



Admin & User Documentation

August 2024 - R20

Data Navigation System™	5
Glossary of Terms	5
Datamap	5
Data Explorer	5
Answers	6
Answer Builder	6
Answer Explorer	6
Preview	6
Samples	6
Reasoner	6
Tag	7
Domains	7
Trino	7
User Accounts & RBAC	7
Users	8
Tenant Management	8
Navigate to User Management	8
Creating a User	9
Deleting a User	11
Roles	12
Deleting a role	13
Domains	14
Domains & RBAC	15
Domain Creation	16
Domain Assignment	17
Domain Searching & Filtering	18
Data Contracts	19



Creating a Data Contract	20
Applying the Data Contract	26
Audit Log	27
Ranger	29
Creating a User	30
Creating a Masking Rule	32
Creating a Role-Level Filter Rule	37
Cluster Management	40
Cluster Creation	41
System Management	44
Data Sources	45
Redshift	51
Connect Redshift	51
FAQs	53
Teradata	53
Connect Teradata	53
FAQs	55
Athena	55
Connect Athena	55
FAQs	57
Snowflake	57
Connect Snowflake	58
FAQs	59
PostgreSQL	59
Connect PostgreSQL	60
FAQs	61
Oracle	62
Connect Oracle	62
FAQs	63
Trino	64
Connect Trino	64
FAQs	65
MS SQL Server	66
Connect MS SQL Server	66
FAQs	67
Impala	68
Connect Impala	68



FAQs	70
Greenplum	70
Connect Greenplum	71
FAQs	72
ADLS	72
Connect ADLS	72
FAQs	74
BigQuery	74
Connect BigQuery	74
FAQs	75
DB2	75
Connect DB2	75
FAQs	77
Databricks	77
Connect Databricks	77
FAQs	79
GenericJDBC	79
Connect GenericJDBC	79
FAQs	81
Hive	81
Connect Hive	81
FAQs	83
Iceberg	83
Connect Iceberg	83
FAQs	85
MongoDB	85
Connect MongoDB	85
FAQs	86
MySQL	86
Connect MySQL	86
FAQs	88
Netezza	88
Connect Netezza	88
FAQs	90
SAPHANA	90
Connect SAPHANA	90
FAQs	92
SalesForce	92
Connect SalesForce	92
FAQs	93



ServiceNow	93
Connect ServiceNow	93
FAQs	95
Google Sheets	95
Connect Google Sheets	95
FAQs	96
API	97
Connect API	97
FAQs	98
ClickHouse	98
Connect ClickHouse	98
FAQs	99
AWS S3	100
Connect S3	103
FAQs	104
Data Explorer	104
Saved Filters	106
Filter and Sort Tables	113
Reasoner	114
Preview Data	115
Overview	115
Columns	117
Samples	118
Usage	119
Comments	121
Lineage	122
Endorsement	123
Datamaps	123
No-Code	125
SQL	132
Ask Reasoner	133
Data Pipelines	136
Pipeline Scheduling	137
Data Marketplace	138
Publishing a Data Product	139
Promotions	141
My Orders	143
Checkout	145



Answers	147
Search for Answers	148
Filter and Sort Answers	151
Requesting an Answer	152
Select an Existing Dataset	153
Explore	155
Interrogate	155
Query	157
Build a New Dataset	159
SQL	160
Dashboard	161
Embed Dashboard	164
Chatbot	164

Data Navigation System™

Data Navigation System™ helps Data and Analytics teams significantly reduce the time and effort needed to answer business questions with data.

Data Navigation System™ is an all-in-one data management solution that helps you discover, prepare, query, and visualize data quickly and with less effort.

Glossary of Terms

This section describes, at a high level, the various terms used throughout the Promethium product.

Datamap

A Datamap is the term for a query created within Promethium. It consists of the dataset, SQL, metadata, and operations reporting.

As you build a Datamap, the SQL will be auto-generated for you using our no-code flows, with the option to edit the SQL code directly.

Data Explorer

Data Explorer is your base of operations for exploring and managing your data sources and the data within those sources.

The default view is a table that can be filtered by a wide range of modifiers such as Domains, Source, Type (Table, View, Datamap), Tag, Reasoner Score, Recency, etc. In addition to these filter modifiers, you can also search directly via our advanced free text search.



Answers

Promethium is designed to help you answer business questions with data. Our Answers feature is where you will leverage the work of your Data Team by:

1. Creating a business related question
2. Finding or exploring questions that have been previously answered

Answer Builder

You can leverage the Reasoner to build on existing data sets or “Request an Answer” to raise a new question and assign it to a member of your team for collaboration. When you are happy with the output of your Answer you can publish it, along with any comments you have, for consumption by your peers.

Answer Explorer

Explore available answers built by collaborators and peers. If an Answer is not available, then the ability to request an Answer is an option. Typical Answers include a resultant, or “golden”, dataset and a visualization. An Answer may contain multiple visualizations and datasets.

The default view is a list of Answers that can be filtered by a wide range of modifiers such as Domain, Status, View Count, Recency, etc. In addition to these filter modifiers, you can also search directly via our advanced free text search.

Preview

Run SQL code to render results before publishing.

Samples

When exploring data it is helpful to see it in context. “Samples” does just this, allowing you to ‘sample’ the first 100 rows (this is adjustable to 200, 500, and 1000) of a table, Datamap, or View. These results are searchable.

To improve the readability of sparse data sets, a filter to remove NULL values is available.

Reasoner

Some of the largest barriers in knowledge extraction come about because only a few users are intimately familiar with their data landscape and the various database schemas therein.

The Promethium Reasoner solves this complexity by leveraging AI and ML to analyze Datamaps, tables, and views and make suggestions of what objects can be logically connected. For example, when the Reasoner is joining two tables, it can suggest rows to join upon.



The selection candidates are displayed in the number of drop-down boxes on the Reasoner UI and users are required to make a choice. To help users make an informed decision, we display additional information and sort candidates according to their probability of a match.

Tag

Tags are searchable terms associated with various components of Promethium. They elevate the NLP searching capabilities and the useability of the Datamaps. Tags can be applied at the column level, table level, Answers, and Datamaps.

The rules for a tag are:

- Single word (no spaces)
- "-" and "_" allowed
- Numbers allowed
- No special characters

Domains

Domains are a powerful labeling and organizational concept within Promethium that integrates with Role Based Access Control. You can add a domain to any object to better group, filter, or convey access rights.

Trino

Trino is our distributed SQL query engine. It is deployed as a cluster and utilizes file storage and metastore for persistence. Trino is a distributed system that utilizes an architecture similar to massively parallel processing (MPP) databases. Like many other big data engines, there is a coordinator node that then manages multiple worker nodes to process all the work that needs to be done in a K8S cluster.

User Accounts & RBAC

This section describes the types of user accounts that exist within Promethium, their usage, and the default permissions sets that are created with a new installation.

Promethium operates a fine-grained, role-based, access control mechanism.

Users and roles are system-wide objects; they do not belong to specific data sources. Users are independent accounts with login access, while roles define groups of users that accrue privileges through role membership.

Every role inherits one or more sets of permissions. A permission set can include any combination of permissions across every object within the system.



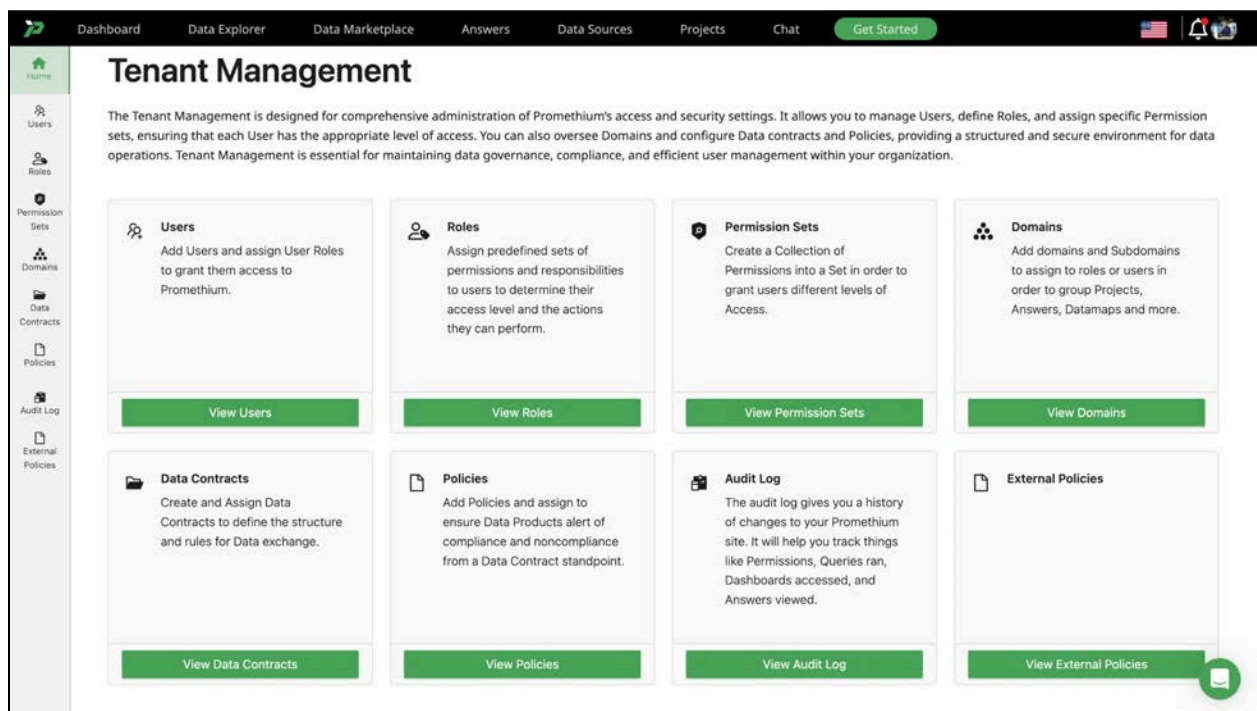
After creating roles and users, you can grant or revoke permissions with immediate effect. You can grant access privileges on various objects to one or more users or roles.

Users

Users can be created, edited,, and deleted by users with the appropriate rights. By default, this permission is associated only with the Tenant Admin role.

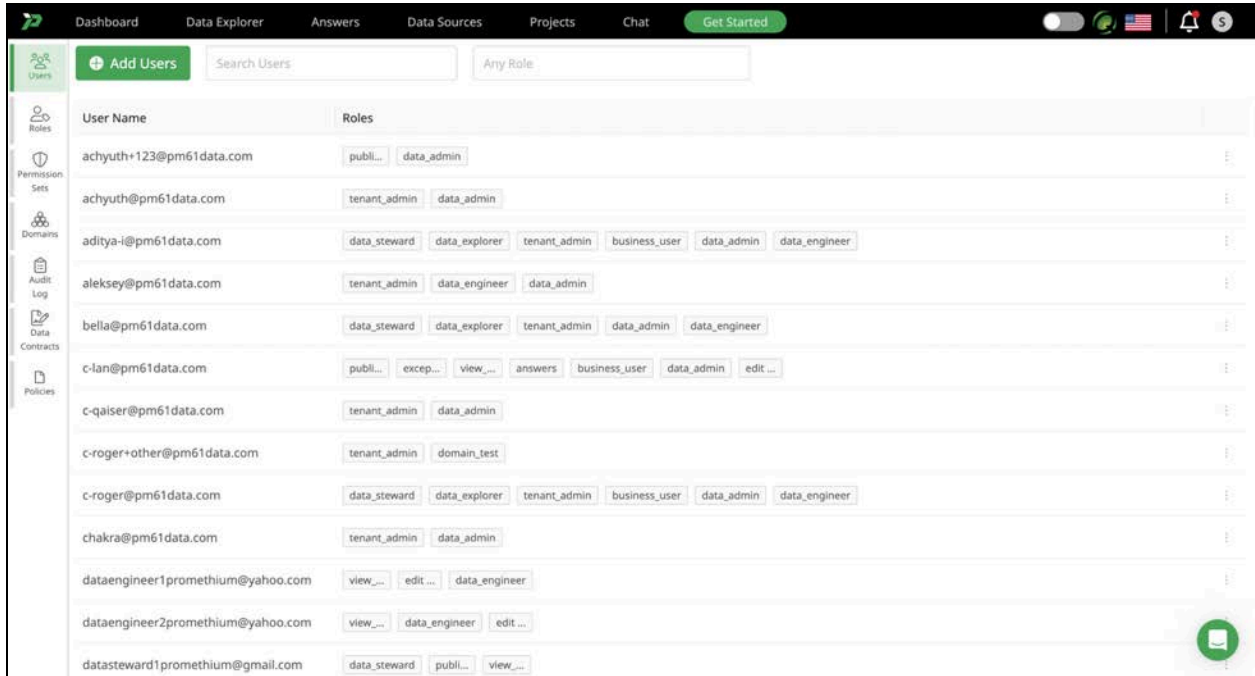
Tenant Management

The Tenant Management is designed for comprehensive administration of Promethium’s access and security settings. It allows you to manage Users, define Roles, and assign specific Permission sets, ensuring that each User has the appropriate level of access. You can also oversee Domains and configure Data contracts and Policies, providing a structured and secure environment for data operations. Tenant Management is essential for maintaining data governance, compliance, and efficient user management within your organization.



Navigate to User Management

Click on your initial or avatar on the top navigation bar. Select Tenant Management from the menu. Searching for a User.



All currently registered users within the system will be displayed. You can search by username or role to filter a large list quickly. Searching by name will filter the list immediately with each letter entered. Searching by role will allow you to limit the list to users who have that role

Creating a User

From the User Management screen click **+ Add Users**.

Two different types of users can be created, “Regular User” and “Service User”

For the “Regular User”, begin by creating a username, followed by selecting at least one role for the user in the dropdown, and finally create a password for the user account.



Dashboard Data Explorer Answers Data Sources Projects Chat

Add Users

User Type: Regular User Service User

SSO Only:

Email:

Separate with commas, semicolons, or new lines.

* Roles:

Send Welcome Email:

Dashboard Data Explorer Answers Data Sources Projects Chat

Add Users

User Type: Regular User Service User

User Name:

* Roles:

* Password :

* Verify Password :

For the “Service User”

It is often helpful to alert a new user to the availability of their account and this option is enabled by default. If you wish to not send an email, uncheck the box **Send Welcome Email**. Click Submit to create the account and return to the main user's view.



Hi, your account is ready.

Hi,

Your account is ready to go! Getting started is fast and easy. Click the following link to login with you email and temporary password.

[Activate Account](#)

Link expires in 7 days

Your username: bella@pm61data.com

Your temporary password: hF5tz2pJwz

Your organization's sign in page is <https://qa-gcp-trino.qa.pm61data.io>

You will also receive a welcome email with tips to help you get started, including how to contact support.

Regards,

The Promethium Team

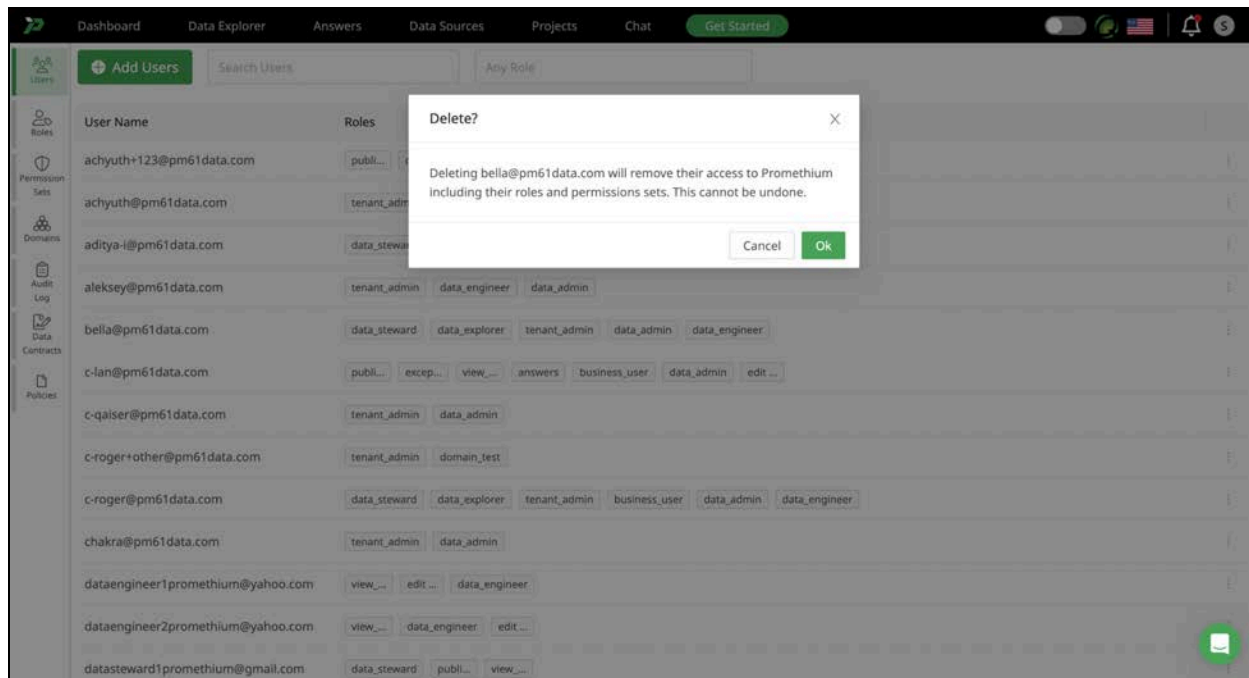
promethium.ai

You received this email because someone added your information into Promethium. If you have any questions or require assistance, [contact us](#).

Deleting a User

Users can be deleted by users with the appropriate rights. By default, this permission is only associated with the Tenant Admin role.

Select Users, select the user in question, and click on the three vertical dots to the right of the table. Select Delete.



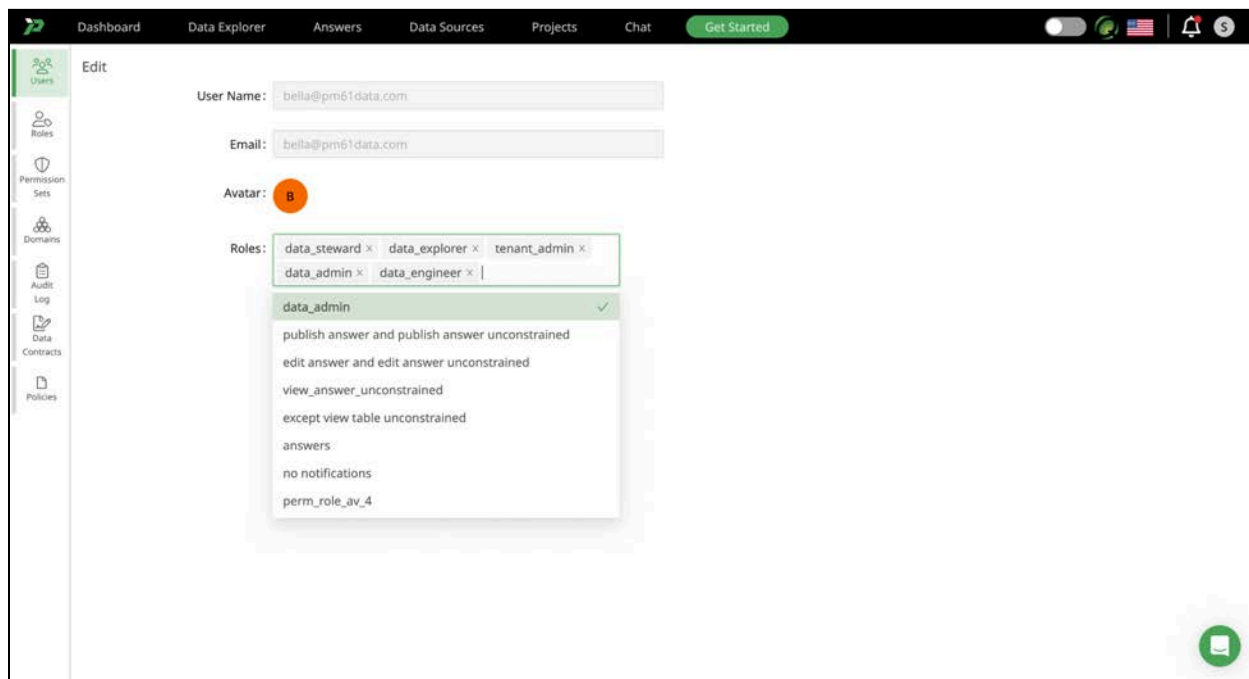
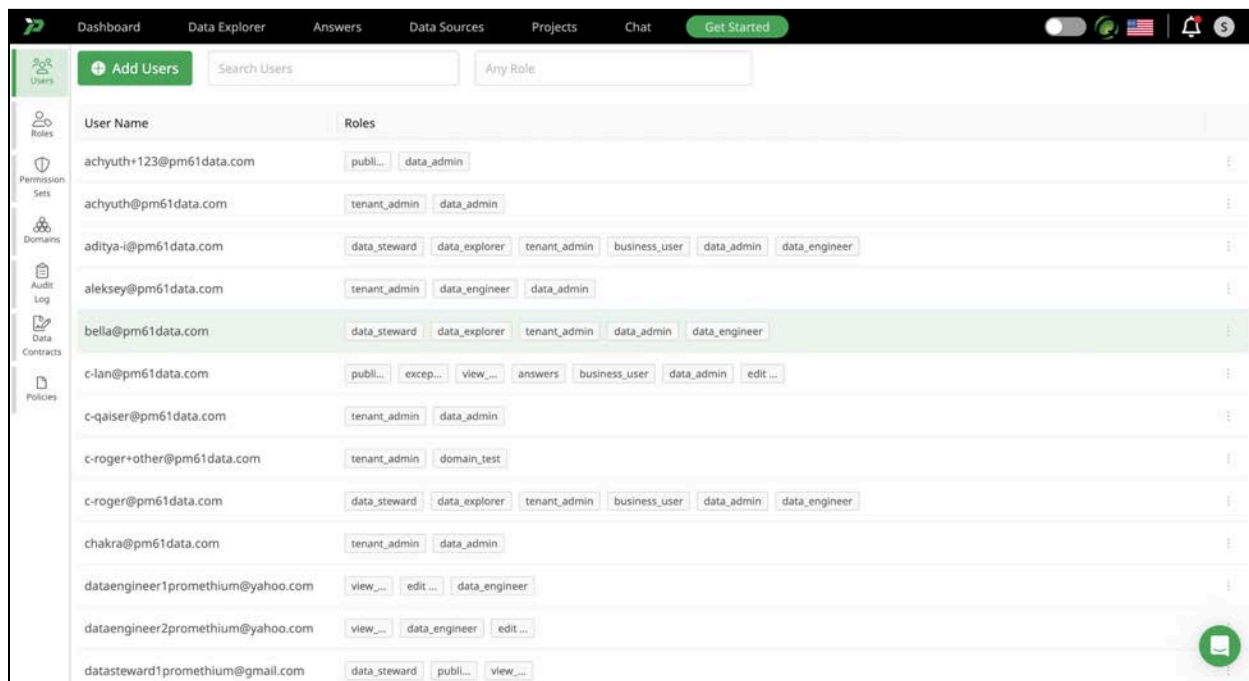
Roles

6 pre-configured roles ship with every installation.

Role	Purpose
Business User	This user needs to ask questions, view answers, and experiment with data.
Data Engineer	This user is responsible for creating reports and dashboards to help business users answer questions they can't self-serve.
Data Steward	This user is responsible for ensuring data access complies with corporate policies and other regulations. This user will often act as an analytics project manager to ensure requests and responses are passed appropriately between users and roles.
Data Explorer	This user has basic access rights - access to public, published answers and dashboards with no create or edit permissions.
Tenant Admin	Full control over the creation, update, and deletion of tenant accounts
Data Admin	Full control over the creation, updating, and deletion of data and data sources. Note that removing data sources does not delete the underlying database. This role is restricted from tenant modifications.



Multiple roles can be assigned to a single user. Select Users, select the user in question, and add or remove roles within the Roles: field.

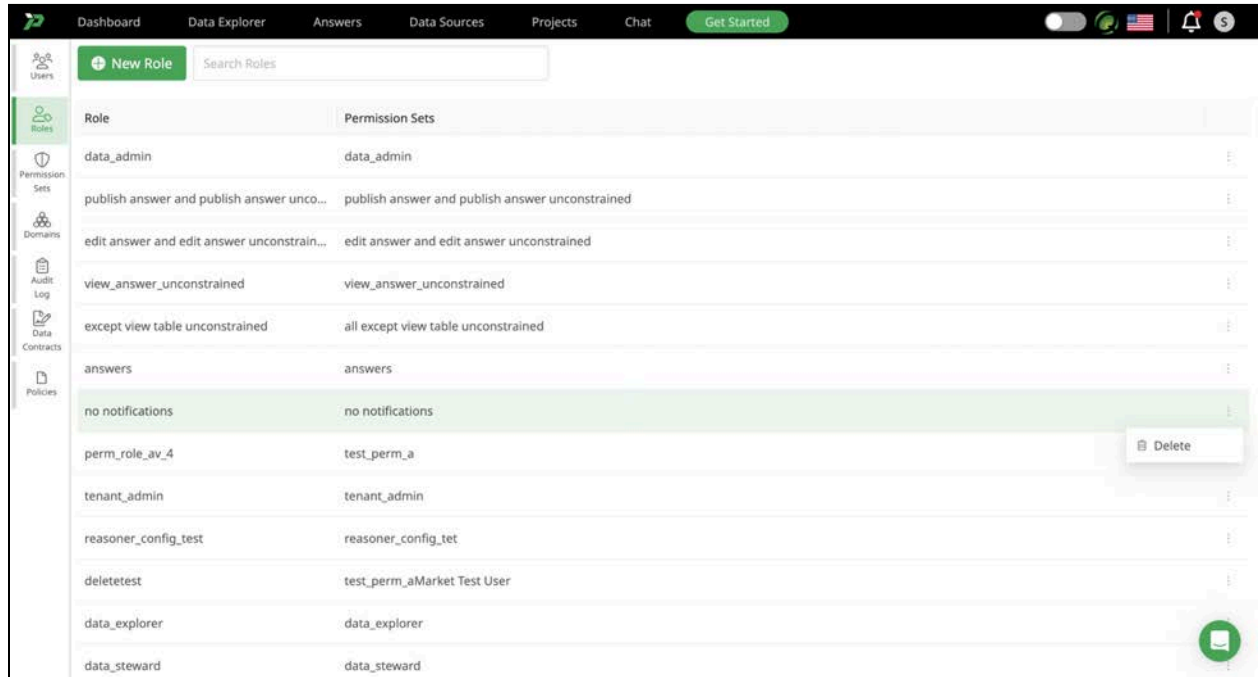


Deleting a role

Roles can be deleted by users with the appropriate rights. By default, this permission is associated with the Data Steward and Tenant Admin roles.



Select Roles, select the role in question, and click on the three vertical dots to the right of the table. Select Delete.



Domains

To navigate to Domain Management click on your initial or avatar on the top navigation bar. Select Tenant Management from the menu. Select Domain from the left navigation menu.

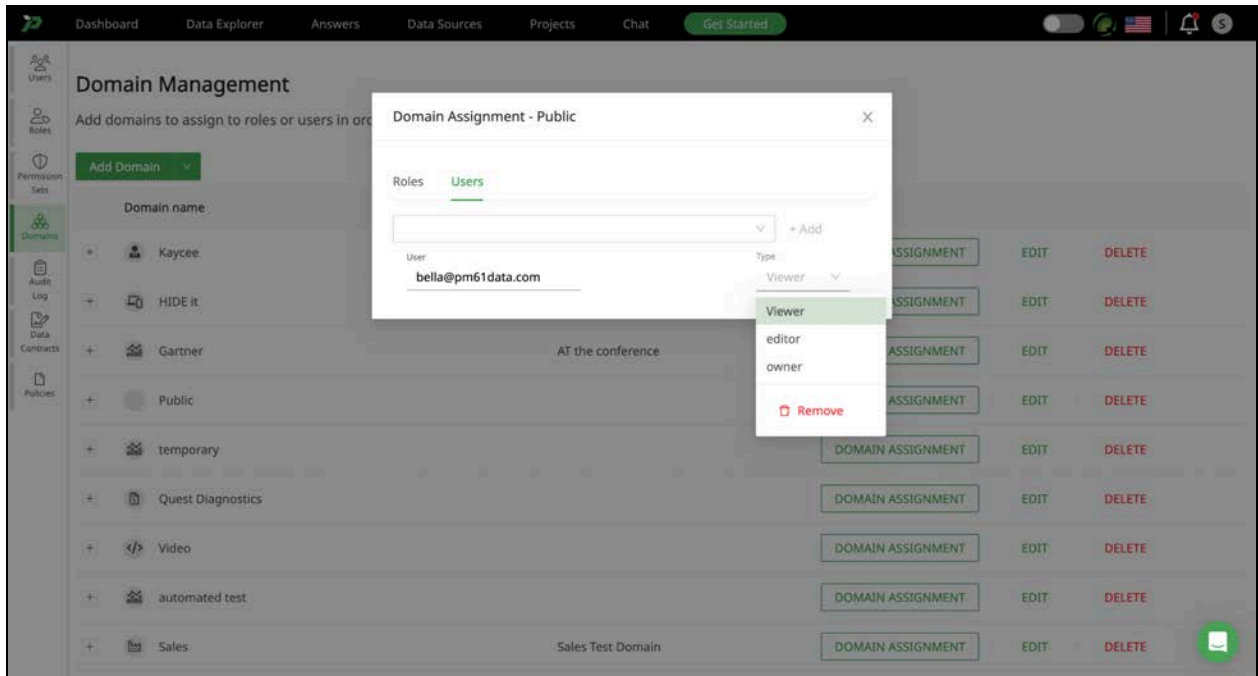
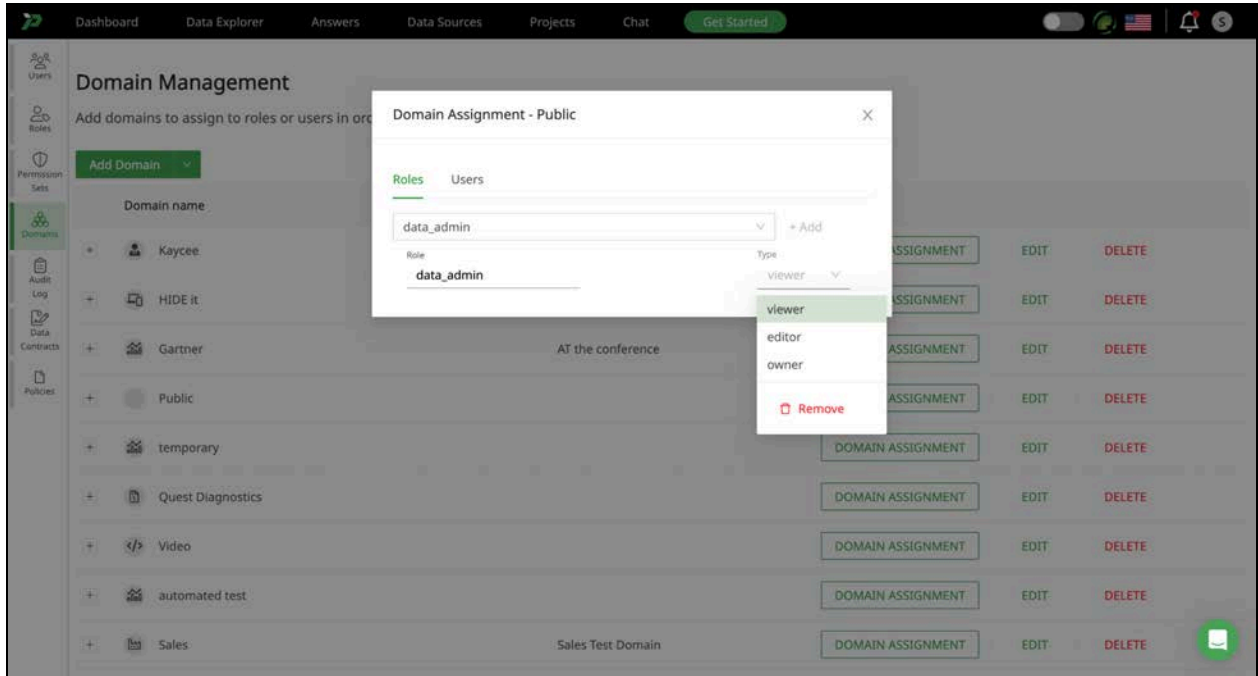


Domain name	Description	Action
+ Kaycee	Don't Mess	DOMAIN ASSIGNMENT EDIT DELETE
+ HIDE it		DOMAIN ASSIGNMENT EDIT DELETE
+ Gartner	AT the conference	DOMAIN ASSIGNMENT EDIT DELETE
+ Public		DOMAIN ASSIGNMENT EDIT DELETE
+ temporary		DOMAIN ASSIGNMENT EDIT DELETE
+ Quest Diagnostics		DOMAIN ASSIGNMENT EDIT DELETE
+ Video		DOMAIN ASSIGNMENT EDIT DELETE
+ automated test		DOMAIN ASSIGNMENT EDIT DELETE
+ Sales	Sales Test Domain	DOMAIN ASSIGNMENT EDIT DELETE

Domains & RBAC

Domains can be assigned to one or more Roles and/or Users to grant viewer, editor, or owner rights. Any object that is associated with that Domain will inherit those user or role restrictions.

To assign a Domain to a specific, pre-defined, role or user click **Domain Assignment** to render the assignment modal. Select either Role or User as appropriate and select the pre-defined role or user from the dropdown menu. Click **+Add**. From the type drop down menu select viewer, editor, or owner as required. Note: you can remove an assignment by clicking remove from the type dropdown next to the role or user you wish to unassign.



Domain Creation

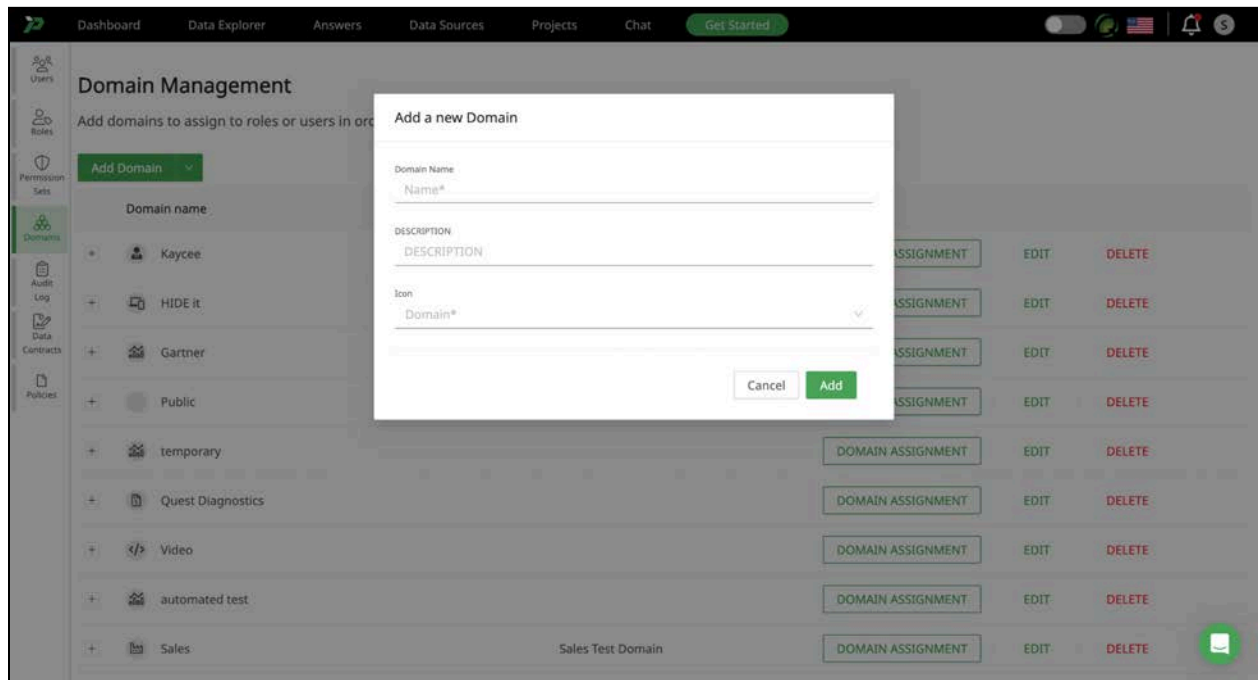
Domains can be created, edited, and deleted by users with the appropriate rights. By default, this permission is associated only with the Data Steward and Data Admin roles.

A user with the appropriate rights can create any number of Domains and Subdomains.



To Create a Domain click the green Add Domain button from the Domain Management page. To create a Subdomain click the arrow next to Add Domain and select Add Subdomain.

You must provide a domain name and select an icon. If you are creating a subdomain you must also select a pre-existing domain as the parent domain. The description field is optional but recommended so that other users can understand the purpose of your Domain at a later date. To create your domain click the green Add button.



Domain Assignment

You can assign any object to a domain or subdomain - e.g. data source, answers, or datamaps.

Assigning an object to a domain occurs from the object page. e.g. select a datamap or table from the Data Explorer flow. From the description modal you can assign or modify the current assignment directly.



The screenshot shows a web application interface with a dark top navigation bar containing 'Dashboard', 'Data Explorer', 'Answers', 'Data Sources', 'Projects', 'Chat', and a 'Get Started' button. On the left, a 'Filters' sidebar includes 'My Items', 'Favorites', 'PII', and several expandable categories like 'Duplicates', 'Domains (24)', 'Data Sources (18)', 'Endorsements (3)', 'Data Product Owner', and 'Subject Matter Expert'. The main area displays a table of items with columns for ID, status, data source, and name. The 'agents' item is highlighted. On the right, a 'DETAILS' panel for 'agents' shows settings for 'Private', 'Public', 'Sales', and 'temporary' categories, along with a 'Subdomain' dropdown set to 'Private' and a 'Tags' field. A vertical toolbar on the far right contains icons for 'Usages', 'Comments', 'Insights', 'Endorsements', 'Lineage', 'Relationships', and 'Data Contracts'. At the bottom of the details panel are 'Close', 'Analyze', and 'New Datamap' buttons.

Domain Searching & Filtering

When searching or browsing you can use filtering to limit the result set to that Domain or set of Domains.

Select the Domains filter and click the Domains and subdomains you want to restrict your searches to. Your Domain filter will remain in effect until you modify directly or select Clear all filters.



The screenshot shows a dashboard with a navigation bar at the top containing 'Dashboard', 'Data Explorer', 'Answers', 'Data Sources', 'Projects', and 'Chat'. A 'Get Started' button is also present. On the left, there is a 'Filters' sidebar with options for 'My Items', 'Favorites', 'PII', and 'Duplicates'. The 'Domains (24)' filter is selected, showing a search bar and a list of domain names such as 'automated test', 'Boomi', 'business', 'Demo', 'sub-domain manage...', 'domain1-viewer', 'domain2-editor', 'domain3-owner', 'domain4-nothing', 'Gartner', 'engineering', and 'qa'. The main area displays a table of data sources with columns for ID, source type, name, table name, status, and time. The table lists various data sources like 'orders', 'customer_personal_record', 'opportunities', 'raw_orders', 'keywords', 'master_plan_wdates', 'test_null_hiding', 'stg_customers', 'sirius_xm_dataset_1692123904890...', 'toy_products', and 'agents'. A pagination bar at the bottom shows page 1 of 89.

The screenshot shows the 'Data Contracts' management page. The navigation bar is the same as in the previous screenshot. On the left, a sidebar contains icons for 'Users', 'Roles', 'Permission Sets', 'Domains', 'Audit Log', 'Data Contracts', and 'Policies'. The 'Data Contracts' section is active. At the top right, there are buttons for 'Upload JSON' and 'Create Data Contract'. The main content area features a table with columns for 'Column Name', 'Description', and 'Actions'. The table lists several data contracts, including 'Contract - Answer fulfillment 3 hours', 'data_contract_for_bq', 'Financial Reporting', 'test contracts', 'Contract with all policies', 'test contract all policies qs', and 'Financial Analysis for Marketing'. Each row includes 'Edit' and 'Delete' action buttons.

Data Contracts

The Data Contract feature offers a fully customizable framework to set precise guidelines for your data management needs. Customers can tailor the refresh frequency, execution time, and answer fulfillment time according to their specific requirements. This customizable documentation serves to eliminate uncertainties and prevent undocumented assumptions about data, providing a clear and consistent understanding of data expectations.



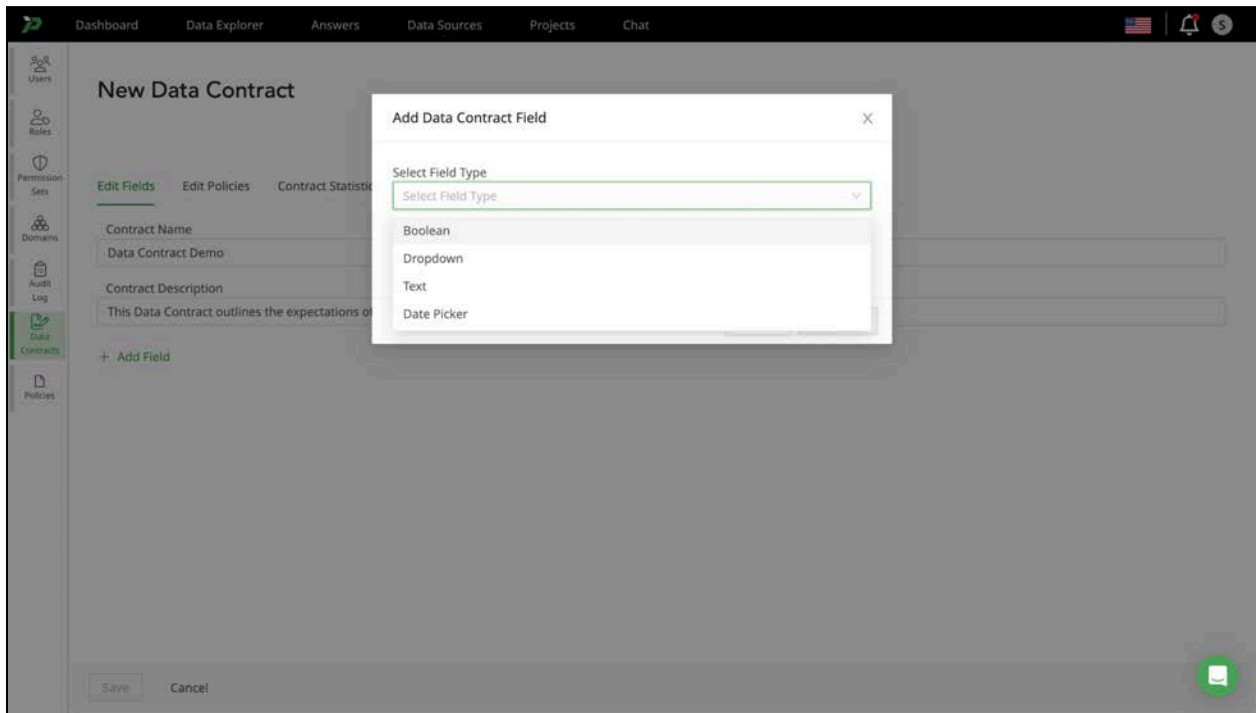
Creating a Data Contract

Navigate to the **Create Data Contract** button on the right side of the Data Contracts tab. In the New Data Contract page, the first initial fields you will see are Contract Name and Contract Description.

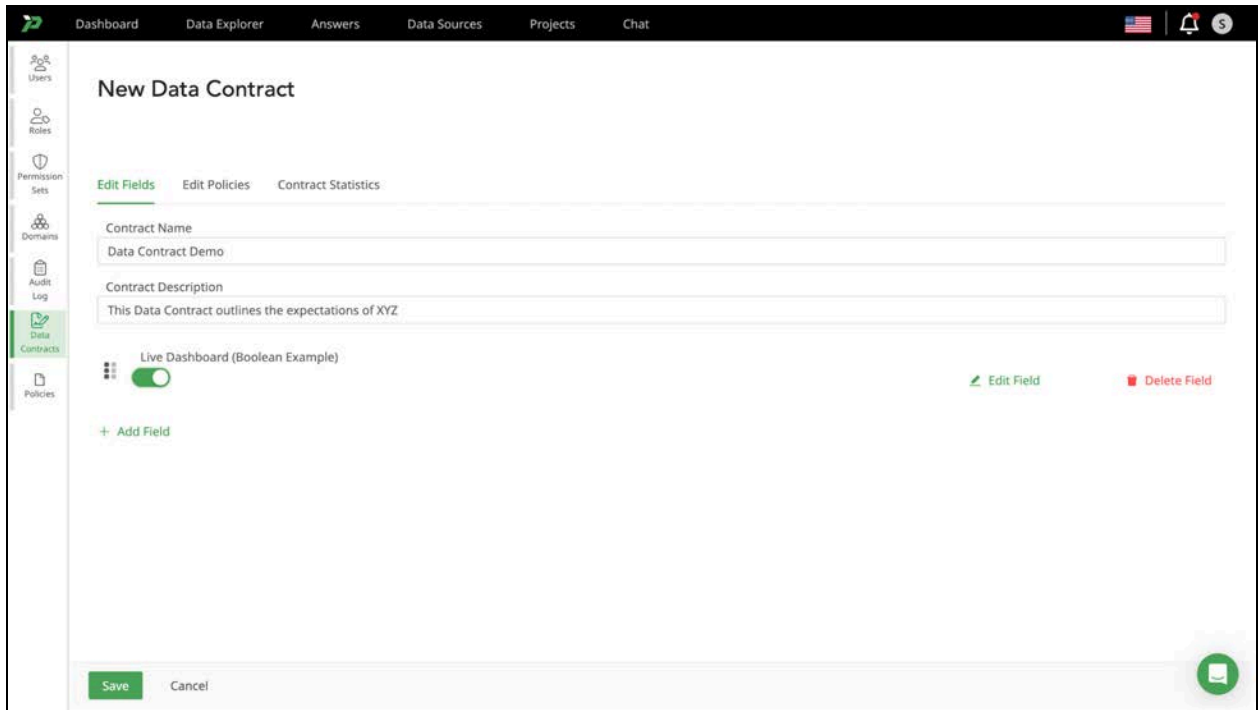
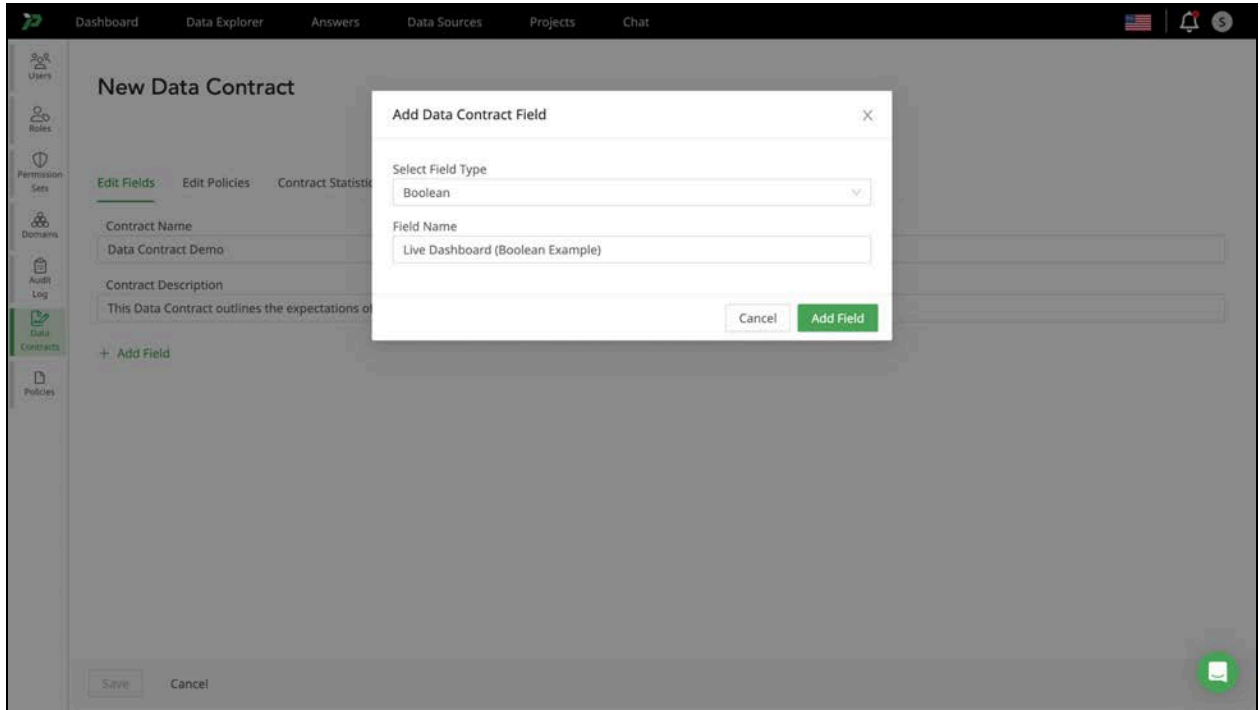
The screenshot shows a web application interface for creating a new data contract. The top navigation bar includes 'Dashboard', 'Data Explorer', 'Answers', 'Data Sources', 'Projects', and 'Chat'. A sidebar on the left contains icons for 'Users', 'Roles', 'Permission Sets', 'Domains', 'Audit Log', 'Data Contracts' (highlighted), and 'Policies'. The main content area is titled 'New Data Contract' and features three tabs: 'Edit Fields' (selected), 'Edit Policies', and 'Contract Statistics'. The 'Edit Fields' tab contains two text input fields: 'Contract Name' with the value 'Data Contract Demo' and 'Contract Description' with the value 'This Data Contract outlines the expectations of XYZ'. Below these fields is a '+ Add Field' button. At the bottom of the form are 'Save' and 'Cancel' buttons, and a green circular icon with a white document symbol.



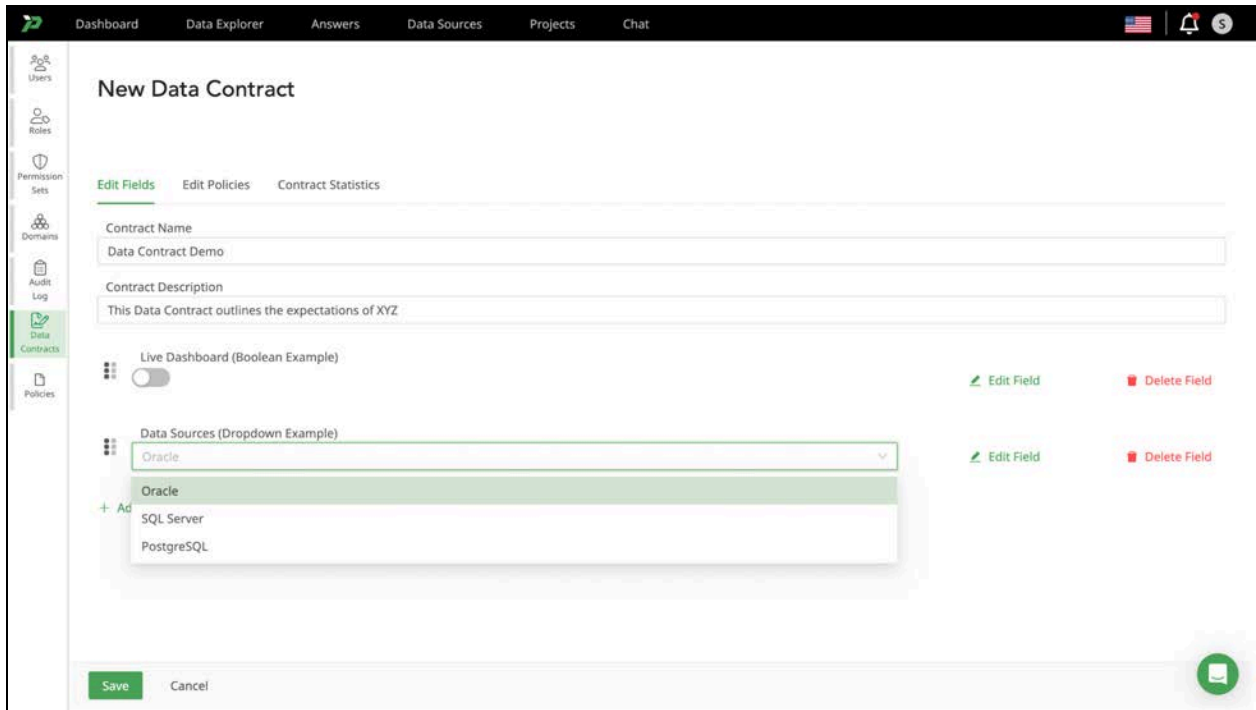
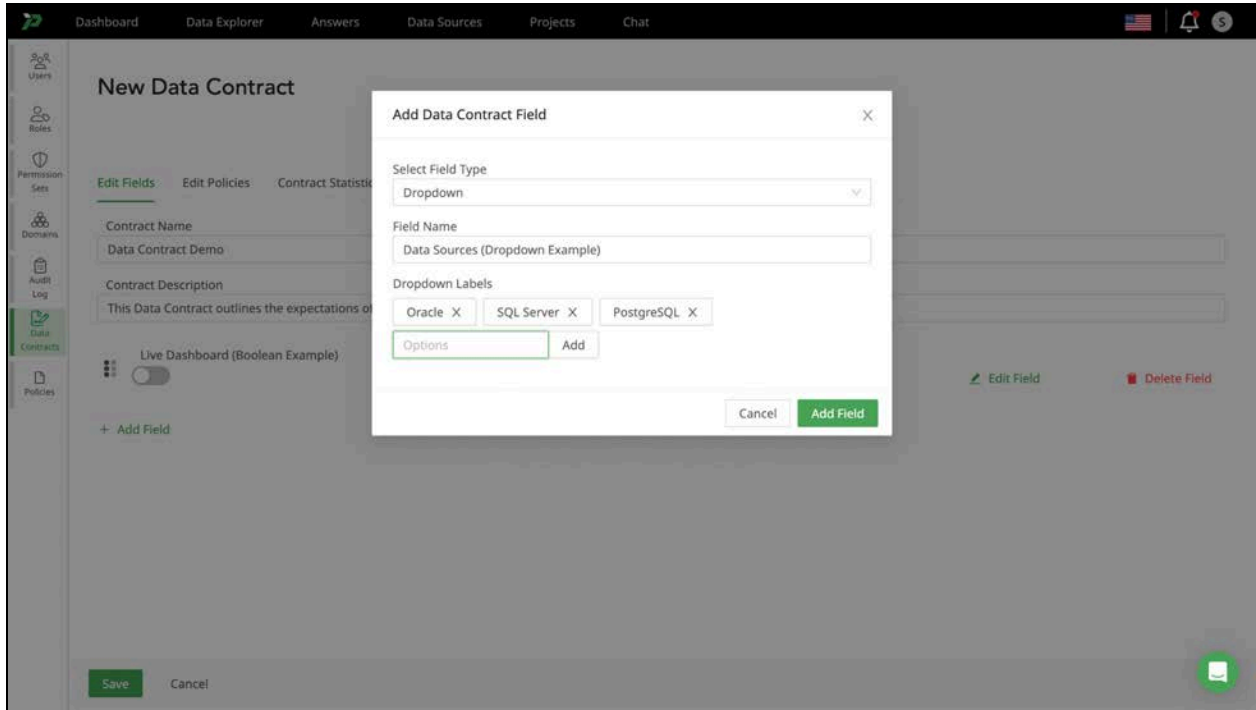
Click Add Field to add additional fields, a dialog box will appear with Select Field Type and Field Name.



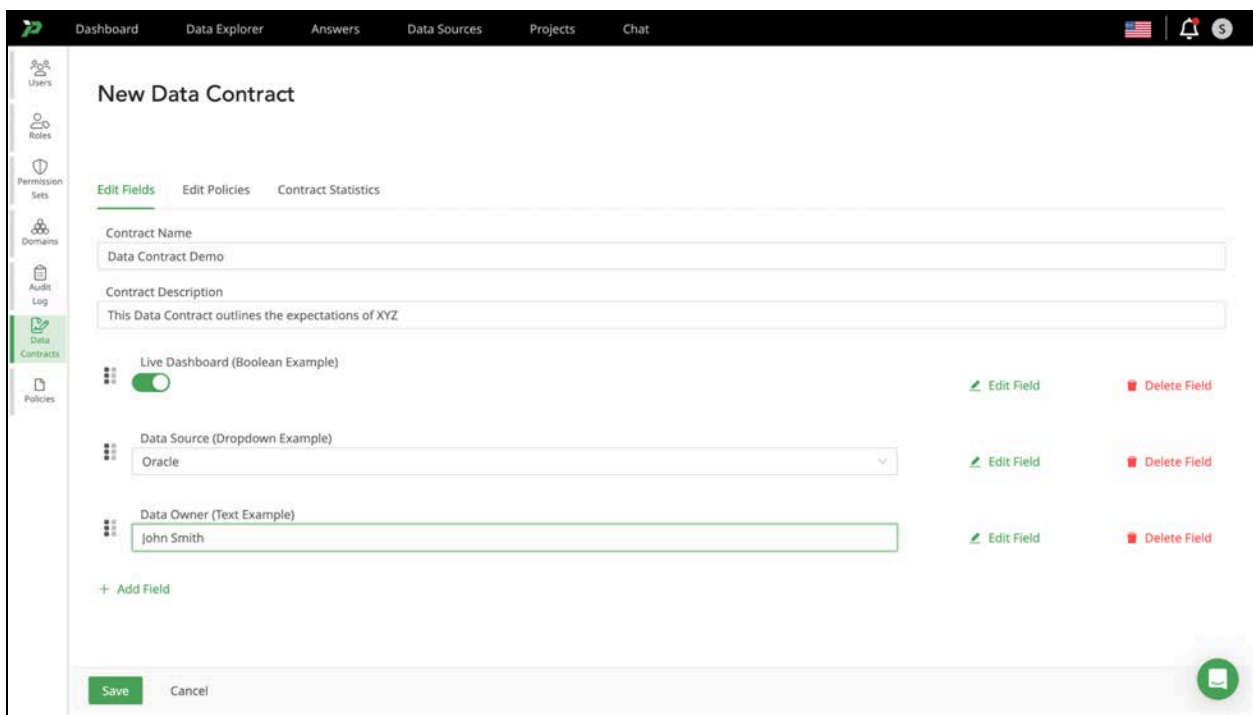
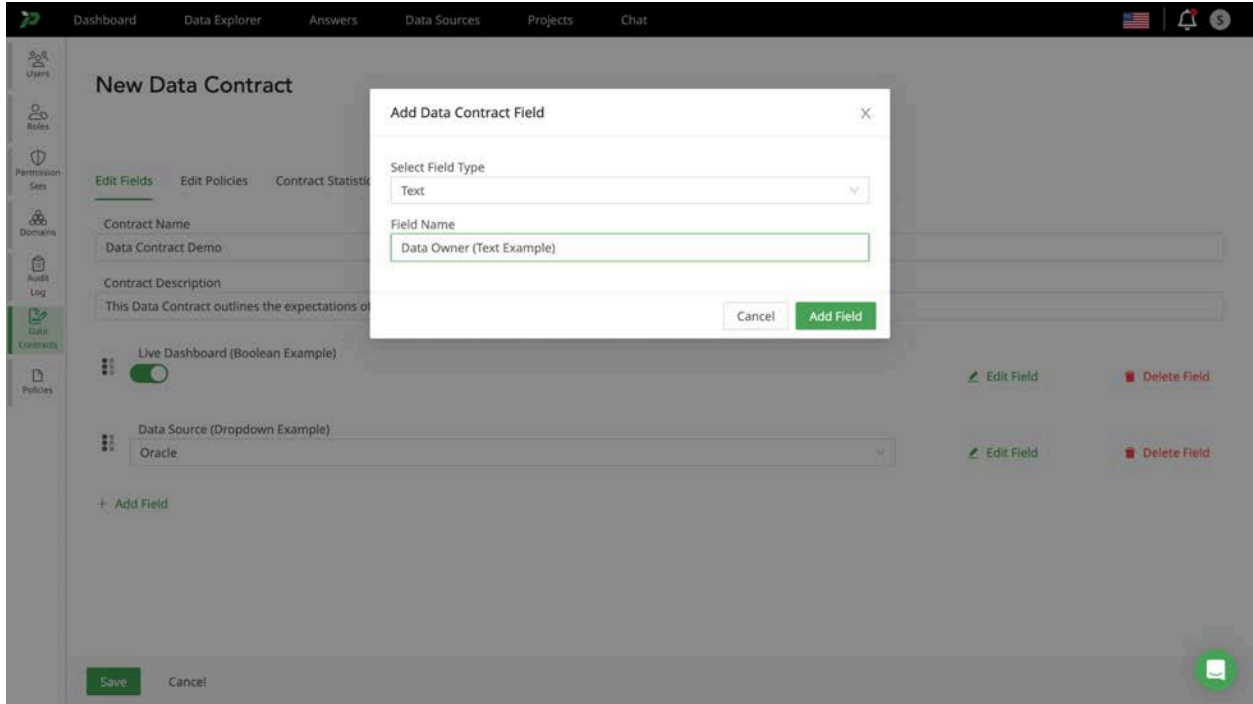
- Boolean - Name the field and then Add Field. A toggle switch will appear.



- Dropdown - Add all the options to be included in the dropdown, then Add Field. From there, you can select the appropriate option.



- Text - Name the field and then Add Field. A text box will be provided for you to enter any relevant information or comments as needed.



- Date Picker - Name the field and then Add Field. You will be able to select a date off a calendar view and specify a time to the second, time is in a 24-hour format.



Dashboard Data Explorer Answers Data Sources Projects Chat

New Data Contract

Add Data Contract Field

Select Field Type
Date Picker

Field Name
Date Effective

Cancel Add Field

Contract Name
Data Contract Demo

Contract Description
This Data Contract outlines the expectations of

Live Dashboard (Boolean Example)

Data Source (Dropdown Example)
Oracle

Data Owner (Text Example)
John Smith

+ Add Field

Save Cancel

Dashboard Data Explorer Answers Data Sources Projects Chat

New Data Contract

Edit Fields Edit Policies Contract Statistics

Contract Name
Data Contract Demo

Contract Description
This Data Contract outlines the expectations of XYZ

Jan 2020							16:00:00		
Su	Mo	Tu	We	Th	Fr	Sa	16	00	00
29	30	31	1	2	3	4	17	01	01
5	6	7	8	9	10	11	18	02	02
12	13	14	15	16	17	18	19	03	03
19	20	21	22	23	24	25	20	04	04
26	27	28	29	30	31	1	21	05	05
2	3	4	5	6	7	8	22	06	06
							23	07	07

Now

2020-01-01 16:00:00

Ok

+ Add Field

Save Cancel



Once you are satisfied with the fields added, click Save. Fields can be conveniently edited or removed by selecting the links on the right, which will prompt a dialog box for the necessary changes.

Applying the Data Contract

With your newly created Data Contract, you can now apply it to an Answer page and/or a Datamap. Go to your selected Answer page or Datamap, and then switch to the Data Contract tab. Choose your preferred option from the dropdown menu, click 'Save,' and the Data Contract



will be set in place with all related policies actively enforced.

The screenshot shows the 'Answer Demo' configuration page. The top navigation bar includes 'Dashboard', 'Data Explorer', 'Answers', 'Data Sources', 'Projects', and 'Chat'. The page title is 'Answer Demo' with two tabs: 'answer x' and 'demo x'. On the right, there are fields for 'Status' (In Progress), 'Initial Requester' (support@pm61...), 'Data Product Owner' (Not Selected), and 'Subject Matter Expert' (Not Selected). The main content area is titled 'Data Contract' and contains several fields: 'Select Data Contract' (Data Contract Demo), 'Live Dashboard (Boolean Example)' (checked), 'Data Source (Dropdown Example)' (Oracle), 'Data Owner (Text Example)' (John Smith), and 'Date Effective' (2020-01-01 16:00:00). A 'Save' button is in the top right. A 'Policy Enforcement History' link is at the bottom left.

The screenshot shows the 'data_contract_datamap_demo' configuration page. The top navigation bar is the same as the previous screenshot. The page title is 'data_contract_datamap_demo'. The main content area is titled 'Data Contract' and contains the same fields as the previous screenshot: 'Select Data Contract' (Data Contract Demo), 'Live Dashboard (Boolean Example)' (checked), 'Data Source (Dropdown Example)' (Oracle), 'Data Owner (Text Example)' (John Smith), and 'Date Effective' (2020-01-01 16:00:00). A 'Save' button is in the top right. A 'Policy Enforcement History' link is at the bottom left.

Audit Log

The Audit Log provides a historical record of modifications made to your Promethium site, allowing you to monitor activities such as changes in permissions, executed queries, accessed dashboards, and viewed answers.



The screenshot displays the 'Audit Log' section of a web application. The top navigation bar includes 'Dashboard', 'Data Explorer', 'Data Marketplace', 'Answers', 'Data Sources', 'Projects', and 'Chat'. The left sidebar contains navigation icons for 'Users', 'Roles', 'Permission Sets', 'Domains', 'Data Contracts', 'Policies', 'Audit Log', and 'External Policies'. The main content area is titled 'Audit Log' and includes a search bar, 'Filters', 'Export', and 'Settings' buttons. Below this is a table with the following data:

Affected Item	Date	Item Type	Event Type	User
User [All]	At 03/22/2024 9:05:24	User	Viewed	support@pm61data.com
cytest	At 03/22/2024 9:04:52	Domain	Viewed	support@pm61data.com
chakra_datamap_1710864047882_mq	At 03/22/2024 9:04:51	Catalog	Viewed	support@pm61data.com
cytest	At 03/22/2024 9:04:51	Domain	Viewed	support@pm61data.com
Catalog	At 03/22/2024 9:04:50	Catalog	Searched	support@pm61data.com

The expanded view for the 'chakra_datamap_1710864047882_mq' event shows the following details:

- Affected Item: [chakra_datamap_1710864047882_mq](#)
- Date: 03-22-2024 9:04:51
- Item Type: Catalog
- Event Type: Viewed
- Details: assigned to domains: cytest, description (chakra_datamap_1710864047882_mqt)
- User: support@pm61data.com

The Audit Log will link items that have been affected where applicable.

- **Affected Item:** Provides information on a specific item that was interacted with and will show up as a link, i.e. title of datamap or answer, unless the full catalog of a page is open, it will be signified with [All].
- **Date:** The logged time and date the interaction occurred.
- **Item Type:** Item associated with the interaction - Project, System, Cluster, Answer, Data_explorer, Datamap, Dbt model, View, Table, Data_Source, Chat, 3rd party application, Job management, Tenant management, Data marketplace, Data product, System, User.
- **Event Type:** Action associated with interaction - Schedule, Unschedule, View, Run, Update, Delete, Create, Search, Login, Start, Publish, Unpublish, Stop, Contextualize.
- **User:** The user associated with logged action.
- **+/-** : Expand or Collapse log for additional information and important details



The screenshot displays the Promethium Audit Log interface. The main content area shows a table of audit events with the following columns: Affected Item, Date, Item Type, and Event Type. Below the table, a detailed view of an event is shown, including the affected item name, date, item type, event type, and user information. A 'Filters' sidebar on the right allows filtering by Date, User, Item Type, and Event Type.

Affected Item	Date	Item Type	Event Type
User [All]	At 03/22/2024 9:05:24	User	Viewed
cytest	At 03/22/2024 9:04:52	Domain	Viewed
chakra_datamap_1710864047882_mq	At 03/22/2024 9:04:51	Catalog	Viewed
cytest	At 03/22/2024 9:04:51	Domain	Viewed
Catalog	At 03/22/2024 9:04:50	Catalog	Searched

The Filters button allows you to refine the results in the Audit Log based on when the event was logged, the user, item type, and action type.

Ranger

Apache Ranger is a well established enterprise data level security tool. Promethium includes Ranger, pre-configured with Trino, within our Intelligent Edge (IE) to ensure data security.

This section describes the architecture and how to create masking and row-level filter rules in Ranger.

Promethium includes several core Ranger components -

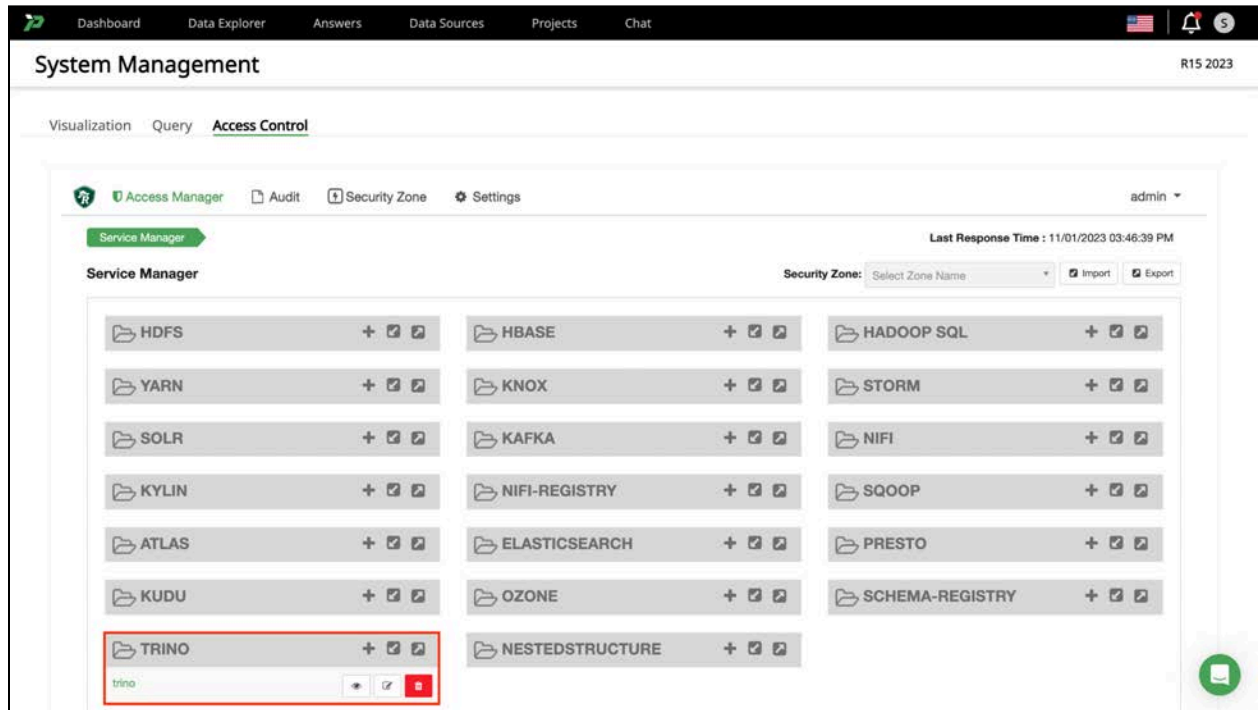
Ranger Admin - This is the Ranger agent. It manages all configurations, provides a UI, and allows the management of policies and permissions. This UI will be made available, embedded, directly within the System Management area of Promethium

Trino Ranger Plugin - The Trino Ranger plugin allows Trino to validate policies before executing queries.

Ranger PostgreSQL - Ranger requires a separate database for the storage of users and policies.



Creating a User



Click **System Management** and then select **Ranger**. This will render the Ranger Admin UI within Prometheus.

Click on **Trino** to enter the policy manager for Trino.



To apply a rule to a user that does not exist in Ranger, hover over the **Settings** button in the navigation bar and click on **Users/Groups/Roles**.

The screenshot shows the Ranger System Management interface. The top navigation bar includes Dashboard, Data Explorer, Answers, Data Sources, Projects, and Chat. The main header is "System Management" with the date "R15 2023". Below the header, there are tabs for Visualization, Query, and Access Control. The "Access Manager" section is active, showing a "Service Manager" view. A dropdown menu is open under "Settings", highlighting "Users/Groups/Roles". The main content area displays a grid of service icons such as HDFS, YARN, SOLR, KYLIN, ATLAS, KUDU, TRINO, HBASE, KNOX, KAFKA, NIFI-REGISTRY, ELASTICSEARCH, OZONE, NESTEDSTRUCTURE, HADOOP SQL, STORM, NIFI, SQOOP, PRESTO, and SCHEMA-REGISTRY. Each icon has a plus sign and a checkmark. The "Last Response Time" is 11/01/2023 03:46:39 PM. The URL at the bottom is https://qa-gcp-ranger.qa.pm61data.io/#/users/usertab.

Click **Add New User**

The screenshot shows the Ranger System Management interface with the "Users/Groups/Roles" section selected. The "Users" tab is active. The "User List" section features a search bar and buttons for "Add New User", "Set Visibility", and a trash icon. Below is a table listing users with columns for User Name, Email Address, Role, User Source, Sync Source, Groups, Visibility, and Sync Details.

	User Name	Email Address	Role	User Source	Sync Source	Groups	Visibility	Sync Details
<input type="checkbox"/>	admin		Admin	Internal	--	--	Visible	--
<input type="checkbox"/>	rangerusersync		Admin	Internal	--	--	Visible	--
<input type="checkbox"/>	rangertagsync		Admin	Internal	--	--	Visible	--
<input type="checkbox"/>	promethium		User	External	--	--	Visible	--
<input type="checkbox"/>	support@pm61data.com	support@pm61data.com	Admin	Internal	--	--	Visible	--
<input type="checkbox"/>	ravi@pm61data.com	ravi@pm61data.com	Admin	Internal	--	--	Visible	--
<input type="checkbox"/>	enrique-automation@pm61data.com	enrique@pm61data.com	Admin	Internal	--	--	Visible	--
<input type="checkbox"/>	kaycee@pm61data.com	kaycee@pm61data.com	Admin	Internal	--	--	Visible	--



Fill out the form. The username field must be the full email address of the user registered in Promethium. The password does not matter - Ranger and Promethium will match the user.

Click **Save**. You will now be able to apply masking and row-level rules to this user.

The screenshot shows the 'User Create' form in the Promethium interface. The form fields are:

- User Name *: bella@promethium.ai
- New Password *: [masked]
- Password Confirm *: [masked]
- First Name *: Bella
- Last Name *: Baek
- Email Address: bella@pm61data.com
- Select Role *: Admin
- Group: Please select (with a '+' button)

The interface includes a top navigation bar with 'Dashboard', 'Data Explorer', 'Answers', 'Data Sources', 'Projects', and 'Chat'. A sidebar on the left shows 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The top right shows 'admin' and a 'Last Response Time' of 11/01/2023 04:33:14 PM.

Creating a Masking Rule

Navigate to the **Masking** tab. If you don't see this tab, click the **Ranger** logo to move to the top level navigation, select **Trino** and you should see the correct tab options.



Click **Add New Policy**.

Policy ID	Policy Name	Policy Labels	Status	Audit Logging	Roles	Groups	Users	Action
1	all - trinouser	--	Enabled	Enabled	--	--	promethium (USER) support@pm61data.com admin	👁️ ✓ 🚫
2	all - catalog	--	Enabled	Enabled	--	--	promethium (USER) admin support@pm61data.com	👁️ ✓ 🚫
3	all - function	--	Enabled	Enabled	--	--	promethium (USER) admin support@pm61data.com	👁️ ✓ 🚫
4	all - catalog, sessionproperty	--	Enabled	Enabled	--	--	promethium (USER) admin support@pm61data.com	👁️ ✓ 🚫
5	all - catalog, schema, procedure	--	Enabled	Enabled	--	--	promethium (USER) admin support@pm61data.com	👁️ ✓ 🚫

Name the policy in a manner that accurately describes your intended policy rule.

In this example, we will be applying a masking rule to a datamap in Trino.

To apply the masking rule to a datamap in Trino, set the Catalog as **hive** and the schema as **promethium**.

We'll be applying a masking rule on a datamap **nick_ranger_demo** and column **plan_name**.



Ignore any errors that pop up on the right of the screen claiming that “Resource lookup failed”. This does not impact the application masking rule.

The screenshot shows a configuration form for a policy. At the top, there are navigation tabs: Dashboard, Data Explorer, Answers, Data Sources, Projects, and Chat. The policy is of type "Access" with ID "15". The Policy Name is "mask_plan_name" and is currently "Enabled". There are fields for Policy Label and Description. Audit Logging is set to "Yes". A green button "Add Validity Period" is in the top right. The Resources section contains three entries: "catalog" with "hive", "schema" with "promethium", and "table" with "nick_ranger_demo", each with an "Include" toggle.

This screenshot shows the same form with an additional resource added: "column" with "plan_name". Below the resources is a green "+ Add Resource" button. The "Allow Conditions" section features a table with columns for "Select Role", "Select Group", "Select User", "Permissions", and "Delegate Admin".

Select Role	Select Group	Select User	Permissions	Delegate Admin
Select Roles	Select Groups	nick_bella@promethium.ai	Alter	<input checked="" type="checkbox"/>

At the bottom of the form are **Mask Conditions**. Here you can define which users and rules you want to apply to the specified datamap.



Click in the **Select User** field and either type or scroll to find one of the existing users.

Click the **+** button in **Permissions** and choose **Alter**.

Click the **+** button in **Delegate Admin** and select the option you want to apply. For this documentation, we are going to check the box.

When you are satisfied with the rules click **Add** to save the rule at the very bottom of the page. The rule will instantly take effect. You will be returned to the full list of masking rules and will see the policy you created at the bottom of the table.

ID	Name	Status	Users
2	all - catalog	Enabled	promethium [USER] admin support@pm61data.com
3	all - function	Enabled	promethium [USER] admin support@pm61data.com
4	all - catalog, sessionproperty	Enabled	promethium [USER] admin support@pm61data.com
5	all - catalog, schema, procedure	Enabled	promethium [USER] admin support@pm61data.com
6	all - catalog, schema, table	Enabled	promethium [USER] admin support@pm61data.com
7	all - systemproperty	Enabled	promethium [USER] admin support@pm61data.com
8	all - catalog, schema, table, column	Enabled	promethium [USER] admin support@pm61data.com
9	all - catalog, schema	Enabled	promethium [USER] admin support@pm61data.com
10	ranger-user-catalog-policy	Enabled	promethium [USER] admin support@pm61data.com
15	mask_plan_name	Enabled	bella@promethium.ai

Licensed under the Apache License, Version 2.0



ID	Name	Status	Status	User	Actions
2	all - catalog	Enabled	Enabled	promethium (USER) admin support@pm61data.com	👁️ 🗑️
3	all - function	Enabled	Enabled	promethium (USER) admin support@pm61data.com	👁️ 🗑️
4	all - catalog, sessionproperty	Enabled	Enabled	promethium (USER) admin support@pm61data.com	👁️ 🗑️
5	all - catalog, schema, procedure	Enabled	Enabled	promethium (USER) admin support@pm61data.com	👁️ 🗑️
6	all - catalog, schema, table	Enabled	Enabled	promethium (USER) admin support@pm61data.com	👁️ 🗑️
7	all - systemproperty	Enabled	Enabled	promethium (USER) admin support@pm61data.com	👁️ 🗑️
8	all - catalog, schema, table, column	Enabled	Enabled	promethium (USER) admin support@pm61data.com	👁️ 🗑️
9	all - catalog, schema	Enabled	Enabled	promethium (USER) admin support@pm61data.com	👁️ 🗑️
10	ranger-user-catalog-policy	Enabled	Enabled	promethium (USER) admin support@pm61data.com	👁️ 🗑️
15	mask_plan_name	Enabled	Enabled	bella@promethium.al	👁️ 🗑️

Licensed under the Apache License, Version 2.0

To test the rule, return to **Data Explorer**, search for the datamap **nick_ranger_demo**, and click on it. Select the **Samples** tab to load the data. You will see the column you specified masked per your rule.

plan_id	plan_name	mer	disc	package	len	rollup1	asp	startdate
249	-	11.000000000	Discount	All Access	Annual	All Access Discount Annual	Y	2017-11-01
250	-	5.000000000	Discount	Mostly Music	Monthly	Mostly Music Discount Monthly	N	2019-02-01
251	-	15.740000000	Discount	All Access	Monthly	All Access Monthly - 25% Off	N	2019-01-01
252	-	0.000000000	TBD	Select	Semi Annual	Select Semi	TBD	2015-01-05
253	-	6.660000000	TBD	TBD	TBD	Other	TBD	2015-01-05
254	-	6.660000000	TBD	TBD	TBD	Other	TBD	2015-01-05
255	-	6.660000000	Full Price	TBD	TBD	Other	TBD	2015-01-05
256	-	6.660000000	TBD	TBD	TBD	Other	TBD	2015-01-05

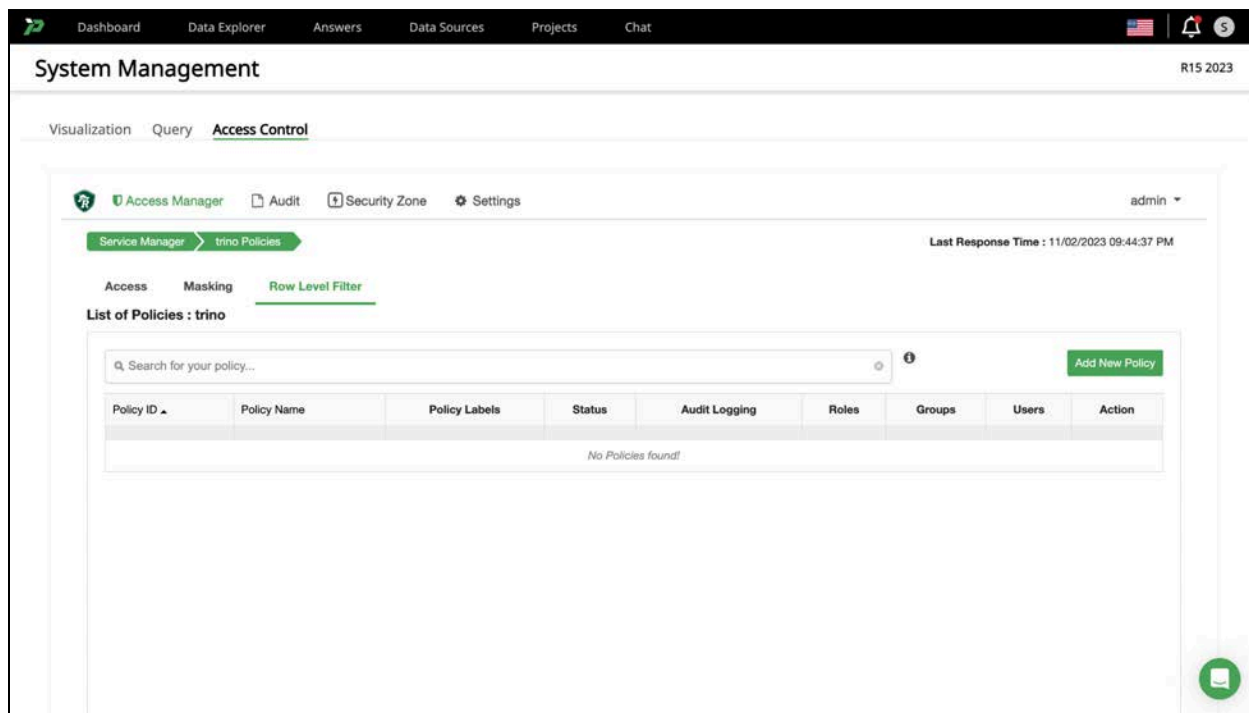
Cancel Edit Datamap



Creating a Role-Level Filter Rule

Navigate to the **Row Level Filter** tab. If you don't see this tab, click the **Ranger** logo to move to the top level navigation, select **Trino** and you should see the correct tab options.

Click **Add New Policy**.



Name the policy in a manner that accurately describes your intended policy rule.

In this example, we will be applying a row-level filter rule to a datamap in Trino.

To apply the masking rule to a datamap in Trino, set the Catalog as **hive** and the schema as **promethium**.

We'll be applying a masking rule on a datamap **nick_ranger_demo** and the column **package**.

Ignore any errors that pop up on the right of the screen claiming that "Resource lookup failed". This does not impact the application masking rule.



Dashboard Data Explorer Answers Data Sources Projects Chat Last Response Time : 11/02/2023 09:45:29 PM

Service Manager trino Policies Create Policy

Create Policy

Please ensure that users/groups listed in this policy have access to the table via an Access Policy. This policy does not implicitly grant access to the table.

Policy Details:

Policy Type: **Row Level Filter** Add Validity Period

Policy Name *: package_is_all_access Enabled Normal

Policy Label: Policy Label

Description: filter to all_access packages only.

Audit Logging: Yes

Resources :

Trino Catalog *: hive

Dashboard Data Explorer Answers Data Sources Projects Chat Last Response Time : 11/02/2023 09:45:29 PM

Service Manager trino Policies Create Policy

Trino Schema *: promethium

Trino Table *: nick_ranger_demo

+ Add Resource

Row Filter Conditions:

Select Role	Select Group	Select User	Access Types	Row Level Filter
Select Roles	Select Groups	bella@promethium.ai	Select	package = 'All Access'

+ Add Cancel

At the bottom of the form are **Row Filter Conditions**. Here you can define which users and rules you want to apply to the specified datamap.

Click in the **Select User** field and either type or scroll to find one of the existing users.

Click the **+** button in **Access Types** and choose **Select**.



Click the **+** button in **Row Level Filter** and enter the filter you want to apply. The filter expression must be a valid WHERE clause for the table or view. To limit the scope of this documentation, we are filtering to display only those rows which include 'all access'.

When you are satisfied with the rules click **Add** to save the rule. The rule will instantly take effect. You will be returned to the full list of masking rules and will see the policy you created at the bottom of the table.

The screenshot shows the Ranger Access Manager interface. At the top, there are navigation tabs for 'Access Manager', 'Audit', 'Security Zone', and 'Settings'. The user is logged in as 'admin'. The main content area is titled 'trino Policies' and includes a search bar and an 'Add New Policy' button. Below this is a table with the following data:

Policy ID	Policy Name	Policy Labels	Status	Audit Logging	Roles	Groups	Users	Action
31	ranger-user-row-level	--	Enabled	Enabled	--	--	ranger, himangshu@pm61data.com, support@pm61data.com	[Eye] [Clipboard] [Delete]
43	sumof_mer-is-7	--	Enabled	Enabled	--	--	support@pm61data.com	[Eye] [Clipboard] [Delete]
45	package_is_all_access	--	Enabled	Enabled	--	--	nick@promethium.ai	[Eye] [Clipboard] [Delete]

To test the rule, return to **Data Explorer**, search for the datamap **nick_ranger_demo**, and click on it. Select the **Samples** tab to load the data. You will see the row level mask you specified ensures only rows matching the mask are displayed.



nick_ranger_demo
Mar 16, 2023

Overview Columns **Samples** Usage Insights Comments Endorsement

Limit Rows: 100 Filter Null 100 rows, 39 KB fetched and 797 ms runtime Search fields

plan_id	plan_name	mer	disc	package	len	rollup1	asp	startdate
249	Promo - All Access - 12mo - 132.00 (+\$1.00 per month)	11.000000000	Discount	All Access	Annual	All Access Discount Annual	Y	2017-11-01
251	All Access - 1mo - wActv Veteran 25% OffA1	15.740000000	Discount	All Access	Monthly	All Access Monthly - 25% Off	N	2019-01-01
257	Introductory - All Access - 12mo w/2mo Savings - wActv	18.990000000	Full Price	All Access	Annual	All Access Annual	Y	2015-01-05
271	Promo - All Access - Semi-Annual - 25% off	15.740000000	Full Price	All Access	Semi Annual	All Access Semi - 25% off	N	2018-02-14

Cancel Edit Datamap ▼

Cluster Management

The Cluster Management page allows you to Create Cluster and run the clusters on an as-needed basis. The purpose of this feature is to give customers flexibility around price, performance, and security when it comes to query and pipeline execution. Rather than having all queries and pipelines consume the same Trino cluster, users now have the option of choosing which Trino cluster they want to use. Customers may choose different Trino clusters because of the following reasons:

- Security:
 - Customers may want to isolate data from different users to only consume a specific Trino cluster so that the data is completely separate and isolated from other users / systems.
- Performance:
 - Customers may want to pin specific DataMaps / Queries to a specific Trino cluster to optimize for performance.
 - In order to maintain the performance of SLA, customers will also want to prevent other DataMaps / Queries from being able to consume specific Trino clusters.
- Cost:
 - Not all DataMaps / Queries should be able to or need to consume the same resources.
 - By being able to create lower powered and power cost Trino clusters, users can now assign lower priority queries to these clusters that have less resources.



- In the event that an organization applies chargeback to their businesses to consume cloud resources, this option allows customers to assign Trino clusters of high performance to those costumes who are willing and able to to pay for them.

The information available will give you insight on the name of clusters available, the status whether the clusters are **Running**, **Starting**, or **Stopped**, actions to **Start** or **Stop** running the cluster, the number of catalogs in each cluster, and the cluster size.

Name	Status	Actions	catalogs	Size
test3	Stopped	▶ START	0 catalogs	free
test8	Stopped	▶ START	5 catalogs	free
test7	Stopped	▶ START	3 catalogs	free

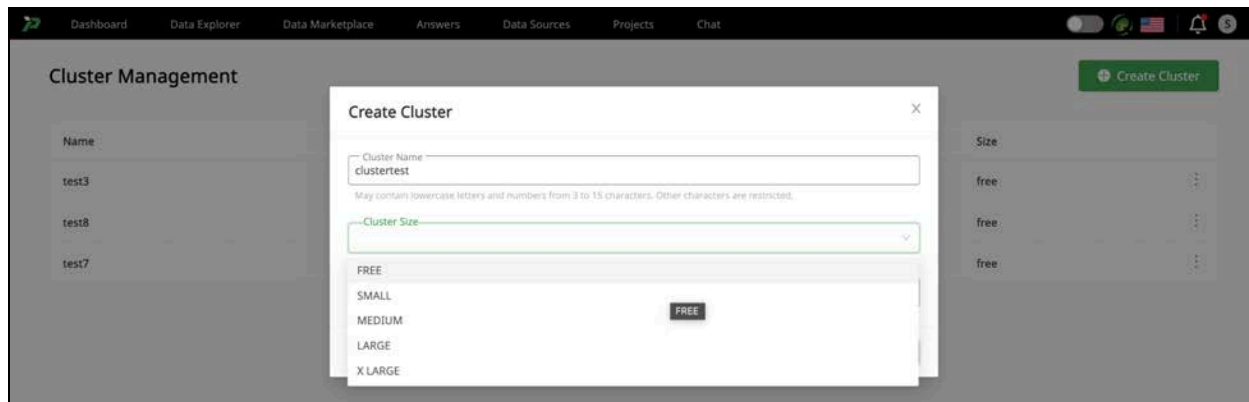
Cluster Creation

Clusters can be created by clicking on the + **Create Cluster** button, where a dialogue box appears with three fields: **Cluster Name**, **Cluster Size**, and **Idle shutdown time**.

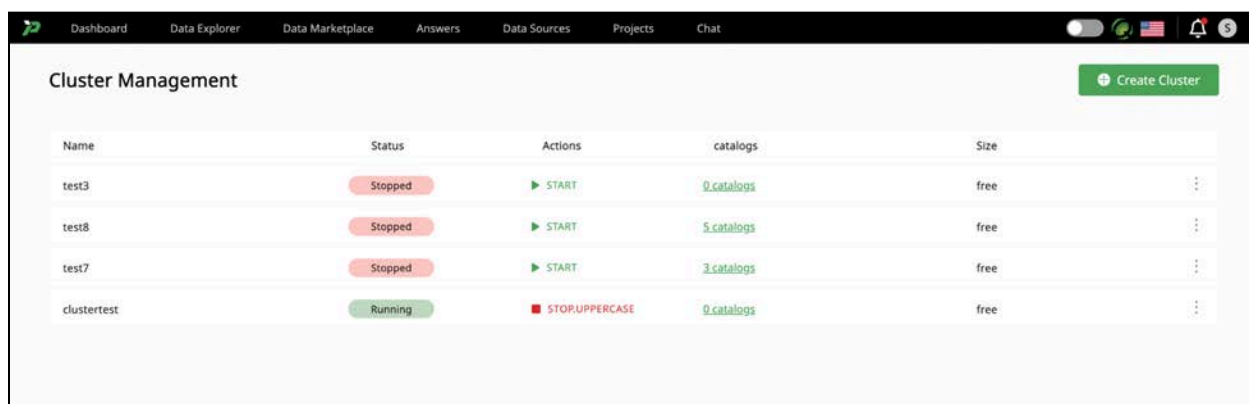
The dialog box titled "Create Cluster" contains the following fields:

- Cluster Name:** A text input field with a note: "May contain lowercase letters and numbers from 3 to 15 characters. Other characters are restricted."
- Cluster Size:** A dropdown menu.
- Idle shutdown time:** A dropdown menu.

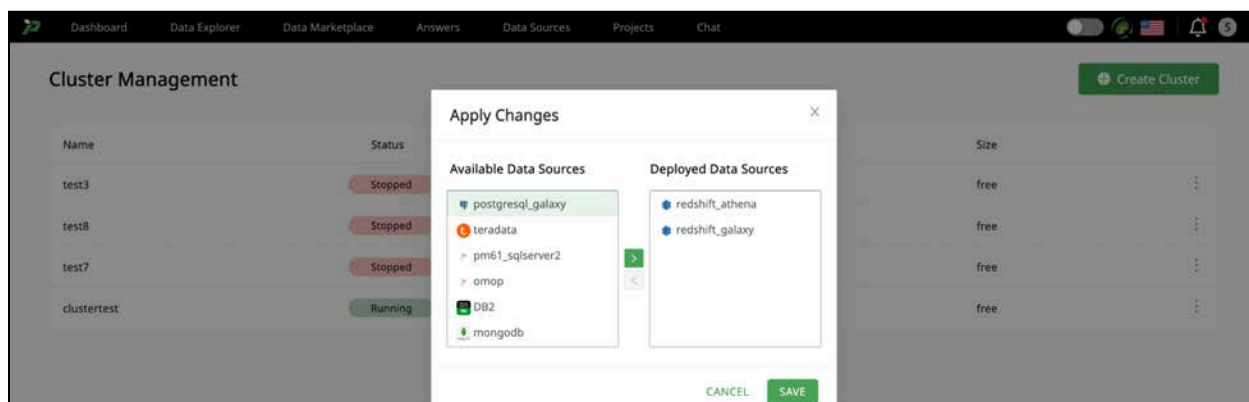
Buttons: CANCEL, CREATE



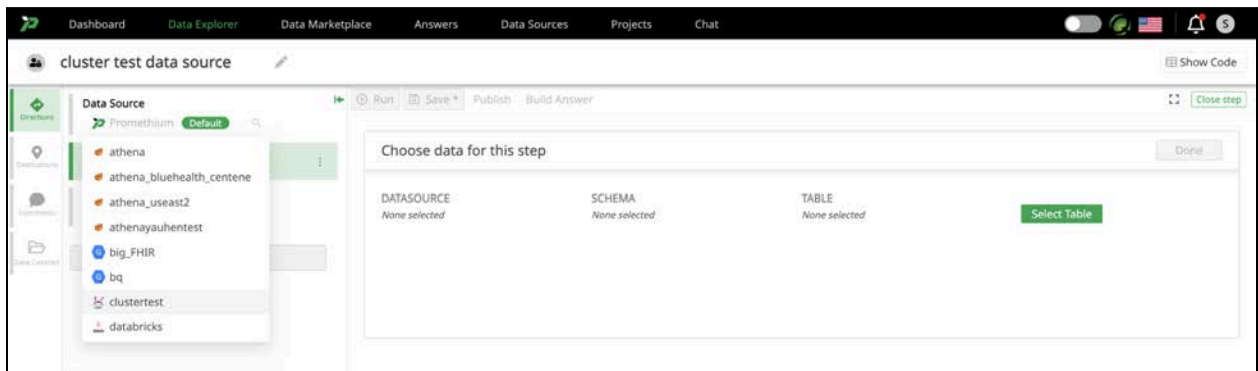
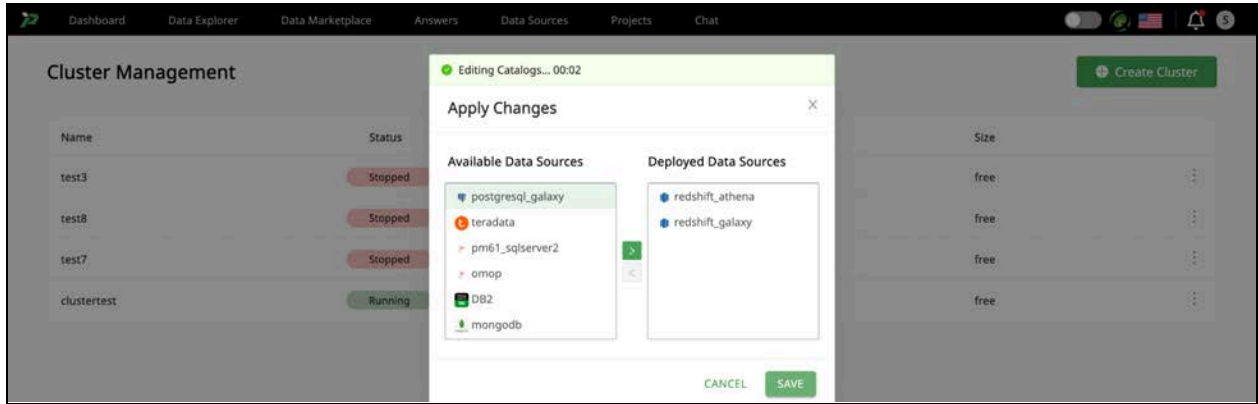
Name the cluster appropriately and select the cluster size from the drop down list, the Idle shutdown time will be autoselected to shut down after 15 or 30 minutes. Once the cluster is created, it will automatically start running.



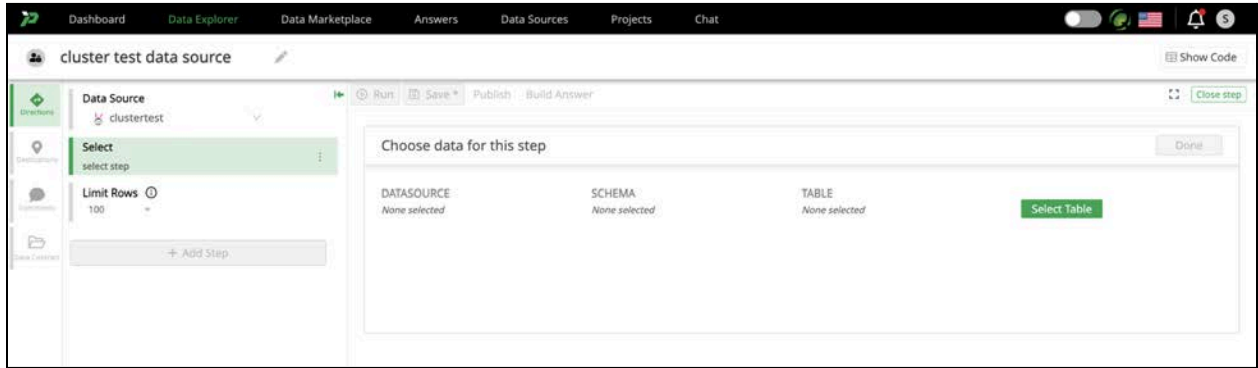
To add catalogs to your cluster, click **0 catalogs** and a dialog box will appear.



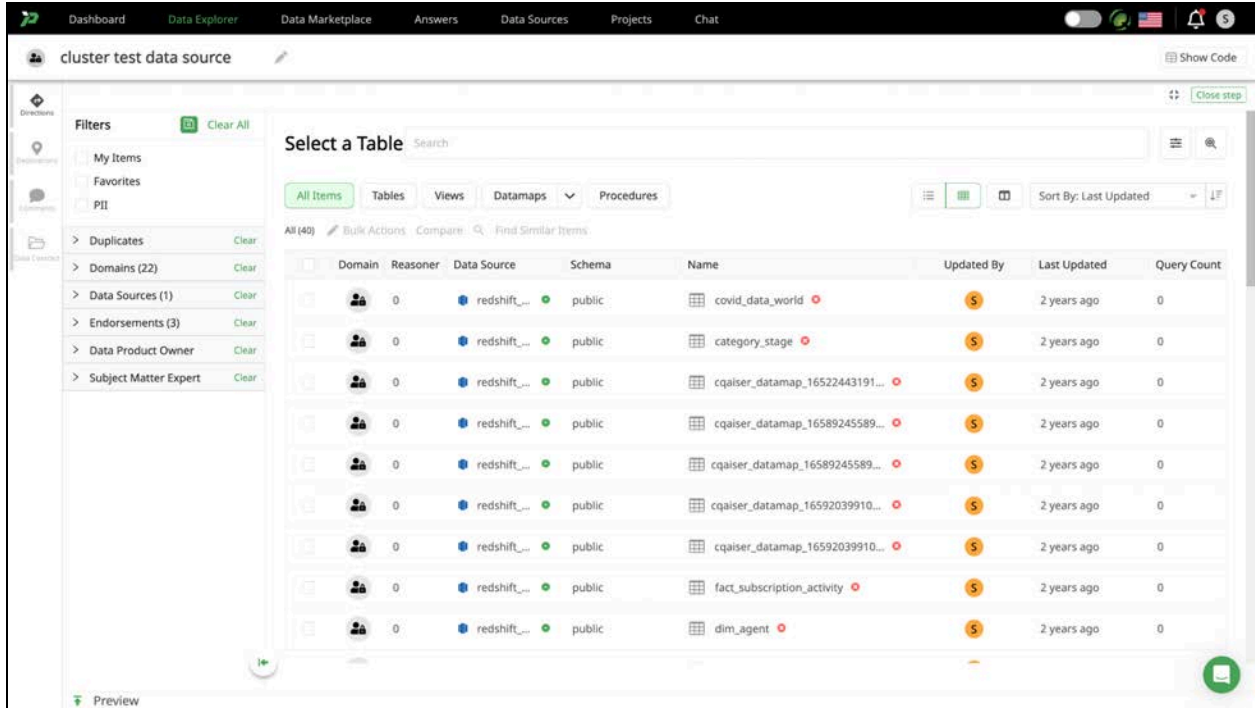
From here, select the multiple data sources and click >. Once you are satisfied with your selection, click **Save** and the cluster will update automatically.



When creating a new datamap, you have the option to select your newly created cluster as the data source.



From the **Select** step, you will be limited to tables in the data sources chosen in the cluster.



System Management

The System Management page houses realtime information on some of the major components within Promethium - Redash, Trino, and Ranger. Additional component summary pages will be added in subsequent releases.

The data presented is mostly informational but can be used to diagnose performance issues, see real time environment statistics, and even restart services.

To navigate to System Management click on your initial or avatar on the top navigation bar. Select System Management from the menu. Select the major system component of interest from the tabbed list at the top of the page.



The screenshot shows the 'System Management' dashboard. At the top, there are navigation tabs: Dashboard, Data Explorer, Answers, Data Sources, Projects, Chat, and a 'Get Started' button. The main content area is titled 'System Management' and includes a 'Query' tab and an 'Auto Scale' toggle. The central focus is the 'CLUSTER OVERVIEW' section, which displays several key performance indicators (KPIs) in a grid format:

- RUNNING QUERIES:** 0
- ACTIVE WORKERS:** 4
- ROWS/SEC:** 0.00
- QUEUED QUERIES:** 0
- RUNNABLE DRIVERS:** 0.00
- BYTES/SEC:** 0
- BLOCKED QUERIES:** 0
- RESERVED MEMORY (B):** 0
- WORKER PARALLELISM:** 0.00

Additional information at the top right of the cluster overview includes: VERSION: B0B6F52-DIRTY, ENVIRONMENT: PRODUCTION, and UPTIME: 7.84d. Below the KPIs, there is a 'QUERY DETAILS' section with a table of query information.

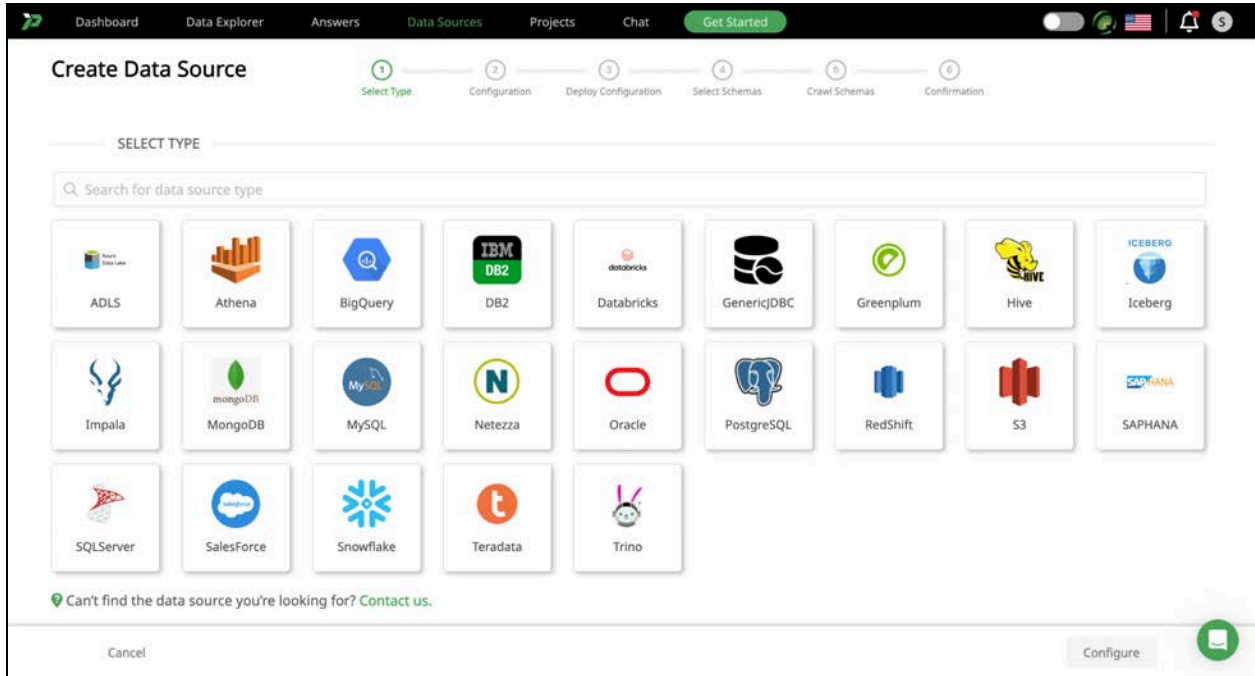
Data Sources

Start by navigating to Data Sources from the menu. The Data Sources page is displayed.

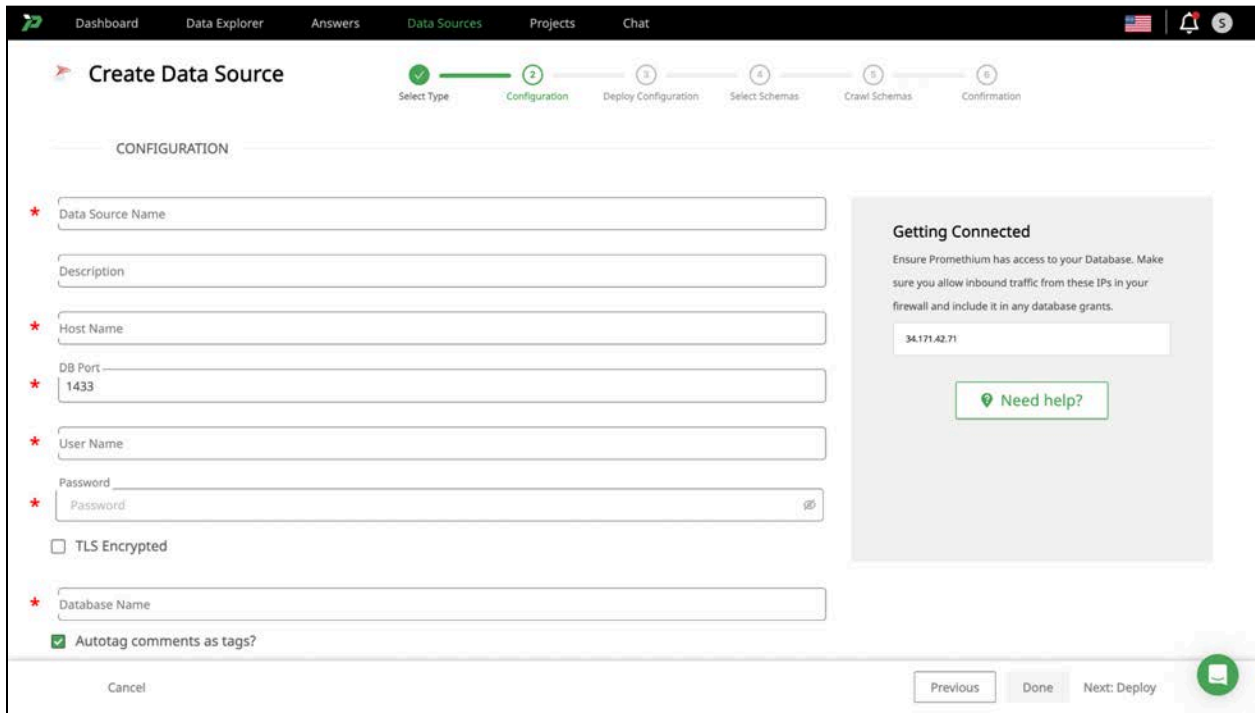
The screenshot shows the 'Data Sources' page. At the top, there are navigation tabs: Dashboard, Data Explorer, Answers, Data Sources, Projects, Chat, and a 'Get Started' button. The main content area is titled 'Data Sources' and includes a search bar and a '+ New Data Source' button. Below the search bar, there is a 'Sort By: Last Updated' dropdown menu. The page displays a grid of data source cards, each representing a different data source:

- Upload VQL Files:** Max file size is 50MB
- postgresql:** Don't delete, Last updated a day ago by [user], Connected
- snowflake:** Last updated 8 days ago by [user], Connected
- postgres_1229_001:** Last updated 13 days ago by [user], Connected
- Automation_DB2_122320...:** Last updated 16 days ago by [user], Connected
- Automation_PostgreSQL_...:** Last updated 16 days ago by [user], Connected
- Automation_PostgreSQL_...:** Last updated 16 days ago by [user], Connected
- Automation_MongoDB_1...:** Last updated 16 days ago by [user], Connected

To add a new Data Source click the **+ New Data Source** button.



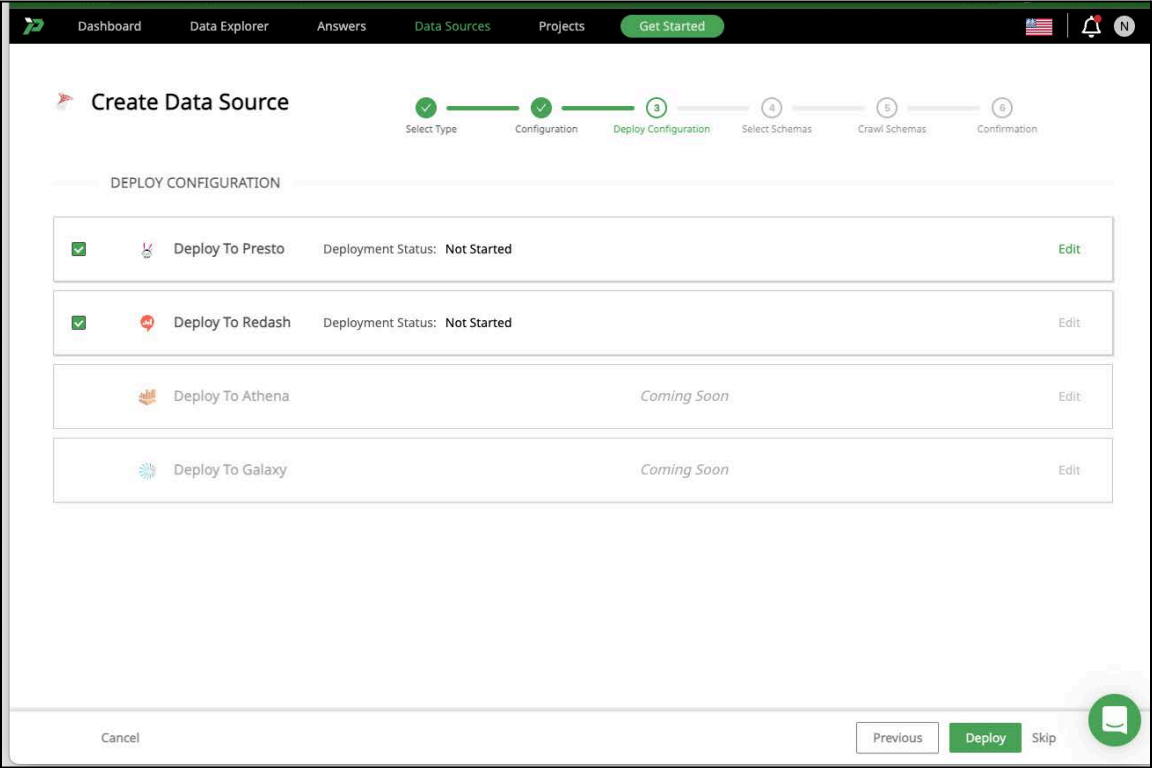
Choose a data source type from the tiles presented, or search for a specific type using the text search box. If you choose to search for a data source, each character typed will filter the available options.



Configure your chosen data source by completing the required fields. Required fields are indicated by the presence of a red asterisk to the left of the field. Optional Advanced

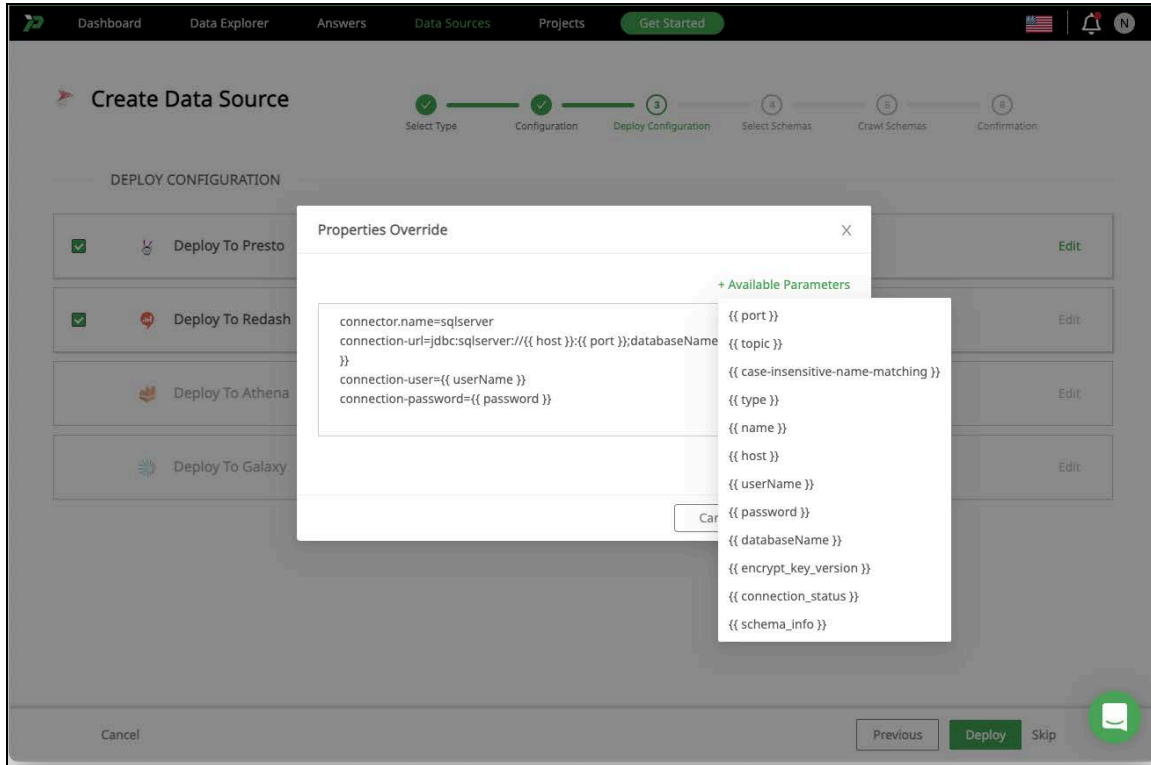


Parameters, if available for your chosen data source, can be entered by clicking on the Advanced Parameters > link at the bottom of the form. When you are satisfied with your configuration details click the green Next: Deploy text at the bottom right of the page.

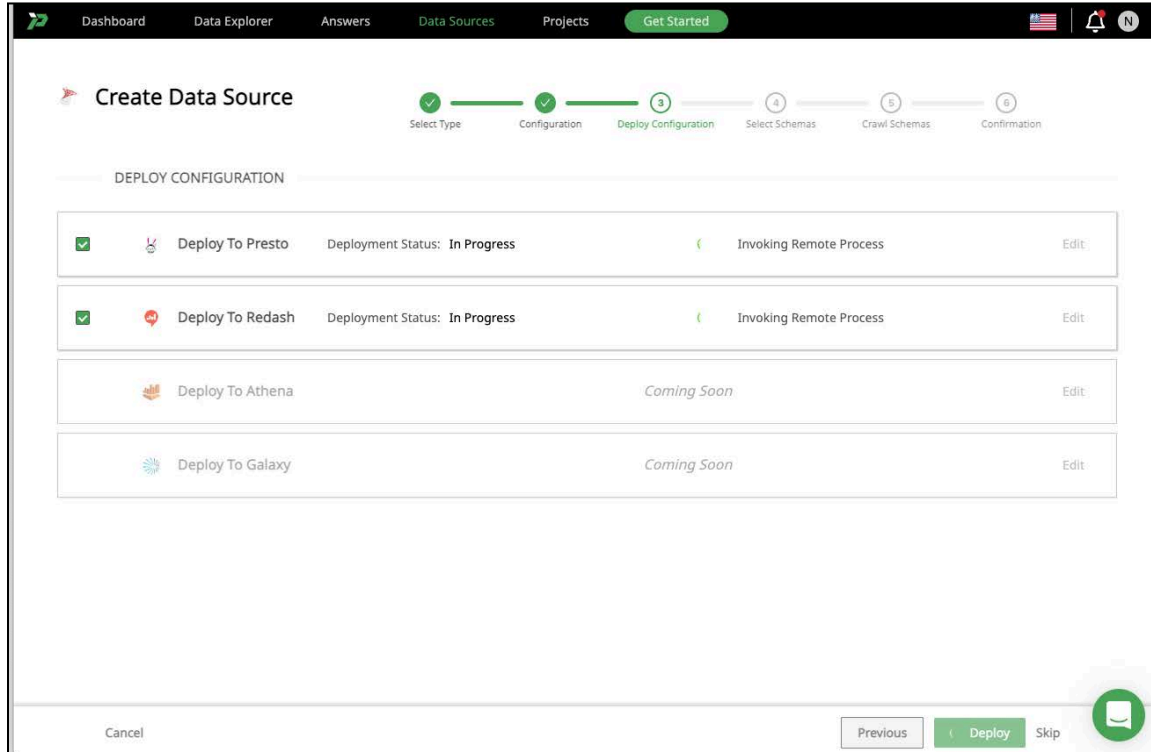


To use the data source, we must deploy the configuration to one or more query engines. You can select each that you wish to use.

The default connector parameters required to configure access will be inferred from your configuration on the previous screen, however, if you need to modify or add additional parameters, you can click Edit next to the query engine and modify the connector.



When you have completed modifications and selected your query engines, click the green Deploy button at the bottom right of the screen.



A status will read Invoking Remote Process next to each selected query engine. The deployment will take up to 2 minutes.



Dashboard Data Explorer Answers Data Sources Projects **Get Started**

Create Data Source

Select Type Configuration Deploy Configuration **Select Schemas** Crawl Schemas Confirmation

Fetching Schema information from sqlserver_demo 00:04

Cancel Previous Next: Crawl

When completed a message will appear informing you that that schema information for your chosen database is being fetched.

Create Data Source

Select Type Configuration Deploy Configuration **Select Schemas** Crawl Schemas Confirmation

SELECT SCHEMA FOR CRAWLING

- promethium
 - db_denydatawriter 0 Tables
 - sales_publish 23 Tables
 - db_ddladmin 0 Tables
 - dbo 50 Tables
 - DB_CRAWLER_IT 1 Tables
 - sales 15 Tables
 - db_datareader 0 Tables
 - dbo_sales 1 Tables
 - db_datawriter 0 Tables
 - willowred 10 Tables
 - db_owner 0 Tables
 - production 4 Tables
 - db_denydatareader 0 Tables
 - cdmDatabaseSchema 44 Tables
 - saleslt 13 Tables
 - db_backupoperator 0 Tables
 - db_securityadmin 0 Tables
 - guest 0 Tables
 - db_accessadmin 0 Tables

Cancel Previous Next: Crawl



You can now select the schemas of interest for crawling. In the example above, the schema 'sales' has been selected. To continue select the green text Next: Crawl at the bottom right of the screen.

The screenshot shows the 'Create Data Source' interface with a progress bar at the top indicating the current step is 'Crawl Schemas'. Below the progress bar, there is a 'CRAWL SCHEMA' section with a 'Crawl Options' icon. The main content is a table with columns for Catalog Name, Schema Name, Status, and a sub-section for Tables (Total Count, Retrieved Count, Error Count). The 'sales' schema is highlighted in green and has a status of 'Retrieved' with a green checkmark.

Catalog Name	Schema Name	Status	Tables		
			Total Count	Retrieved Count	Error Count
promethium	db_denydatawriter	Skipped	0	0	0
promethium	sales_publish	Skipped	23	0	0
promethium	db_ddladmin	Skipped	0	0	0
promethium	dbo	Skipped	50	0	0
promethium	DB_CRAWLER_IT	Skipped	1	0	0
promethium	sales	Retrieved ✓	15	15	0
promethium	db_datareader	Skipped	0	0	0
promethium	dbo_sales	Skipped	1	0	0
promethium	db_datawriter	Skipped	0	0	0

At the bottom right, there are buttons for 'Cancel', 'Previous', 'Confirm', and a green 'Done' button with a checkmark icon.

Your selected schemas will be crawled and the page will refresh to show a confirmation of the data source, schemas, and other table stats. You can either click the green Done button at the bottom right or take action immediately to view your new data within Data Explorer or create a Project.



Create Data Source

CONFIRMATION

Data Source Status	Data Source Name	Data Source Type	Schema(s) Crawled	Schema(s) skipped	Table(s) Retrieved	Table(s) Failed	Size of Metadata Crawled
Connected	sqlserver_demo	SQLServer	1	18	15	0	0

[View in Data Explorer](#) [Create Project from Data Source](#)

Cancel Previous Done

For speed and security Promethium only stores metadata, it does not read or store any actual data.

Redshift

The Redshift connector allows data to be cataloged, queried, and transformed and to create tables in Redshift. This can be used to quickly and easily join data between different systems like Redshift and Hive, or between different Redshift instances.

Connect Redshift

Connecting to Redshift is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Redshift and enter the information for your Redshift instance.



* Data Source Name

Description

* Host Name

* DB Port

* User Name

* Password Password 

* Database Name

Advanced Parameters >

Redshift connection details can be case-sensitive. Ensure that data is entered with the correct case for your Redshift instance.

Field	Description
Data Source Name	Give the Redshift connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Host Name	Enter the hostname for your Redshift database. For example: db.company.com
DB Port	Enter the port number for your Redshift database.
User Name	Enter the username associated with your Redshift account.
Password	Enter the password associated with your Redshift account.
Database Name	Enter the database containing the data you want to retrieve.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Redshift environment, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Redshift environment.



The screenshot shows the Data Explorer interface with a search bar and a 'New Datamap' button. Below the search bar are tabs for 'All Items', 'Tables', 'Views', 'Datamaps', and 'Procedures'. The 'Datamaps' tab is selected. The interface shows a list of 180 items, with the first 10 items displayed in a table. The table has columns for Domain, Reasoner, Data Source, Schema, Name, Updated By, Last Updated, Query Count, and Row Count. The items are sorted by 'Last Updated'.

Domain	Reasoner	Data Source	Schema	Name	Updated By	Last Updated	Query Count	Row Count
	0	redshift	promethium	support_datamap_1704299738040		8 days ago	0	N/A
	0	redshift	promethium	support_datamap_1704298326422		8 days ago	0	N/A
	0	redshift	cutomer_demo_p...	support_datamap_170429834557...		8 days ago	0	N/A
	0	redshift	cutomer_demo_p...	revenue_by_plan_by_vendor_in_2...		9 days ago	0	N/A
	0	redshift	promethium	support_datamap_1704195514401		9 days ago	0	N/A
	0	redshift	cutomer_demo_p...	support_datamap_170419561633...		9 days ago	0	N/A
	0	redshift	promethium	support_datamap_1703610633029		16 days ago	0	N/A
	0	redshift	cutomer_demo_p...	support_datamap_170361065910...		16 days ago	0	N/A
	0	redshift	promethium	support_datamap_1703359597105		19 days ago	0	N/A

FAQs

- Is my data moved or copied out of Redshift?
 - No. Your data stays in Redshift and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, as a table in Redshif, or as a dbt model.
- Is this a Production-level connector?
 - Yes. The Redshift connector is designed for all-sized Datasets

Teradata

The Teradata connector allows data to be cataloged, queried, transformed, and to create tables in Teradata. This can be used to quickly and easily join data between different systems like Teradata and Hive, or between different Teradata instances.

Connect Teradata

Connecting to Teradata is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Teradata and enter the information for your Teradata.



* Data Source Name

Description

* Host Name

* DB Port

* User Name

Password

* Password

Data Source Topic

Private

Advanced Parameters >

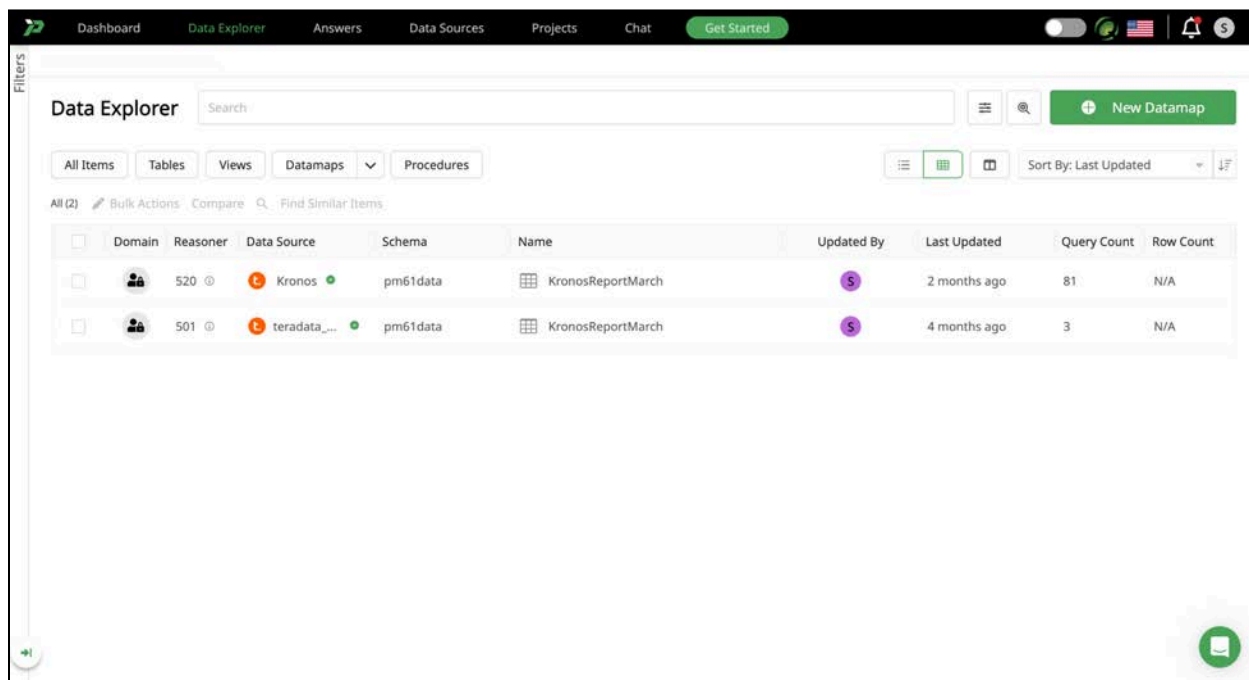
Teradata connection details can be case-sensitive. Ensure that data is entered with the correct case for your Teradata environment.

Field	Description
Data Source Name	Give the Teradata connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Host Name	Enter the hostname for your Teradata database. For example: db.company.com
DB Port	Enter the port number for your Teradata database.
User Name	Enter the username associated with your Teradata account.
Password	Enter the password associated with your Teradata account.
Data Source Domain	Optional. Categorize the Teradata connection. For example, Marketing.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Teradata, then make the data easily searchable and usable.



Click Data Explorer from the menu and begin exploring the data in your Teradata database.



FAQs

- Is my data moved or copied out of Teradata?
 - No. Your data stays in Teradata and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, as a table in Teradata, or as a dbt model.
- Is this a Production-level connector?
 - Yes. The Teradata connector is designed for all-sized Datasets.

Athena

The Athena connector allows data to be cataloged, queried, transformed, and more in Athena. This can be used to quickly and easily join data between different systems like Athena and Hive, or between different Athena instances.

Connect Athena

Connecting to Athena is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Athena and enter the information for your Athena environment.



CONFIGURATION

* Data Source Name

Description

AWS Access Key

* AWS Access Key

AWS Secret Key

* AWS Secret Key

* S3 Output Location

* AWS Region

* Schema Name

Data Source Topic

Private

Advanced Parameters >

Athena connection details can be case-sensitive. Ensure that data is entered with the correct case for your Athena instance.

Field	Description
Data Source Name	Give the Athena Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
AWS Access Key	Unique Key for AWS Athena
AWS Secret Key	Unique Secret for AWS Athena
S3 Output Location	Indicate the location where the query output files are downloaded. Athena will cache all query results in this location.
AWS Region	Enter the region endpoint of your Athena instance. To find the region code from a region name.
Schema Name	Enter the schema containing the data you want to retrieve.



Data Source Domain	Optional. Categorize the Athena connection. For example, Marketing.
--------------------	---

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Athena, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Athena instance.

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table, or as a dbt model.
- Is this a Production-level connector?
 - Yes. The Athena connector is designed for all-sized Datasets.

Snowflake

The Snowflake connector allows data to be cataloged, queried, transformed, and to create tables in Snowflake Data Cloud. This can be used to quickly and easily join data between different systems like Snowflake and Hive, or between different Snowflake instances.



Connect Snowflake

Connecting to Snowflake is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Snowflake and enter the information for your Snowflake environment.

CONFIGURATION

* Data Source Name

Description

* Snowflake Account Name

* User Name

Password

* Warehouse

* Database Name

Data Source Topic

Advanced Parameters >

Snowflake connection details can be case sensitive. Ensure that data is entered with the correct case for your Snowflake.

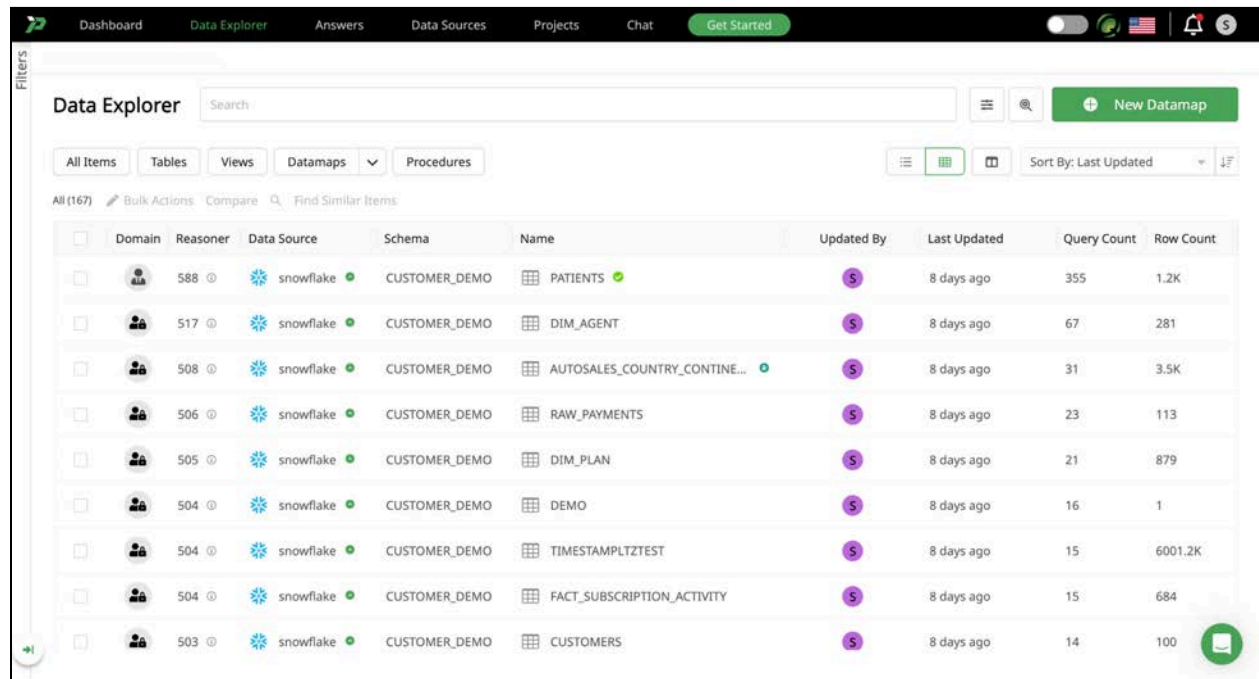
Field	Description
Data Source Name	Give the Snowflake connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Snowflake Account Name	Enter the account name for your Snowflake account.
User Name	Enter the username associated with your Snowflake account.
Password	Enter the password associated with your Snowflake account.
Warehouse	Enter the warehouse containing the data you want to retrieve.



Database Name	Enter the database containing the data you want to retrieve.
Data Source Domain	Optional. Categorize the Snowflake connection. For example, Marketing.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Snowflake, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Snowflake Data Cloud.



FAQs

- Is my data moved or copied out of Snowflake?
 - No. Your data stays in Snowflake and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, as a table in Snowflake or as a dbt model.
- Is this a Production-level connector?
 - Yes. The Snowflake connector is designed for all-sized Datasets.

PostgreSQL

The PostgreSQL connector allows data to be cataloged, queried, transformed, and to create tables in PostgreSQL. This can be used to quickly and easily join data between different systems like PostgreSQL and Hive, or between different PostgreSQL instances.



Connect PostgreSQL

Connecting to PostgreSQL is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select PostgreSQL and enter the information for your PostgreSQL environment.

CONFIGURATION

* Data Source Name

Description

* Host Name

DB Port

* User Name

Password

* Database Name

Data Source Topic

Advanced Parameters >

PostgreSQL connection details can be case sensitive. Ensure that data is entered with the correct case for your PostgreSQL instance.

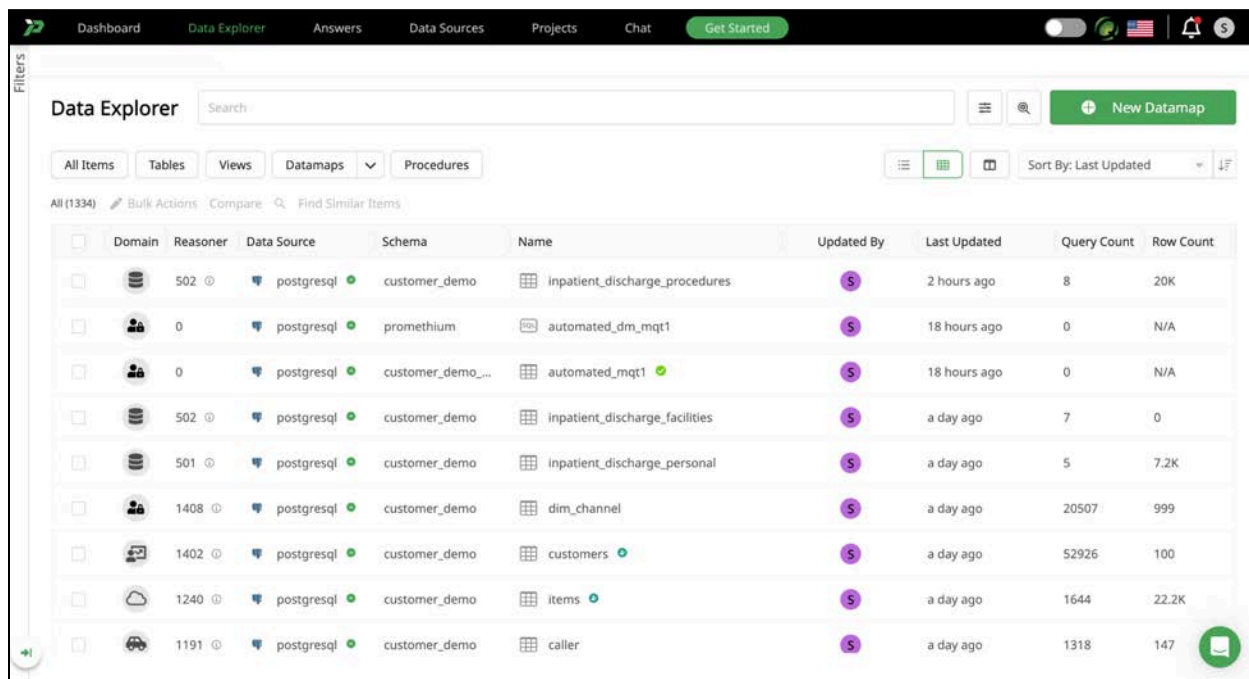
Field	Description
Data Source Name	Give the PostgreSQL connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Host Name	Enter the hostname for your PostgreSQL database. For example: db.company.com
DB Port	Enter the port number for your PostgreSQL database.



User Name	Enter the username associated with your PostgreSQL account.
Password	Enter the password associated with your PostgreSQL account.
Database Name	Enter the database containing the data you want to retrieve.
Data Source Domain	Optional. Categorize the PostgreSQL connection. For example, Marketing.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your PostgreSQL, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your PostgreSQL environment.



FAQs

- Is my data moved or copied out of PostgreSQL?
 - No. Your data stays in PostgreSQL and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, as a table in PostgreSQL or as a dbt model.
- Is this a Production-level connector?
 - Yes. The PostgreSQL connector is designed for all-sized Datasets.



Oracle

The Oracle connector allows data to be cataloged, queried, transformed, and to create tables in Oracle. This can be used to quickly and easily join data between different systems like Oracle and Hive, or between different Oracle instances.

The Oracle connector is available for all Promethium editions, including Free, Professional, Premium, and Enterprise.

Connect Oracle

Connecting to Oracle is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Oracle and enter the information for your Oracle environment.

CONFIGURATION

* Data Source Name

Description

* Host Name

* DB Port

* SID

* User Name

* Password

Data Source Topic

Advanced Parameters >

Oracle connection details can be case-sensitive. Ensure that data is entered with the correct case for your Oracle instance.

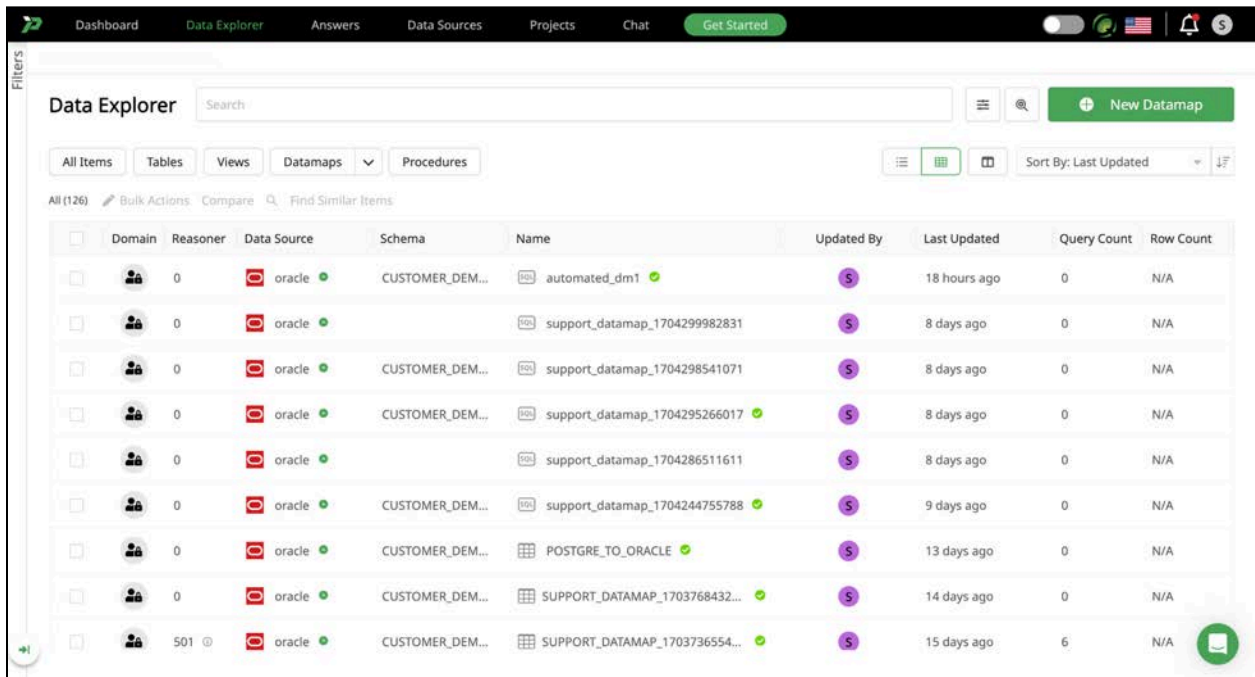
Field	Description
-------	-------------



Data Source Name	Give the Oracle connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Host Name	Enter the hostname for your Oracle database. For example: db.company.com
DB Port	Enter the port number for your Oracle database.
SID	A unique name that uniquely identifies the database instance.
User Name	Enter the username associated with your Oracle account.
Password	Enter the password associated with your Oracle account.
Data Source Domain	Optional. Categorize the Oracle connection. For example, Marketing.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Oracle, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Oracle database.



FAQs

- Is my data moved or copied out of Oracle?
 - No. Your data stays in Oracle and isn't moved to Promethium.



- What options are available for publishing datasets?
 - Datasets can be published as a view, as a table in Oracle or as a dbt model.
- Is this a Production-level connector?
 - Yes. The Oracle connector is designed for all-sized Datasets.

Trino

The connector for Trino allows data sources connected to Trino to be cataloged, queried, transformed, and more. This can be used to quickly and easily join data between different systems.

Note: PrestoSQL was renamed Trino in December 2020.

Connect Trino

Connecting to Trino is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Presto and enter the information for your Trino environment.

CONFIGURATION

* Data Source Name

Description

* Host Name

DB Port

User Name

Password

* Catalog Name

* Schema Name

Advanced Parameters >

Trino connection details can be case sensitive. Ensure that data is entered with the correct case for your Trino environment.



Field	Description
Data Source Name	Give the PrestoSQL Trino connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Host Name	Enter the hostname for your PrestoSQL Trino database. For example: db.company.com
DB Port	Enter the port number for your PrestoSQL Trino.
User Name	Enter the username associated with your PrestoSQL Trino account.
Password	Enter the password associated with your PostgreSQL account.
Catalog Name	Enter the PrestoSQL Trino catalog name.
Schema Name	Enter the PrestoSQL Trino schema name.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Trino environment, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Trino environment.

The screenshot shows the 'Data Explorer' interface. At the top, there's a navigation bar with 'Dashboard', 'Data Explorer', 'Answers', 'Data Sources', 'Projects', 'Chat', and a 'Get Started' button. Below the navigation bar, there's a search bar and a 'New Datamap' button. The main area displays a table of data items. The table has columns: Domain, Reasoner, Data Source, Schema, Name, Updated By, Last Updated, Query Count, and Row Count. There are 17 items listed, all updated 4 or 5 months ago. The items are categorized as 'Tables' and 'Views'.

Domain	Reasoner	Data Source	Schema	Name	Updated By	Last Updated	Query Count	Row Count
500	Promethi...	promethium	new_model_dnnf4fb1	4 months ago	2	N/A		
500	Promethi...	promethium	new_model_ih4tg6un	4 months ago	2	N/A		
500	Promethi...	promethium	new_model_kq1w5oom	4 months ago	2	N/A		
500	Promethi...	promethium	new_model_qt0579il	4 months ago	2	N/A		
500	Promethi...	promethium	new_model_xm5cbgef	4 months ago	2	N/A		
500	Promethi...	default	new_model_8mq8f6ty	4 months ago	2	N/A		
0	Promethi...	default	new_model_qul64knz	4 months ago	0	N/A		
500	Promethi...	promethium	new_model_y8031zsc	4 months ago	2	N/A		
0	Promethi...	default	new_model_zp1ixjgy	5 months ago	0	N/A		

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.



- What options are available for publishing datasets?
 - Datasets can be published as a view, as a table, or as a dbt model.
- Is this a Production-level connector?
 - Yes. The PrestoSQL Trino connector is designed for all-sized Datasets.

MS SQL Server

The MS SQL Server connector allows data to be cataloged, queried, transformed, and to create tables in MS SQL Server. This can be used to quickly and easily join data between different systems like MS SQL Server and Hive, or between different MS SQL Server instances.

Connect MS SQL Server

Connecting to MS SQL Server is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select MS SQL Server and enter the information for your MS SQL Server.

The screenshot shows a configuration form for connecting to MS SQL Server. It includes the following fields and options:

- Data Source Name**: Required field (marked with a red asterisk).
- Description**: Optional text field.
- Host Name**: Required field (marked with a red asterisk).
- DB Port**: Required field (marked with a red asterisk), currently set to 1433.
- User Name**: Required field (marked with a red asterisk).
- Password**: Required field (marked with a red asterisk), currently masked with a red dot.
- Database Name**: Required field (marked with a red asterisk).
- Data Source Topic**: Optional dropdown menu, currently set to "Private" with a lock icon.
- Advanced Parameters**: A link to expand the configuration options.

MS SQL Server connection details can be case-sensitive. Ensure that data is entered with the correct case for your MS SQL Server.

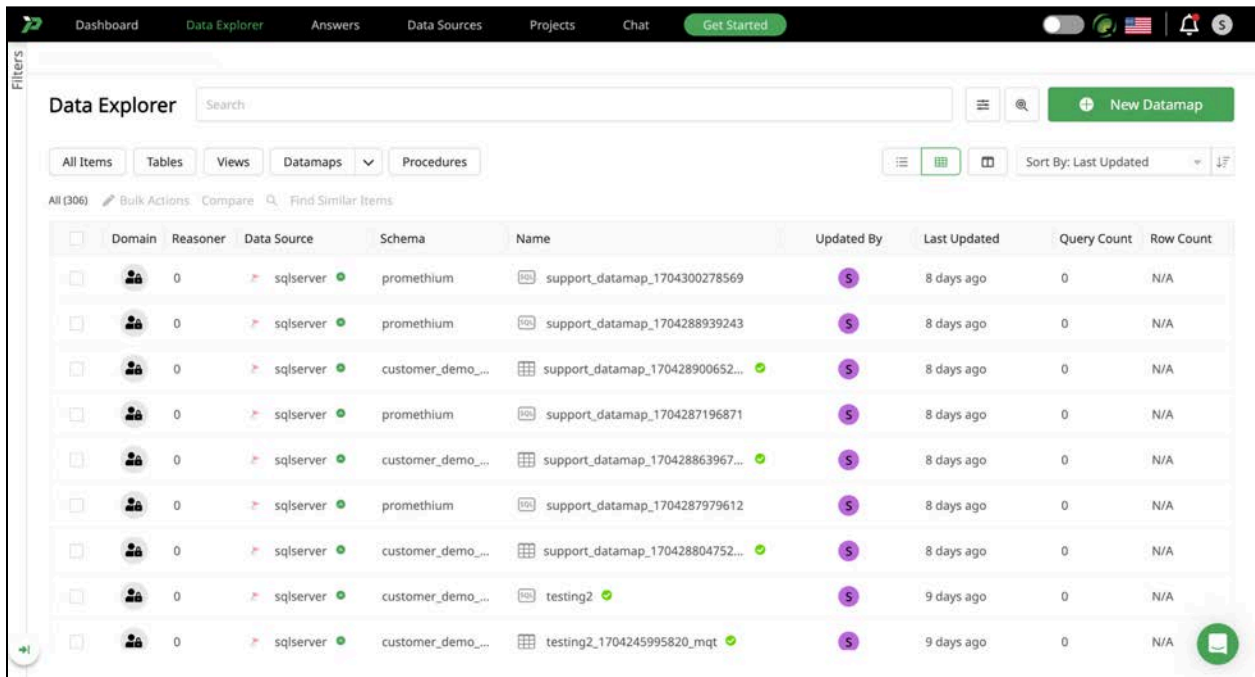
Field	Description
-------	-------------



Data Source Name	Give the MS SQL Server connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Host Name	Enter the hostname for your MS SQL Server database. For example: db.company.com
DB Port	Enter the port number for your MS SQL Server database.
User Name	Enter the username associated with your MS SQL Server account.
Password	Enter the password associated with your MS SQL Server account.
Database Name	Enter the database containing the data you want to retrieve.
Data Source Domain	Optional. Categorize the MS SQL Server connection. For example, Marketing.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your MS SQL Server, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your MS SQL Server.



FAQs

- Is my data moved or copied out of MS SQL Server?
 - No. Your data stays in MS SQL Server and isn't moved to Promethium.
- What options are available for publishing datasets?



- o Datasets can be published as a view, as a table in MS SQL Server, or as a dbt model.
- Is this a Production-level connector?
 - o Yes. The MS SQL Server connector is designed for all-sized Datasets.

Impala

The Impala connector allows data to be cataloged, queried and transformed. This can be used to quickly and easily join data between different systems like Impala and HDFS, or between different Impala instances.

Connect Impala

Connecting to Impala is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Impala, select New connection, and enter the information for your Impala instance.



CONFIGURATION

* Data Source Name

Description

* Host Name

* DB Port

User Name

Password

* Database Name

Use HTTP Transport

Timeout

Http Path

Auth Mechanism

Use SSL

Data Source Topic

Advanced Parameters >

Impala connection details can be case-sensitive. Ensure that data is entered with the correct case for your Impala environment.

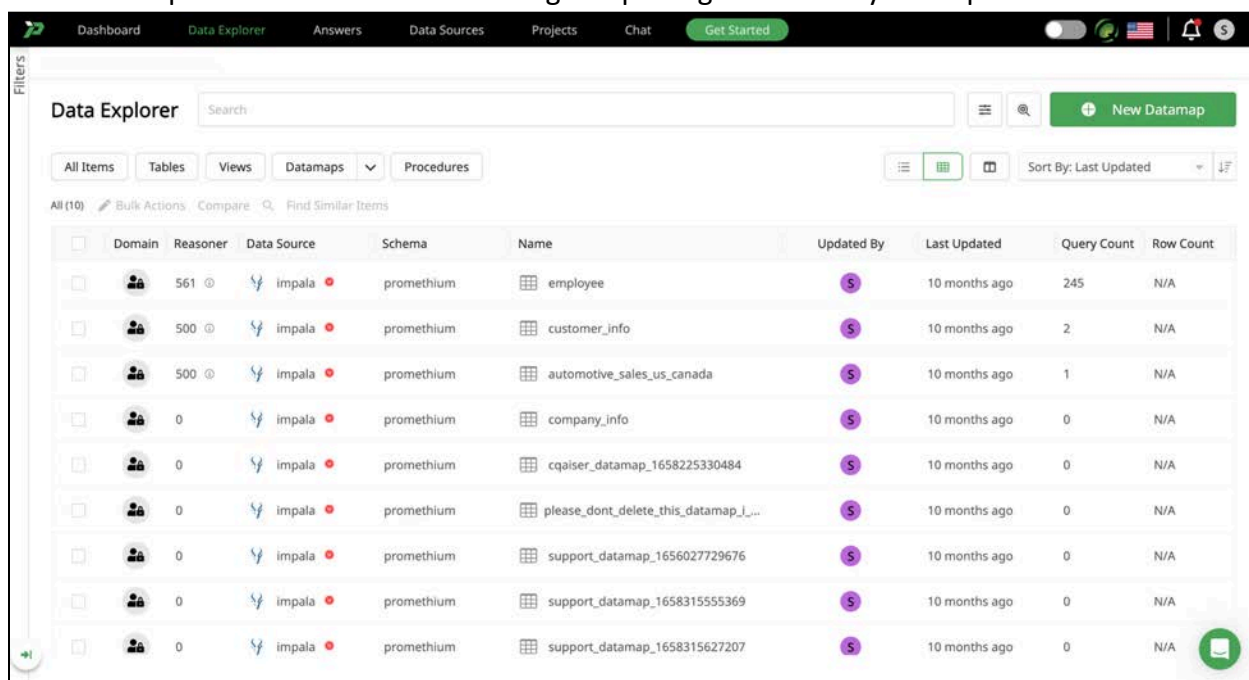
Field	Description
Data Source Name	Give the Impala connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Hostname	The name of the server where your data is located.
DB Port	Enter the port number for your Impala.
Username	Enter the username associated with your Impala account.



Password	Enter the password associated with your Impala account.
Database Name	Enter the database containing the data you want to retrieve.
Use HTTP Transport	Select this checkbox to use HTTP for your Impala connection.
Use secure sockets layer (SSL)	Select this checkbox to connect using SSL.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Impala environment, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Impala data source.



FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, as a table, or as a dbt model.
- Is this a Production-level connector?
 - Yes. The Impala connector is designed for all-sized Datasets.

Greenplum

The Greenplum connector allows data to be cataloged, queried, transformed and to create tables in Greenplum. This can be used to quickly and easily join data between different systems like Greenplum and Hive, or between different Greenplum instances.



Connect Greenplum

Connecting to Greenplum is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Greenplum and enter the information for your Greenplum.

* Data Source Name

Description

* Host Name

* DB Port

* User Name

Password

* Password

* Database Name

Data Source Topic

Private

Advanced Parameters >

Greenplum connection details can be case-sensitive. Ensure that data is entered with the correct case for your Greenplum environment.

Field	Description
Data Source Name	Give the Greenplum connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Host Name	Enter the hostname for your Greenplum database. For example: db.company.com
DB Port	Enter the port number for your Greenplum database.
User Name	Enter the username associated with your Greenplum account.
Password	Enter the password associated with your Greenplum account.



Database Name	Enter the database containing the data you want to retrieve.
Data Source Domain	Optional. Categorize the Greenplum connection. For example, Marketing.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Greenplum, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Greenplum database.

FAQs

- Is my data moved or copied out of Greenplum?
 - No. Your data stays in Greenplum and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, as a table in Greenplum, or as a dbt model.
- Is this a Production-level connector?
 - Yes. The Greenplum connector is designed for all-sized Datasets.

ADLS

The ADLS connector allows data to be cataloged, queried, transformed and more in ADLS. This can be used to quickly and easily join data between different systems like ADLS and Hive, or between different ADLS instances.

Connect ADLS

Connecting to ADLS is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select ADLS and enter the information for your ADLS environment.



Account key
Azure Data Lake Storage URI
Data Source Domain
[▶ Advanced Parameters](#)

ADLS connection details can be case-sensitive. Ensure that data is entered with the correct case for your ADLS instance.

Field	Description
Data Source Name	Give the ADLS Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Account Key	Account key associated with your ADLS account.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your ADLS, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your ADLS instance.



FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table, or as a dbt model.
- Is this a Production-level connector?
 - Yes. The ADLS connector is designed for all-sized Datasets.

BigQuery

The BigQuery connector allows data to be cataloged, queried, transformed and more in BigQuery. This can be used to quickly and easily join data between different systems like BigQuery and Hive, or between different BigQuery instances.

Connect BigQuery

Connecting to BigQuery is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select BigQuery and enter the information for your BigQuery environment.

Autotag comments as tags?
Data Source Domain
[▶ Advanced Parameters](#)



BigQuery connection details can be case-sensitive. Ensure that data is entered with the correct case for your BigQuery instance.

Field	Description
Data Source Name	Give the BigQuery Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Credentials Key JSON	Enter the credentials associated with your BigQuery account.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your BigQuery, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your BigQuery instance.

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The BigQuery connector is designed for all-sized Datasets.




DB2

The DB2 connector allows data to be cataloged, queried, transformed and more in DB2. This can be used to quickly and easily join data between different systems like DB2 and Hive, or between different DB2 instances.

Connect DB2

Connecting to DB2 is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select DB2 and enter the information for your DB2 environment.



Data Source Name
Description
Host Name
DB Port 50000
User Name
Password Password 
DB2 Version 10.5
Database Name
<input checked="" type="checkbox"/> Autotag comments as tags?
Data Source Domain  Private 
▶ Advanced Parameters

DB2 connection details can be case-sensitive. Ensure that data is entered with the correct case for your DB2 instance.

Field	Description
Data Source Name	Give the DB2 connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Host Name	Enter the hostname for your DB2 database. For example: db.company.com
DB Port	Enter the port number for your DB2.
User Name	Enter the username associated with your DB2 account.
Password	Enter the password associated with your DB2 account.
DB2 Version	Enter the version associated with your DB2 account.



Database Name	Enter the database containing the data you want to retrieve.
---------------	--

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your DB2, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your DB2 instance.

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The DB2 connector is designed for all-sized Datasets.

Databricks

The Databricks connector allows data to be cataloged, queried, transformed and more in Databricks. This can be used to quickly and easily join data between different systems like Databricks and Hive, or between different Databricks instances.

Connect Databricks

Connecting to Databricks is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Databricks and enter the information for your Databricks environment.



Data Source Name

Description

JDBC Url

Http Path
Http Path

Database

UID
token

Personal access token
Personal access token

Autotag comments as tags?

Data Source Domain
 Private

▶ **Advanced Parameters**

Databricks connection details can be case-sensitive. Ensure that data is entered with the correct case for your Databricks instance.

Field	Description
Data Source Name	Give the Databricks Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
JDBC URL	URL associated with the Databricks database to connect to.
Http Path	Enter the http containing the data you want to retrieve.
Database	Enter the database name containing the data you want to retrieve.



Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Databricks, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Databricks instance.

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The Databricks connector is designed for all-sized Datasets.

GenericJDBC

The GenericJDBC connector allows data to be cataloged, queried, transformed and more in GenericJDBC. This can be used to quickly and easily join data between different systems like GenericJDBC and Hive, or between different GenericJDBC instances.

Connect GenericJDBC

Connecting to GenericJDBC is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select GenericJDBC and enter the information for your GenericJDBC environment.



Autotag comments as tags?
Data Source Domain

[▶ Advanced Parameters](#)

GenericJDBC connection details can be case-sensitive. Ensure that data is entered with the correct case for your GenericJDBC instance.

Field	Description
Data Source Name	Give the GenericJDBC Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
JDBC URL	Credentials associated with your GenericJDBC account.
User Name	Enter the username associated with your GenericJDBC account.
Password	Enter the password associated with your GenericJDBC account.



Additional Credential File path	Enter the file path associated with your GenericJDBC account.
Driver class Name	Enter the driver name associated with your GenericJDBC account.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your GenericJDBC, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your GenericJDBC instance.

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The GenericJDBC connector is designed for all-sized Datasets.

Hive

The Hive connector allows data to be cataloged, queried, transformed and more in Hive. This can be used to quickly and easily join data between different systems like Hive, or between different Hive instances.

Connect Hive

Connecting to Hive is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Hive and enter the information for your Hive environment.




Data Source Name

Description

Host Name


DB Port

User Name

Password
Password 

Database Name

Autotag comments as tags?

Data Source Domain
 Private

▶ [Advanced Parameters](#)

Hive connection details can be case-sensitive. Ensure that data is entered with the correct case for your Hive instance.

Field	Description
Data Source Name	Give the Hive Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Host Name	Credentials associated with your Hive account.
DB Port	Enter the port number for your Hive.
User Name	Enter the username associated with your Hive account.
Password	Enter the password associated with your Hive account.
Database Name	Enter the database name you want to connect to.



Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Hive, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Hive instance.

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The Hive connector is designed for all-sized Datasets.



Iceberg

The Iceberg connector allows data to be cataloged, queried, transformed and more in Iceberg. This can be used to quickly and easily join data between different systems like Iceberg and Hive, or between different Iceberg instances.

Connect Iceberg

Connecting to Iceberg is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Iceberg and enter the information for your Iceberg environment.



Data Source Name
Metastore Type glue
Description
User Name
Password Password 
Glue Region
Data Source Domain  Private
▶ Advanced Parameters

Iceberg connection details can be case-sensitive. Ensure that data is entered with the correct case for your Iceberg instance.

Field	Description
Data Source Name	Give the Iceberg Connection a unique name for easy identification.
Metastore Type	glue
Description	Optional. A meaningful description of this connection.
User Name	Enter the username associated with your Iceberg account.
Password	Enter the password associated with your Iceberg account.
Glue Region	Enter the region associated with your Iceberg account.



Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Iceberg, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Iceberg instance.

FAQs


- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The Iceberg connector is designed for all-sized Datasets.

MongoDB

The MongoDB connector allows data to be cataloged, queried, transformed and more in MongoDB. This can be used to quickly and easily join data between different systems like MongoDB and Hive, or between different MongoDB instances.

Connect MongoDB

Connecting to MongoDB is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select MongoDB and enter the information for your MongoDB environment.



▶ **Advanced Parameters**

MongoDB connection details can be case-sensitive. Ensure that data is entered with the correct case for your MongoDB instance.



Field	Description
Data Source Name	Give the MongoDB Connection a unique name for easy identification.
MongoDB JDBC URL	Enter the URL to the MongoDB database you want to connect to.
Description	Optional. A meaningful description of this connection.
User Name	Enter the username associated with your MongoDB account.
Password	Enter the password associated with your MongoDB account.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your MongoDB, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your MongoDB instance.

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The MongoDB connector is designed for all-sized Datasets.

MySQL

The MySQL connector allows data to be cataloged, queried, transformed and more in MySQL. This can be used to quickly and easily join data between different systems like MySQL and Hive, or between different MySQL instances.

Connect MySQL

Connecting to MySQL is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select MySQL and enter the information for your MySQL environment.



Autotag comments as tags?

[▶ Advanced Parameters](#)

MySQL connection details can be case-sensitive. Ensure that data is entered with the correct case for your MySQL instance.

Field	Description
Data Source Name	Give the MySQL Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Host Name	Enter the host name associated with your MySQL account.
DB Port	3306
User Name	Enter the username associated with your MySQL account.
Password	Enter the password associated with your MySQL account.
Database Name	Enter the database name you want to connect to.



Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your MySQL, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your MySQL instance.

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The MySQL connector is designed for all-sized Datasets.

Netezza

The Netezza connector allows data to be cataloged, queried, transformed and more in Netezza. This can be used to quickly and easily join data between different systems like Netezza and Hive, or between different Netezza instances.

Connect Netezza

Connecting to Netezza is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Netezza and enter the information for your Netezza environment.




Data Source Name

Description

Host Name

DB Port


User Name

Password
Password 

Database Name

Schema

Autotag comments as tags?

Data Source Domain
 Private

▶ **Advanced Parameters**

Netezza connection details can be case-sensitive. Ensure that data is entered with the correct case for your Netezza instance.

Field	Description
Data Source Name	Give the Netezza Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Host Name	Enter the host name associated with your Netezza account.
DB Port	Enter the port number you want to access.
User Name	Enter the username associated with your Netezza account.
Password	Enter the password associated with your Netezza account.
Database Name	Enter the Netezza database name you want to connect to.



Schema	Enter the Netezza schema name you want to connect to.
--------	---

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Netezza, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Netezza instance.

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The Netezza connector is designed for all-sized Datasets.

SAPHANA

The SAPHANA connector allows data to be cataloged, queried, transformed and more in SAPHANA. This can be used to quickly and easily join data between different systems like SAPHANA and Hive, or between different SAPHANA instances.

Connect SAPHANA

Connecting to SAPHANA is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select SAPHANA and enter the information for your SAPHANA environment.



Autotag comments as tags?
Data Source Domain

[▶ Advanced Parameters](#)

SAPHANA connection details can be case-sensitive. Ensure that data is entered with the correct case for your SAPHANA instance.

Field	Description
Data Source Name	Give the SAPHANA Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Connection URL	Enter the URL to the SAPHANA database you want to connect to.
User Name	Enter the username associated with your SAPHANA account.
Password	Enter the password associated with your SAPHANA account.
Database Name	Enter the SAPHANA database name you want to connect to.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your SAPHANA, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your SAPHANA instance.



FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The SAPHANA connector is designed for all-sized Datasets.

SalesForce

The SalesForce connector allows data to be cataloged, queried, transformed and more in SalesForce. This can be used to quickly and easily join data between different systems like SalesForce and Hive, or between different SalesForce instances.

Connect SalesForce

Connecting to SalesForce is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select SalesForce and enter the information for your SalesForce environment.

jdbc:dbschema:salesforce://tables=all

default
 Autotag comments as tags?
[▶ Advanced Parameters](#)

SalesForce connection details can be case-sensitive. Ensure that data is entered with the correct case for your SalesForce instance.



Field	Description
Data Source Name	Give the Salesforce Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Salesforce Host	jdbc:dbschema:salesforce://tables=all
User Name	Enter the username associated with your Salesforce account.
Password	Enter the password associated with your Salesforce account.
Token	Enter the token key associated with your Salesforce account.
Database Name	Enter the Salesforce database name you want to connect to.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Salesforce, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Salesforce instance.

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The Salesforce connector is designed for all-sized Datasets.



ServiceNow

The ServiceNow connector allows data to be cataloged, queried, transformed and more in ServiceNow. This can be used to quickly and easily join data between different systems like ServiceNow and Hive, or between different ServiceNow instances.

Connect ServiceNow

Connecting to ServiceNow is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select ServiceNow and enter the information for your ServiceNow environment.



Data Source Name
Description
Connection URL
User Name
Password Password 
Database Name
<input checked="" type="checkbox"/> Autotag comments as tags?
Data Source Domain  Private
▶ Advanced Parameters

ServiceNow connection details can be case-sensitive. Ensure that data is entered with the correct case for your ServiceNow instance.

Field	Description
Data Source Name	Give the ServiceNow Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Connection URL	Enter the URL to the ServiceNow database you want to connect to.
User Name	Enter the username associated with your ServiceNow account.
Password	Enter the password associated with your ServiceNow account.
Database Name	Enter the ServiceNow database name you want to connect to.



Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your ServiceNow, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your ServiceNow instance.

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The ServiceNow connector is designed for all-sized Datasets.

Google Sheets

The Google Sheets connector allows data to be cataloged, queried, transformed and more in Google Sheets. This can be used to quickly and easily join data between different systems like Google Sheets and Hive, or between different Google Sheets instances.

Connect Google Sheets

Connecting to Google Sheets is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select Google Sheets and enter the information for your Google Sheets environment.



Autotag comments as tags?

[▶ Advanced Parameters](#)

Google Sheets connection details can be case-sensitive. Ensure that data is entered with the correct case for your Google Sheets instance.

Field	Description
Data Source Name	Give the Google Sheets Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
Credentials Key JSON	Enter the JSON credentials for access to Google Sheets.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your Google Sheets, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your Google Sheets instance.

FAQs

- Is my data moved or copied?




- o No. Your data stays where it is and isn't moved to Prometheus.
- What options are available for publishing datasets?
 - o Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - o Yes. The Google Sheets connector is designed for all-sized Datasets.


API

The API connector allows data to be cataloged, queried, transformed and more in API. This can be used to quickly and easily join data between different systems like API and Hive, or between different API instances.

Connect API

Connecting to the API data source is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select API and enter the information for your API environment.



 
▶ Advanced Parameters

API connection details can be case-sensitive. Ensure that data is entered with the correct case for your API instance.



Field	Description
Data Source Name	Give the API Connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
API Spec Location	Where the documentation or definition files of an API are hosted.
Authentication ClientID	Enter the ClientID/username associated with the API datasource.
Authentication Client Secret	Enter the Client Secret/password associated with the API datasource.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your API, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your API instance.

FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - Yes. The API connector is designed for all-sized Datasets.

ClickHouse

The ClickHouse connector allows data to be cataloged, queried, transformed and more in ClickHouse. This can be used to quickly and easily join data between different systems like ClickHouse and Hive, or between different ClickHouse instances.

Connect ClickHouse

Connecting to ClickHouse is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select ClickHouse and enter the information for your ClickHouse environment.



Private ▼

▶ Advanced Parameters

ClickHouse connection details can be case-sensitive. Ensure that data is entered with the correct case for your ClickHouse instance.

Field	Description
Data Source Name	Give the ClickHouse Connection a unique name for easy identification.
JDBC URL	URL associated with the ClickHouse database to connect to.
Description	Optional. A meaningful description of this connection.
Username	Enter the username associated with your ClickHouse account.
Password	Enter the password associated with your ClickHouse account.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your API, then make the data easily searchable and usable.

Click Data Explorer from the menu and begin exploring the data in your ClickHouse instance.

FAQs

- Is my data moved or copied?



- o No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - o Datasets can be published as a view, table or as a dbt model.
- Is this a Production-level connector?
 - o Yes. The ClickHouse connector is designed for all-sized Datasets.

AWS S3

Promethium needs read access to the S3 bucket from Intelligent Edge (IE). This should be setup and tested after installing IE in your Cloud. Contact us for more information or assistance.

In order for Promethium to discover your data in S3, the Promethium Intelligent Edge (IE) needs to have access to it:

1. Navigate to your Bucket on S3, then make sure it is accessible from outside by changing the “Block all public access” to off (click on the Edit button on the right side to change this, as shown below in the screenshot):
2. Go to Bucket Policy
3. In order to provide this access, you need to add the following JSON with the correct values:



aws Services Resource Groups

Amazon S3 > s3mediadata

s3mediadata

Overview Properties **Bucket Policy** Access points

Block public access Access Control List **Bucket Policy** CORS configuration

Bucket policy use JSON-based access policy language to manage advanced permission to your Amazon S3 resources.

Block public access (bucket settings)

Public access is granted to buckets and objects through access control lists (ACLs), bucket policies, access point policies, or all. In order to ensure that public access to all your S3 buckets and objects is blocked, turn on Block all public access. These settings apply only to this bucket and its access points. AWS recommends that you turn on Block all public access, but before applying any of these settings, ensure that your applications will work correctly without public access. If you require some level of public access to your buckets or objects within, you can customize the individual settings below to suit your specific use cases. [Learn more](#)

Amazon S3 > s3mediadata

s3mediadata

Overview Properties **Permissions** Management Access points

Block public access Access Control List Bucket Policy CORS configuration

Block public access (bucket settings)

Public access is granted to buckets and objects through access control lists (ACLs), bucket policies, access point policies, or all. In order to ensure that public access to all your S3 buckets and objects is blocked, turn on Block all public access. These settings apply only to this bucket and its access points. AWS recommends that you turn on Block all public access, but before applying any of these settings, ensure that your applications will work correctly without public access. If you require some level of public access to your buckets or objects within, you can customize the individual settings below to suit your specific storage use cases. [Learn more](#)

Block all public access Off Edit

- Block public access to buckets and objects granted through new access control lists (ACLs) Off
- Block public access to buckets and objects granted through any access control lists (ACLs) Off
- Block public access to buckets and objects granted through new public bucket or access point policies Off
- Block public and cross-account access to buckets and objects through any public bucket or access point policies Off

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AclForTenantGlue",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::{{ IE's AWS Account ID }}:root"
      },
      "Action": [
        "s3:Get*",
        "s3:Put*"
      ]
    }
  ]
}
```



```
"s3:ListBucket",
"s3:ListBucketMultipartUploads",
"s3:ListBucketVersions",
"s3:ListMultipartUploadParts"
],
"Resource": [
  "arn:aws:s3:::{{ Customer's bucket name }}",
  "arn:aws:s3:::{{ Customer's bucket name }}/*"
]
}
]
}
```

The items highlighted in bold text are:

1. IE's AWS Account Id: The AWS Account ID where the Intelligent Edge is installed. Find it in the admin side of your AWS account, which should be in this URL: <https://console.aws.amazon.com/billing/home?#/account>
2. Customer's bucket name: The name of the bucket containing the data. In the example below, "s3mediadata"

The image shows two screenshots from the AWS console. The top screenshot shows the 'Amazon S3' service page with a bucket named 's3mediadata' highlighted by a blue box. The bottom screenshot shows the 'Account Settings' page with the 'Account Id' field highlighted by a blue box, containing the value '734236616923'. Other fields shown include 'Seller: AWS Inc.', 'Account Name', and 'Password: *****'.

Once these JSON with the correct values have been added, you can test access using the command line from the machine where Intelligent Edge is installed. \$\$\$



```

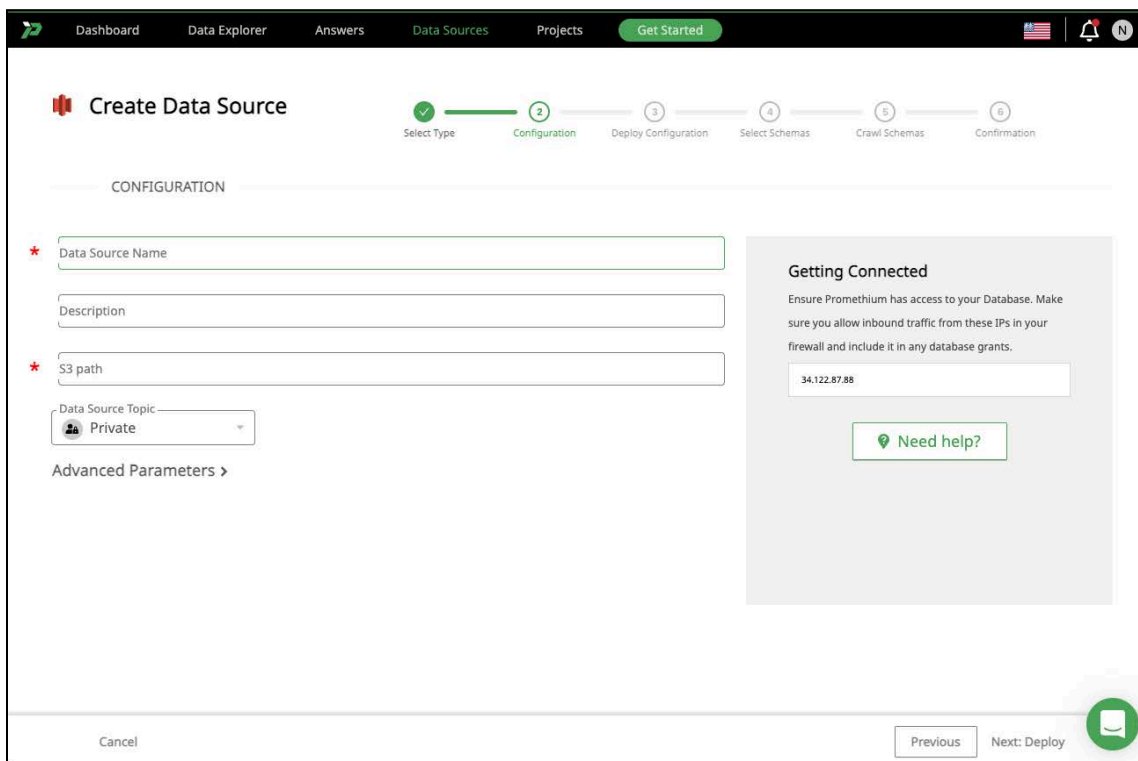
1. bash
Last login: Mon Aug 10 15:08:49 on ttys001
Antonios-MBP-2:~ antionielasnieves$ aws s3 ls s3://infawavedemo
2020-08-10 14:32:11          3992 Account.csv
Antonios-MBP-2:~ antionielasnieves$

```

An example is provided below.

Connect S3

Connecting to S3 is fast and easy. Simply click Data Sources from the menu, then click + New Data Source. Select S3 and enter the information for your AWS environment.



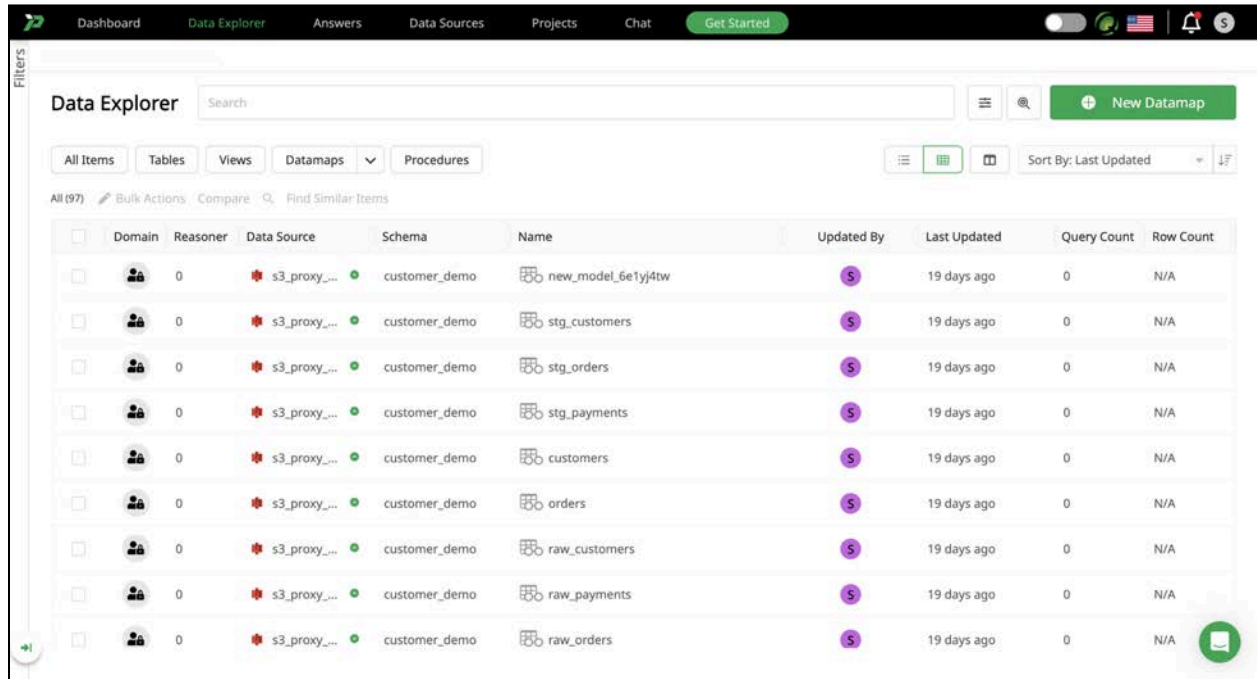
Trino connection details can be case sensitive. Ensure that data is entered with the correct case for your Trino environment.

Field	Description
Data Source Name	Give the AWS S3 connection a unique name for easy identification.
Description	Optional. A meaningful description of this connection.
S3 path	Enter the path to your s3 bucket.

Click the Connect button then select Schemas. After successful connection and schema selection Promethium will automatically crawl and catalog data in your S3 environment, then make the data easily searchable and usable.



Click Data Explorer from the menu and begin exploring the data in your S3 environment.



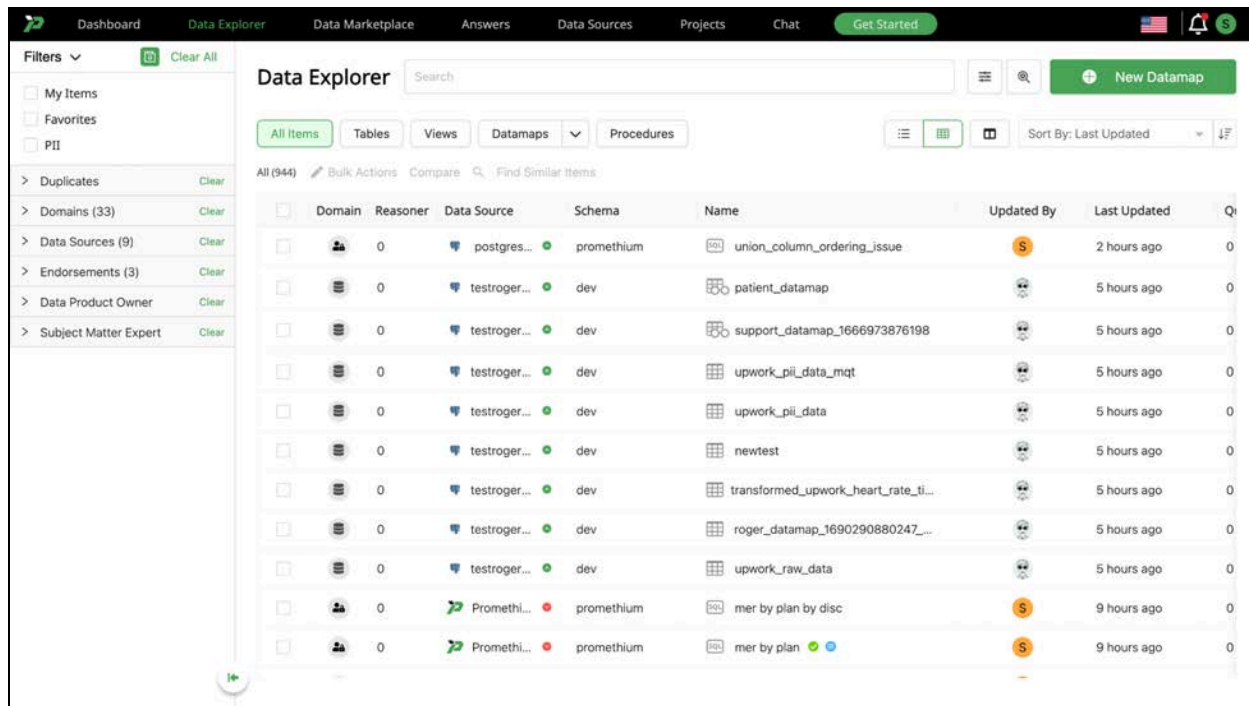
FAQs

- Is my data moved or copied?
 - No. Your data stays where it is and isn't moved to Promethium.
- What options are available for publishing datasets?
 - Datasets can be published as a view, as a table or as a dbt model.

Data Explorer

Data Explorer is the fast and easy way to explore data across distributed data sources.

Use Data Explorer to find, tag and preview data. Open Data Explorer from the menu. The list of data sources and tables is displayed.



A number of powerful result filters are present immediately above the table of results.

- **My Items** - Items that have been created by your user account.
- **Favorites** - Tables, views, and datamaps that are favorited marked by a yellow star.
- **Duplicates** - Any items that have been flagged as a copy.
- **Domains** - Selecting a specific domain or domains and the items associated.
- **Endorsements** - View items by endorsement statuses from “Endorsements,” “Warning,” and “Deprecation.”
- **Data Product Owner** - Filter items by specific data product owner.
- **Subject Matter Expert** - Filter by assigned SME.
- **All Items** - Default filter that shows all available tables, views, and datamaps.
- **Tables** - This filter selection only shows tables.
- **Views** - This filter selection only shows views.
- **Datamaps** - This filter shows only datamaps and can specify if the datamaps are “Published” or “In Progress” status.

You can combine one or all of these filters to rapidly narrow in on your required data.



Saved Filters

The screenshot shows the Data Explorer interface. A dropdown menu for 'Filters' is open, displaying a 'RECENTLY USED' section with three filters: 'tst', 'test tes ttes', and 'master filter sfg'. Below this is a 'View all filters' link. The main interface shows a table of data items with columns: Domain, Reasoner, Data Source, Schema, Name, Updated By, Last Updated, and Q. The table contains 14 rows of data, including datamaps and tables.

Domain	Reasoner	Data Source	Schema	Name	Updated By	Last Updated	Q
	0	postgres...	promethium	union_column_ordering_issue	S	2 hours ago	0
	0	testroger...	dev	patient_datamap		6 hours ago	0
	0	testroger...	dev	support_datamap_1666973876198		6 hours ago	0
	0	testroger...	dev	newtest		6 hours ago	0
	0	testroger...	dev	transformed_upwork_heart_rate_ti...		6 hours ago	0
	0	testroger...	dev	roger_datamap_169029088024_...		6 hours ago	0
	0	testroger...	dev	upwork_pii_data_mqt		6 hours ago	0
	0	testroger...	dev	upwork_pii_data		6 hours ago	0
	0	testroger...	dev	upwork_raw_data		6 hours ago	0
	0	Promethi...	promethium	mer by plan by disc	S	10 hours ago	0
	0	Promethi...	promethium	mer by plan	S	10 hours ago	0

Clicking the downward arrow next to Filters, you will be able to access and discover specific datamaps and tables promptly with the use of the recently used saved filters.

To create your own saved filter, click View all filters.



Name	Description	Actions
tst		Apply Filter Edit Delete
test tes ttes		Apply Filter Edit Delete
master filter sfg	sfgsdfg	Apply Filter Edit Delete
test tes test1	desc	Apply Filter Edit Delete
One more filter	my items + data sources(4)	Apply Filter Edit Delete
new filter 12345		Apply Filter Edit Delete
new filter after refactoring1	my favorites	Apply Filter Edit Delete
Check action 2		Apply Filter Edit Delete
my items and no duplicates [test for very long name: dataexplorer]	test descriptions	Apply Filter Edit Delete
datasources test	prom - 2	Apply Filter Edit Delete

In the Saved Filters page, you will see all the filters created or create one if you don't see one that is relevant.

Filters Clear All

- My Items
- Favorites
- PII

> Duplicates Clear

> Domains (33) Clear

> Data Sources (11) Clear

> Endorsements (3) Clear

> Data Product Owner Clear

> Subject Matter Expert Clear

Create filter

All Items Tables Views Datamaps Procedures

Filter Name
Name here

Search Bar Input
Input goes here

Description
Description Here

Submit Cancel

Like the filters on the Data Explorer page, selections made in the left side pane will be included in your saved filter.



Start by naming the new saved filter, specify which tables to search, and then provide a description of the expected results when the filter is applied.

Additionally, you can select if you only want to apply this filter to Tables, Views, Datamaps, and/or Procedures.

The screenshot shows the 'Create filter' interface. On the left, there is a 'Filters' sidebar with a 'Clear All' button and a list of filter categories: My Items, Favorites, PII, Duplicates, Domains (33), and Data Sources (11). Below these is a search bar for data sources and a list of data sources including Athena, BigQuery, GoogleSheets, MySQL, Oracle, PostgreSQL (checked), Promethium, and RedShift. The main area is titled 'Create filter' and has a tabbed interface with 'Tables' selected. It contains three input fields: 'Filter Name' (Patient information PostgreSQL), 'Search Bar Input' (patient), and 'Description' (Tables that contain patient information). At the bottom, there are 'Submit' and 'Cancel' buttons.

Click Submit when you are ready.



Filter has been saved successfully.

Saved Filters

Create, Edit, or Use Saved Search Parameters to save time on your Discovery process.

Search...

Name	Description	Actions
Patient information PostgreSQL	Tables that contain patient information	Apply Filter Edit Delete
tst		Apply Filter Edit Delete
test tes ttes		Apply Filter Edit Delete
master filter sfg	sfgsdfg	Apply Filter Edit Delete
test tes test1	desc	Apply Filter Edit Delete
One more filter	my Items + data sources(4)	Apply Filter Edit Delete
new filter 12345		Apply Filter Edit Delete
new filter after refactoring1	my favorites	Apply Filter Edit Delete
Check action 2		Apply Filter Edit Delete
my items and no duplicates [test for very long name: dataexplorer]	test descriptions	Apply Filter Edit Delete

Your newly created filter has been saved successfully, click Apply Filter and you will be taken back to the Data Explorer with results relevant to your saved filter.

Filters: Clear All

- My Items
- Favorites
- PII

> Duplicates Clear

> Domains (33) Clear

> Data Sources (9) Clear

> Endorsements (3) Clear

> Data Product Owner Clear

> Subject Matter Expert Clear

Data Explorer

patient

All Items Tables Views Datamaps Procedures

Sort By: Last Updated

All (18) Bulk Actions Compare Find Similar Items

	Domain	Reasoner	Data Source	Schema	Name	Updated By	Last Updated	QI
		1276	postgres...	customer_demo	patients	S	19 days ago	34
		1238	postgres...	customer_demo	patient_encounters	S	19 days ago	41
		1234	postgres...	customer_demo	patient_conditions	S	19 days ago	22
		1232	postgres...	customer_demo	patient_medications	S	19 days ago	7
		1131	postgres...	customer_demo	inpatient_discharge_personal	S	19 days ago	2
		1058	postgres...	customer_demo	encounters_1700671559471_mqt	S	19 days ago	16
		1056	postgres...	customer_demo	observations_1700671501744_...	S	19 days ago	16
		1056	postgres...	customer_demo	medications_1700671463864_...	S	19 days ago	16
		1048	postgres...	customer_demo	conditions_1700671408892_mqt	S	19 days ago	11
		1031	postgres...	customer_demo	claims_1700671197276_mqt	S	19 days ago	1
		931	postgres...	customer_demo	patients_demographics	S	19 days ago	0

You can use the search bar in conjunction with these filters, as desired, to search for data. Enter the search term in the search bar and click Search. The search results are displayed.



The screenshot shows the Data Explorer interface with a search bar containing 'agent'. Below the search bar, there are tabs for 'All Items', 'Tables', 'Views', 'Datamaps', and 'Procedures'. The 'All Items' tab is selected, showing a table of 277 items. The table has columns for Domain, Reasoner, Data Source, Schema, Name, Updated By, Last Updated, Query Count, and Row Count. The results are sorted by 'Last Updated'.

Domain	Reasoner	Data Source	Schema	Name	Updated By	Last Updated	Query Count	Row Count
	0	Promethi...	promethium	support_datamap_1698866380514	S	8 hours ago	7	N/A
	0	Promethi...	promethium	support_datamap_169886653069...	S	8 hours ago	0	N/A
	0	Promethi...	promethium	support_datamap_1698866027773	S	8 hours ago	8	N/A
	0	Promethi...	promethium	support_datamap_169886617273...	S	8 hours ago	1	N/A
	0	Promethi...	promethium	support_datamap_1698864580530	S	8 hours ago	15	N/A
	0	Promethi...	promethium	support_datamap_169886472673...	S	8 hours ago	1	N/A
	1800	postgresql	customer_demo	agents	S	8 hours ago	2	100
	900	postgresql	customer_demo	dim_agent	S	8 hours ago	180	281
	0	postgresql	customer_demo	quotes	S	8 hours ago	0	150

You can access our Advanced Search dialog by selecting the slider icon immediately to the right of the search box. This makes it easy to perform field specific text searches or even numeric count searches such as limiting the result set to those with a query count greater than or equal to 500.

The screenshot shows the Data Explorer interface with the Advanced Search dialog open. The dialog has a title bar 'Advanced Search' and a close button. It is divided into two sections: 'Text Search' and 'Count Search'. The 'Text Search' section has fields for Name, Tags, Field Tags, Fields, Description, and Updated By, each with a 'Contains' dropdown and an input field. The 'Count Search' section has fields for Tag Count, Field Count, and Query Count, each with a dropdown and an input field. The 'Query Count' field is set to 'Greater than or equal to' and '500'. There are 'Clear Advanced Search', 'Cancel', and 'Search' buttons at the bottom of the dialog.

To make the result set more readable, you can add or remove columns from the Data Explorer table with the column selector in the top right corner.



The screenshot shows the Data Explorer interface with a search for 'agent'. The table list includes columns for Domain, Reasoner, Data Source, Schema, Name, and Update. A filter menu is open, showing options like bulkEditorCheckbox, Domain, Vendor, Data Source, Catalog, Schema, Type, Name, Updated By, Last Updated, Reasoner Score, Queries, Row Count, Tags, Fields, and Location. The table data is as follows:

Domain	Reasoner	Data Source	Schema	Name	Update	Query Count	Row Count
	0	postgresql	customer_demo	fact_subscription_activity		783974	684
	900	redshift	customer_demo	dim_agent		772333	281
	0	oracle	CUSTOMER_DEMO	MASTER_PLAN_TABLE_WDATES		907935	1K
	0	oracle	CUSTOMER_DEMO	FACT_SUBSCRIPTION_ACTIVITY		38931	314
	0	redshift	customer_demo	fact_subscription_activity		5175	684
	0	sqlserver	customer_demo	fact_subscription_activity		5374	959
	900	mysql	customer_demo	dim_agent		67654	281
	0	mysql	customer_demo	fact_subscription_activity		4312	681
	900	bq	pm61data	dim_agent	4 days ago	3310	281

To find a similar table, begin by selecting a table and then clicking “Find Similar Items.” The “Parent Item” is the originally selected table where you wish to find similar items, you can select a maximum of 4 items. Once you are satisfied with the selections made, click “Compare.” This view facilitates side-by-side comparison of up to four items, providing insights into each item's column and row count, associated database, schema, description, endorsements status, item



type, classified domain/subdomain, and the latest update date.

The screenshot shows the Data Explorer interface with a search for 'orders'. A 'Compare Items (Up to 4)' dialog is open, displaying four items for comparison:

- orders (postgresql | customer_demo): Table | 9 Columns | 99 Rows | 0 Answers | Query count: 0 | Last Updated a day ago by S
- ORDERS (snowflake2 | CUSTOMER_DEMO): Table | 9 Columns | 99 Rows | 0 Answers | Query count: 0 | Last Updated 7 days ago by S
- orders (Max_redshift | customer_demo): Table | 9 Columns | 99 Rows | 0 Answers | Query count: 0 | Last Updated a month ago by S
- orders (postgres2 | customer_demo): Table | 9 Columns | 99 Rows | 0 Answers | Query count: 0 | Last Updated 4 months ago by S

The 'Compare Items' dialog provides a detailed comparison of the four items:

Item Details	orders (postgresql customer_demo)	ORDERS (snowflake2 CUSTOMER_DEMO)	orders (Max_redshift customer_demo)	orders (postgres2 customer_demo)
Type	TABLE	TABLE	TABLE	TABLE
Domain	Private	Private	Private	Private
Subdomain	-	-	-	-
Row #	99	99	99	99
Column #	9	9	9	9
Queries	0	0	0	0
Last Update	Last Updated a day ago by S	Last Updated 7 days ago by S	Last Updated a month ago by S	Last Updated 4 months ago by S



The screenshot shows the Data Explorer interface with a search for 'order'. The left sidebar contains filters for 'My Items', 'Favorites', 'Duplicates', 'Domains (9)', 'Data Sources (18)', 'Endorsements (3)', 'Data Product Owner', and 'Subject Matter Expert'. The main area displays a table of results under the heading 'Data Explorer' with a search bar containing 'order'. Below the search bar are tabs for 'All Items', 'Tables', 'Views', 'Datamaps', and 'Procedures'. A 'Sort By: Reasoner Score' dropdown is visible. The table shows one 'Parent Item' and several 'Similar Items'.

Domain	Reasoner	Data Source	Schema	Name	Updated By	Last Updated	Q
1662	postgres...	customer_demo	orders		S	a day ago	0

Domain	Reasoner	Data Source	Schema	Name	Updated By	Last Updated	Q
100	postgres2	customer_demo	orders		S	4 months ago	0
100	sqlserver	customer_demo	orders		S	3 months ago	0
100	sqlserver	cdmDatabaseSc...	orders		S	2 months ago	0
100	snowflake	CUSTOMER_DE...	ORDERS		S	3 months ago	0
100	redshift	customer_demo	orders		S	a month ago	0
100	Max_red...	customer_demo	orders		S	a month ago	0

Filter and Sort Tables

To help find and explore the data available you can select from a number of Sort By refiners from the drop down at the top right. The button to the right of this selection field will invert the selected sort. e.g. Sort by Name from A-Z or Z-A.

The screenshot shows the Data Explorer interface with a search for 'agent'. The left sidebar contains filters for 'My Items', 'Favorites', 'Duplicates', 'Domains (9)', 'Data Sources (18)', 'Endorsements (3)', 'Data Product Owner', and 'Subject Matter Expert'. The main area displays a table of results under the heading 'Data Explorer' with a search bar containing 'agent'. Below the search bar are tabs for 'All Items', 'Tables', 'Views', 'Datamaps', and 'Procedures'. A 'Sort By: Last Updated' dropdown menu is open, showing options: 'Sort By: Reasoner Score', 'Sort By: Name', 'Sort By: Last Updated', 'Sort By: Updated By', 'Sort By: Query Count', 'Sort By: Field Count', 'Sort By: Row Count', and 'Sort By: Tag Count'. The table shows several results.

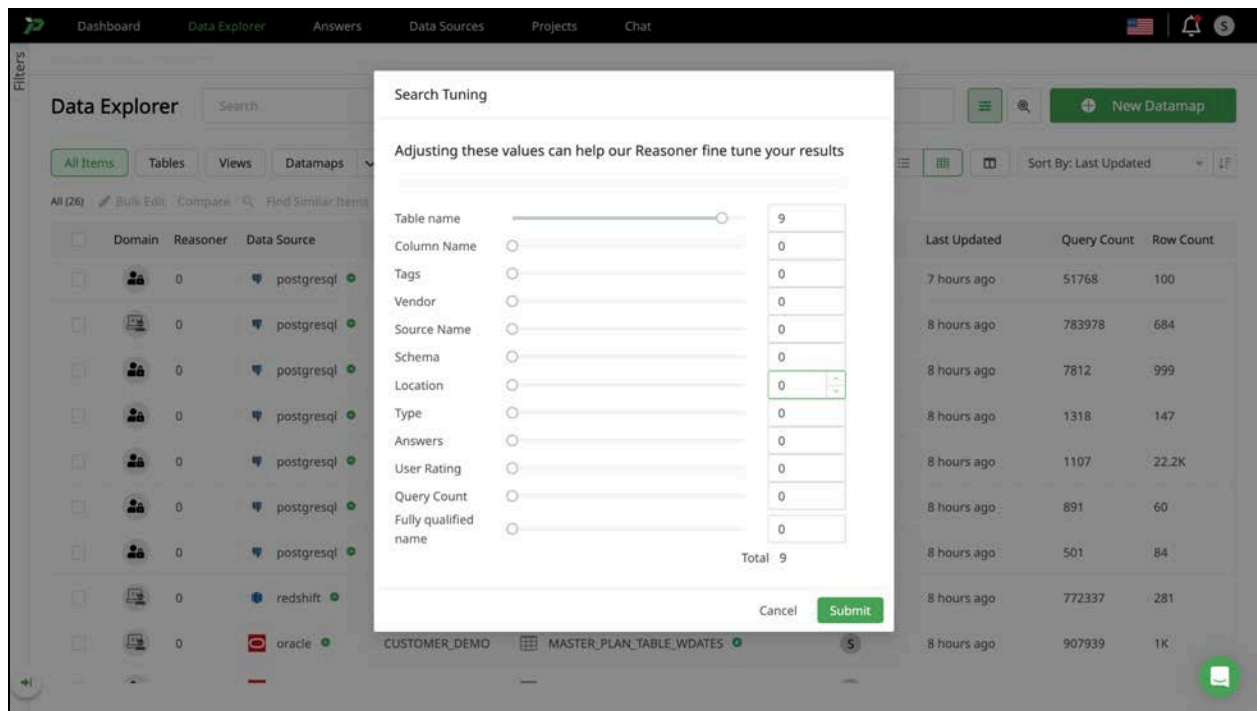
Domain	Reasoner	Data Source	Schema	Name	Updated By	Last Updated	Q
0	postgresql	customer_demo	fact_subscription_activity		S	8 hours ago	
900	redshift	customer_demo	dim_agent		S	8 hours ago	
0	oracle	CUSTOMER_DEMO	MASTER_PLAN_TABLE_WDATES		S	8 hours ago	
0	oracle	CUSTOMER_DEMO	FACT_SUBSCRIPTION_ACTIVITY		S	4 days ago	38931 314
0	redshift	customer_demo	fact_subscription_activity		S	4 days ago	5175 684
0	sqlserver	customer_demo	fact_subscription_activity		S	4 days ago	5374 959
900	mysql	customer_demo	dim_agent		S	4 days ago	67654 281
0	mysql	customer_demo	fact_subscription_activity		S	4 days ago	4312 681
900	bq	pm61data	dim_agent		S	4 days ago	3310 281



Feature	Description
Reasoner™ Score	Reasoner™ uses machine intelligence to recommend the best match
Name	To help locate specific Table, Datamap or Views
Last Updated	Sort by the date of the last update
Updated By	Sort by the username whom last updated
Query Count	Sort by the query volume (e.g. how popular or frequently a table, datamap or view is queried)
Field Count	Sort by the number of fields
Tag Count	Sort by the number of tags. The count of tags is an indicator of data usage and trust in the underlying data.

Reasoner

By default, the result table is sorted by Reasoner™ score. The Reasoner score is tunable by you at any time. Simply click on the magnifying glass icon to the right of the search bar to bring up the Search Tuning dialog.



Adjusting the sliders or numeric values will alter the weight applied to the concept e.g. a value of 10 for Table Name will place more importance on the name of a table when searching for matches to your query. Tables with names that match your query will have a higher Reasoner score than those with a tag matching your query (assuming the tag value is set lower than 10).



Preview Data

To explore the metadata on a table, datamap or view, and to preview data and columns, mouse over a row of interest. The row will be highlighted indicating that a click anywhere on the row will open a detailed view.

The screenshot shows the Data Explorer interface with a table of data items. The table has the following columns: Domain, Reasoner, Data Source, Schema, Name, Updated By, Last Updated, Query Count, and Row Count. The data is sorted by Last Updated.

Domain	Reasoner	Data Source	Schema	Name	Updated By	Last Updated	Query Count	Row Count
	0	postgresql	customer_demo	customers	S	7 hours ago	51768	100
	0	postgresql	customer_demo	fact_subscription_activity	S	8 hours ago	783978	684
	0	postgresql	customer_demo	dim_channel	S	8 hours ago	7812	999
	0	postgresql	customer_demo	caller	S	8 hours ago	1318	147
	0	postgresql	customer_demo	Items	S	8 hours ago	1107	22.2K
	0	postgresql	customer_demo	stores	S	8 hours ago	891	60
	0	postgresql	customer_demo	item_categories	S	8 hours ago	501	84
	0	redshift	customer_demo	dim_agent	S	8 hours ago	772337	281
	0	oracle	CUSTOMER_DEMO	MASTER_PLAN_TABLE_WDATES	S	8 hours ago	907939	1K

The detail view is grouped into six navigational areas:

Overview

The overview screen details some key information about the data.

- Data Source - e.g. Oracle. Snowflake etc.
- Schema - The schema name
- Type - The type of data being explored e.g. Table, View or Datamap
- Columns - The number of columns in the data
- Rows - The number of rows that will be automatically populated in the preview, the default is 100 rows if data contains over 100 rows. Preview limits can be changed from 100 to 200, 500, and 1,000.

Immediately under these data facts is the fully qualified name of the data followed by a rich text Description wiki editor where users can enter information useful to themselves and others in the future. The description, tags, and even who last updated this information are metadata that is used by the Reasoner and our advanced text search.



The screenshot shows the Data Explorer interface. The main table lists datamaps with columns for Source, Schema, Name, Updated By, Last Updated, and Q. The overview panel on the right shows details for a selected datamap, including Data Source (postgrestd), Schema (customer_demo_publish), Type (Table), Columns (11 columns), and Rows (0 rows). It also includes fields for Domain (Private), Subdomain, and Tags.

On the right of the overview are quick visual summaries of data represented in detail on the other tabs e.g. how many up or downvotes have been received. Here you can also assign a Domain and subdomain and add or remove tags.

The screenshot shows the Data Explorer interface with the Access Control panel open. The panel displays Domain Access Rules and a table of User/role_name and Access. The Access column shows 'viewer' for several users.

User/role_name	Access
palwinder+1234@pm61data.com	
enrique+su1@pm61data.com	
enrique+testtrino6@pm61data.com	
avinokurov@itzap.com	
enrique@pm61data.com	
ravi@pm61data.com	
c-roger@pm61data.com	
enrique+datasteward2@pm61data.com	
enrique+ssoonly@pm61data.com	
support@pm61data.com	
enrique+automation@pm61data.com	
chakra+test2@pm61data.com	
enrique+testtrino5@pm61data.com	
c-roger+test23@pm61data.com	
enrique+testtrino4@pm61data.com	viewer
enrique+su4@pm61data.com	
enrique+testtrino3@pm61data.com	
enrique+su3@pm61data.com	
chakra+test1@pm61data.com	

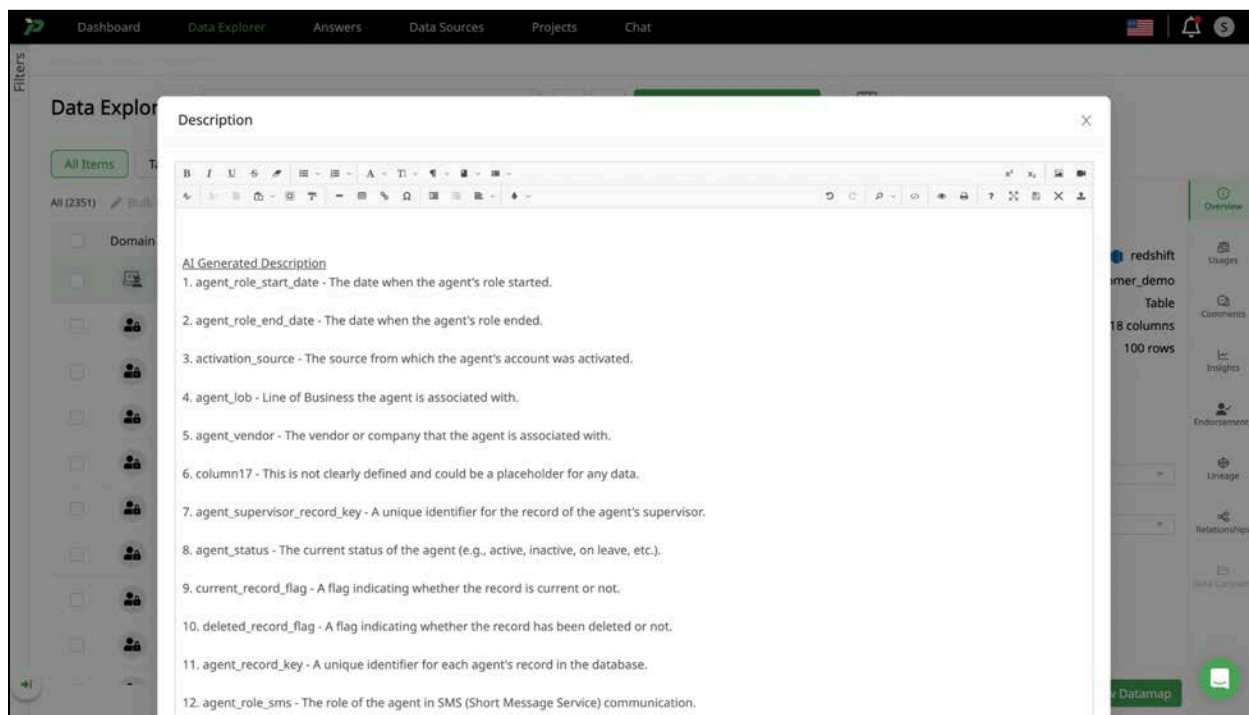


The Access Control tab shows who has access to the data product by user or role, based on assigned Domain.

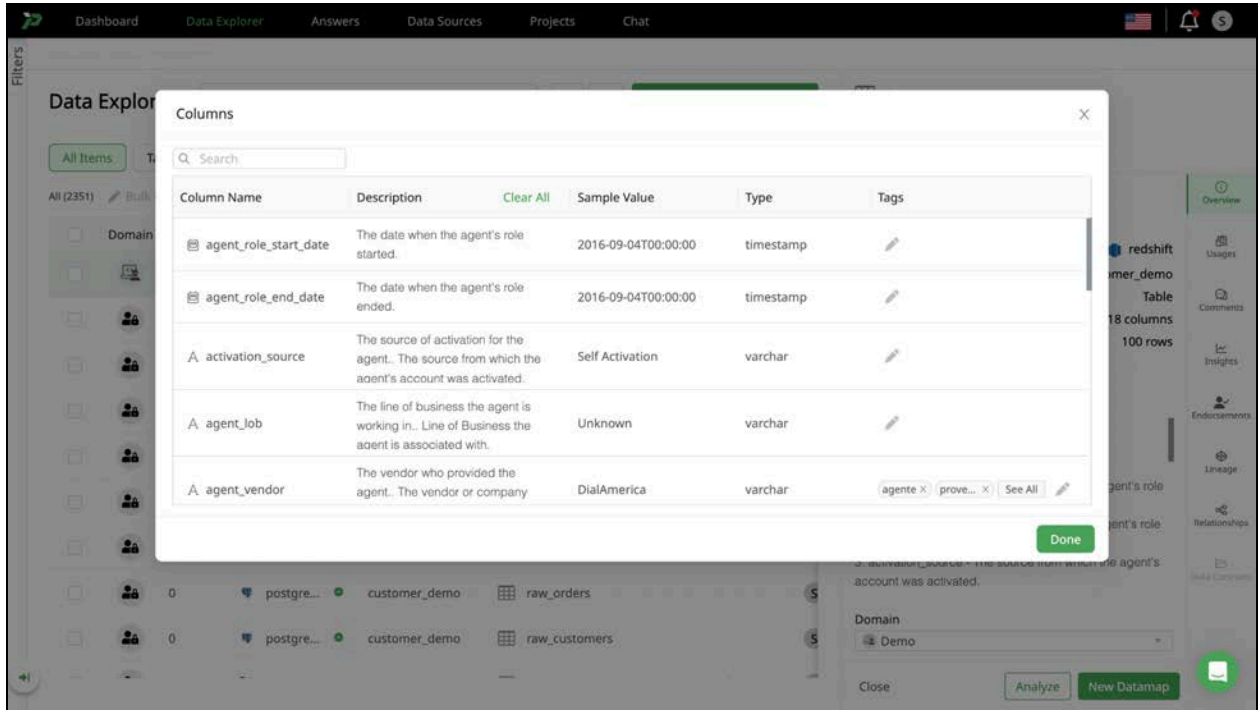
Columns

The Columns tab details each column within the data.

- **Column Name** - The column name is displayed along with an indication of the data type within e.g. A for varchar, #.# for decimal. The details of the data type are located in the Type column of this table.
- **Description** - As with the description wiki editor on the Overview tab, this field is a rich user-editable wiki controlled field. Hovering over the area will display an edit pencil icon. Click on the icon to load the wiki editor. Changes are not saved automatically, it is important to click on the save icon (floppy disk) when you are satisfied with your modifications. Any description here is included within the text search and the Reasoner score. Additionally, there's the functionality to generate descriptions using AI capabilities by clicking "Add AI Description." The description will be generated automatically and note that it was AI generated.



- **Sample Value** - Randomly returned sample data from a row within the column.
- **Type** - The data type represented by the column.
- **Tags** - Tags can be added to columns, tables, and questions. These tags are column specific and are not displayed on the Overview tab. They are part of the free text search and Reasoner score. Tagging enhances the search experience and adds context with the use of synonyms.



Samples

The Samples tab details a limited results set pulled live from the data source.

The default sample size is 100 rows. You can modify this to a different, preset value (100,200,500,1000) by selecting the Limit Rows carat at the top left of the screen.

You can search the returned data using the search box at the top right. Any column with matching information will be highlighted. Note: This search is limited to the returned sample data, not the full result set.



Dashboard Data Explorer Answers Data Sources Projects Chat

dim_agent support@pm61data.com • Nov 1, 2023

Overview Columns (18) Samples (100) Usages (49) Insights Endorsements Comments Lineage

Limit Rows: 100 Filter Null 100 rows, 85 KB fetched and 327 ms runtime Search Filters

agent_role...	agent_role...	activation_s...	agent_job	agent_vendor	column17	agent_super...	agent_status	current_rec...	deleted_rec...	agent_recor...	agent_role
2016-09-04T00:0	2016-09-04T00:0	Self Activation	Unknown	DialAmerica		-1	I	N	N	1	Agent
2016-09-15T00:0	2016-09-15T00:0	Self Activation	Unknown	DialAmerica		-1	I	N	N	2	Agent
2016-09-10T00:0	2016-09-10T00:0	Self Activation	Unknown	DialAmerica		-1	I	N	N	3	Agent
2016-09-24T00:0	2016-09-24T00:0	Self Activation	Unknown	DialAmerica		-1	I	N	N	4	Agent
2016-09-07T00:0	2016-09-07T00:0	Self Activation	Unknown	TeleServicesDire		-1	I	N	N	5	Agent
2016-08-30T00:0	2016-08-30T00:0	Self Activation	Unknown	TeleServicesDire		-1	I	N	N	6	Agent
2016-09-03T00:0	2016-09-03T00:0	Self Activation	Outbound Sales	DialAmerica		-1	I	N	N	7	Agent
2016-09-01T00:0	2016-09-01T00:0	Outbound Sales	INDIA	DialAmerica	N	-1	I	Self Activation	N	8	Agent
2016-09-08T00:0	2016-09-08T00:0	Self Activation	Unknown	DialAmerica		-1	I	N	N	9	Agent
2016-09-25T00:0	1999-12-31T00:0	Self Activation	Unknown	DialAmerica		-1	I	Y	N	10	Agent
2018-02-22T00:0	2018-02-24T00:0	Self Activation	Service	Alorica		1837853	I	N	N	11	SavesAgent
2016-09-02T00:0	2016-09-02T00:0	Self Activation	Unknown	Results		-1	I	N	N	12	Agent
2017-03-30T00:0	1999-12-31T00:0	Self Activation	Outbound Sales	DialAmerica		-1	A	Y	N	13	SalesOutbo
2016-09-14T00:0	2016-09-14T00:0	Unknown	INDIA	DialAmerica	N	-1	I	Self Activation	N	14	Agent
2015-11-09T00:0	1999-12-31T00:0	Self Activation	Unknown	DialAmerica		-1	I	Y	N	15	SavesAgent
2016-09-09T00:0	2016-09-09T00:0	Self Activation	Unknown	DialAmerica		-1	I	N	N	16	Agent
2016-09-16T00:0	2016-09-20T00:0	Self Activation	Unknown	Convergys		-1	I	N	N	17	Agent
2016-09-16T00:0	2016-09-20T00:0	Self Activation	Unknown	TeleServicesDire		-1	I	N	N	18	Agent
2016-08-31T00:0	2016-08-31T00:0	Self Activation	Unknown	DialAmerica		-1	I	N	N	19	Agent
2016-09-04T00:0	2016-09-04T00:0	Self Activation	Unknown	DialAmerica		-1	I	N	N	20	Agent
2016-09-22T00:0	2016-09-22T00:0	Self Activation	Unknown	Alorica		-1	I	N	N	21	Ann

Usage

The Usage tab details where the data is being used, and how.

Dashboard Data Explorer Answers Data Sources Projects Chat

dim_agent support@pm61data.com • Nov 1, 2023

Overview Columns (18) Samples (100) Usages (49) Insights Endorsements Comments Lineage

Datamaps (49)

Domain	Data Source	Schema	Name	Last Updated	Query #	Row #	User Interactions
	Promethium	promethium	support_datamap_1698558726647	4 days ago	S	0	0
	redshift	promethium	support_datamap_1698566911209	4 days ago	S	0	0
	redshift	promethium	support_datamap_1698567039067	4 days ago	S	0	0
	redshift	promethium	support_datamap_1698572920123	4 days ago	S	0	0
	Promethium	promethium	revenue_subscription_by_plan_by_vendor_from_20...	3 days ago	P	0	0
	Promethium	promethium	support_datamap_1698623084873	3 days ago	S	0	0
	Promethium	promethium	support_datamap_1698623366725	3 days ago	S	0	0
	Promethium	promethium	support_datamap_1698680826930	3 days ago	S	0	0
	Promethium	promethium	support_datamap_1698682359958	2 days ago	S	0	0
	Promethium	promethium	revenue by vendor by term by plan	2 days ago	P	0	0
	Promethium	promethium	support_datamap_1698683019270	2 days ago	E	0	0

The screen is grouped by Datamaps and Answers and lists how many of each concept utilize the data.



Clicking on any of the linked Datamaps or Answers will open that item in a new window. No modifications can be made directly from this view, it simply links you to the underlying usage where modifications can be made.

The screenshot shows a modal dialog titled "Add External Dashboard" with the following fields:

- Dashboard Name:
- Dashboard Public Url:
- Dashboard Type: (dropdown menu)

The dropdown menu is open, showing the following options:

- Tableau
- Power BI
- Maximo

The background interface shows a table with columns: Domain, Title, Status, Last Updated, Requested By, User Interactions. The table is currently empty, displaying "No Data".

The screenshot shows the "Usages" tab selected in the interface. The table below is empty, displaying "No Data".

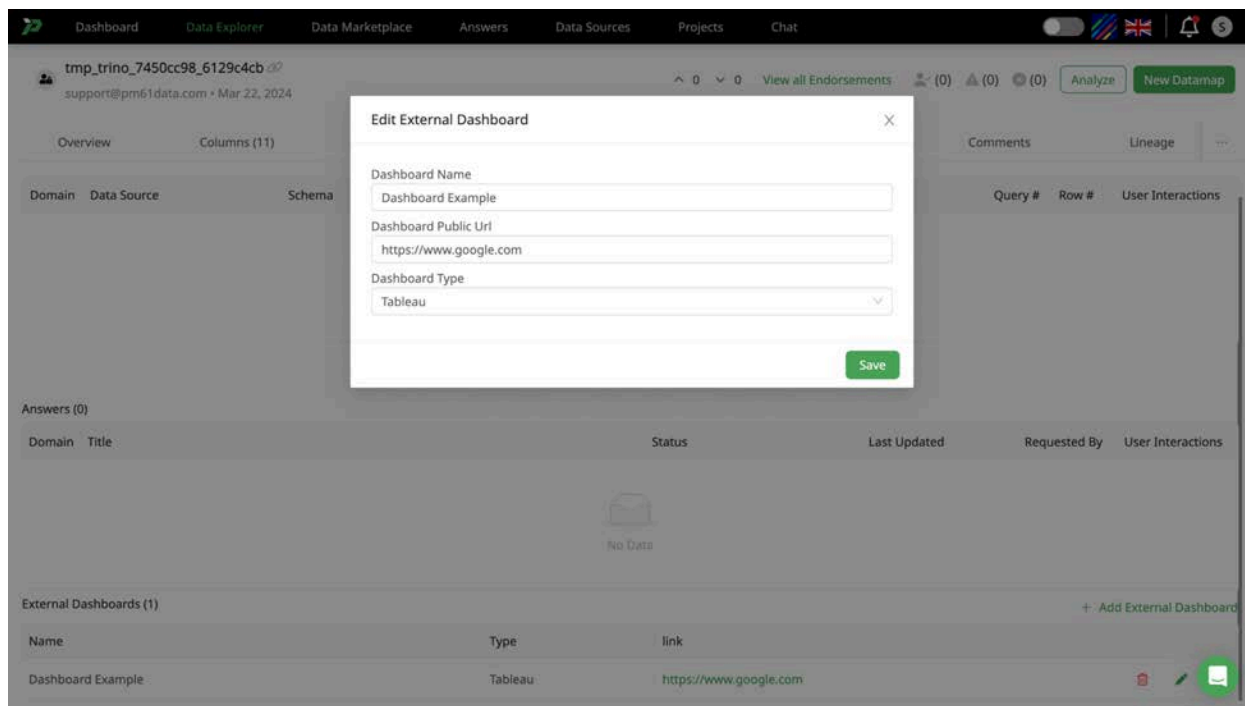
Domain	Title	Status	Last Updated	Requested By	User Interactions
No Data					

External Dashboards (1)

Name	Type	link
Dashboard Example	Tableau	https://www.google.com



Now, users can gain a clearer insight into the downstream implications of a Data Product through the integration of External Dashboards, including options like Tableau, Maximo, and PowerBI.



To edit External Dashboard usage entry, click the pencil icon on the right to make any changes and then Save.

Comments

The Comments tab shows all user comments by name and date. The comments are ordered from oldest to newest.

Comments can be added by entering the comment in the text field and clicking Add comment. Comments can be deleted or edited by selecting the appropriate icon next to an existing comment.

Comments are not included in our Advanced Search at this time.



The screenshot shows the 'Comments' tab for a data map named 'dim_agent'. The interface includes a navigation bar with 'Dashboard', 'Data Explorer', 'Answers', 'Data Sources', 'Projects', and 'Chat'. Below the navigation bar, there are tabs for 'Overview', 'Columns (18)', 'Samples (100)', 'Usages (49)', 'Insights', 'Endorsements', 'Comments', and 'Lineage'. The 'Comments' tab is active, displaying a large empty area with a 'No Comments' message and a comment input field at the bottom. The input field contains the text 'Add a Comment' and a 'Add comment' button.

Lineage

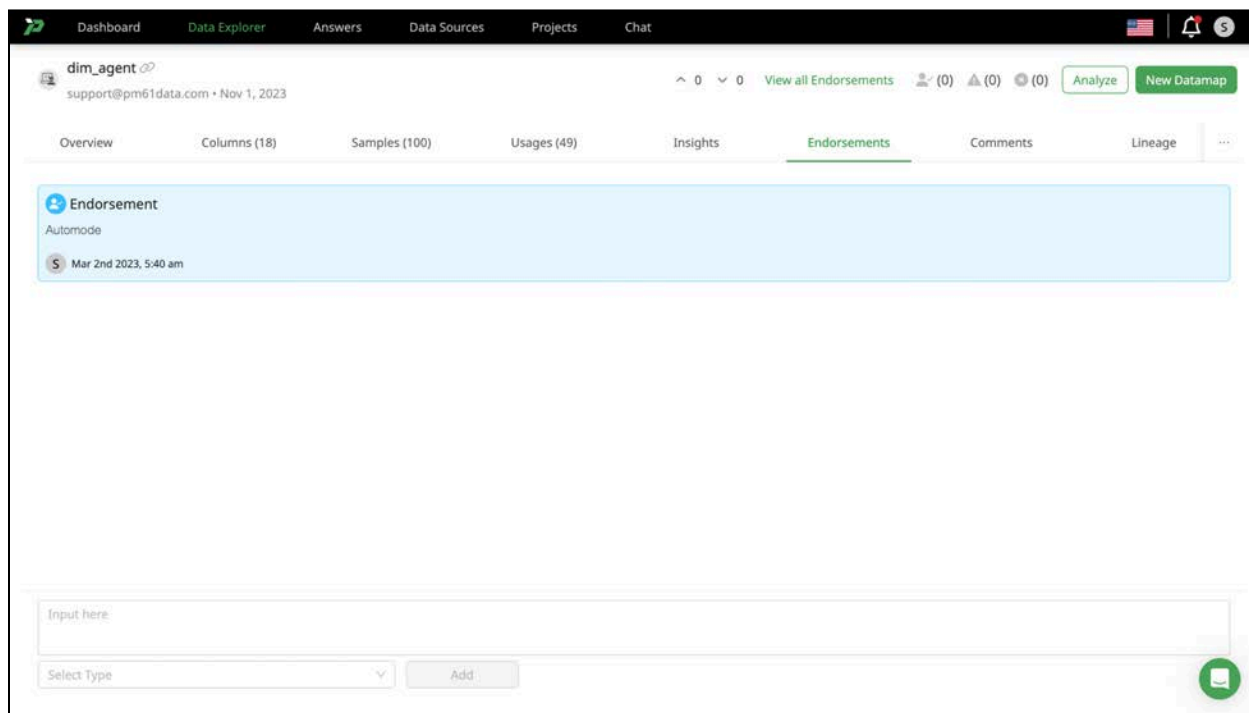
The Lineage feature displays the tables underlying the datamap. By expanding the '+' symbol, you gain access to detailed information, including utilized columns, joins, group bys, unions, and more. By hovering over a table name, you will be able to see the table's data source and schema information.

The screenshot shows the 'Lineage' feature for a query titled 'total revenue by vendor by plan in 2016'. The interface includes a navigation bar with 'Dashboard', 'Data Explorer', 'Answers', 'Data Sources', 'Projects', and 'Chat'. Below the navigation bar, there are tabs for 'Run', 'Save', 'Republish', 'Build Answer', and 'data map out of sync'. The 'Lineage' feature is active, displaying a query flow diagram. The diagram shows a sequence of operations: 'FACT_SUBSCRIPTION...', 'INNER JOIN', 'INNER JOIN', 'FILTER', 'GROUP BY', and 'TOTAL REVENUE BY V...'. Each operation is represented by a box with a '+' symbol. The 'TOTAL REVENUE BY V...' box is expanded, showing a table with columns 'VENDOR', 'PLANNAME', and 'REVENUE'. The 'Data Source: Promethium Schema: default' is also visible. The interface includes a 'Show Nested Lineage' toggle and a 'Run the SQL query to get your results' button at the bottom.



Endorsement

The Endorsement tab shows all user endorsements, warnings, or notes on deprecation.



Endorsements are ordered from oldest to newest and include the username and any comment they included along with their endorsement.

Datamaps

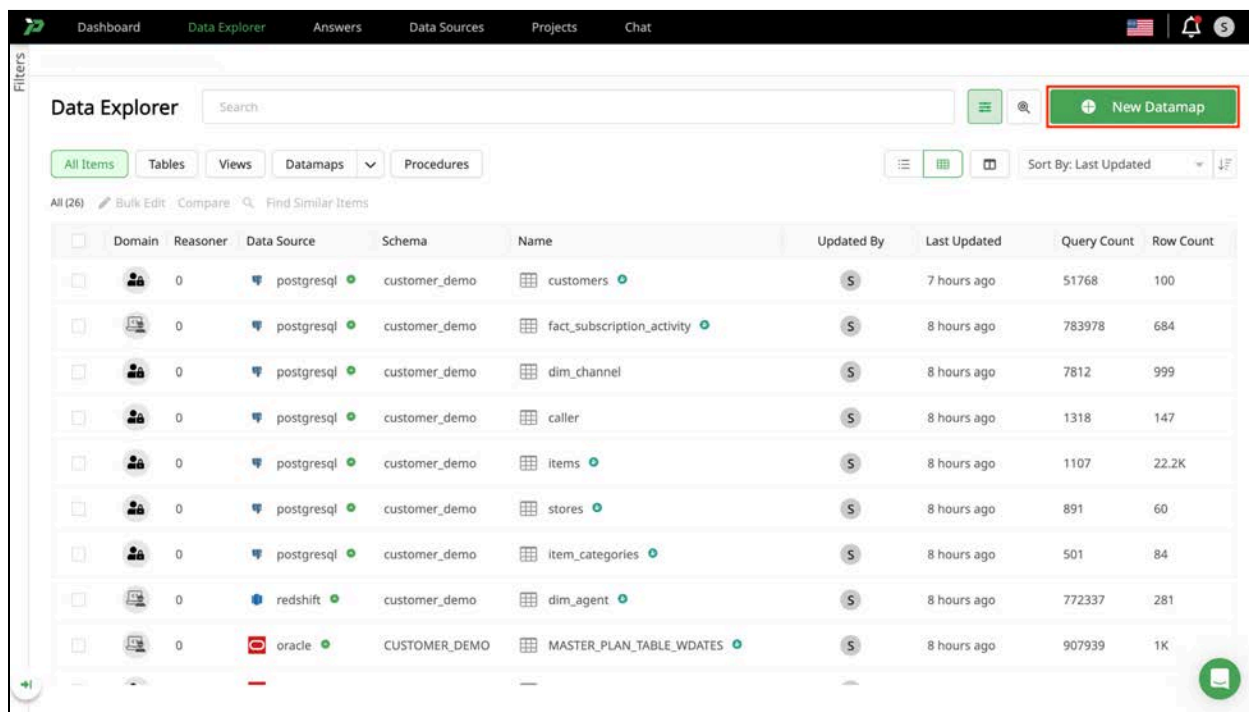
A datamap is, in essence, an augmented dataset. It is the term used to describe a dataset and all the additional metadata and user enrichment available through Promethium.

The key steps for building a Datamap include:

- Add tables, and add and edit joins
- Select and rename columns
- Add filter, Group By, and Order By

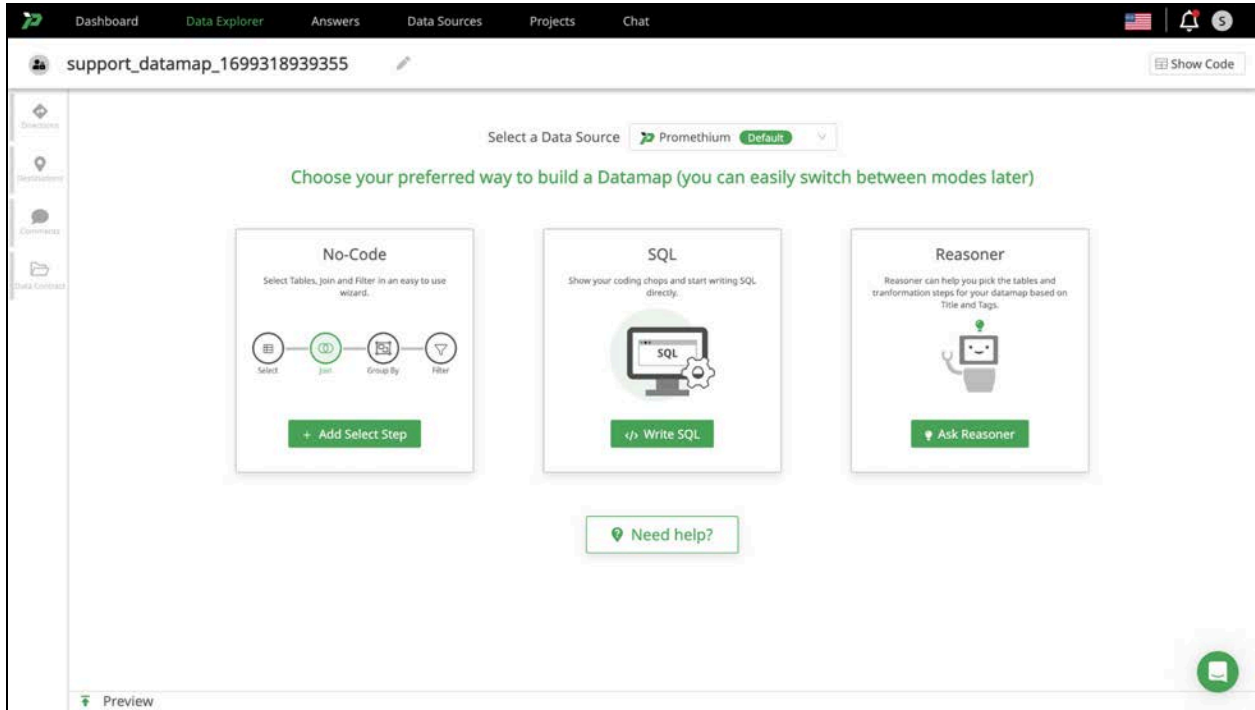


From Data Explorer, click + New Datamap.



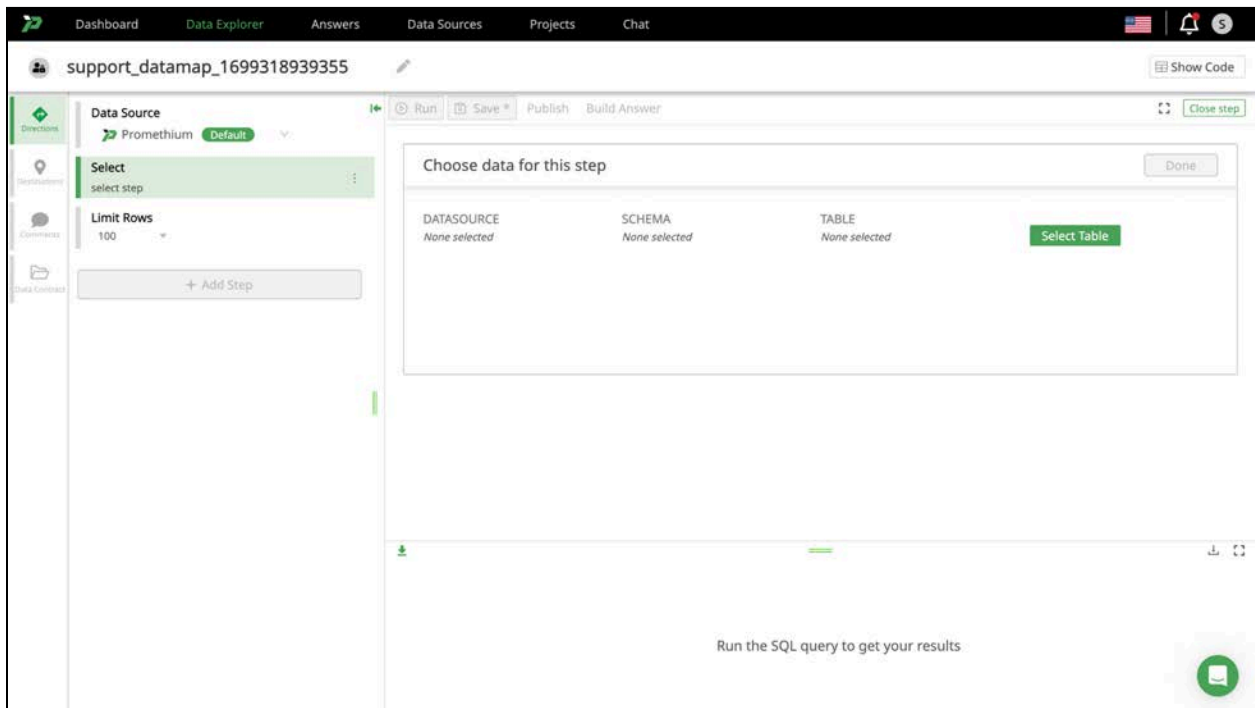
You can select one of three different methods to build a datamap. It's possible to switch between these modes as you go.

- No-Code - a wizard based flow intended for those with little to no SQL knowledge.
- SQL - a full web based IDE that allows you to write unrestricted SQL code directly.
- Reasoner - an AI powered builder that will make recommendations on tables and transformations based on the title and description you provide.



No-Code

Click + Add Select Step to enter the No-Code datamap creation flow.

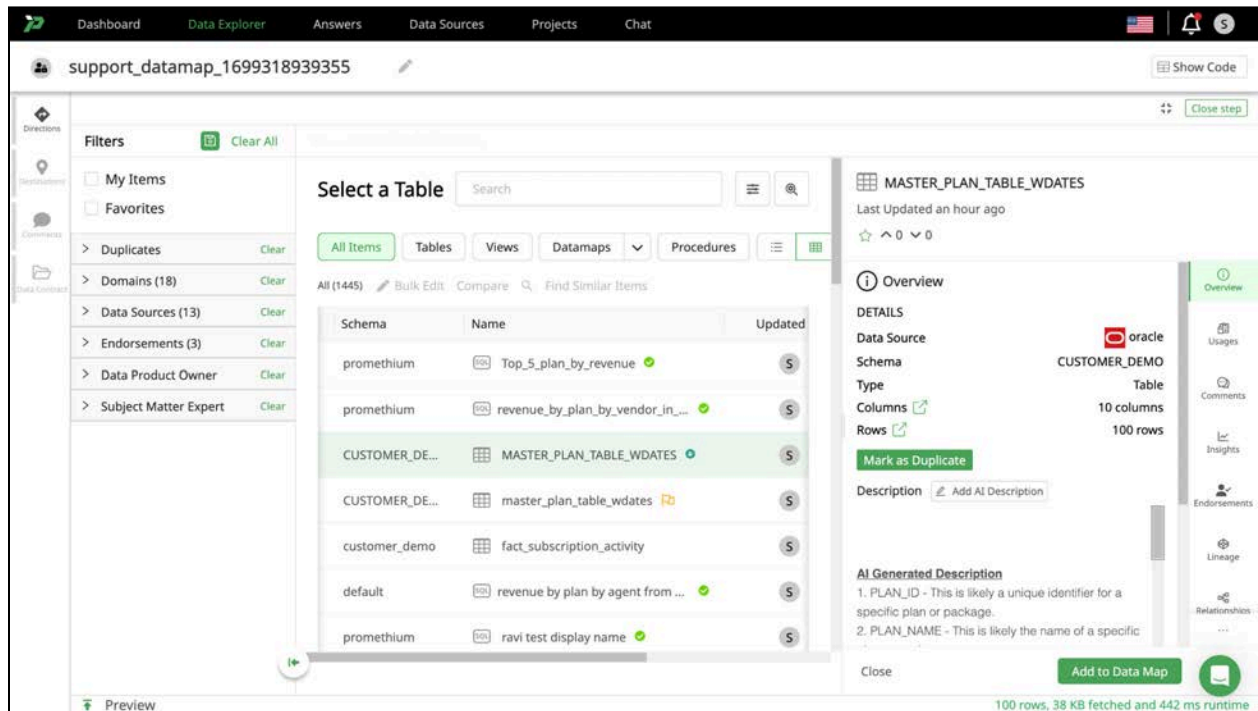


A default Domain and title will be assigned. You can modify these now or accept the defaults and modify them at a later time. You can add tags that will be used by the advanced search to help users locate your datamap.



To swap to the SQL mode click on Show Code at the top right of the page. You can swap back to the no-code flow as required.

The first task is to select a data source, schema, and table. Click Select Table. The Direction tab will display a tabular layout almost identical to the Data Explorer page. From here you can use search and filters to locate and select the table, view, or even existing datamap you want to add to your datamap.



Click the object you want to add and from the modal click Add to Datamap.



DATASOURCE	SCHEMA	TABLE
oracle	CUSTOMER_DEMO	MASTER_PLAN_TABLE_W...

Now we can select the field or fields we are interested in from this table. Click the green button labeled Select Fields Next >.

Column Name	Tags	Description	Expression	Type	Alias	PII
<input checked="" type="checkbox"/> ## PLAN_ID	planid	This is likely a unique identifier for a specific package.		Long	PLAN_ID	<input type="checkbox"/>
<input checked="" type="checkbox"/> A PLAN_NAME	3 tags	This is likely the name of a specific plan or package.		String	PLAN_NAME	<input type="checkbox"/>
<input checked="" type="checkbox"/> ## MER	2 tags	This could stand for 'Management Expense Ratio', a term		Long	MER	<input type="checkbox"/>
<input checked="" type="checkbox"/> A DISC		This could be an abbreviation for 'Discount'. It might		String	DISC	<input type="checkbox"/>
<input checked="" type="checkbox"/> A PACKAGE	2 tags	This could refer to the specific package or bundle of		String	PACKAGE	<input type="checkbox"/>
<input checked="" type="checkbox"/> A LEN	2 tags	This could stand for 'Length'. It might refer to the duration of a		String	LEN	<input type="checkbox"/>
<input checked="" type="checkbox"/> A ROLLUP1		This is likely a specific type of data		String	ROLLUP1	<input type="checkbox"/>

Select one or more fields to add to the datamap by clicking the check box to the left of the row(s) of interest. At this time you can modify the alias of the field for easier identification, flag the field as containing PII, choose to modify the datatype (e.g. cast from date to string), add specific tags or a description, and even create a custom expression.



When you have made your selections and modifications select Done.

Dashboard Data Explorer Answers Data Sources Projects Chat

support_datamap_1699318939355 Show Code

Run Save * Publish Build Answer Close step

Data Source Promethium Default

Select MASTER_PLAN_TABLE_WDATES

Limit Rows 100 + Add Step

Source DATASOURCE SCHEMA TABLE
oracle CUSTOMER_DEMO MASTER_PLAN_TABLE_WDATES Edit

Fields 10 Selected Fields Edit

Filters Add Filters (Optional) Edit

Preview 100 rows, 38 KB fetched and 442 ms runtime

Lastly, we can add any optional filters. Filters can be combined as required. Select a column from the drop down list, select an operator, and finally enter any value as required. When you're happy with your filter click the check mark next to your rule to apply it.

Dashboard Data Explorer Answers Data Sources Projects Chat

support_datamap_1699318939355 Show Code

Run Save * Publish Build Answer Close step

Data Source Promethium Default

Select MASTER_PLAN_TABLE_WDATES

Limit Rows 100 + Add Step

Add filter conditions Done

Column Name	Operator	Value
STARTDATE	is between	2015-01-01 → 2018-12-31

Add Filter

<< < Nov 2018 Dec 2018 >> >>

Su	Mo	Tu	We	Th	Fr	Sa	Su	Mo	Tu	We	Th	Fr	Sa
28	29	30	31	1	2	3	25	26	27	28	29	30	1
4	5	6	7	8	9	10	2	3	4	5	6	7	8
11	12	13	14	15	16	17	9	10	11	12	13	14	15
18	19	20	21	22	23	24	16	17	18	19	20	21	22
25	26	27	28	29	30	1	23	24	25	26	27	28	29
2	3	4	5	6	7	8	30	31	1	2	3	4	5

Preview 100 rows, 38 KB fetched and 442 ms runtime



When you are satisfied with your selections and filters, click the green Run button at the top of the screen to execute your datamap. This will execute the query and render a row limited result set to allow you to confirm the output.

The screenshot shows the 'Data Explorer' interface for a datamap named 'support_datamap_1699318939355'. The configuration is as follows:

- Data Source:** Promethium (Default)
- Select:** MASTER_PLAN_TABLE_WDATES
- Limit Rows:** 100
- Source:** DATASOURCE: oracle, SCHEMA: CUSTOMER_DEMO, TABLE: MASTER_PLAN_TABLE_WDATES
- Fields:** 10 Selected Fields
- Filters:** STARTDATE is between 2015-01-01, 2018-12-31

The 'Run' button is highlighted in green. The status bar at the bottom indicates '100 rows, 38 KB fetched and 442 ms runtime'.

If you are satisfied with the output we must save the datamap to make it available for use. Click the Save button next to the Run button previously clicked.

The screenshot shows the same 'Data Explorer' interface, but now the 'Save' button is highlighted in green. Below the configuration, a table of results is displayed:

PLAN_ID	PLAN_NAME	MER	DISC	PACKAGE
249	Promo - All Access - 12mo - 132.00 - (+\$1.00 per month)	11	Discount	All Acces
252	Fixed Demo - Select - 6mo - (Used Car)	0	TBD	Select
253	NavTraffic + NavWeather - 1mo - wActv	6.66	TBD	TBD
254	Marine Offshore - 6mo - w	6.66	TBD	TBD

The status bar at the bottom indicates '100 rows, 37 KB fetched and 970 ms runtime'.



Custom Expressions - you can create custom expressions with an alias and defined output type.

Column Name	Tags	Description	fx Expression	Type	Alias	PII
## PLAN_ID	planid	This is likely a unique identifier for a specific package.		Long	PLAN_ID	
A PLAN_NAME	3 tags	This is likely the name of a specific plan or package.		String	PLAN_NAME	
## MER	2 tags	This could stand for 'Management Expense Ratio'.		Long	MER	

PLAN_ID	PLAN_NAME	MER	PACKAGE
249	Promo - All Access - 12mo - 132.00 - (+\$1.00 per month)	11	Discount
252	Fixed Demo - Select - 6mo - (Used Car)	0	TBD
253	NavTraffic + NavWeather -	6.66	TRD

Type conversions - you can select a specific data type for the fields within a column (e.g. cast a date field to a string)

Column Name	Tags	Description	fx Expression	Type	Alias	PII
A PLAN_NAME	3 tags	This is likely the name of a specific plan or package.	LOWER("PLAN_NAME")	String	PLAN_NAME	
## MER	2 tags	This could stand for 'Management Expense Ratio'.			MER	
A DISC		This could be an abbreviation for (Discount).			DISC	

PLAN_ID	PLAN_NAME	MER	DISC	PACKAGE
249	Promo - All Access - 12mo - 132.00 - (+\$1.00 per month)	11	Discount	All Accer
252	Fixed Demo - Select - 6mo - (Used Car)	0	TBD	Select
253	NavTraffic + NavWeather -	6.66	TRD	



Alias - you can set a specific alias for a field.

Column Name	Description	fx Expression	Type	Alias	PII
<input checked="" type="checkbox"/> A PLAN_NAME	This is likely the name of a specific plan or package.	LOWER("PLAN_NAME")	String	Plan	<input type="checkbox"/>
<input checked="" type="checkbox"/> ## MER	This could stand for 'Management Expense Ratio'. a		Long	MER	<input type="checkbox"/>
<input checked="" type="checkbox"/> A DISC	This could be an abbreviation for 'Discount'. (right)		String	DISC	<input type="checkbox"/>

PLAN_ID	PLAN_NAME	MER	DISC	PACKAGE
249	Promo - All Access - 12mo - 132.00 - (+\$1.00 per month)	11	Discount	All Acces
252	Fixed Demo - Select - 6mo - (Used Car)	0	TBD	Select
253	NavTraffic + NavWeather - 1mo - wactv	6.66	TBD	TBD

You can modify your Fields as much as required until you are satisfied. After each change click the Run button to execute the query. When you're satisfied with the results, click Save.

Source	Datasource	Schema	Table
Source	oracle	CUSTOMER_DEMO	MASTER_PLAN_TABLE_WDATES

Fields
10 Selected Fields

Filters
STARTDATE is between 2015-01-01, 2018-12-31

PLAN_ID	Plan	MER	DISC	PACKAGE
249	promo - all access - 12mo - 132.00 - (+\$1.00 per month)	11	Discount	All Acces
252	fixed demo - select - 6mo - (used car)	0	TBD	Select
253	navtraffic + navweather - 1mo - wactv	6.66	TBD	TBD
254	marine offshore - 6mo - w	6.66	TBD	TBD



SQL

Clicking the <> Write SQL button launches the advanced SQL mode. Swapping between modes (clicking the Hide Code button) can cause some customizations to be lost. It is far more common to lose customizations when moving from the advanced views to the less advanced.

The screenshot shows the SQL editor interface with a query and its results. The query is as follows:

```
1 WITH
2 select_step1 as (
3 SELECT
4 "redshift"."customer_demo"."orders"."order_id" AS "order_id",
5 "redshift"."customer_demo"."orders"."customer_id" AS "customer_id",
6 CAST(
7 "redshift"."customer_demo"."orders"."order_date" AS DATE
8 ) AS "order_date",
9 "redshift"."customer_demo"."orders"."status" AS "status",
10 "redshift"."customer_demo"."orders"."credit_card_amount" AS "credit_card_amount",
11 "redshift"."customer_demo"."orders"."coupon_amount" AS "coupon_amount",
12 "redshift"."customer_demo"."orders"."bank_transfer_amount" AS "bank_transfer_amount",
13 "redshift"."customer_demo"."orders"."gift_card_amount" AS "gift_card_amount",
14 "redshift"."customer_demo"."orders"."amount" AS "amount"
15 FROM
16 "redshift"."customer_demo"."orders"
17 ),
18 join_step2 as (
19 SELECT
20 select_step1."order_id" AS "order_id",
21 select_step1."customer_id" AS "customer_id",
22 CAST(select_step1."order_date" as DATE) AS "order_date",
23 select_step1."status" AS "status",
24 select_step1."credit_card_amount" AS "credit_card_amount",
25 select_step1."coupon_amount" AS "coupon_amount",
26 select_step1."bank_transfer_amount" AS "bank_transfer_amount",
27 select_step1."gift_card_amount" AS "gift_card_amount",
28 select_step1."amount" AS "amount"
```

The results table shows the following data:

order_id	customer_id	order_date	status	credit_card
1	1	2024-07-25	redshift	10
9	53	2018-01-12	completed	0

The screenshot shows the SQL editor interface with an "Auto-complete Details" dialog box open. The dialog box has the following options:

- Off
- Across Entire Catalog
- From a Subset of the Catalog (Selected Domain)
- From Tables added to this Datamap
- Across selected Datasource Catalog

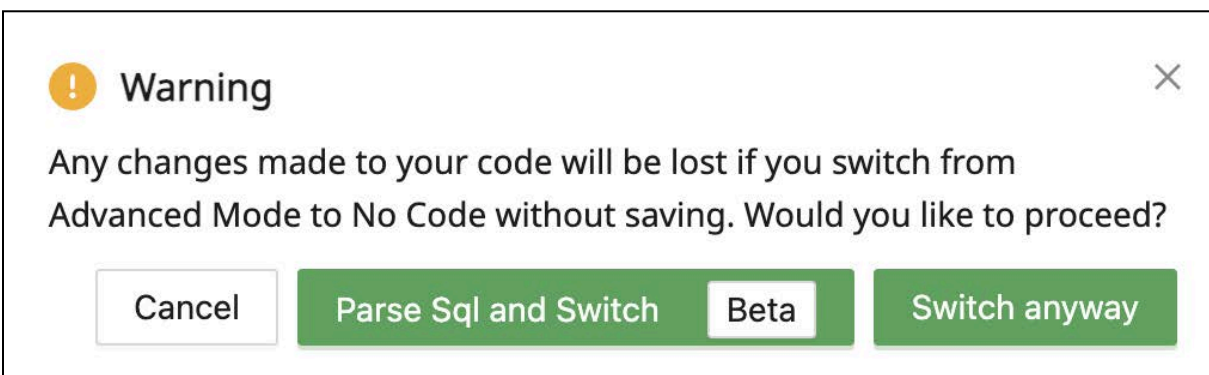
The dialog box also has "Cancel" and "Save" buttons. The background shows the same query and results as the previous screenshot.



The advanced SQL mode is a web based SQL IDE wherein you can modify and execute any changes you need.

The Auto-complete feature allows for the auto-completing for the names of schema, tables, and columns.

This view is intended for SQL experts. SQL is a powerful language and it's easy to encounter errors. You can modify your SQL as often as required and click the Run button to execute the query. When you're satisfied with the results, click Save.



Transitioning from Code mode to No-Code, the 'Parse SQL and Switch' option automatically tracks and visualizes the data flow in your SQL queries.

**Please note: Currently only available for datamaps created initially in Code Mode.*

Ask Reasoner

Clicking the Ask Reasoner button, a window will pop up. Select the appropriate Domain to narrow the search for a more relevant dataset. Give the datamap a title, once you are satisfied, click "Ask Reasoner" on the bottom right. When using Reasoner, Prometheus uses key phrases used in the Datamap title to create your datamap, including parameters when asking a question.



Dashboard Data Explorer Answers Data Sources Projects Chat **Get Started**

support_datamap_1699296837476

Create Reasoner Assisted Datamap ✕

Data Source
Promethium

Domain Selection
Perfect Order

Datamap Title
What is the top 5 quantity by customer by status where status is Imperfect

Optional Description
Customers with the most imperfect orders

Use Generative AI

Close **Ask Reasoner**


Preview



Dashboard Data Explorer Answers Data Sources Projects Chat Get Started

support_datamap_1699296837476 Show Code

Create Reasoner Assisted Datamap



Reasoner is crunching numbers based on your Title and Tags...

- Finding matching tables...
- Finding matching tags...
- Finding columns...
- Finding joins...

Close Ask Reasoner

Preview

Dashboard Data Explorer Answers Data Sources Projects Chat Get Started

What_is_the_top_5_quantity_by_customer_by_status_where_status_is_... Show Code

Run Save Publish Build Answer

Data Source: Promethium Default

Select: orders_usa_flagged_mqt


Filter: filter step

Group By: group_by step

Order By: order_by step

Limit Rows: 5

+ Add Step



Run the SQL query to get your results



quantity	customer_name	status
28	MBD DISTRIBUTING	Imperfect
66	CORE-MARK - PORTLAND	Imperfect
49	EBY BROWN COMPANY LLC	Imperfect
84	VISTAR DISTRIBUTORS INC EAST	Imperfect
2	HOSTESS BRANDS LLC	Imperfect

Data Pipelines

After a datamap has been saved, a new option of Publish will be available as will the Destinations tab on the left of the screen.

Promethium supports publishing a View, Table, or dbt Model. For each option you can configure the destination, selecting a data source and schema to publish to. The view name will be the name you assigned to your datamap when you saved it in the previous step.

test_run_pipeline_notif

Destination Configuration

Republish Done

The Destination table has been created successfully. Do not forget to Run/Schedule the Pipeline next to get rows into MQT.

Datamaps built on Promethium can only be published to datasources deployed in Promethium

DATASOURCE	SCHEMA	TABLE NAME
Promethium	promethium	test_run_pipeline_notif_1705012951786_mqt

Status: Published Last Published: January 11, 2024, 14:42 by S

Next: Run Pipeline >

Activities

No jobs found. Comeback later to check for scheduled jobs!

Refresh



Select your desired datasource, schema and specify a table name. Click Publish.

Your pipeline will execute.

Pipeline Scheduling

You can schedule pipeline execution for any published Table or dbt Model.

From the Destination tab, click the dbt Model or Table you wish to schedule. Click Run to run the data pipeline.

User	Status	Start Time	Type
S	SUCCESS	January 11, 2024, 15:14	datamap
S	SUCCESS	January 11, 2024, 14:43	datamap

On the data pipeline page, you can select a time based schedule. You can execute a one time pipeline execution by selecting non as your schedule. Alternatively, you can configure your pipeline to execute every specified number of minutes, hours, days, weeks, or months. Daily, weekly, or monthly selections will require you to specify a time of day in addition to the specific day or date.

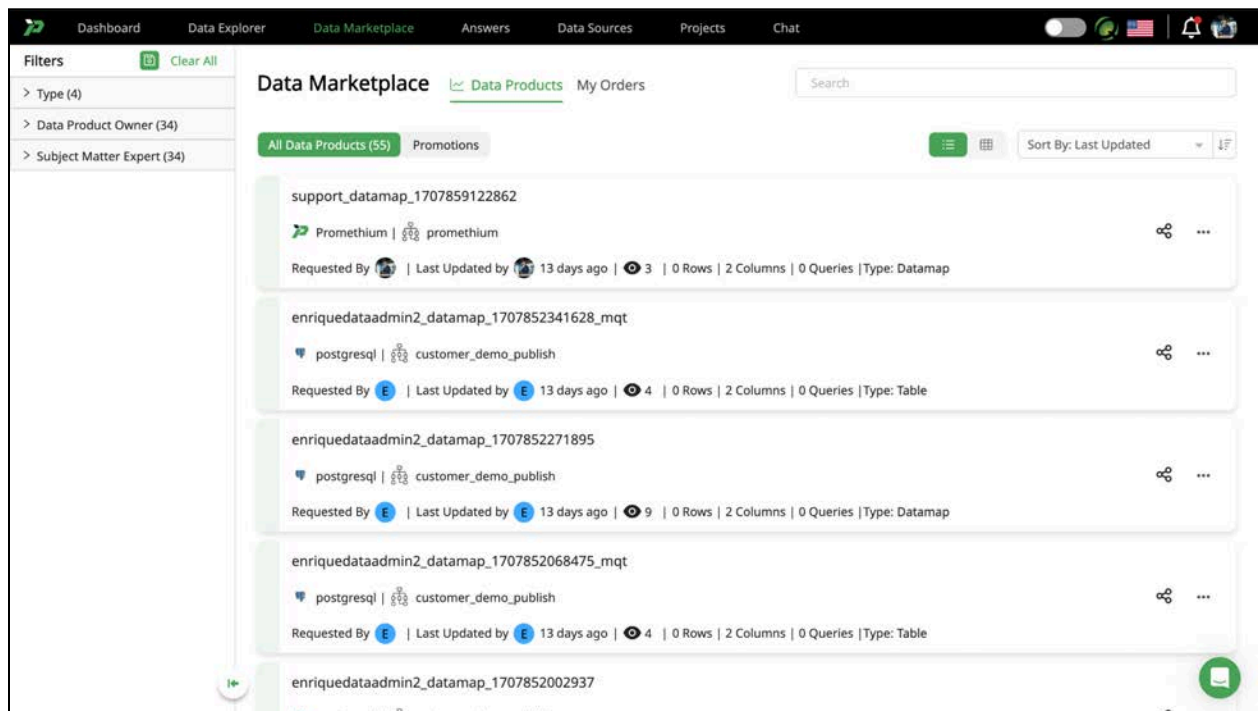
You can enforce the continuous execution of your pipeline on your schedule by leaving the default choice of Ends At Never selected. Alternatively, you can specify a specific end date/time to limit the pipeline execution to a specific schedule.

When you are satisfied with your pipeline schedule click the green Schedule button to save your selections and trigger the pipeline schedule to execute.



Data Marketplace

The Data Marketplace is a centralized resource for procuring data products, including data maps, answers, views, and tables, tailored to meet your specific needs.



Explore data products available for purchase, presented in either a list or grid format.

On the left side, you can filter the available data products by type, Data Product Owner, and Subject Matter Expert.



The screenshot shows the Data Marketplace interface. On the left, there are filters for 'Type (4)' (Answers, Tables, Views, Datamaps) and 'Data Product Owner (34)' with a search bar and a list of email addresses. The main area displays a grid of data products under the heading 'Data Marketplace' and 'Data Products'. Each product card includes the product name, update time, owner, subject matter expert, type (Table or Datamap), source (e.g., postgresql), and promotion status. A search bar is located at the top right of the main area.

Publishing a Data Product

A data product must be published into the Data Marketplace in order for it to be available. Once published, you will be unable to make changes to the data product unless the data product is unpublished and then republished.

The screenshot shows the 'data marketplace publish test' configuration window. On the left, there are sections for 'View', 'Table' (with a green checkmark), and 'dbt Model'. The main area is titled 'Destination Configuration' and contains a table with columns 'DATASOURCE', 'SCHEMA', and 'TABLE NAME'. The table shows 'postgresqlyauhent...', 'customer_demo', and 'data_marketplace_publish_test_1708966555395_mqt'. Below the table, it indicates 'Status: Published' and 'Last Published: February 26, 2024, 08:56 by S'. There are buttons for 'Republish', 'Publish to Data Marketplace', and 'Done'. At the bottom, there is an 'Activities' section with the message 'No jobs found. Comeback later to check for scheduled jobs!' and a 'Refresh' button.

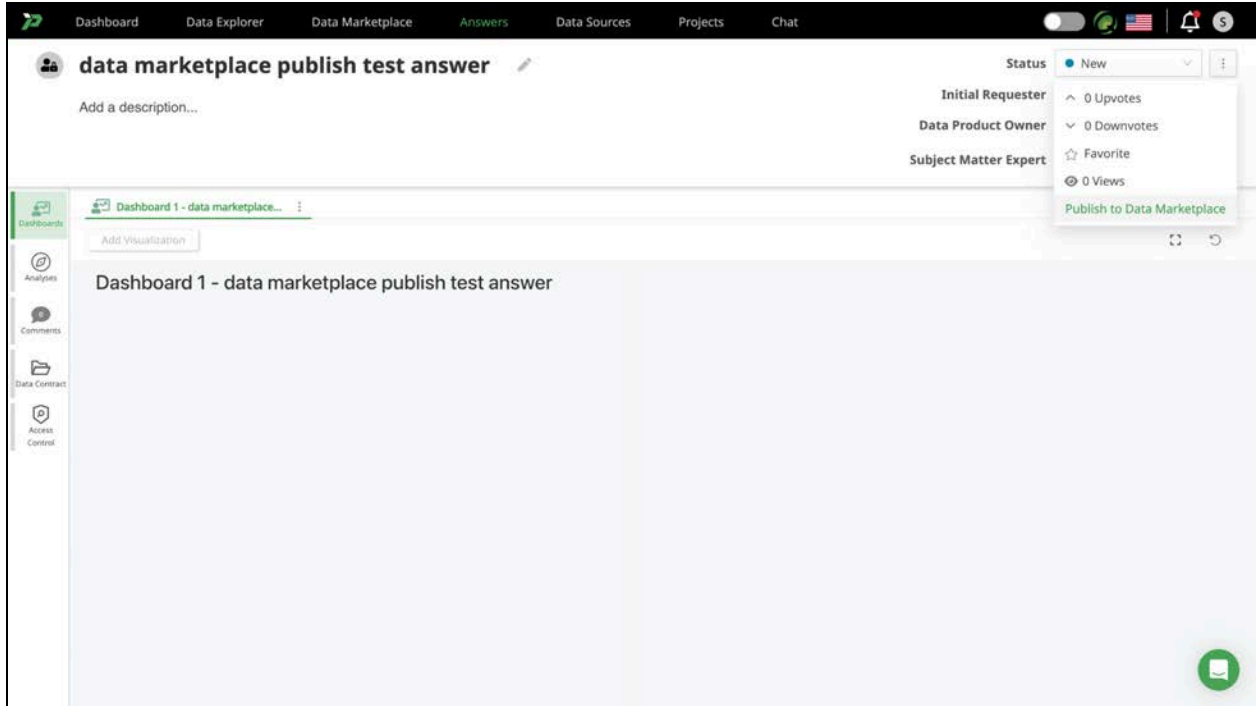


The screenshot shows the 'data marketplace publish test' page. A modal dialog box is open with the following text: "You will not be able to make edits once you publish to the Data Marketplace. To make edits you will first have to unpublish. Are you sure you'd like to publish to the Data Marketplace?". The dialog has 'Cancel' and 'OK' buttons. In the background, the 'Publish to Data Marketplace' button is visible, along with a 'Done' button. Below the dialog, the status of the data product is shown as 'Published', with a 'Last Published' timestamp of 'February 26, 2024, 08:56' by user 'S'. A 'Next: Run Pipeline >' button is also present. The 'Activities' section at the bottom indicates 'No jobs found. Comeback later to check for scheduled jobs!' with a 'Refresh' button.

The screenshot displays the 'Data Marketplace' interface, specifically the 'Data Products' section. The page title is 'Data Marketplace' with sub-sections 'Data Products' and 'My Orders'. A search bar is located at the top right. The main content area shows a list of data products under the heading 'All Data Products (44) Promotions'. The products listed are:

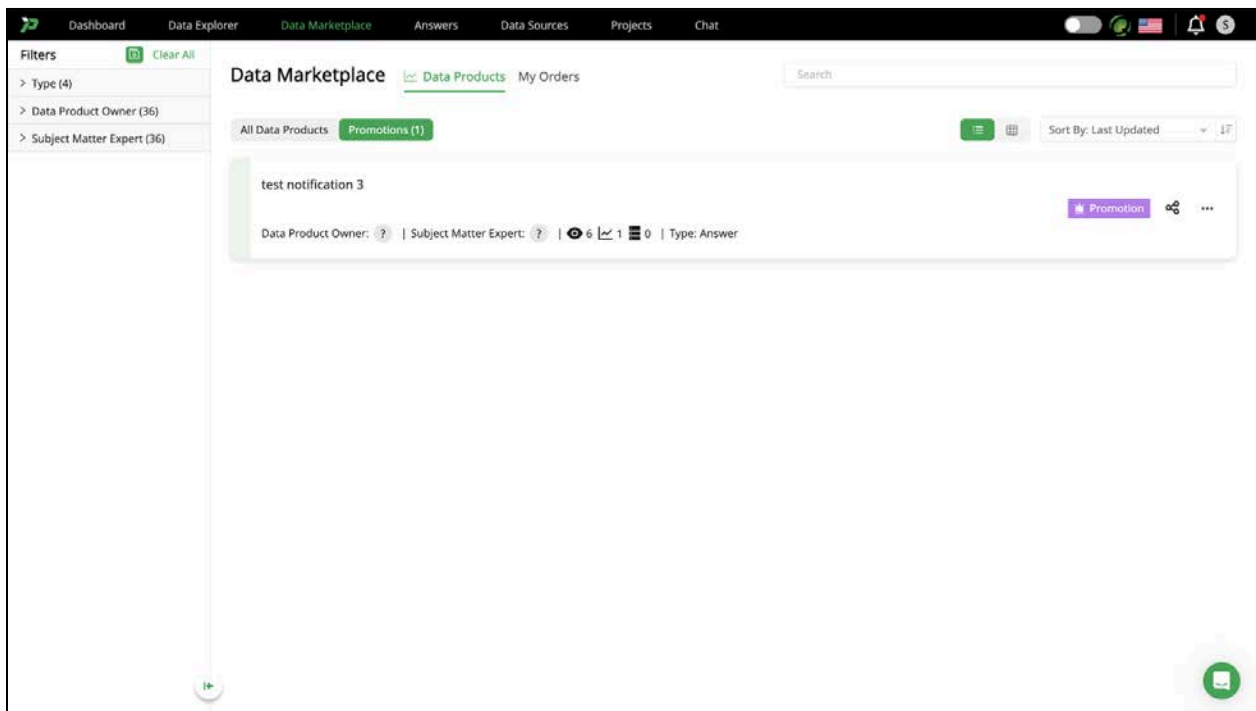
- data marketplace publish test answer**: Data Product Owner: ? | Subject Matter Expert: ? | 0 | 1 | 0 | Type: Answer
- data_marketplace_publish_test_1708966555395_mqt**: Requested By S | Last Updated by S a minute ago | 1 | 0 Rows | 5 Columns | 0 Queries | Type: Table. Includes an 'Unpublish' button.
- Test 3**: Data Product Owner: ? | Subject Matter Expert: ? | 3 | 1 | 1 | Type: Answer
- Test 2**: Data Product Owner: ? | Subject Matter Expert: ? | 1 | 1 | 1 | Type: Answer
- Test 1**

Each product entry includes a share icon and a menu icon. A 'Filters' sidebar on the left shows categories like 'Type (4)', 'Data Product Owner (35)', and 'Subject Matter Expert (35)'. The page is sorted by 'Last Updated'.



Promotions

In the main Data Marketplace tab, promoted data products are distinctly highlighted. To exclusively view these promoted items, simply navigate to the “Promotions” section.





Clicking on a promoted data product provides an overview including the promotion duration, data product type, description, tags, creation details including who and when, the count of associated dashboards, queries, and visualizations, as well as information on the assigned Data Product Owner and Subject Matter Expert.

On the right side of the page, you will be able to see similar published data products in the Data Marketplace.

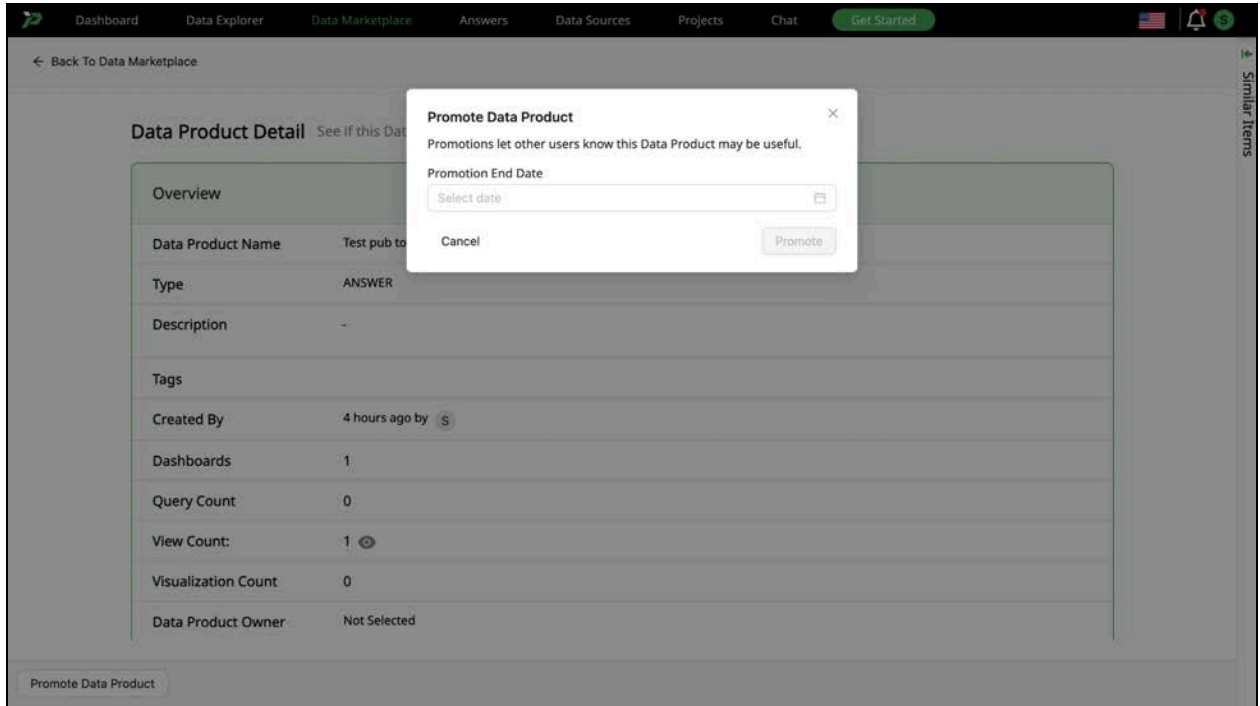
If you are the data admin of the data product, you can edit the end date of the promotion by clicking **“Edit Data Product Promotion”**.

The screenshot displays the Data Marketplace interface. The top navigation bar includes Dashboard, Data Explorer, Data Marketplace, Answers, Data Sources, Projects, Chat, and a Get Started button. The main content area is titled "Data Product Detail" and includes a sub-header "See if this Data Product fits your needs by exploring its details." Below this is a table with the following information:

Overview	
Data Product Name	Test pub to market
Type	ANSWER
Description	-
Tags	
Created By	3 hours ago by
Dashboards	1
Query Count	0
View Count:	1
Visualization Count	0
Data Product Owner	Not Selected

At the bottom left of the detail view is a "Promote Data Product" button. On the right side, there is a "Similar Items" section with two items:

- Test Pub1 test pub** (View)
 - Private
 - Subject Matter Expert: Not Selected
 - Data Product Owner: Not Selected
 - Last Updated: 2 months ago by
 - Description
- data marketplace publish test answer** (View)
 - Private
 - Subject Matter Expert: Not Selected
 - Data Product Owner: Not Selected
 - Last Updated: 2 months ago by
 - Description
 - AI Generated Description**
 - 1. **Hospital Service Area** - A broader ge
 - 2. **Permanent Facility ID** - A unique ide



My Orders

Under the “**My Orders**” tab, you can see a breakdown of all your data product orders by their status: Pending Orders, Approved Orders, and Rejected Orders.



Dashboard Data Explorer Data Marketplace Answers Data Sources Projects Chat

Filters Clear All

- > Type (4)
- > Data Product Owner (35)
- > Subject Matter Expert (35)

Data Marketplace

Data Products My Orders Search

All Orders (6) Pending Orders Approved Orders Rejected Orders Sort By: Last Updated LF

- test notification 3 Approved
Data Product Owner: ? | Subject Matter Expert: ? | 7 | 1 | 0 | Type: Answer Promotion
- c-qaizer answer for market Approved
Data Product Owner: ? | Subject Matter Expert: ? | 7 | 1 | 1 | Type: Answer
- newtestforpublish Approved
Data Product Owner: ? | Subject Matter Expert: ? | 6 | 1 | 0 | Type: Answer
- Answer to reject Approved
Data Product Owner: ? | Subject Matter Expert: ? | 8 | 1 | 1 | Type: Answer
- Max Answer for Market Approved
Data Product Owner: C | Subject Matter Expert: C | 31 | 1 | 1 | Type: Answer

Dashboard Data Explorer Data Marketplace Answers Data Sources Projects Chat

Filters Clear All

- > Type (4)
- > Data Product Owner (35)
- > Subject Matter Expert (35)

Data Marketplace

Data Products My Orders Search

All Orders **Pending Orders (1)** Approved Orders Rejected Orders Sort By: Last Updated LF

- test md12 Pending
Data Product Owner: S | Subject Matter Expert: C | 19 | 1 | 1 | Type: Answer
allow temp description newMD xxxxxxxxxxxxxxx... editing



The screenshot shows the 'Data Marketplace' interface with the 'My Orders' tab selected. The left sidebar contains filters for 'Type (4)', 'Data Product Owner (35)', and 'Subject Matter Expert (35)'. The main content area displays a list of five orders, all marked as 'Approved'. The orders are:

- test notification 3 (8 views, 1 like, 0 dislikes)
- c-qaiser answer for market (7 views, 1 like, 1 dislike)
- newtestforpublish (6 views, 1 like, 0 dislikes)
- Answer to reject (8 views, 1 like, 1 dislike)
- Max Answer for Market (31 views, 1 like, 1 dislike)

Each order card includes a 'Promotion' button and a share icon. The top navigation bar includes 'Dashboard', 'Data Explorer', 'Data Marketplace', 'Answers', 'Data Sources', 'Projects', and 'Chat'.

The screenshot shows the 'Data Marketplace' interface with the 'My Orders' tab selected. The left sidebar contains filters for 'Type (4)', 'Data Product Owner (35)', and 'Subject Matter Expert (35)'. The main content area displays a message: 'Oops! There are no products yet'. The top navigation bar includes 'Dashboard', 'Data Explorer', 'Data Marketplace', 'Answers', 'Data Sources', 'Projects', and 'Chat'.

Checkout

Once a user completes the checkout process for an item, it appears instantly in the Data Explorer and on the Answer page, even if it doesn't match their existing permissions or the domains they can access. Although the item isn't linked to any specific domain, it's labeled as



'Private,' allowing the user unique access to utilize the Data Product based on their individual needs.

The Data Product Owner can decide to approve or deny access. Once they approve, the user can access the item they wanted, found under the “My Orders” tab in the Data Marketplace. It's a good idea, though not required, for users to tell the data product owner how they plan to use the item, helping the owner decide whether to allow access.

Checkout Data Product Request ✕

The Data Steward must approve your checkout before you can use this Data Product. Adding a note about what you will be using it for can help them decide.

Cancel Request



Notifications(324) CLOSE

System Notifications User Notifications ⁹⁹⁺

- E** enrique+businessuser1@pm61data.com has pending Data Product checkout request
8 days ago
example
[See details](#)
- C** c-lan@pm61data.com has pending Data Product checkout request
8 days ago
revenue by plan
[See details](#)
- E** enrique+businessuser1@pm61data.com has pending Data Product checkout request
8 days ago
support_dmp_for_business_user_1708016203358_mqt
[See details](#)
- E** enrique+businessuser1@pm61data.com has pending Data Product checkout request
8 days ago
revenue by plan
[See details](#)
- E** enrique+businessuser1@pm61data.com has pending Data Product checkout request
8 days ago
dmb_for_buser_1708013213251_mqt
[See details](#)

Dashboard Data Explorer Data Marketplace Answers Data Sources Projects Chat

Notifications(326) CLOSE

System Notifications User Notifications ⁹⁹⁺

Data Product Check Out Request

C c-lan@pm61data.com 8 days ago
No message provided.

Overview

Data Product Name	support_dmp_for_business_user_1708016203358_mqt
Type	TABLE
Description	support_dmp_for_business_user_1708016203358_mqt
Tags	
Created By	8 days ago by S
View Count:	3
Data Product Owner	Not Selected
Subject Matter Expert	Not Selected

Accept Request
 Reject Request

Cancel Done

enrique+businessuser2@pm61data.com has pending Data Product checkout request

Answers

All in one convenient location, questions can be searched, created, and answered.



Search for Answers

Searching for and filtering Answers is easy and uses keywords and tags associated with the Answer to make search results highly relevant.

Using the search bar, you can quickly find results matching your query. Matches will be directly displayed below the search bar. You can select the most appropriate answer directly or hit return to see a more detailed results page from which to make your selection.

The screenshot shows a web application interface for searching answers. The top navigation bar includes 'Dashboard', 'Data Explorer', 'Answers', 'Data Sources', 'Projects', and 'Chat'. The 'Answers' section is active, displaying a search for 'revenue'. On the left, there is a 'Filters' sidebar with options like 'My Items', 'Favorites', 'Assigned to Me', and 'Has Stories'. Below these are expandable sections for 'Domains (17)', 'Data Product Owner (30)', 'Status (7)', and 'Subject Matter Expert (30)'. The main content area shows a table of search results with columns for 'ANSWER', 'STATUS', 'REQUESTED BY', and 'LAST UPDATED'. The results are filtered to show 'All (158)' items. A dropdown menu is open over the first few rows, showing details for each answer, including the requester's name and the time it was asked. To the right of the table, there is a 'Request Answer' button and a dropdown menu for sorting results by 'Last Updated'. Below the table, there are several answer cards, each with a title, requester information, and a status indicator (e.g., 'In Progress').

ANSWER	STATUS	REQUESTED BY	LAST UPDATED
What is revenue for last year	In Progress	S	Asked 6 days ago
What is revenue by plan by vendor from 2018 to...	In Progress	S	Asked 5 days ago
What is the revenue subscription by plan by ven...	In Progress	S	Asked 4 days ago
What is the revenue subscription by plan by ven...	In Progress	S	Asked 4 days ago
what is the average revenue by vendor location ...	In Progress	S	Asked 4 days ago



Dashboard | Data Explorer | **Answers** | Data Sources | Projects | Chat

Filters: Clear All

- My Items
- Favorites
- Assigned to Me
- Has Stories

> Domains (17)

> Data Product Owner (30)

> Status (7)

> Subject Matter Expert (30)

Answers

 revenue + Request Answer

All (158) Sort By: Last Updated

Card Title	Requester	Views	Dashboards	Queries	Status
testing_copy_master	Demo	0	1	1	In Progress
master plan analyze test	Demo	0	1	1	In Progress
support_datamap_1698923113903	Private	0	1	1	In Progress

The detailed results page shows tags, status, and who originally requested the Answer alongside usage metrics such as view count, the number of dashboards that have been created, and the number of queries. Available to view results in either a list or card view.

Dashboard | Data Explorer | **Answers** | Data Sources | Projects | Chat

Filters: Clear All

- My Items
- Favorites
- Assigned to Me
- Has Stories

> Domains (17)

> Data Product Owner (30)

> Status (7)

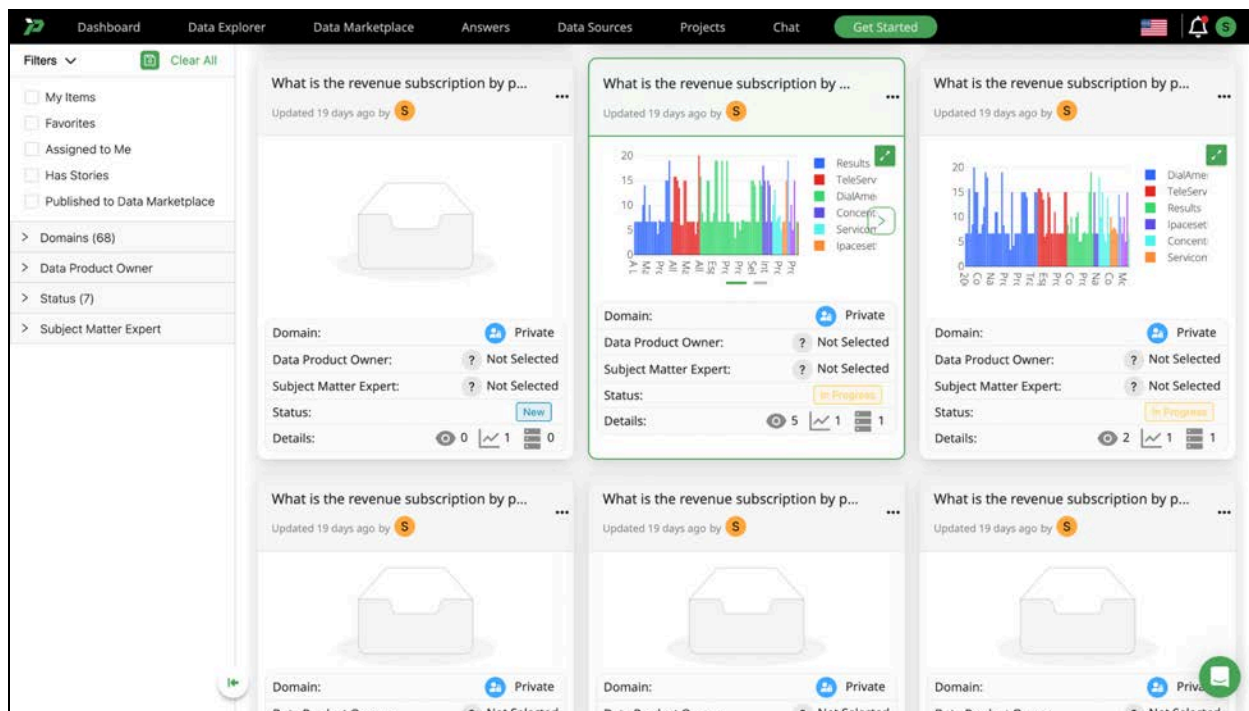
> Subject Matter Expert (30)

Answers

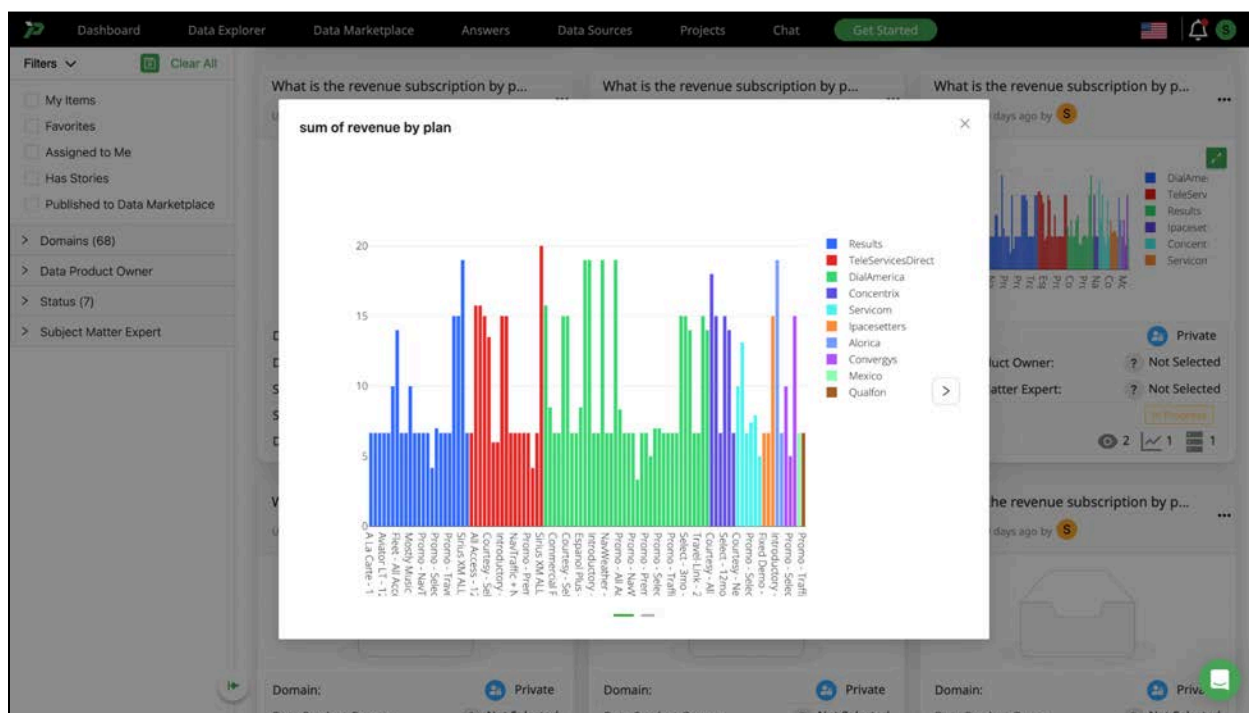
 revenue + Request Answer

All (55) Sort By: Last Updated

Card Title	Requester	Views	Dashboards	Queries	Status
Revenue By Year	Sales	0	0	0	In Progress
revenue by plan by package	Private	2	1	1	In Progress
revenue by plan by package	Private	0	1	1	In Progress



For viewing Answers with multiple visualizations. Seamlessly switch between various visual representations within a single answer to enhance your data exploration. This provides deeper insights and a dynamic view of the full context and value of the Answer.





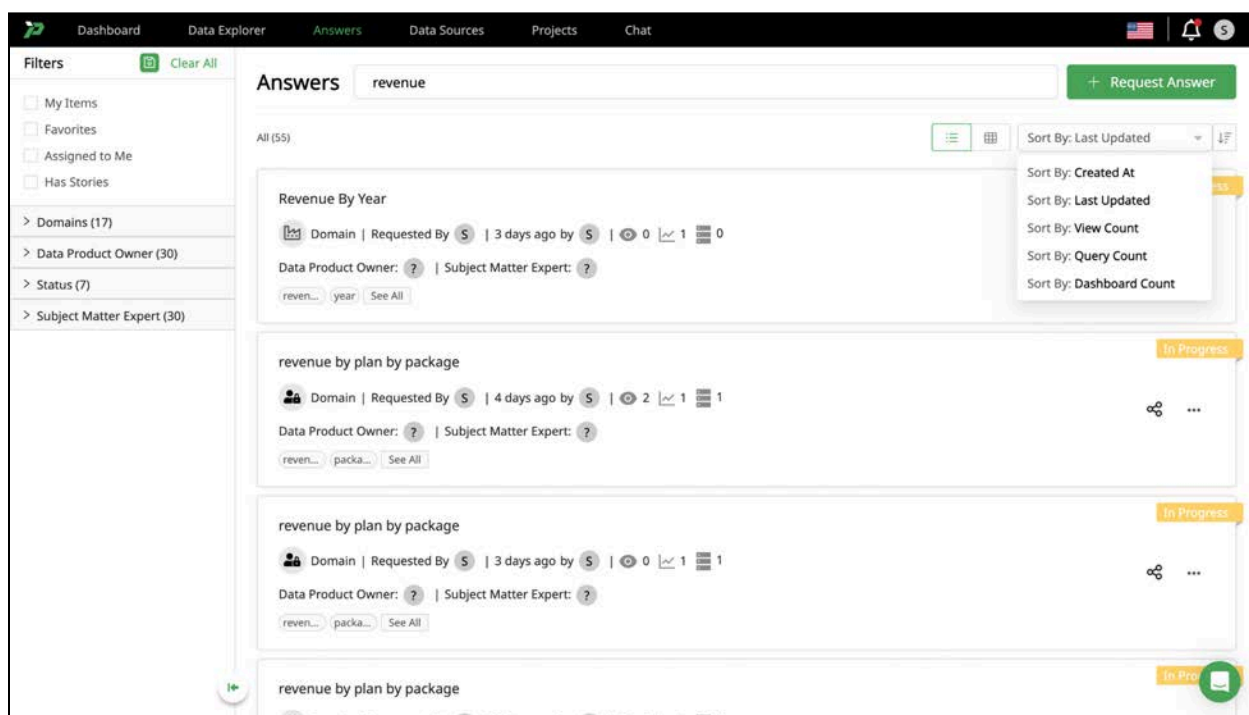
Filter and Sort Answers

As with the Data Explorer, you can rapidly filter your results using the powerful result filters that are present immediately above the table of results.

- **Domains** - Search or browse for an existing Domain to limit the result set to that Domain or set of Domains.
- **Status** - Filter by the status set by the owner of the Answer - Published, New, In Progress, or Assigned to Me (Answered or Unanswered).
- **Has Stories** - Has Stories will limit the result set to those that have Stories. Clicking the filter will toggle it on or off.
- **Favorites** - Favorites will limit the result set to those that you as a user have favorited. Clicking the filter will toggle it on or off.

You can combine one or all of these filters to rapidly narrow in on your required data.

To help find and explore the Answers available you can select from several Sort By refiners from the drop down at the top right.



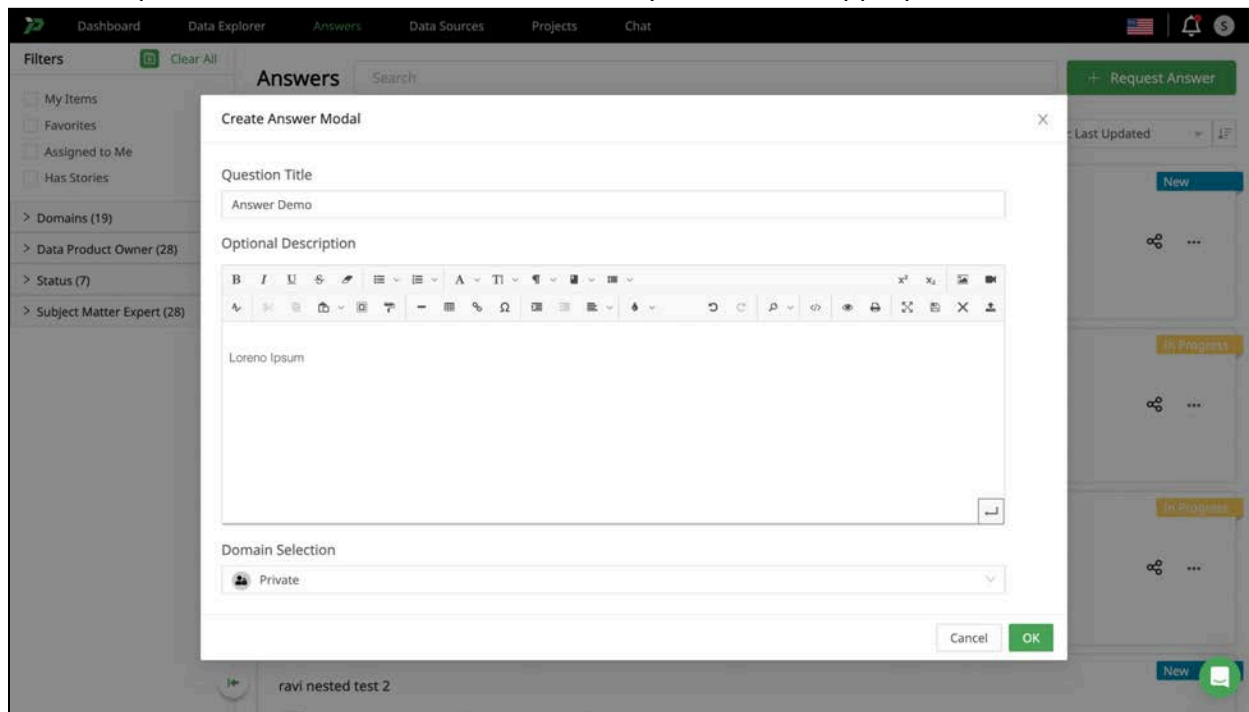
Feature	Description
Last Updated	Sort by the date of the last update
View Count	Sort by the number of times the Answer has been viewed (e.g. how popular the Answer is)
Dashboard Count	Sort by the number of Dashboards present within the Answer. A higher number may indicate a more valuable Answer.



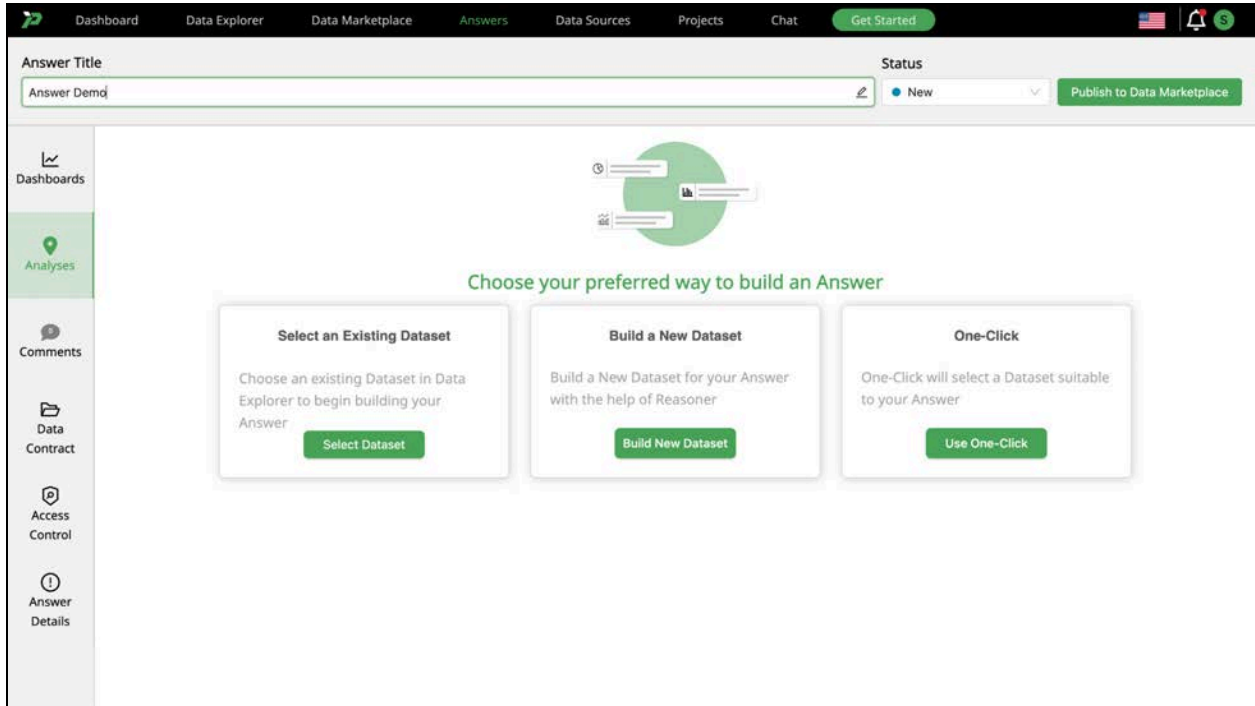
Should you still not have found an Answer meeting your requirements you can quickly select Request Answer to begin that flow.

Requesting an Answer

If you start this flow by clicking Request Answer from a results page, the search term that you failed to find an answer for will be used as the title for this new Answer. If you want to edit this, select the pencil icon next to the name and modify it to a more appropriate title.



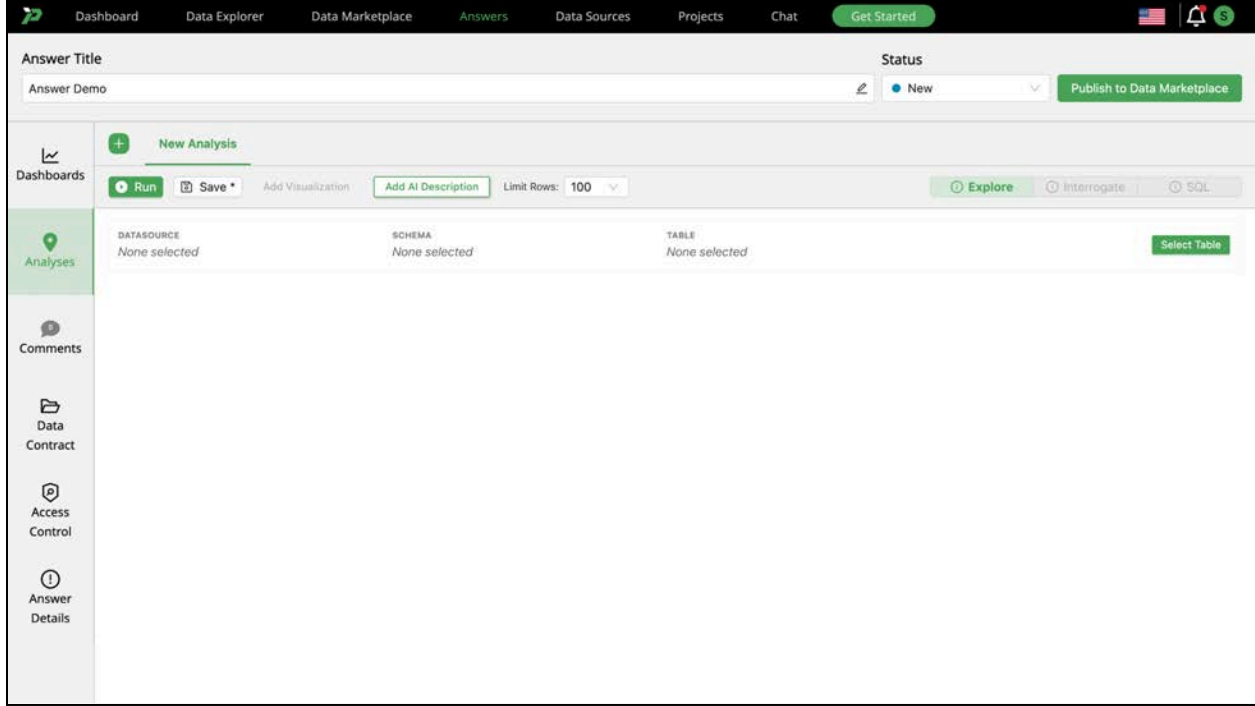
If you start this flow by clicking Request Answer with no text in the search box you will be given an auto generated title consisting of your username and a date / timestamp. You can add/remove tags and even add a description to help the data team understand the context behind your need and answer the question completely. Tagging data and Answers enhances the search experience for all users. To add tags to an Answer click the Add button next to the title.



The three major steps to create an Answer are to build, select an existing datamap, or use One-Click to create visualizations.

Select an Existing Dataset

To build an Answer using a pre-existing datamap choose Select an Existing Dataset.





First, we must select at least one Datasource, Schema, and Table. Click Select Table which will render the familiar Data Explorer user interface. Here, as with the Data Explorer tab you can search, filter, and refine the available data sources to find what you need.

If you find what you need, select it and choose the green Analyze button

The screenshot shows the Data Explorer interface. At the top, there's a navigation bar with 'Dashboard', 'Data Explorer', 'Data Marketplace', 'Answers', 'Data Sources', 'Projects', 'Chat', and 'Get Started'. Below this, the 'Answer Title' is 'Answer Demo' and the status is 'New'. A 'Publish to Data Marketplace' button is visible. The main area is titled 'New Analysis' and contains a search bar, a 'New Datamap' button, and a table of results. The table has columns for 'Reasoner', 'Data Source', 'Schema', and 'Name'. One result is selected: 'postgres...' (Data Source), 'customer_demo' (Schema), and 'car_sales_revenue_us_canada' (Name). To the right, a detailed view of the selected table is shown, including its data source, schema, type, columns, rows, and a description field. At the bottom right, there are 'New Datamap' and 'Analyze' buttons.

You will be returned to the Answer flow where you can search and filter the data available from your selected Datamap.

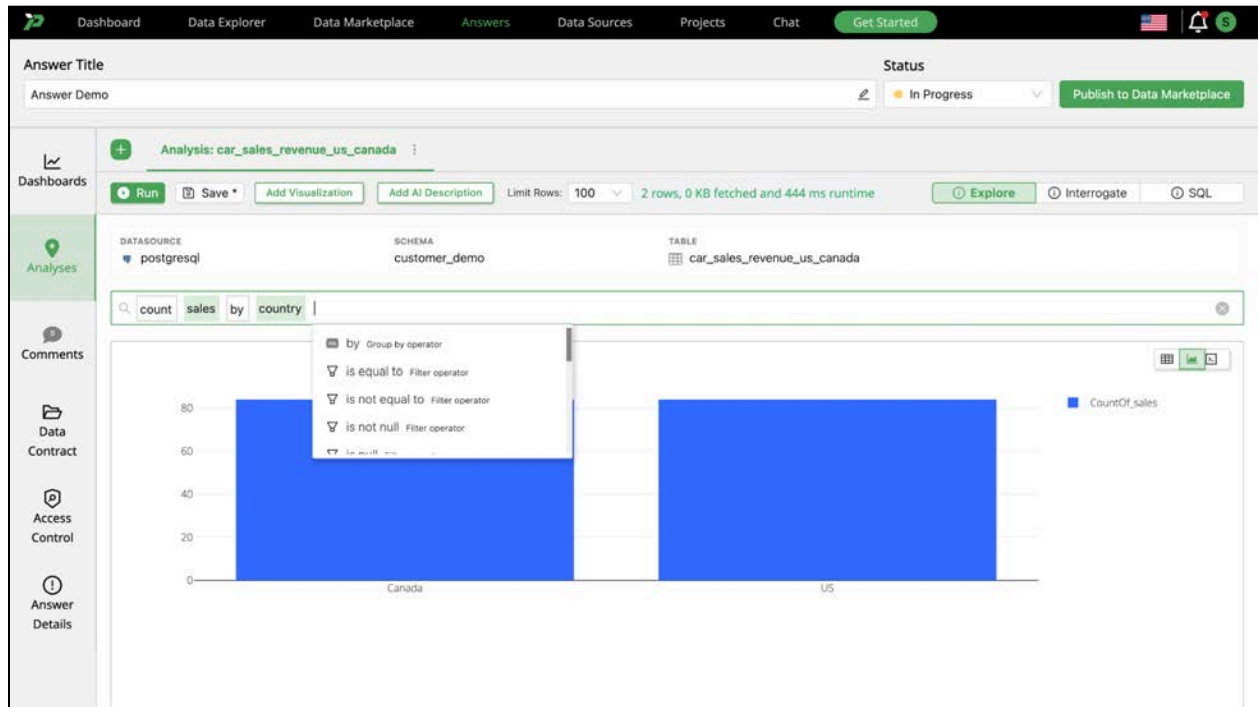
The screenshot shows the Data Explorer interface after clicking 'Analyze'. The status is now 'In Progress'. The main area is titled 'Analysis: car_sales_revenue_us_canada'. It includes a 'Run' button, a 'Save' button, and an 'Add Visualization' button. Below this, there are buttons for 'Add AI Description', 'Limit Rows: 100', '100 rows, 9 KB fetched and 472 ms runtime', 'Explore', 'Interrogate', and 'SQL'. The data table is displayed with columns for 'country', 'sales', 'date', 'year', and 'month'. The data is as follows:

country	sales	date	year	month
US	1006049	2014-01-01T00:00:00	2014	1
US	1186038	2014-02-01T00:00:00	2014	2
US	1529348	2014-03-01T00:00:00	2014	3
US	1381300	2014-04-01T00:00:00	2014	4
US	1598730	2014-05-01T00:00:00	2014	5
US	1412565	2014-06-01T00:00:00	2014	6



Explore

Clicking the search field will trigger the display of all the available search verbs and nouns that are appropriate in context such as Sum, Count, etc. Simply by selecting the appropriate suggestions you can build a powerful query that will attempt to render using an appropriate format such as a table or graph.



Interrogate

The Interrogate feature allows you to analyze data through the use of natural language.

With the help of the Generative AI option alongside the Interrogate feature, you can easily and efficiently extract deeper insights and uncover important correlations within their datasets.



Dashboard Data Marketplace Answers Data Sources Projects Chat Get Started

Answer Title: Answer Demo Status: In Progress Publish to Data Marketplace

Analysis: car_sales_revenue_us_canada

Run Save Add Visualization Add AI Description Limit Rows: 100 2 rows, 0 KB fetched and 444 ms runtime Explore Interrogate SQL

DATASOURCE	SCHEMA	TABLE
postgresql	customer_demo	car_sales_revenue_us_canada

Search your data

Use Generative AI

country	CountOf_sales
US	84
Canada	84

Comments: Data Contract, Access Control, Answer Details

Answer Title: Genai Interrogate Analysis Example

Answer Details

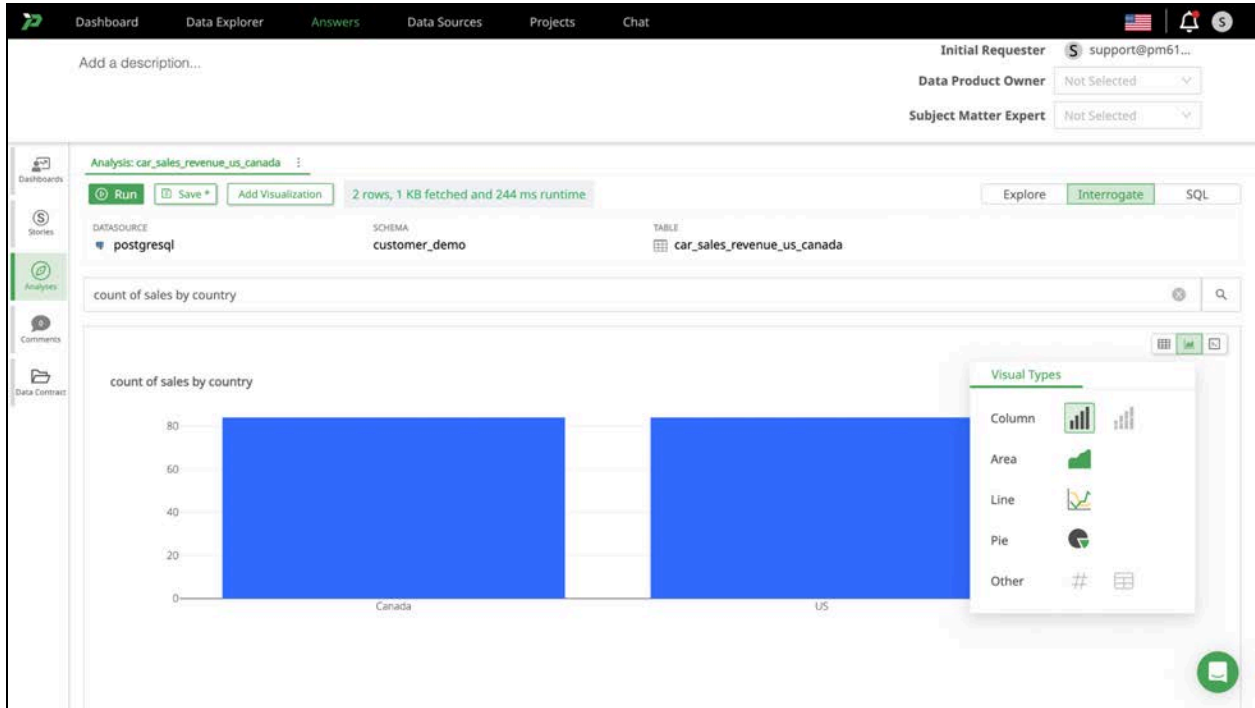
Answer Title: Genai Interrogate Analysis Example

Description

AI Generated Description

1. "number_of_orders" - Indicates the number of orders associated with each first name.
2. "CountOf_customer_id" - Represents the count of unique customer IDs for each combination of first name and number of orders.
3. "first_name" - The first name of the customer.

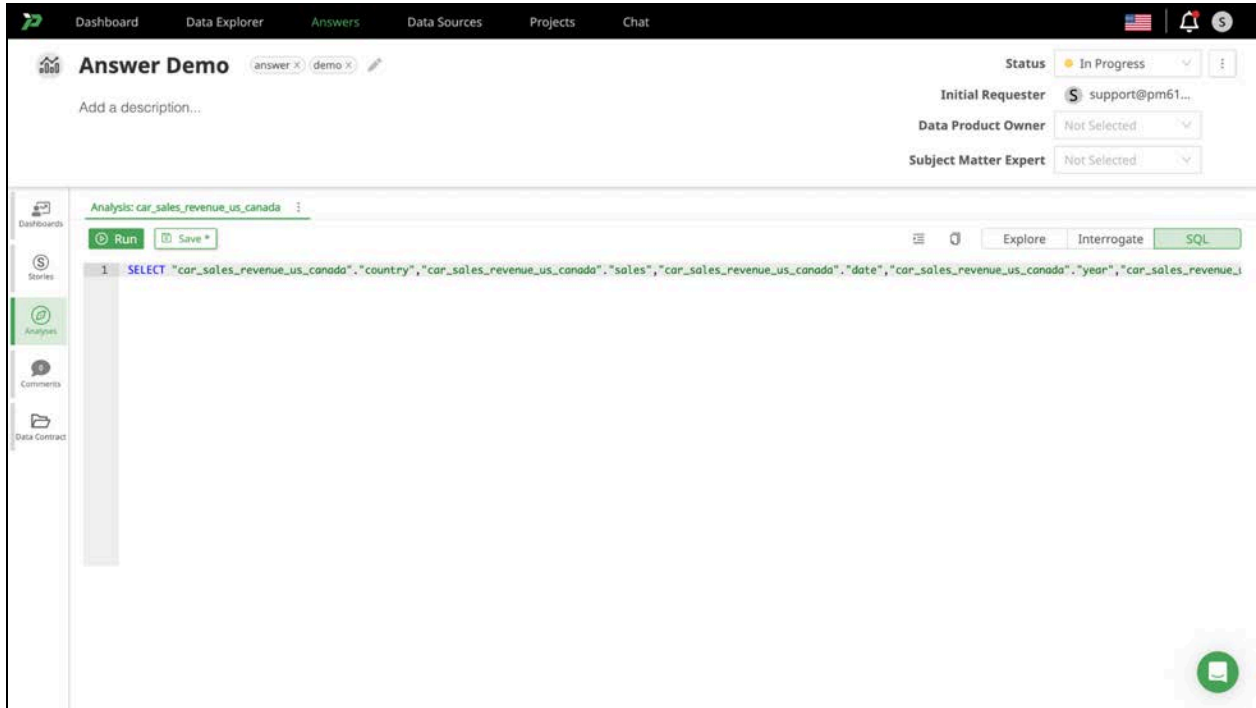
Dashboards, Analyses, Comments, Data Contract, Access Control



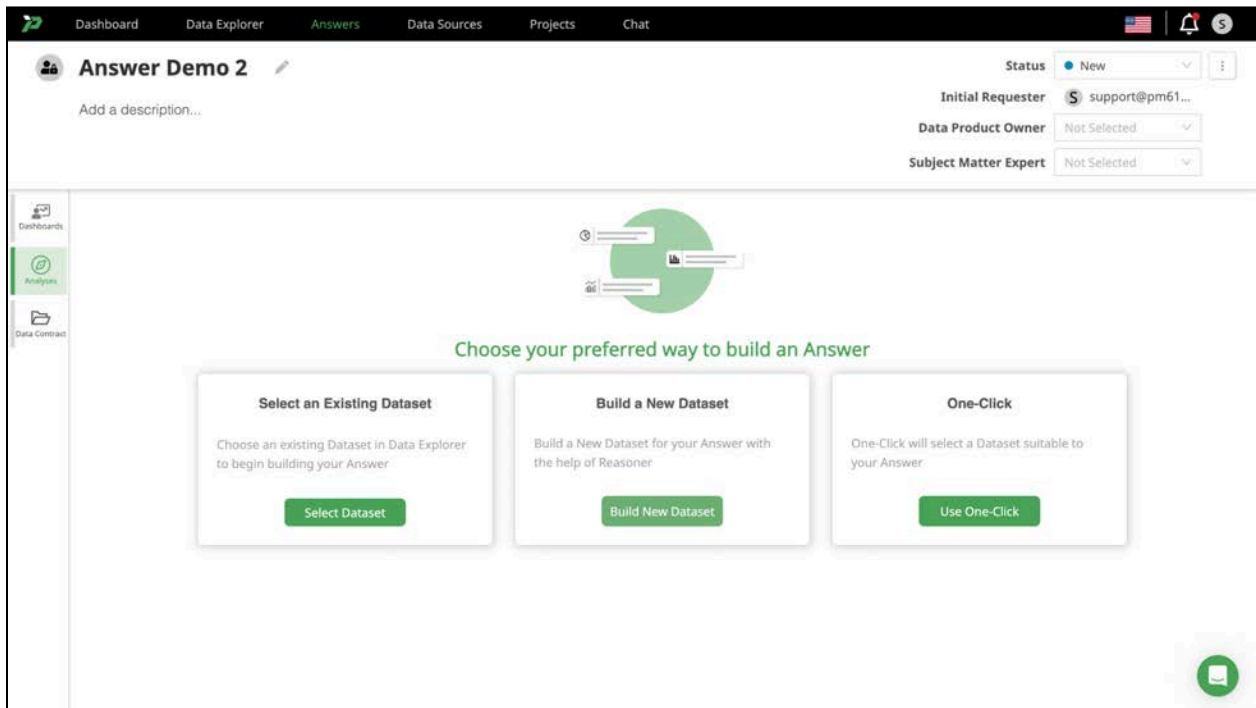
When you have a visualization appropriate to the Answer you can click Save and it will be added to your Dashboard.

Query

Clicking the Query button at the top right will switch to the Query mode. Swapping between modes can cause some customizations to be lost. It is far more common to lose customizations when moving from the advanced views to the less advanced e.g. from Query to Explorer. Ensure you have saved and added any critical analysis before switching modes to prevent having to recreate it.



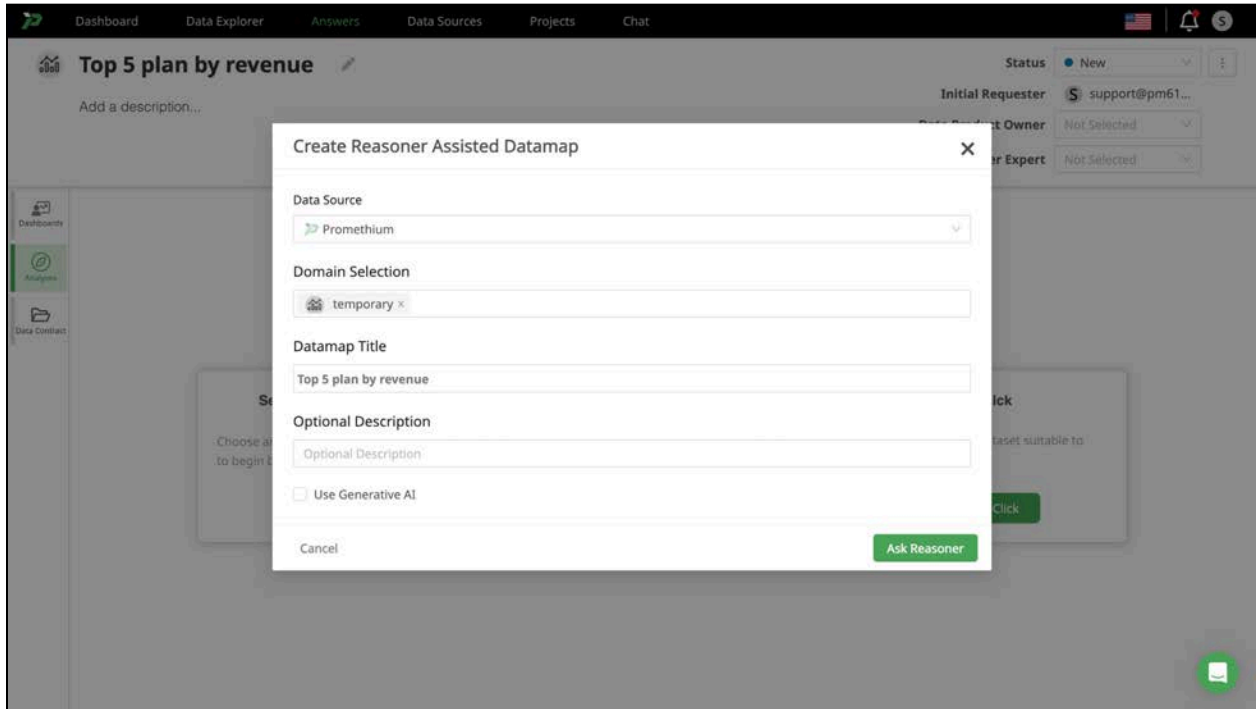
The Query view allows the user control over Source, Fields, Filters, Group By, and Order By. Each can be selected and edited by clicking Edit.



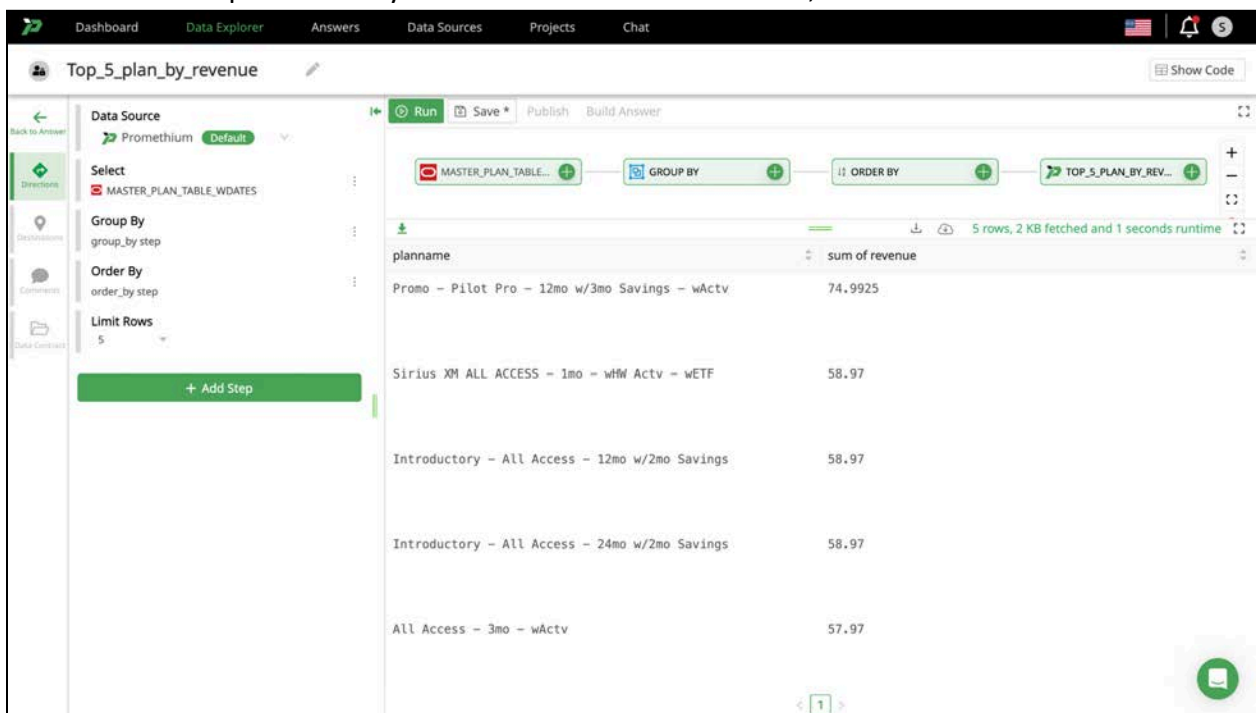


Build a New Dataset

When choosing to Build a New Dataset, it is important to remember when using Reasoner, Prometheus uses key phrases used in the **Datamap** title to create your datamap, including parameters when asking a question.

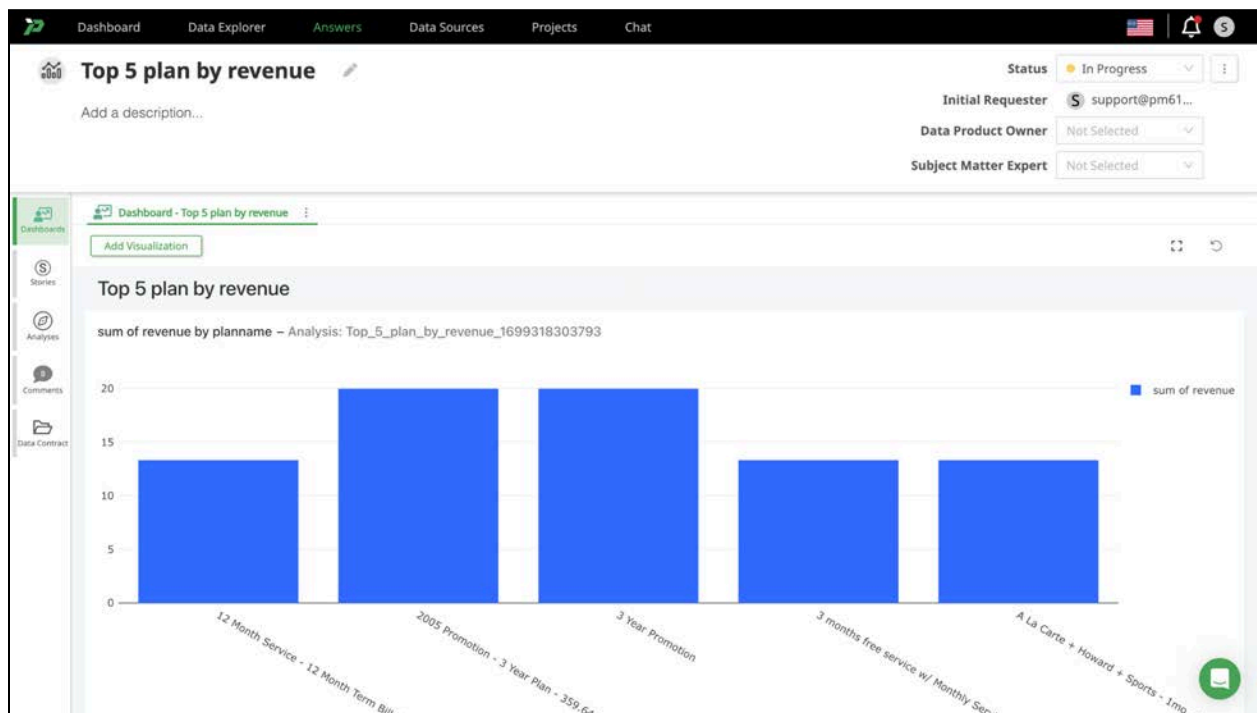


Once the Datamap is run and you are satisfied with the results, Save and Publish.



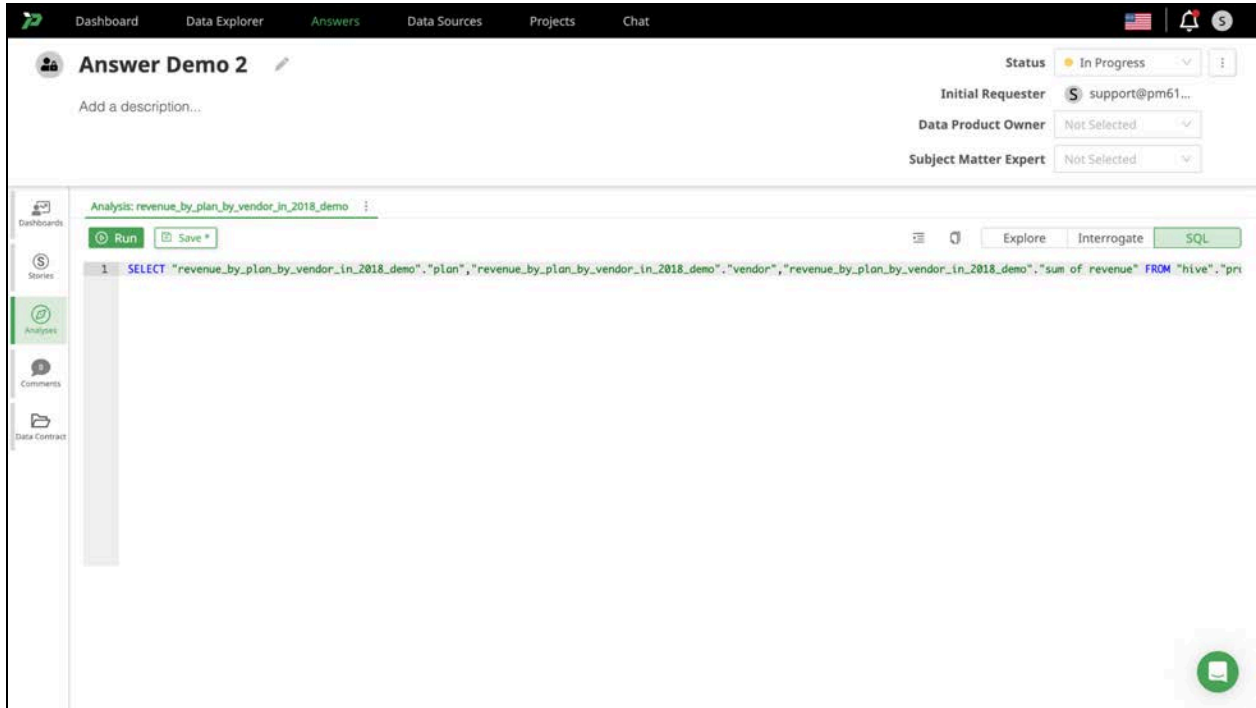


When you return to the Answer page, your dashboard will display visualizations based on the data.



SQL

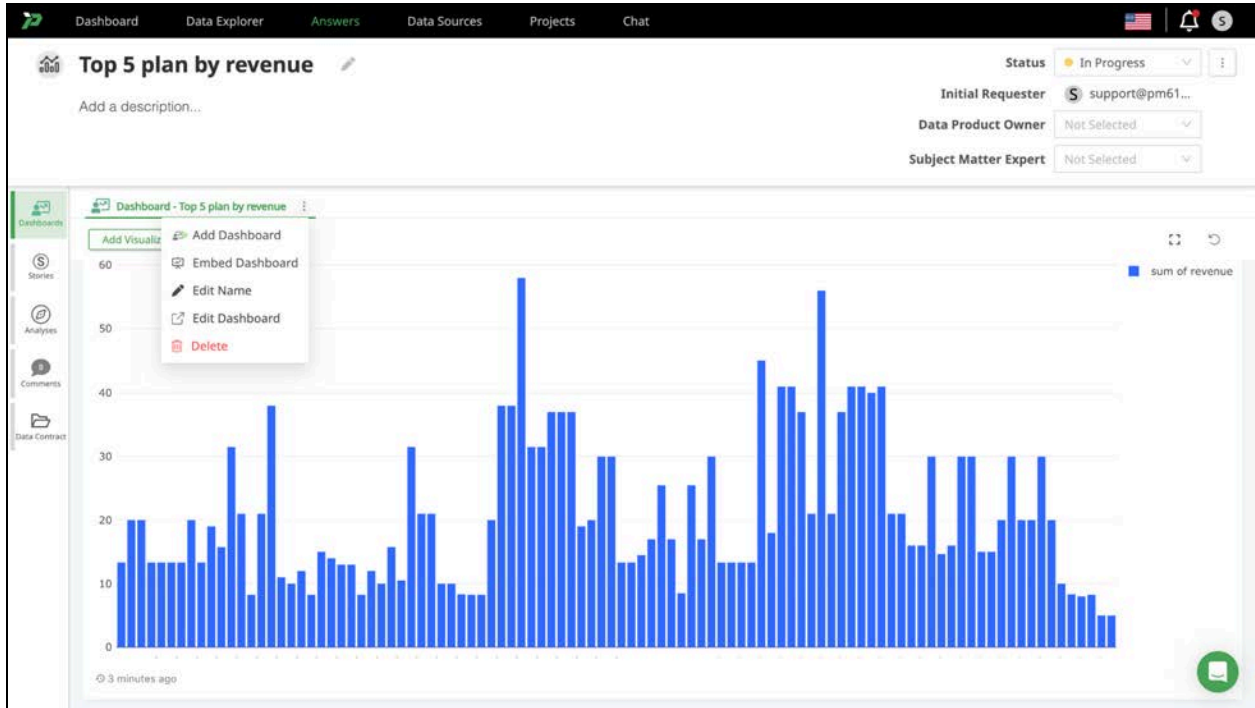
Clicking the SQL button at the top right will switch to the advanced SQL mode. Swapping between modes can cause some customizations to be lost. It is far more common to lose customizations when moving from the advanced views to the less advanced e.g. from Query to Explorer. Ensure you have saved and added any critical analysis before switching modes to prevent having to recreate it. The advanced SQL mode is a web based SQL IDE wherein you can modify and execute any changes you need to the SQL generated so far. The SQL shown below has been automatically generated through our use of the Explore and Query tabs so far.



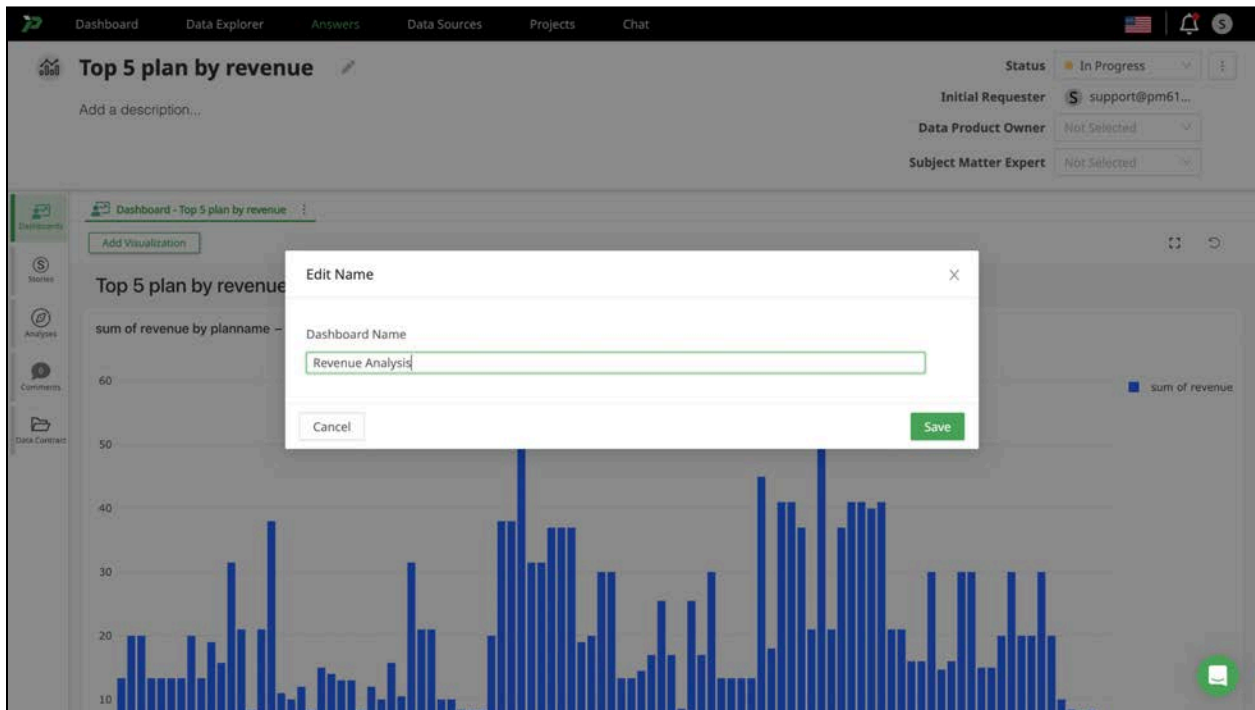
This view is intended for SQL experts. SQL is a powerful language and it's easy to encounter errors. Within this flow, SQL error messages are limited in detail. The example below was caused by the simple removal of the final word in the previous example 'ONLY'. You can modify your SQL as often as required and click the Run button to execute the query. When you're satisfied with the results, click Save.

Dashboard

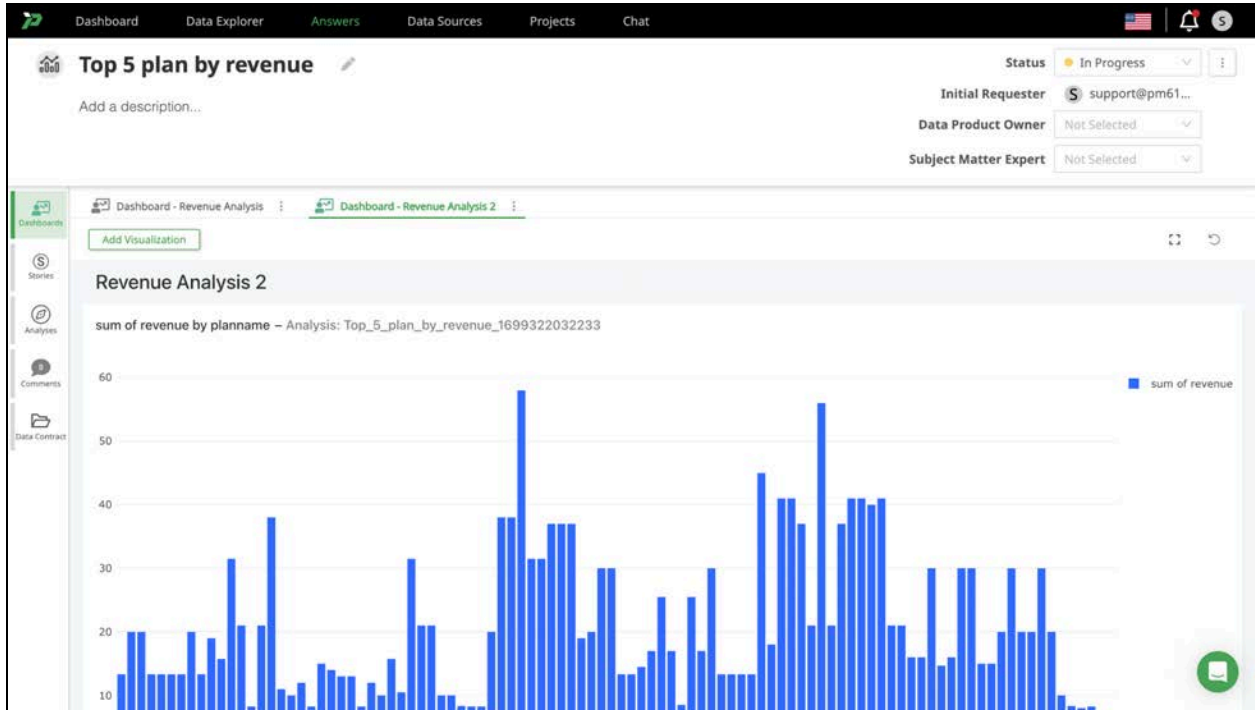
Everything you have saved from the Explore, Interrogate, or SQL flow will be available on the default dashboard available within the Dashboards tab on the left of the screen.



You can rename the default dashboard by selecting the three dots to the right of the name and choosing Edit Name

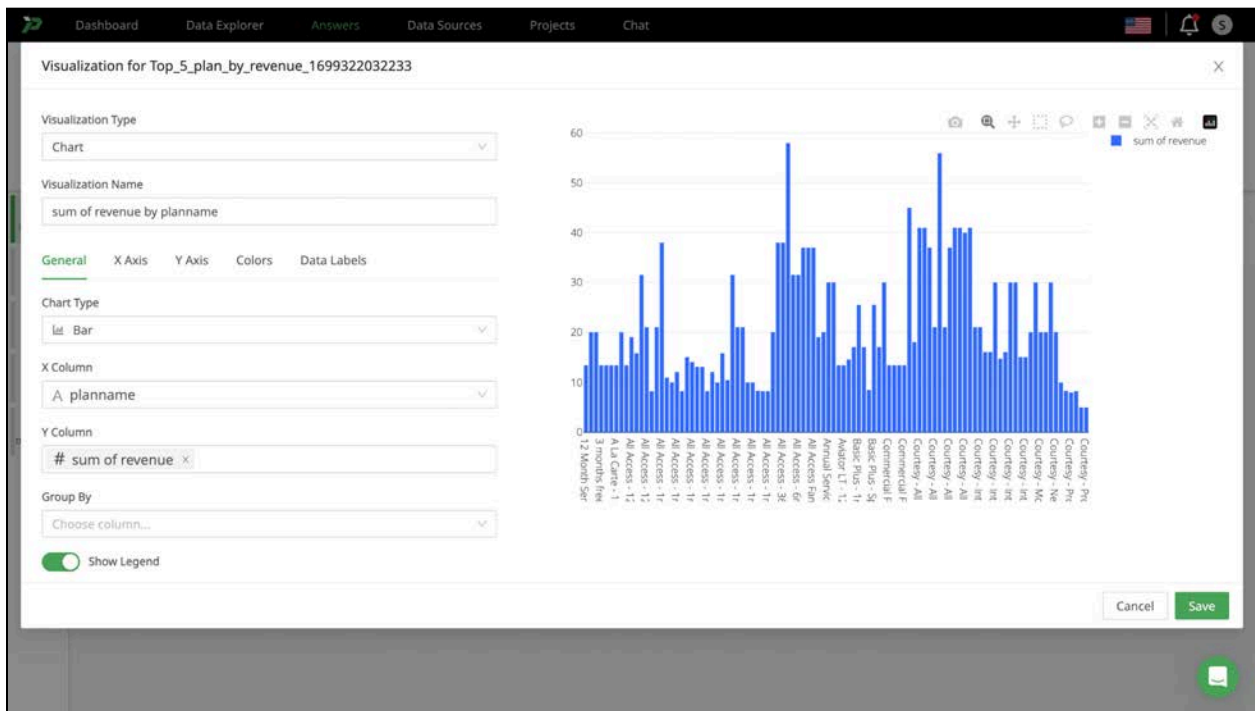


- Edit Dashboard opens the full Redash interface where you can create data alerts, modify the visual layout of your graphs, or create new dashboards directly.
- Add Dashboard creates a new Dashboard tab on which you can add additional visualizations. There is currently a limit of 5 dashboards per Answer.



- You can also choose to embed a 3rd party dashboard. See Embed Dashboard below for details.

From the Dashboards page, you can select **Add Visualization** to build a custom visualization (pivot table or graph) from data powering your dashboard.

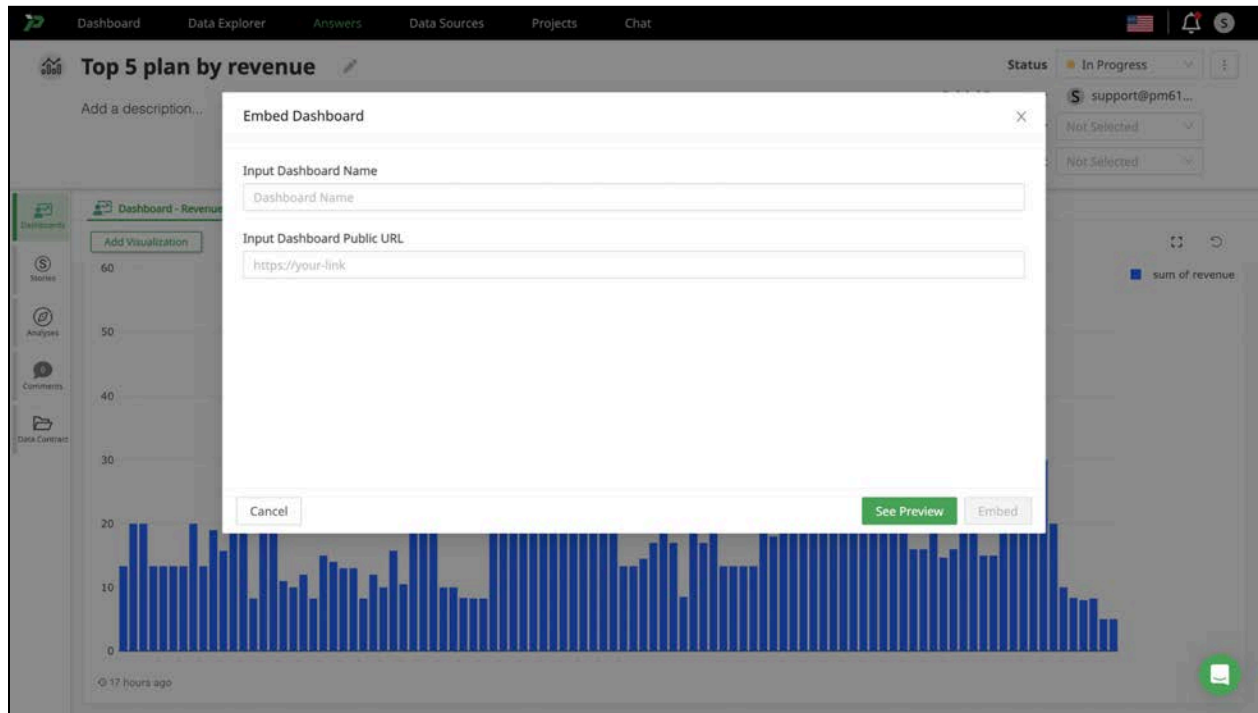




Embed Dashboard

If you have an existing dashboard or visualization created in a BI tool, such as Tableau, you can select Embed Dashboard and inform Promethium of the name and location.

The dashboard you connect must be available via a public URL and cannot require login to access. This link will be resolved and the dashboard rendered live to any future consumer of the Answer.



Chatbot

By selecting Chat in the top navigation, you will be directed to Promethium's Chatbot. This feature enables you to interact using natural language to find the right datamaps/answers, construct datamaps, and generate answers. You'll have the option to preview outcomes before finalizing your choice, and any products created by the Chatbot will be accessible on the Data Explorer or Answers page.



Dashboard Data Explorer Answers Data Sources Projects Chat

Reasoner Chat

hello, build a new datamap

Reasoner: Please ask Promethium chat-bot your question

build me a new dataset that will show me revenue by plan by agent from 2015 to 2018

Reasoner: Would you like to search for a datamap that can answer your question?

yes

Reasoner: There are multiple data maps matching your request please select the one you want to use.

- Select Preview Dataset revenue_by_plan_by_agent_from_2015_to_2018_1698792442549
- Select Preview Dataset average_revenue_by_plan_by_agent_by_discount_from_2018_to_2020_1698710997230
- Select Preview Dataset revenue_by_plan_by_vendor_from_2017_to_2019
- Select Preview Dataset What_is_revenue_by_plan_by_vendor_from_2018_to_2019_1698527094585
- Select Preview Dataset average_revenue_by_agent_by_discount_from_2018_to_2020_1698701626847

Ask a Question

Dashboard Data Explorer Answers Data Sources Projects Chat

revenue by plan by agent from 2015 to 2018

fact_subscription_activity

MASTER_PLAN_TABLE_WDATES

dim_agent

Run Save Republish Build Answer

100 rows, 20 KB fetched and 3 seconds runtime

plan	agente	sum of revenue
Fixed Demo - Select - 6mo - (Used Ca	Servicom	0
NavTraffic + NavWeather - 1mo - wActv	TeleServicesDirect	6.66
Promo - Select - Semi-Annual - (25% o	Servicom	11.24
Promo - Voyager + All Access - 6mo -	TeleServicesDirect	11.99
Promo - All Access - 6mo (49.99) - 1X	DialAmerica	8.3317



Dashboard Data Explorer Answers Data Sources Projects Chat

+ New Conversation

Find a Conversation...

Conversation 1

Reasoner Chat

Search Conversation...

Find me the top 5 plans with the highest revenue

Reasoner

There are multiple answers matching your request please select the one you want to use.

- Select Preview Answer revenue_by_plan_by_vendor_in_2018
- Select Preview Answer What_is_the_revenue_subscription_by_vendor_by_package_from_2013_to_2019
- Select Preview Answer Active Radio Subscriptions
- None are Applicable

Ask a Question



Dashboard Data Explorer Answers Data Sources Projects Chat

+ New Conversation Reasoner Chat Search Conversation...

Find a Conversation...

Conversation 1

Revenue_by_plan_by_vendor_in_2018 Preview

revenue_by_plan_by_vendor_in_2018

Analysis: revenue_by_plan_by_vendor_in_2018

plan	vendor	SumOf_sum of revenue
Promo - Voyager + All Access Family Friendly - 6mo - (Half-off) - 1X	DialAmerica	11.99
Long Term Plan - All Access - 5 yrs (WBOVCORP)	Results	3.32
Promo - All Access - Semi-Annual - (25% off)	Transparent	15.74
Fleet - All Access - 1mo - wActv - 13.49	Ipacesetters	13.49
Introductory - Pilot Preferred - 24mo w/2mo Savings - wActv	DialAmerica	15.99
Promo - All Access - 12mo - 119.88 - wETF (6mo50) - AMZ DOT-G3 - (OLPT)	Results	9.99
Promo - All Access - 6mo - 49.99 - wETF (6mo50) - AMZ DOT-G3 - (OLPT)	TeleServicesDirect	8.33
Promo - Select - 6mo - 54.00 - (\$9.00/mo)	Results	9.00
Promo - Select - Semi-Annual - (25% off)	Servicom	11.24
Promo - All Access - 12mo - 179.88	Results	14.99

Done

Dashboard Data Explorer Answers Data Sources Projects Chat

+ New Conversation Reasoner Chat Search Conversation...

Find a Conversation...

Conversation 1

Find me the top 5 plans with the highest revenue

Reasoner

There are multiple answers matching your request please select the one you want to use.

Select Preview Answer revenue_by_plan_by_vendor_in_2018

Select Preview Answer What_is_the_revenue_subscription_by_vendor_by_package_from_2013_to_2019

Select Preview Answer Active Radio Subscriptions

None are Applicable

revenue_by_plan_by_vendor_in_2018

Reasoner

You selected [revenue_by_plan_by_vendor_in_2018] answer. - Does this answer your question?

Ask a Question



Dashboard Data Explorer Answers Data Sources Projects Chat

revenue_by_plan_by_vendor_in_2018

Add a description...

Status: Published
Initial Requester: ravi@pm61 data...
Data Product Owner: Not Selected
Subject Matter Expert: Not Selected

Dashboard - revenue_by_plan_by_vendor_in_2018

revenue_by_plan_by_vendor_in_2018

Analysis: revenue_by_plan_by_vendor_in_2018

plan	vendor	SumOf_sum of revenue
Promo - Voyager + All Access Family Friendly - 6mo - (Half-off) - 1X	DialAmerica	11.99
Long Term Plan - All Access - 5 yrs (WBOVCCORP)	Results	3.32
Promo - All Access - Semi-Annual - (25% off)	Transparent	15.74
Fleet - All Access - 1mo - wActv - 13.49	lpacesetters	13.49
Introductory - Pilot Preferred - 24mo w/2mo Savings - wActv	DialAmerica	15.99
Promo - All Access - 12mo - 119.88 - wETF (6mo50) - AMZ DOT-G3 - (OLPT)	Results	9.99
Promo - All Access - 6mo - 49.99 - wETF (6mo50) - AMZ DOT-G3 - (OLPT)	TeleServicesDirect	8.33
Promo - Select - 6mo - 54.00 - (\$9.00/mo)	Results	9.00
Promo - Select - Semi-Annual - (25% off)	Servicom	11.24
Promo - All Access - 12mo - 179.88	Results	6.00



Dashboard Data Explorer Answers Data Sources Projects Chat

+ New Conversation

Find a Conversation...

Conversation 1

Reasoner Chat

Search Conversation...

build me a new dataset, revenue by plan by agent from 2015 to 2017

Reasoner: Would you like me to run automode for question [revenue by plan by agent from 2015 to 2017] and build new datamap?

yes

Reasoner: Here is the datamap I built for you. Shall I build an answer using this datamap?

Select Preview Dataset revenue by plan by agent from 2015 to 2017

This is not what I want

Reasoner: Here is your Answer:

Select Preview Answer revenue by plan by agent from 2015 to 2017

This is not what I want

Ask a Question

Dashboard Data Explorer Answers Data Sources Projects Chat

+ New Conversation

Reasoner Chat

Search Conversation...

Revenue By Plan By Agent From 2015 To 2017 Preview

revenue_by_plan_by_agent_from_2015_to_2017_1699377352058

Analysis: revenue_by_plan_by_agent_from_2015_to_2017_1699377352058

planname	agente	sum of revenue
Marine Offshore - 24mo w/3mo Free - wActv	Results	6.66
Promo - Traffic - 5mo - 10.00 (5FOR10)	DialAmerica	6.66
Mostly Music - 1mo - wActv	Results	9.99
Promo - Premier - 24mo	DialAmerica	6.66
Promo - Marine Offshore - 6mo - (Half-off) - 1X	TeleServicesDirect	14.99
Promo - Select Plus SXiR - 12mo - 89.00 - (Half-Off) - 1X	TeleServicesDirect	7.41
Promo - All Access - 6mo (49.99) - (OAC) - 1X	TeleServicesDirect	8.33
Long Term Plan - Select - 3 yrs	Servicom	2.75
Travel-Link - 12mo w/1mo Free	DialAmerica	6.66
News-Sports-Talk - 3mo	Results	10.99

Done



Dashboard Data Explorer Answers Data Sources Projects Chat

Reasoner Chat

build me a new dataset, revenue by plan by agent from 2015 to 2017

Reasoner: Would you like me to run automode for question [revenue by plan by agent from 2015 to 2017] and build new datamap?

yes

Reasoner: Here is the datamap I built for you. Shall I build an answer using this datamap?

Select Preview Dataset revenue by plan by agent from 2015 to 2017

This is not what I want

Reasoner: Here is your Answer:

Select Preview Answer revenue by plan by agent from 2015 to 2017

This is not what I want

Reasoner: Great, Promethium chat-bot is happy to help! How else can I help?

Ask a Question

Dashboard Data Explorer Answers Data Sources Projects Chat

revenue by plan by agent from 2015 to 2017

date x agent x reven... x See All

Status: In Progress

Initial Requester: support@pm61...

Data Product Owner: Not Selected

Subject Matter Expert: Not Selected

Dashboard - revenue_by_plan_by_agent_from_...

revenue_by_plan_by_agent_from_2015_to_2017_1699377352058

Analysis: revenue_by_plan_by_agent_from_2015_to_2017_1699377352058

planname	agente	sum of revenue
Introductory - Travel-Link + Traffic - 12mo w/1mo Savings - wActv	DialAmerica	6.66
Fleet - All Access - 12mo - wActv - 167.88	Results	13.99
Promo - Select - 6mo - 29.94SelfPay	DialAmerica	4.99
Aviator - 1mo - wActv	Results	6.66
Promo - Select Family Friendly - 6mo w/3mo Free - (Half-Off) - 1X	DialAmerica	7.00
Promo - Select - 6mo - 24.99 (6for25) - WBFB	TeleServicesDirect	4.17
Commercial Family - 1mo - (DSTR) - wActv	DialAmerica	6.66
Marine Offshore - 6mo - wActv	TeleServicesDirect	6.66
Commercial Family - 1mo - (D2B) - wActv	Results	6.66
Mostly Music - 3mo	Convergys	6.66
All a Carte Child - 1mo - wActv	Results	6.66

