



1010data Insights Platform User's Guide

Contents

About this guide.....	10
Documentation conventions.....	10
Related publications.....	10
Documentation feedback.....	10
Copyright.....	11
Getting started.....	12
System requirements.....	12
Log in.....	13
Reset your password.....	14
Understanding the user interface.....	16
Workspace menu.....	18
Window tray.....	18
Windows.....	19
Window controls.....	21
Rename a window.....	21
Keyboard shortcuts.....	22
Keyboard navigation.....	26
Obtaining technical assistance.....	28
Help menu.....	28
Submit a support request.....	29
Contact 1010data Support.....	30
Log out.....	31
User basics.....	32
Versions.....	32
Sessions.....	32
Workspace.....	33
Save your workspace.....	34
Open a saved workspace.....	36
Add a new workspace.....	37
Workspace Settings.....	37
Appearance.....	38
Associations.....	39
Object Manager.....	41
Libraries.....	42
Customize your workspace appearance.....	42
Define object associations.....	43
Select default metadata columns.....	44
Specify automatically imported libraries.....	45
Virtual memory.....	46
Clear the cache.....	48
Rows and columns.....	48
Rows.....	48
Columns.....	48
Column information.....	49
Data types and display formats.....	52
Data types.....	53

Integer.....	54
Big integer.....	54
Decimal.....	54
Text.....	55
Dates and time.....	55
Display formats.....	56
Currency codes (ISO 4217).....	59
Dashboard.....	70
Recent Workspaces.....	71
Open a recent workspace.....	72
Favorites.....	72
Remove a favorite.....	73
User Profile.....	75
Profile.....	75
Profile Information.....	76
Password.....	77
Password requirements.....	78
Settings.....	79
Localization.....	79
Session Preferences.....	80
Data Preferences.....	81
Edit your user profile.....	83
Change your password.....	84
Change your user settings.....	85
Object Manager.....	87
Object browser.....	88
Object icons.....	89
Permission icons.....	90
My Data and Other People's Data.....	90
Metadata columns.....	91
Object actions by type.....	93
Choose the displayed metadata columns.....	96
Sort the object list.....	98
Open the right-side object browser pane.....	100
Search.....	102
Search for objects.....	103
Information dialog.....	104
View object information.....	107
Edit object information.....	108
Share an object.....	109
Add an object to your favorites.....	110
Delete an object.....	112
Move an object.....	112
Reorder objects in a folder.....	113
Add a subfolder.....	114
Refresh an object.....	115
Select multiple objects.....	116
Open an object.....	117

Trillion-Row Spreadsheet.....	119
Analyze.....	120
Object browser.....	121
Favorites list.....	122
<table> panel.....	122
Foreign database panel.....	123
Open an existing table or query.....	124
Open a favorite table or query.....	124
Create a temporary table.....	125
Open a foreign table.....	127
Analysis Timeline.....	129
New operation.....	131
Operation controls.....	133
Insert a new operation.....	135
Delete an operation.....	135
View the previous operation state.....	136
View the next operation state.....	136
Edit an operation using Macro Language XML.....	136
Add an operation using Macro Language XML.....	137
Submit controls.....	138
Expression Editor.....	139
Filter a list of columns.....	140
Selecting rows.....	141
Select rows panel.....	142
Simple comparisons relationship options.....	143
Use simple comparisons to select rows.....	145
Specifying a list of values in a selection.....	147
Select rows by entering a selection expression.....	149
Performing a Quick selection using the grid.....	151
Select rows based on cell value.....	151
Select rows based on cell location.....	152
Select a group of cells.....	153
Sorting rows.....	155
Use the timeline to sort rows.....	156
Sort rows from the grid column header.....	157
Sort rows within the grid.....	158
Arranging columns.....	159
Change the column order.....	159
Rearrange columns using the grid.....	160
Use the timeline to hide or show columns.....	161
Hide a column within the grid.....	162
Linking tables and worksheets.....	163
Link tables panel.....	163
Include or exclude column options.....	166
Link types.....	167
Link in a table.....	168
Link in an existing worksheet.....	172
Link in a new worksheet.....	175
Computed column.....	180
Create a computed column.....	181
Summarizations and tabulations.....	183
Tabulation panel.....	184
Perform a quick summary.....	186
Summarization options in the grid menu.....	189

Perform a tabulation.....	190
Perform a cross tabulation.....	193
Change the tabulation function summarization.....	197
Types of summarizations.....	199
Use a rollup in a tabulation.....	204
Format the tabulation results.....	205
Tabulation result options.....	207
Drilldown analysis.....	209
Set up a drilldown hierarchy.....	214
Drilldown controls in the grid.....	215
Merging tables and worksheets.....	216
Merge tables panel.....	217
Merge in a table.....	219
Merge in an existing worksheet.....	220
Merge in a new worksheet.....	222
Amending values.....	225
Amend column panel.....	225
Amend values in a column using an expression.....	226
Amend a single value within the grid.....	227
Amend values panel.....	229
Running custom code.....	229
Code panel.....	230
Use custom code in your analysis.....	230
Miscellaneous operations.....	231
Expand rows.....	232
Split and expand.....	233
Perform a string replacement.....	234
Convert dates and times.....	235
Use a group metric.....	237
Types of metrics.....	240
Query.....	242
Edit a query using Macro Language XML code.....	244
View.....	244
Select a view.....	245
Visualize.....	246
Chart Builder.....	247
Data columns.....	247
Chart parameters.....	248
Chart types.....	250
Chart data.....	257
Interactions.....	258
Settings.....	259
Customization settings.....	269
General.....	270
Title.....	270
Axes.....	271
Ticks.....	275
Data Series.....	276
Grid.....	276
Legend.....	276
Info.....	277
Create a chart.....	277
Download a chart.....	278
Develop.....	279
Edit a query in the Macro Language Workshop.....	280
Save an analysis as a Quick Query.....	281

Consolidate timeline.....	283
Export.....	284
Delimited text export settings.....	285
Excel workbook export settings.....	287
PDF export settings.....	288
Download the data as a delimited text file.....	288
Download the data as an Excel workbook file.....	289
Download the data as a PDF file.....	290
Results pane.....	291
View bar.....	292
Grid view.....	294
Column header menu options.....	296
Editing column metadata information.....	297
Column metadata dialog.....	297
Change the metadata information of a column.....	299
Resize a column within the grid.....	300
Modify panel.....	301
Draw a selection rectangle.....	302
Grid menu options.....	302
Copy values in the grid.....	305
Copy menu options.....	306
Analyze column data.....	308
Quick analysis menu options.....	308
Quick statistics.....	309
Distribution.....	311
Median and mode.....	311
Finding rows in the grid.....	312
Find rows dialog.....	313
Find rows using a selection expression.....	313
Find rows based on cell value.....	315
Find rows by searching for a value within a column.....	319
Scroll directly to a specific row.....	323
Scroll directly to a specific column.....	324
Single-row view.....	325
Clone the TRS window.....	326
Table Editor.....	329
Create a table.....	330
Edit a table.....	331
View Query Results.....	332
Uploading tables.....	333
Supported file formats.....	333
Text file format.....	333
Prerequisites.....	334
Transfer a file to your FTP account.....	334
Upload a table (simple mode).....	337
Upload a table (detail mode).....	339
Structure.....	344
Headers.....	346
Headers section menu options.....	347
Fields.....	348
Destination.....	350
Spec XML.....	352

Load.....	354
Upload a table (advanced upload).....	355
File Information.....	360
Table Information.....	363
Column Information.....	364
Save Data to FTP.....	370
File.....	371
Query.....	373
Data.....	373
Manage Files.....	374
Send the data to your 1010data FTP directory.....	374
Delete a file from your 1010data FTP directory.....	375
Save Data as Table.....	377
Metadata.....	378
Access.....	379
Append.....	380
Structure.....	380
Query.....	381
Data.....	382
Save the data to a table.....	382
Quick Queries/PQQs.....	385
Block setup panel.....	385
Add parameters to an operation.....	387
Edit parameters in an operation.....	390
Edit a parameter with a large value selection.....	392
Remove all parameters from an operation.....	395
Edit a Parameterized Quick Query in MLW.....	397
Run a Quick Query.....	399
Run a Parameterized Quick Query.....	400
Legacy Quick Queries/PQQs.....	402
Edit and convert a legacy Parameterized Quick Query.....	402
Edit a legacy Parameterized Quick Query in MLW.....	405
Run a legacy Parameterized Quick Query.....	407
Macro Language Workshop.....	410
Menu bar.....	411
File.....	412
Code.....	415
Run.....	417
Debug.....	418
History icons.....	418
Open the Macro Language Workshop.....	419
Specify a base table.....	419
Open a saved query.....	420
Render a query.....	422
Save a query.....	422
Post-hoc Performance Analyst.....	424

Finished queries.....	424
MDB state cache.....	426
Temporary table cache.....	428
Physical tables.....	429
Gadgets.....	432
Memo Pad.....	432
Calculator.....	433
Evaluate a scalar expression.....	435
Tag Inspector.....	436
Image Encoder.....	438
Scheduler.....	440
Existing Scheduled Jobs and Selected Job's Run History.....	441
Scheduled Job List.....	442
Selected Job's Run History.....	443
Job Details.....	444
Basic Info.....	446
Base Table and XML.....	446
Job Timing.....	447
Output Options.....	449
Email Contents.....	450
Advanced settings.....	450
Monthly Run Summary.....	451
Run limit.....	452
Create a job.....	452
Find a job.....	456
Edit a job.....	456
Run a job.....	457
Clone a job.....	458
Delete a job.....	459
Download a completed job report.....	460
Stop a currently running job.....	460
Unsubscribe from a multi-user job.....	461
Deactivate a job.....	461
Admin Tool.....	463
User Manager.....	463
User Info.....	464
Account Settings.....	466
GUI Preferences.....	468
Edit your user information.....	473
Create a new user (admin only).....	473
Create multiple users (admin only).....	473
Bulk edit file format.....	474
Find existing users (admin only).....	478
Edit an existing user (admin only).....	478
Edit multiple existing users (admin only).....	479
Deactivate an existing user (admin only).....	479
Group Manager.....	480
Create a new group.....	480
Find an existing group.....	481
Edit an existing group.....	481

Delete an existing group.....	482
Company Manager (admin only).....	482
Edit company information (admin only).....	483
System Statistics (admin only).....	485
Appendix A: Missing values.....	487
Automatic handling of N/A values.....	488
Manual handling of missing or invalid values.....	489

About this guide

The *1010data Insights Platform User's Guide* covers many of the concepts and tasks that are helpful to the beginner and advanced user alike.

Welcome to the *1010data Insights Platform User's Guide*. This is where you can find answers to many of the questions you may have while working with the web interface of the 1010data Insights Platform.

Documentation conventions

This guide uses specific typographical conventions to describe how to use 1010data products.

Please be aware of the following conventions when using this guide.

Table 1: Typographic conventions

Convention	Example	Description
Bold text	Click Submit or press Enter .	The names of menus, buttons, and fields in the user interface, or keys on the keyboard, display in bold typeface.
Right angle bracket (>) character	Select File > New from the menu.	The menu path and menu selection display in bold typeface with embedded right angle bracket characters.
Courier Font	The Store Master table file name is: <code>certification.retail.stores.</code>	Names of tables (the table file name) and information you must enter into the 1010data Insights Platform display in Courier font.

Related publications

The 1010data Insights Platform core documentation set includes several resources to help you learn and use the platform.

The Insights Platform core documentation set includes:

- *1010data Insights Platform User's Guide*
- *1010data Reference Manual*
- *Trillion-Row Spreadsheet Getting Started Guide*
- *1010Cookbook*
- *1010data Glossary*

To obtain a related document, visit the [1010data Documentation Center](#).

Documentation feedback

We welcome your feedback and will use it to improve future versions of this guide.

It is the goal of 1010data to provide you with accurate and useful documentation. If you do not see the information you are looking for, let us know, and we'll do our best to add the topic in a future revision.

To provide feedback, send an email message with your comments to docs@1010data.com.

Please include the 1010data product name and version number, documentation topic title, and a brief description of the suggestion.

Copyright

© 2020 1010data, Inc. All rights reserved. 1010data and the 1010data logo are trademarks of 1010data, Inc. 1010data is registered in the United States. All trademarks identified by ®, ™, or SM are trademarks and service marks of 1010data, Inc. All other trademarks are the property of their respective owners.

1010data, Inc.
750 Third Avenue, 4th Floor
New York, NY 10017
(212) 405-1010
www.1010data.com

The following terms are trademarks, registered trademarks, or service marks of 1010data: 1010data®, 1010data Facts™, 1010data Shelf Monitor™, 1010data Supplier Portal™, 1010data Test Monitor™, Make Sense of It All®, QuickApps™, R1010™, TenDash™, TenUp™, Trillion-Row Spreadsheet®.

Getting started

Information to help new users get started with the 1010data Insights Platform.

System requirements

To optimize system performance and provide a high level of security, the 1010data Insights Platform requires adherence to specific set up protocols and security standards.

The following list outlines requirements to use the full functionality available in the platform.

Internet browser

We encourage you to use the latest version of a supported browser.

Supported browsers

The latest versions of the following browsers are supported:

- Apple Safari
- Google Chrome
- Microsoft Internet Explorer 11
- Mozilla Firefox

For the best experience, we recommend using the latest version of Google Chrome with the Insights Platform.

Browser requirements

Configure your Internet browser to accommodate the following:

- Enable JavaScript
- Allow cookies. See the 1010data [Privacy Policy](#) for information about our use of cookies.
- Allow pop-ups

Additional software

While not required to use the platform itself, you need additional software to work with files that you upload to and export from the system. For example, you can transfer files through FTP to upload into the Insights Platform or export data as a PDF file.

Note: 1010data does not recommend or support specific third-party software.

File compression software

Any file compression software. For example, WinZip.

PDF reader

Any PDF reader. For example, Adobe Acrobat Reader.

FTP client

Any third-party FTP client. For example, FileZilla.

Spreadsheet application

Any spreadsheet application. For example, Microsoft Excel.

Screen resolution

The 1010data Insights Platform requires a minimum screen resolution of 1100px x 800px.

Note: If your browser zoom level is set to anything other than actual size (100%), the platform user interface may not appear as intended.

Blocking

Check with the technical support department at your company to verify the following email addresses and domains are not blocked.

Domains	1010data.com
Email addresses	support@1010data.com

Log in

You must log in to the 1010data Insights Platform to access and analyze your data.

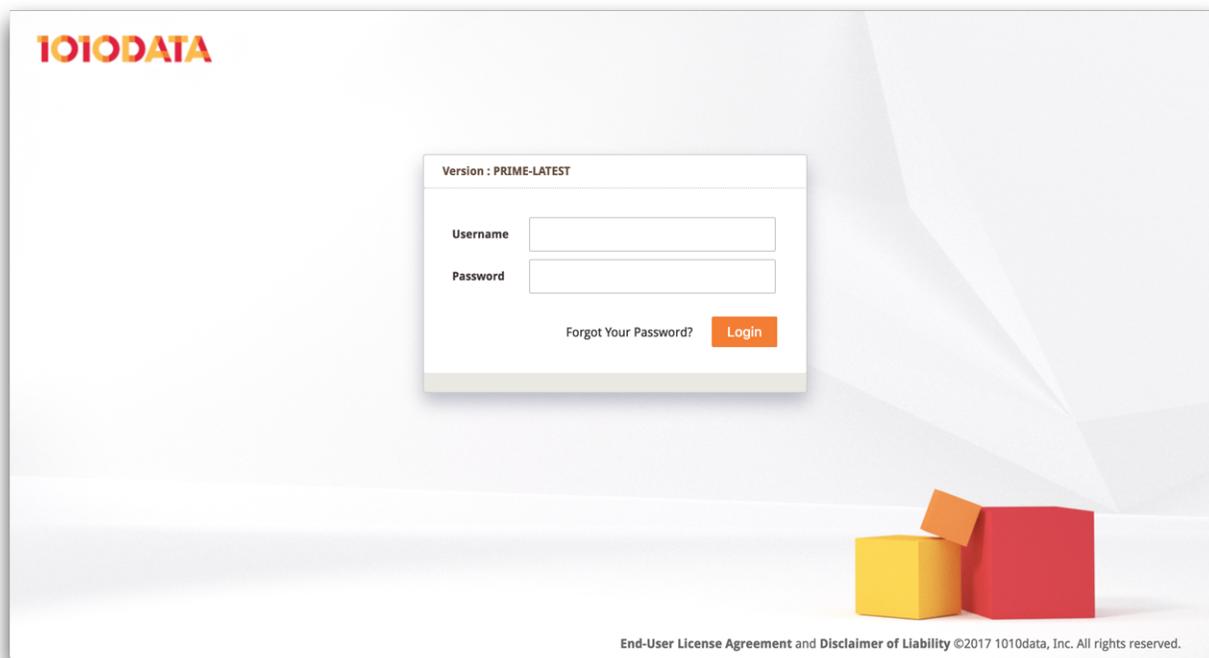
To log in:

1. In your Internet browser address bar, enter <https://www2.1010data.com/prime-latest/> and press **Enter**.

This URL logs you in to the version of the platform configured in your user account. If your user account is not set to use a specific version, the version configured for your organization is used by default. To log in to a different version of the system, you can specify it in the URL. For more information, see [Versions](#) on page 32.

Note: If you cannot log in with the URL above, your organization may have a custom Insights Platform URL. If you are unsure of the URL you should use, contact the 1010data Insights Platform administrator at your company or 1010data Support.

Your browser displays the Insights Platform login page.



2. Complete the following fields:

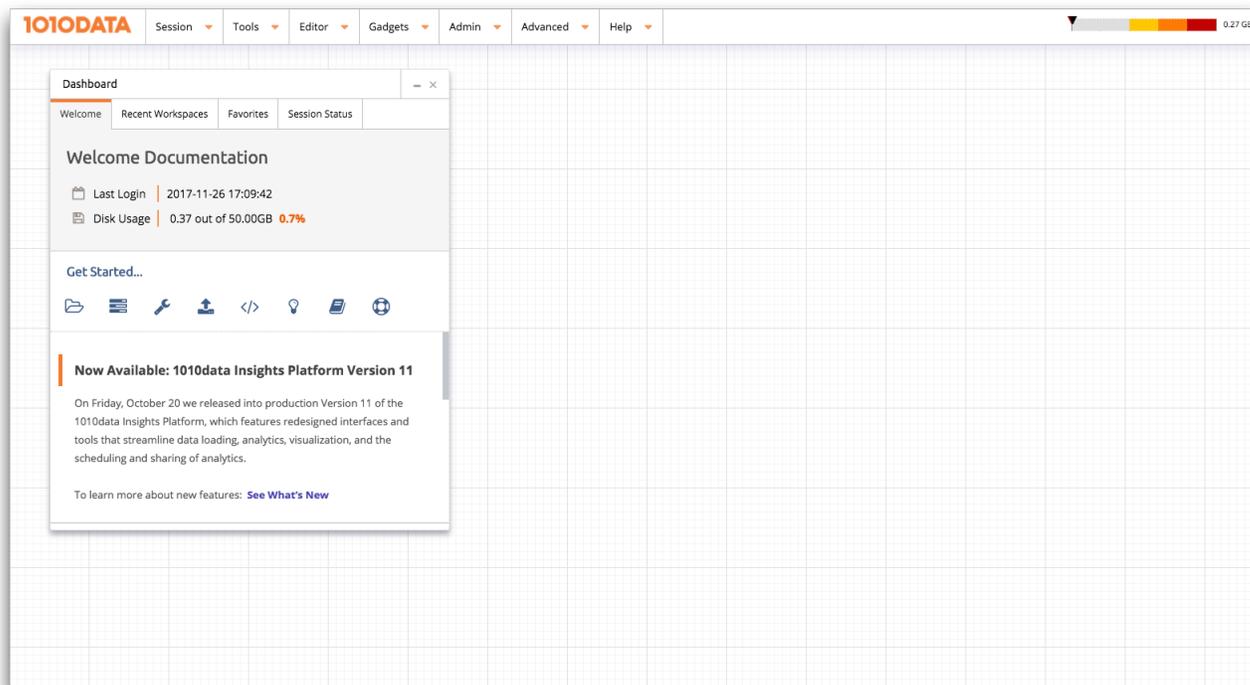
Username Enter your Insights Platform username.
If you do not have a username, contact 1010data Support.

Password Enter your Insights Platform password.
If you cannot remember your password, you can reset it. Click **Forgot your password?** and follow the instructions on the **Password Reset** page. For detailed instructions, see [Reset your password](#) on page 14.

3. Click **Login**.

Note: If you already have an existing Insights Platform session, you are prompted to either re-enter or end the existing session. For more information, see [Sessions](#) on page 32.

Your browser displays the Insights Platform workspace.



Reset your password

If you have forgotten your 1010data Insights Platform password, the login page provides a way for you to reset it.

You can reset your password from the login page.

To reset your password:

1. From the Insights Platform login page, click **Forgot Your Password?**

A screenshot of the 1010data Insights Platform login page. At the top, it says 'Version : PRIME-LATEST'. Below this, there are two input fields: 'Username' and 'Password'. At the bottom of the form, there are two buttons: 'Forgot Your Password?' and 'Login'. The 'Forgot Your Password?' button is highlighted with a red border.

Your browser displays the **Password Reset** dialog on the login page.

Version : PRIME-LATEST

Password Reset

If you have forgotten your 1010data password, simply enter your username and the corresponding email address and we will send you a link to reset your password.

Username:

Email:

Reset My Password

2. Complete the following fields:

Username Enter your Insights Platform username.

If you do not have a username, or have forgotten it, contact 1010data Support. For more information, see [Contact 1010data Support](#) on page 30

Email Enter the email address associated with your Insights Platform account.

3. Click **Reset My Password**.

After validating your information, an email from support@1010data.com is sent to you containing a link to reset your password.

4. In the password reset email, click the password reset link.

Your browser opens a new login page and displays the **Password Reset** dialog where you can enter your new password.

PRIME-LATEST

Password Reset

Please enter a new password.

- New password must be different from current password.
- Password must contain 8 to 16 characters.
- Password must contain at least one letter and one number.
- Password must not contain spaces.
- Password may contain symbols, except for: & % " =
- Password is case sensitive.

Username

New Password

Confirm Password

Change Password

5. Complete the following fields:

Username This field lists your username; do not change the contents of this field.

New Password Enter a new Insights Platform account password.

Password information and requirements:

- Your new password must be different from your current password

- Passwords must contain 8 to 16 characters
- Password must contain at least one letter and one number
- Passwords must not contain spaces
- Passwords may contain symbols, except for & % " =
- Passwords are case sensitive

Confirm Password

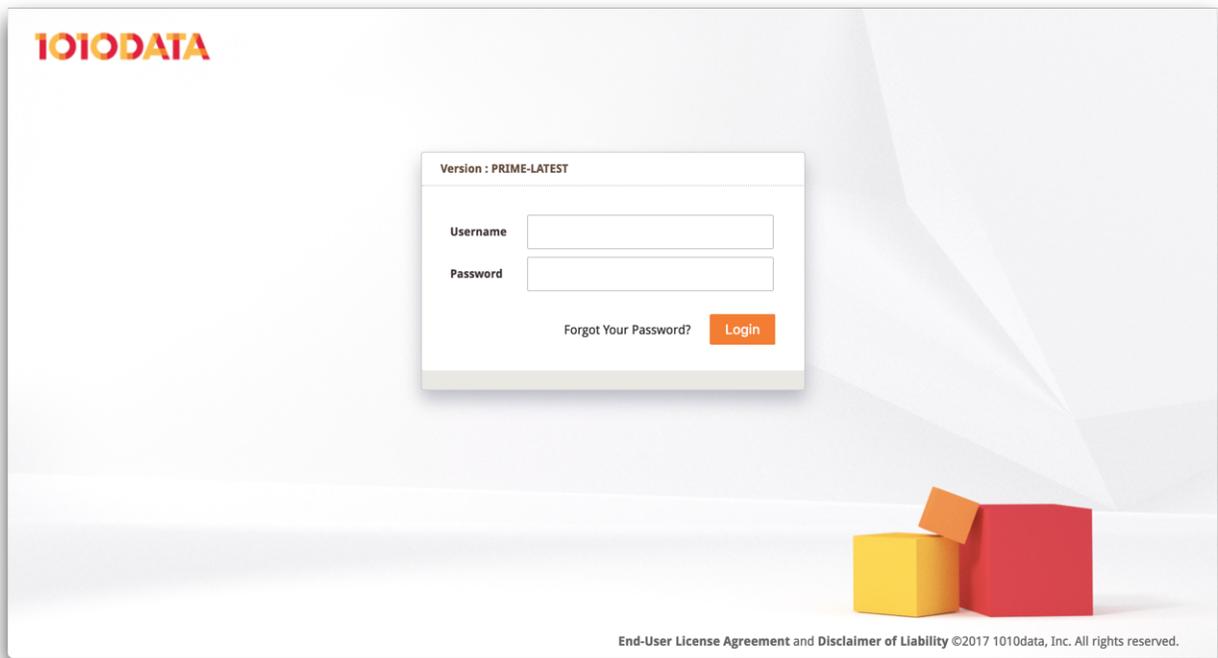
Confirm your new password.

Note: The values in the **New Password** and **Confirm Password** fields must match.

6. Click **Change Password**.

7. Click **Return to Login**.

Your browser displays the Insights Platform login page.



From the login page, use your new password to log in. For instructions, see [Log in](#) on page 13.

Understanding the user interface

Learn how to use some of the most common user interface elements in the 1010data Insights Platform.

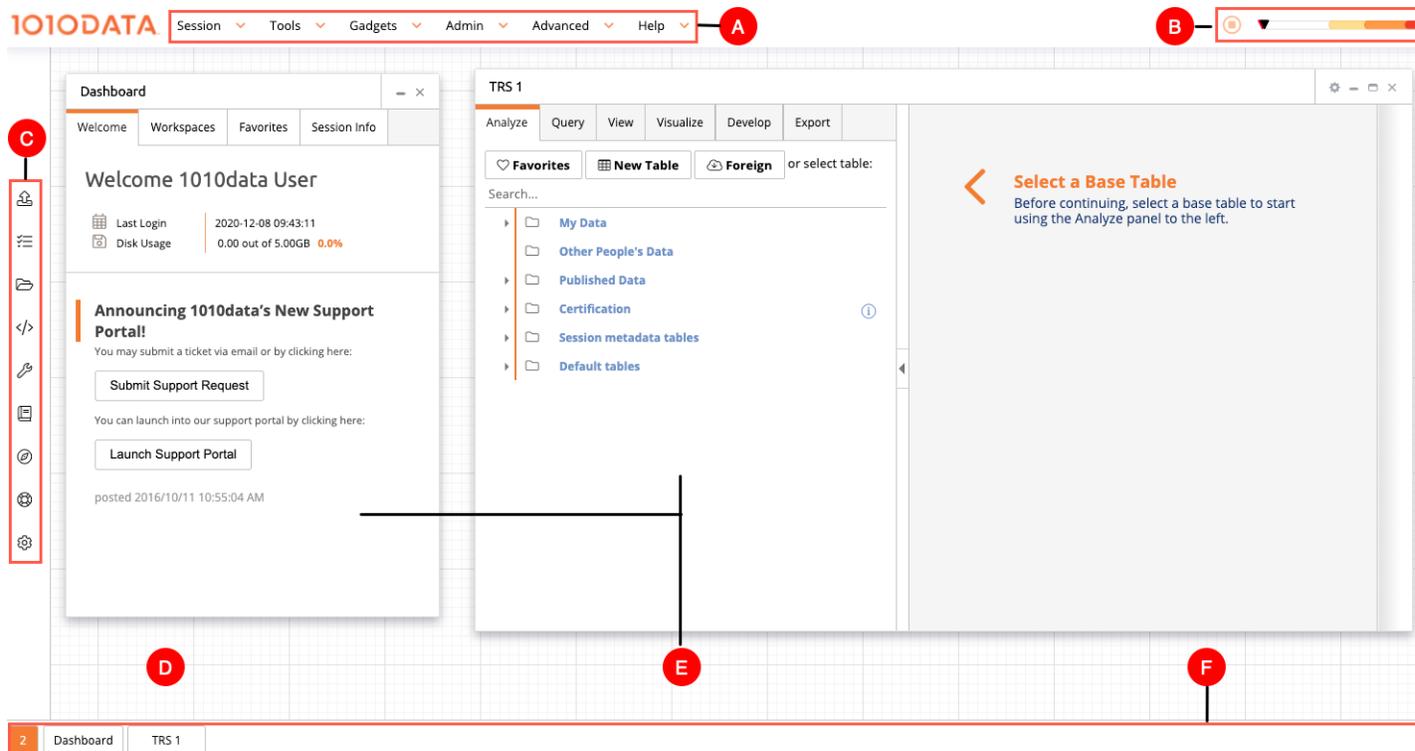


Figure 1: The 1010data Insights Platform user interface

A. Workspace menu

The workspace menu, located at the top left side of the screen, provides access to the features and functionality available to you in the Insights Platform. The menu is customizable to each individual user ID and, as such, you may not see or have access to all of the features and options discussed in this guide.

B. Session memory meter

The session memory meter, located at the top right side of the screen, displays the amount of virtual memory you have used in the session. For more information, see [Virtual memory](#) on page 46.

C. Launch bar

The launch bar offer easy access to the most frequently-used 1010data Insights Platform tools, such as the Trillion-Row Spreadsheet and the Macro Language Workshop. Hover the mouse over each icon to view the tool name. You can customize whether you want the launch bar to appear at all times or only when you hover the mouse on the far left of the screen. For more information, see [Workspace Settings](#) on page 37.

D. Canvas

Data and tools used in your analysis display in individual windows on the canvas. The canvas expands when you move or resize a window so that a portion of the window extends beyond the right or bottom borders of the canvas. When this happens, the Insights Platform displays the canvas scroll bars. You can customize the appearance of the canvas of each workspace in your session. For more information, see [Workspace Settings](#) on page 37.

E. Windows

Most of the information in the Insights Platform is displayed in windows. You can move, resize, minimize, and close individual windows.

The **Dashboard** displays the **Welcome** tab each time you log in or open a new workspace. For more information, see [Dashboard](#) on page 70.

F. Window tray

By default, the window tray always appears at the bottom of the workspace. The window tray displays a button for every window in your workspace, regardless of whether the window is open or minimized. If you minimize a window, you can restore the window to its original size and location by clicking the window's button in the window tray.

Optionally, you can choose to display the window tray only when there is at least one minimized window in your workspace. Open windows will not appear in the window tray. As with the default option, you can restore the window to its original size and location by clicking the window's button in the window tray. See [Appearance](#) on page 38 for window tray options.

Workspace menu

The workspace menu provides access to the features and functionality available to you in the 1010data Insights Platform.

Located at the top left side of the workspace, the menu is customizable to each individual user ID and, as such, you may not see or have access to all of the features and options discussed in this guide.



Figure 2: Workspace menu

Session

From this menu you can perform tasks related to your current Insights Platform session. For example, you can view the amount of virtual memory used, customize the user interface, and save, load, or open a new workspace.

Tools

This menu provides access to the various tools that are available for analyzing your data. For example, you can access the Trillion-Row Spreadsheet to perform ad hoc analysis, upload data with the Quick or Custom Uploader tools, and develop applications using the Macro Language Workshop.

Gadgets

Gadgets are simplified tools that support your data analysis session. For more information, see [Gadgets](#) on page 432.

Admin

You can manage your Insights Platform user account from this menu. For example, you can change your password and manage other profile settings. Depending on your role and access rights, you may have additional options. For example, you may manage users and groups.

Help

The help menu provides access to a variety of resources to assist you while using the Insights Platform. For more information, see [Help menu](#) on page 28.

Window tray

Depending on your chosen workspace settings, the window tray either holds all windows open in your session or your minimized windows only.

Located at the bottom of your workspace, the window tray provides a way to organize all of your windows. You can choose how the window tray behaves.

By default, the window tray is always displayed. The window tray displays a button for each open and minimized window in your workspace. For reference, the counter indicates the total number of open and minimized windows.

You can change the default setting so that only minimized windows appear in the window tray. For reference, the counter indicates the number of minimized windows.

For information about changing the behavior of the window tray, see [Appearance](#) on page 38.

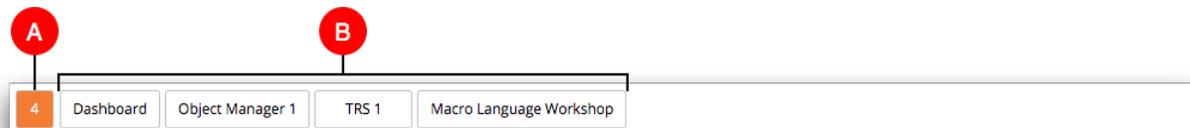


Figure 3: Window tray

A. Window tray counter

By default, the counter indicates the number of all open and minimized windows in the workspace.

If you choose to display minimized windows only, the counter indicates the number of minimized windows in the workspace.

B. Window buttons

By default, the window tray displays a button for each open and minimized window in your session. Buttons in the window tray are listed in the order that windows are opened. Click a button in the window tray to make the selected window active. If a window is minimized, clicking the window button restores the window to the size and location it was in before it was minimized.

If you choose to display minimized windows only in the window tray, when you minimize a window, the Insights Platform moves it to the window tray and displays it as a button. To learn how to minimize a window, see [Window controls](#) on page 21. In this situation, buttons in the window tray are listed in the order that windows are minimized. Click a button in the window tray to restore the window to the size and location it was in before it was minimized.

Windows

Data and tools used in your analysis display in individual windows on the canvas.

You can move, resize, minimize, and close individual windows.

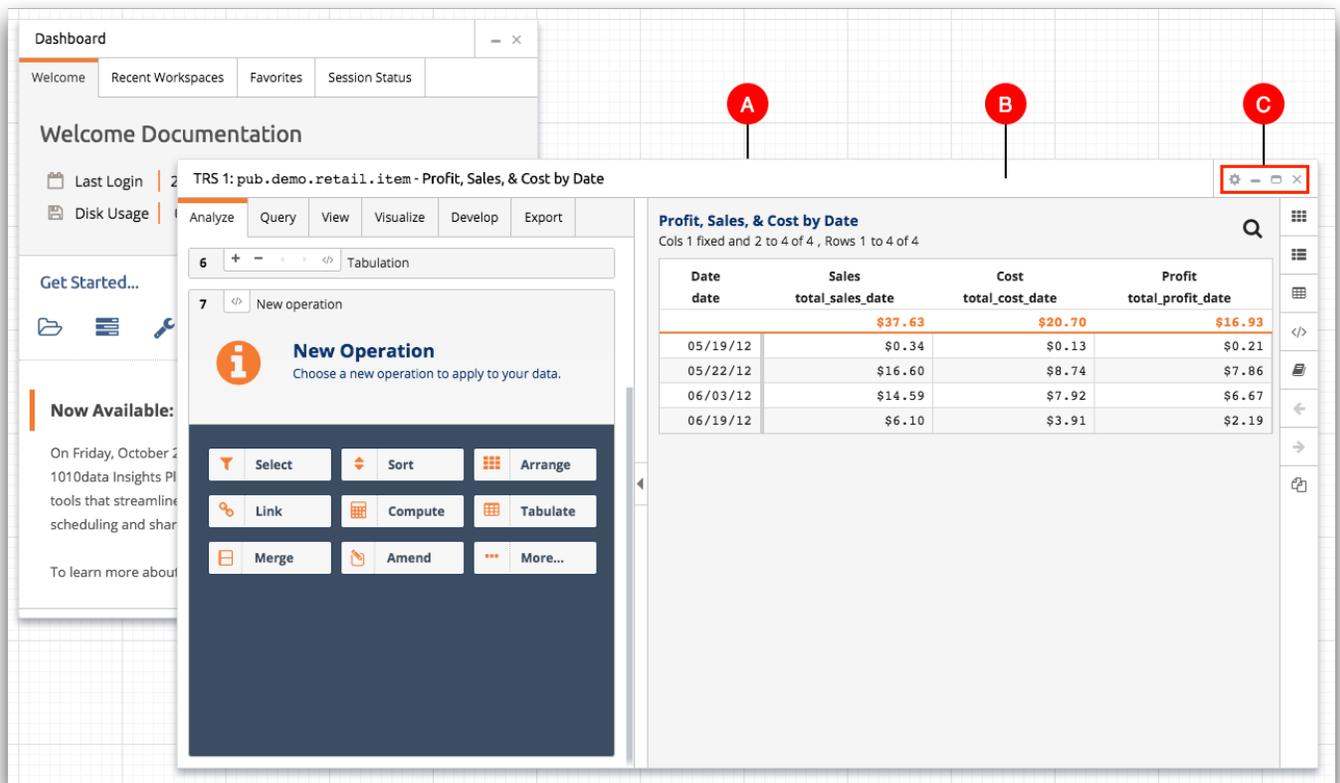


Figure 4: Open windows on the canvas

A. Active window

When you have multiple windows open on the canvas, the active window is displayed on top of the other, inactive, windows.

B. Title bar

The title bar displays the name of the window. Some windows hide the title bar when the window is not active. To display a hidden title bar, either place your pointer over the window or make the window active.

C. Window controls

Each window has a set of icons in the upper right corner that allow you to control certain aspects of the window size and placement. In addition, the title of some windows can be customized.

Move or resize a window

As needed, you can move or resize a window. The canvas expands when you move or resize a window so that a portion of the window extends beyond the right or bottom borders of the canvas.

Move

Click the title bar in the window and drag the window to the location where you want to place it.

You can also use the arrow keys on your keyboard to move an active window.

Resize

Drag the border or corner of a window to resize it.

Note: Not all windows are resizable.

Switch between windows

You can switch between open windows to make a window active.

- Click a window on the canvas.
- Click a window button in the window tray.
- Press **Ctrl**+` (PC) or **Control**+` (Mac) to cycle through open windows in the window manager.

Note: On most U.S. keyboards, the backtick key (`) is located directly to the number one (1) key.

After the window manager appears, continue holding down the **Ctrl** key (PC) or **Control** key (Mac) to keep the window manager on the screen.

Each time the backtick key is pressed, the next window in the window manager is highlighted. When the window you want to make active is highlighted, release the **Ctrl** or **Control** key.

Window controls

Icons on the right side of the window title allow you to control certain aspects of the window size, placement, and appearance.

Similar to other software applications that use a windowed layout, the 1010data Insights Platform allows you to maximize, minimize, restore, and close individual windows. In addition, you can edit the title of certain windows customizing them to your needs.

The following window control icons are available for your use.

Note: The controls described below may not be available in all windows.

Table 2: Window control icons

Icon	Name	Description
	Edit	Click the Edit () icon to change the window title. For instructions, see Rename a window on page 21.
	Minimize	Click the Minimize () icon to remove the window from the canvas. Minimizing a window does not close it or delete its contents—it temporarily removes it from the canvas. Click the minimized window button in the window tray to place it back on the canvas. The window will be restored to the size and location it was before it was minimized. For more information, see Window tray on page 18.
	Maximize	Click the Maximize () icon to enlarge the window so that it fills the available canvas space as displayed on your screen.
	Restore	When a window is maximized, click the Restore () icon to return the window to the size and location it was in before it was maximized.
	Close	Click the Close () icon to remove a window from the canvas.

Rename a window

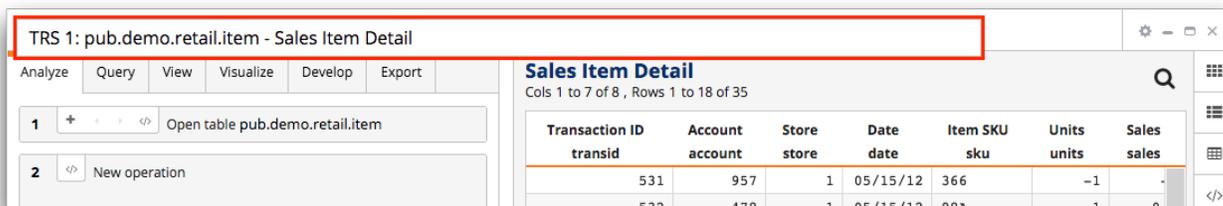
You can edit the title of certain windows customizing them to your needs.

Certain 1010data Insights Platform windows, such as the Trillion-Row Spreadsheet and Macro Language Workshop, can be renamed. You can change the title of a window if it has the **Edit** (⚙️) icon in the window controls.

Note: In the steps below, a Trillion-Row Spreadsheet window is used for illustrative purposes.

To rename a window:

1. Click the **Edit** (⚙️) icon in the window you want to rename. The Insights Platform displays the window title in an editable field.

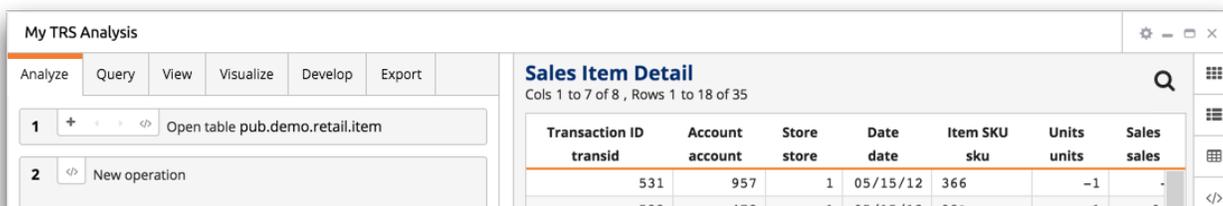


2. Edit the text in the title field and then do one of the following:

- Click outside of the title field.
- Press **Enter** (PC).
- Press **Return** (Mac).

Note: To discard the changes to the window title, press **Esc**.

The Insights Platform closes the field and saves the window title.



Keyboard shortcuts

Keyboard shortcuts help to streamline your experience by giving you quick access to some of the most common actions in the 1010data Insights Platform.

You can use keyboard shortcuts in many areas of the Insights Platform. The tables below group shortcuts by the location of the system in which they are available.

Note: Shortcuts may differ based on the operating system of your computer. Where applicable, differences between shortcuts for PC and Mac are indicated.

General

The following keyboard shortcuts are available no matter where you are in the web interface.

Table 3: Window manager shortcuts

Action	Keyboard shortcut	Description
Switch active window	<ul style="list-style-type: none"> • Ctrl+ (PC) • Control+ (Mac) 	Displays the window manager which allows you to cycle through open windows.

Action	Keyboard shortcut	Description
		<p>Note: On most U.S. keyboards, the backtick (`) key is located directly to the left of the number one (1) key.</p> <p>After the window manager appears, continue holding down the Ctrl key (PC) or Control key (Mac) to keep the window manager on the screen.</p> <p>Each time the backtick key is pressed, the next window in the window manager is highlighted. When the window you want to make active is highlighted, release the Ctrl or Control key.</p>

Macro Language Workshop

The following keyboard shortcuts are available within the code editor pane in the **MLW** window.

Table 4: Code editor pane shortcuts

The cursor must be within the code editor pane to use these keyboard shortcuts.

Action	Keyboard shortcut	Description
Submit code	<ul style="list-style-type: none"> • Ctrl+Enter (PC) • Command+Return (Mac) • Control+Return (Mac) 	<p>Renders the set of actions that appear in the code editor pane.</p> <p>This keyboard shortcut executes the last render command selected from the Run menu.</p> <p>If a render command has not yet been selected, this keyboard shortcut is the same as selecting Run > Render from the menu.</p>
Clone the MLW	<ul style="list-style-type: none"> • Ctrl+Shift+Enter (PC) • Shift+Command+Return (Mac) 	<p>Creates a new MLW window with only the selected code from the current code editor. If the code is a valid query, it is run immediately.</p> <p>This keyboard shortcut is the same as selecting File > Export > and run selection in MLW from the Run menu.</p> <p>Note: Pressing Ctrl+Shift+Enter (Return) or Shift+Enter (Return) with no text selected will open a blank MLW window.</p> <p><i>(Available as of version 11.27)</i></p>
Code help	<ul style="list-style-type: none"> • Ctrl+Shift+? (PC) • Shift+Command+? (Mac) 	<p>Opens the reference pane and displays help content about the Macro Language code associated with the location of the cursor in the code editor.</p> <p><i>(Available as of version 11.27)</i></p>
Find	<ul style="list-style-type: none"> • Ctrl+F (PC) • Command+F (Mac) 	<p>Displays the search options to locate specified search text in the code editor pane.</p>

Action	Keyboard shortcut	Description
Find next	<ul style="list-style-type: none"> • Ctrl+G (PC) • Command+G (Mac) 	Locates the next instance of search text in the code editor pane.
Replace	<ul style="list-style-type: none"> • Ctrl+Shift+F (PC) • Option+Command+F (Mac) 	Replaces each occurrence of specified text in the code editor pane with replacement text. When using this option, you must confirm the replacement of each occurrence.
Replace all	<ul style="list-style-type: none"> • Ctrl+Shift+R (PC) • Option+Shift+Command+F (Mac) 	Replaces all occurrences of specified text in the code editor pane with replacement text. When using this option, no confirmation is required before the replacement of all occurrences.

Trillion-Row Spreadsheet

The following keyboard shortcuts are available in the **TRS** window.

Table 5: New operation shortcuts

The **New operation** panel in the Analysis Timeline must be open and selected to use these keyboard shortcuts.

Action	Keyboard shortcut	Description
Add computed column	Equal sign (=) (Mac and PC)	This shortcut creates a new computed column in the Analysis Timeline. It is the equivalent of the Compute icon in the New operation panel.
Add selection	Question mark (?) (Mac and PC)	This shortcut creates a new selection in the Analysis Timeline. It is the equivalent of the Select icon in the New operation panel.
Add code	Less than sign (<) (Mac and PC)	This shortcut creates a new code panel in the Analysis Timeline. It is the equivalent of the Code icon in the New operation panel.

Table 6: TRS timeline shortcuts

The **Analyze** view of TRS must be in focus to use these keyboard shortcuts.

Action	Keyboard shortcut	Description
Undo	<ul style="list-style-type: none"> • Ctrl+Z (PC) • Command+Z (Mac) 	Undoes the last step in the TRS analysis timeline. You can undo every step of the timeline back to the beginning.
Redo	<ul style="list-style-type: none"> • Ctrl+Shift+Z (PC) • Command+Shift+Z (Mac) 	Redoes any steps in the timeline that were previously undone. You can continue using this shortcut until all undone operations are restored.

Table 7: Computed column panel shortcuts

The **Computed column** panel in the Analysis Timeline must be open and selected to use these keyboard shortcuts.

Action	Keyboard shortcut	Description
Submit operation	<ul style="list-style-type: none"> • Enter (PC) • Return (Mac) 	<p>Provided that the entered expression is complete and valid, pressing Enter (PC) or Return (Mac) submits the operation.</p> <p>Using this shortcut performs the same action as clicking the Submit operation (✓) icon.</p>
Submit and run query to the end	<ul style="list-style-type: none"> • Ctrl+Enter (PC) • Control+Return (Mac) 	<p>Provided that the entered expression is complete and valid, pressing Ctrl+Enter (PC) or Control+Return (Mac) submits the operation and runs the query to the end of the timeline.</p> <p>Using this shortcut performs the same action as clicking the Submit and run query to end (▶) icon.</p>

Table 8: Panel Macro Language view shortcuts

The Macro Language view of a panel in the Analysis Timeline must be displayed and the cursor must be within the code editor field to use these keyboard shortcuts.

Action	Keyboard shortcut	Description
Find	<ul style="list-style-type: none"> • Ctrl+F (PC) • Command+F (Mac) 	Displays the search options to locate specified search text in the Macro Language view code editor field.
Find next	<ul style="list-style-type: none"> • Ctrl+G (PC) • Command+G (Mac) 	Locates the next instance of search text in the Macro Language view code editor field.
Replace	<ul style="list-style-type: none"> • Ctrl+Shift+F (PC) • Option+Command+F (Mac) 	<p>Replaces each occurrence of specified text in the Macro Language view code editor field with replacement text.</p> <p>When using this option, you must confirm the replacement of each occurrence.</p>
Replace all	<ul style="list-style-type: none"> • Ctrl+Shift+R (PC) • Option+Shift+Command+F (Mac) 	<p>Replaces all occurrences of specified text in the Macro Language view code editor field with replacement text.</p> <p>When using this option, no confirmation is required before the replacement of all occurrences.</p>
Undo	<ul style="list-style-type: none"> • Ctrl+Z (PC) • Command+Z (Mac) 	Undoes the last text typed in the Macro Language view code editor.
Redo	<ul style="list-style-type: none"> • Ctrl+Shift+Z (PC) • Command+Shift+Z (Mac) 	Redoes any text undone by the undo command.

Table 9: Select rows panel shortcuts

The **Select rows** panel in the Analysis Timeline must be open and selected to use these keyboard shortcuts.

Action	Keyboard shortcut	Description
Enter a value expression	Equal sign (=)	Displays the expression editor. When the Simple Comparisons option in the Select rows panel is selected, press the equal sign (=) key to displays the expression editor. Use the expression editor to enter a value expression.
Select to the end	Ctrl+Shift+Right arrow (PC and Mac)	In the expression editor, this shortcut selects all text from the cursor to the end of the expression.
Select to the beginning	<ul style="list-style-type: none"> • Ctrl+Shift+Left arrow (PC and Mac) 	In the expression editor, this shortcut selects all text from the cursor to the beginning of the expression.
Submit operation	<ul style="list-style-type: none"> • Enter (PC) • Return (Mac) 	Provided that the expression entered in the expression editor is complete and valid, pressing Enter (PC) or Return (Mac) submits the operation. Using the shortcut performs the same action as clicking the Submit operation (✓) icon.
Submit and run query to the end	<ul style="list-style-type: none"> • Ctrl+Enter (PC) • Control+Return (Mac) 	Provided that the expression entered in the expression editor is complete and valid, pressing Ctrl+Enter (PC) or Control+Return (Mac) submits the operation and runs the query to the end of the timeline. Using this shortcut performs the same action as clicking the Submit and run query to end (▶) icon.

Keyboard navigation

Keyboard navigation allows you to perform certain actions without needing to use the mouse.

You can use keyboard navigation in many areas of the 1010data Insights Platform. The tables below group keyboard navigation by the location of the system in which they are available.

This is not an exhaustive list of all keyboard navigation available in the Insights Platform. Many standard computer operating system and web browser keyboard navigation actions may be used throughout the platform.

Note: Keyboard navigation may differ based on the operating system of your computer. Where applicable, differences between keyboard navigation for PC and Mac are indicated.

General

The following keyboard navigation actions are available in multiple areas of the web interface.

Table 10: Object browser keyboard navigation

The object browser appears throughout the web interface; each instance uses the same keyboard navigation actions.

Action	Keyboard navigation	Description
Move up the tree	Up Arrow	Moves the selector up through the object browser tree. Objects in expanded folders are included as the selector moves through the tree.
Move down the tree	Down Arrow	Moves the selector down through the object browser tree. Objects in expanded folders are included as the selector moves through the tree.
Expand a folder	Right Arrow	Expands a collapsed folder. When the selector is on an expanded folder, the Right Arrow key moves the selector to the first object within the folder.
Collapse a folder	Left Arrow	Collapses an expanded folder. When the selector is on a collapsed folder, or on a the first object within a folder, the Left Arrow key returns the selector to the previous folder level, if any.
Open an object	<ul style="list-style-type: none"> • Enter (PC) • Return (Mac) 	Opens an object in the window defined for the object.

Trillion-Row Spreadsheet

The following keyboard shortcuts are available in the **TRS** window.

Table 11: Expression editor keyboard navigation

The expression editor is available in both the **Select rows** and **Computed column** panels in the Analysis Timeline.

Action	Keyboard navigation	Description
Cycle between menu options	<ul style="list-style-type: none"> • Up Arrow • Down Arrow 	Cycles between options in the autocomplete menu.
Select menu option	Tab	Inserts the selected item in the autocomplete menu into the expression. Note: If nothing is selected, the Tab key inserts the first option available, if any.
Close menu	Esc	Closes the autocomplete menu. Note: The autocomplete menu will reappear when you begin entering text in the expression editor field.

Table 12: Tabulation Function menu keyboard navigation

The **Tabulation Function** menu appears after clicking a tabulation function in the **Tabulation** panel.

Action	Keyboard navigation	Description
Move up the menu	Up Arrow	Moves the selector up through the menu.

Action	Keyboard navigation	Description
Move down the menu	Down Arrow	Moves the selector down through the menu.
Select menu item	<ul style="list-style-type: none"> • Enter (PC) • Return (Mac) 	Confirms the menu selection.
Previous level	Esc	Returns the selector to the previous menu level, if any.

Obtaining technical assistance

1010data provides a variety of resources to help you have the best experience possible as you use our system.

You can find help via the **Help** menu, submit a support request, or contact 1010data Support for technical assistance.

Help menu

Help is available directly from within the user interface.

The following items are located under **Help** in the workspace menu:

Support

Support gives you two options: You can either launch the 1010data Support Portal or open the **Help & Support** dialog.

The 1010data Support Portal offers an area to manage your support tickets and create a new support request. The Support Portal also offers a knowledge base of useful support information.

The **Help & Support** dialog allows you to request help, report bugs, suggest features, and communicate with 1010data Support without logging in to the Support Portal. For more information, see [Submit a support request](#) on page 29.

Documentation

Documentation gives you several options for accessing documentation.

- **Getting Started Guide** opens the *Trillion-Row Spreadsheet Getting Started Guide*.
- **User's Guide** launches the most recently released edition of the *1010data Insights Platform User's Guide*.
- **Reference Manual** launches the most recently released edition of the *1010data Reference Manual*.
- **Documentation Center** launches the **1010data Documentation Center**. This webpage contains the online repository of 1010data documentation and training materials.

Training

Launches your personal 1010data Lessonly webpage, where you can complete 1010data training.

Changelog

Launches the **Changelog** webpage for the version of the system you are using.

This webpage catalogs all of the fixes and enhancements contained within each release, listed reverse chronologically by release date.

Technical Interfaces Download Page

Launches the 1010data **Technical Interfaces** webpage.

This webpage provides download links to 1010data APIs, SDKs, drivers, and tools.

Submit a support request

Send a support request to 1010data using the **Help & Support** window.

You can submit a support request directly from within the 1010data Insights Platform.

To submit a support request:

1. From the workspace menu, click **Help > Support > Contact Support**.

The Insights Platform displays the **Help & Support** window.

2. Complete the following fields:

Requester This field lists the email address associated with your Insights Platform account. If the address is missing or incorrect, update your email address in the **User Profile** window. For instructions, see [Edit your user profile](#) on page 83.

Your Name Enter your first and last name. This field appears if you enter or edit your email address in the **Requester** field.

Subject Enter the subject of your support request.

3. In the message field, enter a detailed explanation of the question, problem, or request.

We recommend including the following information:

- Copy and paste the full error message as text
- Describe the steps taken prior to receiving an error
- Include a link with the linking tool
- Where possible, include the Macro Language code of your query

Note: Do NOT include your Insights Platform password in your support request.

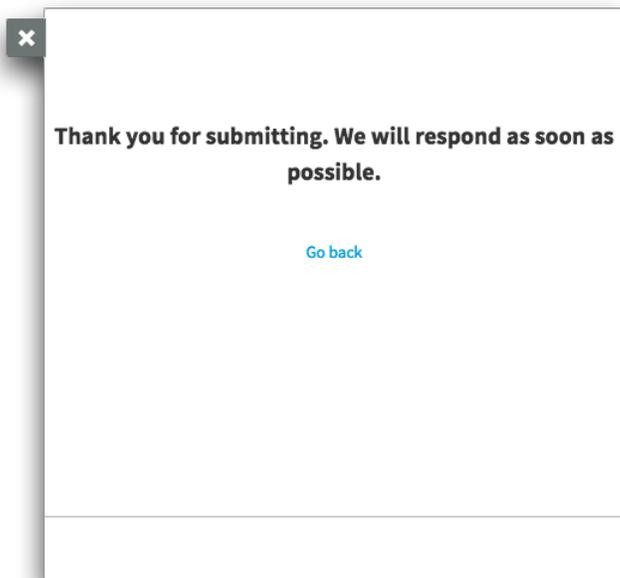
4. From the **Priority** drop-down list, select one of the following options:

- **Low**
- **Medium**
- **High**
- **Urgent**

The rest of the information is for 1010data Support and cannot be edited.

5. Click **Submit**.

The Insights Platform submits your request and displays a confirmation screen.



In addition, you will receive an email message from support@1010data.com with your support ticket number.

6. Click the **Close** () icon to return to the workspace.

Contact 1010data Support

You can contact 1010data Support directly either through email or by phone.

As a starting point for all questions, refer to this *User's Guide* and other Insights Platform documentation first. Questions can usually be resolved by reading and following the instructions in these documents.

If after reviewing this guide you still have questions, or if you encounter any system problems as you use the platform, contact 1010data Support.

Email support@1010data.com

Phone 1-866-405-DATA (3282)

In addition, the 1010data Support Portal contains a knowledge base of information.

To access the 1010data Support Portal, do one of the following:

- Log in to the 1010data Support Portal directly at <https://support.1010data.com>.
- From the workspace menu in the 1010data Insights Platform, click **Help** > **Support** > **Launch Support Portal**.

Log out

Log out to close your active session.

You should log out each time you finish using the 1010data Insights Platform. Logging out ensures your identity cannot be used without your permission after you leave your computer.

Note: Simply closing your browser window does not log you out of the platform.

When you log out of your Insights Platform session, all open workspaces are closed and the session is ended. The next time you log in, you can restore your workspace from the **Recent Workspaces** tab in the **Dashboard**.

To log out:

1. From the workspace menu, click **Session > Log out**.
The Insights Platform displays a confirmation message.
2. Click **OK**.

The Insights Platform logs you out, closes your active session, and displays the login page.

Note: If a session is also in use with the Insights Platform Excel Add-in, the session in the Excel Add-in will also end when you log out.

User basics

Learn about the key concepts and terminology that provide the foundation for the 1010data Insights Platform.

When you first start using the Insights Platform, you should familiarize yourself with the basics that make up the system.

Versions

The 1010data Insights Platform is available in both a stable production version of the system (Prime) and a version that provides access to the latest features (Beta).

The Insights Platform follows an iterative release cycle along three tracks: Prod, Prime, and Beta. Prime changes infrequently, one or two major updates a year, and focuses on stability. Beta changes frequently, roughly once a week, and focuses on providing the latest features of the system for testing and experimentation. The Prod track is no longer under active development, but when necessary may receive important updates such as security patches or bug fixes.

To access the newest version of each track, use the following URLs:

Prod	<code>https://www2.1010data.com/prod-latest</code>
Prime	<code>https://www2.1010data.com/prime-latest</code>
Beta	<code>https://www2.1010data.com/beta-latest</code>

For each release track, it is possible to access older versions. Each version of the system is assigned a version number. To access a specific version in each release track, use the following URL format:

Prod	<code>https://www2.1010data.com/prod-[VERSION_NUMBER]</code>
Prime	<code>https://www2.1010data.com/prime-[VERSION_NUMBER]</code>
Beta	<code>https://www2.1010data.com/beta-[VERSION_NUMBER]</code>

For example, to access Prime version 10.48, enter the following URL:

```
https://www2.1010data.com/prime-10.48
```

It should be noted that versions reflect a snapshot of features, functions, and other system processes. However, versions do not affect the folders, tables, and queries a user can access. These items persist across all versions of the system, but are specific to each Insights Platform URL.

Sessions

A session in the 1010data Insights Platform is an established period of communication between a user and the system. A single user ID can only have one session active at a time.

In the Insights Platform, a session refers to the period of time that a user ID is connected to the system. The time period begins when you log in and ends when you either log out or are logged out due to inactivity.

Sessions are an important part of the Insights Platform. One of the most important benefits is that sessions permit the caching of operations. Whenever an operation is performed in the Insights Platform, the operation and its results are stored in the session cache. If the operation needs to be run again at a later time, the results are already calculated and readily accessible. This design feature allows for a compositional, orderly approach to data analysis.

A session cache has physical limitations. Each user account has a maximum amount of memory available for both cached information and current processes.

If you attempt to connect to the Insights Platform when you already have an established session, you are given the option of either re-entering the existing session, or ending the existing session and creating a new one.

The screenshot shows a login form with the following elements:

- Version: PRIME-LATEST
- Warning: User is already logged in. Please choose to:
 - [Re-enter existing session](#)
 - [End existing session](#)
- Username: user_name
- Password: [masked]
- Forgot Your Password? [Login](#)

Table 13: Session options

Option	Description
Re-enter existing session	This option logs you in to the active session and gives you access to the operations and data in the cache.
End existing session	This option ends the existing session and starts a new session with a clear cache.

Note: You cannot re-enter a session that was started through the Insights Platform API.

Workspace

The workspace is an individual instance of the 1010data Insights Platform user interface.

It is important to understand that the term *workspace* is used differently in the new web interface version of the Insights Platform than in the legacy interface. In the legacy interface, workspace refers to the amount of virtual memory allocated to your session. In the new web interface, the workspace is the digital work area of the user interface where you interact with your data.

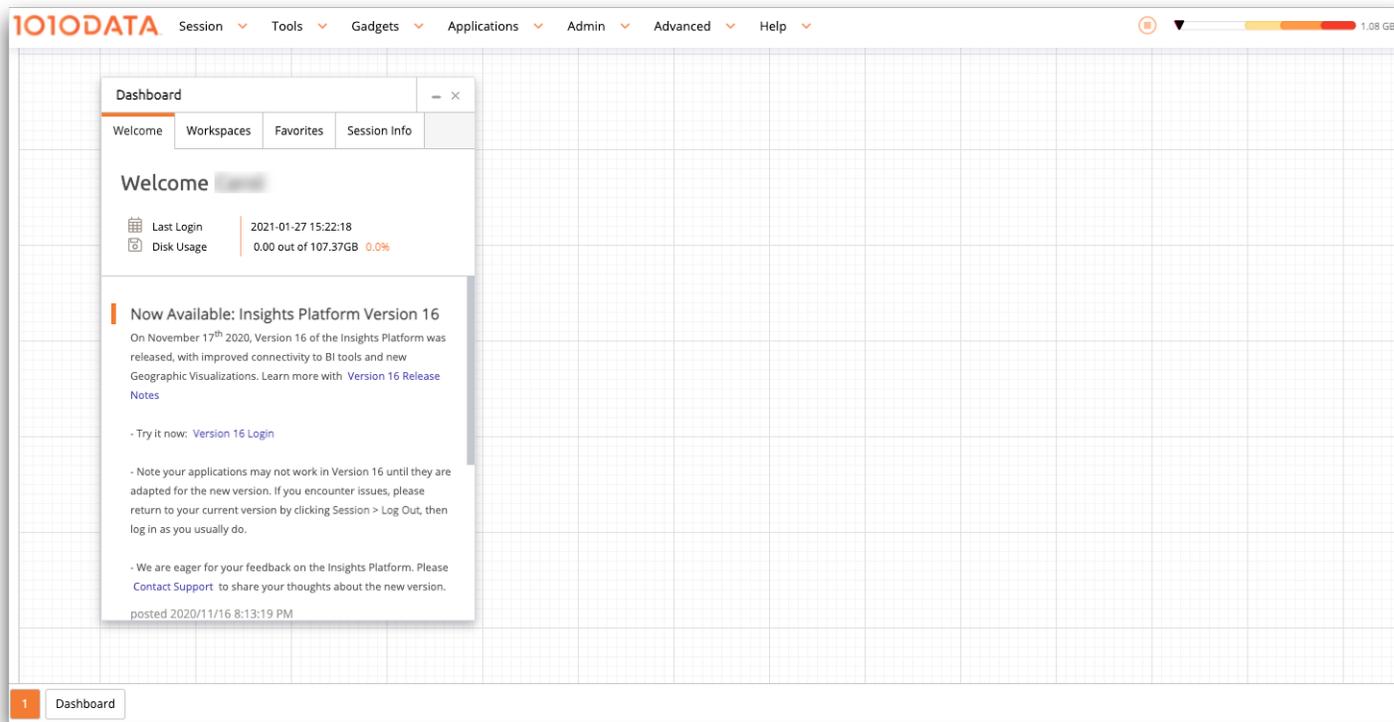


Figure 5: Workspace of the Insights Platform

You can have multiple workspaces open within a single session. Each workspace is located in a separate browser tab which allows you to perform multiple, but separate, analyses at the same time. However, because workspaces are contained within a single session, they share your single allocated amount of session memory.

Your workspace is automatically saved each time the Insights Platform interacts with the server. Automatically saved workspaces are called recent workspaces and can be opened from the **Recent workspaces** window. For more information, see [Recent Workspaces](#) on page 71. Recent workspaces are temporary and are limited to the last 10 interactions with the server. However, you can save a permanent copy of your workspace and return to it later or share it with others.

When you log out of your session, all open workspaces are closed and the session is ended.

Save your workspace

Save your workspace so you can return to it later or share it with others.

When you save your workspace, the 1010data Insights Platform creates an exact copy that you can return to and continue working in later. Workspaces are saved as QuickApps and can be opened from the Object Manager.

To save your workspace:

1. Select **Session > Save Workspace**.

The Insights Platform displays the **Save Workspace as QuickApp** window.

Save Workspace as QuickApp

Folder: My Data Name:

Replace object if it already exists Add to Favorites after saving

You may choose a destination folder above (a name will be supplied automatically when uploading to the My Data folder). You may also choose an existing object to replace in the Folder browser.

Title of saved workspace

Short description of saved workspace

Save current workspace

2. Click the **Folder** field.
The Insights Platform displays the file browser dialog.

Save Workspace as QuickApp

Folder: My Data Name:

- ▶ My Data
- ▶ Other People's Data
- ▶ Published Data
- ▶ Retail Demo Data
- ▶ Certification

Grid Search... Search

Save current workspace

3. Select the location where you want to save your workspace.

Note: You can only save your workspace in a folder that you own (🔑) or have permission to add to (+).

4. Complete the following fields and options:

Name

Enter a name for your workspace.

The workspace name is the file name of your saved workspace. It must begin with a letter and can only contain numbers, letters, and

underscores. The name cannot contain any spaces or other special characters.

Note: If you save your workspace in the **My Data** folder, you cannot enter anything in the **Name** field. Instead, the Insights Platform uses a system-generated name (e.g., t571707446_yourusername).

Replace object if it already exists	Select this option to replace an existing workspace. To replace a workspace, the new workspace must have the same name and be saved in the same folder as the existing workspace.
Add to Favorites after saving	Select this option to add the saved workspace to your Favorites.
Title of saved workspace	Enter a title for your workspace. The title is used to help describe your workspace (e.g., <i>Sales Analysis Workspace</i>) and is used to identify your workspace in the user interface. The title may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters. Note: If you leave this field blank, the text area next to the workspace icon in the file browser and the Favorites tab will be blank.
Short description of saved workspace	Enter a brief description about your workspace.

5. Click **Save current workspace**.

The Insights Platform saves your workspace.

Use the Object Manager or the **Favorites** tab in the Dashboard (if applicable) to open a saved workspace.

Open a saved workspace

Open a workspace you previously saved.

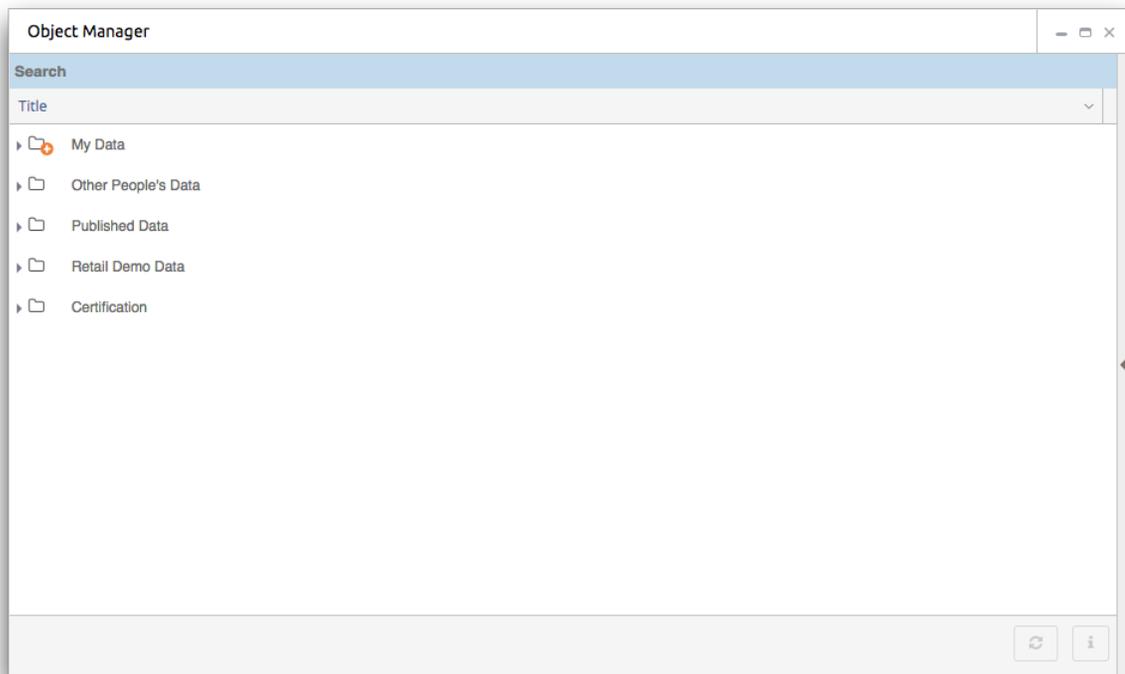
Use the Object Manager to open a previously saved workspace. The saved workspace opens in a new browser tab within your current 1010data Insights Platform session.

Note: TRS windows in the saved workspace display the **Open table** panel in the timeline regardless of the panel visible when the workspace was saved. To view the results of all operations in the timeline, click the **Submit and run query to the end** (▶) icon.

To open a saved workspace:

1. Select **Tools > Object Manager**.

The Insights Platform displays the **Object Manager** window.



2. In the object browser, locate the saved workspace and double-click it.
The Insights Platform opens the saved workspace in a new browser tab.

Add a new workspace

Add a new workspace to your existing session.

You can perform separate analyses concurrently by adding additional workspaces to your 1010data Insights Platform session. When you add a new workspace, the Insights Platform opens it in a separate browser tab. Each additional workspace shares the memory allocated to your session.

To add a new workspace:

Select **Session > New Workspace**.

The Insights Platform adds a new workspace in a separate tab of your browser.

Workspace Settings

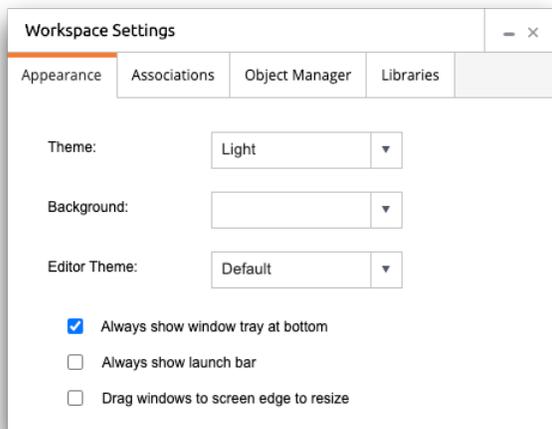
Change the appearance of your workspace and set the default options for objects.

The 1010data Insights Platform allows you to customize the appearance of your workspace. For example, you can change the background image on the canvas or choose a different color theme.

In addition, you can choose how objects behave when opened from the Object Manager. For example, you can choose to have tables open in the Trillion-Row Spreadsheet by default.

The **Workspace Settings** window is accessed by selecting **Session > Workspace Settings** from the workspace menu.

Figure 6: Workspace Settings



The **Workspace Settings** window contains the following tabs:

Appearance

Settings for customizing the appearance of your workspace.

Associations

Settings for defining how objects behave when opened from the Object Manager.

Object Manager

Settings for defining the metadata columns that appear by default in the Object Manager.

Libraries

Settings for specifying the libraries that are automatically loaded into your workspace when you log in.

Appearance

Settings for customizing the appearance of your workspace.

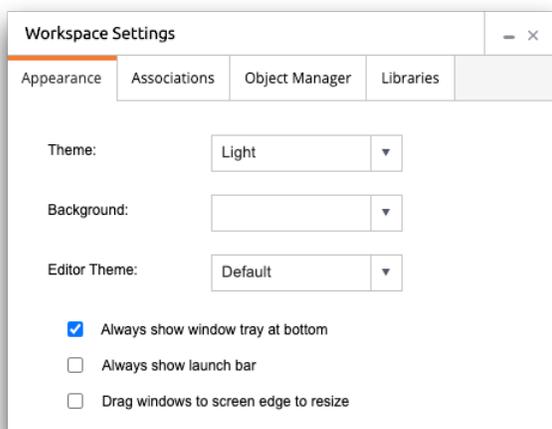


Figure 7: Appearance tab in the Workspace Settings window

The **Appearance** tab contains the following options:

Theme

Color theme options for your workspace. You can choose from light and dark themes.

Light

This is the default theme.

Dark

This theme is useful in low-light situations.

Background

Background image options for your workspace.

Editor Theme

Color theme options for the code editor. Code editors are located in the Macro Language Workshop, in the Trillion-Row Spreadsheet, and in XML editors created by a textbox widget.

When the editor theme is changed, any open instances of the code editor are updated immediately.

Always show window tray at bottom

This option allows you to choose whether or not the window tray is always displayed at the bottom of your workspace. For more information, see [Window tray](#) on page 18.

By default, this option is selected and the window tray always appears at the bottom of the workspace. The window tray displays both open and minimized windows in the workspace.

Clear this option if you want to see minimized windows only, not open windows, in the window tray.

Always show launch bar

Select this option if you want the launch bar on the far left of the workspace to display at all times. The launch bar offers quick access to the most frequently-used 1010data Insights Platform tools. By default, the launch bar displays only when you move the mouse to the far left of the workspace.

Drag windows to screen edge to resize

This option allows you to resize and "snap" a window into place by dragging it to the edge or the corners of the canvas. Additionally, you can maximize a window by dragging it to the top edge and minimize the window by dragging it to the bottom edge or window tray. By default, this option is not selected.

Note: The **Drag windows to screen edge to resize** option only applies to resizable windows.

When this option is selected, the following actions are available:

Drag to corner

Resizes the window to $\frac{1}{4}$ of the canvas area and snaps it into place.

Drag to left- or right-side edge

Resizes the window to $\frac{1}{2}$ of the vertical canvas area and snaps it into place.

Drag to top edge

Maximizes the window.

Drag to bottom edge or window tray

Minimizes the window.

Note: If a window has length or width size limits, the actions above do not resize the window beyond those limits.

Associations

Settings for defining how objects behave when opened from the Object Manager.

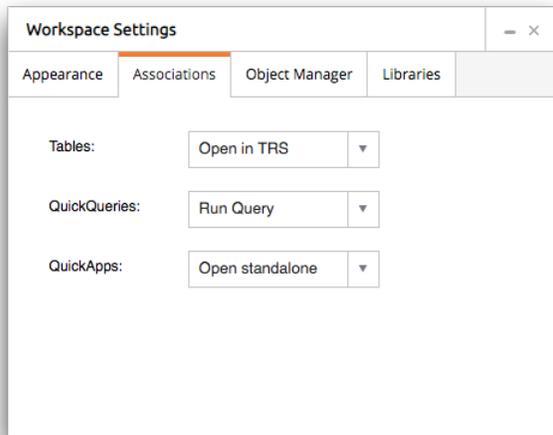


Figure 8: Associations tab in the Workspace Settings window

The **Associations** tab contains the following options:

Tables

This drop-down list defines where a table opens when double-clicked in the Object Manager.

Open in TRS

Double-clicking a table in the Object Manager opens the object in the Trillion-Row Spreadsheet. This is the default setting.

As base in MLW

Double-clicking a table in the Object Manager opens the object in the Macro Language Workshop as the base table.

QuickQueries

This drop-down list defines where a Quick Query/Parameterized Quick Query/legacy PQQ opens when double-clicked in the Object Manager.

Run in TRS

Opens the Quick Query in the view-only mode of the Trillion-Row Spreadsheet. If the query contains parameters, the parameter input fields and options are displayed. Otherwise, the results of the query are displayed in the grid.

The **Run in TRS** option is the default setting.

Run in Query Runner

Opens the query in the **Query Runner** window.

Open in TRS

Opens the query in the Trillion-Row Spreadsheet as the base table and displays the results of the query in the grid.

A parameterized timeline will have all steps loaded into TRS when opened. A non-parameterized timeline will have its individual operations hidden initially. To view the individual operations of the query, click **Import Hidden Actions** below the object browser in the **Open table** panel.

Edit in MLW

Double-clicking a Quick Query in the Object Manager opens the object in the Macro Language Workshop. The Macro Language Workshop displays the individual operations of the query as Macro Language XML code.

QuickApps

This drop-down list defines where a table opens when double-clicked in the Object Manager.

Open standalone

Double-clicking a QuickApp in the Object Manager opens the object as a standalone QuickApp in a new browser tab. This is the default setting.

Open in window

Double-clicking a QuickApp in the Object Manager opens the object in a window within the current workspace.

Open in MLW

Double-clicking a QuickApp in the Object Manager opens the object in the Macro Language Workshop. The Macro Language Workshop displays the individual operations of the query as Macro Language XML code.

Object Manager

Settings for defining the metadata columns that appear by default in the Object Manager.

Although you are able to choose the displayed metadata columns directly from within the Object Manager, those selections apply only to that single instance of the **Object Manager** window. Conversely, this tab is used to set select the metadata columns you want to see each and every time a new **Object Manager** window is opened.

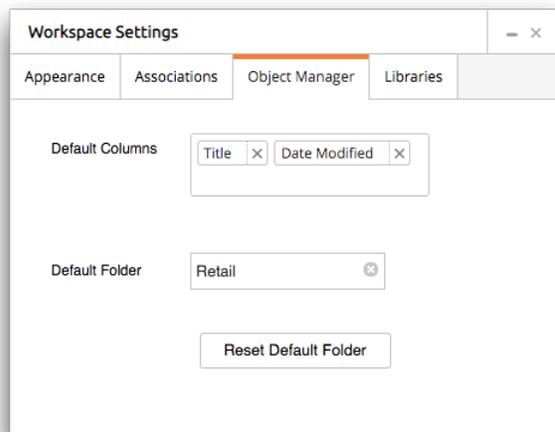


Figure 9: Object Manager tab in the Workspace Settings window

The **Object Manager** tab contains the following:

Default Columns

This field is used to select the metadata columns that appear by default in a new **Object Manager** window. For more information, see [Metadata columns](#) on page 91. Currently selected metadata columns are listed within the field.

You can add or remove metadata columns as necessary. For instructions, see [Select default metadata columns](#) on page 44.

Default Folder

This field is used to select the folder that is open by default in a new **Object Manager** window. The name of the currently selected default folder is listed in the field.

When the field is empty, all folders at the root level are displayed in a new **Object Manager** window.

Reset Default Folder

Click this button to remove the folder listed in the **Default Folder** field.

Libraries

Settings for specifying the libraries that are automatically loaded into your workspace when you log in to the 1010data Insights Platform.

If you regularly import libraries into your workspace, consider customizing your workspace to have those libraries loaded for you automatically. This is useful, for example, if you use custom extended operations in the Trillion-Row Spreadsheet or often refer to library content in Macro Language code.

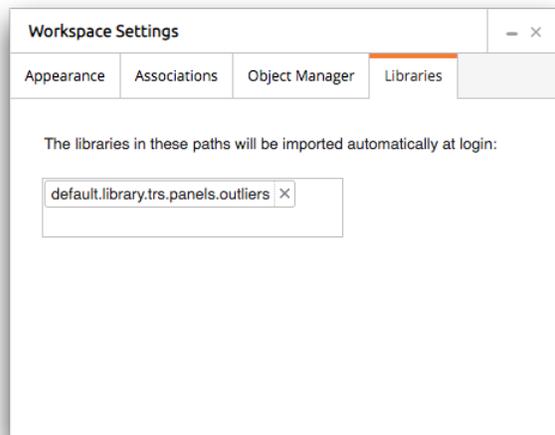


Figure 10: Libraries tab in the Workspace Settings window

The **Libraries** tab contains the following:

The libraries in these paths will be imported automatically at login

This field lists the libraries which are automatically imported when you log in. Each library is identified by its full path (path and name).

You can add or remove libraries as necessary. For instructions, see [Specify automatically imported libraries](#) on page 45.

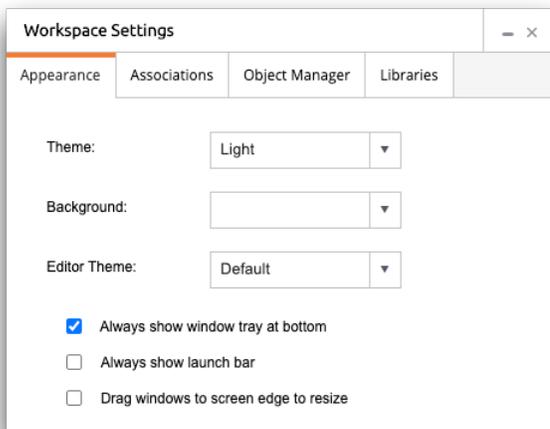
Customize your workspace appearance

Customize the appearance of your workspace.

You can choose how your 1010data Insights Platform workspace looks. The appearance options include color themes and canvas background images.

To customize your workspace appearance:

1. From the workspace menu, select **Session > Workspace Settings**.
The Insights Platform displays the **Appearance** tab on the **Workspace Settings** window.



2. As desired, adjust the workspace appearance settings.

For a description of the available option, see [Appearance](#) on page 38.

Note: After an option is selected, the Insights Platform automatically saves the setting.

Define object associations

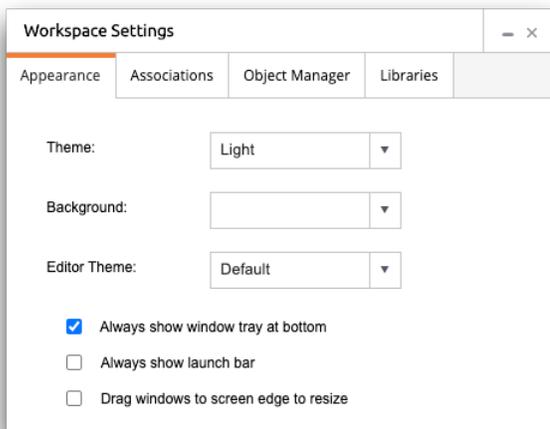
Choose how objects behave when opened from the Object Manager.

You can choose the default behavior for how 1010data Insights Platform objects behave when double-clicked in the Object Manager. For example, you can choose to have tables open in the Trillion-Row Spreadsheet by default.

To define object associations:

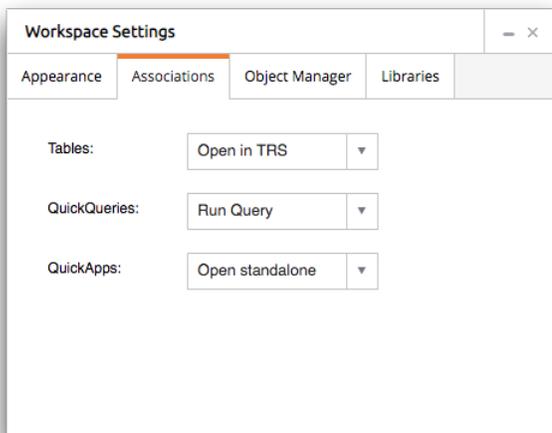
1. From the workspace menu, select **Session > Workspace Settings**.

The Insights Platform displays the **Appearance** tab on the **Workspace Settings** window.



2. Click the **Associations** tab.

The Insights Platform displays the association options.



3. As desired, adjust the object association settings.

For a description of the available option, see [Associations](#) on page 39.

Note: After an option is selected, the Insights Platform automatically saves the setting.

Select default metadata columns

Choose the metadata columns that are displayed by default in a new **Object Manager** window.

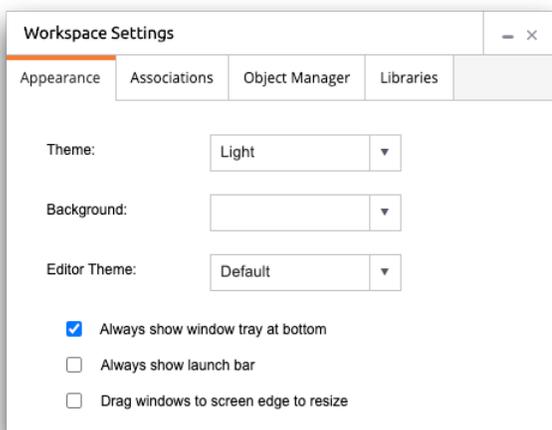
You can choose the columns of metadata you want to see each time a new **Object Manager** window is opened. Changes apply to new **Object Manager** windows and do not affect any **Object Manager** windows that are currently open.

Note: Defining default metadata columns does not prevent you from showing or hiding the metadata columns directly from within the Object Manager. This setting only applies to the metadata columns that are shown by default in a new **Object Manager** window.

To select default metadata columns:

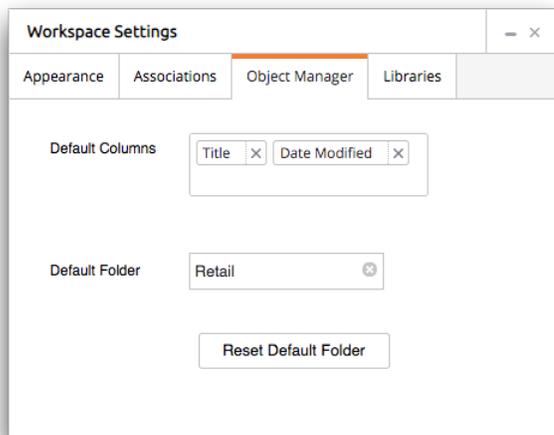
1. From the workspace menu, select **Session > Workspace Settings**.

The 1010data Insights Platform displays the **Appearance** tab on the **Workspace Settings** window.



2. Click the **Object Manager** tab.

The Insights Platform displays the Object Manager options.



3. Choose the default metadata columns by doing the following:

Option

To display a column by default

Description

Click the **Default Columns** field and then select the appropriate columns from the drop-down list.

Note: Although you can add metadata columns in any order, you cannot change the order in which metadata columns are displayed in the Object Manager.

To hide a column by default

In the **Default Columns** field, click the **Remove** (**×**) icon next to its column label.

For a description of the available metadata columns, see [Metadata columns](#) on page 91.

Note: After an option is selected, the Insights Platform automatically saves the setting.

Specify automatically imported libraries

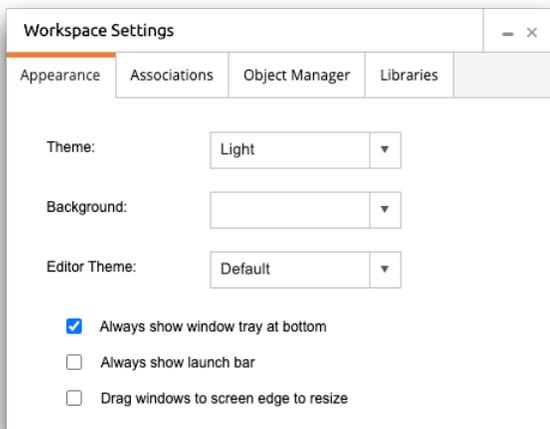
Choose the libraries that are automatically loaded into your workspace when you log in to the 1010data Insights Platform.

You can choose the libraries you want imported automatically each time you log in to the Insights Platform. Changes apply to new sessions and new workspaces; they do not affect your current workspace.

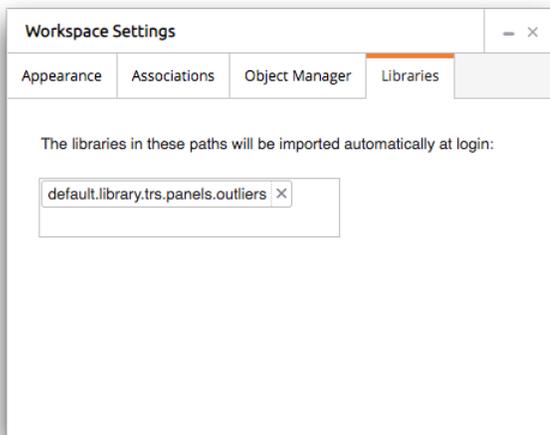
Note: Defining automatically imported libraries does not prevent you from importing libraries manually. This setting only applies to the libraries that are automatically imported when you log in to the platform.

To specify automatically imported libraries:

1. From the workspace menu, select **Session > Workspace Settings**.
The Insights Platform displays the **Appearance** tab on the **Workspace Settings** window.



2. Click the **Libraries** tab.
The Insights Platform displays the libraries setting.



3. Choose the libraries you want automatically imported by doing the following:

Option	Description
To add a library	In the field, type the full path (path and name) of the library and press Enter .
To remove a library	In the field, click the Remove (X) icon next to the library.

Note: After you add or remove a library, the Insights Platform automatically saves the setting.

Virtual memory

The quantity of data that your session can handle at the same time is determined by the amount of virtual memory assigned to your user ID.

The amount of virtual memory allocated to your 1010data Insights Platform session is displayed to the right of the session memory meter.

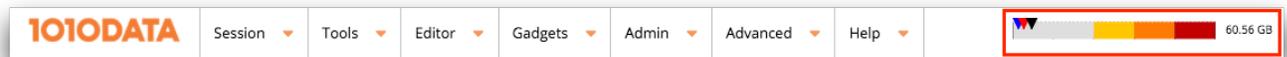


Figure 11: Session memory meter

As you perform your analyses in the Insights Platform, the meter reflects how much virtual memory you have used. You can think of this value as a current high-water mark.

There are limits on the size of a process's virtual memory (i.e., on the amount of data that the session can handle at once). If you work with relatively small tables, your process(es) use small amounts of virtual memory, and you should never even come close to the limit. But if you are dealing with large tables with hundreds of millions of rows, it is possible that you may run into the limit *under certain circumstances*. This qualification is important: Just because you are dealing with large tables doesn't mean that you will have a problem.

The Insights Platform is designed to process data piecewise, dealing with relatively small amounts of data at a time rather than all of it at once. But there are cases when it may be necessary to process large amounts of data at one time, either because there is no practical way to do it piecewise or because doing it piecewise would seriously impact performance.

Troubleshooting memory issues

If your current session starts to behave strangely (e.g., you see cryptic error messages) or stops responding altogether, it may be because you are approaching or have exceeded your virtual memory limit. After the virtual memory limit is surpassed, the system can behave unpredictably.

Note: If you see any message containing the word `wsfull`, it is almost certainly a sign that you have hit the virtual memory limit.

Even if you do run out of virtual memory, the system as a whole is unaffected. There is generally no impact on other users, and you won't affect any of the data on the system.

You always have the option of logging in again. That will end your first session and set things right. But it is sometimes possible to recover without losing the session. Also, even if you start a new session, you will have to do things a bit differently the second time around; otherwise, you will end up in the same place.

Here are some steps we recommend:

- **Clear the cache**

First, see if you can recover without logging in again by clearing the cache. See [Clear the cache](#) on page 48 for details.

If the virtual memory value goes down, your session is OK; if it is still too high, you will have to log in again.

Note: Remember to save your work before starting a new session.

- **Restructure your query**

There is usually more than one way to do an analysis in the Insights Platform, and some methods use significantly less memory (and run more quickly) than others.

Unfortunately, it is essentially impossible to monitor virtual memory usage at every step in a process or to predict how much virtual memory a given query will use. The only way to ensure that you won't hit the limit would be to impose significant restrictions on all queries. However, such restrictions would end up disallowing many viable queries as well. Rather than imposing such arbitrary restrictions, we put the full power of the Insights Platform under your control. This gives you the flexibility to run a greater variety of analyses using different methods to get the results you need using the most efficient queries possible. The downside, of course, is that you can encounter virtual memory problems; but, as noted above, you can recover and you cannot affect anyone else. However, keeping an eye on your virtual memory usage can help you steer clear of the limit before hitting it.

Clear the cache

Clearing the cache cleans out your virtual memory so that you have the maximum amount available to your session.

You can clear the 1010data Insights Platform cache from the menu bar or from the **Session Status** tab in the **Dashboard**.

Note: Clearing the cache clears it from all open workspaces.

To clear the cache:

Perform one of the following actions:

- From the menu bar, select **Session > Clear Caches**.
- From the **Dashboard**, click **Session Status** and then click **Clear caches**.

The Insights Platform clears your cache and displays a confirmation message.

Rows and columns

Within the 1010data Insights Platform, you will manipulate rows and columns to view your data and perform new computations to gain further insight into your data.

Rows

Each row contains a set of related values. You can find or select rows that meet certain criteria and sort the rows of a table.

A table may contain any number of rows. Using the 1010data Insights Platform web interface, you can analyze any amount of data quickly and easily.

Columns

Every table consists of columns, which each contain data of a particular type. You can rearrange the order of columns, hide columns, show meta information about columns, or create computed columns.

Columns can be identified by either their column label or their column name. The *column label* may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters. If you want to have a multi-line column label, you can use the backtick character (`) to separate the lines (e.g., "Percentage of `Total Sales (%)"). By default, the column heading appears at the top of the column in the 1010data web interface.

Sales Item Detail										
Cols 1 to 11 of 11 , Rows 1 to 18 of 34										
Transaction ID	Account	Store	Date	Item SKU	Units	Sales	Cost	Total Sales	Transaction Total	Percentage of Total Sales (%)
532	478	1	05/15/12	98A	1	0.5	0.25	28.19	1.6	6%
532	478	1	05/15/12	3B7	1	1.1	0.56	28.19	1.6	6%
534	738	1	05/16/12	A96	2	6	2.9	28.19	8.25	29%
534	738	1	05/16/12	65B	1	2.25	1.35	28.19	8.25	29%
535	709	2	05/15/12	CB7	1	1.65	1.1	16.31	3.85	24%
535	709	2	05/15/12	96A	1	1.1	1	16.31	3.85	24%
535	709	2	05/15/12	969	1	1.1	1	16.31	3.85	24%
536	748	2	05/17/12	3A4	3	1.02	0.39	16.31	6.02	37%
536	748	2	05/17/12	366	1	5	1.8	16.31	6.02	37%
537	523	3	05/15/12	CB7	1	1.65	1.1	20.35	10.35	51%
537	523	3	05/15/12	A96	2	6	2.9	20.35	10.35	51%
537	523	3	05/15/12	98A	1	0.5	0.25	20.35	10.35	51%
537	523	3	05/15/12	96A	1	1.1	1	20.35	10.35	51%
537	523	3	05/15/12	3B7	1	1.1	0.56	20.35	10.35	51%
538	668	1	05/18/12	CB7	1	1.65	1.1	28.19	6.95	25%
538	668	1	05/18/12	98A	4	2	1	28.19	6.95	25%
538	668	1	05/18/12	96A	1	1.1	1	28.19	6.95	25%

The *column name* is a way of referring to the column internally in value expressions and selection expressions. The column name may only contain alphanumeric characters or underscores and must begin with an alphabetic character (e.g., `percent_total_sales`). It may not contain any spaces or other special characters.

Sales Item Detail										
Cols 1 to 11 of 11 , Rows 1 to 18 of 34										
transid	account	store	date	sku	units	sales	cost	total_sales	trans_total	percent_total_sales
532	478	1	05/15/12	98A	1	0.5	0.25	28.19	1.6	6%
532	478	1	05/15/12	3B7	1	1.1	0.56	28.19	1.6	6%
534	738	1	05/16/12	A96	2	6	2.9	28.19	8.25	29%
534	738	1	05/16/12	65B	1	2.25	1.35	28.19	8.25	29%
535	709	2	05/15/12	CB7	1	1.65	1.1	16.31	3.85	24%
535	709	2	05/15/12	96A	1	1.1	1	16.31	3.85	24%
535	709	2	05/15/12	969	1	1.1	1	16.31	3.85	24%
536	748	2	05/17/12	3A4	3	1.02	0.39	16.31	6.02	37%
536	748	2	05/17/12	366	1	5	1.8	16.31	6.02	37%
537	523	3	05/15/12	CB7	1	1.65	1.1	20.35	10.35	51%
537	523	3	05/15/12	A96	2	6	2.9	20.35	10.35	51%
537	523	3	05/15/12	98A	1	0.5	0.25	20.35	10.35	51%
537	523	3	05/15/12	96A	1	1.1	1	20.35	10.35	51%
537	523	3	05/15/12	3B7	1	1.1	0.56	20.35	10.35	51%
538	668	1	05/18/12	CB7	1	1.65	1.1	28.19	6.95	25%
538	668	1	05/18/12	98A	4	2	1	28.19	6.95	25%
538	668	1	05/18/12	96A	1	1.1	1	28.19	6.95	25%
538	668	1	05/18/12	969	1	1.1	1	28.19	6.95	25%

Note: You can control whether the column label or the column name is displayed at the top of the column. For more information, see [Data Preferences](#) on page 81.

Column information

Column information is provided throughout the web interface in the 1010data Insights Platform.

There are several locations in the web interface where you can view information about a column in a table. This topic lists the most common ways you can view column information.

To learn about the different column data types available in the Insights Platform, see [Data types](#) on page 53.

Grid View

The Grid view is available from the results pane in both the Trillion-Row Spreadsheet and Macro Language Workshop. When viewing a table in the Grid view, place the pointer over a column header to view information about the column.

The screenshot shows a table titled "Sales Item Detail" with 8 columns: Transaction ID, Account, Store, Date, Item SKU, Units, Sales, and Cost. The "Account" column header is highlighted, and a tooltip is displayed over it. The tooltip contains the following information:

- Type: Integer (Custom Format)
- Origin: Base pub.demo.retail.item
- Description: This is the loyalty card number of the person who made this transaction

The table data is as follows:

Transaction ID	Account	Store	Date	Item SKU	Units	Sales	Cost
transid	account	store	date	sku	units	sales	cost
366				366	-1	-5	-1.84
98A				98A	1	0.5	0.25
3B7				3B7	1	1.1	0.56
A96				A96	2	6	2.9
65B				65B	1	2.25	1.35
CB7				CB7	1	1.65	1.1
96A				96A	1	1.1	1
535	709	2	05/15/12	969	1	1.1	1
536	748	2	05/17/12	3A4	3	1.02	0.39
536	748	2	05/17/12	366	1	5	1.8
537	523	3	05/15/12	CB7	1	1.65	1.1
537	523	3	05/15/12	A96	2	6	2.9
537	523	3	05/15/12	98A	1	0.5	0.25
537	523	3	05/15/12	96A	1	1.1	1
537	523	3	05/15/12	3B7	1	1.1	0.56
538	668	1	05/18/12	CB7	1	1.65	1.1
538	668	1	05/18/12	98A	4	2	1
538	668	1	05/18/12	96A	1	1.1	1

Figure 12: Column information in the Grid view

Column information is contextual and can include the following:

Column name

The name of the column. The column name is the name used to refer to a column in Macro Language code. The column name is displayed only if it is not shown in the column header.

Column label

The label of the column. The column label is an optional descriptive column title. The column label is displayed only if it is not shown in the column header.

Type

The kind of data contained in the column and its display format. For more information, see [Data types and display formats](#) on page 52.

Origin

Information about the source of the column. The origin information varies and is based on whether the column is local, linked, computed, or tabulated.

Local

If the column is physical column, the path to the base table containing the column is listed.

Linked

If the column is a linked column, the path of the foreign table containing the column is listed.

Computed

If the column is a result of a computed column, a definition of the computation is listed.

Tabulated

If the column is a result of a tabulation column break, an indication of whether the column is a row or column break is listed.

If the result of the tabulation is a `<tbl>` column, an indication of whether the column is a table, row, or cell aggregation is listed.

Description

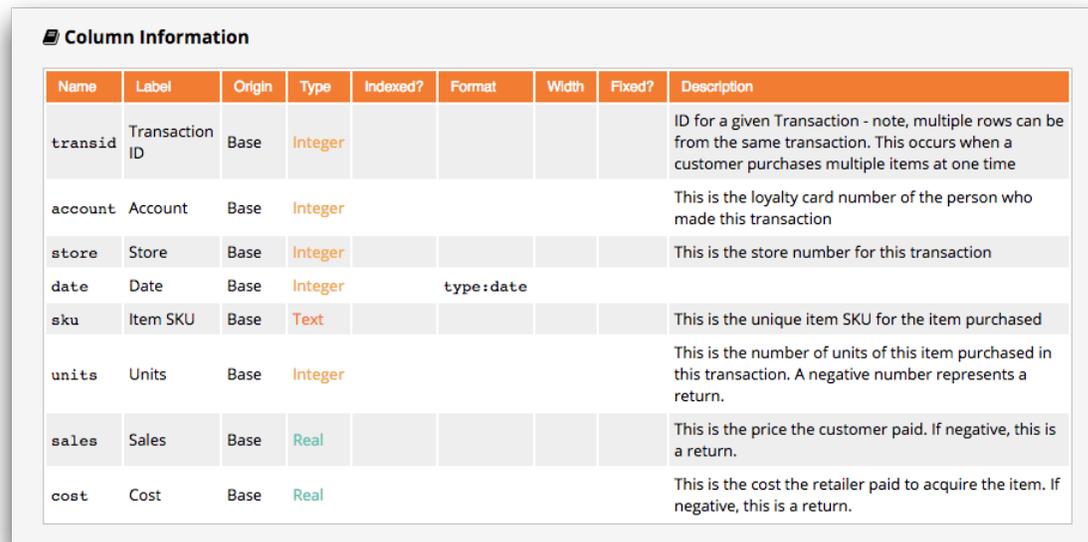
An explanation of the column values and their meaning.

For more information about the Grid view, see [Grid view](#) on page 294.

Note: In addition to the Grid view in the results pane of the Trillion-Row Spreadsheet and Macro Language Workshop, a table may appear as a grid in other areas of the system. You can perform the same actions in any table displayed as a grid to view column information.

Data Dictionary

The data dictionary is available from the results pane in both the Trillion-Row Spreadsheet and Macro Language Workshop.



Name	Label	Origin	Type	Indexed?	Format	Width	Fixed?	Description
transid	Transaction ID	Base	Integer					ID for a given Transaction - note, multiple rows can be from the same transaction. This occurs when a customer purchases multiple items at one time
account	Account	Base	Integer					This is the loyalty card number of the person who made this transaction
store	Store	Base	Integer					This is the store number for this transaction
date	Date	Base	Integer		type:date			
sku	Item SKU	Base	Text					This is the unique item SKU for the item purchased
units	Units	Base	Integer					This is the number of units of this item purchased in this transaction. A negative number represents a return.
sales	Sales	Base	Real					This is the price the customer paid. If negative, this is a return.
cost	Cost	Base	Real					This is the cost the retailer paid to acquire the item. If negative, this is a return.

Figure 13: Column information in the Data Dictionary

Column biscuits

In certain panels in the Analysis Timeline, and within the data columns panel of the Chart Builder, columns are displayed as biscuits. Each column biscuit is color coded so that you can quickly determine its data type. In addition, you can also place the pointer over a biscuit in the timeline to view column information. For more information about the timeline, see [Analysis Timeline](#) on page 129. For more information about the Chart Builder, see [Chart Builder](#) on page 247.

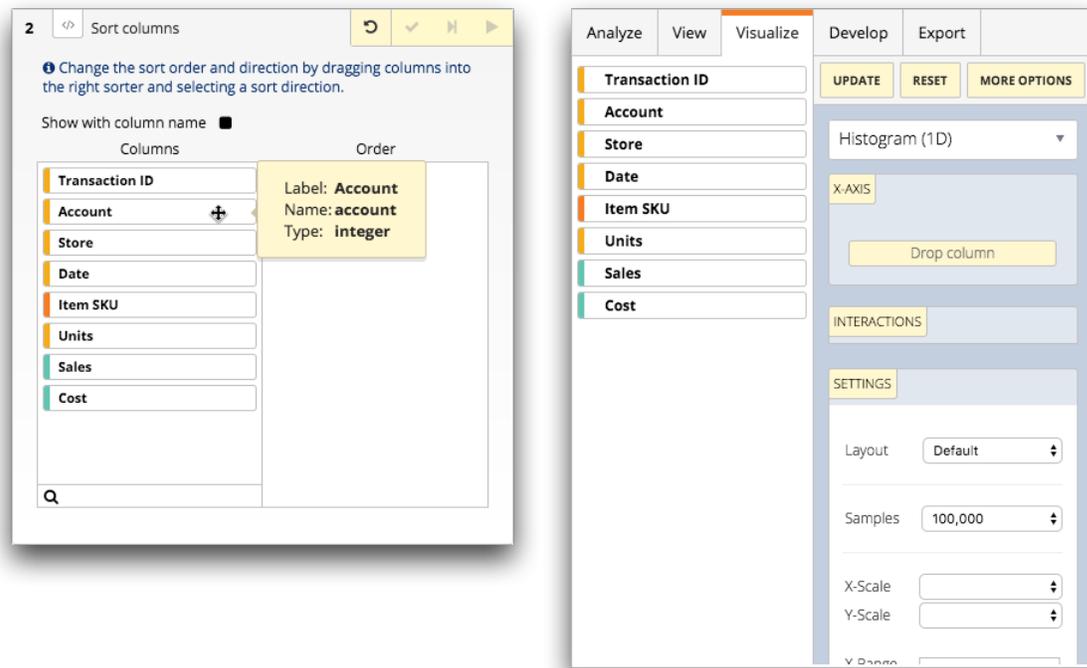


Figure 14: Column information in biscuits

Table 14: Column data type by color

Color	Data Type
Yellow	integer
Purple	big integer <i>(Available as of version 11.25)</i>
Teal	decimal
Orange	text

Data types and display formats

Data types govern how data is stored internally, and *display formats* control how that data is displayed.

Each column in a 1010data Insights Platform table contains data of a particular category, such as numbers, dates, times of day, or alphanumeric text. The type of values determines the data type of the column.

In the Insights Platform, these values (numbers, dates, times, alphanumeric values) are stored in one of the simple *data types*: integer, big integer, decimal, and text. The way in which these values are displayed is defined by the *display format* specified for the column in which those values reside. The way the values are displayed is not necessarily similar to or the same as how the values are stored.

For example, the value 1,234.56 contains a comma, but the internal value is just a decimal number (1234.56). Commas may aid how the values are read by users, but they are not included in the data that is stored. Furthermore, although two decimal places are displayed, the stored value may be more precise. Though the number appears to be 1234.56, it may actually be 1234.55574. Displaying the number rounded to two decimal places is part of the column's *display format*.

It may or may not be easy to tell what the data type of a column is by looking at it. For instance, 123 might be an integer, text, or a decimal number with a display format that shows no decimal places, i.e., `dec:0`.

Another example that shows the difference between how a value is represented in the 1010data Insights Platform and how it appears in the table view is dates and time. A date is simply stored as an integer in the *date* form (YYYYMMDD). The display format identifies the value as a date and displays it as such. For instance, a number such as 20031115 may be displayed as:

- 11/15/03 using the `date` display format.
- 11/15/2003 using the `date4y` display format.

Time is also stored as an integer in the *time* form (HHMMSS). A number like 224556 can be displayed as:

- 10:45:56p using the `hms12` display format.
- 22:46 using the `hm24` display format.

Because dates and times are stored as integers, the platform technically allows these columns to be summed or added to other numeric columns; however, the results are generally meaningless. To do arithmetic with such columns, use the functions provided for this purpose, e.g., to add a number of days to a date, use `shift(X;Y)`. For more information, see [shift\(X;Y\)](#).

Data types

There are four simple data types that represent how information is internally stored in the 1010data Insights Platform.

The simple data types in the platform are:

- integer
- big integer
- decimal
- text

The integer data type is often used to hold information like transaction IDs and store numbers. Integers are also used to represent date or time values, using the *date*, *month*, *quarter*, and *time* forms. For instance, the date 12/10/2013 is stored as the integer 20131210. The time 10:25:30 is stored as the integer 102530. For more information, see [Integer](#) on page 54.

The big integer data type supports whole number values that exceed the range that is supported by the integer data type. Most commonly, big integer columns are used in links, selections, and tabulations, since data such as keys and IDs are often stored in columns of this type. For more information, see [Big integer](#) on page 54.

The decimal data type may contain values such as sales prices, unemployment rates, interest rates, and floating-point values. In the 1010data Insights Platform, a decimal data type can also hold a value that contains a combination of date and time information. For more information, see [Decimal](#) on page 54.

The text data type is used for alphanumeric values that may include symbols, and spaces, such as customer IDs, product names, sizes, ZIP codes, or geographic locations. For more information, see [Text](#) on page 55.

In the Insights Platform, the data types are represented in various ways.

Data type	Type name	Sigil	Description
integer	integer int	i	Whole numbers between -2,147,483,646 and 2,147,483,646
big integer	big integer bigint int64	j	Whole numbers between -9,223,372,036,854,775,806 and 9,223,372,036,854,775,806
decimal	decimal dec real	f	64-bit floating-point values

Data type	Type name	Sigil	Description
	double		
text	text alpha	a	String values that contain uppercase and lowercase letters, numbers, spaces, and symbols

The data type of a column is referred to by its type name or its sigil in different contexts. Different interfaces may tend to use one or the other or both.

Data types are different from *display formats*, which control how values are displayed within columns.

Integer

The integer data type (`int`) is used to represent whole numbers that can be stored within 32-bits.

Columns of type integer can store information such as:

- Transaction IDs (3026049)
- Identification numbers (24)
- Dates (20161201)
- Time (121837)

Dates and time are represented by the integer data type. For more information, see [Dates and time](#) on page 55.

Valid integer values are -2,147,483,646 to 2,147,483,646.

Note: If the values, or the results of mathematical functions applied to these values, will be outside of this range, use the decimal data type or the big integer data type instead.

In some places in the 1010data Insights platform, the integer data type (`int`) is also represented as `i`.

Big integer

The big integer data type (`bigint`) is used to represent whole numbers that are outside the range of the integer data type and can be stored within 64 bits.

In most cases, columns of type `bigint` are used in links, selections, and tabulations. Columns of type `bigint` can also be specified as the `G` argument in `g_` functions such as `g_sum(G;S;O;X)`.

Columns of type big integer can store information that cannot be represented within 32 bits such as:

- Keys
- Universally unique identifiers

Values in a column of type big integer are -9,223,372,036,854,775,806 to 9,223,372,036,854,775,806.

In some places in the 1010data Insights platform, the big integer data type (`bigint`) is also represented as `j`.

(Available as of version 11.25)

Decimal

The decimal data type (`dec`) is used to represent 64-bit floating point values.

Columns of type decimal can store information such as:

- Sales prices (6.95)
- Unemployment rates (7.1)
- Interest rates (5.35)

Columns of type decimal are often used to represent values that contain a fractional part, e.g., 1.25, or numbers that fall outside the valid range for columns of type integer.

Use decimal data types when the values need to be represented contain a fractional part, or if they fall outside the valid range for integers, which is -2,147,483,646 to 2,147,483,646. For example, although UPC codes do not have a fractional part, they should be stored as a decimal data type instead of an integer because they are typically 12-digit numbers. Decimal values may be positive or negative.

Numbers cast to decimal data type are computed & stored as IEEE 754 double floating point values. In most cases, the system holds a minimum 15 places of decimal precision. The maximum is roughly 1.7976931348623158e308.

The decimal data type also holds the floating point value that represents the date and time together as a *date+time* value. For more information, see [Date+Time](#) on page 56.

In some places in the 1010data Insights platform, the decimal data type (`dec`) is also represented as `f`.

Text

The text data type (`text`) is used to represent values that contain uppercase and lowercase letters, numbers, spaces, and symbols.

Columns of type text can store information such as:

- Customer IDs (`9a90e2f6`)
- Product names (`GARLIC GUACAMOLE ZESTY`)
- Sizes (`500 ML`)
- Geographic locations (`Boulder, CO`)

Columns of type text may have values that appear to be integers, decimals, or big integers. It may be difficult to distinguish these values from other numeric values in columns of type decimal or type integer. Text type values cannot be summed (via a quick summary) or added to another column (using a computed column). There are some numeric values that should be stored in columns of type text. For example:

- Social security numbers
- Zip codes (`00616`)
- SKUs (`94873`)

These numeric values are not added together, are not added to other values, and may begin with a 0, which would be removed in an integer, decimal, or big integer column, even though it is a valid and necessary part of that value.

Columns of type integer, decimal, or big integer should be used when mathematical functions need to be performed on the numeric values.

In some places in the 1010data Insights platform, the text data type (`text`) is also represented as `a` or `alpha`.

Dates and time

While dates and times are typically stored in columns that have an integer or decimal data type, they must be specified in a valid form to produce the expected results.

Most date and time values are represented in columns that have an integer data type. The *date+time* form, which contains both date and time information in one value, is unique, and it is represented by a column that has a decimal data type. Functions that operate on dates and times sample the input values in the column to make sure they are in the correct data type (integer or decimal). As long as the data type is correct, the function is performed on the input. Note that the values in the column are not checked to make sure that they are valid. If the values are not valid, the function returns an erroneous or illogical output value.

Dates

Values that contain date-related information can be specified in the following forms:

Form	Syntax	Example	Value
<i>date</i>	YYYYMMDD	11/5/2017	20171105
<i>month</i>	YYYYMM	11/2017	201711
<i>quarter</i>	YYYYQ	4Q2017	20174
<i>year</i>	YYYY	2017	2017

For dates to be considered valid and produce a correct result, they must appear in the format above, and the values for *MM*, *DD*, and *Q* must stay within the expected range. *MM* values range from 01-12, *DD* values from 01-31, and *Q* 1-4. If a function is expecting a date value as input and the value passed to the function is outside the valid range or supported format, the function returns an erroneous or illogical output.

Note: The values YYYYMMDD and YYYYMM are "tolerant of decoration", meaning that you can add slashes or other decoration. For example, 20171105 can be displayed as 11/5/2017.

Time

Values that contain time-related information can be specified in following format:

Form	Syntax	Example	Value
<i>time</i>	HHMMSS	12:46:18	124618

For times to be considered valid and produce a correct result, they must appear in the syntax above, and the values for each portion of *HH*, *MM*, and *SS* must stay within the expected range. *HH* values range from 00-23, and *MM* and *SS* values range from 00-59. If a function is expecting a time value as input and that value is outside the valid range or supported syntax, the function returns an erroneous or illogical output.

Date+Time

Values that contain a combination of date and time information can be specified in the following form:

Form	Syntax	Example	Value
<i>date+time</i>	XXXX.YYYYYYYYYYYY	12/10/13_10:25:30	-7691.56562470746

The *date+time* input uses a convention similar to the Julian date, but uses 1/1/2035 as the *epoch* (or reference) point. Specifically, the value XXXX.YYYYYYYYYYYY is a signed float representing the number of days before or after the epoch. The value is typically negative because the count uses the epoch as the starting point and counts backwards to the given *date+time*:

- XXXX is the number of full days from the given date to the epoch
- .YYYYYYYYYYYY is the fraction of the current day remaining until midnight

So, for the above example, on 12/10/2013 at 10:25:30 there is still .56562470746 of the day remaining until midnight and 7691 full days until midnight on 1/1/2035.

Values in the *date+time* form are stored in columns that have a decimal data type.

Display formats

Display formats control how values are displayed within columns in the 1010data Insights Platform.

Syntax

The syntax for using display formats is:

```
[type:value;] [width:value;] [dec:value]
```

where the *value* for each parameter is as follows:

parameter	value
type	Any Display format listed in the tables below (e.g., <code>nocommas</code> , <code>monthshort4y</code>)
width	The number of characters to display in the column. Valid values are 1 through 100. If the contents of a cell exceeds the specified width, the value is truncated. In this case, the number of visible characters is reduced by two, and " >" is appended to the value. For example, for <code>width:4</code> , the string <code>Example</code> would be displayed as <code>Ex ></code> . Clicking on the > displays the full contents of the cell in a separate window.
dec	The number of decimal places to show (valid values are 0 through 9) Note: For integer values, a decimal point and that many 0's will be appended to the integer. The <code>dec</code> parameter does not affect text values.

A single `type`, `width`, or `dec` format may each be applied to a particular column.

Note: Multiple parameters of the same format will be removed; only the first instance will be retained. For example, if the specified display format is `format="type:num;type:char;width:3"`, it will be changed to `format="type:num;width:3"` when the query is submitted.

Examples

The `secs (X)` function returns the number of seconds since midnight for the time given as input. With no display format applied, a computed column using this function would use the default display format, `num`, and show values like `54,163` and `49,525`.

To display the values in the computed column without the commas, use the `nocommas` display format:

```
<willbe name="secsexample" value="secs(thistime)" label="Secs since `midnight"
format="type:nocommas"/>
```

Those same values would now be displayed as `54163` and `49525`.

Display formats can also show only a certain number of decimal places for a column. For example, to show two decimal places for the numbers in a particular column, specify the following display format for that column:

```
<willbe name="margin" value="sales-cost" label="Margin"
format="type:num;dec:2"/>
```

Number display formats

The following display formats can be applied to values with the integer or decimal data type:

Display format	Example value	Displayed value
<code>num</code>	1234567.89	1,234,567.89
<code>nocommas</code>	11,121,314.15	11121314.15
<code>pct</code>	0.9583	95.83%

Currency display formats

Currency formatting uses an additional property: `unit:[VALUE]`, where the `[VALUE]` is a valid ISO 4217 currency letter code, as shown below:

```
format="type:currency;width:10;unit:EUR"
```

For the full list of currency codes, see [Currency codes \(ISO 4217\)](#) on page 59.

The default value if no code is provided is USD. The number of decimal places shown by default is determined by the standard for the particular currency code used. If a value is provided to the `dec` property, that value will override the default for that currency.

Currency display formats can be applied to columns of integer or decimal type:

Display format	Example unit value	Example value	Displayed value
currency	EUR	123456.789	€123,456.79

Date display formats

The following display formats can be applied to integers in the *date* form (YYYYMMDD):

Display format	Example value	Displayed value
date	19981015	10/15/98
date4y	19981015	10/15/1998
ansidate	19981015	1998-10-15

Note: If the date display formats are applied to numbers that are not in the *date* form, the behavior is unexpected.

Month display formats

The following display formats can be applied to integers in the *month* form (YYYYMM):

Display format	Example value	Displayed value
month	200310	10/03
month4y	200310	10/2003
monthshort	200310	Oct03
monthshort4y	200310	Oct2003
monthshortdash	200310	Oct-03
monthshortdash4y	200310	Oct-2003
monthlong	200310	October03
monthlong4y	200310	October2003
monthlongdash	200310	October-03
monthlongdash4y	200310	October-2003

Note: If the month display formats are applied to numbers that are not in the *month* form, the behavior is unexpected.

Quarter display formats

The following display formats can be applied to integers in the *quarter* form (YYYYQ):

Display format	Example value	Displayed value
quarter	20032	2Q03
quarter4y	20032	2Q2003

Note: If the quarter display formats are applied to numbers that are not in the *quarter* form, the behavior is unexpected.

Time display formats

The following display formats can be applied to integers in the *time* form (*HHMMSS*):

Display format	Example value	Displayed value	Notes
hms24	224556	22:45:56	
hms12	224556	10:45:56p	
hm24	224556	22:46	The time is rounded to the nearest minute.
hm12	224556	10:46p	The time is rounded to the nearest minute.

Note: If the time display formats are applied to numbers that are not in the *time* form, the behavior is unexpected.

Date+Time display formats

The following display formats can be applied to values with the decimal data type in the *date+time* form (*XXXX.YYYYYYYYYYYY*):

Display format	Example value	Displayed value
datehms24	-7691.56562470746	12/10/13_10:25:30
ansidatetime	-7691.56562470746	2013-12-10 10:25:30

Note: If the *date+time* display format is applied to numbers that are not in the *date+time* form, the behavior is unexpected.

Text display formats

The following display format can be applied to text:

Display format	Example value	Displayed value	Notes
char	10/15/98	10/15/98	
url	www.1010data.com	www.1010data.com	Creates an <code>http:</code> link of the text. If the text is not a valid URL, the behavior is unexpected.
email	info@1010data.com	info@1010data.com	Creates a <code>mailto:</code> link of the text. If the text is not a valid email address, the behavior is unexpected.

Currency codes (ISO 4217)

The following table provides a list of valid 4217 letter codes.

Table 15: ISO 4217 Codes

Entity	Currency	Code	Minor Unit
AFGHANISTAN	Afghani	AFN	2
ÅLAND ISLANDS	Euro	EUR	2
ALBANIA	Lek	ALL	2
ALGERIA	Algerian Dinar	DZD	2
AMERICAN SAMOA	US Dollar	USD	2
ANDORRA	Euro	EUR	2
ANGOLA	Kwanza	AOA	2
ANGUILLA	East Caribbean Dollar	XCD	2
ANTARCTICA	No universal currency		
ANTIGUA AND BARBUDA	East Caribbean Dollar	XCD	2
ARGENTINA	Argentine Peso	ARS	2
ARMENIA	Armenian Dram	AMD	2
ARUBA	Aruban Florin	AWG	2
AUSTRALIA	Australian Dollar	AUD	2
AUSTRIA	Euro	EUR	2
AZERBAIJAN	Azerbaijani Manat	AZN	2
BAHAMAS	Bahamian Dollar	BSD	2
BAHRAIN	Bahraini Dinar	BHD	3
BANGLADESH	Taka	BDT	2
BARBADOS	Barbados Dollar	BBD	2
BELARUS	Belarussian Ruble	BYR	0
BELGIUM	Euro	EUR	2
BELIZE	Belize Dollar	BZD	2
BENIN	CFA Franc BCEAO	XOF	0
BERMUDA	Bermudian Dollar	BMD	2
BHUTAN	Ngultrum	BTN	2
BHUTAN	Indian Rupee	INR	2
BOLIVIA, PLURINATIONAL STATE OF	Boliviano	BOB	2
BOLIVIA, PLURINATIONAL STATE OF	Mvdol	BOV	2
BONAIRE, SINT EUSTATIUS AND SABA	US Dollar	USD	2

Entity	Currency	Code	Minor Unit
BOSNIA AND HERZEGOVINA	Convertible Mark	BAM	2
BOTSWANA	Pula	BWP	2
BOUVET ISLAND	Norwegian Krone	NOK	2
BRAZIL	Brazilian Real	BRL	2
BRITISH INDIAN OCEAN TERRITORY	US Dollar	USD	2
BRUNEI DARUSSALAM	Brunei Dollar	BND	2
BULGARIA	Bulgarian Lev	BGN	2
BURKINA FASO	CFA Franc BCEAO	XOF	0
BURUNDI	Burundi Franc	BIF	0
CAMBODIA	Riel	KHR	2
CAMEROON	CFA Franc BEAC	XAF	0
CANADA	Canadian Dollar	CAD	2
CABO VERDE	Cabo Verde Escudo	CVE	2
CAYMAN ISLANDS	Cayman Islands Dollar	KYD	2
CENTRAL AFRICAN REPUBLIC	CFA Franc BEAC	XAF	0
CHAD	CFA Franc BEAC	XAF	0
CHILE	Unidad de Fomento	CLF	4
CHILE	Chilean Peso	CLP	0
CHINA	Yuan Renminbi	CNY	2
CHRISTMAS ISLAND	Australian Dollar	AUD	2
COCOS (KEELING) ISLANDS	Australian Dollar	AUD	2
COLOMBIA	Colombian Peso	COP	2
COLOMBIA	Unidad de Valor Real	COU	2
COMOROS	Comoro Franc	KMF	0
CONGO	CFA Franc BEAC	XAF	0
CONGO, DEMOCRATIC REPUBLIC OF THE	Congolese Franc	CDF	2
COOK ISLANDS	New Zealand Dollar	NZD	2
COSTA RICA	Costa Rican Colon	CRC	2
CÔTE D'IVOIRE	CFA Franc BCEAO	XOF	0
CROATIA	Croatian Kuna	HRK	2
CUBA	Peso Convertible	CUC	2
CUBA	Cuban Peso	CUP	2

Entity	Currency	Code	Minor Unit
CURAÇAO	Netherlands Antillean Guilder	ANG	2
CYPRUS	Euro	EUR	2
CZECH REPUBLIC	Czech Koruna	CZK	2
DENMARK	Danish Krone	DKK	2
DJIBOUTI	Djibouti Franc	DJF	0
DOMINICA	East Caribbean Dollar	XCD	2
DOMINICAN REPUBLIC	Dominican Peso	DOP	2
ECUADOR	US Dollar	USD	2
EGYPT	Egyptian Pound	EGP	2
EL SALVADOR	El Salvador Colon	SVC	2
EL SALVADOR	US Dollar	USD	2
EQUATORIAL GUINEA	CFA Franc BEAC	XAF	0
ERITREA	Nakfa	ERN	2
ESTONIA	Euro	EUR	2
ETHIOPIA	Ethiopian Birr	ETB	2
EUROPEAN UNION	Euro	EUR	2
FALKLAND ISLANDS (MALVINAS)	Falkland Islands Pound	FKP	2
FAROE ISLANDS	Danish Krone	DKK	2
FIJI	Fiji Dollar	FJD	2
FINLAND	Euro	EUR	2
FRANCE	Euro	EUR	2
FRENCH GUIANA	Euro	EUR	2
FRENCH POLYNESIA	CFP Franc	XPF	0
FRENCH SOUTHERN TERRITORIES	Euro	EUR	2
GABON	CFA Franc BEAC	XAF	0
GAMBIA	Dalasi	GMD	2
GEORGIA	Lari	GEL	2
GERMANY	Euro	EUR	2
GHANA	Ghana Cedi	GHS	2
GIBRALTAR	Gibraltar Pound	GIP	2
GREECE	Euro	EUR	2
GREENLAND	Danish Krone	DKK	2
GRENADA	East Caribbean Dollar	XCD	2
GUADELOUPE	Euro	EUR	2

Entity	Currency	Code	Minor Unit
GUAM	US Dollar	USD	2
GUATEMALA	Quetzal	GTQ	2
GUERNSEY	Pound Sterling	GBP	2
GUINEA	Guinea Franc	GNF	0
GUINEA-BISSAU	CFA Franc BCEAO	XOF	0
GUYANA	Guyana Dollar	GYD	2
HAITI	Gourde	HTG	2
HAITI	US Dollar	USD	2
HEARD ISLAND AND McDONALD ISLANDS	Australian Dollar	AUD	2
HOLY SEE (VATICAN CITY STATE)	Euro	EUR	2
HONDURAS	Lempira	HNL	2
HONG KONG	Hong Kong Dollar	HKD	2
HUNGARY	Forint	HUF	2
ICELAND	Iceland Krona	ISK	0
INDIA	Indian Rupee	INR	2
INDONESIA	Rupiah	IDR	2
INTERNATIONAL MONETARY FUND (IMF)	SDR (Special Drawing Right)	XDR	N.A.
IRAN, ISLAMIC REPUBLIC OF	Iranian Rial	IRR	2
IRAQ	Iraqi Dinar	IQD	3
IRELAND	Euro	EUR	2
ISLE OF MAN	Pound Sterling	GBP	2
ISRAEL	New Israeli Sheqel	ILS	2
ITALY	Euro	EUR	2
JAMAICA	Jamaican Dollar	JMD	2
JAPAN	Yen	JPY	0
JERSEY	Pound Sterling	GBP	2
JORDAN	Jordanian Dinar	JOD	3
KAZAKHSTAN	Tenge	KZT	2
KENYA	Kenyan Shilling	KES	2
KIRIBATI	Australian Dollar	AUD	2
KOREA, DEMOCRATIC PEOPLE'S REPUBLIC OF	North Korean Won	KPW	2

Entity	Currency	Code	Minor Unit
KOREA, REPUBLIC OF	Won	KRW	0
KUWAIT	Kuwaiti Dinar	KWD	3
KYRGYZSTAN	Som	KGS	2
LAO PEOPLE'S DEMOCRATIC REPUBLIC	Kip	LAK	2
LATVIA	Euro	EUR	2
LEBANON	Lebanese Pound	LBP	2
LESOTHO	Loti	LSL	2
LESOTHO	Rand	ZAR	2
LIBERIA	Liberian Dollar	LRD	2
LIBYA	Libyan Dinar	LYD	3
LIECHTENSTEIN	Swiss Franc	CHF	2
LITHUANIA	Euro	EUR	2
LUXEMBOURG	Euro	EUR	2
MACAO	Pataca	MOP	2
MACEDONIA, THE FORMER YUGOSLAV REPUBLIC OF	Denar	MKD	2
MADAGASCAR	Malagasy Ariary	MGA	2
MALAWI	Kwacha	MWK	2
MALAYSIA	Malaysian Ringgit	MYR	2
MALDIVES	Rufiyaa	MVR	2
MALI	CFA Franc BCEAO	XOF	0
MALTA	Euro	EUR	2
MARSHALL ISLANDS	US Dollar	USD	2
MARTINIQUE	Euro	EUR	2
MAURITANIA	Ouguiya	MRO	2
MAURITIUS	Mauritius Rupee	MUR	2
MAYOTTE	Euro	EUR	2
MEMBER COUNTRIES OF THE AFRICAN DEVELOPMENT BANK GROUP	ADB Unit of Account	XUA	N.A.
MEXICO	Mexican Peso	MXN	2
MEXICO	Mexican Unidad de Inversion (UDI)	MXV	2

Entity	Currency	Code	Minor Unit
MICRONESIA, FEDERATED STATES OF	US Dollar	USD	2
MOLDOVA, REPUBLIC OF	Moldovan Leu	MDL	2
MONACO	Euro	EUR	2
MONGOLIA	Tugrik	MNT	2
MONTENEGRO	Euro	EUR	2
MONTserrat	East Caribbean Dollar	XCD	2
MOROCCO	Moroccan Dirham	MAD	2
MOZAMBIQUE	Mozambique Metical	MZN	2
MYANMAR	Kyat	MMK	2
NAMIBIA	Namibia Dollar	NAD	2
NAMIBIA	Rand	ZAR	2
NAURU	Australian Dollar	AUD	2
NEPAL	Nepalese Rupee	NPR	2
NETHERLANDS	Euro	EUR	2
NEW CALEDONIA	CFP Franc	XPF	0
NEW ZEALAND	New Zealand Dollar	NZD	2
NICARAGUA	Cordoba Oro	NIO	2
NIGER	CFA Franc BCEAO	XOF	0
NIGERIA	Naira	NGN	2
NIUE	New Zealand Dollar	NZD	2
NORFOLK ISLAND	Australian Dollar	AUD	2
NORTHERN MARIANA ISLANDS	US Dollar	USD	2
NORWAY	Norwegian Krone	NOK	2
OMAN	Rial Omani	OMR	3
PAKISTAN	Pakistan Rupee	PKR	2
PALAU	US Dollar	USD	2
PALESTINE, STATE OF	No universal currency		
PANAMA	Balboa	PAB	2
PANAMA	US Dollar	USD	2
PAPUA NEW GUINEA	Kina	PGK	2
PARAGUAY	Guarani	PYG	0
PERU	Nuevo Sol	PEN	2
PHILIPPINES	Philippine Peso	PHP	2

Entity	Currency	Code	Minor Unit
PITCAIRN	New Zealand Dollar	NZD	2
POLAND	Zloty	PLN	2
PORTUGAL	Euro	EUR	2
PUERTO RICO	US Dollar	USD	2
QATAR	Qatari Rial	QAR	2
RÉUNION	Euro	EUR	2
ROMANIA	New Romanian Leu	RON	2
RUSSIAN FEDERATION	Russian Ruble	RUB	2
RWANDA	Rwanda Franc	RWF	0
SAINT BARTHÉLEMY	Euro	EUR	2
SAINT HELENA, ASCENSION AND TRISTAN DA CUNHA	Saint Helena Pound	SHP	2
SAINT KITTS AND NEVIS	East Caribbean Dollar	XCD	2
SAINT LUCIA	East Caribbean Dollar	XCD	2
SAINT MARTIN (FRENCH PART)	Euro	EUR	2
SAINT PIERRE AND MIQUELON	Euro	EUR	2
SAINT VINCENT AND THE GRENADINES	East Caribbean Dollar	XCD	2
SAMOA	Tala	WST	2
SAN MARINO	Euro	EUR	2
SAO TOME AND PRINCIPE	Dobra	STD	2
SAUDI ARABIA	Saudi Riyal	SAR	2
SENEGAL	CFA Franc BCEAO	XOF	0
SERBIA	Serbian Dinar	RSD	2
SEYCHELLES	Seychelles Rupee	SCR	2
SIERRA LEONE	Leone	SLL	2
SINGAPORE	Singapore Dollar	SGD	2
SINT MAARTEN (DUTCH PART)	Netherlands Antillean Guilder	ANG	2
SISTEMA UNITARIO DE COMPENSACION REGIONAL DE PAGOS "SUCRE"	Sucre	XSU	N.A.
SLOVAKIA	Euro	EUR	2
SLOVENIA	Euro	EUR	2

Entity	Currency	Code	Minor Unit
SOLOMON ISLANDS	Solomon Islands Dollar	SBD	2
SOMALIA	Somali Shilling	SOS	2
SOUTH AFRICA	Rand	ZAR	2
SOUTH GEORGIA AND THE SOUTH SANDWICH ISLANDS	No universal currency		
SOUTH SUDAN	South Sudanese Pound	SSP	2
SPAIN	Euro	EUR	2
SRI LANKA	Sri Lanka Rupee	LKR	2
SUDAN	Sudanese Pound	SDG	2
SURINAME	Surinam Dollar	SRD	2
SVALBARD AND JAN MAYEN	Norwegian Krone	NOK	2
SWAZILAND	Lilangeni	SZL	2
SWEDEN	Swedish Krona	SEK	2
SWITZERLAND	WIR Euro	CHE	2
SWITZERLAND	Swiss Franc	CHF	2
SWITZERLAND	WIR Franc	CHW	2
SYRIAN ARAB REPUBLIC	Syrian Pound	SYP	2
TAIWAN, PROVINCE OF CHINA	New Taiwan Dollar	TWD	2
TAJIKISTAN	Somoni	TJS	2
TANZANIA, UNITED REPUBLIC OF	Tanzanian Shilling	TZS	2
THAILAND	Baht	THB	2
TIMOR-LESTE	US Dollar	USD	2
TOGO	CFA Franc BCEAO	XOF	0
TOKELAU	New Zealand Dollar	NZD	2
TONGA	Pa'anga	TOP	2
TRINIDAD AND TOBAGO	Trinidad and Tobago Dollar	TTD	2
TUNISIA	Tunisian Dinar	TND	3
TURKEY	Turkish Lira	TRY	2
TURKMENISTAN	Turkmenistan New Manat	TMT	2
TURKS AND CAICOS ISLANDS	US Dollar	USD	2
TUVALU	Australian Dollar	AUD	2

Entity	Currency	Code	Minor Unit
UGANDA	Uganda Shilling	UGX	0
UKRAINE	Hryvnia	UAH	2
UNITED ARAB EMIRATES	UAE Dirham	AED	2
UNITED KINGDOM	Pound Sterling	GBP	2
UNITED STATES	US Dollar	USD	2
UNITED STATES	US Dollar (Next day)	USN	2
UNITED STATES MINOR OUTLYING ISLANDS	US Dollar	USD	2
URUGUAY	Uruguay Peso en Unidades Indexadas (URUIURUI)	UYI	0
URUGUAY	Peso Uruguayo	UYU	2
UZBEKISTAN	Uzbekistan Sum	UZS	2
VANUATU	Vatu	VUV	0
VENEZUELA, BOLIVARIAN REPUBLIC OF	Bolivar	VEF	2
VIET NAM	Dong	VND	0
VIRGIN ISLANDS (BRITISH)	US Dollar	USD	2
VIRGIN ISLANDS (U.S.)	US Dollar	USD	2
WALLIS AND FUTUNA	CFP Franc	XPF	0
WESTERN SAHARA	Moroccan Dirham	MAD	2
YEMEN	Yemeni Rial	YER	2
ZAMBIA	Zambian Kwacha	ZMW	2
ZIMBABWE	Zimbabwe Dollar	ZWL	2
ZZ01_Bond Markets Unit European_EURCO	Bond Markets Unit European Composite Unit (EURCO)	XBA	N.A.
ZZ02_Bond Markets Unit European_EMU-6	Bond Markets Unit European Monetary Unit (E.M.U.-6)	XBB	N.A.
ZZ03_Bond Markets Unit European_EUA-9	Bond Markets Unit European Unit of Account 9 (E.U.A.-9)	XBC	N.A.
ZZ04_Bond Markets Unit European_EUA-17	Bond Markets Unit European Unit of Account 17 (E.U.A.-17)	XBD	N.A.

Entity	Currency	Code	Minor Unit
ZZ06_Testing_Code	Codes specifically reserved for testing purposes	XTS	N.A.
ZZ07_No_Currency	The codes assigned for transactions where no currency is involved	XXX	N.A.
ZZ08_Gold	Gold	XAU	N.A.
ZZ09_Palladium	Palladium	XPD	N.A.
ZZ10_Platinum	Platinum	XPT	N.A.
ZZ11_Silver	Silver	XAG	N.A.

Dashboard

Quickly access information and data that is relevant to you.

The **Dashboard** opens each time you log in to a new session *sticky* or open a new workspace. It is divided into tabs that provide log in information, announcements, and the status of virtual memory used in your session. In addition, you can access recent workspaces and your favorite objects.

The **Dashboard** can also be accessed by selecting **Session** > **Dashboard** from the workspace menu.

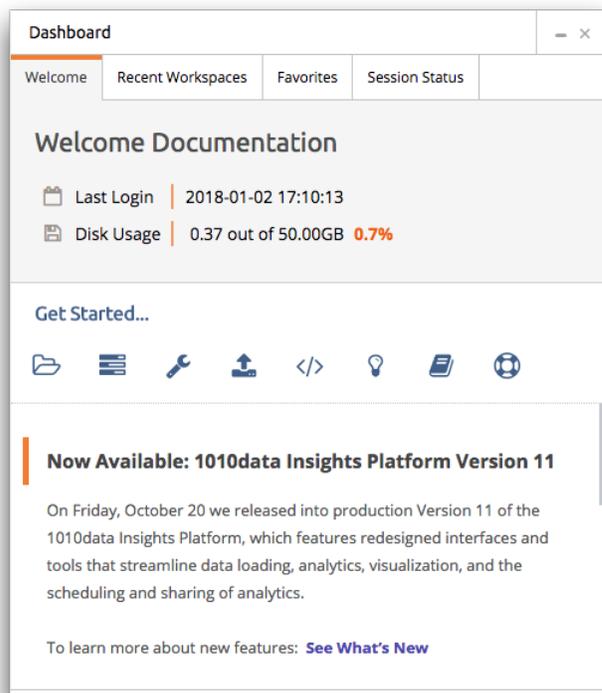


Figure 15: Dashboard window

Welcome

This tab lists your previous log in information and indicates how much space of your total disk quota you are using.

The **Announcements** section contains recent announcements from 1010data pertaining to system-wide changes and updates. It may also include more specific information about your data, queries, and so forth.

Recent Workspaces

This tab lists your 10 most recent saved workspaces.

Favorites

This tab lists your favorite objects such as tables, folders, and queries. You can add favorites in the **Object Manager**. For instructions, see [Add an object to your favorites](#) on page 110.

Session Status

This tab allows you to monitor your virtual memory usage and to clear your cache.

Recent Workspaces

A recent workspace is a saved workspace that is automatically created for you.

Each time you log in to the 1010data Insights Platform, or open a separate workspace from within an existing session, a recent workspace is created. As you use the platform, the saved workspace is automatically updated each time the Insights Platform interacts with the server. This helps ensure that your work is not lost should your browser close unexpectedly. Recent workspaces are also helpful if you want to revisit a recent analysis that you did not save.

Note: Standalone QuickApps (a QuickApp that runs in a separate browser tab) are not saved in your Recent Workspaces.

When you log out, or are logged out due to inactivity, updates to that session's recent workspace cease. The next time you log in, a new recent workspace is created. Unlike a workspace you save yourself, recent workspaces are not permanent; you are limited to the last 10 workspaces. To permanently save a workspace, see [Save your workspace](#) on page 34.

The **Recent Workspaces** tab in the **Dashboard** window lists your automatically saved workspaces. Access it directly by selecting **Session > Recent Workspaces** from the workspace menu.

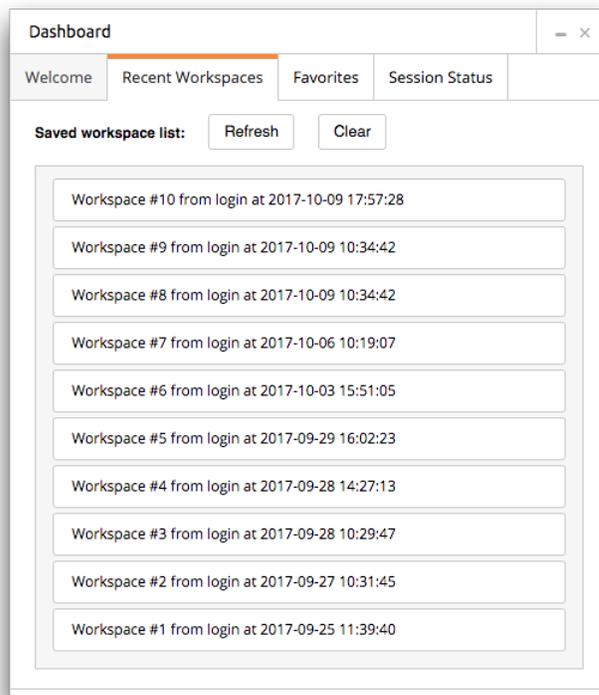


Figure 16: Recent Workspaces tab in the Dashboard

Saved workspace list

Each line in the list represents an automatically saved workspace. Place the pointer over a saved workspace to view additional details. Click a saved workspace to open it in a new browser tab.

Refresh

Although your workspace is automatically saved each time the Insights Platform interacts with the server, the list is not updated while the **Recent Workspaces** tab is open. Click **Refresh** to display any additional saved workspaces.

Clear

This button removes saved workspaces from the list.

Open a recent workspace

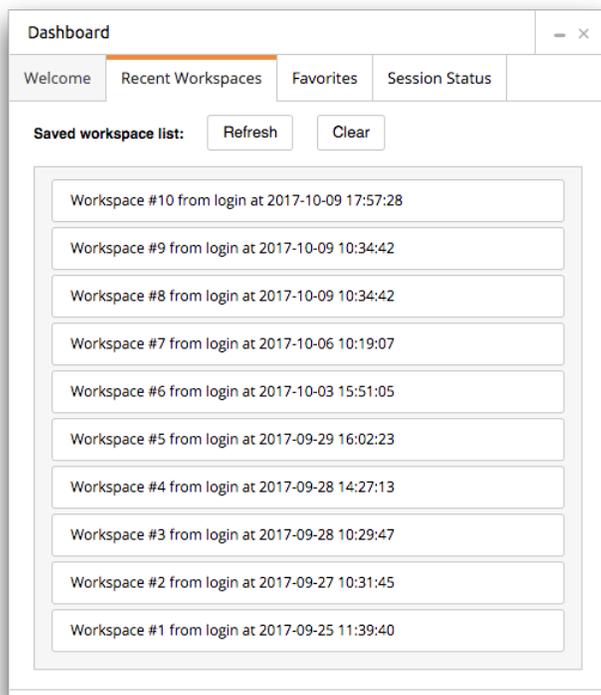
Open an automatically saved recent workspace in a new browser tab.

Each time you log in to the 1010data Insights Platform, or open a separate workspace from within an existing session, a recent workspace is created. As you use the platform, the saved workspace is automatically updated each time the Insights Platform interacts with the server. This helps ensure that your work is not lost should your browser close unexpectedly. Recent workspaces are also helpful if you want to revisit a recent analysis that you did not save.

To open a recent workspace:

1. From the workspace menu, select **Session > Recent Workspaces**.

The Dashboard displays the last 10 automatically saved workspaces in the **Recent Workspaces** tab.



2. Click the recent workspace you want to open.

The Insights Platform opens the recent workspace in a new browser tab.

Favorites

View and interact with objects saved to your list of favorites.

The **Favorites** tab lists your favorite objects such as tables, folders, and queries. You can add favorites in the **Object Manager**. For instructions, see [Add an object to your favorites](#) on page 110.

Note: The first time you open the **Favorites** tab, or when no favorites are saved, the Dashboard provides the option of importing your favorites saved in the legacy 1010data Insights Platform user interface.

You can easily access objects you use often and choose where to open the object. For example, you might want to open a saved query in the Trillion-Row Spreadsheet for ad hoc analysis or open it in the Macro

Language Workshop for development. As necessary, you can reorder objects in the list or remove an object altogether.

Access the **Favorites** tab in the Dashboard directly from the workspace menu by selecting **Session > Favorites**.

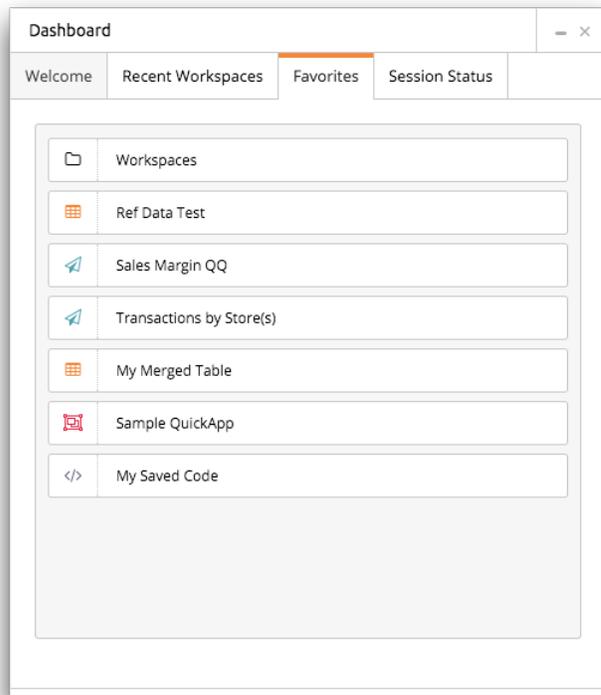


Figure 17: Favorites tab in the Dashboard

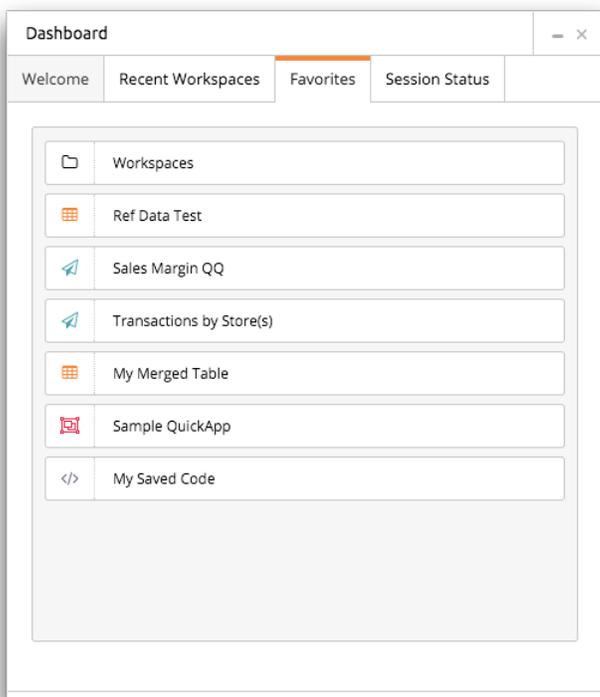
Remove a favorite

Remove an object from your favorites that you no longer use very often.

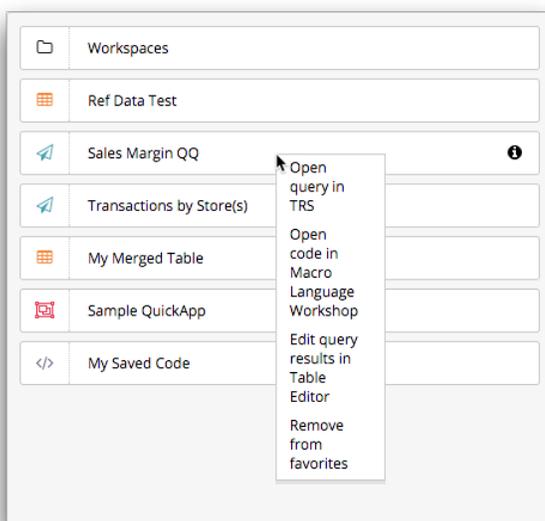
As necessary, you can remove an object from your favorites list in the Dashboard.

To remove a favorite:

1. From the workspace menu, select **Session > Favorites**.
The Dashboard displays your saved favorites in the **Favorites** tab.



2. Right-click the object you want to remove.
The Dashboard displays a menu of options.



3. Select **Remove from favorites**.
The Dashboard removes the object from your favorites list.

User Profile

The **User Profile** window allows you to view your user profile information, change your password, and edit your account settings.

The User Profile is accessed by selecting **Admin > User Profile** from the workspace menu.

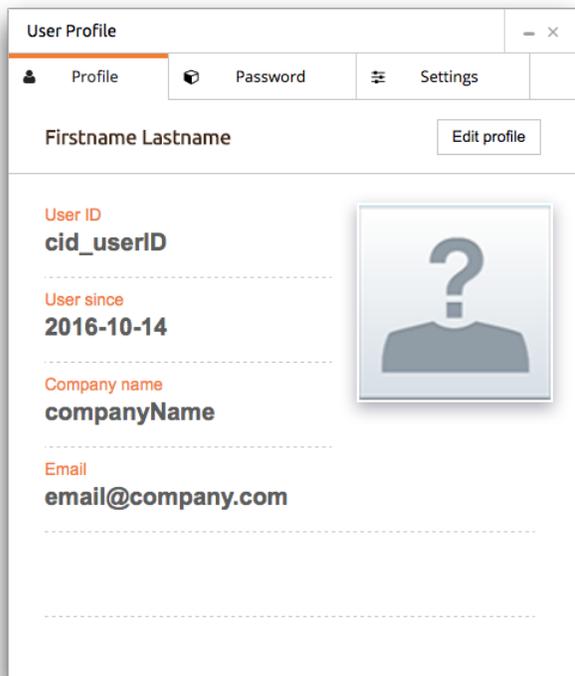


Figure 18: User Profile window

The **User Profile** window contains three tabs:

Profile

The **Profile** tab displays your information, such as your **User ID** and **Company name**.

Password

The **Password** tab allows you to change your password.

Settings

The **Settings** tab allows you to change your settings, such what column headers display or what version of the interface you log into by default.

Profile

The **Profile** tab displays your user information.

The **Profile** tab is displayed by default when the **User Profile** window opens.

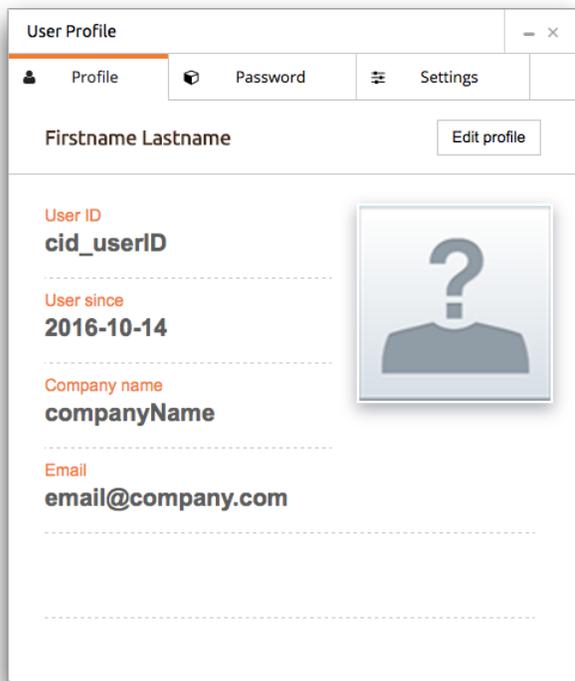


Figure 19: The Profile tab

The **Profile** tab contains:

User ID

This field displays the user's username. Typically, usernames start with your company ID, followed by an underscore (_) and a user specific ID.

User since

Displays the date your user credentials were created.

Company name

Displays the name of your company. This is different from your company ID.

Email

Displays the email address associated with your account. Password reset emails are sent to this address.

Edit profile

Opens a view where you can edit your user information. For more information, see [Edit your user profile](#) on page 83.

Profile picture

Your user profile picture.

Profile Information

The **Edit profile** button brings up the editable user information fields.

Users can edit their profile information. For more information, see [Edit your user profile](#) on page 83.

First name

User's first name

The user's first name is displayed on the **Welcome** tab in the Dashboard. The user's full name is displayed on the **Profile** tab in the User Profile.

Middle initial

User's middle initial

The user's full name is displayed on the **Profile** tab in the User Profile.

Last name

User's last name

The user's full name is displayed on the **Profile** tab in the User Profile.

Email

User's email address

The user's email address is displayed on the **Profile** tab in the User Profile.

Phone

User's phone number

The user's phone number is not displayed in the user interface.

Address

User's address

The user's address is not displayed in the user interface.

City

User's city

The user's city is not displayed in the user interface.

State

User's state

The user's state is not displayed in the user interface.

Postal code

User's ZIP code

The user's ZIP code is not displayed in the user interface.

Country

User's country

The user's country is not displayed in the user interface.

Profile pic URL

A URL that points to a user's profile picture.

The URL must link to an external source, not to a local machine or to a private server. The profile picture is displayed on the **Profile** tab in the User Profile.

Password

The **Password** tab allows you to change your account password.

The **Password** tab is the middle tab in the **User Profile** window. This password change form is different from the **Password Reset** page.

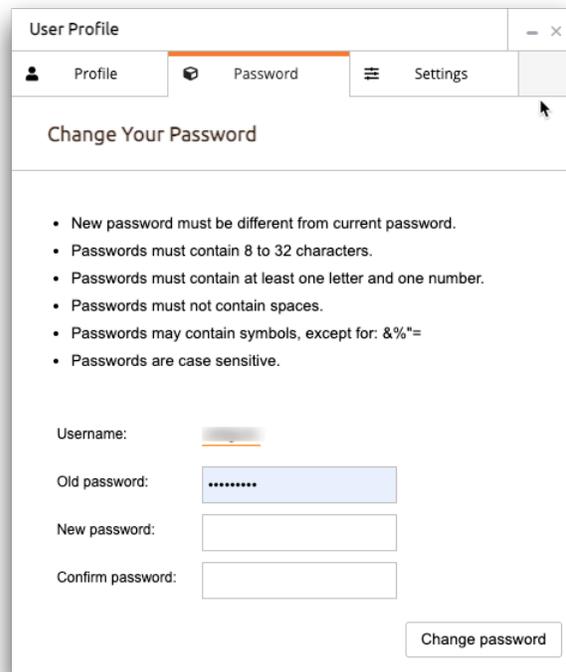


Figure 20: The Password tab

The **Password** tab contains:

List of password criteria

A list of the current criteria to create a valid 1010data Insights Platform user password. For more information, see [Password requirements](#) on page 78.

Username

Username

Old password

Current password

New password

New password, which must comply with the password criteria.

Change password

Click to complete your password change after filling out the required fields.

Password requirements

A list of requirements for a 1010data Insights Platform password.

Users can change their password in the **User Profile** or reset their password if they cannot remember their current password. For more information, see [Change your password](#) on page 84 or [Reset your password](#) on page 14.

A new, valid Insights Platform password:

- must be different from the user's current password.
- must contain 8-32 characters.
- must contain at least one letter and one number.
- may contain symbols. The following symbols are valid:

! @ # \$ % ^ * () _ + - : ; { } [] | < > , . ? ' ~ ` \ /

- must not contain spaces.

Insights Platform passwords are case sensitive.

Settings

The **Settings** tab allows you to change settings related to your session.

Changes made in the **Settings** tab are persistent between sessions.

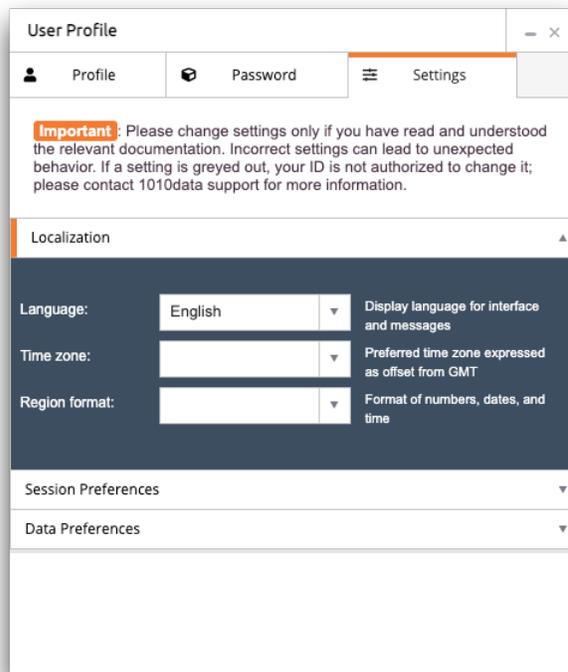


Figure 21: Settings tab in the User Profile window

The **Settings** tab is divided into three categories:

Localization

Specifies how the platform displays information that can be different based on your location, such as your time zone. For more information, see [Localization](#) on page 79.

Session Preferences

Specifies how platform sessions behave, such as how often a keepalive transaction is sent to a proxy server. For more information, see [Session Preferences](#) on page 80.

Data Preferences

Specifies how data is displayed, such as what is displayed in the column header. For more information, see [Data Preferences](#) on page 81.

Localization

Specifies how the platform displays information that can be different based on your location, such as your time zone.

The **Localization** section contains the following settings:

Language

Select the language used to display interface text and messages.

Note: Currently, the only language option is English.

Time zone

Select the time zone used for the date and time on reports.

The options are presented as offsets to the Greenwich Mean Time.

For example, New York is in the Eastern Time Zone of the United States, and Eastern Standard Time (EST) is five hours behind Greenwich Mean Time. Therefore, the appropriate selection for New York is **GMT-5**.

Region format

Select the desired format for dates and times.

This setting determines the order in which dates are presented (or downloaded as formatted values), when using the `date`, `date4y`, and `datehms24` display formats as well as how numbers (i.e., integers and decimals) are displayed. Region format options include the following:

	Date	Date+Time	Integer	Decimal
USA	<i>MM/DD/YY</i>	<i>MM/DD/YY_HH:MM:SS</i>	10,000	0.001
Europe	<i>DD.MM.YY</i>	<i>DD.MM.YY_HH:MM:SS</i>	10.000	0,001
United Kingdom	<i>DD/MM/YY</i>	<i>DD/MM/YY_HH:MM:SS</i>	10,000	0.001
ANSI	<i>YYYY-MM-DD</i>	<i>YYYY-MM-DD HH:MM:SS</i>	10,000	0.001

Users can change their localization settings in the **User Profile**. For more information, see [Change your user settings](#) on page 85.

Session Preferences

Specifies how platform sessions behave, such as how often a keepalive transaction is sent to a proxy server.

The **Session Preferences** section contains the following settings:

Platform default

This controls which version of the platform the system logs into when a version is not specified. The options are:

Default

The default version is generally set by 1010data when a user profile is created. If a default version is not specified, the company's default version is the default.

Prime

Logs the user into the latest version of Prime.

Beta

Logs the user into the latest version of Beta.

V9

Logs the user into the latest version of Prod for version 9.

Idle timeout

The number of minutes a session must be idle before you are logged out.

Keepalive interval

This is for use with accounts that connect to the Insights Platform using a proxy. This setting defines the number of seconds before a query is sent to keep the proxy connection active.

Users can change their session preferences in the **User Profile**. For more information, see [Change your user settings](#) on page 85.

Data Preferences

Specifies how data is displayed, such as what is displayed in the column header.

Column headers

This controls what is displayed in column headers, which are at the top of each column in a table. For more information, see [Columns](#) on page 48.

There are three display options:

Labels

Selecting **Labels** displays only the column's label in the column header.

Labels are optional descriptive column titles. They may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters. Column labels may contain any combination of:

- Uppercase and lowercase letters (a-z, A-Z)
- Numbers (0-9)
- Spaces
- Special characters: ! @ # \$ % ^ & * () _ + - = : ; { } [] | < > , . ? ' " ~ ` \ /

For example, `Product ID, 98430 Total, and *Stores` are all valid column labels.

You can have a multi-line column label, use the backtick character (`) to separate the lines, e.g., `Percentage`Total Sales (%)`.

If the column does not have a label, the column's name is displayed instead.

Names

Selecting **Names** displays only the column's name in the column header.

Names are required, single word identifiers for columns. They are used to identify the column in value expressions and selection expression. Column names:

- Must be a single word consisting of alphanumeric characters (a-z, 0-9) or underscores (_).
- Must start with a letter.
- Must be lower case.
- Must not end with an underscore.
- Must not contain any spaces or other special characters.

For example, `productid` or `product_id` are valid column names, but `Product Id` is not.

Both labels and names

Selecting **Both labels and names** displays both the column name and the column label in the column header.

Repeated Values

This controls what is displayed by default if the same value in a column is repeated in consecutive rows. There are five options:

Value

Display the repeated value in each cell.

Transaction ID	Account	Store
537	523	3
537	523	3
537	523	3
537	523	3
538	668	1
538	668	1
538	668	1
538	668	1

Ditto mark (all columns)

Display the ditto mark (") in place of a repeated value and applies this setting to all columns.

Transaction ID	Account	Store
537	523	3
"	"	"
"	"	"
"	"	"
538	668	1
"	"	"
"	"	"
"	"	"

Ditto mark (fixed columns only)

Display the ditto mark (") in place of a repeated value, and apply this setting to fixed columns only. If a column is not fixed, the repeated value is displayed.

Transaction ID	Account	Store
537	523	3
"	523	3
"	523	3
"	523	3
"	523	3
538	668	1
"	668	1
"	668	1
"	668	1
"	668	1

Blank (all columns)

Display a blank cell in place of a repeated value, and apply this setting to all columns.

Transaction ID	Account	Store
537	523	3
538	668	1
539	25	2

Blank (fixed columns only)

Display a blank cell in place of a repeated value, and apply this setting to fixed columns only. If a column is not fixed, the repeated value is displayed.

Transaction ID	Account	Store
537	523	3
	523	3
	523	3
	523	3
	523	3
538	668	1
	668	1
	668	1
	668	1
	668	1

Users can change their data preferences in the **User Profile**. For more information, see [Change your user settings](#) on page 85.

Edit your user profile

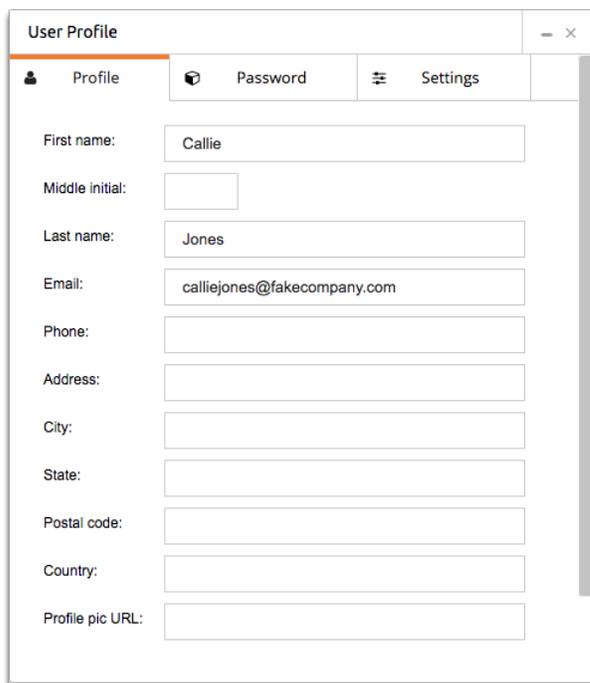
You can edit or update your user information in the **User Profile** window.

Your user information, such as your name, phone number, and profile picture can be changed from the **User Profile** window.

To edit your user profile:

1. Open the **User Profile** window.
2. Click **Edit profile**.

The **Profile** tab changes to display editable user information. For more information, see [Profile Information](#) on page 76.



The screenshot shows a window titled "User Profile" with three tabs: "Profile", "Password", and "Settings". The "Profile" tab is active. The form contains the following fields:

First name:	<input type="text" value="Callie"/>
Middle initial:	<input type="text"/>
Last name:	<input type="text" value="Jones"/>
Email:	<input type="text" value="calliejones@fakecompany.com"/>
Phone:	<input type="text"/>
Address:	<input type="text"/>
City:	<input type="text"/>
State:	<input type="text"/>
Postal code:	<input type="text"/>
Country:	<input type="text"/>
Profile pic URL:	<input type="text"/>

3. Edit your profile information.

4. Click **Save changes**.

Your changes are saved and the **Profile** tab returns to the default display.

Change your password

You can change your password from the **User Profile** window.

To change your password:

1. Open the **User Profile** window.

2. Click **Password**.

The Insights Platform displays the **Password** tab.

The screenshot shows a 'User Profile' window with three tabs: Profile, Password, and Settings. The 'Password' tab is selected. The main content area is titled 'Change Your Password'. It contains a list of requirements for a new password:

- New password must be different from current password.
- Passwords must contain 8 to 32 characters.
- Passwords must contain at least one letter and one number.
- Passwords must not contain spaces.
- Passwords may contain symbols, except for: &%"' =
- Passwords are case sensitive.

Below the requirements are four input fields:

- Username:** A text field with a default value.
- Old password:** A password field with masked characters (dots).
- New password:** An empty password field.
- Confirm password:** An empty password field.

A 'Change password' button is located at the bottom right of the form.

3. Complete the following fields:

Username

Your username is filled in by default.

While this field is editable, you can only change your password.

Old password

Enter your current password.

New password

Enter your new password.

4. Click **Change password**.

The Insights Platform changes your password and briefly displays a confirmation message.

Change your user settings

The user profile contains the user settings that are persistent between sessions.

If you need to change your user settings, such as your time zone, default platform version, or default interface, you can use the User Profile window.

1. Click **Admin** in the menu bar.

2. Click **User Profile**.

The User Profile window opens.

User Profile

Profile Password Settings

Firstname Lastname Edit profile

User ID
cid_userID

User since
2016-10-14

Company name
companyName

Email
email@company.com

3. Click the **Settings** tab.

User Profile

Profile Password Settings

Important : Please change settings only if you have read and understood the relevant documentation. Incorrect settings can lead to unexpected behavior. If a setting is greyed out, your ID is not authorized to change it; please contact 1010data support for more information.

Localization

Language: Display language for interface and messages

Time zone: Preferred time zone expressed as offset from GMT

Region format: Format of numbers, dates, and time

Session Preferences

Data Preferences

4. Expand the category you want to edit.

5. Edit the user setting. The change is saved automatically.

For information about the settings you can change, see [Settings](#) on page 79.

Object Manager

The Object Manager allows you to view and traverse the hierarchy of folders, tables, and other objects to which you have access, and to perform actions on those objects.

The Object Manager is accessed by selecting **Tools > Object Manager** from the workspace menu.

Note: Depending on your access and permissions, the folders you see and have access to may be different from those in the example below.

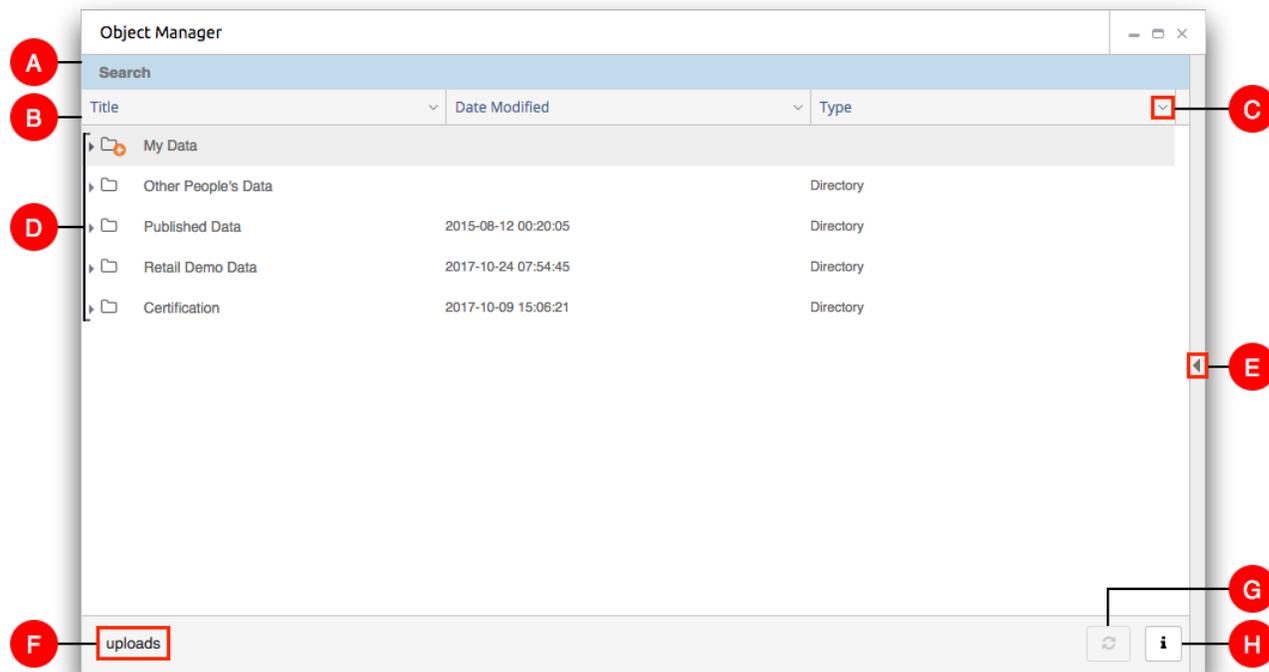


Figure 22: Object Manager

A. Search

Use this field to search for tables, folders, and other objects. For instructions, see [Search for objects](#) on page 103.

B. Metadata columns

Metadata columns display object information. Information for each metadata item is displayed in its own column. This allows you to display only the metadata information that you need. In addition, you can sort the object list by each metadata column. For more information, see [Metadata columns](#) on page 91.

C. Options menu

This drop-down arrow allows you to define the sort order of objects and to choose the object information you want displayed in the object browser.

D. Object browser

The object browser is a file browser that provides access to the objects you have permission to use.

Traverse the folder hierarchy by clicking the **Expand** (▸) icon to the left of a folder to open it and the **Collapse** (▹) icon to close it.

Objects in the file browser are represented by icons. For a list and description of these icons, see [Object icons](#) on page 89.

If you own an object, or have permission to save or upload an object into a folder, permission icons appear next to the object in the browser. For more information, see [Permission icons](#) on page 90.

When an object is selected, it is highlighted in gray and its path is displayed in the lower left-side of the window. In addition, as appropriate, the **Refresh** and **Information** icons are enabled.

E. Show/Hide

Click the **Show** (◀) icon to display the right-side object browser pane in the Object Manager. Viewing both the left- and right-side panes at the same time can help make some tasks easier. For example, when moving an object from one folder to another, you can view each folder, side-by-side, in the two panes.

When both object browser panes are visible, click the **Show** (◀) icon to hide the left-side pane or click the **Hide** (▶) icon to hide the right-side. Additionally, you can resize the panes using the split bar handle located between the **Show** and **Hide** icons.

F. Path

When an object is selected, its full path is displayed. The full path includes the path to, and the name of, the object.

G. Refresh

Click this icon to update an object in your session with data from the database. For more information, see [Refresh an object](#) on page 115.

H. Information

This icon opens the information dialog which displays additional information about a selected object. If you have permission, you can also modify the properties of the selected object. For more information, see [Information dialog](#) on page 104.

Object browser

The object browser is a file browser that lists all of the objects to which you have access.

The folders and objects visible to you in the 1010data Insights Platform are displayed in the object browser. You can open objects such as tables and queries, which are organized in folders.

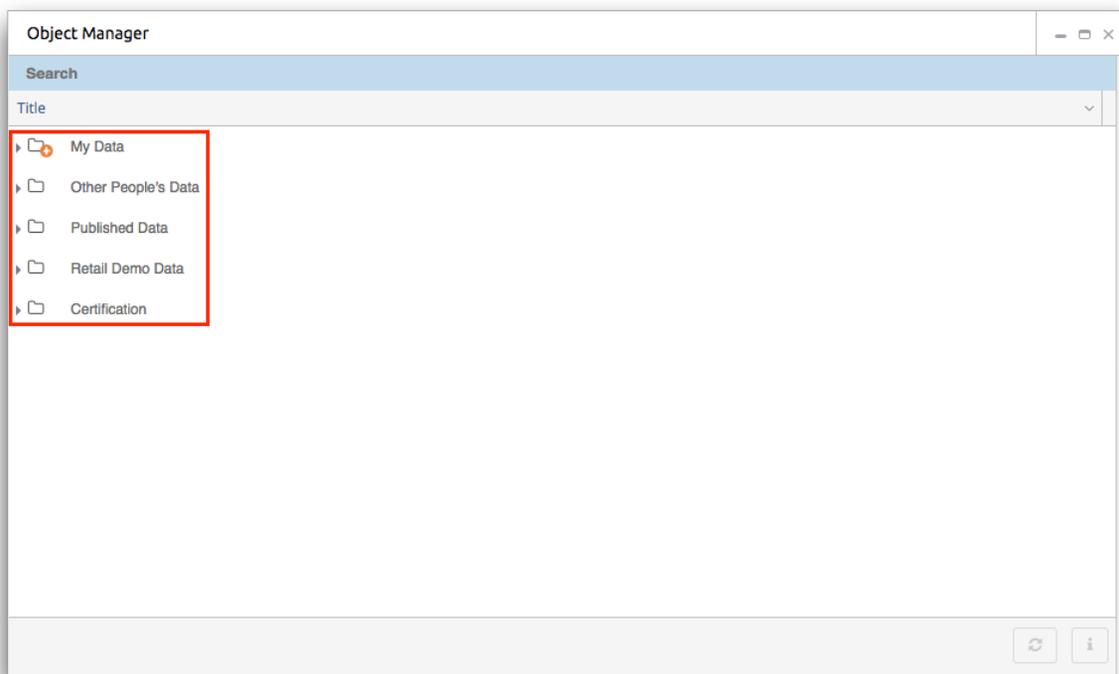


Figure 23: Object browser in the Object Manager

To expand a folder, click the **Expand** (▸) icon. To collapse a folder, click the **Collapse** (◀) icon.

To open an object, simply double-click anywhere in that object's row. The Object Manager will open the object in the Insights Platform tool designated in your association settings. For example, you can choose to have tables open in the Trillion-Row Spreadsheet by default. For more information, see [Associations](#) on page 39.

Note: If the object does not open after you double-click it, the object association for the object type may not be defined. To set the default object association, see [Define object associations](#) on page 43.

You can also right-click an object in the object browser to open a menu of other actions you can perform on the object. For example, you can add an object to your favorites list or choose to open the object in a specific Insights Platform tool. Additional actions available in the right-click menu vary and are based upon the selected object. For a list of additional actions by type, see [Object actions by type](#) on page 93.

Note: Some actions, such as deleting an object or adding a subfolder within a folder, are available only for objects you own. For more information, see [Permission icons](#) on page 90.

Object icons

Descriptions of object icons in the object browser.

Objects in the object browser are represented by icons which are described in the table below.

Table 16: Object icons

Icon	Name	Description
	Folder	<p>Folders represent directories where saved objects are stored.</p> <p>To expand a folder, click the Expand (▸). To collapse a folder, click the Collapse (◀) icon.</p>

Icon	Name	Description
	Table	A table is the permanent, unchanging version of the data that is saved on the server. When you open a table in the Insights Platform and perform an action or query, it becomes a worksheet. A worksheet is a working copy of the table. In worksheets, you can work with and manipulate your data.
	Quick Query or Parameterized Quick Query	A Quick Query is a saved query that can be rerun any time. Quick Queries are used to save your work so that you can use it later or share it with others. You can also add parameters to a Quick Query to create a more flexible Parameterized Quick Query (PQQ).
	Legacy Parameterized Quick Query	A legacy Parameterized Quick Query (PQQ) is a legacy Quick Query object from Insights Platform version 9 and prior that allows simple parameterization. Parameterization provides the ability to choose different inputs when running the saved query. You are now able to run a legacy PQQ in TRS. See Run a legacy Parameterized Quick Query on page 407 for more information. You may also open and resave a legacy PQQ as a version 10+ PQQ. See Edit and convert a legacy Parameterized Quick Query on page 402.
	Merged table	A merged table is the result of combining the rows of two or more tables or worksheets together.
	QuickApp	A QuickApp is an interactive application that provides a custom front-end interface to the 1010data Insights Platform.
	Saved Code	The saved code object represents saved, non-executable Macro Language XML code such as a <code><library></code> , <code><defblock></code> , and <code><defop></code> . For example, you might save a library to store multiple block statements for reuse. You can open and edit saved code objects in the Macro Language Workshop.

Permission icons

Descriptions of permission icons in the object browser.

If you own an object, or have permission to save or upload an object into a folder, permission icons appear next to the object in the browser.

Table 17: Permission icons

Icon	Name	Description
	Owner	Indicates that you own this object.
	Uploader	Indicates that you can save or upload objects into this folder.

My Data and Other People's Data

My Data is a folder where you may save your data, such as tables, queries, and reports. Other People's Data is a folder where you can access objects in other users' **My Data** folders that they have shared with you.

Each user has their own **My Data** folder, which can be found at the top of the object browser.

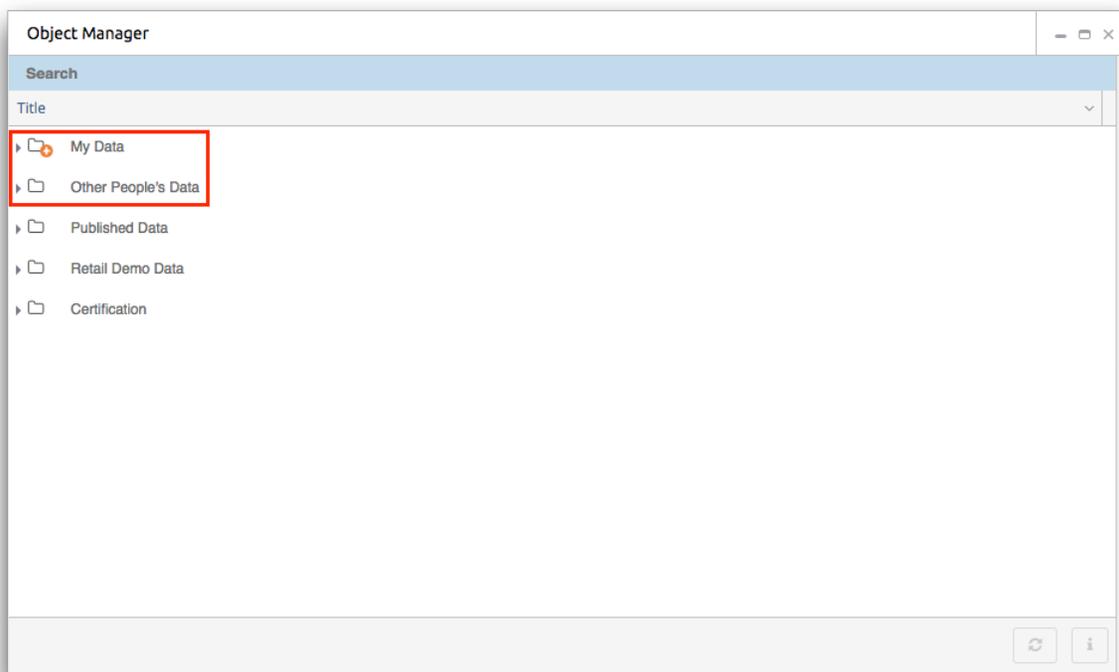


Figure 24: My Data and Other People's Data folders in the object browser

The name of this folder is `uploads`. This folder is intended for users to store temporary tables, worksheets, queries, and reports. The **My Data** folder has limited space; therefore, it is not suggested that users attempt to save large tables here.

If you share something in your **My Data** folder with someone else, they can access it through their own **Other People's Data** folder. Conversely, if someone shares something with you in their **My Data** folder, you can access it via your **Other People's Data** folder.

Note: You cannot create subfolders in **My Data**, which means you would not be able to use the inherit permissions feature that the 1010data Insights Platform provides. You would need to explicitly give permission to each user for those objects you want to share in your **My Data** folder.

In most cases, users will also have a user folder that their company creates for them where they can save their data and in which they can create subfolders. They can then share those subfolders with other users or groups, which is the method that 1010data recommends.

Metadata columns

Descriptions of metadata columns available in the object browser.

You can choose to view or hide specific columns of metadata within the object browser. For instructions, see [Choose the displayed metadata columns](#) on page 96. To choose the columns of metadata information that are displayed by default each time a new **Object Manager** window is opened, see [Select default metadata columns](#) on page 44.

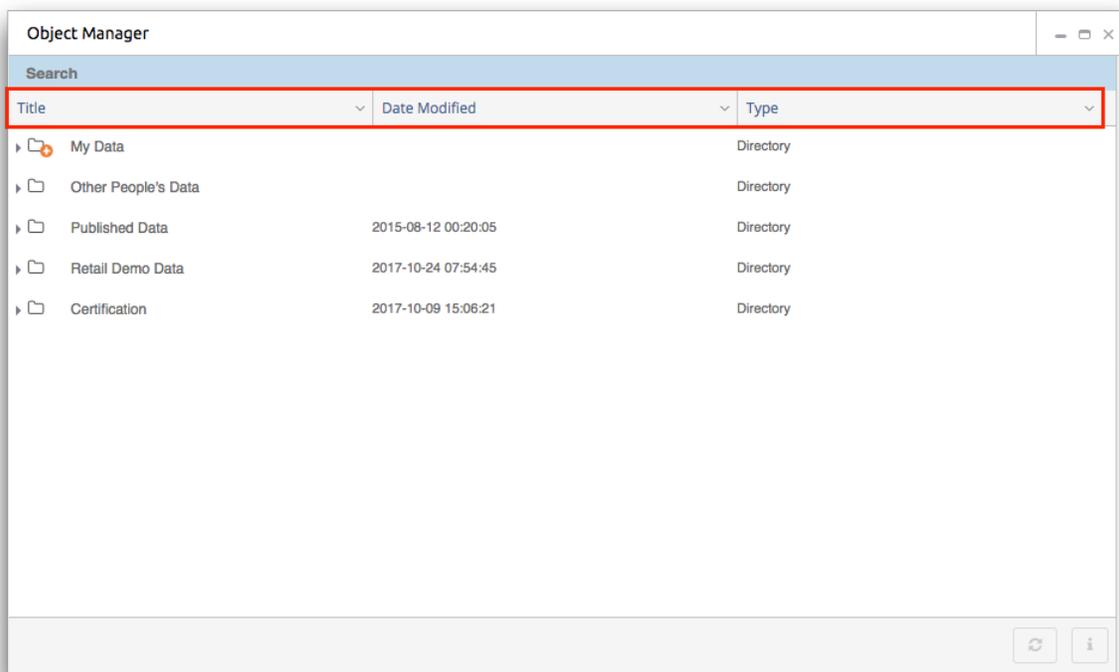


Figure 25: Metadata columns shown in the object browser

A list and description of available columns is provided below.

Title

This column displays the icon representing the type of object and the title of the object.

Note: The name and title of an object are different. To find the name of the object, use the information dialog. For more information, see [View object information](#) on page 107.

Date Modified

This column displays the date and time the object was last modified.

Type

This column indicates an object's type. For a list and description of object types, see [Object icons](#) on page 89.

Short Description

This column displays a concise description about the object.

Long Description

This column displays a more detailed explanation, which may include the Macro Language XML of the query.

Full Path

This column displays the full path of the object. The full path consists of the file path to the folder in which the saved object is stored and the name of the object itself.

Object Count

This column displays the number of objects within a folder.

Object actions by type

By default, the Object Manager opens an object in its designated tool. You also have the option of viewing and interacting with an object in other ways.

You can open an object in its designated 1010data Insights Platform tool or choose to open and interact with it in other ways. Additional actions available in the right-click menu of the object browser vary and are based upon the selected object. The table below provides a list and description of additional actions by object type.

Note: Some actions, such as deleting an object or adding a subfolder within a folder, are available only for objects you own. For more information, see [Permission icons](#) on page 90.

Table 18: Additional object actions by type

Object type	Available menu options
Folder	<p>Add Subfolder Creates a new folder within the selected folder.</p> <p>Delete Deletes the selected folder including any objects contained within.</p> <p>Open Folder in Other Pane Displays the contents of the folder in the other pane of the Object Manager.</p> <p>Show Object Info Opens the Information Dialog for the selected folder.</p> <p>Toggle favorites Adds the selected folder to your favorites in the Dashboard. If the selected folder is already a favorite, this option removes the folder from the Favorites list.</p>
Table	<p>Open in TRS Opens the selected table in the Trillion-Row Spreadsheet.</p> <p>Open in Macro Language Workshop Opens the selected table in the Macro Language Workshop.</p> <p>Edit in Table Editor Opens the selected table in the Table Editor.</p> <p>Delete Deletes the selected table.</p> <p>Show Object Info Opens the Information Dialog for the selected table.</p> <p>Toggle favorites Adds the selected table to your favorites in the Dashboard. If the selected table is already a favorite, this option removes the table from the Favorites list.</p>
Quick Query	<p>Open in TRS</p>

Object type	Available menu options
	<p>Opens the selected Quick Query/Parameterized Quick Query (PQQ) in the Trillion-Row Spreadsheet. The results of the query are used as the base table on which to base your analysis.</p> <p>Note: Parameterized timelines (including legacy PQQs) will now have all steps loaded into TRS when opened. Non-parameterized timelines will still be hidden initially. You can click Import Hidden Actions below the object browser in the Open table panel to view the entire timeline.</p> <p>Run query in TRS</p> <p>Opens the selected Quick Query/PQQ in the view only mode of the Trillion-Row Spreadsheet. This is the default action when you double-click a Quick Query in the Object Manager.</p> <p>If the query contains parameters (PQQ), parameter input fields and options are displayed. Otherwise, the results of the query are displayed in the grid.</p> <p>Run query in Query Runner</p> <p>Runs the selected Quick Query/PQQ in the Query Runner.</p> <p>Edit query in Macro Language Workshop</p> <p>Opens the selected Quick Query in the Macro Language Workshop. The Macro Language Workshop displays the individual operations of the query as Macro Language XML code.</p> <p>Edit query results in Table Editor</p> <p>Opens the results of the selected Quick Query in the Table Editor.</p> <p>Delete</p> <p>Deletes the selected Quick Query.</p> <p>Show Object Info</p> <p>Opens the Information Dialog for the selected Quick Query.</p> <p>Toggle favorites</p> <p>Adds the selected Quick Query/PQQ/legacy PQQ to your favorites in the Dashboard. If the selected Quick Query is already a favorite, this option removes the Quick Query from the Favorites list.</p>
Legacy Parameterized Quick Query	<p>Run query in TRS</p> <p>Opens the selected legacy Quick Query/PQQ in the view only mode of the Trillion-Row Spreadsheet. This is the default action when you double-click a legacy Quick Query in the Object Manager.</p> <p>If the query contains parameters (PQQ), parameter input fields and options are displayed. Otherwise, the results of the query are displayed in the grid.</p> <p>Open query in TRS</p> <p>Opens the selected legacy Parameterized Quick Query in the Trillion-Row Spreadsheet. The results of the query are used as the base table on which to base your analysis.</p>

Object type	Available menu options
	<p>Note: You are now able to view and edit the individual operations of a legacy Parameterized Quick Query. If you edit the legacy Quick Query, you can save it as a version 10+ Quick Query/PQQ.</p> <p>For more information, see Legacy Quick Queries/PQQs on page 402.</p> <p>Run Query in Query Runner</p> <p>Opens the selected legacy Quick Query as a QuickApp in the Converted QuickQuery window.</p> <p>Edit query in Macro Language Workshop</p> <p>Opens the selected legacy Quick Query in the Macro Language Workshop. The Macro Language Workshop displays the individual operations of the query as Macro Language XML code.</p> <p>Delete</p> <p>Deletes the selected legacy Parameterized Quick Query.</p> <p>Toggle favorites</p> <p>Adds the selected Quick Query/PQQ/legacy PQQ to your favorites in the Dashboard. If the selected Quick Query is already a favorite, this option removes the Quick Query from the Favorites list.</p>
Merged table	<p>Open in TRS</p> <p>Opens the selected merged table in the Trillion-Row Spreadsheet.</p> <p>Open in Macro Language Workshop</p> <p>Opens the selected merged table in the Macro Language Workshop.</p> <p>Edit in Table Editor</p> <p>Opens the selected merged table in the Table Editor.</p> <p>Delete</p> <p>Deletes the selected merged table.</p> <p>Show Object Info</p> <p>Opens the Information Dialog for the selected merged table.</p> <p>Toggle favorites</p> <p>Adds the selected merged table to your favorites in the Dashboard. If the selected merged table is already a favorite, this option removes the merged table from the Favorites list.</p>
QuickApp	<p>Open QuickApp in a Window</p> <p>Opens the selected QuickApp in its own window within the current workspace.</p> <p>Open QuickApp in a new browser tab</p> <p>Opens the selected QuickApp as a standalone QuickApp in a new browser tab.</p> <p>Edit code in Macro Language Workshop</p> <p>Displays the code for the selected QuickApp in the Macro Language Workshop.</p>

Object type	Available menu options
	<p>Delete Deletes the selected QuickApp.</p> <p>Show Object Info Opens the Information Dialog for the selected QuickApp.</p> <p>Toggle favorites Adds the selected QuickApp to your favorites in the Dashboard. If the selected QuickApp is already a favorite, this option removes the QuickApp from the Favorites list.</p>
Saved Code	<p>Edit code in Macro Language Workshop Displays the selected saved code in the Macro Language Workshop.</p> <p>Export to Workspace Exports the selected saved code and installs it into your current workspace.</p> <p>Delete Deletes the selected QuickApp.</p> <p>Show Object Info Opens the Information Dialog for the selected QuickApp.</p> <p>Toggle favorites Adds the selected saved code to your favorites in the Dashboard. If the selected saved code is already a favorite, this option removes the saved code from the Favorites list.</p>

Choose the displayed metadata columns

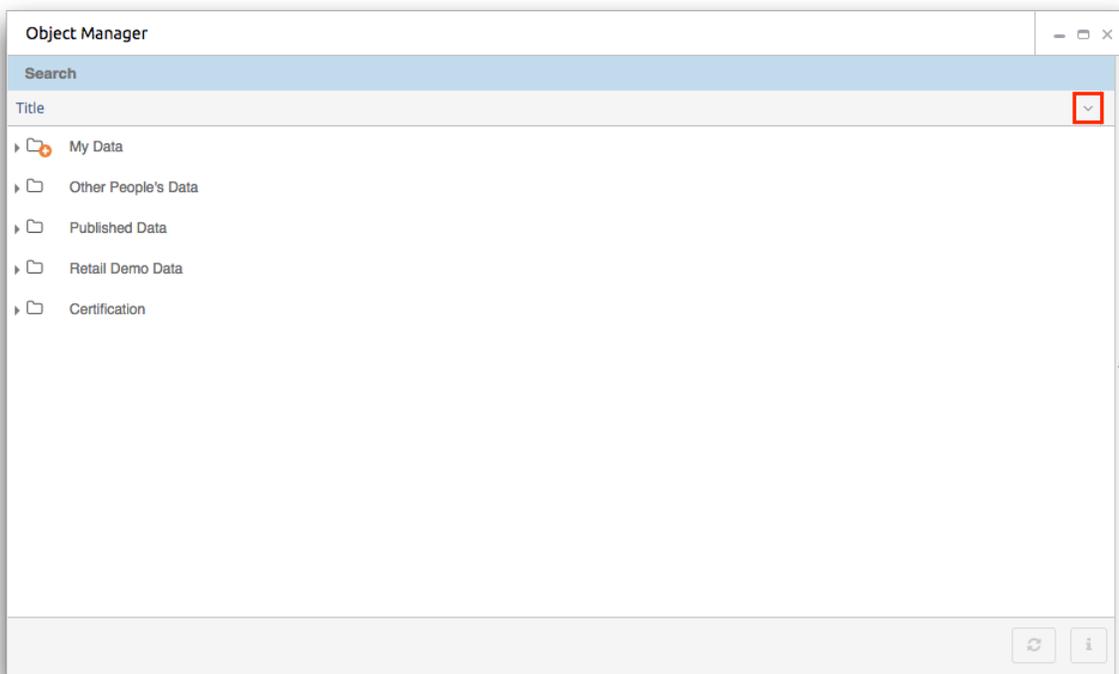
Add or remove various columns of metadata information displayed in the object browser.

You can choose the object metadata you want to see in the object browser. Information for each metadata item is displayed in its own column. This allows you to display only the metadata information that you need. In addition, you can sort the object list by each metadata column.

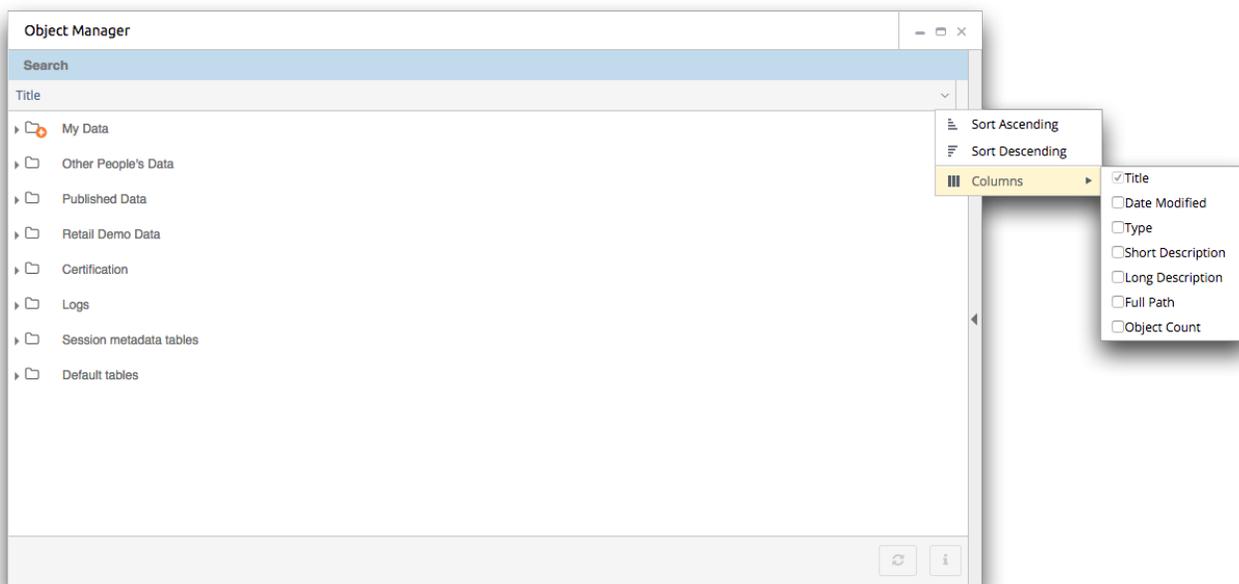
Note: These selections do not apply to other **Object Manager** windows and are not saved when the window is closed. To choose the columns of metadata information that are displayed by default each time a new **Object Manager** window is opened, see [Select default metadata columns](#) on page 44.

To choose the displayed metadata columns:

1. In the Object Manager, click the drop-down arrow within a column header.
For example, the drop-down arrow in the **Name** column is outlined in red below.



2. In the menu that appears, point to **Columns**.
The Object Manager displays the list of available metadata columns.



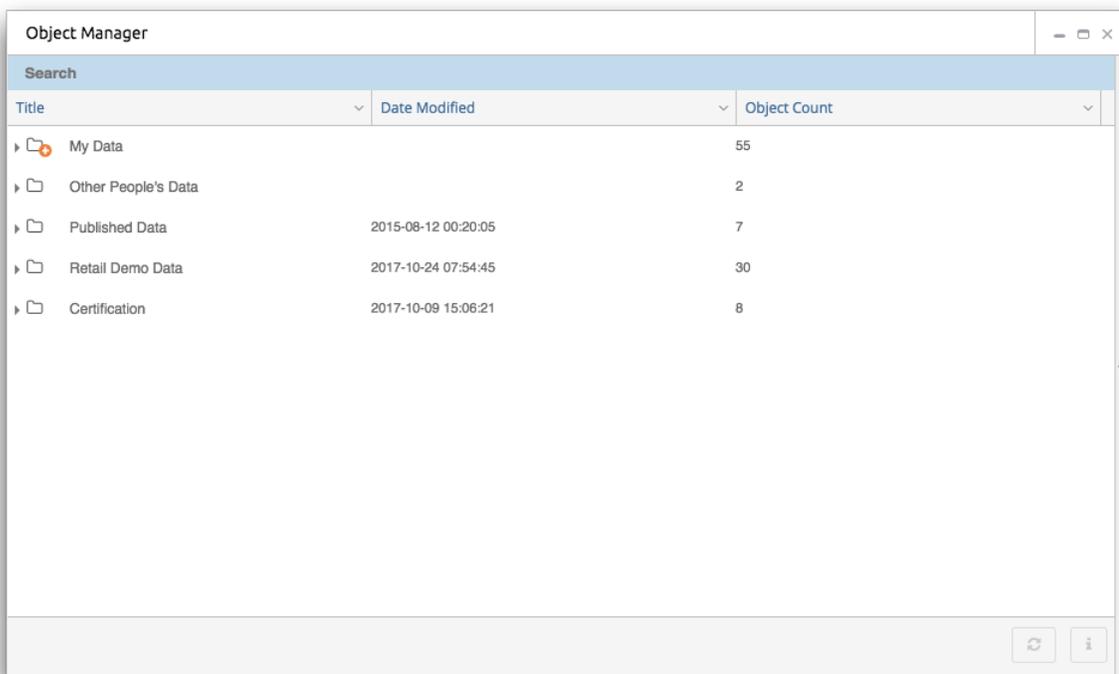
3. Choose the metadata columns you want displayed in the object browser by doing the following:

Option	Description
To display a column	Select the option next to its column label.
To hide a column	Clear the option next to its column label.

For a description of the available metadata columns, see [Metadata columns](#) on page 91.

Note: You cannot hide all metadata columns; at least one column must be selected.

The Object Manager displays or hides the metadata columns in the object browser.



Sort the object list

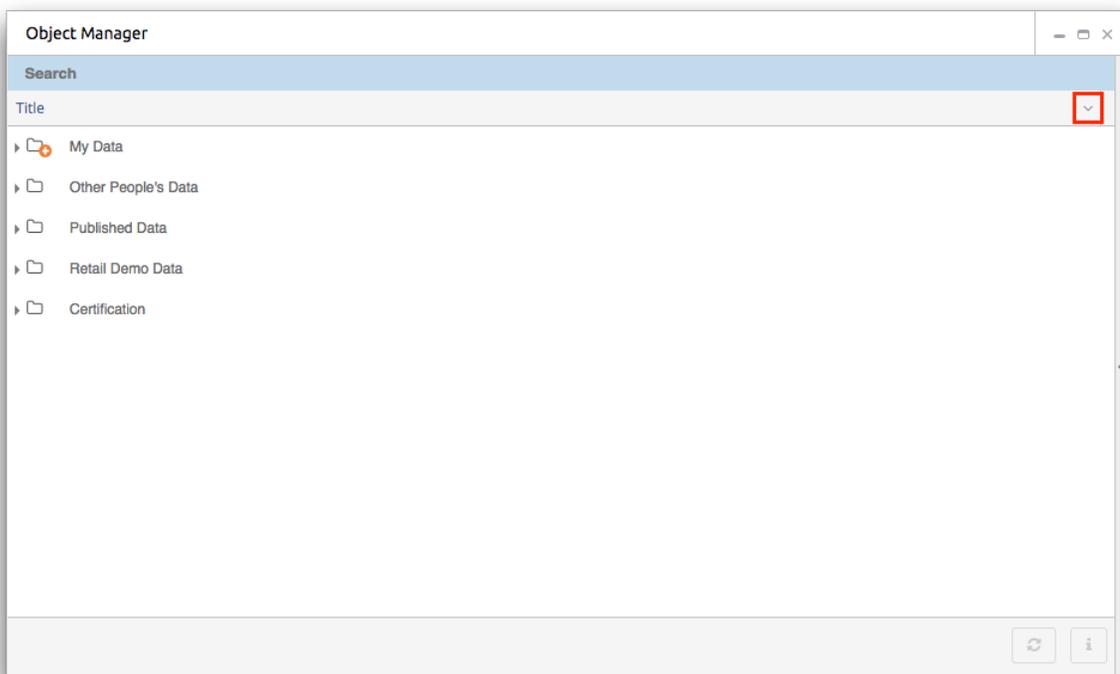
Sort the list of objects in the object browser in ascending or descending order.

You can reorder the list of objects in the object browser such that the values in a particular metadata column are in ascending or descending order. For example, you can sort the object list by name, date last modified, or type.

To sort the object list:

1. In the Object Manager, click the drop-down arrow within the metadata column by which you want to base the sort.

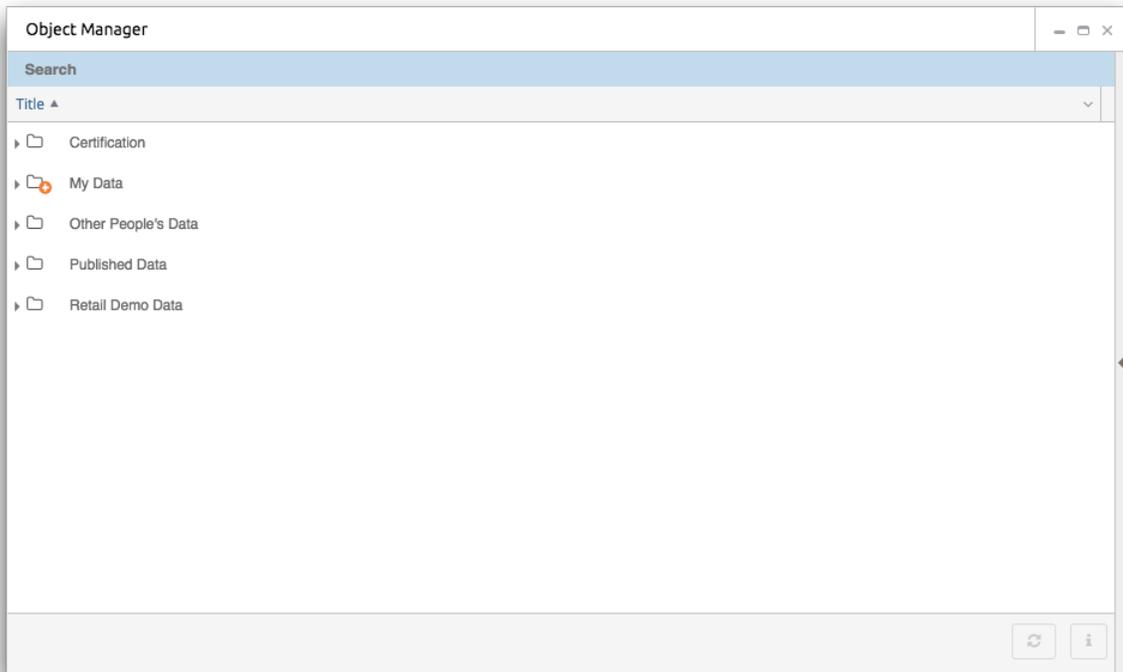
For example, to sort the list of objects by name, click the drop-down arrow in the **Name** column.



2. In the menu that appears, click one of the following:

Option	Description
Sort Ascending	Sorts the list of objects in ascending order based on the selected metadata column.
Sort Descending	Sorts the list of objects in descending order based on the selected metadata column.

The Object Manager sorts the list of objects based on the selected metadata column. In addition, an arrow icon appears within the column indicating the direction of the sort.



Note: To reverse the sort order, click the header of the sorted metadata column one time. To clear the sort order, click the header metadata column a second time.

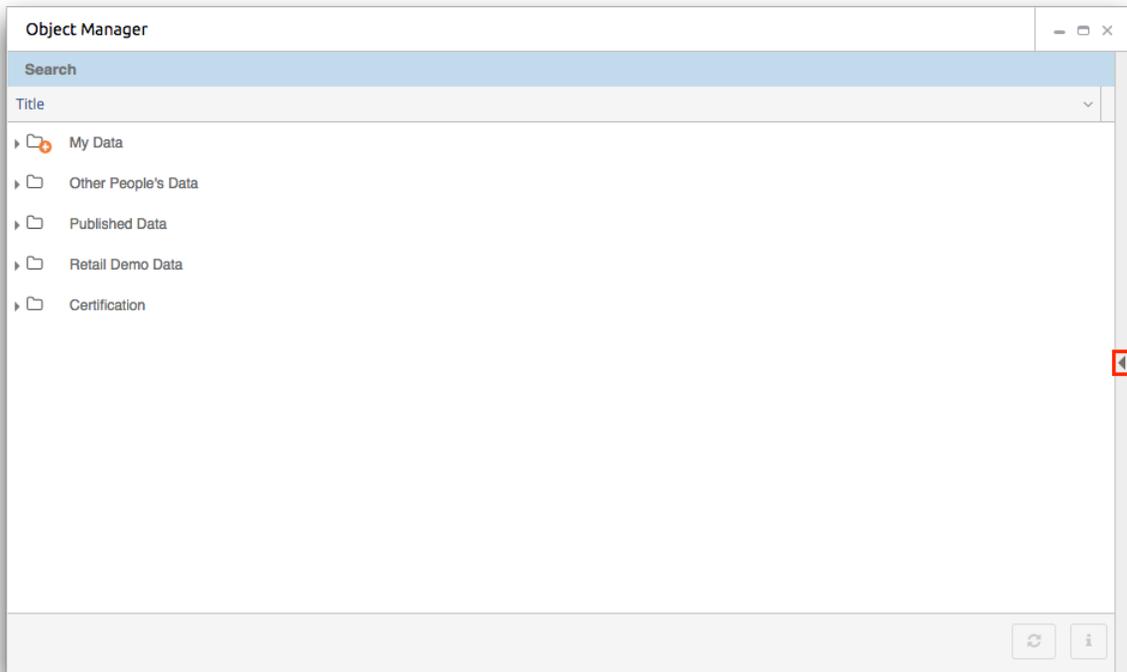
Open the right-side object browser pane

You can view both object browser panes at the same time.

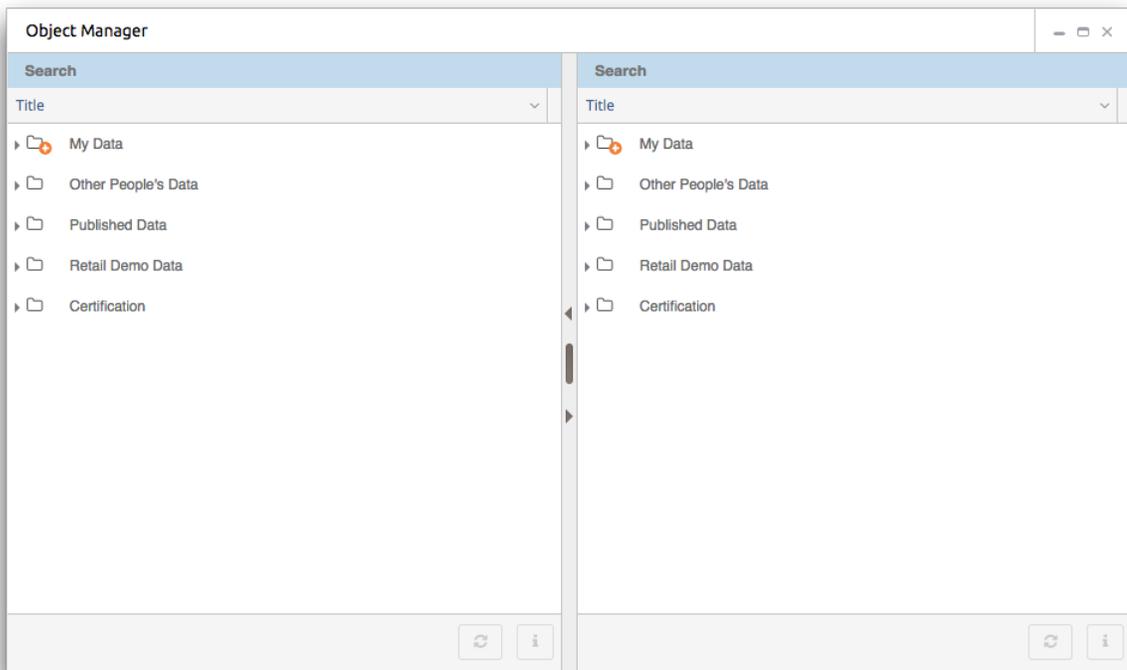
Viewing both the left- and right-side panes at the same time can help make some tasks in the Object Manager easier. For example, when moving an object from one folder to another, you can view each folder, side-by-side, in the two panes.

To open the right-side object browser pane:

In the Object Manager split bar, located on the right-side of the window, click the **Show** (📄) icon.



The Object Manager displays the right-side object browser pane.



When both object browser panes are visible, click the **Show** (◀) icon to hide the left-side pane or click the **Hide** (▶) icon to hide the right-side. Additionally, you can resize the panes using the split bar handle located between the **Show** and **Hide** icons.

Search

Use the search functionality in an object browser to find objects in the 1010data Insights Platform.

Although this topic explains the search functionality within the Object Manager, search is available throughout the 1010data Insights Platform anytime an object browser is displayed. While the functionality is the same, depending on the tool or window being used, the **Search** field may appear either above or below the object browser.

In the Object Manager, the **Search** field is located above the object browser.

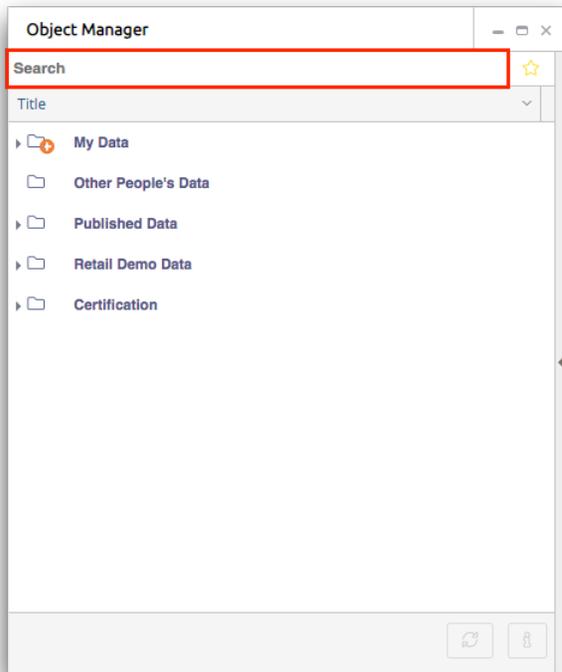


Figure 26: Search field in the Object Manager

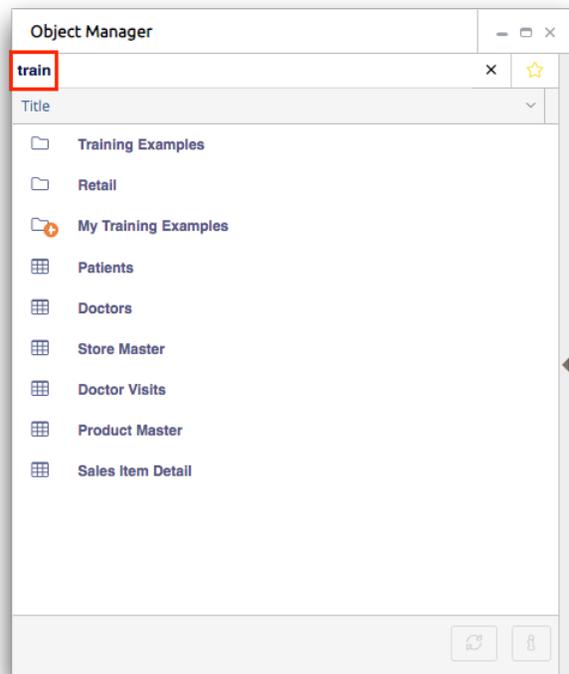
The Object Manager will display all of the folders, tables, queries, and reports whose titles, names, or paths match the text you entered in the **Search** field. For an explanation of the differences between titles, names, and paths, see [Information dialog](#) on page 104.

Note: The search is not case sensitive, and partial matches are listed.

See the example below for further details about search results.

Example

The below example shows the results of entering `train` in the **Search** field.



All of the tables and folders that contain `train` either in their names, titles, or paths are displayed in the object browser. The reasons are detailed below:

Training Examples	This folder is listed because both its title (<code>Training Examples</code>) and its name (<code>training</code>) contain the search string <code>train</code> .
Retail	This folder is listed because it is a subfolder of the Training Examples folder, and therefore its path (<code>training.retail</code>) contains the search string <code>train</code> , even though its title does not.
My Training Examples	This folder is listed because its title (<code>My Training Examples</code>) contains the search string <code>train</code> , even though its name and path do not (<code>retaildemo.randd.username.searchexample</code>).
Patients Doctors Doctor Visits	These tables are listed because they are found in the Training Examples folder, and therefore their paths contain the search string <code>train</code> (e.g., <code>training.patients</code> , <code>training.doctors</code> , <code>training.docvisits</code>), even though their titles and names do not.
Store Master Product Master Sales Item Detail	These tables are listed because they are found in the Retail folder, which is a subfolder of the Training Examples folder, and therefore their paths contain the search string <code>train</code> (e.g., <code>training.retail.store</code> , <code>training.retail.prod</code> , <code>training.retail.object</code>), even though their titles and names do not.

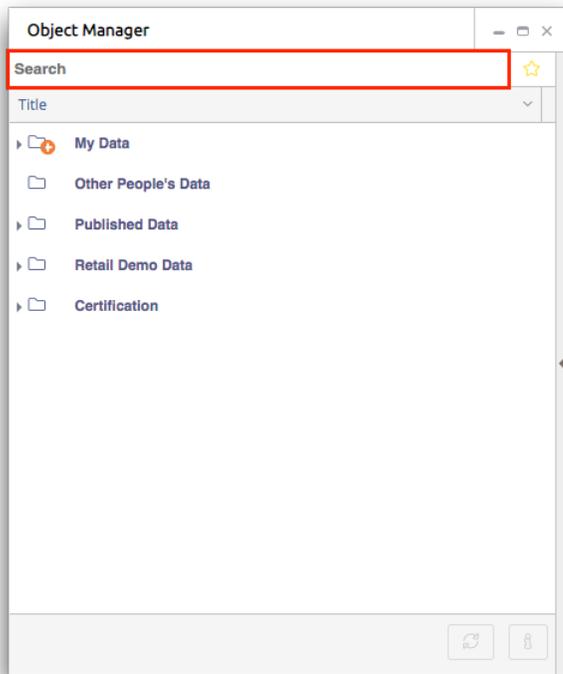
Search for objects

Search for objects such as tables, queries, and folders in the Object Manager.

Note: Although this topic explains how to search for objects within the Object Manager, search functionality is available throughout the 1010data Insights Platform any time an object browser is displayed. While the functionality is the same, depending on the tool or window being used, the **Search** field may appear either above or below the object browser.

To search for objects:

1. In the Object Manager **Search** field, enter the text for the objects you want to find.



2. Do one of the following:

- To automatically submit the search, stop typing for at least one second.
- To submit the search immediately, press **Enter**.

The Object Manager displays the objects with titles or names that match the text you entered.

Note: The search is not case sensitive, and partial matches are listed. For more information about search results, see [Search](#) on page 102.

3. To clear the search results, click the **Clear** (×) icon.

The icon is located in the **Search** field at the far right side.

Information dialog

Descriptions for fields and options in the information dialog of the Object Manager.

To access the information dialog, select an object in the Object Manager and then click the **Information** () icon. You can also right-click an object and select **Show Object Info** from the menu.

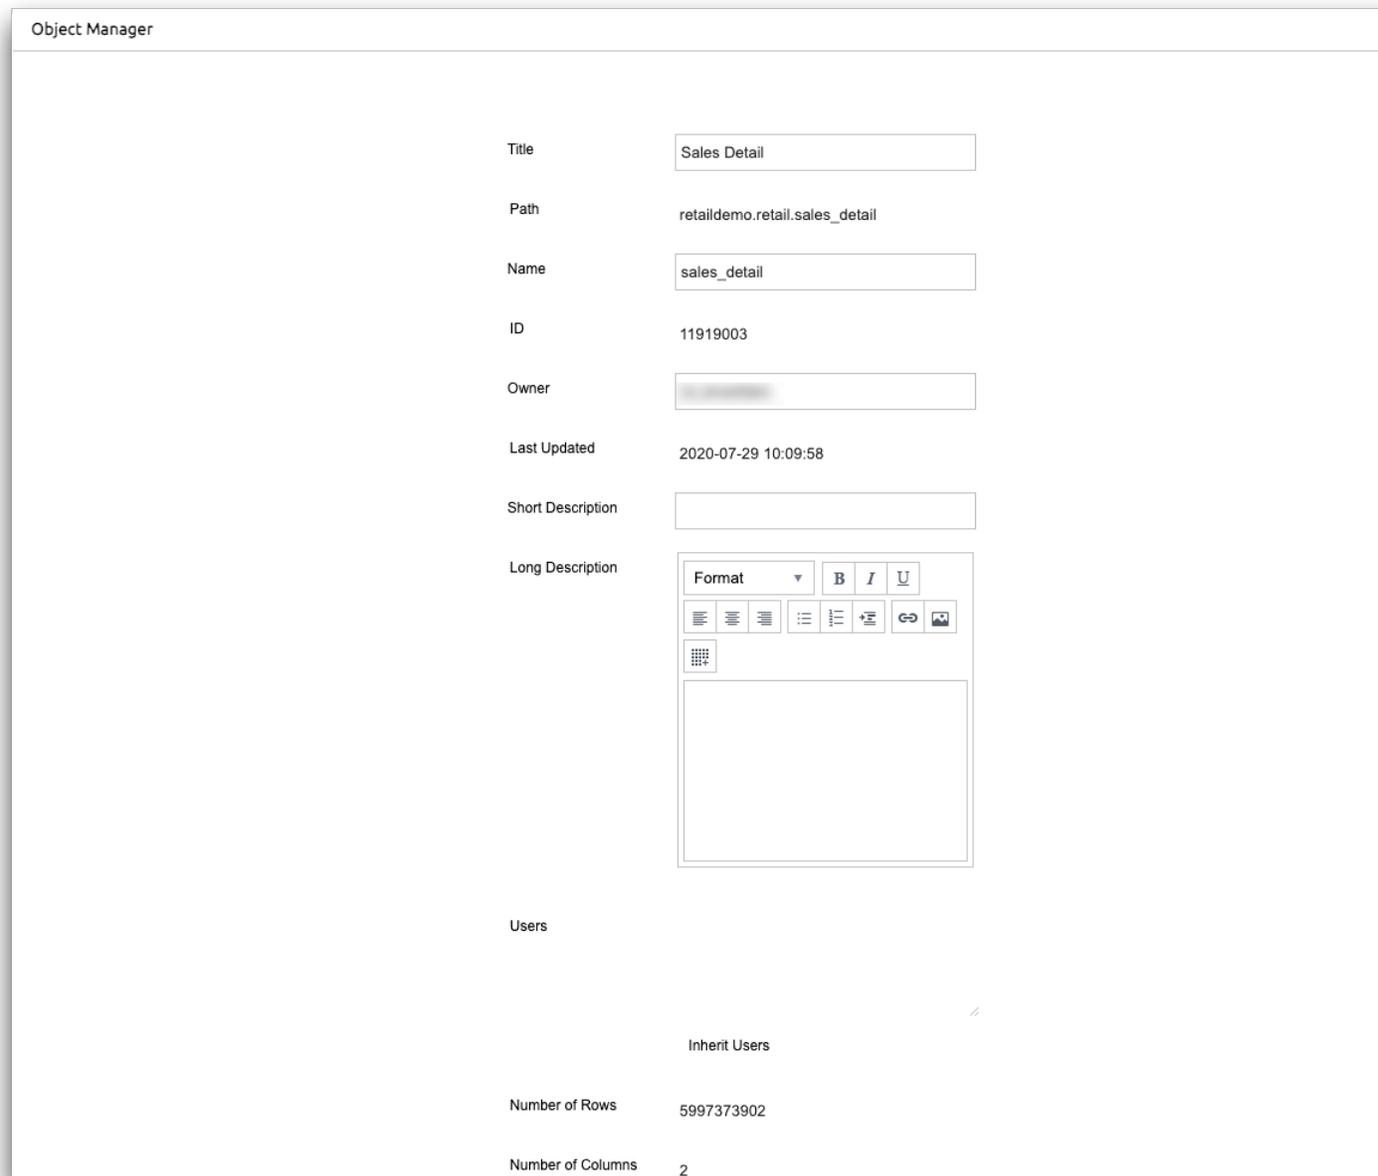


Figure 27: Information dialog in the Object Manager

Depending on the selected object, a subset of the following information is available.

Title

The title of the selected object.

The *title* of an object is different from its *name*. In the image above, the table *title* is **Sales Detail** and the table *name* is `salesdetail_customer`.

Path

The full path of the selected object.

The full path includes the path to, and the name of, the object. In the image above, the full path of the object is `certification.retail.salesdetail_customer`. The *path* portion is `certification.retail` and the *name* portion is `salesdetail_customer`.

Name

The name of the selected object.

The *name* of an object is different from its *title*. In the image above, the table *name* is `salesdetail_customer`, and the table *title* is **Sales Detail**.

Note: You cannot rename objects in the **My Data** folder.

ID

The internal identification number of the selected object.

Owner

The username of the selected object's owner.

Last Updated

The date and time the selected object was last updated.

Short Description

A concise description about the selected object.

Long Description

A more detailed explanation, which may include the Macro Language XML, of the selected object. The long description also includes formatting options, such as bulleted lists, links, and images.

Users

The list of users and groups (separated by a comma, space, or new line), who have access to the selected object.

Note: Users and groups listed in this field are ignored when the **Inherit Users** option is selected.

Inherit Users

When this option is selected, users who have access to the parent folder (the folder in which the selected object is located) will have access to the selected object. Users and groups listed in the **Users** field are ignored when this option is selected.

Note: You cannot use the inherit permissions feature for objects saved in your **My Data** folder. To share an object saved in your **My Data** folder, you must explicitly give permission to each user with whom you want to share the object.

Uploaders

The list of users and groups who have permission to upload to the selected folder.

Note: Users and groups listed in this field are ignored when the **Inherit Uploaders** option is selected.

Inherit Uploaders

When this option is selected, users who have upload permission to the parent folder (the folder in which the selected folder is located) will have upload permission to the selected folder. Users and groups listed in the **Uploaders** field are ignored when this option is selected.

Number of Rows

The number of rows in the selected object (when available).

Number of Columns

The number of columns in the selected object (when available).

View object information

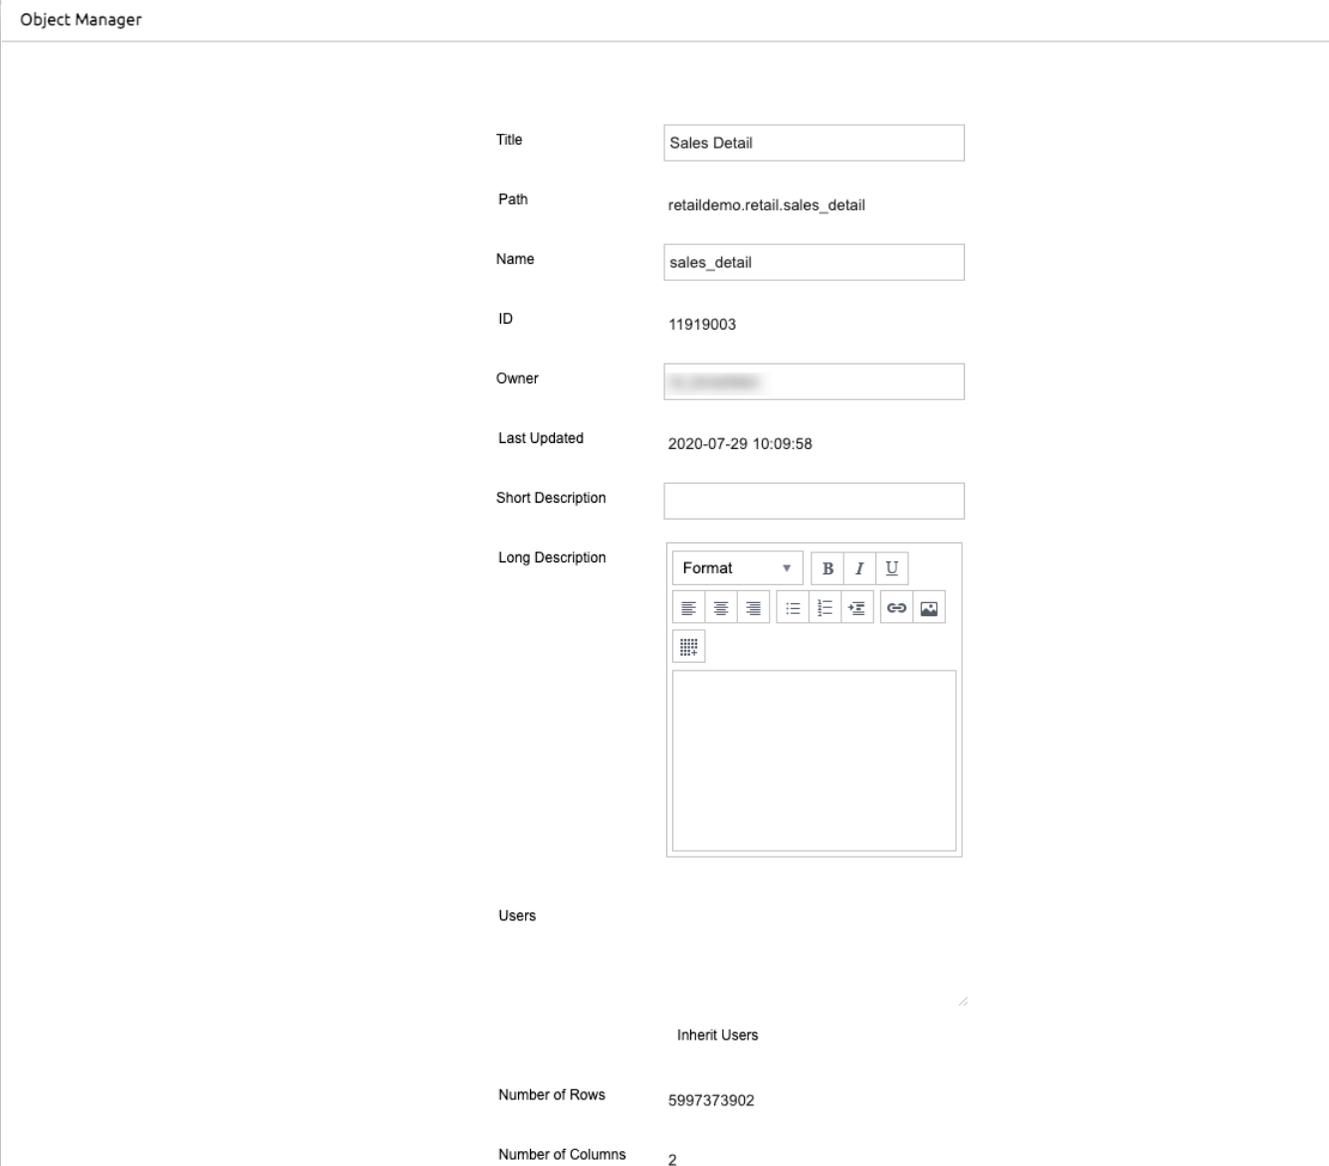
You can view the metadata information for objects such as tables, folders, queries, and reports.

Metadata information is available for 1010data Insights Platform objects in the Object Manager.

To view object information:

1. In the Object Manager, navigate to the desired object.
2. Do one of the following:
 - Click the object and then click the **Information** () icon.
 - Right-click the object and select **Show Object Info** from the menu.

The Object Manager displays the information dialog.



The screenshot shows the 'Object Manager' window with the following information displayed:

Title	Sales Detail
Path	retaildemo.retail.sales_detail
Name	sales_detail
ID	11919003
Owner	[Redacted]
Last Updated	2020-07-29 10:09:58
Short Description	[Empty text box]
Long Description	<div style="border: 1px solid #ccc; padding: 5px;"> <p>Format ▼ B I U</p> <p>          </p> <p></p> <div style="border: 1px solid #ccc; height: 100px; width: 100%;"></div> </div>
Users	[Empty list]
Inherit Users	<input type="checkbox"/>
Number of Rows	5997373902
Number of Columns	2

Relevant information for the selected object is displayed. For a list of possible fields and options, see [Information dialog](#) on page 104.

3. To close the information dialog, do one of the following:

- Click the **Close dialog** () icon.
- Click **Cancel**.

Edit object information

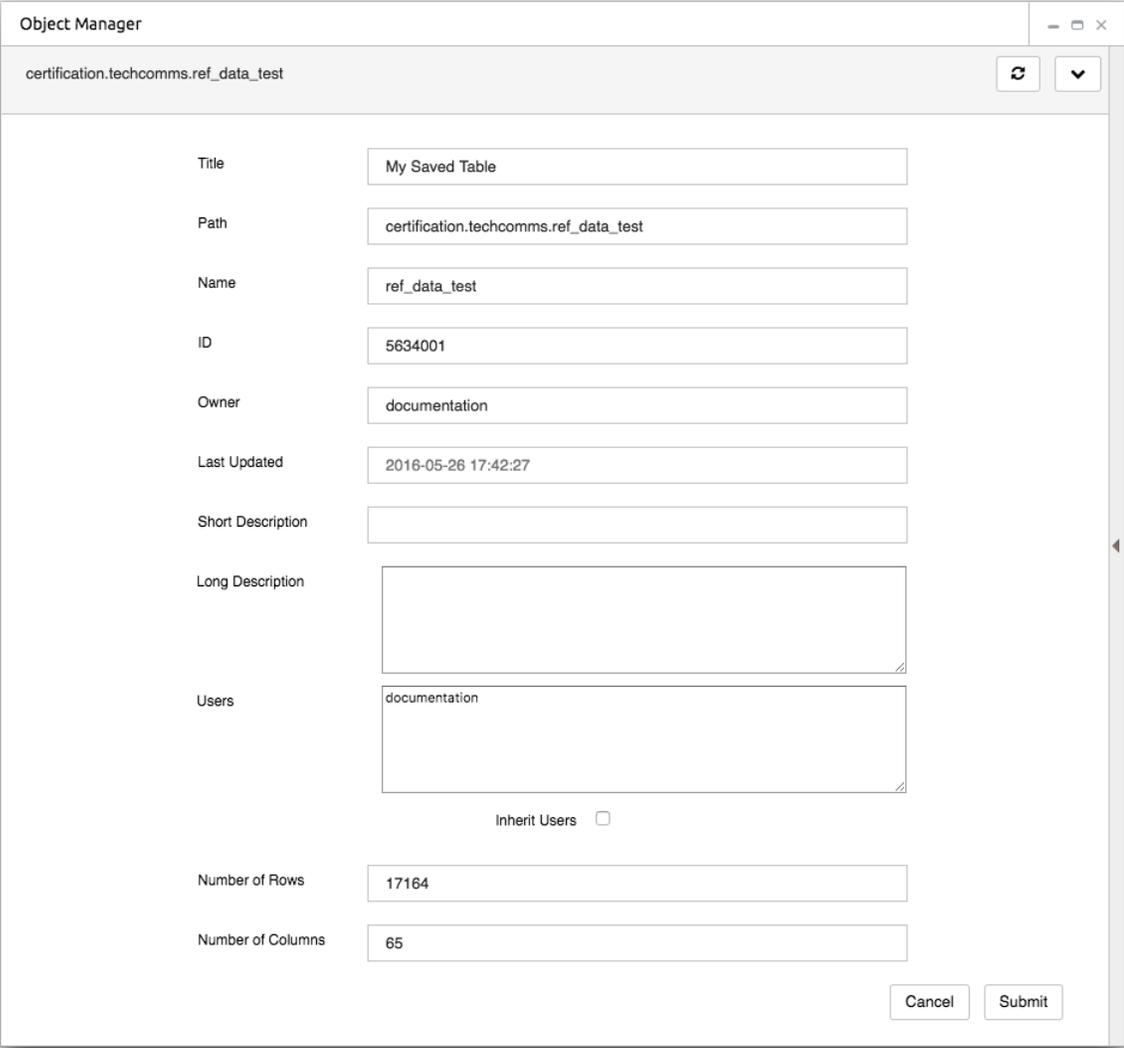
You can edit the meta information for a folder, table, query, or any other object you own.

You can edit information for 1010data Insights Platform objects that you own. Objects you own have the **Owner** () icon next to them in the object browser.

To edit object information:

1. In the Object Manager, navigate to the object you want to edit.
2. Do one of the following:
 - Click the object and then click the **Information** () icon.
 - Right-click the object and select **Show Object Info** from the menu.

The Object Manager displays the information dialog.



The screenshot shows the 'Object Manager' window with the 'Show Object Info' dialog open for the object 'certification.techcomms.ref_data_test'. The dialog contains the following fields and controls:

- Title:** My Saved Table
- Path:** certification.techcomms.ref_data_test
- Name:** ref_data_test
- ID:** 5634001
- Owner:** documentation
- Last Updated:** 2016-05-26 17:42:27
- Short Description:** (empty text box)
- Long Description:** (empty text area)
- Users:** documentation
- Inherit Users:**
- Number of Rows:** 17164
- Number of Columns:** 65

At the bottom right of the dialog are 'Cancel' and 'Submit' buttons.

3. Change the desired information.

Relevant information for the selected object is displayed. For a list of possible fields and options, see [Information dialog](#) on page 104.

4. Click **Submit**.

The Object Manager saves your changes and closes the information dialog.

Share an object

You can share objects that you own with others, giving them permission to see your tables or run your queries.

If you need to share objects with multiple users, consider creating a group.

To share an object:

1. In the Object Manager, navigate to the object you want to share.

Note: You can only share objects that have an **Owner** icon () next to them in the object browser, which indicates that you own those objects.

2. Do one of the following:

- Click the object and then click the **Information** () icon.
- Right-click the object and select **Show Object Info** from the menu.

The Object Manager displays the information dialog.

The screenshot shows the 'Object Manager' dialog box with the following fields and values:

Title	My Saved Table
Path	certification.techcomms.ref_data_test
Name	ref_data_test
ID	5634001
Owner	documentation
Last Updated	2016-05-26 17:42:27
Short Description	
Long Description	
Users	documentation
Inherit Users	<input type="checkbox"/>
Number of Rows	17164
Number of Columns	65

Buttons: Cancel, Submit

3. In the **Users** field, add the names of the 1010data Insights Platform users or groups (separated by a comma, space, or new line) with whom you want to share the selected object.

If you select **Inherit Users**, the Insights Platform users that have access to the parent folder (the folder this object is in) will have access to this object.

Note: **Inherit Users** takes precedence over the **Users** list for this object.

4. Click **Submit**.

The Object Manager updates the permissions and closes the information dialog.

Add an object to your favorites

If you often use a specific object, you may want to add it to your favorites for quick and easy access.

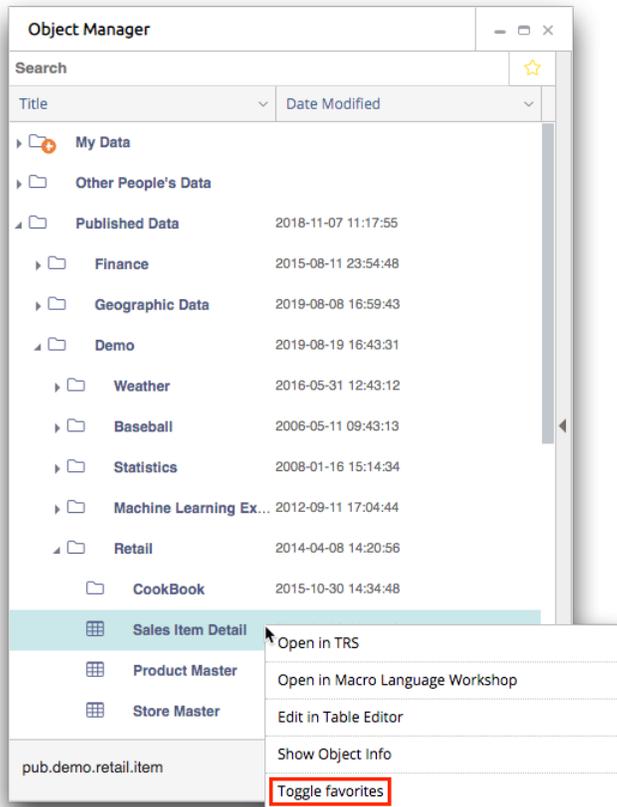
When you add a 1010data Insights Platform object to your favorites, it is available from the **Favorites** tab in the **Dashboard** window for easy access.

To add an object to your favorites:

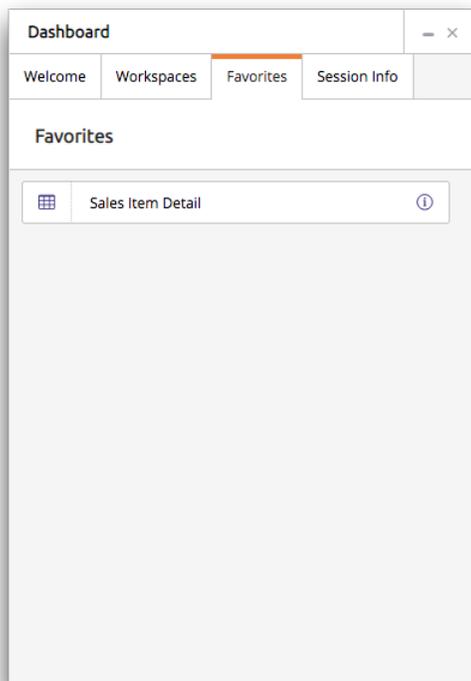
1. In the Object Manager, select one or more objects.

For instructions on selecting multiple objects, see [Select multiple objects](#) on page 116.

2. Right-click the object you want to add and then select **Toggle Favorites** from the menu.



The Insights Platform adds the object to the **Favorites** tab in the **Dashboard**.



After you add an object to your favorites, you can find it quickly and easily by selecting **Session > Favorites** from the workspace menu. For more information, see [Favorites](#) on page 72.

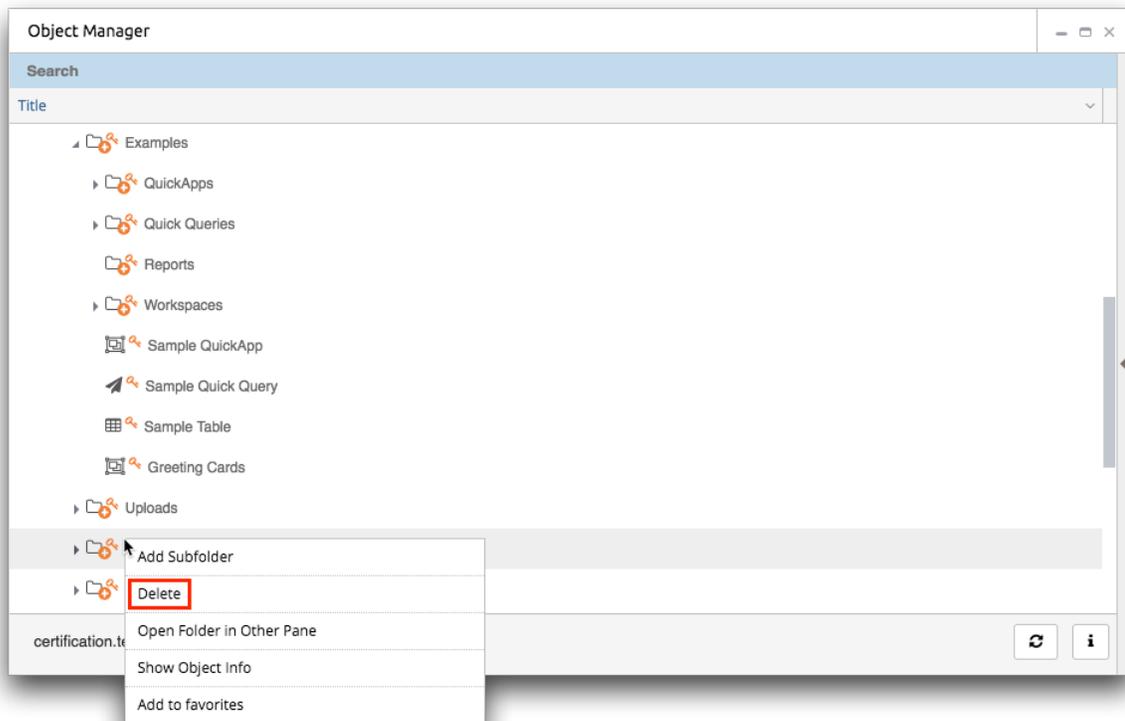
Delete an object

You can delete a folder, table, query, or any other object you own.

You can delete 1010data Insights Platform objects that you own. Objects you own have the **Owner** (🔑) icon next to them in the object browser.

To delete an object:

1. In the Object Manager, select one or more objects.
For instructions on selecting multiple objects, see [Select multiple objects](#) on page 116.
2. Right-click the object you want to delete and then select **Delete** from the menu.



The Object Manager displays a dialog confirming that you want to delete the object.

3. Click **OK**.
The Object Manager deletes the object.

Move an object

You can move a folder, table, query, or any other object you own to a different folder.

You can move 1010data Insights Platform objects that you own to another folder you own or have permission to modify. Objects you own have the **Owner** (🔑) icon next to them in the object browser. Folders you have permission to modify have the **Uploader** (👤) icon next to them in the object browser.

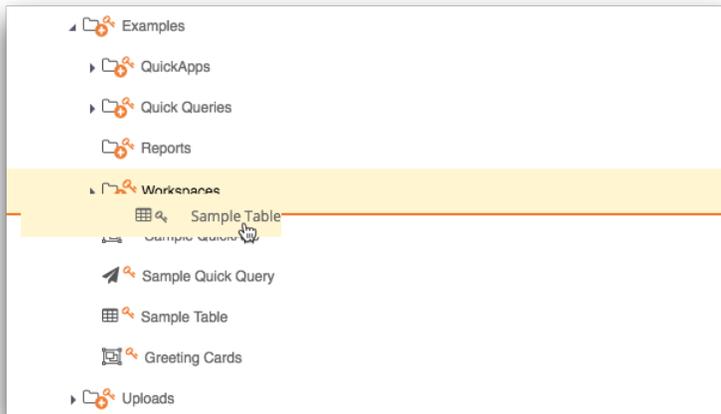
Note: You cannot move objects into the **My Data** folder.

To move an object:

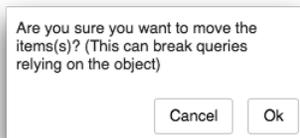
1. In the Object Manager, select one or more objects.
For instructions on selecting multiple objects, see [Select multiple objects](#) on page 116.

2. Drag the object to the desired folder.

As you drag an object, the Object Manager highlights the folder into which the object will be moved. In the image below, **Sample Table** will be moved into the **Workspaces** folder.



The Object Manager displays a dialog confirming that you want to move the selected object.



3. Click **OK**.

The Object Manager moves the object.

Reorder objects in a folder

Reorder the contents of a folder in the 1010data Insights Platform.

You can reorder objects within a folder that you own. Folders you own have the **Owner** (🔑) icon next to them in the object browser.

Note: You cannot reorder objects in the **My Data** folder.

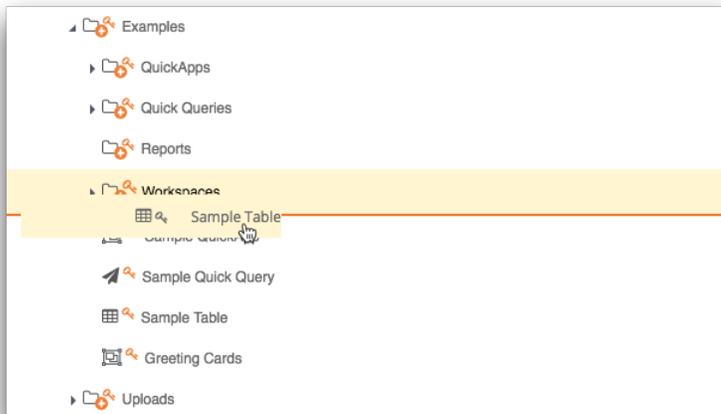
To reorder objects in a folder:

1. In the Object Manager, select one or more objects.

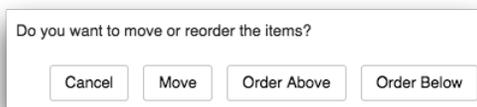
For instructions on selecting multiple objects, see [Select multiple objects](#) on page 116.

2. Drag the object to the desired location within the folder.

As you drag an object, the Object Manager displays an orange line indicating where the object will be placed. In the image below, **Sample Table** will be reordered within the **Examples** folder between **Workspaces** and **Sample QuickApp**.



The Object Manager displays a dialog confirming that you want to reorder the object.



3. Click one of the following:

Option	Description
Order Above	Reorders the object so that it is above the highlighted object in the object manager.
Order Below	Reorders the object so that it is below the highlighted object in the object manager.

The Object Manager saves the new object order.

Add a subfolder

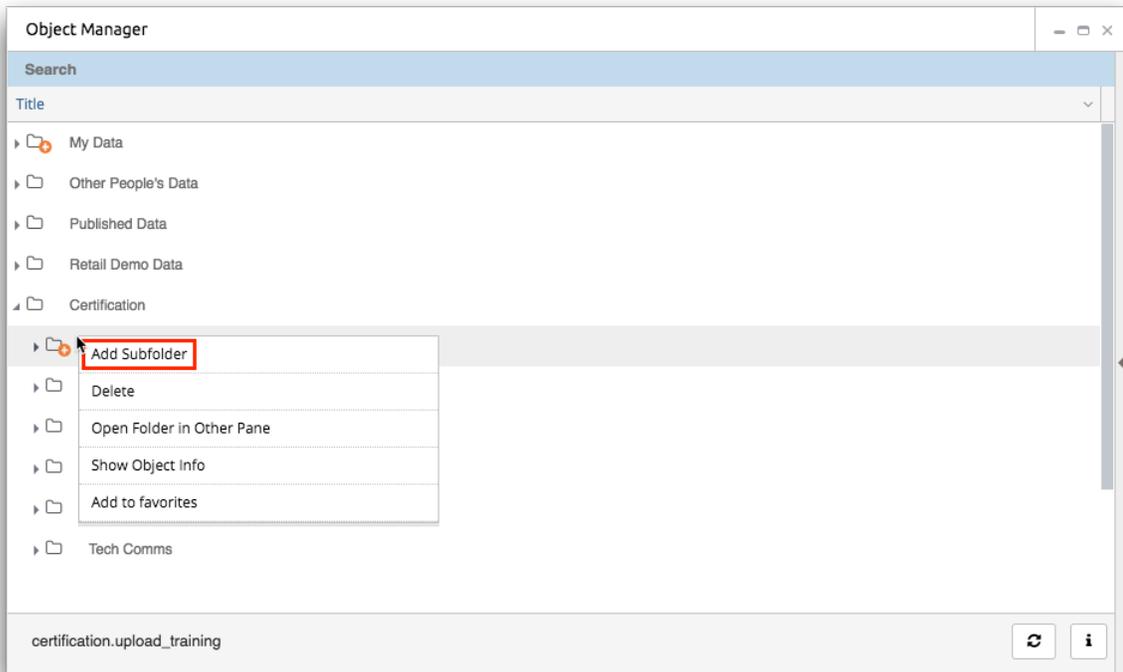
Add a subfolder to a folder.

You can add a subfolder to a folder you own or have permission to modify in the 1010data Insights Platform. Folders you own have the **Owner** (🔑) icon next to them in the object browser. Folders you have permission to modify have the **Uploader** (⊕) icon next to them in the object browser.

Note: You cannot add a subfolder to the **My Data** folder.

To add a subfolder:

1. In the Object Manager, navigate to the folder in which you want to add a subfolder.
2. Right-click the folder and then select **Add Subfolder** from the menu.



The Object Manager displays a dialog used to define the subfolder details.

To add a new subfolder to **certification.upload_training** select a title and name. *Names must be alphanumeric, they are the ending of the path to the object.*

Title

Name

3. In the **Title** field, enter the title of the subfolder.
The title is used to help describe the contents of the subfolder (e.g., *Shared Queries and Tables*). The title may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters. If you leave this field blank, a default title will be used (e.g., *New Folder*).
4. In the **Name** field, enter the subfolder name.
The subfolder name must begin with a letter and can only contain numbers, letters, and underscores. It cannot contain any spaces or other special characters. If you leave this field blank, a system-generated name will be used (e.g., *t662518159_yourusername*). The path to the parent folder will be automatically prepended to the **Full Path**.
5. Click **Submit**.
The Object Manager adds the subfolder.

Refresh an object

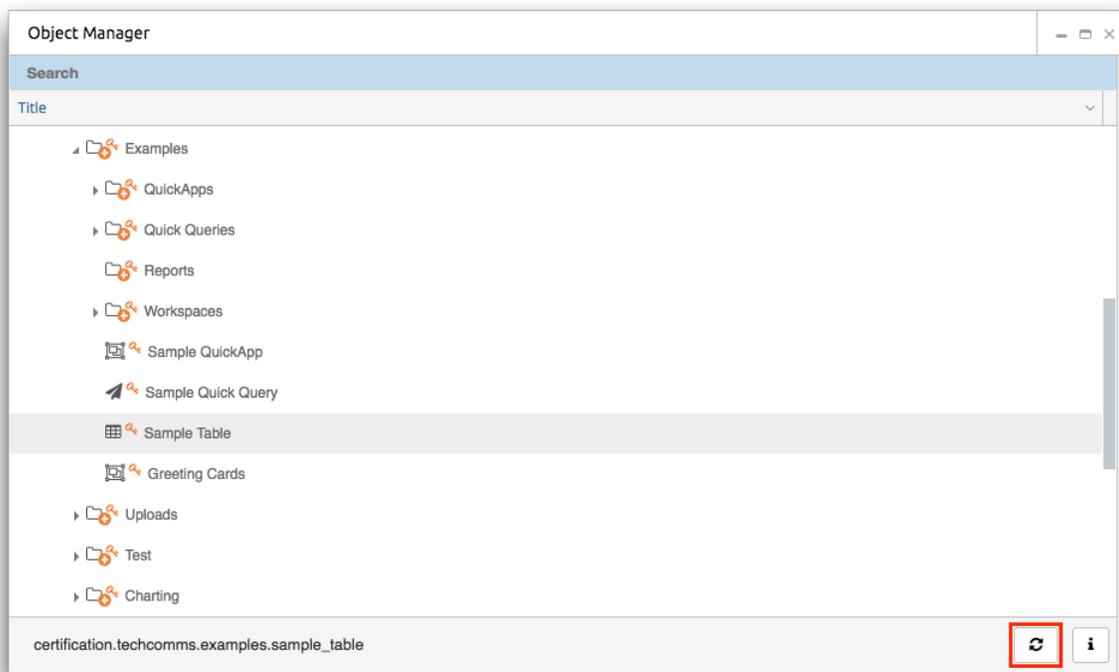
Update an object in your session with data from the database.

Because the 1010data Insights Platform is session-based, when an object is updated in the database, the object is not automatically updated in your session. You can manually refresh an object in your session to update it with the changes to the data in the database.

Note: It is important to note that the refresh functionality does not act recursively through a folder; refresh only updates child objects. If a folder contains other folders, the child folders are not updated when a refresh is performed on the parent. Child objects, however, will be updated.

To refresh an object:

1. In the Object Manager, select one or more objects.
For instructions on selecting multiple objects, see [Select multiple objects](#) on page 116.
2. Click the **Refresh** (↻) icon located in the lower right corner of the window.



The Object Manager refreshes the selected object with any changes from the database.

Select multiple objects

You can select multiple objects in the Object Manager and then perform common tasks, such as moving or deleting, to all of the selected objects at the same time.

There may be times when you want to perform an action on multiple objects at the same time. For example, you might want to move a group of objects to another folder or delete several unneeded objects at once. Using standard operating system selection methods, you can select multiple objects in the Object Manager.

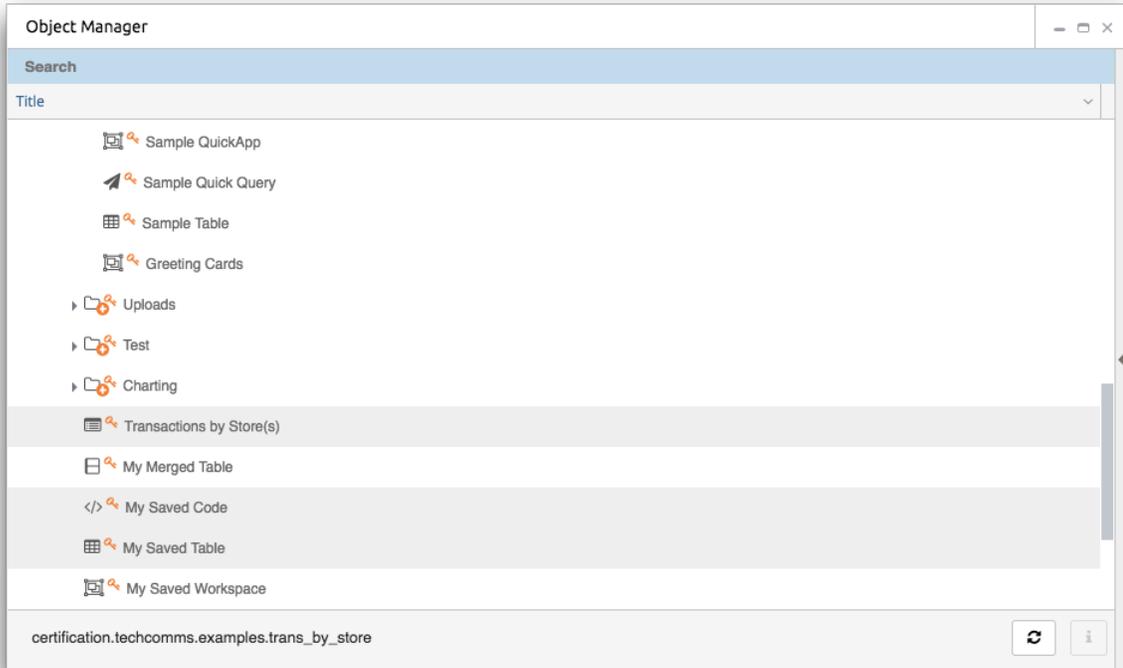
To select multiple objects:

In the Object Manager, select the desired objects by doing one or more of the following:

Task	Action
Select non-consecutive objects	Press and hold the Ctrl (PC) or Control (Mac) key, then click the desired objects.

Task	Action
Select a consecutive group of objects	Click the first object, press and hold the Shift key, and then click the last object.
Exclude objects from your selection	Press and hold the Ctrl (PC) or Control (Mac) key, then click the objects.

Selected objects are highlighted in gray.



After you select the desired objects, you can either right-click a highlighted object and select an available option from the menu or drag the objects.

Open an object

You can open an object in its designated 1010data Insights Platform tool or choose to open and interact with it in other ways.

The Insights Platform provides many options for viewing and interacting with objects. By default, the Object Manager opens an object in its designated tool. You also have the option of viewing and interacting with an object in other ways. For example, tables may be configured to open in the Trillion-Row Spreadsheet by default. However, as needed, you could choose to open a table in the Macro Language Workshop instead.

To open an object:

1. In the Object Manager, navigate to the object you want to open.
2. Do one of the following:
 - Double-click anywhere in that object's row.

The Object Manager will open the object in the Insights Platform tool designated in your association settings. For more information, see [Associations](#) on page 39.

Note: If the object does not open after you double-click it, the object association for the object type may not be defined. To set the default object association, see [Define object associations](#) on page 43.

- Right-click the object and select the appropriate option from the menu.

Options in the menu vary and are based on the object type. For a list of available menu options, see [Object actions by type](#) on page 93.

The Insights Platform opens the object in the appropriate tool.

Trillion-Row Spreadsheet

The Trillion-Row Spreadsheet (TRS) allows you to visually interact with your data.

The 1010data Insights Platform gives you several options to interact with the system and analyze your data, but when you first get started, you are likely going to spend most of your time in the Trillion-Row Spreadsheet. The TRS is a browser-based web interface that works similarly to many spreadsheet applications, such as Microsoft Excel. The interface uses many elements you may already be familiar with, such as folders, navigation icons, and menus.

The Trillion-Row Spreadsheet is accessed by selecting **Tools > Trillion-Row Spreadsheet** from the workspace menu.

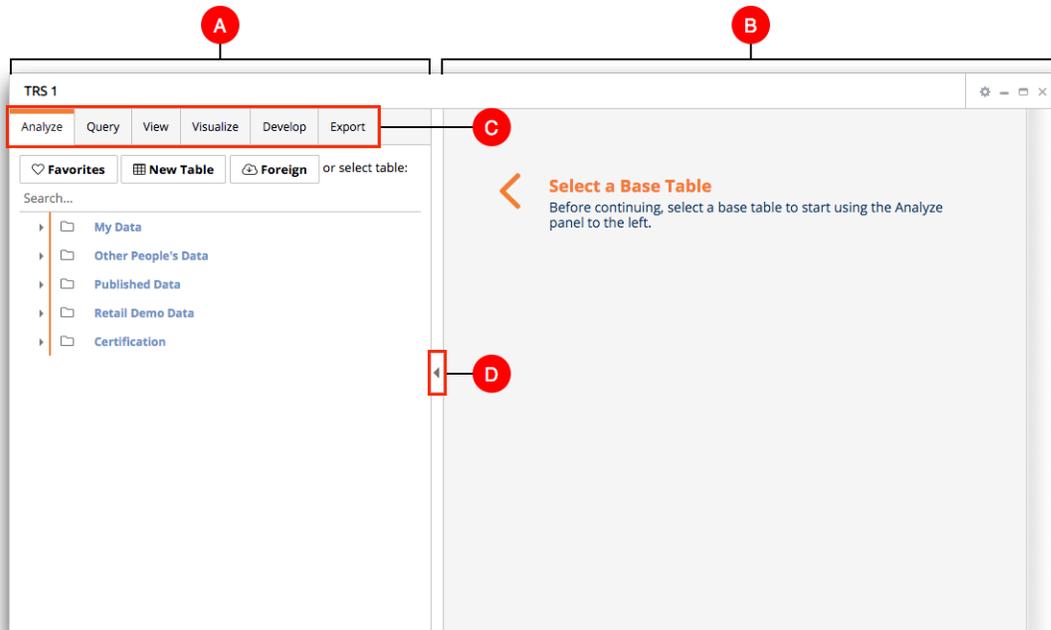


Figure 28: TRS window

A. Analysis pane

The analysis pane occupies the left side of the **TRS** window. This pane is used to perform an analysis on your data.

B. Results pane

The results pane is displayed on the right side of the **TRS** window. This pane displays the results of your analysis in spreadsheet format. In addition, this pane also displays charts created in the **Visualize** tab.

C. Tab bar

The **Tab bar** contains the options used to perform an analysis and other actions on your data.

D. Hide/Show

To hide the analysis pane, click the **Hide** (◀) icon on the right side of the pane. This is helpful when you need more space to view your data in the results pane. Click the **Show** (▶) icon to reveal the analysis pane when it is hidden.

Table 19: Tab bar

Tab	Description
Analyze	In a new TRS window, this tab displays options for opening objects such as tables and queries. After an object is opened, the Analyze tab displays the Analysis Timeline which provides quick access to perform operations such as summaries, tabulations, and cross tabulations.
Query	The Query tab displays the current timeline query as Macro Language XML code. It also provides the ability to undo and redo operations in the timeline.
View	Use this tab to select from among the various ways to view and interact with the data resulting from your analysis.
Visualize	Create charts based on your data analysis.
Develop	This tab allows you to save your query and provides options to further develop your analysis. For example, you can clone the Trillion-Row Spreadsheet and explore different scenarios without losing your original query.
Export	Export options to save the results of your analysis in a variety of file types such as CSV and Microsoft Excel.

Analyze

The **Analyze** tab displays the object browser and the Analysis Timeline.

In a new Trillion-Row Spreadsheet (TRS) window, the **Analyze** tab displays options for opening objects such as tables and queries. After an object is opened, the **Analyze** tab displays the Analysis Timeline which provides quick access to perform operations such as summaries, tabulations, and cross tabulations.

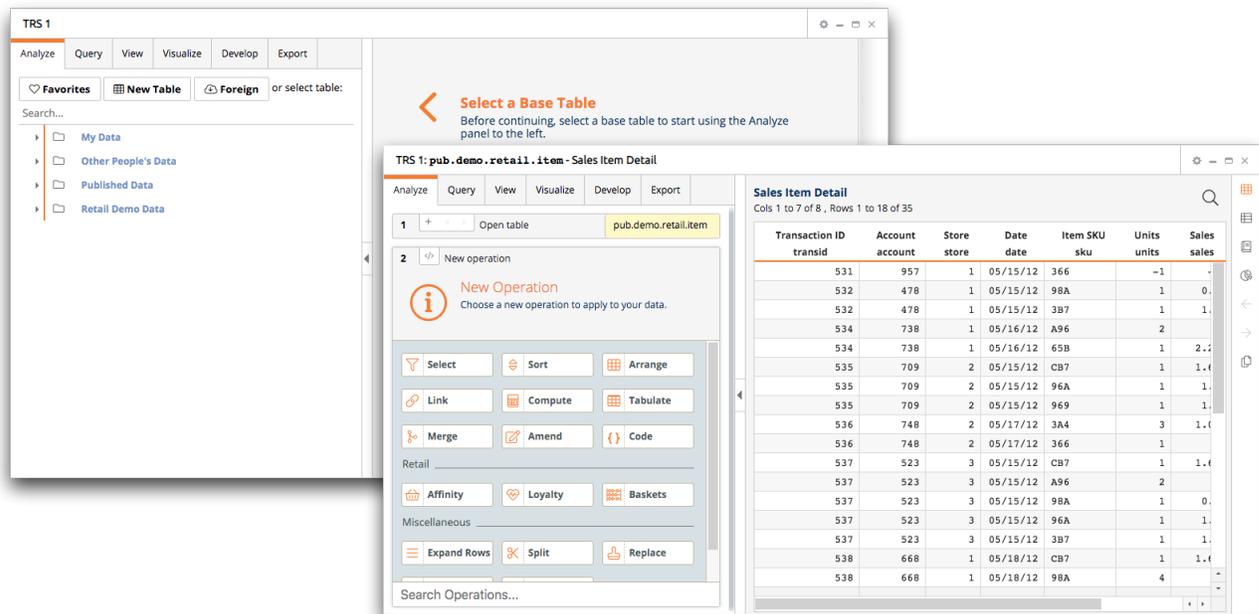


Figure 29: Analyze tab

Object browser

The object browser provides a hierarchy of folders and objects to which you have access. Use the object browser to select the base table or query for your analysis.

When a new **TRS** window is opened, the object browser is displayed in the **Analyze** tab.

The object browser is used to browse for and select the table or query on which to base your data analysis. A search feature is provided to help you locate existing folders, tables, and other objects.

Additional options allow you to choose a table or query from your list of favorites, create a new temporary table, or select a table from a foreign database.

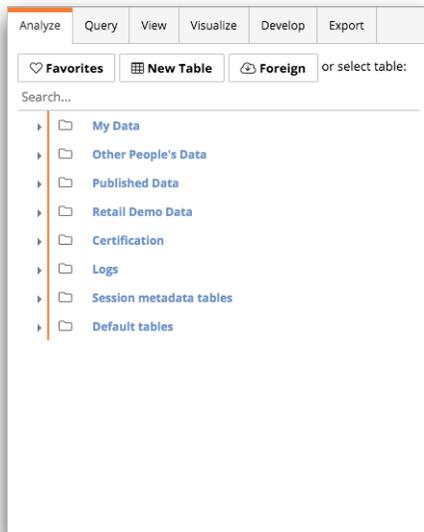


Figure 30: Object browser

Favorites

The **Favorites** button displays a list of your favorite tables and queries. If you have not added any tables or queries to your favorites list, this button is not available.

New Table

The **New Table** button opens the `<table>` panel which allows you to create a new temporary table using Macro Language code.

Foreign

The **Foreign** button opens the **Foreign database** panel which allows you to select a table from a connected database outside of the 1010data Insights Platform.

Search

The search field helps you locate objects such as folders, tables, and queries. For more information, see [Search](#) on page 102.

Object browser

The object browser, located below the search field, allows you to browse for objects while maintaining a hierarchical perspective. This is the default view in a new **TRS** window.

To expand a folder, either click the **Expand** (▸) icon or double-click the folder. To collapse a folder, either click the **Collapse** (◀) icon or double-click the folder.

The object browser displays the title of the object. Place the pointer over the **Information** (i) icon to view the name of the object.

Favorites list

The favorites list displays a list of your favorite tables and queries.

You can use the favorites list to quickly and easily open tables and queries in the Trillion-Row Spreadsheet that you use often.

The Favorites list is accessed by clicking **Favorites** from the **Analyze** tab in a new TRS window.

Note: The **Favorites** button is not available if you have not previously added any tables or queries to your favorites list. For instructions on adding a favorite, see [Add an object to your favorites](#) on page 110.

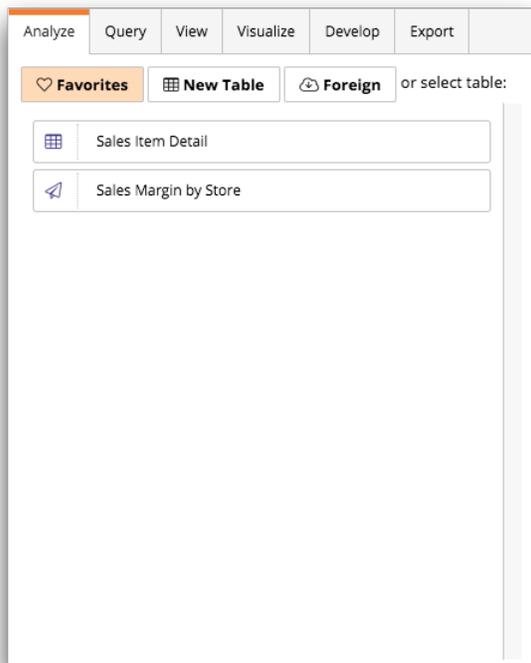


Figure 31: Favorites list in the object browser

<table> panel

The <table> panel allows you to create a new temporary table on which to base your analysis.

The code editor within the <table> panel allows you to create a temporary table using Macro Language code.

The <table> panel is accessed by clicking **New Table** from the **Analyze** tab in a new TRS window. It does not currently have its own TRS panel. The <table> panel contains sample data to get you started creating the temporary table.

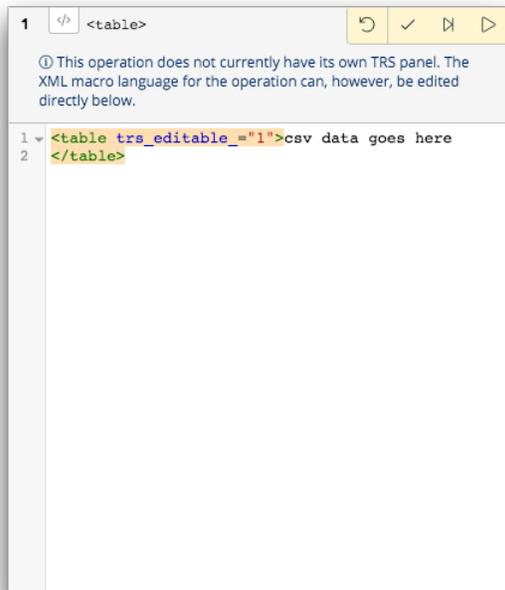


Figure 32: <table> panel

Foreign database panel

The **Foreign database** panel allows you to select a table from a connected database outside of the 1010data Insights Platform to use as the basis of your analysis.

Before you can use a foreign table as the basis of your analysis, a connection from the 1010data Insights Platform to a foreign database must first be configured. Contact 1010data Support for assistance with setting up a foreign database connection.

After a foreign database connection is configured, you can use the table browser in the **Foreign database** panel to browse for and select a table that exists outside of the Insights Platform.

The **Foreign database** panel is accessed by clicking **Foreign** from the **Analyze** tab in a new **TRS** window.

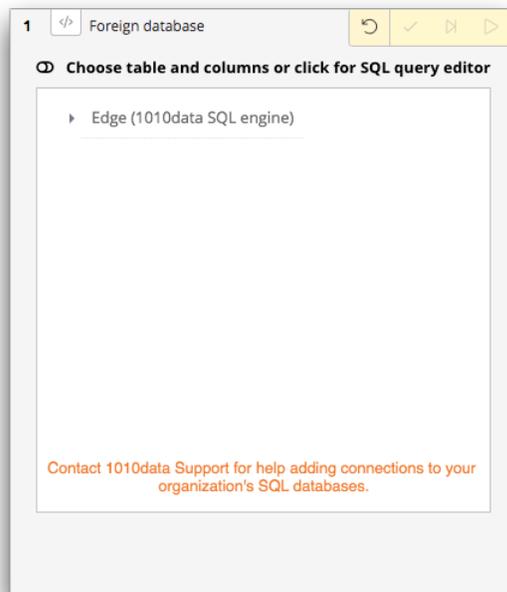


Figure 33: Foreign database panel

Open an existing table or query

Open an existing 1010data Insights Platform table or query in the Trillion-Row Spreadsheet (TRS).

When you open an existing Insights Platform table or query in TRS, the table or query is opened as a worksheet and is displayed in the results pane.

To open an existing table or query:

1. In the object browser of a new **TRS** window, locate the desired table or query by doing one of the following:
 - Use the tree to navigate to the table or query.
 - Enter the name of a table or query in the **Search** field.
 - Click **Favorites** and select a table or query from your list of Favorites.
2. Click the table or query you want to open.
The Trillion-Row Spreadsheet opens the table or query in the results pane and displays the **New operation** panel in the timeline.

Open a favorite table or query

Open a favorite 1010data Insights Platform table or query in the Trillion-Row Spreadsheet (TRS).

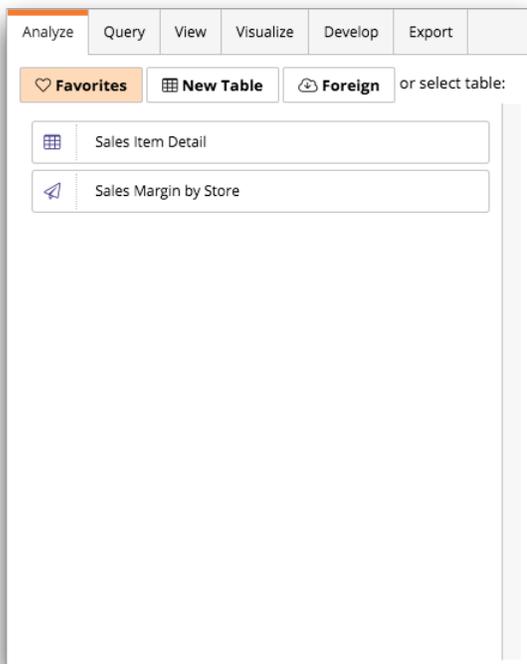
When you open a favorite Insights Platform table or query in TRS, the table or query is opened as a worksheet and is displayed in the results pane.

To open a favorite table or query:

1. In the object browser of a new **TRS** window, click **Favorites**.

Note: The **Favorites** button is not available if you have not previously added any tables or queries to your favorites list.

The Trillion-Row Spreadsheet displays a list of your favorite tables and queries.



2. Click the table or query you want to open.

The Trillion-Row Spreadsheet opens the table or query in the results pane and displays the **New operation** panel in the timeline.

Create a temporary table

You can create a temporary table in the Trillion-Row Spreadsheet (TRS) on which to base your analysis.

At times, you might have data that you want to analyze in TRS that does not already exist as a 1010data Insights Platform table. For example, you might have data in a CSV file that you want to use for your analysis. Rather than create a new table by uploading the file into the Insights Platform, you can create a temporary table from the data directly in the Trillion-Row Spreadsheet using Macro Language code.

To create a temporary table:

1. In the object browser of a new **TRS** window, click **New Table**.
The Trillion-Row Spreadsheet displays the `<table>` panel.

```

1 <table trs_editable='1'>csv data goes here
2 </table>

```

2. Within the `<table>` element in the code editor, enter the data for use in the temporary table. The following is an example of temporary table code.

```

1 <table title="Ladybug Picnic"
  cols="picnic_table_1, picnic_table_2,
  picnic_table_3" labels="Picnic Table 1, Picnic
  Table 2, Picnic Table 3">1, 2, 3; 4, 5, 6; 7, 8,
  9; 10, 11, 12
2 </table>

```

For more information, see `<table>` in the *1010data Reference Manual*.

3. Click the **Submit operation** (✓) icon. The Trillion-Row Spreadsheet displays the temporary table in the results pane and the **New operation** panel in the timeline.

The screenshot displays the 1010data Insights Platform interface. On the left, a panel titled "TRS 1: <table> - Temporary table t14" contains a menu with options: Analyze, Query, View, Visualize, Develop, and Export. Below the menu, there are two numbered sections: "1" with a plus icon and "Edit <table> op" and "2" with a code icon and "New operation". A "New Operation" section follows, featuring an information icon and the text "Choose a new operation to apply to your data." Below this is a grid of operation buttons: Select, Sort, Arrange, Link, Compute, Tabulate, Merge, Amend, Code, Expand Rows, Split, Replace, Date/Time, and Group Metric. A search bar labeled "Search Operations..." is at the bottom of this panel.

On the right, a panel titled "Ladybug Picnic" shows a table with 3 columns and 4 rows. The columns are labeled "Picnic Table 1", "Picnic Table 2", and "Picnic Table 3". The table contains numerical data:

Picnic Table 1	Picnic Table 2	Picnic Table 3
1	2	3
4	5	6
7	8	9
10	11	12

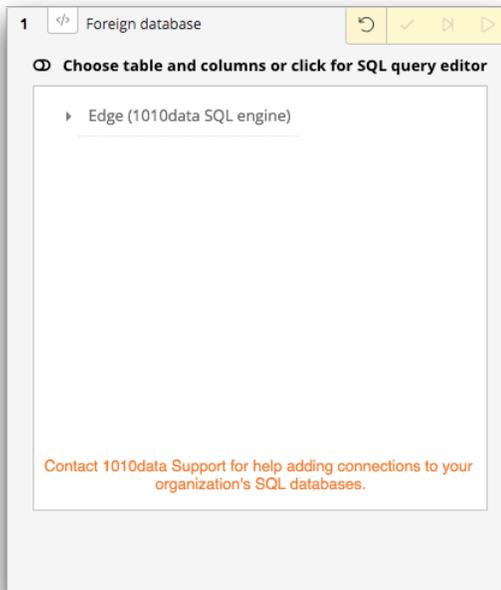
Open a foreign table

Open a table from a connected database outside of the 1010data Insights Platform in the Trillion-Row Spreadsheet (TRS).

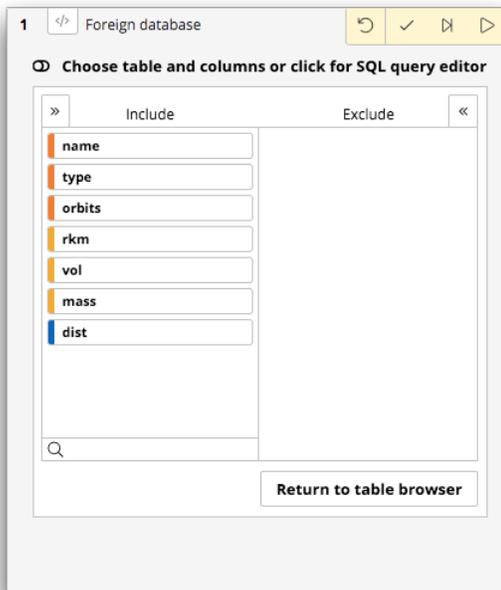
When you open a foreign table in TRS, the table is opened as a worksheet and is displayed in the results pane.

To open a foreign table:

1. In the object browser of a new **TRS** window, click **Foreign**.
The Trillion-Row Spreadsheet displays the **Foreign database** panel.



- Using the table browser, navigate to and click the foreign table you want to open. The Trillion-Row Spreadsheet displays a list of columns contained in the foreign table.



- Perform any of the following actions to choose the columns you want included in the table:

Option	Description
To exclude a column	Drag the column from the Include section to the Exclude section. Note: You can also double-click a column in the Include section to move it to the Exclude section.
To include an excluded column	Drag the column from the Exclude section to the Include section.
To exclude all columns	Click the Move All (») icon above the Include section.

Option	Description
To include all excluded columns	Click the Move All («) icon above the Exclude section.

- Click the **Submit operation** (✓) icon.
The Trillion-Row Spreadsheet displays the foreign table in the results pane and the **New operation** panel in the timeline.

The screenshot shows the TRS interface with the 'Analyze' tab selected. The Analysis Timeline on the left contains two operations: 'Foreign database query' and 'New operation'. The 'New operation' panel is active, displaying a 'New Operation' message and a grid of operation buttons: Select, Sort, Arrange, Link, Compute, Tabulate, Merge, Amend, Code, Affinity, Loyalty, Baskets, Expand Rows, Split, and Replace. The main results pane on the right displays a table with 7 columns: name, type, orbits, rkm, vol, mass, and dist. The table contains 19 rows of data, including celestial bodies like Sun, Jupiter, Saturn, Uranus, Neptune, Earth, Venus, Mars, Ganym, Titan, Mercury, Calli, Io, Moon, Europa, Triton, Pluto, and Eris.

name	type	orbits	rkm	vol	mass	dist
Sun	star	-	696324	1414300000	1988550000	
Jupiter	planet	Sun	69911	1431280	1898600	7783
Saturn	planet	Sun	58232	827130	568460	14270
Uranus	planet	Sun	25362	68340	86832	28710
Neptune	planet	Sun	24622	62540	102430	44971
Earth	planet	Sun	6371	1083	59736	1496
Venus	planet	Sun	6051	928	4868	1082
Mars	planet	Sun	3389	163	642	2279
Ganym >	moon	Jupi >	2634	76	148	10
Titan	moon	Saturn	2576	71	134	12
Mercury	planet	Sun	2439	60	330	579
Calli >	moon	Jupi >	2410	58	108	18
Io	moon	Jupi >	1821	25	89	4
Moon	moon	Earth	1737	21	74	3
Europa	moon	Jupi >	1560	16	48	6
Triton	moon	Nept >	1353	10	21	3
Pluto	dwarf planet	Sun	1186	7	13	59130
Eris	dwarf planet	Sun	1163	6	17	140000

Analysis Timeline

The Analysis Timeline displays a visual representation of your data analysis in the Trillion-Row Spreadsheet.

The Analysis Timeline in the Trillion-Row Spreadsheet (TRS) provides an easy, yet powerful way to visually manipulate a query and the operations it contains.

The operations that define a data analysis query in the TRS are displayed as a sequence of individual panels in the Analysis Timeline. Panels appear in the order that the associated operations are performed on the data. The functionality of each panel varies, and is based on the selected operation. For example, the **Select rows** panel narrows down the data in a table by creating a select operation.

Since each panel represents a separate operation, in a sequence of operations, you can select any panel in the timeline and immediately view the results of the analysis up to that point of the query. This allows you to easily see the state of the data at each step of an analysis.

An operation can be added to, or removed from, the timeline at any point in the query. Furthermore, the individual actions within each operation are editable. Changes to the actions of an operation are tracked in the panel allowing you to undo and redo changes at will. As operations in the timeline are added, removed, or edited, the results of the changes are immediately displayed. This flexibility makes it easy to explore and discover new insights in your data.

After opening a table or query in the TRS, the Analysis Timeline is displayed in, and is accessed from, the **Analyze** tab.

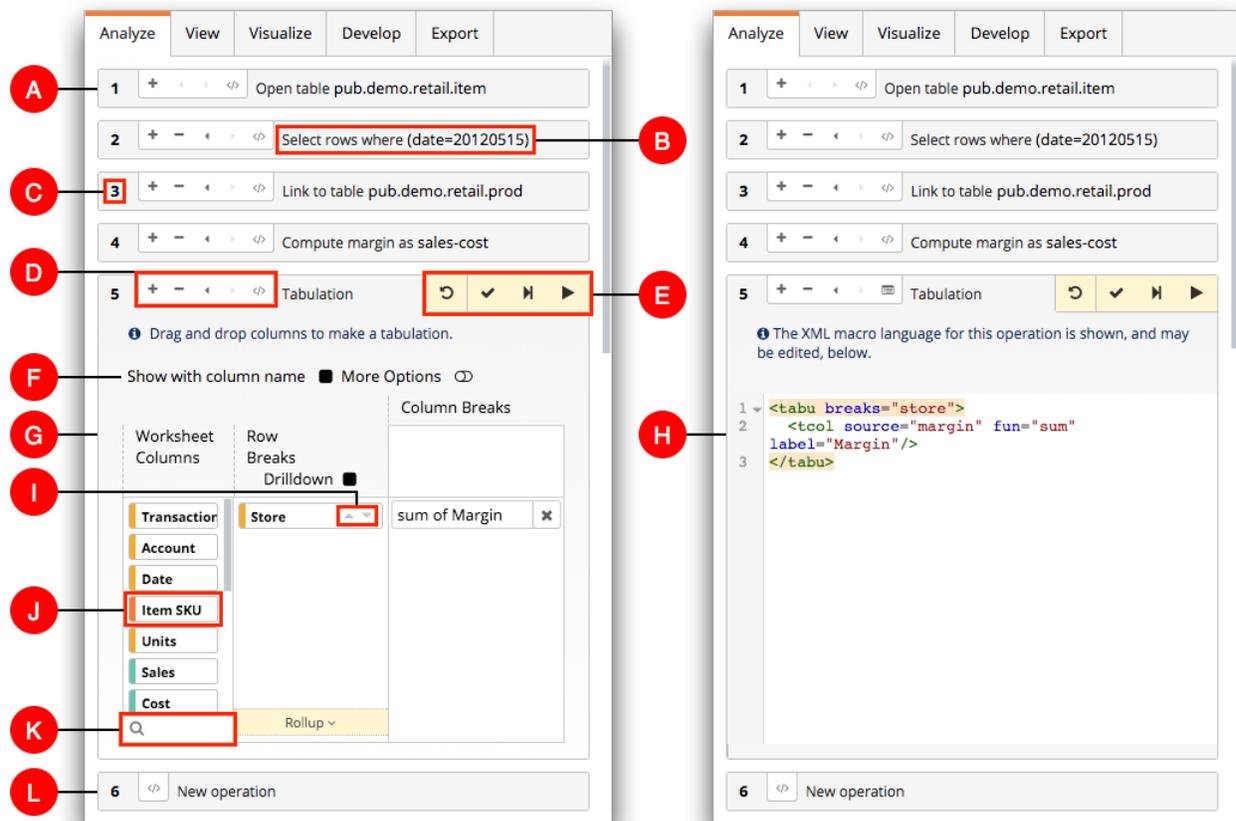


Figure 34: Analysis Timeline in the Trillion-Row Spreadsheet

A. Panel

Each operation in the timeline is contained within a separate panel. The look and functionality of each panel varies and is based on the selected operation. However, the first panel in the timeline always lists the table or query on which the analysis is based.

B. Description

The description provides high-level information about the operation in the panel. This can help you see, at a glance, what each operation in the analysis is doing.

C. Operation number

Each panel in the timeline is assigned an operation number. The operation number helps with keeping track of the panels in the timeline and indicates the order in which operations are performed. In the image above, the link operation is the third operation performed in the analysis.

D. Operation controls

Operation controls allow you to add, remove, and view the history of changes made to an operation. You can also choose to interact with an operation in either the edit panel view or the Macro Language view. For more information, see [Operation controls](#) on page 133.

E. Submit controls

You can submit an operation or cancel any changes made to an operation. When submitting an operation, you have the choice to remain on the same panel, go to the next operation, or run the query to the end. For more information, see [Submit controls](#) on page 138.

F. Show with column name

This option controls whether the column name or column label is displayed in the column biscuits. This option is available in all panels except for the **Select rows** and **Computed column** panels.

When this option is selected, the column name is displayed. When the option is cleared, the column label is displayed. By default, this option is cleared in a new panel. For a description about the differences between the *column name* and *column label*, see [Columns](#) on page 48.

G. Edit panel view

In the timeline image above on the left, the **Tabulation** operation is open in the edit panel view. Use the edit panel view to configure a new operation or make changes to an existing operation using the web interface.

H. Macro Language view

In the timeline image above on the right, the same analysis is shown with the **Tabulation** operation open in the Macro Language view. Use the Macro Language view to configure a new operation or make changes to an existing operation using Macro Language XML code.

I. Sort controls

The sort control icons in column biscuits set the sort order of column data in the results. This option is available in the **Sort columns** and **Computed column** panels.

Click the **Sort Up** () icon to sort the results in the column up. Click the **Sort Down** () icon to sort the results in the column down. When a column sort is set, the selected direction arrow is colored orange (). To remove a set sort order, click the orange direction arrow icon. If no sort order is selected, data in the result columns are sorted up.

J. Column biscuits

In certain panels in the timeline, columns are displayed as biscuits. Drag a column biscuit within a panel to perform tasks such as sorting and linking. Each column biscuit is color coded so that you can quickly determine its data type. For more information about determining the data type of a column, see [Column information](#) on page 49.

K. Column filter

The column filter field limits the number of column biscuits that are displayed in a panel section. This is particularly useful in locating a specific column in a table that has many columns. The column filter field is available in the **Sort columns**, **Arrange columns**, **Link tables**, and **Tabulation** timeline panels. For instructions, see [Filter a list of columns](#) on page 140.

L. New operation panel

The last panel in the timeline is always the **New operation** panel. This panel allows you to add a new operation to your analysis at the end of the timeline. For more information, see [New operation](#) on page 131.

New operation

Add or insert a new operation to your analysis in the timeline.

The **New operation** panel allows you to select from the available operations, organized into groups, that you can use in your analysis. It is always the last panel in the timeline and also appears when you insert a new operation before an existing operation in the timeline. For instructions, see [Insert a new operation](#) on page 135.

Note: In addition to the basic operations, you may have access to additional operations. For example, if you are a retailer, you may have access to operations specific to retail data analysis. Further, the 1010data Insights Platform provides the ability to create custom, also called extended, operations for use in the Trillion-Row Sheet timeline. If you or your organization have custom operations, those operations are also displayed in the **New operation** panel.

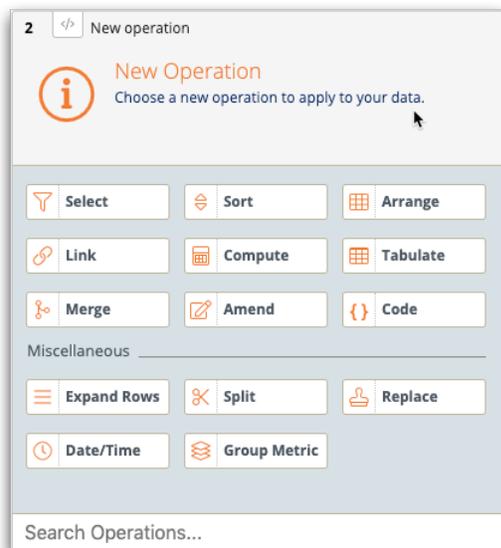


Figure 35: New operation panel

Basic operations

Select

Adds or inserts a **Select rows** panel.

Use row selection to isolate data that is important to your analysis and create more efficient queries. For more information, see [Selecting rows](#) on page 141.

Tip: The keyboard shortcut **?** also adds or inserts a **Select rows** panel.

Sort

Adds or inserts a **Sort columns** panel.

You can reorder the rows of a table such that the values in a particular column are in ascending or descending order. For more information, see [Sorting rows](#) on page 155.

Arrange

Adds or inserts an **Arrange columns** panel.

You can display a subset of the columns in a table and specify the order in which they appear. For more information, see [Arranging columns](#) on page 159.

Link

Adds or inserts a **Link tables** panel.

Use linking to combine tables or worksheets into a single, bigger table with columns from both. For more information, see [Linking tables and worksheets](#) on page 163.

Compute

Adds or inserts a **Compute** panel.

Computed columns are columns that you add to a table, whose values are typically based on one or more existing columns. For more information, see [Computed column](#) on page 180.

Tip: The keyboard shortcut **=** also adds or inserts a **Compute** panel.

Tabulate

Adds or inserts a **Tabulation** panel.

Perform a tabulation to group the values in a column based on the values in another column. For more information, see [Tabulation panel](#) on page 184.

Merge

Adds or inserts a **Merge tables** panel.

You can merge two or more tables or worksheets to combine their rows together into a single, larger worksheet. For more information, see [Merging tables and worksheets](#) on page 216.

Amend

Adds or inserts an **Amend column** panel.

Use amend to replace values in a table or worksheet. For more information, see [Amending values](#) on page 225.

Code

Adds or inserts a **Code** panel.

Use the **Code** panel to run custom SQL, R, or K3 code in your analysis. For more information, see [Use custom code in your analysis](#) on page 230.

Tip: The keyboard shortcut **<** also adds or inserts a **Code** panel.

Miscellaneous operations

Expand Rows

Repeat the number of times that each row appears in the table based on an expression. For more information, see [Expand rows](#) on page 232.

Split

Split and expand a string column on a given separator. For more information, see [Split and expand](#) on page 233.

Replace

Find and replace one or more patterns within a string column. For more information, see [Perform a string replacement](#) on page 234.

Date/Time

Convert strings into dates, times, or intervals, or format time and dates into strings. For more information, see [Convert dates and times](#) on page 235.

Group Metric

Compute an aggregation, grouping rows by one or more columns. For more information, see [Use a group metric](#) on page 237

Search operations

The **Search operations** field narrows the set of available operations based on their label, description, category, or other metadata.

Operation controls

The operation controls allow you to make changes to any operation in your data analysis at any point in the timeline.

You can insert or delete an operation in your analysis, step backward and forward through changes made to an individual operation, and choose the appropriate view for interacting with an operation. Additionally, you can create and manage operation parameters. The operation controls are located at the top left side in each operation panel.

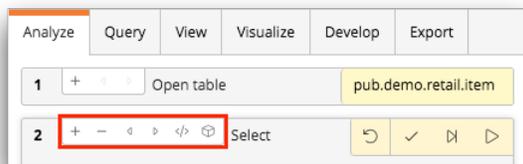


Table 20: Operation control icons

Icon	Name	Description
+	Insert operation	<p>Inserts a New operation panel in the timeline directly after the current operation panel. For more information, see New operation on page 131.</p> <p>For instructions on inserting an operation, see Insert a new operation on page 135.</p>
-	Delete operation	<p>Deletes the operation panel from the timeline.</p> <p>For instructions, see Delete an operation on page 135.</p>
◀	Previous operation state	<p>Go back one step in this operation.</p> <p>Note: If the icon is inactive, no previous steps exist.</p> <p>For instructions, see View the previous operation state on page 136.</p>
▶	Next operation state	<p>Go forward one step in this operation.</p> <p>This icon becomes active after clicking the Previous operation state icon to go back one step.</p> <p>Note: If the icon is inactive, no next steps exist.</p> <p>For instructions, see View the next operation state on page 136.</p>
<>	Macro Language view	<p>Displays the Macro Language view for the operation. This view allows you to use Macro Language XML to edit the operation.</p> <p>Note: This icon is displayed only when the operation is in the edit panel view.</p> <p>For instructions, see Edit an operation using Macro Language XML on page 136.</p>
☰	Edit panel view	<p>Displays the edit panel view for the operation.</p> <p>Note: This icon is displayed only when the operation is in the Macro Language view.</p>
📦	Create block/Edit block	<p>Displays the Block setup view for the operation. This view allows you to create or edit parameters for the operation.</p> <p>Parameters can be used when creating a Quick Query. Queries containing parameters are called Parameterized Quick Queries. For more information, see Quick Queries/PQs on page 385.</p>

Icon	Name	Description
+	Expand block	<p>Clears the parameters in an operation and displays the operation edit panel view.</p> <p>Note: This icon is displayed in the Parameters view of the operation panel indicating that the operation contains parameters. The Expand block icon is also displayed when creating or editing parameters in the Block setup view.</p>

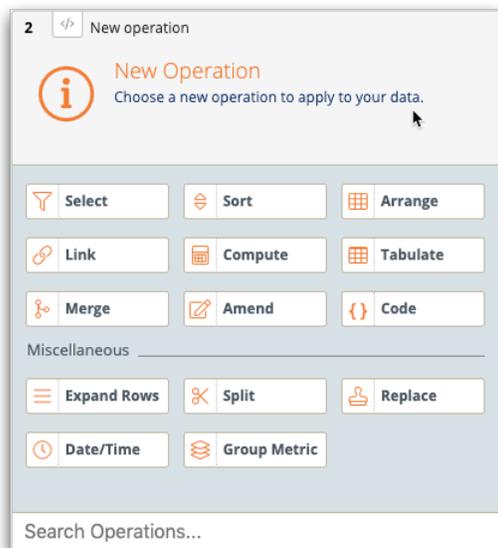
Insert a new operation

Add an operation to your analysis in the timeline by inserting a new operation.

You can insert a **New operation** panel at any location in the Trillion-Row Spreadsheet (TRS) Analysis Timeline to make changes to your query. A new operation is inserted after the currently selected operation in the timeline.

To insert a new operation:

1. In the TRS, click the operation in the timeline where you want the new operation to follow.
2. Click the **Insert operation** (+) icon.
The TRS inserts the **New operation** panel after the currently selected operation in the timeline.



Delete an operation

Delete an operation from your analysis in the timeline.

You can delete an operation panel in the Trillion-Row Spreadsheet (TRS) Analysis Timeline to make changes to your query.

To delete an operation:

1. In the TRS, click the operation in the timeline that you want to delete.
2. Click the **Delete operation** (-) icon.
You can also use the keyboard shortcut **Ctrl+Z** (PC) or **Command+Z** (Mac) to undo the *last* operation in the timeline.
The TRS deletes the selected operation and displays the next operation panel in the timeline.

View the previous operation state

Move backward through the actions applied to an operation in the Analysis Timeline.

You can move backward through the history of actions executed within an operation in the timeline of the Trillion-Row Spreadsheet (TRS). With each step back, the results of the data displayed in the grid are updated accordingly. You can repeat this all the way to the initial state of the operation.

After moving backward through the actions, you can move forward through those same actions to execute them once again.

To view the previous operation state:

1. In the TRS timeline, click the operation for which you want to view the previous operation state.
2. Click the **Previous operation state** (◀) icon.

Note: If the icon is inactive, no previous steps exist.

The TRS moves backward one step in the history of actions and displays the state of the operation in the panel. In addition, the grid is updated with any changes to the data resulting from the step backward.

View the next operation state

Move forward through the actions previously applied to an operation in the Analysis Timeline.

After moving backward through the history of actions executed within an operation in the timeline of the Trillion-Row Spreadsheet (TRS), you can move forward through those same actions to execute them once again. With each step forward, the results of the data displayed in the grid are updated accordingly. You can repeat this all the way to the final state of the operation.

To view the next operation state:

1. In the TRS timeline, click the operation for which you want to view the next operation state.
2. Click the **Next operation state** (▶) icon.

Note: If the icon is inactive, no next steps exist.

The TRS moves forward one step in the history of actions and displays the state of the operation in the panel. In addition, the grid is updated with any changes to the data resulting from the step forward.

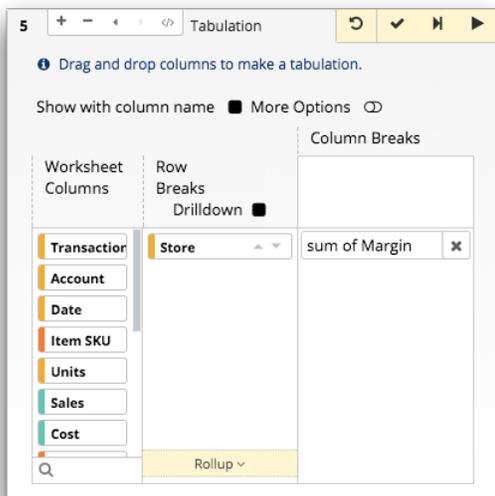
Edit an operation using Macro Language XML

Use Macro Language XML code to edit an operation in the Analysis Timeline.

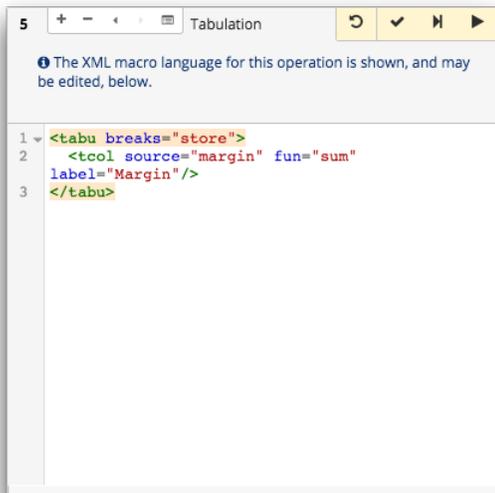
You can view and edit the Macro Language XML code of an operation directly from the timeline panel using the Macro Language view.

To edit an operation using Macro Language XML:

1. In the TRS timeline, click the operation that you want to edit using Macro Language XML code. The TRS displays the edit panel view of the operation.



2. Click the **Macro Language view** (`</>`) icon.
The TRS displays the Macro Language view of the operation.



3. In the code editor, edit the operation Macro Language XML code.
4. Click the **Submit operation** () icon.
The Insights Platform updates the operation and displays the results in the grid.

Note: To return to the edit panel view of the operation, click the **Edit panel view** () icon.

Add an operation using Macro Language XML

Use Macro Language XML code to add an operation to your analysis in the timeline.

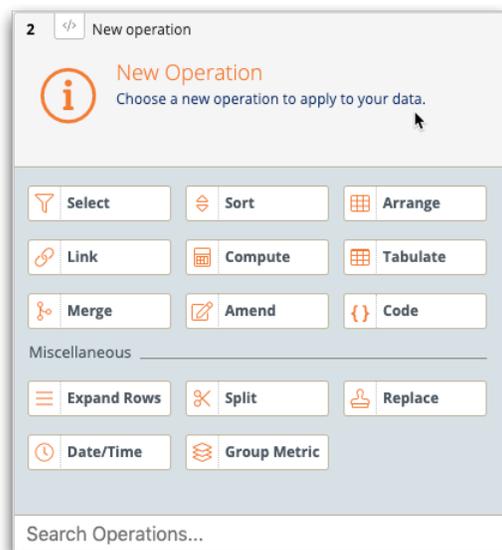
You can use Macro Language XML code to add a new operation to the Trillion-Row Spreadsheet (TRS) Analysis Timeline.

To add an operation using Macro Language XML:

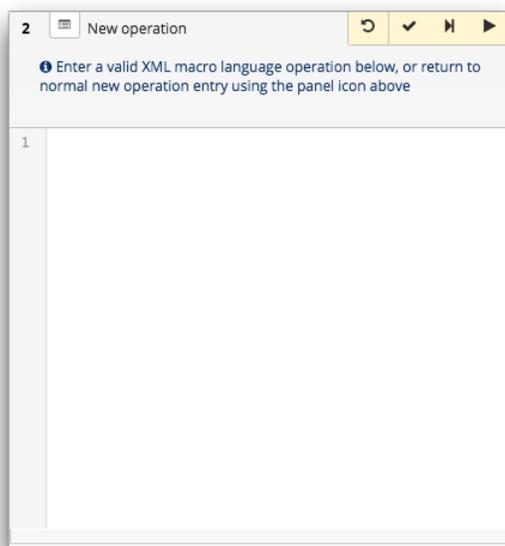
1. Do one of the following:
 - Insert a new operation in the timeline.

- Click the **New operation** panel at the end of the timeline.

The TRS displays the **New operation** panel.



2. Click the **Macro Language view** (`</>`) icon.
The TRS displays the Macro Language view of the **New operation** panel.



3. In the code editor, add the Macro Language XML code for the new operation.
4. Click the **Submit operation** () icon.
The Insights Platform adds the operation and displays the results in the grid.

Note: To return to the edit panel view of the operation, click the **Edit panel view** () icon.

Submit controls

The submit controls allow you to submit or cancel changes made to an operation.

You can submit an operation or cancel any changes made to an operation. When submitting an operation, you have the choice to remain on the same panel, go to the next operation, or run the query to the end. The submit controls are located at the top right side in each operation panel.

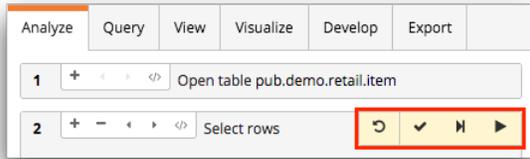


Table 21: Submit controls icons

Icon	Name	Description
	Cancel changes	<p>Cancels any changes made to an operation in a panel.</p> <p>In a new operation that has not yet been submitted, clicking Cancel changes displays the New operation panel.</p> <p>In a previously submitted operation, clicking Cancel changes cancels changes made in the panel and displays the operation in the last state it was submitted.</p>
	Submit operation	Submits the operation and stays on the same operation panel.
	Submit and go to next operation	Submits the operation and goes to the next operation in the timeline.
	Submit and run query to end	Submits the query and runs all operations to the end of the timeline.

Expression Editor

Use the Expression Editor to build valid selection and value expressions in the Trillion-Row Spreadsheet.

The Expression Editor assists in composing 1010data Insights Platform expressions when selecting rows or creating computed columns in the Trillion-Row Spreadsheet (TRS).

The Expression Editor is available in the following locations of the TRS:

- Expression Editor view of the **Select rows** panel
- **Computed column** panel
- **Amend column** panel
- **Find rows** dialog in the grid

In the images below, a partial sample expression is entered in the Expression Editor for illustrative purposes.

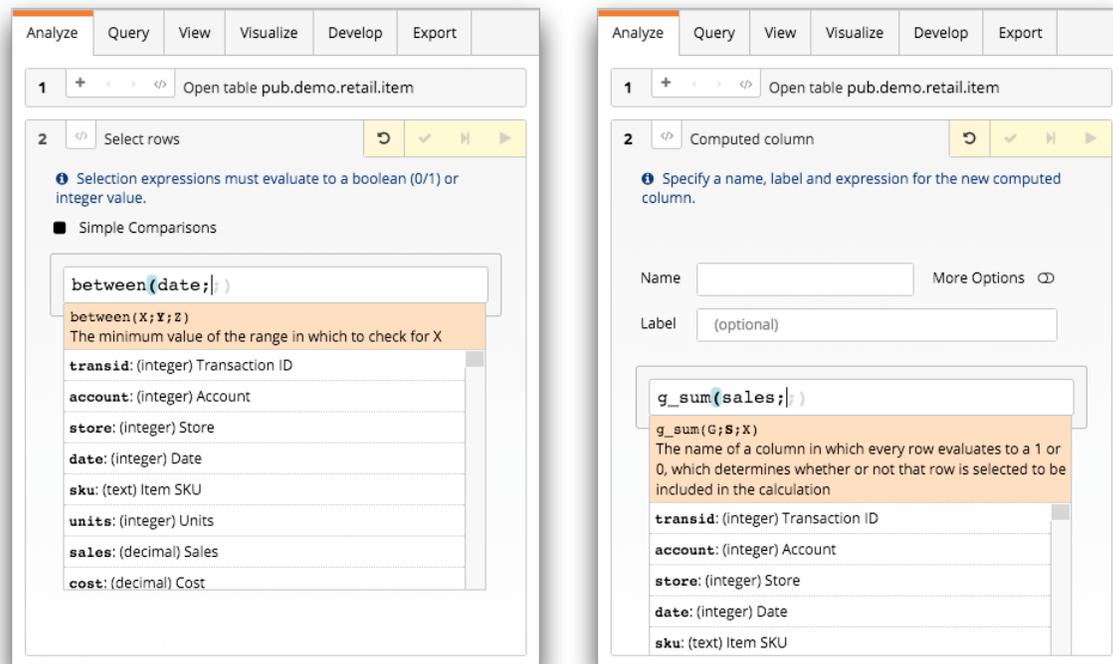


Figure 36: The Expression Editor in TRS timeline panels

The Expression Editor helps you build valid 1010data Insights Platform expressions within the TRS by providing:

- Suggestion of columns, variables, and functions. Suggestions include descriptions of all system variables. Use the **Up Arrow** and **Down Arrow** keys or the mouse to navigate the list of suggestions. Press the **Tab** key or click the menu item to insert it.
- An inline function reference. While editing the subexpressions representing any function's arguments, the top portion of the context menu indicates the argument index and a description of that argument's constraints and purpose.
- Detection and highlighting of mismatched delimiters (e.g., brackets, braces, quotation marks, and parentheses).
- Highlighting of the current enclosing delimiters.
- Suggestion of closing delimiters and the implied semicolon separators for function calls. To accept the suggestions, press the **Right Arrow** key.
- Coloring of quotations, so they stand out easily.
- Context-aware insertion and replacement. If the cursor is positioned within an existing column, variable, or function, choosing a replacement from the context menu replaces the item with the new selection.
- Convenient text selection. To select text from the location of the cursor to the beginning of the entry field, press **Ctrl+Shift+Left Arrow**. To select text from the location of the cursor to the end of the entry field press **Ctrl+Shift+Right Arrow**.

For more information about expressions in the Insights Platform, see [Writing Expressions](#) in the *1010data Reference Manual*.

Filter a list of columns

Limit the number of column biscuits that are displayed in an operation panel section.

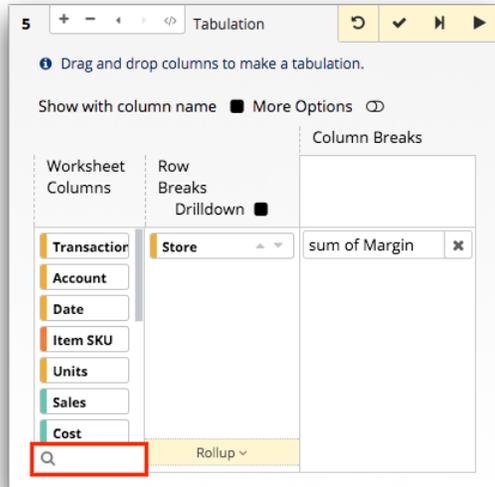
The column filter field limits the number of column biscuits that are displayed in a Trillion-Row Spreadsheet (TRS) operation panel section. Filtering is particularly useful in locating a specific column in a table that has many columns.

The column filter field is available in the **Sort columns**, **Arrange columns**, **Link tables**, and **Tabulation** timeline panels.

To filter a list of columns:

1. In the TRS, select a panel that contains the **Column filter** field.

For example, the **Column filter** field, outlined in red, is shown in the **Tabulation** panel below.



2. In the **Column filter** field, enter the name or label of a column.

As you enter text, the TRS filters the list and displays only the columns that match your entry.

Note: To clear the filter, click the **Clear Filter** (🗑️) icon.

Selecting rows

Selecting rows is one of the most basic tasks in the 1010data Insights Platform. Use row selection to isolate data that is important to your analysis and create more efficient queries.

Whether you are working with a table that is very small or very large, the first step to answering an analytical question is to narrow down the data you are working with so you see only the pieces of information that apply. While there are many ways to do this, the most basic is a row selection.

When you work with very large data sets like those typically found in the Insights Platform, the order in which you perform operations on the data makes a difference in how fast the system can complete those operations. If you make the largest selection that eliminates the most data first, then the subsequent operations you perform finish faster because the platform does not need to work with as much information. As a best practice, if your data contains columns with date- or time-related information, that is generally the best place to start. However, the largest selection may be something other than date or time information and is based on your data and the analysis you want to perform.

If you perform several row selections in succession, each selection is applied to the previously selected rows. With each selection, the number of rows will decrease (or possibly stay the same); it will never increase. The Analysis Timeline displays every selection applied to the table or worksheet and the number of rows remaining after these selections are listed at the top of the grid. Only those rows that meet your criteria are displayed. Also, if you summarize the data, only the selected rows will be included in the summary.

Note: The selection applies only to your current workspace and does not affect the original table.

There are several ways to perform a row selection in the Trillion-Row Spreadsheet. You can specify the selection criteria using the Simple Comparisons options, enter a selection expression in the Expression Editor, or select rows within the grid itself.

Select rows panel

The **Select rows** panel allows you to specify the selection criteria using the simple comparisons options or enter a selection expression in the Expression Editor.

The **Select rows** panel is accessible from the following locations in the 1010data Insights Platform:

- Analysis Timeline in the Trillion-Row Spreadsheet
- Right-click menu in the grid

It is important to note that the **Select rows** panel available in the grid is visually different from the panel in the Analysis Timeline. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the **Select rows** panel as it appears in the timeline.

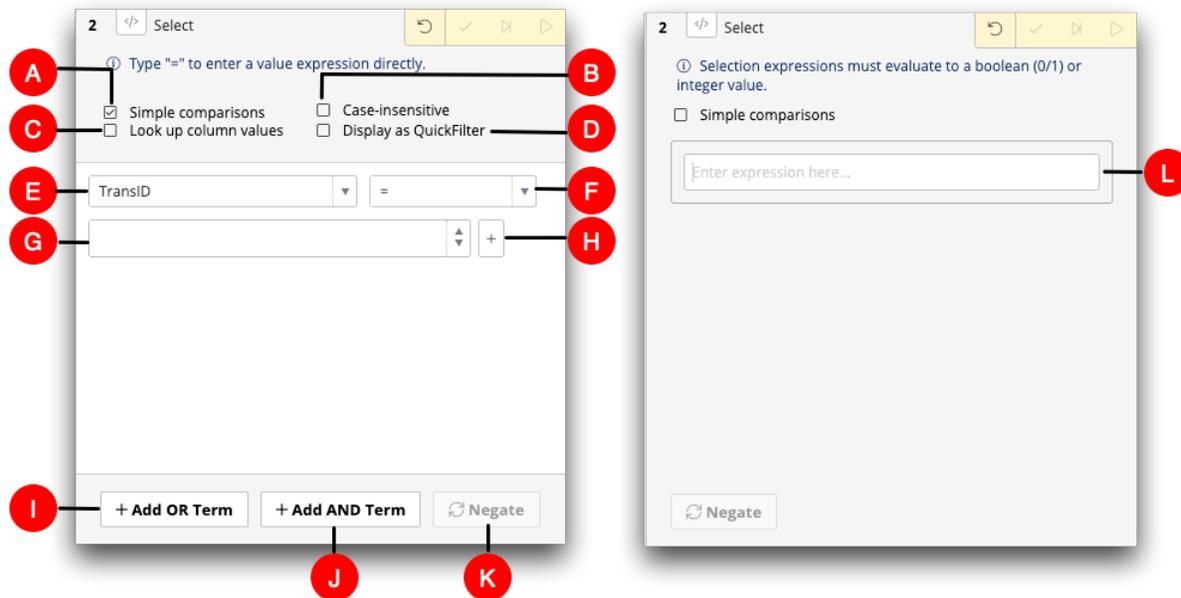


Figure 37: Select rows panel in the Analysis Timeline

A. Simple Comparisons

In a new **Select rows** panel, the **Simple Comparisons** option is selected by default. When selected, fields and options used to specify the criteria for a row selection are displayed in the Simple Comparisons view. Clear this option to display the Expression Editor view.

B. Case-insensitive

When the selected column is a text column, select this option to ignore case in your search. This option does not appear for non-text columns.

C. Look up column values

If you select this option, you can start entering a value, and the **Select Row** panel will display relevant values in a drop-down list. You can then select a value from the drop-down list.

D. Display as QuickFilter

If you select this option, the search criteria will appear in the results pane along with the resulting table.

E. Column

This drop-down list contains the columns in the table or worksheet. By default, the first column in the table or worksheet is displayed first in the list. When the list is open, you can filter the list of columns

by entering text in the search field. This is helpful if your table or worksheet contains a large number of columns.

F. Relationship

This drop-down list contains the options used to specify the relationship (e.g., equal to, greater than, contains, etc.) between a column and a value. For a list of relationship options, see [Simple comparisons relationship options](#) on page 143.

G. Value

This field is used to enter or select the value for which you want to search in the table or worksheet.

The combination of the column, relationship, and value are called a *term* and define the criteria required for a row selection. As an example, if you wanted to select all rows in a table where the word "milk" appears in a cell within a specific column, you would select the desired column from the **Column** drop-down list, select **contains** from the **Relationship** drop-down list, and enter `milk` in the **Value** field.

H. Add Value

This icon adds an additional **Value** field to the selection criteria.

Adding additional **Value** fields creates an *or* selection. This type of selection includes only those rows that meet *any* of the entered value criteria.

To remove a **Value** field from the selection criteria, click the **Remove Value**  icon.

I. Add 'or' Term

This button adds an additional set of selection criteria, called a *term*, to the row selection. A term is the combination of the column, relationship, and value.

Adding an additional term creates an *or* selection. This type of selection includes rows that meet *any* of the of the entered value criteria.

J. Add 'and' Term

This button creates an *and* selection. This type of selection includes rows that meet *all* of the entered value criteria.

K. Negate

Click this button to negate the relationship selected in the **Relationship** drop-down list. For example, a relationship of "has the value" (=) will become "does not have the value" (≠), and a relationship of "contains" will become "does not contain".

L. Expression Editor

The Expression Editor is available from the Expression Editor view in the **Select rows** panel. To access the Expression Editor, clear the **Simple Comparisons** option. For an easy shortcut, simply type "=" to view the Expression Editor. For more information about the Expression Editor, see [Expression Editor](#) on page 139.

The Expression Editor allows you to enter a selection expression to perform a row selection. For instructions, see [Select rows by entering a selection expression](#) on page 149.

Simple comparisons relationship options

When performing a simple comparison in the **Select rows** panel, the relationship options define the association between a column and a value in the resulting selection expression.

The **Relationship** drop-down list contains the options used to specify the association (e.g., equal to, greater than, contains, etc.) between a column and a value. The options available in the list are dependent on the data type of the selected column. For more information about data types, see [Data types](#) on page 53.

When the select operation is submitted, the 1010data Insights Platform searches the selected column for the specified value, or values, and displays the appropriate rows in the grid.

Note: Values in text columns are case sensitive.

The following relationship options are available when performing a simple comparison in the **Select rows** panel.

Table 22: Relationship options

Relationship	Applicable data types	Description
= (has the value)	<ul style="list-style-type: none"> integer big integer decimal text 	Rows where data in the selected column is identical to the specified value are displayed in the grid.
≠ (does not have the value)	<ul style="list-style-type: none"> integer big integer decimal text 	Rows where data in the selected column contains anything other than the specified value are displayed in the grid.
> (is greater than)	<ul style="list-style-type: none"> integer big integer decimal text 	Rows where the data in the selected column is greater than the specified value are displayed in the grid.
≥ (is greater than or equal to)	<ul style="list-style-type: none"> integer big integer decimal text 	Rows where the data in the selected column is greater than or equal to the specified value are displayed in the grid.
< (is less than)	<ul style="list-style-type: none"> integer big integer decimal text 	Rows where the data in the selected column is less than the specified value are displayed in the grid.
≤ (is less than or equal to)	<ul style="list-style-type: none"> integer big integer decimal text 	Rows where the data in the selected column is less than or equal to the specified value are displayed in the grid.
is between	<ul style="list-style-type: none"> integer big integer decimal text 	Rows where the data in the selected column is between, or equal to, the specified values are displayed in the grid.
contains	text	Rows where the data in the selected column contain the entered text string are displayed in the grid.

Relationship	Applicable data types	Description
does not contain	text	Rows where data in the selected column does not contain the entered text string are displayed in the grid.
begins with	text	Rows where data in the selected column begins with the entered text string are displayed in the grid.
ends with	text	Rows where data in the selected column ends with the entered text string are displayed in the grid.
is N/A	<ul style="list-style-type: none"> integer big integer decimal text 	Rows where the selected column contains an N/A value are displayed in the grid.
is not N/A	<ul style="list-style-type: none"> integer big integer decimal text 	Rows where the selected column contains anything other than an N/A value are displayed in the grid.
is one of	<ul style="list-style-type: none"> integer big integer decimal text 	<p>Rows where data in the selected column is identical to any one of the specified values are displayed in the grid.</p> <p>When selected, this option allows you specify multiple values in the same Value field. For more information, see Specifying a list of values in a selection on page 147.</p>
is not one of	<ul style="list-style-type: none"> integer big integer decimal text 	<p>Rows where data in the selected column contains anything other than the specified values are displayed in the grid.</p> <p>When selected, this option allows you specify multiple values in the same Value field. For more information, see Specifying a list of values in a selection on page 147.</p>

Use simple comparisons to select rows

Row selection allows you to focus on a subset of rows.

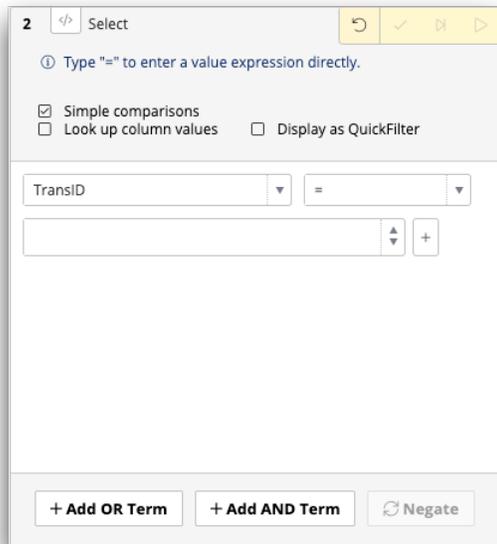
To use simple comparisons to select rows:

1. Open the **Select rows** panel by doing one of the following:

- In the **New operation** panel, click **Select**.
- In the grid, right-click a cell within the column on which you want to base the selection and click **Select on [COLUMN_LABEL]**.

Note: In this topic, *[COLUMN_LABEL]* represents the label of the column in which the selected cell is located.

The Trillion-Row Spreadsheet displays the **Simple Comparisons** fields and options in the **Select rows** panel.



Note: The **Select rows** panel available in the grid is visually different from the panel in the Analysis Timeline. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the **Select rows** panel as it appears in the timeline.

- From the **Column** drop-down list, select a column.

This drop-down list contains the columns in the table or worksheet. By default, the first column in the table or worksheet is displayed first in the list.

Note: When the list is open, you can filter the list of columns by entering text in the search field. This is helpful if your table or worksheet contains a large number of columns.

- From the **Relationship** drop-down list, select a relationship.

This drop-down list contains the options used to specify the relationship (e.g., equal to, greater than, contains, etc.) between a column and a value. For a list of relationship options, see [Simple comparisons relationship options](#) on page 143.

- In the **Value** field, enter or select a value.

Note: You can specify multiple values in the same **Value** field when either **is one of** or **is not one of** is selected from the **Relationship** drop-down list. For more information, see [Specifying a list of values in a selection](#) on page 147.

- Optionally, add an additional value to the row selection.

Adding additional **Value** fields creates an *or* selection. This type of selection includes those rows that meet any of the entered value criteria for the selected column.

Note: You can specify more than one value when any of the following relationship options are selected: **=** (has the value), **≠** (does not have the value), **contains**, **does not contain**, **begins with**, or **ends with**.

- Click the **Add Value** (+) icon.
 - In the new **Value** field, specify the additional value.
- Optionally, add an additional term to the row selection.

A *term* is the combination of a column, relationship, and value. You can choose whether to add an **OR** term or an **AND** term.

- Click **Add AND Term**.

The Insights Platform adds an additional set of selection criteria (a term) to the row selection.

3 Select

Type "=" to enter a value expression directly.

Simple comparisons
 Look up column values

Store = 1

AND

Trans ID = 1400053529

+ Add AND Term - Remove Term Negate

b) Specify the column, relationship, and value for the additional term.

7. Click the **Submit operation** (✓) icon.

The Insights Platform displays the results of your selection in the grid.

Specifying a list of values in a selection

Select rows based on multiple values without the need to specify each one individually.

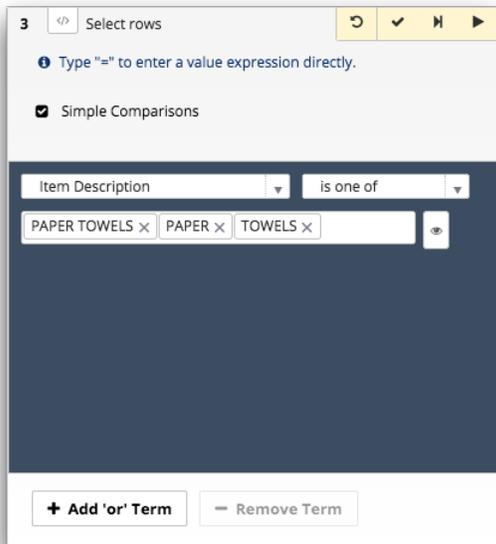
There may be times when you want to use an existing list of values in a new row selection. For example, you may have data located in a separate file, such as within an email or a spreadsheet, that you want to use in your selection. Rather than manually specifying each of those values separately, you can copy the list and paste it into the **Select rows** panel to specify the values all at the same time.

Note: To specify a list of values in a selection, either **is one of** or **is not one of** must be selected from the **Relationship** drop-down list in the **Simple Comparisons** options of the **Select rows** panel. For more information about either of these options, see [Simple comparisons relationship options](#) on page 143.

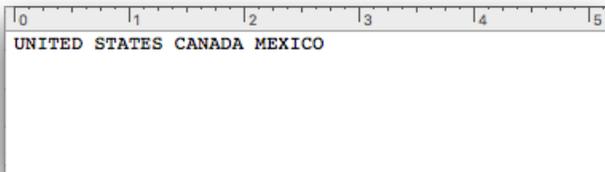
When a list of values is pasted in the **Select rows** panel, the 1010data Insights Platform determines how to split the list based on whether the list of values was copied from a spreadsheet or a text file. In a spreadsheet file, data in each cell is treated as a single value. For example, the spreadsheet file below contains a list of values separated by cell.

	A	B	C
1	PAPER TOWELS		
2	PAPER		
3	TOWELS		
4			
5			

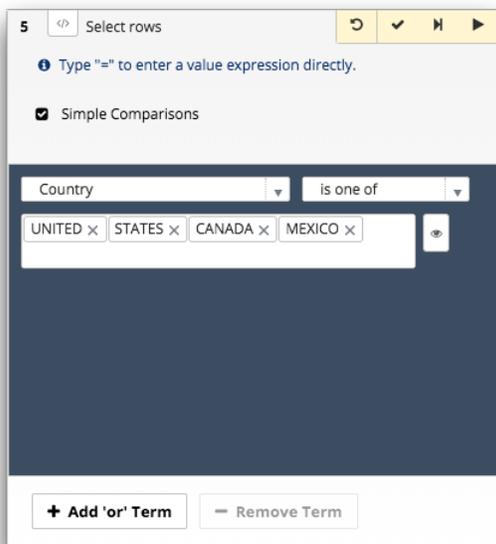
When this list of values is pasted into the **Value** field in the **Simple Comparisons** options of the **Select rows** panel, the Insights Platform splits the list of values by cell.



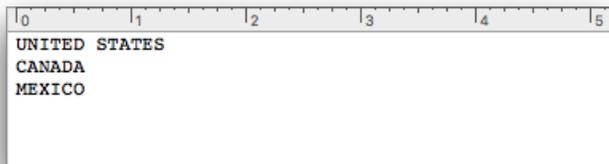
In a text file, the Insights Platform splits data in one of two ways. If no line breaks exist in the file, values are split by spaces. However, if line breaks exist in the file, spaces are ignored and values are split by line breaks. For example, the text file below contains a list of values separated by spaces without line breaks.



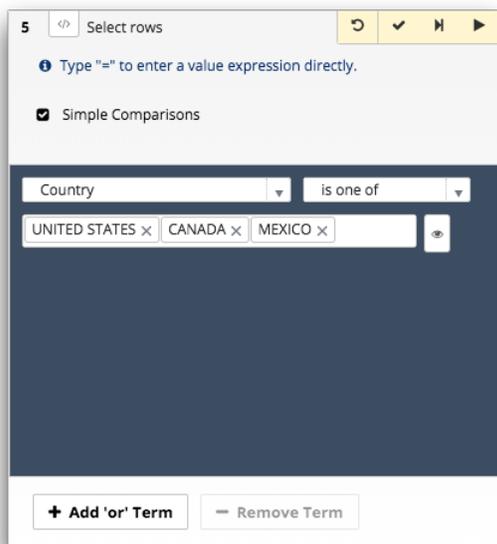
When this list of values is pasted into the **Value** field in the **Select rows** panel, the Insights Platform splits the list of values by spaces.



Because the text file did not contain line breaks, UNITED STATES was split into two separate values: UNITED and STATES. For comparison, the text file below contains the same list of values, but separated by line breaks.



When this list of values is pasted into the **Value** field in the **Select rows** panel, the Insights Platform splits the list of values by line breaks.



Even though the text file contained spaces, UNITED STATES was retained as a single value.

In addition to pasting a list of values, you can also manually enter multiple values into the **Value** field in the **Select rows** panel. To do so, after entering a value, press **Enter**. Manually entered values are split by the **Enter** key.

Select rows by entering a selection expression

You can define a subset of rows using a selection expression, a mathematical formula that evaluates to 1 (true) or 0 (false).

Enter a selection expression to define a subset of rows.

Only those rows for which the selection expression evaluates to true will be displayed. Also, if you summarize the data, only the selected rows will be included in the summary.

Note: The selection applies only to your current workspace and does not affect the original table.

To select rows by entering a selection expression:

1. Open the **Select rows** panel by doing one of the following:
 - In the **New operation** panel, click **Select**.
 - In the grid, right-click a cell within the column on which you want to base the selection and click **Select on [COLUMN_LABEL]**.

Note: In this topic, *[COLUMN_LABEL]* represents the label of the column in which the selected cell is located.

The Trillion-Row Spreadsheet displays the **Simple Comparisons** fields and options in the **Select rows** panel.

Note: The **Select rows** panel available in the grid is visually different from the panel in the Analysis Timeline. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the **Select rows** panel as it appears in the timeline.

2. Clear the **Simple Comparisons** checkbox.

Tip: Simply type "=" to switch to the Expression Editor view, and then enter your expression.

The Trillion-Row Spreadsheet displays the Expression Editor view.

3. In the Expression Editor, enter a selection expression.

The selection expression may refer to one or more columns and may include standard arithmetic, relational, and logical operators as well as any Insight Platform functions.

Note: When entering a column in an expression, use the column *name* as opposed to the column *heading*.

For example, to select rows where the transaction ID (`transid`) is between 535 and 540, use the selection expression `between(transid;535;540)`.

For more information, see [Writing Expressions](#) in the *1010data Reference Manual*. For more information about the Expression Editor, see [Expression Editor](#) on page 139.

4. Click the **Submit operation** () icon.
The Trillion-Row Spreadsheet displays the results of your selection in the grid.

Performing a Quick selection using the grid

Use the grid to perform a Quick selection in a table or worksheet.

The Trillion-Row Spreadsheet allows you to perform a Quick select directly from within the grid. Options include the ability to select rows based on the value within a cell or to select rows based on their location within the table or worksheet.

A Quick select can be based on a single cell or on a group of cells. To choose a group of cells, you draw a selection rectangle in the grid around the cells you want included in the selection.

Select rows based on cell value

Use a value within a cell in the grid to perform a row selection.

You can quickly perform a basic row selection directly from a table or worksheet within the grid by right-clicking any cell and then choosing a selection option in the menu. These options let you use the value in the selected cell as the basis for the row selection.

Note: In this topic, `[COLUMN_LABEL]` represents the label of the column in which the selected cell is located and `[CELL_VALUE]` represents the value contained in the selected cell.

To select rows based on cell value:

1. In the Grid view of an open table or worksheet, right-click the cell on which you want to base the row selection and point to **Quick select where**.
The Trillion-Row Spreadsheet displays a list of selection options.

Sales Item Detail
Cols 1 to 8 of 8 , Rows 1 to 19 of 35

Transaction ID transid	Account account	Store store	Date date	Item SKU sku	Units units	Sales sales	Cost cost
531	957	1	05/18/12	96A	1	-5	-1.84
532	478	1				0.5	0.25
532	478	1				1.1	0.56
534	738	1				6	2.9
534	738	1				2.25	1.35
535	709	2					
535	709	2					
535	709	2					
536	748	2					
536	748	2					
537	523	3					
537	523	3				6	2.9
537	523	3				0.5	0.25
537	523	3				1.1	1
537	523	3				1.1	0.56
538	668	1				1.65	1.1
538	668	1	05/18/12	98A	4	2	1
538	668	1	05/18/12	96A	1	1.1	1

2. Click the desired selection option.

Available options are specific to the selected cell and include the following:

Selection command	Description
[COLUMN_LABEL] has the value [CELL_VALUE]	Selects rows that have the same value, within the same column, as the cell you right-clicked.
[COLUMN_LABEL] does not have the value [CELL_VALUE]	Selects rows that do not have the same value, within the same column, as the cell you right-clicked.
[COLUMN_LABEL] is N/A	Selects rows, within the same column as the cell you right-clicked, that contain an N/A value.
[COLUMN_LABEL] is not N/A	Selects rows, within the same column as the cell you right-clicked, that contain anything other than an N/A value.
[COLUMN_LABEL] is greater than [CELL_VALUE]	Selects rows, within the same column as the cell you right-clicked, that have a larger value.
[COLUMN_LABEL] is less than [CELL_VALUE]	Selects rows, within the same column as the cell you right-clicked, that have a smaller value.

The Trillion-Row Spreadsheet displays the results of your selection in the grid and adds the select operation to the Analysis Timeline.

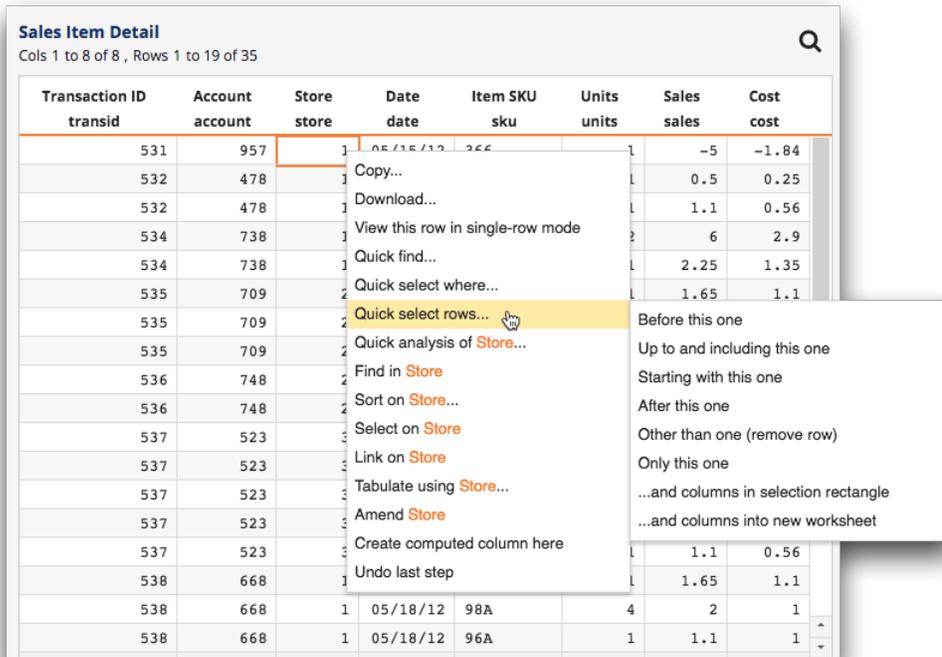
Select rows based on cell location

Select rows in the table or worksheet based on the location of the row in the grid.

You can quickly perform a basic row selection directly from a table or worksheet within the grid by right-clicking any cell and then choosing a selection option in the menu. These options use the location of the selected cell as the basis for the row selection.

To select rows based on cell location:

- In the Grid view of an open table or worksheet, right-click the cell on which you want to base the row selection and point to **Quick select rows**.
The Trillion-Row Spreadsheet displays a list of selection options.



- Click the desired selection option.

Available options include the following:

Selection command	Description
Before this one	Rows listed before the row containing the cell you right-clicked are selected.
Up to and including this one	Rows listed before, and including, the row containing the cell you right-clicked are selected.
Starting with this one	Rows listed after, and including, the row containing the cell you right-clicked are selected.
After this one	Rows listed after the row containing the cell you right-clicked are selected.
Other than this one (remove row)	All rows other than the row containing the cell you right-clicked are selected.
Only this one	Only the row containing the cell you right-clicked is selected.

The Trillion-Row Spreadsheet displays the results of your selection in the grid and adds the select operation to the Analysis Timeline.

Select a group of cells

Draw a selection rectangle within the grid to select a group of cells from a table or worksheet.

You can select a group of cells within a table or worksheet by drawing a selection rectangle around a collection of cells in the grid. The Trillion-Row Spreadsheet then uses two separate operations to reduce the table or worksheet to only those cells contained within the rectangle.

First the Trillion-Row Spreadsheet utilizes the select operation to select rows within the rectangle. Second, any columns that fall outside of the rectangle are hidden from the grid with the arrange operation. The results of your selection are displayed in the grid and the two operations are added to the Analysis Timeline.

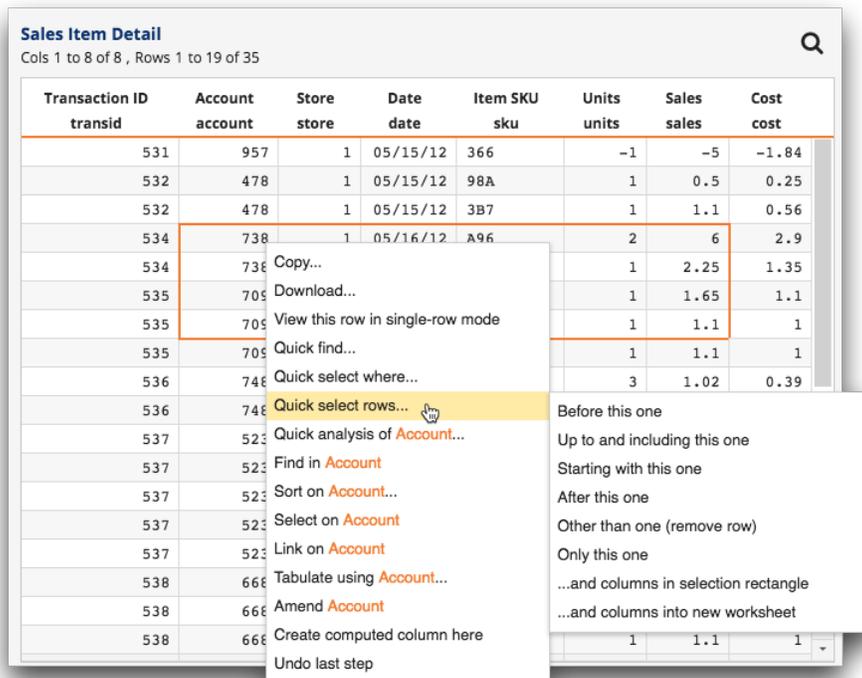
You have the option of viewing the results of the selection within the same **TRS** window or as a new worksheet in a separate **TRS** window.

To select a group of cells:

1. In the Grid view of an open table or worksheet, draw a selection rectangle around the cells you want to select.

For instructions, see [Draw a selection rectangle](#) on page 302.

2. Right-click any cell within the selection rectangle and point to **Quick select rows** in the menu. The Trillion-Row Spreadsheet displays a list of selection commands.



3. Click one of the following options:

Selection command	Description
and columns in selection rectangle	This option performs the selection and displays the results within the same TRS window.
and columns into new worksheet	This option performs the selection and displays the results as a new worksheet in a separate TRS window.

The Trillion-Row Spreadsheet displays the results of your selection in the grid and adds the select and arrange operations to the Analysis Timeline.

The screenshot shows the 'Sales Item Detail' table with the following data:

Account account	Store store	Date date	Item SKU sku	Units units	Sales sales
738	1	05/16/12	A96	2	6
738	1	05/16/12	65B	1	2.25
709	2	05/15/12	CB7	1	1.65
709	2	05/15/12	96A	1	1.1

The 'Arrange columns' panel on the left shows a list of columns: Account, Store, Date, Item SKU, Units, Sales, Transaction ID, and Cost. The 'Transaction ID' and 'Cost' columns are currently in the 'Hidden' section.

Sorting rows

You may reorder the rows of a table such that the values in a particular column are in ascending or descending order.

You must be viewing multiple rows of a table in order to perform a sort. You will not be able to sort if you are only viewing one row at a time.

To reorder the table based on more than one column, simply sort by one column and then by the others. For instance, sometimes you may want to sort the table so that one column is the primary sort "key" and the second column is the secondary key. In this case, the values in the first column would be sorted across all the rows of the table, and the values in the second column would be sorted in order *only for those rows whose values in the first column are the same*. To get this effect, sort first by the *second* column and then by the first.

The 'Sort columns' panel shows a list of columns: Transaction ID, Account, Store, Date, Item SKU, Units, Sales, and Cost. The 'Transaction ID' column is currently selected for sorting.

Figure 38: Sort columns panel

Show with column name

By default, columns are identified by the column label. Select this option to display the column name instead of the label.

Columns

The **Columns** section shows the columns that are currently visible in your table, in the order they are displayed from left to right.

Order

The **Order** section shows the column or columns by which the table is sorted.

Use the timeline to sort rows

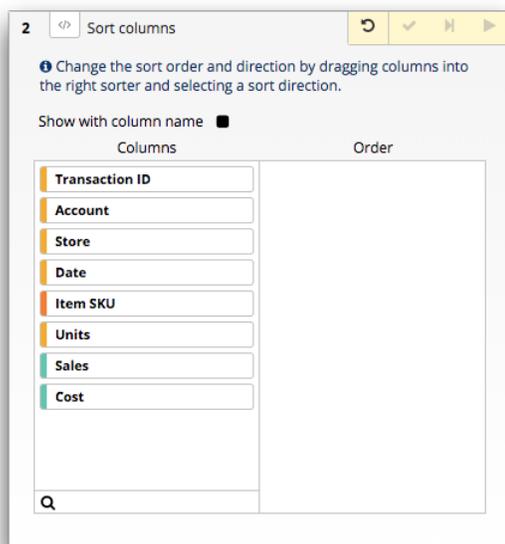
Sort the table according to one or more columns in ascending or descending order.

Change the sort order and direction by dragging columns into the right sorter and selecting a sort direction.

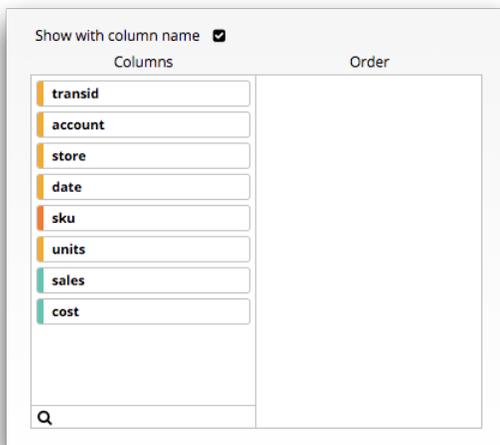
To use the timeline to sort rows:

1. In the **New operation** panel, click **Sort**.

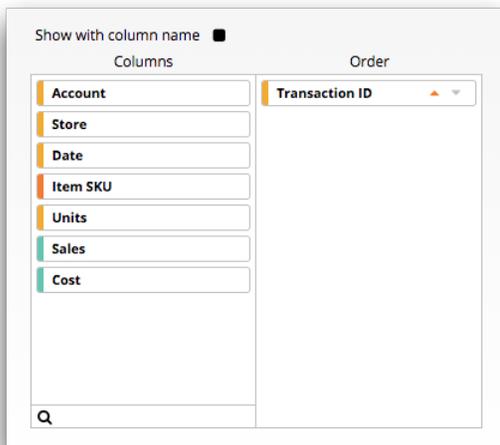
The Trillion-Row Spreadsheet displays the **Sort columns** panel.



2. Optionally, to view the column name instead of the column label, select the **Show with column name** option.



3. Drag the column on which you want to base the sort from the **Columns** section to the **Order** section. The Trillion-Row Spreadsheet displays sorting arrows within the ordered column.



4. Click one of the following icons:

Option	Description
Sort Up (^)	Sorts the rows in the table such that the values in the column are in ascending order.
Sort Down (v)	Sorts the rows in the table such that the values in the column are in descending order.

5. Optionally, to reorder the table based on more than one column, drag additional columns to the **Order** field and then define the sort order for each column.
6. Click the **Submit operation** (✓) icon. The Trillion-Row Spreadsheet sorts the table based on the column in ascending or descending order.

Sort rows from the grid column header

Use the grid column header menu to easily sort the table based on a single column in ascending or descending order.

You may reorder the rows of a table such that the values in a particular column are in ascending or descending order.

To sort rows from the grid column header:

1. In the Grid view of an open table or worksheet, place the pointer in the heading of the column on which you want to base the sort.
The Trillion-Row Spreadsheet displays a drop-down arrow.

Date	Item SKU	Units	Sales	Cost
05/15/12	366			
05/15/12	98A			
05/15/12	3B7			
05/16/12	A96			
05/16/12	65B			
05/15/12	CB7			
05/15/12	96A			
05/15/12	969	1	1.1	1

2. Click the drop-down arrow (▼) and then click one of the following:

Option	Description
Sort up	Sorts the rows in the table such that the values in the column are in ascending order.
Sort down	Sorts the rows in the table such that the values in the column are in descending order.

The Trillion-Row Spreadsheet sorts the table based on the column in ascending or descending order and adds the sort operation to the Analysis Timeline.

Sort rows within the grid

Use the right-click menu in the grid to easily sort the table based on a single column in ascending or descending order.

Reorder the rows of a table such that the values in a particular column are in ascending or descending order.

To sort rows within the grid:

1. In the Grid view of an open table or worksheet, right-click a cell in the column on which you want to base the sort.
2. From the menu, select one of the following:

Option	Description
up	Sorts the rows in the table such that the values in the column are in ascending order.
up with N/As at end	Sorts the rows in the table such that the values in the column are in ascending order. Rows with N/A (missing) values are sorted such that they are located at the bottom of the worksheet instead of at the top.
down	Sorts the rows in the table such that the values in the column are in descending order.
down with N/As at beginning	Sorts the rows in the table such that the values in the column are in descending order. Rows with N/A (missing) values are sorted such that they are located at the top of the worksheet instead of at the bottom.

The Trillion-Row Spreadsheet sorts the table based on the column in ascending or descending order and adds the sort operation to the Analysis Timeline.

Arranging columns

You can display a subset of the columns in a table and specify the order in which they appear.

You can rearrange, hide, and show columns. This changes only how your worksheet is displayed and does not affect the original table in any way.

Note: Fixed columns will always appear before all other columns.

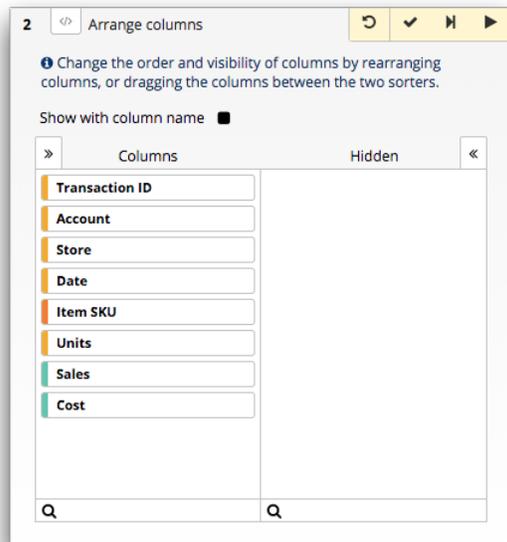


Figure 39: Arrange columns panel

Show with column name

By default, columns are identified by the column label. Select this option to display the column name instead of the label.

Hide All

Click the **Hide All** (») icon to move all of the columns to the **Hidden** section.

Columns

The **Columns** section shows the columns that are currently visible in your table, in the order they are displayed from left to right, including computed columns and columns that have been linked in from other tables or worksheets.

Hidden

The **Hidden** section shows any additional columns contained in your worksheet that are not currently displayed.

Show All

Click the **Show All** icon («) to move all of the columns to the **Columns** section.

Change the column order

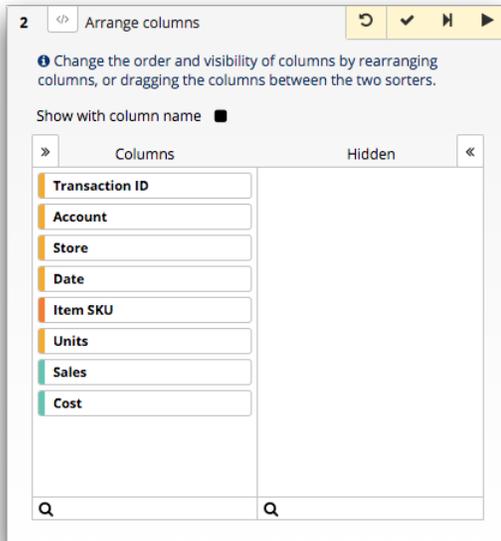
You can specify the order in which columns appear in your table or worksheet.

Using this panel, you can rearrange and hide columns in your table or worksheet.

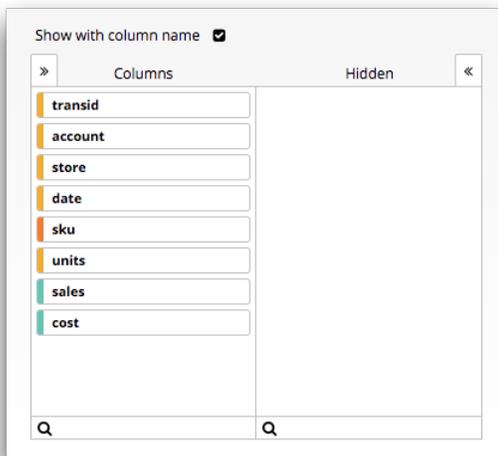
To change the column order:

1. In the **New operation** panel, click **Arrange**.

The Trillion-Row Spreadsheet displays the **Arrange columns** panel.



2. Optionally, to view the column name instead of the column label, select the **Show with column name** option.



3. In the **Columns** section, drag the column above or below other columns to reorder it in the table.
The order in which columns appear from top to bottom in the **Columns** section is the order the columns are displayed from left to right in the grid.
4. Click the **Submit operation** (✓) icon.
The Trillion-Row Spreadsheet rearranges the columns in the table.

Rearrange columns using the grid

You can drag columns within the grid to rearrange the order in which they appear.

To rearrange columns using the grid:

In the Grid view of an open table or worksheet, click the heading of a column and drag it to the position you want the column to be displayed.

As you drag the column heading, the Trillion-Row Spreadsheet displays a line indicating the location the column will be placed.

Sales Item Detail
Cols 1 to 8 of 8 , Rows 1 to 19 of 35

Transaction ID transid	Account account	Store store	Date date	Cost cost	Item SKU ku	Units units	Sales sales	Cost cost
531	957	1	05/15/12		98A	-1	-5	-1.84
532	478	1	05/15/12			1	0.5	0.25
532	478	1	05/15/12		3B7	1	1.1	0.56
534	738	1	05/16/12		A96	2	6	2.9
534	738	1	05/16/12		65B	1	2.25	1.35
535	709	2	05/15/12		CB7	1	1.65	1.1
535	709	2	05/15/12		96A	1	1.1	1
535	709	2	05/15/12		969	1	1.1	1
536	748	2	05/17/12		3A4	3	1.02	0.39
536	748	2	05/17/12		366	1	5	1.8
537	523	3	05/15/12		CB7	1	1.65	1.1
537	523	3	05/15/12		A96	2	6	2.9
537	523	3	05/15/12		98A	1	0.5	0.25
537	523	3	05/15/12		96A	1	1.1	1
537	523	3	05/15/12		3B7	1	1.1	0.56
538	668	1	05/18/12		CB7	1	1.65	1.1
538	668	1	05/18/12		98A	4	2	1
538	668	1	05/18/12		96A	1	1.1	1

The Trillion-Row Spreadsheet moves the column and adds the arrange operation to the Analysis Timeline.

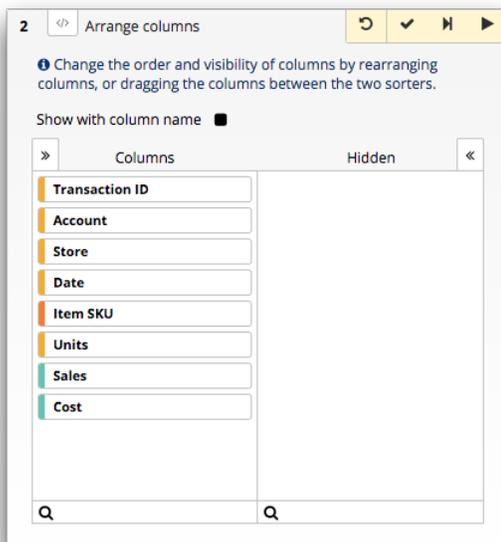
Use the timeline to hide or show columns

You can hide or show hidden columns in your table.

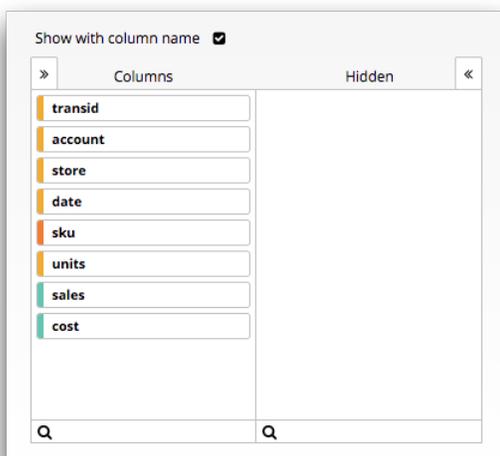
You can remove columns from view in a table, which can be particularly helpful for tables that have many columns. In addition, you can choose to show previously hidden columns.

To hide or show columns:

1. In the **New operation** panel, click **Arrange**.
The Trillion-Row Spreadsheet displays the **Arrange columns** panel.



- Optionally, to view the column name instead of the column label, select the **Show with column name** option.



- Perform any of the following actions to customize your table view:

Option	Description
To hide a column	Drag the column from the Columns section to the Hidden section.
To show a hidden column	Drag the column from the Hidden section to the Columns section.
To hide all columns	Click Hide All .
To show all columns	Click Show All .

- Click the **Submit operation** (✓) icon.
The Trillion-Row Spreadsheet hides or shows the appropriate columns in the table.

Hide a column within the grid

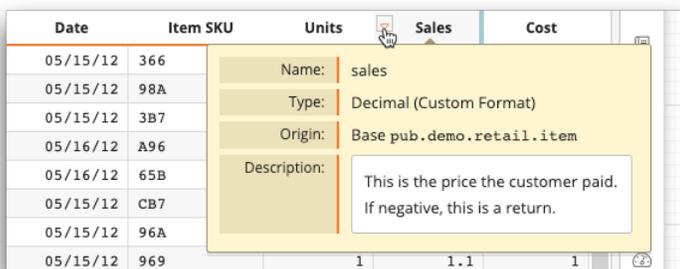
You can hide a column directly from within the grid.

In addition to the **Arrange** panel in the Trillion-Row Spreadsheet, you can also hide a column directly from the grid. This is an easy way to customize the results that are shown in the grid without needing to manually create an arrange operation in the Analysis Timeline.

Note: You cannot show a hidden column from within the grid. To show a hidden column, see [Use the timeline to hide or show columns](#) on page 161.

To hide a column within the grid:

- In the Grid view of an open table or worksheet, place the pointer in the heading of the column you want to hide.
The Trillion-Row Spreadsheet displays a drop-down arrow.



- Click the drop-down arrow (▾) and then click **Hide column**.
The Trillion-Row Spreadsheet hides the column and adds the arrange operation to the Analysis Timeline.

Linking tables and worksheets

Use linking to combine tables or worksheets into a single, bigger table with columns from both.

When you work a lot with data, you will often find yourself in a position where you want to use two separate tables, which contain different but related data, to answer a question. The 1010data Insights Platform has a highly effective feature that allows you join the information from related data tables together in a straightforward manner. It is called linking, and it is one of the things the Insights Platform does best.

When you link tables (or worksheets) together, you essentially combine them into a single, bigger table with all the columns from both. The change applies only to your current workspace and does not affect the original tables or other people in any way. For example, if one table contains addresses and telephone numbers for a group of people and another table contains their ages, heights, and weights, linking the two tables results in a table that has all the information about each person.

Linking is quite important for analysis. For example, using the first table described above, you could easily find all the people who live in a given city. Similarly, with the second table you could find all the people older than a given age. Both are simple row selections. But suppose you want to find all people older than a given age who also live in a given city? Link the tables and it's just another row selection.

Note: When you link to a worksheet (a table with operations applied to it), the full worksheet is pulled into memory, so there is a limit on how large the worksheet can be (roughly about 250,000,000 cells). When you link to a table, no such restriction applies.

If, for a particular row in the table that you are viewing, there is no information in the foreign table, N/A values are filled in for that row. In many cases, you may not be interested in such rows; after all, they have a significant amount of missing information. One way to eliminate such rows from the analysis is to do the link and then select only those rows that have non-N/A values in the columns you have linked in. Another, more efficient, method is to *link-and-select*, which combines the link and the selection in one operation. The result is the same as a simple link, except that only those rows for which there are matches in the foreign table are linked in.

Note that this is another way of doing row selections and may be useful in certain circumstances. Suppose, for example, that you are looking at a **Sales** table that has millions of rows involving thousands of customer IDs. Now suppose that you have a list of a few hundred customer IDs that you would like to select for analysis and that this list is in a text file or Excel spreadsheet. An easy way of proceeding in the Insights Platform would be:

- Upload the list of customer IDs as a new table.
- Go to the **Sales** table and perform a link-and-select to your uploaded table, linking on the customer ID.

Your view of the **Sales** table will now reflect only those customers in your list.

The result of a link is a dynamic table that may be manipulated just like any other table. You may select rows, sort rows, rearrange columns, create computed columns, link in other tables, perform summarizations or tabulations, or apply any other kind of analysis.

Link tables panel

Descriptions of user interface elements that are available when linking tables and worksheets in the Trillion-Row Spreadsheet and Macro Language Workshop.

The **Link tables** panel is accessible from the following locations in the 1010data Insights Platform:

- Analysis Timeline in the Trillion-Row Spreadsheet
- Right-click menu in the results grid of the Trillion-Row Spreadsheet and Macro Language Workshop.

It is important to note that the **Link tables** panel available in the grid is visually different from the panel in the Analysis Timeline. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the **Link tables** panel as it appears in the timeline.

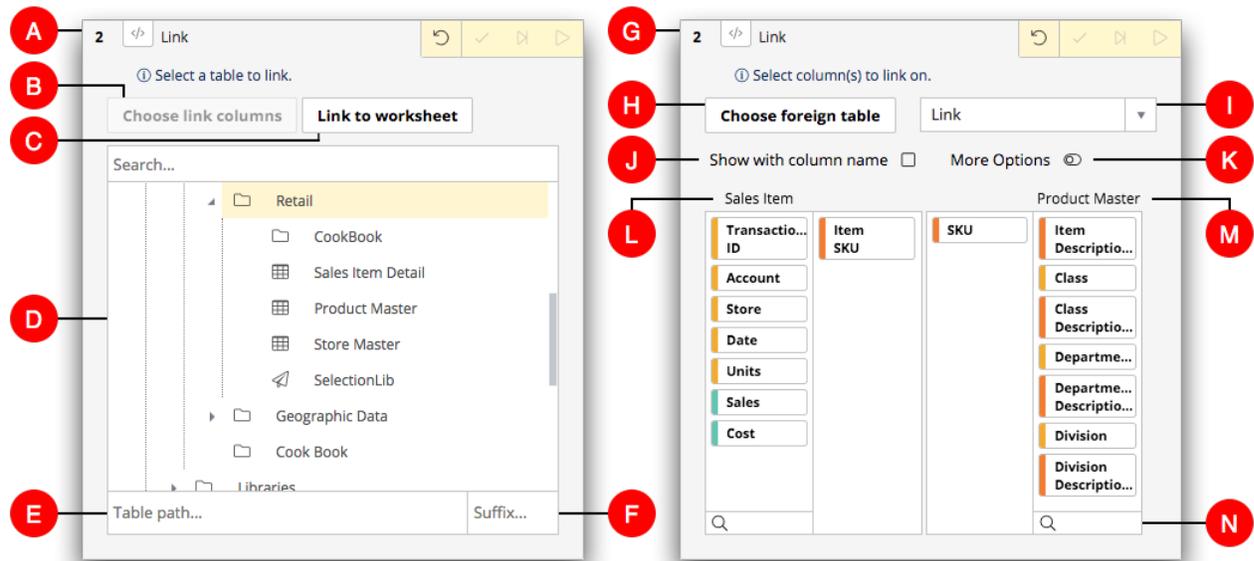


Figure 40: Link tables panel in the Trillion-Row Spreadsheet

A. Object browser view

In the image above on the left, the **Link tables** panel is open in the object browser view.

The object browser view of the **Link tables** panel appears after clicking **Link** in the **New operation** panel. It also appears after clicking **Choose foreign table** in the column view.

Use this view to select the foreign table or worksheet that you want to link into the current table or worksheet.

B. Choose link columns

This button is available after clicking **Choose foreign table** in the column view of the **Link tables** panel.

Click this button to return to column view of the **Link tables** panel without selecting a different table or worksheet.

C. Link to worksheet

This button displays a list of existing worksheets currently open in your session. The worksheet can be in either a TRS or a MLW window. This button also provides an option for creating a new worksheet to use in the link operation.

Click this button to select either an existing or new worksheet to link into the current table or worksheet.

D. Object browser

The object browser is used to select the foreign table that you want to link into the current table or worksheet. It provides a hierarchy of folders and objects to which you have access. For more information, see [Object browser](#) on page 121.

Use the **Search** field in the object browser to find objects such as folders, tables, and queries. For more information, see [Search](#) on page 102.

E. Table path

The **Table path** field provides an alternate way of selecting the foreign table that you want to link into the current table or worksheet.

If you know the path and name of the foreign table, simply enter it in the **Table path** field and press **Enter**.

F. Suffix

The **Suffix** field is used to differentiate the column names from the two tables or worksheets that you are linking. In the resultant worksheet, the suffix will be appended to the column names of the foreign table or worksheet, which you are linking in, and is useful when linking tables or worksheets that contain similarly-named columns.

The suffix may be any alphanumeric value and may contain underscores.

G. Column view

In the image above on the right, the **Link tables** panel is open in the column view.

The column view of the **Link tables** panel appears after selecting the foreign table or worksheet that you want to link into the current table or worksheet.

Use this view to choose how the columns from each table or worksheet are matched.

H. Choose foreign table

This button displays the object browser view of the **Link tables** panel.

Click this button to select a different table or worksheet to link into the current table or worksheet.

I. Link type

This drop-down menu allows you to choose from the various types of links you can perform.

For more information, see [Link types](#) on page 167.

J. Show with column name

Select this option to display the column name instead of the column label within the column biscuits in the panel. Column labels are displayed by default in the column view of the **Link tables** panel.

K. More Options

This switch controls whether the column view or the more options view of the **Link tables** panel is displayed.

When the switch is off (to the left), the column view is displayed. When the switch is on (to the right), the more options view is displayed. By default, the switch in the column view is off. You can submit a link operation when either the column view or more options view is displayed.

The more options view allows you to include or exclude specific columns from the foreign table or worksheet in the resultant worksheet. For more information, see [Include or exclude column options](#) on page 166.

L. Current table or worksheet

This is the label of the current table or worksheet. The two sections under this label (the two on the left side of the panel) show the columns in the current table or worksheet.

Of these two sections, the one on the left displays the columns in the current table that are available for linking. The section directly to the right shows the columns in the current table that are matched to columns in the foreign table.

In the image above, the **Item SKU** column from the **Sales Item Detail** table (the current table) is linked to the **SKU** column from the **Product Master** table (the foreign table).

M. Foreign table or worksheet

This is the label of the foreign table or worksheet. The two sections under this label (the two on the right side of the panel) show the columns in the foreign table or worksheet.

Of these two sections, the one on the right displays the columns in the foreign table that are available for linking. The section directly to the left shows the columns in the foreign table that are matched to columns in the current table.

In the image above, the **SKU** column from the **Product Master** table (the foreign table) is linked to the **Item SKU** column from the **Sales Item Detail** table (the current table).

N. Column filter

The column filter fields limit the number of column biscuits that are displayed in the panel section. This is useful in locating a specific column within a table or worksheet that has many columns. For instructions, see [Filter a list of columns](#) on page 140.

Include or exclude column options

The more options view of the **Link tables** panel allows you to choose columns in the foreign table that you want to include in, or exclude from, the worksheet resulting from a link operation.

The more options view is accessed by clicking the **More Options** switch on the **Link tables** panel.

It is important to note that the more options view of the **Link tables** panel available in the grid is visually different from the panel in the Analysis Timeline. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the more options view of the **Link tables** panel as it appears in the timeline.

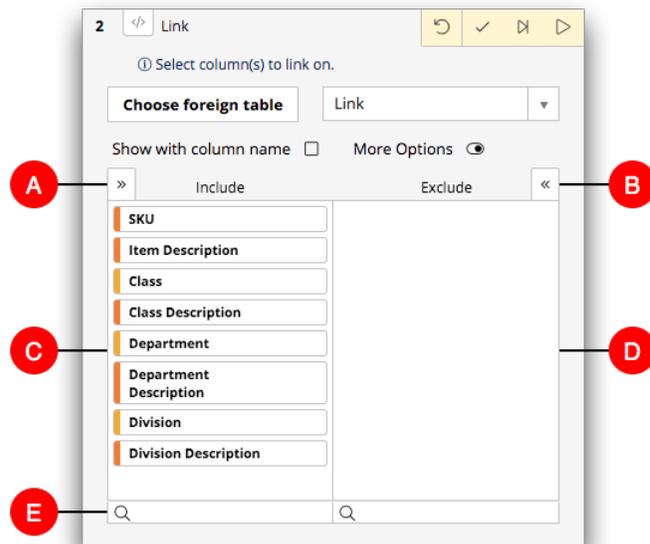


Figure 41: More options view of the Link tables panel

A. Move All (»)

The **Move All** icon (») in the **Include** section allows you to move all columns from the foreign table or worksheet to the **Exclude** section at once.

B. Move All («)

The **Move All** icon («) in the **Exclude** section allows you to move all columns from the foreign table or worksheet to the **Include** section at once.

C. Include section

This section displays the columns from the foreign table or worksheet that will be included in the resultant worksheet.

D. Exclude section

This section displays the columns from the foreign table or worksheet that will be excluded from the resultant worksheet.

E. Column filter

The column filter fields limit the number of column biscuits that are displayed in the panel section. This is useful in locating a specific column within a table or worksheet that has many columns. For instructions, see [Filter a list of columns](#) on page 140.

Link types

Descriptions of the link type options that are available when linking tables and worksheets in the Trillion-Row Spreadsheet.

The **Link type** drop-down list in the **Link tables** panel allows you to define the type of link you want to perform. Link types include the following:

Link

A row in the current table or worksheet is matched to the first row in the foreign table or worksheet that has exactly the same values in the linked columns. Column information from the foreign table or worksheet is appended to matching rows in the base table or worksheet. All rows in the resultant worksheet are retained.

Link and select

A row in the current table or worksheet is matched to the first row in the foreign table or worksheet that has exactly the same values in the linked columns. Column information from the foreign table or worksheet is appended to matching rows in the current table or worksheet. Only those rows in the resultant worksheet that have a match between the two are retained.

Link and exclude

A row in the current table or worksheet is matched to the first row in the foreign table or worksheet that has exactly the same values in the linked columns. However, column information from the foreign table or worksheet is not appended to matching rows in the base table or worksheet. Only those rows in the base table or worksheet that do not have a match between the two are retained.

Link and include

A row in the current table or worksheet is matched to the first row in the foreign table or worksheet that has exactly the same values in the linked columns. However, column information from the foreign table or worksheet is not appended to matching rows in the base table or worksheet. Only those rows in the base table or worksheet that have a match between the two are retained.

Link and expand

A row in the current table or worksheet is matched to all rows in the foreign table or worksheet that have exactly the same values in the linked columns. As necessary, a row in the current table or worksheet is duplicated so that it appears as many times as there are matches to that row in the foreign table or worksheet. Column information from the foreign table or worksheet is appended to matching rows in the current table or worksheet. All rows in the resultant worksheet are retained.

Link closest match before

The **Link closest match** options allow you to match rows in a less strict way when two tables are linked together. **Link closest match** can be performed only on date/time columns.

For a particular row in the current table, if no exact match can be made in the foreign table based on the value in the link column, the row in the current table or worksheet is matched to the first row in the foreign table or worksheet that has the *next* closest value to the row in the current table or worksheet. Column information from the foreign table or worksheet is appended to matching rows in the current table or worksheet. All rows in the resultant worksheet are retained.

If you are linking on more than one column (for example, product and date), only the last column (the date/time column) is linked with closest match. The other columns are linked exactly.

In the foreign table, the column that is linked with closest match must be sorted in ascending order within the remaining link columns, or it may be sorted in ascending order for the entire table. For example, if you link on product and date, then values can be sorted by date in ascending order for each product, or the entire table can be sorted in ascending order by date.

Link closest match after

For a particular row in the current table, if no exact match can be made in the foreign table based on the value in the link column, the row in the current table or worksheet is matched to the first row in the foreign table or worksheet that has the *previous* closest value to the row in the current table or worksheet. Column information from the foreign table or worksheet is appended to matching rows in the current table or worksheet. All rows in the resultant worksheet are retained.

If you are linking on more than one column (for example, product and date), only the last column (the date/time column) is linked with closest match. The other columns are linked exactly.

In the foreign table, the column that is linked with closest match must be sorted in ascending order within the remaining link columns, or it may be sorted in ascending order for the entire table. For example, if you link on product and date, then values can be sorted by date in ascending order for each product, or the entire table can be sorted in ascending order by date.

Link in a table

You can link in another 1010data Insights Platform table to the current table or worksheet.

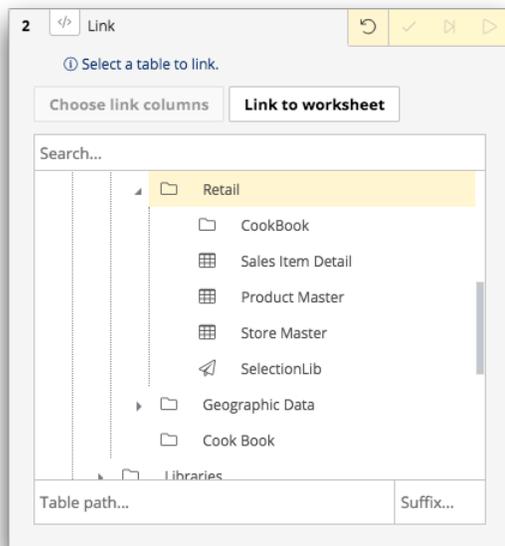
To link in a table:

1. Open the **Link tables** panel by doing one of the following:

- In the **New operation** panel, click **Link**.
- In the grid, right-click a cell within the column on which you want to base the link and click **Link on [COLUMN_LABEL]**.

Note: In this topic, *[COLUMN_LABEL]* represents the label of the column in which the selected cell is located.

The Trillion-Row Spreadsheet displays the object browser in the **Link tables** panel.



Note: The **Link tables** panel available in the grid is visually different from the panel in the Analysis Timeline. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the **Link tables** panel as it appears in the timeline.

2. Optionally, in the **Suffix** field at the bottom of the object browser, enter a suffix.

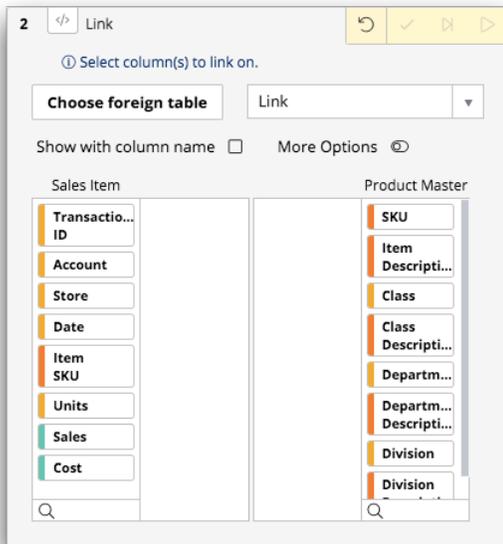
While not required, this field is highly recommended as it differentiates the column names from the two tables you are linking.

The suffix may be any alphanumeric value and may contain underscores. In the resultant worksheet, the suffix will be appended to the column names of the foreign table, which you are linking in, and is useful when linking tables or worksheets that contain similarly-named columns. For example, if you are viewing your **Sales Item Detail** table and want to link in your **Product Master** table, you might append a suffix of `_PM` to the columns from **Product Master** in the resultant table.

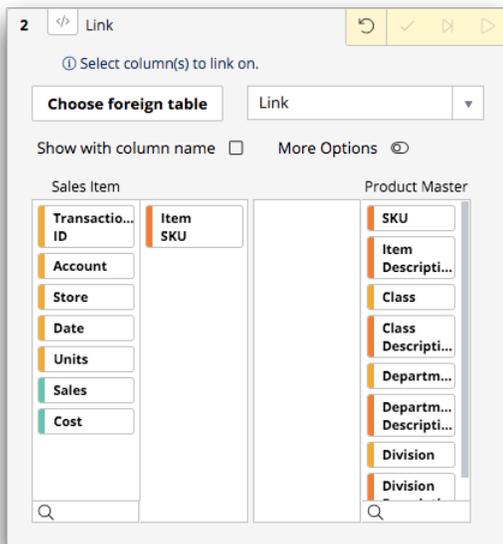
In this example, a column named `class` in the **Product Master** table is named `class_PM` in the resultant worksheet. Any reference to these columns (e.g., when creating a computed column) would use the column names with the suffix.

3. Select the foreign table by doing one of the following:
 - In the object browser, locate and select the foreign table.
 - In the **Table path** field, enter the table name and press **Enter**.

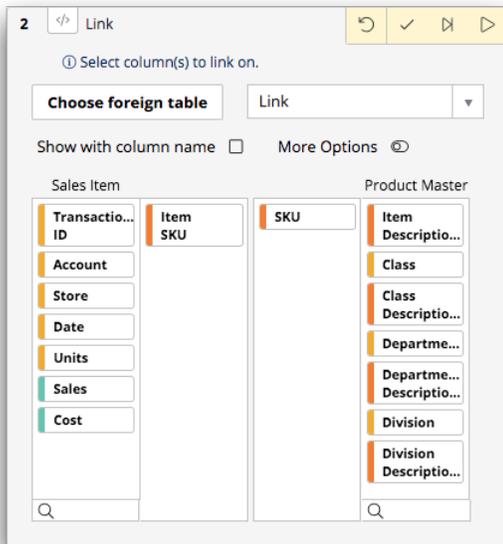
The Trillion-Row Spreadsheet displays the column view of the **Link tables** panel.



4. In the **Link type** drop-down list, select the type of link you want to perform.
For more information, see [Link types](#) on page 167.
5. From the first section, which correspond to and lists the current table or worksheet title, drag the column or columns you want to link on into the section directly to the right.
This is the column from the table or worksheet you currently have open.



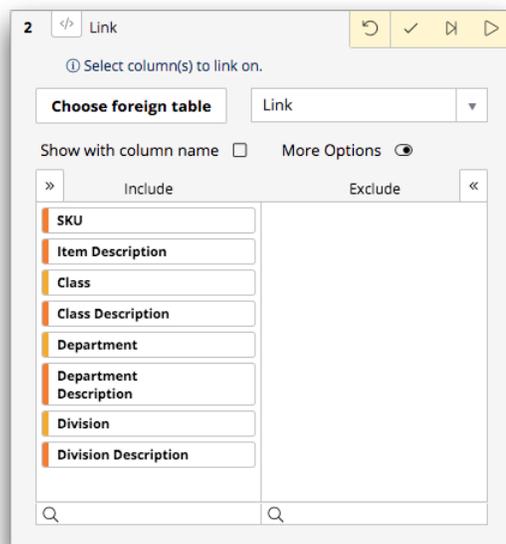
6. From the last section, which correspond to and lists the foreign table, drag the column or columns you want to match to those specified in the previous step to the section directly to the left.
This is the column from the foreign table.



7. Optionally, to include only specific columns from the foreign table in the resultant worksheet, complete the following:

a) Click the **More Options** switch.

The Trillion-Row Spreadsheet displays the more options view of the **Link tables** panel.



b) Move the columns you want to exclude from the **Include** section to the **Exclude** section.

For a list of fields and options, see [Include or exclude column options](#) on page 166.

8. Click the **Submit operation** (✓) icon.

The Trillion-Row Spreadsheet links the foreign table into the current table or worksheet and displays the columns from the two combined.

The result of a link is a dynamic table that may be manipulated just like any other table. You may select rows, sort rows, rearrange columns, create computed columns, link in other tables, perform summarizations or tabulations, or apply any other kind of analysis.

Link in an existing worksheet

You can link the resultant worksheet from a separate existing TRS or MLW window into the table or worksheet in your current TRS window.

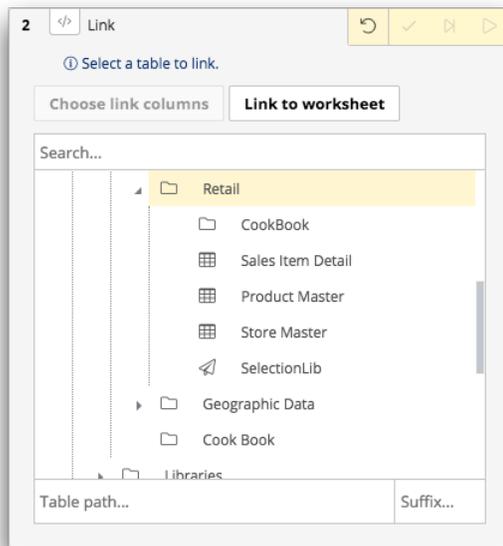
To link in an existing worksheet:

1. Open the **Link tables** panel by doing one of the following:

- In the **New operation** panel, click **Link**.
- In the grid, right-click a cell within the column on which you want to base the link and click **Link on [COLUMN_LABEL]**.

Note: In this topic, [COLUMN_LABEL] represents the label of the column in which the selected cell is located.

The Trillion-Row Spreadsheet displays the object browser in the **Link tables** panel.



Note: The **Link tables** panel available in the grid is visually different from the panel in the Analysis Timeline. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the **Link tables** panel as it appears in the timeline.

2. Optionally, in the **Suffix** field at the bottom of the object browser, enter a suffix.

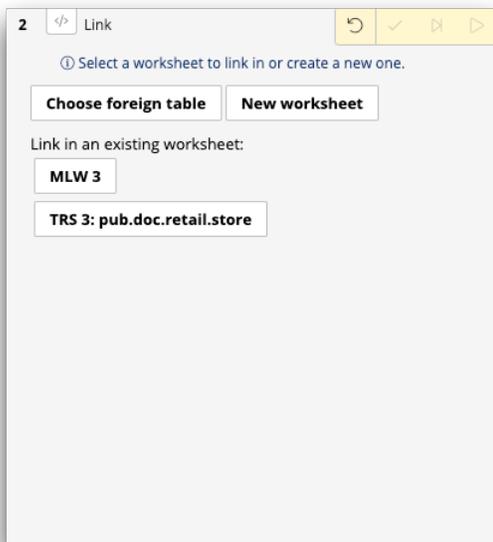
While not required, this field is highly recommended as it differentiates the column names from the two tables you are linking.

The suffix may be any alphanumeric value and may contain underscores. In the resultant worksheet, the suffix will be appended to the column names of the foreign worksheet, which you are linking in, and is useful when linking tables or worksheets that contain similarly-named columns. For example, if you are viewing your **Sales Item Detail** table and want to link in your **Product Master** table, you might append a suffix of `_PM` to the columns from **Product Master** in the resultant table.

In this example, a column named `class` in the **Product Master** table is named `class_PM` in the resultant worksheet. Any reference to these columns (e.g., when creating a computed column) would use the column names with the suffix.

3. Click **Link to worksheet**.

The Trillion-Row Spreadsheet displays options to link to a table, a new worksheet, or an existing worksheet.

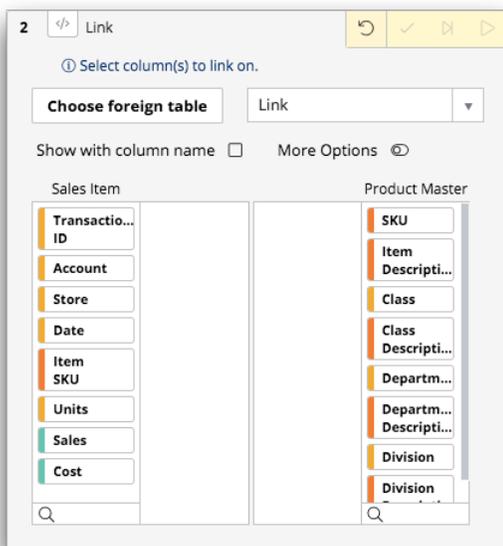


Note: If you do not have other TRS or MLW window(s) open, the Trillion-Row Spreadsheet instead displays the object browser in the **Link to worksheet** window. This new window allows you to link in a new worksheet. For instructions, see [Link in a new worksheet](#) on page 175.

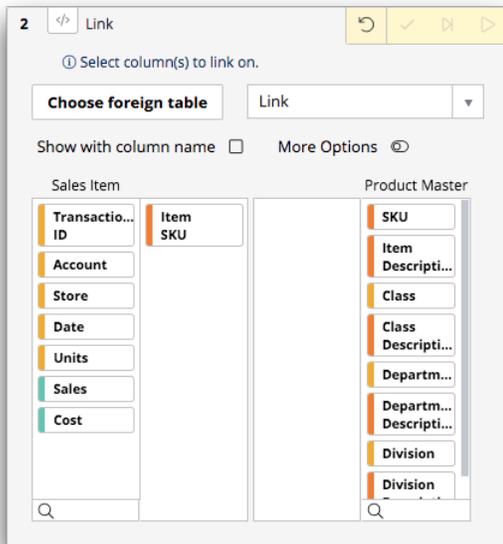
4. Below **Link in an existing worksheet**, click the existing worksheet you want to link in. The worksheet can exist in a TRS or a MLW window.

You can hover over the existing worksheet in the **Link** panel for help in locating the other worksheet in the workspace. The worksheet will be highlighted in the workspace and in the window tray at the bottom of the workspace.

The Trillion-Row Spreadsheet displays the column view of the **Link tables** panel.

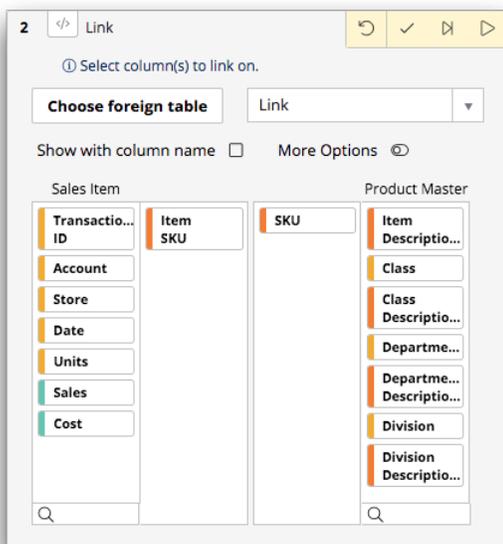


5. In the **Link type** drop-down list, select the type of link you want to perform. For more information, see [Link types](#) on page 167.
6. From the first section, which correspond to and lists the current table or worksheet title, drag the column or columns you want to link on into the blank section directly to the right. This is the column from the table you currently have open.



7. From the last section, which correspond to and lists the foreign worksheet title, drag the column or columns you want to match to those specified in the previous step to the blank section directly to the left.

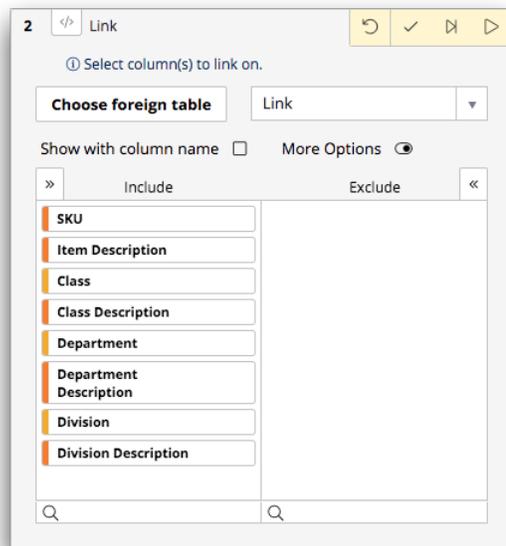
This is the column from the foreign worksheet.



8. Optionally, to include only specific columns from the foreign worksheet in the resultant worksheet, complete the following:

- a) Click the **More Options** switch.

The Trillion-Row Spreadsheet displays the more options view of the **Link tables** panel.



b) Move the columns you want to exclude from the **Include** section to the **Exclude** section.

For a list of fields and options, see [Include or exclude column options](#) on page 166.

9. Click the **Submit operation** (✓) icon.

The Trillion-Row Spreadsheet links the foreign worksheet into the current table or worksheet and displays the columns from the two combined.

The result of a link is a dynamic table that may be manipulated just like any other table. You may select rows, sort rows, rearrange columns, create computed columns, link in other tables, perform summarizations or tabulations, or apply any other kind of analysis.

Note: Changing the foreign worksheet *after* you link it in will not change the resultant worksheet.

Link in a new worksheet

You can open a new TRS window and perform various actions on a table or query. You can then link those results into the exiting table or worksheet in the current TRS window.

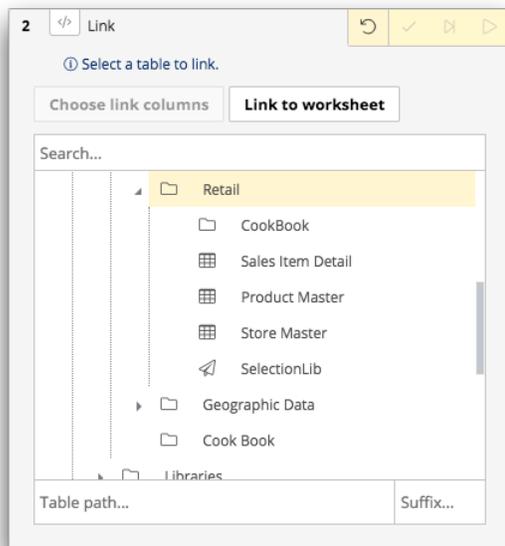
To link in a new worksheet:

1. Open the **Link tables** panel by doing one of the following:

- In the **New operation** panel, click **Link**.
- In the grid, right-click a cell within the column on which you want to base the link and click **Link on [COLUMN_LABEL]**.

Note: In this topic, *[COLUMN_LABEL]* represents the label of the column in which the selected cell is located.

The Trillion-Row Spreadsheet displays the object browser in the **Link tables** panel.



Note: The **Link tables** panel available in the grid is visually different from the panel in the Analysis Timeline. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the **Link tables** panel as it appears in the timeline.

2. Optionally, in the **Suffix** field at the bottom of the object browser, enter a suffix.

While not required, this field is highly recommended as it differentiates the column names from the two tables you are linking.

The suffix may be any alphanumeric value and may contain underscores. In the resultant worksheet, the suffix will be appended to the column names of the foreign worksheet, which you are linking in, and is useful when linking tables or worksheets that contain similarly-named columns. For example, if you are viewing your **Sales Item Detail** table and want to link in your **Product Master** table, you might append a suffix of `_PM` to the columns from **Product Master** in the resultant table.

In this example, a column named `class` in the **Product Master** table is named `class_PM` in the resultant worksheet. Any reference to these columns (e.g., when creating a computed column) would use the column names with the suffix.

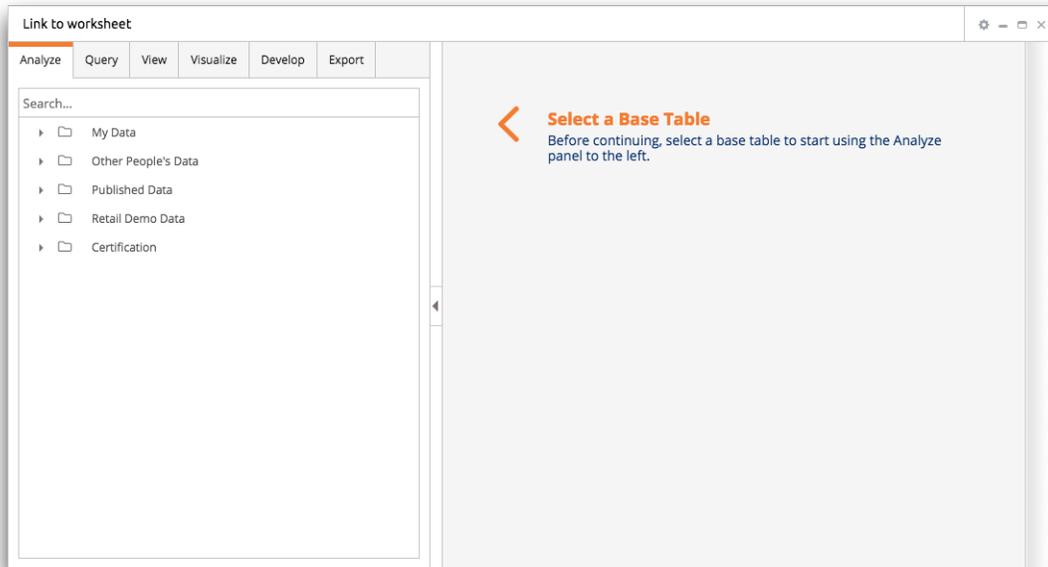
3. Click **Link to worksheet**.

The Trillion-Row Spreadsheet displays the object browser in the **Link to worksheet** window.

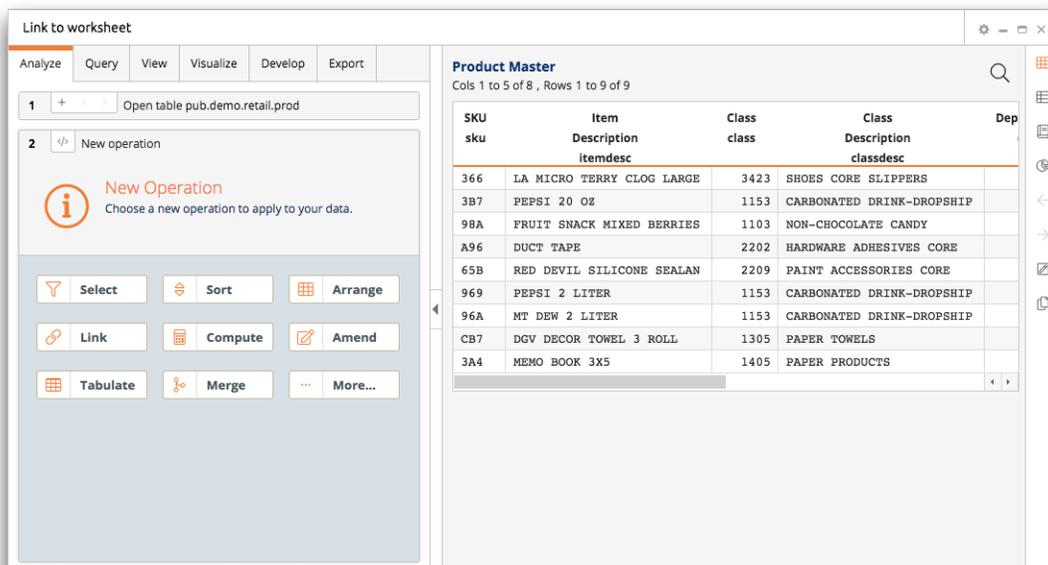
Note: If you have multiple TRS windows open, the Trillion-Row Spreadsheet instead displays options to link to a table, a new worksheet, or an existing worksheet.

4. If the Trillion-Row Spreadsheet displayed link options, click **New worksheet**.

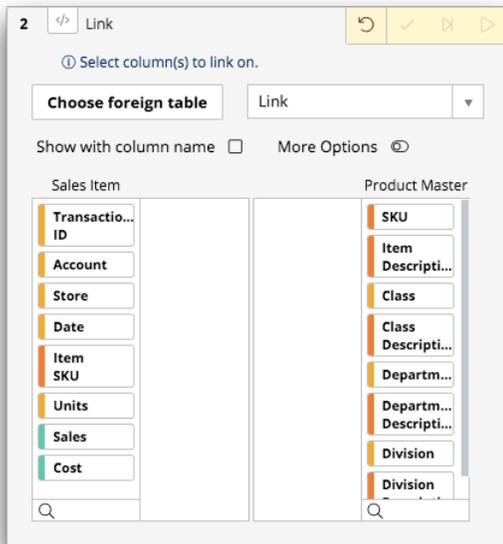
The Trillion-Row Spreadsheet displays the object browser in the **Link to worksheet** window.



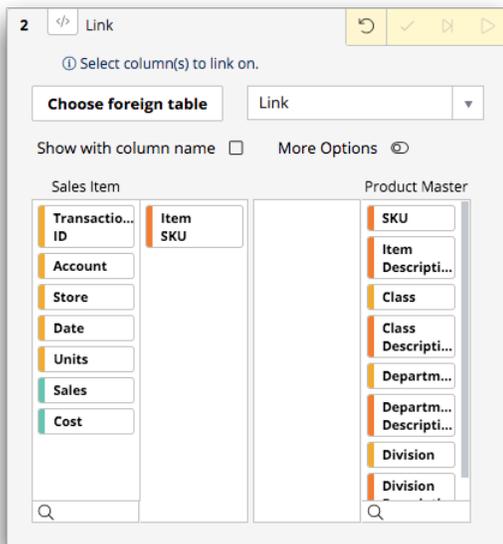
5. Navigate to and select the foreign table on which you want to create a new worksheet. The Trillion-Row Spreadsheet displays the **New Operation** panel in the **Analyze** tab on the left of the window and the table on the right.



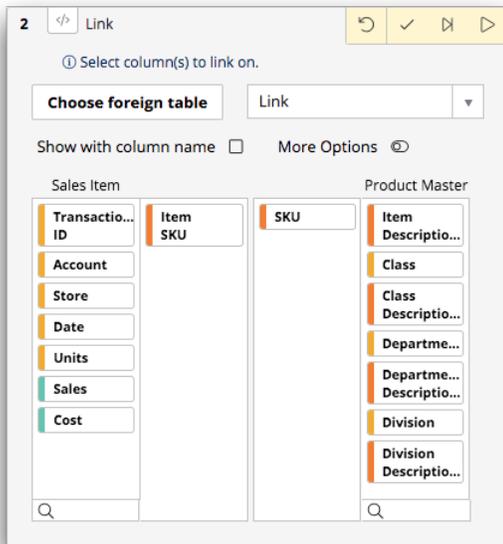
6. After performing the desired analysis, return to the **TRS** window that contains the base table you want to link the new worksheet into. The Trillion-Row Spreadsheet displays the column view of the **Link tables** panel.



7. In the **Link type** drop-down list, select the type of link you want to perform.
For more information, see [Link types](#) on page 167.
8. From the first section, which correspond to and lists the current table or worksheet title, drag the column or columns you want to link on into the blank section directly to the right.
This is the column from the table you currently have open.



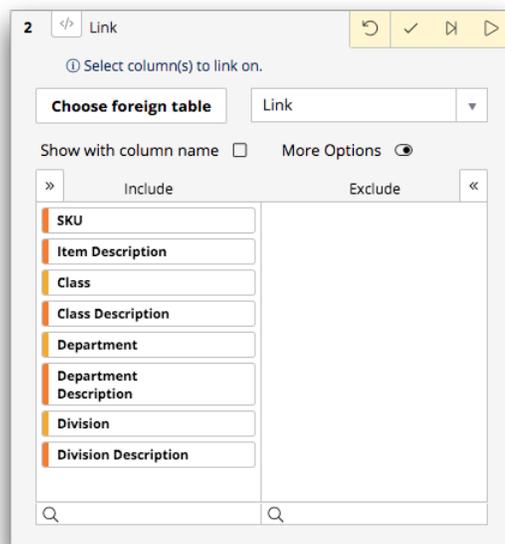
9. From the last section, which correspond to and lists the foreign worksheet title, drag the column or columns you want to match to those specified in the previous step to the blank section directly to the left.
This is the column from the foreign worksheet.



10. Optionally, to include only specific columns from the foreign worksheet in the resultant worksheet, complete the following:

a) Click the **More Options** switch.

The Trillion-Row Spreadsheet displays the more options view of the **Link tables** panel.



b) Move the columns you want to exclude from the **Include** section to the **Exclude** section.

For a list of fields and options, see [Include or exclude column options](#) on page 166.

11. Click the **Submit operation** (✓) icon.

The Trillion-Row Spreadsheet links the foreign worksheet into the current table or worksheet and displays the columns from the two combined.

The result of a link is a dynamic table that may be manipulated just like any other table. You may select rows, sort rows, rearrange columns, create computed columns, link in other tables, perform summarizations or tabulations, or apply any other kind of analysis.

Note: Changing the foreign worksheet *after* you link it in will not change the resultant worksheet.

Computed column

A computed column is a column that you add to a table, whose values are typically based on one or more existing columns.

A computed column is determined by a given value expression. The value expression may refer to one or more columns and may include standard arithmetic, relational, and logical operators.

To create a column in the Trillion-Row Spreadsheet, use the Expression Editor in the **Computed column** panel. For more information about the Expression Editor, see [Expression Editor](#) on page 139.

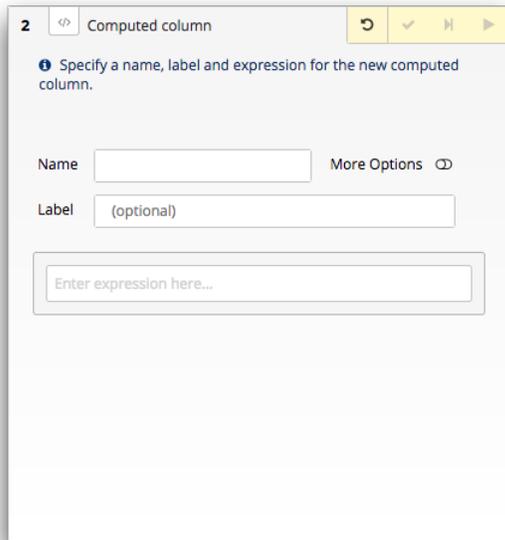


Figure 42: Computed column panel

For example, suppose a table has the columns **Sales** and **Cost**. You can create a new column **Margin**, which is **Sales** minus **Cost**.

In the Expression Editor, the value expression to create this computed column would be:

```
sales-cost
```

Note: Column names (e.g., `sales`), not column headings (e.g., **Sales**) are used in value expressions.

You can also create a computed column using the 1010data Insights Platform Macro Language XML code. Macro Language code can be entered in the Macro Language view of the **Computed column** panel and in the Macro Language Workshop.

In the Macro Language, the code to create this same computed column would be:

```
<willbe name="margin" value="sales - cost" label="Margin"/>
```

In addition to simple arithmetic operators, there are more advanced mathematical functions, functions that perform mathematical operations on dates (e.g., find the number of days between the dates in two columns), functions to manipulate strings, and categorization functions (e.g., is the value in the column between 0 and 10, between 10 and 100, or greater than 100?).

For example, to find the sum of sales for each store in the **Sales Item Detail** table, you can create a computed column that uses the `G_Function` `g_sum(G;S;X)`. In Macro Language, the code to create this computed column would be:

```
<willbe name="sum_of_sales" value="g_sum(store;;sales)" label="Sum of Sales"/>
```

In this example, *G* is the column we're grouping on (in our example, *store*); *S* is a selection column (by omitting a value for this, we're telling the *G_Function* to select all rows when computing the sum); and *x* is the column we are acting on (in our example, we want the sum of *sales*).

An important feature of computed columns is that missing ("N/A") values are handled automatically. For example, if **Sales** is N/A on one or more rows, then **Margin (Sales minus Cost)** will also be N/A on those rows. See the description of each operator and function for details about the way it handles N/A values.

Create a computed column

Use a computed column to add a column of information to a table or worksheet typically using the information from one or more existing columns.

Computed columns are columns you create yourself in many cases using information that is already in the table with which you are working. For example, suppose you have a table that contains all the sales data for a chain of stores. This table has one column called **Sales**, which is the purchase price paid by the customer. It has another column called **Cost**, which is the cost of the item to the retailer. With a computed column, it is easy to create a third column called **Margin** that contains the difference between the sale price and the cost.

To create a computed column:

1. In the **New operation** panel, click **Compute**.

The Trillion-Row Spreadsheet displays the **Computed column** panel.

2. Complete the following fields:

Name

Enter the name that you want to assign to the computed column.

This is the name the 1010data Insights Platform uses to interact with the column. It is also used when writing more advanced value expressions and queries.

The column name must be unique. It may only contain alphanumeric characters or underscores and must begin with an alphabetic character (e.g., `percent_total_sales`). It may not contain any spaces or other special characters.

This is a required field.

Note: When entered, the Trillion-Row Spreadsheet (TRS) automatically removes special characters, replaces spaces with underscores, and displays a visual indication that a change was made.

When TRS makes these changes, and if nothing has been entered in the **Label** field, the original text is automatically copied into the **Label** field.

Label

Enter the text that you want to appear at the top of the computed column when the table is displayed.

This is the column heading that displays by default at the top of a column in the user interface.

The column label may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters. If you want to have a multi-line column label, use the backtick character (`) to separate the lines (e.g., "Percentage of `Total Sales (%)").

Note: On most standard United States keyboards, the **backtick** (`) character is immediately to the left of **1** key.

While not required, this field is recommended.

3. In the Expression Editor field, enter a value expression.

The value expression may refer to one or more columns and may include standard arithmetic, relational, and logical operators as well as any of the Insights Platform's functions.

Note: When entering a column in an expression, use the column *name* as opposed to the column *heading*.

For more information, see [Writing Expressions](#) in the *1010data Reference Manual*. For more information about the Expression Editor, see [Expression Editor](#) on page 139.

4. Optionally, to define the format of the computed column, complete the following:

- a) Click the **More Options** switch.

The Trillion-Row Spreadsheet displays the formatting options in place of the Expression Editor.

The screenshot shows a configuration panel for a column. At the top, there is a 'Name' input field containing the text 'margin' and a 'More Options' toggle switch that is currently turned on. Below this is a 'Label' input field containing the text 'Margin'. Underneath the label field are three drop-down menus: 'Display Format' (set to 'Default'), 'Column Width' (set to 'Default'), and 'Decimal Places' (set to 'Default').

- b) As necessary, select the appropriate option from the following drop-down lists:

Display Format

This drop-down list instructs the Trillion-Row Spreadsheet how to display the value in the computed column. For example, as a date, a number, as text, and so forth.

Note: If set to **Default**, the format is determined by the data.

Column Width

This drop-down list specifies how many place values to hold in the column. In other words, the width of the column by the number of characters.

You can select from **1** to **20** characters.

Note: If set to **Default**, the width is determined by the data.

Decimal Places This drop-down list sets how many numbers are displayed after the decimal point.

You can select from **0** to **9** decimal places.

Note: If set to **Default**, the number of decimal places is determined by the data, and may be different for each row.

- Click the **Submit operation** (✓) icon.
The Trillion-Row Spreadsheet creates the computed column.

Summarizations and tabulations

Some of the basic ways to analyze your data include quick summaries, tabulations, and cross tabulations.

Oftentimes, you may want to determine the total of all the values in a particular column or the average of those values. Or perhaps you want to find the highest or lowest value in a column. A quick summary allows you to perform calculations, called summarizations, on one or more columns in a 1010data Insights Platform table to ascertain these quantities very easily.

Using tabulations, you can summarize the values in a column (or columns) based on the values in another column (or columns) and display those results in a tabular format.

For example, suppose a table has a row for every employee and columns **Name**, **Title**, **Department**, and **Salary**. Using a tabulation, you can compute the number of people and average salary for each department. The result may look something like the example below:

Department	Count	Average Salary
	228	73,946
Finance	35	61,778
Marketing	41	87,562
Sales	67	92,962
Engineering	34	69,276
Technology	51	49,481

The cells in the top row beneath the column headings are column aggregates. In this example, the total number of employees (which corresponds to the number of rows in the original table) is 228, and the average salary across all departments is 73,946.

A cross tabulation allows you to summarize the values in a column based on the values in two or more other columns and display the result as a matrix.

In the above example, you could use a cross tabulation to compute the average salary in each department for each title. This would result in a table with one row for each department and one column for each title:

	Title	VP	AVP	Manager	Grunt
Department	73,946	143,356	118,979	85,345	85,345
Finance	61,778	95,774	85,812	70,553	48,944
Marketing	87,562	127,344	127,664	100,359	70,686
Sales	92,962	155,863	125,268	100,945	74,314
Engineering	69,276	116,181	94,189	78,594	53,997

	Title	VP	AVP	Manager	Grunt
Technology	49,481	79,178	67,080	56,285	39,657

In this example, you can see that the average salary for AVPs in Engineering is 94,189. The cells in the first row beneath the column headings and the first column next to the row headings are aggregates. The first row of data contains the average salaries across all departments for each title (e.g., the average salary for all VPs is 143,356). Similarly, the first column of data contains the average salaries across all titles for each department (e.g., the average salary for all people in the Finance department is 61,778). The top left-most cell of data is the average for all employees across all departments and titles, which is 73,946.

Certain types of summarizations require a second column to perform the computation. For instance, a weighted average requires two columns: the column on which you are calculating the average and the column that contains the weight. Other examples of summarizations that involve two columns of data include the correlation, dot product, or covariance of two columns.

Missing values in the data are handled automatically in the Insights Platform. For example, if you are computing the average value of a column that has some missing (N/A) values, those values are ignored when computing the average.

Quick summaries, tabulations, and cross tabulations can be performed easily using the web interface. In addition, the Macro Language XML code functions (in particular, G_Functions) provide a comprehensive set of tools for performing similar and more advanced operations on data in the Insights Platform. For instance, you can obtain the same information from the tabulation in the above example using a very simple G_Function. To find the average salary for each department, you could create a computed column and use the `g_avg` function as its value expression. The results of the cross tabulation could be determined in a similar way.

Tabulation panel

The **Tabulation** panel allows you to perform any of the three basic types of tabulations (quick summaries, tabulations, and cross tabulations) to analyze your data.

The **Tabulation** panel is accessible from the following locations in the 1010data Insights Platform:

- Analysis Timeline in the Trillion-Row Spreadsheet
- Right-click menu in the grid

It is important to note that the **Tabulation** panel available in the grid is visually different from the panel in the Analysis Timeline. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the **Tabulation** panel as it appears in the timeline.

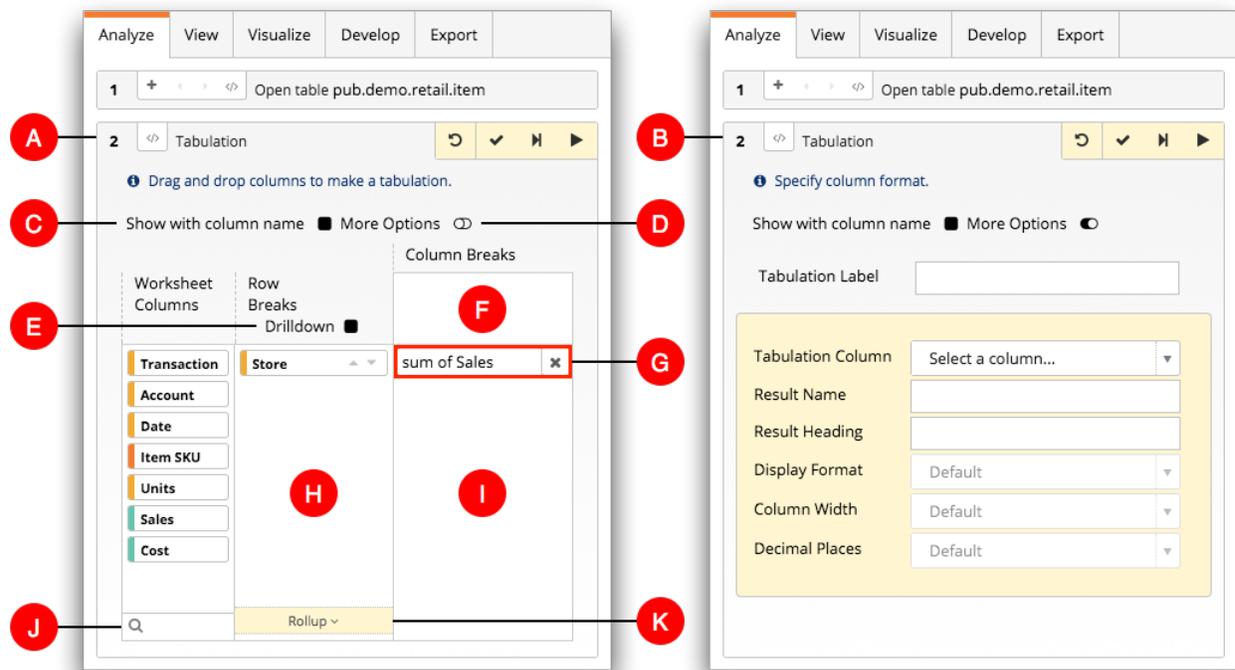


Figure 43: Tabulation panel

A. Edit panel view

In the image above on the left, the **Tabulation** operation is open in the edit panel view.

Use the edit panel view to configure a new tabulation or make changes to an existing tabulation in the timeline. By default, a new **Tabulation** panel opens in the edit panel view.

B. More options view

In the image above on the right, the **Tabulation** operation is open in the more options view. For a list of format fields and options, see [Tabulation result options](#) on page 207.

This view displays the options available to format the tabulation results. For example, you can give tabulated columns meaningful names and define the various format options of the tabulation. For instructions, see [Format the tabulation results](#) on page 205.

C. Show with column name

Select this option to display the column name instead of the column label within the column biscuits and the **Tabulation Function** menu in the panel. Column labels are displayed by default in a new **Tabulation** panel.

D. More Options

This switch controls whether the edit panel view or the more options view of the **Tabulation** panel is displayed.

When the switch is off (to the left), the edit panel view is displayed. When the switch is on (to the right), the more options view is displayed. By default, the switch in a new **Tabulation** panel is off. You can submit a tabulation operation when either panel view is displayed.

E. Drilldown

This option allows you to add a drilldown hierarchy to a tabulation or cross tabulation. The drilldown hierarchy is defined by the order of the specified columns in the **Row Breaks** section. After the

operation is submitted, you can use the Grid view to perform a drilldown analysis on the results. For more information, see [Drilldown analysis](#) on page 209.

F. Upper area of the Column Breaks section

The upper area of the **Column Breaks** section is used to create a cross tabulation.

When at least one column biscuit is in the **Row Breaks** section and at least one tabulation function is in the lower area of the **Column Breaks** section, drag at least one column biscuit to the upper area to perform a cross tabulation. For instructions, see [Perform a cross tabulation](#) on page 193.

G. Tabulation function

The tabulation function field defines the calculation, called a summarization, in a quick summary, tabulation, or cross tabulation. You can add multiple summarizations to any of the three types of tabulations.

Clicking the tabulation function field displays the **Tabulation Function** menu which allows you to select the type of summarization you want to perform. For a list of available tabulation functions, see [Types of summarizations](#) on page 199.

H. Row Breaks section

The **Row Breaks** section is used to create a tabulation.

Drag at least one column biscuit into this section to create a tabulation. For instructions, see [Perform a tabulation](#) on page 190.

I. Lower area of the Column Breaks section

The lower area of the **Column Breaks** section is used to set the tabulation function for the quick summary, tabulation, or cross tabulation.

Note: By default, a new **Tabulation** panel contains the **cnt** (count) tabulation function.

When a column biscuit is dragged into the lower area of the **Column Breaks** section, a tabulation function is created for that column. By default, columns placed in the lower area of the **Column Breaks** section are assigned the **sum of [COLUMN]** summarization. If the data type of the column is text, it is assigned the **last of [COLUMN]** summarization. To remove a tabulation function, click the **Delete Function** (✕) icon in the tabulation function field.

Drag at least one column biscuit to the lower area of the **Column Breaks** section, without placing any columns in either the **Row Breaks** or **Column Breaks** sections, to perform a quick summary. For instructions, see [Perform a quick summary](#) on page 186.

J. Column filter

The column filter field is used to limit the number of column biscuits that appear in the **Worksheet columns** section. This is useful when locating a specific column in a table or worksheet that has many columns. For instructions, see [Filter a list of columns](#) on page 140.

K. Rollup bar

The **Rollup** bar is used to group like values when performing a tabulation with row breaks. A rollup makes it easy to aggregate data at multiple levels in a tabulation. For instructions, see [Use a rollup in a tabulation](#) on page 204.

Note: The rollup feature is not available in a cross tabulation.

Perform a quick summary

A quick summary is a type of tabulation that allows you to easily perform calculations, called summarizations, on one or more columns.

Oftentimes, you may want to determine the total of all the values in a particular column or the average of those values. Or perhaps you want to find the highest or lowest value in a column. A quick summary

allows you to perform summarizations on one or more columns in a table or worksheet to ascertain these quantities.

To perform a quick summary:

1. Open the **Tabulation** panel by doing one of the following.

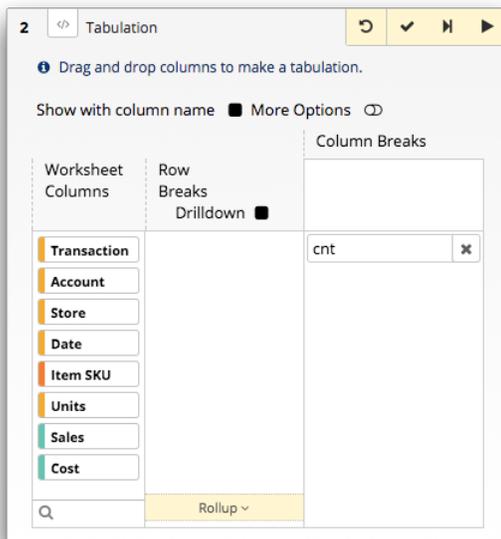
- In the **New operation** panel, click **Tabulate**.
- In the grid, right-click a cell within the column on which you want to base the quick summary, point to **Tabulate using [COLUMN_LABEL] > as a source for**, and select a summarization option.

For a list and description of options available from the grid, see [Summarization options in the grid menu](#) on page 189.

Note: In this topic, [COLUMN_LABEL] represents the label of the column in which the selected cell is located.

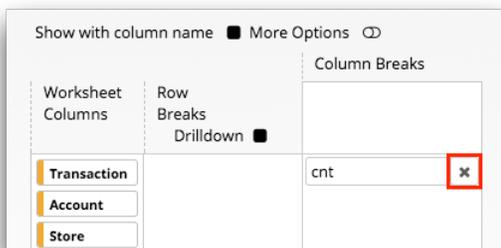
The Trillion-Row Spreadsheet displays the **Tabulation** panel.

Note: The **Tabulation** panel available in the grid is visually different from the panel in the **Analysis Timeline**. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the **Tabulation** panel as it appears in the timeline.



Note: By default, when opened from the **New operation** panel, the **Tabulation** panel contains the **cnt** (count) tabulation function. Otherwise, the **New operation** panel contains the summarization selected from the grid menu.

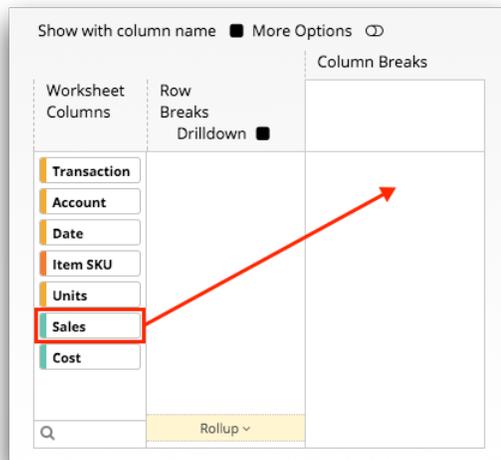
2. If applicable, remove the default tabulation function by clicking the **Delete Function** (**x**) icon in the field.



3. For each column of data that you want to summarize, complete the following:

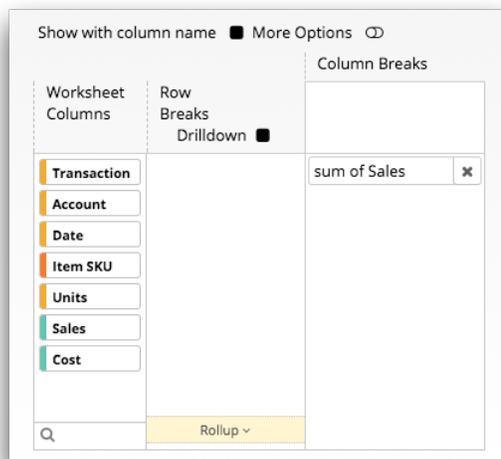
- a) Drag the column from the **Worksheet Columns** section to the lower area of the **Column Breaks** section.

This selection chooses the column of data you want to summarize.



The Trillion-Row Spreadsheet adds a tabulation function field and returns the column to the **Worksheet Columns** section.

Note: By default, columns placed in the lower area of the **Column Breaks** section are assigned the **sum of [COLUMN]** summarization. If the data type of the column is text, it is assigned the **last of [COLUMN]** summarization.



- b) If applicable, change the summarization.

For instructions, see [Change the tabulation function summarization](#) on page 197.

By default, columns in the resultant worksheet are named t0, t1, t2, and so on. You can give tabulated columns more meaningful names within the **More Options** view of the **Tabulation** panel. In addition, you can define the various formats of the tabulation results.

4. Optionally, define the format of the tabulation results.

For instructions, see [Format the tabulation results](#) on page 205.

5. Click the **Submit operation** (✓) icon.

The Trillion-Row Spreadsheet displays the results of your quick summary.

Sales	
qs_sum_sales	59.85

Summarization options in the grid menu

The grid menu provides a list of common calculations, called summarizations, from which to choose when creating a quick summary directly in the grid.

You can choose from the following types of summarizations when performing a quick summary in the grid.

Note: In the list below, *[COLUMN]* indicates either the column name or the column label of the selected column. By default, the column label is displayed. However, when the **Show with column names** option is selected in the **Tabulation** panel, the column name is displayed instead.

Transaction ID	Account	Store	Date	Item SKU	Units	Sales
transid	account	store	date	sku	units	sales
531	957	1	05/15/12	366	-1	-
532	47	Copy...			1	0.
532	47	Download...			1	1.
534	73	View this row in single-row mode			2	
534	73	Quick find...			1	2.
535	70	Quick select where...			1	1.
535	70	Quick select rows...			1	1.
535	70	Quick analysis of Account...			1	1.
536	74	Find in Account			3	1.
536	74	Sort on Account...			1	
537	52	Select on Account			1	1.
537	52	Link on Account			2	
537	52	Tabulate using Account...				
537	52	Amend Account				
537	52	Create computed column here			1	
537	52	Undo last step			1	1.
538	66				1	1.
538	668	1	05/18/12	98A	4	

Figure 44: Grid menu summarization options

sum

The subtotal of the column for each group. (*numeric columns only*)

- If a group contains N/A values, those rows are ignored (i.e., the result is the sum of the non-N/A values).
- If a group contains only N/A values, the result for the group is 0.

Note: This option adds the **sum of [COLUMN]** tabulation function field to the lower area of the **Column Breaks** section in the **Tabulation** panel.

average

The average of the column for each group. (*numeric columns only*)

- If a group contains N/A values, those rows are ignored (i.e., the result is the average of the non-N/A values).
- If a group contains only N/A values, the result for the group is N/A.

Note: This option adds the **avg of [COLUMN]** summarization to the lower area of the **Column Breaks** section in the **Tabulation** panel.

valid count

The number of non-N/A values in the column for each group.

Note: This option adds the **valcnt of [COLUMN]** summarization to the lower area of the **Column Breaks** section in the **Tabulation** panel.

N/A count

The number of N/A values in the column for each group.

Note: This option adds the **nacnt of [COLUMN]** summarization to the lower area of the **Column Breaks** section in the **Tabulation** panel.

unique count

The number of unique values in the column for each group.

- If a group contains N/A values, those rows are ignored (i.e., the result is the number of unique non-N/A values).
- If a group contains only N/A values, the result for the group is 0.

Note: This option adds the **ucnt of [COLUMN]** summarization to the lower area of the **Column Breaks** section in the **Tabulation** panel.

lowest value

The lowest number in the column for each group. (*numeric columns only*)

- If a group contains N/A values, those rows are ignored (i.e., the result is the lowest non-N/A value).
- If a group contains only N/A values, the result for the group is N/A.

Note: This option adds the **lo of [COLUMN]** summarization to the lower area of the **Column Breaks** section in the **Tabulation** panel.

highest value

The highest number in the column for each group. (*numeric columns only*)

- If a group contains N/A values, those rows are ignored (i.e., the result is the highest non-N/A value).
- If a group contains only N/A values, the result for the group is N/A.

Note: This option adds the **hi of [COLUMN]** summarization to the lower area of the **Column Breaks** section in the **Tabulation** panel.

Perform a tabulation

A tabulation summarizes large amounts of data into a small, easy-to-read table. Perform a tabulation to group the values in a column based on the values in another column.

A tabulation allows you to calculate, or summarize, various metrics, such as the sum, average, or highest value of a particular column, grouping those calculations based on common values in one or more other columns. For instance, in a table containing weather data, you could calculate the highest temperature during the year for every unique combination of ZIP Code and month in the table.

To perform a tabulation:

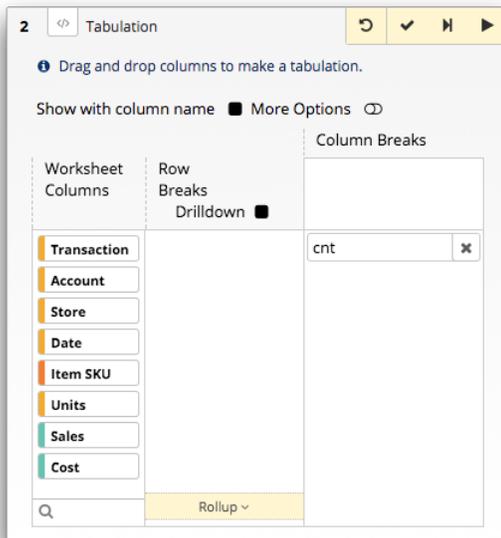
1. Open the **Tabulation** panel by doing one of the following.

- In the **New operation** panel, click **Tabulate**.
- In the grid, right-click a cell within the column on which you want to base the tabulation and, from the menu, select the **Tabulate using [COLUMN_LABEL] > as a break column** option.

Note: In this topic, [COLUMN_LABEL] represents the label of the column in which the selected cell is located.

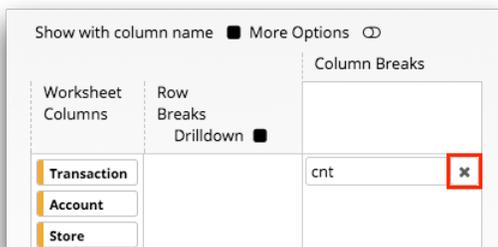
The Trillion-Row Spreadsheet displays the **Tabulation** panel.

Note: The **Tabulation** panel available in the grid is visually different from the panel in the Analysis Timeline. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the **Tabulation** panel as it appears in the timeline.



Note: By default, when opened from the **New operation** panel, the **Tabulation** panel contains the **cnt** (count) tabulation function. Otherwise, the **New operation** panel contains just the column of the selected cell in the **Row Breaks** section.

2. If applicable, remove the default tabulation function by clicking the **Delete Function** (**x**) icon in the field.



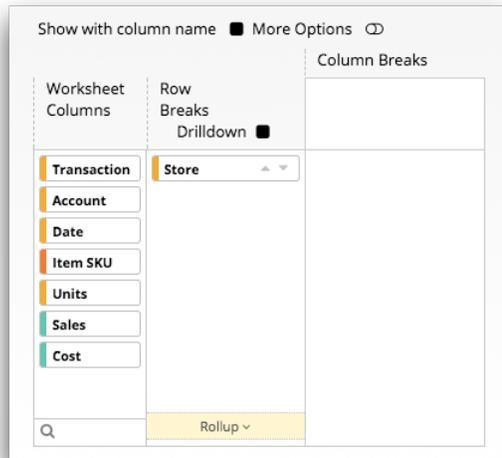
3. For each column by which you want to group the records, complete the following:

- a) Drag the column from the **Worksheet Columns** section to the **Row Breaks** section.

This selection groups the data by creating a separate row in the resultant worksheet for each unique value within the chosen column. Grouping is a way of pooling all the records for a single entity or value into a single entry in the worksheet.

Note: In most cases, the order by which you choose to group the data is important because it can affect the results of the tabulation.

The Trillion-Row Spreadsheet moves the column to the **Row Breaks** section.



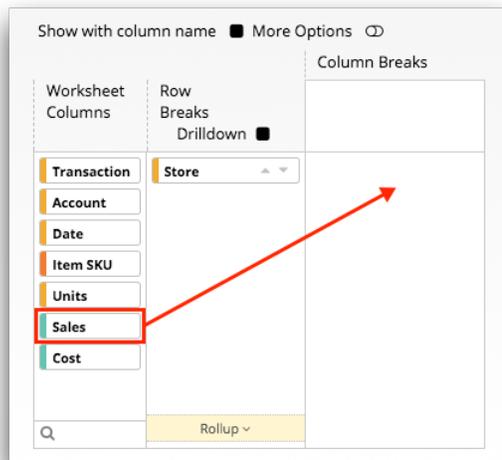
- b) Optionally, set the sort order of the column by either clicking the **Sort Up** (^) or **Sort Down** (^) icon.

All rows that have the same values for all of the columns specified will be considered part of one group.

4. For each column of data that you want to summarize, complete the following:

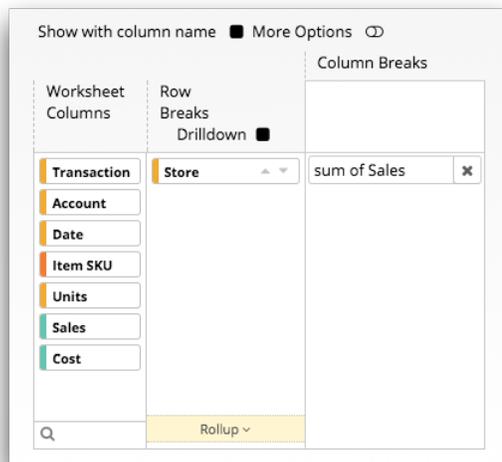
- a) Drag the column from the **Worksheet Columns** section to the lower area of the **Column Breaks** section.

This selection chooses the column of data you want to summarize.



The Trillion-Row Spreadsheet adds a tabulation function field and returns the column to the **Worksheet Columns** section.

Note: By default, columns placed in the lower area of the **Column Breaks** section are assigned the **sum of [COLUMN]** summarization. If the data type of the column is text, it is assigned the **last of [COLUMN]** summarization.



b) If applicable, change the summarization.

For instructions, see [Change the tabulation function summarization](#) on page 197.

5. Optionally, use a rollup in your tabulation.

A rollup provides a subtotal of the summarizations as a row in the resultant worksheet. For instructions, see [Use a rollup in a tabulation](#) on page 204.

6. Optionally, create a drilldown analysis from the tabulation results.

A drilldown analysis provides an interactive way to explore data points and view row-level data in the grid without changing your underlying query. For instructions, see [Set up a drilldown hierarchy](#) on page 214.

By default, summarization columns in the resultant worksheet are named t_0 , t_1 , t_2 , and so on. You can give the resultant columns more meaningful names within the **More Options** view of the **Tabulation** panel. In addition, you can define the various formats of the tabulation results.

7. Optionally, define the format of the tabulation results.

For instructions, see [Format the tabulation results](#) on page 205.

8. Click the **Submit operation** (✓) icon.

The Trillion-Row Spreadsheet displays the results of your tabulation.

Sales by store		Sum of Sales
Cols 1 fixed and 2 to 2 of 2 , Rows 1 to 3 of 3		
Store		sum_sales
		59.85
	1	23.19
	2	16.31
	3	20.35

Perform a cross tabulation

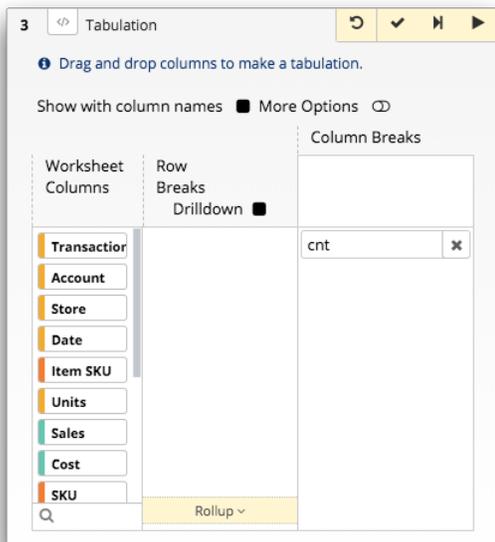
A cross tabulation allows you to calculate, or summarize, the values in a column based on the values in two or more other columns and display the result as a matrix.

A cross tabulation can help you gain granularity from a summary without losing the highest level of data summarization. For example, in addition to finding the total amount of sales for each store in a grocery chain, you can also obtain the sales figures for the individual departments within each store and how they compare to the total.

To perform a cross tabulation:

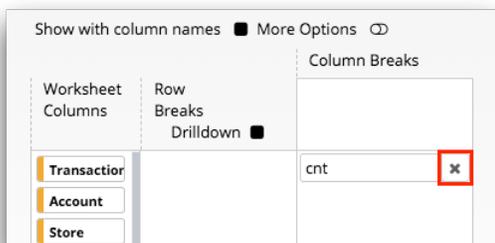
1. In the **New operation** panel, click **Tabulate**.

The Trillion-Row Spreadsheet displays the **Tabulation** panel.



Note: By default, a new **Tabulation** panel contains the **cnt** (count) tabulation function.

2. If applicable, remove the default tabulation function by clicking the **Delete Function** (**x**) icon in the field.

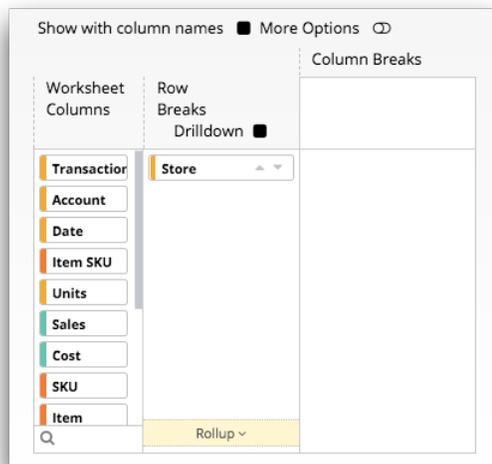


3. For the *rows* in the resultant cross tabulation, specify each column by which you want to group the data.
 - a) Drag the column from the **Worksheet Columns** section to the **Row Breaks** section.

This selection groups the data by creating a separate row in the resultant worksheet for each unique value within the chosen column. Grouping is a way of pooling all the records for a single entity or value into a single entry in the worksheet.

Note: In most cases, the order by which you choose to group the data is important because it can affect the results of the cross tabulation.

The Trillion-Row Spreadsheet moves the column to the **Row Breaks** section.



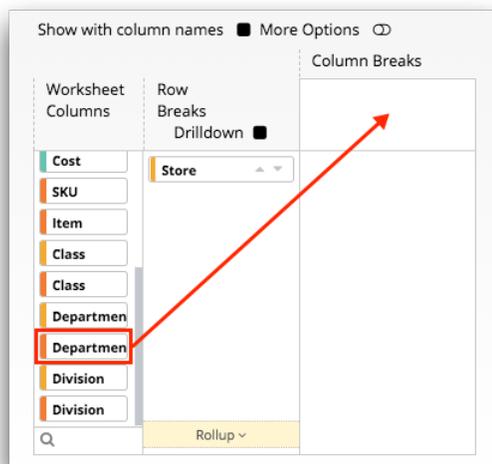
- b) Optionally, set the sort order of the column by either clicking the **Sort Up** (^) or **Sort Down** (v) icon.

All rows that have the same values for all of the columns specified will be considered part of one group.

4. For the *columns* in the resultant cross tabulation, specify each column by which you want to group the data.

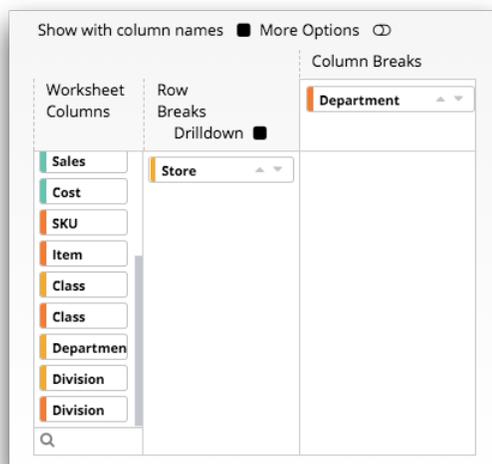
- a) Drag the column from the **Worksheet Columns** section to the upper area of the **Column Breaks** section.

This selection groups the data by creating a separate column in the resultant worksheet for each unique value within the chosen column. Grouping is a way of pooling all the records for a single entity or value into a single entry in the worksheet.



Note: In most cases, the order by which you choose to group the data is important because it can affect the results of the cross tabulation.

The Trillion-Row Spreadsheet moves the column to the upper area of the **Column Breaks** section.



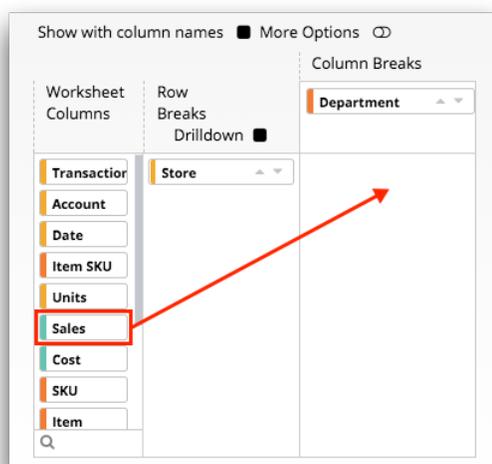
- b) Optionally, set the sort order of the column by either clicking the **Sort Up** (^) or **Sort Down** (v) icon.

All rows that have the same values for all of the columns specified will be considered part of one group.

5. For each column of data that you want to summarize, complete the following:

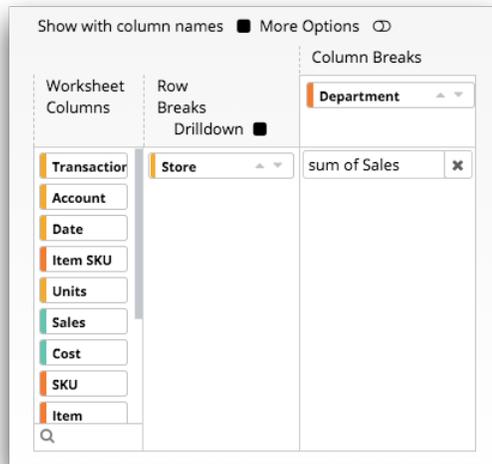
- a) Drag the column from the **Worksheet Columns** section to the lower area of the **Column Breaks** section.

This selection chooses the column of data you want to summarize.



The Trillion-Row Spreadsheet adds a tabulation function field and returns the column to the **Worksheet Columns** section.

Note: By default, columns placed in the lower area of the **Column Breaks** section are assigned the **sum of [COLUMN]** summarization. If the data type of the column is text, it is assigned the **last of [COLUMN]** summarization.



b) If applicable, change the summarization.

For instructions, see [Change the tabulation function summarization](#) on page 197.

6. Optionally, create a drilldown analysis from the cross tabulation results.

A drilldown analysis provides an interactive way to explore data points and view row-level data in the grid without changing your underlying query. For instructions, see [Set up a drilldown hierarchy](#) on page 214.

By default, columns in the resultant worksheet are named t_0 , t_1 , t_2 , and so on, for each summarization column and m_0 , m_1 , m_2 , and so on, for each tabulated column break. You can give the resultant columns more meaningful names within the **More Options** view of the **Tabulation** panel. In addition, you can define the various formats of the cross tabulation.

7. Optionally, define the format of the cross tabulation results.

For instructions, see [Format the tabulation results](#) on page 205.

8. Click the **Submit operation** (✓) icon.

The Trillion-Row Spreadsheet displays the results of your cross tabulation.

Tabulation on Sales Item Detail								
Cols 1 fixed and 2 to 7 of 7, Rows 1 to 3 of 3								
Department		FOOTWEAR	SNACKS	TOOLS/AUTO	DISPOSABLES	OFFICE SUPPLIES		
Description	dentdesc							
Store	store	59.85	0.00	24.30	27.25	6.60	1.70	
1	23.19	-5.00	12.30	12.25	3.30	0.34		
2	16.31	5.00	3.30	5.00	1.65	1.36		
3	20.35	0.00	8.70	10.00	1.65	0.00		

Change the tabulation function summarization

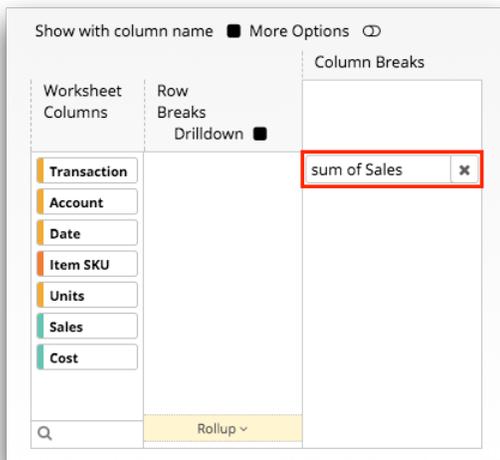
The tabulation function field defines the calculation, called a summarization, in a quick summary, tabulation, or cross tabulation. Change the summarization to select the type of calculation you want to perform.

When a column biscuit is dragged into the lower area of the **Column Breaks** section, a tabulation function is created for that column. By default, columns placed in the lower area of the **Column Breaks** section are assigned the **sum of [COLUMN]** summarization. If the data type of the column is text, it is assigned the **last of [COLUMN]** summarization.

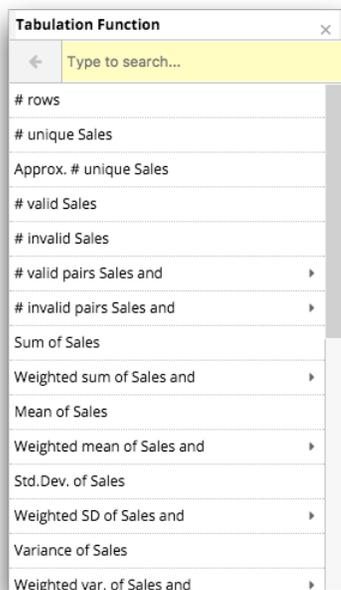
To change the tabulation function summarization:

1. In the lower area of the **Column Breaks** section, click the tabulation function field containing the summarization that you want to change.

A tabulation function field is outlined in red in the image below.



The Trillion-Row Spreadsheet displays the **Tabulation Function** menu.

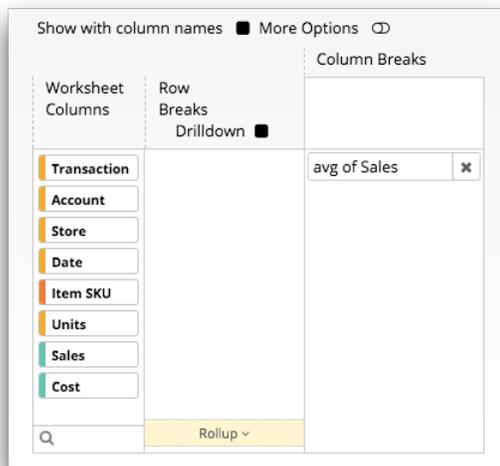


2. Select the type of summarization you want to perform.

For information about the types of summarizations available, see [Types of summarizations](#) on page 199.

Note: If the type of summarization you are performing requires a second column (a reference column), the available columns are displayed after selecting the summarization. Select the reference column from the **Tabulation Function** menu.

The Trillion-Row Spreadsheet updates the tabulation function field with the selected summarization.



Types of summarizations

The **Tabulation Function** menu in the **Tabulation** panel contains a number of calculations, called summarizations, from which you can choose.

You can choose from the following types of summarizations when performing a quick summary, tabulation, or cross tabulation.

Note: In the list below, *[COLUMN]* indicates either the column name or the column label of the selected column. By default, the column label is displayed. However, when the **Show with column names** option is selected in the **Tabulation** panel, the column name is displayed instead.

# rows	<p>The number of rows in each group.</p> <ul style="list-style-type: none"> The number of rows includes those with N/A values.
# unique [COLUMN]	<p>The number of unique values in the Column for each group.</p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the number of unique non-N/A values). If a group contains only N/A values, the result for the group is 0.
Approx. # unique [COLUMN]	<p>The approximate number of unique values in the Column for each group. This can be much faster than # unique [COLUMN] for very large cardinality columns, but provides an approximate number that can be off by a few percent.</p> <ul style="list-style-type: none"> Uses the hyper-log-log algorithm. If a group contains N/A values, those rows are ignored (i.e., the result is the number of unique non-N/A values). If a group contains only N/A values, the result for the group is 0. <p>Note: While this function can be used for small tables, it is generally used to compute the approximate number of distinct values in a column with a large number of unique values, or across a table with a very large number of rows.</p>
# valid [COLUMN]	<p>The number of non-N/A values in the Column for each group.</p>
# invalid [COLUMN]	<p>The number of N/A values in the Column for each group.</p>

# valid pairs [COLUMN] and	The number of rows where both the Column and the Reference Column are not N/A, for each group.
# invalid pairs [COLUMN] and	The number of rows where either the Column or the Reference Column is N/A, for each group.
Sum of [COLUMN]	The subtotal of the Column for each group. <i>(numeric columns only)</i> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the sum of the non-N/A values). If a group contains only N/A values, the result for the group is 0.
Weighted sum of [COLUMN] and	The dot product of the Column and the Reference Column , for each group. <i>(numeric columns only)</i> <ul style="list-style-type: none"> If, for a particular group, either the Column or the Reference Column contains N/A values, those rows are ignored (i.e., the result for the group is the dot product for those rows where neither the Column nor the Reference Column is N/A). If there are no such rows, the result for the group is 0.
Mean of [COLUMN]	The average of the Column for each group. <i>(numeric columns only)</i> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the average of the non-N/A values). If a group contains only N/A values, the result for the group is N/A.
Weighted mean of [COLUMN] and	The average of the Column , weighted by the Reference Column , for each group. <i>(numeric columns only)</i> <ul style="list-style-type: none"> If, for a particular group, either the Column or the Reference Column contains N/A values, those rows are ignored (i.e., the result for the group is the weighted average for those rows where neither the Column nor the Reference Column is N/A). If there are no such rows, the result for the group is N/A.
Std. Dev. of [COLUMN]	The standard deviation of the Column for each group. <i>(numeric columns only)</i> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the standard deviation of the non-N/A values). If a group contains only N/A values, the result for the group is N/A.
Weighted SD of [COLUMN] and	The standard deviation of the Column , weighted by the Reference Column , for each group. <i>(numeric columns only)</i> <ul style="list-style-type: none"> If, for a particular group, either the Column or the Reference Column contains N/A values, those rows are ignored (i.e., the result for the group is the weighted standard deviation for those rows where neither the Column nor the Reference Column is N/A). If there are no such rows, the result for the group is N/A.
Variance of [COLUMN]	The variance of the Column for each group. <i>(numeric columns only)</i> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the variance of the non-N/A values). If a group contains only N/A values, the result for the group is N/A.
Weighted var. of [COLUMN] and	The variance of the Column , weighted by the Reference Column , for each group. <i>(numeric columns only)</i>

	<ul style="list-style-type: none"> If, for a particular group, either the Column or the Reference Column contains N/A values, those rows are ignored (i.e., the result for the group is the weighted variance for those rows where neither the Column nor the Reference Column is N/A). If there are no such rows, the result for the group is N/A.
Covariance of [COLUMN] and	<p>The covariance of the Column and the Reference Column, for each group. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If, for a particular group, either the Column or the Reference Column contains N/A values, those rows are ignored (i.e., the result for the group is the covariance for those rows where neither the Column nor the Reference Column is N/A). If there are no such rows, the result for the group is N/A.
Correlation of [COLUMN] and	<p>The correlation of the Column and the Reference Column, for each group. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If, for a particular group, either the Column or the Reference Column contains N/A values, those rows are ignored (i.e., the result for the group is the correlation for those rows where neither the Column nor the Reference Column is N/A). If there aren't at least two such rows in the group, the result is N/A.
Sum of squares [COLUMN]	<p>The sum of squares of the Column for each group. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the sum of squares of the non-N/A values). If a group contains only N/A values, the result for the group is 0.
Highest value of [COLUMN]	<p>The highest number in the Column for each group. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the highest non-N/A value). If a group contains only N/A values, the result for the group is N/A.
Lowest value of [COLUMN]	<p>The lowest number in the Column for each group. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the lowest non-N/A value). If a group contains only N/A values, the result for the group is N/A.
First value of [COLUMN]	<p>The first value in the Column for each group. <i>(alphanumeric and numeric columns)</i></p> <p>This is the first value based on the current sort order of the table (i.e., the way the data is shown when the table is viewed multiple rows at a time).</p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the first of the non-N/A values). If a group contains only N/A values, the result for the group is N/A.
Last value of [COLUMN]	<p>The last value in the Column for each group. <i>(alphanumeric and numeric columns)</i></p> <p>This is the last value based on the current sort order of the table (i.e., the way the data is shown when the table is viewed multiple rows at a time).</p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the last of the non-N/A values).

	<ul style="list-style-type: none"> If a group contains only N/A values, the result for the group is N/A.
Median of [COLUMN]	<p>The median of the Column for each group. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the median of the non-N/A values). If there are an even number of values, the median is the average of the two middle values. If a group contains only N/A values, the result for the group is N/A.
Mode of [COLUMN]	<p>The mode of the Column for each group. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the mode of the non-N/A values). If, for a given group, there is no unique mode, the result is arbitrary. If a group contains only N/A values, the result for the group is N/A.
F(mode) of [COLUMN]	<p>The frequency of the mode of the Column for each group. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the frequency of the mode of the non-N/A values). If a group contains only N/A values, the result for the group is 0.
# modes of [COLUMN]	<p>The number of modes in the Column for each group. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the number of modes within the non-N/A values). If a group contains only N/A values, the result for the group is 0.
Rank among	<p>all values of</p> <p>The rank of each group, in terms of its number of records, with respect to all other groups.</p> <ul style="list-style-type: none"> The rank includes rows with N/A values. <p>column values of</p> <p>The rank of each group, in terms of its number of records, with respect to the other groups in the same column of the cross tabulation result.</p> <ul style="list-style-type: none"> The rank includes rows with N/A values. <p>row values of</p> <p>The rank of each group, in terms of its number of records, with respect to the other groups on the same row of the cross tabulation result.</p> <ul style="list-style-type: none"> The rank includes rows with N/A values.
Percent of	<p>grand total of</p> <p>The subtotal of the Column for each group as a percentage of the grand total. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the sum of the non-N/A values in the group as a percentage of the sum of the non-N/A values in the entire column). If a group contains only N/A values, the result for the group is 0. <p>Note: The result is a whole number, not a fractional value. For instance, a result of 20% would be represented as 20, not 0.20.</p>

	<p>column total of</p> <p>The subtotal of the Column for each group as a percentage of the total for that column of the cross tabulation result. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> • If a group contains N/A values, those rows are ignored (i.e., the result is the sum of the non-N/A values in the group as a percentage of the sum of the non-N/A values in that column of the cross tabulation result). • If a group contains only N/A values, the result for the group is 0. <p>Note: The result is a whole number, not a fractional value. For instance, a result of 20% would be represented as 20, not 0.20.</p> <p>row total of</p> <p>The subtotal of the Column for each group as a percentage of the total for that row of the cross tabulation result. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> • If a group contains N/A values, those rows are ignored (i.e., the result is the sum of the non-N/A values in the group as a percentage of the sum of the non-N/A values in that row of the cross tabulation result). • If a group contains only N/A values, the result for the group is 0. <p>Note: The result is a whole number, not a fractional value. For instance, a result of 20% would be represented as 20, not 0.20.</p>
<p>Fraction of</p>	<p>grand total of</p> <p>The subtotal of the Column for each group as a fraction of the grand total. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> • If a group contains N/A values, those rows are ignored (i.e., the result is the sum of the non-N/A values in the group as a fraction of the sum of the non-N/A values in the entire column). • If a group contains only N/A values, the result for the group is 0. <p>Note: The result is a fractional value. For instance, a result of 20% would be represented as 0.20.</p> <p>column total of</p> <p>The subtotal of the Column for each group as a fraction of the total for that column of the cross tabulation result. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> • If a group contains N/A values, those rows are ignored (i.e., the result is the sum of the non-N/A values in the group as a fraction of the sum of the non-N/A values in that column of the cross tabulation result). • If a group contains only N/A values, the result for the group is 0. <p>Note: The result is a fractional value. For instance, a result of 20% would be represented as 0.20.</p> <p>row total of</p> <p>The subtotal of the Column for each group as a fraction of the total for that row of the cross tabulation result. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> • If a group contains N/A values, those rows are ignored (i.e., the result is the sum of the non-N/A values in the group as a fraction of the sum of the non-N/A values in that row of the cross tabulation result). • If a group contains only N/A values, the result for the group is 0.

Note: The result is a fractional value. For instance, a result of 20% would be represented as 0.20.

Use a rollup in a tabulation

When performing a tabulation with row breaks, you can choose to group like values using the rollup feature.

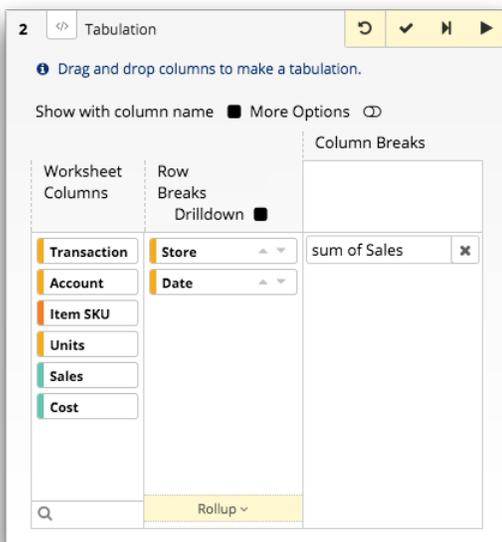
A rollup performs an additional tabulation by grouping like values on columns in the **Row Breaks** section of the **Tabulation** panel. This makes it easy to aggregate data at multiple levels in a tabulation.

Note: The rollup feature is not available in a cross tabulation.

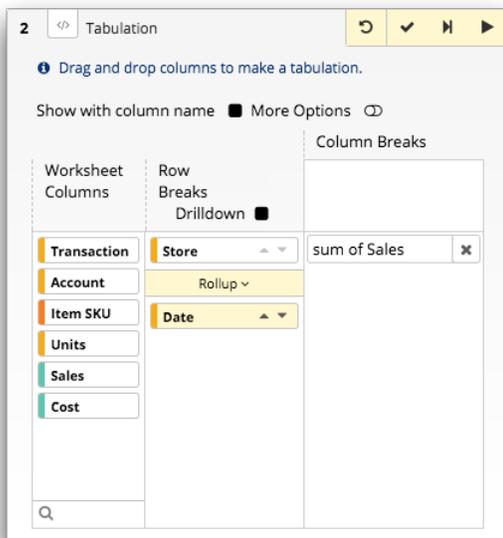
To use a rollup in a tabulation:

1. Perform a tabulation with at least one row break.

For instructions, see [Perform a tabulation](#) on page 190.



2. In the **Row Breaks** section, drag the **Rollup** bar above the columns you want included in the rollup. Columns above the bar are not included in the rollup. To include all columns in the rollup, drag the bar to the top of the section. The Trillion-Row Spreadsheet changes the color of the biscuits included in the rollup to yellow.



3. Click the **Submit operation** (✓) icon.
The Trillion-Row Spreadsheet displays the results of your tabulation with the selected rollup columns.

Tabulation on Sales Item Detail
Cols 2 fixed and 3 to 3 of 3 , Rows 1 to 13 of 13

Store	Date	Sales
store	date	t0
		59.85
1	05/15/12	-3.40
1	05/16/12	8.25
2	05/15/12	3.85
2	05/17/12	6.02
3	05/15/12	10.35
1	05/18/12	6.95
2	06/19/12	6.10
2	05/19/12	0.34
3	05/22/12	10.00
1	06/03/12	11.39
1		23.19
2		16.31
3		20.35

Format the tabulation results

Define the format of data in columns resulting from any of the three basic types of tabulations (quick summaries, tabulations, and cross tabulations).

By default, columns in the resultant worksheet are named t0, t1, t2, and so on, for each summarization column and m0, m1, m2, and so on, for each tabulated column break. You can give the resultant columns more meaningful names and define other various formatting options.

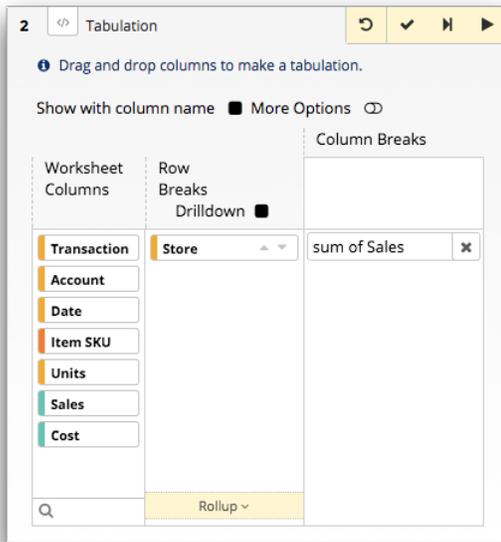
To format the tabulation results:

1. Perform a quick summary, tabulation, or cross tabulation.

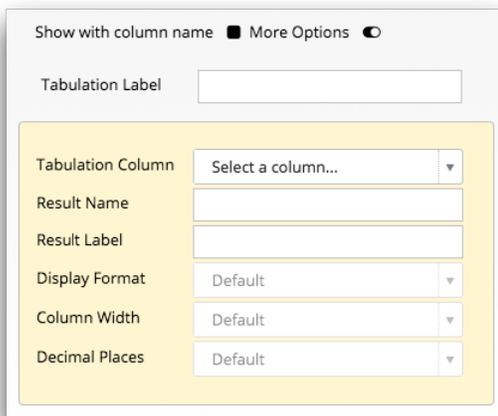
For instructions, see the following:

- [Perform a quick summary](#) on page 186

- [Perform a tabulation](#) on page 190
- [Perform a cross tabulation](#) on page 193



2. Click the **More Options** switch.
The Trillion-Row Spreadsheet displays the tabulation result options.



3. In the **Tabulation Label** field, enter a label for the tabulation.
This optional field allows you to replace the table title with a label that describes the tabulation.
4. For each summary result column you want to format, complete the following:
 - a) In the **Tabulation Column** drop-down list, select a summarization.
The Trillion-Row Spreadsheet enables the fields and options for the selected summary result column.

Show with column name More Options

Tabulation Label

Tabulation Column

Result Name

Result Label

Display Format

Column Width

Decimal Places

b) Complete the appropriate fields and options.

For a list of fields and options, see [Tabulation result options](#) on page 207.

5. Click the **Submit operation** (✓) icon.
The Trillion-Row Spreadsheet displays the results of your tabulation with the defined formatting.

Sales by store
Cols 1 fixed and 2 to 2 of 2 , Rows 1 to 3 of 3

Store	Total Sales
store	sum_sales
	\$59.85
1	\$23.19
2	\$16.31
3	\$20.35

Tabulation result options

Fields and options for defining the format of the tabulation results.

The fields and options available in the more options view of the **Tabulation** panel allow you to define the format of data in columns resulting from any of the three basic types of tabulations (quick summaries, tabulations, and cross tabulations).

Show with column name More Options

Tabulation Label

Tabulation Column

Result Name

Result Label

Display Format

Column Width

Decimal Places

Figure 45: More options view of the Tabulation panel

Tabulation Label

This optional field allows you to replace the table title with a label that describes the tabulation.

Tabulation Column

This drop-down list contains the set tabulation functions. Select a tabulation function to define the formatting options for the resulting column data.

After a tabulation function is selected, the Trillion-Row Spreadsheet enables the remaining fields and options.

Result Name

This field defines the name to assign to the result column.

This is the *column name*. The column name is used by the 1010data Insights Platform to interact with the column. It is also used when writing more advanced value expressions and queries.

The column name must be unique. It may only contain alphanumeric characters or underscores and must begin with an alphabetic character (e.g., `percent_total_sales`). It may not contain any spaces or other special characters.

Note: When entered, the Trillion-Row Spreadsheet (TRS) automatically removes special characters, replaces spaces with underscores, and displays a visual indication that a change was made.

Result Label

This field defines the label to assign to the result column.

This is the *column label*. For easier readability, column labels may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters (e.g., "Percentage of Total Sales (%)"). The backtick (`) special character is used to break the column label text line into multiple lines.

Note: On most standard United States keyboards, the **backtick** (`) character is immediately to the left of **1** key.

While not required, this field is recommended.

Display Format

This drop-down list instructs the Trillion-Row Spreadsheet how to display the value in the result column. For example, as a date, a number, as text, and so forth. For more information, see [Display formats](#) on page 56.

Note: If **Default** is selected, the format is determined by the data.

Column Width

This drop-down list specifies how many place values to hold in the column. In other words, the width of the result column by the number of characters.

You can select from **1** to **20** characters.

Note: If **Default** is selected, the width is determined by the data.

Decimal Places

This drop-down list sets how many numbers are displayed after the decimal point in the result column data.

You can select from **0** to **9** decimal places.

Note: If **Default** is selected, the number of decimal places is determined by the data, and may be different for each row.

Drilldown analysis

A drilldown analysis provides an interactive way to explore data points and view row-level data in the grid without changing your underlying query.

When performing a tabulation or cross tabulation, you have the option of creating a drilldown analysis from the results. To create a drilldown analysis, you must first set up a drilldown hierarchy in the **Tabulation** panel within the Trillion-Row Spreadsheet timeline. After the operation is submitted, the drilldown analysis is displayed in the Grid view. You can then use the drilldown controls in the grid to easily traverse the results of your analysis to look at different aspects of the data.

See the example below for an explanation of how a drilldown hierarchy could be used in an analysis.

Example

This example illustrates the use of a simple drilldown hierarchy, set up in the **Tabulation** panel of the Trillion-Row Spreadsheet timeline, to perform a drilldown analysis within the Grid view.

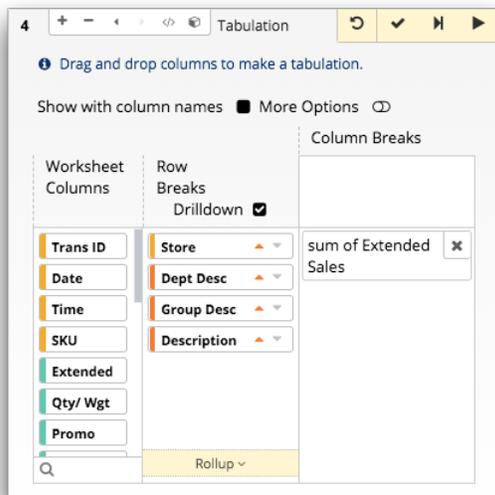
For example, you may want to "drill down" into your sales item detail data to see which product most contributed to the overall sales at a store in a grocery chain. This is accomplished by adding a drilldown hierarchy to a tabulation operation. The drilldown hierarchy in the **Tabulation** panel is set up by selecting the **Drilldown** option and then dragging the desired columns into the **Row Breaks** section.

In this example, these four columns are used and are arranged in the following order:

- **Store**
- **Dept Desc**
- **Group Desc**
- **Description**

The order in which columns are arranged in the **Row Breaks** section determine their level within the hierarchy. Each level of the drilldown hierarchy (from top to bottom) provides greater detail than the level above it. In this case, the least amount of detail is at the store level (the first item in the hierarchy) with the greatest amount of detail at the product description level (the last item in the hierarchy).

As each item is added to the drilldown, the Trillion-Row Spreadsheet automatically sets the column sort controls to display the results in ascending order.



After defining the drilldown hierarchy, setting the tabulation function, and submitting the operation, the Trillion-Row Spreadsheet updates the grid to show the total amount of sales for each store. This occurs because **Store** was set as the first item in the drilldown hierarchy.

The screenshot displays the 'Analyze' tab in the Trillion-Row Spreadsheet. The timeline shows four steps: 1. Open table pub.doc.retail.salesdetail; 2. Link to table pub.doc.retail.product; 3. Select rows where (between(trans_date;20140101;20140201)); 4. Tabulation. The tabulation configuration shows 'Store' as the primary column break, with 'Dept Desc', 'Group Desc', and 'Description' as secondary breaks. The resulting grid, titled 'Tabulation on Sales Detail', shows a total of 9,570,416.82 for 'Extended Sales' across 19 rows. The 'Store' column contains blue links, and the 'Extended Sales' column is labeled 't0'.

Store	Extended Sales
store	t0
	9,570,416.82
1	338,196.48
3	431,970.72
4	764,244.22
6	824,972.49
7	986,379.71
8	492,204.24
9	297,584.79
10	424,385.36
11	631,413.33
12	257,702.18
13	323,935.82
14	507,887.88
15	567,888.26
16	403,790.72
17	665,996.62
18	510,316.22
19	534,356.85

Values in the **Store** column are blue, signifying that each value is a link. The presence of links indicates that an additional level of detail is available for the value. In addition, the buttons located at the top of the grid display the remaining tabulation drilldown levels.

Clicking a link in the **Store** column allows you to "drill down" into the next level of detail in the tabulation. When you drill down into another level, the Trillion-Row Spreadsheet adds the appropriate select operation to the timeline, just before the tabulation, and updates the results in the grid. For example, clicking 3 updates the grid to show the total amount of sales for each department (the second level in the drilldown hierarchy) at store 3.

The screenshot displays the 'Analyze' tab in the Trillion-Row Spreadsheet. The timeline shows five steps: 1. Open table pub.doc.retail.salesdetail; 2. Link to table pub.doc.retail.product; 3. Select rows where (between(trans_date;20140101;20140201)); 4. Select rows where store='3'; 5. Tabulation. The tabulation configuration shows 'Dept Desc' and 'Description' as column breaks. The resulting grid, titled 'Tabulation on Sales Detail', shows a total of 431,970.72 for 'Extended Sales' across 17 rows. The 'Dept Desc' column contains blue links, and the 'Extended Sales' column is labeled 't0'.

Store	Dept Desc	Extended Sales
store	deptdesc	t0
		431,970.72
3	BAKERY	16,342.45
3	BEER	26,271.70
3	BEVERAGE	40,433.91
3	BOOKS/MAGAZINES	1,229.56
3	DAIRY	33,143.36
3	DAIRY DELI	59,275.89
3	DISTILLED SPIRITS	12,596.35
3	FROZEN	38,477.94
3	GARDEN	381.50
3	GM	10,010.07
3	LOTTO SALES	365.42
3	MEAT	19,050.25
3	MEAT DELI	30,777.92
3	POSTAGE STAMPS	486.58
3	PRODUCE	115,030.03
3	SEAFOOD	6,649.92
3	SERVICE DELI	13,752.18

From the results, it is easy to see that the produce department contributed most to the overall sales at the store.

Like the values in the **Store** column discussed above, links in **Dept Desc** indicate that another level of detail is available for each department. The buttons above the grid show that the next level in the drilldown is the group description.

Clicking **PRODUCE** updates the grid and shows the total amount of sales for each group (the third level in the drilldown hierarchy) in the produce department.

The screenshot displays the 'Tabulation on Sales Detail' interface. The control panel on the left shows filters for 'Store' (value: 3) and 'Dept Desc' (value: PRODUCE). The data grid on the right shows the following data:

Store	Dept	Group Desc	Extended Sales
store	Desc	groupdesc	t0
			deptdesc
			115,030.030
3	PRODUCE	BEANS & RICE	3,909.070
3	PRODUCE	DRIED FRUIT NUTS & SEEDS	2,786.470
3	PRODUCE	FRUIT	32,418.900
3	PRODUCE	POTATOES	5,619.600
3	PRODUCE	PRODUCE RELATED	1,328.290
3	PRODUCE	SALAD	7,761.310
3	PRODUCE	SPICES	1,639.680
3	PRODUCE	TOMATOES	9,243.060
3	PRODUCE	TROPICAL	16,808.050
3	PRODUCE	UNSCANNED PRODUCE	541.460
3	PRODUCE	VEGETABLES PRODUCE	32,974.140

Although the totals are close, the results show that the vegetables group made up the bulk of the sales amount in the produce department; the fruit group was a close second. To see which product contributed most to the store sales, you can drill into the vegetable group.

Clicking **VEGETABLES PRODUCE** updates the grid and shows the total amount of sales for each product (the last level in the drilldown hierarchy) in the vegetable group.

Tabulation on Sales Detail
Cols 1 to 5 of 5, Rows 1 to 17 of 117

Store	Dept	Group Desc	Description	Extended Sales
store	Desc	groupdesc	description	t0
deptdesc				32,974.14
3	PRODUCE	VEGETABLES PRODUCE	BROCCOLI CROWN CUT-WIC	2,082.20
3	PRODUCE	VEGETABLES PRODUCE	CELERY STALK	1,903.31
3	PRODUCE	VEGETABLES PRODUCE	SQUASH ITALIAN-WIC	1,614.55
3	PRODUCE	VEGETABLES PRODUCE	CUCUMBER LUG - WIC	1,568.61
3	PRODUCE	VEGETABLES PRODUCE	PEPPERS BELL GREEN - WIC	1,160.32
3	PRODUCE	VEGETABLES PRODUCE	CAULIFLOWER - WIC	1,110.22
3	PRODUCE	VEGETABLES PRODUCE	ONIONS YELLOW - WIC	938.20
3	PRODUCE	VEGETABLES PRODUCE	ONIONS WHITE - WIC	859.82
3	PRODUCE	VEGETABLES PRODUCE	ONION GREEN - WIC	851.77
3	PRODUCE	VEGETABLES PRODUCE	CABBAGE GREEN - WIC	847.70
3	PRODUCE	VEGETABLES PRODUCE	FRESH CILANTRO	760.33
3	PRODUCE	VEGETABLES PRODUCE	PEPPERS CHILI JALAPENO >	706.99
3	PRODUCE	VEGETABLES PRODUCE	CARROTS BABY CELLO 1 LB	690.53
3	PRODUCE	VEGETABLES PRODUCE	ASPARAGUS LARGE - WIC	668.76
3	PRODUCE	VEGETABLES PRODUCE	SQUASH ITALIAN GREY (ME >	636.96
3	PRODUCE	VEGETABLES PRODUCE	CARROTS CELLO 2 LB - WIC	636.45
3	PRODUCE	VEGETABLES PRODUCE	CARROTS BABY - WIC	634.76

After sorting the results by the **Extended Sales** column, you can see that broccoli is the top contributor to the total sales at store 3.

Tabulation on Sales Detail
Cols 1 to 5 of 5, Rows 1 to 17 of 117

Store	Dept	Group Desc	Description	Extended Sales
store	Desc	groupdesc	description	t0
deptdesc				32,974.14
3	PRODUCE	VEGETABLES PRODUCE	BROCCOLI CROWN CUT-WIC	2,082.20
3	PRODUCE	VEGETABLES PRODUCE	CELERY STALK	1,903.31
3	PRODUCE	VEGETABLES PRODUCE	SQUASH ITALIAN-WIC	1,614.55
3	PRODUCE	VEGETABLES PRODUCE	CUCUMBER LUG - WIC	1,568.61
3	PRODUCE	VEGETABLES PRODUCE	PEPPERS BELL GREEN - WIC	1,160.32
3	PRODUCE	VEGETABLES PRODUCE	CAULIFLOWER - WIC	1,110.22
3	PRODUCE	VEGETABLES PRODUCE	ONIONS YELLOW - WIC	938.20
3	PRODUCE	VEGETABLES PRODUCE	ONIONS WHITE - WIC	859.82
3	PRODUCE	VEGETABLES PRODUCE	ONION GREEN - WIC	851.77
3	PRODUCE	VEGETABLES PRODUCE	CABBAGE GREEN - WIC	847.70
3	PRODUCE	VEGETABLES PRODUCE	FRESH CILANTRO	760.33
3	PRODUCE	VEGETABLES PRODUCE	PEPPERS CHILI JALAPENO >	706.99
3	PRODUCE	VEGETABLES PRODUCE	CARROTS BABY CELLO 1 LB	690.53
3	PRODUCE	VEGETABLES PRODUCE	ASPARAGUS LARGE - WIC	668.76
3	PRODUCE	VEGETABLES PRODUCE	SQUASH ITALIAN GREY (ME >	636.96
3	PRODUCE	VEGETABLES PRODUCE	CARROTS CELLO 2 LB - WIC	636.45
3	PRODUCE	VEGETABLES PRODUCE	CARROTS BABY - WIC	634.76

In the column header row of the grid, orange column labels and column names are drill-up links. You can click the link to return, or "drill up," to the previous level of detail in the results. This allows you to easily traverse the drilldown hierarchy to look at different aspects of the data.

When you drill up into a previous level, the Trillion-Row Spreadsheet removes the appropriate select operation from the timeline and updates the results in the grid. For example, at this point of the analysis, clicking **Description** in the column header row updates the grid to, once again, show the total amount of sales for each group in the produce department.

Tabulation on Sales Detail rolling up: **Description**
Cols 3 fixed and 4 to 4 of 4, Rows 1 to 11 of 11

Store	Dept Desc	Group Desc	Extended Sales t0
			115,030.030
3	PRODUCE	BEANS & RICE	3,909.070
3	PRODUCE	DRIED FRUIT NUTS & SEEDS	2,786.470
3	PRODUCE	FRUIT	32,418.900
3	PRODUCE	POTATOES	5,619.600
3	PRODUCE	PRODUCE RELATED	1,328.290
3	PRODUCE	SALAD	7,761.310
3	PRODUCE	SPICES	1,639.680
3	PRODUCE	TOMATOES	9,243.060
3	PRODUCE	TROPICAL	16,808.050
3	PRODUCE	UNSCANNED PRODUCE	541.460
3	PRODUCE	VEGETABLES PRODUCE	32,974.140

Because the fruit group was a close second, you might be interested to learn which item in that group contributed most to the store sales. To see which product contributed the second most to the store sales, you can drill into the fruit group.

Clicking **FRUIT** updates the grid and shows the total amount of sales for each product in the fruit group.

Tabulation on Sales Detail
Cols 4 fixed and 5 to 5 of 5, Rows 1 to 17 of 110

Store	Dept Desc	Group Desc	Description	Extended Sales t0
				32,418.90
3	PRODUCE	FRUIT	APPLE GRANNY SMITH APPL >	81.57
3	PRODUCE	FRUIT	APPLE HONEY CRISP-WIC	19.49
3	PRODUCE	FRUIT	APPLES AMBROSIA-WIC	36.81
3	PRODUCE	FRUIT	APPLES FUJI CELLO 5 LB BAG	708.03
3	PRODUCE	FRUIT	APPLES FUJI LARGE - WIC	1,813.70
3	PRODUCE	FRUIT	APPLES GALA CELLO -WIC	465.21
3	PRODUCE	FRUIT	APPLES GALA LARGE - WIC	963.70
3	PRODUCE	FRUIT	APPLES GALA SMALL - WIC	1,396.54
3	PRODUCE	FRUIT	APPLES GOLD DEL LARGE - >	430.50
3	PRODUCE	FRUIT	APPLES GOLDEN DEL 3 LB BAG	127.78
3	PRODUCE	FRUIT	APPLES GRANNY CELLO	142.31
3	PRODUCE	FRUIT	APPLES GRANNY SMI LARGE >	560.76
3	PRODUCE	FRUIT	APPLES RED DELIC LARGE - >	682.69
3	PRODUCE	FRUIT	APPLES RED.CELLO 5 LB BAG	161.47
3	PRODUCE	FRUIT	APPLES ROMES LOOSE-WIC	21.13
3	PRODUCE	FRUIT	APRICOTS - WIC	359.92
3	PRODUCE	FRUIT	AVOCADOS BAG 4 CT	397.35

After sorting the results by the **Extended Sales** column, you can see that avocados are the next top contributor to the total sales at store 3.

Tabulation on Sales Detail
Cols 4 fixed and 5 to 5 of 5 , Rows 1 to 17 of 110

Store store	Dept Desc deptdesc	Group Desc groupdesc	Description description	Extended Sales t0
				32,418.90
3	PRODUCE	FRUIT	AVOCADOS HASS 60s - WIC	2,689.40
3	PRODUCE	FRUIT	ORANGES LG 48/56S	2,314.10
3	PRODUCE	FRUIT	GRAPES SEEDLESS WHITE - >	1,982.84
3	PRODUCE	FRUIT	APPLES FUJI LARGE - WIC	1,813.70
3	PRODUCE	FRUIT	GRAPES RED SEEDLESS - WIC	1,706.45
3	PRODUCE	FRUIT	ORANGES CELLO LARGE	1,496.43
3	PRODUCE	FRUIT	APPLES GALA SMALL - WIC	1,396.54
3	PRODUCE	FRUIT	PEACHES - WIC	1,207.49
3	PRODUCE	FRUIT	GRAPES RED GLOBE. - WIC	1,126.83
3	PRODUCE	FRUIT	PEARS DANJOU - WIC	1,021.11
3	PRODUCE	FRUIT	SMALL JAZZ APPLES	978.51
3	PRODUCE	FRUIT	APPLES GALA LARGE - WIC	963.70
3	PRODUCE	FRUIT	LIMES - WIC	866.65
3	PRODUCE	FRUIT	TANGERINES CLEMENTINE 5 >	811.04
3	PRODUCE	FRUIT	MELON CANTALOUPE - WIC	788.64
3	PRODUCE	FRUIT	APPLES FUJI CELLO 5 LB BAG	708.03
3	PRODUCE	FRUIT	APPLES RED CELLO LARGE - >	692.69

Set up a drilldown hierarchy

Set up a drilldown hierarchy in a tabulation or cross tabulation operation to create a drilldown analysis in the Grid view of the Trillion-Row Spreadsheet.

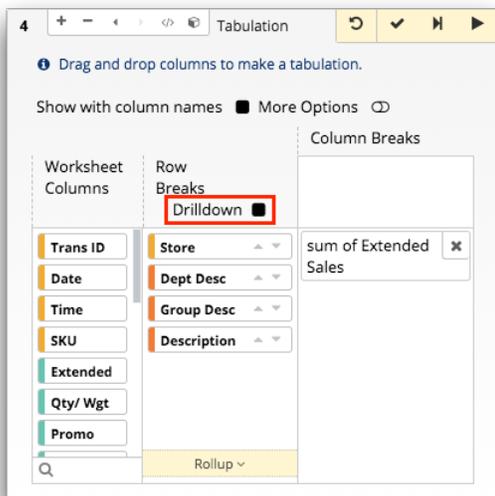
Before setting up a drilldown hierarchy, first create a tabulation or cross tabulation. For instructions, see [Perform a tabulation](#) on page 190 or [Perform a cross tabulation](#) on page 193.

After the drilldown hierarchy is set up and the operation is submitted, the drilldown analysis is displayed in the Grid view. You can then use the drilldown controls in the grid to easily traverse the results of your analysis to look at different aspects of the data. For more information, see [Drilldown controls in the grid](#) on page 215.

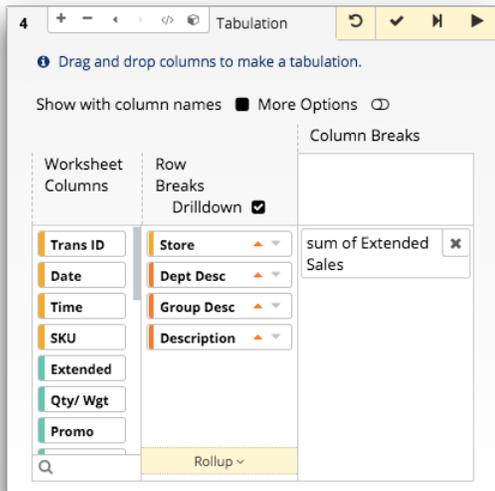
To set up a drilldown hierarchy:

1. In the **Row Breaks** section of the **Tabulation** panel, select the **Drilldown** option.

This option allows you to set up a hierarchy, defined by the order of the columns in the **Row Breaks** section.



The Trillion-Row Spreadsheet sets the column sort controls to display the results in ascending order.



2. As necessary, rearrange the columns in the **Row Breaks** section to create the drilldown hierarchy.

The order in which columns are arranged in the **Row Breaks** section determine their level within the hierarchy. Each level of the drilldown hierarchy (from top to bottom) provides greater detail than the level above it. For example, in the image above, the least amount of detail is at the store level (the first item in the hierarchy) with the greatest amount of detail at the description level (the last item in the hierarchy).

Drilldown controls in the grid

The drilldown controls in the grid allow you to easily traverse the drilldown hierarchy to look at different aspects of the data.

The results of a tabulation or cross tabulation containing a drilldown hierarchy display in the Grid view of the Trillion-Row Spreadsheet. Controls in the grid allow you to view more general or more detailed information within a predefined drilldown hierarchy.

Tabulation on Sales Detail rolling up: **Group Desc** Description

Cols 2 fixed and 3 to 3 of 3, Rows 1 to 17 of 18

	Store store	Dept Desc deptdesc	Extended Sales t0
			431,970.72
3		BAKERY	16,342.45
3		BEER	26,271.70
3		BEVERAGE	40,433.91
3		BOOKS/MAGAZINES	1,229.56
3		DAIRY	33,143.36
3		DAIRY DELI	59,275.89
3		DISTILLED SPIRITS	12,596.35
3		FROZEN	38,477.94
3		GARDEN	381.50
3		GM	10,010.07
3		LOTTO SALES	365.42
3		MEAT	19,050.25
3		MEAT DELI	30,777.92
3		POSTAGE STAMPS	486.58
3		PRODUCE	115,030.03
3		SEAFOOD	6,649.92
3		SERVICE DELI	13,752.18

Figure 46: Drilldown controls in the Grid view

A. Drilldown hierarchy button

The buttons located at the top of the grid display the remaining tabulation drilldown levels.

B. Drill-up link

In the column header row, column names and labels in orange are drill-up links. The presence of drill-up links indicates that a previous level of detail is available.

Click the contents of the column header row to return, or "drill up," to the previous level of detail in the results. When you drill up into the previous level of detail, the Trillion-Row Spreadsheet removes the appropriate select operation from the timeline and updates the results in the grid.

C. Drilldown link

In a column, values in blue are drilldown links. The presence of drilldown links indicate that an additional level of detail is available for the value.

Click a link to drill down into the next level of detail in the tabulation. When you drill down into another level, the Trillion-Row Spreadsheet adds the appropriate select operation to the timeline, just before the tabulation, and updates the results in the grid.

Merging tables and worksheets

You can merge two or more tables or worksheets to combine their rows together into a single, larger worksheet.

Merging two or more tables or worksheets combines their rows together into a single, larger worksheet. The rows from the foreign table(s) are appended to the end of the base table in the order in which they are specified. The results apply only to your session; the original tables or worksheets are not affected. The resulting table may contain all columns in either table, all common columns, or only a subset of the common columns.

Note: Merging differs from linking. Merging combines the *rows* of two or more tables or worksheets together whereas linking combines the *columns* of two tables or worksheets together. Links align worksheets side by side, while merges combine worksheets vertically.

Merge tables panel

Descriptions of user interface elements that are available when merging tables and worksheets in the Trillion-Row Spreadsheet.

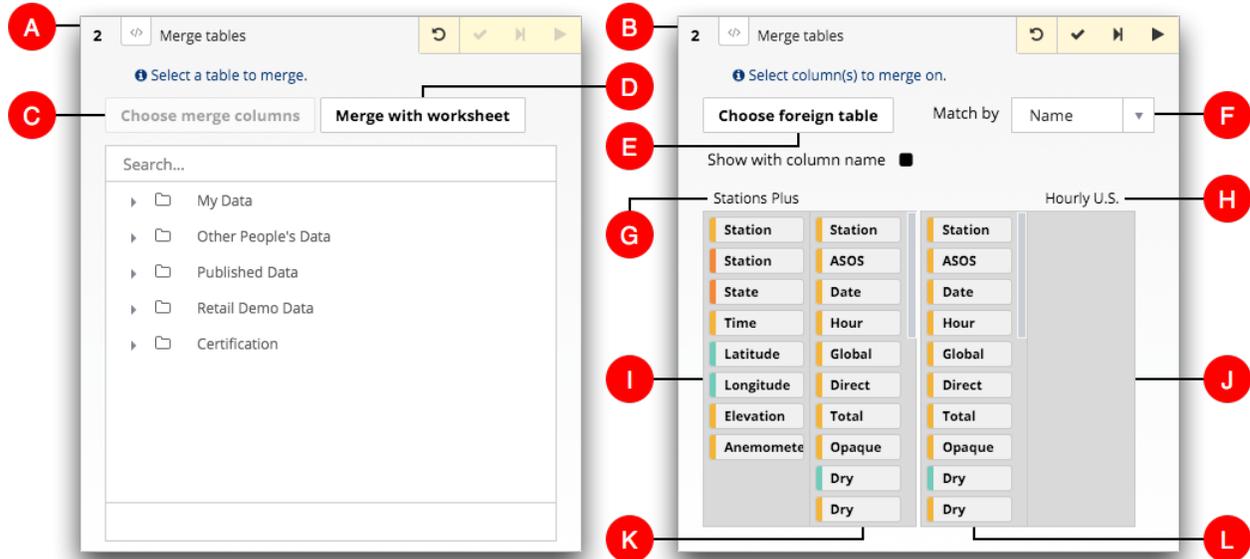


Figure 47: Merge tables panel in the Trillion-Row Spreadsheet

A. Object browser view

The object browser view of the **Merge tables** panel appears after clicking **Merge** in the **New operation** panel. It also appears after clicking **Choose foreign table** in the column view.

Use this view to select the foreign table or worksheet that you want to merge into the current table or worksheet.

B. Column view

The column view of the **Merge tables** panel appears after selecting the foreign table or worksheet that you want to merge into the current table or worksheet.

Use this view to choose how the columns from each table or worksheet are matched. After choosing how columns are matched, this view displays a visual indication of how columns are matched.

C. Choose merge columns

This button is available after clicking **Choose foreign table** in the column view of the **Merge tables** panel.

Click this button to return to column view of the **Merge tables** panel without selecting a different table or worksheet.

D. Merge with worksheet

This button displays a list of existing worksheets currently open in your session. It also provides an option for creating a new worksheet to use in the merge operation.

Click this button to select either an existing or new worksheet to merge into the current table or worksheet.

E. Choose foreign table

This button displays the object browser view of the **Merge tables** panel.

Click this button to select a different table or worksheet to merge into the current table or worksheet.

F. Match by

This drop-down list defines how the columns in each table or worksheet are matched in the merge operation.

You can choose to include all columns in either table or worksheet, all common columns, or only a subset of common columns in the merged worksheet.

Name

This option matches columns in the current table or worksheet with columns of the same name in the foreign table or worksheet. Each table must have at least one column with the same name. Matching columns must have the same data type. The merged table will contain only those columns that are common to both tables; all non-matching columns from each table or worksheet are excluded.

Name with padding

This option matches columns in the current table or worksheet with columns of the same name and data type in the foreign table or worksheet. Columns that match appear once in the merged worksheet with data from both tables. Unmatched columns are "padded" with the appropriate N/A value.

Order

This option matches columns in the current table or worksheet with the corresponding columns (by order, regardless of name) in the foreign table or worksheet. The columns must have the same data type and must be in the same order. In addition, there must be the same number of columns in both tables or worksheets.

G. Current table or worksheet

This is the label of the current table or worksheet.

The two sections under this label (the two on the left side of the panel) show the columns in the current table or worksheet.

Note: Unlike the column biscuits in other Trillion-Row Spreadsheet panels, you cannot drag the columns to move them in the **Merge tables** panel. Biscuits in this panel show how the columns from the current and foreign table or worksheet are matched. The display order and location of column biscuits is defined by the option selected in the **Match by** drop-down list.

H. Foreign table or worksheet

This is the label of the foreign table or worksheet.

The two sections under this label (the two on the right side of the panel) show the columns in the foreign table or worksheet.

Note: Unlike the column biscuits in other Trillion-Row Spreadsheet panels, you cannot drag the columns to move them in the **Merge tables** panel. Biscuits in this panel show how the columns from the current and foreign table or worksheet are matched. The display order and location of column biscuits is defined by the option selected in the **Match by** drop-down list.

I. Excluded current columns

This section displays the columns from the current table or worksheet that are excluded from the merged worksheet.

In the example image above, some columns from the current worksheet are excluded from the merged worksheet.

J. Excluded foreign columns

This section displays the columns from the foreign table or worksheet that are excluded from the merged worksheet.

In the example image above, all of the columns from the foreign table are included in the merged worksheet.

K. Matched current columns

This section displays the columns from the current table or worksheet that are matched with columns from the foreign table or worksheet.

Note: When the **Name with padding** match option is selected, columns in the foreign table or worksheet that do not exist in the current table or worksheet are indicated by the **Blank** column biscuit in this section.

L. Matched foreign columns

This section displays the columns from the foreign table or worksheet that are matched with columns from the current table or worksheet.

Note: When the **Name with padding** match option is selected, columns in the current table or worksheet that do not exist in the foreign table or worksheet are indicated by the **Blank** column biscuit in this section.

Merge in a table

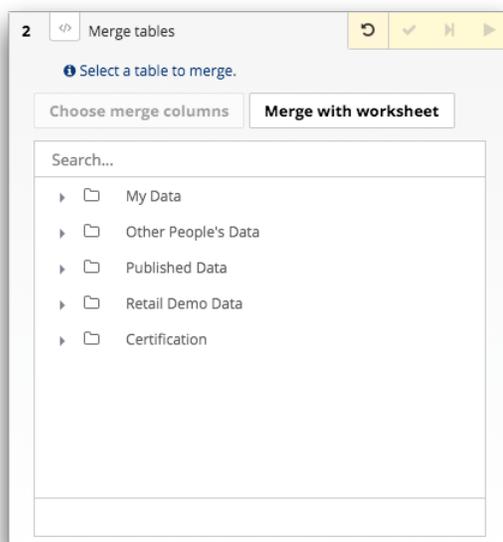
Merge a foreign table into the current table or worksheet.

Merging two or more tables or worksheets combines their rows together into a single, larger worksheet.

To merge in a table:

1. In the **New operation** panel, click **Merge**.

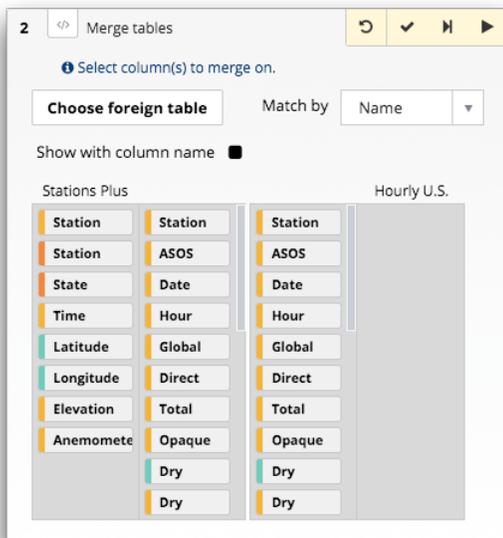
The Trillion-Row Spreadsheet displays the object browser view of the **Merge tables** panel.



2. Select the foreign table by doing one of the following:

- In the table browser, locate and select the foreign table.
- In the **Search** field, enter the table name and press **Enter**.

The Trillion-Row Spreadsheet displays the columns view of the **Merge tables** panel.



3. In the **Match by** drop-down list, select how the columns in each table or worksheet are matched in the merge operation.
For a description of the available options, see [Merge tables panel](#) on page 217.
4. Click the **Submit operation** (✓) icon.
The Trillion-Row Spreadsheet merges the foreign table into the current base table or worksheet and displays the columns from the two tables combined.

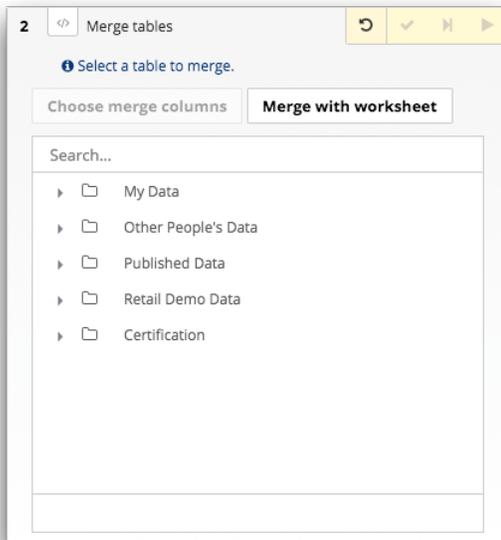
Merge in an existing worksheet

Merge a foreign existing worksheet into the current table or worksheet.

Merging two or more tables or worksheets combines their rows together into a single, larger worksheet.

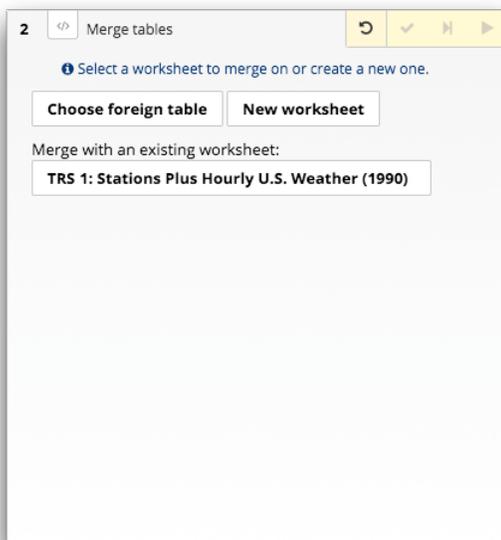
To merge in an existing worksheet:

1. In the **New operation** panel, click **Merge**.
The Trillion-Row Spreadsheet displays the object browser view of the **Merge tables** panel.

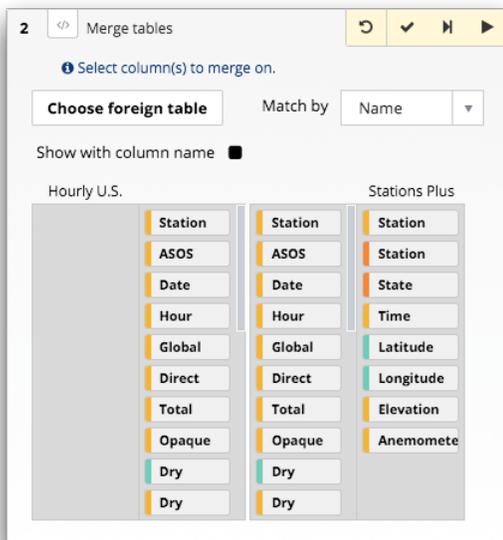


2. Click **Merge with worksheet**.

The Trillion-Row Spreadsheet displays a list of existing worksheets currently open in your session.



3. Click the existing worksheet you want to merge into the current table or worksheet.
The Trillion-Row Spreadsheet displays the columns view of the **Merge tables** panel.



- In the **Match by** drop-down list, select how the columns in each table or worksheet are matched in the merge operation.

For a description of the available options, see [Merge tables panel](#) on page 217.

- Click the **Submit operation** (✓) icon.
The Trillion-Row Spreadsheet merges the foreign existing worksheet into the current base table or worksheet and displays the columns from the two tables combined.

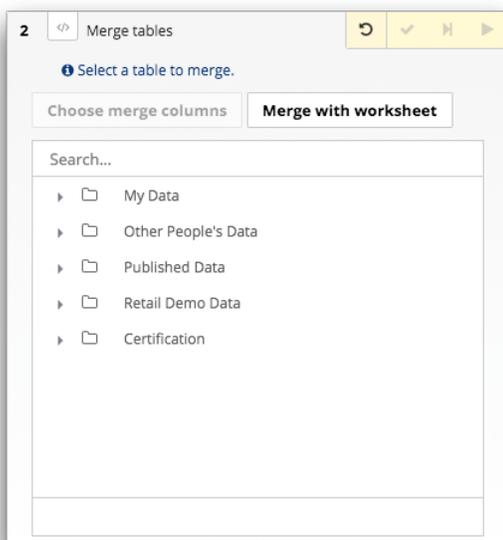
Merge in a new worksheet

Merge a foreign new worksheet into the current table or worksheet.

Merging two or more tables or worksheets combines their rows together into a single, larger worksheet.

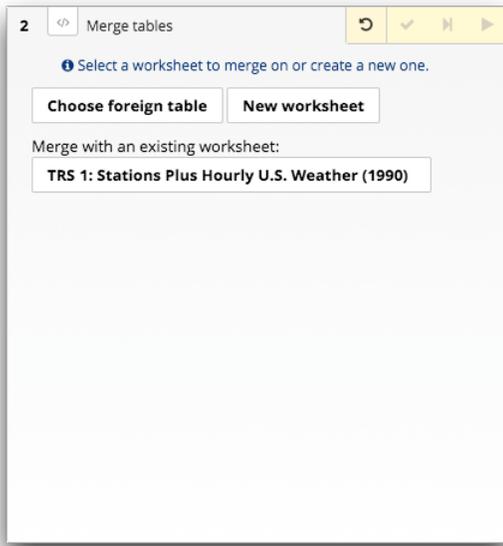
To merge in a new worksheet:

- In the **New operation** panel, click **Merge**.
The Trillion-Row Spreadsheet displays the object browser view of the **Merge tables** panel.



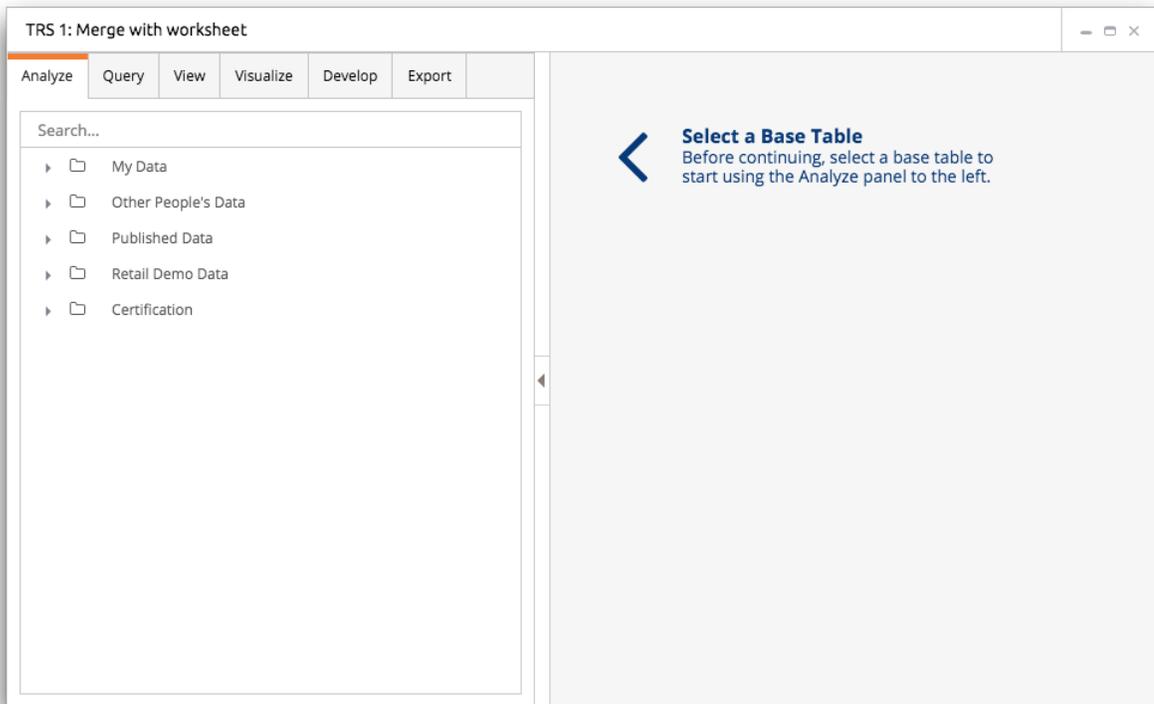
2. Click **Merge with worksheet**.

The Trillion-Row Spreadsheet displays a list of existing worksheets currently open in your session.



3. Click **New worksheet**.

The Trillion-Row Spreadsheet displays the object browser in the **Merge with worksheet** window.



4. Browse to and select a table or query on which you want to create a new worksheet.

The Trillion-Row Spreadsheet displays the **New Operation** panel in the **Analyze** tab on the left of the window and the table on the right.

The screenshot shows the Trillion-Row Spreadsheet interface. On the left, the 'New Operation' panel is open, displaying various data manipulation options: Select, Sort, Arrange, Link, Compute, Tabulate, Merge, Amend, and More... The main window displays a table titled 'Hourly U.S. Weather (1991)' with the following columns: Station ID, Date, Hour, ASOS Flag, Global Radiation (0.1 w/m2 grad), Direct Radiation (0.1 w/m2 drad), and Total Sky Cover tsky. The table contains 15 rows of data for Station ID 3103 on 01/01/91, with hours 1 through 15. The ASOS Flag is 0 for all rows, and the radiation values are consistently 9,999, except for the last row where it is 9.

Station ID	Date	Hour	ASOS Flag	Global Radiation (0.1 w/m2 grad)	Direct Radiation (0.1 w/m2 drad)	Total Sky Cover tsky
3103	01/01/91	1	0	9,999	9,999	99
3103	01/01/91	2	0	9,999	9,999	99
3103	01/01/91	3	0	9,999	9,999	99
3103	01/01/91	4	0	9,999	9,999	99
3103	01/01/91	5	0	9,999	9,999	99
3103	01/01/91	6	0	9,999	9,999	10
3103	01/01/91	7	0	9,999	9,999	10
3103	01/01/91	8	0	9,999	9,999	8
3103	01/01/91	9	0	9,999	9,999	7
3103	01/01/91	10	0	9,999	9,999	8
3103	01/01/91	11	0	9,999	9,999	10
3103	01/01/91	12	0	9,999	9,999	10
3103	01/01/91	13	0	9,999	9,999	10
3103	01/01/91	14	0	9,999	9,999	10
3103	01/01/91	15	0	9,999	9,999	9

5. After performing the desired analysis, return to the **TRS** window that contains the base table into which you want to link the new worksheet. The Trillion-Row Spreadsheet displays the columns view of the **Merge tables** panel.

The screenshot shows the 'Merge tables' panel. It includes a 'Select column(s) to merge on.' section with a 'Choose foreign table' dropdown and a 'Match by' dropdown set to 'Name'. There is a checkbox for 'Show with column name'. Below, two columns of tables are shown, both labeled 'Hourly U.S.'. Each column contains a list of columns: Station, ASOS, Date, Hour, Global, Direct, Total, Opaque, Dry, and another Dry. The 'Dry' columns are highlighted in teal.

6. In the **Match by** drop-down list, select how the columns in each table or worksheet are matched in the merge operation. For a description of the available options, see [Merge tables panel](#) on page 217.
7. Click the **Submit operation** (✓) icon. The Trillion-Row Spreadsheet merges the foreign existing worksheet into the current base table or worksheet and displays the columns from the two tables combined.

Amending values

Replace the values in a table or worksheet using the amend operation.

In a 1010data Insights Platform session, you work from a "snapshot" of the data as it existed when the table was first opened. This prevents the data in your session from changing out from underneath you. You can, however, replace the values of the underlying data within your session. This is helpful in making corrections to the data or to see how different values affect your query results.

In the Trillion-Row Spreadsheet (TRS), the amend operation allows you to replace all values in a column using a value expression, apply the changes to only certain rows in the column using a selection expression, or edit the values of individual cells directly from within the grid. As with other 1010data Insights Platform operations performed in the TRS, the amend operation is displayed as a panel in the Analysis Timeline.

When you amend the values, the changes are not permanent and apply only to the worksheet in your session. You can save the value changes for future use by saving your query. To make any value changes permanent, you must materialize the data to a 1010data Insights Platform table.

Amend column panel

The **Amend column** panel allows you to choose a column and enter a value expression to replace values in a table or worksheet.

The **Amend column** panel is accessible from the following locations in the 1010data Insights Platform:

- Analysis Timeline in the Trillion-Row Spreadsheet
- Right-click menu in the grid

It is important to note that the **Amend column** panel available in the grid is visually different from the panel in the Analysis Timeline. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the **Amend column** panel as it appears in the timeline.

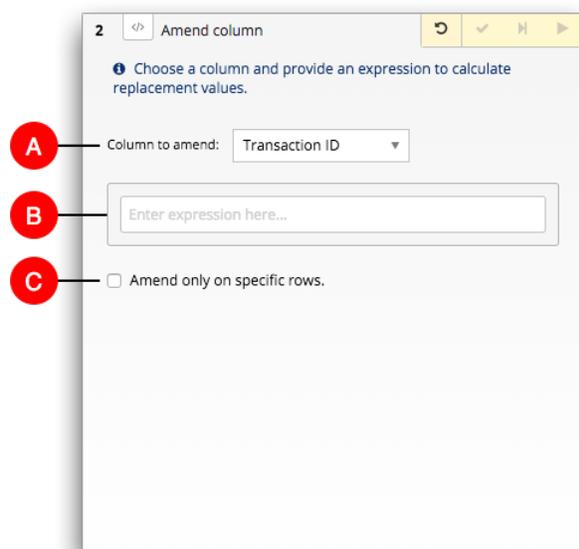


Figure 48: Amend column panel

A. Column to amend

This drop-down list contains the columns in the table or worksheet. By default, the first column in the table or worksheet is listed first in the drop-down list.

Use this drop-down list to select the column containing the values you want to amend.

B. Expression Editor

The Expression Editor field is used to enter a value expression.

Enter a value expression to recompute the values in the selected column. For more information, see [Writing Expressions](#) in the *1010data Reference Manual*. For more information about the Expression Editor, see [Expression Editor](#) on page 139.

C. Amend only on specific rows

Select this option to apply the value expression to only certain rows in the column.

When the **Amend only on specific rows** option is selected, a second Expression Editor field appears in the panel. This additional Expression Editor field allows you to define the specific rows to which you want to apply the amend operation.

Amend values in a column using an expression

Recompute the value of a column or specific rows of a column by using a value expression.

The amend operation allows you to recompute the value of a column or specific rows of a column. Although generally this can be done with a computed column, the amend operation makes it easier, particularly if you only want to recompute specific rows.

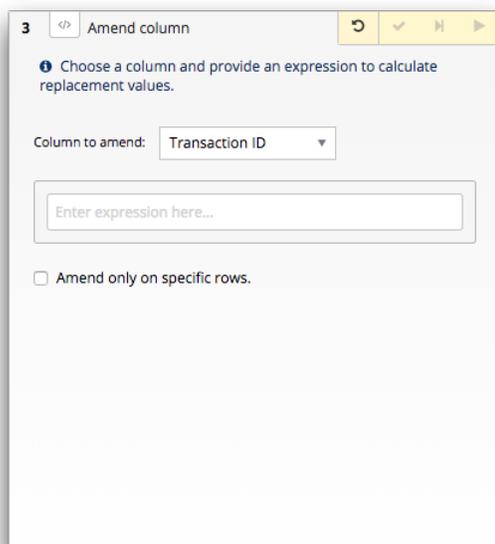
To amend values in a column using an expression:

1. Open the **Amend column** panel by doing one of the following:

- In the **New operation** panel, click **Amend**.
- In the grid, right-click a cell within the column in which you want to amend values and click **Amend [COLUMN_LABEL]**.

Note: In this topic, *[COLUMN_LABEL]* represents the label of the column in which the selected cell is located.

The Trillion-Row Spreadsheet displays the **Amend column** panel.



Note: The **Amend column** panel available in the grid is visually different from the panel in the Analysis Timeline. While the panels vary slightly from one another, the available fields and functionality is identical. For illustration purposes, this topic shows images of the **Amend column** panel as it appears in the timeline.

- In the **Column to amend** drop-down list, select the column containing the values you want to amend.

Note: If you opened the **Amend column** panel from the grid, the **Column to amend** drop-down list should already contain the correct column as it is based on the cell selected in the grid.

- In the Expression Editor field, enter a value expression.

For more information, see [Writing Expressions](#) in the *1010data Reference Manual*. For more information about the Expression Editor, see [Expression Editor](#) on page 139.

- Optionally, to amend only specific rows, complete the following:

- Select the **Amend only on specific rows** option.

The Trillion-Row Spreadsheet displays a second Expression Editor field.

- In the **Amend only on those rows where** Expression Editor field, enter a selection expression to select a subset of rows.

- Click the **Submit operation** (✓) icon.

The Trillion-Row Spreadsheet displays the results of the amend in the grid.

Amend a single value within the grid

Change the value of a cell directly from within the grid.

You can quickly edit the value of a cell in a table or worksheet directly from within the grid in the Trillion-Row Spreadsheet. This allows you to easily update data without needing to write an expression.

To amend a single value within the grid:

- In the Grid view of an open table or worksheet, double-click the cell containing the value you want to amend.

The Trillion-Row Spreadsheet changes the cell to an editable field.

Sales Item Detail
Cols 1 to 8 of 8 , Rows 1 to 19 of 35

Transaction ID transid	Account account	Store store	Date date	Item SKU sku	Units units	Sales sales	Cost cost
531	957	1	05/15/12	366	-1	-5	-1.84
532	478	1	05/15/12	98A	1	0.5	0.25
532	478	1	05/15/12	3B7	1	1.1	0.56
534	738	1	05/16/12	A96	2	6	2.9
534	738	1	05/16/12	65B	1	2.25	1.35
535	709	2	05/15/12	CB7	1	1.65	1.1
535	709	2	05/15/12	96A	1	1.1	1
535	709	2	05/15/12	969	1	1.1	1
536	748	2	05/17/12	3A4	3	1.02	0.39
536	748	2	05/17/12	366	1	5	1.8
537	523	3	05/15/12	CB7	1	1.65	1.1
537	523	3	05/15/12	A96	2	6	2.9
537	523	3	05/15/12	98A	1	0.5	0.25
537	523	3	05/15/12	96A	1	1.1	1
537	523	3	05/15/12	3B7	1	1.1	0.56
538	668	1	05/18/12	CB7	1	1.65	1.1
538	668	1	05/18/12	98A	4	2	1
538	668	1	05/18/12	96A	1	1.1	1

2. Make the desired changes to the value and then do one of the following:

- Press **Enter** (PC).
- Press **Return** (Mac).
- Click anywhere outside of the field.

The Trillion-Row Spreadsheet updates the cell value in the grid and adds the **Amend values** panel to the Analysis Timeline.

The screenshot shows the Trillion-Row Spreadsheet interface. On the left, the 'Amend values' panel is active, displaying the following information:

The following amendments are in effect.

Row	Column	Value
1	account	1957

On the right, the 'Sales Item Detail' table is shown with the following data:

Transaction ID transid	Account account	Store store	Date date	Item SKU sku	Units units	Sales sales	Cost cost
531	1957	1	05/15/12	366	-1		
532	478	1	05/15/12	98A	1	0	
532	478	1	05/15/12	3B7	1	1	
534	738	1	05/16/12	A96	2		
534	738	1	05/16/12	65B	1	2..	
535	709	2	05/15/12	CB7	1	1..	
535	709	2	05/15/12	96A	1	1	
535	709	2	05/15/12	969	1	1	
536	748	2	05/17/12	3A4	3	1..	
536	748	2	05/17/12	366	1		
537	523	3	05/15/12	CB7	1	1..	
537	523	3	05/15/12	A96	2		
537	523	3	05/15/12	98A	1	0	
537	523	3	05/15/12	96A	1	1	
537	523	3	05/15/12	3B7	1	1	
538	668	1	05/18/12	CB7	1	1..	
538	668	1	05/18/12	98A	4		

The **Amend values** panel displays information about the value amended in the grid. Specifically, the panel lists the row and column of the amended cell and the new value. For more information, see [Amend values panel](#) on page 229.

Amend values panel

The **Amend values** panel appears in the Analysis Timeline after a single value is amended within the Grid view of the Trillion-Row Spreadsheet.

The **Amend values** panel displays information about a cell amended within the grid. Specifically, the panel identifies the location of the amended cell, indicated by row number and column name, and lists the cell's new value. A new line of information is created for each cell amendment performed in the grid.

The **Amend values** panel also allows you to edit an existing single value amendment. For example, if you accidentally made an amendment to the wrong cell in the grid, you can make the correction by changing the appropriate information in the panel.

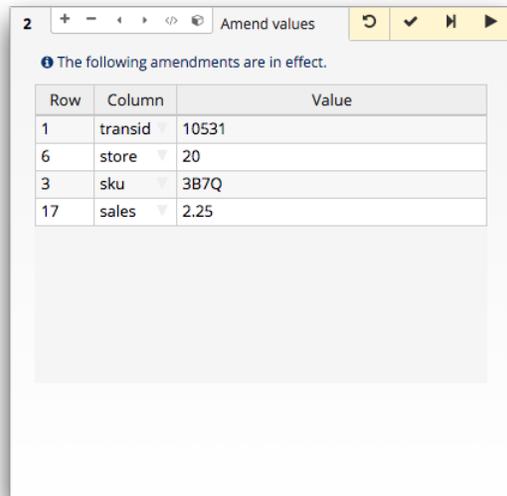


Figure 49: Amend values panel

Row

Displays the row number in which the amended cell is located.

To change the row number of the amended value, double-click the cell, edit the number, and then submit the operation.

To restore a single amended value to its original state, right-click the appropriate row in the **Amend values** panel and then click **Remove row** from the menu.

Column

Displays the name of the column in which the amended cell is located.

To change the column location of the amended value, click the drop-down arrow in the cell, select the column, and then submit the operation.

Value

Displays the new value of the amended cell.

To change the value of the cell, double-click the cell, edit the value, and then submit the operation.

Note: If you attempt to enter a value that does not match the data type of the column (e.g., entering text in an integer column), the cell containing the invalid value is highlighted and the panel prevents you from submitting the value change.

Running custom code

You can run custom SQL, R, or K3 code within the TRS Analysis Timeline.

The **Code** panel within the TRS Analysis Timeline allows you to add SQL, R, or K3 code as part of your analysis. You can use the native language and the Macro Language Workshop will provide the correct syntax to interpret the code correctly within the 1010data Insights Platform.

Code panel

The **Code** panel allows you to insert custom code from SQL, R, or K3 into your TRS Analysis Timeline.

The **Code** panel is accessible from the Analysis Timeline in the Trillion-Row Spreadsheet.

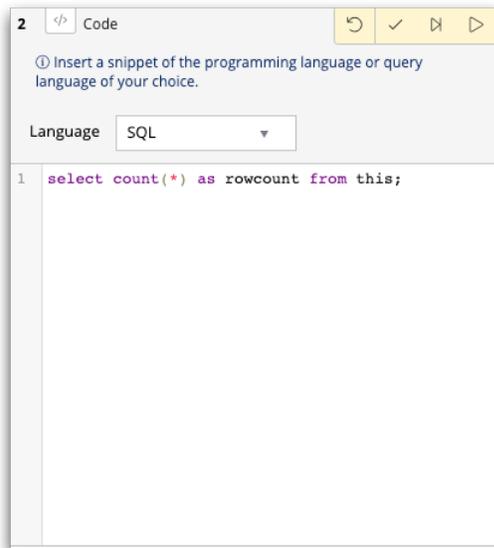


Figure 50: Code panel

Language

Choose one of the following languages for your custom code from the drop-down list:

SQL

This is the equivalent of `<code language_="sql">` in Macro Language code.

R

This is the equivalent of `<code language_="r">` in Macro Language code. For more information about using the R language within the 1010data Insights Platform, see the *1010data R Language Guide*.

Note: Your account must have access to 1010data's latest cloud environments, and must be specifically enabled to make use of the R language, in order to use this feature.

K3

This is the equivalent of `<code language_="k">` in Macro Language code.

For more information, see [code](#) in the *1010data Reference Manual*.

Code Editor

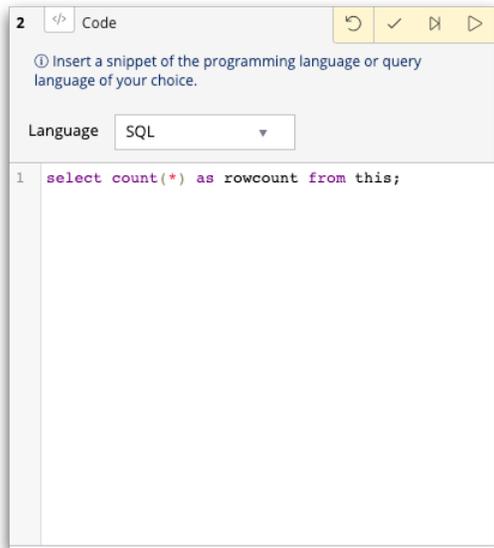
This section contains an editor for the language of your choice from the drop-down list. Each selection contains a code snippet in the selected language. You can use the code snippet as a starting point, or delete it and add your own code.

Use custom code in your analysis

Use the **Code** panel to add custom code to your TRS Analysis Timeline.

To use the TRS Analysis Timeline to insert code:

1. In the **New Operation** panel of the timeline, click **Code**.
The **Code** panel appears in the timeline.



2. Select the desired language from the drop-down list.
The sample code in the code editor changes based on the language you choose.
3. Edit the code in the code editor. Use the syntax of SQL, R, or K3, as you normally would.
Note: Your account must have access to 1010data's latest cloud environments in order to use the R language. For more information about R functions that can be used in the 1010data Insights Platform, see the *1010data R Language Guide*.
4. Click the **Submit operation** (✓) icon.

Miscellaneous operations

In addition to the five basic operations, you can choose to use a miscellaneous operation in the Trillion-Row Spreadsheet.

Miscellaneous operations provide additional functionality to help simplify common analytical tasks.

The following miscellaneous operations are available:

Expand rows

Choose the number of times that each row appears in the table.

Split

Split and expand a string column on a given separator.

Replace

Find and replace one or more patterns within a string column.

Date/Time

Convert strings into dates, times, or intervals and format times and dates into strings.

Group Metric

Compute an aggregation, grouping by rows by one or more columns.

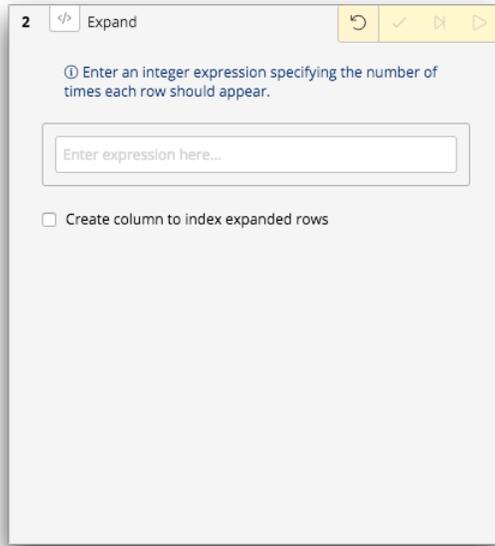
Expand rows

Choose the number of times that each row appears in the table.

The expand rows operation allows you to duplicate rows in your table. You can choose the number of times that each row should appear in the table. You can also create a computed column to index expanded rows, which can make it easier when querying the subsequent data.

To expand rows:

1. In the **New operation** panel, click **Expand Rows**.
The Trillion-Row Spreadsheet displays the **Expand** panel.



2. In the Expression Editor field, enter a value expression.

For more information, see [Writing Expressions](#) in the *1010data Reference Manual*. For more information about the Expression Editor, see [Expression Editor](#) on page 139.

3. Optionally, to create a computed column to index expanded rows, complete the following:
 - a) Select the **Create column to index expanded rows** option.
The Trillion-Row Spreadsheet displays indexing fields and options.

2 Expand

Enter an integer expression specifying the number of times each row should appear.

Enter expression here...

Create column to index expanded rows

Index column name

Starting index 1

- b) In the **Index column name** field, enter a name for the new column.
 - c) In the **Starting index** field, enter or select the number at which to start indexing the expanded rows.
4. Click the **Submit operation** (✓) icon.
The Trillion-Row Spreadsheet displays the results of the expand operation.

Split and expand

Split and expand a string column on a given separator.

The split operation splits each string in a given column based on a specified separator and expands the table so that there is one row for each portion of the split string. A new column contains the individual portions.

For example, the **Address** column in the following sample table contains string data:

Store ID	Store Number	Address	City	State	Zip Code
store_id	Number	addr	city	state	zip
1	136	123 MAIN ST.	CITY	NY	00000

Expanding the **Address** column by a space separator results in the following table:

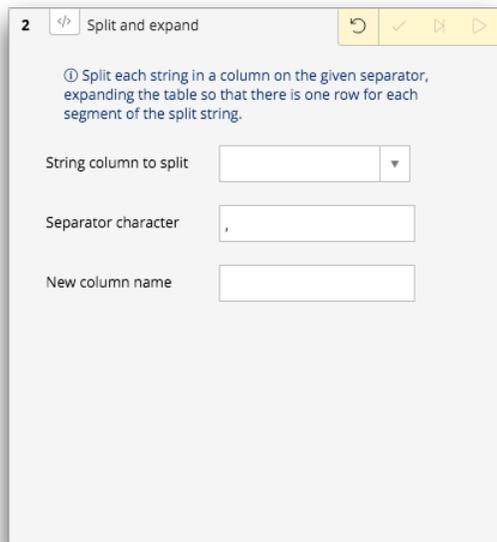
Store ID	Store Number	Address	City	State	Zip Code	expanded
store_id	Number	addr	city	state	zip	expanded
1	136	123 MAIN ST.	CITY	NY	00000	123
1	136	123 MAIN ST.	CITY	NY	00000	MAIN
1	136	123 MAIN ST.	CITY	NY	00000	ST.

Each portion of the address is placed in its own row within the new column, indicated in red above.

To perform a split and expand:

1. In the **New operation** panel, click **Split**.

The Trillion-Row Spreadsheet displays the **Split and expand** panel.



2. Complete the following fields and options:

String column to split

Select the column containing string data that you want to split into multiple rows.

Separator character

Enter the character by which you want to split the data.

New column name

Enter the name for the new column.

3. Click the **Submit operation** (✓) icon.

The Trillion-Row Spreadsheet displays the results of string expansion in the grid.

Perform a string replacement

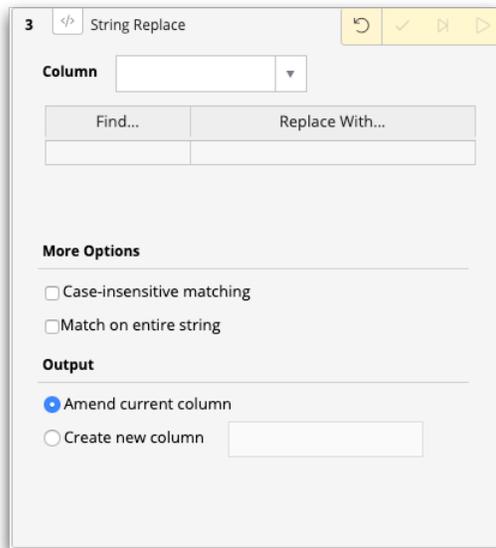
Find and replace one or more string patterns within a string column.

The string replacement operation allows you to select one or more string patterns within a given string column, and replace it with another string.

To replace a string:

1. In the **New operation** panel, click **Replace**.

The Trillion-Row Spreadsheet displays the **String Replace** panel.



2. Select the column where you want to replace text from the drop-down list (text columns only).
3. Under **Find...**, enter the text that you want to replace, and under **Replace With...**, enter the replacement text.
4. Enter additional **Find...** and **Replace With...** text patterns as necessary.
5. Under **More Options**, select **Case-insensitive matching**, if desired.
6. Select **Match on entire string**, if desired. If not selected, the partial text will be replaced with the new text.
7. Under **Output**, select when to **Amend** the text in the existing column or create a new column with the replacement text. If you create a new column, name the new column.
8. Click the **Submit operation** (✓) icon.
The Trillion-Row Spreadsheet displays the results of the string replace operation.

Convert dates and times

Convert strings into dates, times, or intervals, or format time and dates into strings.

You can perform a bulk change of a date field (date, time, or date+time) into a string field or a string field into a date field.

To perform date and time conversion:

1. In the **New operation** panel, click **Date/Time**.
The Trillion-Row Spreadsheet displays the **Date and Time Conversion** panel.

2. From the Column drop-down list, select the column you wish to convert. The lower part of the **Data and Time Conversion** panel changes depending on which column you selected. If you selected an integer column, the panel assumes you have selected a date format and will present string conversion options. If you selected a text column, the panel assumes you want to convert the string to a date format and will present date conversion options.

3. If converting a date/time to string format: Select the format for your date string from the **Format** drop-down list.

Note: The drop-down list also contains an example of each option.

4. If converting a string to a date/time format:
 - a) Select the datatype from the **Datatype** drop-down list.
 - b) Select the order hint from the drop-down list, if applicable.

Note: The order hint helps the operation parse the date correctly if there is any ambiguity to the date string. For example, if the order hint is **Month/Day/Year** and the string is "1/2/3", the operation would parse the date as January 2, 2003.

Note: For detailed information about date/time formats in the 1010data Insights Platform, see [Dates and time](#) on page 55.

5. Enter the name and label for the new column containing the converted data.

Name

This is the name the 1010data Insights Platform uses to interact with the column. It is also used when writing more advanced value expressions and queries.

The column name must be unique. It may only contain alphanumeric characters or underscores and must begin with an alphabetic character (e.g.,

percent_total_sales). It may not contain any spaces or other special characters.

This is a required field.

Note: When entered, the Trillion-Row Spreadsheet (TRS) automatically removes special characters, replaces spaces with underscores, and displays a visual indication that a change was made.

When TRS makes these changes, and if nothing has been entered in the **Label** field, the original text is automatically copied into the **Label** field.

Label

This is the column heading that displays by default at the top of a column in the user interface.

The column label may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters. If you want to have a multi-line column label, use the backtick character (`) to separate the lines (e.g., "Percentage of `Total Sales (%)").

Note: On most standard United States keyboards, the **backtick** (`) character is immediately to the left of **1** key.

While not required, this field is recommended.

- Click the **Submit operation** () icon.
The Trillion-Row Spreadsheet displays the results of the date conversion operation. In this example, the **Date** column is converted from an integer column to a string column called **String Date** with the format "YYYY-MM-DD".

Store	Date	Item SKU	Units	Sales	Cost	String Date
1	05/15/12	366	-1	-5	-1.84	2012-05-15
1	05/15/12	98A	1	0.5	0.25	2012-05-15
1	05/15/12	3B7	1	1.1	0.56	2012-05-15
1	05/16/12	A96	2	6	2.9	2012-05-16
1	05/16/12	65B	1	2.25	1.35	2012-05-16
2	05/15/12	CB7	1	1.65	1.1	2012-05-15
2	05/15/12	96A	1	1.1	1	2012-05-15
2	05/15/12	969	1	1.1	1	2012-05-15
2	05/17/12	3A4	3	1.02	0.39	2012-05-17
2	05/17/12	366	1	5	1.8	2012-05-17
3	05/15/12	CB7	1	1.65	1.1	2012-05-15
3	05/15/12	A96	2	6	2.9	2012-05-15
3	05/15/12	98A	1	0.5	0.25	2012-05-15
3	05/15/12	96A	1	1.1	1	2012-05-15
3	05/15/12	3B7	1	1.1	0.56	2012-05-15
1	05/18/12	CB7	1	1.65	1.1	2012-05-18
1	05/18/12	98A	4	2	1	2012-05-18
1	05/18/12	96A	1	1.1	1	2012-05-18

Use a group metric

Compute an aggregation, grouping rows by one or more columns.

The **Group Metric** operation panel performs functions similar to basic group functions, or `g_functions`, in Macro Language, with a simple user interface. Group functions are used to perform various kinds of calculations across a particular set of rows, grouping by unique values, within one or more columns of a

table. For example, you can calculate total sales by store or average temperature by city. Group functions are similar to tabulations, but they do not result in loss of granularity of the data.

For more information about `g_` functions, see [Group Functions](#), as well as the particular group function, such as `g_sum` and `g_avg`, in the *1010data Reference Manual*.

To perform an aggregation, or group metric:

1. In the **New operation** panel, click **Group Metric**.

The Trillion-Row Spreadsheet displays the **Group Metric** panel.

2. From the **Group by** drop-down list, select the column you wish to group by. You can also select **Check All** to group by all columns.
3. From the **Metric** drop-down list, select the metric by which you want to aggregate the group. For example, you may want to count all the members of the group, in which case you would select **Count**. You may want a sum, such as total sales per store. In this case, you would select the metric of **Sum**.

Note: For the definitions of the different metrics, see [Types of metrics](#).

Depending on the metric chosen, the **Group Metric** panel may display one or more additional options. For example, if you choose **Sum** as the metric, you will have an additional **Data** drop-down list.

4. Choose the data field on which to perform the metric on the data, such as performing a metric of **Sum** on **Sales**.

2 </> Group Metric

Group by: Store

Metric: Sum

Data: Sales

More Options

Filter rows by: None

Select first of each group

Output

New Column:

5. There may be additional drop-down lists, depending on the metric chosen. For example, if you chose **Weighted Average**, you would select the column to apply the weighted average and the column supplying the weights.
6. Optionally, you can choose to filter rows. The filter row you select must contain boolean values, that is, the value in each row must be either 0 (omit) or 1 (include). If you select the default value of **None**, all rows are considered in the group metric.
7. Optionally, you can display only the first value of each group.
8. Enter a name for the new column containing the results of the group metric.
9. Click the **Submit operation** (✓) icon.
The Trillion-Row Spreadsheet displays the results of the group metric operation. In this example, **Group by** is **Store**, the **Metric** is **Sum**, and the **Data** field is **Sales**. The new column is for the group metric is **sales_by_store**.

Sales Item Detail

Cols 3 to 9 of 9, Rows 1 to 19 of 35

Store	Date	Item SKU	Units	Sales	Cost	sales_by_store
1	05/15/12	366	-1	-5	-1.84	23.19
1	05/15/12	98A	1	0.5	0.25	23.19
1	05/15/12	3B7	1	1.1	0.56	23.19
1	05/16/12	A96	2	6	2.9	23.19
1	05/16/12	65B	1	2.25	1.35	23.19
2	05/15/12	CB7	1	1.65	1.1	16.31
2	05/15/12	96A	1	1.1	1	16.31
2	05/15/12	969	1	1.1	1	16.31
2	05/17/12	3A4	3	1.02	0.39	16.31
2	05/17/12	366	1	5	1.8	16.31
3	05/15/12	CB7	1	1.65	1.1	20.35
3	05/15/12	A96	2	6	2.9	20.35
3	05/15/12	98A	1	0.5	0.25	20.35
3	05/15/12	96A	1	1.1	1	20.35
3	05/15/12	3B7	1	1.1	0.56	20.35
1	05/18/12	CB7	1	1.65	1.1	23.19
1	05/18/12	98A	4	2	1	23.19
1	05/18/12	96A	1	1.1	1	23.19

In addition, the name of the panel changes from **Group Metric** to the more descriptive **g_sum by store**.

Types of metrics

The **Metric** drop-down list in the **Group Metric** panel contains a number of group metrics from which you can choose.

You can choose from the following types of group metrics in the **Metric** drop-down list. These metrics correspond to the most commonly used `g_` functions in Macro Language.

Average	<p>The average of the data in the selected Data column for each group. This is similar to using the <code>g_</code> function <code>g_avg</code>. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the average of the non-N/A values). If a group contains only N/A values, the result for the group is N/A.
Weighted average	<p>The average of the data in the selected Data column for each group, weighted by the Weights column for each group. This is similar to using the <code>g_</code> function <code>g_wavg</code>. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If, for a particular group, either the Data column or the Weights column contains N/A values, those rows are ignored (i.e., the result for the group is the weighted average for those rows where neither the Data column nor the Weights column is N/A). If there are no such rows, the result for the group is N/A.
Sum	<p>The subtotal of the Data column for each group. This is similar to using the <code>g_</code> function <code>g_sum</code>. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the sum of the non-N/A values). If a group contains only N/A values, the result for the group is 0.
Sum of squares	<p>The sum of squares of the Data column for each group. This is similar to using the <code>g_</code> function <code>g_sumsqr</code>. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the sum of squares of the non-N/A values). If a group contains only N/A values, the result for the group is 0.
Product	<p>The product of the Data column for each group. This is similar to using the <code>g_</code> function <code>g_prod</code>. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the product of the non-N/A values). If a group contains only N/A values, the result for the group is 1.
Maximum value	<p>The highest number in the Data column for each group. This is similar to using the <code>g_</code> function <code>g_hi</code>. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the highest non-N/A value). If a group contains only N/A values, the result for the group is N/A.
Minimum value	<p>The lowest number in the Data column for each group. This is similar to using the <code>g_</code> function <code>g_lo</code>. <i>(numeric columns only)</i></p>

	<ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the lowest non-N/A value). If a group contains only N/A values, the result for the group is N/A.
Median value	<p>The median of the Data column for each group. This is similar to using the g_function <code>g_median</code>. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the median of the non-N/A values). If there are an even number of values, the median is the average of the two middle values. If a group contains only N/A values, the result for the group is N/A.
Mode value	<p>The mode of the Data column for each group. This is similar to using the g_function <code>g_mode</code>. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the mode of the non-N/A values). If, for a given group, there is no unique mode, the result is arbitrary. If a group contains only N/A values, the result for the group is N/A.
Standard deviation	<p>The standard deviation of the Data column for each group. This is similar to using the g_function <code>g_std</code> <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the standard deviation of the non-N/A values). If a group contains only N/A values, the result for the group is N/A.
Weighted Standard deviation	<p>The standard deviation of the Data column, weighted by the Weights column, for each group. This is similar to using the g_function <code>g_wstd</code>. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If, for a particular group, either the Data column or the Weights column contains N/A values, those rows are ignored (i.e., the result for the group is the weighted standard deviation for those rows where neither the Data column nor the Weights column is N/A). If there are no such rows, the result for the group is N/A.
Variance	<p>The variance of the Data column for each group. This is similar to using the g_function <code>g_var</code>. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the variance of the non-N/A values). If a group contains only N/A values, the result for the group is N/A.
Weighted variance	<p>The variance of the Data column, weighted by the Weights column, for each group. This is similar to using the g_function <code>g_wvar</code>. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If, for a particular group, either the Data column or the Weights column contains N/A values, those rows are ignored (i.e., the result for the group is the weighted variance for those rows where neither the Data column nor the Weights column is N/A). If there are no such rows, the result for the group is N/A.
Correlation	<p>The correlation between Column A and Column B for each group. This is similar to using the g_function <code>g_cor</code>. <i>(numeric columns only)</i></p> <ul style="list-style-type: none"> If, for a particular group, either Column A or Column B contains N/A values, those rows are ignored (i.e., the result for the group is the correlation for those rows where neither Column A nor Column B is N/A).

	<ul style="list-style-type: none"> If there aren't at least two such rows in the group, the result is N/A.
Covariance	<p>The covariance of Column A and Column B for each group. This is similar to using the g_function <code>g_cov</code>. (<i>numeric columns only</i>)</p> <ul style="list-style-type: none"> If, for a particular group, either Column A or Column B contains N/A values, those rows are ignored (i.e., the result for the group is the covariance for those rows where neither Column A nor Column B is N/A). If there are no such rows, the result for the group is N/A.
Count	<p>The number of rows in each group. This is similar to using the g_function <code>g_cnt</code>.</p> <ul style="list-style-type: none"> The number of rows includes those with N/A values.
Count of unique values	<p>The number of unique values in the Data column for each group. This is similar to using the g_function <code>g_ucnt</code>.</p> <ul style="list-style-type: none"> If a group contains N/A values, those rows are ignored (i.e., the result is the number of unique non-N/A values). If a group contains only N/A values, the result for the group is 0.
Count of non-N/A values	<p>The number of non-N/A values in the Data column for each group. This is similar to using the g_function <code>g_valcnt</code>.</p>
Count of N/A values	<p>The number of N/A values in the Data column for each group. This is similar to using the g_function <code>g_nacnt</code>.</p>
Dot product	<p>The dot product of Column A and Column B for each group. This is similar to using the g_function <code>g_dot</code>.</p> <ul style="list-style-type: none"> If, for a particular group, either Column A or Column B contains N/A values, those rows are ignored (i.e., the result for the group is the dot product for those rows where neither Column A nor Column B is N/A). If there are no such rows, the result for the group is 0.
Logical AND	<p>Returns a boolean value indicating whether all values in the Data column within a given group are 1 (true). This is similar to using the g_function <code>g_and</code>.</p> <ul style="list-style-type: none"> The Data column must contain a boolean value (0 or 1). N/A values are ignored (i.e., the result is whether all non-N/A values within a given group are true). If all values for a group are N/A, the result is 1 (true).
Logical OR	<p>Returns a boolean value indicating whether any value in the Data column within a given group is 1 (true). This is similar to using the g_function <code>g_or</code>.</p> <ul style="list-style-type: none"> The Data column must contain a boolean value (0 or 1). N/A values are ignored (i.e., the result is whether any non-N/A values within a given group are true). If all values for a group are N/A, the result is 0 (false).

Query

The **Query** tab displays the current timeline query as Macro Language XML code. It also provides the ability to undo and redo operations in the timeline.

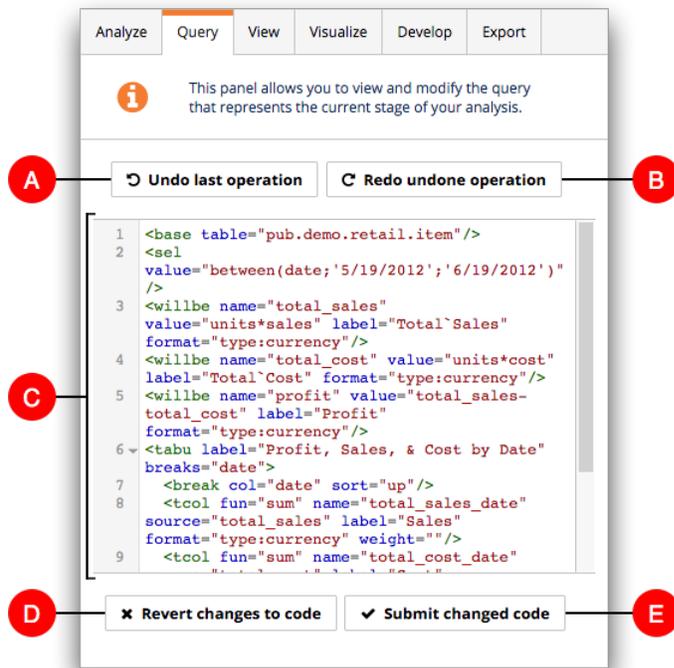


Figure 51: Query tab in TRS

A. Undo last operation

Click this button to undo the last operation added to the timeline.

When the Trillion-Row Spreadsheet is the active window, you can also use the following keyboard shortcuts to undo the last operation:

- **Ctrl+Z** (PC)
- **Command+Z** (Mac)

All operations, back to the initial state of the Trillion-Row Spreadsheet, can be undone.

B. Redo undone operation

Click this button to redo the last undone operation in the timeline.

You can redo an undone operation if no other changes to the timeline have occurred since the undo action.

When the Trillion-Row Spreadsheet is the active window, you can also use the following keyboard shortcuts to redo the last undone operation:

- **Ctrl+Shift+Z** (PC)
- **Shift+Command+Z** (Mac)

If multiple operations have been undone, they can be restored by performing multiple redo actions (as long as no other changes to the timeline have occurred).

C. Code editor

The Code editor displays the query in the Trillion-Row Spreadsheet timeline as Macro Language XML code. Regardless of the step selected in the timeline, the full query is shown in the Code editor and is updated whenever the timeline changes.

You can use the Code editor to make changes to the query using Macro Language XML code.

D. Revert changes to code

Click this button to undo any changes made to the Macro Language XML code.

This button becomes active after changes are made to the query code in the Code editor.

E. Submit changed code

Click this button to submit the changes made to the Macro Language XML code.

This button becomes active after changes are made to the query code in the Code editor.

Edit a query using Macro Language XML code

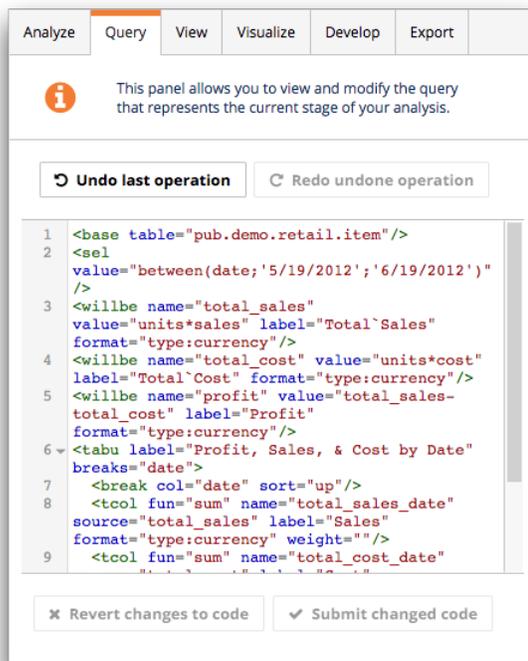
Modify the Macro Language XML query code that represents your analysis in the Trillion-Row Spreadsheet timeline.

When you perform an analysis in the Trillion-Row Spreadsheet, the **Query** tab displays the analysis in the timeline as Macro Language XML code. You can use the Code editor in the **Query** tab to edit the timeline query using Macro Language code.

To edit a query using Macro Language XML code:

1. In the Trillion-Row Spreadsheet, click the **Query** tab.

The Trillion-Row Spreadsheet displays the Macro Language XML code that represents your query.



2. In the Code editor, modify the query as desired.

3. Click **Submit changed code**.

The Trillion-Row Spreadsheet updates the query in the timeline and displays the results in the currently selected view.

View

Select among the ways to view and interact with the data resulting from your analysis.

By default, the results from your data analysis are displayed in the grid inside the results pane. You can choose to view the data in different formats, some of which also provide additional ways to interact with the data.

You can also select a view from the View bar in the analysis pane of the Macro Language Workshop. For more information, see [View bar](#) on page 292.

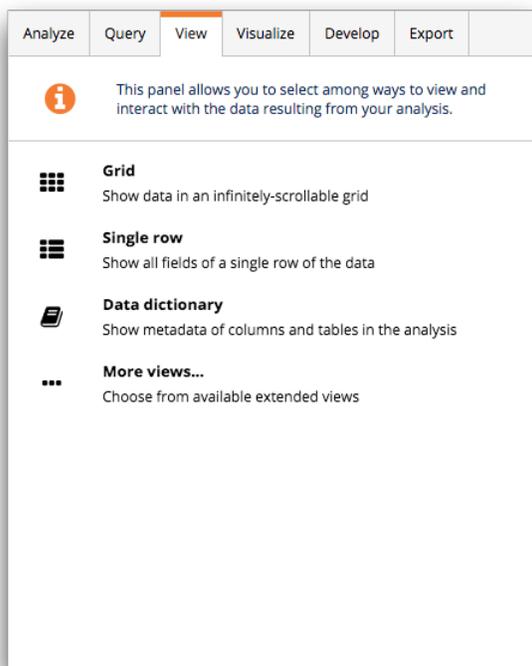


Figure 52: View tab in TRS

Grid

Displays data in an infinitely-scrollable grid. This is the default view.

For more information, see [Grid view](#) on page 294.

Single row

Displays all fields of a single row of the data arranged vertically. This view is particularly useful for viewing information in a table that has many columns.

For more information, see [Single-row view](#) on page 325.

Data dictionary

Displays the metadata of columns and tables in the analysis.

More views

Displays the extended views available to you.

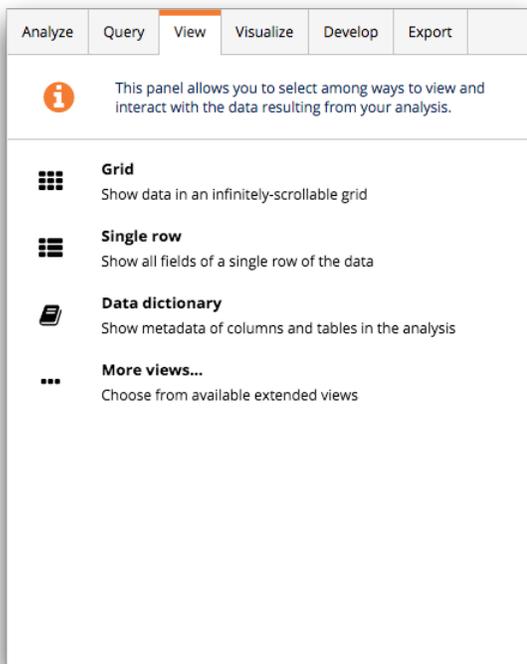
Select a view

Choose how data is displayed in the results pane of the Trillion-Row Spreadsheet.

When using the Trillion-Row Spreadsheet, you can choose how the data resulting from your analysis is displayed in the results pane.

To select a view:

1. In the Trillion-Row Spreadsheet, click the **View** tab.
The Trillion-Row Spreadsheet displays the available view options.



2. Click the desired view.

For a description of the available views, see [View](#) on page 244.

The Trillion-Row Spreadsheet displays the results of your analysis in the selected view within the results pane.

Visualize

The 1010data Insights Platform provides a mechanism to generate high-quality charts, which in turn can support the notion of visual data discovery.

Charting can help you make sense of large amounts of data that might otherwise be too overwhelming to comprehend. Imagine a table with 20 columns and 100,000 rows of data. With 2 million pieces of important information right in front of you, it might be difficult to get a clear picture of what it all means. With charting, you can more easily see relationships among the data points so you can quickly glean insights and use them to improve your business.

Numerous chart types are supported in the Insights Platform, including two-dimensional bar charts, pie charts, histograms, and bubble charts, as well as three-dimensional scatter charts and surface charts. In addition, stacked, stepped, and percentile bar chart types are provided. A number of chart types for visualizing financial data are also available, including Kagi, Renko, and candlestick charts. All chart types allow full aspect ratio rendering, so that the chart will fill the whole area when you resize the chart window.

Useful charts can be created with minimal effort using the web interface by applying one of the provided chart types to the current table or worksheet. You can use the Chart Builder to select the columns of data to be plotted and provide other chart-related parameters, such as the number of samples for a chart. You can select one of the predefined themes to display your chart or graph with a particular color palette, or you can modify individual attributes such as the color, style, or size of the title font; the degree of rotation for the ticks along a particular axis; or whether a legend should be included and where it should appear in relation to the chart.

You can also download any of your charts in PNG format, making your data available for internal or external business reports.

Chart Builder

The Chart Builder allows you to create new charts or modify the appearance of existing charts using a drag-and-drop interface and customizable settings.

The Chart Builder is presented in the **Visualize** tab of the **TRS** window. It consists of multiple sections in the analysis pane that allow you to configure the type of chart, the columns of data to be plotted, and a number of customizable settings such as the size of the chart, its background color, title, theme, axes, ticks, and legend. Any changes with respect to these settings are reflected immediately in the chart. The chart is displayed in the results pane on the right.

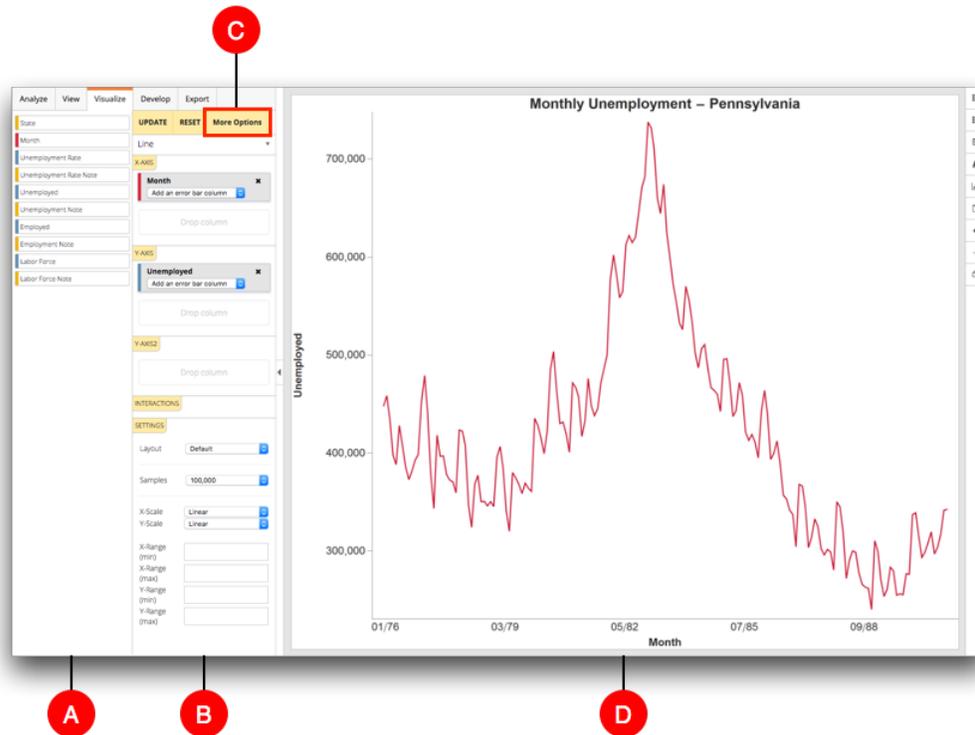


Figure 53: Chart Builder

A. Data columns

Lists all of the columns from the current table or worksheet that are available for plotting in a chart.

B. Chart parameters

Displays the options to select the chart type and the data to be plotted.

C. Customization settings

When the data columns and chart parameters panels are displayed, the customization settings are not visible. To access the customization settings panel, click the **More Options** button.

Customization settings allow you to configure the appearance of the chart.

D. Chart

The chart is displayed in the results pane of the **TRS** window.

Data columns

Columns available for plotting in a chart are displayed in the data columns panel of the Chart Builder.

The data columns panel, outlined in red below, is located on the left side of the analysis pane in the **Visualize** tab. The panel contains a list of all the columns from the current table or worksheet.

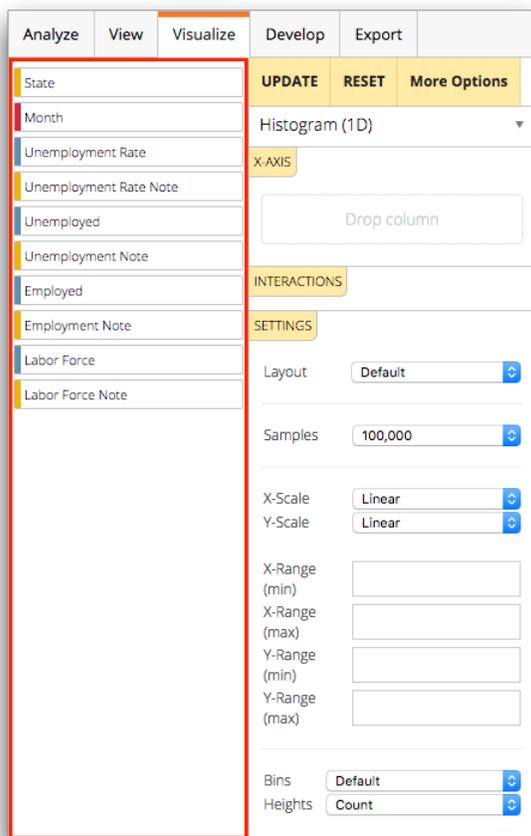


Figure 54: Data columns panel

Each column in the panel is identified by its column label. The color on the left side of each column indicates the type of data values in that column. For more information about determining the data type of a column, see [Column information](#) on page 49.

Any column in the data columns panel can be dragged to a **Drop column** area in the chart parameters panel for plotting. For more information, see [Chart data](#) on page 257.

Chart parameters

Chart-specific settings are displayed in the chart parameters panel of the Chart Builder.

The chart parameters panel, outlined in red below, appears on the right side of the analysis pane in the **Visualize** tab. This panel is used to select the chart type, specify chart-specific settings, specify the columns of data to be plotted, and enable chart interactions.

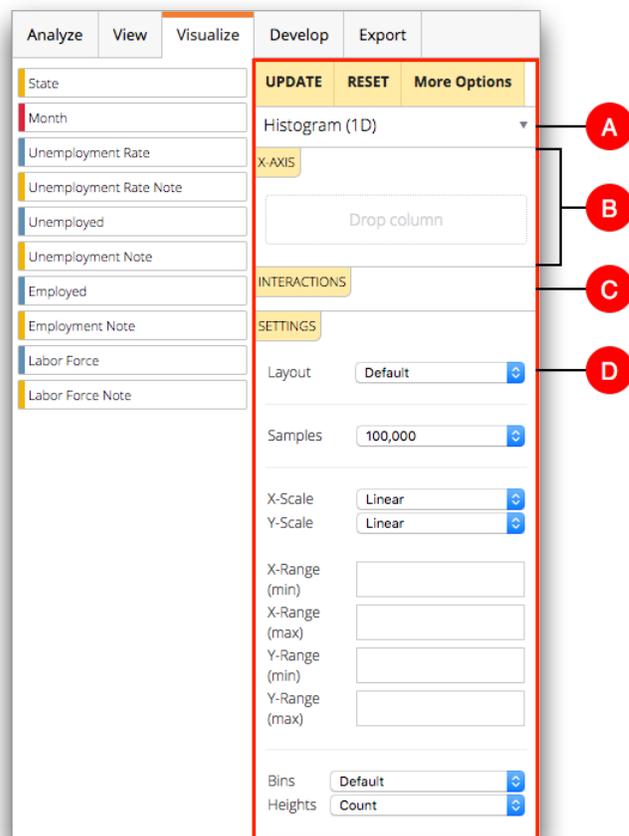


Figure 55: Chart parameters panel

The chart parameters panel consists of the following:

A. Chart type

The chart type drop-down list is used to select the type of chart you want to create.

The available fields and options in the other sections of the chart parameters panel are dependant on the chart type selected.

B. Chart data

Depending on the selected chart type, one or more chart data sections are displayed. These sections are labeled with the type of data required to plot the selected chart. (e.g., **X-AXIS**, **BARS**, **LABELS**, etc.)

C. Interactions

Chart interactions allow you to enable certain mouse actions in the chart. For example, you can zoom, recenter, or select data points.

Note: Chart interactions are provided only for the **Line** and **Scatter** chart types.

D. Settings

This section displays a number of general parameters related to the chart, such as the number of samples to include in the chart or whether to use a linear or logarithmic scale for a particular axis. Additional options specific to the selected chart type may also appear.

Chart types

Numerous chart types are supported in the 1010data Insights Platform, including two-dimensional line, bar, pie, and bubble charts; three-dimensional scatter and surface charts; and chart types for visualizing financial data such as Kagi, Renko, and Candlestick charts.

When you first open the **Visualize** tab, the **Histogram (1D)** chart type is selected by default.

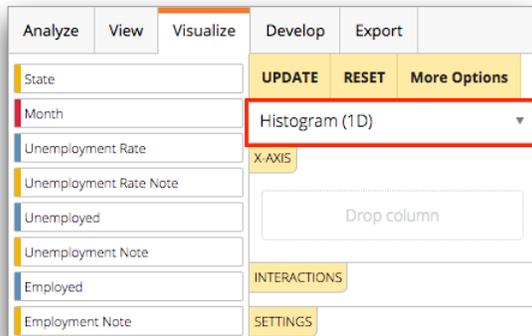


Figure 56: Chart type drop-down list

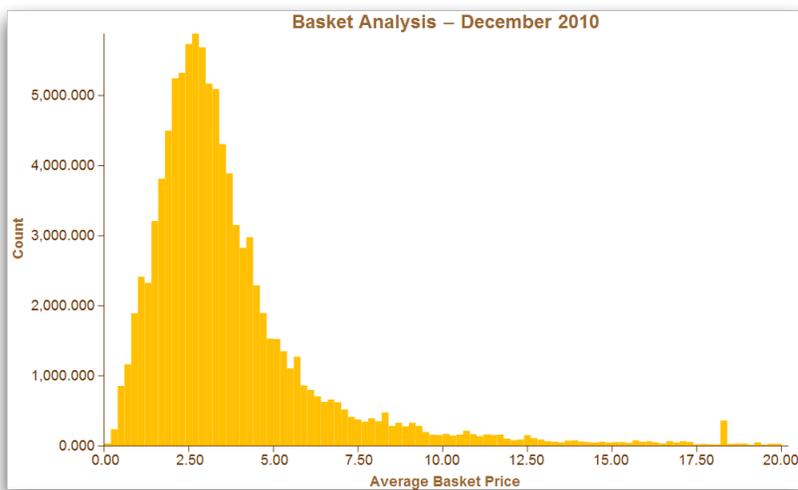
The chart type drop-down list allows you to choose the type of chart you want to create. After selecting a chart type, fields and options related to the chart type appear in the chart data, **Interactions**, and **Settings** sections. The available fields and options in these sections are dependant on the chart type selected.

Information about each of the different chart types available in the Chart Builder is provided below.

Histogram (1D)

A chart that shows the distribution of data using contiguous columns (bins).

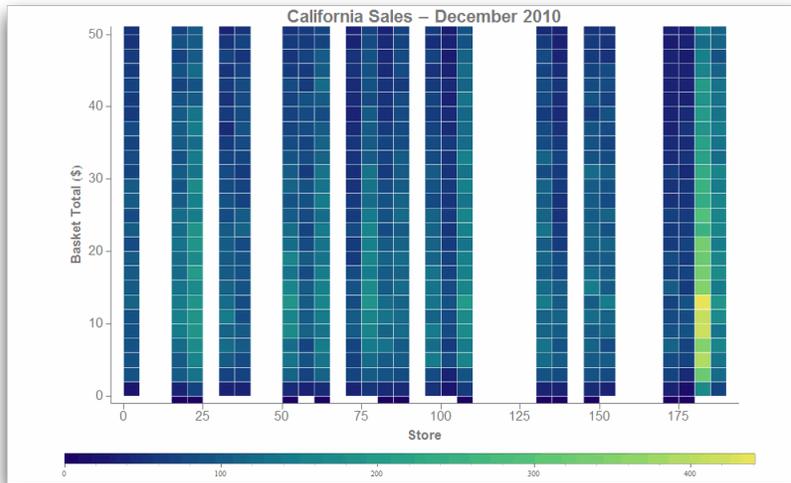
The height of each bin, which represents a grouping of data, indicates the frequency of occurrence of that particular group within the given data set.



Histogram (2D)

A chart that uses bins of equal size to show the distribution of data.

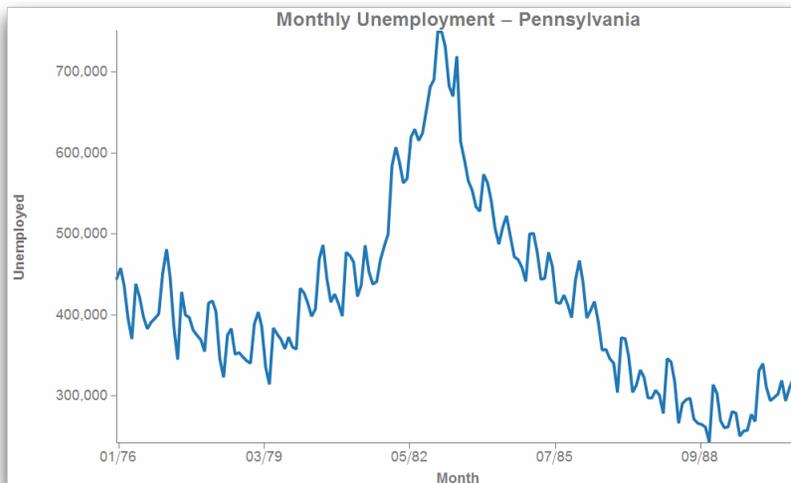
The density of the distribution is indicated by the coloring of the bins.



Line

A two-dimensional chart that displays a series of data points interconnected by line segments.

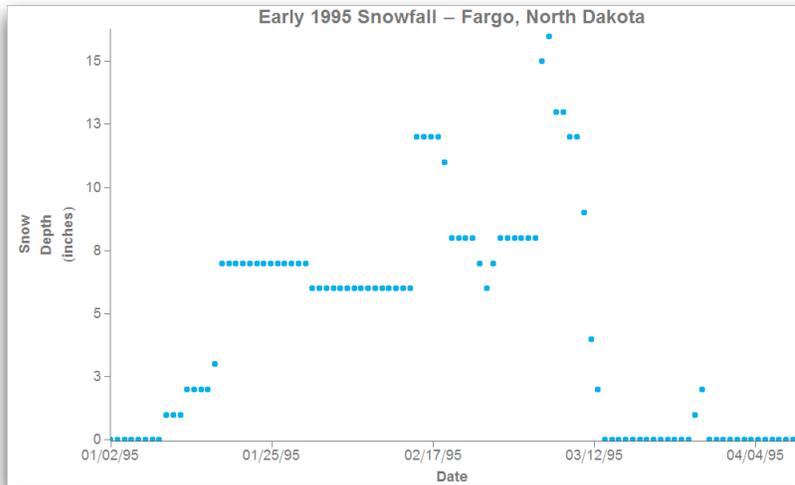
Line charts can show the change in a particular variable over time and are often used to identify trends in data.



Scatter

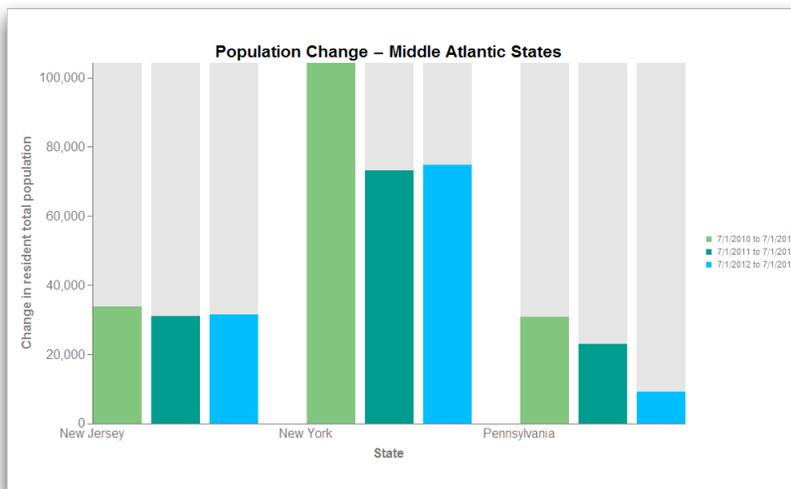
A two-dimensional chart that plots a set of related values as a collection of individual data points.

Scatter charts are often used to show the non-linear relationship between variables. They help to visualize the distribution of data points within a particular data set and can help to identify any outliers.



Bar

A chart with either horizontal or vertical bars that proportionally represent the values of various subsets of data.

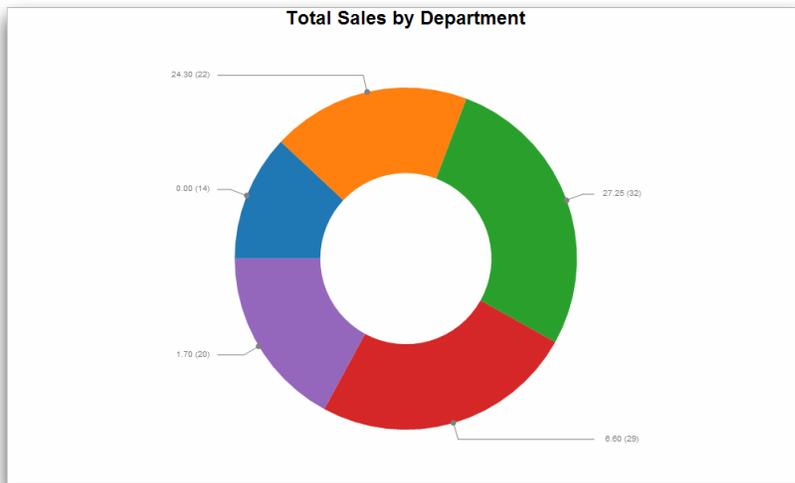


Pie

A circular chart that is divided into sections representing various subsets of data in relation to the whole.

Pie charts are used to show the proportional relationship between various subsets of data and are most effective when the number of subsets is relatively small.

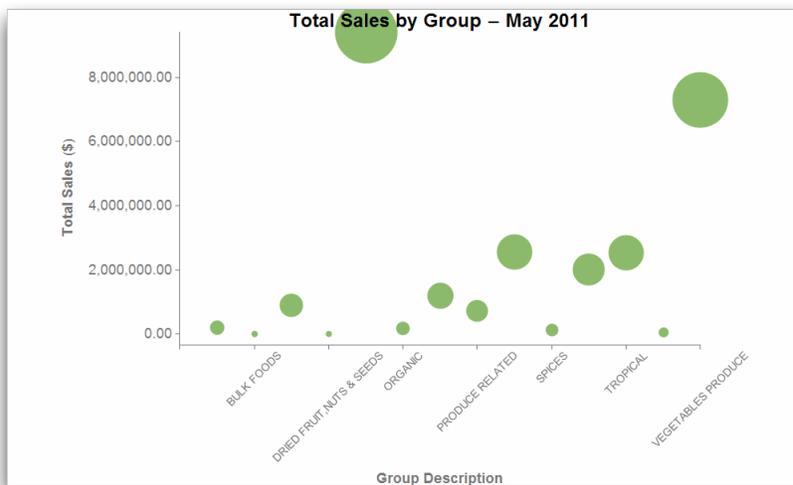
A *donut chart* is a variation of the pie chart, with a blank center in the middle of the chart.



Bubble

A chart that displays each data point as a bubble.

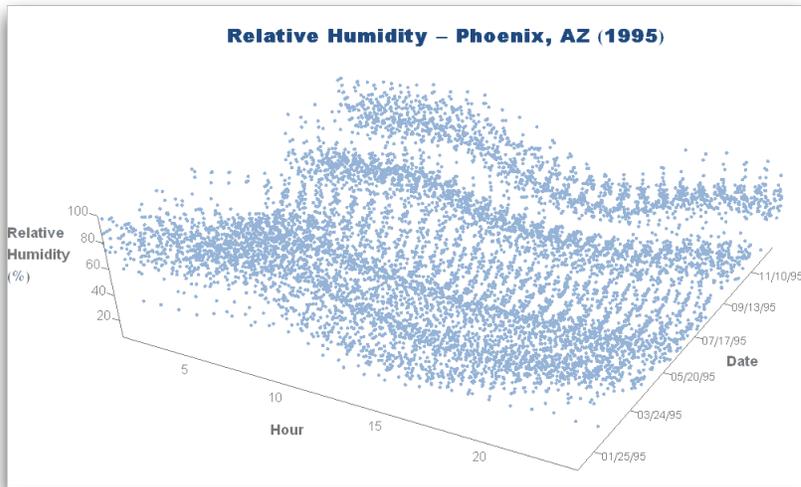
The position of the bubble represents two dimensions of the data, and the size of the bubble corresponds to a third dimension.



Scatter (3D)

A three-dimensional chart that is often used to show the change in variables over time.

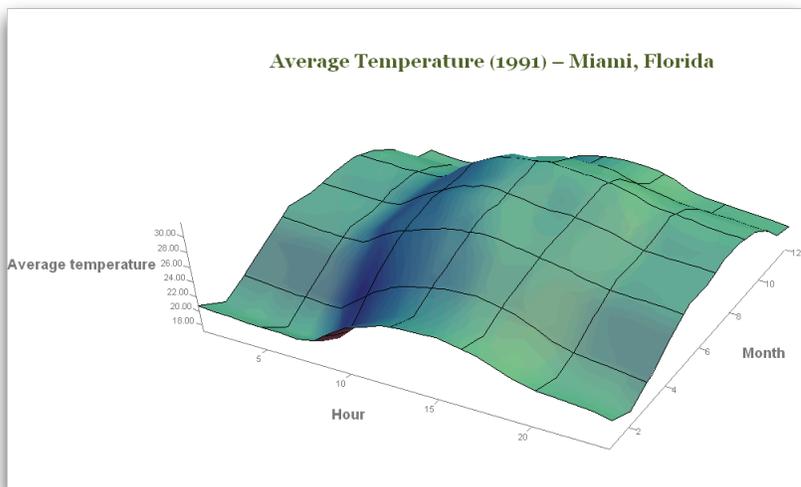
Scatter (3D) charts show concentrations of data points and patterns within a given data set.



Surface

A chart that displays a three-dimensional graphic of interrelated data points.

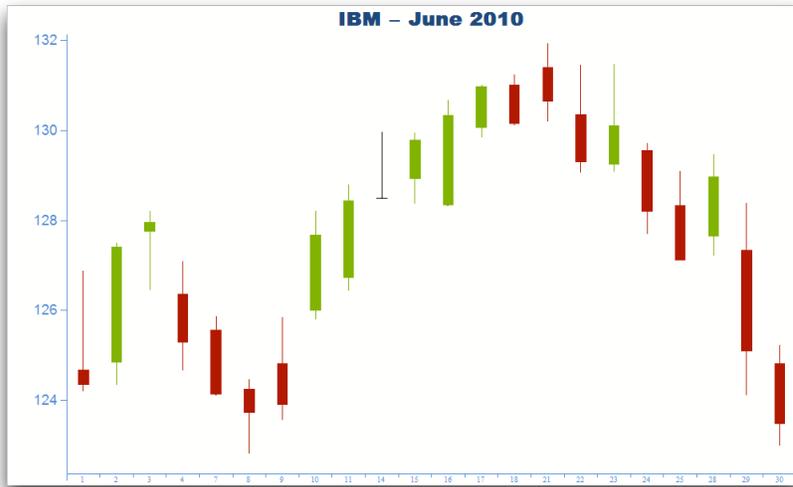
Surface charts use wireframes and shading to show relationships among the variables.



Candlestick

A chart that is typically used to illustrate the price movement of a security, currency, or derivative over time.

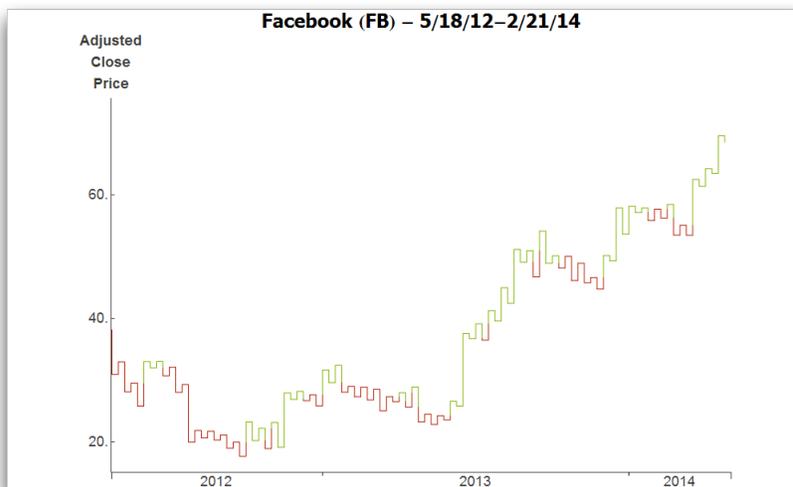
Candlestick charts consist of entries that use the candle "body" to represent the opening and closing price and the "wick" to represent the highest and lowest trades of the security on a particular day.



Kagi

A chart that shows the trend of a financial asset, irrespective of its day-to-day fluctuations.

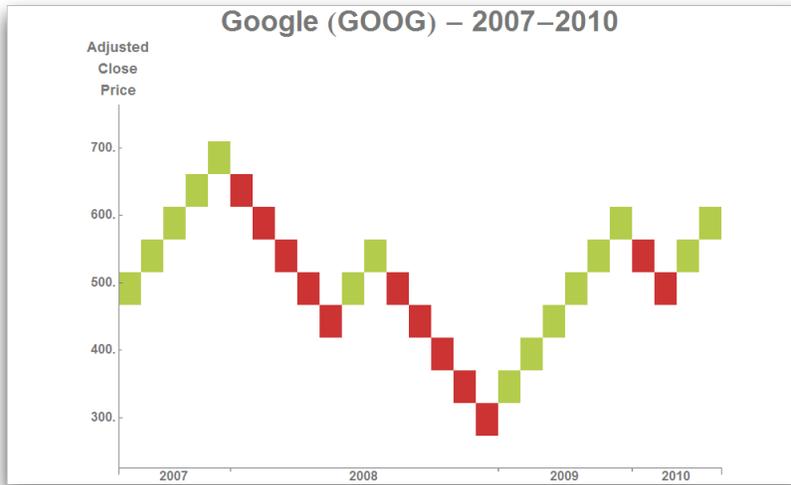
The vertical lines in a Kagi chart depend on price action; horizontal lines connect changes in the direction of its movement.



Renko

A chart that uses "bricks" of a predetermined size to show the price movement of a stock or other financial asset.

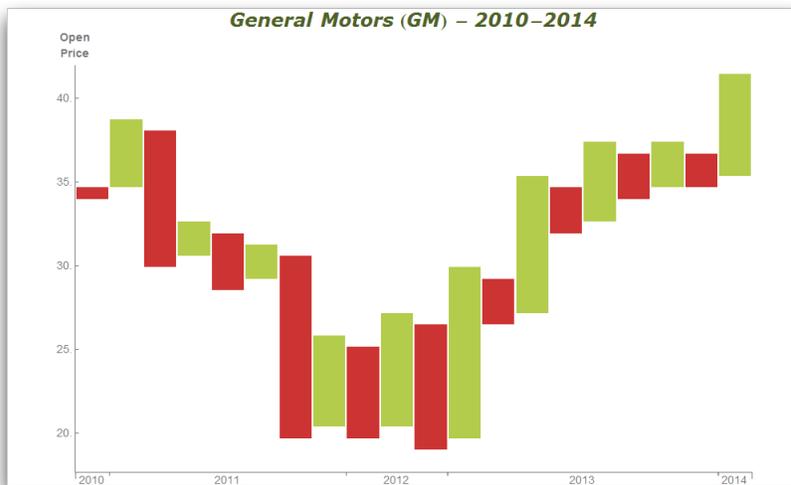
A new brick is added when the price surpasses the top or bottom value of the previous brick. Because Renko charts are time independent, the time axis is not constantly spaced.



Point & Figure

A chart that consists of alternating green and red columns, which indicate the rising and falling price movements of a financial asset over a given date range.

The column changes when the price changes direction by a set value.



Line Break

A chart that is useful for detecting changes in trends for a particular security over a specified time interval.

A new line is drawn when the closing price exceeds the previous day's high or low. Consecutive green lines indicate a trend in higher closing prices; consecutive red lines show a trend in lower closing prices.

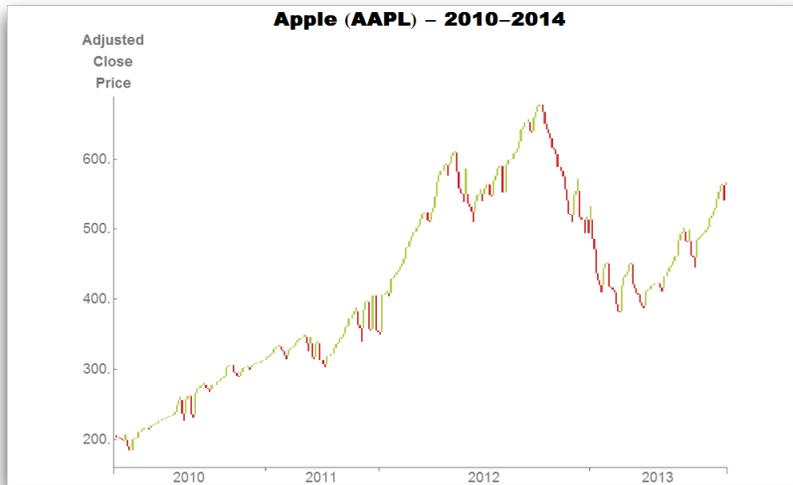


Chart data

Depending on the chart selected in the chart type drop-down list, one or more chart data sections appear.

In the figure below, the red box outlines the **X-AXIS**, **Y-AXIS**, and **Y-AXIS2** sections. These sections are available for plotting the selected **Line** chart type and indicate how that data is plotted.

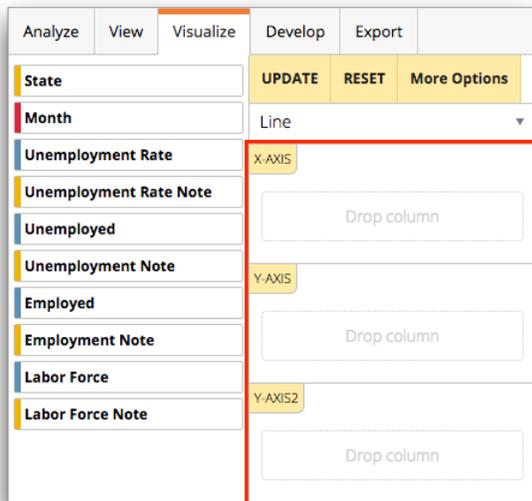


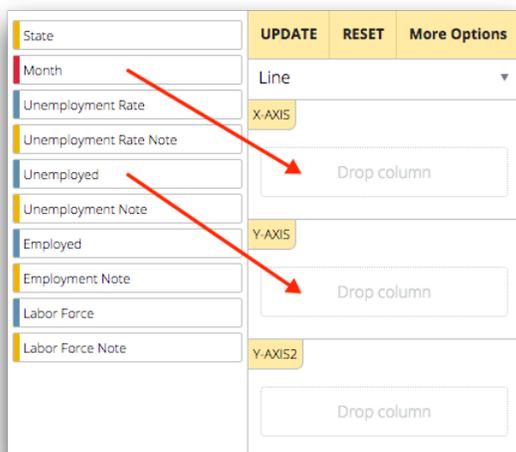
Figure 57: Chart data sections

Depending on the selected chart type, different sections are displayed. For example, a line chart has areas for **X-AXIS** and **Y-AXIS** while a bar chart has areas for **LABELS** and **BARS**. The **Drop column** area in each chart data section displays the data required to plot the selected chart type.

Any column in the data columns panel can be dragged to a **Drop column** area in the chart parameters panel for plotting.

Note: Columns cannot be moved from one chart data section to another.

In the following example, if you wanted to plot the values associated with the column heading **Month** along the X-axis, you would drag that column into the **Drop column** area under the **X-AXIS** section in the chart parameters panel. If you wanted to plot the values associated with the column heading **Unemployed** along the Y-axis, you would drag that column into the **Drop column** area under **Y-AXIS**.



Interactions

Chart interactions allow you to enable certain mouse actions in the chart.

Chart interactions allow you to zoom in or out of the current chart or select data points individually or as a range in the current chart. You can also center the chart around a certain point by double-clicking in the desired spot on the chart.

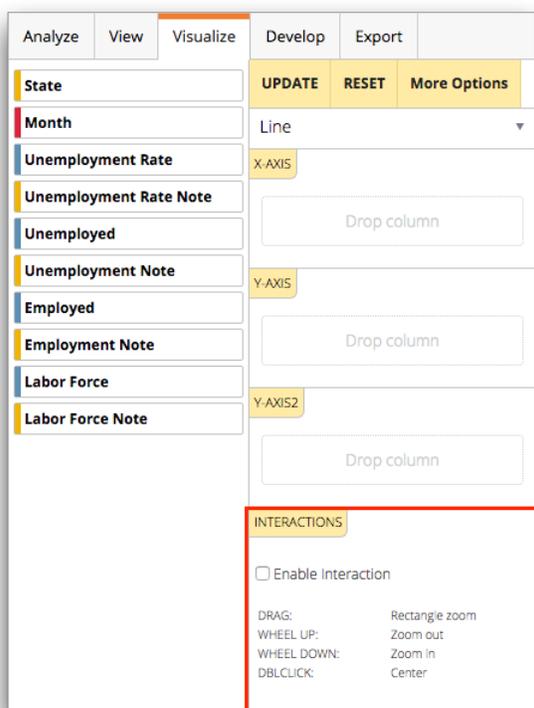


Figure 58: Interactions section

Once chart interactions are enabled, you can perform the following mouse actions on the chart:

Action	Description
Drag	Drag the pointer over an area to zoom into the corresponding rectangular region.

Action	Description
Wheel up	Rotate the wheel button forward to zoom out of the chart.
Wheel down	Rotate the wheel button backward to zoom in to the chart.
Double-click	Double-click at a particular spot on a chart to center the chart at that location.

Note: Chart interactions are provided only for the **Line** and **Scatter** chart types.

Settings

The **Settings** section displays a number of general parameters related to the chart, such as the number of samples to include in the chart or whether to use a linear or logarithmic scale for a particular axis.

Fields and options in the **Settings** section are specific to the chart type selected in the chart type drop-down list. For instance, if you select the **Bar** chart type, you will see settings for **Bar orientation** and **Layout**.

The image shows a software interface with a top navigation bar containing 'Analyze', 'View', 'Visualize', 'Develop', and 'Export'. Below this is a list of data fields on the left, including 'State', 'Month', 'Unemployment Rate', 'Unemployment Rate Note', 'Unemployed', 'Unemployment Note', 'Employed', 'Employment Note', 'Labor Force', and 'Labor Force Note'. To the right of these fields is a configuration panel with tabs for 'UPDATE', 'RESET', and 'More Options'. The 'Visualize' tab is active, showing a 'Line' chart type and three axes labeled 'X-AXIS', 'Y-AXIS', and 'Y-AXIS2', each with a 'Drop column' placeholder. Below these are 'INTERACTIONS' settings, including an 'Enable Interaction' checkbox and a list of actions: DRAG (Rectangle zoom), WHEEL UP (Zoom out), WHEEL DOWN (Zoom in), and DBLCLICK (Center). The 'SETTINGS' section, highlighted with a red border, includes:

- Layout: Default (dropdown)
- Samples: 100,000 (dropdown)
- X-Scale: Linear (dropdown)
- Y-Scale: Linear (dropdown)
- X-Range (min): [input field]
- X-Range (max): [input field]
- Y-Range (min): [input field]
- Y-Range (max): [input field]

Figure 59: Settings section

For a detailed listing of the general and chart-specific parameters for each chart type, see the appropriate section below.

Histogram (1D)

Fields and options available for the **Histogram (1D)** chart type include:

Layout

Default

Stacked

Samples

Specify the number of samples to use when plotting a chart.

X-Scale

Linear

Use a linear scale for the X-axis.

Log

Use a logarithmic scale for the X-axis.

Y-Scale**Linear**

Use a linear scale for the Y-axis.

Log

Use a logarithmic scale for the Y-axis.

X-Range (min)

Specify the minimum value to be plotted along the X-axis.

X-Range (max)

Specify the maximum value to be plotted along the X-axis.

Y-Range (min)

Specify the minimum value to be plotted along the Y-axis.

Y-Range (max)

Specify the maximum value to be plotted along the Y-axis.

Bins

You can select a particular binning method or choose a specific bin size.

The following are the available binning methods:

Default**Sturges**

Compute the number of bins based on the length of data

Scott

Asymptotically minimize the mean square error

Freedman-Diaconis

Twice the interquartile range divided by the cube root of sample size

Knuth

Balance likelihood and prior probability of a piecewise uniform model

Wand

One-level recursive approximate Wand binning

The following are the available bin sizes:

10

20

50

100

Heights

Count

The number of values lying in each bin

Cumulative Count

Cumulative counts

Survival Count

Survival counts

Probability

Fraction of values lying in each bin

CDF

Cumulative distribution function

SF

Survival function

HF

Hazard function

CHF

Cumulative hazard function

Histogram (2D)

Fields and options available for the **Histogram (2D)** chart type include:

Samples

Specify the number of samples to use when plotting a chart.

X-Scale

Linear

Use a linear scale for the X-axis.

Log

Use a logarithmic scale for the X-axis.

Y-Scale

Linear

Use a linear scale for the Y-axis.

Log

Use a logarithmic scale for the Y-axis.

X-Range (min)

Specify the minimum value to be plotted along the X-axis.

X-Range (max)

Specify the maximum value to be plotted along the X-axis.

Y-Range (min)

Specify the minimum value to be plotted along the Y-axis.

Y-Range (max)

Specify the maximum value to be plotted along the Y-axis.

Bins

You can select a particular binning method or choose a specific bin size.

The following are the available binning methods:

Default

Sturges

Compute the number of bins based on the length of data

Scott

Asymptotically minimize the mean square error

Freedman-Diaconis

Twice the interquartile range divided by the cube root of sample size

Knuth

Balance likelihood and prior probability of a piecewise uniform model

Wand

One-level recursive approximate Wand binning

The following are the available bin sizes:

10

20

50

100

Heights

Count

The number of values lying in each bin

Cumulative Count

Cumulative counts

Survival Count

Survival counts

Probability

Fraction of values lying in each bin

CDF

Cumulative distribution function

SF

Survival function

HF

Hazard function

CHF

Cumulative hazard function

Line

Fields and options available for the **Line** chart type include:

Layout**Default**

No fill under the plotted line.

Filled

Fill the area below the plotted line with color.

Samples

Specify the number of samples to use when plotting a chart.

X-Scale**Linear**

Use a linear scale for the X-axis.

Log

Use a logarithmic scale for the X-axis.

Y-Scale**Linear**

Use a linear scale for the Y-axis.

Log

Use a logarithmic scale for the Y-axis.

X-Range (min)

Specify the minimum value to be plotted along the X-axis.

X-Range (max)

Specify the maximum value to be plotted along the X-axis.

Y-Range (min)

Specify the minimum value to be plotted along the Y-axis.

Y-Range (max)

Specify the maximum value to be plotted along the Y-axis.

Scatter

Fields and options available for the **Scatter** chart type include:

Layout**Default**

Scatter points appear as separate entities on the chart.

Filled

A line extends from each scatter point to the X-axis.

Samples

Specify the number of samples to use when plotting a chart.

X-Scale**Linear**

Use a linear scale for the X-axis.

Log

Use a logarithmic scale for the X-axis.

Y-Scale**Linear**

Use a linear scale for the Y-axis.

Log

Use a logarithmic scale for the Y-axis.

X-Range (min)

Specify the minimum value to be plotted along the X-axis.

X-Range (max)

Specify the maximum value to be plotted along the X-axis.

Y-Range (min)

Specify the minimum value to be plotted along the Y-axis.

Y-Range (max)

Specify the maximum value to be plotted along the Y-axis.

Bar

Fields and options available for the **Bar** chart type include:

Bar orientation**Default**

Bars are displayed vertically in the chart.

Horizontal

Bars are displayed horizontally in the chart.

Layout**Default**

Each grouping in a particular data set is displayed as a separate bar in the chart.

Stacked

All of the related groupings of data in a particular data set are stacked on top of one another contiguously in a single bar.

Stepped

All of the related groupings of data in a particular data set are displayed next to one another in an ascending series of separate bars.

Percentile

All of the related groupings of data in a particular data set are stacked on top of one another contiguously in a single bar; the height of each grouping in the bar corresponds to its value in relation to the other groupings in the same data set.

Samples

Specify the number of samples to use when plotting a chart.

Y-Scale

Linear

Use a linear scale for the Y-axis.

Log

Use a logarithmic scale for the Y-axis.

Y-Range (min)

Specify the minimum value to be plotted along the Y-axis.

Y-Range (max)

Specify the maximum value to be plotted along the Y-axis.

Pie

Fields and options available for the **Pie** chart type include:

Layout

Default

Pie slices extend into the center of the chart.

Donut

The center of the chart is empty, providing the data in the shape of a ring.

Samples

Specify the number of samples to use when plotting a chart.

Bubble

Fields and options available for the **Bubble** chart type include:

Samples

Specify the number of samples to use when plotting a chart.

X-Scale

Linear

Use a linear scale for the X-axis.

Log

Use a logarithmic scale for the X-axis.

Y-Scale

Linear

Use a linear scale for the Y-axis.

Log

Use a logarithmic scale for the Y-axis.

X-Range (min)

Specify the minimum value to be plotted along the X-axis.

X-Range (max)

Specify the maximum value to be plotted along the X-axis.

Y-Range (min)

Specify the minimum value to be plotted along the Y-axis.

Y-Range (max)

Specify the maximum value to be plotted along the Y-axis.

Scatter (3D)

Fields and options available for the **Scatter (3D)** chart type include:

Layout

Default

Filled

Joined

Samples

Specify the number of samples to use when plotting a chart.

X-Scale

Linear

Use a linear scale for the X-axis.

Log

Use a logarithmic scale for the X-axis.

Y-Scale

Linear

Use a linear scale for the Y-axis.

Log

Use a logarithmic scale for the Y-axis.

Z-Scale

Linear

Use a linear scale for the Z-axis.

Log

Use a logarithmic scale for the Z-axis.

X-Range (min)

Specify the minimum value to be plotted along the X-axis.

X-Range (max)

Specify the maximum value to be plotted along the X-axis.

Y-Range (min)

Specify the minimum value to be plotted along the Y-axis.

Y-Range (max)

Specify the maximum value to be plotted along the Y-axis.

Z-Range (min)

Specify the minimum value to be plotted along the Z-axis.

Z-Range (max)

Specify the maximum value to be plotted along the Z-axis.

Surface

Fields and options available for the **Surface** chart type include:

Samples

Specify the number of samples to use when plotting a chart.

X-Scale**Linear**

Use a linear scale for the X-axis.

Log

Use a logarithmic scale for the X-axis.

Y-Scale**Linear**

Use a linear scale for the Y-axis.

Log

Use a logarithmic scale for the Y-axis.

Z-Scale**Linear**

Use a linear scale for the Z-axis.

Log

Use a logarithmic scale for the Z-axis.

X-Range (min)

Specify the minimum value to be plotted along the X-axis.

X-Range (max)

Specify the maximum value to be plotted along the X-axis.

Y-Range (min)

Specify the minimum value to be plotted along the Y-axis.

Y-Range (max)

Specify the maximum value to be plotted along the Y-axis.

Z-Range (min)

Specify the minimum value to be plotted along the Z-axis.

Z-Range (max)

Specify the maximum value to be plotted along the Z-axis.

Candlestick

Fields and options available for the **Candlestick** chart type include:

Samples

Specify the number of samples to use when plotting a chart.

Kagi

Fields and options available for the **Kagi** chart type include:

Samples

Specify the number of samples to use when plotting a chart.

Renko

Fields and options available for the **Renko** chart type include:

Samples

Specify the number of samples to use when plotting a chart.

Point & Figure

Fields and options available for the **Point & Figure** chart type include:

Samples

Specify the number of samples to use when plotting a chart.

Line Break

Fields and options available for the **Line Break** chart type include:

Samples

Specify the number of samples to use when plotting a chart.

Customization settings

The Chart Builder provides a number of customizable settings that control the way a chart is displayed, including the size of the chart, background color, appearance of the title and the axes, and placement of the legend.

The customization settings panel, outlined below on the right, is accessed by clicking **More Options** at the top of the chart parameters panel.

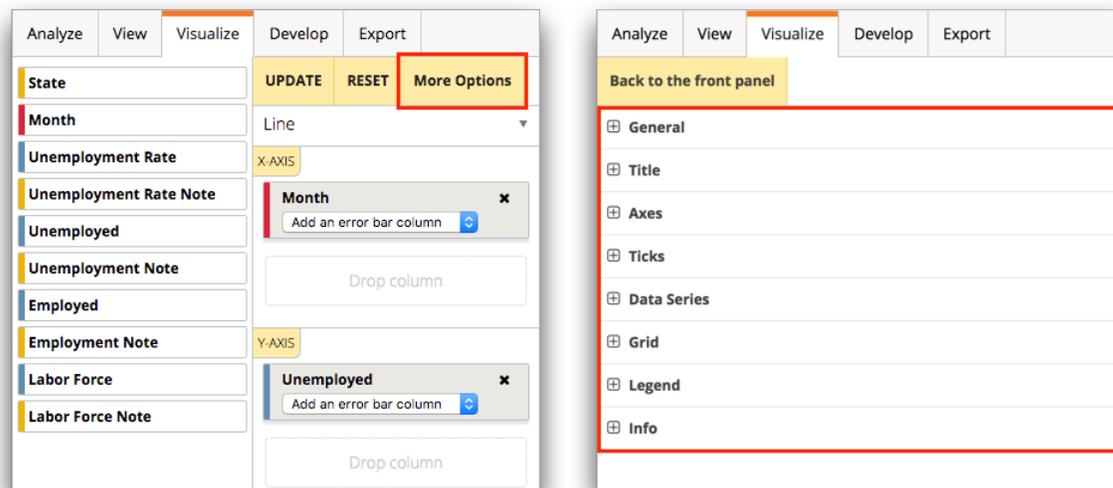


Figure 60: Customization settings panel

Any changes with respect to these settings are reflected immediately in the chart.

To view the data columns and chart parameters panels, click **Back to the front panel**.

General

These settings determine the general appearance of the chart and allow you to select a predefined theme as a basis for the color palette.

Chart width

The width of the chart in pixels.

Chart height

The height of the chart in pixels.

Background color

The color to be used for the background of the chart.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Theme

Predefined color palettes that determine the color of the graph elements.

Available themes include:

- **Default**
- **Happy**
- **Excel**
- **Maximum**

Title

These settings apply to the appearance of the title for the chart.

Show title?

If **Yes**, the title will be displayed above the chart; otherwise, the chart will appear without a title.

If no **Title text** is specified, but **Show title?** is set to **Yes**, the title of the base table will be used.

Title text

The text to appear above the chart.

Title size

The font size of the title text.

Background color

The color to be used for the background of the title.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Color

The color to be used for the title font.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Font family

The font family specification for the title.

Font weight

The weight or thickness of the title.

Font slant

The type of emphasis to be applied to the title.

Axes

These settings determine the appearance of the X, Y, and Z axes and their associated labels and markers.

X-AXIS

The following fields and options are available in the **X-AXIS** tab.

Show X-axis

If **Yes**, the X-axis will be displayed for the chart; otherwise, the chart will appear without an X-axis.

X-axis color

The color of the X-axis and X-axis markers.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

X-axis size

The font size of the X-axis markers.

Font family

The font family specification for the X-axis markers.

Font weight

The weight or thickness of the X-axis markers.

Font slant

The type of emphasis to be applied to the X-axis markers.

Show X-label

If **Yes**, the X-axis label will be displayed for the chart; otherwise, the chart will appear without an X-axis label.

If no **X-label** is specified, but **Show X-label?** is set to **Yes**, the column heading associated with the X-axis will be used by default. If multiple columns are specified for the X-axis, no default label will be displayed.

X-label

The text to appear under the X-axis.

X-label size

The font size of the X-axis label.

X-label color

The color to be used for the X-axis label.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

X-label background

The color to be used for the background of the X-axis label.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Font family

The font family specification for the X-axis label.

Font weight

The weight or thickness of the X-axis label.

Font slant

The type of emphasis to be applied to the X-axis label.

Y-AXIS

The following fields and options are available in the **Y-AXIS** tab.

Show Y-axis

If **Yes**, the Y-axis will be displayed for the chart; otherwise, the chart will appear without a Y-axis.

Y-axis color

The color of the Y-axis and Y-axis markers.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Y-axis size

The font size of the Y-axis markers.

Font family

The font family specification for the Y-axis markers.

Font weight

The weight or thickness of the Y-axis markers.

Font slant

The type of emphasis to be applied to the Y-axis markers.

Show Y-label

If **Yes**, the Y-axis label will be displayed for the chart; otherwise, the chart will appear without a Y-axis label.

If no **Y-label** is specified, but **Show Y-label?** is set to **Yes**, the column heading associated with the Y-axis will be used by default. If multiple columns are specified for the Y-axis, no default label will be displayed.

Y-label

The text to appear next to the Y-axis.

Y-label size

The font size of the Y-axis label.

Y-label color

The color to be used for the Y-axis label.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Y-label background

The color to be used for the background of the Y-axis label.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Font family

The font family specification for the Y-axis label.

Font weight

The weight or thickness of the Y-axis label.

Font slant

The type of emphasis to be applied to the Y-axis label.

Y-AXIS2

The following fields and options are available in the **Y-AXIS2** tab.

Show Y-axis

If **Yes**, the Y-axis will be displayed for the chart; otherwise, the chart will appear without a Y-axis.

Y-axis color

The color of the Y-axis and Y-axis markers.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Y-axis size

The font size of the Y-axis markers.

Font family

The font family specification for the Y-axis markers.

Font weight

The weight or thickness of the Y-axis markers.

Font slant

The type of emphasis to be applied to the Y-axis markers.

Show Y-label

If **Yes**, the Y-axis label will be displayed for the chart; otherwise, the chart will appear without a Y-axis label.

If no **Y-label** is specified, but **Show Y-label?** is set to **Yes**, the column heading associated with the Y-axis will be used by default. If multiple columns are specified for the Y-axis, no default label will be displayed.

Y-label

The text to appear next to the Y-axis.

Y-label size

The font size of the Y-axis label.

Y-label color

The color to be used for the Y-axis label.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Y-label background

The color to be used for the background of the Y-axis label.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Font family

The font family specification for the Y-axis label.

Font weight

The weight or thickness of the Y-axis label.

Font slant

The type of emphasis to be applied to the Y-axis label.

Z-AXIS

The following fields and options are available in the **Z-AXIS** tab.

Show Z-axis

If **Yes**, the Z-axis will be displayed for the chart; otherwise, the chart will appear without a Z-axis.

Z-axis color

The color of the Z-axis and Z-axis markers.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Z-axis size

The font size of the Z-axis markers.

Font family

The font family specification for the Z-axis markers.

Font weight

The weight or thickness of the Z-axis markers.

Font slant

The type of emphasis to be applied to the Z-axis markers.

Show Z-label

If **Yes**, the Z-axis label will be displayed for the chart; otherwise, the chart will appear without a Z-axis label.

If no **Z-label** is specified, but **Show Z-label?** is set to **Yes**, the column heading associated with the Z-axis will be used by default. If multiple columns are specified for the Z-axis, no default label will be displayed.

Z-label

The text to appear next to the Z-axis.

Z-label size

The font size of the Z-axis label.

Z-label color

The color to be used for the Z-axis label.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Z-label background

The color to be used for the background of the Z-axis label.

The color can be selected using the color picker or can be specified as any valid HTML color name, an RGB value, or a hex value.

Font family

The font family specification for the Z-axis label.

Font weight

The weight or thickness of the Z-axis label.

Font slant

The type of emphasis to be applied to the Z-axis label.

Ticks

These settings allow the ticks along the axes to be rotated.

X-tick rotation

Rotate the ticks along the X-axis by the selected number of degrees.

Y-tick rotation

Rotate the ticks along the Y-axis by the selected number of degrees.

Z-tick rotation

Rotate the ticks along the Z-axis by the selected number of degrees.

Data Series

These settings control the appearance of different arrays of values within the current chart.

Colors

Thicknesses

Opacities

Line strokes

Show line markers?

Line marker types

Line marker colors

Line marker sizes

Show scatter markers?

Scatter marker types

Scatter marker colors

Scatter marker sizes

Scatter point sizes

Grid

This setting determines whether the chart is displayed with a grid or not.

Show grid?

This determines whether a grid should be displayed with the corresponding chart.

Default

The default setting.

Yes

The grid is displayed.

No

The grid is not displayed.

Legend

These settings control whether a legend should be displayed with the corresponding chart, and if so, where it should appear.

Show legend?

This determines whether a legend should be displayed with the corresponding chart.

Default

The default setting.

Yes

The legend is displayed according to the specified **Position**.

No

The legend is not displayed.

Position

If the legend is to be displayed, this determines its position relative to the chart.

Bottom

The legend appears below the chart.

Top

The legend appears above the chart.

Left

The legend appears to the left of the chart.

Right

The legend appears to the right of the chart.

Info

This sections displays the running time of the chart and information about sampling.

Running time

The running time of the chart.

Is chart sampled?

Displays whether the chart was sampled or not.

Samples

If the chart was sampled, this displays how many samples.

Sampling type

If the chart was sampled, this displays how the chart was sampled.

Create a chart

Create a chart for the current table or worksheet by selecting a chart type and then specifying the data columns to be plotted.

It can sometimes be easier to find trends and see differences more easily in a chart. For example, it may be quicker and easier to find the day of sales with the highest profit margin from a chart than from a table of numbers listing the same information. Charting can help you make sense of large amounts of data that might otherwise be too overwhelming to comprehend.

Note: A chart is not retained when you save your analysis as a Quick Query or export it as a QuickApp. To save the chart with your query, save your analysis as a workspace. For instructions, see [Save your workspace](#) on page 34.

To create a chart:

1. In an open table or worksheet, click the **Visualize** tab.
The 1010data Insights Platform displays the Chart Builder.
2. In the drop-down list below the **Update**, **Reset**, and **More Options** buttons, select a chart type.
For information about the available chart types, see [Chart types](#) on page 250.

The Insights Platform displays settings related to the selected chart type in the chart data, **Interactions**, and **Settings** sections. Within the chart data sections, **Drop column** areas are provided.

3. For each column that you want to plot, drag it from the data columns panel into the **Drop column** area in the chart data section associated with the axis where you want that data plotted.

Note: You cannot drag columns from one chart data section to another. If you want to remove a column from a chart data section, click the **Delete** (✖) icon.

4. To enable certain mouse actions in the chart, select the **Enable Interactions** option.

Chart interactions allow you to zoom, recenter, and select data points individually or as a range in the current chart. For more information, see [Interactions](#) on page 258.

5. Make any desired changes in the **Settings** section.

For instance, if you selected a **Pie** chart type, you can display it as a donut chart by changing the **Layout** setting to **Donut**.

For a list of available fields and options, see [Settings](#) on page 259.

6. Make any desired changes in the customization settings panel.

The customization settings panel is accessed by clicking **More Options** at the top of the chart parameters panel and allows you to customize your chart. For example, you can enter a title, change the tick rotation for a particular axis, select a different color for the background, or change the position of the legend.

For a list of available fields and options, see [Customization settings](#) on page 269.

7. Click **Update**.

Note: If the customization settings panel is displayed, you must first click **Back to the front panel**.

The chart is displayed in the results pane of the **TRS** window.

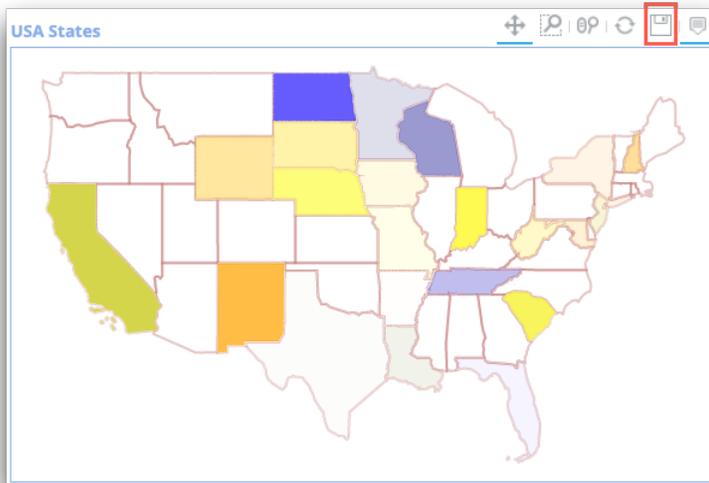
Note: If you make further changes in the **Settings** section or add/remove columns from any of the chart data sections, you must click **Update** to see those changes reflected in the chart. Modifications made in the customization settings panel take effect immediately.

Download a chart

You can download the chart you are currently viewing to your computer as a PNG file.

To download a chart:

In the chart viewer, click the save () icon.



The chart is saved as `chart.png` in your designated downloads folder.

Note: Depending on the browser you are using, you may be prompted to open or save the file.

Develop

Use the query that represents your analysis as a base for further development.

The **Develop** tab allows you to send the query to the Macro Language Workshop (MLW) for editing and further development, save the operations from your analysis as a Quick Query, or consolidate your timeline into as few blocks as possible.

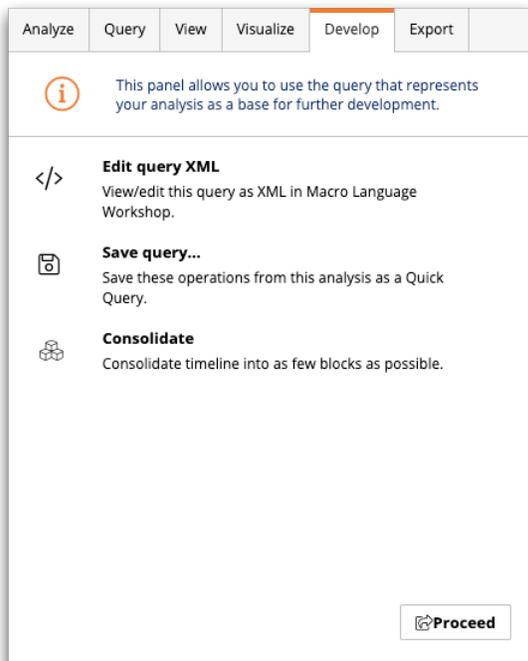


Figure 61: Develop tab in the Trillion-Row Spreadsheet

Edit query XML

View and edit your query as XML in the Macro Language Workshop.

When this option is selected, the query from your analysis in the Trillion-Row Spreadsheet is sent to the Macro Language Workshop and is shown as XML code. From there, you can develop your query further.

For instructions, see [Edit a query in the Macro Language Workshop](#) on page 280.

Save query

Save the operations from your analysis as a Quick Query.

A saved Quick Query allows you to return to the analysis at a later time or share the analysis with other 1010data Insight Platform users.

For instructions, see [Save an analysis as a Quick Query](#) on page 281.

Consolidate

Consolidate the timeline into as few blocks as possible.

For instructions, see [Consolidate timeline](#) on page 283.

Edit a query in the Macro Language Workshop

View and edit a Trillion-Row Spreadsheet query as XML code in the Macro Language Workshop.

You can send an analysis performed in the Trillion-Row Spreadsheet to the Macro Language Workshop to view and edit the query as Macro Language XML code.

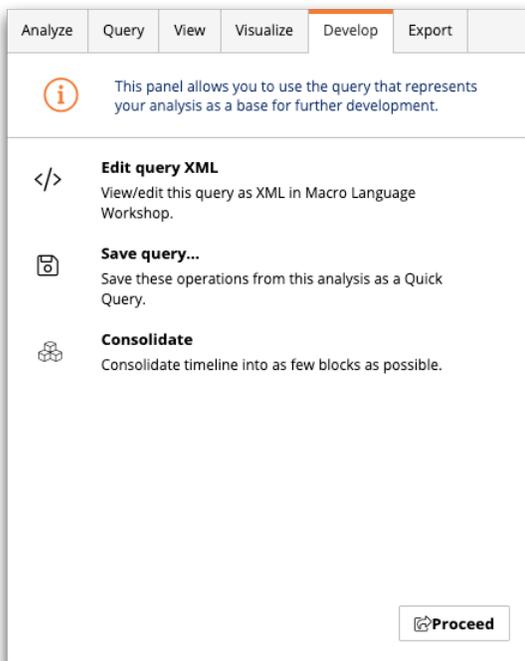
When you do so, the 1010data Insights Platform creates a dynamic link between the Trillion-Row Spreadsheet (TRS) and the Macro Language Workshop (MLW). As the query changes in the TRS, the Macro Language XML code is automatically updated in the MLW.

Note: This dynamic link is automatic in only one direction. If changes are made to the XML code in the MLW, the linked query in the TRS is not automatically updated. To update the linked query in the TRS with the code changes in the MLW, you must manually re-export those changes.

To edit a query in the Macro Language Workshop:

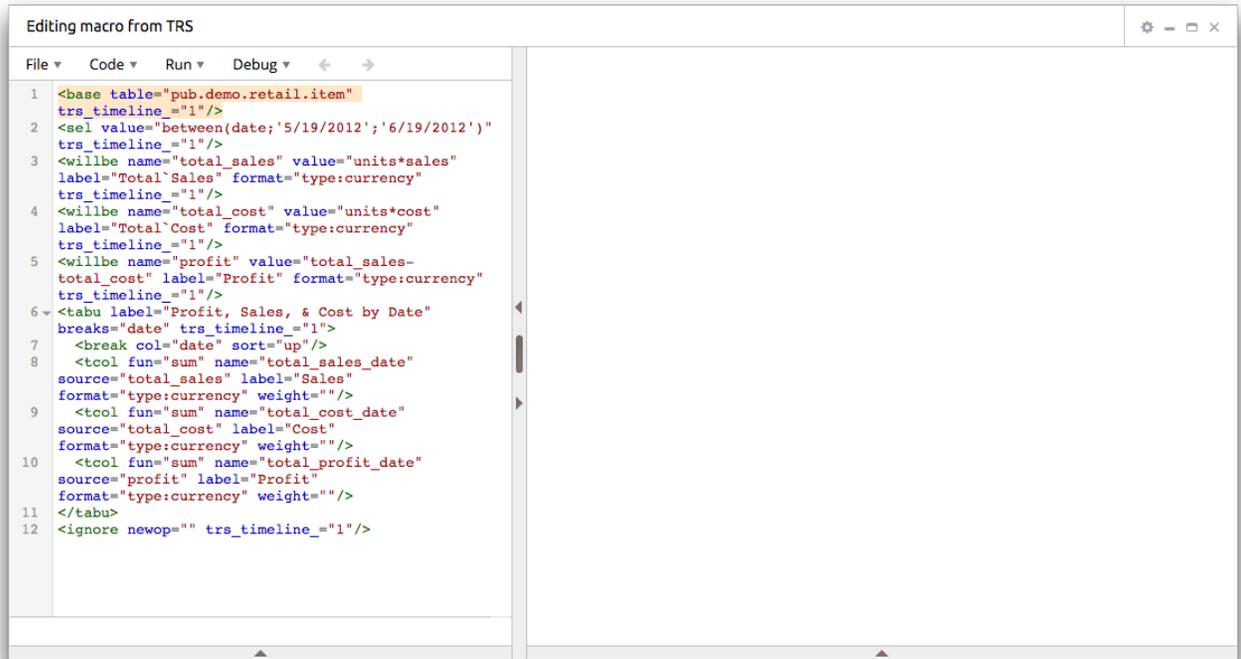
1. In the Trillion-Row Spreadsheet, click the **Develop** tab.

The Trillion-Row Spreadsheet displays the options available for further developing your TRS analysis.



2. Click **Edit query XML** and then click **Proceed**.

The Insights Platform sends your analysis to the Macro Language Workshop and displays the query as XML code in the **Editing macro from TRS** window.



```

1 <base table="pub.demo.retail.item"
  trs_timeline_"1"/>
2 <sel value="between(date;'5/19/2012';'6/19/2012')"
  trs_timeline_"1"/>
3 <willbe name="total_sales" value="units*sales"
  label="Total`Sales" format="type:currency"
  trs_timeline_"1"/>
4 <willbe name="total_cost" value="units*cost"
  label="Total`Cost" format="type:currency"
  trs_timeline_"1"/>
5 <willbe name="profit" value="total_sales-
  total_cost" label="Profit" format="type:currency"
  trs_timeline_"1"/>
6 <tabu label="Profit, Sales, & Cost by Date"
  breaks="date" trs_timeline_"1">
7   <break col="date" sort="up"/>
8   <col fun="sum" name="total_sales_date"
  source="total_sales" label="Sales"
  format="type:currency" weight=""/>
9   <col fun="sum" name="total_cost_date"
  source="total_cost" label="Cost"
  format="type:currency" weight=""/>
10  <col fun="sum" name="total_profit_date"
  source="profit" label="Profit"
  format="type:currency" weight=""/>
11 </tabu>
12 <ignore newop="" trs_timeline_"1"/>

```

3. View/edit your Macro Language XML code in this window.

4. If desired, you can click **Run > Re-export** to send the Macro Language code back to the TRS where it originated.

Note:

When a Macro Language Workshop window is created as a child of a Trillion-Row Spreadsheet instance, *Re-export* updates the Trillion-Row Spreadsheet with a modified query from the Macro Language Workshop. The option has been moved from the "File" menu to the "Run" menu, and choosing it will now also update the "render" shortcut button. Likewise, once this option has been explicitly chosen once, using the `control+enter/return` keyboard shortcut will re-re-export. These changes make it easier to achieve workflows such as using a child Macro Language Workshop instance as a floating "query editor" in concert with a fullscreened Trillion-Row Spreadsheet.

For more information on running queries, see [Run](#) on page 417 in the Macro Language Workshop section.

Save an analysis as a Quick Query

Save the operations of an analysis in the Trillion-Row Spreadsheet as a Quick Query.

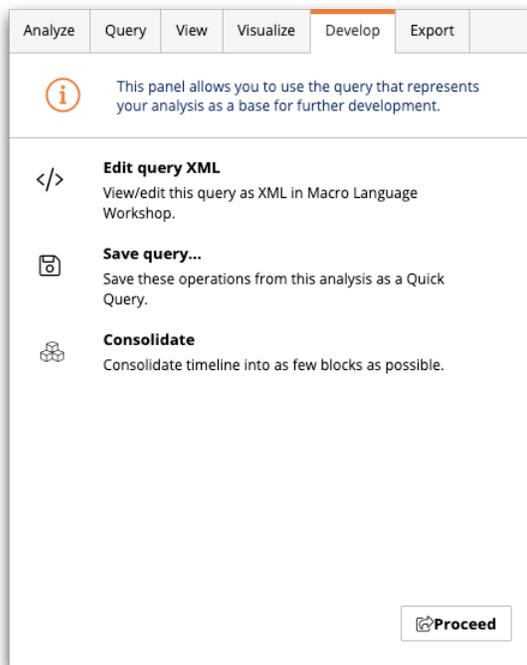
After an analysis in the Trillion-Row Spreadsheet (TRS) is saved as a Quick Query, you can return to the analysis at a later time or share the Quick Query with other 1010data Insight Platform users.

You can save the analysis as a new Quick Query or replace an existing Quick Query that you own.

To save your analysis as a Quick Query:

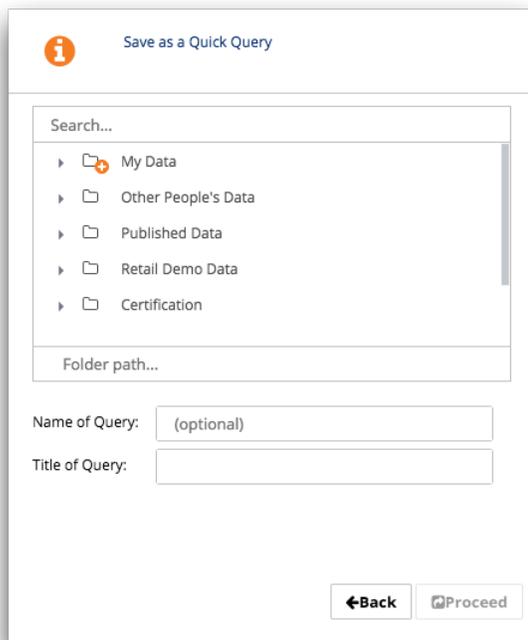
1. In the Trillion-Row Spreadsheet, click the **Develop** tab.

The Trillion-Row Spreadsheet displays the options available for further developing your TRS analysis.



2. Click **Save query**.

The Trillion-Row Spreadsheet displays the object browser and save option fields.



3. Do one of the following:

Option

Save a new Quick Query

Description

Select the folder in which you want to save the new Quick Query.

Option	Description
	<p>Note: You can only save the Quick Query to a folder you own (🔑) or have permission to add to (+).</p>
<p>Replace an existing Quick Query</p>	<p>Select an exiting Quick Query that you want to replace.</p> <p>Note: You can only replace a Quick Query that you own (🔑).</p>

4. Complete the following fields:

Name of Query

Optionally, enter a name for the Quick Query.

The name is the file name of the saved Quick Query. It must begin with a letter and can only contain numbers, letters, and underscores. The name cannot contain any spaces or other special characters. If you leave this field blank, a system-generated name will be used (e.g., t662528755_yourusername).

Note: If you save the Quick Query in the **My Data** folder, you cannot enter anything in the Name field. Instead, the Insights Platform uses a system-generated name.

Title of Query

Enter a title for the Quick Query.

The title is used to help describe the behavior of the Quick Query (e.g., *Transactions by Store*) and may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters.

5. Click **Proceed**.

The Trillion-Row Spreadsheet saves the Quick Query and displays a confirmation message.

Consolidate timeline

Consolidate the TRS timeline into as few blocks as possible.

Consolidating a timeline has the effect of grouping the steps in the timeline. For example, a timeline with no parameterized steps may be grouped into a single block. Parameterized steps will be consolidated, if possible, so that their prompts appear in a single panel. However, if the prompt of a step depends on the prompt of a previous step, the steps will not be consolidated. The creator of the parameterization can also specify explicitly in the **Block setup** panel not to consolidate a particular step (the option **Force new step when consolidated**). See [Block setup panel](#) on page 385 for more information.

Once the timeline is consolidated, it is not possible to edit the parameterization of the original panels, except by undoing the consolidation.

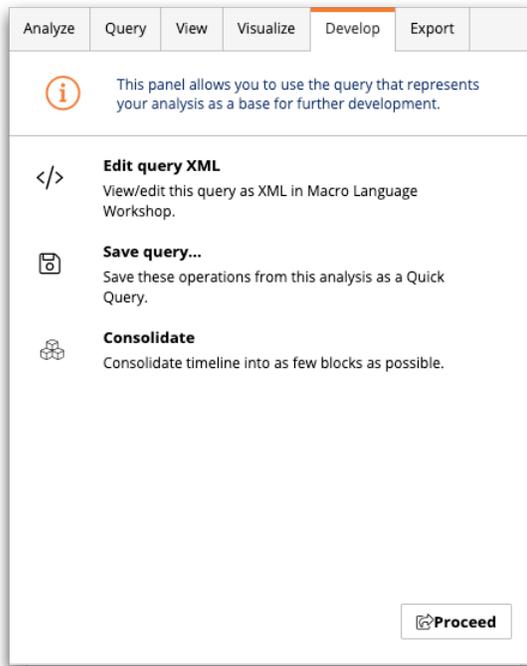
For more information about parameterized steps, see [Add parameters to an operation](#) on page 387.

Typically, consolidation would be the last step before saving the timeline as a Quick Query. See [Save an analysis as a Quick Query](#) on page 281.

To consolidate the timeline:

1. In the Trillion-Row Spreadsheet, click the **Develop** tab.

The Trillion-Row Spreadsheet displays the options available for further developing your TRS analysis.



2. Click **Consolidate**.
3. Click **Proceed**.
The TRS Analysis Timeline groups steps into blocks, where possible.
4. If desired, save the consolidated query as a Quick Query.

Export

Export the data resulting from your analysis to various destinations.

You can choose to download the data as delimited text, Microsoft Excel workbook, or PDF file. In addition, you can export the data to your 1010data FTP directory or save the data to a 1010data Insights Platform table.

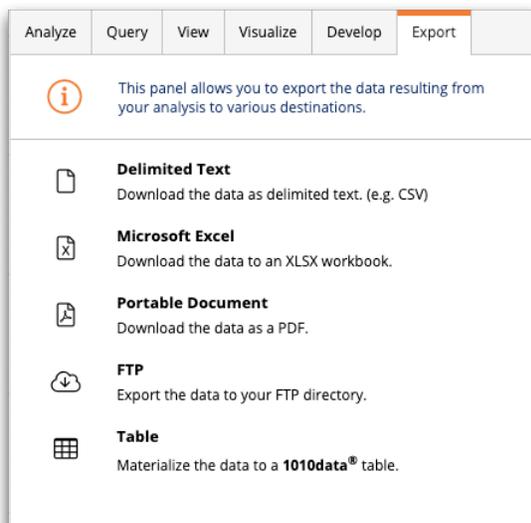


Figure 62: Export tab in the Trillion-Row Spreadsheet

Delimited Text

Download the data as a delimited text file, such as a comma-separated values (CSV) file. Selecting this option displays the delimited text export settings.

For more information, see [Delimited text export settings](#) on page 285.

Microsoft Excel

Download the data to an Excel workbook (XLSX) file. Selecting this option displays the Excel workbook export settings.

For more information, see [Excel workbook export settings](#) on page 287.

Portable Document

Download the data to a portable document (PDF) file. Selecting this option displays the PDF export settings.

For more information, see [PDF export settings](#) on page 288.

FTP

Export the data to your 1010data FTP directory. Selecting this option opens the **Save Data to FTP** window.

For more information, see [Save Data to FTP](#) on page 370.

Table

Save the data resulting from your analysis to an Insights Platform table. Selecting this option opens the **Save Data as Table** window.

For more information, see [Save Data as Table](#) on page 377.

Delimited text export settings

Fields and options for downloading the data as a delimited text file.

The delimited text export settings allow you to customize an exported file to suit your needs. Depending on the defined fields and options, the file extension of the exported file will generally be either `.csv`, `.txt`, or `.zip`.

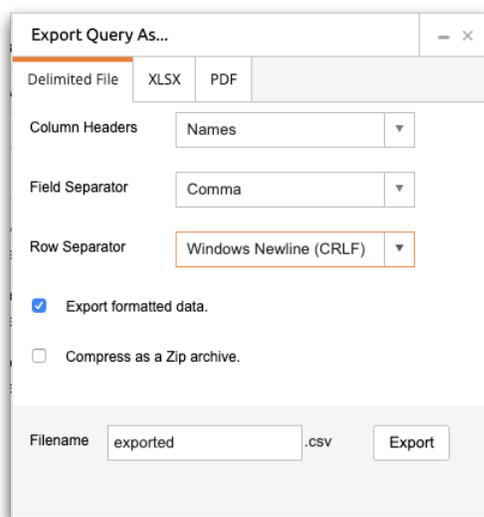


Figure 63: Export settings for delimited files

Column Headers

Names

The column header in the text file includes only the column name for each column.

Labels

The column header in the text file includes only the column label for each column.

Note: If this option is selected and a column does not have a label, the column name is used in its place.

Names and Labels

The column header in the text file includes both column name and column label for each column, in that order.

Labels and Names

The column header in the text file includes both column label and column name for each column, in that order.

No Headers

The column header in the text file does not include column headers.

Field Separator

The options in this section define the character used to separate adjacent columns in the text file.

Comma

Columns in each row of the table are separated by a comma (,) in the text file.

Tab

Columns in each row of the table are separated by a tab in the text file.

Space

Columns in each row of the table are separated by a space in the text file.

Colon

Columns in each row of the table are separated by a colon (:) in the text file.

Semicolon

Columns in each row of the table are separated by a semicolon (;) in the text file.

Vertical Bar

Columns in each row of the table are separated by a vertical bar (|) in the text file.

Row Separator

The options in this section define the record (row) separator character. The separator is used to indicate the end of each row of the table in the text file.

The row separator is a sequence of one or more characters used to specify the boundary between separate, independent regions in plain text files. In general, the two most commonly used record delimiters in the Insights Platform are CRLF (Windows newline) and LF (Unix/OS X Newline).

Windows Newline (CRLF)

The last column in each row of the table is indicated by the Carriage Return Line Feed (CRLF) row separator in the text file.

The CRLF ('`\r\n`', 0x0D 0x0A) row separator is used in Microsoft Windows, DOS (MS-DOS, PC DOS, etc.), DEC TOPS-10, RT-11, CP/M, MP/M, Atari TOS, OS/2, Symbian OS, Palm OS, Amstrad CPC, and most other early non-Unix and non-IBM OSes. Select this option for PC files.

Unix/OS X Newline (LF)

The last column in each row of the table is indicated by the Line Feed (LF) row separator in the text file.

The LF ('`\n`', `0x0A`) row separator is used in Multics, Unix and Unix-like systems (Linux, OS X, FreeBSD, AIX, Xenix, etc.), BeOS, Amiga, RISC OS, and others.

macOS 9 Newline (CR)

The last column in each row of the table is indicated by the Carriage Return (CR) row separator in the text file.

The CR ('`\r`', `0x0D`) row separator is used in macOS 9.

File-type settings

Export formatted data

Compress as a Zip archive

The exported text file is compressed into a `.zip` file. The **Filename** extension changes to `zip` automatically.

Filename

By default, exported delimited `.csv` files are named `exported.csv`. You can choose to name the exported file something else by entering a new name in this field.

Export

Click **Export** to export the query with the selected settings.

Excel workbook export settings

Fields and options for downloading the data to a Microsoft Excel workbook file.

The Excel workbook export settings allow you to customize an exported file to suit your needs. The file extension of the exported file is `.xlsx`.

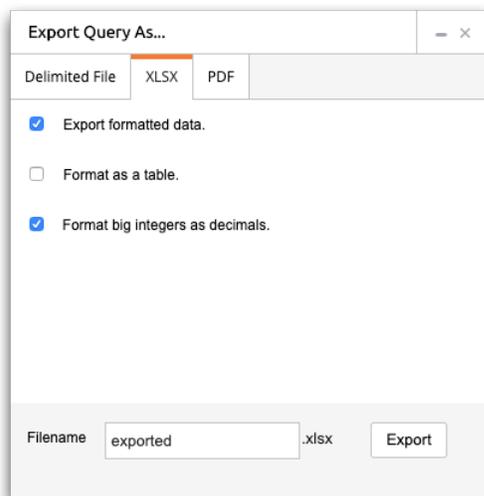


Figure 64: Export settings for Excel workbook files

Export formatted data

Format as a table

If selected, the exported Excel workbook file is saved as an Excel Table object and formatted. If not selected, the exported Excel workbook file is not formatted and appears as a plain, non-styled worksheet.

Format big integers as decimals

Filename

By default, exported `.xlsx` files are named `exported.xlsx`. You can choose to name the exported file something else by entering a new name in this field.

Export

Click **Export** to export the query with the selected settings.

PDF export settings

Fields and options for downloading the data to a portable document (PDF) file.

The PDF export settings allow you to customize an exported PDF file to suit your needs. The file extension of the exported file is `.pdf`.

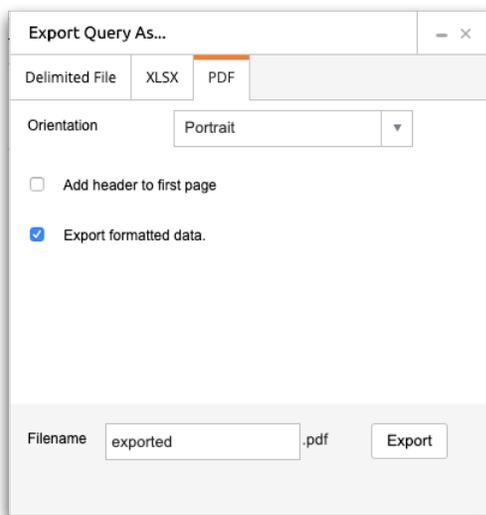


Figure 65: Export settings for portable document (PDF) files

Orientation

Choose between **Portrait** and **Landscape** modes.

Add header to first page

Add column headers to the first page of the PDF file.

Export formatted data

Filename

By default, exported `.pdf` files are named `exported.pdf`. You can choose to name the exported file something else by entering a new name in this field.

Export

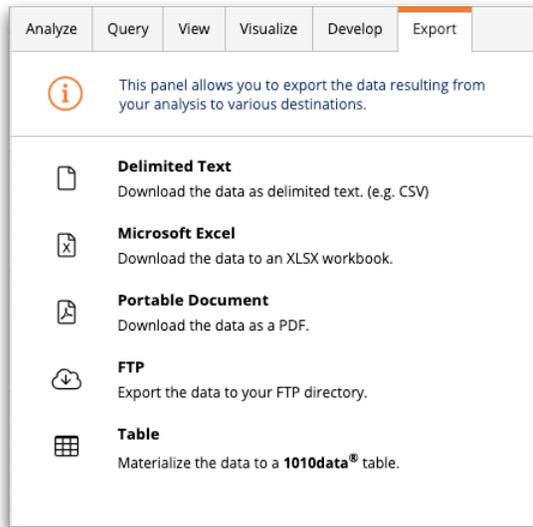
Click **Export** to export the query with the selected settings.

Download the data as a delimited text file

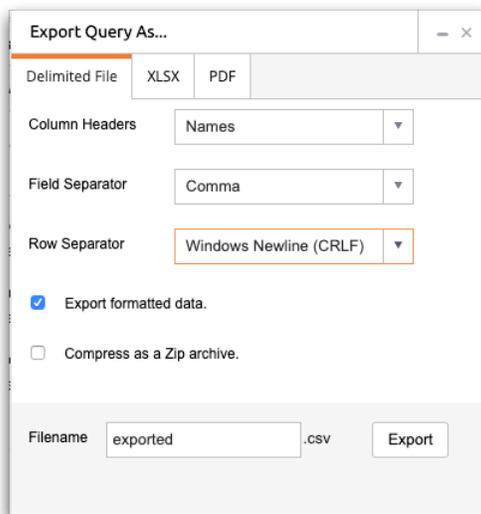
Download the data resulting from a query as a delimited text file.

To download the data as a delimited text file:

1. In the Trillion-Row Spreadsheet, click the **Export** tab.
The Trillion-Row Spreadsheet displays the export options.



2. Click **Delimited Text**.
The Trillion-Row Spreadsheet displays the delimited text export settings.



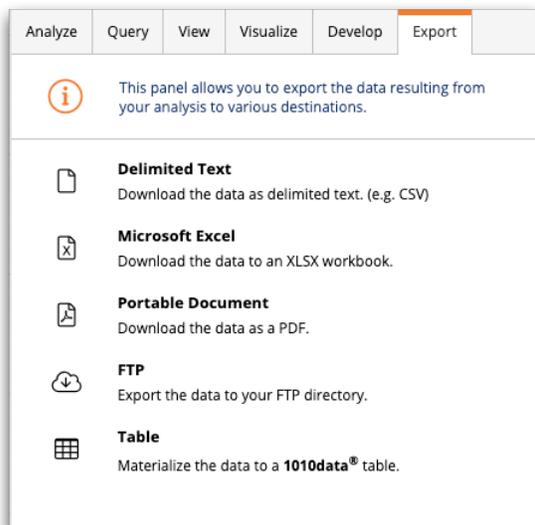
3. Complete the appropriate fields and options.
For a list of fields and options, see [Delimited text export settings](#) on page 285.
4. Click **Export**.
The data is saved to your computer with the file name you entered.

Download the data as an Excel workbook file

Download the data resulting from a query as a Microsoft Excel workbook file.

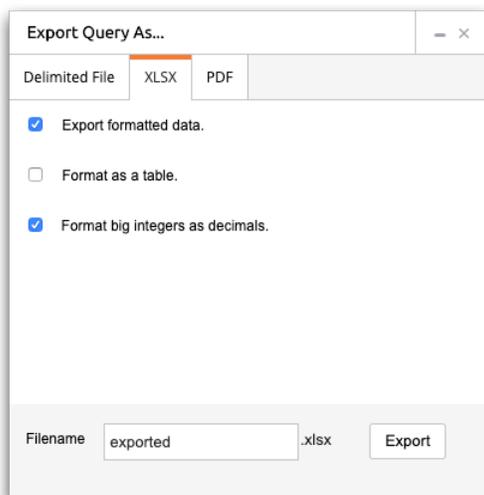
To download the data as an Excel workbook file:

1. In the Trillion-Row Spreadsheet, click the **Export** tab.
The Trillion-Row Spreadsheet displays the export options.



2. Click Microsoft Excel.

The Trillion-Row Spreadsheet displays the Excel workbook export settings.



3. Select the appropriate options.

For a list of options, see [Excel workbook export settings](#) on page 287.

4. Click Export.

The data is saved to your computer as an Excel workbook with the file name you entered.

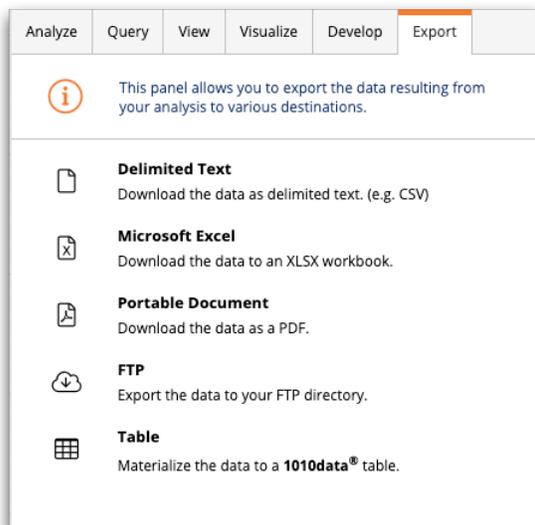
Download the data as a PDF file

Download the data resulting from a query as a portable document (PDF) file.

To download the data as a PDF file:

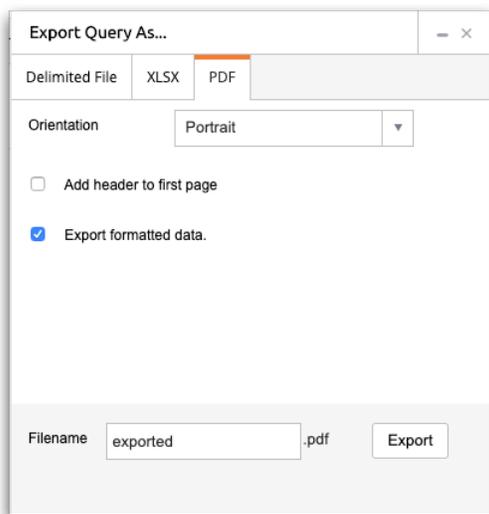
1. In the Trillion-Row Spreadsheet, click the Export tab.

The Trillion-Row Spreadsheet displays the export options.



2. Click **Portable Document.**

The Trillion-Row Spreadsheet displays the PDF export settings.



3. Select the appropriate options.

For a list of options, see [PDF export settings](#) on page 288.

4. Click **Export.**

The data is saved to your computer as a PDF file with the file name you entered.

Results pane

The results pane is displayed on the right side of the Trillion-Row Spreadsheet window and Macro Language Workshop.

This pane displays the results of your analysis. In the Trillion-Row Spreadsheet, this pane also displays charts created in the **Visualize** tab. The results pane within the Trillion-Row Spreadsheet and Macro Language Workshop have similar functionality.

Sales Detail
Cols 1 fixed and 2 to 9 of 16 , Rows 1 to 19 of 4,481,814

Trans ID	Date	Time	Store	SKU	Extended Sales	Qty/Wgt	Promo	Cost
1400 >	01/01/14	07:17:31	18	225501	0.94	1.00	0	0.
1400 >	01/01/14	07:17:31	18	310659	2.34	1.00	0	1
1400 >	01/01/14	07:17:31	18	127509	0.05	1.00	0	0.
1400 >	01/01/14	07:17:31	18	254062	1.44	1.00	0	0
1400 >	01/01/14	07:17:31	18	310659	2.34	1.00	0	1
1400 >	01/01/14	07:17:31	18	127509	0.05	1.00	0	0.
1400 >	01/01/14	07:17:31	18	75452	0.28	1.00	0	0.
1400 >	01/01/14	07:17:31	18	254062	1.44	1.00	0	0
1400 >	01/01/14	07:17:31	18	479314	1.1	1.00	0	0.
1400 >	01/01/14	07:17:31	18	479314	1.1	1.00	0	0.
1400 >	01/01/14	07:17:31	18	479314	1.1	1.00	0	0.
1400 >	01/01/14	07:17:31	18	254062	1.44	1.00	0	0
1400 >	01/01/14	07:17:31	18	306726	7.63	1.00	0	7.
1400 >	01/01/14	07:17:31	18	288594	1.22	1.00	0	1.
1400 >	01/01/14	07:17:31	18	227891	3.3	1.00	0	1.
1400 >	01/01/14	07:17:31	18	62891	3.47	1.00	0	2.
1400 >	01/01/14	07:17:31	18	247593	7.39	2.00	>	5.
1400 >	01/01/14	07:17:31	18	285981	-1 >	1.00	0	-1
1400 >	01/01/14	07:22:28	18	99429	0.78	1.00	0	0.

Figure 66: Results Pane in TRS and MLW

A. Grid view

The grid view displays the results of your analysis as a grid. See [Grid view](#) on page 294 for more information.

B. View bar

The view bar contains icons showing the different options available for displaying your results in the results pane. See [View bar](#) on page 292 for the available options.

C. Grid menu options

You can perform many actions on data in the grid, by right-clicking the appropriate cell in grid view. See [Grid menu options](#) on page 302 for more information.

View bar

Choose how data is displayed in the results pane of the Trillion-Row Spreadsheet and Macro Language Workshop.

By default, the results from your data analysis are displayed in the grid inside the results pane. You can choose to view the data in different formats, some of which also provide additional ways to interact with the data.

The view bar in the Trillion-Row Spreadsheet (TRS) and Macro Language Workshop (MLW) results pane provides quick access to the available views.

You can also select a view from the **View** tab in the analysis pane in the Trillion-Row Spreadsheet. For more information, see [View](#) on page 244.

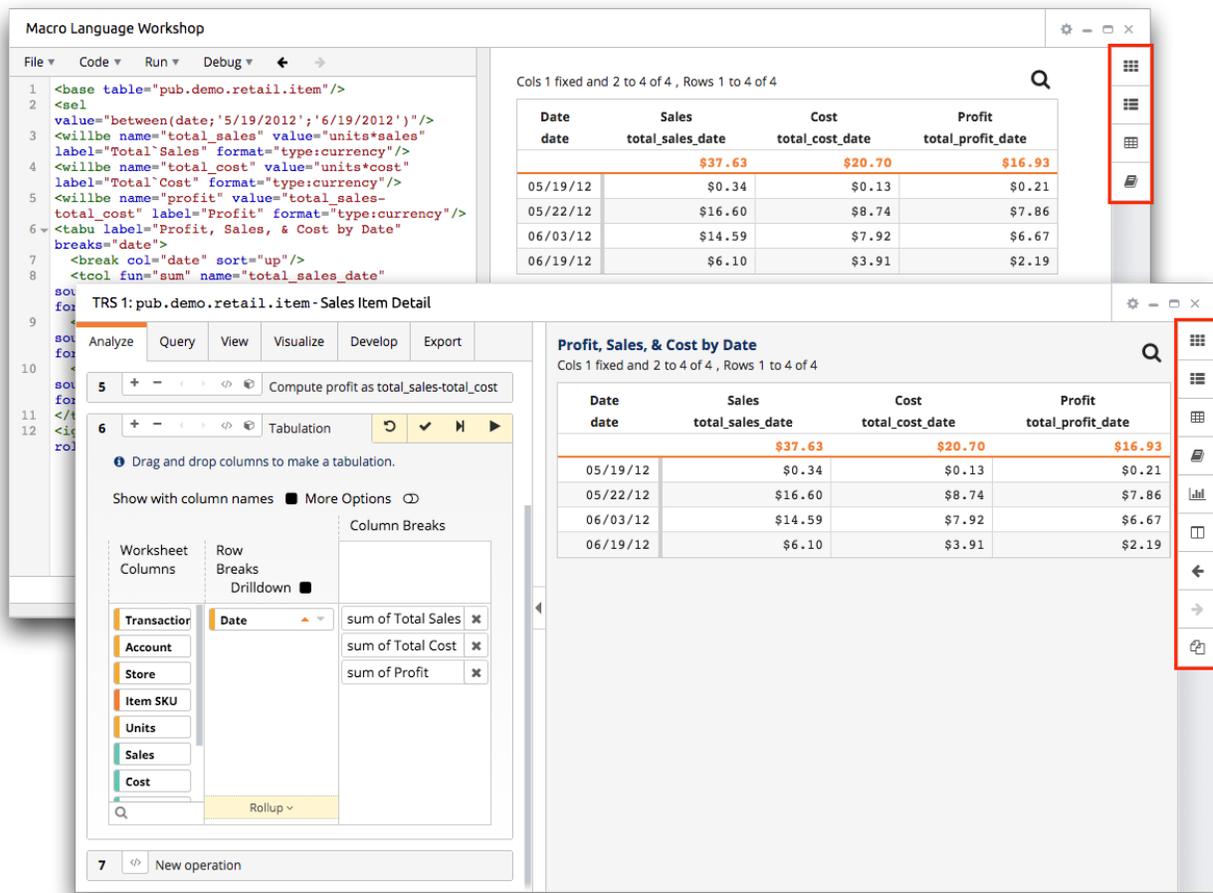


Figure 67: View bar in the results pane of the MLW and TRS windows

Table 23: View bar icons

Icon	Name	Description
	Grid View	Displays data in an infinitely-scrollable grid. This is the default view. For more information, see Grid view on page 294.
	Single-row View	Displays all fields of a single row of the data arranged vertically. This view is particularly useful for viewing information in a table that has many columns. For more information, see Single-row view on page 325.
	Data Dictionary	Displays the metadata of columns and tables in the analysis.
	Visualization	Displays the chart created from the data analysis. This icon appears after a chart is created for the data analysis using the Chart Builder in the Visualize tab of the Trillion-Row Spreadsheet. For more information, see Visualize on page 246. This view is only available in the Trillion-Row Spreadsheet.

Icon	Name	Description
	Grid and visualization	Displays both the grid and the chart created from the data analysis. This icon appears after a chart is created for the data analysis using the Chart Builder in the Visualize tab of the Trillion-Row Spreadsheet. For more information, see Visualize on page 246. This view is only available in the Trillion-Row Spreadsheet.
	Previous view	Displays the previously selected view. This icon appears after the first time a different view is selected. This view is only available in the Trillion-Row Spreadsheet.
	Next view	Displays the next view in the view history. This icon appears after clicking the Previous view icon and allows you to return the next view. This view is only available in the Trillion-Row Spreadsheet.
	Clone view	Clones the current view to a new Trillion-Row Spreadsheet window. This view is only available in the Trillion-Row Spreadsheet. For instructions, see Clone the TRS window on page 326.

Grid view

The Grid view displays data in an infinitely-scrollable, spreadsheet-like interface in the results pane of the Trillion-Row Spreadsheet.

By default, data in the Trillion-Row Spreadsheet (TRS) is displayed in the Grid view. The Grid view is an infinitely-scrollable, spreadsheet-like view of a table or worksheet.

In addition to viewing table or query data, you can interact with the data directly from within the grid. For example, you can use the grid to hide and rearrange columns, sort rows, perform a row selection, and analyze column data.

The Grid view in the results pane is accessed by clicking the **Grid View** () icon in the View bar.

Sales Item Detail
Cols 1 to 7 of 8 , Rows 1 to 18 of 35

Transaction ID transid	Account account	Store store	Date date	Item SKU sku	Units units	Sales sales
531	957	1	05/15/12	366	-1	-
532	478	1	05/15/12	98A	1	0.
532	478	1	05/15/12	3B7	1	1.
534	738	1	05/16/12	A96	2	
534	738	1	05/16/12	65B	1	2.2
535	709	2	05/15/12	CB7	1	1.6
535	709	2	05/15/12	96A	1	1.
535	709	2	05/15/12	969	1	1.
536	748	2	05/17/12	3A4	3	1.0
536	748	2	05/17/12	366	1	
537	523	3	05/15/12	CB7	1	1.6
537	523	3	05/15/12	A96	2	
537	523	3	05/15/12	98A	1	0.
537	523	3	05/15/12	96A	1	1.
537	523	3	05/15/12	3B7	1	1.
538	668	1	05/18/12	CB7	1	1.6
538	668	1	05/18/12	98A	4	

Figure 68: Grid view in the results pane of the TRS

A. Table title

The table title, also referred to as the table label, identifies the table or worksheet. The table *title* is different from the table *name*.

B. Column and row counts

The column and row counts show the total number of column and rows in the table or worksheet. It also indicates which columns and rows are currently displayed.

The number of columns and rows shown is dependent on the size of the **TRS** window. You can resize the window to change the number of visible rows and columns.

If the entire table cannot be displayed in the window, the TRS displays scroll bars on the side or bottom of the grid.

C. Find rows

This icon opens the **Find rows** dialog where you can locate rows in the table or worksheet by entering any valid selection expression. For instructions, see [Find rows using a selection expression](#) on page 313.

D. Column header row

The column header row in the grid displays the column header which contains the column name, column, label, or both. You can choose which is displayed by adjusting your user profile settings. For instructions, see [Change your user settings](#) on page 85.

When you place the pointer over a column header, the TRS displays additional information about the column. For details, see [Column information](#) on page 49. In addition, a drop-down arrow (▼) icon and a blue column header handle appear.

Use the options in the drop-down menu to sort, hide, resize, or modify a column. For a description of each option, see [Column header menu options](#) on page 296. Drag the blue handle to resize a column. For more information, see [Resize a column within the grid](#) on page 300.

Finally, you can drag a column header to rearrange columns in the table or worksheet. For instructions, see [Rearrange columns using the grid](#) on page 160.

E. Grid

The grid displays data from the table or worksheet in an infinitely-scrollable, spreadsheet-like interface.

You can right-click a cell in the grid to access a menu of options available to perform on the data. For example, you can copy the contents of a cell, perform a row selection based on the cell data, and analyze the column data.

F. Fixed columns

Fixed columns, when they exist in a table or worksheet, are always visible and are separated from the other columns by a thicker vertical line. Fixed columns do not move, regardless of how far the table is scrolled in either direction. This is useful when a column has key information that should always be shown.

G. Scroll bars

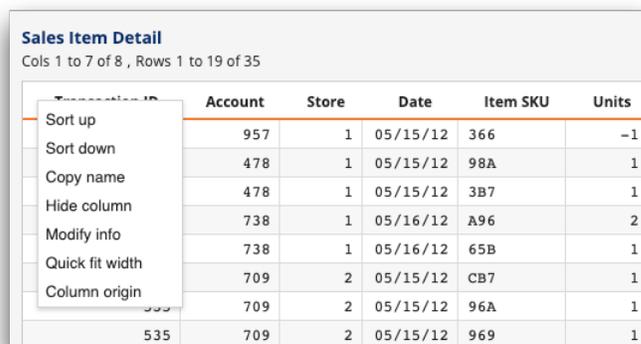
Scroll bars appear on the side or bottom of the grid when the entire table or worksheet cannot be displayed in the **TRS** window. Use the scroll bars to navigate through the rows and columns in the grid.

You can also go directly to a specific row or column by right-clicking the appropriate scroll bar and specifying the location in the dialog. For instructions, see [Scroll directly to a specific row](#) on page 323 and [Scroll directly to a specific column](#) on page 324.

Column header menu options

The column header menu provides quick access to commonly used column settings and commands.

The column header menu is accessed by clicking the drop-down arrow (▾) icon in the column header row within the Grid view of the Trillion-Row Spreadsheet or the Macro Language Workshop.



	Account	Store	Date	Item SKU	Units	
Sort up	957	1	05/15/12	366	-1	
Sort down	478	1	05/15/12	98A	1	
Copy name	478	1	05/15/12	3B7	1	
Hide column	738	1	05/16/12	A96	2	
Modify info	738	1	05/16/12	65B	1	
Quick fit width	709	2	05/15/12	CB7	1	
Column origin	709	2	05/15/12	96A	1	
	535	709	2	05/15/12	969	1

Figure 69: Column header menu in the grid

Sort up

Sorts the rows in the table such that the values in the column are in ascending order.

For instructions, see [Sort rows within the grid](#) on page 158.

Sort down

Sorts the rows in the table such that the values in the column are in descending order.

For instructions, see [Sort rows within the grid](#) on page 158.

Copy name

Copies the selected column name to the clipboard. This is useful when you want to copy and paste column names into your query.

Hide column

Hides the column from the grid.

For instructions, see [Hide a column within the grid](#) on page 162.

Modify info

Opens a dialog which allows you to edit the metadata information of the column.

For more information, see [Column metadata dialog](#) on page 297.

Quick fit width

Adjusts the width of the column to fully display the widest row of data in the column.

Column origin

Opens the timeline panel from which the column originates.

For example, if the column is the result of a computed column, the Trillion-Row Spreadsheet opens the relevant **Computed column** panel. If the column is derived from a link to a foreign table, the appropriate **Link tables** panel is opened. If the column is the result of a tabulation, the applicable **Tabulation** panel is opened. If the column is a local physical column, the **Open table** panel of the base table is opened.

Editing column metadata information

Change column information such as the label, format type, and description, directly from the grid.

You can change the column metadata information directly from within the Grid view of the Trillion-Row Spreadsheet and the Macro Language Workshop.

When you edit column metadata information in the grid, the changes are not permanent and apply only to the worksheet in your session. You can save the metadata changes for future use by saving your query. To make any metadata changes permanent, you must materialize the data to a 1010data Insights Platform table.

Column metadata dialog

The **Column metadata** dialog contains fields and options for editing the metadata information of a column.

The **Column metadata** dialog is accessible from the Grid view of a table or worksheet within the Trillion-Row Spreadsheet and the Macro Language Workshop.

Figure 70: Column metadata dialog

Label

This field allows you to edit the column label.

The column label can contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters. You can create a multi-line column label by using the backtick character (`) to separate the lines (e.g., "Percentage of `Total Sales (%)").

Note: The column *label* is an optional descriptive column title and is different from the column *name*. For more information see [Columns](#) on page 48.

Display Format

This drop-down list allows you to change how values are displayed within the column.

For example, you can choose to display data in a column as a date, a number, as text, and so forth. For more information, see [Display formats](#) on page 56.

Width

This field allows you to set the width of the column.

Column width is specified by the number of place values to hold in the column. In other words, the width of the column is determined by the number of characters it can hold. You can specify up to 100 characters.

Note: You can also drag a column header in the grid to resize a column. For instructions, see [Resize a column within the grid](#) on page 300.

Decimals

This drop-down list allows you to choose how many numbers are displayed after the decimal point.

You can specify up to 9 decimal places.

Fixed column

This option allows you to determine whether or not the column is fixed.

A fixed column is always visible and is separated from the other columns in the table or worksheet by a thicker vertical line. Fixed columns do not move, regardless of how far the table is scrolled in either direction. A fixed column is useful when the column has key information that should always be shown.

Description

This field allows you to edit the explanation of the values in the column. For example, you might want to describe how the values in the column are calculated or explain what the column values represent.

You can customize the appearance of the description text by using the provided rich text editor controls. These controls are located above the **Description** field and include the following:

Format

This drop-down list contains pre-formatted text styles. For example, you can choose different heading levels to organize the description text.

Bold

This icon applies bold formatting to the selected text.

Italic

This icon makes the selected text italic.

Underline

This icon allows you to underline the selected text.

Insert hyperlink

This icon opens the **Insert hyperlink** dialog which allows you to create a link in the description text.

Insert image

This icon opens the **Insert image** dialog which allows you to add an image in the description text.

Cancel

Click this button to close the dialog window without applying any changes to the column.

Submit

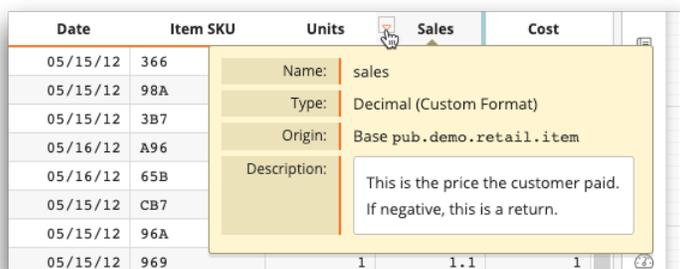
Click this button to apply the changes to the column and close the dialog window.

Change the metadata information of a column

As necessary, you can edit the metadata information of a column directly from the Grid view of a table or worksheet.

To change the metadata information of a column:

1. In the Grid view of an open table or worksheet, place the pointer in the heading of the column for which you want to change the metadata information. The Trillion-Row Spreadsheet displays a drop-down arrow. If you hover your mouse on the arrow, a tooltip displays information about the column.



Date	Item SKU	Units	Sales	Cost
05/15/12	366			
05/15/12	98A			
05/15/12	3B7			
05/16/12	A96			
05/16/12	65B			
05/15/12	CB7			
05/15/12	96A			
05/15/12	969	1	1.1	1

2. Click the drop-down arrow (▼) and then click **Modify info**. The Insights Platform displays the **Column metadata** dialog for the appropriate column.

Column metadata for sa1es

Label: Sales

Display Format: Number: 1,234,567.89

Width: [] Decimals: []

Fixed column:

Description:

Format: [] **B** *I* U ↺ 🖼️

This is the price the customer paid. If negative, this is a return.

Cancel Submit

3. Edit the appropriate column metadata information.

For a list of fields and options, see [Column metadata dialog](#) on page 297.

4. Click **Submit**.

The Trillion-Row Spreadsheet applies the changes to the column metadata information and adds a **Modify** panel to the timeline.

pub.demo.retail.item

Analyze Query View Visualize Develop Export

1 + - < > Open table pub.demo.retail.item

2 + - < > Modify ↺ ✓ ⏪ ⏩

ⓘ The following column metadata changes are in effect. An empty field indicates no change.

Column	Label	Width	Format	Fix?	Desc
sales	Sales		type:currency	<input type="checkbox"/>	This is the price the customer paid. If negative, this is a return.

3 < > New operation

Sales Item Detail
Cols 1 to 8 of 8, Rows 1 to 19 of 35

Transaction ID	Account	Store	Date	Item SKU	Units	Sales	Cost
531	957	1	05/15/12	366	-1	\$-5.00	-1.84
532	478	1	05/15/12	98A	1	\$0.50	0.25
532	478	1	05/15/12	3B7	1	\$1.10	0.56
534	738	1	05/16/12	A96	2	\$6.00	2.9
534	738	1	05/16/12	65B	1	\$2.25	1.35
535	709	2	05/15/12	CB7	1	\$1.65	1.1
535	709	2	05/15/12	96A	1	\$1.10	1
535	709	2	05/15/12	969	1	\$1.10	1
536	748	2	05/17/12	3A4	3	\$1.02	0.39
536	748	2	05/17/12	366	1	\$5.00	1.8
537	523	3	05/15/12	CB7	1	\$1.65	1.1
537	523	3	05/15/12	A96	2	\$6.00	2.9
537	523	3	05/15/12	98A	1	\$0.50	0.25
537	523	3	05/15/12	96A	1	\$1.10	1
537	523	3	05/15/12	3B7	1	\$1.10	0.56
538	668	1	05/18/12	CB7	1	\$1.65	1.1
538	668	1	05/18/12	98A	4	\$2.00	1
538	668	1	05/18/12	96A	1	\$1.10	1
538	668	1	05/18/12	969	1	\$1.10	1

Resize a column within the grid

Drag a column header within the grid to make the column wider or narrower.

You can resize a column directly from within the grid by dragging its column header.

To resize a column within the grid:

- In the Grid view of an open table or worksheet, place the pointer in the heading of the column that you want to resize.
The Trillion-Row Spreadsheet displays a blue handle indicating the right border of the column.

Sales Item Detail
Cols 1 to 8 of 8, Rows 1 to 19 of 35

Transaction ID transid	Account account	Store store	Date date	Item SKU sku	Units units	Sales sales	Cost cost
531	957	1	05/15/12	366	-1	-5	-1.84
532	478	1	05/15/12	98A	1	0.5	0.25
532	478	1	05/15/12	3B7	1	1.1	0.56

- Do one of the following:
 - Drag the handle to the right to make the column wider.
 - Drag the handle to the left to make the column narrower.

As you drag the handle, the Trillion-Row Spreadsheet indicates what the column width in characters will be when the handle is released.

Note: A column must be at least one character wide.

After releasing the handle, the Trillion-Row Spreadsheet resizes the column and adds the **Modify** panel to the timeline.

pub.demo.retail.item

Analyze Query View Visualize Develop Export

1 + Open table pub.demo.retail.item

2 + - < > Modify

The following column metadata changes are in effect. An empty field indicates no change.

Column	Label	Width	Format	Fix?	Desc
transid		10	width:10	<input type="checkbox"/>	

3 </> New operation

Sales Item Detail
Cols 1 to 7 of 8, Rows 1 to 19 of 35

Transaction ID	Account	Store	Date	Item SKU	Units	Sales
531	957	1	05/15/12	366	-1	-
532	478	1	05/15/12	98A	1	0.
532	478	1	05/15/12	3B7	1	1.
534	738	1	05/16/12	A96	2	
534	738	1	05/16/12	65B	1	2.2
535	709	2	05/15/12	CB7	1	1.6
535	709	2	05/15/12	96A	1	1.
535	709	2	05/15/12	969	1	1.
536	748	2	05/17/12	3A4	3	1.0
536	748	2	05/17/12	366	1	
537	523	3	05/15/12	CB7	1	1.6
537	523	3	05/15/12	A96	2	
537	523	3	05/15/12	98A	1	0.
537	523	3	05/15/12	96A	1	1.
537	523	3	05/15/12	3B7	1	1.
538	668	1	05/18/12	CB7	1	1.6
538	668	1	05/18/12	98A	4	
538	668	1	05/18/12	96A	1	1.

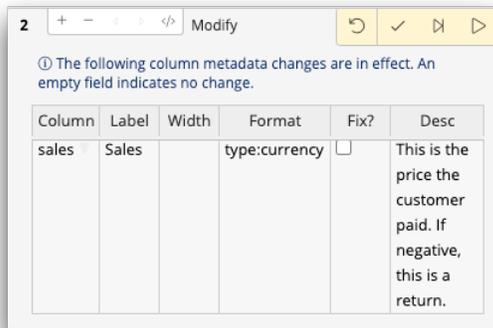
In the example above, the **Modify** panel displays the new column width of 10. For more information about this panel, see [Modify panel](#) on page 301.

Modify panel

The **Modify** panel appears in the timeline after the metadata of a column is changed within the grid.

The **Modify** panel displays the changes to the column metadata information, such as a display format change or a column description change.

Figure 71: Modify panel



In the above example, the display format for the **Sales** column was changed to the currency format.

Draw a selection rectangle

Choose a group of cells by drawing a selection rectangle around a collection of cells in the grid.

You can choose a group of cells in the Grid view by drawing a selection rectangle around a collection of cells. This is useful when you want to perform a selection on a group of cells or copy only a portion of the data in the table or worksheet.

To draw a selection rectangle in the grid:

In the Grid view of an open table or worksheet, completing the following:

- Click the cell that you want as the upper left corner of the rectangle.
- Press **Shift** and then click the cell that you want as the bottom right corner.

The Trillion-Row Spreadsheet surrounds the selected cells with an orange rectangle.

Sales Item Detail
Cols 1 to 8 of 8, Rows 1 to 19 of 35

Transaction ID transid	Account account	Store store	Date date	Item SKU sku	Units units	Sales sales	Cost cost
531	957	1	05/15/12	366	-1	-5	-1.84
532	478	1	05/15/12	98A	1	0.5	0.25
532	478	1	05/15/12	3B7	1	1.1	0.56
534	738	1	05/16/12	A96	2	6	2.9
534	738	1	05/16/12	65B	1	2.25	1.35
535	709	2	05/15/12	CB7	1	1.65	1.1
535	709	2	05/15/12	96A	1	1.1	1
535	709	2	05/15/12	969	1	1.1	1
536	748	2	05/17/12	3A4	3	1.02	0.39
536	748	2	05/17/12	366	1	5	1.8
537	523	3	05/15/12	CB7	1	1.65	1.1
537	523	3	05/15/12	A96	2	6	2.9
537	523	3	05/15/12	98A	1	0.5	0.25
537	523	3	05/15/12	96A	1	1.1	1
537	523	3	05/15/12	3B7	1	1.1	0.56
538	668	1	05/18/12	CB7	1	1.65	1.1
538	668	1	05/18/12	98A	4	2	1
538	668	1	05/18/12	96A	1	1.1	1

Grid menu options

The grid menu provides options for interacting with the data in a table or worksheet directly from within the grid itself.

The grid menu is accessed by right-clicking a cell in the Grid view of the results pane of either TRS or MLW. Options available in the grid menu may vary and are dependent on table metadata and applied query operations.

When you apply an operation from the grid menu, the Insights Platform adds the operation to the timeline. In the case of TRS, the operation is inserted after the open step in the Analysis Timeline. In the case of MLW, the operation is always added as XML code to the end of your code.

Note: In this topic, *[COLUMN_LABEL]* represents the label of the column in which a cell is located.

Transaction ID transid	Account account	Store store	Date date	Item SKU sku	Units units	Sales sales
531	957	1	05/15/12	366	-1	-
532	478	1				0.
532	478	1				1.
534	738	1				2.
535	709	2				1.
535	709	2				1.
535	709	2				1.
536	748	3				1.
536	748	2				
537	523	3				1.
537	523	2				
537	523	3				0.
537	523	3				1.
537	523	3				1.
538	668	1				1.
538	668	1				

Figure 72: Grid menu options

Copy

Copies data from a single cell, a group of cells, or all of the cells within a row or column.

For instructions, see [Copy values in the grid](#) on page 305.

Download

Downloads all values, or a group of selected values, from the grid as a CSV, Microsoft Excel workbook, or PDF file. Files are saved as `download.[file extension]`.

View this row in single-row mode

Displays the row of the selected cell in the Single-row view.

For more information, see [Single-row view](#) on page 325.

Analyze underlying data for this cell

Displays the rows of data used to calculate the value within the selected cell in a new **TRS** window.

The timeline in the new Trillion-Row Spreadsheet represents the query up to, but not including, the tabulation and any additional selections based on the tabulation breaks for the selected cell.

The **Analyze underlying data for this cell** option appears in the menu if the displayed worksheet is the result of a quick summary, tabulation, or cross-tabulation.

Quick find

Locates rows in the grid based on the value of the selected cell.

For instructions, see [Find rows based on cell value](#) on page 315.

Quick select where

Selects rows in the grid based on the value of the selected cell.

For instructions, see [Select rows based on cell value](#) on page 151.

Quick select rows

Selects rows in the grid based on the location of the selected cell.

For instructions, see [Select rows based on cell location](#) on page 152.

Quick link

Links the base table or worksheet to the foreign table or worksheet selected in the submenu.

The **Quick link** option appears in the menu if the current table or worksheet is either prelinked to one or more foreign tables you can access, or has a metadata attribute specifying a conventional link to one or more foreign tables.

For more information, see [Prelinks](#) in the *1010data Cookbook*.

Quick analysis of [COLUMN_LABEL]

Displays statistical information about the data values located with the column of the selected cell.

For instructions, see [Analyze column data](#) on page 308.

Find in [COLUMN_LABEL]

Locates rows in the grid by searching for a value within the column of the selected cell.

For instructions, see [Find rows by searching for a value within a column](#) on page 319.

Sort on [COLUMN_LABEL]

Reorders the rows of a table such that the values in the column of the selected cell are in ascending or descending order.

For instructions, see [Sort rows within the grid](#) on page 158.

Select on [COLUMN_LABEL]

Displays the **Select rows** panel which allows you to create a select operation directly from within the grid. You can specify the selection criteria using the simple comparisons options or enter a selection expression in the Expression Editor.

For more information, see [Select rows panel](#) on page 142.

Link on [COLUMN_LABEL]

Displays the **Link tables** panel which allows you to create a link operation directly from within the grid.

For more information, see [Link tables panel](#) on page 163.

Tabulate using [COLUMN_LABEL]

Displays the **Tabulation** panel which allows you to perform any of the three basic types of tabulations (quick summaries, tabulations, and cross tabulations) directly from within the grid. For more information, see [Tabulation panel](#) on page 184.

Select the **as a break column** option from the sub menu to perform a tabulation using the column in which the selected cell is located as the row break. For instructions, see [Perform a tabulation](#) on page 190.

Select an option from the **as a break column** sub menu to perform the desired quick summary on the column in which the selected cell is located. For instructions, see [Perform a quick summary](#) on page 186.

Amend [COLUMN_LABEL]

Displays the **Amend column** panel which allows you to perform an amend operation directly from within the grid.

For more information, see [Amend column panel](#) on page 225.

Create computed column here

Displays the **Computed column** panel which allows you to create a computed column directly from within the grid.

For more information, see [Create a computed column](#) on page 181.

Undo last step

Undoes the last operation, or action applied to an operation, in the Analysis Timeline.

This option appears in the menu after a change is made to the timeline.

Redo last step

Reapplies the last previously undone operation, or action applied to an operation, in the Analysis Timeline.

This option appears in the menu after an operation or action is undone.

Copy values in the grid

In the Grid view, you have the option of copying data from a single cell, a group of cells, or all of the cells within a row or column. You can copy up to 100,000 cells of data at a time.

When copying values in the grid, you can choose to retain the formatting of the data or to copy only the underlying value. For example, a date appearing as 06/26/69 is copied as "19690626" (its underlying value). With formatting, the data is copied as it appears ("06/26/69"). For more information about formats, see [Display formats](#) on page 56.

To copy values in the grid:

1. In the Grid view of an open table or worksheet, select the cell or cells that you want to copy.

Depending on the values you want to copy, do the following:

To copy	Do the following
A single value	Right-click the cell you want to copy.
A group of values	First, draw a selection rectangle around a collection of cells you want to copy. For instructions see Draw a selection rectangle on page 302. Then, right-click any cell within the selection rectangle.
All values in a column	Right-click any cell within the column that you want to copy.
All values in a row	Right-click any cell within the row that you want to copy.

The Trillion-Row Spreadsheet displays the grid menu.

2. From the grid menu, point to **Copy** and then click the appropriate option.

For a list and description of the available options, see [Copy menu options](#) on page 306.

If more than one cell is selected, the Trillion-Row Spreadsheet displays a confirmation message. Otherwise, the selected value is copied.

The screenshot shows a window titled "Sales Item Detail" with a search icon in the top right. Below the title bar, it says "Cols 1 to 7 of 8, Rows 1 to 19 of 35". The main area is a table with the following columns: Transaction ID, Account, Store, Date, Item SKU, Units, and Sales. A confirmation dialog box is overlaid on the table, displaying the text "Confirm copy of 12 values" and two buttons: "Copy" (highlighted in orange) and "Cancel". The table data is as follows:

Transaction ID	Account	Store	Date	Item SKU	Units	Sales
531	957	1	05/15/12	366	-1	
532	478					0.
532	478					1.
534	738					2.2
534	738					1.6
535	709					1.
535	709	2	05/15/12	96A	1	1.
535	709	2	05/15/12	969	1	1.
536	748	2	05/17/12	3A4	3	1.0
536	748	2	05/17/12	366	1	
537	523	3	05/15/12	CB7	1	1.6
537	523	3	05/15/12	A96	2	
537	523	3	05/15/12	98A	1	0.
537	523	3	05/15/12	96A	1	1.
537	523	3	05/15/12	3B7	1	1.
538	668	1	05/18/12	CB7	1	1.6
538	668	1	05/18/12	98A	4	
538	668	1	05/18/12	96A	1	1.

- If a confirmation is displayed, click **Copy**.
The Trillion-Row Spreadsheet copies the selected values.

Copy menu options

The copy menu provides multiple options for copying values directly from within the grid.

The copy menu options are accessed by right-clicking a cell in the Grid view and then pointing to **Copy** in the menu.

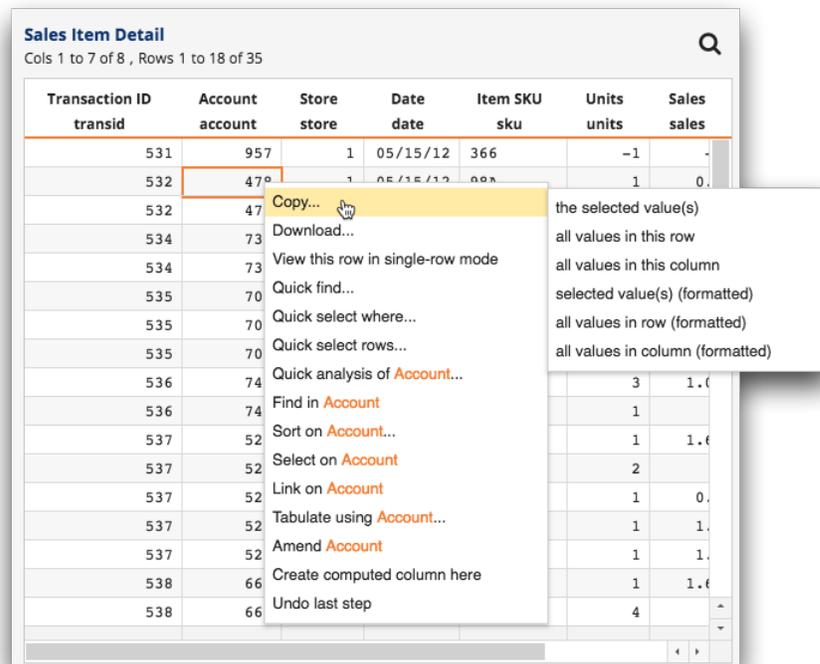


Figure 73: Copy menu options in the grid

the selected value(s)

Copies the underlying values selected in the grid.

You can also use the following keyboard shortcuts to copy the selected values:

- **Ctrl+C** (PC)
- **Command+C** (Mac)

all values in this row

Copies all of the underlying values in the selected row.

Copied row values are separated by the tab character.

all values in this column

Copies all of the underlying values in the selected column.

Copied column values are separated by the line feed (LF) newline character.

selected value(s) (formatted)

Copies the values selected in the grid. If a value is formatted, the display format is also copied.

Display formats control how values are displayed within columns in the 1010dat Insights Platform. For more information, see [Display formats](#) on page 56.

all values in row (formatted)

Copies all of the values in the selected row. If a value is formatted, the display format is also copied.

Display formats control how values are displayed within columns in the 1010dat Insights Platform. For more information, see [Display formats](#) on page 56.

Copied row values are separated by the tab character.

all values in column (formatted)

Copies all of the values in the selected column. If the column is formatted, the column display format is also copied.

Display formats control how values are displayed within columns in the 1010dat Insights Platform. For more information, see [Display formats](#) on page 56.

Copied column values are separated by the line feed (LF) newline character.

Analyze column data

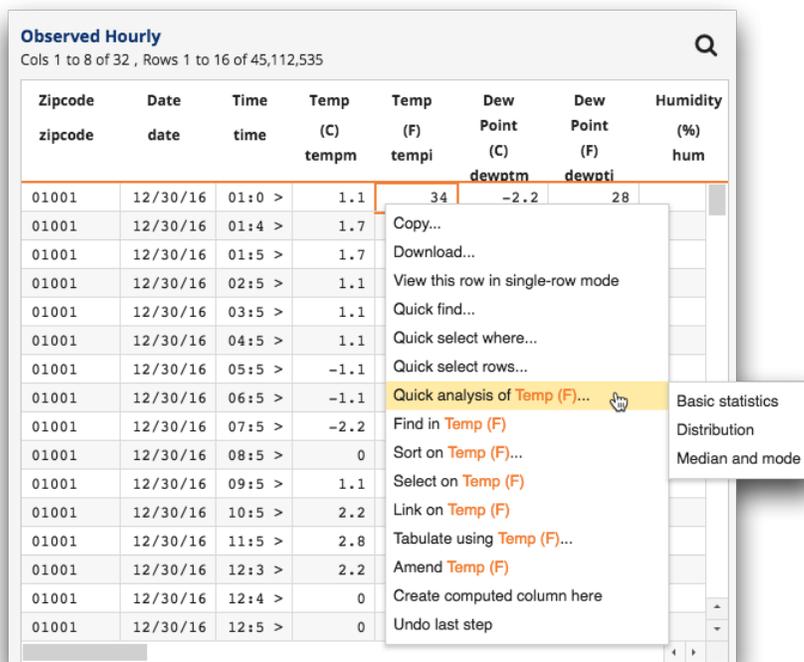
View statistics about the data values located within a column in the grid.

To analyze column data:

1. In the grid view of an open table or worksheet, right-click a cell within the column that you want to analyze and point to **Quick analysis of [COLUMN_LABEL]**.

Note: In this topic, [COLUMN_LABEL] represents the label of the column in which the selected cell is located.

The Trillion-Row Spreadsheet displays a list of column analysis options.



2. Click the desired column analysis option.

For a list and description of the available options, see [Quick analysis menu options](#) on page 308.

The Trillion-Row Spreadsheet displays the selected analysis view.

Quick analysis menu options

Display information about a particular column in the current table or worksheet.

The Quick analysis menu options are accessed by right-clicking a cell in the Grid view and then pointing to **Quick analysis of [COLUMN_LABEL]**.

Note: In this topic, [COLUMN_LABEL] represents the label of the column in which the selected cell is located.

Observed Hourly
Cols 1 to 8 of 32 , Rows 1 to 16 of 45,112,535

Zipcode	Date	Time	Temp	Temp	Dew	Dew	Humidity
zipcode	date	time	(C)	(F)	Point	Point	(%)
			tempm	tempi	(C)	(F)	hum
					dewotm	dewoti	
01001	12/30/16	01:0 >	1.1	34	-2.2	28	
01001	12/30/16	01:4 >	1.7				
01001	12/30/16	01:5 >	1.7				
01001	12/30/16	02:5 >	1.1				
01001	12/30/16	03:5 >	1.1				
01001	12/30/16	04:5 >	1.1				
01001	12/30/16	05:5 >	-1.1				
01001	12/30/16	06:5 >	-1.1				
01001	12/30/16	07:5 >	-2.2				
01001	12/30/16	08:5 >	0				
01001	12/30/16	09:5 >	1.1				
01001	12/30/16	10:5 >	2.2				
01001	12/30/16	11:5 >	2.8				
01001	12/30/16	12:3 >	2.2				
01001	12/30/16	12:4 >	0				
01001	12/30/16	12:5 >	0				

Figure 74: Quick analysis menu in the grid

Basic statistics

Displays the Quick Statistics view which provides summary statistics about the data in a selected column. For more information, see [Quick statistics](#) on page 309.

Distribution

Displays the Distribution view which lists the frequency, or count of the occurrences, of values in a selected column. For more information, see [Distribution](#) on page 311.

Median and mode

Displays the Median and mode view which provides median and mode statistics for a column selected in a selected column. For more information, see [Median and mode](#) on page 311.

Quick statistics

The Quick statistics view provides summary statistics about the data in a selected column.

The **Quick statistics** view in the results pane is accessed by right-clicking a cell in the grid and then selecting **Quick analysis of [COLUMN_LABEL] > Basic statistics** from the menu.

Note: In this topic, [COLUMN_LABEL] represents the label of the column in which the selected cell is located.

Count t0	45,112,535
Low t1	-77.8
High t2	167
# of Valid Records t3	44,906,134
# of N/A's t4	206,401
Average t5	45.89
Std. Deviation t6	22.18

Figure 75: Quick statistics view in the results pane of the TRS

The following information is available in the Quick Statistics view:

Count

The total number of values (both N/A and non-N/A) in the selected column.

Low

The minimum value located in the selected column.

High

The maximum value located in the selected column.

of Valid Records

The number of non-N/A values located in the selected column.

of N/A's

The number of N/A values located in the selected column.

Average

The average of the values in the selected column. (numeric columns only)

- If the column contains N/A values, those rows are ignored (i.e., the result is the average of the non-N/A values).
- If the column contains only N/A values, the result is N/A.

Std. Deviation

The standard deviation of the values in the selected column. (numeric columns only)

- If the column contains N/A values, those rows are ignored (i.e., the result is the standard deviation of the non-N/A values).
- If the column contains only N/A values, the result is N/A.

#ofUniqs

The number of unique values in the selected column.

- If the column contains N/A values, those rows are ignored (i.e., the result is the number of unique non-N/A values).
- If the column contains only N/A values, the result is 0.

Distribution

The Distribution view displays the frequency, or count of the occurrences, of values in a selected column.

The **Distribution** view in the results pane is accessed by right-clicking a cell in the grid and then selecting **Quick analysis of [COLUMN_LABEL] > Distribution** from the menu.

Note: In this topic, [COLUMN_LABEL] represents the label of the column in which the selected cell is located.

Temp (F)	Count t0	Count as % Ratio of Grand Count t1
temp1	45,112,535	100.0000%
	206,401	0.4575%
-77.8	4	0.0000%
-43.6	7	0.0000%
-41.8	9	0.0000%
-40	7	0.0000%
-38.2	64	0.0001%
-36.4	165	0.0004%
-34.6	289	0.0006%
-34.1	27	0.0001%
-33	9	0.0000%
-32.8	654	0.0014%
-32.1	36	0.0001%
-31	661	0.0015%
-30.3	5	0.0000%
-30.1	5	0.0000%
-29.9	86	0.0002%

Figure 76: Distribution view in the results pane of the TRS

The following information is available in the Distribution view:

[COLUMN_LABEL]

Lists each unique value located in the selected column.

Count

The number of times each unique value appears in the selected column.

Count as % Ratio of Grand Count

The number of times each unique value appears in the selected column represented as a percentage of the total count.

Median and mode

Displays the median and mode statistics for a column selected in the grid.

The **Median and mode** view in the results pane is accessed by right-clicking a cell in the grid and then selecting **Quick analysis of [COLUMN_LABEL] > Median and mode** from the menu.

Note: In this topic, *[COLUMN_LABEL]* represents the label of the column in which the selected cell is located.



Median t0	48
Mode t1	32
Frequency of mode t2	926,569

Figure 77: Median and mode view in the results pane of the TRS

The following information is available in the Median and mode view:

Median

The median, or middle, value located in the selected column.

This middle value separates the lower half of the values in the column from the higher half.

Mode

The value that appears most often in the selected column.

Frequency of mode

The number of times the mode value appears in the selected column.

Finding rows in the grid

Locate rows in a table or worksheet without eliminating the rows that do not match your criteria.

Finding rows in the Grid view of a table or worksheet is similar to using a select operation to isolate rows. However, unlike the select operation, rows are not eliminated from the table or worksheet. Instead, found rows are highlighted in the grid so that they are easy to see amongst the rest of the data. If desired, you can create a select operation from the results. The outcome is a new worksheet containing only the rows that were previously highlighted.

There are multiple ways of finding rows in a table or worksheet. You can enter any valid selection expression in the **Find rows** dialog, select an existing cell value in the grid to perform a Quick find, or search for a value within a specified column.

Find rows dialog

The **Find rows** dialog allows you to identify rows in a table or worksheet by entering any valid selection expression, scroll through found rows in the grid, and create a select operation from the results.

The **Find rows** dialog is accessible from the Grid view by clicking the **Find rows** (🔍) icon located at the top of the grid.

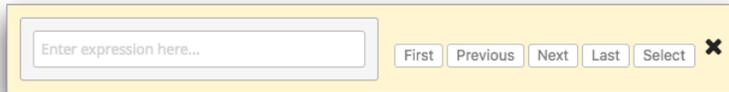


Figure 78: Find rows dialog in the grid

Expression editor

Enter a selection expression to find rows in the table or worksheet.

For more information, see [Writing Expressions](#) in the *1010data Reference Manual*.

First

Displays the first row of the results at the top of the grid.

Previous

Displays the previous row of the results at the top of the grid.

Next

Displays the next row of the results at the top of the grid.

Last

Displays the last row of the results at the top of the grid.

Select

Creates a new select operation from the rows highlighted in the grid. The outcome is a new worksheet containing only the rows that were previously highlighted.

Hide the find dialog

Click the **Hide the find dialog** (✕) icon to close the **Find rows** dialog.

Find rows using a selection expression

Enter a selection expression in the Grid view of the Trillion-Row Spreadsheet to find rows within a table or worksheet.

Use a selection expression to identify rows in a table or worksheet. Found rows are highlighted in the grid so that you can easily see them amongst the rest of the data.

To find rows using a selection expression:

1. In the Grid view of an open table or worksheet, click the **Find rows** (🔍) icon.

The **Find rows** icon is located at the top of the grid.

The Trillion-Row Spreadsheet displays the **Find rows** dialog above the grid.

The screenshot shows a dialog box with an input field containing "Enter expression here..." and buttons for "First", "Previous", "Next", "Last", and "Select" with a close button (X). Below the dialog is a data grid with the following columns: transaction ID, Account, Store, Date, Item SKU, Units, and Sales. The grid contains 20 rows of data.

transaction ID	Account	Store	Date	Item SKU	Units	Sales
transid	account	store	date	sku	units	sales
531	957	1	05/15/12	366	-1	
532	478	1	05/15/12	98A	1	0.
532	478	1	05/15/12	3B7	1	1.
534	738	1	05/16/12	A96	2	
534	738	1	05/16/12	65B	1	2.2
535	709	2	05/15/12	CB7	1	1.6
535	709	2	05/15/12	96A	1	1.
535	709	2	05/15/12	969	1	1.
536	748	2	05/17/12	3A4	3	1.0
536	748	2	05/17/12	366	1	
537	523	3	05/15/12	CB7	1	1.6
537	523	3	05/15/12	A96	2	
537	523	3	05/15/12	98A	1	0.
537	523	3	05/15/12	96A	1	1.
537	523	3	05/15/12	3B7	1	1.
538	668	1	05/18/12	CB7	1	1.6
538	668	1	05/18/12	98A	4	

2. In the Expression Editor, enter a selection expression.

The selection expression may refer to one or more columns and may include standard arithmetic, relational, and logical operators as well as any Insight Platform functions.

For more information, see [Writing Expressions](#) in the *1010data Reference Manual*.

3. Do one of the following:

- Press **Enter** (PC).
- Press **Return** (Mac).
- Click anywhere outside of the Expression Editor field.

The Trillion-Row Spreadsheet validates the expression and enables the buttons in the dialog.



4. Click the appropriate button in the dialog to view the results.

For a list and description of available buttons, see [Find rows dialog](#) on page 313.

The Trillion-Row Spreadsheet highlights the found rows in yellow, based on the button clicked, and displays the appropriate row at the top of the grid.

units>1

First Previous Next Last Select X

transaction ID	Account	Store	Date	Item SKU	Units	Sales
transid	account	store	date	sku	units	sales
534	738	1	05/16/12	A96	2	
534	738	1	05/16/12	65B	1	2.2
535	709	2	05/15/12	CB7	1	1.6
535	709	2	05/15/12	96A	1	1.1
535	709	2	05/15/12	969	1	1.1
536	748	2	05/17/12	3A4	3	1.0
536	748	2	05/17/12	366	1	
537	523	3	05/15/12	CB7	1	1.6
537	523	3	05/15/12	A96	2	
537	523	3	05/15/12	98A	1	0.1
537	523	3	05/15/12	96A	1	1.1
537	523	3	05/15/12	3B7	1	1.1
538	668	1	05/18/12	CB7	1	1.6
538	668	1	05/18/12	98A	4	
538	668	1	05/18/12	96A	1	1.1
538	668	1	05/18/12	969	1	1.1
538	668	1	05/18/12	3B7	1	1.1

Find rows based on cell value

Use the value of a cell in the Grid view of the Trillion-Row Spreadsheet to find rows in a table or worksheet.

Select a cell in the grid and then use its value to identify rows in a table or worksheet. This allows you to find rows based on simple conditions without needing to enter an expression. Found rows are highlighted in the grid so that you can easily see them amongst the rest of the data.

To find rows based on cell value:

1. In the Grid view of an open table or worksheet, right-click the cell on which you want to base the find and then point to **Quick find**.
The Trillion-Row Spreadsheet displays a list of row location options.

Sales Item Detail
Cols 1 to 7 of 8 , Rows 1 to 18 of 35

Transaction ID transid	Account account	Store store	Date date	Item SKU sku	Units units	Sales sales
531	957	1	05/15/12	366	-1	-
532	478		Copy...		1	0.
532	478		Download...		1	1.
534	738		View this row in single-row mode		2	
534	738		Quick find...			
535	709		Quick select where...			
535	709		Quick select rows...			
535	709		Quick analysis of Store...			
536	748		Find in Store			
536	748		Sort on Store...			
537	523		Select on Store		1	1. €
537	523		Link on Store		2	
537	523		Tabulate using Store...		1	0.
537	523		Amend Store		1	1.
537	523		Create computed column here		1	1.
538	668		Undo last step		1	1. €
538	668	1	05/18/12	98A	4	

2. Point to the desired row location option.

For a list and description of the available options, see [Row location options in the Quick find menu](#) on page 317.

The Trillion-Row Spreadsheet displays a list of value criteria options.

Sales Item Detail
Cols 1 to 7 of 8 , Rows 1 to 18 of 35

Transaction ID transid	Account account	Store store	Date date	Item SKU sku	Units units	Sales sales
531	957	1	05/15/12	366	-1	-
532	478		Copy...		1	0.
532	478		Download...		1	1.
534	738		View this row in single-row mode		2	
534	738		Quick find...			
535	709		Quick select where...			
535	709		Quick select rows...			
535	709		Quick analysis of Store...			
536	748		Find in Store			
536	748		Sort on Store...			
537	523		Select on Store		1	1. €
537	523		Link on Store		2	
537	523		Tabulate using Store...		1	0.
537	523		Amend Store		1	1.
537	523		Create computed column here		1	1.
538	668		Undo last step		1	1. €
538	668	1	05/18/12	98A	4	

next row where...
prev. row where...
first row where...
last row where...
again

Store has the value 1
Store does not have the value 1
Store is N/A
Store is not N/A
Store is greater than 1
Store is less than 1

3. Click the desired value criteria option.

Available value criteria options are specific to the selected cell. For a list and description of the available options, see [Value criteria options in the Quick find menu](#) on page 318.

The Trillion-Row Spreadsheet highlights the results and scrolls to the appropriate location in the grid so that the appropriate row is displayed at the top. In addition, the selection expression representing your search is displayed within the Expression Editor, in the **Find rows** dialog, above the grid.

transaction ID transid	Account account	Store store	Date date	Item SKU sku	Units units	Sales sales
535	709	2	05/15/12	CB7	1	1.6
535	709	2	05/15/12	96A	1	1.0
535	709	2	05/15/12	969	1	1.0
536	748	2	05/17/12	3A4	3	1.0
536	748	2	05/17/12	366	1	1.0
537	523	3	05/15/12	CB7	1	1.6
537	523	3	05/15/12	A96	2	1.0
537	523	3	05/15/12	98A	1	0.0
537	523	3	05/15/12	96A	1	1.0
537	523	3	05/15/12	3B7	1	1.0
538	668	1	05/18/12	CB7	1	1.6
538	668	1	05/18/12	98A	4	1.0
538	668	1	05/18/12	96A	1	1.0
538	668	1	05/18/12	969	1	1.0
538	668	1	05/18/12	3B7	1	1.0
539	25	2	06/19/12	65B	1	1.0
540	361	2	06/19/12	3B7	1	1.0

Row location options in the Quick find menu

The **Quick find** menu provides options for finding rows based on the location of a cell in the grid.

The find row location options are accessed by right-clicking a cell in the Grid view and then pointing to **Quick find** in the menu.

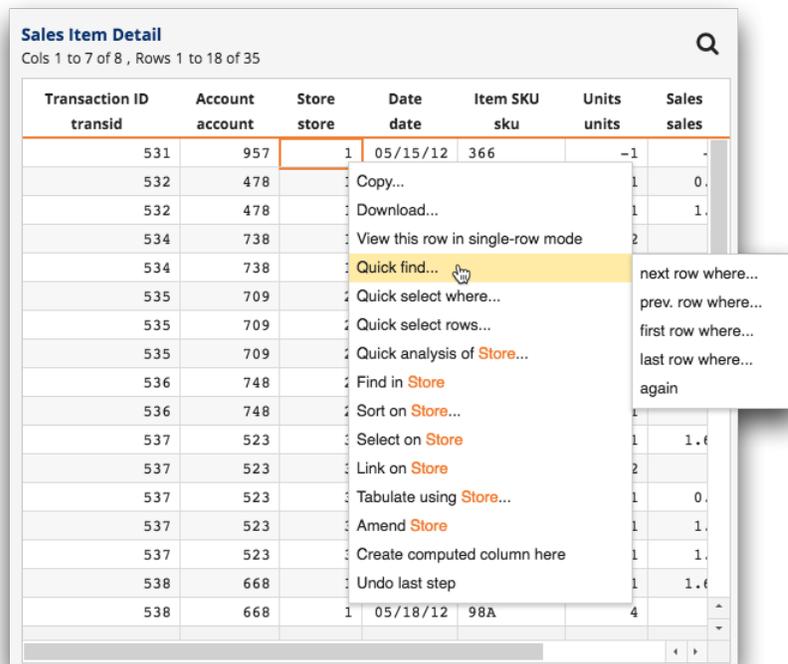


Figure 79: Row location options in the Quick find menu

The options in this menu allow you to find a row in the grid based on the location of the currently selected cell. For example, you can choose to find the next row in the table or worksheet, in relation to the currently selected cell, that matches the value criteria. The found row will be displayed at the top of the grid.

The following row location options are available:

next row where

Displays the next row in the table or worksheet, in relation to the location of the currently selected cell, that matches the value criteria.

prev. row where

Displays the previous row in the table or worksheet, in relation to the location of the currently selected cell, that matches the value criteria.

first row where

Displays the first row in the table or worksheet that matches the value criteria.

last row where

Displays the last row in the table or worksheet that matches the value criteria.

again

Repeats the previous search, if any, starting from the current row.

Value criteria options in the Quick find menu

The **Quick find** menu provides options for finding rows based on the value of a cell in the grid.

The find row value criteria options are accessed by right-clicking a cell in the Grid view, pointing to **Quick find**, and then pointing to a row location option in the menu.

Note: In this topic, *[COLUMN_LABEL]* represents the label of the column in which the selected cell is located and *[CELL_VALUE]* represents the value contained in the selected cell.

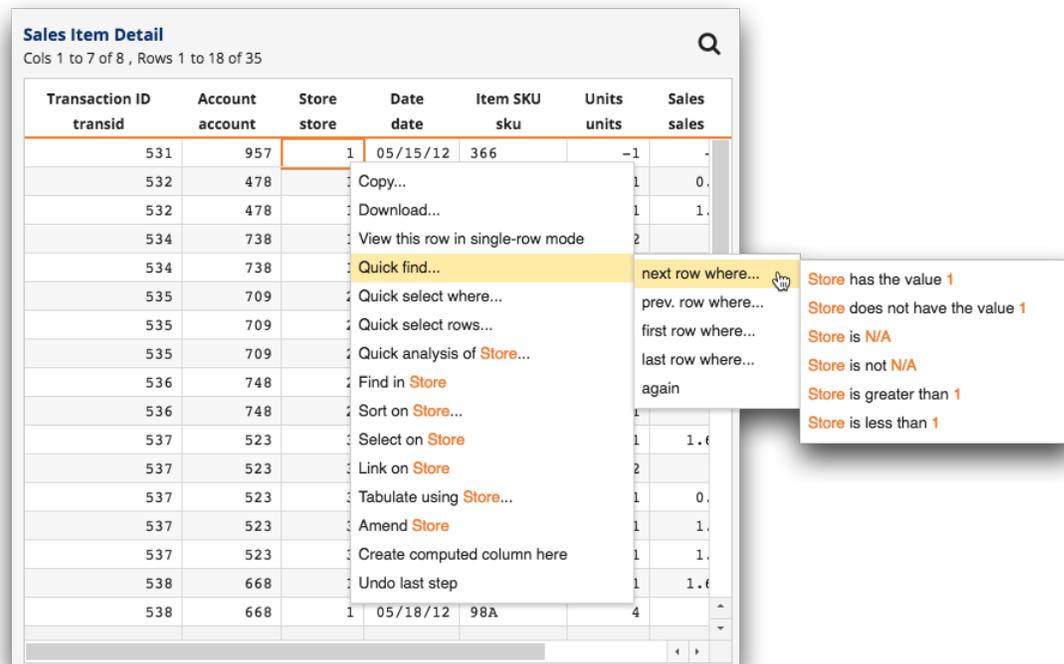


Figure 80: Value criteria options in the Quick find menu

The options in this menu allow you to find a row in the grid based on the cell value. Rows that match the selected value criteria option will be highlighted in the grid.

Note: Available value criteria options are specific to the selected cell.

The following find row value criteria options are available:

[COLUMN_LABEL] has the value [CELL_VALUE]

Highlights rows in the grid that have the same value, within the same column, as the selected cell.

[COLUMN_LABEL] does not have the value [CELL_VALUE]

Highlights rows in the grid that do not have the same value, within the same column, as the selected cell.

[COLUMN_LABEL] is N/A

Highlights rows in the grid, within the same column as the selected cell, that contain an N/A value.

[COLUMN_LABEL] is not N/A

Highlights rows in the grid, within the same column as the selected cell, that contain anything other than an N/A value.

[COLUMN_LABEL] is greater than [CELL_VALUE]

Highlights rows in the grid that have a larger value, within the same column, as the selected cell.

[COLUMN_LABEL] is less than [CELL_VALUE]

Highlights rows in the grid that have a smaller value, within the same column, as the selected cell.

Find rows by searching for a value within a column

Search for a value within a column to find rows in a table or worksheet.

You can find rows in a table or worksheet by searching for a value within a column. While similar to finding rows based on an existing cell value, this feature allows you to enter any value to use in the search.

To find rows by searching for a value within a column:

1. In the Grid view of an open table or worksheet, right-click a cell within the column that you want to search.
2. From the menu, click **Find in [COLUMN_LABEL]**.

The [COLUMN_LABEL] is the label of the column in which the cell you right-clicked is located.

The Trillion-Row Spreadsheet displays the **Find in column** dialog.

The screenshot shows a spreadsheet titled "Sales Item Detail" with columns: Transaction ID, Account, Store, Date, Item SKU, Units, and Sales. A right-click context menu is open over the 'Store' column, showing the option "is equal to". The table data is as follows:

Transaction ID	Account	Store	Date	Item SKU	Units	Sales
transid	account	store	date	sku	units	sales
531	957	1	05/15/12	366	-1	-
532	478	1	05/15/12	98A	1	0.
532	478	1	05/15/12	3B7	1	1.
534	738	1	05/16/12	A96	2	
534	738	1	05/16/12	65B	1	2.
535	709	2	05/15/12	CB7	1	1.
535	709	2	05/15/12	96A	1	1.
535	709	2	05/15/12	969	1	1.
536	748					
536	748					
537	523	3	05/15/12	CB7	1	1.
537	523	3	05/15/12	A96	2	
537	523	3	05/15/12	98A	1	0.
537	523	3	05/15/12	96A	1	1.
537	523	3	05/15/12	3B7	1	1.
538	668	1	05/18/12	CB7	1	1.
538	668	1	05/18/12	98A	4	

3. In the drop-down list, select a find option.

Available options are specific to the data type of the column. For a list and description of the options by data type, see [Find in column options](#) on page 321.

4. Enter a value in the field and then do one of the following:

- Press **Enter** (PC).
- Press **Return** (Mac).

The Trillion-Row Spreadsheet highlights the results and scrolls to the appropriate location in the grid so that the appropriate row is displayed at the top. In addition, the selection expression representing your search is displayed within the Expression Editor, in the **Find rows** dialog, above the grid.

The screenshot shows a search interface with a text input field containing "store='2'", navigation buttons (First, Previous, Next, Last, Select), and a close button (X). Below the search bar is a data grid with the following columns: transaction ID (transid), Account (account), Store (store), Date (date), Item SKU (sku), Units (units), and Sales (sales). The grid contains 20 rows of data, with rows where the 'store' value is 2 highlighted in yellow.

transaction ID transid	Account account	Store store	Date date	Item SKU sku	Units units	Sales sales
535	709	2	05/15/12	CB7	1	1.6
535	709	2	05/15/12	96A	1	1.1
535	709	2	05/15/12	969	1	1.1
536	748	2	05/17/12	3A4	3	1.0
536	748	2	05/17/12	366	1	1.1
537	523	3	05/15/12	CB7	1	1.6
537	523	3	05/15/12	A96	2	1.1
537	523	3	05/15/12	98A	1	0.9
537	523	3	05/15/12	96A	1	1.1
537	523	3	05/15/12	3B7	1	1.1
538	668	1	05/18/12	CB7	1	1.6
538	668	1	05/18/12	98A	4	1.1
538	668	1	05/18/12	96A	1	1.1
538	668	1	05/18/12	969	1	1.1
538	668	1	05/18/12	3B7	1	1.1
539	25	2	06/19/12	65B	1	1.1
540	361	2	06/19/12	3B7	1	1.1

Find in column options

Descriptions of options that are available when finding rows by searching for a value within a column.

The **Find in column** dialog is accessible from the Grid view by right-clicking a cell and then clicking **Find in [COLUMN_LABEL]** from the menu.

Note: In this topic, *[COLUMN_LABEL]* represents the label of the column in which the selected cell is located.

Sales Item Detail
Cols: 1 to 7 of 8 , Rows: 1 to 18 of 35

Transaction ID transid	Account account	Store store	Date date	Item SKU sku	Units units	Sales sales
531	957	1	05/15/12	366	-1	-
532	478	1	05/15/12	98A	1	0.
532	478	1	05/15/12	3B7	1	1.
534	738	1	05/16/12	A96	2	
534	738	1	05/16/12	65B	1	2.
535	709	2	05/15/12	CB7	1	1.
535	709	2	05/15/12	96A	1	1.
535	709	2	05/15/12	969	1	1.
536	748					
536	748					
537	523	3	05/15/12	CB7	1	1.
537	523	3	05/15/12	A96	2	
537	523	3	05/15/12	98A	1	0.
537	523	3	05/15/12	96A	1	1.
537	523	3	05/15/12	3B7	1	1.
538	668	1	05/18/12	CB7	1	1.
538	668	1	05/18/12	98A	4	

is equal to

Figure 81: Find in column dialog

The options in the dialog drop-down list are dependant on the data type of the selected column. See the appropriate section below for a list of options based on the data type.

Integer, Big integer, Decimal

The following options are available for columns assigned the integer, big integer, or decimal data type.

is equal to

Highlights rows in the grid where data in the selected column is the same as the entered value.

is not equal to

Highlights rows in the grid where data in the selected column is anything other than the entered value.

is less than

Highlights rows in the grid where data in the selected column is smaller than the entered value.

is greater than

Highlights rows in the grid where data in the selected column is larger than the entered value.

is not less than

Highlights rows in the grid where data in the selected column is larger than or equal to the entered value.

is not greater than

Highlights rows in the grid where data in the selected column is smaller than or equal to the entered value.

Text

The following options are available for columns assigned the text data type.

contains

Highlights rows in the grid where data in the selected column contains the entered text string.

does not contain

Highlights rows in the grid where data in the selected column does not contain the entered text string.

begins with

Highlights rows in the grid where data in the selected column begins with the entered text string.

ends with

Highlights rows in the grid where data in the selected column ends with the entered text string.

matches pattern

Highlights rows in the grid where data in the selected column contains the entered text string pattern. The text string pattern includes text characters plus special "wildcard" characters.

For a list of available "wildcard" characters, and examples of how to use them, see [sm \(x;y\)](#) in the *1010data Reference Manual*.

contains (case-sensitive)

Highlights rows in the grid where data in the selected column contains the exact (case-sensitive) entered text string.

matches (case-sensitive)

Highlights rows in the grid where data in the selected column contains the exact (case-sensitive) entered text string pattern. The text string pattern includes text characters plus special "wildcard" characters.

For a list of available "wildcard" characters, and examples of how to use them, see [sm \(x;y\)](#) in the *1010data Reference Manual*.

is exactly (case-sensitive)

Highlights rows in the grid where data in the selected column is identical (case-sensitive) to the entered text string.

is not (case-sensitive)

Highlights rows in the grid where data in the selected column contains anything other than the exact (case-sensitive) entered text string.

matches regex

Highlights rows in the grid where data in the selected column matches the regular expression contained in the entered text string.

This option supports Perl compatible regular expressions (PCRE). See [PCRE - Perl Compatible Regular Expressions](#) for more information.

Scroll directly to a specific row

You can go directly to a specific row in the grid by right-clicking the vertical scroll bar and entering the row number in the dialog.

In addition to scrolling through the rows in the grid, the scroll bar allows you to go directly to a specific row number. You can also use this option to scroll through a page of rows at once.

To scroll directly to a specific row:

1. In the Grid view of an open table or worksheet, right-click the vertical scroll bar. The Trillion-Row Spreadsheet displays the **Scroll to row** dialog.

Observed Hourly
Cols 1 to 8 of 32 , Rows 1 to 16 of 45,112,535

Zipcode	Date	Time	Temp	Temp	Dew	Dew	Humidity
zipcode	date	time	(C)	(F)	Point	Point	(%)
			tempm	tempi	dewptm	dewpti	hum
01001	12/30/16	01:0 >	1.1	34	-2.2	28	
01001	12/30/16	01:4 >	1.7	35.1	-2.2	28	
01001	12/30/16	01:5 >	1.7	35.1	-2.2	28	
01001	12/30/16	02:5 >	1.1	34	-2.2	28	
01001	12/30/16	03:5 >	1.1	34	-2.2	28	
01001	12/30/16	04:5 >	1.1	34	-2.2	28	
01001	12/30/16	05:5 >	-1.1	30	-3.9	25	
01001	12/30/16	06:5 >	-1.1	30	-3.9	25	
01001	12/30/16	07:5 >	-2.2	28	-4.4	24.1	
01001	12/30/16	08:5 >	0	32	-3.3	26.1	
01001	12/30/16	09:5 >	1.1	34	-2.2	28	
01001	12/30/16	10:5 >	2.2	36	-3.3	26.1	
01001	12/30/16	11:5 >	2.8	37	-4.4	24.1	
01001	12/30/16	12:3 >	2.2	36	-4.4	24.1	
01001	12/30/16	12:4 >	0	32	-1.7	28.9	
01001	12/30/16	12:5 >	0	32	-2	28.4	

2. In the **Scroll to row** dialog, do one of the following:

Action

Instruction

Go to a specific row number

Enter a row number and press **Enter** (PC) or **Return** (Mac).

Go up one page

Click the **Page up** (▲) icon.

Go down one page

Click the **Page down** (▼) icon.

The Trillion-Row Spreadsheet displays the appropriate row or page in the grid.

Note: To dismiss the **Scroll to row** dialog, click anywhere outside of the dialog.

Scroll directly to a specific column

You can go directly to a specific column in the grid by right-clicking the horizontal scroll bar and choosing a column in the dialog.

In addition to scrolling through the columns in the grid, the scroll bar allows you to go directly to a specific column in the table or worksheet.

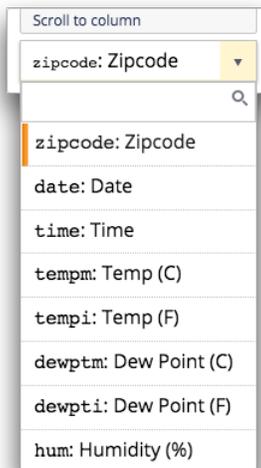
To scroll directly to a specific column:

1. In the Grid view of an open table or worksheet, right-click the horizontal scroll bar.
The Trillion-Row Spreadsheet displays the **Scroll to column** dialog.

Observed Hourly
Cols 1 to 8 of 32 , Rows 1 to 16 of 45,112,535

Zipcode	Date	Time	Temp (C)	Temp (F)	Dew Point (C)	Dew Point (F)	Humidity (%)
zipcode	date	time	tempm	tempi	dewptm	dewpti	hum
01001	12/30/16	01:0 >	1.1	34	-2.2	28	
01001	12/30/16	01:4 >	1.7	35.1	-2.2	28	
01001	12/30/16	01:5 >	1.7	35.1	-2.2	28	
01001	12/30/16	02:5 >	1.1	34	-2.8	27	
01001	12/30/16	03:5 >	1.1	34	-3.3	26.1	
01001	12/30/16	04:5 >	1.1	34	-3.9	25	
01001	12/30/16	05:5 >	-1.1	30	-3.9	25	
01001	12/30/16	06:5 >	-1.1	30	-3.9	25	
01001	12/30/16	07:5 >	-2.2	28	-4.4	24.1	
01001	12/30/16	08:5 >	0	32	-3.3	26.1	
01001	12/30/16	09:5 >	1.1	34	-2.2	28	
01001	12/30/16	10:5 >	2.2	36	-3.3	26.1	
01001	12/30/16	11:5 >				24.1	
01001	12/30/16	12:3 >				24.1	
01001	12/30/16	12:4 >				28.9	
01001	12/30/16	12:5 >	0	32	-2	28.4	

- In the **Scroll to column** dialog, click the drop-down (▼) icon. The dialog displays a column filter field and a list of the columns contained in the table or worksheet.



- Optionally, in the column filter field, enter the name or label of a column. As you enter text, the dialog filters the list and displays only the columns that match your entry.
- In the drop-down list, select a column. The Trillion-Row Spreadsheet displays the appropriate column in the grid.

Note: To dismiss the **Scroll to column** dialog without selecting a column, click anywhere outside of the dialog.

Single-row view

View the cells in a single row of data, arranged vertically.

When a table or worksheet has many columns, it can sometimes be helpful to use the Single-row view. This view allows you to scroll through the data vertically allowing you to more easily see all of the data within a row of the table or worksheet.

The Single-row view is accessible in the following ways:

- Clicking the Single-row view () icon in the View bar
- Right-clicking a cell in the grid and then selecting **View this row in single-row mode** from the menu
- Clicking **Single row** in the **View** tab of the Trillion-Row Spreadsheet (TRS)

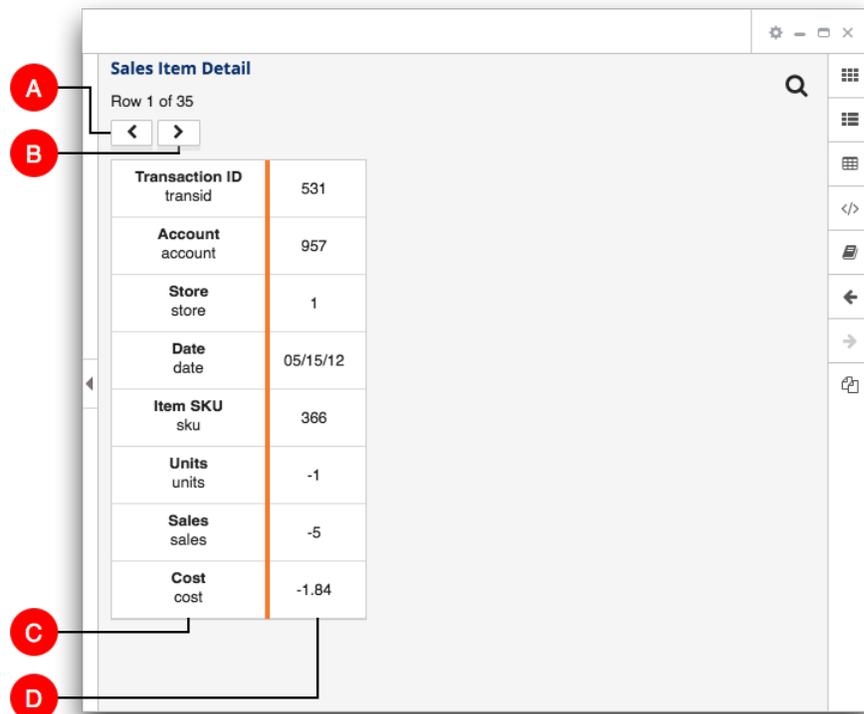


Figure 82: Single-row view in the results pane of the TRS

A. Previous row

Click the **Previous Row** () icon to view the previous row in the table or worksheet.

B. Next row

Click the **Next Row** () icon to view the next row in the table or worksheet.

C. Column header

The column header displays the column name, column, label, or both. You can choose which is displayed by adjusting your user profile settings. For instructions, see [Change your user settings](#) on page 85.

D. Row data

Data in the table or worksheet is displayed a single row at a time, arranged vertically.

Clone the TRS window

Clone the current view in the Trillion-Row Spreadsheet to a new **TRS** window.

You can copy, or clone, a **TRS** window to create a duplicate of your current analysis. Cloning allows you to fork your analysis from any point in the timeline without losing your prior work.

When you clone a **TRS** window, you choose the point in the timeline that you want included in the new window. The operations up to, and including, the selected operation are copied to the timeline in the new window. Any operations in the timeline that are after the selected operation are not included in the cloned window.

To clone the **TRS** window:

1. In the timeline of the **TRS** window that you want to clone, click the operation that you want included as the last operation in the new window.

The operations up to, and including, the selected operation are copied to the timeline in the new window. Any operations in the timeline that are after the selected operation are not included in the cloned window.

TRS 1: pub.demo.retail.item - Profit, Sales, & Cost by Date

Analyze Query View Visualize Develop Export

5 + - < > Compute profit as

6 + - < > Tabulation

Drag and drop columns to make a tabulation.

Show with column name More Options

Worksheet Columns: Transaction, Account, Store, Item SKU, Units, Sales, Cost

Row Breaks: Date

Column Breaks: sum of Total Sales, sum of Total Cost, sum of Profit

7 < > New operation

Profit, Sales, & Cost by Date

Cols 1 fixed and 2 to 4 of 4, Rows 1 to 4 of 4

Date	Sales	Cost	Profit
date	total_sales_date	total_cost_date	total_profit_date
	\$37.63	\$20.70	\$16.93
05/19/12	\$0.34	\$0.13	\$0.21
05/22/12	\$16.60	\$8.74	\$7.86
06/03/12	\$14.59	\$7.92	\$6.67
06/19/12	\$6.10	\$3.91	\$2.19

2. Optionally, choose the view in which to display the data.

For more information, see [View bar](#) on page 292.

3. In the View bar, click the **Clone view** () icon. The 1010data Insights Platform clones the **TRS** window and displays the new window with the analysis pane hidden.

Clone of TRS 1

Profit, Sales, & Cost by Date

Cols 1 fixed and 2 to 4 of 4, Rows 1 to 4 of 4

Date date	Sales total_sales_date	Cost total_cost_date	Profit total_profit_date
	\$37.63	\$20.70	\$16.93
05/19/12	\$0.34	\$0.13	\$0.21
05/22/12	\$16.60	\$8.74	\$7.86
06/03/12	\$14.59	\$7.92	\$6.67
06/19/12	\$6.10	\$3.91	\$2.19

Table Editor

The table editor allows you to create and edit tables in the 1010data Insights Platform interface.

Select a table to load data for editing, or select a folder or table to save to or replace.

(Browse) Open temp copy in TRS Folder: Name: Run query Load table Open saved table in TRS Title: Save

	c1	B
1	1	
2		

Object Browser (Browse)

The object browser provides a hierarchy of folders and objects to which you have access. Use the object browser to select a folder, query, or table.

Open as temp copy in TRS

Open an unsaved table or a table with unsaved changes in a TRS window.

Note: This feature is not yet available.

Folder

This field displays the folder path where you are saving the new or edited table. You can enter a table path instead of navigating with the object browser.

Name

This field displays the table name. If you are creating a new table, you can enter a table name in this field unless you are saving to the My Data folder.

Run query

This button runs an existing query in order to load and edit or copy the results.

Load table

This button loads an existing table or query results.

Open saved table in TRS

This button opens the saved version of the table in a **TRS** window. This option is not available for new tables.

Table title

This field displays the title of the table. You can enter a new title here.

Save/Replace

This button changes depending on the action. You can save a new table or replace an existing table with a table containing your edits.

Table pane

The table pane is the bottom portion of the **Table Editor** window and contains the new or existing table. When creating a new table, you can manually enter data into the cells or paste a copy of another table. For example, you could paste in a table from Excel.

Create a table

The Table Editor allows you to create new tables.

You can create new tables by manually entering data or by pasting copied data from an existing spreadsheet.

To create a new table with the Table Editor.

1. In the object browser, navigate to a folder that you have write access to. You can also enter the folder path in the **Folder** field.

2. Enter a table name in the **Table** field.

If you don't enter a table name, the Insights Platform automatically assigns one. You cannot enter a table name if you are saving to the **My Data** folder.

3. Enter a table title.

This is optional, but if you do not enter a table title, the entry for your new table in Object Manager will be blank.

The screenshot shows the 'Table Editor' window. At the top, it says 'Folder uploads exists and you can save to it. A table name will be assigned automatically.' Below this, there is a 'My Data' folder selected in an orange box. To its right is a button 'Open temp copy in TRS'. Further right, the 'Folder:' field contains 'uploads' and the 'Name:' field is empty. Below these, there are buttons for 'Run query', 'Load table', and 'Open saved table in TRS'. The 'Title:' field contains 'New Table' and there is a 'Save' button on the far right.

4. In the table pane, add data to the new table.

There are two ways to add data to a new table. You can manually enter the table data or paste in data that you've copied from a spreadsheet.

5. Right click on the table and select **Show Column Info**.

The Insights Platform reveals the column information fields.

	Employee ID	First Name	Last Name
Name	id	first	last
Label	Employee ID	First Name	Last Name
Type	int	text	text
Format			
1	43	Tanitansy	Cathrae

6. Enter the column information.

For more information, see [Columns](#) on page 48.

7. To save the new table, do one of the following:
 - Click **Save**.
 - Right-click and select **Save table**.
8. If you want to view the saved table in TRS, click **Open saved table in TRS**. The Insights Platform opens a **TRS** window.

Edit a table

The Table Editor allows you to edit existing tables that you have write access to.

To edit an existing table:

1. In the object browser, navigate to the table that you want to edit. The Table Editor populates the **Folder**, **Table**, and **Table title** fields and loads the table into the table pane.

2. Click **Load**.

All of the table data is downloaded to your browser, so there is a limit on how large a table you can edit with the Table Editor. You can't edit a table larger than 100,000 cells and the table will not load. A larger table may impact performance.

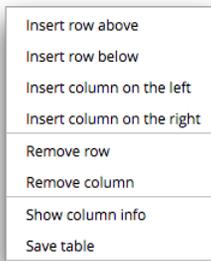
Note: If you open the Table Editor from the Object Manager, the table is loaded automatically.

The 1010data Insights Platform loads the table.

	Store	Division	Subdivision	zip	Division Description	Sub Division Description	Manager	Selling Sq Ft	Format	Cit
1	1	4	5	16063	West	Mountain Pass	Oswaldo Denton	12,529	Classic	Burbank
2	2	3	1	24637	East	Outer Banks	Barbie Turvey	9,408	Grocery Fuel Super Center	Mobile
3	3	1	8	33541	North	City Center	Arletha Blomberg	10,073	City Tenant	Detroit
4	4	1	10	4024	North	Unified Springfields	Kam Samayoa	9,503	Classic	Bridgepor
5	6	3	1	4629	East	Outer Banks	Spencer Kinzel	6,955	Shopping Center	Columbu:

3. Edit the table.

Right-clicking on the table reveals more table editing options including inserting new rows and columns, removing rows and columns, and displaying or hiding the column information.



4. Click **Replace**.

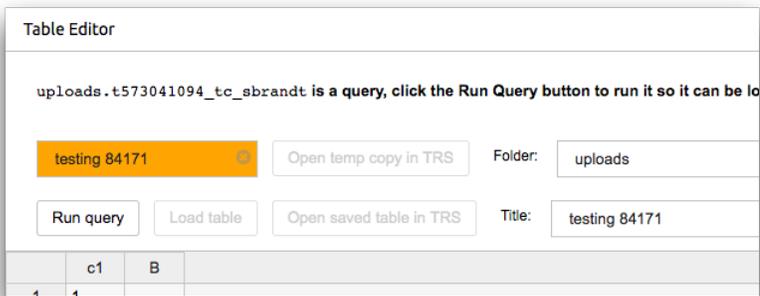
The Insights Platform treats the edited table as a copy of the original table, and it replaces the original table with the edited copy.

View Query Results

The table editor allows you to view query results.

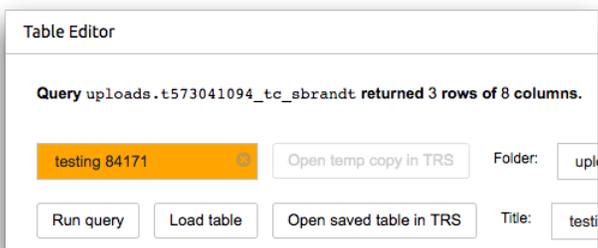
To view query results:

1. In the Object Browser, navigate to the query.



2. Click **Run query**.

The 1010data Insights Platform runs the query so that you can load the results into the table editor.



3. Click **Load table**.

The Insights Platform loads the results into the table pane. The results can be modified, but the modifications cannot be saved.

Uploading tables

You can create a new table by uploading your own data to the 1010data Insights Platform.

To upload a table in the Insights Platform web interface, data must be in a supported file format. Certain restrictions apply to the type of files supported. These restrictions are based on the method used to upload the file and whether the file will be transferred to 1010data using an FTP server. For more information, see [Supported file formats](#) on page 333.

When you perform an upload, you have the option of choosing either a local file, a file in your 1010data FTP account, or a file from Amazon S3 buckets. To choose a file from your FTP account, you must first use a third-party FTP client to transfer the file to 1010data. For instructions, see [Transfer a file to your FTP account](#) on page 334. To choose files from Amazon S3 buckets, you must add S3 API keys. To do this, you must manually edit the `s3_keys` dictionary in your user preferences or invoke the `addkey` api call.

You must be authorized to add tables to a folder. Folders that you have permission to add tables to are marked with the **Uploader** (👤) or **Owner** (🔑) icons in the **Object Manager**. Once your table is uploaded into the Insights Platform, it will be visible to only those people to whom you give permission. If you need to grant permission to multiple users, consider creating a group.

Supported file formats

Before uploading a table using the 1010data Insights Platform web interface, data from your current system must be exported into a file format that the platform can read.

The supported file formats vary based on whether you perform an upload or an advanced upload. Regardless of the type of upload you perform, you can upload tables in comma-separated (`.csv`), tab-separated (`.tsv`), and text file (`.txt`) formats. However, to upload a compressed version of these files, the compressed file must first be transferred to your 1010data FTP account.

Note: Delimited files must contain only printable ASCII characters; UTF-8 strings with multi-byte characters are not allowed.

Compatible file formats include the following:

<code>.csv</code>	A comma-separated values file.
<code>.tsv</code>	A tab-separated values file.
<code>.txt</code>	A standard comma-, tab-, pipe-, or space-separated text file. For details about what the text file may contain, see Text file format on page 333.
<code>.zip, .7z, .rar, .gz</code>	A compressed file containing either a <code>.csv</code> , <code>.tsv</code> , or <code>.txt</code> file. Note: Compressed file formats are only supported when using the FTP option while performing an upload.

Note: Microsoft Excel spreadsheet file formats (e.g., `.xls` and `.xlsx`) are not supported for upload in the platform user interface. To upload an Excel spreadsheet file, save the file as a comma-, tab-, pipe-, or space-separated text file or use the [1010data Excel Add-in](#) to upload the file directly from Excel.

Text file format

Using the 1010data Insights Platform web interface, you can upload data in a comma-, tab-, pipe-, or space-separated text file.

You can use commas, tabs, spaces, or pipes as delimiters within a text file. If you use spaces, multiple contiguous spaces are considered to be the same as a single space. Use the same separator throughout the file.

Here is a sample file where values are separated by spaces:

```
10      abc      3.5
20      def      7.8
25      gh       5.1
```

Note: For large files, it is better to use commas, tabs, pipes, or *single* spaces to separate values. Using multiple spaces, as in the example above, makes the files much larger, and it will take that much longer for you to transmit the data.

Prerequisites

Certain prerequisites are required to transfer a file to 1010data through FTP or to use the advanced upload feature to upload your own data into the 1010data Insights Platform.

FTP

The following table lists the items you must have in place to use FTP when performing an upload.

1010data FTP account	<p>In order to use FTP to transfer files to 1010data, you must first contact 1010data Support to set up your 1010data FTP account. This FTP account provides the access needed to transfer a file through FTP to 1010data before uploading it into the 1010data Insights Platform. Your 1010data FTP account uses the same username and password as your Insights Platform account.</p> <p>Note: If you change your Insights Platform account password, your 1010data FTP account password is automatically updated. However, this change may not take place immediately and can take up to one hour to complete. Continue to use your old platform account password for your FTP account until the change has taken place.</p>
Third-party FTP client	<p>If you intend to use FTP to transfer files to 1010data, you need a third-party FTP client installed on your computer.</p> <p>There are many free FTP clients available. While 1010data does not recommend or support a specific FTP client, many of our upload users choose to use a free, open-source FTP client called FileZilla.</p>

Advanced Uploader

You must have the following in place to perform an upload using the advanced upload feature:

Insights Platform account with API access	<p>To use the Advanced Uploader, you need to have API access added to your Insights Platform account.</p> <p>If you are not sure whether you have API access, contact 1010data Support. For instructions, see Submit a support request on page 29.</p>
--	--

Transfer a file to your FTP account

Using a third-party FTP client, you can transfer a file to your 1010data FTP account.

Before you begin, make sure you have the following items in place:

- 1010data Insights Platform account with API access
- 1010data FTP account
- Third-party FTP client

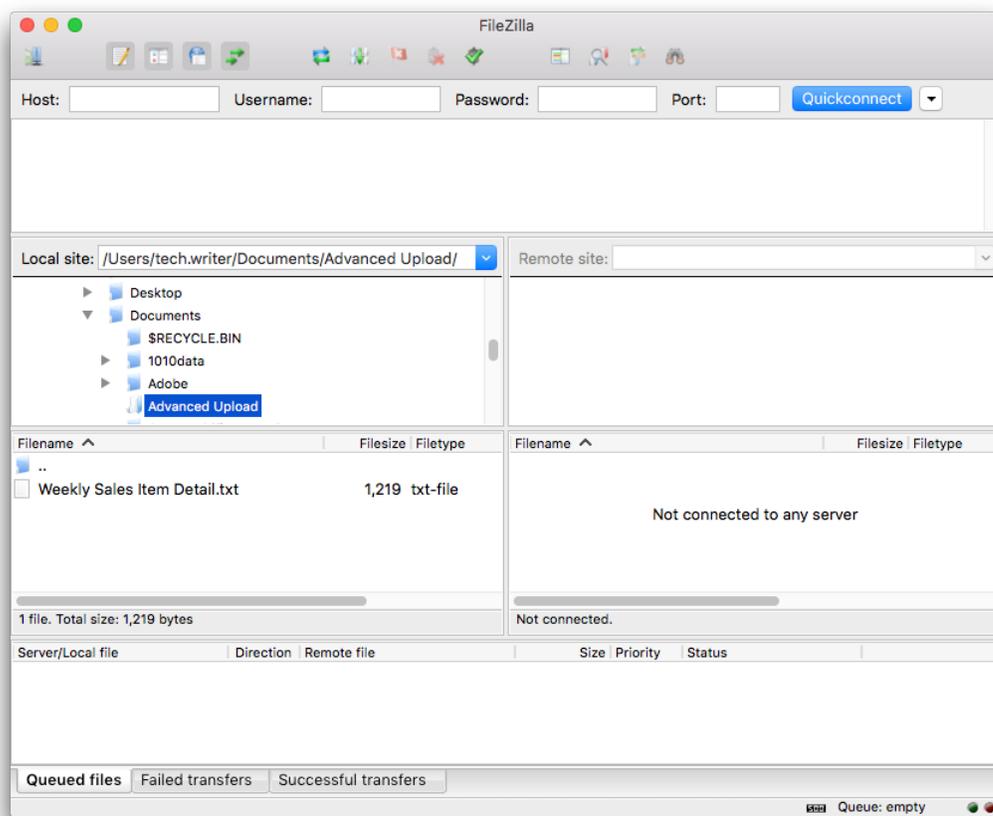
For more information, see [Prerequisites](#) on page 334.

Files transferred to 1010data using FTP are stored in your 1010data FTP account. When you perform an upload, you have the option of choosing either a local file or a file in your 1010data FTP account.

To transfer a file to your 1010data FTP account:

1. Download and install a third-party FTP client.

There are many free FTP clients available. While 1010data does not recommend or support a specific FTP client, many upload feature users choose to use a free, open-source FTP client called [FileZilla](#). For this reason, the FileZilla interface, shown below, is used for illustrative purposes.



2. Configure your third-party FTP client to connect to the 1010data FTP server.

Regardless of the FTP client you use, all FTP clients require and display the same basic information. The following connection information is required.

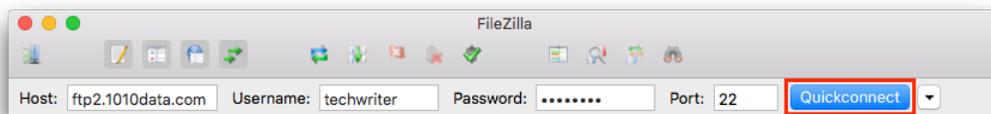
FTP connection information	Description	Value
Host	The URL of the 1010data FTP server. Note: Some 1010data customers may use a different FTP server. Contact 1010data Support if you cannot	ftp2.1010data.com

FTP connection information	Description	Value
	connect to the default server.	
Username	Your 1010data FTP account username. Your 1010data FTP account uses the same username as your 1010data Insights Platform account.	<i>Your 1010data Insights Platform account username</i>
Password	Your 1010data FTP account password. Your 1010data FTP account uses the same password as your 1010data Insights Platform account. Note: If you change your 1010data Insights Platform password, your FTP account password is automatically updated. However, this change may not take place immediately and can take up to one hour to complete. Continue to use your old 1010data Insights Platform account password for your FTP account until the change has taken place.	<i>Your 1010data Insights Platform account password</i>
Port	The TCP/IP port your computer uses to communicate with the 1010data FTP server. Note: 1010data uses the secure FTP protocol.	22

Note: The field labels above may differ slightly from those displayed in other third-party FTP client software.

3. Connect to the 1010data FTP server.

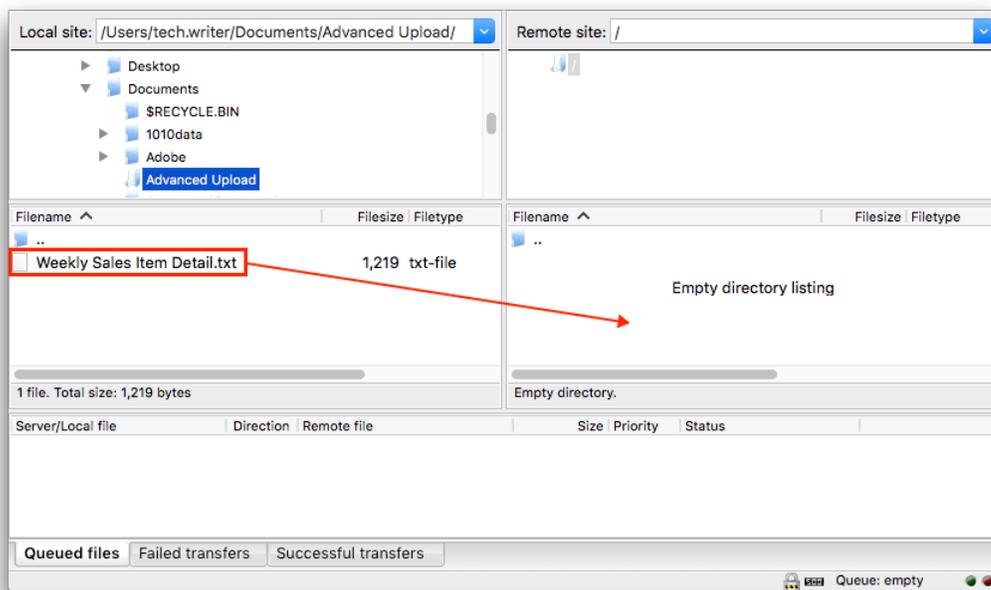
In FileZilla, this is accomplished by clicking **Quickconnect**.



Note: If you experience problems connecting to the 1010data FTP server, contact your IT department to ensure that FTP file transfers are not blocked by your corporate firewall or proxy server.

4. Locate and transfer the desired file to your 1010data FTP account.

Most FTP clients look and operate similarly. In FileZilla, the left side of the screen displays local folders and files. The right side displays your 1010data FTP account folder on the server. To transfer a file to your 1010data FTP account, simply drag the file from the left pane to the right pane.



The FTP client transfers the file to your 1010data FTP account.

Once the file is transferred to your 1010data FTP account, you can use the upload feature to upload it into the 1010data Insights Platform.

Upload a table (simple mode)

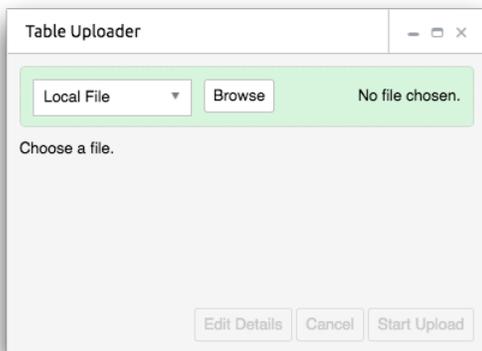
Quickly and easily upload a table directly into your **My Data** folder.

Before you begin, ensure you have completed the necessary prerequisites and that the table is in a supported file format. For more information, see [Prerequisites](#) on page 334 and [Supported file formats](#) on page 333.

The simple mode of the Table Uploader is one of the quickest and easiest ways of uploading a table into the 1010data Insights Platform. In the simple mode, tables are saved into your **My Data** folder and are automatically configured based on the data in the file.

To upload a table using the simple mode:

1. In the workspace menu, select **Tools > Table Uploader**.
The Insights Platform displays the **Table Uploader** window.

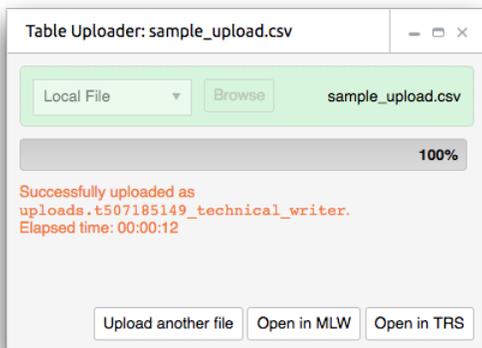


2. Depending where the file you want to upload is located, do one of the following:

Option	Description
<p>Local file</p>	<p>In the drop-down list, select Local File, click Browse, and then select the file containing the table you want to upload.</p> <p>Alternatively, drag the file from your operating system to the Insights Platform workspace.</p>
<p>1010data FTP account</p>	<p>In the drop-down list, select FTP. In the drop-down list that appears, select the file containing the table you want to upload.</p> <p>Note: You must have already transferred the file to your 1010data FTP account. For instructions, see Transfer a file to your FTP account on page 334.</p>

3. Click **Start Upload**.

The Table Uploader uploads the file and creates a new table. After the process is complete, a confirmation message is displayed.



Note: If you close the **Table Uploader** window before the upload process is complete, the Table Uploader will abort any uploads that are in progress.

After the table has been successfully uploaded, you can choose to upload another file, open the table in the Macro Language Workshop, or open the table in the Trillion-Row Spreadsheet.

Upload a table (detail mode)

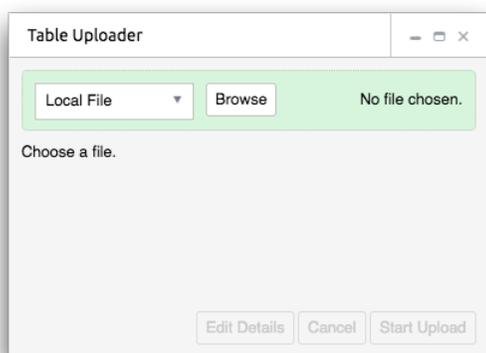
Upload a table to the 1010data Insights Platform.

Before you begin, ensure you have completed the necessary prerequisites and that the table is in a supported file format. For more information, see [Prerequisites](#) on page 334 and [Supported file formats](#) on page 333.

The detail mode of the Table Uploader is workflow-based and visually guides you through an iterative uploading process. As you progress through the steps of the upload, the platform displays a preview of the data so you can make any necessary adjustments along the way. Additionally, the detail mode provides the ability to edit the specification (`.spec`) file during the upload process.

To upload a table using the detail mode:

1. In the workspace menu, select **Tools > Table Uploader**.
The Insights Platform displays the **Table Uploader** window.



2. Depending where the file you want to upload is located, do one of the following:

Option	Description
Local file	In the drop-down list, select Local File , click Browse , and then select the file containing the table you want to upload. Alternatively, drag the file from your operating system to the Insights Platform workspace.
1010data FTP account	In the drop-down list, select FTP . In the drop-down list that appears, select the file containing the table you want to upload. Note: You must have already transferred the file to your 1010data FTP account. For instructions, see Transfer a file to your FTP account on page 334.

3. Click **Edit Details**.

The Table Uploader configures the settings for the table based on the data in the file and displays the **Structure** section. This section contains a preview of the table along with fields and options for defining the divisions between columns and rows in the data.

Table Uploader: sample_upload.csv

Structure **8 Columns, 35 Rows in sample** Raw View

We must identify the divisions between columns and rows in your data.

The *delimiter* separates fields in each row. The *end of record* character separates rows.

Delimiter:

End of Record:

Do your files begin with a preamble section before the data which should be ignored?

Skip Rows:

531	957	1	05/15/12	366	-1	-5	-1.84
532	478	1	05/15/12	98A	1	0.5	0.25
532	478	1	05/15/12	3B7	1	1.1	0.56
534	738	1	05/16/12	A96	2	6	2.9
534	738	1	05/16/12	65B	1	2.25	1.35
535	709	2	05/15/12	CB7	1	1.65	1.1
535	709	2	05/15/12	96A	1	1.1	1
535	709	2	05/15/12	969	1	1.1	1
536	748	2	05/17/12	3A4	3	1.02	0.39
536	748	2	05/17/12	366	1	5	1.8
537	523	3	05/15/12	CB7	1	1.65	1.1
537	523	3	05/15/12	A96	2	6	2.9
537	523	3	05/15/12	98A	1	0.5	0.25
537	523	3	05/15/12	96A	1	1.1	1
537	523	3	05/15/12	3B7	1	1.1	0.56
538	668	1	05/18/12	CB7	1	1.65	1.1
538	668	1	05/18/12	98A	4	2	1
538	668	1	05/18/12	96A	1	1.1	1
538	668	1	05/18/12	969	1	1.1	1
538	668	1	05/18/12	3B7	1	1.1	0.56
539	25	2	06/19/12	65B	1	5	3.35
540	361	2	06/19/12	3B7	1	1.1	0.56

4. As necessary, edit the settings in the **Structure** section.

For more information, see [Structure](#) on page 344.

5. Click **Headers**.

The Table Uploader displays the **Headers** section. This section allows you to define the names and labels of the columns in your data.

Table Uploader: sample_upload.csv

Structure **8 Columns, 35 Rows in sample**

Headers

1010data column *names* must begin with a letter, may contain only lowercase letters, digits, and underscores (`_`), and may not end in an underscore. Column *labels* may contain any character.

Here you can perform bulk changes to the names and labels of the columns in your data. We have included data from skipped rows and the leading rows of your data which look like they may contain usable column names or labels.

Name	Label	Skipped Row 1	Placeholders
transaction_id	Transaction ID	Transaction ID	c1
account	Account	Account	c2
store	Store	Store	c3
date	Date	Date	c4
item_sku	Item SKU	Item SKU	c5
units	Units	Units	c6
sales	Sales	Sales	c7
cost	Cost	Cost	c8

6. As necessary, edit the names and labels of columns in the **Headers** section.

For more information, see [Headers](#) on page 346.

7. Click **Fields**.

The Table Uploader displays the **Fields** section. This section contains a preview of the table and allows you to define the data type and display format of the columns in your data.

Table Uploader: sample_upload.csv

Structure ▾ 8 Columns, 35 Rows in sample Raw View ⌵

Headers ▾

Fields ▲

Select columns in the lower grid to adjust formatting as needed. The *raw view* shows data as it appears in the input file.

Destination ▾

Spec XML ▾

Load ▾

Transaction ID	Account	Store	Date	Item SKU	Units	Sales	Cost
531	957	1	05/15/12	366	-1	-5	-1.84
532	478	1	05/15/12	98A	1	0.5	0.25
532	478	1	05/15/12	3B7	1	1.1	0.56
534	738	1	05/16/12	A96	2	6	2.9
534	738	1	05/16/12	65B	1	2.25	1.35
535	709	2	05/15/12	CB7	1	1.65	1.1
535	709	2	05/15/12	96A	1	1.1	1
535	709	2	05/15/12	969	1	1.1	1
536	748	2	05/17/12	344	2	1.02	0.39

Column Settings:

Name	Label	Input Format	Display Format	Invalid / Empty
transaction_id	Transaction ID	Integer	Number	0 / 0
account	Account	Integer	Number	0 / 0
store	Store	Integer	Number	0 / 0
date	Date	Text	Text	0 / 0
item_sku	Item SKU	Text	Text	0 / 0
units	Units	Integer	Number	0 / 0
sales	Sales	Float	Number	0 / 0
cost	Cost	Float	Number	0 / 0

8. As necessary, edit the data type and display format of columns in the **Fields** section.

For more information, see [Fields](#) on page 348.

9. Click **Destination**.

The Table Uploader displays the **Destination** section. This section allows you to select the save location, describe the table, and define who can access the uploaded table.

The screenshot shows the 'Table Uploader: sample_upload.csv' window. On the left, a sidebar contains a menu with 'Destination' selected. The main area is titled 'Output Table Configuration' and contains several fields: 'Title' (with 'Table Title' entered), 'Path' (empty), 'Short Description' (empty), 'Long Description' (with a rich text editor containing 'Format' and bold/italic/underline buttons), and 'Users' (with 'Private' selected). A note below the Path field states: 'By default, uploads are assigned a unique filename.'

10. Complete the appropriate fields and options in the **Destination** section.

For more information, see [Destination](#) on page 350.

11. Click **Spec XML**.

The Table Uploader displays the **Spec XML** section. This section contains the automatically generated specification (.spec) file which defines the upload configuration for the table.

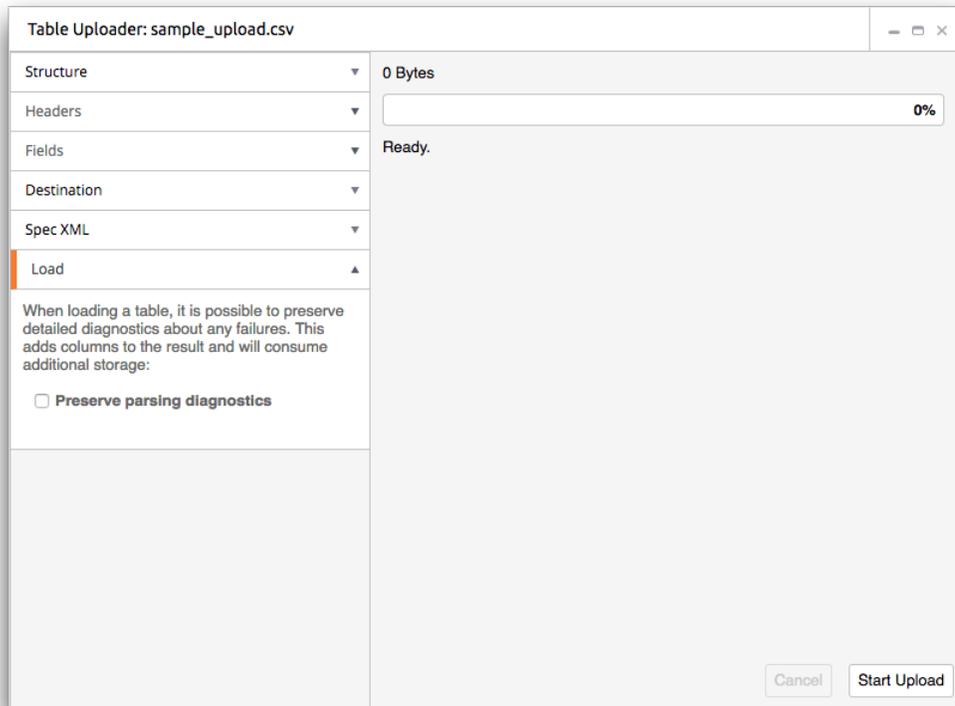
The screenshot shows the 'Table Uploader: sample_upload.csv' window with the 'Spec XML' section selected in the sidebar. The main area displays the generated XML specification code, numbered 1 through 31. The code defines two columns: 'transaction_id' (integer, width 3) and 'account' (integer, width 3). A note below the sidebar states: 'Shown to the right is a .spec file which can be used for automated loads with the current configuration. This tool analyzes the structure of a small sample of an uploaded dataset. If your dataset has a very large number of columns, it may be necessary to examine a larger sample to make sense of the data. Note: working with a larger sample will be slower.' A dropdown menu shows '32 Kilobytes'.

12. As appropriate, load, edit, or save the specification file for the table.

For more information, see [Spec XML](#) on page 352.

13. Click **Load**.

The Table Uploader displays the **Load** section. This section allows you to choose whether to preserve the parsing diagnostics of the upload process and to upload the input file.



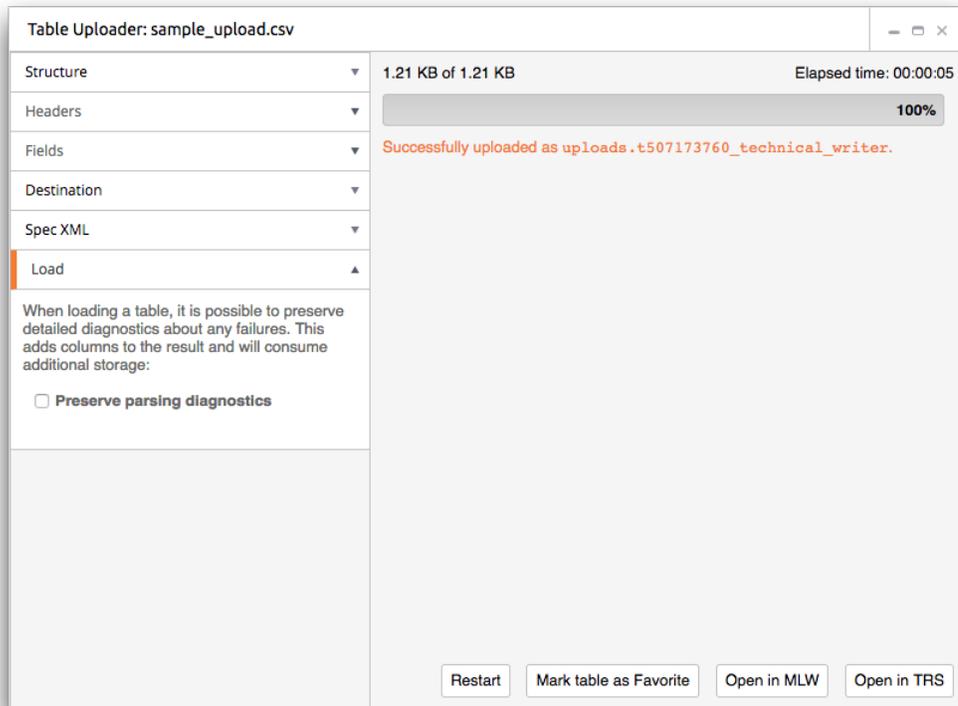
14. If desired, select the **Preserve parsing diagnostics** option.

This option preserves detailed diagnostics about any failures encountered during the upload process.

Note: The diagnostic information is included as additional columns in the resulting table and uses additional storage space.

15. Click **Start Upload**.

The Table Uploader uploads the file and creates a new table. After the process is complete, a confirmation message is displayed.



Note: If you close the **Table Uploader** window before the upload process is complete, the Table Uploader will abort any uploads that are in progress.

After the table is successfully uploaded, additional options are available in the right side of the window. For example, you can mark the table as a favorite, open the table in the Trillion-Row Spreadsheet or Macro Language Workshop, or choose to restart the upload process for the input file. For more information, see [Load](#) on page 354.

Structure

Descriptions of fields, options, and views in the **Structure** section of the Table Uploader.

The **Structure** section allows you to define the divisions between columns and rows in the data. Based on the file type and the data contained within, the Table Uploader automatically chooses the structure options. However, if the Table Uploader did not correctly interpret the data, you can change these settings.

The right side of the **Table Uploader** window displays a preview table of the data based on the settings in the left side of the window. The table preview view is shown in the image below on the left.

You can also choose to view the underlying data of the table in raw text format. The raw view displays the input file data in plain text format with delimiter and end of record characters. The raw view is shown in the image below on the right.

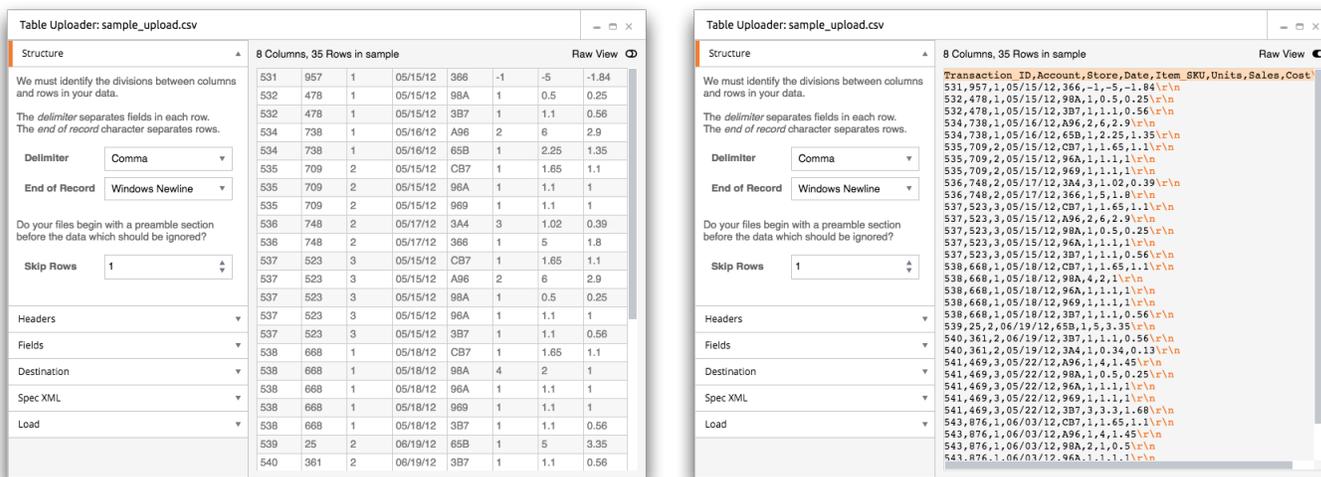


Figure 83: Structure section in the Table Uploader

Delimiter

The character in the source data used to separate the value in each column.

Comma

A comma (,) is used to separate the columns in each row of the table.

Tab

A tab stop is used to separate the columns in each row of the table.

Note: In the raw view, a tab stop is indicated by the (\t) character.

Space

A space is used to separate the columns in each row of the table.

Note: In the raw view, a space is indicated by the underscore (_) character.

Colon

A colon (:) is used to separate the columns in each row of the table.

Semicolon

A semicolon (;) is used to separate the columns in each row of the table.

Vertical Bar

A pipe (|) is used to separate the columns in each row of the table.

End of Record

The character or characters in the source data used to indicate the end of each row.

Windows Newline

The Carriage Return Line Feed (CRLF) row delimiter is used to indicate the end of the row. In the raw view, this is indicated by the (\r\n) characters.

Unix/OSX Newline

The New Line (NL) row delimiter is used to indicate the end of each row. In the raw view, this is indicated by the (\n) character.

MacOS 9 Newline

The Carriage Return (CR) row delimiter is used to indicate the end of each row. In the raw view, this is indicated by the (\r) character.

Skip Rows

The number of rows in the file that are skipped when the table is uploaded.

For example, the first row of a table usually contains column names and not data. To exclude the first row of the table in the upload, enter 1. When the file is uploaded, the Table Uploader will ignore the first row in the table and start uploading the data beginning with the second row.

Note: In the raw view, skipped rows are highlighted in orange.

Raw View

This switch controls whether the table view or the raw view is displayed in the right side of the window. The table view shows a preview of the data based on the settings in the left side of the window. The raw view displays the input file data in plain text format with delimiter and end of record characters.

Headers

Descriptions of fields and options in the **Headers** section of the Table Uploader.

In the **Headers** section, you define the names and labels of the columns in your data. You can change column names or labels individually or in bulk. If data from skipped and leading rows looks like usable column names or labels, it is displayed. This data can then be used to easily define all of the column names, labels, or both.

You can right-click a cell in the table on the right side of the **Table Uploader** window to access a menu of options. These options allow you to update column names and labels in bulk. For more information, see [Headers section menu options](#) on page 347.

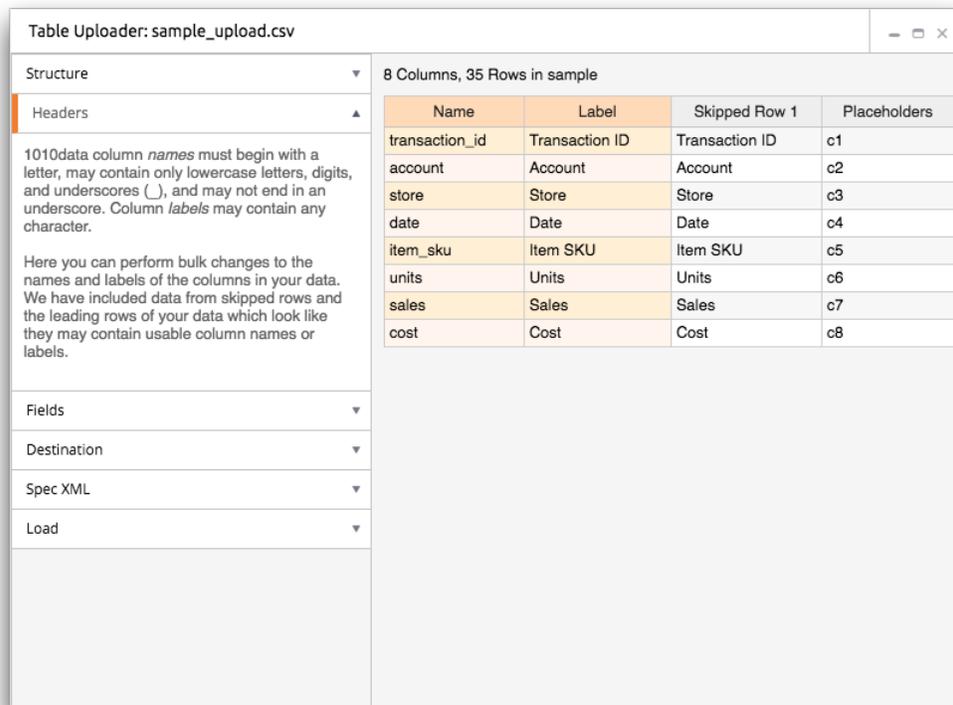


Figure 84: Headers section in the Table Uploader

Name

This column lists the name of each column in the table. The name provides a unique identifier for each column of a table in the 1010data Insights Platform.

To edit a column name, double-click the appropriate cell in the **Name** column and enter the desired text.

The name may only contain alphanumeric characters or underscores and must begin with an alphabetic character (e.g., `percent_total_sales`). It may not contain any spaces or other special characters.

Label

This column lists the label of each column in the table. The column label is displayed by default at the top of a column in the web interface.

To edit a column label, double-click the appropriate cell in the **Label** column and enter the desired text.

The label may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters (e.g., **Percentage of Total Sales (%)**).

Skipped Row *[N]*

This column lists the contents of a skipped row of data in the input file.

Note: *[N]* represents the row number of the skipped row. There can be as many **Skipped Row** columns as the number of rows skipped.

This column is displayed only when a row of data is skipped *and* the contents in the skipped row of data looks like usable column names or labels.

Data Row *[N]*

This column lists the contents of a row of data that immediately follows the skipped rows.

Note: *[N]* represents the first or second row that immediately follows the skipped rows. There can be up to two **Data Row** columns.

This column is displayed only when the contents in the first or second rows of data that immediately follow the skipped rows looks like usable column names or labels.

Placeholders

This column lists the default column name placeholders. When column names are not defined, the Insights Platform automatically names the columns `c1`, `c2`, `c3`, and so forth.

Headers section menu options

The **Headers** section menu provides options for updating column names and labels in bulk.

The menu in the **Headers** section is accessed by right-clicking a cell in the table on the right side of the **Table Uploader** window.

8 Columns, 34 Rows in sample

Name	Label	Skipped Row 1	Placeholders
transaction_id	Transaction ID	Transaction ID	c1
account	Account	A	Use this column as Names
store	Store	S	Use this column as Labels
date	Date	D	Convert Selection to UPPER CASE
item_sku	Item SKU	I	Convert Selection to lower case
units	Units	U	Convert Selection to Proper Case
sales	Sales	Sales	c7
cost	Cost	Cost	c8

Figure 85: Menu options in the Headers section of the Table Uploader

Use this column as Names

Assigns the contents of the column in which you right-clicked as the column names in the table.

Note: Column data used as names is automatically updated to follow proper naming conventions.

Use this column as Labels

Assigns the contents of the column in which you right-clicked as the column labels in the table.

Convert Selection to UPPER CASE

Changes all letters in the selected cells to capital letters.

Convert Selection to lower case

Changes all letters in the selected cells to small letters.

Convert Selection to Proper Case

Changes the content in the selected cells as follows:

- Removes leading and trailing spaces
- Reduces multiple spaces between characters to a single space
- Changes the first letter following a space to a capital letter
- Changes all other letters to small letters

Fields

Descriptions of fields, options, and views in the **Fields** section of the Table Uploader.

In the **Fields** section, you can define the data type and display format of the columns in your data. In addition, this section can help you quickly locate any problems or issues with the data; when detected, the Table Uploader informs you of invalid and empty cells.

The **Fields** section of the **Table Uploader** window contains two grids. The upper grid shows a preview of the data based on the settings in the lower grid. The lower grid contains fields with drop-down menus which allow you select the data type and display format for a column.

The upper grid has two possible views controlled by the **Raw View** switch. The image below on the left shows the formatted table view in the upper grid while the image on the right shows the unformatted (raw) table view.

When you click a cell in either of the two grids, the Table Uploader highlights the column in the upper grid and the corresponding row in the lower grid. This can help you see how the settings in the lower grid affect the data in the upper grid.

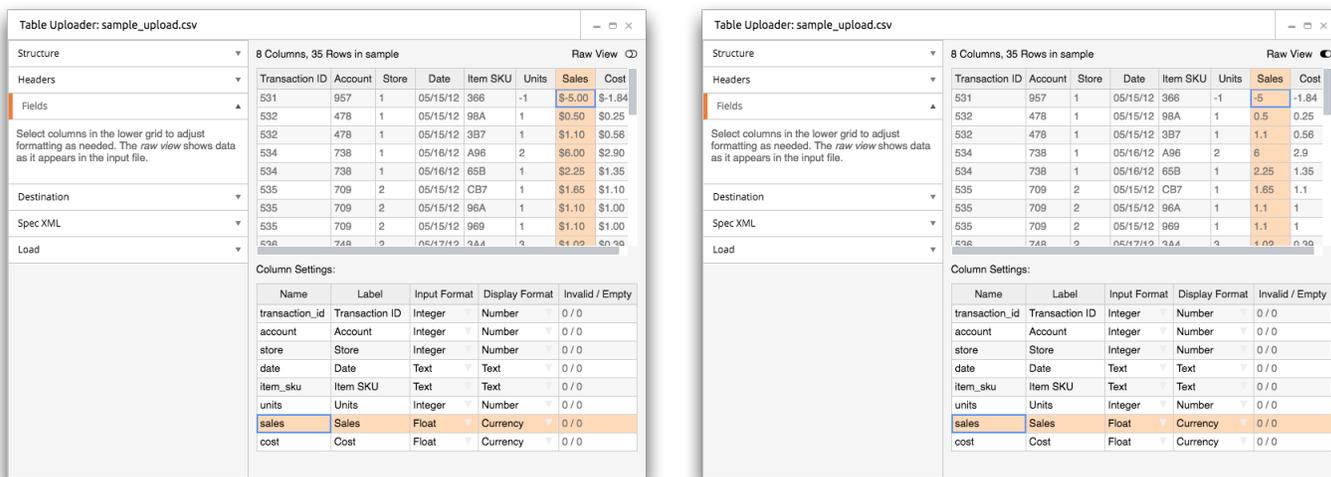


Figure 86: Fields section in the Table Uploader

Raw View

This switch controls whether the formatted table view or the unformatted (raw) table view is displayed in the upper grid on the right side of the window.

The formatted table view shows a preview of the data based on the format settings in the lower grid. The unformatted (raw) table view shows a preview of the data as it appears in the input file. The unformatted (raw) table view ignores any selected format options chosen in the lower grid.

For example, in the images above, the display format for the `sales` and `cost` columns is set to **Currency**. In the image on the left, this is reflected in the formatted table view in the upper grid. In the image on the right, the display format settings are ignored as shown in the unformatted (raw) table view in the upper grid.

Name

This column lists the name of each column in the table. The name provides a unique identifier for each column of a table in the 1010data Insights Platform.

To edit a column name, double-click the appropriate cell in the **Name** column and enter the desired text.

The name may only contain alphanumeric characters or underscores and must begin with an alphabetic character (e.g., `percent_total_sales`). It may not contain any spaces or other special characters.

Label

This column lists the label of each column in the table. The column label is displayed by default at the top of a column in the web interface.

To edit a column label, double-click the appropriate cell in the **Label** column and enter the desired text.

The label may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters (e.g., **Percentage of Total Sales (%)**).

Input Format

Also referred to as data types, this field is used to set the kind of data contained in the column. The most common options include the following:

Text

The column contains text, as opposed to numbers.

The text data type is used to represent values that contain uppercase and lowercase letters, numbers, spaces, and symbols. For more information, see [Text](#) on page 55.

Float

The column contains numbers with place values to the right of a decimal point.

Also referred to as decimal, float is used to represent 64-bit floating point values. For more information, see [Decimal](#) on page 54.

Integer

The column contains only whole numbers.

The integer data type is used to represent whole numbers that can be stored within 32-bits. For more information, see [Integer](#) on page 54.

Big Integer

The column contains 64-bit integer data representing very large and very small whole numbers.

The big integer data type is used to represent whole numbers that are outside the range of the integer data type and can be stored within 64 bits. For more information, see [Big integer](#) on page 54.

Note: In addition to the types explained above, there are also numerous options for date and time data. These data formats apply Insight Platform date formatting measures to the date data in the source file. Dates in the source file can be formatted with dashes ('-'), forward-slashes ('/'), and spaces (' '). Data uploaded with just date information are created as *integers* in the Insights Platform, with time-handling formats applied as specified in the **Input Format** drop-down menu. Data with date and time data are created as *floats* in the platform, with the applicable time-handling formats applied as specified in the **Input Format** drop-down menu.

To edit the data type of a column, click the drop-down arrow (▼) within the appropriate cell in the **Input Format** column and select an option from the menu.

Display Format

Defines how the data is displayed in the web interface. For example, you can choose to exclude commas in numbers or change the way a date is formatted. Choices made in the **Display Format** do not change the type of data in the column, only the way the data is displayed. For a list of options and descriptions, see [Display formats](#) on page 56.

To edit the display format of a column, click the drop-down arrow (▼) within the appropriate cell in the **Display Format** column and select an option from the menu.

Invalid/Empty

Indicates the number of invalid or empty cells in the table. This can help you locate and correct issues with the input file. It can also help alert you to incorrectly set data types.

Destination

Descriptions of fields and options in the **Destination** section of the Table Uploader.

In the **Destination** section, you have the option of selecting a folder to upload the table into or to select an existing table to replace. This section also provides fields which allow you to describe the table and to define the users who can access the uploaded table.

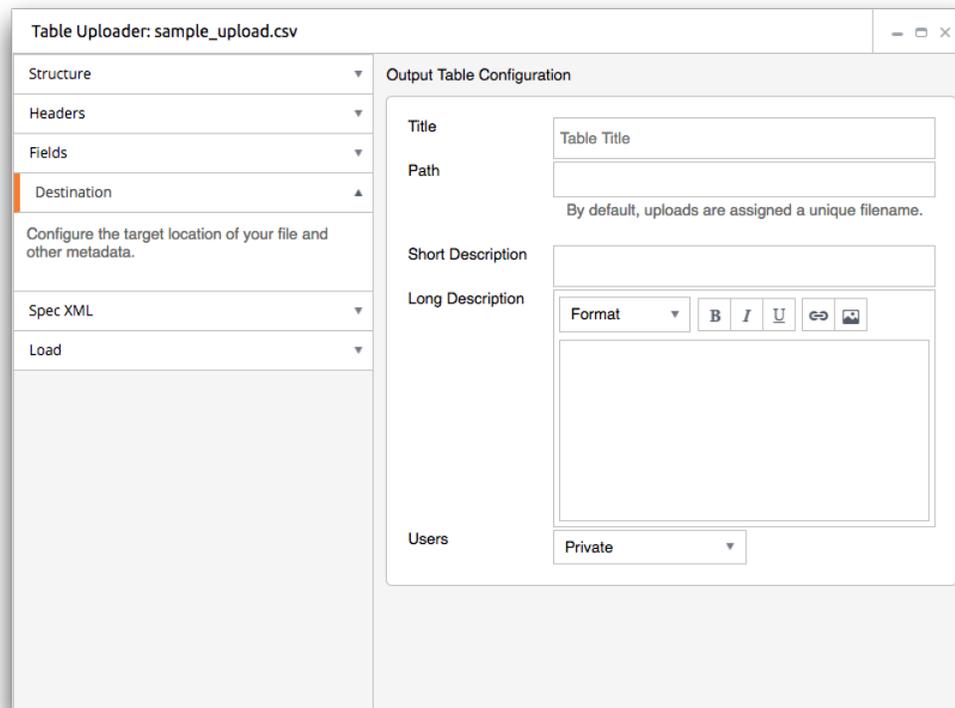


Figure 87: Destination section in the Table Uploader

Title

This field is used to set the title of your table.

The title is used to help describe the contents of a table (e.g., *Sales Detail by Customer*). This is used for recognition of the table and does not necessarily have to be unique.

The title may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters.

Path

This field is used to select the location to where the table is uploaded.

When this field is clicked, the Table Uploader displays an object browser dialog. The object browser allows you to select a folder to upload the table into or to select an existing table to replace.

If an existing table is selected, this field displays the path and the name of the selected table. If a folder is selected, this field contains the path to the selected folder and the **Name** field appears.

Name

This field is used to set the name of your table. This is the unique name of the table used for identification in the 1010data Insights Platform.

The name must begin with a letter and can only contain numbers, letters, and underscores. It cannot contain any spaces or other special characters. If you leave this field blank, a system-generated name will be used (e.g., `t662528755_yourusername`).

Note: This field is not visible when an existing table or the **My Data** folder is selected in the object browser. If you are saving to the **My Data** folder, a system-generated name is used.

Short Description

This field is used to provide a brief description of the data contained in your table.

Long Description

This field is used to provide a detailed explanation of the data contained in your table.

You can customize the appearance of the description text by using the provided rich text editor controls. These controls are located above the **Long Description** field and include the following:

Format

This drop-down list contains pre-formatted text styles. For example, you can choose different heading levels to organize the description text.

Bold

This icon applies bold formatting to the selected text.

Italic

This icon makes the selected text italic.

Underline

This icon allows you to underline the selected text.

Insert hyperlink

This icon opens the **Insert hyperlink** dialog which allows you to create a link in the description text.

Insert image

This icon opens the **Insert image** dialog which allows you to add an image in the description text.

Users

This field is used to define the users who have permission to view the table. You can select from the following options:

Private

When this option is selected, only you will have access to the table.

This is the default setting.

Inherit

When this option is selected, users who have access to the parent folder (the folder the table is being uploaded to) will have access to the table.

Note: Although this option is available, you cannot use the inherit permissions feature for tables saved in your **My Data** folder. To share a table saved in your **My Data** folder, you must explicitly give permission to each user with whom you want to share the table.

List

When this option is selected, the Table Uploader displays a text field below the drop-down list.

Enter the names of the 1010data Insight Platform users or groups, separated with a space, with whom you want to share the table.

Spec XML

Descriptions of fields and options in the **Spec XML** section of the Table Uploader.

When uploading a table, the Table Uploader analyzes the structure of a small sample of the input file and automatically generates a specification (`.spec`) file. The specification file defines the upload configuration for the table and is visible in the **Spec XML** section.

The code editor on the right side of the window displays the specification file in XML format. As you customize the upload configuration in the various sections of the Table Uploader, the specification file is updated to match. If necessary, you can also make changes directly to the specification file in the code editor before uploading the table.

You can save a specification file so it can be used again. Saving the specification file is helpful if you regularly upload files with the same data format.

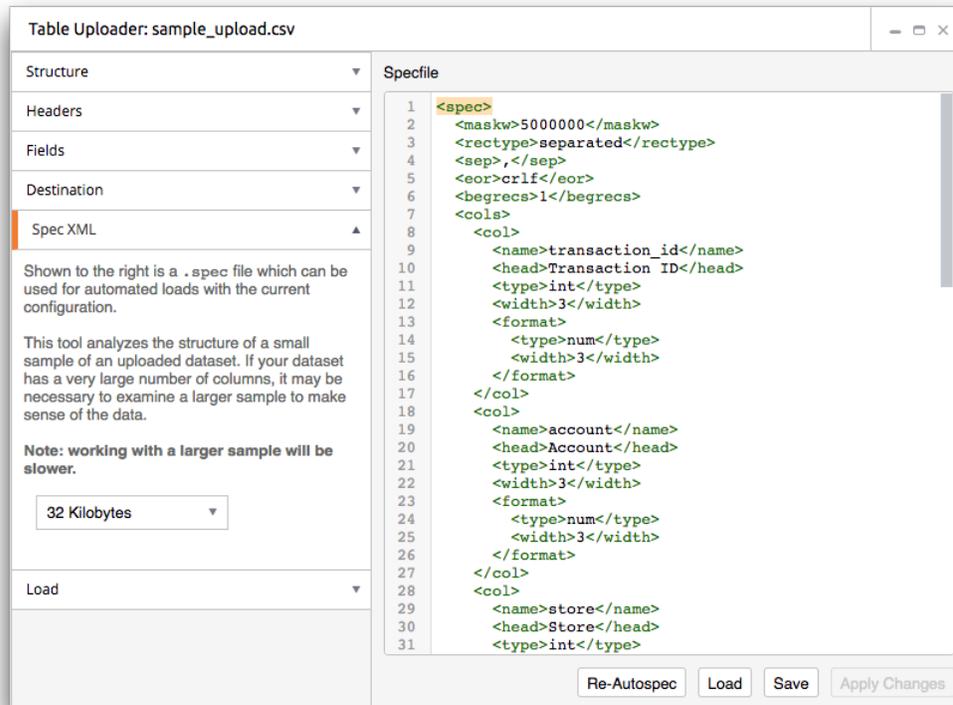


Figure 88: Spec XML section in the Table Uploader

Sample size

Located on the left side of the window, this drop-down list allows you to adjust the sample size of the input data used to generate the specification file. If your dataset has a very large number of columns, it may be necessary to examine a larger sample to make sense of the data.

To change the sample size, click the **Sample size** drop-down list and chose a size from the list of options.

Note: Working with a larger sample will be slower.

Specfile

Located on the right side of the window, the **Specfile** code editor displays the specification (.spec) file for the table.

As necessary, you can edit the specification file directly in the code editor.

Re-Autospec

This button regenerates the specification file based on the data in the input file.

Click this button to remove any edits made to the a specification file either in the code editor or in the various sections of the Table Uploader. This button can also be used to change a loaded specification file to an automatically generated version.

Load

This button allows you to choose a previously saved specification file to configure the settings for the table.

This is helpful if you regularly add tables to the Insights Platform with the same data format. For example, if you upload weekly sales data, the specification file can automatically complete all of the configuration settings of the table for you.

Save

This button saves a copy of the specification file displayed in the code editor to your computer.

You can save your customized upload configuration in a specification file and use it to save time when uploading tables with the same data format in the future.

Apply Changes

This button applies the changes made to the specification file in the code editor.

For changes to take affect, you must apply them to the input file before uploading the table.

Load

Descriptions of options and views in the **Load** section of the Table Uploader.

In the **Load** section, you can choose to preserve the parsing diagnostics of the upload process and upload the input file.

After the table is successfully uploaded, additional options are available in the right side of the window. For example, you can mark the table as a favorite, open the table in the Trillion-Row Spreadsheet or Macro Language Workshop, or choose to restart the upload process for the input file.

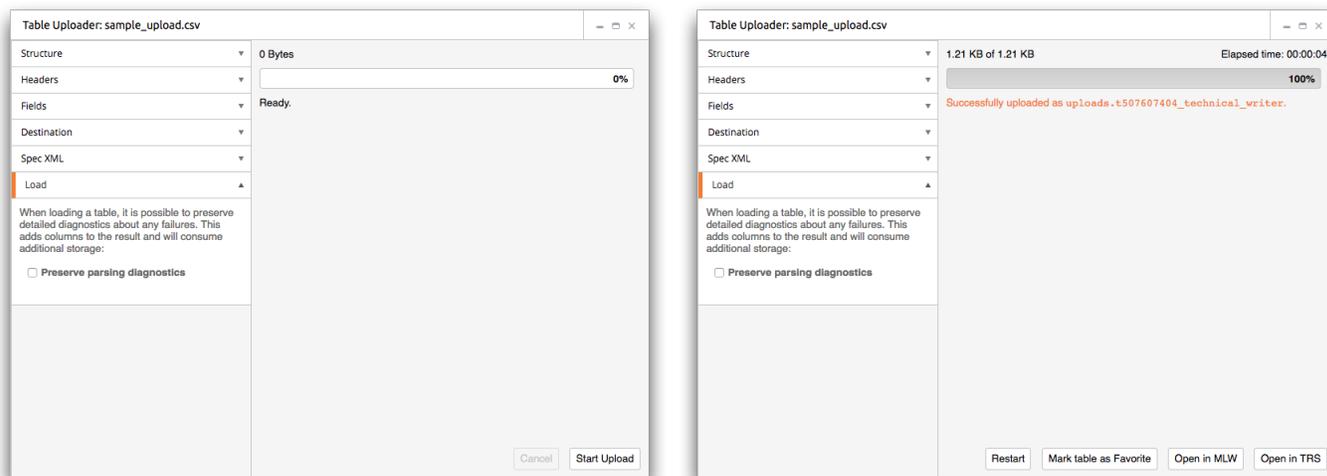


Figure 89: Load section in the Table Uploader

Preserve parsing diagnostics

Located on the left side of the window, this option preserves detailed diagnostics about any failures encountered during the upload process.

The diagnostic information is included as additional columns in the resulting table and uses additional storage space.

Cancel

Allows you to cancel an upload in progress.

Start Upload

Submits the file and starts the upload.

As the input file is uploaded, the status bar (on the right side of the window) provides feedback on the upload progress.

Restart

Displays the initial view of the **Load** section in the window.

This allows you to restart the upload.

Mark table as Favorite

Marks the table as a favorite in the Object Manager and adds the table to the **Favorites** tab in the Dashboard.

This allows for quick and easy access of the uploaded table.

Open in MLW

Opens the uploaded table in the Macro Language Workshop.

Open in TRS

Opens the uploaded table in the Trillion-Row Spreadsheet.

Upload a table (advanced upload)

Using the Advanced Uploader, you can upload a table to the 1010data Insights Platform with a finer grain of control.

Before you begin, ensure you have completed the necessary prerequisites and that the table is in a supported file format. For more information, see [Prerequisites](#) on page 334 and [Supported file formats](#) on page 333.

The Advanced Uploader provides many options for fine-grained control over the data you are uploading. Additionally, you can save your customized upload configuration in a specification file and use it to save time when uploading tables with the same data format.

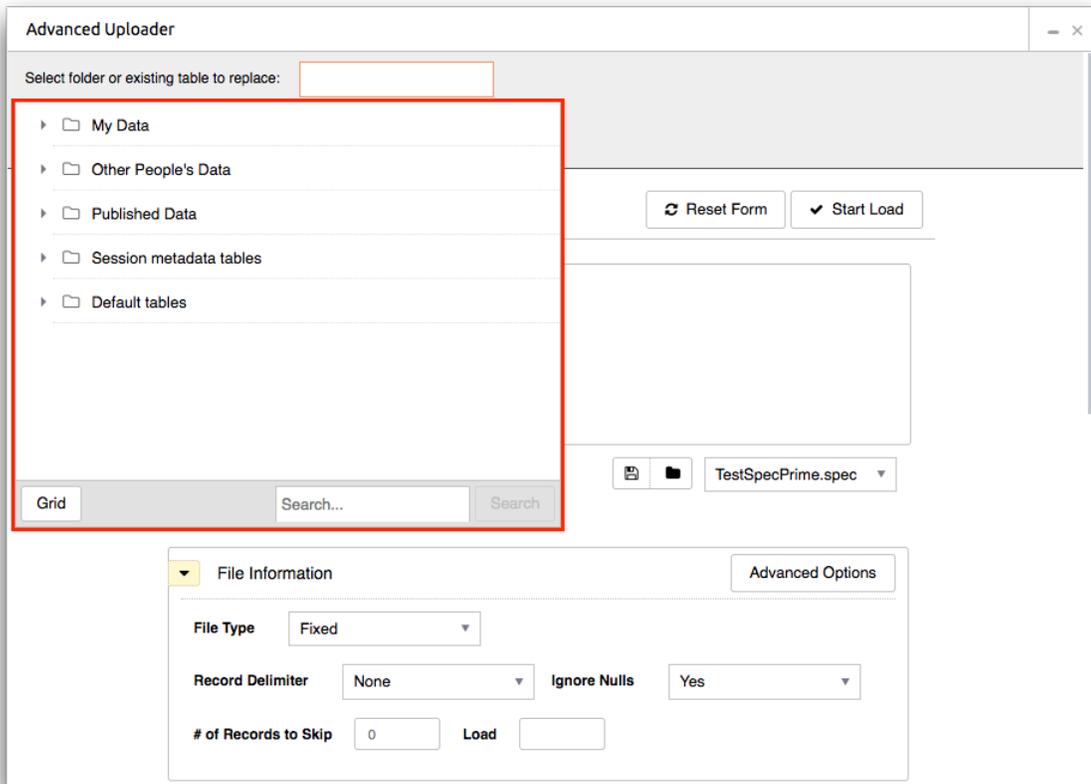
To upload a table using the Advanced Uploader:

1. In the workspace menu, select **Tools > Advanced Uploader**.

The Insights Platform displays the **Advanced Uploader** window. If your user ID has FTP access, the **FTP** option is selected by default and any files transferred to your 1010data FTP account are listed.

The screenshot shows the 'Advanced Uploader' window. At the top, there is a header with the title 'Advanced Uploader' and a close button. Below the header, there is a section for selecting a folder or existing table to replace, with a text input field. Underneath, there are fields for 'Folder:' (containing 'uploads') and 'Name:'. The main area features a 'Source' section with two buttons: 'FTP' (highlighted in orange) and 'Local File'. To the right of these are 'Reset Form' and 'Start Load' buttons. A large text area displays a list of files: 'Product Master.zip', 'Sales Item Detail.txt', and 'Store Master.csv'. Below this list is an 'Auto Detect File Specs' button and a file selection area showing 'TestSpecPrime.spec'. At the bottom, there is a 'File Information' section with a dropdown menu, and an 'Advanced Options' button. The 'File Information' section includes a 'File Type' dropdown set to 'Fixed', a 'Record Delimiter' dropdown set to 'None', an 'Ignore Nulls' dropdown set to 'Yes', and a '# of Records to Skip' input field with the value '0' and a 'Load' button.

2. Near the top of the window, click the **Select folder or existing table to replace** field. The Advanced Uploader displays the object browser in the window.



3. Do one of the following:

- Select the folder where you want to upload your table.
- Select an existing table you want to replace.

The Advanced Uploader completes the **Folder** field.

4. Depending where the file you want to upload is located, do one of the following:

Option	Description
1010data FTP account	In the field below the Source options, click the file containing the table you want to upload. Note: You must have already transferred the file to your 1010data FTP account. For instructions, see Transfer a file to your FTP account on page 334.
Local file	Click Local File , then click the Choose Files button and select the file containing the table you want to upload.

5. In the **Name** field near the top of the window, enter the table name.

The table name must begin with a letter and can only contain numbers, letters, and underscores. It cannot contain any spaces or other special characters. If you leave this field blank, a system-generated name will be used (e.g., t662528755_yourusername). The path to the parent folder will be automatically prepended to the **Full Path**.

Note: If you are uploading to the **My Data** folder, the path `uploads` will be automatically prepended to the full path in the **Folder** field, and a system-generated table name will be used (e.g., `uploads.t662528755_yourusername`).

6. Do one of the following:

- Click **Auto Detect File Specs**.

Click this button to have the Table Uploader automatically configure the settings for the table based on the data in the file. If the file is formatted and delineated properly, the Table Uploader will identify the number of columns in your table, create a section for each column to provide the necessary information, and complete as much information as possible from the data in the file. While you may need to make some adjustments afterward, this option can save you a lot of time and effort compared to manually configuring the file yourself.

Note: The **Auto Detect File Specs** button cannot be used on a compressed file. If a saved specification file does not exist for the compressed file, you can either manually configure the upload file or you can create a new specification file.

- Select a previously saved specification file from the drop-down list and then click the **Load Spec** () icon.

Select this option to use a previously saved specification to configure the settings for the table. This option is best used if you regularly use the custom upload feature to add tables to the Insights Platform with the same data format. For example, if you upload weekly sales data, the specification file can automatically complete all of the configuration settings of the table for you.

The Advanced Uploader identifies each column in the table and completes the appropriate **Column Information** fields.

▼
Advanced Options

File Information

File Type: Delimited | Field Delimiter: Tab

Record Delimiter: CRLF | Ignore Nulls: Yes

of Records to Skip: 1 | Load:

▼

Table Information

Title: Table Title

Short Description: Short Desc.
Choose a title and description you want displayed in the 1010data filesystem view.

Long Description: Long Desc.
Long description you want displayed under table information.

Link Header: Link Header
Link header is used to denote foreign columns as the result of a link operation.

Time Series: No

Redundancy: Fully Mirrored

User: technical_writer +

▼
+ [table icon] ↑ ↓ Advanced Options

Column Information

Name: transaction_id | Header: transaction_id

Type: Integer | Format Type: Number (w/ comm...)

Display Width: 3

▼
+ [table icon] ↑ ↓ × Advanced Options

Column Information

Name: account | Header: account

Type: Integer | Format Type: Number (w/ comm...)

Display Width: 3

▼
+ [table icon] ↑ ↓ × Advanced Options

Column Information

Name: store | Header: store

Type: Integer | Format Type: Number (w/ comm...)

Display Width: 1

▼
+ [table icon] ↑ ↓ × Advanced Options

Column Information

Name: date | Header: date

Type: Text | Force Case: N

Display Width: 8

▼
+ [table icon] ↑ ↓ × Advanced Options

Column Information

Name: item_sku | Header: item_sku

Type: Text | Force Case: N

Display Width: 3

▼
+ [table icon] ↑ ↓ × Advanced Options

Column Information

Name: units | Header: units

Type: Integer | Format Type: Number (w/ comm...)

Display Width: 2

▼
+ [table icon] ↑ ↓ × Advanced Options

Column Information

Name: sales | Header: sales

Type: Float | Format Type: Number (w/ comm...)

Display Width: 4 | Dec. Places: Decimal F

▼
+ [table icon] ↑ ↓ × Advanced Options

Column Information

Name: cost | Header: cost

Type: Float | Format Type: Number (w/ comm...)

Display Width: 5 | Dec. Places: Decimal F

7. Complete or edit the **File Information**, **Table Information**, and **Column Information** fields and options.

For details, see [File Information](#) on page 360, [Table Information](#) on page 363, and [Column Information](#) on page 364.

8. As necessary, you can add, clone, move, or delete a column before uploading the file.

For details, see [Table 28: Column Information icons](#) on page 367.

Note: After configuring the fields and options in your file, you can create a specification file so the settings can be used again. This is helpful if you regularly upload files with the same data format. To save the settings, click the **Save Spec** (📄) icon, name the file, and click **Save**. The specification file is added to the drop-down list.

9. At the top right side of the **Advanced Uploader** window, click **Start Load**.

The Advanced Uploader uploads the file and creates a new table. Once the process is complete, a confirmation message and a preview of the uploaded table is displayed.

Note: If you close the **Table Uploader** window before the upload process is complete, the Table Uploader will abort any uploads that are in progress.

After the table has been successfully uploaded, you can share the table with other Insight Platform users.

File Information

Descriptions for fields and options in the **File Information** section of the **Advanced Uploader** tab.

The tables in this topic explain the fields and options available in the **File Information** section of the **Advanced Uploader** tab.

Figure 90: File Information section

Table 24: File Information fields and options

Field	Description
File Type	<p>The format of data in the file (table).</p> <p>Fixed</p> <p>Fields (columns) in the file have a fixed length. For example, in each record (row) of the table, the first column is always 10 characters, the second is three characters, and the third is 20 characters.</p> <p>Delimited</p> <p>Fields (columns) in the file may be different lengths; a character is used to separate each field in the record (row) of the table. One of the most common file formats is CSV (comma-separated values). For example, each field in a CSV delimited file is separated by a comma (,) character.</p>

Field	Description
Field Delimiter	<p>The character used to separate each field (column) in the record (row) of the file (table). This field displays when Delimited is selected from the Type drop-down list.</p> <p>Comma</p> <p>A comma (,) is used to separate the columns in each row of the table.</p> <p>Tab</p> <p>A tab stop is used to separate the columns in each row of the table.</p> <p>Pipe</p> <p>A pipe () is used to separate the columns in each row of the table.</p> <p>Other</p> <p>A character other than a comma, tab, or pipe is used to separate the columns in each row of the table.</p> <p>When Other is selected, the 1010data Insights Platform displays the Other field next to the Field Delimiter drop-down list. You must enter the character that is used to separate the columns in each row of the table.</p> <p>Note: Only ASCII special characters may be used as a field delimiter.</p>
Other Delimiter	<p>The character other than a comma, tab, or pipe that is used to separate each field (column) in the record (row) of the file (table).</p> <p>This field only displays if Other is selected in the Field Delimiter drop-down list.</p>
Record Delimiter	<p>The delimiter used to indicate the end of a each record (row) in the file (table). A delimiter is a sequence of one or more characters used to specify the boundary between separate, independent regions in plain text files. In general, the two most commonly used record delimiters in the Insights Platform are CRLF and LF.</p> <p>None</p> <p>The last column of each row does not have a delimiter. This option is used for fixed-width files.</p> <p>CRLF</p> <p>The Carriage Return Line Feed (CRLF) ('<code>\r\n</code>', 0x0D 0x0A) record delimiter is used in Microsoft Windows, DOS (MS-DOS, PC DOS, etc.), DEC TOPS-10, RT-11, CP/M, MP/M, Atari TOS, OS/2, Symbian OS, Palm OS, Amstrad CPC, and most other early non-Unix and non-IBM OSes. Select this option for PC files.</p> <p>NL</p> <p>The New Line (NL) (0x15) record delimiter is used in EBCDIC systems—mainly IBM mainframe systems, including z/OS (OS/390) and i5/OS (OS/400).</p> <p>CR</p>

Field	Description
	<p>The Carriage Return (CR) (' \r ', 0x0D) row delimiter is used in Mac OS files.</p> <p>LF</p> <p>The Line Feed (LF) (' \n ', 0x0A) record delimiter is used in Multics, Unix and Unix-like systems (Linux, OS X, FreeBSD, AIX, Xenix, etc.), BeOS, Amiga, RISC OS, and others.</p> <p>CRNL</p> <p>The Carriage Return New Line (CRNL) (0x0D 0x15) record delimiter is used for files delimited by the EBCDIC carriage return character followed by the EBCDIC new line character. The CRNL record delimiter is rarely used.</p>
Ignore Nulls	<p>Determines how the Insights Platform treats a null character in an upload file.</p> <p>Yes</p> <p>The Insights Platform ignores the null character and continues to upload the data included after the null character. An N/A is inserted in place of the null.</p> <p>No</p> <p>Unless masked, the Insights Platform will interpret a null character as the end of the file. Any data in the file after the null character will not be included in the upload.</p>
# of Records to Skip	<p>The number of records (rows) to skip when uploading the file (table).</p> <p>For example, the first row of a table usually contains names and not data. To exclude the first row of the table in the upload, enter 1. When the file is uploaded, the Insights Platform will ignore the first row in the table and start uploading the data beginning with row 2.</p>
Load	<p>The total number of records (rows) in the file (table) to upload.</p> <p>For example, to include only the first 100 rows of the table in the upload, enter 100. The Insights Platform will upload the first 100 rows in the table and then ignore any rows that follow.</p>

Advanced Options

The advanced options provide additional settings for input file handling and data structuring in the 1010data Insights Platform.

To view these additional fields, click **Advanced Options**.

Table 25: Advanced Options fields

Field	Description
Masking Character	<p>The character used to encapsulate fields that may contain the delimiting character. By default, the Insights Platform uses the quotation mark (") character.</p> <p>For example, a field of text data may also contain a comma—a common field delimiter. To prevent the Insights Platform from splitting the data field</p>

Field	Description
	<p>at the comma, that field needs to be surrounded by the indicated masking character.</p> <p>Note: If this field is left blank, the Insights Platform uses the quotation mark (") for the masking character. To set the masking character to none, enter \0.</p>
Masking Width	<p>The maximum number of characters within a masked field that the Insights Platform will consider when identifying field delimiter characters to ignore. The default maximum masking width is 1000 characters.</p> <p>For example, if a masked field is 2000 characters in length, but the maximum masking width is set at 1000, only the field delimiter characters that occur within the first 1000 characters of the field are ignored. Any field delimiter characters after the first 1000 characters will be treated as a standard field delimiter.</p>
Rows Per Seg	<p>The maximum number of records (rows) in the file (table) to include in a single segment. The default is 5 million rows and, in general, should not be adjusted.</p> <p>However, if the uploaded table will be used for aggressive expanding purposes, the maximum number of rows in each segment may need to be reduced. Contact 1010data Support to discuss your needs and we will help you determine the appropriate setting.</p>

Table Information

Descriptions for fields and options in the **Table Information** section of the **Advanced Uploader** tab.

The table in this topic explains the fields and options available in the **Table Information** section of the **Advanced Uploader** tab.

The screenshot shows the 'Table Information' section of the 'Advanced Uploader' tab. It contains the following fields and options:

- Title:** A text input field with the value 'Table Title'.
- Name:** A text input field with the value 'uploads' and a 'Table Name' button.
- Short Description:** A text input field with the value 'Short Desc.' and a note: 'Choose a title and description you want displayed in the 1010data filesystem view.'
- Long Description:** A text input field with the value 'Long Desc.' and a note: 'Long description you want displayed under table information.'
- Link Header:** A text input field with the value 'Link Header' and a note: 'Link header is used to denote foreign columns as the result of a link operation.'
- Time Series:** A dropdown menu with the value 'No'.
- Redundancy:** A dropdown menu with the value 'Fully Mirrored'.
- User:** A text input field with the value 'tc_cmliller' and a '+' button.

Figure 91: Table Information section

Table 26: Table Information fields and options

Field	Description
Title	The title of your table. This is for your own recognition of the table and does not necessarily have to be unique.

Field	Description
	The title may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters.
Name	<p>The name of your table. This is the <i>unique</i> name of the table used for identification in the 1010data Insights Platform.</p> <p>The name must begin with a letter and can only contain numbers, letters, and underscores. It cannot contain any spaces or other special characters.</p>
Short Description	A brief description of the data contained in your table.
Long Description	A detailed explanation of the data contained in your table.
Link Header	Denotes columns that were not originally part of your table but have made part of the table as the result of linking another table.
Time Series	<p>Enables the functionality to group together, or <i>segment</i>, like information in the uploaded table. A table must be segmented before g_functions or Time Series functions can be used to analyze the data in the table.</p> <p>When Yes is selected, the Insights Platform displays the Time Series Break Order field in the Column Information section for each column in the table. The Time Series Break Order field is used to segment the table. For more information, see Column Information on page 364.</p>
Redundancy	<p>Configures how segments are stored on 1010data servers that are provisioned to a customer.</p> <p>Fully Mirrored</p> <p>All segments are stored on all available customer-provisioned 1010data servers. This option provides full data redundancy and is the recommended setting.</p> <p>Striped</p> <p>Each segment is stored on at least two customer-provisioned servers. This option reduces the amount of space used on each server, but provides only partial data redundancy.</p>
User	Shows the users who have permission to view the table. You can specify additional users you want to have access to the table in this field.

Column Information

Descriptions for fields, options, and icons in the **Column Information** section of the **Advanced Uploader** tab.

For every column in your table, you need to provide the information required by the 1010data Insights Platform. If your data file was formatted and delineated properly, the **Auto Detect File Specs** button should have identified the number of columns in your table, created a section for you to provide the necessary information, and populated as much information as possible from the data in the file. The following tables provide descriptions of all fields, options, and icons available in the **Column Information** sections of the **Advanced Uploader** tab.

The screenshot shows the 'Column Information' configuration window. At the top, there are fields for 'Name' (containing 'Column Name') and 'Header' (containing 'Table Title'). Below these are 'Type' (set to 'Expression'), 'Force Case' (set to 'N'), and 'Format Type'. A 'Display Width' section contains 'Display W', 'Dec. Places', and 'Time Series Break Order' (set to 'Break Nut'). The main area has a 'Description' field with the text 'Column Help' and a note: 'Meta information that describes the context of the column data.' Below that is an 'Expression' field with '1010Data Expression' and a note: 'Write a 1010data expression to modify the data as it loads.' There are four checkboxes: 'Fixed Column', 'Hide Column', 'Destroy Column', and 'Custom Compression'. At the bottom, there are three dropdown menus: 'Type' (set to 'Default'), 'Method' (set to 'Default'), and 'Enumerate' (set to 'Default').

Figure 92: Column Information section

Table 27: Column Information fields and options

Field	Description
Name	<p>Also known as the column name, the Name provides a unique identifier for each column of a table in the web interface.</p> <p>The Name may only contain alphanumeric characters or underscores and must begin with an alphabetic character (e.g., percent_total_sales). It may not contain any spaces or other special characters.</p>
Header	<p>Also known as the column heading, the Header is the label of the column that displays by default at the top of a column in the web interface.</p> <p>The Header may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters (e.g., "Percentage of Total Sales (%)").</p>
Type	<p>The kind of data contained in the column. The most common options include the following:</p> <p>Text</p> <p>The column contains text, as opposed to numbers. When this option is selected, the Format Type drop-down changes to the Force Case drop-down.</p> <p>Integer</p> <p>The column contains only whole numbers. For example, 1, 53, and 1,234,597 are all integers.</p> <p>Float</p> <p>The column contains numbers with place values to the right of a decimal point. For example, 1.1, 23.845 and 0.37383636833930 are all floating point numbers.</p> <p>Expression</p> <p>Allows you to derive the value of a column with a Insight Platform Macro Language expression.</p> <p>Big Integer</p>

Field	Description
	<p>The column contains 64-bit integer data representing very large and very small whole numbers.</p> <p>Note: In addition to the types explained above, there are also numerous options for date and time data. These data formats apply Insight Platform date formatting measures to the date data in the source file. Dates in the source file can be formatted with dashes ('-'), forward-slashes ('/'), and spaces (' '). Data uploaded with just date information are created as <i>integers</i> in the Insights Platform, with time-handling formats applied as specified in the Type drop-down menu. Data with date and time data are created as <i>floats</i> in the platform, with the applicable time-handling formats applied as specified in the Type drop-down menu.</p>
Force Case	<p>Allows you to convert the text data in the column to all upper or all lower case letters. Select N to keep the text data as is.</p> <p>Note: This option is displayed only when Text or Expression is selected from the Type drop-down list.</p>
Format Type	<p>Defines how the data is displayed on the screen. For example, you can choose to exclude commas in numbers or change the way a date is formatted. Choices made in the Format Type do not change the type of data in the column, only the way the data is displayed.</p> <p>Note: This option is displayed only when Integer, Float, Big Integer, or Expression is selected from the Type drop-down list.</p>
Display Width	<p>The width of the column in the table. This field is used for fixed-width files.</p>
Dec Places	<p>The number of digits to display after a decimal point.</p> <p>Note: This option is displayed only when Float or Expression is selected from the Type drop-down list.</p>
Time Series Break Order	<p>Allows you to group together, or <i>segment</i>, like information in the uploaded table. A table must be segmented before g_functions or Time Series functions can be used to analyze the data in the table.</p> <p>The Time Series Break Order field accepts whole number sequential values starting with 1.</p> <p>For Time Series functions, a minimum of three columns must be identified. The first column, indicated by entering a 1 in the Time Series Break Order field, is the data that will be grouped together in the same segment. The second, up to the last column, indicated by entering a 2 (and then sequential whole numbers for each additional column) in Time Series Break Order field, identifies the grouping arguments. The last column, indicated by entering the next highest sequential number in the Time Series Break Order field, is used to order the data.</p> <p>For example, if you enter a 1 for the AccountID column and a 2 for the Store column and a 3 for the Date column, the Insights Platform will break the data up into groups ensuring that all records with a unique combination of AccountID and Store are in the same group and will sort the table by Date.</p> <p>The Insights Platform will always sort the table being loaded by the column with the highest value in the Time Series Break Order field. If</p>

Field	Description
	<p>in the previous example a 3 was not entered for the Date column, the platform would group all unique AccountID values in the same segment and sort the table by the Store column.</p> <p>Note: This field is displayed only when YES is selected from the Time Series drop-down list in the Table Information section.</p>

Column Information icons

Above the fields in each **Column Information** section are the following icons:



The table below contains a description of the function of each of these icons.

Table 28: Column Information icons

Icon	Function
+	Add a new column to the table. All the values in the new column will be blank.
	Clone a column so that an exact copy, with all associated values, is created.
↑	Move this column up in the order. This will move the column to the left in the final table.
↓	Move this column down in the order. This will move the column to the right in the final table.
×	Delete this column.

Advanced Options

The advanced options provide additional settings for columns.

To view these additional fields, click **Advanced Options**.

Table 29: Advanced Options fields

Field	Description
Description	<p>Explains the values and their meaning in a given column.</p> <p>The text entered in this field is displayed in the following locations within the Trillion-Row Spreadsheet and the Macro Language Workshop:</p> <ul style="list-style-type: none"> Description field under the Columns tab in the Data Dictionary Description field in the column information that appears when the pointer is placed over a column header within the Grid view of a table or worksheet <p>For more information, see Column information on page 49.</p>
Expression	<p>Macro Language expression for computed columns. Also helps control formatting for date and time data.</p> <p>Note: Expressions are built from a subset of the 1010data Insights Platform Macro Language. All functions from the Macro</p>

Field	Description
	Language are available for use in expressions except for g_ (group), r_ (row), and ts_ (time-series) functions.
Fixed Column	When you select this option, the column will remain in its current position when you scroll horizontally through the table columns in the Trillion-Row Spreadsheet.
Hide Column	Creates the column in the Insights Platform, but hides it. The column will be available in the Arrange columns panel in the Trillion-Row Spreadsheet Analysis Timeline. Hidden columns are also available for use in expressions.
Destroy Column	Excludes the column from the new table. The column will not be available in the base table after it is created.
Custom Compression	<p>Allows for custom compression settings.</p> <p>When selected, the Type, Method, and Enumerate drop-down lists are enabled.</p> <p>In general, the default compression setting should be used. However, if the table is extremely large and the column contains unique text values, the compression settings may need to be changed. Contact 1010data Support to discuss your needs and we will help you determine the appropriate settings.</p>
Type	<p>Determines the Default values for Method and Enumerate.</p> <p>This drop-down list is enabled when the Custom Compression option is selected.</p> <p>Default</p> <p>Dynamic is the default setting.</p> <p>Static</p> <p>Allows you to manually select the options for Method and Enumerate.</p> <p>Dynamic</p> <p>The Insights Platform will determine the best Method and Enumerate settings to use based on the data in the column.</p>
Method	<p>The type of compression to use.</p> <p>This drop-down list is enabled when the Custom Compression option is selected.</p> <p>Default</p> <p>The result of this option changes based on whether Static or Dynamic is selected from the Type drop-down list. If Dynamic is selected, leave Method set to Default. This allows the Insights Platform to determine the best type of compression to use based on the data in the column. If Static is selected, and Method is set to Default, lz4 compression is used.</p> <p>bitpack</p>

Field	Description
	<p>Very fast compression and decompression speeds. Only for compact integer vectors with few unique row values. Integer vectors cannot have 0I, -0I, or 0N infinity values.</p> <p>lzo</p> <p>Moderate compression ratios. Fast compression and very fast decompression speeds. This is the recommended compression method.</p> <p>bzip2</p> <p>Greater compression ratio than lzo compression. Slower compression and decompression speeds than lzo compression.</p> <p>lzma</p> <p>Excellent compression ratio. Very slow compression and moderately fast decompression speeds.</p>
Enumerate	<p>Enumeration is faster for string sets that have several repeated values because instead of storing the same string over and over, the repeated string is stored once and indexed. In general, leave Enumerate set to Default.</p> <p>This drop-down list is enabled when the Custom Compression option is selected.</p> <p>Default</p> <p>The result of this option changes based on the type of data selected in the <i>column data Type</i> drop-down list (not the <i>compression Type</i> drop-down list). If Text is selected and Enumerate is set to Default, enumeration is used. If Integer or Float is selected and Enumerate is set to Default, enumeration is not used.</p> <p>yes</p> <p>Use enumeration. This setting should be selected when there are a lot of repeated strings.</p> <p>Note: Static strings are always enumerated.</p> <p>no</p> <p>Enumeration is not used. This setting should be selected when there are no repeated row values in a very large string column.</p>

Save Data to FTP

Export the data to your 1010data FTP directory.

The **Save Data to FTP** window allows you to export the data as a compressed `.txt` file to your 1010data FTP directory.

Note: There are certain prerequisites you must have in place to save the data to your 1010data FTP directory. See the list of FTP items under [Prerequisites](#) on page 334 for more information.

In addition, you must configure a third-party FTP client to connect to the 1010data FTP server. For instructions, see [Transfer a file to your FTP account](#) on page 334.

The various tabs in this window allow you to customize the exported FTP file to suit your needs. The exported text file is compressed, and the file extension is `.zip`.

The **Save Data to FTP** window is accessible from the following locations:

- The **FTP** option under the **Export** tab in the Trillion-Row Spreadsheet
- The **File > Save Data... > ...to FTP** option in the Macro Language Workshop menu bar
- The **FTP** button in the **Converted QuickQuery** window

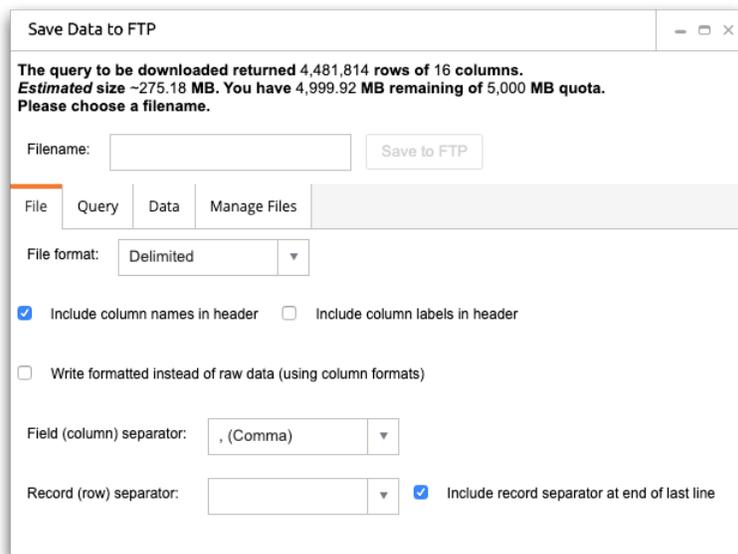


Figure 93: Save Data to FTP window

Filename

The name of the saved file exported to your 1010data FTP directory.

File

This tab contains settings for customizing the file data.

For more information, see [File](#) on page 371.

Query

This tab displays the query used to obtain the data results.

For more information, see [Query](#) on page 373.

Data

This tab shows the results of the query as a table.

For more information, see [Data](#) on page 373.

Manage Files

This tab allows you to view and delete files in your 1010data FTP directory.

For more information, see [Manage Files](#) on page 374.

File

Settings for customizing the file data.

Figure 94: File tab in the Save Data to FTP window

File format

Delimited

Currently, the only file format option is a delimited text file.

Include column names in header

When this option is selected, the column header in the file includes the column name for each column.

When this option is cleared, the column header in the file does not include column names.

Note: You can choose to include the column name, column label, neither, or both in the file column header.

Include column labels in header

When this option is selected, the column header in the file includes the column label for each column.

When this option is cleared, the column header in the file does not include column labels.

Note: You can choose to include the column name, column label, neither, or both in the file column header.

Write formatted instead of raw data (using column formats)

When this option is selected, the data is saved without any formatting. For example, in the case of a date, the data would appear in the date form as YYYYMMDD.

When this option is cleared, the data is saved with the display format applied to the column. In the same date example as above, the date may appear as MM/DD/YYYY.

Field (column) separator

The options in this drop-down list define the column separator character used to separate adjacent columns in the text file.

, (Comma)

Columns in each row of the table are separated by a comma (,) in the text file.

(Tab)

Columns in each row of the table are separated by a tab in the text file.

| (Stile/pipe)

Columns in each row of the table are separated by a pipe (|) in the text file.

: (Colon)

Columns in each row of the table are separated by a colon (:) in the text file.

; (Semicolon)

Columns in each row of the table are separated by a semicolon (;) in the text file.

/ (Slash)

Columns in each row of the table are separated by a slash mark (/) in the text file.

\ (Backslash)

Columns in each row of the table are separated by a backslash (\) in the text file.

*** (Asterisk)**

Columns in each row of the table are separated by an asterisk (*) in the text file.

Record (row) separator

The options in this drop-down list define the record (row) separator character. The separator is used to indicate the end of each row of the table in the text file.

The row separator is a sequence of one or more characters used to specify the boundary between separate, independent regions in plain text files. In general, the two most commonly used record delimiters in the Insights Platform are CRLF and LF.

LF (Newline)

The last column in each row of the table is indicated by the Line Feed (LF) row separator in the text file.

The LF ('`\n`', 0x0A) row separator is used in Multics, Unix and Unix-like systems (Linux, OS X, FreeBSD, AIX, Xenix, etc.), BeOS, Amiga, RISC OS, and others.

CR (Carriage return)

The last column in each row of the table is indicated by the Carriage Return (CR) row separator in the text file.

The CR ('`\r`', 0x0D) row separator is used in macOS.

CRLF (Carriage return + newline)

The last column in each row of the table is indicated by the Carriage Return Line Feed (CRLF) row separator in the text file.

The CRLF ('`\r\n`', 0x0D 0x0A) row separator is used in Microsoft Windows, DOS (MS-DOS, PC DOS, etc.), DEC TOPS-10, RT-11, CP/M, MP/M, Atari TOS, OS/2, Symbian OS, Palm OS, Amstrad CPC, and most other early non-Unix and non-IBM OSes. Select this option for PC files.

Include record separator at end of last line

Some operating systems may require a record separator at the end of the last row to indicate the end of the file.

When this option is selected, the last column in the last row of the table is indicated by the row separator selected above.

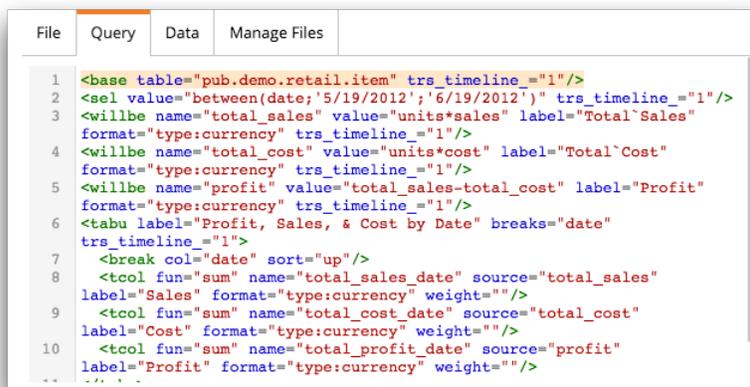
When this option is cleared, the last column of the last row does not have a row separator.

Query

View the query used to obtain the data results.

The query used to obtain the data results is displayed in the **Query** tab.

When you save the results of an analysis to your 1010data FTP directory, only the *results* are saved. The query used to obtain the results is not saved with the data.



```

1 <base table="pub.demo.retail.item" trs_timeline="1"/>
2 <sel value="between(date;'5/19/2012';'6/19/2012')" trs_timeline="1"/>
3 <willbe name="total_sales" value="units*sales" label="Total`Sales"
4 format="type:currency" trs_timeline="1"/>
5 <willbe name="total_cost" value="units*cost" label="Total`Cost"
6 format="type:currency" trs_timeline="1"/>
7 <willbe name="profit" value="total_sales-total_cost" label="Profit"
8 format="type:currency" trs_timeline="1"/>
9 <tabu label="Profit, Sales, & Cost by Date" breaks="date"
10 trs_timeline="1">
11 <break col="date" sort="up"/>
12 <tcot fun="sum" name="total_sales_date" source="total_sales"
13 label="Sales" format="type:currency" weight=""/>
14 <tcot fun="sum" name="total_cost_date" source="total_cost"
15 label="Cost" format="type:currency" weight=""/>
16 <tcot fun="sum" name="total_profit_date" source="profit"
17 label="Profit" format="type:currency" weight=""/>

```

Figure 95: Query tab in the Save Data to FTP window

A query is a list of the actions performed on a table to achieve a particular result or to perform a specific analysis. In other words, a query is the sequence of transformations that are performed on a particular base table.

Data

View the results of the query as a table.

The results of the query are displayed as a table in the **Data** tab.

Date	Sales	Cost	Profit
date	total_sales_date	total_cost_date	total_profit_date
	\$37.63	\$20.70	\$16.93
05/19/12	\$0.34	\$0.13	\$0.21
05/22/12	\$16.60	\$8.74	\$7.86
06/03/12	\$14.59	\$7.92	\$6.67
06/19/12	\$6.10	\$3.91	\$2.19

Figure 96: Data tab in the Save Data to FTP window

Manage Files

View and delete files in your 1010data FTP directory.

Sort by:	Name	Delete selected
0.000530 MB (file)	From Sales Item Detail.zip	
0.001284 MB (file)	Example Spec File.spec	
0.001219 MB (file)	Weekly Sales Item Detail.txt	

Figure 97: Manage Files tab in the Save Data to FTP window

Sort by

Sorts the list of files in your 1010data FTP directory by file name or file size.

Delete selected

Deletes the selected file from your 1010data FTP directory.

Send the data to your 1010data FTP directory

With FTP access, you can save data files to your 1010data FTP directory for downloading.

Your 1010data FTP directory allows you to download large data files.

For large data files, it may take the 1010data Insights Platform some time to create the file. Normally, once the file is created, the platform will notify you. However, if your browser loses the connection before you receive the notification (e.g., in the event of a proxy-server or firewall timeout), the Insights Platform will continue to create the file.

Note: Logging into the user interface to start a new session will end your previous session and stop the file creation.

To send the data to your 1010data FTP directory:

- Depending on the Insights Platform tool you are using, do one of the following:

Option	Description
Trillion-Row Spreadsheet	In the Export tab, click FTP .
Macro Language Workshop	From the menu bar, select File > Save Data... > ...to FTP .

The Insights Platform displays the **Save Data to FTP** window.

- In the **Filename** field, enter a name for the file.
- On the **File** tab, select the appropriate options to customize the file to suit your needs. For a description of the available options, see [File](#) on page 371.
- To view the query used to obtain the data, click the **Query** tab. For more information, see [Query](#) on page 373.
- To view the results of the query as a table, click the **Data** tab. For more information, see [Data](#) on page 373.
- Click **Save to FTP**. The Insights Platform saves the data to your FTP directory as a compressed text file.

Delete a file from your 1010data FTP directory

View and delete files from your 1010data FTP directory within the 1010data Insights Platform.

You can delete a file located in your 1010data FTP directory from the **Manage Files** tab in the **Save Data to FTP** window.

To delete a file from your 1010data FTP directory:

- Depending on the Insights Platform tool you are using, do one of the following:

Option	Description
Trillion-Row Spreadsheet	In the Export tab, click FTP .
Macro Language Workshop	From the menu bar, select File > Save Data... > ...to FTP .

The Insights Platform displays the **Save Data to FTP** window.

2. Click the **Manage Files tab.**

The Insights Platform displays a list of files in your 1010data FTP directory on the **Manage Files** tab.

3. Optionally, in the **Sort by drop-down list, select an option to sort the file list.**

The file list can be sorted by file name or file size.

The Insights Platform sorts the file list.

4. Select the file or files you want to delete from your 1010data FTP directory.

You can select multiple files using the **Shift** and **Control** keys.

5. Click **Delete selected.**

The Insights Platform deletes the selected files from your 1010data FTP directory.

Save Data as Table

Save the data resulting from your analysis to a 1010data Insights Platform table.

The **Save Data as Table** window allows you to save a new table, replace a table, or append data to an existing table. To save data as a table, you must be authorized to add tables to a folder. To replace or append data to a table, you must own the table.

When you save results of an analysis as a new table, replace a table with the results of an analysis, or append a table with the results of an analysis, only the *results* are saved, not the actions to the original table that led to the results.

After the table is saved, it is visible only to users to whom you give permission. If you need to grant permission to multiple users, consider creating a group.

The **Save Data as Table** window is accessible from the following locations:

- The **Table** option under the **Export** tab in the Trillion-Row Spreadsheet
- The **File > Save Data... > ...as table** option in the Macro Language Workshop menu bar
- The **Save** button in the **Converted QuickQuery** window

Figure 98: Save Data as Table window

Select folder to save a new table into or an existing table to replace/append to

When this field is clicked, the Insights Platform displays the object browser in the window.

The object browser allows you to select a folder to save the table into or to select an existing table. When an existing table is selected, you can replace the table or append the data to it.

After a folder or table is selected, this field displays the name of that object.

Folder

This field displays the full folder path of the object selected in the object browser. You can also enter the full path of a folder in this field instead of selecting it from the object browser.

Name

When an existing table is selected for replacement or to append data, this field displays the name of the table selected in the object browser. It also allows you to enter a name for a new table that you are saving.

The table name must begin with a letter and can only contain numbers, letters, and underscores. It cannot contain any spaces or other special characters. If you leave this field blank, a system-generated name will be used (e.g., t662528755_yourusername).

Note: If you are saving to the **My Data** folder, you cannot enter text in this field. Instead, a system-generated table name is used.

Metadata

This tab contains fields for entering a title and describing the table.

For more information, see [Metadata](#) on page 378.

Access

This tab provides settings for defining the users who can access the table.

For more information, see [Access](#) on page 379.

Append

This tab contains the option used to append the data to an existing table.

For more information, see [Append](#) on page 380.

Structure

This tab provides the fields and options used to segment the table.

For more information, see [Structure](#) on page 380.

Query

This tab displays the query used to obtain the data results.

For more information, see [Query](#) on page 381.

Data

This tab shows the results of the query as a table.

For more information, see [Data](#) on page 382.

Metadata

Fields for entering a title and describing the table.

Figure 99: Metadata tab in the Save Data as Table window

Table title

The title is used to help describe the contents of a table (e.g., `Sales Detail by Customer`). This is used for recognition of the table and does not necessarily have to be unique.

The title may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters.

Description

A brief description of the data contained in the table.

Long description (table help)

A detailed explanation of the data contained in the table.

You can use this field to provide detailed help information about using the table.

Access

Settings for defining the users who can access the table.

Figure 100: Access tab in the Save Data as Table window

Inherit permissions from parent folder

When this option is selected, users who have access to the parent folder (the folder the table is being saved in) will have access to the table. Users and groups listed in the **Authorized users and/or groups** field are ignored when this option is selected.

Note: You cannot use the inherit permissions feature for tables saved in your **My Data** folder. To share a table saved in your **My Data** folder, you must explicitly give permission to each user with whom you want to share the table.

Authorized users and/or groups

Enter the names of the 1010data Insight Platform users or groups, separated with a space, with whom you want to share the table. When this field is left blank, only you will have access to the table.

Note: Users and groups listed in this field are ignored when the **Inherit permissions from parent folder** option is selected.

Append

Setting used to append the data to an existing table.

An append will add every row in the current worksheet to the selected table only for columns that exist in the selected table. If the selected table contains columns that do not appear in your worksheet, or if the data type of a column name that appears in both the worksheet do not match, an error will occur. If your worksheet contains columns that do not appear in the selected table, those columns will not be included.

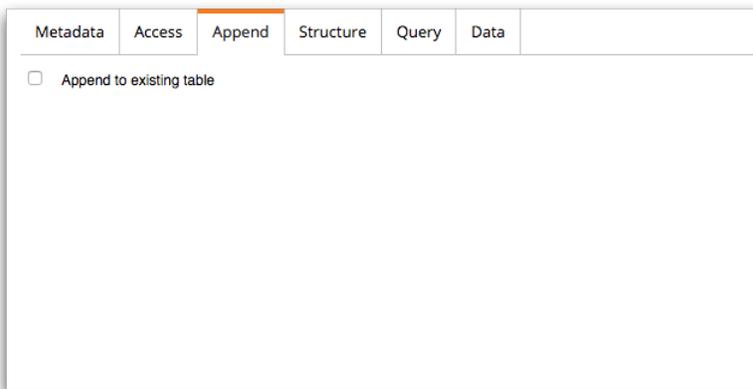


Figure 101: Append tab in the Save Data as Table window

Append to existing table

If an existing table is selected when saving data as a table, you can choose to replace the table or append the data to it.

When this option is selected, the data is appended to the table. When this option is cleared, the existing table is replaced.

Structure

Fields and options used to segment the table.

You can group together, or *segment*, like information in the saved table. A table must be segmented before [g_functions](#) or [Time Series functions](#) can be used to analyze the data in the table. For Time Series functions, a minimum of three columns must be identified.

Figure 102: Structure tab in the Save Data as Table window

Target segment size

The maximum number of rows in the table to include in a single segment. The default is 4 million rows and, in general, should not be adjusted.

However, if the saved table will be used for aggressive expanding purposes, the maximum number of rows in each segment may need to be reduced. Contact [1010data Support](#) to discuss your needs and we will help you determine the appropriate setting.

Columns to segment by

Specifies one or more columns that will be used to segment the table.

This type of segmentation is referred to as *segby*. If a table is segby a given column, no unique value of the column can be found in more than one segment.

For more information, see [segby](#).

Sort segments by

Specifies one or more columns that will be used to sort and segment the new table.

This type of segmentation is referred to as *sortseg*. If a table is sortseg on a particular column, not only are unique column values not allowed to be found in different segments, the segments themselves are internally sorted on the sortseg column.

For more information, see [sortseg](#).

Sort within segments by

Specifies one or more columns that will be sorted within each segment.

Columns to be indexed

Specifies which columns in the resulting table should be indexed.

Indexing provides increased selection performance at the cost of storage.

Query

View the query used to obtain the data results.

The query used to obtain the data results is displayed in the **Query** tab.

When you save the results of an analysis as a table, only the *results* are saved. The query used to obtain the results is not saved with the table.

```

1 <base table="pub.demo.retail.item" trs_timeline="1"/>
2 <sel value="between(date;'5/19/2012';'6/19/2012')" trs_timeline="1"/>
3 <willbe name="total_sales" value="units*sales" label="Total`Sales"
  format="type:currency" trs_timeline="1"/>
4 <willbe name="total_cost" value="units*cost" label="Total`Cost"
  format="type:currency" trs_timeline="1"/>
5 <willbe name="profit" value="total_sales-total_cost" label="Profit"
  format="type:currency" trs_timeline="1"/>
6 <tabu label="Profit, Sales, & Cost by Date" breaks="date"
  trs_timeline="1">
7   <break col="date" sort="up"/>
8   <tcot fun="sum" name="total_sales_date" source="total_sales"
  label="Sales" format="type:currency" weight=""/>
9   <tcot fun="sum" name="total_cost_date" source="total_cost"
  label="Cost" format="type:currency" weight=""/>
10  <tcot fun="sum" name="total_profit_date" source="profit"
  label="Profit" format="type:currency" weight=""/>

```

Figure 103: Query tab in the Save Data as Table window

A query is a list of the actions performed on a table to achieve a particular result or to perform a specific analysis. In other words, a query is the sequence of transformations that are performed on a particular base table.

Data

View the results of the query as a table.

The results of the query are displayed as a table in the **Data** tab.

Date	Sales	Cost	Profit
date	total_sales_date	total_cost_date	total_profit_date
	\$37.63	\$20.70	\$16.93
05/19/12	\$0.34	\$0.13	\$0.21
05/22/12	\$16.60	\$8.74	\$7.86
06/03/12	\$14.59	\$7.92	\$6.67
06/19/12	\$6.10	\$3.91	\$2.19

Figure 104: Data tab in the Save Data as Table window

Save the data to a table

Save a new table, replace a table, or append data to an existing table.

After you have completed an analysis, you can save the results as a new table on the 1010data Insights Platform. You will then be able to access the results at a later time and share them with others. In addition, you can replace a table or append data to an existing table that you own.

To save the data to a table:

1. Depending on the Insights Platform tool you are using, do one of the following:

Option	Description
Trillion-Row Spreadsheet	In the Export tab, click Table .

Option	Description
Macro Language Workshop	From the menu bar, select File > Save Data... > ...as table .

The Insights Platform displays the **Save Data as Table** window.

- Depending on whether you want to save a new table, replace an existing table, or append data to a table, do one of the following:

Option	Description
Save a new table	Specify the folder in which you want to save the new table and, optionally, give the table a name.
Replace an existing table	Select the table you want to replace.
Append data to a table	Select the table to which you want to append data.

For more information, see [Save Data as Table](#) on page 377.

- On the **Metadata** tab, enter the appropriate table information.
For a description of the available fields, see [Metadata](#) on page 378.
- To give other users access to the table, complete the appropriate fields and options on the **Access** tab.
For a description of the available fields and options, see [Access](#) on page 379.
- To append data to an existing table, select the option on the **Append** tab.
For more information, see [Append](#) on page 380.
- To segment the table, complete the appropriate fields on the **Structure** tab.
For a description of the available fields, see [Structure](#) on page 380.
- To view the query used to obtain the data, click the **Query** tab.
For more information, see [Query](#) on page 381.
- To view the results of the query as a table, click the **Data** tab.
For more information, see [Data](#) on page 382.
- Click one of the following:

Option	Description
Append	<p>Click this button to append the data to an existing table.</p> <p>This button appears when an existing table is chosen and the Append to existing table option is selected on the Append tab.</p>
Replace	<p>Click this button to replace an existing table.</p> <p>This button appears when an existing table is chosen.</p>
Save	<p>Click this button to save the data to a new table.</p> <p>This button appears when a folder is selected.</p>

The Insights Platform saves the data to a table.

Quick Queries/PQQs

Quick Queries can be used to save a set of actions so that they may be rerun at a later time.

A Quick Query is a set of saved actions that can be used by anyone on your team. To really harness the power of Quick Queries, add parameters to create a Parameterized Quick Query (PQQ). Each time the user runs the query, the user is prompted for specific input. Since the user can make selections and see different results, the PQQ becomes a flexible tool for data analysis.

If you applied certain actions to a table (such as row selections, sorting, linking, computed columns, summarizations), you may specify which parameters may be input each time the Quick Query is run. For example, if you performed a row selection such as **Column1 is greater than 100**, you may want to use something other than 100 the next time you run the query. Using Quick Queries, you can make that value a user input. Similarly, you may wish to have the user select a different column (e.g., to use **Column2** instead of **Column1**) or even a different relationship (e.g., **is less than** instead of **is greater than**). Each of these parameters of the selection (value, column, and relationship) may be made into a user selection. Similarly, for tabulations, the choice of grouping columns, result columns, and summarization methods may all be made user selections, as can the parameters of most other actions. You can also prompt the user to select the table to which the analysis is applied. Parameters are added to operations in the TRS timeline.

Note the difference between saving actions as a Quick Query or Parameterized Quick Query and saving the results as a new table. When you save your actions as a Quick Query, only the *actions* are saved, not the results. When the Quick Query is run at a later time, the actions are applied to the table at that time. If the table had been updated or corrected, the results of the Quick Query would reflect those changes. On the other hand, when you save the results as a new table, the actual *results* are saved. If the original table were modified, the new table that you had previously saved would not reflect the changes.

You can run a Quick Query in TRS by right-clicking the query and selecting **Run query in TRS**. Depending on your workspace settings, you may also be able to run the Quick Query by double-clicking it in the Object Manager. See [Define object associations](#) on page 43 for information on configuring your default settings for double-clicking an object.

You also have the option of viewing and interacting with a Quick Query in other ways. For more information, see [Object actions by type](#) on page 93.

Note:

Some customers have legacy PQQs from 1010data's legacy environment (version 9 and prior). Legacy PQQs have some differences, which are discussed in [Legacy Quick Queries/PQQs](#) on page 402.

Quick Queries are identified by the icons:

-  - denotes a Quick Query created in version 10 and later.
-  - denotes a legacy Quick Query

In addition, a Quick Query allows you to share your analysis with others, since you may give others permission to access it. When you share a Quick Query with another user or group, they must also have access to the parent folder of the Quick Query (and all of the folders in its full path) as well as any tables that the Quick Query references. View information about a particular Quick Query to find out its full path and the tables it references. For more information, see [Share an object](#) on page 109.

Block setup panel

The **Block setup** panel allows you to add parameters to your analysis in TRS.

The **Block setup** panel is accessible from the Analysis Timeline in the Trillion-Row Spreadsheet. You must first run an operation before you can parameterize it. After you run the operation, click the **Create Block** (📦) icon to view the **Block setup** panel.

Note: The options in the **Block setup** panel depend on the operation you are parameterizing. The following is a **Block setup** panel for a simple **Select** operation. For more information about the **Select** panel, see [Select rows panel](#) on page 142.

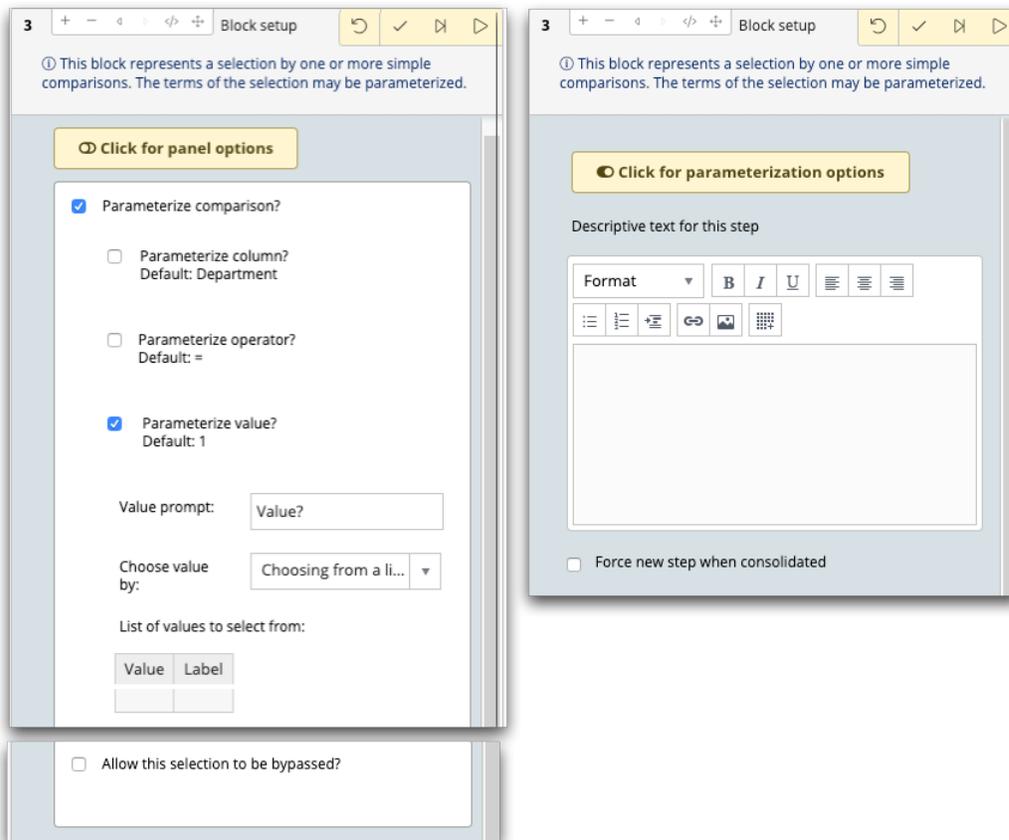


Figure 105: Block setup panel

Parameterize comparison?

Select this option to parameterize the select statement. When you save the query as a Parameterized Quick Query, your users can run the query and perform their own selections. You can parameterize one, two, or all three of the following. Parameterizing all three gives the user the most flexibility when running the PQQ.

Parameterize column?

This option allows the user to select from different columns in the resulting PQQ. If you do not select this option, the user can make a selection from only the default column (in this case **Department**) in the PQQ.

Parameterize operation?

This option allows the user to select from different operators in the resulting PQQ, such as equal to (=), greater than (>), or less than (<). If you do not select this option, the user can use only the default operation (in this case =) in the PQQ.

Parameterize value?

This option allows the user to select different values for the selected column. If you do not select this option, the user can use only the default option (in this case 1) in the PQQ.

For each part of **Select** that you choose to parameterize, you will have additional options. The options below assume that you are parameterizing **value** in the **Select** operation.

Value prompt

This text prompts the user to input or select the required information. For example, if you want the user to choose from a list of department numbers, the prompt could be `Department number?`.

Choose value by

Select how you want the user to input or select the values for the PQQ.

Choosing from a list of values

The user can choose from a list of values that you can manually enter in the **List of values to select from** field.

Entering a value directly

The user enters the value manually. If the value does not exist for the column, the result of the PQQ is an empty table.

Choosing from values in column

The user can choose from a list of all the possible column values in the table.

Entering a value with autocomplete

As the user enters a value, the system finds appropriate values that the user can select from a drop-down list.

Allow this selection to be bypassed?

If you select this option, the user can choose to ignore this step. You can prompt the user with the option of skipping this step.

Parameterization options (click slider to toggle these options)

Descriptive text for this step

If you would like to add more description for users when they run this step of your Quick Query, add it here.

Force new step when consolidated

If you choose to consolidate your timeline into as few blocks as possible, but you want to force this to step to be a new step, select this option. See [Consolidate timeline](#) on page 283 for more information.

Add parameters to an operation

Adding parameters to an operation in your analysis allows you to create a Parameterized Quick Query (PQQ). Parameterized Quick Queries allow users to specify inputs each time the query is run.

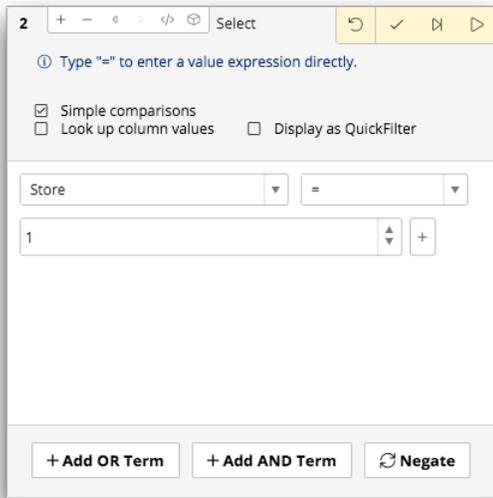
After you have completed an analysis, you may add parameters to an operation in the timeline.

Note: Available parameterization options vary based on the operation type (e.g., select, sort, tabulate). For illustration purposes, images of the select operation are used in this topic.

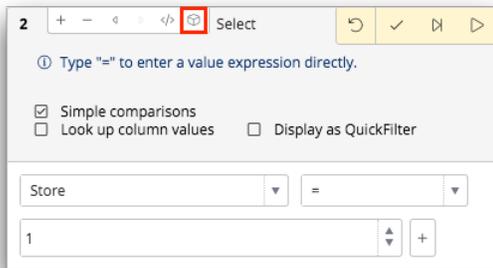
To add parameters to an operation:

1. In the timeline, click the panel of the operation to which you want to add parameters. The Trillion-Row Spreadsheet displays the operation panel.

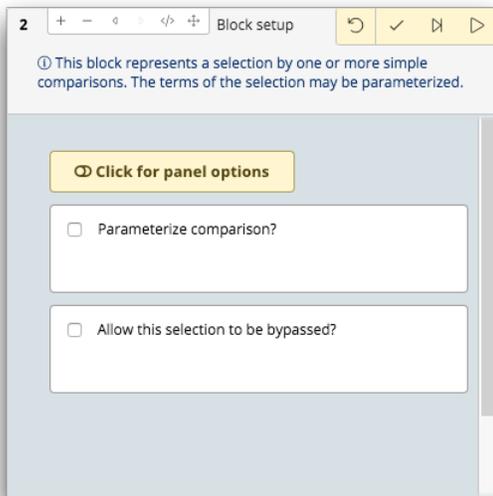
Note: The panel below shows a select operation that was already created and run in the timeline.



2. Click the **Create block** (🧩) icon.



The Trillion-Row Spreadsheet displays the **Block setup** view of the operation panel.



Note: See [Block setup panel](#) on page 385 for a more detailed description of the options available in this panel.

3. Optionally, to add descriptive text for this step of the analysis, click the **Click for panel options** switch.

The descriptive text is displayed for this step of the analysis when the Parameterized Quick Query is run.

4. Select the **Parameterize** option.

The **Parameterize** option varies based on the operation type. For example, in the select operation shown in the image below, the **Parameterize comparison** option is selected to add parameters to the operation.

The Trillion-Row Spreadsheet displays parameters available for the operation.

5. For each parameter you want to prompt for input when the Parameterized Quick Query is run:

- a) Select the desired parameter.
- b) Enter the text that will be displayed to prompt for input.
- c) Select the input type for the parameter from the drop-down list.

The input type depends on the type of parameter (e.g., column, value, relation) and the action associated with the parameter (e.g., select, sort, tabulate).

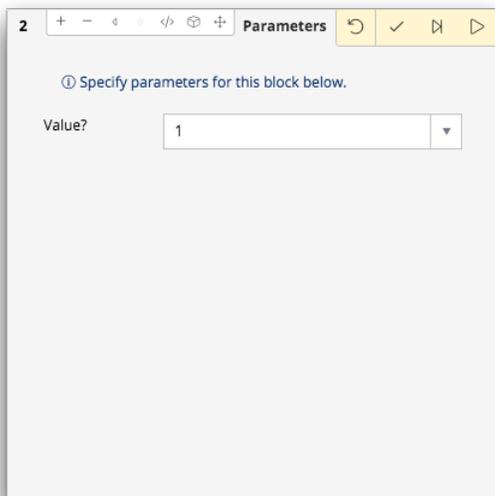
- d) When manually defining a list of values, enter the values in the table at the bottom of the **Parameterize** option section.

6. Optionally, to allow this step of the analysis to be bypassed, complete the following:

- a) Select the bypass option.
- b) Enter the text that will be displayed to prompt for input.

7. Click the **Submit operation** (✓) icon.

The Trillion-Row Spreadsheet saves the parameters and displays the **Parameters** view of the operation panel.



After adding parameters to the desired operations in your analysis, you can save it as a new Parameterized Quick Query or replace an existing one. For instructions, see [Save an analysis as a Quick Query](#) on page 281.

Edit parameters in an operation

Edit parameters previously added to an operation.

As necessary, you can update a Parameterized Quick Query by editing parameters previously added to an operation.

Note: Available parameterization options vary based on the operation type (e.g., select, sort, tabulate). For illustration purposes, images of the select operation are used in this topic.

To edit parameters in an operation:

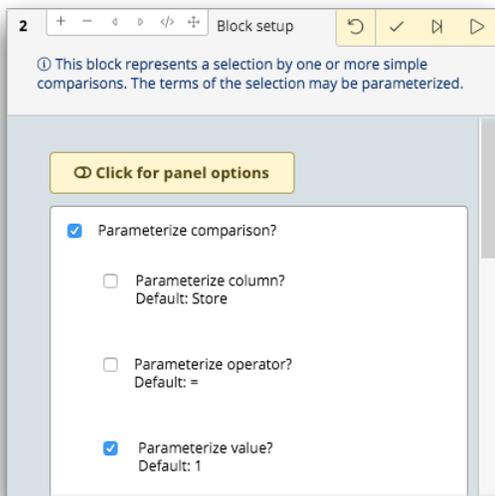
1. If necessary, open the saved Parameterized Quick Query in the Trillion-Row Spreadsheet by completing the following:
 - a) Open the **Object Manager** and locate the Parameterized Quick Query.
 - b) Right-click the Parameterized Quick Query and choose **Open in TRS** from the menu.The 1010data Insights Platform opens the Parameterized Quick Query in the Trillion-Row Spreadsheet.

Transaction ID	Account	Store	Date	Item SKU	Units	Sales
531	957	1	05/15/12	366	-1	-5
532	478	1	05/15/12	98A	1	0.5
532	478	1	05/15/12	3B7	1	1.1
534	738	1	05/16/12	A96	2	6
534	738	1	05/16/12	65B	1	2.25
538	668	1	05/18/12	CB7	1	1.65
538	668	1	05/18/12	98A	4	2
538	668	1	05/18/12	96A	1	1.1
538	668	1	05/18/12	969	1	1.1
538	668	1	05/18/12	3B7	1	1.1
543	876	1	06/03/12	CB7	1	1.65
543	876	1	06/03/12	A96	1	4
543	876	1	06/03/12	98A	2	1
543	876	1	06/03/12	96A	1	1.1
543	876	1	06/03/12	969	1	1.1
543	876	1	06/03/12	3B7	2	2.2
543	876	1	06/03/12	3A4	1	0.34

- In the timeline, click the panel of the operation containing the parameters you want to edit. The Trillion-Row Spreadsheet displays the **Parameters** view of the operation panel.

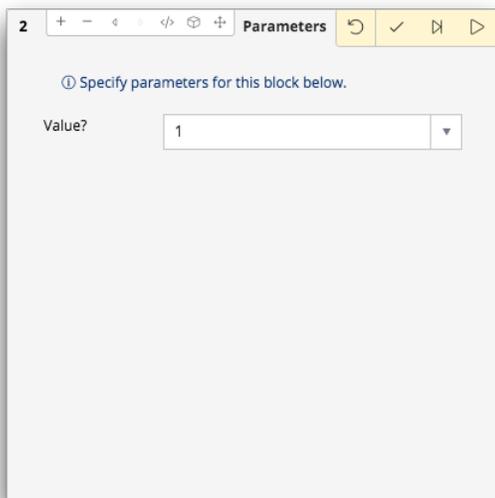
- Click the **Edit block** (🔧) icon.

The Trillion-Row Spreadsheet displays the **Block setup** view of the operation panel.



Note: See [Block setup panel](#) on page 385 for a more detailed description of the options available in this panel.

4. Edit the appropriate fields and options for each parameter that you want to update.
5. Click the **Submit operation** (✓) icon.
The Trillion-Row Spreadsheet saves the changes and displays the **Parameters** view of the operation panel.



After editing the desired parameters in the analysis, you can save it as a new Parameterized Quick Query or replace an existing one. For instructions, see [Save an analysis as a Quick Query](#) on page 281.

Edit a parameter with a large value selection

You can use the **Query** tab of the TRS timeline to edit a parameter with a large value selection.

To edit a parameterization with a large value selection:

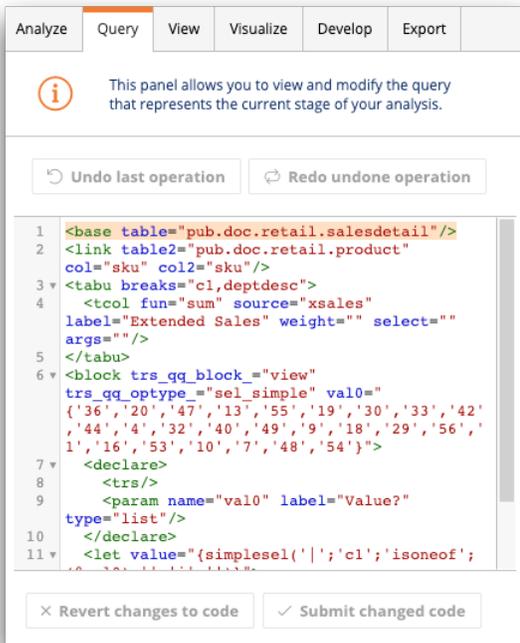
1. If necessary, open the saved Parameterized Quick Query in the Trillion-Row Spreadsheet by completing the following:
 - a) Open the **Object Manager** and locate the Parameterized Quick Query.

- b) Right-click the Parameterized Quick Query and choose **Open in TRS** from the menu.
The 1010data Insights Platform opens the Parameterized Quick Query in the Trillion-Row Spreadsheet.

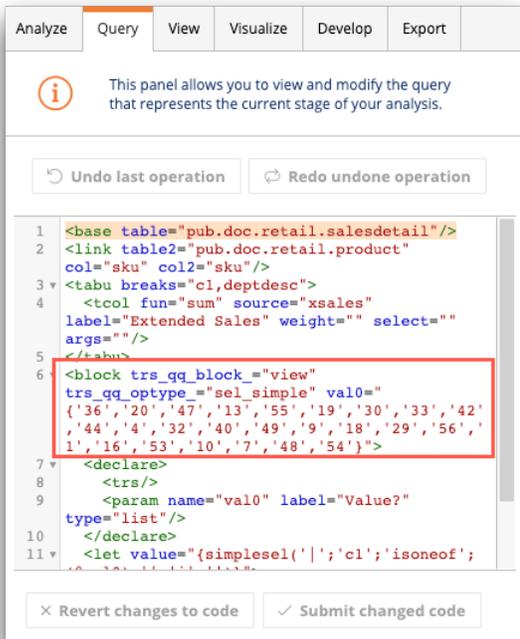
Department	Dept Desc	Extended Sales
		18,620,973.11
36	PRODUCE	3,042,735.19
20	DAIRY DELI	1,678,241.62
47	MEAT DELI	753,677.59
13	BAKERY	670,978.91
55	DAIRY	978,335.52
19	BEVERAGE	980,450.13
30	FROZEN	1,328,196.52
33	GM	532,193.47
42	BEER	606,907.55
44	MEAT	404,284.35
4	DISTILLED SPIRITS	490,779.92
32	SERVICE DELI	1,070,810.43
40	WINE	565,786.50
49	SEAFOOD	367,355.41
9	PACIFIC COAST CAFE	29,127.71
18	BOOKS/MAGAZINES	101,445.19
29	GARDEN	114,697.08

2. If you aren't already at that step in the timeline, click the panel of the selection operation with the large value selection.
The Trillion-Row Spreadsheet displays the **Parameters** view of the operation panel with the set of values.

3. Click the **Query** tab.
TRS displays the PQQ as Macro Language code.

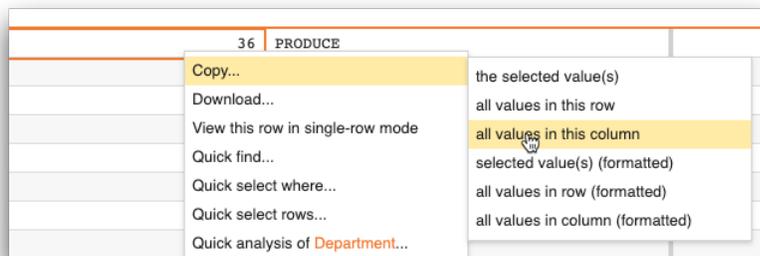


4. Locate the values within the parameterized block (in the `<block>` section of the code).

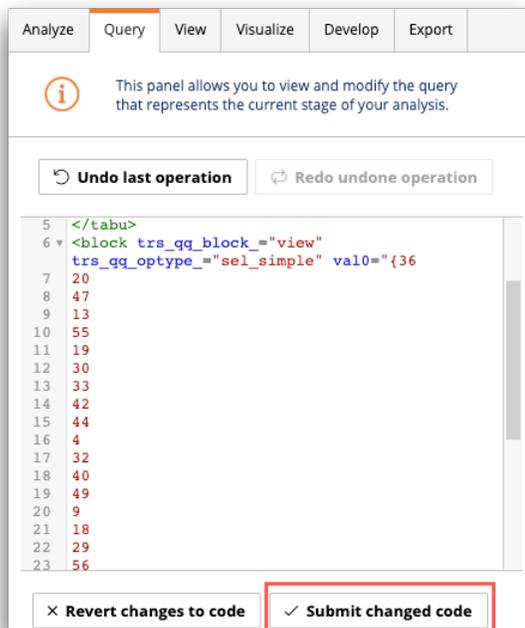


5. Replace the values with the new values you want.

Note: You can use the copy function from the context menu of the results pane to copy values and paste them into the macro code.



6. When you are finished with your edits, click **Submit changed code**.



7. Return to the **Analyze** tab.

8. Click the **Submit operation** (✓) icon for the edited operation.

After editing the desired parameters in the analysis, you can save it as a new Parameterized Quick Query or replace an existing one. For instructions, see [Save an analysis as a Quick Query](#) on page 281.

Remove all parameters from an operation

Remove all parameters previously added to an operation.

As necessary, you can remove all parameters previously added to an operation.

Note: Available parameterization options vary based on the operation type (e.g., select, sort, tabulate). For illustration purposes, images of the select operation are used in this topic.

To remove all parameters from an operation:

1. If necessary, open the saved Parameterized Quick Query in the Trillion-Row Spreadsheet by completing the following:
 - a) Open the **Object Manager** and locate the Parameterized Quick Query.
 - b) Right-click the Parameterized Quick Query and choose **Open in TRS** from the menu.

The 1010data Insights Platform opens the Parameterized Quick Query in the Trillion-Row Spreadsheet.

The screenshot shows the 1010data Insights Platform interface. On the left, a parameterized query is displayed in the 'Parameters' view. The query is titled 'pub.demo.retail.item' and has a single parameter 'Value?' with a value of '1'. The right side of the interface shows the 'Sales Item Detail' spreadsheet, which displays 17 rows of data. The columns are Transaction ID, Account, Store, Date, Item SKU, Units, and Sales. The data is as follows:

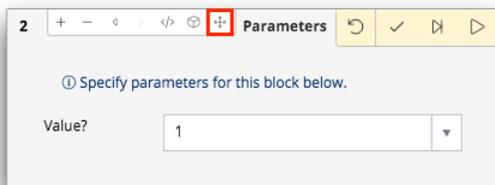
Transaction ID	Account	Store	Date	Item SKU	Units	Sales
531	957	1	05/15/12	366	-1	-5
532	478	1	05/15/12	98A	1	0.5
532	478	1	05/15/12	3B7	1	1.1
534	738	1	05/16/12	A96	2	6
534	738	1	05/16/12	65B	1	2.25
538	668	1	05/18/12	CB7	1	1.65
538	668	1	05/18/12	98A	4	2
538	668	1	05/18/12	96A	1	1.1
538	668	1	05/18/12	969	1	1.1
538	668	1	05/18/12	3B7	1	1.1
543	876	1	06/03/12	CB7	1	1.65
543	876	1	06/03/12	A96	1	4
543	876	1	06/03/12	98A	2	1
543	876	1	06/03/12	96A	1	1.1
543	876	1	06/03/12	969	1	1.1
543	876	1	06/03/12	3B7	2	2.2
543	876	1	06/03/12	3A4	1	0.34

The Trillion-Row Spreadsheet displays all of the operations of the query in the timeline.

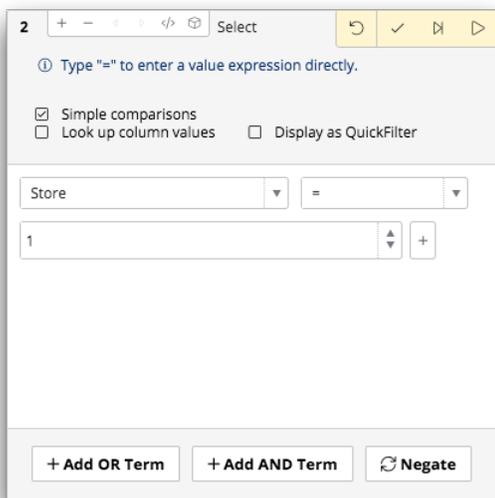
- In the timeline, click the panel of the operation containing the parameters you want to remove. The Trillion-Row Spreadsheet displays the **Parameters** view of the operation panel.

This is a close-up screenshot of the 'Parameters' view of the operation panel. It shows a text input field labeled 'Value?' with the value '1' entered. The panel title is 'Parameters' and it includes a 'Specify parameters for this block below.' instruction.

- Click the **Expand block** (+) icon.



The Trillion-Row Spreadsheet removes the parameters and displays the edit panel view of the operation panel.



After removing the desired parameters in the analysis, you can save it as a new Parameterized Quick Query or replace an existing one. For instructions, see [Save an analysis as a Quick Query](#) on page 281.

Edit a Parameterized Quick Query in MLW

You can open and edit a PQQ in Macro Language Workshop.

There are cases where you will want to edit a Quick Query or PQQ in Macro Language Workshop. However, you can add a new parameterization to your PQQ only in TRS.

Note: You can also edit legacy (pre version 10) Quick Queries and PQQs in MLW. See [Edit a legacy Parameterized Quick Query in MLW](#) on page 405.

To edit a parameterization in Macro Language Workshop:

1. Open the **Object Manager** and locate the Parameterized Quick Query.
2. Right-click the Quick Query or Parameterized Quick Query and choose **Edit query in Macro Language Workshop** from the menu.

The 1010data Insights Platform opens the Parameterized Quick Query in the Macro Language Workshop.

```

MLW: uploads.t447913045 - PQQ in MLW
File Code Run Debug render
1 <base table="pub.demo.retail.item"/>
2 <block trs_qq_block="view" trs_qq_optype="block"
  dept="{ '14', '22', '29' }">
3 <declare>
4   <param type="list" name="dept" label="dept"/>
5 </declare>
6 <block name="block" dept="{@dept}">
7   <link table2="pub.demo.retail.prod" col="sku"
  col2="sku" type="select"/>
8   <sel value="dept={@dept}"/>
9   <tabu breaks="dept,div">
10    <col name="units" source="units" fun="sum"
  label="Total Units"/>
11  </tabu>
12 </block>
13 </block>

```

3. Perform edits to your query.

Note: You can edit your Quick Query or PQQ in MLW, but you cannot parameterize a selection in MLW. For that, you must use the timeline in TRS. See [Add parameters to an operation](#) on page 387 for more information.

In this simple example, we changed the default stores and the selection prompt in the PQQ.

```

MLW: uploads.t447913045_ - PQQ in MLW
File Code Run Debug render
1 <base table="pub.demo.retail.item"/>
2 <block trs_qq_block="view" trs_qq_optype="block"
  dept="{ '22', '29' }">
3 <declare>
4   <param type="list" name="dept" label="Enter
  Department Number(s)"/>
5 </declare>
6 <block name="block" dept="{@dept}">
7   <link table2="pub.demo.retail.prod" col="sku"
  col2="sku" type="select"/>
8   <sel value="dept={@dept}"/>
9   <tabu breaks="dept,div">
10    <col name="units" source="units" fun="sum"
  label="Total Units"/>
11  </tabu>
12 </block>
13 </block>

```

4. Click **render** to run the query.

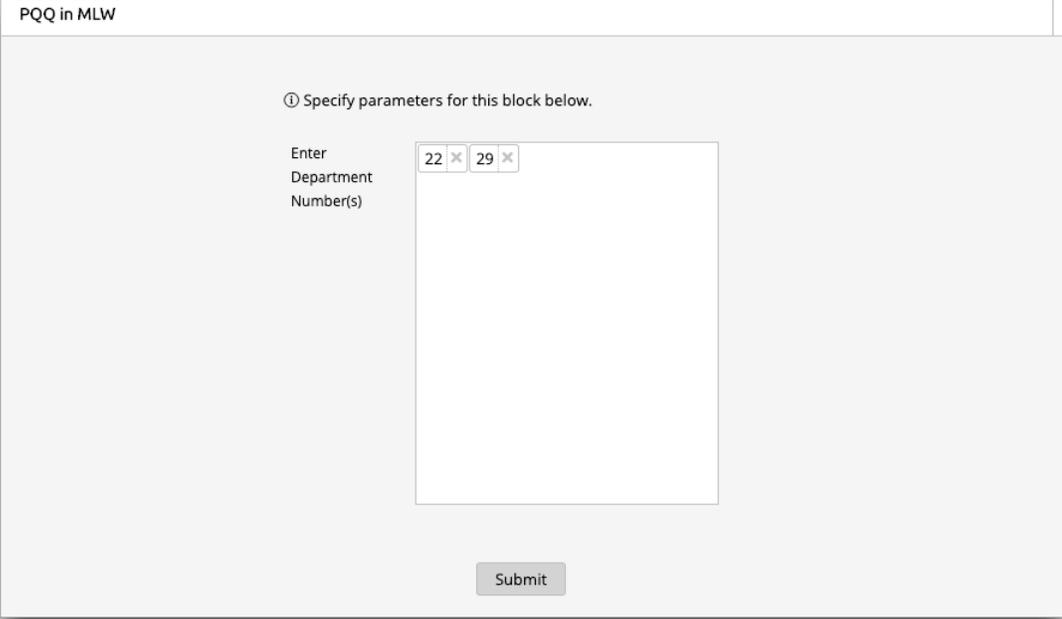
The query results appear in the right panel of MLW.

Cols 2 fixed and 3 to 3 of 3 , Rows 1 to 2 of 2

Department	Division	Total Units
		31
22	2	27
29	2	4

5. Click **File > Save Query** to overwrite the existing PQQ or **File > Save Query As** to save as a new PQQ.

6. To test your changes, right-click your PQQ in the Object Manager and choose **Run query in TRS**.



The screenshot shows a window titled "PQQ in MLW". Inside the window, there is a heading "Specify parameters for this block below." followed by a list of parameters: "Enter", "Department", and "Number(s)". To the right of these parameters is a large text input area. At the top of this input area, the numbers "22" and "29" are entered, each followed by a small "x" icon. At the bottom center of the window is a "Submit" button.

Run a Quick Query

Run a Quick Query or legacy Quick Query to view the results of a set of saved actions.

By default, when you run either a Quick Query or a legacy Quick Query, the 1010data Insights Platform opens the query in the view only mode of the Trillion-Row Spreadsheet and displays the results in the grid.

To run a Quick Query:

1. Open the **Object Manager** and locate the Quick Query you want to run.

Note: If you use a Quick Query often, you may want to add it to your favorites. You can then find it quickly and easily in the **Favorites** tab on the **Dashboard**.

2. Right-click the Quick Query and choose **Run Query in TRS** from the menu.

Note: Depending on your workspace association settings, you may also double-click a Quick Query to run it. For instruction on setting your workspace associations, see [Define object associations](#) on page 43.

The Insights Platform runs the Quick Query and displays the results.

Sales Margin

Sales Item Detail
Cols 1 to 9 of 9, Rows 1 to 19 of 35

Transaction ID	Account	Store	Date	Item SKU	Units	Sales	Cost	Margin
539	25	2	06/19/12	65B	1	5	3.35	1.65
540	361	2	05/19/12	3A4	1	0.34	0.13	0.21
540	361	2	06/19/12	3B7	1	1.1	0.56	0.54
541	469	3	05/22/12	96A	1	1.1	1	0.1
541	469	3	05/22/12	969	1	1.1	1	0.1
541	469	3	05/22/12	98A	1	0.5	0.25	0.25
541	469	3	05/22/12	3B7	3	3.3	1.68	1.62
541	469	3	05/22/12	A96	1	4	1.45	2.55
532	478	1	05/15/12	98A	1	0.5	0.25	0.25
532	478	1	05/15/12	3B7	1	1.1	0.56	0.54
537	523	3	05/15/12	96A	1	1.1	1	0.1
537	523	3	05/15/12	98A	1	0.5	0.25	0.25
537	523	3	05/15/12	3B7	1	1.1	0.56	0.54
537	523	3	05/15/12	CB7	1	1.65	1.1	0.55
537	523	3	05/15/12	A96	2	6	2.9	3.1
538	668	1	05/18/12	96A	1	1.1	1	0.1
538	668	1	05/18/12	969	1	1.1	1	0.1
538	668	1	05/18/12	3B7	1	1.1	0.56	0.54
538	668	1	05/18/12	CB7	1	1.65	1.1	0.55

The view bar on the right side of the results pane offers the following options for interacting with your results:

- If you interact with the data in the grid after running the Quick Query (such as sorting the data), you can undo and redo these operations with the **Go Back** (↶) and **Go Forward** (↷) icons. (These icons will not appear unless you interact with the data.)
- View the results in grid (📊) or single-row (📄) view.
- Download the results as a .csv (📄), .xlsx (📄), or .pdf (📄) file.
- Open a full TRS timeline (📅) and perform further analysis on your data.

Run a Parameterized Quick Query

Run a Parameterized Quick Query to view the results of set of saved actions with user inputs.

By default, when you run a Parameterized Quick Query, the 1010data Insights Platform opens the query in the view only mode of the Trillion-Row Spreadsheet and displays the parameter input fields and options.

To run a Parameterized Quick Query:

1. Open the **Object Manager** and locate the Parameterized Quick Query you want to run.

Note: If you use a Parameterized Quick Query often, you may want to add it to your favorites. You can then find it quickly and easily in the **Favorites** tab on the **Dashboard**. See [Add an object to your favorites](#) on page 110 for instructions.

2. Right-click the Parameterized Quick Query and choose **Run Query in TRS** from the menu.

Note: Depending on your workspace association settings, you may also double-click a Parameterized Quick Query to run it. This is the default setting. For instruction on setting your workspace associations, see [Define object associations](#) on page 43.

The Insights Platform displays the Parameterized Quick Query parameter input fields and options in a new window.

3. Select or enter any parameter inputs that the Parameterized Quick Query contains.

4. Click **Submit**.

The Insights Platform runs the Parameterized Quick Query and displays the results.

Transaction ID	Account	Store	Date	Item SKU	Units	Sales	Cost
531	957	1	05/15/12	366	-1	-5	-1.84
532	478	1	05/15/12	98A	1	0.5	0.25
532	478	1	05/15/12	3B7	1	1.1	0.56
534	738	1	05/16/12	A96	2	6	2.9
534	738	1	05/16/12	65B	1	2.25	1.35
538	668	1	05/18/12	CB7	1	1.65	1.1
538	668	1	05/18/12	98A	4	2	1
538	668	1	05/18/12	96A	1	1.1	1
538	668	1	05/18/12	969	1	1.1	1
538	668	1	05/18/12	3B7	1	1.1	0.56
543	876	1	06/03/12	CB7	1	1.65	1.1
543	876	1	06/03/12	A96	1	4	1.45
543	876	1	06/03/12	98A	2	1	0.5
543	876	1	06/03/12	96A	1	1.1	1
543	876	1	06/03/12	969	1	1.1	1
543	876	1	06/03/12	3B7	2	2.2	1.12
543	876	1	06/03/12	3A4	1	0.34	0.13

The view bar on the right side of the results pane offers the following options for interacting with your results:

- Click the **Go Back** (↶) icon to return to the PQQ input screen and specify different parameters in the PQQ. In addition, if you interact with the data in the grid after running the PQQ (such as sorting the data), you can undo and redo these operations with the **Go Back** (↶) and **Go Forward** (↷) icons.
- View the results in grid (📊) or single-row (📄) view.
- Download the results as a .csv (📄), .xlsx (📄), or .pdf (📄) file.
- Open a full TRS timeline (☰) and perform further analysis on your data.

Legacy Quick Queries/PQQs

It is now possible to edit and run legacy Quick Queries and PQQs (created in the legacy web interface of Insights Platform version 9 and prior) in TRS and interact with the resulting data.

Legacy Quick Queries/PQQs are identified by the  icon in the Object Manager.

You can now open a legacy PQQ from the Object Manager or TRS and edit it in TRS. Parameterized operations are converted into parameterized blocks in the timeline. Once you save the legacy PQQ in TRS, it becomes an Insights Platform version 10+ PQQ (identified by the  icon in the Object Manager). All steps of a PQQ are now loaded into TRS; it is no longer necessary to click **Import Hidden Actions**. (Non-parameterized timelines are still hidden initially.) See [Edit and convert a legacy Parameterized Quick Query](#) on page 402 for more information.

You can now run a legacy Quick Query/PQQ the same way as you would a version 10+ Quick Query/PQQ. See [Run a legacy Parameterized Quick Query](#) on page 407 for more information.

Edit and convert a legacy Parameterized Quick Query

Edit a legacy Parameterized Quick Query and convert it to a version 10+ Parameterized Quick Query.

You can now run a legacy Parameterized Quick Query the same as you would a version 10+ Parameterized Quick Query. However, if you need to make changes to your legacy PQQ, you can open and edit the PQQ, but you will then need to save it as a version 10+ PQQ. Compatibility Mode is no longer available as of version 15 of the Insights Platform.

To edit and convert a legacy Parameterized Quick Query:

1. Open the **Object Manager** and locate the legacy Parameterized Quick Query you want to edit. The legacy PQQs are identified by the  icon.



Note: If you use a Parameterized Quick Query often, you may want to add it to your favorites. You can then find it quickly and easily in the **Favorites** tab on the **Dashboard**.

2. Right-click the legacy Parameterized Quick Query and choose **Open in TRS** from the menu.

Note: Depending on your workspace association settings, you may also double-click a legacy Parameterized Quick Query to open it. For instruction on setting your workspace associations, see [Define object associations](#) on page 43.

Alternatively, you can launch a new **TRS** window and double-click the legacy Parameterized Quick Query within TRS.

The legacy Parameterized Quick Query timeline appears in TRS.

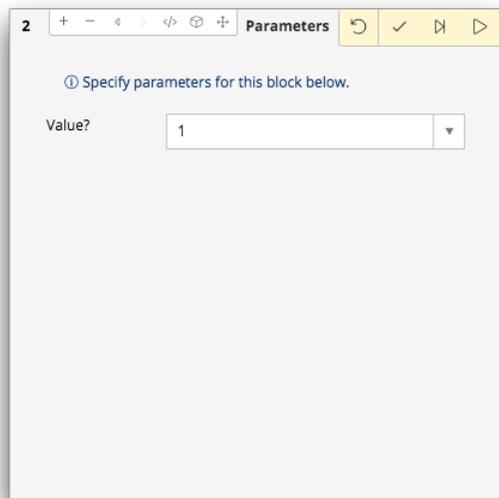
The screenshot shows the 1010data Insights Platform interface. The main window displays a 'Sales Detail' table with columns: Trans ID, Date, Time, Store, SKU, Extended Sales, Qty/Wgt, and Promo. The table contains 18 rows of data. On the left, a 'Parameters' panel is open, showing a 'Select a Store Number' input field. The interface includes a top navigation bar with 'Analyze', 'Query', 'View', 'Visualize', 'Develop', and 'Export' tabs. A sidebar on the right contains various tool icons.

Trans ID	Date	Time	Store	SKU	Extended Sales	Qty/Wgt	Promo
1400 >	01/01/14	00:00:00	11	242846	2.31	1.00	
1400 >	01/01/14	00:00:00	11	105755	3.75	1.00	
1400 >	01/01/14	00:00:00	11	449508	1.37	1.00	
1400 >	01/01/14	00:00:00	11	247560	1.34	2.20	
1400 >	01/01/14	00:00:00	11	118708	1.13	1.00	
1400 >	01/01/14	00:00:00	11	405883	3.15	2.00	
1400 >	01/01/14	00:01:15	11	395395	2.82	1.00	
1400 >	01/01/14	00:04:11	11	104639	1.85	1.00	
1400 >	01/01/14	00:04:11	11	329421	0.1	1.00	
1400 >	01/01/14	00:04:11	11	480160	1.13	1.00	
1400 >	01/01/14	00:04:11	11	472981	3.88	1.00	
1400 >	01/01/14	00:04:11	11	513835	0.1	1.00	
1400 >	01/01/14	00:04:11	11	141231	9.09	1.00	
1400 >	01/01/14	00:04:11	11	220697	1.86	1.00	
1400 >	01/01/14	00:04:11	11	127509	0.14	3.00	
1400 >	01/01/14	00:04:11	11	177084	3.72	2.00	
1400 >	01/01/14	00:24:04	11	268506	2.68	1.00	

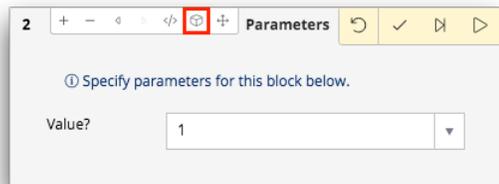
3. To edit operations that are not parameterized selections, do the following for each operation:
- Click the operation to edit.
 - Make the necessary changes.

The screenshot shows the 'Parameters' panel for a query operation. It includes a 'Select' dropdown menu, a 'Type "=" to enter a value expression directly.' instruction, and two checked options: 'Simple comparisons' and 'Look up column values'. The panel contains two comparison terms: 'Date ≥ 20150101' and 'Date ≤ 20150131'. At the bottom, there are buttons for '+ Add AND Term', '- Remove Term', and 'Negate'.

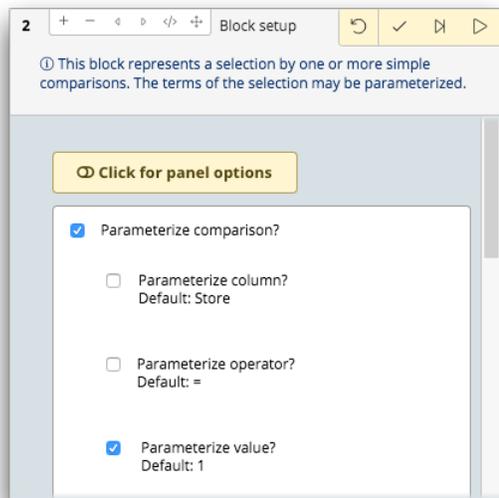
- Click the **Submit operation** (✓) icon.
4. To edit parameterized selections, do the following for each parameterized selection:
- In the timeline, click the panel of the operation containing the parameters you want to edit. The Trillion-Row Spreadsheet displays the **Parameters** view of the operation panel.



- b) Click the **Edit block** (🛠️) icon.



The Trillion-Row Spreadsheet displays the **Block setup** view of the operation panel.

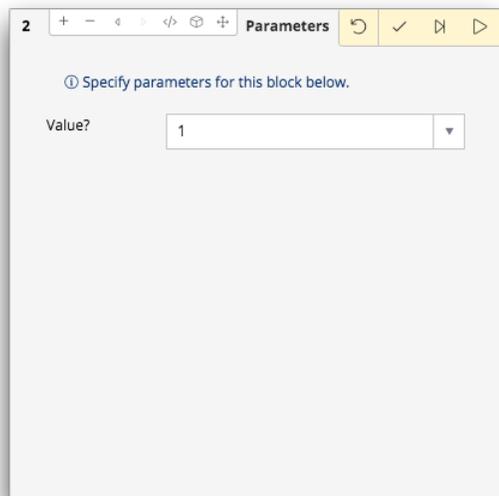


- c) Edit the appropriate fields and options for each parameter that you want to update.

Note: Make sure there is a default value for the parameter, unless you select **Allow this selection to be bypassed**. This option allows for an empty selection.

See [Block setup panel](#) on page 385 for more information.

- d) Click the **Submit operation** (✅) icon.
The Trillion-Row Spreadsheet saves the changes and displays the **Parameters** view of the operation panel.



After editing the desired parameters in the analysis, you can save it as a new Parameterized Quick Query or replace the existing one. For instructions, see [Save an analysis as a Quick Query](#) on page 281. The Parameterized Quick Query will be saved as a version 10+ PQQ (with a  icon). Test your converted PQQ by running it from the Object Manager. See [Run a Parameterized Quick Query](#) on page 400 for more information.

Edit a legacy Parameterized Quick Query in MLW

You can open and edit a legacy PQQ in Macro Language Workshop.

There are cases where you will want to edit a legacy Quick Query or PQQ in Macro Language Workshop. However, you can add a new parameterization to your PQQ only in TRS.

To edit a legacy PQQ in Macro Language Workshop:

1. Open the **Object Manager** and locate the legacy Parameterized Quick Query.
2. Right-click the legacy Quick Query or Parameterized Quick Query and choose **Edit query in Macro Language Workshop** from the menu.

The 1010data Insights Platform opens the legacy Parameterized Quick Query in the Macro Language Workshop.

```

1 <base table="pub.doc.retail.salesdetail" />
2 <sel value="(trans_date='01/01/15')&(trans_date<='01/31/15')" />
3 <sel value="(store=22) [[ \ \ mult_optout:Select a Store Number]]" />

```

Note: In MLW, the legacy PQQs retain the same `[[\ \ ...]]` syntax as that used in **Edit Actions XML** in the legacy interface.

3. Perform edits to your query.

Note: You can edit your legacy Quick Query or PQQ in MLW, but you cannot parameterize a selection in MLW. For that, you must use the timeline in TRS. See [Add parameters to an operation](#) on page 387 for more information.

In this simple example, we changed the date selection, the default store, and the selection prompt in the PQQ.

MLW: uploads.t449038101 - Legacy PQQ Example

```

File Code Run Debug render
1 <base table="pub.doc.retail.salesdetail"/>
2 <sel value="(trans_date>='01/01/14')&
  (trans_date<='01/31/14')"/>
3 <sel value="(store=11) [[ \ \ mult_optout:Select
  Store Number(s)]] "/>

```

- Click **render** to run the query.
The query results appear in the right panel of MLW.

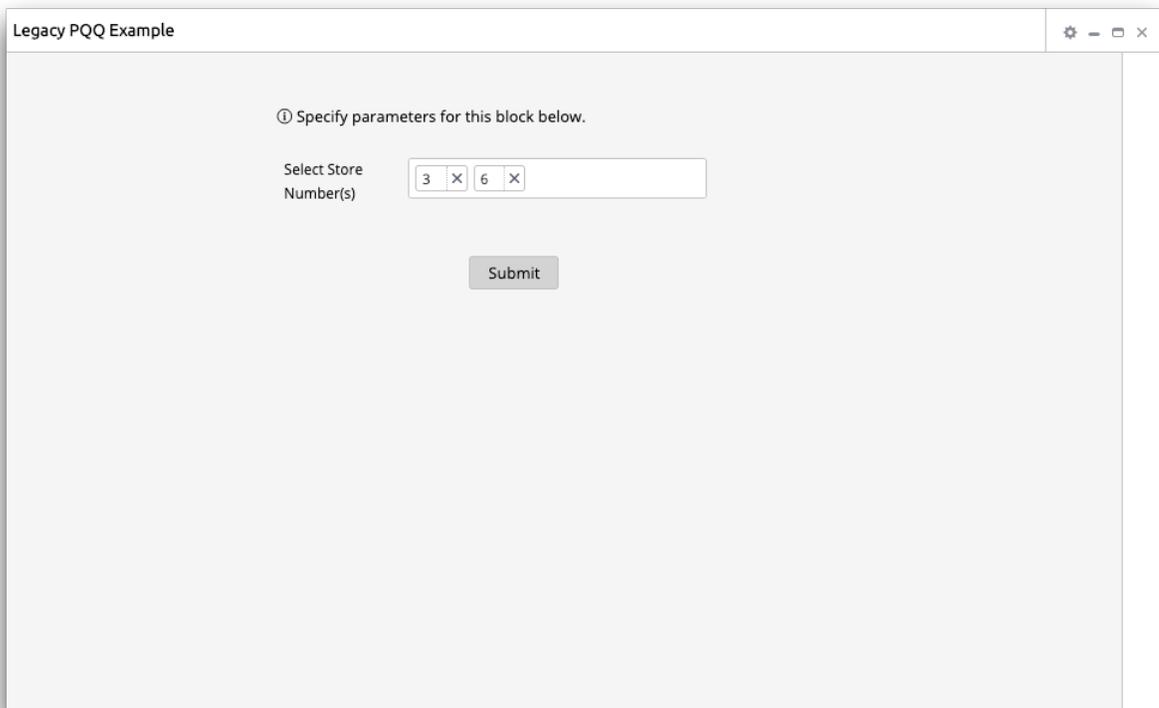
Cols 1 fixed and 2 to 7 of 16 , Rows 1 to 18 of 238,470

Trans ID	Date	Time	Store	SKU	Extended Sales	Qty/Wgt
1400 >	01/01/14	00:00:00	11	242846	2.31	1.0
1400 >	01/01/14	00:00:00	11	105755	3.75	1.0
1400 >	01/01/14	00:00:00	11	449508	1.37	1.0
1400 >	01/01/14	00:00:00	11	247560	1.34	2.2
1400 >	01/01/14	00:00:00	11	118708	1.13	1.0
1400 >	01/01/14	00:00:00	11	405883	3.15	2.0
1400 >	01/01/14	00:01:15	11	395395	2.82	1.0
1400 >	01/01/14	00:04:11	11	104639	1.85	1.0
1400 >	01/01/14	00:04:11	11	329421	0.1	1.0
1400 >	01/01/14	00:04:11	11	480160	1.13	1.0
1400 >	01/01/14	00:04:11	11	472981	3.88	1.0

- Click **File > Save Query** to overwrite the existing PQQ or **File > Save Query As** to save as a new PQQ.

Note: If you edit a legacy Quick Query in MLW and either overwrite the existing query or save as a new query, the Quick Query will remain in the legacy format and have the  icon. To convert your legacy PQQ to a version 10+ PQQ, you must open and save the PQQ in TRS. See [Edit and convert a legacy Parameterized Quick Query](#) on page 402 for more information.

- To test your changes, right-click your legacy PQQ in the Object Manager and choose **Run query in TRS**.



Run a legacy Parameterized Quick Query

Run a legacy Query Query/Parameterized Quick Query to view the results of a set of saved actions with user inputs.

To run a legacy Parameterized Quick Query:

1. Open the **Object Manager** and locate the legacy Parameterized Quick Query you want to run. The legacy PQQs are identified by the  icon.



Note: If you use a Parameterized Quick Query often, you may want to add it to your favorites. You can then find it quickly and easily in the **Favorites** tab on the **Dashboard**.

2. Right-click the legacy Parameterized Quick Query and choose **Run query in TRS** from the menu.

Note: Depending on your workspace association settings, you may also double-click a legacy Parameterized Quick Query to run it. For instruction on setting your workspace associations, see [Define object associations](#) on page 43.

The Insights Platform displays the Parameterized Quick Query parameter input fields and options in a new window.

Legacy PQQ Example

Specify parameters for this block below.

Select a Store Number

Submit

3. Select or enter any parameter inputs that the legacy Parameterized Quick Query contains.
 4. Click **Submit**.
- The Insights Platform runs the legacy Parameterized Quick Query and displays the results.

Legacy PQQ Example

Cols 1 fixed and 2 to 12 of 16, Rows 1 to 19 of 236,504

Trans ID	Date	Time	Store	SKU	Extended Sales	Qty/Wgt	Promo	Cost	Customer	Department	Group
1400 >	01/01/14	00:00:00	4	283807	2.68	1.00	0	2.68		19	221
1400 >	01/01/14	00:00:00	4	186243	0.05	1.00	0	0.05		19	46
1400 >	01/01/14	00:00:00	4	51199	2.44	1.00	0	2.44		30	137
1400 >	01/01/14	00:01:31	4	390955	2.44	1.00	0	1.96		30	137
1400 >	01/01/14	00:01:31	4	293366	0.86	1.00	0	0.71		30	137
1400 >	01/01/14	00:01:31	4	354905	4.38	1.00	0	3.49		30	200
1400 >	01/01/14	00:01:31	4	30631	1.04	1.00	0	0.79		30	200
1400 >	01/01/14	00:01:31	4	288641	1.1	1.00	0	1.1		20	448
1400 >	01/01/14	00:01:31	4	140199	0.08	1.00	0	0.08		20	448
1400 >	01/01/14	00:03:33	4	163050	21 >	1.00	0	19 >		42	23
1400 >	01/01/14	00:03:33	4	376756	1.1	1.00	0	1.1		42	23
1400 >	01/01/14	00:08:13	4	37576	11 >	1.00	0	10 >		42	23
1400 >	01/01/14	00:08:13	4	195389	0.55	1.00	0	0.55		42	23
1400 >	01/01/14	00:15:55	4	431196	1.83	1.00	0	1.01		19	27
1400 >	01/01/14	00:15:55	4	513835	0.1	1.00	0	0.1		19	46
1400 >	01/01/14	00:27:51	4	165413	1.62	1.00	0	1.62		19	27
1400 >	01/01/14	00:27:51	4	186243	0.05	1.00	0	0.05		19	46
1400 >	01/01/14	00:27:51	4	407358	0.84	2.00	0	1.52		36	341

The view bar on the right side of the results pane offers the following options for interacting with your results:

- Click the **Go Back** (↶) icon to specify different parameters in the PQQ and view the results.
- View the results in grid (📊) or single-row (📄) view.
- Download the results as a .csv (📄), .xlsx (📄), or .pdf (📄) file.
- Open a full TRS timeline (☰) and perform further analysis on your data.

Note: Running a legacy PQQ does not convert the PQQ to a version 10+ PQQ. If desired, you can edit and save the legacy PQQ as a version 10+ PQQ. See [Edit and convert a legacy Parameterized Quick Query](#) on page 402.

Macro Language Workshop

The Macro Language Workshop offers a rich set of tools for writing, executing, and debugging 1010data Insights Platform Macro Language queries and QuickApps in the new web interface.

The Macro Language Workshop allows you to write Insights Platform queries and QuickApps and render the results.

The Macro Language Workshop is accessed by selecting **Tools > Macro Language Workshop** from the workspace menu.

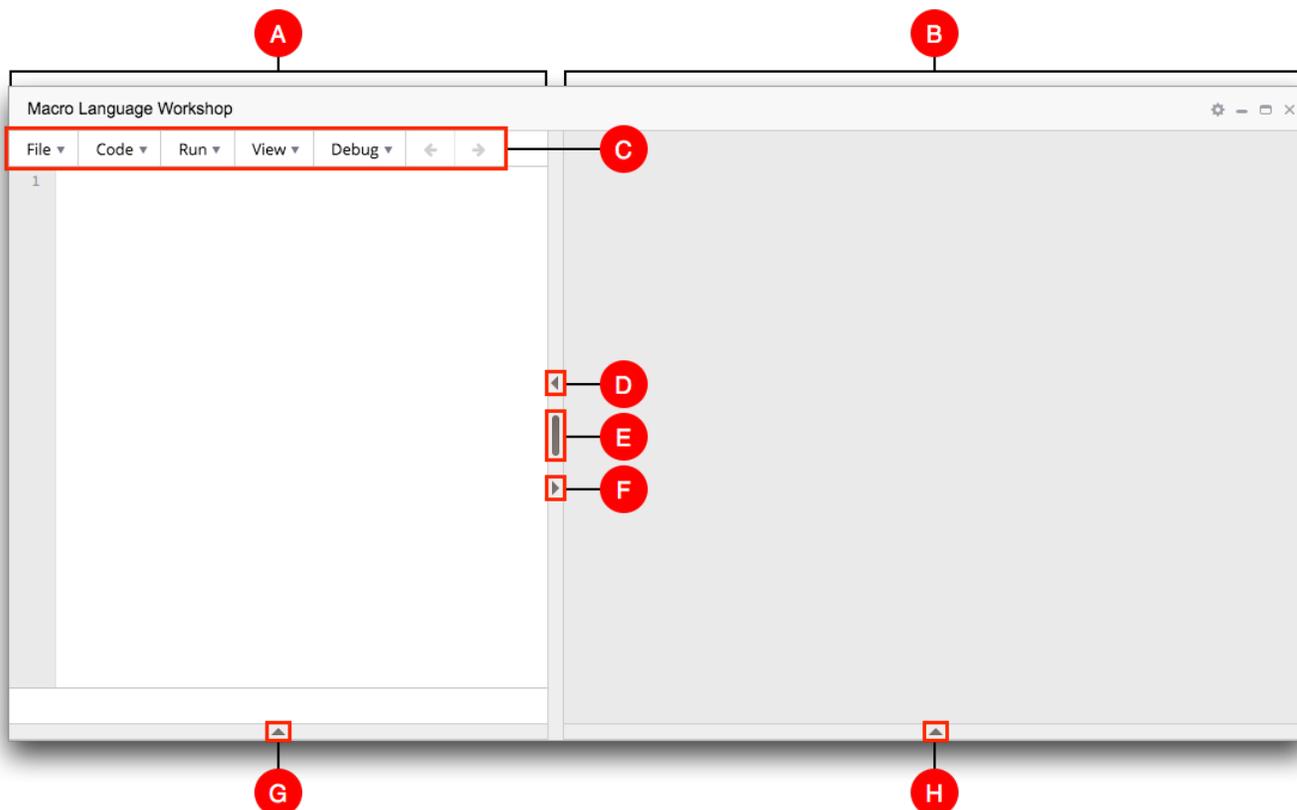


Figure 106: Macro Language Workshop window

A. Analysis pane

The analysis pane occupies the left side of the **Macro Language Workshop** window. This pane is used to enter and modify Macro Language XML code.

B. Results pane

The results pane is displayed on the right side of the **Macro Language Workshop** window. This pane displays the results of your Macro Language XML code.

The results pane of TRS and MLW are very similar. Both allow you to manipulate the results in grid view. See [Results pane](#) on page 291 for a description of how to analyze your results after you run your code.

C. Menu bar

The Macro Language Workshop menu bar provides commands for interacting with the Macro Language XML code.

D. Hide

When both the analysis pane and results pane are visible, click the **Hide** (▲) icon to hide the analysis pane. This is helpful when you need more space to view your results in the results pane. Click the **Show** (▼) icon to reveal the analysis pane when it is hidden.

E. Split bar handle

The split bar handle indicates that the analysis and results panes are resizable. Click anywhere on the split bar and drag it to resize the panes.

F. Show

When both the analysis pane and results pane are visible, click the **Show** (▼) icon to hide the results pane. This is helpful when you need more space to view the Macro Language XML code in the analysis pane. Click the **Hide** (▲) icon to reveal the results pane when it is hidden.

G. Expand/Collapse

Click the **Expand** (▲) icon to display the reference pane. When expanded, a split bar handle appears indicating that the analysis and reference panes are resizable. Click anywhere on the split bar and drag it to resize the panes. Click the **Collapse** (▼) icon to hide the information pane.

Note: The reference pane has a minimum height beyond which the split bar cannot be dragged.

H. Expand/Collapse

Click the **Expand** (▲) icon to display the monitor pane. When expanded, a split bar handle appears indicating that the results and monitor panes are resizable. Click anywhere on the split bar and drag it to resize the panes. Click the **Collapse** (▼) icon to hide the monitor pane.

Note: The monitor pane has a minimum height beyond which the split bar cannot be dragged.

Menu bar

The Macro Language Workshop menu bar provides access to the features and functionality available to you in the Macro Language Workshop.

Located at the top of the analysis pane in the Macro Language Workshop, the menu bar provides options for interacting with the Macro Language XML code. Additionally, the history icons allow you to step backward and forward through changes to the code.

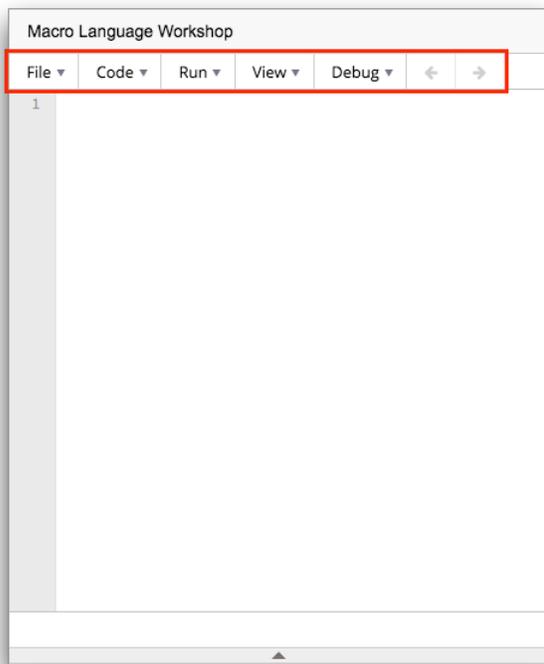


Figure 107: Menu bar

File

This menu allows you to perform actions such as loading, saving, and exporting tables and queries.

Code

This menu provides access to tools used to review and modify Macro Language XML code. In addition, you can upload and download query text and access context-sensitive help.

Run

This menu allows you to run the Macro Language code in the code editor and select from various ways to display the results.

Debug

This menu provides tools to help you manage problems and identify issues in the Macro Language code.

History icons

The history icons allow you to step backward (←) and forward (→) through changes made to the Macro Language code.

File

You can perform actions such as loading, saving, and exporting tables and queries from this menu.

The **File** menu contains the following menu items:

Load

This menu item allows you to select a base table for use in the code editor or to load the Macro Language XML code of a saved object.

When a table is selected, it is used as the base table in the code editor.

When a saved object is selected, you choose whether to use it as the base table in the code editor or to load the Macro Language code of the object. If the saved object uses a query as the base table, and you choose to load the object, the operations of the saved object are displayed in the code editor and a `<base>` operation pointing to the original base query is inserted before all other operations.

Load resolving all ops

Note: Prior to version 11.27, the **Load resolving all ops** menu option was named **Load in full**.

This menu item allows you to select a base table for use in the code editor or to load the Macro Language XML code of a saved object.

When a table is selected, it is used as the base table in the code editor.

When a saved object is selected, you choose whether to use it as the base table in the code editor or to load the Macro Language code of the object. If the saved object uses a query as the base table, and you choose to load the object, one of the following occurs:

- If the query was saved in the legacy 1010data Insights Platform and does not contain a `<base>` operation, the operations of the base query in addition to the operations of the saved object are displayed in the code editor.
- If the query was saved in version 10 or later, the operations of the saved object are displayed in the code editor and a `<base>` operation pointing to the original base query is inserted before all other operations.

Note: This menu item discards saved text source. The text source includes any extra white space such as line breaks or indentations added in the Macro Language code.

Convert Quick Query

Converts a Quick Query to a QuickApp.

While you can use this menu item to convert any Quick Query to a QuickApp, it is most useful in converting a Parameterized Quick Query into a QuickApp. A Parameterized Quick Query is a type of Quick Query available in the legacy version of the Insights Platform that is not supported beginning with version 10. To retain the functionality of a Parameterized Quick Query, you must convert it to a QuickApp.

Save Query

Saves changes made to the Macro Language code of an existing object.

Save Query as

Displays a dialog which allows you to save a new object, make and save changes to the options of an existing object, or save a copy of an existing object. The saved object becomes the new associated file in the **MLW** window.

Save Data

Saves the data resulting from this Macro Language code as a table or as a text file for your 1010data FTP directory.

This submenu contains the following menu items:

as table

Opens the **Save Data as Table** window which allows you to save the data resulting from your analysis to an Insights Platform table.

For more information, see [Save Data as Table](#) on page 377.

to FTP

Opens the **Save Data to FTP** window which allows you to export the results of the query as a compressed text file to your 1010data FTP directory.

For more information, see [Save Data to FTP](#) on page 370.

Filters

(Available as of version 11.27)

Note: Prior to version 11.27, the filter menu options were located in the **File** menu instead of the **File > Filters** menu.

This menu item allows you to load and save table filters. A filter is a query that runs at the time a table is loaded. For example, filters saved to a table can be used to limit the table data shown to a specific user or group.

This submenu contains the following menu items:

Load filter

Displays the Macro Language code of any filters that are saved to a selected table. You can edit, remove, or save the filter to another table.

Save filter

This menu item allows you to save a new filter to a table or save changes made to an existing filter.

Revert

This menu item restores the Macro Language code to its most recently saved state. Revert undoes all code changes at once compared to one at a time using the history icons. Revert is available after opening a saved object and making changes in the code editor.

Export

(Available as of version 11.27)

Note: Prior to version 11.27, the export menu options were located in the **File** menu instead of the **File > Export** menu.

You can export the results of your query for further development and analysis, to create a new table, or to extract data from the Insights Platform for use elsewhere.

This submenu contains the following menu items:

to new TRS

Exports the Macro Language code to a new **TRS** window.

to new MLW

Exports the Macro Language code to a new **MLW** window.

(Available as of version 11.27)

and run selection in MLW

Exports Macro Language code selected (highlighted) in the code editor to a new **MLW** window and, if the selection is a valid query, runs the code immediately using the **Run > Render** menu option.

Selected code can also be exported to a new **MLW** window using the following keyboard shortcuts:

- **Ctrl+Shift+Enter** (PC)
- **Shift+Command+Return** (Mac)

(Available as of version 11.27)

library to workspace

Exports library XML code displayed in the code editor and installs it into your current workspace. The exported XML should consist of `<library>` and/or `<import>` forms.

For example, after a library is installed in the workspace, the library can be used in other Macro Language Workshop windows without needing to import it in the code. This helps make the development of new components and tools in the Insights Platform more fluid.

(Available as of version 11.27)

to file

Opens the **Export Query As** window to export the Macro Language code results to one of the following file options: delimited file, Excel Worksheet, or portable document (PDF).

For more information about the settings for the various export options, see the following:

- [Delimited text export settings](#) on page 285
- [Excel workbook export settings](#) on page 287
- [PDF export settings](#) on page 288

(Available as of version 15.21)

Clear

Deletes all content in the MLW code editor.

Code

This menu provides access to tools used to review and modify Macro Language XML code. In addition, you can upload and download query text and access context-sensitive help.

This menu contains the following menu items:

Check and clean up XML

Analyzes the Macro Language XML code and lists the first syntax error that prevents the code from parsing. In addition, this option restores the default XML formatting of the code. For example, extra spacing is removed and indentation levels are adjusted.

Fix smart punctuation

Replaces all smart (curly) quotation marks with straight quotation marks throughout the Macro Language code.

TRS post-processing

(Available as of version 11.27)

Note: Prior to version 11.27, the **Remove decorations** menu option was named **Remove TRS decorations** and was located in the **Code** menu instead of the **Code > TRS post-processing** menu.

This menu allows you to remove or enable the XML code specific to queries created in the Trillion-Row Spreadsheet.

This submenu contains the following menu items:

Remove decorations

Removes TRS-specific attributes from the Macro Language code.

When an analysis is created in the Trillion-Row Spreadsheet, specific attributes are generated and added to the underlying XML code. For example, the `trs_timeline_="1"` attribute indicates that the operation was created in the TRS and is part of the timeline. These attributes are not displayed in the Trillion-Row Spreadsheet and are only visible when an analysis created in TRS is opened in the Macro Language Workshop.

Enable all operations

Restores ignored TRS timeline operations in the Macro Language code.

When an analysis in the Trillion-Row Spreadsheet is sent to the Macro Language Workshop, operations that follow the panel selected in the timeline are ignored in the Macro Language code. This is indicated in the code by the `<ignore>` tag surrounding an ignored operation. When the last panel in the timeline is selected at the time the TRS analysis is sent to the Macro Language Workshop, no operations are ignored.

(Available as of version 11.27)

Expand in place

Expands `<block>` code in the code editor. Expanding the block replaces variables in the code with values.

Get query hashes

Generates a query hash (`@queryhash_`) and a special hash (`@specialhash_`) for the query in the code editor.

Hashing is the transformation of a string of characters into a usually shorter fixed-length value or key that represents the original string. In the 1010data Insights Platform, query hashes are commonly used to filter tables and encrypt or identify values.

A generated query hash is based on every character in the code of the query where a special hash is generated using only import statements and the parameter names (not values) in the query.

Check query compatibility

Evaluates expressions in the query for operator precedence issues that may affect cross-version compatibility.

In version 10 of the Insights Platform, a more traditional operator precedence for interpreting expressions was implemented. If a query contains an ambiguous expression that could be interpreted differently between the current operator precedence and the old precedence, the platform returns an error. This menu item checks for those differences.

Download query text

Saves the contents of the code editor to your computer as a text file.

(Available as of version 11.27)

Upload query text

Uploads a text file from your computer into the Macro Language Workshop code editor.

(Available as of version 11.27)

Context-sensitive code help

Displays the reference pane below the analysis pane. The reference pane provides context-sensitive help for the Macro Language XML code that is selected (highlighted) in the code editor. In addition, the reference pane provides click-to-insert access to objects such as folders, tables, and columns, functions, and Macro Language tags.

The reference pane can also be opened by using the following keyboard shortcuts:

- **Ctrl+Shift+?** (PC)
- **Shift+Command+?** (Mac)

(Available as of version 11.27)

Run

This menu allows you to run the Macro Language XML code in the code editor and select from various ways to display the results.

Render

Runs the query and displays the results in the results pane. This menu option allows for full interaction with any open tool or window in your session.

Note: If your Macro Language XML code contains global attributes, rendering the code with this option applies those attributes to your entire session.

Render in <iframe>

Runs the query and displays the results in an inline frame within the results pane.

This menu option prevents global attributes in the Macro Language code from applying to your entire session. An inline frame is helpful when developing a QuickApp because it creates a standalone like environment for the application.

Render as window in workspace

Runs the query and displays the results within a separate window in your workspace.

Other than the location of the displayed results, this menu item functions identically to **Run > Render**.

(Available as of version 11.27)

Render standalone

Runs the query and displays the results in a new browser tab.

Other than the location of the displayed results, this menu item functions identically to **Run > Render in <iframe>**.

Render to static HTML target

Runs the query and displays the results as static HTML in a new browser tab.

(Available as of version 11.27)

Render to PDF target

Runs the query and saves the results to your computer as `download.pdf`.

Render to XLSX target

Runs the query and saves the results to your computer as `download.xlsx`.

Render to data target

Runs the query and saves the results to your computer as `download.csv`.

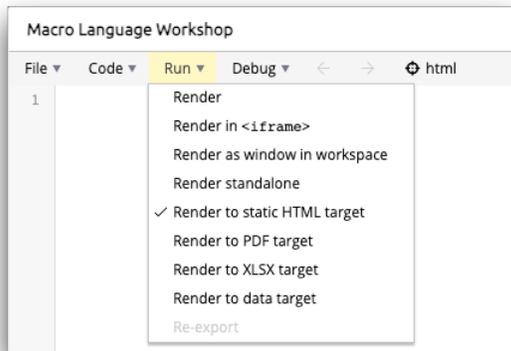
Re-export

Sends the Macro Language code back to the **TRS** window from which it originated.

This menu item is only available in the **Editing macro from TRS** window. This window opens after a query in the Trillion-Row Spreadsheet is sent to the Macro Language Workshop. For instructions, see [Edit a query in the Macro Language Workshop](#) on page 280.

Once you select **Re-export** explicitly from the menu, you can re-export again by clicking the **re-export** button or by using the keyboard shortcut **Ctrl+Enter (Return)**.

Insights Platform places a check mark next to the last-used render target in the **Run** menu for reference.



You can also run the Macro Language XML code in the code editor using the following keyboard shortcuts:

- **Ctrl+Enter** (PC)
- **Command+Return** (Mac)
- **Control+Return** (Mac)

Note: These keyboard shortcuts executes the last render command selected from the **Run** menu, indicated with a check mark.

If a render command has not yet been selected, these keyboard shortcuts are the same as selecting **Run > Render** from the menu.

Debug

This menu provides tools to help you manage problems and identify issues in the Macro Language XML code.

This menu contains the following menu items:

Debug expansion

Displays debugging controls and information in the results pane. The top portion of the results pane provides controls for stepping through your code. The bottom portion of the results pane lists information about the code as it is expanded.

Remove <dynamic> state

Removes the <dynamic> state from your session.

The <dynamic> state is added to your session when a QuickApp is rendered.

Clear all caches

Clearing the cache cleans out your virtual memory so that you have the maximum amount available to your session.

History icons

The history icons allow you to step backward and forward through changes made to the Macro Language code.

Go back

Click the **Go back** (↶) icon to revert the code back to the previous state before it was rendered. This icon is available after making a change to the code and rendering the results.

Go forward

Click the **Go forward** (→) icon to restore undone changes. This icon is available after clicking **Go back** to revert the code to previous state.

Open the Macro Language Workshop

The Macro Language Workshop in the new Insights Platform web interface is easily accessed via the **Tools** menu.

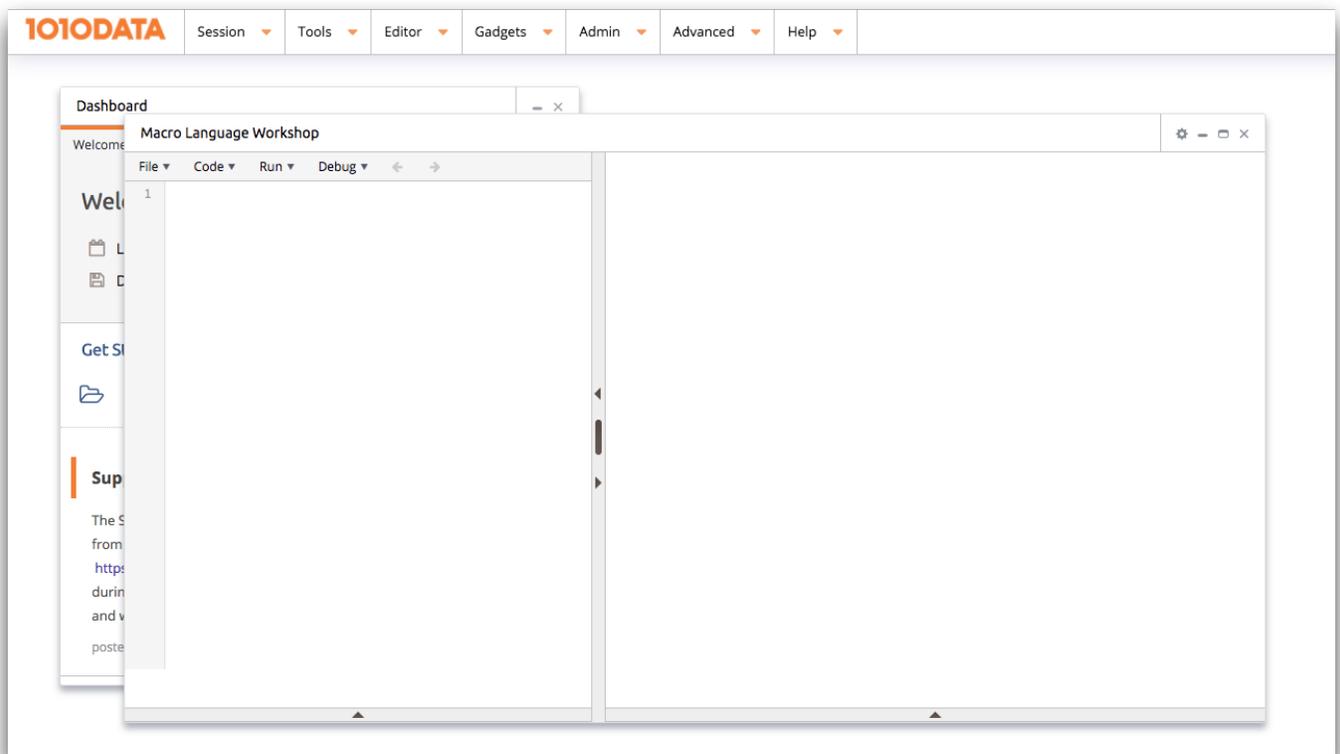
Log in to the 1010data Insights Platform.

Use the **Tools** menu to open the Macro Language Workshop.

To open the Macro Language Workshop:

In the workspace menu, click **Tools > Macro Language Workshop**.

A **Macro Language Workshop** window opens in the current workspace.



Specify a base table for the current **Macro Language Workshop** window in the left pane.

Specify a base table

It is necessary to specify a base table when writing a new query in the Macro Language Workshop.

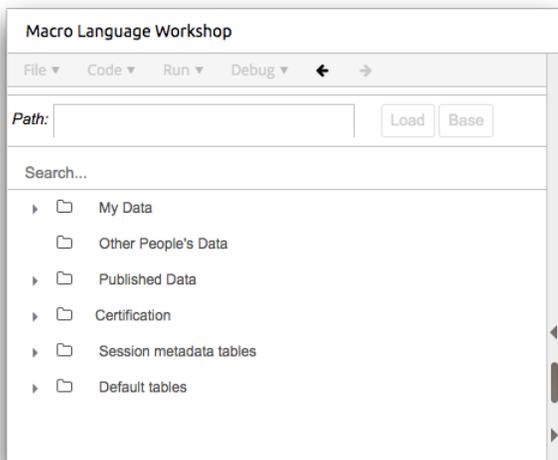
Log in to the 1010data Insights Platform and open a **Macro Language Workshop** window.

After opening a **Macro Language Workshop** window, it is necessary to specify a base table for the query. While there are several methods for completing this task, the most straightforward is to use the **Load** option in the **File** menu.

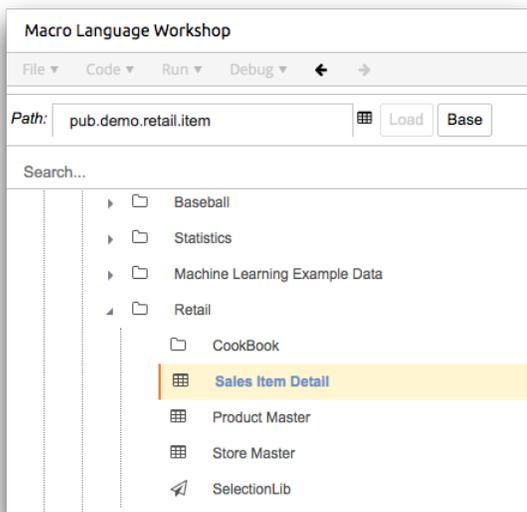
To specify a base table:

1. Click **File > Load** .

The left pane of the **Macro Language Workshop** displays a file browser.



2. Browse to the table that will be used as the base table for the query and click it. The table object will be highlighted and the **Base** button will be available.



3. Click **Base**.

The following code is entered in the left pane of the **Macro Language Workshop**.

```
<base table="pub.demo.retail.item"/>
```

Begin writing an Insights Platform query in the left pane of the **Macro Language Workshop** by adding operations under the `<base>` operation.

Open a saved query

Queries saved in the 1010data Insights Platform can be opened directly from the **Macro Language Workshop**.

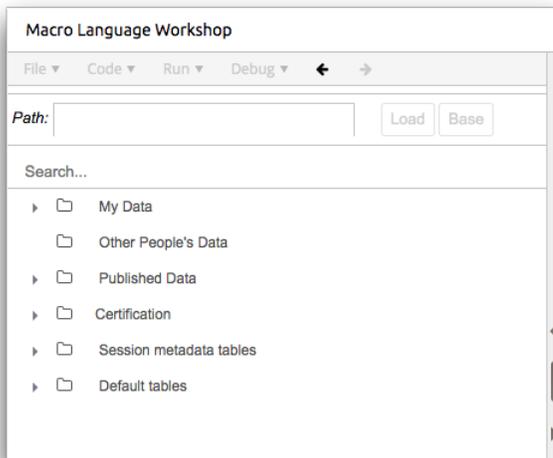
Log in to the Insights Platform and open a **Macro Language Workshop** window.

A query can be saved in the Insights Platform and accessed as an object via the **Object Manager**, the **Trillion-Row Spreadsheet**, and the **Macro Language Workshop**.

To save a query:

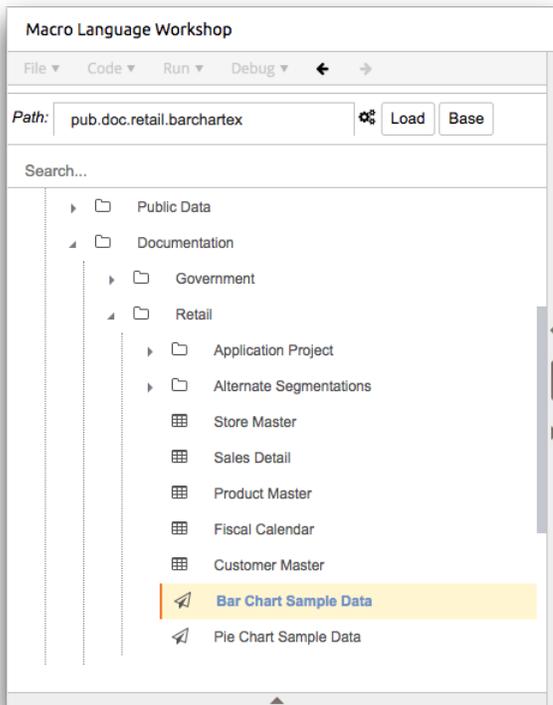
1. Click **File > Load**.

The left pane of the **Macro Language Workshop** displays a file browser.



2. Browse to the saved query and click it.

The query object is highlighted and the **Load** and **Base** buttons are available.



3. Open the saved query by doing one of the following:

- Click **Load** to open the XML for the query in the left pane of the **Macro Language Workshop**. The full text of the query will be available in the editor.
- Click **Base** to use the results of the saved query as a base table. If this option is selected the full XML text of the query will not be shown in the left pane of the **Macro Language Workshop**.

The saved query is loaded in the **Macro Language Workshop** and can be modified.

Run the query.

Render a query

To see the results of a query, it is necessary to select a render option.

Log in to the 1010data Insights Platform and open a **Macro Language Workshop** window. Then, specify a base table, open a saved query, or write a new Insights Platform query in the current **Macro Language Workshop** window.

To render a query:

Click **Run** on the menu bar.

- Click **Render** to display the result of the current query in the right pane of the current **Macro Language Workshop** window.
- Click **Render in <i>iFrame</i>** to display the result of the current query in the right pane of the current **Macro Language Workshop** window as part of an iFrame. This option is important for QuickApps that make use of themes.
- Click **Render standalone** to display the result of the current query as a standalone QuickApp in a new browser tab.
- Click **Render to XLSX target** to download the result of the current query to a Microsoft Excel file.
- Click **Render to data target** to download the result of the current query as a comma-separated text file.

The results of the query that was submitted will be rendered as follows:

Render

The query results are displayed in the right pane.

Render in <i>iFrame</i>

The query results are displayed in the right pane.

Render standalone

The query results are rendered as a standalone QuickApp in a new browser tab.

Render to XLSX target

The query results are rendered and downloaded as a Microsoft Excel workbook.

Render data target

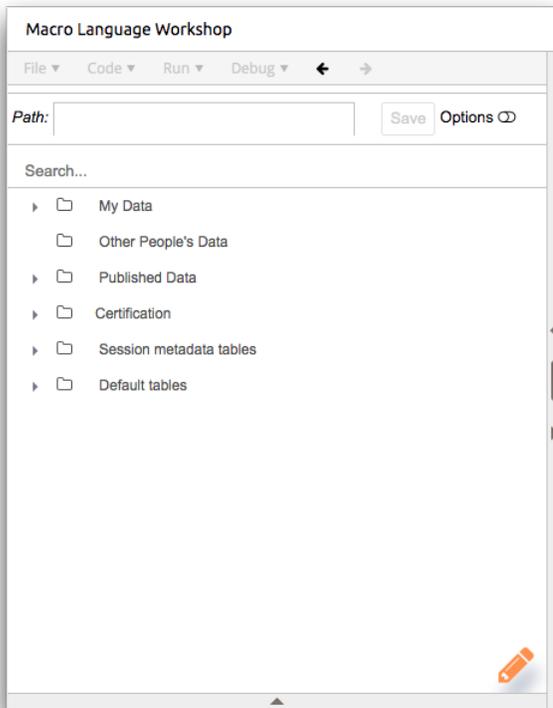
The query results are rendered and downloaded as a delimited text file.

Save a query

1010data Insights Platform Macro Language queries can be saved directly from the Macro Language Workshop. Follow the steps in this task to save a query.

Login to the Insights Platform and open a **Macro Language Workshop** window. Then, write a query or open an existing query and modify it.

1. On the menu bar above the left pane of the active **Macro Language Workshop** window, click **File > Save as**.
The left pane of the Macro Language Workshop displays an object browser for saving the current query:



Note: In order to save a query, you must have upload permissions to the folder where you want to save the query.

2. Browse to folder where the query should be saved and single-click it.

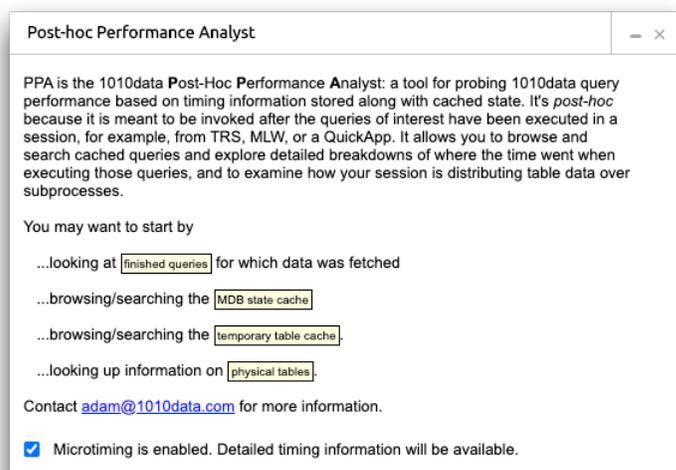
Post-hoc Performance Analyst

The Post-hoc Performance Analyst (PPA) is a tool for probing 1010data query performance *after* the queries of interest have been executed in a session.

The PPA takes advantage of the timing information stored along with the cached state of queries. The analysis is *post-hoc* because it is meant to be invoked after the queries of interest have been executed in a session, for example, from TRS, MLW, or a QuickApp. It allows you to browse and search cached queries and explore detailed breakdowns of where the time went when executing those queries, and to examine how your session is distributing table data over subprocesses.

The Post-hoc Performance Analyst is accessed by selecting **Tools > Post-hoc Performance Analyst** from the workspace menu.

Figure 108: Post-hoc Performance Analyst



Finished queries

Click here to view finished queries for which data was fetched.

MDB state cache

Click here to browse the cache in which MDB stores completed queries.

Temporary table cache

Click here to browse the cache in which the MDB stores temporary in-memory tables, such as worksheets in links.

Physical tables

Click here to view information about a specific physical table.

Microtiming is enabled

Select this to display detailed timing information.

Finished queries

The **Finished queries** view of the PPA allows you to view performance statistics associated with the completed queries from which data was fetched.

You can also access the **Finished queries** view of the PPA from TRS or from MLW. After you run a query in TRS or MLW, click the **Performance** (🕒) icon to view performance information for your query.

Trans ID	Date	Time	Store	SKU	Extended Sales	Qty/Wgt	Promo
1400 >	01/01/14	00:00:00	15	247560	0.44	0.53	
1400 >	01/01/14	00:00:00	15	456371	0.64	0.72	
1400 >	01/01/14	00:00:00	11	242846	2.31	1.00	
1400 >	01/01/14	00:00:00	11	105755	3.75	1.00	
1400 >	01/01/14	00:00:00	11	449508	1.37	1.00	
1400 >	01/01/14	00:00:00	11	247560	1.34	2.20	
1400 >	01/01/14	00:00:00	11	118708	1.13		
1400 >	01/01/14	00:00:00	11	405883	3.15		

Figure 109: Finished queries

Post-hoc Performance Analyst

Help | Close view

Find by base table: pub.doc.retail.sale... | Total time threshold: (None) | Find by raw query text:

Displaying #2 of 2 queries found

Completed query

Physical table `pub.doc.retail.salesdetail`
 4,481,814 rows in 2 segments on 2 machines | Table info | Show data

```
<base table="pub.doc.retail.salesdetail"/>
<sel simple="1" value="(trans_date=20140101)"/>
```

Cached: 58 ms | Cached state | Get: 58 ms | Results | Performance | Total: 116 ms

Find by base table

Select a base table from the drop-down list.

Total time threshold

Select a query by time threshold. This is helpful if you want to focus on slower queries.

Find by raw query text

Enter text that is in the query.

Arrow keys

Browse through the queries that have completed.

Completed query

This section shows details about the completed query, including the following information.

- Whether the table is physical or temporary
- The name of the table
- The total number of rows of the table
- The number of segments (physical tables only)
- The number of machines (physical tables only)
- Click **Table info** for a more detailed breakdown of machine names and segments (physical tables only). See [Physical tables](#) on page 429 for more information.
- Click **Show data** for the actual table data. You can then open the table in TRS, if desired.
- The query itself
- Click **Cached state** to view the MDB state cache (ops yielding a cached state only). See [MDB state cache](#) on page 426 for more information.
- Click **Results** for detailed data retrieval statistics. You can export the result set to TRS.
- Click **Performance** for a performance details view. This view includes query coordination at accumulator and query execution performance. You can export raw performance data to TRS.

MDB state cache

The **MDB state cache** view of the PPA allows you to browse the cache in which MDB stores completed queries, along with the new database states that they engender, and performance information.

The queries in the cache reflect preprocessing that takes place before the database engine begins executing the query, and therefore queries may not appear in their original form, as submitted or embedded in an application. For example, a `<link>` to a worksheet is processed as two separate queries, one for the worksheet (inner query) and one for the query linking to it (outer query). The outer query will appear in the cache with a link to a *temporary table*, and the inner query will appear separately in the cache. Similarly, `<merge>` is decomposed into separate subqueries and the result is placed in a temp table, which may appear in the cache as the base for subsequent operations. Also, block code, and forms such as `<loop>`, `<do>`, and `<dynamic>` are processed before the database engine is invoked and therefore will not appear in the state cache, although queries run by these forms will.

The cache may be browsed in two modes. When **Telescope queries** is not selected, each *section* of a query that corresponds to a new database state will be displayed separately. For context, the *previously-cached* section(s), if any, for a given section will also be displayed, dimmed, before the section of interest. For example, a query consisting of only two selections and a tabulation would correspond to three separate entries in the cache: one with the `<base>` operation and the first `<sel>`, the second with the second `<sel>` (with the first two operations displayed dimmed), and the third with the `<tabu>` (with the first 3 operations displayed dimmed). Note that not every operation corresponds to an independent new state.

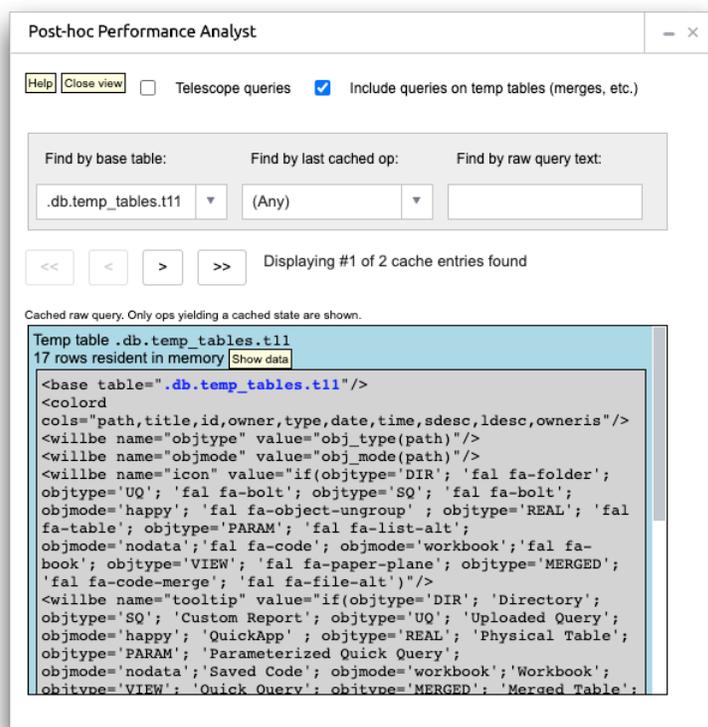
When **Telescope queries** is selected, all sections of a query are displayed at the same time, in separate boxes. In this mode, even if two separate queries were run, one of which was a identical to the beginning of the second, only the second (longer) query will be shown. This mode can reduce the number of cache entries that need to be reviewed. All data and performance information is available for every query section in either mode.

By default, cached queries based on in-memory temporary tables are included in the display. It is possible to eliminate these from the display by unchecking the checkbox. This will prevent, for example, operations subsequent to a `<merge>` operation from appearing. PPA automatically eliminates most, though not necessarily all, system queries (e.g. those issued by the GUI, by tools like the Object Manager, and by PPA itself) to avoid superfluous clutter. The displayed queries may be filtered on three criteria. The **Find by base table** drop-down list contains a list of all the paths that appear as base tables for the cached queries. Selecting one of these restricts the displayed queries to those based on the specified path. The **Find by last cached op** drop-down list contains a list of all "final" operations (e.g. `sel`, `tabu`, `sort`), i.e. those

appearing at the end of the cached queries. Selecting an operation restricts the displayed queries to those that end with the specified operation.

Important note: When **Telescope queries** is selected, this will only find the longest queries ending in the specified operation, not shorter initial portions of those queries. Depending on your objective, you may want to turn off telescope mode before using **Find by last cached op**. It is also possible to search all queries by any arbitrary text string (operations, column names, pathnames, attributes, etc.) by entering the text in the **Find by raw query text** field. Again, remember that cached queries are preprocessed; you won't find them by entering, for example, block code or tags like `<merge>` that appear in your original submitted query. Instead, try using expressions, column names, tables, etc. from your query. Use the navigation buttons `<<`, `<`, `>`, and `>>` ("first," "previous," "next," and "last," respectively) to browse through all queries found with the given filter criteria.

Figure 110: MDB state cache



Telescope queries

When selected, finds the longest queries ending in the specified operation, not shorter initial portions of those queries.

Include queries on temp tables (merges, etc.)

When selected, include both physical and temporary tables in your search.

Find by base table

Select a base table from the drop-down list.

Find by last cached op

Select a "final" operation (those appearing at the end of the cached queries) from the drop-down list. Examples include `sel`, `tabu`, and `sort`.

Find by raw query text

Enter text that is in the query.

Arrow keys

Browse through the cache entries.

Cached raw query

This section shows details about the cached raw query, including the following information.

- Whether the table is physical or temporary
- The name of the table
- The total number of rows of the table
- The number of segments (physical tables only)
- The number of machines (physical tables only)
- Click **Table info** for a more detailed breakdown of machine names and segments (physical tables only). See [Physical tables](#) on page 429 for more information.
- Click **Show data** to view the actual data that the query returns. You can then open the table in TRS, if desired.
- The cached query itself
- Click **Performance** for a performance details view. This view includes query coordination at accumulator and query execution performance. You can export raw performance data to TRS.
- The resulting number of rows

Temporary table cache

The **Temporary table cache** view of the PPA allows you to browse the cache in which MDB stores temporary in-memory tables representing, for example, worksheets in links, merged worksheets, and the results of various operations such as `<table>`, `<transpose>`, and `<loop>` (for some modes).

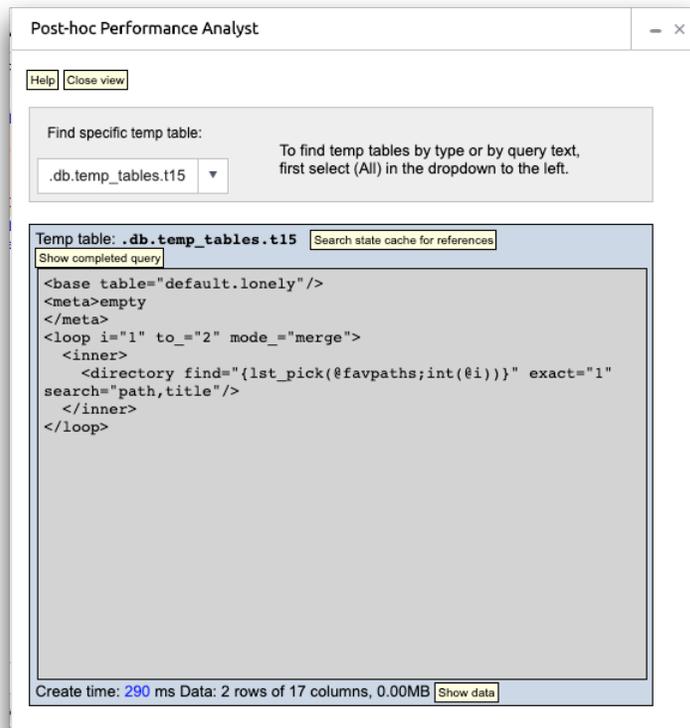
Temporary tables are designated in the system with a "path" (more accurately a "handle," as they may not be referenced directly from user-submitted queries) in the form `.db.temp_tables.tn` for integer $n \geq 0$. They may appear in cached MDB operations as the base table of a query or, for example, in the `table2=` attribute of a `<link>` operation. To search for queries in the MDB state cache which reference a temporary table displayed in the current view, click the **Search state cache for references** button. Conversely, temporary table handles are highlighted throughout the PPA (in the **MDB state cache** view and in the **Temporary table cache** view itself). To look up a handle, simply click the highlighted link.

If invoked via a link, the **Temporary table cache** view will show the operations for a specific temp table, along with the total time used to create the table and its size (rows, columns, and actual memory used in MB). The **Find specific temp table** drop-down list contains the handles of cached temp tables and can be used to select a different table. In addition, if the drop-down list is reset to **(All)**, then two additional filtering options are available. The **Find by temp table type** drop-down list will filter the tables by type: *worksheet* for tables generated as the worksheet inside, e.g., a `<link>`, and other types which in general indicate the operation that created the temp table (e.g. *merge*, *table*, *loop*, *transpose*, etc). Furthermore, the **Find by query text** field may be used to search for all temp tables whose query contains arbitrary text. When the view is being used in all tables/filter mode, you can use the navigation buttons `<<`, `<`, `>`, and `>>` ("first," "previous," "next," and "last," respectively) to browse the set of matching tables.

For some temporary tables (e.g. worksheets in links) that were generated directly from MDB ops, a button **Show ops in state cache** appears in the temp table display. This will show the cached operations in the **MDB state cache** view that were run to generate the temporary table, and performance of these operations can be further investigated from that view. It's important to understand that the query in the state cache may not reflect the full sequence of operations that generated the temporary table, because ordinarily only operations that create new states (e.g. selections and tabulations) yield a new cache state. So, for example, for a worksheet that consists of a selection and a tabulation followed by some computed columns, the query shown in the state cache view may only show the selection and the tabulation. Also, the time to create a temporary table will in general be strictly longer than the time associated with the cached operations, both because of these potential additional, uncached operations, and because it may

include the time to get the data from the subprocesses if the worksheet did not include a tabulation or similar operation.

Figure 111: Temporary table cache



Find specific temp table

Select a base table from the drop-down list.

Find by temp table type

Select the type of temporary table from the drop-down list.

Find by query text

Enter text that is in the query.

Arrow keys

Browse through the temporary tables.

Temp table

This section shows details about the temporary table cache, including the following information.

- Temp table name
- Click **Search state cache for references** to find references to the temporary table.
- Click **Show data** to view the data of the temporary table. You can then view the data in TRS, if desired.

Physical tables

The **Physical tables** view allows you to browse details of the physical tables in the 1010data database.

The **Physical tables** view is available from the main PPA window and from the **Finished queries** and **MDB state cache** views. It offers a concise overview of performance-related information about physical tables, including segby, sortseg, ts breaks, indexed columns, prelinks, and distribution of data over segments, subprocesses, and machines as actually seen by the current session.

Figure 112: Physical tables

The screenshot shows the 'Post-hoc Performance Analyst' window. At the top, there are 'Help' and 'Close view' buttons. Below them, the 'Table' field is set to 'pub.doc.retail.salesde' with a refresh icon. To the right, the 'Title' is 'Sales Detail' and the 'Type' is 'REAL'. The main area displays the following information:

- # of segments:** 2
- Segmented by:** customer, transid
- Sortseg:**
- TS breaks:**
- Indexed columns:**
- Prelinks:** N.B. greyed-out prelinks are invalid

Below this information is a dropdown menu for 'Show data distribution aggregated:' set to 'By server'. A dark redacted bar is present below the dropdown. At the bottom, a table shows the distribution of data across machines:

Machine	#processes	#segments	# rows
	2	2	4,481,814
nj4198	1	1	2,916,768
nj4199	1	1	1,565,046

Table

Select the physical table to analyze.

Title

The title of the selected table.

Type

The type of the selected table.

of segments

The number of segments in the physical table.

Segmented by

The columns by which the table is segmented.

Sortseg

If the table is sortseg, the sortseg column(s).

TS breaks

The time series breaks in the table, if any.

Indexed columns

The columns that are indexed, if any.

Prelinks

The precalculated linkage between two tables on a specific set of columns, if any exist.

Show data distribution aggregated

Choose a data distribution aggregation from the drop-down list.

By server

The distribution of the physical table on servers, including the number of processes, segments, and rows on each server.

By process**By segment**

Gadgets

A 1010data Insights Platform gadget is a simplified tool that supports your data analysis session.

The various gadgets are accessed by selecting **Gadgets** from the workspace menu bar. The available gadgets are:

Memo Pad

The Memo Pad allows you to draft text, such as a query or Macro Language code, create temporary tables, and add images for use during your session.

Calculator

The Calculator allows you to test scalar expressions and mathematical operations to explore the results without adding them to a query or associating them with a table.

Memo Pad

The Memo Pad allows you to draft text, such as a query or Macro Language code, create temporary tables, and add images for use during your session.

The Memo Pad is a simple workspace version of a scratch pad or scrap paper. You can write notes and lists, copy and paste information from tables or Macro Language code, or add information you want to reference as you work. The contents of your Memo Pad will remain on a saved workspace exactly as it is when the workspace is saved if the window remains open or minimized.

Memo Pad contents cannot be saved individually and are not retrievable after the **Memo Pad** window is closed or the session of a non-saved workspace ends.

The Memo Pad is accessed by selecting **Gadgets > Add Memo Pad** from the menu bar. The **Memo Pad** window contains options similar to a standard text editor.

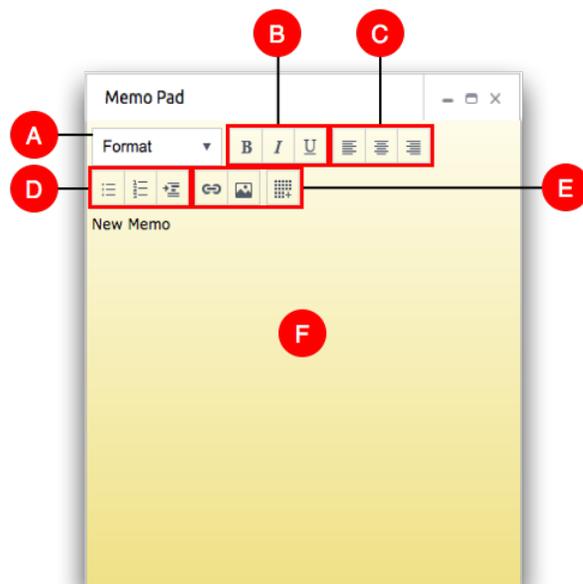


Figure 113: Memo Pad

A. Style options

The style options drop-down contains a list of pre-formatted text styles. For example, you can choose different heading levels to organize the text.

B. Text decorations

The text decoration options change the style of the text. For example, you can format text to be bold, italicized, or underlined.

C. Text alignment

The text alignment options determines the positioning of the text. For example, you can choose left, center, or right justification.

D. List Options

The list options control the type of list and the list item indentation. For example, you can create bulleted or numbered lists.

E. Insert options

The insert options add links, images, or tables.

F. Text field

The text field is where text and inserted objects and lists are displayed.

Note: The text "New Memo" is not overwritten automatically. You may want to delete this text before you begin entering content in the text field.

Calculator

The Calculator allows you to test scalar expressions and mathematical operations. This allows you to explore the results without adding them to a query or associating them with a table.

Since the Calculator is not attached to any table, the Calculator's expression editor can only evaluate mathematical operations and scalar expressions. A scalar expression is an expression that results in an individual value such as an integer or string. It can also be a value containing multiple components that can be referenced individually or as a whole such as a package or list-value. For more information on scalar expressions, see [scalar expression](#). For more information on writing expressions, see [Writing Expressions](#) in the *1010data Reference Manual*.

The expression and results appear as entries in the feed above the Expression Editor field. Each entry is added to the feed, and you can scroll through all of the results if the feed is not cleared and the **Calculator** window remains open. The results in the feed cannot be saved. If you click **Clear**, or close the window, all of the entries in the feed cannot be retrieved.

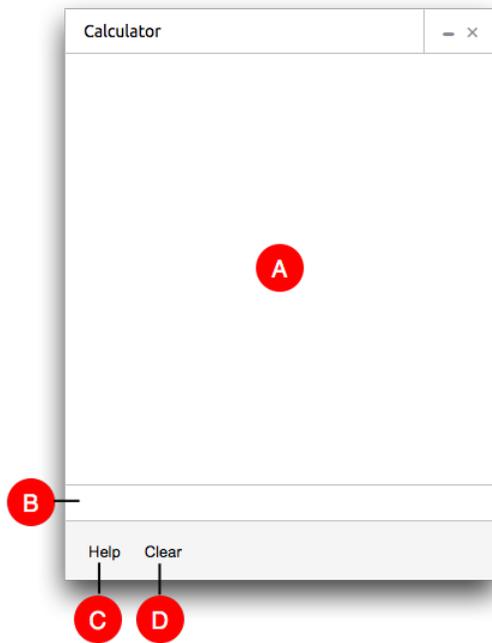


Figure 114: Calculator

The **Calculator** window includes:

A. Feed

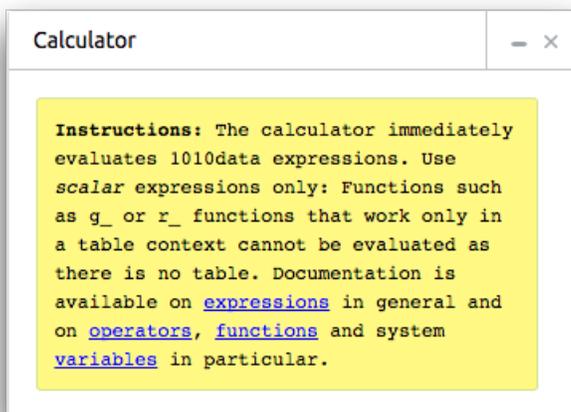
Displays the scalar expressions, mathematical operations, and their results. These are not retrievable if the window is closed or the feed is cleared.

B. Expression Editor

Allows you to enter scalar expressions and mathematical operations for evaluation.

C. Help

Prints the Calculator instructions to the feed.



D. Clear

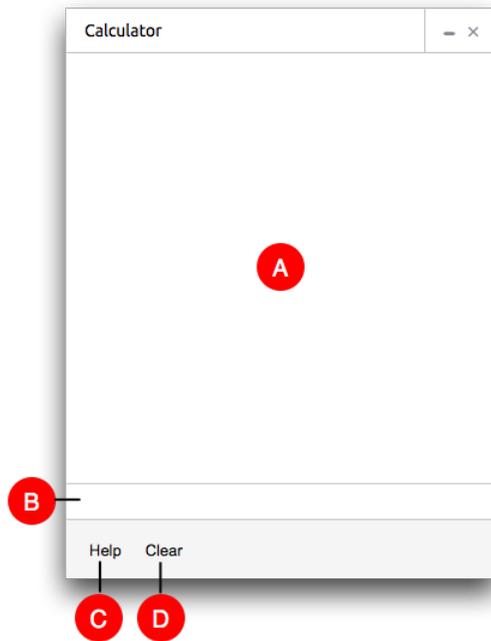
Clears the feed. Text in the feed cannot be recovered after it is cleared.

Evaluate a scalar expression

Use the Calculator to evaluate a scalar expressions.

To evaluate a scalar expression.

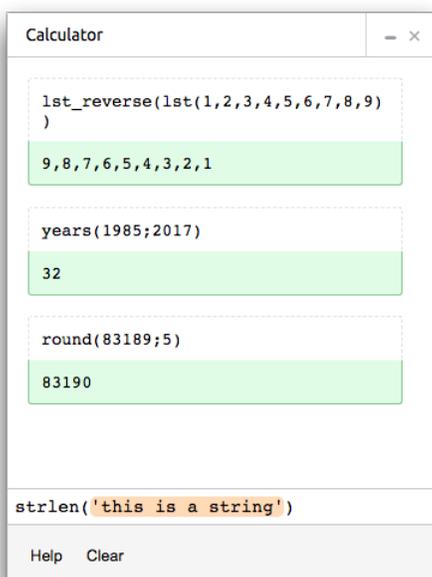
1. Open the Calculator from the menu bar by selecting **Gadgets > Add Calculator**
The Insights Platform displays the **Calculator** window.



2. In the expression editor, enter a valid scalar expression.

For more information on scalar expressions, see [scalar expression](#). For more information on writing expressions, see [Writing Expressions](#).

The Calculator displays the result of the expression in the feed.



Tag Inspector

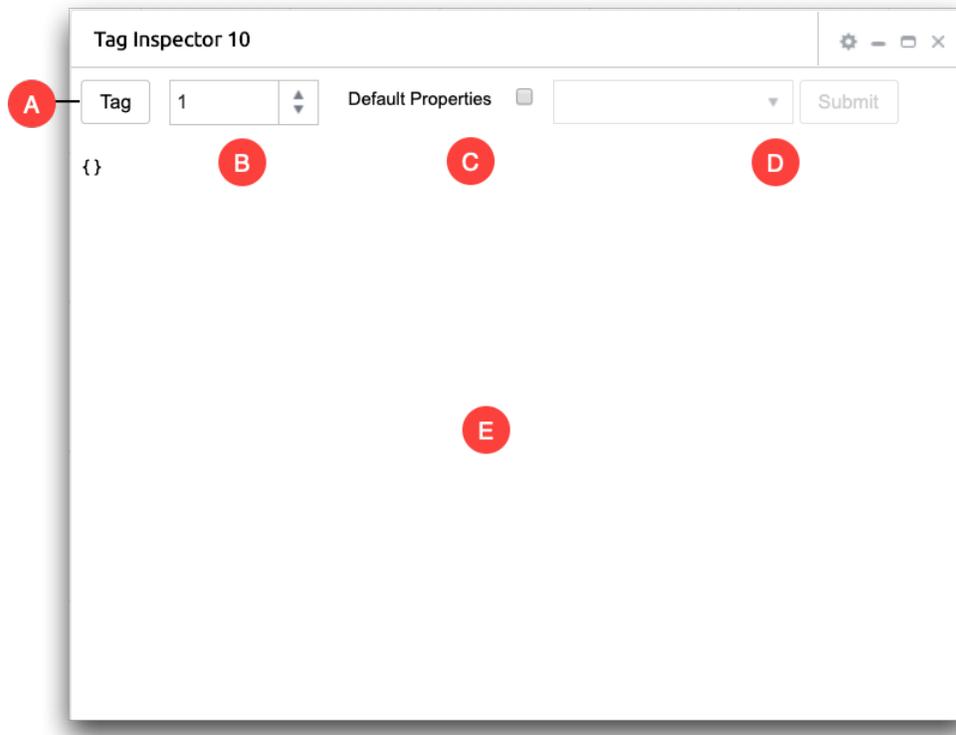
The Tag Inspector allows you to inspect the tags in the currently-running QuickApp and is a useful debugging tool.

The Tag Inspector contains a "scope" widget pointed at the currently-selected QuickApp tag.

The Tag Inspector is accessed by selecting **Gadgets > Add Tag Inspector** from the menu bar. Alternatively, you can launch the Tag Inspector from the results pane of a QuickApp in Macro Language

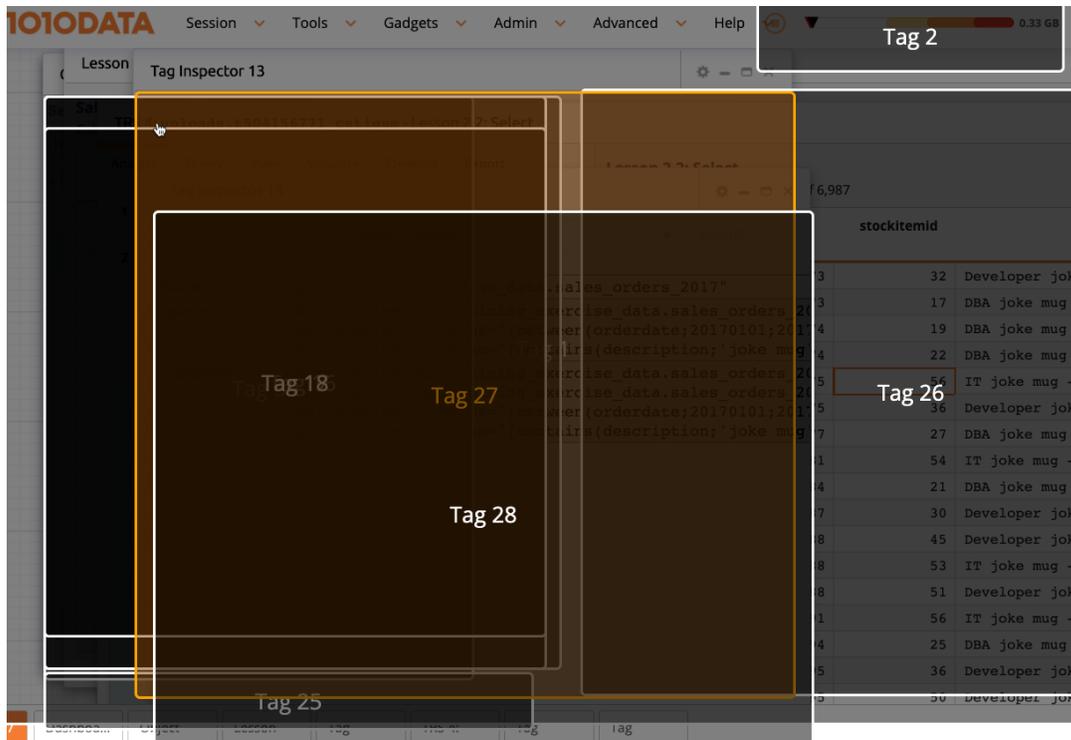
Workshop by clicking the **Tag Inspector** (🔍) icon.

When you first launch the Tag Inspector, the following screen appears:



A. Tag

Click this button to select a tag. When clicked, the Tag Inspector visually shows the tags as overlapping rectangles in your workspace. The borders of the selected tag are highlighted, as well as the number of the tag. In the image below, Tag 27 is selected.



Click a tag to view its contents.

B. Tag selector

Use the up and down arrows to locate the desired tag to inspect, instead of using the Tag button. Click the return key to view the tag details for that numbered tag.

C. Default Properties

Select **Default Properties** if you wish to display all relevant system variables. This is the equivalent of `gui20_="1"` in the `scope` widget.

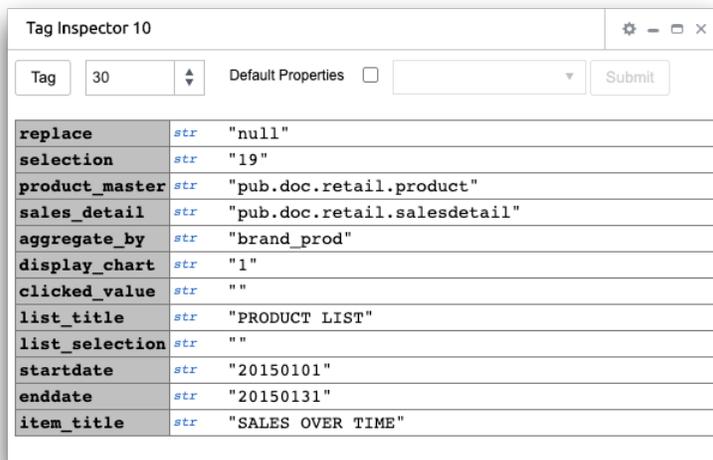
D. Drop-down and Submit

Depending on the tag, there may be options to select from the drop-down list. Select the option, such as `clearcache`, and click **Submit**.

E. Tag contents

The remainder of the screen contains a list of the properties and contents of the tag.

The following is a Tag Inspector window for a simple QuickApp:



You can interact with the Tag Inspector and change values.

Image Encoder

The Image Encoder is a tool for converting images (or other binary files) into data URLs. A data URL has the prefix `data:` and is a convenient way to embed small files inline in documents.

The Image Encoder allows you to encode the actual images instead of having to dynamically fetch remote resources (that can be removed without notice). You can drag and drop the image you wish to convert onto the Insights Platform workspace. The image will appear in a new **Image Encoder** window. Alternatively, you can access the Image Encoder by selecting **Gadgets > Add Image Encoder** from the menu bar.

See https://developer.mozilla.org/en-US/docs/Web/HTTP/Basics_of_HTTP/Data_URLs for more information about data URLs.

Note: For best results, run your image through an optimizer to minimize the file size before encoding the image. The Image Encoder is most suitable for small logos, icons, and favicons.

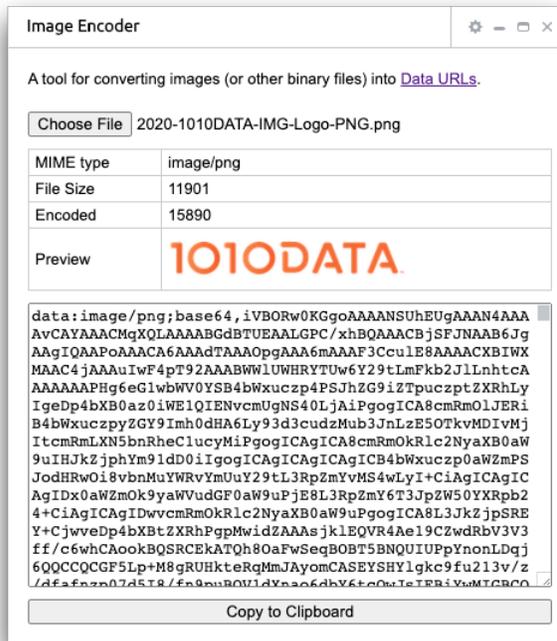


Figure 115: Image Encoder

Choose File

Select the binary file to convert to a data URL.

MIME type

The type of data; for example, image/jpeg, image/png, text/html.

File Size

The original file size in bytes.

Encoded

The file size after it is encoded.

Preview

A preview of the converted image. The remainder of the screen shows the image converted to a data URL.

Copy to Clipboard

Click to copy the data URL into your code.

Scheduler

The Scheduler allows you to create, clone, edit, delete, and manually run report jobs.

You can use the Scheduler to create a report job that runs a query and delivers results on a regular basis, such as every hour or every day at a particular time, to any number of recipients at your company. These reports can be provided in a variety of formats, such as comma-separated text files, PDF documents, and Excel workbooks. You have the option to provide the report as a link or as an attached file. You also have the option to compress the attached file.

You are allotted a certain number of runs per month. These runs can be used for any number of scheduled jobs. For more information, see [Run limit](#) on page 452.

The Scheduler is accessed by selecting **Admin > Scheduler** from the workspace menu.

The screenshot shows the Scheduler interface with several key components labeled A through G:

- A:** Search Jobs input field and filters for Owned Only and Active Only.
- B:** Existing Scheduled Jobs table with columns for Job Name, Job ID, and Status.
- C:** Selected Job's Run History table with columns for Run Date, Scheduled Time (Local), Status, and Run Time (min).
- D:** Job Name and Recipients input fields.
- E:** Monthly Run Summary progress bar showing [Previous: 27] [Scheduled: 14] [Projected: 7] [Allowable: 31].
- F:** Schedule button.
- G:** Overwrite button.

Job Name	Job ID	Status
Job1	466580869	Active
Job2	466580856	Inactive
Job3	466580846	Active

Run Date	Scheduled Time (Local)	Status	Run Time (min)
03/19/20	11:48:00p	Ended	0
03/20/20	04:18:05p	Ended	0
03/20/20	04:18:53p	Ended	0
03/20/20	04:36:22p	Ended	0
03/20/20	11:48:00p	Ended	0
03/21/20	11:48:00p	Ended	0
03/22/20	11:48:00p	Ended	0
03/23/20	11:48:00p	Ended	0
03/24/20	11:48:00p	Scheduled	

Figure 116: Scheduler window

A. Search Jobs

Search for existing jobs. You have the option to select only the jobs you own and/or only active jobs.

B. Existing Scheduled Jobs

View the details of, and perform certain actions on, an existing job. For example, you can run a job, delete a job that is no longer necessary, or reactivate an inactive job.

For more information, see [Existing Scheduled Jobs and Selected Job's Run History](#) on page 441.

C. Selected Job's Run History

Select a job in **Existing Scheduled Jobs** and view the run history and status for the selected job in this section.

For more information, see [Selected Job's Run History](#) on page 443.

D. Job Toolbar

The toolbar contains the components of the job: **Basic Info**, **Base Table**, **Job Timing**, **Output Options**, **Email Contents**, and **Advanced**.

E. Job Details

This section of the Scheduler contains information relevant to the component of the job. Click the Job Toolbar to view the various components.

For more information, see [Job Details](#) on page 444.

F. Monthly Run Summary

This section contains details about the monthly runs, including previous, scheduled, and projected.

For more information, see [Monthly Run Summary](#) on page 451.

G. Schedule/Overwrite Buttons

Click **Schedule** to schedule a new job or **Overwrite** to overwrite an existing job.

Existing Scheduled Jobs and Selected Job's Run History

The **Existing Scheduled Jobs** and **Selected Job's Run History** sections of the Scheduler allows you to find, view, run, stop, delete, activate/deactivate, clone, and subscribe/unsubscribe from jobs.

You can manage your existing scheduled jobs from the **Existing Scheduled Jobs** section. For more instructions see:

- [Find a job](#) on page 456
- [Edit a job](#) on page 456
- [Run a job](#) on page 457
- [Stop a currently running job](#) on page 460
- [Delete a job](#) on page 459
- [Clone a job](#) on page 458
- [Unsubscribe from a multi-user job](#) on page 461
- [Deactivate a job](#) on page 461

The screenshot shows the Scheduler interface. At the top, there is a search bar and filters for 'Owned Only' and 'Active Only'. Below this is the 'Existing Scheduled Jobs' section, which contains a table with columns for Job Name, Job ID, and Status. Three jobs are listed: Job1, Job2, and Job3. Job1 is highlighted, and its Job ID is shown as '_466580869_ >'. Below the table, there are buttons for 'View Raw Data', 'Run', 'Delete', and 'Stop'. The 'Selected Job's Run History' section is also visible, showing a table with columns for Run Date, Scheduled Time (Local), Status, and Run Time (min). The run history for Job1 shows several runs that have ended.

Job Name	Job ID	Status
Job1	_466580869_ >	Active
Job2		
Job3		

Run Date	Scheduled Time (Local)	Status	Run Time (min)
03/19/20	01:32:33p	Ended	0
03/19/20	01:33:19p	Ended	0
03/19/20	11:54:00p	Ended	0
03/20/20	11:54:00p	Ended	0
03/21/20	11:54:00p	Ended	0

Figure 117: Existing Scheduled Jobs and Selected Job's Run History

Existing scheduled jobs list

Search for and view existing scheduled jobs.

If you are an administrator, or a non-administrative user who is subscribed to a multi-user job, you can choose to view all jobs or just the jobs you own. You can also choose active jobs only.

For more information, see [Scheduled Job List](#) on page 442.

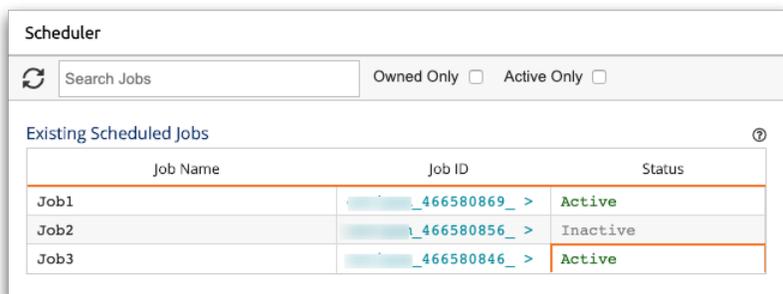
Selected Job's Run History

Displays the run history of a selected job.

For more information, see [Selected Job's Run History](#) on page 443.

Scheduled Job List

The scheduled job list shows all of the existing scheduled jobs and provides options for searching through them.



Job Name	Job ID	Status
Job1	466580869	Active
Job2	466580856	Inactive
Job3	466580846	Active

Figure 118: Scheduled job list and related display options

Refresh () Icon

Refreshes the list of scheduled jobs to display any changes.

Search Jobs

Searches all job names and IDs; you cannot search job owners. For instructions, see [Find a job](#) on page 456.

Owned Only

If selected, displays all the jobs that you own.

If not selected, displays all your jobs, including multi-user jobs to which you are subscribed but do not own.

Active Only

Filters the list to display only active jobs.

Job Name

Displays the name of the job.

This is the job name entered when the job is created.

Job ID

Displays the job ID.

This is generated by the Scheduler. To perform any actions on an existing job, you have to click the job ID.

Status

You can click the **Status** field to toggle its value. The status of the job is one of the following:

Active

The job's end date is in the future.

Inactive

The job's end date is in the past. If you reactivate a job, the end date will be set 30 days from the current day by default. No user will receive results while the job is inactive.

Subscribed

You are subscribed to a multi-user job that an administrator has established on your behalf. The job is run under your user ID and counts toward your monthly limits. You can click here to unsubscribe from the job.

Unsubscribed

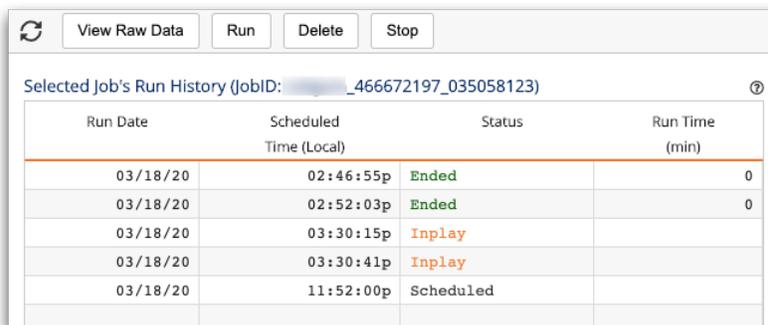
You are unsubscribed from a multi-user job. You can click here to resubscribe to the job.

Additional -Multi(*)

The -Multi suffix indicates that the scheduled job is a multi-job. The * indicates who owns the multi-job.

Selected Job's Run History

Displays the history of a selected job, including when it was run and if it was successful.



Run Date	Scheduled Time (Local)	Status	Run Time (min)
03/18/20	02:46:55p	Ended	0
03/18/20	02:52:03p	Ended	0
03/18/20	03:30:15p	Inplay	
03/18/20	03:30:41p	Inplay	
03/18/20	11:52:00p	Scheduled	

Figure 119: Job run history

Refresh (↻)

Click this button to refresh the run history for the selected job.

View Raw Data

Click this button to view an XML file containing job properties and job history.

Run

Click this button to manually run a job.

Delete

Click this button to delete the job.

Note: The **Delete** button deletes the job and the job's entire run history. You can choose to make a job inactive instead. See [Deactivate a job](#) on page 461 for more information.

Stop

Click this button to stop an in-progress job.

This button is only displayed when a job is in progress. For more information, see [Stop a currently running job](#) on page 460.

Job ID

The job for which you are displaying the run details.

Run Date

The date the job ran or will be run.

Scheduled Time (Local)

The time that the job started or will start.

The displayed time is based on the selected time zone in your user profile settings. For more information, see [Localization](#) on page 79.

Status

The status of the job.

Ended (Succeeded)

The job completed without error and results sent to recipients.

Abended (Failed)

The job ran and failed. The error is recorded. Click in the grid to see the error message.

Inplay (In Progress)

The job has begun and the query is being processed. The results will be sent to recipients once completed.

Closed (Timed Out)

The job reached the third timeout window and did not explicitly fail, but did not complete. All delay notifications have been sent. The job will not be run further.

The timeout windows and delay notifications are defined when the job is created. For more information, see [Advanced settings](#) on page 450.

Scheduled

The job is scheduled to run, but is not in the process of running.

Run Time (min)

The length of time that the job ran.

Job Details

The Job Details section of the Scheduler allows you to create a job, view the details of an existing job, and clone an existing job.

The screenshot displays the Scheduler application interface. On the left, under 'Existing Scheduled Jobs', a table lists Job1 through Job4. Job4 is highlighted, and a tooltip shows its full Job ID: 'catlgun_466571465_783442690'. Below this, the 'Selected Job's Run History' table shows a series of runs from 03/19/20 to 03/23/20, with statuses ranging from 'Ended' to 'Scheduled'. On the right, the 'Job Details' section is active, showing 'Job Name: Job4', a 'Recipients' field, and a 'Monthly Run Summary' section with a progress bar and 'Schedule' and 'Overwrite' buttons.

Figure 120: Job Details section of the Scheduler

The Job Details section of the Scheduler is divided into the following six detail screens, represented by the six icons in the Job Toolbar:

Basic Info

This screen is where you enter the job name and recipients. You can also clone a job here by selecting an existing job and changing the job name.

For more information, see [Basic Info](#) on page 446.

Base Table and XML

This screen is where you select the table for your job, or create a custom Macro Language query for your job.

For more information, see [Base Table and XML](#) on page 446.

Job Timing

This screen is where you select the whether the job has a regular frequency or is triggered by another job, the time of the job, the frequency of the job, and the start and end dates.

For more information, see [Job Timing](#) on page 447.

Output Options

This screen is where you select the format for your job results. Possible formats are Excel, CSV, PDF, QuickApp, and Table.

For more information, see [Output Options](#) on page 449.

Email Contents

This screen is where you enter the email options for your job, including the email subject and body.

For more information, see [Email Contents](#) on page 450.

Advanced

This screen is where you control advanced settings, such as email options for jobs that have timed out.

For more information, see [Advanced settings](#) on page 450.

At the bottom of the Job Details section is a Monthly Run Summary. The Monthly Run Summary keeps track of your monthly scheduled and projected job runs. For more information, see [Monthly Run Summary](#) on page 451.

Basic Info

The Basic Info () section of the Scheduler is where you enter a job name and recipients for your job.

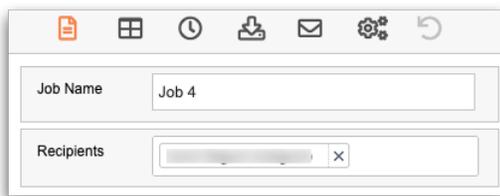


Figure 121: Basic Info in the Scheduler

Job Name

The name of the job. The job name is blank when you launch Scheduler. If you select an existing job ID in **Existing Scheduled Jobs**, the job name appears here.

Recipients

This field automatically lists your user ID. You can add other 1010data Insights Platform users in your company by entering their user IDs in the field. You cannot enter other email addresses.

Base Table and XML

The Base Table and XML () section of the Scheduler is where you select a base table or create an XML query for your job.

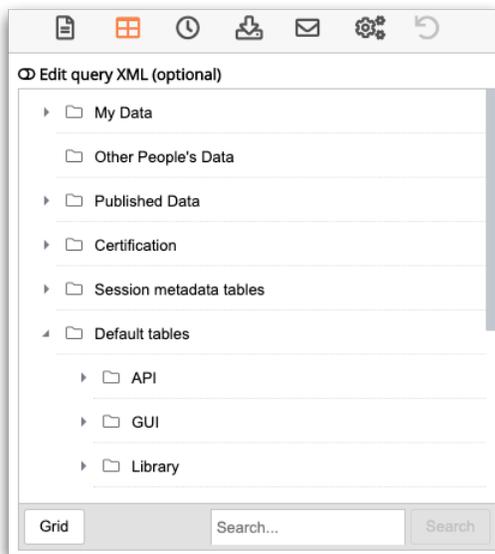


Figure 122: Base Table and XML in the Scheduler

Base Table/Edit query XML

This is the toggle to switch between selecting a table or creating an XML query. You can select a table or a default table such as **Empty base table** or **Default base table**.

Switching to **Edit query XML** provides an edit window similar to Macro Language Workshop. This is where you can create a custom query for your job.

Grid/Tree button

Click **Grid** to view the tables as a grid, and **Table** to view the tables as a tree.

Note: If you are viewing the tables as a grid, select the table and click **Select**.

Search

Enter a table name in the **Search...** box and click **Search**.

Job Timing

The Job Timing (🕒) section of the Scheduler lets you select report frequencies, start date and end date, and days before expiration to send a notification.

Figure 123: Job Timing in the Scheduler

Job Type

Regular Frequency

Report runs at the interval determined by the user-defined settings in this section.

Dependent Trigger

The selected job is triggered when another job is run. Select the **Job Trigger** from the drop-down list.

Time

The time at which the job is run.

Note: If **Hourly** is selected for the frequency, only the minutes and time zone selection are available.

Timezone

The U.S. timezone for the time the report is sent.

Frequency

Hourly

You can choose to run the query and receive a report every hour, or every customized number of hours. Select the minutes after the hour and AM or PM in the minutes part of the **Time** section. The job will run at the first hour of the day and will repeat at the chosen frequency. For example, if the minutes is set at :30 PM and the report is run every 12 hours, the report will first be run at 12:30 PM, the next report will be run at 12:30 AM, and so on every 12 hours.

Every Day

You can choose to run the query each day, each weekday, or every 1-15 days independent of month. For example, if you select 7 days, the report job first runs at the first of the month, and then every 7 days for the rest of the month.

Every Week

You can choose to run the query on a specific day or days of the week.

Every Month

You can choose to run the query on a specific date of each month.

For example, if you chose on day 3 of each month, the job would run on the third of each month.

You can choose to run the query on the first or last day or weekday of each month.

Start Date

The first day the job is scheduled to run.

End Date

The last day the job is scheduled to run. Select **None** to run the job indefinitely.

Days before expiration to send notification

Select **None** through **10** from the drop-down list for the number of days before the end date to send a notification to report recipients.

Output Options

The Output Options (📁) section of the Scheduler allows you to select the file format in which the results are created, and whether to compress the results file.

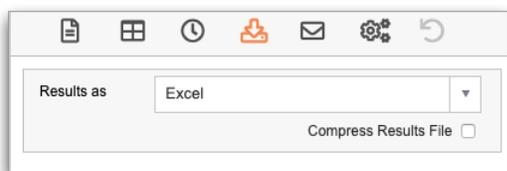


Figure 124: Output Options in the Scheduler

Results as

The file format in which the report results are saved.

Excel

The report results are saved as a Microsoft Excel workbook file.

CSV

The report results are saved as a comma-separated values (CSV) text file.

PDF

The report results are saved as a PDF file.

QuickApp

The report results are saved as a QuickApp.

Table

The report results are saved as a 1010data Insights Platform table.

None

This option sends the email confirming that the job succeeded or failed, but there is no attached report. This option is good for report jobs that build tables.

Save to Path for QA or Table

When either **QuickApp** or **Table** is selected, this field lists the path of the folder where the QuickApp or table is saved.

Compress Results File

Select to compress the results file when saving the file as an attachment to an email.

Email Contents

The Email Contents (✉) section of the Scheduler is where you enter the email subject and body, attach results as a file, and include the log file.

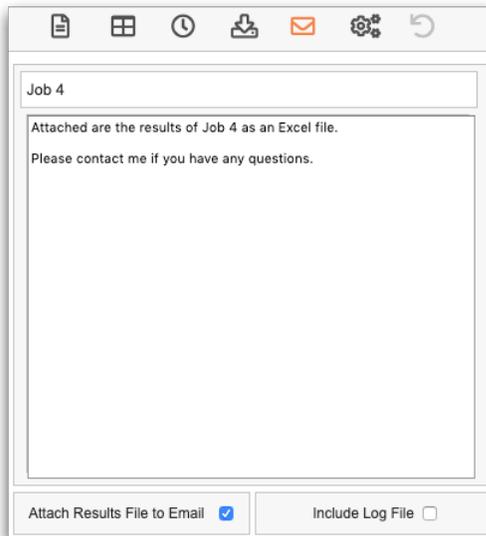


Figure 125: Email Contents in the Scheduler

Subject (opt)

An optional custom subject line for your email containing the results of the job. If you leave this field blank, the subject will contain the status of the job followed by the name of the job, for example, Completed: Job 4.

Email body (opt)

An optional email body. This text appears before the boilerplate text in the email. The boilerplate text contains a link to the report, run limit information, job name, job ID, scheduled run time, actual run time, and end time.

Attach Results File to Email

All emails contain a link to the report, but if you select this option, the email will also contain a file attachment in the format specified in the [Output Options](#) section.

Include Log File

The email also includes an attached log file.

Advanced settings

Select multi-jobs, specify the version of the platform in which to run the job, and configure time-out email messages.

The **Advanced** (⚙️) section of the Scheduler allows you to select multi-jobs and configure whether email messages are sent, and to whom, if the job runs longer than expected. These time-out email messages differ from the results email messages that are sent when a job is complete.

Additionally, the **Advanced** section allows you to choose the platform version in which to run the job.

Figure 126: Advanced setting fields and options

Click the **Advanced** (⚙️) icon in the right pane of the Scheduler to display these options.

Job Runner

This is the user ID used to run the scheduled job. For multiple job runners, select **Multi-job**.

Note: When **Multi-job** is selected, each user ID listed in the **Enter Recipients** field of the **Basic Info** screen runs the job and receives personal results. This is in contrast with the job being run by a single ID and the listed recipients all receiving the same results. Users can unsubscribe from the multi-job if they no longer wish to run the job.

Version

This is the version you are scheduling the job to run in. It does not have to be the version you are currently using. For example, you could be working in Prime and schedule a job for Beta.

Email 1

If the job reaches this time limit, an email is sent. The job continues to run.

Email 1

The email address of the person who should receive notification if the job is running too long.

You can specify multiple email addresses that are space separated.

Send email alert if runtime (min) exceeds

Select the time limit, in minutes, when the notification email is sent.

Email 2

If you reach this time limit, another email is sent. The job continues to run.

Email 3

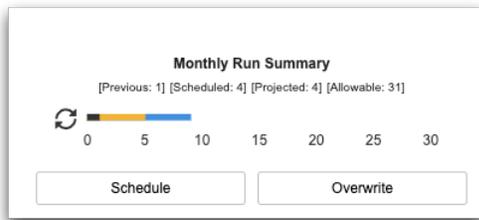
If you reach this time limit, a final email is sent. The job is stopped.

Monthly Run Summary

Use the monthly run summary to monitor your previous, scheduled, and allowable job runs.

For more information about run limits, see [Run limit](#) on page 452.

Figure 127: Monthly Run Summary



Previous

Displays the number of runs already completed this month.

Scheduled

Displays the number of runs scheduled from existing jobs that are scheduled to be performed this month.

Projected

Displays the number of runs a new scheduled job will include.

Note: By default, the gauge displays the number of future runs created by a job that is executed daily. This is the number of runs that would be created using the default frequency options settings.

Allowable

Displays the maximum number of allotted runs per month.

Refresh

The **Refresh** (↻) icon updates the gauge to include any changes to the job settings.

Run limit

Every user is assigned a specific number of job runs that they can execute each month.

In the past, users were assigned a number of scheduled jobs they could set. A single job may have run once a month or once an hour, and regardless of how often it ran, it counted against your job limit. Run limits allowed you to have any number of jobs. Now, you are limited by the specific number of job runs that you are assigned. This number is displayed by the **Monthly Run Summary**. For more information, see [Monthly Run Summary](#) on page 451.

When you create a scheduled job, you select the frequency that it is executed. Each execution is a job run. For example, if you create a scheduled job with hourly frequency, 24 job runs will be executed in one day. Each job run deducts from your allowed monthly runs. A new job, or change in the frequency of an existing job, may not trigger enough runs to go over your remaining run limit for the current month, but may exceed the run limit of the next month when combined with all of your jobs.

We recommend checking the monthly run summary at the beginning of each month to verify the planned number of jobs for the month. If your scheduled jobs produce more runs than your run limit, the remaining job runs for the month are not executed, and you are sent an email for each run that is not performed.

Create a job

Create a new job in the **Scheduler**.

To create a job:

1. On the right pane of the **Scheduler**, click the **Basic Info** (📄) icon. For **Job Name**, enter a name for your new job.

A screenshot of a text input field labeled "Job Name" with a light gray border and a white background.

2. Enter the recipients.

A screenshot of a text input field labeled "Recipients" with a light gray border and a white background.

The **Enter Recipients** field automatically lists your user ID. You can add other 1010data Insights Platform users in your company by entering their user IDs in the field. You cannot enter other email addresses.

3. Click the **Base Table and XML** (🗃️) icon. Select your table or query from the object browser.

A screenshot of a dialog box titled "Edit query XML (optional)". It contains a tree view with the following folders: "My Data", "Other People's Data", "Published Data", "Certification", "Session metadata tables", and "Default tables". At the bottom, there is a "Grid" button, a "Search..." text box, and a "Search" button.

You must select a table. If your query defines the table you want to use, you can select that table, or you can select a default table in the **Default tables** folder, such as **Empty base table** or **Default base table**. The table defined in the query takes precedence over the selected table.

4. Click **Edit query XML (optional)** if you want a query to return specific information from the table.

A screenshot of an XML query editor. The text area contains the following XML code:

```

1 <base table="pub.doc.retail.salesdetail"/>
2 <sel value="store=15"/>
3 <sel value="trans_date=20180101"/>
4 <sel value="dept=36"/>
5

```

5. Click the **Job Timing** (🕒) icon. Choose the **Job Type**. You can choose a **Job Type** of **Regular Frequency** or **Dependent Trigger**.

Job Type	Regular Frequency		
Time	12	:	00
	<input checked="" type="radio"/> AM <input type="radio"/> PM		
Timezone	Eastern Time Zone (5:00)		
Frequency	Daily		
	<input checked="" type="radio"/> Each Day <input type="radio"/> Each Weekday <input type="radio"/> Repeating		
Start Date	03/17/2020		
End Date	06/15/2020	<input type="checkbox"/> None	
Days before expiration to send notification	7		

Choose **Regular Frequency** if you want to run your job at specific intervals, or choose **Dependent Trigger** if you want your job to run only when another job is triggered.

6. If you chose a **Job Type** of **Regular Frequency**, select the time at which the job should run and the time zone to use.

Job Type	Regular Frequency		
Time	2	:	00
	<input type="radio"/> AM <input checked="" type="radio"/> PM		
Timezone	Eastern Time Zone (5:00)		
Frequency	Daily		
	<input checked="" type="radio"/> Each Day <input type="radio"/> Each Weekday <input type="radio"/> Repeating		
Start Date	03/17/2020		
End Date	06/15/2020	<input type="checkbox"/> None	
Days before expiration to send notification	7		

7. Choose the frequency of the job.

Job Type	Regular Frequency		
Time	2	:	00
	<input type="radio"/> AM <input checked="" type="radio"/> PM		
Timezone	Eastern Time Zone (5:00)		
Frequency	Hourly		
Every	12		
Start Date	03/17/2020		
End Date	06/15/2020	<input type="checkbox"/> None	
Days before expiration to send notification	7		

Each frequency option has its own selection options, which are displayed below the chosen option. For more information, see [Job Timing](#) on page 447.

After you select the report frequency, the Monthly Run Summary section displays the amount of runs the new job will produce for the remainder of the month. For more information, see [Monthly Run Summary](#) on page 451.

Note: Your new job may exceed the monthly number of runs. For more information, see [Run limit](#) on page 452.

8. Choose the time frame for the job.

Job Type: Regular Frequency

Time: 2 : 00 AM PM

Timezone: Eastern Time Zone (5:00)

Frequency: Hourly

Every: 12

Start Date: 03/17/2020

End Date: 06/15/2020 None

Days before expiration to send notification: 7

Select **None** if there is no end date.

9. Select the number of days before job expiration to send an email notification to your report recipients, or keep the default number of 7.
10. Click the **Output Options** () icon. Select the format for your results and whether you want to compress your results file.

Results as: Excel

Compress Results File

For more information, see [Output Options](#) on page 449.

Note: If either **TABLE** or **QuickApp** is selected, you must enter the path of the folder where you want the table or QuickApp saved.

11. Click the Email Contents () icon. Enter an optional subject and body of your email.
12. Select whether you want to include the results file in the email, and whether to include the log file. If you do not include the results file, the email will contain only a link to the results file.
For more information, see [Email Contents](#) on page 450.
13. Optional: To schedule multi-jobs, change the version of Insights Platform on which to run the job, or send warnings that the job is running too long, click the **Advanced** () icon.

Job Runner:

Multi-job

Version: prime-15.36

Email 1: joejohnson@democompany.com

Send email alert if runtime (min) exceeds: 30

Email 2: joejohnson@democompany.com

Send email alert if runtime (min) exceeds: 60

Email 3: cjoejohnson@democompany.com

Send email alert if runtime (min) exceeds: 90

For more information about these options, see [Advanced settings](#) on page 450.

14. Click **Schedule** at the bottom of the right pane of the Scheduler.

You will receive a confirmation message that the job was created. You can view your job in **Existing Scheduled Jobs** in the left pane of the Scheduler. You may need to click the refresh (🔄) icon.

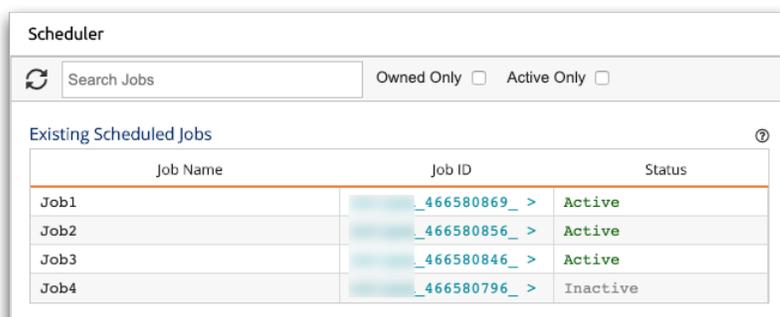
Note: By default, only four jobs can exist for a user at any given time. Delete any inactive jobs in order to create new jobs.

Find a job

You can find previously scheduled jobs by Job Name or Job ID. If you are a non-administrative user, you will see only jobs you own. If you are a company administrator, you can see your jobs and jobs owned by any of the other users in your company.

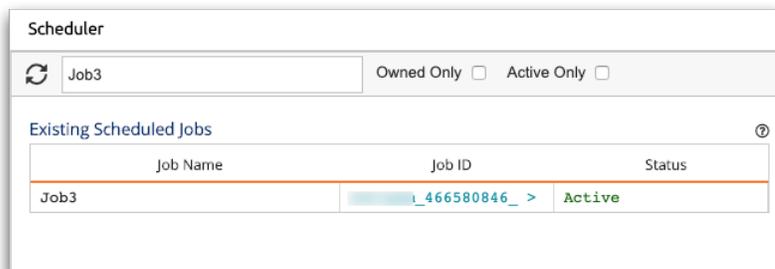
To find a job:

1. In the **Existing Scheduled Jobs** section, click the search field and enter either all or a portion of the job name or job ID.



If you are an administrator, or if you are subscribed to a multi-user job, you can select **Owned Only** to view only the jobs you own. You can also limit your search for active jobs by selecting **Active Only**.

2. Press **Enter** or click outside the search field. The Scheduler displays the results.



Edit a job

You can edit an existing scheduled job.

You can change the details of a job up until it is scheduled to run.

To edit a job:

1. In the **Existing Scheduled Jobs** section, select the job that you want to edit by clicking the Job ID. If you are an administrator, or if you are subscribed to a multi-user job, you can view all jobs, or select **Owned Only** to view only jobs you own. You can also select **Active Only** jobs. The full text for the Job ID appears, and the right pane of the Scheduler displays the job details.

The screenshot shows the Scheduler application window. At the top, there is a search bar and filters for 'Owned Only' and 'Active Only'. Below this is a table of 'Existing Scheduled Jobs' with columns for Job Name, Job ID, and Status. Job 2 is selected, and its Job ID is highlighted. A tooltip shows the full Job ID: '_466580856_783442690'. Below the table are buttons for 'View Raw Data', 'Run', 'Delete', and 'Stop'. The bottom section shows the 'Selected Job's Run History' for Job ID '_466580856_783442690' with a table of run dates, scheduled times, and statuses. To the right, there is a 'Job Name' field (Job2), a 'Recipients' field, and a 'Monthly Run Summary' section with a progress bar and 'Schedule' and 'Overwrite' buttons.

- Using the icons at the top of the pane, select the portion of the job that you want to edit. These options are the same as when you create a job. For more information, see:
 - [Basic Info](#) on page 446
 - [Base Table and XML](#) on page 446
 - [Job Timing](#) on page 447
 - [Output Options](#) on page 449
 - [Email Contents](#) on page 450
 - [Advanced settings](#) on page 450
- Edit the scheduled job. Changes you make to the job frequency results in a change to the Monthly Run Summary.
- Click **Overwrite** to save the edited job without changing the name of the job.

Run a job

You can manually run an existing scheduled job.

Manually running an existing job does not change when the job is scheduled to run. The job will still run at its scheduled time. Manually run jobs count against your run limit. For more information, see [Run limit](#) on page 452.

To run a job:

- On the **Existing Scheduled Jobs** tab, select the job that you want to run by clicking the Job ID. If you are an administrator, or if you are subscribed to a multi-user job, you can select **Owned Only** to view only jobs you own. For more information, see [Existing Scheduled Jobs and Selected Job's Run History](#) on page 441. The Scheduler displays the selected job's run history.

The screenshot shows the Scheduler interface with the following components:

- Search Jobs:** Owned Only Active Only
- Existing Scheduled Jobs:**

Job Name	Job ID	Status
Job1	466580869	Active
Job2	466580856	Active
Job3		
Job4		
- Buttons:** View Raw Data, Run, Delete, Stop
- Selected Job's Run History (JobID: 466580856_783442690):**

Run Date	Scheduled Time (Local)	Status	Run Time (min)
03/19/20	11:52:00p	Ended	0
03/20/20	11:52:00p	Scheduled	
- Job Details (Right Panel):**
 - Job Name: Job2
 - Recipients: [Input field]
 - Monthly Run Summary: [Previous: 11] [Scheduled: 44] [Projected: 11] [Allowable: 31]
 - Progress bar: 0 to 30
 - Buttons: Schedule, Overwrite

2. Click **Run**.
The Scheduler runs the job immediately.

Clone a job

You can clone an existing job.

To clone a job:

1. On the **Existing Scheduled Jobs** tab, select the job that you want to clone by clicking the Job ID. If you are an administrator, or if you are subscribed to a multi-user job, you can select **Owned Only** to view only jobs you own. For more information, see [Existing Scheduled Jobs and Selected Job's Run History](#) on page 441.

The Scheduler displays the selected job in the job details section.

This screenshot is identical to the one above, showing the Scheduler interface with Job 2 selected and its details displayed in the right-hand section.

2. Click the various icons in the job details section to view and/or edit the components of the job, such as base table, timing, output options, and email contents.

- Enter a unique job name in the **Job Name** field of the Basic Info (📄) screen of the job details section.

The screenshot shows a form with two main sections. The top section is labeled 'Job Name' and contains a text input field with the value 'NewJob'. Below this is a section labeled 'Recipients' with a dropdown menu and a close button (X).

- Click **Schedule** to schedule the new cloned job. The cloned job is added to your list of existing scheduled jobs. You can edit the cloned job by clicking **Job ID**. For more information, see [Edit a job](#) on page 456.

To see the change in the Monthly Run Summary, click the refresh button.

Delete a job

You can delete a scheduled job from the **Scheduled Jobs** list.

To delete a job:

- On the **Existing Scheduled Jobs** tab, select the job that you want to run by clicking the Job ID. If you are an administrator, or if you are subscribed to a multi-user job, you can select **Owned Only** to view only jobs you own. For more information, see [Existing Scheduled Jobs and Selected Job's Run History](#) on page 441.

The Scheduler displays the selected job's run history and job details.

The screenshot shows the Scheduler interface. On the left, there is a table titled 'Existing Scheduled Jobs' with columns for Job Name, Job ID, and Status. Job2 is selected, and its Job ID is highlighted. Below the table are buttons for 'View Raw Data', 'Run', 'Delete', and 'Stop'. On the right, there is a 'Selected Job's Run History' section for Job2, showing a table with columns for Run Date, Scheduled Time (Local), Status, and Run Time (min). The table shows two rows: one for 03/19/20 with status 'Ended' and run time 0, and one for 03/20/20 with status 'Scheduled'. Below the table is a 'Monthly Run Summary' section with a progress bar and buttons for 'Schedule' and 'Overwrite'.

Job Name	Job ID	Status
Job1	466580869	Active
Job2	466580856	Active
Job3		
Job4		

Run Date	Scheduled Time (Local)	Status	Run Time (min)
03/19/20	11:52:00p	Ended	0
03/20/20	11:52:00p	Scheduled	

- Confirm that you've selected the correct job to delete by using the job details. If you delete a job, you cannot recover it.
- Click **Delete**. The Scheduler opens a dialog box to confirm that you want to delete the job.
- Click **Confirm** to confirm the deletion. The Scheduler deletes the job. The Monthly Run Summary should update automatically. If it does not, click the refresh button.

Download a completed job report

After a job runs (scheduled or manually executed), the Scheduler automatically sends an email to the recipients specified by the job details.

You can only download a completed job report via the email that is sent after the job is finished. If you need access to the report, you must be added as an email recipient. You cannot download the report from the Scheduler. By default, you are added as an email recipient to any new job that you create.

Single Sign On users must use the same browser they use to access the 1010data Insights Platform.

To download a completed job report:

1. Open the email containing the link to the completed job report.

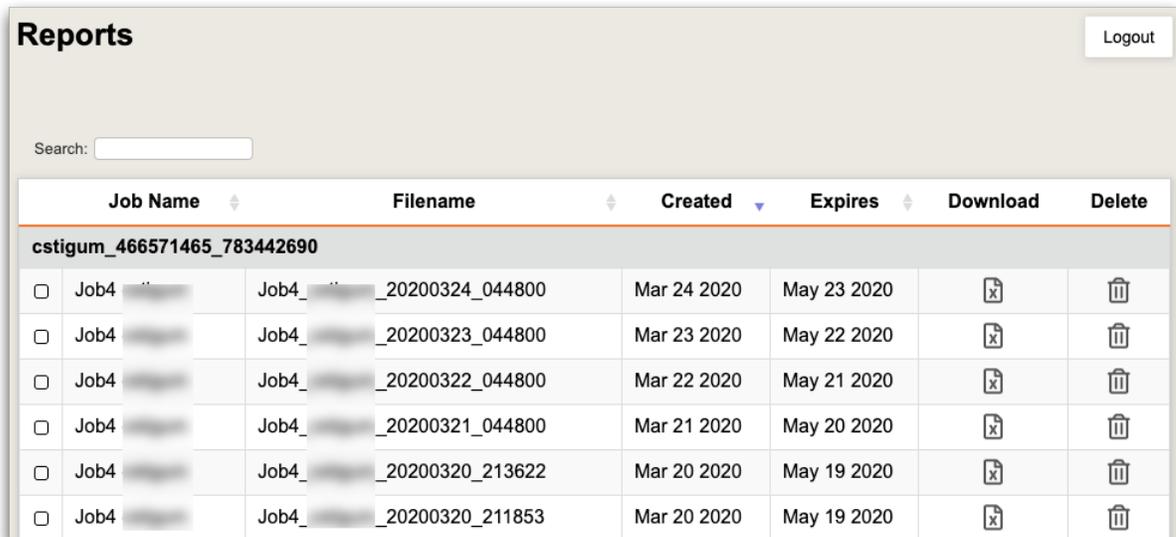
This email is sent by noreply@1010data.com (or a similar, environment-specific email) and contains the default subject: "Completed: [Job_Title]" or a user-defined email subject.

If you selected **Attach Results File to Email** in the Job Details, the email contains a link to download the report, as well as the report itself as an attachment.

2. If prompted, enter your 1010data Insights Platform credentials.

The first time you download a report, you are required to validate your credentials. After you've validated your credentials, the browser should remember them.

The file downloads automatically. You can also view a list of reports that were generated previously, and download them, if desired.



Reports						Logout
Search: <input type="text"/>						
Job Name	Filename	Created	Expires	Download	Delete	
cstigum_466571465_783442690						
<input type="checkbox"/>	Job4 [redacted]	Job4_[redacted]_20200324_044800	Mar 24 2020	May 23 2020		
<input type="checkbox"/>	Job4 [redacted]	Job4_[redacted]_20200323_044800	Mar 23 2020	May 22 2020		
<input type="checkbox"/>	Job4 [redacted]	Job4_[redacted]_20200322_044800	Mar 22 2020	May 21 2020		
<input type="checkbox"/>	Job4 [redacted]	Job4_[redacted]_20200321_044800	Mar 21 2020	May 20 2020		
<input type="checkbox"/>	Job4 [redacted]	Job4_[redacted]_20200320_213622	Mar 20 2020	May 19 2020		
<input type="checkbox"/>	Job4 [redacted]	Job4_[redacted]_20200320_211853	Mar 20 2020	May 19 2020		

Note: To stop receiving emails regarding the scheduled job, click the **Unsubscribe** link at the bottom of the email and then, on the web page that is subsequently opened in your browser, click **Confirm** to unsubscribe.

Stop a currently running job

You can stop a job that is currently running.

You can stop a job that is in progress, which you can see in the **Selected Job's Run History**. The **Stop** button is available until the screen is refreshed. If the job has completed, you may still see the **Stop** button. If you click the button after the job has completed, you may receive an error.

To stop a currently running job:

1. On the **Existing Scheduled Jobs** tab, select the job that you want to stop by clicking the Job ID.

If you are an administrator, or if you are subscribed to a multi-user job, you can select **Owned Only** to view only jobs you own. For more information, see [Existing Scheduled Jobs and Selected Job's Run History](#) on page 441.

The Scheduler displays the selected job's run history.

The screenshot shows the Scheduler interface. At the top, there is a search bar and checkboxes for 'Owned Only' and 'Active Only'. Below this is a table titled 'Existing Scheduled Jobs' with columns for Job Name, Job ID, and Status. The table contains three rows: Job1 (Active), Job2 (Inactive), and Job3 (Active). Below the table are buttons for 'View Raw Data', 'Run', 'Delete', and 'Stop'. At the bottom, there is a section for 'Selected Job's Run History (JobID: ...)' with a table showing run dates, scheduled times, statuses, and run times. The last row in this table has a status of 'Inplay'.

Job Name	Job ID	Status
Job1	..._466580869_ >	Active
Job2	..._466580856_ >	Inactive
Job3	..._466580846_ >	Active

Run Date	Scheduled Time (Local)	Status	Run Time (min)
03/20/20	11:54:00p	Ended	0
03/21/20	11:54:00p	Ended	0
03/22/20	11:54:00p	Ended	0
03/23/20	11:54:00p	Ended	0
03/24/20	11:54:00p	Ended	0
03/25/20	01:44:18p	Ended	0
03/25/20	01:46:04p	Inplay	

2. Select the job run with the status **Inplay**, and click the **Stop** button.

The Scheduler stops the job.

Note: There is no way to restart the stopped job. Additionally, when the job is stopped, no confirmation message is displayed, but the status of the job changes to **Abended (Failed)**.

You can always run the job again from the job details. For instructions, see [Run a job](#) on page 457.

Unsubscribe from a multi-user job

You can unsubscribe from a multi-user job that you are added to by an administrator.

You cannot edit a job that you are not the owner of, but you can unsubscribe from a multi-user job that you were added to by an administrator.

When you unsubscribe, the job continues to appear in the Scheduled Jobs list but it is not run by your user ID while you are unsubscribed. You can resubscribe to an unsubscribed multi-user job at any time.

To unsubscribe from a multi-user job:

1. On the **Existing Scheduled Jobs** tab, find the job that you want to unsubscribe from. Note that the status of the job is **Subscribed**.
2. Click once on the status of **Subscribed**.
A dialog box appears, asking you to confirm unsubscribing from the job.
3. Click **Confirm** in the dialog box.
The status of the job changes to **Unsubscribed**.

Note: To resubscribe to the job, simply click the **Unsubscribed** status and confirm resubscribing to the job.

Deactivate a job

You can deactivate a job before the job's expiration date.

When you deactivate a job, the job continues to appear in the **Existing Scheduled Jobs** list but it is not run and users will not receive results. If **Active Only** is selected, the inactive jobs will be hidden. You can reactivate an inactive job at any time.

Note: You can only deactivate jobs that you own. You can unsubscribe from multi-jobs you do not own. See [Unsubscribe from a multi-user job](#) on page 461.

To deactivate a job:

1. On the **Existing Scheduled Jobs** tab, find the job that you want to deactivate. Note that the status of the job is **Active**.
2. Click once on the status of **Active**.
A dialog box appears, asking you to confirm deactivating the job.
3. Click **Confirm** in the dialog box.
The status of the job changes to **Inactive**.

Note: To reactivate the job, simply click the **Inactive** status and confirm reactivating the job. The end date for the reactivated job will be set by default to 30 days from the present date.

Admin Tool

Administrative functions are accessible from the Admin Tool.

You can use the Admin Tool to access User Manager, Group Manager, Company Manager, and Role Manager.

Note: The administrative functions you have access to depends on your account permissions. You may not see every administration option.

The Admin Tool is accessed by selecting **Admin > Admin Tool** from the workspace menu. The window is labeled **Admin Panel 1.0**.

The screenshot shows the 'Admin Panel 1.0' window. At the top, there's a navigation bar with 'Account Administration' and four icons: User Manager, Group Manager, Company Manager, and Role Manager. Below this is a search box with the prompt 'Find a user by entering: First Name, Last Name, User ID, Email Address, or Company Name' and a 'Find Users' button. Underneath are three buttons: 'Create New User', 'Create User Template', and 'Clear Form'. A section titled 'Create a new user, or search for a user to edit' contains three tabs: 'User Info', 'Account Settings', and 'GUI Preferences'. The 'User Info' tab is active, showing a form with the following fields:

- Personal Information:**
 - *User ID: [text box]
 - *First Name: [text box]
 - *Last Name: [text box]
 - Middle Initial: [text box]
 - *Email Address: [text box]
 - Email List: [text box]
 - Password: [text box]
 - Company Name: [text box]
 - Role: [dropdown menu]
 - Phone Number: [text box]
 - Street Address: [text box]
 - City: [text box]
 - State: [text box]
 - Country: [text box]
 - Zip Code: [text box]
 - Profession: [text box]
- ID Information:**
 - Save changes to User Groups?
 - User Groups: [text box]

User Manager

The User Manager allows users to manage information related to their 1010data accounts. Company administrators can also create, find, edit, and deactivate users.

The **User Manager** page has three tabs:

- **User Info**
- **Account Settings**
- **GUI Preferences**

In addition, company administrators also see a text box in which to enter a first name, last name, user ID, email address, or company name and a **Find Users** button, which is used to initiate the search based on the entered criteria.

There are also four buttons below the text box that are available to company administrators:

Create New User

Create a new 1010data user.

Create Multiple Users

Create multiple 1010data users by uploading a CSV file that contains the users' information.

Save Multiple Users

Change settings for multiple 1010data users at the same time.

Clear Form

Clear the values from the fields on all of the **User Manager** tabs.

User Info

The **User Info** tab of the **User Manager** page contains personal information about a particular user as well as information related to that user's ID, password, and expiration date.

Personal Information

User ID

A unique combination of alphanumeric characters and underscores. *(required)*

The user ID must:

- Begin with a letter
- Contain only lowercase letters

The recommended form for a user ID is:

`<compid>_<first initial of first name><last name>`

Each user must have a user ID that is unique across all companies which use 1010data. For instance, even if you only have one John Smith within your company, another company may already be utilizing the user ID `jsmith`. If you try creating a user with a user ID that already exists, you will get an error message telling you that the user ID is already in use. It is recommended that you add the company ID as a prefix to each user ID that you create to help ensure the uniqueness of your company's user IDs (e.g., `rd_jsmith`).

Note: This field is editable only by company administrators.

First Name

The user's first name. *(required)*

Last Name

The user's last name. *(required)*

Middle Initial

The user's middle initial.

Email Address

A valid email address for the user. *(required)*

This setting must be set correctly. It is used for a user to reset his/her password and to receive important notifications from 1010data. Do not create an ID without a valid email address.

Email List

A space-separated list of additional email addresses associated with the user.

Password

An alphanumeric value containing between 6-15 characters.

The password can be set when creating a user, but it is not required. If a password is not specified, the system will generate a random one. If the administrator prefers to set the password, it should be set to a random string and should not be shared with the user. The existing password is not displayed due to security reasons. The correct way to provide the password to the user is to request that they click **Forgot Your Password?** on www2.1010data.com.

Note: Company administrators can set/change a password for any user. The user that is currently logged into 1010data may also change his or her own password using this field.

Company Name

The name of the company the user works for.

Note: This field is editable only by company administrators.

Phone Number

The phone number for the user.

Street Address

The street address for the user.

City

The city for the user.

State

The state for the user.

Country

The country for the user.

Zip Code

The zip code for the user.

Profession

The job title of the user.

ID Information

Save changes to User Groups?

Select this checkbox to modify a user's group membership.

A company administrator can add or remove users to and from groups by adding or deleting groups from **User Groups**.

Note: Changes to **User Groups** will not take effect unless this checkbox has been selected.

User Groups

A space-separated list of all the groups to which this user belongs.

Idle Time

The amount of time (in minutes) the user can stay logged in without performing an action before they are logged out of the system.

This is important to set, as a user with a large amount of memory mapped can sit in the system and take up space, even though they are done with their work. A good default value for dormant time is 60 minutes, but this value can be changed to something smaller or slightly larger, if need be.

Note: This field is editable only by company administrators.

Deactivate

If this is set to **Yes**, the user is deactivated and can no longer log in to 1010data. If **No**, the user can log in and access the system.

Note: This field is available only to company administrators.

User Expiration Date

The date (in YYYYMMDD form) when this user account expires.

Note: After this date, the user will be denied access to the system.

Note: This field is editable only by company administrators.

Password Set Date

The user will be asked to change their password 45 days after this date. (The date appears in YYYYMMDD form.)

The default cycle of 45 days can be modified by setting the **Password Expiration Cycle**.

Note: When creating a user, it is recommended to set this to a date well in the past (e.g., one year before the date of creation) in order to make sure a user changes their default password when their user ID is given to them.

To turn off the password reset feature, clear the date in the **Password Set Date** field and save the changes for the user.

Note: This field is editable only by company administrators.

Password Expiration Cycle

The number of days the user will be asked to change their password after the **Password Set Date**.

Note: This field is editable only by company administrators.

User Creation Time

The date (in YYYYMMDD form) that this user was created in 1010data. *(read only)*

User Administrator

The user ID of the administrator that created this user in 1010data. *(read only)*

Previous Login Time

The date (in YYYYMMDD form) of this user's last login. *(read only)*

Previous Time User was Updated

The date (in YYYYMMDD form) this user's information was last updated. *(read only)*

Previous User to Update ID

The user ID that last updated this user's information. *(read only)*

Previous Version to Update ID

The version of 1010data that was used when this user's information was last updated. *(read only)*

Account Settings

The **Account Settings** tab of the **User Manager** page contains account-related information for a particular user, including the default version of 1010data they log in to and the amount of local disk space they're allowed to use.

Version

The version of 1010data this user is configured to use.

Possible values are:

- beta-latest
- beta-x.yz
- prime-latest

where *x.yz* is the specific version (e.g.,beta-15.17)

Note: If a version is not specified for a user, the company version is the default. See [Company Manager \(admin only\)](#) on page 482 for more information.

Restricted IPs

The IP addresses from which this user is restricted to log into 1010data.

Note: This field is editable only by company administrators.

API Access

This specifies if the user has access to the 1010data API. (*read only*)

Client time-out period

When running long queries, your browser may time out after several minutes. You should be able to eliminate many of the time-outs by changing this setting to the time-out interval (or lower).

For example, if you are timing out after five minutes, set this to five minutes. If you are timing out after seven-and-a-half minutes, set it to seven minutes.

Please be aware that when a time-out period is set, your browser's progress bar may not function properly.

Note: Time-outs are usually caused by your proxy server or firewall. If you experience such time-outs, first speak to your technology support person to see if they can change your proxy server or firewall settings. If that's not possible, then try changing this setting.

Maximum Download (Cells)

The number of cells a user can download in a single transaction.

Note: This field is editable only by company administrators.

Unlimited

When this checkbox is selected, there is no limit on the number of cells a user can download in a single transaction.

Note: This field is available only to company administrators.

Maximum Download (Rows)

The number of rows a user can download in a single transaction.

Note: This field is editable only by company administrators.

Unlimited

When this checkbox is selected, there is no limit on the number of rows a user can download in a single transaction.

Note: This field is available only to company administrators.

Total Disk Space (GB)

The amount of disk a user can use for saving tables.

A default of ~5GB is good as it will prevent large save tables that might be unintentional but will allow a regular user to make small tables. This is a ceiling for the disk space a user can use, but be aware that the value this compares against is not computed with 100% accuracy.

Disk Space Used (GB)

The total amount of disk space a user has already used for saving tables. *(read only)*

Total Disk Space for FTP storage (GB)

The amount of disk space a user can use for uploading tables via FTP. *(read only)*

To request FTP access, please email support@1010data.com.

Worksheet Limit

The maximum size of the foreign table when linking worksheets. *(read only)*

Note: This field is available only to company administrators.

Vendor ID

The ID used to identify the user's company in the vendor portal.

Vendor Name

The description of the user's company in the vendor portal.

Vendor Access Level

The access level of the user's company in the vendor portal.

GUI Preferences

The **GUI Preferences** tab of the **User Manager** page contains information related to a particular user's preferences in the 1010data web interface.

General

Favorite Folders and Tables (ids)

A space-separated list of numeric table IDs that the user has flagged as favorites.

Allow 0 row tables when the macro is annotated with empty

When selected, this allows selections resulting in tables with 0 rows, if the macro contains `<meta>empty<meta>`.

If this is not selected (or if this is selected, but the macro does not contain the empty `<meta>` element), the user will receive an error saying that no rows were selected.

Location

Time Zone

The time zone that is used when showing the date and time on reports.

The options are presented as offsets to the Greenwich Mean Time.

For example, New York is in the Eastern Time Zone of the United States, and Eastern Standard Time (EST) is 5 hours behind Greenwich Mean Time. Therefore, the appropriate selection for New York would be GMT-5.

Region Format

This determines the order in which dates are presented (or downloaded as formatted values), when using the `date`, `date4y`, and `datehms24` display formats as well as how numbers (i.e., integers and decimals) are displayed.

	Date	Date+Time	Integer	Decimal
USA	<i>MM/DD/YY</i>	<i>MM/DD/YY_HH:MM:SS</i>	10,000	0.001
Europe	<i>DD.MM.YY</i>	<i>DD.MM.YY_HH:MM:SS</i>	10.000	0,001
United Kingdom	<i>DD/MM/YY</i>	<i>DD/MM/YY_HH:MM:SS</i>	10,000	0.001
ANSI	<i>YYYY-MM-DD</i>	<i>YYYY-MM-DD HH:MM:SS</i>	10,000	0.001

User Interface

Return focus to browser window

When selected, the browser window regains focus after a query completes. If the browser window was behind other windows or was minimized, it becomes the active window.

By default, in certain operating systems (e.g., pre-Windows XP), whenever you get a response from the system, the browser window grabs the focus. You can disable this behavior by clearing this check box.

Enable keyboard shortcuts

When selected, keyboard shortcuts for switching tabs, editing actions, showing column information, and searching for items are enabled.

Show "Run Macro" in the "File" menu

In earlier versions of 1010data, queries could be saved as macros and run in this way. This option is provided for legacy reasons.

Display progress bar

When selected, a dialog containing a progress bar will appear while a query is processing. The progress bar shows the percentage of total operations completed within the current query.

Show advanced query error information

When selected, more detailed information related to query errors is presented.

Show column headings as

The column heading, name, or a combination of both can be displayed at the top of each column in a table.

the label

The column heading will be displayed at the top of each column (e.g., **Unemployment Rate**).

the name

The column name will be displayed at the top of each column (e.g., `unemp_rate`).

both the name and label

Both the column heading and the column name will be displayed at the top of each column.

Handling of repeated values within a column

This determines how repeated values will be displayed within a particular column in a table.

the value

If a particular cell has the same value as the cell in the row above it, the same value will be displayed in that cell. This is the default setting.

a ditto mark (all columns)

If a particular cell has the same value as the cell in the row above it, a ditto mark ("") will be displayed in that cell.

a ditto mark (fixed columns only)

Within a fixed column, if a particular cell has the same value as the cell in the row above it, a ditto mark ("") will be displayed in that cell. Otherwise, if it is not a fixed column, the same value will be displayed.

a blank (all columns)

If a particular cell has the same value as the cell in the row above it, the cell will appear blank.

a blank (fixed columns only)

Within a fixed column, if a particular cell has the same value as the cell in the row above it, the cell will appear blank. Otherwise, if it is not a fixed column, the same value will be displayed.

In selection boxes, show columns in

This determines how column headings will appear within the drop-down lists in row selection dialogs.

the order that they appear in the table

Column headings will appear within the drop-down lists in the same order that they appear in the table.

alphabetical order

Column headings will appear in the drop-down lists in alphabetical order and will include column names for any linked tables.

alphabetical order but with columns from linked tables last

Column headings will appear in the drop-down lists in alphabetical order; however, column headings for any linked tables will appear separately (in alphabetical order) at the bottom of the list.

Use combo boxes instead of select boxes

When selected, combo boxes will be used instead of select boxes for the drop-down lists in all dialogs.

Number of decimal places

Specifies the number of decimal places to display for numbers in the 1010data grid.

Number of Summarizations for Quick Summary

The number of summarizations available The default is 10.

Number of Grouping Columns for Tabulation

The number of grouping columns available for a tabulation. The default is 3.

Number of Summarizations for Tabulation

The number of summarizations available in a tabulation. The default is 10.

Number of Row Dimensions for Cross Tabulation

The number of rows of results available in cross tabulation. The default is 3.

Number of Column Dimensions for Cross Tabulation

The number of columns of results available in a cross tabulation. The default is 3.

Number of Columns for Links

The number of available columns to match when linking tables or worksheets using the link dialogs. The default is 3.

Default folder path to save a quick query or new table

The path to the folder where quick queries and new tables are saved by default. If no path is specified, the user's **My Data** folder is the default location.

Macro Language

Enable syntax highlighting

When selected, Macro Language elements, attributes, and values are color coded within the Macro Language Workshop.

Wrap text on Edit Actions page

When selected, text that is too long to fit on one line in Macro Language Workshop is wrapped to the next line.

Note: This option is only available when **Enable syntax highlighting** is selected.

Automatically indent on Edit Actions page

When selected, the line following an opening tag for a particular Macro Language element is indented by the number of spaces specified by **Indent spaces for nested elements**.

If this option is not selected, the line following the opening `<tabu>` tag is not indented.

Enable autocomplete

When selected, a drop-down menu of context-related options is presented as you type in Macro Language element names, attributes, and values for certain attributes (such as functions) in Macro Language Workshop.

For instance, if you type `<` in Macro Language Workshop, a list of context-dependent element names will be presented.

If you type `fun="` when you are specifying a tabulation column (`<tblcol>`) element, a list of summarization functions will appear in the drop-down list.

Note: This option is only available when **Enable syntax highlighting** is selected.

Indent spaces for nested elements

The number of spaces to indent nested Macro Language elements. The default is 2.

Put each attribute on a separate line

When selected, each attribute for a particular Macro Language element will appear on its own line.

The default is for all attributes to appear on the same line.

Break up long selection and value expressions

When selected, this breaks up lengthy `value` expressions in `<sel>` statements over multiple lines, with each selection criterion on a separate line, prefaced by an `&`.

In a `<link>`, show the linked table ID instead of its name

Select whether a table ID or a table name (path) will be displayed as the value of the `table2` attribute for a `<link>` element in the Macro Language.

name

In Macro Language Workshop, the value of the `table2` attribute for a `<link>` element is displayed as the table's name.

ID

In Macro Language Workshop, the value of the `table2` attribute for a `<link>` element is displayed as the table's ID.

Automatically add comments for the following

When selected, a `<note>` element containing a comment is automatically added for the table that the macro is being applied to and the table that is being linked to.

Downloads

Excel Downloads

Choose whether tables downloaded in Excel format should be decorated with colorful borders, fills, and fonts (and include tabulation totals), or if they should appear as plain data.

Decorated with totals

Tables will be decorated with colorful borders, fills, and fonts. Also, totals will appear at the top of tabulation columns.

Plain

Tables will not be decorated with any special formatting, and totals will not appear at the top of tabulation columns.

End-of-record delimiter

Select whether a line feed, or carriage return and line feed, are used to signify the end of a record.

LF

A line feed will be used as the end-of-record delimiter.

CRLF

A carriage return and line feed will be used as the end-of-record delimiter.

Include an end-of-record delimiter after the last record

Select whether an end-of-record delimiter should be included after the last record.

Advanced

Display NA for error or "infinity" values

Some computations (e.g., dividing by zero) can result in errors or "infinity" values (`0I`, `-0I`, `0i`, or `-0i`). Selecting this option can help reduce or eliminate this behavior.

Specifically:

- Division by zero ($X/0$) returns `0i` or `-0i` if the numerator is not zero.

Checking this box causes `N/A` to be returned instead of `0i` and `-0i`.

- Exponentiation (X^Y) can generate error messages if X is negative.

Checking this box mostly eliminates such messages and causes `N/A` to be returned in the event of an error.

- The range functions (`range1`, `range1f`, etc.) return `0I`, `-0I`, `0i`, or `-0i` for column values that lie outside the specified ranges.

Checking this box causes `N/A` to be returned instead.

Do step-wise aggregation

Select this option to help alleviate virtual memory problems when working with large tables. Using this feature saves memory but is slower.

Blocking Level

Select lower numbers to help alleviate virtual memory problems when doing row selections and tabulations on large tables. Lower numbers save memory but may cause queries to run slower.

Edit your user information

You can change user information, account settings, and GUI preferences associated with the 1010data user ID that you used to log into the current session.

To edit the information associated with your 1010data user ID:

1. Under the drop-down menu corresponding to your username in the top right corner of your 1010data session, click **Admin**.

The **Account Administration** page is opened in a new tab.

2. On the **Account Administration** toolbar, click the **User Manager** icon ().

Note: If the **User Manager** icon () is not available in the **Account Administration** toolbar, contact support@1010data.com to have it enabled. If you cannot access the User Manager, you can still change certain items in your user profile. See [User Profile](#) on page 75 for more information.

3. Make your changes to the desired fields under **User Info**, **Account Settings**, and **GUI Preferences**.
4. Click **Save User**.

If your changes have been successfully saved, you will see the message: `Changes were saved!`

Create a new user (admin only)

A company administrator can create a new user in 1010data.

Note: If you have a large number of users to create, or simply know of a user whose preferences and settings may already match your new user, you can use that existing user as a template for creating the new user. Simply find the existing user you want to use as the template, select that user (as if you were going to edit the user's information), modify the fields with the new user's information (at minimum, you should update all the required fields), and click **Create New User**.

Alternatively, you can add multiple users by providing a CSV file containing all the relevant information for each user. See [Create multiple users \(admin only\)](#) on page 473 for more information.

To create a new user:

1. Under the drop-down menu corresponding to your username in the top right corner of your 1010data session, click **Admin**.

The **Account Administration** page is opened in a new tab.

2. On the **Account Administration** toolbar, click the **User Manager** icon ().
3. Enter the information in the **User Info**, **Account Settings**, and **GUI Preferences** tabs for the new user.

Note: Required fields are denoted by a red asterisk.

4. Click **Create New User**.

If you receive a message similar to the following: `Transaction failed: requid is not authorized to assign this server`, contact support@1010data.com.

Create multiple users (admin only)

A company administrator can create multiple 1010data users by uploading a CSV file containing the users' information.

You need specific access to the bulkedit/bulkadd functionality. To request this functionality, please email support@1010data.com.

In addition, you must have a CSV file containing the users' information. For details on what that file may contain as well as an example, see [Bulk edit file format](#) on page 474.

To create multiple users:

1. Under the drop-down menu corresponding to your username in the top right corner of your 1010data session, click **Admin**.

The **Account Administration** page is opened in a new tab.

2. On the **Account Administration** toolbar, click the **User Manager** icon ().
3. Click **Create Multiple Users**.
4. Click **Choose File** (or **Browse**) and select the CSV file containing the users' information.
5. Click **Upload**.

This uploads the file to your **My Data** folder, reads the data, and creates the new ids.

During this process, you'll see some of the following status messages appear:

1. Your file is uploading...
2. Upload Successful! or Upload failed: *error*
3. Reading table...
4. Table Read Successful! or Failed to read uploaded table: *error*
5. Creating users...

If all users were created successfully, the status changes to `Users Created!`, and the newly created user IDs are displayed in an area below the search box. You are then automatically placed in bulk edit mode, where you can make changes to the IDs you just created and can click **Save Multiple Users** to save those changes. To exit this mode, click **Clear Form**.

If there were any errors with the creations, the status changes to `BULKEDIT: Some users failed`, see above error msg(s), and any error messages are displayed in an area below the search box.

If you receive a message similar to the following: `Transaction failed: requid is not authorized to assign this server`, contact support@1010data.com.

Bulk edit file format

Company administrators can create multiple users at one time by providing a CSV file containing the information for those users.

User Fields

The following table contains all available user fields along with their descriptions, valid values, and sample inputs.

For more detailed information about each of these fields, see:

- [User Info](#) on page 464
- [Account Settings](#) on page 466

Field Name	Description	Type	Sample Input
uid	User ID	text	mycomp_jsmith
first	First name	text	John
last	Last name	text	Smith
mi	Middle initial	text	R

Field Name	Description	Type	Sample Input
email	Email address	text	jsmith@mycomp.com
emaillist	Email list	text	johnsmith@gmail.com
company	Company name	text	My Company
phone	Phone number	text	123-456-7890
street	Street address	text	10 Main Street
city	City	text	New York
state	State	text	NY
country	Country	text	USA
zip	Zip code	text	10101
profess	Profession	text	Analyst
dormantmin	Idle time (minutes)	integer	60
inactive	Deactivate	Valid values are: <ul style="list-style-type: none"> • 0 • 1 	0
expiredate	User expiration date	integer (YYYYMMDD)	20201231
pwdate	Password set date	integer (YYYYMMDD)	20010101
pwdatecycle	Password expiration cycle	integer	45
version	Version	text	prod-9.53
ip	Restricted IPs	text	127.0.0.1
timeout	Client time-out period (seconds)	integer	600
maxdown	Maximum download (cells)	integer	100000
maxrows	Maximum download (rows)	integer	100000
totdisk	Total disk space (GB)	integer	5000000
timezone	Time zone (GMT offset)	integer	-5
region_fmt	Region format	text	US
favorites	Favorite folders and tables (IDs)	space-separated list of integers	174638 374757
err_stack_info	Show advanced query error information	Valid values are: <ul style="list-style-type: none"> • 0 • 1 	0
focus	Focus browser on query completion	Valid values are: <ul style="list-style-type: none"> • 0 • 1 	1
shortcut	Enable keyboard shortcuts	Valid values are:	1

Field Name	Description	Type	Sample Input
		<ul style="list-style-type: none"> • 0 • 1 	
allow_empty_tables	Allow 0 row tables when macro is annotated with empty	Valid values are: <ul style="list-style-type: none"> • 0 • 1 	1
grid_labels	Grid labels	Valid values are: <ul style="list-style-type: none"> • name • label • both 	both
showids	Show table IDs instead of paths in Macro Language code	Valid values are: <ul style="list-style-type: none"> • 0 • 1 	0
ditto	Handling of repeated values within column	Valid values are: <ul style="list-style-type: none"> • value • mark • markfx • blank • blankfx 	blank
sortcols	Column order in selections	Valid values are: <ul style="list-style-type: none"> • 0 • 1 • 2 	0
combo	Use combo boxes	Valid values are: <ul style="list-style-type: none"> • 0 • 1 	0
inf2na	Reduce the number of 0i's	Valid values are: <ul style="list-style-type: none"> • 0 • 1 	0
excel_dl	Excel downloads	Valid values are: <ul style="list-style-type: none"> • new • old 	old
down_eor	End-of-record delimiter in downloads	Valid values are: <ul style="list-style-type: none"> • lf • crlf 	crlf
down_last	Include a record delimiter at end of file downloads	Valid values are: <ul style="list-style-type: none"> • 0 • 1 	1
tab_condense	Do step-wise aggregation	Valid values are:	1

Field Name	Description	Type	Sample Input
		<ul style="list-style-type: none"> 0 1 	
s_res_n	Number of summarizations for quick summary	integer	20
maxrows_blk	Tabulation block size	integer	2
t_breaks_n	Number of grouping columns for tabulation	integer	6
t_res_n	Number of summarizations for tabulation	integer	20
x_rbreaks_n	Number of row dimensions for cross tabulation	integer	10
x_cbreaks_n	Number of column dimensions for cross tabulation	integer	10
l_cols_n	Number of columns for links	integer	20
syntax_h	Enable syntax highlighting in Edit Actions (XML) dialog	Valid values are: <ul style="list-style-type: none"> 0 1 	1
oe_wrap	Wrap text on Edit Actions (XML) dialog	Valid values are: <ul style="list-style-type: none"> 0 1 	1
progressbar	Show the progress bar	Valid values are: <ul style="list-style-type: none"> 0 1 	1
auto_hint	Enable autocomplete in Edit Actions (XML) dialog	Valid values are: <ul style="list-style-type: none"> 0 1 	1
nestmacro	Number of spaces to indent nested elements in Macro Language code	integer	3
splitattr	Show each attribute on a separate line in Macro Language code	Valid values are: <ul style="list-style-type: none"> 0 1 	0
splitexpr	Show each ANDed expression on a separate line in Macro Language code	Valid values are: <ul style="list-style-type: none"> 0 1 	0
autocomment	Autocomment Macro Language code	Valid values are: <ul style="list-style-type: none"> base link base link 	link

Example

The following is an excerpt from a sample CSV file:

```
uid,first,last,email,dormantmin,expiredate,progressbar
mycomp_jsmith,John,Smith,jsmith@mycomp.com,60,20201231,0
mycomp_manderson,Maggy,Anderson,manderson@mycomp.com,60,,0
```

Find existing users (admin only)

A company administrator can search for existing 1010data users based on certain criteria such as user ID, first name, last name, email, or company name. They could also get a list of all users by entering no criteria at all.

To find existing users:

1. Under the drop-down menu corresponding to your username in the top right corner of your 1010data session, click **Admin**.

The **Account Administration** page is opened in a new tab.

2. On the **Account Administration** toolbar, click the **User Manager** icon ().
3. In the text box at the top of the form, enter the first name, last name, ID, email, or company name of the user(s) you want to find.

Note: To see a list of all the users you have permission to access, leave the text box empty.

4. Click **Find Users** (or press **Enter**).

A list of users that match your search criteria is presented beneath the search text box. If no users are found matching your search criteria, you will see the message: "No users found".

Note: The search is not case sensitive, and partial matches are listed.

In the results list, click the desired **User ID** to populate the **User Manager** with its information, or click anywhere outside the results list to dismiss it.

Edit an existing user (admin only)

A company administrator can edit an existing 1010data user's information.

To edit an existing user's information in 1010data:

1. Under the drop-down menu corresponding to your username in the top right corner of your 1010data session, click **Admin**.

The **Account Administration** page is opened in a new tab.

2. On the **Account Administration** toolbar, click the **User Manager** icon ().
3. In the text box at the top of the form, enter the first name, last name, ID, email, or company name of the user(s) you want to find.

Note: To see a list of all the users you have permission to access, leave the text box empty.

4. Click **Find Users** (or press **Enter**).

A list of users that match your search criteria is presented beneath the search text box. If no users are found matching your search criteria, you will see the message: "No users found".

Note: The search is not case sensitive, and partial matches are listed.

5. In the search results list, click the username whose information you want to edit.
6. Make your changes to the desired fields in the **User Info**, **Account Settings**, and **GUI Preferences** tabs.
7. Click **Save User**.

If you receive a message similar to the following: Transaction failed: requid is not authorized to assign this server, contact support@1010data.com.

If your changes have been successfully saved, you will see the message: `Changes were saved!`

Edit multiple existing users (admin only)

A company administrator can edit the information for more than one existing 1010data user at the same time. This is helpful when you may want to make changes to certain fields or settings and apply those changes to a number of users.

You need specific access to the `bulkedit/bulkadd` functionality. To request this functionality, please email support@1010data.com.

To edit multiple existing 1010data users' information:

1. Under the drop-down menu corresponding to your username in the top right corner of your 1010data session, click **Admin**.

The **Account Administration** page is opened in a new tab.

2. On the **Account Administration** toolbar, click the **User Manager** icon ().
3. Specify the users whose information you want to edit by performing one of the following actions.
 - Press **Save Multiple Users** and enter a space- or newline-separated list of the user IDs you want to modify.
 - Find the users you want to modify (see [Find existing users \(admin only\)](#) on page 478), click the checkbox associated with each user whose information you want to edit (or select the **All Users** checkbox if you want to select all of the users that appear in the search results list), and click **Submit**.

Note: You can repeat this step to add more user IDs to the set that you are modifying.

4. Make your changes to the desired fields in the **User Info**, **Account Settings**, and **GUI Preferences** tabs.
5. Click **Save Multiple Users**.

A notification dialog is displayed that says any populated field on the User Manager tabs will replace the corresponding field in the selected users.

6. Click **OK** if you wish to continue.

If your changes have been successfully saved, you will see the message: `Users Updated!`

Deactivate an existing user (admin only)

A company administrator can deactivate an existing user in 1010data.

To deactivate an existing user:

Note: 1010data does not allow you to delete users, only deactivate them.

1. Under the drop-down menu corresponding to your username in the top right corner of your 1010data session, click **Admin**.

The **Account Administration** page is opened in a new tab.

2. On the **Account Administration** toolbar, click the **User Manager** icon ().
3. In the text box at the top of the form, enter the first name, last name, ID, email, or company name of the user(s) you want to find.

Note: To see a list of all the users you have permission to access, leave the text box empty.

4. Click **Find Users** (or press **Enter**).

A list of users that match your search criteria is presented beneath the search text box. If no users are found matching your search criteria, you will see the message: `"No users found"`.

Note: The search is not case sensitive, and partial matches are listed.

5. In the search results list, click the username that you want to deactivate.

6. Click **Deactivate User**.

You will be presented with a dialog confirming that you want to deactivate the user.

7. Click **OK**.

The user will be deactivated and will no longer be able to log in to 1010data and access the system. On their next attempt to log in, they will be notified that their user ID has been suspended and will be prompted to contact support@1010data.com.

You can also deactivate a user by editing that user's information and setting the **Deactivate** field in the **User Info** tab to **Yes**, then clicking **Save User**. Conversely, you can reactivate a user by setting the **Deactivate** field to **No**.

Group Manager

The Group Manager allows you to create, find, edit, and delete your own groups in 1010data.

The Group Manager provides a centralized place where you can manage your groups. A group is a set of users and/or other groups. Groups can help you manage the permission you give users to view and edit such items as your folders, tables, and queries. You can share access with a group, thereby eliminating the need to share access with each individual in the group.

Note: You can only manage groups that you own; however, company administrators can manage any group owned by any ID in the company.

Create a new group

Create a new group of 1010data users.

To create a new group in 1010data:

1. Under the drop-down menu corresponding to your username in the top right corner of your 1010data session, click **Admin**.

The **Account Administration** page will be opened in a new tab.

2. On the **Account Administration** toolbar, click the **Group Manager** icon ().
3. In the **Group ID** text box, enter the name of the group. (*required*)

The **Group ID** can only contain letters, numbers, and underscores. It cannot contain any spaces or other special characters.

Note: You cannot change the name of the group once it has been created.

4. Select which value to use for the **Group ID prefix**.

The selected value will be automatically prepended to the value you enter in the **Group ID** text box:

- **User ID** - Your user ID will be automatically prepended to the value you enter in the **Group ID** text box (e.g., rd_jsmith_marketingteam).
- **Company ID** - Your company ID will be automatically prepended to the value you enter in the **Group ID** text box (e.g., acme_supportgroup).

5. In the **Owner** box, enter the name of the owner of the group. (*required*)

This field will contain your user ID by default, but it may be changed so that you can assign a different owner to the group. This must be a valid user ID; otherwise, you will receive an error when you try to create the group.

6. In the **Users** box, enter a space-separated list of usernames that you would like to include in this group. (*optional*)

If there are any invalid usernames (i.e., do not exist, do not have the same company ID), an error message will be returned when you try to create the new group.

7. In the **Title** text box, enter a title to help identify the group. *(optional)*

The title may contain any combination of uppercase and lowercase letters, numbers, spaces, and special characters.

8. In the **Description** text box, enter a description about the group. *(optional)*

9. Click **Create New Group**.

The group is created.

The following values in the **Group Manager** are populated:

- **Group Creation Time**
- **Previous Time Group was Updated**
- **User to Create Group**
- **Previous User to Update Group**
- **Previous Version to Update Group**

Note: These values cannot be changed by the user.

Find an existing group

Find an existing 1010data group.

To find an existing group in 1010data:

1. Under the drop-down menu corresponding to your username in the top right corner of your 1010data session, click **Admin**.

The **Account Administration** page will be opened in a new tab.

2. On the **Account Administration** toolbar, click the **Group Manager** icon ().
3. In the text box at the top of the form, enter the group title, group owner, or group ID associated with the group you want to find, or leave the text box empty to see a list of all the groups you own.
4. Click **Find Groups** (or press **Enter**).

A list of groups that match your search criteria is presented beneath the search text box. If no groups are found matching your search criteria, you will see the message: "No groups found".

Note: The search is not case sensitive, and partial matches are listed.

In the results list, click on a group to populate the **Group Manager** with its information, or click anywhere outside the results list to dismiss it.

Edit an existing group

Edit the details of an existing 1010data group.

To edit an existing group in 1010data:

1. Under the drop-down menu corresponding to your username in the top right corner of your 1010data session, click **Admin**.

The **Account Administration** page will be opened in a new tab.

2. On the **Account Administration** toolbar, click the **Group Manager** icon ().
3. Find the group that you want to edit.
4. Make your desired changes:

- a) In the **Users** box, enter a space-separated list of usernames that you would like to include in this group. *(optional)*

If there are any invalid usernames (i.e., do not exist, do not have the same company ID), an error message will be returned when you try to save the changes.

- b) In the **Title** text box, enter a title to help identify the group. *(optional)*
- c) In the **Description** text box, enter a description about the group. *(optional)*

Note: You cannot change the **Group ID** or the **Owner** once the group has been created.

5. Click Save Group.

If your changes have been successfully saved, you will see the message: `Changes were saved!`

Delete an existing group

Delete an existing 1010data group.

You must be the owner of a group to delete it.

To delete an existing group in 1010data:

1. Under the drop-down menu corresponding to your username in the top right corner of your 1010data session, click **Admin**.

The **Account Administration** page will be opened in a new tab.

2. On the **Account Administration** toolbar, click the **Group Manager** icon ().
3. Find the group that you want to delete.
4. Click **Delete Group**.

You will be presented with a dialog confirming that you want to delete the group.

5. Click **OK**.

Company Manager (admin only)

The Company Manager allows a company administrator to change company-related information within 1010data.

Company ID

A unique combination of alphanumeric characters and underscores specifying the ID of the company. *(read only)*

Company

The name of the company.

Administrator ID

The user ID of the company administrator(s). *(read only)*

Administrator

The full name of the company administrator.

Maximum # of IDs

The maximum number of active user IDs that a company is allowed to own. *(read only)*

Maximum # of concurrent logins

The maximum number of users in the company that can be logged in at the same time. *(read only)*

Users (read only)

A listing of all user IDs owned by the company. *(read only)*

Active Users count

The number of active user IDs owned by the company.

The **Calculate** button initiates this calculation.

Note: This calculation can be resource intensive. If the company owns many user IDs, this action could take some time and possibly lock the session.

Street Address

City

State

Zip Code

Country

The full address of the company.

Email Address

The primary email address associated with the company.

Email List

A space-separated list of additional email addresses associated with the company.

Phone Number

The phone number of the company.

Version

The version of 1010data the company (and, by default, all users within the company) is configured to use.

Possible values are:

- `prod-latest`
- `beta-latest`
- `prod-x.yz`
- `beta-x.yz`

where `x.yz` is the specific version (e.g., `prod-6.92`, `beta-7.36`)

Note: The company version is used by default for users who do not have a specific version set for their particular ID.

Company Creation Time

The date (in `YYYYMMDD` form) that this company was originally created in 1010data. *(read only)*

Previous Time Company was Updated

The date (in `YYYYMMDD` form) this company's information was last updated. *(read only)*

User to Create Company

The user ID of the administrator who created this company in 1010data. *(read only)*

Previous User to Update Company

The user ID of the administrator who previously updated this company's information. *(read only)*

Previous Version to Update Company

The version of 1010data that was used when this company's information was last updated. *(read only)*

Edit company information (admin only)

A company administrator can edit an existing company's information in 1010data.

To edit a company's information in 1010data:

1. Under the drop-down menu corresponding to your username in the top right corner of your 1010data session, click **Admin**.

The **Account Administration** page will be opened in a new tab.

2. On the **Account Administration** toolbar, click the **Company Manager** icon ().
3. Make your changes to the desired fields.
4. Click **Save Company**.

If your changes have been successfully saved, you will see the message: `Changes were saved!`

System Statistics (admin only)

System Statistics allows you view the current memory usage of your company.

System Statistics is a tool for administrators to view the current memory usage and number of active users on the 1010data Insights Platform.

The System Statistics window is accessed by selecting **Admin > System Statistics** from the workspace menu.

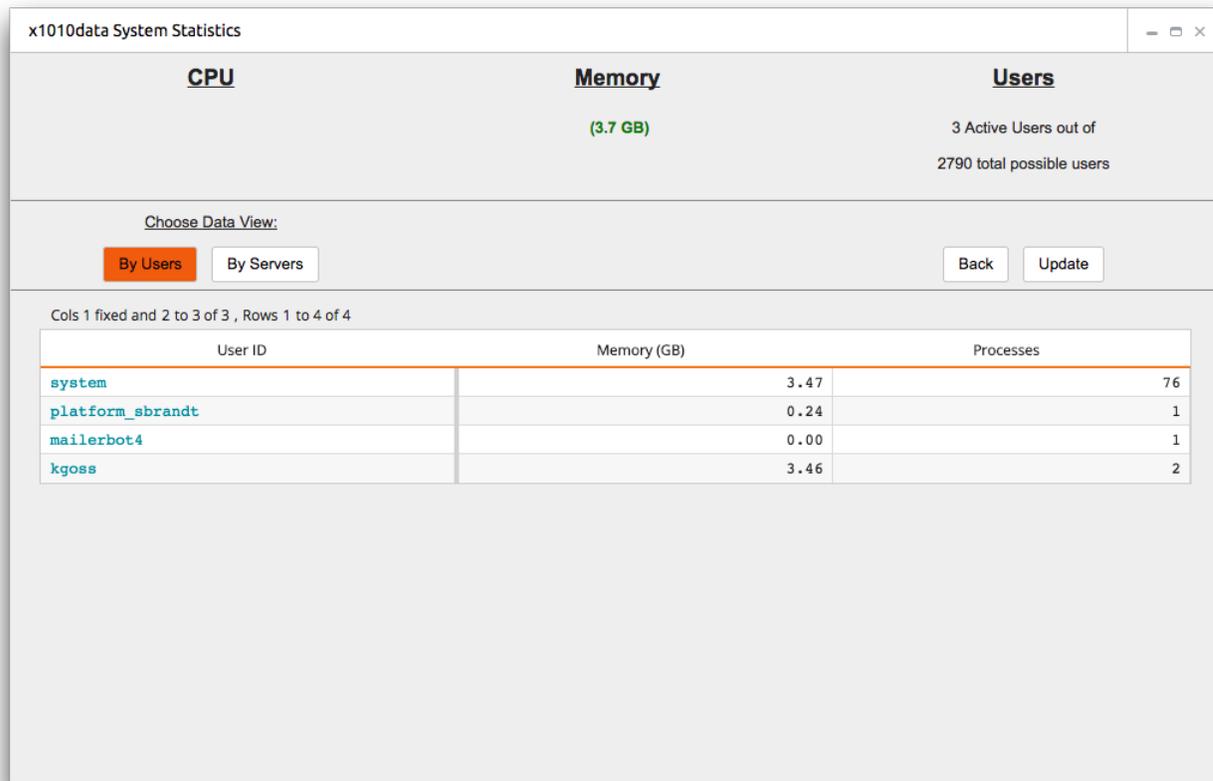


Figure 128: The System Statistics Window

CPU

This statistic is not currently available.

Memory

Your company's current memory usage

Users

The current number of active users and the total number of possible users

Choose Data View

Filter what data is viewed in the table

You can view the statistics in the table broken down by users or broken down by what servers are being used.

Statistics table

The information about the current system usage broken down by user or server

The statistics table displays data about memory usage and the number of processes currently running.

Appendix A: Missing values

Most databases contain at least some missing data values. The values may be missing due to non availability, non applicability, data processing errors, or even intentional omission. Missing values are also referred to as "N/A" values (for "not available" or "not applicable").

N/A values can also be created as a result of certain analyses. Here are some examples:

- A cross tabulation produces results for every combination of values from the specified grouping columns. If there is no data for certain combinations, those results are N/A. For example, suppose we are analyzing a company's sales. The database shows the details of each transaction, including the product sold, the location in which it was sold, and dollar value of the sale. What we would like to know is the average sale size in dollars for each product in each location, a natural cross tabulation. But suppose that a particular product is not sold in a particular location. The average sale size for that product and location is N/A and shows as blank in the result of the cross tabulation.
- When tables are linked, if a particular row in the current table has no match in the "foreign" table, the row is padded with N/A values. For example, suppose we are looking at employee data and we have two tables that give information about each employee. The first shows the address for each employee and the second shows salary-related information. If we were looking at the address table and linked in the salary table, we would have a table that showed both address information and salary information. But suppose that, for some reason, there is no salary information for a particular employee; the salary for that employee shows as blank (N/A) in the combined table.
- Computed columns can have N/A values. This usually happens when one or more of the referenced columns contain N/As, but certain computations can produce N/A results even if the inputs are not N/A.

Accurate analysis requires understanding the nature and source of N/A values and the appropriate ways of dealing with them. The following sections describe how N/A values are represented in the system and how they are handled in selection and value expressions and tabulations. Some additional techniques for handling such values are discussed as well.

Representing N/A values

The first thing we should say is that missing values aren't really missing. That may sound like a contradiction in terms, but here is what we mean: When you see a blank value, it isn't that there is no value there at all, rather there is a special value that indicates that meaningful data is missing. The special value may display as blank, but in the database there really is a value. That is why we prefer to refer to a missing value as an N/A value.

Special values are commonly used to represent missing or meaningless data, although different people use different values. The number 999, for example, it often used as a numeric N/A value, especially in a column containing whole numbers. Take the table:

Name name	Age age
Susan	31
Tom	25
Dick	40
Harry	999
Jane	34

Since it is unlikely that anyone will actually be 999 years old, it is safe to use 999 as the N/A value for age. 999 also has the nice property that it is easily recognizable. Similarly, values like 999.99 or 999.999 are often used for decimal numbers.

In the 1010data Insights Platform, we use our own special values. For example, in columns that contain text, we use ' ' (the empty string). In numeric columns, we use one of two other values, depending on

whether the column contains whole numbers or decimal numbers. What these values are isn't really important; what is important is that they are treated differently than other values.

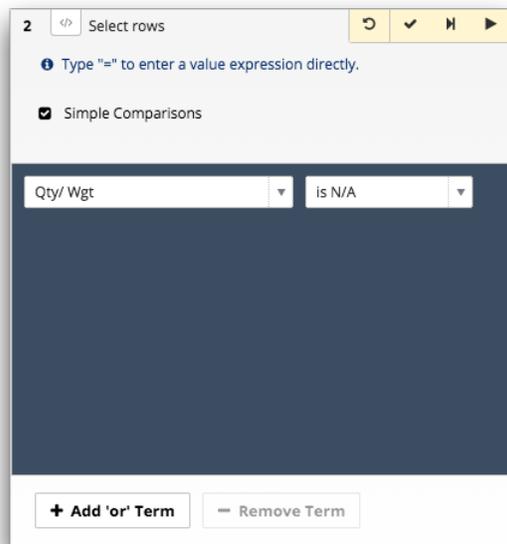
1. They display as blank.
2. They are treated specially in selection and value expressions and tabulations. For example, if a computed column's value expression is `col1+col2` and `col1` (or `col2`) is N/A on a particular row, the computed column will also be N/A on that row. See the following sections about N/A handling in selection and value expressions and tabulations.

Because the system automatically treats these values specially, we try to use them to represent N/As whenever we can. When loading data into a table, if the source data contains other types of N/A values (999, 999.99, NA, etc.), we usually convert those values into our N/A values. We do *not* convert such values unless the technical documentation for the data specifically says the values are in fact N/A values or if it is patently obvious that is the case. It is therefore entirely possible that some columns in some tables may contain 999, for instance. Please be aware of this when doing computations; if `col1` has 999s in it, and you compute `col1+col2`, the result may not be terribly meaningful!

Looking for N/A values

Before using a database, it is only prudent to see how many N/A values there are in various columns. There are many ways of doing this.

You can use the **Select rows** panel to select rows where a particular column is N/A. For example:



Or you can tabulate on the column to get all its unique values and look for N/A in the result.

But perhaps the most convenient method is to use the tabulation summary designed for this very purpose: **# invalid [COLUMN]**. In a quick summary, this allows you to determine the number of N/As in each of several columns in one operation and, in a tabulation or cross tabulation, to determine the number of N/As for each group. Related summarization methods include **# valid [COLUMN]**, **# valid pairs [COLUMN]**, and **# invalid pairs [COLUMN]**.

Automatic handling of N/A values

N/A values are handled automatically in selection and value expressions and tabulations.

N/A values in selection and value expressions

Assuming that missing values are represented by our special N/A values, what happens when these values are used in a computation? Specifically, if we create a computed column that references one or more other columns, and those columns contain N/As, what will be the value of the computed column?

For example, if `col3` is `col1+col2` and one or both of `col1` and `col2` are N/A, what will `col3` be? What if `col3` is `col1>col2` or `round(col1;10)`?

For most functions and operators, if any argument is N/A, the result is N/A, *but this is not always true*. See the description of each operator and function for details about N/A handling.

N/A values in tabulations

What happens when N/A values are encountered in a quick summary or tabulation?

For example, when summing the values in a particular group of rows, what if some of the values are N/A?

The answer is that, as a general rule of thumb, N/A values are ignored. In the case of sum, for example, the result is the sum of the non-N/A values. But there are exceptions to the rule. See the description of each summary type for details about N/A handling in quick summaries, tabulations, and cross tabulations.

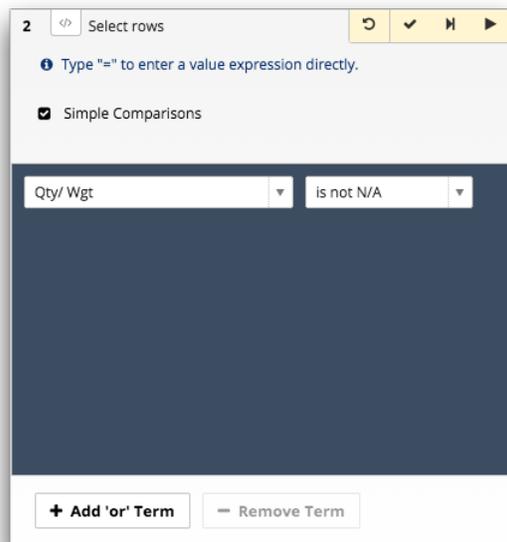
Manual handling of missing or invalid values

The 1010data Insight Platform's built-in N/A handling often provides satisfactory results, but there are cases where it makes sense to deal directly with N/As or invalid values.

These are some that are commonly used techniques for dealing with N/A or invalid values (i.e., values that are outside an expected range or have other problems).

Discarding Bad Values

The most obvious technique is to select only rows which have no bad (N/A or invalid) values before doing further analysis. For example, using the **Select rows** panel:



In Macro Language code, you may also use `NA` in the selection expression. For example, `quantity<>NA`.

Discarding all bad values has the advantage of simplicity and may work well with databases with relatively few such values, but it may not always be the best thing to do. Consider the case where a computation is performed using five columns. If, on a particular row, even one of those columns has a bad value, you will

be throwing out the entire row and losing the good values of the other four columns. In some cases, in fact, you could be throwing out the more important data just because you are missing some less important data.

Fixing Bad Values

Instead of throwing out good data, it may be possible to "fix" the bad data so that the analysis can be completed using as much information as possible. Specifically, it may be possible, using computed columns, to replace many of the bad values with plug values or values derived from other columns. Take the following table:

First Name	Last Name	Enrollment Year	Age at Enrollment	Program Type (# of years)	High School Graduation Year
first_name	last_name	enroll_year	enroll_age	program	hsgrad_year
Susan	Smith	2002	18	4	2002
Tom	Doe	2001	19	5	2001
Dick	Miller	2002	20		2000
Harry	Jefferson	2003	18	4	2003
Jane	Smith	2003		4	1990

Note that there are two pieces of information missing.

Suppose we want to answer the following question: What will be the average age of a student upon completion of his or her program? This is a simple enough calculation. First, define a computed column, `grad_age`, that gives the required age for each student. The value expression for this column is simply: `enroll_age+program`. Then use a quick summary to get the average for `grad_age`.

Let's see what happens when we do this. First we add the computed column:

First Name	Last Name	Enrollment Year	Age at Enrollment	Program Type (# of years)	High School Graduation Year	Age at College Graduation
first_name	last_name	enroll_year	enroll_age	program	hsgrad_year	grad_age
Susan	Smith	2002	18	4	2002	22
Tom	Doe	2001	19	5	2001	24
Dick	Miller	2002	20		2000	
Harry	Jefferson	2003	18	4	2003	22
Jane	Smith	2003		4	1990	

Because we were missing some information, two of the five results are N/A. This doesn't bode well, but we push on. Taking the average of the new column gives a final result of 22.7 years.

Can we do better? Well, consider that most college students are enrolled in four-year programs. (Let's assume this is a four-year college.) We will probably not be very wrong to assume that Dick Miller is enrolled in this type of program as well. So it would make sense to "fill in" the blank with 4. How do we do this? One way is to add a computed column called `adjusted_program` with the value expression:

```
if( program<>NA ; program ; 4 )
```

In other words, `adjusted_program` will have the same value as `program` unless `program` is N/A, in which case `adjusted_program` will have the value 4. Applying this to the original table gives:

College Students
Cols 1 to 8 of 8 , Rows 1 to 5 of 5

First Name	Last Name	Enrollment Year	Age at Enrollment	Program Type (# of years)	High School Graduation Year	Age at College Graduation	Adjusted Program
first_name	last_name	enroll_year	enroll_age	program	hsgrad_year	grad_age	adjusted_program
Susan	Smith	2002	18	4	2002	22	4
Tom	Doe	2001	19	5	2001	24	5
Dick	Miller	2002	20		2000		4
Harry	Jefferson	2003	18	4	2003	22	4
Jane	Smith	2003		4	1990		4

We have thus effectively filled in one of the missing values with a reasonable plug value.

How about the second missing value? It would be a shame to discard the data for Jane Smith, since she is clearly an older student. Ignoring her would give a distorted final result. Fortunately, we can apply another good guess to further clean the data. While people enroll in college at all sorts of ages, most people graduate high school at around age 18. (There are exceptions, of course, but suppose we know enough about this subject to feel comfortable with that assumption.) Given that we know when Jane graduated from high school (1990) and enrolled in college (2003), we can make a pretty good guess that she was 31 when she enrolled in college ($31=18+2003-1990$). So let's add another computed column, `adjusted_enroll_age`, with the value expression:

```
if( enroll_age > NA ; enroll_age ; 18+enroll_year-hsgrad_year )
```

College Students
Cols 1 to 9 of 9 , Rows 1 to 5 of 5

First Name	Last Name	Enrollment Year	Age at Enrollment	Program Type (# of years)	High School Graduation Year	Age at College Graduation	Adjusted Program	Adjusted Enrollment Age
first_name	last_name	enroll_year	enroll_age	program	hsgrad_year	grad_age	adjusted_program	adjusted_enroll_age
Susan	Smith	2002	18	4	2002	22	4	18
Tom	Doe	2001	19	5	2001	24	5	19
Dick	Miller	2002	20		2000		4	20
Harry	Jefferson	2003	18	4	2003	22	4	18
Jane	Smith	2003		4	1990		4	31

Now that we have essentially filled in all the missing values, let's get back to our original question: specifically, what will be the average age of a student upon completion of his or her program? As before, we will add a computed column, `adjusted_grad_age`, that gives the required age for each student, but this time we will use columns `adjusted_program` and `adjusted_enroll_age` instead of `program` and `enroll_age`. The value expression for this column is: `adjusted_enroll_age + adjusted_program` and the result is:

College Students
Cols 1 to 10 of 10 , Rows 1 to 5 of 5

First Name	Last Name	Enrollment Year	Age at Enrollment	Program Type (# of years)	High School Graduation Year	Age at College Graduation	Adjusted Program	Adjusted Enrollment Age	Adjusted Age at College Graduation
first_name	last_name	enroll_year	enroll_age	program	hsgrad_year	grad_age	adjusted_program	adjusted_enroll_age	adjusted_grad_age
Susan	Smith	2002	18	4	2002	22	4	18	22
Tom	Doe	2001	19	5	2001	24	5	19	24
Dick	Miller	2002	20		2000		4	20	24
Harry	Jefferson	2003	18	4	2003	22	4	18	22
Jane	Smith	2003		4	1990		4	31	35

Finally, we get the average for the last column and come up with 25.4. Compare this to the 22.7 that we got earlier and note that this is probably a more accurate result even if our assumptions aren't perfect. Even if, in reality, Dick is enrolled in a five-year program and Jane was 17 or 19 when she graduated high school, our new result of 25.4 is closer to the truth than 22.7 was.

Of course, it isn't always possible to fill in missing data, but when it can be done, it can be quite beneficial or even essential.

Index

1010data
log out [31](#)

A

active window
switch [22](#)
adding multiple users [473](#)
add subfolder [114](#)
admin [463](#), [482](#)
advanced upload [355](#)

B

bulkadd [473](#)
bulk add [473](#)
bulkedit [479](#)
bulk edit [479](#)

C

cache
clear [46](#), [48](#)
clone
MLW window [23](#)
code
submit [23](#)
column headings [48](#)
column labels [48](#)
column names [48](#)
Company Manager [482](#)
computed column
submit [25](#)
submit and run to end [25](#)
Computed column panel
shortcuts [25](#)
custom report
icon [89](#)

D

delete
folder [112](#)
object [112](#)
query [112](#)
table [112](#)
delimiter [333](#)

E

edit information
folder [108](#)
object [108](#)
query [108](#)
table [108](#)
Excel spreadsheets [333](#)

expression editor
keyboard navigation [27](#)
shortcuts [26](#), [26](#), [26](#)

F

favorites
add [110](#)
find
Macro Language view [25](#)
MLW code editor [23](#)
find next
Macro Language view [25](#)
MLW code editor [24](#)
folder
delete [112](#)
edit information [108](#)
icon [89](#)
move [112](#)
reorder [113](#)
search for [103](#)
view information [107](#)

G

Group Manager [480](#)
groups
create [480](#), [480](#)
delete [480](#), [482](#)
edit [480](#), [481](#)
find [480](#), [481](#)
GUI preferences [463](#)

I

Insights Platform
log in [13](#)
URLs [13](#), [32](#)
web addresses [13](#), [32](#)

K

keyboard navigation
expression editor [27](#)
object browser [27](#)
tabulation menu [27](#)
keyboard shortcuts [22](#)

L

login [13](#)

M

Macro Language view
shortcuts [25](#)

- memory
 - utilization [46](#)
 - virtual [46, 48](#)
- merge
 - table [219](#)
 - worksheet, existing [220](#)
 - worksheet, new [222](#)
- merged table
 - icon [89](#)
- MLW code editor
 - reference pane [23](#)
 - shortcuts [23](#)
- MLW window
 - clone [23](#)
- move
 - folder [112](#)
 - object [112](#)
 - query [112](#)
 - table [112](#)
- My Data [90, 112](#)

N

- navigation
 - keyboard [26](#)
- new operation
 - code [24](#)
 - compute [24](#)
 - select [24](#)
- New Operation panel
 - shortcuts [24](#)

O

- object
 - add to favorites [110](#)
 - delete [112](#)
 - edit information [108](#)
 - move [112](#)
 - reorder [113](#)
 - search for [103](#)
 - share [109](#)
 - view information [107](#)
- object browser
 - icons [89](#)
 - keyboard navigation [27](#)
 - Object Manager [88](#)
 - Trillion-Row Spreadsheet [121](#)
- Object Manager
 - object browser [88](#)
 - search [103](#)
- Other People's Data [90](#)
- owner
 - icon [90](#)

P

- Parameterized Quick Query
 - icon [89](#)
- password
 - forgot [14](#)
 - requirements [14](#)

- reset [14](#)
- permission [109](#)
- permissions
 - owner [90](#)
 - uploader [90](#)
- preferences
 - GUI [463](#)

Q

- queries
 - restructure [46](#)
- query
 - delete [112](#)
 - edit information [108](#)
 - move [112](#)
 - reorder [113](#)
 - search for [103](#)
 - view information [107](#)
- QuickApp
 - icon [89](#)
- Quick Query
 - icon [89](#)

R

- redo
 - Analyze view [24](#)
 - Macro Language view [25](#)
- reference pane
 - MLW code editor [23](#)
- releases
 - beta [32](#)
 - prime [32](#)
 - production [32](#)
- rename
 - window [21, 21](#)
- reorder
 - folder [113](#)
 - object [113](#)
 - query [113](#)
 - table [113](#)
- replace
 - Macro Language view [25](#)
 - MLW code editor [24](#)
- replace all
 - Macro Language view [25](#)
 - MLW code editor [24](#)

S

- saved code
 - icon [89](#)
- search
 - folder [103](#)
 - object [103](#)
 - Object Manager [103](#)
 - query [103](#)
 - table [103](#)
- select rows
 - submit [26](#)
 - submit and run to end [26](#)

- Select rows panel
 - shortcuts [26](#)
- session
 - close [31](#)
 - end existing [32](#)
 - re-enter existing [32](#)
- settings
 - account [463](#)
- share
 - object [109](#)
- shortcuts
 - Computed column panel [25](#)
 - keyboard [22](#)
 - Macro Language view [24, 25](#)
 - MLW code editor [23](#)
 - New Operation panel [24](#)
 - Select rows panel [26](#)
 - window manager [22](#)
- spreadsheets
 - Excel [333](#)
- subfolder
 - add [114](#)
- submit code
 - MLW code editor [23](#)

T

- table
 - delete [112](#)
 - edit information [108](#)
 - icon [89](#)
 - merge [219](#)
 - move [112](#)
 - reorder [113](#)
 - search for [103](#)
 - upload [337, 339](#)
 - upload, advanced [355](#)
 - view information [107](#)
- tabulation menu
 - keyboard navigation [27](#)

U

- undo
 - Analyze view [24](#)
 - Macro Language view [25](#)
- upload
 - table [337, 339](#)
 - table, advanced [355](#)
- uploader
 - icon [90](#)
- URLs [32](#)
- User Manager [463](#)
- users
 - adding [473](#)
 - adding multiple [473](#)
 - create [463](#)
 - deactivate [463, 479](#)
 - delete [479](#)
 - edit [463, 478, 479](#)
 - manage [463](#)
 - multiple [473, 479](#)

- search for [463](#)

V

- view information
 - folder [107](#)
 - object [107](#)
 - query [107](#)
 - table [107](#)
- virtual memory [46, 48](#)

W

- window
 - rename [21, 21](#)
 - switch active [22](#)
- window manager
 - shortcuts [22](#)
- worksheet
 - merge existing [220](#)
 - merge new [222](#)
- wsfull [46](#)