
Aginity-Pro Documentation

Release 1

Jeff Schlitt

Dec 11, 2019

Contents

1	Installing Aginity Pro	3
1.1	Mac Install	3
1.2	Windows Install	4
1.3	Opening for the First Time	4
2	Configuring Connections in Aginity Pro	7
2.1	Downloading Drivers	8
2.1.1	Automatic Download	8
2.1.2	Manually Add Drivers	9
2.2	Setting up Netezza Connections	9
2.2.1	Using SSL with Netezza Connections	11
2.2.2	Connection Timeout	11
2.2.3	Advanced Parameters	11
2.3	Setting up Redshift Connections	12
2.3.1	Using SSL Mode with Redshift Connections	12
2.3.2	Advanced Parameters	13
2.4	Setting up Snowflake Connections	13
2.4.1	Warehouse	14
2.4.2	Role	14
2.4.3	Advanced Parameters	15
2.5	Setting up Hive Connections	15
2.5.1	Picking Subprotocol	17
2.5.2	Setting up Kerberos connections	17
2.5.3	Advanced Parameters	17
2.6	Setting up SQL Server Connections	17
2.7	Setting up DB2 Warehouse Server Connections	18
2.7.1	Advanced Parameters	19
2.8	Setting up Greenplum and Postgres Connections	19
2.8.1	Choosing Subprotocol (Only for Greenplum)	20
2.8.2	Advanced Parameters	21
2.9	Import and Exporting Connections	21
3	Using the Database Explorer in Aginity Pro	25
3.1	Viewing the Database Data Dictionary	25
3.2	Database Object Support by Platform	26
3.3	Quick Actions from Explorer	27
3.3.1	SQL Statement Generation	27

3.3.2	Selective SQL Generation on a Set of Columns	28
3.3.3	Drag and Drop Capabilities	29
3.3.4	Copy JDBC Connection String	29
4	Using the Query Panel in Aginity Pro	31
4.1	Writing your First Query	31
4.2	Executing to a File	32
4.3	Using Parameters in your Query	33
4.4	Auto-Complete Capabilities	34
4.5	Snippet support	35
4.6	Common Actions in Query Panel	35
4.6.1	Renaming Tabs	35
4.6.2	Commenting SQL	35
4.6.3	Keeping Connection	36
4.6.4	Find and Replace	36
4.7	Interacting with Query Results	36
4.7.1	Reordering the Result Grid	36
4.7.2	Sorting the Result Grid	37
4.7.3	Filtering the Result Grid	38
4.7.4	Exporting results to CSV	39
5	Managing Files in Aginity Pro	41
5.1	Working with File Projects	42
6	Using the Catalog in Aginity Pro	47
6.1	Catalog Overview	47
6.2	Calling Catalog Items from SQL	52
6.2.1	Single SQL SELECT Statement	53
6.2.2	Using Catalog Item to Facilitate Relationships	53
6.2.3	Creating and Reusing Calculations	53
6.2.4	Executing a Multi-Step script	54
6.2.5	Using Parameters in Catalog	54
6.3	Adding Items to Catalog	55
6.4	Import and Export Items from Catalog	56
6.4.1	Export Steps	56
6.4.2	Import Steps	57
7	Using Discovery Panel in Aginity Pro	61
7.1	Searching Query Execution History	61
7.2	Searching the Amp Pro Catalog	62
7.3	Expanding the results	62
7.4	Doing More with the Discovery Panel	62
8	Using Tools in Aginity Pro	63
8.1	Using the Data Import Wizard	63
8.1.1	Choosing a File - Step 1	63
8.1.2	Redshift Data Upload - Step 2	64
8.1.3	Choosing Columns and Rows - Step 3	65
8.1.4	Choosing Columns to Include - Step 4	66
8.1.5	Configuring Output Columns - Step 5	67
8.1.6	Setting Output Details and Loading - Step 6	68
9	Reference	71
9.1	Video Reference	71
9.2	Shortcut Key Reference	71

9.3	Frequently Asked Questions	71
9.3.1	What Control do I have with Row Limit Preference?	72
9.4	SQL Reference Guides for Supported connections	72
9.5	Driver Support Reference	72
10	Troubleshooting	73
11	End User License for Aginity Pro	75

Aginity Pro is our new SQL analysis tool. With cool new features and functions such as an active personal catalog, smart SQL discovery assistant, and more—beware, you may get spoiled. For more awesome features you can expect to see in Aginity Pro, check out the [feature roadmap](#).

This documentation provides you with conceptual, procedural, and reference information to help you install and use Aginity Pro and find answers to questions you may have about it.

The documentation consists of the following sections:

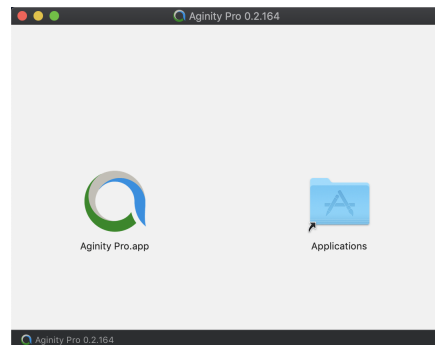
CHAPTER 1

Installing Aginity Pro

The following sections will discuss installing Aginity Pro on both Windows and Mac. To start download the software if you have not from the [Aginity Web Site](#)

1.1 Mac Install

To start the Mac install click on the Pro “dmg” file you downloaded.



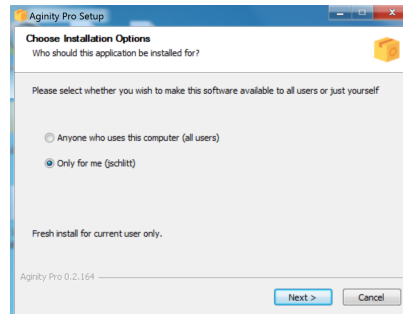
Drag the Aginity Pro.app icon into your Applications Directory. Once complete you can open the application from your Launchpad icon.




1.2 Windows Install

To start the Windows install click on the Pro “exe” file you downloaded.

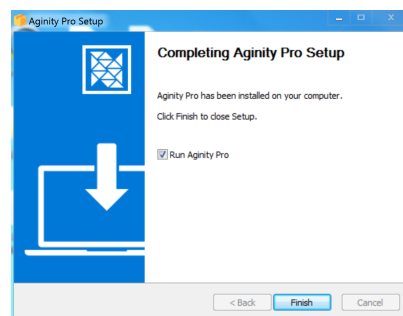
You will be prompted to first choose whether to install as **Only for me** or **Anyone who uses this computer (all users)** on the machine. If you are installing the application on a VM or other shared infrastructure you should choose all users.



Secondly, pick the installation directory for the software.

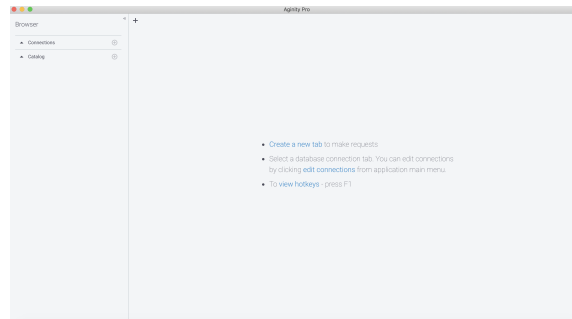


Once complete you should be able to launch Aginity Pro.



1.3 Opening for the First Time

Once you have opened the application you will be presented with a clean instance of Aginity Pro.



- Create a new tab to make requests
- Select a database connection tab. You can edit connections by clicking **edit connections** from application main menu.
- To view **hotkeys** - press F1

Once you open Aginity Pro you will commonly do these Actions.

See also:

Add Connections

Configuring Connections in Aginity Pro

Aginity Pro allows you to set up connections to a variety of different database platforms. In order to do this, click the + symbol next to the “Connections” header in the left navigation pane of Aginity Pro. You will be presented with a list of existing connections, as well as options for adding another connection. Clicking the “New Connection” button will allow you to enter the necessary information for the connection you wish to create.

The screenshot displays the 'CONNECTIONS' section of the Aginity Pro interface. On the left, a list of existing connections is shown, including 'Castor Hive', 'Corp Dev RS Redshift', 'Delivery POC Aginity Enterprise', 'Demonstration Databases ... Hive', 'GCP - Ascena Sandbox Aginity Enterprise', 'Snowflake Prod Snowflake', 'TF12 Netezza', and a 'New connection * Aginity Enterprise' at the bottom. A 'New Connection' button with a plus icon is located at the top of this list. To the right, a modal dialog box for creating a new connection is open. It contains the following fields: 'Connection name' (with a 'New connection' placeholder), 'Database provider' (set to 'Aginity Enterprise'), 'Host (required)' (with 'Enter host' placeholder), 'User (required)' (with 'Enter username' placeholder), 'Password (required)' (with 'Enter password' placeholder), and 'Schema' (with a 'Select Schema' dropdown). At the bottom of the dialog are four buttons: 'Remove' (red), 'Undo changes' (grey), 'Save' (blue), and 'Test' (grey).

Below is a quick start video of creating connections in Aginity Pro

Note: Aginity Pro uses JDBC to connect to database platforms. It is recommended you refer to the JDBC guides for our supported platforms to learn about specific variables to use within the **Advance Parameters** section of Aginity Pro. We will provide links to that documentation where available.

The video below provides a more detailed explanation of how to use advanced parameters to illicit the desired behavior and successful connection to the database platform.

2.1 Downloading Drivers

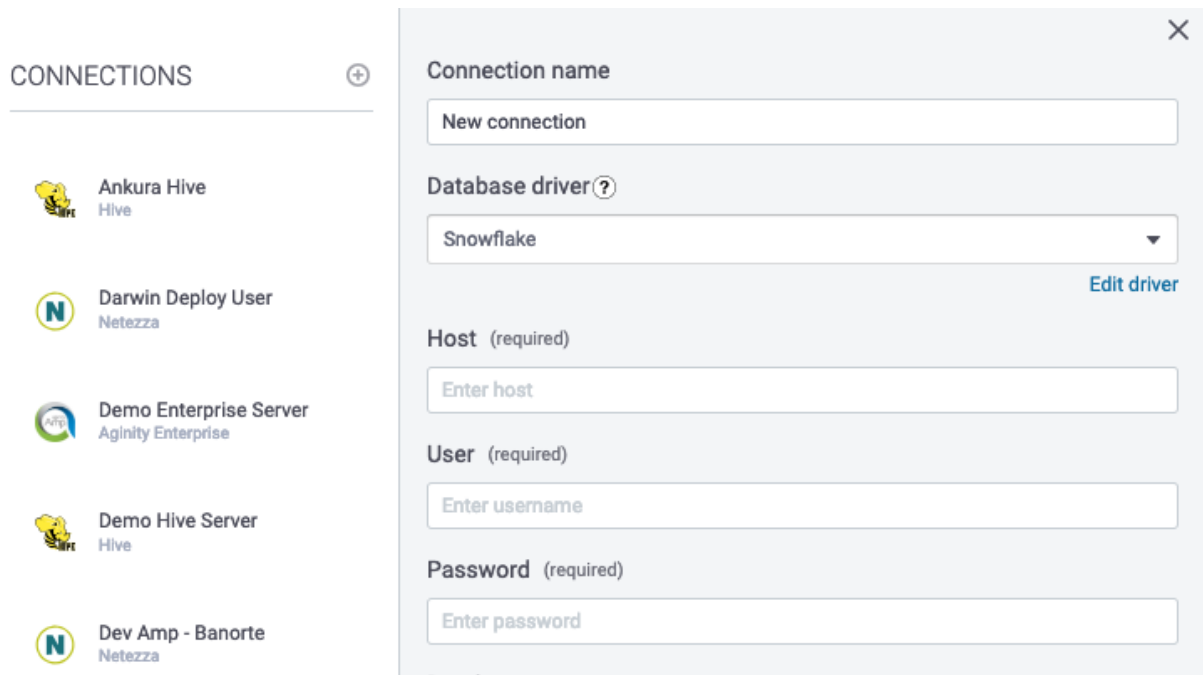
The first step in configuring a connection is to download the JDBC driver for the platform you will be using. In most cases, we will provide access to the drivers which have been certified to work with Aginity Pro. In some cases such as when the vendor does not allow us to distribute or you want to use a different version, you will have to manually add the JDBC driver.

The video below provides a quick illustration of how to manage drivers.

See also:

Database Driver Reference Guide

To get started, in the Edit or Add Connection box you will click on **Edit Drivers** as shown below

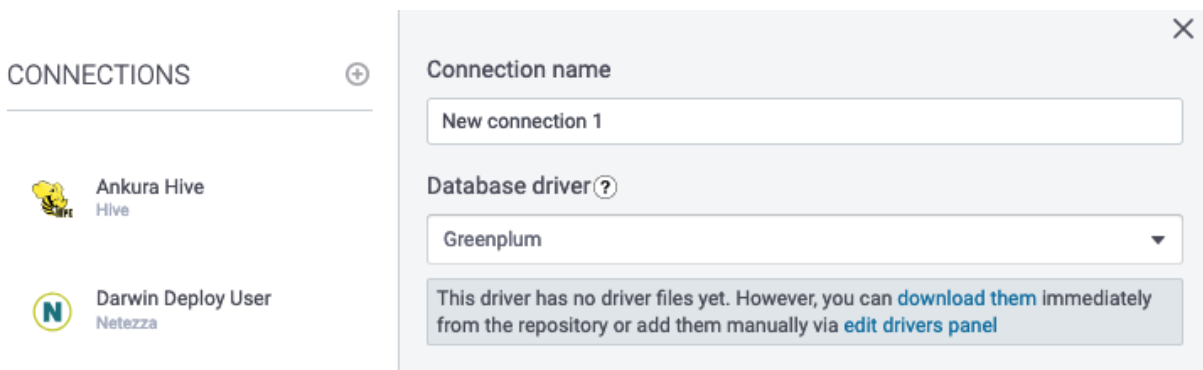


The screenshot shows the 'Edit Drivers' dialog box in the Aginity Pro interface. On the left, a sidebar titled 'CONNECTIONS' lists several connections: Ankura Hive (Hive), Darwin Deploy User (Netezza), Demo Enterprise Server (Aginity Enterprise), Demo Hive Server (Hive), and Dev Amp - Banorte (Netezza). The main dialog box has a close button (X) in the top right corner. It contains the following fields and options:

- Connection name:** A text input field with the value 'New connection'.
- Database driver:** A dropdown menu with 'Snowflake' selected. A question mark icon is next to the label. An 'Edit driver' link is visible to the right of the dropdown.
- Host (required):** A text input field with the placeholder 'Enter host'.
- User (required):** A text input field with the placeholder 'Enter username'.
- Password (required):** A text input field with the placeholder 'Enter password'.

2.1.1 Automatic Download

As shown in the image below, if we are allowed to download drivers you will be provided a link to do so and then be presented with one or more drivers to use and your done.



The screenshot shows the 'Edit Drivers' dialog box in the Aginity Pro interface. On the left, a sidebar titled 'CONNECTIONS' lists two connections: Ankura Hive (Hive) and Darwin Deploy User (Netezza). The main dialog box has a close button (X) in the top right corner. It contains the following fields and options:

- Connection name:** A text input field with the value 'New connection 1'.
- Database driver:** A dropdown menu with 'Greenplum' selected. A question mark icon is next to the label.
- Message:** A light blue box with the text: 'This driver has no driver files yet. However, you can [download them](#) immediately from the repository or add them manually via [edit drivers panel](#)'.

In some cases you may get an error message if your organization does not allow you to download drivers (jar files) typically supplied through a Maven repository.

In this case please review the section below Manually Add Drivers.

2.1.2 Manually Add Drivers

If you have to manually add a driver it starts with first acquiring the JDBC jar for that platform and saving it somewhere on your machine. Once you acquire the JDBC jar you can upload it to Aginity Pro.

To do this you will hit the **Edit Driver or Edit Drivers Panel** section as shown above and you will be presented with the following screen.

You can name the Driver anything you want and provide a description.

The class name if not provided must be looked up typically in the JDBC reference manual from the platform vendor.


Finally if you click on the **Add File** button you will navigate to where you saved the Jar file and then click **Save Changes**. You are now ready to configure your connection.


2.2 Setting up Netezza Connections

You can set up a connection to a Netezza (IBM PureData for Analytics) appliance here. Simply populate the requested data into the form.

Connection name

Dev Amp - Banorte

Database driver 

Netezza 

[Edit driver](#)

Host (required)

lnxsnz1d.dev.unix.banorte.com

Port

5480


User (required)

A3716616


Password (required)

.....

Database

SYSTEM 

Security Level



Only Unsecured 

☒ Autocommit

Connection timeout (sec)

120

Advanced Properties

Note: Netezza does not allow us to distribute their drivers so you must obtain it from your DBA or directly from IBM Fix Central.

2.2.1 Using SSL with Netezza Connections

Aginity Pro supports SSL when setting up a Netezza connection. While setting up your Netezza connection simply select the level of security to enable from the dropdown selector. Your choices include:

- **Only Unsecured** - Specify this option when you want an unsecured connection to the Netezza system. If the Netezza system requires a secured connection, the connection is rejected.
- **Preferred Unsecured** - This argument is the default value. Specify this option when you would prefer an unsecured connection, but you accept a secured connection if the Netezza system requires one.
- **Preferred Secured** - Specify this option when you want a secured connection to the Netezza system, but you accept an unsecured connection if the Netezza system is configured to use only unsecured connections.
- **Only Secured**- Specify this option when you want a secured connection to the Netezza system. If the Netezza system accepts only unsecured connections, or if you are attempting to connect to a Netezza system that is running a release before release 4.5, the connection is rejected.

See also:

For more information on SSL refer to the link [Netezza SSL](#)

2.2.2 Connection Timeout

The connection timeout will tell the session when to disconnect. Typically set this to 10 minutes or 3600 seconds.

2.2.3 Advanced Parameters

For Netezza you can supply additional JDBC parameters using the **Advanced Properties** section.

See also:

Here is the link to other configurable parameters supported within [Netezza JDBC's driver](#)

2.3 Setting up Redshift Connections

✕

Connection name

Database driver ?

Redshift
▼

[Edit driver](#)

Host (required)

Port

User (required)

Password (required)

Database

aginity_demo
▼

SSL Mode

Disable
▼

Advanced Properties

⊖

⊕

2.3.1 Using SSL Mode with Redshift Connections

Aginity Pro supports SSL when setting up a Redshift connection. While setting up your Redshift connection simply select the level of security to enable from the dropdown selector. Your choices include:

- **Disable** - SSL is disabled and the connection is not encrypted.
- **Prefer** - SSL is used if the server supports it. Amazon Redshift supports SSL, so SSL is used when you set

sslmode to prefer.

- **Allow** - SSL is used if the server requires it.
- **Require** - SSL is required.

See also:

For more information on Redshift SSL refer to this [link](#)

2.3.2 Advanced Parameters

For Redshift you can supply additional JDBC parameters using the **Advanced Properties** section.

See also:

Here is the link to other configurable parameters supported within [Redshift JDBC's driver](#)

2.4 Setting up Snowflake Connections

You can set up a connection to a Snowflake here. Simply populate the requested data into the form.

✕

Connection name

Database driver?

Host (required)

User (required)

Password (required)

Database

Warehouse

Role

Advanced Properties

⊖

⊕

2.4.1 Warehouse

A virtual warehouse, often referred to simply as a “warehouse”, is a cluster of compute resources in Snowflake. A warehouse provides the required resources, such as CPU, memory, and temporary storage, to perform the following operations in a Snowflake session. If you do not know which specific Warehouse to use speak to your Snowflake DBA.

2.4.2 Role

This is a security role that may have been setup for you to use. You can leave this field blank if you do not know what Role to connect as.

2.4.3 Advanced Parameters

For Snowflake most connection detail is handled through the command line additions to the JDBC URI. You can supply additional JDBC parameters using the **Advanced Properties** section.

See also:

Here is the link to other configurable parameters supported within [Snowflake JDBC's driver](#)

2.5 Setting up Hive Connections

You can set up a connection to a Hive databases here. Simply populate the requested data into the form.

Connection name

Demo Hive Server

Database driver?

Host (required)

demo-spark.aginity.com

Port

10000

Subprotocol (required)

Hive 2

User

jschlitt

Password

.....

Database

amp_demo_jschlitt

☒ Use Kerberos for authentication

Realm (required)

Enter realm

Kerberos Data Center FQDN (required)

Enter host fqdn

Service Name (required)

Enter service name

Advanced Properties

⊖

⊕

2.5.1 Picking Subprotocol

You can choose either Hive 2 or Spark SQL syntax when executing against Hive databases. You can choose either value in this drop down.

2.5.2 Setting up Kerberos connections

In a lot of cases organizations require the use of Kerberos with Hive so by clicking on the radio button for **Use Kerberos for authentication** you will be prompted to enter the following:

- Realm
- FQDN
- Service name

Note: The value of these parameters are typically set by your Hadoop administrator and it is best to work with them to obtain the appropriate values.

2.5.3 Advanced Parameters

You can supply additional Hive JDBC parameters using the **Advanced Properties** section.

See also:

Here is the link to other configurable parameters supported within [Hive JDBC's driver](#)

2.6 Setting up SQL Server Connections

You can set up a connection to Microsoft SQL Server databases here. Simply populate the requested data into the form.

✕

Connection name

Database driver (?)

MS SQL
▼

[Edit driver](#)

Host (required)

Port

User (required)

Password (required)

Database

AginityBilling
▼

Advanced Properties

⊖

⊕

Note: We only support SQL Server compatibility levels 150 and 140 which is SQL Server 2017 and greater.

2.7 Setting up DB2 Warehouse Server Connections

You can set up a connection to IBM DB2 warehouse databases here. Simply populate the requested data into the form.

CONNECTIONS

AG - Dev Environment
DB2 Warehouse Sandbox
Greenplum - Support Sandbox
Hive - Snoopy
Jeff Hive - Transient
MSSQL - SF Jira DB
MSSQL - Support Sandbox
Netezza Sandbox
Postgres - Support Sandbox
Redshift Aginity Demo
Redshift Snoopy
Rollins Production
Sample GBQ Connection
Snowflake Demonstration DB
WB Licensing Server

Connection name

DB2 Warehouse Sandbox

Database JDBC driver

DB2 Warehouse

Edit JDBC driver

Host (required)

54.89.3.141

Port

50001

User (required)

bluadmin

Password (required)

.....

Database

BLUDB

Schema

AGINITY_DEMO

SSL Connection

☒

Autocommit

☒

Connection timeout (sec)

Enter connection timeout (0 if left empty, means no timeout)

Advanced Properties

2.7.1 Advanced Parameters

You can supply additional Hive JDBC parameters using the **Advanced Properties** section.

See also:

Here is the link to other configurable parameters supported within [SQL Server JDBC's driver](#)

2.8 Setting up Greenplum and Postgres Connections

You can set up a connection to Greenplum and Postgres databases here. Simply populate the requested data into the form. Greenplum and Postgres share the same drivers so we will speak of them together.

Connection name

New connection

Database driver (?)

Greenplum

Edit driver

Host (required)

Enter host

Port

5432

Subprotocol (required)

Postgre SQL

User (required)

Enter username

Password (required)

Enter password

Database

Select database ("postgres" if left empty)

SSL Mode

Disable

Advanced Properties

⊖

⊕

2.8.1 Choosing Subprotocol (Only for Greenplum)

You can choose either generic Postgres or Pivotal Greenplum SQL syntax when executing against Greenplum databases. You can choose either value in this drop down.

2.8.2 Advanced Parameters

You can supply additional Postgres JDBC parameters using the **Advanced Properties** section.

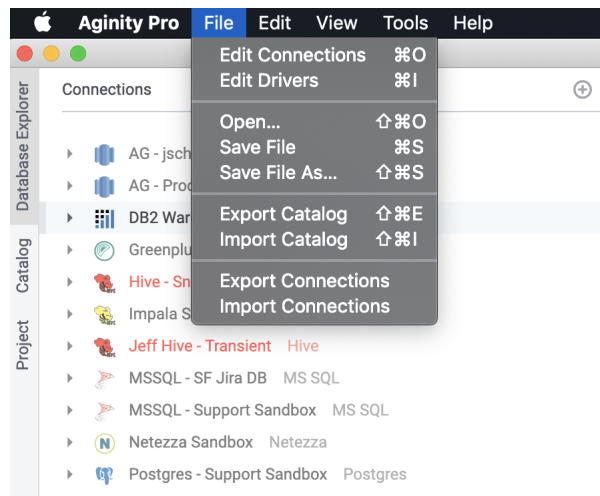
See also:

Here is the link to other configurable parameters supported within [Postgres JDBC's driver](#)

2.9 Import and Exporting Connections

A new feature in Aginity Pro is the ability to import and export your connection templates and share them with other users. A good deal of inefficiency with SQL IDE's is the time it takes to get connected. No longer. With the connection templates you can save all the relevant detail minus the login credentials and distribute them through an easy import and export feature.

To get started have a DBA or power user establish the connection details once and export the connection as shown in the image below.



You will be presented with a selection list of connections as shown below. Select the connections you want to export and click **Export selected**

Note: It is important to either check the radio button **include sensitive data** or to uncheck it. This will include the username and password or exclude it from the connection export.

EXPORT CONNECTIONS

AllNoneInvert

☒ AG - jschlitt73 Dev Environment
 ☐ AG - Prod Environment
 ☐ DB2 Warehouse Sandbox
 ☒ Greenplum - Support Sandbox
 ☐ Hive - Snoopy
 ☐ Impala Sample
 ☐ Jeff Hive - Transient
 ☒ **MSSQL - SF Jira DB**
☐ MSSQL - Support Sandbox
 ☐ Netezza Sandbox
 ☐ Postgres - Support Sandbox
 ☐ Redshift Aginity Demo
 ☐ Redshift Snoopy
 ☐ ROLN Production
 ☐ Snowflake Demonstration DB
 ☐ WB Licensing Server
 ☐ Yodlee Impala

Connection name

MSSQL - SF Jira DB

Host (required)

AWPD-WIN-SQL001.INT.AGINITY.COM

Port

1433

User (required)

reporting

Password (required)

Database

AginityBilling

Advanced Properties

☐ Include sensitive data

Cancel export

Export selected

You will be presented with a file menu window to name and pick the folder location for the export file as shown below.

Save As: connection_info_export_2019-12-10T14_

Tags:

<>≡📁

AG Queries

Search

Favorites

Box Sync

Jeff's Personal...

Applications

Desktop

Documents

Downloads

Pictures

Music

jschlitt

Previous 7 Days

ScoringGAWorkingFile.xlsx	Dec 2, 2019 at 12:03 PM	34 KB	Micros...k (.xlsx)
GA-ChemRad-Sample.csv	Dec 2, 2019 at 4:06 PM	272 KB	Comm...et (.csv)
MasterAnalytesw-Limits.xlsx	Dec 2, 2019 at 7:43 PM	13 KB	Micros...k (.xlsx)
SampleWaterScores3PWSid.xlsx	Dec 2, 2019 at 9:12 PM	25 KB	Micros...k (.xlsx)
SampleWaterScoresGA1252019.xlsx	Dec 4, 2019 at 9:32 PM	10 KB	Micros...k (.xlsx)
AG_Data Quality Check.sql	Dec 6, 2019 at 2:50 PM	327 bytes	TextWr...cument
AG - Load Facts.sql	Dec 7, 2019 at 11:07 AM	12 KB	TextWr...cument

Previous 30 Days

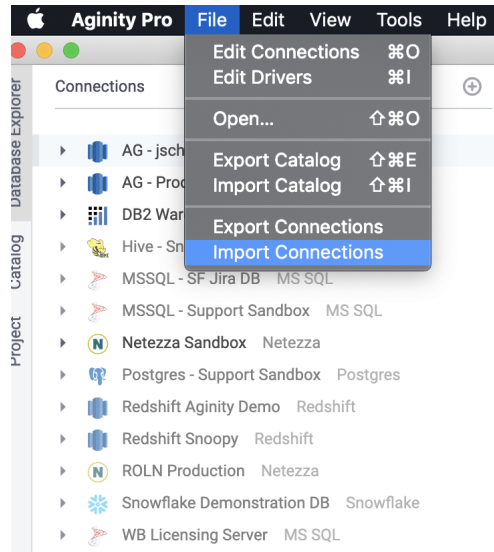
PCF - AG Final Appendix.docx	Nov 18, 2019 at 10:11 PM	22 KB	Micros... (.docx)
PCF - AG Final Appendix.pdf	Nov 18, 2019 at 10:12 PM	77 KB	PDF Document
AGDatabaseCreate.sql	Nov 19, 2019 at 8:18 AM	136 bytes	TextWr...cument
GAPWSids.csv	Nov 20, 2019 at 4:08 PM	3 KB	Comm...et (.csv)
DWW_Master_Scr...-ColiMicrobial.phj	Nov 21, 2019 at 8:12 AM	114 KB	Document

New Folder

Cancel

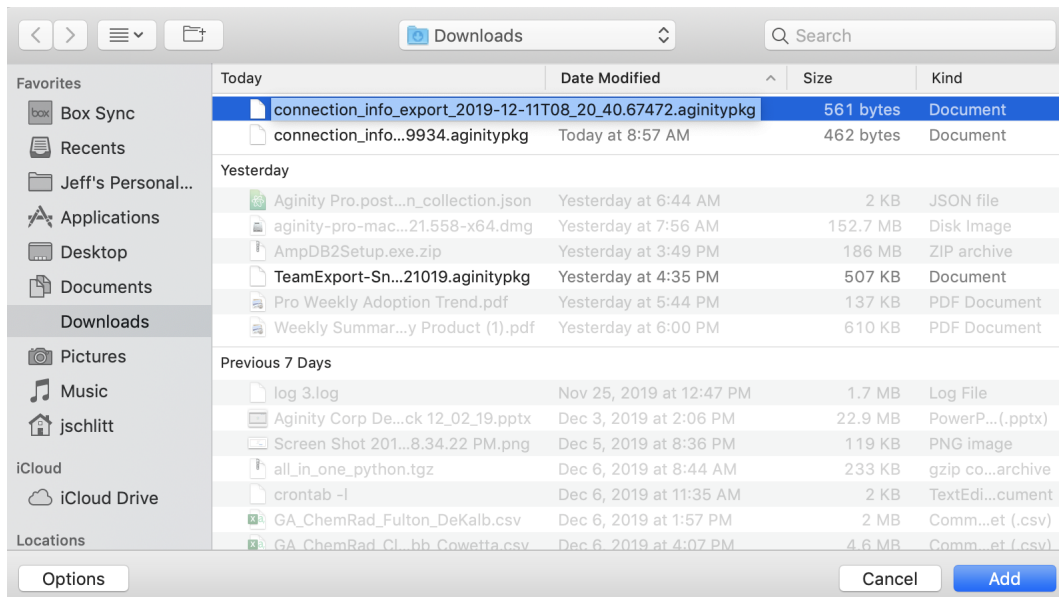
Save

Once you have exported the connections you can send them to Aginity Pro users and have them import them. To do that choose the **Import Connections** menu item as shown below.



You will be presented with a file selection. Choose the appropriate Aginity connection package.

Note: By default Aginity packages for both Catalog and Connections are prefaced with .aginitypkg file types.




Once you choose a file select the connections you want to import and choose **Import Selected**.

IMPORT CONNECTIONS

All

None

Invert

☒

Impala Sample (New)

Connection name

Impala Sample

Host (required)

https://something.com

Port

21050

Subprotocol (required)

Select Subprotocol

User

jschlitt

Password

Database

Select database ("default" if left empty)

Use Kerberos for authentication

☒

Realm (required)

aginity.com


Kerberos Data Center FQDN (required)


blaha

Service Name (required)

blah

Advanced Properties

sslEnabled	true	
------------	------	---



Cancel import

Import selected

See also:

Navigating Database Explorer


Using the Database Explorer in Aginity Pro

The Database Explorer panel in Aginity Pro can be toggled by using the **Database Explorer** and **Catalog** tabs on the left hand pane of the application as show below. The **Database Explorer** is used to navigate database objects based on the **Connections** defined in Aginity Pro. The **Catalog** browser is used to navigate SQL Analytics that you have personally stored.

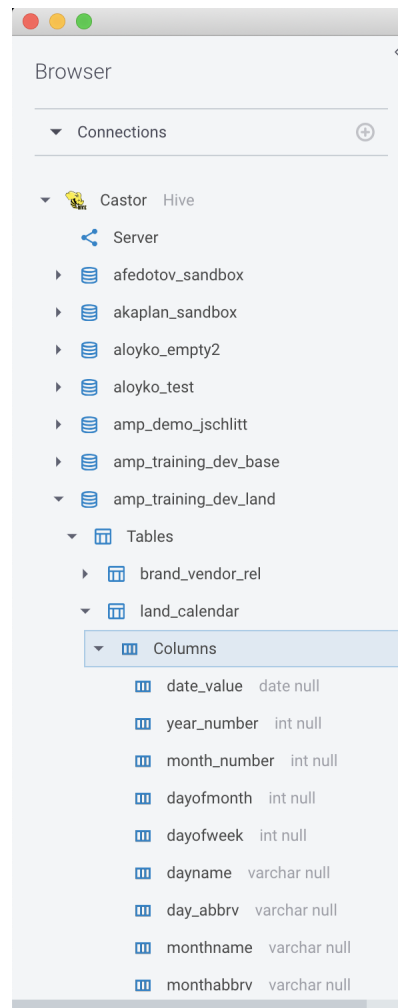
Below is a quick start video of using the Database Explorer in Aginity Pro

3.1 Viewing the Database Data Dictionary

Once you have established a connection to a database you can browse the database dictionary catalogs. You can toggle to the **Database Explorer** browser by clicking on the tab as shown below. The database data dictionaries typically contains the tables, columns, views and other database objects available to query from the query panel.



Once you have selected the **Database Explorer** tab you can drill into each connection as shown below.



3.2 Database Object Support by Platform

The following image details what objects are visible from the browser for all of our currently supported platforms. We do releases every two weeks with Aginity Pro so if there is a feature you need please contact us at support@aginity.com.

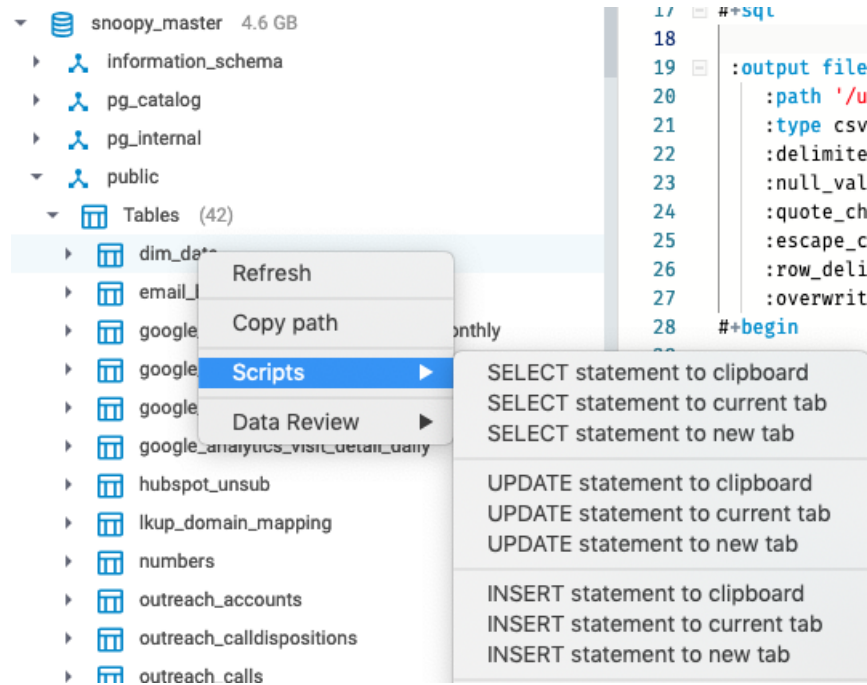
	Netezza	Hive	Redshift	Snowflake	Postgres	Greenplum	Db2	MSSQL	Azure DB	Azure ADW
Database	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
Schema	YES	NO	YES	YES	YES	YES	YES	YES	YES	YES
Table	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
External Table	YES	NO	YES	NO	NO	YES	NO	NO	NO	NO
System Tables	YES	NO	YES	YES	NO	NO	NO	NO	NO	NO
Columns	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES
View	YES	YES	YES	YES	YES	YES	NO	YES	YES	YES
Materialized View	YES	NO	N/A	NO	NO	N/A	YES	N/A	N/A	NO
System Views	YES	NO	YES	YES	NO	NO	NO	NO	NO	NO
Sort Key	N/A	N/A	YES	N/A	NO	N/A	NO	N/A	N/A	N/A
Cluster Key	N/A	N/A	N/A	NO	N/A	N/A	N/A	N/A	N/A	N/A
Distribution Key	YES	N/A	YES	N/A	N/A	NO	NO	N/A	N/A	NO
Constraints(PK, FK, Unique, Check, NOT NULL)	YES	NO	YES	PARTIAL (PK, FK)	NO	PARTIAL (PK, FK)	NO	YES	PARTIAL (PK, FK)	YES (PK, FK)
Index	N/A	NO	N/A	N/A	NO	NO	NO	YES	YES	YES
Functions (Table, Scalar, Aggregate)	YES	NO	YES	YES	NO	YES	NO	NO	NO	N/A
Procedures	YES	NO	YES	NO	NO	NO	NO	NO	NO	NO
Triggers	N/A	NO	N/A	N/A	NO	NO	NO	YES	YES	N/A
Sequence	YES	NO	YES	YES	NO	YES	NO	NO	NO	N/A
Users	YES	NO	YES	YES	NO	YES	NO	NO	NO	NO
Roles	N/A	NO	N/A	YES	NO	NO	NO	NO	NO	NO
Groups	YES	N/A	YES	N/A	NO	YES	NO	N/A	N/A	N/A
Warehouse	N/A	N/A	N/A	YES	N/A	N/A	N/A	N/A	N/A	N/A
Partition	N/A	NO	N/A	N/A	N/A	YES	N/A	N/A	N/A	N/A
External Schema	N/A	NO	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Share	N/A	N/A	N/A	Yes	N/A	N/A	N/A	N/A	N/A	N/A
Library	YES	N/A	YES	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Synonym	YES	NO	N/A	N/A	NO	NO	NO	YES	YES	N/A
Statistic	NO	NO	NO	NO	NO	NO	NO	YES	YES	NO
Custom Datatype	N/A	N/A	N/A	NO	NO	NO	NO	N/A	N/A	N/A
References	NO	NO	NO	N/A	NO	NO	NO	NO	NO	NO
Zone Map	YES	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Polybase	N/A	N/A	N/A	N/A	N/A	N/A	N/A	NO	N/A	NO
Linked Server	N/A	N/A	N/A	N/A	N/A	N/A	N/A	YES	N/A	N/A
Operators	NO	NO	NO	NO	NO	YES	NO	NO	NO	NO

3.3 Quick Actions from Explorer

Currently you can do the following quick actions from the **Database Explorer**

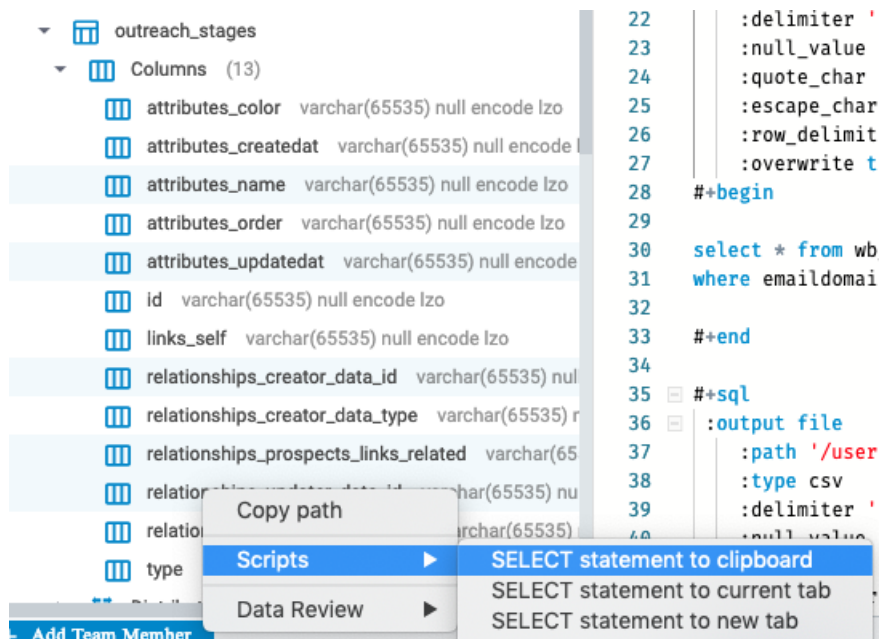
3.3.1 SQL Statement Generation

To generate typical SQL statements from tables or views you can right click on the object, navigate to scripts and choose the appropriate script as shown below.



3.3.2 Selective SQL Generation on a Set of Columns

To generate SQL statements based on a selection of columns within a Table or View, multi-select the columns by holding Ctrl (Windows) or Command (Mac) and selecting the columns as shown below. Then generate the SQL statement.



Hint: Make sure you don't have the table name selected and column names at the same time. If you do it will generate the SQL for the whole table.

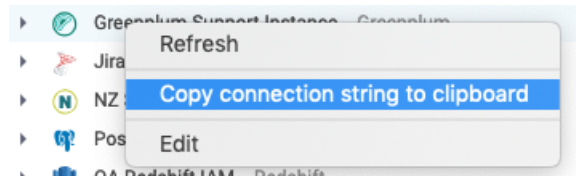
3.3.3 Drag and Drop Capabilities

You can drag and drop any object from the DB Explorer into the Query Panel

Note: Double-clicking expands the DB Explorer tree

3.3.4 Copy JDBC Connection String

If you select any Connection (Top level of the tree) and select “Copy Connection string to Clipboard” you will get the full JDBC URI which is helpful for debugging. See this [Knowledge Base Article](#) for help with debugging your Aginity Pro connections.



See also:


- *Use Catalog Browser*
- *Use Query Panel*

Using the Query Panel in Aginity Pro

You will spend most of your time in the Query Panel writing SQL and saving it to your catalog in Aginity Pro. The video below will give you an overview of creating your first query

4.1 Writing your First Query

Each tab in the Aginity Pro application can be set to different database connections. Once you establish that connection as shown below you can use the query panel like a text editor using SQL syntax specific to that database platform.



Once you write your code in Aginity Pro you have options to run it using the following scope:



Execution Type	Description
Run at cursor	When executing the query it will run the statement where the cursor is placed up until the semicolon
Run all parts of the script in a sequence	Each statement in the window will be run sequentially and the output of each statement will be shown in the output and if the statement was a SELECT then a results panel will be displayed
Run entire script as a batch	
Abort Query	After executing a query you can direct Aginity Pro to send a request to the database to cancel the query.

See also:

Note one of the greatest challenges in SQL is learning the different functional syntax from different database vendors. We are trying to make that easy by giving you links to the SQL reference manuals we support

4.2 Executing to a File

In version .17 of Aginity Pro we released the ability to use our **editor commands** to spool the results of a SQL statement(s) to a file(s).

To do this you must issue the following commands

1. Use the `#+sql` command to tell the editor you will be writing SQL in this section.
2. You will then use the `:output file` command to tell the editor you will spool the results to an output file
3. In the table below is a series of parameters and descriptions to be used when spooling
4. Use the `#+begin` statement to tell the editor you from here until the `#+end` command you will have SQL to execute.
5. Write your SQL
6. End with `#+end`
7. Use the Run commands as you would any normal SQL statement as described in the Run Scope section above.

The code below is an example you can use to build off of.

```

1  -- This will write the SQL statement between #+begin and #+end to the
2  ↪file specified.
3
4  #+sql
5  :output file
6  :path '/users/jschlitt/Downloads/jeff1'
7  :delimiter ','
8  :null_value 'NULL'
9  :quote_char '"'
10 :overwrite true
11
12 #+begin
13 select * from wb_pro_marketing_list
14 where emaildomain = 'aginity.com'
15 #+end

```

Note: We only support today the file type of csv. You can still change the delimiter but all files with have a csv extension

The table below illustrates the current values accepted in execute to file commands

Parameter	Options
path	On Mac this should be in the format “/dir/dir/filename”. On Windows it should be in the format “drive:directoryfilename”.
delimiter	You can use any value but we recommend , or or ^ or t for (tab)
null_value	enter a character field for null typically ‘’ (empty string) or ‘NULL’
quote_char	enter how to enclose character string typically the quotation (“)
overwrite	whether or not to allow overwrite if file exists: true or false
header	true or false on whether you want to display headers in file
encoding	the encoding of the file UTF_8 is the only value supported currently.
compression	we can compress the file after it is written. Values are NONE, ZIP and GZIP.

See also:

We have a great [knowledge base article](#) on how to use the catalog to facilitate using the export as file editor feature.

4.3 Using Parameters in your Query

Parameters can be used in Aginity Pro by using the \$ in any part of a SQL statement. When you execute a query the Aginity Pro application will ask you to fill in the value of the parameter before it sends the query to the database.

```

1 --sample parameter query
2
3 select *
4 from landing_dim_customer
5 where gender = $fill_in_gender_at_runtime
6 limit 100;
```

You will have datatype choices at runtime as shown in the table below.

Parameter Data Type	Action Taken
As Is	Aginity Pro will send the value exactly as written with no casting. This is often used when using the parameter to fill in database table or column names. See example 2 below.

Hint: You can specify default values with parameters by using this syntax with your parameter

```
$fill_in_gender_at_runtime{ 'M' }
```

When you run the query you will see the little gear icon as shown below. Clicking that will populate the default value.

Run query with parameters

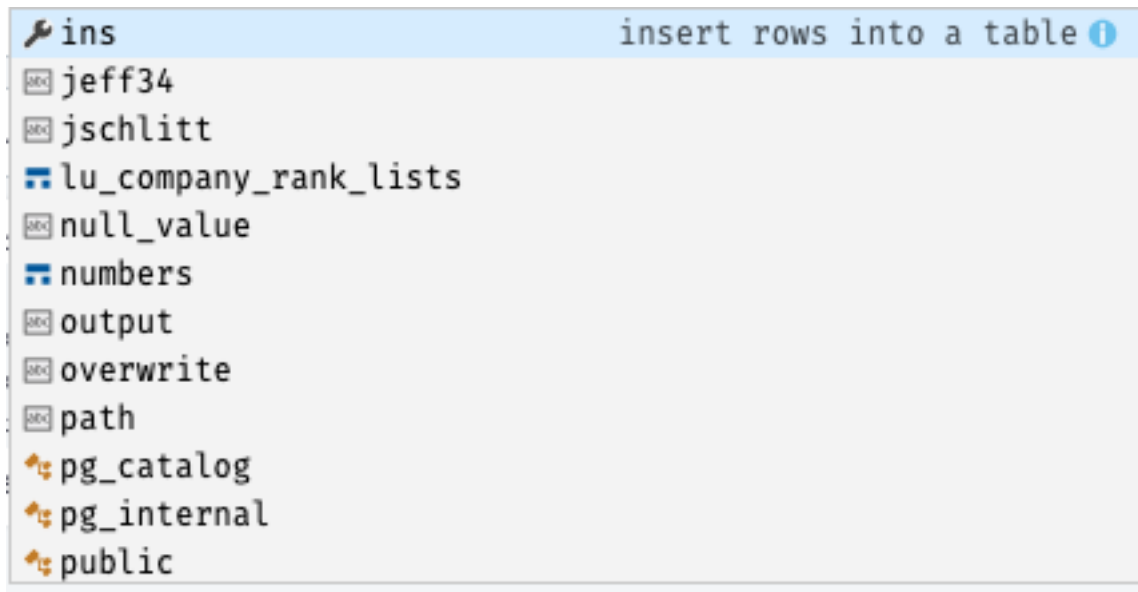
\$fill_in_gender_at_runtime
'M'
⚙

Run
Cancel

4.4 Auto-Complete Capabilities

In Aginity Pro we have implemented auto-complete on several database platforms and will be deploying across all over time. The auto-complete capability has three capabilities.

1. Display any words on the editor screen from above where your current cursor is
2. Display SQL Syntax like SELECT, WHERE, FROM all the way to database platform specific functions like CASE, DATEDIFF or CURRENT_DATE()
3. Display the database, schema, table, column as you are typing



Hint: To initiate just the database browser which shows database objects plus snippets hit the **Ctrl+Space** button otherwise as you type if auto-complete is supported it will automatically popup.

Note: Currently we support auto-complete on Netezza, Redshift and Snowflake

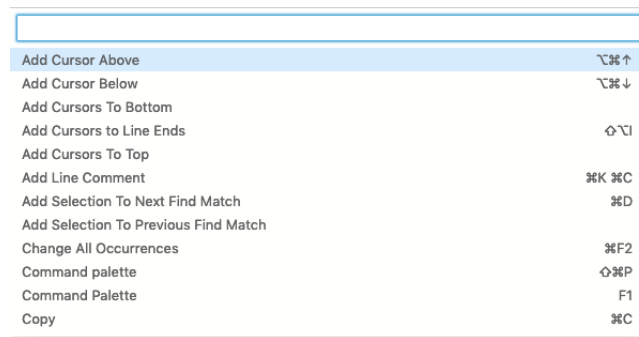
4.5 Snippet support

We have pre-wired some special snippets within the editor as well to make writing SQL simpler. The table below details those snippets that are currently available.

Snippet	Action
selall	will generate generic SELECT FROM
selwh	will generate generic SELECT FROM WHERE
upd	will generate an UPDATE statement
ins	will generate an INSERT statement
del	will generate a DELETE statement
expf	will generate a generic export to file statement


4.6 Common Actions in Query Panel

We will highlight common functions heavily used in the editor here. There is a complete list of available editor functions which can be accessed by hitting the **Command+Shift+P** on Mac or the **Ctrl+Shift+P** button within the editor or from the **View->Command Palette** menu item.



4.6.1 Renaming Tabs

You can rename the query tabs by double clicking and typing in a name as shown below.



4.6.2 Commenting SQL

In Aginity Pro you can use the typical SQL syntax to comment your code as shown below

```

1  -- This is one way to comment your SQL code
2
3  SELECT AVG(LineTotal)
4  FROM sales.SalesOrderDetail
5
6  /* This is another way to comment your SQL code */

```

Hint: If highlight multiple lines of code you can also use the Shortcut Keys **Ctrl+K+C** on Windows or **Command+K+C** on Mac to comment them out. To uncomment use **Ctrl+K+U** or **Command+K+U**.


4.6.3 Keeping Connection

The **Keep Connection** checkbox as show below will instruct Aginity Pro to maintain the connection to the database if the default behavior of the database is to disconnect after the execution of a query. You are still subject to an overall database connection timeout if you session connection exceeds the permitted database setting.

☐ Keep connection

4.6.4 Find and Replace

Ctrl+F (Windows) or Command+F (Mac) will bring up the find and replace box as shown below.



4.7 Interacting with Query Results

Once your query has been executed you can use the features of the result grid to interact with the result set. The typical actions you will do are listed below.

- Reordering the Result Grid
- Sorting the Result Grid
- Filtering the Result Grid
- Exporting results to CSV

4.7.1 Reordering the Result Grid

Once a query has executed successfully, you may rearrange the columns in the Result Grid in order to more effectively perform subsequent analyses. Rearranging the Result Grid is simple: simply click, hold, and drag the column header for the column you wish to move. As you move it left or right, the grid will automatically adjust to reposition the column.

🕒 Output	📊 Result 1	
landing_dim_channel.chann...	landing_dim_channel.chann...	landing_dim_channel.chann...
1	In-Person	In-Person (as in a store)
2	Mobile	Mobile (as in a mobile app)
3	Web	Web (as in a web page)
4	Call	Call (as in a phone call)
5	Chat	Chat (as in an instance mess...

Once you are satisfied with the new position of the column, you may release the drag and the column will lock into place.

🕒 Output	📊 Result 1	
landing_dim_channel.chann...	landing_dim_channel.chann...	landing_dim_channel.chann... ≡
1	In-Person (as in a store)	In-Person
2	Mobile (as in a mobile app)	Mobile
3	Web (as in a web page)	Web
4	Call (as in a phone call)	Call
5	Chat (as in an instance mess...	Chat

4.7.2 Sorting the Result Grid

After successfully executing a query, if the query has returned any rows you may sort the results in order to permit various types of analysis of your results.

🕒 Output	📊 Result 1	
landing_dim_channel.chann...	landing_dim_channel.chann...	landing_dim_channel.chann...
1	In-Person	In-Person (as in a store)
2	Mobile	Mobile (as in a mobile app)
3	Web	Web (as in a web page)
4	Call	Call (as in a phone call)
5	Chat	Chat (as in an instance mess...

In order to sort your results, simply click on the header row in the results grid. If the results are already being sorted by that column, then the sort order will toggle between ascending and descending.

🕒 Output	📄 Result 1	
landing_dim_channel.chann...	landing_dim_channel.ch... ↑ ≡	landing_dim_channel.chann...
4	Call	Call (as in a phone call)
5	Chat	Chat (as in an instance mess...
1	In-Person	In-Person (as in a store)
2	Mobile	Mobile (as in a mobile app)
3	Web	Web (as in a web page)

4.7.3 Filtering the Result Grid

Data in the Result Grid may be filtered to make it easier to locate and analyze certain subsets of the resultant data, without requiring you to submit an entirely separate query.

To filter the result grid by values in a particular column, bring your mouse cursor to the column header for the column to which you wish to apply a filter. In the right-hand side on the header row the filter icon (three horizontal bars) will appear. Click on the filter icon and you will be presented with a pop-up menu that will allow you to enter your filter criteria.

🕒 Output	📄 Result 1	
landing_dim_channel.ch... ↑	landing_dim_channel.chann... ≡	landing_dim_channel.chann...
1	In-Person	Contains store)
2	Mobile	Filter... bile app)
3	Web	Web (as in a web page)
4	Call	Call (as in a phone call)
5	Chat	Chat (as in an instance mess...

By default, the filter criteria will select a “Contains” filter, but several other filter types exist and you may select any of them. Once you have selected the filter type you wish to apply, type the filter into the text box. The results will immediately begin to filter.

🕒 Output	📄 Result 1	
landing_dim_channel.ch...	landing_dim_channel.ch...	landing_dim_channel.chann...
4	Call	Contains call)
5	Chat	C AND OR Contains Filter...

When a column is being filtered, a funnel icon will appear in its header. To remove the filter, you may re-enter the filter mode by clicking on the filter icon (three bars) in the header row, and simply delete the filter text.

Multiple columns may have active filters at the same time. To accomplish this, simply perform these steps on each column.

4.7.4 Exporting results to CSV

Once you have successfully executed a query you can elect to export those results in Comma Separated Value (CSV) format to a file outside of Aginity Pro. To do this, right-click anywhere in the results panel. Doing this will display a context menu from which you will select the option “Export to CSV”.

🕒 Output	📄 Result 1	
landing_dim_channel.chann...	landing_dim_channel.chann...	landing_dim_channel.chann...
1	In-Person	In-Person (as in a store)
2	Mobile	Mobile (as in a mobile app)
3	Web	Web (as in a web page)
4	Call	Call (as in a phone call)
5	Chat	Chat (as in an instance mess...

After selecting this option, you will be presented with a screen that permits you to select the location to which to save your results. Once you have identified the proper location, press the Save button to complete the export.

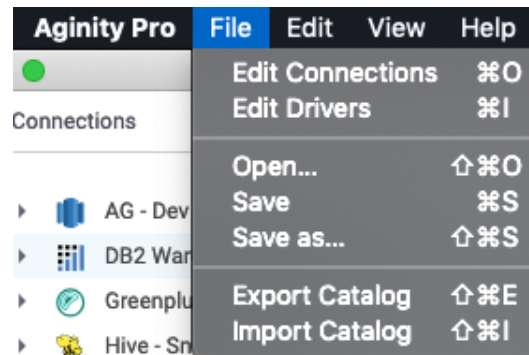
CSV files may subsequently be imported into a large number of different programs for further review and analysis.

See also:

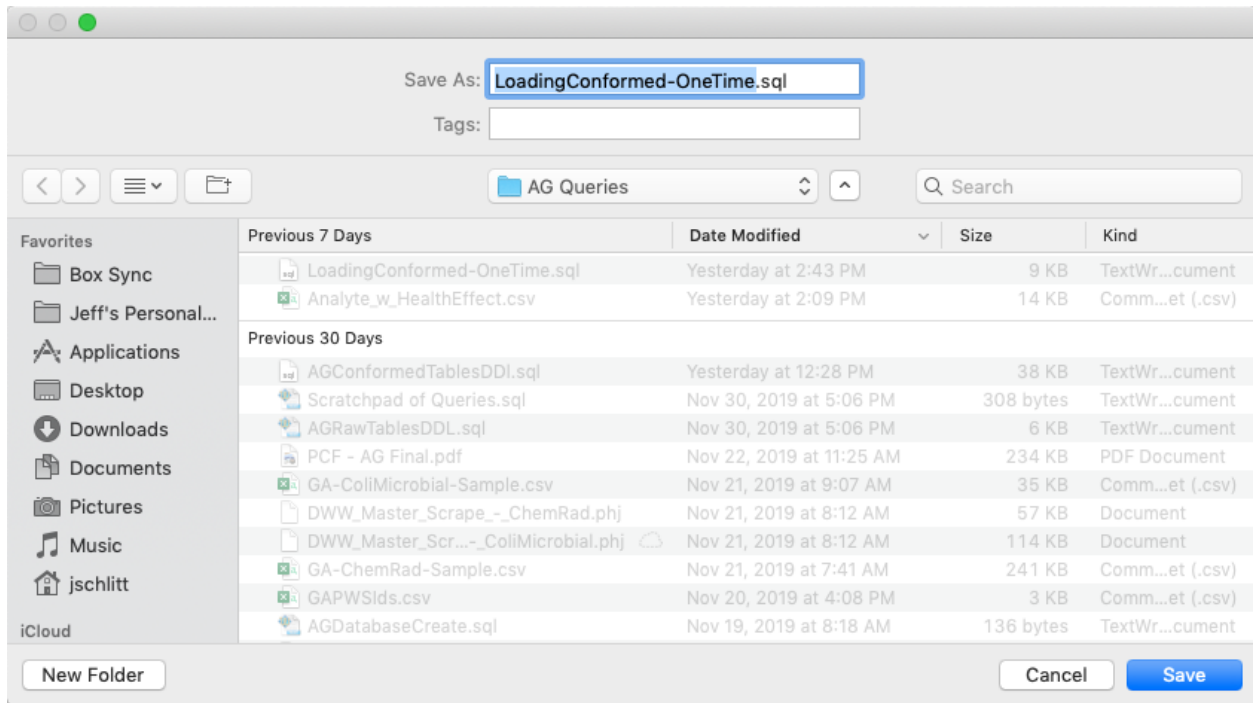
- [Adding Query to Catalog](#)
- [Using Discovery Panel](#)

Managing Files in Aginity Pro

You can now open files with Aginity Pro! You'll notice in the File menu commands for Open, Save and Save As. This commands will let you interact with files directly on your hard drive. By doing this you can now easily open your existing SQL files as well as start to govern your SQL in both files and in the Aginity Catalog. At Aginity we tend to use files to store SQL that we are in the process of developing and then use the Catalog as a place to store finished SQL assets for personal or collaborative use.



Once you click on a file menu item you will be presented with a typical file explorer for your O/S as shown below.

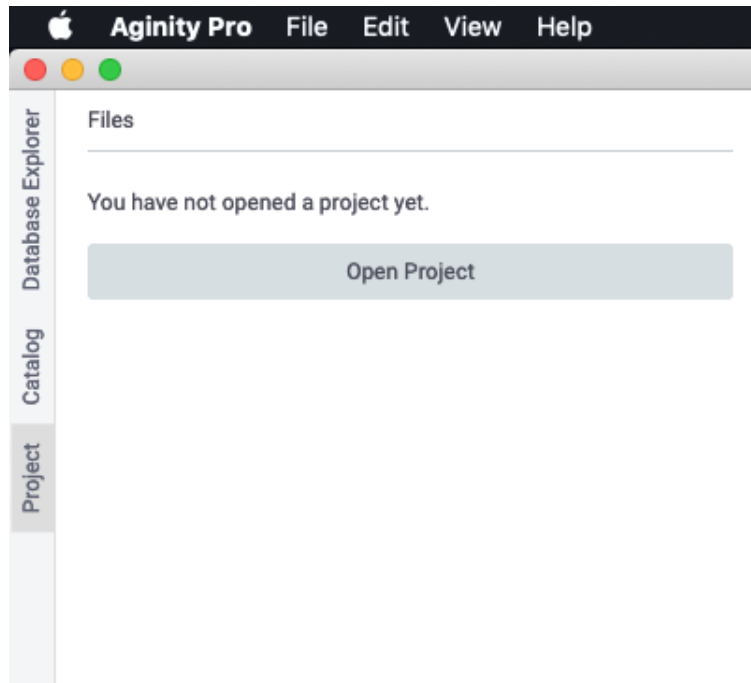


In the Query Panel the tab header will display with a symbol as shown below to let you know you are working with a file. Note the symbol * will denote that you have made changes and need to save the file. Even if you don't save the file and exit Aginity Pro, you'll changes will be there since we maintain state of all tabs regardless. You do need to actually save though if you open the SQL file in another editor.



5.1 Working with File Projects

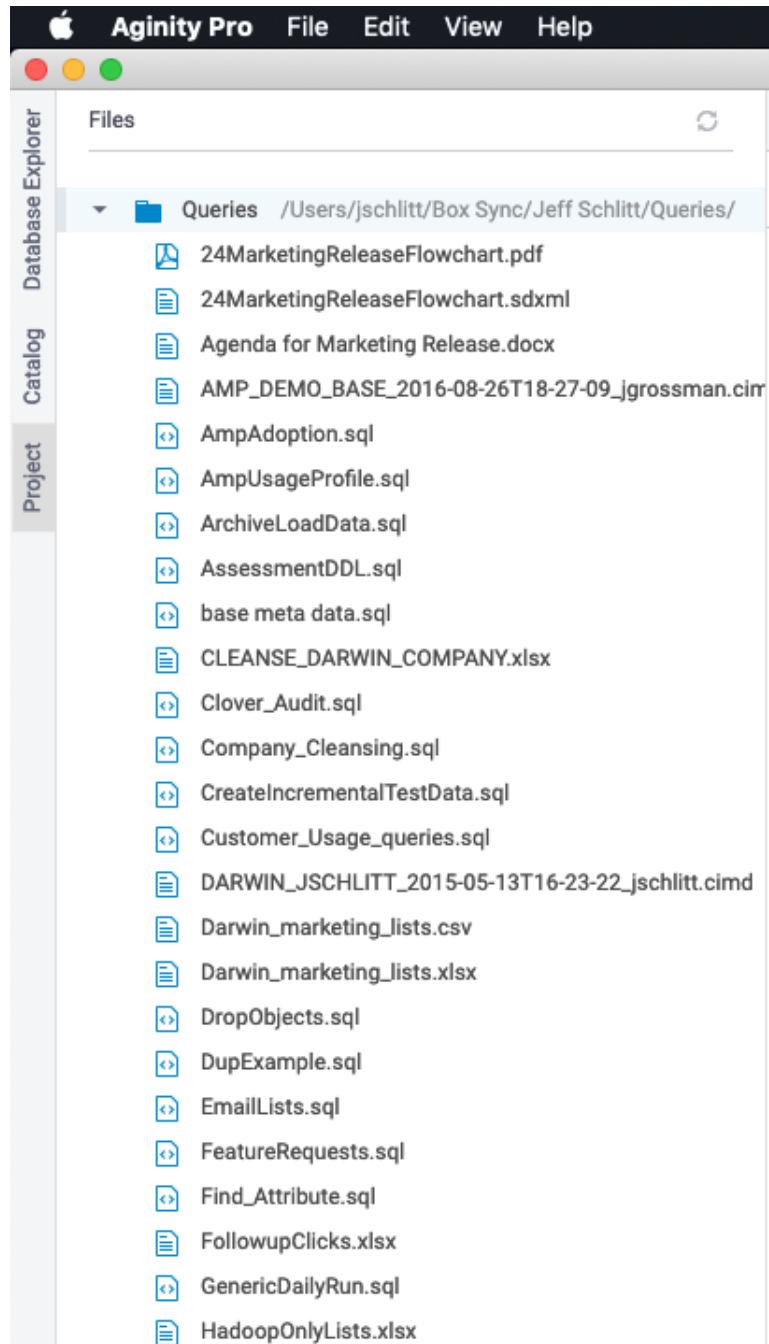
On the left hand side of the object browser you'll see a new tab for **Project**. This allows you to open an entire folder off your hard drive and manage the SQL Files within them. By default you will see a button prompting you to Open Project as shown below.



When you click Open you'll be prompted with your O/S specific file explorer asking you to open a folder to displayed

Note: This is not asking you to Open a SQL File but instead an entire folder. If you want to just open a SQL File use the File - Open menu item.

Once you open the folder you will see the entire folder structure in the object browser as show below.



To open and start editing a file in Aginity Pro just double click.

Right clicking on any file within here presents you with several options listed and described below.

- New file - Will create a new file in the enclosed directory
- New directory - Will create a new folder in the enclosed directory
- Cut file - Will use clipboard to cut the file
- Copy file - Will use clipboard to copy the file
- Rename - Allows you to rename the file
- Delete - Delete the file; note this is an actual physical delete from disk not just within Aginity Pro

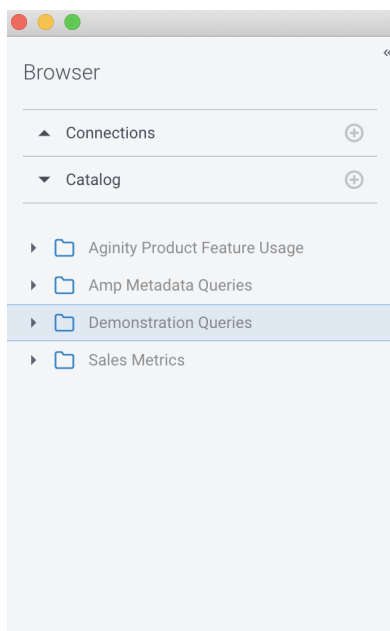
- Open with encoding - If you file uses a specific encoding you can't specify that here
- Copy file path - Copy the file path

Using the Catalog in Aginity Pro

6.1 Catalog Overview

The query catalog is a unique capability of Aginity Pro allowing you to easily save, manage and use items stored in it. Let's start with browsing the catalog.

The catalog is organized by any number of directories and sub-directories of queries as shown below.



See also:

Click [Read this article](#) for some ideas on how to best organize your catalog

Queries are given friendly names and can be used or accessed in several ways. The first way is to drag the catalog query into a **query panel** as shown in the two images below

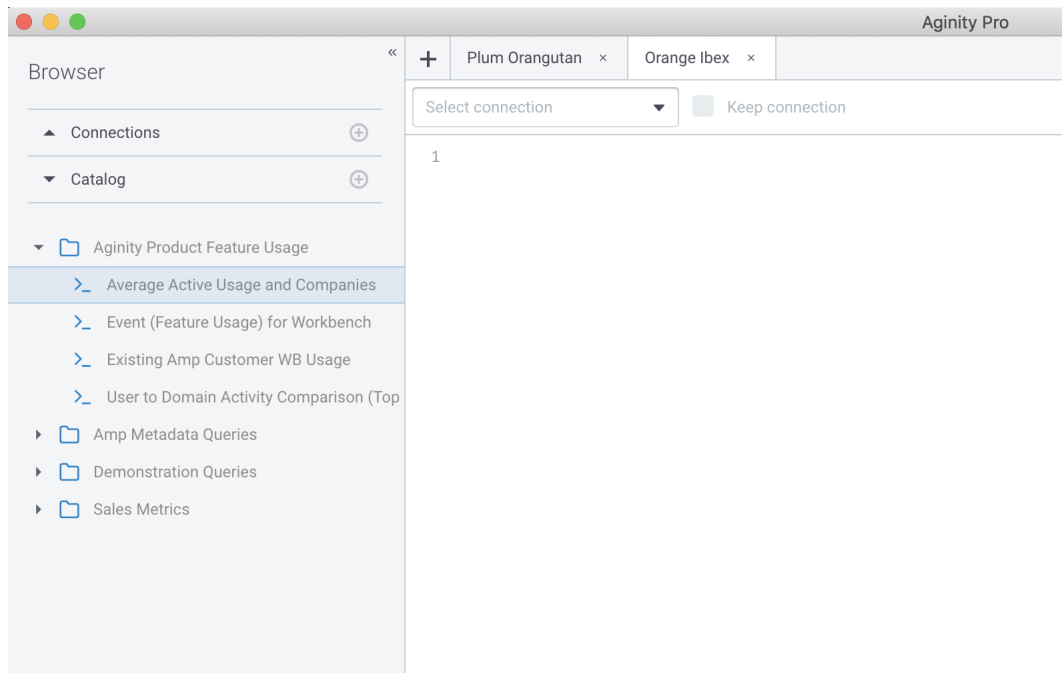


Fig. 1: You can drag the query into the empty query panel

You can also edit queries within a query panel. By right clicking on a catalog item and choosing **edit code in tab** you can change the SQL itself.

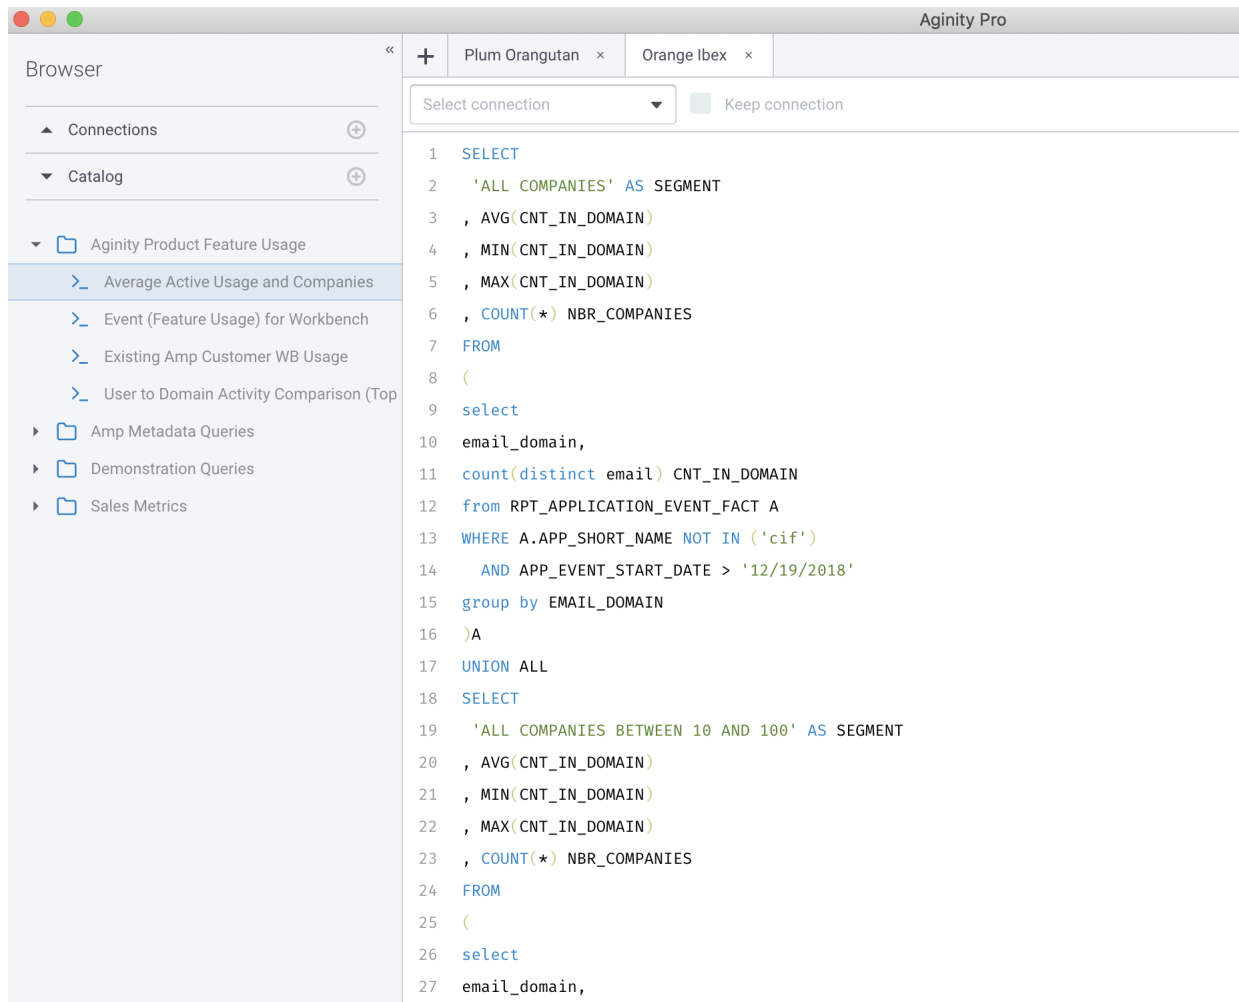
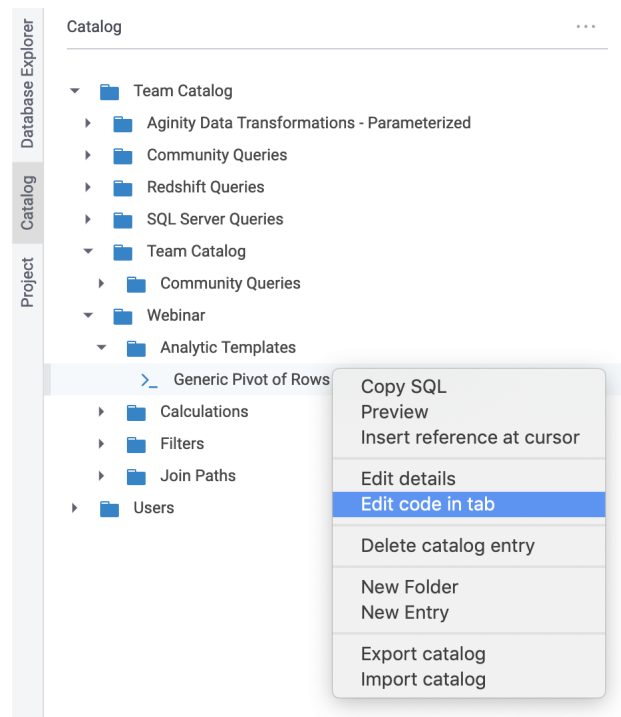
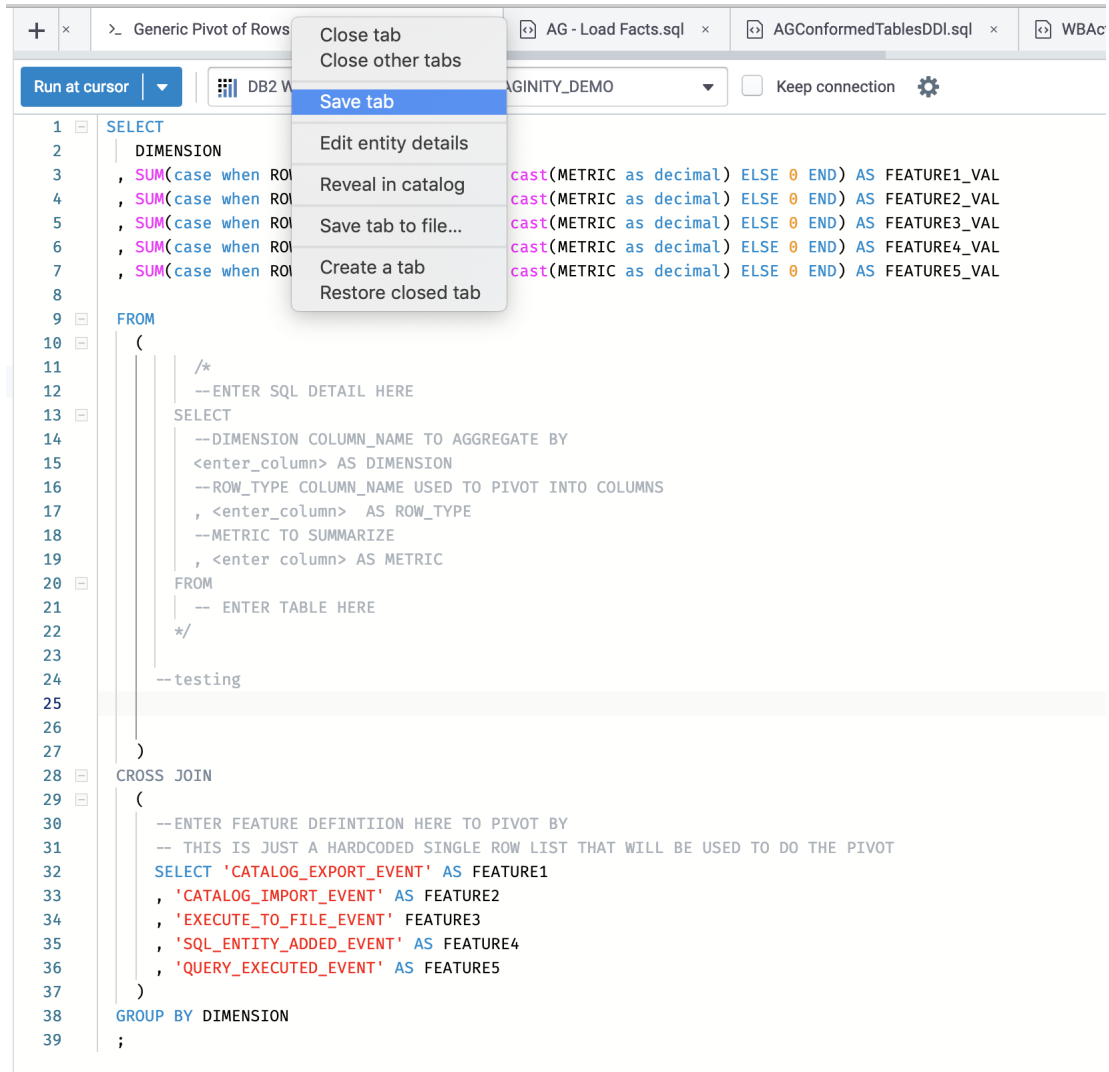


Fig. 2: After you drag the item it will present the full catalog query for edit or execution.

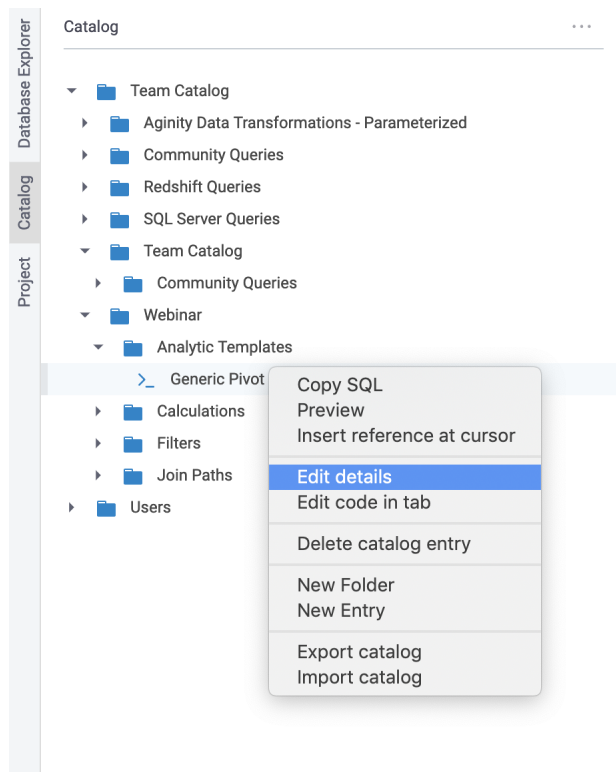


You can edit and save the catalog query in a query panel worksheet as shown below.



Hint: If you see a * next to the tab name in the query panel it means a change has not been saved.

By editing the details of a catalog query you can see key metadata items and statistics about the catalog.



Editing catalog item

Name

Description

Directory

Type

6.2 Calling Catalog Items from SQL

One of the most innovative features of Aginity Pro is the ability to call catalog items from within SQL statements. The video and following examples below will give you basic syntax required to call a catalog item.

6.2.1 Single SQL SELECT Statement

In this example we have a single SQL SELECT statement in the catalog called **OrderAverage**.

This is the code in the Catalog.

```
1 SELECT AVG(LineTotal)
2 FROM sales.SalesOrderDetail
```

To call this from the query panel you would issue the following command. Note the highlighted reference to the object in the catalog.

```
1 SELECT *
2 FROM (@{OrderAverage}) a;
```

This is the code passed to the database.

```
1 SELECT *
2 FROM (
3     SELECT AVG(lineTotal)
4     FROM Sales.SalesOrderDetail
5     ) a;
```

6.2.2 Using Catalog Item to Facilitate Relationships

In this example we will predefine a join in the catalog called **ClientAddress** and then use it in the context of another SQL statement.

This is the code in the Catalog.

```
1 clients c
2 JOIN addresses a
3 ON c.clientsid = a.clientsid
```

To call this from the query panel you would issue the following command. Note the highlighted reference to the object in the catalog.

```
1 SELECT *
2 FROM @{ClientAddress};
```

This is the code passed to the database.

```
1 SELECT *
2 FROM clients c
3 JOIN addresses a
4 ON c.clientsid = a.clientsid;
```

6.2.3 Creating and Reusing Calculations

In this example we will predefine a Calculations in the catalog called **Residual** and then use it in the context of another SQL statement.

This is the code in the Catalog.

```
1 1-total
```

To call this from the query panel you would issue the following command. Note the highlighted reference to the object in the catalog.

```
1 SELECT item,  
2 @{Residual}  
3 FROM table;
```

This is the code passed to the database.

```
1 SELECT item,  
2 1-total  
3 FROM table;
```

6.2.4 Executing a Multi-Step script

In this example we will predefine a series of SQL statements in the catalog called **DailyRun** and then execute it by a single call. Any DML, (SELECT, CREATE, UPDATE, INSERT, DELETE) can be used.

This is the code in the Catalog.

```
1 CREATE TABLE xyz AS...;  
2 CREATE TABLE abc AS...;  
3 SELET * FROM xyz JOIN abc...;
```

To call this from the query panel you would issue the following command. Note the highlighted reference to the object in the catalog.

```
1 @{DailyRun}
```

This is the code passed to the database.

```
1 CREATE TABLE xyz AS...;  
2 CREATE TABLE abc AS...;  
3 SELECT * FROM xyz JOIN abc...;
```

6.2.5 Using Parameters in Catalog

In this example we will use a parameter within a catalog item called **CitySelect**.

This is the code in the Catalog. Note the highlighted line has the parameter in it.

```
1 SELECT *  
2 FROM addresses  
3 WHERE city = $city;
```

To call this from the query panel you would issue the following command. Note the highlighted reference to the object in the catalog.

```
1 @{CitySelect}
```

This is the code passed to the database.

```

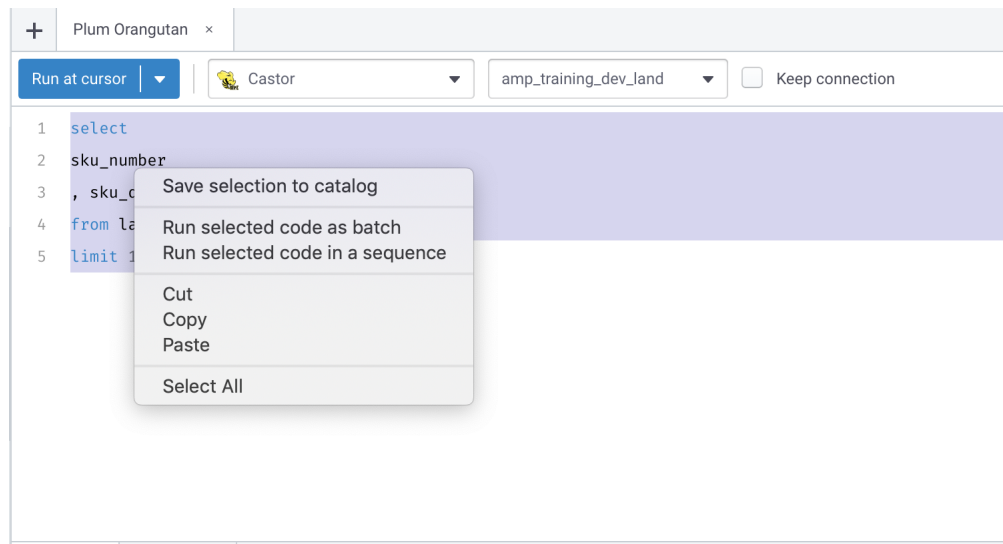
1 SELECT *
2 FROM addresses
3 WHERE city = 'Chicago';

```

Note: Upon execution of the command Aginity Pro will ask you to supply a parameterized value. For more information on parameters *see Using Parameters*

6.3 Adding Items to Catalog

To add a new query to the catalog, highlight any part of a SQL statement(s), right click and select **Save selection to catalog** as shown below.



Once you save you will be asked to provide metadata about the query you are adding.

✕

New catalog item

Name

New snippet

Description

Created from editor

Directory

/

Type

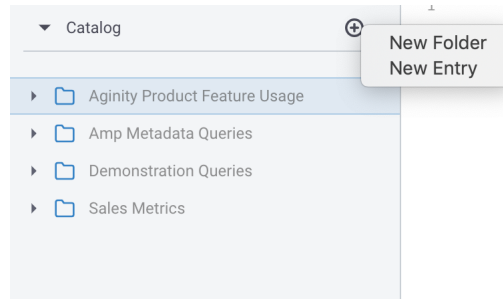
Query

Create


Cancel

Note: It is important to fill out the description metadata as it will be used in searching. The more information you provide about the context of usage, key business terms and intended usage the better search results you will get.

You can also easily create new folders in the catalog or add SQL entries directly from the **Catalog Browser** by clicking the (+) to the right of the Catalog header.



When adding a new folder you will be presented with a basic text box to create your folder as shown below.



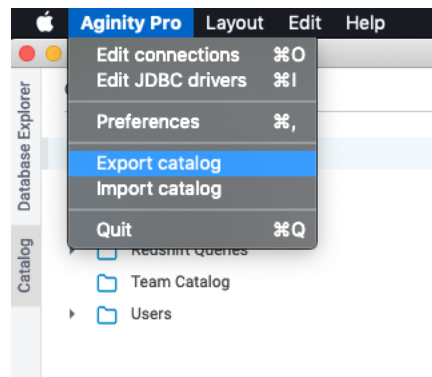
6.4 Import and Export Items from Catalog

You have the ability to Import and Export both your full catalog as well as any folder selected within the catalog.

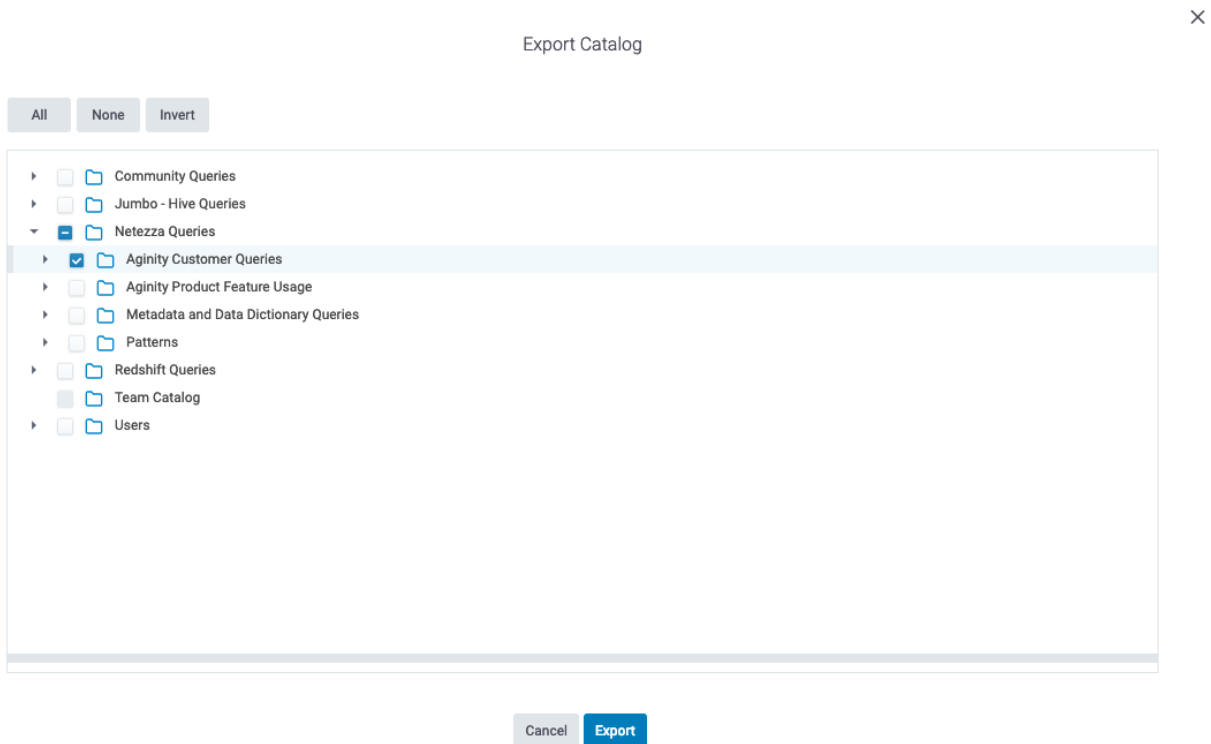
Note: We highly recommend taking regular exports of your catalog to disk to prevent accidental loss of queries.

6.4.1 Export Steps

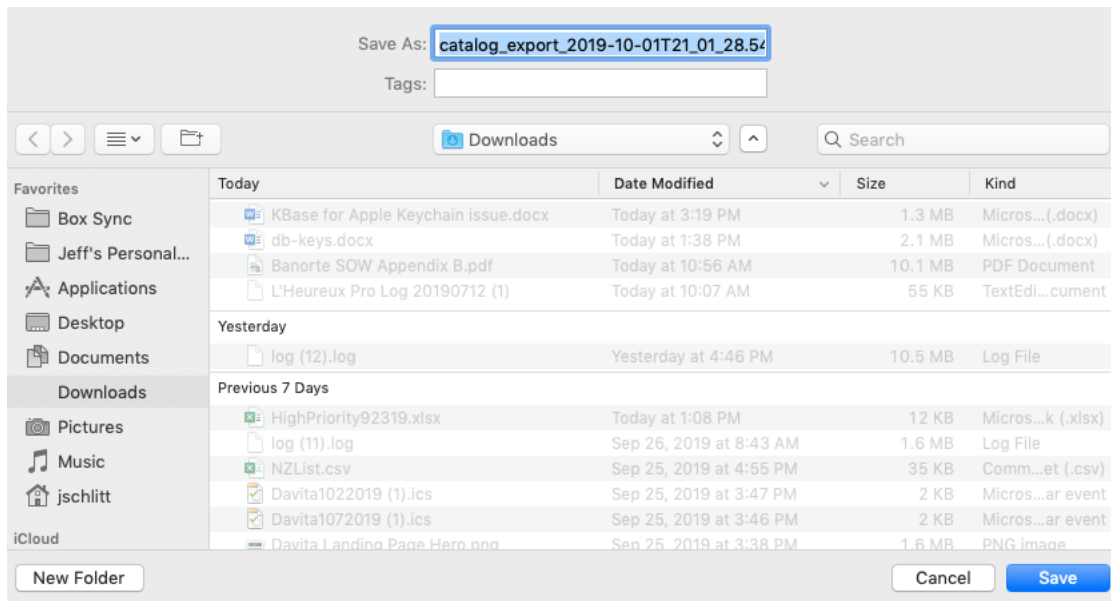
To Export your catalog start by either **Right Clicking** on a Catalog Folder or Using the Aginity Pro -> Export Catalog menu as shown below.



Your first step is to save either an entire folder or a selection of items within a folder.



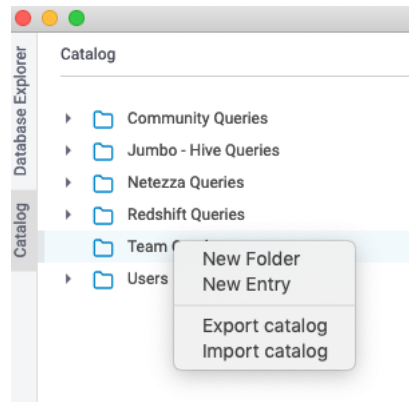
The final step in exporting is to name the file and location you will save to as shown below.



6.4.2 Import Steps

Importing is valuable for several reasons. One is to restore SQL catalog items you may have lost or more importantly the ability to collaborate with other SQL developers and share your SQL making your a more efficient and consistent SQL developer.

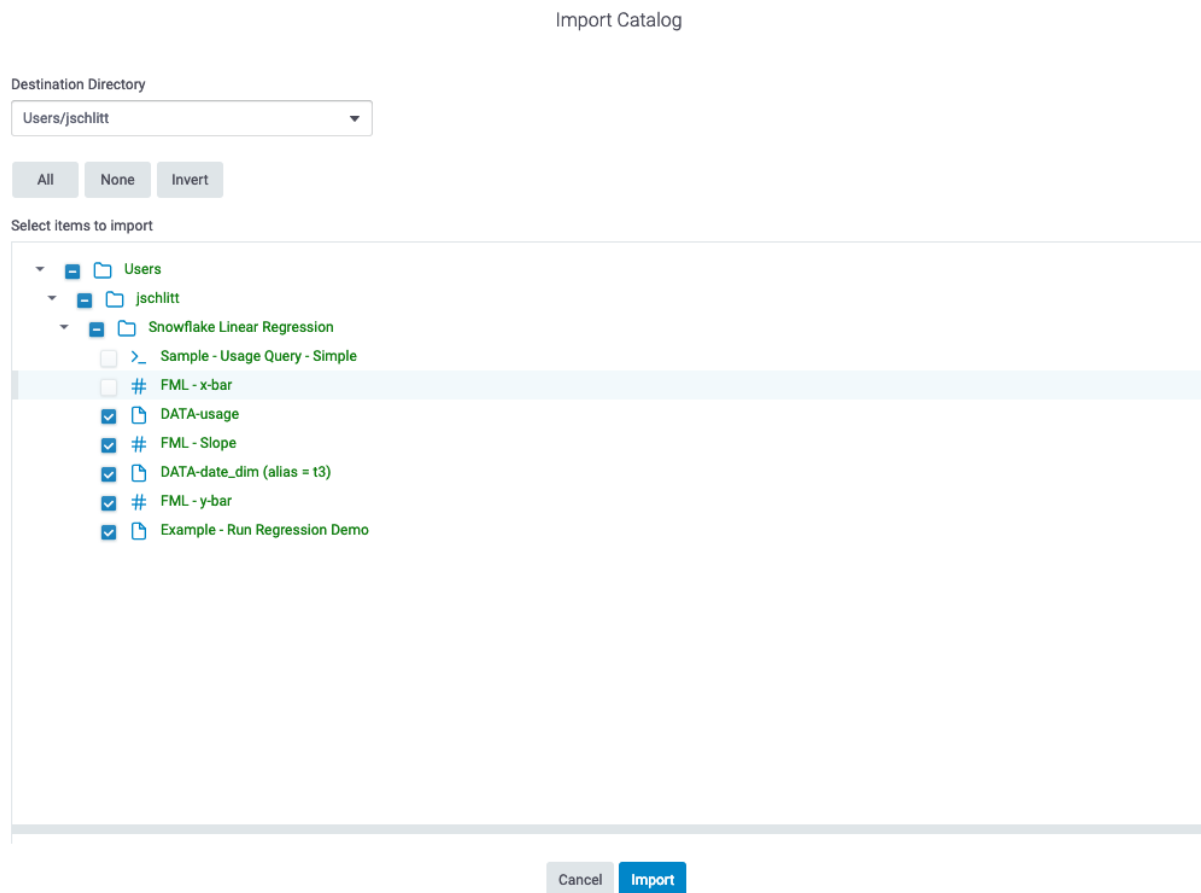
The first step is to **Right Click** on a folder you want to import to and Choose Import Catalog as shown below.



Note: You can also use the Aginity Pro Menu -> Import Catalog. The main difference between using these two capabilities is the context of where you are importing items. If you choose to import from the Aginity Pro menu you must set the context where to import. If you **Right Click** on a folder the context is set to that folder.

After choosing to import you will choose the file to import. By default Aginity catalog files have the extension **.aginitypkg**.

After choosing your file you can specify the folder to import to and what to import as shown below.



Hit the **Import** button and your queries should be imported!

See also:

Aginity also maintains a Github repository of Catalog queries to be used. These queries are either administrative queries for specific platforms, data engineering queries that do powerful transformations against generic data or analytic queries showing how to do cool things in SQL.

Access the repositories [here](#)

See also:

[Using Query Panel](#)

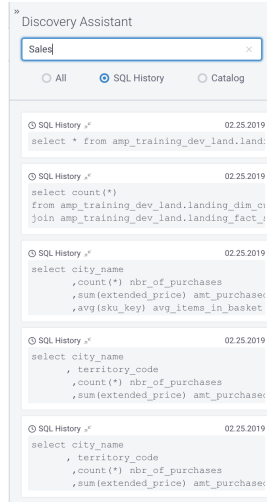
Using Discovery Panel in Aginity Pro

The Discovery Panel is a feature that is unique to Aginity Pro and which allows you to quickly identify and reuse code that you've previously executed or saved. Using the Discovery Panel begins with a search. Watch the video below to get a sense for how this feature works.

By default, the Discovery Panel allows you to search both your query execution history and your Pro Catalog at the same time. However, by selecting one of the radio buttons in the Discovery Panel you can choose to focus on just one or the other if you prefer.

7.1 Searching Query Execution History


Start by typing in any free-text you want such as **Sales** as shown Below



The search engine will do a full text search and display all results.

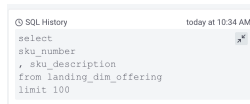
7.2 Searching the Amp Pro Catalog

By clicking on the **Catalog** radio button you will restrict your search as shown below.



7.3 Expanding the results

Once you have performed a search, Amp Pro will dynamically update the results in the Discovery Panel to match the search term you have entered. You'll be able to view a small amount of the query by default, but if more is required you can click on the arrows in the corner of the preview in order to expand the snippet further.



7.4 Doing More with the Discovery Panel

A place holder for advanced filtering capabilities

Using Tools in Aginity Pro


We are adding tools to Aginity Pro to make you more efficient when doing your job. One example of this is the Data Import wizard that is available for Redshift right now.

8.1 Using the Data Import Wizard

The data import wizard is a tool to help you take a local data file, upload it to a temporary location on the database server whether that is in the Cloud or local and then finally load the data to a table.

Note: As of December 2019 we are only supporting uploading to Redshift. Over the next few releases we will be adding support for other database platforms.

To start the data importer go to the **Tools** menu as shown below.



8.1.1 Choosing a File - Step 1

When you start the wizard you will be presented with a six-step process in which to identify a file to upload, describe the behavior you want and then perform the load to a table. In the first screen the following activities are performed.

1. Pick a database connection where you will upload data to.
2. Add a local file using the **Add CSV file** button.
3. Choose a delimiter used to separate fields in the file.

4. If you click the **With header** check box the parser will assume the first row in the file is a header if not it will treat the first row as a data row.
5. You can click Update preview to show the first few rows in the file

✕

Import wizard (step 1 of 6)

Choose connection

Redshift Snoopy

Add CSV file

/Users/jschlitt/Downloads/dim_channel_db2.csv

Delimiter

Pipe (|)

Encoding

Unicode (UTF-8)(utf-8, code page 65001)

With header
☒

Update preview

1	*In-Person*	*In-Person (as in a store)*
2	*Mobile*	*Mobile (as in a mobile app)
3	*Web*	*Web (as in a web page)*
4	*Call*	*Call (as in a phone call)*
5	*Chat*	*Chat (as in an instance me

Cancel

Next

When complete choose Next.

In the next screen you will provide your security credentials specific to the database you are uploading to. Reference the sections below for specific information on each platform supported.

8.1.2 Redshift Data Upload - Step 2

When working with Amazon Redshift the file you upload will first be moved to a S3 bucket of your choosing prior to being loaded into a Redshift table. You must provide the following items.

1. Your Access Key
2. Your Secret Key
3. The Region Endpoint
4. Once you enter the first three you should be able to choose the S3 bucket to use to store data.

The second section within Redshift lets you control the behavior of the file you upload to S3.

1. If you select the check box labeled **Leave the intermediate CSV file on the cloud after importing data** you will leave a version of the file on S3.
2. You can enter a filename of choice for the file.
3. If you want the file compressed in S3 you can choose from NONE, ZIP or GZIP options.
4. You can control file behavior if the file exists. You can either overwrite or abort the process.

See also:

Please refer to the [AWS Bucket Naming conventions article](#) for proper naming standards.

Import wizard (step 2 of 6)

Select Amazon S3 bucket for uploading the data to ⓘ

Access Key AKIARCE3GH0V2PGUUXUT

Private Key

Region Endpoint US East (N. Virginia)

S3 Bucket snoopy-ent-warehouse

☒ Leave the intermediate CSV file on the cloud after importing data

File name _forjeff.txt

Compression NONE

Existing file action Overwrite

1	*In-Person*	*In-Person (as in a store)*
2	*Mobile*	*Mobile (as in a mobile app
3	*Web*	*Web (as in a web page)*
4	*Call*	*Call (as in a phone call)*
5	*Chat*	*Chat (as in an instance me

Cancel Back Next

8.1.3 Choosing Columns and Rows - Step 3

In step 3, you can choose a numeric limiter for both the number of columns you want to parse and then the number of rows to load. This step is valuable if you have very wide and/or very deep files and you want to work with less columns or less rows.

Hint: It is helpful with very large files to load the first 100 rows to ensure everything works well versus letting it run for an hour only to find out it didn't work.

Import wizard (step 3 of 6)

Specify data range

Start column

1

End column

3

Import all rows

☒

Start row

1

End row

0

1	*In-Person*	*In-Person (as in a store)*
2	*Mobile*	*Mobile (as in a mobile app
3	*Web*	*Web (as in a web page)*
4	*Call*	*Call (as in a phone call)*
5	*Chat*	*Chat (as in an instance me

Cancel

Back

Next

8.1.4 Choosing Columns to Include - Step 4

This step may seem redundant to the prior step for columns but it allows you to better see more about the columns. Let's say you have 200 columns and you know you only want the first 10 in Step 3 you can choose columns 1 to 10 and in Step 4 you can better preview and determine the order of the columns you want to load.

Import wizard (step 4 of 6)

Specify columns to include

Available columns

Include columns

1
In-Person
In-Person (as in a store)

1	*In-Person*	*In-Person (as in a store)*
2	*Mobile*	*Mobile (as in a mobile app)
3	*Web*	*Web (as in a web page)*
4	*Call*	*Call (as in a phone call)*
5	*Chat*	*Chat (as in an instance me

Cancel

Back

Next

8.1.5 Configuring Output Columns - Step 5

The next step lets you control the column name, data type and nullability.

×

Import wizard (step 5 of 6)

Specify physical data types

#	Name	Data Type	Precision	Scale	Null
1	1	BIGINT			<input type="checkbox"/>
2	*In-Person*	VARCHAR	8		<input type="checkbox"/>
3	*In-Person (as in a store)*	VARCHAR	42		<input type="checkbox"/>

1	*In-Person*	*In-Person (as in a store)*	
2	*Mobile*	*Mobile (as in a mobile app	
3	*Web*	*Web (as in a web page)*	
4	*Call*	*Call (as in a phone call)*	
5	*Chat*	*Chat (as in an instance me	

Cancel

Back

Next

8.1.6 Setting Output Details and Loading - Step 6

In step 6 you will supply the database you'll load to, schema, table name and other behaviors that will be followed during a load execution.

×

Import wizard (step 6 of 6)

Specify the target database and table name

Database

snoopy_master

Schema

public

Table

SAMPLE_TABLE

Distribution

< none, use random distribution >

If Table Exist

ABORT

1	*In-Person*	*In-Person (as in a store)*	
2	*Mobile*	*Mobile (as in a mobile app	
3	*Web*	*Web (as in a web page)*	
4	*Call*	*Call (as in a phone call)*	
5	*Chat*	*Chat (as in an instance me	

Cancel

Back

Submit

9.1 Video Reference

The following is a reference list of videos about Aginity Pro and how to use its features.

-
-
-
-
-
-

9.2 Shortcut Key Reference

Command	Action
Command K + C (Mac) or Ctrl K + C (Windows)	Comment lines of code
Command K + U (Mac) or Ctrl K + U (Windows)	Uncomment lines of code
Command F (Mac) or Ctrl F (Windows)	Find and Replace Menu

9.3 Frequently Asked Questions

The section below is dedicated to commonly asked questions that we compile and provide answers to. Over time as things change we will keep this table up to date.

9.3.1 What Control do I have with Row Limit Preference?

By default we cap the row limit preference to 100K rows to discourage users from spooling too much data across your network. Trust us your network infrastructure folks do appreciate it. Fear not though, you have options. First, if you set the preference to 0, it will spool as much as you want until either you run out of memory or your query completes. Alternatively, we recommend spooling the output to a file if you are doing returns greater than 100K rows. Refer to the documentation on *execute query to file*.

9.4 SQL Reference Guides for Supported connections

As SQL developers ourselves we thought putting links to the databases we support in our documentation was a good idea so here they are!

-
-
-
-

9.5 Driver Support Reference

The table below is a cross reference of driver versions by platform

Platform	Driver Version	Class Name
Netezza	Not distributed with software (nzjdbc.jar)	org.netezza.Driver
SQL Server	7.2.0 jre8	com.microsoft.sqlserver.jdbc.SQLServerDriver
Redshift	1.2.16.1027	com.amazon.redshift.jdbc.Driver
Greenplum and Postgres	42.2.5	org.postgresql.Driver
Hive	1.2.1	org.apache.hive.jdbc.HiveDriver
Snowflake	3.6.23, 3.6.28, 3.5.5	net.snowflake.client.jdbc.SnowflakeDriver
IBM DB2 Warehouse	db2jcc4	com.ibm.db2.jcc.DB2Driver

CHAPTER 10

Troubleshooting

This page is dedicated to common issues or errors you may encounter when using Aginity Pro. We recommend you also post any new issues or questions to our Aginity Pro Forum.

There is a great resource for posting questions, searching for common solutions to features of Pro or just learning how to be efficient as a SQL analytic developer through samples, best practices and collaboration with other Pro users!

Issue Number	Is this a Defect	Source of Error	Description of Error	Resolution
1	No	Cannot Download Drivers via Connection Screen	In most cases your organization is blocking Aginity's ability to download drivers from a Maven repository	Follow directions in the Configure Connections documentation to Manually Add Drivers.

CHAPTER 11

End User License for Aginity Pro

[Click here for end users license](#)