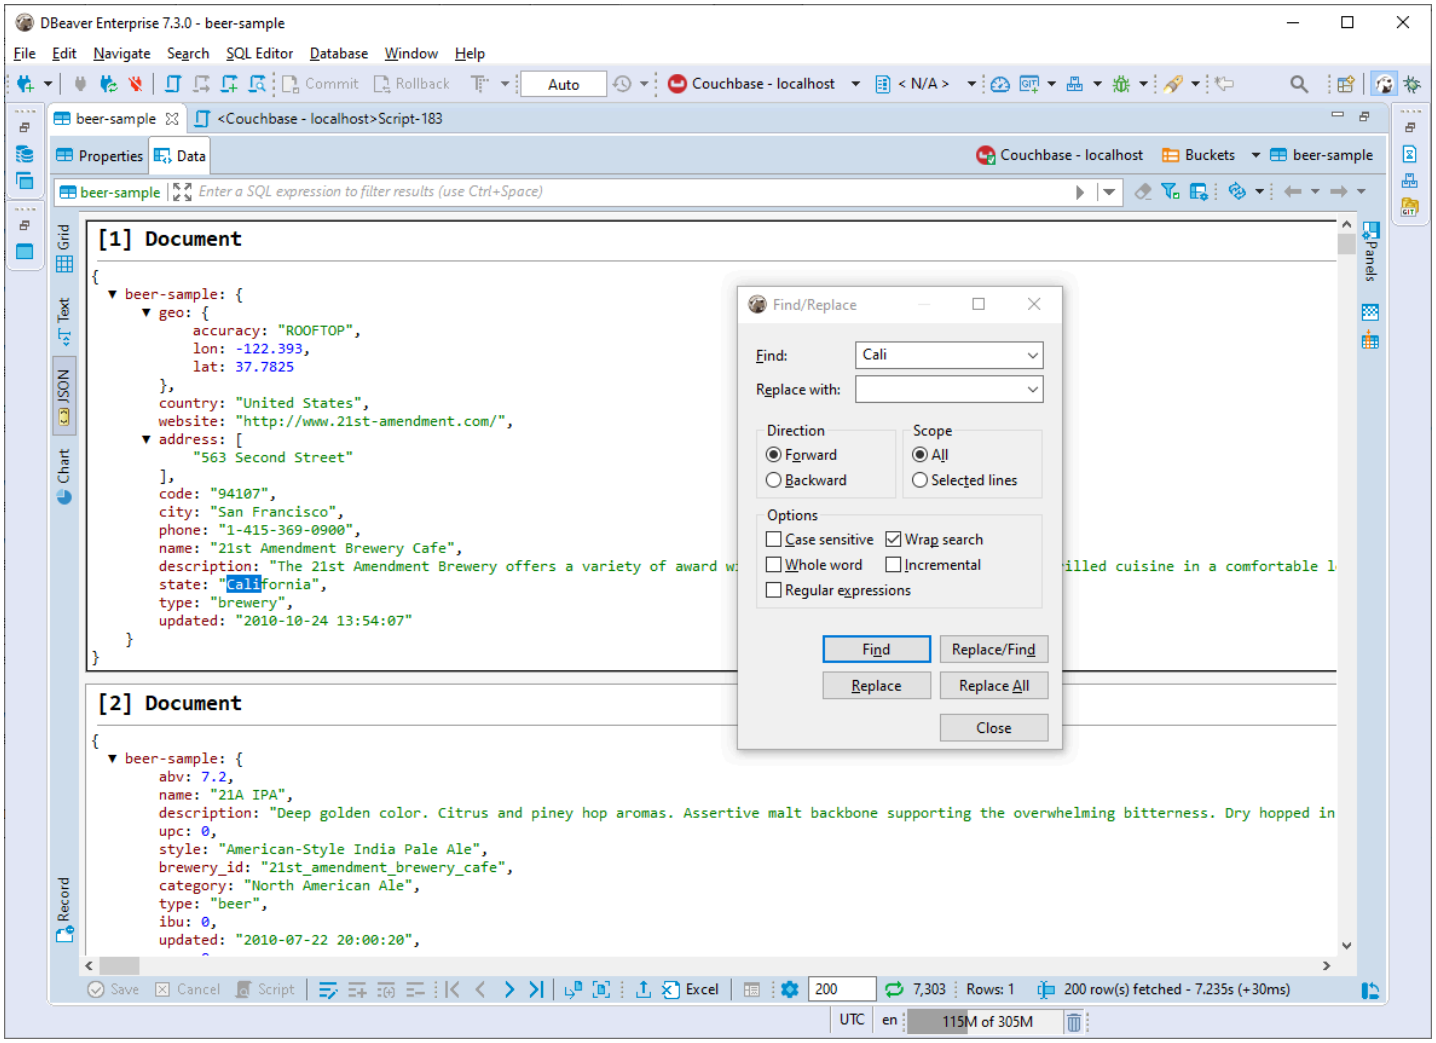

Executing Couchbase queries

Couchbase uses N1QL language for queries. It is very similar to the standard SQL language.

```
SELECT country FROM `travel-sample` WHERE name = "Excel Airways";
```

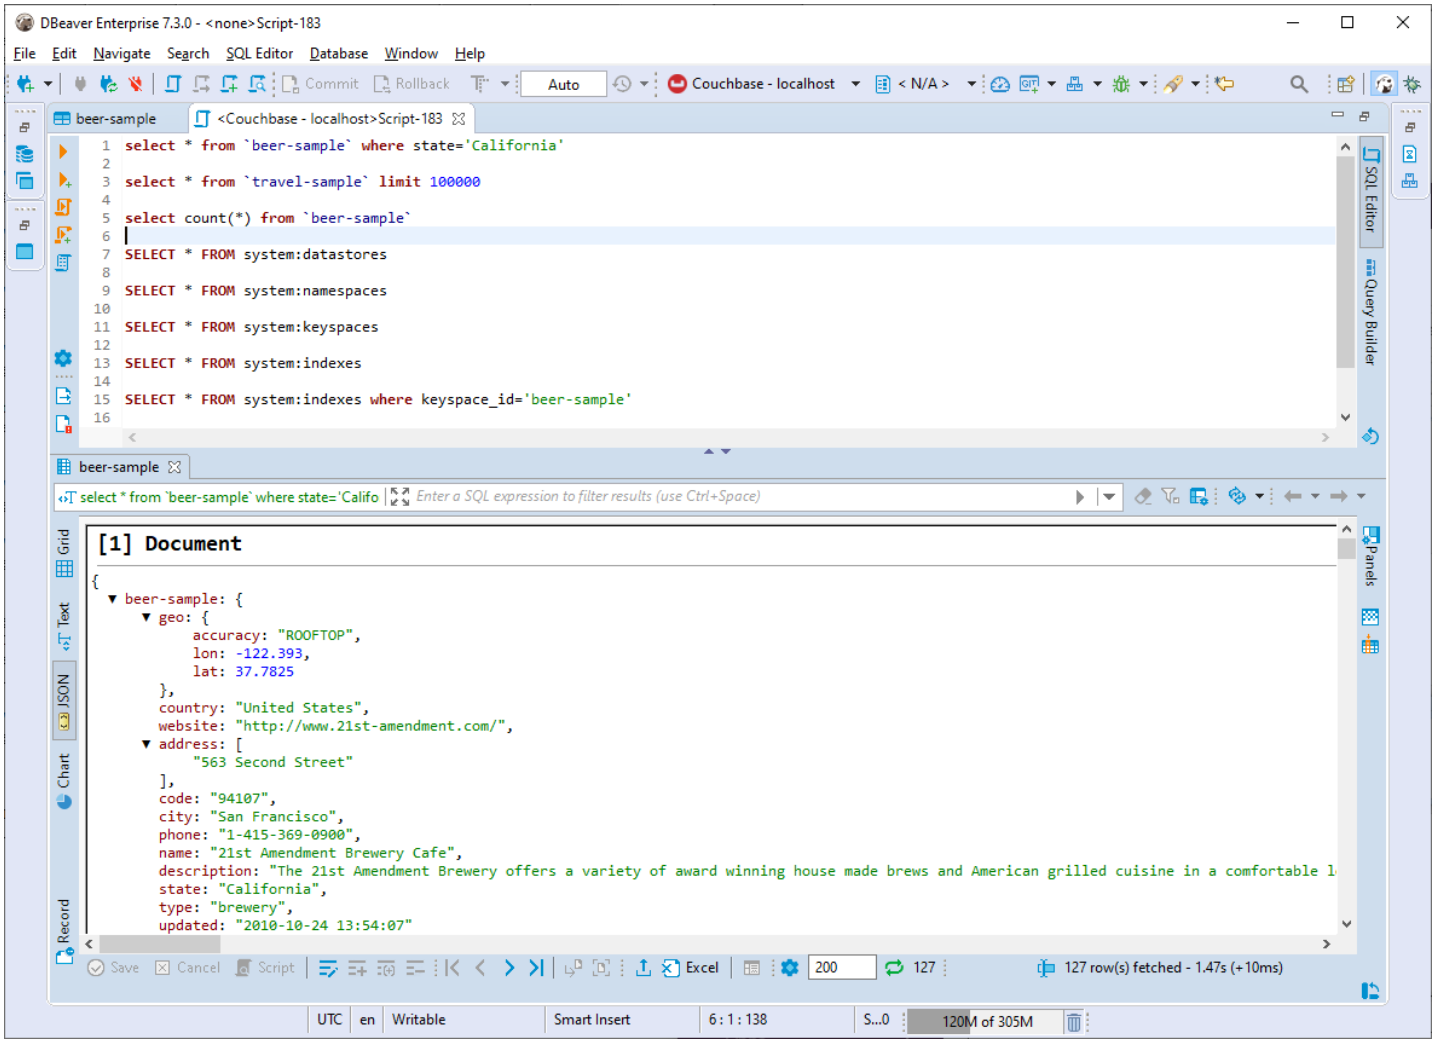
DBeaver Enterprise 7.3.0 - <none>Script-183
File Edit Navigate Search SQL Editor Database Window Help
Commit Rollback Auto Couchbase - localhost < N/A >
beer-sample
<Couchbase - localhost>Script-183
1 select * from `beer-sample` where state='California'
3 select * from `travel-sample` limit 100000
5 select count(*) from `beer-sample`
7 SELECT * FROM system:datastores
9 SELECT * FROM system:namespaces
11 SELECT * FROM system:keyspaces
13 SELECT * FROM system:indexes
15 SELECT * FROM system:indexes where keyspace_id='beer-sample'
SQL Editor
Query Builder
beer-sample
select * from `beer-sample` where state='Califo
Enter a SQL expression to filter results (use Ctrl+Space)
Grid
Text
JSON
Chart
Record
Panels
[1] Document
{
beer-sample: {
geo: {
accuracy: "ROOFTOP",
lon: -122.393,
lat: 37.7825
},
country: "United States",
website: "http://www.21st-amendment.com/",
address: [
"563 Second Street"
],
code: "94107",
city: "San Francisco",
phone: "1-415-369-0900",
name: "21st Amendment Brewery Cafe",
description: "The 21st Amendment Brewery offers a variety of award winning house made brews and American grilled cuisine in a comfortable l
state: "California",
type: "brewery",
updated: "2010-10-24 13:54:07"
Save Cancel Script Excel 200 127
127 row(s) fetched - 1.47s (+10ms)
UTC en Writable Smart Insert 6 : 1 : 138 S...0 120M of 305M

Apache Hive/Spark/Impala

Apache Hive

Hive is a Hadoop-based storage system. Hive uses a special SQL dialect (HiveQL) to operate with data and metadata. Generally, it is quite similar to SQL.

There is multiple implementation of storage systems which utilize Hive on server-side - including Apache Spark, Impala, etc. Most of them support standard Hive JDBC driver which is used in DBeaver to communicate with the server.

DBeaver uses so-called Hive JDBC Uber Jar driver (https://github.com/timveil/hive-jdbc-uber-jar) which includes all necessary dependencies. You don't need to download anything - DBeaver will download everything automatically (if you have internet access).

Connection setup

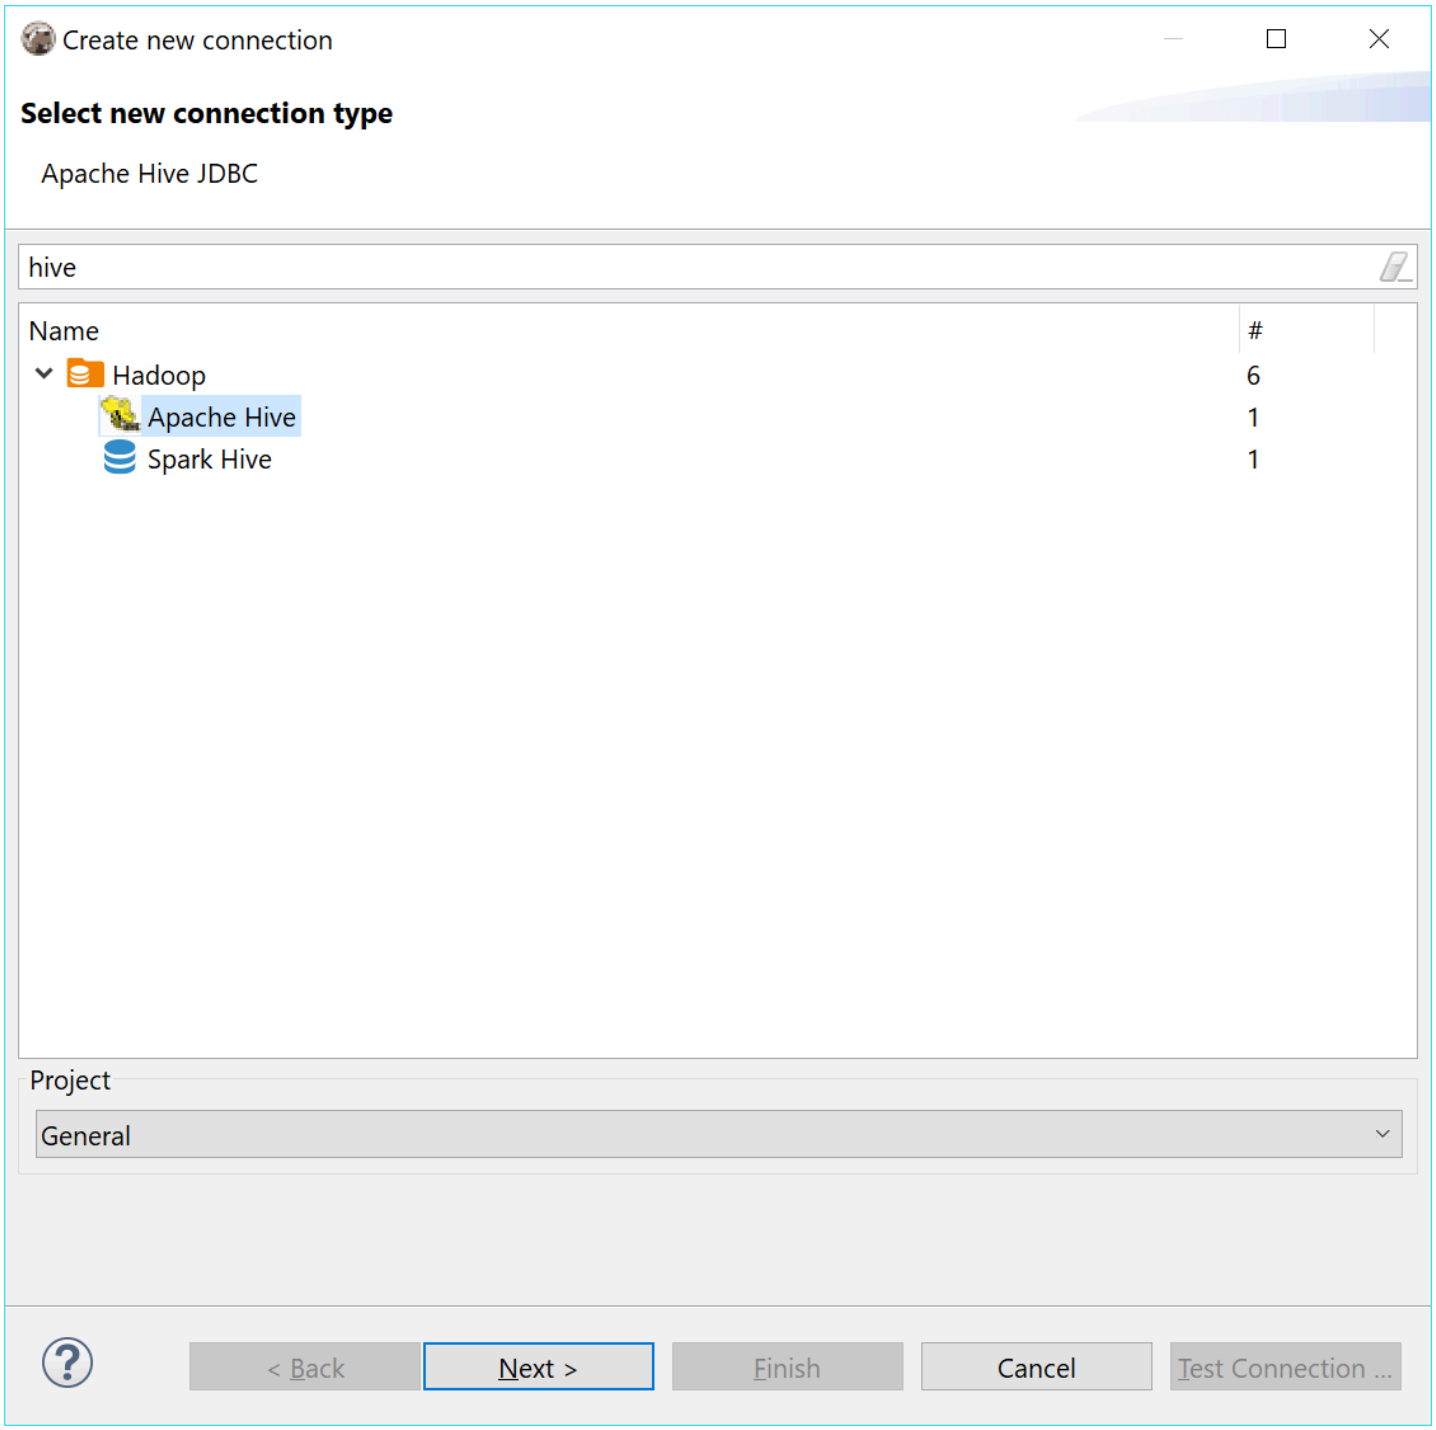
Create new connection
Select new connection type
Apache Hive JDBC
hive
Name
#
Hadoop
6
Apache Hive
1
Spark Hive
1
Project
General
< Back
Next >
Finish
Cancel
Test Connection ...

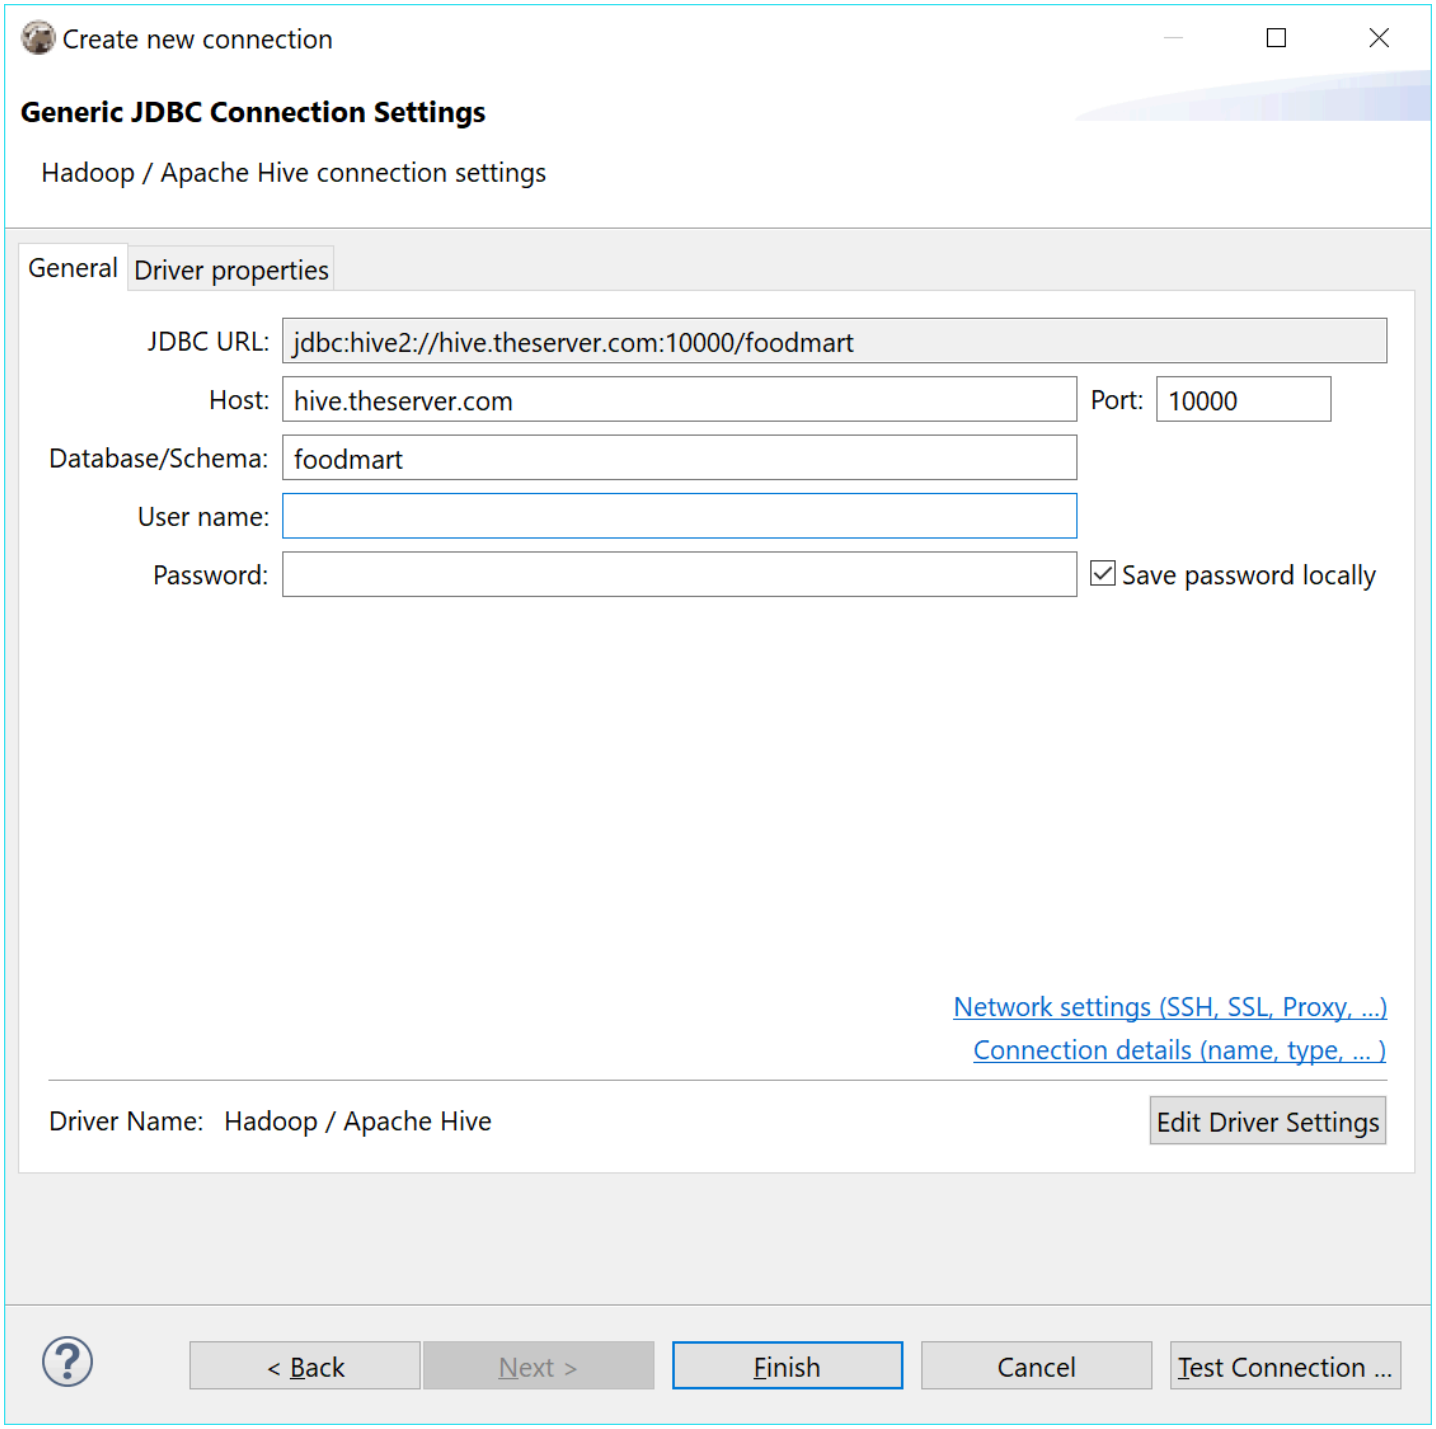

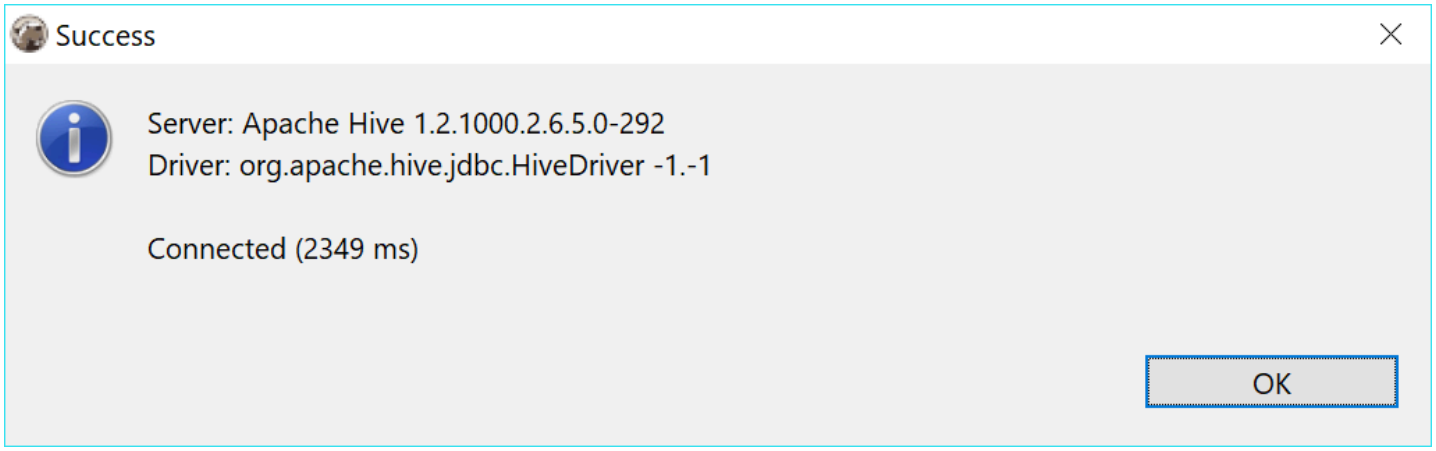

Schema/data browser

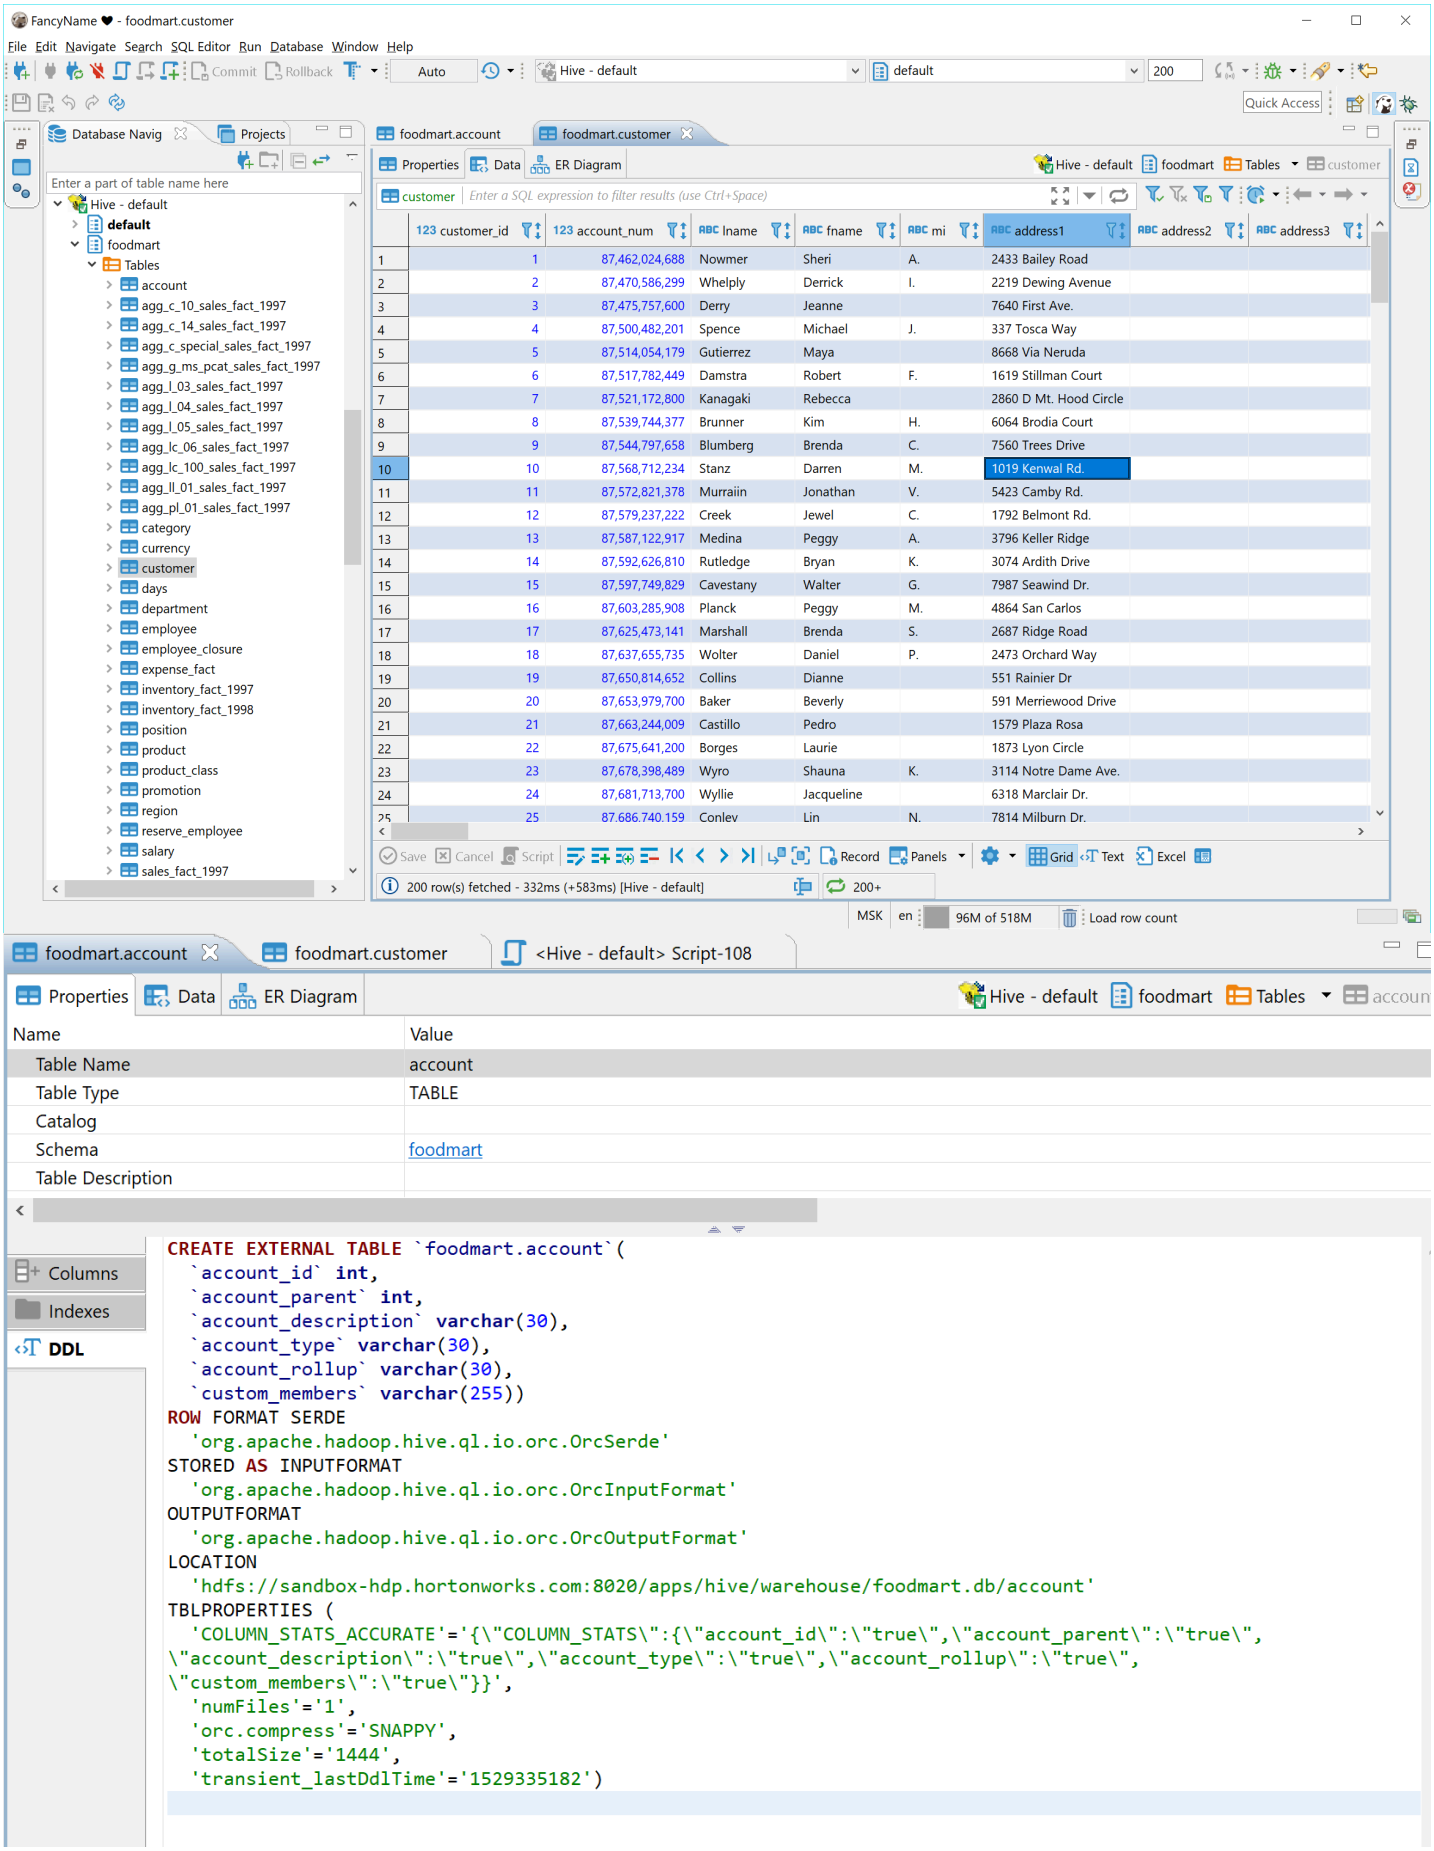

Limitations

Hie doesn't support referential integrity so you won't see primary keys or foreign keys. ER diagrams also don't make much sense.

Installing extensions - Themes, version control, etc

You can install a lot of optional extensions (plugins) in DBeaver. Most of extensions can be found on Eclipse Marketplace website.

Popular extensions for DBeaver

- Darkest Dark theme - the best Dark theme for DBeaver
- Eclipse Color Theme - use it if you don't like Darkest Dark theme by some reason
- Git support - Git version control integration
- Subversion support - Subversion integration
- Embedded Shell - Allows to run shell commands directly from DBeaver
- Editor vertical indents - Adds vertical indents to all text editors

DBeaver-specific extensions

- Office formats support (XLSX)
- Vector graphics support (SVG)
- SSHJ and advanced cryptography
- SQL debugger

Install Process

In DBeaver EE you can use drag-n-drop from Marketplace web site (see button `Install`) into DBeaver main window. This will launch Marketplace installation wizard automatically. In DBeaver Community or other DBeaver-based products which do not include marketplace clients you can use following instructions:

Extension installation in CE version:

1. Copy URL of extension update site:

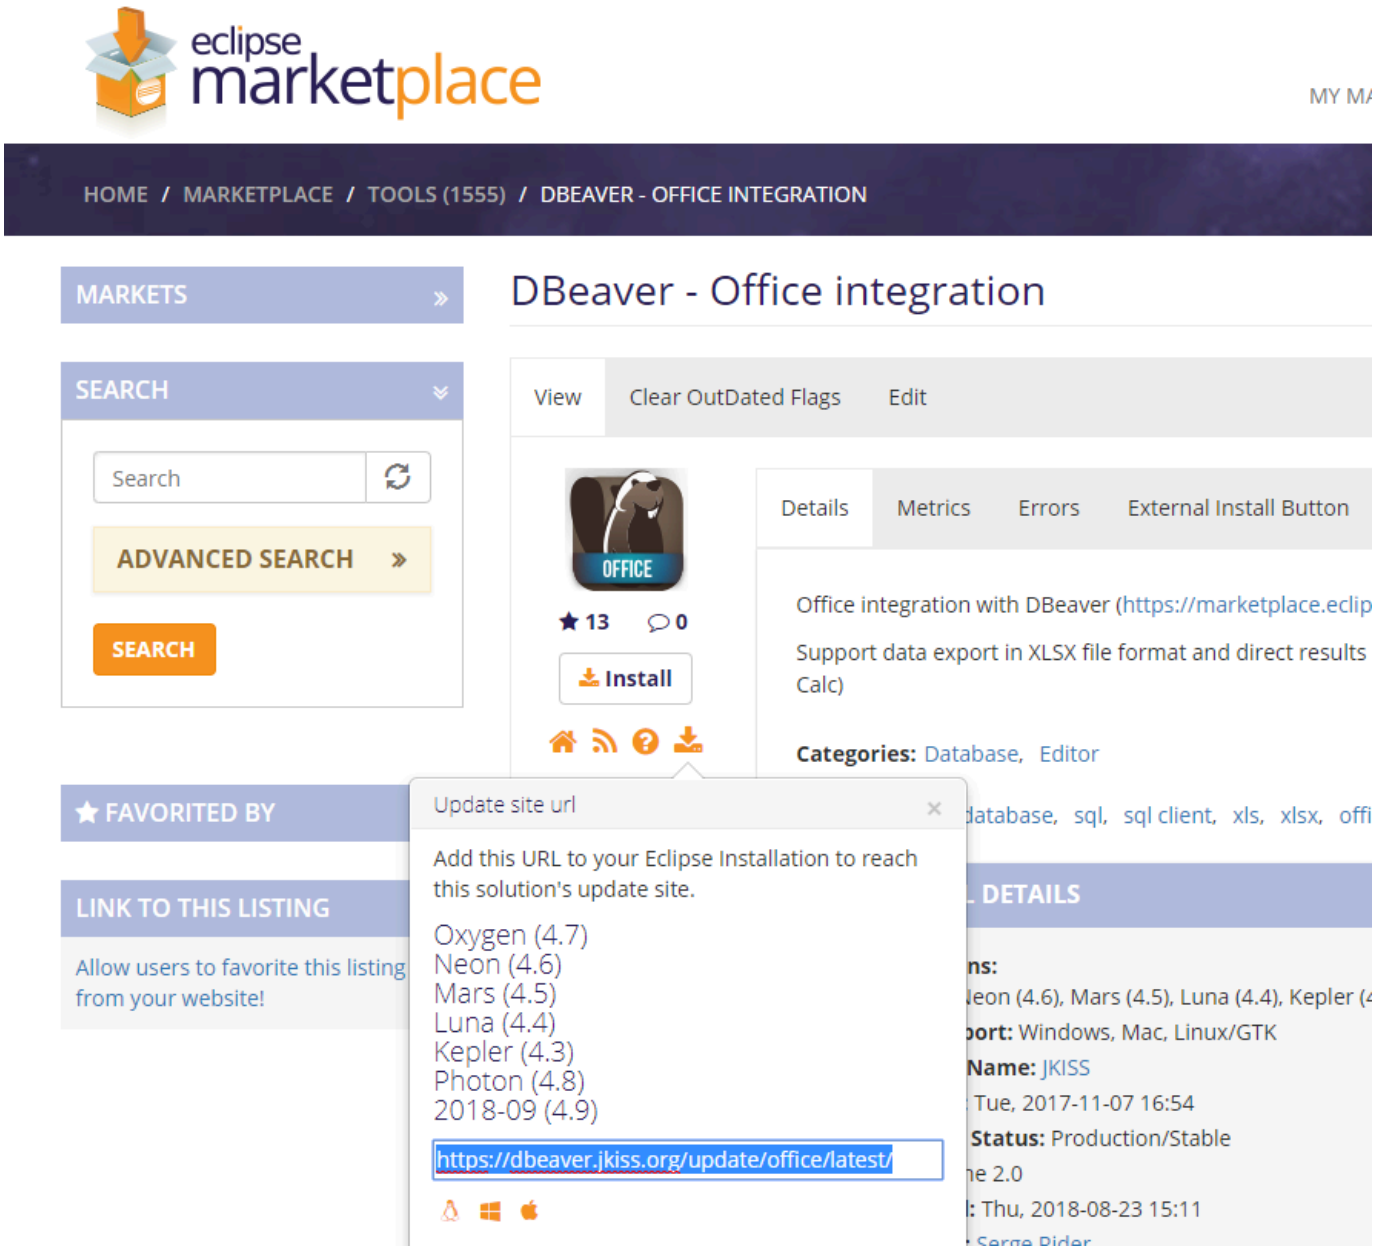

2. In DBeaver main menu open `Help -> Install New Software`
3. Paste update site URL into `Work with` field and press `Enter`

4. Check items you wish to install (in most cases just all items)

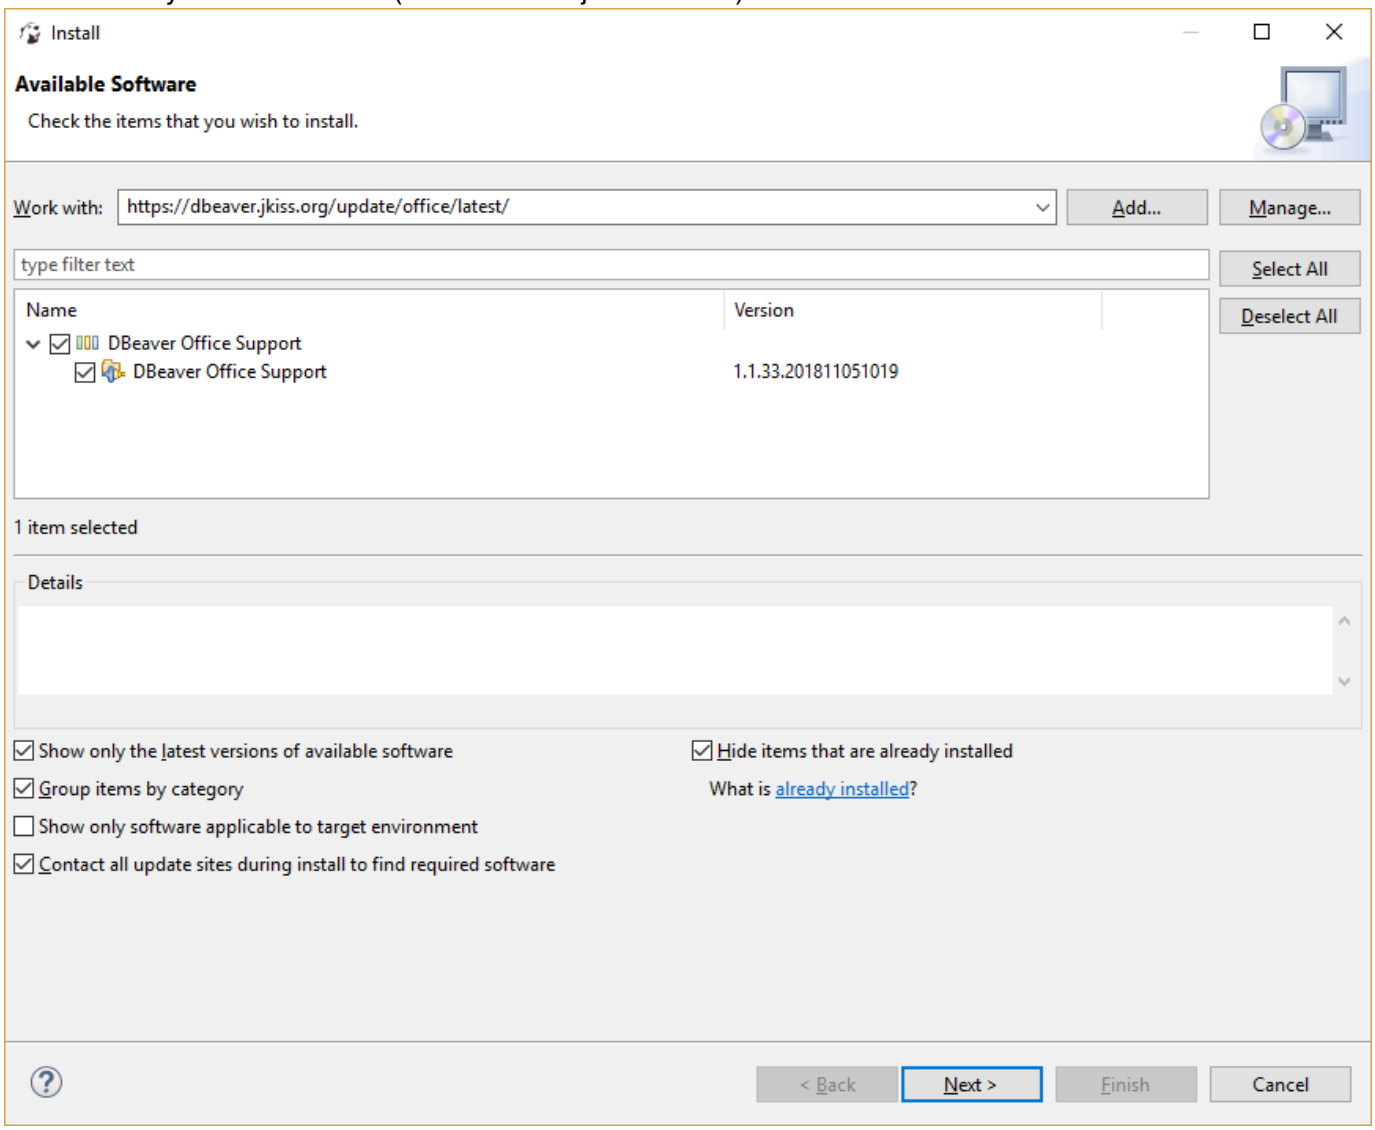

5. Click Next. You may need to accept extension license before installation

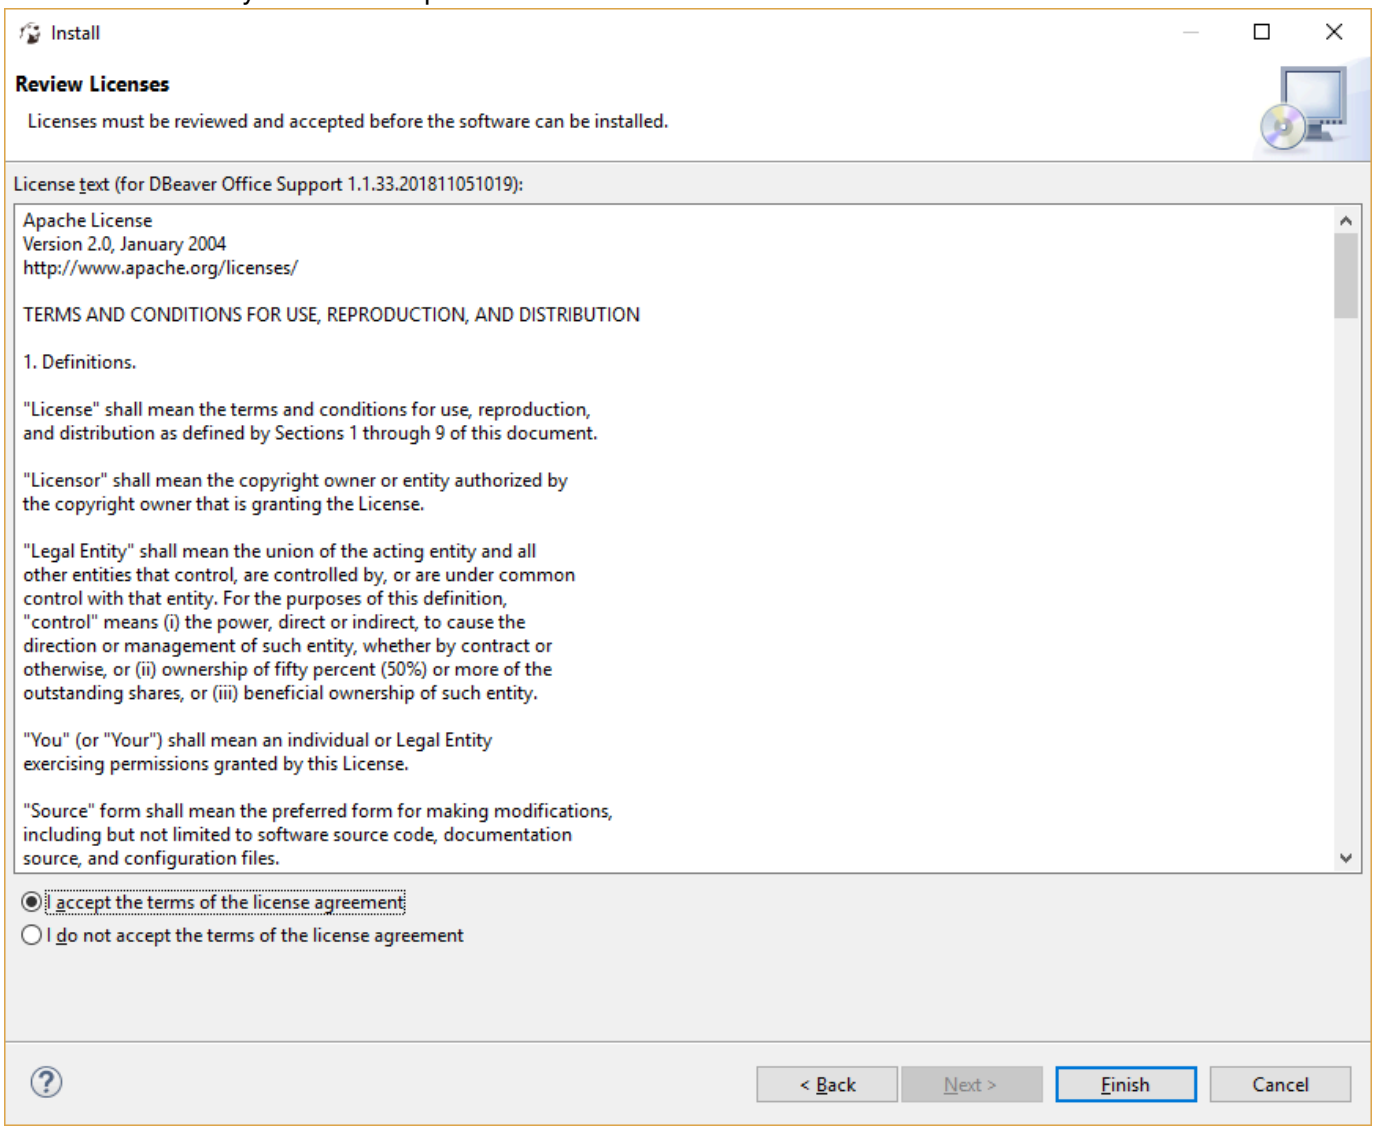

6. Some extensions may contain unsigned bundles. Install such extensions only if you really trust author.

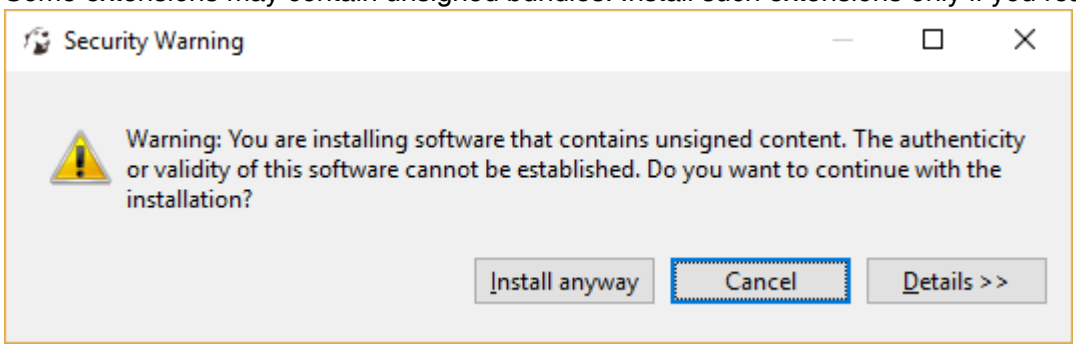

7. Click Next->Finish. Installation will take some time. Restart DBeaver.

Command line parameters

Command line parameters might be passed directly to dbeaver[.exe] executable.
On Windows you also can use `dbeaver-cli.exe` executable (it doesn't spawn new window so you can see output messages).

Also you can add parameters in the `dbeaver.ini` configuration file (in the beginning, each parameter on its own line).

DBeaver control

Name	Value	Example
-help	Prints help message	
-stop	Quits DBeaver	
-dump	Prints DBeaver thread dump	
-f	Opens file in DBeaver UI	`-f c:\some-path\some-file.sql`
-con	Opens database connection in DBeaver UI	See connection parameters table
-closeTabs	Closes all open editor tabs	
-disconnectAll	Closes all open connections	
-reuseWorkspace	Force reuse of single workspace by multiple DBeaver instances	
-newInstance	Force new DBeaver instance creation (do not try to reuse already running one)	
-var EE	Custom variables for runTask. You can change existing variables in the task, you cannot add new task variables with this parameter. You can add several parameters at once to the command line, each starting with "-var". Used right before -runTask. Template: "-var variableName=variableValue"	`-var film=sakila.film` `-var actor=sakila.actor` `-runTask "exportFromSakila"` EE version only.
-runTask EE	Executes specified task	`-runTask "@projectName:taskName"` . EE version only. See task scheduler.
-license EE	Path to the EE license file	`-license "/etc/licenses/dbeaver.txt"` . EE version only.

System parameters

Name	Value	Example

-nl	Locale	en_US
-data	Workspace path	c:\ProgramData\MyWorkspace
-nosplash	Omits splash screen	true
-clean	Clears all Eclipse caches. Use it if DBeaver fails to start after version upgrade.	
-vmargs	VM parameters	See VM arguments table

VM arguments

You can pass any advanced Java parameters supported by your local JVM (Oracle, OpenJDK, IBM, etc).
Parameters supported by Oracle JVM (11): https://docs.oracle.com/en/java/javase/11/tools/java.html

Parameters supported by all JVMs:

Name	Value	Example
-Xms	Sets initial memory available for DBeaver	`-Xmx1000m`
-Xmx	Sets maximum memory available for DBeaver	`-Xmx4000m`

Connection parameters

All connection parameters must be supplied as a single command line argument, parameters are divided by pipe (`|`). Parameter name and value are divided by `=` .
Example: -con driver=sqlite|database=C:\db\SQLite\Chinook.
db|name=SQLiteChin|openConsole=true|folder=SQLite

Name	Value	Example
name	Connection name	`Test connection`
driver	Driver name or ID	`driver=sqlite` , `driver=mysql` , etc
url	Connection URL. Optional (JDBC URL may be constructed by driver from other parameters)	`url=jdbc:sqlite:C:\db\SQLite\Chinook.db`
host	Database host name (optional)	`host=localhost`
port	Database port number (optional)	`port=1534`
server	Database server name (optional)	`server=myserver`
database	Database name or path (optional)	`database=db-name`
user	User name (optional)	`user=root`
password	User password (optional)	`password=mysecret`
auth	Authentication model ID. See Auth models	`auth=postgres_pgpass`
authProp.propName	Custom authentication parameters (depends on driver and auth model)	`authProp.oracle.net.wallet_location=C:/temp/ora-wallet`
savePassword	Do not ask use password on connect	`savePassword=true`

showSystemObjects	Show/hide system schemas, tables , etc	`showSystemObjects=true`
showUtilityObjects	Show/hide utility schemas, tables ,etc	`showUtilityObjects=true`
folder	Put new connection in a folder	`folder=FolderName`
autoCommit	Sets connection auto commit flag (default value depends on driver)	`autoCommit=true`
prop.propName	Advanced connection parameters (depend on driver)	`prop.connectTimeout=30`
id	Connection id	`oracle_thin-16a88e815bd-70598e648cedd28c` (useful in conjunction with `create=false`)
connect	Connect to this database	`connect=false`
openConsole	Open SQL console for this database (sets `connect` to true)	`openConsole=true`
create	Create new connection	`create=false` (true by default). If set to false then existing connection configuration will be used. name or id parameter must be specified.

Reset UI settings

Sometimes, usually after multiple version and /or upgrades/incorrect shutdowns DBeaver UI may become corrupted. Extra toolbar elements, missing menu items, broken keyboard shortcuts, broken localization strings and other glitches may happen.

To reset DBeaver UI just delete file `workbench.xmi` in DBeaver workspace/.metadata. By default workbench.xmi file locations is:

- Windows: `%APPDATA%\DBeaverData\workspace6\.metadata\.plugins\org.eclipse.e4.workbench\workbench.xmi`
- MacOS: `~/Library/DBeaverData/workspace6/.metadata/.plugins/org.eclipse.e4.workbench/workbench.xmi`
- Linux: `$XDG_DATA_HOME/DBeaverData/workspace6/.metadata/.plugins/org.eclipse.e4.workbench/workbench.xmi`

To reset settings:

1. Close DBeaver
2. Delete workbench.xmi from Explorer/Finder or open terminal and run `del` (Windows) or `rm` (Linux/MacOS) followed by workbench.xmi path.
3. Start DBeaver

Reset workspace

Sometimes (especially after multiple DBeaver versions upgrade) workspace become messy.
Some keyboard shortcuts may stop working, toolbars layout may be broken, etc, etc.
To reset all UI settings (this includes menus, shortcuts, view and toolbar layouts):

1. Shutdown DBeaver
2. Go to the default workspace folder `.metadata\.plugins\org.eclipse.e4.workbench\`
    - Windows: `Win+R` , enter `%APPDATA%\DBeaverData\workspace6\.metadata\.plugins\org.eclipse.e4.workbench\`
    - MacOS: `open ~/Library/DBeaverData/workspace6/.metadata/.plugins/org.eclipse.e4.workbench/`
    - Linux: `cd $XDG_DATA_HOME/DBeaverData/workspace6/.metadata/.plugins/org.eclipse.e4.workbench/`
3. Delete file `workbench.xmi`
4. Start DBeaver

If that doesn't help then you can try to remove `.metadata/.plugins/org.eclipse.core.resources` folder.

If that doesn't help then remove `.metadata` folder. It will erase all your UI settings (but all connections, settings and scripts will remain as is).

That's it.

Posting issues

A few tips.

- Check existing issues for your issue (including closed ones). Duplicating an issue is slower for both parties so search through open and closed issues to see if what you're running into has been addressed already.
- Be clear about what your problem is: what was the expected outcome, what happened instead? Detail how someone else can recreate the problem.
- If you posting a bug report check "Error Log" view. If there are any errors related to your bug then post a complete stacktrace. Sometimes there are no errors in Error Log - then try to find them in log files.
- If your issue is related to a database data or metadata management - check the Query Manager view. It contains information about all queries DBeaver executes (explicitly or implicitly). To see more detailed information you can configure Query Manager in Preferences.

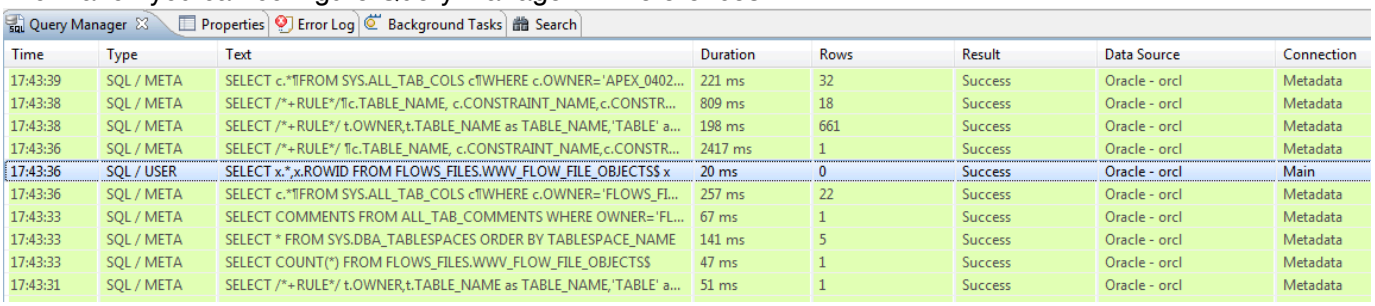

- Depending on the nature of your bug report provide information about:
  - Operating system
  - Window manager (for Linux)
  - Database (name and version)
  - Database driver (name and version)
- Do not write issue type in the issue title (like Feature Request:, Bug: etc). We'll review your issue and assign a corresponding label.

Log files

Error Log view

There is Error Log view (main menu Window->Show View->Error Log) which contains all errors occurred during DBeaver runtime.
You can double click on warning/error in the log viewer and see error stacktrace. Please attache it to the bug report.
Also you can open full log (all error messages) if you need:

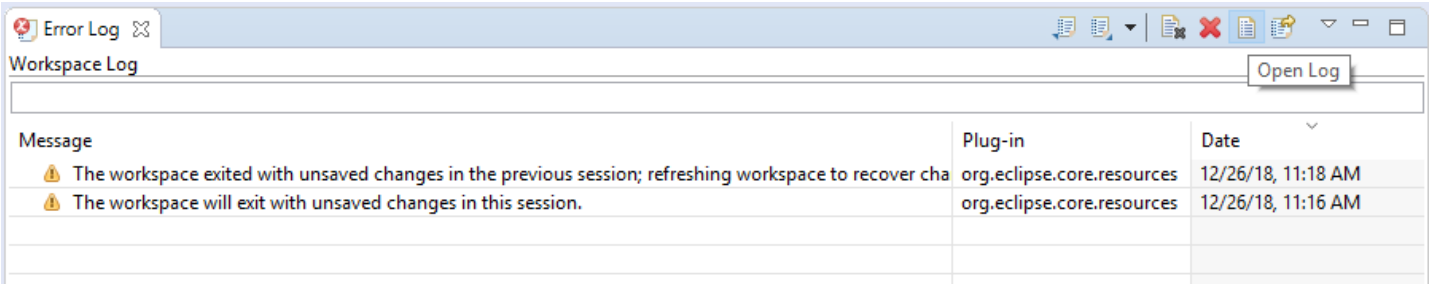

Log files

DBeaver writes different log files. Most of them are Eclipse logs.
Usually log files reside in the workspace/workspace6/.metadata .

- On Windows open Explorer and paste path `%APPDATA%\DBeaverData\workspace6\.metadata`.
- On Linux just type `cd $XDG_DATA_HOME/DBeaverData/workspace6/.metadata`
- On MacOS open path `~/Library/DBeaverData/workspace6/.metadata` in Finder.
    - To view hidden folders press `Cmd+Shift+.` in the folder view.

Two standard log files:

- workspace/workspace6/.metadata/.log - all warnings and errors which happens during normal work
- workspace/workspace6/.metadata/dbeaver-debug.log - the same as `.log` plus debug information

In special cases log files can be written in other directories. Special case is an emergency situation when DBeaver can't start and there is no workspace. Two typical places to find emergency logs:

- `<install-path>/configuration`
- `${HOME}/.eclipse/org.jkiss.dbeaver.product_<dbeaver-version>`

If you are reporting about some error please attach logs (not complete file but valuable part of it).
Logs are very useful, big number of errors can't be reproduced/fixed without full error stacktrace.

Java fatal logs

In rare cases DBeaver process dies and doesn't leave any valuable logs. This caused by Java VM crash. JVM creates a fatal log file for each crash (log gile `hs_err_PID.log`). This log usually resides in the same directory where dbeaver launcher is (e.g. dbeaver.exe).
But in some cases it is a write-protected directory and log file will be created in other folder.
Instruction how to find Java fatal log file: https://docs.oracle.com/javase/9/troubleshoot/fatal-error-log.htm

JDBC trace

In some cases custom JDBC drivers work incorrectly in DBeaver - shows wrong metadata like table columns, constraints or foreign keys.
Usually it happens because driver isn't compliant with JDBC API specification and DBeaver can't correctly interpret metadata provided by driver.

To understand what is going on inside driver you can enable JDBC tracing:

1. Find `dbeaver.ini` file (it is located in the same folder where DBeaver is installed)
2. Add line `-Ddbeaver.jdbc.trace=true` in the end of `dbeaver.ini`
3. Restart DBeaver
4. Connect to your database and browse metadata in the database navigator/object editors.
5. In DBeaver Workspace go to `.metadata` folder
6. File `jdbc-api-trace.log` contains all JDBC API invocations and all queries with results.

Analyzing contents of `jdbc-api-trace.log` you can understand what is wrong with metadata. Attach piece of trace file in GitHub ticket if you think that something is wrong on DBeaver side.

WARNING: disable JDBC tracing in your regular work. Enable it only for debugging. Trace generation decreases application performance and may produce huge log files.

Thread dump

Sometimes (due to some bug) DBeaver UI hangs, freezes or works incorrectly. Usually it is impossible to find the reason of such problem without thread dump. Thread dump is the information about internal execution state of Java program. To get thread dump:

Mac and Linux

Run the following on your terminal:

```
jstack $(ps aux | grep -m1 dbeaver | awk '{print $2}') > thread-dump.txt
```

Windows

Just open task manager (CTRL+Escape), find DBeaver in the process list and copy process ID value. On Windows 8+ you need to switch to "Details" tab. Run

```
jstack <PID> > thead-dump.txt
```

in Command Prompt.

Now you can attach thread-dump.txt to the GitHub issue.

Proxy configuration

External resources access

Sometimes DBeaver needs to access external internet resources for such tasks as:

- 3rd party JDBC drivers download
- Information about new DBeaver version
- Connect to remote databases outside of your corporate network
- Subscription license activation (commercial version)
- License information update (commercial version)

If you are behind some corporate firewall which restricts access to external internet resources then it may become a real problem.
Sometimes corporate firewalls allows to access external resources using web browser but restricts this for all other applications.

How to configure proxy for drivers download

You need ask your network administrator about proxy parameters.
Then go to Preferences->Connections->Drivers.

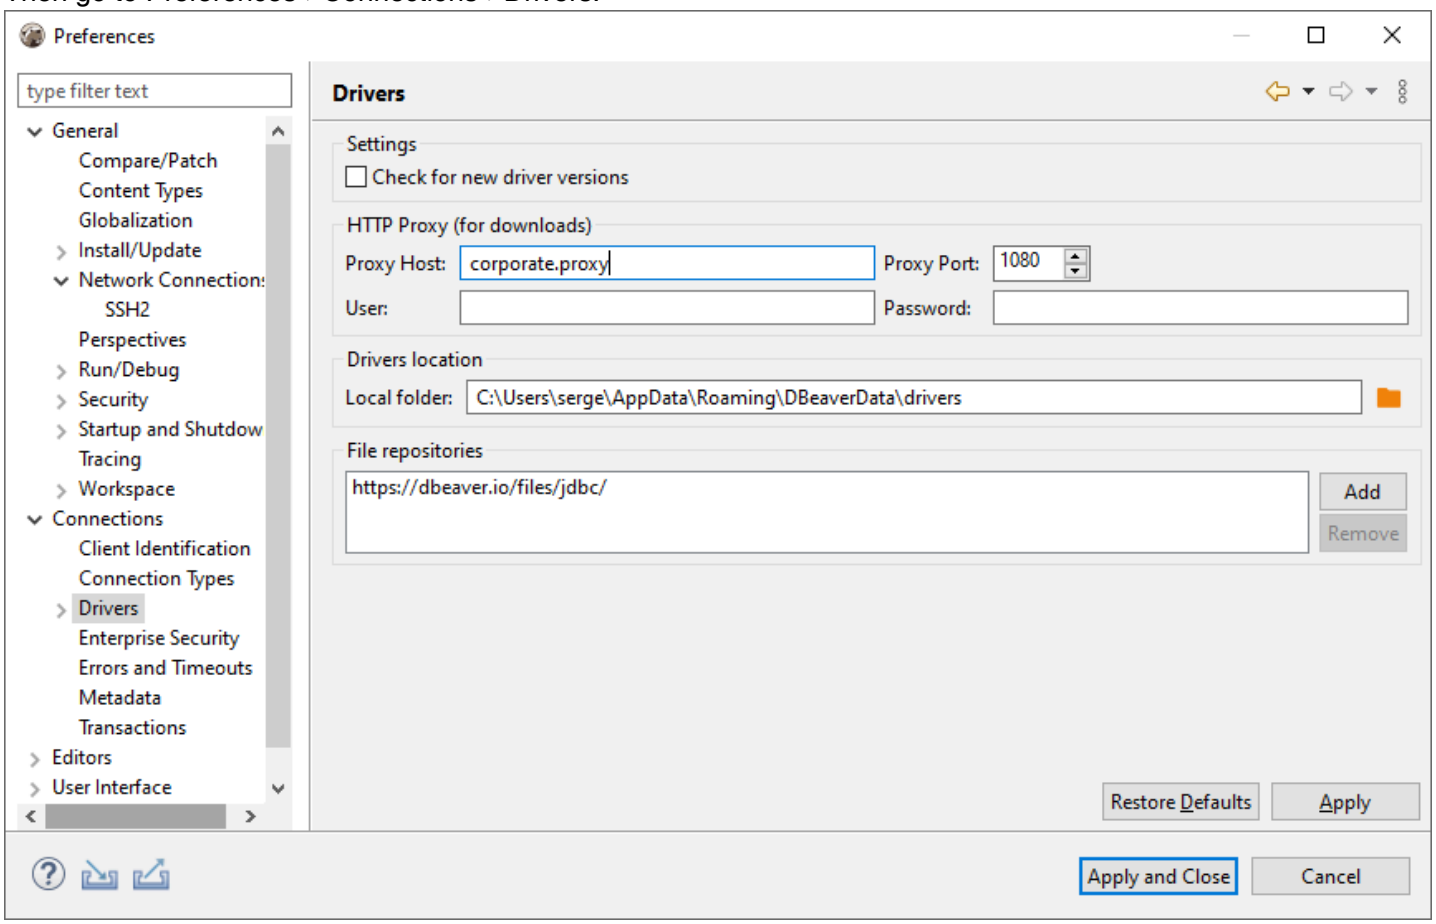

You can enter proxy host/port and (optionally) user/password here. It will be used for drivers download only.

Drivers are usually downloaded from maven.org web site. You also may ask your network admin to add `maven.org` to the list of allowed external domains.

How to configure network for license activation

You need to configure global proxy server.
If you can't activate your subscription license then you first need to use trial version to start DBeaver and configure proxy.

Go to Preferences->General->Network Connections:

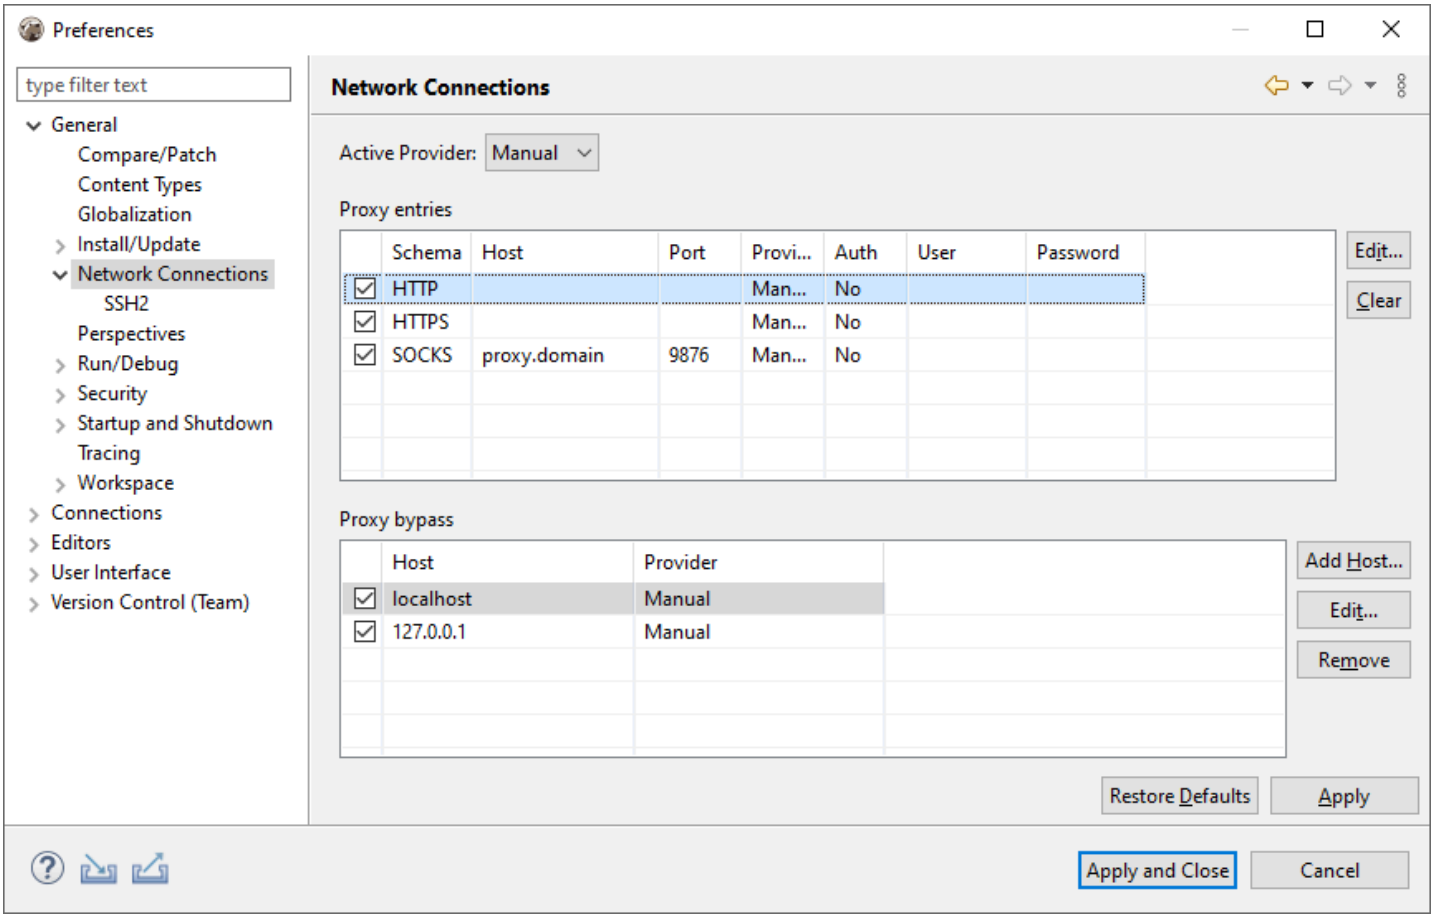

Switch to Manual or Native proxy (native proxy settings will use active web browser proxy configuration).
Note: in order to activate/update license DBeaver needs access to web site `dbeaver.com` only. You may ask your network administrator to add `dbeaver.com` to the white list.

How to configure proxy for external databases access

You can configure proxy settings for individual connection.
You may set proxy settings manually or use active OS/web browser settings:

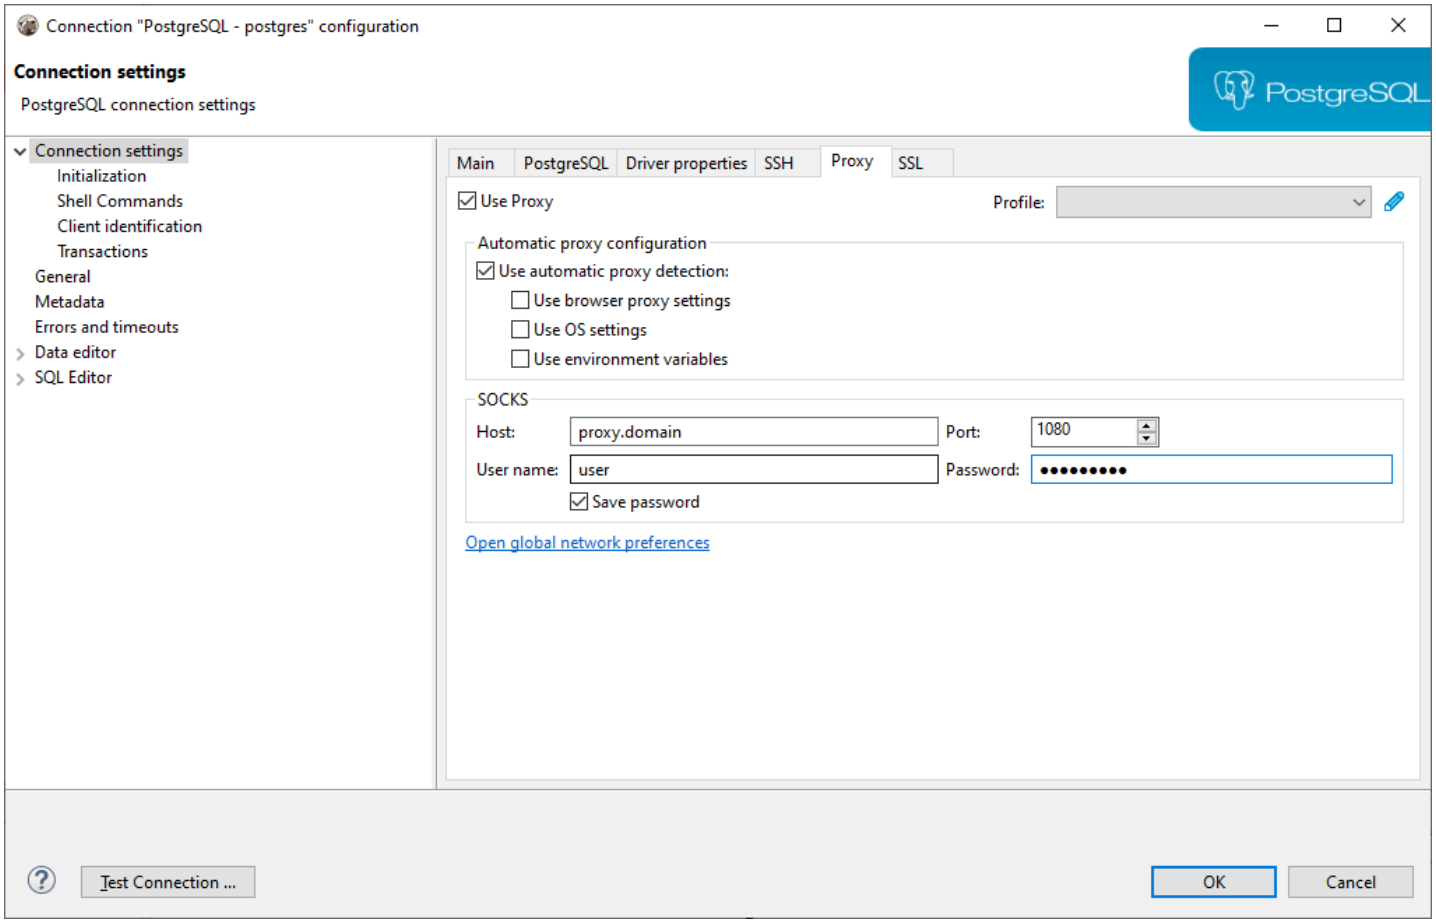
Connection "PostgreSQL - postgres" configuration
Connection settings
PostgreSQL connection settings
PostgreSQL
Connection settings
Initialization
Shell Commands
Client identification
Transactions
General
Metadata
Errors and timeouts
Data editor
SQL Editor
Main
PostgreSQL
Driver properties
SSH
Proxy
SSL
Use Proxy
Profile:
Automatic proxy configuration
Use automatic proxy detection:
Use browser proxy settings
Use OS settings
Use environment variables
SOCKS
Host:
proxy.domain
Port:
1080
User name:
user
Password:
Save password
Open global network preferences
Test Connection ...
OK
Cancel

Managing connections

This guide describes how to manage/secure DBeaver database connections. It is designed for System administrators. Regular users should check this guide.

Provide predefined connections

DBeaver keeps connections information in project folder. By default all projects reside in workspace.
Default project folder is workspace\workspace6\General.

DBeaver 6.1.3+

DBeaver keeps information about project connections in file `.dbeaver/data-sources.json`.
All secured information (user name, password, secret keys, etc) is stored in the encrypted file `.dbeaver/credentials-config.json`.

DBeaver can load multiple connection files. Any files in project folder matching `.dbeaver/data-sources*.json` pattern will be loaded on startup. So you can create a file, say, `.dbeaver/data-sources-2.json` in the project folder and DBeaver will see it.

DBeaver < 6.1.3

DBeaver keeps information about project connections in file `dbeaver-data-sources.xml`.

DBeaver can load multiple connection files. Any files in project folder matching `.dbeaver-data-sources*.xml` pattern will be loaded on startup. So you can create a file, say, `.dbeaver-data-sources-2.xml` in the project folder and DBeaver will see it.

Importing connections from CSV/XML

You can import connection from CSV or XML files.

CSV file must have a header row (first line of file) with column names (see list of supported columns below). XML file should contain top-level element and a set of nested elements. Connections config must be specified in attributes of nested elements. Attribute names are the same as CSV column names.

Supported names:

Name	Meaning
name	Connection name
url	JDBC URL
host	Database server host name
port	Database server port
database	Database/schema name

user	User name
password	User password

You can specify just URL or set host/port/etc setting.
User name/password are options.

Sample CSV

```
name,host,port,server,database,url,user,password,type
Postgre Import XML 1,localhost,5432,,postgres,jdbc:postgresql://localhost:5432/postgres,postgres,postgres,
Postgre Import XML 2,localhost,5432,,postgres2,jdbc:postgresql://localhost:5432/postgres2,postgres2,postgr
```

Sample XML

```
<connections>
        <connection name="Postgre Import XML 1" host="localhost" port="5432" server="" database="postgres"
        <connection name="Postgre Import XML 2" host="localhost" port="5432" server="" database="postgres"
</connections>
```

Secure connections from editing

It is possible to make connection settings read-only (protected by password)

- Generate MD5 hash of your password. You can do it from command line using Linux utility md5sum (`md5sum <<<"your password"`) or you can do it online - just google "MD5 hash online".
- Add field `lockPassword` in connection descriptor (in `.dbeaver/data-sources.json` in `connections` element. So it will look like this:

```
postgres-jdbc-161537836e8-3e0957d039995715": {
   "provider": "postgresql",
   "driver": "postgres-jdbc",
   "name": "PostgreSQL - postgres",
   "save-password": true,
   "show-system-objects": true,
   "read-only": false,
   "folder": "PG",
   "lockPassword": "2ba81a47c5512d9e23c435c1f29373cb"
...
}
```

- Now if user will try to change connection settings he/she will be asked for password

Managing drivers

Configure drivers with pre-installed jars

You can customize drivers configuration in workspace/.metadata/.plugins/org.jkiss.dbeaver.core/drivers.xml file. If you have some pre-installed jar files you can reference them in drivers.xml. Example:

```
<library type="jar" path="absolute-jar-folder-path\driver-jar.jar" custom="true" />
```

Also in drivers.xml you can use following variables to specify relative paths:

Variable	Meaning
drivers_home	Standard DBeaver drivers location - ($workspace/drivers by default)
dbeaver_home	DBeaver installation folder
home	User home folder
workspace	DBeaver workspace path

For instance:

```
<library type="jar" path="${workspace}\drivers\my-driver.jar" custom="true" />
```

Full drivers.xml example:

```
<?xml version="1.0" encoding="UTF-8"?>
<drivers>

        <provider id="postgresql">
                <driver id="postgres-jdbc" custom="false" embedded="false" name="PostgreSQL" class="org.po
                        <library type="jar" path="maven:/org.postgresql:postgresql:RELEASE" custom="false"
                                <file id="org.postgresql:postgresql" version="42.2.3" path="${drivers_home
                        </library>
                        <library type="jar" path="maven:/net.postgis:postgis-jdbc:RELEASE" custom="false"
                                <file id="net.postgis:postgis-jdbc" version="2.2.1" path="${drivers_home}/
                        </library>
                        <library type="jar" path="maven:/net.postgis:postgis-jdbc-jtsparser:RELEASE" custo
                                <file id="net.postgis:postgis-jdbc-jtsparser" version="2.2.1" path="${driv
                        </library>
                </driver>
        </provider>

</drivers>
```

Provide predefined drivers configuration

In some cases you may need to provide drivers configuration or driver jar files for a number of DBeaver installations automatically.
This can be done by adding special parameter in the `dbeaver.ini` file:

```
-Ddbeaver.drivers.configuration-file=c:\some-path\dbeaver-drivers-config.xml
```

This file has the same structure as `drivers.xml` file (see above) and it will be loaded *before* `drivers.xml`.

You can specify partial driver configuration. For example if you need to configure only jar path then it may look like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<drivers>
  <provider id="generic">
    <driver id="netezza">
      <library type="lib" path="X:\jdbc-drivers\netezza-jdbc.jar" />
    </driver>
  </provider>
</drivers>
```

Windows Silent Install

It is possible to install DBeaver in silent mode using Windows Installer command line parameters.
This might be very useful for mass install automation (SSCM and other similar systems).
Installer was improved in DBeaver 5.3.3, special thanks to https://github.com/Drizin/NsisMultiUser team.

Parameters

Command line parameters supported by DBeaver installer:

Parameter	Description
/S	silent mode, requires /allusers or /currentuser, case-sensitive
/D=path	(installer only) set install directory, must be last parameter, without quotes, case-sensitive
/allusers	(un)install for all users, case-insensitive
/currentuser	(un)install for current user only, case-insensitive
/uninstall	(installer only) run uninstaller, requires /allusers or /currentuser, case-insensitive

In order to install with /allusers parameter current user must have administrator permissions.

Installer return codes (decimal):

Code	Meaning
0	normal execution (no error)
1	(un)installation aborted by user (Cancel button)
2	(un)installation aborted by script
666660	invalid command-line parameters
666661	elevation is not allowed by defines
666662	uninstaller detected there's no installed version
666663	executing uninstaller from the installer failed
666666	cannot start elevated instance
other	Windows error code when trying to start elevated instance"

License Administration

Note: This functionality is available only in Enterprise Edition.

Manual license import

DBeaver EE asks user to import license file if this file cannot be found in the local license storage. For individual users this is the most simple and convenient way to import product license.

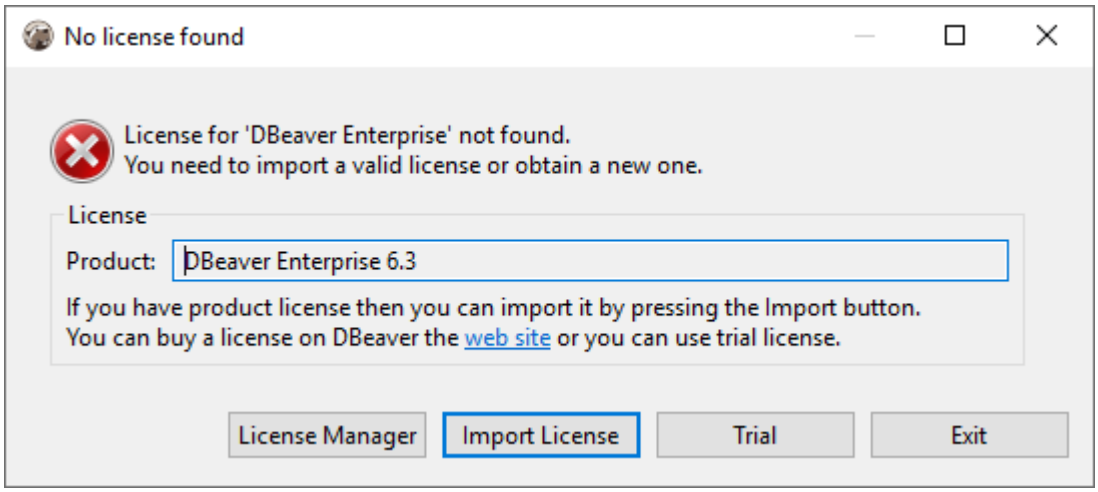

License management automation

There are several ways to automate license management process. It makes sense for multi-user environment.

Put the license file to the predefined locations

On start DBeaver will look for license file in the following locations:

- Windows
  - `%HOMEPATH%\.dbeaver-ee-license.dat`
  - `%APPDATA%\DBeaverData\workspace6\.metadata\.dbeaver-ee-license.dat`
- MacOS X
  - `~/.dbeaver-ee-license.dat`
  - `~/Library/DBeaverData/workspace6/.metadata/.dbeaver-ee-license.dat`
- Linux
  - `~/.dbeaver-ee-license.dat`
  - `$XDG_DATA_HOME/DBeaverData/workspace6/.metadata/.dbeaver-ee-license.dat`

Passing license file through command line

You can add command line parameter `license <license-path>` to the DBeaver EE shortcut. Also you can add this parameter to `dbeaver.ini` config file.

Command line reference.

How to Import License

- Import from email
- Import from the personal account
- Insert the License key to License Manager
- Import of Subscription license
- Import of License extension
- License Manager

To start using DBeaver EE you can

- request a Trial license for 2 weeks;
- request an Academic license if you are a student or a teacher;
- buy Subscription license, Standard DBeaver EE license or DBeaver EE license extension.

After purchase DBeaver EE license or getting Trial/Academic license you receive a License Key by email. Also, it is available in your personal account on our site. The License Key is a unique identification number that contains License ID e. g. DB-821MPZFO-ZA8W, Start date and license owner's name and company name. It is very important to import your License Key correctly.

Import from email

You can just copy-paste License Key to import the license into the License Manager. Please, note that you need to copy-paste full license text (not just license ID). The license text starts with "–" and ends with "==" characters.

Hello,

Thank you for your interest to DBeaver.

We are glad to see you among users of the enterprise version. Please, find your license below:

```
-- DBeaver EE LICENSE - DB-1Q1MQZDB-Z5K5
-- Starts at Tue Mar 02 07:46:42 UTC 2021 to DBeaver Corp. / Joe Test //
SpQZvwdkBF6fTHE4k5JKPTbZPpI/8wWqP+qAVA0ymn/E/mPihcGNSvQ9I75HqApTDTKzrX8Roldp
QfSn+janxCJJjUaSrX55gfiNWyMFWABMLJ+JZPS7v/0cc3f5j0LtcPK/ap66wZQHJ7gNCiBn2QcG
hJFyVh93PzlLenCvIMWabISPVlynJ6RANNa+DTaj5hbAdw6IVe7oh3uTHulf+irs8V/wJxEJaRC3
AQ1OrcotCUV1ijJ82Ea1Cmysx3Vnxqz8fPqP9Q1odjE+B5XNIgjLxwse3P5f7pG5u8Yko+B1WffL
q4eQ/af0e9VF1lqwfQ32t2GxMQ/KDzXE/FXisg==
```

Enjoy your database research with DBeaver and good luck!

Regards,

DBeaver team

Sometimes an email client can corrupt the formatting of the License Key that can cause an error.

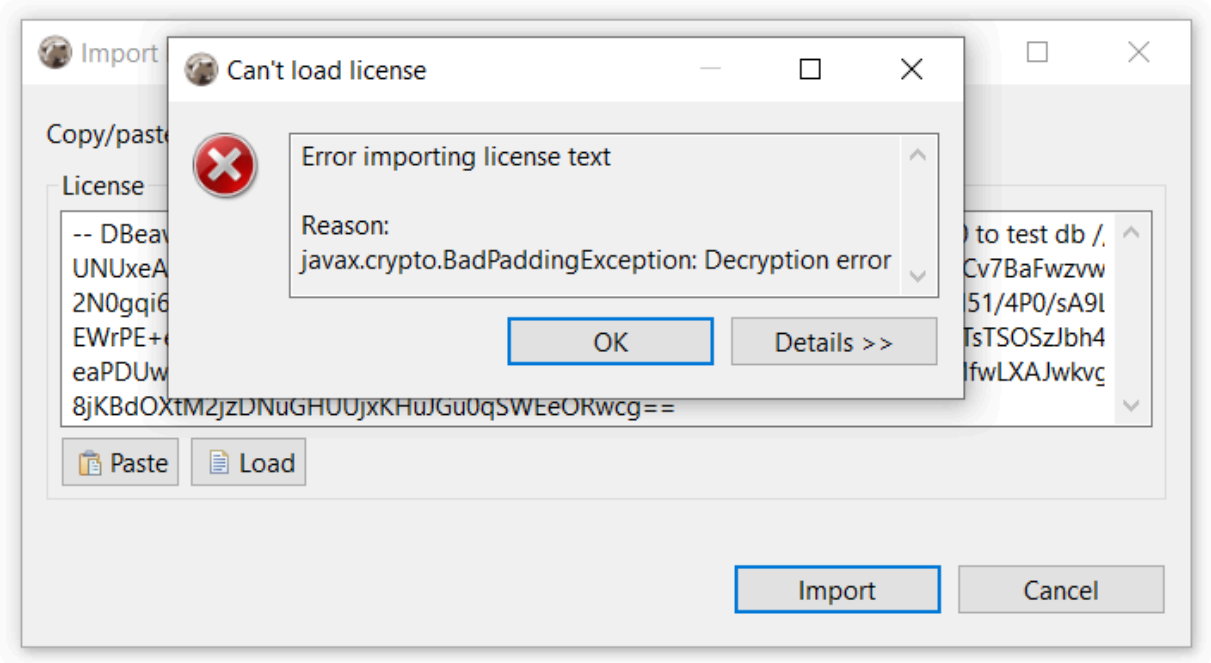

Therefore, you need to import your License Key from your personal account on our site https://dbeaver.com/.

Import from the personal account

Firstly, you need to Sign in.

Secondly, you should open the **Licenses** tab, where you can find all your licenses.

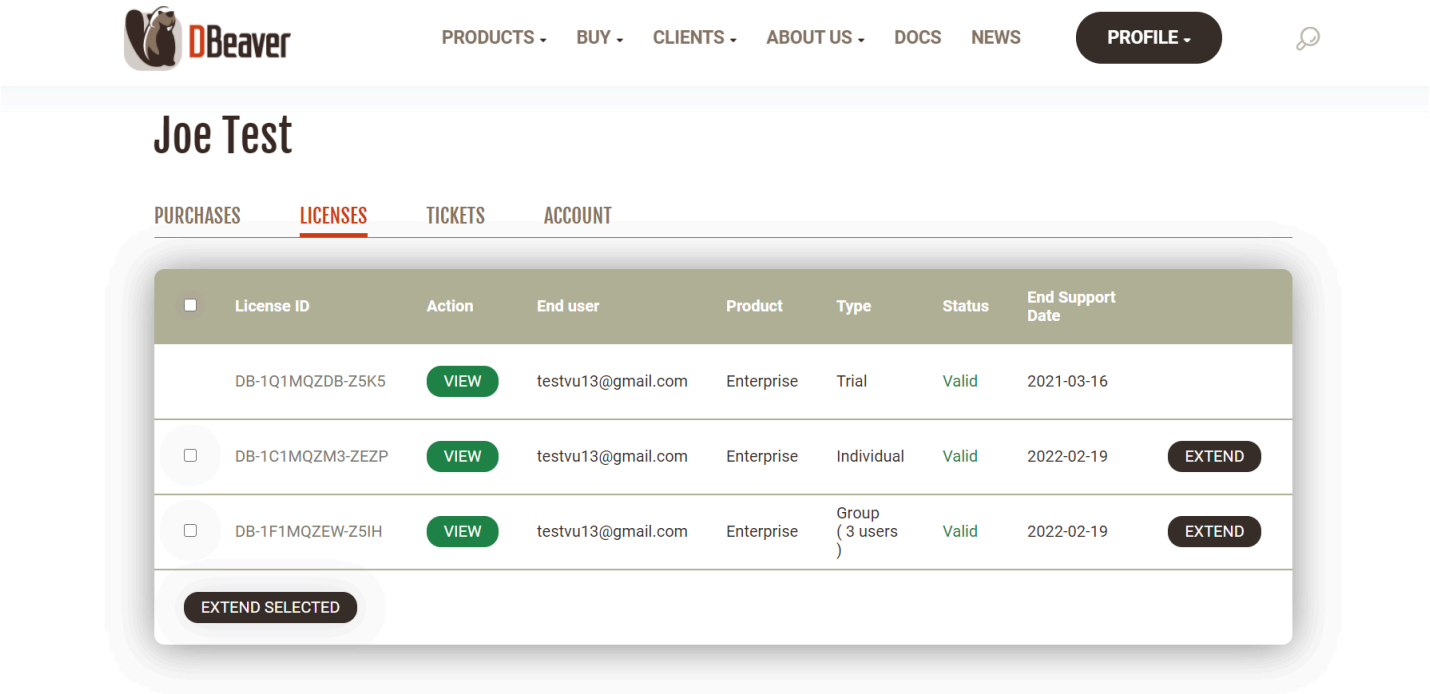

To open the License details and copy the License Key you press the **View** button. You can see the License details where you can find your license status, type, maintenance period, end support date. Also, you can reassign the license to another user.

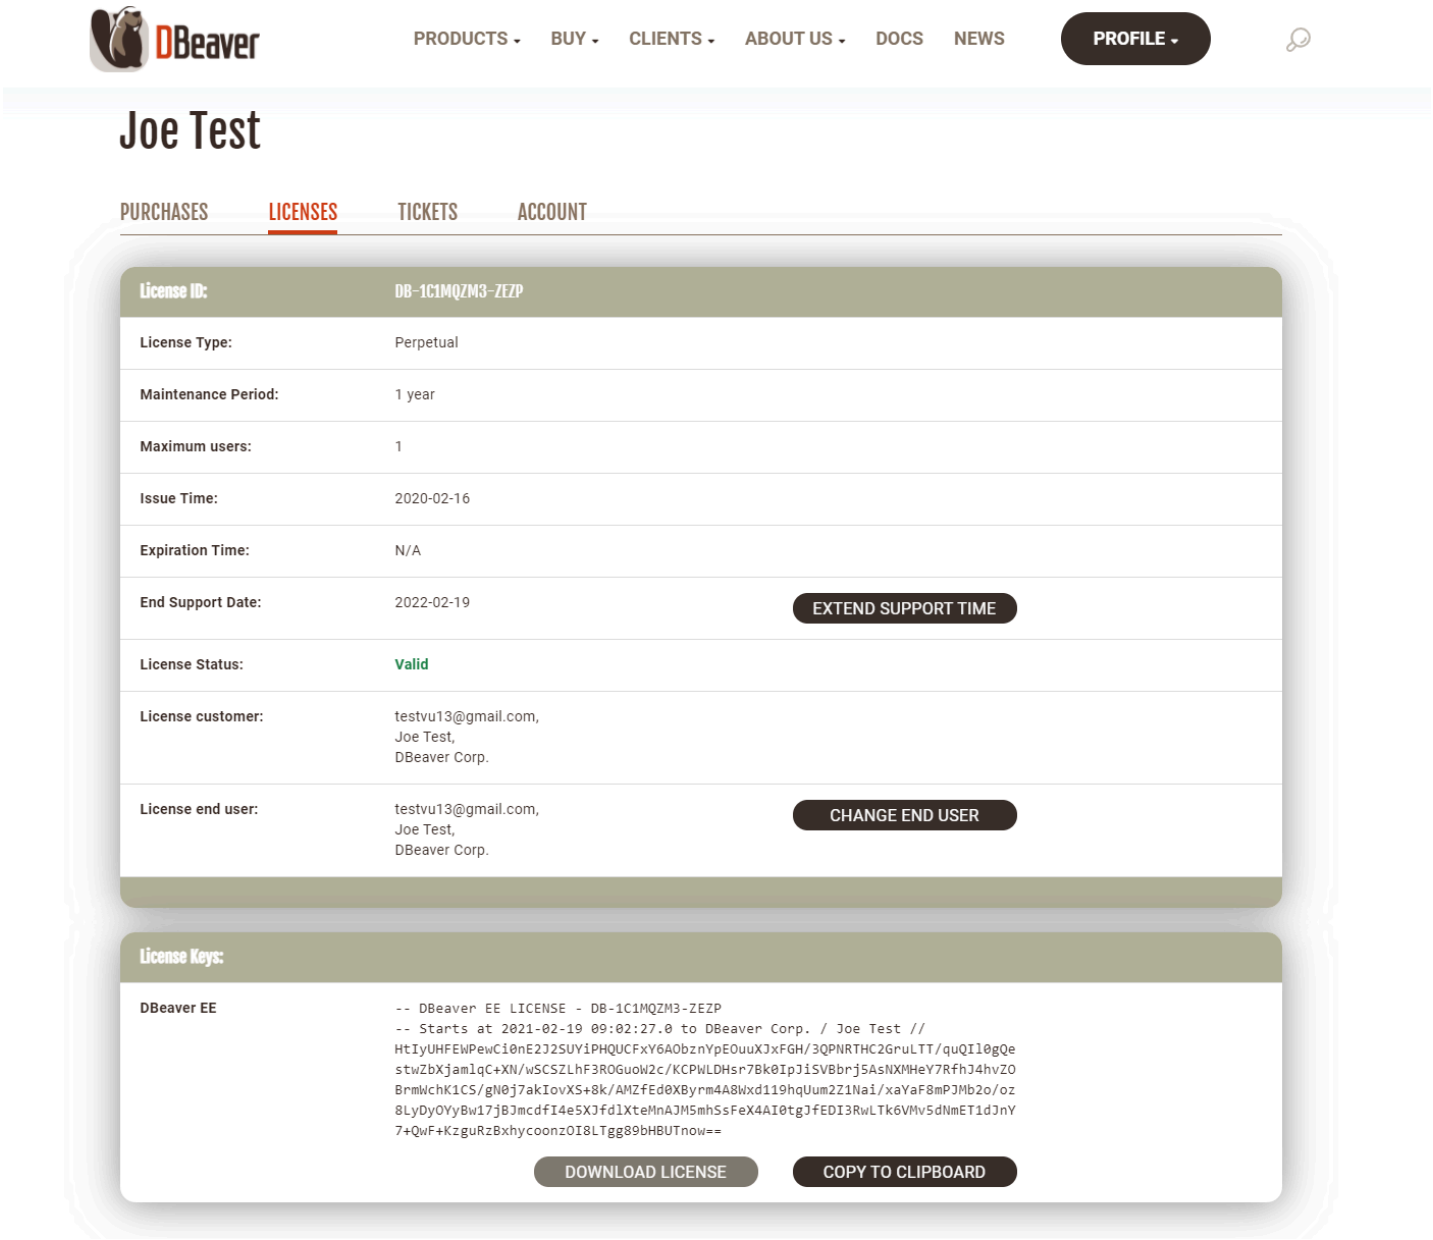

At the bottom of the page you can find the License Key required to start using DBeaver EE. There are two options how to copy your License Key from the personal account:

1)Press the **COPY TO CLIPBOARD** button, then press OK. The license text will be copied to clipboard.

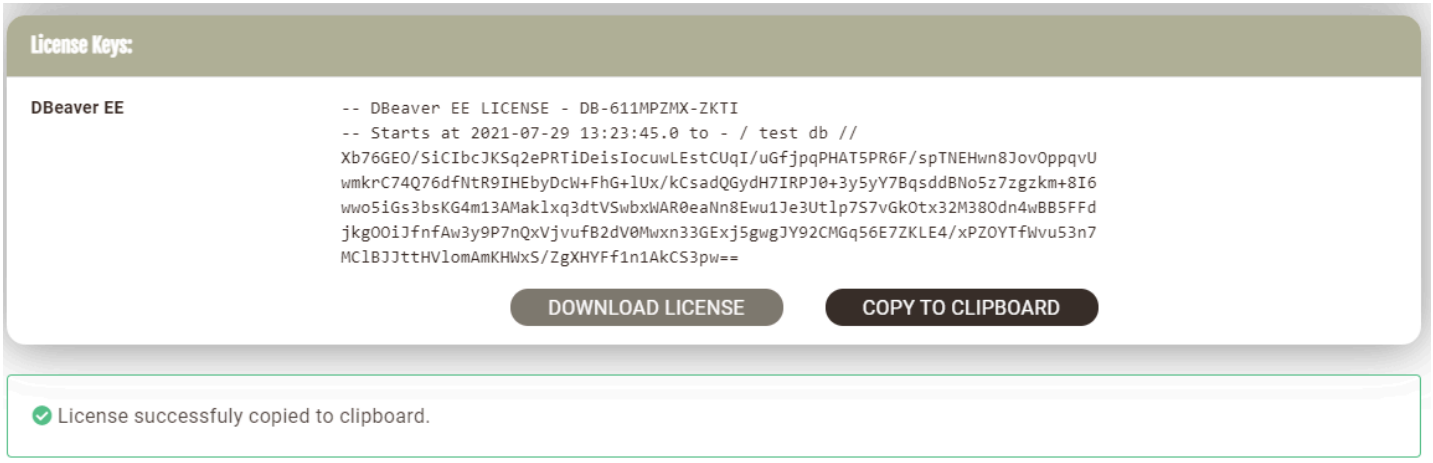

2)Press the **DOWNLOAD LICENSE** button, then press OK.

.txt file with your License Key will be downloaded to your download folder. The file name is License ID, e. g. DB-821MPZFO-ZA8W.

Then you need to insert copied License Key to License Manager in Dbeaver EE.

Insert the License key to License Manager

To start using DBeaver EE with your License Key you need to open License Manager in DBeaver EE: Help -> DBeaver License Info

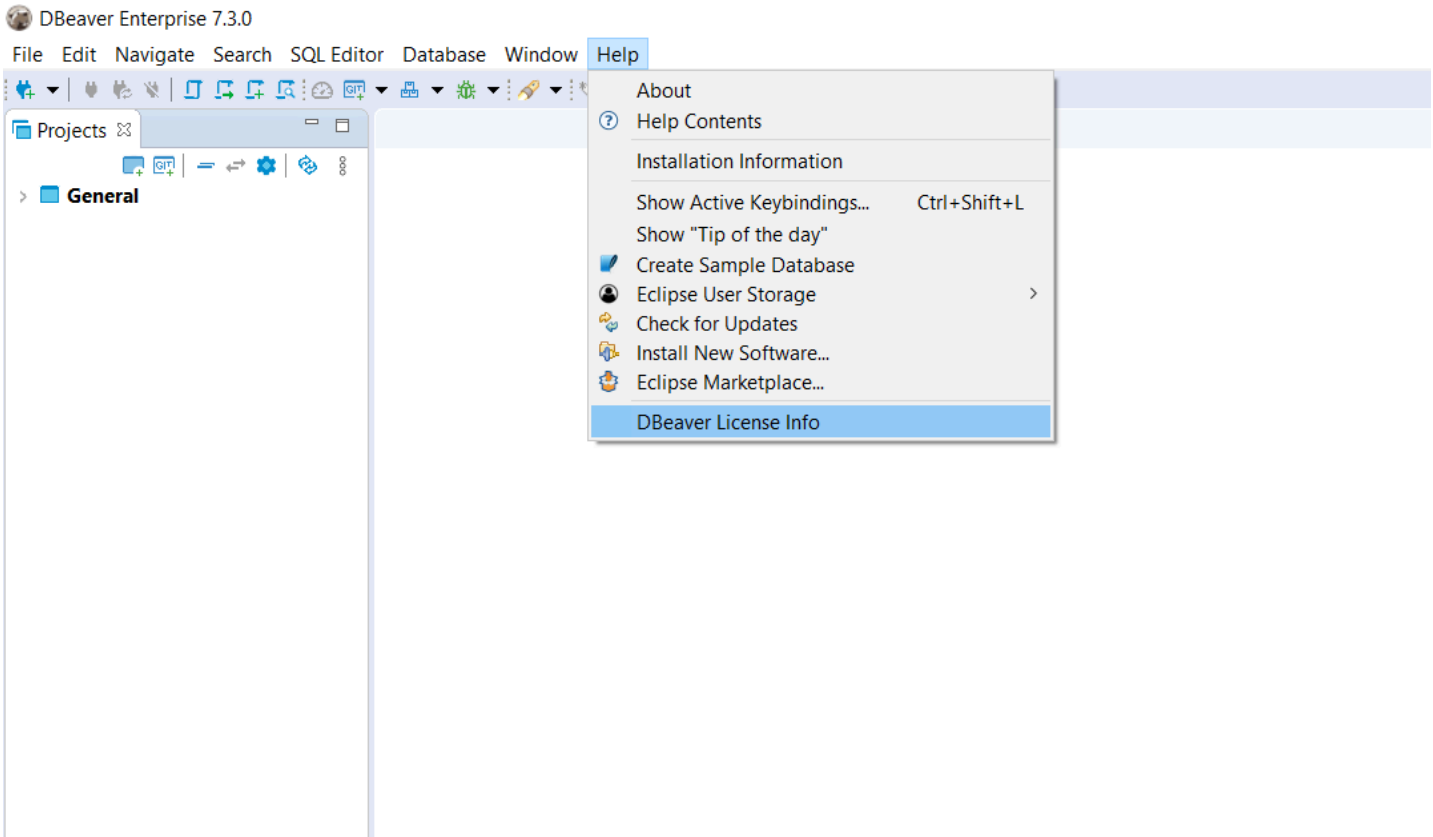

The License information window can look different depending on whether have you already a valid license or not.

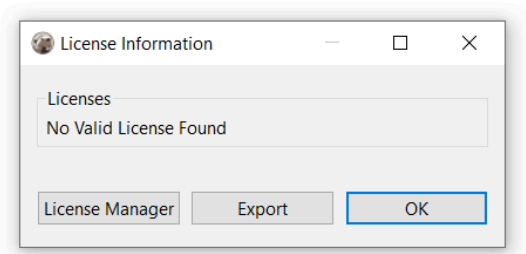

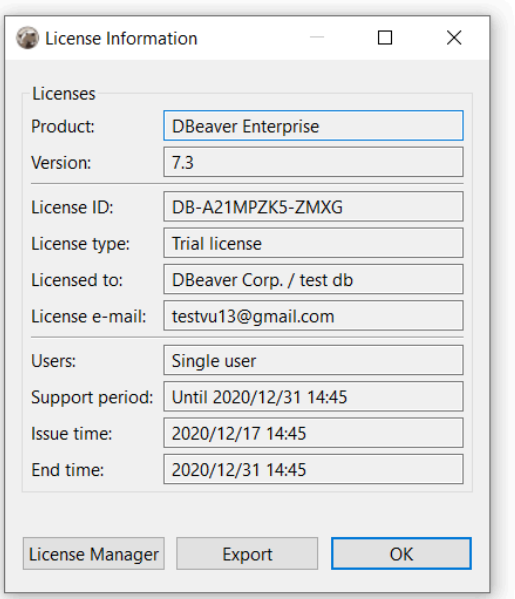

Then you click **License Manager** where you can find all your imported licenses and information about them.

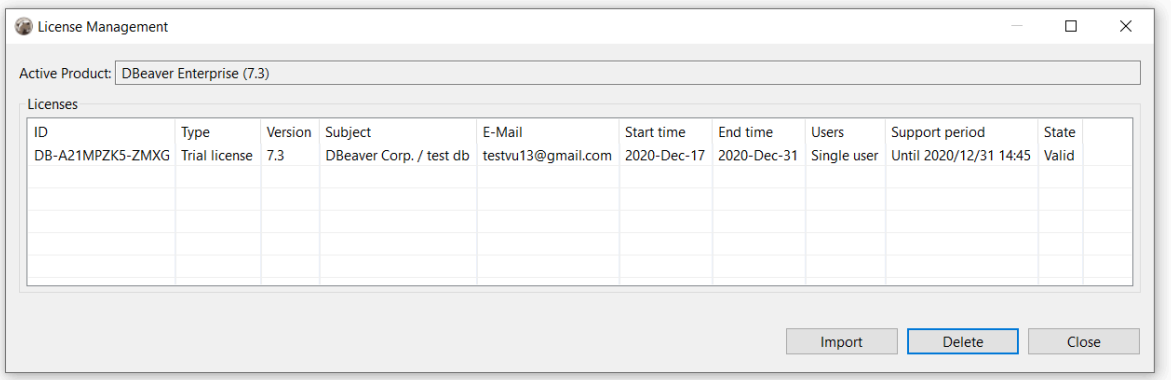

Next you press the **Import** button to paste your License Key.

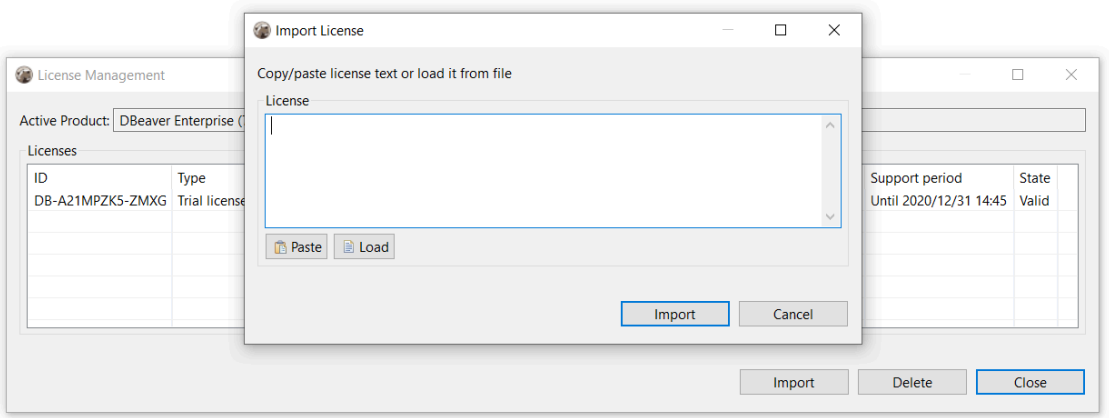

If you copied the License Key to clipboard, press the **Paste** button and then **Import**.

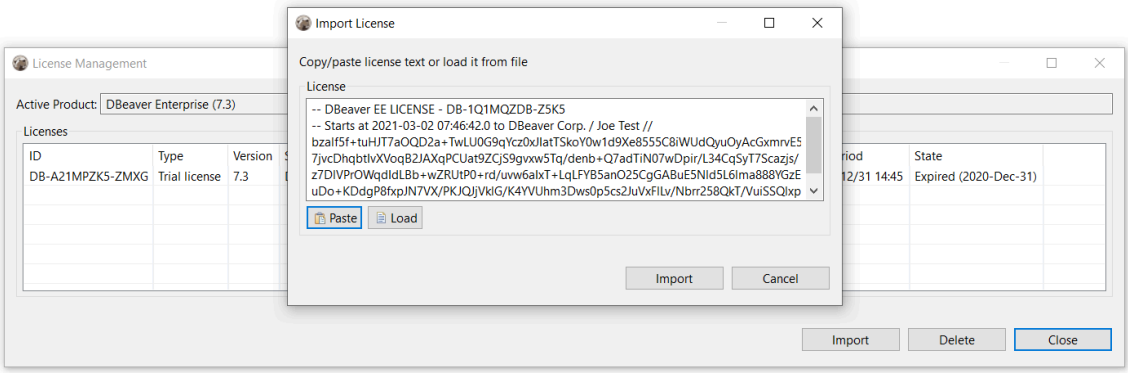

If you downloaded the .txt file with the License Key, press the **Load** button and then select the file from Downloads. The License Key will be pasted automatically.

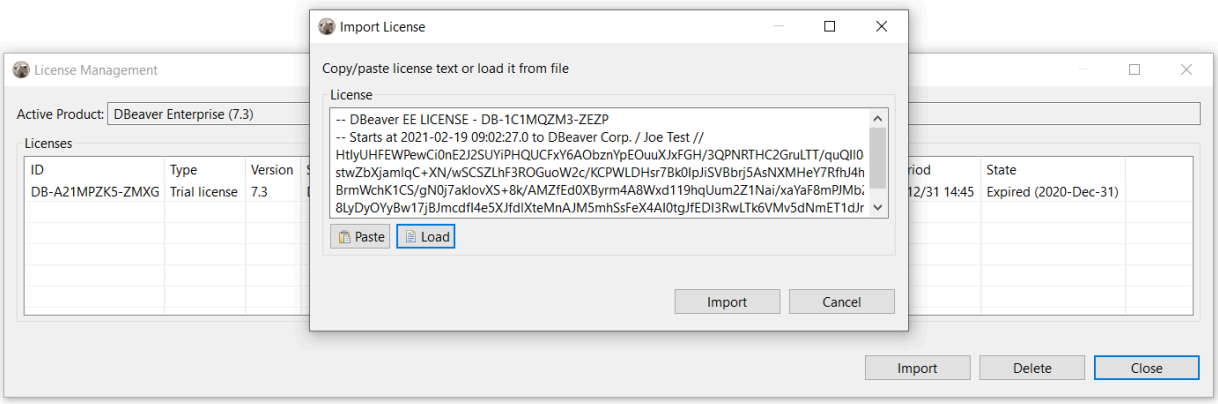

Then press the **Import** button and your license will be added in License Manager.

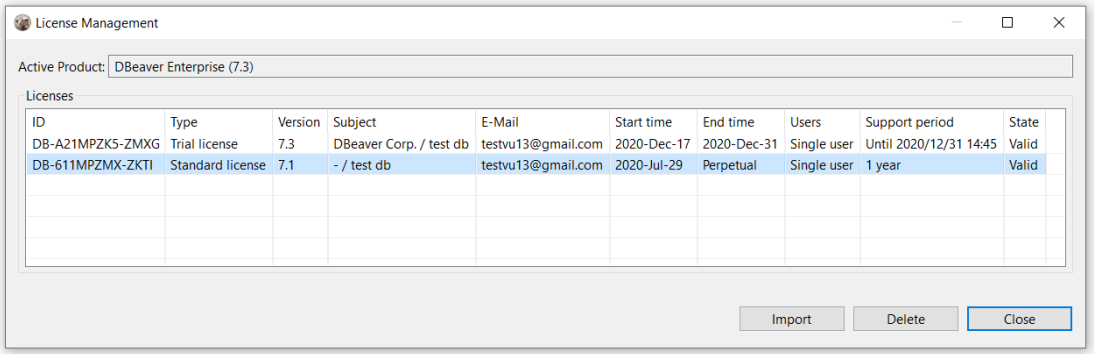

You have successfully imported your license. Now you can close License Manager and start using Dbeaver EE.

Import of Subscription license

Subscription license requires an internet access on the workstation for the first activation and each prolongation.

If you do not have an active internet connection or work behind a corporate firewall while importing the Subscription license, the following error can appear:

Invalid subscription

Can`t find subscription information for license 'DB-821MPZFO-ZA8W'.

Check your internet connection and/or firewall settings and restart application.

In this case you need to check that DBeaver EE has an internet access or to configure firewall.

License extension

Standard DBeaver EE license is a perpetual license with a limited period of support (1 year or 2 years).

After the end of the selected support period you can continue to use DBeaver Enterprise Edition without support and updates or buy a license extension or a new license.

If you buy DBeaver EE license extension and Dbeaver has an internet access, the license in DBeaver EE will be updated automatically. Otherwise, you have to import the license key from the personal account once again.

License Manager

License Manager provides you with the following information about your licenses:

- License ID e. g. DB-821MPZFO-ZA8W;
- License type: Trial/Academic/Subscription/Standard;
- Version;
- License owner`s name and company name;
- License owner`s email;
- Start time is the date the license was received;
- End time is the date of license expiration (standard perpetual licenses don`t have it)
- Number of users: single user or multiuser for group licenses;
- Support period is a period while you have access to the internal support system on the site and a possibility to download new product versions;
- State: valid or expired.

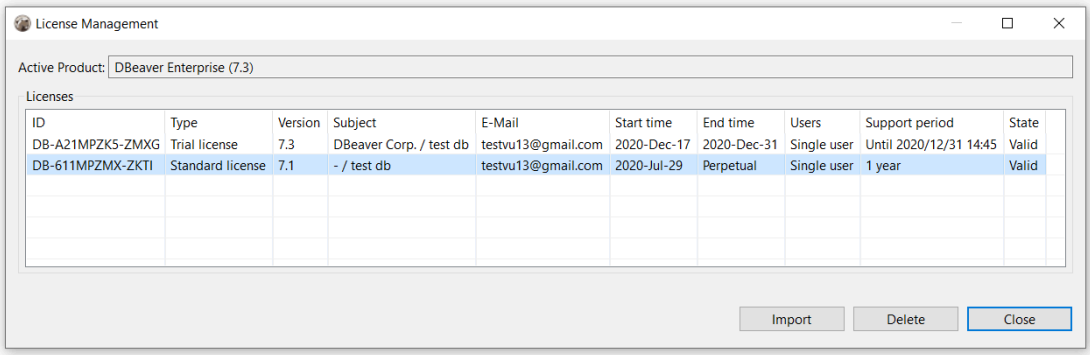
License Management
Active Product: DBeaver Enterprise (7.3)
Licenses
ID Type Version Subject E-Mail Start time End time Users Support period State
DB-A21MPZK5-ZMXG Trial license 7.3 DBeaver Corp. / test db testvu13@gmail.com 2020-Dec-17 2020-Dec-31 Single user Until 2020/12/31 14:45 Valid
DB-611MPZMX-ZKTI Standard license 7.1 - / test db testvu13@gmail.com 2020-Jul-29 Perpetual Single user 1 year Valid
Import
Delete
Close

How to Reassign License

After purchase a bunch of DBeaver EE licenses you have to assign each license to an end user.

If an employee subsequently is leaving the company or the team that is using DBeaver, the license admin may need to reassign the license to another employee.

You can reassign the license to another user in your personal account.

Firstly, you need to Sign in.

Secondly, you should open the **Licenses tab**, where you can find all your licenses.

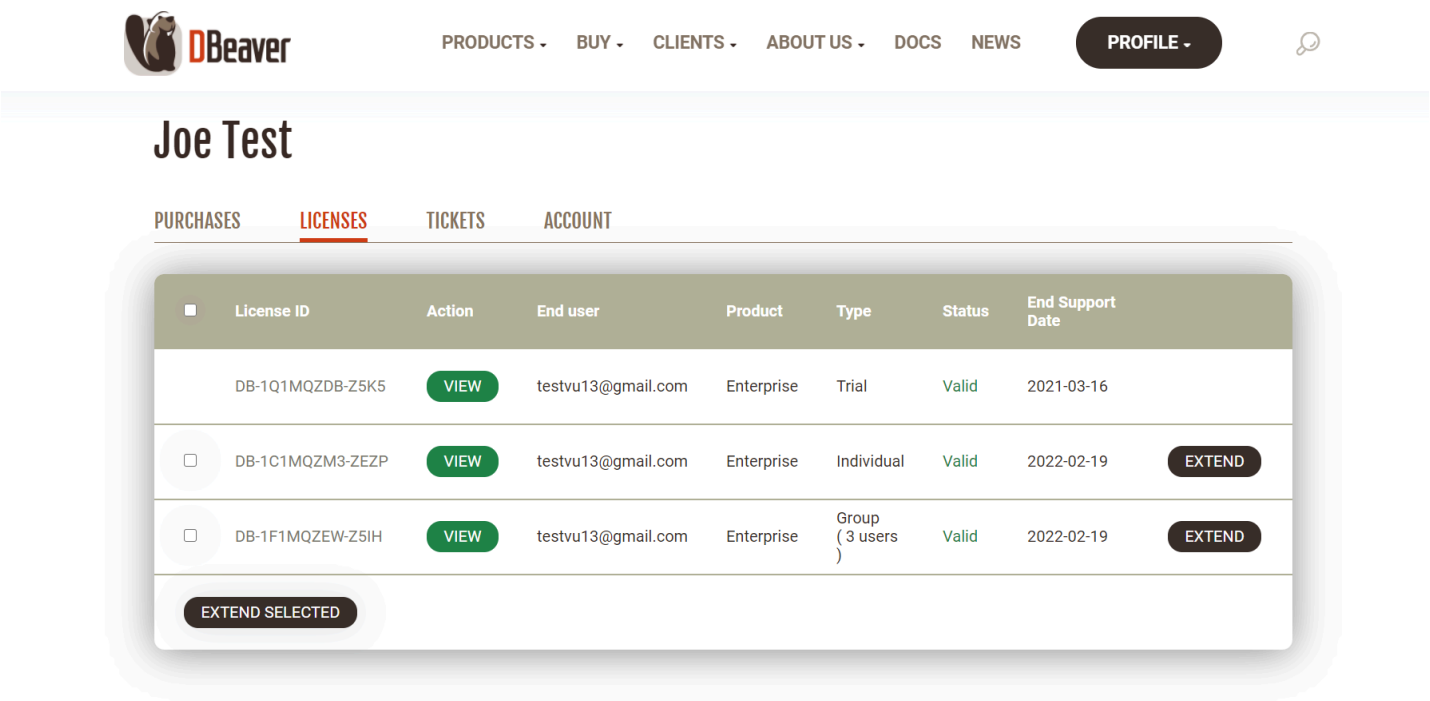

You select which license you need to reassign and press the **VIEW** button near its license ID. You can see the License details where you can find your license status, type, maintenance period, end support date.

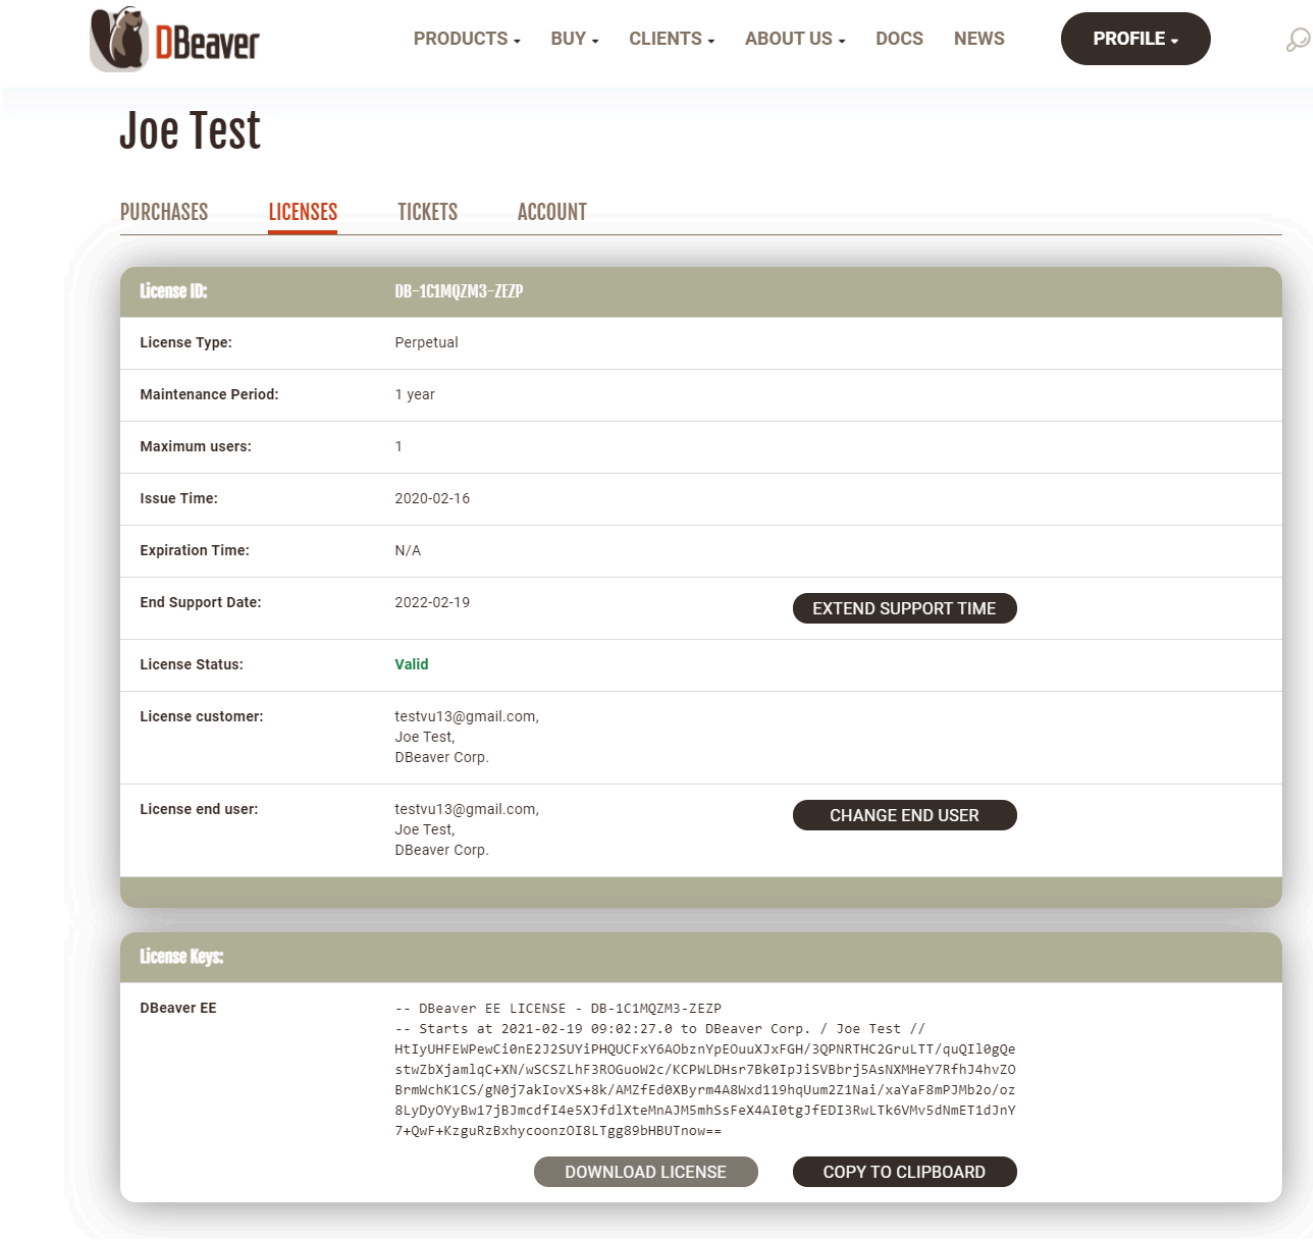

Also, there is the license end user field that contains license end user's details. To reassign the license, you need to press the **CHANGE END USER** button near this field.

A pop-up opens, and you can enter new license end user data: email, first and last name, company. After filling the form, you need to press the **SAVE** button and the license owner will be changed. The license key that contains license end user's name and company name will be changed too.

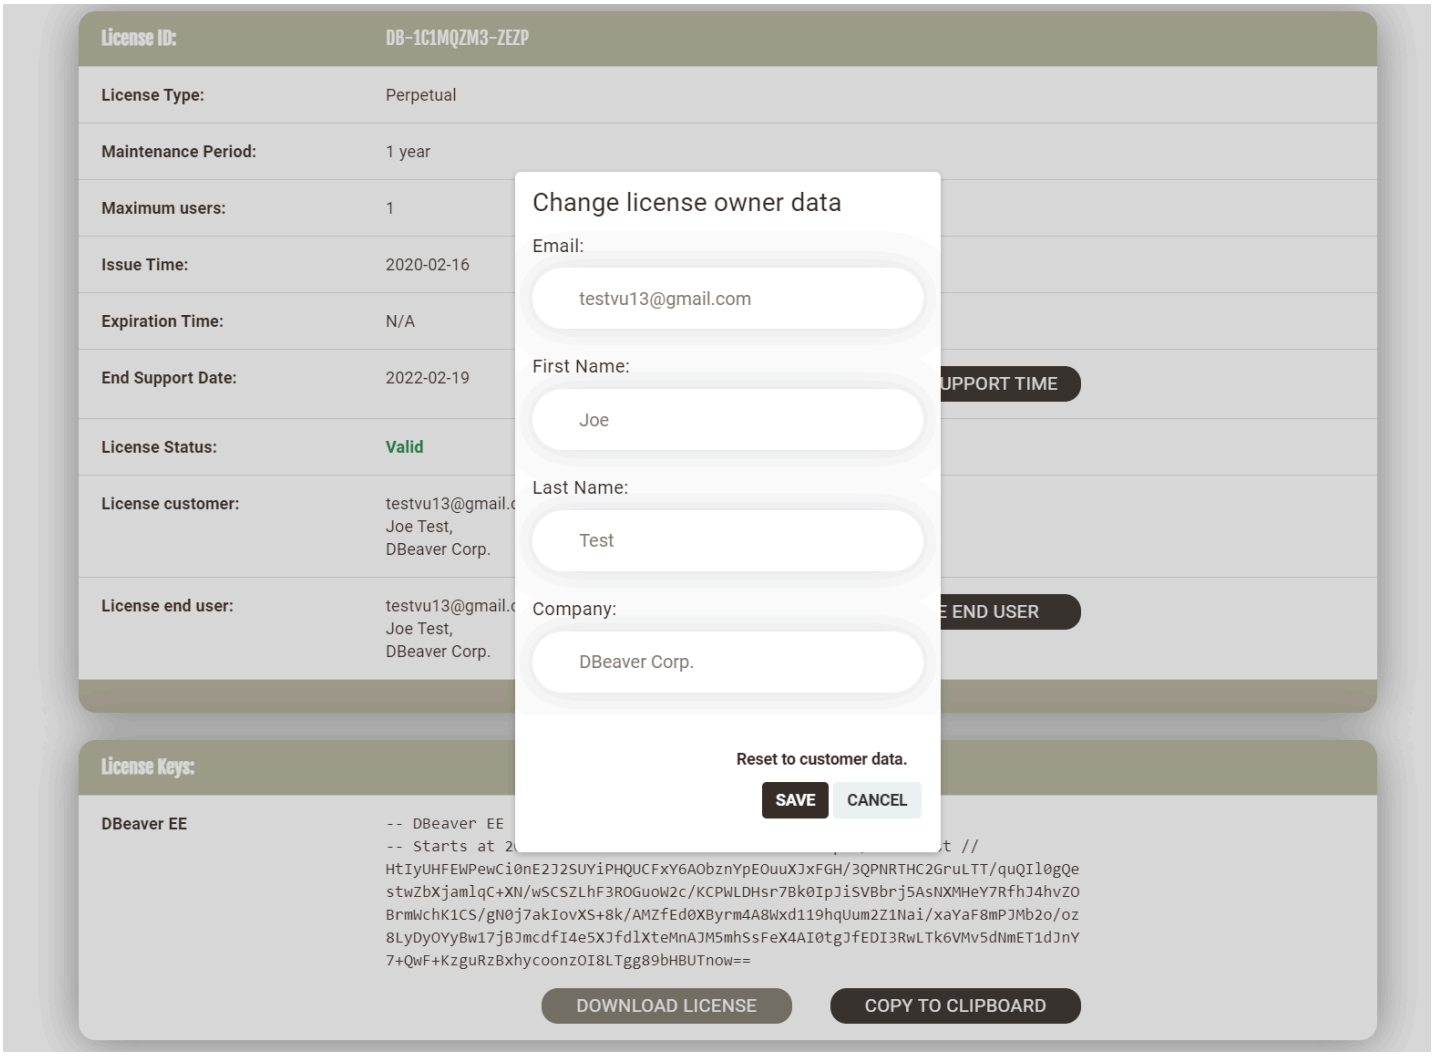

The license will remain in your personal account. The end user's email will be in the end user column on the Licenses tab.

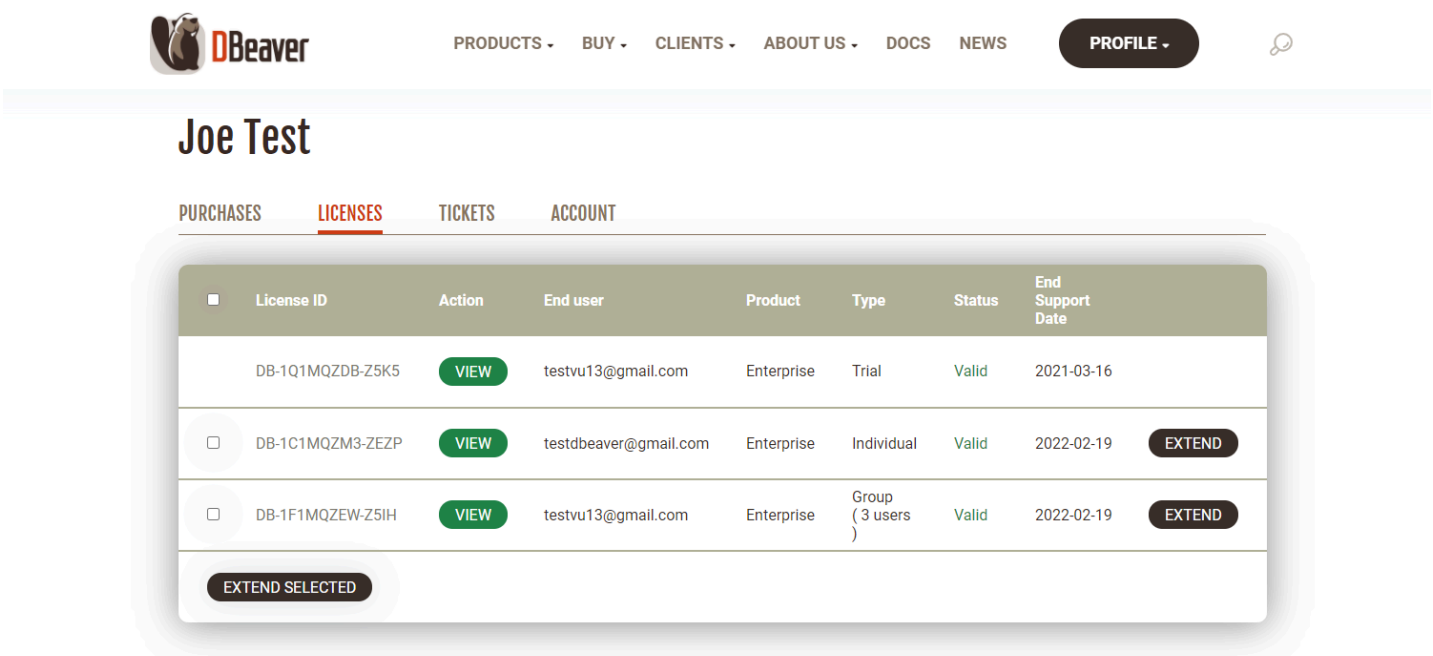

Also, new license end user will be able to find the license in the personal account and to import it. There will be no the CHANGE END USER button since only the license customer can assign a license to the end user. If new license end user hasn't been signed up on our website, the account will be created automatically. New user will receive a welcome email with a link to set up a password.