



TIBCO® Data Virtualization

Web User Interface User Guide

Version 8.7 | October 2023

Contents

Contents	2
Introduction	4
System Requirements	4
How to Access TDV Web UI	5
Accessing from TDV Studio	5
Logging in to TDV Web UI	5
User Management	5
The Login Screen	5
Overview of the Data Catalog	8
Overview of the Flow Editor	9
Navigation	10
Signing out of TDV Web UI	11
Accessing Online Help	11
Application Selector	11
Data Catalog	13
Data Catalog Navigation	13
Published Data	13
Integration	14
Data Stores	14
Resource Browser	14
Data Catalog Overview	15
Data Catalog Details	16
Searching the Data Catalog	18
Search Scope	18
Free Text Search Utility	18
Search Results	21

Creating a New Flow	23
Data Workbench	24
Flow Editor	24
Creating a Flow	24
The Operations Palette	26
Expression Editor	34
Publishing a Flow	36
TIBCO Documentation and Support Services	39
Legal and Third-Party Notices	42

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.

The TDV Web UI provides an agile way to transform and move data between different data sources. These data sources can be located anywhere, even in the cloud.

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.

The TDV Web UI provides an agile way to transform and move data between different data sources. These data sources can be located anywhere, even in the cloud.

System Requirements

You will need to install and run TDV Server 8.4.0 (or later) 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.

The supported Web browsers are:

- Chrome v95 on Windows 10
- Microsoft Edge v95 on Windows 10
- Mozilla Firefox v94 on Windows 10

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.

TIBCO® Data Virtualization

Username

Password

Domain

Sign In

OR

my-oauth2

Copyright © 2022 TIBCO Software Inc. All Rights Reserved.

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.

Note: TDV Web UI does not support Kerberos authentication.

If an OAuth 2.0 domain has been set up and access tools privileges given to the user, then you can login to TDV Web UI by giving the OAuth 2.0 domain name (that has been setup in Web Manager) in the Domain field. Click on the domain name button that is displayed in the login screen. You will not need to provide username and password while using the OAuth 2.0 domain. The option to login to the OAuth2 domain will be available only if the domain is defined in the Web Manager.

You need to be actively logged into your IDP account, before attempting to login to any of the TDV client tools using the OAuth2 domain. If not, you will be redirected to your IDP login page. Once logged in, click on the Oauth2 domain in TDV Studio, Web Manager or Web UI to successfully login to the TDV client.

In order to allow access to Web UI, the TDV Web UI URL must be added to the Redirect Urls list in your Identity Provider. For example,

`http://localhost:9400/webui`

`https://localhost:9402/webui`

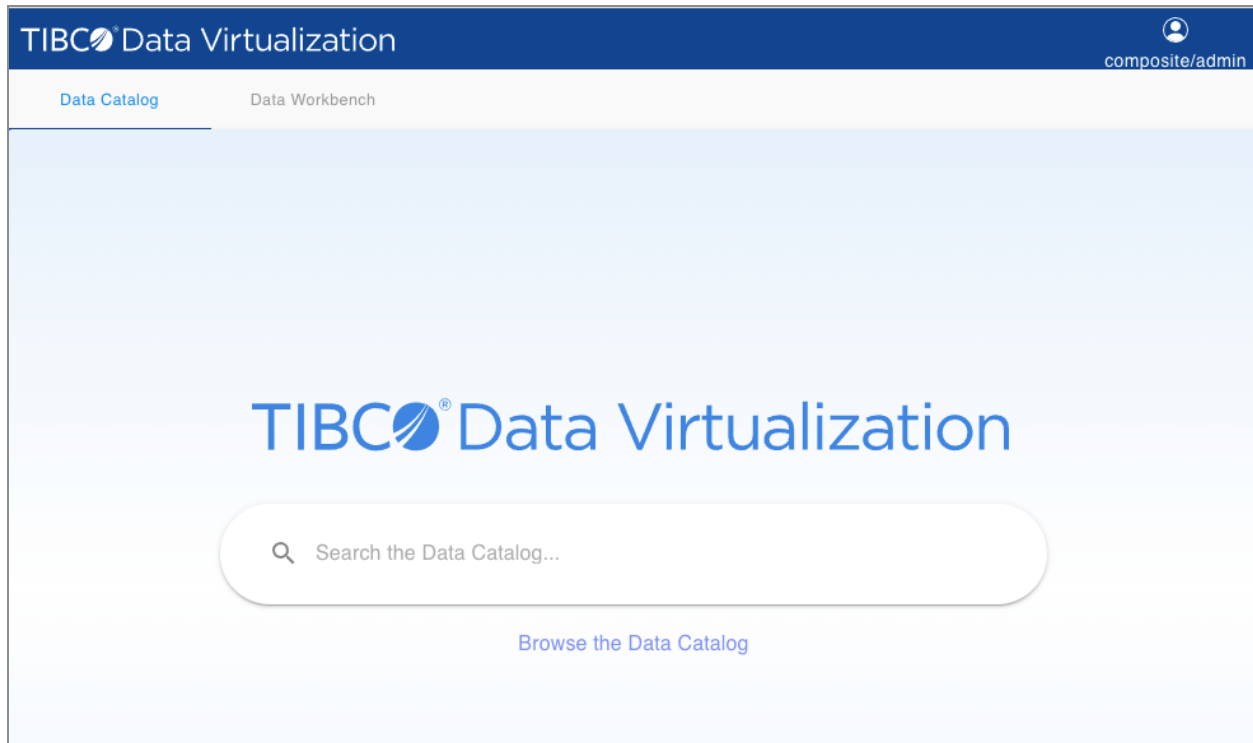
You may also use a domain address in the redirect url list. For example,

`http://xyz.com/webui`

Field	Description
Username	<p>The username required to login to TDV Web UI. Obtain a valid username from your TDV administrator.</p> <p>Note: Username that contains '?' and '&' is not supported.</p> <p>If using an OAuth 2 domain, this field can be left blank.</p>
Password	<p>The password required to login to TDV Web UI. Obtain a valid password from your TDV administrator.</p> <p>If using an OAuth 2 domain, this field can be left blank.</p>
Domain	<p>The domain to which the user belongs.</p>

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

3. The following Data Catalog Search page is displayed.



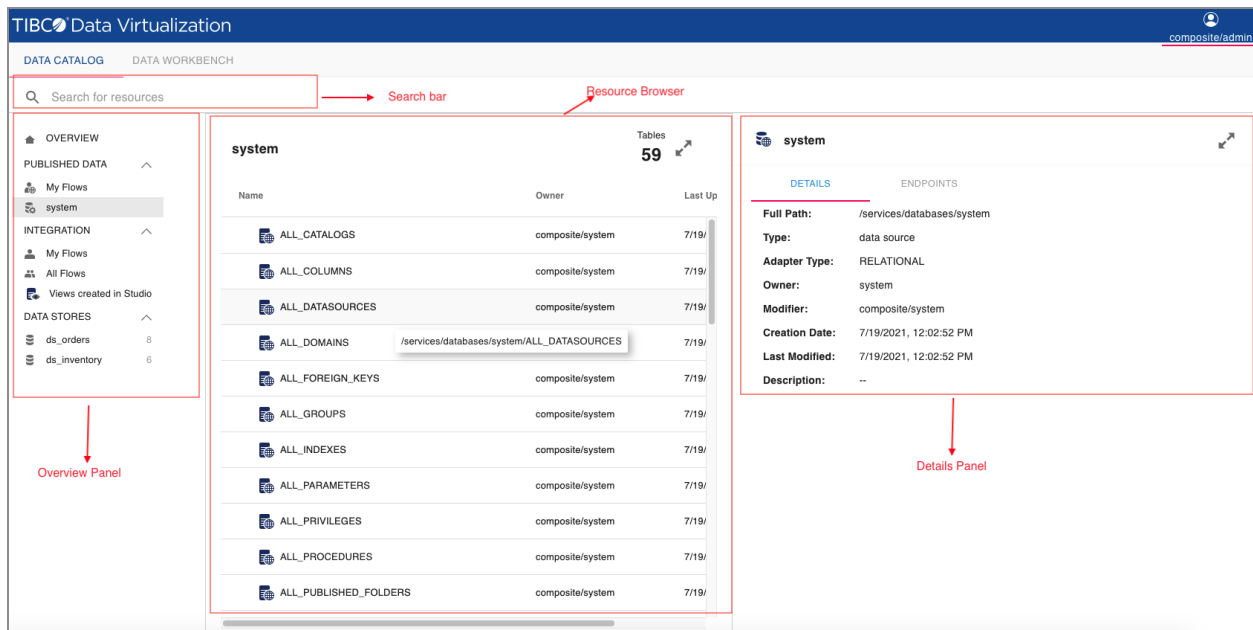
The Search tool in the Data Catalog tab on the landing page allows you to perform a simple keyword search for any resource that you have access to view within the TDV system. Refer to the section [Searching the Data Catalog](#) for more details about the Search tool,

Click on the Browse the Data Catalog link to open the Data Catalog screen. Refer to the section [Data Catalog Navigation](#) for details about the different categories in the Data Catalog Screen.

The Data Workbench tab on the landing page lists all the recently used flows as well as a link to [Creating a New Flow](#).

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.



The search tool in the Data Catalog lets you search for the resources. Refer to [Searching the Data Catalog](#) for details on how to use the Search feature of TDV Web UI.

The Overview Panel on the lists the TDV resources. Refer to [Data Catalog Navigation](#) for details on how the resources are grouped into different categories.

The Resource Browser gives you the ability to browse the resources in each category.

The Resource Browser Details Panel in the far right gives all the details about the resource selected in the Resource Browser. Refer to [Data Catalog Details](#) for more information.

The following actions can be performed from the Data Catalog screen:

- View the details of the available resources.
- Browse the existing flows and data stores.
- Create a new flow.
- Edit an existing flow.
- Search the TDV resources.

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](#) for more information.

Navigation

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

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

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:

From any screen within the TDV Web UI, click on the User icon in the top right corner

1. 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:

From any screen within the TDV Web UI, click on the User icon in the top right corner

1. Choose the option Application Help.
2. The Online help for TDV Web UI is displayed in a separate tab.

Application Selector

To access other TDV Applications from the Web UI portal, click on the Application Selector icon that is displayed on the top right-corner.

Application	Description
Web UI	Navigates to the Web UI in a new browser tab.
Web Manager	Navigates to the Studio Web Manager in a new browser tab. Web Manager provides a thin client Web-page based interface for managing TDV. Refer to the TDV Administration Guide chapter <i>System Management with Manager</i> for more information.
Deployment Manager	Navigates to the Deployment Manager in a new browser tab. Deployment Manager is a web-based tool you can use to manage and streamline the development life cycle of

Application	Description
	TIBCO® Data Virtualization. Refer to the TDV Administration Guide chapter <i>Deployment Manager</i> for more information.
Download TDV Studio Drivers	Navigates to the TIBCO eDelivery site in a new tab. You can download the TDV Studio and Drivers for all the distribution types and platform.

You can also navigate to the following Help and Support tools from this Application Selector menu.

- Getting Started Videos
- TDV Product Documentation
- Training and Certification
- TDV Community
- TDV Architecture
- TIBCO Support

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 sections are covered in this topic:

- [Data Catalog Overview](#)
- [Data Catalog Navigation](#)
- [Searching the Data Catalog](#)
- [Data Catalog Details](#)

Data Catalog Navigation

The Navigation Panel of the Data Catalog displays all the TDV resources and is grouped in the following categories:

Published Data

The Published Data section of the Data Catalog lists all of the virtual databases, published flows and the system tables. To view the contents of each resource in the Details panel, choose a specific category from Published data section and then click on the resource listed in the Resource Browser. The Details panel in the far right displays content of the resource, the columns, endpoints where the resource can be accessed by downstream applications and other details about the resource.

Integration

The Integration section of the Data Catalog lists all the flows created in the TDV Web UI as well as the views created in TDV Studio.

Note: The flows that you have created will be listed under the category "My flows". The flows you have been given access to will be listed under the "All Flows" category along with the flows created by you. The full listing of all non-published views across all of TDV that the user has access to will be listed under the "Views created in Studio" category.

Data Stores

The Data Stores section of the Data Catalog screen lists all the data sources that have been introspected and that has tables. To view the contents of the data store in the Details panel, click on the data store listed in the Resource Browser. The Details panel also displays the contents in the data store, columns and the details of the selected data store such as the full path, owner and date created/modified.

Resource Browser

The Resource Browser of the Data Catalog lists all the resources available for the selected category in the Navigation Panel on the left. For example, choose the category "system" under the "Published Data" to list all System Tables.

The screenshot shows the TIBCO Data Virtualization interface. The top navigation bar includes 'DATA CATALOG' and 'DATA WORKBENCH'. The left sidebar lists various navigation options. The main panel displays a table of resources for the 'system' catalog. The table has columns for Name, Owner, and Last Up. The resource 'ALL_DATASOURCES' is highlighted, showing its full path as '/services/databases/system/ALL_DATASOURCES'. The right panel shows the details for the selected resource, including Full Path, Type, Adapter Type, Owner, Modifier, Creation Date, Last Modified, and Description.

Name	Owner	Last Up
ALL_CATALOGS	composite/system	7/19/
ALL_COLUMNS	composite/system	7/19/
ALL_DATASOURCES	composite/system	7/19/
ALL_DOMAINS	/services/databases/system/ALL_DATASOURCES	7/19/
ALL_FOREIGN_KEYS	composite/system	7/19/
ALL_GROUPS	composite/system	7/19/
ALL_INDEXES	composite/system	7/19/
ALL_PARAMETERS	composite/system	7/19/
ALL_PRIVILEGES	composite/system	7/19/
ALL_PROCEDURES	composite/system	7/19/
ALL_PUBLISHED_FOLDERS	composite/system	7/19/

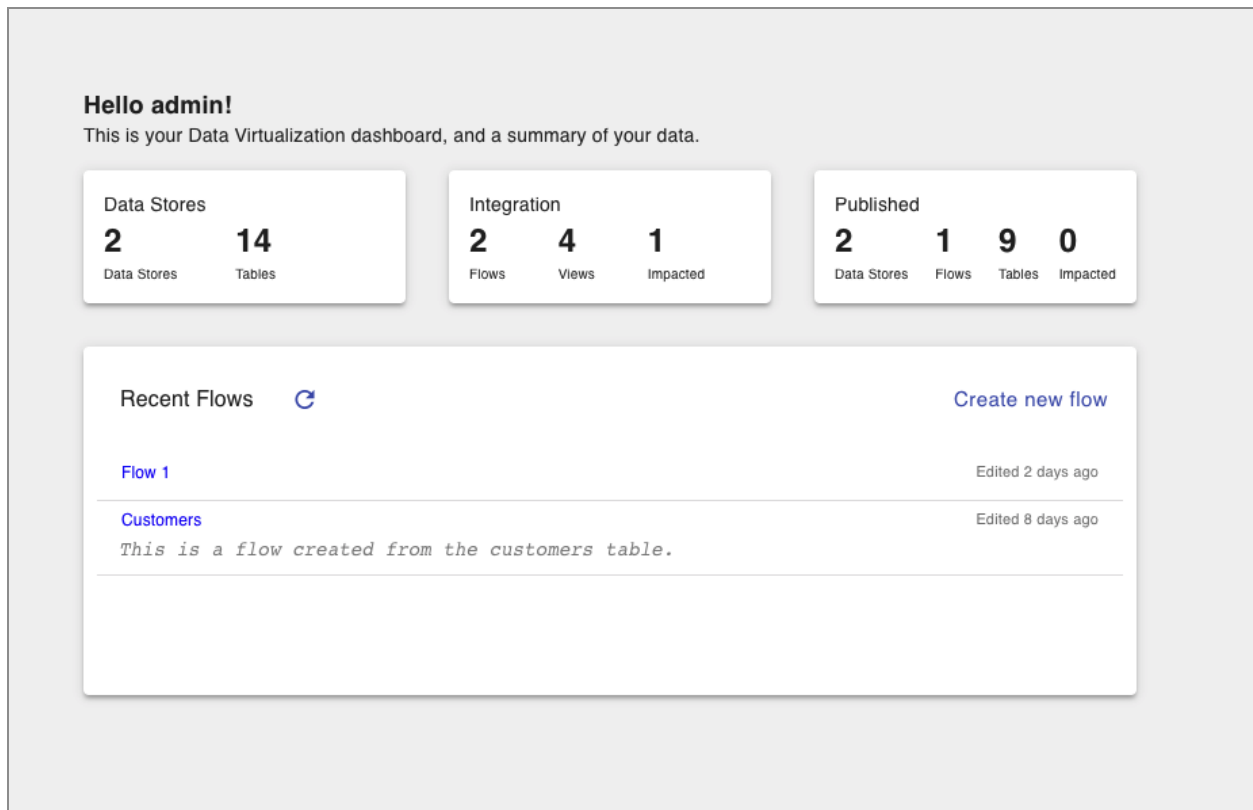
system Details:

- Full Path: /services/databases/system
- Type: data source
- Adapter Type: RELATIONAL
- Owner: system
- Modifier: composite/system
- Creation Date: 7/19/2021, 12:02:52 PM
- Last Modified: 7/19/2021, 12:02:52 PM
- Description: ..

Click on a resource in the resource browser to view the details in the Resource Details panel on the right. Refer to [Data Catalog Details](#) for more information.

Data Catalog Overview


The Data Catalog Overview panel displays a summary section with the number of resources that the logged in user has access to and a list of recent flows.



The summary section is grouped into 3 sections - Data Stores, Published and Integration. Any flow/view that may not have a valid definition, is flagged as an *impacted* resource.

Click on the flow name in the Recent Flows section to edit an exiting flow in the Flow Editor or click on the link "Create New Flow" to create a new flow.

Data Catalog Details

The Resource Browser Details panel displays the details of the selected resource. You can expand/restore the Details panel by using the  button. From the Details panel, click on the Create flow button to open the flow editor with a blank canvas.

To navigate to the parent folder of a resource in the Search Results

1. Select the resource in the Search Results.
2. Click on the Browse button on the top right corner of the Resource Details panel.

Preview Tab: The contents of the selected resource is displayed in the preview tab. You can use the Search bar of the Preview tab to filter data by providing the Search clause. For example, in the Supplier table, to filter all records for the Supplier "Cedar Systems", enter the text **SupplierName='Cedar Systems'** in the Search bar. To sort on any column, simply click on the Column heading.

Columns Tab: The Columns tab lists the columns in the selected resource. Use the Search bar to search for a specific column.

Details Tab: The details tab lists the important details about the selected resource, such as the physical location (Full Path), resource type (table/view), Source of the view, owner of the resource, a description of the resource and the date created/modified.

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

Searching the Data Catalog

The Search utility in Data Catalog allows you to perform a free text search across all TDV resources (data and system tables, flows, data stores and views).

Search Scope

Search results may be different for users with different permissions. Data Catalog search results depend on which TDV resource you have access to. You need to have at least a read-access to the TDV resource. Refer to the Administration Guide for more information about User Access and Privileges and how to grant access to TDV resources.

The Search result set in TDV WebUI Data Catalog does not include TDV procedures. The Search procedure can be directly run from TDV Studio to include procedures in your result set. For information on how to use the Search packaged query from TDV Studio, refer to the *API Guide*, Chapter *Build-in Procedures*.

Free Text Search Utility

From the Data Catalog screen, you can perform single or multiple key words search. The result set includes any or all of the key words. Data tables, system tables, data sources, flows and views are included in the Search.

Search Fields Types

The given key word is matched with any/all of the following field types in all the resources that the user has access to. The result set is displayed based on a Search algorithm that includes a weighted rank for each of these field types.

Field Type	Description
Name	The name of the resource.
Annotation	The description of a resource.

Field Type	Description
Column	The columns of a resource.
Definition	The definition (query SQL) of a View.

Search Filter Groups

The result set are grouped into 4 Search Filter Groups - Common, Resource Types, Search by and Data Stores. You will be able to filter your search results further, based on any of these 4 groups, but can do so only after performing an initial search.

Filter

Common

☐ Published Only

☐ Owned by Me

☐ Not Owned by Me

Resource Types

☒ Data Source

☒ Table

☒ View

Search by

☒ annotation

☒ column

☒ definition

☒ name

Data Stores

Select All | Clear All

☒ ds_inventory

☒ ds_orders

☒ system

Common

The Common filter group has the following categories. You can select any one or all of these categories. By default none of these are selected.

Published Only: Displays all published resources, including tables, views and flows that match the search keyword.

Owned by Me: Displays all flows owned by the user that is logged in and that match the search keyword

Not Owned by Me: Displays all flows not owned by the user that is logged in and that match the search keyword.

Resource Types

The Resource Type filter group displays the result set grouped by the resource type. Following are the categories in this filter group. By default all the resource types are selected. You can select all or one to further narrow the search results.

Database Table: Displays all the database tables that matches the search text and that the user has access to read.

Data Source: Displays all the data sources that matches the search text and that the user has access to read.

Flow: Displays all the flows that matches the search text and that the user has access to read/write.

System Table: Displays all the system tables that matches the search text and that the user has access to read.

View: Displays all the views that matches the search text and that the user has access to read.

Search by

The Search by filter group displays the result set grouped by the field type. By default, the result set includes data that matches any/all these categories. You can select all or one to further narrow the result set. The categories in this filter group are:

- **Name:** Displays all matches with the name of the resource.

- **Annotation:** Displays all matches with the resource description.
- **Column:** Displays all matches with the column names of the resources.
- **Definition:** Displays all matches with the resource definition.

Data Stores

The Data Stores filter group displays the result set grouped by individual data source. By default all the data sources that match the search keyword are included in the result set. You can further narrow your result, by including one or few of the displayed data sources.

Search Results

The Free text search results includes all the resources that have at least one matching field type. The results are grouped into 3 main categories - Published, Integration and Data Stores. The same resource may occur in more than one category. For example, if the search text is "customerid" and if there is a published flow called "customer" that has a matching column, then the flow may appear in more than one category as shown below:

Search Results

Resources
6

Name	Owner	Last Modified
Published (1 of 1)		
ALL_FUNCTIONS This table exposes all the built-in, native and custom functions in TDV. Column Composit... Published Data / system / model Search hit: annotation	composite/system	19 hours ago
Integration (2 of 2)		
ViewOrder Integration / Views created in Studio / examples Search hit: column x3, definition	composite/admin	9 years ago
CompositeView Integration / Views created in Studio / examples Search hit: column x3	composite/admin	9 years ago
Data Stores (3 of 3)		
customers Data Stores / ds_orders / tutorial Search hit: name, column	composite/admin	19 hours ago

Name of the resource

Field types that match the search text

Location of the resource

Resource Details

The details of the selected resource displays on the right pane. The details panel consists of a data preview, list of columns, location information and the endpoints where the resource is available for consumption by the client applications.

The screenshot shows the 'Search Results' page for the 'ALL_INDEXES' resource. The left sidebar lists three resources: 'ALL_INDEXES' (composite/system, 18 hours ago), 'ALL_INDEXES' (composite/system, 18 hours ago), and 'SYS_REQUESTS' (composite/system, 18 hours ago). The main panel displays the 'ALL_INDEXES' resource details, including a description, a search hit, and a 'Load More Results' button. Below the details, there is a table titled 'Viewing 100 rows 20 columns'.

123 INDE	abc INDE	abc INDE	123 COLU	abc COLU	123 ORDI	abc SORT	123 TABL
8487	categories_...	PRIMARY_KEY	8485	categoryid	1	A	8484
8437	employees_...	PRIMARY_KEY	8431	employeeid	1	A	8430
8424	inventorytra...	PRIMARY_KEY	8414	transactionid	1	A	8413
8480	products_pk...	PRIMARY_KEY	8472	productid	1	A	8471
8465	purchaseor...	PRIMARY_KEY	8454	purchaseor...	1	A	8453
8451	suppliers_pk...	PRIMARY_KEY	8440	supplierid	1	A	8439
8603	customers_...	PRIMARY_KEY	8591	customerid	1	A	8590
8553	employees_...	PRIMARY_KEY	8547	employeeid	1	A	8546
8578	orderdetails...	PRIMARY_KEY	8571	orderdetailid	1	A	8570
8621	orders_pkey	PRIMARY_KEY	8606	orderid	1	A	8605
8527	shippingme...	PRIMARY_KEY	8525	shippingme...	1	A	8524
1433	SYSTEM_CA...	OTHER	[NULL]	to_tsvector('...	1	[NULL]	1419
1429	SYSTEM_CA...	OTHER	[NULL]	catalog_name	1	A	1419
1435	SYSTEM_CA...	OTHER	[NULL]	to_tsvector('...	1	[NULL]	1419
1437	SYSTEM_CA...	OTHER	[NULL]	datasource_id	1	A	1419
1437	SYSTEM_CA...	OTHER	[NULL]	catalog_name	2	A	1419
1431	SYSTEM_CA...	PRIMARY_KEY	[NULL]	catalog_id	1	A	1419

To navigate to the parent folder of a resource in the Search Results

1. Select the resource in the Search Results.
2. Click on the Browse button on the top right corner of the Resource Details panel.

Refer to the section [Data Catalog Details](#) for more information about the Resource Details panel.

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](#) for more information on creating a new flow or editing an existing flow.

Data Workbench

The Data Workbench provides an interface to view all the recent flows that you have access to. From the Data Workbench screen, the flow editor can be accessed to create new flows and edit existing flows.

Click on the Create New Flow button to create a new flow. You can edit an existing flow by choosing the flow from the Recent Flows list.

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.

The following topics are explained in detail:

[Creating a Flow](#)

[The Operations Palette](#)

[Expression Editor](#)

[Publishing a Flow](#)

Creating a Flow

The Canvas of the Data Workbench 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 Workbench screen. You can access the Data Workbench either by clicking on the Create New Flow button on the landing page of the Data Workbench , or through the Create Flow option on the Data Catalog screen (by choosing an existing flow or datastore).

1. By default the flow is named as “Flow n” (where n is a sequential number). To change the name of the flow, double -click on the name of the flow in the Data Workbench and modify it.

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

2. Click on the Dataset icon from the Operations palette.
3. The Add Dataset window is displayed with a list of data sources. The list includes existing Flows created using the Web UI along with other published resources. You can choose one or many datasets using nthe SHIFT OR CTRL buttons in your keyboard.

Note: You will not be able to select multiple datasources from the Search results when you use search option.

4. As many Dataset nodes are created in the canvas as you chose in the Add Dataset window.

Note: You can individually choose a dataset node and configure it to a different data source.

5. When you choose a Dataset in your canvas, the dataset configuration panel is displayed below the canvas.
 - 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.
6. 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](#) for more details on each of the operations.
7. 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. To display the Change History list, click on the Change History icon located on the top right corner of the canvas. You can revert to an earlier state of the flow by simply choosing a state from the Change History list and clicking on the "Restore this version" button. The history is only maintained for the session that the flow is active in the canvas. Note that when you click on a specific version from the Change History list, you will be able to view the flow details (dataset definition and any operations used in the flow). However its only a read-only display and you cannot edit the flow from the Change History state. You will have to restore your flow to that particular version and then make any changes needed.

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

Icon	Description
	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

Icon	Description
	query. See INTERSECT Operation .
	Publish you flow. See Publishing a Flow .
	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 add and configure a dataset, follow these steps:

1. Click on the Dataset icon on the Operations Palette
2. The “Add dataset” window is displayed. Choose the datasource from the list that is displayed. The list includes existing Flows created using the Web UI along with other published resources. Note that you can choose multiple datasources by using the SHIFT button (to choose a continuous list) or the CTRL button in the keyboard.

Note: You will not be able to select multiple datasources from the Search results when you use search option.

3. Click the SELECT button to select the datasources into the canvas. If multiple datasources are chosen, multiple datasets will be created on the canvas.
4. In the canvas, when you choose a dataset,

- All the available columns in the data set are displayed in the Columns tab of the Data Configuration pane.
- The Dataset tab of the Data Configuration pane displays the details of the dataset such as the full path, a description, type and owner of the resource, created/modified date, etc.
- A preview of the dataset is displayed in the Preview Pane. The columns can be sorted by clicking on the column header.

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:

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.

1. 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.
2. 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.

3. 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.
4. The Column Type can be changed by clicking on the type displayed in the Type field.
5. Click on the Source column to open the Expression Editor. Using the Expression Editor, you can view and edit the expression of the column.
6. 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.
7. 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: 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.

1. Filter operations can be connected to another operation by dragging the filter operation towards any other operation you wish to apply the filter.
2. Once the Filter operation is added to your data flow diagram, the Filter Configuration pane is displayed below the canvas.
3. 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

4. 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.
5. 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.
6. 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:
 Drag and drop the JOIN operation into your canvas.

1. 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.
2. 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.
3. 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.
4. 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: Drag and drop the UNION operation into your canvas.

1. 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.
2. 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”.
3. 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:

Drag and drop the EXCEPT operation into your canvas.

1. 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.
2. 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”.
3. 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:

Drag and drop the INTERSECT operation into your canvas.

1. 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.
2. 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”.

3. The data preview pane on the right displays the result set data.

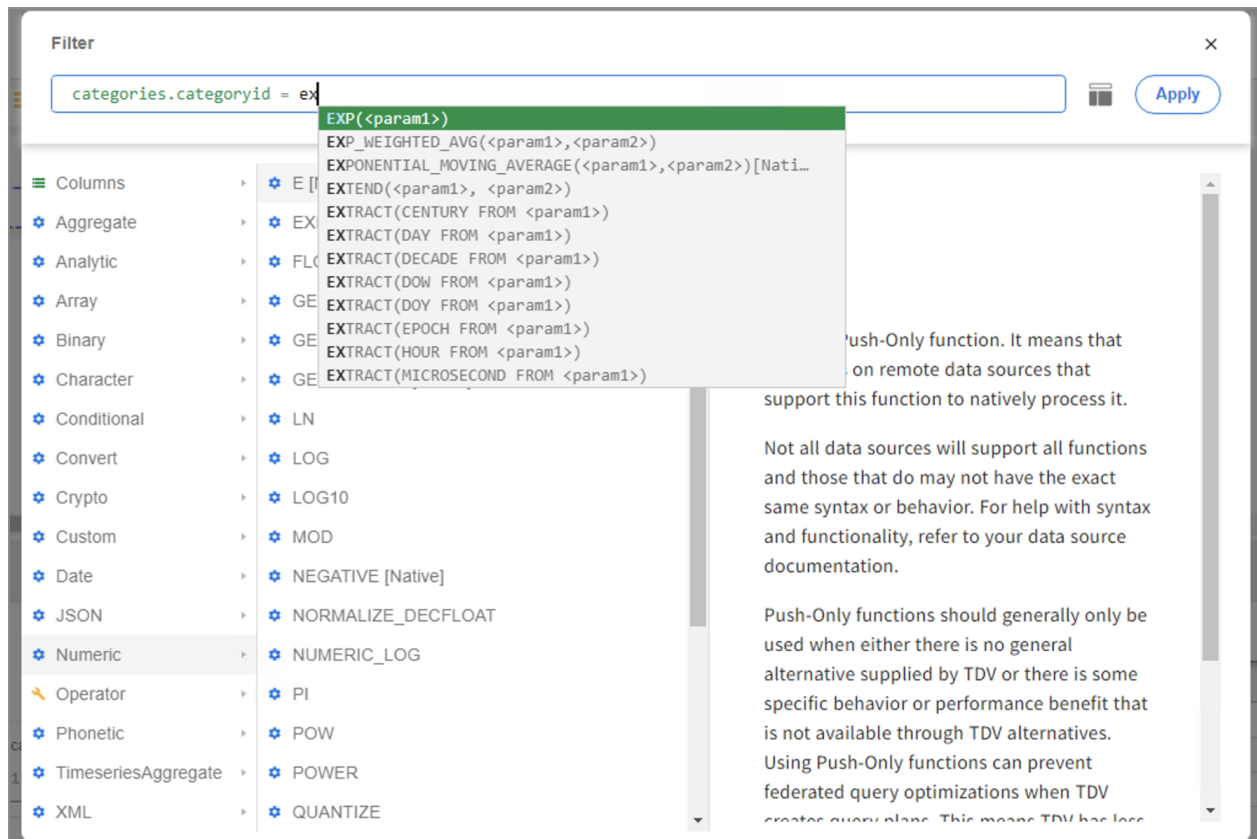
Features and Limitations of the INTERSECT Operation

- 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 TDV Web UI Expression Editor is a full-fledged IDE editing capability, including IntelliSense, syntax coloring, parameter information and error information, among other features. The compiler validates the expression after it's entered. If the expression is invalid, an error icon is displayed. The editor can also be expanded further to display syntax help for all the TDV functions organized by function categories. Click on the icon <> to the right of the Expression Editor text box to display the help.

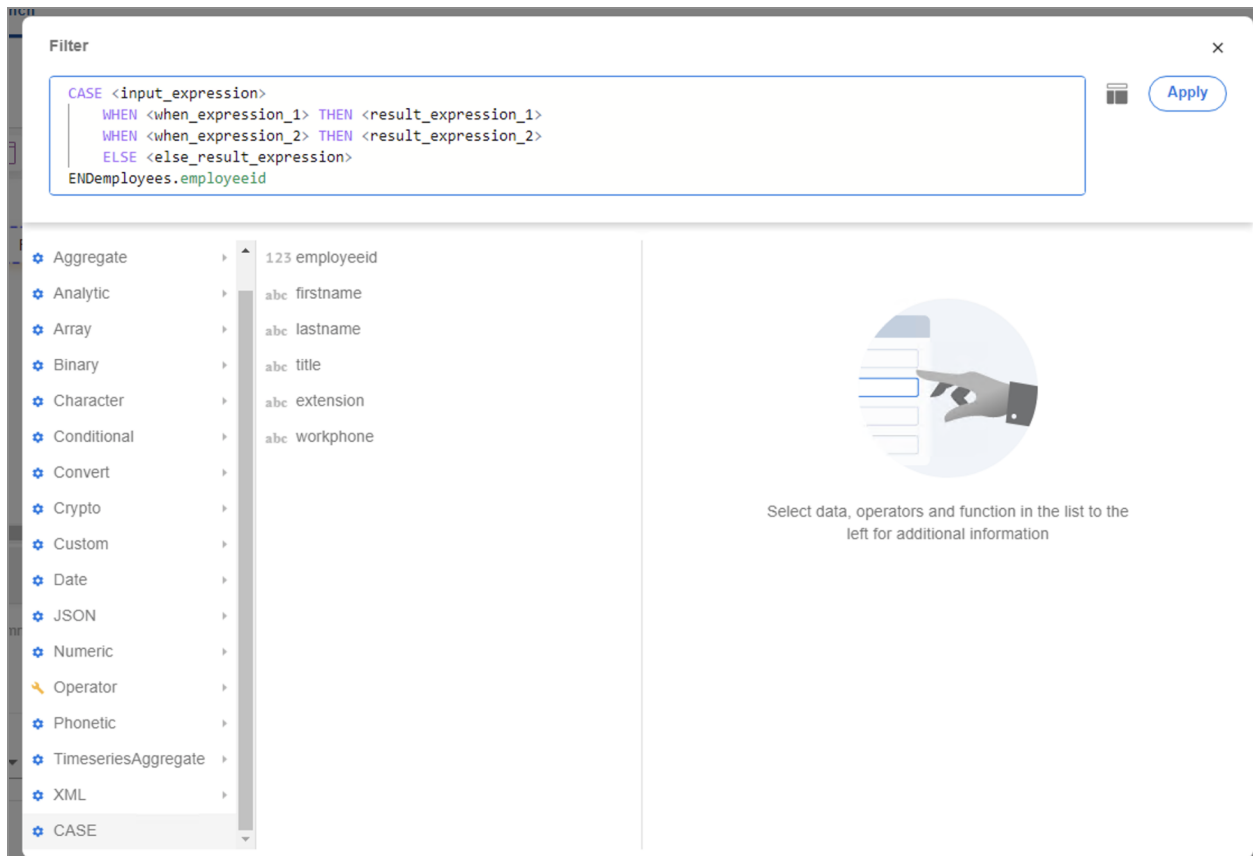
The Expression editor can be accessed through the Filter, Columns or JOIN operators. You can build your own arbitrary expressions using the Expression editor. The following window is displayed when you choose to write your own expression from the Columns/Filter operator's configuration pane.



A list of all the supported functions and a syntax help is provided in the Expression Editor window. Follow these steps to manually 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. Double-click on 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.

In the list of functions that is available to you to choose from, there is also a CASE statement that you can select which will generate an empty structure to give syntax assistance.



Publishing a Flow

To publish a data flow, follow these steps:

Drag and drop the Publish node from the operations palette into the canvas.

1. 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:

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.

1. 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.

It is recommended that you do not create an empty or invalid flow.

TIBCO Documentation and Support Services

For information about this product, you can read the documentation, contact TIBCO Support, and join TIBCO Community.

How to Access TIBCO Documentation

Documentation for TIBCO products is available on the [Product Documentation website](#), mainly in HTML and PDF formats.

The [Product Documentation website](#) is updated frequently and is more current than any other documentation included with the product.

Product-Specific Documentation

Documentation for TIBCO® Data Virtualization is available on the [TIBCO® Data Virtualization Product Documentation](#) page.

To directly access documentation for this product, double-click the following file:

`TIBCO_HOME/release_notes/TIB_TDV_8.7_docinfo.html` where `TIBCO_HOME` is the top-level directory in which TIBCO products are installed. On Windows, the default `TIBCO_HOME` is `C:\tibco`. On UNIX systems, the default `TIBCO_HOME` is `/opt/tibco`.

The following documents for this product can be found in the TIBCO Documentation site:

Users

TDV Getting Started Guide

TDV User Guide

TDV Client Interfaces Guide

TDV Web UI Users 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 Contact Support for TIBCO Products

You can contact the Support team in the following ways:

- To access the Support Knowledge Base and getting personalized content about products you are interested in, visit our [product Support website](#).
- To create a Support case, you must have a valid maintenance or support contract with a Cloud Software Group entity. You also need a username and password to log in to the our [product Support website](#). If you do not have a username, 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 [TIBCO Community](#).

Legal and Third-Party Notices

SOME CLOUD SOFTWARE GROUP, INC. (“CLOUD SG”) SOFTWARE AND CLOUD SERVICES EMBED, BUNDLE, OR OTHERWISE INCLUDE OTHER SOFTWARE, INCLUDING OTHER CLOUD SG SOFTWARE (COLLECTIVELY, “INCLUDED SOFTWARE”). USE OF INCLUDED SOFTWARE IS SOLELY TO ENABLE THE FUNCTIONALITY (OR PROVIDE LIMITED ADD-ON FUNCTIONALITY) OF THE LICENSED CLOUD SG SOFTWARE AND/OR CLOUD SERVICES. THE INCLUDED SOFTWARE IS NOT LICENSED TO BE USED OR ACCESSED BY ANY OTHER CLOUD SG SOFTWARE AND/OR CLOUD SERVICES OR FOR ANY OTHER PURPOSE.

USE OF CLOUD SG SOFTWARE AND CLOUD SERVICES IS SUBJECT TO THE TERMS AND CONDITIONS OF AN AGREEMENT FOUND IN EITHER A SEPARATELY EXECUTED AGREEMENT, OR, IF THERE IS NO SUCH SEPARATE AGREEMENT, THE CLICKWRAP END USER AGREEMENT WHICH IS DISPLAYED WHEN ACCESSING, DOWNLOADING, OR INSTALLING THE SOFTWARE OR CLOUD SERVICES (AND WHICH IS DUPLICATED IN THE LICENSE FILE) OR IF THERE IS NO SUCH LICENSE AGREEMENT OR CLICKWRAP END USER AGREEMENT, THE LICENSE(S) LOCATED IN THE “LICENSE” FILE(S) OF THE SOFTWARE. USE OF THIS DOCUMENT IS SUBJECT TO THOSE SAME TERMS AND CONDITIONS, AND YOUR USE HEREOF SHALL CONSTITUTE ACCEPTANCE OF AND AN AGREEMENT TO BE BOUND BY THE SAME.

This document is subject to U.S. and international copyright laws and treaties. No part of this document may be reproduced in any form without the written authorization of Cloud Software Group, Inc.

TIBCO, TIBCO logo, TIBCO O logo, ActiveSpaces, Enterprise Messaging Service, Spotfire, TERR, S-PLUS, and S+ are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or other countries.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle Corporation and/or its affiliates.

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. You acknowledge that all rights to these third party marks are the exclusive property of their respective owners. Please refer to Cloud SG’s Third Party Trademark Notices (<https://www.cloud.com/legal>) for more information.

This document includes fonts that are licensed under the SIL Open Font License, Version 1.1, which is available at: <https://scripts.sil.org/OFL>

Copyright (c) Paul D. Hunt, with Reserved Font Name Source Sans Pro and Source Code Pro.

Cloud SG 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 a specific version of Cloud SG software on a specific operating system platform.

THIS DOCUMENT 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 DOCUMENT 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 DOCUMENT. CLOUD SG MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S), THE PROGRAM(S), AND/OR THE SERVICES DESCRIBED IN THIS DOCUMENT AT ANY TIME WITHOUT NOTICE.

THE CONTENTS OF THIS DOCUMENT MAY BE MODIFIED AND/OR QUALIFIED, DIRECTLY OR INDIRECTLY, BY OTHER DOCUMENTATION WHICH ACCOMPANIES THIS SOFTWARE, INCLUDING BUT NOT LIMITED TO ANY RELEASE NOTES AND "README" FILES.

This and other products of Cloud SG may be covered by registered patents. For details, please refer to the Virtual Patent Marking document located at <https://www.tibco.com/patents>.

Copyright © 2002-2023. Cloud Software Group, Inc. All Rights Reserved.