

TIBCO WebFOCUS®

Reporting Server and TIBCO® Data Migrator Release Notes

*Release 8207
April 2021
DN4501691.0321*



Contents

1. New Features	11
Server Enhancements	11
Global WebFOCUS Server Monitoring Console.....	11
Server Text Editor Supports Jump to Next or Previous Difference in Compare and Merge. .	14
Upload Message Added in the Edaprint Log.....	15
Displaying the Number of Times an Agent Has Been Used.....	16
Using Server Variables to Load Third-Party Drivers.....	16
Comparing or Merging Server Files.....	17
Comparing Files.....	18
Merging Files.....	19
Using Custom Folding in the Text Editor.....	19
Using Text Editor Tab Key Navigation.....	20
Sharing Application Directories and Files.....	21
Using the Git Diff Tool.....	24
All Adapters	24
Generating Table and Column Names With DBMS-Specific Length Limits.....	24
SQL Adapters	25
Cross-Join Performance Improvements.....	25
Cross Database Join Performance Improvements.....	30
Direct Bulk Load Support.....	33
Navigating Joins Between Cluster Synonyms.....	34
SQL Adapters: Optimizing OUTPUTLIMIT.....	38
Adapter for Amazon Athena: Creating Files in Parquet or ORC Format.....	39
Adapter for Apache Drill: Using Parquet as a Transfer File to HDFS.....	40
Adapter for Google BigQuery: Support for SQL Immediate Commands.....	42
Adapter for Google BigQuery: Support for Service Account Authentication.....	42
Adapters for Hive and Impala: Using Parquet as a Transfer File to HDFS.....	43
Adapter for Hive: Optimize Load With Bulk Load Options.....	45
Adapter for Hive: Configuring Bulk Load for HDFS Transfer Protocol.....	46
Adapter for Hive: Bulk Loading Data With Non-Printable Characters.....	46
Adapter for Hive: Support for New Transactions With ORC-Formatted Records.....	49
Adapter for Hyperstage (PG): Installation Setting for Heap Size.....	51

Adapters for Impala and Crossdata: SET DEFAULTSCHEMA.....	52
Adapter for MySQL: Support for Running Under an SSH Tunnel.....	52
Adapter for Oracle: Support for Instant Client Basic Package.....	53
Adapter for Oracle: Support for Version 19c.....	53
Adapter for Oracle: Support for Autonomous Data Warehouse Cloud.....	53
Adapter for Oracle: Support for HINT on Segment Level.....	53
Adapter for Stratio Crossdata: Support for TIMESTAMP Data Type.....	53
Adapter for TIBCO SnappyData.....	54
Adapter for SQL Server: Version 2019 Support.....	54
ERP Adapters.....	54
Adapter for Microsoft Dynamics CRM: New Mappings for Long String/Memo and Image Data Types.....	54
Adapter for Microsoft Dynamics CRM: Passing Client Credentials in the Body.....	56
Adapter for Microsoft Dynamics CRM: Support for Token Parameters With OAuth Grant Password Authentication.....	57
Adapter for Microsoft Dynamics CRM: Support for OAuth Grant Password Authentication..	58
Adapter for MS Dynamics CRM: Support for Printing LOOKUP Fields.....	60
Adapter for Salesforce.com: Creating Custom Objects.....	60
Adapter for Salesforce.com: Saving Intermediate and Log Files.....	60
Procedures Adapters.....	62
Adapter for REST: Populating XDEFAULT Attributes in a JSON Synonym for a POST Request.....	62
Adapter for REST: Additional Token Parameters for OAuth Client Credentials.....	66
Adapter for REST: Support for Token Parameters With OAuth Grant Password Authentication.....	66
Adapter for REST: Passing Multiple Values to a Parameterized URL in a WHERE Phrase..	67
Adapter for REST: Support for CSV Format in a REST Response.....	68
Adapter for REST: Support for No Value in a Label in the JSON Body.....	72
Adapter for REST: Adding Scope to OAuth Access Token Request.....	73
Adapter for REST: Encoding Special Characters When Creating a Synonym.....	73
Adapter for REST: Sending Credentials in the POST Request Body.....	75
Adapter for REST: Sending Multipart/Form Data in a POST Request.....	76
Adapter for REST: Support for STRING Format.....	77

Adapter for REST: Setting a Default Field Length..... 77

Adapter for REST: Support for Passing Text and Binary in a REST POST Request..... 78

Search Engine Adapters 79

 Adapter for ElasticSearch..... 80

Social Media Adapters 80

 Adapter for Google Analytics: Support for Service Account Authentication..... 80

 Adapter for Google Drive..... 81

 Adapter for Google Sheets..... 81

 Generating a Google Sheets Spreadsheet..... 81

 Adapters for Google Sheets and Google Drive: Service Account Authentication.....81

XML Adapters 82

 Adapters for SharePoint and SharePoint Drive: Passing Client Credentials in the Body... 82

 Adapters for SharePoint and SharePoint Drive: Support for Token Parameters With
 OAuth Grant Password Authentication.....85

 Adapters for SharePoint and SharePoint Drive: Support for OAuth Grant Password
 Authentication..... 86

 Adapter for SharePoint: Create Synonym Support for Multiple Sites in List Mode 90

 Adapter for SharePoint.....91

 Adapter for OData: Passing Client Credentials in the Body..... 92

 Adapter for OData: Additional Token Parameters for OAuth Client Credentials..... 93

 Adapter for OData: Support for Token Parameters With OAuth Grant Password
 Authentication..... 93

 Adapter for OData: Support for Date-Time Functions in Filters..... 94

 Adapter for OData: Support for Reformatting Fields..... 96

 Adapter for OData: Support for MISSING..... 96

 Adapter for OData: Turn Off Aggregation for ENTITY_SET Requests..... 97

 Adapter for OData: Starting Retrieval With the Child Object in a Static Join..... 97

Remote Storage Adapters 101

 Adapter for AWS S3.....101

 Adapter for Amazon Web Services: Loading a File From AWS S3 Directly to Amazon RDS 101

 Adapter for Google Drive: Support for Google Slides..... 103

 Adapter for SharePoint Drive..... 105

General Data Flow Enhancements 105

Viewing Flow Properties on the Data Flow Canvas.....	106
Joining to a Variable in a Data Flow.....	107
Viewing Checkpoint Files Generated by a Change Data Capture (CDC) Process.....	112
Migrating Data Flows.....	118
Showing Whether a Column Supports Nulls in the Select, Merge, and Synonym Editors..	123
Controlling Format Conversion Error Processing in a Data Flow.....	125
Enhancements to the Union Editor.....	128
Showing the Confidence Level for Adding a Union to a Data Flow.....	133
Support for Dragging Targets Into a Data Flow.....	134
Creating Advanced Filters.....	136
Distinguishing Node Status Using Ports and Connectors.....	139
Using Undo and Redo in Join, Union, and Business View Editors.....	142
Function Assist For Mathematical Operators.....	143
Displaying Database-Specific Icons for Sources in a Data Flow.....	144
Displaying Database-Specific Icons for Data Sources in the Join Editor.....	145
Expression Editor Shows Object Types.....	146
View Source Opens in Full Screen.....	147
Editing Joins in Full Screen.....	148
Union, Join, and Select Editors Open as Nearly Full-Screen Windows.....	149
Limiting the Number of Automatic Join Pairs in Tables Without Keys.....	150
Automatically Deriving a Field Name When Creating an Expression.....	151
Selecting Join or Union When Dragging a Source Onto a Flow.....	152
Choosing a Join or Union for a Data Flow From a Menu.....	152
Previewing the Display Format for a Function.....	154
Improving Performance of UNION and UNION ALL in a Data Flow.....	154
Selecting the Type of UNION in a Data Flow.....	155
Generate Flows Enhancements	156
Generate Flows Automatically Completes Target Names.....	157
Generate Flows Saves the Last Used Selections.....	158
Generate Flows Shows the Confidence Level for Existing Targets.....	160
Generate Flows Supports Template Flows With CDC Sources.....	162
Generate Flows Supports Template Flows With an SCD Target.....	163
Enhancements to Generating Flows Using a Template Flow.....	164

Generating Data Flows Using a Template Flow.	176
Upload, Connection, and Synonym Enhancements	185
Classifying Metadata During Upload and Using it in a Data Flow.	185
Prerequisites for Metadata Classification.	186
Using Metadata Classifications in a Union in a Data Flow.	189
New Setting for Controlling Display of Unreferenced Business View Fields.	194
Empty Application Message in the Server Console.	195
Enhanced Options for Creating Cluster Joins.	196
Displaying the Remote Servers List Collapsed on the Get Data Advanced Page.	199
Directly Accessing the Advanced Mode of Get Data.	200
Descriptions and Examples Now Shown in the Decompose Date Dialog Box.	202
Supporting Excel Targets for Upload and Data Flow.	202
Setting the Maximum Size of Files for Upload.	204
Supporting Full Outer Joins When the Data Source Does Not.	205
Changing Target Settings When Uploading Files.	206
Redesign of Business View Icons.	207
Create Synonym Panel Redesign in Get Data Advanced Mode.	208
Previewing a File When Uploading.	209
Creating and Renaming a New Folder in a Business View.	210
Creating a Business View for a Target Data Source.	212
Pivoting Multiple Columns on Upload.	214
Target and Load Options Enhancements	215
Adapter for Apache Hive: Creating Avro Files.	215
Using a Delimited Target When No Relational Adapter is Available.	217
Selecting Excel as the Default Target Adapter in a Flow.	218
Support for New Data Flow Targets With No Key Fields.	219
Displaying Messages at the Top of the Load Options Screen.	221
Show Confidence Level When Selecting an Existing Target in a Flow.	222
Using Change Data Capture/Slowly Changing Dimensions Load Type on a Flow.	222
Supporting Direct Bulk Load for a Flow.	223
Data Profiling, Analysis, Sample Data, and Sampling Enhancements	227
Using Bulk Load to Insert Sample Data.	228
Enhanced Sampling on Joined Tables.	228

Key Analysis Using Sample Data.....	230
Using Stratified Sampling.....	234
Sample Data Shows Warning Messages.....	236
Sampling Shows Warning Messages.....	237
Specifying a Location for Sampling Data.....	239
Staging All Sources When Sampling in a Data Flow.....	240
Profiling Numeric Data Using Standard Deviation.....	241
SQL Statement and Functions Enhancements.....	244
Returning an Ordered Answer Set of a Limited Number of Rows in a Sub-Select.....	244
Support for GROUP BY On Column Position or AS Name.....	248
Using SQL Analytic Functions.....	250
Data Migrator Enhancements.....	255
Assigning Slowly Changing Dimension (SCD) Columns in the Web Console.....	255
Adapter for PostgreSQL: Support for Change Data Capture (CDC).....	260
Preparing the Environment for CDC With PostgreSQL.....	261
Using Logical Decoding and Monitoring Slots.....	263
Configuring the Adapter for PostgreSQL and Using CDC in a Flow.....	266
Support for Direct Load Flows to Load AWS S3 Data to PostgreSQL.....	267
Creating a Synonym in a Data Flow Using Variables for the Existing Excel File Name.....	270
New Option to Save Data Files When Bulk Load is Used in a Flow.....	273
Slowly Changing Dimensions (SCD): Re-run for Missing Date.....	273
Slowly Changing Dimensions: Enhanced Processing for Type I Columns.....	276
Change Data Capture (CDC) on a Data Flow.....	276
Slowly Changing Dimensions Activation Flag Enhancement.....	281
Slowly Changing Dimensions on a Data Flow.....	283
Reporting Language Enhancements.....	284
International System (SI) Numeric Format Abbreviation Options.....	284
New Functions for Date-time Conversion Between Local and UTC Time.....	285
DT_TOUTC: Converting Local Time to UTC Time.....	287
DT_TOLOCAL: Converting UTC Time to Local Time.....	289
New Regular Expression Pattern Matching Functions.....	290
REGEXP_COUNT: Counting the Number of Matches to a Pattern in a String.....	291
REGEXP_INSTR: Returning the First Position of a Pattern in a String.....	291

REGEXP_REPLACE: Replacing All Matches to a Pattern in a String..... 292

REGEXP_SUBSTR: Returning the First Match to a Pattern in a String..... 293

Support for Functions Used in ODBC Connector Client Tools..... 294

Support for Standard Deviation in PARTITION_AGGR..... 298

Raised Limit for Join Fields.....301

Standard Deviation Prefix Operators: STDP. and STDS.....301

Enhancements to DATE_ORDER..... 303

2. Dropped Support 305

 Dropped Support: Adapter Releases 305

 Dropped Support: Operating System Releases306

 DataMigrator: Dropped Support for Perforce306

3. Upgrade Notes 307

 Golden Key License Mode 307

 Converting an Existing Server Installation to the Golden Key Configuration308

 Software Branding309

 Technical Content Branding311

 Cloud312

 Microsoft Azure SQL Data Warehouse Renamed 312

 Running Scheduled Flows 312

 SSL: New OpenSSL Configuration Requirement313

 Data Flow: Warning When Closing Join, Union, SQL, and Target Business View Editors 314

 Data Preparation: Properties Panel Redesign314

 Data Preparation: Server Console Menu 315

 Data Profiling Charts and Reports Display in the Output Window 316

 Data Preparation Upgrade Considerations 316

4. Known Issues 317

 New Licenses For Second Level Geographies Not Yet Supported317

 Client Silent Installation Failures on Windows 317

5. Fixes 319

 Entering Network Path in File Picker for Excel Causes Error319

Legal and Third-Party Notices 321

Additional Third-Party Legal Notices 323

New Features

The following New Features were added to the TIBCO WebFOCUS® Reporting Server and TIBCO® Data Migrator in Release 8207.

In this chapter:

- [Server Enhancements](#)
 - [All Adapters](#)
 - [SQL Adapters](#)
 - [ERP Adapters](#)
 - [Procedures Adapters](#)
 - [Search Engine Adapters](#)
 - [Social Media Adapters](#)
 - [XML Adapters](#)
 - [Remote Storage Adapters](#)
 - [General Data Flow Enhancements](#)
 - [Generate Flows Enhancements](#)
 - [Upload, Connection, and Synonym Enhancements](#)
 - [Target and Load Options Enhancements](#)
 - [Data Profiling, Analysis, Sample Data, and Sampling Enhancements](#)
 - [SQL Statement and Functions Enhancements](#)
 - [Data Migrator Enhancements](#)
 - [Reporting Language Enhancements](#)
-

Server Enhancements

This section describes new features for server monitoring and administration.

Global WebFOCUS Server Monitoring Console

Cluster Manager (CLM) enables you to balance workload across a number of servers, by configuring any number of remote server nodes under one cluster name. CLM has polling and monitoring capabilities to identify which servers are up-and-running, average query response time, and the number of concurrent tasks running. CLM can also identify which remote servers are down, and appropriately re-route requests, as well as route a user connection to the most efficient remote server to enhance overall performance.

However, in prior releases, you had to open each remote server node separately in its own browser window in order to monitor and manage its agents. Now, CLM provides a Global Monitoring Console for all remote server nodes in a cluster, as well as remote servers defined dynamically.

A dynamic Remote Server is defined with the following setting in its edaserve.cfg file.

```
clm = http://host:port
```

where:

```
host:port
```

Is the hostname and port for connecting to the server that will have this server on its Remote Servers list.

To see the remote servers that are defined, navigate to the Workspace page of the Server Web Console, right-click *Cluster Manager*, and click *Remote Servers*. The list of cluster servers and dynamic servers opens, as shown in the following image.

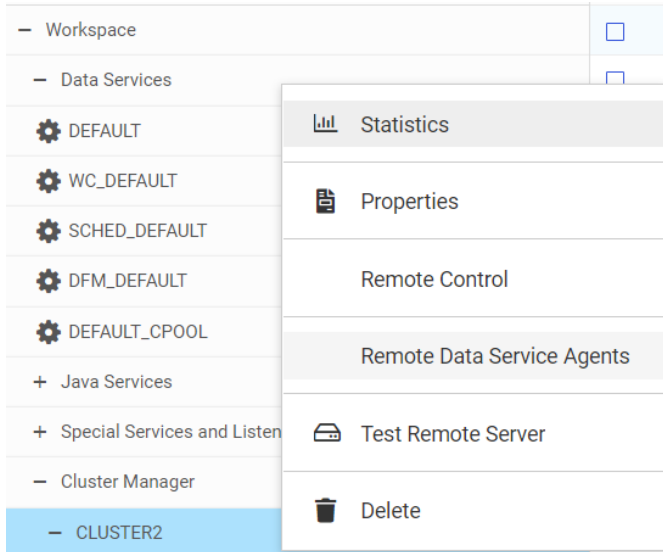
Type	Status	Cluster/Server	Agents Now	Agents Peak	Connections Total	Connections Completed	Avg. Response Time	Running Avg. Response Time
		CLUSTER2	8	6	8	8	0.128 seconds	0.000 seconds
	running	SOLSPARC	2	1	2	2	0.075 seconds	0.000 seconds
	running	SOLSPC11	2	1	2	2	0.281 seconds	0.000 seconds
	running	aixppc72.8981	2	2	2	2	0.105 seconds	0.000 seconds
	running	Inxx64r8.lbi.com.8981	2	2	2	2	0.051 seconds	0.000 seconds

In this image, the first two servers are part of CLUSTER2. The other two are dynamic remote servers.

You can open the Global Monitoring Console from this list, by right-clicking the cluster node and clicking *Remote Data Service Agents*, as shown in the following image.

Type	Status	Cluster/Server	Agents Now	Agents Peak	Connections Total	Connections Completed	Avg. Response Time	Running Avg. Response Time	Queued Now	Queued Peak
		CLUSTER2	8	6	8	8	0.110 seconds	0.000 seconds	0	0
	running	SOLSPARC	2	1	2	2	0.086 seconds	0.000 seconds	0	0
	running	SOLSPC11	2	1	2	2	0.193 seconds	0.000 seconds	0	0
	running	Inxx64r8.lbi.com.8981	2	2	2	2	0.054 seconds	0.000 seconds	0	0
	running	aixppc72.8981	2	2	2	2	0.107 seconds	0.000 seconds	0	0

You can also access the Global Monitoring Console from the Resources Pane. Expand *Cluster Manager*, right-click the cluster name, and click *Remote Data Service Agents*, as shown in the following image.



The Global Monitoring Console is shown in the following image. CLUSTER2 has two remote server nodes configured, both of which can be monitored from the Global Monitoring Console. The two dynamic remote server nodes that are not part of the cluster can also be monitored.

The screenshot shows the 'Remote Data Service Agents' tab in the Global Monitoring Console. It displays a table of agents grouped by cluster. Each cluster has a 'Kill Selected Agents' button and a search field.

Cluster	Tscomid	Service	State	Command	Focus I/O	DBMS I/O	DBMS Time	Query Time	CPU Time	Memory Usage (KB)	Disk Usage (KB)	Process ID
SOLSPARC	1	DEFAULT	idle	TABLE	0	1	0:00:00.000	Nov 09 11:30:14.131	0:00:00.296	99336	---	12254
	2	DEFAULT	idle	---	---	---	---	---	0:00:00.218	84248	---	12255
	3	WC_DEFAULT	idle	---	---	---	---	---	0:00:00.001	04740	---	12256
SOLSPC11	2	DEFAULT	idle	---	---	---	---	---	0:00:00.143	94784	---	23644
	3	WC_DEFAULT	idle	---	---	---	---	---	0:00:00.141	94784	---	23645
	4	WC_DEFAULT	idle	---	---	---	---	---	0:00:00.143	04784	---	23646
	5	WC_DEFAULT	idle	---	---	---	---	---	0:00:00.143	04784	---	23647
alxppc72	1	DEFAULT	idle	TABLE	0	1	0:00:00.000	Nov 09 11:28:31.043	0:00:00.000	17640	---	32375242
	2	DEFAULT	idle	TABLE	0	1	0:00:00.000	Nov 09 11:28:31.141	0:00:00.000	17640	---	18481906
	3	WC_DEFAULT	idle	---	---	---	---	---	0:00:00.000	00000	---	18481907
lmxx64r8	2	DEFAULT	idle	---	---	---	---	---	0:00:00.030	314144	---	37067

Each node is represented in its own table, the same interface used for individual data services, so you can kill selected agents for any node using this console. In addition, you can track a request from client to server, no matter which node the request is routed to.

If a remote server is stopped or becomes unavailable, the Global Monitoring Console will reflect the change.

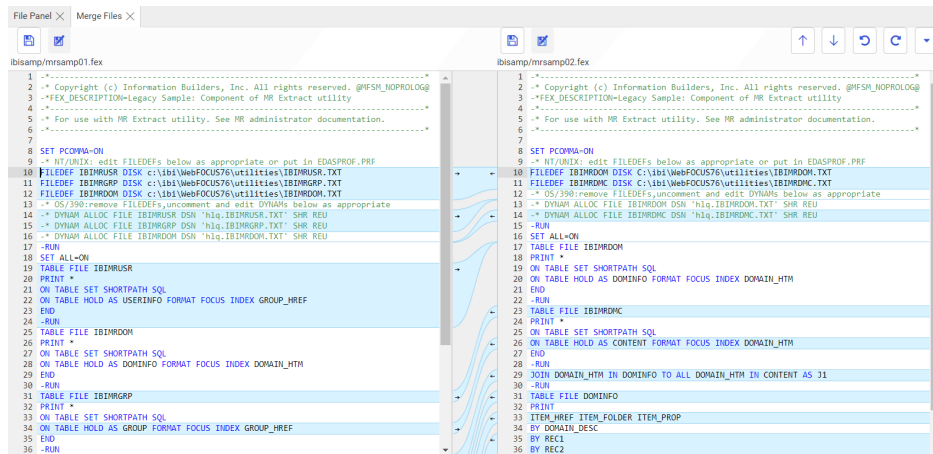
Server Text Editor Supports Jump to Next or Previous Difference in Compare and Merge

In the Server Console, the Text Editor Compare Files and Merge Files pages now have *Jump to Next Difference* and *Jump to Previous Difference* buttons.

These buttons enable you to see connected differences easily, without scrolling.

To compare or merge files, multi-select two files, right-click, and select either *Compare Files* or *Merge Files*. Although both can show the differences, only Merge Files enables you to move content between the files. Compare Files is read-only.

The files open with left and right scrolling locked to the first difference, aligned in the middle of the window (when the file sizes allow), as shown in the following image for Merge Files.



You can use the following buttons to move between differences and control the merge process. The arrows for the current difference will be aligned and displayed with left and right scrolling locked.

Button Image	Description
	<p>Save File (Merge only)</p> <p>Each file has its own Save icon (Save Left and Save Right), so that you can save the file when you are finished with the merge edits for each.</p>
	<p>Save File As (Merge only)</p> <p>Each file has its own Save As icon (Save Left As and Save Right As), so that you can save the file in a new folder and/or under a new name when you are finished with the merge edits for each.</p>
	<p>Jump to Previous Difference</p> <p>Locks and aligns the left and right scroll bars to the previous difference, if there is one.</p>
	<p>Jump to Next Difference</p> <p>Locks and aligns the left and right scroll bars to the next difference, if there is one.</p>
	<p>Undo (Merge Only)</p> <p>Undoes the previous merge action.</p>
	<p>Redo (Merge Only)</p> <p>Redoes the previously undone merge action.</p>

Upload Message Added in the Edaprint Log

An informational message similar to the following has been added to the edaprint.log file for each upload request:

```
I request to upload file=station_zip.csv,size=24398,u=PTH\user1
```

This enables you to generate a list of uploaded files from the edaprint.log showing the user ID (with security provider), name, size, and date and time of each upload.

Displaying the Number of Times an Agent Has Been Used

The Data Service Agents page, available on the Workspace page of the Server Console, now has a column named *Connections* that shows the number of connections that have been used for each agent, as shown in the following image.

Tscomid	Service	State	Connections	User	Group	Client Address	Scheduler Job ID	Command	Procedure	Top Level Procedure	Master File	Focus ID	DBMS
11	WC_DEFAULT	idle	16	--	--	--	--	APP	--	cltbike/foctrans2	HOLD	0	715
12	WC_DEFAULT	idle	1	--	--	--	--	TABLEF	--	--	HOLD	0	74
16	WC_DEFAULT	idle	1	--	--	--	--	--	--	--	--	0	0
13	SCHED_DEFAULT	idle	0	--	--	--	--	--	--	--	--	--	--
14	SCHED_DEFAULT	idle	0	--	--	--	--	--	--	--	--	--	--
18	DEFAULT	idle	0	--	--	--	--	--	--	--	--	--	--
3	DEFAULT	idle	0	--	--	--	--	--	--	--	--	--	--
4	DEFAULT	idle	0	--	--	--	--	--	--	--	--	--	--

You can compare this number to the maximum number set for each service. The maximum number of agents allowed for each service can be changed by right-clicking the service and clicking *Properties*.

Using Server Variables to Load Third-Party Drivers

Server environment variables can now be used to add third-party API and Driver DLL locations to the server library path. You can declare them as environment variables before server startup (using applicable operating system syntax and then exporting them, if applicable to your environment), or add them to the server environment variable configuration file (edaenv.cfg, using name=value syntax, with one variable per line).

Normally, only a single directory would need to be assigned to each of these variables. If necessary, you can assign multiple directories. You must separate the directories using the native path directory separator for your platform, with no spaces before or after each separator character. The separator character is a semi-colon (;) for Windows, and a colon (:) for all other platforms.

The following variables are supported:

- IBI_LOAD_SSL_FROM** can be set to an alternate SSL and CRYPTO DLL location.
- IBI_LOAD_MSSQL_ODBC_DRIVER_FROM** can be set to the Microsoft SQL Server and Azure ODBC driver location.
- IBI_LOAD_ORACLE_CLIENT_FROM** can be set to the Oracle Client DLL library location.

- IBI_LOAD_REDSHIFT_ODBC_DRIVER_FROM** can be set to the Redshift ODBC driver location.
- IBI_LOAD_EXASOL_ODBC_DRIVER_FROM** can be set to the EXASolution ODBC driver location.
- IBI_LOAD_DB2_CLI_CLIENT_FROM** can be set to the Db2 CLI Client DLL library location.

Note: In prior releases, the use of the LD_LIBRARY_PATH, LIBPATH or IBI_LIBPATH variable was recommended to specify the library path for loading third-party DLLs, depending on whether system security was in use. While these methods continue to work, they are deprecated in favor of using these new variables.

If you are setting any of these variables in the edaenv.cfg file, perform the following steps:

1. Go to the Workspace page on the Server Console.
2. Expand *Configuration Files*, then *Miscellaneous*.
3. Right-click *Environment - edaenv.cfg*, and click *Edit*.
4. When you have made the edits, click the *Save* button.
5. Restart the server in order to have these settings take effect.

To restart the server, on the Workspace page click *Server Actions*, then *Restart*.

Comparing or Merging Server Files

Using the server Compare Files and Merge Files options, you can compare or merge two files in the server text editor.

To open the diff tool:

1. Multi-select two files by clicking one file, holding down the Ctrl key (or the Shift key if the files are contiguous) and clicking the second file.
2. Right-click and select:
 - Compare Files* to compare the files.
 - Merge Files* to copy lines between the files.

The two files open side-by-side in the text editor.

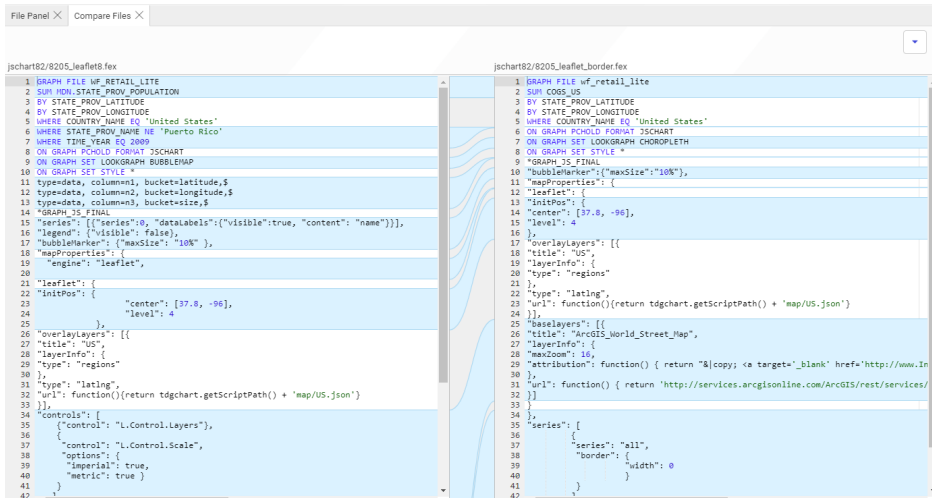
Note: If the two files represent synonyms, selecting:

- Compare Files* compares the Master Files.
- Compare Access Files* compares the Access Files.

- ❑ Merge Files merges the Master Files.
- ❑ Merge Access Files merges the Access Files.

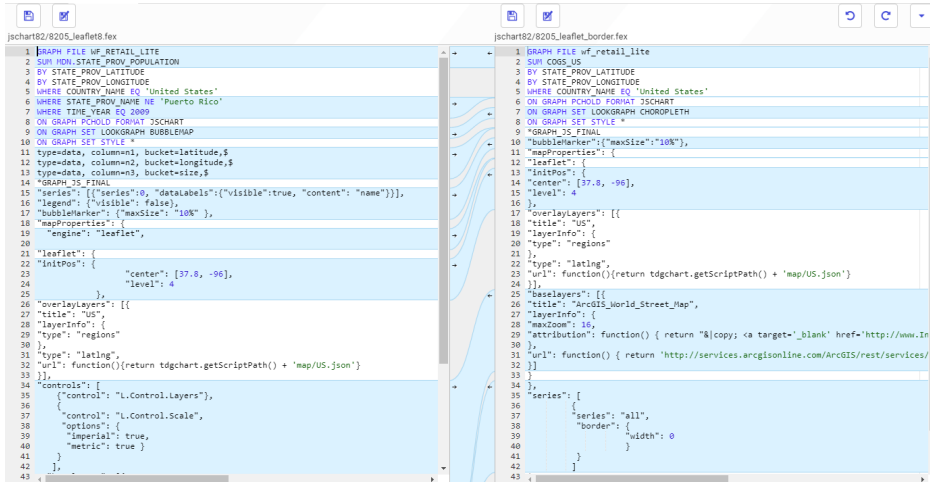
Comparing Files

When you select Compare Files, the files open in two panes in the text editor. Lines that are different are highlighted in light blue, and curves drawn between the two panes show insertions and deletions, as shown in the following image.



Merging Files

When you select Merge Files, the files open in two panes in the text editor. Lines that are different are highlighted in light blue, and curves drawn between the two panes show insertions and deletions. In addition, arrows within the curves enable you to copy lines between the files, as shown in the following image.



If you click an arrow to copy a line from one file to the other, the highlighting and curve disappear once the line is copied.

Each pane has a Save As button, so you can save your changes, under the same name or a different name.


There are undo and redo buttons so you can reverse your changes, if necessary. If you undo a change, the highlighting and curve reappears.

Using Custom Folding in the Text Editor

Custom folding enables you to choose the type of commands that you want to remain unfolded. This is especially useful when you open the Session Log, which contains all of the commands generated in your session.

To enable custom folding, click the down arrow on the text editor toolbar, and click *Custom Folding...*

The Custom Folding dialog box opens, as shown in the following image.



Check the boxes for the commands that you *do not* want to be folded. For example, DEFINE is checked for the following image that shows a file with comments, SET commands, a DEFINE command, and TABLE request.

```
1  
8 * DEFINE FILE retail8205/bv_namespace_off/wf_retail_lite_off ADD  
9 TIME_YEAR_GROUP/16 TITLE "Year,group"=IF TIME_YEAR IN (2008, 2009, 2010, 2011) THEN 0 ELSE IF TIME_YEAR IN (2012, 2013, 2014) THEN 1 ELSE TIME_YEAR;  
10 END  
11 *
```

If you previously implemented custom folding and want to remove it, click *Remove Custom Folding*. You can also unfold any folded text by clicking it.

Click *OK*.

Using Text Editor Tab Key Navigation

The text editor Tab key navigation supports Section 508 Accessibility Compliance by enabling users to navigate through the interface using the Tab key on the keyboard.

To implement Tab key navigation, either press the *Esc* key when in a file open in the text editor, or click the Tab Navigation button on the toolbar that is shown in the following image:



Once you have enabled Tab key navigation, the icon on the button becomes hollow in the center, as shown in the following image.



With Tab key navigation enabled, pressing Tab cycles between the features on the page.


- If you press *Tab* from within a file open in the text editor, the focus moves to the URL bar at the top of the web page, then to the actual URL.
- Continuing to press Tab cycles to the buttons on the ribbon, with a change in color intensity identifying which button has the focus.
- After the ribbon buttons, the focus moves to the data tree, with a cursor or outline showing the position of the focus.
- From there, the focus moves to the text editor toolbar, indicated by a rounding of the edges of the button that has the focus.
- Pressing *Enter* selects the option that has the focus.

Sharing Application Directories and Files

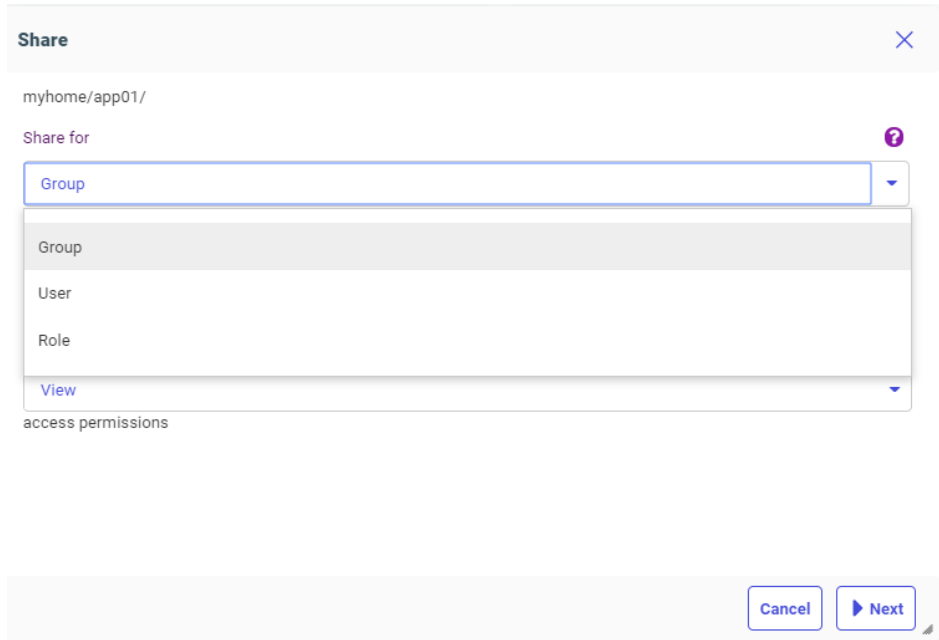
Users can share their personal content under their *myhome* application with other users, groups, or roles.

To share a resource under your myhome application:

1. Right-click the resource you want to share, and click *Share* on the context menu. The *Manage Share Privileges* page opens, as shown in the following image.



2. Click the plus sign (+) and select *User*, *Group*, or *Role* from the *Share for* drop-down list, as shown in the following image.

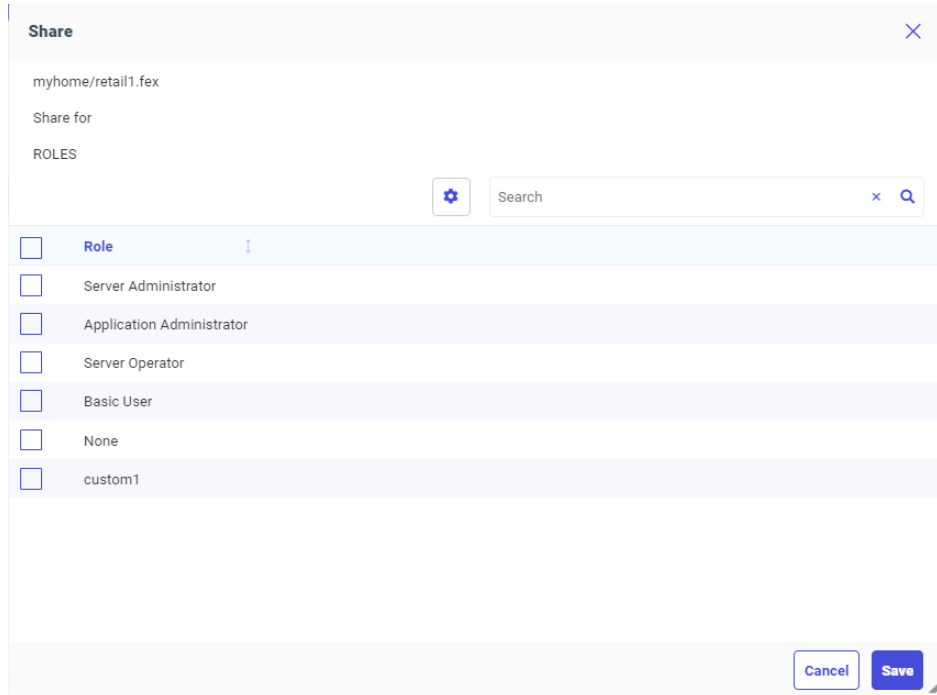


3. Select the security provider from the *Security Provider* drop-down list.
4. Select the type of access allowed, either *View* or *Edit*, from the *Access* drop-down list.

- Click *Next*.

The *Share* dialog box opens, as shown in the following image.

- Check the boxes for the users, groups, or roles that will share this resource, as shown in the following image, and click *Save*.



- To manage the shared resource, right-click it and click *Share*. The *Manage Share Privileges* page opens, showing the share properties that have been configured.

You can edit or delete existing share properties, add additional share properties, save your updates, or reset the page.

When a user with share privileges signs into the server, an application named *shared with me* (*users home*) is added to the application tree and provides access to the shared resources.

Using the Git Diff Tool


Once you have configured the Adapter for Git, you can use the Git Diff tool to open two versions of the same file in side-by-side panels, or generate a list of changes to a folder.

To open a file in the Git Diff tool, right-click the file, point to *Git* on the context menu, click *Diff*, and select one of the following options:


- Not staged.** Opens the staged version and modified (Working Directory) version.
- Staged.** Opens the staged version and last committed version.
- With Last Commit.** Opens the previously committed version and last committed version

The two versions of the file open in side-by-side panels. The differences are identified by a distinctive background color.

The following image shows the staged and modified version of the file prc01 in the diff tool.



The following image shows changes to the folder named test.



If there were no changes to the folder, the following message displays:

`There are no differences`

All Adapters

This section describes new features that apply to all adapters.

Generating Table and Column Names With DBMS-Specific Length Limits

When you use Get Data to upload files, a Data Flow to load data to a target, or Quick Copy to copy data from one DBMS to another, the column names generated may be truncated in order to account for variations in DBMS-specific name length limits.

The name length limit used has been increased to the maximum length supported by the DBMS. In most cases this is 128 characters, although some have shorter limits, including the following:

- ❑ **SAP HANA, Amazon Redshift.** 127 bytes.
- ❑ **MySQL, MariaDB.** 64 bytes.
- ❑ **PostgreSQL, Greenplum, Hyperstage.** 63 bytes.
- ❑ **Progress.** 32 bytes.
- ❑ **Oracle.** 30 bytes.
- ❑ **DB2 on z/OS.** 30 bytes.

SQL Adapters

This section provides descriptions of new features for SQL adapters.

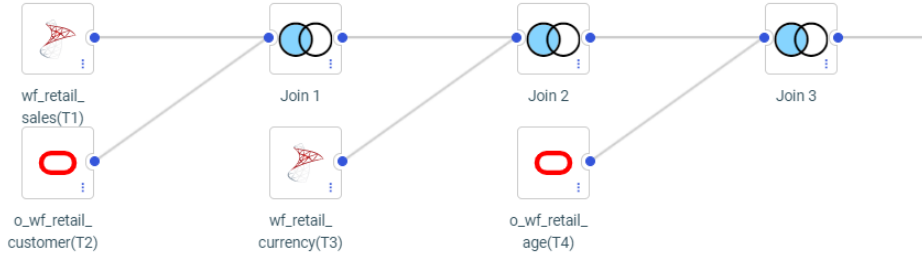
Cross-Join Performance Improvements

The following retrieval performance improvements have been made for joining multiple tables in different databases.

- ❑ If there is more than one cluster of tables from the same DBMS in the FROM clause, joins between those tables are passed to the corresponding DBMS.
- ❑ If there is a WHERE condition on fields from the outer table in a left or right outer join, and this condition always fails for null values, that join is converted to an inner join.
- ❑ If there is WHERE condition on fields from the left, right, or both tables in a full outer join, and this condition always fails for null values, that join is converted to a left, right, or inner join.
- ❑ Merge sub-select phrases *if from*, *where*, *group by*, and *having* are the same and re-use the temporary table created for the merge sub-select statement.

Example: Improving Performance With Left Outer Joins

The following flow creates three left outer joins between Microsoft SQL Server tables and Oracle tables. The Oracle synonyms start with the characters o_.



The following SQL statement corresponds to the joins generated by the flow.

```
SQL
SELECT
  T1.ID_SALES ,
  T1.ID_STORE ,
  T1.ID_CURRENCY ,
  T1.ID_CUSTOMER ,
  T1.ID_DISCOUNT ,
  T1.ID_PRODUCT ,
  T1.ID_TIME ,
  T1.COGS_LOCAL ,
  T1.COGS_US ,
  T1.DISCOUNT_LOCAL ,
  T1.DISCOUNT_US ,
  T1.GROSS_PROFIT_LOCAL ,
  T1.GROSS_PROFIT_US ,
  T1.MSRP_LOCAL ,
  T1.MSRP_US ,
  T1.QUANTITY_SOLD ,
  T1.REVENUE_LOCAL ,
  T1.REVENUE_US ,
  T2.ID_AGE ,
  T2.ID_EDUCATION ,
  T2.ID_GEOGRAPHY ,
  T2.ID_INCOME ,
  T2.ID_INDUSTRY ,
  T2.ID_MARITAL_STATUS ,
  T2.ID_OCCUPATION ,
  T2.ID_TIME_MIN ,
  T2.ID_TIME_MAX ,
  T2.EMAIL_ADDRESS ,
  T2.FIRSTNAME ,
  T2.FULLNAME ,
  T2.GENDER ,
  T2.LASTNAME ,
  T2.INCOME ,
  T3.CURRENCY_NAME ,
  T3.CURRENCY_RATE ,
  T4.AGE ,
  T4.AGE_RANGE ,
  T4.AGE_GROUP
```

```
FROM
  ((("ibisamp/facts".wf_retail_sales T1
    LEFT OUTER JOIN /*Join 1*/
    "ibisamp/dimensions".o_wf_retail_customer T2
    ON
    T1.ID_CUSTOMER = T2.ID_CUSTOMER )
    LEFT OUTER JOIN /*Join 2*/
    "ibisamp/dimensions".wf_retail_currency T3
    ON
    T1.ID_CURRENCY = T3.ID_CURRENCY )
    LEFT OUTER JOIN /*Join 3*/
    "ibisamp/dimensions".o_wf_retail_age T4
    ON
    T2.ID_AGE = T4.ID_AGE )
WHERE
  T1.REVENUE_US > 600 AND
  T2.LASTNAME LIKE 'C%' AND
  T3.CURRENCY_NAME = 'US Dollar' AND
  T4.AGE BETWEEN 35 AND 40
;
END
```

The Session Log shows the joins that were generated. The left outer joins were converted to inner joins, as shown in the following partial listing.

```

JOIN INNER
  SQLAPP01.ID_AGE
  IN 'ibisamp/dimensions/o_wf_retail_customer' TAG SQLAPP01
TO ALL
  ID_AGE
  IN 'ibisamp/dimensions/o_wf_retail_age' TAG SQLJTG01
AS SQLJNM01
END
DEFINE FILE 'ibisamp/dimensions/o_wf_retail_customer' TEMP
  SQLDEF01/I1 WITH SQLJTG01.O_WF_RETAIL_AGE.ID_AGE = 1;
END
TABLEF FILE 'ibisamp/dimensions/o_wf_retail_customer'
  PRINT
    SQLAPP01.ID_AGE
    SQLAPP01.ID_EDUCATION
    SQLAPP01.ID_GEOGRAPHY
    SQLAPP01.ID_INCOME
    SQLAPP01.ID_INDUSTRY
    SQLAPP01.ID_MARITAL_STATUS
    SQLAPP01.ID_OCCUPATION
    SQLAPP01.ID_TIME_MIN
    SQLAPP01.ID_TIME_MAX
    SQLAPP01.EMAIL_ADDRESS
    SQLAPP01.FIRSTNAME
    SQLAPP01.FULLNAME
    SQLAPP01.GENDER
    SQLAPP01.LASTNAME
    SQLAPP01.INCOME
    SQLJTG01.AGE
    SQLJTG01.AGE_RANGE
    SQLJTG01.AGE_GROUP
    SQLAPP01.ID_CUSTOMER
    SQLDEF01 AS (, 'SQL$$HIDDEN01',)
    SQLJTG01.ID_AGE NOPRINT

WHERE ( SQLAPP01.LASTNAME LIKE 'C%' ) ;
WHERE ( ( SQLJTG01.AGE FROM 35 TO 40
  AND SQLJTG01.AGE NE MISSING) ) ;
ON TABLE SET HOLDATTRS ON
ON TABLE SET CARTESIAN ON
ON TABLE SET ASNAMES MIXED
ON TABLE SET HOLDLIST PRINONLY
ON TABLE HOLD
  AS SQLHLD01
  FORMAT DATREC
END

```

In addition, joins that use the same DBMS are passed to that DBMS, as shown in the following partial listing.

```
SELECT
T1."ID_CUSTOMER" ,
T1."ID_AGE" ,
T1."ID_EDUCATION" ,
T1."ID_GEOGRAPHY" ,
T1."ID_INCOME" ,
T1."ID_INDUSTRY" ,
T1."ID_MARITAL_STATUS" ,
T1."ID_OCCUPATION" ,
T1."ID_TIME_MIN" ,
T1."ID_TIME_MAX" ,
T1."EMAIL_ADDRESS" ,
T1."FIRSTNAME" ,
T1."FULLNAME" ,
T1."GENDER" ,
T1."LASTNAME" ,
T1."INCOME" ,
T2."ID_AGE" ,
T2."AGE" ,
T2."AGE_RANGE" ,
T2."AGE_GROUP"
FROM
wf_retail_customer_t T1,
wf_retail_age_t T2
WHERE
(T2."ID_AGE" = T1."ID_AGE") AND
(T1."LASTNAME" LIKE 'C%') AND
(T2."AGE" IS NOT NULL) AND
(T2."AGE" BETWEEN 35 AND 40);
```

Cross Database Join Performance Improvements

Retrieval performance has been improved under certain conditions when you join tables from different Relational database systems.

The performance improvements result from extracting data from the cross-referenced table prior to performing the join or issuing a sub-select. You can disable this optimization process by issuing the following command:

```
SQL SET HOLDSQLJOIN = OFF
END
```

By default, this parameter is ON.

The following performance improvements have been implemented:

- For a non-aggregated query, the cross-referenced (TO) table is saved as a file in an internal binary format. This is faster than joining to a table in a different database system.

- ❑ For an aggregated query, for the cross-referenced table joined from any aggregation functions (min, max, sum, avg, count), the retrieval is passed to the relational database in a sub-select. This can result in retrieving a much smaller answer set, which improves performance.
- ❑ For a request with a clause that tests if two columns are equal or both are NULL, the TO table is held in an internal binary format, also improving performance.

You can view the generated query in either the Session Log or the trace file.

The following SQL request joins a Microsoft SQL Server named citibike_mssql table to a MySQL table named station_zip_mysql.

```
SQL
SELECT
    T1.TRIPDURATION,
    T1.START_STATION_NAME ,
    T1.END_STATION_NAME ,
    T1.BIKEID ,
    T1.BIRTH_YEAR ,
    T1.GENDER ,
    T1.STARTTIME ,
    T1.STOPTIME ,
    T1.USERTYPE ,    T2.ZIP_CODE ,
    T2.COUNTY ,
    T2.CITY
FROM
    citibike.citibike_mssql T1
    INNER JOIN /*Join 1*/
    citibike.station_zip_mysql T2
    ON
    T1.START_STATION_ID = T2.STATION_ID
;
TABLE
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

When this request is run, the Session Log shows the generated request. First the station_zip_mysql table is held in an internal format:

```
TABLEF FILE STATION_ZIP_MYSQL
PRINT
  ZIP_CODE
  COUNTY
  CITY
  STATION_ID
ON TABLE SET CARTESIAN ON
ON TABLE SET ASNAMES ON
ON TABLE SET HOLDLIST PRINONLY
ON TABLE HOLD
  AS SQLHLD01
  FORMAT DATREC
END
```

Next the citibike_mssql table is joined to the HOLD file.

```
JOIN INNER
  CITIBIKE_MSSQL2.START_STATION_ID
  IN CITIBIKE_MSSQL
TO ALL
  E04
  IN SQLHLD01
AS SQLJNM01
END
```

Finally, the joined structure is queried to produce the output.

```
TABLEF FILE CITIBIKE_MSSQL
PRINT
  CITIBIKE_MSSQL.TRIPDURATION
  CITIBIKE_MSSQL.START_STATION_NAME
  CITIBIKE_MSSQL.END_STATION_NAME
  CITIBIKE_MSSQL.BIKEID
  CITIBIKE_MSSQL.BIRTH_YEAR
  CITIBIKE_MSSQL.GENDER
  CITIBIKE_MSSQL.STARTTIME
  CITIBIKE_MSSQL.STOPTIME
  CITIBIKE_MSSQL.USERTYPE
  SQLHLD01.E01
  SQLHLD01.E02
  SQLHLD01.E03
  SQLHLD01.E04 NOPRINT
ON TABLE SET CARTESIAN ON
ON TABLE SET ASNAMES ON
ON TABLE SET HOLDLIST PRINONLY
END
```


Partial output is shown in the following image.

tripduration	start_station_name	end_station_name	bikeid	year	gender	starttime	stoptime	usertype	ZIP_CODE	COUNTY	CITY
897	W 45 St & 6 Ave	E 51 St & 1 Ave	18340	1966	1	2019/07/01 00:00:00.132000	2019/07/01 00:14:58.004000	Subscriber	10036	New York County	New York
267	5 Ave & E 78 St	W 82 St & Central Park West	21458	1996	1	2019/07/01 00:00:05.178000	2019/07/01 00:04:32.450000	Customer	.	.	.
2201	E 6 St & Avenue B	India St & West St	39874	1986	1	2019/07/01 00:00:05.213000	2019/07/01 00:36:46.749000	Subscriber	0	New York County	New York
1660	Harrison St & Hudson St	Washington Pl & 6 Ave	38865	1988	1	2019/07/01 00:00:08.601000	2019/07/01 00:27:48.805000	Subscriber	10013	New York County	New York
109	W 113 St & Broadway	W 110 St & Amsterdam Ave	30256	1997	1	2019/07/01 00:00:12.158000	2019/07/01 00:02:01.567000	Subscriber	10025	New York County	New York
106	31 St & Broadway	Crescent St & Broadway	16875	1988	1	2019/07/01 00:00:12.668000	2019/07/01 00:01:59.606000	Subscriber	11106	Queens County	Astoria New York
550	Park Ave & E 124 St	Frederick Douglass Blvd & W 112 St	34139	1992	1	2019/07/01 00:00:14.494000	2019/07/01 00:09:25.397000	Subscriber	10035	New York County	New York

Direct Bulk Load Support

Using DataMigrator to load data from a file to a database table using Extended Bulk Load always extracted data from the source file to create a file in the expected format of the bulk load program for the database to be loaded. In those cases where the file was already in the required format, this required additional processing time.

Now for certain data sources, when Extended Bulk Load is used and Optimize Load is enabled, data is loaded directly to the target database without the additional extract and transformation steps, provided the following conditions are met:

- The data source is a delimited flat file.
- The file attributes including Header, Enclosure, Delimiter and Record Delimiter are supported by the target database.
- Only insert operations are used (no updates).
- No aggregation is specified.
- No transformations are used.

Under these conditions, direct bulk load is used, reducing processing time, and the following message appears in the DataMigrator log.

(ICM18637) Load Operation will use Direct Bulk Load Feature

The supported databases and data stores are:

- Apache Hive
- Apache Impala
- Apache Spark
- Amazon RedShift

- ExaSQL
- Jethro
- MariaDB
- MS SQL Server
- MongoDB
- MySQL
- PostgreSQL
- Snowflake
- Salesforce.com
- Sybase ASE
- Vertica

Navigating Joins Between Cluster Synonyms

By default, when joining cluster synonyms, a hierarchy of segments is constructed from all of the joined files, and the resulting hierarchy is navigated in top-to-bottom, left-to-right order.

Therefore, if a left outer join is specified from a host synonym to a cluster that has an inner join, the inner join will be performed last and may remove rows from the host file, counteracting the purpose of the left outer join. Using the `SET FOCTRANSFORM = NESTED_CLUSTERS/ON` command, you can force the joins in the target cluster to be performed prior to the join between the host and target synonyms. When you use this setting, SQL scripts are used to join the tables in the target cluster prior to implementing the join to the host file. The left outer join will be performed last and will retain all rows in the host synonym.

The syntax is:

```
SET FOCTRANSFORM = {NESTED\_CLUSTERS/OFF | NESTED\_CLUSTERS/ON}
```

where:

[NESTED_CLUSTERS/OFF](#)

Maintains the left-to-right, top-to-bottom order of segment navigation. This is the default value.

NESTED_CLUSTERS/ON

Performs the joins in the target cluster synonym prior to joining the host synonym to the result.

Reference: Usage Notes for Joins to Cluster Synonyms

- ❑ This feature requires that the joins be optimized. The command SET SHORTPATH = SQL must be in effect for combinations of inner and outer joins with the setting FOCTRANSFORM = NESTED_CLUSTERS/OFF, in order for the request to be optimized. The SHORTPATH = SQL setting has no effect on optimization with the setting FOCTRANSFORM = NESTED_CLUSTERS/ON.
- ❑ You cannot join to a non-root segment of a cluster synonym. If you issue a join to a non-root segment, the following message displays and the request terminates:

```
(FOC906) JOIN TO NON-ROOT SEGMENT segname IS NOT ALLOWED FOR
NESTED_CLUSTERS
```

Example: Navigating Joins Between Cluster Synonyms

This example uses SQL Server data sources generated from a file of citibike trips uploaded from <https://www.citibikenyc.com/system-data>, and from a file of zip codes for the stations used for the trips (you can download this file from https://techsupport.informationbuilders.com/public/station_zip.csv).

A cluster synonym named station_trip_cls joins the station zip data source to a data source containing partial trip data (with only a few rows). The following shows the inner join defined in the cluster synonym:

```
FILENAME=STATION_TRIP_CLS, $
SEGMENT=STATION_ZIP_OLEDB, CRFILE=CITIBIKE/STATION_ZIP_OLEDB,
CRINCLUDE=ALL, $
SEGMENT=CITIBIKE_PARTIAL_OLEDB, SEGTYPE=KU, PARENT=STATION_ZIP_OLEDB,
CRFILE=CITIBIKE/CITIBIKE_PARTIAL_OLEDB, CRINCLUDE=ALL, CRJOINTYPE=INNER,
JOIN_WHERE=STATION_ID EQ START_STATION_ID; $
```

The following request issues a left outer join from a larger version of the trip data file to the cluster:

```

SET FOCTRANSFORM = NESTED_CLUSTERS/&VALUE
SET SHORTPATH = SQL
JOIN LEFT OUTER START_STATION_ID IN CITIBIKE_TRIPDATA TAG T1 TO ALL
STATION_ID IN STATION_TRIP_CLS TAG T2 AS J1
TABLE FILE CITIBIKE_TRIPDATA
" NESTED_CLUSTERS/&VALUE"
" "
SUM CNT.T1.START_STATION_ID AS T1,Station CNT.ZIP_CODE
CNT.T2.START_STATION_ID AS T2,Station
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END

```

Running the request with &VALUE set to OFF generates the following trace:

```

SELECT
  COUNT(T1."START_STATION_ID"),
  COUNT(T2."ZIP_CODE"),
  COUNT(T3."START_STATION_ID")
FROM
  ( ( citibike_tripdata_mssqloledb T1
    LEFT OUTER JOIN
    station_zip_oledb T2
    ON T2."STATION_ID" = T1."START_STATION_ID" )
  INNER JOIN
  citibike_partial_msoledb T3
  ON (T3."START_STATION_ID" = T2."STATION_ID" ) );

```

The output is shown in the following image. The inner join was done last, reducing the number of stations in the host file to the same number as in the cluster.

NESTED_CLUSTERS/OFF		
<u>T1</u>	<u>ZIP_CODE</u>	<u>T2</u>
<u>Station</u>	<u>COUNT</u>	<u>Station</u>
165	158	165

Running the request with &VALUE set to ON generates the following trace. Two SQL scripts are generated, one for the host file and one for the join in the cluster. Then, the left outer join is performed against the result of the inner join:

```

SELECT
  T1."START_STATION_ID" AS "SK001_START_STATION_ID",
  COUNT(T1."START_STATION_ID") AS "VB001_CNT_START_STATION_ID"
  FROM
  citibike_tripdata_mssqloledb T1
  GROUP BY
  T1."START_STATION_ID";
(FOC2546) SQL SCRIPT
__CITIBIKE_TRIPDATA_OLEDB_CITIBIKE_TRIPDATA_OLEDB.SQL CREATED SUCCESSFULLY
(BUT NOT EXECUTED)
_EDATEMP/___citibike_tripdata_oledb_citibike_tripdata_oledb HELD AS
SQL_SCRIPT
SELECT
  T1."STATION_ID" AS "SK001_STATION_ID",
  COUNT(T1."ZIP_CODE") AS "VB001_CNT_ZIP_CODE",
  COUNT(T2."START_STATION_ID") AS "VB002_CNT_START_STATION_ID"
  FROM
  station_zip_oledb T1,
  citibike_partial_msoledb T2
  WHERE
  (T2."START_STATION_ID" = T1."STATION_ID")
  GROUP BY
  T1."STATION_ID";
(FOC2546) SQL SCRIPT
__CITIBIKE_TRIPDATA_OLEDB_STATION_PARTIAL_OLEDB_CLS.SQL CREATED
SUCCESSFULLY (BUT NOT EXECUTED)
_EDATEMP/___citibike_tripdata_oledb_station_partial_oledb_cls HELD AS
SQL_SCRIPT

```

```

SELECT
SUM(T1."VB001_CNT_START_STATION_ID"),
SUM(T2."VB001_CNT_ZIP_CODE"),
SUM(T2."VB002_CNT_START_STATION_ID")
FROM
(
( /* vvv */
SELECT
T1."START_STATION_ID" AS "SK001_START_STATION_ID",
COUNT(T1."START_STATION_ID") AS
"VB001_CNT_START_STATION_ID"
FROM
citibike_tripdata_mssqloledb T1
GROUP BY
T1."START_STATION_ID"
) /* ^^ */ T1
LEFT OUTER JOIN
( /* vvv */
SELECT
T1."STATION_ID" AS "SK001_STATION_ID",
COUNT(T1."ZIP_CODE") AS "VB001_CNT_ZIP_CODE",
COUNT(T2."START_STATION_ID") AS
"VB002_CNT_START_STATION_ID"
FROM
station_zip_oledb T1,
citibike_partial_msoledb T2
WHERE
(T2."START_STATION_ID" = T1."STATION_ID")
GROUP BY
T1."STATION_ID"
) /* ^^ */ T2
ON T2."SK001_STATION_ID" = T1."SK001_START_STATION_ID" );

```

The output is shown in the following image. The left outer join was done last, maintaining the original number of stations in the host file.

NESTED_CLUSTERS/ON		
<u>T1</u>	<u>ZIP_CODE</u>	<u>T2</u>
<u>Station</u>	<u>COUNT</u>	<u>Station</u>
6680	8	9

SQL Adapters: Optimizing OUTPUTLIMIT

An OUTPUTLIMIT filter in a TABLE request is now, by default, passed to some SQL data sources as FETCH FIRST n ROWS. When needed, it can be suppressed using the optimization setting FEATOPT OUTPUTLIMIT OFF.

For example, the following request against a Db2 data source contains an OUTPUTLIMIT filter:

```
SET TRACEUSER = ON
SET TRACEOFF = ALL
SET TRACEON = STMTRACE//CLIENT
TABLE FILE WF_RETAIL_LITE
SUM COGS_US
BY PRODUCT_CATEGORY
IF OUTPUTLIMIT IS 50
END
SET TRACEUSER=OFF
```

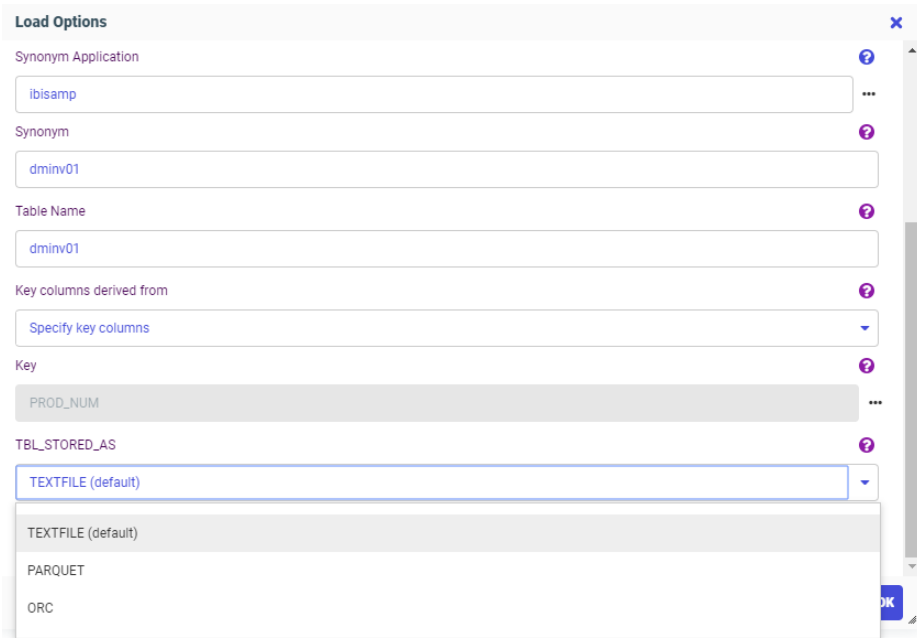
The generated SQL request contains the FOR FETCH FIRST 50 ROWS clause:

```
SELECT
T1."ID_PRODUCT",
T1."COGS_US",
T7."ID_PRODUCT",
T7."PRODUCT_CATEGORY"
FROM
( wrd_wf_retail_sales T1
LEFT OUTER JOIN
wrd_wf_retail_product T7
ON T7."ID_PRODUCT" = T1."ID_PRODUCT" )
ORDER BY
T7."PRODUCT_CATEGORY"
FETCH FIRST 50 ROWS ONLY
FOR FETCH ONLY;
```

Adapter for Amazon Athena: Creating Files in Parquet or ORC Format

When using DataMigrator or a data flow to load data to S3 using the Adapter for Athena, files can now be created in the ORC or Parquet column store formats, providing better performance when reading from them.

In Target Properties or Load Options, select from the TBL_STORED_AS drop-down list, as shown in the following image.



Adapter for Apache Drill: Using Parquet as a Transfer File to HDFS

In prior releases, when you used DataMigrator or a Data Flow to create a file in the Hadoop Distributed File System (HDFS) in Parquet format, an intermediate text file was created.

Now, a file in Parquet format can be created in the local file system and transferred to HDFS without the need to create an intermediate file.

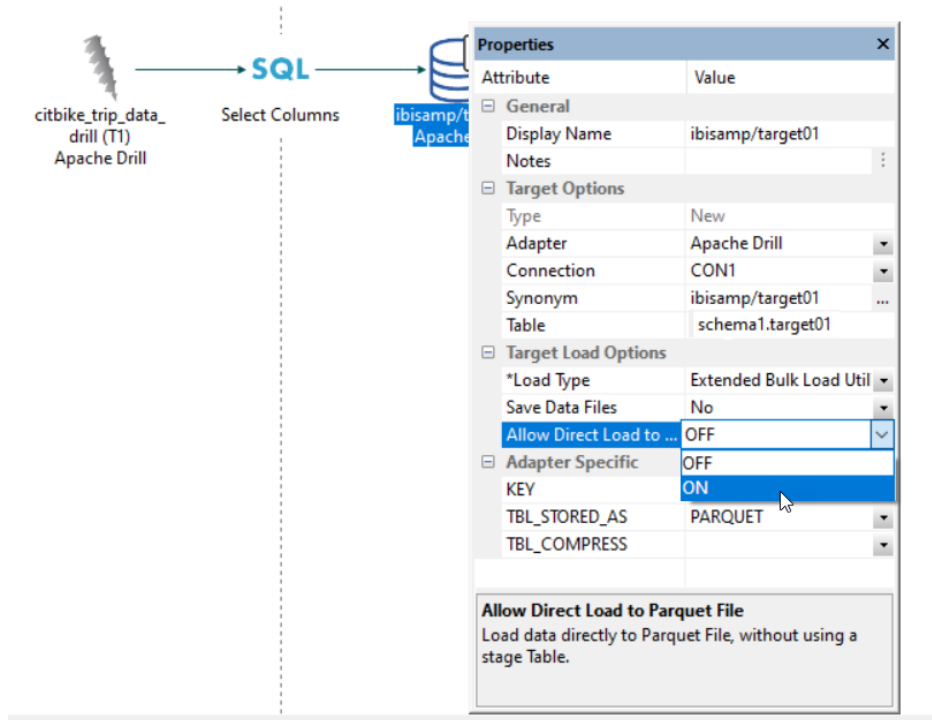
This option can be enabled in the Load Options for a flow by setting *Allow Direct Load to Parquet File* to ON.

The following image shows this option in the Web Console.

The image shows a 'Load Options' configuration panel with the following settings:

- New/Replace:** New/Replace
- Adapter:** Apache Drill
- Synonym Application:** ibisamp
- Synonym:** citbike_trip_data01
- Table Name:** schema1.citbike_trip_data01
- Bulk Load:** Bulk Load
- Allow Direct Load to Parquet File:** OFF - Default (highlighted with a blue circle, dropdown menu is open showing 'PARQUET')
- Save Data Files:** NO - Default
- Key columns derived from:** Key columns in source table / sorted columns in query
- TBL_STORED_AS:** default (dropdown menu is open showing 'PARQUET')

The following image shows this option in the Data Management Console (DMC).



Adapter for Google BigQuery: Support for SQL Immediate Commands

The SQL Passthru commands INSERT, UPDATE, and DELETE are now supported for the Adapter for Google BigQuery.

For example:

```
SQL SQLGBQ
DELETE FROM 'focus-100020.informationbuilders_tables.car_file' WHERE
country="ENGLAND"
END
```

Adapter for Google BigQuery: Support for Service Account Authentication

In prior releases, only OAuth20 authentication was supported for the Adapter for Google BigQuery. Service Account Authentication support has now been added. A service account is a non-human user that needs to authenticate to Google and be authorized to access data in Google APIs, for example, for running reports distributed by ReportCaster.

To configure Service Account authentication, select *Service Account* from the *Security* drop-down list, as shown in the following image.

The screenshot shows a configuration window titled "Add Google BigQuery to Configuration" with a "Prerequisites" link. Under the "Connect parameters" section, there are several input fields: "Connection Name" (containing "CON01"), "Google BigQuery URL" (containing "https://www.googleapis.com/bigquery/v2/projects/"), and "Project ID" (empty). The "Security" dropdown menu is open, showing "OAuth" as the selected option and "Service Account" as an available option below it. There are also help icons (question marks) next to each field.

The following fields appear for configuring Service Account authentication.

Key File Location

Is the location of the JSON Key File that was created in the Google Development environment. For example:

`C:\ibi\apps\googlebigquery\webfocus-555555-874feowfwojoe7.json`

Client Email Address

Is the client email address defined in the JSON Key File. For example:

`xxxxxxxxxxxx-compute@developer.gserviceaccount.com`

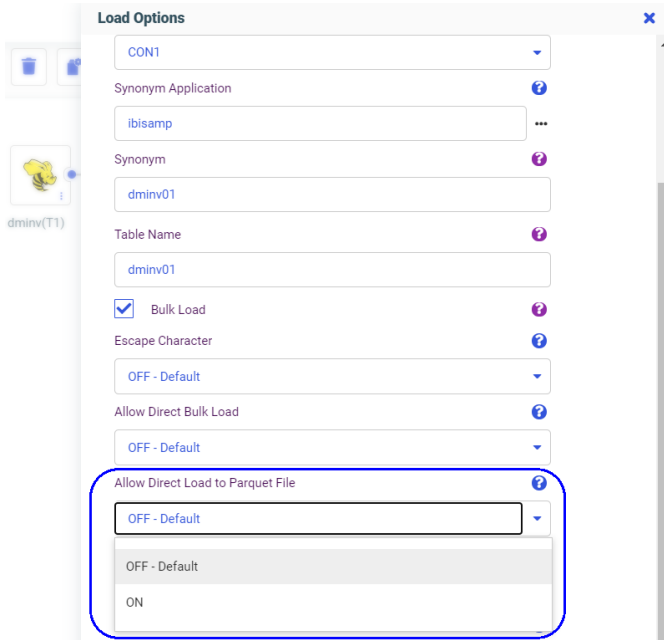
Adapters for Hive and Impala: Using Parquet as a Transfer File to HDFS

In prior releases, when you used DataMigrator or a Data Flow to create a file in the Hadoop Distributed File System (HDFS) in Parquet format, an intermediate text file was created.

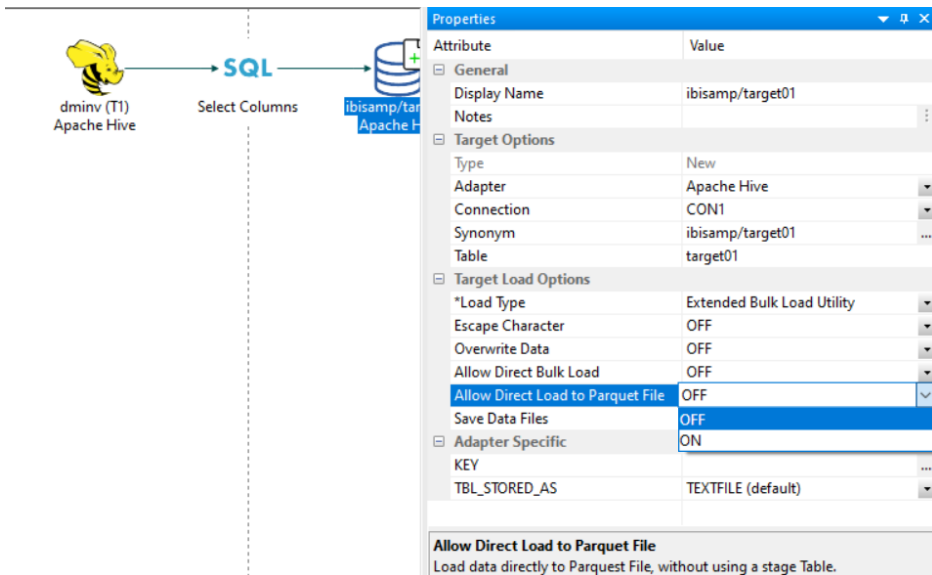
Now, a file in Parquet format can be created in the local file system and transferred to HDFS without the need to create an intermediate file.

This option can be enabled in the Load Options for a flow by setting *Allow Direct Load to Parquet File* to ON.

The following image shows this option in the Web Console.



The following image shows this option in the Data Management Console (DMC).



Adapter for Hive: Optimize Load With Bulk Load Options

In prior releases, when you used DataMigrator or Data Preparation with *Optimize Load* enabled, the ability to specify Load Options was not available.

Now, when loading data to Hive, options such as "Escape Character" are available when Optimize Load is enabled, as shown in the following image.

File Panel X Quick Copy for ibisamp/brokers X

On this page specify the options for where and how you want the target table created. See the onscreen help for the individual parameters

Select Target Parameters

Load Option ?
New/Replace

Adapter ?
Apache Hive

Connection ?
CON01

Synonym Application ?
ibisamp

Synonym ?
brokers

Table Name ?
brokers

Bulk Load ?

Escape Character ?
OFF - Default

Overwrite data ?
OFF - Default

Run Save

Adapter for Hive: Configuring Bulk Load for HDFS Transfer Protocol

Using DataMigrator or Data Preparation with the adapters for Apache Hive or Impala, you can configure Bulk Load to use transfer protocol HDFS to write directly to a Hadoop cluster, when running on an Edge node of the cluster. This eliminates the need to use [S]FTP to transfer the data first.

Adapter for Hive: Bulk Loading Data With Non-Printable Characters

When using Extended Bulk Load on a flow to load to Hive, a setting `SQLHIVE SET ESCAPE ON` can be used to correctly process non-printable ASCII characters such as Linefeed or Carriage Return characters. With this setting an intermediate JSON file is created to load into Hive.

You can set the `BLK_ESCAPE` parameter on the Change Settings page for Hive, as shown in the following image.

Change Settings for Apache Hive

Save settings in
 Profile

Select
 edasprof

Bulk Load Service

BULKLOAD ⓘ
 ON - Default

Use Extended Bulk Load

BULKCHECK ⓘ
 1000000

Commit every row(s). Default: 1000000

BLK_ESCAPE ⓘ
 OFF - Default

OFF - Default
 ON

BLK_PERM_AS_STAGE ⓘ
 ON - Default

Use Permanent Table as Stage

Customize data type mappings ▾

Diagnostics ▾

Metadata ▾

Miscellaneous settings ▾

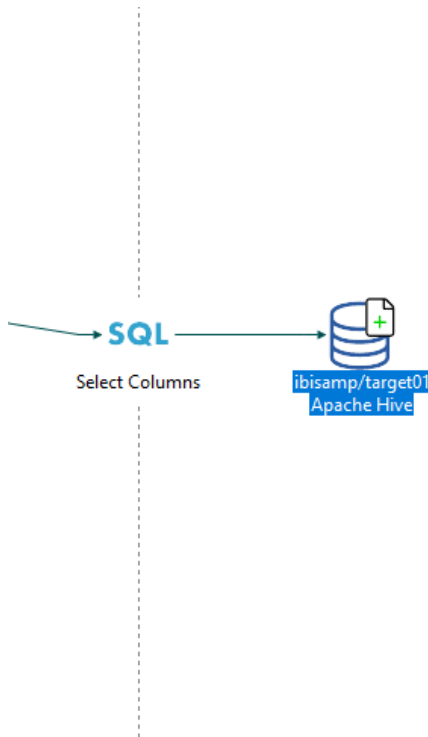
Read/Write Optimization ▾

Reset to defaults Cancel Save

By default, the value is Off.

Select *On* to enable bulk load of non-printable characters.

You can also set this parameter on the Properties panel for a Hive target in a data flow, as shown in the following image.



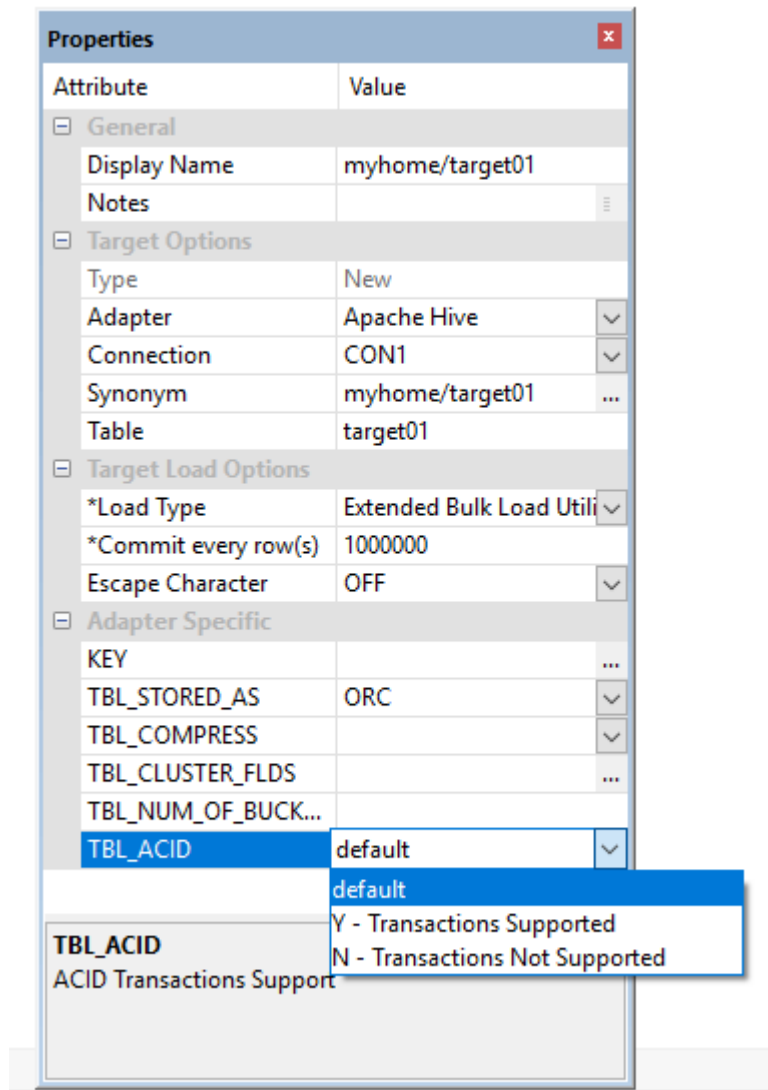
The diagram illustrates a data flow from a source labeled 'SQL' to a target labeled 'ibisamp/target01 Apache Hive'. A dashed vertical line separates the source from the target. An arrow points from 'SQL' to the target, with the text 'Select Columns' positioned below the arrow. The target is represented by a database icon with a plus sign and a label 'ibisamp/target01 Apache Hive'.

Attribute	Value
General	
Display Name	ibisamp/target01
Notes	
Target Options	
Type	New
Adapter	Apache Hive
Connection	CON1
Synonym	ibisamp/target01
Table	target01
Target Load Options	
*Load Type	Extended Bulk Load
*Commit every row(s)	1000000
Escape Character	OFF
Adapter Specific	
KEY	
TBL_STORED_AS	TEXTFILE (default)

Escape Character
Escape Character
Specifies the single character that is used for escape sequences such as ,, or @; and for escaping data characters that might otherwise be interpreted as row or column delimiters

Adapter for Hive: Support for New Transactions With ORC-Formatted Records

When you configure the properties for a Hive target in a data flow, you can specify that new transactions will be allowed for ORC-formatted records, as shown in the following image.



The values for TBL_ACID are:

- default.** Respects the value set within Hive.

- Y.** Transactions are supported.
- N.** Transactions are not supported. To modify the target, it must be recreated.

Adapter for Hyperstage (PG): Installation Setting for Heap Size

By default, Hyperstage reserves 75% of memory for use in its functionality. If other processes are running during the Hyperstage installation, Hyperstage may not be able to reserve that much memory, and the installation process may not complete properly.

The amount of memory that Hyperstage reserves can be adjusted during installation using the `ServerMainHeapSize` setting on the *Set Initial Settings* page, as shown in the following image.

The screenshot shows the 'Select Initial Settings' dialog box. The 'Hyperstage ServerMainHeapSize Configuration (optional)' section is highlighted with a red border. This section contains the following text:

Please enter size in MegaBytes that will be reserved for the Hyperstage engine.

The `ServerMainHeapSize` should normally be set to 60-80% of Total Physical Memory. Memory used by other Reporting Server Processes, including the Bulk Data Loader, should be deducted in the calculation of Available Memory.

Total Physical Memory: 8078
 Current Available Memory: 1831
 Recommended value for `ServerMainHeapSize`: 605 MB

The `ServerMainHeapSize` will be stored in the Hyperstage configuration file (`infobright.cnf`).

Hyperstage `ServerMainHeapSize`:

At the bottom of the dialog, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a blue border.

The setting will be added in the `infobright.cnf` file within the `HyperstagePG\ib_data` directory.

Adapters for Impala and Crossdata: SET DEFAULTSCHEMA

The adapters for Impala and Crossdata have a new SET DEFAULTSCHEMA setting that can be issued in the server profile. The setting allows specifying the schema, but delays doing so until a connection has been established, in order to improve connection stability and eliminate initialization errors at startup.

The syntax for Impala is:

```
SQL {SQLIMP} SET DEFAULTSCHEMA schemaname
```

The syntax for Crossdata is:

```
SQL {SQLCRD} SET DEFAULTSCHEMA schemaname
```

where:

SQLIMP

Identifies the Adapter for Impala. You can omit this if you issued the SET SQLENGINE command for the adapter.

SQLCRD

Identifies the Adapter for Crossdata. You can omit this if you issued the SET SQLENGINE command for the adapter.

schemaname

Is the schema to use once the connection is established.

Adapter for MySQL: Support for Running Under an SSH Tunnel

The Adapter for MySQL can be configured to connect to a Secure Shell (SSH) Tunnel. SSH tunneling transports data over an encrypted SSH connection. Local forwarding is used to forward a port from the client machine to the server machine. The SSH client listens for connections on a configured port, and when it receives a connection, it tunnels the connection to an SSH server. The server connects to a configured destination port, which can be on a different machine than the SSH server.

The following steps configure tunneling:

1. The System Administrator opens a tunnel on the port to MySQL running on a given box.
2. The server user on any box runs port forwarding on session1.
3. A user on session2 (on the same box as in Step 2) starts the server and configures the Adapter for MySQL with a connection to SSH.

Adapter for Oracle: Support for Instant Client Basic Package

The Adapter for Oracle can now be configured to work with the Oracle Instant Client Basic package, in addition to the Database Client supported in prior releases.

Adapter for Oracle: Support for Version 19c

Read/Write support has been added for Oracle Version 19c.

Adapter for Oracle: Support for Autonomous Data Warehouse Cloud

The Adapter for Oracle supports Oracle Autonomous Data Warehouse (ADW) in the cloud, running versions 19c, 18c, or 12c.

Adapter for Oracle: Support for HINT on Segment Level

Hints in Oracle enable the user to affect the execution plan for an SQL statement, using knowledge of the data. The ENGINE SQLORA SET HINT command enables the user to specify hints at the SQL statement level. This feature enables the user to add hints for each segment in the Access File using the following syntax.

```
SQLHINT = 'hint_clause'
```

where:

```
'hint_clause'
```

Is any valid Oracle HINT (see the vendor list of hints for the appropriate release), enclosed in single quotation marks, for example:

```
SEGNAME = EMP_C, SQLHINT = 'INDEX(EMP PK_EMP)', $
SEGNAME = DEPT_C, SQLHINT = 'ALL_ROWS', $
```

The Adapter for Oracle adds the appropriate hint syntax (/ *+ hint_1 hint_2 .. hint_n */) to the SQL statement generated for the TABLE request. Each hint_n will be properly separated by a space and the combination of hints will be enclosed in valid comment tags according to Oracle rules.

The user is responsible for the proper hint clause syntax. Oracle ignores invalid HINT syntax without reporting an error. Therefore, the only way to check whether the Optimizer accepted the HINT is to run the Oracle explain command.

Adapter for Stratio Crossdata: Support for TIMESTAMP Data Type

The Adapter for Stratio Crossdata now supports the TIMESTAMP data type with microseconds, if that data type is supported by the version of Crossdata being run.

Adapter for TIBCO SnappyData

The Adapter for TIBCO SnappyData is new in this release and can be found in the SQL group of adapters.

SnappyData (TIBCO ComputeDB) is a distributed, in-memory optimized, analytics database that delivers high throughput and concurrency.

Adapter for SQL Server: Version 2019 Support

The Adapter for Microsoft SQL Server now supports SQL Server Version 2019.

ERP Adapters

The following features were added for ERP adapters.

Adapter for Microsoft Dynamics CRM: New Mappings for Long String/Memo and Image Data Types

In prior releases of the Adapter for Microsoft Dynamics CRM, a column defined as StringType with a MaxLength greater than 1024, or a column defined as MemoType, had metadata generated with format attributes USAGE=TX80L and ACTUAL=TX. In addition, columns defined as ImageType were not included in the generated metadata.

Now, StringType columns with MaxLength greater than 1024, and MemoType columns, are created with USAGE=STRING and ACTUAL=STRING. STRING is a character data type with no length specification. Also, ImageType columns are created with USAGE=BLOB and ACTUAL=BLOB.

Images stored as BLOB columns can be added to report output using the following StyleSheet declaration.

```
TYPE=DATA, COLUMN=columnname, IMAGE=(columnname), SIZE=(w h), $
```

where:

columnname

Is the name of the column that contains the image.

w

Is the width of the image specified in the units identified by the UNITS parameter.

h

Is the height of the image specified in the units identified by the UNITS parameter.

Example: Displaying an Image Column on a Report

The ACCOUNT Master File has an ImageType column named ENTITYIMAGE generated with USAGE=BLOB and ACTUAL=BLOB.

```
FIELDNAME=ENTITYIMAGE, ALIAS=entityimage, USAGE=BLOB, ACTUAL=BLOB,
MISSING=ON,
REFERENCE=ACCOUNT, $
```

The following request prints the image stored in the ENTITYIMAGE field.

```
TABLE FILE ACCOUNT
PRINT
  NAME
  ENTITYIMAGE AS 'PICTURE'
WHERE NAME EQ 'Josh James'
ON TABLE SET PAGE NOPAGE
_*****
-* Lines between asterisk lines required for BLOB image support
-* for HTML and DHTML formats.
ON TABLE SET HTML EMBEDDING AUTO
-* Required to support IE8 with images larger than 32K
ON TABLE SET HTML ARCHIVE ON
-* Required for image positioning in subheads in HTML reports
ON TABLE SET HTML CSS ON
_*****
ON TABLE PCHOLD FORMAT HTML
ON TABLE SET STYLE *
TYPE=REPORT, COLOR=BLUE, FONT=ARIAL, GRID=OFF, $
TYPE=HEADING, SIZE = 18, COLOR=RED, $
TYPE=DATA, COLUMN=ENTITYIMAGE, IMAGE= (ENTITYIMAGE) , SIZE= (1 1) , $
ENDSTYLE
END
```

The output is shown in the following image.



Adapter for Microsoft Dynamics CRM: Passing Client Credentials in the Body

You now have the option to pass the OAuth Grant Password Client ID and Client Secret in the Body of the Token request for the Adapter for Microsoft Dynamics CRM.

When configuring a connection, select the *Client Credentials in Body* check box, as shown in the following image.

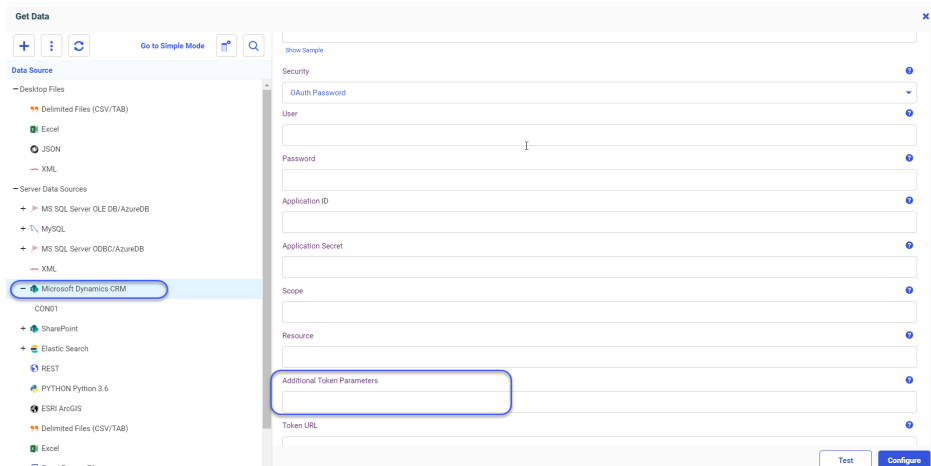
The image shows a configuration window titled "Change Connect Parameters for Microsoft Dynamics CRM". The window contains several fields and a checkbox. A red arrow points to the "Client Credentials in Body" checkbox, which is checked. The fields include:

- Connection Name: CON01
- Dynamics CRM URL: https://xxxxxxxxxxxxx.crm.dynamics.com/api/data/v9.1
- Security: OAuth Password
- Client Credentials in Body: (highlighted with a red arrow)
- User: [Empty field]
- Password: [Empty field]
- Application ID: [Empty field]
- Application Secret: [Empty field]

At the bottom right, there are three buttons: "Test", "Cancel", and "Configure".

Adapter for Microsoft Dynamics CRM: Support for Token Parameters With OAuth Grant Password Authentication

When you configure a connection for the Adapter for Microsoft Dynamics CRM and select OAuth Grant Password authentication, you can add additional token parameters, as shown in the following image.



You can use this field to enter token parameters that may be required to connect to a specific Dynamics CRM service, when the parameter is not part of the Scope.

Adapter for Microsoft Dynamics CRM: Support for OAuth Grant Password Authentication

In prior releases of the Adapter for Microsoft Dynamics CRM, the only type of OAuth authentication that was supported was OAuth Authorization Code authentication. Now OAuth Grant Password authentication is also supported, as shown in the following image.

The screenshot shows a configuration window titled "Add Microsoft Dynamics CRM to Configuration". It has several sections:

- Connect parameters**: A sub-header with an upward arrow.
- Connection Name**: A text input field containing "CON01".
- Dynamics CRM URL**: A text input field containing a sample URL: "https://xxxxxxxxxxxx.crm.dynamics.com/api/data/v9.1". Below it is a "Hide Sample" link.
- Security**: A dropdown menu currently showing "OAuth Password". A red circle highlights this option in the expanded list. Other options in the list are "NTLM", "OAuth Authorization Code", and "Trusted".
- Application Secret**: An empty text input field.
- Scope**: An empty text input field.

 At the bottom right, there are two buttons: "Test" and "Configure".

Configure the following parameters for OAuth Grant Password authentication.

Connection Name

Is the logical name used to identify this particular set of connection attributes.

The default value is CON01.

Dynamics CRM URL

Is the URL of the Dynamics CRM API request.

For example:

<https://mycrmdomain/api/data/v8.1>

Security

Select *OAuth Password* from the *Security* drop-down list.

User

Is the user ID used to authenticate to Dynamics CRM, when using OAuth Password authentication.

Password

Is the password used to authenticate to Dynamics CRM, when using OAuth Password authentication.

Application Secret

Is the Application Secret for the application created in the Azure Active Directory, when using OAuth Password authentication.

Obtain the Application Secret using the following steps:

1. Go to:

<https://portal.azure.com>

2. Click *Azure Active Directory* in the left panel.
3. Click *App registrations*.
4. Click the registered application.
5. Click *Certificates & secrets*.
6. Click *+ New client secret* to create a Client Secret.

The Application Secret is the value for the created Client Secret.

Scope

Is the scope sent as part of the OAuth Token request, when using OAuth Password authentication (optional). For example, *openid*.

Resource

Is the resource sent as part of the OAuth Token request, when using OAuth Password authentication (optional). For example:

<https://mycrmdomain>

where:

mycrmdomain

Is the domain when accessing Microsoft Dynamics, for example, <https://mycompany.crm.dynamics.com>.

Application ID

Is the value that identifies your application to Dynamics CRM when using OAuth Authentication.

Obtain this value using the following steps:

1. Go to:

<https://portal.azure.com>

2. Click *Azure Active Directory* in the left panel.
3. Click *App registrations*.
4. Click the Application for the Dynamics CRM Adapter that was previously created.
5. Use the value from the Application (client) ID.

Token URL

Is the URL used to obtain a Token to access Dynamics CRM.

Obtain this URL using the following steps:

1. Go to:
<https://portal.azure.com>
2. Click *Azure Active Directory* in the left panel.
3. Click *App registrations*.
4. Click *Endpoints*.
5. Use the URL from the OAuth 2.0 token endpoint (v1).

Adapter for MS Dynamics CRM: Support for Printing LOOKUP Fields

In prior releases, lookup fields were returned as missing values in requests using the Adapter for Microsoft Dynamics CRM. Now, the values for these fields are returned. It is recommended that you recreate the metadata prior to issuing a request with a lookup field.

Adapter for Salesforce.com: Creating Custom Objects

In prior releases, you could load data into a Salesforce.com custom object in a data flow as an *Existing Target*, but the custom object first had to be created in Salesforce.com. Now, you can create a Salesforce.com custom object as a *New Target* in a data flow and load it

Adapter for Salesforce.com: Saving Intermediate and Log Files

When loading data to Salesforce.com using the Bulk API, you can retain an intermediate file that is created, as well as the log (response) file. You can then use these files in conjunction with the batch results to see which records were loaded successfully (with the internal id for those records), and which were not (with the error message for those records).

The files are saved in the same application directory as the flow. Each is generated with a synonym that you can use for reporting. The file type of the intermediate HTML file is .ftm and the response file is .log.

This option can be set for the adapter or for an existing target in a data flow.

To implement this option for the adapter, right-click the adapter on the Get Data page, and click *Change Settings*. The *Change Settings for Salesforce.com* page opens, as shown in the following image.

Bulk Services

BULKLOAD ?
 OFF - Default
 Use Extended Bulk Load

BULKCHECK ?
 10000
 Commit every row(s). Default: 10000

BULKQUERY ?
 OFF - Default
 Enable Bulk Query for Data Retrieval

PKCHUNKING ?
 OFF - Default
 Enable Bulk Query PK Chunking for Data Retrieval

PKCHUNKSIZE ?
 0
 Bulk Query PK Chunk Size for Data Retrieval. Default: 0

DIRECT_BULK_LOAD ?
 OFF - Default
 Allow Direct Bulk Load

BLK_SAVE_DATA_FILES ?
 NO - Default
 Save REST API Source Data Files

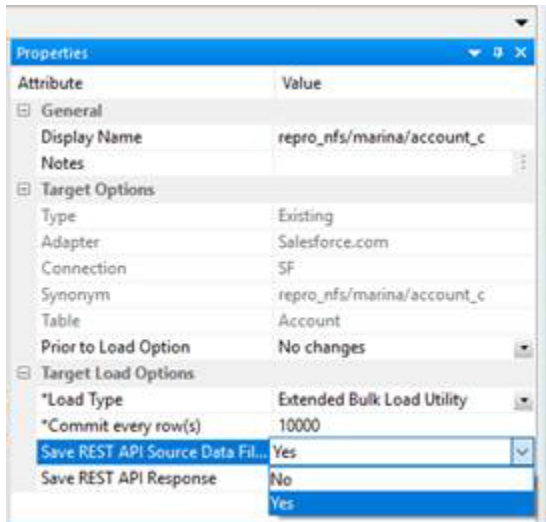
BLK_SAVE_RESPONSE_LOG ?
 NO - Default
 Save REST API Response

Reset to defaults Cancel Save

The parameter `BLK_SAVE_DATA_FILES` is set to No by default. Change it to Yes to save the intermediate HTML file. The log file will also be saved.

The parameter `BLK_SAVE_RESPONSE_LOG` is set to No by default. Change it to Yes to save the log file when you are not saving the intermediate HTML file.

The following image shows these properties for an existing Salesforce.com target.



You can also issue the following command in a stored procedure to save the data and log files:

```
ENGINE SFDC SET BLK_SAVE_DATA_FILES YES
```

You can issue the following command in a stored procedure to save only the log file:

```
ENGINE SFDC SET BLK_SAVE_RESPONSE_LOG YES
```

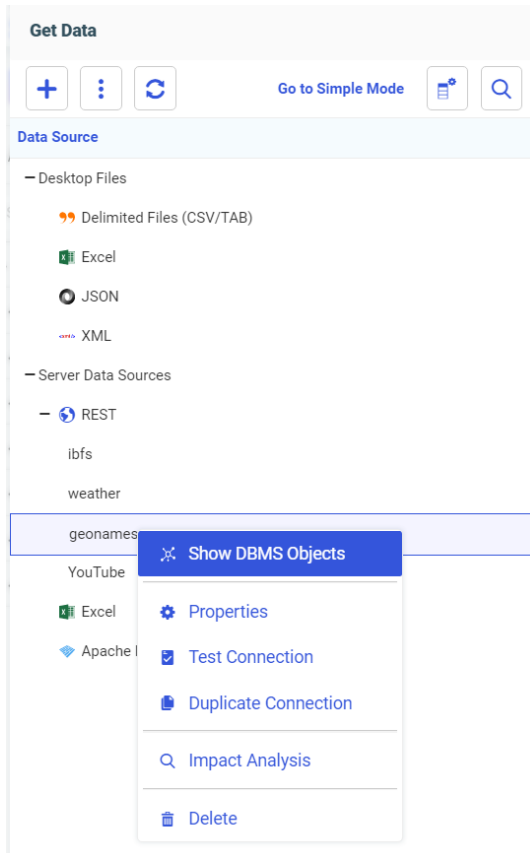
Procedures Adapters

This section provides descriptions of new features for Procedures adapters.

Adapter for REST: Populating XDEFAULT Attributes in a JSON Synonym for a POST Request

XDEFAULT attributes are now populated in the metadata for the Adapter for REST when JSON is used in the Document Sample for a POST Create Synonym request.

The following example shows how to start the synonym creation process for the geonames connection.



The synonym is created for a POST request, as shown in the following image.

The image shows a web interface for creating a REST synonym. The title is "Create Synonym for REST (geonames)". Below the title is a section "Create Synonym options" with a blue arrow icon. The form contains the following fields and controls:

- Select REST Operation:** A dropdown menu with "Post" selected.
- Service URL Extension:** A text input field containing "postalCodeSearchJSON". To its right, there is a "Sample:" label with "search.json" below it, and a "Parameterize" checkbox which is currently unchecked.
- Service URL Parameters:** A text area containing "10024&country=us&maxRows=10&username: [REDACTED]". A sample text "(Sample: q=debate&page=1&rpp=100&lang=en)" is shown above the text area.
- Provide document sample:** A large empty text area. The label indicates valid formats are JSON, XML, and Encoded.
- Custom Headers:** A large empty text area. The label indicates valid formats are Content-Type='CDF', PARAM1='ABC'.

An "Add" button is located at the bottom right of the form.

After you click *Add*, the following Master File is generated.


```

FILENAME=POSTALCODESEARCH1, SUFFIX=REST      , BV_NAMESPACE=OFF, $
  SEGMENT=M6ILO, SEGTYPE=S0, $
    GROUP=HEADER, ALIAS=Header, ELEMENTS=4, $
      FIELDNAME=10024&COUNTRY, ALIAS=10024&country, USAGE=A30, ACTUAL=A30,
ACCESS_PROPERTY=(NEED_VALUE),
  XDEFAULT='us', $
    FIELDNAME=COUNTRY, ALIAS=country, USAGE=A30, ACTUAL=A30,
ACCESS_PROPERTY=(NEED_VALUE),
  XDEFAULT='us',
    GEOGRAPHIC_ROLE=COUNTRY, $
    FIELDNAME=MAXROWS, ALIAS=maxRows, USAGE=A30, ACTUAL=A30,
ACCESS_PROPERTY=(NEED_VALUE),
  XDEFAULT='10', $
    FIELDNAME=USERNAME, ALIAS=username, USAGE=A30, ACTUAL=A30,
ACCESS_PROPERTY=(NEED_VALUE),
  XDEFAULT='xxxxxx', $
    FIELDNAME=__RESPONSE, USAGE=TX80L, ACTUAL=TX,
ACCESS_PROPERTY=(INTERNAL), $
  SEGMENT=RESPONSE, SEGTYPE=S0, SEGSUF=JSON      , PARENT=M6ILO,
POSITION=__RESPONSE, $
  FIELDNAME=JSON_DUMMY_EL, ALIAS=JSON_DUMMY_EL, USAGE=A1, ACTUAL=A1,
ACCESS_PROPERTY=(INTERNAL),
  PROPERTY=ELEMENT, $
SEGMENT=POSTALCODES, SEGTYPE=S0, PARENT=RESPONSE, $
  FIELDNAME=POSTALCODES, ALIAS=postalCodes, USAGE=A55, ACTUAL=A55,
  REFERENCE=JSON_DUMMY_EL, PROPERTY=ELEMENT,
  GEOGRAPHIC_ROLE=POSTAL-CODE, $
  DEFINE POSTALCODES__CNTRY_/A50V WITH POSTALCODES='United States';
  TITLE='Country for Postal Code',
  GEOGRAPHIC_ROLE=COUNTRY, $

```

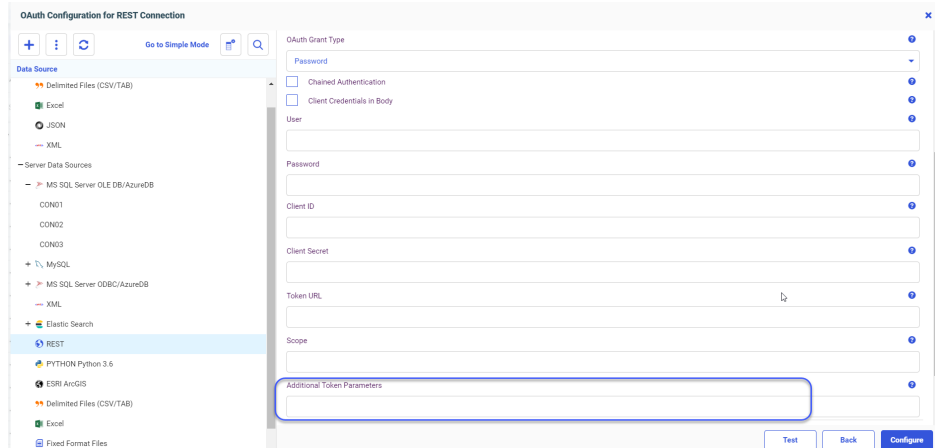
Adapter for REST: Additional Token Parameters for OAuth Client Credentials

When you configure the Adapter for REST with OAuth Client Credentials authentication, you can add additional token parameters to be sent in the token request, as shown in the following image.



Adapter for REST: Support for Token Parameters With OAuth Grant Password Authentication

When you configure a connection for the Adapter for REST and select OAuth Grant Password authentication, you can add additional token parameters, as shown in the following image.



You can use this field to enter token parameters that may be required to connect to a specific REST service, when the parameter is not part of the Scope.

Adapter for REST: Passing Multiple Values to a Parameterized URL in a WHERE Phrase

If a WHERE phrase includes multiple values for a field that is defined as a parameter in the REST URL, multiple calls will be made to the REST service, one for each value.

For example, the following Master File, postalcodesearch1.mas, returns information for the default postal code 10024.

```
FILENAME=M6ILO, SUFFIX=REST, $
SEGMENT=M6ILO, SEGTYPE=S0, $
GROUP=HEADER, ALIAS=Header, ELEMENTS=1, $
FIELDNAME=ID1, ALIAS=ID1, USAGE=A30, ACTUAL=A30,
ACCESS_PROPERTY=(NEED_VALUE), $
FIELDNAME=POSTALCODE, ALIAS=postalcode, USAGE=A30, ACTUAL=A30,
ACCESS_PROPERTY=(NEED_VALUE), XDEFAULT='10024', $
FIELDNAME=COUNTRY, ALIAS=country, USAGE=A30, ACTUAL=A30,
ACCESS_PROPERTY=(NEED_VALUE), XDEFAULT='us', $
FIELDNAME=MAXROWS, ALIAS=maxRows, USAGE=A30, ACTUAL=A30,
ACCESS_PROPERTY=(NEED_VALUE), XDEFAULT='10', $
FIELDNAME=USERNAME, ALIAS=username, USAGE=A30,
ACTUAL=A30, ACCESS_PROPERTY=(NEED_VALUE), XDEFAULT='xxxxxx', $
FIELDNAME=__RESPONSE, USAGE=TX80L, ACTUAL=TX,
ACCESS_PROPERTY=(INTERNAL), $
```

The corresponding Access File, postalcodesearch1.acx, has a parameter.

```
SEGNAME=M6ILO,
CONNECTION=geonames,
OBJECT=&ID1&,
HEADER=HEADER,
SERVICETYPE=REST,
HTTPMETHOD=POST,
RESTRESPONSE=JSON,
HTTPBODY=ENCODE, $
```

The following procedure connects to the REST URL <http://api.geonames.org>, and has a WHERE phrase with two values.

```
SET END_OF_TEXT= ' '
ENGINE REST SET CONNECTION_ATTRIBUTES
geonames_ged/, : 'http://api.geonames.org'
TABLE FILE baseapp/POSTALCODESEARCH1
PRINT
__RESPONSE
WHERE ID1 EQ 'postalCodeSearch' OR 'postalCodeSearchJSON'
ON TABLE HOLD AS BASEAPP/OUTPUT
END
```

The following output is returned. It contains the response with the two values specified in the WHERE phrase, 'postalCodeSearch' and 'postalCodeSearchJSON'.

```

<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<geonames>
  <totalResultsCount>1</totalResultsCount>
  <code>
    <postalcode>10024</postalcode>
    <name>New York</name>
    <countryCode>US</countryCode>
    <lat>40.78645</lat>
    <lng>-73.97638</lng>
    <adminCode1 ISO3166-2="NY">NY</adminCode1>
    <adminName1>New York</adminName1>
    <adminCode2>061</adminCode2>
    <adminName2>New York</adminName2>
    <adminCode3/>
    <adminName3/>
  </code>
</geonames>

{"postalCodes":[
  {
    "adminCode2":"061",
    "adminCode1":"NY",
    "adminName2":"New York",
    "lng":-73.976385,
    "countryCode":"US",
    "postalcode":"10024",
    "adminName1":"New York",
    "ISO3166-2":"NY",
    "placeName":"New York",
    "lat":40.786446
  }
]}

```

Adapter for REST: Support for CSV Format in a REST Response

The Adapter for REST can create metadata for an XML or JSON Web Service response.

If a Web Service response is in CSV format, use the following process to create the response metadata.

1. To create the initial Adapter for REST metadata, follow the process described in the section named *Managing RESTful Web Services Metadata* in the chapter named *Using the Adapter for REST* of the *Adapter Administration* manual.
2. Comment out the SEGMENT=RESPONSE line in the Master file, by adding a dollar sign character (\$) in the first column. For example:

```
$ SEGMENT=RESPONSE, SEGTYPE=S0, PARENT=M6ILO, POSITION=__RESPONSE, $
```

3. Create and run a procedure to save the entire response of the Web Service call in a file, passing any required parameter values using WHERE statements.

The following is a sample procedure that saves the CSV response for a Web Service request to file *outfile.csv* in the baseapp application.

```
SET END_OF_TEXT=' '
FILEDEF OUTFILE DISK baseapp/outfile.csv
TABLE FILE CSV_RESPONSE
PRINT
    __RESPONSE
ON TABLE HOLD AS OUTFILE
END
```

4. Using the Adapter for Delimited Files (CSV/TAB), create metadata against the file *outfile.csv* that you created in step 3.

For information about creating metadata for a delimited file, see the section named *Managing Metadata for Fixed-Format and Delimited Files* in the chapter named *Using the Adapters for Fixed-Format and Delimited Files of the Adapter Administration manual*.

The following image shows the Create Synonym screen for Delimited Files.

The screenshot shows the 'Create Synonym for Delimited Files (CSV/TAB)' configuration window. The 'Create Synonym options' section includes checkboxes for 'Scan All rows?' (unchecked) and 'Header row?' (checked). Below these are dropdown menus for 'Field Delimiter?' (set to '(comma)'), 'Field Enclosure?' (set to '(double quote)'), 'CODEPAGE?' (set to '1252 - Windows (Latin 1) - Server Default'), and 'CDN?' (set to '3,045,000.76 - digit group separator - comma, decimal mark - dot - Server Default'). The 'Advanced' section has 'Application?' set to 'baseapp'. A table lists synonyms with columns for 'Default Synonym Name' and 'Data File'. The 'outfile' synonym is selected, showing its data file as 'outfile.csv'. At the bottom, there are 'Back' and 'Add' buttons.

Default Synonym Name	Data File
baseapp3330000b43e2d8abab	baseapp3330000b43e2d8abab.csv
csv_response	csv_response.csv
outfile	outfile.csv
source_offix	source_offix.csv

5. Make the following modifications to the REST Master File.
 - a. Uncomment the `SEGMENT=RESPONSE` statement, and add `SEGSUF=DFIX`.
 - b. Copy all of the `FIELDNAME` metadata from the Master file created for the delimited file in step 4, and paste it under the `SEGMENT=RESPONSE` line in the REST Master file.

The following is a sample the Master File with these edits. The highlighted `SEGMENT=RESPONSE` declaration has `SEGSUF=DFIX`, and is followed by the FIELD definitions from the metadata created for outfile.csv in step 4.

```

FILENAME=M6ILO, SUFFIX=REST      , $
  SEGMENT=M6ILO, SEGTYPE=S0, $
  GROUP=HEADER, ALIAS=Header, ELEMENTS=1, $
  FIELDNAME=FORMAT, ALIAS=format, USAGE=A30, ACTUAL=A30,
ACCESS_PROPERTY=(NEED_VALUE),
  XDEFAULT='csv', $
  FIELDNAME=__RESPONSE, USAGE=TX80L, ACTUAL=TX,
ACCESS_PROPERTY=(INTERNAL), $
SEGMENT=RESPONSE, SEGTYPE=S0, PARENT=M6ILO, SEGSUF=DFIX,
POSITION=__RESPONSE, $
  FIELDNAME=YEAR1, ALIAS=year, USAGE=A152V, ACTUAL=A152VB,
  MISSING=ON,
  TITLE='year', $
  FIELDNAME=WEIGHT, ALIAS=weight, USAGE=D33.5, ACTUAL=A64V,
  MISSING=ON,
  TITLE='weight', $
  FIELDNAME=GENDER, ALIAS=gender, USAGE=I11, ACTUAL=A11V,
  MISSING=ON,
  TITLE='gender', $
  FIELDNAME=RACE1R, ALIAS=race1R, USAGE=I11, ACTUAL=A11V,
  MISSING=ON,
  TITLE='race1R', $
  FIELDNAME=HISPANIC, ALIAS=hispanic, USAGE=I11, ACTUAL=A11V,
  MISSING=ON,
  TITLE='hispanic', $
  FIELDNAME=ETHNIC1R, ALIAS=ethnic1R, USAGE=I11, ACTUAL=A11V,
  MISSING=ON,
  TITLE='ethnic1R', $

FIELDNAME=AGER, ALIAS=ager, USAGE=I11, ACTUAL=A11V,
  MISSING=ON,
  TITLE='ager', $
FIELDNAME=MARITAL2, ALIAS=marital2, USAGE=I11, ACTUAL=A11V,
  MISSING=ON,
  TITLE='marital2', $
FIELDNAME=HINCOME, ALIAS=hincome, USAGE=I11, ACTUAL=A11V,
  MISSING=ON,
  TITLE='hincome', $
FIELDNAME=POPSIZE, ALIAS=popsiz, USAGE=I11, ACTUAL=A11V,
  MISSING=ON,
  TITLE='popsiz', $
FIELDNAME=REGION, ALIAS=region, USAGE=I11, ACTUAL=A11V,
  MISSING=ON,
  TITLE='region', $
FIELDNAME=MSA, ALIAS=msa, USAGE=I11, ACTUAL=A11V,
  MISSING=ON,
  TITLE='msa', $
FIELDNAME=DIREL, ALIAS=direl, USAGE=I11, ACTUAL=A11V,
  MISSING=ON,
  TITLE='direl', $

```

```

FIELDNAME=NOTIFY, ALIAS=notify, USAGE=I11, ACTUAL=A11V,
MISSING=ON,
TITLE='notify', $
FIELDNAME=WEAPON, ALIAS=weapon, USAGE=I11, ACTUAL=A11V,
MISSING=ON,
TITLE='weapon', $
FIELDNAME=WEAPCAT, ALIAS=weapcat, USAGE=I11, ACTUAL=A11V,
MISSING=ON,
TITLE='weapcat', $
FIELDNAME=NEWCRIME, ALIAS=newcrime, USAGE=I11, ACTUAL=A11V,
MISSING=ON,
TITLE='newcrime', $
FIELDNAME=NEWOFF, ALIAS=newoff, USAGE=I11, ACTUAL=A11V,
MISSING=ON,
TITLE='newoff', $
FIELDNAME=SERIOUSVIOLENT, ALIAS=seriousviolent,
USAGE=I11, ACTUAL=A11V, MISSING=ON,
TITLE='seriousviolent', $
FIELDNAME=INJURY, ALIAS=injury, USAGE=I11, ACTUAL=A11V,
MISSING=ON,
TITLE='injury', $
FIELDNAME=TREATMENT, ALIAS=treatment, USAGE=I11, ACTUAL=A11V,
MISSING=ON,
TITLE='treatment', $
FIELDNAME=VICSERVICES, ALIAS=vicservices, USAGE=I11, ACTUAL=A11V,
MISSING=ON,
TITLE='vicservices', $
FIELDNAME=LOCATIONNR, ALIAS=locationnr, USAGE=I11, ACTUAL=A11V,
MISSING=ON,
TITLE='locationnr', $

```

6. Make the following modifications to the REST Access File.

- a. Insert `RESTRESPONSE=DFIX`, before the dollar sign (\$).
- b. Add a line with `SEGNAME=RESPONSE`.
- c. Copy all of the lines under `SEGMENT`, except `CONNECTION`, from the Access file created for the delimited file `outfile.csv` in step 4, and paste it under the `SEGMENT=RESPONSE` line you added in the REST Access file.
- d. Add an `RDELIMITER` line.

The value of the attribute will be `0x0A` for ASCII platforms such as Windows or Linux (`RDELIMITER=0x0A`).

The value of the attribute will be `0x25` for OS400 platforms (`RDELIMITER=0x25`).

The value of the attribute will be `0x15` for z/OS platforms (`RDELIMITER=0x15`).

- e. Add a comma and dollar sign (`, $`) after the `RDELIMITER` value.

The following is a sample Access File with these edits, on a Windows or Linux platform.

```
SEGNAME=M6ILO,  
  CONNECTION=bjsncvs,  
  OBJECT=personal/2019,  
  HEADER=HEADER,  
  SERVICETYPE=REST,  
  HTTPMETHOD=GET,  
  RESTRESPONSE=DFIX,$  
SEGNAME=RESPONSE,  
  DELIMITER=',',  
  ENCLOSURE=',',  
  HEADER=YES,  
  CDN=COMMAS_DOT,  
  RDELIMITER=0x0A,$
```

Adapter for REST: Support for No Value in a Label in the JSON Body

The Adapter for REST can now pass a label with no value within the JSON Body for a POST request.

For example, WHERE COUNTRY EQ "" will be shown as "country": "" within the JSON Body.

Adding OMIT_MISSING=NO in the Access (.acx) File will produce the following properties in the JSON Body:

- Fields defined in the Master File with MISSING=ON will send a null value, if there is no WHERE clause defined for that field.
- Fields with MISSING=OFF defined in the Master File will send the property "value":"" for a non-numeric field, and the property "value":0 for a numeric field, if there is no WHERE clause defined for that field.

Example: Passing a Label With No Value in the JSON Body

The following Access File specifies a JSON POST response, with OMIT_MISSING=NO.

```
SEGNAME=M6ILO,  
  CONNECTION=geonames,  
  OBJECT=postalCodeSearchJSON,  
  SERVICETYPE=REST,  
  HTTPMETHOD=POST,  
  RESTRESPONSE=JSON,  
  OMIT_MISSING=NO,  
  HTTPBODY=JSON, $
```

The Master File has the following field definition.

```
FIELDNAME=ADMINCODE1, ALIAS=adminCode1, USAGE=A55, ACTUAL=A55,  
MISSING=OFF, ACCESS_PROPERTY=(NEED_VALUE), $
```


With these attributes defined, the JSON Body will automatically include "adminCode1": "", if there is no WHERE or IF clause for ADMINCODE1 defined in the WebFOCUS report request.

Adapter for REST: Adding Scope to OAuth Access Token Request

When you configure a connection for the Adapter for REST using OAuth authentication, a *Scope* parameter has been added when you select the Password or Client Credentials Grant Type.

When there is a value configured for Scope, the Scope parameter will be sent in the Token request.

Adapter for REST: Encoding Special Characters When Creating a Synonym

In prior releases of the Adapter for REST, users had to manually enter encoded values for special characters within parameter values in the *Create Synonym Service URL Parameters* and *Document Sample* text boxes.

The Adapter for REST now automatically encodes special characters within parameter values when creating a synonym.

For example, the synonym named *postalsearch1* was generated using Service URL Extension *postalCodeSearch* and Service URL Parameters *placename=Thézan-des-Corbières&maxRows=20&username=demo*. The default *placename* value contains special characters. The Base URL in the connection attributes is <http://api.geonames.org>.

The following is the generated Master File.

```
FILENAME=M6ILO, SUFFIX=REST, $
  SEGMENT=M6ILO, SEGTYPE=S0, $
    GROUP=HEADER, ALIAS=Header, ELEMENTS=3, $
      FIELDNAME=PLACENAME, ALIAS=placename, USAGE=A30, ACTUAL=A30,
ACCESS_PROPERTY=(NEED_VALUE),
  XDEFAULT='Thézan-des-Corbières', $
    FIELDNAME=MAXROWS, ALIAS=maxRows, USAGE=A30, ACTUAL=A30,
ACCESS_PROPERTY=(NEED_VALUE),
  XDEFAULT='1000', $
    FIELDNAME=USERNAME, ALIAS=username, USAGE=A30, ACTUAL=A30,
ACCESS_PROPERTY=(NEED_VALUE),
  XDEFAULT='ibiibi', $
    FIELDNAME=__RESPONSE, USAGE=TX80L, ACTUAL=TX,
ACCESS_PROPERTY=(INTERNAL), $
```

```

SEGMENT=RESPONSE, SEGTYPE=S0, SEGSUF=XML      , PARENT=M6ILO,
POSITION=__RESPONSE, $
  FIELDNAME=GEONAMES, ALIAS=geonames, USAGE=A1, ACTUAL=A1,
ACCESS_PROPERTY=(INTERNAL),
  PROPERTY=ELEMENT, $
  FIELDNAME=TOTALRESULTSCOUNT, ALIAS=totalResultsCount, USAGE=P32,
ACTUAL=A32,
  REFERENCE=GEONAMES, PROPERTY=ELEMENT, $
  FIELDNAME=CODE, ALIAS=code, USAGE=A1, ACTUAL=A1,
ACCESS_PROPERTY=(INTERNAL),
  REFERENCE=GEONAMES, PROPERTY=ELEMENT, $
  FIELDNAME=POSTALCODE, ALIAS=postalcode, USAGE=P32, ACTUAL=A32,
  REFERENCE=CODE, PROPERTY=ELEMENT, $
  FIELDNAME=NAME, ALIAS=name, USAGE=A55, ACTUAL=A55,
  REFERENCE=CODE, PROPERTY=ELEMENT, $
  FIELDNAME=COUNTRYCODE, ALIAS=countryCode, USAGE=A55, ACTUAL=A55,
  REFERENCE=CODE, PROPERTY=ELEMENT, $
  FIELDNAME=LAT, ALIAS=lat, USAGE=P20.3, ACTUAL=A20,
  REFERENCE=CODE, PROPERTY=ELEMENT, $
  FIELDNAME=LNG, ALIAS=lng, USAGE=P20.3, ACTUAL=A20,
  REFERENCE=CODE, PROPERTY=ELEMENT, $
  FIELDNAME=ADMINCODE1, ALIAS=adminCode1, USAGE=A30, ACTUAL=A30,
  REFERENCE=CODE, PROPERTY=ELEMENT, $
  FIELDNAME=ISO3166_2, ALIAS=ISO3166-2, USAGE=A30, ACTUAL=A30,
  REFERENCE=ADMINCODE1, PROPERTY=ATTRIBUTE, $
  FIELDNAME=ADMINNAME1, ALIAS=adminName1, USAGE=A55, ACTUAL=A55,
  REFERENCE=CODE, PROPERTY=ELEMENT, $
  FIELDNAME=ADMINCODE2, ALIAS=adminCode2, USAGE=P32, ACTUAL=A32,
  REFERENCE=CODE, PROPERTY=ELEMENT, $
  FIELDNAME=ADMINNAME2, ALIAS=adminName2, USAGE=A55, ACTUAL=A55,
  REFERENCE=CODE, PROPERTY=ELEMENT, $
  FIELDNAME=ADMINCODE3, ALIAS=adminCode3, USAGE=P32, ACTUAL=A32,
  REFERENCE=CODE, PROPERTY=ELEMENT, $
  FIELDNAME=ADMINNAME3, ALIAS=adminName3, USAGE=A55, ACTUAL=A55,
  REFERENCE=CODE, PROPERTY=ELEMENT, $

```

The following is the generated Access File.


```

SEGNAME=M6ILO,
CONNECTION=CON02,
OBJECT=postalCodeSearch,
HEADER=HEADER,
SERVICETYPE=REST,
HTTPMETHOD=GET,
RESTRESPONSE=XML, $

```

The following request prints the PLACENAME field, where the special characters were encoded.

```
TABLE FILE postalcodesearch1
PRINT PLACENAME
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF, $
ENDSTYLE
END
```



Adapter for REST: Sending Credentials in the POST Request Body

When OAuth authentication is used for the Adapter for REST, and the OAuth Grant Type is *Password*, the configuration parameter *Client Credentials in Body* has been added to the connection configuration panel, as shown in the following image.

The image shows the 'Add Connection for REST OAuth' configuration panel. The 'Connect parameters' section is expanded. The 'Client Credentials in Body' checkbox is checked, and a mouse cursor is pointing at it. Other visible options include 'Chained Authentication' (unchecked), 'User', 'Password', 'Client ID', and 'Client Secret'. The 'Test', 'Back', and 'Configure' buttons are at the bottom right.

Selecting this option will add the attribute *bdyclt:true* in the connection string.

By default, client credentials for OAuth Grant Type *Password* are sent Base64 encoded in the HTTP header. However, if the web service being used does not support sending these credentials in the HTTP header, you can use this feature to send them in the body instead.

Adapter for REST: Sending Multipart/Form Data in a POST Request

You can send a single part of content type Multipart/Form data in a POST request using the Adapter for REST.

You need the following declarations in the synonym.

- ❑ **Define the body of the POST request to be sent as multipart/form-data.**

Access File Declaration:

```
HTTPBODY=MULTIPART, $
```

Result. Adds Content-Type: multipart/form-data to the HTTP Header of the REST request.

- ❑ **Define the field name that contains the value for *Content-Disposition* within the Part Header.**

Example:

```
Content-Disposition: form-data
```

Access File Declaration:

```
FIELD=part1_content_disposition, FORMAT=PARTHEADER, $
```

Master File Declaration:

```
FIELDNAME=part1_content_disposition, ALIAS=Content-Disposition,  
USAGE=A30, ACTUAL=A30, XDEFAULT='form-data', $
```

- ❑ **Define the field name that contains the value for *name* within the Part Header.**

Example:

```
name="file"
```

Access File Declaration:

```
FIELD=part1_name, FORMAT=PARTHEADER, $
```

Master File Declaration:

```
FIELDNAME=part1_name, ALIAS=name,  
USAGE=A30, ACTUAL=A30, XDEFAULT='file', $
```

- ❑ **Define the field name that contains the value for *filename* within the Part Header.**

Example:

```
filename="reportcaster_version.xlsx"
```

Access File Declaration:

```
FIELD=part1_filename, FORMAT=PARTHEADER, $
```

Master File Declaration:

```
FIELDNAME=part1_filename, ALIAS=filename, USAGE=A100, ACTUAL=A100, $
```

- ❑ **Define the field name that contains the value for *Content-Type* within the Part Header.**

If the value is not supplied, Content-Type will be set based on the extension of the value for *filename*.

Example:

```
Content-Type: application/vnd.openxmlformats-officedocument.spreadsheetml.sheet
```

Access File Declaration:

```
FIELD=part1_content_type, FORMAT=PARTHEADER, $
```

Master File Declaration:

```
FIELDNAME=part1_content_type, ALIAS=Content-Type,
      USAGE=A200, ACTUAL=A200, $
```

Adapter for REST: Support for STRING Format

The Adapter for REST now supports STRING USAGE and ACTUAL formats in both the request and response. Since STRING format supports an unlimited string, PDF, binary, and graph output can be retrieved from a variety of Web Services using this format.

Adapter for REST: Setting a Default Field Length

The Adapter for REST SET FIELDLENGTH command sets a default USAGE and ACTUAL length for alphanumeric fields to be used when creating a synonym.

The syntax is:

```
ENGINE REST SET FIELDLENGTH nnn
```

where:

nnn

Is the default length.

For example, the following command sets the default USAGE and ACTUAL lengths for alphanumeric fields to 100.

```
ENGINE REST SET FIELDLENGTH 100
```

Adapter for REST: Support for Passing Text and Binary in a REST POST Request

The Adapter for REST now can read text and binary data from a file and pass it in the body of a POST request.

The following Access File attribute is used for passing text data to another REST request.

```
HTTPBODY=TEXT
```

The following Access File attribute is used for passing binary data to another REST request.

```
HTTPBODY=BINARY
```

The following Access File attribute is used to define the field used to pass the location of the file to be read for input to the REST call.

```
FIELD=DATAFILE, INPUT=FILENAME
```

Support has also been added for reading the binary response from a REST Web Service request and saving the output to a file. The USAGE and ACTUAL formats for the field containing the Binary output must be defined as STRING.

The following Access File attribute is used for reading binary data from a REST request.

```
RESTRESPONSE=BINARY
```

The following command is used to HOLD the binary output:

```
ON TABLE HOLD FORMAT BINFILE DATASET filename
```

where:

filename

Is the name of the file where the output is to be stored.

For example:

```
ON TABLE HOLD FORMAT BINFILE DATASET baseapp/wfretail.xlsx
```

You can also join the binary response from one REST Web Service request to the input of a subsequent REST Web Service request. The fieldname containing the binary response and the fieldname for the binary input both must be defined with USAGE=STRING and ACTUAL=STRING.

Example: **Reading Binary Data from a REST POST Request and Storing the Output to a File**

The following Master File describes the GET_LIBRARYVERSION REST call. The __RESPONSE field is described with USAGE and ACTUAL formats STRING:

```
FILENAME=M6ILO, SUFFIX=WFCREST      , $
SEGMENT=M6ILO, SEGTYPE=S0, $
  FIELDNAME=IBIRS_ACTION, ALIAS=IBIRS_action, USAGE=A30, ACTUAL=A30,
  ACCESS_PROPERTY=(NEED_VALUE),
  XDEFAULT='run', $
  FIELDNAME=IBIRS_SERVICE, ALIAS=IBIRS_service, USAGE=A30, ACTUAL=A30,
  ACCESS_PROPERTY=(NEED_VALUE),
  XDEFAULT='ibfs', $
  FIELDNAME=IBIRS_PATH, ALIAS=IBIRS_path, USAGE=A100, ACTUAL=A100,
  ACCESS_PROPERTY=(NEED_VALUE), $
  FIELDNAME=__RESPONSE, USAGE=STRING, ACTUAL=STRING,
  ACCESS_PROPERTY=(INTERNAL), $
```

Following is the associated Access File:

```
SEGNAME=M6ILO,
CONNECTION=WIN10HYPER1,
OBJECT=rs,
SERVICETYPE=REST,
TIMEOUT=60,
HTTPMETHOD=POST,
HTTPBODY=ENCODE, $
```

The following request issues the REST call and stores the result in a file named holdversion.xlsx in the baseapp application:

```
SET EQTEST=EXACT
TABLE FILE GET_LIBRARYVERSION
PRINT
  __RESPONSE
WHERE IBIRS_PATH EQ
  'IBFS:/WFC/Repository/Binary_Data_TSCQ/L1e7us4uhe01.lib$(1)'
ON TABLE HOLD FORMAT BINFILE DATASET baseapp/holdversion.xlsx
END
```

Search Engine Adapters

This section provides descriptions of new features for Search Engine adapters.

Adapter for ElasticSearch

The Adapter for ElasticSearch is new in this release and can be found in the Search Engines group folder.

ElasticSearch is a highly-scalable search engine that performs fast searches and analytics against large amounts of data. It centrally stores all types of data, including textual, numerical, geospatial, structured, and unstructured. The Adapter for ElasticSearch is used to send search requests to the ElasticSearch engine and report against the information returned.

Social Media Adapters

This section provides descriptions of new features for Social Media adapters.

Adapter for Google Analytics: Support for Service Account Authentication

In prior releases, only OAuth20 authentication was supported for the Adapter for Google Analytics. Service Account Authentication support has now been added. A service account is a non-human user that needs to authenticate to Google and be authorized to access data in Google APIs, for example, for running reports distributed by ReportCaster.

To configure Service Account authentication, select *Service Account* from the *Security* drop-down list, as shown in the following image.

The screenshot shows a configuration window titled "Add Google Analytics to Configuration". It contains several input fields and a dropdown menu. The fields are: "Connection Name" (containing "CON01"), "Google Analytics URL" (containing "https://www.googleapis.com/analytics/v3/data"), and "Web Profile ID" (empty). The "Security" dropdown menu is open, showing two options: "OAuth" and "Service Account". The "Service Account" option is highlighted.

The following fields appear for configuring Service Account authentication.

Key File Location

Is the location of the JSON Key File that was created in the Google Development environment. For example:

```
C:\ibi\apps\googleanalytics\webfocus-555555-874feowfwojoe7.json
```


Client Email Address

Is the client email address defined in the JSON Key File. For example:

```
xxxxxxxxxxxx-compute@developer.gserviceaccount.com
```

Adapter for Google Drive

The Adapter for Google Drive is the WebFOCUS Adapter used to integrate with Google Drive. It is also used in conjunction with the Adapter for Google Sheets to create spreadsheets in the Google Sheets environment from a WebFOCUS report.

Output formats such as PDF, JSON, and XML from WebFOCUS reports can be stored in Google Drive.

Adapters, such as JSON and XML, can be used to report against their respective documents stored in Google Drive.

Adapter for Google Sheets

The Adapter for Google Sheets is now available and can be used to report off of spreadsheet data resident in the Google Sheets environment.

Generating a Google Sheets Spreadsheet

When you have the Adapter for Google Sheets and the Adapter for Google Drive configured, you can generate a Spreadsheet in the Google Sheets environment from a WebFOCUS request using the HOLD FORMAT GGLSHTS command.

The syntax is:

```
ON TABLE HOLD FORMAT GGLSHTS AS GoogleDriveApplication/Name
```

where:

GoogleDriveApplication

Is the name of the WebFOCUS Server application mapped to Google Drive.

Name

Is the name of the Google Spreadsheet that is to be stored in Google Sheets.

Adapters for Google Sheets and Google Drive: Service Account Authentication

Service Account Authentication is available for Google Sheets and Google Drive. A service account is a non-human user that needs to authenticate to Google and be authorized to access data in Google APIs, for example, for running reports distributed by ReportCaster.

To configure Service Account authentication, select *Service Account* from the *Security* drop-down list, as shown in the following image.

Add Google Sheets to Configuration

Connect parameters ^

Connection Name ?

CON01

Base Uri ?

<https://sheets.googleapis.com/v4/spreadsheets>

Security ?

Service Account v

OAuth

Service Account

The following fields appear for configuring Service Account authentication.

Key File Location

Is the location of the JSON Key File that was created in the Google Development environment. For example:

```
C:\ibi\apps\googlesheets\webfocus-555555-874feowfwojoe7.json
```

Client Email Address

Is the client email address defined in the JSON Key File. For example:

```
xxxxxxxxxxxx-compute@developer.gserviceaccount.com
```

XML Adapters

This section describes new features for the XML adapters.

Adapters for SharePoint and SharePoint Drive: Passing Client Credentials in the Body

You now have the option to pass the OAuth Grant Password Client ID and Client Secret in the Body of the Token request for the Adapters for SharePoint and SharePoint Drive.

When configuring a connection, select the *Client Credentials in Body* check box, as shown in the following image for the Adapter for SharePoint Drive.

The image shows a dialog box titled "Change Connect Parameters for SharePoint Drive". It contains several sections for configuration:

- Connect parameters**: Includes "Connection Name" (CON01) and "SharePoint URL" (https://xxxxxxxxx.sharepoint.com/_api/v2.1).
- Security**: Includes "OAuth Password" and a checked checkbox for "Client Credentials in Body", which is highlighted with a red arrow.
- User**: Includes a text input field.
- Password**: Includes a text input field.
- Application ID**: Includes a text input field.
- Application Secret**: Includes a text input field.

At the bottom right, there are three buttons: "Test", "Cancel", and "Configure".

The following image shows the *Client Credentials in Body* option selected for the Adapter for SharePoint.

Change Connect Parameters for SharePoint

Connect parameters

Connection Name ?
CONT

Graph API URL ?
https://graph.microsoft.com/v1.0
[Show Sample](#)

Security ?
OAuth Password

Client Credentials in Body ? ←

User ?
[Empty text box]

Password ?
[Empty text box]

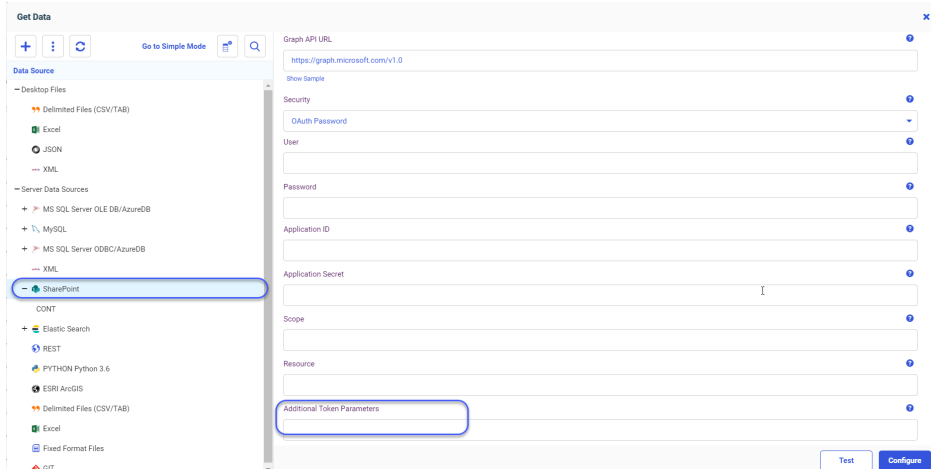
Application ID ?
[Empty text box]

Application Secret ?
[Empty text box]

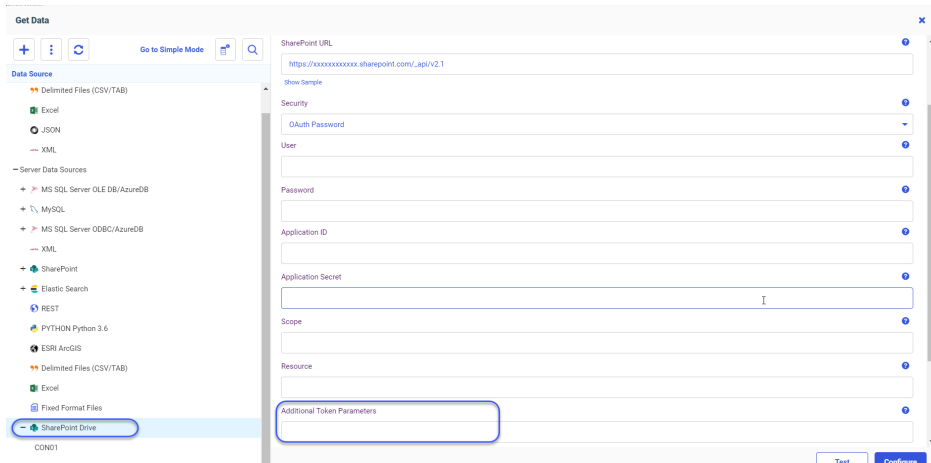
Test Cancel **Configure**

Adapters for SharePoint and SharePoint Drive: Support for Token Parameters With OAuth Grant Password Authentication

When you configure a connection for the Adapter for SharePoint or SharePoint Drive and select OAuth Grant Password authentication, you can add additional token parameters, as shown in the following image for the Adapter for SharePoint.



The following image shows the connection parameters for the Adapter for SharePoint Drive, with OAuth Grant Password authentication.



You can use this field to enter token parameters that may be required to connect to a specific SharePoint service, when the parameter is not part of the Scope.

Adapters for SharePoint and SharePoint Drive: Support for OAuth Grant Password Authentication

In prior releases of the Adapters for SharePoint and SharePoint Drive, the only type of OAuth authentication that was supported was OAuth Authorization Code authentication. Now OAuth Grant Password authentication is also supported, as shown in the following image for the Adapter for SharePoint Drive.

The image shows a configuration window titled "Add SharePoint Drive to Configuration". It contains several input fields and a dropdown menu. The "Security" dropdown is open, and "OAuth Password" is selected and highlighted with a red circle. Other options in the dropdown include "NTLM", "OAuth Authorization Code", and "Trusted". The "Connection Name" field contains "CON01". The "SharePoint URL" field contains "https://xxxxxxxxxxx.sharepoint.com/_api/v2.1". The "Application ID" field is empty. The "Refresh Token" field is empty. There are "Test" and "Configure" buttons at the bottom right.

The following image shows the Adapter for SharePoint configuration screen.

The screenshot shows the 'Add SharePoint to Configuration' interface. It includes a 'Connect parameters' section with the following fields:

- Connection Name:** CON01
- Graph API URL:** https://graph.microsoft.com/v1.0
- Security:** A dropdown menu is open, showing options: OAuth Authorization Code, NTLM, OAuth Authorization Code, **OAuth Password** (highlighted with a red circle), and Trusted.
- Application ID:** (Empty field)
- Refresh Token:** (Empty field) with a 'Get Refresh Token' link.

Buttons for 'Test' and 'Configure' are located at the bottom right.

Configure the following parameters for OAuth Grant Password authentication.

Connection Name

Is the logical name used to identify this particular set of connection attributes.

The default value is CON01.

SharePoint URL (for SharePoint Drive)

Is the URL of the SharePoint API request.

For example:

`https://xxxxxxxxxxxxx.sharepoint.com/_api/v2.1`

Graph API URL (for SharePoint)

Is the URL of the SharePoint request.

For example:

`https://graph.microsoft.com/v1.0`

Security

Select *OAuth Password* from the *Security* drop-down list.

User

Is the user ID used to authenticate to SharePoint, when using OAuth Password authentication.

Password

Is the password used to authenticate to SharePoint, when using OAuth Password authentication.

Application Secret

Is the Application Secret for the application created in the Azure Active Directory, when using OAuth Password authentication.

Obtain the Application Secret using the following steps:

1. Go to:
<https://portal.azure.com>
2. Click *Azure Active Directory* in the left panel.
3. Click *App registrations*.
4. Click the registered application.
5. Click *Certificates & secrets*.
6. Click + *New client secret* to create a Client Secret.

The Application Secret is the value for the created Client Secret.

Scope

Is the scope sent as part of the OAuth Token request, when using OAuth Password authentication (optional). For example, *openid*.

Resource

Is the resource sent as part of the OAuth Token request, when using OAuth Password authentication (optional). For example:

<https://shareptdomain>

where:

shareptdomain

Is the domain when accessing SharePoint, for example, <https://mycompany.sharepoint.com>.

Application ID

Is the value that identifies your application to SharePoint when using OAuth Authentication.

Obtain this value using the following steps:

1. Go to:

<https://portal.azure.com>

2. Click *Azure Active Directory* in the left panel.
3. Click *App registrations*.
4. Click the Application for the SharePoint Adapter that was previously created.
5. Use the value from the Application (client) ID.

Token URL

Is the URL used to obtain a Token to access SharePoint.

Obtain this URL using the following steps:

1. Go to:

<https://portal.azure.com>

2. Click *Azure Active Directory* in the left panel.
3. Click *App registrations*.
4. Click *Endpoints*.
5. Use the URL from the OAuth 2.0 token endpoint (v1).

Adapter for SharePoint: Create Synonym Support for Multiple Sites in List Mode

In prior releases, when you created a synonym using the Adapter for SharePoint and selected Lists mode, only one site displayed. Now, if multiple sites are available, they all appear, as shown in the following image.

Create Synonym for SharePoint (SharePoint_Online)

Mode: Lists Entities

Sites	↓ Drill Down
Information Builders Inc. Team Site	→
teamsitetest	→
PointPublishing Hub Site	→
Community	→
Test Team #2	→
Test Team	→

You can click the arrow to the right of any site in the Drill Down column to open a table of lists for that site, as shown in the following image.

Create Synonym for SharePoint (SharePoint_Online)

Mode: Lists Entities

Search

Select	Lists
<input checked="" type="radio"/>	AppPages
<input type="radio"/>	Hub Settings
<input type="radio"/>	PortalSiteList

Back Next

Select a list and click *Next*.

The Create Synonym page opens for the selected list, as shown in the following image.

Create Synonym for SharePoint (SharePoint_Online) : AppPages

Customize data type mappings

Miscellaneous settings

Application ? Synonym name ?

Synonym Field Names Processing Options

Validate Make Unique

Back Add

Enter or select the appropriate options, as described in the chapter named *Using the Adapter for SharePoint* in the *Adapter Administration* manual, and click *Add* to create the synonym.

Adapter for SharePoint

The Adapter for SharePoint is new in this release. It is used to report against SharePoint Lists residing in the SharePoint Online/Modern environment.

The Adapter for SharePoint is in the XML Based adapter category.

Adapter for OData: Passing Client Credentials in the Body

You now have the option to pass the OAuth Grant Password Client ID and Client Secret in the Body of the Token request for the Adapter for OData.

When configuring a connection, select the *Client Credentials in Body* check box, as shown in the following image.

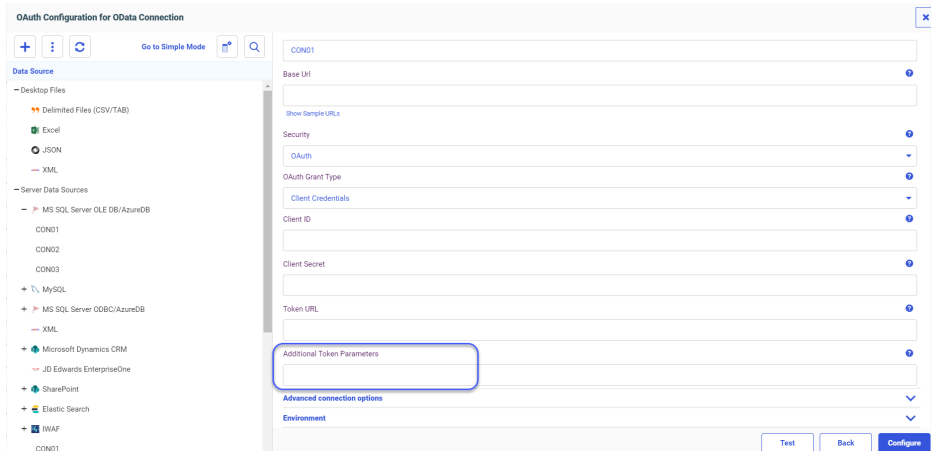
The image shows a configuration window titled "Add Connection for OData OAuth". It contains several fields and a checkbox. A red arrow points to the "Client Credentials in Body" checkbox, which is checked. The fields include:

- Connect parameters (header)
- Connection Name: CON01
- Base Url: https://services.odata.org/V4/Northwind/Northwind.svc/
- Security: OAuth
- OAuth Grant Type: Password
- Client Credentials in Body: (highlighted with a red arrow)
- User: (empty field)
- Password: (empty field)
- Client ID: (empty field)

At the bottom right, there are three buttons: "Test", "Back", and "Configure".

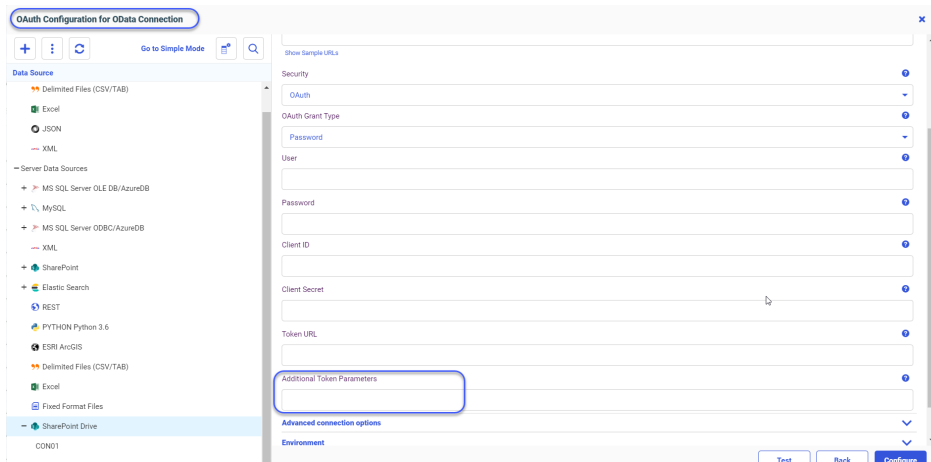
Adapter for OData: Additional Token Parameters for OAuth Client Credentials

When you configure the Adapter for OData with OAuth Client Credentials authentication, you can add additional token parameters to be sent in the token request, as shown in the following image.



Adapter for OData: Support for Token Parameters With OAuth Grant Password Authentication

When you configure a connection for the Adapter for OData and select OAuth Grant Password authentication, you can add additional token parameters, as shown in the following image.



You can use this field to enter token parameters that may be required to connect to a specific OData service, when the parameter is not part of the Scope.

Adapter for OData: Support for Date-Time Functions in Filters

When a field is defined with OData data type Date or DateTime, you have the option to enable Date Functions processing. With this option enabled, decomposed dates in WHERE/IF statements will be processed by the OData Service, rather than the Reporting Server engine. If *Enable Date Functions* is not selected, and *Decompose Date fields into components* is On, DEFINE fields will be created in the metadata for the Date and DateTime fields, and these will be processed by the Reporting Server.

In the following image of the Create Synonym screen, the Orders entity is selected, with *Enable Date Functions* also selected.

The screenshot shows the 'Create Synonym for OData (NorthWind)' interface. The 'Create Synonym options' section includes a 'Metadata type?' dropdown set to 'Entity'. The 'Miscellaneous settings' section shows 'Application?' set to 'baseapp', with 'Prefix?' and 'Suffix?' fields empty. Under 'Synonym Field Names Processing Options', both 'Validate' and 'Make Unique' are checked. A table lists various entities with their corresponding 'Function' and 'Enable Date Functions' checkboxes. The 'Orders' entity is selected, and its 'Enable Date Functions' checkbox is checked.

Default Synonym Name	Entity	Function	Enable Date Functions
<input type="checkbox"/>	Customer_and_Suppliers_by_Cities	Customer_and_Suppliers_by_Cities	<input type="checkbox"/>
<input type="checkbox"/>	Customers	Customers	<input type="checkbox"/>
<input type="checkbox"/>	Employees	Employees	<input type="checkbox"/>
<input type="checkbox"/>	Invoices	Invoices	<input type="checkbox"/>
<input type="checkbox"/>	Order_Details	Order_Details	<input type="checkbox"/>
<input type="checkbox"/>	Order_Details_Extended	Order_Details_Extended	<input type="checkbox"/>
<input type="checkbox"/>	Order_Subtotals	Order_Subtotals	<input type="checkbox"/>
<input checked="" type="checkbox"/>	Orders	Orders	<input checked="" type="checkbox"/>
<input type="checkbox"/>	Orders_Qries	Orders_Qries	<input type="checkbox"/>
<input type="checkbox"/>	Product_Sales_for_1997	Product_Sales_for_1997	<input type="checkbox"/>
<input type="checkbox"/>	Products	Products	<input type="checkbox"/>
<input type="checkbox"/>	Products_Above_Average_Prices	Products_Above_Average_Prices	<input type="checkbox"/>

The following image shows the Create Synonym screen, with *Decompose Date fields into components* set *On*.

The screenshot shows the 'Create Synonym for OData (NorthWind)' configuration window. The 'Customize data type mappings' section is expanded, and the 'Decompose Date fields into components?' option is set to 'On'. Other options include 'Date Order?' (set to 'Adapter specific: default'), 'Activate GEOGRAPHIC_ROLE assignment?' (set to 'On'), 'Set data type DOUBLE for numeric decimal columns?' (set to 'On'), and 'Set data type ANV (On) / STRING for character columns?' (set to 'On (ANV)'). The 'Miscellaneous settings' section shows 'Application?' set to 'baseapp', 'Prefix?' and 'Suffix?' fields, and 'Synonym Field Names Processing Options' with 'Validate' and 'Make Unique' checked. A table below lists default synonym names and their mappings.

Default Synonym Name	Entity	Function	Enable Date Functions
Customer_and_Suppliers_by_Cities	Customer_and_Suppliers_by_Cities		<input type="checkbox"/>
Customers	Customers		<input type="checkbox"/>
Employees	Employees		<input type="checkbox"/>
Invoices	Invoices		<input type="checkbox"/>
Order_Details	Order_Details		<input type="checkbox"/>

The following request uses a synonym created with *Decompose Date fields into components* set *On*. ORDERDATE_YEAR is the Year component of the ORDERDATE field, and is used in a WHERE phrase.

```
TABLE FILE ORDERS
PRINT
SHIPNAME
ORDERDATE
WHERE ORDERDATE_YEAR EQ '1996'
END
```

Enable Date Functions not selected

With *Enable Date Functions* not selected and *Decompose Date fields into components* set to *On*, the following DEFINE field is generated in the Master File:

```
DEFINE ORDERDATE_YEAR/I4 MISSING ON ALL=DTPART(ORDERDATE, YEAR);
TITLE= 'ORDERDATE, Year', $
```

The following parameters are generated for the OData request, which return the entire date. The year component is then calculated on the Reporting Server, using the DEFINE field.

```
$select=OrderID, OrderDate, ShipName
```

Enable Date Functions selected

With *Enable Date Functions* selected and *Decompose Date fields into components* set to *On*, the following field definition is generated in the Master File, which contains the OData year function in the ALIAS value, so the processing can be passed to the OData request:

```
FIELDNAME=ORDERDATE_YEAR, ALIAS='year (OrderDate) ', USAGE=I11, ACTUAL=A11,
MISSING=ON, REFERENCE=ORDERS, $
```

The following parameters are generated for the OData request, which filter the data returned to the Reporting Server by the year component.

```
$select=OrderID,OrderDate,ShipName&$filter=year(OrderDate)%20eq%201996
```

Adapter for OData: Support for Reformatting Fields

In prior releases, the Adapter for OData did not pass aggregation to the OData engine if SUM or BY fields were reformatted to a different USAGE format in a request. The Adapter for OData can now pass aggregation to OData when SUM or BY fields are reformatted in a request.

Adapter for OData: Support for MISSING

The Adapter for OData now translates the filters *field EQ MISSING* and *field NE MISSING* in a report request to *field eq null* and *field ne null* in the generated OData aggregated or non-aggregated request.

Only the root segment and unique child segments (SEGTYPE=U) are supported for filtering on MISSING in the OData request.

The following request is an example of a TABLE request that filters on MISSING.

```
TABLE FILE CUSTOMERS
PRINT
CUSTOMERS.CUSTOMERS.COMPANYNAME
CUSTOMERS.CUSTOMERS.COUNTRY
IF CUSTOMERS.CUSTOMERS.COUNTRY NE MISSING
END
```

The following is the generated OData request, with the filter translated to a filter on null.

```
https://services.odata.org/V4/Northwind/Northwind.svc/Customers?
$select=CustomerID,CompanyName,Country&$filter=Country%20ne%20null
```


Adapter for OData: Turn Off Aggregation for ENTITY_SET Requests

In prior releases, aggregation requests were always sent for aggregate functions (such as CNT.DST.) when Aggregation was set on in the Connection string (EXT_AGGR:Y) and retrieval was started at the child segment by adding the ENTITY_SET attribute in the Child Segment definition of the Access File (for example, ENTITY_SET=Assets).

Now, if no fields from the root segment are referenced in the request, aggregation will not be performed. The Session Log will contain a message indicating that aggregation was turned off for the request.

Adapter for OData: Starting Retrieval With the Child Object in a Static Join

The Adapter for OData can now access data directly from a Related Entity rather than using the \$expand parameter in the OData request.

A related Entity is described in the Master File as a child segment of the parent.

Add the following attribute in the child segment definition of the Access file to create an OData request directly against the defined Entity, assuming all requested fields are from that child segment.

ENTITY_SET=entityname

where:

entityname

Is the name of the Related Entity.

Example: Accessing Data Directly From a Related Entity

The following is the Master File customers.mas, which has three segments. The parent segment is CUSTOMERS. The child segment from which data will be retrieved is ORDERS.

```

FILENAME=M6ILO, SUFFIX=ODATAV4 , $
  SEGMENT=CUSTOMERS, SEGTYPE=S0, $
    FIELDNAME=CUSTOMERS, ALIAS=value, USAGE=A1, ACTUAL=A1,
ACCESS_PROPERTY=(INTERNAL), $
  FIELDNAME=CUSTOMERID, ALIAS=CustomerID, USAGE=A5, ACTUAL=A5,
  MISSING=ON,
  REFERENCE=CUSTOMERS, $
  FIELDNAME=COMPANYNAME, ALIAS=CompanyName, USAGE=A40, ACTUAL=A40,
  MISSING=ON,
  REFERENCE=CUSTOMERS, $
  FIELDNAME=CONTACTNAME, ALIAS=ContactName, USAGE=A30, ACTUAL=A30,
  MISSING=ON,
  REFERENCE=CUSTOMERS, $
  FIELDNAME=CONTACTTITLE, ALIAS=ContactTitle, USAGE=A30, ACTUAL=A30,
  MISSING=ON,
  REFERENCE=CUSTOMERS, $
  FIELDNAME=ADDRESS, ALIAS=Address, USAGE=A60, ACTUAL=A60,
  MISSING=ON,
  REFERENCE=CUSTOMERS, $
  FIELDNAME=CITY, ALIAS=City, USAGE=A15, ACTUAL=A15,
  MISSING=ON,
  REFERENCE=CUSTOMERS, $
  FIELDNAME=REGION, ALIAS=Region, USAGE=A15, ACTUAL=A15,
  MISSING=ON,
  REFERENCE=CUSTOMERS, $
  FIELDNAME=POSTALCODE, ALIAS=PostalCode, USAGE=A10, ACTUAL=A10,
  MISSING=ON,
  REFERENCE=CUSTOMERS, $
  FIELDNAME=COUNTRY, ALIAS=Country, USAGE=A15, ACTUAL=A15,
  MISSING=ON,
  REFERENCE=CUSTOMERS, $
  FIELDNAME=PHONE, ALIAS=Phone, USAGE=A24, ACTUAL=A24,
  MISSING=ON,
  REFERENCE=CUSTOMERS, $
  FIELDNAME=FX, ALIAS=Fax, USAGE=A24, ACTUAL=A24,
  MISSING=ON,
  REFERENCE=CUSTOMERS, $
SEGMENT=ORDERS, SEGTYPE=S0, PARENT=CUSTOMERS, $
  FIELDNAME=ORDERS, ALIAS=Orders, USAGE=A1, ACTUAL=A1,
ACCESS_PROPERTY=(INTERNAL), $
  REFERENCE=CUSTOMERS, $
  FIELDNAME=ORDERID, ALIAS=OrderID, USAGE=I11, ACTUAL=A11,
  MISSING=ON,
  REFERENCE=ORDERS, $
  FIELDNAME=CUSTOMERID1, ALIAS=CustomerID, USAGE=A5, ACTUAL=A5,
  MISSING=ON,
  REFERENCE=ORDERS, $
  FIELDNAME=EMPLOYEEID, ALIAS=EmployeeID, USAGE=I11, ACTUAL=A11,
  MISSING=ON,
  REFERENCE=ORDERS, $

```

```

FIELDNAME=ORDERDATE, ALIAS=OrderDate, USAGE=HYMDm, ACTUAL=A35,
MISSING=ON,
REFERENCE=ORDERS, $
FIELDNAME=REQUIREDDATE, ALIAS=RequiredDate, USAGE=HYMDm, ACTUAL=A35,
MISSING=ON,
REFERENCE=ORDERS, $
FIELDNAME=SHIPPEDDATE, ALIAS=ShippedDate, USAGE=HYMDm, ACTUAL=A35,
MISSING=ON,
REFERENCE=ORDERS, $
FIELDNAME=SHIPVIA, ALIAS=ShipVia, USAGE=I11, ACTUAL=A11,
MISSING=ON,
REFERENCE=ORDERS, $
FIELDNAME=FREIGHT, ALIAS=Freight, USAGE=P19.4, ACTUAL=A19,
MISSING=ON,
REFERENCE=ORDERS, $
FIELDNAME=SHIPNAME, ALIAS=ShipName, USAGE=A40, ACTUAL=A40,
MISSING=ON,
REFERENCE=ORDERS, $
FIELDNAME=SHIPADDRESS, ALIAS=ShipAddress, USAGE=A60, ACTUAL=A60,
MISSING=ON,
REFERENCE=ORDERS, $
FIELDNAME=SHIPCITY, ALIAS=ShipCity, USAGE=A15, ACTUAL=A15,
MISSING=ON,
REFERENCE=ORDERS, $
FIELDNAME=SHIPREGION, ALIAS=ShipRegion, USAGE=A15, ACTUAL=A15,
MISSING=ON,
REFERENCE=ORDERS, $
FIELDNAME=SHIPPOSTALCODE, ALIAS=ShipPostalCode, USAGE=A10, ACTUAL=A10,
MISSING=ON,
REFERENCE=ORDERS, $
FIELDNAME=SHIPCOUNTRY, ALIAS=ShipCountry, USAGE=A15, ACTUAL=A15,
MISSING=ON,
REFERENCE=ORDERS, $
SEGMENT=CUSTOMERDEMOGRAPHICS, SEGTYPE=S0, PARENT=CUSTOMERS, $
FIELDNAME=CUSTOMERDEMOGRAPHICS, ALIAS=CustomerDemographics,
USAGE=A1, ACTUAL=A1, ACCESS_PROPERTY=(INTERNAL),
REFERENCE=CUSTOMERS, $
FIELDNAME=CUSTOMERTYPEID, ALIAS=CustomerTypeID, USAGE=A10, ACTUAL=A10,
MISSING=ON,
REFERENCE=CUSTOMERDEMOGRAPHICS, $
FIELDNAME=CUSTOMERDESC, ALIAS=CustomerDesc, USAGE=A30, ACTUAL=A30,
MISSING=ON,
REFERENCE=CUSTOMERDEMOGRAPHICS, $

```

The corresponding Access File, customers.acx, has the attribute *ENTITY_SET=Orders* in the ORDERS segment definition.

```

SEGNAME=CUSTOMERS,
  CONNECTION=NorthWind,
  TABLENAME=Customers,
  OBJECT_URL=CUSTOMERID,
  EXTCALL=YES, $
FIELD=CUSTOMERID,
  TYPE=ID,
  IN_URL=TRUE, $
SEGNAME=ORDERS,
  OBJECT_URL=ORDERID,
  EXTCALL=NO,
  ENTITY_SET=Orders, $
FOREIGN_KEY=Orders,
  PRIMARY_KEY_TABLE=Order,
  FOREIGN_KEY_COLUMN=OrderID,
  PRIMARY_KEY_COLUMN=OrderID, $
FIELD=ORDERID,
  TYPE=ID,
  IN_URL=TRUE, $
SEGNAME=CUSTOMERDEMOGRAPHICS,
  OBJECT_URL=CUSTOMERTYPEID,
  EXTCALL=NO,
  ENTITY_SET=CustomerDemographics, $
FOREIGN_KEY=CustomerDemographics,
  PRIMARY_KEY_TABLE=CustomerDemographic,
  FOREIGN_KEY_COLUMN=CustomerTypeID,
  PRIMARY_KEY_COLUMN=CustomerTypeID, $
FIELD=CUSTOMERTYPEID,
  TYPE=ID,
  IN_URL=TRUE, $

```

Consider the following request.

```

TABLE FILE CUSTOMERS
BY CUSTOMERS.ORDER.SHIPNAME
IF CUSTOMERS.ORDER.SHIPNAME NE MISSING
END

```

When it is run with *ENTITY_SET=Orders* in the Access File, it creates the following OData request, which references only the Orders entity.

[https://services.odata.org/V4/Northwind/Northwind.svc/Orders?\\$select=OrderID,ShipName&\\$orderby=ShipName](https://services.odata.org/V4/Northwind/Northwind.svc/Orders?$select=OrderID,ShipName&$orderby=ShipName)

If it were run *without* ENTITY_SET=Orders in the Access File, it would create the following OData request, which references the CUSTOMERS segment and performs \$expand to reference the ORDERS fields.

```
https://services.odata.org/V4/Northwind/Northwind.svc/Customers?\$select=CustomerID&\$expand=Orders\(\$select=OrderID,ShipName;\$orderby=ShipName\)
```

This would read every CustomerID from the Customers Entity, even though only data from the ORDERS child segment is requested.

Adding a field from the CUSTOMERS segment, such as COMPANYNAME, would cause the OData request to go against the Customer entity and perform \$expand on the Orders entity.

```
https://services.odata.org/V4/Northwind/Northwind.svc/Customers?\$select=CustomerID,CompanyName&\$expand=Orders\(\$select=OrderID,ShipName;\$orderby=ShipName\)
```

Remote Storage Adapters

This section describes new features for the Remote Storage Adapters.

Adapter for AWS S3

The Adapter for AWS S3 (Amazon Web Services Simple Storage Service) is new in this release. It is the WebFOCUS Adapter used to integrate with AWS S3. Output formats such as PDF, JSON, and XML generated from WebFOCUS reports can be stored in a Bucket within the AWS S3 environment. Adapters such as JSON and XML can be used to report against their respective documents stored in AWS S3.

AWS S3 is in the Remote Storage Adapter category.

For information about configuring and using this adapter, see the chapter named *Using the Adapter for AWS S3* in the *Adapter Administration* manual.

Adapter for Amazon Web Services: Loading a File From AWS S3 Directly to Amazon RDS

Using DataMigrator or a Data Flow, a delimited file (.csv) stored on Amazon Web Services (AWS) S3 can now be loaded directly to Amazon Relational Database Service (RDS), stored in PostgreSQL.

In prior releases, you had to use the Adapter for Athena to read the file.

In order to load the file directly:

- ❑ The adapter for AWS S3 must be configured.

For information about configuring the Adapter for AWS S3, see the chapter named *Using the Adapter for AWS S3* in the *Adapter Administration* manual.

- ❑ The Adapter for PostgreSQL must be configured to point to the PostgreSQL database in Amazon RDS, as shown in the following image that points to the example URL `jdbc:postgresql://database-1.xxxxxxxxxxxxxx.rds.amazonaws.com:9999/postgres`.

Add PostgreSQL to Configuration [Prerequisites](#)

Connect parameters ⤴

Connection Name ?

URL ?

[Show Sample](#)

Security ?

User ?

Password ?

Advanced Parameters ⤵

Environment ⤵

Select profile ?

The "Test" option for the connection is only available after initial configuration is complete.


[Configure](#)

For information about configuring the Adapter for PostgreSQL, see the chapter named *Using the Adapter for PostgreSQL* in the *Adapter Administration* manual.

- ❑ An application directory must be added that is mapped to the AWS S3 connection.

For information about mapping an application directory to an AWS S3 connection, see the chapter named *Using the Adapter for AWS S3* in the *Adapter Administration* manual.

In the following image, aws01 is the application directory mapped to the AWS S3 connection, and the Data Flow loads a delimited file from the AWS S3 repository to PostgreSQL.



Adapter for Google Drive: Support for Google Slides

You can now create Google Slides using the Adapter for Google Drive.

The syntax is:

```
ON TABLE HOLD AS mappedapp/filename FORMAT PPTX
```

where:

mappedapp

Is a mapped Google Drive application.


filename

Is the name of the Google Slide file to be generated.


In the following example, the application directory `gdr01` is the application directory mapped to Google Drive. The following request uses the `HOLD FORMAT PPTX` command in the procedure.

```
TABLE FILE employe2
SUM CURR_SAL
BY LAST_NAME
BY DEPARTMENT
ON TABLE HOLD AS gdr01/empslide FORMAT PPTX
END
```

Sometimes, the page setup of the Google slide needs to be adjusted, as shown in the following image.



In this case, you can modify the Google slide page setup to have a different aspect ratio, as shown in the following image.



The following image shows the Google slide with the layout updated.

The screenshot shows a Google Slides presentation titled 'empslide' with a menu bar (File, Edit, View, Insert, Format, Slide, Arrange, Tools, Add-ons, Help) and a toolbar. The main content area displays a table with the following data:

LAST_NAME	DEPARTMENT	CURR_SAL
BANKING	PRODUCTION	\$29,700.00
BLAKWOOD	HR	\$21,700.00
CROSS	HR	\$27,062.00
GREENSPAN	HR	\$9,500.00
IRVING	PRODUCTION	\$26,862.00
JONES	HR	\$19,480.00
MCCOY	HR	\$15,480.00
MONTGOMERY	PRODUCTION	\$24,100.00
ROMANS	PRODUCTION	\$21,120.00
SUETH	HR	\$22,200.00
STEVEN	PRODUCTION	\$9,800.00
STEVEN	PRODUCTION	\$12,000.00

At the bottom of the slide, there is a text prompt: 'Click to add speaker notes'.

Adapter for SharePoint Drive

The Adapter for SharePoint Drive is the WebFOCUS Adapter used to integrate with SharePoint Drive.

Instructions for configuring the Adapter for SharePoint Drive are included in the chapter *Using the Adapter for SharePoint Drive* of the *Adapter Administration* manual.

Output formats such as PDF, JSON, and XML generated from WebFOCUS reports can be stored in SharePoint Drive.

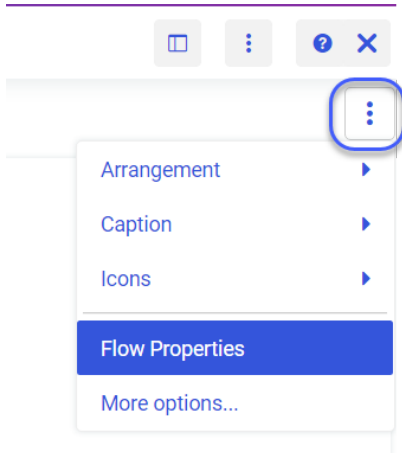
Adapters, such as JSON and XML, can be used to report against their respective documents stored in SharePoint Drive.

General Data Flow Enhancements


This section describes new features for data flows.

Viewing Flow Properties on the Data Flow Canvas

A new *Flow Properties* item has been added to the View Options menu on the Data Flow canvas, as shown in the following image.



Clicking this option opens a Flow Properties dialog box. The properties displayed depend on the type of Data Flow that is open in the canvas. For example, the following image shows a Data Flow that splits the source data into two target tables.



You can make selections, or change options, and click *Apply* to implement them. For example the following image shows how to select restart options for a flow.

The screenshot shows the 'Flow Properties' dialog box with the 'Restart' section expanded. The 'Restart from' dropdown menu is open, showing 'Beginning' and 'Last commit' options. The 'Number of attempts' is set to 2. The 'Execution' section has two checked options: 'Continue processing when a partial answer set is returned' and 'Stop if 0 rows selected'. The 'General' section shows flow name 'dflow01', creator 'edalias', and last modified date '20210309 18.47.23'. The description is 'flow with new target' and the comment is 'my first flow with WEB ETL'. 'Cancel' and 'Apply' buttons are at the bottom right.

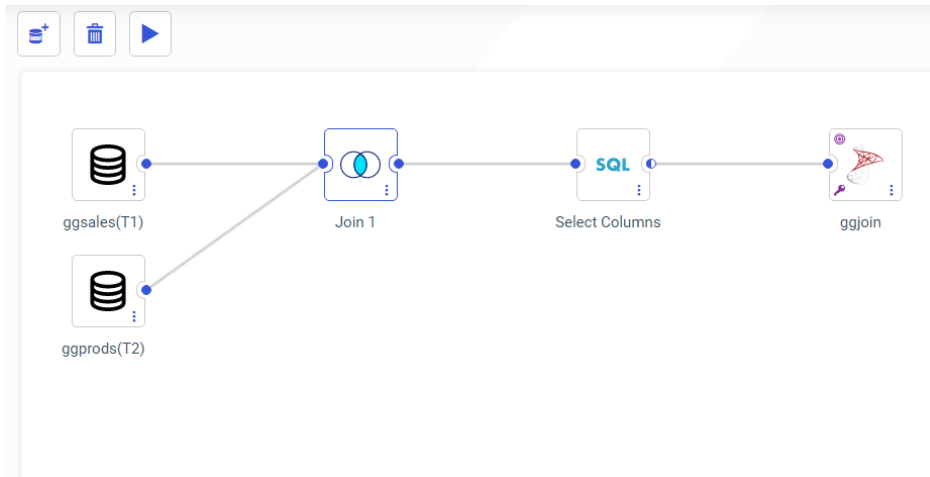
Once you click *Apply*, your settings are saved, and the last modified date and creator ID are displayed. The dialog box closes automatically.

To close without applying any changes, click the Close button (x) to close the dialog box.

Joining to a Variable in a Data Flow

You will now be prompted to enter a value, if you create a join condition that compares a field to a variable, when that variable has no default value.

For example, the Data Flow in the following image joins the ggsales and ggprods data sources.



To open the Join Editor, right-click the Join object and click *Join Editor*.

To add a custom expression, click the *+Expression* button, as shown in the following image.

Edit Join from SALES01 to PRODS01

Configure 'Join 1'

Join Type


Inner Left Outer Right Outer Full Outer

Join Clauses

ggsales (T1)		ggprods (T2)	
Product ID	=	Product,Code	

+ + Expression Suggestions Clear All

The Expression Editor opens. Enter an expression that compares a field to a new variable that has not been assigned a default value, as shown in the following image.



Click OK to return to the Join Editor. The expression is now one of the Join conditions, as shown in the following image.

Edit Join from SALES01 to PRODS01

↶
▼
↷
▼

Configure 'Join 1' ⋮

Join Type

Inner

Left Outer

Right Outer

Full Outer

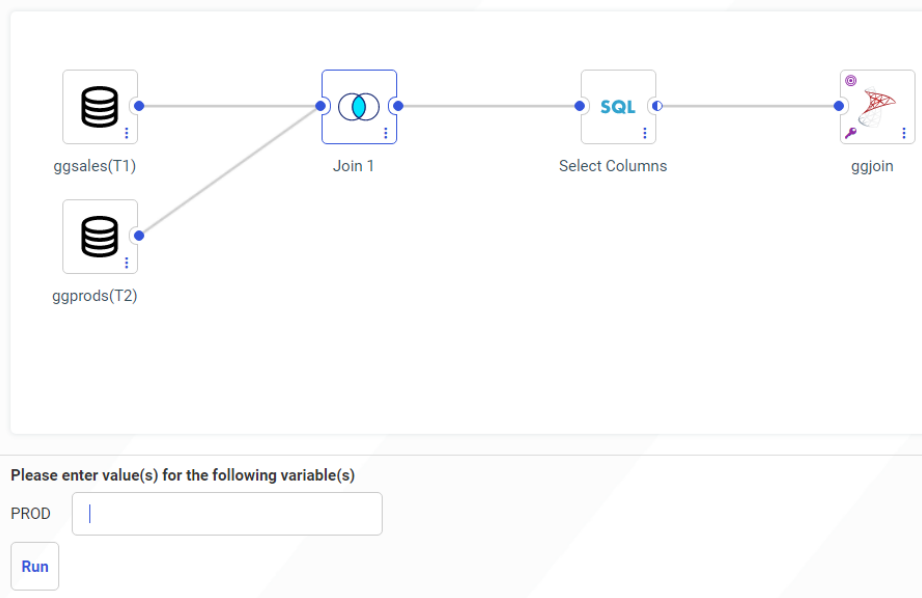
Join Clauses

ggsales (T1)		ggprods (T2)	
Product ID	=	Product,Code	✎ 🗑
T1.Product	=	'&PROD'	✎ 🗑

+
+ Expression
💡 Suggestions

🗑 Clear All

When you click *OK* and return to the Data Flow canvas, you are prompted to enter a value for the variable, as shown in the following image.



Enter a value for the variable, and click *Run*.

Viewing Checkpoint Files Generated by a Change Data Capture (CDC) Process

When you run a flow, a Checkpoint file is generated or, if it already exists, is updated. The Checkpoint file is in a proprietary internal binary format, so users cannot read it to see what happened during the processing. Now, a viewer for checkpoint files has been implemented that will show the checkpoint file contents in readable JSON format.

Note: The contents of a checkpoint file are specific to the DBMS being used.

In order for CDC to be enabled, the tables used must be log-enabled. The Data Flow is then run using the Log file as the source and Change Data Capture as the load option.

1. Create a synonym for the Table Log Records, as shown in the following image. You can create a new application folder to contain the log synonym and log file.

Note: This example uses Db2 tables on a server running on IBM i.

Create Synonym for DB2/DB2 Warehouse (CON1)

Object Type Library Supply value to avoid system-wide search, sample: S999

Table Name

Miscellaneous Parameters Customize data type mappings

Create: Cluster Synonym Cluster Join for Star Schema Base Synonym

Application Prefix Suffix

SYNONYM CANDIDATES Row Limit: 50 Search

Select	Default Synonym Name	Table Name	Library/Schema	Type
<input checked="" type="radio"/>	c1_tmix91b	C1_IMIX91B	DD999	Log Records
<input type="radio"/>	c1_tmix91c	C1_TMIX91C	DD999	Log Records
<input type="radio"/>	c1_tmix91d	C1_TMIX91D	DD999	Log Records
<input type="radio"/>	c1_tmix91e	C1_TMIX91E	DD999	Log Records
<input type="radio"/>	c1_tmix91m	C1_TMIX91M	DD999	Log Records
<input type="radio"/>	c1_trdddc1	C1_TRDDDC1	DD999	Log Records
<input type="radio"/>	c1_trdddc2	C1_TRDDDC2	DD999	Log Records
<input type="radio"/>	c1_trdxmla	C1_TRDXMLA	DD999	Log Records

2. Select the table for which you want to create a synonym. If necessary, add a prefix or suffix that will identify this as a synonym for a log table, optionally change the default synonym name, and click *Add*.

The Synonym Editor opens.

3. Make any changes that are necessary, and click the right arrow on the ribbon to save the synonym.

The synonym contains variables that control the CDC processing. You can edit them in the Synonym Editor or in the Text Editor. To open the synonym in the Text Editor, right-click the synonym on the Resources tree, point to *Metadata Management*, and click *Edit as Text*.

The synonym opens as text, as shown in the following image.

```

FILENAME=C1_TRDCDC1_LOG, SUFFIX=DB2      , $
VARIABLE NAME=&&CDC_POLLING,
  PROMPT='DBMS log polling interval (in seconds) for LUW',
  USAGE=I5, DEFAULT='1', $
VARIABLE NAME=&&CDC_TIMEOUT,
  PROMPT='Timeout interval (in seconds) to listen DBMS log for LUW',
  USAGE=I5, DEFAULT='1', $
VARIABLE NAME=&&CDC_START,
  PROMPT='Starting point in reading log records',
  DEFAULT='CHKPT', ACCEPT=
  DECODE(<'CHKPT - After last LUW retained in checkpoint file',
  'CHKPT'>, <'CUR_TRAN - First LUW after DBMS log reading started',
  'CUR_TRAN'>, <'CUR_LOG - First available LUW in DBMS log',
  'CUR_LOG'>), $
VARIABLE NAME=&&CDC_MAXLUWS,
  PROMPT='Maximum number of LUWs processed in the request',
  USAGE=I10, DEFAULT='1', $
VARIABLE NAME=&&CDC_CHKPT_SAVE,
  PROMPT='Retain last processed LUW in checkpoint file',
  DEFAULT='YES', ACCEPT=
  DECODE(<'YES - Retain last processed LUW in checkpoint file', 'YES'>,
  <'NO - Don't retain LUWs in checkpoint file', 'NO'>), $
VARIABLE NAME=&&CDC_CHKPT_FILE,
  PROMPT='Checkpoint file location', USAGE=A99, DEFAULT=' ', $
VARIABLE NAME=&&CDC_LOG_NAME,
  PROMPT='DBMS log file name', DEFAULT=' ', $
VARIABLE NAME=&&CDC_LOG_LOCATION,
  PROMPT='DBMS log file location', DEFAULT=' ', $
VARIABLE NAME=&&CDC_COMMIT_MODE,
  PROMPT='Transactions Commit Mode',
  DEFAULT=' ', ACCEPT=
  DECODE(<'ON - Transactions Committed', 'ON'>,
  <'OFF - Transactions Autocommit', 'OFF'>), $

```

```

SEGMENT=C1_TRDCDC1_LOG, SEGTYPE=S0, $
  FIELDNAME=CDC_OPER, ALIAS=__C1, USAGE=A1, ACTUAL=A1, $
  FIELDNAME=CDC_TID, ALIAS=__C2, USAGE=A32, ACTUAL=A32, $
  FIELDNAME=CDC_TIMES, ALIAS=__C3, USAGE=A26, ACTUAL=A26, $
  FIELDNAME=FA01INT, ALIAS=FA01INT, USAGE=I11, ACTUAL=I4, $
  FIELDNAME=FA02INT, ALIAS=FA02INT, USAGE=I11, ACTUAL=I4,
  MISSING=ON, $
  FIELDNAME=FA03REAL, ALIAS=FA03REAL, USAGE=D20.2, ACTUAL=D8,
  MISSING=ON, $
  FIELDNAME=FA04NUMERIC16X4, ALIAS=FA04NUMERIC16X4,
  USAGE=P19.4, ACTUAL=P9, MISSING=ON, $
  FIELDNAME=FA05NUMERIC, ALIAS=FA05NUMERIC, USAGE=P16, ACTUAL=P8,
  MISSING=ON, $
  FIELDNAME=FA06CHAR_5, ALIAS=FA06CHAR_5, USAGE=A5, ACTUAL=A5,
  MISSING=ON, $
  FIELDNAME=FA07VARCHAR_5, ALIAS=FA07VARCHAR_5, USAGE=A5V, ACTUAL=A5V,
  MISSING=ON, $
  FIELDNAME=FA08CHAR_10, ALIAS=FA08CHAR_10,
  USAGE=A10, ACTUAL=A10, MISSING=ON, $
  FIELDNAME=FA09VARCHAR_10, ALIAS=FA09VARCHAR_10,
  USAGE=A10V, ACTUAL=A10V, MISSING=ON, $
  FIELDNAME=FA10DATE, ALIAS=FA10DATE, USAGE=YYMD, ACTUAL=DATE,
  MISSING=ON, $
  FIELDNAME=FA11TSTAMP, ALIAS=FA11TSTAMP, USAGE=HYMDm, ACTUAL=HYMDm,
  MISSING=ON, $
  FIELDNAME=FA12TIME, ALIAS=FA12TIME, USAGE=HHIS, ACTUAL=HHIS,
  MISSING=ON, $

```

4. Edit the variable values, as follows:

- &&CDC_MAXLUWS.** Optionally, change this to zero (0) in order to add all LUWs to the checkpoint file.
- &&CDC_CHKPT_SAVE.** Make sure this is set to YES, in order to retain the last processed checkpoint in the checkpoint file.
- &&CDC_CHKPT_FILE.** Optionally, add the checkpoint file name. The extension must be .chp.

For example:

```

VARIABLE NAME=&&CDC_MAXLUWS,
  PROMPT='Maximum number of LUWs processed in the request',
  USAGE=I10, DEFAULT='0', $
VARIABLE NAME=&&CDC_CHKPT_SAVE,
  PROMPT='Retain last processed LUW in checkpoint file',
  DEFAULT='YES', ACCEPT=
  DECODE(<'YES - Retain last processed LUW in checkpoint file',
  'YES'>,
  <'NO - Don't retain LUWs in checkpoint file', 'NO'>), $
VARIABLE NAME=&&CDC_CHKPT_FILE,
  PROMPT='Checkpoint file location',
  USAGE=A99, DEFAULT='mycheck.chp', $

```

5. Run the flow with the log table as the source and your existing target, as shown in the following image.



Change the Load Option to Change Data Capture, as shown in the following image.

Load Options [X]

Load Option [Change Data Capture] ?

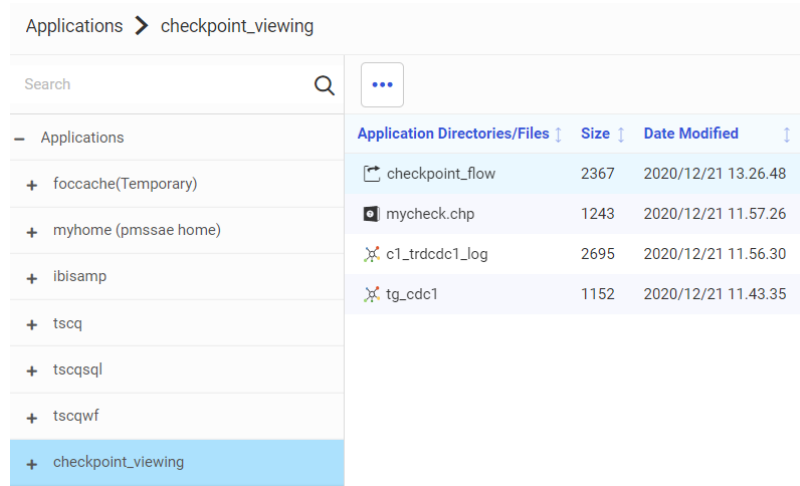
Prior to Load Option [None] ?

Select Target Synonym [checkpoint_viewing/tg_cdc1.mas] ?

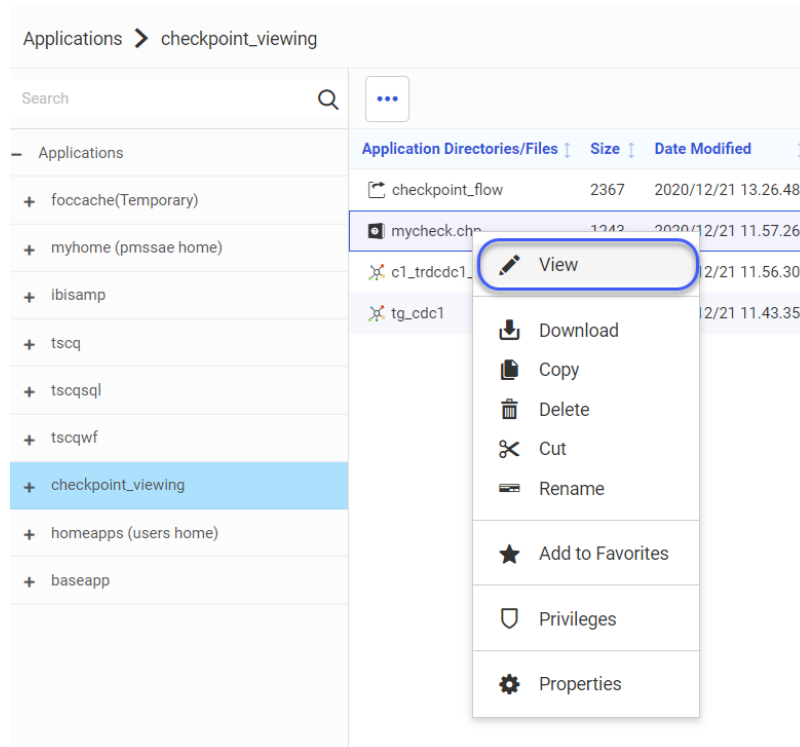
Control column [CDC_OPER] ?

[Cancel] [OK]

The log file is generated in the application folder, as shown in the following image.



- To view the checkpoint file, either double-click the file or right-click the file and click *View*, as shown in the following image.



The checkpoint file is displayed in JSON format, as shown in the following image.

```
{
  "checkpoint": [
    {
      "name": "file",
      "value": "checkpoint_viewing/mycheck.chp",
      "description": "check point file full name"
    },
    {
      "name": "engine",
      "value": "DB2",
      "description": "data access adapter"
    },
    {
      "name": "catalog",
      "value": "*LOCAL",
      "description": "catalog / database name"
    },
    {
      "name": "table",
      "value": "DD999.C1_TRDCDC1",
      "description": "table name"
    },
    {
      "name": "chkp_id",
      "value": "###00000000000000000000",
      "description": "checkpoint ID / min entry LSN"
    },
    {
      "name": "commit_id",
      "value": "QSQRN0002##DD999    ##00000000000000199545",
      "description": "commit ID/LSN"
    },
    {
      "name": "tran_id",
      "value": "00000000000000199533",
      "description": "current transaction ID/LSN"
    },
    {
      "name": "timestamp",
      "value": "2020-12-21 06:27:58.630656000",
      "description": "transaction timestamp"
    }
  ]
}
```

Migrating Data Flows

When you create a Data Flow, an SQL request is generated that retrieves data from Source databases, applies transformations, if necessary, and loads the data to Target databases.

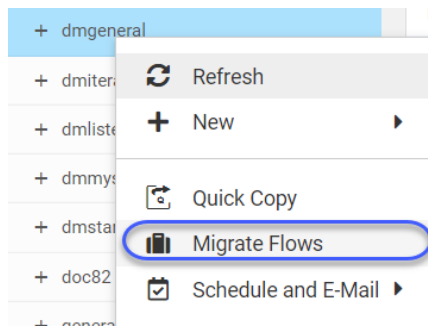
In new releases of DataMigrator, improvements are made, when possible, to the code that is generated for the Flows. The new code uses new back-end algorithms for reading and loading the data, allowing an improved performance.

If a Data Flow created in a prior release no longer runs, or seems to run too slowly, you can migrate the flow to the current release. This process will update the code for the generated Flow to the latest syntax, and take advantage of performance and optimization improvements available in the current release of DataMigrator.

Procedure: How to Migrate Data Flows

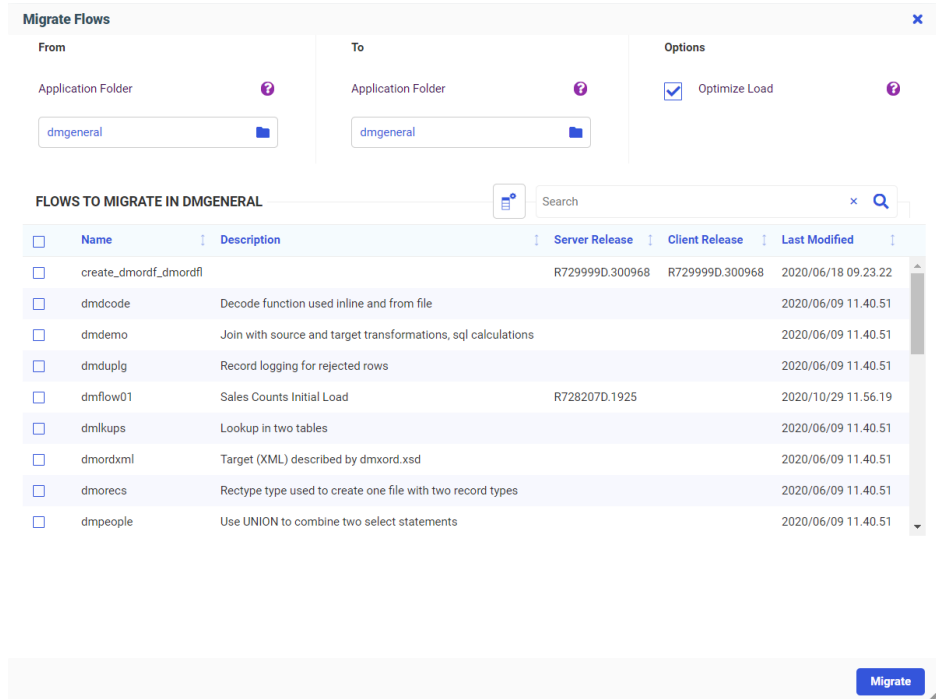
You can migrate flows in both the Web Console and the Data Management Console (DMC).

1. Right-click an application directory, and click *Migrate Flows*, as shown in the following image.

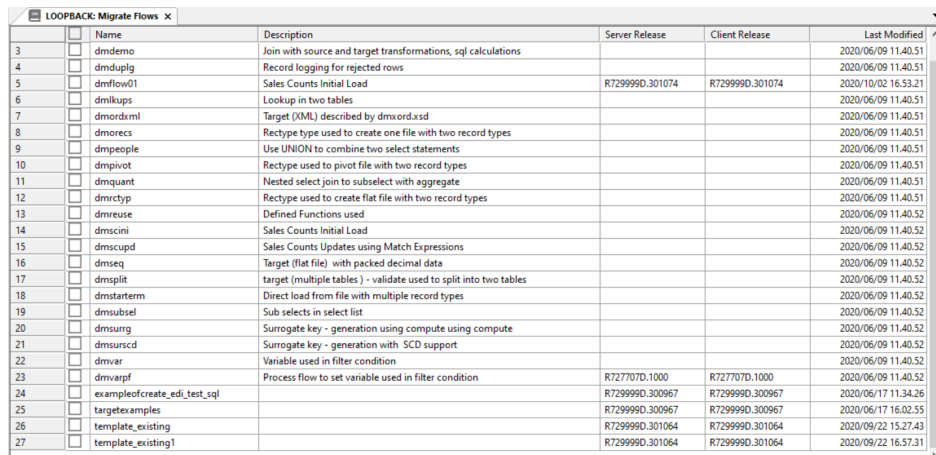


Note: In the DMC, if you do not see the menu item *Migrate Flows* on the Application Directory context menu, you can enable it from the Preferences pane. Right-click the server node name and select *Preferences* from the context menu to open the Preferences pane. Expand the Technical Preview Features section. Select Yes from the *Enable Migrate Flows* drop-down list.

In the Web Console, the *Migrate Flows* window opens, as shown in the following image.



In the DMC, the Migrate Flows panel opens in a new tab, as shown in the following image.



2. If you are using the Web Console, the Migrate Flows window enables you to change the application directories for the migration process and to deselect the Optimize Load option.

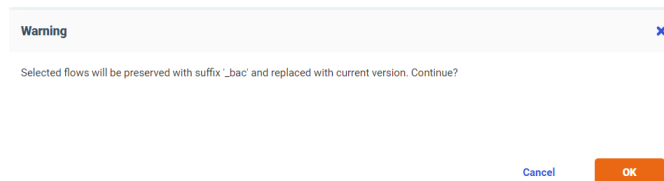
Initially, both the *From* and *To* application folders are set to the application directory from which you initiated *Migrate Flows*.

- Click in the *Application Folder* entry field to open the *Select Application Folder* dialog box to change one or both of the folders.
- By default, *Optimize Load* is selected. If you want to disable this feature, deselect the check box.

Note: The DMC version of *Migrate Flows* does not support changing the application folders or deselecting *Optimize Load*. In the DMC, *Optimize Load* will be used whenever it is supported for the Target DBMS.

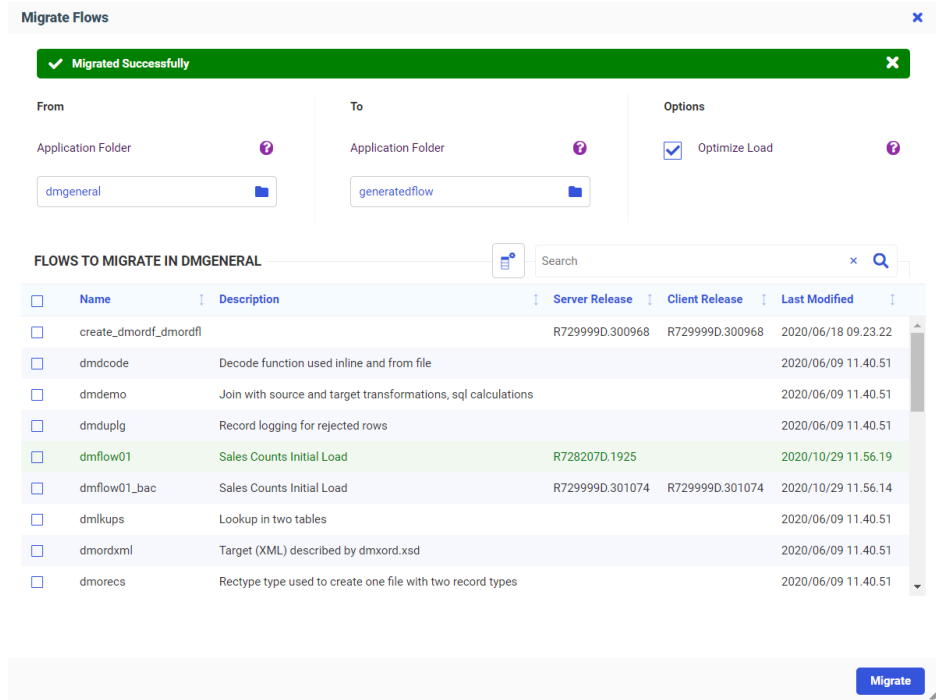
3. Select the check boxes next to the flows you want to migrate, and click *Migrate*.

If the *From* and *To* folders are the same, the original versions of the selected flows will be saved with the suffix `_bac` appended to the flow name. The migrated flows will have the original names. The following message displays to inform you of this.



If you receive this message, either click *OK* to migrate the flows and save the original versions with names that have the suffix `_bac`, or click *Cancel* to cancel the migration.


In the Web Console, the Migrate Flows window has either a green banner and success message at the top, if the flows were migrated successfully, or a red banner with error messages, if they were not. A successful migration is shown in the following image.



The flows on the list that were selected for migration are also highlighted.

4. Close the Migrate Flows window by clicking the X at the top right, if you are using the Web Console.

In the DMC, the original flows (with the _bac suffix) and migrated flows (with the original names) will appear on the panel, as shown in the following image.




In the Web Console:

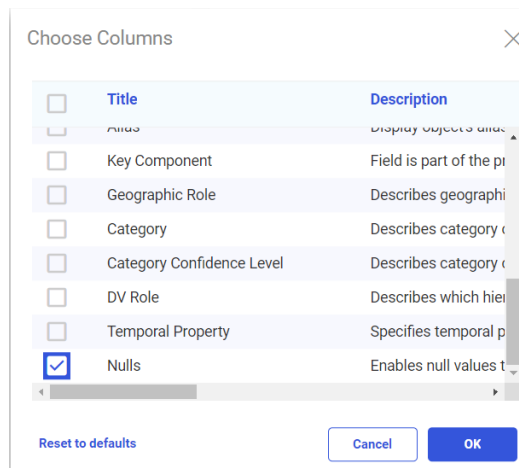
- if the *From* and *To* application folders were the same, the original flows (with the _bac suffix) and migrated flows will appear in the application folder.

- ❑ If the *From* and *To* folders were different, the original flows will be in the *From* folder, with their original names, and the migrated flows will be generated in the *To* folder, with the same names.

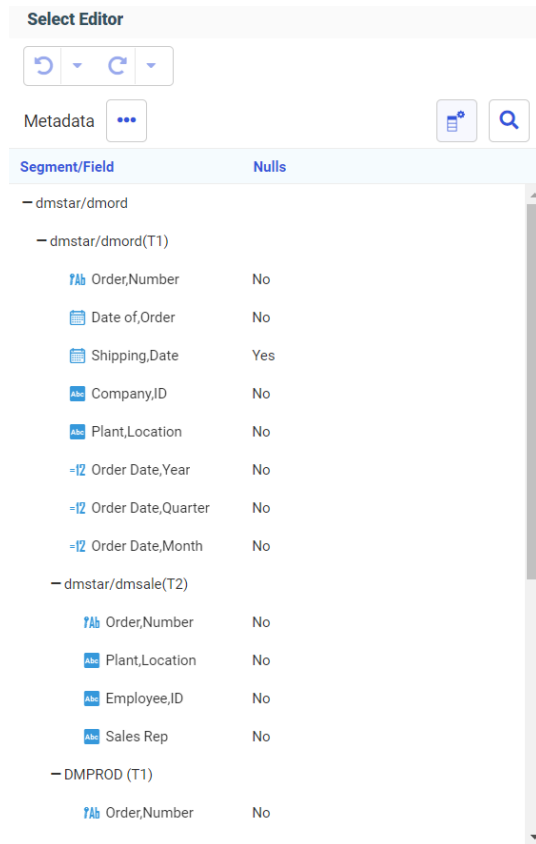
Showing Whether a Column Supports Nulls in the Select, Merge, and Synonym Editors

The Choose Columns dialog box in the Select and Merge Editors in a Data Flow, and in the Synonym Editor, now enables you to see whether each column listed in the editor supports null values.

Click the *Choose Columns* icon () and select the column named *Nulls*, as shown in the following image.



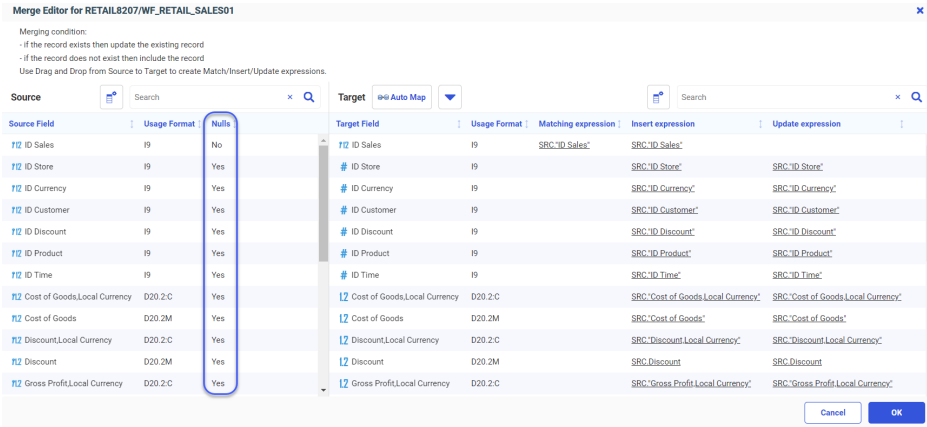
After you click *OK*, you can see whether each field supports null values, as shown in the following image for the Select Editor.



The screenshot shows the 'Select Editor' interface. At the top, there are navigation buttons (refresh, back, forward) and a 'Metadata' section with a search icon. Below this is a table with two columns: 'Segment/Field' and 'Nulls'. The table lists various segments and their fields, along with whether they support null values.

Segment/Field	Nulls
- dmstar/dmord	
- dmstar/dmord(T1)	
Order,Number	No
Date of,Order	No
Shipping,Date	Yes
Company,ID	No
Plant,Location	No
Order Date,Year	No
Order Date,Quarter	No
Order Date,Month	No
- dmstar/dmsale(T2)	
Order,Number	No
Plant,Location	No
Employee,ID	No
Sales Rep	No
- DMPROD (T1)	
Order,Number	No

The following image shows the Nulls column selected in the Merge Editor.



Controlling Format Conversion Error Processing in a Data Flow

A format conversion error occurs when a field described with a numeric USAGE format contains character values. In prior releases, format conversion errors in a Data Flow were displayed, but not counted or logged. Records with errors were discarded.

Two parameters have been added that control format conversion error processing in a Data Flow.

- CNVERR, which instructs the data engine to replace values that cause data truncation or conversion errors with the default value (null if the column is nullable, zero otherwise).

The default value, REJECT, rejects the record. The value IGNORE replaces the values causing errors with default values, and continues processing the record.

- CNVERRLIMIT, which sets a limit, up to 999999, for the number of records with warnings about data conversion errors.

The default value (0) allows unlimited warning messages. When the limit is reached, processing terminated and the system variable FOCERRNUM is set to 1346.

The following are the warning messages counted for this limit.

```
(FOC1277) RECORD %1. FIELD %2 HAS FORMAT ERROR, VALUE%3: %4
```

```
(FOC1130) FORMAT CONVERSION ERROR FIELD/KEY %1%2%3%4
```

An error occurred attempting to convert the specified field from its ACTUAL to USAGE format. Either the ACTUAL format is described incorrectly or the data may be invalid. Check the MFD.

```
(FOC1346) %1%2%3%4
```

This message provides the record number and position in this record when data conversion interruption occurs.

Note: These settings apply when the source is a delimited or Excel file.

These parameters are available in the Web Console and in the Source Properties for a Data Flow in the Data Management Console (DMC).

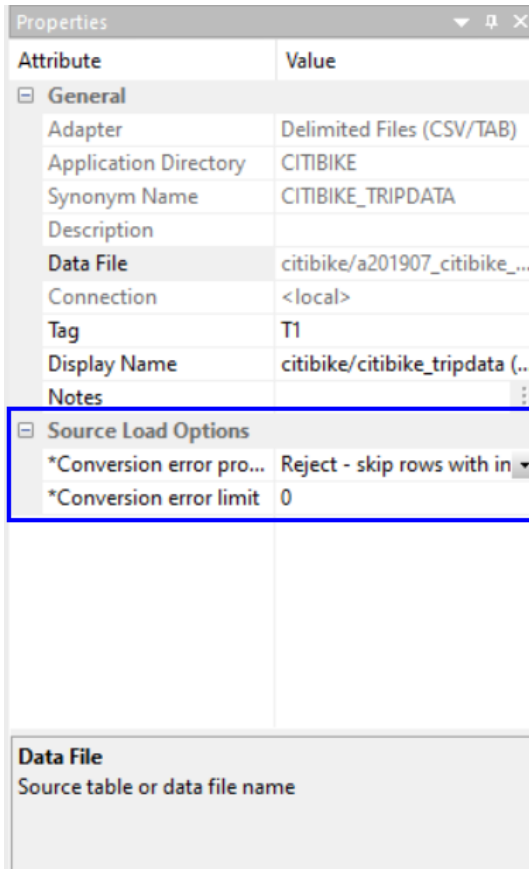
- ❑ In the Web Console, navigate to the Workspace page. Click *Settings*, point to *FOCUS Sets and Info*, and click *Core Engine Settings*.

Select *Data Utility Service*, and click *Next*. The Data Utility Service page opens, as shown in the following image.

Parameter	Value	Description
ALPHACLEAN ?	OFF	Conversion behavior for external ACTUAL= low value hexadecimal 0.
NODATA ?	.	Character string display to indicate missing data in a report.
OVERFLOWCHAR ?		Character to display during numeric overflow conditions, * (asterisk) by default.
EXTENDNUM ?	AUTO	Extend numeric usages to avoid format overflow conditions.
DISPLAYROUND ?	ON	To avoid .499 problem, add ULP (unit of least precision) to display doubles
NEG-ZERO ?	OFF	Display negative values that round to zero as minus zero.
VMSFLOATCONV ?	IEEE	OpenVMS Float conversion from G_FLOAT to IEEE for external sources.
VMS_TIMESTAMP ?	INTERNAL	OpenVMS Timestamp conversion between VMS format and internal format for external sources.
CNVERR ?	REJECT	Data conversion / truncation error processing.
CNVERRLIMIT ?	0	Limit number of records with warnings about data conversion error, 0 - unlimited.
PHONETIC_ALGORITHM ?	METAPHONE	Phonetic algorithm for indexing of strings in function PHONETIC.
ALLOWCVTERR ?	OFF	Date conversion behavior for external ACTUAL= to USAGE= data errors.

- ❑ In the DMC, right-click a delimited, fixed-format, or Excel source, and click *Properties*.

The Properties panel opens, as shown in the following image.



Reference: CNVERR: Controlling Processing of Records With Format Conversion Errors in a Data Flow

```
SET CNVERR = {REJECT|IGNORE}
```

where:

REJECT

Does not process rows with incorrect data, and issues a warning message. This is the default value.

IGNORE

Accepts rows with incorrect data, and substitutes default values (null if the column is nullable, zero otherwise).

Reference: CNVERRLIMIT: Setting a Limit for the Number of Format Conversion Errors in a Data Flow

```
SET CNVERRLIMIT = {0|nnnnnn}
```

where:

nnnnnn

Is an integer value up to 999999. When the limit is reached, processing is terminated and the system variable FOCERRNUM is set to 1346. The default value, zero (0), allows unlimited format conversion errors.

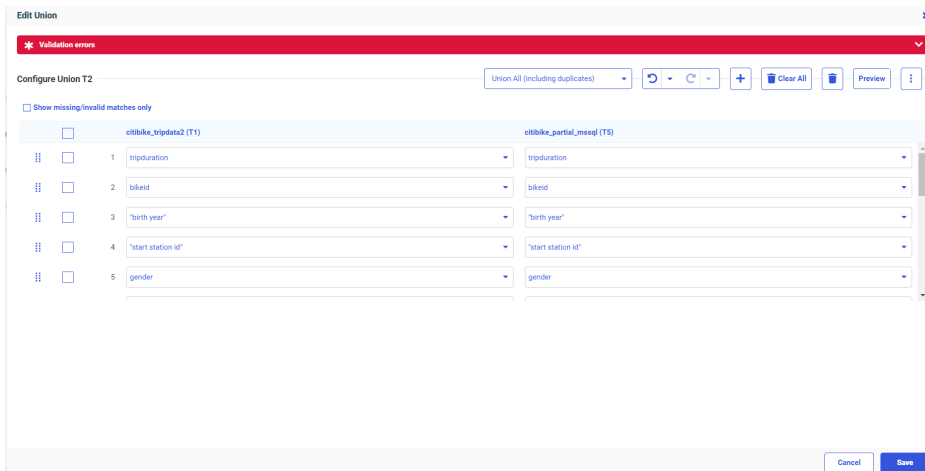
Enhancements to the Union Editor

The union editor has been redesigned to show clearly what needs to be corrected, if there are validation problems.

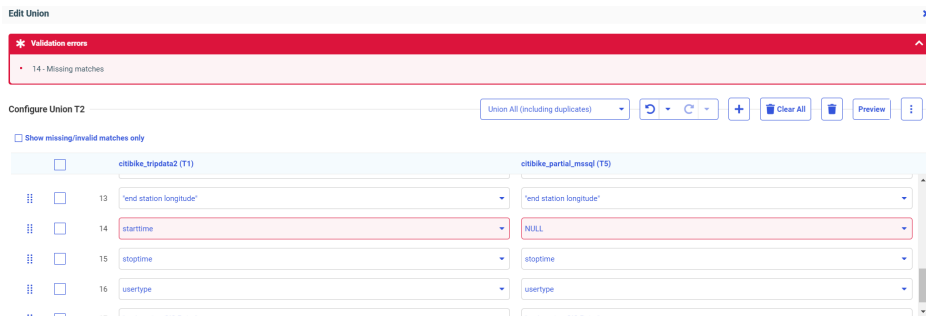
Reference: Correcting Validation Errors and Editing the Matches

The Union Editor has two areas. The top area shows the matches between the fields in the two tables. The bottom area contains a preview of the data in the union. When you first open the Union Editor, the Preview area is hidden.

If any validation errors exist, a message appears at the top of the window, as shown in the following image.



You can click the down arrow in the message area to see a list of problems. In the following image, line 14 has a missing match and, when you scroll down to that line, it is outlined in red.



You can also select the check box *Show missing/invalid matches only*, to display only those lines with missing or invalid matches.

To correct a missing match, you can:

- Click the down arrow next to the right table and select a different field that matches the field from the left table.
- Click the down arrow next to the left table and select a different field that matches the field from the right table.
- Select the check box next to the line and click the *Delete selected columns* button (trash can) at the top of the page.

You can select multiple columns and delete them all at once.

- Click the *Clear All* button (trash can with the text *Clear All*) at the top of the page to remove all columns.

You will then need to add pairs manually using the *Add a pair (+)* button.

- Click the *Add a pair (+)* button to add a matching pair.


Clicking this button adds a line with NULL selected for each table. Select fields from the left and right tables using the drop-down lists, to create a matching pair.

- Change the type of union generated, if that is appropriate. For information, see [Selecting the Type of Union](#) on page 130.

When all validation problems are resolved, the message area disappears.

Reference: Selecting the Type of Union

By default, Union All, including duplicates, is selected. However, you can change the type of union that is generated by selecting an option from the drop-down list at the top of the page, as shown in the following image.








A union combines the output of two SELECT statements into one answer set, where the number of columns in each is the same, and the data types are compatible. You can select the following types of union:


- Union (excluding duplicates).** Returns the rows from both SELECT statements with duplicate rows removed.
- Union All (including duplicates).** Returns all rows from both SELECT statements, including duplicates. This is the default union type.
- Except (excluding duplicates).** Returns distinct rows from the left SELECT statement that are not returned by the right SELECT statement.
- Except All (including duplicates).** Returns all rows from the left SELECT statement that are not returned by the right SELECT statement.

- Intersect (excluding duplicates).** Returns matching rows from the left SELECT statement and the right SELECT statement based on matching all columns. Duplicate rows are removed.
- Intersect All (including duplicates).** Returns matching rows from the left SELECT statement and the right SELECT statement based on matching all columns. Duplicate rows are retained.

Reference: Editing the Union Properties and View

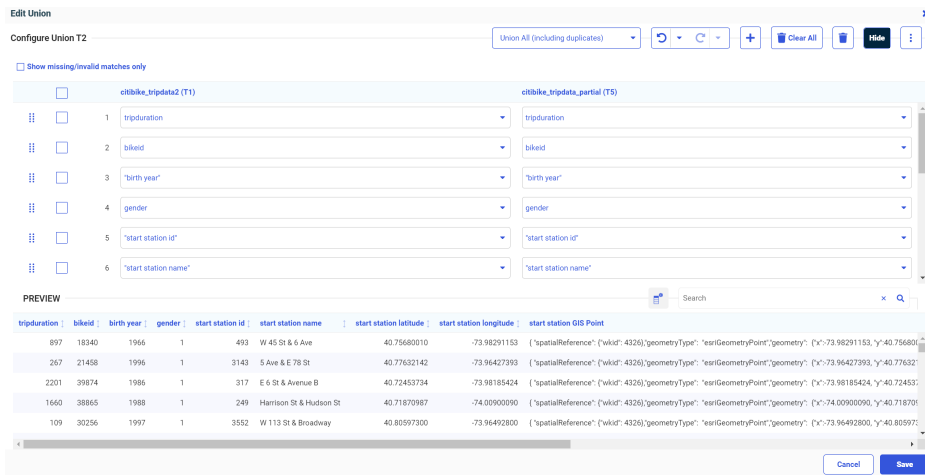
The following table list changes you can make to the union properties and view.

Operation	Icon	Description
Move columns		Click to drag a pair up or down on the list.
Undo		Undoes the most recent change made to the union properties. You can click the down arrow next to the button to see a list of the changes that have been made, and select from the list. Selecting an item undoes that change and all prior changes.
Redo		Redoes the most recent change made to the union properties. You can click the down arrow next to the button to see a list of the changes that have been made and, select from the list. Selecting an item redoes that change and all prior changes.
View		Provides the following options. <ul style="list-style-type: none"> <input type="checkbox"/> Show title. Displays the field title. This is the default view. <input type="checkbox"/> Show Name. Displays the field name. <input type="checkbox"/> Show source name. Displays the name from the source file. <input type="checkbox"/> Diagnostics. Has two options, Session Log, which opens the Session Log, and Clear Session Log, which clears the Session Log.
Choose Columns		In the Preview area, enables you to select the columns that display.

Operation	Icon	Description
Sort		In the Preview area, for each column, changes the sort order of the rows based on the data in that column. The sort toggles between the original order, descending, and ascending.

Reference: Viewing Sample Data

Click the *Preview* button at the top of the page to display sample data for the union, as shown in the following image.



The *Preview* button changes to a *Hide* button, which you can click to hide the sample data.

You can enter characters in the *Search* box to display lines with data values that match those characters.

You can use the *Choose Columns* button to select which columns display in the preview.

You can use the *Sort* button in a column to sort the rows based on that column.


Reference: Saving the Union

When you have resolved all validation errors and made any edits you need, click *Save*.

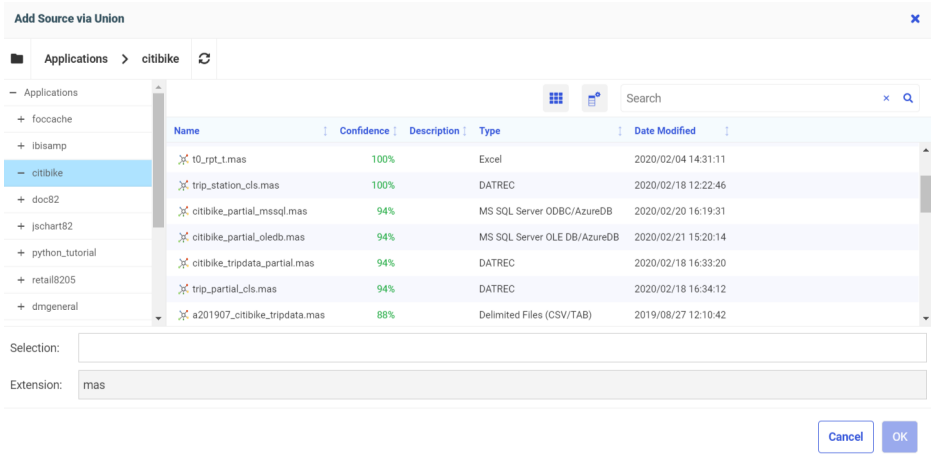
Click *Cancel* if you do not want to save your changes.

Showing the Confidence Level for Adding a Union to a Data Flow

You can right-click a source in a data flow and add a new source as a union, as shown in the following image.



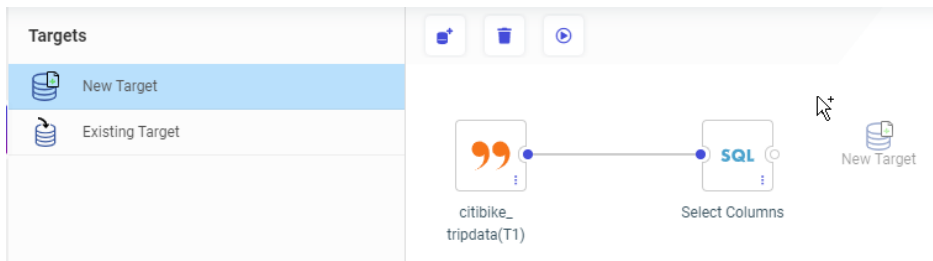
A list of tables displays for choosing the new source to add. This table now includes a confidence level, based on matching column names and data types for each prospective new source, showing how compatible it is with the existing source, as shown in the following image.



Support for Dragging Targets Into a Data Flow

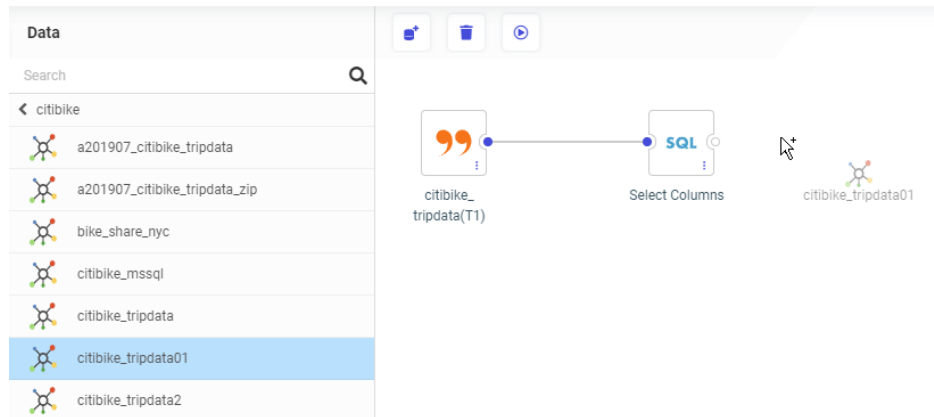
You can now drag new or existing targets onto the data flow canvas.

For a new target, click *Targets* and drag the New Target icon onto the Data Flow canvas, as shown in the following image.



Once you drop the icon onto the data flow canvas, it will be configured as the default target type with the default load options. You can change these attributes by right-clicking the target and clicking *Load Options*.

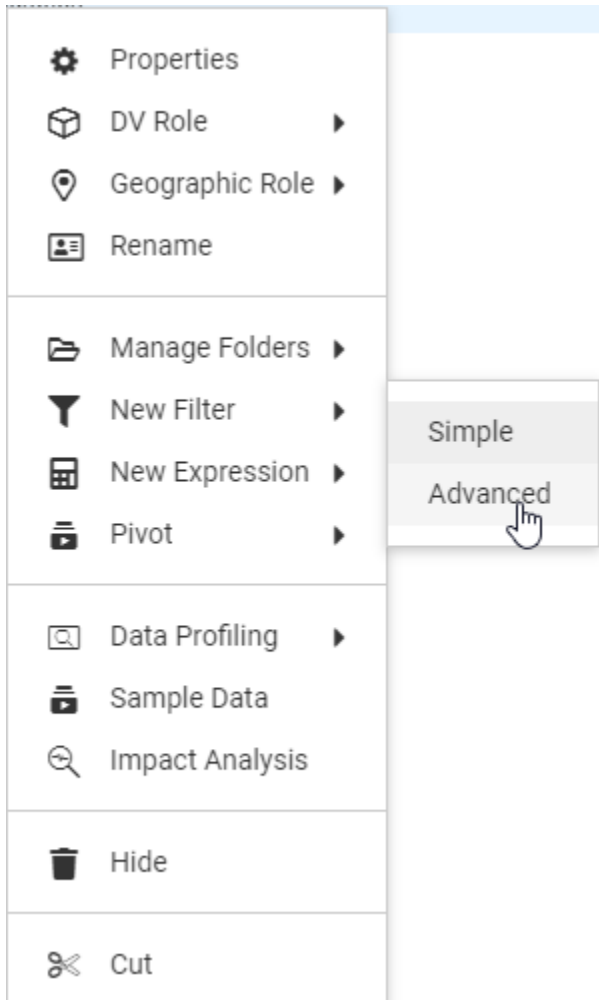
There are now two ways of dragging an existing target. You can drag the *Existing Target* icon from the Targets palette onto the canvas, which then opens a dialog box where you can select the target. You can also drag a synonym from the Resources pane to the right side of the SQL object, as shown in the following image.



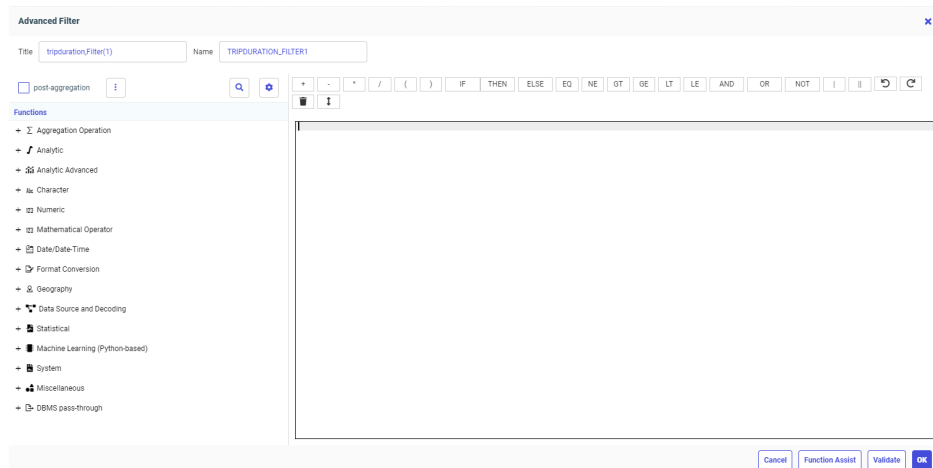
Once you drop the existing target onto the data flow canvas, it will be represented by the database-specific icon and will have the load options specified when the target was first created, with the default merge option. You can edit the options for merging the data into the target by right-clicking the target and clicking *Load Options*. You can also edit the merge expressions and conditions by right-clicking the target and clicking *Merge Editor*.

Creating Advanced Filters


You can now create a complex filter in a data flow or the Synonym Editor using the expression calculator. To create an advanced filter, right-click the field you want to filter, click *New Filter*, and click *Advanced*, as shown in the following image.



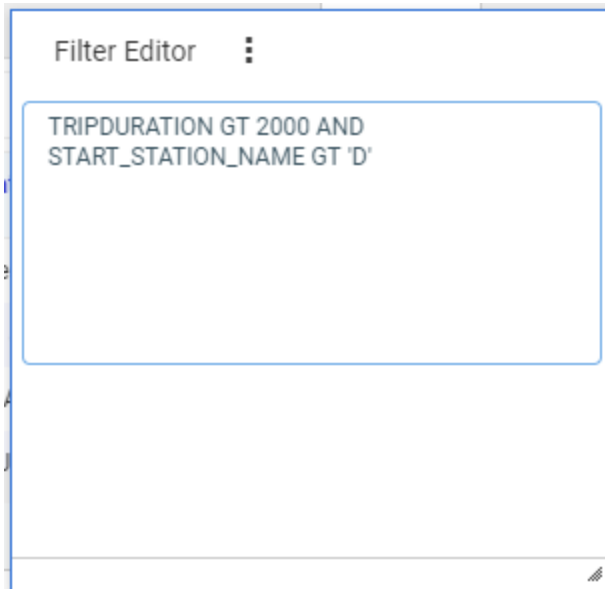
The Advanced Filter Expression Calculator opens. Unlike the calculators that open for DEFINE and COMPUTE, there is no Format entry field, as shown in the following image.



When you edit a filter, you also have the choice of *Simple* or *Advanced*, as shown in the following image.



If you edit an advanced filter but click *Simple* instead of *Advanced*, the filter opens in read-only mode in the simple filter tool, as shown in the following image:



Distinguishing Node Status Using Ports and Connectors


The Data Flow and Join nodes now have ports that indicate if they can have inputs (on the left) or outputs (on the right). The ports also indicate if all required inputs and outputs are configured.

- An empty node that is not yet considered missing is represented by a hollow circle.
- A node that is present is represented by a solid blue circle.
- A required node that is missing is represented by a circle half-filled in red and outlined in red.
- A node with a menu has an ellipsis (...) at the bottom right. You can click the ellipsis or right-click anywhere in the node to open the menu.
- The indicator that there are warning messages is now an exclamation point in a red circle, which you can hover over to display the warning messages.


In the following image of a new flow, a source is required, but not present, indicated by the half-filled red circle on the left of the SQL node. No target has been added, so the circle on the right of the SQL node is hollow.




The cross join node has warning messages, identified by an exclamation point in a red circle, which you can hover over to display the warning messages, as shown in the following image.




In the following image, a source has been added, so the circles between the source and the SQL node are solid blue. No target has been added, so the circle on the right of the SQL node is hollow.



In addition, targets are now represented by database-specific icons, as shown in the following image, which has a MySQL target.

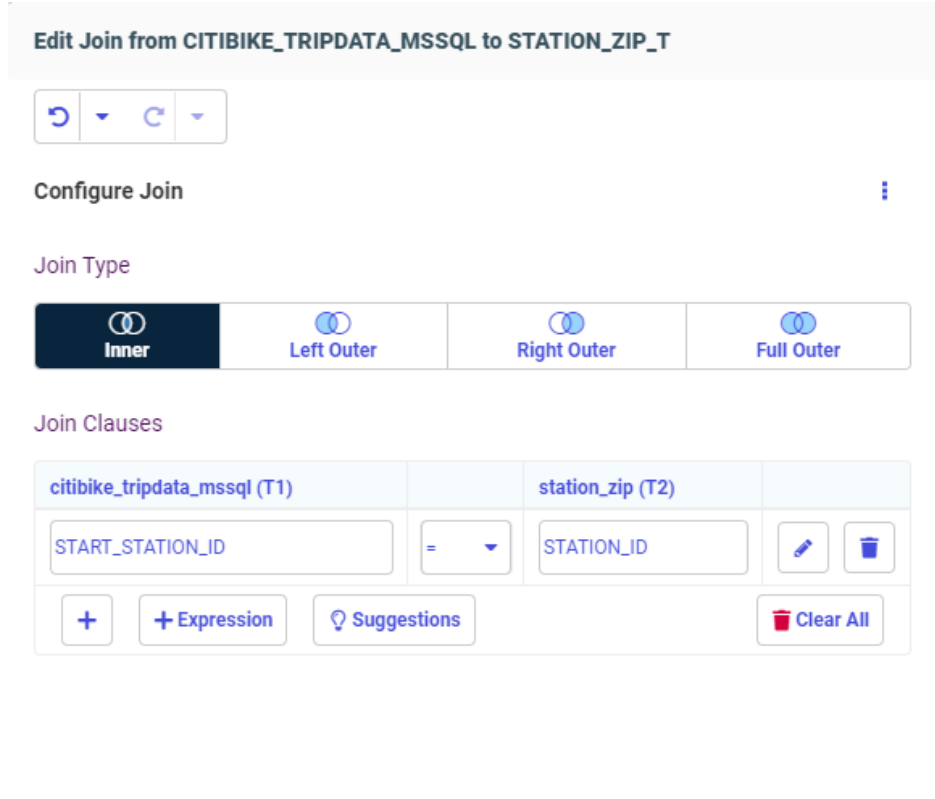


If a new target is added, a plus sign displays in the target node. If the target has one or more keys, a key icon displays in the target node, as shown in the following image of a new target with a key.

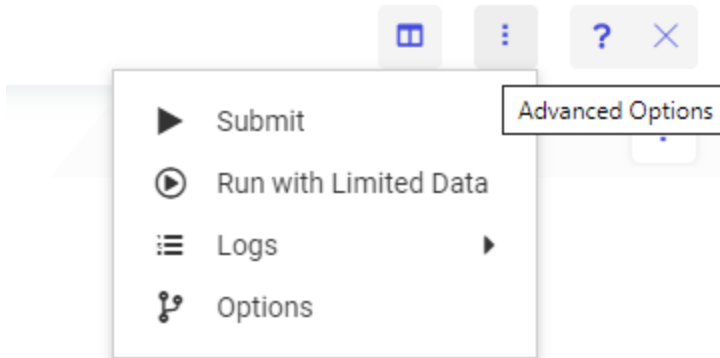


Using Undo and Redo in Join, Union, and Business View Editors

The Union, Join and Business View editors now have Undo and Redo buttons, as shown in the following image.



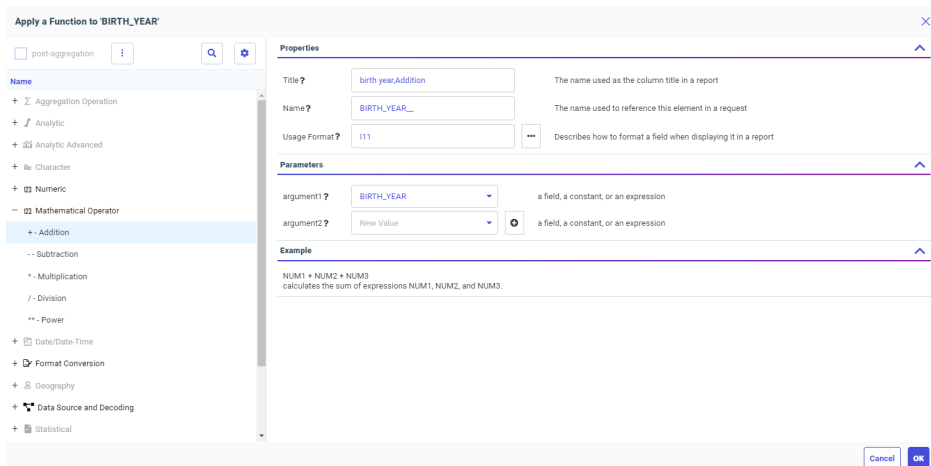
By default, 50 undo/redo operations are supported. You can change this limit by clicking the *Advanced Options* menu and selecting *Options*, as shown in the following image.



This opens the Advanced Options dialog box. In the Limits section, you can change the number in the Undo/Redo Limit entry field and click *OK*.

Function Assist For Mathematical Operators

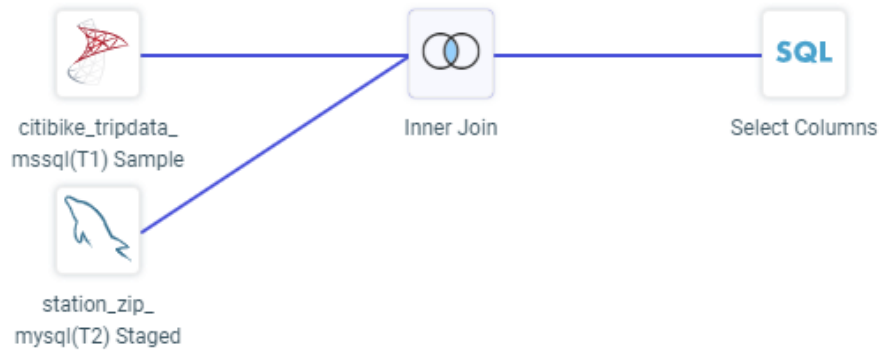
Creating a New Expression to perform a mathematical operation (add, subtract, multiply, divide, or raise to a power) on two numbers previously required a creating an Advanced Expression. Now you can select *Apply Function*. There is a new Mathematical Operator folder from which you can select these operators and use Function Assist to supply values, as shown in the following image.



Displaying Database-Specific Icons for Sources in a Data Flow

When you create a data flow, the data sources are represented by database-specific icons, so you can easily identify the types of sources involved in the flow.

For example, the following image shows a flow using a SQL Server source and a MySQL source.



A generic icon is now only used when a source synonym references multiple tables with disparate sources.

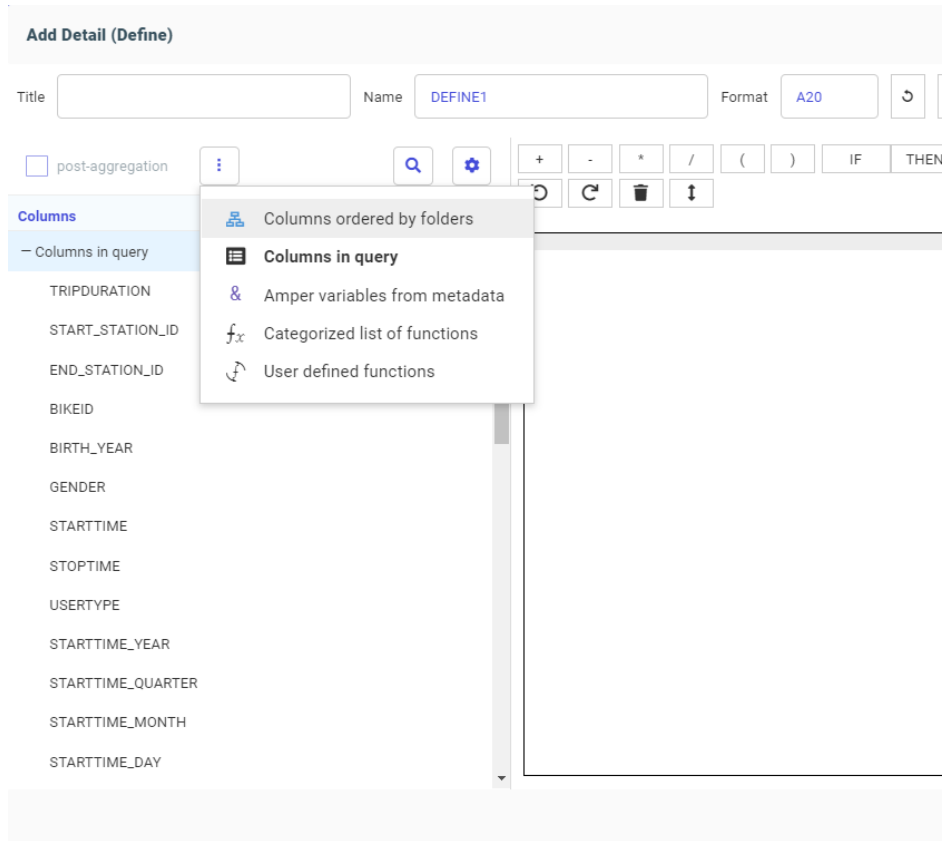
Displaying Database-Specific Icons for Data Sources in the Join Editor

In the Join Editor using the Flow view, data sources are represented by the icons that represent those data sources, as shown in the following image that joins Microsoft SQL Server, MySQL, and delimited data sources.



Expression Editor Shows Object Types

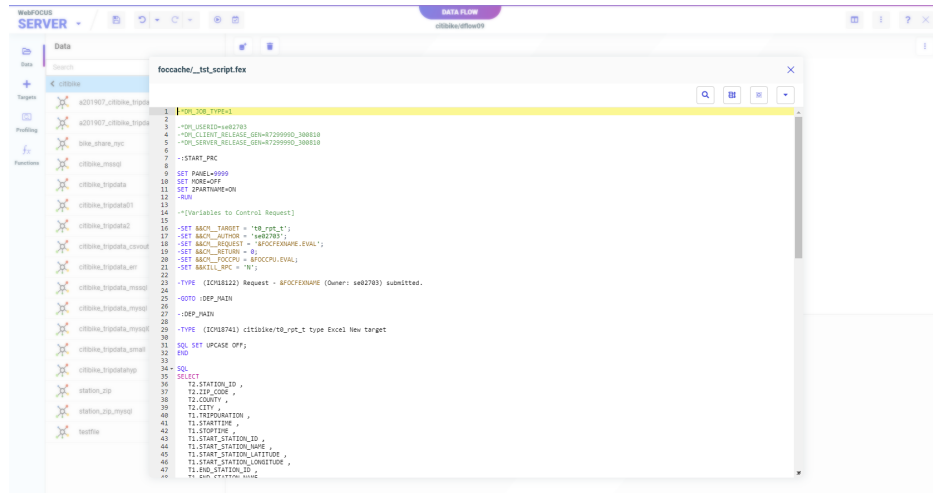
The Expression Editor provides lists of functions, columns, or variables that can be used in editing an expression. You can select which one to use from an options menu, as shown in the following image.



The selected option is highlighted on the options list and is now displayed at the top of the left pane of the Expression Editor.

View Source Opens in Full Screen

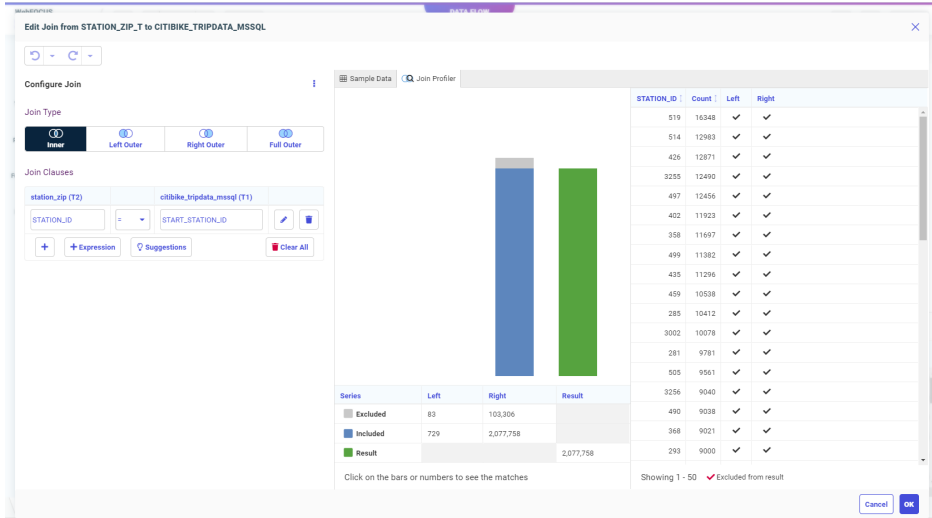
When you select View Source from the Data tab in Designer or when editing a Data Flow (by clicking the *Advanced Options* menu, clicking *Diagnostics*, and selecting *View Source*), the viewer now opens in a nearly full screen pop-up instead of a window at the bottom of the screen, as shown in the following image.



This makes the source easier to read.

Editing Joins in Full Screen

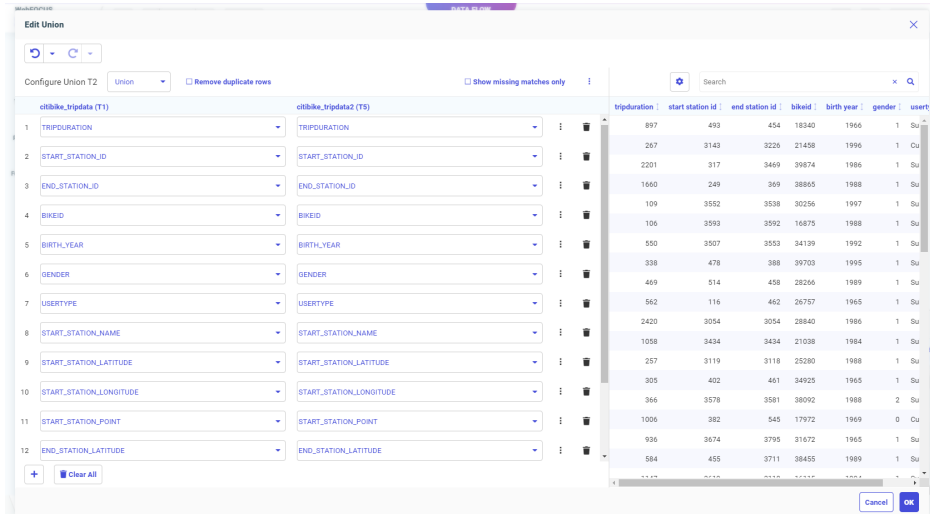
When you edit a join from the Designer Data tab or a Data Flow, the Join Configuration tool now opens in an almost full screen window, as shown in the following image.



This provides more room to see the join fields, profile chart, and values report.

Union, Join, and Select Editors Open as Nearly Full-Screen Windows

In prior releases, the Join Editor, Union Editor, and Select Editor opened as panes on the Data Flow Workspace. Now, they open as popup windows on top of the Data Flow Workspace and are almost full-screen size, as shown in the following image.



This provides more room to see and edit the on-screen elements than when they were confined to a small pane on the bottom of the screen.

Limiting the Number of Automatic Join Pairs in Tables Without Keys

When you join files or tables without keys, the join created automatically based on matching field names is now limited to three pairs of join fields, to avoid attempting a join on an excessive number of fields, especially when joining a table to itself. For example, the following image shows the default automatic join pairs for two tables with the same fields.

Edit Join from CITIBIKE_TRIPDATA to CITIBIKE_TRIPDATA2

↶ ▾ ↷ ▾

Configure Join ⋮

Join Type

⊖
Inner

⊕
Left Outer

⊕
Right Outer

⊕
Full Outer

Join Clauses

citibike_tripdata (T1)		citibike_tripdata2 (T2)	
TRIPDURATION	= ▾	TRIPDURATION	✎ 🗑
START_STATION_ID	= ▾	START_STATION_ID	✎ 🗑
END_STATION_ID	= ▾	END_STATION_ID	✎ 🗑

+
+ Expression
💡 Suggestions

🗑 Clear All

You can always add additional pairs or remove pairs, as needed.

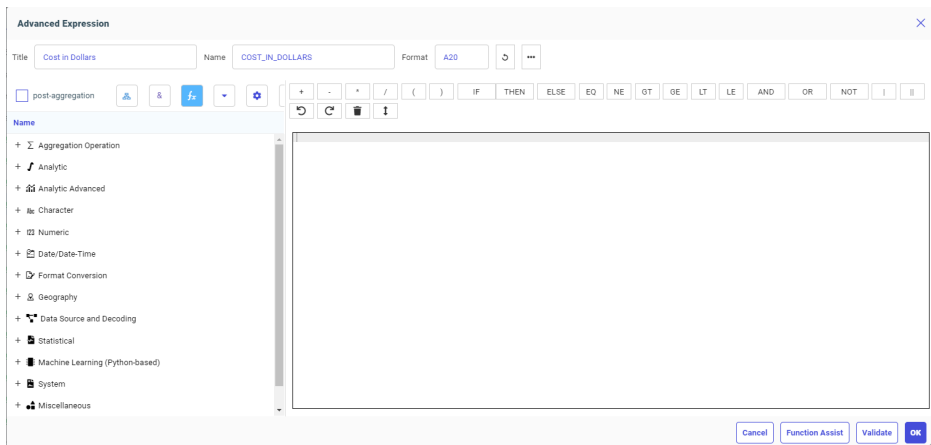
This limit does not affect joins on tables with keys, where all of the key fields are joined.

Automatically Deriving a Field Name When Creating an Expression

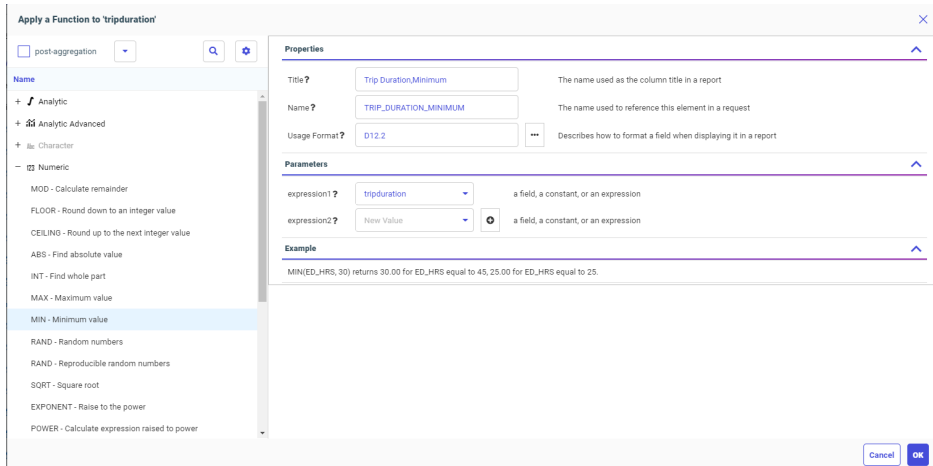
When you add a new field as an Advanced Expression or using Apply Function and enter a Title for the field, a suggested field name is automatically generated when you click outside the field or press the Tab or Enter key on the keyboard.

Titles have fewer restrictions than field names and can include special characters. Suggested field names are created with capital letters and special characters converted to underscores.

For example, if you enter Cost in Dollars as the title, the suggested field name becomes COST_IN_DOLLARS, as shown in the following image.

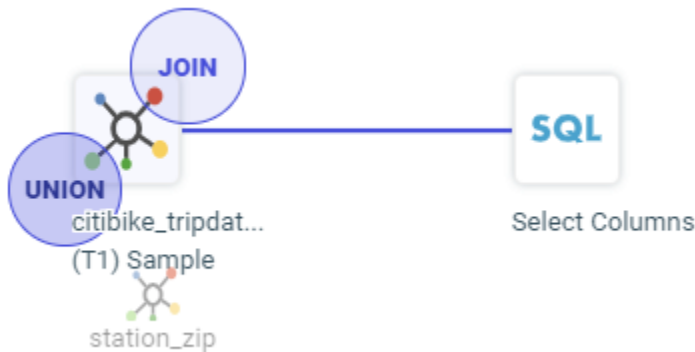


In the Apply Function dialog box, entering the Title Trip Duration,Minutes and clicking outside the field generates the field name TRIP_DURATION_MINUTES, as shown in the following image.



Selecting Join or Union When Dragging a Source Onto a Flow

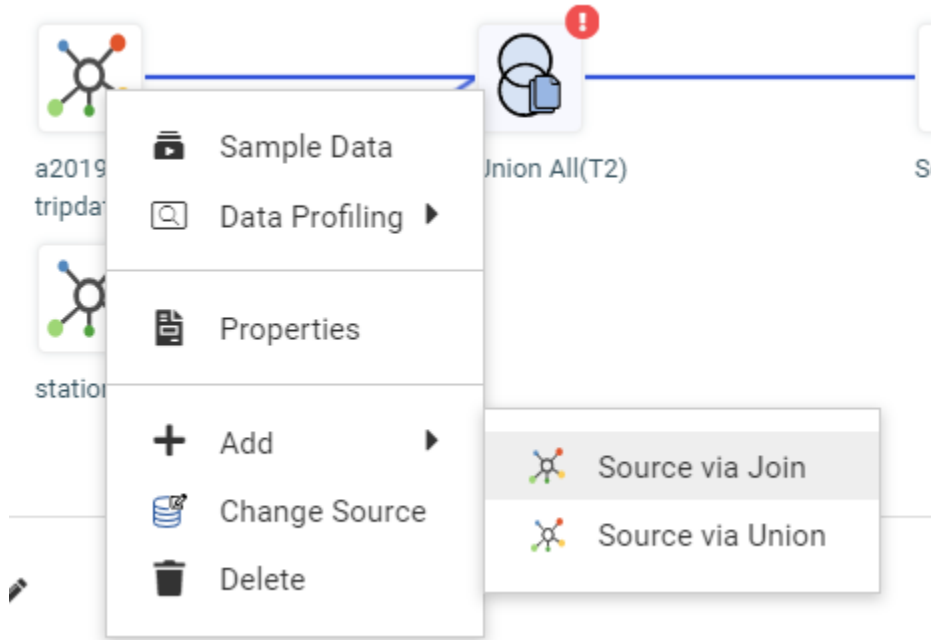
In prior releases, in order to add a source as a union, you had to first add a union to the flow, and then add sources to it. Now, when you add an additional source to a data flow by dragging the new source onto an existing one, you can add the new source as a JOIN or a UNION by dropping it on the appropriate bubble, as shown in the following image.



Choosing a Join or Union for a Data Flow From a Menu

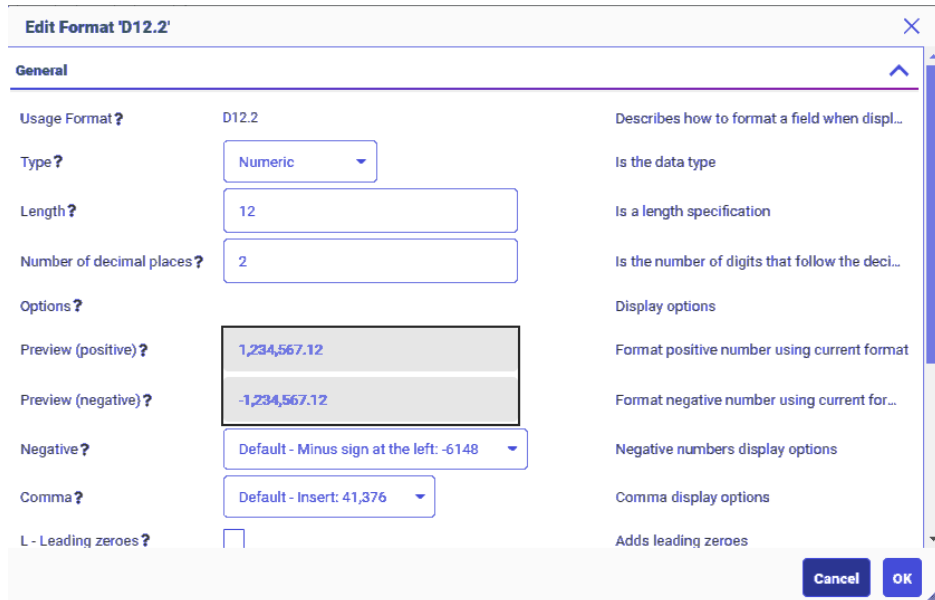
You can add a data source to a data flow as a JOIN or a UNION to existing data sources.

In addition to dragging a new source onto an existing source and selecting Join or Union, you can right-click an existing source, click *Add* from the context menu, then *Source via Join* or *Source via Union*, as shown in the following image.



Previewing the Display Format for a Function

When you create an expression using a function and change the display format of the function, you now see an example of how the field will be displayed in reports, as shown in the following image.

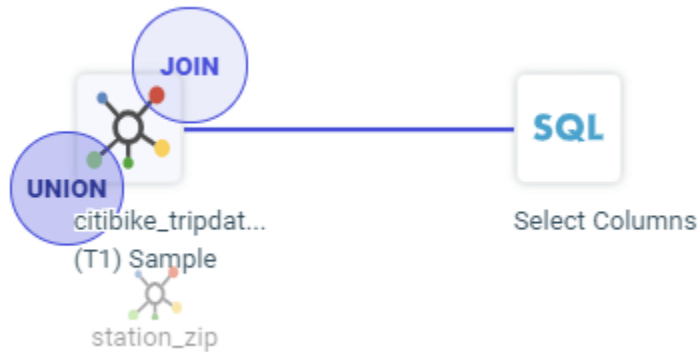


Improving Performance of UNION and UNION ALL in a Data Flow

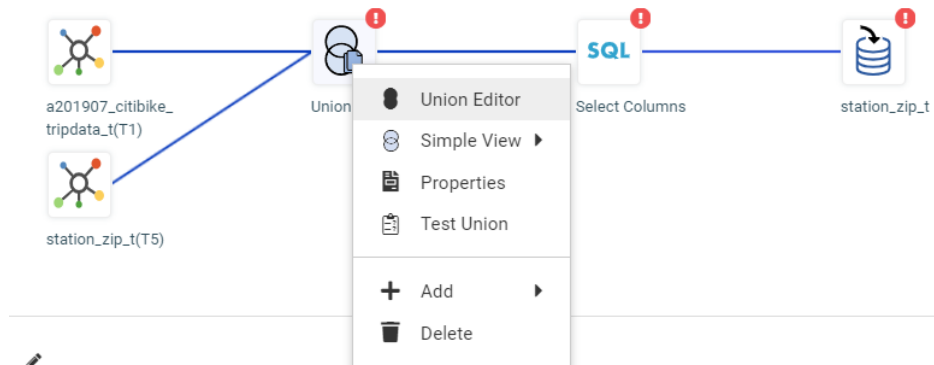
The performance of the UNION and UNION ALL processing in a data flow has been improved by passing more of the union processing to the DBMS in the generated SELECT statements. This reduces the number of intermediate files generated that need to be processed by WebFOCUS after data retrieval.

Selecting the Type of UNION in a Data Flow

When adding a source to a data flow that already has one or more sources, you can select a union by dragging the source to the UNION bubble, as shown in the following image.

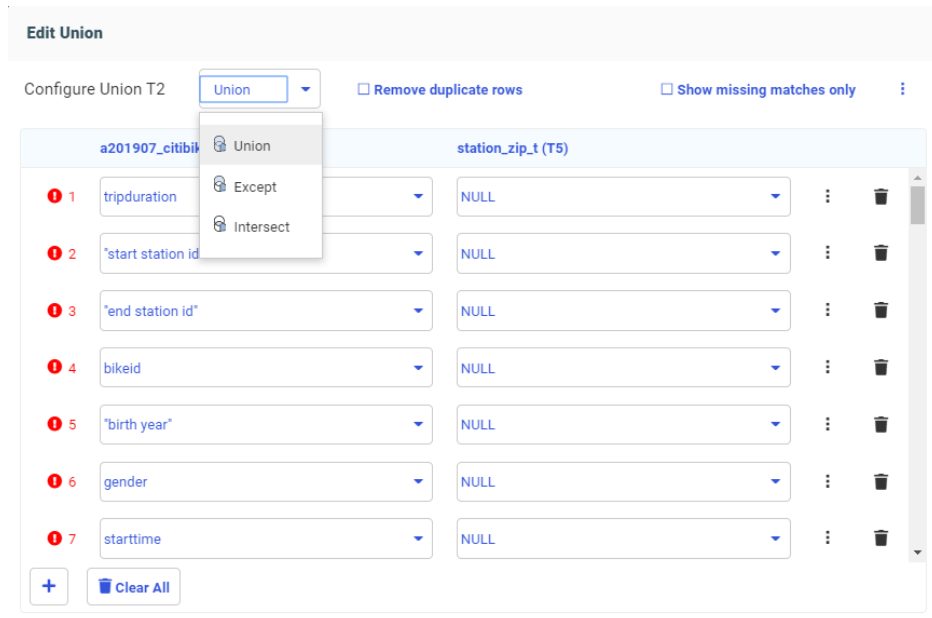


By default, UNION ALL is selected. However, you can open the Union Editor to change the type of union by right-clicking the union and clicking *Union Editor*, as shown in the following image.



Generate Flows Enhancements

The Edit Union panel opens and you can select the type of Union to perform (UNION, INTERSECT, or EXCEPT) from the Configure Union pull-down menu, as shown in the following image.



The following types of unions can be configured.

- Union.** Combines the output of two SELECT statements into one answer set, where the number of columns in each is the same, and the data types are compatible.
- Except.** Returns distinct rows from the left SELECT statement that are not returned by the right SELECT statement.
- Intersect.** Returns matching rows from the left SELECT statement and the right SELECT statement based on matching all columns.

By default UNION ALL is used unless you check *Remove duplicate rows*.

Generate Flows Enhancements

This section describes new features for Generate Flows.

Generate Flows Automatically Completes Target Names

On the Generate Flows screen, you can enter prefix and suffix characters to be removed from the source name when creating the target name, as well as prefix and suffix characters to add to the target name. When you enter one or more prefix and suffix values and move the focus away from the entered values, the target names are automatically updated on the synonym selection panel at the bottom of the screen.

On the following Generate Flows screen, if the prefix *dm* is found in the source synonym name, it is to be removed from the target synonym and table name. The suffix *_tgt* is to be added to all target synonym and table names. The target names appear with these adjustments, as shown in the following image.

Generate Flows based on foccache/genflow/template01

Configure flow | **Template information** | Schedule information

Source

Application Folder: ibisamp

Prefix: dm

New target

Application Folder: ibisamp

Prefix: dm

Suffix: _tgt

Apply Prefix and Suffix To: Target Name and Target Table Name

Flow

Save Location: generatedflow

Process Flow Name: template01_g

Parallel execution

AVAILABLE SOURCES

Source Name	Target Name	Target Table Name
database_of_company	database_of_company_tgt	companydatabase/_search_tgt
dmcomp	comp_tgt	comp_tgt
dmhr	hr_tgt	hr_tgt
dmhrinp	hrinp_tgt	hrinp_tgt
dminv	inv_tgt	inv_tgt
dminv_log	inv_log_tgt	marina.dbo.dminv_tgt
dminva	inva_tgt	inv_tgt
dminvd	invd_tgt	invd_tgt
dminvf	invf_tgt	inv_tgt

Cancel | **Generate & Save**

Generate Flows Saves the Last Used Selections

When you use a template flow to generate new flows, on the Generate Flows screen you can change the default application directory locations, specify source and target prefixes and suffixes, change the Process Flow name, select parallel execution, and select the objects for which to apply the prefixes and suffixes. Once you click *Generate and Save* or *Generate and Submit*, these selections are saved and appear the next time you start the Generate Flows process using the same template flow.

The following image shows the default Generate Flows screen for the template flow named template01.

Generate Flows based on foccache/genflow/template01

Configure flow | Template information | Schedule information

Source

Application Folder:

Prefix: Suffix:

New target

Application Folder:

Prefix: Suffix:

Apply Prefix and Suffix To:

Flow

Save Location:

Process Flow Name:

Parallel execution

AVAILABLE SOURCES

Source Name	Target Name	Target Table Name
<input type="checkbox"/> citibike_tripdata01	citibike_tripdata01	citibike_tripdata01

Cancel | Generate & Save

The following image shows the changes made to the application directories, source prefix field, and target suffix field prior to generating the flows.

Generate Flows based on foccache/genflow/template01

Configure flow | **Template information** | Schedule information

Source

Application Folder:

Prefix: Suffix:

New target

Application Folder:

Prefix: Suffix:

Apply Prefix and Suffix To:

Flow

Save Location:

Process Flow Name:

Parallel execution

AVAILABLE SOURCES

Source Name	Target Name	Target Table Name
<input checked="" type="checkbox"/> dmhr	hr_tgt	hr_tgt
<input type="checkbox"/> dmhrinp	hrinp_tgt	hrinp_tgt
<input checked="" type="checkbox"/> dminv	inv_tgt	inv_tgt
<input type="checkbox"/> dminv_log	inv_log_tgt	marina.dbo.dminv_tgt
<input type="checkbox"/> dminva	inva_tgt	inv_tgt
<input type="checkbox"/> dminvd	invd_tgt	invd_tgt
<input type="checkbox"/> dminvf	invf_tgt	inv_tgt
<input checked="" type="checkbox"/> dmord	ord_tgt	oinf_tgt

Cancel Generate & Save

Note that if you change the Process Flow name, the *Parallel execution* selection, or the selection in the *Apply Prefix and Suffix To* drop-down list, those selections will also be saved.

Once you generate the flows using either *Generate and Save* or *Generate and Submit*, the selections are saved. If you then start the Generate Flows process using the same template flow, your prior selections appear on the screen, as shown in the following image.

Generate Flows based on foccache/genflow/template01

Configure flow | Template information | Schedule information

Source

Application Folder: ibisamp

Prefix: dm | Suffix:

New target

Application Folder: ibisamp

Prefix: | Suffix: .tgt

Apply Prefix and Suffix To: Target Name and Target Table Name

Flow

Save Location: generatedflow

Process Flow Name: template01_g

Parallel execution

AVAILABLE SOURCES

Source Name	Target Name	Target Table Name
database_of_company	database_of_company_tgt	companydatabase/_search_tgt
dmcomp	comp_tgt	comp_tgt
dmhr	hr_tgt	hr_tgt
dmhrinp	hrinp_tgt	hrinp_tgt
dminv	inv_tgt	inv_tgt
dminv_log	inv_log_tgt	marina.dbo.dminv_tgt
dminva	inva_tgt	inv_tgt
dminvd	invd_tgt	invd_tgt
dminvf	invf_tgt	inv_tgt

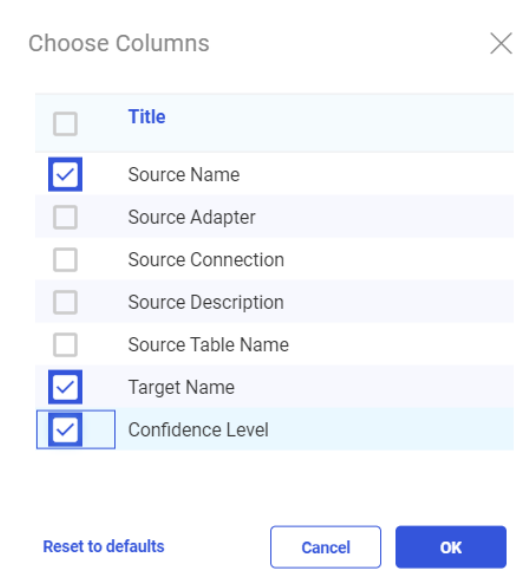
Cancel | Generate & Save

Generate Flows Shows the Confidence Level for Existing Targets

Generate Flows uses a template flow to generate multiple Data Flows and a Process Flow that will contain all of the generated Data Flows. In prior releases, only a new target was supported with Generate Flows.

Now, an existing target is supported with Generate Flows, and the row for each existing target has a confidence level associated with it. The confidence level is expressed as a percentage and identifies the compatibility between the source and target based on matching column names and data types.

By default, the Confidence Level column does not display on the Generate Flows screen. To show it, click the Choose Columns icon, select *Confidence Level*, and click *OK*, as shown in the following image.



The following image shows the Generate Flows screen with the Confidence Level column added.

Generate Flows based on foccache/genflow/template_existing

Configure flow | Template information | Schedule information

Source

Application Folder: dmgeneral

Prefix: Suffix:

Existing target

Application Folder: ibisamp

Prefix: Suffix:

Flow

Save Location: generatedflow

Process Flow Name: template_existing_g

Parallel execution

AVAILABLE SOURCES

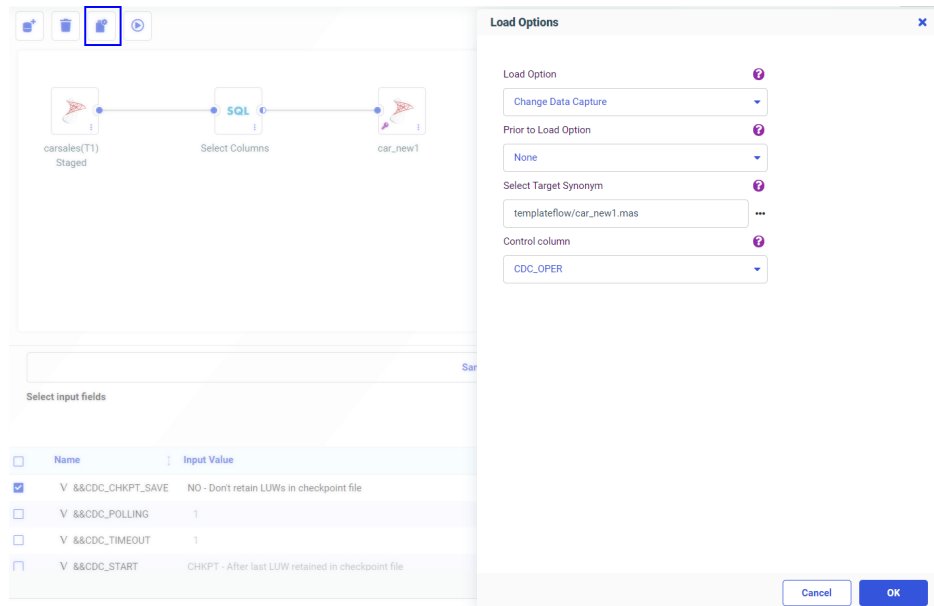
Source Name	Target Name	Confidence Level
dmduplg_dupl	dmduplg_dupl	-
dmhr	dmhr	100 %
dmhrinp	dmhrinp	100 %
dminv	dminv	100 %
dminva	dminva	85 %
dminvd	dminvd	100 %
dmlocs	dmlocs	-
dmord	dmord	100 %
dmordf	dmordf	87 %

Cancel | Generate & Save

Generate Flows Supports Template Flows With CDC Sources

Generate Flows can be used with a template flow that uses a Source Log synonym (Change Data Capture source) and one target that has Load Option *Change Data Capture*.

The following image shows a Data Flow with a source log synonym named carsales and a target with Load Option Change Data Capture.




Note that this flow can be used as a template flow, as the *Generate Flows* button is available. You can also right-click the flow in the application tree and select *Generate Flows* from the context menu.

Generate Flows Supports Template Flows With an SCD Target

Generate Flows can be used with a template Data Flow that has one source and one target, where the target synonym has Slowly Changing Dimensions (SCDs), and in which the Load Option is *Slowly Changing Dimensions*.

The following image shows a Data Flow with one source and one target. The target synonym (dimproduct_scd) has SCD fields defined, and the Load Type is Slowly Changing Dimensions.



Note that this flow can be used as a template flow, as the *Generate Flows* button is available. You can also right-click the Data Flow in the application tree and click *Generate Flows*.

Enhancements to Generating Flows Using a Template Flow

Generate Flows automatically generates a Process Flow and multiple Data Flows based on a template flow that has one source and one new target. This feature enables you to easily move any number of single sources to new targets.

The *Generate Flows* process has been enhanced to support prefixes and suffixes for target table names, as well as prefixes and suffixes for target synonym names, to have an advanced verification dialog, and to provide the *Generate Flows* button on the flow canvas for a flow that conforms to the requirements of a valid template flow.

In addition, there are now two options for generating flows:

- Generate & Save**, which saves the generated Process Flow and Data Flows. Once they are saved, you can run or submit them manually. This option gives you the ability to edit generated flows prior to submitting them.
- Generate & Submit**, which automatically submits the generated Process Flow, but does not save the Process Flow or Data Flows, which are generated in the temporary foccache application and deleted when the session ends. Only the target synonyms and tables are saved.

Procedure: How to Generate Data Flows

The template data flow and the generated data flows can each have one source and one new target. In addition, WHERE filters are not supported in the SQL statement in the template Data Flow.

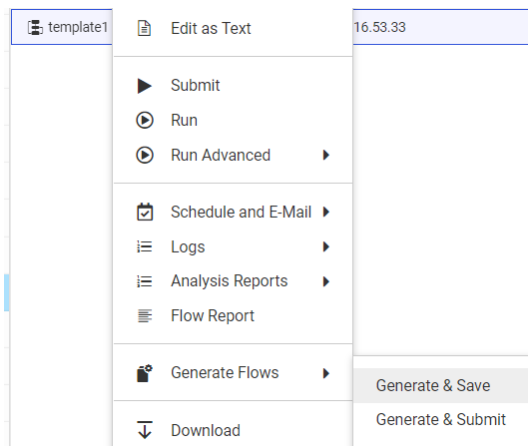
1. Create and save a data flow that has one source and one new target to use as the template.

The template flow is a normal data flow that can have any name and be saved in any application directory.

For example, the following image shows a flow named template1 that has a delimited source and new SQL Server target.




2. Right-click the template flow and point to *Generate Flows*, as shown in the following image.



Note: The *Generate Flows* option appears as long as the flow conforms to the requirements for a template flow. That is, it can have one source and one new target, and no WHERE filter.

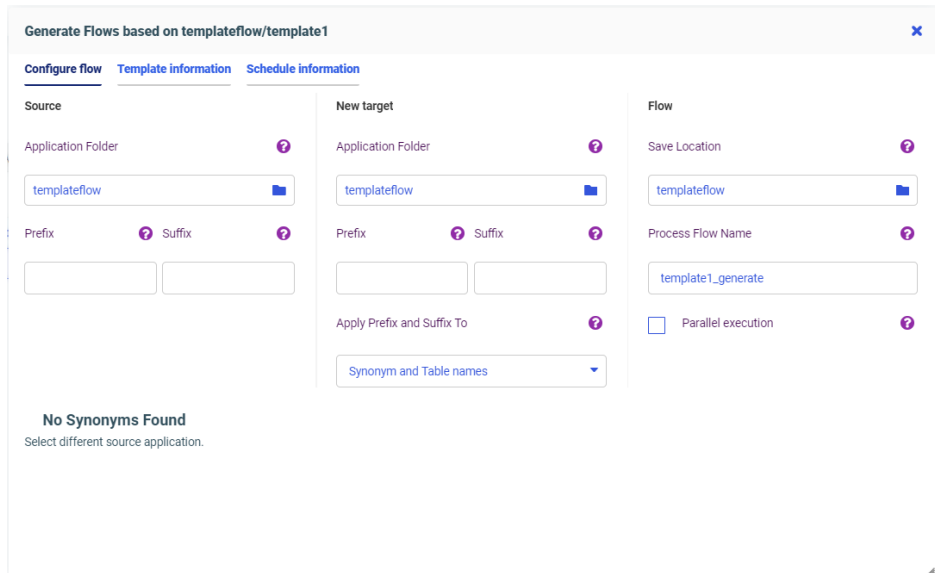
When a Data Flow that conforms to the requirements for a template flow is open in the flow canvas, there is a *Generate Flows* button at the top of the window, as shown in the following image.



Click either:


- Generate & Save* to save the generated flows but not submit them automatically.
- Generate & Submit* to submit the generated flows. The flows are not saved, they are generated in a temporary location. Only the target synonyms and tables are saved.

The *Generate Flows* window opens, as shown in the following image.



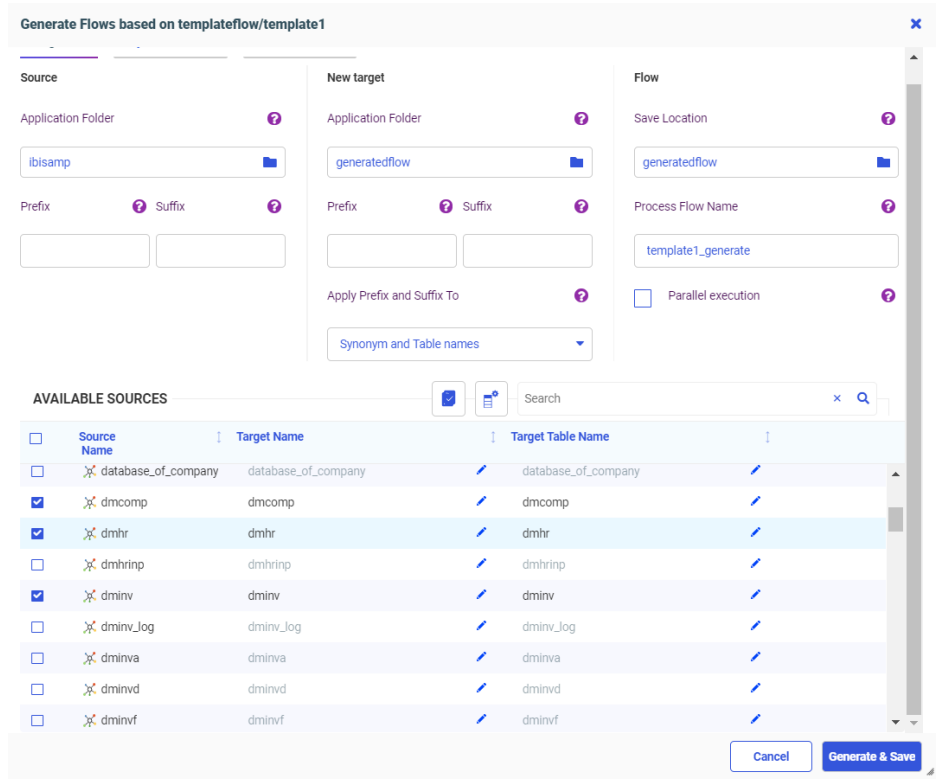
3. Select the source, target, and, if you are saving the flows, the flow application folders.

If the source application folder does not have the sources you want to use, click the application name and navigate to the application you need, as shown in the following image.



Do the same for the New Target application and the Flow application (if you are saving the flows). The source and target applications must be different if the target names match the source names. The New Target application will contain the new target synonyms. The Flow application will contain the generated data flows and process flow. The Target and Flow application folders can be the same or different.

- Select the sources for the new flows, as shown in the following image.



The sources do not have to be the same type as the source in the template flow. Any source type can be selected.

- Optionally, use one of the following techniques to edit generated target synonym and table names.
 - You can type over the default target names to create new names.
 - You can use the Source *Prefix* text box to specify characters that will be removed from the start of the source synonym names when creating the target names, and the Source *Suffix* text box to specify characters that will be removed from the end of the source synonym names when creating the target names.
 - You can use the Target *Prefix* text box to specify characters to be added to the start of the target names, and the Target *Suffix* text box to specify characters to be appended to the end of the target names.

By default, the Prefix and Suffix transformations will be applied to both target synonym and table names. However, you can use the *Apply Prefix and Suffix To* drop-down list to select whether to apply them to:

- Target Synonym and Table names (the default).
- Target Synonym names only.
- Target Table names only.

The following image shows that the characters *dm* will be removed from the beginning of the source synonym name when creating the target synonym name, and that the characters *_tgt* will be appended to the synonym name for the new target.

Source Name	Target Name	Target Table Name
database_of_company	database_of_company	database_of_company
dmcomp	dmcomp	dmcomp
dmhr	dmhr	dmhr
dmhrinp	dmhrinp	dmhrinp
dminv	dminv	dminv
dminv_log	dminv_log	dminv_log
dminva	dminva	dminva
dminvd	dminvd	dminvd
dminvf	dminvf	dminvf


For example, the source synonym name *dminv* will generate a target synonym named *inv_tgt*.

6. Optionally, type over the target table names to generate the table names you need.
7. Optionally, type over the default process flow name to generate the process flow name you need.

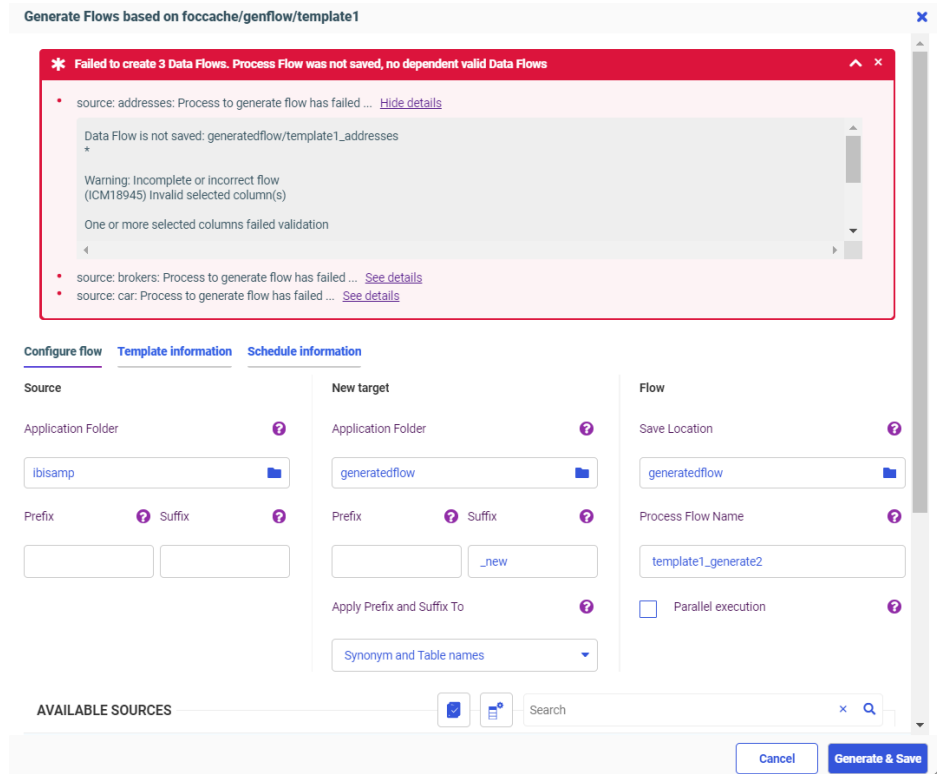
8. Optionally, select the *Parallel execution* check box to create a group with the data flows inside.
9. Once you have made all of your changes, click the *Verify* button to check that your selections are valid, as shown in the following image.



If there were any problems with your choices, warning messages will display in orange and error messages will display in red at the top of the window, identifying the issues. If a message is long, a *See details* link appears, as shown in the following image.




If you click a *See details* link, the Message box expands to show the details. If the details are long, scroll bars enable you to scroll through all of the details, as shown in the following image, in which the sources selected are not valid for the transformations in the flow.



You can click *Hide details* to close the details box.

Note that the *Generate & Save* button is still available, so you can save the erroneous flow.

If your choices are valid, a successful verification message displays in green, as shown in the following image.



10. If there were problems identified, correct them.

Examples of choices that could prevent verification are:

- The application folders for source and target are the same, with the source names matching the target names. The application folders must be different, unless you change the target names.
- The target synonym name is not valid. It should not exceed 64 characters, and the general guidelines for a synonym name should be followed.
- The target adapter is not available. It must be available (not removed accidentally) at the time when *Generate Flows* is in progress.

- ❑ Transformations you made in the template data flow are not valid for the generated flows. Any transformations you made in the template data flow must be valid for the generated flows. If the source tables do not have the columns necessary to apply the transformations, the flows will not be successfully verified.
11. Once the verification is successful, click either *Generate & Save* or *Generate & Submit*, depending on your selection when you started the Generate Flows process.

Some additional verification processes will be run. If problems occur, messages will display in red at the top of the window, identifying the problems. If no problems occur, a message will display in green, indicating that the save was successful and identifying the objects that were generated, as shown in the following image,

Generate Flows based on templateflow/template1

✓ Saved 3 Data Flows, and a Process Flow

Configure flow | Template information | Schedule information

Source

Application Folder: ibisamp

Prefix: dm

New target

Application Folder: generatedflow

Prefix: _tgt

Apply Prefix and Suffix To: Synonym name

Flow

Save Location: generatedflow

Process Flow Name: template2_generate

Parallel execution

AVAILABLE SOURCES

Source Name	Target Name	Target Table Name
<input type="checkbox"/> database_of_company	database_of_company	database_of_company
<input checked="" type="checkbox"/> dmcomp	dmcomp	dmcomp
<input checked="" type="checkbox"/> dmhr	dmhr	dmhr
<input type="checkbox"/> dmhrinp	dmhrinp	dmhrinp
<input checked="" type="checkbox"/> dminv	dminv	dminv
<input type="checkbox"/> dminv_log	dminv_log	dminv_log





Cancel Generate & Save

If additional problems were identified, correct them.

If you selected *Generate & Submit* when you started the Generate Flows process, the flows will be submitted, but not saved. Only the target synonyms and target tables will be saved.

12. Close the *Generate Flows* window by clicking *Cancel*.

13. If you selected *Generate & Save*, the generated data flows and the process flow are generated in the flow application directory you selected, as shown in the following image.

Application Directories/Files ↓	Size ↓	Date Modified ↓
 template1_dmcomp	2197	2020/07/07 10.42.47
 template1_dmhr	2979	2020/07/07 10.42.47
 template1_dminv	2299	2020/07/07 10.42.47
 template1_generate	1698	2020/07/07 10.42.47

The flow named `template1_generate` is the generated process flow. The other flows are the generated data flows, one for each source synonym (identified by having the source synonym name at the end of the flow name).

The following image shows the data flow generated for the `dminv` source synonym. The target synonym name was generated by removing the characters `dm` from the start of the source synonym name and adding the characters `_tgt` to the end of the target synonym name, as specified on the Generate Flows window.




If you selected *Synonym and Table names* from the *Apply Prefix and Suffix To* drop-down list, the prefix and suffix transformations will be applied to the target table name as well as the synonym name, as shown in the following image.

The screenshot shows a 'Load Options' dialog box with the following configuration:


- Load Option:** New/Replace
- Adapter:** MS SQL Server OLE DB/AzureDB
- Connection:** CON01
- Synonym Application:** generatedflow
- Synonym:** inv_tgt
- Table Name:** inv_tgt

The data flows are generated with the same load options as the template flow, but, you can now open any of the flows and edit them as needed.

To open the process flow, you will have to toggle the flow view by clicking the *Toggle Flow View* button at the top of the window, as shown in the following image.



The following image shows the generated process flow.



14. If you selected *Submit & Save*, once you have finished your edits to the flows, if any, you can submit the generated process flow by right-clicking it and selecting *Submit*.

The targets are generated in the New Target application folder (which, in this example, is the same as the Flow application folder), as shown in the following image.

Application Directories/Files ↓	Size ↓	Date Modified ↓
template1_dmcomp	2197	2020/07/07 10.58.56
template1_dmhr	2979	2020/07/07 10.58.56
template1_dminv	2299	2020/07/07 10.58.56
template1_generate	1698	2020/07/07 10.58.56
hr_tgt	3485	2020/07/07 10.59.15
inv_tgt	985	2020/07/07 10.59.16

Generating Data Flows Using a Template Flow

You can now use a template data flow that has one source and one new target to generate multiple new data flows. A process flow will be created that contains all of the generated data flows. This feature enables you to easily move any number of single sources to new targets.


Procedure: How to Generate Data Flows

The template data flow and the generated data flows can each have one source and one new target.

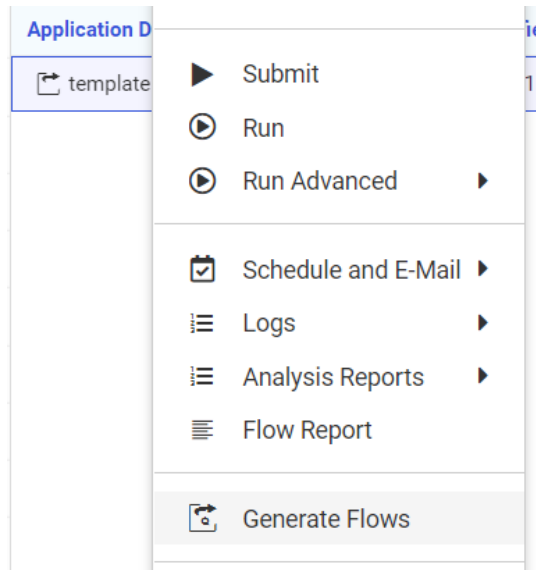
1. Create and save a data flow that has one source and one new target to use as the template.

The template flow is a normal data flow and can have any name and be saved in any application directory.

For example, the following image shows a flow named template1 that has a delimited source and new SQL Server target.

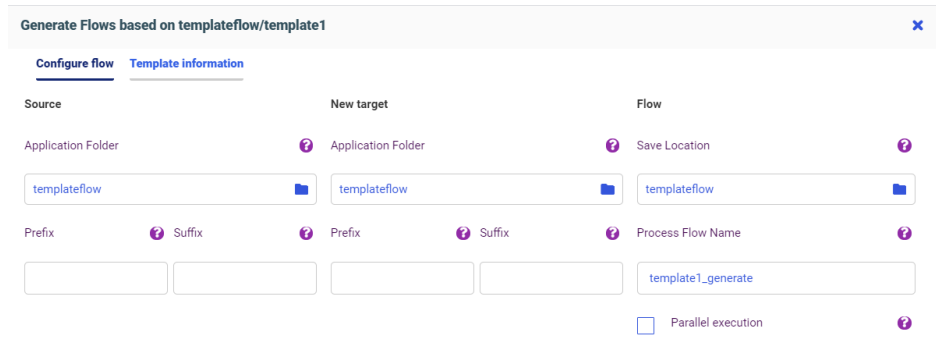


2. Right-click the template flow and click *Generate Flows*, as shown in the following image.



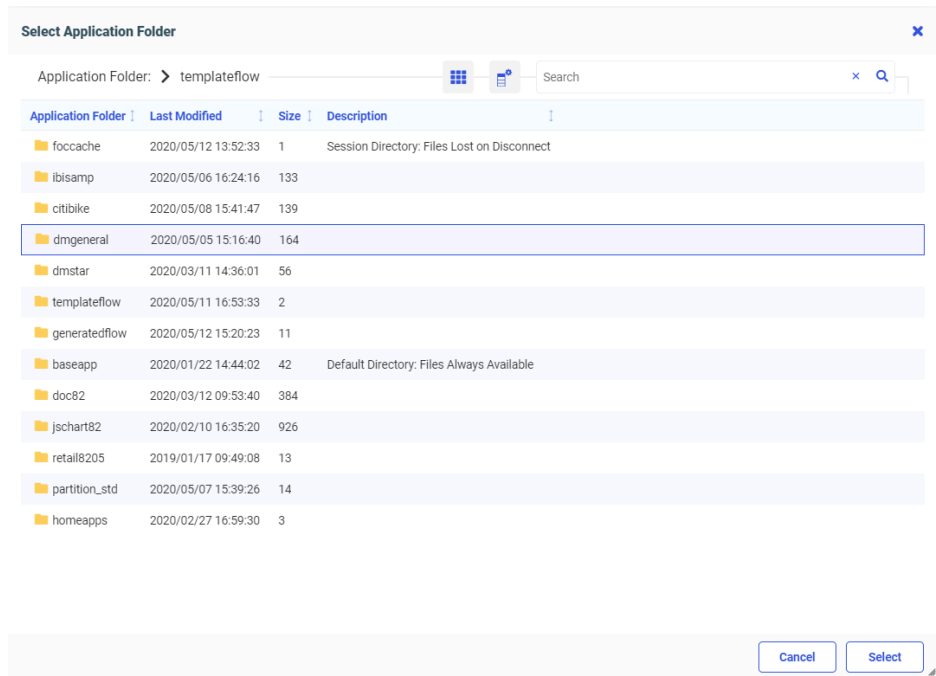
Note: The *Generate Flows* option appears as long as the flow conforms to the requirements for a template flow (one source, one new target).

The *Generate Flows* window opens, as shown in the following image.



3. Select the source, target, and flow application folders.

If the source application folder does not have the sources you want to use, click the application name and navigate to the application you need, as shown in the following image.



Do the same for the New Target application and the Flow application. The source and target applications must be different. The New Target application will contain the new target synonyms. The Flow application will contain the generated data flows and process flow. The Target and Flow application folders can be the same or different.

4. Select the sources for the new flows, as shown in the following image.

Generate Flows based on templateflow/template1

Configure flow | **Template information**

Source | **New target** | **Flow**

Application Folder: dmgeneral | Application Folder: generatedflow | Save Location: generatedflow

Prefix: | Suffix: | Prefix: | Suffix: | Process Flow Name: template1_generate

Parallel execution

AVAILABLE SOURCES

<input type="checkbox"/>	Source Name	Target Name	Target Table Name
<input checked="" type="checkbox"/>	dmcntr	dmcntr	dmcntr
<input checked="" type="checkbox"/>	dmcomp	dmcomp	dmcomp
<input checked="" type="checkbox"/>	dmhr	dmhr	dmhr
<input type="checkbox"/>	dmhrinp	dmhrinp	dmhrinp
<input checked="" type="checkbox"/>	dminv	dminv	dminv
<input type="checkbox"/>	dminva	dminva	dminva

Submit | Cancel | Save

The sources do not have to be the same type as the source in the template flow. Any source type can be selected.

5. Optionally, use one of the following techniques to edit generated target synonym names.
 - You can type over the default target names to create new names.
 - You can use the Source *Prefix* text box to specify characters that will be removed from the start of the source synonym name when creating the target synonym name, and the Source *Suffix* text box to specify characters that will be removed from the end of the source synonym name when creating the target synonym name.

- You can use the Target *Prefix* text box to specify characters that will be added to the start of the target synonym name, and the Target *Suffix* text box to specify characters that will be added to the end of the target synonym.

The following image shows that the characters *dm* will be removed from the beginning of the source synonym name when creating the target synonym name, and that the characters *_tgt* will be appended to the synonym name for the new target.

Generate Flows based on templateflow/template1

Configure flow | Template information

Source | New target | Flow

Application Folder: dmgeneral | generatedflow | Save Location: generatedflow

Prefix: dm | Suffix: | Prefix: | Suffix: _tgt | Process Flow Name: template1_generate

Parallel execution

AVAILABLE SOURCES

Source Name	Target Name	Target Table Name
<input checked="" type="checkbox"/> dmcntr	dmcntr	dmcntr
<input checked="" type="checkbox"/> dmcomp	dmcomp	dmcomp
<input checked="" type="checkbox"/> dmhr	dmhr	dmhr
<input type="checkbox"/> dmhrinp	dmhrinp	dmhrinp
<input checked="" type="checkbox"/> dminv	dminv	dminv
<input type="checkbox"/> dminva	dminva	dminva

Submit | Cancel | Save

For example, the source synonym name *dminv* will generate a target synonym named *inv_tgt*.

- Optionally, type over the target table names to generate the table names you need.
- Optionally, type over the default process flow name to generate the process flow name you need.
- Optionally, select the *Parallel execution* check box to create a group with the data flows inside.

9. Once you have made all of your changes, click the *Verify* button to check that your selections are valid, as shown in the following image.

<input type="checkbox"/>	Source Name	Target Name	Target Table Name
<input type="checkbox"/>			

If there were any problems with your choices, messages will display in red at the top of the window, identifying the problems. If your choices are valid, a successful verification message displays in green, as shown in the following image.

Generate Flows based on templateflow/template1

✓ Verified Successfully

Configure flow | Template information

Source	New target	Flow
Application Folder	Application Folder	Save Location
dmgeneral	generatedflow	generatedflow
Prefix	Prefix	Process Flow Name
dm	_tgt	template1_generate
		<input type="checkbox"/> Parallel execution

AVAILABLE SOURCES

<input type="checkbox"/>	Source Name	Target Name	Target Table Name
<input checked="" type="checkbox"/>	dmcntr	dmcntr	dmcntr
<input checked="" type="checkbox"/>	dmcomp	dmcomp	dmcomp
<input checked="" type="checkbox"/>	dmhr	dmhr	dmhr
<input type="checkbox"/>	dmhrinp	dmhrinp	dmhrinp
<input checked="" type="checkbox"/>	dminv	dminv	dminv

Submit Cancel Save

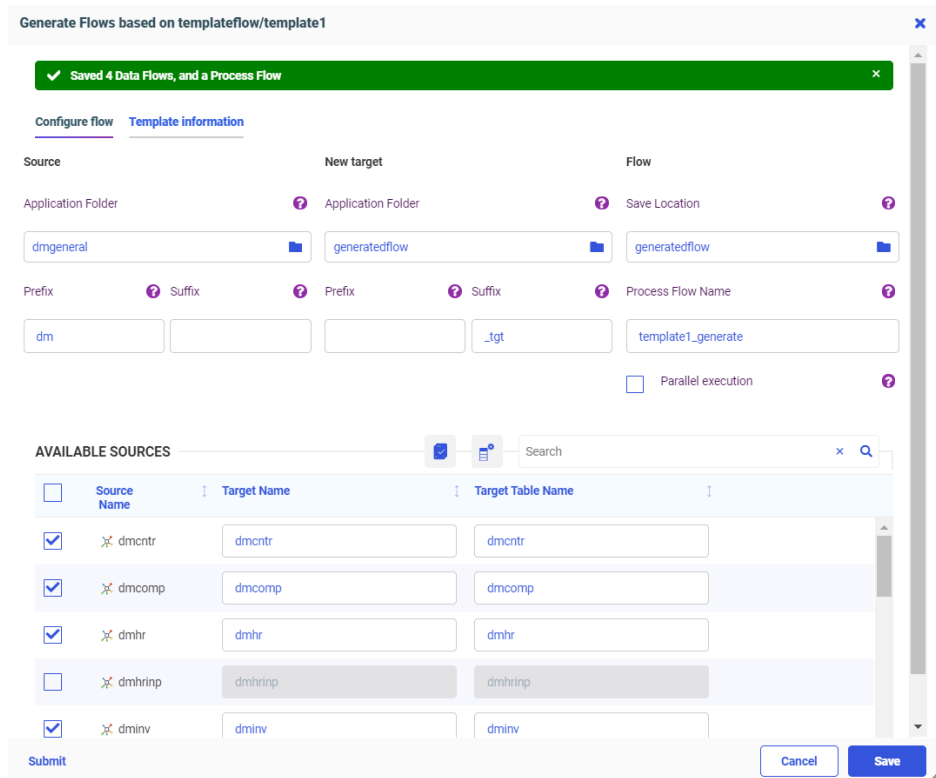
If there were problems identified, correct them. Examples of choices that could prevent verification are:

- The application directories for source and target are the same. They must be different.
- The target synonym name is not valid. It should not exceed 64 characters, and the general guidelines for a synonym name should be followed.

- ❑ The target adapter is not available. It must be available (not removed accidentally) at the time when *Generate Flows* is in progress.
- ❑ Transformations you made in the template data flow are not valid for the generated flows. Any transformations you made in the template data flow must be valid for the generated flows. If the source tables do not have the columns necessary to apply the transformations, the flows will not be successfully verified.

10. Once the verification is successful, click Save.

Some additional verification processes will be run. If problems occur, messages will display in red at the top of the window, identifying the problems. If no problems occur, a message will display in green, indicating that the save was successful and identifying the objects that were generated, as shown in the following image,








If there were problems, correct them.

You can submit the flows at this point, by clicking the *Submit* button.

11. Close the *Generate Flows* window by clicking *Cancel*.

The generated data flows and the process flow are generated in the application directories you selected, as shown in the following image.

Application Directories/Files ↓	Size ↓	Date Modified ↓
 template1_dmcntr	2146	2020/05/12 13.43.43
 template1_dmcomp	2233	2020/05/12 13.43.43
 template1_dmhr	3015	2020/05/12 13.43.43
 template1_dminv	2335	2020/05/12 13.43.43
 template1_generate	2031	2020/05/12 13.43.43

The flow named `template1_generate` is the generated process flow. The other flows are the generated data flows, one for each source synonym (identified by having the source synonym name at the end of the flow name).

The following image shows the data flow generated for the `dminv` source synonym. The target synonym name was generated by removing the characters `dm` from the start of the source synonym name and adding the characters `_tgt` to the end of the target synonym name, as specified on the Generate Flows window.

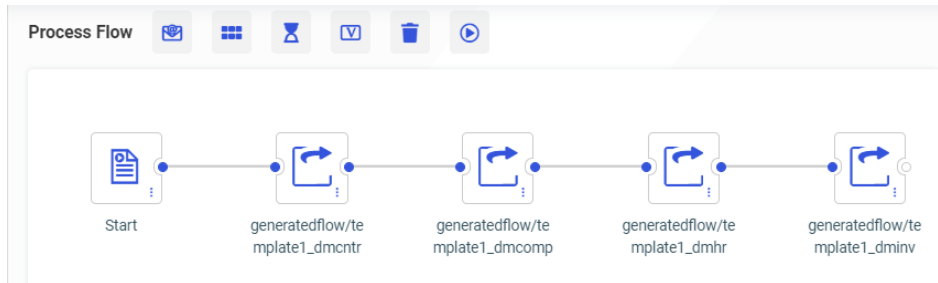


The data flows are generated with the same load options as the template flow, but you can now open any of the flows and edit them as needed.

To open the process flow, you will have to toggle the flow view by clicking the *Toggle Flow View* button at the top of the window, as shown in the following image.



The following image shows the generated process flow.



12. You can add a schedule, an email node, variables, a wait node, or a parallel group to the process flow using the buttons at the top of the page.
13. Once you have finished your edits to the flows, if any, you can submit the generated process flow by right-clicking it and selecting *Submit*.

The targets are generated, as shown in the following image.

The screenshot shows the WebFOCUS SERVER interface. The top navigation bar includes a search icon, a '+ Get Data' button, a 'Filter' button, and a 'Manage' button. Below the navigation bar, the breadcrumb 'Applications > generatedflow' is visible. The main content area is split into two panes: 'File Panel' and 'Application Path Configuration'. The 'File Panel' on the left shows a tree view of applications, with 'generatedflow' selected. The 'Application Path Configuration' pane on the right displays a table of application directories/files.

Application Directories/Files	Size	Date Modified
template1_dmcntr	2146	2020/05/12 13.43.43
template1_dmcomp	2233	2020/05/12 13.43.43
template1_dmhr	3015	2020/05/12 13.43.43
template1_dminv	2335	2020/05/12 13.43.43
template1_generate	2031	2020/05/12 13.43.43
cntr_tgt	225	2020/05/12 14.42.03
comp_tgt	569	2020/05/12 14.42.04
hr_tgt	3485	2020/05/12 14.42.05
inv_tgt	985	2020/05/12 14.42.06

Upload, Connection, and Synonym Enhancements

This section describes new features for the Get Data page for uploading files, connecting to data, and creating synonyms, and to the Synonym Editor (Data Assist) page for editing synonyms and Business Views.

Classifying Metadata During Upload and Using it in a Data Flow

Metadata Classification examines your data and assigns classifications to the columns, which can then be used to match columns from separate data sources. In this release, you can classify data that you upload, and use it to match fields in a Union in a Data Flow.

When uploading a data file, you can now view the recommended metadata classification values for each character-valued column. You can choose to keep the recommended values, or to change them.

Including classification values in your data improves the accuracy of mapping column tables correctly to each other. This is useful if you are integrating data from multiple sources, or if integrating data into a system with a predefined hierarchy.

The new Data Classification engine is a machine learning application, provided as a service, that will base class assignments on examining actual data values. It is trained by being given access to large volumes of data gathered throughout the world. This application will continually learn more about data and how to classify it, and new classes will be added as needed.

Prerequisites for Metadata Classification

In prior releases, automatic matching of columns from multiple data sources was based either on primary key/foreign key relationships in relational tables or, if those were not available, column names and formats. The new Data Classification engine is a machine learning application, provided as a service, that will base class assignments on examining actual data values. It is trained by being given access to large volumes of data gathered throughout the world. This application will continually learn more about data and how to classify it, and new classes will be added as needed.

A customer site wanting to use the Data Classification application will have to submit proprietary data to the classification API. To ensure the safety and integrity of proprietary data, the API will be installed as a service in the customer environment, and will be accessible only to the customer. Once the data is examined, feature vectors with encoded values will be sent to another service, available in the WebFOCUS or Omni environment, which will use these vectors to recommend class assignments for columns. If a classification cannot be recommended based on the known classes, an attempt will be made to classify columns based on column names and formats. You will be able to view and change the classifications prior to saving them in a synonym. Once they are stored in a synonym, they will be used to suggest column matches when you create a Union in a Data Flow on the server. In future releases, the functionality will be extended to other environments.

Instructions for installing the services are included your WebFOCUS Installation manual. Please make sure your installation team has installed them prior to using this feature.

After the services are installed, you must add the paths to the services in the `edaserv.cfg` file on the server.

Changes Needed in the `edaserv.cfg` Configuration File

The URLs pointing to the services (APIs) must be added to the `edaserv.cfg` configuration file on the WebFOCUS Server.

1. On the Server Console, navigate to the Workspace page using the Tools menu.
2. On the Resources Tree, expand *Configuration Files*.

- Right-click *Workspace - edaserve.cfg* and click *Edit*.

The edaserve.cfg file opens in the Text Editor.

- Enter values for the following keywords:

```
lblfeat_url = url_to_service_that_encodes_values
lblpred_url = url_to_service_calculates_labels_based_on_encoded_values
```

where:

url_to_service_that_encodes_values

Is the URL to the service that resides in the customer environment. For example:

```
lblfeat_url = https://myserver/api/1.0.0
```

This service is called by passing values of rows in a column, and it passes back a feature vector that describes the patterns of values in the column.

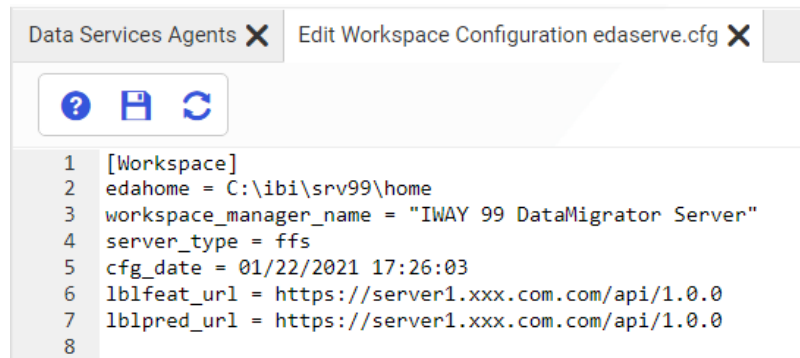
url_to_service_calculates_labels_based_on_encoded_values

Is the URL to the service that resides in the WebFOCUS and Omni environment. For example:

```
lblpred_url = https://ibiserver/api/1.0.0
```

This service is called using the feature vector created by the first service and returns the classification for the column.

The Text Editor should look similar to the following.



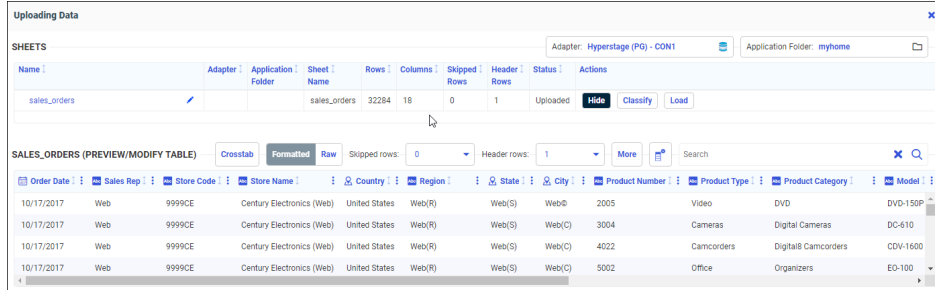
- Click the Save icon and close the tab.

Procedure: How to Change Metadata Classification Values When Uploading Data Files

You can change the metadata classification values for your data during the upload process.

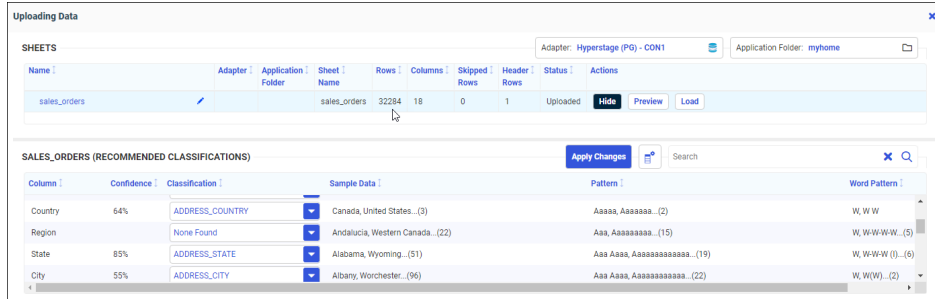
1. From the Home Page or Server Get Data page, upload an Excel or delimited file using Get Data.

A preview page opens, as shown in the following image.



2. Click the *Classify* button.

The Recommended Classifications page opens, as shown in the following image.



When a recommended classification can be made, the columns that were classified have a Confidence Level associated with them, which indicates the probability that the classification is correct.

You can select from a column list of classifications, or enter your own classification, by typing into the text box next to the column. Clicking a column drop-down arrow shows the list of classifications, and indicates which is recommended, as shown in the following image.

The screenshot shows the 'Uploading Data' window with the 'SALES_ORDERS (RECOMMENDED CLASSIFICATIONS)' section. A dropdown menu is open for the 'City' column, showing a list of classifications. The 'ADDRESS_STATE' classification is highlighted as 'Recommended'. The 'NEXT STEPS' section is also visible.

Column	Confidence	Classification	Sample Data	Pattern	Word Pattern
Country	64%	ADDRESS_COUNTRY	Canada, United States... (3)	Aaaa, Aaaaaa... (2)	W, W W
Region	None Found	None Found	Andalucia, Western Canada... (22)	Aa, Aaaaaaaa... (15)	W, W-W-W-W... (5)
State	85%	ADDRESS_STATE	Alabama, Wyoming... (51)	Aa Aaaa, Aaaaaa... (19)	W, W-W-W (.)... (6)
City	55%	ADDRESS_1	lary, Worcester... (96)	Aa Aaaa, Aaaaaa... (22)	W, W(W)... (2)

3. Click *Apply Changes*.

Your changes are applied, as indicated by a green message bar at the top of the page, as shown in the following image.

The screenshot shows the 'Uploading Data' window with a green message bar at the top indicating 'Classification changes successfully applied'. The 'SALES_ORDERS (RECOMMENDED CLASSIFICATIONS)' section is visible below, showing the same data as the previous screenshot.

Once your data is loaded, with the proper permissions to edit synonyms, you can leverage further capabilities by accessing the server. For more information, see [Technical Preview: Using Metadata Classifications in a Union in a Data Flow](#).

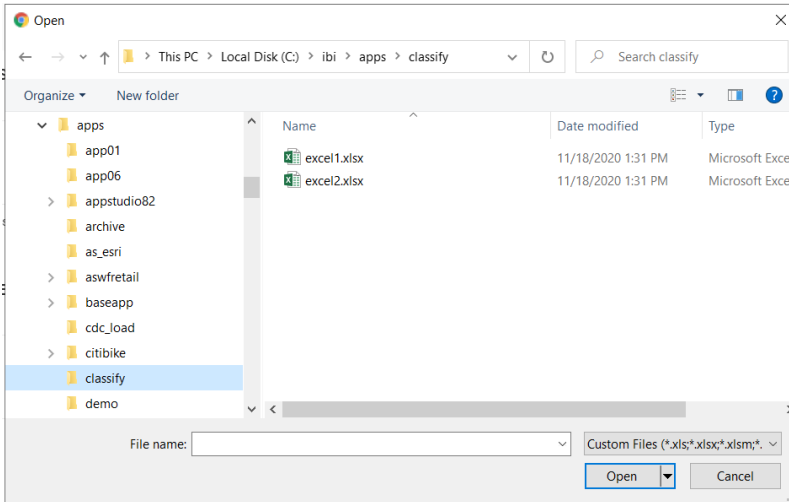
Using Metadata Classifications in a Union in a Data Flow

In this release, the classifications are only available when you upload a file from the Home Page or the Get Data page of the Server Console. We will upload two Excel files, and examine the classifications, then use the two Excel files in a Union.

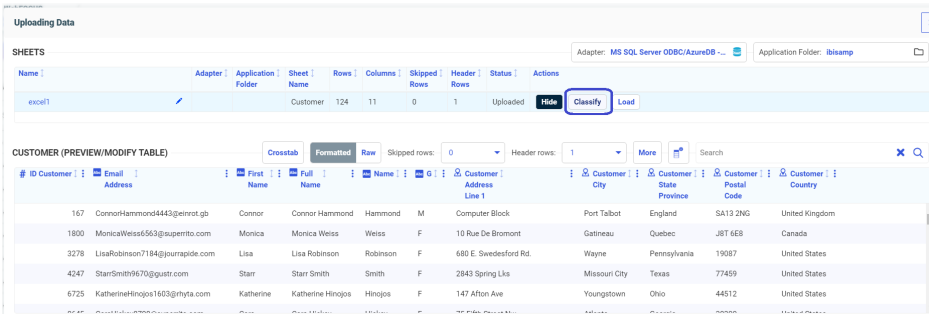
In the server environment, you can upload a file by clicking *Get Data* on the Applications page of the Server Console.

Classifying Excel1

Click *Excel* in the Local Files section, and a file picker dialog box opens, as shown in the following image.



Selecting *Excel1* and clicking *Open* opens the Upload Data page. If the Metadata Classification feature has been installed and configured, the Upload Data page has a *Classify* button, as shown in the following image.



Clicking the *Classify* button opens the Recommended Classifications page, as shown in the following image.

Column	Confidence	Classification	Sample Data	Pattern	Word Pattern
ID Customer			167.374849..(121)		
Email_Address		None Found	AaronConway8532@teleworm.ca, WolfgangBumgarner4230@enrot.de..(121)	AaAaaa9999@aaaaaa.aaa_AaaaaaaaAaaaaaa999@aaaaaaa.aa..(98)	WlG W W
FirstName	85%	FIRSTNAME	Aaron, Wolfgang..(117)	Aa_Aaaaaaaa..(6)	W
FullName		None Found	Aaron Conway, Wolfgang Bumgarner..(120)	Aa_Aaaaa_Aaaaaaaa_Aaaaaaaa..(42)	W W W W W W
Name		None Found	Allen, Winther..(116)	A Aaaa_Aaaaaaaa..(11)	W W W W
G	88%	GENDER	F, M	A	W
CustomerAddress_Line 1	54%	ADDRESS_1	1 Queen's Crescent, Zitone Airport..(121)	9 A 99Aa Aa_Aaaaaaaa..(112)	1 1 W W W W W
CustomerCity	76%	ADDRESS_CITY	Abbeyleix, Zwolle..(109)	Aa-Aaaaaaa_Aaaaaaaa..(27)	W W W W..(5)
CustomerState_Province	83%	ADDRESS_STATE	Alaska, Zurich..(71)	AaAaaa_Aaaaaaaa..(30)	W W W W..(4)
CustomerPostal_Code		None Found	..W1F 0UJ..(66)	..AA99 9AA..(7)	..WWW W W..(7)
CustomerCountry	85%	ADDRESS_COUNTRY	Austria, United States..(18)	Aaaaa_Aaaaaaaa..(6)	W W W

Note that not all columns were classified. You can select from the list of classifications for a column, or enter your own classification, by typing it into the text box next to the column. The columns that were classified have a Confidence Level associated with them. The confidence level represents the probability that the classification is correct.

In this case, select FULLNAME for Full,Name and LASTNAME for Name, as shown in the following image.

Column	Confidence	Classification	Sample Data	Pattern	Word Pattern
ID Customer			167.374849..(121)		
Email_Address		None Found	AaronConway8532@teleworm.ca, WolfgangBumgarner4230@enrot.de..(121)	AaAaaa9999@aaaaaa.aaa_AaaaaaaaAaaaaaa999@aaaaaaa.aa..(98)	WlG W W
FirstName	85%	FIRSTNAME	Aaron, Wolfgang..(117)	Aa_Aaaaaaaa..(6)	W
FullName		FULLNAME	Aaron Conway, Wolfgang Bumgarner..(120)	Aa_Aaaaa_Aaaaaaaa_Aaaaaaaa..(42)	W W W W W W
Name		LASTNAME	Allen, Winther..(116)	A Aaaa_Aaaaaaaa..(11)	W W W W
G	88%	GENDER	F, M	A	W
CustomerAddress_Line 1	54%	ADDRESS_1	1 Queen's Crescent, Zitone Airport..(121)	9 A 99Aa Aa_Aaaaaaaa..(112)	1 1 W W W W W
CustomerCity	76%	ADDRESS_CITY	Abbeyleix, Zwolle..(109)	Aa-Aaaaaaa_Aaaaaaaa..(27)	W W W W..(5)
CustomerState_Province	83%	ADDRESS_STATE	Alaska, Zurich..(71)	AaAaaa_Aaaaaaaa..(30)	W W W W..(4)
CustomerPostal_Code		None Found	..W1F 0UJ..(66)	..AA99 9AA..(7)	..WWW W W..(7)
CustomerCountry	85%	ADDRESS_COUNTRY	Austria, United States..(18)	Aaaaa_Aaaaaaaa..(6)	W W W

Click *Apply Changes*. A green message bar at the top of the page shows that the changes were successfully applied.

Click *Load*.

The file is uploaded, and the synonym is generated with the classifications added, as shown in the following partial synonym.

```
FILENAME=EXCEL1, SUFFIX=MSODBC , BV_NAMESPACE=OFF, $
SEGMENT=EXCEL1, SEGTYPE=S0, $
  FIELDNAME=ID_CUSTOMER, ALIAS=ID_CUSTOMER, USAGE=I9, ACTUAL=I4,
  MISSING=ON,
  TITLE='ID Customer', $
  FIELDNAME=EMAIL__ADDRESS, ALIAS=EMAIL__ADDRESS, USAGE=A42V,
ACTUAL=A42V,
  MISSING=ON,
  TITLE='Email ,Address', $
FIELDNAME=FIRST_NAME, ALIAS=FIRST_NAME, USAGE=A11V, ACTUAL=A11V,
MISSING=ON,
TITLE='First,Name',
CATEGORY='FIRSTNAME', CONFIDENCE=0.853071, $
FIELDNAME=FULL_NAME, ALIAS=FULL_NAME, USAGE=A26V, ACTUAL=A26V,
MISSING=ON,
TITLE='Full,Name',
CATEGORY='FULLNAME', $
FIELDNAME=NAME, ALIAS=NAME, USAGE=A15V, ACTUAL=A15V,
MISSING=ON,
TITLE='Name',
CATEGORY='LASTNAME', $
FIELDNAME=G, ALIAS=G, USAGE=A1V, ACTUAL=A1V,
MISSING=ON,
TITLE='G',
CATEGORY='GENDER', CONFIDENCE=0.883651, $
FIELDNAME=CUSTOMER_ADDRESS__LINE_1, ALIAS=CUSTOMER_ADDRESS__LINE_1,
USAGE=A57V, ACTUAL=A57V,
MISSING=ON,
TITLE='Customer,Address ,Line 1',
GEOGRAPHIC_ROLE=ADDRESS_LINE,
CATEGORY='ADDRESS_1', CONFIDENCE=0.547940, $
```

Classifying Excel2

Follow the same steps for file Excel2. The Classification Screen is shown in the following image.

Column	Confidence	Classification	Sample Data	Pattern	Word Pattern
ID Customer		None Found	167.374849..(121)		
Email Address		None Found	AaronConway8532@teleworm.ca, WolfgangBumgarner4230@enrot.de..(121)	AaAaaa9999@aaaaaa.aaa, AaaaaaaaAaaaaaa999@aaaaaaa.a..(98)	W@ W W
Name	85%	FIRSTNAME	Aaron, Wolfgang..(117)	Aa, Aaaaaaaa..(6)	W
Full Name		None Found	Aaron Conway, Wolfgang Bumgarner..(120)	Aa Aaaaa, Aaaaaaaa Aaaaaaa..(42)	W W, W W W W
Last Name		None Found	Allen, Winther..(116)	AaAaaa, Aaaaaaaa..(11)	W, W W W
Gender	88%	GENDER	F, M	A	W
Customer Address Line 1	54%	ADDRESS_1	1 Queen's Crescent, Zitone Airport..(121)	9 A 95Aa Aa, Aaaaaaaa..(112)	W W W W, W W
Customer City	76%	ADDRESS_CITY	Abbeyleix, Zwolle..(109)	Aa-Aaaaaaa, Aaaaaaaa..(27)	W, W W W..(5)
Customer State / Province	83%	ADDRESS_STATE	Alaska, Zurich..(71)	AaAaaaa, Aaaaaaaa..(30)	W, W W W..(4)
Customer Postal Code		None Found	..W1F 00U..(66)	..AA99 9AA..(7)	..W1W (W1)..(7)
Customer Country	85%	ADDRESS_COUNTRY	Austria, United States..(18)	Aaaaa, Aaaaaaaa..(6)	W, W W W

Note that the first name column is called Nme in this file (it was called Name in Excel1), and the gender column is called GENDER (it was called G in Excel1).


Again, select FULLNAME for Full,Name and LASTNAME for Name, click *Apply Changes*, and then *Load*.

Creating a Union

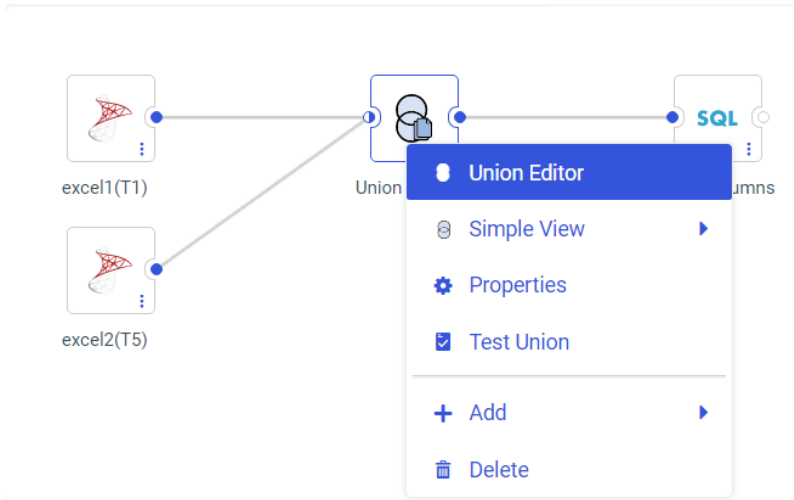
Now that the files have been uploaded and the synonyms have been generated with the classifications, click the New (+) button on the Server Console ribbon, and click *Flow*.

The Data Flow canvas opens.

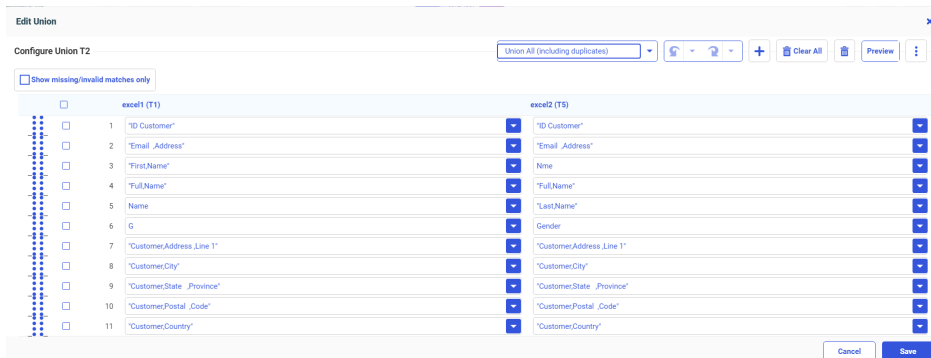
Drag Excel1 to the flow canvas, then drag Excel2 to the canvas and select Union, as shown in the following image.



A Union object is added to the flow. Right-click the Union object, and click *Union Editor*, as shown in the following image.



The Union Editor opens, as shown in the following image.



The classified fields are correctly matched. Fields that were not classified, such as Customer ID and Email Address, are matched based on column names and formats. If any matches are missing or incorrect, you can correct them using the drop-down lists.

New Setting for Controlling Display of Unreferenced Business View Fields

In prior releases, when a synonym had a Business View defined, the Object Inspector in the WebFOCUS tools, such as Designer, displayed only those fields that were referenced in a Business View folder.


A new setting is available in the Server Console that enables you to show all fields, even those not referenced in a folder.

To access this setting, do the following.

1. Navigate to the Workspace page in the Server Console.
2. Click *Settings*, then *Settings for Web Console Preferences*.

The *Change Settings for Web Console Preferences* page opens.

3. Expand the *Data Assist* section.
4. Scroll down to find the `NTM_METADATA_FIELDS` setting, as shown in the following image.



By default, it is set to the value *BV*, to show only those fields referenced in Business View folders.

5. If you want the WebFOCUS tools to display all fields, change the value to *All* and click *Save*.

Empty Application Message in the Server Console


When the Server Console opens to the Applications page by default, only the top level folders are displayed, and when an application folder is selected, only files in that folder are shown.

If you select an application folder that has no files, a message telling you to select a different application displays on the empty canvas, as shown in the following image.

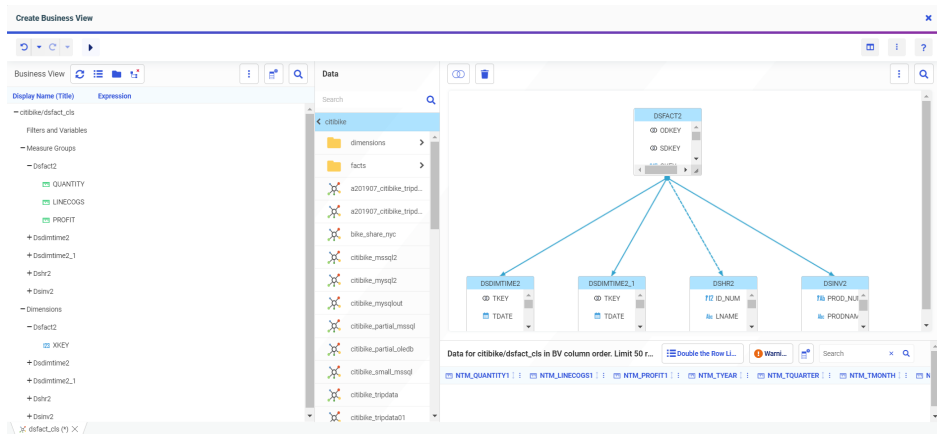


Enhanced Options for Creating Cluster Joins


When you use *Get Data* in Advanced Mode to create a synonym for multiple tables, two options are now available, as shown in the following image.



Cluster Synonym. All selected tables are added, and join columns can be selected as needed, as shown in the following image.



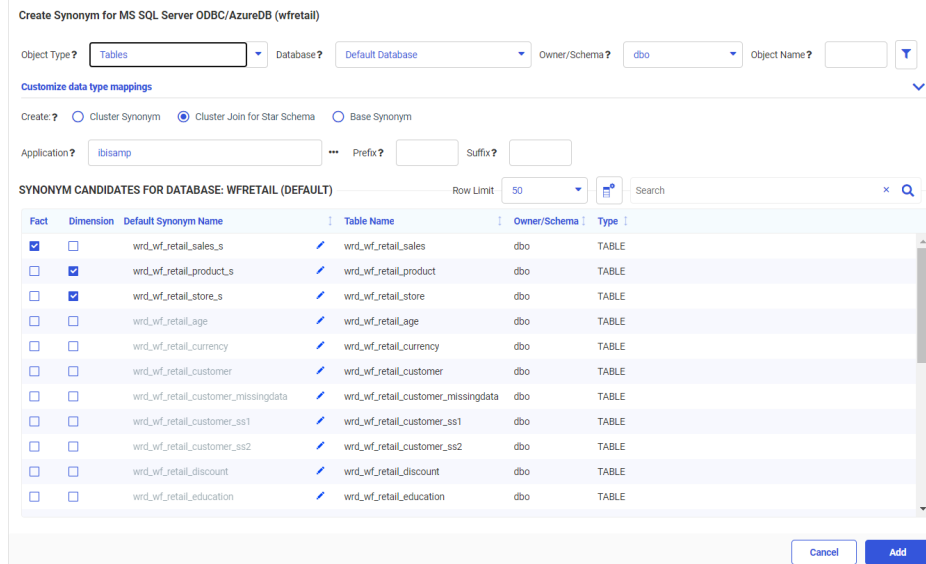
The dotted line indicates where no join condition could be identified. You can click the Warning in the Sample Data area to see the details, as shown in the following image.



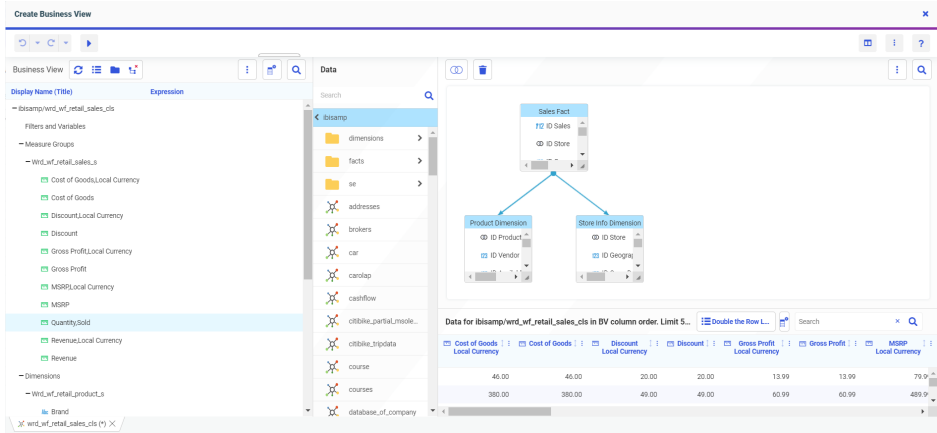
You can create a join condition in the Join Editor, if possible, or remove the table from the cluster.

- Cluster Join for Star Schema.** Only creates join conditions when the tables have Primary Key/Foreign Key relationships.


The following image shows the fact table and two related dimension tables selected to generate a cluster join for a star schema.



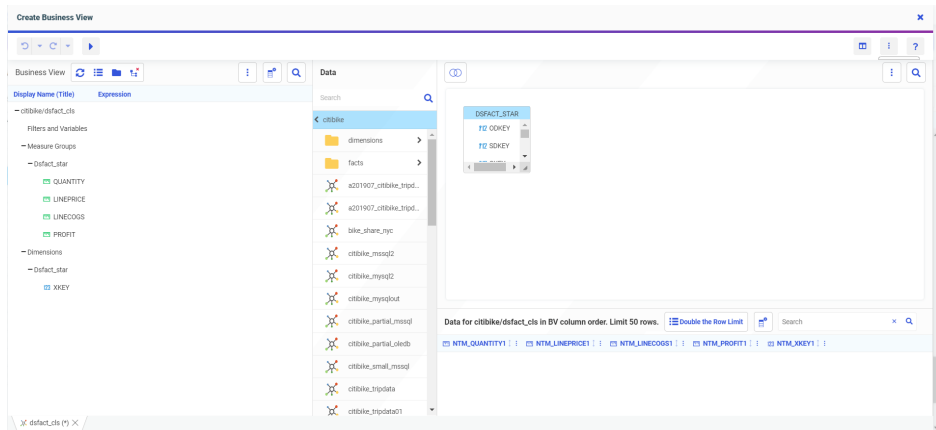
The following image shows the cluster join that was created.



If no primary key/foreign key relationship is found, the corresponding dimension table is not added to the cluster, and a diagnostic message is issued, as shown in the following image.



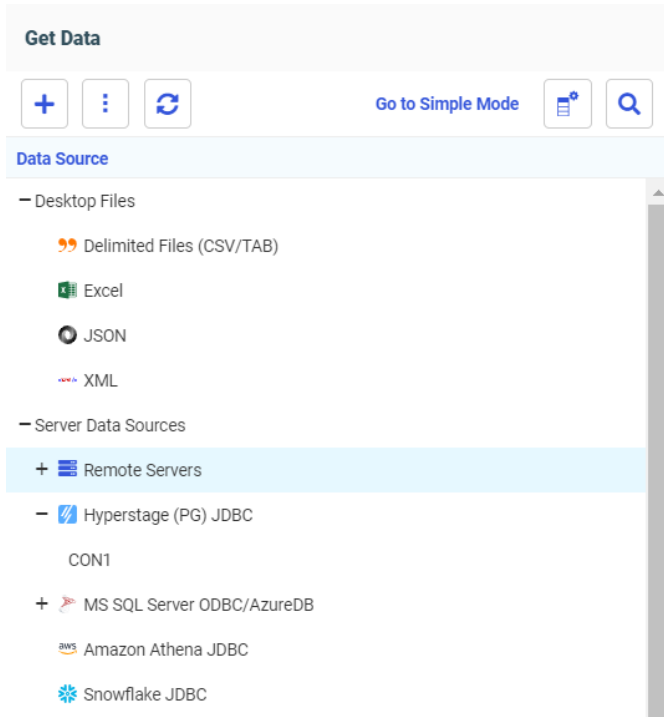
The dimensions were not added because of the lack of Primary Key/Foreign Key relationships, as shown in the following image.



Displaying the Remote Servers List Collapsed on the Get Data Advanced Page

In prior releases, the Remote Servers folder was shown expanded on the Get Data Advanced Mode page of the Server Console. When this list was long, a lot of scrolling was required to get past this list and see the configured adapters.

Now, the Remote Servers folder is collapsed by default, and the first adapter folder is open, as shown in the following image.



Directly Accessing the Advanced Mode of Get Data

Using *Get Data* from the server console, or from the WebFOCUS home page option *Prepare and Manage Data*, opens a standard user interface with the most common options. A link on that page for *Advanced mode* provides additional options.

Now, users who always want to use advanced mode can enable the `GET_DATA_ADV_DEFAULT` setting. To enable this setting do the following:

1. Navigate to the Workspace page on the server console.
2. Click *Settings*, then click *Settings for Web Console Preferences*.

The *Change Settings for Web Console Preferences* page opens.

3. Expand the *Data Assist* section and scroll down, as shown in the following image.

The screenshot shows a configuration window with several settings. The settings are as follows:

Setting Name	Value
NTM_SUPPORT_BV_PLUS_FOR_CUBE	Yes
AUTORUN_COMPOUND	Run on Change
USE_MONOTREE	Yes
AUTO_REDO_PREPARE	No
NTM_DEFAULT_BV_INCLUDE_KEYS	No
GET_DATA_ADV_DEFAULT	No

The **GET_DATA_ADV_DEFAULT** setting is highlighted with a red box, and its dropdown menu is open, showing the options "No" and "Yes". The "Yes" option is selected. Below the settings, there are three buttons: "Reset to defaults", "Cancel", and "Save".

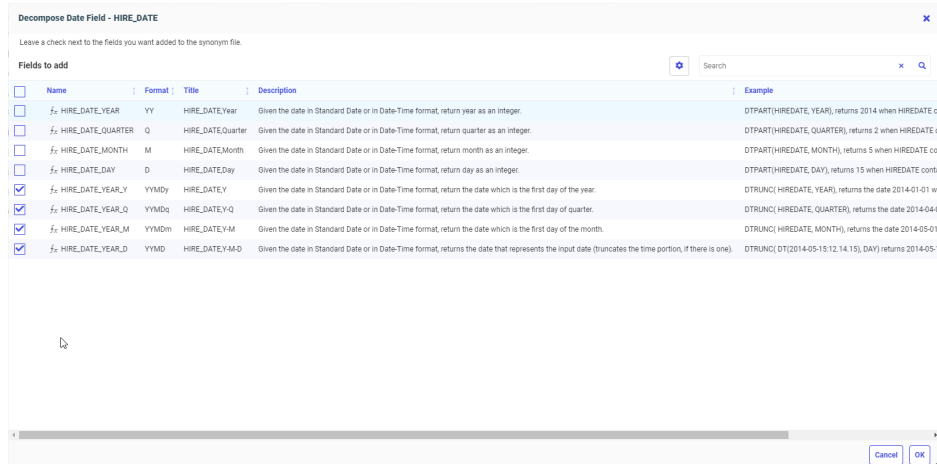
4. Select Yes for *GET_DATA_ADV_DEFAULT*.

5. Click Save.

With this setting in place, *Get Data* or *Prepare and Manage Data* will open the advanced mode of Get Data without first opening the simple mode.

Descriptions and Examples Now Shown in the Decompose Date Dialog Box

In the Synonym Editor or a Data Flow, the option to add or replace a date or date-time field with "decomposed dates" opens a dialog box that previously showed the names and titles that would be created for the new fields. Now, a description of what will be returned for each option and an example of its value have been added, as shown in the following image.



You can choose whether these additional items show in the Decompose Date dialog box by clicking the *Choose Columns* (gear) icon in the dialog box.

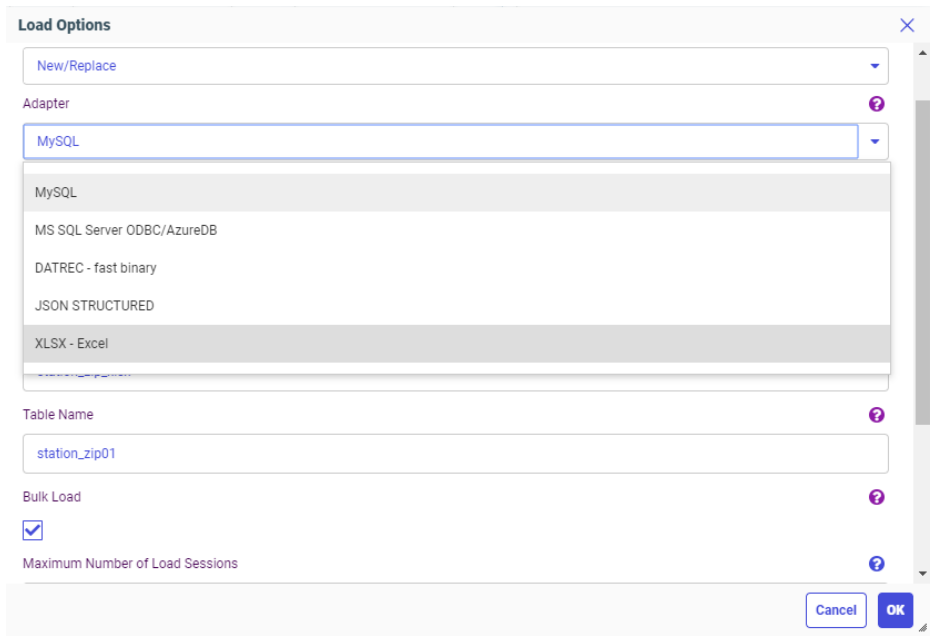
Supporting Excel Targets for Upload and Data Flow

In prior releases, using DataMigrator to create an Excel worksheet required configuring the Adapter for Fixed format files and, on Target Properties, selecting the Adapter for Formatted Files and the format Excel.

Now, when the Adapter for Excel (via Direct Retrieval) is configured, it can be selected as the target adapter. Files are created with the extension .xlsx. Additionally when this option is used, now a synonym is automatically created with the name and location specified, and with the suffix DIREXCEL.

Note that the obsolete format EXL2K that created files with the extension .xlht has been removed as an option.

The following image shows the Load Options dialog box when the Adapter for Excel (via Direct Retrieval) is configured.



The following image shows the flow with the Excel target configured.



The following Master File was generated for this target.

```
FILENAME=STATION_ZIP_XLSX, SUFFIX=DIREXCEL,
DATASET=citibike/station_zip_xlsx.xlsx, $
SEGMENT=STATION_ZIP_XLSX, SEGTYPE=S0, $
  FIELDNAME=STATION_ID, ALIAS=STATION_ID, USAGE=I7, ACTUAL=A11V,
  MISSING=ON,
  TITLE='STATION_ID', $
  FIELDNAME=ZIP_CODE, ALIAS=ZIP_CODE, USAGE=A6V, ACTUAL=A6V,
  MISSING=ON,
  TITLE='ZIP_CODE',
  GEOGRAPHIC_ROLE=POSTAL-CODE, $
  FIELDNAME=COUNTY, ALIAS=COUNTY, USAGE=A18V, ACTUAL=A18V,
  MISSING=ON,
  TITLE='COUNTY',
  GEOGRAPHIC_ROLE=COUNTY, $
  FIELDNAME=CITY, ALIAS=CITY, USAGE=A20V, ACTUAL=A20V,
  MISSING=ON,
  TITLE='CITY',
  GEOGRAPHIC_ROLE=CITY, $
  DEFINE ZIP_CODE__CNTRY__/A50V WITH STATION_ID='United States';
  TITLE='Country for ZIP_CODE',
  GEOGRAPHIC_ROLE=COUNTRY, $
```

The following Access File was generated for this target.

```
SEGNAME=STATION_ZIP_XLSX,
WORKSHEET=Sheet1,
HROWS=1,
NUMDATA=RAW, $
```

Setting the Maximum Size of Files for Upload

The setting `upload_maxsize` has been added to the Application Settings page to limit the size of files that can be uploaded. By default, file size for upload is not limited.

To change this setting, click *Manage*, then *Settings* on the Server Console Applications page or, on the Workspace page, click *Settings*, then *Workspace SET*, then *Application Settings*. The Application Settings pane opens, as shown in the following image.

The screenshot shows the 'Application Settings' pane with the following fields and values:

- focache_dir**: An empty text input field.
- focache_maxage**: A text input field containing the value '180'.
- nested_app**: A dropdown menu showing the value '5'.
- upload_allowed**: A text input field containing a list of file extensions: 'html,htm,css,xmi,txt,dot,ftm,wp,grf,jpg,jpeg,png,bmp,svg,pdf,zip,csv,tsv,tab,docx,doc,ppt,pptx,xls,xlsx,json'.
- upload_maxsize**: A text input field containing the value '0'.
- homeapps**: A text input field containing the path 'C:\lib\srvt2\homeapps'.

Below the 'upload_maxsize' field, there is a note: 'Home Applications feature allows each connected user to keep private versions of hold files, synonyms and other WebFocus objects. Location of home application root can be on Disk (type or use File picker to set physical directory) or SQL (type SQL/repository_prefix). Note: nested_app must not be 'n' to enable this setting.'

By default, `upload_maxsize` is zero (0), which supports unlimited file size for upload.

To limit file size for upload, enter a number of kilobytes to be the maximum supported, and click *Save and Restart Server*.


Once a limit is set, an attempt to upload a file larger than the maximum will generate the following message, and the upload will terminate.

The size of the uploaded file exceeds the limit of *nnnn* kilobytes

Supporting Full Outer Joins When the Data Source Does Not

In prior releases, when using the Synonym Editor to create a cluster join, it was only possible to create a full outer join if the underlying relational database supported that capability. Now, the server will generate multiple retrievals and perform the full outer join for those data sources.

For example, MySQL does not support full outer joins. However, you can implement a full outer join between MySQL tables in the Synonym Editor, as shown in the following image.



Changing Target Settings When Uploading Files

When using Get Data to upload a file, with appropriate permission, you can override the default Adapter and Connection settings by selecting an alternative configured adapter from your current profile.

When using Get Data to upload a file and create a synonym, with appropriate permission, you can override the default Application setting where the synonym is created by selecting an application folder from your current application path.

The screenshot shows the 'Uploading Data' dialog box. The 'SHEETS' section is expanded to show a table with columns: Name, Adapter, Application Folder, Sheet Name, Rows, Columns, Skipped Rows, Header Rows, Status, and Actions. The table contains one row for 'retail_sales' with the following values: Adapter: MySQL - mysql, Application Folder: ~pth_webfocus, Sheet Name: Retail Sales, Rows: 43247, Columns: 14, Skipped Rows: 0, Header Rows: 1, Status: Uploaded, and Actions: Hide, Load.


Name	Adapter	Application Folder	Sheet Name	Rows	Columns	Skipped Rows	Header Rows	Status	Actions
retail_sales	MySQL - mysql	~pth_webfocus	Retail Sales	43247	14	0	1	Uploaded	Hide Load

For more information see [Changing a Target Connection When Uploading Files](#).






For more information see [Changing a Target Application When Uploading Files](#).

Redesign of Business View Icons




When editing a Business View, the buttons now have icons that are more standard and clear, and are grouped in a more logical configuration, as shown in the following image.



The first group of buttons control the Business View:

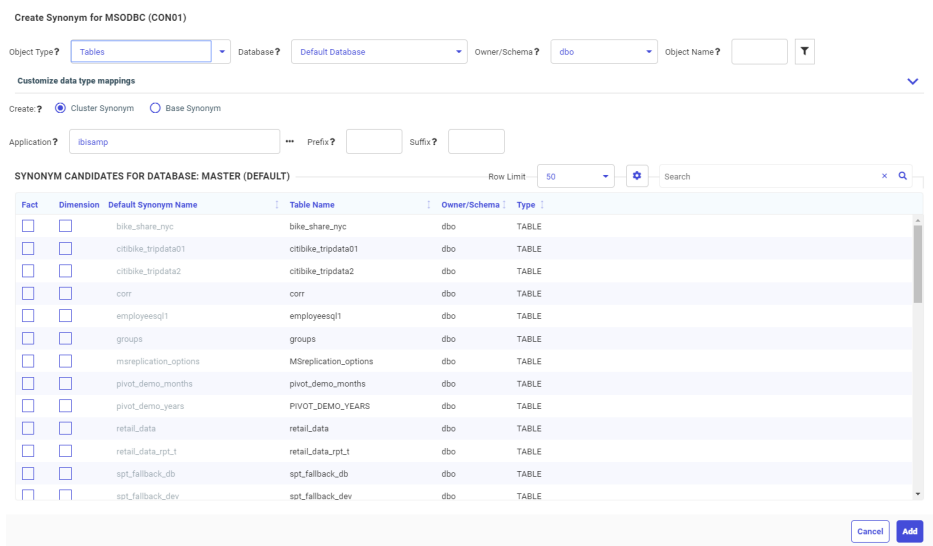
- ❑ The *Reset Default* button () recreates the default Business View.
- ❑ The *Flatten* button () lists all measure fields under a single measure group and all dimension fields under a single dimension, with no children.
- ❑ The *Create Template Folders* button () creates empty basic folders when there is no Business View already defined.
- ❑ The *Hide All* button () removes the Business View. A *Create Default* button () appears so you can create the default Business View.


The second group of buttons provide options for the window:

- ❑ The *Display* button () lets you show Folders and Columns or just Folders, or opens a window with a modeling view of the folders.
- ❑ The *Find* button () opens a search box in which you can enter characters. All component names that contain those characters are highlighted.
- ❑ The *Choose Columns* button () has two options. *Choose Columns* opens a dialog box for selecting the attributes in the synonym that you want to display. *Reset Columns* returns the display to the default columns.

Create Synonym Panel Redesign in Get Data Advanced Mode

On the WebFOCUS Server console, when using Get Data in the advanced mode to create metadata objects (synonyms) for tables in a data source, the user interface and on-screen help have been improved to make it easier to find and select tables, as shown in the following image.



- The selection criteria are on one line.
- You can enter Owner/Schema or select from a list. If you enter a name, you must click the *Filter* button () to search for the candidates. If you select from the list, the search happens automatically.
- In the Object Name field, you can use the percent character (%) as a wildcard.

Previewing a File When Uploading

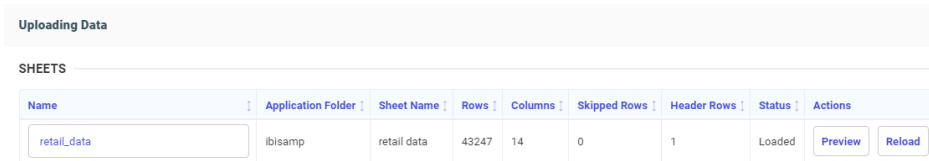
When you upload a file, you can preview the file to make sure that you uploaded the one that you wanted. You can also make some adjustments to the data before loading it to a target, as shown in the following image.

The screenshot shows the 'Uploading Data' window. At the top, there's a 'SHEETS' table with columns: Name, Application Folder, Sheet Name, Rows, Columns, Skipped Rows, Header Rows, Status, and Actions. Below that is the 'RETAIL DATA (PREVIEW)' section, which includes a 'Crosstab' button, 'Formatted' (selected) and 'Raw' buttons, 'Skipped rows: 0', and 'Header rows: 1'. A 'More' button and a search bar are also present. The main data table has columns: PRODUCT_CATEGORY, PRODUCT_SUBCATEGORY, STORE_TYPE, STORE_NAME, COUNTRY, STATE, CITY, REGION, SALE_DATE, REVENUE, COST_OF_GOODS, and GROS. A context menu is open over the 'CITY' column header, showing options: Geographic Role, Data Type, Continent Name, Country Name, State Name, Rename, Remove, County Name, City Name, Postal code, Address, Advanced, and (None).

When preparing the data:

- The headers and sample content are shown. By default, formatted data is shown, with the generated data types identified. You can click the *Raw* button to see unformatted data.
- The number of header rows is determined by scanning the file, but it can be changed as required.
- For an Excel Worksheet, there is an option to read the file as a cross-tab.
- The derived data type is shown for each field. You can change the basic data type.
- You can, optionally, change or assign a geographic role.
- A field name is shown for each field, either from the header, when available, or generated as the name FIELD with an index (FIELDx). You have the option to rename the field.
- You can also change the name of the metadata object (synonym) that is created and, for workbooks with multiple worksheets, choose which ones to load.
- You can click the *More* button to access additional file options.

After you load the data, the page that opens has a *Preview* button, as shown in the following image.



Name	Application Folder	Sheet Name	Rows	Columns	Skipped Rows	Header Rows	Status	Actions
retail_data	lbisamp	retail data	43247	14	0	1	Loaded	Preview Reload

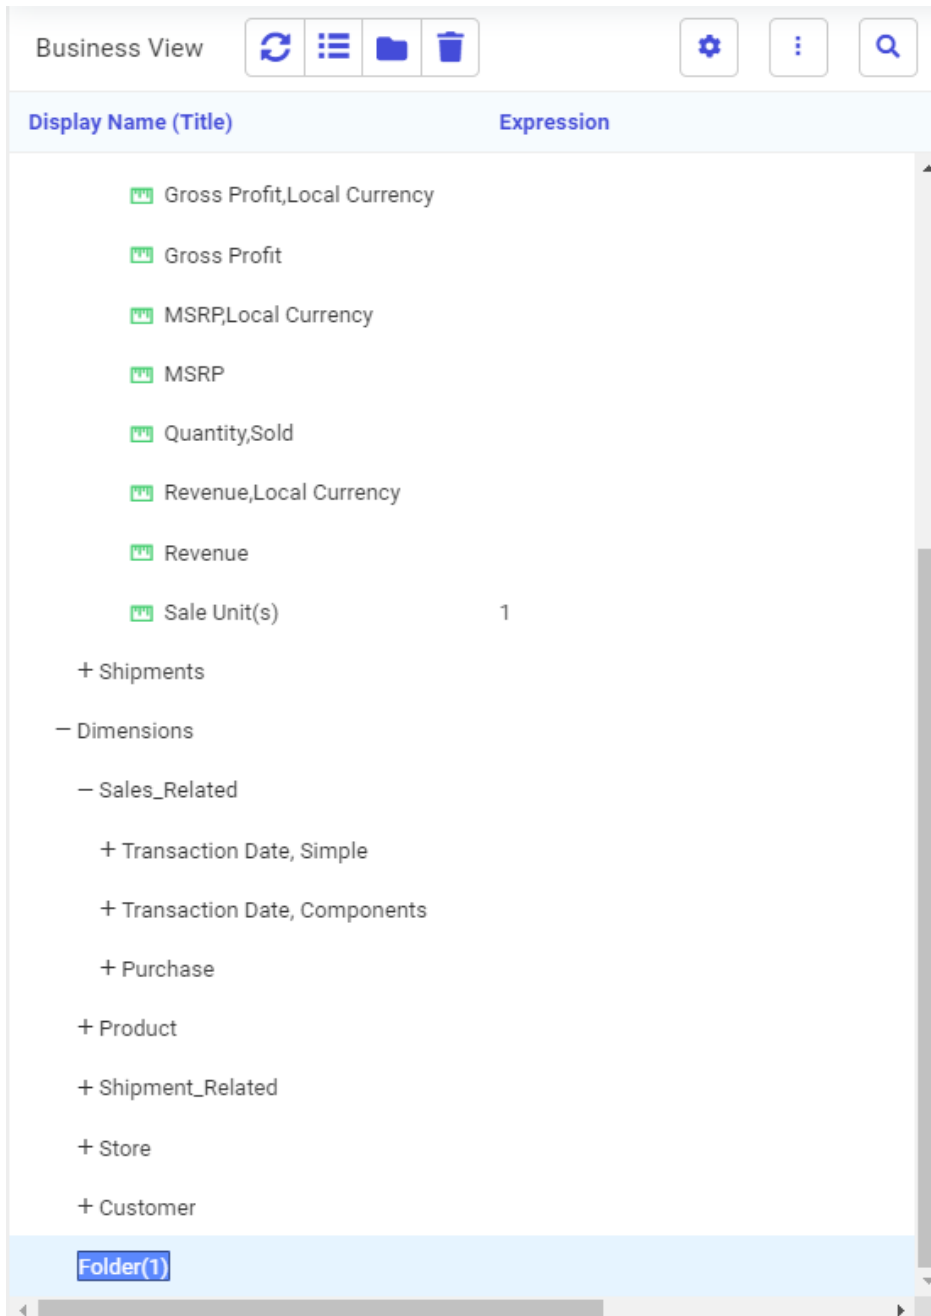
If you click *Preview*, you open the same preview pane, with the same options that were available in the Prepare step.

You can click *Reload*, to reload the data, if necessary.

Creating and Renaming a New Folder in a Business View

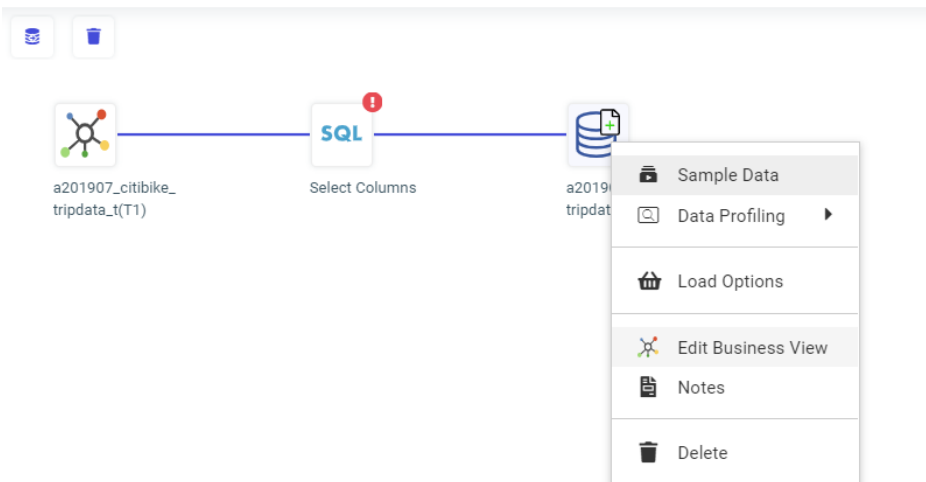
In prior releases, adding a new folder to a Business View in the Synonym editor created the new folder with a default name. To rename this folder, you then had to right-click the folder, click *Rename* from the context menu, and enter a new name.

Now, when you add a new folder, it is generated with the default name selected for renaming, and you can type the new name in place, as shown in the following image.

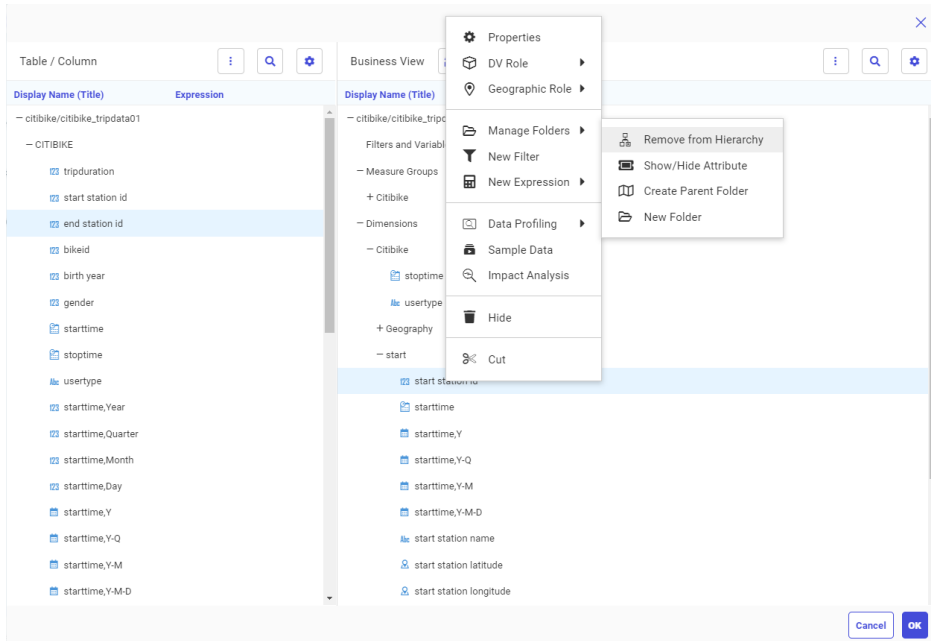


Creating a Business View for a Target Data Source




When you create and load a target table in a data flow, you can right-click the target after adding it to the flow and click *Edit Business View*, as shown in the following image.





A default Business View is created for the target table with default folders based on the fields in the synonym. A page opens on which you can edit the Business View, as shown in the following image.



The following are the buttons on the Business View toolbar:

Button	Definition	Description
	Reset	Resets to the initial default Business View
	Flatten	Shows the default folders only.
	Create Template Folders	In an empty Business View, creates the Filters and Variables, Measure Groups, and Dimensions folders with nothing in them.

Button	Definition	Description
	Hide All	Removes the Business View
	Create Default	If you removed the Business View, you can use this button to recreate the default Business View.

The following options are available on the Business View menu:

- Folders/Columns.** Shows the folders and columns within the folders.
- Folders.** Shows only folders.
- Modeling View of Folders.** Shows a graphical representation of the folders and the connections between them.
- Autorun.** Has the following options:
 - Data Preview.** Has options *Off* and *On*, where *Off* is the default. When Data Preview is *On* an additional panel is displayed that shows the data values in the order the fields appear in the Business View.
 - Off.** Does not automatically refresh the Data Preview panel when a change is made to the Business View. This is the default value.
 - On.** Refreshes the Data Preview panel for each change in the Business View.
- Diagnostics.** You can view the Master File or Access File, Open the Session Log, Clear the Session Log, and view the layout.
- Options.** Opens the data flow Advanced Options dialog box.

For more information about Business Views, including managing folders and hierarchies and specifying DV roles, see Chapter 5, *Metadata*, in the *Server Administration* manual.

Pivoting Multiple Columns on Upload

When you upload a spreadsheet or delimited file with columns of repeating values (for example each column contains an amount for a year) you can multi-select columns and click the context menu choice *Pivot to rows* to transpose the columns to rows.

For example, the following Excel file has a column for each year from 1960 to 2012. This file, `pivot_demo.xlsx`, is generated as part of the WebFOCUS Retail tutorial. If you want to try uploading this file using the *Pivot to rows* option, generate the WebFOCUS Retail tutorial, open the `uploads` folder in the application directory that contains the WebFOCUS Retail sample files, right-click `pivot_demo`, and click *Download*. This will download the file to your Downloads folder on Windows so you can upload it to the server.

To multi-select all of the individual year columns, click in the data portion of the first column, hold down the shift key, and click in the data portion of the last column. Then click the menu button in one of the columns, and click *Multiple columns to rows*, as shown in the following image.

	2004	2005	2006	2007	2008	2009	2010	2011	2012
06973	35,662.211822			45,181.4675944	49,679.1107045	45,872.2209766	44,723.2039437	49,338.7606551	46,642.2903265
61430	511.2929789	532.6109983	557.2324469	632.3608696	739.2561276	712.6165270	689.6483313	745.9154748	751.9163168
07426	1,490.3800560	1,731.1252346	2,069.3436307	2,651.2601212	3,413.5886614	3,749.2724237	4,433.3612200	5,447.3414216	6,091.0143517
05204	400.4724749	421.1233114	427.2911670	467.1363614	537.6384971	597.7118029	664.0642317	731.8941678	752.1560525
43873	326.1973133	340.9078669	365.4970833	413.4814386	474.3894790	464.5132313	456.5634087	498.7800669	482.6697266
45510	121.8593036	143.7835295	158.3050232	162.8275094	186.8717489	194.8966661	219.5297995	246.9143271	251.0145230
70756	371.2179868	406.9988043	422.8448175	474.5899781	569.6392934	553.0423743	592.6074699	649.9252300	634.3209746

The year columns become rows, as shown in the following image.

Country Name	Country Code	Indicator Name	Indicator Code	Pivot value	Pivot key
Afghanistan	AFG	GDP per capita (current US\$)	NY.GDPPCAPCD	61.28913206	1960
Afghanistan	AFG	GDP per capita (current US\$)	NY.GDPPCAPCD	61.30409306	1961
Afghanistan	AFG	GDP per capita (current US\$)	NY.GDPPCAPCD	59.79869332	1962
Afghanistan	AFG	GDP per capita (current US\$)	NY.GDPPCAPCD	80.42300207	1963
Afghanistan	AFG	GDP per capita (current US\$)	NY.GDPPCAPCD	83.79481181	1964
Afghanistan	AFG	GDP per capita (current US\$)	NY.GDPPCAPCD	103.08910311	1965
Afghanistan	AFG	GDP per capita (current US\$)	NY.GDPPCAPCD	140.13838324	1966

Target and Load Options Enhancements

This section describes new features for targets in a Data Flow and Load Options for targets.

Adapter for Apache Hive: Creating Avro Files

Using the Adapter for Apache Hive, files can now be created in the Avro self-describing binary file format. Avro is a portable (language neutral) format that can be created in the Hadoop Distributed File System (HDFS).

To create Avro files, create a Data Flow with a new Hive target. In the Load Options dialog box, select Avro for the TBL_STORED_AS option, as shown in the following image.

The image shows a 'Load Options' dialog box with the following configuration:

- Synonym: employe201
- Table Name: employe201
- Bulk Load:
- Escape Character: OFF - Default
- Allow Direct Bulk Load: OFF - Default
- Allow Direct Load to Parquet File: OFF - Default
- Overwrite data: OFF - Default
- Save Data Files: NO - Default
- Key columns derived from: Key columns in source table / sorted columns in qt
- TBL_STORED_AS: TEXTFILE (default)

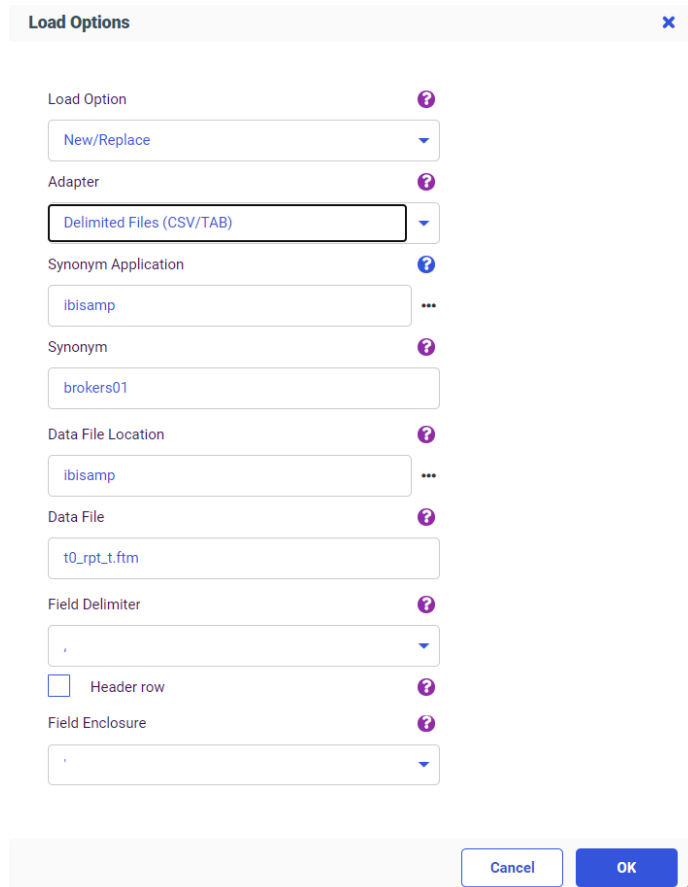
The dropdown menu for TBL_STORED_AS is open, showing the following options:

- PARQUET
- ORC
- AVRO

Buttons for 'Cancel' and 'OK' are visible at the bottom right of the dialog.

Using a Delimited Target When No Relational Adapter is Available

If you create a Data Flow in the Web Console or the Data Management Console (DMC), and no relational adapter is configured, the default Adapter will be Delimited Files, as shown in the following image.



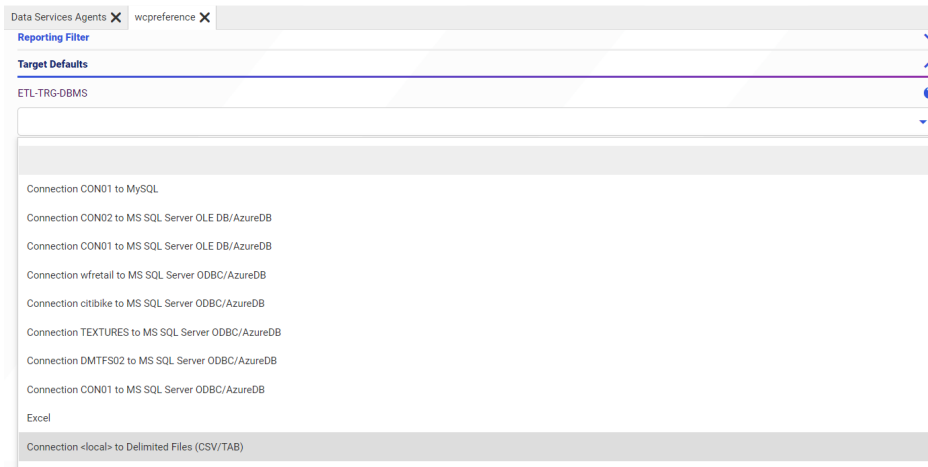
The image shows a 'Load Options' dialog box with the following fields and values:

- Load Option:** New/Replace
- Adapter:** Delimited Files (CSV/TAB)
- Synonym Application:** ibisamp
- Synonym:** brokers01
- Data File Location:** ibisamp
- Data File:** t0_rpt_t.ftm
- Field Delimiter:** ,
- Header row:**
- Field Enclosure:** ' (single quote)

Buttons: Cancel, OK

In addition, when a connection for the Adapter for Delimited Files is configured, you can select it as the default target adapter.

To set the connection to the Adapter for Delimited Files as the default target adapter, navigate to the Workspace page of the Web Console (or expand Workspace in the DMC). Click *Settings*, then *Settings for Web Console Preferences*. Scroll down to the *Target Defaults* section, and select the connection for the Adapter for Delimited Files in the *ETL-TRG-DBMS* drop-down list, as shown in the following image.



Selecting Excel as the Default Target Adapter in a Flow

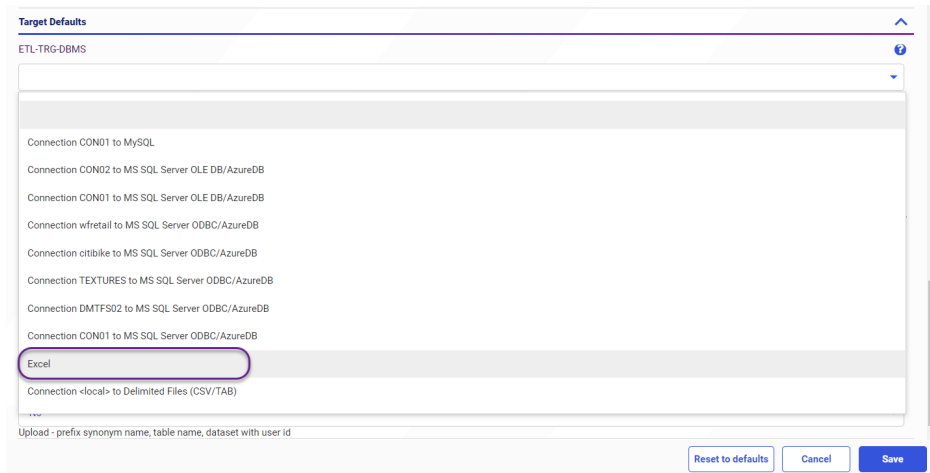
In prior releases, Excel could not be selected as the default target adapter for Data Flows. Now, if the Adapter for Excel is configured, Excel can be selected as the default adapter.

To select Excel as the default target adapter, navigate to the Workspace page in the Web Console, or expand Workspace in the Data Management Console (DMC).

Click *Settings*, then *Settings for Web Console Preferences*. The *Change Settings for Web Console Preferences* panel opens.

Scroll down, and expand the Target Defaults section.

Select *Excel* in the ETL-TRG-DBMS drop-down list, as shown in the following image.



Click Save.

Support for New Data Flow Targets With No Key Fields

When you create a Data Flow with a new target, you can now specify in the load options that the target has no key fields.

Target and Load Options Enhancements

To specify that the target has no key columns, select *No Key* from the *Key columns derived from* drop-down list, as shown in the following image.

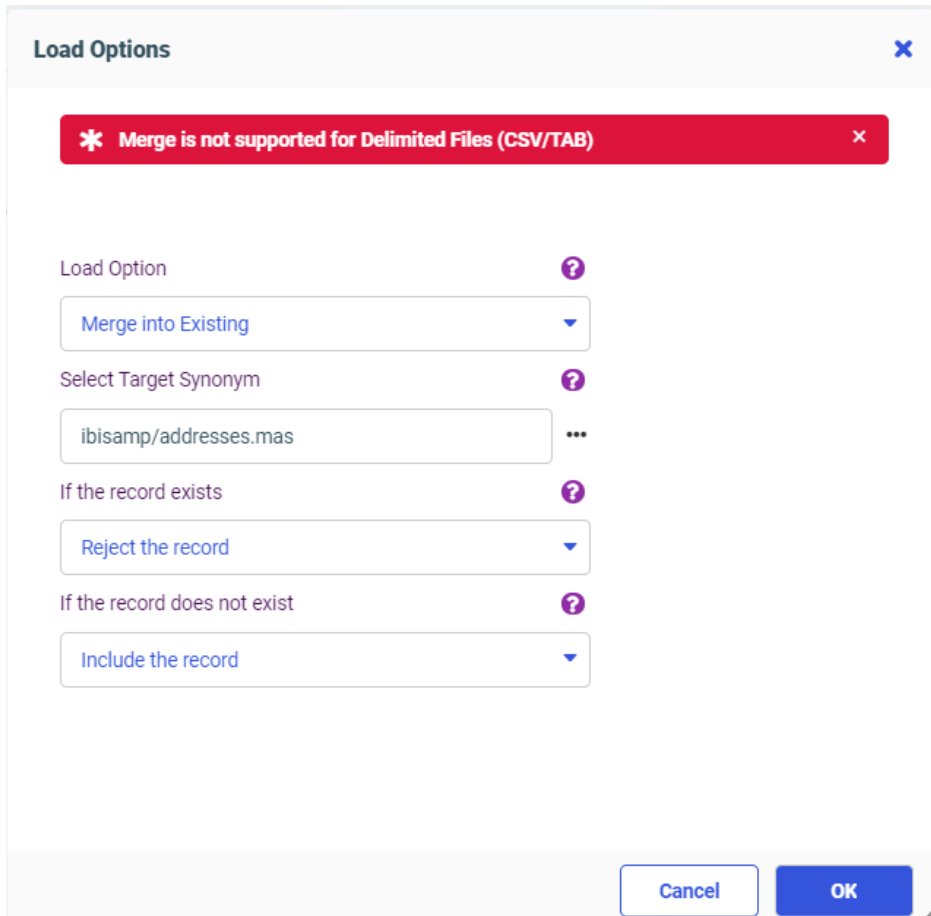
The image shows a 'Load Options' dialog box with the following fields and settings:

- Synonym Application:** citibike
- Synonym:** citibike_mssql1
- Table Name:** citibike_mssql1
- Bulk Load**
- Maximum Number of Load Sessions:** 10
- Maximum Number of Load Session Restarts:** 10
- Column Delimiter:** TAB
- Allow Direct Bulk Load:** OFF - Default
- Save Data Files:** NO - Default
- Key columns derived from:** No Key

The 'Key columns derived from' dropdown is open, showing a list with 'No Key' selected. The list also includes 'Key columns in source table / sorted columns in query'. At the bottom right of the dialog are 'Cancel' and 'OK' buttons.

Displaying Messages at the Top of the Load Options Screen

When you edit the load options for a data flow, error messages, if there are any, are now displayed in red on the top of the screen where they can more easily be seen, as shown in the following image.



The screenshot shows a dialog box titled "Load Options" with a close button (X) in the top right corner. At the top of the dialog, there is a red error message bar that reads: "* Merge is not supported for Delimited Files (CSV/TAB)" with a close button (X) on the right. Below the error message, there are four configuration options, each with a question mark icon to its right:


- Load Option:** A dropdown menu with "Merge into Existing" selected.
- Select Target Synonym:** A text input field containing "ibisamp/addresses.mas" and a three-dot menu icon to its right.
- If the record exists:** A dropdown menu with "Reject the record" selected.
- If the record does not exist:** A dropdown menu with "Include the record" selected.

At the bottom right of the dialog, there are two buttons: "Cancel" and "OK".

When you correct the problems, the messages disappear.

Show Confidence Level When Selecting an Existing Target in a Flow

When selecting an existing target in a data flow, the file picker shows the confidence level, based on matching column names and data types for each prospective new target, showing how compatible it is with the existing source, as shown in the following image.



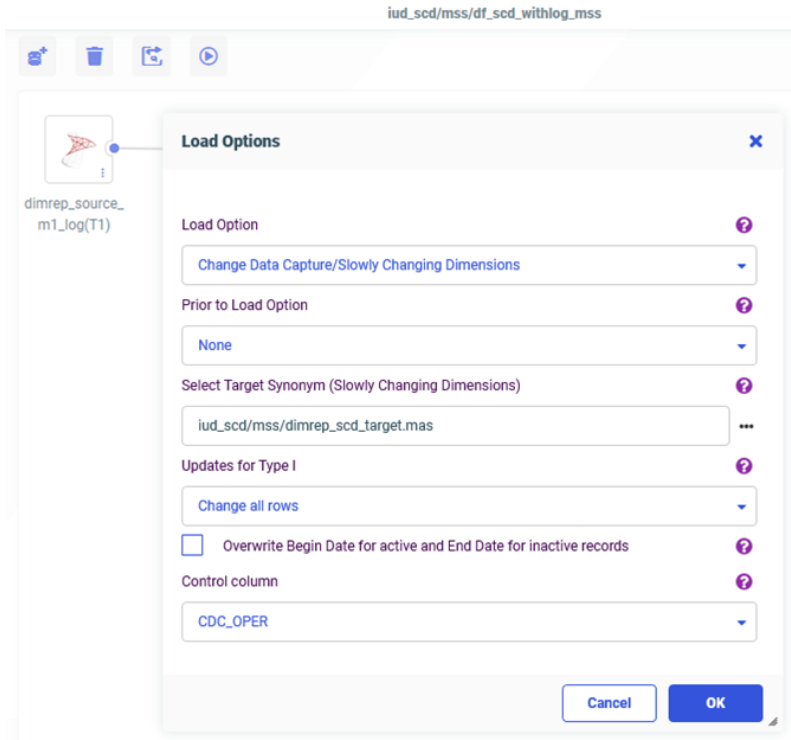
Using Change Data Capture/Slowly Changing Dimensions Load Type on a Flow

In the Web Console, the Load Option Change Data Capture/Slowly Changing Dimensions can be used to load a source file that supports Change Data Capture (CDC) to an existing target that has Slowly Changing Dimensions (SCD).

When you use a source log file for an existing relational table in a flow (the Access File contains the attribute DATA_ORIGIN=DBMSLOG), only records that have changed are passed to the load process.

When an existing target has Slowly Changing Dimensions (the Access File contains SCD_TYPE fields), Type I records are overwritten, and Type II records are processed using either the Activation Flag or End Date/Begin Date columns to track the history of the changed rows.

This load option is shown in the following image.



Supporting Direct Bulk Load for a Flow

Direct Bulk Load is a feature of DataMigrator and Data Flow that can significantly shorten load times when a source file is in the expected format of the bulk load program for the database to be loaded. The faster load times result from loading the file directly, without creating an intermediate file.

In prior releases, Direct Bulk Load could only be enabled as an adapter setting, meaning that all flows using that adapter would use it. This could be problematic if some input files met the requirements of a database and some did not.

Now you can enable or disable Direct Bulk Load on each data flow. However, in order for Direct Bulk Load to be enabled for a flow, Extended Bulk Load must be used and Optimize Load must be enabled.

When Extended Bulk Load is used and Optimize Load is enabled, data is loaded directly to the target database without the additional extract and transformation steps provided the following conditions are met:

- The data source is a delimited flat file.
- The file attributes including Header, Enclosure, Delimiter and Record Delimiter are supported by the target database.
- Only insert operations are used (no updates).
- No filters are used.
- No aggregation is specified.
- No transformations are used.

When Direct Bulk Load is used, the following message appears in the DataMigrator log:

`(ICM18637) Load Operation will use Direct Bulk Load Feature`

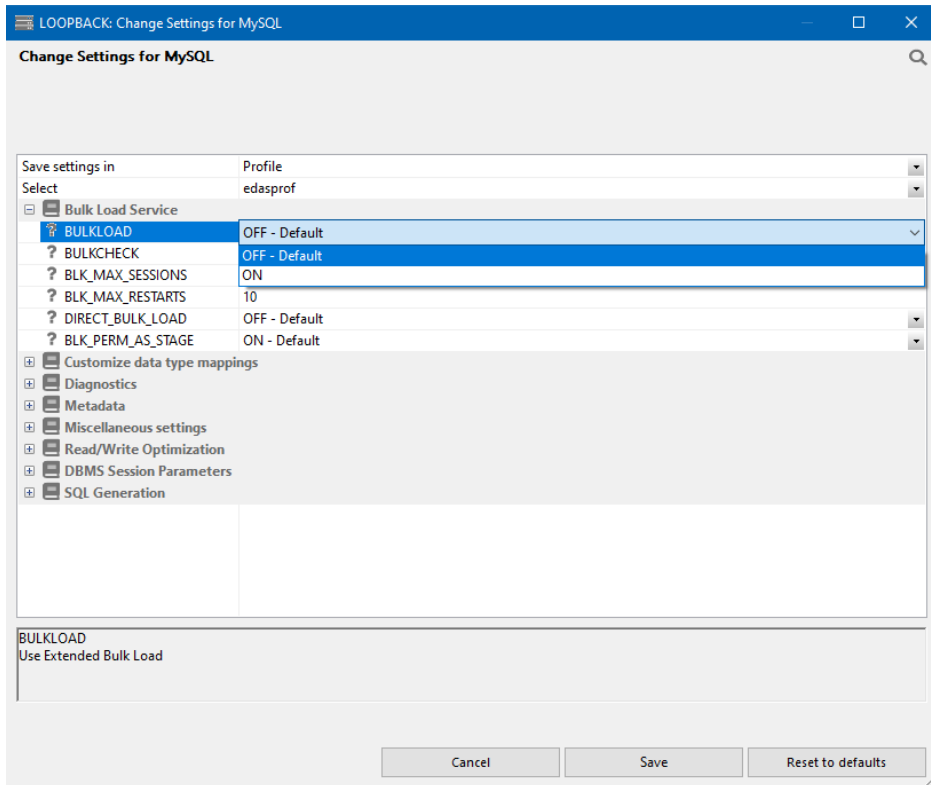
The following are the supported databases and data stores:

- Apache Hive
- Apache Impala
- Apache Spark
- Amazon RedShift
- ExaSQL
- Jethro
- MariaDB
- MS SQL Server
- MongoDB
- MySQL
- PostgreSQL
- Snowflake
- Salesforce.com

Sybase ASE

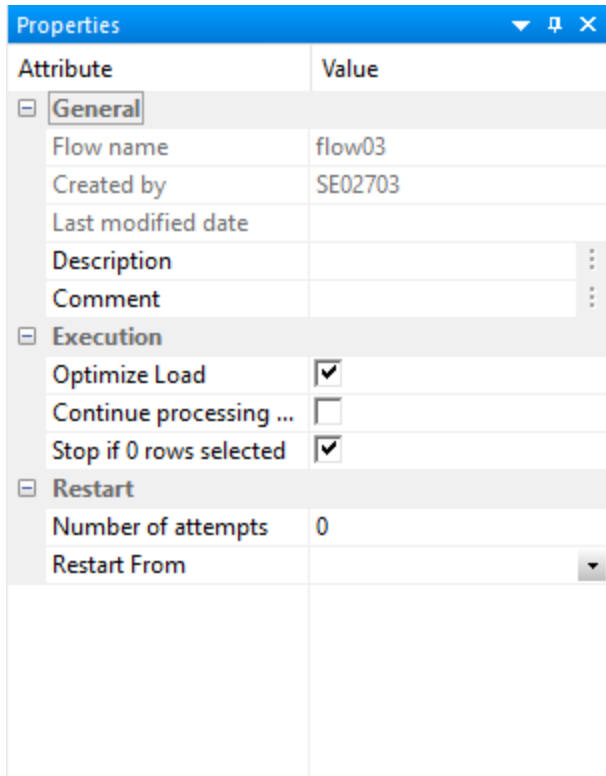
Vertica

To enable Extended Bulk Load, right-click an adapter that supports bulk load, and select Change Settings, as shown in the following image.

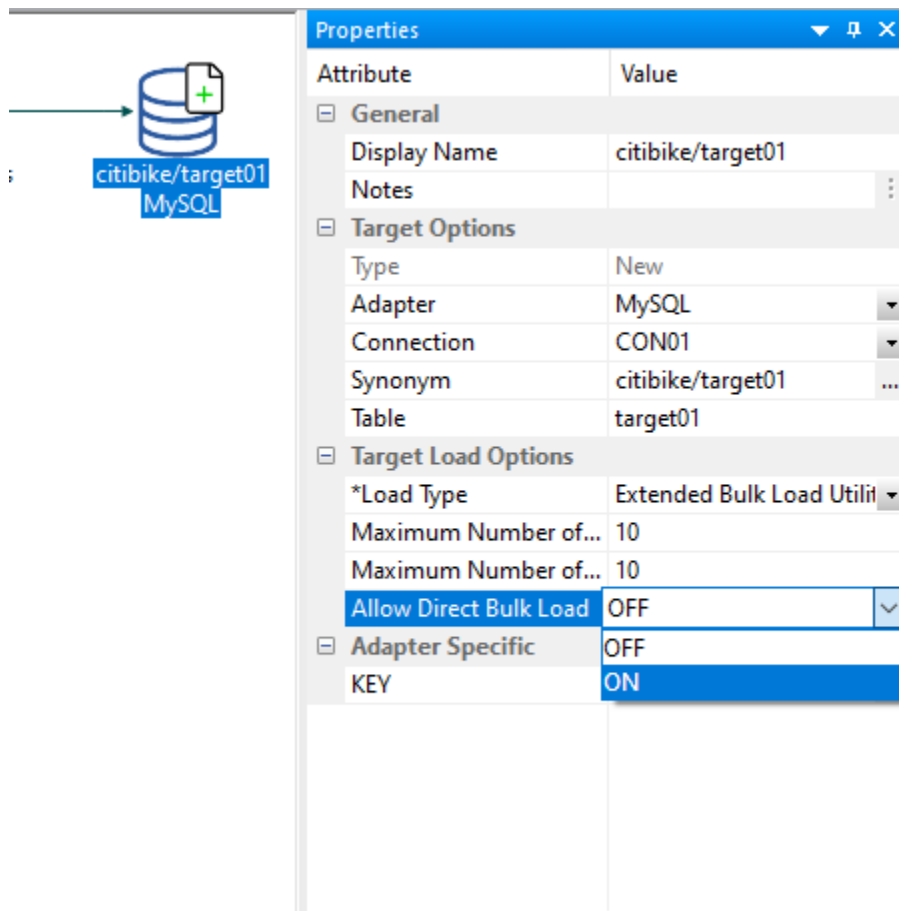


Select ON and click Save.

The following image shows Optimize Load enabled in the Flow Properties (right-click the canvas to the right of the SQL object and select Flow Properties).



The following image shows the option to enable Direct Bulk Load in the Load Options for a flow.

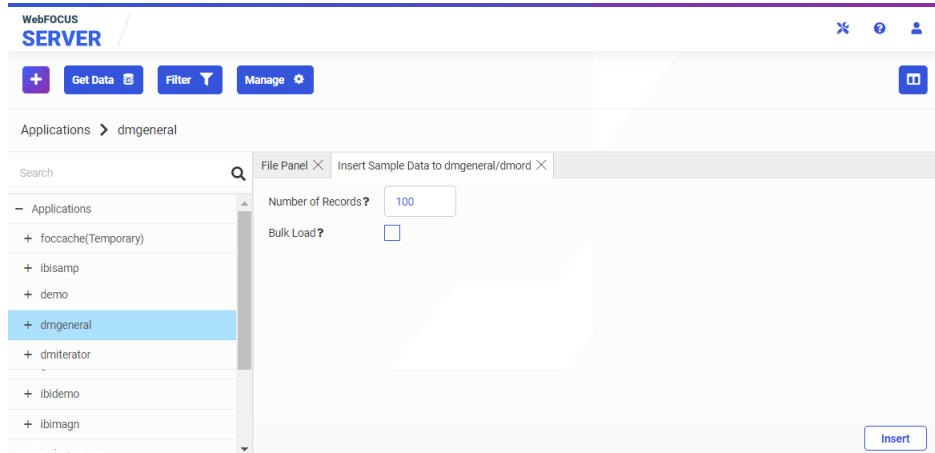


Data Profiling, Analysis, Sample Data, and Sampling Enhancements

This section describes new features for data profiling and analysis reports, sampling, and sample data.

Using Bulk Load to Insert Sample Data

On the DMC and the Server Web Console, *Insert Sample Data* is an option if you right-click a synonym and point to *Data Management*. If the synonym is for an adapter that supports Bulk Load, a *Bulk Load* check box is available in the Insert Sample Data dialog, as shown in the following image.



If you select the Bulk Load check box and click *Insert*, the sample data will be inserted using Bulk Load. This reduces load times when a large number of sample rows are inserted.

Enhanced Sampling on Joined Tables

When Representative Sampling is enabled, you can prioritize joins in the sampling by setting the new parameter `SMPL_PRIORITIZE_JOIN` to Yes.

When preparing data using very large data sets, the performance of many tasks can be greatly improved by working with a representative sample instead of the entire data set.

However, when data sets are joined, sampling can be problematic when not all values of a joined column are present in both samples.

For example, when joining from a (large) sales fact table to a foreign key in a customer dimension table, sampling each table separately may not include rows from the customer table that match every customer key in the sample of the fact table.

When the new `SMPL_PRIORITIZE_JOIN` parameter is enabled, the first (fact) table selected is sampled first. Then, when the sample of the second (dimension) table is taken, every customer that was selected in the fact table is included in the sample of the dimension table. This ensures that every row in the fact table is matched.

To prioritize joins when sampling, do the following:

1. Navigate to the *Workspace* page of the Web Console.
2. Click *Settings*, then *Settings for Web Console Preferences*.

The *Change Settings for Web Console Preferences* page opens.

3. Expand the *Data Assist (Representative Sampling)* section, as shown in the following image.

Data Assist (Representative Sampling)

ENABLE_SAMPLING ⓘ

 Controls Representative Sampling in Data Flow

SAMPLING_APPROOT ⓘ

 Application for representative samples. Default: foccache

SHOW_SAMPLING_INFO ⓘ

 Show Sampling Info

SMPL_ETL_TRG_DBMS ⓘ

 Use ETL-TRG-DBMS for Sampling

SMPL_PRIORITIZE_JOIN ⓘ

 Prioritize join when creating sampling

4. Select *On* for *ENABLE_SAMPLING*. This enables sampling wherever it is available.
5. Select *Yes* for *SMPL_PRIORITIZE_JOIN*.

This ensures that the sampling for a dimension table in a join will contain all of the key values from the sample rows generated for the fact table to which it is joined. The default value is *No*.

You can also change the following parameters.

SAMPLING_APPROOT

Is the application where the sampling files will be generated. The default is *foccache*.

SHOW_SAMPLING_INFO

The default value is *Yes*. This shows sampling information on the data flow canvas.

SMPL_TRG_DBMS

Use the adapter set as *ETL_TRG_DBMS* for sampling. The default value is *No*.

6. Click *Save*.

Key Analysis Using Sample Data

The Key Analysis report helps you select one or more fields that uniquely identify each record in a data set, so that these fields can be used as the key when loading the data to a relational database table.

Using Key Analysis, you select a synonym and then select fields that you think may be good candidates to become key fields. The server will then evaluate the fields, separately and in groups of two, three, and so on, to generate statistics that enable you to select an optimal set of key fields.


While using the entire data set ensures total accuracy in determining the uniqueness of the fields, for large data sets this could take a long time to process. Therefore, if you enable sampling in your Web Console preferences, a representative sample of the data will be used, enhancing performance while providing a reasonable assessment of the uniqueness of the selected fields.

After determining which fields should be key fields using the sample, rerun the analysis for those fields using the full set of data, in order to be sure that they uniquely identify each row.

The Key Analysis report is most useful when run against a data set that does not already have keys defined in the metadata. Most relational tables already have key definitions.

Procedure: How to Enable Sampling

1. Go to the Workspace page of the Web Console, click *Settings*, then click *Settings for Web Console Preferences*.
The *Change Settings for Web Console Preferences* page opens.
2. Expand the *Data Assist (Representative Sampling)* section, as shown in the following image.



3. Select *On* for *ENABLE_SAMPLING*.
This enables sampling everywhere it is available.

You can also change the following parameters. Not all of them are related to Key Analysis.

SAMPLING_APPROOT

Is the application where the sampling files will be generated. The default is foccache.

SHOW_SAMPLING_INFO

The default value is Yes. This shows sampling information on the data flow canvas.

SMPL_TRG_DBMS

Use the adapter set as ETL_TRG_DBMS for sampling. The default value is No.

SMPL_PRIORITIZE_JOIN

Prioritize the join when creating sampling. When this is set to Yes, the sampling for a dimension table in a join will contain all of the key values from the sample rows generated for the fact table to which it is joined. The default value is No.

4. Click Save.

Procedure: How to Run the Key Analysis Report

The samples shown in this section use a .csv file uploaded from the Citibike bike share web site that contains data about bike rides during a specific month.

1. Right-click a synonym, point to *Data Profiling*, and click *Key Analysis*.

The Key Analysis page opens, as shown in the following image.

Key Analysis for citibike/citibike_tripdata_csv

<input type="checkbox"/>	Segment Name	Display Field Name	Data Type	Nulls	Description
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	tripduration	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	start station id	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	end station id	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	bikeid	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	birth year	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	gender	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime	Date and Time	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	stoptime	Date and Time	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	usertype	Character	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Year	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Quarter	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Month	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Day	Integer	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Y	Date (Y)	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Y-Q	Date (YQ)	Yes	
<input type="checkbox"/>	CITIBIKE_TRIPDATA_CSVOUT	starttime,Y-M	Date (YM)	Yes	

Analyze (All Data) Analyze

2. Select the check boxes next to the columns you want to have analyzed, and click *Analyze*.

The key analysis report opens, showing statistics for all combinations of the columns you selected, as shown in the following image.

Key Analysis for "start station id"/"end station id"/bikeid/starttime/stoptime/usertype (CITIBIKE/CITIBIKE_TRIPDATA_CSV)

Segment	Name	Data Type	Elements	Count	Distinct Count	Distinct Percent	Duplicate Count	Duplicate Percent
CITIBIKE_TRIPDATA_CSVOUT	<u>start station id</u>	Integer	1	16466	777	4.72	15689	95.28
CITIBIKE_TRIPDATA_CSVOUT	<u>end station id</u>	Integer	1	16466	777	4.72	15689	95.28
CITIBIKE_TRIPDATA_CSVOUT	<u>bikeid</u>	Integer	1	16466	9096	55.24	7370	44.76
CITIBIKE_TRIPDATA_CSVOUT	<u>starttime</u>	Date and Time	1	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>stoptime</u>	Date and Time	1	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>usertype</u>	Character	1	16466	2	.01	16464	99.99
CITIBIKE_TRIPDATA_CSVOUT	<u>start station id/end station id</u>	Integer/Integer	2	16466	13715	83.29	2751	16.71
CITIBIKE_TRIPDATA_CSVOUT	<u>start station id/bikeid</u>	Integer/Integer	2	16466	16419	99.71	47	.29
CITIBIKE_TRIPDATA_CSVOUT	<u>start station id/starttime</u>	Integer/Date and Time	2	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>start station id/stoptime</u>	Integer/Date and Time	2	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>start station id/usertype</u>	Integer/Character	2	16466	1421	8.63	15045	91.37
CITIBIKE_TRIPDATA_CSVOUT	<u>end station id/bikeid</u>	Integer/Integer	2	16466	16432	99.79	34	.21
CITIBIKE_TRIPDATA_CSVOUT	<u>end station id/starttime</u>	Integer/Date and Time	2	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>end station id/stoptime</u>	Integer/Date and Time	2	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>end station id/usertype</u>	Integer/Character	2	16466	1403	8.52	15063	91.48
CITIBIKE_TRIPDATA_CSVOUT	<u>bikeid/starttime</u>	Integer/Date and Time	2	16466	16466	100.00	0	.00
CITIBIKE_TRIPDATA_CSVOUT	<u>bikeid/stoptime</u>	Integer/Date and Time	2	16466	16466	100.00	0	.00

Clicking a column heading sorts on that column, toggling between ascending and descending order.

The following columns are included in the report.

Segment

Specifies the segment in the synonym that contains the field.

Name

Is the field name or title of the field or combination of fields, depending on your name display strategy.

You can right-click a row in the Name column and select one of the following:

- Duplicate values.** Shows the count and percent of duplicate values for each combination of field values.
- Duplicate values chart.** Available for single fields only. Displays a horizontal bar chart showing the count of duplicate values for each field value.
- Distribution chart.** Available for single fields only. Displays a vertical bar chart showing the count of values for each range of field values.
- Duplicate values pie chart.** Available for single fields only. Displays a pie chart showing the count of duplicate values for each field value.

Data Type

Is the data type or combination of data types for the fields listed on each row.

Elements

Is the number of fields included in the analysis for each row.

Count

Is the number of instances of the field or combination of fields for each row, based on the sample.

Distinct Count

Is the number of distinct instances of the field or combination of fields for each row, based on the sample.

Distinct Percent

Is the percent of distinct instances of the field or combination of fields for each row, based on the sample.

Duplicate Count

Is the number of duplicate instances of the field or combination of fields for each row, based on the sample.

You can right-click a row in the Duplicate Count column and select one of the following:

- Duplicate values.** Shows the count and percent of duplicate values for each combination of field values.
- Duplicate values chart.** Available for single fields only. Displays a horizontal bar chart showing the count of duplicate values for each field value.
- Distribution chart.** Available for single fields only. Displays a vertical bar chart showing the count of values for each range of field values.
- Duplicate values pie chart.** Available for single fields only. Displays a pie chart showing the count of duplicate values for each field value.

Duplicate Percent

Is the percent of duplicate instances of the field or combination of fields for each row, based on the sample.

3. Look for fields or combinations of fields on the report that provide 100% distinct values.

In order to be a key, the combination of fields must provide 100% distinct values. Select the smallest combination of fields that provide distinct values and that are good candidates to be key fields.

- If no combination of fields provides 100% distinct values, either add additional fields to the Key Analysis or create a field (such as a row number) that will provide a distinct value for each row.
- If multiple fields or combinations of fields on the report provide 100% distinct values, use your knowledge of the data to select a good combination.

For example, in the sample report shown, both the starttime and endtime fields provide 100% distinct values, but that may be an anomaly due to the sample chosen. There is always the possibility that multiple bike rides will start or end at the same time. Similarly, the combination of starttime and endtime may not always provide unique values. The combination of bikeid with either starttime or endtime is probably unique, as the same bike could not be used for multiple rides at the same exact date and time. Also, combinations that include more than two fields are not as desirable as combinations that include only two fields.

4. Once you have chosen a key, rerun the Key Analysis report, selecting only those candidate fields, and this time clicking *Analyze (All Data)*.

Check to make sure that this selection provides a unique key for all of the data, not just the sample. In this example, the combination bikeid/starttime provides a unique key, as shown in the following image.

Key Analysis for bikeid/starttime (CITIBIKE/CITIBIKE_TRIPDATA_CSV)

Segment	Name	Data Type	Elements	Count	Distinct Count	Distinct Percent	Duplicate Count	Duplicate Percent
CITIBIKE_TRIPDATA_CSVOUT	<u>bikeid</u>	Integer	1	2181064	14356	.66	<u>2166708</u>	99.34
CITIBIKE_TRIPDATA_CSVOUT	<u>starttime</u>	Date and Time	1	2181064	2179730	99.94	<u>1334</u>	.06
CITIBIKE_TRIPDATA_CSVOUT	<u>bikeid/starttime</u>	Integer/Date and Time	2	2181064	2181064	100.00	<u>0</u>	.00

Using Stratified Sampling

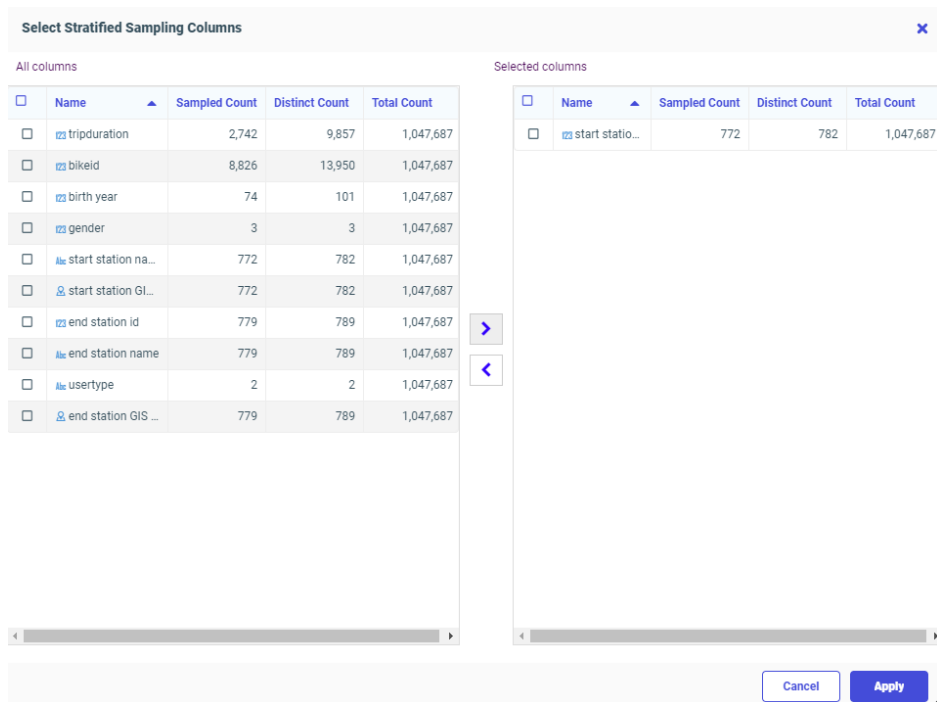
When preparing data and viewing the distribution profile of the data, Representative Sampling can be used. This provides better responsiveness, especially for large data sets.

However, there may be times when you want to ensure that the sample includes every distinct value of one or more fields. For example, if a file has a field with the value State, and some states have very few records, a representative sample may not include all of them.

With Stratified Sampling, you can select one or more fields, and the sampling process will include records with every unique value of the selected fields, as long as the number of values does not exceed the sample size.


When sampling is enabled, this option can be selected by right-clicking a source node and clicking *Recreate Stratified Sampling*. The Stratified Sampling dialog box opens. Select the fields to use from a list of available fields, and move them to the list of selected fields using the arrow buttons. The list shows the count of unique and total values for each field. Key fields are omitted from this list, as are floating point fields and timestamp fields with precision of less than a second, as they cannot be used.

The following image shows that the `start_station_id` field is selected for stratified sampling. This ensures that every start station will be represented in the sample, even those that are rarely used and may be omitted from a random sample.



When the required fields are selected, click *Apply*.


When Stratified Sampling is applied, the label of the node says *Stratified Sampling*, as shown in the following image.



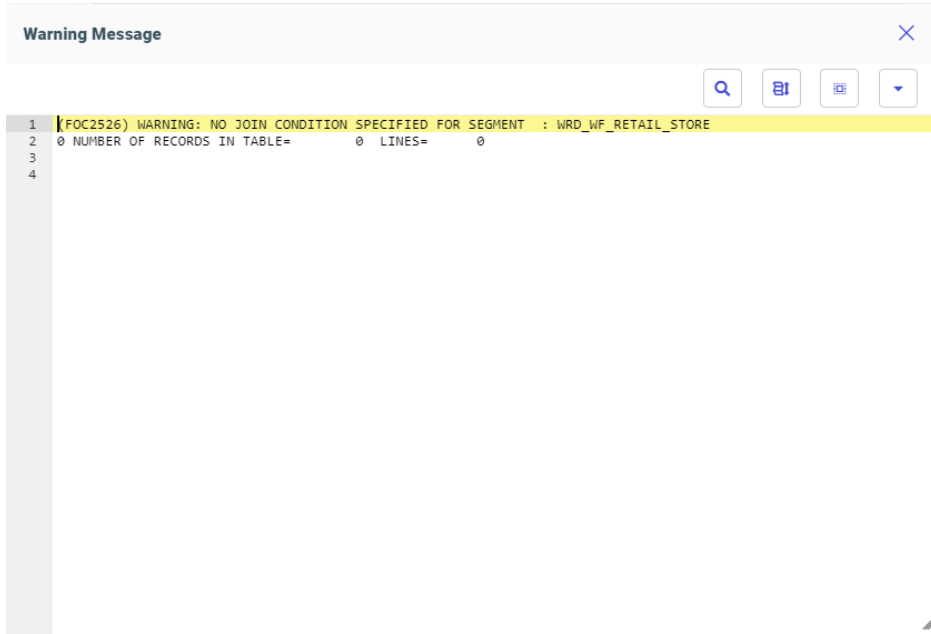
To remove stratified sampling and return to representative sampling, right-click the node and click *Recreate Sampling*.

Sample Data Shows Warning Messages

When generating sample data in the Synonym Editor or a Data Flow, if problems occur reading the data, an indicator that there are messages displays on the screen, as shown in the following image.




If you click this indicator, you can see the actual messages, as shown in the following image.

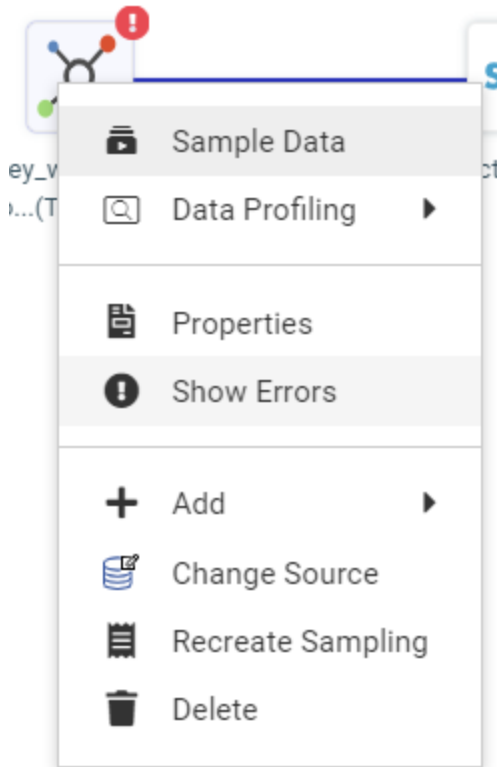


Sampling Shows Warning Messages

When sampling is enabled in a data flow and any errors occur reading the data, a warning indicator displays on the synonym node that has the sampling errors. Hovering over the indicator shows the messages, or a partial list, depending on the length and number of the messages, as shown in the following image.



To open a window with the complete list of messages, you can right-click the synonym and click *Show Errors*, as shown in the following image.



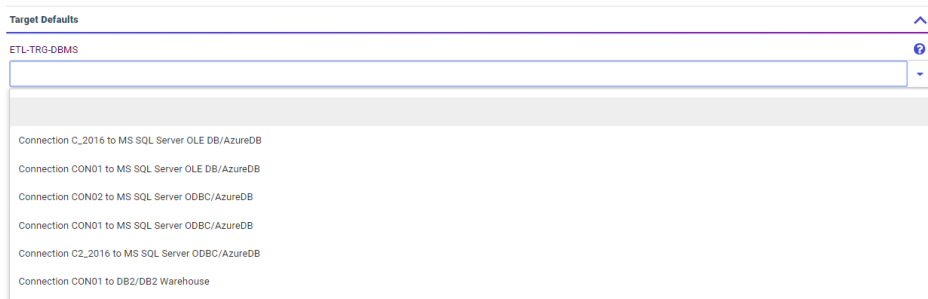
A window opens showing the complete list of messages.

The Reports pane also shows the indicator.

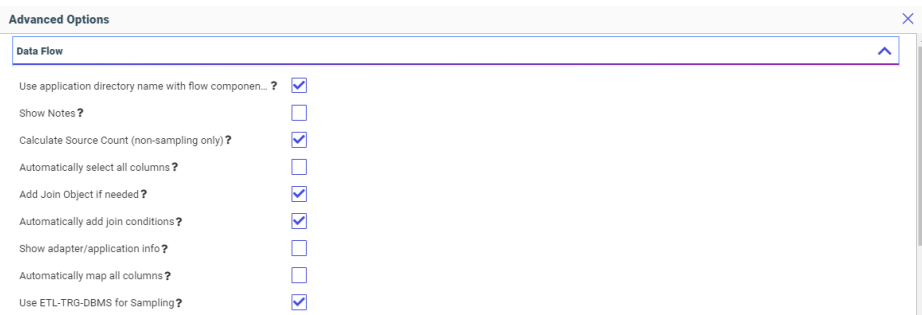
If you click that indicator, the message window opens.

Specifying a Location for Sampling Data


When representative sampling was enabled in prior releases, the samples would be stored in a binary file on disk. Now, the samples can also be stored in relational database tables using an adapter and connection you specify as the default target for flows. You can set a target for data flows on the Workspace page by clicking *Settings*, then *Settings for Web Console Preferences*, then *Target Defaults*. You can then select a connection and target DBMS for the ETL-TRG-DBMS parameter, as shown in the following image.



Then, in the Data Flow parameters of the Advanced Options dialog box, you can check *Use ETL-TRG-DBMS for Sampling*, in order to use the same target for the samples, as shown in the following image.



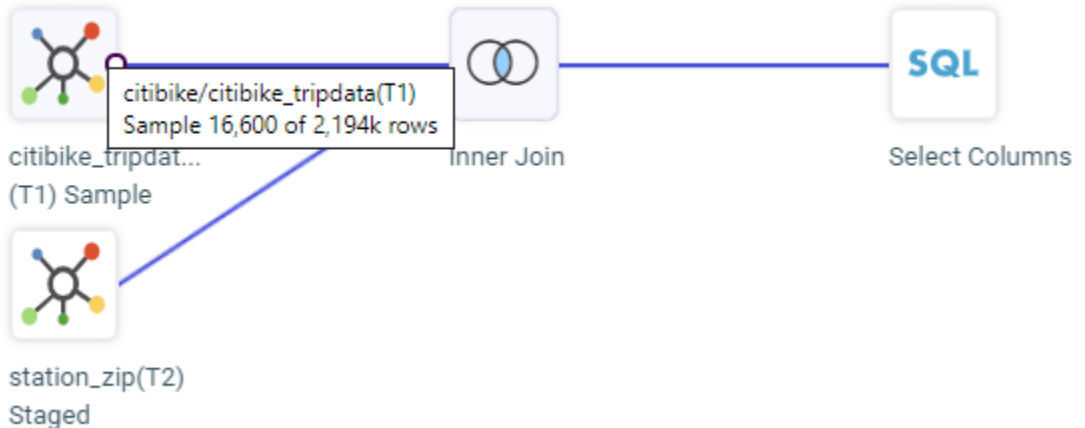
When you hover over the source icons on the canvas, the location of the sample data is displayed, as shown in the following image.




Staging All Sources When Sampling in a Data Flow

When sampling is enabled for a flow, even sources that are too small to support sampling are staged in the same staging target as the large sources, to eliminate cross joins and, therefore, improve join performance.

In the data flow shown in the following image, a sample of the first data source has been staged, as shown in the following image.



The entire second data source has been staged, as shown in the following image.

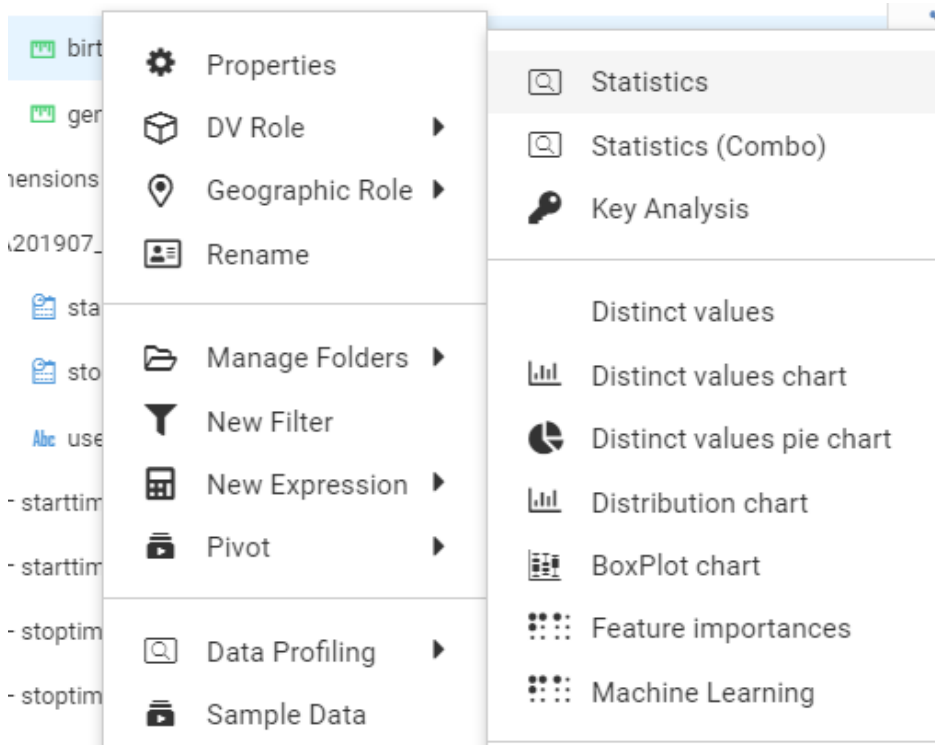


By default, sources are staged as DATREC files. However, you can configure them to be staged in the same relational data source as the target using the `SMPL_ETL_TRG_DBMS` parameter in *Settings for Web Console Preferences* on the Workspace page Settings menu.


Profiling Numeric Data Using Standard Deviation

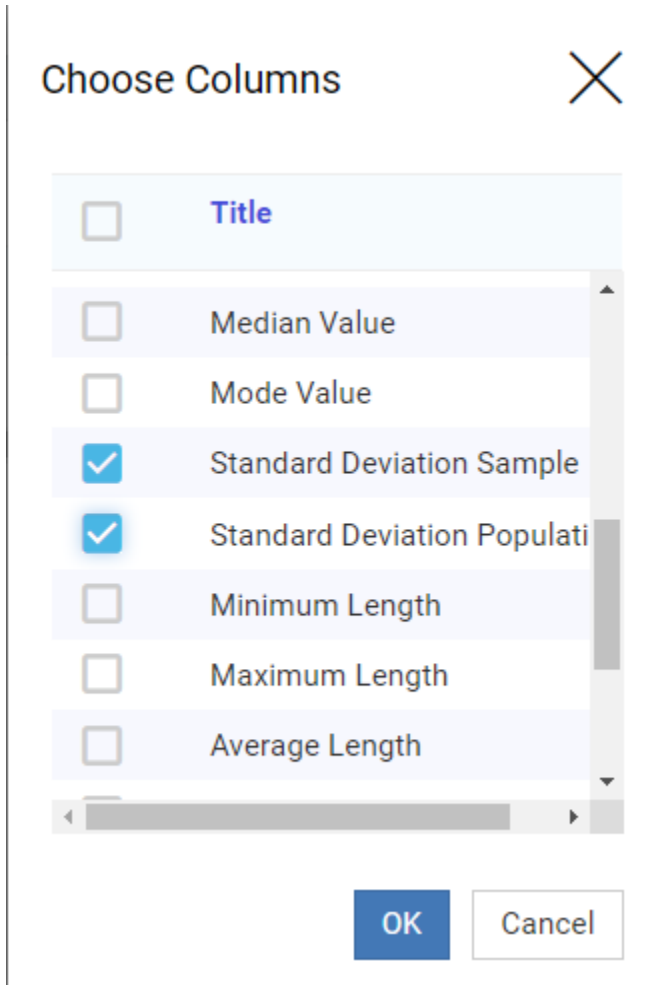
When using the Synonym or Flow editor and selecting a numeric field, two new data profiling statistics are available, Standard Deviation (Population) and Standard Deviation (Sample).

Right-click a numeric field, click *Data Profiling*, then *Statistics*, as shown in the following image.

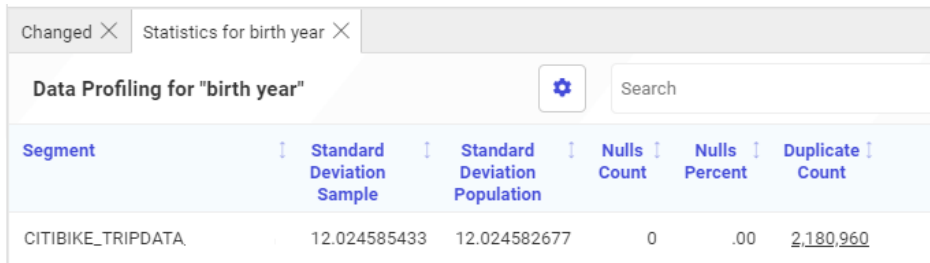


Each statistic is represented as a column on the resulting data profiling report. Some statistics are selected by default. You can control which statistics to display.

Click the *Choose Columns* gear icon () to select the columns for the report, as shown in the following image.



When you have made your selections, click *OK* to display them on the data profiling report, as shown in the following image.



The screenshot shows a data profiling report for the 'birth year' column. The report is titled 'Data Profiling for "birth year"' and includes a search bar and a settings icon. The data is presented in a table with the following columns: Segment, Standard Deviation Sample, Standard Deviation Population, Nulls Count, Nulls Percent, and Duplicate Count. The data row shows the following values: CITIBIKE_TRIPDATA, 12.024585433, 12.024582677, 0, .00, and 2,180,960.

Segment	Standard Deviation Sample	Standard Deviation Population	Nulls Count	Nulls Percent	Duplicate Count
CITIBIKE_TRIPDATA	12.024585433	12.024582677	0	.00	2,180,960

SQL Statement and Functions Enhancements

This section describes new features for SQL statements and SQL Translator functions.

Returning an Ordered Answer Set of a Limited Number of Rows in a Sub-Select

A sub-select can now contain the `TOP n`, `ORDER BY`, and `FETCH FIRST n ROWS ONLY` phrases to return an ordered answer set with a limited number of rows.

The syntax is:

```
ORDER BY column ... FETCH FIRST n ROWS ONLY
```

or:

```
SELECT TOP n ... ORDER BY column
```

where:

column

Is a column in the subquery used to sort the rows. The column can be expressed as an index number, column name, or AS name.

n

Is the number of rows in the sub-select to retrieve.

Example: Using FETCH FIRST n ROW ONLY in a Subquery

The following request retrieves the first 5 rows from dmsale, ordered by plant in descending order, and then retrieves rows from dminv that have plant values retrieved from the sub-select.

```
SQL
SELECT
    order_num,
    order_date,
    plant
FROM
    dmord

WHERE
    plant IN (
        SELECT
            plant
        FROM
            dmsale

ORDER BY
1 DESC
FETCH FIRST 5 ROWS ONLY
    )
;
TABLE
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

The output is shown in the following image.

<u>Order Number</u>	<u>Date of Order</u>	<u>Plant Location</u>
93677	2004/03/30	STL
93705	2005/02/07	STL
93706	2006/02/14	STL
93708	2005/03/08	STL
93710	2005/02/14	STL
93711	2005/02/14	STL
93712	2005/02/14	STL
93713	2005/03/06	STL
93714	2005/02/23	STL
93717	2005/02/14	STL
93718	2006/02/14	STL
93719	2005/02/14	STL
93727	2005/02/14	STL
93729	2005/02/09	STL
93734	2005/02/14	STL
93743	2005/02/14	STL
93753	2006/02/22	STL
93755	2005/02/15	STL
93770	2005/02/15	STL
93774	2006/02/14	STL
93779	2006/02/14	STL
93780	2006/02/14	STL
93791	2006/02/14	STL
93793	2005/02/14	STL
93797	2005/02/14	STL
93812	2005/02/20	STL
93816	2006/03/06	STL
93822	2005/03/06	STL
93825	2005/03/06	STL
93829	2006/02/20	STL
93843	2006/02/20	STL
93847	2005/02/20	STL
93848	2006/02/20	STL
93855	2006/02/20	STL
93856	2005/02/15	STL
93858	2005/02/15	STL
93860	2005/02/15	STL

If the ORDER BY phrase had been ascending instead of descending, the plant location would have been BOS.

Example: **Using SELECT TOP n in a Subquery**

The following request retrieves the top 2 plant values from dmsale, ordered by plant in ascending order, and then retrieves rows from dminv that have plant values retrieved from the sub-select.

```
SQL
SELECT
    order_num,
    order_date,
    plant
FROM
    dmord

WHERE
    plant IN (
        SELECT TOP 2
            plant
        FROM
            dmsale
ORDER BY 1
    )
;
TABLE
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

The output is shown in the following image.

<u>Order Number</u>	<u>Date of Order</u>	<u>Plant Location</u>
93674	2005/02/15	BOS
93675	2005/02/15	BOS
93693	2006/02/10	BOS
93700	2006/02/14	BOS
93715	2005/02/10	BOS
93730	2006/02/14	BOS
93702	2006/02/14	BOS
93744	2006/02/14	BOS
93742	2005/02/14	BOS
93754	2005/02/15	BOS
93775	2005/02/15	BOS
93776	2005/02/20	BOS
93777	2006/02/20	BOS
93787	2005/02/14	BOS
93788	2005/02/16	BOS
93769	2006/02/20	BOS
93796	2005/02/14	BOS
93803	2005/03/06	BOS
93805	2006/02/20	BOS
93818	2005/02/20	BOS
93839	2005/02/20	BOS
93849	2005/02/20	BOS
93850	2005/02/20	BOS
93853	2005/02/15	BOS
93861	2005/02/15	BOS

If the ORDER BY phrase had been descending instead of ascending, the plant location would have been STL.

Support for GROUP BY On Column Position or AS Name

In prior releases, a GROUP BY clause in an SQL SELECT statement could only use a column name or expression. Now, the GROUP BY clause can also use the column position (such 1 to indicate the first column) or an AS name from the SELECT list.

This syntax is commonly generated by some third-party tools and can also be used in DataMigrator in a custom SQL flow, or in WebFOCUS in an SQL query.

Example: Using an AS Name in a GROUP BY Clause

The following SQL request uses the AS name CATEGORY in the GROUP BY clause.

```
SQL
SELECT
PRODCAT AS CATEGORY ,
COUNT(QTY_IN_STOCK)
FROM
DMINV T1
GROUP BY CATEGORY
;
TABLE
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

The output is shown in the following image.

<u>CATEGORY</u>	<u>Quantity in Stock</u>
Camcorders	6
Cameras	4
CD Players	1
Digital Tape Recorders	1
DVD	2
PDA Devices	2
VCRs	1

Example: Using a Column Position in a GROUP BY Clause

The following SQL request uses the column position 1 in the GROUP BY clause.

```
SQL
SELECT
PRODCAT,
COUNT(QTY_IN_STOCK)
GROUP BY 1
FROM
DMINV T1
GROUP BY 1
;
TABLE
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

The output is shown in the following image.

<u>Product</u> <u>Category</u>	<u>Quantity</u> <u>in Stock</u>
Camcorders	6
Cameras	4
CD Players	1
Digital Tape Recorders	1
DVD	2
PDA Devices	2
VCRs	1

Using SQL Analytic Functions

SQL Analytic functions compute an aggregate value based on a group of rows, called a partition. They return multiple rows for each group. For each row, a range of surrounding rows can be defined, called a window. The calculation for the current row is performed using the defined window, as long as the window stays within the partition. The order of rows within each partition is controlled by an optional ORDER BY clause which affects both which rows are included in the window and the result value. The calculations restart when a new partition starts. If a window is not defined, the entire partition is used in the calculation. Analytic functions can appear only in the SELECT list or ORDER BY clause.

The following SQL Analytic functions have been added.

- ❑ **AVG.** Calculates the average over a group of rows defined by a partitioning expression, an ORDER BY clause, and a sliding window.
- ❑ **COUNT.** Calculates a count over a group of rows defined by a partitioning expression, an ORDER BY clause, and a sliding window.
- ❑ **DENSE_RANK.** Calculates a dense rank (each rank number is the next sequential integer, even when the same rank is assigned to multiple data values) over a group of rows defined by a partitioning expression and an ORDER BY clause.
- ❑ **FIRST_VALUE.** Retrieves the first result from an ordered set of rows defined by a partitioning expression, an ORDER BY clause, and a sliding window.
- ❑ **LAG.** Retrieves the value from a previous row defined by a partitioning expression, an ORDER BY clause, an offset, and a default.
- ❑ **LAST_VALUE.** Retrieves the last result from an ordered set of rows defined by a partitioning expression, an ORDER BY clause, and a sliding window.
- ❑ **LEAD.** Retrieves the value from a following row defined by a partitioning expression, an ORDER BY clause, an offset, and a default.
- ❑ **MAX.** Calculates the maximum over a group of rows defined by a partitioning expression, an ORDER BY clause, and a sliding window.
- ❑ **MEDIAN.** Calculates the median over a group of rows defined by a partitioning expression, an ORDER BY clause, and a sliding window.
- ❑ **MIN.** Calculates the minimum over a group of rows defined by a partitioning expression, an ORDER BY clause, and a sliding window.
- ❑ **MODE.** Calculates the mode over a group of rows defined by a partitioning expression, an ORDER BY clause, and a sliding window.
- ❑ **PERCENT_RANK.** Calculates a percent rank over a group of rows defined by a partitioning expression and an ORDER BY clause,.
- ❑ **RANK.** Calculates a sparse rank (each rank number is the integer value of the prior rank plus the number of values assigned to the prior rank) over a group of rows defined by a partitioning expression and an ORDER BY clause,.
- ❑ **ROW_NUMBER().** Lists the current row number over a group of rows defined by a partitioning expression and an ORDER BY clause,.

- ❑ **STDDEV_POP.** Calculates the standard deviation of the population over a group of rows defined by a partitioning expression, an ORDER BY clause, and a sliding window.
- ❑ **STDDEV_SAMP.** Calculates the standard deviation of a sample over a group of rows defined by a partitioning expression, an ORDER BY clause, and a sliding window.
- ❑ **SUM.** Calculates the sum over a group of rows defined by a partitioning expression, an ORDER BY clause, and a sliding window.

The syntax varies by function, and may vary depending on the SQL supported for the database being used. The basic structure of the syntax follows.

```
analytic_function_name ( [ argument_list ] )
  OVER (
    [ PARTITION BY partition_expression_list ]
    [ ORDER BY expression [{ ASC | DESC } ] [, ...] ]
    [ window_frame_clause ]
  )
```

where:

```
analytic_function_name ( [ argument_list ] )
```

Is the name of the function and its calling arguments. Can be one of the following:

```
AVG(expression)
COUNT(expression)
DENSE_RANK()
FIRST_VALUE(expression)
LAG(expression, offset, default)
LAST_VALUE(expression)
LEAD(expression, offset, default)
MAX(expression)
MIN(expression)
MEDIAN(expression)
MODE(expression)
PERCENT_RANK()
RANK()
ROW_NUMBER()
STDDEV_POP(expression)
STDDEV_SAMP(expression)
SUM(expression)
```

```
PARTITION BY partition_expression_list
```

Divides the rows into partitions.

```
ORDER BY expression
```

Specifies the row order within each partition.

window_frame_clause

Defines the sliding window within each partition (starting row and ending row for the window). The window frame clause defines a frame around the current row within a partition over which the analytic function is evaluated. Both physical window frames (defined by ROWS) and logical window frames (defined by RANGE) are allowed. It is your responsibility to know the syntax for your environment. The ordering of rows within the partition (ORDER BY clause) affects which rows are included in the sliding window.

The window frame clause is not supported for functions DENSE_RANK, LAG, LEAD, PERCENT_RANK, RANK, and ROW_NUMBER.

Basic syntax for the window frame clause follows:

```
{ROWS|RANGE}
{
  {UNBOUNDED PRECEDING|numeric_expression PRECEDING|CURRENT ROW} |
  {BETWEEN boundary_start AND boundary_end}
}
```

The basic syntax for the start of the boundary is:

```
{UNBOUNDED PRECEDING|numeric_expression PRECEDING|CURRENT ROW}
```

The basic syntax for the end of the boundary is:

```
{UNBOUNDED FOLLOWING|numeric_expression {PRECEDING|FOLLOWING} |
CURRENT ROW}
```

Example: Returning the Maximum Trip Duration

The following call to the MAX SQL analytic function partitions the rows by zip code and returns the maximum trip duration within the partition windowed by the preceding 1 row, the current row, and the following 1 row, ordered by the start station name.

```
SQL
SELECT
  T2.STATION_ID ,
  T2.ZIP_CODE ,
  T1.TRIPDURATION ,
  MAX(T1.TRIPDURATION ) OVER (PARTITION BY T2.ZIP_CODE
    ORDER BY T1.START_STATION_NAME
    ROWS BETWEEN 1 PRECEDING AND 1 FOLLOWING) AS MAXDURATION ,
  T1.START_STATION_NAME
FROM
  (station_zip T2
  INNER JOIN
  citibike_tripdata T1
  ON
  T2.STATION_ID = T1.START_STATION_ID )
  WHERE T2.ZIP_CODE IS NOT NULL
  ORDER BY T2.ZIP_CODE DESC
;
TABLE
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

The following image shows partial output.

<u>STATION_ID</u>	<u>ZIP_CODE</u>	<u>tripduration</u>	<u>MAXDURATION</u>	<u>start station name</u>
343	11251	1671	1671	Clinton Ave & Flushing Ave
3092	11249	413	413	Berry St & N 8 St
3092	11249	204	1588	Berry St & N 8 St
3092	11249	1588	1588	Berry St & N 8 St
3092	11249	608	1588	Berry St & N 8 St
3092	11249	422	608	Berry St & N 8 St
3092	11249	402	1409	Berry St & N 8 St
3092	11249	1409	1409	Berry St & N 8 St
3092	11249	566	1409	Berry St & N 8 St
3092	11249	687	781	Berry St & N 8 St
3092	11249	781	781	Berry St & N 8 St
389	11249	274	781	Broadway & Berry St
389	11249	397	397	Broadway & Berry St

Row 1 is in its own partition, as the zip code changes in row 2, so MAXDURATION is the same as tripduration.

For Row 2, there is no preceding row within the partition, so MAXDURATION is the maximum of tripduration for rows 2 and 3, which is 413.

For Row 3, the preceding and following rows are within the partition, so MAXDURATION is the maximum of tripduration for rows 2, 3, and 4, which is 1588.

Data Migrator Enhancements

This section describes enhancements for Data Migrator and the Data Management Console.

Assigning Slowly Changing Dimension (SCD) Columns in the Web Console

A Slowly Changing Dimension (SCD) is a dimension in which the data does not change frequently or on a regular schedule. Slowly Changing Dimensions are common in data warehouses.

When a Synonym has SCD columns defined, the changes to the records in the database can be handled following Type I and Type II (Type 1 and Type 2) implementation techniques. The SCD implementation techniques are applied when such a Synonym is being used as a target in a Data Flow, and the Load Option on this target is set to *Slowly Changing Dimensions*.

For example, customer addresses change only when a customer moves. A customer who moves may change spending habits at the new address, so tracking customer data for that customer may depend on the address. Therefore, it could be useful to keep historical records of those changes.

Other types of changes may not require keeping a record of the change.

WebFOCUS supports two methodologies for handling Slowly Changing Dimensions.

- ❑ **Type I Processing.** SCD Type I processing does not store historical data in the dimension table. This method overwrites the old data in the dimension table with the new data. One common use for this is to correct data errors in the dimension, for example, to correct a misspelling. This type of processing is easy to maintain, and it does not add additional rows to the data.

- ❑ **Type II Processing.** SCD Type II processing stores the entire history the data in the dimension table. With this type of processing, unlimited history can be tracked. With Type II processing, WebFOCUS supports storing the data in two different ways, which can be used separately or in combination with each other. They are:
 - ❑ **Flagging.** With flagging, an activation flag column in the data identifies whether a row is active (current) or inactive (historical). Every time a change is made, a new row is created, with the flag value identifying it as active (for example, the value 1), and the record that was just replaced has its flag column changed to the inactive value (for example, 0).
 - ❑ **Effective Date.** With effective date processing, the period of the change is tracked using start_date and end_date columns in the dimension table. The current row has NULL in the end_date column.

You can assign the required SCD column values in the Web Console.

Reference: Summary of SCD Columns

The following columns define SCD processing.

Surrogate Key (required)

Is an Integer column that uniquely identifies the row in the dimension. The Surrogate Key column must be assigned.

Logical Key (required)

Is the Source database key, which can consist of multiple columns. The Logical Key column must be assigned.

Type I

Defines columns whose values are overwritten with new values, no history preserved. Either Type I and/or Type II columns must be assigned.

Type II

Defines columns whose older values are preserved with inactive status, and whose new values are added with active status, in order to preserve the history. Either Type I and/or Type II columns must be assigned.

Begin Date

Is the Date or Date-Time column used to set the time when the row becomes the active row. DataMigrator sets a default value.

The format of this column should be identical with the format of the End Date column.

The Begin/End Date pair and/or Activation Flag mark the active row. Either the Begin/End Date columns pair and/or Activation Flag column must be assigned.

End Date

Is the Date or Date-Time column used to set the time when the row becomes an inactive row. The user can set a value, otherwise DataMigrator sets a default value.

The format of this column should be identical with the format of the Begin Date column.

The Begin/End Date pair and/or Activation Flag mark the active row. Either the Begin/End Date columns pair and/or Activation Flag column must be assigned.

Activation Flag

Is an Integer or Alphanumeric column that indicates if the row is current. The user can set a value, otherwise DataMigrator sets a default value.

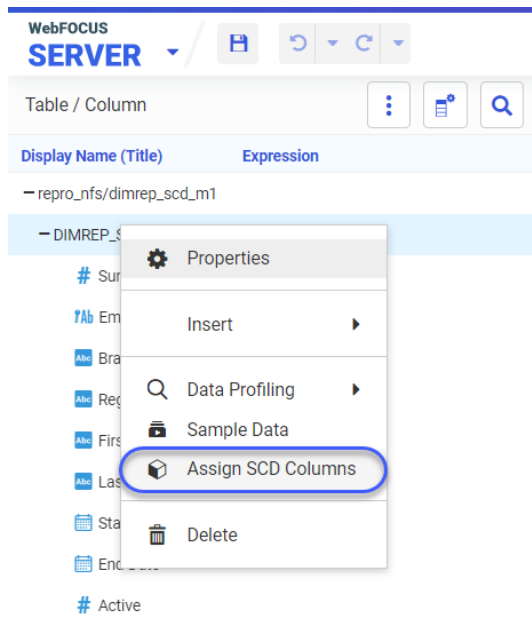
The Begin/End Date pair and/or Activation Flag mark the active row. Either the Begin/End Date columns pair and/or Activation Flag column must be assigned.

Change Flag

Is a Date or Date-Time column that is used to set the time when the value of a Type I column changes. The user can set a value, otherwise DataMigrator sets a default value.

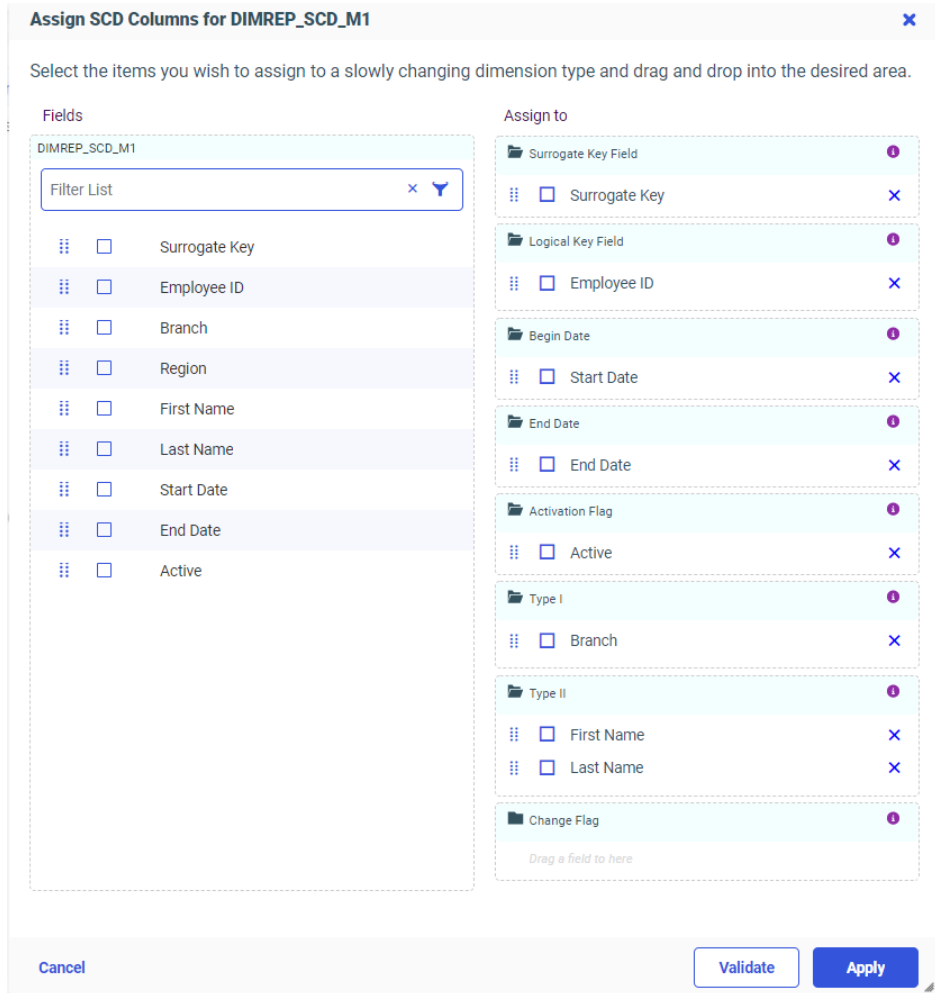
Procedure: How to Assign SCD Columns

1. Open the synonym that has SCD columns in the Synonym Editor, by double-clicking the synonym, or by right-clicking and selecting *Open*.
2. Right-click the segment that has the SCD columns, and click *Assign SCD Columns*, as shown in the following image.



The Assign SCD Columns pane opens.

3. Drag the appropriate fields from the segment fields list to the SCD columns in the *Assign to* pane, as shown in the following image.



You can click the Information icon in each row, if you require help with the specific column, as shown in the following image.

Assign SCD Columns for DIMREP_SCD_M1 ✕

Select the items you wish to assign to a slowly changing dimension type and drag and drop into the desired area.

Fields

DIMREP_SCD_M1

Filter List ✕ ▾

<input checked="" type="checkbox"/>	Surrogate Key
<input type="checkbox"/>	Employee ID
<input type="checkbox"/>	Branch
<input type="checkbox"/>	Region
<input type="checkbox"/>	First Name
<input type="checkbox"/>	Last Name
<input type="checkbox"/>	Start Date
<input type="checkbox"/>	End Date
<input type="checkbox"/>	Active

Assign to

Surrogate Key Field

Surrogate Key ✕

Logical Key Field

Employee ID ✕

Begin Date

Start Date ✕

End Date

End Date ✕

Activation Flag

Active ✕

Type I

Branch ✕

Type II

First Name ✕

Last Name ✕

Change Flag

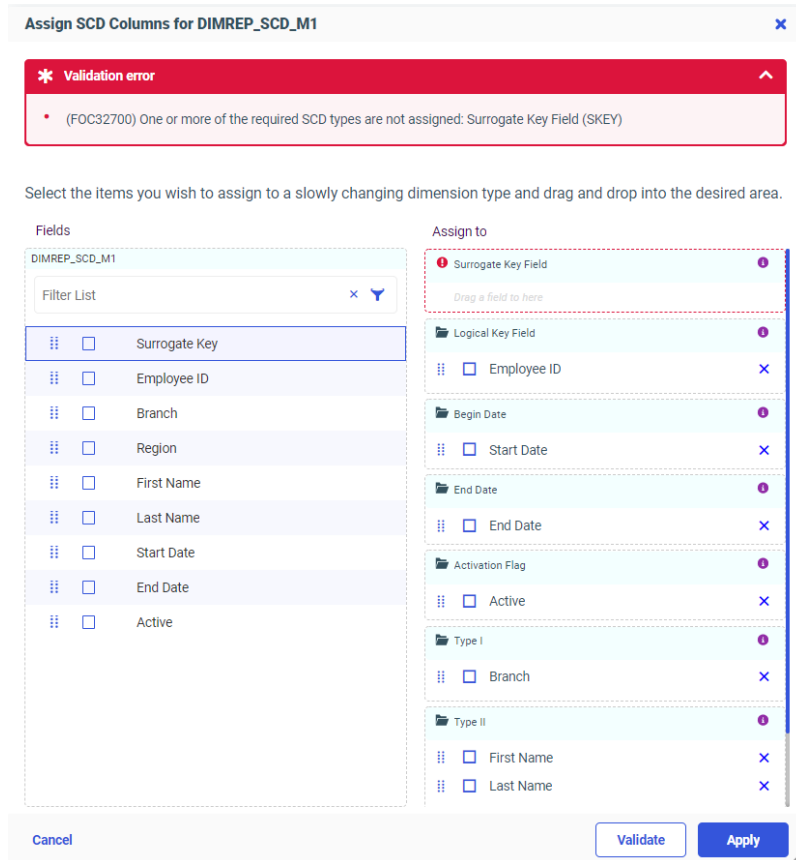
Drag a field to here

Cancel

Validate
Apply

4. When you have finished, click *Validate* to check that the columns you configured are valid.

If there are invalid selections, messages are displayed at the top of the pane in red, as shown in the following image.



5. If necessary, correct any issues.
6. Click *Apply*.

Adapter for PostgreSQL: Support for Change Data Capture (CDC)

In DataMigrator, CDC updates the target using a log file that contains the list of changes to a source, which avoids having to process the entire source file.

Now, DataMigrator can support CDC with PostgreSQL through the consumption of Logical decoding in PostgreSQL.

Preparing the Environment for CDC With PostgreSQL

In order to use CDC with PostgreSQL, you must install the following on the system where the DataMigrator Server is installed.

- Java.
- PostgreSQL.
- The JDBC Driver.
- The output plugin wal2json.

Reference: Configuring Java and the JDBC Driver

To use CDC with the Adapter for PostgreSQL, you must have Java installed. The System Environment Variable `JAVA_HOME` or `JDK_HOME` must be specified, and the `JAVA_HOME` location must be added to the `PATH` Environment Variable.

For example, if your DataMigrator Server is on Linux, add the following lines to your profile:

```
export JAVA_HOME=/usr/lib/jvm/jre-1.8.0_91
export JDK_HOME=/usr/lib/jvm/jdk-1.8.0_91
export PATH=$JAVA_HOME/bin:$PATH
```

The JDBC Driver jar file is needed in order to use CDC with PostgreSQL. You should download the jar file and save it to your path since you will be entering it in the `IBI_CLASSPATH` entry field during the configuration process. You can download the JAR file from the following site:

<https://www.postgresql.org/download/>

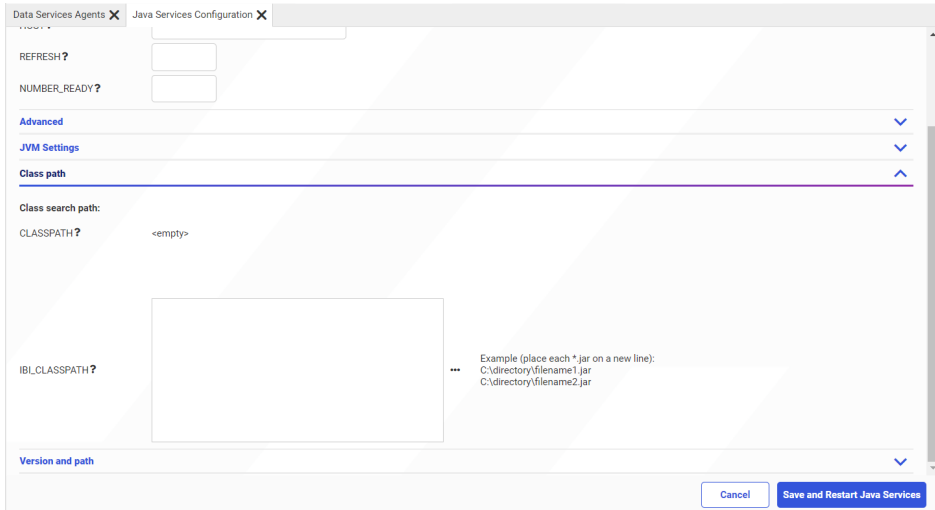
Copy the JDBC Driver jar file to any location on your system.

To specify the location of the jar file in the `IBI_CLASSPATH` variable of a DataMigrator Server, perform the following steps.

1. In the Web Console, navigate to the Workspace page or, in the Data Management Console (DMC), expand Workspace. Then expand the Java Services folder.
2. Right-click *DEFAULT* and click *Properties*.

The Java Services Configuration pane opens.

3. Expand the Class path section, as shown in the following image.



4. In the IBI_CLASSPATH field, enter the full path file name to the JDBC Driver jar file.
5. Click *Save and Restart Java Services*.

Reference: Configuring Use of Write-Ahead Logs

Logical decoding in PostgreSQL uses an output plugin called *wal2json*. This plugin converts the PostgreSQL Write-Ahead Log (WAL) to a readable format.

The plugin is available for download from the site <https://github.com/eulerto/wal2json>.

Follow the instructions on the site to install the plugin.

The following is an example installation command for Red Hat Enterprise Linux (RHEL)/CentOS 7 or 8:

```
sudo yum install wal2json12
```

The following is an example installation command for Debian:

```
sudo dnf install wal2json12
```

The following is an example installation command for Ubuntu:

```
sudo apt-get install postgresql-12-wal2json
```

When using Windows, follow the instructions from the same site and build the wal2json.dll with the following considerations.

- ❑ Use non-debug mode.
- ❑ Use the same operating system where PostgreSQL is installed.
- ❑ During the build, target the same version of PostgreSQL as the version used with the DLL.

Copy the created wal2json.dll to the pkglibdir location. To find this location, open a command prompt and navigate to the PostgreSQL bin directory. Use the following command, where PGDATA is an environmental variable that points to the *data* directory.

```
cd %PGDATA%\..\bin
```

From this location, enter the following command.

```
pg_config --pkglibdir
```

The following is a sample path on Windows.

```
C:\Program Files\PostgreSQL\12
```

In order to use the write-ahead logs, you must edit the postgresql.conf configuration file and set the options for the PostgreSQL *wal_level* parameter, which determines the information that is written to the WAL. This configuration file is in the data area directory:

```
%PGDATA%/postgresql.conf
```

The *wal_level* parameter should be set to the value *logical*, in order to support logical decoding.

If you connect to the database as user *postgres*, you can edit the postgresql.conf file using the following SQL statement.

```
ALTER SYSTEM SET wal_level = logical;
```

After making the changes, you need to recycle the server as user *postgres*. For example:

```
%PGDATA%\..\bin pg_ctl restart
```

To verify the value of the *wal_level* parameter, you can issue the following SQL statement.

```
SHOW wal_level;
```

Using Logical Decoding and Monitoring Slots

Using a replication slot requires PostgreSQL replication privileges.

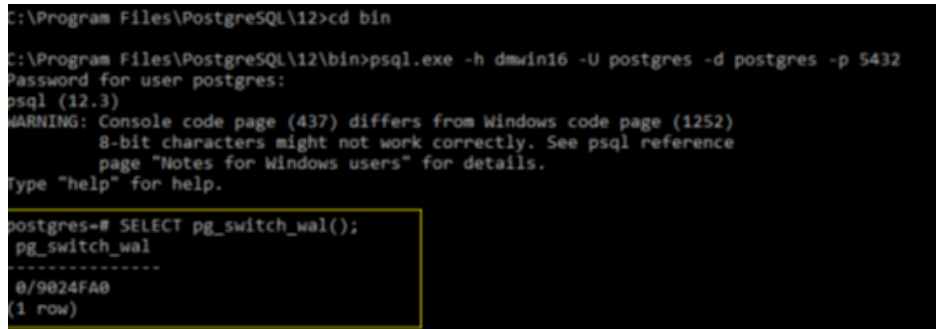
Each replication slot in PostgreSQL represents a stream of changes from a single database.

Procedure: How to Create a Replication Slot

1. Connect to the database with a user that has a SUPERUSER permissions.
2. To archive the current write-ahead log file and to switch to a new log file, enter the following command:

```
SELECT pg_switch_wal();
```

The following image shows issuing this command in Windows.



```
C:\Program Files\PostgreSQL\12>cd bin
C:\Program Files\PostgreSQL\12\bin>psql.exe -h dwin16 -U postgres -d postgres -p 5432
Password for user postgres:
psql (12.3)
WARNING: Console code page (437) differs from Windows code page (1252)
         8-bit characters might not work correctly. See psql reference
         page "Notes for Windows users" for details.
Type "help" for help.

postgres=# SELECT pg_switch_wal();
 pg_switch_wal
-----
 0/9024FA0
(1 row)
```

3. When a slot is no longer needed, drop it.

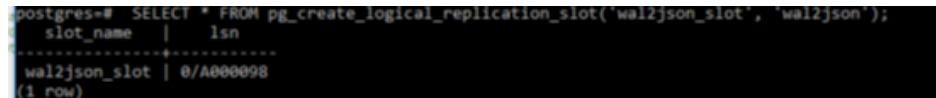
According to the PostgreSQL documentation, the logical replication slots are consumed continuously, and you need to drop a slot as soon as the replication slot is not needed anymore. Use the following command:

```
SELECT pg_drop_replication_slot('wal2json_slot');
```

4. Add new replication slots as needed, using the following command.

```
SELECT * FROM pg_create_logical_replication_slot('wal2json_slot',
'wal2json');
```

The following image shows issuing this command in Windows.



```
postgres=# SELECT * FROM pg_create_logical_replication_slot('wal2json_slot', 'wal2json');
 slot_name | lsn
-----
 wal2json_slot | 0/A000098
(1 row)
```


Reference: Monitoring Slots

The following commands are used for logical decoding and monitoring the replication slots.

Listing Replication Slots

To list all replication slots in the primary server, issue the following command:

```
SELECT * FROM pg_replication_slots;
```

The following image shows issuing this command in Windows.

```
postgres=# SELECT * FROM pg_replication_slots;
 slot_name | plugin | slot_type | datoid | database | temporary | active | active_pid | xmin | catalog_xmin | restart_lsn | confirmed_flush_lsn
-----
 test_decoding_slot | test_decoding | logical | 16383 | main | f | f | | | | | 672 | 0/1111000
 wal2json | wal2json | logical | 13310 | postgres | f | f | | | | | 672 | 0/58A2940
 0/58A2978
 wal2json_slot | wal2json | logical | 13310 | postgres | f | f | | | | | 672 | 0/A000000
 0/A000000
(3 rows)
```

Listing the Current WAL Write Location

To list the current write-ahead log write location, issue the following command:

```
SELECT pg_current_wal_lsn();
```

The following image shows issuing this command in Windows.

```
postgres=# SELECT pg_current_wal_lsn();
 pg_current_wal_lsn
-----
 0/A000148
(1 row)
```

Reference: Creating a User ID for Connection to PostgreSQL

A user ID on a connection to a PostgreSQL database must have the LOGIN and REPLICATION permissions. Please note that it is not advisable to add a SUPERUSER permission to this user, since it is insecure. This role allows you to bypass all permission checks. The following is sample syntax to create the user:

```
CREATE ROLE cdcuser1 WITH
    LOGIN
    NOSUPERUSER
    CREATEDB
    REPLICATION
    PASSWORD 'password';
```

Use the following statement to grant privileges to the user on the schema that will use the replications:

```
GRANT ALL PRIVILEGES ON ALL TABLES IN SCHEMA cdcschema1 TO cdcuser1;
```

Use the following statement to make the user an owner of the schema:

```
ALTER SCHEMA cdcschema1 OWNER TO cdcuser1;
```

Configuring the Adapter for PostgreSQL and Using CDC in a Flow

You must create a connection for the Adapter for PostgreSQL that accesses the environment you prepared.

Reference: **Connection Attributes for the Adapter for PostgreSQL**

The following image shows the configuration parameters for the Adapter for PostgreSQL.

The screenshot shows a configuration window titled "Add PostgreSQL to Configuration" with a "Prerequisites" link. Below the title is a "Connect parameters" section with an upward arrow. The configuration fields are: "Connection Name" with a value of "CON01"; "URL" with a "Show Sample" link below it; "Security" set to "Explicit" with a dropdown arrow; "User" and "Password" fields, all with help icons to their right.

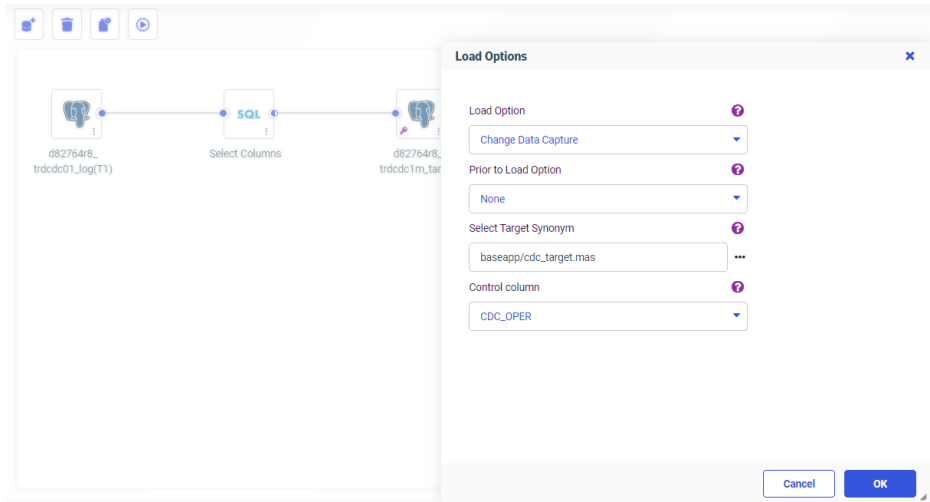
From your PostgreSQL DBA, obtain the values for the node name, port number and database name, then enter them in the URL entry field in the following format.

```
jdbc:postgresql://nodename:port/dbname
```

Leave Security as *Explicit*, and enter the User ID and Password associated with the role that you configured.

Reference: Using CDC With PostgreSQL

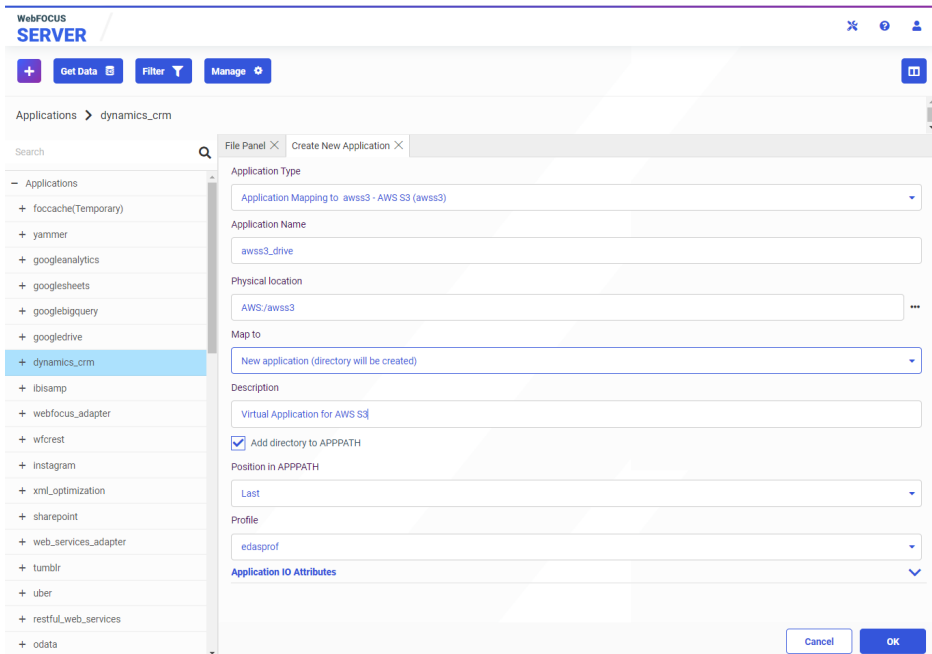
Once you have configured PostgreSQL for logging, you can create a data flow with Load Option Change Data Capture using a PostgreSQL log file as the source and an existing PostgreSQL target, as shown in the following image.



Support for Direct Load Flows to Load AWS S3 Data to PostgreSQL

If you have configured both the Adapter for AWS S3 and the Adapter for PostgreSQL, you can create a direct load flow from a delimited file in the AWS S3 repository to a PostgreSQL target.

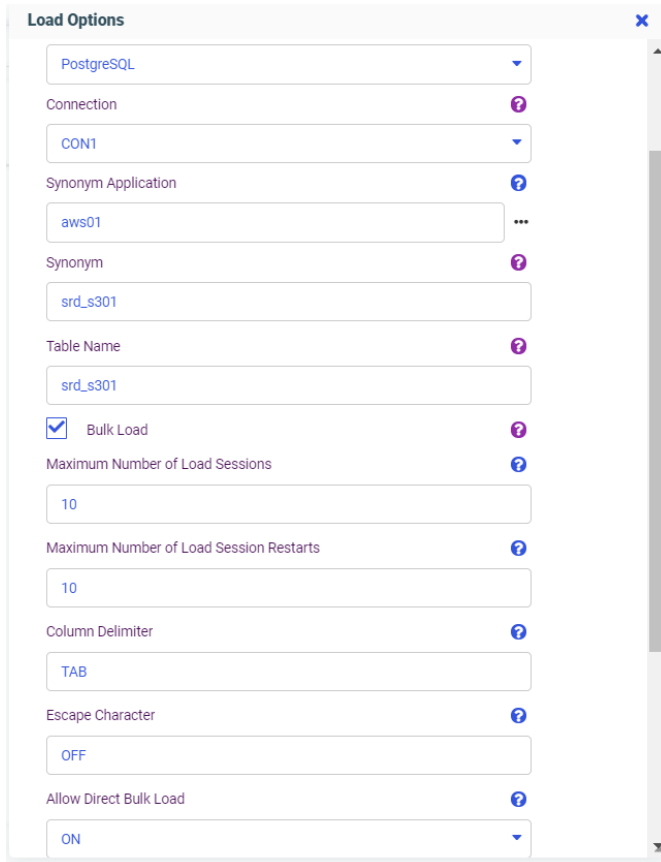
When you configure the Adapter for AWS S3, you map the AWS S3 repository to an application on the server, as shown in the following image.



Then, you can drag a file from that application directory to the Data Flow canvas as a source for the flow. The following image shows a Data Flow in the Server Web Console with an AWS S3 source and a PostgreSQL target.



In the Load Options dialog box, select *PostgreSQL* from the *Adapter* drop-down list, select the *Bulk Load* check box, and select *On* from the *Allow Direct Bulk Load* drop-down list, as shown in the following image.



The screenshot shows the 'Load Options' dialog box with the following settings:

- Adapter: PostgreSQL
- Connection: CON1
- Synonym Application: aws01
- Synonym: srd_s301
- Table Name: srd_s301
- Bulk Load:
- Maximum Number of Load Sessions: 10
- Maximum Number of Load Session Restarts: 10
- Column Delimiter: TAB
- Escape Character: OFF
- Allow Direct Bulk Load: ON

When you run the flow, the AWS S3 data will be loaded directly to the PostgreSQL table.

Creating a Synonym in a Data Flow Using Variables for the Existing Excel File Name

Using the Data Management Console (DMC), you can now create a synonym in a flow for an existing Excel target using variables for the data file name, as shown in the following image.

Properties	
Attribute	Value
<div style="border: 1px solid #ccc; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px;"> [-] General </div> <div style="padding: 2px;"> Display Name: <code>ibisamp/target01</code> </div> <div style="padding: 2px;"> Notes: ⋮ </div> </div>	
<div style="border: 1px solid #ccc; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px;"> [-] Target Options </div> <div style="padding: 2px;"> Type: <code>New</code> </div> <div style="padding: 2px;"> Adapter: <code>Formatted File</code> </div> <div style="padding: 2px;"> Format: <code>EXCEL</code> </div> <div style="padding: 2px;"> Synonym: <code>ibisamp/target01</code> </div> <div style="padding: 2px;"> Data File: <code>ibisamp/&TARGET_FILE.x...</code> </div> </div>	
<div style="border: 1px solid #ccc; padding: 2px;"> <div style="background-color: #f0f0f0; padding: 2px;"> [-] Target Load Options </div> <div style="padding: 2px;"> *Load Type: <code>Load via Hold</code> </div> </div>	

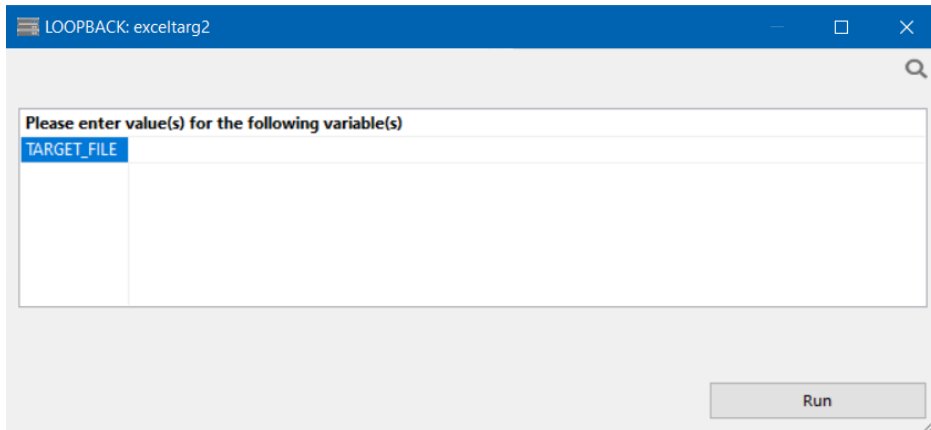
The Formatted File adapter is used with the format EXCEL.

While the Excel file must exist prior to running the flow, the synonym does not have to exist prior to running the flow.

The following flow loads the existing Excel target file with the data from the dminv table created as part of the DataMigrator - General tutorial.



Running the flow opens the following dialog box for entering the name of the existing Excel file.



Entering the name `excel01` and running the flow loads the table and generates the synonym.

The generated Master File, `target01.mas`, is:

```
FILENAME=TARGET01, SUFFIX=DIREXCEL,
DATASET=ibisamp/excel01.xlsx, $
SEGMENT=EXCEL01, SEGTYPE=S0, $
  FIELDNAME=PRODUCT_NUMBER, ALIAS='Product,Number', USAGE=A5V, ACTUAL=A5V,
  MISSING=ON,
  TITLE='Product,Number', $
  FIELDNAME=PRODUCT_NAME, ALIAS='Product,Name', USAGE=A37V, ACTUAL=A37V,
  MISSING=ON,
  TITLE='Product,Name', $
  FIELDNAME=QUANTITY_IN_STOCK, ALIAS='Quantity,in Stock',
  USAGE=I8, ACTUAL=A11V,
  MISSING=ON,
  TITLE='Quantity,in Stock', $
  FIELDNAME=SALE_PRICE, ALIAS='Sale,Price', USAGE=D9.2, ACTUAL=A64V,
  MISSING=ON,
  TITLE='Sale,Price', $
  FIELDNAME=OUR_COST, ALIAS='Our,Cost', USAGE=D9.2, ACTUAL=A64V,
  MISSING=ON,
  TITLE='Our,Cost', $
  FIELDNAME=PRODUCT_CATEGORY, ALIAS='Product,Category',
  USAGE=A27V, ACTUAL=A27V,
  MISSING=ON,
  TITLE='Product,Category', $
  FIELDNAME=PRODUCT_TYPE, ALIAS='Product,Type', USAGE=A8V, ACTUAL=A8V,
  MISSING=ON,
  TITLE='Product,Type', $
```

The generated Access File, target01.acx is:

```
SEGNAME=TARGET01,
WORKSHEET=Sheet1,
HROWS=1,
NUMDATA=RAW, $
```

The following image shows the Excel target, excel01.xlsx, after running the flow.

	A	B	C	D	E	F	G
1	Product Number	Product Name	Quantity in Stock	Sale Price	Our Cost	Product Category	Product Type
2	1004	2 Hd VCR LCD Menu	43,068	179.00	129.00	VCRs	Analog
3	1006	Combo Player - 4 Hd VCR + DVD	13,527	399.00	289.00	DVD	Digital
4	1008	DVD Upgrade Unit for Cent. VCR	199	199.00	139.00	DVD	Digital
5	1010	750SL Digital Camcorder 300 X	10,758	999.00	750.00	Camcorders	Digital
6	1012	650DL Digital Camcorder 150 X	2,972	899.00	710.00	Camcorders	Digital
7	1014	340SX Digital Camera 65K P	990	249.00	199.00	Cameras	Digital
8	1016	330DX Digital Camera 1024K P	12,707	279.00	199.00	Cameras	Digital
9	1018	250 8MM Camcorder 40 X	60,073	399.00	320.00	Camcorders	Analog
10	1020	150 8MM Camcorder 20 X	5,961	319.00	240.00	Camcorders	Analog
11	1022	120 VHS-C Camcorder 40 X	2,300	399.00	259.00	Camcorders	Analog
12	1024	110 VHS-C Camcorder 20 X	4,000	349.00	249.00	Camcorders	Analog
13	1026	AR3 35MM Camera 10 X	12,444	129.00	95.00	Cameras	Analog
14	1028	AR2 35MM Camera 8 X	11,499	109.00	79.00	Cameras	Analog
15	1030	QX Portable CD Player	22,000	169.00	99.00	CD Players	Digital
16	1032	R5 Micro Digital Tape Recorder	1,990	89.00	69.00	Digital Tape Recorders	Digital
17	1034	ZT Digital PDA - Commercial	21,000	499.00	349.00	PDA Devices	Digital
18	1036	ZC Digital PDA - Standard	33,000	299.00	249.00	PDA Devices	Digital
19	2002	Memory Stick	1,000	80.00	40.00	PDA Devices	Analog
20	2003	Memory Stick Pro	1,000	200.00	100.00	PDA Devices	Analog
21							
22							
23							

New Option to Save Data Files When Bulk Load is Used in a Flow

When Bulk Load is selected for a target in a Data Flow, the option *Save Data Files* appears, as shown in the following image.

The screenshot shows the 'Load Options' configuration window for a target named 'citibike_tripdata01'. The 'Bulk Load' checkbox is checked. The 'Save Data Files' dropdown menu is open, showing 'NO - Default' as the selected option and 'YES' as an alternative option. Other options include 'Maximum Number of Load Sessions' (10), 'Maximum Number of Load Session Restarts' (10), 'Column Delimiter' (TAB), and 'Allow Direct Bulk Load' (ON).


The default value, *No*, deletes intermediate *.ftm files from the directory for temporary files (for example, edatemp) after the corresponding portion of data is successfully committed to the target table during the Bulk Load Flow execution.

If you select *Yes*, intermediate *.ftm files are deleted from the temporary directory (for example, edatemp) after the job completes execution.

Slowly Changing Dimensions (SCD): Re-run for Missing Date

Using the Web Console and Slowly Changing Dimensions, you can re-run an SCD flow for a missed day, after subsequent days have already been processed.

When the Data Flow has an SCD target, the Load Options dialog box has the option *Overwrite Begin Date for active and End Date for inactive records*, as shown in the following image.



Check this box to select Begin Date and End Date options. Both Begin Date and End Date must be described with the same data type. The choices for overwriting the dates for Begin Date and End Date follow.

Begin Date

A Calendar Control opens, so you can select a beginning date, as shown in the following image.

The screenshot shows the 'Load Options' dialog box with the following fields and options:

- Prior to Load Option:** None
- Select Target Synonym (Slowly Changing Dimensions):** repro_nfs/demo/customer.mas
- Updates for Type I:** Change all rows
- Overwrite Begin Date for active and End Date for inactive records**
- Begin Date for active records:** 2020/06/10

A calendar control is open, showing the month of June 2020. The date 10 is selected. The calendar has a header with 'Su', 'Mo', 'Tu', 'We', 'Th', 'Fr', 'Sa' and a grid of dates from 1 to 30. Navigation arrows and dropdowns for month and year are visible at the top of the calendar.

End Date

The choices for End Date depend on the data type:

- DATE format (for example, YYMD).** The options are:

The same as Begin Date for active record

Day prior to Begin Date for active record

These options are shown in the following image.

Load Options

Prior to Load Option: None

Select Target Synonym (Slowly Changing Dimensions): repro_nfs/demo/customer.mas

Updates for Type I: Change all rows

Overwrite Begin Date for active and End Date for inactive records

Begin Date for active records: 2020/06/25

End Date for inactive records: The same as Begin Date for active records

Dropdown options for End Date for inactive records:
The same as Begin Date for active records
The same as Begin Date for active records
Day prior to Begin Date for active records

Date-Time (Timestamp) format (for example, HYYMDS). Only one option is available:

[The same as Begin Date for active record](#)

Slowly Changing Dimensions: Enhanced Processing for Type I Columns

In prior releases, if a source row had Type I field values that were the same as the existing row, the Changed Flag was restored and the record was updated with the same value.

Now, when DataMigrator Slowly Changing Dimension processing is used to load a dimension table with Type I columns, only rows in which the source and target rows have different values for the Type I columns are updated. This can result in reduced processing time, especially for large dimension tables.

Change Data Capture (CDC) on a Data Flow

In the Web Console, when a table is log-enabled, you can use the load option *Change Data Capture* along with the log file to update only the rows that have changed since the last load.

Procedure: How to Use CDC on a Data Flow

In order to use CDC on a Data Flow, you need a synonym for the log file, a synonym for the existing target table, a Data Flow, and, optionally, a synonym for the source table (which is not used in this type of load). When changes are made to the source table, the log file is updated. If you submit a flow that has Change Data Capture as its load option, the target table will be updated using the information in the log.

1. Create a synonym for the log table.
 - a. Right-click the connection for the adapter with the log-enabled tables, and select *Show DBMS Objects*.

The Create Synonym page opens.

- b. Select *Table Log Records* from the Object Type drop-down list, as shown in the following image.

Create Synonym for MS SQL Server ODBC/AzureDB (DMTFS02)

Object Type? Table Log Records Database? Default Database Owner/Schema? Object Name?

Customize data Tables

Create? Tables, Views and Other Objects

Application? External SQL Scripts Stored Procedures Prefix? Suffix?

SYNONYM CA Table Log Records (DEFAULT) Row Limit 50 Search


Fact	Dimension	Default Synonym Name	Table Name	Owner/Schema	Type
<input type="checkbox"/>	<input type="checkbox"/>	m1_orders	M1_ORDERS	dbo	Log Records
<input type="checkbox"/>	<input type="checkbox"/>	dimrep_s	dimrep_s	dbo	Log Records
<input type="checkbox"/>	<input type="checkbox"/>	lobvarbin	LOBVARBIN	dbo	Log Records
<input type="checkbox"/>	<input type="checkbox"/>	car_new_mss	car_new_mss	dbo	Log Records
<input type="checkbox"/>	<input type="checkbox"/>	carsales	CARSALES	dbo	Log Records
<input type="checkbox"/>	<input type="checkbox"/>	carsales_enriched	carsales_enriched	dbo	Log Records
<input type="checkbox"/>	<input type="checkbox"/>	lobtest	LOBTEST	dbo	Log Records

Cancel Add

- c. Click the *Base Synonym* button.
 - d. Either accept the default application in the *Application* text box, enter an application name, or click the ellipsis (...) next to the text box to select the application name.
 - e. Select the check boxes for one or more log records.

If you want to distinguish the synonym name for the log record from the synonym name for the base table, you can add a prefix or suffix using the *Prefix* or *Suffix* entry fields, or you can change the synonym name by typing over the default synonym name. (If the *Default Synonym Name* column is not showing, click the *Choose Columns* button and select that column.)

The following image shows one log table record named *dimrep_s* selected, with the synonym name changed to *dmrep_s_log* (the characters *_log* are entered in the *Suffix* field).



- f. Click *Add*.
2. You should already have a synonym for the existing target table and, optionally, the source table. If you do not have a synonym for the target, create it now.
3. Create a Data Flow using the log file synonym as the source and the existing target as the Target, as shown in the following image.



4. Right-click the target synonym and select *Load Options*.

The Load Options dialog box opens.

5. Select *Change Data Capture* from the Load Option drop-down list, as shown in the following image.


The screenshot shows a dialog box titled "Load Options" with a close button (X) in the top right corner. The dialog contains the following fields:

- Load Option:** A dropdown menu with "Change Data Capture" selected.
- Prior to Load Option:** A dropdown menu with "None" selected.
- Select Target Synonym:** A text input field containing "cdc_load/dimrep.mas" and a search icon (three dots).
- Control column:** A dropdown menu with "CDC_OPER" selected.

At the bottom right of the dialog are two buttons: "Cancel" and "OK".


6. Click *OK*.

The Merge Editor opens, as shown in the following image.



Missing matches display in red.

7. Drag source columns to the Insert Expression field for the target where needed to eliminate missing matches, and click **OK**.
8. Save the flow.
9. If the flow is scheduled to run at set intervals, once the source table and log have been updated, the target will be updated the next time the flow runs. You can also right-click the flow and select *Submit*, as shown in the following image.



You can see the number of changes that were applied by right-clicking the flow, clicking *Logs*, and selecting *Last Log*.

Slowly Changing Dimensions Activation Flag Enhancement

Synonyms for data sources that have Type I and Type II Slowly Changing Dimensions (SCD) have an activation flag that identifies whether a record is part of the history or is currently active.

In prior releases, the flag had to be an integer value, where the value 1 identified the active record and the value zero (0) identified inactive (historical) records.

Starting in this release, the activation flag can be a user-defined alphanumeric value up to 20 characters. For example:

```
FIELDNAME=ACTIVE, ALIAS=ACTIVE, USAGE=A20V, ACTUAL=A20V, $
```

When this field described as a character field, you can still use the values 1 and 0 as in prior releases, or you can assign other values. To assign values, you can use the Set Variables object on a Process Flow tab and set values for the global variables `CM_SCDACT` and `CM_SCDINACT`.

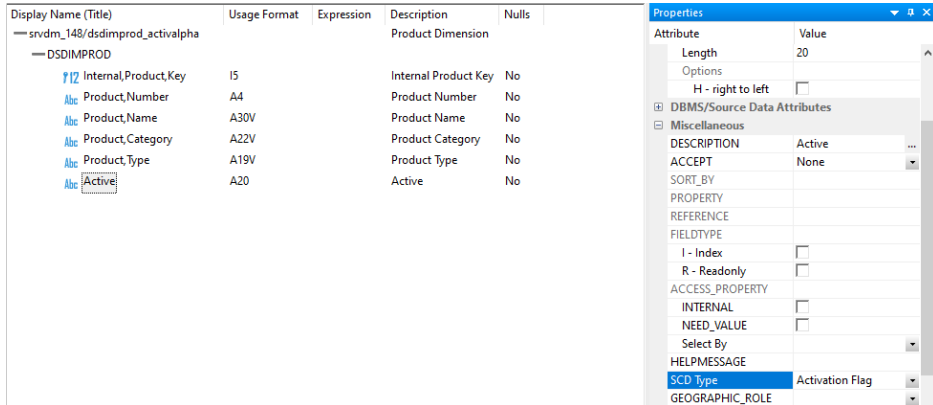
Example: Using an Alphanumeric Activation Flag for a Slowly Changing Dimension

The Master File **`dsdimprod_activealpha.mas`** has a field named `PROD_NAME` defined as SCD Type II, as shown in the following image.

Display Name (Title)	Usage Format	Expression	Description	Nulls
— srvdm_148/dsdimprod_activalpha			Product Dimension	
— DSDIMPROD				
↑ Internal_Product_Key	I5		Internal Product Key	No
Product_Number	A4		Product Number	No
Product_Name	A30V		Product Name	No
Product_Category	A22V		Product Category	No
Product_Type	A19V		Product Type	No
Active	A20		Active	No

Attribute	Value
FIELDNAME	PRODNAME
TITLE	Product,Name
Usage Format	A30V
Type	Character
Length	30
DBMS/Source Data Attributes	
Miscellaneous	
DESCRIPTION	Product Name
ACCEPT	None
SORT_BY	
PROPERTY	
REFERENCE	
FIELDTYPE	
I - Index	<input type="checkbox"/>
R - Readonly	<input type="checkbox"/>
ACCESS_PROPERTY	
INTERNAL	<input type="checkbox"/>
NEED_VALUE	<input type="checkbox"/>
Select By	
HELPMESSAGE	
SCD_Type	Type II
GEOGRAPHIC_ROLE	

The field named ACTIVE is the activation flag, as shown in the following image.



The field ACTIVE has format A20.

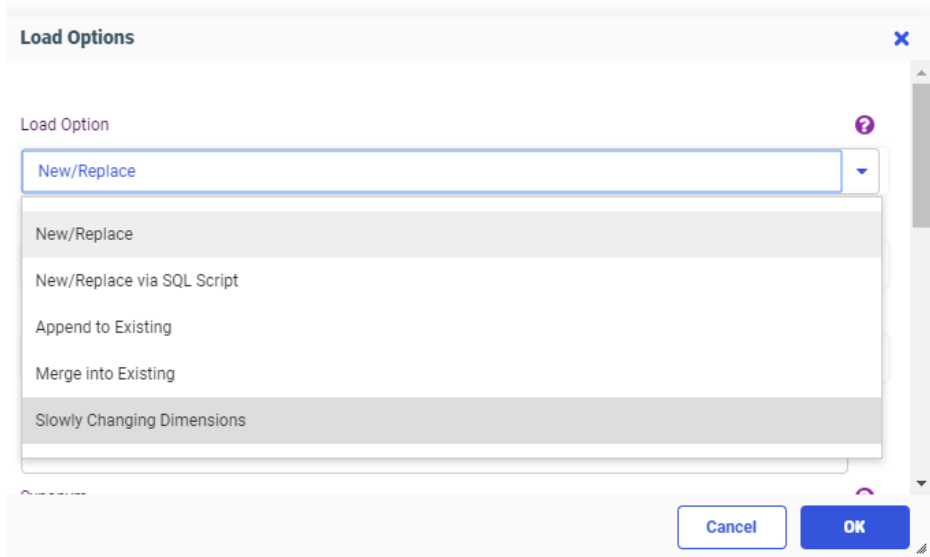
The following image shows how the records are stored when a new source comes in for product number 1006 with a non-matching product name.

	Internal Product Key	Product Number	Product Name	Product Category	Product Type	Active
1	1	1004	2 Hd VCR LCD Menu	VCRs	Analog	active
2	2	1006	producta	PDA Devices	Analog	NOT ACTIVE
3	18	1006	productn	PDA Devices	Analog	active
4	3	1008	DVD Upgrade Unit for Cent. VCR	DVD	Digital	active
5	4	1010	750SL Digital Camcorder 300 X	Camcorders	Digital	active
6	5	1012	650DL Digital Camcorder 150 X	Camcorders	Digital	active

The activation flag value for the active record is *active*. The activation flag value for the inactive record is *NOT ACTIVE*.


Slowly Changing Dimensions on a Data Flow

When loading to a target in a data flow on the Web Console that supports Slowly Changing Dimensions (SCDs), you can select the load type *Slowly Changing Dimensions* from the Load Options drop-down list, as shown in the following image.



If your target synonym has columns with:

- Type II SCDs, those columns will be loaded using the rules established for Type II slowly changing dimensions.
- Type I SCDs, you can choose whether to update all rows (the default) or only active rows, as shown in the following image.



Note: Although this option appears whether the synonym has Type I SCDs or not, it is ignored for Type II SCDs.

Reporting Language Enhancements

This section describes reporting language enhancements.

International System (SI) Numeric Format Abbreviation Options

The International System standard provides numeric abbreviations for very large and very small numbers.

WebFOCUS supports the following SI-compliant numeric abbreviations. The SI-compliant format uses a two-character display code that consists of a lowercase *n* followed by the SI abbreviation.

Prefix	WebFOCUS Format Code	Size	Example	English Name (American/British)
yotta	nY	10**24	1000000000000000000000000	septillion/quadrillion
zetta	nZ	10**21	100000000000000000000000	sextillion/trilliard
exa	nE	10**18	100000000000000000000000	quintillion/trillion
peta	nP	10**15	100000000000000000000000	quadrillion/billiard
tera	nT	10**12	100000000000000000000000	trillion/billion
giga	nG	10**9	1000000000	billion/milliard
mega	nM	10**6	1000000	million
kilo	nK	10**3	1000	thousand
milli	nm	10**(-3)	0.001	thousandth
micro	nu	10**(-6)	0.000001	millionth
nano	nn	10**(-9)	0.000000001	billionth/milliardth
pico	np	10**(-12)	0.000000000001	trillionth/billionth
femto	nf	10**(-15)	0.000000000000001	quadrillionth/billiardth
atto	na	10**(-18)	0.000000000000000001	quintillionth/trillionth
zepto	nz	10**(-21)	0.0000000000000000001	sextillionth/trilliardth
yocto	ny	10**(-24)	0.0000000000000000000001	septillionth/quadrillionth

The following request uses the mega and giga format options. The decimal precision is controlled by the format which, in this case, is a reformat specified in the SUM command.

```
DEFINE FILE GGSales
NEWDOll/D12.2 = DOLLARS * 100;
END
TABLE FILE GGSales
SUM DOLLARS NEWDOll/D12.5nM AS Millions NEWDOll/D12.3nG AS Billions
BY CATEGORY
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

The output is shown in the following image.

<u>Category</u>	<u>Dollar Sales</u>	<u>Millions</u>	<u>Billions</u>
Coffee	17231455	1,723.14550M	1.723G
Food	17229333	1,722.93330M	1.723G
Gifts	11695502	1,169.55020M	1.170G

New Functions for Date-time Conversion Between Local and UTC Time

Coordinated Universal Time (UTC) is the time standard commonly used around the world. To convert UTC time to a local time, a certain number of hours must be added to or subtracted from the UTC time, depending on the number of time zones between the locality and Greenwich, England (GMT).

The following functions convert date-time values between UTC time and local time.

- ❑ **DT_TOUTC.** Converts local time to UTC time.
- ❑ **DT_TOLOCAL.** Converts UTC time to local time.

Converting timestamp values from different localities to a common standard time enables you to sort events into the actual event sequence.

These functions require IANA (Internet Assigned Numbers Authority) time zone database names (expressed as 'Area/Location') as parameters. You can find a table of IANA TZ database names on Wikipedia at https://en.wikipedia.org/wiki/List_of_tz_database_time_zones, as shown in the following image.

Legend [\[edit \]](#)

UTC offsets (columns 6 and 7) are positive east of UTC and negative west of UTC. The *UTC DST offset* is different from the *UTC offset* for zones where *daylight saving time* is observed (see *Individual time zone pages* for details). The UTC offsets are for the current or upcoming rules, and may have been different in the past.

The "Status" field means:

- Canonical - The primary, preferred zone name.
- Alias - An alternative name, which may fit better within a particular country.
- Deprecated - An older style name, left in the tz database for backwards compatibility, which should generally not be used.

List [\[edit \]](#)

Country code	Latitude, longitude ±DDMM(SS) ±DDMM(SS)	TZ database name	Portion of country covered	Status	UTC offset ±hh:mm	UTC DST offset ±h:mm	Notes
CI	+0519-00402	Africa/Abidjan		Canonical	+00:00	+00:00	
GH	+0533-00013	Africa/Accra		Canonical	+00:00	+00:00	
ET	+0902+03842	Africa/Addis_Ababa		Alias	+03:00	+03:00	Link to Africa/Nairobi
DZ	+3647+00303	Africa/Algiers		Canonical	+01:00	+01:00	
ER	+1520+03853	Africa/Asmara		Alias	+03:00	+03:00	Link to Africa/Nairobi
ML	+1239-00800	Africa/Bamako		Alias	+00:00	+00:00	Link to Africa/Abidjan
CF	+0422+01835	Africa/Bangui		Alias	+01:00	+01:00	Link to Africa/Lagos
GM	+1328-01639	Africa/Banjul		Alias	+00:00	+00:00	Link to Africa/Abidjan
GW	+1151-01535	Africa/Bissau		Canonical	+00:00	+00:00	
MW	-1547+03500	Africa/Biantyre		Alias	+02:00	+02:00	Link to Africa/Maputo
CG	-0416+01517	Africa/Brazzaville		Alias	+01:00	+01:00	Link to Africa/Lagos
BI	-0323+02922	Africa/Bujumbura		Alias	+02:00	+02:00	Link to Africa/Maputo
EG	+3003+03115	Africa/Cairo		Canonical	+02:00	+02:00	

If you do not know what Area and Location corresponds to your time zone, but you do know your offset from GMT, or your legacy time zone name (such as EST), scroll down in the table. There are TZ database names that correspond to these time zone identifiers, as shown in the following image.

	EST	Deprecated	-05:00	-05:00	Choose a zone that currently observes EST without daylight saving time, such as America/Cancun .
	EST5EDT	Deprecated	-05:00	-04:00	Choose a zone that observes EST with United States daylight saving time rules, such as America/New_York .
	Etc/GMT	Canonical	+00:00	+00:00	
	Etc/GMT+0	Alias	+00:00	+00:00	Link to Etc/GMT
	Etc/GMT+1	Canonical	-01:00	-01:00	Sign is intentionally inverted. See the Etc area description .
	Etc/GMT+10	Canonical	-10:00	-10:00	Sign is intentionally inverted. See the Etc area description .
	Etc/GMT+11	Canonical	-11:00	-11:00	Sign is intentionally inverted. See the Etc area description .
	Etc/GMT+12	Canonical	-12:00	-12:00	Sign is intentionally inverted. See the Etc area description .
	Etc/GMT+2	Canonical	-02:00	-02:00	Sign is intentionally inverted. See the Etc area description .
	Etc/GMT+3	Canonical	-03:00	-03:00	Sign is intentionally inverted. See the Etc area description .
	Etc/GMT+4	Canonical	-04:00	-04:00	Sign is intentionally inverted. See the Etc area description .
	Etc/GMT+5	Canonical	-05:00	-05:00	Sign is intentionally inverted. See the Etc area description .

Note: If you use a standard IANA time zone database name in the form "Area/Location" (for example, "America/New_York"), automatic adjustments are made for Daylight Savings Time. If you use a name that corresponds to an offset from GMT or to a legacy time zone name, it is your responsibility to account for Daylight Savings Time.

DT_TOUTC: Converting Local Time to UTC Time

DT_TOUTC takes a local date-time value and an IANA time zone name and converts the local time to UTC time.

Syntax: How to Convert Local Time to UTC Time

```
DT_TOUTC(datetime, timezone)
```

where:

datetime

Date-time

Is a date-time expression representing local time, containing date and time components.

timezone

Alphanumeric

Is a character expression containing the IANA time zone name of the local time, in the form 'Area/Location' (for example, 'America/New_York').

Example: Converting Local Time to UTC Time

The following request converts the current local date-time value for time zone America/New_York to UTC time.

```
TABLE FILE GGSales
SUM DOLLARS NOPRINT
COMPUTE LOCAL1/HYYMDS = DT_CURRENT_DATETIME(SECOND);
COMPUTE UTC1/HYYMDS = DT_TOUTC(LOCAL1, 'America/New_York');
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

The output is shown in the following image.

<u>LOCAL1</u>	<u>UTC1</u>
2020/09/04 14:49:41	2020/09/04 18:49:41

Example: Sorting by UTC Time

The following request retrieves the current date and time into the field LOCALT1 and sets the field TIMEZONE to IANA time zone database names. It then uses DT_TOUTC to convert the same local time, with different time zones, to the UTC time that corresponds to the given time zone, and sorts the output by the generated UTC time.

```
DEFINE FILE GGSales
LOCALT1/HYMDs=DT_CURRENT_DATETIME(SECOND);
TIMEZONE/A30=IF LAST TIMEZONE EQ ' ' THEN 'AMERICA/NEW_YORK'
  ELSE IF LAST TIMEZONE EQ 'AMERICA/NEW_YORK' THEN 'AMERICA/CHICAGO'
  ELSE IF LAST TIMEZONE EQ 'AMERICA/CHICAGO' THEN 'AMERICA/DENVER'
  ELSE IF LAST TIMEZONE EQ 'AMERICA/DENVER' THEN 'ASIA/TOKYO'
  ELSE IF LAST TIMEZONE EQ 'ASIA/TOKYO' THEN 'EUROPE/LONDON'
  ELSE IF LAST TIMEZONE EQ 'EUROPE/LONDON' THEN 'AMERICA/NEW_YORK';
UTCTIME/HYMDs=DT_TOUTC(LOCALT1, TIMEZONE) ;
END
TABLE FILE GGSales
PRINT TIMEZONE LOCALT1 DOLLARS NOPRINT
BY UTCTIME
WHERE PRODUCT EQ 'Thermos'
IF RECORDLIMIT EQ 20
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```


The output is shown in the following image.

<u>UTCTIME</u>	<u>TIMEZONE</u>	<u>LOCALTIME</u>
2020/10/02 06:45:59	ASIA/TOKYO	2020/10/02 15:45:59
	ASIA/TOKYO	2020/10/02 15:45:59
	ASIA/TOKYO	2020/10/02 15:45:59
	ASIA/TOKYO	2020/10/02 15:45:59
2020/10/02 14:45:59	EUROPE/LONDON	2020/10/02 15:45:59
	EUROPE/LONDON	2020/10/02 15:45:59
	EUROPE/LONDON	2020/10/02 15:45:59
	EUROPE/LONDON	2020/10/02 15:45:59
2020/10/02 19:45:59	AMERICA/NEW_YORK	2020/10/02 15:45:59
	AMERICA/NEW_YORK	2020/10/02 15:45:59
	AMERICA/NEW_YORK	2020/10/02 15:45:59
	AMERICA/NEW_YORK	2020/10/02 15:45:59
2020/10/02 20:45:59	AMERICA/CHICAGO	2020/10/02 15:45:59
	AMERICA/CHICAGO	2020/10/02 15:45:59
	AMERICA/CHICAGO	2020/10/02 15:45:59
	AMERICA/CHICAGO	2020/10/02 15:45:59
2020/10/02 21:45:59	AMERICA/DENVER	2020/10/02 15:45:59
	AMERICA/DENVER	2020/10/02 15:45:59
	AMERICA/DENVER	2020/10/02 15:45:59
	AMERICA/DENVER	2020/10/02 15:45:59

DT_TOLOCAL: Converting UTC Time to Local Time

DT_TOLOCAL takes a UTC date-time value and an IANA time zone name and converts the UTC time to local time.

Syntax: How to Convert UTC Time to Local Time

```
DT_TOLOCAL(datetime, timezone)
```

where:

datetime

Date-time

Is a date-time expression representing UTC time, containing date and time components.

timezone

Alphanumeric

Is a character expression containing the IANA time zone name of the local time, in the form 'Area/Location' (for example, 'America/New_York').

Example: **Converting UTC Time to Local Time**

The following request converts the current date-time value from UTC time to local time for time zone America/New_York.

```
TABLE FILE GGSales
SUM DOLLARS NOPRINT
COMPUTE UTC1/HYMD5 = DT_CURRENT_DATETIME(SECOND);
COMPUTE LOCAL1/HYMD5 = DT_TOLOCAL(UTC1, 'America/New_York');
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

The output is shown in the following image.

<u>UTC1</u>	<u>LOCAL1</u>
2020/09/04 15:00:26	2020/09/04 11:00:26

New Regular Expression Pattern Matching Functions

The following FOCUS and SQL functions search for strings that match a pattern expressed as a regular expression.

- REGEXP_COUNT.** Counts the number of matches to a regular expression pattern in a string.
- REGEXP_INSTR.** Returns the first position of a regular expression pattern in a string.
- REGEXP_REPLACE.** Replaces all matches to a regular expression pattern in a string with a replacement string.
- REGEXP_SUBSTR.** Returns the first match to a regular expression pattern in a string.

You can search online for information about the symbols used to create a regular expression pattern. For example, Wikipedia has a good introduction at:

https://en.wikipedia.org/wiki/Regular_expression

REGEXP_COUNT: Counting the Number of Matches to a Pattern in a String

REGEXP_COUNT returns the integer count of matches to a specified regular expression pattern within a source string.

Syntax: How to Count the Number of Matches to a Pattern in a String

```
REGEXP_COUNT(string, pattern)
```

where:

string

Alphanumeric

Is the input string to be searched.

pattern

Alphanumeric

Is the regular expression pattern to match.

Example: Counting the Number of Matches to a Pattern in a String

The following examples use the following Regular Expression symbols.

❑ \$, which searches for a specified expression that occurs at the end of a string.

❑ ^, which searches for a specified expression that occurs at the beginning of a string.

REGEXP_COUNT counts the number of occurrences of the characters 'umpty' that occur at the end of the string 'Humpty Dumpty'.

```
REGEXP_COUNT('Humpty Dumpty', 'umpty$')
```

The result is 1.

REGEXP_COUNT counts the number of occurrences of the characters 'umpty' that occur at the beginning of the string 'Humpty Dumpty'.

```
REGEXP_COUNT('Humpty Dumpty', '^umpty')
```

The result is 0.

REGEXP_INSTR: Returning the First Position of a Pattern in a String

REGEXP_INSTR returns the integer position of the first match to a specified regular expression pattern within a source string. The first character position in a string is indicated by the value 1. If there is no match within the source string, the value 0 is returned.

Syntax: **How to Return the Position of a Pattern in a String**

```
REGEXP_INSTR(string, pattern)
```

where:

string

Alphanumeric

Is the input string to be searched.

pattern

Alphanumeric

Is the regular expression pattern to match.

Example: **Finding the Position of a Pattern in a String**

The following examples use the following Regular Expression symbols.

❑ \$, which searches for a specified expression that occurs at the end of a string.

❑ ^, which searches for a specified expression that occurs at the beginning of a string.

REGEXP_INSTR finds the position of the characters 'umpty' that occur at the end of the string 'Humpty Dumpty'.

```
REGEXP_INSTR('Humpty Dumpty', 'umpty$')
```

The result is 9.

REGEXP_INSTR finds the position of the characters 'umpty' that occur at the beginning of the string 'Humpty Dumpty'.

```
REGEXP_INSTR('Humpty Dumpty', '^umpty')
```

The result is 0.

REGEXP_REPLACE: Replacing All Matches to a Pattern in a String

REGEXP_REPLACE returns a string generated by replacing all matches to a regular expression pattern in the source string with the given replacement string. The replacement string can be a null string.

Syntax: **How to Replace Matches to a Pattern in a String**

```
REGEXP_REPLACE(string, pattern, replacement)
```

where:

string

Alphanumeric

Is the input string to be searched.

pattern

Alphanumeric

Is the regular expression pattern to match.

replacement

Alphanumeric

Is the replacement string.

Example: Replacing Matches to a Pattern in a String

The following example uses the following Regular Expression symbol.

□ ^, which searches for a specified expression that occurs at the beginning of a string.

REGEXP_REPLACE replaces the characters 'ENG' at the beginning of the field COUNTRY with the replacement string 'SCOT'.

```
REGEXP_REPLACE(COUNTRY, '^ENG', 'SCOT')
```

For 'ENGLAND', the result is 'SCOTLAND'.

REGEXP_SUBSTR: Returning the First Match to a Pattern in a String

REGEXP_SUBSTR returns a string that contains the first match to a specified regular expression pattern within a source string. If there is no match within the source string, a null string is returned.

Syntax: How to Returning the First Match to a Pattern in a String

```
REGEXP_SUBSTR(string, pattern)
```

where:

string

Alphanumeric

Is the input string to be searched.

pattern

Alphanumeric

Is the regular expression pattern to match.

Example: Returning the First Match of a Pattern in a String

The following example uses the following Regular Expression symbols.

- ❑ [A-Z], which matches any uppercase letter.
- ❑ \$, which searches for a specified expression that occurs at the end of a string.

REGEXP_SUBSTR searches for a string with any uppercase letter followed by the characters 'umpty' at the end of the string 'Humpty Dumpty'.

```
REGEXP_SUBSTR('Humpty Dumpty', '[A-Z]umpty$')
```

The result is 'Dumpty'.

Support for Functions Used in ODBC Connector Client Tools

FOCUS and the SQL translator now support the following functions used by client tools with the ODBC Connector.

FOCUS and SQL Functions

Function Name	Description	Syntax
ASCII	Returns the ASCII code value of the leftmost character of a character expression.	<i>ASCII(charexp)</i> For example, the following returns the value 65. <i>ASCII('A')</i>
DAYNAME	Returns a character string that contains the data-source-specific name of the day for the day part of a date expression.	<i>DAYNAME(date_exp)</i> For example, the following returns Monday: <i>DAYNAME('August 3, 2020')</i>

Function Name	Description	Syntax
DIFFERENCE	<p>Returns an integer value measuring the difference between the SOUNDEX or METAPHONE values of two different character expressions. Zero (0) represents the least similarity. For SOUNDEX, 4 represents the most similarity, and for METAPHONE, 16 represents the most similarity.</p> <p>The use of SOUNDEX or METAPHONE depends on the PHONETIC_ALGORITHM setting. METAPHONE is the default algorithm.</p>	<p><code>DIFFERENCE(<i>chrexpl</i>, <i>chrexpl2</i>)</code></p> <p>For example, the following returns the value 4 when SOUNDEX is the phonetic algorithm:</p> <p><code>DIFFERENCE('Green', 'Greene')</code></p>
LEFT	<p>Given a character string, or an expression that can be converted to varchar, and an integer number, returns that number of characters from the left of the string.</p>	<p><code>LEFT(<i>chr_exp</i>, <i>int_exp</i>)</code></p> <p>For example, the following returns the value <i>ab</i>:</p> <p><code>LEFT('abcdefg', 2)</code></p>
LOG10	<p>Returns the base-10 logarithm of a numeric expression.</p>	<p><code>LOG10(<i>num_exp</i>)</code></p> <p>For example, the following returns the value 2.161:</p> <p><code>LOG10(145)</code></p>
MONTHNAME	<p>Returns a character string that contains the data-source-specific name of the month for the month part of a date expression.</p>	<p><code>MONTHNAME(<i>date_exp</i>)</code></p> <p>For example, the following returns August:</p> <p><code>MONTHNAME('August 3, 2020')</code></p>

Function Name	Description	Syntax
OVERLAY	Given a starting position, length, source string, and insertion string, replaces the number of characters defined by <i>length</i> in the source string with the insertion string, starting from the starting position.	<pre>OVERLAY(<i>src, ins, start, len</i>)</pre> <p>For example, the following returns SCOTLAND by replacing the first 3 characters in ENGLAND with the characters SCOT:</p> <pre>OVERLAY('ENGLAND', 'SCOT', 1, 3)</pre>
POSITION	Given a search string, a source string, and a starting position, returns the position of the search string within the source string. The search starts at the given starting position. If the string is not found, returns zero (0). The search is case sensitive.	<pre>POSITION(<i>search, source, start</i>)</pre> <p>For example, when CustomerName is Sandra Arzola, the following returns 8:</p> <pre>POSITION('A', CustomerName, 3)</pre>
REPEAT	Given a source string and an integer number, returns a string with the source string repeated that number of times, each repetition separated from the previous one with a space.	<pre>REPEAT(<i>source_str, number</i>)</pre> <p>For example, when FIRST_NAME is MARY, the following returns the string MARY MARY MARY:</p> <pre>REPEAT(FIRST_NAME, 3)</pre>
RIGHT	Given a character string, or an expression that can be converted to varchar, and an integer number, returns that number of characters from the right of the string.	<pre>RIGHT(<i>char_exp, integer_exp</i>)</pre> <p>For example, the following returns the value fg:</p> <pre>RIGHT('abcdefg', 2)</pre>

Function Name	Description	Syntax
ROUND	Given a numeric expression and an integer count, returns the numeric expression rounded to that number of decimal places. If the number of decimal places is negative, it rounds to the left of the decimal point.	<p><code>ROUND(<i>num_exp</i>, <i>count</i>)</code></p> <p>For example, the following returns 1.23500.</p> <p><code>ROUND(1.23456, 3)</code></p>
SIGN	Given a numeric expression, returns the value 1 if it is positive, or -1 if it is negative. SIGN(0) returns 0.	<p><code>SIGN(<i>num_exp</i>)</code></p> <p>For example, the following returns 1.</p> <p><code>SIGN(1.23456)</code></p>
SPACE	Given an integer count, returns a string consisting of that number of spaces.	<p><code>SPACE(<i>count</i>)</code></p> <p>For example, the following returns a string consisting of two spaces.</p> <p><code>SPACE(2)</code></p>
TRUNCATE	Truncates a numeric expression to a given number of decimal places. If the number of decimal places is negative, the number is truncated to the left of the decimal point.	<p><code>TRUNCATE(<i>num_exp</i>, <i>count</i>)</code></p> <p>For example, the following returns 1.23400.</p> <p><code>TRUNCATE(1.23456, 3)</code></p>

SQL Functions

Function Name	Description	Syntax
CHR	Takes a number as an argument and returns the ASCII character.	<p><code>CHR(<i>number</i>)</code></p> <p>For example, the following returns ¾.</p> <p><code>CHR(190)</code></p>

Function Name	Description	Syntax
LOCATE	Given a substring, a source string and a starting position (the default is 1), returns the position of the first occurrence of the substring, starting the search at the starting position. If the substring is not found, returns zero (0). The search is case insensitive.	<pre>LOCATE(<i>substr</i>, <i>source</i> [<i>,start</i>])</pre> <p>For example, when CustomerName is Sandra Arzola, the following returns 6:</p> <pre>LOCATE('a', CustomerName, 3)</pre> <p>The following returns 2:</p> <pre>LOCATE('a', CustomerName)</pre>

Support for Standard Deviation in PARTITION_AGGR

The PARTITION_AGGR function generates rolling calculations based on a block of rows from the internal matrix of a TABLE request. Population Standard Deviation (STDP) and Sample Standard Deviation (STDS) have been added as operations for the rolling calculation.

Note: Using the STDS or STDP aggregation operators requires that the request use the PRINT display command to avoid duplicate aggregation steps.

The syntax is:

```
PARTITION_AGGR([prefix.]measure, reset_key, lower, upper, operation)
```

where:

prefix.

Defines an aggregation operator to apply to the measure before using it in the rolling calculation. Valid operators are:

- SUM.** which calculates the sum of the measure field values. SUM is the default operator.
- CNT.** which calculates a count of the measure field values.
- AVE.** which calculates the average of the measure field values.
- MIN.** which calculates the minimum of the measure field values.
- MAX.** which calculates the maximum of the measure field values.
- FST.** which retrieves the first value of the measure field.

- ❑ **LST.** which retrieves the last value of the measure field.
- ❑ **STDP.** which retrieves the population standard deviation of the measure field.
- ❑ **STDS.** which retrieves the sample standard deviation of the measure field.

Note: The operators PCT., RPCT., TOT., MDN., and DST. are not supported. COMPUTEs that reference those unsupported operators are also not supported.

measure

Is the measure field to be aggregated. It can be a real field in the request or a calculated value generated with the COMPUTE command, as long as the COMPUTE does not reference an unsupported prefix operator.

reset_key

Identifies the point at which the calculation restarts. Valid values are:

- ❑ The name of a sort field in the request.
- ❑ **PRESET**, which uses the value of the PARTITION_ON parameter, as described in *Specify the Partition Size for Simplified Statistical Functions*.
- ❑ **TABLE**, which indicates that there is no break on a sort field.

The sort field may use BY HIGHEST to indicate a HIGH-TO-LOW sort. ACROSS COLUMNS AND is supported. BY ROWS OVER and FOR are not supported.

lower

Identifies the starting point for the rolling calculation. Valid values are:

- ❑ **n, -n**, which starts the calculation *n* rows forward or back from the current row.
- ❑ **B**, which starts the calculation at the beginning of the current sort break (the first line with the same sort field value as the current line).

upper

Identifies the ending point of the rolling calculation. The *lower* row value must precede *upper* row value.

Valid values are:

- ❑ **C**, which ends the rolling calculation at the current row in the internal matrix.
- ❑ **n, -n**, which ends the calculation *n* rows forward or back from the current row.
- ❑ **E**, which ends the rolling calculation at the end of the sort break (the last line with the same sort value as the current row.)

Note: The values used in the calculations depend on the sort sequence (ascending or descending) specified in the request. Be aware that displaying a date or time dimension in descending order may produce different results than those you may expect.

operation

Specifies the rolling calculation used on the values in the internal matrix. The new supported operations are:

- ❑ **STDP.** which calculates a population standard deviation.
- ❑ **STDS.** which calculates a sample standard deviation.

Example: Using **PARTITION_AGGR** to Calculate a Population Standard Deviation

The following request uses the STDP aggregation operator in PARTITION_AGGR to calculate the standard deviation for each category.

```
TABLE FILE ggsales
PRINT DOLLARS
COMPUTE STDP1/D12.2M = PARTITION_AGGR(DOLLARS, CATEGORY, B, E, STDP);
BY CATEGORY
BY PRODUCT
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

Partial output is shown in the following image.

<u>Category</u>	<u>Product</u>	<u>Dollar Sales</u>	<u>STDP1</u>
Coffee	Capuccino	20805	\$6,358.13
		20748	\$6,358.13
		20376	\$6,358.13
		20028	\$6,358.13
		19905	\$6,358.13
...			
Food	Biscotti	18200	\$6,565.19
		18084	\$6,565.19
		17100	\$6,565.19
		16918	\$6,565.19
		16656	\$6,565.19
...			
Gifts	Coffee Grinder	7752	\$4,518.06
		7715	\$4,518.06
		7623	\$4,518.06
		7485	\$4,518.06

Raised Limit for Join Fields

In prior releases, the limit for field pairs in a join was 20. The number of join field pairs has now been raised to 128.

Standard Deviation Prefix Operators: STDP. and STDS.

The standard deviation prefix operators return a numeric value that represents the amount of dispersion in the data. The set of data can be specified as the entire population (STDP.) or a sample (STDS.). The standard deviation is the square root of the variance, which is a measure of how observations deviate from their expected value (mean). If specified as a population, the divisor in the standard deviation calculation (also called degrees of freedom) will be the total number of data points, N. If specified as a sample, the divisor will be N-1.

If x_i is an observation, N is the number of observations, and μ is the mean of all of the observations, the formula for calculating the standard deviation for a population is:

$$\sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - \mu)^2}$$

To calculate the standard deviation for a sample, the mean is calculated using the sample observations, and the divisor is $N-1$ instead of N .

Syntax: **How to Calculate the Standard Deviation Using Prefix Operators**

To calculate the standard deviation for a population, the syntax is:

STDP.field

To calculate the standard deviation for a sample, the syntax is:

STDS.field

where:

field

Numeric

Is the set of observations for the standard deviation calculation.

Example: **Calculating the Standard Deviation of a Population**

The following request calculates the standard deviation of the population of the DOLLARS field converted to double precision.

```
DEFINE FILE ibisamp/ggsales
DOLLARS/D12.2 = DOLLARS;
END
TABLE FILE ibisamp/ggsales
SUM DOLLARS STDP.DOLLARS
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

The output is shown in the following image.

<u>DOLLARS</u>	STDP <u>DOLLARS</u>
46,156,290.00	6,156.997845651

Example: Calculating the Standard Deviation of a Sample

The following request calculates the standard deviation of a sample of the DOLLARS field converted to double precision.

```
DEFINE FILE ibisamp/ggsales
DOLLARS/D12.2 = DOLLARS;
END
TABLE FILE ibisamp/ggsales
SUM DOLLARS STDS.DOLLARS
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

The output is shown in the following image.

<u>DOLLARS</u>	STDS <u>DOLLARS</u>
46,156,290.00	6,157.711080272

Enhancements to DATE_ORDER

The DATE_ORDER parameter for date-time formats is now supported when the format specifies month translation, zero suppression or removal, or the comma option. For some formats with the comma option, reordering the date components may require elimination of the comma.

For example, the following request creates date fields with zero removal and suppression (YYMDoe), month translation and zero suppression (YYMte), and month translation with a comma and space between the month and year (HMTkYY). The DATE_ORDER is set to DMY:

```
-DEFAULT &ORDER=DMY;
SET DATE_ORDER=&ORDER
DEFINE FILE GGSales
ORIGINAL/YYMD=20190704;
YYMDoe/HYYoe=DT(2019/07/04);
YYMte/HYYMte=YYMD;
YYMTDk/HMTkYY=YYMD;
END
TABLE FILE GGSales
SUM ORIGINAL YYMDoe YYMte YYMTDk
BY CATEGORY
ON TABLE SET PAGE NOLEAD
ON TABLE SET STYLE *
GRID=OFF,$
ENDSTYLE
END
```

The output is shown in the following image.

<u>Category</u>	<u>ORIGINAL</u>	<u>YYMDoe</u>	<u>YYMte</u>	<u>YYMTDek</u>
Coffee	04/07/2019	4/7/2019	4 Jul 2019	July, 2019
Food	04/07/2019	4/7/2019	4 Jul 2019	July, 2019
Gifts	04/07/2019	4/7/2019	4 Jul 2019	July, 2019

Dropped Support

This section describes changes in supported adapters, operating systems, and programs.

In this chapter:

- [Dropped Support: Adapter Releases](#)
- [Dropped Support: Operating System Releases](#)
- [DataMigrator: Dropped Support for Perforce](#)

Dropped Support: Adapter Releases

As of Release 8207, support has been dropped for the following adapters.

Adapter Category	Adapter Name	Adapter Release Desupported as of 8207	Affected OS Version	Nearest Adapter Release Supported in 8207
DBMS	Adabas	6.x.x	Windows	None
ERP	Axiom EPM	All Relevant	Many	None
Files	Informix CISAM	7.2.5 and 7.2.6	Windows	None
SQL	C9	(Unicode Optional)	Many	None
SQL	Informix	SDK 3.0 and SDK 3.5	Windows (32bit)	SDK 3.7 on Windows x86_64bit
SQL	Informix	SDK 3.0 and SDK 3.5	AIX, HP-UX, Solaris and Windows 64bit	SDK 3.7

Adapter Category	Adapter Name	Adapter Release Desupported as of 8207	Affected OS Version	Nearest Adapter Release Supported in 8207
SQL	Interplex/ Unisys DMS 1100	All Relevant	Windows	None
SQL	UniVerse	All Relevant	Windows	None

Dropped Support: Operating System Releases

As of Release 8207, support has been dropped for the following operating system releases.

Operating System	OS Version DeSupported in 8207	Nearest OS Version Supported in 8207
HP-UX	HP-UX 11 31 and up IA64 64bit	None
OpenVMS	OpenVMS V84 and up IA64 32on64bit	None
OS400	V7R2 and up iSeries	V7R3 and up iSeries
Windows	Windows (this was a 32bit version)	Windows x86_64 64bit
z/OS™	z/OS 2.1 and up zSeries 64bit	z/OS 2.2 and up zSeries 64bit

DataMigrator: Dropped Support for Perforce

The Perforce source control system is not supported in DataMigrator, starting with release 7709.

Chapter 3

Upgrade Notes

This section describes changes in behaviour in this release.

In this chapter:

- [Golden Key License Mode](#)
 - [Running Scheduled Flows](#)
 - [Converting an Existing Server Installation to the Golden Key Configuration](#)
 - [SSL: New OpenSSL Configuration Requirement](#)
 - [Software Branding](#)
 - [Data Flow: Warning When Closing Join, Union, SQL, and Target Business View Editors](#)
 - [Technical Content Branding](#)
 - [Data Preparation: Properties Panel Redesign](#)
 - [Cloud](#)
 - [Data Preparation: Server Console Menu](#)
 - [Microsoft Azure SQL Data Warehouse Renamed](#)
 - [Data Profiling Charts and Reports Display in the Output Window](#)
 - [Data Preparation Upgrade Considerations](#)
-

Golden Key License Mode

In release 8207.27, the new Golden Key license replaces the legacy License Key as the default product license. This license automatically allows administrators to maintain an unlimited number of users, adapters, and CPUs, grants unlimited use of all features to all users, and imposes no consumption limits. Under this license key, administrators can use security settings and individual roles to limit the access of groups and users to individual products and features.

The default ports assigned when you install a Reporting Server are six consecutive port numbers starting from 8120. The default TCP Listener port is 8120, and the default HTTP Listener port is 8121. You can change the port assignments during the installation process, if you want to use different port numbers. If you configure additional server instances, you will need to specify a different set of ports for each instance, as in prior releases.

As of release 8207.27, the Golden Key license is assigned to all new product installations automatically. Customers who upgrade to 8207.27 from an earlier release continue to maintain their Legacy license. They can change to the Golden Key license with support from the Customer Service team.

Converting an Existing Server Installation to the Golden Key Configuration

In Release 8207.27.0, the new Golden Key license replaces the legacy License Key as the default product license. The Golden Key license is assigned to new product installations automatically, and the installation procedures do not prompt for a license key.

This new license policy is actually a lack of keys for features that previously required them. This also means there is no longer a differentiation between a Full Function Server, a WebFOCUS Server, and a Data Migrator Server.

This does not mean that all features are automatically enabled, as some features require configuration settings in order to be enabled, or they would cause overhead for a feature a customer might not want.

An upgrade of an existing pre-8207.27.0 Reporting Server installation only updates the software portion of a prior installation (files under EDAPHOME), not the configuration portion (files under EDACONF). At Server or Client run-time, the 8207.27.0 (or higher) software will recognize the use of license keys for a prior configuration, and run with the applicable features enabled.

For WebFOCUS, you can manually enable the Golden Key after upgrading, by removing the following files from the /ibi/WebFOCUS82/config default folder:

- license.cfg
- wflicense.key

For the Reporting Server, you can manually convert the EDACONF configuration files to the Golden Key, if you desire.

It is recommended that you do the conversion on an existing 8207.27 or higher release, so that you have a prior release to use if the conversion is not completed. You can do the conversion with the server or client running, but if it is running, it must be restarted when you complete the steps.

Use the following steps to convert the configuration to the Golden Key configuration.

1. Backup the following files under EDACONF, in case of conversion failure.
 - edaserve.cfg in the bin folder.
 - odin.cfg in the etc folder.

2. Edit edaserve.cfg in the bin folder:
 - Remove the following lines:
 - [Licenses] section label line.
 - license = *license_key* line.All license_*feature* = *key* lines, where *feature* is the code for the feature being licensed, and *key* is the license key for that feature.
 - Change workspace_manager_name to one of the following, depending on whether you installed a Reporting Server, a Data Migrator client, or an Open Visualization client.
 - "WebFOCUS 82 Server"
 - "WebFOCUS 82 DM Client"
 - "WebFOCUS 82 ODBC Client"
 - Save the edaserve.cfg file.
3. Edit the odin.cfg file in the etc folder:
 - Remove the ;enable CLM section label line.
 - Remove the LICENSE = *key* line.
 - Save the odin.cfg file.
4. Start the server, or restart it if it is already running.

Software Branding

As of the following releases, ibi software and technical content are now branded under TIBCO® Software Inc.

- WebFOCUS:** 8207.27.0
- FOCUS:** 8207.27.0
- iWay Service Manager:** 8.0.5
- Omni-Gen:** 3.16.0

This change only impacts the names to which these products are referred. For example, WebFOCUS is now known as TIBCO WebFOCUS[®], while iWay DataMigrator is now known as TIBCO[®] Data Migrator. You will begin to see this change throughout the software and corresponding technical content assets, including PDF covers and KnowledgeBase collections, where both new and former product names will be used interchangeably. For a full list of software branding, see below.

Former Product Name	New Product Name
WebFOCUS	TIBCO WebFOCUS [®]
WebFOCUS Client	TIBCO WebFOCUS [®] Client
WebFOCUS Server	TIBCO WebFOCUS [®] Reporting Server
iWay Data Migrator Server	TIBCO WebFOCUS [®] Reporting Server
iWay DataMigrator	TIBCO [®] Data Migrator
WebFOCUS App Studio	TIBCO WebFOCUS [®] App Studio
WebFOCUS Mobile App	TIBCO WebFOCUS [®] Mobile App
WebFOCUS Infographics	TIBCO WebFOCUS [®] Infographics
WebFOCUS Narrative Charts	TIBCO WebFOCUS [®] Narrative Charts
iWay Service Manager	TIBCO iWay [®] Service Manager
Omni-Gen	TIBCO Omni-Gen [®]
Omni Master Data Management	TIBCO Omni-Gen [®] MDM
Omni for Customer	TIBCO Omni-Gen [®] MDM
Omni-HealthData	TIBCO Omni-HealthData [®]
Omni-Insurance	TIBCO Omni-Insurance [™]
Data Quality Edition	TIBCO Omni-Gen [®] DQ Edition
iWay Data Quality Server	TIBCO Omni-Gen [®] DQ Server
FOCUS	TIBCO FOCUS [®]

Technical Content Branding

As of Release 8207.27, ibi products are now branded under TIBCO® Software Inc. You will begin to see this change throughout the software and corresponding technical content assets, including PDF covers and KnowledgeBase collections, where both new and former product names will be used interchangeably. For a full list of software branding, see [Software Branding](#) on page 309.

The following tables provide a quick overview of the PDF titles that have changed to support this release. Note that the way you access these technical content offerings and information remains the same. As we continue the branding effort, you may notice additional changes within our technical content, which we will continue to communicate to you.

Former Title	New Title
TIBCO WebFOCUS®	
<i>WebFOCUS App Studio Installation and Configuration Guide</i>	<i>App Studio Installation and Configuration</i>
<i>WebFOCUS App Studio User's Manual</i>	<i>App Studio User's Manual</i>
<i>Creating HTML5 Charts With WebFOCUS Language</i>	<i>Creating HTML5 Charts With TIBCO WebFOCUS® Language</i>
<i>Creating Reports With WebFOCUS Language</i>	<i>Creating Reports With TIBCO WebFOCUS® Language</i>
<i>Describing Data With WebFOCUS Language</i>	<i>Describing Data With TIBCO WebFOCUS® Language</i>
<i>WebFOCUS Embedded Business Intelligence User's Guide</i>	<i>Embedded Business Intelligence User's Guide</i>
<i>WebFOCUS InfoAssist User's Manual</i>	<i>InfoAssist User's Manual</i>
<i>WebFOCUS and ReportCaster Installation and Configuration for UNIX</i>	<i>Installation and Configuration for UNIX</i>
<i>WebFOCUS and ReportCaster Installation and Configuration for Windows</i>	<i>Installation and Configuration for Windows</i>
<i>WebFOCUS Integrated Installation</i>	<i>Integrated Installation</i>

<i>Using the WebFOCUS Mobile App for iOS and Android</i>	<i>Mobile App for iOS and Android</i>
<i>Server Administration</i>	<i>Reporting Server Administration</i>
<i>Server Installation</i>	<i>Reporting Server Installation</i>
<i>Server Release Notes</i>	<i>Reporting Server and TIBCO® Data Migrator Release Notes</i>
<i>WebFOCUS Security and Administration</i>	<i>Security and Administration</i>
<i>WebFOCUS Security and Administration Best Practices</i>	<i>Security and Administration Best Practices</i>
<i>Using WebFOCUS</i>	<i>User's Guide</i>
TIBCO® Data Migrator	
<i>Server Release Notes</i>	<i>Reporting Server and TIBCO® Data Migrator Release Notes</i>
<i>iWay DataMigrator User's Guide</i>	<i>User's Guide</i>

Cloud

The following is an upgrade consideration for Cloud instances:

- Upon your first login to the Cloud environment or after upgrading to the latest version, you will be prompted to read and accept the End User License Agreement, before you can proceed with the product.

Microsoft Azure SQL Data Warehouse Renamed

Microsoft has renamed its Azure Data Warehouse product to Azure Synapse Analytics. The adapter has been renamed to match the change in the product name.

Running Scheduled Flows

When creating or editing a Data Flow in the Web Console, you have the option to schedule the flow to run in the future.

In order for the flow to run on schedule, it must be saved in a directory on the application path of the scheduler (the user ID set to `sched_scan_id`).

SSL: New OpenSSL Configuration Requirement

OpenSSL libraries are loaded dynamically at run time on UNIX and Windows. You are responsible for installing these libraries and making them available to the Server by specifying their location in the `PATH` or in the `IBI_LOAD_SSL_FROM` or `IBI_LIBPATH` environment variable value at Server startup time.

Open SSL libraries v1.1.1 and v1.0.2 are supported.

The current run-time version depends on the following configuration keyword in `edaserve.cfg`.

```
ssl_lib_vendor = {sslv11|sslv10}
```

The default value for `ssl_lib_vendor` is `sslv11`.

You should also set the environment variable `IBI_LOAD_SSL_FROM` or `IBI_LIBPATH` to point to the appropriate openSSL library.

For example:

```
ssl_lib_vendor = sslv11  
IBI_LIBPATH=/usr/local/ssl64/1.1.1/lib
```


or:

```
ssl_lib_vendor = sslv10  
IBI_LIBPATH=/usr/local/ssl64/1.0.2/lib
```

Note: OpenSSL-1.1.1 is available by default on Red Hat Enterprise Linux (RHEL) 8 and above but must be built from source on RHEL 7 and below.

Data Flow: Warning When Closing Join, Union, SQL, and Target Business View Editors

When using Data Preparation and editing the select statement using the Join, Union, SQL, or Target Business View editor, then clicking the X (Close) icon, any changes were discarded as though you had pressed cancel. Now, if any changes are made, a confirmation prompt opens asking you to confirm whether you want to discard your changes, as shown in the following image.



Data Preparation: Properties Panel Redesign

The Properties Panel for entering parameters when configuring a connection or uploading data has been redesigned with a cleaner and more consistent user interface, to make it easier to use and to provide clearer on-screen help.

A tooltip is available when you hover over the help icon (?) for each property, as shown in the following image.

Add Connection for MS SQL Server OLE DB/AzureDB Prerequisites

Connect parameters

Connection Name: CON02

Server: [Name of the machine where Server is running]

Security: Explicit

User:

Password:

Default Database:

Additional connection string keywords (Optional):

Show Sample

Advanced Parameters

Environment:

Select profile: sqlasmnf

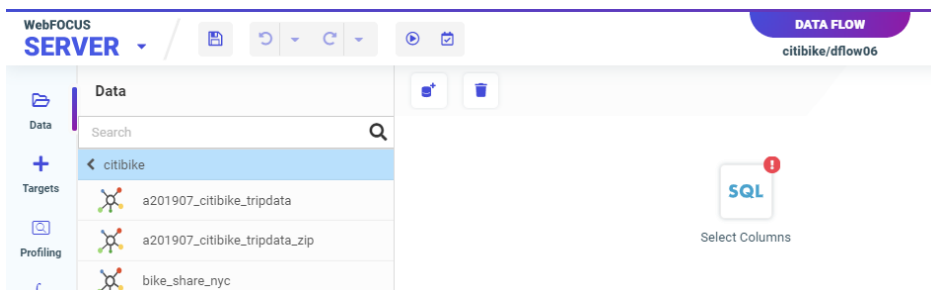
Test Configure

If you click the help icon, more detailed help is displayed.

For some parameters, a sample value is available when you click *Show Sample*.

Data Preparation: Server Console Menu

In the Metadata and Data Flow editors, to indicate that the WebFOCUS Server logo on the Server Console has a menu, a down arrow has been added, as shown in the following image.



Data Profiling Charts and Reports Display in the Output Window

In the Metadata Editor, all data profiling charts and reports now open on a new tab in the output pane. In prior releases, some charts opened in a separate pane.

Data Preparation Upgrade Considerations

The following are release considerations and product changes for WebFOCUS.

- ❑ In prior releases, when you created a synonym containing date or date-time columns, eight additional virtual (DEFINE) fields were automatically generated for each date component and combination of components.

This could result in much larger synonyms that took longer to open and edit. This option is now disabled by default. If you want to have the date components generated, you can enable Decompose Dates on the Get Data page. To enable Decompose Dates so that it is always on when creating a synonym, click the link for *Common Adapter Settings*. Under the *Customize datatype mappings* section, for *DECOMPOSE-DATE* select *On*.

- ❑ The Expression Editor has an area that displays objects you can add to an expression, which are columns, variables, or functions. The selection and display of these was controlled by buttons above the area, but when the size of the area was compressed, some of the buttons were hidden. Now there is a single button with a drop-down menu that allows selecting what objects are displayed.
- ❑ From the Designer Data tab or the Web Console, when *Show title* is selected from the *View* shortcut menu, if a synonym does not have a field title for a column, the default column name is the field name, not the alias. This allows you to see the same column names in Data Preparation as in Designer and InfoAssist.
- ❑ When you join files or tables without keys, the join created automatically based on matching field names is now limited to three pairs of join fields, to avoid attempting a join on an excessive number of fields, especially when joining a table to itself.
- ❑ From the Designer Data tab or the Web Console, when you reload the profile charts after you have horizontally scrolled to and edited a profile chart field, the scroll bar is positioned so that the edited profile chart remains visible. You do not have to scroll to see the edited profile chart.

Known Issues

This section describes known issues in this release.

In this chapter:

- [New Licenses For Second Level Geographies Not Yet Supported](#)
 - [Client Silent Installation Failures on Windows](#)
-

New Licenses For Second Level Geographies Not Yet Supported

In WebFOCUS Release 8207.27.0, new licenses for second level geographies are not yet supported. Existing customers who have already licensed second level geographies can continue to use them.

Client Silent Installation Failures on Windows

The Windows silent installations for `setup_dm_client` (the TIBCO Data Migrator client installation) and `setup_odbc_client` (the Open Visualization client installation) fail when there is a space in the path to the option file.

To avoid this failure, make sure the path to the option file does not contain a blank space.

Fixes

This section describes fixed issues.

In this chapter:

- ❑ [Entering Network Path in File Picker for Excel Causes Error](#)
-

Entering Network Path in File Picker for Excel Causes Error

When using the Data Management Console (DMC) and selecting an application directory, the file picker allowed typing in a file name instead of making a selection. When a network file path was entered, this resulted in an error.

Now, entering a name is not allowed. You must make a selection, and these errors do not occur.

Legal and Third-Party Notices

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 DOCUMENT 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 DOCUMENT IS SUBJECT TO THOSE 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 TIBCO Software Inc.

TIBCO, the TIBCO logo, the TIBCO O logo, FOCUS, iWay, Omni-Gen, Omni-HealthData, and WebFOCUS 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.

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 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. TIBCO SOFTWARE INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENT AT ANY TIME.

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 "READ ME" FILES.

This and other products of TIBCO Software Inc. may be covered by registered patents. Please refer to TIBCO's Virtual Patent Marking document (<https://www.tibco.com/patents>) for details.

Copyright © 2021. TIBCO Software Inc. All Rights Reserved.

Additional Third-Party Legal Notices

Please note that the component dmcgui.exe is subject to the terms of the TIBCO End User License Agreement and is not controlled by the terms of the PostgreSQL License. Please disregard any notice to the contrary.

TIBCO WebFOCUS(R) includes the following open source components, which are subject to the licenses indicated:

Serp 1.15.1

Copyright (c) 2002-2007, A. Abram White
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of 'serp' nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

commons-pool:commons-pool:jar: 1.6
Apache License 2.0 [full text below]

org.apache.openjpa:openjpa: jar:2.4.2
Apache License 2.0 [full text below]

Apache License

Version 2.0, January 2004

<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. **Definitions.**

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, " control " means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or **"Your"**) shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. **Grant of Copyright License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. **Grant of Patent License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. **Redistribution.** You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - a. You must give any other recipients of the Work or Derivative Works a copy of this License; and
 - b. You must cause any modified files to carry prominent notices stating that You changed the files; and
 - c. You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

-
- d. If the Work includes a **"NOTICE"** text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. **Disclaimer of Warranty.** Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

TOC.js

The MIT License (MIT)

Copyright (c) 2013 Patrice Freydiere

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PARTICULAR PURPOSE AND NON INFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE.

JsonCoverters.js

Copyright 2012 Environmental Systems Research Institute, Inc. (Esri)

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at: <http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Apache License - 2.0

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or **"Your"**) shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensors for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensors or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensors for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensors and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensors and subsequently incorporated within the Work.

2. **Grant of Copyright License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. **Grant of Copyright License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
4. **Grant of Patent License.** Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
5. **Redistribution.** You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - a. You must give any other recipients of the Work or Derivative Works a copy of this License; and
 - b. You must cause any modified files to carry prominent notices stating that You changed the files; and

-
- c. You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
- d. If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License. You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.
6. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
7. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
8. **Disclaimer of Warranty.** Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
9. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

10. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

