TIBCO Data Virtualization®

Web UI - User Guide

Version 8.4

Last Updated: *March 9, 2021*

Two-Second Advantage®

TIBCO®

Important Information

SOME TIBCO SOFTWARE EMBEDS OR BUNDLES OTHER TIBCO SOFTWARE. USE OF SUCH EMBEDDED OR BUNDLED TIBCO SOFTWARE IS SOLELY TO ENABLE THE FUNCTIONALITY (OR PROVIDE LIMITED ADD-ON FUNCTIONALITY) OF THE LICENSED TIBCO SOFTWARE. THE EMBEDDED OR BUNDLED SOFTWARE IS NOT LICENSED TO BE USED OR ACCESSED BY ANY OTHER TIBCO SOFTWARE OR FOR ANY OTHER

PURPOSE.

USE OF TIBCO SOFTWARE AND THIS DOCUMENTATION IS SUBJECT TO THE TERMS AND CONDITIONS OF A LICENSE AGREEMENT FOUND IN EITHER A SEPARATELY EXECUTED SOFTWARE LICENSE AGREEMENT, OR, IF THERE IS NO SUCH SEPARATE AGREEMENT, THE CLICKWRAP END USER LICENSE AGREEMENT WHICH IS DISPLAYED DURING DOWNLOAD OR INSTALLATION OF THE SOFTWARE (AND WHICH IS DUPLICATED IN THE LICENSE FILE) OR IF THERE IS NO SUCH SOFTWARE LICENSE AGREEMENT OR CLICKWRAP END USER LICENSE AGREEMENT, THE LICENSE(S) LOCATED IN THE "LICENSE" FILE(S) OF THE SOFTWARE. USE OF THIS DOCUMENTATION IS SUBJECT TO THOSE TERMS AND CONDITIONS, AND YOUR USE HEREOF SHALL CONSTITUTE ACCEPTANCE OF AND AN AGREEMENT TO BE BOUND BY THE SAME.

TIBCO and the TIBCO logo are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or other countries

TIBCO, Two-Second Advantage, TIBCO Spotfire, TIBCO ActiveSpaces, TIBCO Spotfire Developer, TIBCO EMS, TIBCO Spotfire Automation Services, TIBCO Enterprise Runtime for R, TIBCO Spotfire Server, TIBCO Spotfire Web Player, TIBCO Spotfire Statistics Services, S-PLUS, and TIBCO Spotfire S+ are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or other countries.

All other product and company names and marks mentioned in this document are the property of their respective owners and are mentioned for identification purposes only.

THIS SOFTWARE MAY BE AVAILABLE ON MULTIPLE OPERATING SYSTEMS. HOWEVER, NOT ALL OPERATING SYSTEM PLATFORMS FOR A SPECIFIC SOFTWARE VERSION ARE RELEASED AT THE SAME TIME. SEE THE README FILE FOR THE AVAILABILITY OF THIS SOFTWARE VERSION ON A SPECIFIC OPERATING SYSTEM PLATFORM.

THIS DOCUMENTATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS DOCUMENTATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THIS DOCUMENTATION. TIBCO SOFTWARE INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENTATION AT ANY TIME.

Contents

Preface

Documentation for this and other TIBCO products is available on the TIBCO Documentation site. This site is updated more frequently than any documentation that might be included with the product. To ensure that you are accessing the latest available help topics, please visit:

- https://docs.tibco.com

Product-Specific Documentation

The following documents form the TIBCO® Data Virtualization(TDV) documentation set:

- Users
  - TDV Getting Started Guide
  - TDV User Guide
  - TDV Client Interfaces Guide
  - TDV Tutorial Guide
  - TDV Northbay Example
- Administration
  - TDV Installation and Upgrade Guide
  - TDV Administration Guide
  - TDV Active Cluster Guide
  - TDV Security Features Guide
- Data Sources
  - TDV Adapter Guides
  - TDV Data Source Toolkit Guide (Formerly Extensibility Guide)
- References
  - TDV Reference Guide
  - TDV Application Programming Interface Guide

- **Other**

  TDV Business Directory Guide

  TDV Discovery Guide

- *TIBCO TDV and Business Directory Release Notes* Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.

How to Access TIBCO Documentation

Documentation for TIBCO products is available on the TIBCO Product Documentation website mainly in the HTML and PDF formats.

The TIBCO Product Documentation website is updated frequently and is more current than any other documentation included with the product. To access the latest documentation, visit https://docs.tibco.com.

Documentation for TIBCO Data Virtualization is available on https://docs.tibco.com/products/tibco-data-virtualization-server.

How to Contact TIBCO Support

You can contact TIBCO Support in the following ways:

- For an overview of TIBCO Support, visit https://www.tibco.com/services/support.
- For accessing the Support Knowledge Base and getting personalized content about products you are interested in, visit the TIBCO Support portal at https://support.tibco.com.
- For creating a Support case, you must have a valid maintenance or support contract with TIBCO. You also need a user name and password to log in to https://support.tibco.com. If you do not have a user name, you can request one by clicking **Register** on the website.

How to Join TIBCO Community

TIBCO Community is the official channel for TIBCO customers, partners, and employee subject matter experts to share and access their collective experience. TIBCO Community offers access to Q&A forums, product wikis, and best practices. It also offers access to extensions, adapters, solution accelerators, and tools that extend and enable customers to gain full value from TIBCO products. In addition, users can submit and vote on feature requests from within the TIBCO Ideas Portal. For a free registration, go to https://community.tibco.com.

Introduction

The TIBCO Data Virtualization Web UI is a Self-Service data provisioning web user interface that enables you to create and publish your own views that can then be consumed in your 3rd party downstream apps.

Using the easy to use web interface, the data engineers, TDV developers and business analysts can easily search for available datasets in the Data Catalog.

Self-Service users can then create datasets, perform complex SQL queries to manipulate data, export and publish the result set. The complex data processing workflows created are reliable, repeatable and secure. The easy to use graphical interface's drag and drop feature lets you do all these with limited or no knowledge of SQL.

Following topics are covered in this chapter:

System Requirements

You will need to install and run TDV Server 8.4.0 in order to access the Web UI. Refer to the *TDV Installation and Upgrade Guide* for the system requirements and for instructions on how to install TDV Server.

How to Access TDV Web UI

In order to access the TDV Web UI, you need to install TDV Server. Refer to the *TDV Installation and Upgrade Guide* for instructions on how to install TDV Server. Once the TDV Server is installed and running, the TDV Web UI can be accessed at:

http://<server>:<port>/webui/login

If you are running TDV Server locally, then you can access it at:

http://localhost:9400/webui/login

Accessing from TDV Studio

To edit a flow created using the Web UI, from TDV Studio, right-click on the flow and choose the option "Edit in Web UI".

Logging in to TDV Web UI

User Management

You need "Access Tools" user rights to access the TDV Web UI. Use the Web Manager tool to add a new user or modify an existing user's privileges. Refer the *TDV Administration Guide*, chapter *Composite Domain Administration* for details about user management.

The Login Screen

The following Login screen is displayed when:

1. You access the url for TDV Web UI (http://<hostname>:9400/webui).
2. You are in a different screen within the TDV Web UI and you have been logged out due to inactivity.
3. You choose to sign out of TDV Web UI.

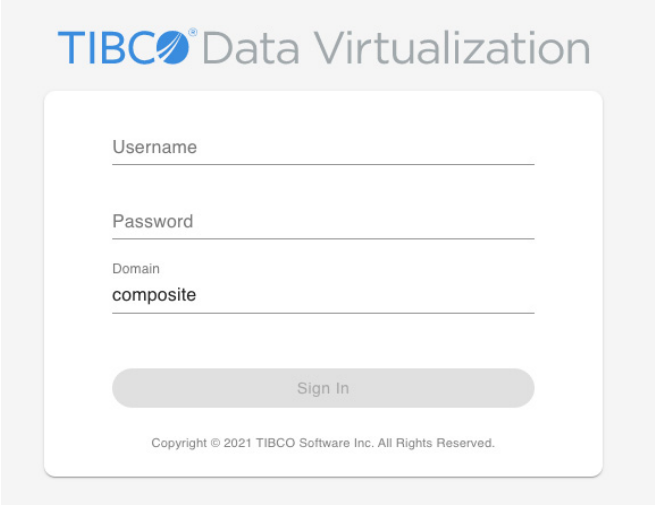

Follow these instructions to login to TDV Web UI:

1. Provide the valid credentials for all the fields in the Login screen.

The table below describes the various fields in the Login screen.

2. Click on Sign in to Login to TDV Web UI.

Field	Description
Username	The username required to login to TDV Web UI. Obtain a valid username from your TDV administrator.
Password	The password required to login to TDV Web UI. Obtain a valid password from your TDV administrator.
Domain	The domain to which the user belongs.

3. The following Web UI Home page is displayed.

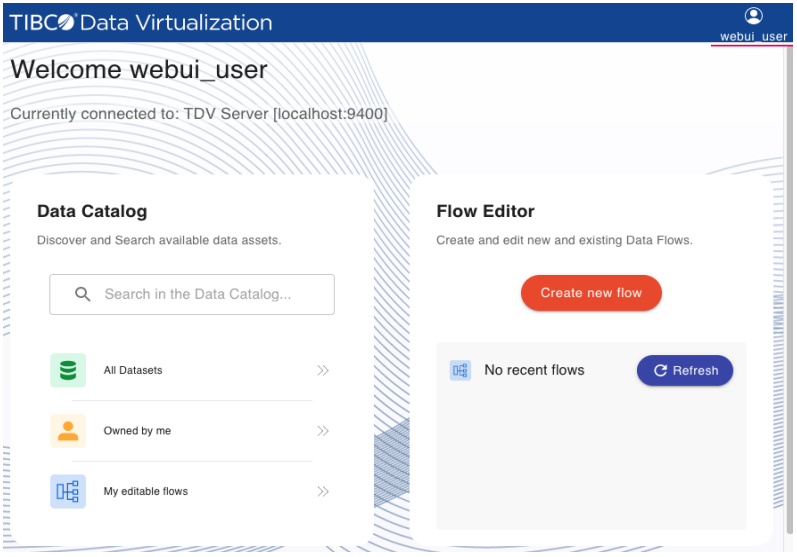

You can either go to the Data Catalog to view or edit an existing flow or create your own flow in the Flow Editor, by clicking on the Create New Flow button.

Overview of the Data Catalog

The Data Catalog lets you browse all the published and non-published TDV tables and views as well as all the tabular resources created using the TDV Web UI. Any resource created within the Web UI can be edited/deleted using the Data Catalog.

The search tool in the Data Catalog lets you search for the resources.

The Resource preview pane in the bottom of the page, gives you the ability to browse through the contents of the tabular data sets.

Refer to the Chapter Data Catalog, page 13 for more information.

Overview of the Flow Editor

The Web UI Flow Editor lets you create and edit data flow diagrams. The dataset node and the operations that can be performed on the dataset can be dragged from the Operations Palette and dropped in the canvas. Connections between the nodes can then be made by gently dragging one node towards the other. Once the data flow diagram is complete, the flow can be published for use in Studio and other applications outside the Web UI.

The Operations Palette displays various operations such as Filter, Join, Union, Except, Intersect. Using drag and drop feature, users can build the data flow diagram on the Canvas.

The Expression Builder lets you manipulate and build the query. Syntax help is provided for every function available in the expression builder.

The Details pane that is displayed below the Canvas provides the user with the list of columns to show/hide, an expression builder to manipulate the data, the result set data with ability to scroll through, and a read-only query view.

Refer to the Chapter Flow Editor, page 19 for more information.

Navigation

The Graphical UI of TDV lets you access any of the pages directly through a Url:

Home Page	http://<hostname>:9400/webui/home
Data Catalog:	All Datasets - http://<hostname>:9400/webui/home Owned by you - http://<hostname>:9400/webui/directoryBrowser Editable Views - http://<hostname>:9400/webui/directoryBrowser?filter=editableViews
Flow Editor	http://<hostname>:9400/webui/view?newFlow=true

While directly accessing the Urls, if there has been no activity, you are automatically taken to the Login screen for re-authentication. Once logged in, you are transferred back to the page you were originally in.

You can also use the browser's Forward and Back buttons to navigate between the different pages.

Signing out of TDV Web UI

To Sign out of TDV Web UI, follow these steps:

1. From any screen within the TDV Web UI, click on the User icon in the top right corner
2. Choose Sign out. You will be logged out of TDV Web UI.

Accessing Online Help

To access the online help for TDV Web UI, follow these steps:

1. From any screen within the TDV Web UI, click on the User icon in the top right corner
2. Choose the option Application Help.
3. The Online help for TDV Web UI is displayed in a separate tab.

Data Catalog

The Data Catalog provides a complete view into all securely accessible datasets, allowing the user to gain a quick understanding of what data is available for their consumption.

Users can analyze on-premise and cloud datasets together through the search and data preview capabilities of the Data Catalog quickly and securely by utilizing TIBCO's patented data virtualization engine.

Following topics are covered in this chapter:

Resource Tree

The Data Catalog gives you the capability to browse the tabular resources created using TDV Studio or the Web UI. The resources that are listed includes published and non-published tables and views.

The resources created are organized into the following folders:

Folder name	Description
My Home\Published	The user's published resources.
My Home\Private	The user's non-published resources.
Published\databases	All TDV published resources.
Shared\	All shared resources of TDV.
Users\	Resources of the user listed by domain name.

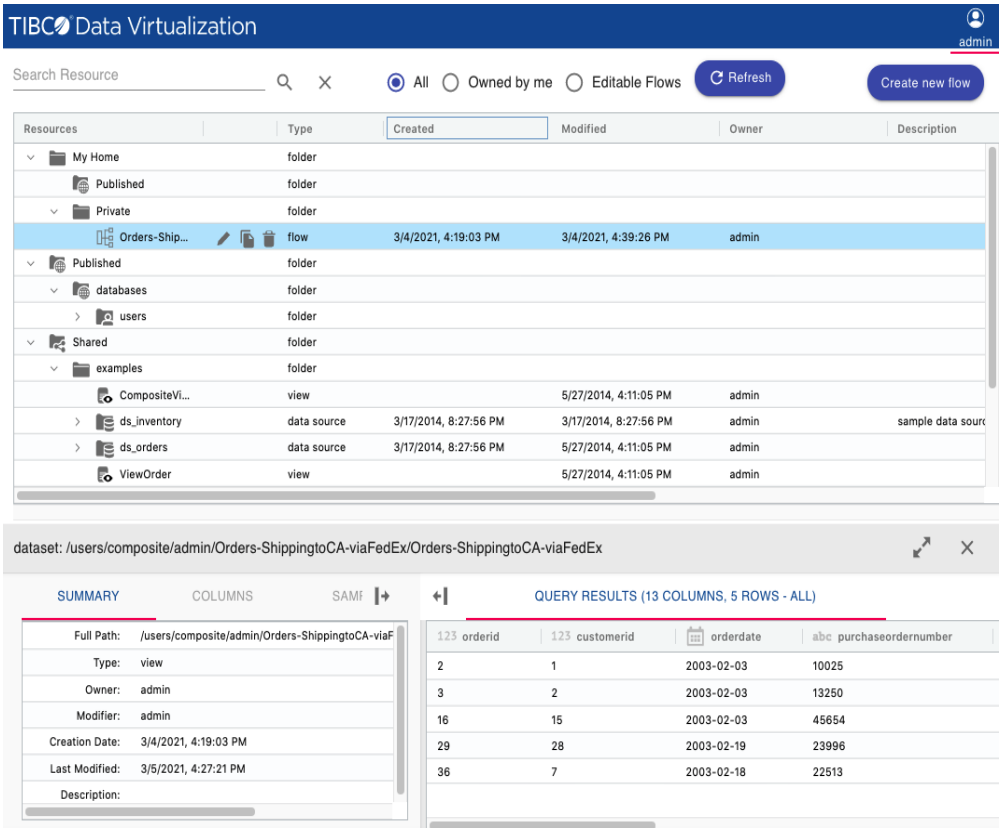

The resources are listed in a hierarchical expandable/collapsible format. The following table explains the details of the Resource Tree grid labels:

Label	Description
Resources	This field displays the following: • An Expand/Collapse icon - click on the ">"icon to expand or collapse a resource (such as Folder, Schema). • A Resource type icon - Indicates the type of resource. • Name of the Resource. • Resource Edit icons - Lets you open the resource in Flow editor, copy the resource into a new flow or delete it. Notes: • Any resource created within the TDV Web UI can be edited, deleted as well as copied into a blank flow. Resources created in TDV Studio can only be viewed in the Data Catalog of the Web UI and cannot be edited. • If you wish to copy/delete the flow using TDV Studio, you must do so at the Flow Folder level.
Type	Type of the Resource (Folder, Schema, Catalog, Database, Table, View or a Flow)
Created Date	Timestamp of when the resource was created.
Modified Date	Timestamp of when the resource was modified.
Owner	The user who created the resource.
Description	Any description of the resource given at the time of creating the resource.
Tooltip	The full path of the resource is displayed in a tooltip as the mouse hovers over the resource name.

Search/Filter Resource

The Data Catalog allows resources to be sorted and filtered. There is also a Search tool using which the user can do a simple keyword search for any table or view that the user has access to view within the TDV system.

The resources listed in the Data Catalog can be filtered to show all resources or just the ones created by the user.

Details of the quick filters are given below

- **Editable Flows** - Used to filter tables/views that have been created within the Web UI.
- **Owned by me** - Used to filter tables/views that have been created by the currently logged in user.
- **All** - Used to clear all filters and display all tables/views that match the keyword search criteria.

Resource Details

The details and data (in case of a table, view or flow) of each resource selected in the resource tree is displayed in a pane below the resource tree.

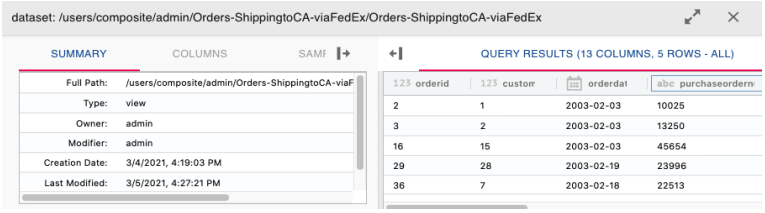

The preview window or the resource details window displays the details of the resource on the left and a data preview (where applicable) on the right.

Resource Details

Following are the details displayed in the Resource Details pane:

- **Summary Tab** - The full path of the resource, the resource type, owner, description and the created/modified dates.
- **Columns** - List of all the columns, the data type of each column and a description. This tab is only visible for certain resources such as tables, views and flows. Using the search utility, you can search for a specific column.

- Sample - The sample tab gives the ability of filtering the data preview by providing a where clause that will be applied to the data displayed. You can also limit the row count of the data displayed. Adjusting the sample data size can improve performance.

Data Preview

The Data Preview pane displays a sample/full dataset of the selected resource. Following are some of the user-friendly features of the Data Preview pane:

- The Data Preview pane gives a preview of the dataset up to 10000 rows.
- The data displayed can also be filtered using a where clause in the sample tab of the Resource Details pane.

  Note: You may need to wrap column names in double-quotes if it contains certain special characters.
- The Data Preview Pane also gives the ability to sort on a column, pin a column, search for a specific column, and auto size the column width.
- To view a larger dataset, you can expand the view to a full screen.

Creating a New Flow

From the Data Catalog, you can create a new flow by clicking on the "Create New Flow" button. You can also copy an existing flow to a new flow by clicking on the Copy icon next to the Resource name in the Resource Tree.

Refer to the chapter Flow Editor, page 19 for more information on creating a new flow or editing an existing flow.

Flow Editor

The Flow Editor provides the user a self-service interface to create a view from tables in different data sources and manipulate the data using the column filters, JOIN, INTERSECT and other query options. The resulting dataset can be published for consumption in TDV Studio and other client tools the user may have access to.

This chapter explains in detail the following topics:

Creating a Flow

The Canvas of the Flow Editor provides a self-service interface to create a SQL-based view with options to publish. The flow diagram that the user creates, provides a quick insight into the logic behind the data lineage.

Follow these steps to create a data flow diagram on the data canvas screen.

1. You can access the Flow Editor either by clicking on the Create New Flow button on the TDV Web UI landing page, the Data Catalog, or by choosing to copy from an existing flow in the Data Catalog.
2. By default the flow is named as “New Flow”. To change the name of the flow, double -click on the name of the flow in the Flow Editor.

Note: Special characters such as double-quote, single-quote, “/”, “<“, and “>” are not allowed in flow names.

3. Drag the Dataset icon from the Operations palette and drop it in the Canvas.
4. Select the dataset node on the canvas, if it is not already selected. The Dataset configuration pane is displayed below the canvas.

Note: You can drag and drop many data sets into your canvas and individually configure those to different data sources.

5. In the Dataset configuration pane, choose a datasource to be associated with the dataset by clicking on the “Choose a Dataset” button. A list of data sources

is displayed in the "Add a Dataset" window. The list includes existing Flows created using the Web UI along with other published resources.

6. Choose the data source and click Ok.

   - All the available columns in the data set are displayed in the Columns tab of the Data Configuration pane.

   - A preview of the dataset is displayed in the Preview Pane. The columns can be sorted by clicking on the column header.

   - A view-only SQL query is displayed in the Query pane.

7. From the Operations palette, drag and drop the operations you want to perform in order to manipulate your data. Refer to the section The Operations Palette, page 21 for more details on each of the operations.

8. The flow is auto-saved and can be viewed in the Data Catalog. Once the data flow diagram is created, you have the option publish the flow or review and make more changes.

Flow Validation

A warning symbol appears on the right side of the node that is on the canvas, if the configuration of the node is incomplete or incorrect. The exact reason of the warning appears in the tool tip as you hover the mouse over the warning symbol.

In addition to the warning on the node, an error message also appears in the Resource Configuration pane that is displayed below the canvas.

Deleting a Node

Any node that is on the canvas can be deleted by simply clicking on the delete icon that appears in the top right corner of the node. The delete icon appears as you hover the mouse over the node.

Rearranging the Nodes

The nodes in the data flow diagram can be rearranged by simply dragging it around and reconnecting it to the desired node. If there is an error in the flow, a warning symbol appears on the right side of the node.

Change History

As you edit a flow, a change history is maintained for the Flow Editor actions. You can revert to the earlier state of the flow by simply choosing a state from the Change History list that is available in the top-right corner of the Canvas. The history is only maintained for the session that the flow is active in the canvas.

The Operations Palette

The table below describes the different elements used to create a data flow diagram on the Canvas.

Icon	Description
	Click on this icon in the palette to add a Dataset to the Canvas. See Adding a Dataset.
	Click on this icon in the palette to choose specific Columns and/or build your expression. See Adding Columns.
	Click on this icon in the palette to add a WHERE clause and build your query. See Adding Filter.
	Click on this icon in the palette to add a JOIN operation to your query. See JOIN Operation.
	Click on this icon in the palette to add a UNION operation to your query. See UNION Operation.
	Click on this icon in the palette to add an EXCEPT operation to your query. See EXCEPT Operation.
	Click on this icon in the palette to add an INTERSECT operation to your query. See INTERSECT Operation.
	Publish you flow. See Publishing a Flow.

Icon	Description
	Clear your work on the Canvas.
	Delete your flow.

Adding a Dataset

A flow is created by adding a dataset to the canvas. In order to do that, drag the "Add Dataset" icon and drop it in the canvas. Select the Dataset icon on the canvas to display the Details pane below the canvas. Follow these steps to configure the dataset.

1. Select the Dataset icon on the canvas. The details pane is displayed.
2. Click on "Select a dataset" button. The "Add a dataset" window is displayed.
3. Choose a datasource to be associated with the dataset by clicking on the "Choose a Dataset" button. A list of data sources is displayed in the "Add a Dataset" window. The list includes existing Flows created using the Web UI along with other published resources.
4. Choose the data source and click Ok.

   - All the available columns in the data set are displayed in the Columns tab of the Data Configuration pane.

   - A preview of the dataset is displayed in the Preview Pane. The columns can be sorted by clicking on the column header.

   - A view-only SQL query is displayed in the Query pane.

Features and Limitations of the Dataset Operation

- There are no inputs to the Dataset node/operation.
- The Dataset node can have only one output.
- By default, the name of the table/datasource is the name of the dataset node. However, this can be edited by clicking on the name. The path of the datasource is also displayed on the node.
- There is no limit to the number of datasets defined in a flow. But there must be at least one dataset node defined.

- To delete a dataset node, hover the mouse on the node and click on the delete icon that is displayed on the top right corner of the node.

Adding Columns

Follow these steps to add a Column node and configure it:

1. Drag and drop the column node from the Operations palette and drop it in the canvas. By default the name of the Column node is "Columns". To modify the name, click on the name and provide a new name.
2. Move the column node towards the dataset node on the canvas. You will see a connector appear between the two nodes, indicating a successful connection.
3. The Column configuration pane is displayed below the Canvas. If the configuration pane is not displayed, make sure the column node is selected in the canvas. All the columns of the dataset are included by default, By clicking on the check box to the right of each column, columns can be selected/deselected (included/excluded) from displaying in the flow.
4. By default, the name of the column is retained from the dataset it was selected from. You can change the display name by clicking on the name displayed in the Column Name field and providing a new name.
5. The Column Type can be changed by clicking on the type displayed in the Type field.
6. Click on the Source column to open the Expression Editor. Using the Expression Editor, you can view and edit the expression of the column.
7. Click on the helper tools icon displayed on the top left corner of the expression editor A list of expressions will be displayed. Click on the function name to display a help window with the syntax of the selected function. Refer to the chapter *SQL Functions* in the *TDV Reference Guide* for a complete list of the supported SQL functions with syntax help and examples.
8. A preview of the dataset is displayed to the right of the Column Configuration pane. The order of the displayed columns can be altered by moving the column header.

Features and Limitations of the Column Operation

- The Column operation can have only one input and one output.
- This operation can be added after any operation except the Publish operation.

Adding Filter

The Filter operation enables users to specify a WHERE clause for the dataset, thereby limiting the number of rows returned. Follow these steps to define a filter in the data flow:

1. Drag the Filter operation from the Operations palette and drop it in the canvas. By default the name of the node is set as "Filter". Click on the name to change it.
2. Filter operations can be connected to another operation by dragging the filter operation towards any other operation you wish to apply the filter.
3. Once the Filter operation is added to your data flow diagram, the Filter Configuration pane is displayed below the canvas.
4. You can add as many filters as needed to construct your WHERE clause and build your query. Click on the + button to add more conditions. The following operators and conjunctions are supported in the Filter node:
   - = (Equal)
   - > (Greater than)
   - < (Less than)
   - <= (Less than or Equal)
   - >= (Greater than or Equal)
   - <> (Not Equal)
   - LIKE
   - BETWEEN
   - IN
   - IS NULL
   - IS NOT NULL
   - AND
   - OR
   - AND NOT
5. Click on the "Write an Expression" option in the field specific drop down list, to write your own expression. The Expression editor window is displayed.
6. Click on the helper tools icon to view the list of functions supported. The Expression editor is a self-service tool that lets you build your expression. Refer to the chapter *SQL Functions* in the *TDV Reference Guide*, for a list of the supported functions.

7. After the WHERE clause is successfully built, the data preview pane on the right displays the result set. Data in the preview pane can be sorted.

Features and Limitations of the Filter Operation

- The Filter operation can have only one input and one output.
- This operation can be added after any operation except the Publish operation.

JOIN Operation

The JOIN operation enables you to manipulate the data coming from different datasources. Follow these steps to add and configure a JOIN operation in the data flow:

1. Drag and drop the JOIN operation into your canvas.
2. JOIN operation should have two data inputs on which the operation can be performed. Drag the JOIN node towards the node to which you wish to connect it to. A connector appears between the two nodes indicating a successful connection.
3. Once the JOIN operation is successfully added to the data flow diagram, the Join Configuration pane is displayed below the canvas. The Join tab in the Join Configuration pane lets you choose a LEFT, RIGHT, INNER or a FULL JOIN and apply it to the data displayed in the preview pane.
4. In addition to specifying the type of JOIN in the Join tab, you can also build your own query expression in the Column tab. Click on the Source field to open the Expression editor. Use the helper tools to build your own expression.
5. The data preview pane on the right displays the result set data.

Features and Limitations of the JOIN Operation

- The JOIN operation must have two inputs and can have only one output.
- This operation can be added after any operation except the Publish operation.

UNION Operation

The UNION operation enables you to manipulate the data coming from different datasources. Follow these steps to add and configure a UNION operation in the data flow:

1. Drag and drop the UNION operation into your canvas.
2. UNION operation should have at least two data inputs on which the operation can be performed. Drag the UNION node towards the node to which you wish to connect it to.
3. Once the UNION operation is successfully added to the data flow diagram, the Union Configuration pane is displayed below the canvas. All the columns from both the data inputs are displayed. Click on the + button next to the column to add the UNION operation and X button to delete it. You also have the option to do a “UNION ALL” as well to “Auto Match the Columns”.
4. The data preview pane on the right displays the result set data.

Features and Limitations of the UNION Operation

- The UNION operation can have unlimited inputs. But there must be at least two inputs and can have only one output.
- This operation can be added after any operation except the Publish operation.

EXCEPT Operation

The EXCEPT operation enables you to manipulate the data coming from different datasources. Follow these steps to add and configure an EXCEPT operation in the data flow:

1. Drag and drop the EXCEPT operation into your canvas.
2. EXCEPT operation should have at least two data inputs on which the operation can be performed. Drag the EXCEPT node towards the node to which you wish to connect it to.
3. Once the EXCEPT operation is successfully added to the data flow diagram, the Except Configuration pane is displayed below the canvas. All the columns from both the data inputs are displayed. Click on the + button next to the column to add the EXCEPT operation and X button to delete it. You also have the option to do a “EXCEPT ALL” as well to “Auto Match the Columns”.
4. The data preview pane on the right displays the result set data.

Features and Limitations of the EXCEPT Operation

- The EXCEPT operation can have unlimited inputs. But there must be at least two inputs and can have only one output.
- This operation can be added after any operation except the Publish operation.

INTERSECT Operation

The INTERSECT operation enables you to manipulate the data coming from different datasources. Follow these steps to add and configure an INTERSECT operation in the data flow:

1. Drag and drop the INTERSECT operation into your canvas.
2. INTERSECT operation should have at least two data inputs on which the operation can be performed. Drag the Intersect node towards the node to which you wish to connect it to.
3. Once the INTERSECT operation is successfully added to the data flow diagram, the Intersect Configuration pane is displayed below the canvas. All the columns from both the data inputs are displayed. Click on the + button next to the column to add the INTERSECT operation and X button to delete it. You also have the option to do a “INTERSECT ALL” as well to “Auto Match the Columns”.
4. The data preview pane on the right displays the result set data.

Features and Limitations of the INTERSECT Operation

- The INTERSECT operation can have unlimited inputs. But there must be at least two inputs and can have only one output.
- This operation can be added after any operation except the Publish operation.

Expression Editor

The Expression editor can be accessed either through the Filter or the Columns operators. You can build your own arbitrary expressions using the Expression editor. The following window is displayed when you choose to write your own expression from the Columns/Filter operator’s configuration pane.

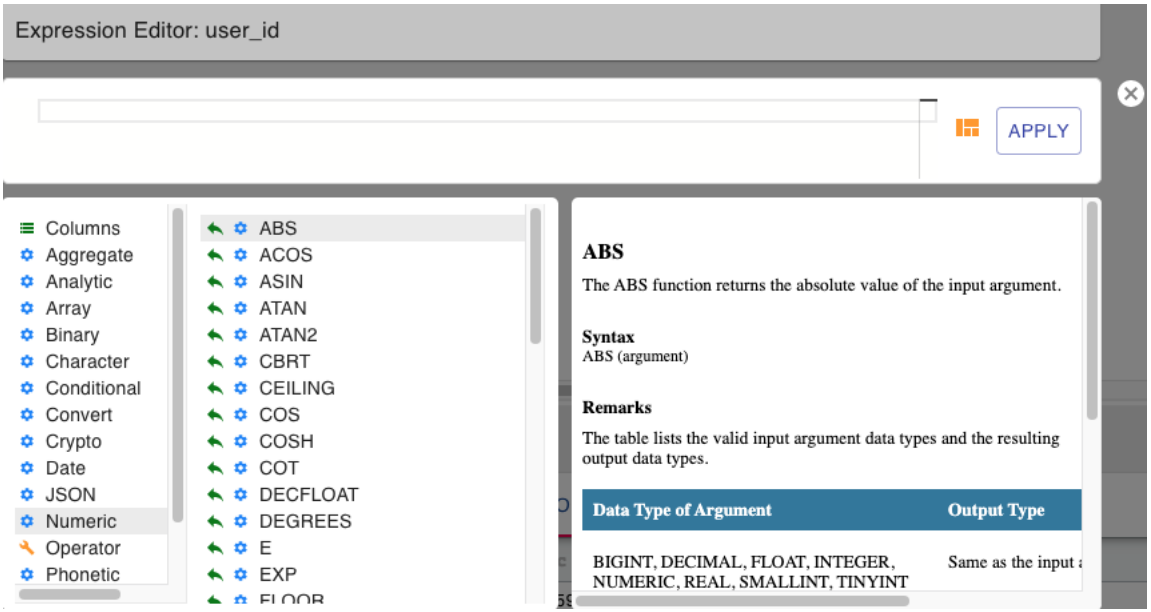

A list of all the supported functions and a syntax help is provided in the Expression Editor window. Follow these steps to add the function and build your expression:

1. Choose the column you want to write the expression for.
2. Choose a function from the list of functions displayed.
3. Click on the green icon next to the function name to insert it into the editor.
4. Complete the expression by providing the required parameter values for the function.
5. Click on Apply. You can validate the expression by reviewing the data displayed in the data preview pane.

Publishing a Flow

To publish a data flow, follow these steps:

1. Drag and drop the Publish node from the operations palette into the canvas.
2. Drag the Publish node onto the last output node of the data flow. By simply having this publish node in the data flow diagram, the flow is now published.

Publishing a flow creates a virtual schema in a user-specific location under the system folder /services/databases/users named like <domain>_<user>. "/services/databases/users" is a system provisioned virtual database and is meant to be used by TDV instances/clusters supporting self-service, data provisioning use cases. Specifically, it allows individual users, such as analysts, to publish data flows, defined within TDV Web UI, in the form of data endpoints that can be consumed by external reporting, visualization or data science tools of their choice.

To see a list of the endpoints, click on the Publish node. A Publish Configuration window displays below the canvas and has the following tabs:

Publish tab: This tab displays the Default Publish Location (non-editable) and an optional editable Description field.

Endpoints tab: The endpoints tab displays a list of client locations where the published flow is available for consumption. This list includes the jdbc, odbc, ado.net and odata endpoints.

Editing a Published Flow

A published flow can be edited if you have created it or if you have privileges to do so.

The TDV Web UI is intended to be used as a self-service tool by analysts to create, modify and publish their own data flows. Within the Web UI, there is no provision to share resources between users. However, users with read/modify all resources rights (typically, an Administrator) or users that have been granted read access to someone else's flows, can view/edit resources owned by other users. You may need to republish the flow during such scenarios as listed below:.

1. When a published flow owned by a different user is renamed/modified by you, then it is important to know that the flow is not automatically copied to your private folder or your private publish folder.
2. A published flow owned by a different user can be copied to your private folder. However, this does not automatically copy the published flow to your private publish folder.

Follow these steps to re-publish the flow to your private publish folder:

- Click on the Publish node.
- Click on the "Fix it" button that appears in the Publish configuration window below the canvas.

Additional Notes

- When a published flow owned by a different user is renamed by you, then it is important to know that when that flow is deleted by the original owner, it will delete the flow from your published folder as well.
- When a non-published flow owned by a different user is renamed/modified by you, then it is important to know that the flow is modified only in the original owner's private folder.