



# Sisense V7.1

---

Technical Documentation

# Table of Contents

<b>New to Sisense .....</b>	<b>13</b>
<b>Sisense Architecture.....</b>	<b>15</b>
<b>Backing Up the Sisense Application Database.....</b>	<b>21</b>
<b>Communication Ports .....</b>	<b>23</b>
<b>Sisense Basic Concepts and Terminology .....</b>	<b>28</b>
<b>Configuring Base URLs.....</b>	<b>36</b>
<b>Customizing the Installation .....</b>	<b>39</b>
<b>Downloading and Installing Sisense .....</b>	<b>42</b>
<b>Importing and Exporting Elasticube Data .....</b>	<b>46</b>
<b>Launching Sisense .....</b>	<b>48</b>
<b>Backing Up and Migrating a Sisense Installation .....</b>	<b>49</b>
<b>Minimum Requirements and Supported Platforms .....</b>	<b>54</b>
Supported Web Browsers .....	54
Supported Operating Systems.....	54
Capacity and Hardware Requirements .....	55
Supported Locales .....	57
<b>Installing Sisense Offline .....</b>	<b>58</b>
<b>Introduction .....</b>	<b>63</b>
Audience and Scope .....	63
Overview of the Sisense Web Application .....	63
Requirements .....	63
A Tour of the Sisense Web Application Interface .....	64
<b>Setting Up Your Production Environment.....</b>	<b>65</b>
<b>Performing a Silent Installation.....</b>	<b>67</b>
<b>Tutorial: Getting Started with Sisense .....</b>	<b>69</b>
<b>Upgrading Sisense.....</b>	<b>88</b>

<b>Sisense User Roles.....</b>	<b>94</b>
<b>Connecting to Access .....</b>	<b>98</b>
<b>Connecting to Bing .....</b>	<b>100</b>
<b>Connecting to Box.....</b>	<b>106</b>
<b>Changing Connectivity Settings for Data Sources.....</b>	<b>111</b>
<b>Connection String Parameters.....</b>	<b>114</b>
<b>Creating an App in Facebook.....</b>	<b>122</b>
<b>Connecting to CSV.....</b>	<b>128</b>
<b>Connecting to CSV.....</b>	<b>133</b>
<b>Customizing Google Analytics and AdWords Tables.....</b>	<b>136</b>
<b>Connecting to DynamicsNav .....</b>	<b>143</b>
<b>Connecting to DynamoDB.....</b>	<b>147</b>
<b>Connecting to eBay .....</b>	<b>152</b>
<b>Connecting to Exact Online .....</b>	<b>160</b>
<b>Connecting to Excel.....</b>	<b>168</b>
<b>Connecting to Microsoft Excel.....</b>	<b>172</b>
<b>Connecting to Facebook Ads.....</b>	<b>176</b>
<b>Facebook Insights.....</b>	<b>189</b>
<b>Connecting to Facebook.....</b>	<b>205</b>
<b>Connecting to Gmail.....</b>	<b>225</b>
<b>Connecting to Google Adwords.....</b>	<b>230</b>
<b>Connecting to Google Analytics .....</b>	<b>234</b>
<b>Connecting to Google BigQuery .....</b>	<b>238</b>
<b>Connecting to Google Spreadsheets.....</b>	<b>251</b>
<b>Connecting to Hadoop via Hive.....</b>	<b>254</b>
<b>Connecting to Heroku Postgres.....</b>	<b>270</b>
<b>Connecting to HubSpot.....</b>	<b>273</b>
<b>Importing Data with Custom Queries.....</b>	<b>283</b>

<b>Connecting to Intacct.....</b>	<b>286</b>
<b>Connecting to MailChimp.....</b>	<b>294</b>
<b>Connecting to MongoDB.....</b>	<b>303</b>
<b>Connecting to Microsoft Active Directory.....</b>	<b>321</b>
<b>Connecting to Microsoft Dynamics CRM.....</b>	<b>329</b>
<b>MySQL Live.....</b>	<b>358</b>
Adding Live Connections.....	358
<b>Connecting to MySQL.....</b>	<b>361</b>
<b>Connecting to MySQL.....</b>	<b>364</b>
<b>Connecting via ODBC Drivers.....</b>	<b>368</b>
<b>Connecting via OLE DB Drivers.....</b>	<b>374</b>
<b>Oracle Live.....</b>	<b>378</b>
Adding Oracle Live Connections.....	378
<b>Connecting to Oracle.....</b>	<b>381</b>
<b>Connecting to Oracle.....</b>	<b>384</b>
<b>Connecting to PayPal.....</b>	<b>387</b>
<b>Connecting to PostgreSQL.....</b>	<b>392</b>
<b>Connecting to PostgreSQL.....</b>	<b>396</b>
<b>PostgreSQL Live.....</b>	<b>400</b>
Adding Live Connections.....	400
<b>Preparing Your Excel Files.....</b>	<b>403</b>
<b>Connecting to Quickbooks Desktop.....</b>	<b>409</b>
<b>Connecting to QuickBooks Online.....</b>	<b>426</b>
<b>Amazon Redshift Live.....</b>	<b>433</b>
Adding Amazon Redshift Live Connections.....	433
<b>Connecting to Redshift.....</b>	<b>436</b>
<b>Connecting to Redshift.....</b>	<b>439</b>
<b>Connecting to Salesforce.....</b>	<b>442</b>

<b>Segmenting Google Analytics Data .....</b>	<b>445</b>
<b>Connecting to ServiceNow.....</b>	<b>450</b>
<b>Connecting to SharePoint .....</b>	<b>456</b>
<b>Connecting to Sisense ElastiCubes .....</b>	<b>461</b>
<b>Connecting to Snowflake.....</b>	<b>472</b>
<b>Connecting to Splunk .....</b>	<b>476</b>
<b>Connecting to SQL.....</b>	<b>479</b>
<b>Connecting to SQL Server.....</b>	<b>482</b>
<b>Adding SQL Server Live Connections.....</b>	<b>486</b>
<b>Connecting to Stripe .....</b>	<b>488</b>
<b>Connecting to Teradata .....</b>	<b>493</b>
<b>Connecting to Twitter .....</b>	<b>496</b>
<b>Using Custom SQL to Import Data.....</b>	<b>502</b>
<b>Using Custom SQL to Import Data.....</b>	<b>504</b>
<b>Connecting to Xero Accounting .....</b>	<b>506</b>
<b>Connecting to YouTube Analytics .....</b>	<b>515</b>
<b>Connecting to YouTube .....</b>	<b>524</b>
<b>Connecting to Zendesk .....</b>	<b>535</b>
<b>Deleting ElastiCubes.....</b>	<b>538</b>
<b>Importing and Exporting ElastiCube Data .....</b>	<b>540</b>
<b>Managing ElastiCubes .....</b>	<b>542</b>
<b>Scheduling Data Loads.....</b>	<b>544</b>
<b>Sisense Server Console .....</b>	<b>546</b>
<b>ElastiCube for Advanced Business Scenarios .....</b>	<b>549</b>
<b>Enhancing Data .....</b>	<b>551</b>
<b>Integrating Data.....</b>	<b>567</b>
<b>Reformatting Data.....</b>	<b>573</b>
<b>Date and Time Functions.....</b>	<b>574</b>

<b>Functions Overview</b> .....	<b>576</b>
<b>Logical Functions</b> .....	<b>577</b>
<b>Mathematical Functions</b> .....	<b>578</b>
<b>Miscellaneous Functions</b> .....	<b>581</b>
<b>String and Text Functions</b> .....	<b>582</b>
<b>Web Functions</b> .....	<b>585</b>
<b>Adding a Custom Column</b> .....	<b>589</b>
<b>Adding a Custom Table</b> .....	<b>594</b>
<b>Editing Field Formulas</b> .....	<b>595</b>
<b>Editing SQL Statements</b> .....	<b>597</b>
<b>Previewing Results</b> .....	<b>599</b>
<b>SQL Reference</b> .....	<b>601</b>
<b>Transforming and Enriching Data</b> .....	<b>605</b>
<b>Chasm and Fan Traps</b> .....	<b>606</b>
<b>Creating and Removing a Relationship between Tables</b> .....	<b>610</b>
<b>Handling Relationship Cycles</b> .....	<b>612</b>
<b>Managing Tables and Columns</b> .....	<b>618</b>
<b>Many-to-Many Relationships</b> .....	<b>622</b>
<b>Navigating the ElastiCube Manager</b> .....	<b>636</b>
<b>Previewing Data from a Table</b> .....	<b>638</b>
<b>Finding Tables and Columns</b> .....	<b>645</b>
<b>Working with Data</b> .....	<b>647</b>
<b>Building ElastiCubes</b> .....	<b>648</b>
<b>Build Settings and Data Accumulation Behavior</b> .....	<b>650</b>
<b>Introduction to Data Sources</b> .....	<b>656</b>
<b>ElastiCube Manager</b> .....	<b>661</b>
Accessing the ElastiCube Manager .....	661
Navigating the ElastiCube Manager .....	662

<b>ElastiCubes .....</b>	<b>665</b>
ElastiCubes – Technical Overview.....	666
<b>Manage Data.....</b>	<b>671</b>
Data Page .....	671
<b>Live Connect.....</b>	<b>674</b>
Adding Live Connections.....	675
Working with Data Sets.....	675
<b>Making Changes to an Existing ElastiCube.....</b>	<b>681</b>
<b>Adding Text Widgets .....</b>	<b>682</b>
<b>Adding Widgets to a Dashboard.....</b>	<b>684</b>
<b>Area Chart.....</b>	<b>691</b>
<b>Bar Chart .....</b>	<b>696</b>
<b>Box &amp; Whisker Plot.....</b>	<b>698</b>
<b>Calendar Heatmap .....</b>	<b>701</b>
<b>Column Chart .....</b>	<b>703</b>
<b>Indicator .....</b>	<b>707</b>
<b>Line Chart.....</b>	<b>710</b>
<b>Managing Widgets.....</b>	<b>715</b>
<b>Pie Chart .....</b>	<b>717</b>
<b>Pivot.....</b>	<b>720</b>
<b>Polar Chart.....</b>	<b>725</b>
<b>Scatter Chart.....</b>	<b>727</b>
<b>Scatter Map.....</b>	<b>730</b>
<b>Sunburst Widget.....</b>	<b>735</b>
<b>Table .....</b>	<b>737</b>
<b>Treemap.....</b>	<b>739</b>
<b>Widget Designer .....</b>	<b>742</b>
<b>Changing the Dashboard’s Color Palette .....</b>	<b>746</b>

<b>Changing a Widget’s Visualization .....</b>	<b>749</b>
<b>Combining Two Types of Visualizations in a Widget .....</b>	<b>750</b>
<b>Creating a Continuous Chart with Missing Date Values .....</b>	<b>752</b>
<b>Adding Drill Hierarchies to Widgets.....</b>	<b>755</b>
<b>Formatting Numbers in Widgets.....</b>	<b>758</b>
<b>Renaming Fields in Widgets .....</b>	<b>761</b>
<b>Selecting Colors in Widgets.....</b>	<b>762</b>
<b>Configuring How Filters Affect the Dashboard and Widgets...766</b>	
<b>Creating Dashboard Filters.....</b>	<b>771</b>
<b>Creating Widget Filters .....</b>	<b>785</b>
<b>Creating Dependent Filters.....</b>	<b>787</b>
<b>Editing and Deleting a Dashboard Filter .....</b>	<b>790</b>
<b>Interacting with Filters as a Designer .....</b>	<b>792</b>
<b>Interacting with Filters as a Viewer.....</b>	<b>793</b>
<b>Limiting Filters to Specific Values with Background Filters ...796</b>	
<b>Locking Filters.....</b>	<b>798</b>
<b>Restoring Filters .....</b>	<b>799</b>
<b>Saving Your Default Filters View.....</b>	<b>800</b>
<b>Switching Filters On and Off.....</b>	<b>802</b>
<b>Introduction to Filters .....</b>	<b>803</b>
<b>Building Formulas with Functions .....</b>	<b>804</b>
<b>Conditional Statements.....</b>	<b>809</b>
<b>Creating and Editing a Formula.....</b>	<b>811</b>
<b>Creating Formulas Based on Criteria and Conditions (Filters)815</b>	
<b>Function Reference.....</b>	<b>820</b>
<b>Introduction to Formulas .....</b>	<b>841</b>
<b>Quick Functions .....</b>	<b>843</b>
<b>Using R in Formulas.....</b>	<b>846</b>

Reusing Formulas.....	856
Formula Editor .....	858
Fiscal Years.....	861
Date and Time Fields.....	865
Time in Sisense.....	867
Changing a Dashboard’s Data Source .....	868
Creating Dashboards.....	869
Customizing the Dashboard Layout .....	873
Dashboards.....	879
Reporting in Sisense.....	882
Managing Dashboards .....	883
Sharing Dashboards .....	886
Changing Sisense’s Language .....	891
Creating PDF Reports.....	894
Customizing PDF Reports .....	896
Deleting Dashboards.....	899
Downloading a Dashboard as an Image.....	900
Downloading a Widget as an Image .....	902
Drilling Down in a Widget .....	903
Email Troubleshooting.....	905
Exporting and Sharing Dashboards .....	914
Exporting Pivot Tables to Excel.....	915
Exporting Table Widgets to PDF .....	918
Exporting Widgets to CSV/Excel.....	920
Exporting Pivot Tables to PDF .....	921
Exporting Widgets.....	924
Interacting with Filters as a ViewerFiltering Dashboards.....	925
Making Selections in a Widget .....	928

<b>Exporting Dashboards .....</b>	<b>929</b>
<b>Working with Dashboards .....</b>	<b>930</b>
<b>Sending Email Dashboard Reports.....</b>	<b>931</b>
<b>Viewing Dashboards.....</b>	<b>934</b>
<b>Advanced Settings .....</b>	<b>935</b>
<b>Creating Build Alerts .....</b>	<b>936</b>
<b>Configuring Additional Options.....</b>	<b>938</b>
<b>Creating Data Alerts .....</b>	<b>940</b>
<b>Customizing Notification Messages.....</b>	<b>943</b>
<b>Managing Tiles in Sisense Pulse .....</b>	<b>946</b>
<b>Notifying Additional Users .....</b>	<b>951</b>
<b>IntroductionSisense Pulse .....</b>	<b>953</b>
<b>Setting Notification Channels.....</b>	<b>955</b>
<b>Sisense Alerts.....</b>	<b>964</b>
<b>Sisense Pulse.....</b>	<b>967</b>
<b>Accessing the Sisense Application Database.....</b>	<b>969</b>
<b>Enabling SSL for the Application Database .....</b>	<b>973</b>
<b>Hosting the Application Database .....</b>	<b>976</b>
<b>Creating a KeyFile .....</b>	<b>977</b>
<b>Limiting Access to the Application Database.....</b>	<b>978</b>
<b>Sisense Application Database .....</b>	<b>979</b>
<b>Managing Your Users Dashboards.....</b>	<b>980</b>
<b>Assigning Security Rights .....</b>	<b>982</b>
<b>Data Security.....</b>	<b>983</b>
<b>Security Bulletin .....</b>	<b>988</b>
<b>Optional Security Hardening for Sisense Web Pages.....</b>	<b>989</b>
<b>Sisense Security Overview .....</b>	<b>993</b>
<b>Security at Sisense.....</b>	<b>998</b>

<b>Setting Up SSL .....</b>	<b>999</b>
<b>Data Management .....</b>	<b>1002</b>
<b>Managing Drill Hierarchies.....</b>	<b>1003</b>
<b>Adding and Removing ElastiCube Servers.....</b>	<b>1006</b>
<b>Sharing ElastiCubes .....</b>	<b>1008</b>
<b>Embedding Dashboards and Widgets .....</b>	<b>1010</b>
<b>Displaying Custom URLs .....</b>	<b>1015</b>
<b>Rebranding Sisense Automated Emails.....</b>	<b>1016</b>
<b>White Labeling Sisense.....</b>	<b>1027</b>
<b>Embedded Analytics.....</b>	<b>1037</b>
<b>Querying ElastiCubes during Builds .....</b>	<b>1040</b>
<b>High Availability in Sisense.....</b>	<b>1060</b>
<b>Replicating the Application Database .....</b>	<b>1082</b>
<b>Routing Queries in ElastiCube Sets .....</b>	<b>1094</b>
<b>Changing the Sisense Web Application’s Language.....</b>	<b>1096</b>
<b>Internationalization and Localization .....</b>	<b>1101</b>
<b>Translating Sisense Metadata .....</b>	<b>1102</b>
<b>Server Settings .....</b>	<b>1109</b>
<b>Setting Up a Custom Email Server.....</b>	<b>1111</b>
<b>System Configuration.....</b>	<b>1115</b>
<b>System Settings .....</b>	<b>1117</b>
<b>Introduction to SSO.....</b>	<b>1118</b>
<b>SSO via JWT.....</b>	<b>1119</b>
<b>SSO via SAML 2.0 .....</b>	<b>1127</b>
<b>Customizing User Roles.....</b>	<b>1134</b>
<b>Integrating Active Directory.....</b>	<b>1142</b>
<b>License Utilization.....</b>	<b>1147</b>
<b>Managing Dashboard Users .....</b>	<b>1149</b>

<b>Managing User Groups.....</b>	<b>1152</b>
<b>Managing User Sessions .....</b>	<b>1156</b>
<b>Managing Sisense Users .....</b>	<b>1158</b>
<b>Administration.....</b>	<b>1159</b>
<b>ElastiCube Manager Management .....</b>	<b>1161</b>
<b>Managing Plug-ins.....</b>	<b>1162</b>
<b>Migrating Sisense across Environments .....</b>	<b>1164</b>
<b>Viewing Dashboards on Mobile Devices.....</b>	<b>1171</b>
<b>Connecting to Sisense via ODBC .....</b>	<b>1191</b>

# New to Sisense

This section describes how Sisense works, how to set Sisense up, and provides a tutorial for building your first dashboard from scratch.

## Learn how Sisense Works

Sisense's end-to-end, single stack solution simplifies complex data in 3 short steps:

1. **Model:** Blend massive amounts of data from a variety of sources into the Sisense ElastiCube.
2. **Analyze:** Once your data is inside the ElastiCube, interact with it in a visual way to get deeper understanding, or zoom out to understand the holistic data landscape.
3. **Impact:** Create dashboards that make it easy to get insights and take action.

The topics below describe in more detail how Sisense works.

- [Sisense Basic Concepts and Terminology](#)
- [Sisense Architecture Technical Overview](#)

## Set Up Sisense

Before you start building ElastiCubes and designing dashboards, you need to set Sisense up, the topics below explain how:

- [Downloading and Installing Sisense](#)
- [Setting Up Your Production Environment](#)
- [Customizing the Installation](#)
- [Performing a Silent Installation](#)
- [Migrating Sisense](#)
- [Upgrading Sisense](#)
- [Launching Sisense](#)

# Give Sisense a Try

Now that you've set Sisense up, it's time to play. See the tutorial below to learn how to import data and build your first dashboard.

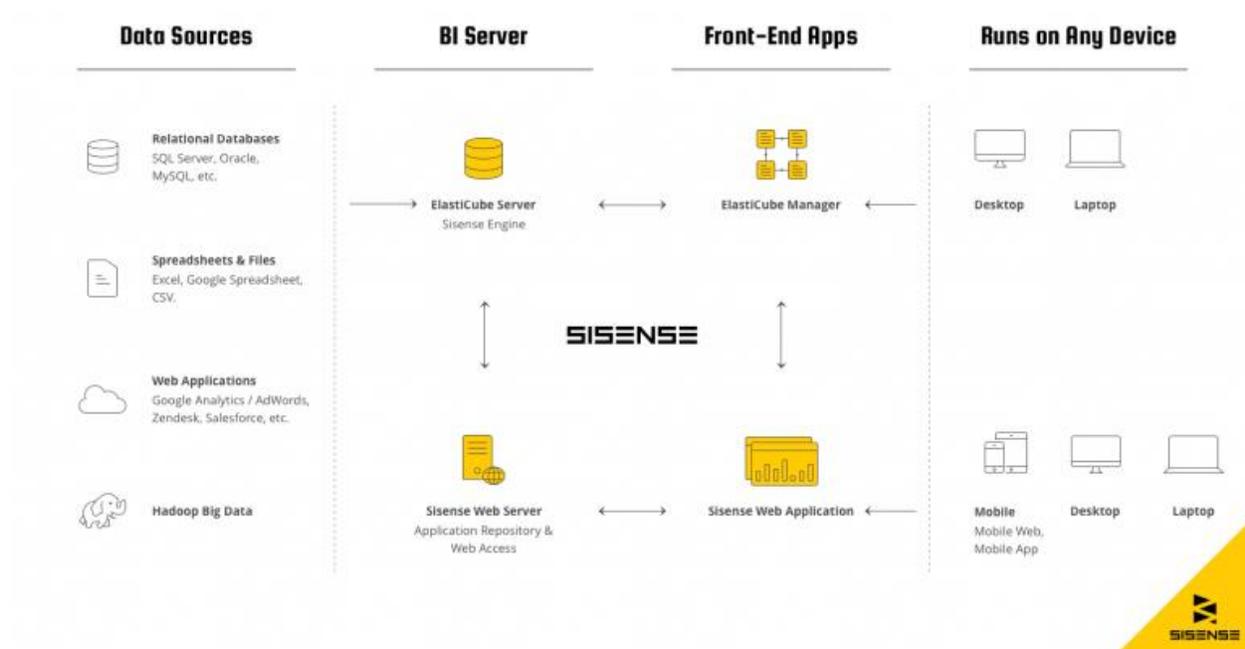
- [Tutorial: Getting Started](#)

See [Minimum Requirements and Supported Platforms](#) to see what you need to support Sisense.

# Sisense Architecture

This topic provides an advanced description of Sisense's architecture for IT managers and Administrators responsible for supporting Sisense in their organization.

For a high-level overview of the Sisense architecture, click [Sisense Architecture](#).

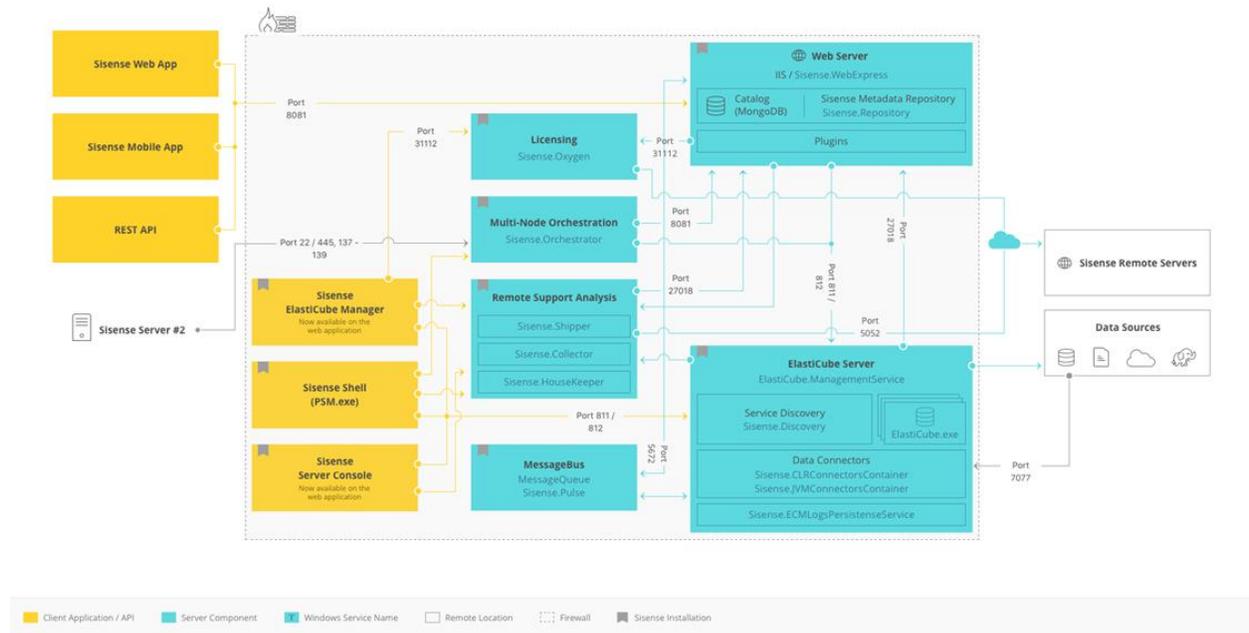


When working with Sisense, there are two possible types of deployments. The first type is a **Single deployment** in which Sisense is installed on a single server. Single deployments can be leveraged to support single integrations or as proof of concept integrations for larger deployments.

The second type is a **Distributed deployment** where Sisense is installed on multiple machines, but only part of the components are active on each one. This type of configuration is used to support heavier traffic loads and [high availability](#).

The Sisense full-stack solution is provided in a single installation process. After installing Sisense on a server, you will have access to all of the Sisense functionality.

The diagram below describes the main components of a Sisense installation and how Sisense supports a full-stack solution:



In this diagram, the yellow components specify Sisense client applications or APIs. The blue components specify components of the Sisense server. The Windows service and application names that you can see in the Windows Task Manager under the **Services** or **Processes** tabs appear in a dark blue font. When Sisense is installed on a single server, the dotted line indicates the boundaries of the server firewall. The lines between the components indicate communication between them, and its direction. When the Sisense Server listens on a specific port for inbound traffic, the relevant ports must be open in the firewall for the communication to succeed.

The diagram above describes the ports used by the services, for more information about communication ports, click [here](#).

# Sisense Architecture

Sisense supports a full stack solution from data, such as transaction details, user information, and machine-generated data, to dashboards accessible by Viewers. There are two sides to this solution, the management of data and the analytic side where users can interact with the data.

This section provides a description of the Sisense architecture as part of the full stack solution, from data sources to dashboards.

## Managing Data

The Sisense Server is responsible for managing data. Before you can begin to manage data however, Sisense communicates with remote servers for importing data sources, licensing information, tracking and similar activities.

### **Sisense Remote Servers**

Sisense remote servers are used for license management, monitoring and software updates.

### **Data Sources**

Data sources may be files or databases located on servers on a local network, or remote locations and web services such as Salesforce and Google AdWords.

Sisense supports a variety of data sources through native connectors, customer REST connectors, and ODBC drivers. These connectors are used to import data into the ElastiCube Server.

### **Sisense Server**

The Sisense Server manages data and supports the Sisense Application Server, which hosts the Sisense Web Application. In a Single deployment, the Sisense Server resides locally on your machine. In a Distributed deployment, the Sisense Server can reside on one of your machines or across many machines depending on your needs. For example, if you want to support a high availability

environment, you may install the ElastiCube Server on one machine, the Application Server on another, and the Multi-Node Orchestration server on another machine.

The Sisense Server is comprised of the following components:

1. **Licensing:** This service checks and validates licensing information.
2. **Multi-Node Orchestration:** The Sisense Orchestrator Service is an automated service that can be configured to synchronize builds across the ElastiCube Set.
3. **Remote Support Analysis:** This service is responsible for controlling logging, monitoring and automatic software updates.
4. **Application Server:** The Sisense Application Server is installed locally on your computer and hosts the Sisense Web Application. For more information, [click here](#).
5. **ElastiCube Server:** The ElastiCube Server manages the process of importing data from various data sources and ETL processes. For more information, see ElastiCube Server below.

## ElastiCube Server

The ElastiCube Server supports the management of ElastiCubes and the data they contain. The Sisense Server also supports the Application Server, which queries the ElastiCube Server where the ElastiCubes are hosted.

Within the ElastiCube Server are the following main components:

1. **ElastiCube Catalog:** The catalog maintains a list of ElastiCubes contained within the ElastiCube Server.
2. **Data Connectors:** The Data Connectors service is responsible for managing both native and 3rd-party connectors.

The ElastiCube Server supports the following applications:

1. **ElastiCube Manager:** A Windows client application for managing data sources. The ElastiCube Manager provides the GUI for creating ElastiCubes, including importing data sources, preparing them, and building the ElastiCubes.
2. **Sisense Shell (PSM.exe):** A command line interface for managing ElastiCubes, such as building an ElastiCube, attaching and detaching ElastiCubes from a server, updating data sources, etc. The Sisense Shell can also be controlled programmatically.

3. **Sisense Server Console:** A Windows application that controls Sisense servers, including attaching ElastiCube folders to Sisense servers and starting and stopping ElastiCubes.

## MessageBus

The Message bus is a standalone component responsible for transporting events across Sisense components.

## Interacting with Data

After you have created your ElastiCubes, your users can interact with data through Sisense dashboards.

Dashboards are hosted on the Sisense Application Server.

### **Application Server**

The Application Server provides access to Sisense dashboards and Sisense Admin console for user management.

The Sisense Application Server hosts the Sisense Web Application, which provides the user interface from which users can design, share, view and explore ElastiCubes and dashboards on the Sisense Web Application and the Sisense Mobile App.

Each dashboard created is maintained within the Application Server catalog and application database, which contains metadata regarding users, groups, dashboards and system settings.

The Application Server is comprised of many additional services required by the Sisense front-end, including a Proxy service for routing application requests, plug-in management, and more.

There are three main client applications that communicate with the Application Server:

1. **Sisense Web Application:** This Sisense Web Application is the client application for dashboard creation and access, and system administration.

2. **Sisense Mobile App:** This Sisense Mobile Application is an Android and Apple-device compatible application for viewing viewing dashboards across devices.
3. **REST API:** The Sisense REST API provides programmatic access to the Sisense Application Server functionality including managing ElastiCubes, user and group management, white labeling, user roles, reporting, and security.

# Backing Up the Sisense Application Database

The Sisense application database installed with Sisense contains metadata for ElastiCubes, dashboards, and users. As your metadata can change frequently, Sisense recommends backing up your application database regularly.

Sisense provides scripts that will back up your application database and restore it if necessary. To back up your application database, you need to create an application database user with WriteUser access and then replace the credentials in the script. Once the credentials have been updated, you can schedule the scripts to back up your database with the Windows Task Scheduler.

This page describes how to back up the application database with scripts prepared by Sisense and schedule the back up with Windows Tasks Manager.

**Note:** This page is relevant for Sisense V6.7 and later.

## Downloading the Scripts

There are two scripts Sisense provides one for backing up the application database and a second for restoring a back up copy of your application database.

You can download the scripts below:

[MongoBackup.bat](#): This script is used for backing up your application database.

[MongoRestore.bat](#): This script is used in case you need to restore the backed up information.

# Activating a WriteUser

To be able to access the application database, you need to have a WriteUser to your application database. See [Accessing the Sisense Application Database](#) for instructions on how to activate a user with WriteUser access to your Sisense application database.

# Changing the Credentials

After you have activated a WriteUser for the Sisense application database, you need to replace the credentials in the relevant script file.

The example below displays where you need to add your WriteUser's credentials. Following **-u**, replace **WriteUser** with your application database user's username and after **-p**, replace **Sisense** with your password.

```
c:\
cd C:\Program Files\Sisense\Infra\MongoDB
mongodump --host localhost:27018 --db prismConfig --out C:\SisenseWebBackups\MongoBackup --authenticationDatabase admin -u WriteUser -p Sisense
mongodump --host localhost:27018 --db prismWebDB --out C:\SisenseWebBackups\MongoBackup --authenticationDatabase admin -u WriteUser -p Sisense
```

# Scheduling Back Ups

The **MongoBackup.bat** is the script you should use to schedule periodic backups of your application database. The **MongoRestore.bat** should not be scheduled as you use it only to restore a back up file.

The easiest way to schedule a back up is with Windows Task Scheduler. For information, see [Schedule a Task](#).

# Communication Ports

Sisense utilizes multiple communication ports. Some of the ports are used for communication between Sisense modules, and others are used for communication with remote locations.

When Sisense is installed as part of a Single machine deployment, ports used for communication between Sisense modules are considered internal ports, and do not need to be open in the firewall. When Sisense is installed in a Clustered deployment, where Sisense modules installed on different computers communicate with each other, ports used for communication between Sisense modules may need to be opened in the firewall. This depends on the specific deployment implemented in your environment.

The tables below lists the ports on which Sisense listens for communication.

## Ports for External Communication

The following table list ports used for communication with external locations. These ports should be open on your firewall to allow for this communication.

Module	Port Number	Protocol	Description
Web Server	8081(*)	HTTP	Sisense Web Server
Sisense.Orchestrator	445, 137-139	SMB	Used in high availability deployments. Otherwise, not used. In high-availability, it copies DB farm to

Module	Port Number	Protocol	Description
			query nodes. The Sisense Orchestrator Service generates inbound (query nodes) and outbound (build node) traffic.
Sisense.Orchestrator	22(*)	SSH	Used in high availability deployments. Otherwise, not used. In high-availability, it copies DB farm to query nodes. The Sisense Orchestrator Service generates inbound (query nodes) and outbound (build node) traffic.
OAuth2 server	7077 In Sisense V6.7, port 7043 is also used.	HTTP In Sisense V6.7, HTTPS is also used.	Used by GenericREST connectors that use oAuth authentication. Otherwise, not used.

(\*) Indicates that the port is configurable.

# Ports for Internal and External Communication

The following table lists all ports in which the type of communication, internal or external, depends on your deployment. In a single deployment, Sisense communicates internally to other services. These ports do not need to be open on your firewall. In a clustered deployment, Sisense communicates across your deployment to other Sisense services. These ports should be open on your firewall to allow for this communication.

Module	Port Number	Protocol	Description
ElastiCube Management Service, QueryServiceTcpPort	812	TCP	In high-availability, it's the ElastiCube query end point
ElastiCube Management Service, QueryServiceHttpPort	16162	HTTP	In high-availability, it's the ElastiCube query end point
ElastiCube Management Service, ManagementServiceTcpPort	811	TCP	Elasticube Management Service endpoint
ElastiCube Management Service, ManagementServiceHttpPort	16161	HTTP	Elasticube Management Service endpoint
Oxygen	31112	TCP	
MongoDB	20718(*)	TCP	

(\*) Indicates that the port is configurable.

## Ports for Internal Communication

The following table lists all ports on which Sisense communicates internally.

These ports do not need to be open on your firewall.

Module	Port Number	Protocol	Description
MonitoringMongoDB	27019	TCP	
Sisense.CLRConnectorsContainer Server	8090	TCP	CLR container server: Data transport between EC and data sources
Sisense.JVMConnectorsContainer Server	Dynamically allocated	TCP	JVM container server: Data transport between EC and data sources
Sisense.JVMConnectorsContainer Management	8095	TCP	Container management - e.g. Adding/Removing/Updating connectors
Sisense.Discovery	2181	TCP	Discovery service for connectors (Zookeeper)
Sisense WebServer (Node.JS)	40001	HTTP	Webserver Note: This module is relevant from Sisense V6.7 onwards.
Sisense.WebQueryService (.Net)	40002	HTTP	WebQueryService Note: This module is

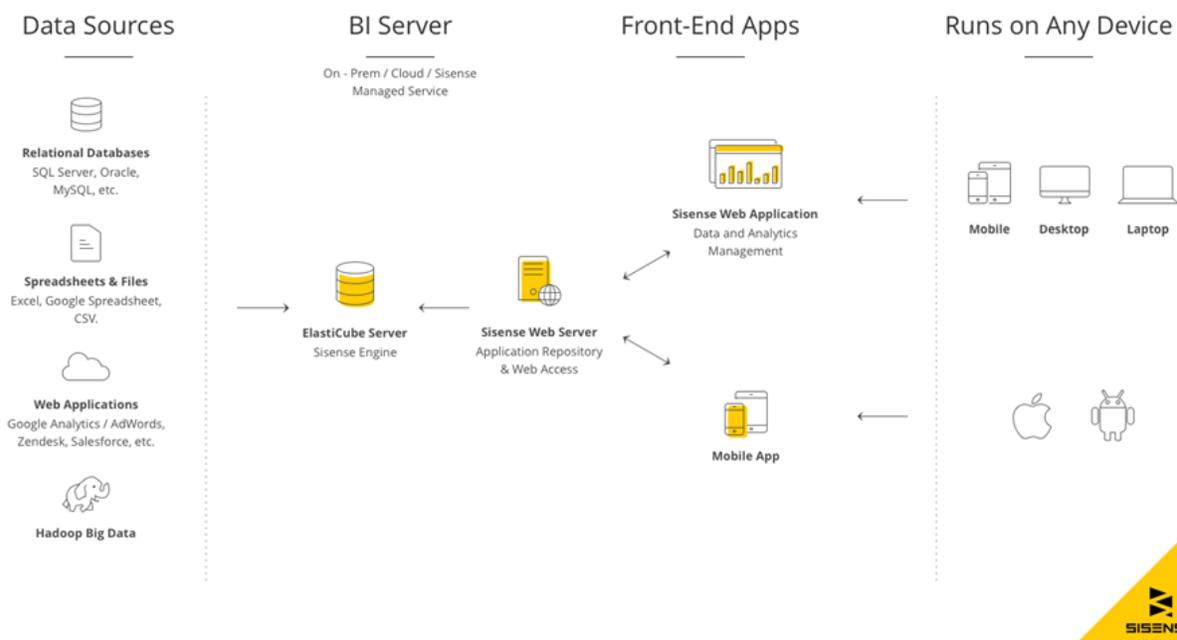
Module	Port Number	Protocol	Description
			relevant fro, Sisense V6.7 onwards.
Streaming Processing Engine	40003	HTTP	Processing Engine (Node.js)
CustomREST MongoDB	27077(*)	TCP	Intermediate and cache db for CustomREST: connector results
Sisense.ECMLogsPersistenceService	5672 and 27018	AMQP	Reads build logs and saves them on the application database

(\*) Indicates that the port is configurable.

# Sisense Basic Concepts and Terminology

## Sisense Architecture

The Sisense system is comprised of the following components.



## ElastiCube Manager

The ElastiCube Manager is a visual environment in which you create ElastiCubes. The ElastiCube Manager is available online or locally on the desktop where Sisense is installed. The ElastiCube Manager enables you to structure and import multiple data sources, create relationships between data, and perform ETL (Extract, Transform, and Load) processes to prepare data for analysis and visualization. For more information, see [ElastiCube Manager](#).

## ElastiCube Server

The ElastiCube Server is installed locally on your computer and enables access to ElastiCubes. Both the Sisense Web Application and the ElastiCube Manager query the ElastiCube Server and receive results.

## Sisense Server Console

The Sisense Server Console provides administration functions for managing ElastiCubes on the ElastiCube Server.

## Sisense Web Application

The Sisense Web application is an interactive web application that provides the user interface in which users can design, share, view and explore dashboards. The Sisense Web application runs in popular web browsers and enables access to the Sisense server. Dashboards can also be viewed on mobile devices.

## Sisense Web Server

The Sisense Web Server is installed locally on your computer and hosts the Sisense Web Application.

## ElastiCubes

ElastiCube are Sisense's proprietary super-fast data stores. ElastiCubes are specifically designed to withstand the extensive querying typically required by business intelligence applications.

ElastiCubes enable you to mash up terabytes of data from a variety of sources, for example:

Traditional relational databases, such as SQL Server, MySQL or Oracle.

File-based data sources, such as spreadsheets (Excel) and CSV files.

Online web services, such as Salesforce.com, Google AdWords, Google Analytics, Zendesk and more.

ElastiCubes are based on Sisense's proprietary In-Chip technology, which leverages the Sisense engine. Read more about the technology behind ElastiCube in Sisense's [technology section](#).

ElastiCubes are created and managed in the ElastiCube Manager.

## ElastiCube Sets

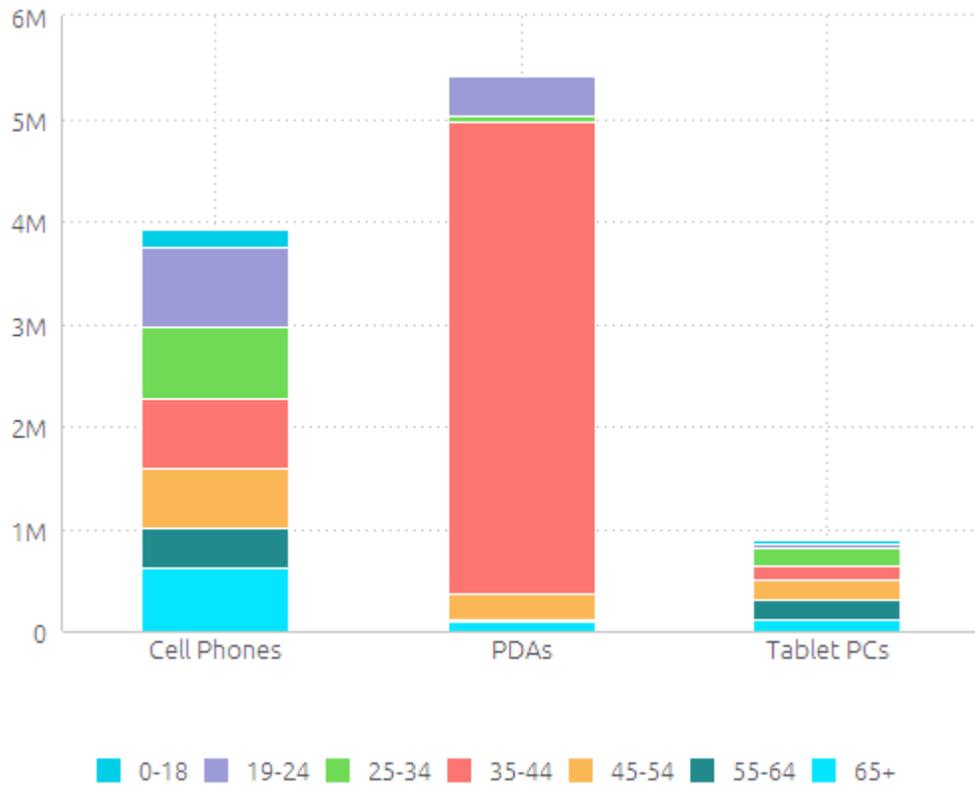
Sisense ElastiCube Sets are collections of ElastiCubes with identical schemas that support high availability by allowing you to query running ElastiCubes within the ElastiCube Set while other ElastiCubes are building.

For more information see [ElastiCube Sets](#).

## Widgets

Each widget is a dynamic visualization of data. You pick the type of data to appear in a widget and you pick the type of visualization (chart type). A few examples of widgets are displayed below:

Column Chart Widget



## Indicator Widget

GROWTH OF VISITORS

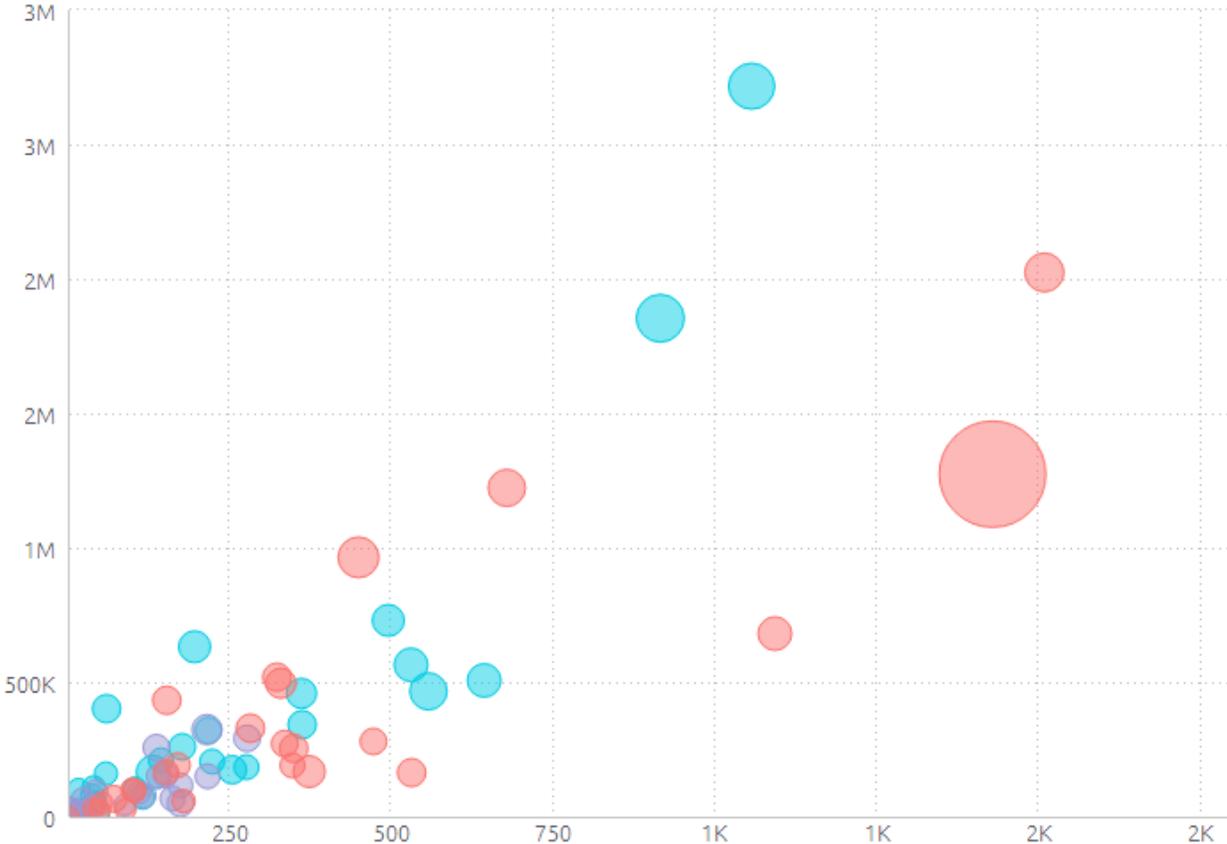
13.2%

# of Visitors 220K

# Pivot Widget

Category	Female		Male	
	Total Revenue	Total Quantity	Total Revenue	Total Quantity
Apple Mac Desktops	1,151	4	26,071	47
Apple Mac Laptops	5,310	18	50,717	110
Calculators	609	15	3,147	38
Camcorders	2,790	23	26,890	124
Camera Flashes	1,168	12	16,140	115
Car Amplifiers	734	6	4,619	32
Car Speakers and Subwoofers	612	9	6,968	90
Cell Phones	94,323	638	605,011	2,122
Digital Cameras	34,045	153	216,707	669
DVD Players	1,488	23	5,309	85

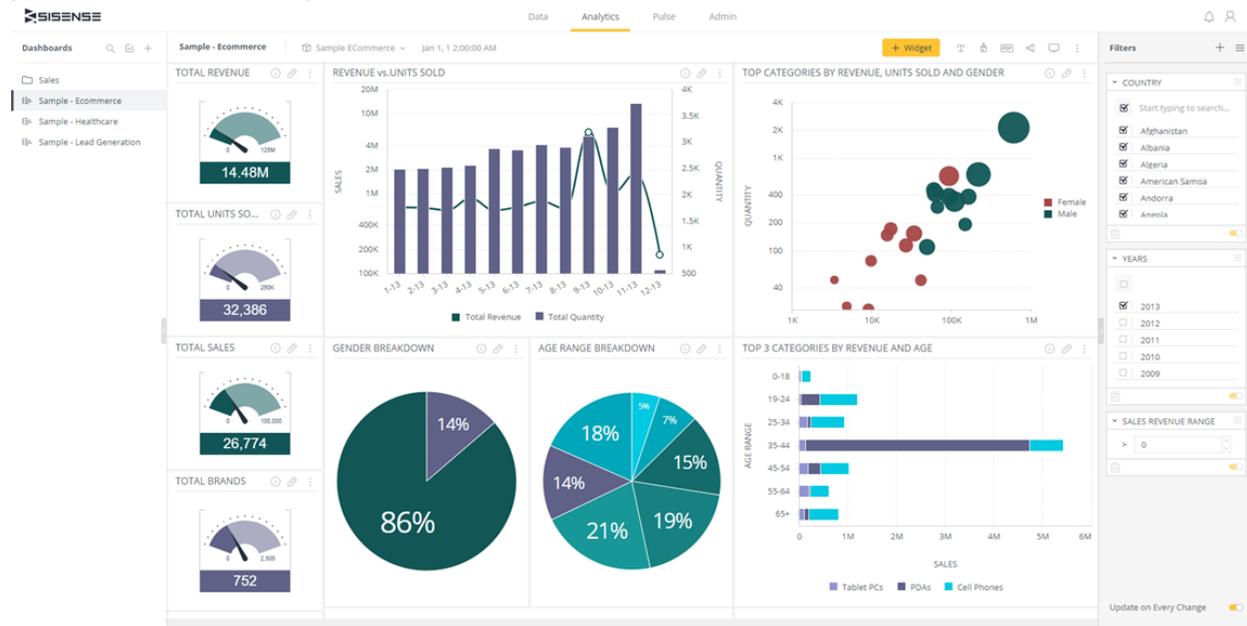
# Scatter Chart Widget



To learn more about the different widget types, [click here](#).

# Dashboards

A dashboard is a collection of one or more widgets that visualize the data that you select and design.



You define which widgets appear in the dashboard, their design, how they are organized and the filtering of the data that appears.

You can access your own dashboards and dashboards others shared with you.

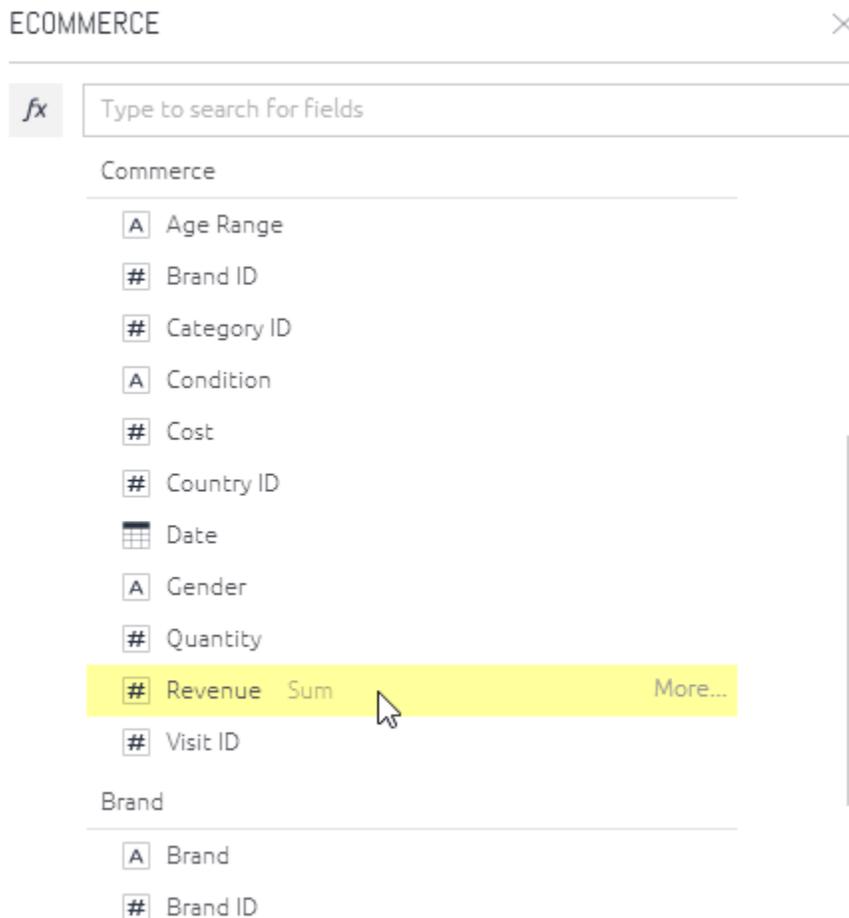
Sisense allows you to create as many dashboards as you need. No limitations!

When you create a dashboard, you are the dashboard's owner. Owners can share dashboards that they have created with other Sisense users.

Sisense provides a variety of built-in automatic dashboard sharing features. You can easily share a dashboard with others in the Sisense environment or have a dashboard automatically delivered to your chosen recipients by email (on a scheduled basis or upon each data update).

# Fields

Widgets are composed of fields that represent the data in your data sources. You create widgets by simply selecting from the fields displayed in a Data Browser, which appears in various places across the product. Each field represents a column of data in the ElastiCube.



When designing widgets, fields can be categorized into three groups:

- **Numeric Fields:** Numbers (quantified data), such as salaries, sales, scores, number of clicks and so on. This is data that you may want to aggregate or calculate. For example, the sum of sales or the average of costs.
- **Date Fields:** Dates can describe both date and time values. Dates can be used to organize your data into hierarchies according to year, quarter, and month or into buckets of time such as by hour or 15 minute intervals. For more information, see Working with Time.

- **Descriptive Fields:** Items used to label and categorize, such as Products, Locations, Categories and so on.

Generally, widgets combine both these types of fields.

For example, to show the sales of a product over time, Sales is the numeric data and both Products and Time are descriptive data.

Descriptive data can be considered as numeric when it is simply a count of the number of items. For example, Activities can be considered numeric when it is a count of the number of Activities. A number can be considered as descriptive when it is an actual entity. For example, an identification ID or phone number.

For information regarding Sisense internals and a more advanced explanation of the Sisense architecture, click [Sisense Architecture Internals](#).

# Configuring Base URLs

Base URLs are the consistent part of a web address for a site or web application. In Sisense, the default address is localhost:8081. Sisense allows you to specify a Sisense hostname as a subfolder of a domain, for example, `baseurl.sisense.com/reporting`.

The result is that if your users would be accessing the Pulse page for example, instead of the default `baseurl.sisense.com/pulse`, they would see the URL `baseurl.sisense.com/reporting/pulse`.

To implement base URLs in Sisense, you need to direct your users to a reverse proxy that is configured to redirect the users to the appropriate base URL that you define in Sisense.

The implementation of a reverse proxy is done on your side. An example reverse proxy configuration file for Nginx is provided below.

Implementing a base URL on Sisense's side requires updating the `default.yaml` configuration file on your Sisense server.

## To configure a new base URL in Sisense:

1. With a text editor, open the `default.yaml` file located at `...\Program Files\Sisense\PrismWeb\vnext\config`
2. Add the property `proxyurl: 'yourbaseURL'` to the `webServer` object. For example:

```
webServer:
  ver: '5.0'
  port: 40002
  iisPort: 8081
  virtualDir: '/api'
  apiVersion: 'Beta'
  oldApi: '/apiv1'
  secure: false
  reportPath: \sisense\reports\
```

```

jobEnable: true
trackEnable: true
dashboardPath:
'http://localhost/app/main#/dashboards/'
proxyurl: https://baseurl.sisense.com/reporting
enableSSL: false
jsonLimit: '25mb'

```

### 3. Save your changes.

There are a variety of reverse proxies you can choose to implement a base URL.

You can see an example Nginx configuration below (nginx.conf):

```

worker_processes 1;
events {
    worker_connections 1024;
}
http {
    include mime.types;
    default_type application/octet-stream;
    sendfile on;
    keepalive_timeout 65;
    proxy_connect_timeout 6000;
    proxy_send_timeout 6000;
    proxy_read_timeout 6000;
    send_timeout 6000;
    server {
        #listen 80;
        server_name site.corporate.com;
        listen 443 ssl;
        ssl_certificate "location of pem file";
        ssl_certificate_key "location of key file";
        ssl_protocols TLSv1 TLSv1.1 TLSv1.2;
        ssl_ciphers HIGH:!aNULL:!MD5;
        rewrite ^/reporting$ /reporting/ last;
        location ~/reporting(/.*)$ {
            set $path $1;
            client_max_body_size 2G;
            proxy_pass
http://sisenseIP:sisensePort$path$sis_args$args;
            proxy_http_version 1.1;
            proxy_set_header Upgrade $http_upgrade;

```

```

        proxy_set_header Connection "upgrade";
        proxy_set_header X-Real-IP $remote_addr;
        proxy_set_header X-Forwarded-For
$remote_addr;
        proxy_set_header Host $host;
        proxy_connect_timeout          6000;
        proxy_send_timeout              6000;
        proxy_read_timeout              6000;
        send_timeout                    6000;
        #if ($http_origin ~*
'^ (http://origin1.com:8081|https://www.origin2.org)$')
{
        # set $cors_origin '${http_origin}';
        # set $cors_cred 'true';
        # set $cors_methods 'GET, POST, DELETE,
PUT, OPTIONS, PATCH, HEAD';
        # set $cors_headers 'Content-Type, X-
Device-Id, X-API-KEY, Origin, Authorization, X-
Requested-With, Accept, Key, Cache-Control
        #}
        add_header Access-Control-Allow-Origin
$cors_origin always;
        add_header Access-Control-Allow-Credentials
$cors_cred always;
        add_header Access-Control-Allow-Methods
$cors_methods always;
        add_header Access-Control-Allow-Headers
$cors_headers always;
        if ($request_method = OPTIONS) {
        return 200;
        }
    }
}
}

```

# Customizing the Installation

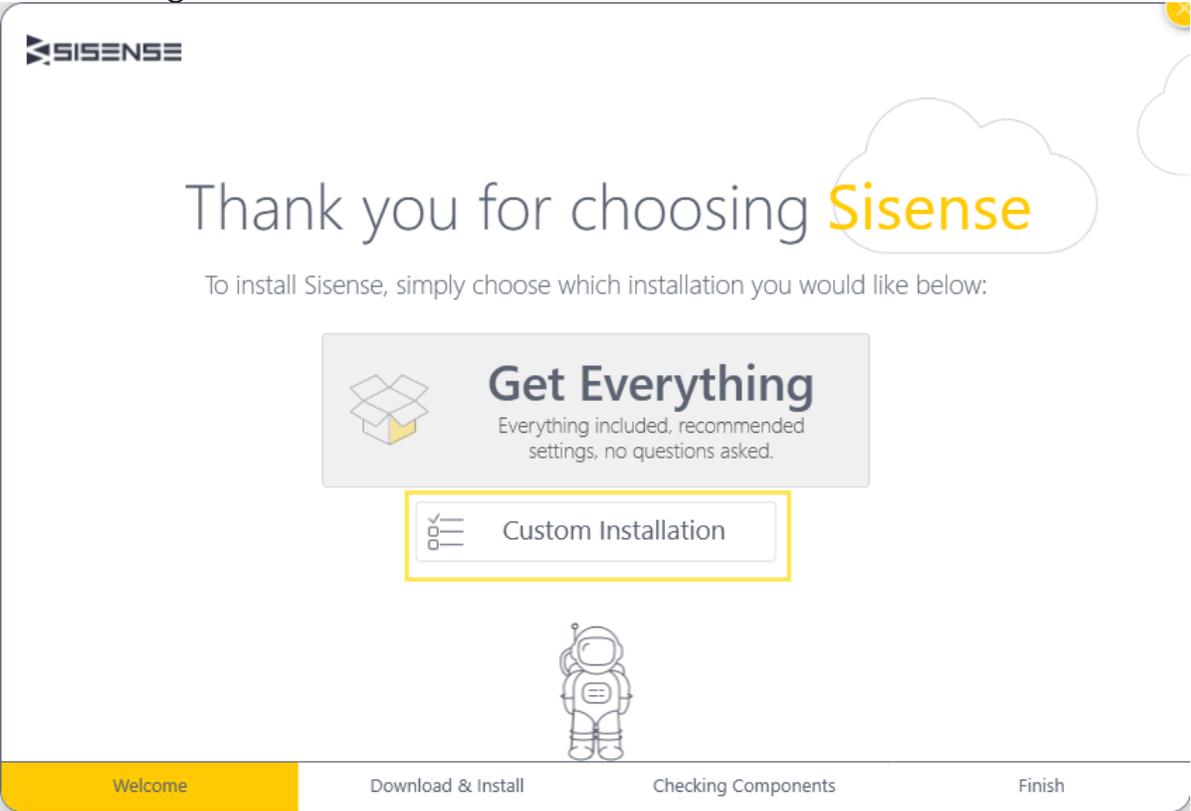
This topic describes how you can customize the settings for your web server application during the installation process or after Sisense is installed.

For example, you can configure your server to use the Long index edition of Sisense if your ElastiCubes are expected to contain a table with over 300 million rows.

**Note:** You cannot customize the settings during an upgrade process. If you upgraded Sisense and want to customize settings, refer to the post-installation option described below.

**To access the customization options during the installation process:**

- While [installing Sisense](#), click **Customize Installation**, and define the settings as described below.



## To access the customization options after you have already installed

Sisense:

1. Open the Control Panel, and go to **Programs > Programs and Features**. Right-click **Sisense**, and click **Change**.
2. Click **Continue > Change Features**, and define the settings as described below.

**SISENSE**

# Custom Installation

ElastiCube Server  Long Index edition [more info](#)

Sisense Web Application

Web Server Configuration :

Port :

Website Name :

Web Server Type :

Microsoft IIS Express

Microsoft IIS

Manual

Back < > Continue

Welcome | Download & Install | Checking Components | Finish

## Settings

### ElastiCube Server

Select the [Long index](#)<sup>1</sup> edition if your ElastiCubes are expected to contain a table with over 300 million rows.

---

<sup>1</sup> The 'Long Index edition' of Sisense determines how Sisense transforms your data when loading it in to an ElastiCube. Sisense builds your ElastiCube using pointers that represent your data. These pointers are used to find the actual values so no information is lost. To reduce

It is not recommended to use this configuration setting for ElastiCubes in which the largest table does not include at least 300 million rows.

### **Sisense Web Application**

You must specify the Sisense Web Application **Port** and **Website Name**. Usually, if no other sites are running on the machine, you can use port 8081.

There are three options for the server type:

- **Microsoft IIS Express:** This is the default web server used to host the web application.
- **Microsoft IIS:** This is recommended for production environments. In case there is an existing site with the same name, it will be overridden.
- **Manual:** This is recommended when you have to manually set up your IIS. If this is selected during upgrades, the installation will keep the existing site details and configuration in the IIS. When using this setting, make sure to define the physical path of the IIS site to the new installation folder (C:\Program Files\Sisense\PrismWeb).

---

memory consumption and improve performance, Sisense builds ElastiCubes with 32-bit pointers reducing the size of the pointers. You may experience “Out of memory” errors when working with big data such as: 300 million rows of numeric fields 2GB of string fields For example, let’s assume that your data contains strings with a length of 1024 characters and 4,194,304 unique rows such that  $1024 * 4,194,304 > 4\text{GB}$  in content after indexing, you should install the Long Index version of Sisense. By installing the Long Index edition, Sisense will use 64-bit pointers. While this is fine for big data, you should use the default installation of Sisense when working data smaller than the figures mentioned above for faster performance.

# Downloading and Installing Sisense

This topic describes how to download Sisense and the installation process for getting Sisense up and running.

## Downloading Sisense

If you want to try Sisense for the first time, sign up for the free trial at [Sisense.com](https://sisense.com).

Otherwise, you can download the latest version of Sisense [here](#).

## Installing Sisense

The entire Sisense solution can be installed from one file, either locally or in a central place in your organization within a couple of minutes.

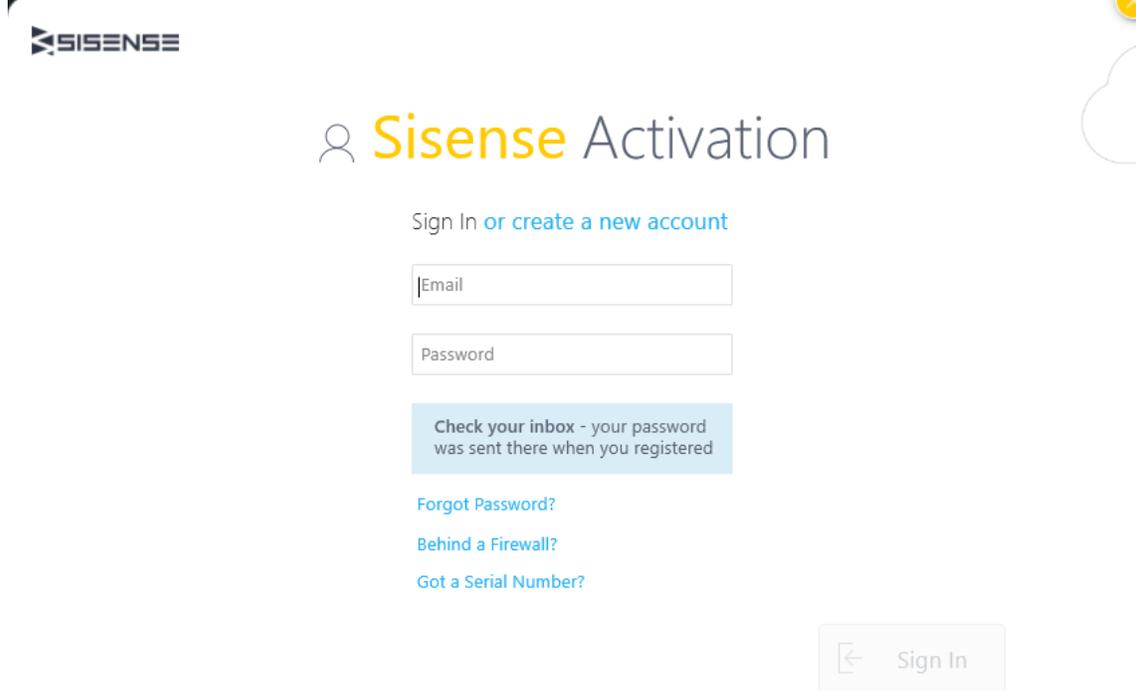
All Sisense components are installed with a default configuration, as follows:

- Sisense Web Application
- ElastiCube Manager
- Sisense Server Console
- ElastiCube Server
- IIS Express as the web server hosting Sisense Web Application platform

**Note:** To install Sisense, you must have administrator privileges on the machine where you are installing Sisense and your environment must meet the [Minimum Requirements and Supported Platforms](#).

**To install Sisense:**

1. Open the downloaded executable file to run the installation.  
**Note:** If your Window Firewall is active, click Allow access to continue installing Sisense.
2. If you are not logged into Sisense, enter your username and password to confirm licensing and begin the installation process. If you are logged in already, the activation screen is not displayed.



3. After entering your login details, select **Get Everything** to run the default installation:



# Thank you for choosing **Sisense**

To install Sisense, simply choose which installation you would like below:



## Get Everything

Everything included, recommended settings, no questions asked.



Custom Installation



Welcome

Download & Install

Checking Components

Finish

If you are upgrading your version, click **Upgrade**.



# Gideon, Welcome Back!



## Continue

Use your current account and continue with Upgrade



My Account

By clicking the button above you agree to our [Terms and Conditions](#)

## Installing Sisense behind a Firewall

The standard installation requires internet access to download the necessary components. For machines behind a firewall, or without internet access, you can download and install the full self-contained installation file.

For detailed instructions, see [Installing Sisense Offline](#).

### **Related Topics:**

If you need to customize your installation, see [Customizing the Installation](#). You can customize your installation during the first installation or after an installation, but not during an upgrade.

If you want to install Sisense on a remote machine, see [Performing a Silent Installation](#).

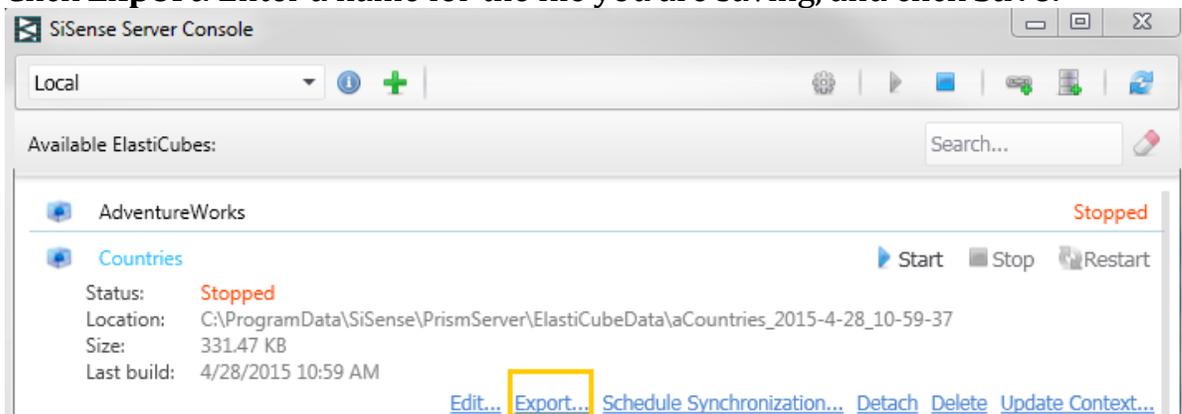
If you are migrating your installation to a different machine, see [Backing Up and Migrating a Sisense Installation](#).

# Importing and Exporting Elasticube Data

ElastiCube data can be transferred to different machines. This is useful when you are deploying on a new machine or need to transfer the data to a different server. The transferred data is stored in a compressed eedata file. The file contains all of the information contained in an ElastiCubeData folder (including the source data), and is used by the Sisense Server Console to build a new ElastiCube.

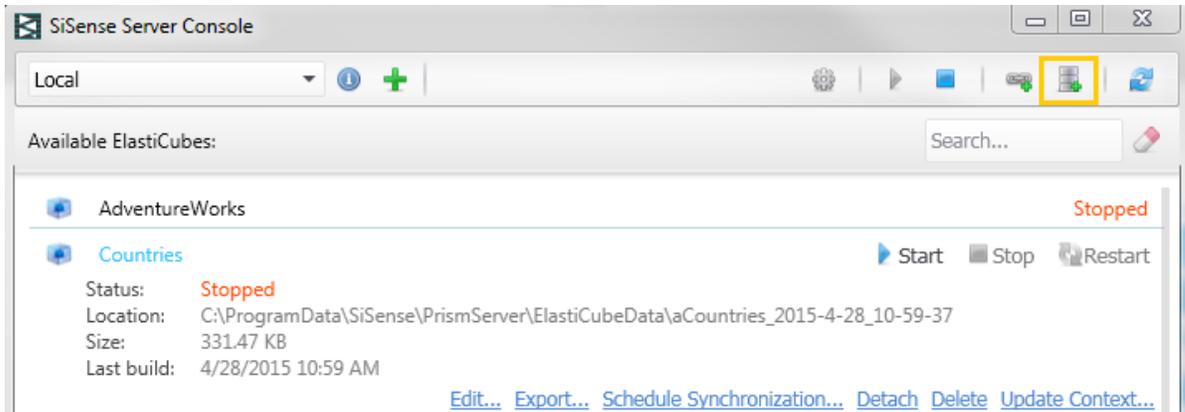
## To export data:

1. Open the Sisense Server Console.
2. Click on the relevant ElastiCube.
3. Click the 'Stop' button to stop the specific ElastiCube (not the entire server).
4. Click **Export**. Enter a name for the file you are saving, and click **Save**.



## To import data:

1. Open the Sisense Server Console.
2. Click on the import eedata file icon .
3. Select the location of the eedata file and click **Open**.
4. A window will appear showing the import progress. Once complete, the new ElastiCubes will be added to the list of existing ElastiCubes. Make sure the ElastiCube is running before accessing it.



**Note:** You can also transfer data by copying the folder where Sisense stores the data locally, instead of exporting and importing the compressed file. To learn more, click [here](#).

# Launching Sisense

This topic describes how to launch the Sisense Web Application and ElastiCube Manage.

## To launch Sisense:

Use one of the following options to launch Sisense:

- If you have Sisense installed on your machine, select Sisense from the **Start** menu.
- From a standard browser, go to the URL of the Sisense environment provided to you by your system administrator.

## To launch the ElastiCube Manager:

Use one of the following options:

**From Windows:** Open the Windows **Start** menu and select Sisense ElastiCube Manager. This opens the desktop version of the ElastiCube Manager.

**From the Sisense Web Application:** Click **Data** in the top menu. This opens the Data page where a list of your ElastiCubes is displayed.

## To log into Sisense:

- If you installed Sisense, you can use the same credentials to log into Sisense.
- If an account has been created for you, you should receive a password activation email. Click on the link to activate your account.

# Backing Up and Migrating a Sisense Installation

This topic describes two separate, but related procedures, backing up a Sisense installation and migrating a Sisense installation from one machine to another. Backing up your installation is recommended before a migration process, but also on a regular basis for securing your data and Sisense configurations.

You may need to migrate your entire setup when moving from a staging/development environment to a production environment, when upgrading hardware, etc.

**Note:** If you have advanced and/or customized web configurations, such as SSL, SSO, redirects, etc., please consult Sisense Support through the [Help Center](#).

## Backing Up a Sisense Installation

The backup process requires backing up your ElastiCubes and the Sisense Web Application.

In addition, if you have any custom plugins located in the directory C:\Program Files\Sisense\PrismWeb\plugins, this folder should be backed up as well.

The procedures below describe how to back up your ElastiCubes and the Sisense Web Application.

### Backing Up Your ElastiCubes

Your ElastiCube Server contains information about your ElastiCubes, including the schema and the data that was imported when the ElastiCube was built. This

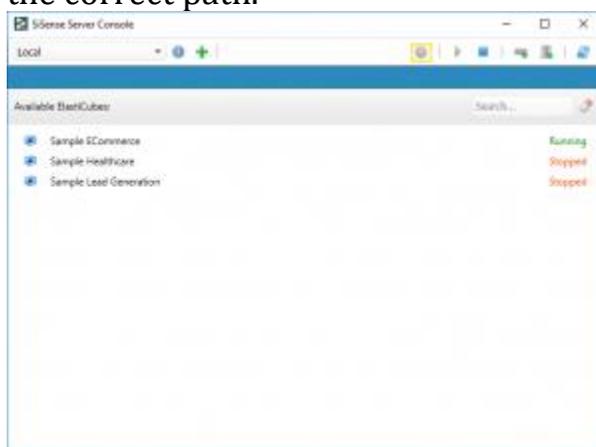
information is saved in a collection of folders stored in the ElastiCube Data folder. The files and folders in the ElastiCube Data must be saved and their structured maintained to ensure your implementation will work as expected on the migrated machine.

### To back your ElastiCubes:

1. Open the directory:

`%ProgramData%\Sisense\PrismServer\ElastiCubeData`

**Note:** This is the default path. To check which path is being used, open the Sisense Server Console and click on the Server Preferences icon  to verify the correct path.



2. Copy the data to your backup location.

## Backing Up the Sisense Web Application

The Sisense Web Application uses information stored in an instance on the Sisense application database where the Sisense Web Application is installed. This instance stores data about users, dashboards, widgets, and other information. This information must be stored to ensure your Sisense Web Application runs properly.

### To back up the Sisense Web Application:

1. Open the directory:  
`%ProgramData%\Sisense\PrismWeb\Repository\DB`
2. Copy the data to your backup location maintaining the same file structure.
3. (Optional) Back up your custom plugins by copying the directory:  
`C:\Program Files\Sisense\PrismWeb\plugins`

and saving it in the appropriate folder the new server.  
For information on how to automate the backup process, see [Automatically Backup Sisense Web Data](#).

## Migrating a Sisense Installation

The migration process involves migrating both the Sisense application database and your ElastiCubes.

Additionally, if you are using plugins, you will need to replace the Plugins folder with your current plugins located at C:\Program Files\Sisense\PrismWeb\plugins.

### Migrating ElastiCubes

The ElastiCube Server contains your ElastiCubes. To migrate ElastiCubes, you must export them as .ecdata files and import them into your new environment. If the data sources that support the ElastiCube are at a different location than they were on your original server, the ElastiCube's connectivity settings must be updated.

#### To migrate your ElastiCubes:

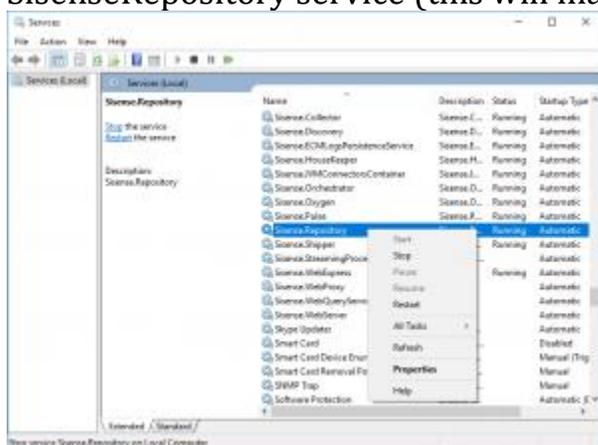
1. Declare downtime to avoid the loss of work during the process and stopping the IIS service. You can stop the IIS service by running the command **iisreset /stop** in the Command Prompt as an Administrator. Do not close the Command Prompt as you will need it in Step 3.
2. Export and import your .ecdata files. See [Importing and Exporting Elasticube Data](#) for more information.
3. In the Command Prompt, run the command **iisreset /start** to restart your IIS Server.
4. If the database location was changed as well, modify the connectivity settings. See [Change Connectivity Settings](#) to learn how.

## Sisense Web Application

To migrate your Sisense Web Application configuration, you must copy several configuration files from the current environment and migrate these files to the new environment in the same directories. These configuration files are used by Sisense to allow communication between the Sisense application database and the Sisense Web Application.

### To migrate the Sisense Web Application:

1. Back up the old environment, open your machine's Control Panel, and go to System and **Security > Administrative Tools > Services**, and stop the **SisenseRepository** service (this will make the website inaccessible).



2. Create a copy of the entire Repository directory located at:  
`%ProgramData%\Sisense\PrismWeb\Repository`
3. Replace the existing **Repository** directory in the new environment with the one you backed up.
4. Create copies of the following configuration files and paste them into the same directories in the new environment:  
`%ProgramFiles%\Sisense\PrismWeb\ECMNext\GraphQL\src\config\default.yaml`  
`%ProgramFiles%\Sisense\PrismWeb\vnext\config\default.yaml`  
`%ProgramFiles%\Sisense\Infra\MongoDB\mongodbconfig.conf`  
`%ProgramFiles%\Sisense\PrismWeb\App_Data\Configurations\db.config`

**Note:** In each of the default.yaml files are the application database user credentials that the Sisense Web Application needs to access the Sisense

application database. In these files, verify the values of user and password match.

```
default.yaml - Notepad
File Edit Format View Help
# protocol: 'https'
server:
  port: 40005
mongo:
  host: 127.0.0.1
  port: 27018
  db: prismWebDB
  configDb: prismConfig
  user: AppUser
  password: jDH1X74ckttW00If75
  key: ThY3Nk31R3AzNUSvZThDMg==
  IV: ePth5UDfIH6YvSR01hwkcQ==
# replicaSet:
#   name: 'rs0'
#   readPreference: ''
#   servers:
#     - host: 127.0.0.1
#       port: 27017

default.yaml - Notepad
File Edit Format View Help
encryption:
  aesKey: 'luVPIFse6rgkYbxUuDKvWg=='
  IV: 'ePth5UDfIH6YvSR01hwkcQ=='
sisense:
  location: 'localhost'
  port: 8081
# protocol: 'https'
server:
  port: 40005
mongo:
  host: 127.0.0.1
  port: 27018
  db: prismWebDB
  configDb: prismConfig
  user: AppUser
  password: jDH1X74ckttW00If75
  key: ThY3Nk31R3AzNUSvZThDMg==
  IV: ePth5UDfIH6YvSR01hwkcQ==
```

5. Start the SisenseRepository service. Your Sisense configuration has now been migrated to the new environment.

# Minimum Requirements and Supported Platforms

The following prerequisites and supported platforms are required for working with Sisense.

## Supported Web Browsers

The Sisense Web Application runs in the following HTML5 supported browsers:

- Internet Explorer 10 and higher
- Google Chrome
- Firefox
- Safari version 7 and higher

**Note:** When embedding iFrames, Sisense supports Safari 10 and higher

The Sisense Web Application also works in mobile phone and tablet browsers that support HTML5. [Click here to learn more about mobile compatibility.](#)

## Supported Operating Systems

ElastiCube Server and ElastiCube Manager can be installed on the following 64-bit operating systems:

- Windows 7 and higher
- Windows Server 2008 R2 through Windows Server 2016

Note: While Sisense supports Windows 7, it is highly recommended that production environments use Windows Server 2008 and later.

# Capacity and Hardware Requirements

Sisense easily scales up to billions of records with typical query response times of split seconds.

This section suggests system requirements for various performance capacities of the ElastiCube Server. Actual capacity requirements are provided after consultation with a Sisense technical representative at [support@Sisense.com](mailto:support@Sisense.com).

Extreme scenarios may require additional resources.

Sisense Dashboard viewers only require an HTML5 compliant Web browser.

These hardware requirements are also true for cloud deployments. Whether you are using AWS/Azure/Rackspace/etc., you must choose a machine that meets the recommended hardware configuration.

	< 10 users		10s of users		100s of users	
# of rows	RAM (GB)	Logical Cores	RAM (GB)	Logical Cores	RAM (GB)	Logical Cores
Up to 100K	8	4	8	4	16	8
up to 500K	8	4	16	4	16	8
up to 1M	8	4	16	4	32	8
up to 10M	16	4	16	8	32	8
up to 50M	16	8	32	8	64	16
up to	25	8	32	8	64	16

<b>100M</b>						
<b>up to 200M</b>	32	8	32	8	64	16
<b>up to 300M</b>	32	8	64	16	64	16
<b>up to 500M</b>	64	16	128	16	256	32
<b>up to 600M</b>	128	16	128	32	256	32
<b>up to 700M</b>	128	16	128	32	512	32
<b>up to 800M</b>	128	16	256	32	512	32
<b>over 800M</b>	256	16	256	32	512	32

\* Users are defined as concurrent users of the system regardless of licenses.

## Sisense Recommendations

Sisense recommends that your Sisense server meets or exceeds the minimum requirements listed above. The actual requirements of your Sisense server may vary depending on the number of concurrent users, builds running in parallel, ElastiCubes hosted on a server, and additional factors specific to your server, for example, non-Sisense applications running on the same server.

For optimal performance, Sisense recommends:

- No more than 40 ElastiCubes on a single Sisense server
- No more than 4 concurrent builds at a time
- No more than 150 concurrent users

For optimal performance, Sisense highly recommends that your servers have a processor that supports AVX (Advanced Vector Extensions), which is leveraged by Sisense for improved query performance and user concurrency.

## Supported Locales

The formats for dates, times and numbers in your dashboards are based on your computer's operating system or browser's locale settings (depending on your browser).

Locales from the following countries are supported:

- Unites States
- United Kingdom
- Israel
- Canada
- South Africa
- Australia
- Netherlands
- Germany
- Ireland
- Mexico
- France
- China
- Brazil

### Notes

- You can override the default locale settings for all users by updating a parameter via the Rest API. [Click here to learn how to override default locale settings.](#)
- To use a locale file that is not included in the above list of countries, you can manually add a locale file in the following location: C:\Program Files\Sisense\PrismWeb\client\resources\base\localization. [Click here to access locale files and view their their codes.](#)

# Installing Sisense Offline

The standard installation of Sisense requires internet access to download additional necessary components. To support machines behind firewalls, or without internet access, you need to download and install the full Sisense [installation file](#). After you have installed Sisense, you need to provide an offline activation key.

The steps below describe how to install Sisense offline and activate it.

## Prerequisites

Before installing Sisense offline, make sure you have the following installed:

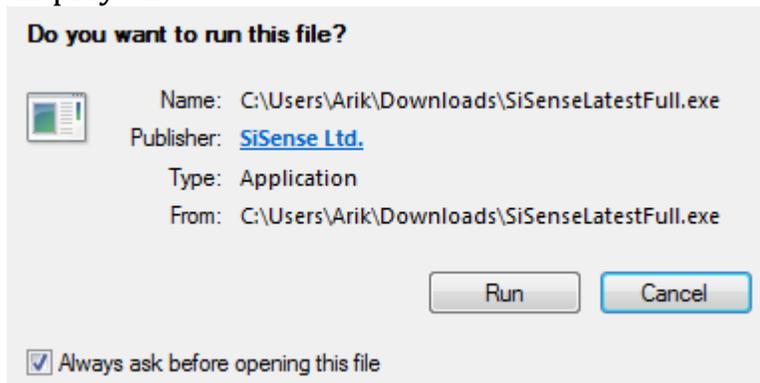
1. [Microsoft .NET 4.6.1 Framework](#)
2. **IIS - Internet information Services** (Microsoft's Web server), see the following instructions per your operating system  
[Windows 7 or 8](#)  
[Windows Server 2008](#)  
[Windows Server 2012](#)

## Installing Sisense

**To install Sisense offline:**

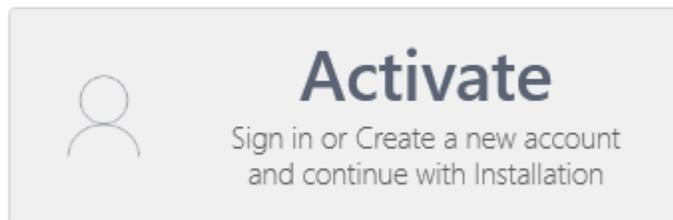
1. Download the [latest full version of Sisense](#).

2. Open the installation file and click **Run**. The Welcome to Sisense window is displayed.



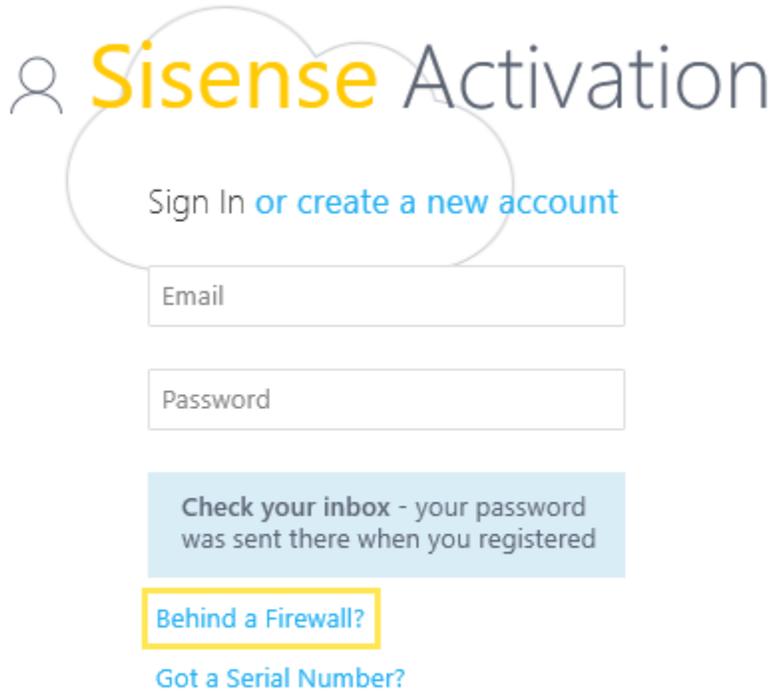
3. In the Welcome to Sisense window, click **Activate**.

# Welcome to **Sisense**



By clicking the button above you agree to our [Terms and Conditions](#)

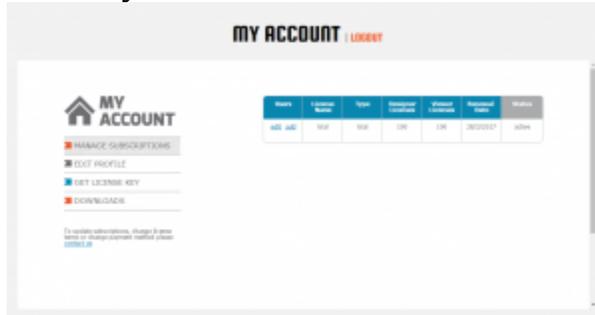
4. In the Sisense Activation window, click **Behind a Firewall**. The product ID is displayed.



5. Copy the product ID.



6. On an online machine, open <http://www.sisense.com/my-account/> and log in with your Sisense credentials. The My Account page is displayed.



7. In the My Account page, from the left menu, select **GET KEY**.



8. Paste your product ID from the Sisense installation into the field and click **GET KEY**. Your Off License key is displayed.

#### GET OFFLINE LICENSE KEY

SiSense Prism occasionally validates its license status via an Internet connection. For situations where an Internet connection is unavailable (or blocked by a firewall), an offline license key can be generated from this page instead.

```
nMaUr4KLZGTXLwTwleoqYD0uqjFUQZ3Mpvt2DTFCxd2eF4EWlafg7fDT85hI2
Nx47we6lk3XR94RGBXs5V9k7FgkNs9TPITiJQgPqkhET5tAmh+QJR9WCU52
1Hcz                                     300DC
St8e                                     75Lng
JjMd                                     i9gfal0
/gsqc                                    OEgjS
cwb6                                     rYOIG
UoTI                                     xlcBcV
Wzyl                                     jdejbD
maE                                       iSzgY
Ve/9                                     SbDyp
515h                                     Vfmjq3
FDh`                                     JWmB
oJyr                                     jz2AB
sE5+                                     Ln+Bk
uub6                                     iarYhr
XrCv                                     !ldwrM
CaXl                                     1qJ5u
zQWN69hvr0IDTbQwx7vgh9r+niK+1jU+P6fN38s9Q6hYsdVTILyIkWIFTDblJR4/
1o99zfnPYwxBzEAgwZ+elwZ2mxcHmuMjLTXfhAEIjXymegYZhC8b3JN1BWdj/
7L
```



9. Copy the key and paste it into the key area in the Sisense Activation window.



10. Click **Install**. Sisense is activated and you can begin to install Sisense.  
11. Click **Custom Installation**.



12. Select **Microsoft IIS** and click **Continue**.



The installation will start running and complete.

# Introduction

## Audience and Scope

This user guide is intended for Sisense customers that use the Sisense Web Application to view and analyze dashboards. It provides an in-depth description of the interface, functionality, and options available in the Sisense Web Application.

## Overview of the Sisense Web Application

The Sisense Web Application is an interactive web application that provides the user interface in which you can view and explore dashboards. The Sisense Web Application runs in popular web browsers and enables access to the Sisense server.

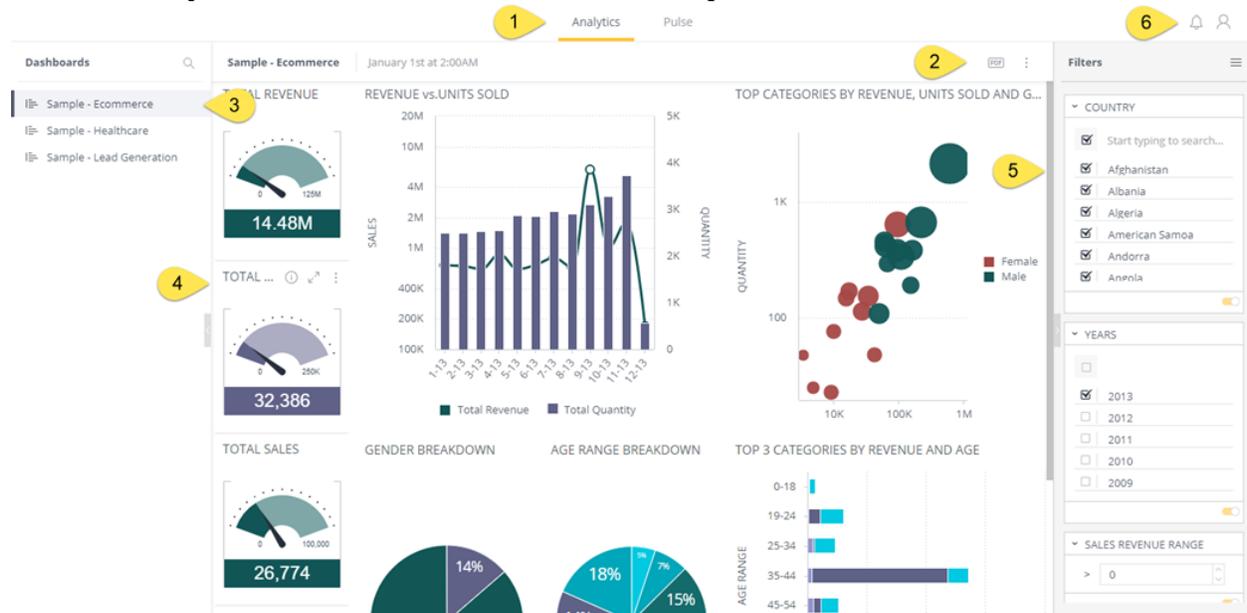
## Requirements

The Sisense Web Application runs in the following HTML5 supported browsers:

- Internet Explorer 10 and higher
- Google Chrome
- Firefox
- Safari version 7 and higher

# A Tour of the Sisense Web Application Interface

The Sisense Web Application screens contain a variety of information and functionality that relate to dashboards. A sample screen is shown below:



1. **Navigation Bar:** The navigation bar contains links to the Analytics page where your dashboards are displayed and the Pulse page where any widgets you have added to Pulse are displayed.
2. **Menu Bar:** The menu bar displays your dashboard's name, the time it was last updated, and menu options for downloading or deleting the dashboard.
3. **Dashboard List:** A list of your dashboards.
4. **Dashboard:** A collection of one or more widgets that visualize your data.
5. **Filters Pane:** The Filters pane displays dashboard filters you can toggle to change the way values are presented in a visualization.
6. **User Profile:** This section displays your Pulse feed where incoming notifications are listed and your user information. This information includes your user name and the Sisense Web Application version number.

# Setting Up Your Production Environment

This topic describes how to set up a production environment for hosting your Sisense implementation.

Sisense recommends installing Sisense in a staging environment, where you can check all of your ElastiCubes and dashboards, as well as test different configurations and customization options. When everything is working as expected, deploy Sisense in your production environment. To learn more about migrating your installation, see [Migrating Sisense](#).

## Providing Remote Access to the Sisense Web Environment

To maximize the Sisense experience, you will want to make your dashboards accessible to external users outside the organization's network. This requires the following configurations:

- Open TCP port (default is 8081) to access the Sisense Web Application internally.
- Ensure that your Sisense machine has a public IP address or domain name associated with a public IP for external access.
- Open port 443 if you are using SSL. Sisense recommends using SSL for secure connections or a VPN solution. For more information about implementing SSL, see [Setting Up SSL](#).

External users who were granted access can view dashboards by entering the machine's IP and the port number, for example: xx.xxx.xx.xx:8081 or my.company.com:8081

Click [here](#) to learn more and see troubleshooting tips.

# Providing Remote Access to the ElastiCube Manager

In some cases, you may want to provide external access to the ElastiCube Manager for designing and modeling your ElastiCubes. In such cases it is necessary to consult your Sisense Success Manager to make sure your licensing terms allow this option.

In addition, make sure that ports 811, 812 on your Sisense machine are open.

**Note:** To avoid exposing your data, do not open ports 811/812 without the proper security measures.

## Base URLs

Base URLs are the consistent part of a web address for a site or web application. In Sisense, the default address is localhost:8081. Sisense allows you to specify a Sisense hostname as a subfolder of a domain, for example, baseurl.sisense.com/reporting.

For more information, see [Configuring Base URLs](#).

# Performing a Silent Installation

This topic describes how to install Sisense silently.

You can install Sisense using the command-line interpreter, for example, when you need to run an installation on a remote machine.

You can also customize the installation using variables, for example, install Sisense without the sample dashboards.

## To perform a silent installation:

1. Download the [latest sisense version](#).
2. Open the command line interpreter and type in the installation commands as follows:

- Full installation (without previous user activation)

```
<Location of downloaded executable file>\SisenseLatestFull.exe -q -username=<> -password=<>
```

- Full installation (for user already activated by Sisense)

```
<Location of downloaded executable file>\SisenseLatestFull.exe -q
```

- Full offline installation (requires a license key from Sisense)

**Note:** To get an offline license, go to [My Account](#) and click **GET KEY**, or contact your Sisense Success Manager.

After attaining your key, copy and paste it after *-offlinelicense=* (replacing *<offline\_license\_key>* in the example below).

```
<Location of downloaded executable file>\SisenseLatestFull.exe -q -offlinelicense=<offline_license_key>
```

- Installation without Sample ElastiCubes and dashboards

```
<Location of downloaded executable file>\SisenseLatestFull.exe -nosamples -q
```

- Choosing Sisense Website name and port

```
<Location of downloaded executable
```

```
file>\SisenseLatestFull.exe -q -  
webname=<website_name> -webport=<website_port>
```

## Silent Customized Installations

To define which Sisense components are installed, you can edit the PrismFeature.xml and copy it to “%programData%\sisense”.

If you are using IIS express, download [PrismFeature.xml](#).

If you are using IIS, download [PrismFeature.xml](#).

**Note:** Create “%programData%\sisense” if it does not exist.

The PrismFeature.xml file contains an XML node named “Feature”. Changing its ‘Selected’ child node value (true/false) determines if the feature is installed or not. For example, in this configuration, the Prism Server Console is installed:

```
<Features>  
<Id>PrismServerConsole_Feature</Id>  
  <Name>Prism Server Console</Name>  
  <Selected>>true</Selected>  
  <Manual>>false</Manual>  
  <Selector>CheckBox</Selector>  
  <Tests />  
</Feature>  
</Features>
```

After editing the XML file, run the installation according to the instructions above.

### To uninstall Sisense using the command-line interpreter:

In the command-line interpreter, enter:

```
<Location of downloaded executable  
file>\SisenseLatestFull.exe /q /uninstall
```

# Tutorial: Getting Started with Sisense

## Welcome!

In this basic tutorial, you will learn how to connect to data, and how to build your first dashboard.

The demo data in this tutorial is based on sample E-commerce data, and a common market analysis scenario will be used.

If you have already connected to data sources, you can skip this step and go directly to Step 2 – Your First Dashboard.

### *Step 1 Connect to Data*

## Introduction

To work with data in Sisense you need ElastiCubes. An ElastiCube is Sisense's proprietary analytical database, which enables you to connect multiple data sources and run complex queries in split seconds.

You can connect to databases (like SQL Server, MySQL, Oracle, etc.), files (text, csv, Microsoft Access, etc.) and online web services (Google AdWords/Analytics, Salesforce, Zendesk, etc.).

In this tutorial you will create your first ElastiCube using some sample ECommerce data.

Please download the following two sample files:

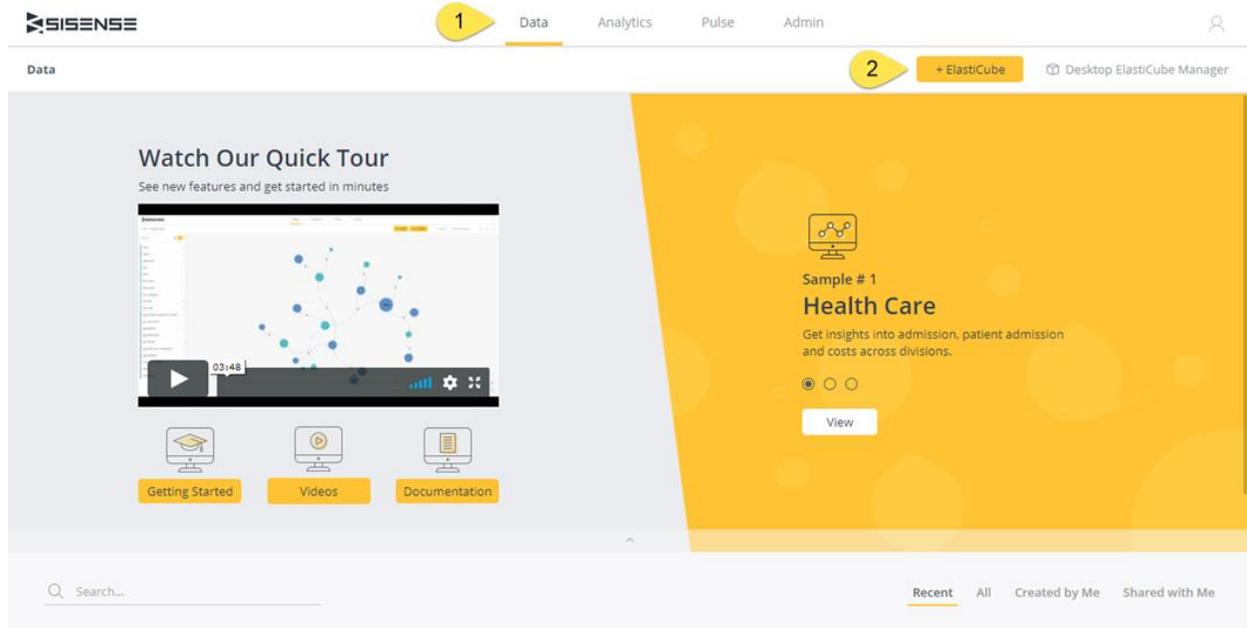
1. [GettingStarted ECommerce.csv](#) – A CSV file with a few hundred thousand entries.

2. [GettingStarted Brands.xlsx](#) – An Excel file with data on brands sold in our ECommerce sample.

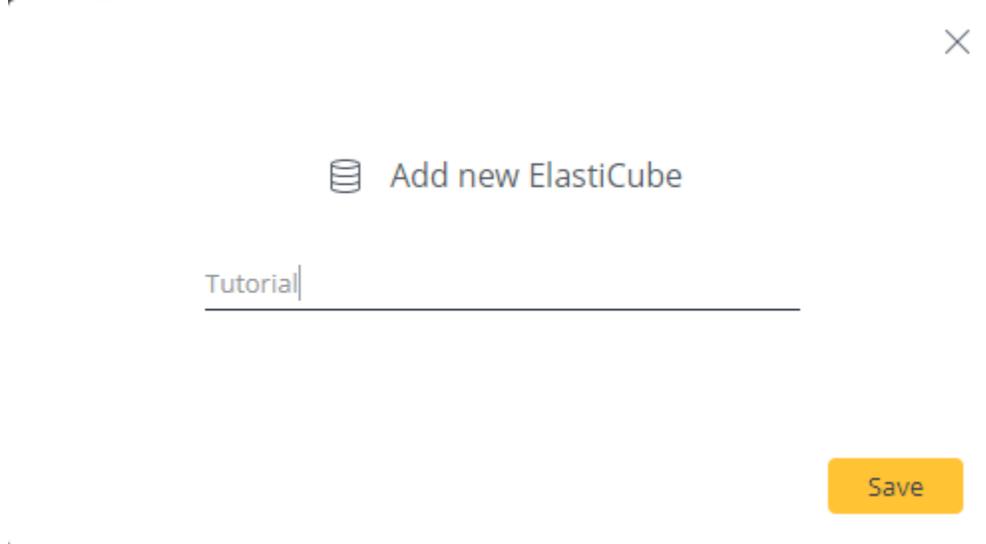
# 1. Open the ElastiCube Manager

Look for **Sisense** in your Windows Start Menu and open it.

From the menu bar, select **Data > + ElastiCube**.



Now, give your ElastiCube the name “Tutorial”:



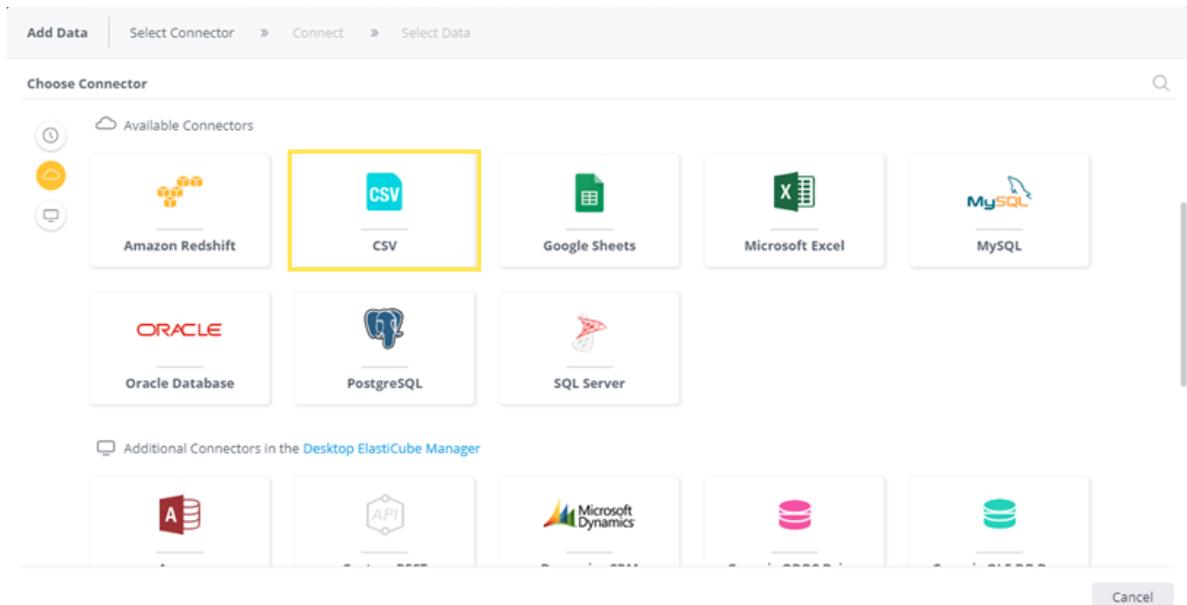
You just created an empty ElastiCube. The next step is to add some data to it.

## 2. Add Data Sources

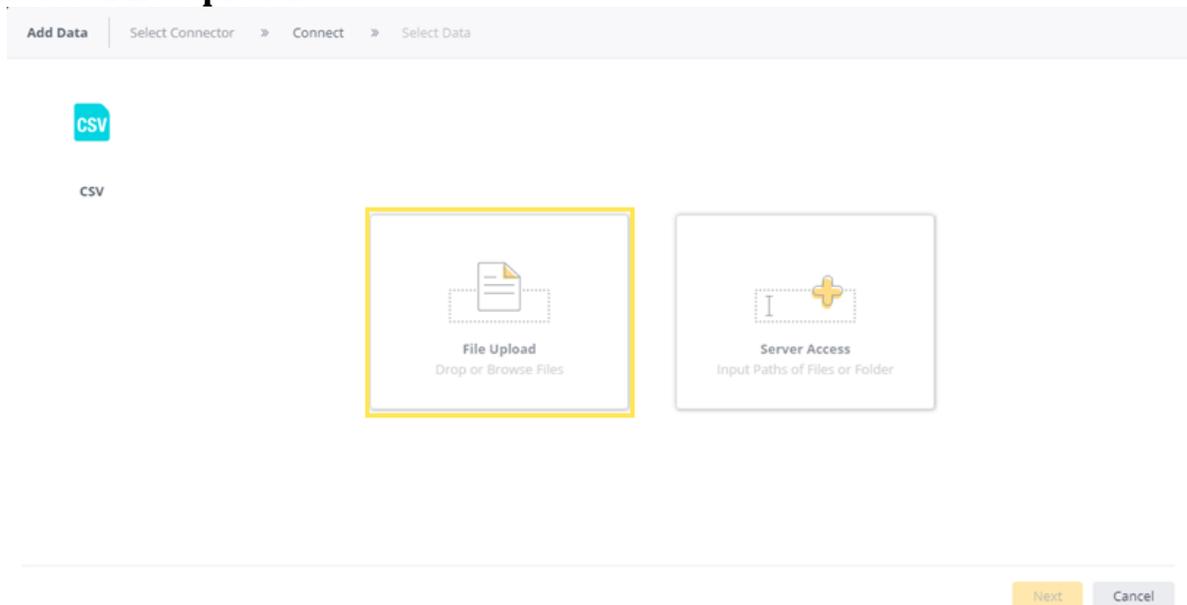
### Table #1 – CSV File

Import the base table, which is in the form of a CSV file. This table contains info about what kind of products have been purchased, and how much they cost.

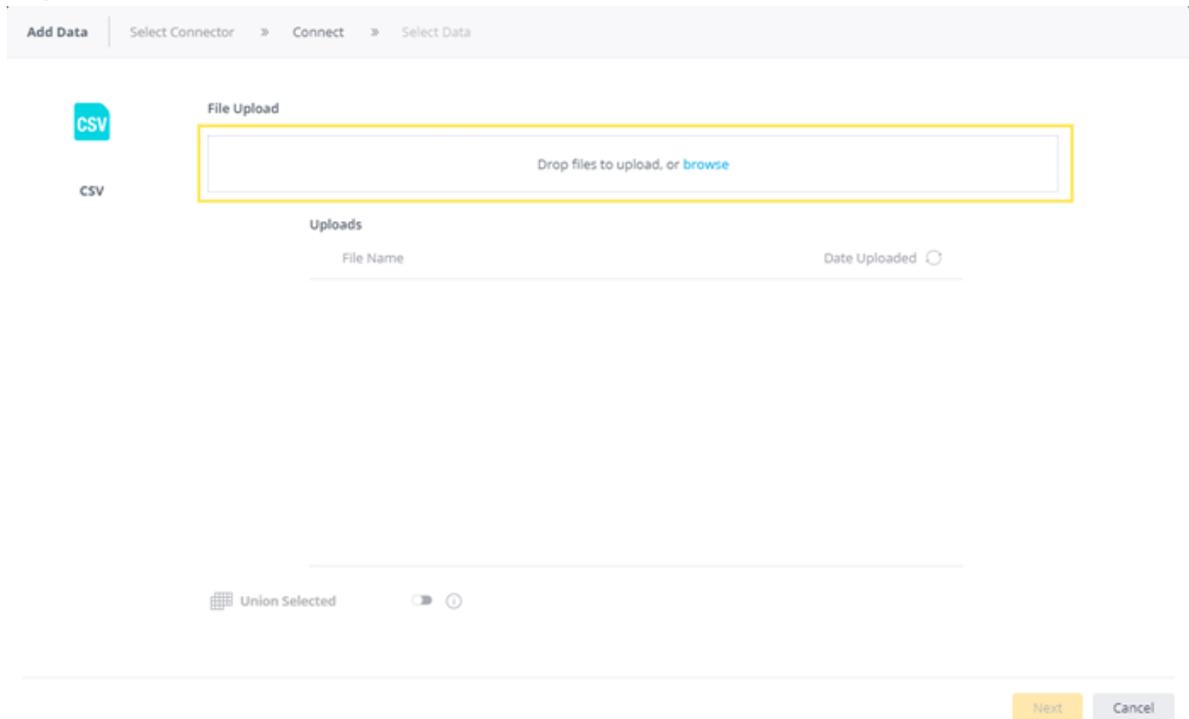
1. Click **+ Data** and select CSV from the list:



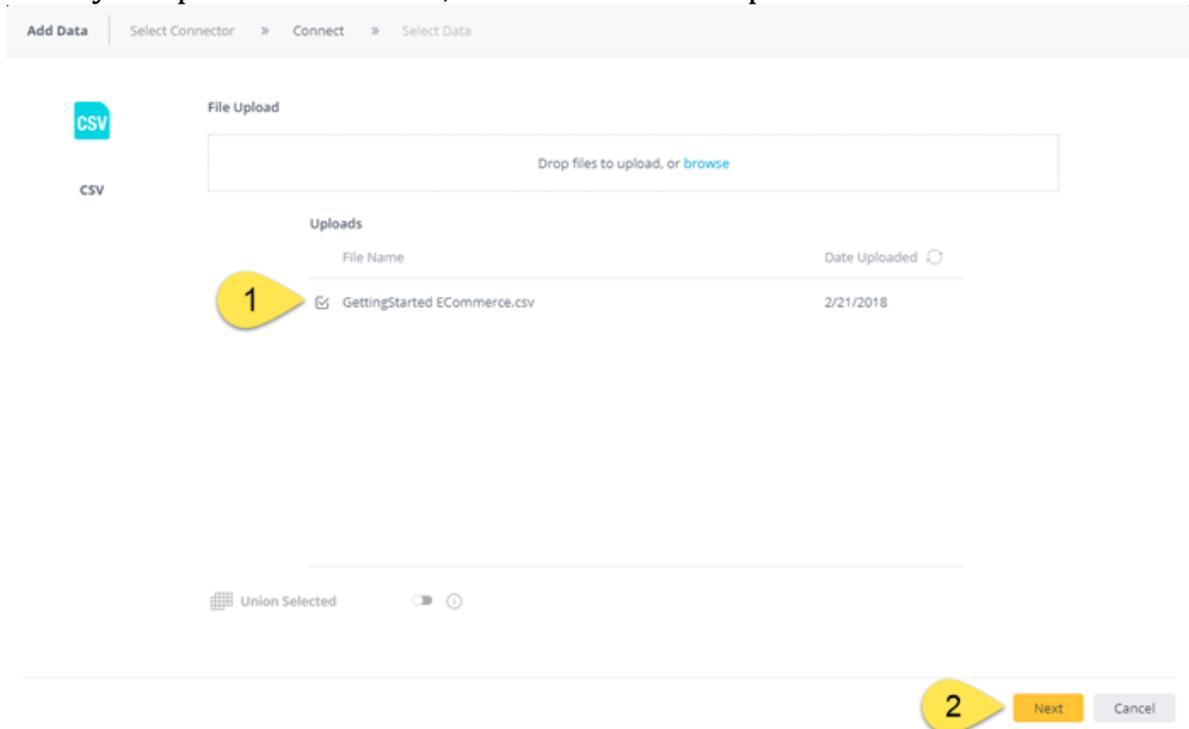
2. You will be prompted to choose how you want to upload your CSV file. Select **File Upload**.



3. Drag the GettingStarted.CSV file to the Upload area or click browse and navigate to the file.



4. After you upload the CSV file, select it from the Uploads list and click **Next**:



5. Select the GettingStarted Ecommerce file from the Select Table list. You can preview the data inside it by clicking . This displays some more settings

you can use to customize your data, but for now, click **Done**.

The screenshot shows the 'Add Data' interface in ElastiCube. At the top, there are navigation steps: 'Add Data', 'Select Connector', 'Connect', and 'Select Data'. Below this is a 'Select Table' section with a search icon and a table of data. The table has columns for Date, Country, Category, Age Range, Brand ID, and Cost. The first row is selected. Below the table is a 'Settings' panel with various options for data import, such as Culture, Text Qualifier, First Row Contains Field Names, Delimit..., Ignore Rows that Start..., Ignore Rows that Contain, and Ignore First Rows. At the bottom right, there are 'Done' and 'Cancel' buttons.

Select Table	Date	Country	Category	Age Range	Brand ID	Cost
<input checked="" type="checkbox"/> GettingStarted ECommerc...	6/7/2010, 3:00...	United States	Digital Cameras	65+		3
	6/18/2010, 3:0...	Germany	Cell Phones	25-34		3
	6/19/2010, 3:0...	United States	Printers	35-44		5
	7/9/2010, 3:00...	United States	Digital Cameras	65+		3
	7/19/2010, 3:0...	United States	Digital Cameras	19-24		5
	7/19/2010, 3:0...	United States	Apple Mac Desktops	25-34		4
	5/11/2010, 3:0...	Canada	Digital Cameras	45-54		3
	5/24/2010, 3:0...	United States	Printers	65+		5
	6/5/2010, 3:00...	United States	Printers	35-44		5
	6/15/2010, 3:0...	Latvia	Monitors	35-44		1,400
	6/26/2010, 3:0...	Spain	Camera Flashes	19-24		100

Settings

Culture English (United States) v

Text Qualifier

First Row Contains Field Names

Delimit... Comma (,)

Ignore Rows that Start ...

Ignore Rows that Contain

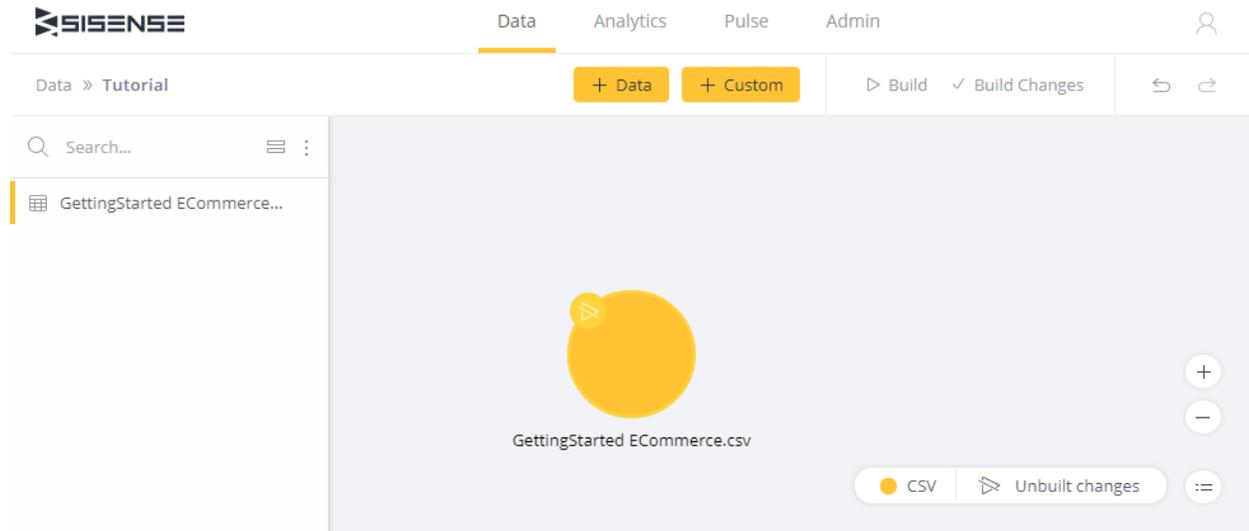
Ignore First Rows

Done Cancel

Congratulations, you have connected to your first data source. You can now see the CSV file in your schema.

The name of the table in the ElastiCube defaults to the file name in this case. You can double-click the title to rename it and remove the CSV extension.

Your ElastiCube should look like this:



## Table #2 – Excel File

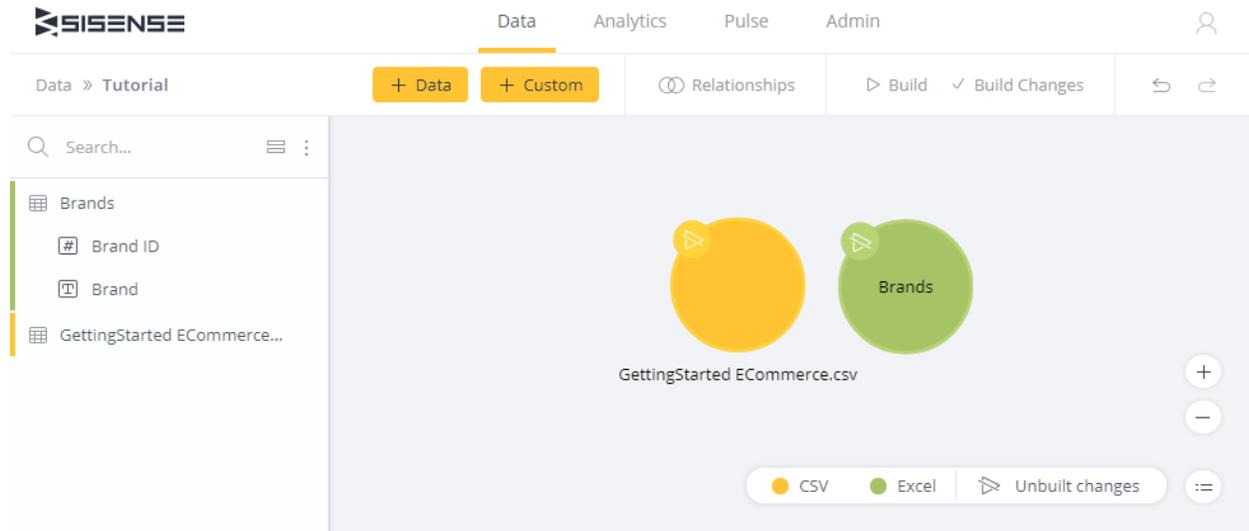
Information about the brands that were sold in the main table exists in another table, this time in an Excel file.

1. Click **Add Data** and select **Microsoft Excel File** from the list.
2. You will be prompted to select an Excel file. Find and select “GettingStarted Brands.xlsx”, which you downloaded in Step 1 (you can also download the file [here](#)).
3. In the new window leave the default settings and click **Add**.

Congratulations, you have now added your second table, an Excel file.

To rename the table, click the table once and select and rename the table to “Brands”.

Your ElastiCube should look like this:



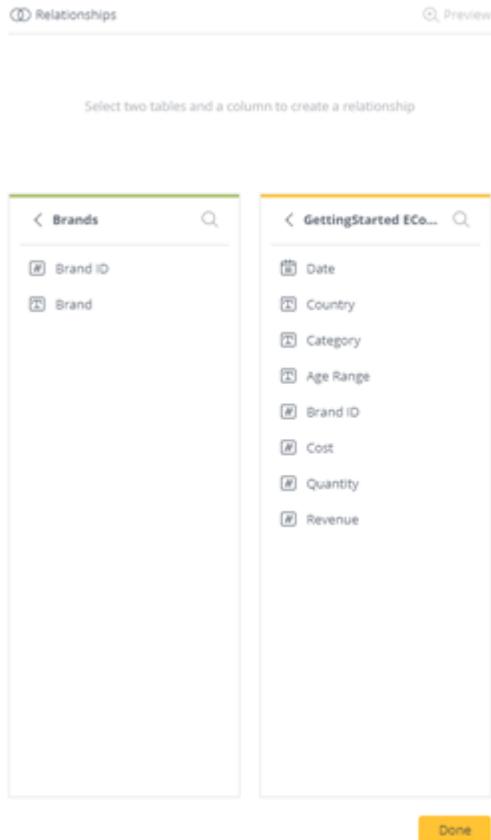
### 3. Connect Different Data Sources

So far you have added two tables from two different data sources. Now let's see how to connect them.

In the ElastiCube, creating and deleting relationships between tables is as simple as drag and drop.

Any two fields of the same type (numeric/text/date) can be connected. In this sample, both tables have a "Brand ID" column. Let's connect them:

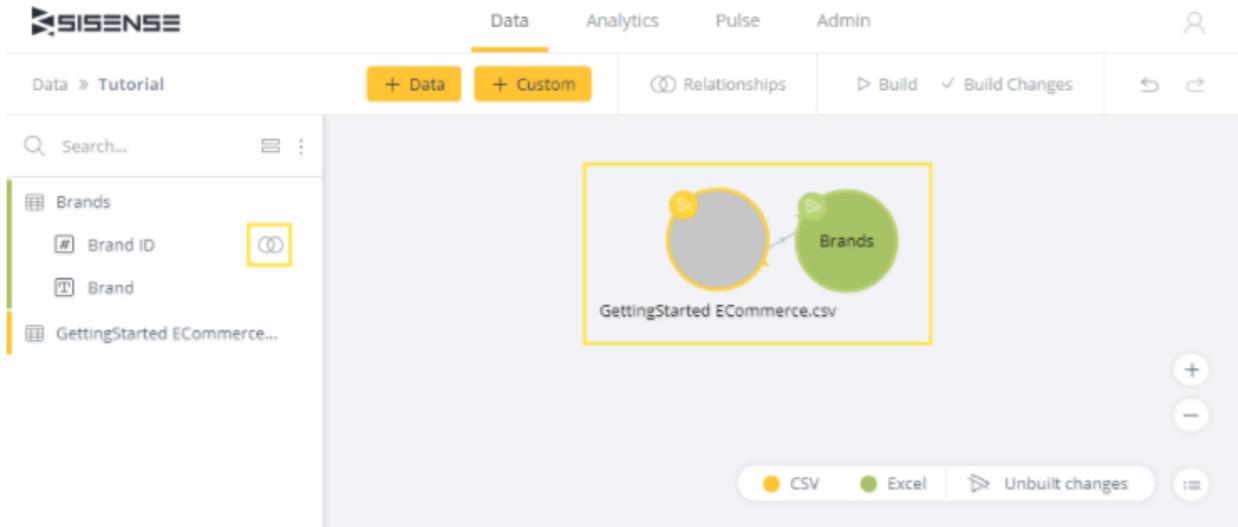
1. Drag the Brands table onto the Ecommerce table. This opens the Relationships pane.



2. Select **Brand ID** from both tables and click **Done**. You should now see the two tables are joined through the column Brand ID.

## 4. Build the ElastiCube

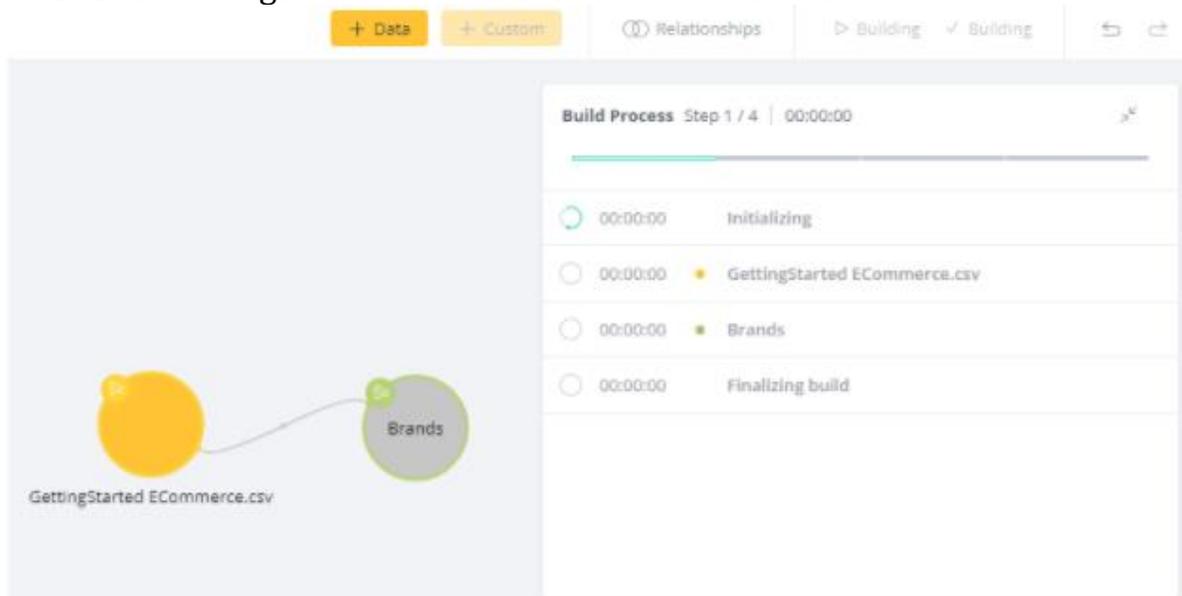
Now that you have defined your ElastiCube, you are ready to build it. This will pull the data from the data sources into the ElastiCube.



1. Click Build in the top menu.



2. The build will start. You will see the progress at the bottom of the ElastiCube Manager window. Wait for the build to finish:



Congratulations! You have successfully built your first ElastiCube.

## 5. Dashboard Time!

Click Analytics in the top menu to open the Sisense Web Application and create a new dashboard.

Data **Analytics** Pulse Admin

### *Step 2 – Your first Dashboard*

Welcome to Step 2 of the Sisense tutorial. In Step 1 – Connect to Data of this tutorial, you connected to data and built your first ElastiCube. Now, you are ready to jump right into creating your first dashboard.

Although it is recommended to complete Step 1 in advance, you can jump right into Step 2 if you wish.

If you're arriving directly from Step 1, you should already be in the process of creating a new dashboard after clicking Analytics in the top menu of Sisense:



## Introduction

The Sisense Web application is designed to let business users easily create powerful and meaningful dashboards, without being a developer or data scientist.

To get a feel for the process, Sisense has put together this tutorial which shows you how to do a quick market analysis to identify what products are meeting customers' needs, and in which markets. You will use appropriate visualizations to create a dashboard that provides answers to four questions:

1. **Sales Trend:** Is annual revenue increasing?
2. **Optimal Product Mix:** What product lines should we grow?
3. **Segmentation:** Which customer segments should we develop?
4. **Market Size and Growth:** Where are our best markets located?

To make things interesting, your data set is based on sales and customer data from an E-commerce site.

## 1 – Create a New Dashboard

1. Click **+** to create a new dashboard on the Analytics page.



2. In the **Data Set** field, select the ElastiCube to which you want to connect. If you've completed [Step 1 – Connect to Data](#), select 'Tutorial' in the Data Set field, otherwise choose 'Sample ECommerce'.
3. A default title name matching the ElastiCube name will be given to the dashboard. Feel free to change it.

New Dashboard / Please select a Data Set to connect to

Data Set: Tutorial

Title: Tutorial

Create

Cancel

4. Click **Create**.  
Now that you have a dashboard, you can start adding widgets to it.

## 2 – Create Your First Widget: Sales Trends

Create your first widget to get a general idea of your sales performance.

1. Click Select Data on the left side of the screen:

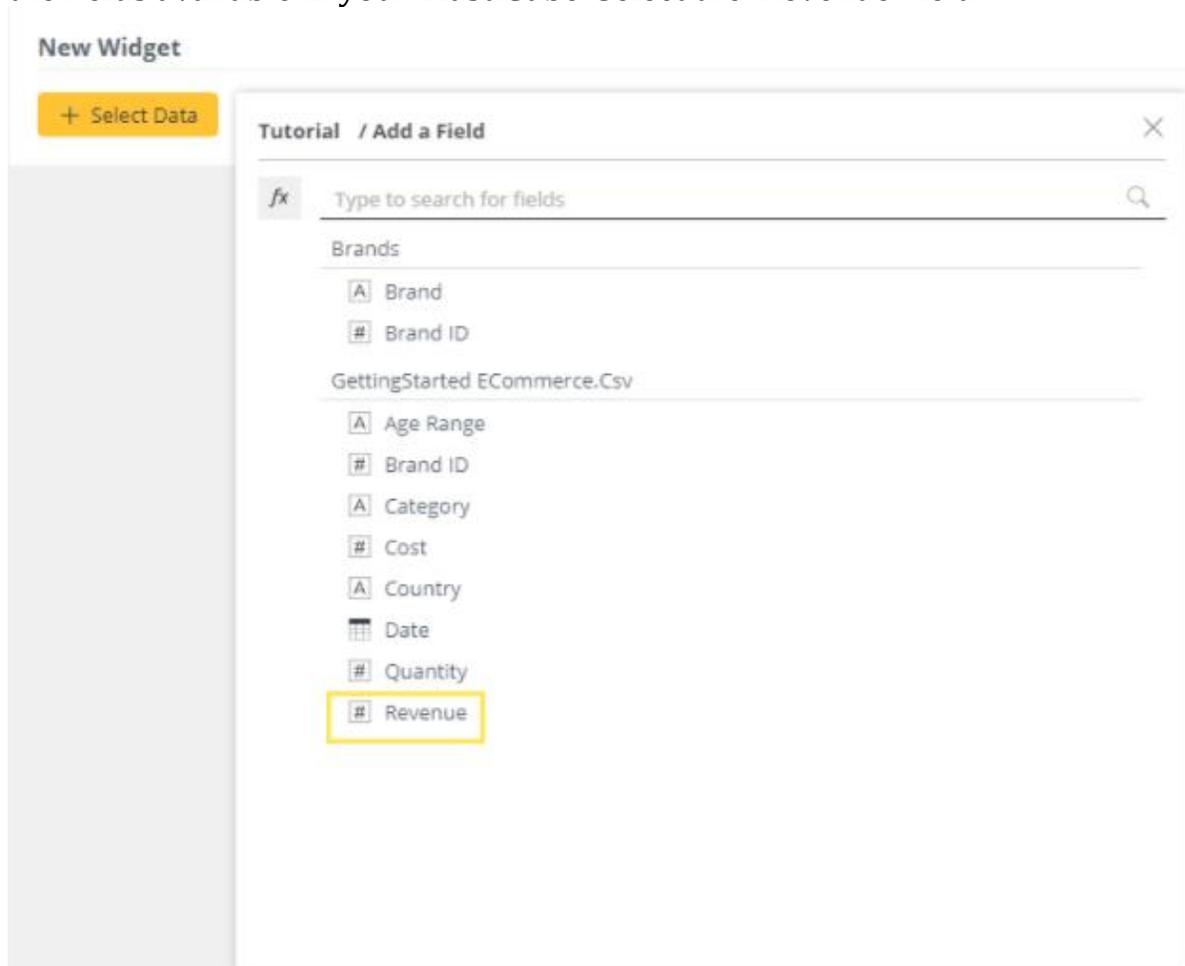
Welcome to your new dashboard!

To create a new widget, first select the data you would like to visualize.

## New Widget

+ Select Data

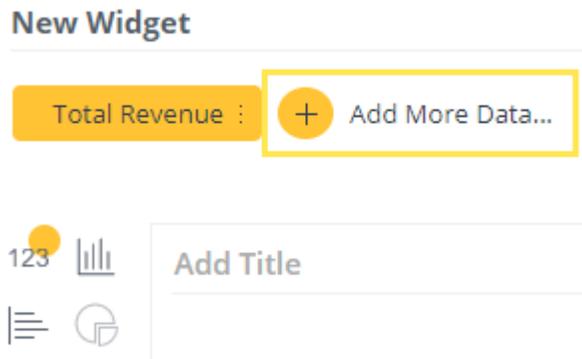
2. A list of available fields will be displayed grouped according to the table they belong to. This window is called 'Data Browser', and lets you easily find the fields available in your ElastiCube. Select the 'Revenue' field:



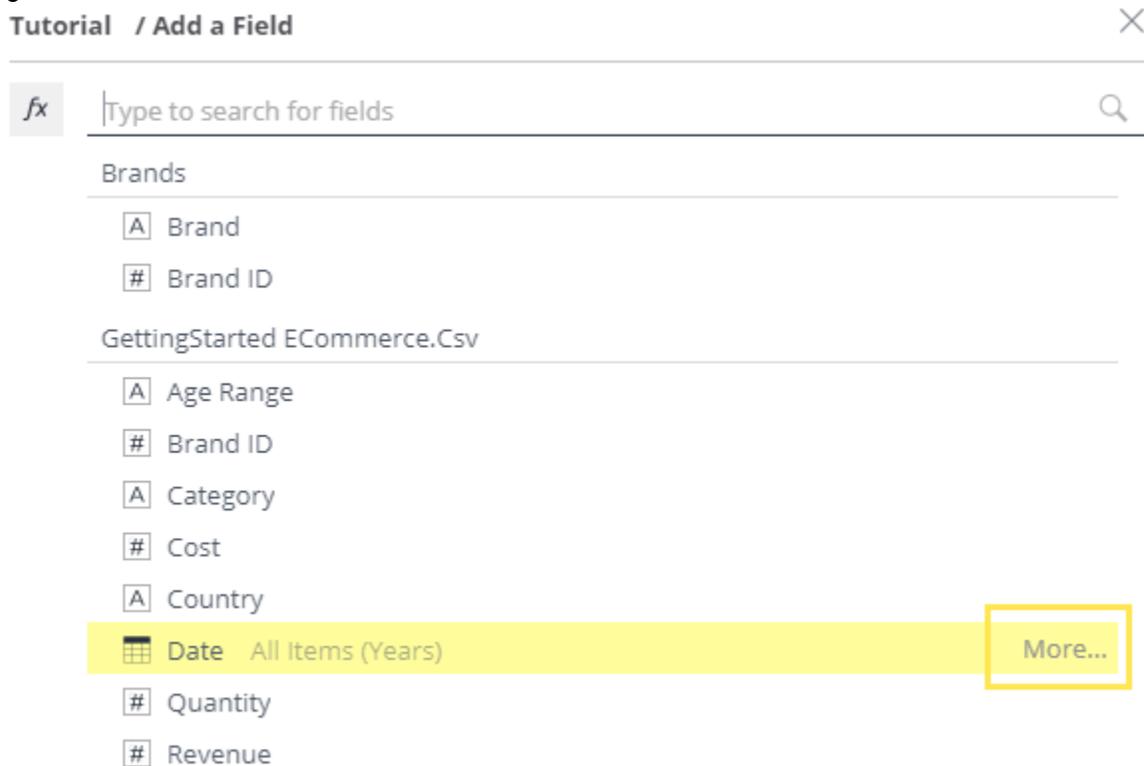
Now you immediately get a ready made widget that shows you all revenue to date. We call this kind of widget an "Indicator".

But to understand revenue over time, add a time field to the mix:

1. Click **+ Add More Data:**



2. In the data browser, hover over the **Date** field and click **More...** Then select **Quarters:**

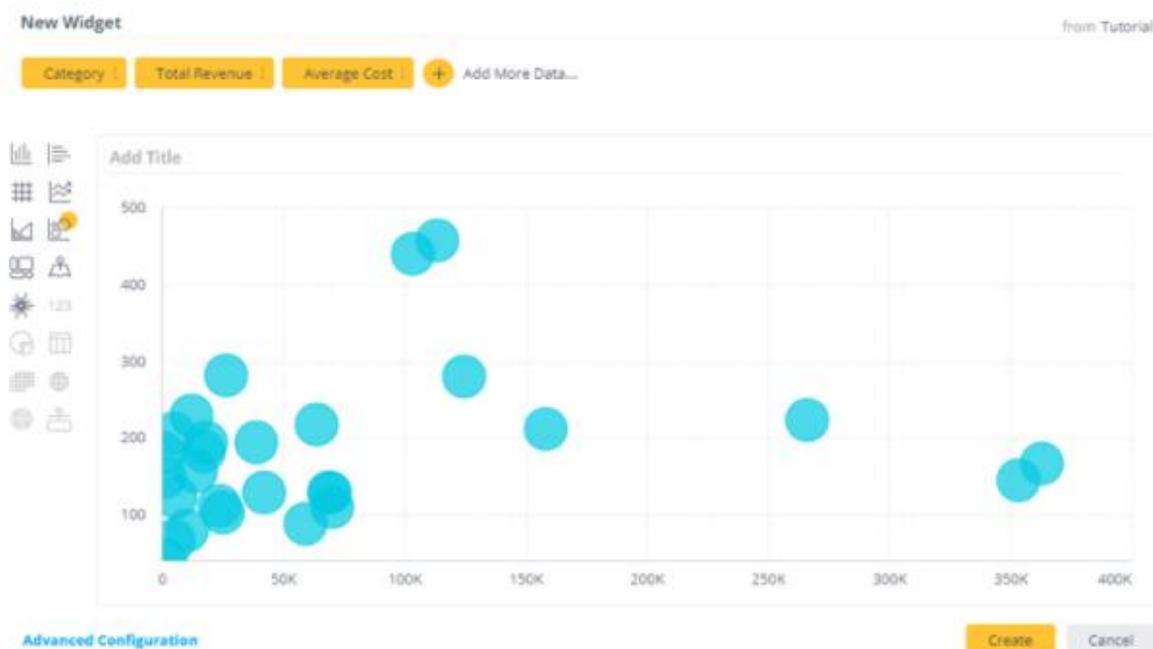


3. Click **Create** to add the line chart widget to the dashboard.  
The revenue trend is also positive showing annual growth, so now go ahead and explore what's driving this growth. It would be interesting to uncover profitable product categories.

## 3 – Create a Scatter Chart to See Cost and Revenue

Now try to understand the interaction of cost and revenue to discover the most profitable product lines.

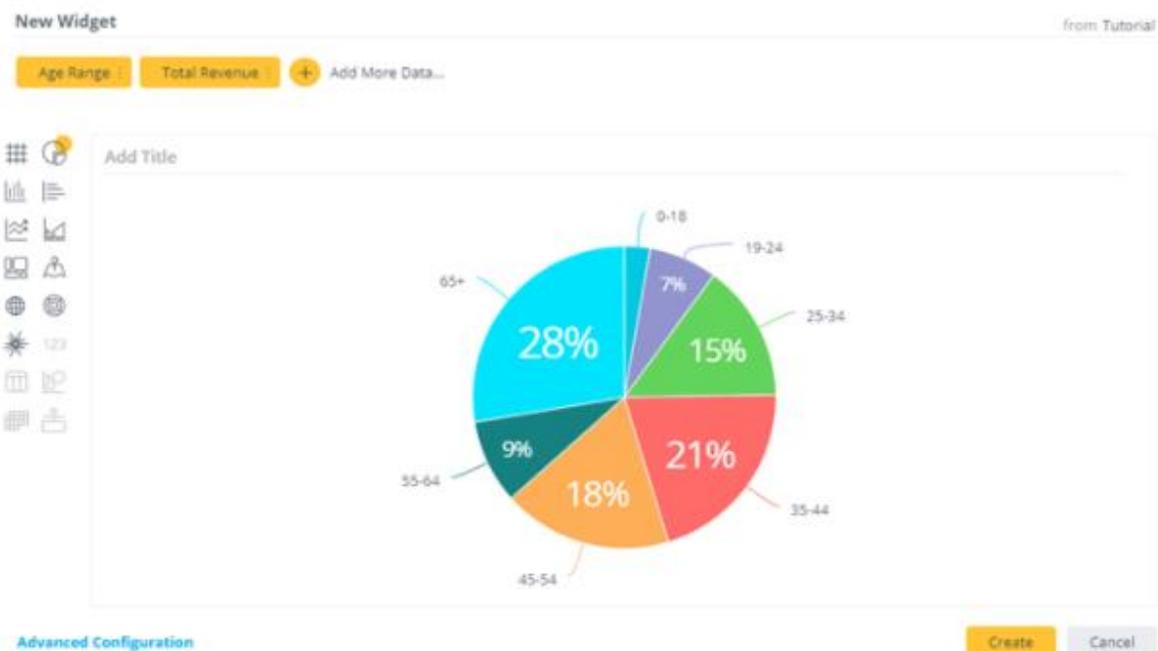
1. On the top of the dashboard click .
2. Click **Select Data** and select the 'Category' field.
3. Click **Add More Data** and select the 'Revenue' field.
4. Click Add More Data again and this time hover over the 'Cost' field and click **More... > Average**.
5. Now select the scatter chart icon  from the widget buttons on the left, and click **Create**.



It's simple to see that product categories on the right generate the most revenue and those products in the lower left have the lowest cost and revenue. Hover over each scatter point to see the category behind it, to discover that Cell Phones and PDAs generate most of the revenue. You now understand which product lines to sell, but to get a complete view of your market, you need to identify optimal customer segments to target in each market – start by creating segments based on age.

## 4 – Creating Pie Charts to Visualize Customer Segmentation

1. On the top of the dashboard, click **New Widget**.
2. Click **Select Data** and select the 'Age Range' field.
3. Click **Add Data** and select the 'Revenue' field.
4. Now select the pie chart icon  from the widget buttons on the left, and click **Create**.



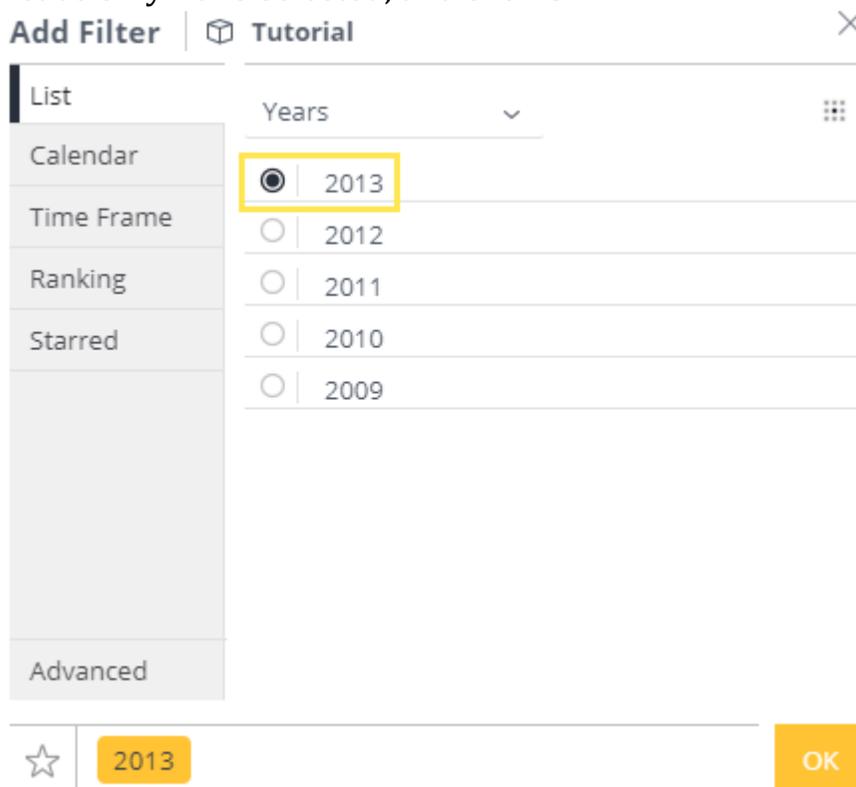
You can clearly see that customers between the ages of 65+ are the most valuable in terms of the revenue they generate. This insight can be used to focus advertising budgets and develop unique campaigns to grow these customer segments. Now, you can turn your attention to identifying which markets are growing at the quickest rate.

## 5 – Creating a Map to See Market Growth

Before you focus on growth, you should focus on a specific year first, by adding a filter.

The **Filter** panel on the right of the dashboards lets you add dynamic filters to the dashboard.

1. Click **Filter Your Dashboard** on the right side of the dashboard.
2. In the Data Browser, click on the **Date** field.
3. Click the multi/single selection button  to make it a single selection filter.
4. Leave only 2013 selected, and click **OK**.

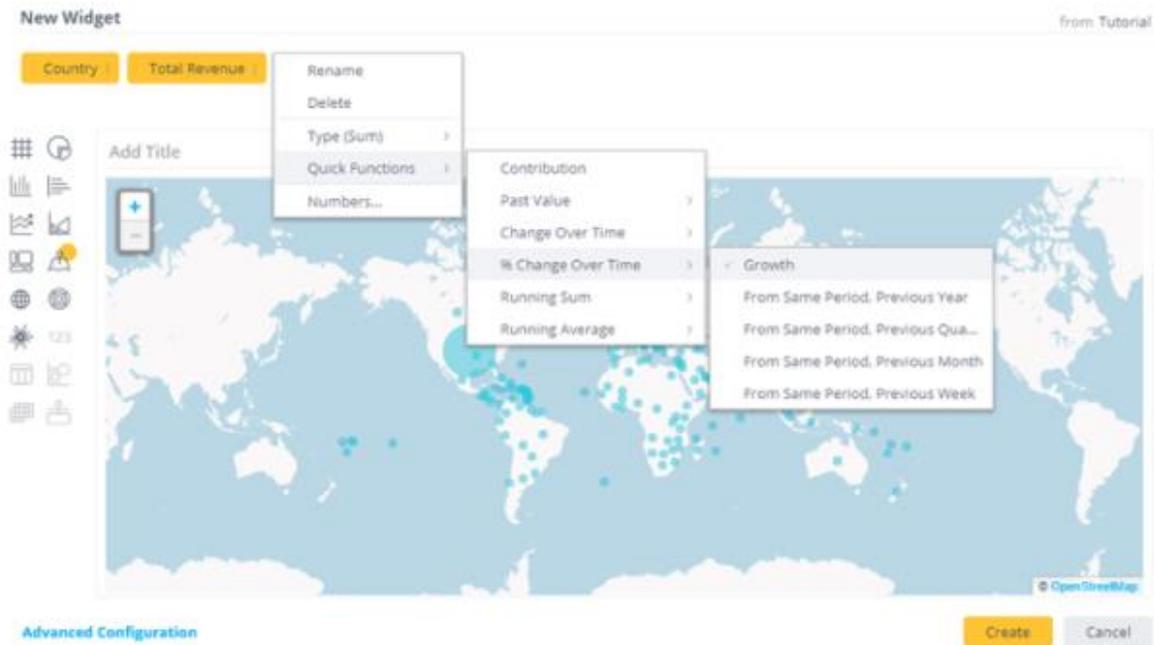


Your dashboard is now filtered to show data only for the year 2013.  
Now create a map:

5. Click  **Widget** at the top of the dashboard.
6. Click **Select Data** and select the **Country** field.
7. Click **Add More Data** again and select the **Revenue** field.
8. Click on the scatter map button type  on the left.

Now you see a map, where the biggest circles represent the countries with the highest revenue. But if you want to concentrate on the markets with the biggest growth in revenue, its easy:

9. Hover and click to open the menu of the Total Revenue field and select **Quick Functions > % Change Over Time > Growth**.



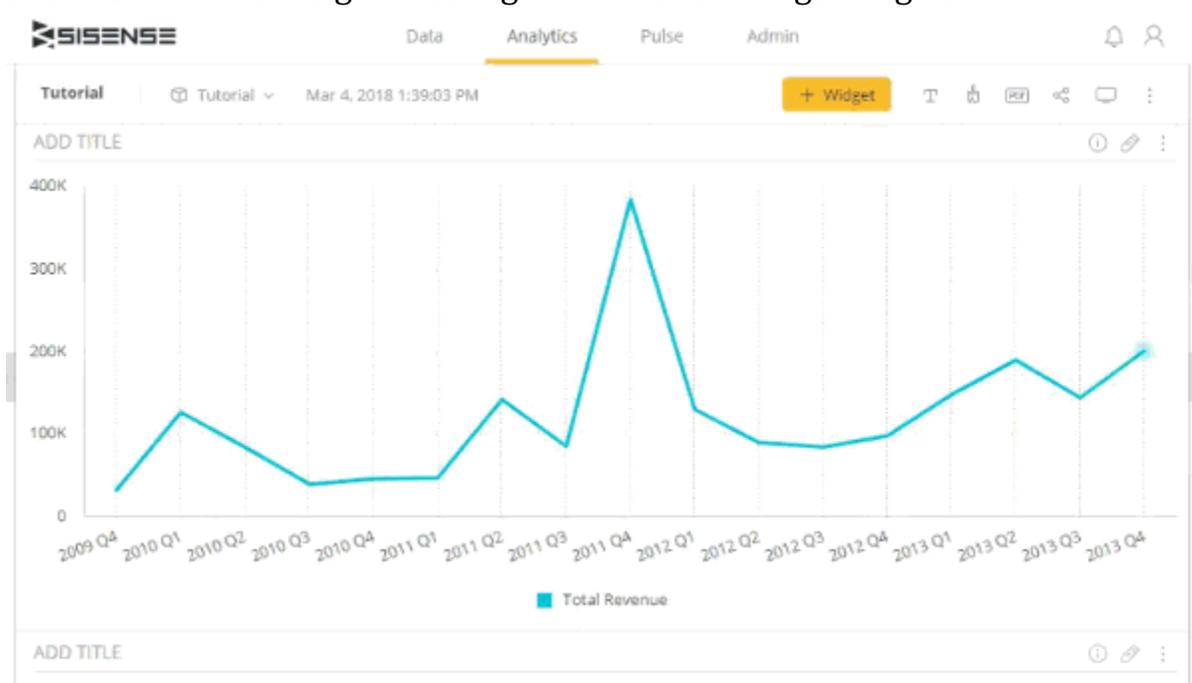
10. Click **Create**.

Now the biggest circles on the map represent countries with highest growth. You can see for example that the United States is very hot right now. You can use this information to focus resources as well as further analyze these markets to understand revenue drivers.

## 6 – Rearrange the Widgets in the Dashboard

You can rearrange the dashboard by dragging and dropping, and resizing the widgets. Dragging one widget on top of another splits the area in the dashboard and allocates space for both widgets.

See this animation to get a feeling of how to rearrange widgets:



Go ahead and rearrange your dashboard however you see fit.

You can now easily see all four visualizations showing sales trends, product mix, customer segments and fast growing markets. You can explore and drill into the details to get further insight.

## 7 – Filtering and Drill Downs

In our pie chart we can clearly see that our strongest age range is 65+. Let's explore this age range a little more:

1. Click the 65+ slice in the pie chart.
2. As a result, a new filter is added to the filter panel on the right, and the entire dashboard is now filtered to focus on this age range.
3. Taking a look at the updated scatter chart, you can see that Monitors are by far the most revenue generating category for this age range.
4. Right-click the 65+ slice in the pie chart and select **Drill Into...** from the menu.
5. A Data Browser will pop up, letting you select any field to drill into. Click on the **Brand** field.

6. You can now see which brands contribute the most to the revenue of your selected age range.  
You can make direct selections by dragging the mouse over the points that interest you in the scatter and line charts as well.

## **8 – So What Did We See Here?**

The dashboard you created enables you to interactively analyze three market priorities – where to sell, what to sell and who to sell to. More importantly it makes it simple to see the interaction of product categories, client segments and revenue, to optimize marketing and sales activities and track performance targets.

This was just a quick glimpse of what you can do. Now it's time to start exploring on your own.

# Upgrading Sisense

This topic describes Sisense's recommended prerequisites for upgrading Sisense and how to execute an upgrade.

Sisense releases several major versions a year, and a few minor versions in between. Each version includes new features and enhancements as well as other improvements.

Sisense allows you to decide when you want to upgrade to the latest version.

There are no automatic updates for the Sisense Server. Before upgrading, you can learn about the contents of each release in the Release Notes.

## Before Your Upgrade

Sisense recommends you follow these guidelines before you begin the upgrade process:

- Install the new version of Sisense on a testing/staging environment and check all of your ElastiCubes and active dashboards. When you are sure that everything is working well, continue with the installation of the new version in your [production environment](#).
- Save all open ElastiCubes and close the ElastiCube Manager.
- Verify that no one is currently designing a dashboard.
- Declare a downtime period and notify users.
- Install the new version with the same user who installed the original Sisense version (with the same administrator privileges).
- If you have a multi-server environment, make sure all of your Sisense servers are running the same version to avoid compatibility issues.

## Upgrading from an Older Version

If you are upgrading your version of Sisense, Sisense recommends you follow these guidelines before you begin the upgrade process:

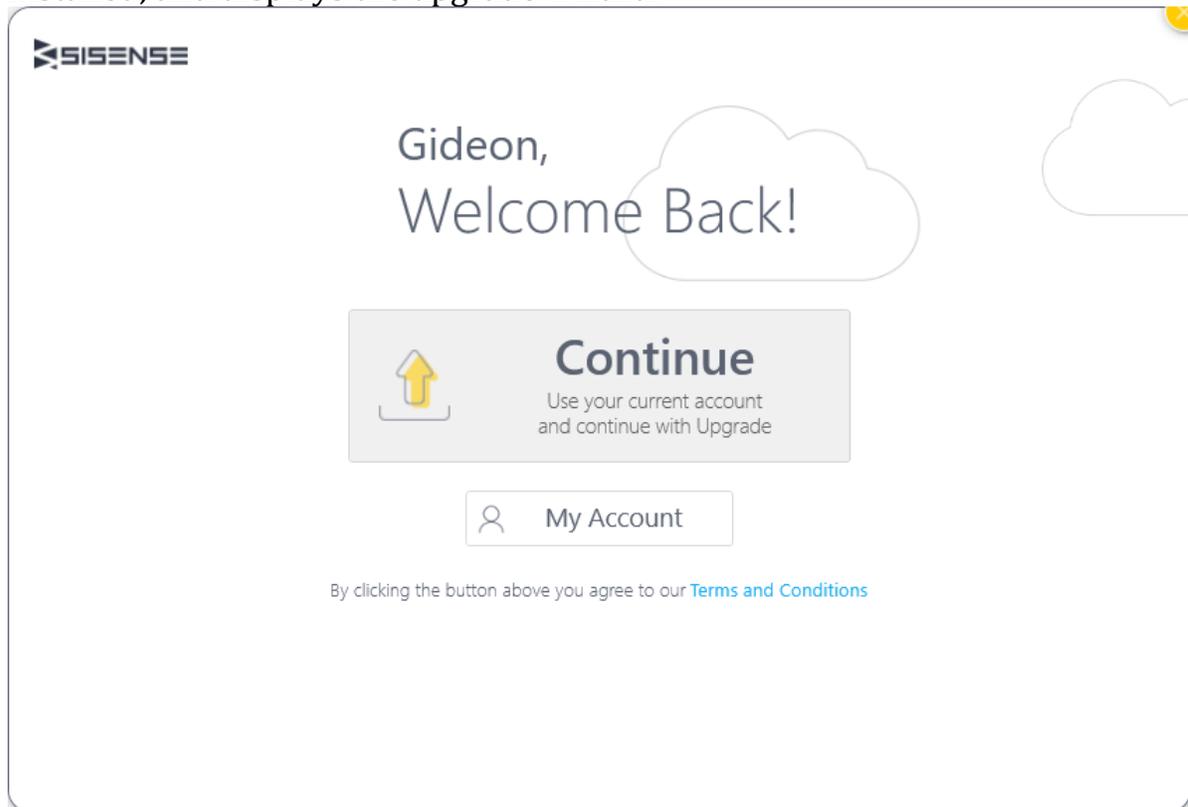
- Back up your ElastiCube data and Sisense Web Application's back-end data on a regular basis, and before upgrading your Sisense version. [Click here to learn more.](#)
- Install Sisense on a testing/staging environment and check all of your ElastiCubes and active dashboards. When you are sure that everything is working well, continue with the installation of the new version in your production environment.
- If you have a multi-server environment, make sure all of your Sisense servers are running the same version to avoid compatibility issues.
- Save all open ElastiCubes and close the ElastiCube Manager.
- Verify that no one is currently designing a dashboard.
- Declare a downtime of approximately 15 minutes and notify users.
- Install the new version with the same user who installed the original Sisense version (with the same administrator privileges).

## Installing the Latest Version

**To install the latest version:**

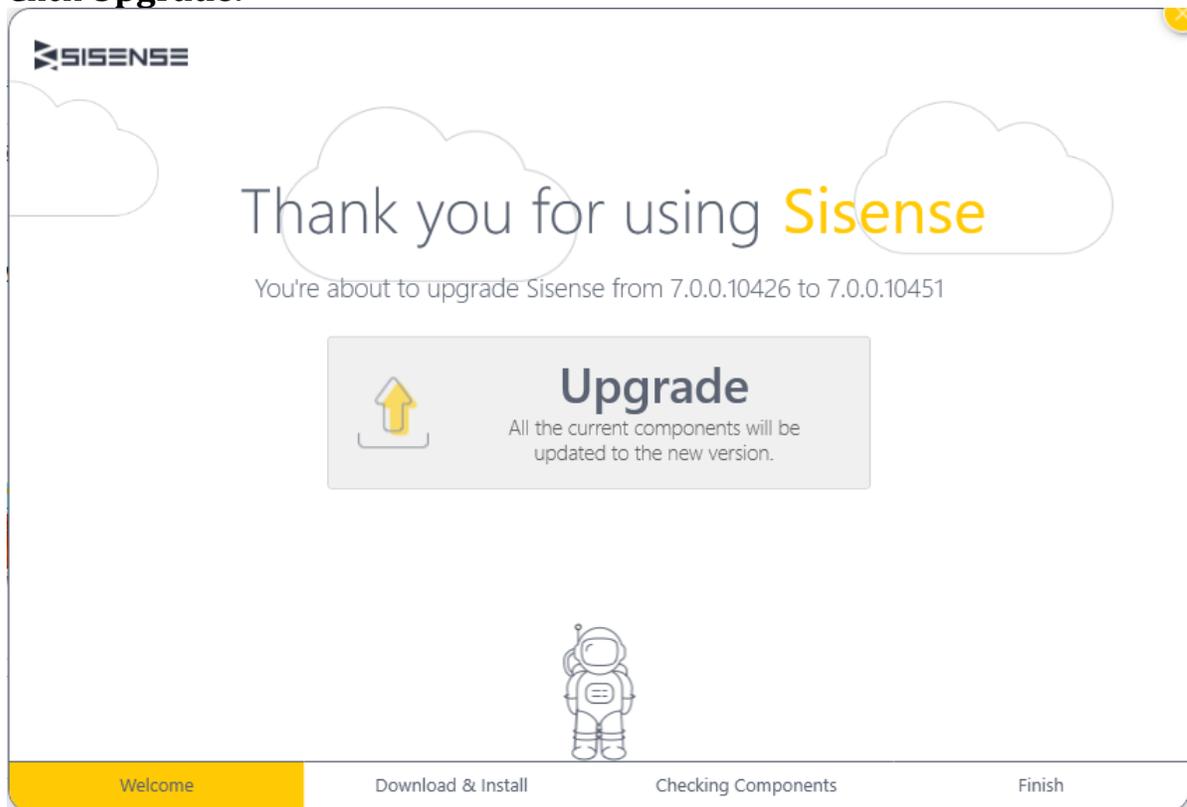
1. Download the latest version of Sisense. For more information, see [Downloading and Installing Sisense](#).

2. Run the installation file. Sisense automatically detects if a version is already installed, and displays the upgrade wizard.



3. Click **Continue**.

#### 4. Click **Upgrade**.



Sisense is upgraded according to the installation settings you selected the last time you installed Sisense. To customize a new installation, you should uninstall Sisense and then install the latest version.

## Downgrading Sisense

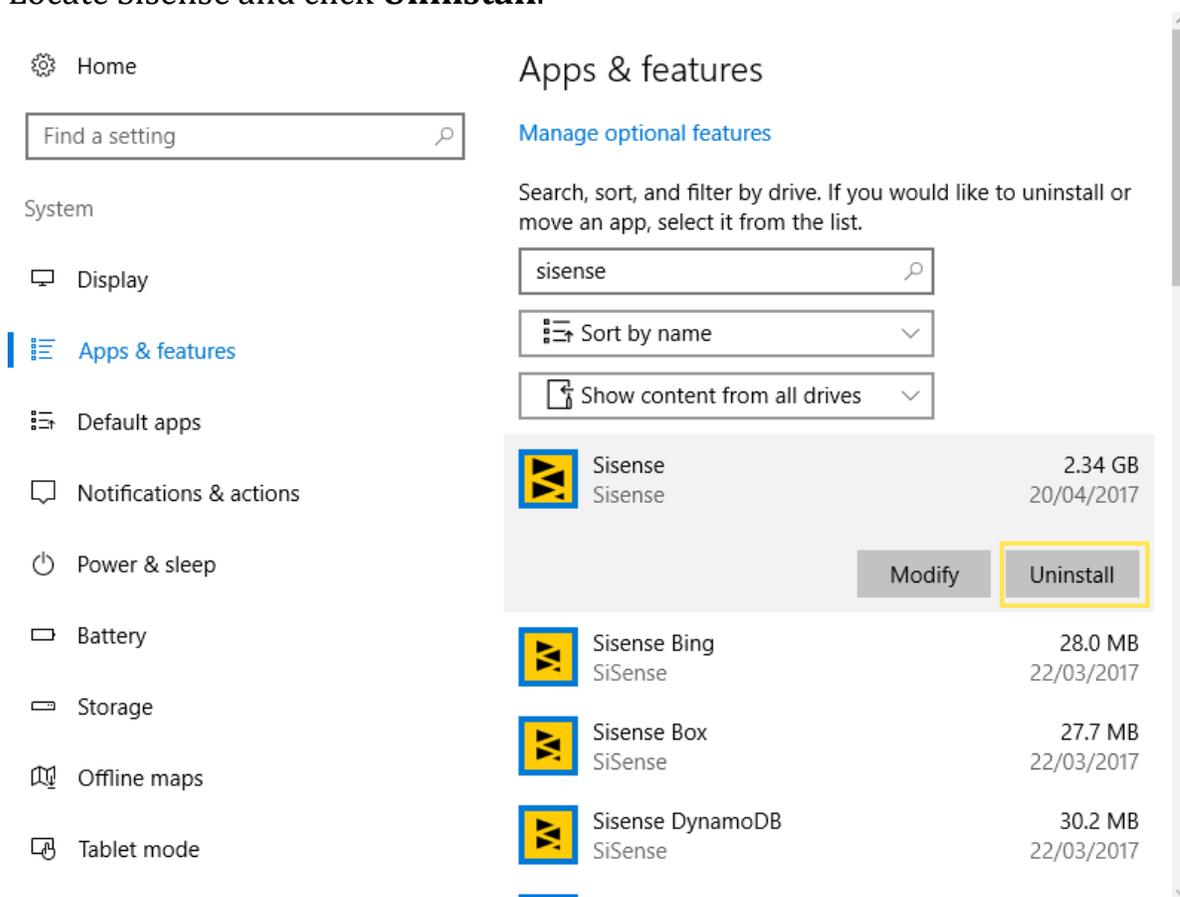
When you upgrade Sisense, a backup copy of your MongoDB instance is automatically created. The MongoDB contains your server settings including your user and dashboard information. This allows you to reinstall earlier versions of Sisense without having to set up your system a second time.

During the uninstall process, you can choose to remove all user data, which deletes the MongoDB instance. The next time you install Sisense, your system will be completely clean, and your dashboard, dashboard list, and user data will be fresh.

Sisense is improving all the time and some versions of Sisense are incompatible with earlier versions. If you upgrade to a version and need to downgrade to an incompatible version, you can restore a backup copy of the MongoDB created automatically when you installed the last version of Sisense. For help downgrading to an earlier version of Sisense that is incompatible with the last version you installed, you can contact Sisense Support who can assist you in restoring a backup copy of your MongoDB instance.

### To downgrade to an earlier version of Sisense:

1. In Windows, goto **Add or Remove Programs**.
2. Locate Sisense and click **Uninstall**.



The Uninstall Wizard is displayed.

3. In the Uninstall Wizard, click **Remove Anyway**. If you select the **Remove User Data** checkbox, Sisense deletes all information stored in the MongoDB and this information must be configured the next time you install Sisense.

# Uninstall Sisense

Last chance to change your mind

Remove User Data

 Remove Anyway

 Keep Sisense



Sisense is uninstalled.

4. Install the previous version of Sisense. The downgrade process is complete.

# Sisense User Roles

**Admin:** Administrators can access the Admin section of the Sisense Web Application from where they can manage ElastiCubes, users, user groups and more. Sisense supports multiple Administrators.

**Data Admin:** Data Administrators can manage all data sources in the system. They have access to the Admin section of the Sisense Web Application where they can manage ElastiCubes and live connections. Data Administrators can also add and manage servers. This is useful for migrating Sisense across environments. Data Administrators do not have access to User Management and System Configuration.

**Data Designer:** Data Designers can create and edit ElastiCubes they created or were shared with them. Data Designers can access the Admin section of the Sisense Web Application where they can manage the ElastiCubes and live Connections they have access to in the Data Sources page. In addition, Data Designers can see servers, but can not add new servers. Data Designers do not have access to User Management and System Configuration.

**Designer:** Designers can create, design, edit and share dashboards. Sisense Designers determine whether the user with who they share a dashboard has editing rights (is a Designer) or only viewing rights (Viewer).

**Viewer:** Viewers can view, explore, drill down, make selections and filter the dashboards that Designers share with them. Viewers only need a standard web browser; no plugins or downloads. They cannot create new dashboards or edit existing ones.

	Admin	Data Admin	Data Designer	Designer	Viewer
System Management					
Access Admin Page	✓	✓	✓		
Edit System Configuration	✓				
Configure SSO	✓				
User Management					
Create/Delete Users	✓				
Create/Delete Groups	✓				
Add Active Directory Users	✓				
ElastiCube Management					
Add/Delete Servers	✓	✓			
View Servers	✓	✓	✓		
Add Live Connections	✓	✓	✓		
Create/Delete ElastiCubes	✓	✓	✓		
Edit ElastiCubes	✓	✓	✓		
Share ElastiCubes	✓	✓	✓		
Edit Servers	✓	✓	✓		
Create ElastiCube Sets	✓	✓			
Delete ElastiCube Sets	✓	✓			
Manage Hierarchies	✓	✓	✓		
Configure Data Access Rights	✓	✓	✓		

	Admin	Data Admin	Data Designer	Designer	Viewer
Stop/Restart Servers	✓	✓			
Stop/Restart ElastiCubes	✓	✓	✓		
Set Up Data Security	✓	✓	✓		
Access Data Page	✓	✓	✓		
Dashboard Design					
Create/Edit/Delete Dashboards	✓	✓	✓	✓	
Copy Dashboard to Another Server	✓	✓	✓	✓	
Restore Dashboards	✓	✓	✓	✓	
Change Color Palettes	✓				
Create/Edit/Delete Widgets	✓	✓	✓	✓	
Rename Widgets	✓	✓	✓	✓	
Duplicate Widgets	✓	✓	✓	✓	
Change Widget Type	✓	✓	✓	✓	
Drill into Widget	✓	✓	✓	✓	✓
Reorder Widgets	✓	✓	✓	✓	
Create/Delete Dashboard Filters	✓	✓	✓	✓	
Toggle Dashboard Filters On/Off	✓	✓	✓	✓	✓

	Admin	Data Admin	Data Designer	Designer	Viewer
Create/Delete Widget Filters	✓	✓	✓	✓	
Toggle Widget Filters On/Off	✓	✓	✓	✓	
Sharing					
Export Dashboards to PDF, Img, CSV	✓	✓	✓	✓	✓
Share Dashboards	✓	✓	✓	✓*	
Migrate Dashboards	✓	✓	✓	✓*	
Pulse					
Access Pulse	✓	✓	✓	✓	✓
Create Data Alert	✓	✓	✓	✓	✓
Create Build Alert	✓	✓			
Delete Alerts *	✓	✓	✓	✓	✓

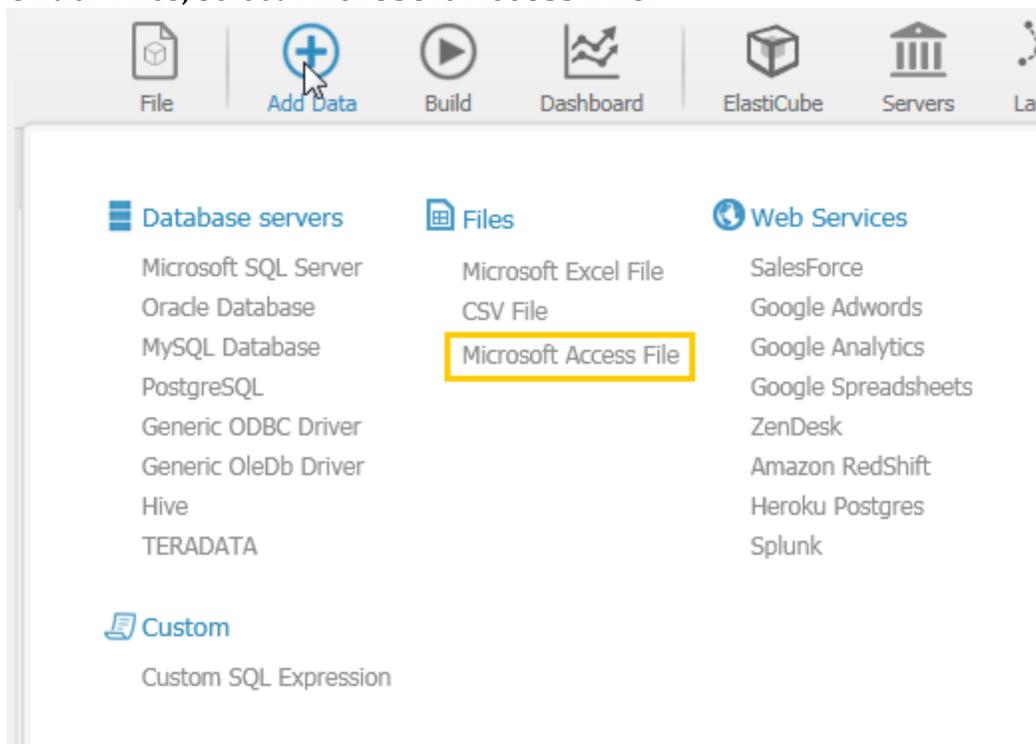
\* The user must be the owner of the dashboard, ElastiCube, or alert.

# Connecting to Access

The ElastiCube Manager enables easy and quick access to tables contained within Microsoft Access files.

## To connect to Microsoft Access:

1. Click **Add data** in the top menu of the Desktop version of the ElastiCube Manager.
2. Under Files, select **Microsoft Access File**.

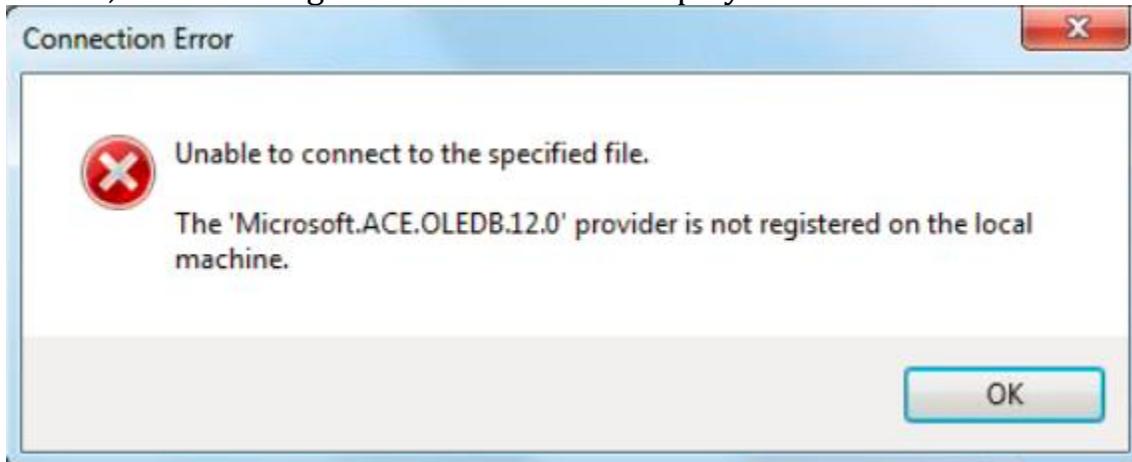


3. Select the Microsoft Access file you want to use. Click **OK**.  
A popup will open with the following options.  
All tables and views associated with the database will appear in a new window. To preview the data contained in a particular table, highlight the table or view in the list and click the preview pane below. To preview the table, select the **Preview** checkbox.
4. Select the checkboxes next to each table or view you want to use.
5. Existing relationships between tables can be automatically replicated in the ElastiCube by selecting the **Automatically create relationships from database** checkbox. Likewise fields with similar names can be linked by

selecting the **Automatically create relationships for fields** with the same name checkbox.

6. Once all relevant tables are selected, click **OK**.

**Note:** Sisense does not support Access files that were created with the 32-bit version of Access. If you try to upload an Access file created in a 32-bit version of Access, the following connection error is displayed.



To upload 32-bit Access files, you can install Microsoft's 64-bit [Access Database Engine](#) to upload your Access files or click [here](#) for more information.

# Connecting to Bing

The Sisense Bing connector is a standalone connector that allows you to import data from Bing's APIs into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to Bing's API. To obtain a connection string, you will need to create a Bing app.

Once you have connected to Bing, you can import a variety of tables from the Bing API.

This page describes how to install the Bing connector, how to connect to Bing with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the Bing Connector](#)
- [Connecting to the Bing REST API](#)
- [Adding Bing Tables to your ElastiCube](#)
- [Bing Tables](#)
- [Versions](#)

## Installing the Bing Connector

Sisense provides the Bing connector as a standalone connector that you can download and add to your list of default Sisense connectors.

### To install the Bing connector:

1. [Download](#) the Bing installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The Bing connector is displayed in the ElastiCube Manager under **Add Data** > **Web Services**.

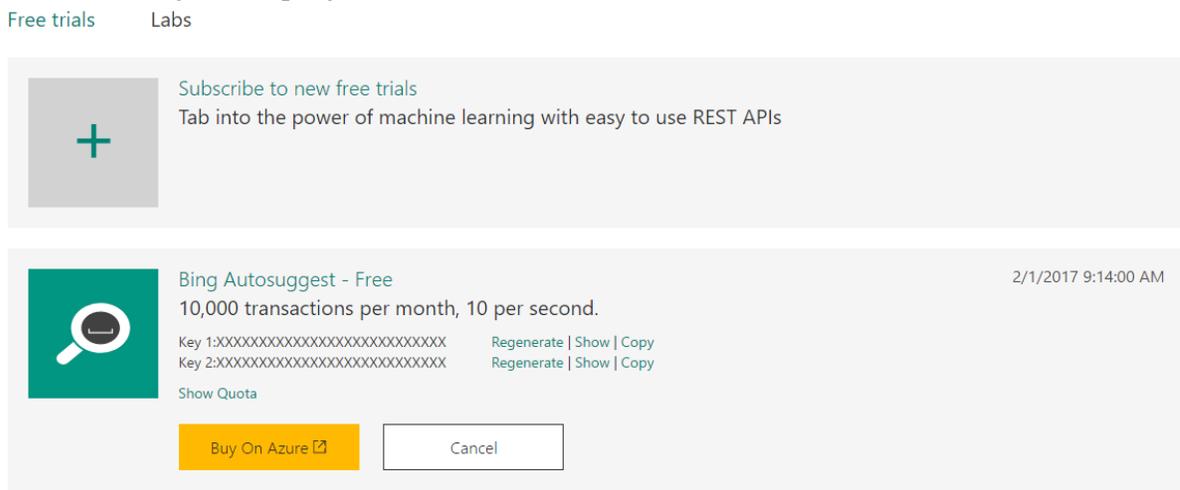
# Connecting to the Bing REST API

To access Bing's REST API from Sisense, you must obtain a valid API key. The API Key is provided by Bing after you subscribe to an API.

## Obtaining an API Key:

1. Log in to <https://www.bing.com/dev>.
2. From the main menu, select the API you want to retrieve data from.
3. If you have already subscribed to it, you will see two subscription types, Free Trial and Labs. Select the relevant subscription type.

The API Key is displayed.



4. Click **Copy** for the relevant key.
5. Save the key to paste it into your connection string.

## Adding Bing Tables to your ElastiCube

Sisense uses connection strings to connect to Bing and import data into the ElastiCube Manager.

The connection string to connect to Bing has the following structure:

```
jdbc:bing:ApiKey=73529b7a4994ytec8d4ca6c33049;InitiateOAuth=GETANDREFRESH;
```

Mandatory properties are ApiKey, and InitiateOAuth.

- **ApiKey:** Application key provided by Microsoft that authenticates your account.

- **InitiateOAuth:** Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken connection property.

**Note:** To switch between accounts, you need to delete the file OAuthsettings.txt file located at ...\\Users\\xxx\\AppData\\Roaming\\CData\\Bing Data Provider.

#### **To add Bing data:**

1. In the ElastiCube Manager, click **Add Data** and then, **Bing**. The Connect to Bing window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. Bing is displayed in the Select Database list.
4. Click **OK**. Sisense connects to Bing and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

## Switching between Accounts

When you connect to the Bing data source, Sisense saves your OAuth values in the file OAuthsettings.txt file located at ...\\Users\\xxx\\AppData\\Roaming\\CData\\Bing Data Provider on your Sisense server. To connect to the Bing data source with another user on the same machine, you must delete the OAuthsettings.txt file. Sisense will then generate a new file for that user.

Another option to support multiple users is to define the location and file name of an OAuthsettings file for each unique user in your connection string through the **OAuthSettingsLocation** parameter. When each user connects to the data source, Sisense generates the OAuth file with the file name you specify in the location you define. In the examples below, two users are allowed to access the Bing data source and for each user, Sisense generates a file that contains that user's OAuth values in the location defined in the string.

```

jdbc:OauthSettingsLocation=C:\Bing\auth\john.txt;OauthClientId=11276856774486;OauthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOauth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
jdbc:OauthSettingsLocation=C:\Bing\auth\sally.txt;OauthClientId=11276856774486;OauthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOauth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;

```

In the example above, two OAuth files are created, one for John and one for Sally in the location C:\Bing\auth\.

This is useful if you support many users who each need to access the Bing data source.

## Bing Tables

The Sisense Bing connector allows you to import the following tables into the ElastiCube Manager.

Name	Description
ImageSearch	<p>Query the Bing image search engine. When querying this table the SearchTerms parameter must be specified. For example, to search Bing Images for the term 'bing' use the following query:</p> <pre>SELECT * FROM ImageSearch where SearchTerms = 'bing'</pre> <p>You can also specify a search safety level:</p> <pre>SELECT * FROM ImageSearch where SearchTerms = 'bing' and SafeSearch = 'Moderate'</pre>
NewsSearch	<p>Query the Bing news search engine.</p> <p>When querying this table the SearchTerms parameter must be</p>

Name	Description
	<p>specified. For example, to search Bing Images for the term 'bing' use the following query:</p> <pre>SELECT * FROM NewsSearch where SearchTerms = 'bing'</pre> <p>You can also specify a search safety level:</p> <pre>SELECT * FROM NewsSearch where SearchTerms = 'bing' and SafeSearch = 'Moderate'</pre>
Search	<p>Query the Bing search engine.</p> <p>When querying this table the SearchTerms parameter must be specified. For example, to search Bing Images for the term 'bing' use the following query:</p> <pre>SELECT * FROM Search where SearchTerms = 'bing'</pre> <p>You can also specify a search safety level:</p> <pre>SELECT * FROM Search where SearchTerms = 'bing' and SafeSearch = 'Moderate'</pre>
VideoSearch	<p>Query the Bing video search engine.</p> <p>When querying this table the SearchTerms parameter must be specified. For example, to search Bing Images for the term 'bing' use the following query:</p> <pre>SELECT * FROM VideoSearch where SearchTerms = 'bing'</pre> <p>You can also specify a search safety level:</p> <pre>SELECT * FROM VideoSearch where SearchTerms = 'bing' and SafeSearch = 'Moderate'</pre>
WebSearch	<p>Query the Bing web search engine.</p>

Name	Description
	<p>When querying this table the SearchTerms parameter must be specified. For example, to search Bing Images for the term 'bing' use the following query:</p> <pre>SELECT * FROM WebSearch where SearchTerms = 'bing'</pre> <p>You can also specify a search safety level:</p> <pre>SELECT * FROM WebSearch where SearchTerms = 'bing' and SafeSearch = 'Moderate'</pre>

### Versions

Version	Release Date	Improvements
6.6.0.16	March 16, 2017	General release
6.5.1.3	February 8, 2017	General Beta release

# Connecting to Box

The Sisense Box connector is a standalone connector that allows you to import data from Box's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the Box APIs. To obtain a connection string, you will need to register a Box app.

Once you have connected to Box, you can import a variety of tables from the Box API.

This page describes how to install the Box connector, how to connect to Box with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the Box Connector](#)
- [Connecting to the Box REST API](#)
- [Adding Box Tables to your ElastiCube](#)
- [Box Tables](#)
- [Limitations](#)

## Installing the Box Connector

Sisense provides the Box connector as a standalone connector that you can download and add to your list of default Sisense connectors.

### To install the Box connector:

1. [Download](#) the Box installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The Box connector is displayed in the ElastiCube Manager under **Add Data** > **Web Services**.

# Connecting to the Box REST API

To access Box's REST API from the ElastiCube Manager, you must provide valid OAuth Box credentials through a connection string. These credentials are provided by Box when you create an application.

After you receive your credentials from Box, you can create the connection string and provide Sisense with it to connect to your data.

## Creating an App

You can follow the procedure below to register an app and obtain the OAuth client credentials, the OAuthClientId and OAuthClientSecret:

1. Log in to your Box developers dashboard and click **Create a Box Application**.
2. If you are making a desktop application, set the Redirect URI to <http://localhost:33333> or a different port number of your choice. If you are making a web application, set the Redirect URI to a page on your Web app you would like the user to be returned to after they have authorized your application. The OAuthClientId and OAuthClientSecret are also displayed in the same page.
3. Select the scope of user permissions your app will request.

## Adding Box Tables to your ElastiCube

Sisense uses connection strings to connect to Box and import data into the ElastiCube Manager. Connection strings contain authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

Box requires the following parameters that you receive when registering an app with Box:

- **OAuthClientId**: The client ID assigned when you register your application with an OAuth authorization server.
- **OAuthClientSecret**: Set this to the App Secret in your app settings.

- **CallbackURL:** The OAuth callback URL to return to when authenticating. This value must match the callback URL you specify in your app settings. Set this to `http://localhost`.
- **InitiateOAuth:** Set this to `GETANDREFRESH`. You can use `InitiateOAuth` to avoid repeating the OAuth exchange and manually setting the `OAuthAccessToken` connection property.
- **Timeout:** The value in seconds until the timeout error is thrown, cancelling the operation. Set to 0 to disable time outs.

For more information, see [Connection String Parameters](#).

Connection strings have the following structure:

```
jdbc:DataSourceName:Property1=Value1;Property2=Value2;
```

The following is an example of a Exact connection string:

```
jdbcbox:OAuthClientId=xxxxxxxxxxxxxx;OAuthClientSecret=xxxxxxxxxxxxxx;Region='United States';Division=xxxx;InitiateOAuth=GETANDREFRESH;OAuthCallbackURL=localhost:12345;
```

#### To add Box data:

1. In ElastiCube Manager, click **Add Data** and then, **Box**. The Connect to Box window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. Box is displayed in the Select Database list.
4. Click **OK**. Sisense connects to Box and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

## Box Tables

Box's RESTful APIs expose the following Box tables that you can import into the ElastiCube Manager through the Sisense Box connector:

#### Available Tables

Name	Description
------	-------------

Name	Description
Collaborations	Create, update, delete, and query the available Collaborations in Box.
Comments	Create, update, delete, and query the available Comments in Box.
Files	Update, delete, and query the available Files in Box.
Folders	Create, update, delete, and query the available Folders in Box.
Groups	Create, update, delete, and query the available Groups in Box.
Memberships	Create, update, delete, and query the available Memberships in Box.
TaskAssignments	Create, update, delete, and query the available Task Assignments in Box.
Tasks	Create, update, delete, and query the available Tasks in Box.
TrashedItems	Delete and query the available TrashedItems in Box.
Users	Create, update, delete, and query the available Users in Box.

#### Available Views

Name	Description
SharedItems	Query the available shared items in Box.
UserEvents	Query the available Events in Box.

## Limitations

Aggregate functions are not supported.

To query the Collaborations table, specify an Id, FolderId, or a GroupId.

```
SELECT * FROM Collaborations WHERE Id = '123'
```

To select from the Tasks table you need to specify an Id or an ItemId.

```
SELECT * FROM Tasks WHERE Id = '123'
```

To select from the TasksAssignments table you need to specify an Id or a TaskId.

```
SELECT * FROM TasksAssignments WHERE Id = '123'
```

To select from the Memberships table you need to specify an Id or a GroupId.

```
SELECT * FROM Memberships WHERE Id = '123'
```

To select from SharedItems specify the Shared Link that points to the item.

```
SELECT * FROM SharedItems WHERE SharedLink =  
'www.url.com'
```

If the SharedItem has a password also specify the password.

```
SELECT * FROM SharedItems WHERE SharedLink =  
'www.url.com' AND SharedLinkPassword = 'password'
```

To query the Comments table specify an Id or a FileId.

```
SELECT * FROM Comments WHERE Id = '123'
```

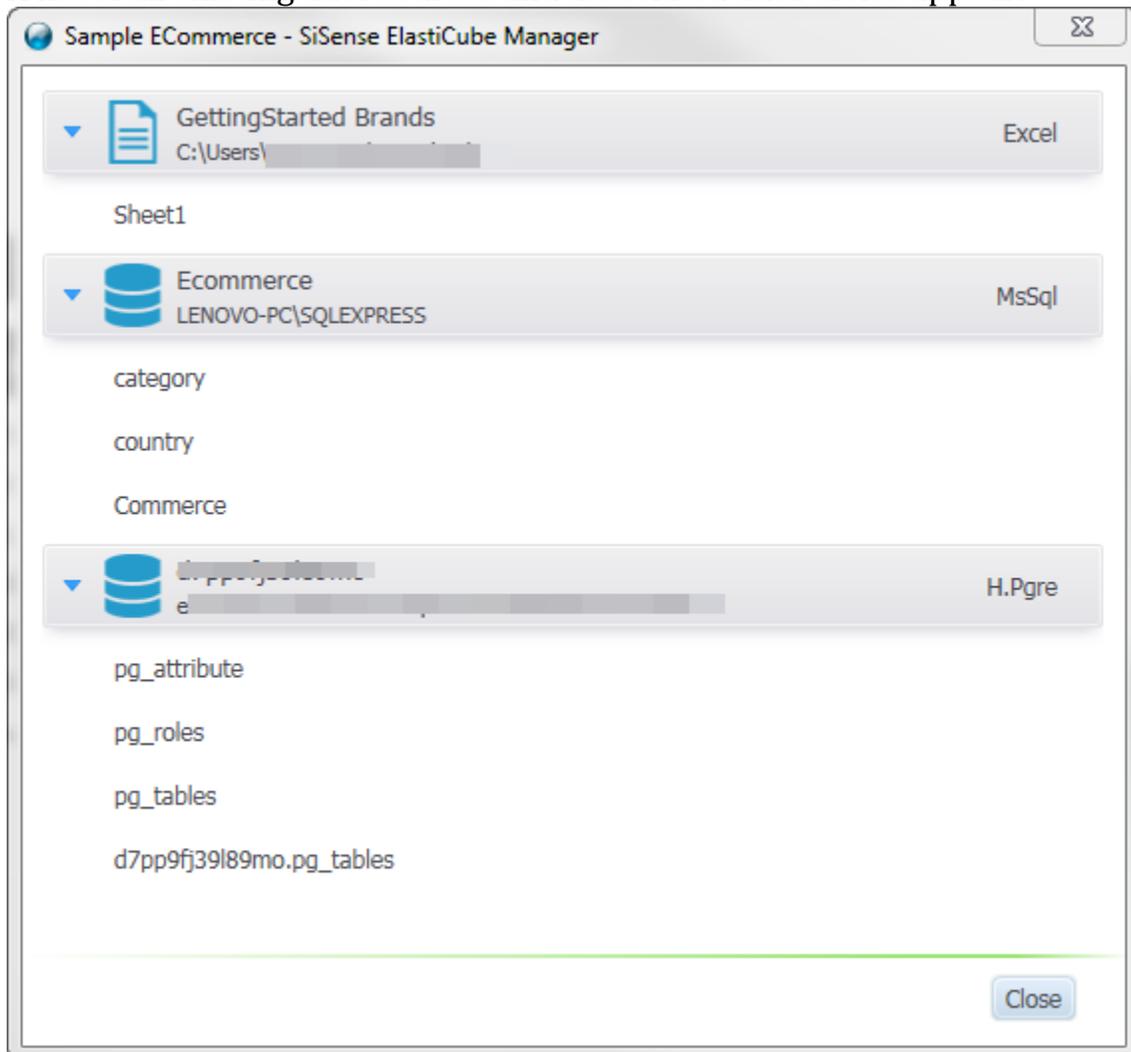
Due to these limitations, the ElastiCube cannot be built without custom SQL.

# Changing Connectivity Settings for Data Sources

**Note:** This page describes how to change connectivity settings in the desktop version of the ElastiCube Manager. This feature will be available online in the near future. For more information see [Sisense V7.0 Backwards Compatibility](#). This section describes how to change the connectivity settings for an existing data source. You can update the source of the data, on both the database and table levels. For example, a database or data file may have been moved to a new location, or updated, and the ElastiCube needs to be updated accordingly.

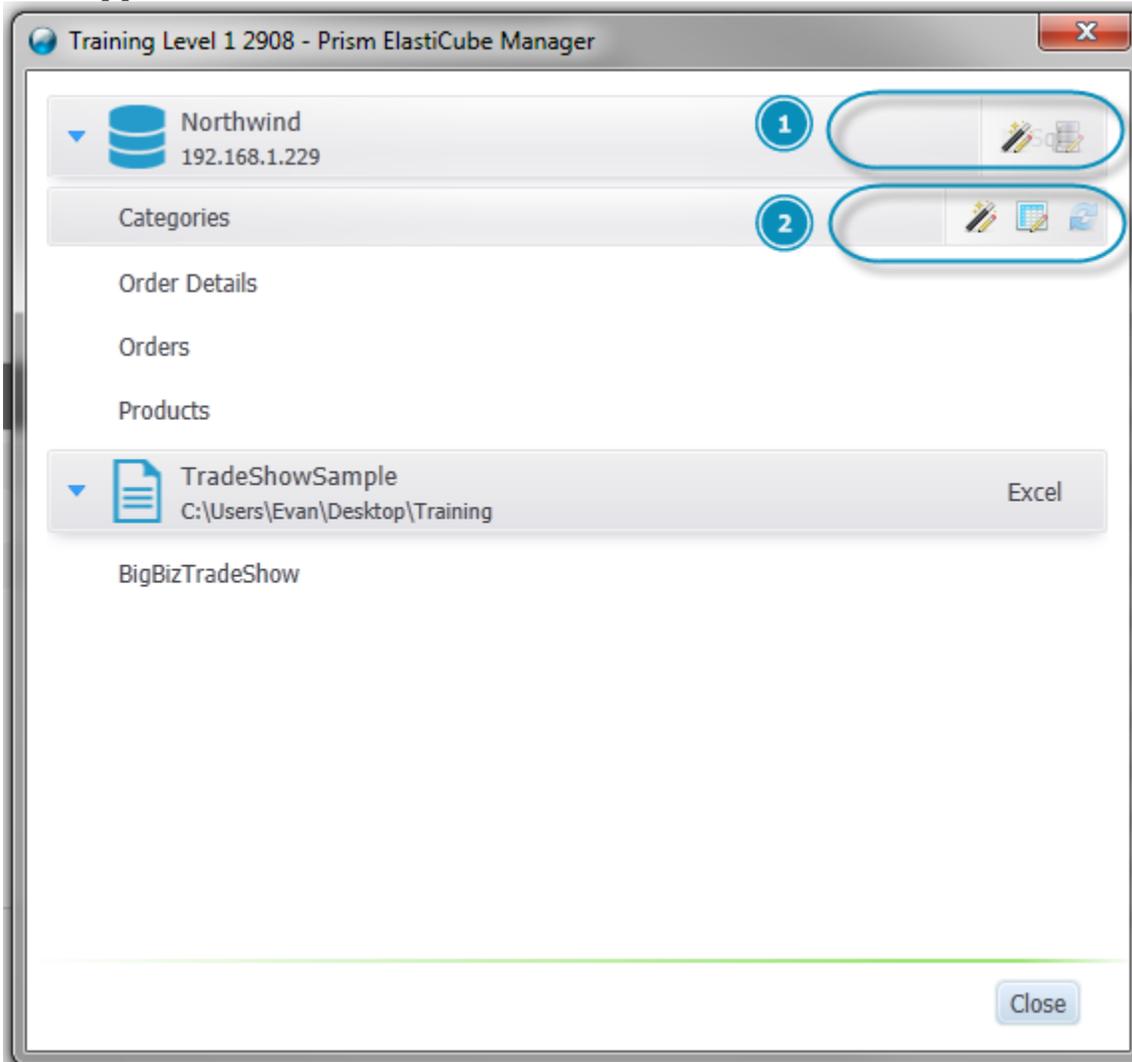
1. In the top menu, Click ElastiCube, and under **Build Settings**, select **Change Connectivity Settings...**

A list of all existing data sources and associated data tables appear.



2. Hover over the data source to reveal the following options:
  - **Change Source Provider:** Click to select a data source connection type, for example, change MS SQL to an Oracle database. When done, click **Close**.
  - **Change Source Database:** Click to edit the login details to the database as well as connections to associated tables.
3. Hover over the data table to reveal the following options:
  - **Change Source Provider:** Click to change the data source connection type, for example, from MS SQL to an Oracle database.
  - **Change Source Table:** Click to change the source tables using the current connection settings.
  - **Refresh:** Refresh the tables' fields based on current settings.

- **Note:** The data source name, fields and field types must remain the same for existing widgets in the dashboard to be able to reference the data correctly.
4. To remove a data source, you must first [delete all related tables](#) from the ElastiCube schema. After deleting all tables, the data source will no longer appear.



# Connection String Parameters

Sisense uses connection strings to connect to some data sources and import data into the ElastiCube Manager.

Connection strings have the following structure:

```
jdbc:DataSourceName:Property1=Value1;Property2=Value2;
```

The following is an example of a Exact connection string:

```
jdbc:exactonline:OAuthClientId=xxxxxxxxxxxxxxxx;OAuthClientSecret=xxxxxxxxxxxxxxxx;Region='United States';Division=xxxx;InitiateOAuth=GETANDREFRESH;OAuthCallbackURL=localhost:12345;
```

This page lists the mandatory and optional parameters required by the data source you want to connect to. Mandatory parameters must be defined and included in your connection string whereas optional parameters are only recommended by Sisense.

In addition, some parameters are required by specific data sources. You can view these parameters by clicking the relevant link below:

[Exact](#)

[Facebook](#)

[HubSpot](#)

[Quickbooks Online](#)

[ServiceNow](#)

[Xero](#)

[YouTube Analytics](#)

## Mandatory Properties

- **OAuthClientId:** The client ID assigned when you register your application with an OAuth authorization server.
- **OAuthClientSecret:** Set this to the App Secret in your app settings.
- **CallbackURL:** The OAuth callback URL to return to when authenticating. This value must match the callback URL you specify in your app settings. Set this to `http://localhost`.
- **InitiateOAuth:** Set this to `GETANDREFRESH`. You can use `InitiateOAuth` to avoid repeating the OAuth exchange and manually setting the `OAuthAccessToken` connection property.
- **Timeout:** The value in seconds until the timeout error is thrown, cancelling the operation. Set to 0 to disable time outs.

## Optional Parameters

- **OAuthSettingsLocation:** The location of the settings file where OAuth values are saved for each user when `InitiateOAuth` is set to `GETANDREFRESH`. See [Switching between Accounts](#).
- **Logfile:** Creates a log file. The value of this parameter is the location on your drive where the log file is to be created. To determine how much detail is included in the log file, see the parameter `Verbosity` below.
- **Verbosity:** The verbosity level that determines the amount of detail included in the log file. Set the verbosity from 1 – 5.
  - 1 will log the query, the number of rows returned by it, the start of execution and the time taken, and any errors.
  - 2 will log everything included in Verbosity 1, cache queries, and HTTP headers.
  - 3 will additionally log the body of the HTTP requests.
  - 4 will additionally log transport-level communication with the data source. This includes SSL negotiation.
  - 5 will additionally log communication with the data source and additional details that may be helpful in troubleshooting problems. This includes interface commands.

## Switching between Accounts

When you connect to the Facebook data source, Sisense saves your OAuth values in the file OAuthsettings.txt file located at

...\Users\xxx\AppData\Roaming\CData\Facebook Data Provider on your Sisense server. To connect to the Facebook data source with another user on the same machine, you must delete the OAuthsettings.txt file. Sisense will then generate a new file for that user.

Another option to support multiple users is to define the location and file name of an OAuthsettings file for each unique user in your connection string through the **OAuthSettingsLocation** parameter. When each user connects to the data source, Sisense generates the OAuth file with the file name you specify in the location you define. In the examples below, two users are allowed to access the Facebook data source and for each user, Sisense generates a file that contains that user's OAuth values in the location defined in the string.

```
jdbc:facebook:OAuthSettingsLocation=C:\facebook\auth\john.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
```

```
jdbc:facebook:OAuthSettingsLocation=C:\facebook\auth\sally.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
```

In the example above, two OAuth files are created, one for John and one for Sally in the location C:\facebook\auth\.

This is useful if you support many users who each need to access the Facebook data source.

# Data Source Parameters

This section lists parameters that are required to be inside connection strings

## Exact

- **Region:** Set this to the region of the Exact Online service you want to connect to.
- **Division:** Set this to the division of the Exact Online administration.

The following is an example of an Exact connection string:

```
jdbc:exactonline:OAuthClientId=xxxxxxxxxxxxxxxx;OAuthClientSecret=xxxxxxxxxxxxxxxx;Region='United States';Division=xxxx;InitiateOAuth=GETANDREFRESH;OAuthCallbackURL=localhost:12345;
```

## Facebook

### Rate Limitations

- **ThrottleSeconds:** Indicates how many seconds to wait until sending another insight request when the ThrottleMaxPercent has been met.
- **ThrottleMaxPercent:** The threshold set for throttling further insight requests. The defaults for ThrottleSeconds and ThrottleMaxPercent are 10 and 95.
- **RetryWaitTime:** The amount of time in milliseconds to wait before retrying to query the Facebook API.
- **MaximumRequestRetries:** The maximum amount of request retries.

The following is an example of a Facebook connection string:

```
jdbc:facebook:OAuthClientId=xxxxxxxxxxxxxxxx;OAuthClientSecret=xxxxxxxxxxxxxxxx;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;Timeout=0;
```

### Optional Parameters

- **Aggregate Format:** The format aggregate or collection columns should return in.
- **Authenticate As Page:** The name or Id of a page to authenticate as when making requests to Facebook.
- **Auto Cache:** The AutoCache property controls automatic caching of data.
- **Cache Connection:** The connection string for the cache database.
- **Cache Driver:** The database driver to be used to cache data.
- **Cache Location:** Specifies the path to the cache when caching to a file.
- **Cache Metadata:** Whether or not to cache the table metadata. The cache will be stored in memory unless CacheConnection or CacheLocation are set.
- **Callback URL:** The OAuth callback URL to return to when authenticating. This value must match the callback URL you specify in your app settings.
- **Firewall Password:** A password, if authentication is required to connect through a firewall.
- **Firewall Port:** The TCP port for the firewall FirewallServer — see the description of the FirewallServer option for details.
- **Firewall Server:** Specify a firewall name or IP address to authenticate requested connections, if necessary.
- **Firewall Type:** The type of firewall to connect through.
- **Firewall User:** The user name to authenticate with the firewall.
- **Initiate OAuth:** Set this property to initiate the process to obtain or refresh the OAuth access token when you connect.
- **Location:** A path to the directory that contains the schema files defining tables, views, and stored procedures.
- **Logfile:** A path to the log file.
- **Max Log File Size:** A string specifying the maximum size in bytes for a log file (ex: 10MB). When the limit is hit, a new log is created in the same folder with the date and time appended to the end.
- **OAuth Access Token:** The access token for connecting using OAuth.
- **OAuth Client Id:** The client Id assigned when you register your application with an OAuth authorization server.
- **OAuth Client Secret:** The client secret assigned when you register your application with an OAuth authorization server.
- **OAuth Settings Location:** The location of the settings file where OAuth values are saved when InitiateOAuth is set to GETANDREFRESH or REFRESH.
- **Offline:** Use offline mode to get the data from the cache instead of the live source.

- **Other:** The other parameters necessary to connect to a data source, such as username and password, when applicable.
- **Pagesize:** The maximum number of results to return per page from Facebook.
- **Proxy Auth:** Scheme The proxy server authorization scheme (default: BASIC).
- **Proxy Auto Detect:** This indicates whether to use the default system proxy settings or not. Set ProxyAutoDetect to FALSE to use custom proxy settings. This takes precedence over other proxy settings.
- **Proxy Password:** A password, if authenticating with a proxy server.
- **Proxy Port:** The TCP port for the proxy ProxyServer (default: 80).
- **Proxy Server:** If a proxy server is given, then the HTTP request is sent to the proxy instead of the specified server.
- **Proxy SSL Type:** The SSL type to use when connecting to the proxy server (default: AUTO).
- **Proxy User:** A user name, if authentication is to be used for the proxy.
- **Pseudo Columns:** Indicates whether or not to include pseudo columns as columns to the table.
- **Query Cache:** The time duration, in seconds, for which the in-memory cached data is reused.
- **Readonly:** You can use this property to enforce read-only access to Facebook from the provider.
- **RTK:** The runtime key used for licensing.
- **Search Terms:** Default SearchTerms if none are specified. Used for some tables, such as Users, where SearchTerms may be specified.
- **SSL Server Cert:** The certificate to be accepted from the server when connecting using TLS/SSL.
- **Support Enhanced SQL:** If set to true, the provider will cache the entire table in memory and issue the current query against the memory cache. Allows queries with aggregate functions that are otherwise unsupported. The entire memory cache is discarded after the query is executed, unlike QueryCache.
- **Tables:** Comma-separated list of tables to be listed. For example: Tables=TableA,TableB,TableC.
- **Target:** A default target if none is specified. Used for some tables, such as Comments, where a target may be specified.
- **Timeout:** The value in seconds until the timeout error is thrown, canceling the operation.

- **Verbosity:** The verbosity level that determines the amount of detail included in the log file.
- **Version:** The Facebook Graph API version to use.

## HubSpot

- **HubId:** Set this to the Hub Id of the HubSpot account you are connecting to.

The following is an example of a HubSpot connection string:

```
jdbc:hubspot:OAuthClientId=xxxxxxxxxxxxxxxx;OAuthClientSecret=xxxxxxxxxxxxxxxx;InitiateOAuth=GETANDREFRESH;HubId=xxxxxxxx;CallbackURL=https://localhost;Timeout=0;
```

## Quickbooks Online

- **CompanyId:** Set this to the ID of the company you want to connect to.

The following is an example of a Quickbooks Online connection string:

```
jdbc:quickbooksonline:OAuthClientId=xxxxxxxxxxxxxxxx;OAuthClientSecret=xxxxxxxxxxxxxxxx;CompanyId=xxxxxxxxxxxxxxxx;CallbackURL=http://localhost/;InitiateOAuth=GETANDREFRESH;Timeout=0;
```

## ServiceNow

- **Password:** Set this to your password.
- **Username:** Set this to your username.
- **Instance:** Set this to your instance.

The following is an example of a ServiceNow connection string:

```
jdbc:servicenow:InitiateOAuth=GETANDREFRESH;Instance=xxxxxx;OAuthClientId=xxxxxxxxxxxxxxxx;OAuthClientSecret=xxxxxxxx;Username=xxxxxxxx;Password=xxxxxxxxxxxxxxxx;
```

## Xero

- **XeroAppAuthentication:** Set this to the type of your application. Allowed values: PUBLIC or PARTNER. Default: PUBLIC.

The following is an example of a Xero connection string:

```
jdbc:xero:OAuthClientId=xxxxxxxxxxxxxxxx;OAuthClientSecret=xxxxxxxxxxxxxxxx;  
InitiateOAuth=GETANDREFRESH;XeroAppAuthentication=PUBLIC;
```

## YouTube Analytics

- **ChannelId:** Set this to the YouTube Analytics Channel (Profile). This can be set to the Id of the channel. If not specified, MINE will be used.
- **ContentOwnerId:** Set this to the Id of the content owner.

The following is an example of a YouTube Analytics connection string:

```
jdbc:youtubeanalytics:OAuthClientId=xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx  
;  
OAuthClientSecret=xxxxxxxxxxxxxxxx;InitiateOAuth=GETANDREFRESH;  
CallbackURL=http://localhost;
```

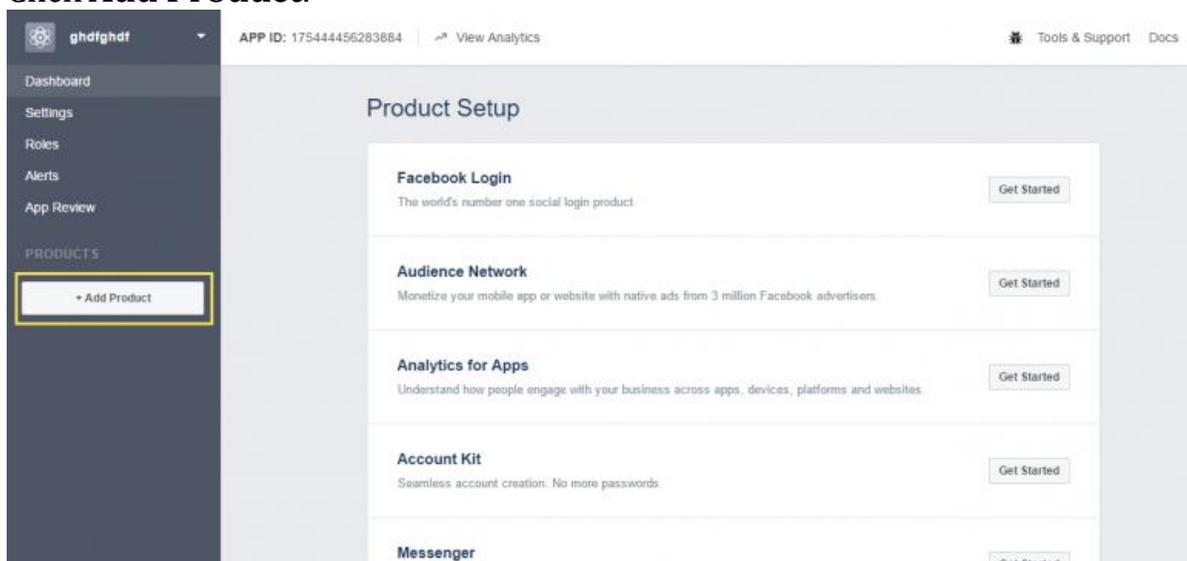
# Creating an App in Facebook

When connecting to Facebook's API through connectors such as [Facebook](#) and [Facebook Ads](#), you need valid credentials that you can obtain when creating a Facebook app.

This page describes how you can create a Facebook app to obtain the necessary credentials you need to import your Facebook data into the ElastiCube Manager. After you have obtained these credentials, you can add them to your connection string to connect to Facebook's APIs.

## To obtain authentication credentials:

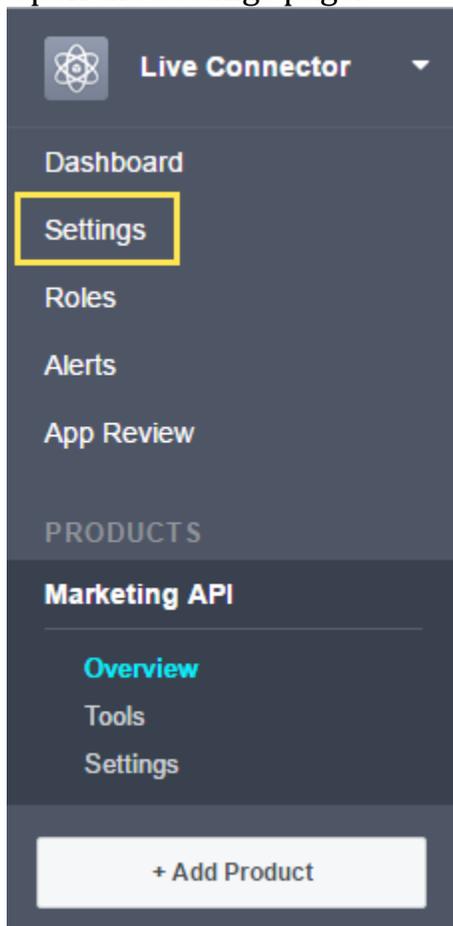
1. Log into Facebook and navigate to <https://developers.facebook.com/apps>.
2. Click **+ Add a New App**.
3. Define your app's name and click **Create App ID**. The Product Setup page is displayed.
4. Click **Add Product**.



5. For the Marketing API, click **Get Started**.



6. Open the Settings page.



7. In the Settings page, the OAuth client credentials, **App Id** and **App Secret**, are displayed.

The screenshot shows the Facebook App Settings page. The 'App Domains' field is highlighted with a yellow box and contains the text 'localhost'. Other fields include 'App ID' (883504611784072), 'App Secret' (masked with dots and a 'Show' button), 'Display Name' (Test), 'Namespace', 'Contact Email' (test@sisense.com), 'Privacy Policy URL' (Privacy policy for Login dialog and App Details), 'Terms of Service URL' (Terms of Service for Login dialog and App Details), 'App Icon' (1024 x 1024), and 'Category' (Books). A '+ Add Platform' button is located at the bottom of the settings area.

8. In the **App Domains** field, enter localhost.
9. In the Settings page, click **+ Add Platform** and select **Website**. The Website area is added to the Settings page.
10. In the **Site URL** field, enter http://localhost. This value represents the callback URL in your connection string where the user will return with the token that verifies that they have granted your app access.

The screenshot shows a 'Website' settings dialog box. The 'Site URL' field contains the text 'http://localhost'. There is a 'Quick Start' button in the top right corner of the dialog.

11. In the bottom-right corner of the Settings page, click **Save Changes**. Facebook requests that you authenticate your account.
12. Enter your password to authenticate your account and click **OK**. Your app is saved.
13. After you have authenticated your account, you must set permissions for the app to retrieve your data. Open the [Facebook Explorer](#) to set your permissions.

14. In Facebook Explorer, click **Get Token > Get User Access Token**.

**Select Permissions** v2.8 ✕

**User Data Permissions**

<input checked="" type="checkbox"/> email ✕	<input type="checkbox"/> user_hometown	<input type="checkbox"/> user_religion_politics
<input type="checkbox"/> publish_actions	<input type="checkbox"/> user_likes	<input type="checkbox"/> user_status
<input type="checkbox"/> user_about_me	<input type="checkbox"/> user_location	<input type="checkbox"/> user_tagged_places
<input checked="" type="checkbox"/> user_birthday ✕	<input type="checkbox"/> user_photos	<input type="checkbox"/> user_videos
<input type="checkbox"/> user_education_history	<input type="checkbox"/> user_posts	<input type="checkbox"/> user_website
<input checked="" type="checkbox"/> user_friends ✕	<input type="checkbox"/> user_relationship_details	<input type="checkbox"/> user_work_history
<input type="checkbox"/> user_games_activity	<input type="checkbox"/> user_relationships	

**Events, Groups & Pages**

<input type="checkbox"/> ads_management	<input type="checkbox"/> pages_messaging	<input type="checkbox"/> read_page_mailboxes
<input type="checkbox"/> ads_read	<input type="checkbox"/> pages_messaging_payments	<input type="checkbox"/> rsvp_event
<input type="checkbox"/> business_management	<input type="checkbox"/> pages_messaging_phone_number	<input type="checkbox"/> user_events
<input type="checkbox"/> manage_pages	<input type="checkbox"/> pages_messaging_subscriptions	<input type="checkbox"/> user_managed_groups
<input type="checkbox"/> pages_manage_cta	<input type="checkbox"/> pages_show_list	<input type="checkbox"/> pages_manage_instant_articles
<input type="checkbox"/> publish_pages		
<input type="checkbox"/> user_actions.books	<input type="checkbox"/> user_actions.music	<input type="checkbox"/> user_actions.video
<input type="checkbox"/> user_actions.fitness	<input type="checkbox"/> user_actions.news	

**Other**

<input type="checkbox"/> read_audience_network_insights	<input type="checkbox"/> read_custom_friendlists	<input type="checkbox"/> read_insights
---	--	--

Public profile included by default

**Get Access Token** Clear Cancel

15. Select the relevant permissions you need and click **Get Access Token**. Your permissions are now set and you can request the relevant data from Facebook.

**Note:** For apps to manage ads, you must be granted the app ads\_management or ads\_read extended permission. Use ads\_read if you only need Ad Insights API access, to pull reporting information. Use ads\_management when you need to read, create and update ads. For more information about permissions, click [here](#).

## Accessing the Facebook API

Facebook has three levels of access to its Marketing API, which determine how you can use and access data from your Facebook app:

- **Development:** Try and test apps with the API.
- **Basic:** Test, iterate, build apps with up to 25 external accounts.

- **Standard:** Support unlimited accounts and be nominated to Facebook Marketing Partner program.

When you create a Facebook, you are assigned a Development level of access by default. Each access level has its own account and rate limitations. You can modify your level of access by submitting a request to Facebook. For more information about these limitations and upgrading your access, click [here](#).

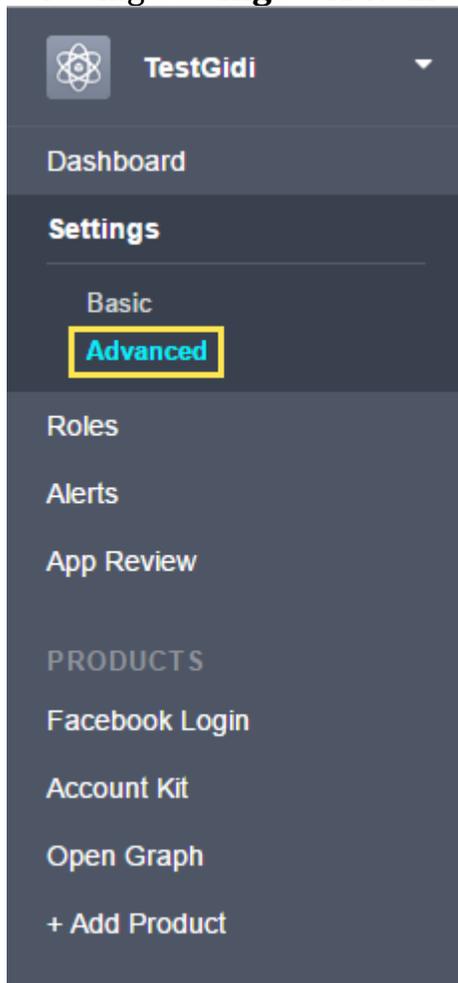
## Adding Advertising Account Permissions

In addition to Facebook permissions, if you want to retrieve data from a Facebook Ads account, that is associated with your Facebook account, you need to add that account's ID to the app through the Authorized Ad Accounts ID field. This field is where you define which ad accounts are associated with your account and allowed to advertise for your app.

### **To associate advertising accounts to your account:**

1. Log into Facebook and navigate to <https://developers.facebook.com/apps>.

2. Selecting **Settings > Advanced**.



3. In the **Authorized Ad Accounts ID** field, enter the Facebook ID of the accounts that can access your Facebook Ads data.

A screenshot of the 'Advertising Accounts' settings screen in the TestGidi app. The screen has a light gray background. At the top left, it says 'Advertising Accounts' and at the top right, there is a gear icon and the text 'Ads API'. Below this, there are two input fields. The first is labeled 'Authorized Ad Account IDs' and the second is labeled 'Authorized Businesses'. Both fields are currently empty.

# Connecting to CSV

The ElastiCube Manager enables easy and quick access to tables contained within CSV files.

There are two options for uploading data, the first is to upload your file to the Sisense Server. Once the file is uploaded, the data will be imported into the ElastiCube as it was at the time the file was uploaded.

The second option is to define the location of your files on the Sisense Server. This is the preferred option if your Excel files frequently change, but remain in the same location on the Sisense Server as the data is taken from the Excel files each time the ElastiCube is built.

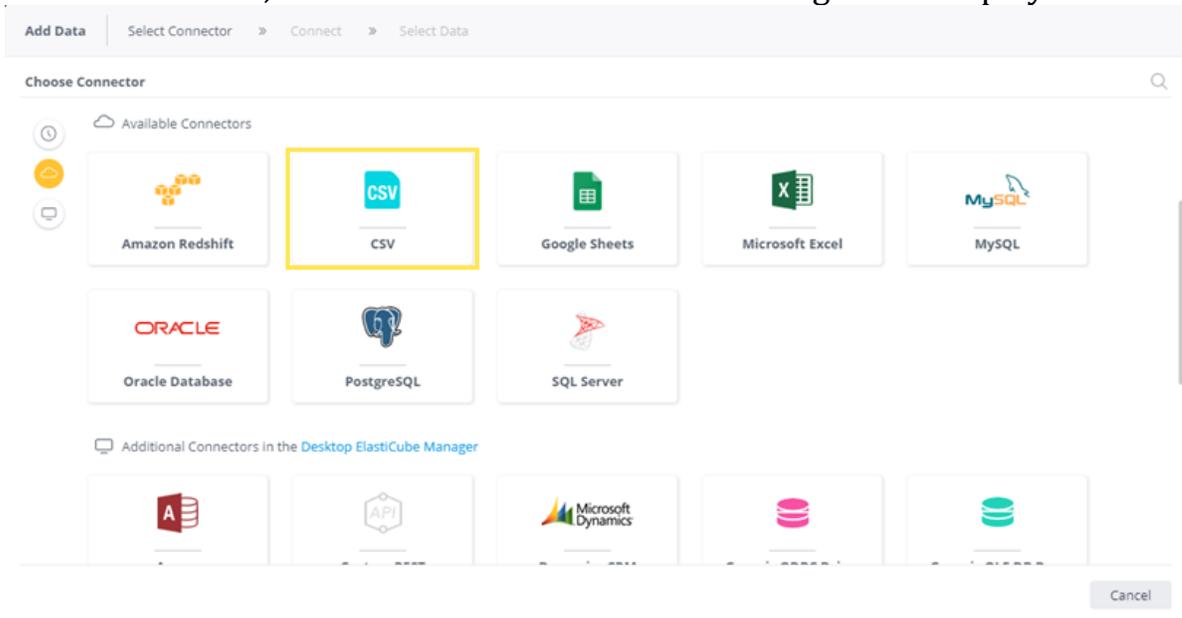
When you import multiple files from a folder and build your ElastiCube, the data is combined together in a single table. To build multiple tables with separate tables for each file, you must repeat the process for each file you want to import as a table.

## Notes:

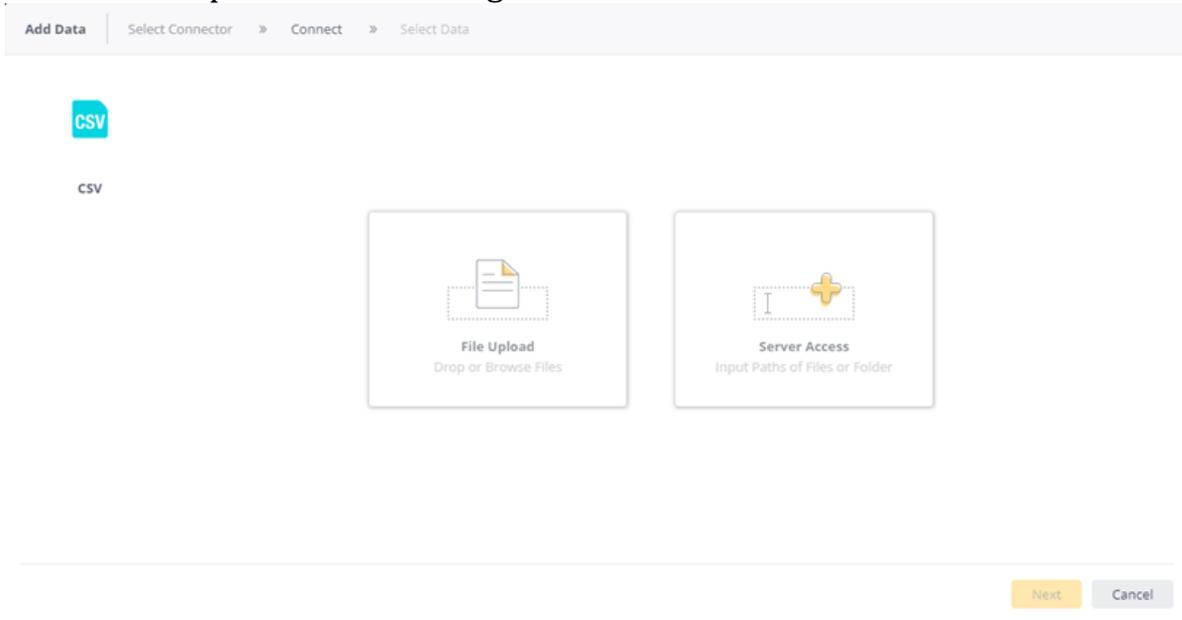
- Make sure that all the files adhere to the same structure, including the presence or absence of headers in the first row. Files will be added to the table based on their file names in alphabetical order.
- When you import multiple files from a folder and build your ElastiCube, the data is combined together in a single table. To build multiple tables with separate tables for each file, you must repeat the process for each file you want to import as a table. If you have any questions about data accumulation between builds, please contact our support team.
- If your CSV file contains special symbols or foreign characters, the import process may fail. To resolve this issue, convert the CSV file to the UTF-8 format. One way to do this is to open the .csv file in an application like Notepad, and then select **File > Save As**. At the bottom of the dialog box, open the Encoding settings, select UTF-8 and save the file as a new file.

**To import data from your CSV files:**

1. In the Data page of ElastiCube Manager, open an ElastiCube or create a new ElastiCube.
2. In the ElastiCube, click **+ Data**. The Add Data dialog box is displayed.



3. Click **CSV** to open the CSV settings.



4. Select the relevant option for importing your data:

### **File Upload**

1. Select this option to import your data from your CSV file. If the file is updated later, you will need to upload it again. To upload the file, click **Browse** and navigate to the file to be uploaded or drag the CSV file to

the File Upload area.

The screenshot shows the 'File Upload' interface. At the top, there are navigation tabs: 'Add Data', 'Select Connector', 'Connect', and 'Select Data'. Below this, there is a 'File Upload' section with a 'Drop files to upload, or browse' prompt. Below this is an 'Uploads' table with columns for 'File Name' and 'Date Uploaded'. A file named 'GettingStarted ECommerce.csv' is listed with a date of '2/21/2018'. At the bottom, there is a 'Union Selected' toggle switch and 'Next' and 'Cancel' buttons.

2. The uploaded file is added to the Uploads list.
3. Select the relevant CSV file(s) to be uploaded.
4. (Optional) Toggle the **Union Selected** switch to append several CSV files together when the data is imported into the ElastiCube. The data in the files must have the same column and data types.
5. After you have finished uploading the relevant CSV files, select them from the upload list and click **Next**. Continue to Step 5.

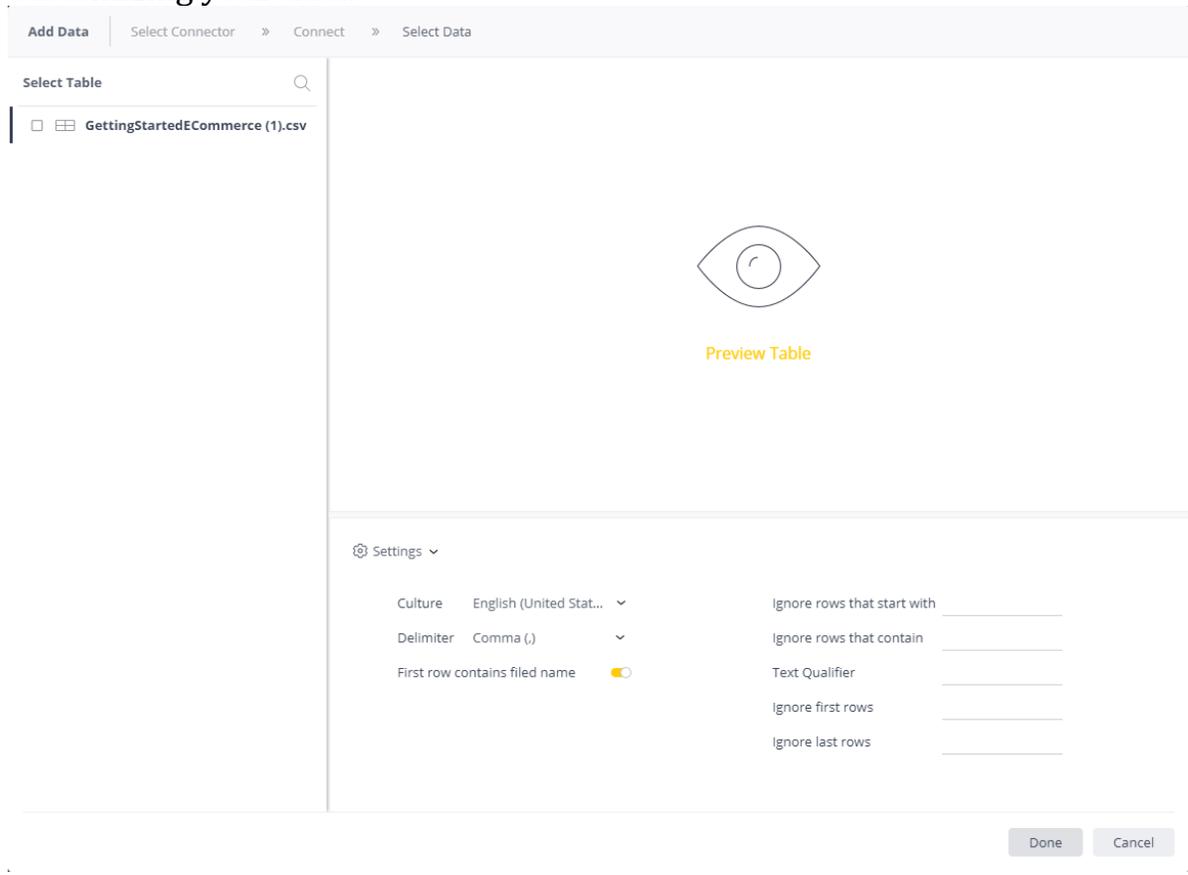
### Server Access

1. Select this option to define the location of your files on the Sisense Server if your CSV files frequently change, but remain in the same location on the Sisense Server as the data is taken from the CSV files each time the ElastiCube is built.
2. Select **Input Folder Path** and enter the full file path with the file name where your Excel files are located. This will display each Excel file in the folder in the next screen where you select what tables to add to the ElastiCube.

OR

Select **Input File Path** and enter the full file path with the file name and its extension of your CSV file. For example, C:\Example.csv. This file display all the tables in the CSV file on the next screen where you select what tables to add to the ElastiCube.

3. (Optional) Toggle the **Union Selected** switch to append several Excel files together when the data is imported into the ElastiCube. The data in the files must have the same column and data types.
4. After you have finished defining the locations of your CSV files, click **Next** and continue to Step 5. A list of CSV files in the directory are displayed.
5. From the **Select Table** list, click  to select preview the columns in the Excel file and display the Settings, which provides more options for customizing your data.



6.

In the **Settings** area, define the following settings:

- **Culture:** Select the culture for your CSV. This defines settings such as the format of the date and time or delimiter (decimal or comma) used in your CSV file.
- **Text Qualifier:** Change the value if necessary.
- **First Row Contains Field Names:** Enables you to specify table column names based on the header in the first row of the spreadsheet.
- **Delimiter:** Allows you to choose the character that separates values within the CSV file.

- **Ignore Rows that Start...:** Specify rows to ignore that start with a specific symbol, value or letter.
  - **Ignore Rows that Contain:** Specify rows to ignore that contain a specific symbol, value or letter.
  - **Ignore First Rows:** Specify a number of first rows to ignore.
7. After you have selected all the relevant tables, click **Done**. The tables are added to your schema in the ElastiCube Manager.

# Connecting to CSV

The ElastiCube Manager enables easy and quick access to tables contained within CSV files.

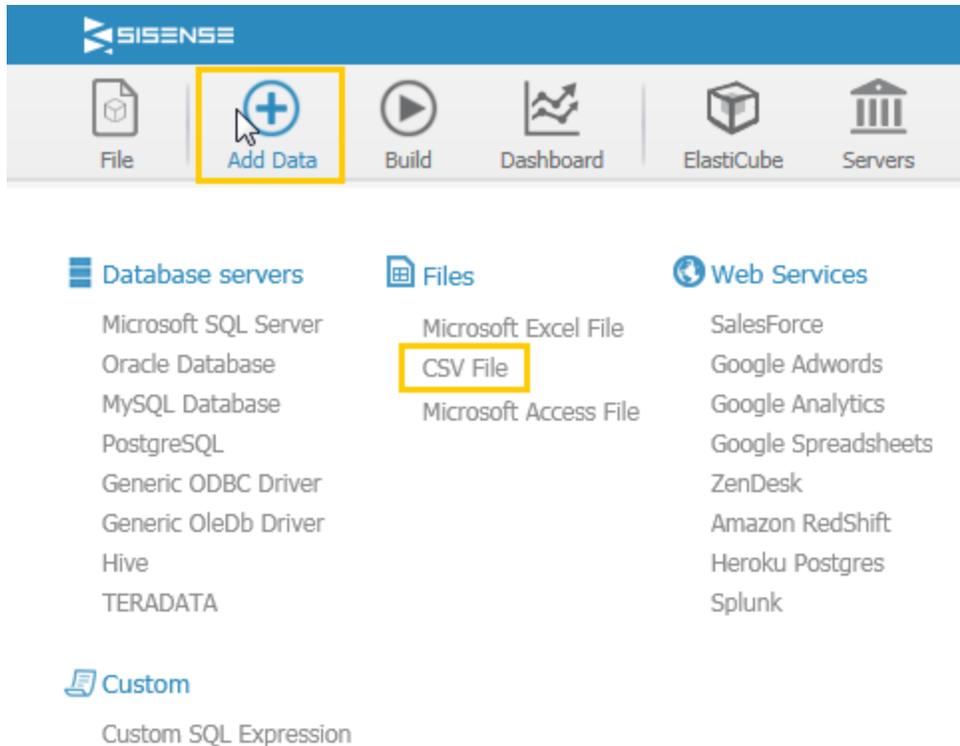
You can add a single CSV file or multiple CSV files to your ElastiCube. A typical use case for importing multiple files is when you want to split a large csv file into multiple smaller files.

## Notes:

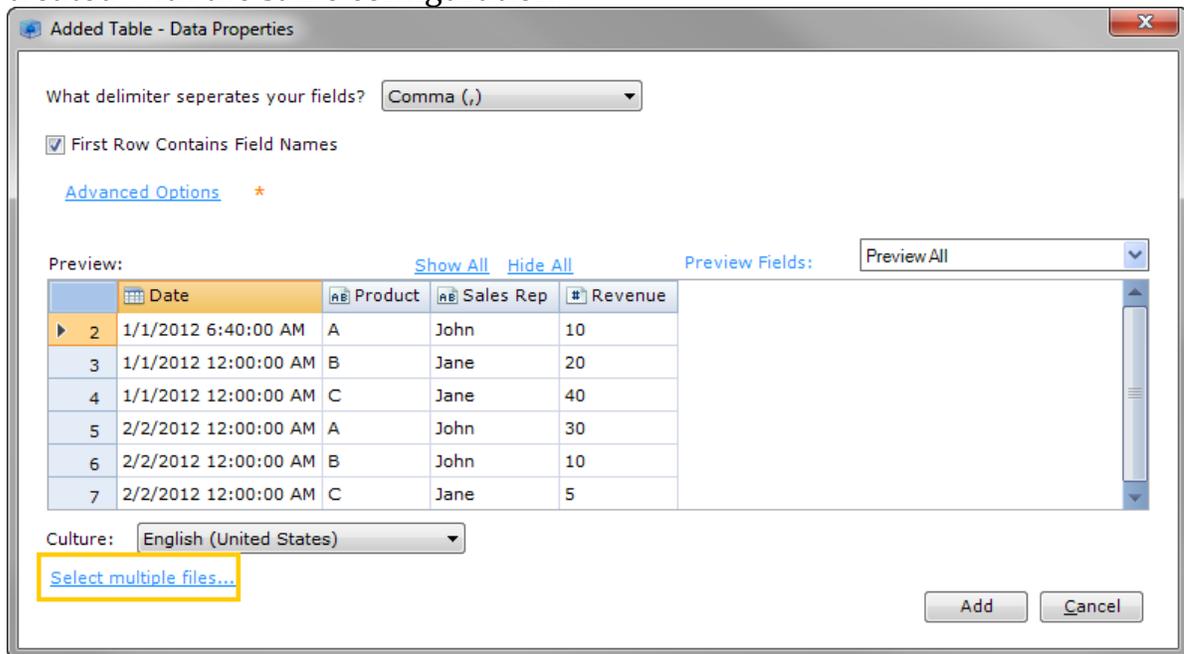
- Make sure that all the files adhere to the same structure, including the presence or absence of headers in the first row. Files will be added to the table based on their file names in alphabetical order.
- When you import multiple files from a folder and build your ElastiCube, the data is combined together in a single table. To build multiple tables with separate tables for each file, you must repeat the process for each file you want to import as a table. If you have any questions about data accumulation between builds, please contact our support team.
- If your CSV file contains special symbols or foreign characters, the import process may fail. To resolve this issue, convert the CSV file to the UTF-8 format. One way to do this is to open the .csv file in an application like Notepad, and then select **File > Save As**. At the bottom of the dialog box, open the Encoding settings, select UTF-8 and save the file as a new file.

## To import data from your CSV files:

1. In the ElastiCube Manager, click **Add Data**.
2. Under the **Files** category select **CSV File**.

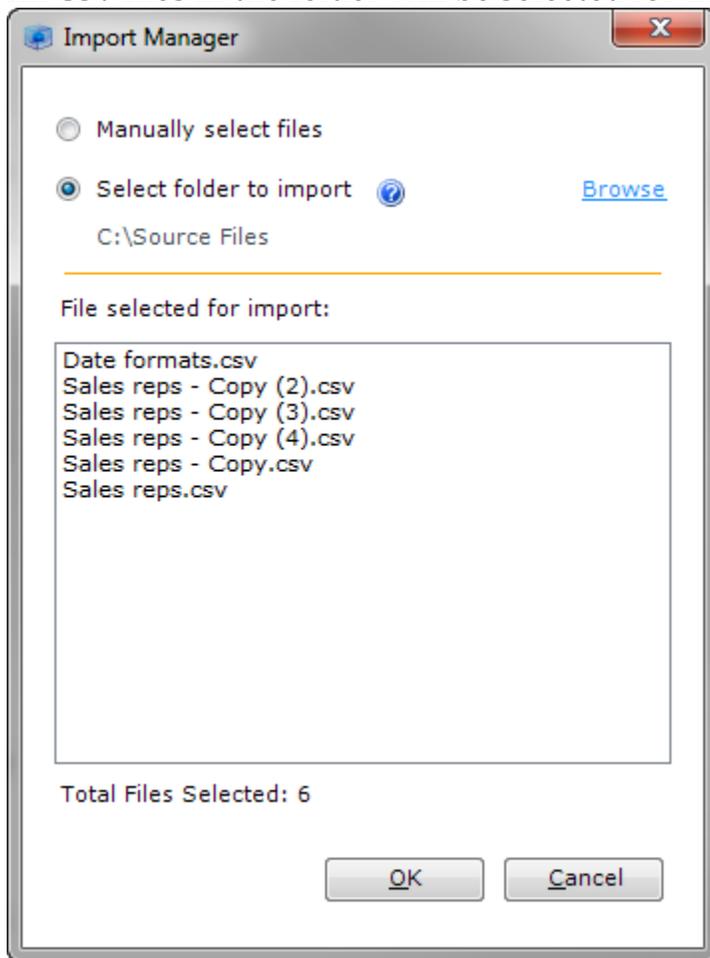


3. Select a file to import. If you want to import multiple files, select just a single file that will be used for previewing the configuration options. However, all files are treated with the same configuration.



4. Define the following settings in the Data Properties window:
- **What delimiter separates your sheets:** Allows you to choose the character that separates values within the csv file.

- **First Row Contains Field Names:** Enables you to specify table column names based on the header in the first row of the spreadsheet.
  - Click **Advanced Options** to reveal more options to filter rows:
    - **Ignore rows that start with:** Specify rows to ignore that start with a specific symbol, value or letter.
    - **Ignore rows that contain:** Specify rows to ignore that contain a specific symbol, value or letter.
    - **Text qualifier:** Change the value if necessary.
    - **Ignore first rows:** Specify a number of first rows to ignore.
5. To add additional files:
  6. Click **Select multiple files...**
  7. Click **Select folder to import.** Locate and select the folder. Click **OK.**
  8. All CSV files in the folder will be selected for import. Click **OK.**



9. Click **Add** to complete the import configuration.

# Customizing Google Analytics and AdWords Tables

## Introduction

Sisense provides a connector to easily access data from both Google Analytics and AdWords. By default certain tables and fields are imported from these sources. A detailed description of these standard tables and fields can be found below.

In addition, you can add customized tables and additional fields. For more details, read the section below on customization.

## Standard Tables from Google Analytics

Below is a list of ready-made reports you can easily access using the Sisense Google Analytics provider.

- **Goals, Daily Report:** Site, date, source, medium, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value, goal2Value, goal3Value, goal4Value
- **Goals, Monthly Report:** Site, year, month, source, medium, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value
- **Goals, Weekly Report:** Site, year, week, source, medium, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value
- **Goals By Keyword, Daily Report:** Site, date, campaign, keyword, hostname, source, medium, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value

- **Goals By Keyword, Daily Report:** Site, year, month, campaign, keyword, hostname, source, medium, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value
- **Goals By Keyword, Daily Report:** Site, year, month, campaign, keyword, hostname, source, medium, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value
- **Goals By Keyword, Weekly Report:** Site, year, week, campaign, keyword, hostname, source, medium, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value
- **Goals By Geography, Daily Report:** Site, date, continent, subcontinent, region, country, city, campaign, keyword, hostname, source, medium, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value
- **Goals By Geography, Monthly Report:** Site, date, continent, subcontinent, region, country, city, campaign, keyword, hostname, source, medium, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value
- **Goals By Geography, Weekly Report:** Site, date, continent, subcontinent, region, country, city, campaign, keyword, hostname, source, medium, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value
- **Goals By Content, Daily Report:** Site, date, pagePath, pageTitle, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value, goal2Value, goal3Value, goal4Value
- **Goals By Content, Monthly Report:** Site, year, month, pagePath, pageTitle, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value, goal2Value, goal3Value, goal4Value
- **Goals By Content, Weekly Report:** Site, year, week, pagePath, pageTitle, goalCompletionsAll, goal1Completions, goal2Completions, goal3Completions, goal4Completions, goalValueAll, goal1Value, goal2Value, goal3Value, goal4Value
- **Traffic, Daily Report:** Site, date, bounces, entrances, exits, newVisits, pageviews, timeOnPage, timeOnSite, visitors, visits
- **Traffic, Monthly Report:** Site, year, month, bounces, entrances, exits, newVisits, pageviews, timeOnPage, timeOnSite, visitors, visits

- **Traffic, Weekly Report:** Site, year, week, bounces, entrances, exits, newVisits, pageviews, timeOnPage, timeOnSite, visitors, visits
- **Traffic By Geography, Daily Report:** Site, date, continent, subcontinent, region, country, city, bounces, entrances, exits, newVisits, pageviews, timeOnPage, timeOnSite, visits
- **Traffic By Geography, Monthly Report:** Site, year, month, continent, subcontinent, region, country, city, bounces, entrances, exits, newVisits, pageviews, timeOnPage, timeOnSite, visits
- **Traffic By Geography, Weekly Report:** Site, year, week, continent, subcontinent, region, country, city, bounces, entrances, exits, newVisits, pageviews, timeOnPage, timeOnSite, visits
- **Traffic By Keyword, Daily Report:** Site, year, date, pagePath, pageTitle, source, medium, bounces, entrances, exits, newVisits, pageviews, timeOnPage, visits
- **Traffic By Keyword, Monthly Report:** Site, year, month, pagePath, pageTitle, source, medium, bounces, entrances, exits, newVisits, pageviews, timeOnPage, visits
- **Traffic By Keyword, Weekly Report:** Site, year, week, pagePath, pageTitle, source, medium, bounces, entrances, exits, newVisits, pageviews, timeOnPage, visits
- **Traffic By Content, Daily Report:** Site, date, pagePath, pageTitle, source, medium, bounces, entrances, exits, newVisits, pageviews, timeOnPage, visits
- **Traffic By Content, Monthly Report:** Site, year, month, pagePath, pageTitle, source, medium, bounces, entrances, exits, newVisits, pageviews, timeOnPage, visits
- **Traffic By Content, Weekly Report:** Site, year, week, pagePath, pageTitle, source, medium, bounces, entrances, exits, newVisits, pageviews, timeOnPage, visits
- **Traffic By Navigation, Daily Report:** Site, date, pagePath, landingPagePath, exitPagePath, secondPagePath, medium, uniquePageviews, bounces, entrances, exits, newVisits, pageviews, timeOnPage
- **Traffic By Navigation, Monthly Report:** Site, year, Month, pagePath, landingPagePath, exitPagePath, secondPagePath, medium, uniquePageviews, bounces, entrances, exits, newVisits, pageviews, timeOnPage
- **Traffic By Content, Navigation Report:** Site, year, week, pagePath, landingPagePath, exitPagePath, secondPagePath, medium,

uniquePageviews, bounces, entrances, exits, newVisits, pageviews, timeOnPage

## Standard Tables from Google AdWords

Below is a list of ready-made reports you can easily access using the Sisense Google AdWords provider.

- **Ad Group Performance Report:** AccountDescriptiveName, Date, Campaign, Campaign Id, Ad Group, Ad Group Id, AD Group Status, Avg Position, Clicks, Avg CPC, Cost, Cost/Conv. (1-per-click), Ad Distribution, Impressions, Conversions, (1-per-click), Device, ClickType
- **Ad Group (Report) Daily Report:** AccountDescriptiveName, Date, Campaign, Campaign Id, Ad Group, Ad Group Id, AD Group Status, Avg Position, Clicks, Avg CPC, Cost, Cost/Conv. (1-per-click), Ad Distribution, Impressions, Conversions, (1-per-click), Device, ClickType
- **Ad Group (Report) Monthly Report:** AccountDescriptiveName, Month, Campaign, Campaign Id, Campaign Status, Ad Group, Ad Group Id, AD Group Status, Avg Position, Clicks, Avg CPC, CTR, Avg CPM, Cost, Cost/Conv. (1-per-click), Ad Distribution, Impressions, Conversions, (1-per-click), Device, ClickType
- **Search Query Performance, Daily Report:** AccountDescriptiveName, Date, Search Query, Search Query Match Type, Avg Position, Clicks, Cost, Impressions, Conversions (1-per-click)
- **Search Query Performance, Monthly Report:** AccountDescriptiveName, Month, Search Query, Search Query Match Type, Avg Position, Clicks, Cost, Impressions, Conversions (1-per-click)
- **Keyword, Daily Report:** AccountDescriptiveName, Date, Campaign, Campaign Id, Campaign Status, Ad Group, Ad Group Id, Ad Group Status, Keyword, Avg Position, Clicks, Cost, Impressions, Conversions (1-per-click), Est. First Page Bid
- **Keyword, Monthly Report:** AccountDescriptiveName, Month, Campaign, Campaign Id, Campaign Status, Ad Group, Ad Group Id, Ad Group Status, Keyword, Avg Position, Clicks, Cost, Impressions, Conversions (1-per-click), Est. First Page Bid
- **Geo Performance Report:** AccountDescriptiveName, Campaign, Campaign Id, Country Criteria Id, City, Metro, Region, Campaign Status, Ad Group, Ad Group Id, Ad Group Status, Clicks, Cost, Impressions

# Adding and Customizing Tables and Fields in Google Analytics and AdWords

You can customize both the tables and the fields that you imported from Google Analytics and AdWords. Using the procedure below, you can also import fields from tables that are not automatically imported by Sisense.

## To customize the imported data:

1. Find the appropriate dimension from the Google API: [Google Analytics API](#) / [Google AdWords API](#).
2. **Note:** Use the dimension name and not the XML attribute.

CREATIVE_CONVERSION_REPORT	
AccountCurrencyCode	The currency of the Customer account.
Dimension name	Type String
	Display Name Currency
<del>XML Attribute</del>	<del>currency</del>
	Behavior Attribute
	Filterable True
	Supports Zero Impressions False

Open the

GAnalytics.Fields and GAnalytics.Tables xml files for Google Analytics, or the AdWords.Fields and AdWords.Tables xml files for Google AdWords. These files can usually be found in the following location after the first successful build through the Google Analytics of Google Adwords connector:

- %ProgramData%\Sisense\PrismServer\LocalRepository2.0\Resources\Google\Analytics
- %ProgramData%\Sisense\PrismServer\LocalRepository2.0\Resources\Google\Adwords

3. **Note:** If you are unable to locate the files in the above directories, look under %AppData%\Roaming\Sisense\Prism\LocalRepository2.0\Resources\Google\Analytics/AdWords, or search for the XML files in your root directory. To add a new field from a different table, open the “GAnalytics.Tables.xml” file or “AdWords.Tables.xml” file, and add a new table element that includes the new field.

**Note:** Do not add the field to an existing table element.

In the following example, AccountTimeZoneId is a field from a table that is not imported by default.

```
<table name="Time Zone Table Example"
report_type="Query" aggregation_type="Monthly">
<fields>
<field id="AccountTimeZoneId"></field>
</fields>
</table>
</tables>
```

4. Add the same field to the relevant GAAnalytics.Fields or AdWords.Fields xml file, as follows:

**Note:** Do not add the field to an existing table element.

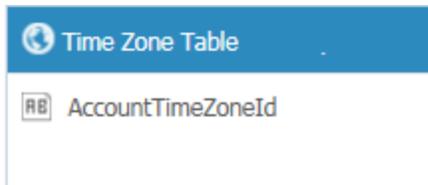
In the following example, AccountTimeZoneId is a field from a table that is not imported by default.

```
<table name="Time Zone Table Example" report_type="Query"
aggregation_type="Monthly">
<fields>
<field id="AccountTimeZoneId"></field>
</fields>
</table>
</tables>
```

Once you save these changes and connect to Google Analytics, you will see the AccountTimeZoneId field in the ‘Time Zone Table’ example.

## Notes

- At least one new table must be added to the Tables.xml file before field changes are made to any standard tables.



- When the **type** attribute is long, the build may fail. If this happens, change the **type** from long to a different type, for example string.

## Adding Custom Dimensions and Values to Google Analytics Tables

To import custom dimensions and values, they must be first defined in the 'Fields' XML, and later included in the relevant reports in the 'Tables' XML.

Open the fields XML from:

```
%ProgramData%\Sisense\PrismServer\LocalRepository2.0\Resources\Google\Analytics
```

Custom Dimension:

```
<field id="ga:dimension1" name="FIELD_NAME"
response_value="ga:dimension1"
```

Custom Metric:

```
<field id="ga:customVarValue1" name="VALUE_NAME"
response_value="ga:customVarValue1"
```

The serial number of the dimension/value must be the same as defined in Google Analytics. The rest of the custom field attributes must be defined according to type; you can use a similar field for reference.

After adding the field, save the XML and add the field to the relevant report in the tables XML.

For more troubleshooting, see our [forum post](#).

# Connecting to DynamicsNav

**This connector is currently in beta and subject to change.**

The Sisense Microsoft Dynamics Nav connector is a standalone connector that allows you to import data from Microsoft Dynamics Nav into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the Microsoft Dynamics Nav tables. To obtain a connection string, you will need to obtain your credentials from Microsoft Dynamics Nav.

Once you have connected to Microsoft Dynamics Nav, you can import a variety of tables from Microsoft Dynamics Nav.

This page describes how to install the Microsoft Dynamics Nav connector, how to connect to Microsoft Dynamics Nav with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the Microsoft Dynamics Nav Connector](#)
- [Connecting to the Microsoft Dynamics Nav](#)
- [Microsoft Dynamics Nav Tables](#)

## Installing the Microsoft Dynamics Nav Connector

Sisense provides the Microsoft Dynamics Nav connector as a standalone connector that you can download and add to your list of default Sisense connectors.

**To install the Microsoft Dynamics Nav connector:**

1. Contact us for the Microsoft Dynamics Nav installation file.
2. Open the installation file and click **Install**.

3. After the installation process is complete, click **Close**.

The Microsoft Dynamics Nav connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.

## Connecting to the Microsoft Dynamics Nav

The connection string used to connect to the Microsoft Dynamics Nav connects to various instances you have in the Microsoft Dynamics Nav. By providing your Microsoft Dynamics Nav credentials as values in the connection string, you can add data from your Microsoft Dynamic Nav tables to the ElastiCube.

The connection string can be set to a series of option=value strings separated by semicolons. If a connection string property value has special characters such as semicolons, single quotes, spaces, etc., then you must quote the value using either single or double quotes.

Connection options are case insensitive.

The connection string to be passed to Dynamics Nav must contain the following parameters:

- **URL:** URL to the DynamicsNAV server organization root. For example, `http://MyServer:7048`.
- **User:** The user who is authenticating to the Dynamics NAV server.
- **Password:** The password used to authenticate to the Dynamics NAV server.
- **ServerInstance:** The instance of the Dynamics NAV server. For example, `DynamicsNAV71`.

In addition, when the **CacheLocation** property is set, the driver will cache metadata to the specified file, ensuring that queries execute more quickly.

CacheLocation specifies the path to the cache when caching to a file.

**Note:** Before the driver can connect with Dynamics NAV, OData services need to be enabled on the server. Once OData Services are enabled, the driver will be able to query any services that are published on the server.

If there is not a Service Default Company for the server, set the Company as well where Company is the company to submit queries against, for example, 'CRONUS Canada, Inc.'. In a multitenant installation, specify the tenant Id in Tenant.

The connection string to connect to Microsoft Dynamics Nav has the following structure:

```
jdbc:dynamicsnav:Property1=Value1;Property2=Value2;
```

The following is an example of a Microsoft Dynamics Nav connection string:

```
jdbc:dynamicsnav:URL=http://myserver:7048;User=myserver\Administrator;Password=admin;ServerInstance=DYNAMICSNAV71;
```

### To add Microsoft Dynamics Nav data:

1. In ElastiCube Manager, click **Add Data** and then, **Microsoft Dynamics Nav**. The Connect to Microsoft Dynamics Nav window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. Microsoft Dynamics Nav is displayed in the Select Database list.
4. Click **OK**. Sisense connects to Microsoft Dynamics Nav and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**.

The tables are displayed in the ElastiCube Manager.

## Microsoft Dynamics Nav Tables

Microsoft Dynamics Nav exposes the following Microsoft Dynamics Nav tables that you can import into the ElastiCube Manager through the Sisense Microsoft Dynamics Nav connector:

### Available Tables

Name	Description
Company	The DynamicsNAV table Company.
Country	The DynamicsNAV table Country.

Name	Description
Currency	The DynamicsNAV table Currency.
Customer	The DynamicsNAV table Customer.
GeneralProductPosting	The DynamicsNAV table GeneralProductPosting.
PurchaseLines	The DynamicsNAV table PurchaseLines.
PurchaseOrder	The DynamicsNAV table PurchaseOrder.
PurchaseOrderPurchLines	The DynamicsNAV table PurchaseOrderPurchLines.
PurchaseOrders	The DynamicsNAV table PurchaseOrders.
ResourceCard	The DynamicsNAV table ResourceCard.
SalesOrder	The DynamicsNAV table SalesOrder.

# Connecting to DynamoDB

The Sisense DynamoDB connector is a standalone connector that allows you to import data from the DynamoDB API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the DynamoDB API. To obtain a connection string, you will need to create a DynamoDB developers account. Once you have connected to DynamoDB, you can import a variety of tables from the DynamoDB API.

This page describes how to install the DynamoDB connector, how to connect to DynamoDB with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the DynamoDB Connector](#)
- [Connecting to the DynamoDB REST API](#)
- [Adding DynamoDB Tables to your ElastiCube](#)
- [DynamoDB Tables](#)
- [Versions](#)

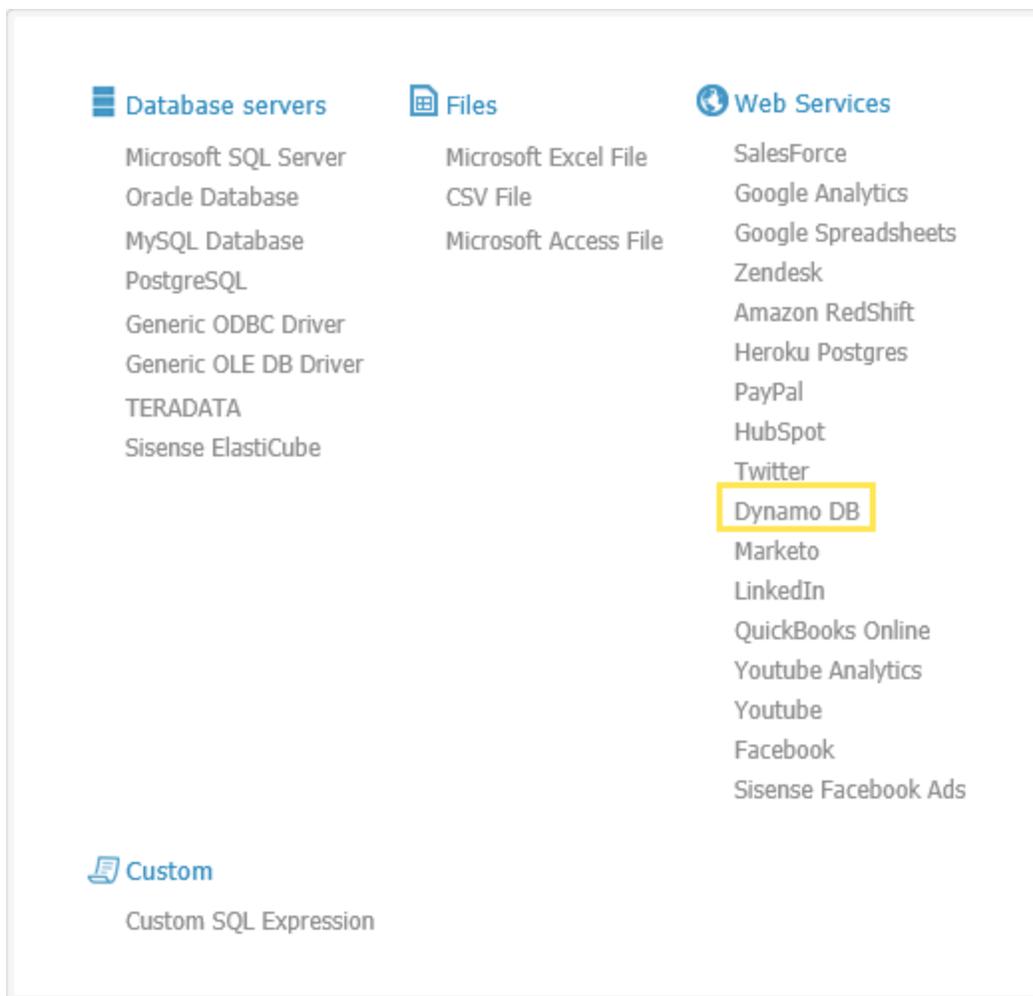
## Installing the DynamoDB Connector

Sisense provides the DynamoDB connector as a standalone connector that you can download and add to your list of default Sisense connectors.

### To install the DynamoDB connector:

1. [Download](#) the DynamoDB installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The DynamoDB connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.



## Connecting to the DynamoDB REST API

The connection to DynamoDB is made using your **AccessKey**, **SecretKey**, and optionally your **Domain** and **Region**. Your AccessKey and SecretKey can be obtained on the security credentials page for your Amazon Web Services account. Your Region will be displayed in the upper left-hand corner when you are logged into DynamoDB.

- **AccessKey:** Your AWS account access key. This value is accessible from your AWS security credentials page.
- **SecretKey:** Your AWS account secret key. This value is accessible from your AWS security credentials page.
- **Domain:** Your AWS domain name. You can optionally choose to associate your domain name with AWS.

- **Region:** The hosting region for your Amazon Web Services. Available values are NORTHERNVIRGINIA, OREGON, NORTHERNCALIFORNIA, IRELAND, SINGAPORE, SYDNEY, TOKYO, and SAOPAULO. The default value is NORTHERNVIRGINIA.

## Adding DynamoDB Tables to your ElastiCube

Sisense uses connection strings to connect to DynamoDB and import data into the ElastiCube Manager.

The connection string to connect to DynamoDB has the following structure:

```
jdbc:DynamoDB:Property1=Value1;Property2=Value2;
```

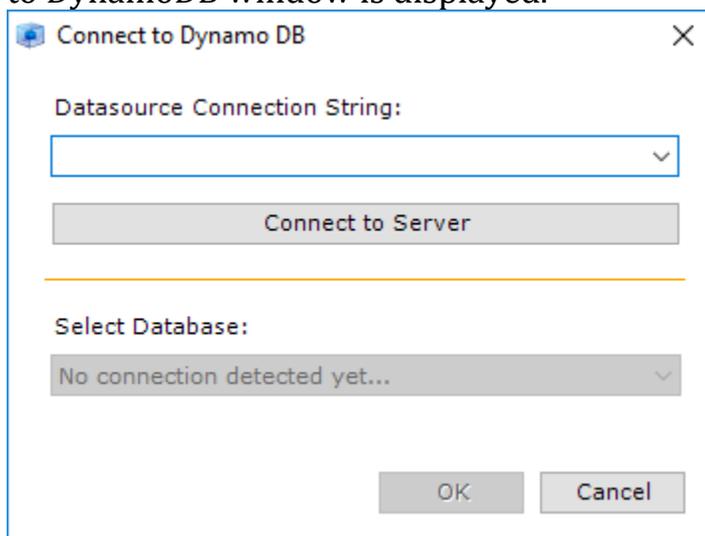
An example of a connection string:

```
jdbc:dynamodb:Access Key=AKIAJH283HDH2932DQ;Secret
Key=vG07bFgSmvfgrdfgHsB6iQU/HVOq09L9g;Domain=amazonaws.c
om;Region=OREGON;
```

**Note:** To switch between accounts, you need to delete the file OAuthsettings.txt file located at ...\Users\xxx\AppData\Roaming\CData\DynamoDB Data Provider.

### To add DynamoDB data:

1. In ElastiCube Manager, click **Add Data** and then, **DynamoDB**. The Connect to DynamoDB window is displayed.



2. In **Datasource Connection String**, enter your connection string.

3. Click **Connect to Server**. DynamoDB is displayed in the Select Database list.
4. Click **OK**. Sisense connects to DynamoDB and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

## Switching between Accounts

When you connect to the DynamoDB data source, Sisense saves your OAuth values in the file OAuthsettings.txt file located at ...\\Users\\xxx\\AppData\\Roaming\\CData\\DynamoDB Data Provider on your Sisense server. To connect to the DynamoDB data source with another user on the same machine, you must delete the OAuthsettings.txt file. Sisense will then generate a new file for that user.

Another option to support multiple users is to define the location and file name of an OAuthsettings file for each unique user in your connection string through the OAuthSettingsLocation parameter. When each user connects to the data source, Sisense generates the OAuth file with the file name you specify in the location you define. In the examples below, two users are allowed to access the DynamoDB data source and for each user, Sisense generates a file that contains that user's OAuth values in the location defined in the string.

```
jdbc:DynamoDB:OAuthSettingsLocation=C:\\DynamoDB\\auth\\john.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
jdbc:DynamoDB:OAuthSettingsLocation=C:\\DynamoDB\\auth\\sally.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
```

In the example above, two OAuth files are created, one for John and one for Sally in the location C:\DynamoDB\auth\.

This is useful if you support many users who each need to access the DynamoDB data source.

## DynamoDB Tables

The Sisense DynamoDB connector allows you to import the following tables into the ElastiCube Manager.

Name	Description
Table0	Test table.
Table1	Test table.
Movies	Sample data file that contains information about a few thousand movies from the Internet Movie Database (IMDb).

### Accumulative Builds

Sisense support accumulative builds for all numeric and dates data types.

However, the data must be sorted before building the ElastiCube.

### Versions

Version	Release Date	Improvements
6.6.0.63	April 20, 2017	Sisense automatically retries a query attempt for certain errors
6.6.0.54	March 21, 2017	TypeDetectionScheme in the connection string works as expected
6.6.0.48	March 3, 2017	Support for Accumulated builds with Sisense V6.6 Beta installed

# Connecting to eBay

This connector is currently in beta and subject to change.

The Sisense eBay connector is a standalone connector that allows you to import data from eBay's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to eBay's API. The connection string contains parameters that define who is trying to access the eBay API.

Once you have connected to eBay, you can import a variety of tables from the eBay API.

This page describes how to install the eBay connector, how to connect to eBay with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the eBay Connector](#)
- [Connecting to the eBay REST API](#)
- [Adding eBay Tables to your ElastiCube](#)
- [eBay Tables](#)

## Installing the eBay Connector

Sisense provides the eBay connector as a standalone connector that you can download and add to your list of default Sisense connectors.

### To install the eBay connector:

1. Contact us for the eBay installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The eBay connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.

# Connecting to the eBay REST API

There are two authentication methods available for connecting to eBay: Using a token generated for accessing only personal data or using tokens that allow other users to access their own data. Both methods require you to join the eBay Developers Program.

## Register and Get Your Keys

After joining the eBay Developers Program, you can obtain the following values from the Application Keys page.

1. AppID
2. DevID
3. CertID

[Home](#) > [Application Keys](#)

## Application Keys

Sandbox ⓘ

SisenseTest ⓘ

[Request another keyset](#)

App ID (Client ID)	Gidibehs-SisenseT-SBX-5c6c7072b	<a href="#">User Tokens</a>   <a href="#">Notifications</a>   <a href="#">API Reports</a>
Dev ID	4812747c-8658-4fb2-80d6-193b4	
Cert ID (Client Secret)	SBX-c6c7072b9a9c-ecfc-464a-96cc	<a href="#">Rotate (Reset) Cert ID</a>

If you're working with your own account, follow the steps below to generate a token:

1. Go to your Developer Account Dashboard and open the **Application Keys** page.
2. In the App Id row for the key set you want to use, click **User Tokens**. Use production keys if you're searching for real items or listing real items to the

live eBay site. Use sandbox keys if you're creating test listings in the sandbox.

3. Click **Get a User Token Here** and log in with credentials for either your sandbox or production eBay account.

The screenshot shows the 'User Tokens (eBay Sign-in)' section of the eBay Developer Tools interface. At the top, there are radio buttons for 'Environment' with 'Sandbox' selected and 'Production' unselected. Below this is a dropdown menu showing 'SisenseTest' and the 'AppID/ClientID: Gidibehs-SisenseT-SBX-5c6c7072b-53800d66' with links for 'Get OAuth Application Token' and 'OAuth Scopes'. The main content area has three tabs: 'User Tokens (eBay Sign-in)', 'Alerts & Notifications', and 'Reports'. Under the 'User Tokens (eBay Sign-in)' tab, there is a blue link 'Get a User Token Here' with a dropdown arrow. Below this is the text 'Get a token for an eBay Sandbox user. Learn more >'. There are two radio buttons: 'Auth'n/Auth' (selected) and 'OAuth (new security)'. A blue button 'Sign in to Sandbox' is visible, along with a link 'Revoke a Token'. At the bottom, there is a link 'Register a new Sandbox user'.

4. Click **Get a Token from eBay via Your Application** and enter a title for your application and the relevant callback and return URLs. From this page you can also retrieve your RuName, which is a mandatory parameter that needs to be included in the connection string.

> [Get a User Token Here](#)

## ▼ [Get a Token from eBay via Your Application](#)

How does this work? [Learn more >](#)

Your eBay Sign-in Settings

[+ Add eBay Redirect URL](#) ⓘ

RuName (eBay Redirect URL name)	Display Title	OAuth Enabled
Gidi_behsdfasd-Gidibehs-Sisens-phgbgxit		✓
Your branded eBay Sandbox Sign In (Auth'n'auth)	https://signin.sandbox.ebay.com/ws/eBayISAPI.dll?SignIn&runame=Gidi_behsdfasd-Gidibehs-Sisens-phgbgxit&sessID=<SESSION_ID>	
Your branded eBay Sandbox Sign In (OAuth)	https://signin.sandbox.ebay.com/authorize?client_id=Gidibehs-SisenseT-SBX-5c6... <a href="#">See all</a>	
Display Title	<input type="text" value="Ebay Connector"/>	
	URLs must be secure (https://)	
Your privacy policy URL	<input type="text" value="https://"/>	
Your auth accepted URL <sup>1</sup>	<input type="text" value="https://localhost:33333"/>	
Your auth declined URL <sup>1</sup>	<input type="text" value="https://"/>	
<sup>1</sup> If you leave these fields blank, the eBay default <a href="#">accept</a> and <a href="#">reject</a> pages will be used.		
<a href="#">Delete</a>	<a href="#">Clone</a>	<input checked="" type="radio"/> Auth'n'Auth <input type="radio"/> OAuth
		<input type="button" value="Revert"/> <input type="button" value="Save"/>
You have to save your new settings to test them.		

5. Click **Save**.

After setting the following connection properties, you are ready to connect.

## Adding eBay Tables to Your ElastiCube

After you have retrieve your API Key or OAuth credentials from eBay, you provide the relevant information in a connection string. Sisense uses connection strings to connect to eBay and import data into the ElastiCube Manager. Each connection string contains a authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

The connection string to eBay is in the following format:

```
jdbc:ebay:MyDevID=xxxxxxxxxxxxxxxxxxxxxxxxxxxx;
```

For OAuth, the connection string to eBay is in the following format:

```
jdbc:ebay:DevID=MyDevID;AppID=MyAppID;CertID=MyCertID;RuName=MyRuName;SiteID=0;OAuthAccessToken=MyOAuthAccessToken
```

- **DevID:** Set this to the Developer ID from your key set.
- **AppID:** Set this to the Application ID from your key set.
- **CertID:** Set this to the Certificate ID from your key set.
- **RuName:** Set this to the Redirect URL name on the Token from eBay via your Application section of the User Tokens page.
- **SiteID:** Set this to the eBay site ID. By default this is 0 (the U.S. site). See Site Mappings, for more options.
- **CallbackURL:** Set this to <http://localhost:33333> or another port of your choice.
- **InitiateOAuth:** Set this to GETANDREFRESH.
- **Sandbox:** If you are using the sandbox environment add `UseSandbox=true`.

#### To add eBay data:

1. In ElastiCube Manager, click **Add Data** and then, **eBay**. The Connect to eBay window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. eBay is displayed in the Select Database list.
4. Click **OK**. Sisense connects to eBay and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

## eBay Tables

The Sisense eBay connector allows you to import the following tables into the ElastiCube Manager.

#### Available Tables

Name	Description
ItemListing	Update, delete, and query Item Listings for sellers.
ItemVariations	Item Variations are multiple similar (but not identical)

Name	Description
	<p>items in a single fixed-price listing. For example, a T-shirt listing could contain multiple items of the same brand that vary by color and size; each variation would specify a combination of one of these colors and sizes. Each variation can have a different quantity and price. You can buy multiple items from one variation at the same time. (That is, one order line item can contain multiple items from a single variation.)</p>

#### Available Views

Name	Parameter
Bidders	Query eBay users that bid on a specified listing, regardless of whether the listing has ended.
Categories	Query eBay Categories.
CategoryFeatures	Query the settings supported by eBay Categories.
Countries	Lists the country code and associated name of the countries supported by the eBay system, regardless of the site specified in the request.
Currencies	Query the supported Currencies in eBay.
DSRSummary	Query rating information from DSR Summary Reports.
HighestBidder	Query the Highest Bidder of the eBay users who bid on a specified listing, regardless of whether the listing has ended.

Name	Parameter
ItemShipping	Query details pertinent to one or more items for which calculated shipping has been offered by the seller, such as package dimension, weight, and packaging/handling costs.
ItemTransactions	Query order line item (transaction) information for a specified listing. For auction listings, order line items are not created until the auction ends with a winning bidder, or if the Buy It Now feature is enabled for the listing and used by the buyer to purchase the item. For fixed-price listings, order line items are created after a buyer has committed to purchase the item(s).
MostWatchedItems	Query items with the highest watch counts for the entire site or for a specific category. The top items are determined by the Watch Count totals from the preceding day. Ranking of the Most Watched Items is calculated with the latest Watch Count information.
ShippingCarriers	Query Shipping Carriers supported by the specified site.
ShippingCategories	A shipping service category supported for the site.
ShippingLocations	Lists the regions and locations supported by eBay's shipping services. Returns all shipping locations supported by eBay, regardless of the site specified in the request.
ShippingServices	Query Shipping Services supported by the specified eBay site.

<b>Name</b>	<b>Parameter</b>
Sites	Lists all available eBay sites and their associated SiteID numbers.
TopSellingProducts	Query your eBay Top Selling Products.

# Connecting to Exact Online

The Sisense Exact connector is a standalone connector that allows you to import data from Exact's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the Exact APIs. To obtain a connection string, you will need to create a Exact app.

Once you have connected to Exact, you can import a variety of tables from the Exact API.

This page describes how to install the Exact connector, how to connect to Exact with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the Exact Connector](#)
- [Adding Exact Tables to your ElastiCube](#)
- [Exact Tables](#)

## Installing the Exact Connector

Sisense provides the Exact connector as a standalone connector that you can download and add to your list of default Sisense connectors.

### To install the Exact connector:

1. [Download](#) the Exact installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The Exact connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.

## Adding Exact Tables to your ElastiCube

Sisense uses connection strings to connect to Exact and import data into the ElastiCube Manager. The default connection string for Exact is:

```
jdbc:exactonline:InitiateOAuth=GETANDREFRESH;Region='United States';Division=5512;
```

After you connect using the default connection string, Sisense redirects you to Exact's Login page where you log in and your account is authenticated. Once the account is authenticated, the relevant tables can be added to your ElastiCube. If you want to customize your connection string, you can by creating an Exact app and passing additional connection string parameters. For more information, see [Registering an App](#) at the end of this page.

### To add Exact data:

1. In ElastiCube Manager, click **Add Data** and then, **Exact**. The Connect to Exact window is displayed.
2. In **Datasource Connection String**, enter the connection string:  
`jdbc:exactonline:InitiateOAuth=GETANDREFRESH;Region='United States';Division=5512;`
3. Click **Connect to Server**. You are redirected to the Exact Login page.
4. Log in to your Exact account. Exact is displayed in the Select Database list in the ElastiCube Manager.
5. Click **OK**. Sisense connects to Exact and displays a list of tables available for you to import.
6. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

## Exact Tables

Exact's RESTful APIs expose the following Exact tables that you can import into the ElastiCube Manager through the Sisense Exact connector:

### Available Tables

- AcceptQuotation

- AccountInvolvedAccounts
- AccountOwners
- Accounts
- Addresses
- BankAccounts
- BankEntries
- BankEntryLines
- CashEntries
- CashEntryLines
- CommunicationNotes
- Complaints
- Contacts
- Costcenters
- CostTransactions
- Costunits
- DepreciationMethods
- DirectDebitMandates
- DocumentAttachments
- DocumentFolders
- DocumentTypeFolders
- Events
- ExchangeRates
- GeneralJournalEntries
- GeneralJournalEntryLines
- GLAccounts
- GoodsDeliveries
- GoodsDeliveryLines
- InvoiceSalesOrders
- InvoiceTerms
- InvolvedUserRoles
- InvolvedUsers
- Items
- ItemWarehouses
- Journals
- Mailboxes
- MailMessageAttachments
- MailMessages
- OperationResources
- Operations

- Opportunities
- PaymentConditions
- PrintedSalesInvoices
- PrintedSalesOrders
- PrintQuotation
- ProductionAreas
- ProjectHourBudgets
- ProjectRestrictionEmployees
- ProjectRestrictionItems
- ProjectRestrictionRebillings
- Projects
- PurchaseEntries
- PurchaseEntryLines
- QuotationLines
- Quotations
- RejectQuotation
- ReopenQuotation
- ReviewQuotation
- SalesEntries
- SalesEntryLines
- SalesInvoiceLines
- SalesInvoices
- SalesItemPrices
- SalesOrderID
- SalesOrderLines
- SalesOrders
- ServiceRequests
- ShopOrderMaterialPlans
- ShopOrderRoutingStepPlans
- ShopOrders
- SolutionLinks
- StockCountLines
- StockCounts
- SubscriptionLines
- SubscriptionRestrictionEmployees
- SubscriptionRestrictionItems
- Subscriptions
- Tasks
- TimeCorrections

- TimeTransactions
- VATCodes
- Warehouses
- Workcenters

### **Available Views**

- AccountantInfo
- AccountClasses
- AccountClassificationNames
- AccountClassifications
- ActiveEmployments
- AddressStates
- AgingOverview
- AgingPayablesList
- AgingReceivablesList
- Assets
- AvailableFeatures
- Banks
- BatchNumbers
- Budgets
- Currencies
- DefaultMailbox
- Departments
- Divisions
- DocumentCategories
- Documents
- DocumentsAttachments
- DocumentTypeCategories
- DocumentTypes
- Employees
- EmploymentContractFlexPhases
- EmploymentContracts
- EmploymentEndReasons
- EmploymentOrganizations
- Employments
- EmploymentSalaries
- FinancialPeriods
- GLClassifications
- GLSchemes

- GLTransactionTypes
- HourCostTypes
- ItemGroups
- ItemVersions
- JobGroups
- JobTitles
- JournalStatusList
- Layouts
- OpportunityContacts
- OutstandingInvoicesOverview
- PayablesList
- PreferredMailbox
- PriceLists
- ProfitLossOverview
- ProjectBudgetTypes
- PurchaseOrderLines
- PurchaseOrders
- ReasonCodes
- ReceivablesList
- RecentCosts
- RecentHours
- ReportingBalance
- Returns
- RevenueList
- SalesPriceListDetails
- Schedules
- SerialNumbers
- ShippingMethods
- StockBatchNumbers
- StockSerialNumbers
- StorageLocations
- SubscriptionLineTypes
- SubscriptionReasonCodes
- SubscriptionTypes
- TaxComponentRates
- TaxEmploymentEndFlexCodes
- TaxScheduleComponents
- TaxSchedules
- TimeAndBillingAccountDetails

- TimeAndBillingActivitiesAndExpenses
- TimeAndBillingEntryAccounts
- TimeAndBillingEntryProjects
- TimeAndBillingEntryRecentAccounts
- TimeAndBillingEntryRecentActivitiesAndExpenses
- TimeAndBillingEntryRecentHourCostTypes
- TimeAndBillingEntryRecentProjects
- TimeAndBillingItemDetails
- TimeAndBillingProjectDetails
- TransactionLines
- Transactions
- Units
- UserRoles
- UserRolesPerDivision
- Users
- VatPercentages

## Registering an App

To access Exact's REST API from the ElastiCube Manager, you must provide valid OAuth Exact credentials through a [connection string](#). These credentials are provided by Exact when you register an application.

After you receive your credentials from Exact, you can create the [connection string](#) and provide Sisense with it to connect to your data. Each connection string contains a authentication parameters that the data source uses to verify your identity and what information you can export to Sisense. To learn more, see [Connection String Parameters](#).

Follow the steps below to obtain the OAuth client credentials, the OAuthClientId and OAuthClientSecret:

1. Create an Exact Online developer account.
2. Log into the App Center and click **Manage Apps > Add a New Application**.
3. Enter the app name to be displayed to users when they are prompted to grant permissions to your app.

4. Enter the Redirect URI. If you are making a desktop application, set the Redirect URI to `http://localhost:portnumber`. If you are making a Web application, set the Redirect URI to a page you would like the user to be returned to after they have granted your application permissions.
5. Click **Edit** for your app. The client credentials, the client Id and client secret, are displayed.

# Connecting to Excel

The ElastiCube Manager enables easy and quick access to tables contained within Excel spreadsheets.

There are two options for uploading data, the first is to upload your file to the Sisense Server. Once the file is uploaded, the data will be imported into the ElastiCube as it was at the time the file was uploaded.

The second option is to define the location of your files on the Sisense Server. This is the preferred option if your Excel files frequently change, but remain in the same location on the Sisense Server as the data is taken from the Excel files each time the ElastiCube is built.

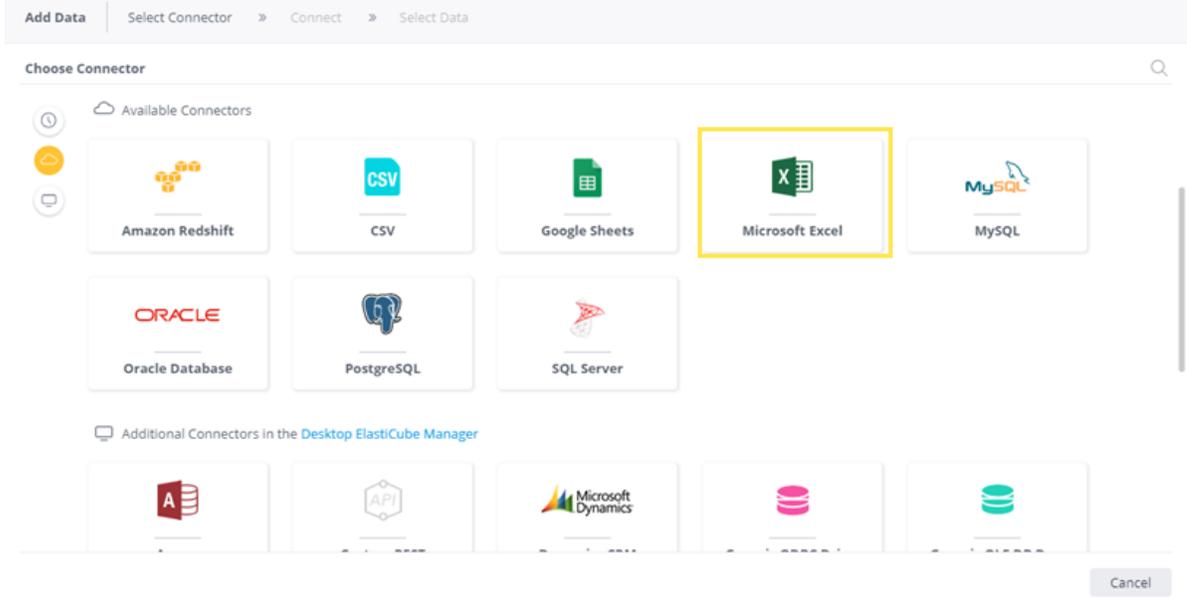
When you import multiple files from a folder and build your ElastiCube, the data is combined together in a single table. To build multiple tables with separate tables for each file, you must repeat the process for each file you want to import as a table. If you have any questions about data accumulation between builds, please contact Sisense Support.

**Note:** Click [here](#) to see tips on how to prepare your Excel files before adding them to your schema in the ElastiCube.

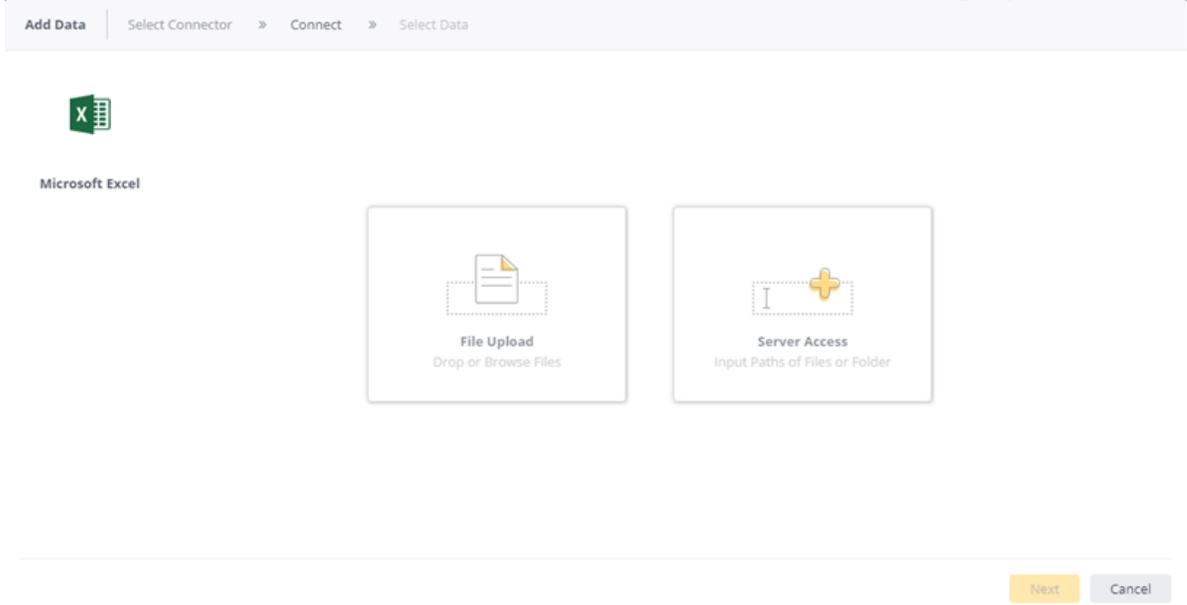
## **To import data from your Excel files:**

1. In the Data page of ElastiCube Manager, open an ElastiCube or create a new ElastiCube.

2. In the ElastiCube, click **+ Data**. The Add Data dialog box is displayed.



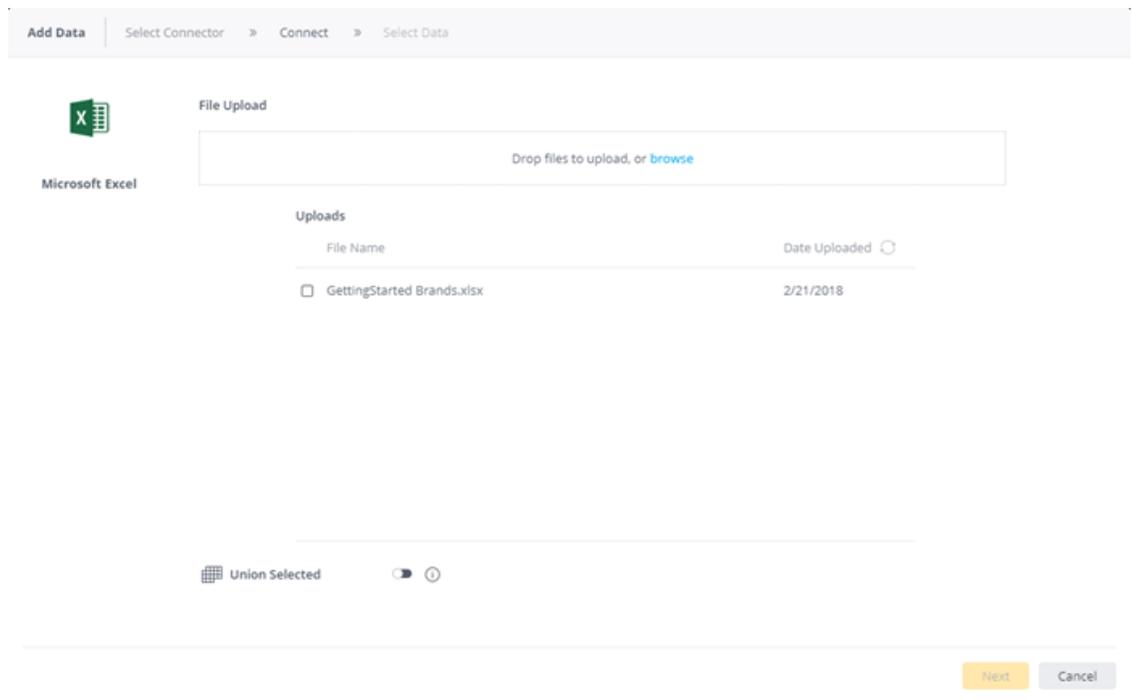
3. Click **Microsoft Excel**. The Microsoft Excel Connect area is displayed.



4. Select the relevant option for importing your data:

### **File Upload**

1. Select this option to import your data from your Excel file. If the file is updated later, you will need to upload it again. To upload the file, click **Browse** and navigate to the file to be uploaded or drag the Excel file to the File Upload area.



The uploaded file is added to the Uploads list.

2. Select the relevant Excel file(s) to be uploaded.
3. (Optional) Toggle the **Union Selected** switch to append several Excel files together when the data is imported into the ElastiCube. The data in the files must have the same column and data types. After you have finished uploading the relevant Excel files, select them from the upload list and click **Next**. Continue to Step 5.

### Server Access

1. Select this option to define the location of your files on the Sisense Server if your Excel files frequently change, but remain in the same location on the Sisense Server as the data is taken from the Excel files each time the ElastiCube is built.
2. Select **Input Folder Path** and enter the full file path with the file name where your Excel files are located. This will display each Excel file in the folder in the next screen where you select what tables to add to the ElastiCube.

OR

Select **Input File Path** and enter the full file path with the file name and its extension of your Excel file. For example, C:\Example.xlsx. This file display all the tables in the Excel file on the next screen where you select what tables to add to the ElastiCube.

3. (Optional) Toggle the Union Selected switch to append several Excel files together when the data is imported into the ElastiCube. The data in the files must have the same column and data types.
4. After you have finished defining the locations of your Excel files, click **Next** and continue to Step 5. A list of Excel files in the directory are displayed.
5. From the **Select Table** list, click  to select preview the columns in the Excel file and display the **Settings**, which provides more options for customizing your data.
6. In the Settings area, define the following settings:
  - **Culture**: Select the culture for your Excel. This defines settings such as the format of the date and time or delimiter (decimal or comma) used in your Excel file.
  - **Fields in First Row**: Enables you to specify table column names based on the header in the first row of the spreadsheet.
  - **Static Range**: Enter two cells, each with a leading \$ sign and a colon as a delimiter. Press **Enter** to preview the selection. Selecting Static Range option enables you to select a specific range of data in the sheet. Data needs to be in a table structure, starting at the top left cell of the range, with field names as the first row.  
For a static range between cells A1 and E10, type in \$A1:\$E10.
7. After you have selected all the relevant tables, click **Done**. The tables are added to your schema in the ElastiCube Manager.

# Connecting to Microsoft Excel

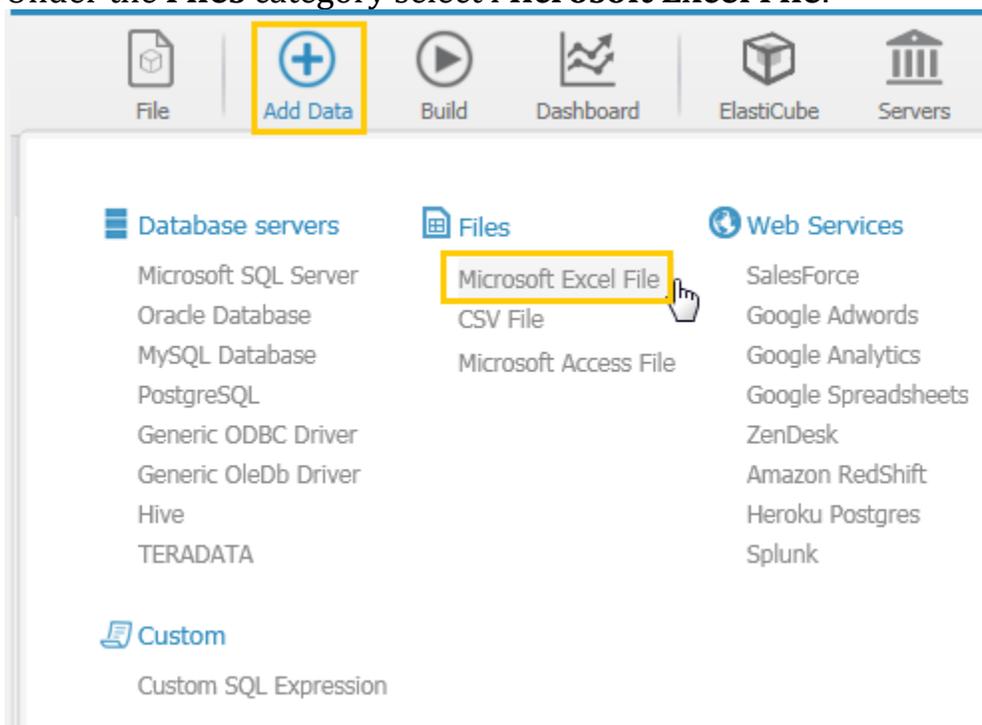
The ElastiCube Manager enables easy and quick access to tables contained within Excel spreadsheets.

You can add multiple Excel files to your ElastiCube. A typical use case for importing multiple files is when you want to split a large file into multiple smaller files.

**Note:** [Click here](#) to see tips on how to prepare your Excel files before adding them to your schema in the ElastiCube.

## To connect to Microsoft Excel:

1. Click **Add Data** in the top menu of the ElastiCube Manager.
2. Under the **Files** category select **Microsoft Excel File**.



3. In the window that opens, locate the Excel file.
4. Select the Excel file and click **OK**.
5. In the window that opens, you have the following options:

- **Fetch cells from:** You can fetch cells from either a sheet or a named range (provided you have named ranges in your file).
- **Available Sheets:** Select a sheet contained in the spreadsheet.
- **Range:** Select a Static Range or Take range from whole sheet.

For a Static Range, enter two cells, each with a leading \$ sign and a colon as a delimiter. Hit Enter to preview the selection. Selecting **Static Range** option enables you to select a specific range of data in the sheet. Data needs to be in a table structure, starting at the top left cell of the range, with field names as the first row.

For a static range between cells A1 and E10, type in \$A1:\$E10

Date	Total	Campaign	CampaignType	Region	AdGroup	Keyword
7/1/2012	Total	Europe LIVE - Search	Live	EMEA	all music	blues songs
7/1/2012	Total	Europe LIVE - Search	Live	EMEA	all music	download jazz
7/1/2012	Total	Europe LIVE - Search	Live	EMEA	all music	free classical music online
7/1/2012	Total	Europe LIVE - Search	Live	EMEA	all music	online classical music free
7/1/2012	Total	Europe LIVE - Search	Live	EMEA	bing music	famous jazz musicians
7/1/2012	Total	Europe LIVE - Search	Live	EMEA	bing music	free blues radio
7/1/2012	Total	Band - Search	Test	Unknown	cheap music downloads	stl blues
7/1/2012	Total	MUSIC	Test	EMEA	classical music online	what is hip hop
7/1/2012	Total	MUSIC	Test	EMEA	country music	most famous classical music
7/1/2012	Total	MUSIC	Test	EMEA	country music videos	blues guitar chords
7/1/2012	Total	MUSIC	Test	EMEA	download free music	classic music
7/1/2012	Total	MUSIC	Test	EMEA	download free music	download classical music
7/1/2012	Total	MUSIC	Test	EMEA	download mp3 music	psychedelic trance music
7/1/2012	Total	MUSIC	Test	EMEA	bing music	famous jazz musicians
7/1/2012	Total	MUSIC	Test	EMEA	bing music	house music mix
7/1/2012	Total	MUSIC	Test	EMEA	download music	good classical music
7/1/2012	Total	MUSIC	Test	EMEA	downloadable music	hip hop music
7/1/2012	Total	MUSIC	Test	EMEA	downloading free music	hip hop new releases
7/1/2012	Total	MUSIC	Test	EMEA	downloading free music	hip hop news and gossip
7/1/2012	Total	MUSIC	Test	EMEA	free country music	hip hop songs free download

Selecting **whole sheet** will import all data within the sheet. Data within the sheet needs to be in a table structure starting at the top left corner (cell A1) with the field names as the first row.

	A	B	C	D	E	F	G
1	Date	Total	Campaign	CampaignType	Region	AdGroup	Keyword
2	7/1/2012	Total	Europe LIVE - Search	Live	EMEA	all music	blues songs
3	7/1/2012	Total	Europe LIVE - Search	Live	EMEA	all music	download jazz
4	7/1/2012	Total	Europe LIVE - Search	Live	EMEA	all music	free classical music online
5	7/1/2012	Total	Europe LIVE - Search	Live	EMEA	all music	online classical music free
6	7/1/2012	Total	Europe LIVE - Search	Live	EMEA	bing music	famous jazz musicians
7	7/1/2012	Total	Europe LIVE - Search	Live	EMEA	bing music	free blues radio
8	7/1/2012	Total	Band - Search	Test	Unknown	cheap music downloads	stl blues
9	7/1/2012	Total	MUSIC	Test	EMEA	classical music online	what is hip hop
10	7/1/2012	Total	MUSIC	Test	EMEA	country music	most famous classical music
11	7/1/2012	Total	MUSIC	Test	EMEA	country music videos	blues guitar chords
12	7/1/2012	Total	MUSIC	Test	EMEA	download free music	classic music
13	7/1/2012	Total	MUSIC	Test	EMEA	download free music	download classical music
14	7/1/2012	Total	MUSIC	Test	EMEA	download mp3 music	psychedelic trance music
15	7/1/2012	Total	MUSIC	Test	EMEA	bing music	famous jazz musicians
16	7/1/2012	Total	MUSIC	Test	EMEA	bing music	house music mix
17	7/1/2012	Total	MUSIC	Test	EMEA	download music	good classical music
18	7/1/2012	Total	MUSIC	Test	EMEA	downloadable music	hip hop music
19	7/1/2012	Total	MUSIC	Test	EMEA	downloading free music	hip hop new releases
20	7/1/2012	Total	MUSIC	Test	EMEA	downloading free music	hip hop news and gossip

- **Preview:** Shows a preview of data contained in the sheet.
- **First Row Contains Field Names:** Enables you to specify table column names based on the header in the first row of the spreadsheet.
- Click **Select multiple files...** to import multiple files.

Added Table - Data Properties

Fetch cells from: Sheet

Available Sheets: Sheet1

Static Range:

First Row Contains Field Names

Take range from whole sheet

Preview: [Show All](#) [Hide All](#) Preview Fields: Preview All

Year	Category	Revenue
1/1/2013 12:00:00 AM	Electronics	10
1/2/2013 12:00:00 AM	Books	10
1/3/2013 12:00:00 AM	Furniture	20
1/4/2013 12:00:00 AM	Furniture	20
1/5/2013 12:00:00 AM	Furniture	20
1/6/2013 12:00:00 AM	Electronics	30

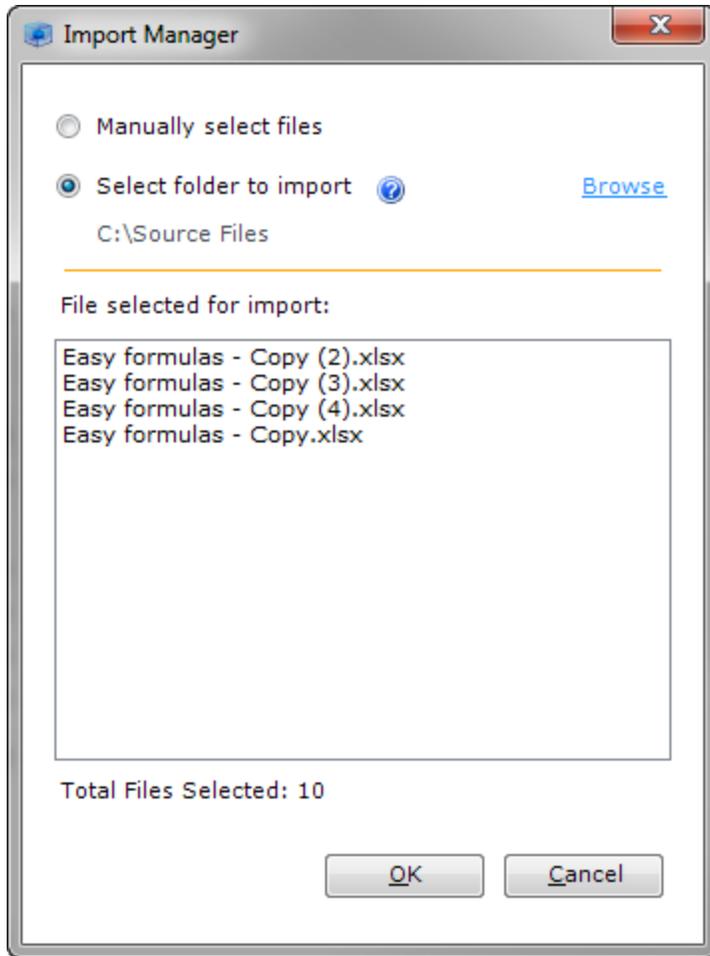
Culture: English (United States)

[Select multiple files...](#)

Add Cancel

**Note:** Make sure that all the files adhere to the same structure, including the presence or absence of headers in the first row. Files will be added to the table based on their file names in alphabetical order.

Select the **Select folder to import** option. All the Excel files in the folder are added to the selected files list.



Click **OK**.

6. Click **Add**.

# Connecting to Facebook Ads

The Sisense Facebook Ads connector is a standalone connector that allows you to import data quickly from your Facebook Ads account into the ElastiCube Manager without custom SQL. This connector includes several features to enable you to quickly and easily import data from Facebook Ads:

- Date range selector from within the ElastiCube Manager
- Support for accumulative builds
- Retry mechanism to improve build success

After you have downloaded and installed the connector, you can connect to Facebook Ads through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to Facebook's APIs.

This page describes how to install the Facebook Ads connector, how to connect to Facebook Ads with a connection string, and what tables you can import into the ElastiCube Manager.

- Prerequisites
- Installing the Facebook Ads Connector
- Connecting to Facebook Ads
- Accessing the Facebook API
- Creating a Connection String
- Adding Facebook Ads Tables to your ElastiCube
- Facebook Ads Tables
- Versions

## **Breaking Changes for the Marketing API Version 2.8**

As of July 25, 2017, Facebook no longer supports their Marketing API version 2.8. Currently, all users whose Facebook app is v2.8 will receive an error. To overcome this issue, add the parameter 'Version=2.10' to the connection string.

For more information about the latest version of the Marketing API, see Facebook's [blog posts](#).

## Prerequisites

- The Facebook Ads connector requires Sisense V6.6.1 and later.

## Installing the Facebook Ads Connector

Sisense provides the Facebook Ads connector as a standalone connector that you can download and add to your list of default Sisense connectors. After you install the Facebook Ads connector, it is added to your list of Web Service connectors.

### To install the Facebook Ads connector:

1. [Download](#) the Facebook Ads installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The Facebook Ads connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.



# Connecting to the Facebook Ads

To access Facebook Ads' data from Sisense, you must provide valid Facebook Ads credentials in a connection string.

This section describes how to create a Facebook app to retrieve the necessary authentication parameters for the connections string and then how to construct the connection string.

**Note:** If you have already created a connection string for the Facebook connector, you can use the same credentials in the Facebook Ads connection string.

## Creating an App

The credentials you need to connect to your data in Facebook are provided by Facebook when you create a Facebook app through the Facebook Developer site. For more information, see [Creating an App in Facebook](#).

## Creating a Connection String

Sisense uses connection strings to connect to Facebook Ads and import data into the ElastiCube Manager. Each connection string contains authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

**Note:** When switching accounts, you must sign out of your account and then sign it with the new account.

After you have obtained the relevant credentials, you can create the connection used to connect to your Facebook account. The Facebook connection string has the following structure:

```
jdbc:DataSourceName:Property1=Value1;Property2=Value2;
```

The following is an example of Facebook connection string:

```
jdbc:facebook:OAuthClientId=112345383774486;OAuthClientSecret=064c524478d712534b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;Timeout=0;ThrottleSeconds=15;ThrottleMaxPercent=90;Other=RetryWaitTime=60000;MaximumRequestRetries=5;
```

The example above includes both mandatory and optional parameters you can provide in the connection. The required parameters are emphasized in bold.

### **Mandatory Parameters**

- **OAuthClientId**: Set this to the App ID in your app settings.
- **OAuthClientSecret**: Set this to the App Secret in your app settings.
- **CallbackURL**: Set this to the Return URL in your app settings.
- **InitiateOAuth**: Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken connection property.

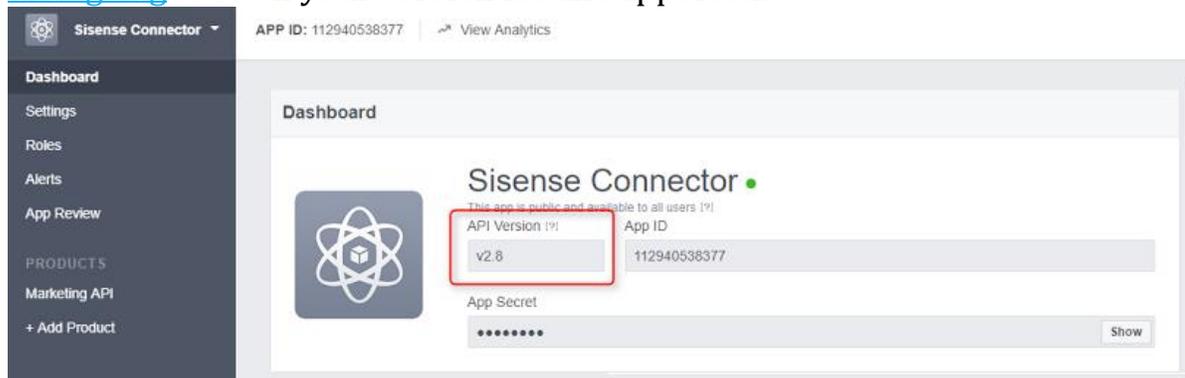
### **Optional Parameters**

- **Aggregate Format**: The format aggregate or collection columns should return in.
- **Authenticate As Page**: The name or Id of a page to authenticate as when making requests to Facebook.
- **Auto Cache**: The AutoCache property controls automatic caching of data.
- **Cache Connection**: The connection string for the cache database.
- **Cache Driver**: The database driver to be used to cache data.
- **Cache Location**: Specifies the path to the cache when caching to a file.
- **Cache Metadata**: Whether or not to cache the table metadata. The cache will be stored in memory unless CacheConnection or CacheLocation are set.
- **Callback URL**: The OAuth callback URL to return to when authenticating. This value must match the callback URL you specify in your app settings.
- **Firewall Password**: A password, if authentication is required to connect through a firewall.
- **Firewall Port**: The TCP port for the firewall FirewallServer — see the description of the FirewallServer option for details.
- **Firewall Server**: Specify a firewall name or IP address to authenticate requested connections, if necessary.
- **Firewall Type**: The type of firewall to connect through.

- **Firewall User:** The user name to authenticate with the firewall.
- **Initiate OAuth:** Set this property to initiate the process to obtain or refresh the OAuth access token when you connect.
- **Location:** A path to the directory that contains the schema files defining tables, views, and stored procedures.
- **Logfile:** A path to the log file.
- **Max Log File Size:** A string specifying the maximum size in bytes for a log file (ex: 10MB). When the limit is hit, a new log is created in the same folder with the date and time appended to the end.
- **MaximumRequestRetries:** The maximum amount of request retries.
- **OAuth Access Token:** The access token for connecting using OAuth.
- **OAuth Client Id:** The client Id assigned when you register your application with an OAuth authorization server.
- **OAuth Client Secret:** The client secret assigned when you register your application with an OAuth authorization server.
- **OAuth Settings Location:** The location of the settings file where OAuth values are saved when InitiateOAuth is set to GETANDREFRESH or REFRESH.
- **Offline:** Use offline mode to get the data from the cache instead of the live source.
- **Other:** The other parameters necessary to connect to a data source, such as username and password, when applicable.
- **Pagesize:** The maximum number of results to return per page from Facebook.
- **Proxy Auth:** Scheme The proxy server authorization scheme (default: BASIC).
- **Proxy Auto Detect:** This indicates whether to use the default system proxy settings or not. Set ProxyAutoDetect to FALSE to use custom proxy settings. This takes precedence over other proxy settings.
- **Proxy Password:** A password, if authenticating with a proxy server.
- **Proxy Port:** The TCP port for the proxy ProxyServer (default: 80).
- **Proxy Server:** If a proxy server is given, then the HTTP request is sent to the proxy instead of the specified server.
- **Proxy SSL Type:** The SSL type to use when connecting to the proxy server (default: AUTO).
- **Proxy User:** A user name, if authentication is to be used for the proxy.
- **Pseudo Columns:** Indicates whether or not to include pseudo columns as columns to the table.

- **Query Cache:** The time duration, in seconds, for which the in-memory cached data is reused.
- **ReadOnly:** You can use this property to enforce read-only access to Facebook from the provider.
- **RetryWaitTime:** The amount of time in milliseconds to wait before retrying to query the Facebook API.
- **RTK:** The runtime key used for licensing.
- **Search Terms:** Default SearchTerms if none are specified. Used for some tables, such as Users, where SearchTerms may be specified.
- **SSL Server Cert:** The certificate to be accepted from the server when connecting using TLS/SSL.
- **Support Enhanced SQL:** If set to true, the provider will cache the entire table in memory and issue the current query against the memory cache. Allows queries with aggregate functions that are otherwise unsupported. The entire memory cache is discarded after the query is executed, unlike QueryCache.
- **Tables:** Comma-separated list of tables to be listed. For example: Tables=TableA,TableB,TableC.
- **Target:** A default target if none is specified. Used for some tables, such as Comments, where a target may be specified.
- **ThrottleSeconds:** Indicates how many seconds to wait until sending another insight request when the ThrottleMaxPercent has been met.
- **ThrottleMaxPercent:** The threshold set for throttling further insight requests. The defaults for ThrottleSeconds and ThrottleMaxPercent are 10 and 95.
- **Timeout:** The value in seconds until the timeout error is thrown, canceling the operation.
- **Verbosity:** The verbosity level that determines the amount of detail included in the log file.
- **Version:** The Facebook Graph API version to use. You can see your version in the Facebook Developer dashboard. See the [Facebook Platform](#)

[Changelog](#) to see if your version is still supported.



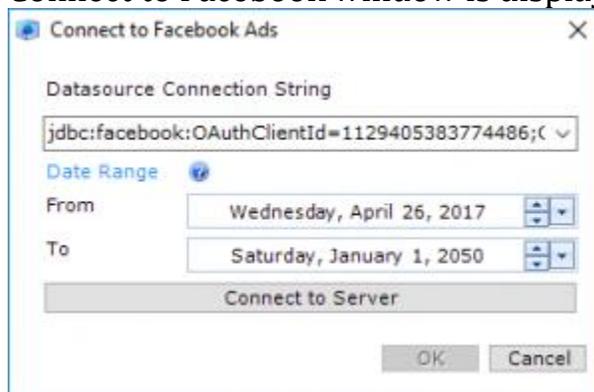
## Adding Facebook Ads Tables to your ElastiCube

After you have changed the connectivity settings of the predefined ElastiCube, or if you want to create your own ElastiCube from scratch, you can add Facebook Ads tables exposed by the Facebook API.

Facebook provides access to data associated with the Facebook account you defined in the connection string. Once you have connected, Sisense displays a list of Facebook Ads tables you can import. For a list of supported tables, see Facebook Ads Tables.

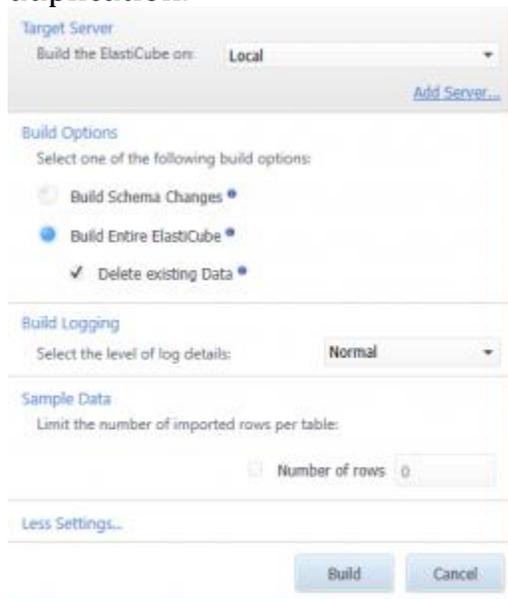
### To add Facebook Ads data:

1. In ElastiCube Manager, click **Add Data** and then, **Facebook Ads**. The Connect to Facebook window is displayed.



2. In the **Date Range fields**, select the date range for the relevant Facebook Ads data you want to import.

**Note:** When performing an accumulative build, if you change the date range from your previous build, select **Delete existing Data** to prevent duplication.



The screenshot shows a 'Build' dialog box with the following sections:

- Target Server:** 'Build the ElastiCube on:' dropdown set to 'Local'. A link 'Add Server...' is visible.
- Build Options:** 'Select one of the following build options:' with three radio buttons: 'Build Schema Changes', 'Build Entire ElastiCube', and 'Delete existing Data' (which is checked).
- Build Logging:** 'Select the level of log details:' dropdown set to 'Normal'.
- Sample Data:** 'Limit the number of imported rows per table:' with a checkbox and a text input field containing '0'.
- Less Settings...** link.
- Buttons:** 'Build' and 'Cancel' buttons at the bottom.

3. In **Datasource Connection String**, enter your connection string.
4. Click **Connect to Server**. Facebook is displayed in the Select Database list.
5. Click **OK**. Sisense connects to Facebook and displays a list of tables available for you to import. For a list of tables you can import, see Facebook Ads Tables.
6. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

**Tip:** Facebook Ads leverages the Facebook Marketing API, however, this API includes user request limitations. To overcome this limitation, try to reduce the date range or query fewer Ad accounts by removing them from your developers profile (**Settings > Advanced > Advanced Ad Account IDs**). For more information, see Facebook's Marketing Rate limitations.

## Facebook Ads Tables

The table below describes the tables you can import from various Facebook APIs. Facebook limits the number of calls and the amount of data that you can retrieve from their APIs. These limitations depend on which Facebook API you are using

and your account type. When you reach a limit, Sisense returns the error provided by Facebook. For example, if you receive any of the following error codes, this usually indicates that you have reached one of Facebook’s limitations:

- error code = 4, CodedException
- error code = 32, CodedException
- error code = 17, User request limit reached

If you encounter one of these errors, refer to the Facebook API that exposed the table for more details regarding that API’s limitations.

#### Available Tables

Table	Description
AdAccounts	The Ad Accounts available for a user. AdAccounts in Facebook are the accounts you have available in Facebook for creating Ads on. Accessing Ad Account information requires the ads_read permission.
AdCreatives	Information about an Ad Creative or the Ad Creatives on a specific Ad Account, Ad Set, or Ad. AdCreatives in Facebook represent a collection of the creatives for specific Ads in Facebook. Accessing Ad Creative information requires the ads_read permission.
AdInsights	Ads Insights provides access for reporting and analytics purposes. Accessing AdInsights information requires the ads_read permission.
AdInsights_Age_Gender	The following are breakdowns of the AdInsights table. Due to storage constraints, Facebook has made only some permutations of breakdowns available. For more information, see <a href="#">Facebook Breakdowns</a> .
AdInsights_Country	
AdInsights_FrequencyValue	
AdInsights_Placement_	

Table	Description
Device	
AdInsights_HStatsByAdvertiserTZ	
AdInsights_HStatsByAudienceTZ	
AdInsights_Placement_Device	
AdInsights_Placement_ImpressionDevice	
AdInsights_ProductId	
AdInsights_PublisherPlatform_ImpressionDevice_PlatformPosition	
AdInsights_Region	
AdInsightsActions	Query an Ad Report. Accessing Ad Report information requires the ads_read permission. AdInsightsActions represents a breakdown of the Actions column from AdInsights.
AdInsightsActions_Age_Gender	The following are breakdowns of the AdInsightsActions table. Due to storage constraints, Facebook has made only some permutations of breakdowns available. For more information, see <a href="#">Facebook Breakdowns</a> .
AdInsightsActions_Country	
AdInsightsActions_FrequencyValue	

Table	Description
AdInsightsActions_HStatsByAdvertiserTZ	
AdInsightsActions_HStatsByAudienceTZ	
AdInsightsActions_Placement_Device	
AdInsightsActions_Placement_ImpressionDevice	
AdInsightsActions_ProductId	
AdInsightsActions_PublisherPlatform_ImpressionDevice_PlatformPosition	
AdInsightsActions_Region	
Ads	Query information about an Ad or the Ads in a specific Ad Set, Campaign, or Ad Account. Accessing Ad Information requires the ads_read permission. Ads in Facebook represent individual ads that have been created.
AdScheduledReports	Query the Scheduled Reports for a given Ad Account. Accessing Scheduled Report information requires the ads_read permission. The scheduled reports are a list of reports scheduled to run every day for your account.

Table	Description
	These can be maintained in the Facebook Ads Manager. Daily runs of these reports are available from ScheduledReportRuns.
AdSets	Query information about an Ad Set or the Ad Sets on a specific Campaign or Ad Account. Accessing Ad Set information requires the ads_read permission. AdSets in Facebook represent a collection of Ads in Facebook.
Campaigns	Query information about a Campaign or the Campaigns on a specific Ad Account. Accessing Campaign information requires the ads_read permission. Campaigns in Facebook represent advertising campaigns that individual Ads or AdSets may be a part of.

#### Versions

Version	Release Date	Improvements
6.6.1.50	June 12, 2017	New version of Graph API v2.9 Changelog New objects: SimpleVideoInsights, VideoInsightsByActionType, VideoInsightsByDistributionType, VideoInsightsByReactionType
6.6.1.45	May 17, 2017	General Release

### Version 6.6.1.45

#### Breaking Changes

In the latest version of the Facebook Graph API, permissions have been largely reworked. All of the friends\* permissions have been removed. In addition, most

permissions require Facebook approval if you wish to request information about other user accounts who are not an administrator of the Facebook App. Sisense recommends reviewing the new permissions listed by Facebook here:

<https://developers.facebook.com/docs/facebook-login/permissions/v2.1>

In addition, users may reject requested permissions in the Authorization dialog on Facebook. If this happens, the specific Permissions requested will state that they are declined under the Status column. The permission may be rerequested by executing `GetOAuthAccessToken` and setting `AuthType` to `rerequest`.

If you want to keep using an old version of the API, you can by simply setting the `Version` connection property in the connection string to the specific version you were using and also setting the `Location` to a folder where you have saved an older copy of the schema files (.rsd and .rsb files).

If you want to update your app to 2.9, then you don't have to do anything

**Note:** Sisense removed few metrics from the `AdInsights / AdInsightsActions` due to their removal by Facebook in 2.9.

A number of various changes have been made across tables due to changes in Facebook. These changes are listed below:

- Facebook no longer allows public posts to be searched. We have had to remove `SearchTerms` from the `Posts` table because of this.
- Checkins have been removed from the Graph API. Due to this, we have removed the `Checkins` table.
- Events may no longer be inserted, updated, or deleted via the Graph API. We have had to change the `Events` table to a view.
- Usernames for Facebook Users may no longer be used as a target and are no longer returned. The Facebook username for Pages is unaffected.
- User ID's for Facebook Users will be different across different Facebook Apps. This is due to the new policy of App-scoped User ID's.

# Facebook Insights

To retrieve data from Facebook Insights, you must make a request that contains several parameters with unique values to a specific table.

Each request includes the name of the table you want to request data from, for example, `InsightsByConsumptionType`, and three unique parameters described below

**Insight Name:** The metric to be returned.

**Target:** The element to retrieve the insights from. When requesting data from a source, all the data is returned for the authorized user who made the request. For example, `target in (select ID from pages)` returns all the pages for the relevant ID for the defined period of time.

**Note:** When requesting data for all posts for all of your pages, you might experience performance issues. Sisense recommends breaking up requests in which large amounts of data are returned by Facebook.

**Period:** The aggregation period. The value 'lifetime' includes all your data up to the time of the request. Any other values such as 'day', 'month' etc. are limited to a range of 90 days as per Facebook's time range limitation.

This page provides example requests and lists the possible values for each table.

## InsightsByConsumptionType

### Example:

```
SELECT EndTime as Date, Target as PageID, Period,
InsightName, VideoPlay, OtherClicks, PhotoView,
LinkClicks, ButtonClicks
FROM InsightsByConsumptionType
```

where target in (select ID from pages) and InsightName = 'PAGE\_CONSUMPTIONS\_BY\_CONSUMPTION\_TYPE' and Period = 'day' and FromDateTime='Date' and ToDateTime='Date'

## Mapping

Insight Name	Target	Period
PAGE_CONSUMPTIONS_BY_CONSUMPTION_TYPE	Page	day, week, days_28
PAGE_CONSUMPTIONS_BY_CONSUMPTION_TYPE_UNIQUE	Page	day, week, days_28
POST_CONSUMPTIONS_BY_TYPE	Post	day, week, days_28
POST_CONSUMPTIONS_BY_TYPE_UNIQUE	Post	day, week, days_28

## InsightsByFeedbackType

### Example:

```
SELECT EndTime as Date, Target as PageID, Period,
InsightName, [Like], Comment, Link, Answer, Claim, Rsvp,
HideAllClicks, HideClicks, UnlikePageClicks,
ReportSpamClicks, XButtonClicks
FROM InsightsByFeedbackType
where target in (select ID from pages) and InsightName =
'PAGE_NEGATIVE_FEEDBACK_BY_TYPE' and Period = 'day' and
FromDateTime='Date' and ToDateTime='Date'
```

## Mapping

Insight Name	Target	Period
PAGE_NEGATIVE_FEEDBACK_BY_TYPE	Page	day, week, days_28
PAGE_NEGATIVE_FEEDBACK_BY_TYPE_UNIQUE	Page	day, week, days_28
PAGE_POSITIVE_FEEDBACK_BY_TYPE	Page	day, week, days_28
PAGE_POSITIVE_FEEDBACK_BY_TYPE_UNIQUE	Page	day, week, days_28
POST_NEGATIVE_FEEDBACK_BY_TYPE	Post	lifetime
POST_NEGATIVE_FEEDBACK_BY_TYPE_UNIQUE	Post	lifetime

## InsightsByLikeSourceType

### Example:

```
SELECT EndTime as Date, Target as PageID, Period,
InsightName, PageSuggestion, PageTimeline, Ads,
MobileAds, Registration, Mobile, WizardSuggestion,
ProfileConnect, ExternalConnect,
RecommendedPages, Favorites, Api, PageBrowser,
MobilePageBrowser, Hovercard, Search, PageProfile,
Ticker, LikeStory, FeedChaining, AllCategoryPym1,
PageSuggestionsOnLiking, MobilePageSuggestionsOnLiking,
FanContextStory, SponsoredStory,
PageInviteEscapeHatchFinch, PageAdminNumPosts,
PageAdminNumPostsByType
FROM InsightsByLikeSourceType
```

where target in (select ID from pages) and InsightName = 'PAGE\_FANS\_BY\_LIKE\_SOURCE' and FromDateTime='Date' and ToDateTime='Date'

### Mapping

Insight Name	Target	Period
PAGE_FANS_BY_LIKE_SOURCE	Page	day
PAGE_FANS_BY_LIKE_SOURCE_UNIQUE	Page	day
PAGE_FANS_BY_UNLIKE_SOURCE_UNIQUE	Page	day

## InsightsByPaidStatus

### Example:

```
SELECT EndTime as Date, Target as PageID, Period,
InsightName, Total, Paid, Unpaid
FROM InsightsByPaidStatus
```

where target in (select ID from pages) and InsightName = 'PAGE\_IMPRESSIONS\_BY\_PAID\_NON\_PAID' and Period = 'day' and FromDateTime='Date' and ToDateTime='Date'

### Mapping

Insight Name	Target	Period
PAGE_IMPRESSIONS_BY_PAID_NON_PAID	Page	day, week, days_28
PAGE_IMPRESSIONS_BY_PAID_NON_PAID_UNIQUE	Page	day, week, days_28
POST_IMPRESSIONS_BY_PAID_NON_PAID	Post	lifetime

Insight Name	Target	Period
POST_IMPRESSIONS_BY_PAID_NON_PAID_UNIQUE	Post	lifetime
PAGE_POSTS_IMPRESSIONS_BY_PAID_NON_PAID	Post	day, week, days_28
PAGE_POSTS_IMPRESSIONS_BY_PAID_NON_PAID_UNIQUE	Post	day, week, days_28

## InsightsByReactionTotals

### Example:

```
SELECT EndTime as Date, Target as PageID, Period,
InsightName, [Like], Love, Wow, Haha, Sorry, Anger,
Target
FROM InsightsByReactionTotals
where target in (select ID from pages)
and InsightName = 'PAGE_ACTIONS_POST_REACTIONS_TOTAL' and
Period = 'day' and FromDateTime='Date' and
ToDateTime='Date'
```

### Mapping

Insight Name	Target	Period
PAGE_ACTIONS_POST_REACTIONS_TOTAL	Page	day

## InsightsByStoryType

### Example:

```

SELECT EndTime as Date,Target as PostID, Period,
InsightName, UserPost, PagePost, Checkin, Fan, Question,
Coupon, Event, Mention, Other
FROM InsightsByStoryType
where target in (select ID from Posts where target in
(select ID from pages))
and InsightName = 'POST_IMPRESSIONS_BY_STORY_TYPE' and
Period = 'day' and FromDateTime='Date' and
ToDateTime='Date'

```

## Mapping

Insight Name	Target	Period
PAGE_STORIES_BY_STORY_TYPE	Page	day, week, days_28
PAGE_STORYTELLERS_BY_STORY_TYPE	Page	day, week, days_28
PAGE_IMPRESSIONS_BY_STORY_TYPE	Page	day, week, days_28
PAGE_IMPRESSIONS_BY_STORY_TYPE_UNIQUE	Page	day, week, days_28
POST_IMPRESSIONS_BY_STORY_TYPE	Post	day, week, days_28
POST_IMPRESSIONS_BY_STORY_TYPE_UNIQUE	Post	day, week, days_28

## InsightsByTabType

### Example:

```

SELECT EndTime as Date,Target as PageID, Period,
InsightName, AllActivity, App, Info, Insights, Likes,
Locations, Photos, PhotosAlbums, PhotosStream, Profile,

```

```

ProfileInfo, ProfileLikes, ProfilePhotos, Timeline,
Events, Videos, Wall
FROM InsightsByTabType
where target in (select ID from pages)
and InsightName = 'PAGE_TAB_VIEWS_LOGIN_TOP_UNIQUE' and
Period = 'day' and FromDateTime='Date' and
ToDateTime='Date'

```

### Mapping

Insight Name	Target	Period
PAGE_TAB_VIEWS_LOGIN_TOP_UNIQUE	Page	day, week
PAGE_TAB_VIEWS_LOGIN_TOP	Page	day, week
PAGE_TAB_VIEWS_LOGOUT_TOP	Page	day

## SimpleInsights

### Example:

```

SELECT EndTime as Date, Target as PageID, Period,
InsightName, Value FROM SimpleInsights where target in
(select ID from pages) and InsightName =
'PAGE_ACTIONS_POST_REACTIONS_ANGER_TOTAL' and Period =
'day' and FromDateTime='@Date' and ToDateTime='@Date'

```

### Mapping

Insight Name	Target	Period
PAGE_STORIES	Page	day, week, days_28
POST_STORIES	Page	lifetime

<b>Insight Name</b>	<b>Target</b>	<b>Period</b>
POST_STORYTELLERS	Page	lifetime
PAGE_IMPRESSIONS	Page	day, week, days_28
PAGE_IMPRESSIONS_UNIQUE	Page	day, week, days_28
PAGE_IMPRESSIONS_PAID	Page	day, week, days_28
PAGE_IMPRESSIONS_PAID_UNIQUE	Page	day, week, days_28
PAGE_IMPRESSIONS_ORGANIC	Page	day, week, days_28
PAGE_IMPRESSIONS_ORGANIC_UNIQUE	Page	day, week, days_28
PAGE_IMPRESSIONS_VIRAL	Page	day, week, days_28
PAGE_IMPRESSIONS_VIRAL_UNIQUE	Page	day, week, days_28
PAGE_IMPRESSIONS_FREQUENCY_DISTRIBUTION	Page	day, week, days_28
PAGE_IMPRESSIONS_VIRAL_FREQUENCY_DISTRIBUTION	Page	day, week, days_28
PAGE_ENGAGED_USERS	Page	day, week, days_28
PAGE_CONSUMPTIONS	Page	day, week,

<b>Insight Name</b>	<b>Target</b>	<b>Period</b>
		days_28
PAGE_CONSUMPTIONS_UNIQUE	Page	day, week, days_28
PAGE_PLACES_CHECKIN_TOTAL	Page	day, week, days_28
PAGE_PLACES_CHECKIN_TOTAL_UNIQUE	Page	day, week, days_28
PAGE_PLACES_CHECKIN_MOBILE	Page	day, week, days_28
PAGE_PLACES_CHECKIN_MOBILE_UNIQUE	Page	day, week, days_28
PAGE_NEGATIVE_FEEDBACK	Page	day, week, days_28
PAGE_NEGATIVE_FEEDBACK_UNIQUE	Page	day, week, days_28
PAGE_FANS_ONLINE	Page	day
PAGE_FANS_ONLINE_PER_DAY	Page	day
PAGE_FANS	Page	lifetime
PAGE_FAN_ADDS	Page	day
PAGE_FAN_ADDS_UNIQUE	Page	day, week
PAGE_FAN_REMOVES	Page	day
PAGE_FAN_REMOVES_UNIQUE	Page	day
PAGE_VIEWS	Page	day

<b>Insight Name</b>	<b>Target</b>	<b>Period</b>
PAGE_VIEWS_UNIQUE	Page	day, week
PAGE_VIEWS_LOGIN	Page	day, week
PAGE_VIEWS_LOGIN_UNIQUE	Page	day, week
PAGE_VIEWS_LOGOUT	Page	day
PAGE_VIDEO_VIEWS	Page	day, week, days_28
PAGE_VIDEO_VIEWS_PAID	Page	day, week, days_28
PAGE_VIDEO_VIEWS_ORGANIC	Page	day, week, days_28
PAGE_VIDEO_VIEWS_AUTOPLAYED	Page	day, week, days_28
PAGE_VIDEO_VIEWS_CLICK_TO_PLAY	Page	day, week, days_28
PAGE_VIDEO_VIEWS_UNIQUE	Page	day, week, days_28
PAGE_VIDEO_REPEAT_VIEWS	Page	day, week, days_28
PAGE_VIDEO_COMPLETE_VIEWS_30S	Page	day, week, days_28
PAGE_VIDEO_COMPLETE_VIEWS_30S_PAID	Page	day, week, days_28
PAGE_VIDEO_COMPLETE_VIEWS_30S_ORGANIC	Page	day, week, days_28

<b>Insight Name</b>	<b>Target</b>	<b>Period</b>
PAGE_VIDEO_COMPLETE_VIEWS_30S_AUTOPLAYED	Page	day, week, days_28
PAGE_VIDEO_COMPLETE_VIEWS_30S_CLICK_TO_PLAY	Page	day, week, days_28
PAGE_VIDEO_COMPLETE_VIEWS_30S_UNIQUE	Page	day, week, days_28
PAGE_VIDEO_COMPLETE_VIEWS_30S_REPEAT_VIEWS	Page	day, week, days_28
PAGE_POSTS_IMPRESSIONS	Post	lifetime
PAGE_POSTS_IMPRESSIONS_UNIQUE	Post	lifetime
PAGE_POSTS_IMPRESSIONS_PAID	Post	lifetime
PAGE_POSTS_IMPRESSIONS_PAID_UNIQUE	Post	lifetime
PAGE_POSTS_IMPRESSIONS_ORGANIC	Post	lifetime
PAGE_POSTS_IMPRESSIONS_ORGANIC_UNIQUE	Post	lifetime
PAGE_POSTS_IMPRESSIONS_VIRAL	Post	lifetime
PAGE_POSTS_IMPRESSIONS_VIRAL_UNIQUE	Post	lifetime
POST_IMPRESSIONS	Post	lifetime
POST_IMPRESSIONS_UNIQUE	Post	lifetime
POST_IMPRESSIONS_PAID	Post	lifetime
POST_IMPRESSIONS_PAID_UNIQUE	Post	lifetime
POST_IMPRESSIONS_FAN	Post	lifetime
POST_IMPRESSIONS_FAN_UNIQUE	Post	lifetime
POST_IMPRESSIONS_FAN_PAID	Post	lifetime

<b>Insight Name</b>	<b>Target</b>	<b>Period</b>
POST_IMPRESSIONS_FAN_PAID_UNIQUE	Post	lifetime
POST_IMPRESSIONS_ORGANIC	Post	lifetime
POST_IMPRESSIONS_ORGANIC_UNIQUE	Post	lifetime
POST_IMPRESSIONS_VIRAL	Post	lifetime
POST_IMPRESSIONS_VIRAL_UNIQUE	Post	lifetime
POST_CONSUMPTIONS	Post	day, week, days_28
POST_CONSUMPTIONS_UNIQUE	Post	day, week, days_28
POST_ENGAGED_USERS	Post	day, week, days_28
POST_NEGATIVE_FEEDBACK	Post	day, week, days_28
POST_NEGATIVE_FEEDBACK_UNIQUE	Post	day, week, days_28
POST_VIDEO_AVG_TIME_WATCHED	Post	lifetime
POST_VIDEO_COMPLETE_VIEWS_ORGANIC	Post	lifetime
POST_VIDEO_COMPLETE_VIEWS_ORGANIC_UNIQUE	Post	lifetime
POST_VIDEO_COMPLETE_VIEWS_PAID	Post	lifetime
POST_VIDEO_COMPLETE_VIEWS_PAID_UNIQUE	Post	lifetime
POST_VIDEO_VIEWS_ORGANIC	Post	lifetime
POST_VIDEO_VIEWS_ORGANIC_UNIQUE	Post	lifetime
POST_VIDEO_VIEWS_PAID	Post	lifetime

Insight Name	Target	Period
POST_VIDEO_VIEWS_PAID_UNIQUE	Post	lifetime

## SimpleVideoInsights

### Example:

```
SELECT Target as PostID, Period, InsightName, Value
FROM SimpleVideoInsights where target in (select id
from videos where target in (select ID from Pages)) and
InsightName = 'TOTAL_VIDEO_VIEWS' and Period =
'lifetime'
```

### Mapping

Insight Name	Target	Period
PAGE_VIDEO_VIEW_TIME	Page	day
TOTAL_VIDEO_VIEWS	Video	lifetime
TOTAL_VIDEO_VIEWS_UNIQUE	Video	lifetime
TOTAL_VIDEO_VIEWS_AUTOPLAYED	Video	lifetime
TOTAL_VIDEO_VIEWS_CLICKED_TO_PLAY	Video	lifetime
TOTAL_VIDEO_VIEWS_SOUND_ON	Video	lifetime
TOTAL_VIDEO_COMPLETE_VIEWS	Video	lifetime
TOTAL_VIDEO_COMPLETE_VIEWS_UNIQUE	Video	lifetime
TOTAL_VIDEO_COMPLETE_VIEWS_AUTO_PLAYED	Video	lifetime
TOTAL_VIDEO_COMPLETE_VIEWS_CLICKED_TO_PLAY	Video	lifetime
TOTAL_VIDEO_10S_VIEWS	Video	lifetime

Insight Name	Target	Period
TOTAL_VIDEO_10S_VIEWS_UNIQUE	Video	lifetime
TOTAL_VIDEO_10S_VIEWS_auto_played	Video	lifetime
TOTAL_VIDEO_10S_VIEWS_CLICKED_TO_PLAY	Video	lifetime
TOTAL_VIDEO_10S_VIEWS_SOUND_ON	Video	lifetime
TOTAL_VIDEO_AVG_TIME_WATCHED	Video	lifetime
TOTAL_VIDEO_VIEW_TOTAL_TIME	Video	lifetime
TOTAL_VIDEO_IMPRESSIONS	Video	lifetime
TOTAL_VIDEO_IMPRESSIONS_UNIQUE	Video	lifetime
TOTAL_VIDEO_IMPRESSIONS_VIRAL_UNIQUE	Video	lifetime
TOTAL_VIDEO_IMPRESSIONS_VIRAL	Video	lifetime
TOTAL_VIDEO_IMPRESSIONS_FAN_UNIQUE	Video	lifetime
TOTAL_VIDEO_IMPRESSIONS_FAN	Video	lifetime

## VideoInsightsByActionType

### Example:

```
SELECT Target as PostID, Period, InsightName, Value
FROM VideoInsightsByActionType
where target in (select id from videos where target in
(select ID from Pages)) and InsightName =
'TOTAL_VIDEO_STORIES_BY_ACTION_TYPE' and Period =
'lifetime'
```

### Mapping

Insight Name	Target	Period
--------------	--------	--------

Insight Name	Target	Period
TOTAL_VIDEO_STORIES_BY_ACTION_TYPE	Video	lifetime

## VideoInsightsByDistributionType

### Example:

```
SELECT EndTime as Date, Target as PageID, Period,
InsightName, Page_Owned, Shared, Crossposted
FROM VideoInsightsByDistributionType
where target in (select ID from pages)
and InsightName =
'TOTAL_VIDEO_VIEWS_BY_DISTRIBUTION_TYPE' and Period =
'lifetime'
```

### Mapping

Insight Name	Target	Period
TOTAL_VIDEO_VIEWS_BY_DISTRIBUTION_TYPE	Video	lifetime
TOTAL_VIDEO_VIEW_TIME_BY_DISTRIBUTION_TYPE	Video	lifetime

## VideoInsightsByReactionType

### Example:

```
SELECT Target as PostID, Period, InsightName, Value
FROM VideoInsightsByReactionType
where target in (select id from videos where target in
(select ID from Pages)) and InsightName =
'TOTAL_VIDEO_REACTIONS_BY_TYPE_TOTAL' and Period =
'lifetime'
```

### Mapping

<b>Insight Name</b>	<b>Target</b>	<b>Period</b>
TOTAL_VIDEO_REACTIONS_BY_TYPE_TOTAL	Video	lifetime

# Connecting to Facebook

The Sisense Facebook connector is a standalone connector that allows you to import data from various Facebook APIs into the ElastiCube Manager without any custom queries. With just a few clicks, you can import tables directly from the Facebook API into your ElastiCube Manager.

Sisense supports two versions of the Facebook connector, the version described on this page and an earlier version described here. The difference between this version and the earlier version is that this version does not require custom SQL when importing tables, but does require SisenseV6.6.1 or higher.

After you have downloaded and installed the connector, you can connect to Facebook through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to Facebook's APIs.

This page describes how to install the Facebook connector, how to connect to Facebook with a connection string, and what tables you can import into the ElastiCube Manager.

- [Prerequisites](#)
- [Installing the Facebook Connector](#)
- [Adding Facebook Tables to your ElastiCube](#)
- [Facebook Tables](#)

## **Breaking Changes for the Marketing API Version 2.8**

As of July 25, 2017, Facebook no longer supports their Marketing API version 2.8. Currently, all users whose Facebook app is v2.8 will receive an error. To overcome this issue, add the parameter 'Version=2.10' to the connection string. For more information about the latest version of the Marketing API, see Facebook's [blog posts](#).

# Prerequisites

The Facebook connector requires Sisense V6.6.1 and later.

This connector is compatible with previous Facebook connectors. The following tables have been moved to the Facebook Ads connector: AdAccounts, AdCreatives, Ads, AdScheduledReports, AdSets, Campaigns, AdInsights, AdInsightsActions to continue to use these connectors in an existing ElastiCube, see [Changing Facebook Connectivity Settings](#).

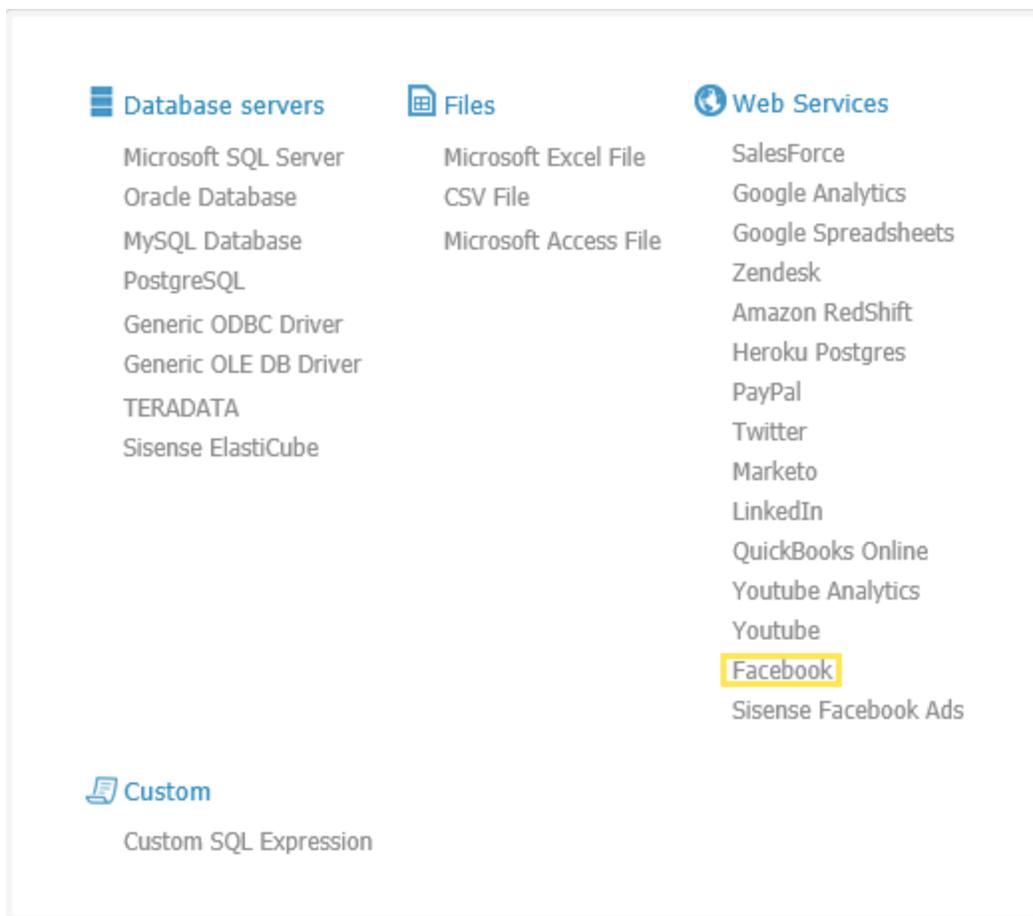
## Installing the Facebook Connector

Sisense provides the Facebook connector as a standalone connector that you can download and add to your list of default Sisense connectors. After you install the Facebook connector, it is added to your list of Web Service connectors.

### To install the Facebook connector:

1. [Download](#) the Facebook installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The Facebook connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.



## Adding Facebook Tables to your ElastiCube

Sisense uses connection strings to connect to Facebook and import data into the ElastiCube Manager. The default connection string for Facebook is:

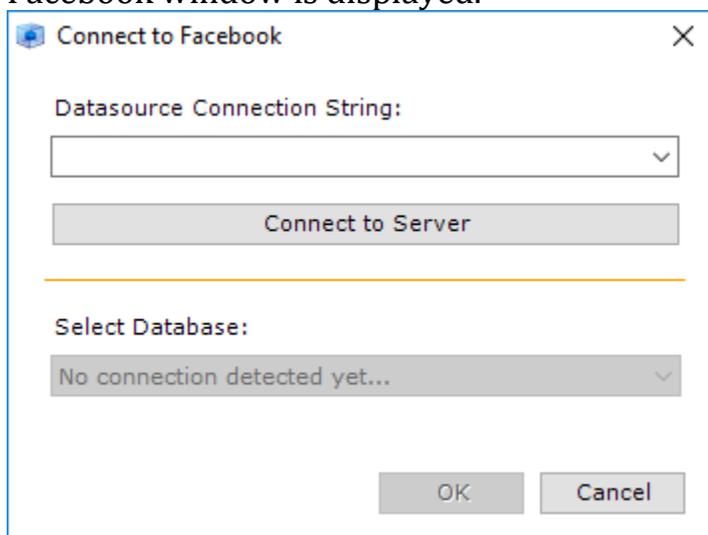
```
jdbc:facebook:InitiateOAuth=GETANDREFRESH;
```

After you connect using the default connection string, Sisense redirects you to Facebook's Login page where you log in and your account is authenticated. Once the account is authenticated, the relevant tables can be added to your ElastiCube. If you want to customize your connection string, you can by creating a Facebook app and passing additional connection string parameters. For more information, see [Creating an App](#) at the end of this page.

Once you have connected to the Facebook datasource, Sisense displays a list of Facebook tables you can import. For a list of supported tables, see [Facebook Tables](#).

### To add Facebook data:

1. In ElastiCube Manager, click **Add Data** and then, **Facebook**. The Connect to Facebook window is displayed.



2. In **Datasource Connection String**, enter the connection string:  
`jdbc:facebook:InitiateOAuth=GETANDREFRESH;`
3. Click **Connect to Server**. Facebook is displayed in the Select Database list.
4. Click **OK**. Sisense connects to Facebook and displays a list of tables available for you to import. For a list of tables you can import, see [Facebook Tables](#).
5. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

**Tip:** Facebook leverages the Facebook Graph API, however, this API includes user request limitations. To overcome these limitations, try to reduce the date range or query fewer tables for a single build. For more information, see [Facebook's Graph API limitations](#).

## Facebook Tables

The table below describes the tables you can import from various Facebook APIs. Facebook limits the number of calls and the amount of data that you can retrieve

from their APIs. These limitations depend on which Facebook API you are using and your account type. When you reach a limit, Sisense returns the error provided by Facebook. For example, if you receive any of the following error codes, this usually indicates that you have reached one of Facebook's limitations:

- error code = 4, CodedException
- error code = 32, CodedException
- error code = 17, User request limit reached

If you encounter one of these errors, refer to the Facebook API that exposed the table for more details regarding that API's limitations.

**Note:** Some Facebook tables require custom SQL, the table below provides examples of custom SQL statements you can use to access the relevant table.

## Available Tables

Name	Description
Groups	The Groups table returns information about the groups of the user defined in the connection string. Groups may require the user_groups permission.
GroupMemberships	The GroupMemberships table returns all members of all groups for the user defined in the connection string. Groups may require the user_groups permission. GroupMemberships in Facebook are the individual members of a given group.
Albums	The Albums table returns a information about albums associated with a user defined in the connection

Name	Description
	string. Accessing Album information typically requires the user_photos permission.
Applications	The Application table returns information about a user's applications.
Books	The Books table returns information regarding books a user is interested in. Accessing Book information typically requires the user_books permissions.
Events	The Events for a user based on either the Target or SearchTerms. May require the user_events permission.
Users	The Users table returns information about a user according to SearchTerms or Id.
Pages	The Pages table returns the pages of a user.
Posts	The Posts table for a user based on either the Target or Id. Posts can also be inserted based on a Target, or deleted based on Id. This table requires authentication.
Page_Posts	The Page_Posts table returns information regarding the posts for all

Name	Description
	pages for a user.
Post_Comments	All Post_Ccomments table returns information regarding all the posts of a user.
Comments_For_Page_Posts	All Comments_For_Page_Posts table returns information regarding all the comments of a page post by a user.
Album_Comments	The Album_Comments table returns information regarding comments for a user's albums.
Photo_Comments	The Photo_Comments table returns information regarding comments for a user's photos.
Video_Comments	The Video_Comments table returns information regarding comments for a user's videos.
MyVideos	The MyVideos table returns information regarding videos for a user. Videos in Facebook are videos that have been uploaded by a user or to a page on Facebook. Normally requires the user_videos permission.
MyLikes	The Likes table returns information regarding a user's likes. Likes in Facebook represent the users that like

Name	Description
	a particular target. The target may be a post, page, picture, or other type of valid entity that may have likes associated with it.
Post_Likes	The Post_Likes table returns information regarding posts a user liked.
Page_Likes	The Page_Likes table returns information regarding pages a user liked.
Photo_Likes	The Photo_Likes table returns information regarding photos a user liked.
Milestones	The Milestones table returns information regarding milestones for the given Page. Milestones may only be inserted, updated, or deleted when authenticating as the target Page.
Friends	The Friends table returns information regarding the friends of a user. Requires the user_friends permission.
Games	The Games table returns information regarding games a user is interested in. Accessing Game information may require the user_likes and

Name	Description
	user_interests permissions.
InsightsByConsumptionType	Allows retrieval of insights by consumption type. For examples and more information regarding this table, see <a href="#">Facebook Insights</a> .
InsightsByFeedbackType	Allows retrieval of insights by feedback type. For examples and more information regarding this table, see <a href="#">Facebook Insights</a> .
InsightsByLikeSourceType	Allows retrieval of insights by like source type. For examples and more information regarding this table, see <a href="#">Facebook Insights</a> .
InsightsByPaidStatus	Allows retrieval of insights by paid status. For examples and more information regarding this table, see <a href="#">Facebook Insights</a> .
InsightsByReactionTotals	Allows retrieval of insights by like source type. For examples and more information regarding this table, see <a href="#">Facebook Insights</a> .
InsightsByStoryType	Allows retrieval of insights by like story type. For examples and more information regarding this table, see <a href="#">Facebook Insights</a> .

Name	Description
InsightsByTabType	Allows retrieval of insights by tab type. For examples and more information regarding this table, see <a href="#">Facebook Insights</a> .
SimpleInsights	Allows the retrieval of simple insights with a single value in the response. For examples and more information regarding this table, see <a href="#">Facebook Insights</a> .
SimpleVideoInsights	Allows the retrieval of simple video insights with a single value in the response. For examples and more information regarding this table, see <a href="#">Facebook Insights</a> .
VideoInsightsByActionType	Allows the retrieval of video insights by story action type. For examples and more information regarding this table, see <a href="#">Facebook Insights</a> .
VideoInsightsByDistributionType	Allows the retrieval of video insights by distribution type. For examples and more information regarding this table, see <a href="#">Facebook Insights</a> .
VideoInsightsByReactionType	Allows the retrieval of video insights by reaction type. For examples and more information regarding this table, see <a href="#">Facebook Insights</a> .

Name	Description
Movies	The Movies table returns information regarding movies a user is interested in. Accessing Movie information requires the user_likes and user_interests permissions.
Music	The Music table returns information regarding music a user is interested in. Accessing Music requires the user_likes and user_interests permissions.
Photos	The Photos table returns information regarding a user's photos. Accessing Photo information typically requires the user_photos permission.
Album_Photos	The Album_Photos table returns information regarding a user's album photos. Accessing Album information typically requires the user_photos permission.
Event_Photos	The Event_Photos table returns information regarding a user's photos from an event. May require the user_events permission.
Page_Photos	The Page_Photos table returns information regarding a user's photos from any of their pages.

Name	Description
User_Photos	The Photos table returns information regarding a user's photos. Accessing Photo information typically requires the user_photos permission.
TaggedBy	Query information about Posts, Statuses, Photos, and other entities that have tagged the User or Page. This view is a derivative of the Wall connection where only entries that have tagged the Target User or Page will be returned. In general it is only available for Pages.
Television	Query the TV shows a User is interested in. Accessing Television information may require the user_likes and user_interested permissions.
Wall	Query Posts from the Wall of a Target.
Permissions	Query the Permissions the User has granted the current application.
Likes_For_Page_Posts	Likes of my posts on pages.

**Note:** For the following tables: InsightsByConsumptionType, InsightsByFeedbackType, InsightsByLikeSourceType, InsightsByPaidStatus, InsightsByTabType, SimpleInsights, the Object Insights period {day, week, month, lifetime} lifetime is not supported for all measures. For more information, see Facebook's API documentation.

## Creating an App

To access Facebook's data from Sisense with a customized connection string, you must provide valid Facebook credentials in a connection string.

This section describes how to create a Facebook app to retrieve the necessary authentication parameters for the connections string and then how to construct the connection string.

**Note:** If you have already created a connection string for the Facebook Ads connector, you can use the same credentials in the Facebook connection string. The credentials you need to connect to your data in Facebook are provided by Facebook when you create a Facebook app through the Facebook Developer site. For more information, see [Creating an App](#).

## Creating a Connection String

Sisense uses connection strings to connect to Facebook and import data into the ElastiCube Manager. Each connection string contains authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

**Note:** When switching accounts, you must sign out of your account and then sign it with the new account.

After you have obtained the relevant credentials, you can create the connection used to connect to your Facebook account. The Facebook connection string has the following structure:

```
jdbc:DataSourceName:Property1=Value1;Property2=Value2;
```

The following is an example of a Facebook connection string:

```
jdbc:facebook:OAuthClientId=112345383774486;OAuthClientSecret=064c524478d712534b7e7e4224fad;
```

```
InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;Timeout=0;ThrottleSeconds=15;ThrottleMaxPercent=90;
```

```
Other=RetryWaitTime=60000;MaximumRequestRetries=5;
```

The example above includes both mandatory and optional parameters you can provide in the connection. The required parameters are emphasized in bold.

### Mandatory Parameters

- **OAuthClientId**: Set this to the App ID in your app settings.
- **OAuthClientSecret**: Set this to the App Secret in your app settings.
- **CallbackURL**: Set this to the Return URL in your app settings.
- **InitiateOAuth**: Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken connection property.

### *Optional Parameters*

- **Aggregate Format**: The format aggregate or collection columns should return in.
- **Authenticate As Page**: The name or Id of a page to authenticate as when making requests to Facebook.
- **Auto Cache**: The AutoCache property controls automatic caching of data.
- **Cache Connection**: The connection string for the cache database.
- **Cache Driver**: The database driver to be used to cache data.
- **Cache Location**: Specifies the path to the cache when caching to a file.
- **Cache Metadata**: Whether or not to cache the table metadata. The cache will be stored in memory unless CacheConnection or CacheLocation are set.
- **Callback URL**: The OAuth callback URL to return to when authenticating. This value must match the callback URL you specify in your app settings.
- **Firewall Password**: A password, if authentication is required to connect through a firewall.
- **Firewall Port**: The TCP port for the firewall FirewallServer — see the description of the FirewallServer option for details.
- **Firewall Server**: Specify a firewall name or IP address to authenticate requested connections, if necessary.
- **Firewall Type**: The type of firewall to connect through.
- **Firewall User**: The user name to authenticate with the firewall.

- **Initiate OAuth:** Set this property to initiate the process to obtain or refresh the OAuth access token when you connect.
- **Location:** A path to the directory that contains the schema files defining tables, views, and stored procedures.
- **Logfile:** A path to the log file.
- **Max Log File Size:** A string specifying the maximum size in bytes for a log file (ex: 10MB). When the limit is hit, a new log is created in the same folder with the date and time appended to the end.
- **MaximumRequestRetries:** The maximum amount of request retries.
- **OAuth Access Token:** The access token for connecting using OAuth.
- **OAuth Client Id:** The client Id assigned when you register your application with an OAuth authorization server.
- **OAuth Client Secret:** The client secret assigned when you register your application with an OAuth authorization server.
- **OAuth Settings Location:** The location of the settings file where OAuth values are saved when InitiateOAuth is set to GETANDREFRESH or REFRESH.
- **Offline:** Use offline mode to get the data from the cache instead of the live source.
- **Other:** The other parameters necessary to connect to a data source, such as username and password, when applicable.
- **Pagesize:** The maximum number of results to return per page from Facebook.
- **Proxy Auth: Scheme** The proxy server authorization scheme (default: BASIC).
- **Proxy Auto Detect:** This indicates whether to use the default system proxy settings or not. Set ProxyAutoDetect to FALSE to use custom proxy settings. This takes precedence over other proxy settings.
- **Proxy Password:** A password, if authenticating with a proxy server.
- **Proxy Port:** The TCP port for the proxy ProxyServer (default: 80).
- **Proxy Server:** If a proxy server is given, then the HTTP request is sent to the proxy instead of the specified server.
- **Proxy SSL Type:** The SSL type to use when connecting to the proxy server (default: AUTO).
- **Proxy User:** A user name, if authentication is to be used for the proxy.
- **Pseudo Columns:** Indicates whether or not to include pseudo columns as columns to the table.
- **Query Cache:** The time duration, in seconds, for which the in-memory cached data is reused.

- **Readonly:** You can use this property to enforce read-only access to Facebook from the provider.
- **RetryWaitTime:** The amount of time in milliseconds to wait before retrying to query the Facebook API.
- **RTK:** The runtime key used for licensing.
- **Search Terms:** Default SearchTerms if none are specified. Used for some tables, such as Users, where SearchTerms may be specified.
- **SSL Server Cert:** The certificate to be accepted from the server when connecting using TLS/SSL.
- **Support Enhanced SQL:** If set to true, the provider will cache the entire table in memory and issue the current query against the memory cache. Allows queries with aggregate functions that are otherwise unsupported. The entire memory cache is discarded after the query is executed, unlike QueryCache.
- **Tables:** Comma-separated list of tables to be listed. For example: Tables=TableA,TableB,TableC.
- **Target:** A default target if none is specified. Used for some tables, such as Comments, where a target may be specified.
- **ThrottleSeconds:** Indicates how many seconds to wait until sending another insight request when the ThrottleMaxPercent has been met.
- **ThrottleMaxPercent:** The threshold set for throttling further insight requests. The defaults for ThrottleSeconds and ThrottleMaxPercent are 10 and 95.
- **Timeout:** The value in seconds until the timeout error is thrown, canceling the operation.
- **Verbosity:** The verbosity level that determines the amount of detail included in the log file.
- **Version:** The Facebook Graph API version to use.

You can see your version in the Facebook Developer dashboard. See the Facebook Platform Changelog to see if your version is still supported.

## Switching between Accounts

When you connect to the Facebook data source, Sisense saves your OAuth values in the file OAuthsettings.txt file located at

...\Users\xxx\AppData\Roaming\CData\Facebook Data Provider on your

Sisense server. To connect to the Facebook data source with another user on the

same machine, you must delete the OAuthsettings.txt file. Sisense will then generate a new file for that user.

Another option to support multiple users is to define the location and file name of an OAuthsettings file for each unique user in your connection string through the OAuthSettingsLocation parameter. When each user connects to the data source, Sisense generates the OAuth file with the file name you specify in the location you define. In the examples below, two users are allowed to access the Facebook data source and for each user, Sisense generates a file that contains that user's OAuth values in the location defined in the string.

```
jdbc:facebook:OAuthSettingsLocation=C:\facebook\auth\john.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
jdbc:facebook:OAuthSettingsLocation=C:\facebook\auth\sally.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
```

In the example above, two OAuth files are created, one for John and one for Sally in the location C:\facebook\auth\.

This is useful if you support many users who each need to access the Facebook data source.

## Changing Facebook Connectivity Settings

Prior to Sisense V6.6.1, Facebook Ads tables were part of the Facebook Connector. Sisense has created a new Facebook Ads connector that supports tables specific to Facebook Ads. These tables include:

AdAccounts, AdCreatives, Ads, AdScheduledReports, AdSets, Campaigns, AdInsights, AdInsightsActions

You can continue to use earlier versions of the Facebook connector to connect to these tables, but this may require custom SQL queries. To connect to these tables quickly and easily without any custom SQL, you need to install the Facebook Ads connector and change the connectivity settings to direct these tables to the Facebook Ads connector.

For more information on how to change connectivity settings after installing the Facebook Ads connector, see [Changing Connectivity Settings](#).

### Accumulative Builds

Sisense support accumulative builds for all numeric and dates data types. However, the data must be sorted before building the ElastiCube.

### Versions

Version	Release Date	Improvements
6.6.1.113	June 12, 2017	New version of Graph API v2.9 Changelog New views added: InsightsByReactionTotals, SimpleVideoInsights, VideoInsightsByActionType, VideoInsightsByDistributionType, VideoInsightsByReactionType New insights added for the InsightsByConsumptionType view: POST_CONSUMPTIONS_BY_TYPE, POST_CONSUMPTIONS_BY_TYPE_UNIQUE For more information about Version 6.6.1.113 including Breaking Changes, see Version 6.6.1.113 below.
6.6.0.74	March 16, 2017	JSON columns in the AdInsights table are available for querying in AdInsightActions by using the ActionCollection column

Version	Release Date	Improvements
6.6.1.65	March 9, 2017	In some scenarios, getMetaData() returned unexpected results
6.6.0.52	March 6, 2017	Dynamic Targets New table: GroupMemberShips

### Version 6.6.1.113

#### Breaking Changes

In the latest version of the Facebook Graph API, permissions have been largely reworked. All of the friends\* permissions have been removed. In addition, most permissions require Facebook approval if you wish to request information about other user accounts who are not an administrator of the Facebook App. Sisense recommends reviewing the new permissions listed by Facebook here:

<https://developers.facebook.com/docs/facebook-login/permissions/v2.1>

In addition, users may reject requested permissions in the Authorization dialog on Facebook. If this happens, the specific Permissions requested will state that they are declined under the Status column. The permission may be rerequested by executing GetOAuthAccessToken and setting AuthType to rerequest.

If you want to keep using an old version of the API, you can by simply setting the Version connection property in the connection string to the specific version you were using and also setting the Location to a folder where you have saved an older copy of the schema files (.rsd and .rsb files).

If you want to update your app to 2.9, then you don't have to do anything

**Note:** Sisense removed few metrics from the AdInsights / AdInsightsActions due to their removal by Facebook in 2.9.

A number of various changes have been made across tables due to changes in Facebook. These changes are listed below:

- Facebook no longer allows public posts to be searched. We have had to remove SearchTerms from the Posts table because of this.
- Checkins have been removed from the Graph API. Due to this, we have removed the Checkins table.
- Events may no longer be inserted, updated, or deleted via the Graph API. We have had to change the Events table to a view.
- Usernames for Facebook Users may no longer be used as a target and are no longer returned. The Facebook username for Pages is unaffected.
- User ID's for Facebook Users will be different across different Facebook Apps. This is due to the new policy of App-scoped User ID's.

# Connecting to Gmail

The Sisense Gmail connector is a standalone connector that allows you to import data from Gmail's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the Gmail APIs. To obtain a connection string, you will need to register a Gmail app.

Once you have connected to Gmail, you can import a variety of tables from the Gmail API.

This page describes how to install the Gmail connector, how to connect to Gmail with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the Gmail Connector](#)
- [Connecting to the Gmail REST API](#)
- [Adding Gmail Tables to your ElastiCube](#)
- [Gmail Tables](#)
- [Limitations](#)

## Installing the Gmail Connector

Sisense provides the Gmail connector as a standalone connector that you can download and add to your list of default Sisense connectors.

### To install the Gmail connector:

1. [Download](#) the Gmail installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The Gmail connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.

# Connecting to the Gmail REST API

To access Gmail's REST API from the ElastiCube Manager, you must provide valid OAuth Gmail credentials through a connection string. These credentials are provided by Gmail when you register an application.

After you receive your credentials from Gmail, you can create the connection string and provide Sisense with it to connect to your data.

## Registering an App

Follow the steps below to obtain the OAuth client credentials, the OAuthClientId and OAuthClientSecret necessary for your connection string:

1. Log into the Google API Console.
2. Click **Create Project** or select an existing project.
3. In the API Manager, click **Credentials > Create Credentials > OAuth Client Id**.
4. Select the application type. If you are making a Web application, select Web application. In the Authorized Redirect URIs box, enter the URL you want to be used as a trusted redirect URL, where the user will return with the token that verifies that they have granted your app access. If you are making a desktop application, select Other.
5. Click **Create**. The OAuthClientId and OAuthClientSecret are displayed.
6. Click **Library > Gmail API > Enable API**.

## Adding Gmail Tables to your ElastiCube

Sisense uses connection strings to connect to Gmail and import data into the ElastiCube Manager.

The connection string to connect to Gmail has the following structure:

```
jdbc:Gmail:Property1=Value1;Property2=Value2;
```

The following is an example of a Gmail connection string:

```
jdbc:gmail:user=xxxx@xxxxx.com;InitiateOAuth=GETANDREFRESH;OAuthClientId
=xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx;OAuthClientSecret=xxxxxxxxxxxxxxxx
xxxxxxxxxxxxxxxx;
```

The connection string should include the following information:

- **InitiateOAuth:** Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken connection property.
- **OAuthClientId:** Set this to the Client ID in your app settings.
- **OAuthClientSecret:** Set this to the Client secret in your app settings.
- **User:** Set this to the user of the Gmail account.

When you connect the driver opens the OAuth endpoint in your default browser. Log in and grant permissions to the application.

**Note:** To switch between accounts, you need to delete the file OAuthsettings.txt file located at ...\\Users\\xxx\\AppData\\Roaming\\CData\\Gmail Data Provider.

### To add Gmail data:

1. In ElastiCube Manager, click **Add Data** and then, Gmail. The Connect to Gmail window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. Gmail is displayed in the Select Database list.
4. Click **OK**. Sisense connects to Gmail and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

## Switching between Accounts

When you connect to the Gmail data source, Sisense saves your OAuth values in the file OAuthsettings.txt file located at

...\\Users\\xxx\\AppData\\Roaming\\CData\\Gmail Data Provider on your Sisense server. To connect to the Gmail data source with another user on the same

machine, you must delete the OAuthsettings.txt file. Sisense will then generate a new file for that user.

Another option to support multiple users is to define the location and file name of an OAuthsettings file for each unique user in your connection string through the **OAuthSettingsLocation** parameter. When each user connects to the data source, Sisense generates the OAuth file with the file name you specify in the location you define. In the examples below, two users are allowed to access the Gmail data source and for each user, Sisense generates a file that contains that user's OAuth values in the location defined in the string.

```
jdbc:Gmail:OAuthSettingsLocation=C:\Gmail\auth\john.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
jdbc:Gmail:OAuthSettingsLocation=C:\Gmail\auth\sally.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
```

In the example above, two OAuth files are created, one for John and one for Sally in the location C:\Gmail\auth\.

This is useful if you support many users who each need to access the Gmail data source.

## Gmail Tables

Gmail's RESTful APIs expose the following Gmail tables that you can import into the ElastiCube Manager through the Sisense Gmail connector:

- Gmail/All Mail
- Gmail/Drafts
- Gmail/Important
- Gmail/Sent Mail

- Gmail/Spam
- Gmail/Starred
- Gmail/Trash
- Inbox
- Test Label

## Limitations

- The default Gmail table names contain a backslash; some also contain spaces. You can use square brackets to escape the table names for the default Gmail folders.
- For example: `SELECT * FROM [Gmail/All Mail]` The default number of messages returned is 100. This can be changed by setting either `LIMIT` or `MaxItems`. If you wish to return all mail within a mailbox, specify a value of 0 or -1.
- Aggregate functions not supported.

# Connecting to Google Adwords

The ElastiCube Manager enables easy and quick access to tables contained within Google Adwords. The steps below describe how to connect to this type of data source.

**Note:** Importing data from Google Adwords is performed using the Sisense developer key and Sisense pays for the data imported from AdWords. For this reason your AdWords account must have access to the 'Campaign Service' and the user needs to have Standard Access. To check the type of access, log into your AdWords account, then go to My Account > Account Access. A list of all associated campaigns and permission levels will be displayed. Increasing the permission level can only be performed by a campaign administrator using the "invite user functionality". See [Campaign Management and Users](#) for more details. In addition, you will need an MCC account. For more information, [click here](#).

1. Click **Add Data** in the top menu of the ElastiCube Manager.
2. Under the **Web Services** category, select **Google AdWords**.
3. You will be prompted to enter the following information:
  - **Email:** Enter email associated with your Google Analytics account.
  - **Auth Code:** Google requires an authentication code to enable access from the ElastiCube to data associated with a Google Analytics account (site name and profile IDs). To attain the Auth code click on the link **Get Auth Code**. The browser will open and you will need to sign in to your Google Analytics account. After signing in, click on the 'Allow access' button. This will generate a temporary code that you must paste into the **Auth Code** text box in the connection window.
  - **Customer ID:** Enter your Google AdWords Customer ID. Your Customer ID appears at the top corner of your AdWords screen (after logging into your AdWords account).adwords1
4. Click **Connect**.

Connect to Google Adwords

Email

Auth Code  [Get Auth Code](#)

Customer ID

Date Range ⓘ

From

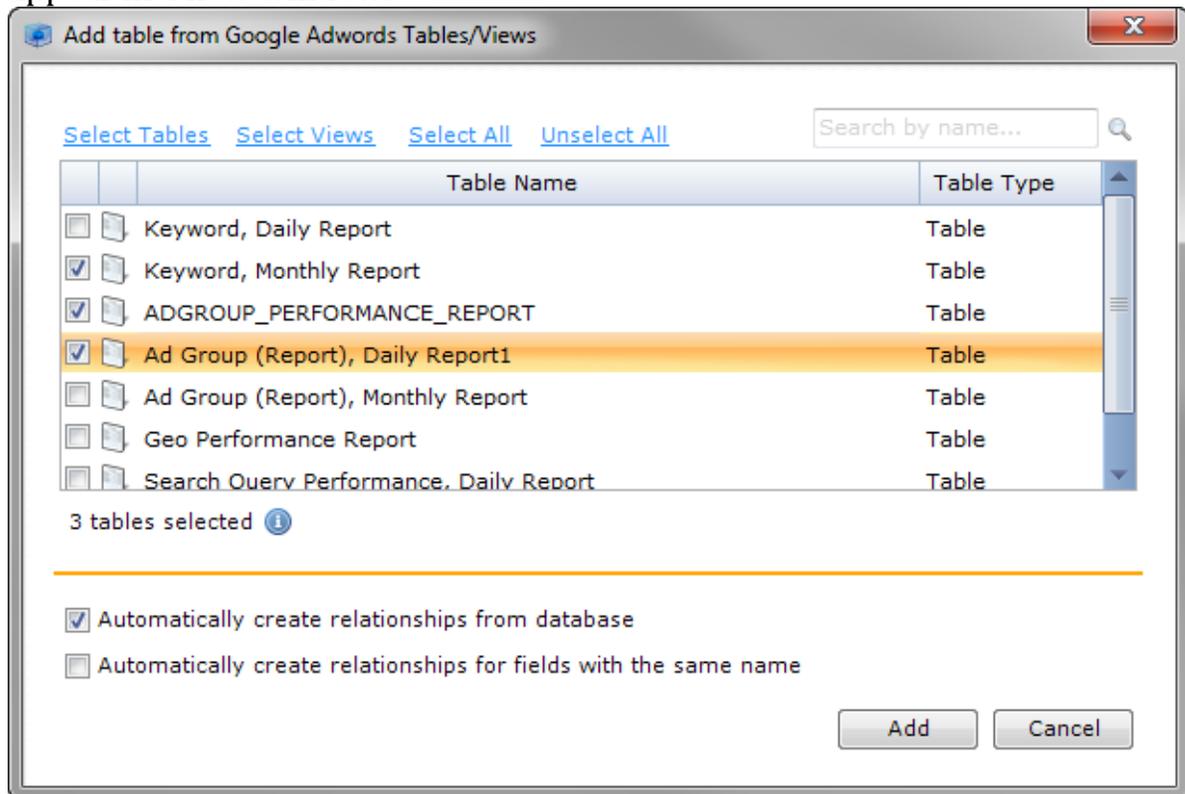
To

[Check All](#) [Uncheck All](#)

Email	Site	Customer
-------	------	----------

5. Select the relevant client email associated with Adwords you want to work with and click **OK**. All tables and views associated with the website will

appear in a new window.



6. Select the checkbox next to each table you want to use.
7. Once all relevant tables are selected, click Add.

**Note about possible build failures:** If you upgrade your version of Sisense to 5.8.1, and the build fails, it may be due to the fact that the newer Sisense version is using a newer version of the Google AdWords API (v201506). If this is the case, then this API version includes changes to some of the AdWords predefined columns. Consequently, it is required for you to click refresh in the 'Change Connectivity Settings...' window next to your AdWords data source and rebuild your ElastiCube. Unless this action is taken, your ElastiCube build will fail due to the change in Google AdWords schema.

See also Customizing Google Analytics and AdWords Tables.

## Versions

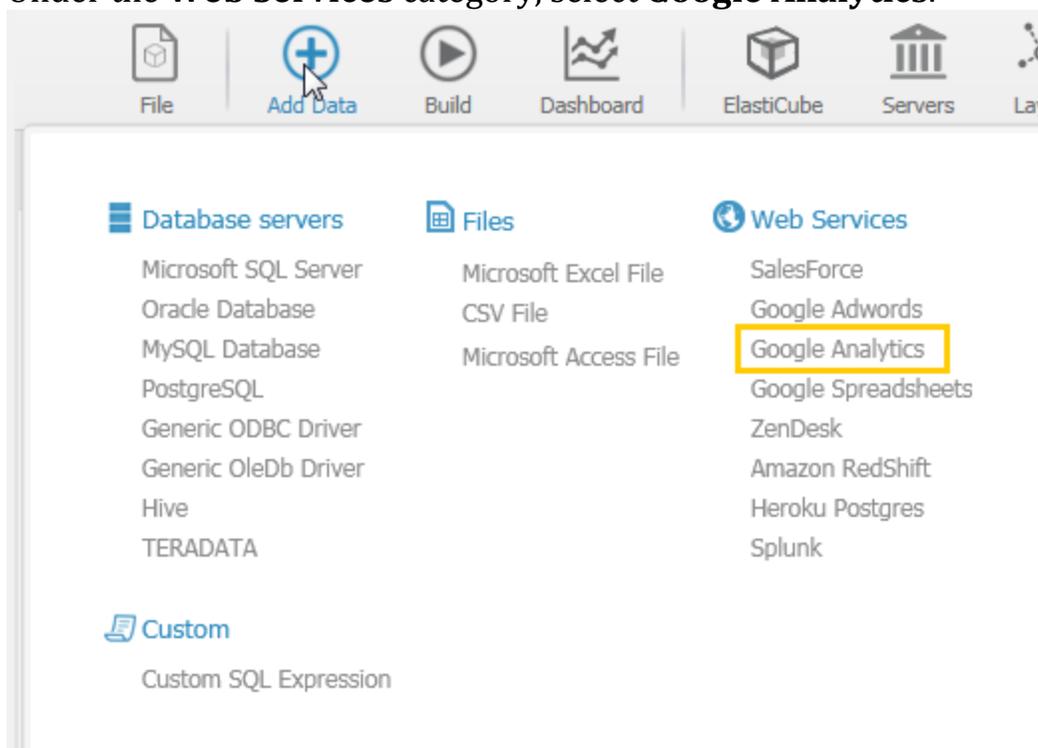
Version	Release Date	Description
6.5.1.15	April 19, 2017	Google Adwords API v201607 supported by Sisense v6.5.1 and earlier has been deprecated

# Connecting to Google Analytics

The ElastiCube Manager enables easy and quick access to tables contained within Google Analytics. The steps below describe how to connect to this type of data source.

**Note:** Google Analytics provides a quota on the free daily API requests. For this reason a developer key must be supplied. Read more details about [Quota Limits](#).

1. Click **Add Data** in the top menu of the ElastiCube Manager.
2. Under the **Web Services** category, select **Google Analytics**.



3. You will be prompted to enter the following information:
  - **Email:** Enter email associated with your Google Analytics account.
  - **Auth Code:** Google requires an authentication code to enable access from the ElastiCube to data associated with a Google Analytics account (site name and profile IDs). To attain the Auth code click on the link

**Get Auth Code.** The browser will open and you will need to sign in to your Google Analytics account. After signing in, click on the 'Allow access' button. This will generate a temporary code that you must paste into the **Auth Code** text box in the connection window.

- **Date Range:** Select the date range of the analytics data you require.

The screenshot shows a dialog box titled "Connect to Google Analytics". It features an "Email" dropdown menu, an "Auth Code" text input field, and a "Get Auth Code" link. A "Connect" button is positioned below the "Auth Code" field. The "Date Range" section includes "From" and "To" date pickers, with "From" set to "Thursday, January 01, 2015" and "To" set to "Monday, December 31, 2029". There are also "Check All" and "Uncheck All" links. At the bottom is a table with a header "Site" and an empty body. At the very bottom are "OK" and "Cancel" buttons.

4. Click **Connect** to display a list of available websites the list box below.
5. Select the relevant website that you want to work with, and click **OK**.
6. All tables and views associated with the website appear in a new window.
7. Select the checkbox next to each table you want to use. Existing relationships between tables can be automatically replicated in the ElastiCube by selecting the **Automatically create relationships from database** checkbox. Likewise fields with similar names can be linked by selecting the **Automatically create relationships for fields with the same name** checkbox.
8. Once all relevant tables are selected, click **Add**.

See [Customizing Google Analytics and AdWords Tables](#) for in-depth info on controlling your Google Analytics tables.

## Avoiding Sampling in Google Analytics

Google Analytics applies sampling whenever more than 500,000 sessions are returned within a given query. For example, if we are asking for a 1-month time period that had 1,000,000 session in it, Google will apply sampling and we will only get 50% of the actual sessions. Whereas in some cases sampling might only slightly affect the results accuracy, in many cases it is crucial.

Sisense supports three configuration options to reduce the chances of sampling: While these configuration options may reduce the chances that Google samples your data, Sisense cannot guarantee that your data will not be sampled as the data is extracted from Google.

1. **Split Intervals:** A Google Analytics table can be configured to be split into a number of given intervals. For example, if intervals=4, then a one month period will be split into 4 query requests equal in days.
2. **Day Intervals:** A Google Analytics table can be configured to be split into multiple intervals where each is the size of the given number of days. For example, if day\_intervals=10, then a 30 day time period will be split into 3 intervals.
3. **Segmenting Data:** A segment is a subset of your data. You can use predefined segments to work around Google's sampling limitations. For more information, see [Segment Google Analytics Data](#).

Since Google limitations are applied separately per request, this reduces the risk of overall sampling.

### Configuring Google Analytics Sampling Intervals

Edit the following XML file (if it doesn't exist, create it):

```
%appdata%\Sisense\PrismServer\LocalRepository2.0\Resources\nAnalytics\GAnalytics.Tables.xml
```

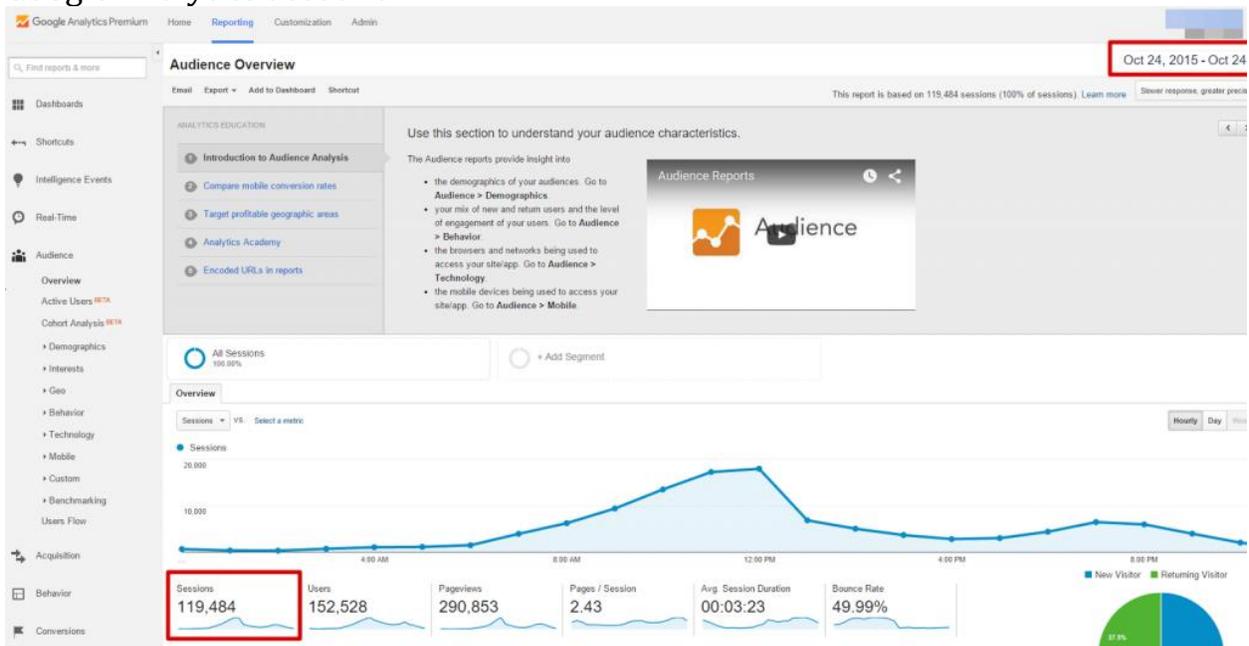
The xml file needs to have a table node for each table, under a tables root:

```
<?xml version="1.0" encoding="utf-8" ?>
<tables>
<table name="MyGoogleAnalyticsTable" days_interval="2"
split_by="10"/>
</tables>
```

For Split Intervals as described above use the “**split\_by**” attribute. For Day Intervals use the “**days\_interval**” attribute.

Note that only one of these attributes should be used. “split\_by” will be used in case both are available.

It is easy to estimate the amount of sessions your account generates from your Google Analytics account:



# Connecting to Google BigQuery

The ElastiCube Manager provides you with quick and easy access to Google BigQuery tables enabling you to generate and analyze your data.

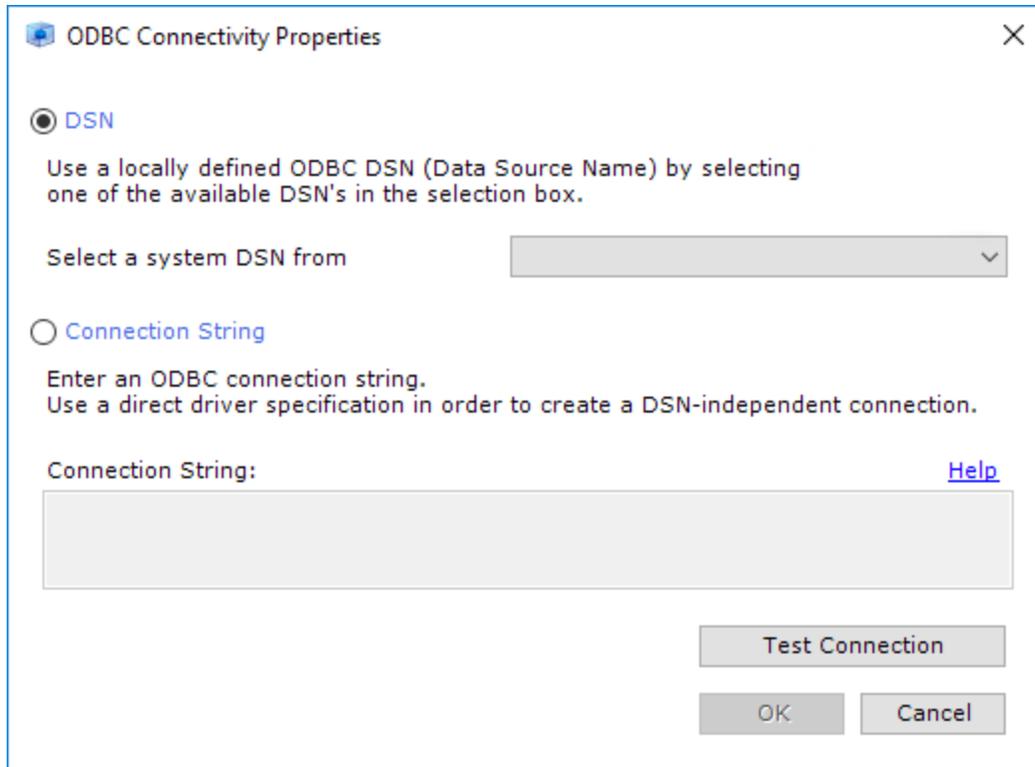
There are two connection types you can configure when connecting to your BigQuery tables, DSN and Connection String (DSN-Less). For more information about the differences between the connection types, click [here](#).

To connect to your BigQuery tables, follow these steps:

1. [Download the Sisense BigQuery ODBC driver and connect the data source in the ElastiCube Manager](#)
2. [Add a DSN](#)
3. [\(Optional\) Configure advanced DSN configurations](#)
4. [\(Optional\) Configure logging settings](#)
5. [Add the BigQuery tables to your project](#)

## Downloading and Connecting the Sisense Google BigQuery ODBC Driver

1. [Download](#) and install the Sisense BigQuery ODBC Driver.
2. In ElastiCube Manager, click **Add Data** and then, **Generic ODBC Driver**.
3. Select **DSN**. If a DSN file has already been created, select it in the dropdown list, and click **Test Connection**. If your connection has been set up, you can click **OK** and review the tables generated by the Sisense BigQuery ODBC driver.



OR

Select Connection String (DSN-Less) and enter your Connection String. Click **Test Connection**. For more information, see [Connecting without a DSN](#).

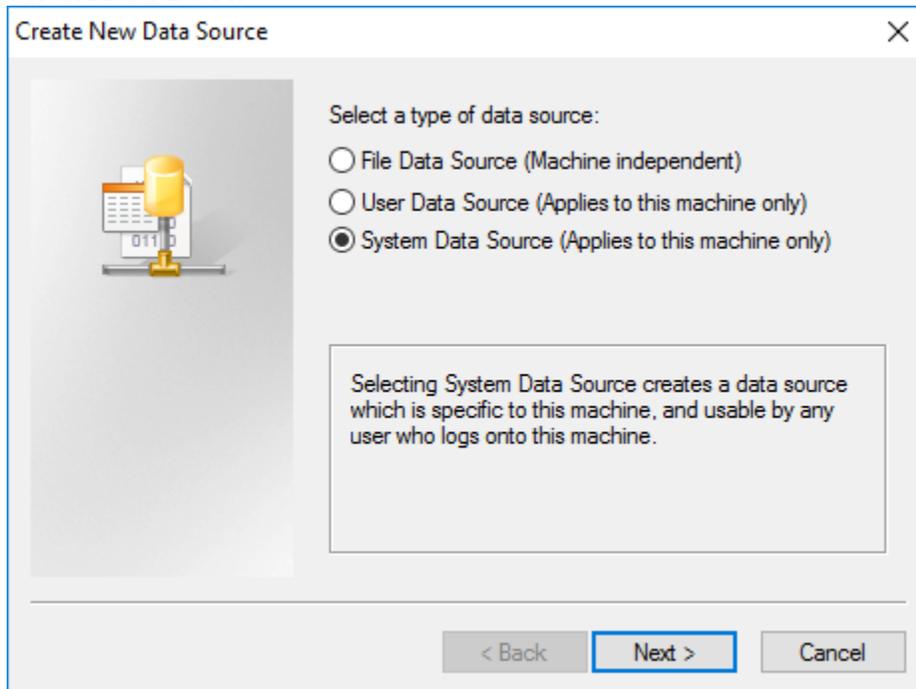
## Adding a DSN

If a DSN has not yet been created, you must add the DSN manually.

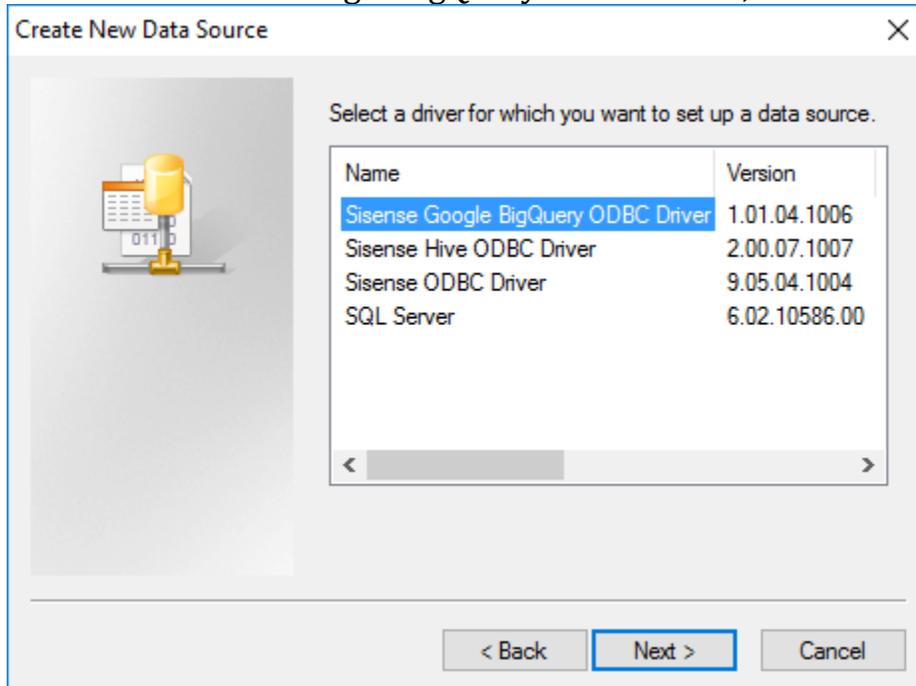
**Note:** To add a DSN you must run the ODBC Data Source Administrator (64-bit) as a Windows administrator.

**To add a DSN:**

1. Click **Add DSN**.



2. Select the **System Data Source** option. The created file applies to all users in a specific machine only. Click **Next**.
3. Select the Sisense Google BigQuery ODBC Driver, and click **Next**.



4. Click **Finish**. The ODBC Driver for Google BigQuery DSN Setup dialog box opens.
5. Enter the following information:

In the Data Source Name field, type a name for your DSN.

Optionally, in the Description field, type relevant details about the DSN.

Configure authentication using one of the following methods:

**To authenticate the connection as a service:**

1. From the OAuth Mechanism list, select Service Authentication.
2. In the **Email** field, type your service account email ID.
3. In the **Key File Path** field, type the path to the .p12 key file that is used to authenticate the service account ID.

**Or, to authenticate the connection as a user:**

1. From the **OAuth Mechanism** list, select **User Authentication**.
2. Click **Sign In**.
3. In the browser that opens, type your credentials for accessing your BigQuery data, and then click **Sign In**.
4. When prompted to allow BigQuery Client Tools to access your data in Google BigQuery, click **Accept**.
5. Copy the code that Google provides, and then paste the code in the **Confirmation Code** field in the Simba ODBC Driver for Google BigQuery DSN Setup dialog box.
6. When you paste the confirmation code, the **Refresh Token** populates automatically when you select the field. The refresh token is used whenever the driver needs to access your BigQuery data. You can save the refresh token in the DSN so that you only need to generate it once.  
**Note:** A confirmation code can only be used once. You must get a new confirmation code from Google whenever you need another refresh token.
6. To allow the driver to access Google Drive so that it can support federated tables that combine BigQuery data with data from Google Drive, select the **Request Google Drive Scope Access** check box.
7. Choose one:
  - To verify the server using the trusted CA certificates from a specific .pem file, specify the full path to the file in the **Trusted Certificates** field and leave the **Use System Trust Store** check box cleared.
  - Or, to use the trusted CA certificates .pem file that is installed with the driver, leave the default value in the **Trusted Certificates** field and the **Use System Trust Store** check box cleared.

- Or, to use the Windows Trust Store, select the Use System Trust Store check box and leave the **Trusted Certificates** field cleared.
8. In the **Catalog (Project)** list, select the name of your BigQuery project. This project is the default project that the ODBC Driver for Google BigQuery queries against, and also the project that is billed for queries that are run using the DSN.
  9. (Optional) To configure advanced driver options, click **Advanced Options**. For more information, see [Advanced Options](#).
  10. (Optional) To configure logging behavior for the driver, click **Logging Options**. For more information, see [Logging Options](#).
  11. Click Test. A pop window is displayed that indicates the connection was successful.
  12. Click **OK**.
  13. In the ODBC Connectivity Properties window, click **Test Connection**.
  14. Click **OK** to save and close the ODBC Data Source Administrator.

## Advanced Options

You can configure advanced driver by clicking **Advanced Options** in the Sisense BigQuery ODBC Driver DSN Setup window. This displays the Advanced Options window from which you can configure the following options:

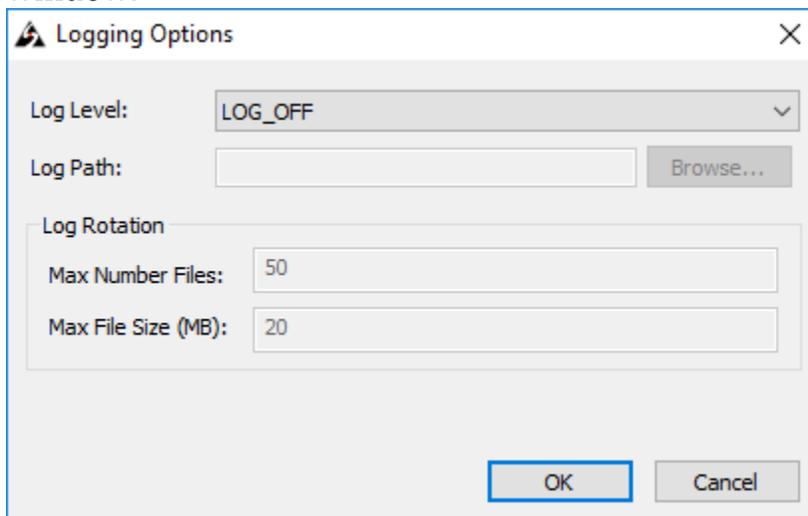
Field	Description
Max Requests	Enter the maximum number of requests that can be made per second.

Field	Description
Per Second (0=unlimited)	<b>Note:</b> To allow an unlimited number of requests per second with no throttling, type 0.
Rows Per Block	Enter the maximum number of rows to fetch for each data request.
Default String Column Length	Enter the maximum data length for String columns.
Dataset Name For Large Result Sets	Enter the name of the BigQuery dataset to use to store temporary tables. <b>Note:</b> The dataset created from the default ID is hidden.
Temporary Table Expiration Time(ms)	Enter the length of time (in milliseconds) that the temporary table exists for.
Use Native Query	Select the Use Native Query checkbox to disable the SQL Connector feature and allow the driver to execute BigQuery SQL directly. When Use Native Query is selected, the driver does not transform the queries emitted by an application, so the native query is used. When Use Native Query is disabled, the driver transforms the queries emitted by an application and converts them into an equivalent form in BigQuery SQL. <b>Note:</b> If an application already emits BigQuery SQL, then enable this option to avoid the extra overhead of query transformation. (Optional)
Enable SQLPrepare Metadata with Native Query	To enable the driver to retrieve metadata during the prepare stage of a query when working in Native Query mode, select the Enable SQLPrepare Metadata with Native Query checkbox. <b>Note:</b> This option is available only if the Use Native Query checkbox is

Field	Description
(slower)	selected.
Allow Large Result Sets	Select to to allow query results larger than 128MB in size.After you select this checkbox, the Dataset Name For Large Result Sets and Temporary Table Expiration Time(ms) fields are enabled.

## Logging Options

You can configure logging options by clicking **Logging Options** in the Sisense BigQuery ODBC Driver DSN Setup window. This displays the Logging Options window.



The ODBC Data Source Administrator provides tracing functionality, which you can activate to help troubleshoot issues.

**Important:** Only enable logging long enough to capture an issue. Logging decreases performance and can consume a large quantity of disk space.

The driver allows you to set the amount of detail included in log files. The table below lists the logging levels provided by the Sisense ODBC Driver with SQL Connector for Google BigQuery, in order from least verbose to most verbose.

In the Logging Options window, you can configure the following information:

Field	Description
Log Level	<p>Select the Log Level. There are seven possible options:</p> <p><b>LOG_OFF:</b> Disables all logging.</p> <p><b>LOG_FATAL:</b> Logs very severe error events that lead the driver to abort.</p> <p><b>LOG_ERROR:</b> Logs error events that might still allow the driver to continue running.</p> <p><b>LOG_WARNING:</b>Logs potentially harmful situations.</p> <p><b>LOG_INFO:</b> Logs general information that describes the progress of the driver.</p> <p><b>LOG_DEBUG:</b> Logs detailed information that is useful for debugging the driver.</p> <p><b>LOG_TRACE:</b> Logs more detailed information than the DEBUG level.</p>
Log Path	<p>Enter the full path to the folder where you want to save log files.</p> <p>OR</p> <p>Click <b>Browse</b> and select the folder where you want to save log files.</p>
Log Rotation	<p>Enter the maximum number of log files to keep in the Max Number Files field.</p> <p><b>Note:</b> After the maximum number of log files is reached, each time an additional file is created, the driver deletes the oldest log file.</p> <p>Enter the maximum size of each log file in megabytes (MB) in the Max File Size field.</p> <p><b>Note:</b> After the maximum file size is reached, the driver creates a new file and continues logging.</p>

## Connecting without a DSN

When you choose to connect without a DSN, you must define the connection string that sets the configuration options.

The following is an example connection string that sets advanced options:

```

DSN=Sample Sisense BigQuery DSN; Catalog=public_data;
MaxRequestsPerSecond=20; UseNativeQuery=0;
AllowLargeResults=0;
LargeResultsDataSetId=_bqodbc_temp_tables;
LargeResultsTempTableExpirationTime=3600000

```

The table below provides a list of possible keys you can add to your string and their descriptions:

Key Name	Mandatory	Description
AllowLargeResults	No	The value of this key is 0 or 1.0: Disabled. The driver returns an error when query results are larger than 128MB in size.1: Enabled. The driver allows query results that are larger than 128MB in size.
ExecCatalog	Yes	The name of the project to bill for queries that are run using the DSN.
Project or Catalog	Yes	The name of your BigQuery project.
LargeResultDataSetId	Yes*	The ID of the BigQuery dataset to use to store temporary tables. <b>Note:</b> This option is available only when the Allow Large Result Sets option is enabled (the AllowLargeResults key is set to 1). The dataset created from the default ID is hidden. *Required if AllowLargeResults is enabled.
DefaultStringColumnLength	No	The maximum data length for STRING columns.

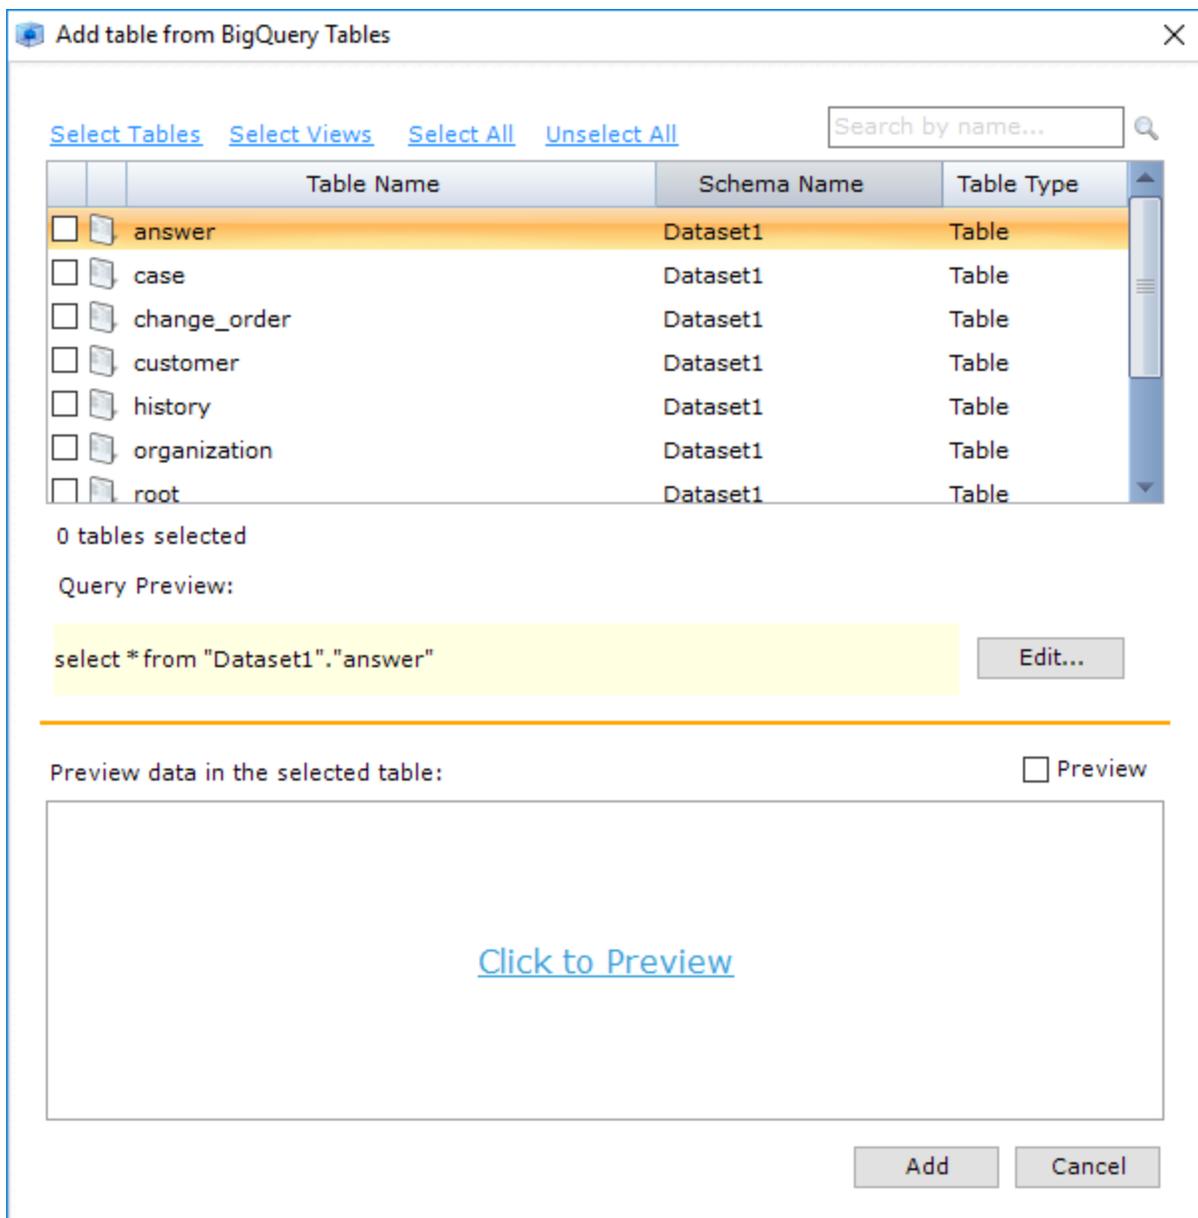
Key Name	Mandatory	Description
UseNativeQueryMetadata	No	The value of this key is 0 or 1. <b>0:</b> Disabled. The driver does not retrieve metadata during the prepare stage. <b>1:</b> Enabled. The driver retrieves metadata during the prepare stage of a query when working in Native Query mode. <b>Note:</b> This option is available only when UseNativeQuery key is set to 1.
MaxRequestsPerSecond	No	The maximum number of requests that can be made per second. To allow an unlimited number of requests per second with no throttling, set the value to 0.
ProxyHost	No	The IP address or hostname of the proxy server. If this key is not set for the Linux or Mac OS X driver, then support for connecting to proxy servers is disabled.
ProxyPwd	Yes*	The password corresponding to the user name provided in the Proxy Username field (the ProxyUid key). *Required if connecting to a proxy server.
ProxyPort	Yes*	The number of the port on which the proxy server is listening. *Required if connecting to a proxy server.
ProxyUid	Yes*	The user name that you use to access the

Key Name	Mandatory	Description
		proxy server. *Required if connecting to a proxy server.
RefreshToken	Yes	The refresh token that you obtain from Google for authorizing access to BigQuery. When you configure a DSN with the Windows driver, the refresh token is generated automatically after you provide the confirmation code. When you configure a DSN with the Linux or Mac OS X versions of the driver, you must use the Refresh Token configuration tool to generate the token.
RowsFetchedPerBlock	No	The maximum number of rows that the driver can fetch for each data request.
LargeResultsTempTableExpirationTime	Yes*	The length of time (in milliseconds) for which a temporary table exists. <b>Note:</b> This option is available only when AllowLargeResults key is set to 1. The default value is 1 hour in milliseconds.*Required if AllowLargeResults is enabled.
UseNativeQuery	No	The value of this key is 0 or 1. <b>0:</b> Disabled. The driver transforms the queries emitted by an application and converts them into an equivalent form in

Key Name	Mandatory	Description
		<p>BigQuery SQL.</p> <p><b>1:</b> Enabled. The driver does not transform the queries emitted by an application, so the native query is used.</p> <p><b>Note:</b> If an application already emits BigQuery SQL, then enable this option to avoid the extra overhead of query transformation.</p>

## Adding BigQuery Tables to your Project

After setting up the DSN or the DSN-less connection, the Add table from BigQuery Tables window is displayed.



From this window, you add your BigQuery tables to your ElastiCube Manager. In addition, you can view the SQL syntax in the Query Preview section and click Edit to customize it.

### To add BigQuery Tables to your Project:

1. Connect to Google BigQuery via the Sisense ODBC Tool.
2. Add your data source.
3. In the Add table from BigQuery Tables window, select the table you want to add to the ElastiCube Manager.
4. Click **Add**. The selected tables are added to your ElastiCube Manager.

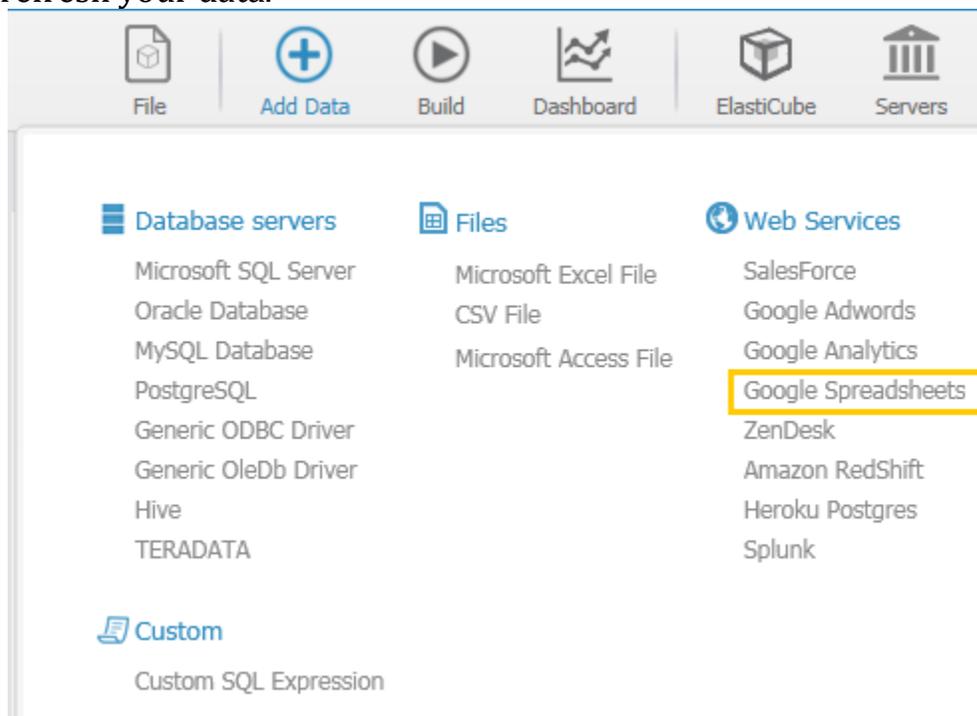
# Connecting to Google Spreadsheets

The ElastiCube Manager enables easy and quick access to tables contained within Google Spreadsheets.

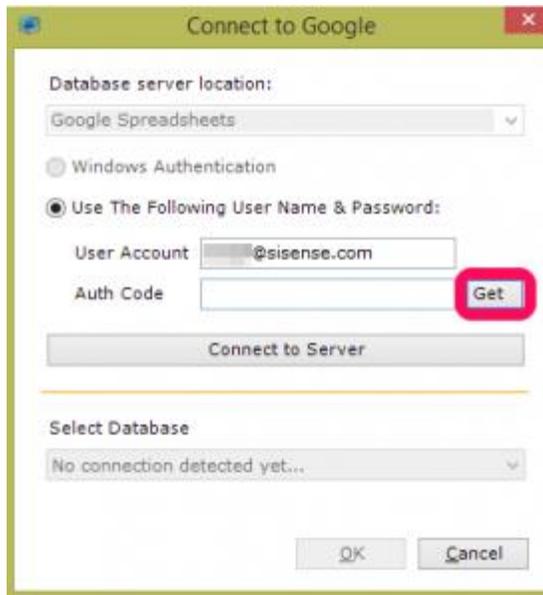
**Note:** To work with Google Spreadsheets, you must have a professional Google account, since Google's API is only available to business accounts.

## To connect to Google Spreadsheets:

1. Click the **Add Data** icon in the top menu of the ElastiCube Manager.
2. Under the **Web Services** category, select Google Spreadsheets.
3. Note: Make sure you are working with Sisense version 5.7.6.71 or above. If you are working with a previous version and cannot connect to Google Spreadsheets, [click here](#) to upgrade your Sisense version and learn how to refresh your data.



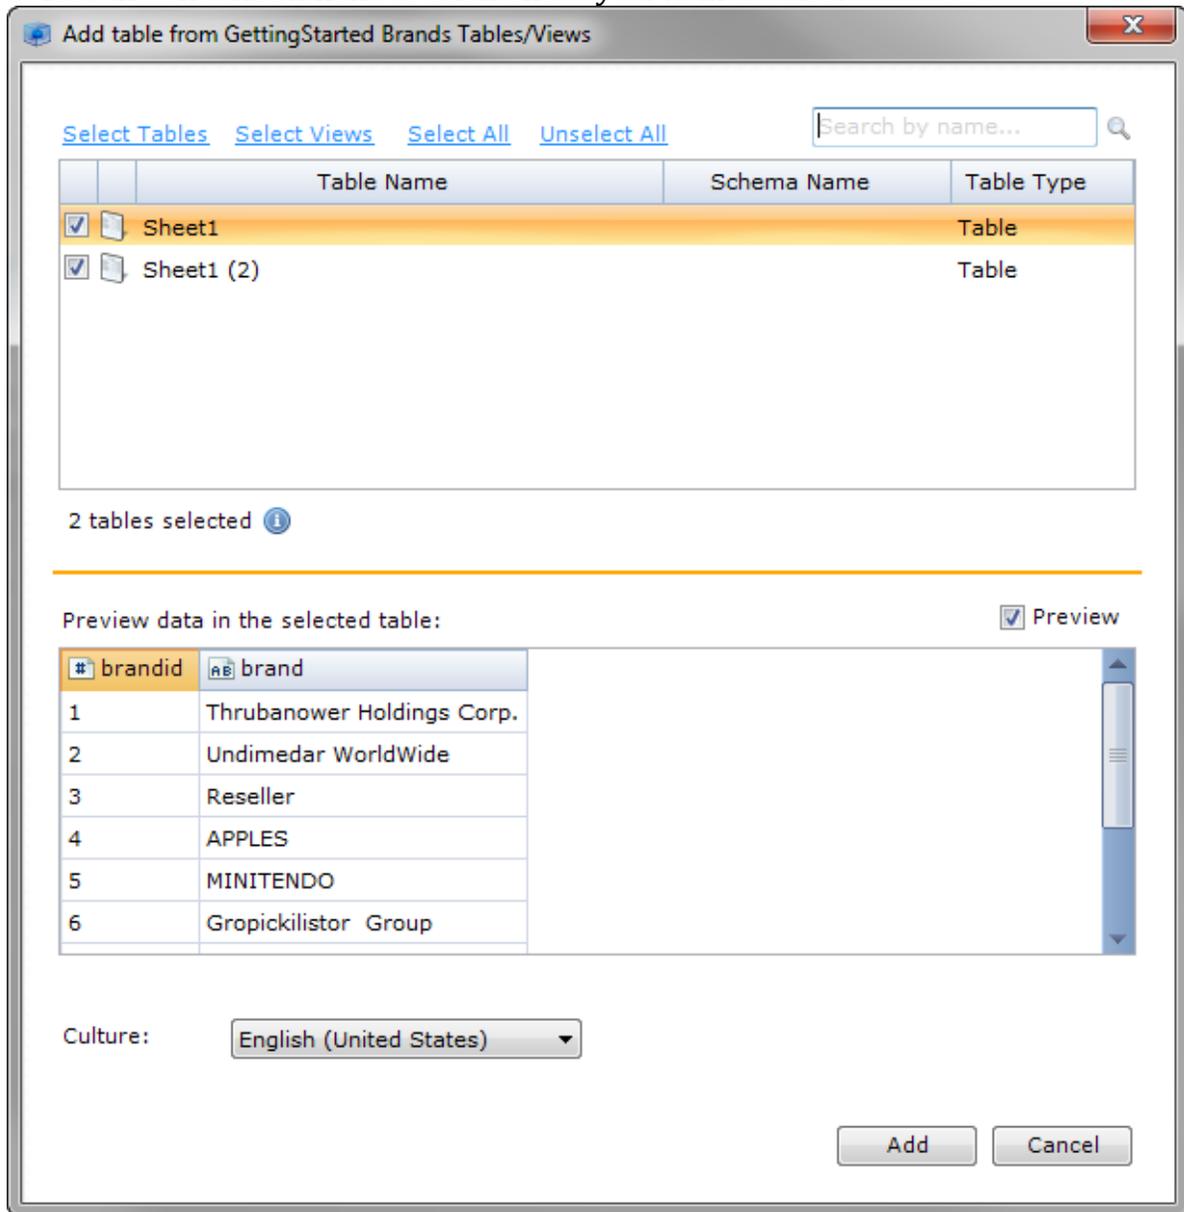
The following screen appears.



In the **User Account** field, enter your user account email.

4. Click **Get** to obtain your Google 2-Step Verification code, and enter it in the **Auth Code** field.
5. Click **Connect to Server**.
6. From the **Select Database** list, select the relevant spreadsheet you want to work with and click **OK**.

7. Select the checkbox next to each sheet you want to use.



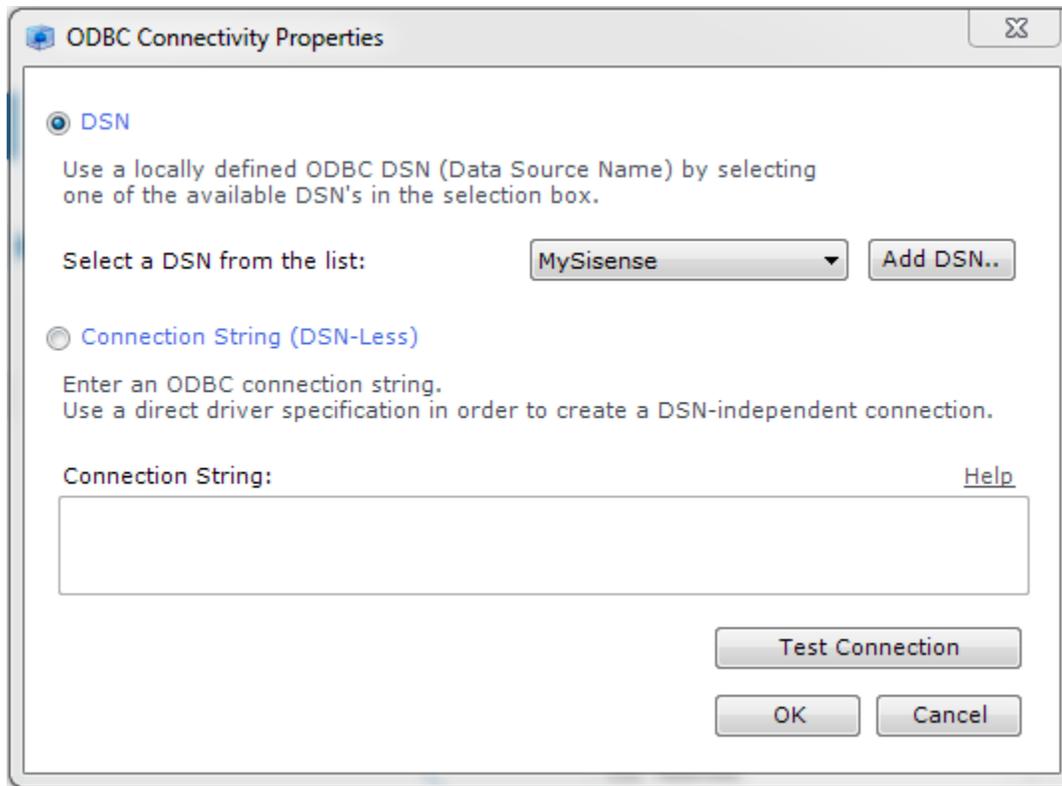
8. Once all relevant tables are selected, click **Add**.

# Connecting to Hadoop via Hive

The Sisense ODBC Driver with SQL Connector for Apache Hive enables standard SQL-92 access directly to Apache Hive distributions. It efficiently maps SQL to HiveQL delivering full SQL application functionality and real-time analytic and reporting capabilities to viewers.

## Downloading and Connecting the Sisense Hive ODBC Driver

1. [Download](#) and install the Sisense Hive ODBC Driver.
2. In ElastiCube Manager, click **Add Data** and then, **Generic ODBC Driver**.
3. Select **DSN**. If a DSN file has already been created, select it in the list, and click Test Connection. If you need to add a DSN, see [Adding a DSN](#). If your connection has been set up, you can click **OK** and review the tables generated by the Sisense Hive ODBC driver.



OR

Select **Connection String (DSN-Less)** and enter your Connection String. Click **Test Connection**. For more information, see [Connecting without a DSN](#).

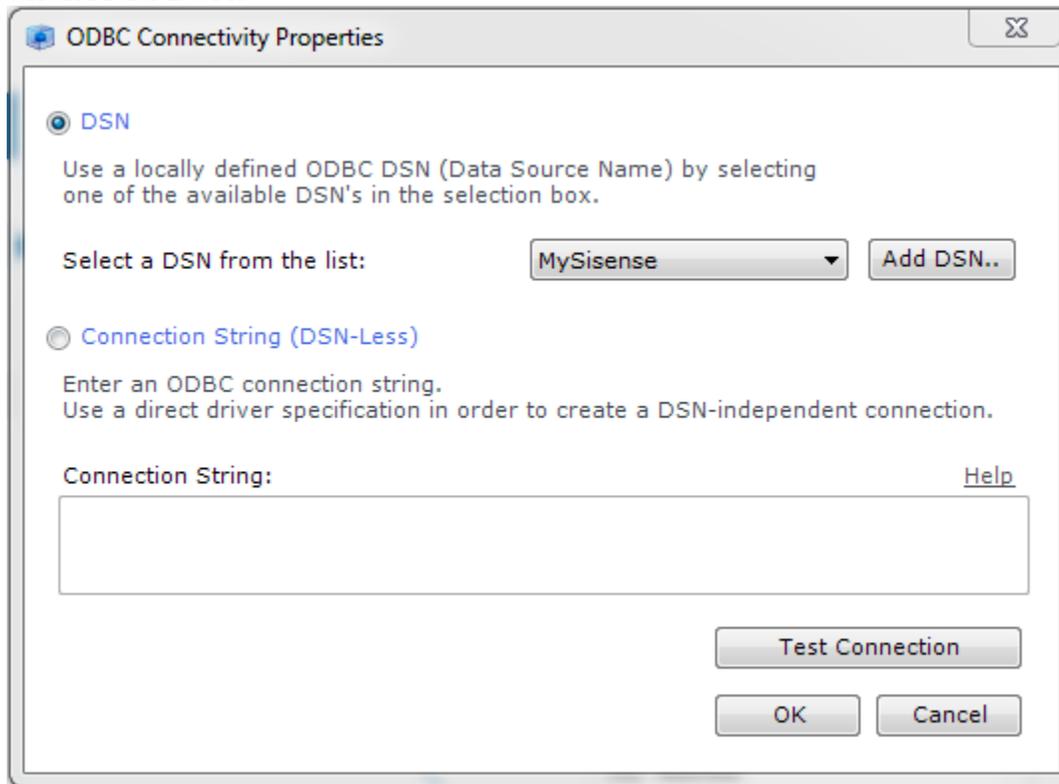
## Adding a DSN

If a DSN has not yet been created, you must add the DSN manually.

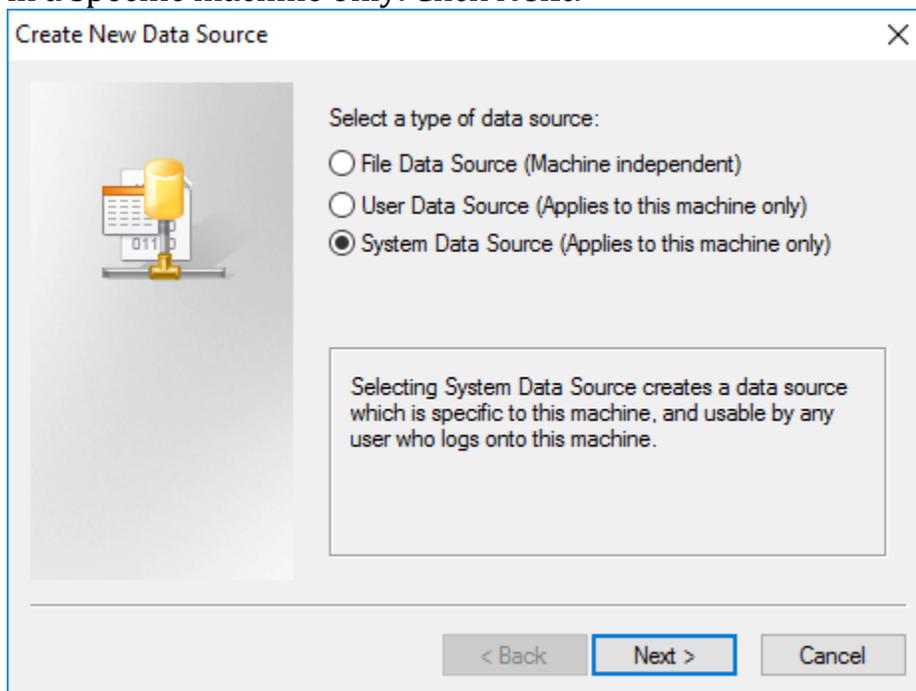
**Note:** To add a DSN you must run the Sisense ElastiCube Manager as an administrator.

**To add a DSN:**

1. Click **Add DSN**.



2. Select the **System Data Source** option. The created file applies to all users in a specific machine only. Click **Next**.



3. Select the Sisense Hive ODBC Driver, and click **Next**.
4. Click **Finish**.

5. In the Sisense Hive ODBC Driver DSN Setup window, do the following:

Field	Procedure
Data Source Name	Enter your name for the DSN that is to be displayed in the DSN dropdown list.

Field	Procedure
Description	Enter your description of the DSN.
Hive Server Type	Select the server type used by your data source. <b>Note:</b> If you are connecting through Apache ZooKeeper, then Hive Server 1 is not supported.
Service Discovery Mode	Select the Service Discovery Mode. There are two possible options you can select: <b>No Service Discovery:</b> The driver connects to Hive without using the ZooKeeper service. <b>ZooKeeper:</b> The driver discovers Hive Server 2 services via the ZooKeeper service.
Host(s)	If you selected No Service Discovery above as the Service Discovery Mode, then enter the IP address or hostname of the Hive server. OR If you selected ZooKeeper as the Service Discovery Mode, enter a comma-separated list of ZooKeeper servers. Use the following format, where zk_host is the IP address or hostname of the ZooKeeper server and zk_port is the number of the port that the ZooKeeper server uses: zk_host1:zk_port1,zk_host2:zk_port2 .
Port	If you selected No Service Discovery above as the Service Discovery Mode, then enter the number of the TCP port on which the Hive server is listening. Otherwise, do not type a value in the field.
Database	Enter the name of the database schema to use when a schema is not explicitly specified in a query. You can still issue queries on other schemas by explicitly specifying the schema in the query. To inspect your databases and determine the appropriate schema to

Field	Procedure
	use, type the show databases command at the Hive command prompt.
ZooKeeper Namespace	If you selected ZooKeeper as the Service Discovery Mode, enter the namespace on ZooKeeper under which Hive Server 2 znodes are added. Otherwise, do not type a value in the field.
Authentication	<p>Some Hive servers are configured to require authentication for access. To connect to a Hive server, you must configure the Sisense Hive ODBC Driver to use the authentication mechanism that matches the access requirements of the server and provides the necessary credentials. Hive Server 1 does not support authentication.</p> <p>Most default configurations of Hive Server 2 require User Name authentication. To verify the authentication mechanism that you need use for your connection, check the configuration of your Hadoop/Hive distribution.</p>
Mechanism	If your Hive server is configured to use authentication, select the appropriate authentication mechanism and provide the necessary credentials and authentication settings in the fields below the list.
Realm	<p>This field is available when Kerberos is selected as your authentication mechanism.</p> <p>Enter the realm of the Hive Server 2 host.</p> <p>If your Kerberos configuration already defines the realm of the Hive Server 2 host as the default realm, then you do not need to configure this option.</p>
Host FQDN	This field is available when Kerberos is selected as your authentication mechanism. Enter the HS2HostFQDN key to the fully qualified domain name of the Hive Server 2 host.

Field	Procedure
Service Name	This field is available when Kerberos is selected as your authentication mechanism. Enter the HS2KrbServiceName key to the service name of the Hive Server 2 host.
User Name	This field is available when User Name, User Name and Password, or Windows Azure is selected as your authentication mechanism. The user name that you use to access Hive Server 2.
Password	This field is available when User Name and Password or Windows Azure is selected as your authentication mechanism. The password corresponding to the user name that you provided in the User Name field (the UID key).
Delegation UID	Enter the name of the user to be delegated in the Delegation UID field. The Delegation UID represents a user that is different than the authenticated user for the connection. This option is applicable only when connecting to a Hive Server 2 instance that supports this feature.
Thrift Transport	Select the transport protocol to use in the Thrift layer. For Hive Server select Binary. When using an authentication mechanism, the <b>Binary</b> transport protocol is not supported. For Hive Server 2, select <b>SASL</b> . If you select <b>HTTP</b> , the HTTP Options button is enabled. See Configuring HTTP Options for more information. To configure client-server verification over SSL, click SSL Options. See SSL Options for more information.

6. (Optional) To configure advanced driver options, click **Advanced Options**. For more information, see Advanced Options.
7. (Optional) To configure logging behavior for the driver, click **Logging Options**. For more information, see Logging Options.

8. Click **Test**. A pop window is displayed that indicates the connection was successful.
9. Click **OK**.
10. In the ODBC Connectivity Properties window, click **Test Connection**.
11. Click **OK** to save and close the ODBC Data Source Administrator.

## Configuring HTTP Options

You can configure options such as custom headers when using the HTTP transport protocol in the Thrift layer.

### To configure HTTP options:

1. In the Sisense Hive ODBC Driver DSN Setup window, select **HTTP** as the HTTP transport protocol from the Thrift Transport list.
2. Click **HTTP Options**. The HTTP Properties window is displayed.

The screenshot shows the 'HTTP Properties' dialog box. At the top, there is a title bar with 'HTTP Properties' and a close button. Below the title bar, there is a text input field labeled 'HTTP Path:'. Underneath that is a section titled 'Custom HTTP Headers' which contains a list box. To the right of the list box are three buttons: 'Add...', 'Edit', and 'Remove'. At the bottom of the dialog box are two buttons: 'OK' and 'Cancel'.

3. In the **HTTP Path** field, enter the partial URL corresponding to the Hive server.
4. Click **Add**. The Edit Property window is displayed.

The screenshot shows the 'Edit Property' dialog box. It has a title bar with 'Edit Property' and a close button. Below the title bar is a section titled 'Server Side Properties'. Inside this section, there are two text input fields: one labeled 'Key:' and one labeled 'Value:'. At the bottom of the dialog box are two buttons: 'OK' and 'Cancel'.

5. In the **Key** and **Value** fields, enter the appropriate keys and values and click **OK**.

6. After you add an HTTP header, you can edit or delete the header by clicking **Edit** or **Delete**.
7. Click **OK**.

## SSL Options

You can configure verification between the client and the Hive server over SSL.

### To configure SSL verification:

1. In the Sisense Hive ODBC Driver DSN Setup window, click **SSL Options**. The SSL Options window is displayed.

The screenshot shows the 'SSL Options' dialog box. It features a title bar with a close button (X). The main area contains the following elements:

- Checkboxes:  Enable SSL,  Allow Common Name Host Name Mismatch,  Allow Self-signed Server Certificate.
- Trusted Certificates: A text box containing 'C:\Program Files\Sisense Hive ODBC Driver\lib\cacerts.pem' and a 'Browse...' button.
- Two-way SSL:  Two-way SSL.
- Client Certificate File: A text box and a 'Browse...' button.
- Client Private Key File: A text box and a 'Browse...' button.
- Client Private Key Password: A text box and a  Save Password (Encrypted) checkbox.
- Buttons: 'OK' and 'Cancel' buttons at the bottom right.

2. Select **Enable SSL**.
3. Select the **Allow Self-signed Server Certificate** checkbox to allow self-signed certificates from the server.
4. Select the **Allow Common Name Host Name Mismatch** checkbox to allow the common name of a CA-issued SSL certificate to not match the host name of the Hive server.
5. In the **Trusted Certificates** field, enter the path or click **Browse** and navigate to the location of your PEM file if you want the driver to load your SSL certificates when verifying the server.

OR

To use the trusted CA certificates PEM file that is installed with the driver, leave the Trusted Certificates field empty.

6. Select the **Two-Way** SSL checkbox if you want to configure two-way SSL verification. The Client authentication fields are enabled.
7. In the **Client Certificate File** field, enter the path or click **Browse** and navigate to the location of your PEM containing the client's certificate.
8. In the **Client Private Key File** field, enter the path or click **Browse** and navigate to the location of your file containing the client's private key.
9. In the **Client Private Key Password** field, if the private key file is protected with a password, enter the password. Select the **Save Password (Encrypted)** checkbox to save the password.
10. Click **OK**.

## Advanced Options

You can configure advanced driver by clicking **Advanced Options** in the Sisense Hive ODBC Driver DSN Setup window.

This displays the Advanced Options window from which you can configure the following options:

Field	Description
<input type="checkbox"/> Use Native Query	Rows fetched per block: 10000
<input type="checkbox"/> Fast SQLPrepare	Default string column length: 255
<input type="checkbox"/> Driver Config Take Precedence	Binary column length: 32767
<input type="checkbox"/> Use Async Exec	Decimal column scale: 10
<input type="checkbox"/> Get Tables With Query	Async Exec Poll Interval (ms): 100
<input type="checkbox"/> Unicode SQL character types	Socket Timeout (s): 30
<input type="checkbox"/> Show System Table	
<input type="checkbox"/> Use only SSPI	

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

Field	Description
Use Native Query	<p>Select to disable the SQL Connector feature.</p> <p>Select Use Native Query to disable the SQL Connector feature and allow the driver to execute HiveQL directly.</p> <p>When Use Native Query is selected, the driver does not transform the queries emitted by an application, so the native query is used.</p> <p>When Use Native Query is disabled, the driver transforms the queries emitted by an application and converts them into an equivalent form for Hive.</p>
Fast SQLPrepare	<p>Select to defer query execution to SQLExecute.</p>
Driver Config Take Precedence	<p>Select to allow driver-wide configurations to take precedence over connection and DSN settings.</p>
Use Async Exec	<p>Select to use the asynchronous version of the API call against Hive for executing a query.</p> <p>This option is applicable only when connecting to a Hive cluster running Hive 0.12.0 or later.</p>
Get Tables With Query	<p>Select to retrieve the names of tables in a database by using the SHOW TABLES query.</p> <p>This option is applicable only when connecting to Hive Server 2.</p>
Unicode SQL character types	<p>Select to enable the driver to return SQL_WVARCHAR instead of SQL_VARCHAR for STRING and VARCHAR columns, and SQL_WCHAR instead of SQL_CHAR for CHAR columns.</p>
Show System Table	<p>Select to enable the driver to return the hive_system table for catalog function calls such as SQLTables and SQLColumns.</p>

Field	Description
Use only SSPI	Select to handle Kerberos authentication using the SSPI plugin instead of Kerberos by default.
Rows fetched per block	Enter the number of rows to be fetched per block.
Default string column length	Enter the maximum data length for STRING columns.
Binary column length	Enter the maximum data length for BINARY columns.
Decimal column	Enter the maximum number of digits to the right of the decimal point for numeric data types.
Async Exec Poll Interval (ms)	Enter the time in milliseconds between each poll for the query execution status. This option is applicable only to HDInsight clusters.
Socket Timeout (s)	Define the amount of time until the connection times out. The default amount of time is 30 seconds.

1. Click **Server Side Properties** to configure the driver to apply configuration properties to the Hive server. For more information, see [Configuring Server Side Properties](#).
2. Click **OK**.

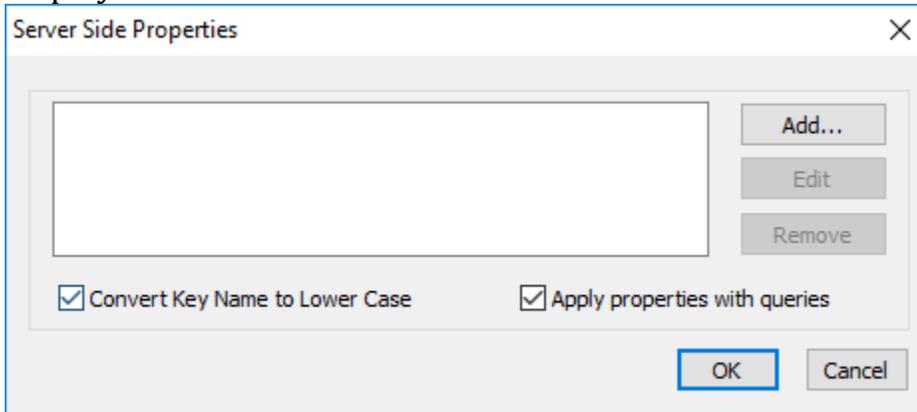
## Configuring Server-Side Properties

You can use the driver to apply configuration properties to the Hive server.

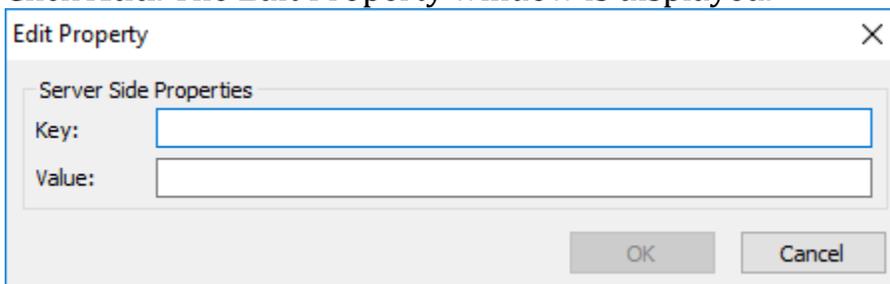
The Sisense ODBC Driver with SQL Connector for Apache Hive allows you to set server-side properties via a DSN. Server-side properties specified in a DSN affect only the connection that is established using the DSN.

**To create a server-side property:**

1. In the Sisense Hive ODBC Driver DSN Setup window, click **Advanced Options > Side Server Properties**. The Server Side Properties window is displayed.



2. Click **Add**. The Edit Property window is displayed.



3. In the **Key** and **Value** fields, enter the appropriate keys and values and click **OK**.

For a list of all Hadoop and Hive server-side properties that your implementation supports, type `set -v` at the Hive CLI command line or Beeline. You can also execute the `set -v` query after connecting using the driver. After you add your properties, you can edit or delete the properties by clicking **Edit** or **Delete**.

4. Click **OK**.
5. Select the **Apply properties with queries** checkbox to configure the driver to apply each server-side property by executing a query when opening a session to the Hive server.

OR

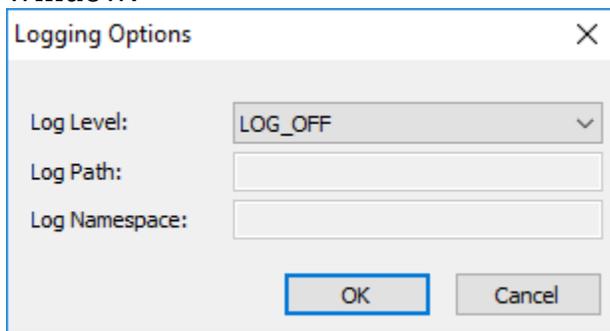
Clear the Apply Server Side Properties with Queries checkbox to configure the driver to use a more efficient method for applying server-side

properties that does not involve additional network round-tripping. The more efficient method is not available for Hive Server 1, and it might not be compatible with some Hive Server 2 builds. If the server-side properties do not take effect when the checkbox is clear, then select the checkbox.

6. Select the **Convert Key Name to Lower Case** checkbox to force the driver to convert server-side property key names to all lower case characters.
7. Click **OK**.

## Logging Options

You can configure logging options by clicking **Logging Options** in the Sisense Hive ODBC Driver DSN Setup window. This displays the Logging Options window.



The ODBC Data Source Administrator provides tracing functionality, which you can activate to help troubleshoot issues.

**Important:** Only enable logging long enough to capture an issue. Logging decreases performance and can consume a large quantity of disk space.

The driver allows you to set the amount of detail included in log files. The table below lists the logging levels provided by the Sisense ODBC Driver with SQL Connector for Hive, in order from least verbose to most verbose.

In the Logging Options window, you can configure the following information:

Field	Description
Log Level	Select the Log Level. There are seven possible options: <b>LOG_OFF:</b> Disables all logging.

Field	Description
	<p><b>LOG_FATAL:</b> Logs very severe error events that lead the driver to abort.</p> <p><b>LOG_ERROR:</b> Logs error events that might still allow the driver to continue running.</p> <p><b>LOG_WARNING:</b> Logs potentially harmful situations.</p> <p><b>LOG_INFO:</b> Logs general information that describes the progress of the driver.</p> <p><b>LOG_DEBUG:</b> Logs detailed information that is useful for debugging the driver.</p> <p><b>LOG_TRACE:</b> Logs more detailed information than the DEBUG level.</p>
Log Path	<p>Enter the full path to the folder where you want to save log files.</p> <p>OR</p> <p>Click <b>Browse</b> and select the folder where you want to save log files.</p>
Log Namespace	<p>Enter a name for the log.</p>

## Connecting without a DSN

Some client applications provide support for connecting to a data source using a driver without a DSN. To configure a DSN-less connection, you can use a connection string.

The following is an example of a connection string for a DSN-less connection:

```
Driver=DriverNameOrFile;HOST=MyHiveServer;PORT=PortNumber; Schema=DefaultSchema;HiveServerType=ServerType
```

The table below provides a list of possible keys you can add to your string and their descriptions:

Key Name	Description
DriverNameOrFile	The absolute path of the shared object file for the driver.

Key Name	Description
MyHiveServer	The IP address or host name of the Hive server.
PortNumber	The number of the port that the Hive server uses.
DefaultSchema	The database schema to use when a schema is not explicitly specified in a query.
ServerType	The server type, either 1 (for Hive Server 1) or 2 (for Hive Server 2).

## Adding Hive Tables to your Project

After setting up the DSN or the DSN-less connection, the Add table from Hive Tables window is displayed.

From this window, you add your Hive tables to your ElastiCube Manager. In addition, you can view the SQL syntax in the Query Preview section and click Edit to customize it.

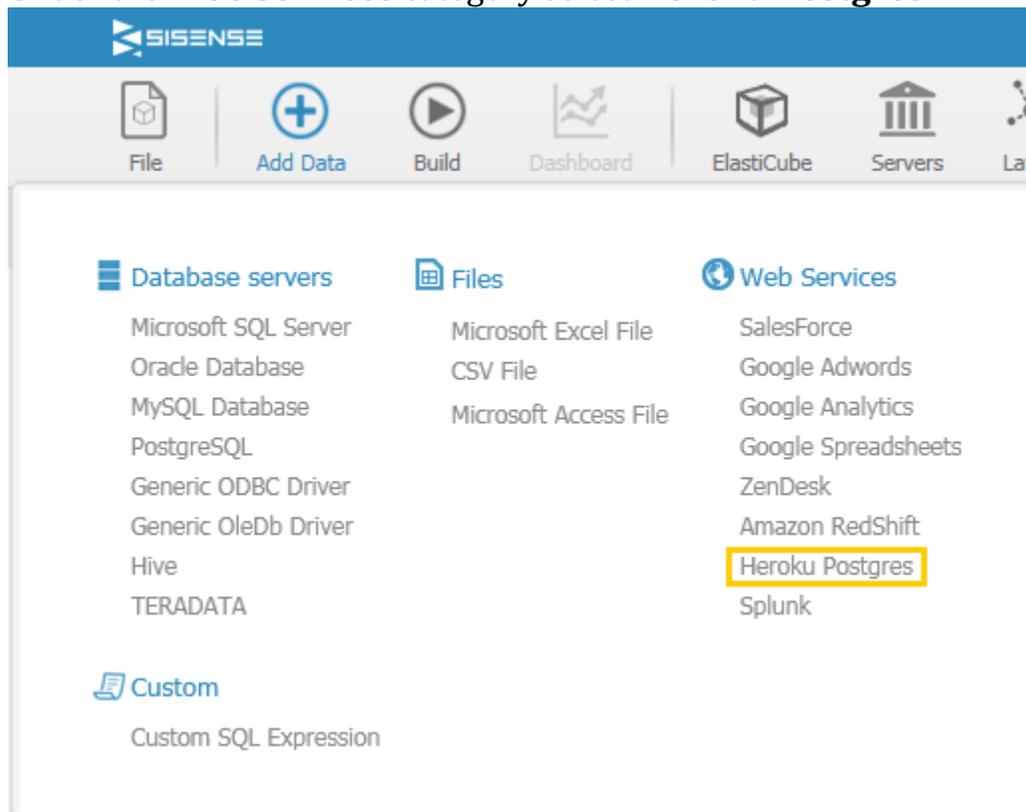
### To add Hive Tables to your Project:

1. Connect to Hive via the Sisense ODBC Tool.
2. Add your data source.
3. In the Add table from Hive Tables window, select the table you want to add to the ElastiCube Manager.
4. Click **Add**. The selected tables are added to your ElastiCube Manager.

# Connecting to Heroku Postgres

The ElastiCube Manager enables easy and quick access to databases, tables and views contained within Heroku Postgres databases. The steps below describe how to connect to this type of data source.

1. Click **Add data** in the top menu of the ElastiCube Manager.
2. Under the **Web Services** category select **Heroku Postgres**.



3. You will be prompted to enter the following information:
  - **Database server location:** Enter the computer/server IP address which holds the database.
  - Either use **Windows Authentication** if configured with the database, or alternatively enter the **Username** and **Password** to connect to the database.

- **Database:** Enter the name of the database to connect to.

Connect to HerokuPostgres

Database server location:

Windows Authentication

Use The Following User Name & Password:

UserName

Password

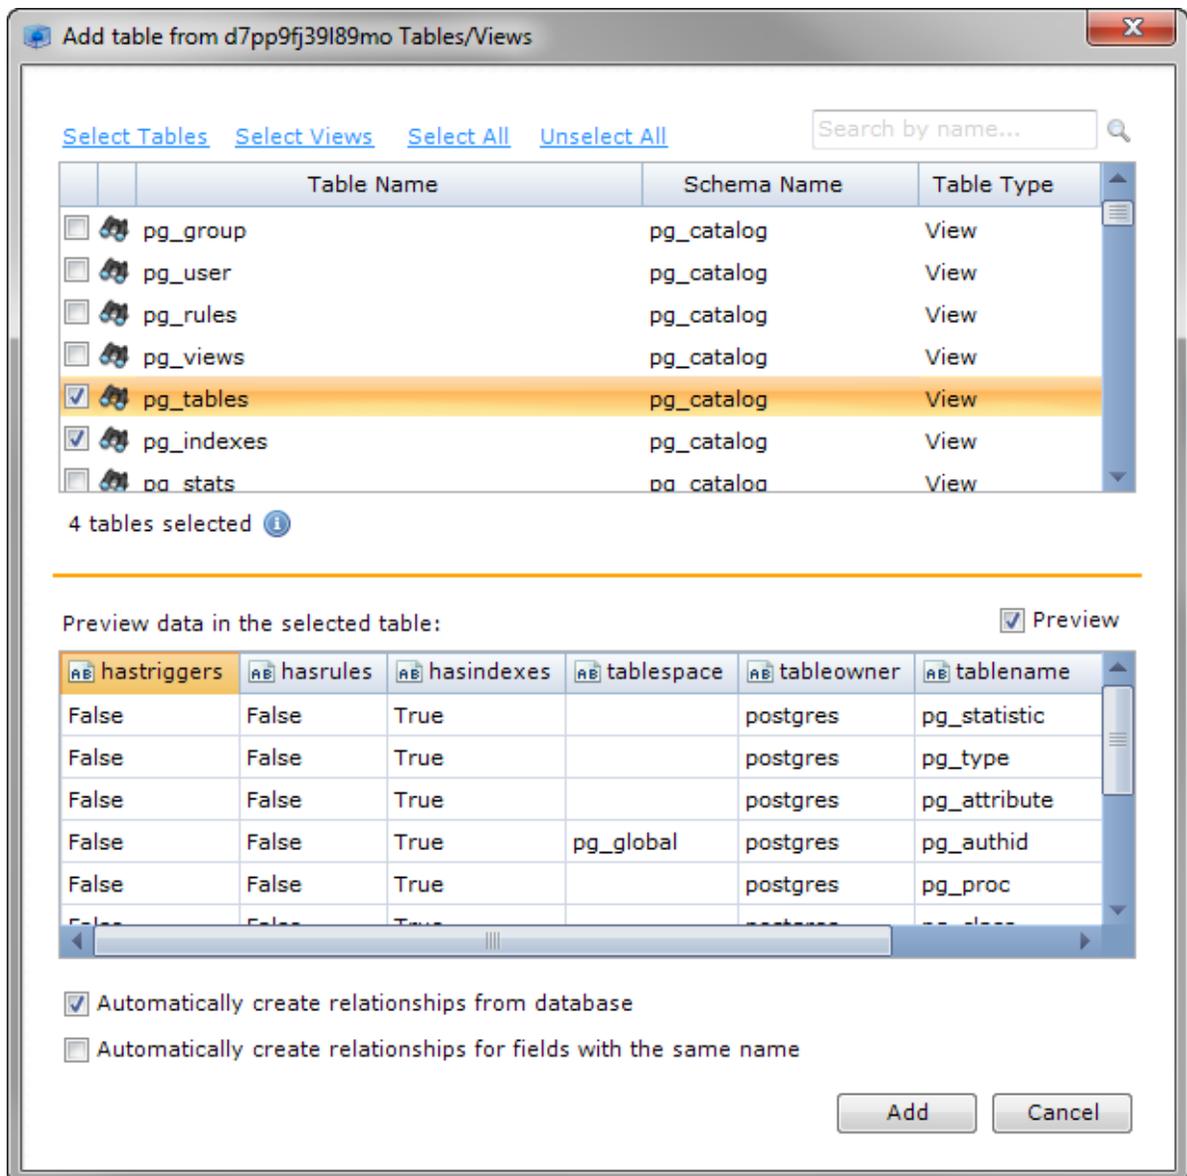
Database

Connect to Database

Cancel

4. Click **Connect to Database**.

All tables and views associated with the database will appear in a new window.



To view a preview of data contained in a particular table, highlight the table or view in the list, and select the **Preview** checkbox.

5. Select the checkbox next to each table or view you want to use. Existing relationships between tables can be automatically replicated in the ElastiCube by selecting the **Automatically create relationships from database** checkbox. Likewise fields with similar names can be linked by selecting the **Automatically create relationships for fields with the same name** checkbox.
6. Once all relevant tables are selected, click **Add**.

# Connecting to HubSpot

The Sisense HubSpot connector is a standalone connector that allows you to import data from HubSpot's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the HubSpot APIs. To obtain a connection string, you will need to create a HubSpot app.

Once you have connected to HubSpot, you can import a variety of tables from the HubSpot API.

This page describes how to install the HubSpot connector, how to connect to HubSpot with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the HubSpot Connector](#)
- [Adding HubSpot Tables to your ElastiCube](#)
- [HubSpot Tables](#)
- [Versions](#)

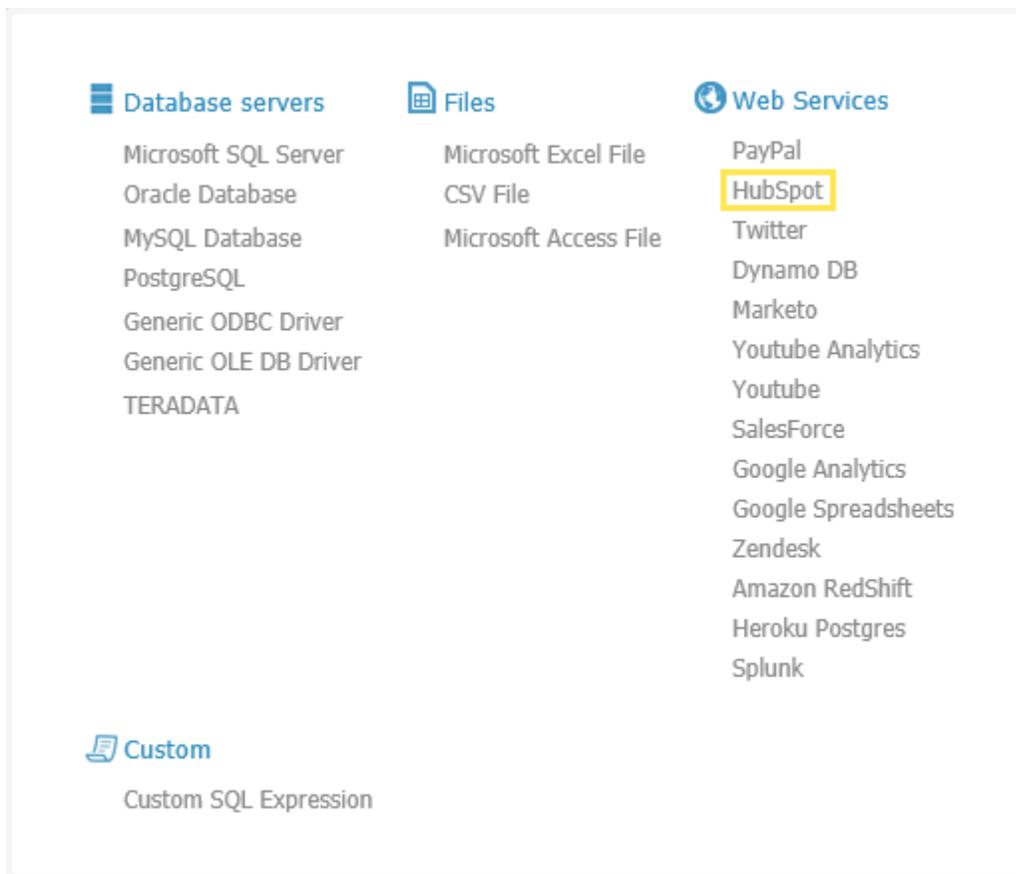
## Installing the HubSpot Connector

Sisense provides the HubSpot connector as a standalone connector that you can download and add to your list of default Sisense connectors.

### To install the HubSpot connector:

1. [Download](#) the HubSpot installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The HubSpot connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.



## Adding HubSpot Tables to Your ElastiCube

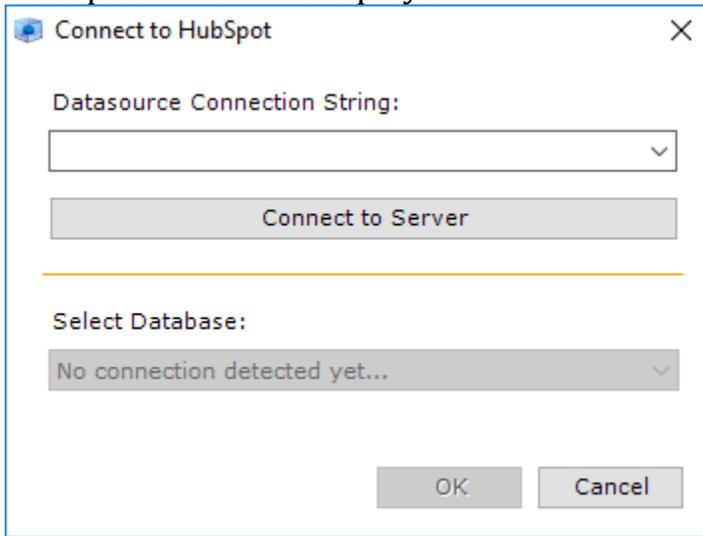
Sisense uses connection strings to connect to HubSpot and import data into the ElastiCube Manager. The default connection string for HubSpot is:

```
jdbc:hubspot:InitiateOAuth=GETANDREFRESH;
```

After you connect using the default connection string, Sisense redirects you to Hubspot's homepage where you log in and your account is authenticated. Once the account is authenticated, the relevant tables can be added to your ElastiCube. If you want to customize your connection string, you can by creating an HubSpot app and passing additional connection string parameters. For more information, see [Creating an App](#) at the end of this page.

**To add HubSpot data:**

1. In ElastiCube Manager, click Add Data and then, HubSpot. The Connect to HubSpot window is displayed.



2. In **Datasource Connection String**, enter the connection string: `jdbc:hubspot:InitiateOAuth=GETANDREFRESH;`
3. Click **Connect to Server**. You are redirected to the HubSpot Login page.
4. Log in to your HubSpot account. HubSpot is displayed in the Select Database list in the ElastiCube Manager.
5. Click **OK**. Sisense connects to HubSpot and displays a list of tables available for you to import.
6. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

## HubSpot Tables

HubSpot's RESTful APIs expose the following HubSpot tables that you can import into the ElastiCube Manager through the Sisense HubSpot connector:

Name	Description
BlogAuthors	Retrieve the available blog authors in HubSpot.
BlogPosts	Retrieve the available blog posts in HubSpot.
Blogs	Retrieve the available blogs in HubSpot.
BlogTopics	Retrieve the available blog topics in HubSpot.
Comments	Retrieve the available comments from your blog in

Name	Description
	HubSpot.
Companies	Companies in HubSpot represent different organizations you are interested in keeping track of for potential marketing opportunities. They are treated similarly like contacts and have a number of custom properties that will be dynamically determined and added to the Companies table from your HubSpot Hub.
CompanyProperties	Company properties are the standard and custom pieces of field data that appear in HubSpot.
CompanyPropertyGroups	Company property groups in HubSpot offer a way of organizing individual types of properties for companies. Each company property must belong to a property group.
Contacts	Contacts in HubSpot represent different contacts you are interested in keeping track of for potential marketing opportunities. Contacts are similar to companies: They both have a number of custom properties that will be dynamically determined and added to the table from your HubSpot Hub.
ContactLists	Contact lists in HubSpot can be used to group together contacts with similar characteristics.
ContactListMemberships	A list of objects that represents the memberships of a contact from their contact lists. This list may be empty if the record is not a member of any lists.
ContactProperties	Contact properties are the standard and custom pieces

Name	Description
	of field data that appear in HubSpot.
ContactPropertyGroups	Contact property groups in HubSpot offer a way of organizing individual types of properties for contacts. Each contact property must belong to a property group.
DealAssociations	The deal can be associated with either contacts or companies with each request.
DealPipelines	Allows you to programmatically access options for the 'deal stage' and 'pipeline' properties within your instance of the HubSpot CRM.>br/> Every portal initially contains a default pipeline with the pipelineId "default". On portals that have only one pipeline, the pipeline property for any deal will be set to "default" automatically. On portals that have multiple pipelines, if you're setting a dealstage that is not in the default pipeline, you'll also need to set the pipeline property that lines up with the dealstage being used.
DealPipelineStages	A list of stages for this specific pipeline.
DealProperties	Deal properties are the standard and custom pieces of field data that appear in HubSpot.
DealPropertyGroups	Contact property groups in HubSpot offer a way of organizing individual types of properties for contacts. Each contact property must belong to a property group.
Deals	Deals in HubSpot represent different organizations you are interested in keeping track of for potential marketing opportunities. They are treated similarly like

Name	Description
	contacts and have a number of custom properties that will be dynamically determined and added to the Deals table from your HubSpot Hub.
Domains	Retrieve the available domains in HubSpot.
EmailCampaigns	Email campaigns in HubSpot allow you to keep track of and update email marketing campaigns. This table allows you to create, update, and delete your email campaigns in HubSpot.
EmailCampaignEvents	The events associated with an email campaign or a recipient.
EmailSubscriptions	The subscription types a given email is subscribed to. An email must be specified to return results. Subscriptions may be removed by deleting them.
EmailSubscriptionTypes	A list of email subscription types for a HubSpot hub.
Engagements	Engagements represent any of a number of different types of engagements you have in HubSpot. These can vary from simple Tasks, to Emails, Calls, Meetings, or others. Due to the different types of engagements available, any individual Engagement will have several columns that come back null due to its type.
Files	Retrieves information about the available files in HubSpot.
Folders	Retrieves information about the available folders in HubSpot.
Forms	HubSpot Forms.

Name	Description
Owners	Retrieve the owners in HubSpot.
Pages	Retrieve the available pages in HubSpot.
SocialMediaChannels	List available social media channels in HubSpot.
Settings	HubSpot settings.
SocialMediaMessages	Social media messages.
Templates	Retrieve the available templates in HubSpot.
UrlMappings	Retrieve the available url mappings in HubSpot.
Workflows	HubSpot workflows.

### Accumulative Builds

Sisense support accumulative builds for all numeric and dates data types. However, the data must be sorted before building the ElastiCube.

### Versions

Version	Release Date	Improvements
6.6.0.61	April 18, 2017	<p>UNIX timestamp was set to local time</p> <p><b>Note:</b> HubSpot has deprecated the V1 method of retrieving the OAuth Access token. If you have problems authenticating, you will need to log into your HubSpot account with the same user you are using as your OAuth Client Id/Secret. After logging into your HubSpot account, go to your application settings and edit the Scopes. Remove any scopes that you might have selected and save the changes. Then, attempt to retrieve your OAuth token again. For more information, see Migrating from OAuth1 to OAuth 2.0</p>

Version	Release Date	Improvements
6.6.0.51	March 16, 2017	Dates are handled as expected
6.5.1.32	January 16, 2017	General release

## Creating an App

To access HubSpot's REST API from the ElastiCube Manager with a customized connection string, you must provide valid OAuth HubSpot credentials through a connection string. These credentials are provided by HubSpot when you create an application.

After you receive your credentials from HubSpot, you can create the connection string and provide Sisense with it to connect to your data.

Each connection string contains authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

The following is an example of a HubSpot connection string and its mandatory parameters:

```
jdbc:hubspot:OAuthClientId=xxxxxxxxxxxxxxxx;OAuthClientSecret=xxxxxxxxxxxxxxxx;InitiateOAuth=GETANDREFRESH;HubId=xxxxxxxx;CallbackURL=https://localhost;Timeout=0;
```

1. **HubId:** Set your Hubspot account ID to HubId in the connection string.
2. **OAuthClientId:** Set this to the client Id in the connection string.
3. **OAuthClientSecret:** Set this to the client Id in the connection string.
4. **InitiateOAuth:** Set to GETANDREFRESH.
5. **CallbackURL:** Following HubSpot's migration to OAuth 2.0, the CallbackURL needs to be SSL compliant so it should be https://localhost.

For more information about connection strings and optional parameters, see

[Connection String Parameters](#).

When you connect to HubSpot, you are redirected to HubSpot's site where your Administrator must grant access to your account.

Follow the steps below to obtain the OAuth client credentials, the OAuthClientId and OAuthClientSecret:

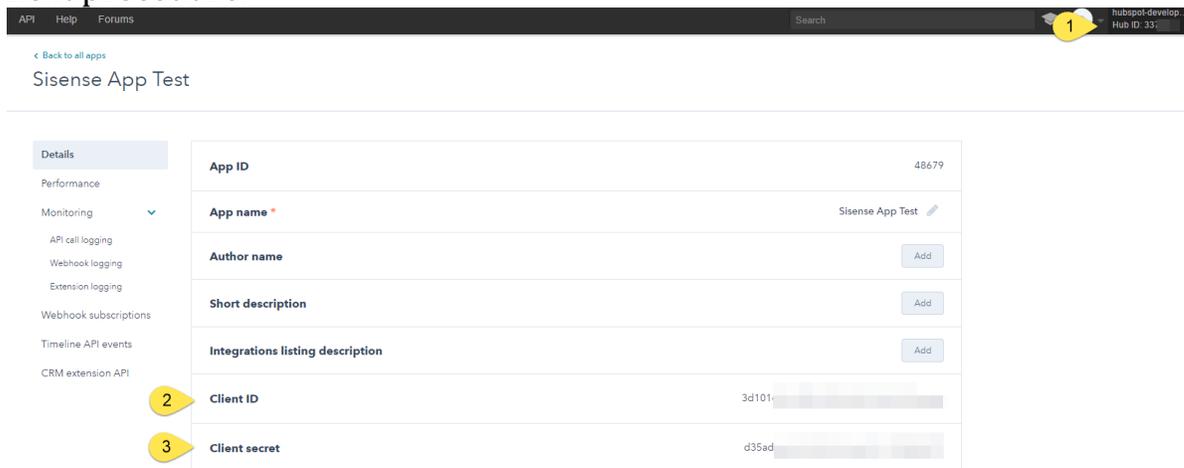
1. Log into your HubSpot developer account.
2. Click **Create App**.
3. In the **App name** field, enter a name and select Public or Private.
4. Click **Create**.
5. In the Apps page, click the name of your app.

Apps [Get HAPIkey](#) [Create app](#)

---

NAME	APP TYPE	DATE CREATED ▼	APP ID	APP INSTALLS
 Sisense App Test	Public	May 25, 2017	48679	0

6. Copy the following values and add them to your connection described in the next procedure.



API Help Forums Search 1 hubspot-developer Hub ID: 33

[Back to all apps](#)

### Sisense App Test

---

**Details**

- Performance
- Monitoring ▼
  - API call logging
  - Webhook logging
  - Extension logging
  - Webhook subscriptions
- Timeline API events
- CRM extension API

<b>App ID</b>	48679
<b>App name *</b>	Sisense App Test <span style="font-size: 0.8em;">✎</span>
<b>Author name</b>	<a href="#">Add</a>
<b>Short description</b>	<a href="#">Add</a>
<b>Integrations listing description</b>	<a href="#">Add</a>
<b>Client ID</b>	3d101... <span style="font-size: 0.8em;">✎</span>
<b>Client secret</b>	d35ad... <span style="font-size: 0.8em;">✎</span>

1. **HubId**
2. **OAuthClientId**
3. **OAuthClientSecret**

**Note:** Do not select any scopes in your account as this will prevent access to your data.

## Switching between Accounts

When you connect to the HubSpot data source, Sisense saves your OAuth values in the file OAuthsettings.txt file located at

...\Users\xxx\AppData\Roaming\CData\HubSpot Data Provider on your Sisense server. To connect to the HubSpot data source with another user on the same machine, you must delete the OAuthsettings.txt file. Sisense will then generate a new file for that user.

Another option to support multiple users is to define the location and file name of an OAuthsettings file for each unique user in your connection string through the OAuthSettingsLocation parameter. When each user connects to the data source, Sisense generates the OAuth file with the file name you specify in the location you define. In the examples below, two users are allowed to access the HubSpot data source and for each user, Sisense generates a file that contains that user's OAuth values in the location defined in the string.

```
jdbc:hubspot:OAuthSettingsLocation=C:\hubspot\auth\john.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
```

```
jdbc:hubspot:OAuthSettingsLocation=C:\hubspot\auth\sally.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
```

In the example above, two OAuth files are created, one for John and one for Sally in the location C:\hubspot\auth\.

This is useful if you support many users who each need to access the HubSpot data source.

# Importing Data with Custom Queries

Sisense allows you to import data from a variety of data sources and manipulate the data that is imported by running custom SQL queries against the database, for example:

- Import selected data only to reduce import and build times when you do not require all the data.
- Perform various procedures in the database such as table joins.

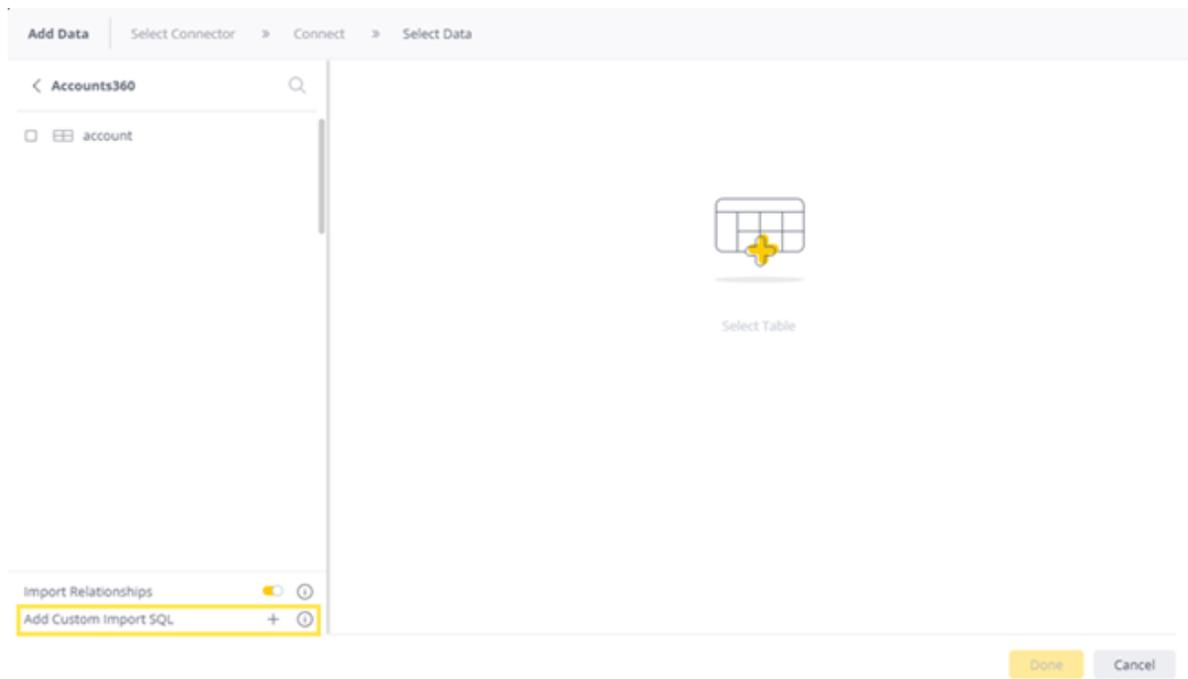
After running your custom query, you can preview the updated tables before importing the data into the ElastiCube.

## Notes:

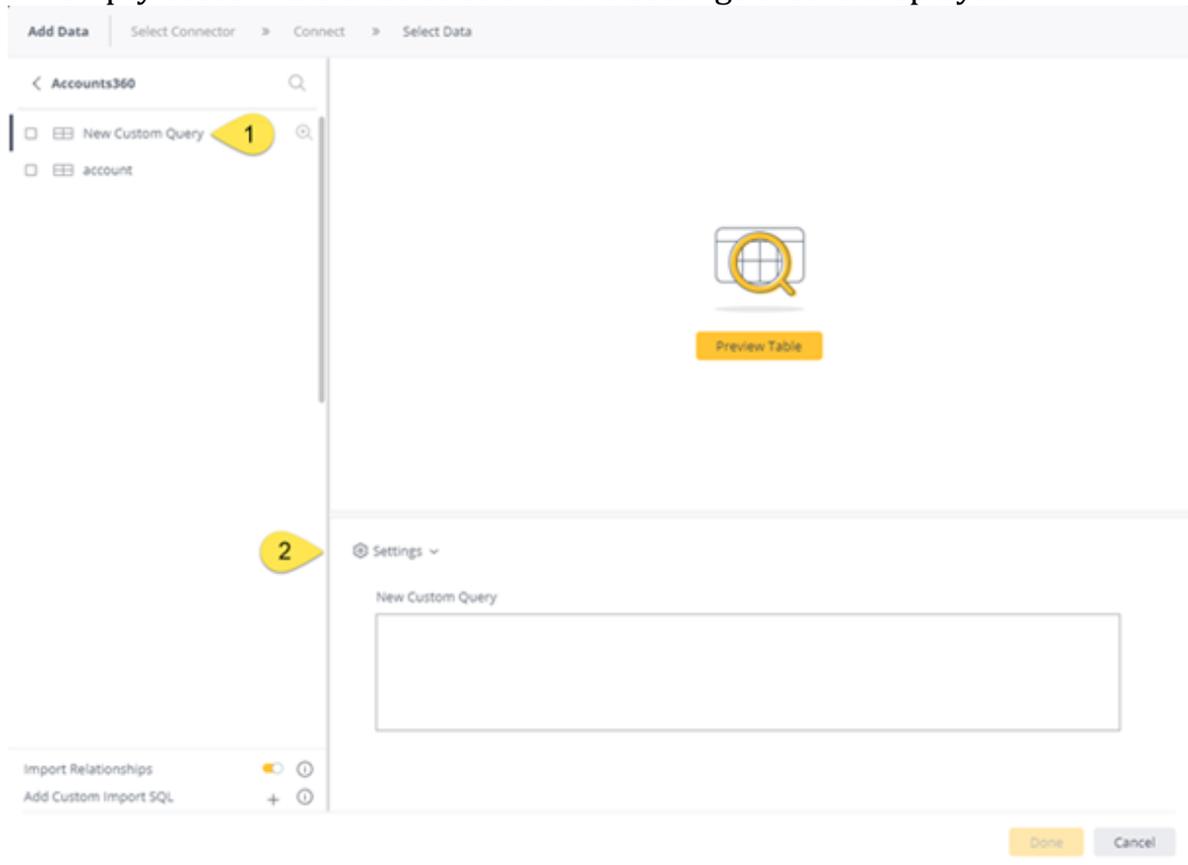
- This feature is available for the following relational databases: Microsoft SQL Server, MySQL, PostgreSQL, Oracle, and Amazon Redshift.
- This feature runs a query in the source database, which may require a lot of processing power from your database, especially for complex queries.

## To add tables with custom queries:

1. After selecting your database, click + at the bottom the list of tables and views in the database.



An empty custom table is added and the Settings area is displayed.



Click + to add additional custom tables.

2. In the Settings area, enter your custom SQL query.
3. Click **Preview Table** to preview the results of your custom query.  
OR  
Click **Done** to import your custom data into the ElastiCube Manager.

# Connecting to Intacct

This connector is currently in beta and subject to change.

The Sisense Intacct connector is a standalone connector that allows you to import data from Intacct into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who want to connect to Intacct.

Once you have connected to Intacct, you can import a variety of tables from the Intacct API.

This page describes how to install the Intacct connector, how to connect to Intacct with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the Intacct Connector](#)
- [Adding Intacct Tables to your ElastiCube](#)
- [Intacct Tables](#)

## Installing the Intacct Connector

Sisense provides the Intacct connector as a standalone connector that you can download and add to your list of default Sisense connectors.

### To install the Intacct connector:

1. Contact us for the Intacct installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The Intacct connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.

# Adding Intacct Tables to Your ElastiCube

After you have installed the Intacct connector, you can import your data from Intacct by connecting with a connection string. The connection string contains your Intacct users credentials. Intacct uses the connection string to authenticate your account and allow you to extract the relevant data for your account.

The connection string to Intacct is in the following format:

```
jdbc:Intacct:APIKey=xxxxxxxxxxxxxxxxxxxxxxxxxxxx;
```

For example:

```
jdbc:intacct:User='myusername';CompanyID='TestCompany';Password='mypassword';SenderID='Test';SenderPassword='abcde123';
```

Your Intacct connection string should include the following parameters, separated by a semicolon ';':

- **User:** Your Intacct username.
- **CompanyID:** Your company's Intacct ID used when you log in.
- **Password:** Your Intacct user password.
- **SenderID:** The Web Services SenderID assigned to you by Intacct.

To connect to a shared child company, add your LocationID parameter to the connection string. To connect to a distributed child company, add your ClientID parameter to the connection string.

## To add Intacct data:

1. In ElastiCube Manager, click **Add Data** and then, **Intacct**. The Connect to Intacct window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. Intacct is displayed in the Select Database list.
4. Click **OK**. Sisense connects to Intacct and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

# Intacct Tables

The Sisense Intacct connector allows you to import the following tables into the ElastiCube Manager.

## Available Tables

- AcctRange
- ACHBank
- ActivityLog
- Aisle
- Allocation
- APAccountLabel
- APAdjustment
- APBill
- APBillBatch
- APPymt
- APRecurBill
- APTerm
- ARAccountLabel
- ARAdjustment
- ARInvoice
- ARInvoiceBatch
- ARPaymentBatch
- ARTerm
- Attendee
- Billing\_Contact
- Bin
- BudgetHeader
- CheckingAccount
- Class
- ClassGroup
- Comments
- Contact
- CreditCard
- Customer
- CustomerEmailTemplate

- CustomerGroup
- CustomerVisibility
- CustType
- Department
- DepartmentGroup
- EEarningType
- EAccountLabel
- EExpenses
- EmailTemplate
- Employee
- EmployeeEntityContacts
- EmployeeGroup
- EmployeeOutOfOffice
- EmployeePositionSkill
- EmployeeRate
- EmployeeType
- ExchangeRate
- ExchangeRateTypes
- ExpenseAdjustments
- ExpensePaymentType
- GAAPAdjJrnl
- GLAccount
- GLAcctGrp
- GLAcctGrpMember
- GLBatch
- GLBudget
- GLCoaCatMember
- GLCompGrpMember
- ICRow
- InvDocument
- InvDocumentParams
- InvPriceList
- Item
- ItemGroup
- ItemTaxGroup
- Journal
- Location
- LocationGroup
- Misc\_Supporting\_Document

- OutOfOffice
- PODocument
- PODocumentParams
- POPriceList
- PositionSkill
- ProductLine
- Project
- ProjectGroup
- ProjectResources
- ProjectStatus
- ProjectTransactionRule
- ProjectType
- RecurGLBatch
- RenewalMacro
- ReportingPeriod
- RevRecTemplate
- RevRecTemplMilestone
- SavingsAccount
- Seminar
- Seminar\_Package
- Seminar\_Registration
- SODocument
- SODocumentParams
- SOPriceList
- StatAccount
- StatJournal
- STKitDocument
- Task
- TaskResources
- TaxGroup
- TerritoryGroup
- Timesheet
- TimeType
- TransactionRule
- TransactionRuleDetail
- TrxCurrencies
- UoM
- UserAdjJrnl
- UserInfo

- UserRoles
- Vendor
- VendorGroup
- VendorVisibility
- VendType
- Warehouse

## Available Views

- AllocationEntry
- APAdjustmentItem
- APBillItem
- APBillPayment
- APIUsageDetail
- APIUsageSummary
- APPayment
- APPaymentItem
- APPaymentRequest
- APRecurBillEntry
- ARAdjustmentItem
- ARInvoiceItem
- ARPayment
- ARPaymentItem
- ARRecurInvoice
- ARRecurInvoiceEntry
- BankFee
- BankFeeEntry
- BillingSchedule
- CCTransaction
- CCTransactionEntry
- ChargePayoff
- ChargePayoffEntry
- CNSAcctBal
- Company
- CreditCardFee
- CreditCardFeeEntry
- DDSJob
- Deposit

- DepositEntry
- EExpensesItem
- EExpensesPayment
- EPPayment
- EPPaymentItem
- EPPaymentRequest
- ExchangeRateEntry
- ExpenseAdjustmentsItem
- FundsTransfer
- FundsTransferEntry
- GLAccountBalance
- GLAcctGrpHierarchy
- GLDetail
- GLEntry
- GLResolve
- InvDocumentEntry
- InvDocumentSubtotals
- InvoiceRun
- InvPriceListEntry
- InvRecurSubtotals
- ItemComponent
- OtherReceipts
- OtherReceiptsEntry
- PODocumentEntry
- PODocumentSubtotals
- POPriceListEntry
- PORecurSubtotals
- PSADocumentEntry
- RecurDocumentEntry
- RecurGLEntry
- RevRecChangeLog
- RevRecSchedule
- RevRecScheduleEntry
- RevRecScheduleEntryTask
- SODocumentEntry
- SODocumentSubtotals
- SOPriceListEntry
- SORecurDocument
- SORecurSubtotals

- STKitDocumentEntry
- TerritoryGrpMember
- TimesheetEntry
- TransTmplBatch
- TransTmplEntry
- UoMDetail
- UserRights
- WfpmPaymentRequest

# Connecting to MailChimp

This connector is currently in beta and subject to change.

The Sisense MailChimp connector is a standalone connector that allows you to import data from MailChimp's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to MailChimp's API. To obtain a connection string, you will need to retrieve an API Key from MailChimp or register a MailChimp app.

Once you have connected to MailChimp, you can import a variety of tables from the MailChimp API.

This page describes how to install the MailChimp connector, how to connect to MailChimp with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the MailChimp Connector](#)
- [Connecting to the MailChimp REST API](#)
- [Adding MailChimp Tables to your ElastiCube](#)
- [MailChimp Tables](#)
- [Versions](#)

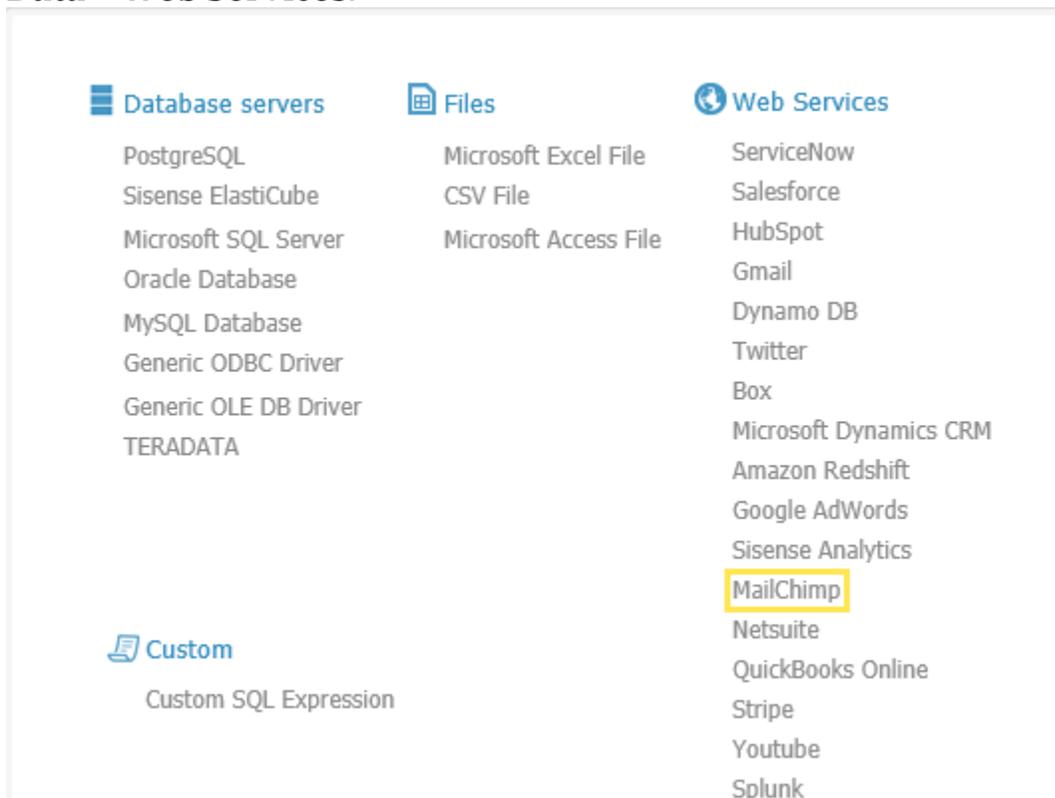
## Installing the MailChimp Connector

Sisense provides the MailChimp connector as a standalone connector that you can download and add to your list of default Sisense connectors.

**To install the MailChimp connector:**

1. [Download](#) the MailChimp installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The MailChimp connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.



## Connecting to the MailChimp REST API

There are two authentication methods available for connecting to MailChimp, through an API Key or through OAuth. Both methods expose the same data, the difference is in what information is required by MailChimp to authenticate your account. For the API Key, all you need is a key provided by MailChimp in your account. For OAuth, you need to register an app with MailChimp, and then pass those details to Sisense.

After you have the relevant information from MailChimp, you create a connection string and connect to MailChimp.

The steps below describe how to retrieve the API Key and OAuth credentials that need to be included in the connection string.

# Using the API Key to Connect to MailChimp

The APIKey grants full access to your MailChimp account. To obtain the APIKey, log into MailChimp and click **Account > Extras > API Keys**.

With the API Key, you can connect to MailChimp by passing the key in a connection string.

# Using OAuth to Connect to MailChimp

To access MailChimp's REST API from Sisense, you must provide valid OAuth MailChimp credentials. These credentials are provided by MailChimp when you register an application.

OAuth requires the authenticating user to interact with MailChimp using the browser. The driver facilitates this in various ways as described below.

## Register Your Application

Follow the steps below to obtain the OAuth client credentials, the OAuthClientId and OAuthClientSecret:

1. Log into your MailChimp account and click **Account > Extras > API Keys > Register and Manage Your Apps**.
2. Enter the information you want to be displayed to users when they are prompted to grant permissions to your application. This information includes your app name, company, and website.
3. If you are making a desktop application, in the Redirect URI box enter `http://127.0.0.1`. If you are making a Web application, in the Redirect URI box enter a URL where you would like users to be redirected after they grant permissions to your application.

# Adding MailChimp Tables to Your ElastiCube

After you have retrieve your API Key or OAuth credentials from MailChimp, you provide the relevant information in a connection string. Sisense uses connection strings to connect to MailChimp and import data into the ElastiCube Manager.

Each connection string contains a authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

For MailChimp, your connection string should include the following parameters:

- **OAuthClientId**: Set this to the consumer key in your app settings.
- **OAuthClientSecret**: Set this to the consumer secret in your app settings.
- **CallbackURL**: Set this to the callback URL you specified in your app settings.
- **InitiateOAuth**: Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken and OAuthAccessTokenSecret.

For API Keys, the connection string to MailChimp is in the following format:

```
jdbc:mailchimp:APIKey=xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx;
```

For OAuth, the connection string to MailChimp is in the following format:

```
jdbc:mailchimp:OAuthClientId=xxxxxxxxxx;OAuthClientSecret=xxxxxxxxxx;CallbackURL=http://127.0.0.1;InitiateOAuth=GETANDREFRESH;
```

**To add MailChimp data:**

1. In ElastiCube Manager, click **Add Data** and then, MailChimp. The Connect to MailChimp window is displayed.

2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. MailChimp is displayed in the Select Database list.
4. Click **OK**. Sisense connects to MailChimp and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

## MailChimp Tables

The Sisense MailChimp connector allows you to import the following tables into the ElastiCube Manager.

### Available Tables

Name	Description
AuthorizedApps	A list of applications authorized to access the account.
AutomationEmailQueues	A summary of the queue for an email in an automation workflow.
AutomationsRemovedSubscribers	A summary of the subscribers removed from

Name	Description
	an automation workflow.
CampaignFeedback	A summary of the comment feedback for a specific campaign.
CampaignFolders	Folders for organizing campaigns
Campaigns	A summary of the campaigns within an account.
ConversationMessages	Messages from a specific conversation. Conversation tracking is a feature available to paid accounts that lets you view replies to your campaigns from inside your MailChimp account.
EcommerceCartLines	A list of an ecommerce cart's lines.
EcommerceCarts	A list of an account's ecommerce carts.
EcommerceCustomers	A list of an account's ecommerce customers.
EcommerceOrderLines	A list of an ecommerce order's lines.
EcommerceOrders	A list of an account's ecommerce orders.
EcommerceProducts	A list of an account's ecommerce products.
EcommerceProductVariants	A list of an ecommerce product's variants.
EcommerceStores	A list of an account's ecommerce stores.
FileManagerFiles	A listing of all available images and files within an account's gallery.
FileManagerFolders	A listing of all variable folders within an account's gallery.

Name	Description
ListInterestCategories	A listing of this list's interest categories.
ListInterests	A list of this category's interests
ListMemberNotes	The last 10 notes for a specific list member, based on date created.
ListMembers	Individuals who are currently or have been previously subscribed to this list, including members who have bounced or unsubscribed.
ListMergeFields	The merge field (formerly merge vars) for a given list. These correspond to merge fields in MailChimp's lists and subscriber profiles.
Lists	A collection of subscriber lists associated with this account. Lists contain subscribers who have opted-in to receive correspondence from you or your organization.
ListSegmentMembers	Individuals who are currently or have been previously subscribed to this list, including members who have bounced or unsubscribed.
ListSegments	A list of available segments.
ListsTwitterLeadGenCards	Twitter Lead Generation Cards for given List
ListsWebhooks	Webhooks configured for the given list.
TemplateFolders	Folders for organizing templates
Templates	A list an account's available templates.

## Available Views

Name	Description
AutomationEmails	A summary of the emails in an automation workflow.
Automations	A summary of the automations within an account.
Conversations	A collection of this account's tracked conversations. Conversation tracking is a feature available to paid accounts that lets you view replies to your campaigns from inside your MailChimp account.
ListAbuse	A collection of abuse complaints for a specific list. An abuse complaint occurs when your recipient clicks to 'report spam' in their email program.
ListActivity	Up to the previous 180 days of daily detailed aggregated activity stats for a given list. Does not include AutoResponder or Automation activity.
ListClients	Top email clients used, as measured by their user-agent string
ListGrowthHistory	A month-by-month summary of a specific list's growth activity.
ListMemberActivity	The last 50 member events for a list.
ListSignupForms	Collection of List Signup Forms
ReportAbuse	A list of abuse complaints for a specific list.
ReportAdvice	A list of feedback based on a campaign's statistics.
ReportClickDetails	A list of URLs and unique IDs included in HTML and plain-text versions of a campaign.

Name	Description
ReportClickDetailsMembers	A collection of members who clicked on a specific link within a campaign.
ReportDomainPerformance	Statistics for the top-performing email domains in a campaign.
ReportEmailActivity	A list of member's subscriber activity in a specific campaign.
ReportLocations	Top open locations for a specific campaign.
Reports	A list of reports containing campaigns marked as Sent.
ReportSentTo	A list of subscribers who were sent a specific campaign.
ReportUnsubscribes	A list of members who have unsubscribed from a specific campaign.

## Versions

Version	Release Date	Description
6.6.0.54	March 28, 2017	Beta release

# Connecting to MongoDB

The ElastiCube Manager enables easy and quick access to databases, tables and views contained within MongoDB databases.

Sisense provides a MongoDB connector for the ElastiCube.

**Note:** Before connecting to MongoDB database with Sisense, please note that MongoDB is an unstructured database, and therefore tables may be flattened with additional tables being created for nested items.

**To connect and define MongoDB, follow these steps:**

- Download the MongoDB ODBC Driver and connect to the data source in the ElastiCube Manager
- Review and edit the DSN schema
- Preview and select the tables generated by the MongoDB ODBC Driver
- (Optional) Change advanced DSN configurations
- (Optional) Connect to MongoDB ODBC Driver using a Connection String (alternative method to step 2 above)

## Upgrading to MongoDB v2.0

The latest version of MongoDB provides enhanced security and improved handling for complex MongoDB models.

The latest version of MongoDB makes significant breaking changes, if you use a version of MongoDB prior to version 2.0, and you are satisfied with your service, you can continue without upgrading.

### When to Upgrade?

In certain circumstances it is recommended or necessary that you upgrade your driver for MongoDB version 2.0, for example:

1. You are new to Sisense and MongoDB
2. You want to connect with the Sisense Native REST Connector

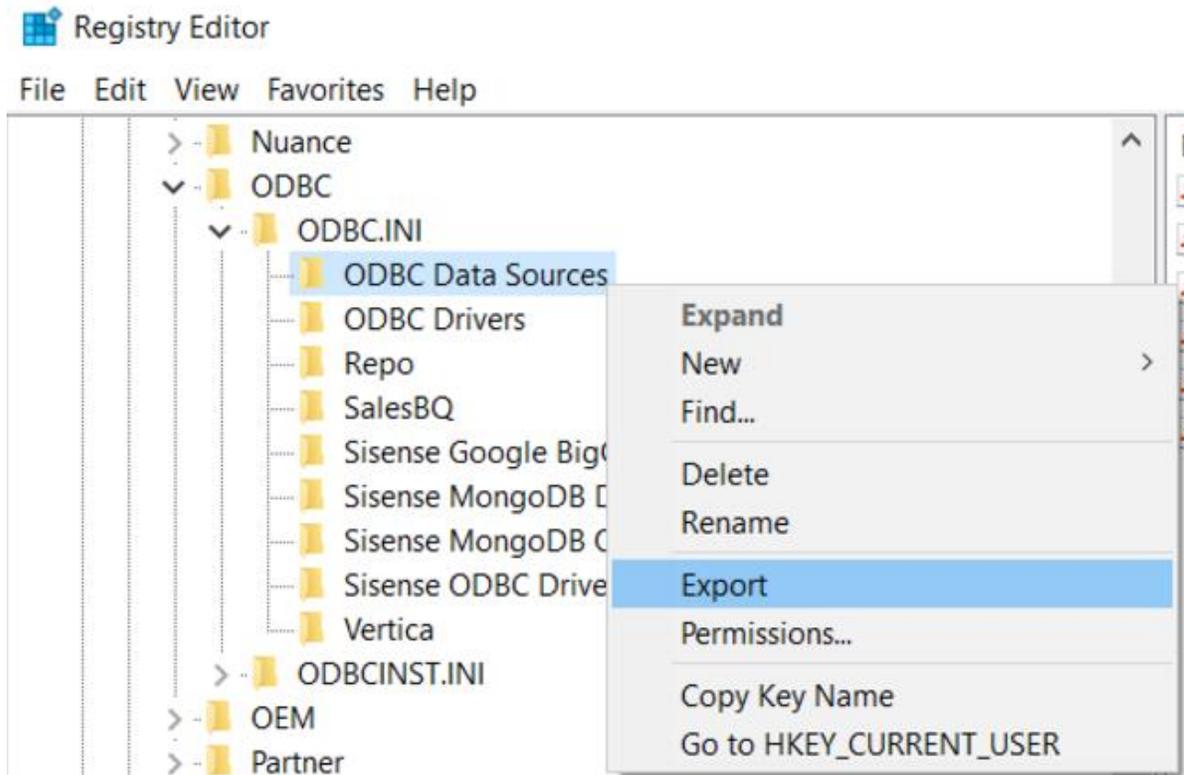
3. You experience security-related issues in MongoDB v1.0
4. You have a really complex MongoDB model that v1.0 can't handle

### Before Upgrading

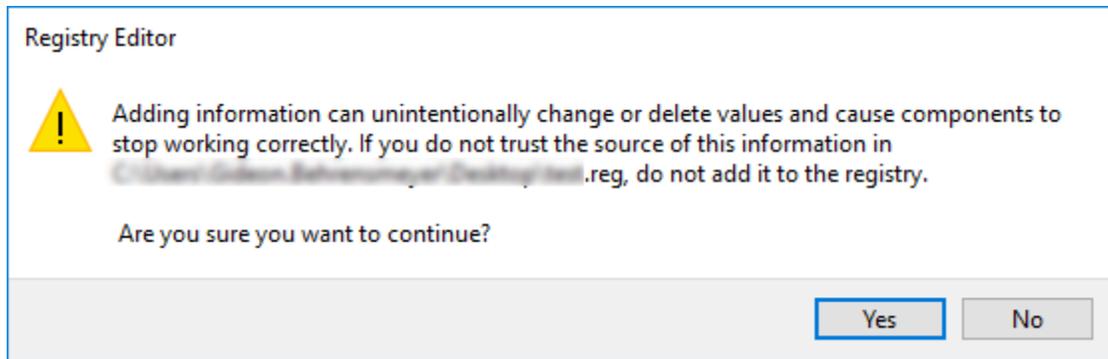
When upgrading to MongoDB v2.0, your current DSN names configured locally may be deleted. As a workaround, you can export the ODBC registry entry and then import after completing the upgrade.

#### To migrate your ODBC Data Source registry file:

1. In Windows, open the Run box with by pressing the Windows key + r.
2. In the Run line, enter regedit.
3. Click **OK**.
4. Navigate to the ODBC Data Source registry entry at the following location:  
`HKEY_LOCAL_MACHINE\SOFTWARE\ODBC\ODBC.INI\ODBC Data Sources`
5. Right click the entry and select Export.



6. Save the registry file on your machine.
7. After installing MongoDB v2, click the exported registry file. The Register Editor warning message is displayed.



8. Click **Yes** to import the entry.

## After Upgrading

After upgrading to MongoDB v2.0, note the following:

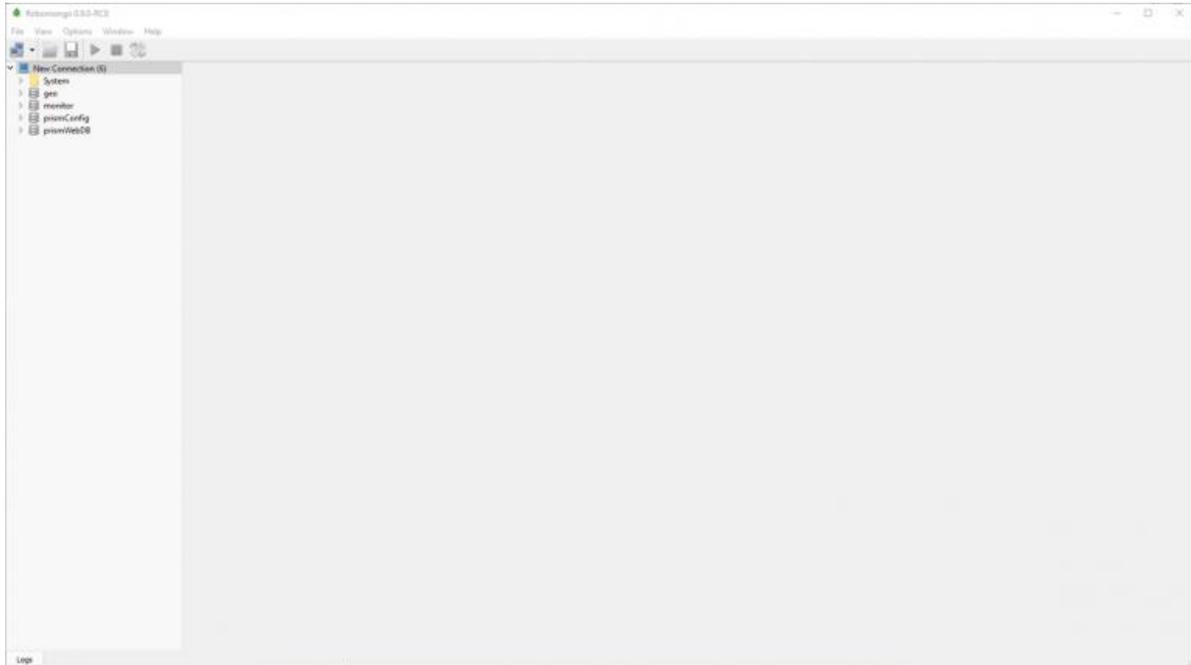
1. The new ODBC schema creation is different than v1.0. MongoDB v2.0 will not create virtual tables (*\_vt\_* tables) in the ElastiCube Manager. As a result, your current ElastiCube schema will not match schemas created before upgrading. Create a new schema using the new v2.0 DSN.
2. The MongoDB v2.0 driver saves a JSON format of the extracted schema, as opposed to the v1.0 driver which uses XML.
3. The MongoDB v2.0 driver stores the schema extracted in MongoDB by default. This allows multiple users to work with the same extracted schema from different servers without extracting a new schema individually each time. However, any changes to the schema affects the schema for everyone else.
4. Storing the schema in the MongoDB means your DB is not **ReadOnly**. Administrators should be aware of this and protect their database with authentication and access control, or communicate this store action and its effects to all users.

## Downloading and Connecting to the MongoDB ODBC Driver

**To connect to MongoDB, complete the following procedures:**

1. [Download](#) and install the MongoDB ODBC Driver.
2. Install the MongoDB driver. **Note:** If you have previously installed an earlier version of the MongoDB driver, you need to uninstall it first, and then install this version.
3. In the ElastiCube Manager, click **Add Data**, and then **Generic ODBC Driver**.

4. Select **DSN**. If a DSN file has already been created, look for it in the dropdown list, and click Test Connection.
5. If you need to add a DSN, see the next section. If your connection has been set up, you can click **OK** and review the tables generated by the MongoDB ODBC Driver.

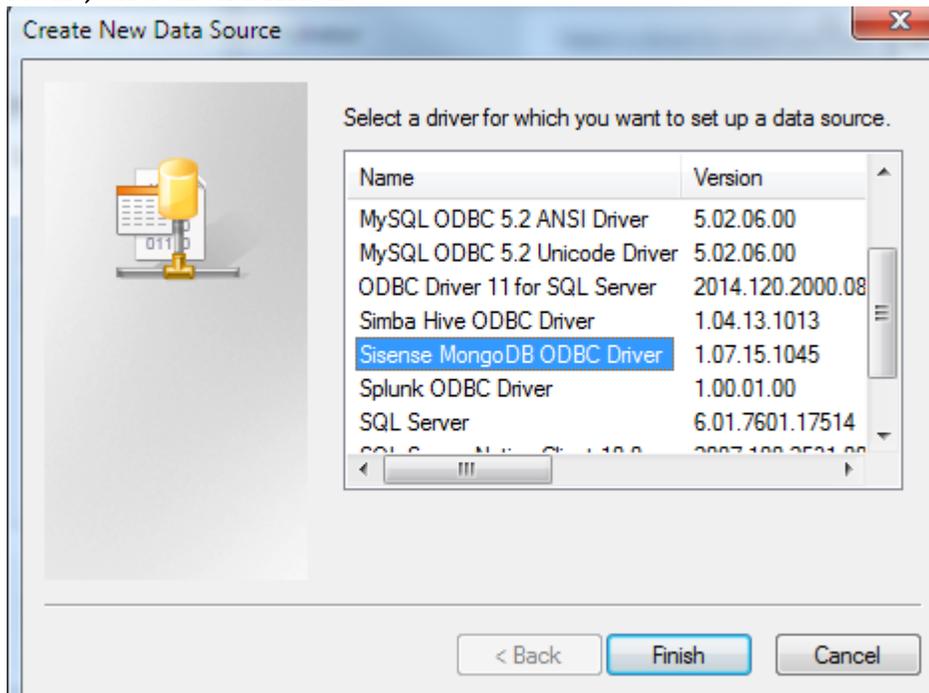


## Adding a DSN

If a DSN has not yet been created, you will need to add it as follows:

1. In Windows, search for and open the application, **ODBC Data Sources**.
2. Under the System DSN tab, click **Add DSN**.

3. Select the previously installed Sisense MongoDB ODBC Driver, and click **Next**, and then **Finish**.



4. Enter the following information:

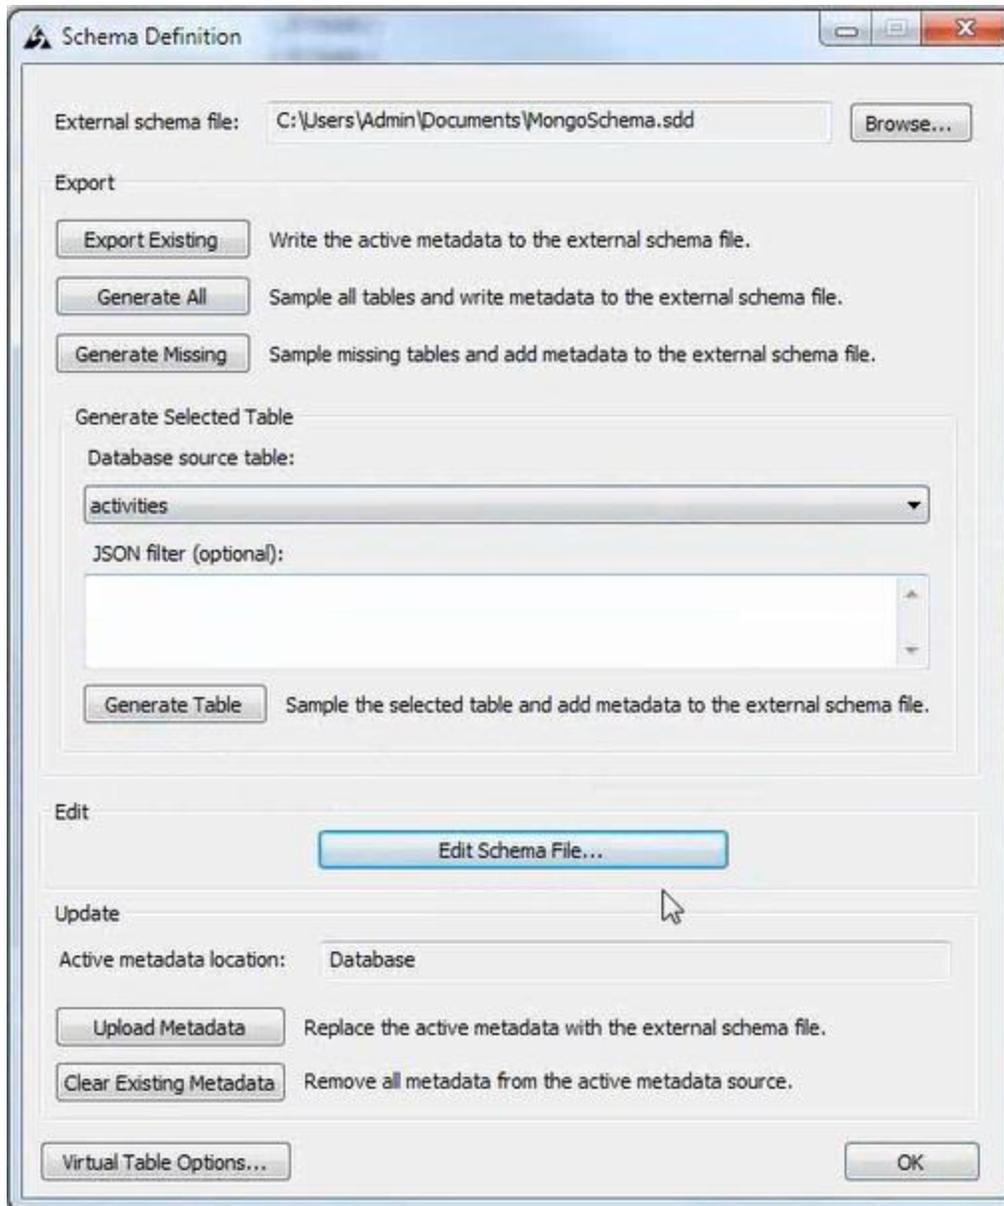
Field	Description
Data Source Name	Enter your name for the DSN that is to be displayed in the DSN dropdown list.
Description	(Optional) Enter your description of the DSN.
Server	Enter the name or IP address of the host where your MongoDB instance is running.
Port	Enter the number of the TCP port that the server uses to listen for client connections.
Database	Enter the name of the database that you want to access. Note: The name of the database is case sensitive.
Connect to Replica Set	If you are connecting to a replica set in your MongoDB implementation, select the Connect to Replica Set check box and then do the following: Replica Set Name: Enter the name of the replica set

Field	Description
	(this is a required field).Secondary Servers: Enter a comma-separated list of the servers in the replica set. You can indicate the TCP port that each server is using to listen for client connections by appending a colon (:) and the port number to the server name or IP address.Read Preference: Select the appropriate option to specify how the driver routes read operations to the members of a replica set.

5. (Optional) If the database that you are connecting to requires authentication, then use the options in the Authentication area to configure authentication as needed. For more information, see [Authentication Options](#).
6. (Optional) To configure advanced driver options, click Advanced Options. For more information, see [Advanced Options](#).
7. (Optional) To configure logging behavior for the driver, click Logging Options. For more information, see [Logging Options](#).
8. Click **Test**. A pop window is displayed that indicates the connection was successful.
9. Click **OK**.
10. In the ODBC Connectivity Properties window, click **Test Connection**.
11. Click **OK** to save and close the ODBC Data Source Administrator.

## DSN Schema Definition

1. Open the MongoDB ODBC Driver from the ODBC Data Source Administrator.
2. Click **Schema Definition** to review the schema and make changes if needed:
  - **External schema file:** The driver automatically generates the schema if not manually defined (path to file is displayed). To edit an existing DSN schema, click **Browse** to locate and open the file in the ODBC Data Source Administrator.
  - **Export options include:**
    - **Export Existing:** Exports metadata that has already been generated.
    - **Generate All:** Exports metadata for all the tables in the database.
    - **Generate Missing:** Exports metadata for tables that have not been included in the schema.



- **Edit Schema File:** Use this option to edit the schema before updating metadata (change field types, hide columns etc.).

Open Save Save As... Close Undo All

Database

- local
  - documents
    - documents\_vt\_main
    - documents\_vt\_any
    - documents\_vt\_lines
    - documents\_vt\_seeBy
    - documents\_vt\_targetDuns
    - documents\_vt\_watchedBy
  - prismWebDB
    - configurations
    - dashboards
      - dashboards\_vt\_main
      - dashboards\_vt\_any
      - dashboards\_vt\_store
      - dashboards\_vt\_columns
      - dashboards\_vt\_elements
      - dashboards\_vt\_filters
      - dashboards\_vt\_filters\_jaqlfilter.exclude.members
      - dashboards\_vt\_filters\_jaqlfilter.members
      - dashboards\_vt\_filters\_levels\_filter.exclude.members
      - dashboards\_vt\_filters\_levels\_filter.members
      - dashboards\_vt\_levels
      - dashboards\_vt\_shares
      - dashboards\_vt\_subcells
    - groups
    - jobs
    - servers
      - servers\_vt\_main
      - servers\_vt\_any
      - servers\_vt\_defaultCubePermissions
    - tags
    - users
    - widgets
      - widgets\_vt\_main
      - widgets\_vt\_any
      - widgets\_vt\_conditions
      - widgets\_vt\_dimensions
      - widgets\_vt\_drill-history
      - widgets\_vt\_items
      - widgets\_vt\_members
      - widgets\_vt\_metadata.drill-history\_throughjaqlfilter.members

Table Properties

ODBC table name: dashboards\_vt\_main Virtual type: Main

ODBC catalog name: prismWebDB Permissions: Read Write

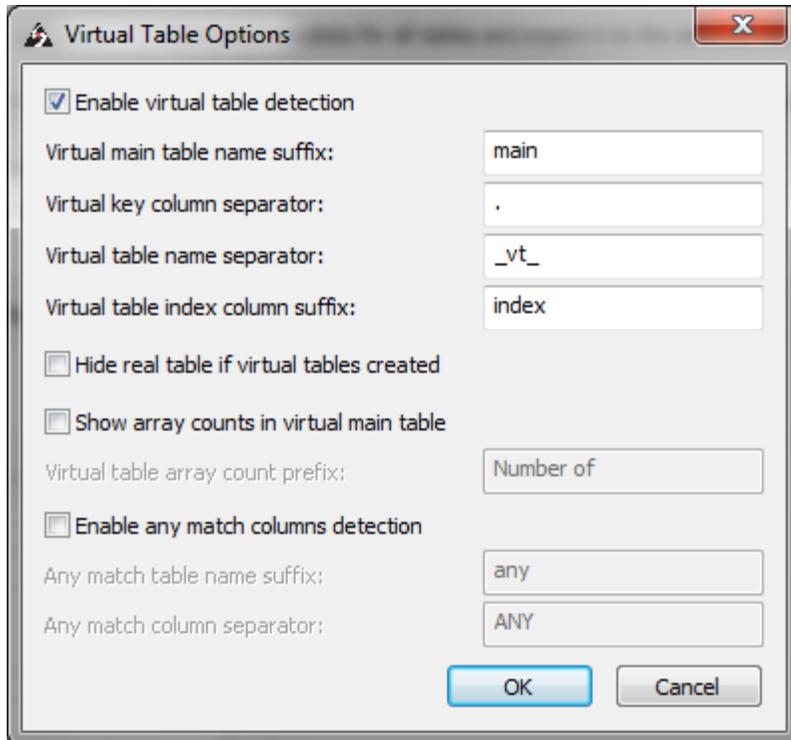
Source table name: dashboards

Source catalog name: prismWebDB

Add a New Column Delete Column(s)

ODBC Column Name	SQL Type	Source Column Name	Source Type	Hide Column	Behavior	Key Type	ODBC Type Hint	Source Nesting Level	Alternative Source Type
dashboards_id	SQL_VARCHAR	_id	ObjectID	<input type="checkbox"/>	Scalar	Unique Fore...	None	0	[String]
created		created	DateTime	<input type="checkbox"/>	Scalar	Not Key	None	0	
datasource	SQL_VARCHAR	datasource	Document	<input checked="" type="checkbox"/>	Container	Not Key	None	0	
datasource_address	SQL_VARCHAR	datasource.address	String	<input type="checkbox"/>	Scalar	Not Key	None	1	
datasource_database	SQL_VARCHAR	datasource.database	String	<input type="checkbox"/>	Scalar	Not Key	None	1	
datasource_id	SQL_VARCHAR	datasource.id	String	<input type="checkbox"/>	Scalar	Not Key	None	1	
datasource_title	SQL_VARCHAR	datasource.title	String	<input type="checkbox"/>	Scalar	Not Key	None	1	
desc	SQL_VARCHAR	desc	String	<input type="checkbox"/>	Scalar	Not Key	None	0	
editing	SQL_BIT	editing	Boolean	<input type="checkbox"/>	Scalar	Not Key	None	0	
Number of filters	SQL_BIGINT	filters	Array	<input checked="" type="checkbox"/>	Container	Not Key	Array Count	0	
instanceType	SQL_VARCHAR	instanceType	String	<input type="checkbox"/>	Scalar	Not Key	None	0	
isPublic	SQL_VARCHAR	isPublic	String	<input type="checkbox"/>	Scalar	Not Key	None	0	
lastUpdated		lastUpdated	DateTime	<input type="checkbox"/>	Scalar	Not Key	None	0	
lastUsed		lastUsed	DateTime	<input type="checkbox"/>	Scalar	Not Key	None	0	
layout	SQL_VARCHAR	layout	Document	<input checked="" type="checkbox"/>	Container	Not Key	None	0	
Number of layout_colu...	SQL_BIGINT	layout.columns	Array	<input checked="" type="checkbox"/>	Container	Not Key	Array Count	1	
layout_type	SQL_VARCHAR	layout.type	String	<input type="checkbox"/>	Scalar	Not Key	None	1	
oid	SQL_VARCHAR	oid	ObjectID	<input type="checkbox"/>	Scalar	Not Key	None	0	
options	SQL_VARCHAR	options	Document	<input checked="" type="checkbox"/>	Container	Not Key	None	0	
options_manual	SQL_BIT	options.manual	Boolean	<input type="checkbox"/>	Scalar	Not Key	None	1	
original	SQL_VARCHAR	original	String	<input type="checkbox"/>	Scalar	Not Key	None	0	
owner	SQL_VARCHAR	owner	ObjectID	<input type="checkbox"/>	Scalar	Not Key	None	0	
parentFolder	SQL_VARCHAR	parentFolder	String	<input type="checkbox"/>	Scalar	Not Key	None	0	
script	SQL_VARCHAR	script	String	<input type="checkbox"/>	Scalar	Not Key	None	0	
Number of shares	SQL_BIGINT	shares	Array	<input checked="" type="checkbox"/>	Container	Not Key	Array Count	0	
source	SQL_VARCHAR	source	String	<input checked="" type="checkbox"/>	Scalar	Not Key	None	0	
style	SQL_VARCHAR	style	String	<input type="checkbox"/>	Scalar	Not Key	None	0	
Number of tags	SQL_BIGINT	tags	Array	<input checked="" type="checkbox"/>	Container	Not Key	Array Count	0	
title	SQL_VARCHAR	title	String	<input type="checkbox"/>	Scalar	Not Key	None	0	

- **Upload Metadata:** Updates metadata after settings are complete. **Current metadata source** displays the source that has been set in **Advanced Options**.
- **Clear Existing Metadata:** Deletes all the metadata that the driver has generated for the MongoDB instance.
- **Virtual Tables Options:** Sisense enables virtual tables creation by default. Virtual tables are created when Sisense detects an array within a field of the main table. Virtual tables can be disabled or custom configured. By default, virtual tables will receive the name `MAINTABLE_vt_FIELD` and the main table will receive the suffix "main". We recommend leaving these settings as is.



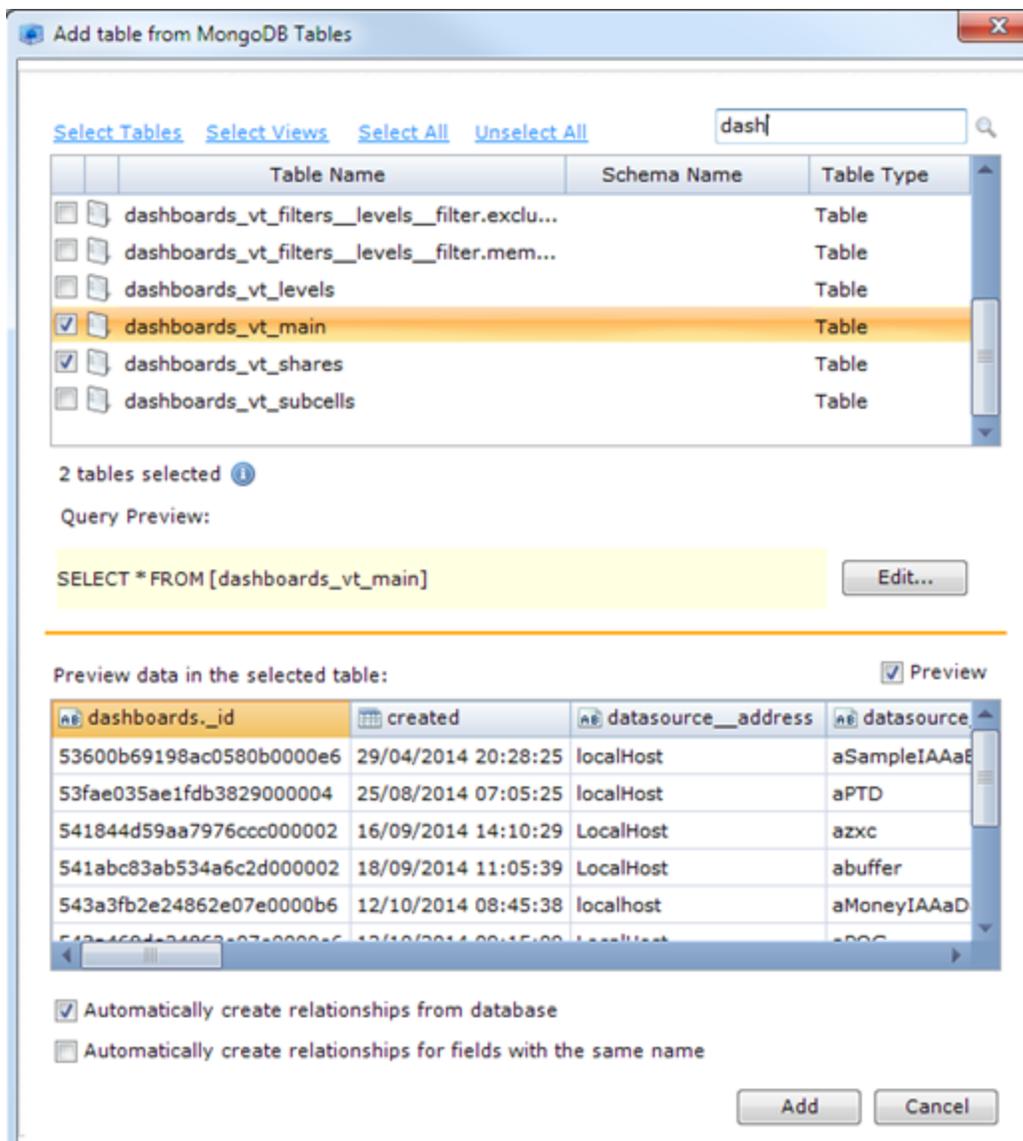
3. When done, click Test Connection, and **OK**.

## Selecting Tables in the ElastiCube Manager

After setting up the DSN, you can preview and customize the query (similar to other ODBC connections) of the imported MongoDB table.

In the ElastiCube Manager You can view the SQL syntax in the Query Preview section. Click Edit to customize it.

Select the tables created by the Sisense MongoDB ODBC Driver (tables including MAINTABLE\_vt\_FIELD), as opposed to the table that includes all the raw data (with suffix main).



## Connecting to Sisense MongoDB ODBC Driver using a Connection String

If you have a driver already defined, then you can connect to Sisense MongoDB ODBC Driver using strings (without using a DSN).

1. Select **Connection String (DSN-Less)**.
2. Type in the connection details as strings, as follows:

**Example with minimum required attributes**

```
Driver=Sisense MongoDB ODBC Driver;
```

```
Host=localhost;
```

```
Port=27017;
```

```
Database=aff;
```

### **Example with with advanced options:**

#### **Driver=Sisense MongoDB ODBC Driver**

```
Host=192.168.100.100;
```

```
Port=27017;
```

```
Database=MyDatabase;
```

```
UID=MyUsername;
```

```
PWD=MyPassword;
```

```
RowsFetchedPerBlock=4096;
```

```
DefaultStringColumnLength=255;
```

```
DefaultContainerColumnLength=511;
```

```
UseSqlWVarchar=0;
```

```
CacheMetadata=1; VirtualTableDetection=1
```

3. When done, click **Test Connection**, and **OK**.

## Authentication Options

Some MongoDB databases require authentication. You can configure the MongoDB ODBC Driver to provide your credentials and authenticate the connection to the database using one of the following methods:

- [Using SCRAM-SHA-1](#)
- [Using Kerberos](#)
- [Using LDAP](#)

### Using SCRAM-SHA-1

You can configure the driver to use the SCRAM-SHA-1 protocol to authenticate the connection. SCRAM-SHA-1 is the default authentication protocol used by MongoDB.

### To configure SCRAM-SHA-1 authentication:

1. To access authentication options, open the ODBC Data Source Administrator where you created the DSN, select the DSN, and then click **Configure**.
2. In the **Mechanism** list, select **MongoDB User Name and Password**.
3. To use a database other than the admin database to check your credentials, type the name of the database in the **Authentication Source** field.
4. In the **Username** field, type an appropriate user name for accessing the MongoDB database.
5. In the **Password** field, type the password corresponding to the user name you typed above.
6. Encrypt your credentials by doing one of the following:
  - If the credentials are used only by the current Windows user, select **Current User Only**.
  - Or, if the credentials are used by all users on the current Windows machine, select **All Users Of This Machine**.
7. To save your settings and close the dialog box, click **OK**.

### Using Kerberos

You can configure the driver to use the Kerberos protocol to authenticate the connection.

Kerberos must be installed and configured before you can use this authentication mechanism. For information about how to install and configure Kerberos, see the MIT Kerberos Documentation: <http://web.mit.edu/kerberos/krb5-latest/doc/>.

### To configure Kerberos authentication:

1. To access authentication options, open the ODBC Data Source Administrator where you created the DSN, select the DSN, and then click **Configure**.
2. In the **Mechanism** list, select **Kerberos**.
3. In the **Service Name** field, type the service name of the MongoDB server.
4. To save your settings and close the dialog box, click **OK**.

## Using LDAP

You can configure the driver to use the LDAP protocol to authenticate the connection.

### To configure LDAP authentication:

1. To access authentication options, open the ODBC Data Source Administrator where you created the DSN, select the DSN, and then click **Configure**.
2. In the **Mechanism** list, select **LDAP**.
3. In the **Username** field, type an appropriate user name for accessing the MongoDB database.
4. In the **Password** field, type the password corresponding to the user name you typed above.
5. Encrypt your credentials by doing one of the following:
  - If the credentials are used only by the current Windows user, select **Current User Only**.
  - Or, if the credentials are used by all users on the current Windows machine, select **All Users Of This Machine**.
6. To save your settings and close the dialog box, click **OK**.

## DSN Advanced Options

You can configure advanced options to modify the behavior of the driver.

### To configure advanced options:

1. To access advanced options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Advanced Options**.
2. To retrieve data using double-buffering instead of single-buffering, select the **Enable Double-Buffering** check box. You can configure the buffer size using the **Documents to fetch per block** field.
3. In the **Documents to fetch per block** field, type the maximum number of documents that a query returns at a time. This setting also determines the buffer size used when double-buffering is enabled.
4. To return MongoDB String data as SQL\_WVARCHAR instead of SQL\_VARCHAR, select the **Expose Strings as SQL\_WVARCHAR** check box.
5. In the **String Column Size** field, type the maximum data length for String columns.

6. To return MongoDB Binary data as SQL\_LONGVARBINARY instead of SQL\_VARBINARY, select the **Expose Binary as SQL\_LONGVARBINARY** check box.
7. In the **Binary Column Size** field, type the maximum data length for Binary columns.
8. To configure the driver to optimize joins between virtual tables and pass filtering and aggregation optimizations to the MongoDB database for handling, select the **Enable Passdown** check box.
9. Use the options in the Metadata area to specify the schema definition to use when connecting to the database:
  - To configure the driver to use a schema definition stored in a JSON file, select **Local File** from the **Mechanism** list, and then click Browse and select the JSON file that you want to use.
  - To configure the driver to use a schema definition stored in the MongoDB database that you are connecting to, in the **Mechanism** list, select **Database**.
10. Use the options in the Sampling area to configure how the driver samples data to generate temporary schema definitions:
  - In the **Sampling Method** list, select the direction in which the driver reads data during sampling. For example, if you select **Forward**, the driver samples data starting from the first record in the database, then samples the next record, and so on.
  - In the **Documents to sample** field, type the maximum number of documents that the driver can sample to generate the schema definition. To sample every document in the database, type 0.
11. In the **Step Size** field, type the interval at which the driver samples a record when scanning through the database. For example, if you set this option to 2, then the driver samples every second record in the database starting from the first record.
12. To configure write-back behavior in the driver, click **Writeback Options**.
  - In the **Batch Size** field, type the maximum number of documents that the driver can handle at one time during a write operation.
  - Use the options in the Write Concern Settings area to configure how the driver reports the success of a write operation:
  - In the **Write Concern** field, type the total number of primary and secondary servers that must acknowledge a write operation in order for the driver to report a successful write operation.
  - In the **Timeout** field, type the maximum number of seconds that the driver waits for a secondary server to acknowledge a write operation before reporting that the operation has failed.

- To require the data to be committed to the journal before a write operation can be acknowledged, select the **Journalled Writes** check box.
13. To save your settings and close the Advanced Options dialog box, click **OK**.
  14. To close the MongoDB ODBC Driver DSN Setup dialog box, click **OK**.

## Logging Options

To help troubleshoot issues, you can enable logging. In addition to functionality provided in the MongoDB ODBC Driver, the ODBC Data Source Administrator provides tracing functionality.

### To enable driver logging:

1. To access logging options, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **Logging Options**.
2. From the **Log Level** list, select the logging level corresponding to the amount of information that you want to include in log files:

Field	Description
Log Level	Select the Log Level. There are seven possible options: <b>LOG_OFF</b> : Disables all logging. <b>LOG_FATAL</b> : Logs very severe error events that lead the driver to abort. <b>LOG_ERROR</b> : Logs error events that might still allow the driver to continue running. <b>LOG_WARNING</b> : Logs potentially harmful situations. <b>LOG_INFO</b> : Logs general information that describes the progress of the driver. <b>LOG_DEBUG</b> : Logs detailed information that is useful for debugging the driver. <b>LOG_TRACE</b> : Logs more detailed information than the DEBUG level.
Log Path	Enter the full path to the folder where you want to save log files. OR Click <b>Browse</b> and select the folder where you want to save log files.

Field	Description
Log Rotation	<p>Enter the maximum number of log files to keep in the Max Number Files field.</p> <p>Note: After the maximum number of log files is reached, each time an additional file is created, the driver deletes the oldest log file.</p> <p>Enter the maximum size of each log file in megabytes (MB) in the Max File Size field.</p> <p>Note: After the maximum file size is reached, the driver creates a new file and continues logging.</p>

3. In the **Log Path** field, specify the full path to the folder where you want to save log files. You can type the path into the field, or click **Browse** and then browse to select the folder.
4. In the **Max Number Files** field, type the maximum number of log files to keep.
5. In the **Max File Size** field, type the maximum size of each log file in megabytes (MB).
6. Click **OK**.
7. Restart your ODBC application to make sure that the new settings take effect.

The MongoDB ODBC Driver produces a log file named `mongodboDBC_driver.log` at the location that you specify in the Log Path field.

## SSL Options

If you are connecting to a MongoDB server that has Secure Sockets Layer (SSL) enabled, then you can configure the driver to connect to an SSL-enabled socket. When connecting to a server over SSL, the driver supports identity verification between the client and the server.

### Configuring an SSL Connection without Identity Verification

You can configure a connection that uses SSL but does not verify the identity of the client or the server.

**To configure an SSL connection without verification:**

1. To access the SSL options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **SSL Options**.
2. Select the **Enable SSL** check box.
3. Select the **Allow Self-Signed Certificates** check box.
4. To save your settings and close the dialog box, click **OK**.

### Configuring One-way SSL Verification

You can configure one-way verification so that the client verifies the identity of the MongoDB server.

#### To configure one-way SSL verification:

1. To access the SSL options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **SSL Options**.
2. Select the **Enable SSL** check box.
3. Choose one:
  - To verify the server using a certificate from a specific PEM file, in the **Certificate Authority File** field, specify the full path of the PEM file.
  - Or, to verify the server using certificates stored in multiple PEM files, in the **Certificate Authority Directory** field, specify the full path to the directory where the PEM files are located.
4. In the **Certificate Revocation List File** field, specify the full path of the PEM file containing the list of revoked certificates.
5. To save your settings and close the dialog box, click **OK**.

### Configuring Two-way SSL Verification

You can configure two-way SSL verification so that the client and the MongoDB server verify each other.

#### To configure two-way SSL verification:

1. To access the SSL options for a DSN, open the ODBC Data Source Administrator where you created the DSN, then select the DSN, then click **Configure**, and then click **SSL Options**.
2. Select the **Enable SSL** check box.
3. In the **PEM Key File** field, specify the full path of the PEM file containing the certificate for verifying the client.
4. If the client certificate is protected with a password, type the password in the **PEM Key Password** field.
5. Choose one:

- To verify the server using a certificate from a specific PEM file, in the **Certificate Authority File** field, specify the full path of the PEM file.
  - Or, to verify the server using certificates stored in multiple PEM files, in the **Certificate Authority Directory** field, specify the full path to the directory where the PEM files are located.
6. In the **Certificate Revocation List File** field, specify the full path of the PEM file containing the list of revoked certificates.
  7. To save your settings and close the dialog box, click **OK**.

## Versions

Version	Release Date	Improvements
6.6.1.2	May 19, 2017	Support added for MongoDB 3.4 For more information, see the complete MongoDB Release Notes

# Connecting to Microsoft Active Directory

The Sisense Microsoft Active Directory connector is a standalone connector that allows you to import data from Microsoft Active Directory's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the Microsoft Active Directory tables.

Once you have connected to Microsoft Active Directory, you can import a variety of tables from the Microsoft Active Directory API.

This page describes how to install the Microsoft Active Directory connector, how to connect to Microsoft Active Directory with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the Microsoft Active Directory Connector](#)
- [Connecting to the Microsoft Active Directory](#)
- [Microsoft Active Directory Tables](#)
- [Limitations](#)

## Installing the Microsoft Active Directory Connector

Sisense provides the Microsoft Active Directory connector as a standalone connector that you can download and add to your list of default Sisense connectors.

**To install the Microsoft Active Directory connector:**

1. [Download](#) the Microsoft Active Directory installation file.

2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The Microsoft Active Directory connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.

## Connecting to the Microsoft Active Directory

Sisense uses connection strings to connect to Microsoft Active Directory and import data into the ElastiCube Manager.

The connection string to connect to Microsoft Active Directory has the following structure:

```
jdbc:activedirectory:Property1=Value1;Property2=Value2;
```

The following is an example of a Microsoft Active Directory connection string:

```
jdbc:activedirectory:User=xxxxxxxxx;Password=xxxxxxxxx;Server=xxxxxxxxx;Port=389;BaseDN=CN=xxxx,DC=xxxx,DC=xxxxxx,DC=xxxxxxxx;
```

To establish a connection, the following properties under the Authentication section must be provided:

- Valid User and Password credentials (e.g., Domain\\BobF or cn=BobF,ou=Employees,dc=Domain).
- Server information, including the IP or host name of the Server and the Port.
- BaseDN will limit the scope of LDAP searches to the height of the distinguished name provided.

**Note:** Specifying a narrow BaseDN may greatly increase performance; for example, cn=users,dc=domain will only return results contained within cn=users and its children.

- If you define your own custom schemas to work with your ActiveDirectory object classes, set Location to the path to the folder containing the schema files.

**Note:** To switch between accounts, you need to delete the file OAuthsettings.txt file located at ...\\Users\\xxx\\AppData\\Roaming\\CData\\Microsoft Active Directory Data Provider.

**To add Microsoft Active Directory data:**

1. In ElastiCube Manager, click **Add Data** and then, Microsoft Active Directory. The Connect to Microsoft Active Directory window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. Microsoft Active Directory is displayed in the Select Database list.
4. Click **OK**. Sisense connects to Microsoft Active Directory and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**.

The tables are displayed in the ElastiCube Manager.

## Switching between Accounts

When you connect to the Microsoft Active Directory data source, Sisense saves your OAuth values in the file OAuthsettings.txt file located at

...\\Users\\xxx\\AppData\\Roaming\\CData\\Microsoft Active Directory Data

Provider on your Sisense server. To connect to the Microsoft Active Directory data source with another user on the same machine, you must delete the OAuthsettings.txt file. Sisense will then generate a new file for that user.

Another option to support multiple users is to define the location and file name of an OAuthsettings file for each unique user in your connection string through

the OAuthSettingsLocationparameter. When each user connects to the data

source, Sisense generates the OAuth file with the file name you specify in the

location you define. In the examples below, two users are allowed to access the

Microsoft Active Directory data source and for each user, Sisense generates a file that contains that user's OAuth values in the location defined in the string.

```
jdbc:activedirectory:OAuthSettingsLocation=C:\\MicrosoftActiveDirectory\\auth\\john.txt;OAuthClientId=11276856774486;
```

```
OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
```

```
jdbc:activedirectory:OAuthSettingsLocation=C:\MicrosoftActiveDirectory\auth\sally.txt;OAuthClientId=11276856774486;
```

```
OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
```

In the example above, two OAuth files are created, one for John and one for Sally in the location C:\Microsoft Active Directory\auth\.

This is useful if you support many users who each need to access the Microsoft Active Directory data source.

## Microsoft Active Directory Tables

Microsoft Active Directory's RESTful APIs expose the following Microsoft Active Directory tables that you can import into the ElastiCube Manager through the Sisense Microsoft Active Directory connector:

### Available Tables

Name	Description
Account	The account object class is used to define entries that represent computer accounts.
ApplicationEntity	X.500 base class for applications: Directory Service only uses subclass MSFT-DSA.
ApplicationProcess	X.500 base class for applications: Exchange only uses

Name	Description
	subclass DSA-Application.
ApplicationSettings	Base class for server-specific application settings.
ApplicationSiteSettings	Contains all site-specific settings.
ApplicationVersion	Can be used by application developers to store version information about their application or its schema.
BuiltinDomain	The container that holds the default groups for a domain.
CertificationAuthority	Represents a process that issues public key certificates, for example, a Certificate Server.
Computer	This class represents a computer account in the domain.
Contact	This class contains information about a person or company that you may need to contact on a regular basis.
Events	Query the Events for a Target based on either the Target or SearchTerms. May require the user_events permission.
CRLDistributionPoint	The object holding Certificate, Authority, and Delta Revocation lists.
DHCPClass	Represents a DHCP Server (or set of servers).
DnsNode	Holds the DNS resource records for a single host.
DnsZone	The container for DNS Nodes. Holds zone metadata.
Domain	Contains information about a domain.
DomainDNS	Windows NT domain with DNS-based (DC=) naming.

Name	Description
DomainPolicy	Defines the local security authority policy for one or more domains.
DomainRelatedObject	The domainRelatedObject object class is used to define an entry that represents a series of documents.
ForeignSecurityPrincipal	The Security Principal from an external source.
Group	Stores a list of user names. Used to apply security principals on resources.
GroupOfNames	Used to define entries that represent an unordered set of names that represent individual objects or other groups of names.
GroupOfUniqueNames	Defines the entries for a group of unique names. In general, used to store account objects.
GroupPolicyContainer	This represents the Group Policy Object. It is used to define group polices.
IpHost	Represents an abstraction of a host or other IP device.
IpNetwork	Represents an abstraction of a network. The distinguished name value of the Common-Name attribute denotes the canonical name of the network.
Organization	Stores information about a company or organization.
OrganizationalPerson	This class is used for objects that contain organizational information about a user, such as the employee number, department, manager, title, office address, and so on.
OrganizationalRole	This class is used for objects that contain information

Name	Description
	that pertains to a position or role within an organization, such as a system administrator, manager, and so on. It can also be used for a nonhuman identity in an organization.
OrganizationalUnit	A container for storing users, computers, and other account objects.
Person	Contains personal information about a user.
PosixAccount	Represents an abstraction of an account with Portable Operating System Interface (POSIX) attributes.
PosixGroup	Represents an abstraction of a group of accounts.
PrintQueue	Contains information about a print queue.
SecurityObject	This is an auxiliary class that is used to identify security principals.
SecurityPrincipal	Contains the security information for an object.
Server	This class represents a server computer in a site.
Site	A container for storing server objects. Represents a physical location that contains computers. Used to manage replication.
Top	The top level class from which all classes are derived.
TrustedDomain	An object that represents a domain trusted by (or trusting) the local domain.
User	This class is used to store information about an employee or contractor who works for an organization. It is also possible to apply this class to long term visitors.

## Limitations

1. Accumulated builds are supported because all tables have string columns.
2. Aggregate functions are not supported

# Connecting to Microsoft Dynamics CRM

The Sisense Microsoft Dynamics CRM connector is a standalone connector that allows you to import data from Microsoft Dynamics CRM's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the Microsoft Dynamics CRM tables. To obtain a connection string, you will need to obtain your credentials from Microsoft Dynamics CRM.

Once you have connected to Microsoft Dynamics CRM, you can import a variety of tables from the Microsoft Dynamics CRM API.

This page describes how to install the Microsoft Dynamics CRM connector, how to connect to Microsoft Dynamics CRM with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the Microsoft Dynamics CRM Connector](#)
- [Connecting to the Microsoft Dynamics CRM](#)
- [Adding Microsoft Dynamics CRM Tables to your ElastiCube](#)
- [Microsoft Dynamics CRM Tables](#)

## Installing the Microsoft Dynamics CRM Connector

Sisense provides the Microsoft Dynamics CRM connector as a standalone connector that you can download and add to your list of default Sisense connectors.

**To install the Microsoft Dynamics CRM connector:**

1. [Download](#) the Microsoft Dynamics CRM installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The Microsoft Dynamics CRM connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.

## Connecting to the Microsoft Dynamics CRM

The connection string used to connect to the Microsoft Dynamics CRM connects to various instances you have in the Microsoft Dynamics CRM. By providing your Microsoft Dynamics CRM credentials as values in the connection string, you can add Microsoft Dynamic CRM tables to the ElastiCube.

### Adding Microsoft Dynamics CRM Tables to your ElastiCube

Sisense uses connection strings to connect to Microsoft Dynamics CRM and import data into the ElastiCube Manager.

The connection string to connect to Microsoft Dynamics CRM has the following structure:

```
jdbc:dynamicscrm:Property1=Value1;Property2=Value2;
```

The following is an example of a Microsoft Dynamics CRM connection string:

```
jdbc:dynamicscrm:User=xxxxxxxxxx;Password=xxxxxxxxxx;URL=https://xxxxxxxxxx.dynamics.com/;CRMVersion=CRM Online Office 365;
```

To connect to your instance, set the User and Password properties, under the Authentication section, to valid Dynamics CRM user credentials and set the Url to a valid Dynamics CRM server organization root.

Additionally, set the CRMVersion property. Supported versions are CRM Online Office 365, CRM 4.0, CRM 2011, CRM 2013, CRM 2015.

**Note:** To switch between accounts, you need to delete the file OAuthsettings.txt file located at ...\\Users\\xxx\\AppData\\Roaming\\CData\\Microsoft Dynamics CRM Data Provider or you can add OAuthSettingsLocation to the connection string. See Switching between Accounts for more information.

**To add Microsoft Dynamics CRM data:**

1. In ElastiCube Manager, click **Add Data** and then, **Microsoft Dynamics CRM**. The Connect to Microsoft Dynamics CRM window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. Microsoft Dynamics CRM is displayed in the Select Database list.
4. Click **OK**. Sisense connects to Microsoft Dynamics CRM and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**.

The tables are displayed in the ElastiCube Managerr.

## Switching between Accounts

When you connect to the Microsoft Dynamics CRM data source, Sisense saves your OAuth values in the file OAuthsettings.txt file located at ...\\Users\\xxx\\AppData\\Roaming\\CData\\Microsoft Dynamics CRM Data Provider on your Sisense server. To connect to the Microsoft Dynamics CRM data source with another user on the same machine, you must delete the OAuthsettings.txt file. Sisense will then generate a new file for that user.

Another option to support multiple users is to define the location and file name of an OAuthsettings file for each unique user in your connection string through the OAuthSettingsLocation parameter. When each user connects to the data source, Sisense generates the OAuth file with the file name you specify in the location you define. In the examples below, two users are allowed to access the Microsoft Dynamics CRM data source and for each user, Sisense generates a file that contains that user's OAuth values in the location defined in the string.

```

jdbc:MicrosoftDynamicsCRM:OAuthSettingsLocation=C:\MicrosoftDynamicsCRM\auth\john.txt;OAuthClientId=11276856774486
;
OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost
/;
jdbc:MicrosoftDynamicsCRM:OAuthSettingsLocation=C:\MicrosoftDynamicsCRM\auth\sally.txt;OAuthClientId=11276856774486
;
OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost
/;

```

In the example above, two OAuth files are created, one for John and one for Sally in the location C:\MicrosoftDynamicsCRM\auth\.

This is useful if you support many users who each need to access the Microsoft Dynamics CRM data source.

## Microsoft Dynamics CRM Tables

Microsoft Dynamics CRM exposes the following Microsoft Dynamics CRM tables that you can import into the ElastiCube Manager through the Sisense Microsoft Dynamics CRM connector:

### Available Tables

Name	Description
Accounts	Create, update, delete, and query Account entities in Dynamics CRM.
ActivityMimeAttachment	Create, update, delete, and query

Name	Description
	ActivityMimeAttachment entities in Dynamics CRM.
ActivityParty	Create, update, delete, and query ActivityParty entities in Dynamics CRM.
ActivityPointer	Create, update, delete, and query ActivityPointer entities in Dynamics CRM.
Annotation	Create, update, delete, and query Annotation entities in Dynamics CRM.
AnnualFiscalCalendar	Create, update, delete, and query annual fiscal calendar entities in Dynamics CRM.
ApplicationFile	Create, update, delete, and query ApplicationFile entities in Dynamics CRM.
Appointment	Create, update, delete, and query Appointment entities in Dynamics CRM.
AsyncOperation	Create, update, delete, and query AsyncOperation entities in Dynamics CRM.
Attachment	Create, update, delete, and query Attachment entities in Dynamics CRM.
AttributeMap	Create, update, delete, and query AttributeMap entities in Dynamics CRM.
Audit	Create, update, delete, and query Audit entities in Dynamics CRM.
BulkDeleteFailure	Create, update, delete, and query

Name	Description
	BulkDeleteFailure entities in Dynamics CRM.
BulkDeleteOperation	Create, update, delete, and query BulkDeleteOperation entities in Dynamics CRM.
BulkOperation	Create, update, delete, and query BulkOperation entities in Dynamics CRM.
BulkOperationLog	This is a table representing the BulkOperationLog entities in Dynamics CRM.
BusinessUnit	This is a table representing the BusinessUnit entities in Dynamics CRM.
BusinessUnitMap	This is a table representing the BusinessUnitMap entities in Dynamics CRM.
BusinessUnitNewsArticle	This is a table representing the BusinessUnitNewsArticle entities in Dynamics CRM.
Calendar	This is a table representing the Calendar entities in Dynamics CRM.
CalendarRule	This is a table representing the CalendarRule entities in Dynamics CRM.
Campaign	This is a table representing the Campaign entities in Dynamics CRM.

Name	Description
CampaignActivity	This is a table representing the CampaignActivityItem entities in Dynamics CRM.
CampaignActivityItem	This is a table representing the CampaignActivityItem entities in Dynamics CRM.
CampaignItem	This is a table representing the CampaignItem entities in Dynamics CRM.
CampaignResponse	This is a table representing the CampaignResponse entities in Dynamics CRM.
ClientUpdate	This is a table representing the ClientUpdate entities in Dynamics CRM.
ColumnMapping	This is a table representing the ColumnMapping entities in Dynamics CRM.
Commitment	This is a table representing the Commitment entities in Dynamics CRM.
Competitor	This is a table representing the Competitor entities in Dynamics CRM.
CompetitorAddress	This is a table representing the CompetitorAddress entities in Dynamics CRM.
CompetitorProduct	This is a table representing the CompetitorProduct entities in Dynamics

Name	Description
	CRM.
CompetitorSalesLiterature	This is a table representing the CompetitorSalesLiterature entities in Dynamics CRM.
Connection	This is a table representing the Connection entities in Dynamics CRM.
ConnectionRole	This is a table representing the ConnectionRole entities in Dynamics CRM.
ConnectionRoleAssociation	This is a table representing the ConnectionRoleAssociation entities in Dynamics CRM.
ConnectionRoleObjectTypeCode	This is a table representing the ConnectionRoleObjectTypeCode entities in Dynamics CRM.
ConstraintBasedGroup	This is a table representing the ConstraintBasedGroup entities in Dynamics CRM.
Contact	This is a table representing the Contact entities in Dynamics CRM.
ContactInvoices	This is a table representing the ContactInvoices entities in Dynamics CRM.
ContactLeads	This is a table representing the ContactLeads entities in Dynamics CRM.

Name	Description
ContactOrders	This is a table representing the ContactOrders entities in Dynamics CRM.
ContactQuotes	This is a table representing the ContactQuotes entities in Dynamics CRM.
Contract	This is a table representing the Contract entities in Dynamics CRM.
ContractDetail	This is a table representing the ContractDetail entities in Dynamics CRM.
ContractTemplate	This is a table representing the ContractTemplate entities in Dynamics CRM.
CustomerAddress	This is a table representing the CustomerAddress entities in Dynamics CRM.
CustomerOpportunityRole	This is a table representing the CustomerOpportunityRole entities in Dynamics CRM.
CustomerRelationship	This is a table representing the CustomerRelationship entities in Dynamics CRM.
Dependency	This is a table representing the Dependency entities in Dynamics CRM.
DependencyNode	This is a table representing the DependencyNode entities in Dynamics CRM.

Name	Description
	CRM.
Discount	This is a table representing the Discount entities in Dynamics CRM.
DiscountType	This is a table representing the DiscountType entities in Dynamics CRM.
DisplayString	This is a table representing the DisplayString entities in Dynamics CRM.
DisplayStringMap	This is a table representing the DisplayStringMap entities in Dynamics CRM.
DocumentIndex	This is a table representing the DocumentIndex entities in Dynamics CRM.
DuplicateRecord	This is a table representing the DuplicateRecord entities in Dynamics CRM.
DuplicateRule	This is a table representing the DuplicateRule entities in Dynamics CRM.
DuplicateRuleCondition	This is a table representing the DuplicateRuleCondition entities in Dynamics CRM.
Email	This is a table representing the Email entities in Dynamics CRM.
EmailHash	This is a table representing the

Name	Description
	EmailHash entities in Dynamics CRM.
EmailSearch	This is a table representing the EmailSearch entities in Dynamics CRM.
EntityMap	This is a table representing the EntityMap entities in Dynamics CRM.
Equipment	This is a table representing the Equipment entities in Dynamics CRM.
Fax	This is a table representing the Fax entities in Dynamics CRM.
FieldPermission	This is a table representing the FieldPermission entities in Dynamics CRM.
FieldSecurityProfile	This is a table representing the FieldSecurityProfile entities in Dynamics CRM.
FilterTemplate	This is a table representing the FilterTemplate entities in Dynamics CRM.
FixedMonthlyFiscalCalendar	This is a table representing the FixedMonthlyFiscalCalendar entities in Dynamics CRM.
Goal	This is a table representing the Goal entities in Dynamics CRM.
GoalRollupQuery	This is a table representing the GoalRollupQuery entities in Dynamics

Name	Description
	CRM.
Import	This is a table representing the ImportData entities in Dynamics CRM.
ImportData	This is a table representing the ImportEntityMapping entities in Dynamics CRM.
ImportEntityMapping	This is a table representing the ImportEntityMapping entities in Dynamics CRM.
ImportFile	This is a table representing the ImportFile entities in Dynamics CRM.
ImportJob	This is a table representing the ImportJob entities in Dynamics CRM.
ImportLog	This is a table representing the ImportLog entities in Dynamics CRM.
ImportMap	This is a table representing the ImportMap entities in Dynamics CRM.
Incident	This is a table representing the Incident entities in Dynamics CRM.
IncidentResolution	This is a table representing the IncidentResolution entities in Dynamics CRM.
IntegrationStatus	This is a table representing the IntegrationStatus entities in Dynamics

Name	Description
	CRM.
InternalAddress	This is a table representing the InternalAddress entities in Dynamics CRM.
InterProcessLock	This is a table representing the InterProcessLock entities in Dynamics CRM.
InvalidDependency	This is a table representing the InvalidDependency entities in Dynamics CRM.
Invoice	This is a table representing the Invoice entities in Dynamics CRM.
InvoiceDetail	This is a table representing the InvoiceDetail entities in Dynamics CRM.
IsvConfig	This is a table representing the IsvConfig entities in Dynamics CRM.
KbArticle	This is a table representing the KbArticle entities in Dynamics CRM.
KbArticleComment	This is a table representing the KbArticleComment entities in Dynamics CRM.
KbArticleTemplate	This is a table representing the KbArticleTemplate entities in Dynamics CRM.

Name	Description
Lead	This is a table representing the Lead entities in Dynamics CRM.
LeadAddress	This is a table representing the LeadAddress entities in Dynamics CRM.
LeadCompetitors	This is a table representing the LeadCompetitors entities in Dynamics CRM.
LeadProduct	This is a table representing the LeadProduct entities in Dynamics CRM.
Letter	This is a table representing the Letter entities in Dynamics CRM.
License	This is a table representing the License entities in Dynamics CRM.
List	This is a table representing the List entities in Dynamics CRM.
ListMember	This is a table representing the ListMember entities in Dynamics CRM.
LookUpMapping	This is a table representing the LookUpMapping entities in Dynamics CRM.
MailMergeTemplate	This is a table representing the MailMergeTemplate entities in Dynamics CRM.
Metric	This is a table representing the Metric

Name	Description
	entities in Dynamics CRM.
MonthlyFiscalCalendar	This is a table representing the MonthlyFiscalCalendar entities in Dynamics CRM.
Notification	This is a table representing the Notification entities in Dynamics CRM.
Opportunity	This is a table representing the Opportunity entities in Dynamics CRM.
OpportunityClose	This is a table representing the OpportunityClose entities in Dynamics CRM.
OpportunityCompetitors	This is a table representing the OpportunityCompetitors entities in Dynamics CRM.
OpportunityProduct	This is a table representing the OpportunityProduct entities in Dynamics CRM.
OrderClose	This is a table representing the OrderClose entities in Dynamics CRM.
Organization	This is a table representing the Organization entities in Dynamics CRM.
OrganizationStatistic	This is a table representing the OrganizationStatistic entities in Dynamics CRM.

Name	Description
OrganizationUI	This is a table representing the OrganizationUI entities in Dynamics CRM.
Owner	This is a table representing the Owner entities in Dynamics CRM.
OwnerMapping	This is a table representing the OwnerMapping entities in Dynamics CRM.
PhoneCall	This is a table representing the PhoneCall entities in Dynamics CRM.
PickListMapping	This is a table representing the PickListMapping entities in Dynamics CRM.
PluginAssembly	This is a table representing the PluginAssembly entities in Dynamics CRM.
plug-intype	This is a table representing the plug-in type entities in Dynamics CRM.
PluginTypeStatistic	This is a table representing the PluginTypeStatistic entities in Dynamics CRM.
PriceLevel	This is a table representing the PriceLevel entities in Dynamics CRM.
PrincipalAttributeAccessMap	This is a table representing the PrincipalAttributeAccessMap entities in

Name	Description
	Dynamics CRM.
PrincipalEntityMap	This is a table representing the PrincipalEntityMap entities in Dynamics CRM.
PrincipalObjectAccess	This is a table representing the PrincipalObjectAccess entities in Dynamics CRM.
PrincipalObjectAttributeAccess	This is a table representing the PrincipalObjectAttributeAccess entities in Dynamics CRM.
Privilege	This is a table representing the Privilege entities in Dynamics CRM.
PrivilegeObjectTypeCodes	This is a table representing the PrivilegeObjectTypeCodes entities in Dynamics CRM.
ProcessSession	This is a table representing the ProcessSession entities in Dynamics CRM.
Product	This is a table representing the Product entities in Dynamics CRM.
ProductAssociation	This is a table representing the ProductAssociation entities in Dynamics CRM.
ProductPriceLevel	This is a table representing the ProductPriceLevel entities in Dynamics

Name	Description
	CRM.
ProductSalesLiterature	This is a table representing the ProductSalesLiterature entities in Dynamics CRM.
ProductSubstitute	This is a table representing the ProductSubstitute entities in Dynamics CRM.
Publisher	This is a table representing the Publisher entities in Dynamics CRM.
PublisherAddress	This is a table representing the PublisherAddress entities in Dynamics CRM.
QuarterlyFiscalCalendar	This is a table representing the QuarterlyFiscalCalendar entities in Dynamics CRM.
Queue	This is a table representing the Queue entities in Dynamics CRM.
QueueItem	This is a table representing the QueueItem entities in Dynamics CRM.
Quote	This is a table representing the Quote entities in Dynamics CRM.
QuoteClose	This is a table representing the QuoteClose entities in Dynamics CRM.
QuoteDetail	This is a table representing the

Name	Description
	QuoteDetail entities in Dynamics CRM.
RecurrenceRule	This is a table representing the RecurrenceRule entities in Dynamics CRM.
RecurringAppointmentMaster	This is a table representing the RecurringAppointmentMaster entities in Dynamics CRM.
RelationshipRole	This is a table representing the RelationshipRole entities in Dynamics CRM.
RelationshipRoleMap	This is a table representing the RelationshipRoleMap entities in Dynamics CRM.
Report	This is a table representing the Report entities in Dynamics CRM.
ReportCategory	This is a table representing the ReportCategory entities in Dynamics CRM.
ReportEntity	This is a table representing the ReportEntity entities in Dynamics CRM.
ReportLink	This is a table representing the ReportLink entities in Dynamics CRM.
ReportVisibility	This is a table representing the ReportVisibility entities in Dynamics CRM.

Name	Description
Resource	This is a table representing the Resource entities in Dynamics CRM.
ResourceGroup	This is a table representing the ResourceGroup entities in Dynamics CRM.
ResourceGroupExpansion	This is a table representing the ResourceGroupExpansion entities in Dynamics CRM.
ResourceSpec	This is a table representing the ResourceSpec entities in Dynamics CRM.
RibbonCommand	This is a table representing the RibbonCommand entities in Dynamics CRM.
RibbonContextGroup	This is a table representing the RibbonContextGroup entities in Dynamics CRM.
RibbonCustomization	This is a table representing the RibbonCustomization entities in Dynamics CRM.
RibbonDiff	This is a table representing the RibbonDiff entities in Dynamics CRM.
RibbonRule	This is a table representing the RibbonRule entities in Dynamics CRM.
RibbonTabToCommandMap	This is a table representing the RibbonTabToCommandMap entities in

Name	Description
	Dynamics CRM.
Role	This is a table representing the Role entities in Dynamics CRM.
RolePrivileges	This is a table representing the RolePrivileges entities in Dynamics CRM.
RoleTemplate	This is a table representing the RoleTemplate entities in Dynamics CRM.
RoleTemplatePrivileges	This is a table representing the RoleTemplatePrivileges entities in Dynamics CRM.
RollupField	This is a table representing the RollupField entities in Dynamics CRM.
SalesLiterature	This is a table representing the SalesLiterature entities in Dynamics CRM.
SalesLiteratureItem	This is a table representing the SalesLiteratureItem entities in Dynamics CRM.
SalesOrder	This is a table representing the SalesOrder entities in Dynamics CRM.
SalesOrderDetail	This is a table representing the SalesOrderDetail entities in Dynamics CRM.
SalesProcessInstance	This is a table representing the

Name	Description
	SalesProcessInstance entities in Dynamics CRM.
SavedQuery	This is a table representing the SavedQuery entities in Dynamics CRM.
SavedQueryVisualization	This is a table representing the SavedQueryVisualization entities in Dynamics CRM.
SdkMessage	This is a table representing the SdkMessage entities in Dynamics CRM.
SdkMessageFilter	This is a table representing the SdkMessageFilter entities in Dynamics CRM.
SdkMessagePair	This is a table representing the SdkMessagePair entities in Dynamics CRM.
SdkMessageProcessingStep	This is a table representing the SdkMessageProcessingStep entities in Dynamics CRM.
SdkMessageProcessingStepImage	This is a table representing the SdkMessageProcessingStepImage entities in Dynamics CRM.
SdkMessageProcessingStepSecureConfig	This is a table representing the SdkMessageProcessingStepSecureConfig entities in Dynamics CRM.
SdkMessageRequest	This is a table representing the

Name	Description
	SdkMessageRequest entities in Dynamics CRM.
SdkMessageRequestField	This is a table representing the SdkMessageRequestField entities in Dynamics CRM.
SdkMessageResponse	This is a table representing the SdkMessageResponse entities in Dynamics CRM.
SdkMessageResponseField	This is a table representing the SdkMessageResponseField entities in Dynamics CRM.
SemiAnnualFiscalCalendar	This is a table representing the SemiAnnualFiscalCalendar entities in Dynamics CRM.
Service	This is a table representing the Service entities in Dynamics CRM.
ServiceAppointment	This is a table representing the ServiceAppointment entities in Dynamics CRM.
ServiceContractContacts	This is a table representing the ServiceContractContacts entities in Dynamics CRM.
ServiceEndpoint	This is a table representing the ServiceEndpoint entities in Dynamics CRM.

Name	Description
SharePointDocumentLocation	This is a table representing the SharePointDocumentLocation entities in Dynamics CRM.
SharePointSite	This is a table representing the SharePointSite entities in Dynamics CRM.
Site	This is a table representing the Site entities in Dynamics CRM.
SiteMap	This is a table representing the SiteMap entities in Dynamics CRM.
Solution	This is a table representing the Solution entities in Dynamics CRM.
SolutionComponent	This is a table representing the SolutionComponent entities in Dynamics CRM.
StatusMap	This is a table representing the StatusMap entities in Dynamics CRM.
StringMap	This is a table representing the StringMap entities in Dynamics CRM.
Subject	This is a table representing the Subject entities in Dynamics CRM.
Subscription	This is a table representing the Subscription entities in Dynamics CRM.
SubscriptionClients	This is a table representing the SubscriptionClients entities in Dynamics

Name	Description
	CRM.
SubscriptionManuallyTrackedObject	This is a table representing the SubscriptionManuallyTrackedObject entities in Dynamics CRM.
SubscriptionSyncInfo	This is a table representing the SubscriptionSyncInfo entities in Dynamics CRM.
SubscriptionTrackingDeletedObject	This is a table representing the SubscriptionTrackingDeletedObject entities in Dynamics CRM.
SystemForm	This is a table representing the SystemForm entities in Dynamics CRM.
SystemUser	This is a table representing the SystemUser entities in Dynamics CRM.
SystemUserBusinessUnitEntityMap	This is a table representing the SystemUserBusinessUnitEntityMap entities in Dynamics CRM.
SystemUserLicenses	This is a table representing the SystemUserLicenses entities in Dynamics CRM.
SystemUserPrincipals	This is a table representing the SystemUserPrincipals entities in Dynamics CRM.
SystemUserProfiles	This is a table representing the SystemUserProfiles entities in Dynamics

Name	Description
	CRM.
SystemUserRoles	This is a table representing the SystemUserRoles entities in Dynamics CRM.
Task	This is a table representing the Task entities in Dynamics CRM.
Team	This is a table representing the Team entities in Dynamics CRM.
TeamMembership	This is a table representing the TeamMembership entities in Dynamics CRM.
TeamProfiles	This is a table representing the TeamProfiles entities in Dynamics CRM.
TeamRoles	This is a table representing the TeamRoles entities in Dynamics CRM.
Template	This is a table representing the Template entities in Dynamics CRM.
Territory	This is a table representing the Territory entities in Dynamics CRM.
TimeZoneDefinition	This is a table representing the TimeZoneDefinition entities in Dynamics CRM.
TimeZoneLocalizedName	This is a table representing the TimeZoneLocalizedName entities in

Name	Description
	Dynamics CRM.
TimeZoneRule	This is a table representing the TimeZoneRule entities in Dynamics CRM.
TransactionCurrency	This is a table representing the TransactionCurrency entities in Dynamics CRM.
TransformationMapping	This is a table representing the TransformationMapping entities in Dynamics CRM.
TransformationParameterMapping	This is a table representing the TransformationParameterMapping entities in Dynamics CRM.
UnresolvedAddress	This is a table representing the UnresolvedAddress entities in Dynamics CRM.
UoM	This is a table representing the UoM entities in Dynamics CRM.
UoMSchedule	This is a table representing the UoMSchedule entities in Dynamics CRM.
UserEntityInstanceData	This is a table representing the UserEntityInstanceData entities in Dynamics CRM.
UserEntityUISettings	This is a table representing the UserEntityUISettings entities in Dynamics CRM.

Name	Description
UserFiscalCalendar	This is a table representing the UserFiscalCalendar entities in Dynamics CRM.
UserForm	This is a table representing the UserForm entities in Dynamics CRM.
UserQuery	This is a table representing the UserQuery entities in Dynamics CRM.
UserQueryVisualization	This is a table representing the UserQueryVisualization entities in Dynamics CRM.
UserSettings	This is a table representing the UserSettings entities in Dynamics CRM.
WebResource	This is a table representing the WebResource entities in Dynamics CRM.
WebWizard	This is a table representing the WebWizard entities in Dynamics CRM.
WizardAccessPrivilege	This is a table representing the WizardAccessPrivilege entities in Dynamics CRM.
WizardPage	This is a table representing the WizardPage entities in Dynamics CRM.
Workflow	This is a table representing the Workflow entities in Dynamics CRM.
WorkflowDependency	This is a table representing the

Name	Description
	WorkflowDependency entities in Dynamics CRM.
WorkflowLog	This is a table representing the WorkflowLog entities in Dynamics CRM.
WorkflowWaitSubscription	This is a table representing the WorkflowWaitSubscription entities in Dynamics CRM.

# MySQL Live

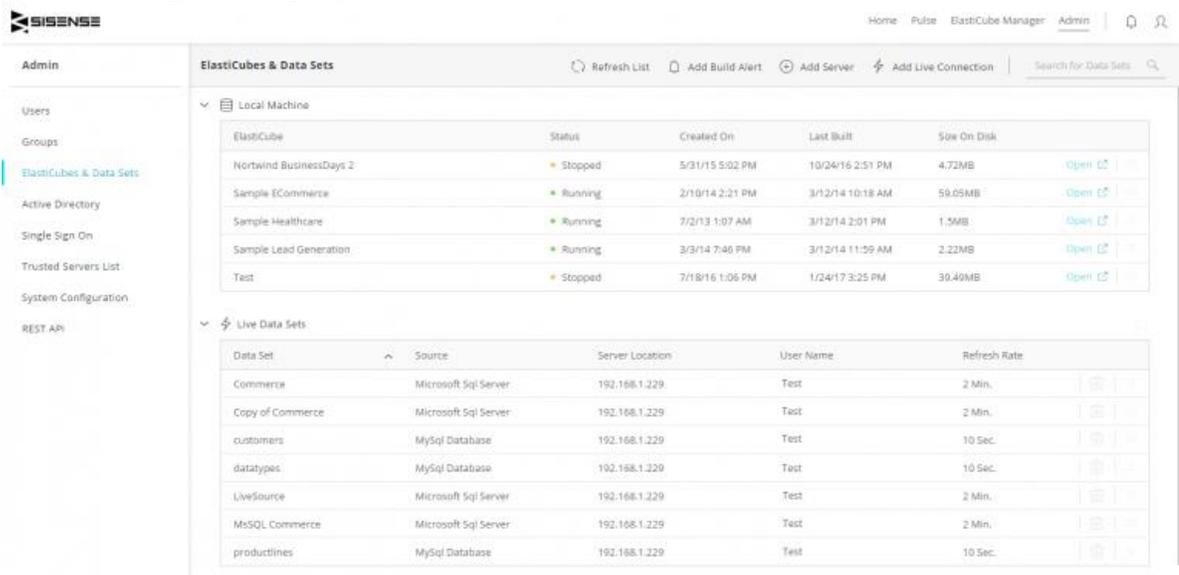
This topic describes how you can add live connections to your MySQL databases from the Sisense Web Application.

For information about live connections, see [Sisense Live Connect](#).

## Adding Live Connections

**To add a live connection:**

1. In the Sisense Web Application, select **Admin > Data Sources**. The Data Sources page is displayed.



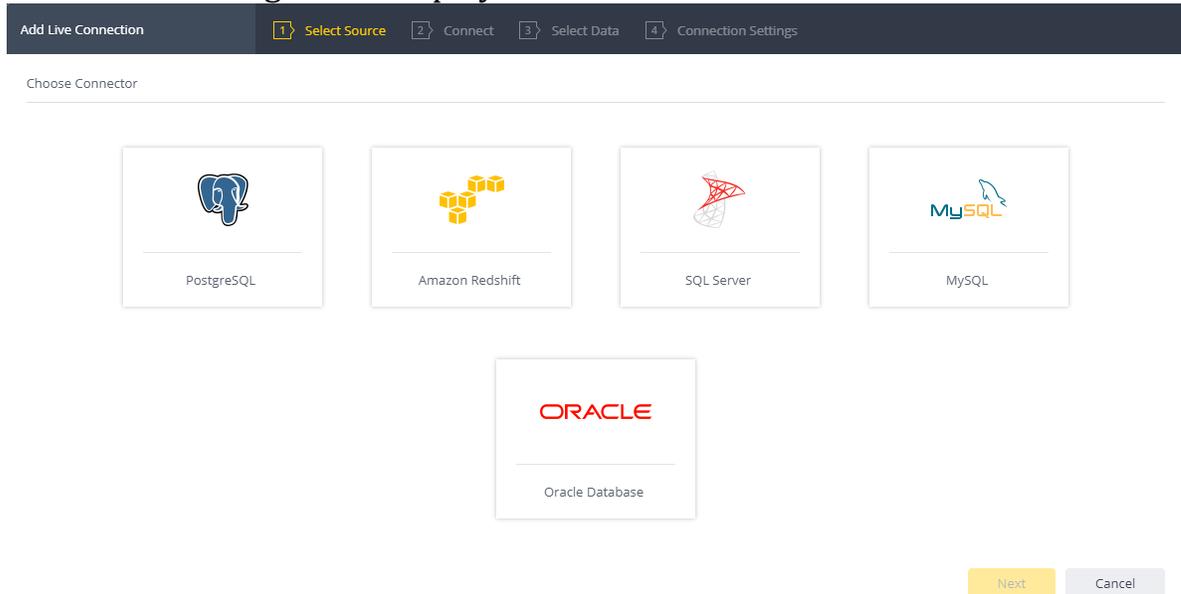
The screenshot shows the Sisense Admin console interface. The left sidebar contains navigation options: Admin, Users, Groups, ElastiCubes & Data Sets (highlighted), Active Directory, Single Sign On, Trusted Servers List, System Configuration, and REST API. The main content area is titled 'ElastiCubes & Data Sets' and includes a search bar and several action buttons: Refresh List, Add Build Alert, Add Server, and Add Live Connection. The page is divided into two sections: 'Local Machine' and 'Live Data Sets'.

ElastiCube	Status	Created On	Last Built	Size On Disk	
Nortwind BusinessDays 2	Stopped	5/31/15 5:02 PM	10/24/16 2:51 PM	4.72MB	Open
Sample eCommerce	Running	2/10/14 2:21 PM	3/12/14 10:18 AM	59.05MB	Open
Sample Healthcare	Running	7/2/13 1:07 AM	3/12/14 2:01 PM	1.5MB	Open
Sample Lead Generation	Running	3/3/14 7:46 PM	3/12/14 11:59 AM	2.23MB	Open
Test	Stopped	7/18/16 1:06 PM	1/24/17 3:25 PM	30.49MB	Open

Data Set	Source	Server Location	User Name	Refresh Rate	
Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	
Copy of Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	
customers	MySql Database	192.168.1.229	Test	10 Sec.	
datatypes	MySql Database	192.168.1.229	Test	10 Sec.	
LiveSource	Microsoft Sql Server	192.168.1.229	Test	2 Min.	
MsSQL Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	
productlines	MySql Database	192.168.1.229	Test	10 Sec.	

2. In the Data Sources page, click **Add Live Connection**. The Add Live Connection dialog box is displayed.



3. In the Add Live Connection dialog box, select your live data source.
4. Under the Connect tab, enter the following credentials:
  - Location:** The IP address of your data source.
  - User Name:** The user name to access data source.
  - Password:** The password to access your data source.
5. Click **Next**.
6. Under the Select Data tab, on the left side is a list of schemas located in your data source. Select the relevant schema. You can find the schema easily by searching for it through the search field at the top of the list to filter the schemas displayed. After selecting the schema, a list of the available tables and views are displayed.
7. Select a table or view. You can find the table easily by searching for it through the search field at the top of the list to filter the tables displayed.
8. Click **Next**.
9. Under the Connection Settings tab, set the following:
  - Connection Name:** Enter a name for your connection. This field defines the name of your data source displayed in the Sisense Web Application.
  - Refresh Rate:** Enter or select the refresh rate for your data source. The shortest refresh rate you can set is 10 seconds. The refresh rate is the time in between queries to the data source. If the previous query has not been completed, the next query is not sent until the previous refresh is complete unless the refresh rate is greater than the time it took to complete the

refresh since the last query. For example, if your refresh rate is 10 seconds, and your query takes 15 seconds, the next query is sent when the previous query is complete as 10 seconds has passed and the previous query is complete. If your refresh rate is 30 seconds, and the previous query is complete, the next refresh will be in 30 seconds from the last query.

**Time Out:** Enter or select the Time Out rate for your data source. The Time Out rate determines how long Sisense waits before cancelling a query that does not receive a response from your data source.

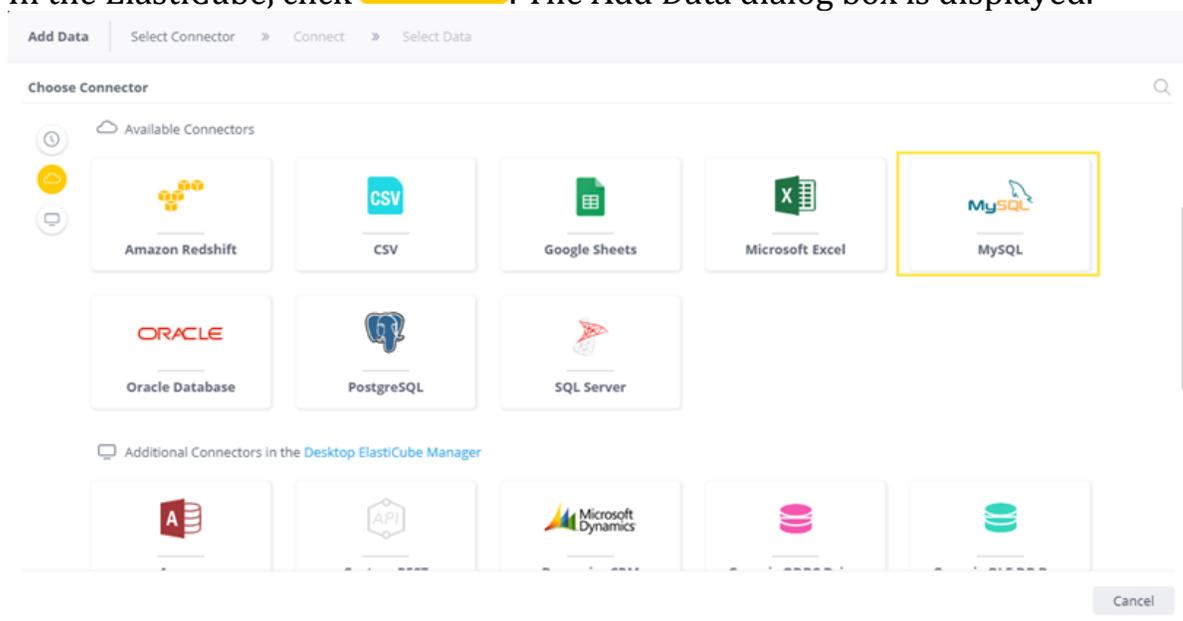
10. Click **Done**. The connection is added as a data set to your Data Sources page.

# Connecting to MySQL

The ElastiCube Manager enables easy and quick access to databases, tables and views contained with MySQL databases.

## To import data from your MySQL database:

1. In the Data page of ElastiCube Manager, open an ElastiCube or click create a new ElastiCube.
2. In the ElastiCube, click **+ Data**. The Add Data dialog box is displayed.



3. Click **MySQL**. The MySQL Connect area is displayed.

Add Data | Select Connector » Connect » Select Data

MySQL

Location  
127.0.0.1

User name

Password

Use SSL

Private Key

Certificate

Authority

Next Cancel

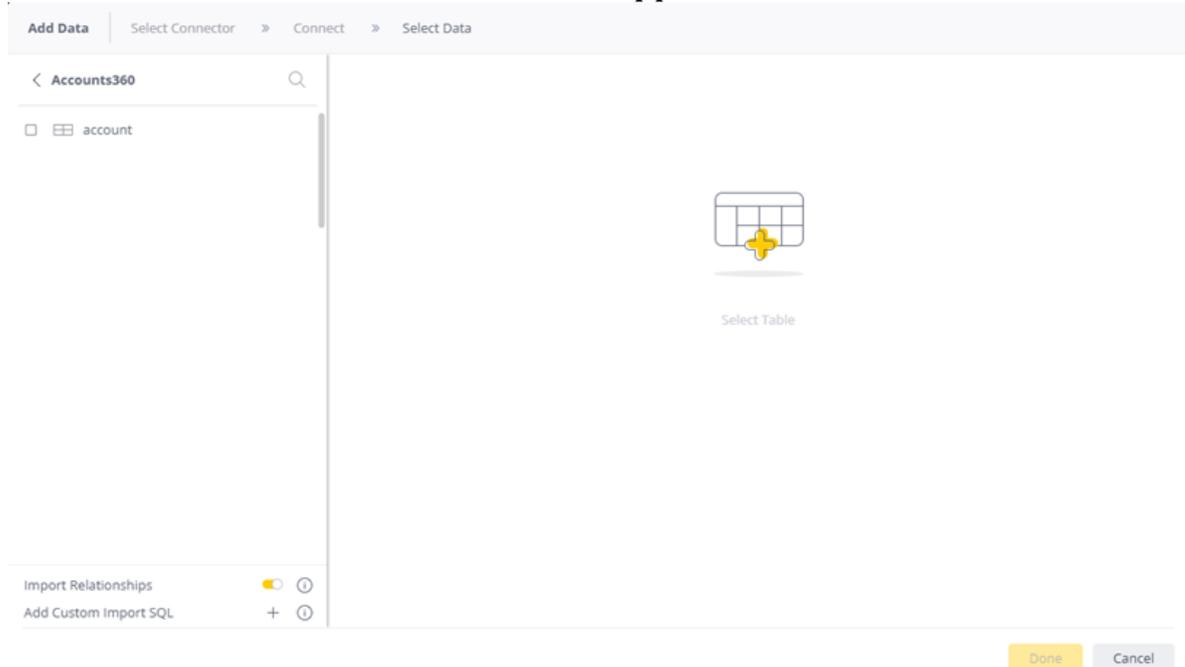
4. Enter the following details:

- **Location:** Enter the computer/server IP address of the database. To connect to a database running on your own computer enter **localhost**.
- If Windows Authentication is configured with the database in the **User Name** and **Password** fields, enter your database credentials.

5. **(Optional)** Select Use SSL if you are connecting to an SSL server. Enter the relevant information to configure verification between the client and the MySQL server over SSL:

- In the **Private Key** field, enter the path of the location of your file containing the your private key.
- In the **Certificate** field, enter the path of the location of your PEM file containing the client's certificate.
- In the **Authority** field, enter the path of the location of your PEM file containing the trusted SSL certificate authority.

6. Click **Next**. A list of tables in the database are displayed. All tables and views associated with the database will appear in a new window.

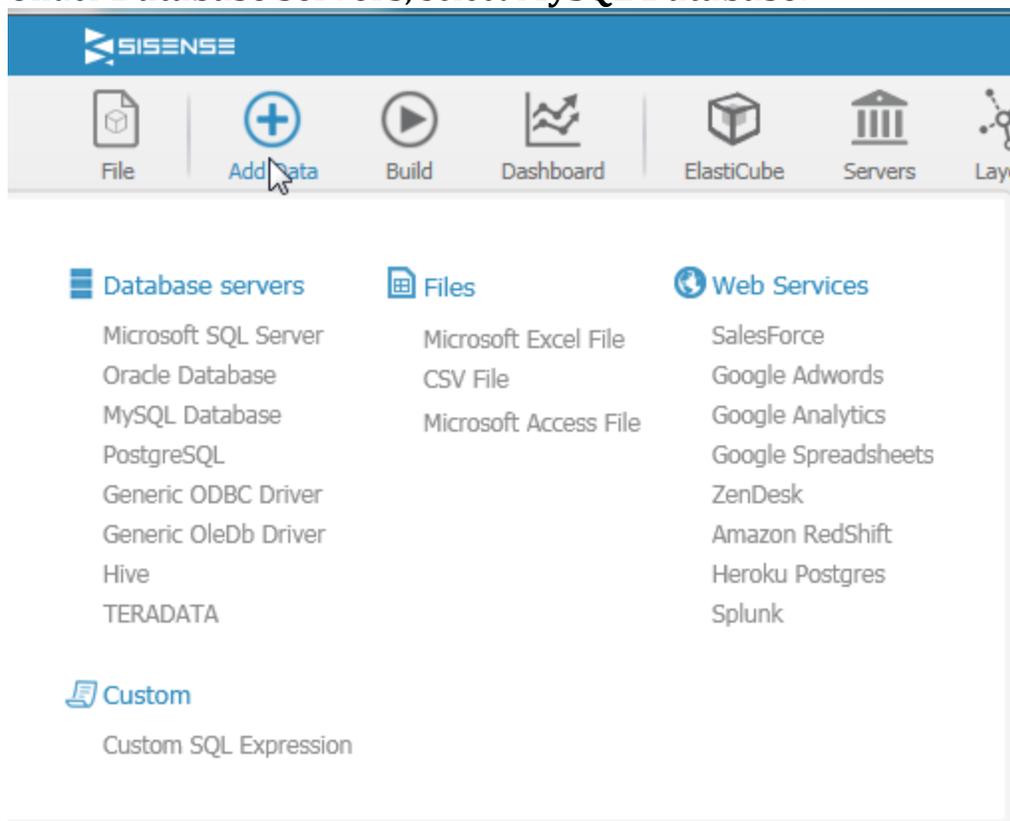


7. From the Tables list, select the relevant table or view you want to work with. You can click  next to the relevant table or view to see a preview of the data inside it. When you select the table or view, two new options are displayed at the bottom of the list, **Import Relationships** and **Add Custom Import SQL**.
8. (Optional) By default, existing relationships between tables are automatically replicated in the ElastiCube. You can disable this by toggling the **Import Relationships** switch.
9. (Optional) Click + to customize the data you want to import with SQL. See [Importing Data with Custom Queries](#) for more information.
10. After you have selected all the relevant tables, click **Done**. The tables are added to your schema in the ElastiCube Manager.

# Connecting to MySQL

The ElastiCube Manager enables easy and quick access to databases, tables and views contained with MySQL databases. The steps below detail how to connect to this type of data source.

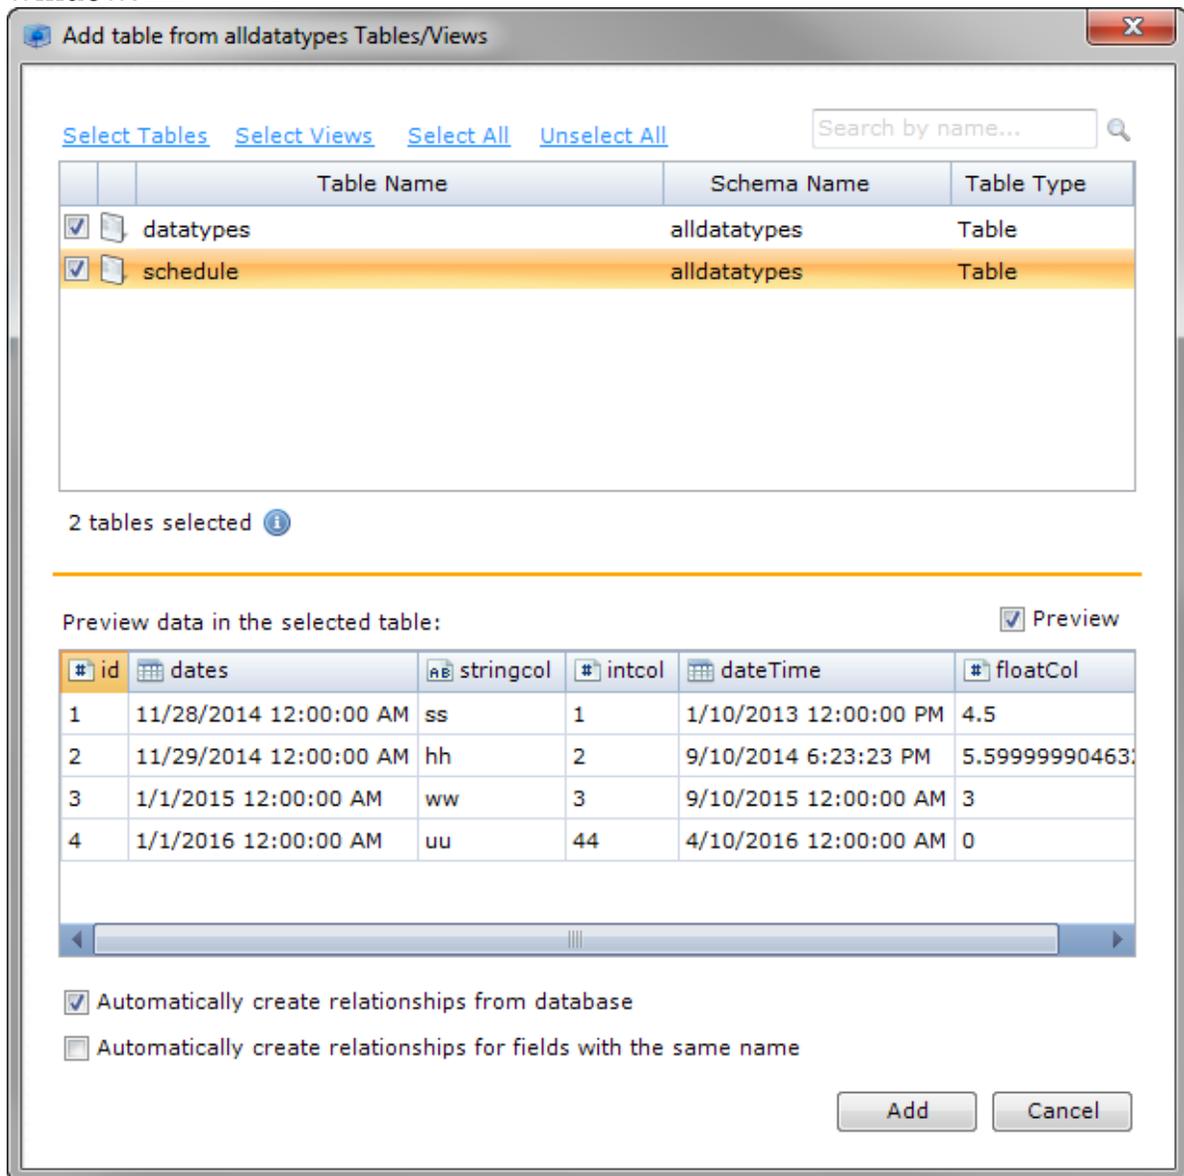
1. Click **Add data** on the top menu of the ElastiCube Manager.
2. Under **Database servers**, select **MySQL Database**.



The Connect to MySQL window is displayed.

3. **Database server location:**  
Enter the computer/server IP address of the database. To connect to a database running on your own computer enter **localhost**.
4. Select either **Windows Authentication** if configured with the database or alternatively, **Use the Following User Name & Password**, and enter the database credentials.
5. (Optional) Select **Use SSL** to if you are connecting to an SSL server. Enter the relevant information to configure verification between the client and the MySQL server over SSL:
  - In the **Private Key** field, enter the path or click  and navigate to the location of your file containing the your private key.
  - In the **Certificate** field, enter the path or click  and navigate to the location of your PEM file containing the client's certificate.
  - In the **Authority** field, enter the path or click  and navigate to the location of your PEM file containing the trusted SSL certificate authority.

6. Alternatively, you can use Sisense's ODBC connectors to connect to secure data sources. For more information, click [here](#).
7. Click **Connect to server**.
8. From the **Select Database** list, select the relevant database you want to work with, and click OK
9. All tables and views associated with the database will appear in a new window.



10. To view a preview of data contained in a particular table, highlight the table or view in the list and in click the preview pane below. To preview the table, select the **Preview** checkbox. Enable the checkbox next to each table or view you would like to use.

11. Existing relationships between tables can be automatically replicated in the ElastiCube by selecting the **Automatically create relationships from database** option. Likewise fields with similar names can be linked by selecting the **Automatically create relationships for fields with the same** name option.
12. If you want to customize the data before importing it into the ElastiCube, you can run a custom SQL query to manipulate the data. This can be useful, for example, when you want to import only a portion of the data, rather than all of the data. [Click here](#) to learn more.
13. Once all relevant tables are selected, click **Add**.

# Connecting via ODBC Drivers

## Open Database Connectivity

The ODBC provider allows access to a data source via an installed driver on the operating system regardless of the architecture of the data source.

In Sisense, it serves as a connector to data providers, such as MS Sql, MySql, Oracle, Hive and so forth.

## DSN vs. Connection-String (DSN-Less)

A connection to the provider is configured by a connection-string. Microsoft windows allows you to store the connection string properties in its registry for further use. It supplies an ODBC configuration wizard and requests an identifying Data Source Name (DSN).

When you add an ODBC table in the ElastiCube Manager, the connection configuration dialog will provide you with two options:

- DSN – choosing a data source name.
- Connection String (DSN-Less) – typing the actual connection string

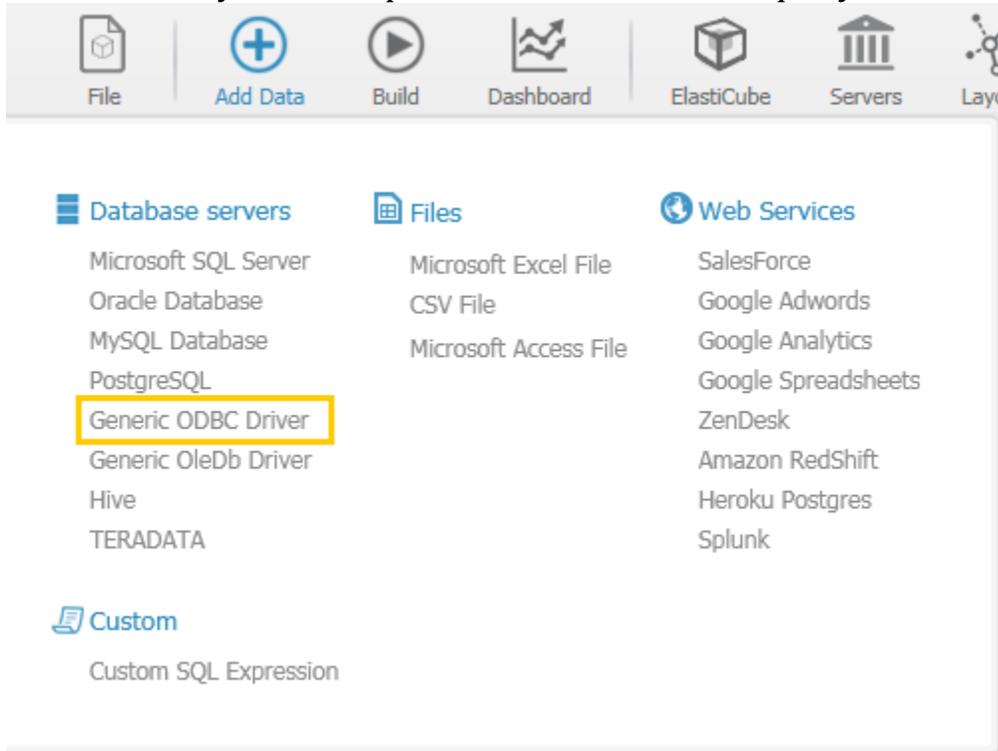
## Adding Tables to the ElastiCube Manager:

1. Click **Add Data** in the top menu of the ElastiCube Manager.
2. Under the **Database servers** category, select **Generic ODBC Driver**.
3. From the ODBC Connectivity Properties dialog choose **DSN** or Connection String (DSN-Less).
  - **DSN:** Choose one of the existing DSNs from the drop-down list (each item is represented by its name), or click **Add DSN...** to open the **Create New Data Source** window. In this window, you'll be asked to select one of the installed ODBC drivers, select a name for the DSN, and configure your

connection using the wizard steps. At the end, the newly added DSN will be added to the drop-down list as the selected item.

- **Connection String:** Type the connection string in the input box. For known connection strings, click the Help link to navigate to Connection Strings (External site from Sisense). Click **Test Connection** to test the connection. Click **OK** to move to the final dialog. From the tables selection dialog, select tables that you want to add and click **Add** to complete the wizard.

**Note:** You may need to update the executed table query.



## Setting the Table Query (Troubleshoot)

ODBC and OLE DB are generic drivers; therefore the SQL syntax being executed against the actual data source may vary from one provider to another.

Sisense sets by default the standard SQL for the ODBC/OLE DB tables which covers most of today's known data providers. For cases where non-standard SQL must be supplied, the table addition dialog provides a quick preview of the current SQL query and an edit button that opens the **Query Properties** dialog.

In **Query Properties**, you can set the following options:

- Auto Query Structure: Choosing the proper SQL syntax delimiter.
- Manual Query Input: Type the desired query.

## Customizing Your ODBC Connection

Though there hundreds of data sources available, the generic ODBC driver can help you connect to just about all of them. However, if the default settings do not fit your specific use-case, Sisense lets you customize your ODBC connection to give you the functionality you need.

When you add an ODBC connection to your Sisense configuration, Sisense saves the details of your connection in a file called `odbcConfig.json`. The `odbcConfig.json` is a configuration file written in JSON syntax located in the directory: `...\Program`

`Files\Sisense\DataConnectors\DotNetContainer\Connectors\GenericODBC` You can edit this file to add functionality such as support for accumulative builds, increase query timeouts, or modify your query patterns.

The `odbcConfig.json` file contains the following objects that you can edit to modify your connectors behavior:

Name	Type	Mandatory	Description
Provider	String	Yes	The name of your data source. This name must be unique.
DriverNames	String	Yes	An array of ODBC driver names. If you don't know a driver name, you could check "Connectors.log".
IsAccumulativeSupported	Boolean	No	Indicates whether accumulative build is supported. Default value:

Name	Type	Mandatory	Description
			false.
IsRelationsSupported	Boolean	No	Enable/disables relations functionality in the ElastiCube Manager.Default value: false.
QueryTimeout	Integer	No	Query timeout value. You can increase this value is your datasource takes longer to establish a connection.Default value: 0.
QuerySettings	QueryConfig*	No	Datasource specific query settings (see QueryConfig Object table below).Default value: default query config values.
Extension	String	No	A name of the assembly with the appropriate extension.Default value: none.

\* The properties for the QueryConfig object are described in the table below:

## QueryConfig Object

Name	Type	Mandatory	Description
ColumnNamePattern	String	No	Pattern which represents column name format (should include appropriate separators).To create a pattern, you can use the following keywords: @schema: Schema name (When schema is empty, table name will be using instead). @table: Table name. @column: Column name..Example: ` @schema ` . ` @column ` .Default value: [@column]
TableNamePattern	String	No	Pattern which represents table name format (should include appropriate separators).To create a pattern, you can use the following keywords: @schema: Schema name (When schema is empty, table name will be using instead). @table: Table name.Example: "@schema"."@table"Default value: [@schema].[@table] or [@table] in the case when schema name is missing.
PreviewSqlPattern	String	No	Pattern which represents preview SQL format (should include appropriate SQL operators).To create a pattern, you can use the following keywords: @rowcount: The number of rows to show. @sql:

Name	Type	Mandatory	Description
			Inner select.Example: SELECT * FROM (@sql) LIMIT @rowcountDefault value: SELECT TOP @rowcount x.* FROM (@sql) x
DateTimeFormat	String	No	Format used to convert DateTime value to string for an accumulative workflow.More info about standard formatshere.More info about defining a custom formathere..Default value: sortable format specifier.

# Connecting via OLE DB Drivers

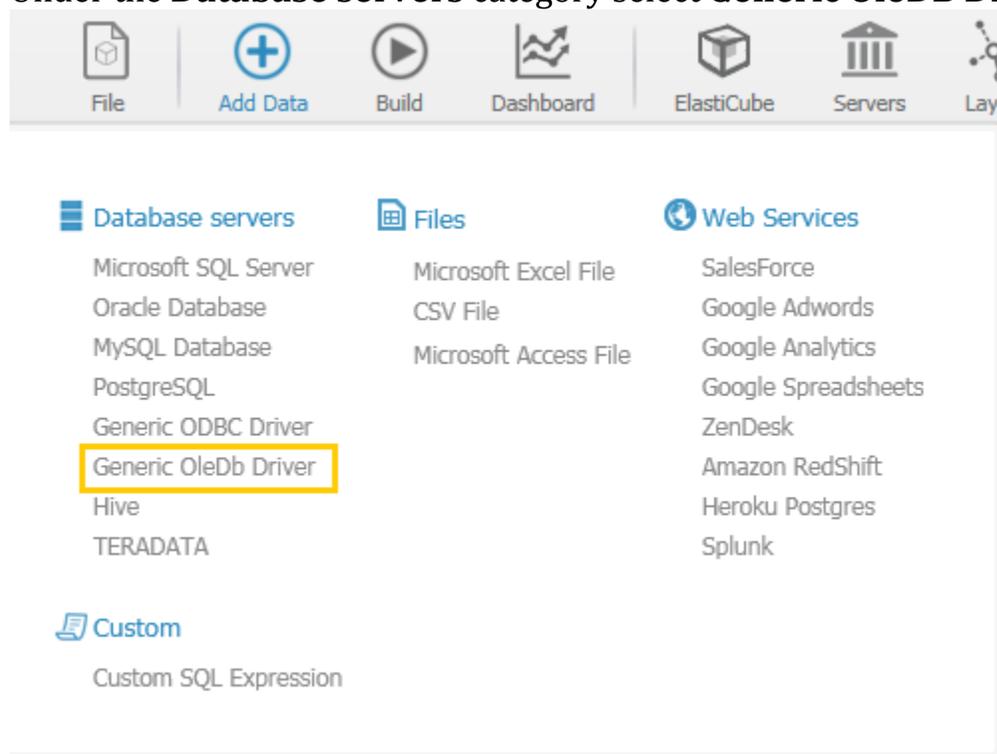
## Object Linking and Embedding Database

The OLE DB provides access to a data source via an installed driver on the operating system regardless of the architecture of the data source.

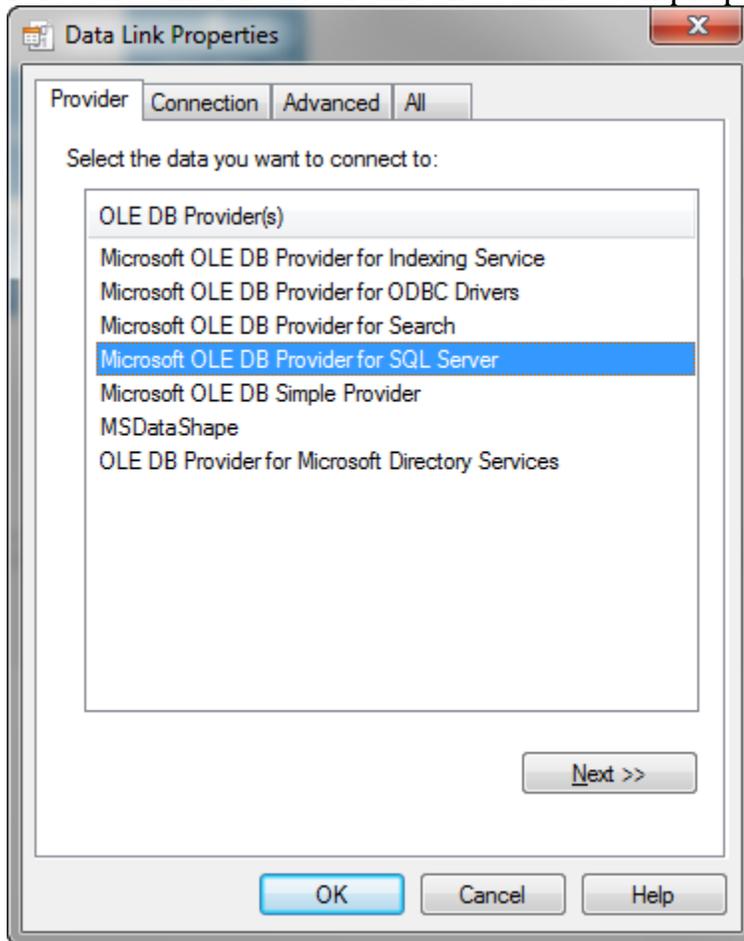
In Sisense, it serves as a connector to data providers that are not included in the common providers list, such as MS SQL, MySQL, Oracle and so forth.

## Adding a table to the ElastiCube Manager

1. Click **Add Data** on the top menu of the ElastiCube Manager.
2. Under the **Database servers** category select **Generic OleDB Driver**.

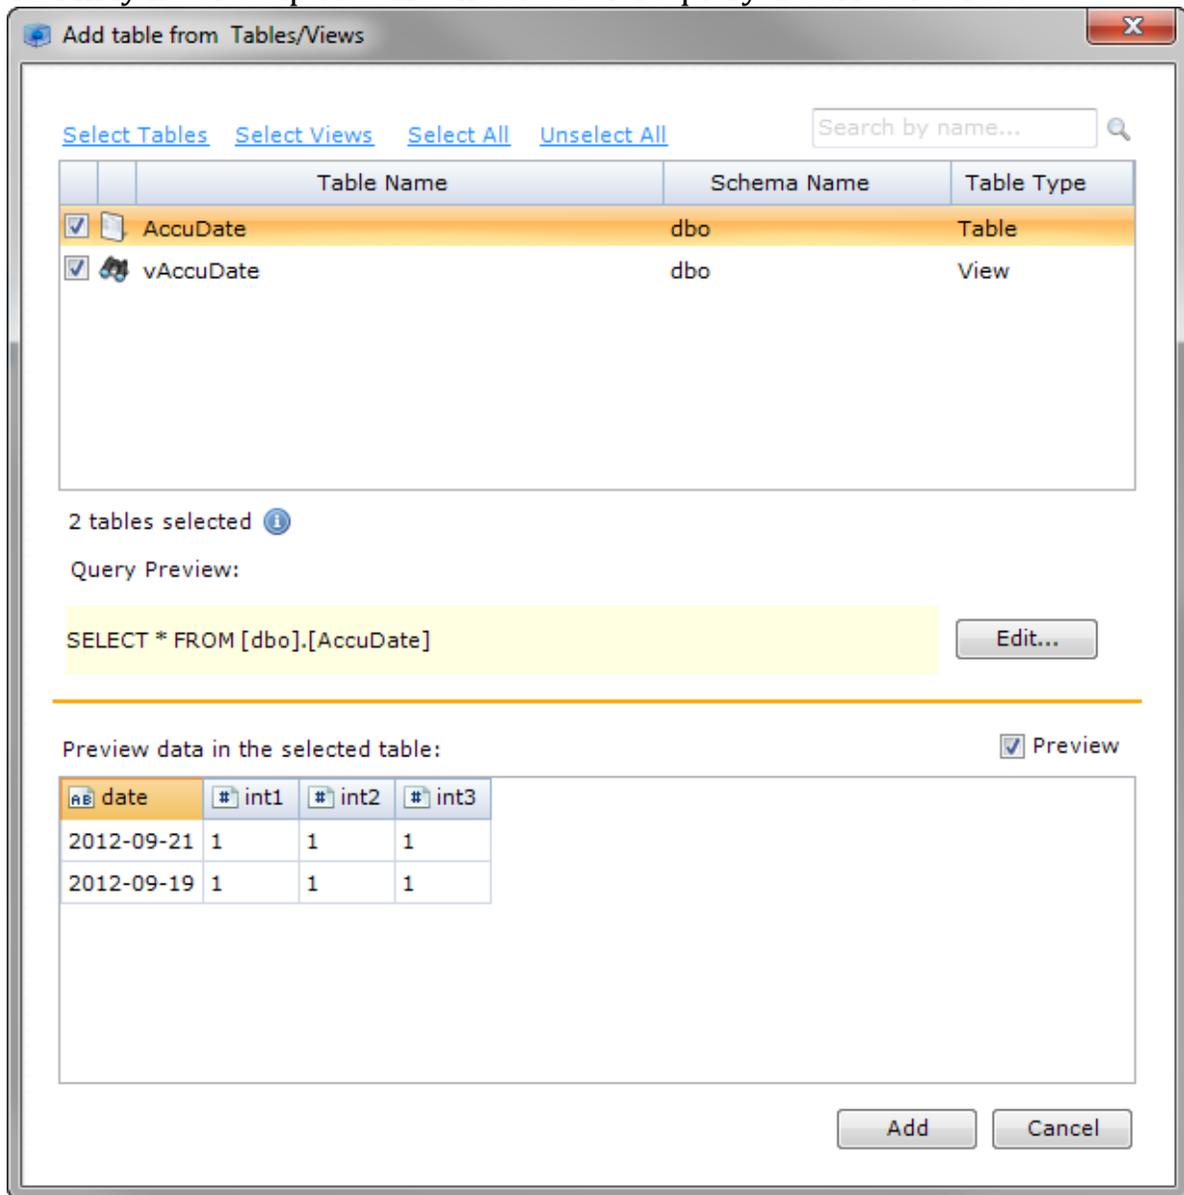


3. In the **Data Link Properties** dialog, select a connection by choosing the relevant OLE DB driver and the connection properties.



4. **Note:** When the data source server requires a login credentials, you must select the Allow saving password checkbox. Follow all wizard steps and click **OK** to move to the final dialog. From the tables selection dialog, select tables that you want to add and click **Add** to complete the wizard. **Note:**

You may need to update the executed table query described below.



## Setting the table query (Troubleshoot)

ODBC and OLE DB are generic drivers; therefore the SQL syntax being executed against the actual data source may vary from one provider to the other.

Sisense sets by default the standard SQL for the ODBC/OLE DB tables which covers most of the known data providers. For cases where non-standard SQL

must be supplied, the table addition dialog provides a quick preview of the current SQL query and an edit button which opens the Query Properties dialog.

In the **Query Properties** dialog, two options are available:

- **Auto Query Structure:** Choose the proper SQL syntax delimiter
- **Manual Query Input:** Type the desired query.

**Note:** The automatic options can be applied on the current table or on all tables, while the manual option defines a single table and can only be applied on the current table. The **Apply All** and **Apply Table** buttons will be enabled according to the selection made.

## Misc Troubleshooting:

- When an ODBC table is added by a DSN connection, the DSN must exist on the target ElastiCube Server machine, or else the server will fail to connect and import data from that table. In general, setting a connection by a DSN (not by connection string) is not recommended practice.
- Regarding ODBC only, ensure that the addressed ODBC driver platform (32 or 64 bit) matches the current installation of Sisense (Manager and Server). Platform mismatches will lead to connectivity errors.
- When setting an OLE DB connection, and the data source server requires login credentials, the Allow saving password checkbox must be selected.

# Oracle Live

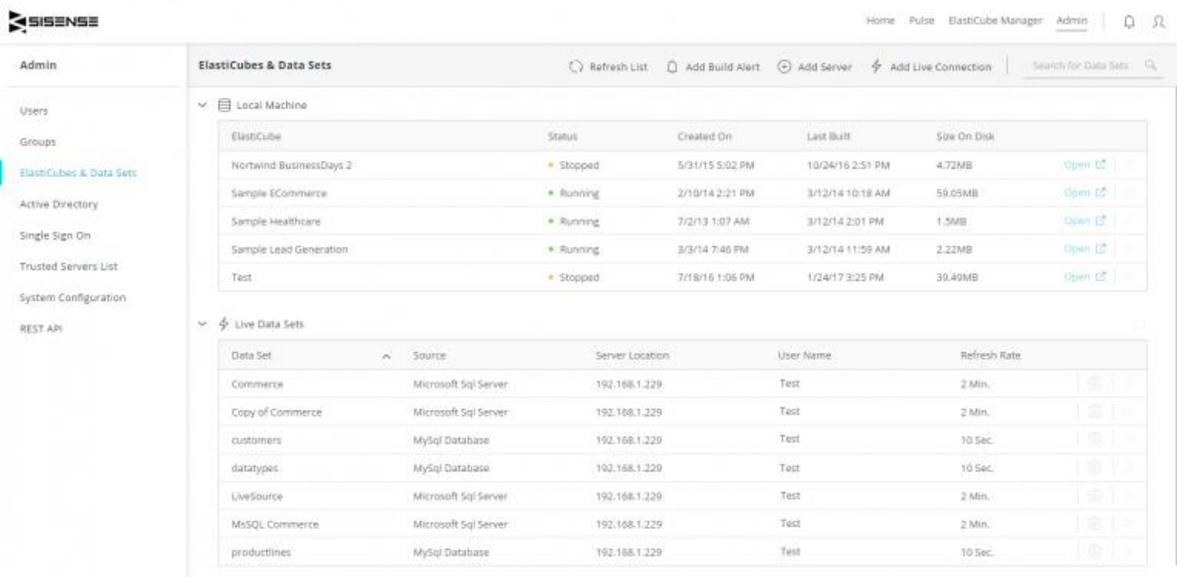
This topic describes how you can add live connections to your Oracle databases from the Sisense Web Application.

For information about live connections, see [Sisense Live Connect](#).

## Adding Oracle Live Connections

**To add a live connection:**

1. In the Sisense Web Application, select **Admin > Data Sources**. The Data Sources page is displayed.



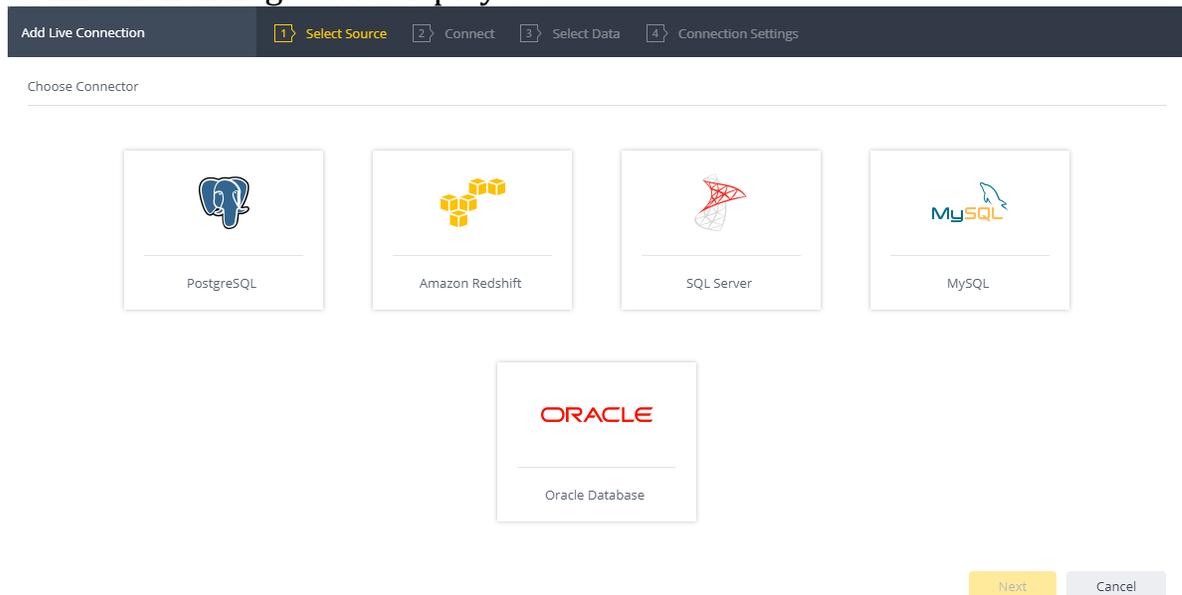
The screenshot shows the Sisense Admin console interface. The left sidebar contains navigation options: Admin, Users, Groups, ElastiCubes & Data Sets (highlighted), Active Directory, Single Sign On, Trusted Servers List, System Configuration, and REST API. The main content area is titled 'ElastiCubes & Data Sets' and includes a search bar and action buttons: Refresh List, Add Build Alert, Add Server, and Add Live Connection. The data is organized into two sections: 'Local Machine' and 'Live Data Sets'.

ElastiCube	Status	Created On	Last Built	Size On Disk	
Nortwind BusinessDays 2	Stopped	5/31/15 5:02 PM	10/24/16 2:51 PM	4.72MB	Open
Sample eCommerce	Running	2/10/14 2:21 PM	3/12/14 10:18 AM	59.05MB	Open
Sample Healthcare	Running	7/2/13 1:07 AM	3/12/14 2:01 PM	1.5MB	Open
Sample Lead Generation	Running	3/3/14 7:46 PM	3/12/14 11:59 AM	2.23MB	Open
Test	Stopped	7/18/16 1:06 PM	1/24/17 3:25 PM	30.49MB	Open

Data Set	Source	Server Location	User Name	Refresh Rate	
Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	Open
Copy of Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	Open
customers	MySql Database	192.168.1.229	Test	10 Sec.	Open
datatypes	MySql Database	192.168.1.229	Test	10 Sec.	Open
LiveSource	Microsoft Sql Server	192.168.1.229	Test	2 Min.	Open
MsSQL Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	Open
productlines	MySql Database	192.168.1.229	Test	10 Sec.	Open

2. In the Data Sources page, click **Add Live Connection**. The Add Live Connection dialog box is displayed.



3. In the Add Live Connection dialog box, select your live data source.
4. Under the Connect tab, enter the following credentials:
  - Location:** The IP address of your data source.
  - User Name:** The user name to access data source.
  - Password:** The password to access your data source.
5. (Optional) You can also use a direct connection to connect. Select **Use Direct Connection**. You will need to enter the **Service ID** and **Port** number to connect.

**Service ID:** Enter your Service ID. This can be found in the TNSNames.ora file. You can find a full description in the following topic: [Service ID](#).

Note that in the file for each tnsname, you can find the service ID and service name.

6. Click **Next**.
7. Under the Select Data tab, on the left side is a list of schemas located in your data source. Select the relevant schema. You can find the schema easily by searching for it through the search field at the top of the list to filter the schemas displayed. After selecting the schema, a list of the available tables and views are displayed.

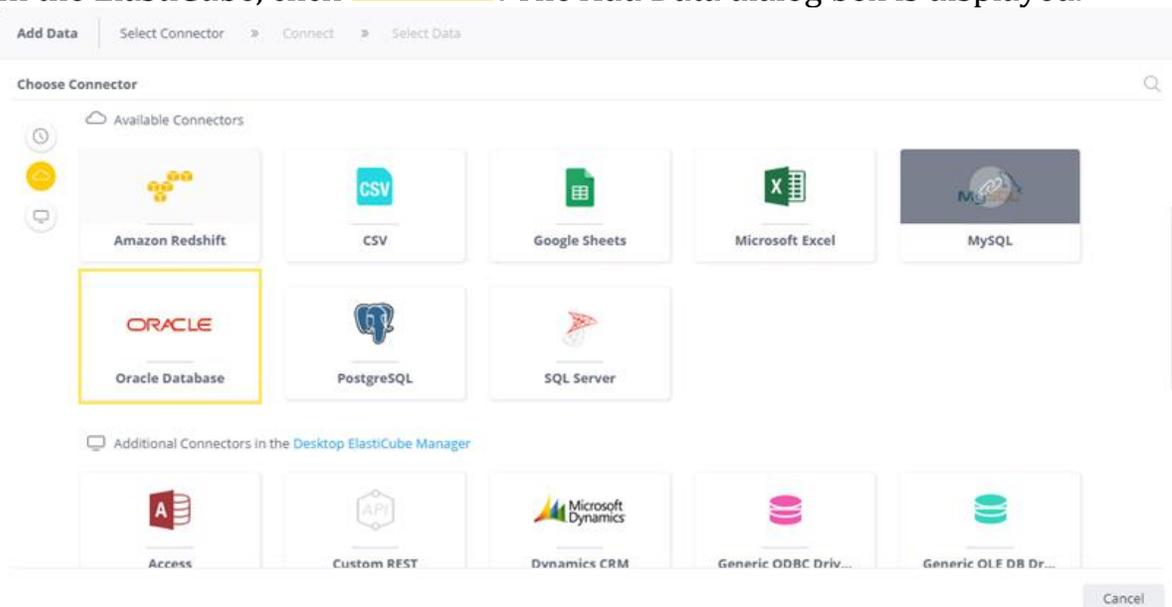
8. Select a table or view. You can find the table easily by searching for it through the search field at the top of the list to filter the tables displayed.
9. Click **Next**.
10. Under the Connection Settings tab, set the following:
  - Connection Name:** Enter a name for your connection. This field defines the name of your data source displayed in the Sisense Web Application.
  - Refresh Rate:** Enter or select the refresh rate for your data source. The shortest refresh rate you can set is 10 seconds. The refresh rate is the time in between queries to the data source. If the previous query has not been completed, the next query is not sent until the previous refresh is complete unless the refresh rate is greater than the time it took to complete the refresh since the last query. For example, if your refresh rate is 10 seconds, and your query takes 15 seconds, the next query is sent when the previous query is complete as 10 seconds has passed and the previous query is complete. If your refresh rate is 30 seconds, and the previous query is complete, the next refresh will be in 30 seconds from the last query.
  - Time Out:** Enter or select the Time Out rate for your data source. The Time Out rate determines how long Sisense waits before cancelling a query that does not receive a response from your data source.
11. Click **Done**. The connection is added as a data set to your Data Sources page.

# Connecting to Oracle

The ElastiCube Manager enables easy and quick access to databases, tables and views contained within Oracle databases. The steps below describe how to connect to this type of data source.

## To import data from Oracle:

1. In the Data page of ElastiCube Manager, open an ElastiCube or create a new ElastiCube.
2. In the ElastiCube, click . The Add Data dialog box is displayed.



### 3. Click **Oracle Database** to open the Oracle settings.

Add Data | Select Connector » Connect » Select Data

ORACLE

Oracle Database

Location: 127.0.0.1

Service ID: ORCL

Port: 1521

User name:

Password:

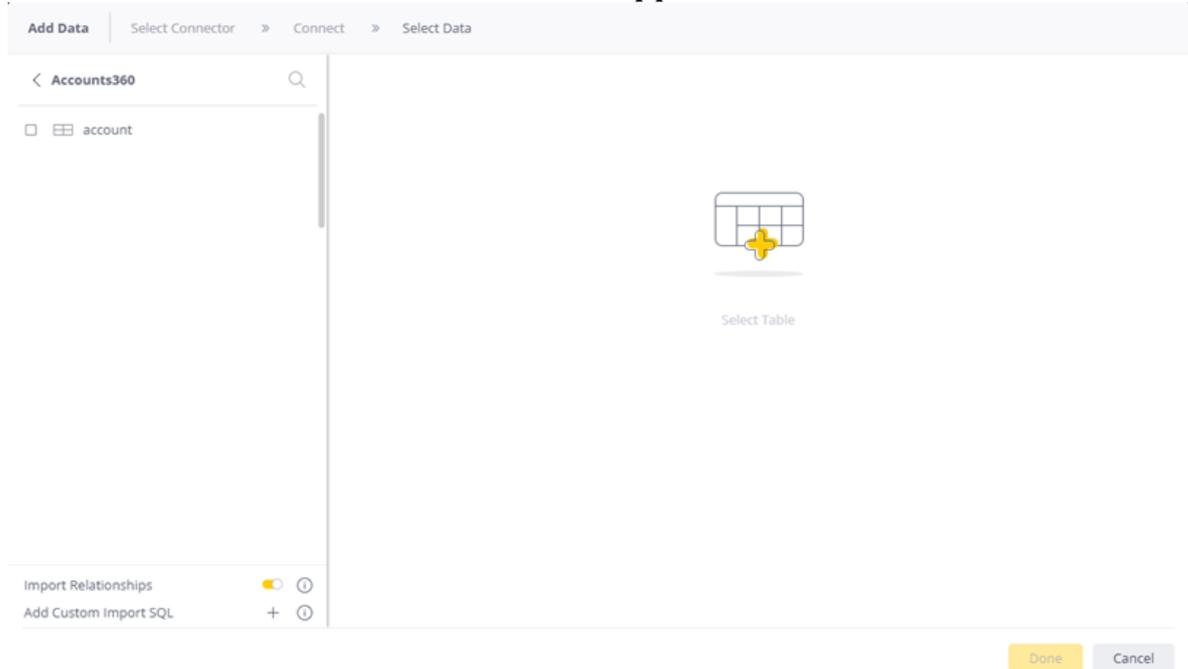
Use Direct Connection

Next Cancel

### 4. Enter the following details:

- **Location:** Enter the computer/server IP address of the database. To connect to a database running on your own computer enter **localhost**.
- If Windows Authentication is configured with the database in the **User Name** and **Password** fields, enter your database credentials.
- **Use Direct Connection:** You can also use a direct connection to connect. Select **Use Direct Connection**. You will need to enter the **Service ID** and **Port** number to connect.
- **Service ID:** Enter your Service ID. This can be found in the TNSNames.ora file. You can find a full description in the following topic: [Service ID](#).
- Note that in the file for each tnsname, you can find the service ID and service name.
- **User name and Password:** Either use your Windows login details if they are configured with the database, or alternatively enter the username and password to connect to the database.

5. Click **Next**. A list of tables in the database are displayed. All tables and views associated with the database will appear in a new window.

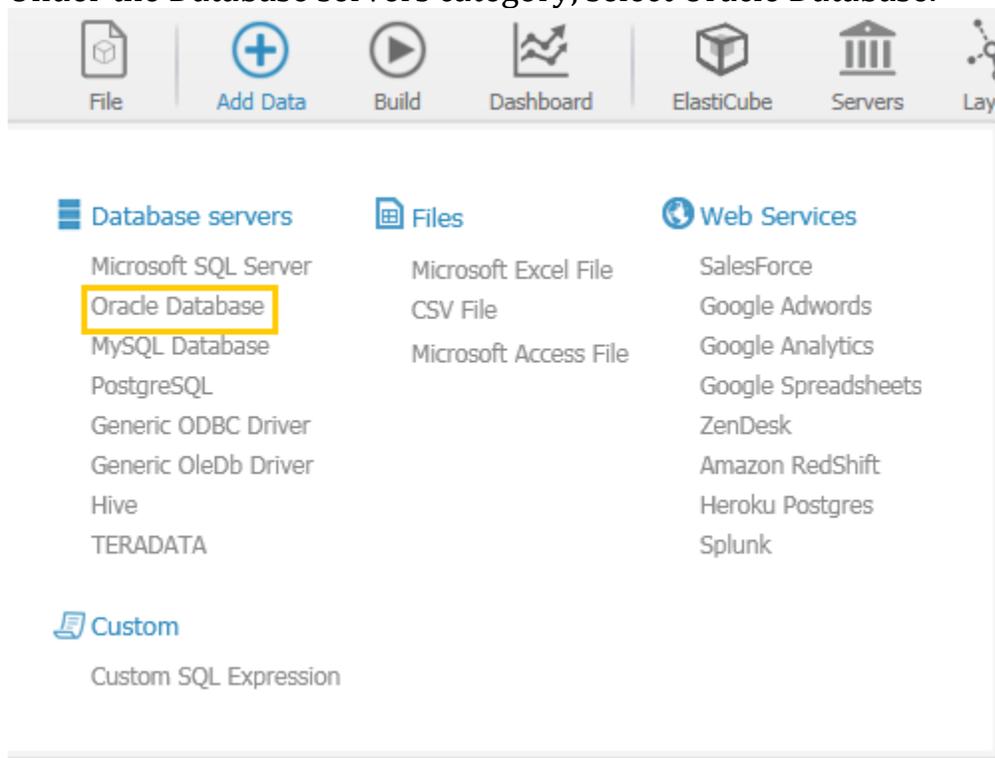


6. From the Tables list, select the relevant table or view you want to work with. You can click  next to the relevant table or view to see a preview of the data inside it. When you select the table or view, two new options are displayed at the bottom of the list, **Import Relationships** and **Add Custom Import SQL**.
7. (Optional) By default, existing relationships between tables are automatically replicated in the ElastiCube. You can disable this by toggling the **Import Relationships** switch.
8. (Optional) Click + to customize the data you want to import with SQL. See [Importing Data with Custom Queries](#) for more information.
9. After you have selected all the relevant tables, click **Done**. The tables are added to your schema in the ElastiCube Manager.

# Connecting to Oracle

The ElastiCube Manager enables easy and quick access to databases, tables and views contained within Oracle databases. The steps below describe how to connect to this type of data source.

1. Click **Add data** in the top menu of the ElastiCube Manager.
2. Under the Database servers category, select Oracle Database.



3. You are prompted to enter the following information:
  - **Database server location:** Enter the computer/server IP address with the database. To connect to a database running on your own computer, type in localhost.
  - **Login details:** Either use your Windows login details if they are configured with the database, or alternatively enter the username and password to connect to the database.
  - **Use Direct Connection:** You can also use a direct connection to connect. You will need to enter the Service ID and Port number to connect.

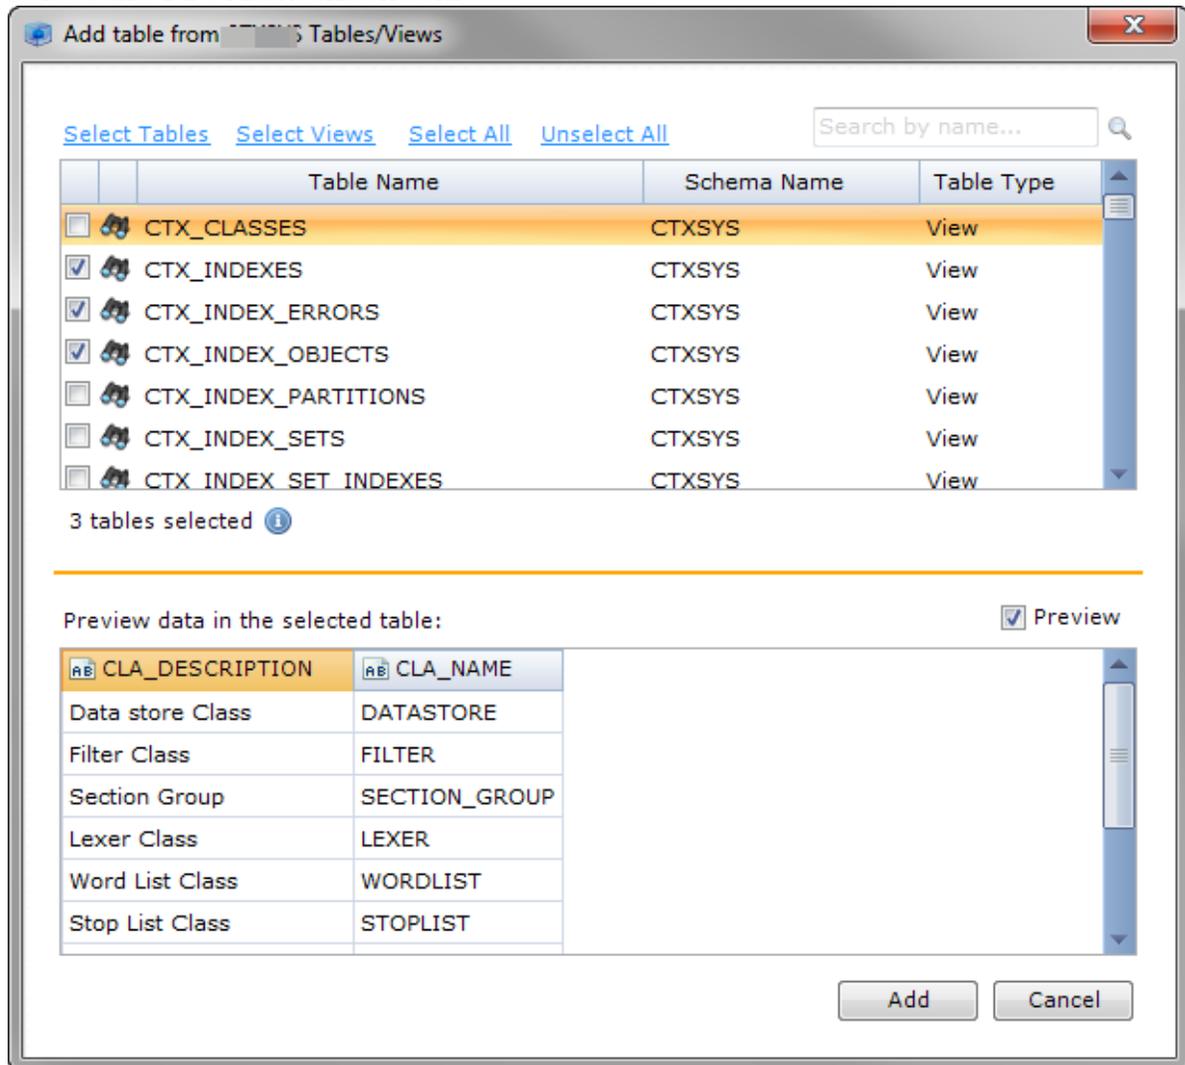
- As long as you choose to work in direct mode to Oracle (this can be defined in the Oracle provider wizard), no client software needs to be installed. If you want to work with an already installed Oracle client, you can use the indirect mode, which assumes you have the Oracle client already installed.
- Service ID can be found in the TNSNames.ora file. You can find a full description in the following topic: Service ID
- Note that in the file for each tnsname, you can find the service ID and service name.

4. Click **Connect to Server**.

5. A list of available databases will appear in the list box below.

6. Select the relevant database you want to work with and click **OK**. All tables and views associated with the database will appear in a new window. To view a preview of data contained in a particular table, highlight the table or view in the list and in click the preview pane below. To preview the table,

select the **Preview** checkbox.



7. Select the check box next to each table or view you want to use.
8. If you want to customize the data before importing it into the ElastiCube, you can run a custom SQL query to manipulate the data. This can be useful, for example, when you want to import only a portion of the data, rather than all of the data. [Click here](#) to learn more.
9. Once all relevant tables are selected, click **Add**.

# Connecting to PayPal

This connector is currently in beta and subject to change.

The Sisense PayPal connector is a standalone connector that allows you to import data from PayPal's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide ElastiCube Manager in the ElastiCube Manager. The connection string is used to authenticate users who connect to PayPal's API. To obtain a connection string, you will need to retrieve an API Key from PayPal or register a PayPal app.

Once you have connected to PayPal, you can import a variety of tables from the PayPal API.

This page describes how to install the PayPal connector, how to connect to PayPal with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the PayPal Connector](#)
- [Connecting to the PayPal REST API](#)
- [Adding PayPal Tables to your ElastiCube](#)
- [PayPal Tables](#)

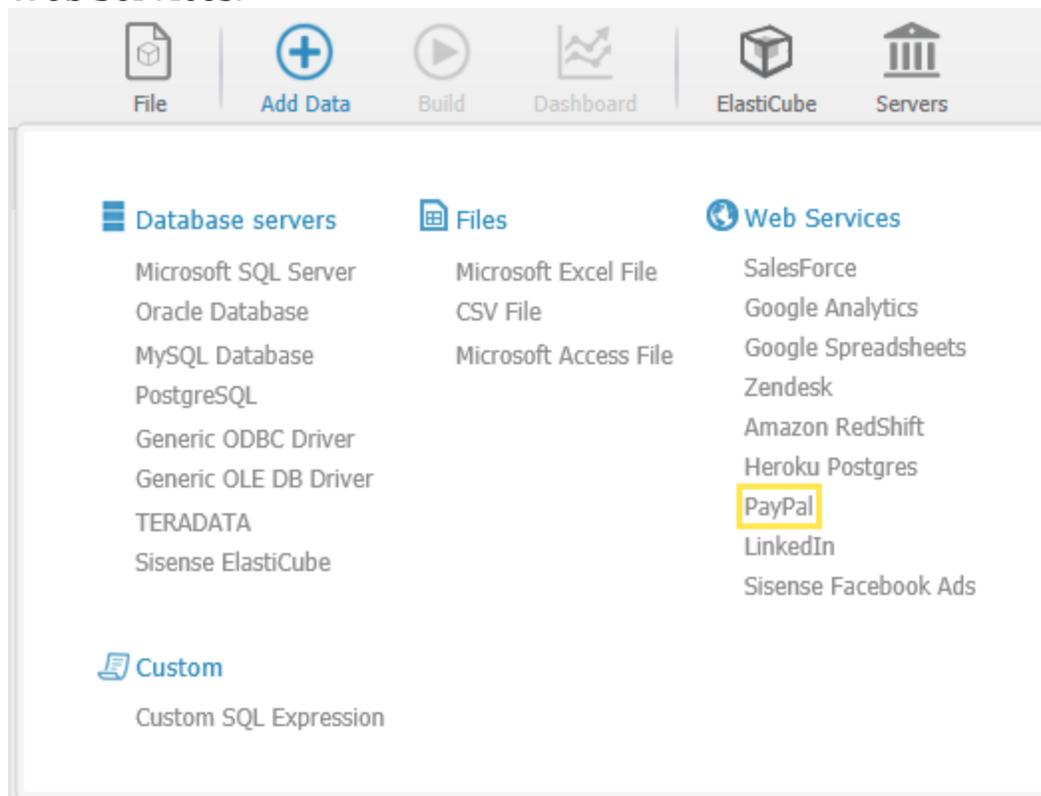
## Installing the PayPal Connector

Sisense provides the PayPal connector as a standalone connector that you can download and add to your list of default Sisense connectors.

**To install the PayPal connector:**

1. Contact us for the PayPal installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The PayPal connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.



## Connecting to the PayPal REST API

There are two authentication methods available for connecting to PayPal, through an API Key or through OAuth. Both methods expose the same data, the difference is in what information is required by PayPal to authenticate your account. For the API Key, all you need is a key provided by PayPal in your account. For OAuth, you need to register an app with PayPal, and then pass those details to Sisense.

After you have the relevant information from PayPal, you create a connection string and connect to PayPal.

To access PayPal's REST API from Sisense, you must provide valid OAuth PayPal credentials. These credentials are provided by PayPal when you create an application.

## Creating an App

You can follow the procedure below to register an app and obtain the OAuth client credentials, the Client Id and Client Secret:

1. Log in to your PayPal developer account and click the **Dashboard** tab.
2. On the REST API apps section, click **Create App**.
3. Enter the App Name.
4. From the list, select the sandbox account that will be used to get data from.

PayPal generates a set of sandbox and live OAuth keys for the application. The sandbox credentials are shown after creating the app. For the live credentials, select the Live option. The sandbox credentials will be used to get data only for the sandbox account selected when creating the App, while the Live credentials will be used to get data only for the developer account used to create this app.

## Adding PayPal Tables to Your ElastiCube

After you have retrieve your API Key or OAuth credentials from PayPal, you provide the relevant information in a connection string. Sisense uses connection strings to connect to PayPal and import data into the ElastiCube Manager. Each connection string contains a authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

The following parameters are mandatory for PayPal's connection string:

- **OAuthClientId**: Set this to the consumer key in your app settings.
- **OAuthClientSecret**: Set this to the consumer secret in your app settings.
- **CallbackURL**: Set this to the callback URL you specified in your app settings.

- **InitiateOAuth:** Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken and OAuthAccessTokenSecret.
- **UseSandbox:** Set this to true if you are using sandbox credentials.

For API Keys, the connection string to PayPal is in the following format:

```
jdbc:paypal:APIKey=xxxxxxxxxxxxxxxxxxxxxxxxxxxx;
```

For OAuth, the connection string to PayPal is in the following format:

```
jdbc:paypal:OAuthClientId=AZhGVsogGRX5uqwtE0B9i115uhFrbvS  
klbu2Hg;
```

```
OAuthClientSecret=EGFkKCToAep3uH9EMRtaj_GHWmm3;InitiateOA  
uth=GETANDREFRESH;
```

### To add PayPal data:

1. In ElastiCube Manager, click **Add Data** and then, **PayPal**. The Connect to PayPal window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. PayPal is displayed in the Select Database list.
4. Click **OK**. Sisense connects to PayPal and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

## PayPal Tables

The Sisense PayPal connector allows you to import the following tables into the ElastiCube Manager.

### Available Tables

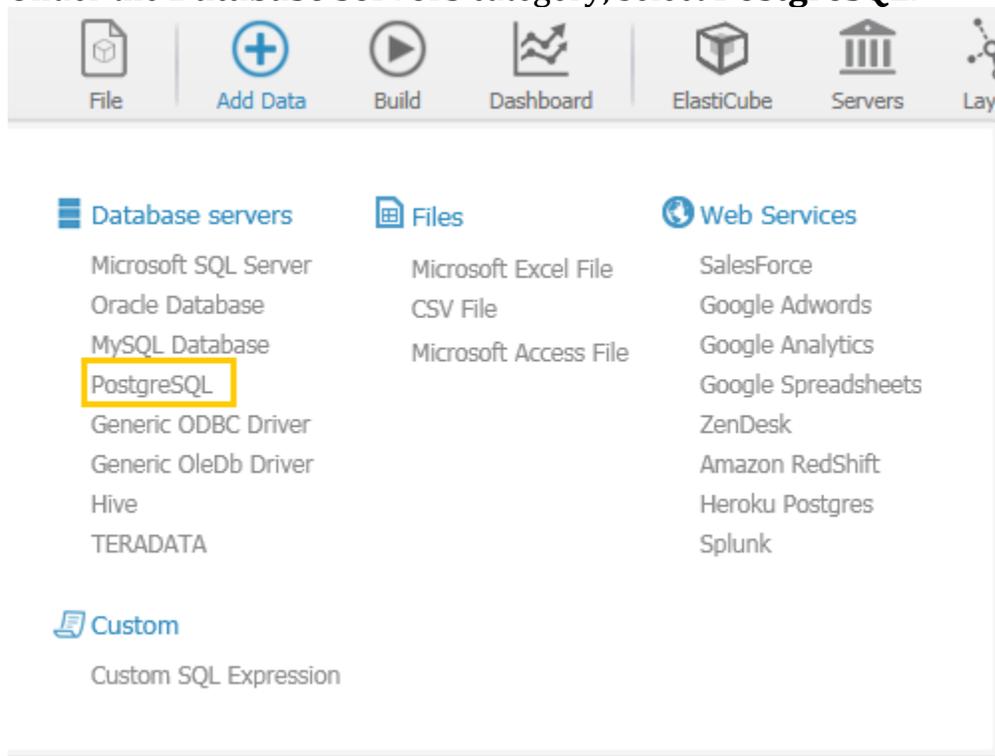
Name	Description
AuthorizationDetails	Query details about Authorizations.
CaptureDetails	Query details about a captured Payment.
CreditCardDetails	Query details about the funding instrument of the payer,

Name	Description
	such as a Credit Card or a token that represents a Credit Card.
Invoices	Query Invoices in PayPal.
Notifications	Query the list of event types that are subscribed to a Webhook.
OrderDetails	Query details about an Order in PayPal.
Payments	Query details about Payments
PaymentTransactions	Query Transaction details including the amount and item details.
RefundDetails	Query details about a specific Refund.
Refunds	Query the available refunds in Stripe.
SaleDetails	Query details about a Sale Transaction.
Webhooks	Query the list of Webhooks.

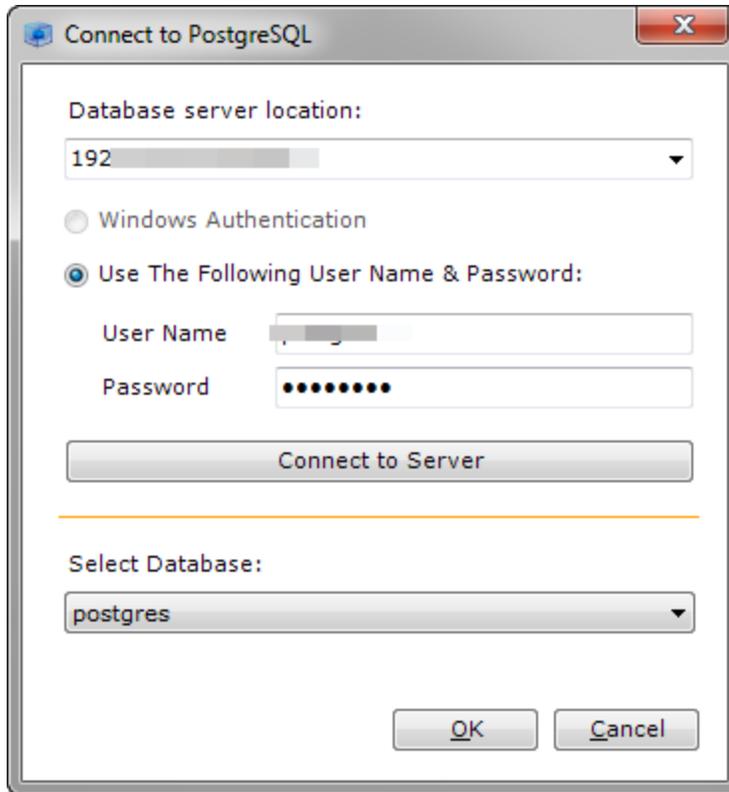
# Connecting to PostgreSQL

The ElastiCube Manager enables easy and quick access to databases, tables and views contained within PostgreSQL databases. The steps below describe how to connect to this type of data source.

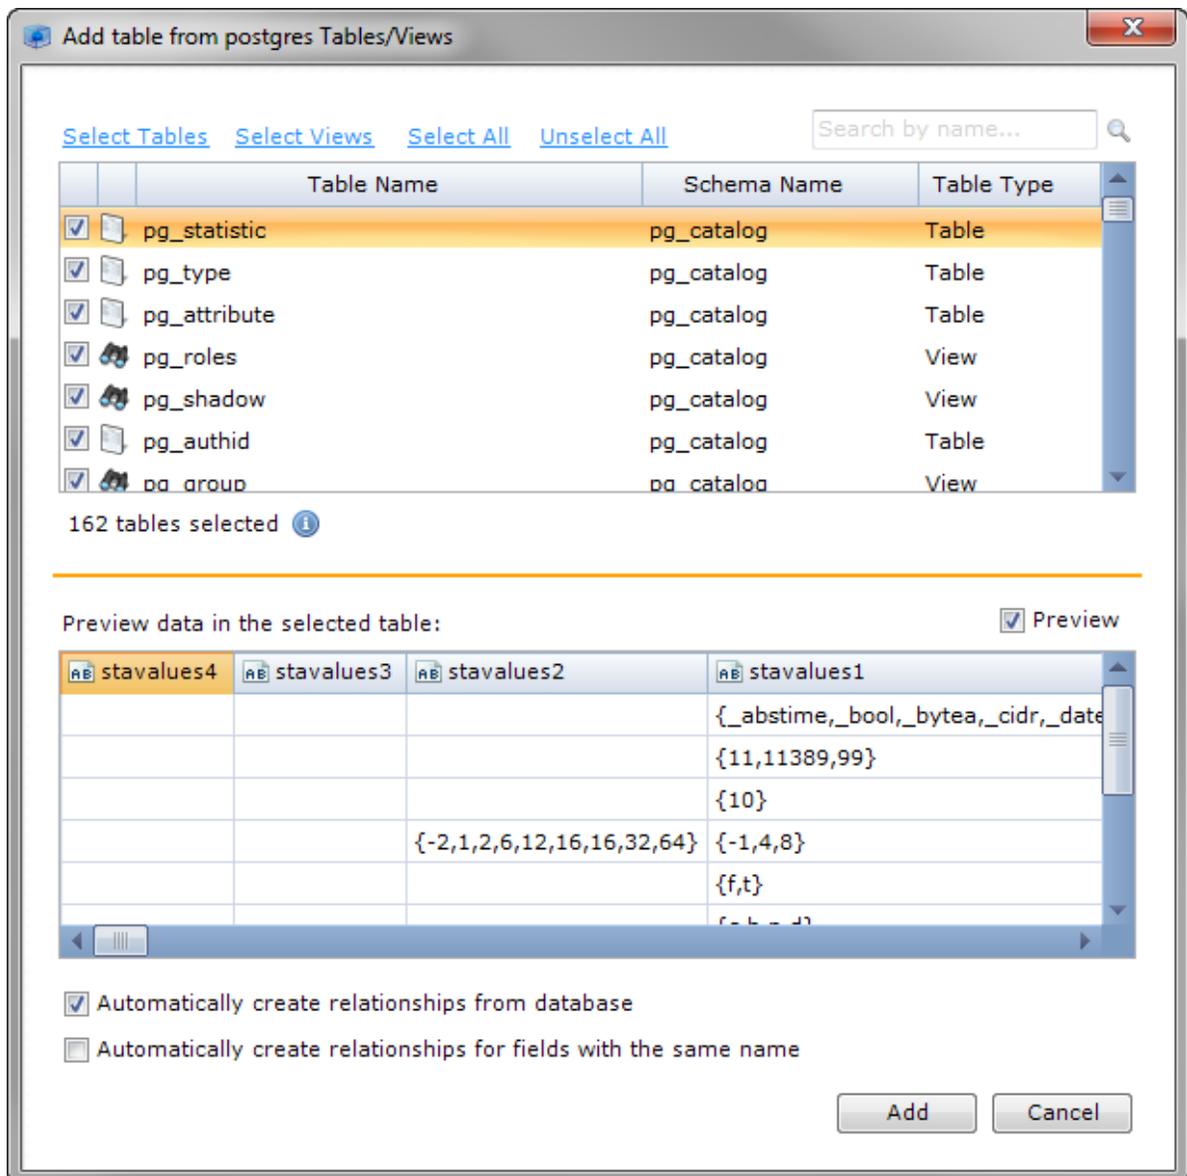
1. Click **Add Data** in the top menu of the ElastiCube Manager.
2. Under the **Database servers** category, select **PostgreSQL**.



3. You will be prompted to enter the following information:
  - **Database server location:** Enter the computer/server IP address where the database is located. To connect to a database on running on your own computer enter **localhost**.
  - **Login details:** Either use your Windows login details if they are configured with the database, or alternatively enter the username and password to connect to the database.



4. Click **Connect to Server**.
5. A list of available databases will appear in the list below.
6. Select the relevant database you want to work with and click **OK**.
7. All tables and views associated with the database will appear in a new window.



To view a preview of data contained in a particular table, highlight the table or view in the list and click the preview pane below. To preview the table, select the **Preview** checkbox.

8. Select the checkboxes next to each table or view you want to use.
9. Existing relationships between tables can be automatically replicated in the ElastiCube by selecting the **Automatically create relationships from database** checkbox. Likewise fields with similar names can be linked by selecting **Automatically create relationships for fields with the same name**. Once all relevant tables are selected, click **Add**.

**Note:** The default connection timeout limit is 30 seconds to PostgreSQL data sources. In Sisense V6.6.1 and later, you can modify this limit by editing the `Sisense.Connectors.PostgreSqlConnector.dll.config` located at `...\Program Files\Sisense\DataConnectors\DotNetContainer\Connectors\PostgreSql`. In the configuration file, under the section `PostgreSettings`, set the value of the parameter `Timeout` to change the connection timeout limit. This is useful in cases where you receive a `Timeout` exception when trying to connect.

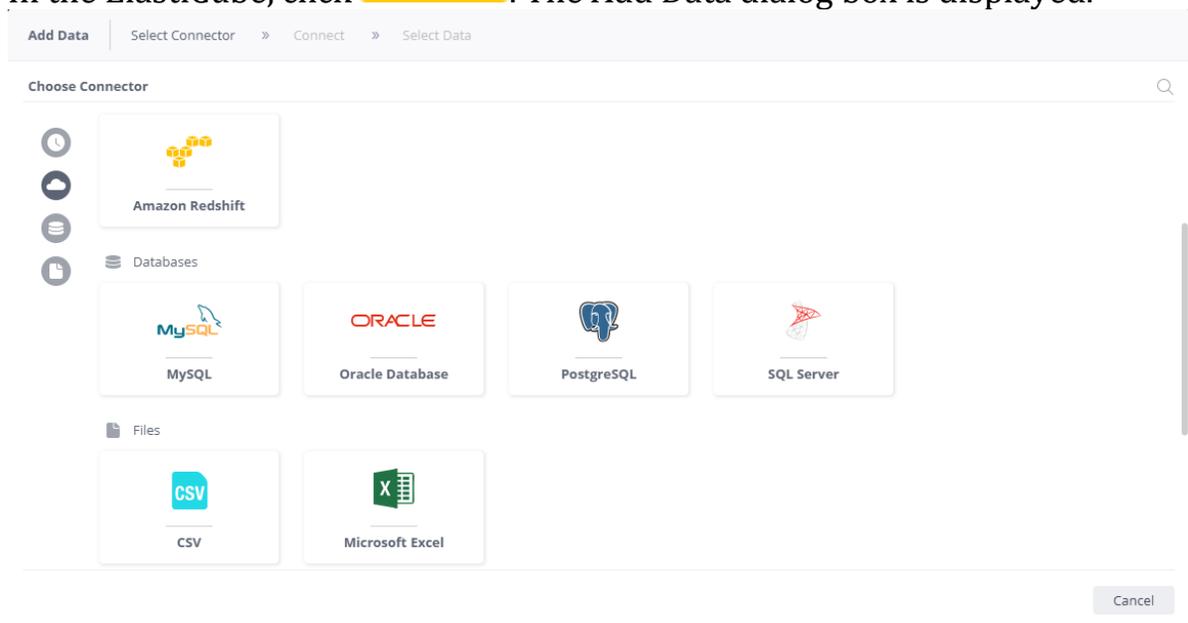
```
1 <?xml version="1.0" encoding="utf-8"?>
2 <configuration>
3   <configSections>
4     <section name="PostgreSettings" type="System.Configuration.DictionarySectionHandler" />
5   </configSections>
6   <startup>
7     <supportedRuntime version="v4.0" sku=".NETFramework,Version=v4.5.2" />
8   </startup>
9   <PostgreSettings>
10    <!--
11     The time to wait (in seconds) while trying to execute a command before terminating the attempt and generating an error.
12     Use 0 for infinity.
13    -->
14    <add key="Timeout" value="30" />
15  </PostgreSettings>
16 </configuration>
```

# Connecting to PostgreSQL

The ElastiCube Manager enables easy and quick access to databases, tables and views contained within PostgreSQL databases. The steps below describe how to connect to this type of data source.

## To import data from PostgreSQL:

1. In the Data page of ElastiCube Manager, open an ElastiCube or create a new ElastiCube.
2. In the ElastiCube, click . The Add Data dialog box is displayed.



3. Click **PostgreSQL** to open the PostgreSQL settings.

Add Data | Select Connector » Connect » Select Data

PostgreSQL

Location 127.0.0.1

User name

Password

Encrypt Connection  
 Trust Server Certificate

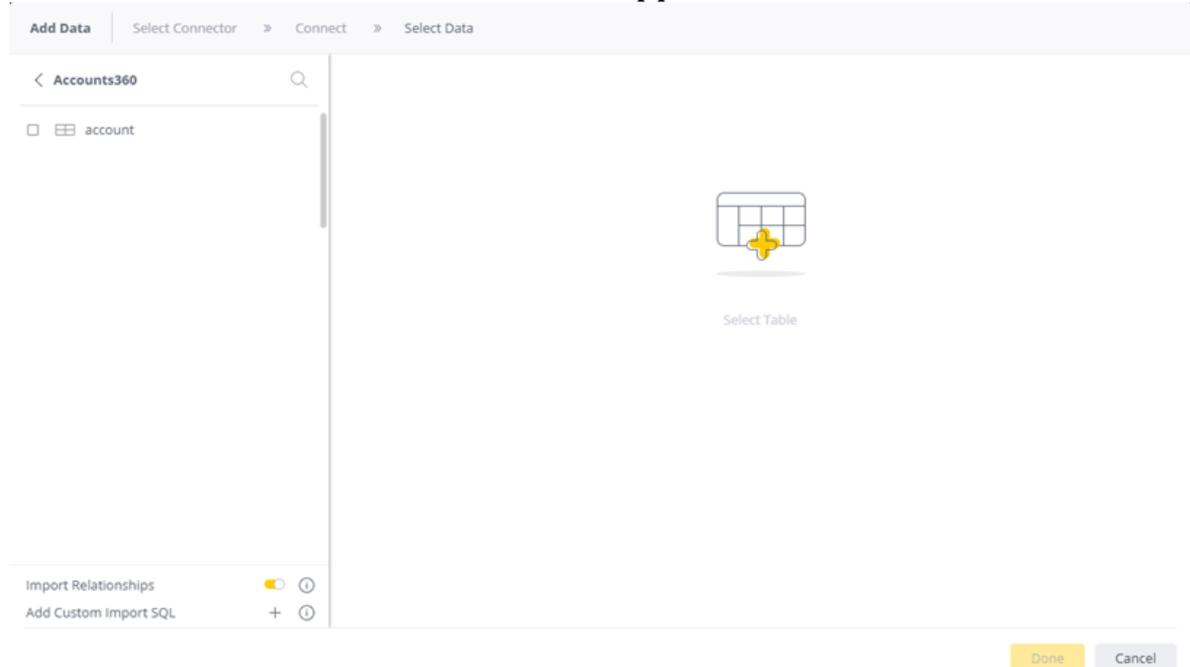
Next Cancel

4. Enter the following details:

- **Location:** Enter the computer/server IP address of the database. To connect to a database running on your own computer enter **localhost**.
- If Windows Authentication is configured with the database in the **User Name** and **Password** fields, enter your database credentials. Name and Password fields, enter your database credentials.

**(Optional)** Select **Encrypt Connection** to configure the driver to encrypt all communication with the PostgreSQL instance before sending it over the network, and then select **Trust Server Certificate**.

5. Click **Next**. A list of tables in the database are displayed. All tables and views associated with the database will appear in a new window.



6. From the Tables list, select the relevant table or view you want to work with. You can click  next to the relevant table or view to see a preview of the data inside it. When you select the table or view, two new options are displayed at the bottom of the list, **Import Relationships** and **Add Custom Import SQL**.
7. (Optional) By default, existing relationships between tables are automatically replicated in the ElastiCube. You can disable this by toggling the **Import Relationships** switch.
8. (Optional) Click + to customize the data you want to import with SQL. See [Importing Data with Custom Queries](#) for more information.
9. After you have selected all the relevant tables, click **Done**. The tables are added to your schema in the ElastiCube Manager.

**Note:** The default connection timeout limit is 30 seconds to PostgreSQL data sources. You can modify this limit by editing the `Sisense.Connectors.PostgreSqlConnector.dll.config` located at `...\Program Files\Sisense\DataConnectors\DotNetContainer\Connectors\PostgreSql`. In the configuration file, under the section `PostgreSettings`, set the value of the parameter `Timeout` to change the connection timeout limit. This is useful in cases where you receive a `Timeout` exception when trying to

## connect.

```
1 <?xml version="1.0" encoding="utf-8"?>
2 <configuration>
3 <configSections>
4 <section name="PostgreSettings" type="System.Configuration.DictionarySectionHandler" />
5 </configSections>
6 <startup>
7 <supportedRuntime version="v4.0" sku=".NETFramework,Version=v4.5.2" />
8 </startup>
9 <PostgreSettings>
10 <!--
11 | The time to wait (in seconds) while trying to execute a command before terminating the attempt and generating an error.
12 | Use 0 for infinity.
13 |-->
14 <add key="Timeout" value="30" />
15 </PostgreSettings>
16 </configuration>
```

# PostgreSQL Live

This topic describes how you can add live connections to your PostgreSQL Server databases from the Sisense Web Application.

For information about live connections, see [Sisense Live Connect](#).

## Adding Live Connections

To add a live connection:

1. In the Sisense Web Application, select **Admin > Data Sources**. The Data Sources page is displayed.

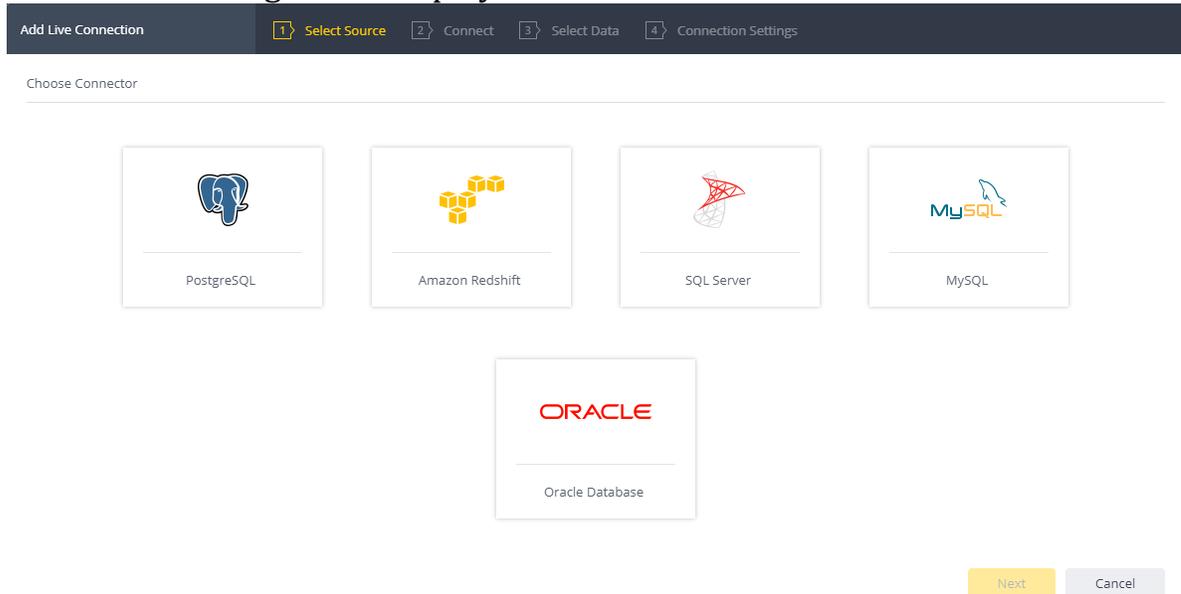
The screenshot displays the Sisense Admin interface. The top navigation bar includes 'Home', 'Pulse', 'ElastiCube Manager', and 'Admin'. The left sidebar shows the 'Admin' menu with 'ElastiCubes & Data Sets' selected. The main content area is titled 'ElastiCubes & Data Sets' and contains two tables.

ElastiCube	Status	Created On	Last Built	Size On Disk	
Nortwind BusinessDays 2	Stopped	5/31/15 5:02 PM	10/24/16 2:51 PM	4.72MB	Open
Sample eCommerce	Running	2/10/14 2:21 PM	3/12/14 10:18 AM	59.05MB	Open
Sample Healthcare	Running	7/2/13 1:07 AM	3/12/14 2:01 PM	1.5MB	Open
Sample Lead Generation	Running	3/3/14 7:46 PM	3/12/14 11:59 AM	2.23MB	Open
Test	Stopped	7/18/16 1:06 PM	1/24/17 3:25 PM	30.49MB	Open

Data Set	Source	Server Location	User Name	Refresh Rate	
Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	Open
Copy of Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	Open
customers	MySql Database	192.168.1.229	Test	10 Sec.	Open
datatypes	MySql Database	192.168.1.229	Test	10 Sec.	Open
LiveSource	Microsoft Sql Server	192.168.1.229	Test	2 Min.	Open
MsSQL Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	Open
productlines	MySql Database	192.168.1.229	Test	10 Sec.	Open

2. In the Data Sources page, click **Add Live Connection**. The Add Live Connection dialog box is displayed.



3. In the Add Live Connection dialog box, select your live data source.
4. Under the Connect tab, enter the following credentials:
  - Location:** The IP address of your data source.
  - User Name:** The user name to access data source.
  - Password:** The password to access your data source.
5. (Optional) Select **Encrypt Connection** to configure the driver to encrypt all communication with the PostgreSQL Server instance before sending it over the network, and then select **Trust Server Certificate**.
6. Click **Next**.
7. Under the Select Data tab, on the left side is a list of schemas located in your data source. Select the relevant schema. You can find the schema easily by searching for it through the search field at the top of the list to filter the schemas displayed. After selecting the schema, a list of the available tables and views are displayed.
8. Select a table or view. You can find the table easily by searching for it through the search field at the top of the list to filter the tables displayed.
9. Click **Next**.
10. Under the Connection Settings tab, set the following:
  - Connection Name:** Enter a name for your connection. This field defines the name of your data source displayed in the Sisense Web Application.
  - Refresh Rate:** Enter or select the refresh rate for your data source. The shortest refresh rate you can set is 10 seconds. The refresh rate is the time

in between queries to the data source. If the previous query has not been completed, the next query is not sent until the previous refresh is complete unless the refresh rate is greater than the time it took to complete the refresh since the last query. For example, if your refresh rate is 10 seconds, and your query takes 15 seconds, the next query is sent when the previous query is complete as 10 seconds has passed and the previous query is complete. If your refresh rate is 30 seconds, and the previous query is complete, the next refresh will be in 30 seconds from the last query.

**Time Out:** Enter or select the Time Out rate for your data source. The Time Out rate determines how long Sisense waits before cancelling a query that does not receive a response from your data source.

11. Click **Done**. The connection is added as a data set to your Data Sources page.

# Preparing Your Excel Files

Use the following guidelines to prepare and verify Excel files before connecting them with the ElastiCube.

Learn more about Sisense's Excel connector.

## Structure Your Excel Sheet

All Excel data must be imported into the ElastiCube Manager in a **tabular format**, to allow for accurate data categorization.

Characteristics of Tabular Data include:

1. Each **sheet** represents a **class**, for example customers, sales, employees.
2. Each **column** represents a **variable**, for example revenue, name, address.
3. Each **row** represents a **single observation**, for example customer X's details.

The screenshot shows an Excel spreadsheet with the following data:

	A	B	C	D
1	CustomerID	Address	City	CompanyName
2	ALFKI	Obere Str. 57	Berlin	Alfreds Futterkiste
3	ANATR	Avda. de la Constitución 2222	México D.F.	Ana Trujillo Emparedados
4	ANTON	Mataderos 2312	México D.F.	Antonio Moreno Taquería
5	AROUT	120 Hanover Sq.	London	Around the Horn
6	BERGS	Berguvsvägen 8	Luleå	Berglunds snabbköp
7	BLAUS	Forsterstr. 57	Mannheim	Blauer See Delikatessen
8	BLONP	24, place Kléber	Strasbourg	Blondesddsl père et fils
9	BOLID	C/ Araquil, 67	Madrid	Bólido Comidas preparada
10	BONAP	12, rue des Bouchers	Marseille	Bon app'
11	BOTTM	23 Tsawassen Blvd.	Tsawassen	Bottom-Dollar Markets
12	BSBEV	Fauntleroy Circus	London	B's Beverages

Callouts in the image:

- 1: Sheet is related class - Customers
- 2: Column represents a single variable of interest - the name of the company
- 3: Row only has details for a single customer

# Remove Blank Rows and Duplicate Headers

Remove blank rows OR rows after the first column that have duplicate headers.

Blank rows may return incorrect calculations and results.

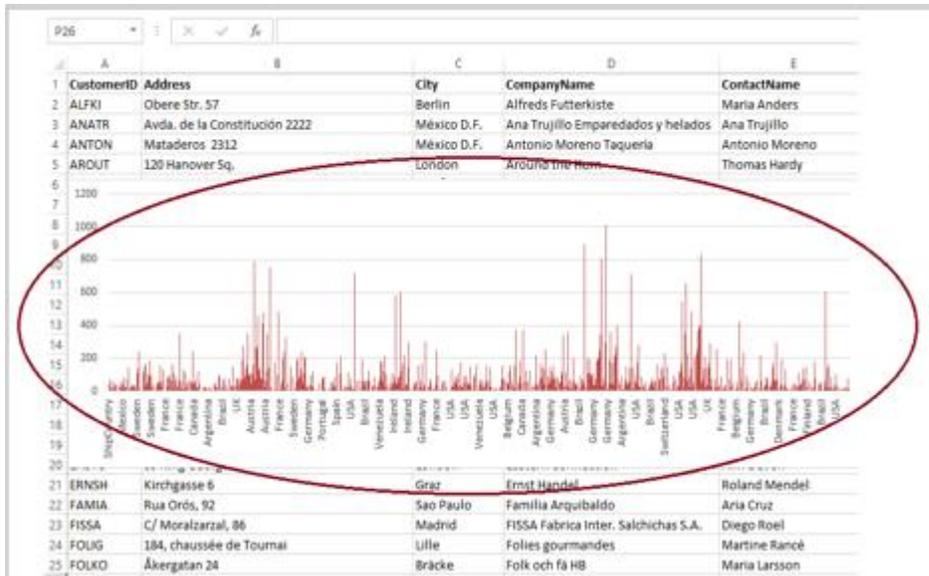
CustomerID	Address	City	CompanyName
ALFKI	Obere Str. 57	Berlin	Alfreds Futterkiste
ANATR	Avda. de la Constitución 2222	México D.F.	Ana Trujillo Emparedados
AROUT	120 Hanover Sq.	London	Around the Horn
BERGS	Berguvsvägen 8	Luleå	Berglunds snabbköp
BLAUS	Forsterstr. 57	Mannheim	Blauer See Delikatessen
BOLID	C/ Araquil, 67	Madrid	Bolido Comidas preparadas
BONAP	12, rue des Bouchers	Marseille	Bon app'
BOTTM	23 Tsawassen Blvd.	Tsawassen	Bottom-Dollar Markets
BSBEV	Fauntleroy Circus	London	B's Beverages
CustomerID	Address	City	CompanyName
CENTC	Sierras de Granada 9993	México D.F.	Centro comercial Moctezuma
CHOPS	Hauptstr. 29	Bern	Chop-suey Chinese
COMMI	Av. dos Lusíadas, 23	Sao Paulo	Comércio Mineiro
CONSH	Berkeley Gardens 12 Brewery	London	Consolidated Holdings
DRACD	Walserweg 21	Aachen	Drachenblut Delikatessen
DUMON	67, rue des Cinquante Otages	Nantes	Du monde entier

**Blank Rows:** Delete rows 4 and 8 if they contain no data.

**Additional Headers:** Delete row 13 as it has the same titles as the main headers. If this heading represented a different class of data, a new sheet should be created.

## Remove All Unstructured Data

Remove all unstructured data from each sheet, such as charts, images or non-data elements. Unstructured data cannot be imported, and may result in other data being obscured.



**Images, Charts and Non-Data Elements:** All non-data elements must be removed from the sheet, such as the chart in the image above. All visualization based on Excel data can be created in a dashboard using Sisense.

## Rename/Remove Duplicate Column Names

Rename or remove columns with the same name. Ensure all columns have a unique name. Columns with the same name cannot be imported into the ElastiCube.

	A	B	C	D	E
1	CustomerID	Address	City	CompanyName	CompanyName
2	ALFKI	Obere Str. 57	Berlin	Alfreds Futterkiste	Maria Anders
3	ANATR	Avda. de la Constitución 2222	México D.F.	Ana Trujillo Emparedados y helados	Ana Trujillo
4	ANTON	Mataderos 2312	México D.F.	Antonio Moreno Taquería	Antonio Moreno
5	AROUT	120 Hanover Sq.	London	Around the Horn	Thomas Hardy
6	BERGS	Berguvsvägen 8	Luleå	Berglunds snabbköp	Christina Berglund
7	BLAUS	Forsterstr. 57	Mannheim	Blauer See Delikatessen	Hanna Moos

**Column Names:** Rename columns with the same name, and ensure all columns have a unique name. For example, in the above example 'CompanyName' appears twice. Once column should be given a different name such as 'Secondary CompanyName'. In the ElastiCube, you can rename any column as you wish.

# Modify Columns for Quicker Analysis

Consolidate columns as much as possible into a single variable. For example, instead of having a column for each month of the year, it is best to have a single column with all dates. Unconsolidated columns limit the ability to analyze and segment data.

OrderID	ShipAddr	ShipCountry	Jan-13	Feb-13	Mar-13	Apr-13	May-13
10952	Obere Str.	Germany	40,42	84,84	109,11	88,4	52,52
10835	Obere Str.	Germany	69,53	23,72	92,69	77,92	14,01
11011	Obere Str.	Germany	1,21	25,36	188,85	97,09	74,44
10702	Obere Str.	Germany	23,94	72,97	13,32	16,16	2,4
10692	Obere Str.	Germany	61,02	4,52	95,75	11,06	48,83
10643	Obere Str.	Germany	29,46	146,32	27,71	63,54	62,89
10759	Avda. de l	Mexico	11,99	3,04	0,15	166,31	243,73
10625	Avda. de l	Mexico	43,9	10,96	44,65	113,15	44,17
10926	Avda. de l	Mexico	39,92	21,74	1,93	16,97	47,42

 Incorrect column structure

OrderID	ShipAddress	ShipCountry	Revenue	OrderDate
10952	Obere Str. 57	Germany	40,42	2011-03-16 00:00
10835	Obere Str. 57	Germany	69,53	2011-01-15 00:00
11011	Obere Str. 57	Germany	1,21	2011-04-09 00:00
10702	Obere Str. 57	Germany	23,94	2010-10-13 00:00
10692	Obere Str. 57	Germany	61,02	2010-10-03 00:00
10643	Obere Str. 57	Germany	29,46	2010-08-25 00:00

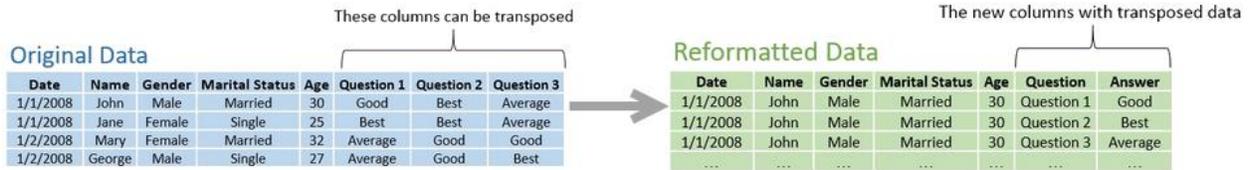
 Correct column structure

**Consolidate Columns:** Data should be consolidated into a single column for each variable. In the example above instead of having revenue reflected separately as a column for each month, it is best to consolidate the data into a single column for date and another for revenue. This logic should be applied to any other segments such as summary by country, or channel, and so forth.

## Additional Methods and Tools for Consolidating (Transposing) Columns

To consolidate or transpose columns in your files, you can use one or both of the following methods:

- Use an Excel macro to transpose data. This is useful when your data is already in Excel or easy to put into Excel.
- Create custom SQL. This method should be used if importing your data into Excel is more difficult, or the amount of data is causing poor performance in Excel. This method requires basic SQL knowledge.



To learn more about both of the above methods, and to download the macro, click [here](#).

## Convert Scientific Formatting to Numeric Formatting

Convert cells with scientific formatting to another number format. Data in scientific format cannot be imported correctly into the ElastiCube.

C	D	E	F
ProductID	ProductName	QuantityPerUnit	ReorderLevel
1	Chai	10 boxes x 20 bags	1,00E+04
2	Chang	24 - 12 oz bottles	2,50E+04
34	Sasquatch Ale	24 - 12 oz bottles	1,50E+04
35	Steeleye Stout	24 - 12 oz bottles	1,50E+04
38	Côte de Blaye	12 - 75 cl bottles	1,50E+04
39	Chartreuse verte	750 cc per bottle	5,00E+03
43	Ipoh Coffee	16 - 500 g tins	2,50E+04
67	Laughing Lumberjack Lager	24 - 12 oz bottles	1,00E+04
70	Outback Lager	24 - 355 ml bottles	3,00E+04
75	Rhönbräu Klosterbier	24 - 0.5 l bottles	2,50E+04
76	Lakkalikööri	500 ml	2,00E+04

**Scientific Format:** Convert all values in scientific format to another number format such as number, currency, accounting, fraction and so forth.

## Troubleshooting

When importing an Excel spreadsheet with the XLSX format, if you get the error: "Unable to connect to the specified file. Corrupt OpenXML document.", try saving the file as XLS and importing it again. If this works, keep in mind that in an XLS

workbook, the row limit is 65,536 (216) and 256 columns (28) which correspond to the column IV. With XLSX workbooks, limits are 1,048,576 rows (220) and 16,384 columns (214).

# Connecting to Quickbooks Desktop

The Sisense Quickbooks Desktop connector is a standalone connector that allows you to import data from Quickbooks Desktop's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the Quickbooks Desktop APIs. To obtain a connection string, you will need to create a Quickbooks Desktop app.

Once you have connected to Quickbooks Desktop, you can import a variety of tables from the Quickbooks Desktop API.

This page describes how to install the Quickbooks Desktop connector, how to connect to Quickbooks Desktop with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the Quickbooks Desktop Connector](#)
- [Connecting to the Quickbooks Desktop REST API](#)
- [Adding Quickbooks Desktop Tables to your ElastiCube](#)
- [Quickbooks Desktop Tables](#)

## Installing the Quickbooks Desktop Connector

Sisense provides the Quickbooks Desktop connector as a standalone connector that you can download and add to your list of default Sisense connectors.

**To install the Quickbooks Desktop connector:**

1. [Download](#) the Quickbooks Desktop installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The Quickbooks Desktop connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.

## Connecting to the Quickbooks Desktop REST API

The Quickbooks Desktop connector makes requests to QuickBooks through the Remote Connector. The Remote Connector is an easy-to-use tool that enables developers to access QuickBooks data remotely. The Remote Connector runs on the same machine as QuickBooks and accepts connections through a lightweight, embedded Web server.

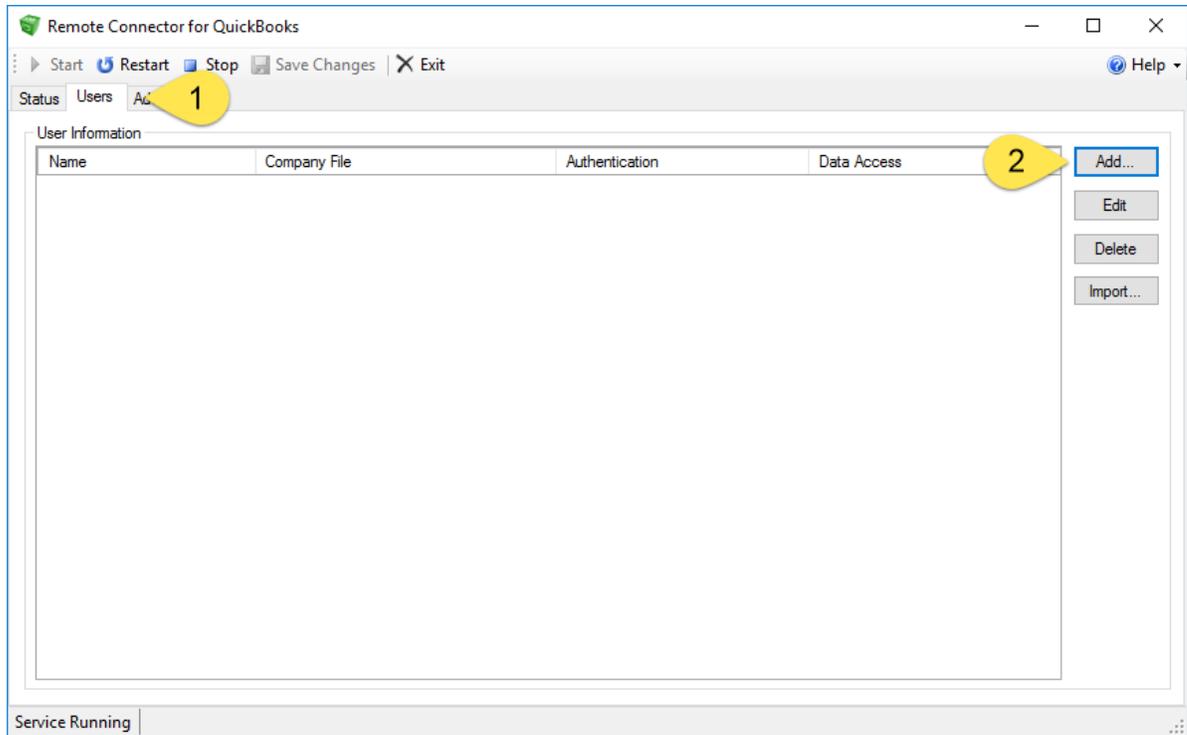
The server supports SSL/TLS, enabling users to connect securely from remote machines. The first time you connect, you will need to authorize the driver with QuickBooks.

The Remote Connector can be used to read and write to QuickBooks in situations where direct COM access to QuickBooks is not available (e.g., ASP.NET, Java, or QuickBooks on a remote machine).

In the Remote Connector, you can define users and their passwords and then use these credentials in the connection string that you create to connect to your QuickBooks Desktop tables.

Follow the procedure below to connect to QuickBooks for the first time through the Remote Connector:

1. Download the Remote Connector from [remotecconnector.com](http://remotecconnector.com) and install the Remote Connector on the machine where QuickBooks is installed.
2. Open the company file you want to connect to in QuickBooks using an administrator account in single-user mode.
3. Open the Remote Connector from the system tray and add a user on the Users tab. Enter a User and Password and select the level of access in the Data Access menu.



**Note:** The Remote Connector does not use the User and Password properties to access QuickBooks; the User and Password properties authenticate the user to the Remote Connector. Authentication to QuickBooks is handled based on the Application Name property.

4. When you first connect, a dialog will appear in QuickBooks prompting you to authorize the application. After authorizing the application, you can then execute commands to QuickBooks. Specify the URL of the Remote Connector and the User and Password. By default, the Remote Connector connects to the currently open company file.
5. If you want to access QuickBooks when QuickBooks is not running, save the company file information for the user. The Remote Connector will then automatically open QuickBooks in the background with the company file for that user.

**Note:** If the QuickBooks UI is open, you can only connect to that company file.

Additionally, note that the user permissions you run the Remote Connector under must match the user permissions you run QuickBooks under. The Remote Connector installation process installs the Remote Connector as a service under the current user account.

# Adding Quickbooks Desktop Tables to your ElastiCube

Sisense uses connection strings to connect to Quickbooks Desktop and import data into the ElastiCube Manager.

The connection string to connect to Quickbooks Desktop has the following structure:

```
jdbc:Quickbooks
```

```
Desktop:Property1=Value1;Property2=Value2;
```

The following is an example of a Quickbooks Desktop connection string:

```
jdbc:quickbooks:User=xxxxx;Password=xxxxxxxxxxx;URL=http://  
/xxxxxxxxxxxxx;
```

- **User:** Set this to username defined in Remote Connector.
- **Password:** Set this to password defined in Remote Connector.
- **URL:** The IP address of the remote machine where Remote Connector was installed and the port. The port is defined under Connection Settings in the Advanced tab of the Remote Connector.
- **OAuthSettingsLocation:** The location of the settings file where OAuth values are saved for each user when InitiateOAuth is set to GETANDREFRESH. See [Switching between Accounts](#) for more information.

**Note:** To switch between accounts, you need to delete the file OAuthsettings.txt file located at

...\Users\xxx\AppData\Roaming\CData\QuickbooksDesktop\DataProvider.

## To add Quickbooks Desktop data:

1. In ElastiCube Manager, click **Add Data** and then, **Quickbooks Desktop**. The Connect to Quickbooks Desktop window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. Quickbooks Desktop is displayed in the Select Database list.
4. Click **OK**. Sisense connects to Quickbooks Desktop and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

## Switching between Accounts

When you connect to the Quickbooks Desktop data source, Sisense saves your OAuth values in the file OAuthsettings.txt file located at ...\\Users\\xxx\\AppData\\Roaming\\CData\\Quickbooks Desktop Data Provider on your Sisense server. To connect to the Quickbooks Desktop data source with another user on the same machine, you must delete the OAuthsettings.txt file. Sisense will then generate a new file for that user.

Another option to support multiple users is to define the location and file name of an OAuthsettings file for each unique user in your connection string through the **OAuthSettingsLocation** parameter. When each user connects to the data source, Sisense generates the OAuth file with the file name you specify in the location you define. In the examples below, two users are allowed to access the Quickbooks Desktop data source and for each user, Sisense generates a file that contains that user's OAuth values in the location defined in the string.

```
jdbc:QuickbooksDesktop:OAuthSettingsLocation=C:\\QuickbooksDesktop\\auth\\john.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
```

```
jdbc:QuickbooksDesktop:OAuthSettingsLocation=C:\\QuickbooksDesktop\\auth\\sally.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/;
```

In the example above, two OAuth files are created, one for John and one for Sally in the location C:\\QuickbooksDesktop\\auth\\.

This is useful if you support many users who each need to access the Quickbooks Desktop data source.

# Quickbooks Desktop Tables

Quickbooks Desktop's RESTful APIs expose the following Quickbooks Desktop tables that you can import into the ElastiCube Manager through the Sisense Quickbooks Desktop connector:

## Available Tables

Name	Description
Accounts	Create, update, delete, and query QuickBooks Accounts. To update Accounts, set the QBXMLVersion to 6.0 or higher.
BillExpenseItems	Create, update, delete, and query QuickBooks Bill Expense Line Items.
BillLineItems	Create, update, delete, and query QuickBooks Bill Line Items.
BillPaymentChecks	Create, update, delete, and query QuickBooks Bill Payment Checks. QBXMLVersion must be set to 6.0 or higher to update a BillPaymentCheck.
BillPaymentChecksAppliedTo	Create, update, delete, and query QuickBooks Bill Payment AppliedTo aggregates. In a Bill Payment, each AppliedTo aggregate represents the Bill transaction to which this part of the payment is being applied.
BillPaymentCreditCards	Create, update, delete, and query QuickBooks Bill Payments.

Name	Description
BillPaymentCreditCardsAppliedTo	Create, update, delete, and query QuickBooks Bill Payment AppliedTo aggregates. In a Bill Payment, each AppliedTo aggregate represents the Bill transaction to which this part of the payment is being applied.
Bills	Create, update, delete, and query QuickBooks Bills.
BuildAssemblies	Insert, Update, delete, and query QuickBooks Build Assembly transactions.
BuildAssemblyLineItems	Create and query QuickBooks Build Assembly transactions.
CheckExpenseItems	Create, update, delete, and query QuickBooks Check Expense Line Items.
Checks	Create, update, delete, and query QuickBooks Check Line Items.
Class	Create, update, delete, and query QuickBooks Classes. QuickBooks requires QBXML version 8.0 or higher for updates to a Class.
CreditCardChargeExpenseItems	Create, update, delete, and query QuickBooks Credit Card Charge Expense Line Items.
CreditCardChargeLineItems	Create, update, delete, and query QuickBooks Credit Card Charge Line Items.
CreditCardCharges	Create, update, delete, and query QuickBooks Credit Card Charges.

Name	Description
CreditCardCreditExpenseItems	Create, update, delete, and query QuickBooks Credit Card Credit Expense Line Items.
CreditCardCreditLineItems	Create, update, delete, and query QuickBooks Credit Card Credit Line Items.
CreditCardCredits	Create, update, delete, and query QuickBooks Credit Card Credits.
CreditMemoLineItems	Create, update, delete, and query QuickBooks Credit Memo Line Items.
CreditMemos	Create, update, delete, and query QuickBooks Credit Memos.
Currency	Create, update, delete, and query QuickBooks Currencies. This table requires QBXML version 8.0 or higher, and you will need to enable multiple Currencies in your QuickBooks company file to use it.
CustomerContacts	Create, update, delete, and query QuickBooks Customer Contacts. This table requires QBXML version 12.0 or higher, and is only available in QuickBooks editions 2016 and above.
CustomerMessages	Create, delete, and query Customer Messages.
CustomerNotes	Create, update, and query QuickBooks Customer Notes. This table requires QBXML version 12.0 or higher.
Customers	Create, update, delete, and query QuickBooks Customers.

Name	Description
CustomerShippingAddresses	Create, update, delete, and query QuickBooks Customer Shipping Addresses. Multiple Customer Shipping Addresses are supported in only QuickBooks 2013 and higher. Additionally, this table requires QBXML version 12.0 or higher.
CustomerTypes	Create, update, delete, and query QuickBooks Customer Types.
DateDrivenTerms	Create, delete, and query QuickBooks Date Driven Terms.
DepositLineItems	Create, update, delete, and query QuickBooks Deposit Line Items. QBXMLVersion must be set to 7.0 or higher to update a deposit.
Deposits	Create, update, delete, and query QuickBooks Deposits. QBXMLVersion must be set to 7.0 or higher to update a deposit.
EmployeeEarnings	Create, update, delete, and query QuickBooks Employee Earnings.
Employees	Create, update, delete, and query QuickBooks Employees.
EstimateLineItems	Create, update, delete, and query QuickBooks Estimate Line Items.
Estimates	Create, update, delete, and query QuickBooks Estimates.
InventoryAdjustmentLineItems	Create and query QuickBooks Inventory

Name	Description
	Adjustment Line Items.
InventoryAdjustments	Create, query, and delete QuickBooks Inventory Adjustments.
InventorySites	Create, update, delete, and query QuickBooks Inventory Sites. Inventory Sites are only available in QuickBooks Enterprise 2010 and above, and only with the Advanced Inventory add-on.
InvoiceLineItems	Create, update, delete, and query QuickBooks Invoice Line Items.
Invoices	Create, update, delete, and query QuickBooks Invoices.
ItemLineItems	Create, update, delete, and query QuickBooks Item Line Items.
ItemReceiptExpenseItems	Create, update, delete, and query QuickBooks Item Receipt Expense Line Items.
ItemReceiptLineItems	Create, update, delete, and query QuickBooks Item Receipt Line Items.
ItemReceipts	Create, update, delete, and query QuickBooks Item Receipts.
Items	Create, update, delete, and query QuickBooks Items.
JobTypes	Create and query QuickBooks JobTypes.
JournalEntries	Create, update, delete, and query QuickBooks

Name	Description
	Journal Entries. Note that while Journal Entry Lines can be created with a new Journal Entry, they cannot be added or removed from an existing Journal Entry.
JournalEntryLines	Create, update, delete, and query QuickBooks Journal Entries. Note that while Journal Entry Lines can be created with a new Journal Entry, they cannot be added or removed from an existing Journal Entry.
OtherNames	Create, update, delete, and query QuickBooks Other Name entities.
PaymentMethods	Create, update, delete, and query QuickBooks Payment Methods.
PayrollNonWageItems	Query QuickBooks Non-Wage Payroll Items.
PayrollWageItems	Create and query QuickBooks Wage Payroll Items.
PriceLevelPerItem	Create and query QuickBooks Price Levels Per Item. Only QuickBooks Premier and Enterprise support Per-Item Price Levels. Note that while Price Levels can be added from this table, you may only add Per-Item Price Levels from this table. Price Levels may be deleted from the PriceLevels table. This table requires QBXML version 4.0 or later.
PriceLevels	Create, delete, and query QuickBooks Price

Name	Description
	Levels. Note that while Price Levels can be added and deleted from this table, you may add only fixed-percentage Price Levels from this table. Per-Item Price Levels may be added via the PriceLevelPerItem table. This table requires QBXML version 4.0 or higher.
PurchaseOrderLineItems	Create, update, delete, and query QuickBooks Purchase Order Line Items.
PurchaseOrders	Create, update, delete, and query QuickBooks Purchase Orders.
ReceivePayments	Create, update, delete, and query QuickBooks Receive Payment transactions. QBXMLVersion must be set to 6.0 or higher to update a ReceivePayment.
ReceivePaymentsAppliedTo	Create, update, and query QuickBooks Receive Payment AppliedTo aggregates. In a Receive Payment, each AppliedTo aggregate represents the transaction to which this part of the payment is being applied. QBXMLVersion must be set to 6.0 or higher to update a ReceivePayment.
SalesOrderLineItems	Create, update, delete, and query QuickBooks Sales Order Line Items.
SalesOrders	Create, update, delete, and query QuickBooks Sales Orders.

Name	Description
SalesReceiptLineItems	Create, update, delete, and query QuickBooks Sales Receipt Line Items.
SalesReceipts	Create, update, delete, and query QuickBooks Sales Receipts.
SalesReps	Create, update, delete, and query QuickBooks Sales Rep entities.
SalesTaxCodes	Create, update, delete, and query QuickBooks Sales Tax Codes.
SalesTaxItems	Create, update, delete, and query QuickBooks Sales Tax Items.
ShippingMethods	Create, update, delete, and query QuickBooks Shipping Methods.
StandardTerms	Create, update, delete, and query QuickBooks Standard Terms.
StatementCharges	Create, update, delete, and query QuickBooks Statement Charges.
TimeTracking	Create, update, delete, and query QuickBooks Time Tracking events.
ToDo	Create, update, delete, and query QuickBooks To Do entries.
TransferInventory	Query and delete QuickBooks Transfer Inventory transactions. Transfer Inventory is available in only QuickBooks Enterprise 2010 and above, and only with the Advanced

Name	Description
	Inventory add-on.
TransferInventoryLineItems	Create and query QuickBooks Transfer Inventory Line Items. Transfer Inventory is available in only QuickBooks Enterprise 2010 and above, and only with the Advanced Inventory add-on.
Transfers	Create, update, and query QuickBooks transfers. Requires QBXML version 12.0 or higher.
UnitOfMeasure	Create, update, delete, and query QuickBooks units of measure. QuickBooks requires QBXML version 7.0 or higher to use this table.
UnitOfMeasureDefaultUnits	Create and query QuickBooks unit-of-measure default units. QuickBooks requires QBXML version 7.0 or higher to use this table.
UnitOfMeasureRelatedUnits	Create and query QuickBooks unit-of-measure related units. QuickBooks requires QBXML version 7.0 or higher to use this table.
VehicleMileage	Create, update, delete, and query QuickBooks Vehicle Mileage entities. QuickBooks requires QBXML version 6.0 or higher to use this table.
VendorCreditExpenseItems	Create, update, delete, and query QuickBooks Vendor Credit Expense Line Items.
VendorCreditLineItems	Create, update, delete, and query QuickBooks Vendor Credit Line Items.

Name	Description
VendorCredits	Create, update, delete, and query QuickBooks Vendor Credits.
Vendors	Create, update, delete, and query QuickBooks Vendors.
VendorTypes	Create, update, delete, and query QuickBooks Vendor Types.
WorkersCompCodeLines	Query QuickBooks Workers Comp Code entries by line. Requires QBXML Version 7.0 or higher.
WorkersCompCodes	Query QuickBooks Workers Comp Code entries. Requires QBXML Version 7.0 or higher.

## Available Views

Name	Description
BillLinkedTransactions	Query QuickBooks Bill Linked Transactions.
CompanyInfo	Query the company information from QuickBooks.
CreditMemoLinkedTransactions	Query QuickBooks Credit Memo Linked Transactions.
DeletedEntities	Query deleted Entities.
DeletedTransactions	Query deleted Transactions.
EstimateLinkedTransactions	Query QuickBooks Estimate Linked transactions.
Host	Query the QuickBooks host process. The

Name	Description
	Host represents information about the QuickBooks process currently being executed.
InvoiceLinkedTransactions	Query QuickBooks Invoice Linked Transactions.
ItemReceiptLinkedTransactions	Query QuickBooks Item Receipt Linked Transactions.
ItemSites	Create, update, delete, and query QuickBooks Item Sites. Item sites are only available in QuickBooks Enterprise 2010 and above, and only with the Advanced Inventory add-on. This table requires a minimum of QBXML version 10.0.
Preferences	Query information about many of the preferences the QuickBooks user has set in the company file.
PurchaseOrderLinkedTransactions	Query QuickBooks Purchase Order Linked Transactions.
SalesOrderLinkedTransactions	Query QuickBooks Sales Order Linked Transactions.
StatementChargeLinkedTransactions	Query QuickBooks Statement Charge Linked Transactions.
Templates	Query QuickBooks templates.
Transactions	Query QuickBooks transactions. You may search the Transactions using a number of

<b>Name</b>	<b>Description</b>
	values including Type, Entity, Account, ReferenceNumber, Item, Class, Date, and TimeModified.
VendorCreditLinkedTransactions	Query QuickBooks Vendor Credit Linked Transactions.

# Connecting to QuickBooks Online

The Sisense QuickBooks Online connector is a standalone connector that allows you to import data from the QuickBooks Online API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect to Quickbooks Online API through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the QuickBooks Online API. To obtain a connection string, you will need to create a QuickBooks Online app.

Once you have connected to QuickBooks Online, you can import a variety of tables from the QuickBooks Online API.

This page describes how to install the QuickBooks Online connector, how to connect to QuickBooks Online with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the QuickBooks Online Connector](#)
- [Connecting to the QuickBooks Online REST API](#)
- [Adding QuickBooks Online Tables to your ElastiCube](#)
- [QuickBooks Online Tables](#)
- [Versions](#)

## Installing the Quickbooks Online Connector

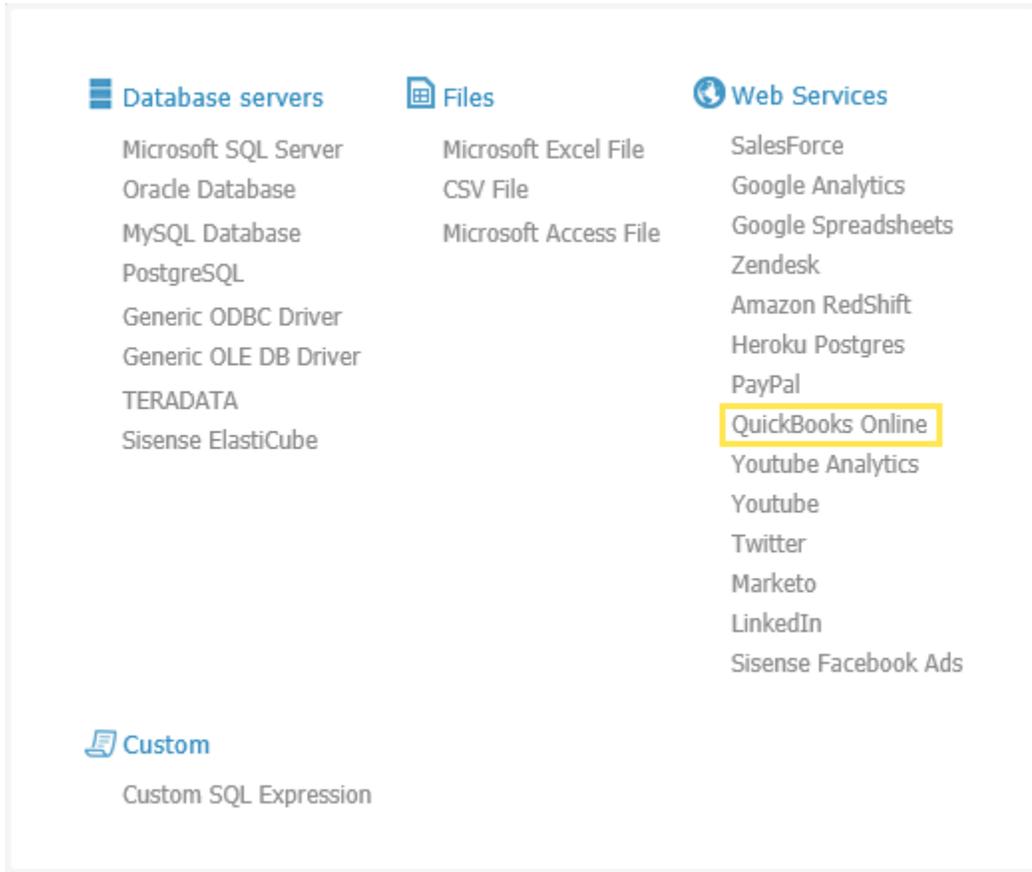
Sisense provides the Quickbooks Online connector as a standalone connector that you can download and add to your list of default Sisense connectors.

**To install the QuickBooks Online connector:**

1. [Download](#) the Quickbooks Online installation file.
2. Open the installation file and click **Install**.

3. After the installation process is complete, click **Close**.

The Quickbooks Online connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.



## Connecting to the Quickbooks Online REST API

To access Quickbooks Online's REST API from Sisense, you must provide valid Oauth Quickbooks Online credentials. These credentials are provided by Quickbooks Online when you create an application.

## Creating an App

You can follow the steps below to create an app and obtain the OAuth client credentials, the consumer key and consumer secret. You will need to create an Intuit developer account if you do not already have one.

1. Log in to <http://developer.intuit.com>.
2. Click **My Apps > Create New App > Select APIs**. Select the **Accounting API**.
3. On the Settings tab, define the Launch URL. If you are making a desktop application, set Launch URL to <http://localhost/>. For a Web application, set Launch URL to a page you would like the user to be returned to after they have granted your application permissions.

Once you have created the app, you will be redirected to a page with information about your app. The Development tab contains your Sandbox credentials and the Production tab contains credentials for use with a production account. The consumer key and consumer secret are displayed on the Keys tab.

## Adding Quickbooks Online Tables to your ElastiCube

Sisense uses connection strings to connect to Quickbooks Online and import data into the ElastiCube Manager. Each connection string contains a authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

The following parameters are mandatory for Quickbooks Online:

- **OAuthClientId**: The client ID assigned when you register your application with an OAuth authorization server.
- **OAuthClientSecret**: Set this to the App Secret in your app settings.
- **CallbackURL**: The OAuth callback URL to return to when authenticating. This value must match the callback URL you specify in your app settings. Set this to <http://localhost>.

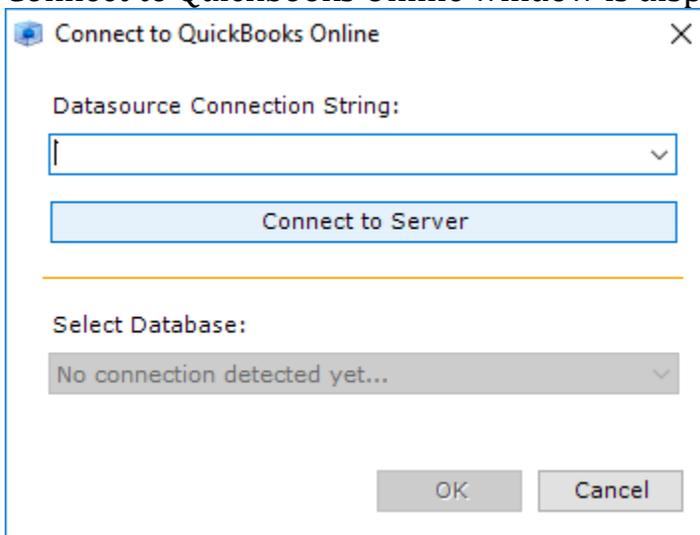
- **InitiateOAuth:** Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken connection property.
- **Timeout:** The value in seconds until the timeout error is thrown, cancelling the operation. Set to 0 to disable time outs.
- **CompanyId:** Set this to the ID of the company you want to connect to.

The following is an example of a Quickbooks Online connection string:

```
jdbc:quickbooksonline:OAuthClientId=xxxxxxxxxxxx;OAuthClientSecret=xxxxxxxxxxxx;CompanyId=xxxxxxxxxxxx;CallbackURL=http://localhost/;InitiateOAuth=GETANDREFRESH;Timeout=0;
```

### To add Quickbooks Online data:

1. In ElastiCube Manager, click **Add Data** and then, **Quickbooks Online**. The Connect to Quickbooks Online window is displayed.



2. In **Datasource Connection**

**String**, enter your connection string.

3. Click **Connect to Server**. Quickbooks Online is displayed in the Select Database list.
4. Click **OK**. Sisense connects to Quickbooks Online and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**.

The tables are displayed in the ElastiCube Manager.

# Quickbooks Online Tables

The Sisense Quickbooks Online connector allows you to import the following tables into the ElastiCube Manager.

Name	Description
Accounts	Query QuickBooks Accounts.
Attachables	Attachables.
BillLineItems	Query QuickBooks Bill Line Items.
BillPaymentLineItems	Query QuickBooks Bill Payment Line Items.
BillPayments	Query QuickBooks Bill Payments.
Bills	Query QuickBooks Bills.
BudgetDetails	Query QuickBooks BudgetDetails.
Budgets	Query QuickBooks Budgets.
Class	Query QuickBooks Classes.
CompanyInfo	Retrieve information about the QuickBooks company.
CreditMemoLineItems	Query QuickBooks Credit Memo Line Items.
CreditMemos	Query QuickBooks Credit Memos.
Customers	Query QuickBooks Customers.
Departments	Query QuickBooks Departments.
DepositLineItems	Query QuickBooks Deposit Line Items.
Deposits	Query QuickBooks Deposits.
Employees	Query QuickBooks Employees.
Entitlements	Retrieves QuickBooks Entitlements.

Name	Description
EstimateLineItems	Query QuickBooks Estimate Line Items.
EstimateLinkedTransactions	Query QuickBooks Estimate Line Items.
Estimates	Query QuickBooks Estimates.
InvoiceLineItems	Query QuickBooks Invoice Line Items.
InvoiceLinkedTransactions	Query QuickBooks Invoice Line Items.
Invoices	Query QuickBooks Invoices.
Items	Query QuickBooks Items.
JournalEntries	Query QuickBooks Journal Entries.
JournalEntryLineItems	Query QuickBooks Journal Entry Line Items.
PaymentLineItems	Query QuickBooks Payment Line Items.
PaymentMethods	Query QuickBooks Payment Methods.
Payments	Query QuickBooks Payments.
Preferences	Query QuickBooks Preferences. The Preferences table contains settings for company-wide preferences, which affect all users.
PurchaseLineItems	Query QuickBooks Purchase Line Items.
PurchaseOrderLineItems	Query QuickBooks Purchase Order Line Items.
PurchaseOrders	Query QuickBooks Purchase Orders.
Purchases	Query QuickBooks Purchases.
RefundReceiptLineItems	Query QuickBooks RefundReceiptLineItems.
RefundReceipts	Query QuickBooks RefundReceipts.
SalesReceiptLineItems	Query QuickBooks Sales Receipt Line Items.

Name	Description
SalesReceipts	Query QuickBooks Sales Receipts.
TaxCodes	Query QuickBooks Sales Tax Codes.
TaxRates	Query QuickBooks Tax Rates.
Terms	Query QuickBooks Terms.
TimeActivities	Query QuickBooks Time Activities.
Transfers	Query QuickBooks Transfers
VendorCredits	Query QuickBooks Vendor Credits.
Vendors	Query QuickBooks Vendors.

## Versions

Version	Release Date	Improvements
6.6.0.44	March 16, 2017	General connector enhancements
6.5.1.23	March 09, 2017	Minor bug fixes
6.5.1.20	January 16, 2017	General release

# Amazon Redshift Live

This topic describes how you can add live connections to your Redshift databases from the Sisense Web Application.

For information about live connections, see [Sisense Live Connect](#).

## Adding Amazon Redshift Live Connections

**To add a live connection:**

1. In the Sisense Web Application, select **Admin > Data Sources**. The Data Sources page is displayed.

The screenshot shows the Sisense Admin console interface. The top navigation bar includes 'Home', 'Pulse', 'ElastiCube Manager', and 'Admin'. The left sidebar contains navigation options: 'Admin', 'Users', 'Groups', 'ElastiCubes & Data Sets' (highlighted), 'Active Directory', 'Single Sign On', 'Trusted Servers List', 'System Configuration', and 'REST API'. The main content area is titled 'ElastiCubes & Data Sets' and contains two expandable sections.

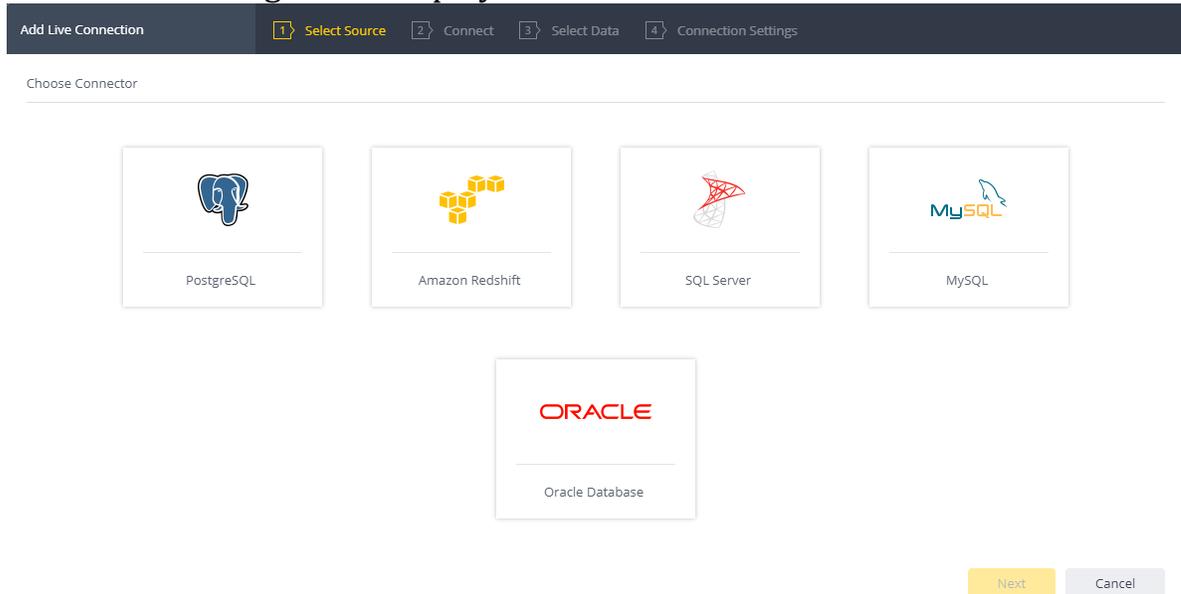
**Local Machine**

ElastiCube	Status	Created On	Last Built	Size On Disk	
Nortwind BusinessDays 2	Stopped	5/31/15 5:02 PM	10/24/16 2:51 PM	4.72MB	Open
Sample eCommerce	Running	2/10/14 2:21 PM	3/12/14 10:18 AM	59.05MB	Open
Sample Healthcare	Running	7/2/13 1:07 AM	3/12/14 2:01 PM	1.5MB	Open
Sample Lead Generation	Running	3/3/14 7:46 PM	3/12/14 11:59 AM	2.23MB	Open
Test	Stopped	7/18/16 1:06 PM	1/24/17 3:25 PM	30.49MB	Open

**Live Data Sets**

Data Set	Source	Server Location	User Name	Refresh Rate	
Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	
Copy of Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	
customers	MySql Database	192.168.1.229	Test	10 Sec.	
datatypes	MySql Database	192.168.1.229	Test	10 Sec.	
LiveSource	Microsoft Sql Server	192.168.1.229	Test	2 Min.	
MSSQL Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	
productlines	MySql Database	192.168.1.229	Test	10 Sec.	

2. In the Data Sources page, click **Add Live Connection**. The Add Live Connection dialog box is displayed.



3. In the Add Live Connection dialog box, select your live data source.
4. Under the Connect tab, enter the following credentials:
  - Location:** The IP address of your data source.
  - User Name:** The user name to access data source.
  - Password:** The password to access your data source.
5. (Optional) Select **Encrypt Connection** to configure the driver to encrypt all communication with the Amazon Redshift instance before sending it over the network, and then select **Trust Server Certificate**.
6. Click **Next**.
7. Under the Select Data tab, on the left side is a list of schemas located in your data source. Select the relevant schema. You can find the schema easily by searching for it through the search field at the top of the list to filter the schemas displayed. After selecting the schema, a list of the available tables and views are displayed.
8. Select a table or view. You can find the table easily by searching for it through the search field at the top of the list to filter the tables displayed.
9. Click **Next**.
10. Under the Connection Settings tab, set the following:
  - Connection Name:** Enter a name for your connection. This field defines the name of your data source displayed in the Sisense Web Application.
  - Refresh Rate:** Enter or select the refresh rate for your data source. The shortest refresh rate you can set is 10 seconds. The refresh rate is the time

in between queries to the data source. If the previous query has not been completed, the next query is not sent until the previous refresh is complete unless the refresh rate is greater than the time it took to complete the refresh since the last query. For example, if your refresh rate is 10 seconds, and your query takes 15 seconds, the next query is sent when the previous query is complete as 10 seconds has passed and the previous query is complete. If your refresh rate is 30 seconds, and the previous query is complete, the next refresh will be in 30 seconds from the last query.

**Time Out:** Enter or select the Time Out rate for your data source. The Time Out rate determines how long Sisense waits before cancelling a query that does not receive a response from your data source.

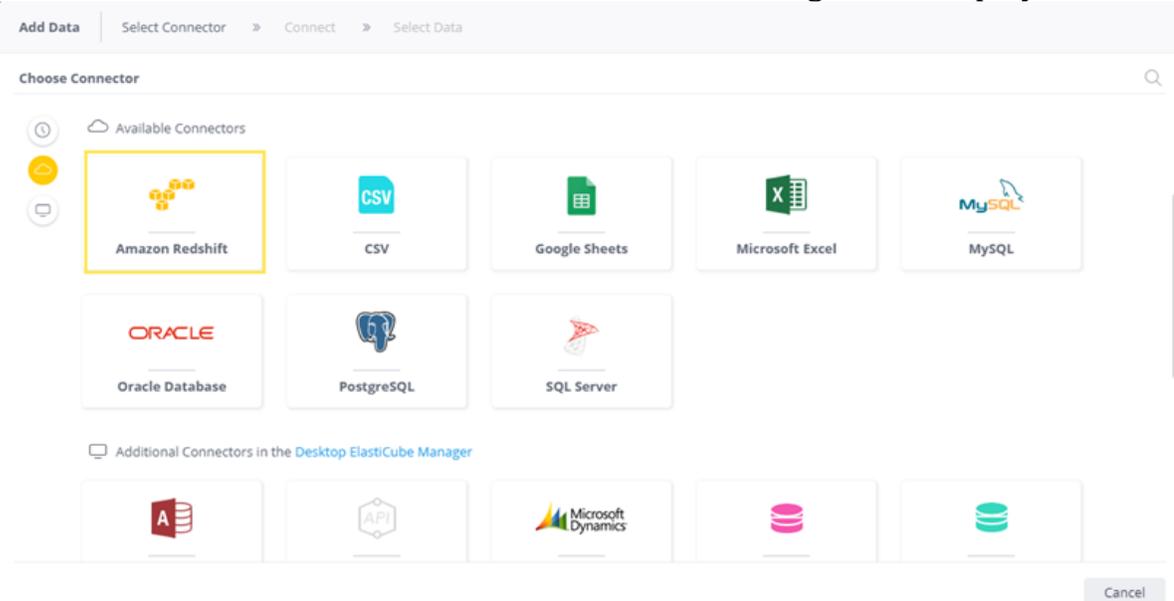
11. Click **Done**. The connection is added as a data set to your Data Sources page.

# Connecting to Redshift

The ElastiCube Manager enables easy and quick access to databases, tables and views contained within Amazon RedShift databases. The steps below detail how to connect to this type of data source.

## To import data from Amazon Redshift:

1. In the Data page of ElastiCube Manager, open an ElastiCube or create a new ElastiCube.
2. In the ElastiCube, click . The Add Data dialog box is displayed.



3. Click **Amazon Redshift** to open the Amazon Redshift settings.

Add Data | Select Connector » Connect » Select Data

 Amazon Redshift

Location

User name

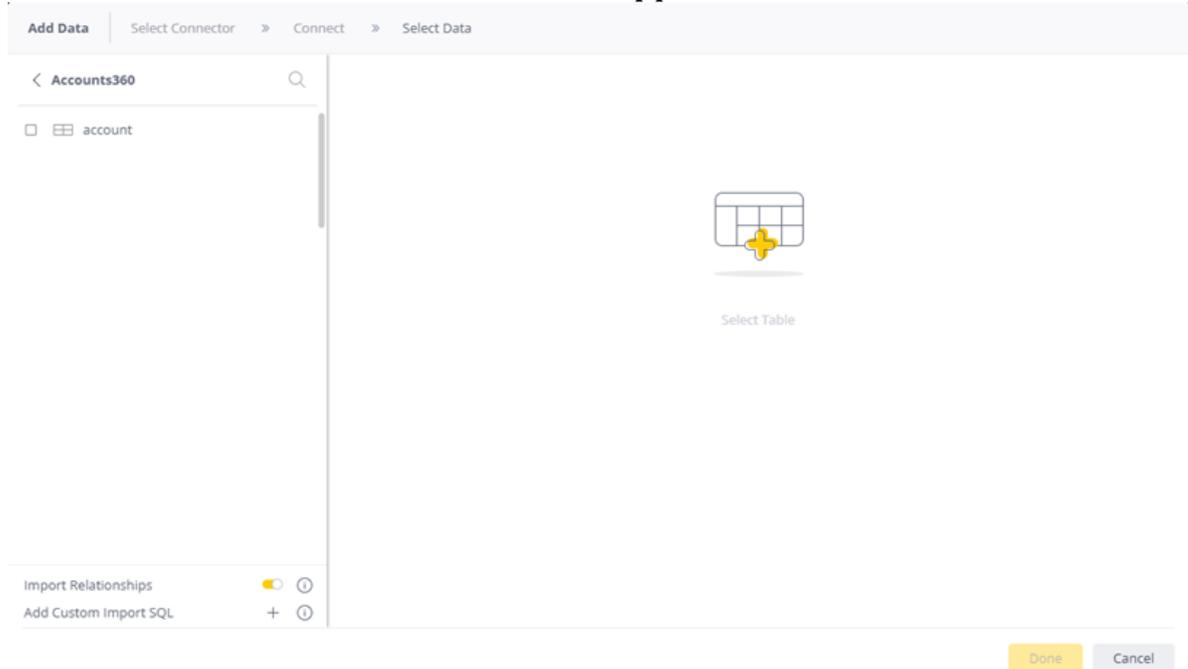
Password

Encrypt Connection  
 Trust Server Certificate

Next Cancel

4. Enter the following details:
  - **Location:** Enter the computer/server IP address of the database. This server address should include the port separated by a colon. By default, the port is 5439, for example: [IP Address]:5439. To connect to a database running on your own computer, enter **localhost**.
  - If Windows Authentication is configured with the database in the **User Name** and **Password** fields, enter your database credentials.
5. **(Optional)** Select **Encrypt Connection** to configure the driver to encrypt all communication with Redshift before sending it over the network, and then select **Trust Server Certificate**.

6. Click **Next**. A list of tables in the database are displayed. All tables and views associated with the database will appear in a new window.



7. From the Tables list, select the relevant table or view you want to work with. You can click  next to the relevant table or view to see a preview of the data inside it. When you select the table or view, two new options are displayed at the bottom of the list, **Import Relationships** and **Add Custom Import SQL**.
8. (Optional) By default, existing relationships between tables are automatically replicated in the ElastiCube. You can disable this by toggling the **Import Relationships** switch.
9. (Optional) Click + to customize the data you want to import with SQL. See [Importing Data with Custom Queries](#) for more information.
10. After you have selected all the relevant tables, click **Done**. The tables are added to your schema in the ElastiCube Manager.

**Note:** If you encounter problems connecting to Amazon RedShift, please see [this article](#) in our support forums.

# Connecting to Redshift

**IMPORTANT:** On October 23, 2017, Amazon Redshift is replacing the SSL certificates on your cluster with AWS Certificate Manager (ACM) issued certificates. ACM is a trusted public certificate authority (CA) that is trusted by most current systems. In order to continue importing data into the ElastiCube Manager with the Sisense Redshift connector, you must apply the new ACM certificate. Amazon provides a Certificate Authority, which needs to be saved on a Windows machine.

## To replace your certificate:

1. Download the [Redshift Certificate Authority Bundle](#).
2. Place the certificates from the bundle in your root.crt file. On Microsoft Windows, the file is %APPDATA%\postgresql\root.crt

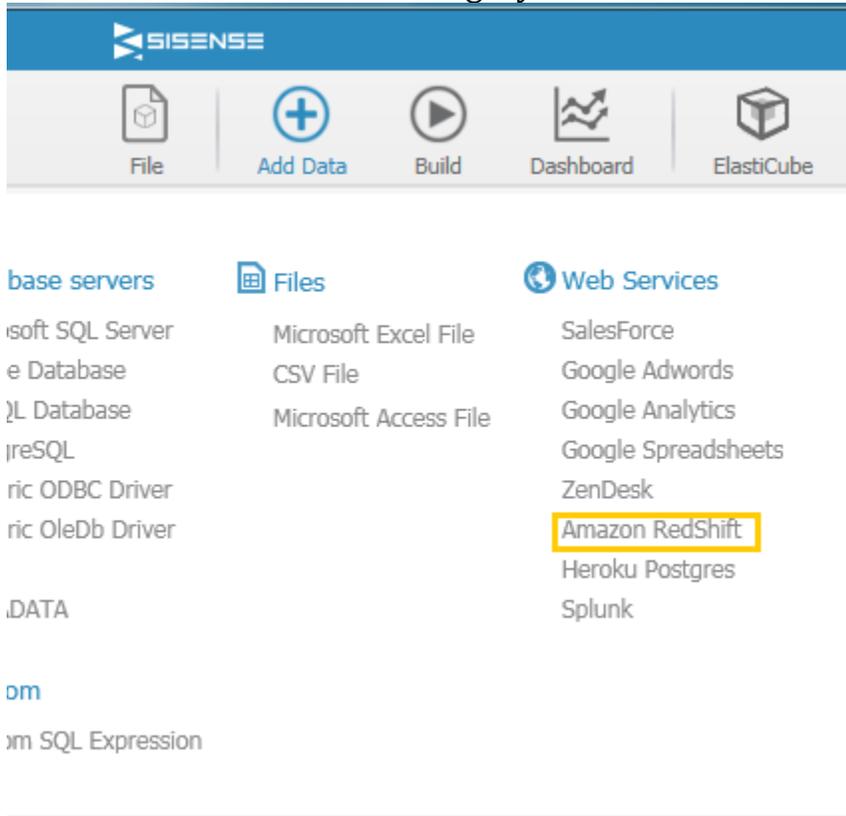
For more information, see Transitioning to [ACM Certificates for SSL Connections](#).

## Introduction

The ElastiCube Manager enables easy and quick access to databases, tables and views contained within Amazon RedShift databases. The steps below detail how to connect to this type of data source.

1. Click **Add Data** in the top menu of the ElastiCube Manager.

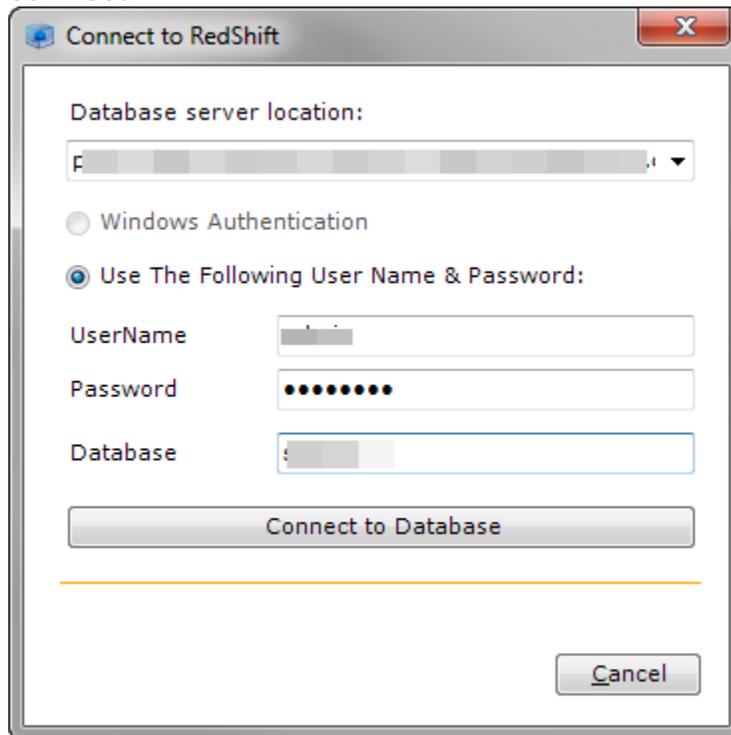
2. Under the **Web Services** category select **Amazon Redshift**.



3. You will be prompted to enter the following information:

- **Database server location:** Enter the computer/server IP address that holds the database. This server address should include the port separated by a colon. By default, the port is 5439, for example: [IP Address]:5439.
- Either use your **Windows Authentication** if configured with the database, or alternatively enter the **Username** and **Password** to connect to the database.

- **Database:** Enter the name of the database to which you want to connect.



4. Click **Connect to Database**.

All tables and views associated with the database will appear in a new window.

To view a preview of data contained in a particular table, highlight the table or view in the list and in click the preview pane below. To preview the table, select the **Preview** checkbox.

Enable the checkbox next to each table or view you would like to use.

Existing relationships between tables can be automatically replicated in the ElastiCube by selecting the **Automatically create relationships from database** checkbox. Likewise fields with similar names can be linked by selecting the **Automatically create relationships for fields with the same name** checkbox.

**Note:** If you encounter problems connecting to Amazon RedShift, please see this article in our support forums.

5. Once all relevant tables are selected click **Add**.

# Connecting to Salesforce

The ElastiCube Manager enables easy and quick access to databases, tables and views contained within Salesforce. The steps below describe how to connect to this type of data source.

If you need to connect to your Salesforce Sandbox environment, read this article in our support forum.

**Note:** Sisense only supports Salesforce Enterprise and above.

## Generating a Security Token

Before getting started, you will need a Salesforce Security Token.

**To generate a token:**

1. Open Salesforce and click **Setup>'My Personal Information'>'Reset Security Token'**.
2. Click the button labeled **'Reset Security Token'**. The token will be sent to your email address.

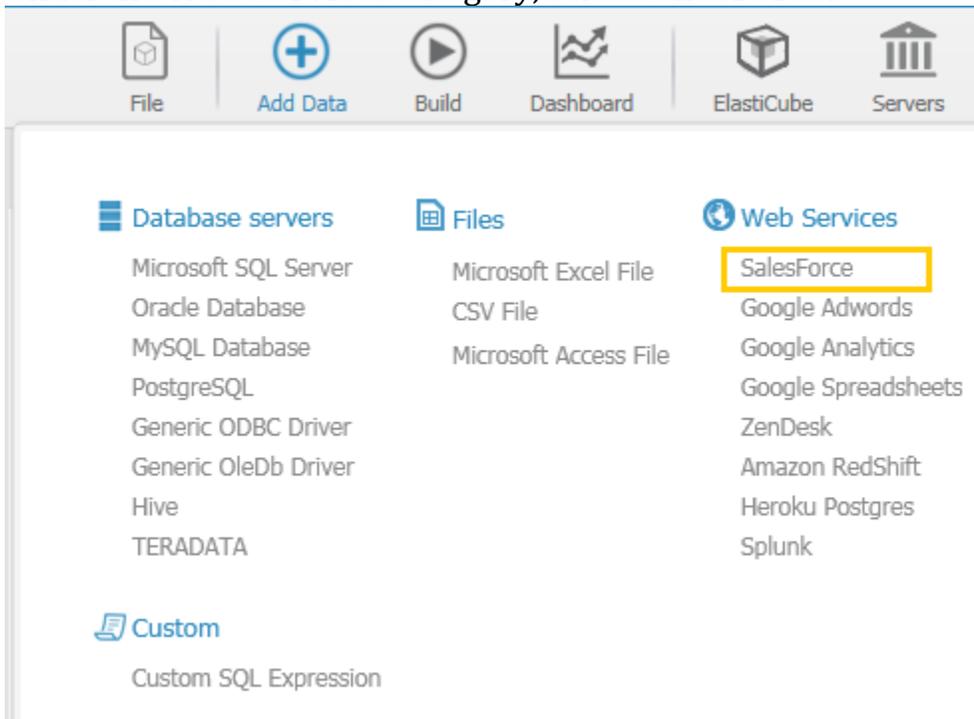
## Importing Salesforce Data

SalesForce can be customized in various ways, therefore its recommend to first review your deployment, and understand in which entity (account, contact etc.) and how the required data is stored. Once understood, connect to Salesforce and import a few tables to review the data. For a general outline of Salesforce entities see [SalesForce Schema](#) and [Table Description](#).

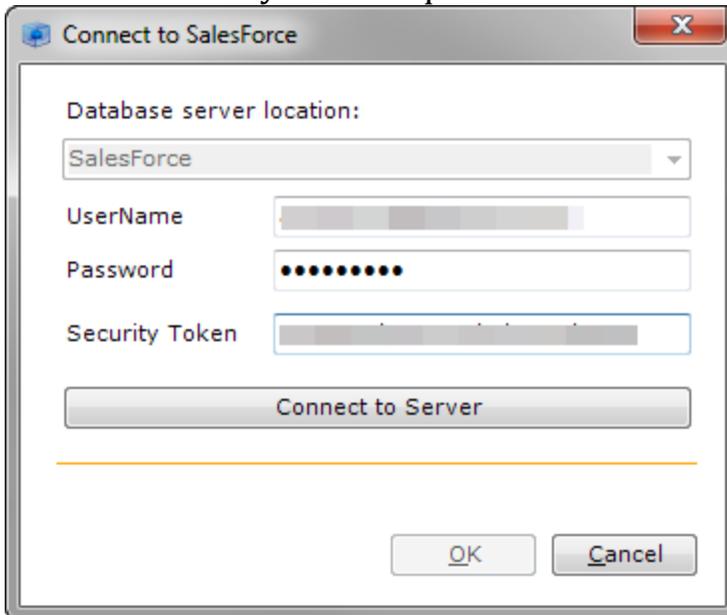
**To import Salesforce data:**

1. Click **Add Data** in the top menu of the ElastiCube Manager.

2. Under the **Web Services** category, select **SalesForce**.

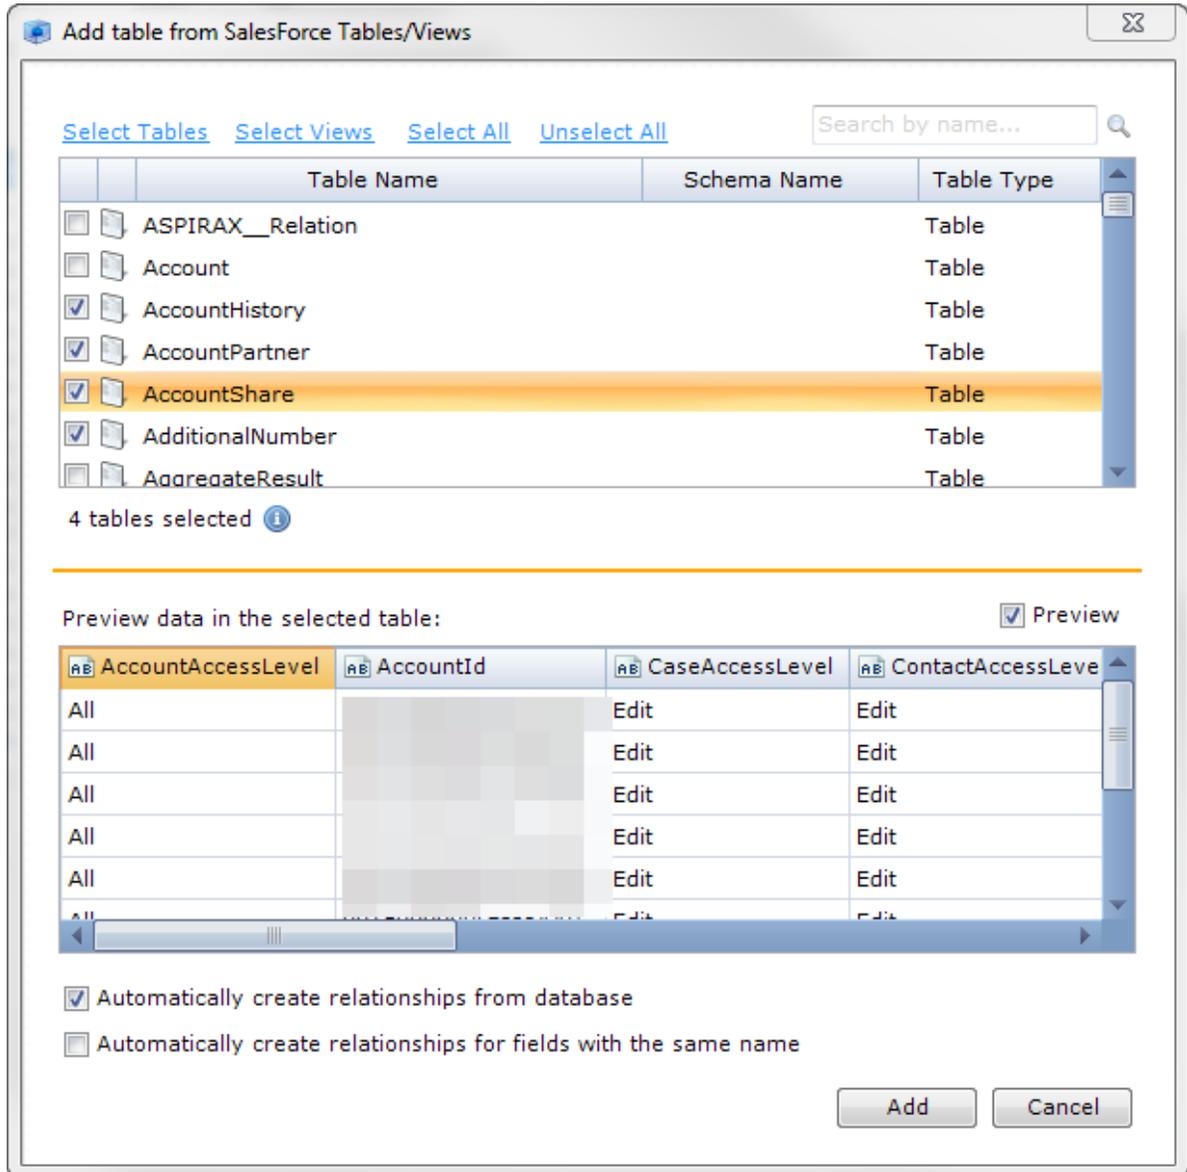


3. Enter the username and password to connect to SalesForce.
4. Enter the security token requested from SalesForce (see note above).



5. Click **Connect to Server**, then click **OK**.  
All tables and views associated with SalesForce will appear in a new window.  
To preview the data contained in a particular table, highlight the table or view in the list and click the preview pane below. To preview the table, select the Preview checkbox.

6. Select the checkboxes next to each table or view you want to use. Existing relationships between tables can be automatically replicated in the ElastiCube by selecting the **Automatically create relationships from database** checkbox. Likewise fields with similar names can be linked by selecting the **Automatically create relationships for fields with the same name** checkbox.



7. Once all relevant tables are selected, click **Add**.

For troubleshooting Salesforce connections, see our [forum post](#).

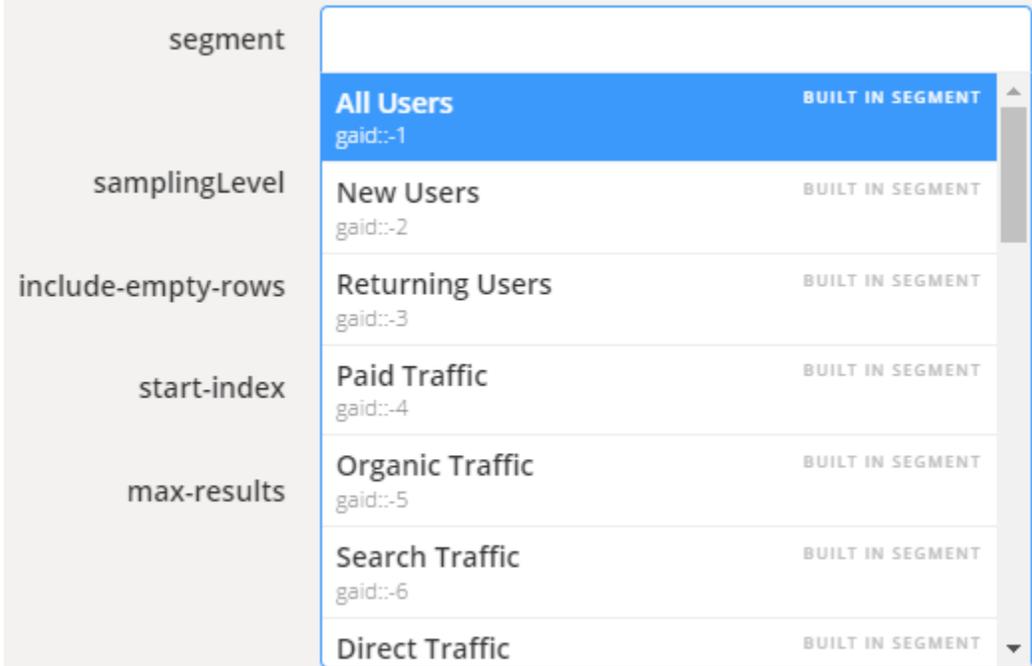
# Segmenting Google Analytics Data

A segment is a subset of your Analytics data. For example, of your entire set of users, one segment might be users from a particular country or city.

You can have predefined segments created by Google or custom segments created by you. You can segment your data when querying Google Analytics. This is useful to prevent Google from sampling your data or return data associated with specific segments.

## To request data associated from a segment:

1. Open the [Google Analytics Query Explorer](#).
2. In the segment field, enter the name of your segment. As you begin to type, all the relevant segments are displayed. Below the segment name is Google Analytics ID.

3. 

the Google Analytics ID (gaid::X).

4. Open the Google Analytics tables.xml file saved in  
%ProgramData%\Sisense\PrismServer\LocalRepository2.0  
\Resources\Google\Analytics
5. In the XML file, paste the following code under the relevant table to be segmented. For Example:

```
<table name="Traffic, Daily Report"  
accumulate_by="ga:date">  
<segment>  
  <![CDATA[gaid::-4]]>  
</segment>  
<fields>  
  <field id="Site"></field>  
  <field id="ga:date"></field>  
  <field id="ga:bounces"></field>  
  <field id="ga:entrances"></field>  
  <field id="ga:exits"></field>  
  <field id="ga:newUsers"></field>  
  <field id="ga:sessionDuration"></field>  
  <field id="ga:users"></field>  
  <field id="ga:sessions"></field>  
</fields>  
</table>
```

6. Save your changes to the XML file. The next time you return data, the segmented data is returned.

## Filtering Google Analytics Data

Sisense allows you to filter your data returned from Google Analytics according to View-level filters. For example, you can use filters to exclude traffic from particular IP addresses, include only data from specific subdomains.

By using a filter in Sisense you increase the number of dimensions you can import to a table as a workaround to Google Analytics 7 dimension limitation for any query.

Filtering your incoming hits permanently includes, excludes, or alters those hits in that view, according to the type of filter. Therefore, you should ALWAYS maintain an unfiltered view of your data so you always have access to your full data set.

**To apply filters to your query:**

1. Open the Google Analytics tables.xml file saved in %ProgramData%\Sisense\PrismServer\LocalRepository2.0\Resources\Google\Analytics
2. In the XML file, paste the following code under the relevant table to be filtered. For Example:

```
<table name="Traffic, Daily Report"
accumulate_by="ga:date">
<filter>
<![CDATA[ga:userType!=New Visitor]]>
</filter>
<fields>
<field id="Site"></field>
<field id="ga:date"></field>
<field id="ga:bounces"></field>
<field id="ga:entrances"></field>
<field id="ga:exits"></field>
<field id="ga:newUsers"></field>
<field id="ga:sessionDuration"></field>
<field id="ga:users"></field>
<field id="ga:sessions"></field>
</fields>
</table>
```

3. Save your changes to the XML file. The next time you return data, the filtered data is returned.

## Sorting Google Analytics Data

You can define the sort order and the direction of data returned from Google Analytics according to A list of metrics and dimensions indicating the sorting order and sorting direction for the returned data.

- Sorting order is specified by the left to right order of the metrics and dimensions listed.
- Sorting direction defaults to ascending and can be changed to descending by using a minus sign (-) prefix on the requested field.

For example, “What are my top countries, and which browsers do they use most?”

```
<sort>
  ga:country,ga:browser
</sort>
```

This example sorts the data by country and browser in ascending order. To return the data in descending order, add the prefix “-” as shown below:

```
<sort>
  -ga:country,-ga:browser
</sort>
```

### To sort your Google Analytics data:

1. Open the Google Analytics tables.xml file saved in  
%ProgramData%\Sisense\PrismServer\LocalRepository2.0  
\Resources\Google\Analytics
2. In the XML file, paste the following code under the relevant table to be sorted. For Example:

```
<table name="Traffic, Daily Report"
  accumulate_by="ga:date">
  <sort>
    -ga:sessions,-ga:date
```

```
</sort>
<fields>
  <field id="Site"></field>
  <field id="ga:date"></field>
  <field id="ga:entrances"></field>
  <field id="ga:exits"></field>
  <field id="ga:newUsers"></field>
  <field id="ga:pageviews"></field>
  <field id="ga:users"></field>
  <field id="ga:sessions"></field>
</fields>
</table>
```

3. Save your changes to the XML file. The next time you return data, the sorted data is returned.

# Connecting to ServiceNow

The Sisense ServiceNow connector is a standalone connector that allows you to import data from ServiceNow's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the ServiceNow APIs. To obtain a connection string, you will need to register a ServiceNow app.

Once you have connected to ServiceNow, you can import a variety of tables from the ServiceNow API.

This page describes how to install the ServiceNow connector, how to connect to ServiceNow with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the ServiceNow Connector](#)
- [Connecting to the ServiceNow REST API](#)
- [Adding ServiceNow Tables to your ElastiCube](#)
- [ServiceNow Tables](#)
- [Versions](#)

## Installing the ServiceNow Connector

Sisense provides the ServiceNow connector as a standalone connector that you can download and add to your list of default Sisense connectors.

### To install the ServiceNow connector:

1. [Download](#) the ServiceNow installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.  
The ServiceNow connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.

# Connecting to the ServiceNow REST API

To access ServiceNow's REST API from the ElastiCube Manager, you must provide valid OAuth ServiceNow credentials through a connection string. These credentials are provided by ServiceNow when you create an application. After you receive your credentials from ServiceNow, you can create the connection string and provide Sisense with it to connect to your data.

## Registering an App

Follow the steps below to obtain the OAuth client credentials, the OAuthClientId and OAuthClientSecret:

1. In the Navigator, click **System OAuth > Application Registry**.
2. Click **New** and then click **Create** an OAuth API Endpoint for External Clients.
3. Select the **Active** check box and enter the details for your app.

## Adding ServiceNow Tables to your ElastiCube

Sisense uses connection strings to connect to ServiceNow and import data into the ElastiCube Manager. Each connection string contains authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

The connection string to ServiceNow is in the following format:

```
jdbc:servicenow:APIKey=xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx;
```

For example:

```
jdbc:servicenow:InitiateOAuth=GETANDREFRESH;Instance=xxxx  
xxxxxx;OAuthClientId=xxxxxxxxxxxxxxxx;  
OAuthClientSecret=xxxxxxxxxx;Username=xxxxxxxxxx;Password=  
xxxxxxxxxxxxxxxx;
```

Your ServiceNow connection string should include the following parameters, separated by a semicolon ';':

- **InitiateOAuth:** Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken connection property.
- **OAuthClientId:** Set this to the Client Id in your app settings.
- **OAuthClientSecret:** Set this to the Client Secret in your app settings.
- **Password:** Set this to your password.
- **Username:** Set this to your username.
- **Instance:** Set this to your instance.

#### To add ServiceNow data:

1. In ElastiCube Manager, click **Add Data** and then, **ServiceNow**. The Connect to ServiceNow window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. ServiceNow is displayed in the Select Database list.
4. Click **OK**. Sisense connects to ServiceNow and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

## ServiceNow Tables

ServiceNow's RESTful APIs expose the following ServiceNow tables that you can import into the ElastiCube Manager through the Sisense ServiceNow connector:

### Available Tables

Name	Description
ast_contract	The ServiceNow table ast_contract.
ast_license_base	The ServiceNow table ast_license_base.
change_request	The ServiceNow table change_request.
cmdb_ci	The ServiceNow table cmdb_ci.

Name	Description
cmdb_metric	The ServiceNow table cmdb_metric.
cmn_building	The ServiceNow table cmn_building.
cmn_context_help	The ServiceNow table cmn_context_help.
cmn_cost_center	The ServiceNow table cmn_cost_center.
cmn_department	The ServiceNow table cmn_department.
cmn_location	The ServiceNow table cmn_location.
cmn_map_page	The ServiceNow table cmn_map_page.
cmn_notif_device	The ServiceNow table cmn_notif_device.
cmn_notif_device_variable	The ServiceNow table cmn_notif_device_variable.
cmn_notif_grmember	The ServiceNow table cmn_notif_group.
cmn_notif_group	The ServiceNow table cmn_notif_message.
cmn_notif_message	The ServiceNow table cmn_notif_message.
cmn_notif_service_provider	The ServiceNow table cmn_notif_service_provider.
cmn_other_schedule	The ServiceNow table cmn_other_schedule.
cmn_relative_duration	The ServiceNow table cmn_relative_duration.
cmn_schedule	The ServiceNow table cmn_schedule.
cmn_schedule_blackout	The ServiceNow table cmn_schedule_blackout.
cmn_schedule_condition	The ServiceNow table cmn_schedule_condition.
cmn_schedule_maintenance	The ServiceNow table cmn_schedule_maintenance.
cmn_schedule_page	The ServiceNow table cmn_schedule_page.
cmn_schedule_span	The ServiceNow table cmn_schedule_span.
cmn_timeline_page	The ServiceNow table cmn_timeline_page.

Name	Description
cmn_timeline_page_style	The ServiceNow table cmn_timeline_page_style.
cmn_timeline_sub_item	The ServiceNow table cmn_timeline_sub_item.
diagrammer_action	The ServiceNow table diagrammer_action.
expert_panel	The ServiceNow table expert_panel.
item_option_new	The ServiceNow table item_option_new.
question	The ServiceNow table question.
sc_category	The ServiceNow table sc_category.
sc_cat_item	The ServiceNow table sc_cat_item.
sla	The ServiceNow table sla.
sysauto	The ServiceNow table sysauto.
sysauto_script	The ServiceNow table sysauto_script.
syslog	The ServiceNow table syslog.
sysrule	The ServiceNow table sysrule.
system_db_object	The ServiceNow table system_db_object.
system_dictionary	The ServiceNow table system_dictionary.
system_documentation	The ServiceNow table system_documentation.
system_import_set_row	The ServiceNow table system_import_set_row.
system_script_client	The ServiceNow table system_script_client.
system_ui_policy	The ServiceNow table system_ui_policy.
system_ui_policy_action	The ServiceNow table system_ui_policy_action.
task	The ServiceNow table task.
v_field_creator	The ServiceNow table v_field_creator.

## Versions

<b>Version</b>	<b>Release Date</b>	<b>Improvements</b>
6.6.1.47	June 15, 2017	400 Bad Request error has been fixed
6.6.0.23	March 16, 2017	General Release
6.5.1.17	January 16, 2017	General Beta release

# Connecting to SharePoint

This connector is currently in beta and subject to change.

The Sisense SharePoint connector is a standalone connector that allows you to import data from SharePoint's into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string contains your credentials and is used by SharePoint to authenticate your account. Once you have connected to SharePoint, you can import a variety of tables from SharePoint.

This page describes how to install the SharePoint connector, how to connect to SharePoint with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the SharePoint Connector](#)
- [Creating A SharePoint Connection String](#)
- [Adding SharePoint Tables to your ElastiCube Manager](#)
- [SharePoint Tables](#)
- [Versions](#)

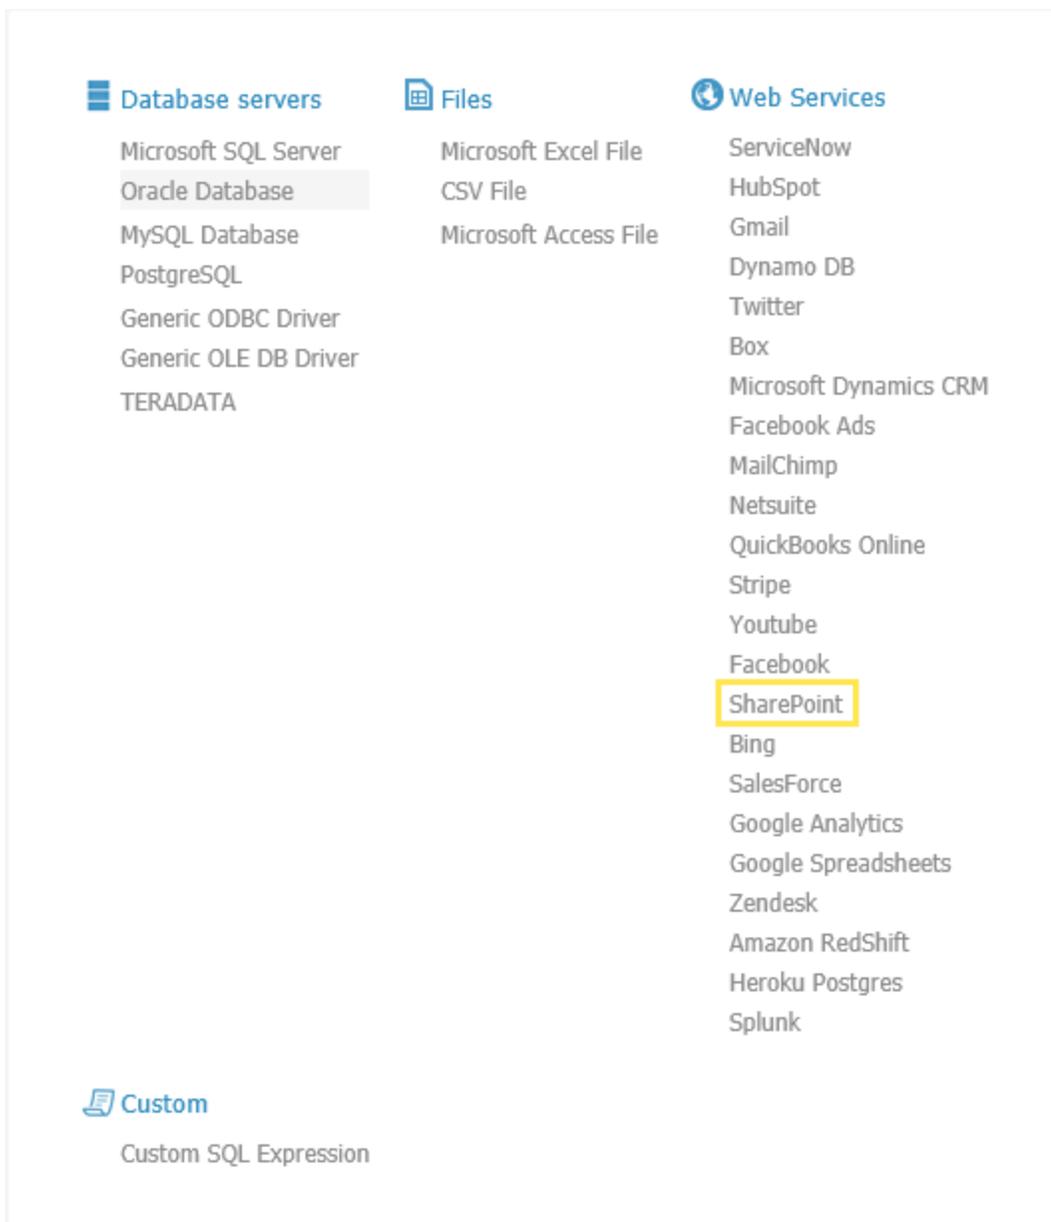
## Installing the SharePoint Connector

Sisense provides the SharePoint connector as a standalone connector that you can download and add to your list of default Sisense connectors.

**To install the SharePoint connector:**

1. [Download](#) the SharePoint installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The SharePoint connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.



## Creating a SharePoint Connection String

To connect to your SharePoint account's data, you provide your SharePoint credentials in a connection string. Sisense uses connection strings to connect to SharePoint and import data into the ElastiCube Manager. Each connection string contains a authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

The following is an example of a SharePoint connection string:

```
jdbc:sharepoint:User=MyUserAccount;Password=MyPassword;Auth Scheme=NTLM;Share Point Edition=SharePoint Online;URL=http://sharepointserver/mysite;
```

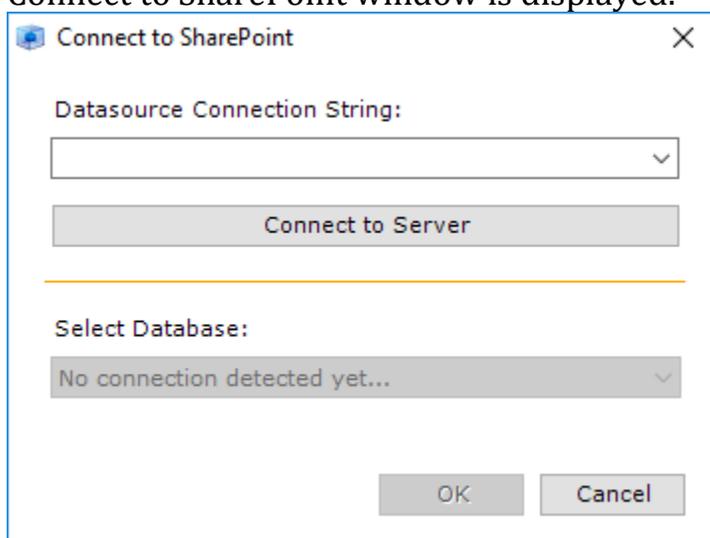
The connection string must start with “jdbc:sharepoint:” and include the following name-value pairs separated with semicolons:

- **User:** The SharePoint user account used to authenticate.
- **Password:** The password used to authenticate the user.
- **Share Point Edition:** The edition of SharePoint being used. Set to either SharePoint Online or SharePoint On-Premise. If you are using SharePoint Online, add the URL parameter. If you are using SharePoint On-Premise, add the Auth Scheme parameter.
- **URL:** The base URL for the site.
- **Auth Scheme:** The scheme used for authenticating to SharePoint On-Premise instances. Accepted entries are NTLM, BASIC, DIGEST, FORMS, NONE, NEGOTIATE, and KERBEROSDELEGATION. NTLM is the default.

## Adding SharePoint Tables to Your ElastiCube

**To import SharePoint tables:**

1. In the ElastiCube Manager, click **Add Data** and then, SharePoint. The Connect to SharePoint window is displayed.



2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. SharePoint is displayed in the Select Database list.

4. Click **OK**. Sisense connects to SharePoint and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

## SharePoint Tables

The Sisense SharePoint connector allows you to import the following tables into the ElastiCube Manager.

### Available Tables

Name	Description
Announcements	Create, update, delete, and query items in Announcement lists.
Attachments	Read or delete Attachments for the specified item on the specified list.
Calendar	Create, update, query, and delete items in SharePoint Calendar lists.
Contacts	Create, update, query, and delete items in SharePoint Contact lists.
Documents	Create, update, delete, and query Documents from SharePoint libraries.
FileVersionns	Lists the versions of files available on SharePoint.
GetValidTerms	Gets a list of valid terms for the specified column on the specified table.
Groups	Create, update, delete, and query Groups from SharePoint.
IssueTracking	Create, update, query, and delete items in SharePoint Issue Tracking lists.

Name	Description
Links	Create, update, query, and delete items in SharePoint Link lists.
Lists	Lists can be used to list the tables in SharePoint. This will only return actual lists in SharePoint and not any special tables associated with the driver.
Permissions	The permissions for a site or list.
Pictures	Create, update, delete, and query documents in a picture library.
Roles	Create, update, delete, and query Roles from SharePoint.
Subsites	This lists the available subsites.
Tasks	Create, update, query, and delete items in a SharePoint Tasks list.
Users	Create, update, delete, and query Users from SharePoint.
Views	Create, update, delete, and query the available lists in SharePoint.

## Versions

Version	Release Date	Improvements
6.6.1.2	May 23, 2017	General Beta release

# Connecting to Sisense ElastiCubes

ElastiCubes can also be connected to a data source using the Sisense ElastiCube connector or the Sisense ODBC, allowing you to import tables across different ElastiCubes. This can be useful for several use cases, such as:

- Separating custom SQL logic that is used for heavy data transformation from the main ElastiCube for easier management and shorter build times.
- Combining tables that are based on different refresh schedules. As every ElastiCube has its own build schedule, it is possible to import a table from an ElastiCube that doesn't refresh as often.

This page describes how to connect to your Sisense ElastiCubes through both methods:

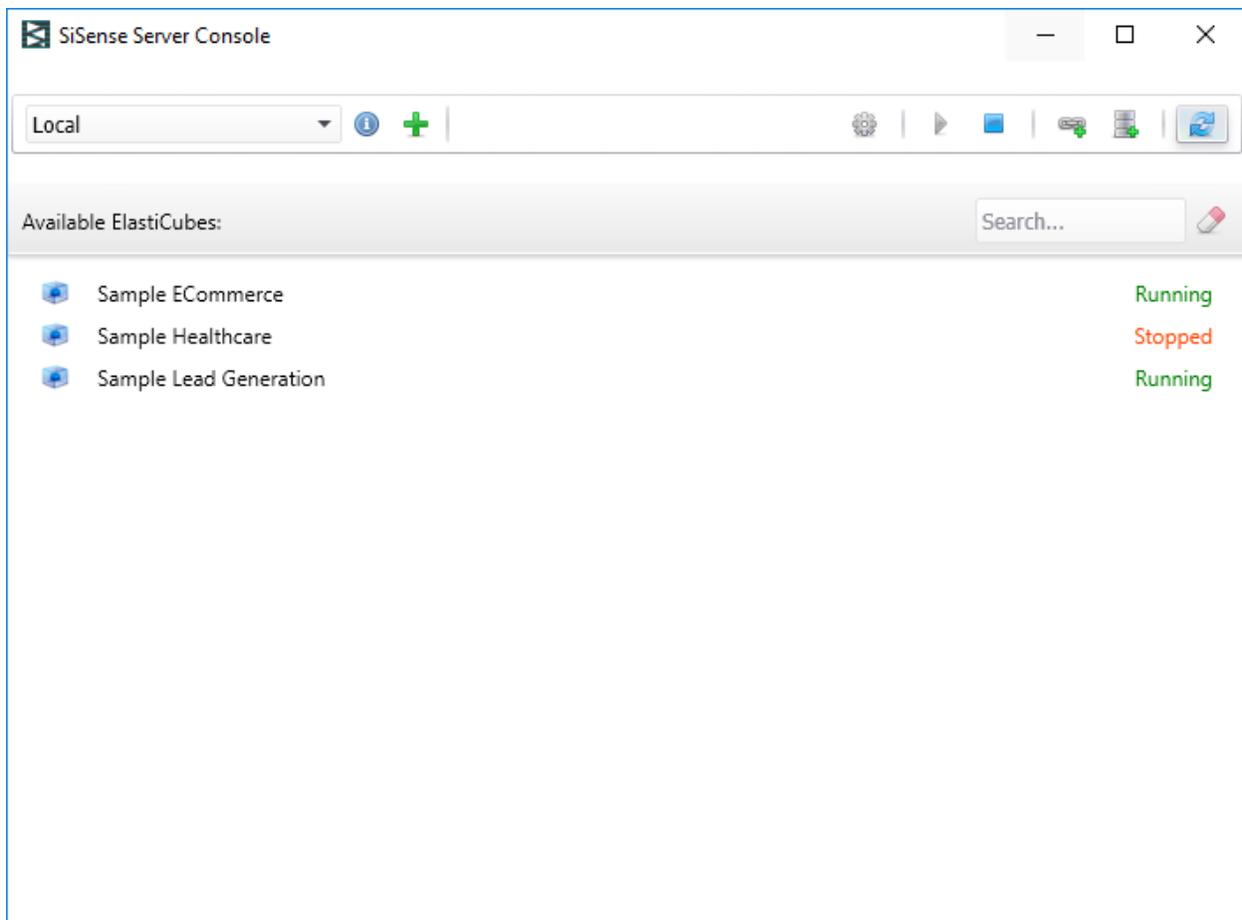
- Sisense ElastiCube Connector
- ODBC Driver

## Sisense ElastiCube Connector

The native Sisense ElastiCube Connector enables you to import data from other running ElastiCubes and provides faster performance over the Sisense ODBC.

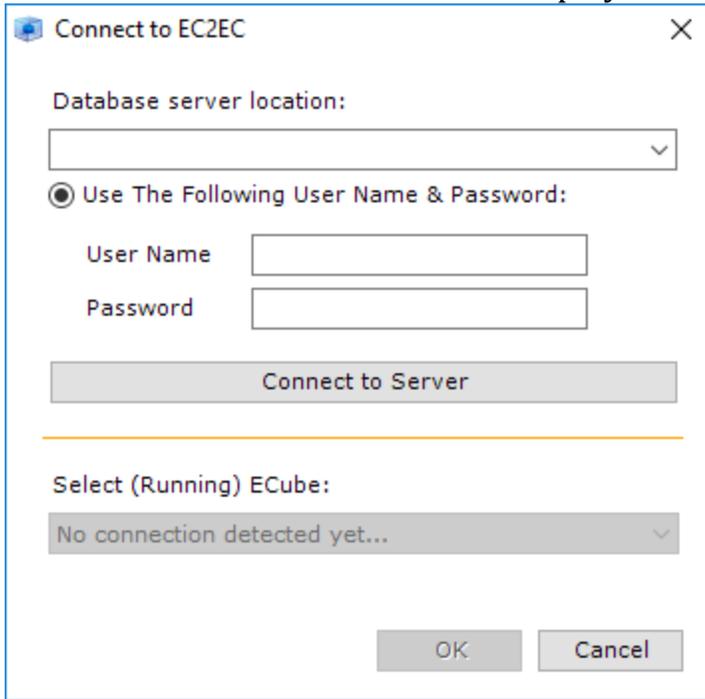
The Sisense ElastiCube Connector does not support SSL connections, however, you can import ElastiCube through SSL connections via the ODBC driver described below.

**Note:** The first time you connect to an ElastiCube, the ElastiCube must be running. For subsequent connections, the connector automatically activates the ElastiCube. You can verify and activate your ElastiCubes from the Sisense Server Console in the ElastiCube Manager Admin.



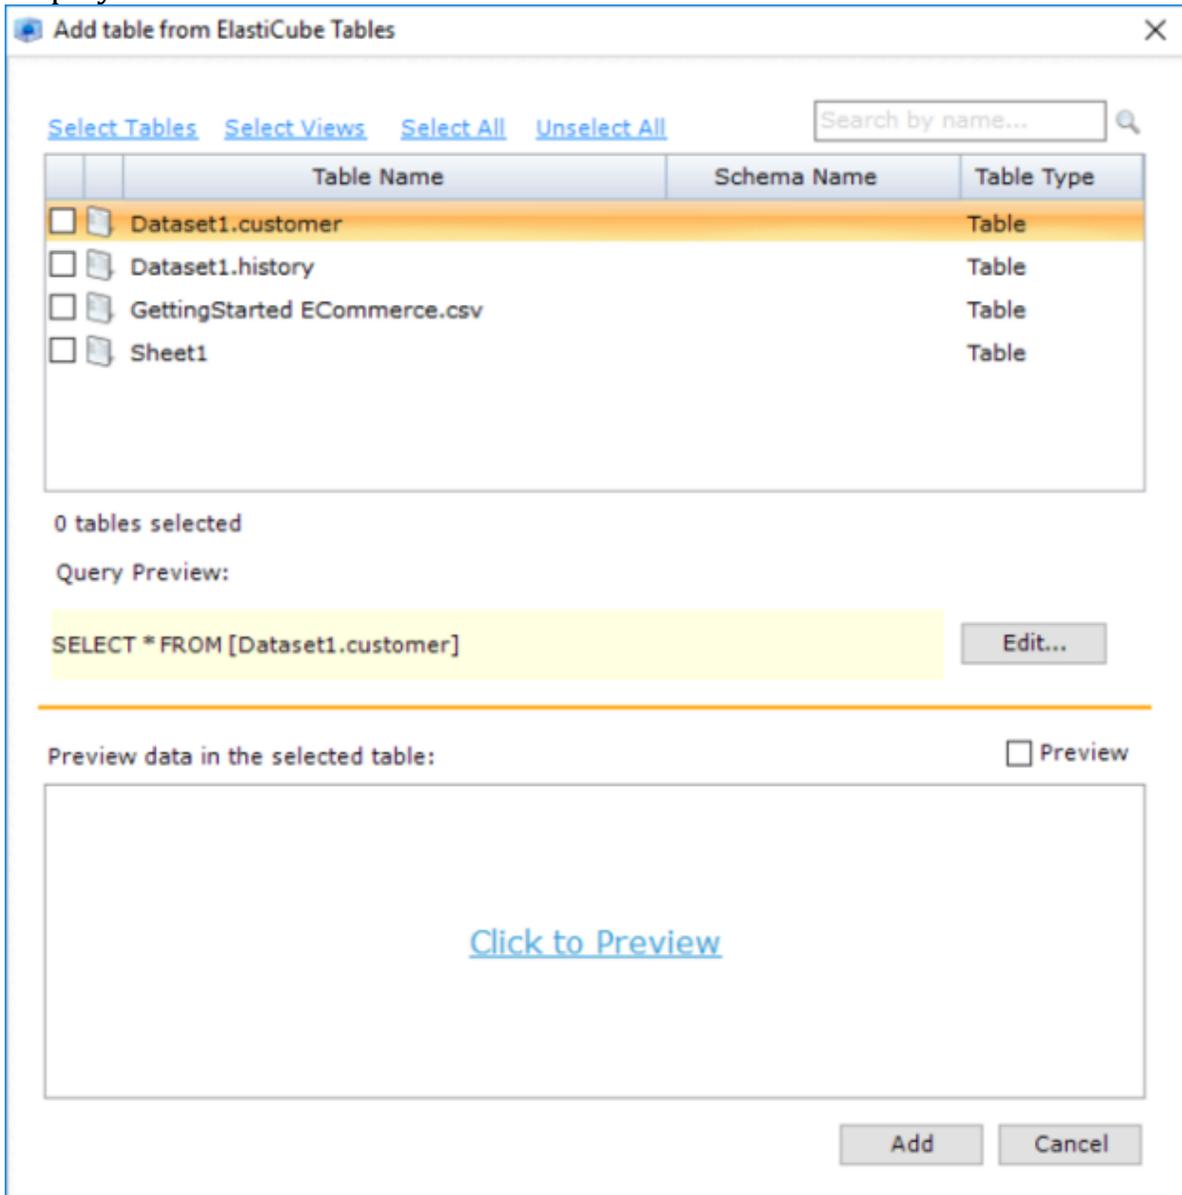
**To import data from ElastiCubes:**

1. In ElastiCube Manager, click **Add Data** and then, **Ecube Database**. The **Connect to ElastiCube** window is displayed.



2. From the **Database server location** list, enter the address of your ElastiCube. If your ElastiCube is hosted locally, the address is localhost:8081, unless you have changed your port in which case, the address will be localhost:(XXXX) where X is the port you defined. If your ElastiCube is hosted on a remote server, the address is the IP address followed by the port (XX.XX.XX.XX:port). For remote connections, open the Sisense Server Console and click  to display the remote ElastiCube's IP address.  
OR  
If you have previously entered the address, select the relevant address.
3. In the **User Name** and **Password** fields, enter your Sisense credentials that you use when logging into Sisense.
4. Click **Connect to Server**. The running ElastiCubes at the location you defined are displayed in the **Select (Running) ECube** list.

5. Select the relevant ElastiCube. The Add table from ElastiCube Tables is displayed.



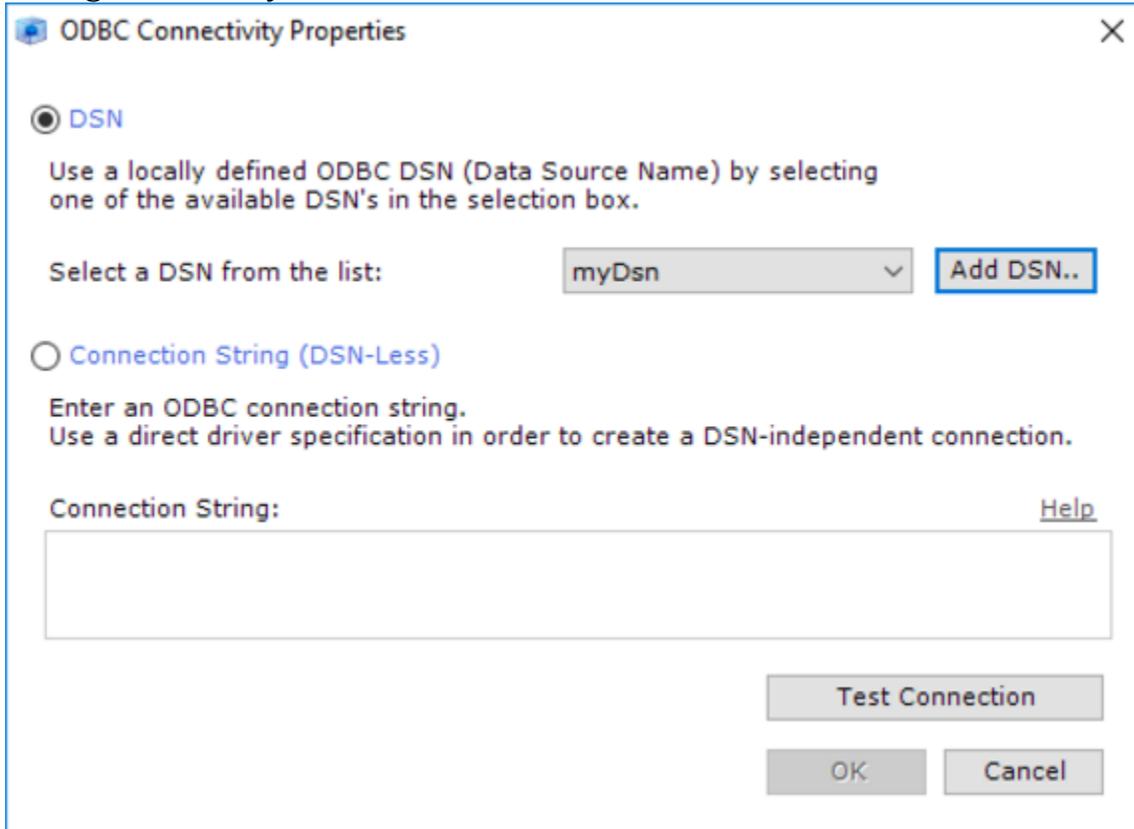
6. Select the relevant tables to import and click **Add**. The tables are displayed in the ElastiCube Manager.

## Sisense ODBC Driver

The following topics explain how you can import ElastiCube data from another ElastiCube.

### Downloading and Connecting the Sisense ODBC Driver

1. Download and install the Sisense ODBC Driver [as described here](#).
2. In ElastiCube Manager, click **Add Data** and then, **Generic ODBC Driver**.
3. Select **DSN**. If a DSN file has already been created, select it in the dropdown list, and click Test Connection. If you need to add a DSN, see Adding a DSN. If your connection has been set up, you can click OK and review the tables generated by the Sisense ODBC driver.



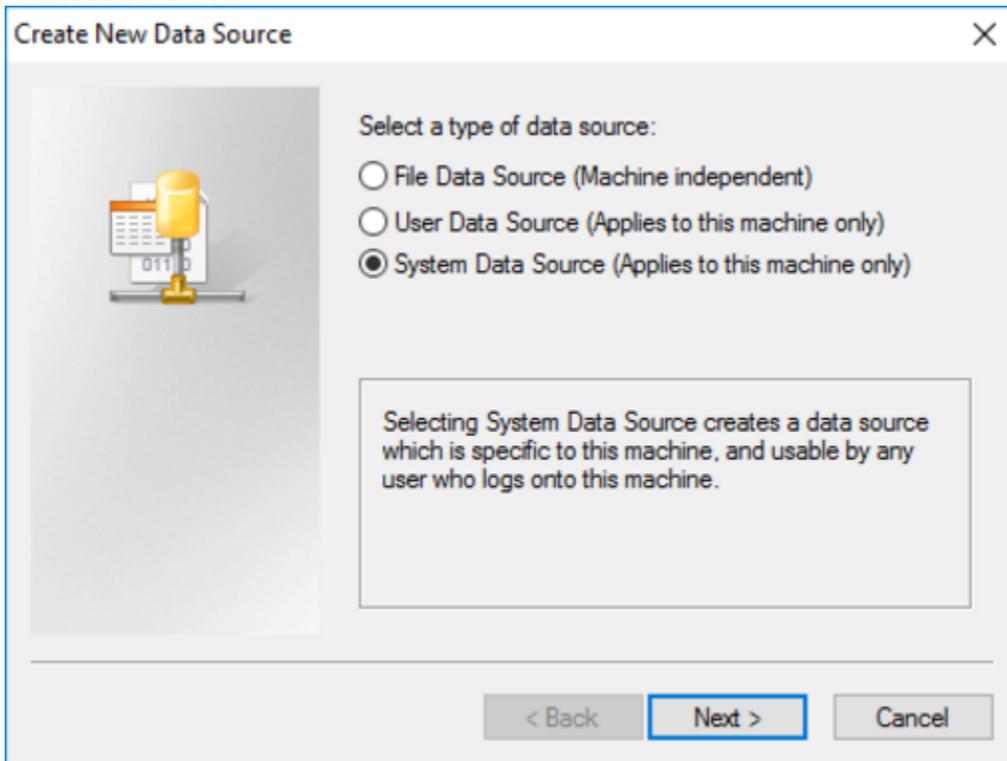
## Adding a DSN

If a DSN has not yet been created, you must add the DSN manually.

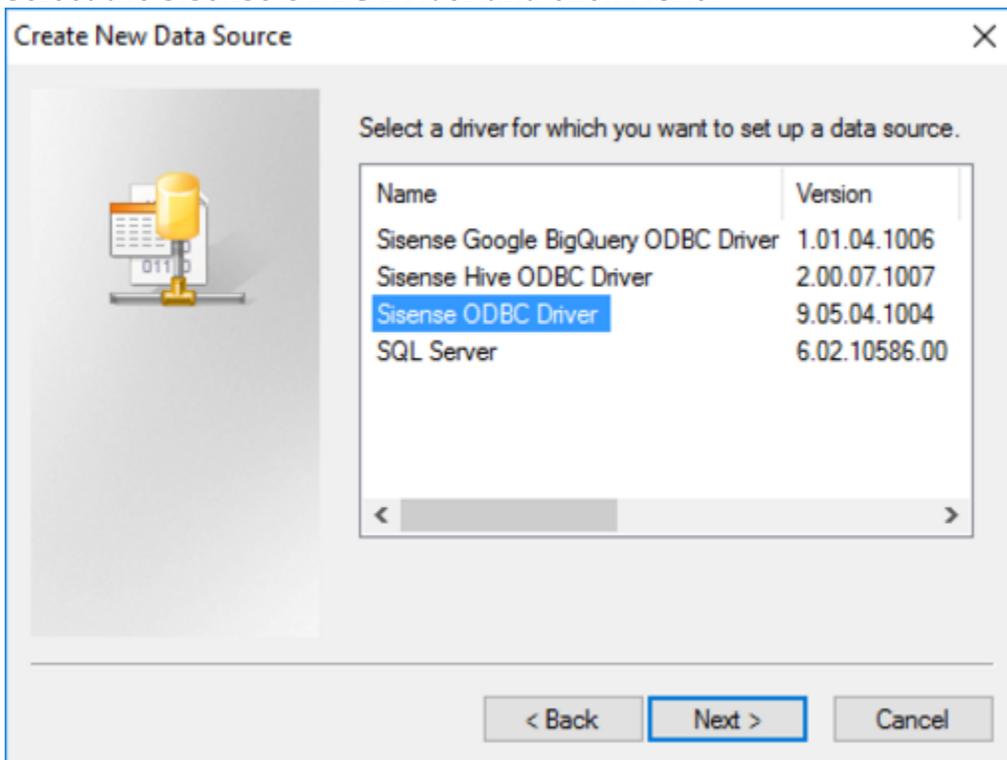
**Note:** To add a DSN you must run the Sisense ElastiCube Manager as an administrator.

**To add a DSN:**

1. Click **Add DSN**.



2. Select the System Data Source option. The created file will apply to all users in a specific machine only. Click **Next**.
3. Select the Sisense ODBC Driver and click **Next**.



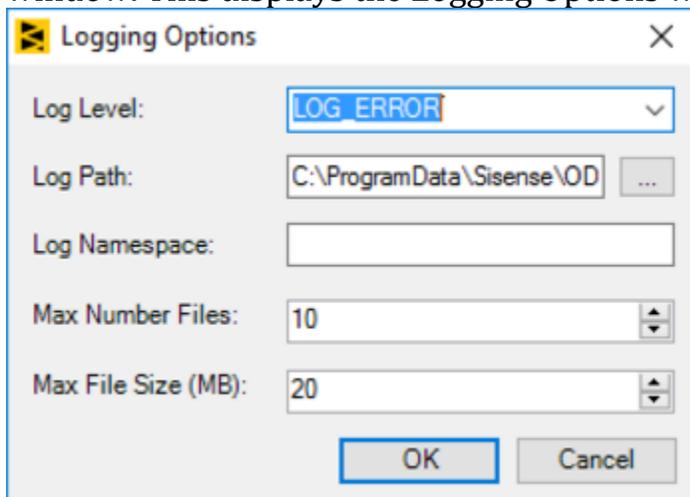
4. Click **Finish**.
5. In the New Sisense Data Source window, do the following:

Field	Procedure
Data Source Name	Enter your name for the DSN that is to be displayed in the DSN dropdown list.
Description	Enter your description of the DSN.
Server	Enter the IP address of your ElastiCube server.
Use secure connection (SSL)	Select to enable a secure SSL connection.
ElastiCube	Select the ElastiCube you want to export. OR If this is your first time configuring a DSN for the Sisense ODBC, click the ElastiCube list. The Login to Sisense window is displayed. Enter your credentials and click <b>OK</b> . Wait several seconds while Sisense returns a list of your ElastiCubes. After the list is displayed in the ElastiCube list, select the ElastiCube you want to import
Log	(Optional) To configure logging behavior for the driver, click <b>Logging Options</b> . For more information, see Logging Options.

6. Click **Test**. A pop window is displayed that indicates the connection was successful.
7. Click **OK**.
8. In the ODBC Connectivity Properties window, click **Test Connection**.
9. Click **OK** to save and close the ODBC Connectivity Properties window.

## Logging Options

You can configure logging options by clicking **Log** in the New Sisense Data Source window. This displays the Logging Options window.



The Sisense ODBC driver provides tracing functionality, which you can activate to help troubleshoot issues.

**Important:** Only enable logging long enough to capture an issue. Logging decreases performance and can consume a large quantity of disk space.

The driver allows you to set the amount of detail included in log files. The table below lists the logging levels provided by the Sisense ODBC Driver, in order from least verbose to most verbose.

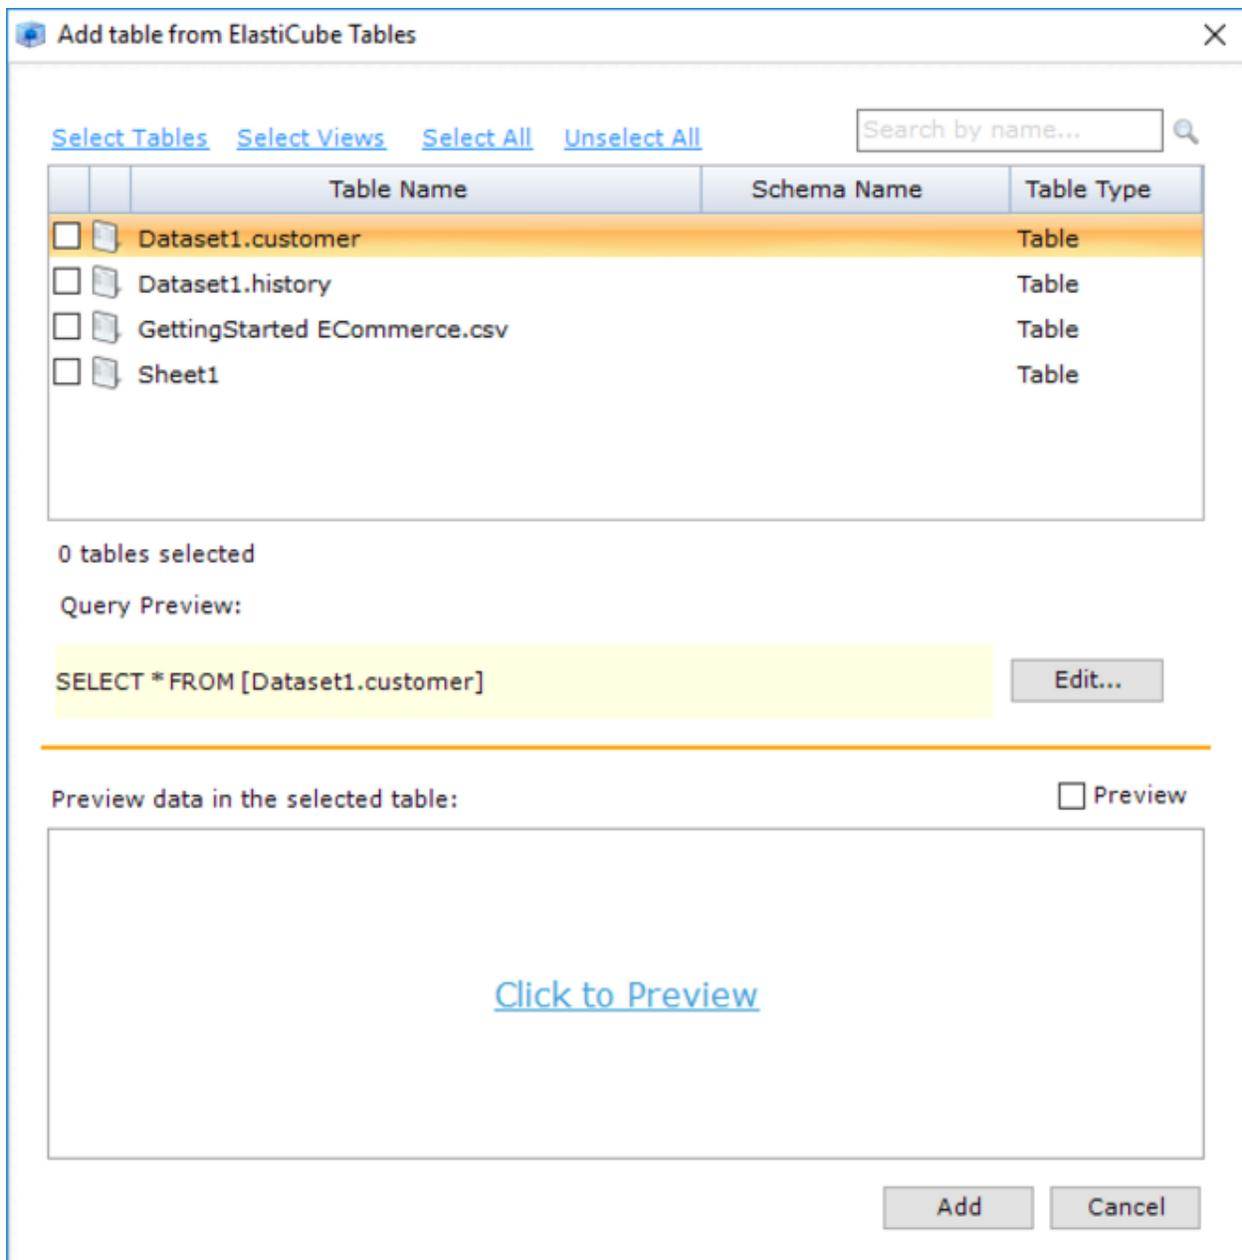
In the Logging Options window, you can configure the following information:

Field	Description
Log Level	Select the Log Level. There are seven possible options: <b>LOG_OFF:</b> Disables all logging. <b>LOG_FATAL:</b> Logs very severe error events that lead the driver to abort. <b>LOG_ERROR:</b> Logs error events that might still allow the driver to continue running. <b>LOG_WARNING:</b> Logs potentially harmful situations.

Field	Description
	<p><b>LOG_INFO:</b> Logs general information that describes the progress of the driver.</p> <p><b>LOG_DEBUG:</b> Logs detailed information that is useful for debugging the driver.</p> <p><b>LOG_TRACE:</b> Logs more detailed information than the DEBUG level.</p>
Log Path	<p>Enter the full path to the folder where you want to save log files.</p> <p>OR</p> <p>Click Browse and select the folder where you want to save log files.</p>
Log Namespace	Enter the filename for your log.
Max Number Files	Enter the maximum number of logs that can be created before the oldest log files are overwritten by newer logs.
Max File Size (MB)	Enter the maximum number of megabytes your log can reach before the log is overwritten by a newer log.

## Adding ElastiCubes Tables to your Project

After setting up the DSN, the Add table from ElastiCube Tables window is displayed.



From this window, you add your ElastiCube tables to your ElastiCube Manager. In addition, you can view the SQL syntax in the Query Preview section and click Edit to customize it.

#### **To add ElastiCube Tables to your Project:**

1. Connect to your ElastiCube tables via the Sisense ODBC Tool.
2. Add your data source.
3. In the Add table from ElastiCube Tables window, select the table you want to add to the ElastiCube Manager.
4. Click **Add**. The selected tables are added to your ElastiCube Manager.



# Connecting to Snowflake

Through the ElastiCube Manager, you can connect and integrate your Snowflake data warehouse quickly and easily to generate and analyze your data. Snowflake provides an ODBC driver for connecting to Snowflake using ODBC-based client applications such as Sisense.

## Downloading and Connecting the Snowflake ODBC Driver

1. Download and install the Snowflake ODBC Driver. For more information, click [here](#).
2. In ElastiCube Manager, click **Add Data** and then, **Generic ODBC Driver**.
3. Select **DSN**. If a DSN file has already been created, select it in the dropdown list, and click **Test Connection**. If you need to add a DSN, see [Adding a DSN](#). If your connection has been set up, you can click **OK** and review the tables generated by the Snowflake ODBC driver.



OR

Select Connection String (DSN-Less) and enter your Connection String. Click **Test Connection**. For more information, see [Connecting without a DSN](#).

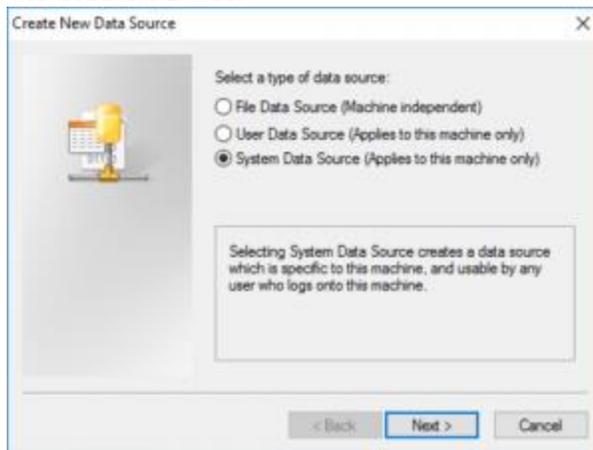
# Adding a DSN

If a DSN has not yet been created, you must add the DSN manually.

**Note:** To add a DSN you must run the Sisense ElastiCube Manager as an administrator.

**To add a DSN:**

1. Click **Add DSN**.



2. Select the **System Data Source** option. The created file applies to all users in a specific machine only. Click **Next**.
3. Select the **Snowflake ODBC Driver**, and click **Next**.
4. Click **Finish**.
5. In the Snowflake ODBC Driver DSN Setup window, enter the following details:



Field	Description
User	Snowflake user name.
Password	Snowflake password.
Server	Specifies the domain name for your account provided to you by Snowflake.
Database	Specifies the default database to use for sessions initiated by the driver.
Schema	Specifies the default schema to use for sessions initiated by the driver.
Warehouse	Specifies the default warehouse to use for sessions initiated by the driver.
Role	Specifies the default role to use for sessions initiated by the driver. The specified role should be a role that has been assigned to the specified user for the driver. If the specified role has not been assigned to the user, the role is not used for sessions initiated by the driver.
Tracing	The level of detail to be logged in the driver trace files: 0 = Disable tracing 1 = Fatal only error tracing 2 = Error tracing 3 = Warning tracing 4 = Info tracing 5 = Debug tracing 6 = Detailed tracing

When entering parameters, note the following:

- Data Source, User and Server are the only parameters required to create a DSN.
- The Password field accepts a value, but does not store the value. This is a security precaution to ensure passwords are never stored directly in the driver.
- All other parameters in the dialog are optional.

6. Click **OK**.

# Adding Snowflake Tables to your Project

After setting up the DSN, the Add table from Snowflake Tables window is displayed.

From this window, you add your Snowflake tables to your ElastiCube Manager. In addition, you can view the SQL syntax in the Query Preview section and click **Edit** to customize it.

## To add Snowflake Tables to your Project:

1. Connect to Snowflake via the Sisense ODBC Tool.
2. Add your data source.
3. In the Add table from Snowflake Tables window, select the table you want to add to the ElastiCube Manager.
4. Click **Add**. The selected tables are added to your ElastiCube Manager.

# Connecting to Splunk

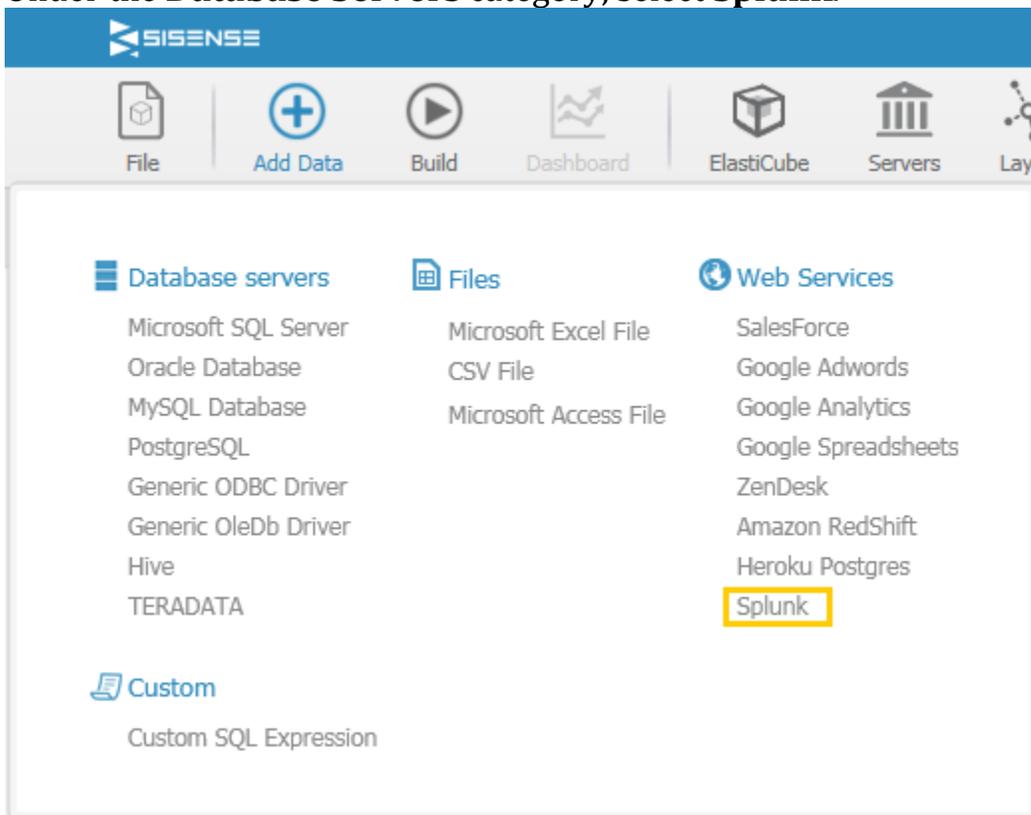
## Note:

The Splunk SDK for C# v1.0.x is deprecated, and has been replaced by the Splunk SDK for C# 2.0. Unlike the Splunk SDK for C# 1.0.x, the Splunk SDK for C# 2.0 abides by .NET guidelines, as well as FxCop and StyleCop rules. In addition, the API client in Splunk SDK for C# 2.0 is a Portable Class Library (PCL), and supports cross-platform development. Version 2 of the SDK is incompatible with version 1, and applications built with Splunk SDK for C# version 1.0.x will not recompile using Splunk SDK for C# version 2.0. See [Migrating from Splunk SDK for C# v1.0.x](#) for more information.

The ElastiCube Manager enables easy and quick access to Splunk. The steps below detail how to connect to this type of data source.

1. Click **Add data** in the top menu of the ElastiCube Manager.

2. Under the **Database Servers** category, select **Splunk**.



3. You will be prompted to enter the following information:

- **Database server location:** Enter the computer/server IP address which holds the Splunk instance.
- Select either **Windows Authentication** if configured with the database or alternatively, **Use the Following User Name &**

**Password**, and enter the database credentials.

The screenshot shows a 'Connect to Splunk' dialog box. It features a title bar with a close button. The main content area includes a 'Database server location:' dropdown menu. Below this are two radio buttons: 'Windows Authentication' (selected) and 'Use The Following User Name & Password:'. The second option has 'User Name' and 'Password' text boxes. A 'Connect to Server' button is below. A horizontal line separates this from the 'Select Database:' section, which has a dropdown menu showing 'No connection detected yet...'. At the bottom are 'OK' and 'Cancel' buttons.

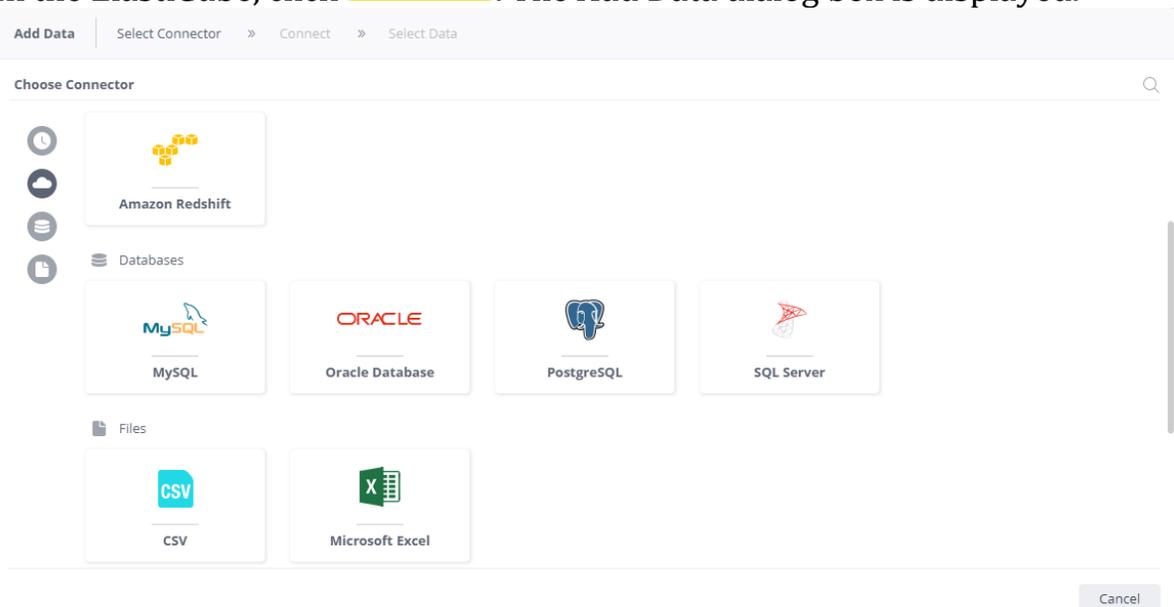
4. Click **Connect to Server**.  
A list of available Splunk instances will appear in the list box below.
5. Select the relevant Splunk instance you want to work with and click **OK**.  
All saved searches associated with the Splunk instance will appear in a new window.  
To preview data contained in a particular Splunk search, highlight the search in the list and in click in the **Preview** pane. To preview the search, select the **Preview** checkbox.
6. Select the checkbox next to each table or view you want to use.  
Fields with similar names can be linked by selecting the **Automatically create relationships for fields with the same name** checkbox.
7. Once all relevant tables are selected, click **Add**.

# Connecting to SQL

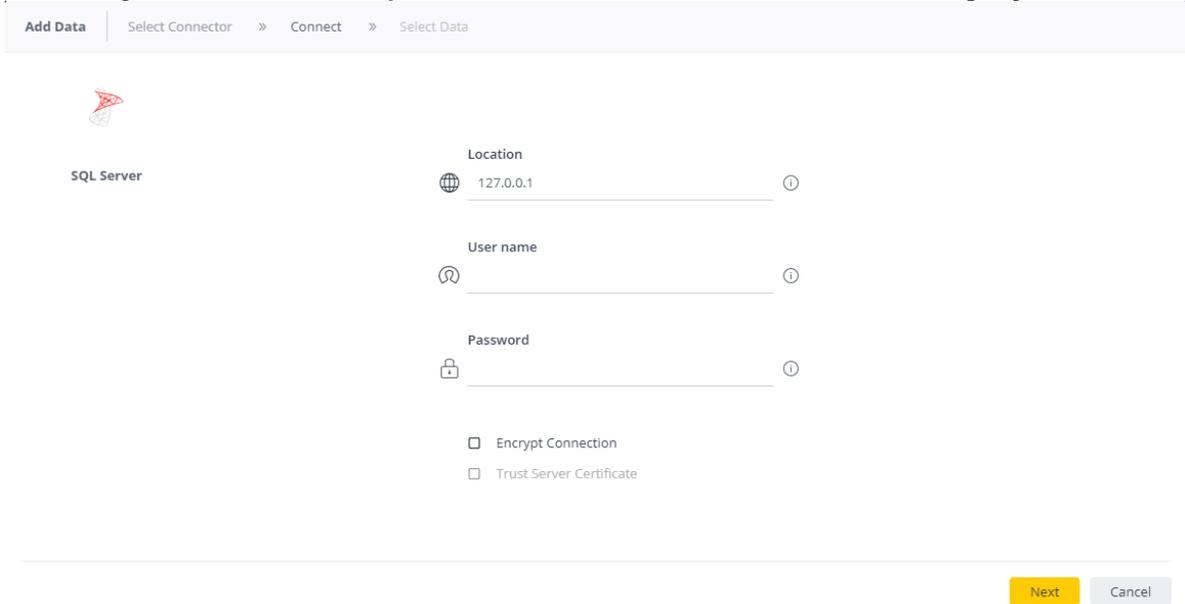
The web-based version of the ElastiCube Manager enables easy and quick access to databases, tables and views contained with SQL databases.

## To import data from your SQL database:

1. In the Data page of ElastiCube Manager, open an ElastiCube or click **+ ElastiCube** to create a new ElastiCube.
2. In the ElastiCube, click **+ Data**. The Add Data dialog box is displayed.



3. Click **SQL Server**. The SQL Server Database Connect area is displayed.



Add Data | Select Connector » Connect » Select Data

SQL Server

Location  
127.0.0.1

User name

Password

Encrypt Connection  
 Trust Server Certificate

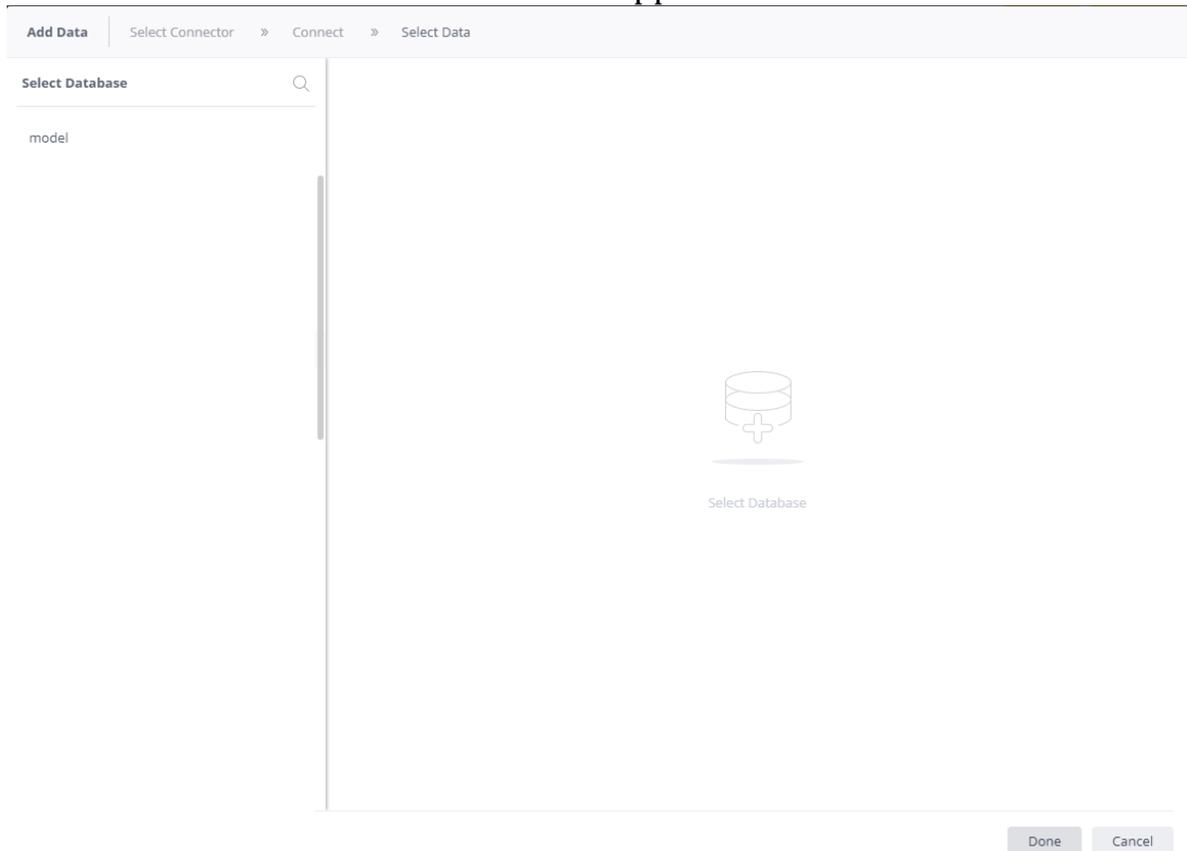
Next Cancel

4. Enter the following details:

- **Location:** Enter the computer/server IP address of the database. To connect to a database running on your own computer enter **localhost**.
- If Windows Authentication is configured with the database in the **User Name** and **Password** fields, enter your database credentials.

5. **(Optional)** Select **Encrypt Connection** to configure the driver to encrypt all communication with the SQL Server instance before sending it over the network, and then select **Trust Server Certificate**.

6. Click **Next**. A list of tables in the database are displayed. All tables and views associated with the database will appear in a new window.

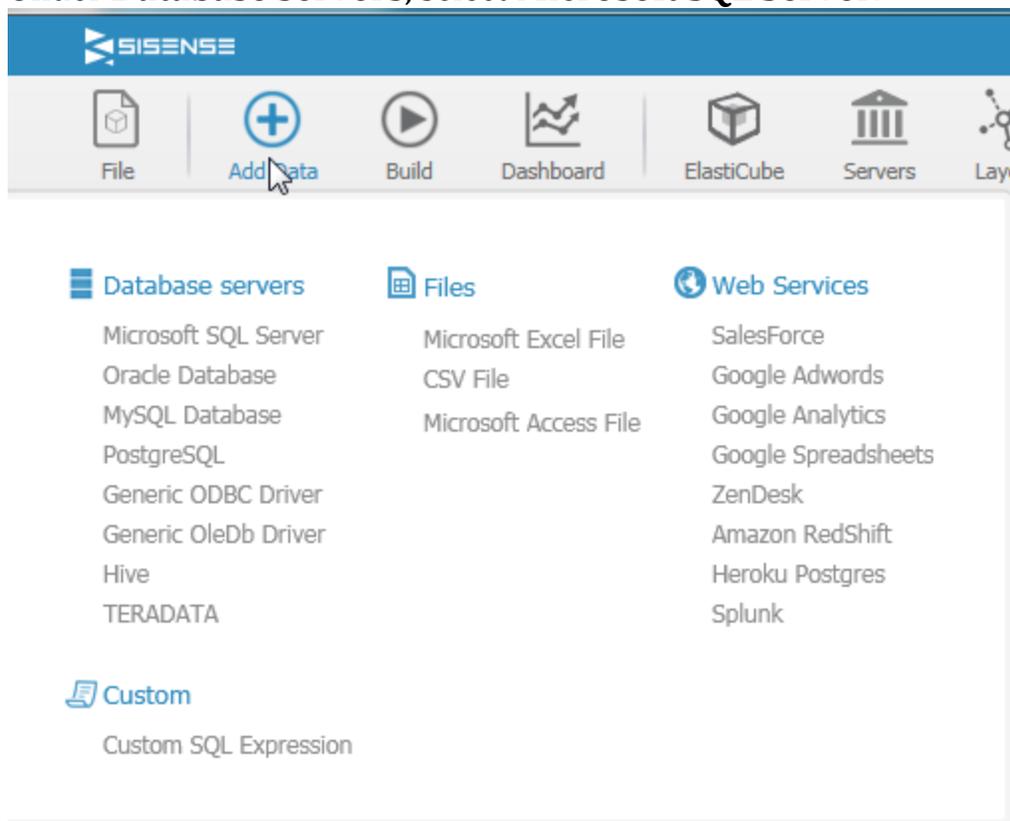


7. From the **Select Table** list, select the relevant database you want to work with, and click **Done**. To view a preview of data contained in a particular table, select the table or view in the list and in click the Preview Table.
8. After you have selected all the relevant tables, click **Done**.

# Connecting to SQL Server

The ElastiCube Manager enables easy and quick access to databases, tables and views contained with Microsoft SQL Server databases. The steps below detail how to connect to this type of data source.

1. Click **Add data** on the top menu of the ElastiCube Manager.
2. Under **Database servers**, select **Microsoft SQL Server**.



The Connect to SQL Server window is displayed.

Connect to SQL Server

Database server location:

Windows Authentication

Use The Following User Name & Password:

User Name

Password

Encrypt Connection

Trust Server Certificate

Connect to Server

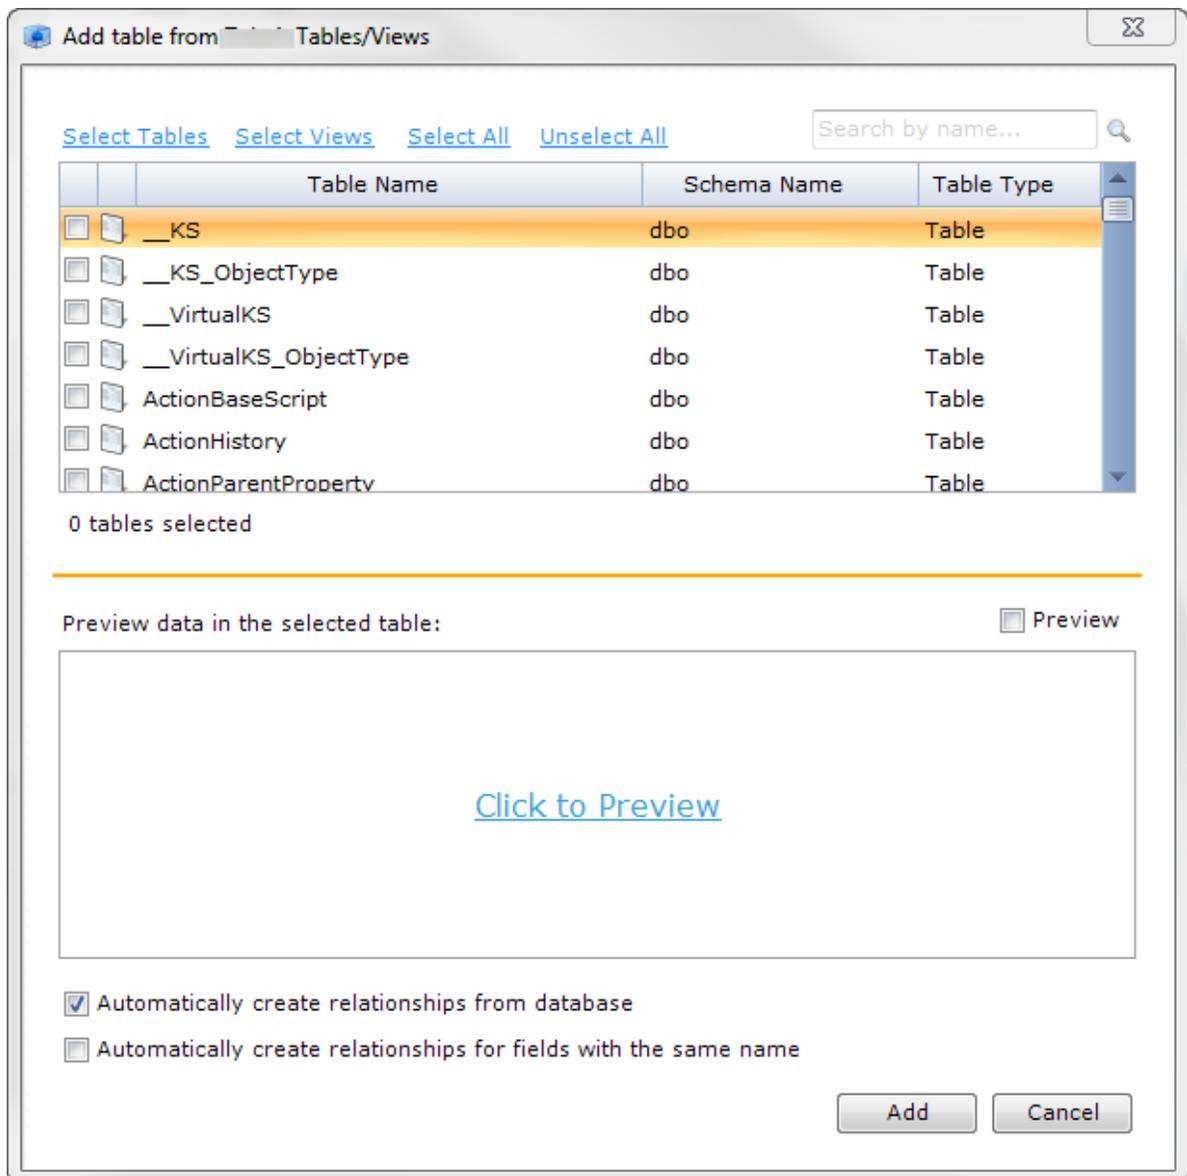
---

Select Database:

No connection detected yet...

OK Cancel

3. **Database server location:** Enter the computer/server IP address of the database. To connect to a database running on your own computer enter **localhost**.
4. Select either **Windows Authentication** if configured with the database or alternatively, **Use the Following User Name & Password**, and enter the database credentials.
5. (Optional) Select **Encrypt Connection** to configure the driver to encrypt all communication with the SQL Server instance before sending it over the network, and then select **Trust Server Certificate**.
6. Click **Connect to server**.
7. From the **Select Database** list, select the database you want to work with, and click **OK**.  
All tables and views associated with the database will appear in a new window.



To view a preview of data contained in a particular table, highlight the table or view in the list and click the preview pane below. To preview the table, select the **Preview** checkbox. Enable the checkbox next to each table or view you would like to use.

Existing relationships between tables can be automatically replicated in the ElastiCube by selecting the **Automatically create relationships from database** option. Likewise fields with similar names can be linked by selecting the **Automatically create relationships for fields with the same name** option.

8. If you want to customize the data before importing it into the ElastiCube, you can run a custom SQL query to manipulate the data. This can be useful,

for example, when you want to import only a portion of the data, rather than all of the data.

9. Once all relevant tables are selected, click **Add**.

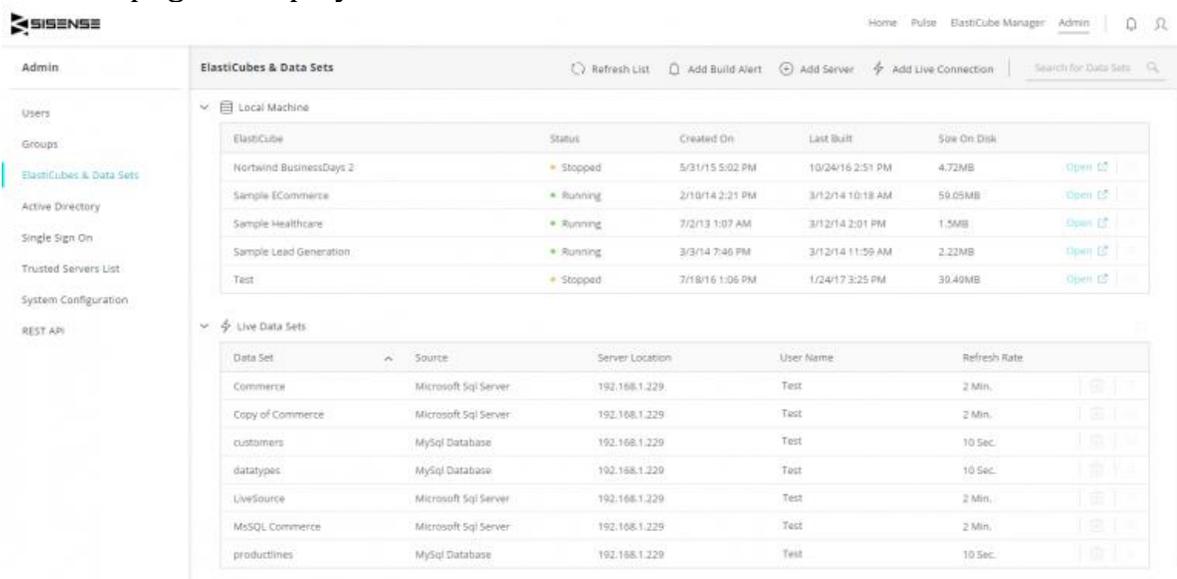
# Adding SQL Server Live Connections

This topic describes how you can add live connections to your SQL Server databases from the Sisense Web Application.

For information about live connections, see [Sisense Live Connect](#).

## To add a SQL Server Live connection:

1. In the Sisense Web Application, select **Admin > Data Sources**. The Data Sources page is displayed.



2. In the Data Sources page, click **Add Live Connection**. The Add Live Connection dialog box is displayed.
3. In the Add Live Connection dialog box, select **SQL Server**.
4. Under the Connect tab, enter the following credentials:

**Location:** The IP address of your data source.

**User Name:** The user name to access data source.

**Password:** The password to access your data source.

5. (Optional) Select **Encrypt Connection** to configure the driver to encrypt all communication with the SQL Server instance before sending it over the network, and then select **Trust Server Certificate**.
6. Click **Next**.
7. Under the Select Data tab, on the left side is a list of schemas located in your data source. Select the relevant schema. You can find the schema easily by searching for it through the search field at the top of the list to filter the schemas displayed. After selecting the schema, a list of the available tables and views are displayed.
8. Select a table or view. You can find the table easily by searching for it through the search field at the top of the list to filter the tables displayed.
9. Click **Next**.
10. Under the Connection Settings tab, set the following:
  - Connection Name:** Enter a name for your connection. This field defines the name of your data source displayed in the Sisense Web Application.
  - Refresh Rate:** Enter or select the refresh rate for your data source. The shortest refresh rate you can set is 10 seconds. The refresh rate is the time in between queries to the data source. If the previous query has not been completed, the next query is not sent until the previous refresh is complete unless the refresh rate is greater than the time it took to complete the refresh since the last query. For example, if your refresh rate is 10 seconds, and your query takes 15 seconds, the next query is sent when the previous query is complete as 10 seconds has passed and the previous query is complete. If your refresh rate is 30 seconds, and the previous query is complete, the next refresh will be in 30 seconds from the last query.
  - Time Out:** Enter or select the Time Out rate for your data source. The Time Out rate determines how long Sisense waits before cancelling a query that does not receive a response from your data source.
11. Click **Done**. The connection is added as a data set to your Data Sources page.

# Connecting to Stripe

The Sisense Stripe connector is a standalone connector that allows you to import data from Stripe's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the Stripe APIs. To obtain a connection string, you will need to create a Stripe app.

Once you have connected to Stripe, you can import a variety of tables from the Stripe API.

This page describes how to install the Stripe connector, how to connect to Stripe with a connection string, and what tables you can import into the ElastiCube

Manager:

- [Installing the Stripe Connector](#)
- [Connecting to the Stripe REST API](#)
- [Adding Stripe Tables to your ElastiCube](#)
- [Stripe Tables](#)
- [Versions](#)
- [Limitations](#)

## Installing the Stripe Connector

Sisense provides the Stripe connector as a standalone connector that you can download and add to your list of default Sisense connectors.

**To install the Stripe connector:**

1. [Download](#) the Stripe installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.  
The Stripe connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.

# Connecting to the Stripe REST API

To access Stripe's REST API from the ElastiCube Manager, you must provide valid OAuth Stripe credentials through a connection string. These credentials are provided by Stripe when you create an application.

After you receive your credentials from Stripe, you can create the connection string and provide Sisense with it to connect to your data.

## Registering an App

Follow the steps below to obtain the OAuth client credentials, the OAuthClientId and OAuthClientSecret:

1. Log into your Stripe dashboard and click **Your Account > Account Settings > Connect > Platform Settings**.
2. Enter a name, description, and other information to be displayed to users when they log in to grant permissions to your app.
3. If you are making a desktop application, set the Redirect URL to `http://localhost:33333` or a different port number of your choice. If you are making a Web application, set the Redirect URL to a page on your Web app you would like the user to be returned to after they have authorized your application.

## Adding Stripe Tables to your ElastiCube

Sisense uses connection strings to connect to Stripe and import data into the ElastiCube Manager. Each connection string contains a authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

The connection string to Stripe is in the following format:

```
jdbc:stripe:APIKey=xxxxxxxxxxxxxxxxxxxxxxxx;
```

For example:

```
jdbc:stripe:InitiateOAuth=GETANDREFRESH;Instance=xxxxxxxxx
xx;OAuthClientId=xxxxxxxxxxxxxx;
OAuthClientSecret=xxxxxxxxxx;
```

Your Stripe connection string should include the following parameters, separated by a semicolon ';':

- **OAuthClientId**: Set this to the client Id in the Connect section of your account settings.
- **OAuthClientSecret**: Set this to a secret key in the API Keys section of your account settings.
- **CallbackURL**: Set this to the Redirect URL in the Connect section of your account settings.
- **InitiateOAuth**: Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken connection property.

#### To add Stripe data:

1. In ElastiCube Manager, click **Add Data** and then, **Stripe**. The Connect to Stripe window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. Stripe is displayed in the Select Database list.
4. Click **OK**. Sisense connects to Stripe and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**.
6. The tables are displayed in the ElastiCube Manager.

## Stripe Tables

Stripe's RESTful APIs expose the following Stripe tables that you can import into the ElastiCube Manager through the Sisense Stripe connector:

### Available Tables

Name	Description
Accounts	Create, update, delete, and query the Accounts you manage in Stripe.

Name	Description
BankAccounts	Create, update, delete, and query the available Bank Accounts in Stripe.
BankAccountTokens	Insert and query the available Bank Account Tokens in Stripe.
Cards	Create, update, and query the available Cards in Stripe.
CardTokens	Create and query the available Card Tokens in Stripe.
Charges	Create, update, and query the available Charges in Stripe.
Coupon	Get and delete the available discount of a Subscription
Customers	Create, update, delete, and query the available Customers in Stripe.
CustomerDiscounts	Get and delete the available discount of a Customer.
Disputes	Update and query the available Disputes in Stripe.
Plans	Create, update, delete, and query the available Plans in Stripe.
Refunds	Query the available refunds in Stripe.
SubscriptionDiscounts	Get and delete the available discount of a Subscription.
SubscriptionItems	Create, update, delete, and query the available subscription items in Stripe.
TransferReversals	Query the available categories in Stripe.
Transfers	Query the available transfers in Stripe.

## Available Views

Name	Description
AvailableBalance	Query the Available Balance in Stripe.
BalanceHistory	Query Balance History in Stripe.
CountrySpecs	Query the available Country Specs in Stripe.
Events	Query the available events in Stripe.
Orders	Query the available orders in Stripe.
PendingBalance	Query the available balance in Stripe.
Products	Query the available products in Stripe.
Skus	Query the available SKUs in Stripe.

## Versions

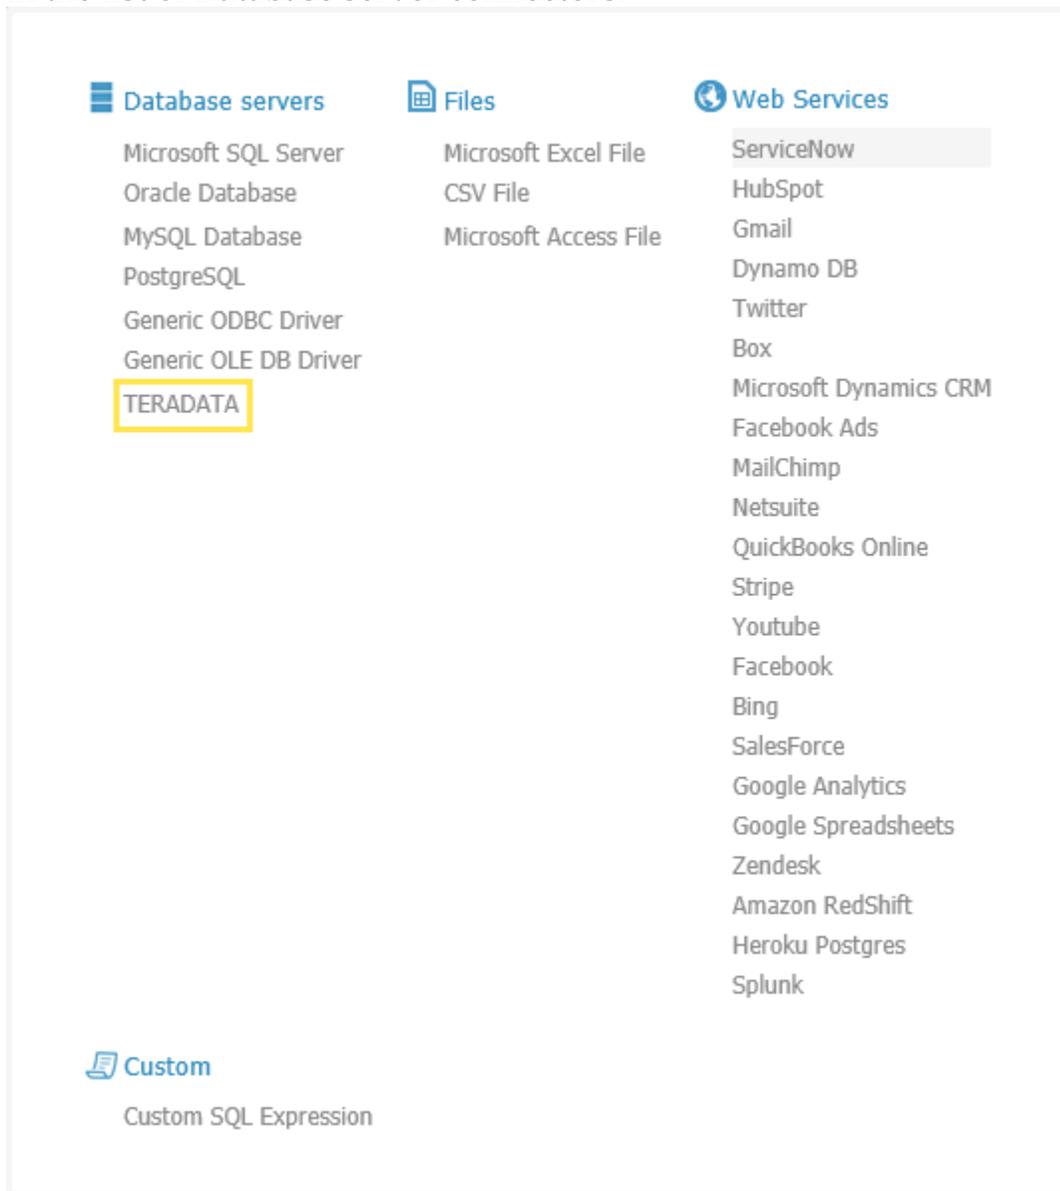
Version	Release Date	Improvements
6.6.1.57	May 18, 2017	New Objects supported: Coupons, CustomerDiscounts, Plans, SubscriptionDiscounts, SubscriptionItems
6.6.0.37	March 16, 2017	General Release
6.5.1.10	December 15, 2016	General Beta release

## Limitations

- Not all fields can be filtered.
- When a field cannot be filtered, you can not perform an accumulative build on that field.

# Connecting to Teradata

You connect to a Teradata database by selecting TERADATA from the list of connectors in the Sisense Add Data window. The Teradata Connector is displayed in the list of Database server connectors.

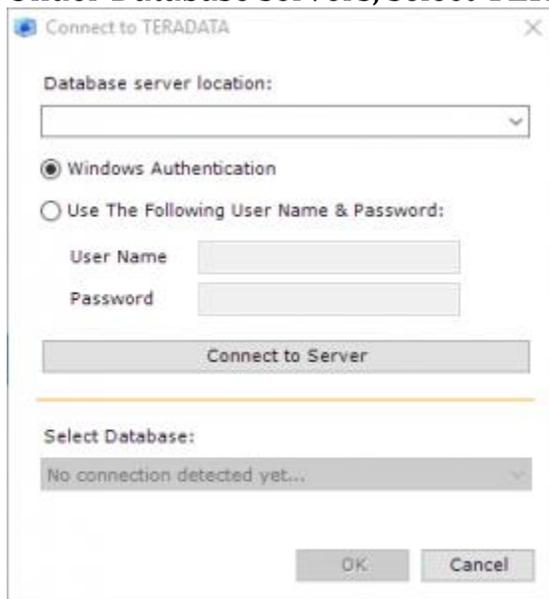


# Importing Data from Teradata

You can add Teradata tables to the ElastiCube Manager by clicking **Add Data**. By entering your Teradata credentials, you can easily connect to your Teradata database and import its tables into the ElastiCube Manager.

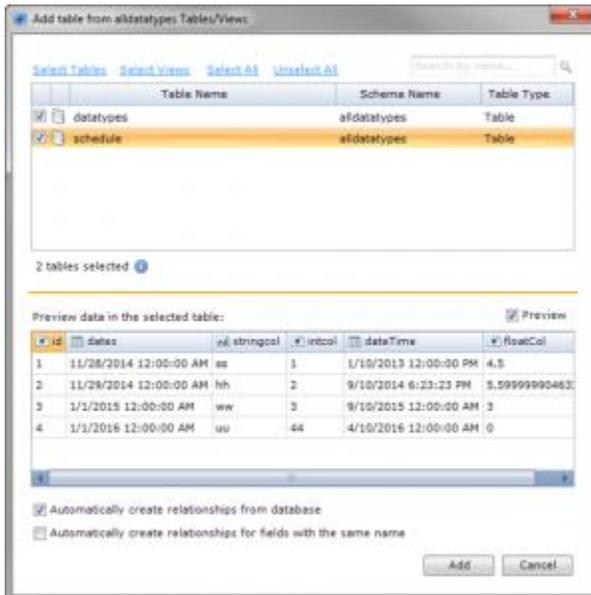
## To add tables from Teradata:

1. Click **Add data** on the top menu of the ElastiCube Manager.
2. Under Database servers, select **TERADATA**.



3. **Database server location:** Enter the computer/server IP address of the database. To connect to a database running on your own computer enter **localhost**.
4. Select either **Windows Authentication** if configured with the database or alternatively, Use the Following **User Name** and **Password** fields, and enter the database credentials.
5. Click **Connect to server**.
6. From the **Select Database** list, select the relevant database you want to work with, and click **OK**.

7. All tables and views associated with the database will appear in a new window.



8. To view a preview of data contained in a particular table, highlight the table or view in the list and in click the preview pane below. To preview the table, select the **Preview** checkbox. Enable the checkbox next to each table or view you would like to use. Existing relationships between tables can be automatically replicated in the ElastiCube by selecting the **Automatically create relationships from database** option. Likewise fields with similar names can be linked by selecting the **Automatically create relationships for fields with the same name** option.
9. Once all relevant tables are selected, click **Add**.

## Limitations

**Custom tables:** This is a feature not yet available, only entire tables can be imported.

# Connecting to Twitter

The Sisense Twitter connector is a standalone connector that allows you to import data from Twitter's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to Twitter's API. To obtain a connection string, you will need to create a Twitter app.

Once you have connected to Twitter, you can import a variety of tables from the Twitter API.

This page describes how to install the Twitter connector, how to connect to Twitter with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the Twitter Connector](#)
- [Connecting to the Twitter API](#)
- [Twitter Tables](#)
- [Versions](#)

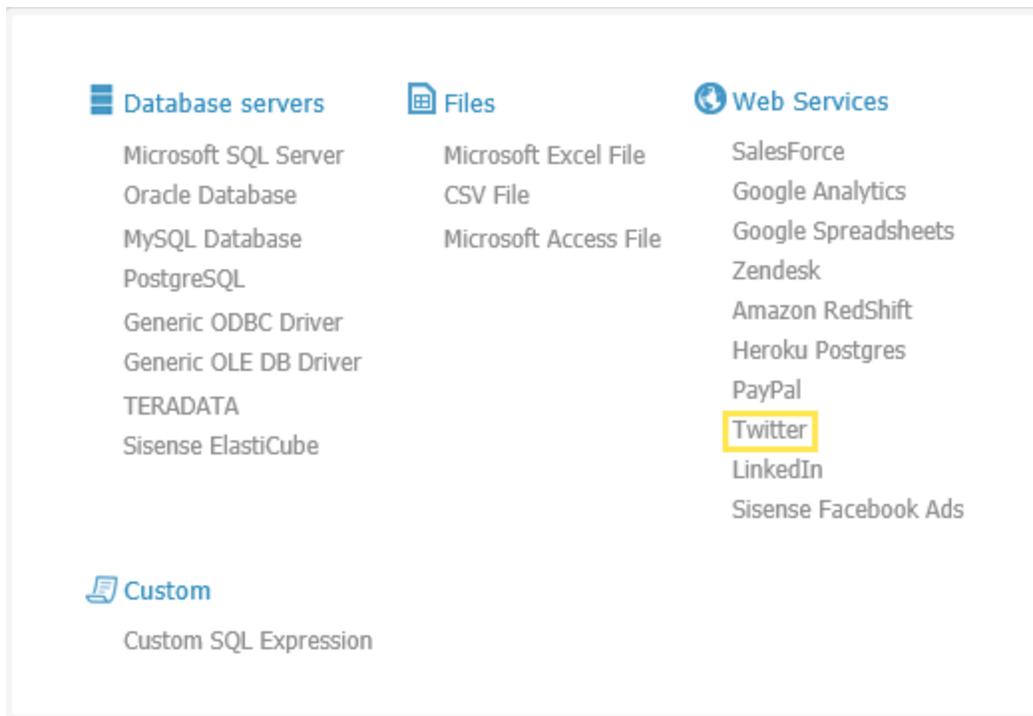
## Installing the Twitter Connector

Sisense provides the Twitter connector as a standalone connector that you can download and add to your list of default Sisense connectors.

### To install the Twitter connector:

1. [Download](#) the Twitter installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.

The Twitter connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.



## Connecting to the Twitter API

### Accessing Ads Data from Twitter

Before requesting data from Twitter, you must apply for access to your data from Twitter by submitting this [request form](#). As part of your request, you must provide your APP ID. [Click here](#) for instructions on how to find your APP ID. After you have created a Twitter app, you can manage it [here](#). This is useful if you need to set permissions so you can access the relevant data.

### Creating an App

To access Twitter's REST API from Sisense, you must provide valid Oauth Twitter credentials. These credentials are provided by Twitter when you create an application. Each connection string contains a authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

The Twitter connection string looks like this:

```
jdbc:twitter:OAuthClientId=xxxxxxxxxxxx;OAuthClientSecret=xxxxxxxxxxxx;InitiateOAuth=GETANDREFRESH;CallbackURL=http://localhost/;Timeout=0;
```

The following parameters are mandatory when connecting with a customized connection string:

- **OAuthClientId:** Set this to the consumer key in your app settings.
- **OAuthClientSecret:** Set this to the consumer secret in your app settings.
- **CallbackURL:** Set this to the callback URL you specified in your app settings.
- **InitiateOAuth:** Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken and OAuthAccessTokenSecret.

You can follow the steps below to obtain the OAuth client credentials:

1. Log in to <https://dev.twitter.com/apps>.
2. Click **Create New App**.
3. Enter the name, description, and website to be displayed to users when they connect.
4. Define the callback URL setting. If you are making a Web app, set the callback URL to the URL you would like users to be returned to after they have granted your application privileges. If you are making a desktop application, set the callback URL to `http://localhost/`.

Once you have created the app, you will be redirected to a page with information about your app.

If you intend to communicate with Twitter only as the currently authenticated user, then you can obtain the OAuthAccessToken and OAuthAccessTokenSecret directly from this page on Twitter. The OAuthAccessToken and OAuthAccessTokenSecret are listed under the OAuth Settings in the Your Access Token section. You can then connect by setting these connection string properties and do not need to continue to the other steps.

Otherwise, if you need to generate access tokens for other user accounts besides the one you used to create the app on Twitter, use the consumer key and consumer secret in the following steps.

## Adding Twitter Tables to your ElastiCube

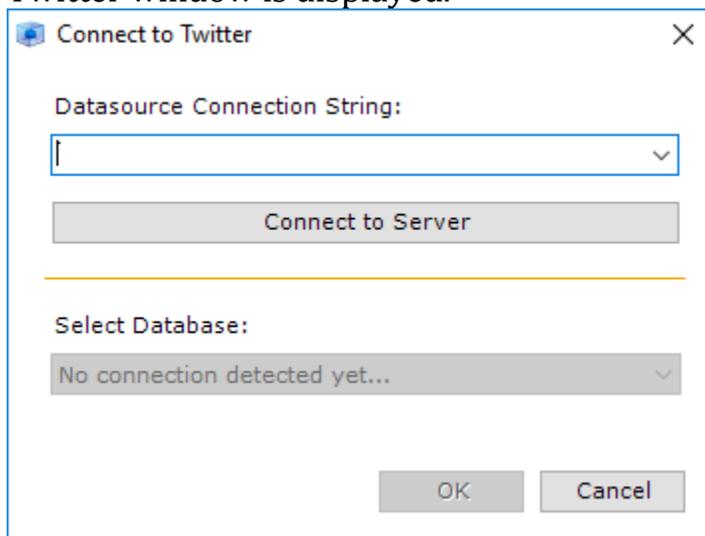
Sisense uses connection strings to connect to Twitter and import data into the ElastiCube Manager. The default connection string for Twitter is:

```
jdbc:twitter:InitiateOAuth=GETANDREFRESH;
```

After you connect using the default connection string, Sisense redirects you to Twitter's Login page where you log in and your account is authenticated. Once the account is authenticated, the relevant tables can be added to your ElastiCube. If you want to customize your connection string, you can by creating a Twitter app and passing additional connection string parameters. For more information, see [Creating an App](#) at the end of this page.

### To add Twitter data:

1. In ElastiCube Manager, click **Add Data** and then, **Twitter**. The Connect to Twitter window is displayed.



2. In **Datasource Connection String**, enter the connection string:  
`jdbc:twitter:InitiateOAuth=GETANDREFRESH;`
3. Click **Connect to Server**. You are redirected to the Twitter Login page.
4. Log in with your credentials. Twitter is displayed in the Select Database list in the ElastiCube Manager.
5. Click **OK**. Sisense connects to Twitter and displays a list of tables available for you to import.
6. Select the relevant tables and click **Add**. The tables are displayed in the ElastiCube Manager.

# Twitter Tables

The Sisense Twitter connector allows you to import the following tables into the ElastiCube Manager.

Name	Description
DirectMessagesReceived	Query and delete direct messages received by the authenticated user.
DirectMessagesSent	Send direct messages and query messages sent by the authenticated user.
Favorites	Create, delete, and query a list of favorite tweets of the authenticated user and allow the user to favorite new tweets or remove existing favorites.
Following	Create, delete, and query a list of users that the current Twitter account is following, otherwise known as friends.
Tweets	Create, delete, and query status updates and tweets from the authenticated user.
AccountSettings	Query account settings about the currently authenticated user.
AdAccounts	Retrieve all of the advertising-enabled accounts the authenticating user has access to.
AdAvailableAudiences	Query available audiences for this account.
AdCampaigns	Retrieve details for some or all campaigns associated with the current account.
AdFundingInstruments	Retrieve some or all funding instruments associated with the account specified in the path.

Name	Description
AdInsights	Provide qualitative information about various Twitter audiences.
AdLineItems	Retrieve some or all funding instruments associated with the account specified in the path.
AdPromotedTweets	Retrieve all of the advertising-enabled tweets the authenticating user has access to.
AdStats	Query information about all available metrics for a certain entity.
Followers	Query a list of users following the current Twitter account.
ListMembers	Query the members of a specified list.
Lists	Query Twitter list information based on a set of criteria.
ListSubscribers	Query the subscribers to a specified list.
Mentions	Query a list of retweets of the authenticated user.
Retweets	Query a list of retweets of the authenticated user.
Trends	Query the daily trending topics from Twitter.
Users	Query a list of users based on the SearchTerms, Id, or Screen_Name.

## Versions

Version	Release Date	Improvements
6.6.0.49	March 20, 2017	AdStats table did not return results as expected
6.5.1.33	February 8, 2017	Ads Marketing table added

# Using Custom SQL to Import Data

Create custom SQL queries to manipulate the data that you want to import, for example:

- Import selected data only, to save import and build times in cases where you do not require all the data.
- Perform various procedures in the database such as table joins.

After pushing your custom query, you can preview the updated tables before importing the data into the ElastiCube.

Notes:

This feature is available for the following relational databases: Microsoft SQL Server, MySQL, PostgreSQL, Oracle, and Amazon Redshift. This feature runs a query in the source database, which may require a lot of processing power from your database, especially for complex queries.

**To create a custom SQL query:**

1. After connecting to your data, click **Edit**.
2. Select **Manual Query Input**, and type in your query.
3. Click **Apply Table**. Your preview will update to reflect the query.
4. You can now click **Add** to import the data based on your SQL query.

Add table from 38RecordsOfInt Tables/Views

Select Tables Select Views Select All Unselect All Search by name...

	Table Name	Schema Name	Table Type
<input checked="" type="checkbox"/>	BigTable3B	dbo	Table
<input type="checkbox"/>	INTSCOUNTRIES	dbo	Table
<input type="checkbox"/>	SEARCHTABLE	dbo	Table
<input type="checkbox"/>	STRINGSCOUNTRIES	dbo	Table

1 table selected

Query Preview:

select \* from BigTable3B 1 Edit...

Preview data in the selected table:  Preview

RowID	Value	Search_Date	QUERY_ENGINE	Query_Type	COU
1	0	1/30/2012 1:08:09 PM	3	1	
2	0	1/2/2009 5:42:03 PM	7	2	
3	1	1/4/2007 10:56:32 PM	9	3	
4	3	11/15/2013 5:51:19 AM	2	4	
5	2	7/14/2006 1:18:01 PM	6	5	

Automatically create relationships from database  
 Automatically create relationships for fields with the same name

Add Cancel

BigTable3B Query Properties

Auto Query Structure

Set query identifiers' delimiters by selecting one of the available delimiter types or set your custom delimiters and watch the query preview as it displays the resulting query.

Delimiter: Square Brackets [ ] Starting: [ Ending: ]

Query Preview:

```
SELECT * FROM [dbo].[BigTable3B]
```

Manual Query Input

Type your query at the text box below.

```
select * from BigTable3B
```

3 Apply All Apply Table Cancel

# Using Custom SQL to Import Data

Sisense allows you to import data from a variety of data sources. For users that want additional control over the data that is imported, you can define custom queries that run custom SQL against the source database. For example:

- Write SQL to import only selected table columns when you do not require all the data.
- Perform various procedures in the database such as table joins and import only the result set.

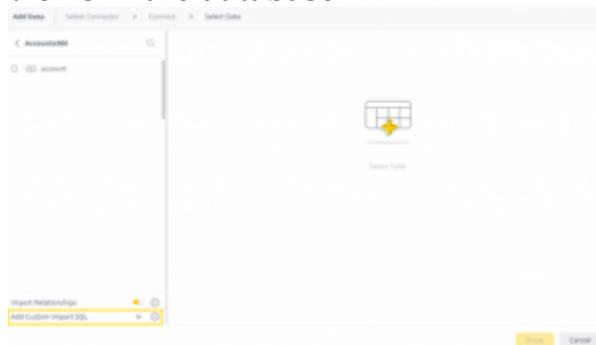
After writing your custom query, you can preview the new table before importing the data into the ElastiCube.

## Notes:

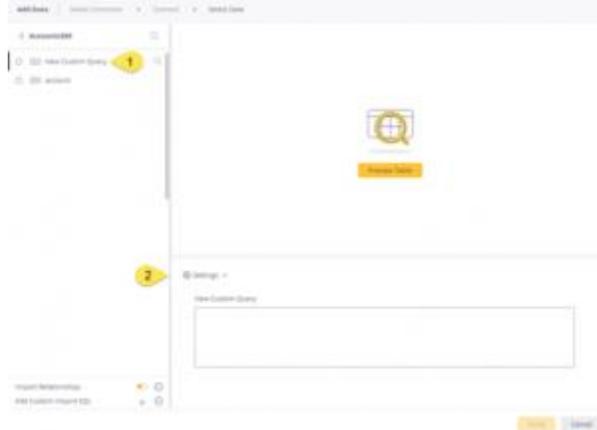
- This feature is available for the following relational databases: Microsoft SQL Server and MySQL.
- This feature runs a query in the source database. The performance will depend on processing power from your database, especially for complex queries.

## To add tables with custom queries:

1. After selecting your database, click + at the bottom the list of tables and views in the database.



An empty custom table is added and the Settings area is displayed.



2. Click + to add additional custom tables.
3. In the Settings area, enter your custom SQL query.
4. Click **Preview Table** to preview the results of your custom query.

OR

Select the table from the Table list and click **Done** to import your custom data into the ElastiCube Manager.

# Connecting to Xero Accounting

The Sisense Xero connector is a standalone connector that allows you to import data from Xero's API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the Xero APIs. To obtain a connection string, you will need to create a Xero app.

Once you have connected to Xero, you can import a variety of tables from the Xero API.

This page describes how to install the Xero connector, how to connect to Xero with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the Xero Connector](#)
- [Connecting to the Xero REST API](#)
- [Xero Tables](#)
- [Limitations](#)

## Installing the Xero Connector

Sisense provides the Xero connector as a standalone connector that you can download and add to your list of default Sisense connectors.

**To install the Xero connector:**

1. [Download](#) the Xero installation file.
2. Open the installation file and click **Install**.

3. After the installation process is complete, click **Close**.  
The Xero connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.

## Connecting to the Xero REST API

To access Xero's REST API from the ElastiCube Manager, you must provide valid OAuth Xero credentials through a connection string. These credentials are provided by Xero when you register an application.

After you receive your credentials from Xero, you can create the connection string and provide Sisense with it to connect to your data.

### Registering an App

Follow the steps below to obtain the OAuth client credentials, the OAuthClientId and OAuthClientSecret:

1. Log in to the Xero Developer Portal.
2. Click **My Applications > Add Application**.
3. Select the **Public** option.
4. Enter a name for your application and the URL of your company. This information is displayed to users when they connect.
5. If you are making a Desktop application, set the Callback Domain to 'localhost'.

If you are making a Web application, set the Callback Domain to the domain name of the URL where the user returns with the token that verifies that they have granted your app access.

After you click Save, you are shown your OAuth credentials, the consumer key, and consumer secret.

### Adding Xero Tables to your ElastiCube

Sisense uses connection strings to connect to Xero and import data into the ElastiCube Manager. Each connection string contains a authentication

parameters that the data source uses to verify your identity and what information you can export to Sisense.

Each connection string contains authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

The connection string to Xero is in the following format:

```
jdbc:xero:OAuthClientId=xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx;
```

For example:

```
jdbc:xero:OAuthClientId=xxxxxxxxxxxxxxxx;OAuthClientSecret=xxxxxxxxxxxxxxxx;
```

```
InitiateOAuth=GETANDREFRESH;XeroAppAuthentication=PUBLIC;
```

Your Xero connection string should include the following parameters, separated by a semicolon

- **OAuthClientId**: Set this to the consumer key in your app settings.
- **OAuthClientSecret**: Set this to the consumer secret in your app settings.
- **XeroAppAuthentication**: Set this to the type of your application. Allowed values: PUBLIC or PARTNER. Default: PUBLIC.
- **InitiateOAuth**: Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken connection property.

#### To add Xero data:

1. In ElastiCube Manager, click **Add Data** and then, **Xero**. The Connect to Xero window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. Xero is displayed in the Select Database list.
4. Click **OK**. Sisense connects to Xero and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**.  
The tables are displayed in the ElastiCube Manager.

# Xero Tables

Xero's RESTful APIs expose the following Xero tables that you can import into the ElastiCube Manager through the Sisense Xero connector:

## Available Tables

Name	Description
Accounts	Create, delete, and query accounts for a Xero organisation.
BankTransactions	Create, update, delete, and query bank transactions for a Xero organisation.
BankTransfers	Usage information for the operation BankTransfers.rsd.
ContactGroups	Create, update, and query contact groups for a Xero organisation.
Contacts	Create, update, and query contacts for a Xero organisation.
CreditNotes	Create, update, delete, and query credit notes for a Xero organisation.
Employees	Create, update, and query employees for a Xero organisation.
ExpenseClaims	Create, update, and query expense claims for a Xero organisation.
Invoices	Create, update, delete, and query invoices for a Xero organisation.

Name	Description
Items	Create, update, delete, and query items for a Xero organisation.
LinkedTransactions	Create, update, and query LinkedTransactions for a Xero organisation.
ManualJournals	Create, update, and query manual journals for a Xero organisation.
Payments	Create and query payments for a Xero organisation.
PurchaseOrders	Create, update, delete, and query purchase orders for a Xero organisation.
Receipts	Create, update, delete, and query receipts for a Xero organisation.
TaxRates	Usage information for the operation TaxRates.rsd.
<b>Payroll US Data Model</b>	
Employees	Create, update, and query employees for a Xero organisation.
PayItemsBenefits	Create, update, and query PayItems for a Xero organisation.
PayItemsDeductions	Create, update, and query PayItems for a Xero organisation.
PayItemsEarnings	Create, update, and query PayItems for a Xero organisation.
PayItemsReimbursements	Create, update, and query PayItems for a Xero organisation.
PayItemsTimeOff	Create, update, and query PayItems for a Xero organisation.

Name	Description
	organisation.
PayRuns	Create, update, and query payruns for a Xero organisation.
PaySchedules	Create, update, and query PaySchedules for a Xero organisation.
Paystubs	Create, update, and query pay stubs for a Xero organisation.
WorkLocations	Create, update, and query work locations for a Xero organisation.
<b>Payroll AUS Data Model</b>	
Employees	Create, update, and query employees for a Xero organisation.
LeaveApplications	Create, update, and query Leave Applications for a Xero organisation.
PayItemsDeductions	Create, update, and query PayItems for a Xero organisation.
PayItemsEarnings	Create, update, and query PayItems for a Xero organisation.
PayItemsLeave	Create, update, and query PayItems for a Xero organisation.
PayItemsReimbursements	Create, update, and query PayItems for a Xero organisation.
PayrollCalendars	Create, update, and query Payroll Calendars for a Xero organisation.

Name	Description
PayRuns	Create, update, and query payruns for a Xero organisation.
SuperFunds	Retrieve, add and update Payroll Super Funds in a Xero organisation.

## Available Views

Name	Description
<b>Accounting Data Model</b>	
AgedPayablesByContact	Query organisation data for a Xero organisation.
AgedReceivablesByContact	Query organisation data for a Xero organisation.
BalanceSheet	Query organisation data for a Xero organisation.
BankSummary	Query organisation data for a Xero organisation.
BrandingThemes	Query branding themes for a Xero organisation.
BudgetSummary	Query organisation data for a Xero organisation.
Currencies	Query currencies for a Xero organisation.
Journals	Query the line items in journals for a Xero organisation.
Organisation	Query organisation data for a Xero organisation.
Overpayments	Create and query Overpayments for a Xero organisation.
Prepayments	Create and query Prepayments for a Xero organisation.
ProfitAndLoss	Query organisation data for a Xero organisation.

Name	Description
RepeatingInvoices	Usage information for the operation RepeatingInvoices.rsd.
TrackingCategories	Query tracking categories for a Xero organisation.
TrialBalance	Query organisation data for a Xero organisation.
Users	Query users for a Xero organisation.
<b>Payroll US Data Model</b>	
TimeOffBalances	Retrieve, add and update an Employee's TimeOff balance in a Xero organisation.
Timesheets	Create, update, and query timesheets for a Xero organisation.
<b>Payroll AUS Data Model</b>	
LeaveBalances	Retrieve, add and update an Employee's Leave balance in a Xero organisation.

## Limitations

To access the AgedPayablesByContact and AgedRecievablesByContact tables, you need to specify a particular ContactID, the unique ID of the Contacts table, in the select query using a WHERE clause. This allows you to retrieve data for a specific contact. These are limitations of the Xero API, as these reports are specific to each contact. Here's an example:

```
SELECT due, paid, total FROM AgedPayablesByContact WHERE
ContactID='xyz'
SELECT duedate, reference FROM AgedRecievablesByContact
WHERE ContactID='1234'
```

Due this limitation, you cannot build an ElastiCube with these tables without custom SQL.

### **Daily Limit**

There is a daily limit of 1000 API calls against a single Xero organisation in a rolling 24-hour period.

In addition to the daily limit, a single access token can only be used up to 60 times in a rolling 60-second period.

### **Encountering a Rate Limit**

If you encounter a rate limit, the Xero API will return an HTTP 503 (Service Unavailable) error, with the following message: “oauth\_problem=rate limit exceeded”.

**Note:** If you encounter a rate limit, do not continue to make requests, as this may continue to add to your limitation. If necessary, you may need to queue requests.

### **Token Limitations**

Each access token will only last for 30 minutes. If you want longer access to the organization, you need the user to re-authorize your application.

# Connecting to YouTube Analytics

The Sisense YouTube Analytics connector is a standalone connector that allows you to import data from the YouTube Analytics API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect to the YouTube Analytics API through a connection string you provide Sisense in the ElastiCube Manager. The connection string is used to authenticate users who connect to the YouTube Analytics API. To obtain a connection string, you will need to create a YouTube Analytics app.

Once you have connected to YouTube Analytics, you can import a variety of tables from the YouTube Analytics API.

This page describes how to install the YouTube Analytics connector, how to connect to YouTube Analytics with a connection string, and what tables you can import into the ElastiCube Manager:

- [Installing the YouTube Analytics Connector](#)
- [Connecting to the YouTube Analytics REST API](#)
- [Adding YouTube Analytics Tables to your ElastiCube](#)
- [YouTube Analytics Tables](#)
- [Versions](#)

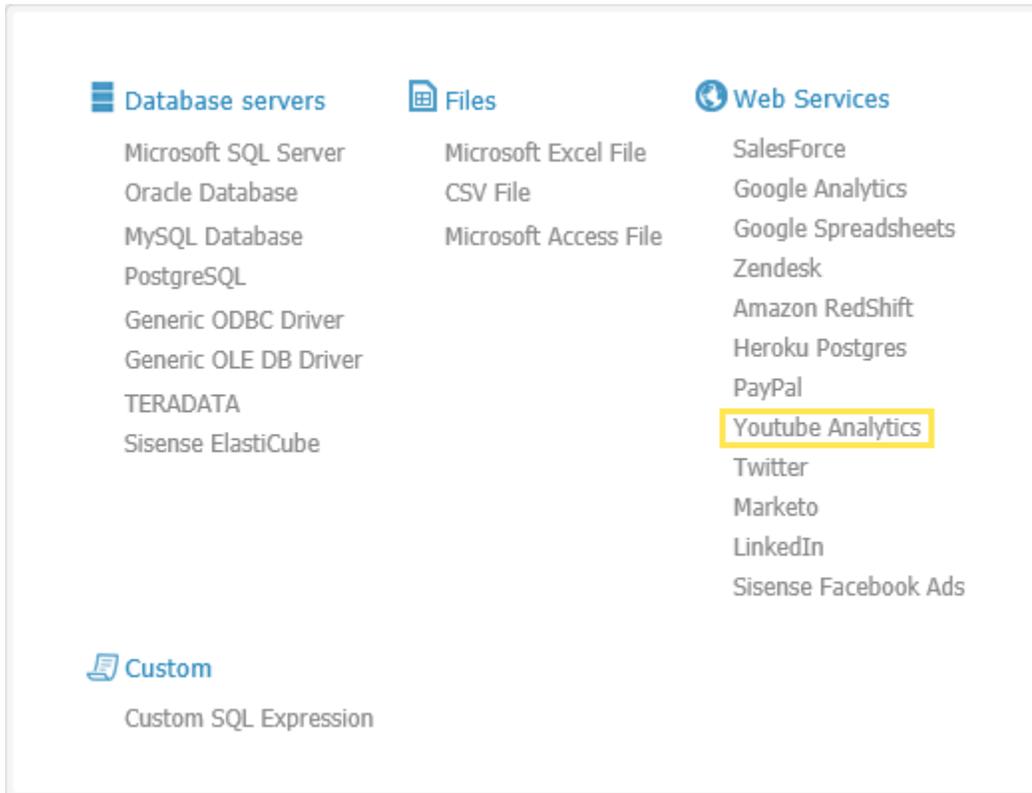
## Installing the YouTube Analytics Connector

Sisense provides the YouTube Analytics connector as a standalone connector that you can download and add to your list of default Sisense connectors.

**To install the YouTube Analytics connector:**

1. [Download](#) the YouTube Analytics installation file.
2. Open the installation file and click **Install**.

3. After the installation process is complete, click **Close**.
4. The YouTube Analytics connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.



## Connecting to the YouTube Analytics REST API

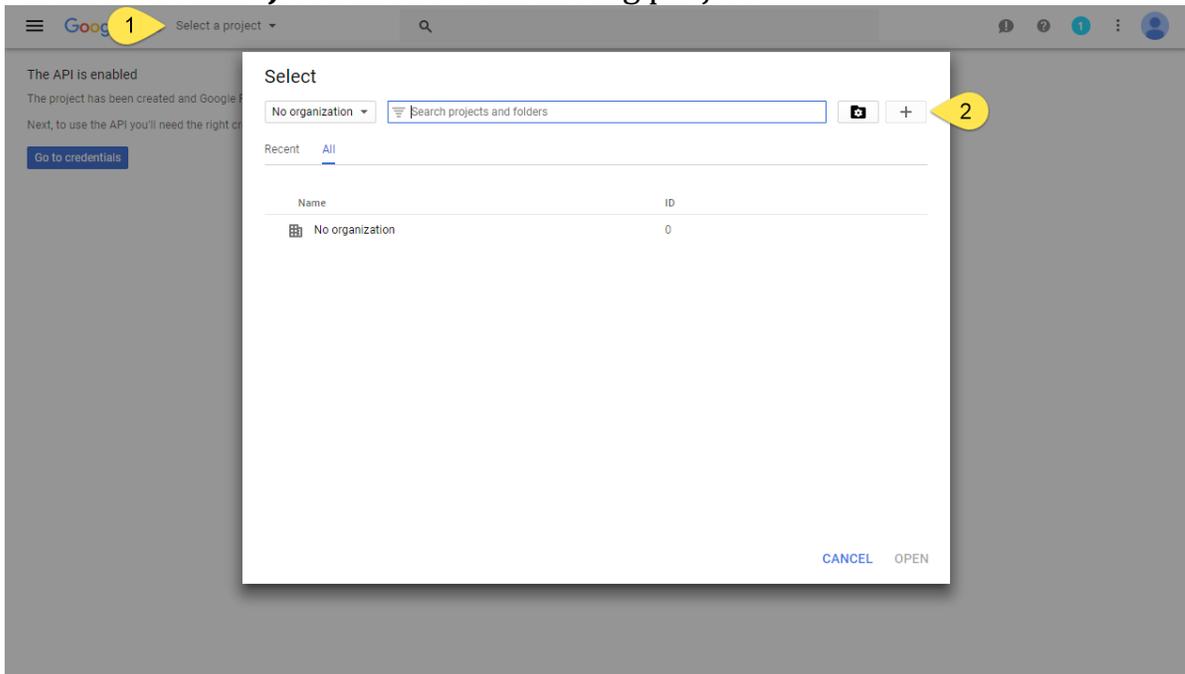
To access YouTube Analytics' REST API from the ElastiCube Manager, you must provide valid Oauth YouTube Analytics credentials through a connection string. These credentials are provided by YouTube Analytics when you create an application.

After you receive your credentials from YouTube Analytics, you can create the connection string and provide Sisense with it to connect to your data.

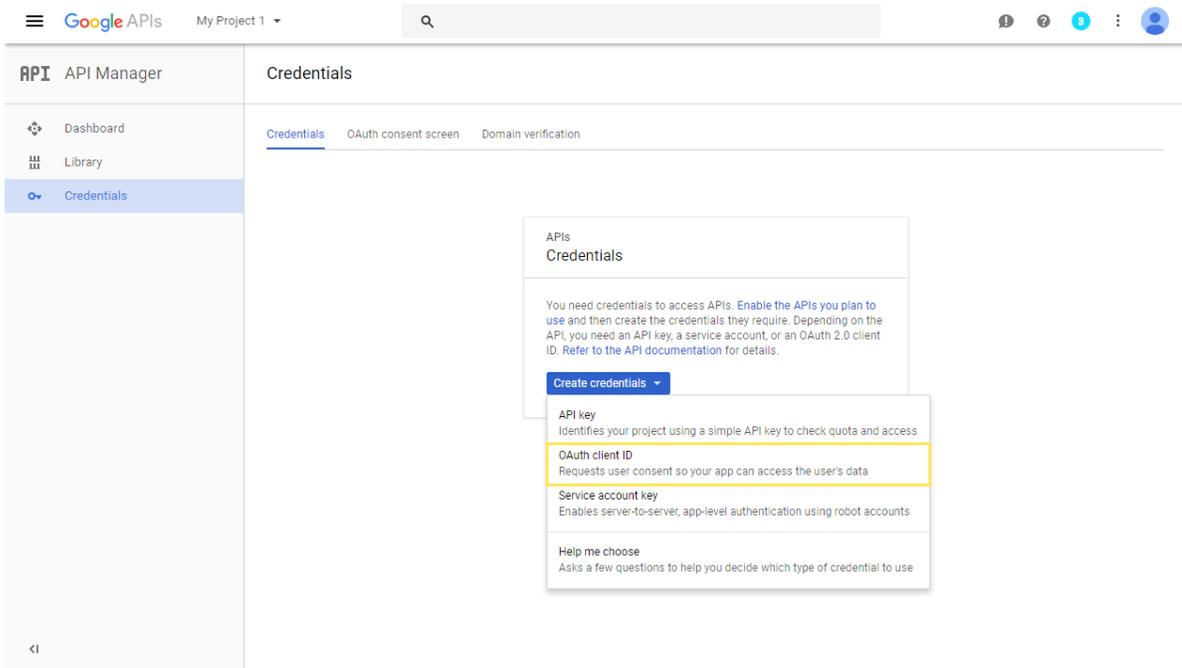
# Creating an App

You can follow the procedure below to register an app and obtain the OAuth client credentials, the OAuthClientId and OAuthClientSecret:

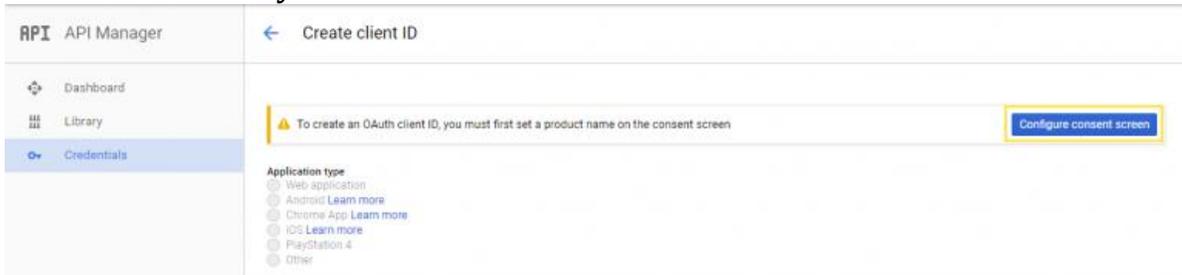
1. Log in to the [Google Developers Console](#).
2. Click **Create Project** or select an existing project.



3. In the API Manager, click **Credentials** > **Create Credentials** > **OAuth Client Id**.

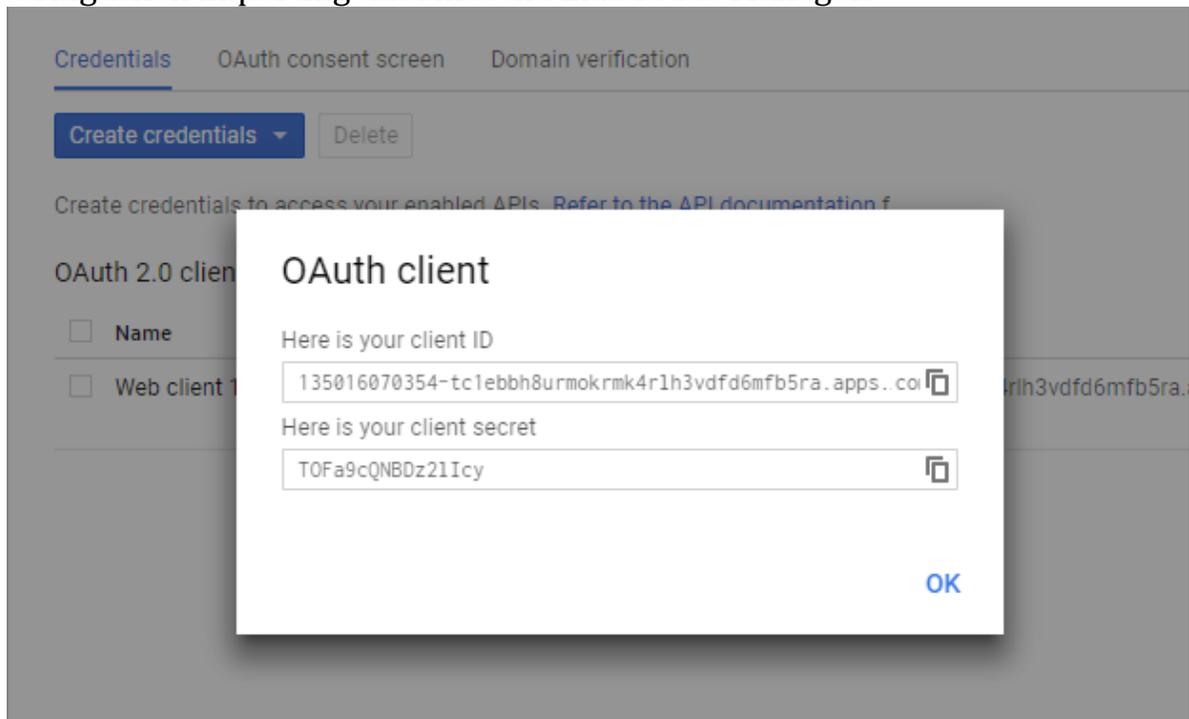


4. Click **Configure Consent Screen** to customize the information displayed to users when they connect.

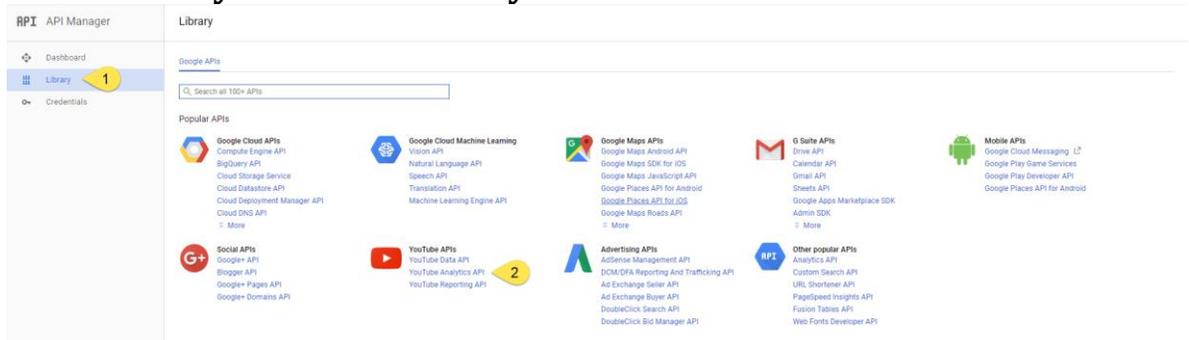


5. If you are connecting from a desktop application, click **Other** in the Application Type section. If you are connecting from a Web application, click the Web Application option. In the Authorized Redirect URIs box, enter the URL you want to be used as a trusted redirect URL, where the user will return with the token that verifies that they have granted your app access.
6. Click **Create**. The OAuthClientId and OAuthClientSecret are displayed. Save these credentials as they need to be passed to YouTube in the connection

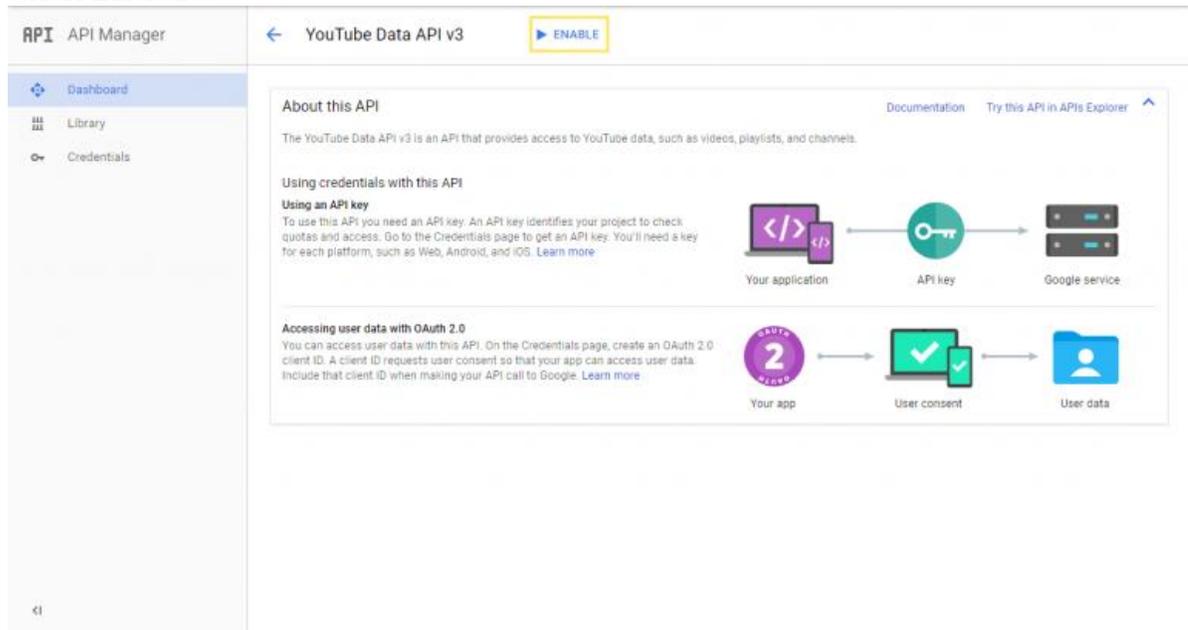
string when importing data into the ElastiCube Manager.



7. Click **OK**.
8. Select **Library > YouTube Analytics API**.



## 9. Click **Enable**.



# Adding YouTube Analytics Tables to your ElastiCube

Sisense uses connection strings to connect to YouTube Analytics and import data into the ElastiCube Manager. Each connection string contains a authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

After you have obtained the relevant credentials, you can create the connection used to connect to your YouTube account. The YouTube Analytics connection string has the following structure:

```
jdbc:DataSourceName:Property1=Value1;Property2=Value2;
```

The following is an example of a YouTube Analytics connection string:

```
jdbc:youtubeanalytics:OAuthClientId=xxxxxxxxxxxxxxxxxxxxxxxxxxxx  
;  
OAuthClientSecret=xxxxxxxxxxxxxxxx;InitiateOAuth=GETANDREFR  
ESH;CallbackURL=http://localhost;
```

The example above includes mandatory parameters you can provide in the connection. The required parameters are emphasized in bold.

### Mandatory Parameters

- **InitiateOAuth**: Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken connection property.
- **OAuthClientId**: Set this to the Client Id.
- **OAuthClientSecret**: Set this to the Client Secret.
- **CallbackURL**: Set this to http://localhost.
- **ChannelId**: Set this to the YouTube Analytics Channel (Profile). This can be set to the Id of the channel. If not specified, MINE will be used.
- **ContentOwnerId**: Set this to the Id of the content owner

### To add YouTube Analytics data:

1. In ElastiCube Manager, click **Add Data** and then, **YouTube Analytics**. The Connect to YouTube Analytics window is displayed.
2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. YouTube Analytics is displayed in the Select Database list.
4. Click **OK**. Sisense connects to YouTube Analytics and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**.  
The tables are displayed in the ElastiCube Manager.

## YouTube Analytics Tables

The Sisense YouTube Analytics connector allows you to import the following tables into the ElastiCube Manager.

Name	Description
GroupItems	Create, delete, and query Items that compose a Group.
Groups	Create, update, delete, and query YouTube Analytics Groups.
AdPerformanceReports	Query Ad Performance Reports. Ad Performance Reports provide impression-based metrics for ads that ran during

Name	Description
	<p>video playbacks. These metrics account for each ad impression, and each video playback can yield multiple impressions. To access AdPerformanceReports, in the WHERE clause, you should specify at least the video, group or a supported combination of uploaderType and claimedStatus. Optionally, the following filters can be specified: country, continent, subcontinent</p> <p>For dimensions, the 'adType' type is required. Optionally, 'day' dimension can be added. An example query similar to yours would be (notice, AdPerformanceReports view does not have a channel column, instead you have to use a video filter. This is a limitation of the YouTube Analytics API):</p> <pre>SELECT AdType, Day, GrossRevenue, StartDate, EndDate FROM AdPerformanceReports WHERE video='vid123' and continent = '019'</pre>
Demographics	Query YouTubeAnalytics Demographics grouped by age groups and gender.
Devices	Query YouTubeAnalytics Devices or Operating System statistics. This table allows you to aggregate viewing statistics based on the manner in which viewers reached your video content. For example, you can identify the number of views that occurred on mobile devices or game consoles.
EngagementReports	Query Engagement Reports info. For example, how many times a user shared a video.
TimeBasedReports	Query time-based info on views, subscribers, etc. For

Name	Description
	example, how many new subscribers were gained at a specific time.
TopVideos	Query the TopVideos for a channel or the TopPlaylists if you are a content owner. You can also use this view to query playback detail and geographic filters.
TrafficSources	Query TrafficSources. Statistics are based on the manner in which viewers reached your video and playlist content.

## Versions

Version	Release Date	Improvements
6.6.1.72	June 20, 2017	Support added for retrieving all videos based on channel <code>SELECT * FROM Videos WHERE ChannelId IN (SELECT Id FROM Channels)</code> Many of the tables have been updated to include additional inputs such as <code>OnBehalfOfContentOwner</code> (string) and <code>ForContentOwner</code> (boolean).
6.6.0.45	March 16, 2017	The <code>AdPerformanceReports</code> was enhanced to support more combinations of columns to query.
6.5.1.20	January 16, 2017	General release

# Connecting to YouTube

The Sisense YouTube connector is a standalone connector that allows you to import data from the YouTube API into the ElastiCube Manager. After you have downloaded and installed the connector, you can connect to the YouTube API through a connection string you provide in the ElastiCube Manager. The connection string is used to authenticate users who connect to YouTube's API. To obtain a connection string, you will need to create a YouTube app.

Once you have connected to YouTube, you can import a variety of tables from the YouTube API.

This page describes how to install the YouTube connector, how to connect to YouTube with a connection string, and what tables you can import into the ElastiCube Manager:

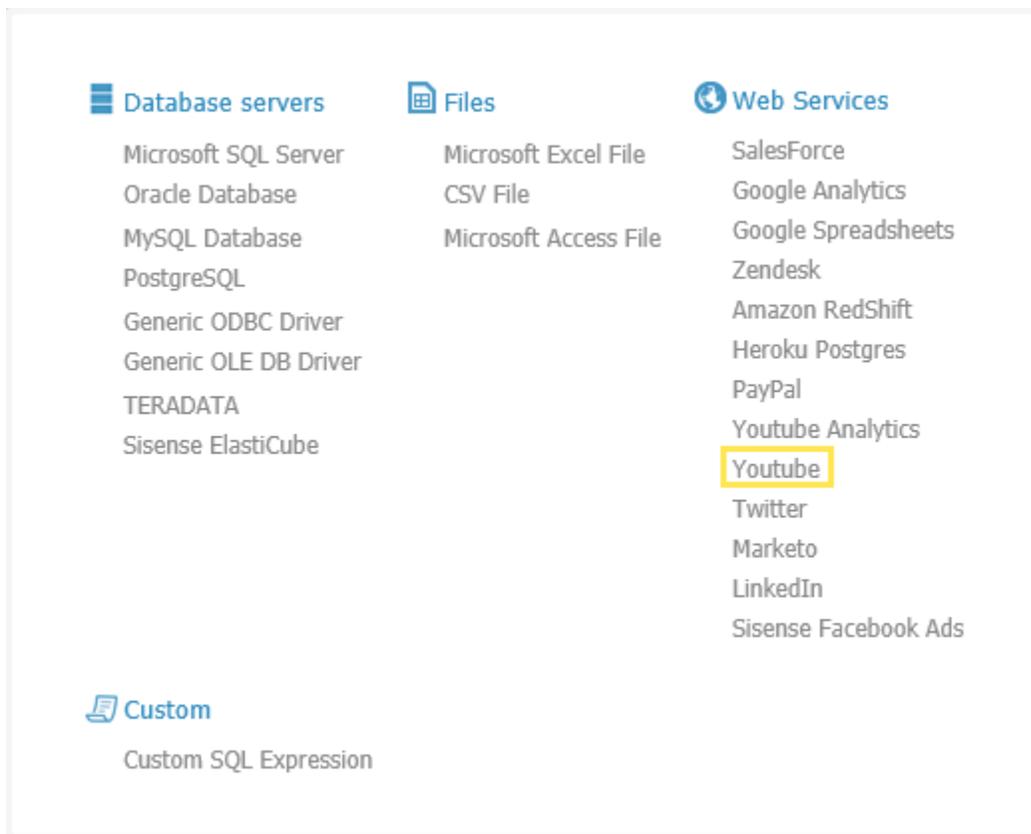
- [Installing the YouTube Connector](#)
- [Connecting to the YouTube REST API](#)
- [Adding YouTube Tables to your ElastiCube](#)
- [YouTube Tables](#)
- [Versions](#)

## Installing the YouTube Connector

Sisense provides the YouTube connector as a standalone connector that you can download and add to your list of default Sisense connectors.

### To install the YouTube connector:

1. [Download](#) the YouTube installation file.
2. Open the installation file and click **Install**.
3. After the installation process is complete, click **Close**.
4. The YouTube connector is displayed in the ElastiCube Manager under **Add Data > Web Services**.



## Connecting to the YouTube REST API

To access YouTube's REST API from the ElastiCube Manager, you must provide valid OAuth YouTube credentials through a connection string. These credentials are provided by YouTube when you create an application.

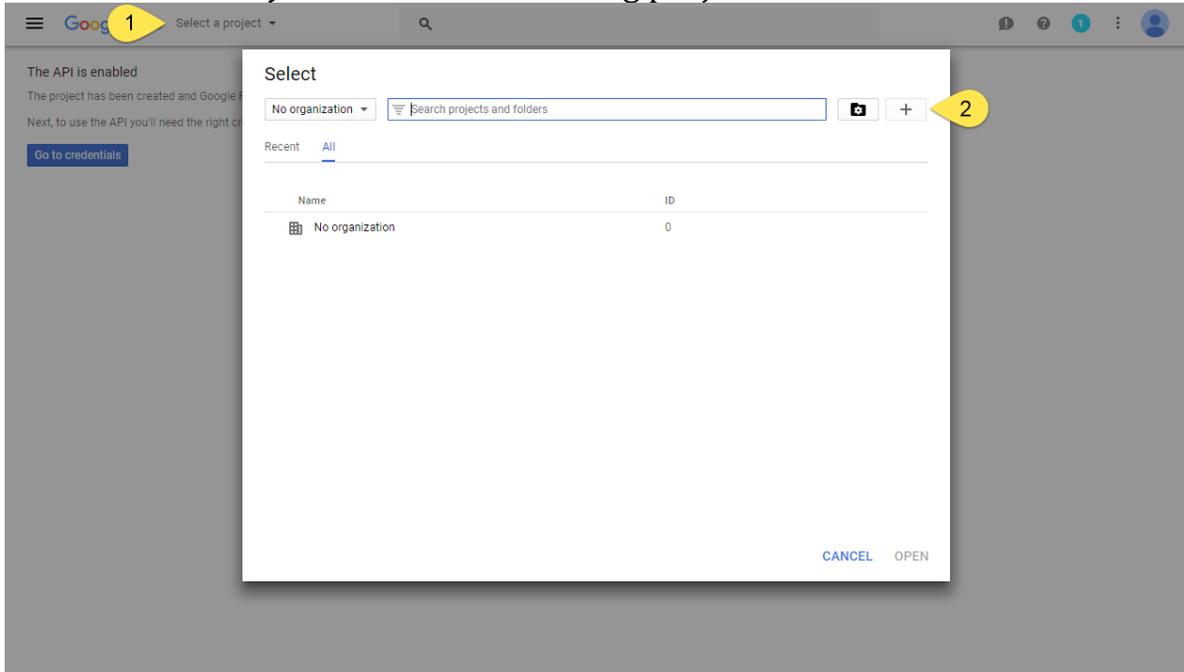
After you receive your credentials from YouTube, you can create the connection string and provide Sisense with it to connect to your data.

### Creating an App

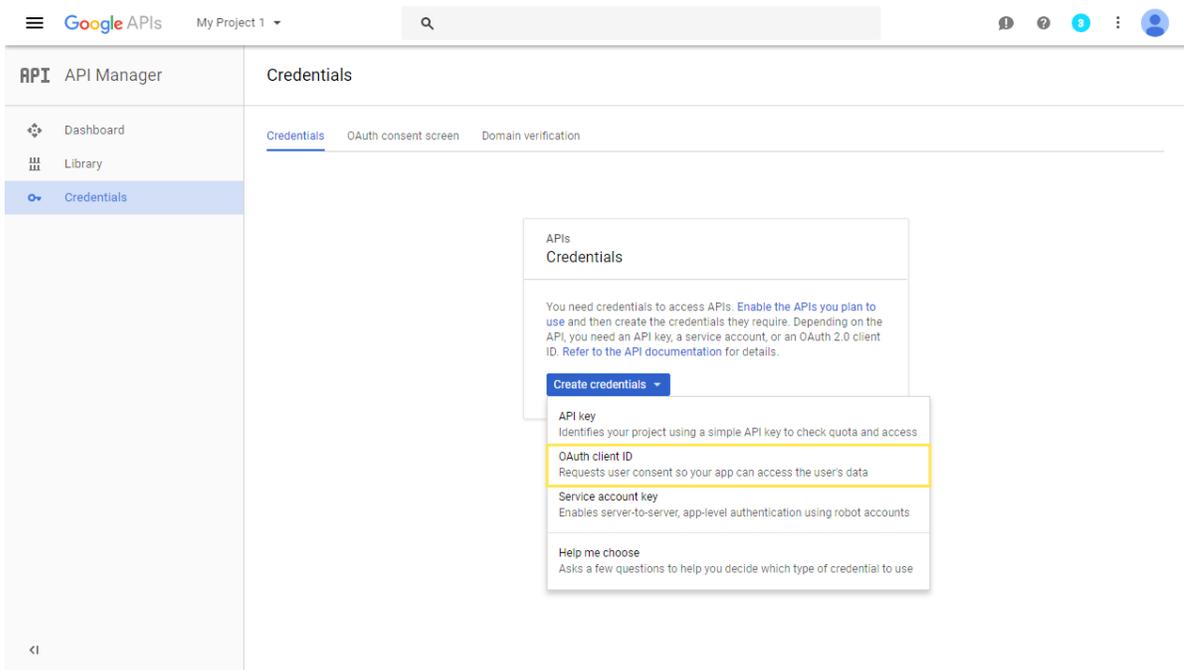
You can follow the procedure below to register an app and obtain the OAuth client credentials, the OAuthClientId and OAuthClientSecret:

1. Log in to the [Google Developers Console](#).

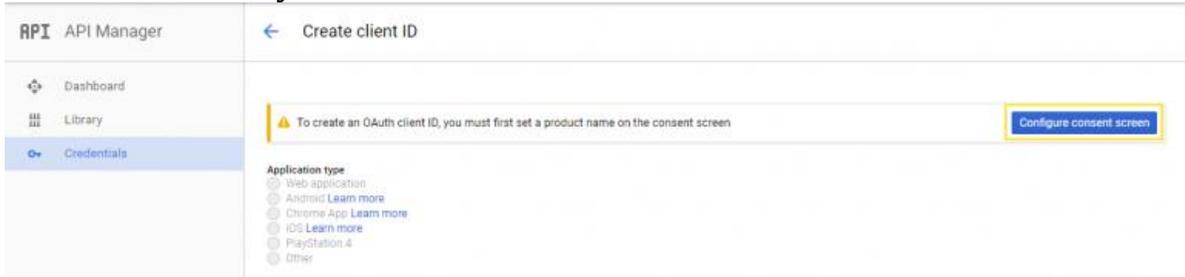
2. Click **Create Project** or select an existing project.



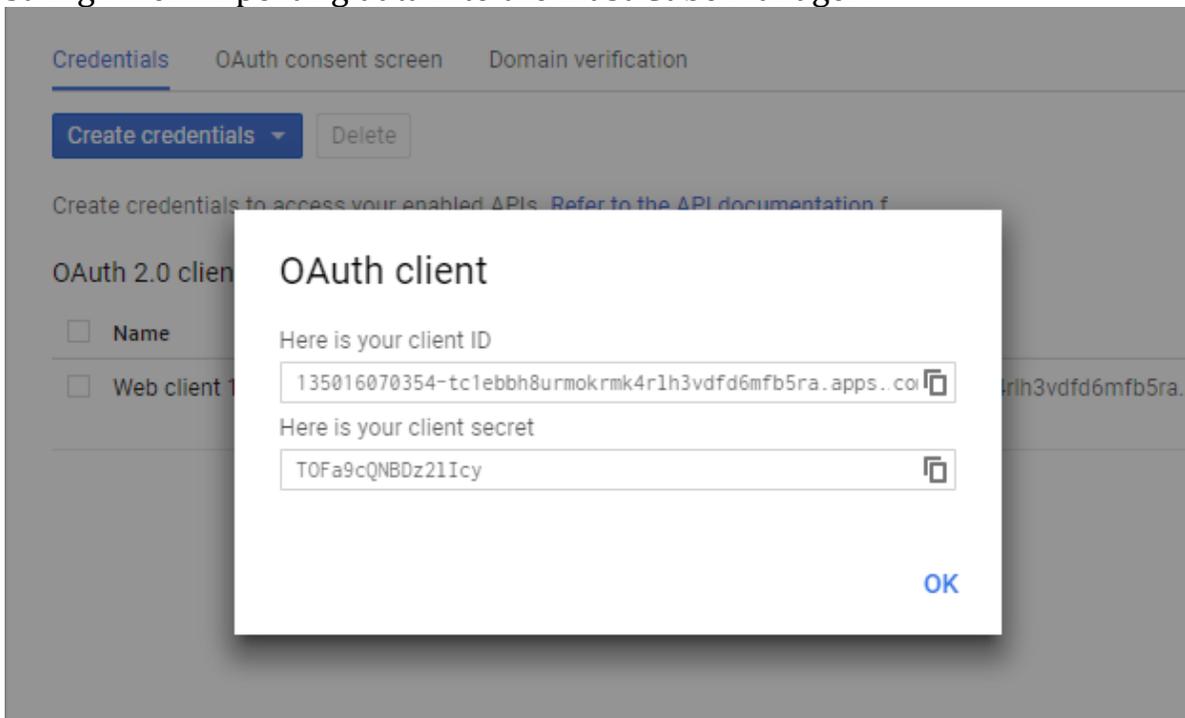
3. In the API Manager, click **Credentials** > **Create Credentials** > **OAuth Client Id**.



4. Click **Configure Consent Screen** to customize the information displayed to users when they connect.

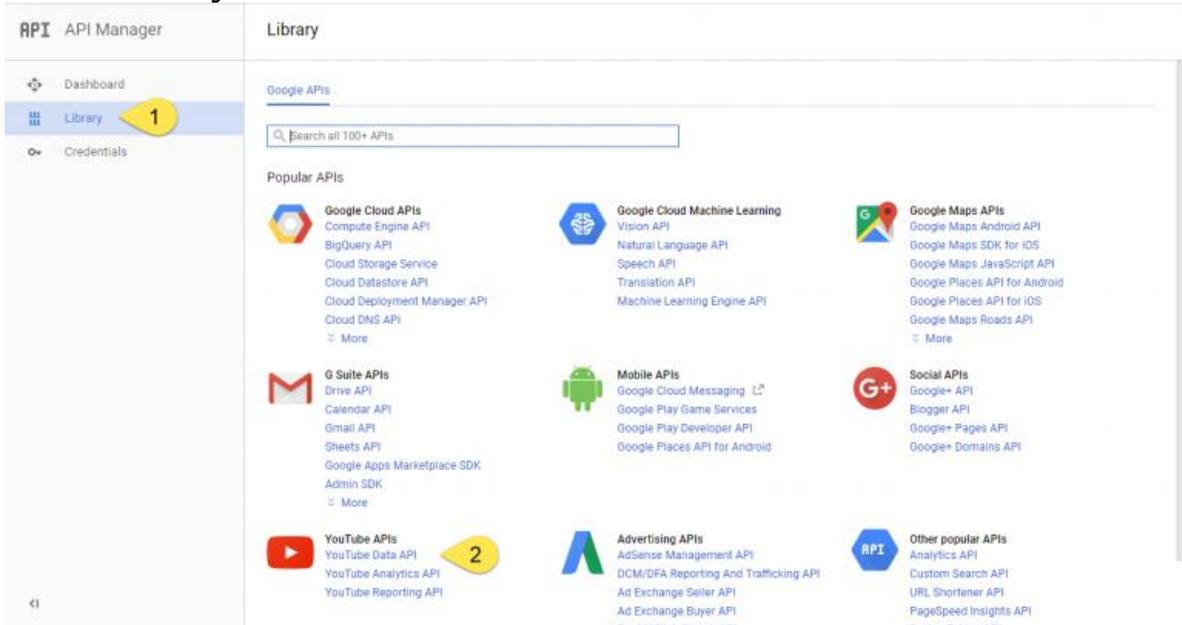


5. If you are connecting from a desktop application, click **Other** in the Application Type section. If you are connecting from a Web application, click the Web Application option. In the Authorized Redirect URIs box, enter the URL you want to be used as a trusted redirect URL, where the user will return with the token that verifies that they have granted your app access.
6. Click **Create**. The OAuthClientId and OAuthClientSecret are displayed. Save these credentials as they need to be passed to YouTube in the connection string when importing data into the ElastiCube Manager.

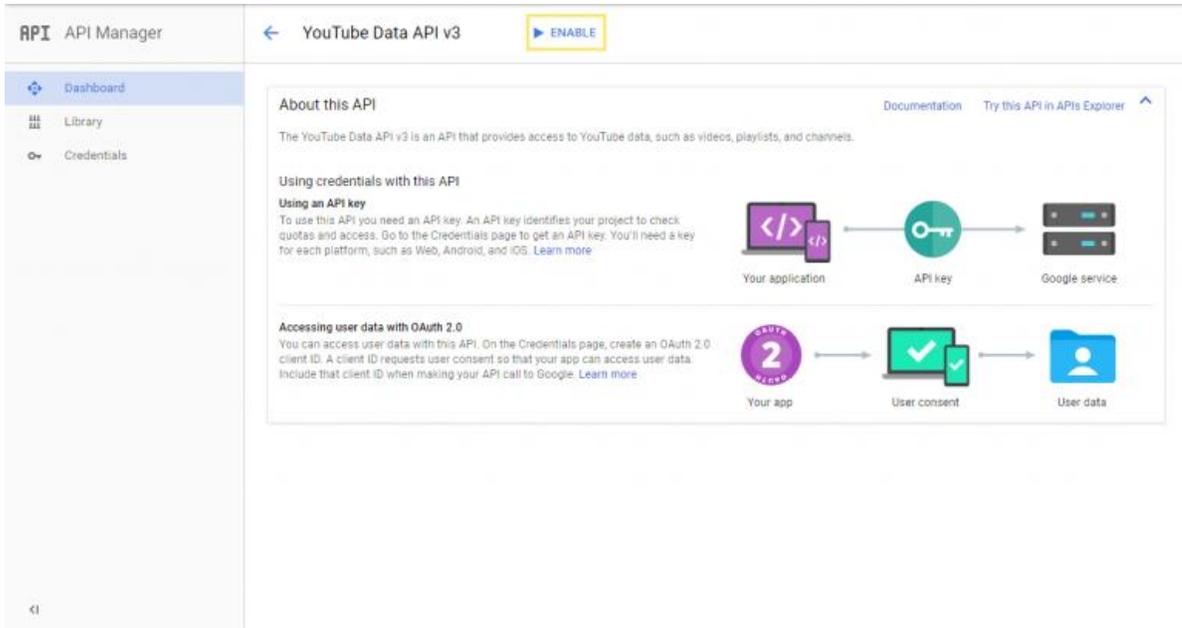


7. Click **OK**.

## 8. Select **Library** > **YouTube Data API**.



## 9. Click **Enable**.



## Creating a Connection String

Sisense uses connection strings to connect to YouTube and import data into the ElastiCube Manager. Each connection string contains authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

**Note:** When switching accounts, you must sign out of your account and then sign it with the new account.

After you have obtained the relevant credentials, you can create the connection used to connect to your YouTube account. The YouTube connection string has the following structure:

```
jdbc:DataSourceName:Property1=Value1;Property2=Value2;
```

The following is an example of a YouTube connection string:

```
jdbc:youtube:OAuthClientId=112345383774486;OAuthClientSecret=064c524478d712534b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.10;CallbackURL=http://localhost/33333;
```

The example above includes mandatory parameters you can provide in the connection. The required parameters are emphasized in bold.

### Mandatory Parameters

- **InitiateOAuth:** Set this to GETANDREFRESH. You can use InitiateOAuth to avoid repeating the OAuth exchange and manually setting the OAuthAccessToken connection property.
- **OAuthClientId:** Set this to the Client Id.
- **OAuthClientSecret:** Set this to the Client Secret.
- **CallbackURL:** Set this to http://localhost.

Alternatively, you can keep the OAuthClientId and the OAuthClientSecret parameters empty to redirect your users to the YouTube authentication page where they can log in to their accounts to access their tables without creating a YouTube app.

## Switching between Accounts

When you connect to the YouTube data source, Sisense saves your OAuth values in the file OAuthsettings.txt file located at

...\Users\xxx\AppData\Roaming\CData\YouTube Data Provider on your Sisense server. To connect to the YouTube data source with another user on the same

machine, you must delete the OAuthsettings.txt file. Sisense will then generate a new file for that user.

Another option to support multiple users is to define the location and file name of an OAuthsettings file for each unique user in your connection string through the OAuthSettingsLocation parameter. When each user connects to the data source, Sisense generates the OAuth file with the file name you specify in the location you define. In the examples below, two users are allowed to access the YouTube data source and for each user, Sisense generates a file that contains that user's OAuth values in the location defined in the string.

```
jdbc:youtube:OAuthSettingsLocation=C:\youtube\auth\john.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/33333;  
jdbc:youtube:OAuthSettingsLocation=C:\youtube\auth\sally.txt;OAuthClientId=11276856774486;OAuthClientSecret=064c70d78567jm2b7e7e4224fad;InitiateOAuth=GETANDREFRESH;Version=2.8;CallbackURL=http://localhost/33333;
```

In the example above, two OAuth files are created, one for John and one for Sally in the location C:\youtube\auth\.

This is useful if you support many users who each need to access the YouTube data source.

## Adding YouTube Tables to your ElastiCube

Sisense uses connection strings to connect to YouTube and import data into the ElastiCube Manager. Each connection string contains a authentication parameters that the data source uses to verify your identity and what information you can export to Sisense.

**To add YouTube data:**

1. In ElastiCube Manager, click **Add Data** and then, **YouTube**. The Connect to YouTube window is displayed.

2. In **Datasource Connection String**, enter your connection string.
3. Click **Connect to Server**. YouTube is displayed in the Select Database list.
4. Click **OK**. Sisense connects to YouTube and displays a list of tables available for you to import.
5. Select the relevant tables and click **Add**.  
The tables are displayed in the ElastiCube Manager.

## YouTube Tables

The Sisense YouTube connector allows you to import the following tables into the ElastiCube Manager.

Name	Description
Activities	Query YouTube Activities. An Activity resource contains information about an action that a particular Channel, or User, has taken on YouTube.
Captions	Query YouTube Captions. A Caption resource represents a YouTube caption track.
Channels	Query YouTube Channels.
ChannelSections	Query YouTube ChannelSections. A ChannelSection

Name	Description
	contains information about a set of videos that a channel has chosen to feature.
Comments	Query YouTube Comments.
CommentThreads	Query YouTube CommentThreads. A CommentThread record contains information about a YouTube comment thread, a top-level comment and replies, if any exist, to that comment.
GuideCategories	Query YouTube GuideCategories. A GuideCategory resource identifies a category that YouTube algorithmically assigns based on a content of a channel or other indicators, such as the popularity of the channel.
Languages	Query YouTube i18nLanguages. An i18nLanguage resource identifies an application language that the YouTube website supports. The application language can also be referred to as a UI language.
PlaylistItems	Query YouTube PlaylistItems. A PlaylistItem resource identifies another YouTube entity, such as a Video, that is included in a Playlist. In addition, the PlaylistItem record contains details about how that entity is used in that Playlist.
PlayLists	Query YouTube PlayLists. A Playlist is a collection of videos that can be viewed sequentially and shared with other users.
Regions	Query YouTube i18nRegions. An i18nRegion resource

Name	Description
	identifies a geographic area that a YouTube user can select as the preferred content region. The content region can also be referred to as a content locale.
Subscriptions	Query YouTube User Subscriptions. A Subscription notifies a User when new Videos are added to a Channel, or when another user takes one of several actions on YouTube, such as uploading a Video, rating a Video, or commenting on a Video.
VideoAbuseReportReasons	Query YouTube VideoAbuseReportReasons. A VideoAbuseReportReason resource contains information about a reason that a Video would be flagged for containing abusive content.
VideoCategories	Query YouTube VideoCategories. A VideoCategory resource identifies a category that has been or could be associated with uploaded Videos.
Videos	Query YouTube Videos.

## Versions

Version	Release Date	Improvements
6.6.1.72	June 20, 2017	Support added for retrieving all videos based on channel <code>SELECT * FROM Videos WHERE ChannelId IN (SELECT Id FROM Channels)</code> Many of the tables have been updated to include additional inputs such as <code>OnBehalfOfContentOwner</code> (string) and <code>ForContentOwner</code> (boolean).

Version	Release Date	Improvements
6.6.0.44	March 16, 2017	PublishedAt column presented correctly Views are now more readable with important fields reordered and unnecessary fields removed
6.5.1.20	January 16, 2017	General release

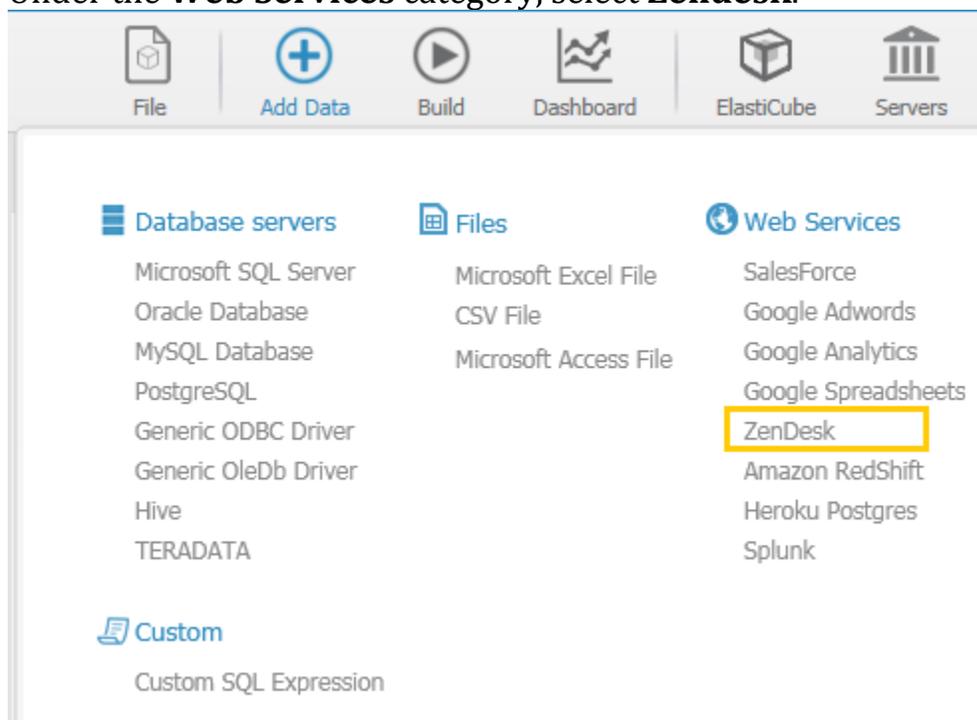
# Connecting to Zendesk

The ElastiCube Manager enables easy and quick access to tables and views contained within Zendesk databases.

**Note:** Only non-archived tickets are supported.

The steps below describe how to connect to this type of data source.

1. Click **Add Data** in the top menu of the ElastiCube Manager.
2. Under the **Web Services** category, select **Zendesk**.



3. You will be prompted to enter the following information:
  - **Web Address:** Enter the web address where your Zendesk deployment is hosted, for example: `http://support.yourcompany.com`.  
**Note:** You must include `http://` in your web address.
  - **Login details:** Either use your Windows login details if they are configured with Zendesk, or alternatively enter the username and password used to connect to Zendesk.

4. Click **Connect to Server**.

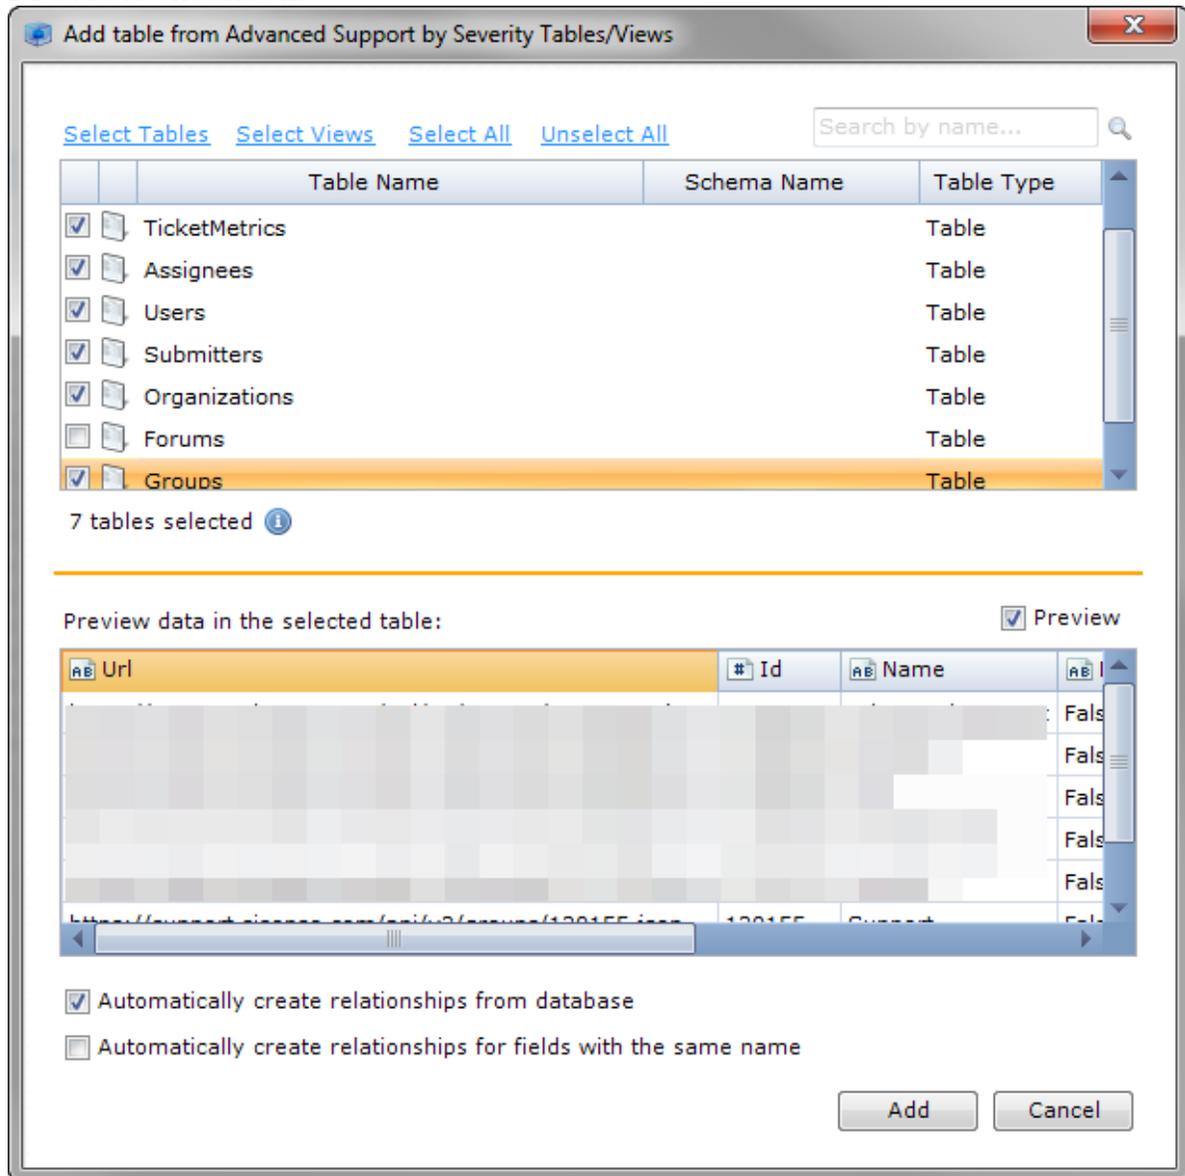
A list of available views will appear in the list box below.

The screenshot shows a dialog box titled "Connect to ZenDesk". It features a "Web Address:" dropdown menu with "https://" selected. There are two radio button options: "Windows Authentication" (unselected) and "Use The Following User Name & Password:" (selected). Below the second option are "User Name" and "Password" text input fields. A "Connect to Server" button is positioned below these fields. A horizontal line separates this section from the "Select View:" dropdown menu, which currently displays "Advanced Support by Severity". At the bottom of the dialog are "OK" and "Cancel" buttons.

5. Select the relevant view you want to work with and click **OK**.

All tables associated with the Zendesk view will appear in a new window. To preview data contained in a particular table, highlight the table or view in the list and click the preview pane below. To preview the table, select the

## Preview checkbox.



6. Select the checkbox next to each table or view you want to use. Existing relationships between tables can be automatically replicated in the ElastiCube by selecting the **Automatically create relationships from database** checkbox. Likewise fields with similar names can be linked by selecting the **Automatically create relationships for fields with the same name** checkbox.
7. Once all relevant tables are selected, click **Add**.

# Deleting ElastiCubes

When you no longer want to use ElastiCube, you can delete the ElastiCube from the Data page.

Deleting an ElastiCube removes it from the server and the ElastiCube's folder saved locally on your hard drive. After an ElastiCube has been deleted, it cannot be restored.

If your ElastiCube has any dependencies, or entities such as dashboards, that are currently connected to the ElastiCube, some of these entities will also be deleted and some you will be prompted to remap to another data source. Entities that are deleted when you delete an ElastiCube include:

- Data security rules
- Build alerts
- Drill hierarchies

Entities that are not deleted, but should be remapped to another data source include:

- Dashboards
- Widgets (from other dashboards)
- Pulse Alerts

Entities that you can remap to another data source are clickable links that when clicked, display a list of all the entities that need to be remapped.



## Delete ElastiCube

"Sample Lead Generation" on LocalHost will be deleted.

The following entities will not be deleted, but should be mapped to a replacement data source:

- [Dashboards](#)

Select new data source or "None" to manually update them later.

 None ▼

Delete

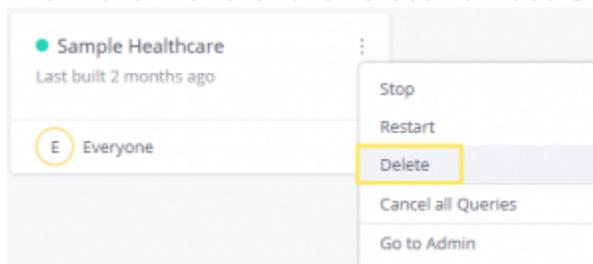
Cancel

You can choose to remap the entities to another data source manually or select the new data source from the list on the bottom of the Delete ElastiCube window. If you do not remap the entities, the dependencies will not work until you apply a new data source.

**Note:** Any temporary folders that were saved locally during failed builds are not deleted when the ElastiCube is deleted and must be removed manually.

### To delete an ElastiCube:

1. In the Sisense Web Application, click **Data**.
2. From the menu of the relevant ElastiCube, select **Delete**.



3. Click **Yes** to confirm you want to delete the ElastiCube.

OR

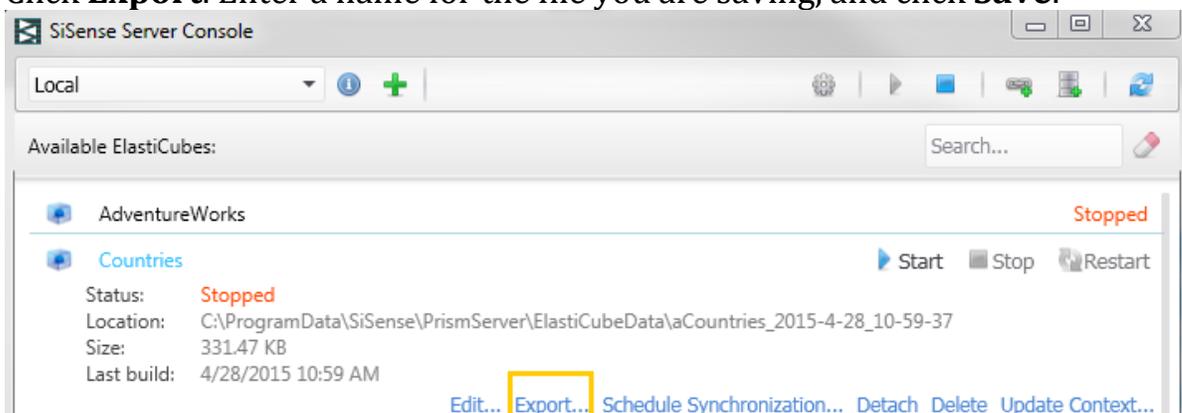
If your ElastiCube has any dependencies, remap those entities to a new ElastiCube or select None to remap them later.

# Importing and Exporting ElastiCube Data

**Note:** This feature is currently available in the desktop version of the ElastiCube Manager. This feature will be migrated to the web-based ElastiCube Manager in the near future. For more information, see [Sisense V7.0 Backwards Compatibility](#). ElastiCube data can be transferred to different machines. This is useful when you are deploying on a new machine or need to transfer the data to a different server. The transferred data is stored in a compressed ecdData file. The file contains all of the information contained in an ElastiCubeData folder (including the source data), and is used by the Sisense Server Console to build a new ElastiCube.

## To export data:

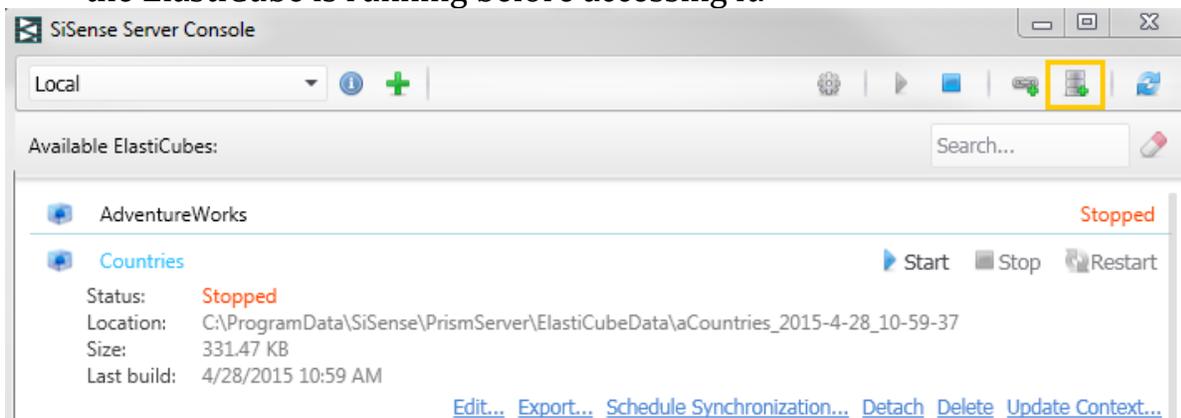
1. Open the Sisense Server Console.
2. Click on the relevant ElastiCube.
3. Click **Stop** to stop the specific ElastiCube (not the entire server).
4. Click **Export**. Enter a name for the file you are saving, and click **Save**.



## To import data:

1. Open the Sisense Server Console.
2. Click on the import ecdData file icon .
3. Select the location of the ecdData file and click **Open**.

4. A window will appear showing the import progress. Once complete, the new ElastiCubes will be added to the list of existing ElastiCubes. Make sure the ElastiCube is running before accessing it.

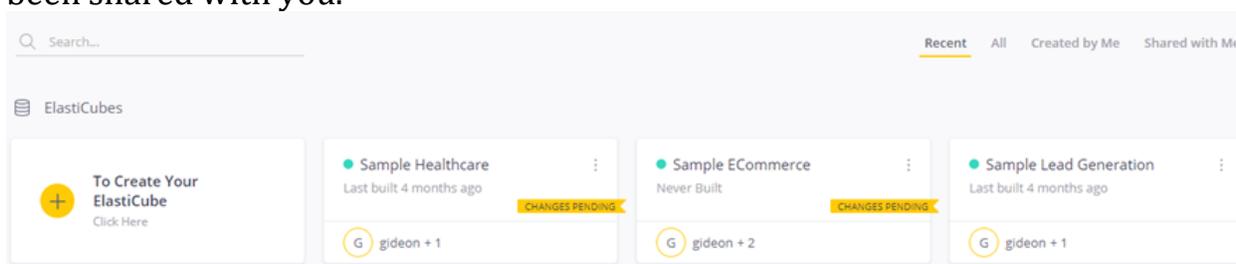


**Note:** You can also transfer data by copying the folder where Sisense stores the data locally, instead of exporting and importing the compressed file.

# Managing ElastiCubes

From the Data page of ElastiCube Manager, you can manage your ElastiCubes and view their current states.

The Data page contains a list of all the ElastiCubes you have created or that have been shared with you.



To locate ElastiCubes easily, on the right side of the page, you can view ElastiCubes you recently worked on, that were created by you, or that were shared with you. In addition, on the left side, you can locate an ElastiCube through the **Search** field.

Each ElastiCube tile contains information about the status of the ElastiCube, its owner, and who its shared with it. You can open your ElastiCube by clicking  $\vdots$  on its tile or you can manage your ElastiCubes by clicking and selecting one of the following options:

- **Stop:** Click to stop the ElastiCube. ElastiCubes that have been stopped cannot be queried. This is useful for troubleshooting ElastiCubes and issues with your dashboard.
- **Restart:** Click to restart an ElastiCube that you have stopped.
- **Dependencies:** Click to see what entities are using the ElastiCube. Entities include data security rules, drill hierarchies, dashboards, widgets (from other dashboards), and Pulse alerts.
- **Schedule Build:** Click to schedule builds to automatically synchronize with the underlying data source to ensure all data is up to date. See [Scheduling Builds](#) for more information.

- **Delete:** Click to delete an ElastiCube. If the ElastiCube does not have any entities dependent (dependencies) on it, then you will be asked to confirm that you want to delete the ElastiCube. Otherwise, you will be prompted to remap your dependencies to another data source. See [Deleting ElastiCubes](#) for more information.
- **Cancel all Queries:** Click to cancel all current queries to the ElastiCube.
- **Go to Admin:** Opens the Data Source page in the Admin section of the Sisense Web Application where you can manage your ElastiCube.

Some features and functionality supported in the desktop version of the ElastiCube Manager are not supported in ElastiCube Manager currently. These features will be added in the near future. In the meantime, you can find workarounds for missing features at [Sisense V7.0 Backwards Compatibility](#).

### Related Topics

- [Sisense Server Console](#)
- [Scheduling Data Loads](#)
- [Importing and Exporting ElastiCube Data](#)

# Scheduling Data Loads

After you build an ElastiCube, you can schedule future builds to automatically synchronize with the underlying data source to ensure all data is up to date.

There are two options you can choose when scheduling a build, by time or by intervals. Builds scheduled by time occur according to the days and hour you define. You can configure Sisense to build an ElastiCube automatically any or every day of the week at the same hour.

Builds scheduled by intervals occur after the defined intervals passes. You can set a single interval for an ElastiCube.

**Note:** Only one option can be selected at a time. The last configuration you set is the configuration that will be used.

If you want to cancel a scheduled build, select **Disable** and click **Save**.

By default, scheduled builds are Entire builds unless any table within the ElastiCube is set to Accumulative or a table is set to be ignored.

Your ElastiCube's build configuration determines how the scheduled build is to be performed. For more information about how to build ElastiCubes click [here](#).

## To schedule a build:

1. In the Data page, open the menu of the relevant ElastiCube and select **Schedule Build**. The Schedule Build window is displayed.

Schedule Build

ElastiCube: Tutorial Last Build: 03/04/2018 1:28 PM

Choose one of the scheduling options:

Disable Daily Interval

Select the days and time to build your ElastiCube automatically and click Save.

Monday  Tuesday  Wednesday  Thursday  
 Friday  Saturday  Sunday

12:00 AM

Save Cancel

2. Select the relevant scheduling option:

**Disable:** Disables a schedule build.

**Daily:** Select the relevant days and define the hour when the build is to occur. For each selected day, the build will occur at the same hour defined.

**Interval:** Select the time in between automatic builds. The first automatic build begins after you save your settings and the interval passes.

3. Click **Save**.

# Sisense Server Console

**Note:** This feature is currently available in the desktop version of the ElastiCube Manager. This feature will be migrated to ElastiCube Manager in the near future. For more information, see [Sisense V7.0 Backwards Compatibility](#).

You can use the Sisense Server Console to perform various management tasks on your ElastiCubes, including:

- [Scheduling Data Loads](#)
- [Importing and exporting data](#)
- [Attaching and detaching directories](#)
- [Changing the Location of your ElastiCube Data Storage](#)
- [Canceling Running Queries](#)

## To open the Sisense Server Console:

- Right-click the systems tray and double-click the console icon .  
or  
Click on the Windows start menu, and select **All Programs** > Sisense > **Sisense Server Console**.

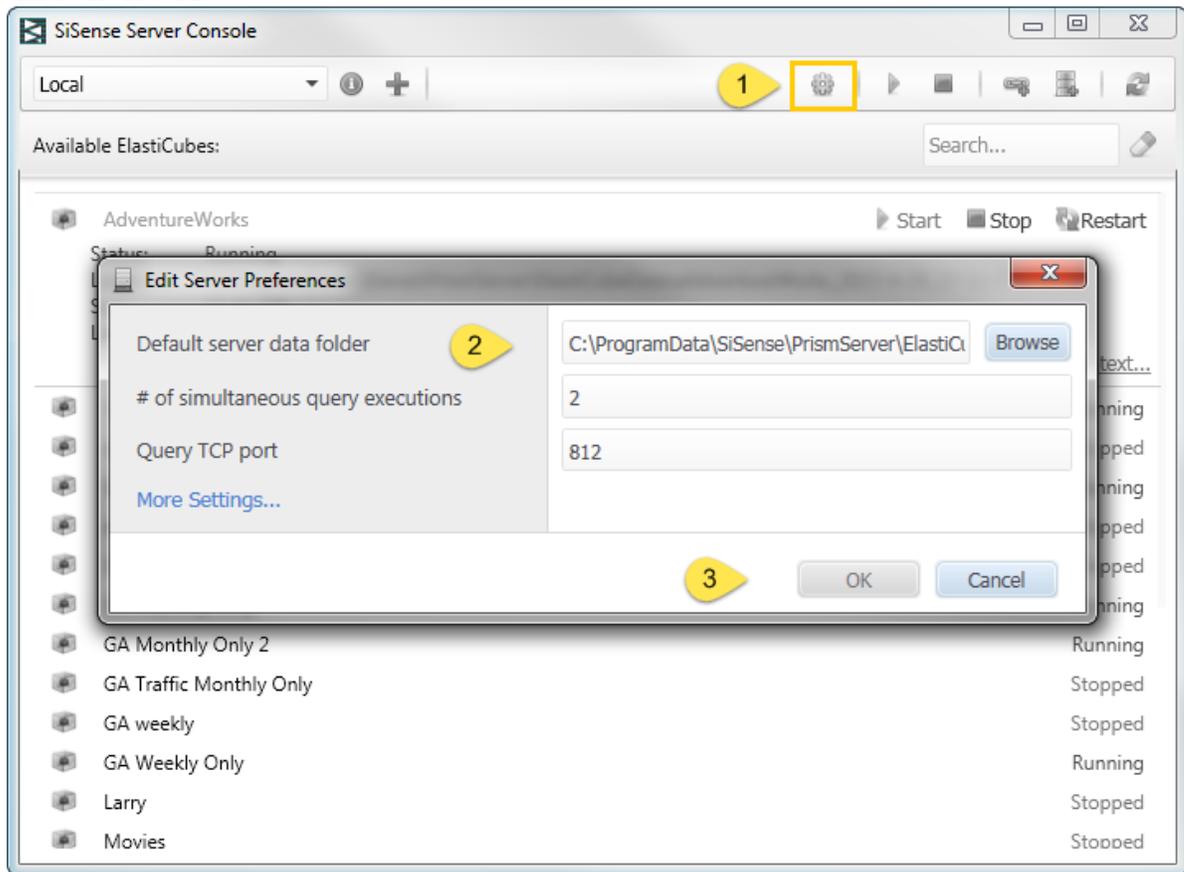
ElastiCube data is stored in a default folder, which you can change in the Sisense Server Console.

## Changing the Location of the ElastiCube Data Storage

### To change the folder:

1. From the Windows start menu, open the Sisense Server Console.
2. In the top panel, click on the cog icon to view/change server preferences.
3. Click **Browse** next to **Default server data folder**.

4. Select the location of the folder where you want to store ElastiCube data, and click **OK**.



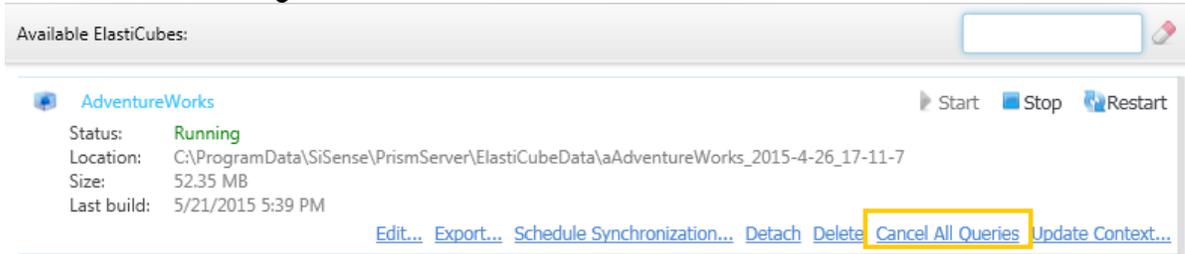
## Canceling Running Queries

You can cancel all running queries in your ElastiCube to free up resources. This is useful in cases where your ElastiCube is consuming too many resources, and as a result, causing the Sisense Web Application to be unresponsive.

### To cancel all queries:

1. Click on the Windows start menu, and select **All Programs > Sisense > Sisense Server Console**.
2. Click the relevant ElastiCube to view the menu panel.

### 3. Click **Cancel All Queries**.



# ElastiCube for Advanced Business Scenarios

This section provides examples of more advanced scenarios, and the recommended methods for implementing the required business logic. The examples are categorized into the following data manipulation methods: Integrating, Formatting and Enhancing.

## Integrating Data

Integrate and merge data from different sources into a single ElastiCube structure by identifying common keys between the different tables. Proper planning is important for merging the data; on the one hand, you need to avoid creating unnecessary relationships, and on the other hand, make sure you don't have any many to many relationships. Examples include:

- [Creating a Common Date Selection](#): Create a common date field from multiple date sets (from multiple data sources), and still keep the ability to use each original date field individually.
- [Financial Reporting](#): Bring in an additional data source to help analyze data from transaction systems. For example, Financial GL data will include all transactions but may not have all the income statement or balance sheet reporting definitions.
- [Looking Up Values](#): Look up a value from one table and bring it into another table. For example, knowing how much a marketing campaign costs versus the sales opportunity amount is an important KPI to measure.

# Reformatting Data

Reformat field data to free space, and make fields more readable and usable. For example, convert a date field to numeric, or reduce the precision of real numbers.

You can reformat fields within the ElastiCube using a custom SQL expression.

- [Numeric Representation of Date Fields](#): Create a date table that is represented by a numeric representation instead of a date field to improve the query performance, as well as provide more flexibility, including the ability to filter a date range.

# Enhancing Data

[Enhance data](#) by adding attributes/records that did not exist in the original data source.

- [Calculating Derived Facts](#): Derived Facts are additional facts that we calculate while importing or delivering the data.
- [Calendar vs. Fiscal Year](#) : Align a fiscal calendar with a Gregorian calendar.
- [Time Zone Conversion](#): Use a source table to convert dates and times from different time zones into a uniform data set.
- [Currency Conversion](#): Convert one currency into another using custom fields and a currency exchange rate table.
- [Current vs. Previous Period for Specific Date Range](#): Compare data such as sales between a current period and a past period.
- [Calculating the Number of Open Orders per Day](#): Check open sales orders where the order has been placed but has not yet been delivered.
- [Slowly Changing Dimensions](#): Transactional data does not usually change, however the data that describes the associated dimensions may change. See how to manage dimensions that may be updated with new values within the data warehouse at different points in time.

# Enhancing Data

**Note:** The image on this page were taken in the desktop version of the ElastiCube Manager, however, the same principles described on this page also apply to the ElastiCube Manager.

The following examples explain how to add attributes and/or records that did not exist in the data source. Examples include:

- [Calculating Derived Facts](#)
- [Calendar vs. Fiscal Year](#)
- [Time Zone Conversion](#)
- [Currency Conversion](#)
- [Current vs. Previous Period for Specific Date Range](#)
- [Calculating the Number of Open Orders per Day](#)
- [Slowly Changing Dimensions](#)

## Calculating Derived Facts

### Business Case

Derived Facts are additional facts that we calculate while importing or delivering the data. For example:

$$\text{Amount} = \text{Qty} * \text{UnitPrice}$$

$$\text{Profit} = \text{Amount} - \text{Cost}$$

$$\text{Inventory Ratio} = \frac{\text{Sum}(\text{Qty Sold})}{\text{Sum}(\text{Qty Ordered})}$$

$$\text{AVG Price} = \frac{\text{Sum}(\text{Qty} * \text{UnitPrice})}{\text{Sum}(\text{Qty})}$$

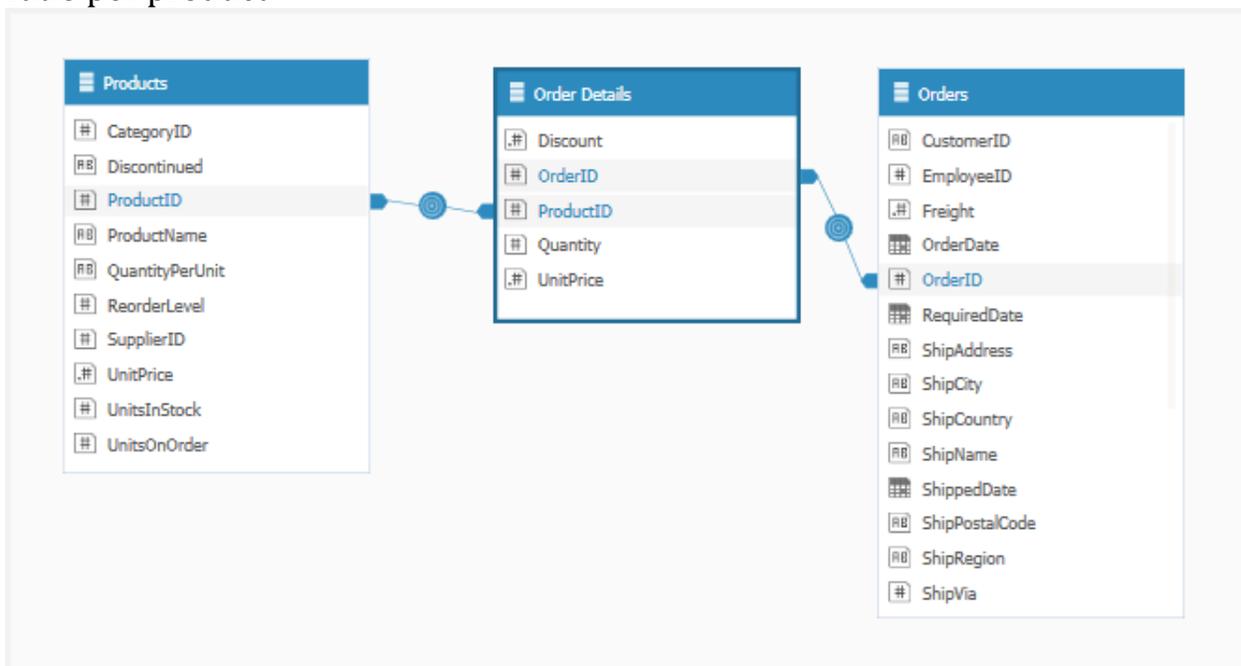
### Modeling Challenge

You must decide whether to calculate the derived facts “on demand”, meaning in the web application, or in advance in the ElastiCube. Take into consideration that

calculating ‘On Demand’ Derived Facts in the web application can enable more dynamic filtering, while calculating them in the ElastiCube stage will save query time when retrieving the data, and enforce calculation consistency, especially with non-trivial facts. This is due to the fact that the dashboard designer/end users will receive consistent results for measures, instead of having to create the complex measures individually, by their own understanding.

### Solution

In the following schema you can create a derived fact to calculate the inventory ratio per product.



Create a [custom table](#) using an SQL Expression that joins the “Order Details” table with the “Products” table and returns the division result of “Quantity” and “UnitOnOrder”, with the following Syntax:

```
SELECT
[Products].ProductID,
toFloat(sum(UnitsOnOrder))/toFloat(sum(Quantity)) AS
InventoryRatio
FROM [Products] JOIN [Order Details]
```

```

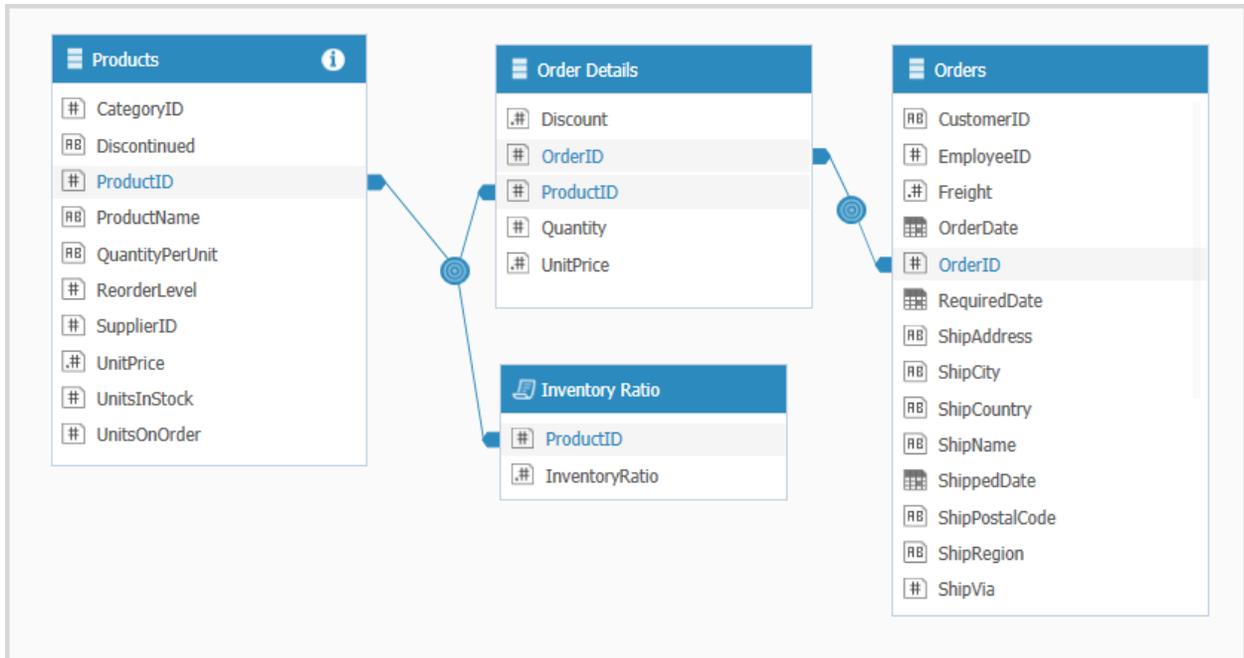
ON [Products].ProductID=[Order Details].ProductID
GROUP BY [Products].ProductID
HAVING
tofloat (sum (UnitsOnOrder) ) / tofloat (sum (Quantity) ) > 0

```

The result table will give the desired results:

ProductID	InventoryRatio
11	1.614731
31	2.555476
49	2.423077
74	0.8754209
2	1.66509
32	2.020202
21	1.535433
37	2.4
70	0.4773562
56	0.3958828
43	0.4827586
68	0.4255319
3	2.560976
64	3.243243
66	3.34728
45	1.929134
48	3.043478

Connect the custom table to the rest of the tables:



**Note:** You can also add the “InventoryRatio” measure to the “Products” table using the Lookup() function by “ProductID”.

## Calendar vs. Fiscal Year

### Business Case

A large number of companies use a fiscal calendar that does not comply with the Gregorian 12-month calendar.

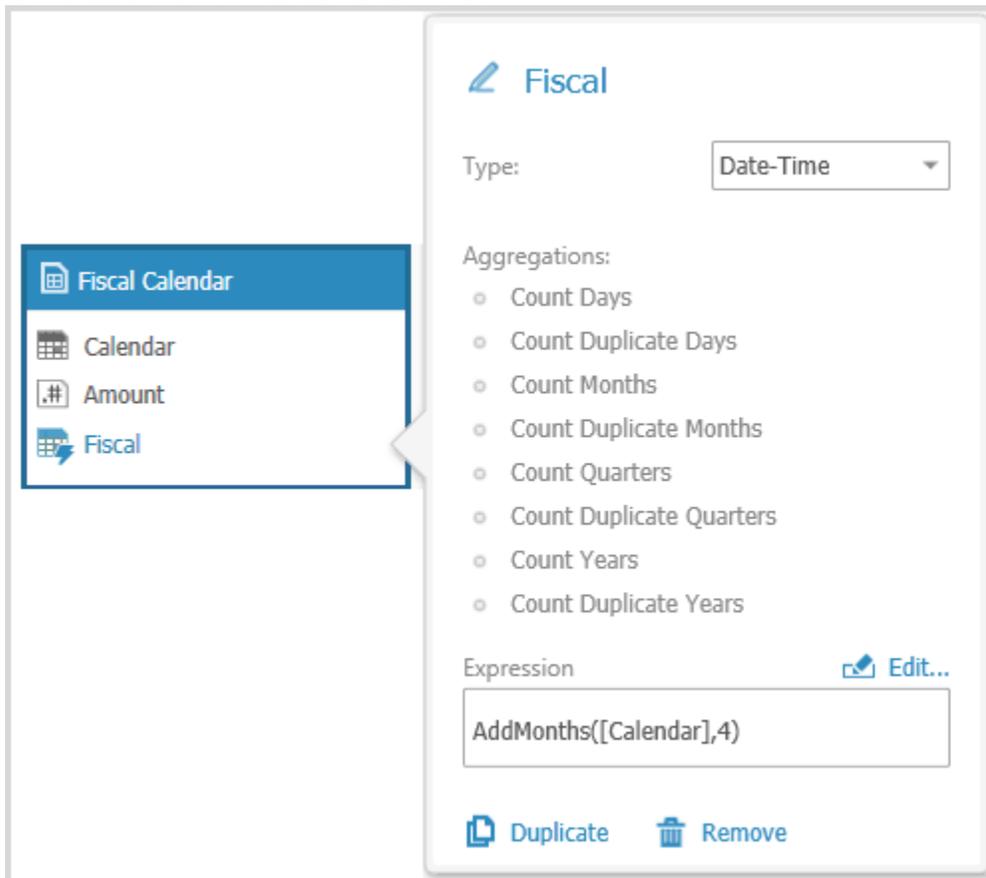
### Modeling Challenge

This requires modeling the data properly so that the data can be reported or analyzed via the normal calendar or via the revised fiscal calendar.

### Solution

In this example, let’s assume that the Fiscal Calendar starts on September 1st. So if we are in the calendar year of 2013, then the fiscal year of 2014 starts

September 1st. To accomplish this, we create a custom field that takes the date field and adds four months to it.



When you create a pivot table in the web application, you will see that the new year (2014) starts in September using the Fiscal field.

Calendar Vs. Fiscal Year ⓘ			
Days in Date	Calendar Year	Fiscal Year	Total Amount
08/29/2013	2013	2013	35,756
08/30/2013	2013	2013	3,687
08/31/2013	2013	2013	35,617
09/01/2013	2013	2014	1,108
09/02/2013	2013	2014	91,387
09/03/2013	2013	2014	11,330
09/04/2013	2013	2014	69,273

# Time Zone Conversion

## Business Case

In many cases, we need to generate reports based on data from different time zones.

## Modeling Challenge

When working with different time zones, the challenge is to store all of the business transactions in an absolute time reference that does not change with the seasons, locations (for instance – GMT), or daylight saving. Therefore, the absolute transition time is a combination of location and date.

## Solution

The aim is to add an “absolute time” field to every business transaction, based on its location and time.

### Step 1 – Create a Reference Source Table

Create a source table (database table / Excel / CSV) that contains the countries and cities that exist in the database, a numeric representation of timestamp range to determine if the transaction belongs to daylight savings time or not (see the this web site), and the UTC to allow the conversion to GMT.

For example:

Country	City	DST_From	DST_To	UTC
USA	Seattle	20120311.2	20121103.1	-7
USA	Seattle	20121103.1	20130310.2	-8
USA	Seattle	20130310.2	20131027.1	-7
USA	Seattle	20131027.1	20140309.2	-8

UK	London	20120325.1	20121028.2	0
UK	London	20121028.2	20130330.1	1
UK	London	20130330.1	20131027.2	0
UK	London	20131027.2	20140330.1	1

## Step 2 – Add a Numeric Representation of the OrderDate

To associate the Order Date with its UTC, create a custom field of type “Decimal” with a numeric representation of the Date timestamp, using this SQL statement:

```
getyear (OrderDate) *10000+getmonth (OrderDate) *100+getday (OrderDate) +ToDouble (gethour (OrderDate) ) /100
```

The result table should look like this:

CustomerID	OrderID	EmployeeID	OrderDate	DateNum	ShipAddress	ShipCity	ShipCountry	Freight
VINET	10248	5	7/4/2012 12:00:00 AM	20120704	59 rue de l'Abbaye	Reims	France	32.38
TOMSP	10249	6	7/5/2012 12:00:00 AM	20120705	Luisenstr. 48	Münster	Germany	11.61
HANAR	10250	4	7/8/2012 12:00:00 AM	20120708	Rua do Paço, 67	Rio de Janeiro	Brazil	65.83
VICTE	10251	3	7/8/2012 12:00:00 AM	20120708	2, rue du Commerce	Lyon	France	41.34
SUPRD	10252	4	7/9/2012 12:00:00 AM	20120709	Boulevard Tirou, 255	Charleroi	Belgium	51.3
HANAR	10253	3	7/10/2012 12:00:00 AM	20120710	Rua do Paço, 67	Rio de Janeiro	Brazil	58.17
CHOPS	10254	5	7/11/2012 12:00:00 AM	20120711	Hauptstr. 31	Bern	Switzerland	22.98
RICSU	10255	9	7/12/2012 12:00:00 AM	20120712	Starenweg 5	Genève	Switzerland	148.33
WELLI	10256	3	7/15/2012 12:00:00 AM	20120715	Rua do Mercado, 12	Resende	Brazil	13.97
HILAA	10257	4	7/16/2012 12:00:00 AM	20120716	Carrera 22 con Ave. Carlos Soublette #8-35	San Cristóbal	Venezuela	81.91
ERNSH	10258	1	7/17/2012 12:00:00 AM	20120717	Kirchgasse 6	Graz	Austria	140.51
CENTC	10259	4	7/18/2012 12:00:00 AM	20120718	Sierras de Granada 9993	México D.F.	Mexico	3.25
OTTIK	10260	4	7/19/2012 12:00:00 AM	20120719	Mehrheimerstr. 369	Köln	Germany	55.09
QUEDE	10261	4	7/19/2012 12:00:00 AM	20120719	Rua da Panificadora, 12	Rio de Janeiro	Brazil	3.05
RATTC	10262	8	7/22/2012 12:00:00 AM	20120722	2817 Milton Dr.	Albuquerque	USA	48.29
ERNSH	10263	9	7/23/2012 12:00:00 AM	20120723	Kirchgasse 6	Graz	Austria	146.06
FOLKO	10264	6	7/24/2012 12:00:00 AM	20120724	Åkerгатan 24	Bräcke	Sweden	3.67
BLONP	10265	2	7/25/2012 12:00:00 AM	20120725	24, place Kléber	Strasbourg	France	55.28
WARTH	10266	3	7/26/2012 12:00:00 AM	20120726	Torikatu 38	Oulu	Finland	25.73
FRANK	10267	4	7/29/2012 12:00:00 AM	20120729	Berliner Platz 43	München	Germany	208.58
GROSR	10268	8	7/30/2012 12:00:00 AM	20120730	5ª Ave. Los Palos Grandes	Caracas	Venezuela	66.29
WHITC	10269	5	7/31/2012 12:00:00 AM	20120731	1029 - 12th Ave. S.	Seattle	USA	4.56
WARTH	10270	1	8/1/2012 12:00:00 AM	20120801	Torikatu 38	Oulu	Finland	136.54
SPLIR	10271	6	8/1/2012 12:00:00 AM	20120801	P.O. Box 555	Lander	USA	4.54
RATTC	10272	6	8/2/2012 12:00:00 AM	20120802	2817 Milton Dr.	Albuquerque	USA	98.03
QUICK	10273	3	8/5/2012 12:00:00 AM	20120805	Taucherstraße 10	Cunewalde	Germany	76.07
VINET	10274	6	8/6/2012 12:00:00 AM	20120806	59 rue de l'Abbaye	Reims	France	6.01
MAGAA	10275	1	8/7/2012 12:00:00 AM	20120807	Via Ludovico il Moro 22	Bergamo	Italy	26.93
TORTU	10276	8	8/8/2012 12:00:00 AM	20120808	Avda. Azteca 123	México D.F.	Mexico	13.84
MORGK	10277	2	8/9/2012 12:00:00 AM	20120809	Heerstr. 22	Leipzig	Germany	125.77

## Step 3 – Join between the Two Tables

The third step includes creating a custom SQL expression that joins between the two tables and creating the “Absolute Time” custom field within it (“GMTDate”). This is to create a synchronization between all the transactions. The custom field will be created using the “add hours” function with the matching UTC value. See the following script:

```
SELECT
[Orders].CustomerID,
[Orders].EmployeeID,
[Orders].Freight,
[Orders].OrderDate,
[Orders].OrderID,
[Orders].ShipAddress,
[Orders].ShipCity,
[Orders].ShipCountry,
AddHours([Orders].OrderDate,[GMT Conversion.csv].UTC)
AS GMTDate
FROM [Orders]
JOIN
[GMT Conversion.csv]
ON
[Orders].ShipCity=[GMT Conversion.csv].City AND
[Orders].ShipCountry=[GMT Conversion.csv].Country AND
[Orders].DateNum>=[GMT Conversion.csv].DST_From AND
[Orders].DateNum<[GMT Conversion.csv].DST_To
```

The result table will look like this:

CustomerID	EmployeeID	Freight	OrderDate	OrderID	ShipAddress	ShipCity	ShipCountry	GMTDate
WHITC	5	4.56	7/31/2012 12:00:00 AM	10269	1029 - 12th Ave. S.	Seattle	USA	7/30/2012 4:00:00 PM
WHITC	5	4.56	7/31/2012 12:00:00 AM	10269	1029 - 12th Ave. S.	Seattle	USA	7/30/2012 5:00:00 PM
WHITC	5	4.56	7/31/2012 12:00:00 AM	10269	1029 - 12th Ave. S.	Seattle	USA	7/30/2012 4:00:00 PM
WHITC	5	4.56	7/31/2012 12:00:00 AM	10269	1029 - 12th Ave. S.	Seattle	USA	7/30/2012 5:00:00 PM
BSBEV	7	22.77	8/26/2012 12:00:00 AM	10289	Fauntleroy Circus	London	UK	8/26/2012 1:00:00 AM
BSBEV	7	22.77	8/26/2012 12:00:00 AM	10289	Fauntleroy Circus	London	UK	8/26/2012 12:00:00 AM
BSBEV	7	22.77	8/26/2012 12:00:00 AM	10289	Fauntleroy Circus	London	UK	8/26/2012 1:00:00 AM
BSBEV	7	22.77	8/26/2012 12:00:00 AM	10289	Fauntleroy Circus	London	UK	8/26/2012 12:00:00 AM
WHITC	4	23.29	11/1/2012 12:00:00 AM	10344	1029 - 12th Ave. S.	Seattle	USA	10/31/2012 4:00:00 PM
WHITC	4	23.29	11/1/2012 12:00:00 AM	10344	1029 - 12th Ave. S.	Seattle	USA	10/31/2012 5:00:00 PM
WHITC	4	23.29	11/1/2012 12:00:00 AM	10344	1029 - 12th Ave. S.	Seattle	USA	10/31/2012 4:00:00 PM
WHITC	4	23.29	11/1/2012 12:00:00 AM	10344	1029 - 12th Ave. S.	Seattle	USA	10/31/2012 5:00:00 PM
SEVES	5	288.43	11/21/2012 12:00:00 AM	10359	90 Wadhurst Rd.	London	UK	11/21/2012 1:00:00 AM
SEVES	5	288.43	11/21/2012 12:00:00 AM	10359	90 Wadhurst Rd.	London	UK	11/21/2012 12:00:00 AM
SEVES	5	288.43	11/21/2012 12:00:00 AM	10359	90 Wadhurst Rd.	London	UK	11/21/2012 1:00:00 AM
SEVES	5	288.43	11/21/2012 12:00:00 AM	10359	90 Wadhurst Rd.	London	UK	11/21/2012 12:00:00 AM
EASTC	1	71.97	11/26/2012 12:00:00 AM	10364	35 King George	London	UK	11/26/2012 1:00:00 AM
EASTC	1	71.97	11/26/2012 12:00:00 AM	10364	35 King George	London	UK	11/26/2012 12:00:00 AM
EASTC	1	71.97	11/26/2012 12:00:00 AM	10364	35 King George	London	UK	11/26/2012 1:00:00 AM
EASTC	1	71.97	11/26/2012 12:00:00 AM	10364	35 King George	London	UK	11/26/2012 12:00:00 AM
SEVES	1	22.21	12/9/2012 12:00:00 AM	10377	90 Wadhurst Rd.	London	UK	12/9/2012 1:00:00 AM
SEVES	1	22.21	12/9/2012 12:00:00 AM	10377	90 Wadhurst Rd.	London	UK	12/9/2012 12:00:00 AM
SEVES	1	22.21	12/9/2012 12:00:00 AM	10377	90 Wadhurst Rd.	London	UK	12/9/2012 1:00:00 AM
SEVES	1	22.21	12/9/2012 12:00:00 AM	10377	90 Wadhurst Rd.	London	UK	12/9/2012 12:00:00 AM
SEVES	2	34.86	12/19/2012 12:00:00 AM	10388	90 Wadhurst Rd.	London	UK	12/19/2012 1:00:00 AM
SEVES	2	34.86	12/19/2012 12:00:00 AM	10388	90 Wadhurst Rd.	London	UK	12/19/2012 12:00:00 AM
SEVES	2	34.86	12/19/2012 12:00:00 AM	10388	90 Wadhurst Rd.	London	UK	12/19/2012 1:00:00 AM

## Step 4 – Make Schema Adjustments

For the next step, do the following:

- Replace the current Orders table with the new one,
- Refer to the new “Absolute Time” custom field (“GMTDate”) as the leading date field
- Make the reference tables (“Orders” and “GMT Conversion.csv”) invisible.

# Currency Conversion

## Business Case

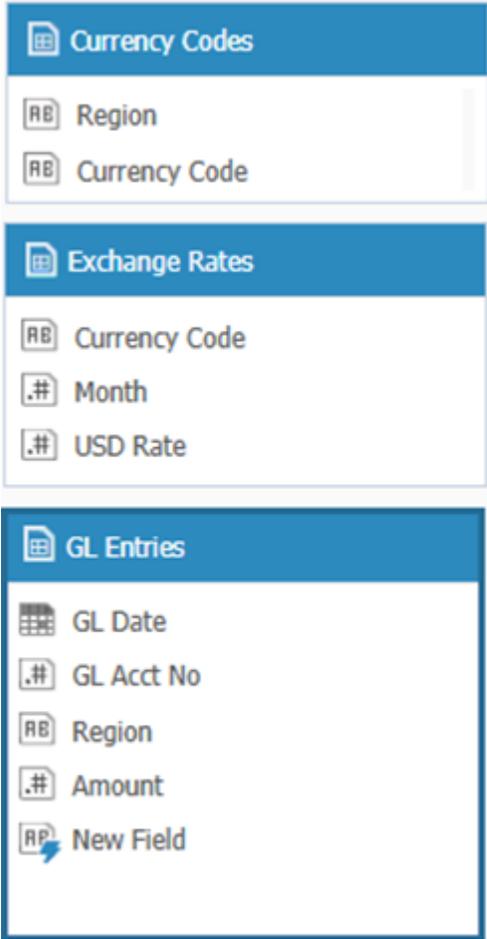
Most data for entities is recorded in their local reporting currency (ie \$ for United States, £ for UK). Here we want to convert all the amounts to USD.

# Modeling Challenge

This requires determining the Currency Rate of the region and then multiplying the value in local currency by the associated Exchange Rate by Month.

## Solution

Create two custom fields in the GL Entries. The first will look up the Currency code of the region. This field will be used along with a month field to link to the Exchange Rates table.



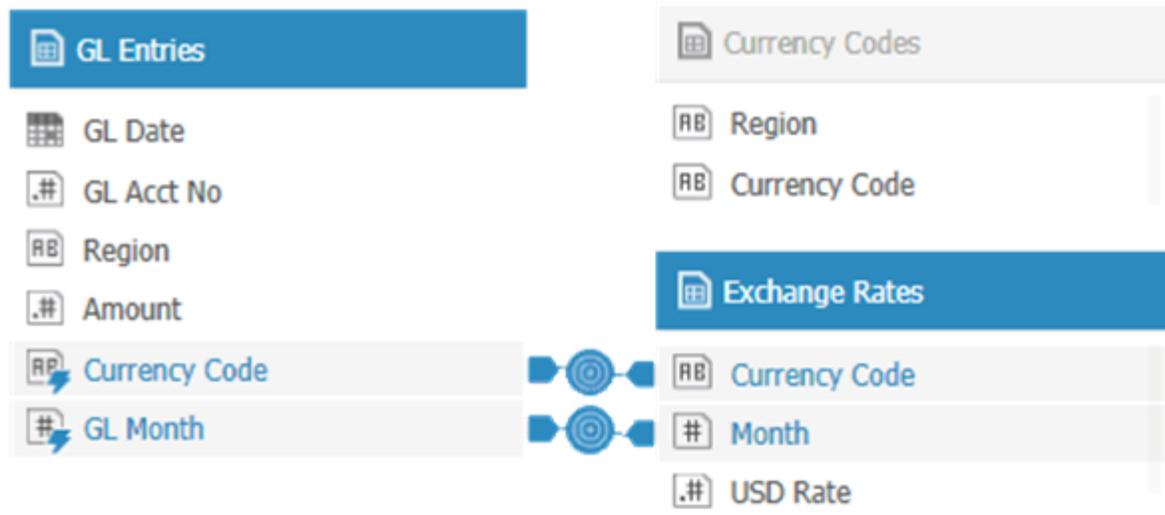
The first field in the GL Entries is created using the lookup function to retrieve values from the Currency Codes table.

```
Lookup([Currency Codes],[Currency Code],Region,Region)
```

Then create a second Custom Field for the Month of the GL Date.

```
GetMonth([GL Date])
```

Next, link the fields together (note that both Month fields were set to Integer and the Currency Codes table to Invisible).



The Local Amount multiplied by the Exchange Rate gives the Converted USD Amount.

## LOCAL and USD

Region	GL Month	Local Amount	Exchange Rate	Converted USD Amount
IL	1	1,303,490	.25	325,873
	2	833,853	.27	225,140
	3	1,915,644	.28	536,380
	4	1,303,288	.27	345,371
	5	1,232,185	.29	357,334
	6	1,743,100	.28	479,353
	7	1,387,776	.28	388,577
	8	17,116	.29	4,964
US	1	939,352	1.00	939,352
	2	752,517	1.00	752,517
	3	604,556	1.00	604,556
	4	648,043	1.00	648,043
	5	769,723	1.00	769,723
	6	884,302	1.00	884,302
	7	550,515	1.00	550,515

## Current vs. Previous Period for Specific Date Range

### Business Case

In many cases we would like to compare our business' performance last week, to the week before, or maybe we would like to see a percentage of sales growth for the current month/quarter compared to the previous month/quarter.

### Modeling Challenge

Since we want the compared time range to be as flexible as possible, the solution has to include both layers – ElastiCube and web application.

## Solution

Create a custom table in the ElastiCube to summarize the totals/counts per day for the source table:

```
SELECT
a.Date,
sum(a.Revenue)AS value
FROM [Accord 2011 Client List] AS a GROUP BY a.Date
```

Create a custom table in the ElastiCube with current vs. previous values, by adjusting the script below:

```
SELECT
curr.Date AS date,
curr.value AS current,
prev.value AS prev
FROM [sum] curr
LEFT JOIN [sum] AS prev
ON curr.Date = addyears(prev.Date,1)
UNION
SELECT
addyears(prev.Date,1) AS date,
curr.value,
prev.value
FROM [sum] prev
LEFT JOIN [sum] AS curr
ON prev.Date= addyears(curr.Date,-1)
```

In the web application, add a 'date range picker' using the days from the custom table. Then add two new numeric indicators. In the first numeric picker add the 'sum of the current value', in the second numeric picker, add the 'sum of the previous value'.

In the date range picker, select the days of interest and you will see the current and previous values.

## Calculating the Number of Open Orders per Day

### Business Case

An open sales order is where the order has been placed but has not yet been delivered. If for example there is an order for 100 items and against this order only 50 items have been delivered (it is partially delivered). A high level of open orders per day may indicate that something is wrong with orders handling.

### Modeling Challenge

We cannot just count the number of orders per day because it will exclude orders that were open on a certain day and are already closed. Therefore, we will need to create a snapshot of the number of open orders per day.

### Solution

1. Import an Excel file with all dates listed in the Orders table into the ElastiCube.
2. To improve query performance, convert all the date fields into numeric representations (for more information, see [Numeric Representation of Date fields](#)).
3. Create the following custom table:

```
SELECT  
s.Dates,  
tm.Created_At,  
tm.Closed_At,  
tm.TicketId
```

```
FROM [All Dates] s LEFT JOIN [Orders] tm
ON s.DateInt >= tm.CreatedAtInt
AND (tm.SolvedAt IS NULL OR s.DateInt <= tm.SolvedAtInt)
```

## Slowly Changing Dimensions

### Business Case

Transactional data typically does not change, however the data that describes the associated dimensions may change. This example demonstrates how to manage dimensions that may be updated with new values within the data warehouse at different points in time.

For example, a customer that was living in NYC and moved to LA earlier this year.

Date Customer City

1/1/2011 John Smith NYC

1/1/2013 John Smith LA

### Modeling Challenge

Following the example above, when the transactions were made last year, the customer was living in NYC. Later this year, the customer moved to LA.

If you decide to refer only to the last city and summarize revenue by city, the NYC transaction will be credited to LA only because the customer currently lives there.

If you connect the “Customer” field within the above table to the “Customer” field in the fact table, you will create a “Many to many” relationship because “Customer” is not a unique identifier of this table.

## Solution

The solution is to change the level of granularity of the “Customer\_City” table and add the “Date” field to the key –

1. Concatenation of the Slowly Changing Dimension table’s unique identifier (for example – Customer\_ID + Date)
2. Creation of the same concatenated field in the transactions table.
3. Merge between the 2 keys.
4. This way, you can associate the [Customer\_ID + Date] key of every transaction with the relevant customer city.

# Integrating Data

**Note:** The images on this page were taken in the desktop version of the ElastiCube Manager, however, the same principles described on this page also apply to the web-based version of the ElastiCube Manager.

The following examples explain how to integrate and merge data from different sources into a single ElastiCube structure. This requires properly planning how to merge the data; to avoid creating unnecessary relationships, while avoiding many to many relationships. Examples in this section:

- [Creating a Common Date Selection](#)
- [Financial Reporting](#)
- [Looking Up Values](#)

## Creating a Common Date Selection

### Business Case

When pulling together data from multiple sources, you will have a number of different dates. Marketing has a Campaign Date, Sales has an Opportunity Date and Finance has a GL Date.

### Modeling Challenge

This type of data leaves us with three sets of dates. Modeling it properly will allow you to select from one common date field while still leaving you the option to choose from one of the three date fields individually.

## Solution

Create a custom table that retrieves a unique list of the dates used between the three tables.

GL Entries	Marketing Campaigns	Sales Opportunities
GL Date	Marketing Campaign ID	Opportunity ID
GL Acct No	Marketing Campaign Date	Opportunity Stage
Region	Campaign Type	Opportunity Date
Amount	Campaign Cost	Amount

Use the following syntax:

```
Select [GL Date] AS [Common Date]
```

```
FROM [GL Entries]
```

```
Union
```

```
Select [Marketing Campaign Date] AS [Common Date]
```

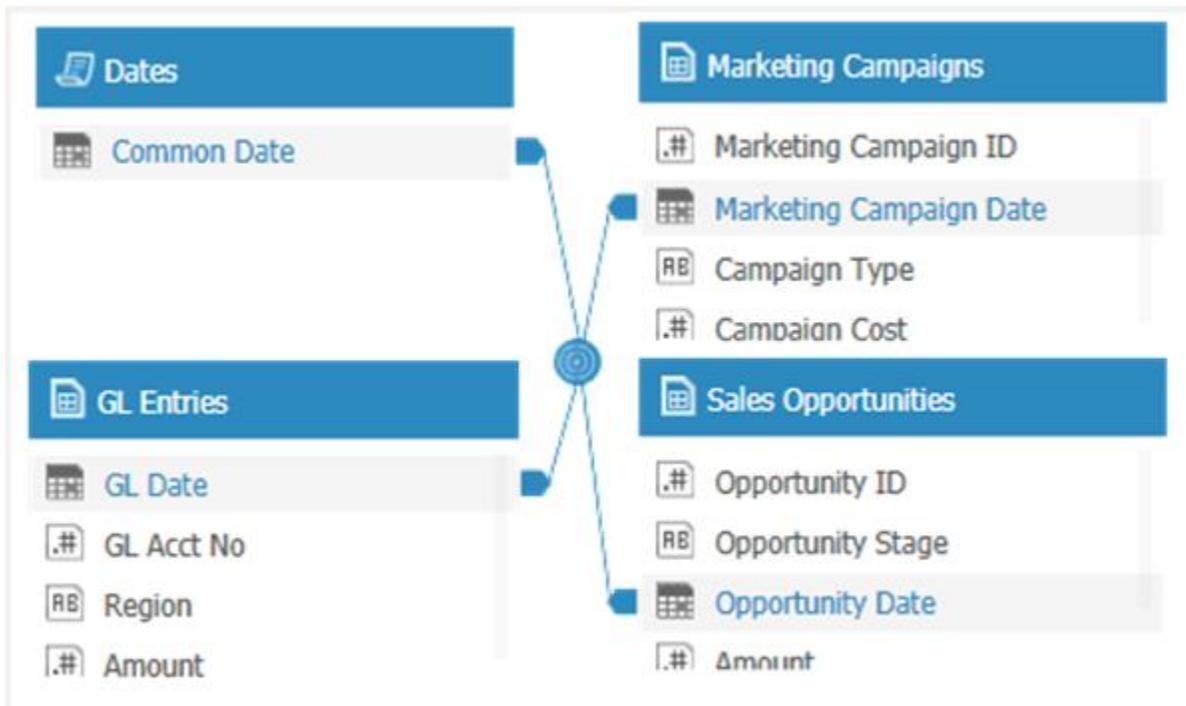
```
FROM [Marketing Campaigns]
```

```
Union
```

```
Select [Opportunity Date] AS [Common Date]
```

```
FROM [Sales Opportunities]
```

Link the four tables together:



## Results

This gives us the common date field to use. This allows, for example, to select a month that will narrow down the selections across all three tables.

# Financial Reporting

## Business Case

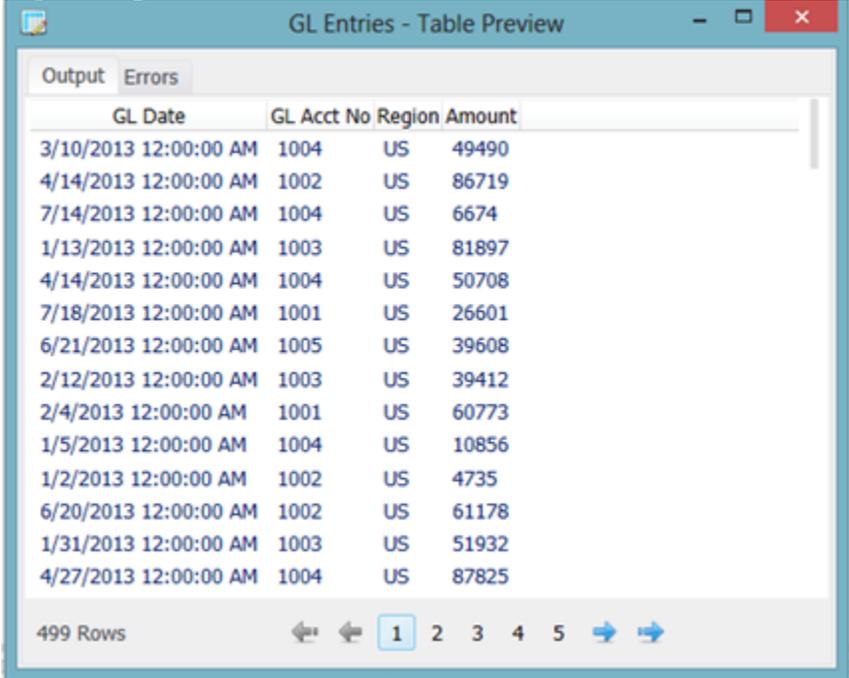
Transactional systems are meant for handling transactions and not for reporting and analysis. For example, Financial GL data will include all the transactions but may not include all income statement or balance sheet reporting definitions.

## Modeling Challenge

This data usually resides in other data sources or tables.

# Solution

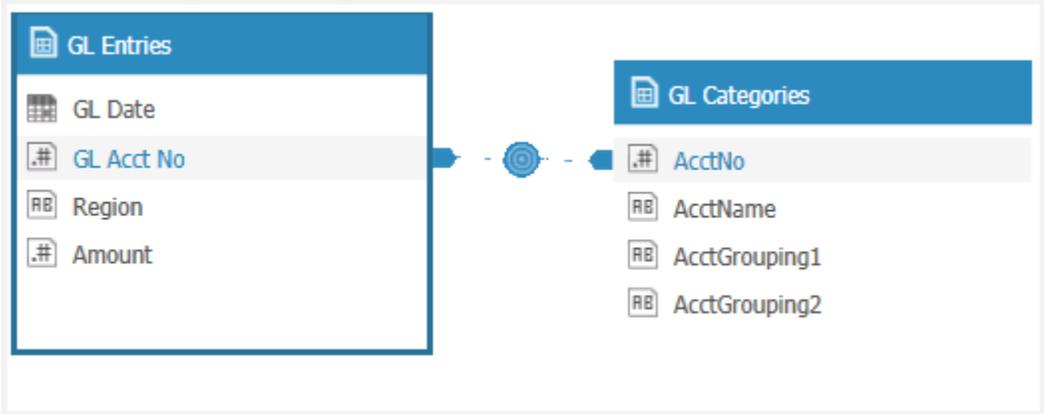
Here is what our sample GL entries look like. It is not very useful to analysis and reporting.



GL Date	GL Acct No	Region	Amount
3/10/2013 12:00:00 AM	1004	US	49490
4/14/2013 12:00:00 AM	1002	US	86719
7/14/2013 12:00:00 AM	1004	US	6674
1/13/2013 12:00:00 AM	1003	US	81897
4/14/2013 12:00:00 AM	1004	US	50708
7/18/2013 12:00:00 AM	1001	US	26601
6/21/2013 12:00:00 AM	1005	US	39608
2/12/2013 12:00:00 AM	1003	US	39412
2/4/2013 12:00:00 AM	1001	US	60773
1/5/2013 12:00:00 AM	1004	US	10856
1/2/2013 12:00:00 AM	1002	US	4735
6/20/2013 12:00:00 AM	1002	US	61178
1/31/2013 12:00:00 AM	1003	US	51932
4/27/2013 12:00:00 AM	1004	US	87825

Another data source is needed to help define how the data will be analyzed and reported.

Start by bringing in another Data Source that contains details about the accounts and how they are categorized:



The two tables are linked on the account number. As there is one record in the GL Categories for multiple rows in the GL Entries, this is known as a One-to-Many Relationship.

The end results show that we can easily view the data in an organized fashion for analysis and reporting.

Revenue, COGs and Expenses ⓘ			
AcctGrouping2	AcctGrouping1	AcctNo	Sum Amount
COGS	COGS	1009	35,434
		1010	231,231
		2001	44,634
	Salaries	3001	3,543,432
		3002	242,423
		3003	523,543
		3004	91,343
Expenses	Operationg Expenses	3009	8,546,214
		3010	8,214,443
	Travel & Entertainment	3005	123,134
		3006	54,354
		3007	234,234
		3008	4,374,573
	Revenue	License	1001
1002			2,264,080
Maintenance		1005	996,276
		1006	1,286,252
		1007	321,347
		1008	6,534,321
Services		1003	3,319,673
		1004	2,335,265
Sum Amount			44,115,437

# Looking Up Values

## Business Case

Sometimes it is necessary to look up a value from one table and bring it into another table. For example, knowing how much a Marketing Campaign costs versus the Sales Opportunity amount.

## Modeling Challenge

These two amounts typically reside in different systems. We need to look up the value from the Marketing system and bring it into the main Sales table.

## Solution

In the Sales Opportunities table, create a custom field that will allow to look up a value from the Marketing Campaigns table based on the Marketing Campaign ID. You can do this using the lookup function.

# Reformatting Data

By reformatting a field, you can create a more readable, and more usable format for analysis, as well as less space consuming in some cases. For example, convert a date field to a numeric field. You can reformat fields within the ElastiCube using a custom SQL expression. An example:

## Numeric Representation of Date Fields

### Business Case

Create a “Date” table using a numeric representation of a Date field instead of the Date field itself to improve query performances. In addition, a numeric representation of a Date can provide the ability of filtering the data by time range.

### Solution

To convert a Date (“4/21/2012 12:36:56 AM”) to an integer (20120421), use the following syntax:

```
10000*getyear(Date)+100*getmonth(Date)+getday(Date)
```

To get the hours and minutes too, use a bigint numeric representation. So, “4/21/2012 12:36:56 AM” will become: 201204210036, using the following syntax:

```
tobigint(100000000*getyear(DateTime)+getmonth(DateTime)*  
1000000+getday(DateTime)*10000+100*gethour(DateTime)+getm  
inute(DateTime))
```

# Date and Time Functions

You can use date and time functions in custom tables and fields. See [Adding a New Custom Table](#) and [Defining and Editing Field Formula](#) for further details.

Below are explanations of available date and time functions.

**ADDYEARS(datetime, number)**

**ADDQUARTERS(datetime, number)**

**ADDMONTHS(datetime, number)**

**ADDDAYS(datetime, number)**

**ADDHOURS(datetime, number)**

**ADDMINUTES(datetime, number)**

**ADDSECONDS(datetime, number)**

Adds a given number of years|quarters|months|days|hours|minutes|seconds to a specified date/time. An example of when to use this is when the fiscal year is not the same as the Gregorian calendar's beginning and end dates.

**CREATEDATE(year,month,day)**

Creates a timestamp from a given year, month and day. Time is set to midnight.

**CURRENTDATE()**

Returns the current date.

**CURRENTTIME()**

Returns the current time.

**CURRENTTIMESTAMP()**

Returns the current timestamp.

**GETYEAR(datetime)**

**GETQUARTER(datetime)**

**GETMONTH(datetime)**

**GETDAY(datetime)**

**GETHOUR(datetime)**

**GETMINUTE(datetime)**

**GETSECOND(datetime)**

Returns a number that represents the  
year|quarter|month|day|hour|minute|second in a given date/time.

**DAYOFWEEK(datetime)**

Returns the day, out of seven days in a week, represented by a given date/time.

**DAYOFYEAR(datetime)**

Returns the day, out of 365 days in a year, represented by a given date/time.

**WEEKOFYEAR(datetime)**

Returns the week, out of 52 weeks in a year, represented by a given date/time.

**DAYNAME(number)**

Returns the name of the day represented by the given number.

**YEARDIFF(end,start)**

**QUARTERDIFF(end,start)**

**MONTHDIFF(end,start)**

**DAYDIFF(end,start)**

**HOURLDIFF(end,start)**

**MINUTEDIFF(end,start)**

**SECONDDIFF(end,start)**

Returns the difference in years|quarters|months|days|hours|minutes|seconds

.

# Functions Overview

Sisense supports many functions that you can use for data calculations.

The following topics describe the functions Sisense supports:

- [Date and Time Functions](#)
- [Logical Functions](#)
- [Mathematical Functions](#)
- [Miscellaneous Functions](#)
- [String and Text Functions](#)
- [Web Functions](#)

# Logical Functions

You can use mathematical functions in custom tables and fields. See [Adding a New Custom Table](#) and [Defining and Editing Field Formula](#) for further details.

Below are explanations of available logical functions.

## **IFBIGINT(condition, true value, false value)**

Returns the first value if the boolean evaluate is true, else returns the second value.

## **IFDATETIME(condition, true value, false value)**

Returns the first value if the boolean evaluate is true, else returns the second value.

## **IFDOUBLE(condition, true value, false value)**

Returns the first value if the boolean evaluate is true, else returns the second value.

## **IFINT(condition, true value, false value)**

Returns the first value if the boolean evaluate is true, else returns the second value.

## **IFSTRING(condition, true value, false value)**

Returns the first value if the boolean evaluate is true, else returns the second value.

## **INRANGE(value,start,end)**

Returns true if a given value is between start and end.

## **ISNULL(value)**

Returns true if a given value is null.

# Mathematical Functions

You can use mathematical functions in custom tables and fields. See [Adding a New Custom Table](#) and [Defining and Editing Field Formula](#) for further details.

Below are explanations of available mathematical functions.

## **ACOS(number)**

Returns the angle, in radians, whose cosine is the given number. Also referred to as arccosine.

## **ASIN(number)**

Returns the angle, in radians, whose sine is a given number. Also referred to as arcsine.

## **ATAN(number)**

Returns the angle in radians whose tangent is a given number. Also referred to as arctangent.

## **CEILING(number)**

Returns the smallest integer that is greater than, or equal to, the a given number.

**COS(number)**  
Returns the trigonometric cosine of a given angle, in radians.

## **COSH(number)**

Returns the hyperbolic cosine of a given number.

## **COT(number)**

Returns the hyperbolic sine of a given number.

## **DIV(number1, number2)**

Returns the number of times number2 fits completely in number1.

## **EXP(number)**

Returns the exponential value of a given number.

**FLOOR(number)**

Returns the largest integer less than or equal to the given numeric expression.

**ISINFINITY(number)**

Returns -1 if the number evaluates to negative infinity, 1 if it evaluates to positive infinity, and 0 otherwise.

**ISINTEGER(number)**

Returns whether a given number is a whole number.

**LOG(number)**

Returns the natural logarithm of a given number.

**LOG10(number)**

Returns the base-10 logarithm of the given float expression

**MAXVAL(number1,number2)**

Returns the highest of two numbers.

**MINVAL(number1,number2)**

Returns the smallest of two numbers.

**MOD(number1,number2)**

Returns the remainder of dividing number1 by number2 (modulo).

**PI()**

Returns the constant value of the mathematical PI.

**POWER(number1,number2)**

Returns the value of the given number raised to a specified power.

**RAND()**

Returns a random number between 0 and 32767

**RANK()**

Returns the rank (i.e. row number) of the current record in a table.

**RANKASC(order\_by\_field)**

Returns the row number ordered by given field in ascending manner.

You can also sort by multiple columns:

`RANKASC ([partitionA], [partitionB]...[partitionN], [orderByField])`

Example: `RANKASC ([EmployeeID], [CategoryID], [Quantity])`

### **RANKDESC(order\_by\_field)**

Returns the row number ordered by given field in descending manner.

You can also sort by multiple columns:

`RANKDESC ([partitionA], [partitionB]...[partitionN], [orderByField])`

Example: `RANKDESC ([EmployeeID], [CategoryID], [Quantity])`

### **ROUND(number,precision)**

Returns the given number, rounded to a specified precision.

### **SIN(number)**

Returns the trigonometric sine of a given angle, in radians.

### **SINH(number)**

Returns the hyperbolic sine of a given number.

### **SQRT(number)**

Returns the square root of a given number.

### **TAN(number)**

Returns the tangent of a given number.

### **TANH(number)**

Returns the hyperbolic tan of a given number.

### **TOBIGINT(string)**

Converts a string representing a valid BigInt value to BigInt type.

### **TODOUBLE(string)**

Converts a string representing a valid double value to Double type.

### **TOINT(string)**

Converts a string representing a valid Int value to Int type.

# Miscellaneous Functions

You can use mathematical functions in custom tables and fields. See [Adding a New Custom Table](#) and [Defining and Editing Field Formula](#) for further details.

The Miscellaneous Functions category currently includes the lookup function.

**LOOKUP(remote\_table,remote\_result\_column,current\_match\_column,remote\_match\_column)**

The Lookup function imports a field from one table into another table, by matching two other corresponding fields from both tables. The result will be the value in remote\_result\_column for which the corresponding remote\_match\_column equals the current\_match\_column.

# String and Text Functions

You can use mathematical functions in custom tables and fields. See [Adding a New Custom Table](#) and [Defining and Editing Field Formula](#) for further details.

Below are explanations of available string and text functions.

## **ASCII(char)**

Returns the Int ascii code of a given character.

## **CHAR(number)**

Returns a character for a given ascii code.

## **CONCAT(string1,string2)**

Concatenates string1 and string2.

## **INDEXOF(string1,string2,n)**

Returns the index of the n-occurrence of string2 in string1.

## **INSERT(string1,string2,overrides,index)**

Inserts string2 into string1 at a specified index, possibly overriding a specified number characters in string1.

## **INSERTSTR(string1,string2,index)**

Inserts string2 into string1 at specified zero-based index.

## **ISNULLOREMPY(string)**

Returns true if the string is null or empty.

## **LEFT(string,length)**

Returns a substring of a specified length, starting at the first character.

## **LENGTH(string)**

Returns the length of a given string.

## **LOWER(string)**

Returns a lowercase version of a given string.

**LTRIM(string)**

Removes whitespaces from the start of the string.

**REMOTEAT(string,index)**

Removes a character at a specified index from a given string.

**REMOVERANGE(string,index,count)**

Removes a specified number of characters from a given string starting at a specified index.

**REPEAT(string,number)**

Returns a concatenation of a string repeated a given number of time.

**REPLACEALL(string,oldstring,newstring)**

Replaces all occurrences of old string with new string.

**RIGHT(string,length)**

Returns a substring of a specified length, end with the last character.

**RTRIM(string)**

Removes whitespaces from the end of the string.

**SOUNDEX(string)**

Returns a four-character (SOUNDEX) code to evaluate the similarity of two objects.

**STRBETWEEN(string, startstring, endstring)**

Returns the substring contained between startstring and endstring.

**STRPARTS(string,delimiter,n)**

Returns the string delimited by a specified delimiter. The search ends after the n-occurrence of the delimiter is found.

**SUBSTRING(string,index,length)**

Returns the substring of a specified length, starting at a specified index.

**TOSTRING(object)**

Converts a given value to a string.

**TRIM(string)**

Removes whitespaces from both ends of a string.

**UPPER(string)**

Returns an uppercase version of a given string.

# Web Functions

You can use mathematical functions in custom tables and fields. See [Adding a New Custom Table](#) and [Defining and Editing Field Formula](#) for further details.

Below are explanations of available web functions.

**Note:** In all circumstances, the URL must include 'http://'. For example, if the URL is www.sisense.com, the syntax should be GetFile ('http://www.sisense.com')

## **DomainIs: true/false DomainIs (URL,compared domain)**

Get whether the given urls domain matches the given domain. For example 'http://www.wikipedia.org' matches org, wikipedia, www.wikipedia.org.

## **GetFile: string GetFile (url)**

Extract the filename portion of the path in the given url and returns it.

The filename is everything after the last slash in the path.

## **GetHost: string GetHost (url)**

Extract host from the given url, for example 'www.wikipedia.org' in 'http://www.wikipedia.org'

## **GetPassword: string GetPassword (url)**

Extract password from the given url.

Specified in the url following the user name, with leading colon (':') and trailing at sign (@), as in 'http://user:password@host/.'

## **GetPath: string GetPath (url)**

Extract path from the given url.

Path is the url portion following the host name, for example

'/wiki/Business\_intelligence' in

'http://www.wikipedia.org/wiki/Business\_Intelligence'

**GetPort: string GetPort (url)**

Extract port from the given url, parsed as integer, for example 123 in 'http://www.wikipedia.org:123/'

**GetQuery: string GetQuery(url)**

Extract the query from the given url.

Specified in the url after the '?' and before (optional) the '#' sign, for example 'par=val' in 'http://www.wikipedia.org/?par=val#ref'.

**GetQueryParamValue (url,param name): string GetQueryParamValue (url,param name)**

Extract param value from the given url's query, by the given param name. Param/Value pairs when supplied, are delimited by an ampersand or a semicolon.

**GetRef: string GetRef (url)**

Extract ref from given url.

Specified in the url starting after last appearance of the '#' sign, for example 'ref' in 'http://www.wikipedia.org/?par=val#ref'

**GetScheme: string GetScheme (url)**

Extract scheme from given url.

Skips leading spaces and followed colon (':') sign, for example 'http://www.wikipedia.org/'

**GetUserName: string GetUserName (url)**

Extract user name from given url.

Specified in the url before the host name and delimited by an '@' as in 'http://user@host/'

**HasHost: bool HasHost (url)**

Get whether the given url specifies a host.

**HasPassword: bool HasPassword (url)**

Get whether the given url specifies a password.

**HasPath: true\fast HasPath (url)**

Get whether the given url specifies a path.

**HasPort: true\fast HasPort (url)**

Get whether the given url specifies a port.

**HasQuery: true\fast HasQuery (url)**

Get whether the given url specifies a query.

**HasRef: true\fast HasRef (url)**

Get whether the given url specifies a ref.

**HasScheme: true\fast HasScheme (url)**

Get whether the given url specifies a scheme.

**HasUsername: true\fast HasUsername (url)**

Get whether the given url specifies a user name.

**HostIsIpAddress: true\fast HostIsIpAddress (url)**

Get whether the given url host name is an ip address.

**IsEmptyUrl: true\fast IsEmptyUrl (url)**

Get whether the given url is empty.

**IsFileUrl: true\fast IsFileUrl (url)**

Get whether the given url is a file url.

**IsSecuredUrl: true\fast IsSecuredUrl (url)**

Get whether the given url scheme indicates a secured connection(https).

**IsStandardScheme: true\fast IsStandardScheme (url)**

Returns true if the scheme for the current urk is a known 'standard' scheme.

Standard schemes have an authority and a path section (including file:).

**IsValidUrl: true\fast IsValidUrl (url)**

Get whether the given url has a valid form.

**PathForRequest: true\fast PathForRequest (url)**

Returns the path that should be sent to the server.

This is the path, parameter, and query portions of the url.

## **SchemeIs: true\fast SchemeIs (url)**

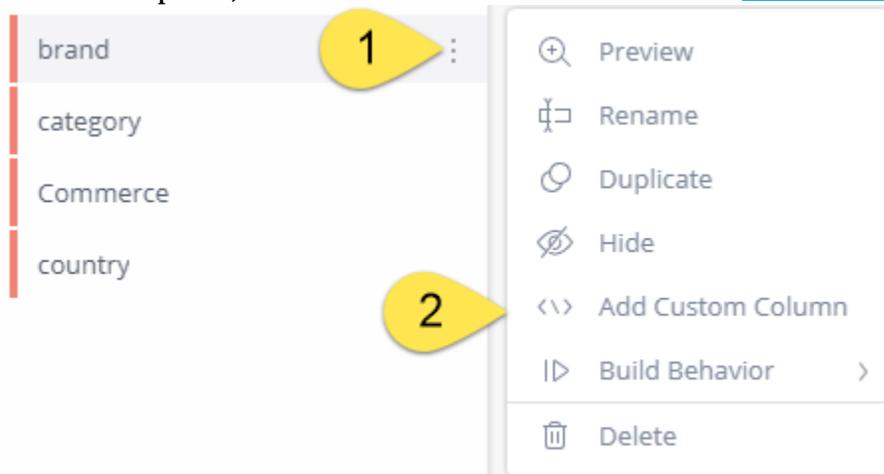
Get whether the given url scheme portion matches the given scheme.

# Adding a Custom Column

The ElastiCube Manager makes it easy for you to add new columns to existing tables. This can be useful if you need to combine data from different existing columns, and when you need to cleanse and prepare data. The new columns also provide designers with additional fields to use as is, or as a basis for even more advanced calculations in their widgets. You can use SQL to customize the values contained within the custom columns.

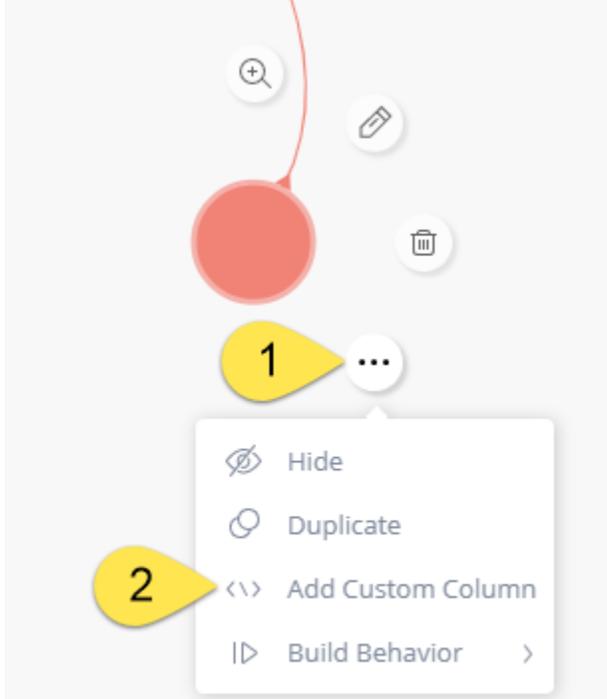
## To add a new column:

1. In the left pane, select **brand** **⋮** **Add Custom Column**. [Click here](#) to learn more.

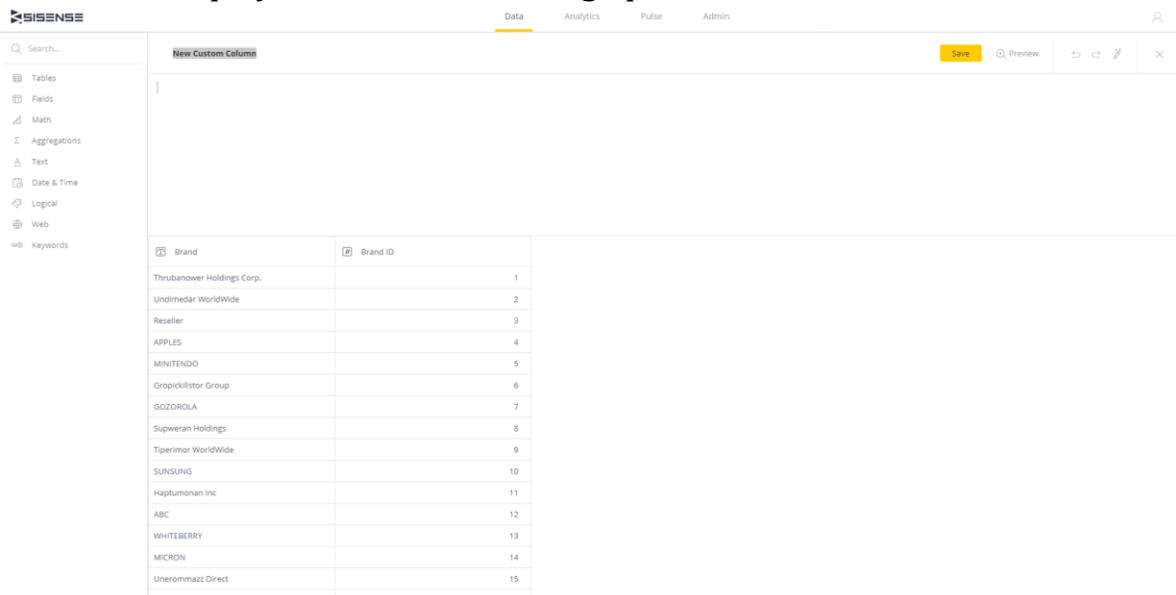


OR

In the schema, select the table and click **...** and **Add Custom Column**.

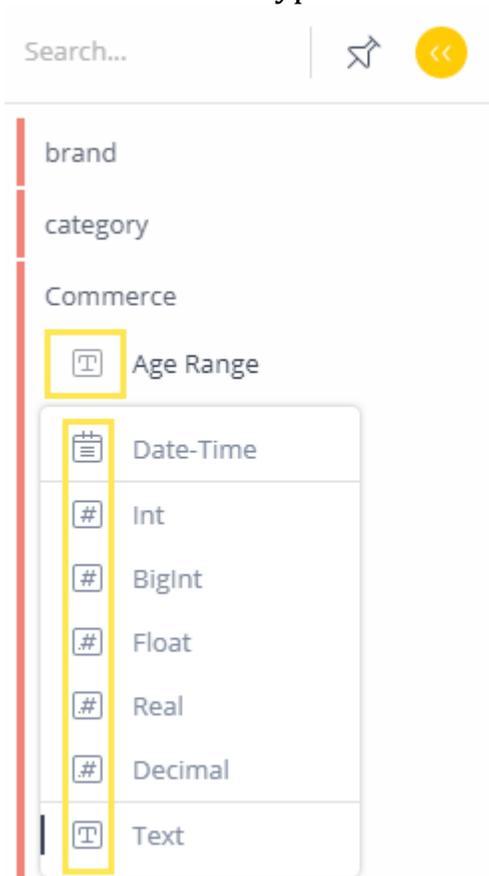


2. The new column will be added to the table, and the column's settings window is displayed with the following options:



- **Name:** Click the column name in the settings window to rename or edit the column name.
- **Type:** When you create the column, it defaults to the type of data you added. From the left pane of your schema, you can modify this by clicking

the field name's Type icon and selecting a new type.



- **Expression:** You can use SQL to customize the data contained in the new field. See [Defining and Editing Column Formulas](#) for further details on developing field expressions.

To delete a custom field, click the field and click **Remove**.

**Note:** Column data can be converted to a different type using functions. For example, a number can be converted to a text field. See the function reference for more details.

## Custom Field Use Cases

Here are some typical scenarios for creating custom fields.

### Row Level Calculations

Calculate revenue from your sales data.

Gross Revenue = (Unit Price \* QuantitySold)

Net Revenue = (Unit Price \* QuantitySold) - (Unit Cost \* QuantityPurchased)

### **Time Difference Calculations**

Work out the arrival time based on due and actual arrival dates.

```
daydiff (DueDate, ReceivedAt)
```

### **Price Comparisons**

Compare your price to the competitor's price.

```
((CompetitorPrice-MyPrice) / ((CompetitorPrice+MyPrice) / 2))
```

### **Buckets**

Create data buckets according to your required business break points.

```
CASE  
WHEN [PriceVsCompetitor%] < -0.05 THEN '-5%'  
WHEN [PriceVsCompetitor%] >= -0.05 AND  
[PriceVsAvgCompetitor%] <= 0.05 THEN '-5%/+5%'  
WHEN [PriceVsCompetitor%] > 0.05 THEN '+5%'  
end
```

### **Attributes**

Create attributes to translate data talk to more coherent categories.

```
CASE  
WHEN Region = 1 THEN 'USA'  
WHEN Region = 2 THEN 'EUR'  
WHEN Region = 3 THEN 'ASIA'  
end
```

### **Date Conversions**

Convert dates into a numeric representation.

```
getyear(Date)*10000+getmonth(Date)*100+getday(Date) AS  
DateNum
```

Convert text into dates.

```
createdate( toint('20'+RIGHT([Timesheet_Date],2)),
toint(LEFT(Timesheet_Date,2)),
toint(StrBetween(Timesheet_Date,'/','/'))
)
```

### **Date Period Comparisons Over Time**

Create month to date and year to date calculations.

[Click here](#) to read our support article.

### **Surrogate Keys**

```
EmployeeID+tostring([DateNum])+tostring(CustomerID)+tostring([Project_ID])
```

### **Importing from Another Table**

Use the lookup function to import a column from a different table. [Click here](#) to read more.

# Adding a Custom Table

With the ElastiCube Manager, you can easily create custom tables, which combine existing data. This is useful for preparing and cleansing data. Follow the steps below to learn how to add custom tables using SQL.

## To create a custom table:

1. In the ElastiCube Manager, open the schema of the ElastiCube you want to add to a custom table to.
2. In the menu bar of the schema, click . The New Custom Table area is displayed.
3. Enter the name of the custom table, and click .
4. Enter SQL statements to access tables and fields that exist in the ElastiCube Manager schema. See also [SQL Reference](#) and [Function Reference](#).
5. To view results based on the SQL statement, click  Preview .  
**Note:** You can save an incomplete or invalid SQL expression, however you will not be able to build the custom expression. At least one build including the relevant base table must be completed in order to preview results.
6. Click **Save** to save your changes.

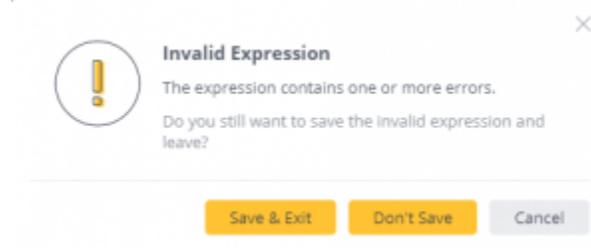
## Tip

- Press Ctrl + Space to complete an SQL statement.
- To make it easier to access table fields, give the table name an alias and reference the alias to bring up the related fields. For example SELECT A. FROM A.
- You can add comments in expressions by using the '-' for a single line or /.../ for blocks.

# Editing Field Formulas

The ElastiCube Manager makes it easy to edit custom tables. You can use SQL to customize the fields contained within a custom table.

When you edit a custom table, any changes you make are saved when you click **Save** unless the expression did not successfully parse. In this case, you are prompted to save a draft of the invalid expression.



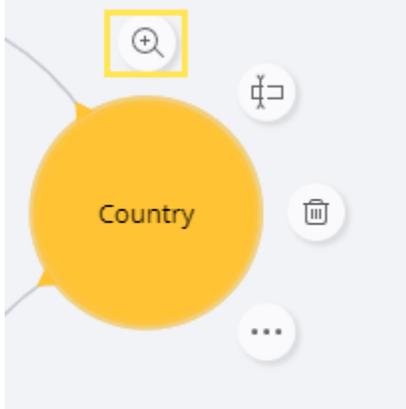
If you save the invalid expression, Sisense continues to use the last valid expression that was saved when building an ElastiCube, but the next time you open a preview of the table, the invalid expression will be displayed. This allows you to continue building working ElastiCubes while letting you save incomplete expressions for later.

If you wish to display the last valid expression, you can click the **Refresh** button in the menu bar and then press **Save**.



**To edit a custom table's SQL expression:**

1. Select the custom table you want to edit and click .



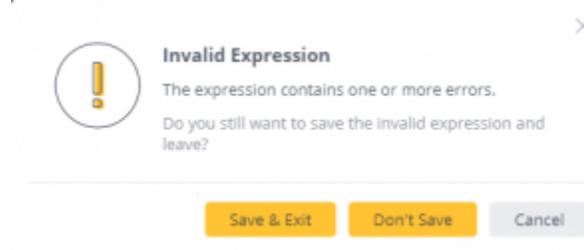
The custom table area is displayed

2. Enter SQL statements to access tables and fields that exist in the schema. See related topics for more details on [SQL](#) and [Function References](#).
3. To view results based on the SQL statement, click  Preview .  
**Note:** At least one build including the relevant base table must have been completed in order to preview results.
4. Click **Save** to save your changes.

# Editing SQL Statements

ElastiCube Manager makes it easy to edit custom tables. You can use SQL to customize the fields contained within a custom table.

When you edit a custom table, any changes you make are saved when you click Save unless the expression did not successfully parse. In this case, you are prompted to save a draft of the invalid expression.



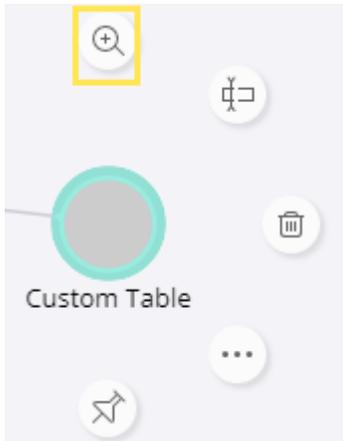
If you save the invalid expression, Sisense continues to use the last valid expression that was saved when building an ElastiCube, but the next time you open a preview of the table, the invalid expression will be displayed. This allows you to continue building working ElastiCubes while letting you save incomplete expressions for later.

If you wish to display the last valid expression, you can click the Refresh button in the menu bar and then press Save.



**To edit a custom table's SQL:**

1. Select the custom table you want to edit and click .



The custom table area is displayed

2. Enter SQL statements to access tables and fields that exist in the schema. See related topics for more details on [SQL References](#) and [Function References](#).
3. To view results based on the SQL statement, click  Preview . (**Note:** At least one build including the relevant base table must have been completed in order to preview results.)
4. Click **Save** to save your changes.

### Tips

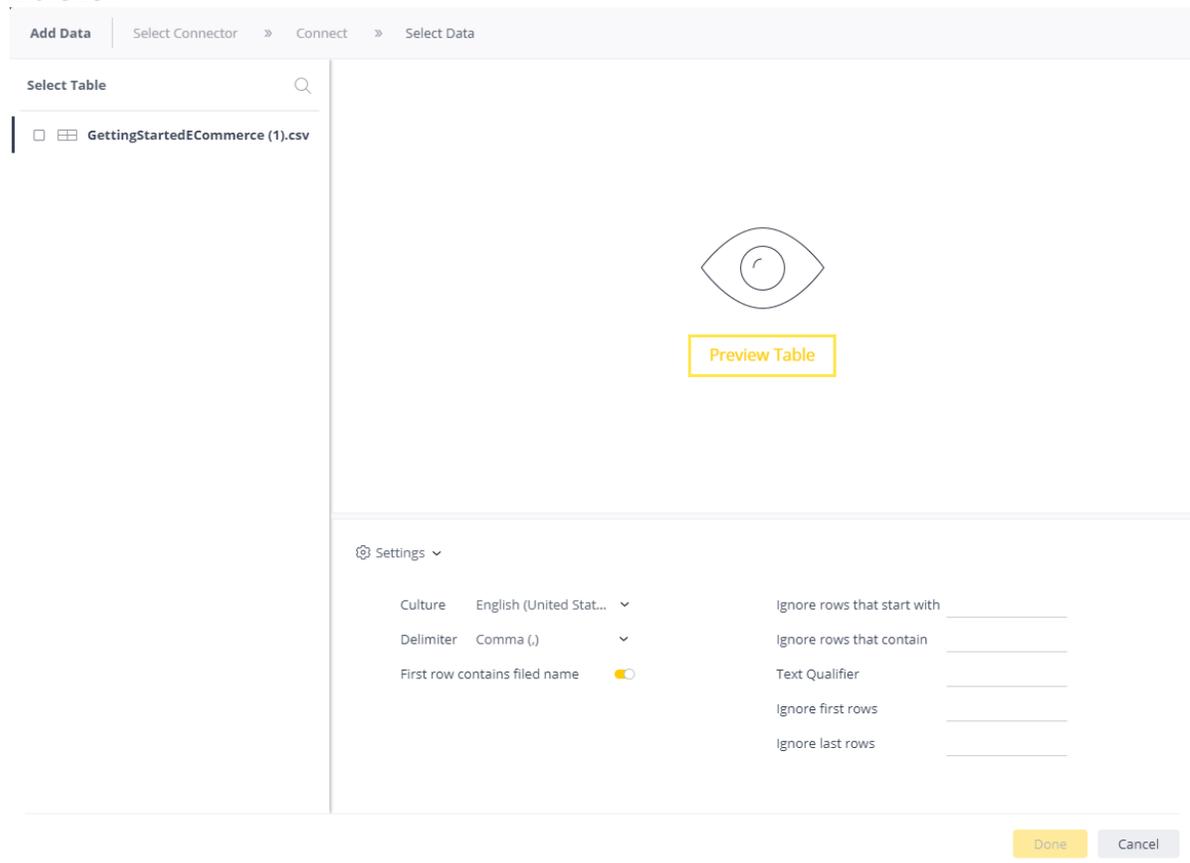
- Press Ctrl + Space to complete an SQL statement.
- To make it easier to access table fields, give the table name an alias, and reference the alias to bring up the related fields. For example Select A. From table as A.
- You can add comments in expressions by using the '-' for a single line or /.../ for blocks.
- A good reference on compatible SQL commands is available on [Microsoft MSDN](#).

# Previewing Results

You can preview results as follows:

## To view results before connecting to a table:

1. In the schema of your ElastiCube, click .
2. Select the relevant connector.
3. Enter your login details and connect to the source data.
4. All tables and views associated with the database will appear in a new window. To preview data contained in a particular table, click **Preview Table**.



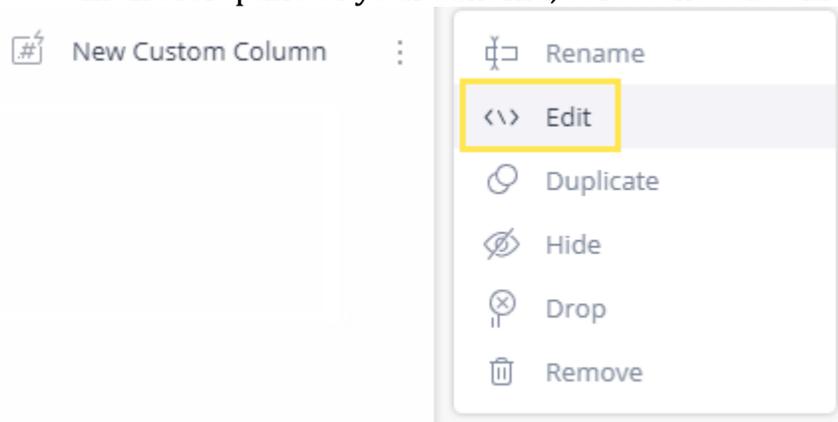
## To view results in a custom table:

- In the schema of your ElastiCube, click the custom table that you want to preview and click .

**Note:** At least one build including the relevant base table must be completed in order to preview results.

**To view results in a custom field:**

- In the left pane of your schema, select the custom field and click **Edit**.



# SQL Reference

You can use SQL to create new tables and fields in the ElastiCube Manager.

SQL Structure	SQL Example
<pre>SELECT FieldName(s), Function(s) FROM TableName 1 INNER JOIN\ LEFT JOIN\RIGHT JOIN\FULL JOIN TableName 2 ON TableName 1. JoinField = TableName 2. JoinField WHERE FieldName Condition AND\OR FieldName Condition GROUP BY FieldName(s) ORDER BY FieldName(s) HAVING FieldName(s) Condition</pre>	<pre>SELECT CustomerName, ContactEmail, Count(Orders) FROM Customer INNER JOIN Order ON Customer.CustomerID = Order.CustomerID WHERE Order.OrderId BETWEEN 10 AND 100 AND Customer.CustomerName IN ('John', 'Mary', 'David') OR Customer.CustomerLastName LIKE 'Harrison' GROUP BY Customer.CustomerName ORDER BY Customer.CustomerLastName HAVING Count(Orders) &gt; 3</pre>

## Basic SQL Syntax Guide

SELECT STATEMENT	
SELECT	SELECT column_name(s)FROM table_name
SELECT *	SELECT *FROM table_name
SELECT DISTINCT	SELECT DISTINCT column_name(s)FROM table_name
Note: Nested SELECT statements will not work with an AND clause in the ElastiCube Manager.	

FUNCTIONS AFTER SELECT CLAUSE	
FUNCTIONS	AVG() – Returns the average value COUNT() – Returns the number of rows MAX() – Returns the largest value MIN() – Returns the smallest value SUM() – Returns the sum

TABLE JOIN FUNCTIONS	
INNER JOIN	SELECT column_name(s)FROM table_name1 T1INNER JOIN table_name2 T2ON T1.column_name=T2.column_name
LEFT JOIN	SELECT column_name(s)FROM table_name1 T1LEFT JOIN table_name2 T2ON T1.column_name=T2.column_name
RIGHT JOIN	SELECT column_name(s)FROM table_name1 T1RIGHT JOIN table_name2 T2ON T1.column_name=T2.column_name
FULL JOIN	SELECT column_name(s)FROM table_name1 T1FULL JOIN table_name2 T2ON T1.column_name=T2.column_name
AS (alias)	SELECT column_name AS column_aliasFROM table_nameorSELECT column_nameFROM table_name AS table_alias

WHERE	SELECT column_name(s)FROM table_nameWHERE column_name operator value
-------	--

FUNCTIONS AFTER THE WHERE CLAUSE	
AND / OR	SELECT column_name(s)FROM table_nameWHERE conditionAND   OR condition
BETWEEN	SELECT column_name(s)FROM table_nameWHERE column_nameBETWEEN value1 AND value2
IN	SELECT column_name(s)FROM table_nameWHERE column_nameIN (value1,value2,..)
LIKE	SELECT column_name(s)FROM table_nameWHERE column_name LIKE pattern

ORDERING AND GROUPING AFTER WHERE CONDITIONS	
GROUP BY	SELECT column_name,aggregate_function(column_name)FROM table_nameWHERE column_name operator valueGROUP BY column_name
ORDER BY	SELECT column_name(s)FROM table_nameORDER BY column_name [ASC   DESC]
HAVING	SELECT column_name,aggregate_function(column_name)FROM table_nameWHERE column_name operator valueGROUP BY column_nameHAVINGaggregate_function(column_name) operator value

## CLAUSES TO COMBINE TABLES

UNION	SELECT column_name(s) FROM table_name1 UNION SELECT column_name(s) FROM table_name2
UNION ALL	SELECT column_name(s) FROM table_name1 UNION ALL SELECT column_name(s) FROM table_name2

## See Also

[Basic SQL Guide](#)

[MSDN SQL Reference](#)

# Transforming and Enriching Data

After you have imported data from your data sources, Data Designers can transform and enrich data in the ElastiCube so Data Designers can prepare easy-to-use data sets for Dashboard Designers. Sisense provides several features that make it easy to prepare, blend, and analyze data that can be fed to the Sisense Web Application and consumed by Viewers.

The following topics describe how you can transform and enrich your data in Sisense:

## Related Topics

- [Adding a Custom Column](#)
- [Editing a Custom Field](#)
- [Adding a Custom Table](#)
- [Editing a Custom Table](#)
- [Previewing Results](#)
- [SQL Reference](#)
- [Function Reference](#)

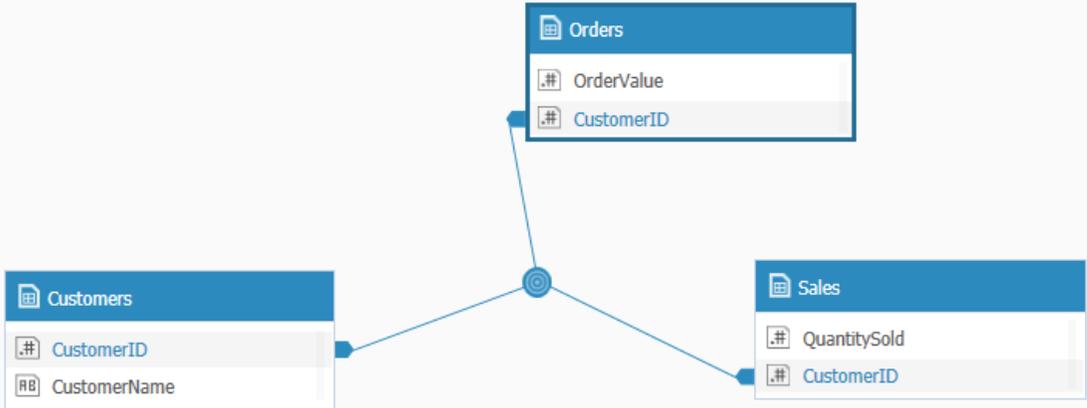
# Chasm and Fan Traps

**Note:** The image on this page were taken in the desktop version of the ElastiCube Manager, however, the same principles described on this page also apply to the web-based ElastiCube Manager.

Chasm and fan traps should be avoided when building your ElastiCube schemas.

## Chasm Traps

A chasm trap occurs when two “Many-to-One” joins converge on a single table, and the query includes measures from both leaf tables. As a result multiple rows are returned from the tables when processing the query.



If you were to calculate both measures (Qty and Value) simultaneously, like in the following example, the values for Customers will be multiplied due to the inner join between the leaf tables, and the results may be incorrect:

```

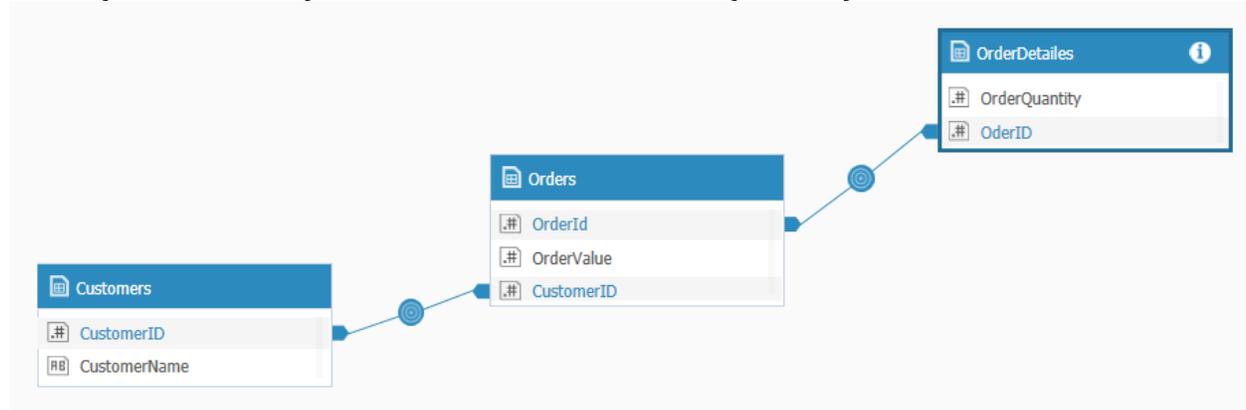
SELECT [Customers].CustomerName,
       sum ([Orders].OrderValue) AS Value,
       sum ([Sales].QuantitySold) AS Qty
FROM [Customers]
JOIN [Orders] ON [Customers].CustomerID=[Orders].CustomerID
JOIN [Sales] ON [Customers].CustomerID=[Sales].CustomerID
GROUP BY [Customers].CustomerName

```

CustomerName	Value	Qty
Jason	1800	180
Daniel	600	60

## Fan Traps

A fan trap occurs when two “many-to-one” joins follow one another in master-detail form (OrderDetails), and the query includes a measure from both the leaf table (OrderDetails) and its immediate master (Orders).



If you try to aggregate both measures simultaneously (using the query below), you will probably get incorrect results:

```

SELECT [Customers].CustomerName,
       sum ([Orders].OrderValue) AS Value,
       sum ([OrderDetails].OrderQuantity) AS Qty
FROM [Customers] JOIN [Orders] ON [Customers].CustomerID= [Orders].CustomerID
join
[OrderDetails] ON [Orders].OrderId=[OrderDetails].OderID
GROUP BY

```

CustomerName	Value	Qty
Jason	700	90
Daniel	400	60
Mike	400	60

The “Qty” measure, corresponding to the leaf measure table (OrderDetails) is calculated correctly, but the “Value” measure, corresponding to the measure held in its master (Orders), is not. This is because we get the “Value” of every OrderID, which may inflate the expected results.

The web application translation module separates the calculations, and unions the results by generating a query for each of the measures’ paths (path in terms of tables to go by). Then the web application translation module will group all the measures with the same tables’ paths into one query and union it with all the other “same path measures” with different paths.

The described “Chasm Trap” can be prevented like this:

```

SELECT [Customers].CustomerName, sum ([Orders].OrderValue) AS Value, 0 AS Qty
FROM [Customers] JOIN [Orders] ON [Customers].CustomerID=[Orders].CustomerID
GROUP BY [Customers].CustomerName
union
SELECT [Customers].CustomerName, 0 AS Value, sum ([Sales].QuantitySold) AS Qty
FROM [Customers] JOIN [Sales] ON [Customers].CustomerID=[Sales].CustomerID
GROUP BY [Customers].CustomerName

```

Output Errors

CustomerName	Value	Qty
Jason	600	0
Daniel	300	0
Mike	300	0
Jason	0	60
Daniel	0	30

And the “Fan Trap” will be prevented this way:

```

SELECT [Customers].CustomerName, 0 AS Value, sum ([OrderDetails].OrderQuantity) AS Qty
FROM [Customers] JOIN [Orders] ON [Customers].CustomerID= [Orders].CustomerID
join
[OrderDetails] ON [Orders].OrderId=[OrderDetails].OderID
GROUP BY
[Customers].CustomerName
union
SELECT [Customers].CustomerName, sum ([Orders].OrderValue) AS Value, 0 AS Qty
FROM [Orders]
JOIN [Customers] ON [Orders].CustomerID=[Customers].CustomerID
GROUP BY CustomerName

```

Output Errors

CustomerName	Value	Qty
Jason	0	90
Daniel	0	60
Mike	0	60
Jason	600	0
Daniel	300	0
Mike	300	0

# Creating and Removing a Relationship between Tables

Relationships specify the connection between tables and enable you to pull data together in meaningful ways. For example, order information is more useful when you know which customer placed each order.

However, you don't need to store both the customer and order information in the same table. Customer and order data can be stored in two related tables with a relationship specified between the two tables to view each order and its corresponding customer information.

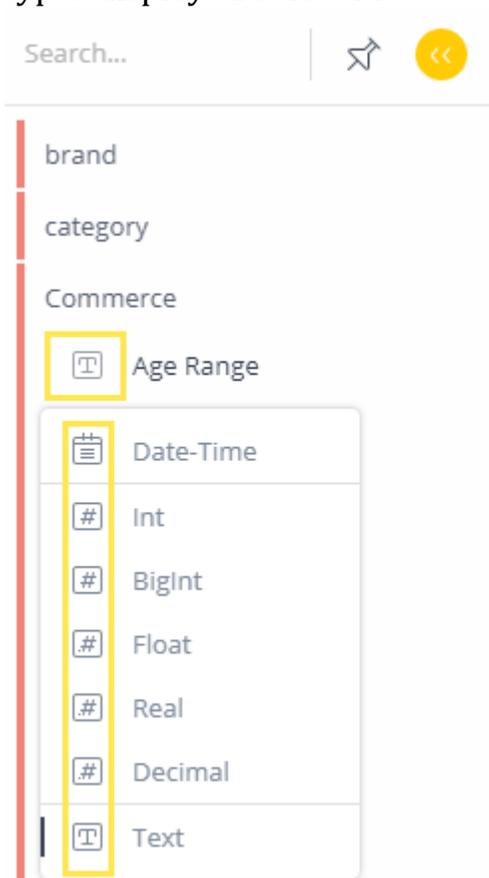
In practical terms relationships allow you to query and combine data from multiple tables in your dashboard.

## **To manually create a relationship between tables in the ElastiCube**

### **Manager:**

1. Drag a table with the relevant field onto another table with the relevant field.  
The Relationship preview window is displayed.
2. Click the columns to be joined from each table. Ensure both fields that you are using to create the relationship are of the same type, for example decimal fields.  
See the video below for a quick example.
3. To check or change the field type, in the left pane, click the field type of the field you want to change and select the new type from the list of possible

types displayed below it.



4. Click **Save** to join the tables.

**To delete an existing relationship:**

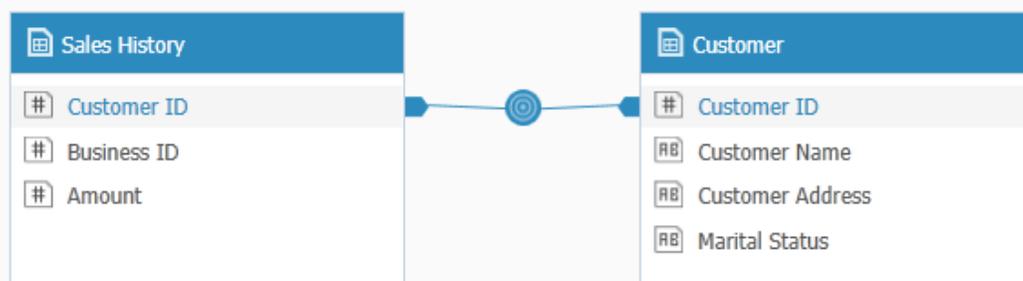
1. Click on any part of the existing relationship. The Delete icon is displayed.
2. Click  to delete the relationship.

# Handling Relationship Cycles

**Note:** The image on this page were taken in the desktop version of the ElastiCube Manager, however, the same principles described on this page also apply to the ElastiCube Manager Online.

Good relationships between tables are the key to pulling in data in logical ways. ElastiCubes are usually created from more than one table. To perform calculations on fields that reside in separate tables, you must define a relationship between the tables so the calculation mechanism knows how to navigate from one field to another.

When a single possible path exists between two fields, there is no issue. For example, performing a query involving the Customer Name and Amount field on the ElastiCube schema below can yield only one result, since there is only one possible path leading from the Customer Name field to the Amount field.

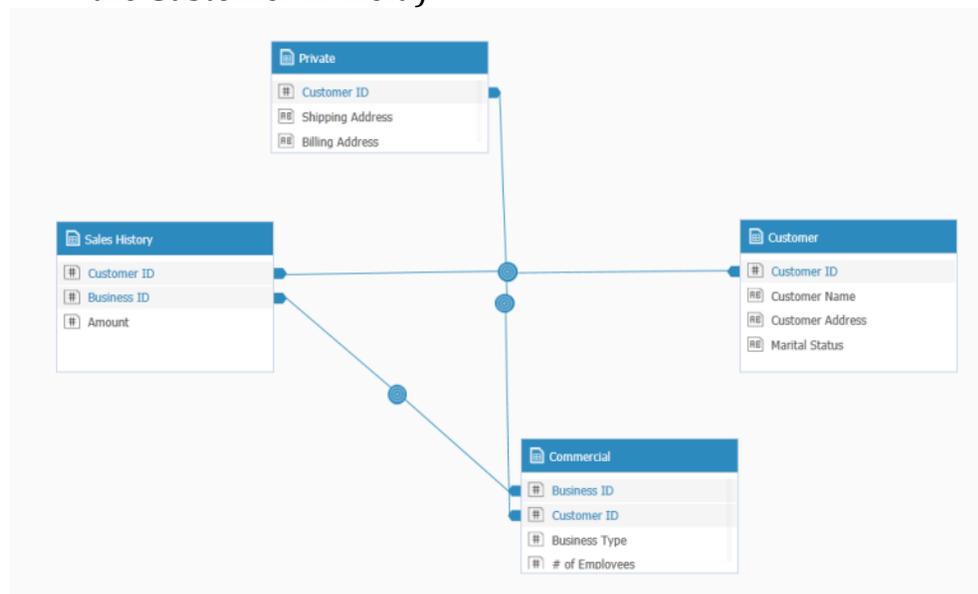


# Relationship Cycles

In some cases, more than one path exists leading from one field to another. Sometimes this is due to poor database design, and other times it is just a realistic necessity due to the way the data is structured. In such cases, there is no absolute way to determine the required path to take to calculate results.

For example, in the ElastiCube schema below, the Customer Name and Amount fields now have multiple paths between them:

- Customer → Sales History (over the Customer ID field)
- Customer → Commercial (over the Customer ID field) → Sales History (over the Business ID field)
- Customer → Private (over the Customer ID field) → Sales History (over the Customer ID field)



## Automatically Dealing with Relationship Cycles

Theoretically, any one of these possible paths could be the correct one in terms of the results you are trying to calculate. When the calculation engine encounters

numerous possible paths, it picks the shortest path, which will have less impact performance-wise on query processing. Below are some general guidelines.

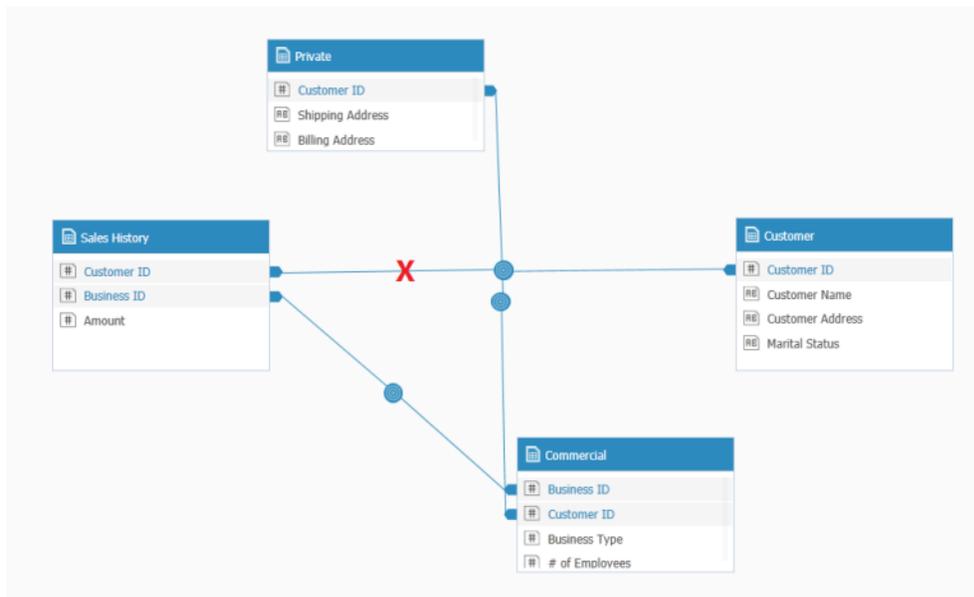
- The calculation engine will always prefer paths that have no many-to-many relationships. If all possible paths contain many-to-many relationships, the path with the least number of this type of relationship is preferred.
- Similarly, the engine will prefer one-to-one relationships over one-to-many relationships.
- If at the end of the first selection process more than one possible path still exists, the engine will prefer the path containing the least number of tables.
- Finally, if numerous paths are still possible, one is picked at random.

## **Breaking Relationship Cycles in the ElastiCube Manager**

To ensure the calculation engine takes the path you want it to take, the cycle must be broken. There are a few techniques to do this depending on the required results.

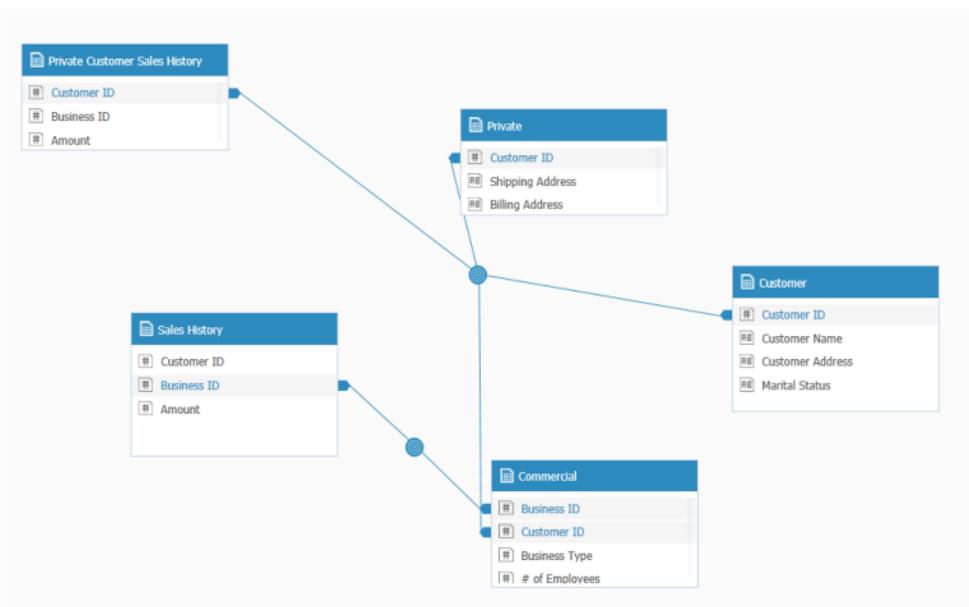
### **Removing Relationships**

One option is to remove one of the relationships in the cycle. For example, removing the Customer ID relationship from the Sales History table will eliminate the cycle. This approach is only applicable when the removed relationship is not used by other queries.



## Duplicating Table Elements

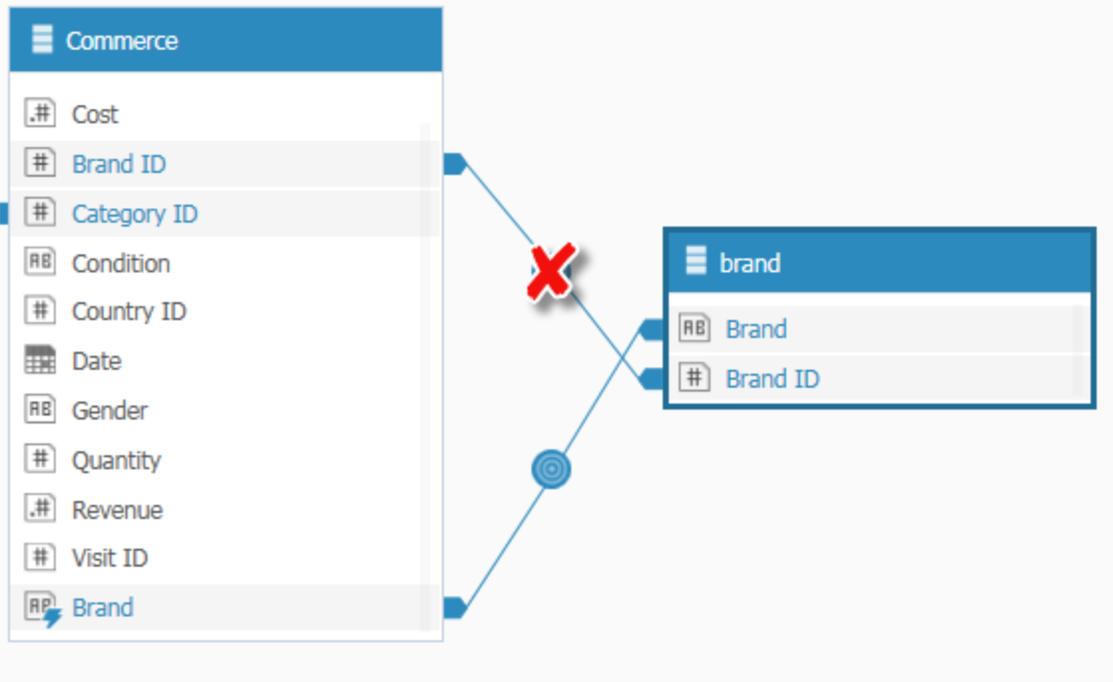
Duplicating table elements is another method to break relationship cycles. In the example below, the Sales History table was duplicated and renamed Private Customers Sales History. In addition, the Customer ID relationship was removed from the original Sales History table. This results in the creation of two separate Amount fields, one for commercial customers residing in the original Sales History table, and one for private customers residing in the new Private Customers Sales History. As the Customer ID relationship was removed from the original table, whichever field you use will determine the correct calculation path.



## Breaking Relationship Cycles in the Sisense Web Application

You can also use the Sisense web app to invoke a specific path to be used during query execution (as opposed to the solutions mentioned above that are defined on the schema level, in the ElastiCube Manager).

Each widget in Sisense has an associated data layout. If you have more than one path connecting corresponding fields between two tables, you can enforce a specific path by filtering the widget, and selecting just the field that you want to use. For example if the Commerce table is connected to Brands via the Brand and Brand ID fields, then add a filter to the widget and select Brand ID to use that path in the query.





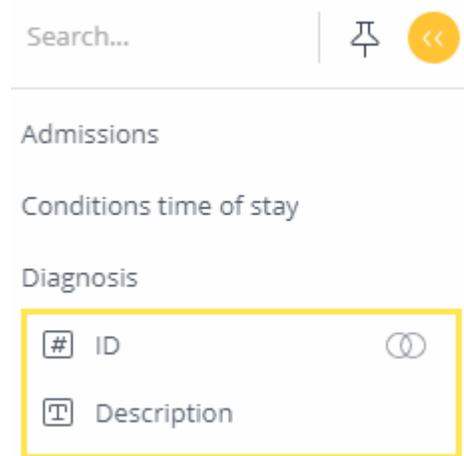
Navigation Pane	Table Icon	Description
Preview & Edit		Displays a preview of your table's data. See <a href="#">Previewing your Data</a> for more information.
Relationships		Displays related tables together side by side. See <a href="#">Creating Relationships</a> for more information.
Rename		Allows you to rename the table.
Duplicate	 Duplicate	Allows you to duplicate the table.
Refresh Schema	 Refresh Schema	Allows you to update a table's columns if there was a change on the data source without having to add the table again.
Hide	 Hide	Allows you to hide a table. Hidden tables are not built with the rest of the ElastiCube.
Add Custom Column	 Add Custom Column	Allows you to add a custom column to the ElastiCube. See <a href="#">Adding Custom Columns</a> for more information.
Build Behavior	 Build Behavior	Allows you to define the build behavior for a specific table. See <a href="#">Build Settings and Data Accumulation Behavior</a> for more information.
Delete		Allows you to delete a table.
--		By default, tables are pinned when you import your data. You can click this icon to unpin

Navigation Pane	Table Icon	Description
		your tables so your tables are organized by Sisense around other pinned tables.

To view your table’s meta, including its name, path, provider, location, and the number of columns it contains, hover over the table in the schema or in the navigation pane.

## Column Level

You can manage tables imported into the ElastiCube through the navigation pane. Columns that are part of a table are displayed below the table. On the left side of the column is an icon that indicates what type of data is in the column, Numeric, Date, or Text. You can click this icon to convert the column into another data type.



On the right side of the column’s name is an icon that indicates if the column is conducted to another column through a relationship. You can hover over the icon to view more details or disconnect the relationship.

For each column, you can hover over its name to display a menu from where you view the following options for managing your columns:

Menu Option	Description
Relationships	Displays related columns together side by side. See Creating Relationships for more information.
Edit Relationship	Allows you to edit a column's relationship. This option is only available when the column is connected to another table.
Rename	Allows you to rename the column.
Duplicate	Allows you to duplicate the column in the table. This is useful if you need the same column with some modifications.
Hide	Allows you to hide a column. Click this option if you want to import the data associated with this column, but do not want this column available to dashboard designers and viewers.
Indexed	Allows you to index the column data for quicker processing in subsequent builds.
Drop	Allows you to remove column from the table.

# Many-to-Many Relationships

**Note:** The image on this page were taken in the desktop version of the ElastiCube Manager, however, the same principles described on this page also apply to the ElastiCube Manager Online.

In databases, a Many-to-Many relationship exists when the value in each field used to create a relationship between tables is included multiple times in each table.

The problem with Many-to-Many relationships is that it can cause duplications in the returned datasets, which can result in incorrect results and might consume excessive computing resources. This section provides solutions and workarounds to common scenarios with many-to-many relationships.

**Example:** A hotel may have a table with reservation data and a table with payment data. Both tables include the name of the guest. A guest can have multiple reservations under their name as well as multiple payments recorded on their name. If a relationship exists between the reservation and payment tables based on the guest's name, a many-to-many relationship is created, as the guest's name appears multiple times in each table.

## Summary of Relationship Types

- **One-to-One Relationship:** In this scenario both sides of the relationship have unique values for every row.
- **One-to-Many Relationship:** In this scenario one side of the relationship will contain unique values for every row, but the other side of the relationship will contain duplicate values for any or all of the corresponding values in the first table.
- **Many-to-Many Relationship:** In this scenario, both sides of the relationship will hold duplicated values, causing excessive calculations for every query run against it.

There are several methods to resolve and bypass a many-to-many relationship. The solution depends on the business model and the logic of the business questions at hand. The following solutions differ by business logic and the schema at hand; each solution can be applied to each schema respectively.

The following sections cover:

- Testing your schema to see if it includes many-to-many relationships
- Understanding which scenario best fits your current schema
- According to your schema logic, applying the respective solution

To check if a relationship is Many-to-Many, you need to check the cardinality of the relationship, and determine the number of unique and duplicate values on each side of the relationship.

## Testing if a Relationship is Many-to-Many

Many-to-Many relationships occur when two tables are joined on a field containing duplicate values on both tables. For example the same guest may have multiple reservations and multiple payments at a hotel, thus joining on the guest between the reservation and payment table would result in a M2M relationship. When testing, if you get the same value for both the unique and duplicate values, then there is no duplication, and this will either be a One-to-Many or a One-to-One relationship. If the number of duplicate values is larger than the number of unique values, then this side of the relationship has duplicated values, and you will need to investigate the other side of the relationship. If the other side of the relationship yields unique values, this is a one-to-many relationship. If not, you have a many-to-many relationship.

Use the following SQL statement to test for potential M2M relationships:

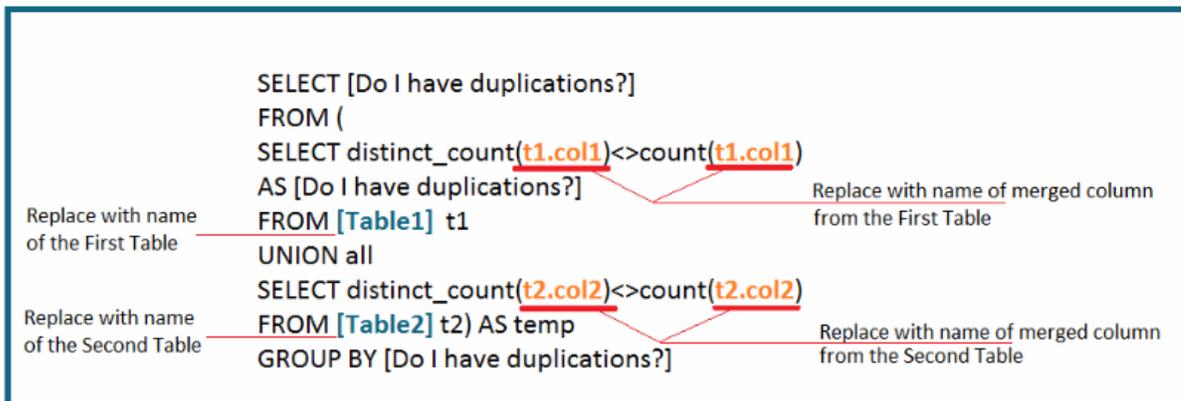
1. In the ElastiCube Manager, open the relevant ecube file.
2. Click **Add Data**> **Custom SQL Expression**.
3. Enter and adjust the SQL statement below.

```
SELECT [Do I have duplications?]  
FROM (
```

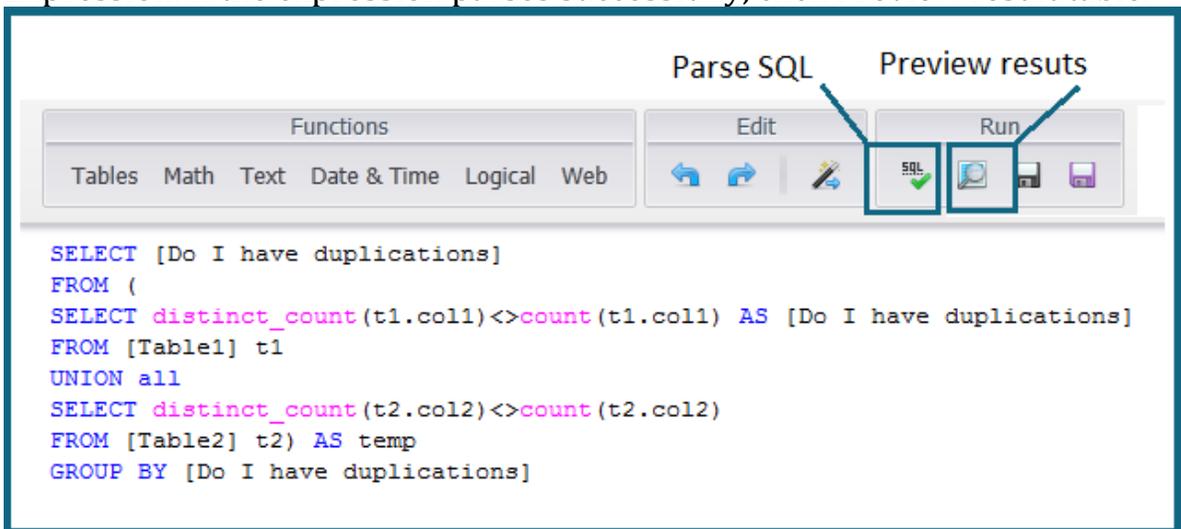
```

SELECT distinct_count(t1.col1)<>count(t1.col1) AS [Do
I have duplications?]
FROM [Table1] t1
UNION all
SELECT distinct_count(t2.col2)<>count(t2.col2)
FROM [Table2] t2) AS temp
GROUP BY [Do I have duplications?]

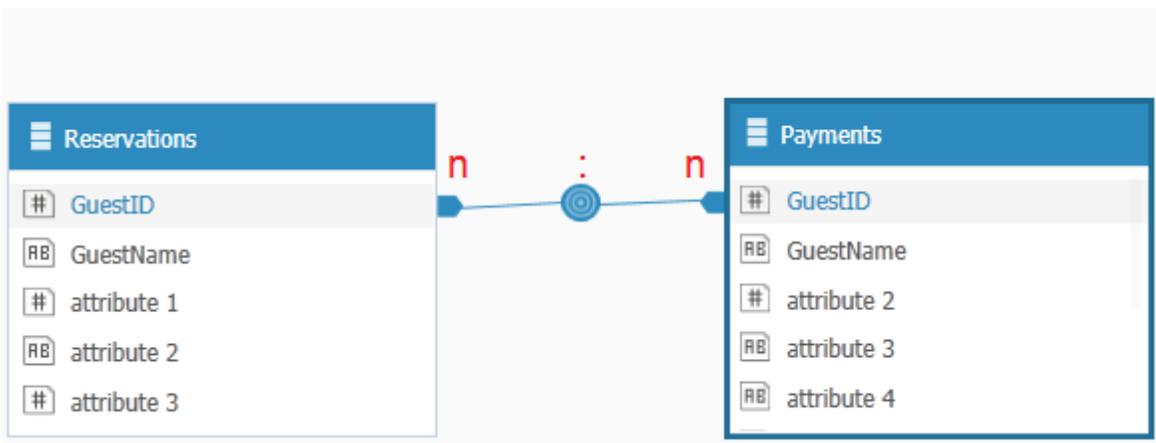
```



4. In the top right of the expression editor window, click Parse SQL Expression. If the expression parses successfully, click Preview result table.



5. If the returned result is 'True' in both lines, a many-to-many relationship exists, and will need to be considered in the ElastiCube design.



## Many-to-Many Relationship Prior to Resolution

If the two values are equal, all guest IDs appear only once, making all values unique. Even if the other side of the relationship has duplicate values for guest ID, this is still a One-To-Many relationship, where the unique values are on the reservations side, and the duplicate values are on the Payments side.

If there are more than two tables connected to this relationship, that is, if there are more than two tables merged on the same field, a few more options exist. The solution for the single many-to-many relationship will be a sub-problem of this scenario. In this case, you'll need to run the test on every table to see the uniqueness or duplication of the merged fields.

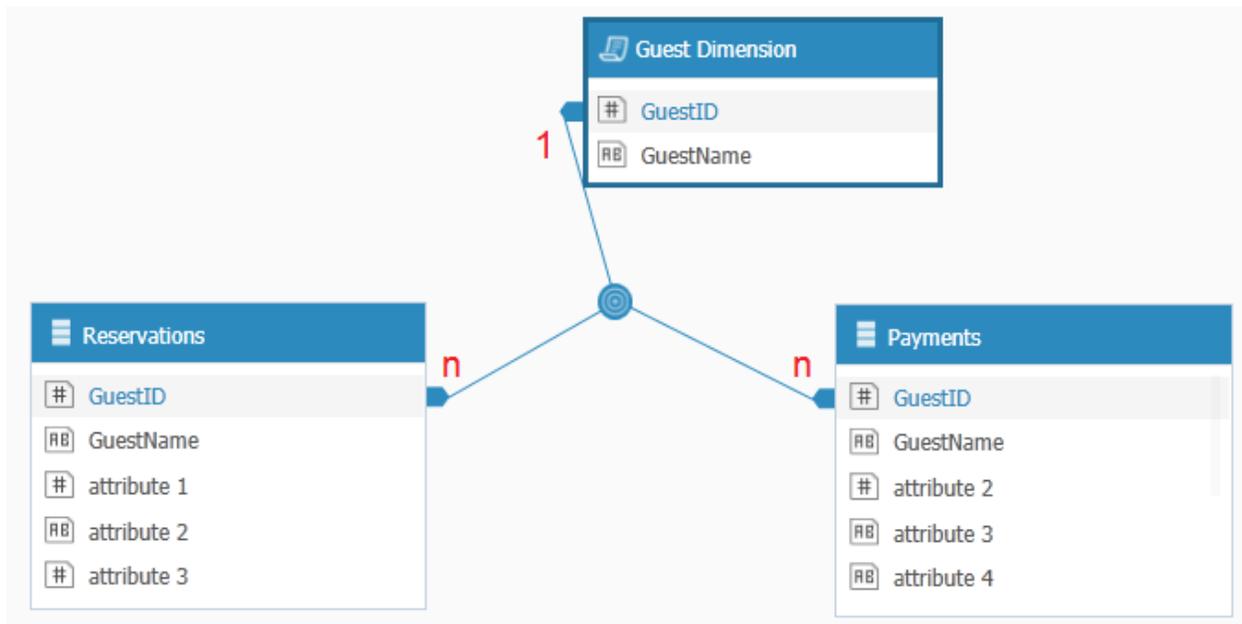
### Two Tables with One Relationship

This section describes two possible workarounds when you have a schema that includes two tables with one relationship:

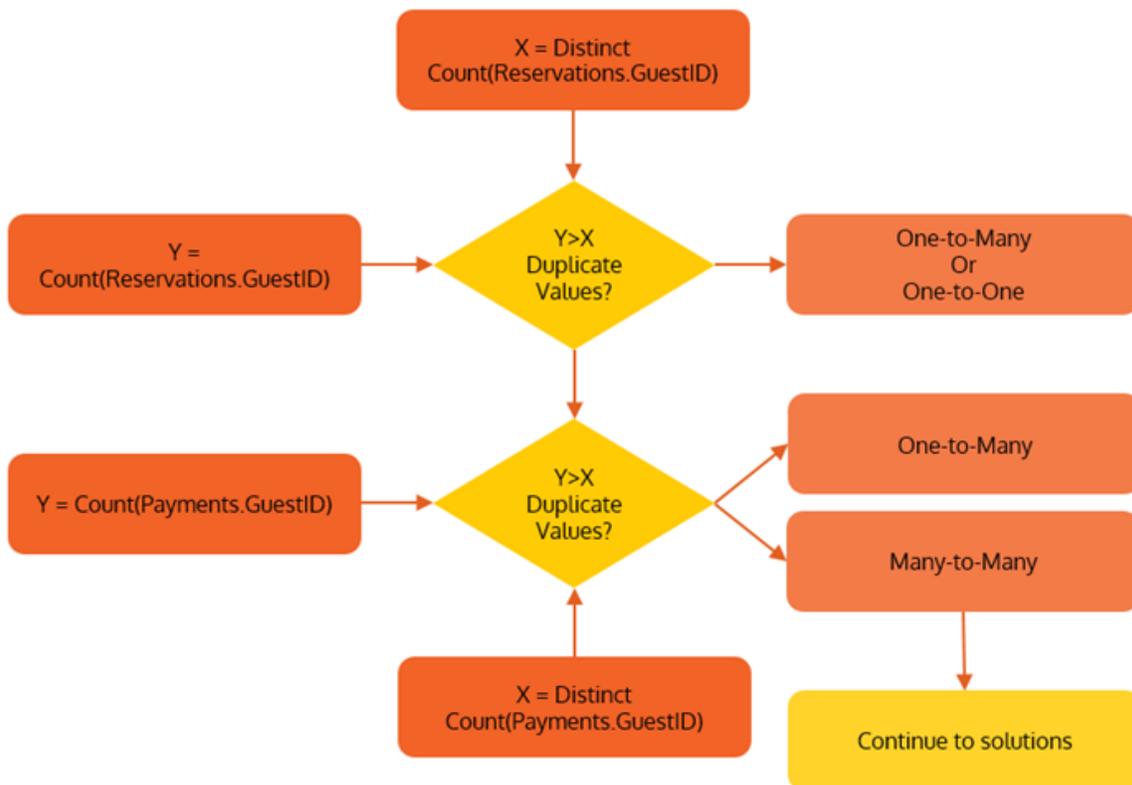
- Two Separate One-to-Many Relationships
- Creating an Aggregated Table

### Two Separate One-to-Many Relationships

The direct solution for such a problem is to break this relationship into two separate one-to-many relationships.



You can use the following diagram to understand the logic behind this testing:



1 – Create a custom SQL expression in the Elasticube. In the expression of this table select all the individual values for the identifier column from both sides. The expression should look like this:

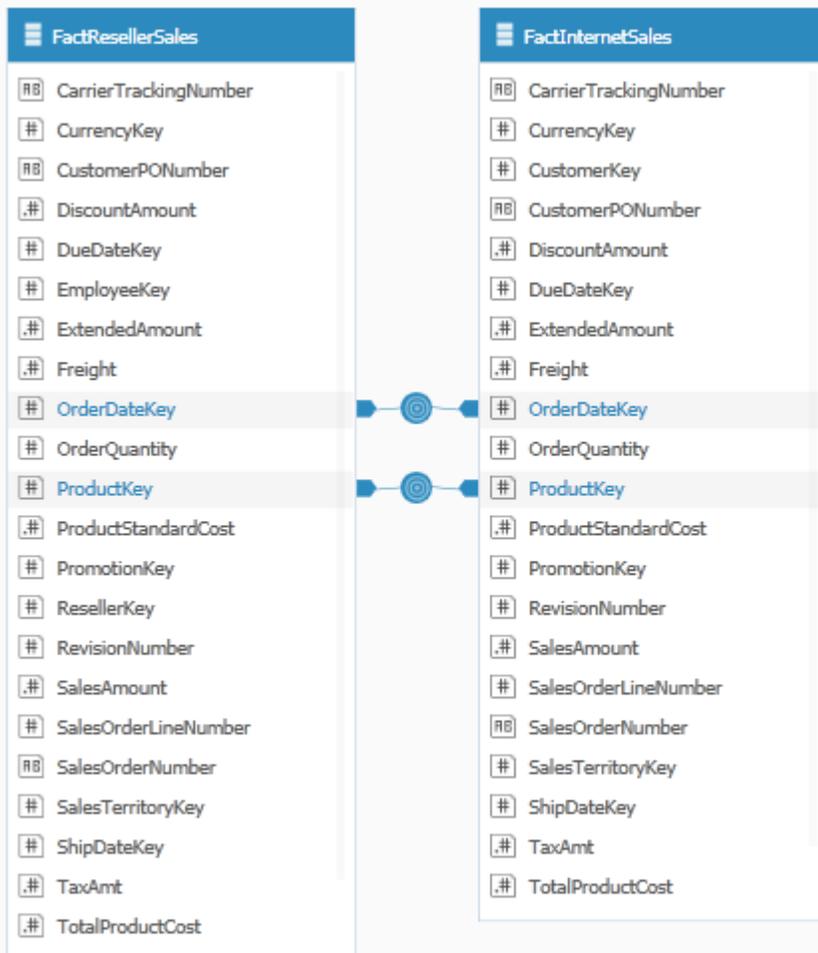
```
SELECT * FROM
(SELECT DISTINCT r.GuestID, r.GuestName
FROM [Reservations] r
UNION
SELECT DISTINCT p.GuestID, p.GuestName
FROM [Payments] p) AS G
```

This query will take all Guest ID values from both tables, and using the UNION statement, will bring in only the unique values from both tables, making this a complete list of all distinct Guest ID values.

2 – Merge the Guest ID field from the new ‘linking’ table to the other two Guest ID fields from the other two tables, thus creating two One-To-Many relationships. You can now use this Guest ID field as the rows or axes elements of a widget, pulling in the unique values from the new Guest Dimension, with measures from the two other tables. See image above.

### **Creating an Aggregated Table**

In situations where you have more than one fact table (a Fact table is a primary table containing the measures or fields used for calculations in the dashboard) in the Elasticube, there are several situations when an aggregated table can resolve a many-to-many relationship.



## Two fact tables

Assuming you want to segment your data according to a few different dimensions, creating relationships directly between these fields can and will create many-to-many relationships in one of two ways, according to the schema:

- Both tables don't hold unique values, and all values from one table are held in the second table. In this scenario either a linked dimension (as described in the first solution – Two Separate One-to-Many Relationships) or an aggregated table can be created which will hold all the unique values and the desired calculations for one of the tables. To create an aggregate table, create a custom SQL expression and aggregate values from the table that includes all values; its own, and the subset present in the other table with the following expression:

```

SELECT i.OrderDateKey, i.ProductKey,
sum(i.DiscountAmount), sum(i.SalesAmount),
avg(i.UnitPriceDiscountPct)
FROM [FactInternetSales] i
GROUP BY i.OrderDateKey, i.ProductKey
SELECT i.OrderDateKey, i.ProductKey, sum(i.DiscountAmount), sum(i.SalesAmount), avg(i.UnitPriceDiscountPct)
FROM [FactInternetSales] i
GROUP BY i.OrderDateKey, i.ProductKey

```

This custom SQL expression will select the distinct OrderDateKeys and their corresponding ProductKeys from the FactInternetSales, grouped by these fields, together with single value aggregations for the different fields, in this case, Discount Amount, Sales Amount and the average unit Price discount. After merging the OrderDateKey and Product Key to the two other tables, you will be able to pull the values from this new table into the rows or axes panel of a widget in the Sisense Web Application with measures and additional aggregations from the two other tables.

**Note:** The non-aggregated table needs to be a subset in terms of the primary fields from the aggregated table.

- Both tables don't include unique values, and there are different values for several fields in both the tables. Resolving this scenario incorporates both solutions mentioned above. In this scenario, create an aggregated table and a dimension table (both described above). The final resolution should look like this:



## Two Fact tables with a date dimension table and an aggregative Products table

### More than Two Tables with More than One Relationship

This section provides two possible workarounds when you have a schema that includes more than two tables with more than one relationship:

Options include:

- Using the Lookup function
- Concatenating two tables into one

### Using the Lookup Function

In most scenarios, you will aggregate values according to a given ID from the unique side of the relationship to the duplicate side. However in specific cases it'll be vice versa.

For example in the following scenario, in which we have three tables, and between them two one-to-many relationships, this can potentially create a many-to-many relationship, if you query the two leaf tables. This means that the query result table will have multiple rows which won't be distinguishable one from the other.



### Two consecutive M-to-M relationships

Using the Lookup Function, you can import values from a remote table by matching values in a different column. This will create a new column in the table where you want to perform an aggregation of a given field(s), with the matching value of the identifying field from the other table. Taking the example of tables T1, T2 and T3, we'd like to run a query that will display aggregations from the duplicate IDs from T1, with a measure from T3. If you run the query as is, you will get multiple values for the query's result set, and we will not be able to run this aggregation. To resolve this, use the Lookup function to import the values from T3 into T2 and then re-run the query only on tables T1 and T2. Using the lookup function, available in the 'Miscellaneous Functions' in the custom SQL editor, you can import the values of 'M3' from the 'T3' table into the 'T2' table. Create a new custom column, and use the Lookup function to import the values of attribute, In this case, the Lookup function should look like this:

```
Lookup([T3],[T3].[M3], [T2].id2,[T3].id2)
```

Running this statement in table T2 will import the matching values of M3 from T3 according to the matching results in ID2 between the two tables.

```
LOOKUP(remote_table,remote_result_column,current_match_column,
remote_match_column)
```

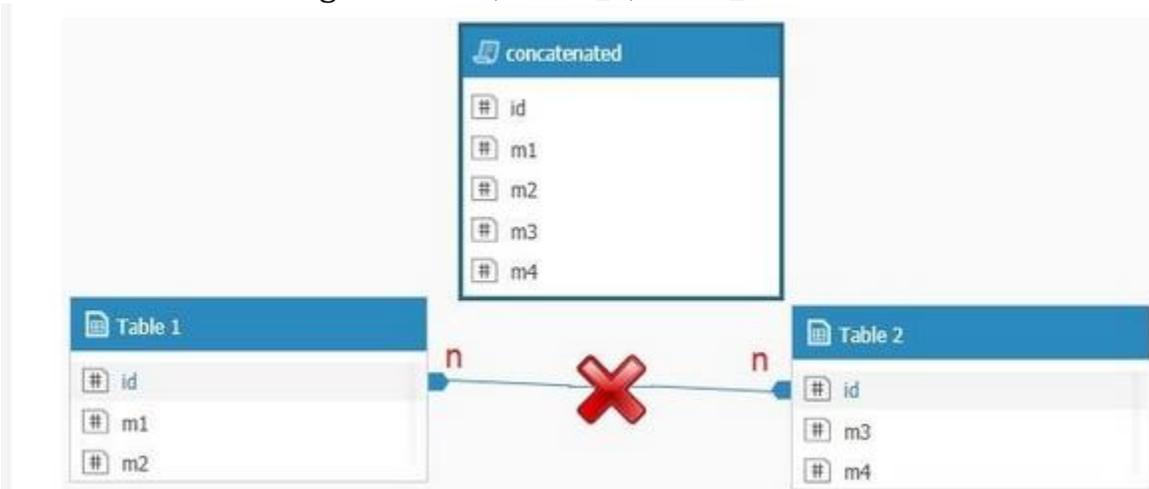
Matches the current value with another value from a remote table. The result will be the value in remote\_result\_column for which the corresponding remote\_match\_column equals the current\_match\_column.



## Two consecutive M-to-O relationships after Lookup fix

### Concatenating Two Tables into One

Assuming you have two separate tables with duplicate ID values in each, and each including different columns for each ID, you can create a new table including all values for every ID, and pull the aggregations from this new table. Note that the two original tables; Table\_1, Table\_2 have different columns.



### Concatenating tables

Using the following SQL statement, you can import the data from both tables, with the IDs and the columns respectively:

```
SELECT s.id AS id, s.m1, s.m2, ToInt( NULL ) m3 , ToInt(
NULL ) m4
FROM [Table 1] s
UNION
```

```
SELECT t.id, ToInt( NULL ) , ToInt( NULL ) , t.m3, t.m4
FROM [Table 2] t
```

This will create a table with five columns:

Id

M1 (from table\_1)

M2 (from table\_1)

M3 (from table\_2)

M4 (from table\_2)

The values missing from each table respectively will be NULL's which will result in the following table.

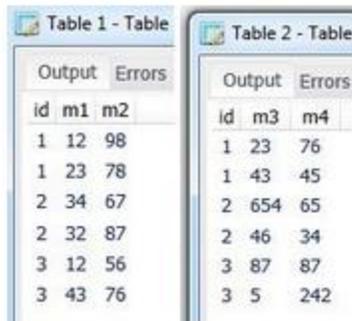
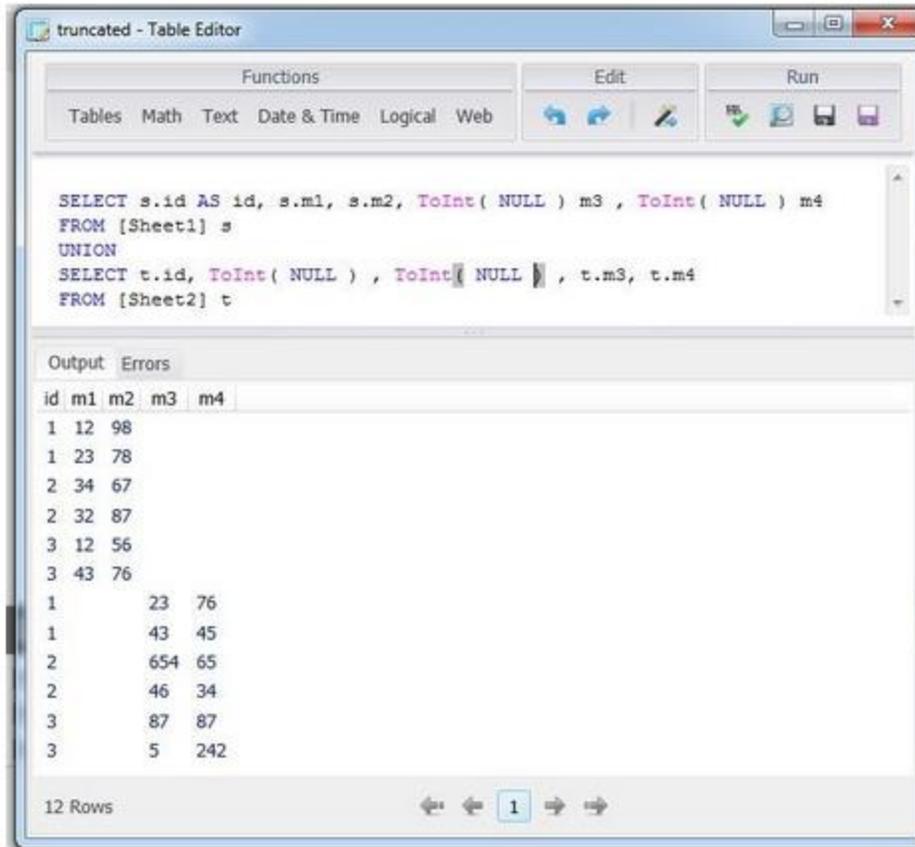
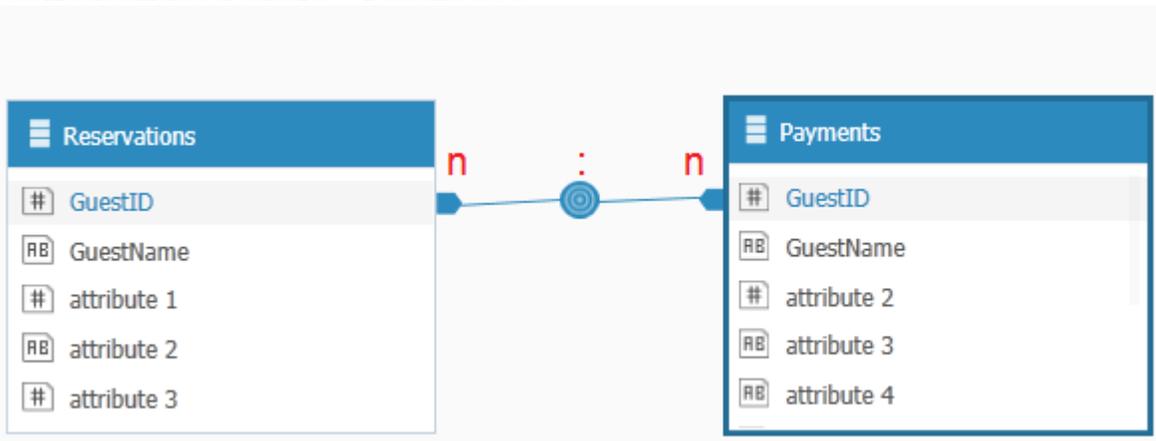


Table 1 - Table		
Output	Errors	
id	m1	m2
1	12	98
1	23	78
2	34	67
2	32	87
3	12	56
3	43	76

Table 2 - Table		
Output	Errors	
id	m3	m4
1	23	76
1	43	45
2	654	65
2	46	34
3	87	87
3	5	242



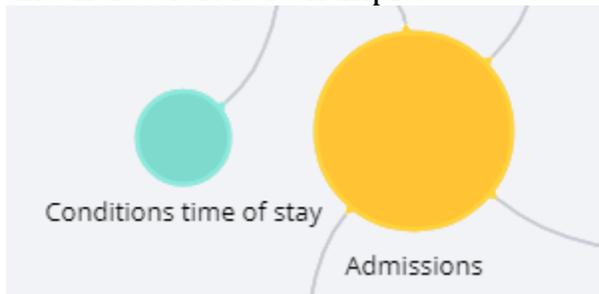
### Concatenated table – result set





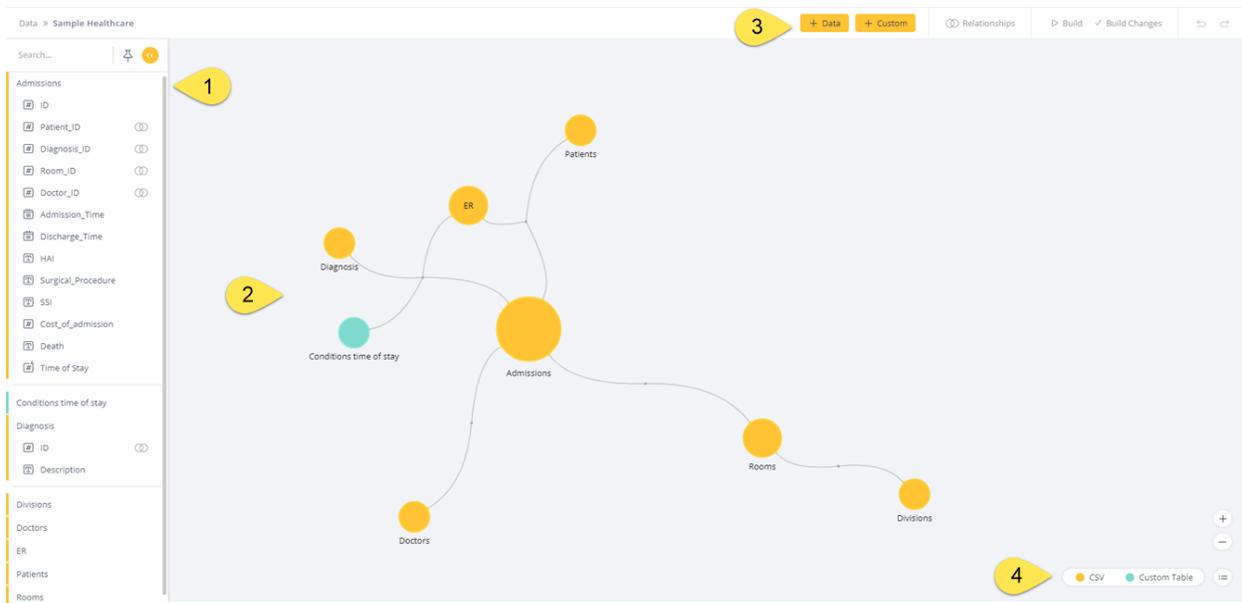
# Navigating the ElastiCube Manager

The first time you create an ElastiCube, its empty until you start adding some data. Once you move some data into the ElastiCube, its represented in the ElastiCube Manager in circular nodes. Each node is a table. The color and size of the table reflects its source and how big the data is in relation to your other tables. So, a small table from one data source will appear as a small circle in one color and another larger table from a different source will appear large and in another color. For example:



The 'Conditions time of stay' table is a custom table and is smaller than the 'Admissions' table from a CSV file.

The ElastiCube Manager contains a variety of information and functionality that lets you prepare your data model as needed. A sample screen is shown below.



- 1. Navigation Pane:** This area contains a list of your tables and their columns. On the left side of the table name is an icon that indicates what type of data is included, Date, Numeric or Text. On the right side is a join icon  when the column is connected to another table. See [Finding Tables and Fields](#) more information.
- 2. Schema:** The schema contains your data model that represents all the data to be added to the ElastiCube and the relationships between the table. Some of your tables might appear with icon attached them. These indicate the status of your build and table. For example,  indicates that the table has changed since the last ElastiCube build while  indicates that a custom expression is currently in draft mode. See [Working with Data](#) for more information.
- 3. Schema Menu:** This menu lists the following buttons:

  - Data:** Click to add a new table to your ElastiCube.
  - Custom:** Click to add a custom table to your ElastiCube.
  - Relationships:** Click to open a preview window where you can join two tables. See [Creating Relationships](#) for more information.
  - Build:** Click **Build** or **Build Changes** to initiate an ElastiCube build.
  - Undo/Redo:** Click these buttons to undo or redo any recent changes to your schema.
- 4. Legend:** The legend lists a description of all your data sources in the ElastiCube and the color that represents them in your schema.

# Previewing Data from a Table

The Preview table lets you see all the values for all the rows and columns in a table. This includes any custom columns, tables, and ETL processes that have been applied to your table.

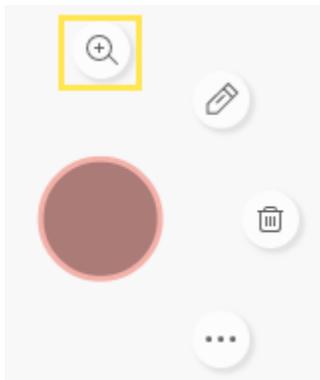
The Preview table makes it easy to verify that the data in your table is accurate and complete.

In addition to displaying a preview of your data, from the Preview table, you can view which columns are connected to other tables and add custom columns.

**Note:** To view underlying table data, you will first need to build your ElastiCube.

## To review the underlying table data:

- Click the relevant table and select the magnifying glass icon to display the Preview table.



**Note:** To view underlying table data, you will first need to build your ElastiCube.

The Preview table contains a variety of information and functionality that relate to your table. A sample preview is displayed below.

The screenshot shows a data table with the following columns: Age Range, Cost, Brand ID, Category ID, Country ID, Date, Quantity, Revenue, and Visit ID. The interface includes a left-hand column filter pane, a top status bar showing 613,002 rows and 11 columns, and a context menu for the 'Category ID' column. Numbered callouts point to: 1. Column Filter (left pane), 2. Connected Table icon (in left pane), 3. Table Details (top bar), 4. Menu Options (context menu), 5. Filter Column (Category ID header), 6. Sort Column (Date header), 7. Add Column (top bar), and 8. Column (top bar).

	Name	Description
1	Column Filter	In the left pane, you can search for columns in your table. As you begin to type, any relevant results are displayed in the left pane. This does not affect the results displayed in the preview.
2	Connected Table	This icon indicates that the column is connected to another table. Click the icon to disconnect the relationship or to view more information about the connection. See <a href="#">Creating and Removing a Relationship between Tables</a> for more information.
3	Table Details	This information describes the number of columns and rows included in the table.
4	Menu Options	This is a list of tasks you can perform on a column. Hover over the column to display the menu and click it to display the options.
5	Filter Column	Click to apply filters to the data in the column. You can filter numeric data and text data. See <a href="#">Filtering Columns</a> for more information.
6	Sort Column	Click to sort a column. Hover over the column name to display the sort icon. You can sort the column by date, numerically, or alphabetically depending on the type of column.
7	Add	Click to add a custom column to the table. See <a href="#">Adding a New</a>

	Name	Description
	Custom Column	<a href="#">Custom Column</a> for more information.
8	Sample Data	Click to see a sample of your data. See <a href="#">Sampling Data</a> for more information.

## Filtering Columns

You can filter data displayed in your Preview table's columns by defining conditions for numeric data and text data. To limit what data is displayed in the preview, you can add multiple conditions that you define. For example, a Text filter can limit what strings are included in the preview.

Filtering the data in the Preview table does not affect the data itself, just what you see in the preview.

Filtering your preview is useful if you need to check specific data to verify its accuracy of your table's completeness.

**To filter entire rows or columns of data from your preview:**

1. In the Preview table, hover over the header in the of the relevant column to display the Filter icon.
2. Click the Filter icon to define the filter's conditions.

Text	Numeric
<div data-bbox="186 394 808 478"> <p>Text Filter <span style="float: right;">Clear</span></p> </div> <div data-bbox="186 478 808 997"> <p>Filter By: Country</p> <p>Containing <span style="margin-left: 10px;">v</span> <span style="margin-left: 10px;">Type here...</span> <span style="float: right;">v</span></p> <p>Add Condition <span style="margin-left: 10px;">+</span></p> </div> <div data-bbox="186 997 808 1104"> <p style="text-align: right;">Apply <span style="margin-left: 20px;">Cancel</span></p> </div>	<div data-bbox="815 394 1433 478"> <p>Values Filter <span style="float: right;">Clear</span></p> </div> <div data-bbox="815 478 1433 892"> <p>Filter By: # Visits</p> <p>&gt; v 0 <span style="float: right;">v v</span></p> <p>Add Condition <span style="margin-left: 10px;">+</span></p> </div> <div data-bbox="815 892 1433 1104"> <p style="text-align: right;">Apply <span style="margin-left: 20px;">Cancel</span></p> </div>

3. Define your filters. Click **Add Condition** to add more conditions to the filter or click **Clear** to erase all the conditions.
4. After you have defined your conditions, click **Apply** to update the preview.

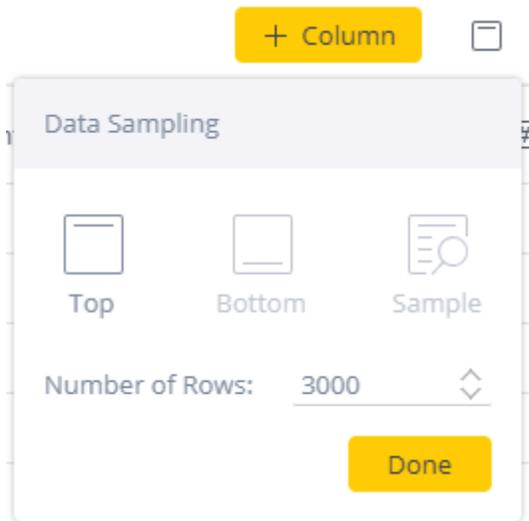
## Sampling Data

If you are working with a large dataset, you may want to restrict the amount of data displayed in your Preview table. Sisense provides three options for displaying a sample of your data:

**Top:** The first rows that are contained at the top of your data.

**Bottom:** The last rows that are contained at the bottom of your data.

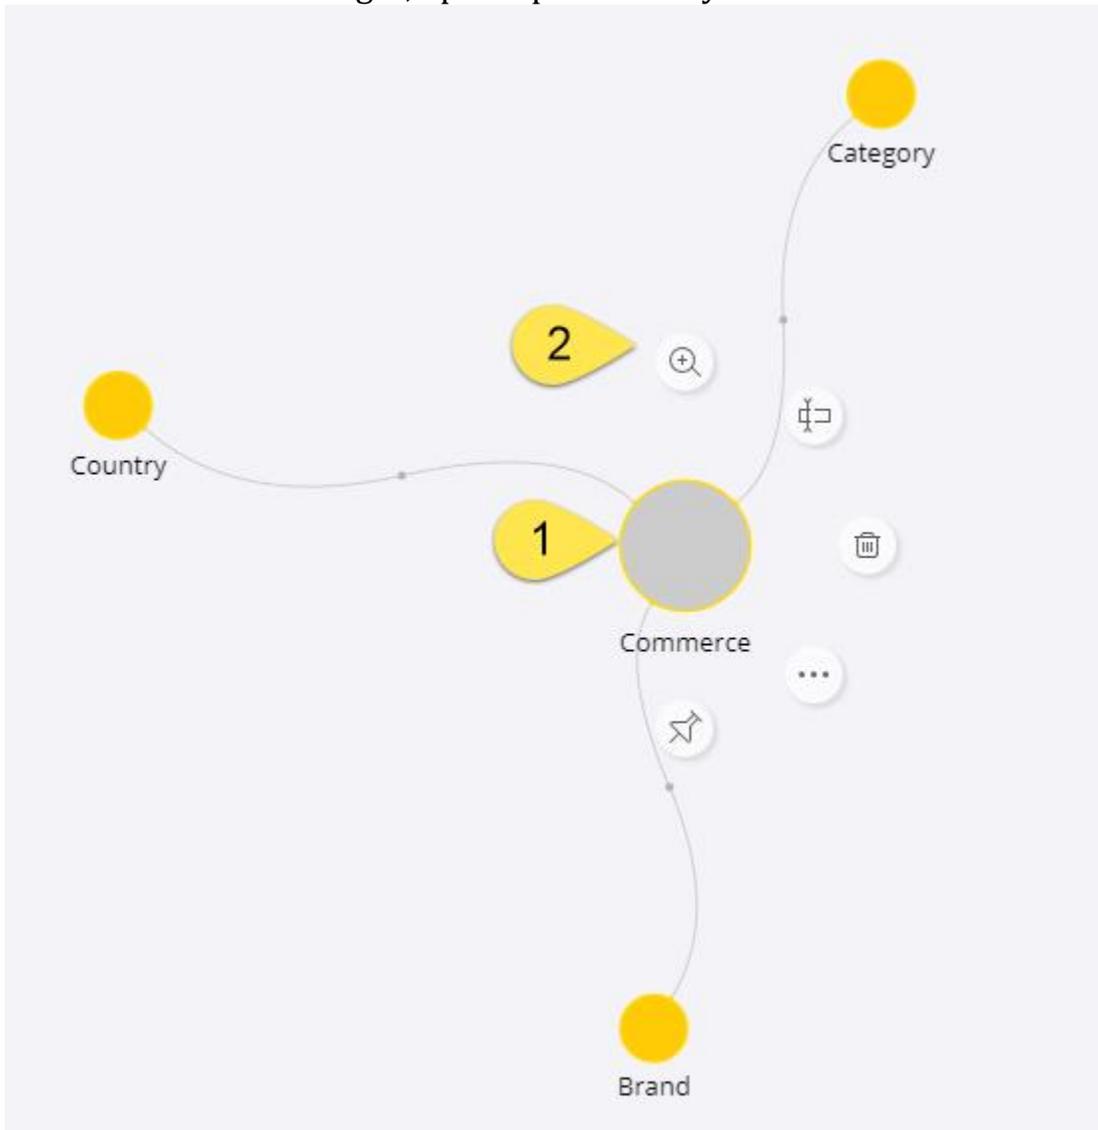
**Sample:** A random selection of rows contained in your table. Sisense selects a sample of your data randomly by rows. If you open the preview again, the same rows will be displayed until your next ElastiCube build.



With each sampling option, you can determine how many rows are display according to the value of the **Number of Rows** field. You can modify this value to display as few as 10 rows or include as many as you like up to the full amount of rows included in the table.

**To view a sample of your data:**

1. In the ElastiCube Manager, open a preview of your table.



2. In the top-right corner of the Preview table, click the **Sampling** Icon.



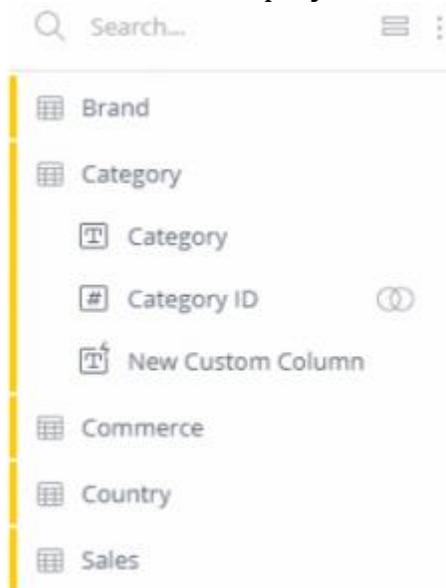
3. Select the relevant sampling option for your data and enter the amount of rows to be included in the sample in the **Number of Rows** field.

The screenshot shows a 'Data Sampling' dialog box. At the top, there is a yellow button labeled '+ Column' and two icons: a square and a square with an arrow. The dialog has a title bar 'Data Sampling'. Below the title bar, there are three options: 'Top', 'Bottom', and 'Sample'. The 'Sample' option is selected, indicated by a magnifying glass icon. Below the options, there is a 'Number of Rows' field with the value '3000' and a dropdown arrow. A yellow 'Done' button is located at the bottom right of the dialog.

4. Click **Done**. Your Preview table is updated.

# Finding Tables and Columns

Through the ElastiCube Search tool, you can locate field tables and fields. The Search tool is displayed in the Navigation pane of the ElastiCube Manager.



When you perform a search, all the relevant tables and columns are returned as you begin to type.

To organize the results displayed in the pane, toggle between the Data Source icon  and the Tables icon .

Each result has an icon on the left side of it. These icons indicate what is being returned.

The meaning of each icon is as follows:

-  Data source
-  Table
-  Text data
-  Numeric data
-  Date data

**To find a table or field in your ElastiCube:**

- On the left-side of the schema, enter the table or field in the **Search** field. As you begin typing the letters or name of the table/field you are searching, the results are displayed below and the relevant table or field is highlighted in the schema.

# Working with Data

Once your data is imported into Sisense, Designers can organize the data in a variety of ways. One of the most common ways is to create a relationship between tables. After you create a relationship, the tables behave as a single table, which can make analyzing your data more easy and accurate.

The following topics describe how you can create relationships with your data and how to avoid some of the pitfalls of working with complex data:

- [Navigating the ElastiCube Manager](#)
- [Finding Tables and Fields](#)
- [Creating and Removing Relationships](#)
- [Previewing Data in a Table](#)
- [Managing Tables and Columns](#)
- [Handling Relationship Cycles](#)
- [Many-to-Many Relationships](#)
- [Chasm and Fan Traps](#)

# Building ElastiCubes

After defining a schema, you will need to ‘build’ your ElastiCube. This will import the data from your data sources that you added to your schema. You must build an ElastiCube at least once before the ElastiCube data can be used in a dashboard.

There are two build options you can select:

▷ **Build** : Imports data for all tables with a default of overriding existing data in the ElastiCube. If tables and fields in your schema are set to import accumulative data, then an additional option **Delete existing Data** exists. If selected, all data will be replaced instead of accumulative data where relevant.

✓ **Build Changes** : Imports data into new tables and fields where a change has been made to the schema since the previous build.

## To build an ElastiCube:

- In the Data page, open the relevant schema and click ▷ **Build** or ✓ **Build Changes**.

A log is displayed that describes the build process. When the build is complete, the following message appears in the logs: Build successfully ended.

## Building ElastiCubes Remotely

**Note:** This functionality is only available in the desktop version of the ElastiCube Manager. This functionality will be available in ElastiCube Manager in the near future.

You can build the ElastiCube locally on your machine or build the ElastiCube remotely from another machine in which Sisense is installed. This is useful for

example when your configuration supports test and production environments.

For more information on setting up production environments, click [here](#).

When connecting remotely to your ElastiCube server must have Sisense installed and a unique Sisense license. While you can have multiple Administrators working on and building an ElastiCube, only one Administrator can work on an ElastiCube at a time to prevent data corruption. In addition, the machine in which the build is performed, must have sufficient resources to be able handle all of your data. See Sisense's minimum requirements for more information.

**To build an ElastiCube remotely:**

1. On the ElastiCube server where the ElastiCube is to be built, open ports 811 and 812 for inbound and outbound access.
2. In the ElastiCube Manager, open the Sisense Server Console by clicking **+**.
3. In the **Connection Settings** window, enter a name for the server and its IP address.
4. Click **Connect**. The list of ElastiCubes for that server are displayed in the Sisense Server Console.
5. In the ElastiCube Manager, click **Build**.
6. Under Target Server, select the new server you have just added.
7. Define your build settings and click **Build**.

# Build Settings and Data Accumulation Behavior

When you build your ElastiCube, data is imported and accumulated based on the following settings:

- **Build Options:** Defines whether to import all data overriding any previous data, or import only data where a change has been made to the schema since the previous build.
- **Field Build Behavior:** Defines a specific integer or date field as an index in the table. The index column will be used to determine whether source records are new, and whether to accumulate data in the table based on a higher index value at the time of the build. New builds will only include values greater than the maximum existing value.
- **Table Build Behavior:** Defines the behavior on the table level. You can replace all data at the time of the build, or add only additional/accumulated data.

## Build Options

There are two primary Build Options, **Build Schema Changes** and **Build Entire ElastiCube**.

To access Build Options, click **Build** in the main menu of the ElastiCube Manager.

See also [Building the ElastiCube](#).

- **Build Schema Changes:** Updates the build only if changes were made to the schema since the previous build, and imports data only if new tables were added. Conditions for updating the build include:
  - New, removed or changes to custom tables
  - New, removed or changes to custom fields
  - New or removed tables
  - New or removed table relationships.

- **Build Entire ElastiCube:** Imports data according to the build behavior defined. By default, Sisense replaces and overrides all existing data in the ElastiCube. While building entire ElastiCubes ensures you have an exact copy of the underlying data source, it can sometimes take a long time and be expensive on the database depending on the size of your data source. You can customize your builds to update only the tables and fields where data was accumulated since the last build. This is useful when you frequently need to refresh a large data source. For example, if you have a data source that is updated daily, rather than rebuild the entire ElastiCube daily, you can just import the new data added each day. This option can significantly reduce the amount of time it takes to complete build depending on the size of your ElastiCube.

While building the ElastiCube with schema changes or entire ElastiCube builds, you can continue to run queries. In the event that the build fails, Sisense restores the original version of the ElastiCube and attempts the build again. While Sisense restores the original ElastiCube, some downtime may occur while Sisense copies and pastes the ElastiCube on your drive. The downtime is the amount of time it takes to copy and paste the ElastiCube locally. If you require high availability for your data, Sisense provides this through ElastiCube Sets. **Note:** When Sisense restores an ElastiCube a copy of the original ElastiCube is created. You should verify that your server can store multiple copies ElastiCubes until the restoration process is complete.

For more information on how to customize your builds, see [Field Build Behavior](#) and [Table Build Behavior](#) below.

## Field Build Behavior

You can select a specific integer or date field that will be used to determine whether to accumulate data at the time of the build. The field will act as an index, and if the index value is greater in a subsequent build then data will be accumulated for the table.

To select the integer or date field within a table as the parameter to accumulate by, click on the field name, and then select **Is Accumulated by**.

Accumulation behavior for integers and for dates are as follows:

## Integer

When you select an integer, only source rows with a value greater than the maximum index value in the ElastiCube table will be inserted. Existing data in the ElastiCube table will not be modified or deleted.

**Note:** Some data sources (for example, CSV files) do not support accumulating data on the field level, and in such cases data will be duplicated. See Accumulative Build Support below.

The following image demonstrates this logic (when supported):

### Accumulate By Index - Integer

Load 1		Load 2	
Source	Ecube Table	Source	Ecube Table
1	1	1	1
3	3	3	3
		2	4
		4	4

After Load 1, the maximum index value is 3. In Load 2, the source index value of 2 is not inserted (since it is less than 3), but the source index value of 4 is inserted (since it is greater than 3).

## Table Build Behavior

A table's Build Behavior lets you define whether to import and replace data for the entire table or bring in only accumulated data at the time of the ElastiCube build.

To select a table's Build behavior, click the table's header and then **Additional Preferences**.

The following options exist:

- **Replace Data (default)**: Replaces all data at the time of the build. This is recommended for dimension tables, for example: store attributes, or dimensions, which describe the objects in a fact table.
- **Accumulate Data**: Adds additional data to existing table data, without making comparisons and without omitting data. This is recommended for detailed fact tables, for example: store quantitative information for analysis. **Note**: If one of the fields in the table is set to accumulate data (see below), then data will be accumulated from the selected integer or date field, and not from the table.
- **Ignore When Unchanged**: Does not import any data unless changes have occurred in the table schema. This is recommended with summary/snapshot fact tables and with data marts (smaller subsets of data, tailored for specific needs).

## Build Options Working Together with Table and Field Behavior

The following table describes how data will be imported using different build options and table behavior combinations.

	<b>Replace Data</b>	<b>Accumulate Data</b>	<b>Ignore when unchanged</b>
Build Entire ElastiCube	Replaces all data at the time of build.	Only adds additional data to existing table data. If a field in the table is marked as Accumulate by, then data is accumulated only for the selected field. Note: If enabled, an additional checkbox exists to delete all	Ignores tables that have remained unchanged, unless changes have occurred in the table schema. Note: If enabled, an additional checkbox exists to delete all

	<b>Replace Data</b>	<b>Accumulate Data</b>	<b>Ignore when unchanged</b>
		existing data during the build.	existing data during the build.
Build Schema Changes	All data remains unchanged during the build.	If a field in the table is marked as Accumulate by, then data is accumulated only for the selected field.	Ignores tables that have remained unchanged, unless changes have occurred in the table schema.

**Note:** Tables with accumulative build settings are highlighted with a dark blue color.

## Accumulative Build Support

The following list displays different data sources and their support for accumulating data on the table and field levels.

<b>Data Source Provider</b>	<b>Accumulate by Table</b>	<b>Accumulate by Field</b>
Oracle	Yes	Yes
MySQL	Yes	Yes
MS SQL	Yes	Yes
PostgreSQL	Yes	Yes
ODBC	Yes	No
OLEDB	Yes	No
Hive	Yes	Yes
MS Excel	Yes	No
CSV file	Yes	No

<b>Data Source Provider</b>	<b>Accumulate by Table</b>	<b>Accumulate by Field</b>
MS Access	Yes	No
Salesforce	Yes	No
Google AdWords	Yes	No
Google Analytics	No	No
Google Spreadsheets	Yes	No
Facebook	Yes	Yes
Facebook Ads	Yes	Yes
HubSpot	Yes	Yes
DynamoDB	Yes	Yes
ZenDesk	Yes	No
Amazon Redshift	Yes	Yes
Heroku Postgres	Yes	Yes
DB2	Yes	Yes
Teradata	Yes	Yes
MongoDB (ODBC)	Yes	Yes

# Introduction to Data Sources

## Sisense Connectors

Connecting to your data is the first step in analyzing your business information. Sisense has native connectors to many SQL databases, NoSQL sources, and popular web applications. Sisense also provides generic frameworks for connecting to many more. Integrating many different data sources enables you to mashup data quickly and easily for more insightful business analytics.

## Sisense Native Connectors

Native connectors are robust connectors to data sources, developed and fully supported by Sisense. Some connectors come pre-installed in the application and others can be downloaded separately for versioning purposes. You can find more information about each source by clicking on the Documentation link.

Connector	Download Link	More Information
Access	Pre-Installed	<a href="#">Documentation</a>
Bing	<a href="#">Download</a>	<a href="#">Documentation</a>
Box	<a href="#">Download</a>	<a href="#">Documentation</a>
CSV	Pre-Installed	<a href="#">Documentation</a>
Custom REST	Pre-Installed	<a href="#">Documentation</a>
DynamoDB	<a href="#">Download</a>	<a href="#">Documentation</a>
Exact Online	<a href="#">Download</a>	<a href="#">Documentation</a>
Excel	Pre-Installed	<a href="#">Documentation</a>

Connector	Download Link	More Information
Generic ODBC Driver	Pre-Installed	<a href="#">Documentation</a>
Generic OLE DB Driver	Pre-Installed	<a href="#">Documentation</a>
Gmail	<a href="#">Download</a>	<a href="#">Documentation</a>
Google AdWords	Pre-Installed	<a href="#">Documentation</a>
Google Analytics	Pre-Installed	<a href="#">Documentation</a>
Google BigQuery	<a href="#">Download</a>	<a href="#">Documentation</a>
Google Sheets	Pre-Installed	<a href="#">Documentation</a>
Heroku Postgres	Pre-Installed	<a href="#">Documentation</a>
Hive	<a href="#">Download</a>	<a href="#">Documentation</a>
Microsoft Active Directory	<a href="#">Download</a>	<a href="#">Documentation</a>
Microsoft Dynamics CRM	<a href="#">Download</a>	<a href="#">Documentation</a>
MongoDB	<a href="#">Download</a>	<a href="#">Documentation</a>
MySQL	Pre-Installed	<a href="#">Documentation</a>
Oracle	Pre-Installed	<a href="#">Documentation</a>
PostgreSQL	Pre-Installed	<a href="#">Documentation</a>
Quickbooks Desktop	<a href="#">Download</a>	<a href="#">Documentation</a>
Quickbooks Online	<a href="#">Download</a>	<a href="#">Documentation</a>
Redshift	Pre-Installed	<a href="#">Documentation</a>
Salesforce	Pre-Installed	<a href="#">Documentation</a>
ServiceNow	<a href="#">Download</a>	<a href="#">Documentation</a>
Sisense ElastiCube	Pre-Installed	<a href="#">Documentation</a>
Snowflake	Click <a href="#">here</a> for instructions.	<a href="#">Documentation</a>

Connector	Download Link	More Information
Splunk	Pre-Installed	<a href="#">Documentation</a>
Stripe	<a href="#">Download</a>	<a href="#">Documentation</a>
SQL Server	Pre-Installed	<a href="#">Documentation</a>
Teradata	Pre-Installed	<a href="#">Documentation</a>
Zendesk	Pre-Installed	<a href="#">Documentation</a>

## Sisense Community Connectors

Community connectors are connectors to certain databases and web applications available to our technical user community. These connectors are not guaranteed by Sisense Support and may require developer expertise to manage or customize. For some connectors, certain limitations on data volume, tables, and refresh will apply.

Connector	Download Link	More Information
Ebay (beta)	<a href="#">Download</a>	<a href="#">Documentation</a>
Facebook	<a href="#">Download</a>	<a href="#">Documentation</a>
Facebook Ads	<a href="#">Download</a>	<a href="#">Documentation</a>
HubSpot	<a href="#">Download</a>	<a href="#">Documentation</a>
Intacct (beta)	<a href="#">Download</a>	<a href="#">Documentation</a>
MailChimp (beta)	<a href="#">Download</a>	<a href="#">Documentation</a>
Microsoft Dynamics NAV (beta)	<a href="#">Download</a>	<a href="#">Documentation</a>
PayPal (beta)	<a href="#">Download</a>	<a href="#">Documentation</a>
Sharepoint (beta)	<a href="#">Download</a>	<a href="#">Documentation</a>
Twitter	<a href="#">Download</a>	<a href="#">Documentation</a>

Connector	Download Link	More Information
Xero	<a href="#">Download</a>	<a href="#">Documentation</a>
Youtube	<a href="#">Download</a>	<a href="#">Documentation</a>
Youtube Analytics	<a href="#">Download</a>	<a href="#">Documentation</a>

## Sisense Live Connect

An alternative to importing data into an ElastiCube is to connect directly with the source through a Live connection. With Live Connect, queries can be processed directly in the data source. This allows you to work with only the result set from a query, rather than doing large imports from a data source. Users may want to leverage this method of connection if they have invested in a high performance database or want to tap into real-time functionality.

Connector	Download Link	More Information
SQL Server	Pre-Installed	<a href="#">Documentation</a>
Amazon Redshift	Pre-Installed	<a href="#">Documentation</a>
Oracle	Pre-Installed	<a href="#">Documentation</a>
MySQL	Pre-Installed	<a href="#">Documentation</a>
PostgreSQL	Pre-Installed	<a href="#">Documentation</a>

Sisense supports Live Connection in the following data sources. All Live Connectors come pre-installed in Sisense\*.

### Connecting to More Data Sources; Frameworks and Data Integration

#### Partners

For connecting to a data source not listed, users can leverage Sisense's generic frameworks such as the Generic ODBC Driver or Custom REST Framework.

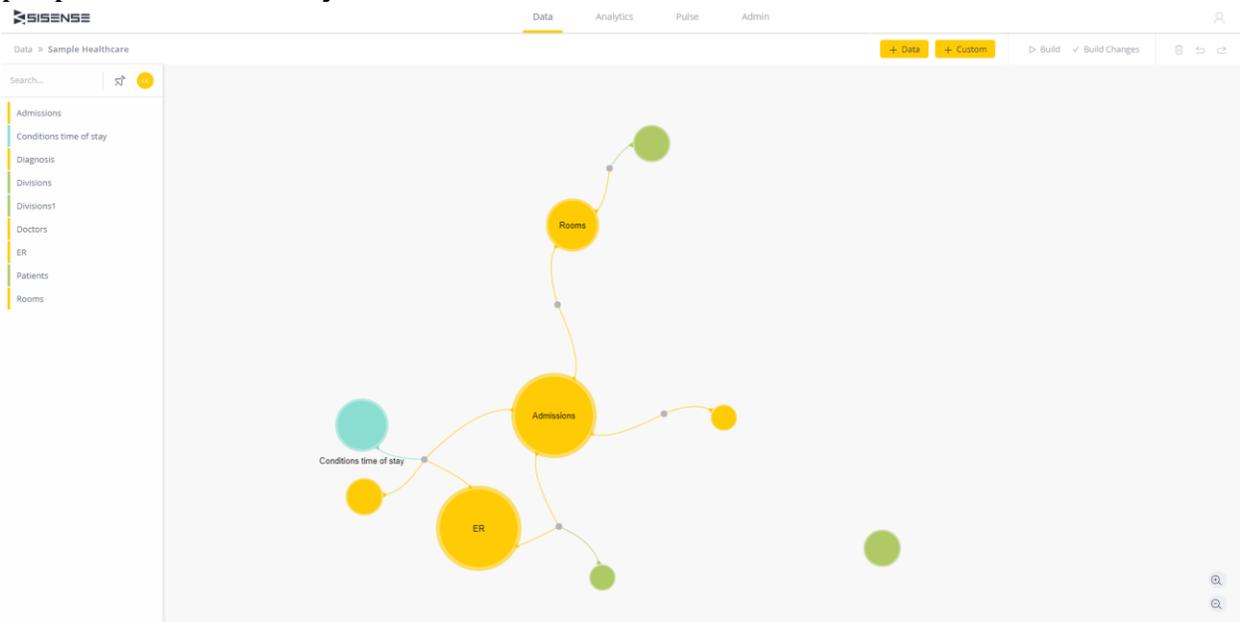
Sisense supports our frameworks and with frequent upgrades and new functionality. However, full connectors to specific data sources created through these frameworks are not supported.

Certain clients may want to use third-party data integration providers, especially if working with large amounts of data from technically intensive web sources.

Our preferred data integration partners include StitchData and Fivetran.

# ElastiCube Manager

Before you can analyze dashboards, you need to add data to an ElastiCube and then build it in the ElastiCube. The ElastiCube Manager is a visual environment where you can import and structure multiple data sources, create relationships between data, and perform ETL (Extract, transform, and load) processes to prepare data for analysis and visualization.



The topics below describe how to access the ElastiCube Manager and how to navigate its interface.

"Accessing the ElastiCube Manager" on page 661

"Navigating the ElastiCube Manager" on page 662

## Accessing the ElastiCube Manager

You can access the ElastiCube Manager through the Data tab displayed at the top of the Sisense Web Application.

Only users with Administrator or Data Administrator privileges can see the Data tab.

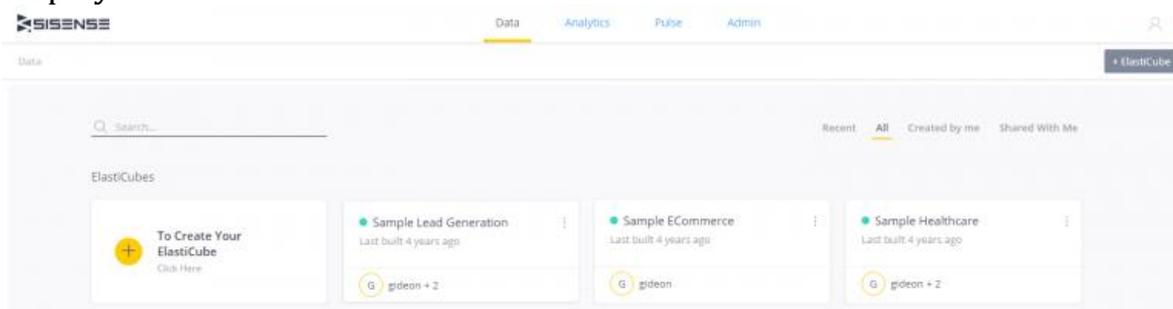
Data      Analytics      Pulse      Admin

There are two ways from the Data page to access the ElastiCube Manager, by creating a new ElastiCube or by editing an existing ElastiCube.

For more information about ElastiCubes, see [Managing ElastiCubes from the Data Page](#).

### To access the ElastiCube Manager:

1. In the top of the Sisense Web Application, click **Data**. The Dashboard list is displayed.



2. For new ElastiCubes, click **+ ElastiCube**.  
For existing ElastiCubes, click the relevant ElastiCube.  
The schema is displayed.

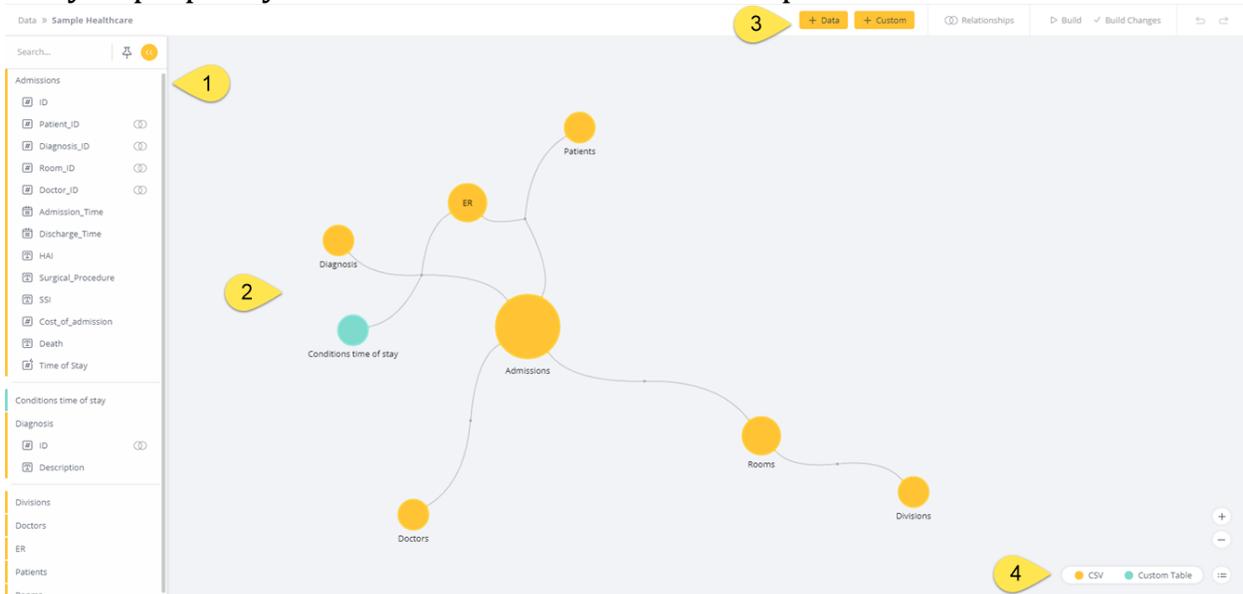
## Navigating the ElastiCube Manager

The first time you create an ElastiCube, its empty until you start adding some data. Once you move some data into the ElastiCube, its represented in the ElastiCube in circular nodes. Each node is a table. The color and size of the table reflects its source and how big the data is in relation to your other tables. So, a small table from one data source will appear as a small circle in a color and another large table from a different source will appear large and in another color. For example:



The Conditions time of stay table is from a custom table and is smaller than the Admissions table that is from a CSV file.

The ElastiCube Manager contains a variety of information and functionality that lets you prepare your data model as needed. A sample screen is shown below.



1. **Navigation Pane:** This area contains a list of your tables and their columns. On the left side of the table name is an icon that indicates what type of data is included, Date, Numeric or Text. On the right side is a join icon  when the column is connected to another table. See [Finding Tables and Fields](#) more information.
2. **Schema:** The schema contains your data model that represents all the data to be added to the ElastiCube and the relationships between the table. Some of your tables might appear with icon attached them. These indicate the status of your build and table. For example,  indicates that the table has changed since the last ElastiCube build while  indicates that a custom expression is currently in draft mode. See [Working with Data](#) for more information.

3. **Schema Menu:** This menu lists the following buttons:
  - Data: Click to add a new table to your ElastiCube.
  - Custom: Click to add a custom table to your ElastiCube.
  - Relationships: Click to open a preview window where you can join two tables. See [Creating Relationships](#) for more information.
  - Build: Click Build or Build Changes to initiate an ElastiCube build.
  - : Click these buttons to undo or redo any recent changes to your schema.
4. **Legend:** The legend lists a description of all your data sources in the ElastiCube and the color that represents them in your schema.

# ElastiCubes

## Introduction

The ElastiCube is Sisense's unique, high-performance analytics database with super-fast data stores that are specifically designed to withstand extensive querying typically required by business intelligence applications.

ElastiCubes allow you to bring in data from multiple sources, and then merge, manipulate and query the data as if it was one consolidated data set. ElastiCubes perform so well, that in most cases the creation of dedicated OLAP cubes and/or optimized data marts are completely unnecessary – even when dealing with hundreds of millions of rows of raw data.

One of the biggest advantages of ElastiCubes is the ability to easily mash up multiple data sources. It is made up of fields where each value in one field has a corresponding value in another field. The data for an ElastiCube can come from one source, multiple sources or even from multiple physical locations. Once the data is inside the ElastiCube, it is all the same and every field coming from every table can be analyzed in the context of any other – quickly.

## Benefits

ElastiCube technology make queries over hundreds of millions of rows of raw data return in seconds, with moderate hardware requirements including standard desktop-class computers with commodity hardware. More importantly, ElastiCubes can do this without having to pre-aggregate and pre-calculate the data ahead of time and store it on the hard-drive, thus radically reducing required import/processing time and storage space.

ElastiCubes are most useful when one or more of the following is true:

- Large amounts of data need to be analyzed
- Data for analysis originates from multiple disparate sources

## ElastiCubes – Technical Overview

Relational databases (RDBMS) like SQL Server, Oracle, MySQL and even Access all store tabular data row-by-row. This structure is best for transactional/operational systems that require large numbers of concurrent insertions. With indexes, it can also provide realistic query response times for row-based queries that do not frequently require aggregations or joining of many tables.

Data analysis often requires aggregation of data as well as merging of data located in multiple disparate tables. When dealing with these types of queries, relational databases reach their limits pretty quickly. The only way to extend these limits is by putting in stronger hardware and pre-aggregating data to reduce the amounts of calculations that occur in real time.

### The ElastiCube Columnar Database

ElastiCube data is held in a Columnar Database Management System (CDBMS) that stores data field-by-field. Each field is individually stored in a memory-mapped file, the same mechanism the Windows operating system pagefile uses for memory dumping and loading.

When a query is executed over an ElastiCube, only fields referenced in the query need to be loaded into memory. This leaves enough space for actually processing the query entirely in memory without any read/write to the hard-drive – the prime reason for poor performance of queries. Once a field is no longer used, it is removed from memory and its consumed space is freed.

This approach has several advantages:

### **Query Response Time**

Queries over data sets containing millions of rows of data return in seconds even under modest hardware configurations such as desktop computers.

### **Materialization Time**

ElastiCubes do not require pre-aggregations and/or creation of indexes to assure fast query response, therefore the actual creation of an ElastiCube takes a fraction of the time of a data mart or an OLAP cube.

### **Storage Space**

Pre-aggregations and the creation of indexes are not needed to assure fast query response, making an ElastiCube's size significantly smaller than a datamart or an OLAP cube.

### **High Compression**

This columnar storage strategy makes the data much more suitable for high levels of compression, without loss of detail or accuracy. This means less hardware is needed; less disk space and less RAM than for an equivalent-sized, traditional Business Intelligence DB.

### **64-bit Support**

Written and designed to natively support 64-bit processing, it vastly increases the amount of memory the system can address at any given time. 64-bit architecture means you can work with virtually unlimited amounts of data.

### **True Multi-User, Multi-Application Architecture**

ElastiCubes are not tightly coupled with the application layer of the system. This frees up a single ElastiCube to handle multiple applications and users. Not having to reproduce your data model for every application saves significant time developing and maintaining your dashboards and reports.

## **Just-In-Time, In-Memory Processing**

### **Smart Cache and Instruction Recycling**

CPU cycles and RAM space are the two most precious resources in any computer, and ElastiCube is designed to use both as efficiently and speedily as possible. Using our sophisticated caching algorithm, the data is only loaded into memory when it's needed. As part of this algorithm, compute- and time-intense calculations are also intelligently cached to further reduce I/O calls.

### **Cache-aware Algorithm**

Additional sophisticated algorithms further increase Sisense's performance. Once data is loaded into memory, the main performance bottleneck becomes CPU cache misses that naturally come with random access. The ElastiCube is specifically designed to minimize these errors by employing a unique cache-aware algorithm, further increasing Sisense's performance by an additional order of magnitude.

### **Compressed Calculations**

Every DB compresses data to save disk space and RAM. ElastiCube is designed to work directly on this compressed data, so that the need for decompression is virtually eliminated, further increasing ElastiCube's performance.

## **Designed with Standard Hardware in Mind**

Just about every new computer on the market—even portables like iPhones and iPads—are built with very powerful multi-core processors, putting several CPUs into one. ElastiCube was built specifically to take advantage of these powerful CPUs, further increasing Sisense's performance on standard hardware, enabling you to run multiple applications and support multiple users.

### **Highly RAM-efficient**

The thing we know for sure about DBs is that they grow. Fast. So no matter how much fancy footwork is done with completely in-memory DBs, eventually you run out of RAM space and need to upgrade—at least your RAM (best case) or your entire hardware platform (worst, very expensive, case). At Sisense we know

this, so we spent years designing the ElastiCube to be able to handle terabytes—billions of rows—of data efficiently and quickly, even on standard PC hardware.

## Unified Analytics Engine

Sisense can execute queries against a wide variety of data sources as if they were all of the same type, essentially making the individual characteristics of each physical data source unimportant. Our Unified Analytics Engine is what makes this possible.

When Sisense imports data, the Unified Analytics Engine creates a metadata layer, or abstraction layer, which is then used to formulate queries across any number of tables from any number of data sources in any number of formats. It even supports the combined querying of resident and external (live) database sources without first loading data into the database!

These capabilities provide the user with unparalleled flexibility and speed in creating, executing and sharing highly complex reports, dashboards, and analytic applications, with any number and variety of data sources.

## Compliant with Industry Standards

### **Supports SQL-92 Standard**

Even with all this advanced technology, we knew that none of it would be any good if our users couldn't access their existing data. So, we built in an SQL layer to the system, which allows users to integrate Sisense to external applications without needing to learn new scripting languages.

### **Seamless Integration with Existing Data Sources**

Got an ODBC/OleDB compliant DB today? Great, we built in the ability to access those, too. ElastiCube will seamlessly connect to those data sources so, again, there is no need to learn a new language or write special code to connect to your

existing data. With ElastiCube there's no need to start over, you just get faster, easier, and more scalable, with minimal need for IT.

Queries over data sets containing millions of rows of data return in seconds even under modest hardware configurations such as desktop computers.

### **Materialization Time**

ElastiCubes do not require pre-aggregations and/or creation of indexes to assure fast query response, therefore the actual creation of an ElastiCube takes a fraction of the time of a data mart or an OLAP cube.

### **Storage Space**

Pre-aggregations and the creation of indexes are not needed to assure fast query response, making an ElastiCube's size significantly smaller than a datamart or an OLAP cube.

### **High Compression**

This columnar storage strategy makes the data much more suitable for high levels of compression, without loss of detail or accuracy. This means less hardware is needed; less disk space and less RAM than for an equivalent-sized, traditional Business Intelligence DB.

### **64-bit Support**

Written and designed to natively support 64-bit processing, it vastly increases the amount of memory the system can address at any given time. 64-bit architecture means you can work with virtually unlimited amounts of data.

### **True Multi-User, Multi-Application Architecture**

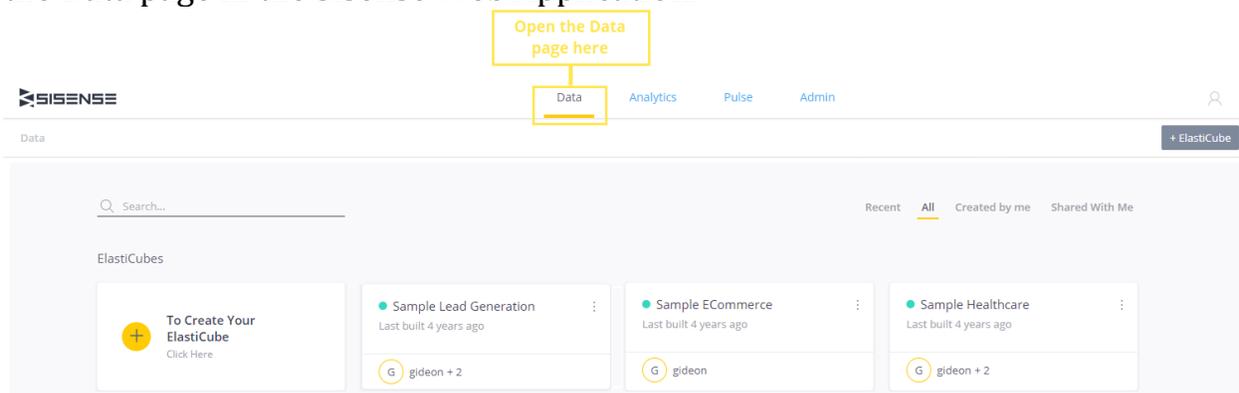
ElastiCubes are not tightly coupled with the application layer of the system. This frees up a single ElastiCube to handle multiple applications and users. Not having to reproduce your data model for every application saves significant time developing and maintaining your dashboards and reports.

# Manage Data

This section describes how to import data into Sisense and avoid common mistakes when preparing it.

## Data Page

In Sisense, Designers can import data into an ElastiCube from a variety of sources. You can see a list of your ElastiCubes and create new ElastiCubes from the Data page in the Sisense Web Application.



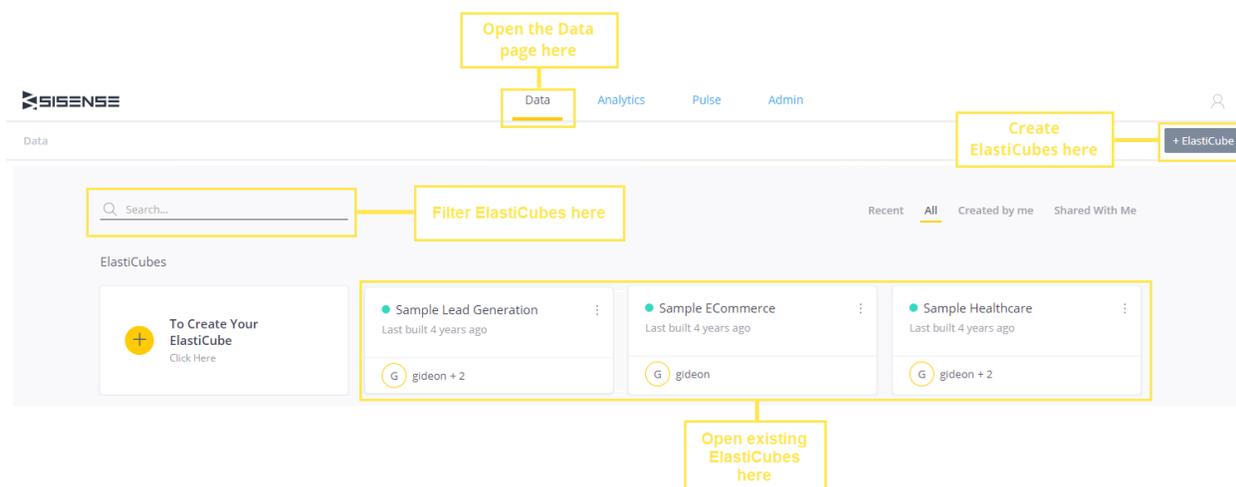
### Related Topics

- [Managing ElastiCubes](#)

## ElastiCubes

The ElastiCube is Sisense's proprietary, high-performance analytical database with specifically designed to withstand extensive querying typically required for your business intelligence application.

You connect to your data sources and import your data into an ElastiCube through the ElastiCube Manager, from where you can then design your ElastiCube's schema.



## Related Topics

- [ElastiCubes](#)
- [ElastiCube Manager](#)

## Schemas

After you create an ElastiCube, the next step is to start adding some data to it. Once your data is there, you need to prepare it before you build the ElastiCube. If your not familiar with working with data, start with Navigating the Schema Editor to learn how your way around.

## Related Topics

- [Navigating the ElastiCube Manager](#)
- [Introduction to Data Sources](#)
- [Sisense Live Connect](#)
- [Working with Data](#)
- [Transforming and Enriching Data](#)
- [Functions Overview](#)

# ElastiCube Builds

After you have designed your schema, it's time to build the ElastiCube. Building an ElastiCube takes all your data from all your sources, and imports it into the ElastiCube.

## Related Topics

- [Building ElastiCubes](#)
- [Build Settings and Data Accumulation Behavior](#)
- [Elasticubes for Advanced Business Scenarios](#)

# Live Connect

Sisense Live Connect queries are run directly against the data source. Sisense performs no additional caching on the data.

Having the right data available at the right time is absolutely critical to making smart business decisions. To ensure that you have the latest data available, Sisense provides live connections to some of the most popular data sources. Traditionally, your data was only as fresh as your latest ElastiCube build. With live connections, you can refresh your dashboards manually or set them to refresh automatically every few seconds.

When you add a live connection, Sisense creates a data set. A data set includes connection and credential details to the data source and defines the table where live data is pulled from. Each data set supports a single table from your data source. After creating a data set, it is displayed in the ElastiCube page of the Admin section from where you can manage all your data sets.

The screenshot shows the Sisense Admin interface. The top navigation bar includes 'Home', 'Pulse', 'ElastiCube Manager', and 'Admin'. The left sidebar contains navigation options: 'Admin', 'Users', 'Groups', 'ElastiCubes & Data Sets', 'Active Directory', 'Single Sign On', 'Trusted Servers List', 'System Configuration', and 'REST API'. The main content area is titled 'ElastiCubes & Data Sets' and features a search bar and buttons for 'Refresh List', 'Add Build Alert', 'Add Server', and 'Add Live Connection'. It is divided into two sections: 'Local Machine' and 'Live Data Sets'.

ElastiCube	Status	Created On	Last Built	Size On Disk	
Nortwind BusinessDays 2	Stopped	5/31/15 5:02 PM	10/24/16 2:51 PM	4.72MB	Open
Sample eCommerce	Running	2/10/14 2:21 PM	3/12/14 10:18 AM	59.05MB	Open
Sample Healthcare	Running	7/2/13 1:07 AM	3/12/14 2:01 PM	1.56MB	Open
Sample Lead Generation	Running	3/3/14 7:46 PM	3/12/14 11:59 AM	2.22MB	Open
Test	Stopped	7/18/16 1:06 PM	1/24/17 3:25 PM	39.49MB	Open

Data Set	Source	Server Location	User Name	Refresh Rate	
Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	
Copy of Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	
customers	MySql Database	192.168.1.229	Test	10 Sec.	
datatypes	MySql Database	192.168.1.229	Test	10 Sec.	
LiveSource	Microsoft Sql Server	192.168.1.229	Test	2 Min.	
MsSQL Commerce	Microsoft Sql Server	192.168.1.229	Test	2 Min.	
productlines	MySql Database	192.168.1.229	Test	10 Sec.	

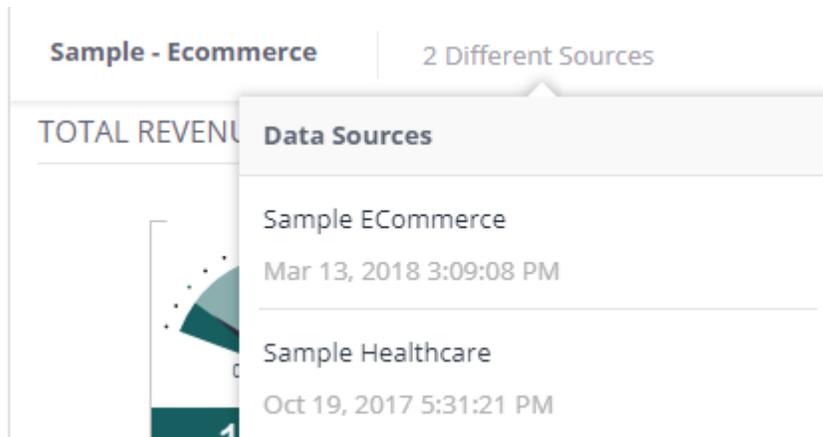
# Adding Live Connections

Each live data source has its own requirements. For documentation for the relevant live data sources, see [Sisense Live Connect](#).

## Working with Data Sets

After you have added your data set, you can begin to create visualizations like you would with any other data source. See [Creating Dashboards](#) and [Creating Filters](#) for more information.

In the Sisense Web Application, live data sources have the ⚡ icon next to their name.



While you can create dashboards from a live data set like you would an ElastiCube, you can also add data sets as a data source to a dashboard that already has a data source such as an ElastiCube. When have multiple data sources working together on a single dashboard, the data source and the widgets built on it operate independently of each other. This allows you to view widgets that visualize time sensitive information in addition to widgets that are refreshed less frequently. For more information about adding multiple data sources, see [Changing a Dashboard's Data Source](#).

In terms of governance, all existing access rights settings and data security settings can be applied for live connections.

## Load Management

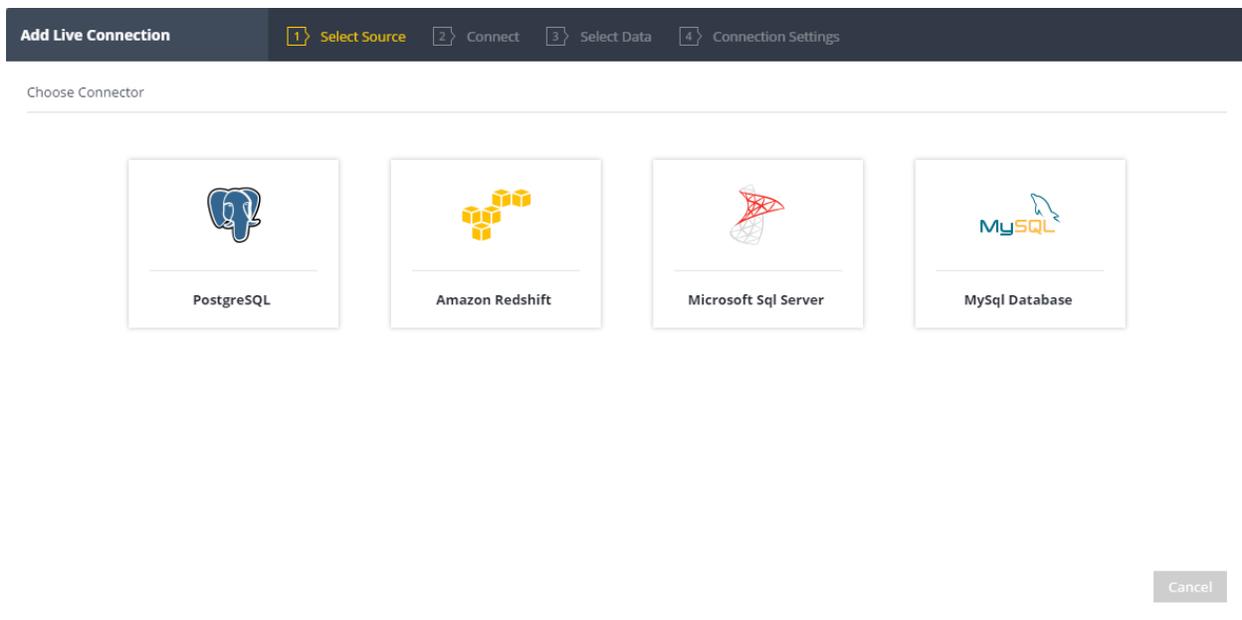
The most significant difference between dashboards built on ElastiCubes and dashboards connected to live data sources is how the load is managed. With an ElastiCube, you refresh the data when you build an ElastiCube and query the ElastiCube, not the data sources that define your ElastiCube. With live connections, queries are run against the data source itself. This means that you don't have to wait for an ElastiCube to finish building before you see the results, however, your data source must be able to handle the potential load of multiple refreshes.

When you refresh the data for a live connection, Sisense uses the connection settings and credentials in the data set to query the live data source. Depending on your Sisense role, there are several options for refreshing the data.

Administrators create the data set. They define the refresh rate in seconds and timeout for each data set. After the refresh rate is set by the Administrator, no user can set a refresh rate shorter than the rate set by the Administrator.

Administrators can manually refresh the data set from the ElastiCube & Data Set page in the Admin section.

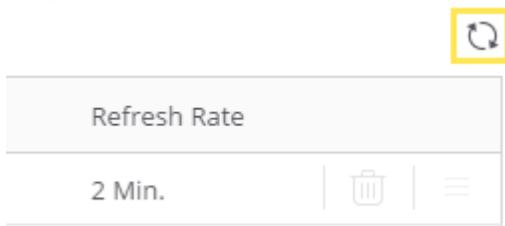
Keep in mind that each time a refresh takes place, your data source is queried, so your data source needs to handle the refresh rate you define. Designers can reduce the frequency of data refreshes or stop automatic refreshes on a per widget level, but they cannot increase the frequency of refreshes beyond the threshold set by an Administrator.



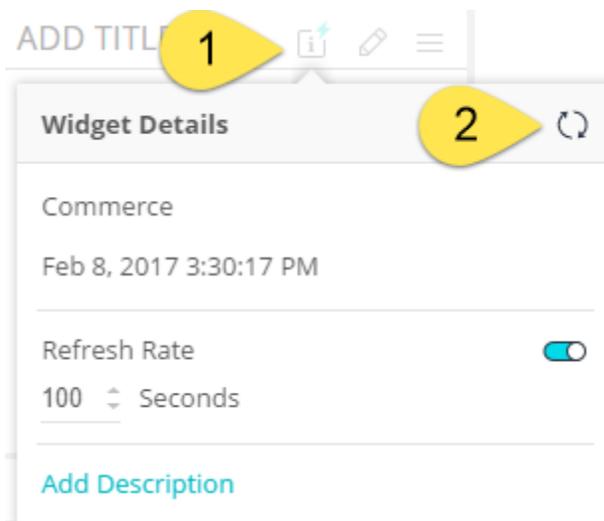
## Refreshing a Data Set

There are several ways in which you can refresh a widget build on a live data set depending on your Sisense role.

Administrators set the default refresh and can refresh the data manually from the ElastiCubes & Data Sets page in the Admin section of the Sisense Web Application.

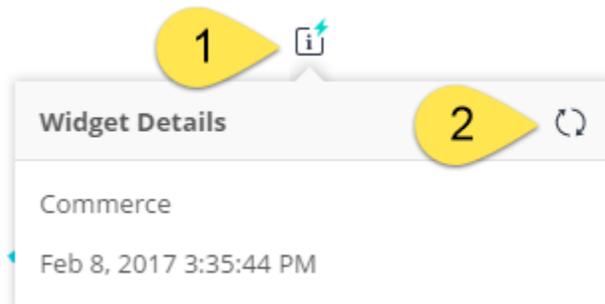


Designers can refresh a widget from the widget's information window .



In addition, Designers can reset the Refresh Rate, however, the rate cannot be lower than the default rate set by the Administrator.

Like Designers, Viewers can refresh a widget from the widget's information window, but they cannot modify the refresh rate.

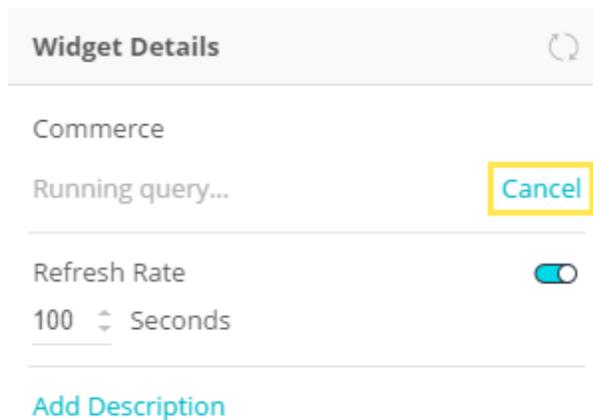


While any users can refresh a widget, Sisense has a minimum 10 second refresh limit. Any refreshes requested within 10 seconds are ignored.

## Canceling a Refresh

If you have a refresh that you want to cancel, you can cancel it by clicking **Cancel** in the widget's information box.

This button is displayed only when a refresh is in progress.



In addition, any time a user leaves a query, for example, to view another page in their browser, Sisense pauses the query.

## Limitations

The current features are not supported by live connections:

1. Pulse alerts
2. Bow & Whisker
3. R is not currently supported
4. Some functions are not currently supported
5. Changing live data sources for dashboards

## Frequently Asked Questions

**Q:** I cannot connect to my data source?

**A:** As the connection is live, the data source must be available for Sisense to connect. Also, if your credentials changes, you must update the Administrator must update their data set so Sisense can continue to connect and refresh the data.

**Q:** No data sources are displayed when trying to add one?

**A:** In Windows Services, verify that the following are running, and if not, restart them:

- Sisense.CLRConnectorsContainer
- Sisense.Discovery
- IIS

**Q:** If I use Windows Authentication with my SQL server, what do I enter in the Username and Password fields?

**A:** Enter the Location and leave the Username and Password fields blank.

**Q:** After entering a data set name, the Done button is not displayed:

**A:** Make sure the data set's name is unique. You cannot create data sets with the same name.

# Making Changes to an Existing ElastiCube

This section describes how to make changes in your ElastiCube.

## To update an ElastiCube:

### In web-based version of the ElastiCube Manager:

- Open the Data page and click the relevant ElastiCube. The schema of the ElastiCube is displayed.

Changes are saved as you work and take affect when you build the ElastiCube.

### In the desktop version of the ElastiCube Manager:

1. Open the ElastiCube Schema file using one of the following methods:
  - In the ElastiCube Manager, click File, and locate the ElastiCube Schema file (\*.ecube), and double click to open.
  - In the **Sisense Server Console**, select the ElastiCube you want to open, and click **Edit**.

The file will open displaying all associated tables, fields and relationships in the ElastiCube Manager.

2. To apply any changes, click **Save > Build the ElastiCube...** in the top panel of the ElastiCube Manager.

# Adding Text Widgets

Add text widgets to create titles and texts that stand out, and when you want to add more descriptive explanations to the dashboard and surrounding visualizations.

In the widget's settings, you can find plenty of options for creating different text styles. Some examples of what you can do include changing the font color, selecting a background color, adding hyperlinks, and defining the text alignment.

## To add a text widget:

1. Click on the Text Widget icon in the top menu.  
A new Text Widget appears in your dashboard.



2. Type in your text. As you begin to type, a tool bar with text and formatting options is displayed.

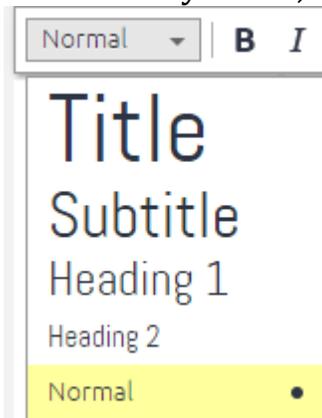


3. Click outside of the widget area to save the text and formatting.

## Examples

### To select a font style:

1. Click anywhere in the relevant paragraph.
2. From the Styles list, select one of the five predefined styles.

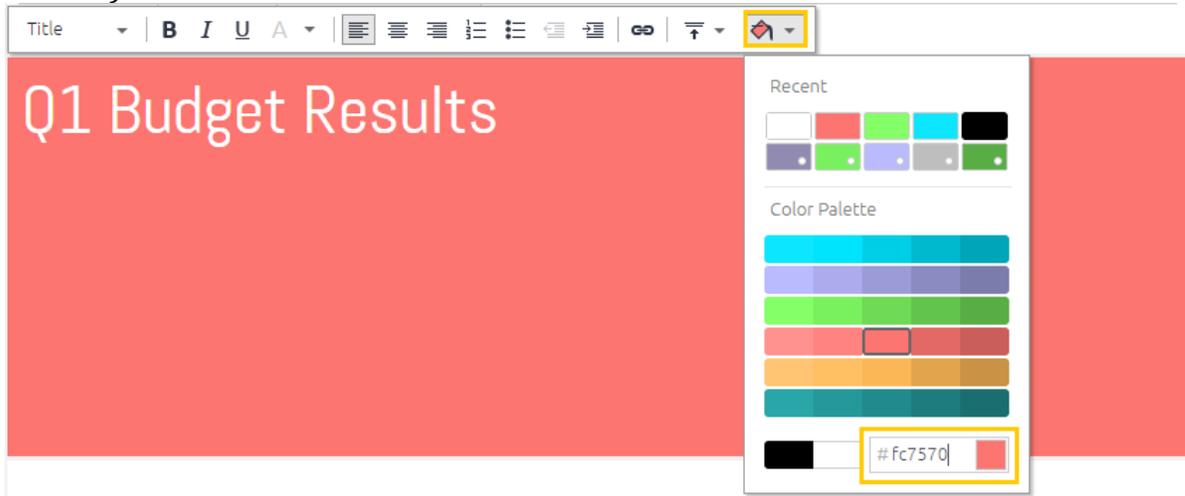


### To add a hyperlink:

1. Select the word or text fragment to which you want to add a link.
2. Click on the hyperlink icon, and enter the hyperlink.
3. You can also select which text to display instead of the selected (linked) text, and whether to open the link in the same window, or in a new window.

**To change the background color:**

1. Click on the paint bucket icon to open the color palette.
2. Select a color from the palette, or enter a custom color (Hexadecimal format).



# Adding Widgets to a Dashboard

This topic describes how to add widgets to your dashboard through the Widget Wizard.

The Widget Wizard automatically guides you through a very simple process of creating a widget while offering the best recommendations for displaying the data that you select.

Almost always, simply accepting the wizard's recommendations will help you create a great widget. However, if you want more control over your widget's appearance and behavior, click the **Advanced Configuration** option in the bottom left of the wizard to access a rich variety of additional design options, as described in [Using the Widget Designer](#).

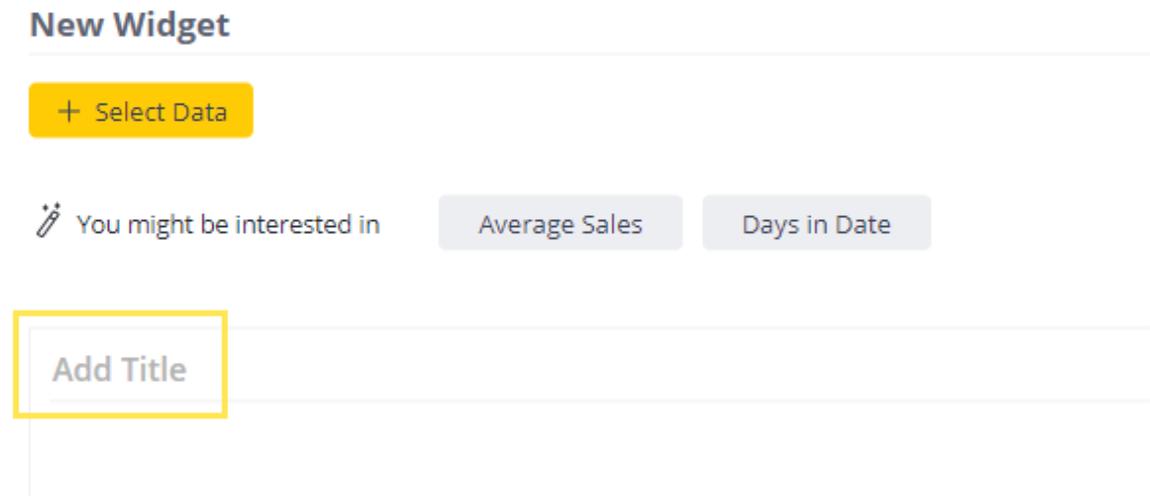
## Using the Widget Wizard

**To add a widget to the dashboard:**

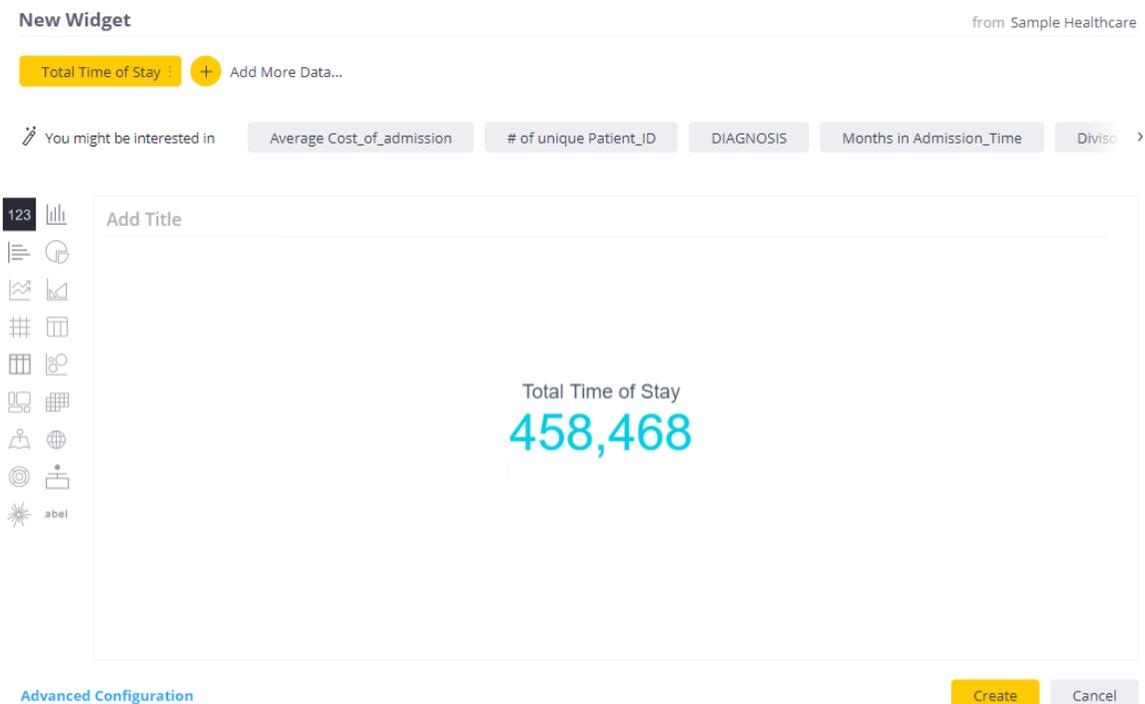
1. If this is your first widget, click . The Sisense Data Browser opens, and from there you can select one of the fields (columns) from this dashboard's ElastiCube. If you already have widgets in your dashboard, then click the  button.

**Note:** To select data from a different ElastiCube, click **from** on the right side of the **New Widget** window, and select an ElastiCube.

2. You can add a title to your widget now or after creating your widget.



3. In the Data Browser, select a field to add to your widget. The field that you select is automatically displayed in a suggested widget. When only a single numeric field is selected, it is displayed as an Indicator widget. A single descriptive field is displayed as a Pivot widget.



## Single Numeric Field

## New Widget

from Commerce Demo

Brand + Add More Data...

Brand
ABC
Adbananor WorldWide
Adcadopazz WorldWide Inc
Addimantor WorldWide
Addimegistor Inc
Addimin
Adericator
Aderor WorldWide
Adhupar
Adhupommazz Corp.
Adkilazz Holdings
Adkilewor WorldWide Corp.
Adnipax Direct

Advanced Configuration

Create

Cancel

## Single Descriptive Field

But your widget is not really interesting yet, because it has only one field!

## New Widget

from Commerce Demo

Brand Total Revenue + Add More Data...

Brand	Total Revenue
ABC	251,866.95
Adbananor WorldWide	4,340.14
Adcadopazz WorldWide Inc	300.17
Addimantor WorldWide	6,726.56
Addimegistor Inc	200
Addimin	56.08
Adericator	99.69
Aderor WorldWide	3,166.59
Adhupar	10.96
Adhupommazz Corp.	13,362.71
Adkilazz Holdings	5,603.10
Adkilewor WorldWide Corp.	2,822.54
Adnipax Direct	2,600.84

Advanced Configuration

Create

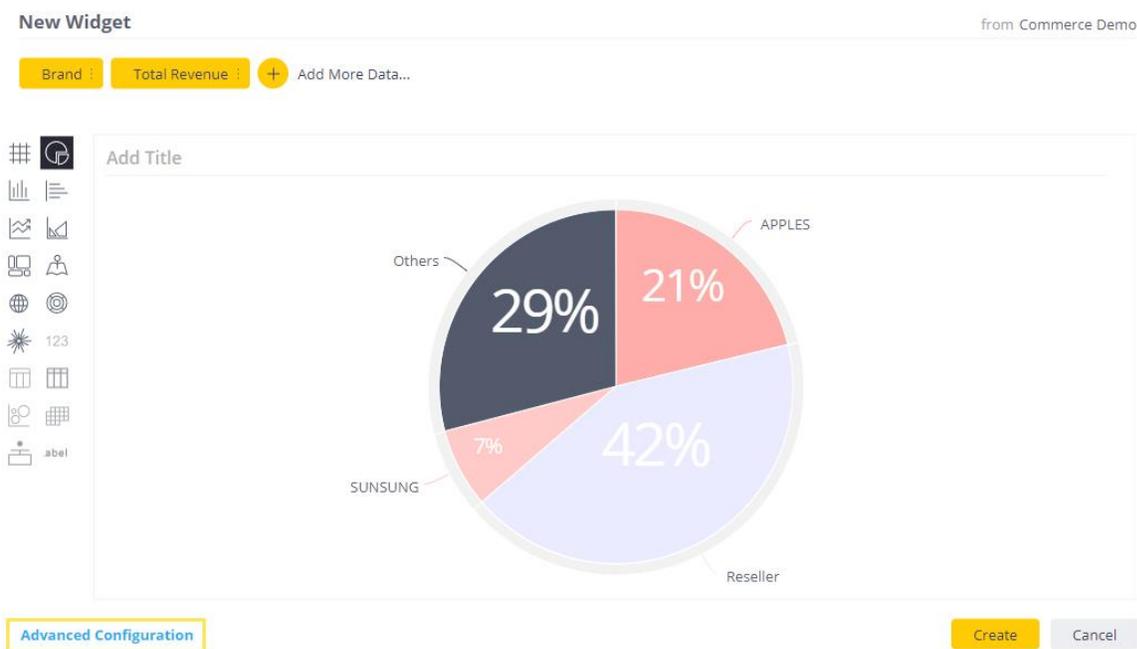
Cancel

Repeat the step above to add more fields to the widget.

As you add fields, the relevant options are displayed as buttons to the left of your selection.

The fields that you select are listed across the top left of the wizard. You can click on each visualization button to display the selected fields in that visualization/chart.

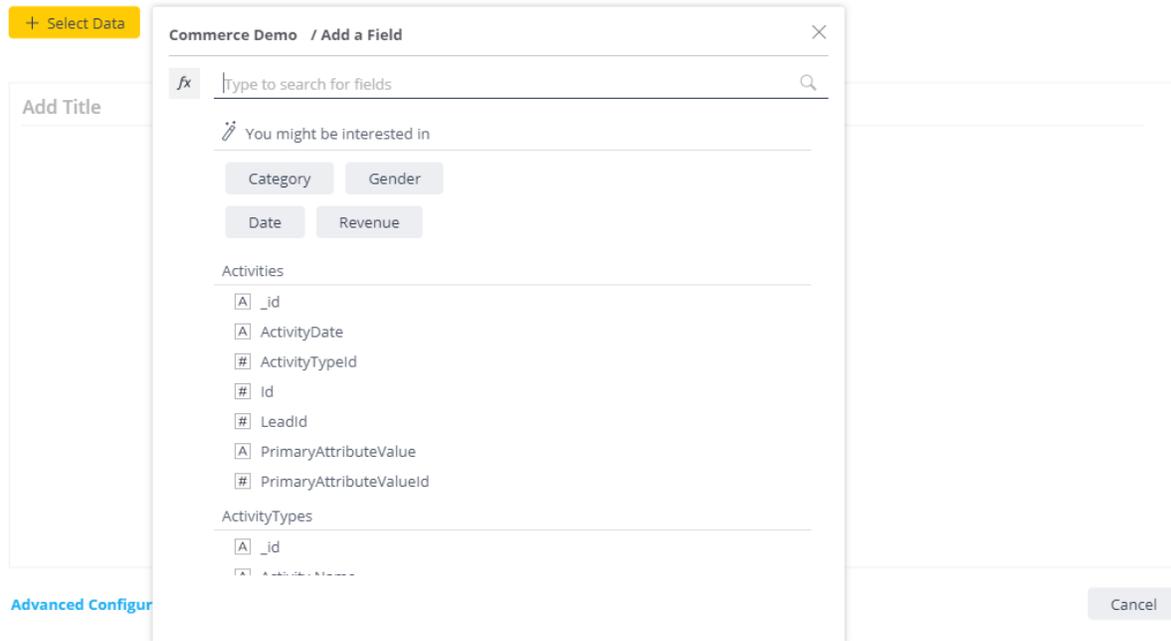
For example, the same Widget as above could be displayed as a Pie Chart:



**Tip:** Sisense provides its best recommendation for your widget design. However, if you want more control, you can click the Advanced Configuration option on the bottom left of the window to provide a rich variety of additional design options, as described in Using the Widget Designer.

## Data Browser

In the Data Browser, you can select and add columns (sometimes called fields) from an ElastiCube to your widget.



An icon appears to the left of each field to indicate its data type:



Date

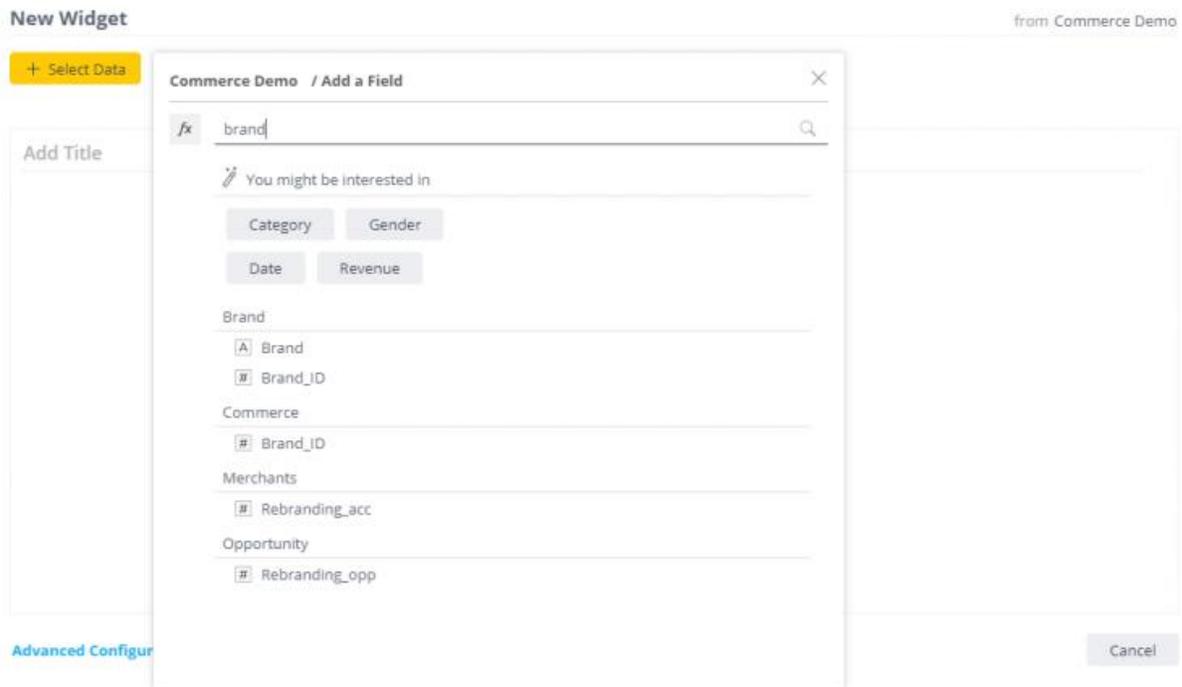


Alphanumeric string



Numeric

As you type into the Data Browser, the list is dynamically filtered to only show the fields that contain the text you typed.



Simply clicking a field to add it to the widget generally provides great results, however there are more options in the Data Browser.

- Hover over a field in the Data Browser and click **More ...** to display additional aggregation (quick functions) and filtering options.
- Click the  button to define formulas (free-form expressions) that define the field values and filters of a widget. A rich variety of functions are provided for you to use in the formula that you define.

Commerce Demo / Add a Field

fx brand

You might be interested in

- Category
- Gender
- Date
- Revenue

Brand

- Brand
- Brand\_ID

Commerce

- Brand\_ID

Merchants

- Rebranding\_acc Sum More...

Opportunity

- Rebranding\_opp

Advanced Configur

Cancel

To learn more about these options, see [Using Formulas](#).

# Area Chart

An Area Chart is very similar to a Line Chart except that the areas under each line are filled in (colored), and it is possible to display them as stacked. The chart is recommended for displaying absolute or relative (stacked) values over a time period.

## Adding Data

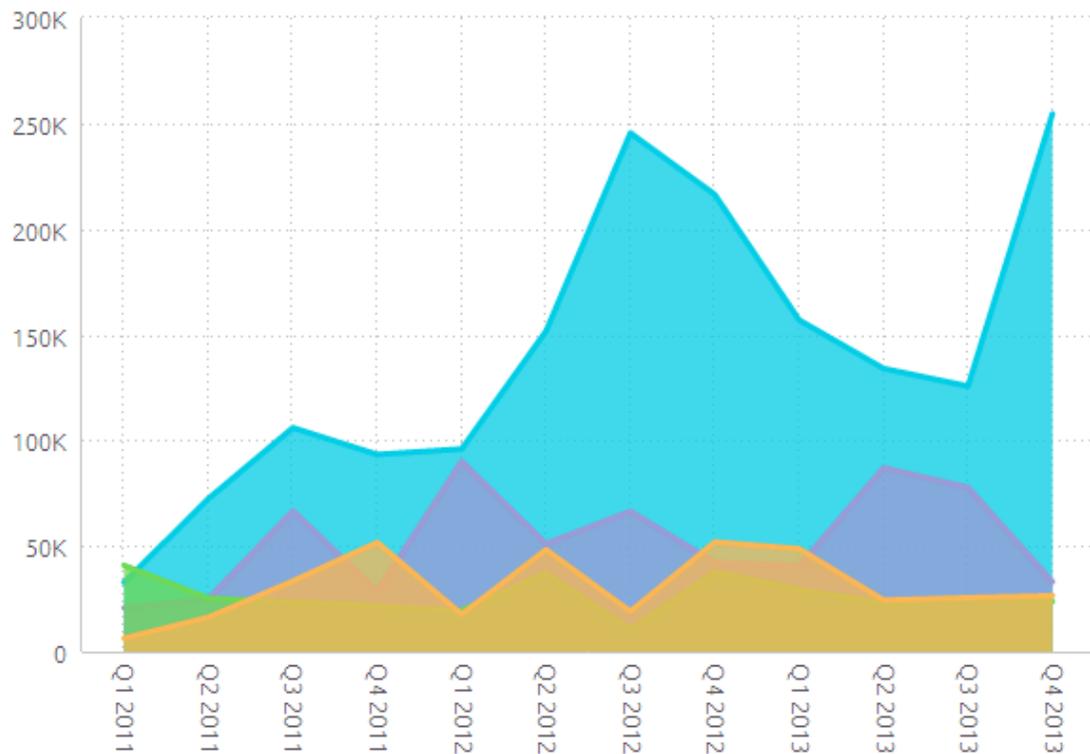
1. Click **Add +** to select the field(s) whose values will be placed on the X-Axis. Typically, the X-Axis of an Area Chart is used to represent time. You must add at least one item to the X-Axis. You can add at most two items. When you add two X-axis items, the chart's two X-Axis are combined.
2. In the **Values** area, click **Add+** to select the field whose value determines the height of the area in the chart. Typically, the Y-Axis of an Areas Chart is used to represent numeric data.  
You must add at least one field to Values. When you add more than one Value, then each item is represented in the chart by its own color and area.
3. **Break by** (optional): Click **Add +** to select a field by which to break up (group) the data represented in this chart.  
Each field added to Break By is represented by its own area in the chart and is automatically added to the chart's legend.  
This is an optional field, and operates under the following conditions:
  - You can add at most one **Break by** field.
  - The **Break by** option is available only when a single field was added to Values (Y-Axis).
  - You can add only a single field to Values, if the **Break by** option was used.

## Designing the Area Chart

Fine-tune the appearance of the Area Chart widget, using the following tools:

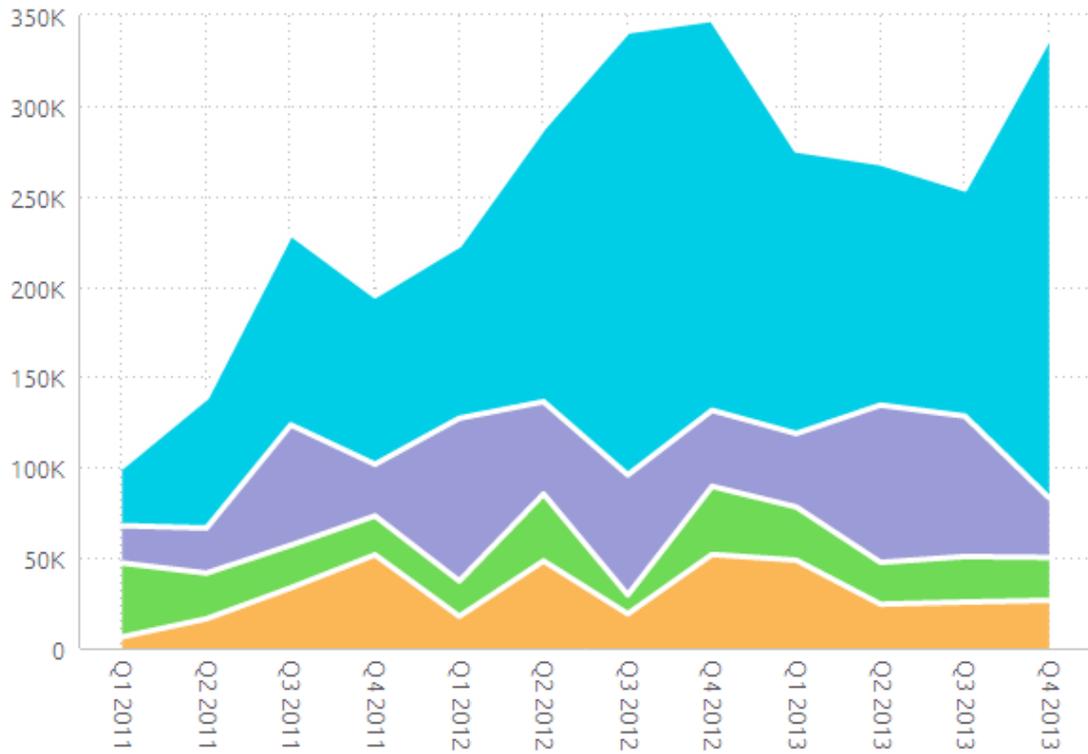
- **Area Types:** Select how areas are represented in the widget:

- **Classic:** Areas overlap, meaning that smaller areas appear on top of larger areas and cover them. The colors indicate the distribution between the **Values**.



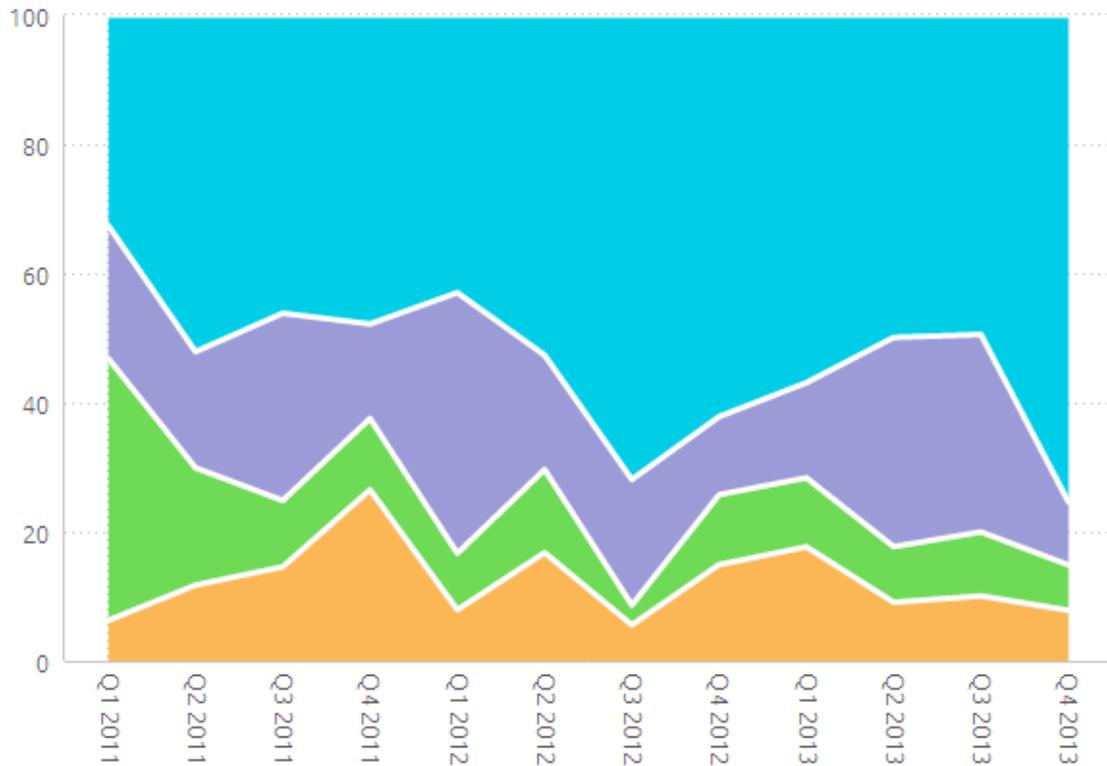
- **Stacked:** Areas are stacked on top of each other and do not overlap. Hover over the columns to see the percentage distribution among the

## Values.

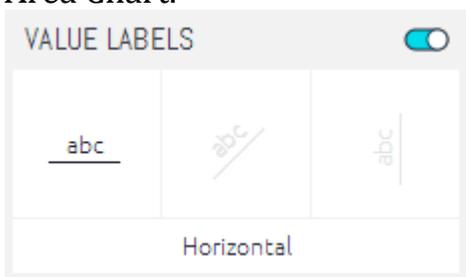


- **Stacked 100:** Areas are stacked on top of each other (but do not overlap). The combined area is stretched to the top in order to represent 100%. This option is most commonly used when the relative

distribution of the values is more important than their aggregation.



- **Line Type:** Select how the line appears in the widget: **Straight** (Default) or **Smooth**.
- **Line Width:** Select either **Thin**, **Bold** (Default) or **Thick**.
- **Legend:** Specify whether to show or hide the Legend and its position.
- **Value Labels:** Specify whether to show or hide labels showing values in the Area Chart.

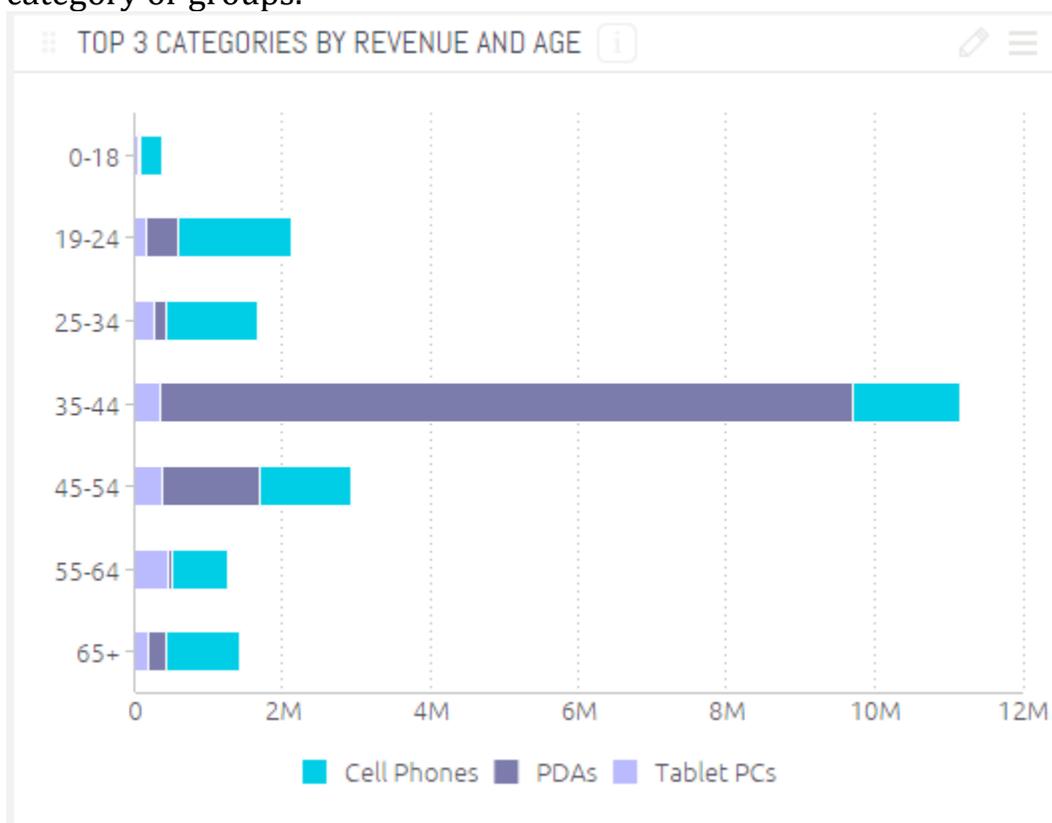


- **Markers:** Specify whether to show or hide markers over data points.
- **X-Axis:**
  - **Grid Lines:** Shows (Default) / hides.
  - **Labels:** Shows (Default) / hides.
  - **Title:** Select the checkbox to display the x-axis title. To edit the title, type in a new title. Click outside the text box to apply the new value.
- **Y-Axis**

- **Grid Lines:** Shows (Default) / hides.
- **Logarithmic:** Displays using orders of magnitude.
- **Labels:** Shows (Default) / hides.
- **Title:** Select the checkbox to display the y-axis title. To edit the title, type in a new title. Click outside the text box to apply the new value.
- **Values on Axis:** To change the default minimum, maximum or interval values on the axis, type in the new values, and click outside the text boxes to apply the new values. To restore any of the values to their default states, click on the reset button. ↺

# Bar Chart

The Bar Chart is commonly used to compare many items. The Bar Chart typically presents categories or items (descriptive data) displayed along the Y-Axis, with their values displayed on the X-Axis. You can also break up the values by another category or groups.



## Adding Data

1. In the **Categories** panel, click **Add+** to select the field(s) whose values will be placed on the Y-Axis.  
Typically, the Y-Axis of a Bar chart is descriptive data. You must add at least one item to **Categories**, and a maximum of two items. When you add two fields to Categories, then the chart's Y-Axis are combined.

2. In **VALUES**, select the fields whose values determine the length of each column. Typically, the values of a Bar Chart are used to represent numeric data.

You must add at least one field to **Values**. When you add more than one value, each item is represented in the chart by its own color and area.

3. (Optional) Select a field by which to break (group) the data represented in this chart by clicking **Add +**.

Each group is represented by a different column and is automatically added to the chart's legend.

Break by is optional, and operates under the following conditions:

- You can add at most one **Break by** field.
- The **Break by** option is available only when a single field was added to **Values** (Y-AXIS).
- You can only add a single field to **Values** if the **Break by** option was used.

## Designing the Bar Chart

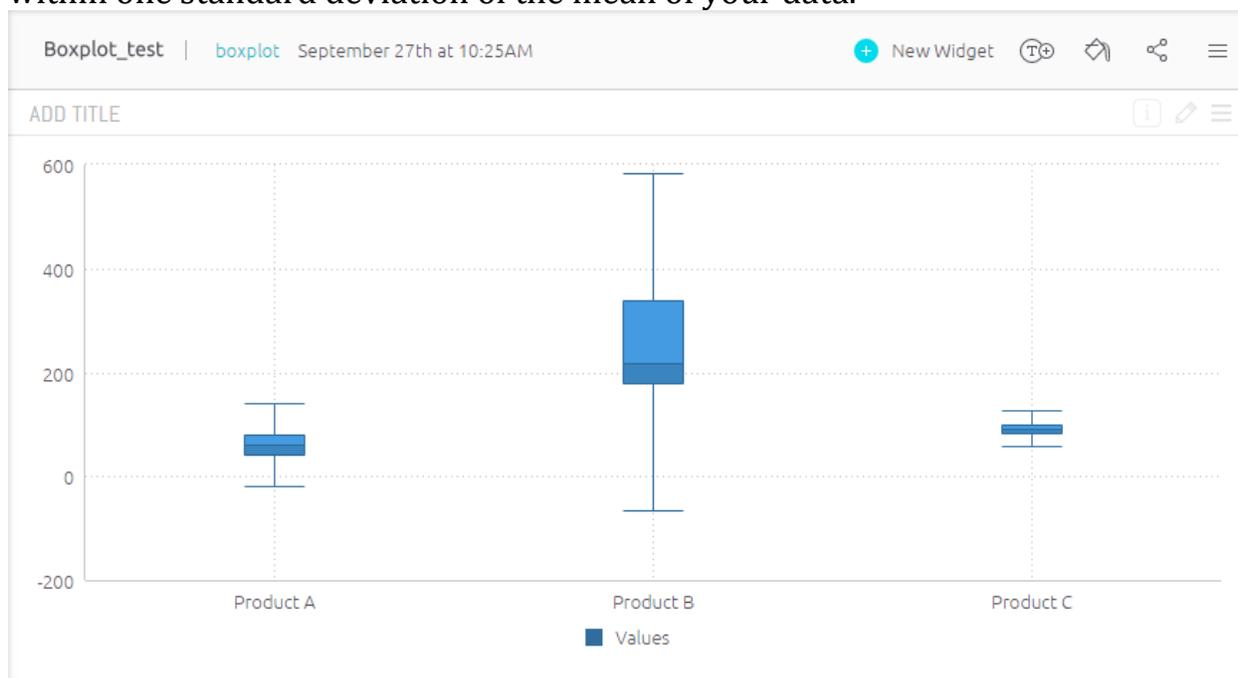
Fine-tune the appearance of the Bar Chart widget. The design options are the same as for [Column Chart](#).

# Box & Whisker Plot

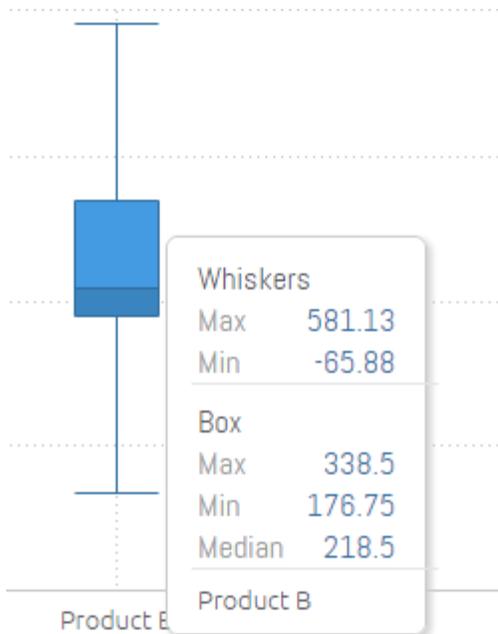
The box & whisker plot, or box plot, widget is a convenient way of visually describing the distribution, variability, and center of a data set along an axis.

Box plots are divided into four quartiles. The middle quartiles are represented by a box that contains 50% of the data and the median value. The upper and lower quartiles contain the maximum and minimum values and the remaining 50% of the data. These quartiles are represented by lines called whiskers.

The maximum and minimum values can be adjusted when defining your box plot. By default, these values are within 1.5 times of the IQR (Interquartile Range), however you can set them to the actual maximum and minimum values or to within one standard deviation of the mean of your data.



Box plots divide five descriptive statistics into four equal quartiles. You can view these statistics by hovering over the box plot.



## Installing the Box Plot Widget

1. [Download](#) the attachment and unzip the contents into your C:\Program Files\Sisense\PrismWeb\plugins\ folder. After those files have been unzipped there, you may also have to restart the web server.
2. Create a new widget. The Box & Whisker Plot chart is displayed in the list of options.



3. Select the Box & Whisker Plot chart from the chart types menu.
4. Click **Create**. The widget is added to the dashboard.

## Adding Data

1. In the **X-Axis** panel, click **Add +** to select the field whose values will be placed on the X-AXIS. The X-AXIS of a box plot is used to represent numeric data.
2. In the **Values** panel, click **Add +** to select

# Design the Box Plot

Fine tune the appearance of the box plot widget with the following tools.

- **Whisker Values:** Defines how the whiskers are calculated and displayed.
  - **IQR (Interquartile Range):** Default setting. The whisker values are calculated within 1.5 times the IQR (the difference between the 1st and 3rd quartiles).
  - **Extremums:** The whiskers values include the maximum and minimum values in the calculation.
  - **Standard Deviation:** The whisker values are calculated to one standard deviation above and below the mean of the data.
- **Box:** Specify whether to display a full or hollow box.
- **Value Labels:** Specify whether to show or hide labels indicating the values of each point (circle). Select the orientation of these labels: **Horizontal** (Default), **Diagonal** or **Vertical**.
- **Legend:** Specify whether to show or hide the Legend, and select its position.
- **Show Outliers:** Specify whether to show or hide labels indicating the values of each point (circle). Select the orientation of these labels: **Horizontal** (Default), **Diagonal** or **Vertical**.
- **X-Axis and Y-Axis:**
  - **Grid Lines:** Shows (Default) / hides.
  - **Labels:** Shows (Default) / hides.
  - **Title:** Select the checkbox to display the axis title. To edit the title, type in a new title. Click outside the text box to apply the new value. To change the default minimum, maximum or interval values on the axis, type in the new values, and click outside the text boxes to apply the new values.
  - **Values on Axis:** To restore any of the values to their default states, click on the reset button.
- **Auto Zoom:** When a widget contains more data than can comfortably be displayed in one view, Auto Zoom will resize the chart to include more data, and in some cases adds a zoom bar under the widget. The zoom bar enables you to scroll right and left and to zoom in/out of different parts of the width of the widget.

# Calendar Heatmap

The Calendar Heatmap widget visualizes values over days in a calendar-like view, making it easy to identify daily patterns or anomalies.

The screenshot displays the Sisense Calendar Heatmap widget. The main visualization is a calendar for the year 2013, showing months from January to December. Each day is represented by a colored square, where the color indicates the value of the data point. The color scale ranges from red (low values) to green (high values). The interface includes a 'DATE' panel on the left with 'Days in Date' selected, and a 'COLOR' panel with '# of Visit ID' selected. The right-hand sidebar contains 'APPLY' and 'CANCEL' buttons, 'FILTERS' and 'DESIGN' tabs, and various configuration options for the calendar view, including 'MONTHS', 'CALENDAR TYPE' (Classic, Week View), 'WEEK BEGINS ON' (Sunday, Monday), 'GRAYOUT WEEKENDS', 'SHOW DAY NUMBER', and 'SHOW DAY NAME'.

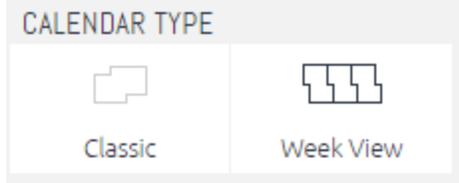
## Adding Data

1. In the **Date** panel, select the date field that you want to be used as the calendar.  
The calendar's range will be between the minimum and maximum date values within the field.
2. In the **Color** panel, click **Add +** to select the field that determines which days to color.

## Designing the Heatmap

Fine-tune the appearance of the Calendar Heatmap widget, using the following tools.

- **Months:** Select how many months should be presented at once: 1 (One Month View), 3 (Quarter View), 6 (Half-Year View) or 12 (Year View). When more months exist than currently visible, paging arrows are displayed to move between months.
- **Calendar Type:** Select between two types of view modes:
  - Classic – Each month is presented in typical calendar format.
  - Week View – Months are ordered by weekdays from top to bottom.



- **Week Begins On:** Select the first day of the week to be represented in the calendar.

## More Configuration Options

- Grey Out Weekends – If on, the weekend days will be colored grey in the calendar.
- Show Day Number – Toggle whether to show the calendar day numbers inside the widget.
- Show Day Name – Toggle whether to show the day name abbreviations on the calendar.

# Column Chart

The column chart can be used in different business scenarios, especially for comparing items, and comparing data over time. The chart can include multiple values on both the X and Y-axis, as well as a break down by categories displayed on the Y-axis.

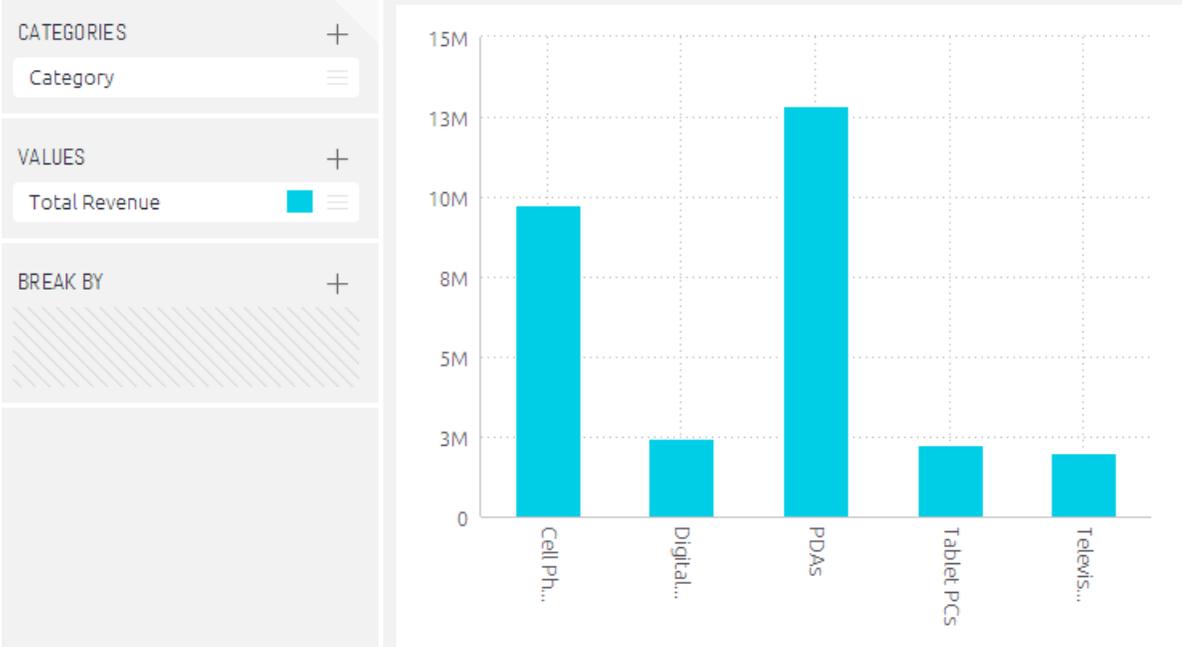
## Adding Data

1. In the **Categories** panel, click **Add +** to select the field(s) whose values will be placed on the X-Axis.

Typically, the X-Axis of a Column Chart is descriptive data. You must add at least one item to **Categories**, and at most, two items.

When two **Category** items are added, then the chart's X-Axes are automatically grouped. To change the order of the categories, drag a category up or down in the Categories list.

### One Category Item



## Two Category Items



2. In the **Values** area, select the field whose value determines the height of the columns by clicking **Add +**. Typically, the Y-Axis of a Column Chart is used to represent numeric data.

You must add at least one field to **Values**. When more than one Value is added, then each item is represented in the chart by its own color and area.

3. **Break by**: Select a field by which to break (group) the data represented in this chart by clicking **Add +**.

Each group is represented by a different column and is automatically added to the chart's legend.

This is an optional field, and operates under the following conditions:

- At most one **Break by** field can be added.
- The **Break by** option is available only when a single field was added to **Values** (Y-Axis).
- Only a single field can be added to **Values**, if the **Break by** option, described below, was used.

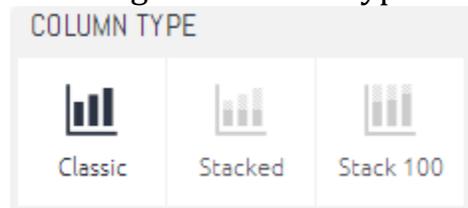
4. Define filters for the widget, as described in [Filtering Dashboards](#).

# Designing the Column Chart

Fine-tune the appearance of the COLUMN CHART Widget, using the following tools:

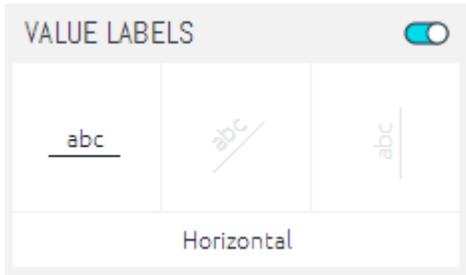
- **Column Type:** Select how columns are represented in the Widget:
  - **Classic:** Columns are displayed side by side.
  - **Stacked:** Columns are stacked on top of each other and do not overlap. The tooltips over the columns show the percentage distribution among the Values.
  - **Stacked 100:** Columns are stacked on top of each other (but do not overlap) and the combined column is stretched to represent 100%. This option is most commonly used when the relative distribution of the values is more important than their aggregation.

Selecting the Column Type



- **Legend:** Specify whether to show or hide the Legend and its position.

- **Value Labels:** Specify whether to show or hide labels showing values in the Column Chart. You can also select the angle of the labels.



- **X-Axis:** Enable or disable the following options:
  - **Grid Lines:** Shows (Default) / hides.
  - **Labels:** Shows (Default) / hides.
  - **Title:** Select the checkbox to display the X-Axis title. To edit the title, type in a new title. Click outside the text box to apply the new value.
- **Y-Axis:** Enable or disable the following options:
  - **Grid Lines:** Shows (Default) / hides.
  - **Logarithmic:** Displays using orders of magnitude.
  - **Labels:** Shows (Default) / hides.
  - **Title:** Select the checkbox to display the y-axis title. To edit the title, type in a new title. Click outside the text box to apply the new value.
  - **Values on Axis:** To change the default minimum, maximum or interval values on the axis, type in the new values, and click outside the text boxes to apply the new values. To restore any of the values to their default states, click on the reset button. reset button
- **Auto Zoom:** When a widget contains more data than can comfortably be displayed in one view, Auto Zoom will resize the chart to include more data, and in some cases adds a zoom bar under the widget. The zoom bar enables you to scroll right and left and to zoom in/out of different parts of the width of the widget.

# Indicator

The Indicator widget provides various options for displaying one or two numeric values as a number, gauge or ticker. It also provides the option to add additional titles and a color-coded indicator icon representing the value, such as a green up arrow or a red down arrow.

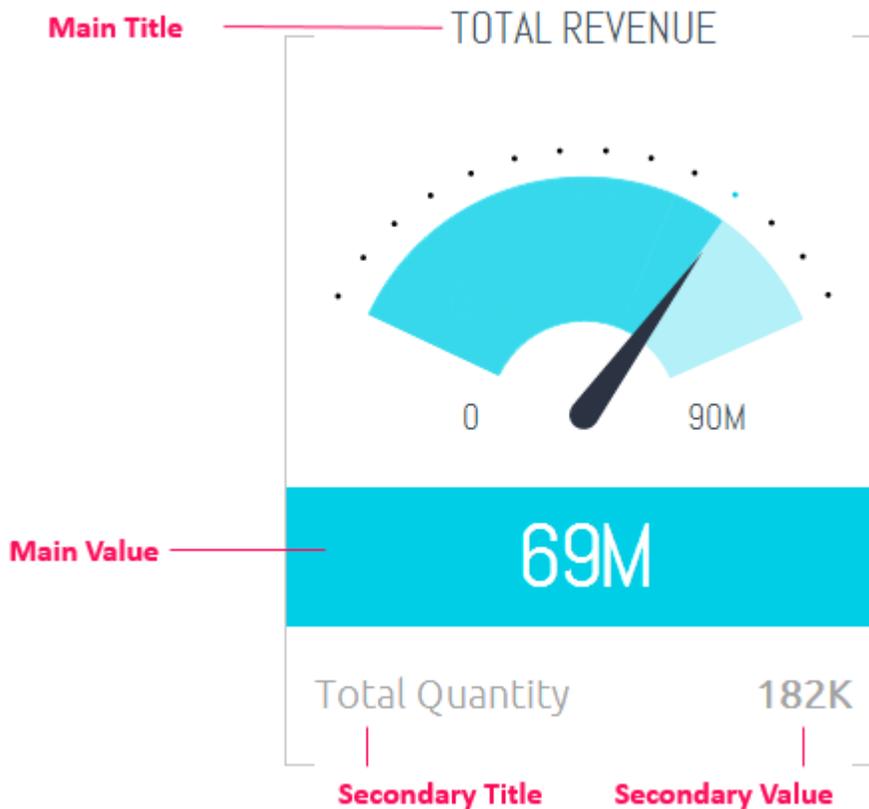
## Adding Data

1. Click **Add +** in the **Value** panel to select one field whose name and value will be shown as the **Main Title** and **Main Value**, as shown below.
2. Click **Add +** in the Secondary panel to select one field whose name and value will be shown as the Secondary Title and Secondary Value, as shown below.

### Numeric Indicator



## Gauge Indicator



3. For a gauge, select the minimum and maximum values that can be represented by the gauge. These values can either be set as a fixed numeric value that you specify or as another numeric field that you select using the Data Browser (thus making the value dynamic).

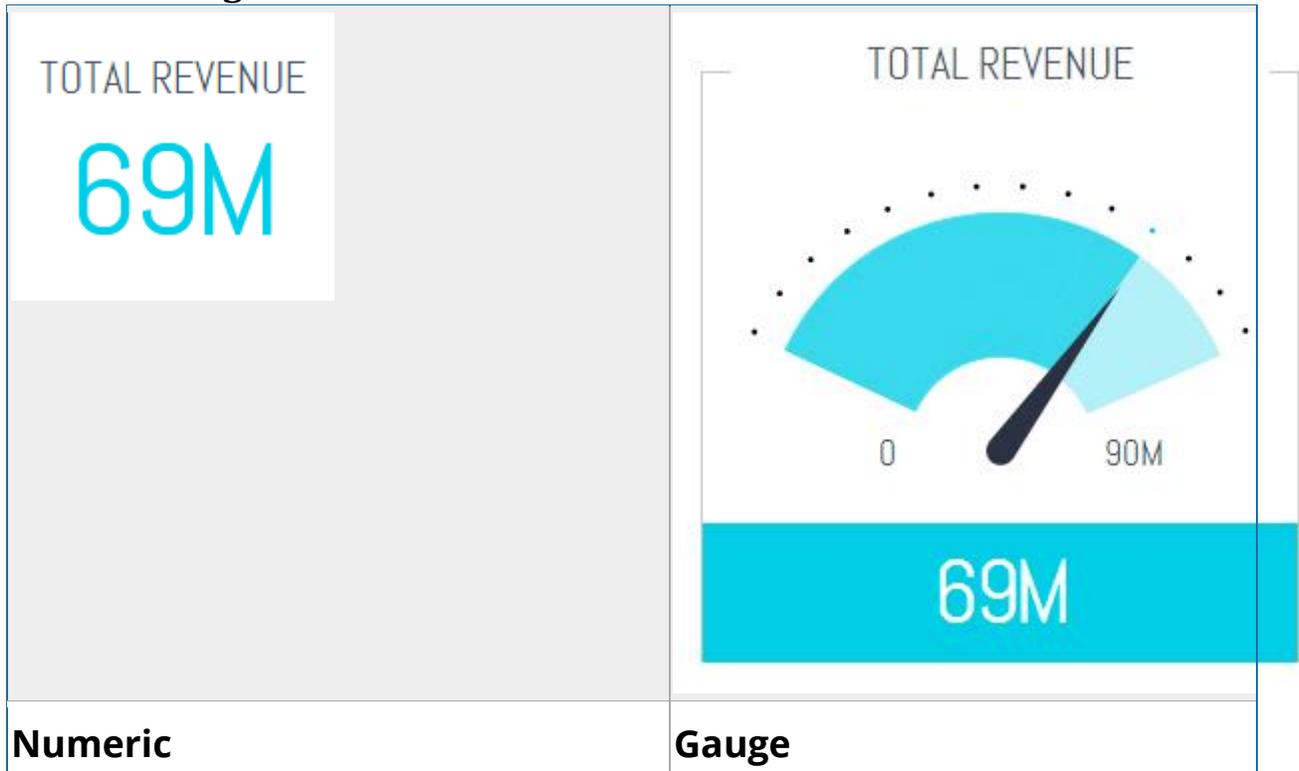
MIN	<input type="text" value="0 (default)"/>
MAX	<input type="text" value="150"/>

In an Indicator widget, you can also define whether an icon appears in the Widget and the icon's color. For example, you can define a red arrow pointing downwards for a negative value. You may refer to [Defining Conditional Coloring - Condition Tab](#) for more information.

# Designing the Indicator

Fine-tune the appearance of the Indicator widget, using the following tools:

- **Indicator Type:** Select whether the Indicator appears in **Numeric** form or as a **Gauge**.



**Note:** If you significantly shorten the height of an Indicator widget in the dashboard, it automatically turns into a Ticker type Widget. You may refer to [Ticker Widgets](#) for more information.

| GROWTH OF VISITORS: 13.2% | # of Visitors: 220K |

- **Skin:** Select the skin of the gauge.
- **Components:** Select which labels to include in the widget.
- **Indicator Sub Type**(for numeric indicator only): You can choose a another skin option.

# Line Chart

The line chart can be used for various business cases, including:

- Comparing data over time, for example: to analyze sales revenue for the past year.
- Comparing changes over the same period of time for more than one group or category. Example: Analyze expenditures of different business units for the past year.

## Adding Data

1. Add data to the X-axis.

Click **Add +** to select the field(s) that will be represented on the X-Axis.

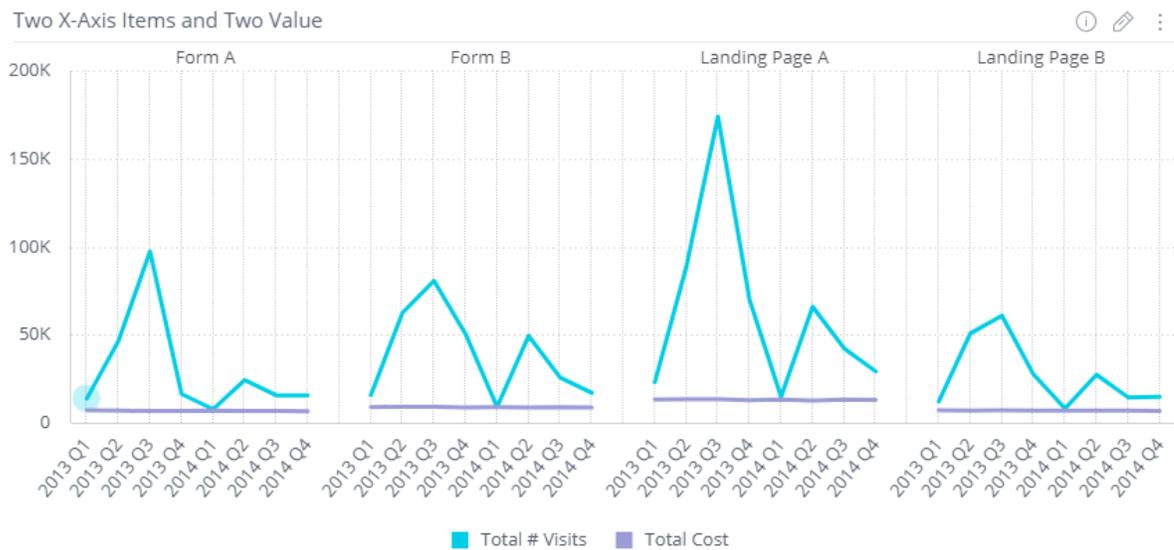
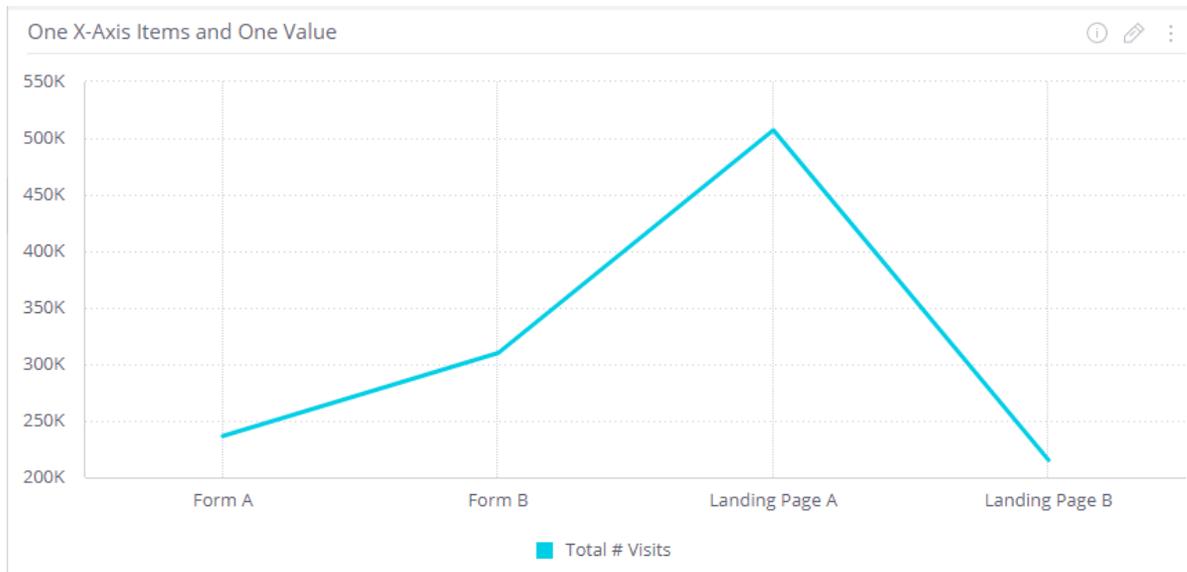
Typically, the x-axis of a Line chart is used to represent time.

You must add at least one item to the x-axis, and at most, two items.

When you add two x-axis items, the chart's x-axes are combined and a separate plot (line) is displayed for each unique value of the second selected x-axis.

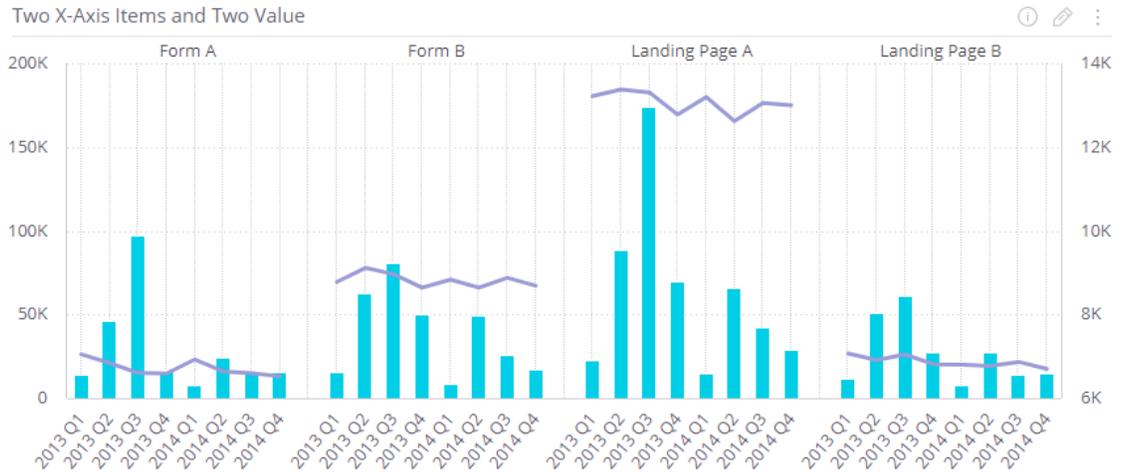
2. Add data to the Y-axis.

In the **Values** area, click **Add +** and select the field whose value determines the height of the lines in the chart. Typically, the Y-Axis of a line chart is used to represent numeric data. You must add at least one field. If you add more than one field, then each item is represented in the chart by its own line and color.

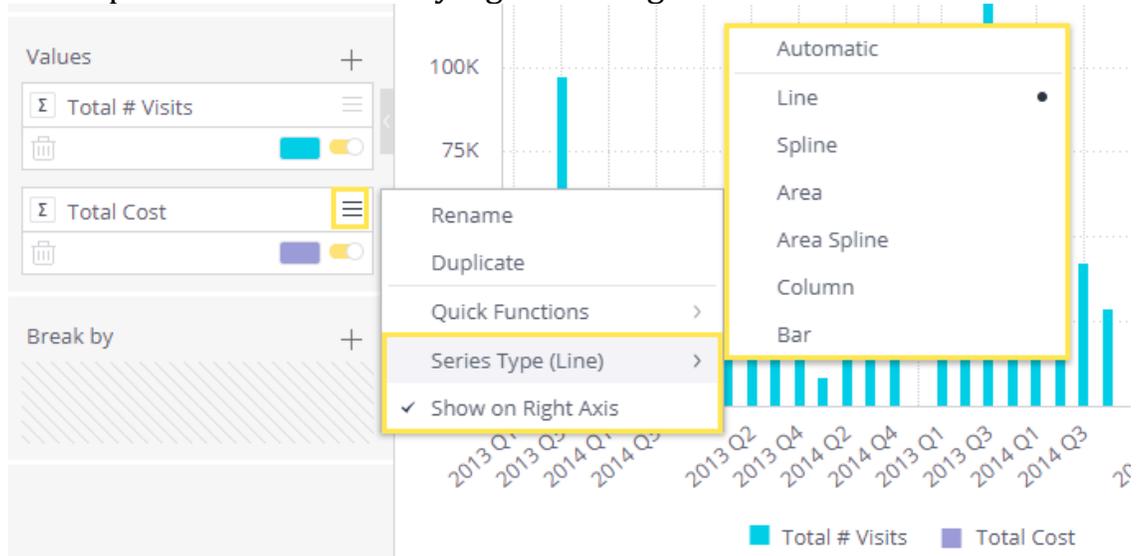


Another version of the same chart includes the following differences:

- The second value has been changed into a column chart.
- The second value has been added to the right axis. This is a useful option when the values are using a considerably different scale.



Both options are available by right-clicking on the value.



- (Optional) Select a field by which to break up the data. Click **Add +**. Each item that is a member of the field added to **Break by** is represented by its own line in the chart, and is automatically added to the chart's legend. You can add only one **Break by** field. The **Break by** option is only available when you have a single field in the Y-Axis.

## Designing the Line Chart

Fine-tune the appearance of the line chart, using the following tools:

- Line Width:** Select either **Thin**, **Bold** (default) or **Thick**.

- **Line Type:** Select how the line appears in the widget: **Straight** (default) or **Smooth**.
- **Legend:** Toggle the legend on or off, and select its position in the chart.
- **Value Labels:** Toggle labels on or off, and define how labels are displayed: **Horizontal**, **Vertical** or **Diagonal**.
- **Markers:** Select whether to show or hide markers over data points. Select **Full** or **Hollow**, and **Small** or **Large** markers.
- **X-Axis:**
  - **Grid Lines:** Show (default) or hide grid lines.
  - **Labels:** Show (default) or hide labels.
  - **Title:** Select the checkbox to display the x-axis title. To edit the title, type in a new title. Click outside the text box to apply the new value.
  - **Values on Axis:** To change the default minimum, maximum or interval values on the axis, type in the new values, and click outside the text boxes to apply the new values. To restore any of the values to their default states, click on the reset button. reset button
- **Y-Axis**
  - **Grid Lines:** Show (default) or hide grid lines.
  - **Logarithmic:** Displays a logarithmic scale instead of a linear scale on the axis.
  - **Labels:** Show (default) or hide labels.
  - **Title:** Select the checkbox to display the y-axis title. To edit the title, type in a new title. Click outside the text box to apply the new value.
  - **Values on Axis:** To change the default minimum, maximum or interval values on the axis, type in the new values, and click outside the text boxes to apply the new values. To restore any of the values to

their default states, click on the reset button. ↻

Y-Axis ☑

<input checked="" type="checkbox"/> Grid Lines
<input type="checkbox"/> Logarithmic
<input checked="" type="checkbox"/> Labels
<input checked="" type="checkbox"/> Title <input type="text" value="Total # Visits"/>
Min <input type="text" value="Auto"/>
Max <input type="text" value="Auto"/>
Interv... <input type="text" value="Auto"/>

# Managing Widgets

This topic describes how you can edit, copy, and delete your dashboard's widgets.

## Editing a Widget

When you first create a new dashboard (as described in [Creating a New Dashboard](#) ) or a new widget (as described in [Adding Widgets to a Dashboard](#)), the Widget Wizard automatically guides you through Sisense Web's best recommendations for visualizing the data that you select.

Then, after a widget is created, you can fine-tune/edit it in the [Widget Designer](#). The Widget Designer provides a variety of options for changing the widget's appearance and behavior.

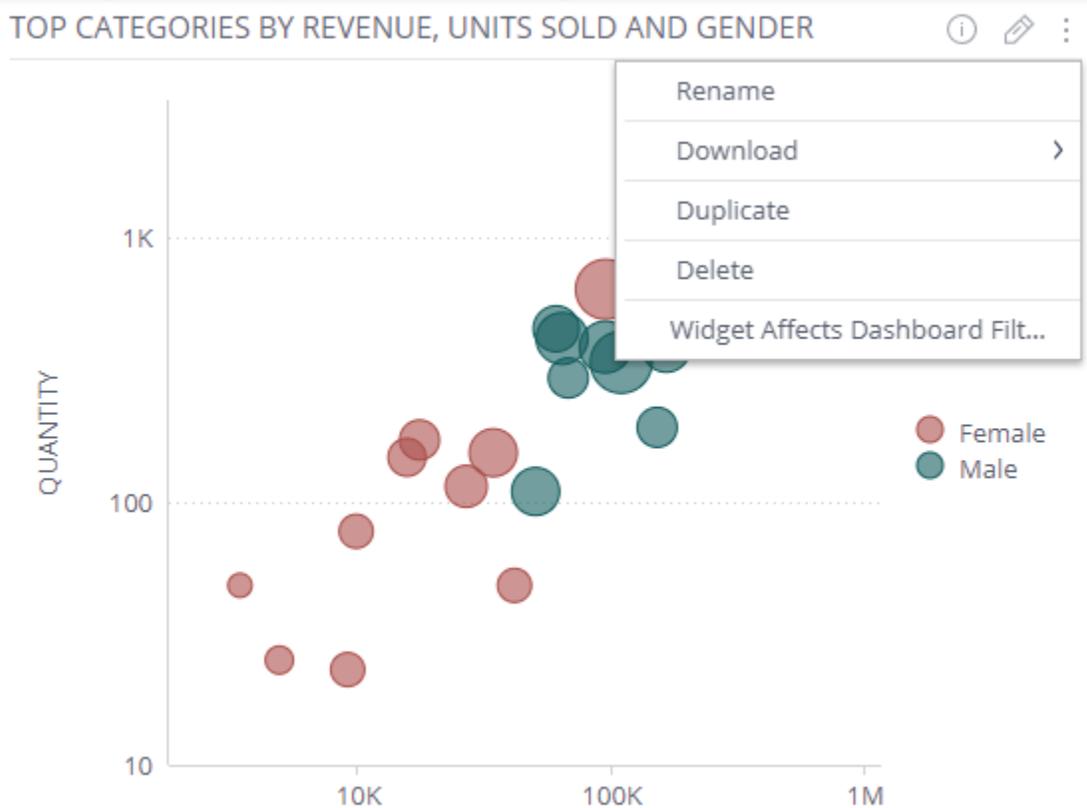
See [Fine-tuning a Widget](#) for more information.

## Copying a Widget

You can create a copy of a widget within the same dashboard or another dashboard in your list. This is useful for saving the original widget in the dashboard before making changes. The new widget is separate from the original so no changes to the new widget affect the original widget.

**To copy a Widget within the same dashboard:**

- In Dashboard view, click the widget menu  $\text{:}$  button in the top-right corner of the widget to be copied and select **Duplicate**.



- If you are editing the widget, click the widget menu  $\text{:}$  button in the top-right corner of the widget to be copied and select Save A Copy.

### To copy a Widget to another dashboard:

- Select the Header of the widget to be copied and drag it to the dashboard you want to copy it to.

**Note:** When you copy a widget to another dashboard whose data source is another ElastiCube, the Dashboard filters of the new dashboard will not affect the widget.

## Deleting a Widget

### To delete a Widget:

- Click the widget menu  $\text{:}$  in the top-right corner of the Widget and select **Delete**. Confirm the deletion.

# Pie Chart

The Pie chart is used to display proportional data, and/or percentages.

## Adding Data

1. In the Categories panel, click **Add +** to select the data that will be displayed in the Pie Chart. You can add only one field.

Typically, the categories of a Pie Chart contain descriptive data.

**Note:** All slices that are smaller or equal to 3% are grouped together into a slice called **Other**. You can hover over this slice to display its breakdown.

2. In the **Values** area, select the field whose values determine the size of the pie slice by clicking **Add +**. Values are typically used to represent numeric data.

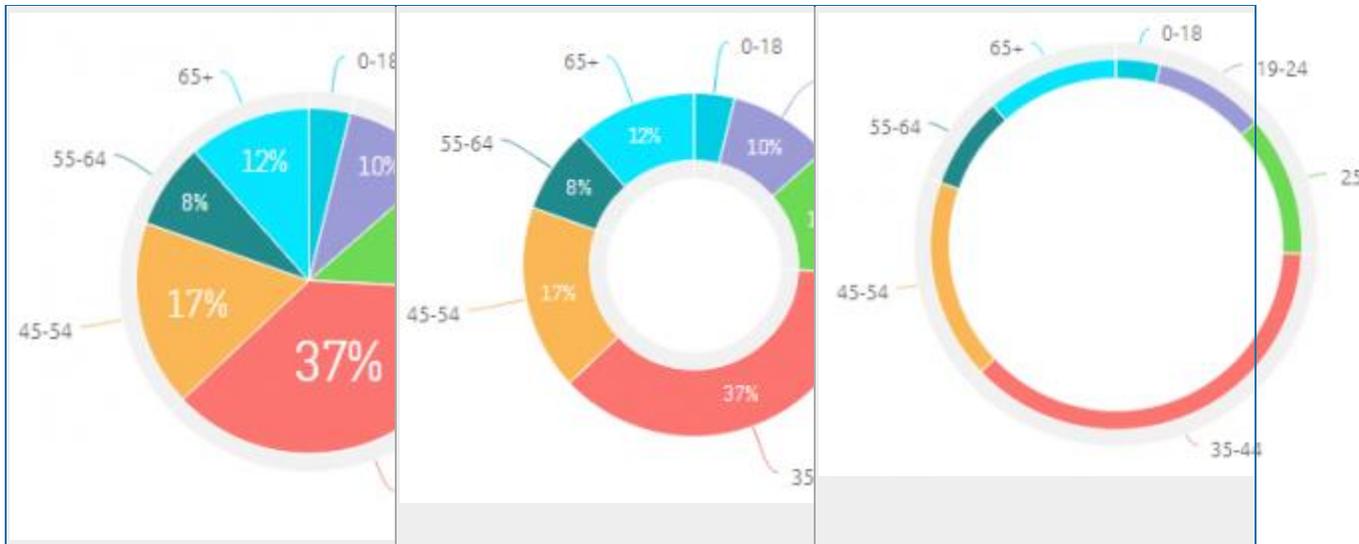
**Note:** You can add only one field.

## Designing the Pie Chart

Fine-tune the appearance of the PIE CHART widget, using the following tools:

- **PIE TYPE:** Select how the pie appears.





- **Legend:** Specify whether to show or hide the Legend, and select its position.
- **Labels:** Select which labels to display on the pie chart, as indicated below:

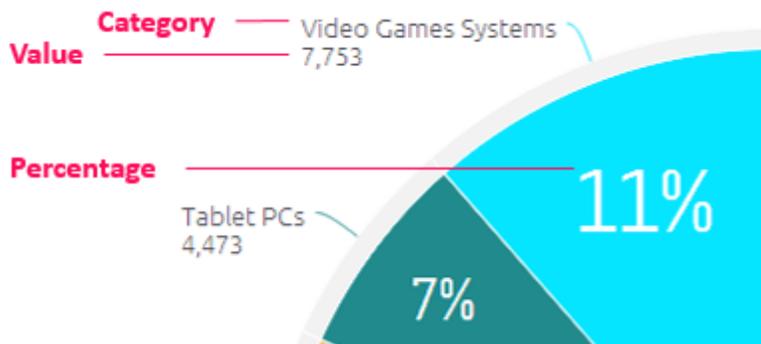
Labels

Categories

Value

Percent

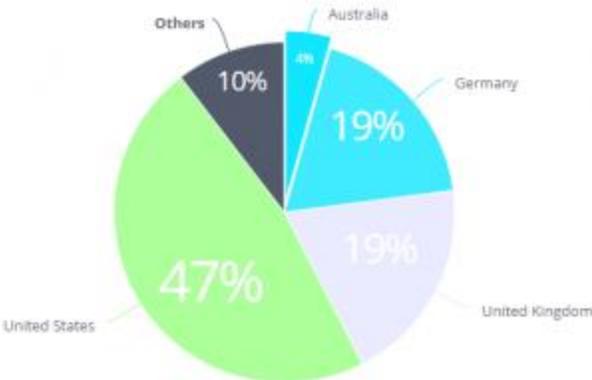
Decimals



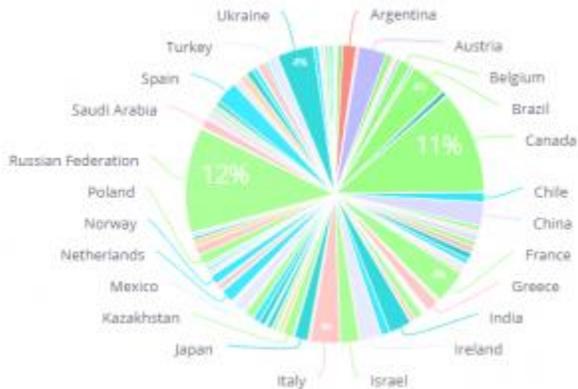
## Exploring Pie Charts

You can [drill down](#) into pie charts just like other widgets. In addition, when you pie chart has so many smaller slices that they are aggregated into a single slice called Other, you can click this slice to display a second pie chart containing a breakdown of the Other slice.

For example, the Pie chart below displays a breakdown of the top 5 slices in a pie chart with many smaller slices included in the Others slice.



By clicking the Others slice, you can see a breakdown of all the slices included in the Others slice.



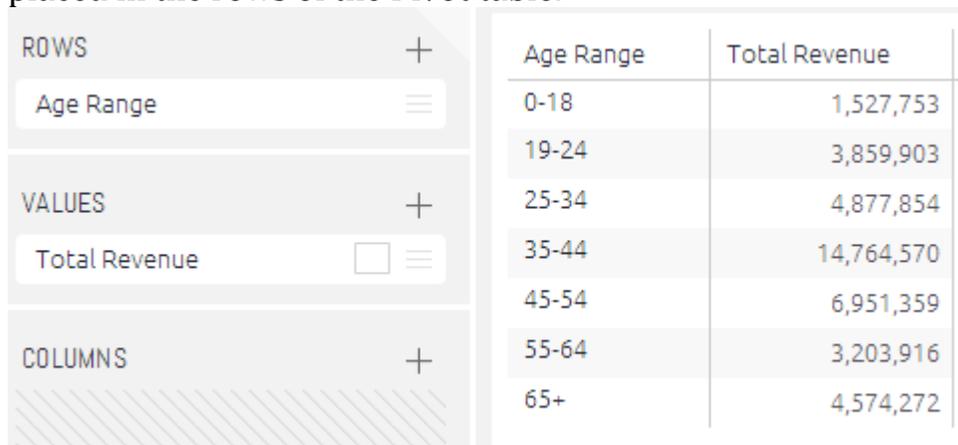
# Pivot

Pivot tables are one of the most useful widgets for visualizing data. They enable you to quickly summarize and analyze large amounts of data.

## Adding Data

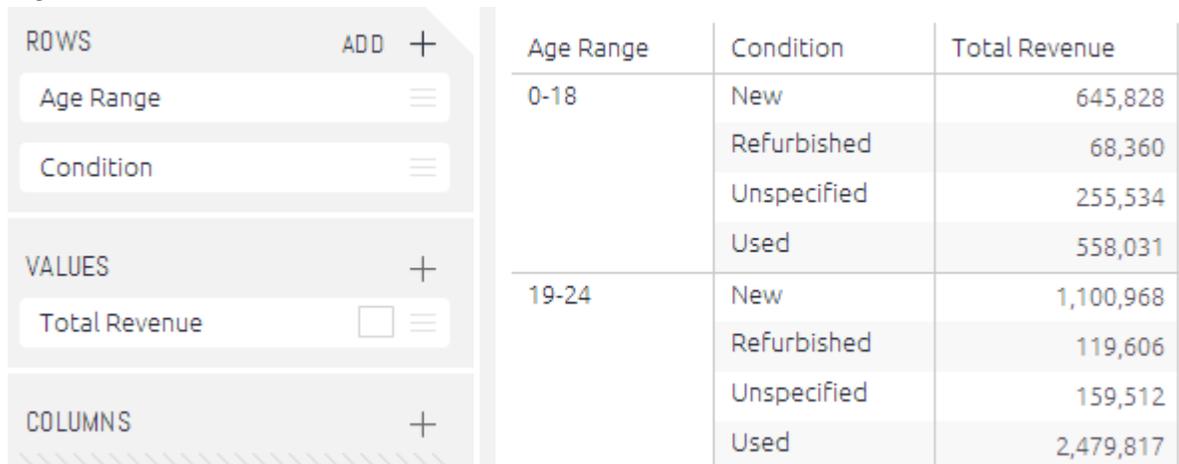
Select the data to appear in the Pivot.

1. In the **Rows** panel, click **Add +** to select the field(s) whose values will be placed in the rows of the Pivot table.



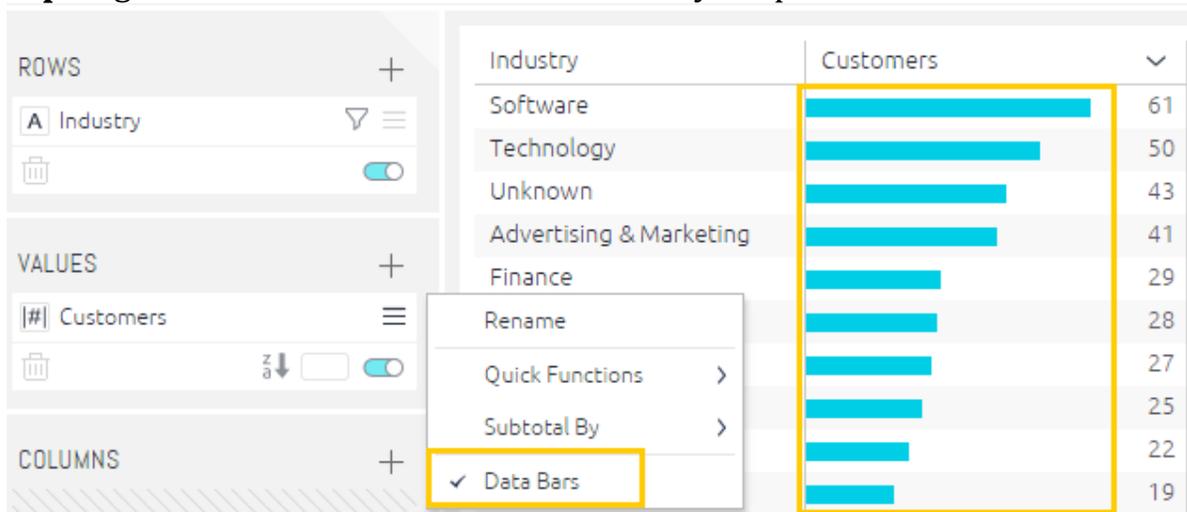
Age Range	Total Revenue
0-18	1,527,753
19-24	3,859,903
25-34	4,877,854
35-44	14,764,570
45-54	6,951,359
55-64	3,203,916
65+	4,574,272

When more than one Row is selected, the Rows are broken into sub-rows (groups). For example, the following shows **Condition** added as a second Row:



Age Range	Condition	Total Revenue
0-18	New	645,828
	Refurbished	68,360
	Unspecified	255,534
	Used	558,031
19-24	New	1,100,968
	Refurbished	119,606
	Unspecified	159,512
	Used	2,479,817

- In the **Columns** panel, click **Add +** to select the field(s) whose values will be placed in the columns of the Pivot table.  
You can drag and drop the fields that you added to **Columns** to the **Rows** area and vice versa.
- In the **Values** panel, select a field whose values will appear in the Rows and Columns of the Pivot table.  
**Tip:** Right-click the value to add data bars to your pivot table.

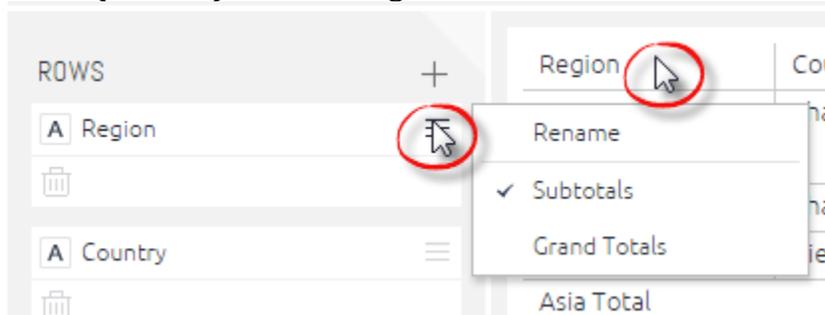


## Subtotals

You can add subtotals to one or more rows, and define how to calculate subtotals.

### To add subtotals to rows:

- Click on the menu of the Row in the Data Panel, or right-click on the row name (header) in the widget, and select Subtotals.



### To define how to calculate the subtotals:

- Click on the menu of the Value in the data panel, or right-click on the value name (header) in the widget, and then click **Subtotal by** and the method.

Region	Country	Sales	Total Sales	Total Revenue	[Total Revenue]/[Total Sales] Revenue per sale
Asia	Thailand	10	10	100	10
		20	20	200	10
	Thailand Total		30	300	10
	Vietnam	5	10	80	8
Asia Total			40	380	9.5
Europe	France	5	5	50	10
		10	10	50	5
	France Total		15	100	6.67
	Germany	5	10	50	5
Europe Total			25	150	6

The default option is Auto, which aggregates all the data. The other custom options (Sum, Min, Max, etc.) calculate only the values in the rows above the subtotal.

Look at the following example:

Region	Country	Total Sales	Total Revenue	Ave. Revenue per sale
Asia	Cambodia	20	300	15
	Thailand	30	300	10
	Vietnam	15	1,050	70
Asia Total		65	1,650	25.38

AUTO →

Region	Country	Total Sales	Total Revenue	Ave. Revenue per sale
Asia	Cambodia	20	300	15
	Thailand	30	300	10
	Vietnam	15	1,050	70
Asia Total		65	1,650	31.67

AVERAGE ↶

In the top example, using the **Auto** option, the subtotal for average revenue aggregates all the sales and revenue data, thus representing a true subtotal of the average revenue for the Asia region.

In the bottom example, **Average** is used to calculate the subtotal, and therefore calculates the average of all the average revenue values in the rows above the subtotal.

## Grand Totals

Grand totals aggregate the rows in the pivot table.

**To add grand totals:**

- Click on the menu of the row in the data panel, or right-click on the row name (header) in the widget, and then click **Grand total**.

Grand totals is effected by subtotals if defined. In the following example, the Grand total for Average Sales is 10, representing the average sales for two regions (Asia and Europe). In the example below, Average Sales has a subtotal (see above) set to Sum, and therefore the Grand total is the total of the rows above.

Region ^	Total Sales	[Average Sales]	Total Revenue
Asia	65	13	1,650
Europe	35	7	250
Grand Total	100	AUTO 10	1,900

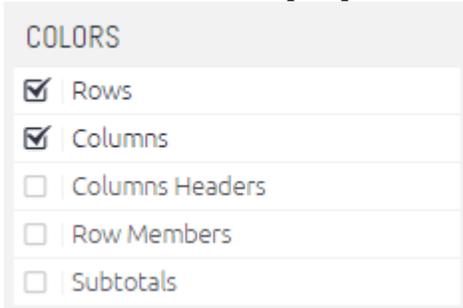
Region ^	Total Sales	[Average Sales]	Total Revenue
Asia	65	13	1,650
Europe	35	7	250
Grand Total	100	SUM 20	1,900

## Designing the Pivot Table

Fine-tune the appearance of the PIVOT table, using the following tools.

- **Page Size:** Specify how many rows appear in each page. Paging options are provided accordingly.

- **Colors:** Select the properties in the table to which you want to add color.



Select how the rows and columns of the Pivot table are highlighted.

## Exception Highlighting

Conditional formatting can be used for exception highlighting in a Pivot table. For example, as shown below:

Age Range	Condition	Total Quantity
0-18	New	1,209
	Refurbished	292
	Unspecified	360
	Used	1,913
19-24	New	2,544
	Refurbished	592
	Unspecified	569
	Used	3,790
25-34	New	5,950
	Refurbished	1,401
	Unspecified	1,407
	Used	8,831

See [Defining Conditional Coloring – Condition](#) for more information.

**Note:** For more styling options, including font style and colors, see [this article](#).

# Polar Chart

Use the polar (radar) chart to compare multiple categories/variables with a spacial perspective in a radial chart.

## Adding Data

1. Click **Add +** to select the field whose values will be represented as columns in the circle. Each Category gets its own data point.  
Typically, the Categories of a Polar Charts are descriptive data.
2. In the **Values** area, click **Add+** to select the field whose values determine the distance of the column from the center of the polar circle. This is typically used to represent numeric data.  
When you select more than one value, they are stacked on top of each other.
3. Select a field by which to break (group) the data in this chart by clicking **Add +**.  
Each field added to Break by is represented by its own column and color in the chart and is automatically added to the chart's legend.  
This is an optional field. You can add at most one **Break by** field. The **Break by** option is available only when a single field was added to **Values** (Y-AXIS).

## Designing the Polar Chart

Fine-tune the appearance of the Polar Chart widget, using the following tools.

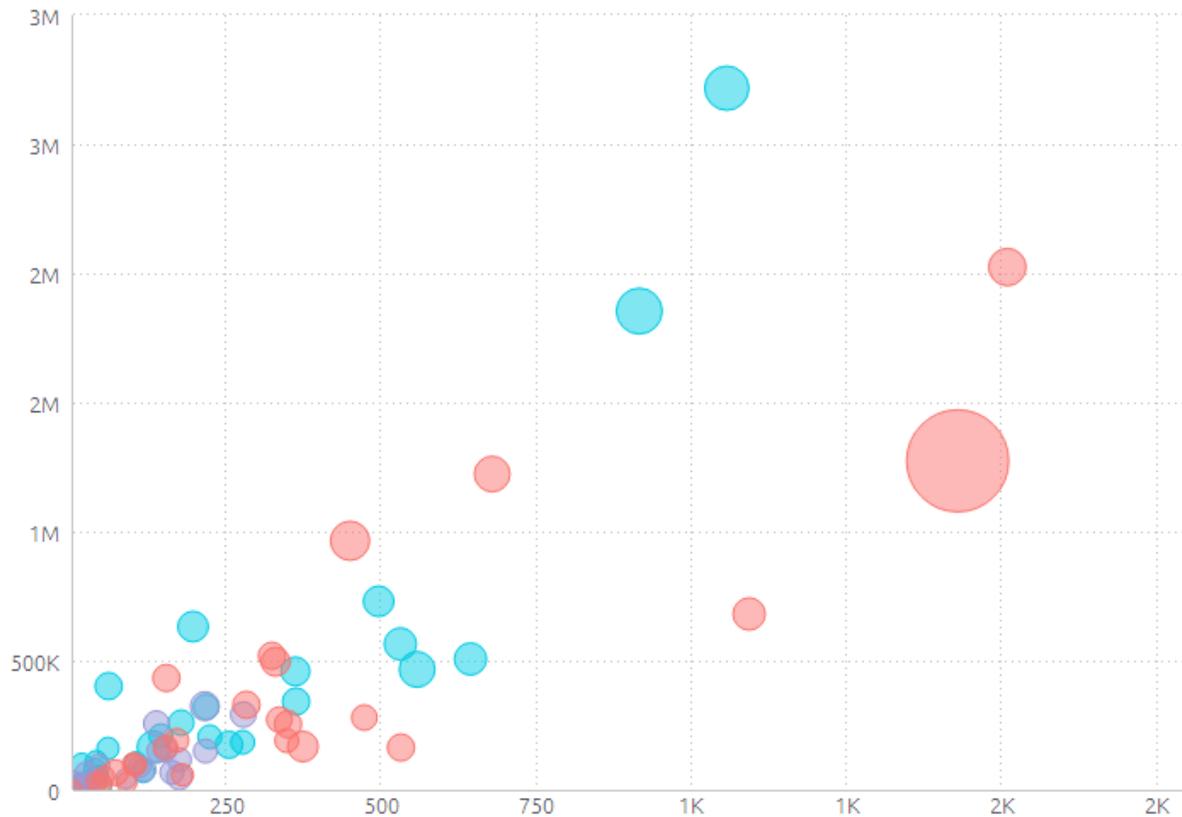
- **Polar Type:** Select the type of chart.
- **Legend:** Specify whether to show or hide the LEGEND, and select its position.
- **Value Labels:** Specify whether to show or hide labels indicating the values. Select the orientation of these labels: **Horizontal** (Default), **Diagonal** or **Vertical**.
- **Categories:**
  - **Grid Lines:** Shows (Default) / hides.

- **Labels:** Shows (Default) / hides.
- **Axis:**
  - **Grid Lines:** Shows (Default) / hides.
  - **Logarithmic:** Displays using orders of magnitude.
  - **Labels:** Shows (Default) / hides.
  - **Title:** Select the checkbox to display the axis title. To edit the title, type in a new title. Click outside the text box to apply the new value.
  - **Values on Axis:** To change the default minimum, maximum or interval values on the axis, type in the new values, and click outside the text boxes to apply the new values. To restore any of the values to their default states, click on the reset button. ↻

# Scatter Chart

The Scatter Chart displays the distribution of two variables on an X-Axis, Y-Axis, and two additional dimensions of data that are shown as colored circles scattered across the chart:

- **Point:** A field that for each of its members a scatter point is drawn.
- **Size:** An optional field represented by the size of the circles. If omitted, all scatter points are equal in size. If used, the circle size is relative to their value.



## Adding Data

1. In the **X-Axis** panel, click **Add +** to select the field whose values will be placed on the X-Axis.

Typically, the X-Axis of a Scatter Chart is used to represent numeric data. Alternatively, Descriptive data is also supported. You can only add one field.

2. In the **Y-Axis** panel, click **Add +** to select the field whose values will be placed on the Y-Axis.

Typically, the Y-Axis of a Scatter Chart is used to represent numeric data. Alternatively, descriptive data is also supported. You can only add one field.

3. In the **Point** panel, click **Add +** to select the point field.  
The point area determines the field to be scattered across the chart as circles; one point (circle) appears on the chart for each member of this field. These fields must be descriptive data (not numeric data).  
You can only add a point if either the X-Axis or the Y-Axis was defined to represent numeric data.

4. **Break by/ Color:** Select a field by which to break (group) the fields in the chart. This must be a descriptive field. If you select descriptive data, then the points in the chart are grouped into the members of the field. Each member is represented by a different color.

OR

Select a field by which to color the fields in the chart. This must be a numeric field. If you select numeric data, then the point color is a gradient where the highest value is the darkest color and the lowest value is the lightest color.

5. In the **Size** panel, select the field that determines the size of the circle. It must be numeric data.

## Designing the Scatter Chart

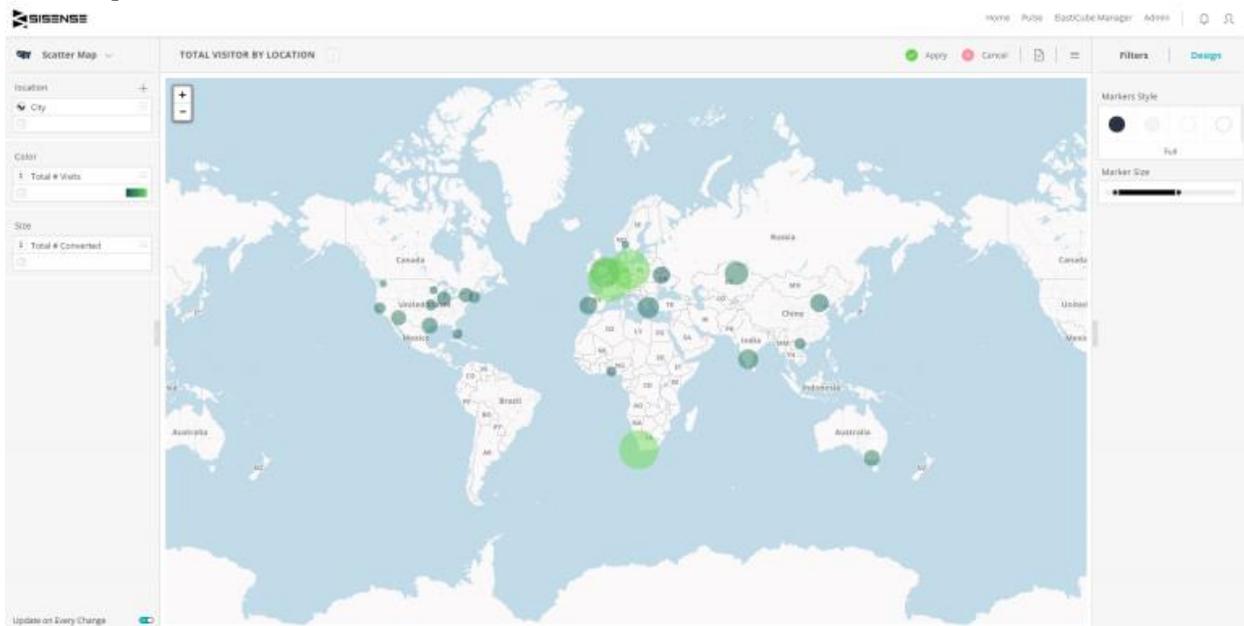
Fine tune the appearance of the SCATTER CHART widget, using the following tools.

- **Legend:** Specify whether to show or hide the Legend, and select its position.
- **Marker Size:** Change the size of the circles. If you are not using a Point value, the slider will change the size of all the points on the map. If you are using the Point value, drag each side of the slider to determine the relative minimum and maximum circle sizes.
- **Value Labels:** Specify whether to show or hide labels indicating the values of each point (circle). Select the orientation of these labels: **Horizontal** (Default), **Diagonal** or **Vertical**.
- **X-Axis** and **Y-Axis:**

- **Grid Lines:** Shows (Default) / hides.
- **Logarithmic:** Displays using orders of magnitude.
- **Labels:** Shows (Default) / hides.
- **Title:** Select the checkbox to display the axis title. To edit the title, type in a new title. Click outside the text box to apply the new value. To change the default minimum, maximum or interval values on the axis, type in the new values, and click outside the text boxes to apply the new values.
- **Values on Axis:** To restore any of the values to their default states, click on the reset button. ↺

# Scatter Map

Scatter Maps allow you to visualize geographical data as data points on a map. The map distinguishes data using different colors and sizes for the data points on the map.



## Adding Data

In the **Location** panel, you can add any field(s) that contain geographic data, and the map will find the most appropriate point. Supported Location types:

- Country
- City
- State/Province
- Latitude/Longitude (for example: 37.5601, -122.50131, but not 37.5601° N, 122.50131° W).

For example, if you have a country field, using it by itself will visualize the countries on the map.

If instead you have a city field, using it will visualize the cities on the map automatically.

## Improving Geographic Accuracy

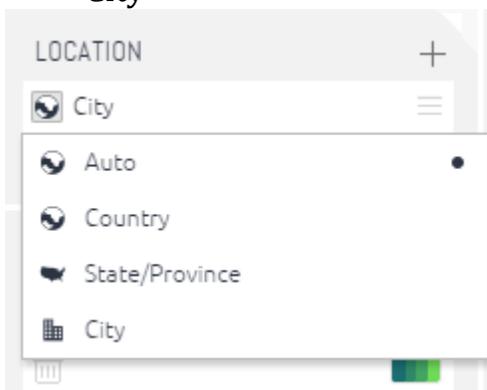
### Use multiple geographic fields

To improve accuracy, it is recommended to use more than one geographic level. For example, if you have both a country and a city field, add both of them to the **Location** panel.

### Define the location type

Clicking on the geographic field's menu will present a location type sub-menu, which contains the following options:

- Auto
- Country
- State/Province
- City



By default, **Auto** is always selected. This option always tries to figure out which fits best. For better accuracy, you can specify the appropriate location type.

## Using Latitude and Longitude data

To visualize latitude and longitude data, you have to add one field containing latitude data, and another field containing longitude data, in this order.

### Color

Select a numeric field here to color the scatter points on the map according to their value and color rule.

The value of the field is displayed in the tooltip of the scatter point on the map.

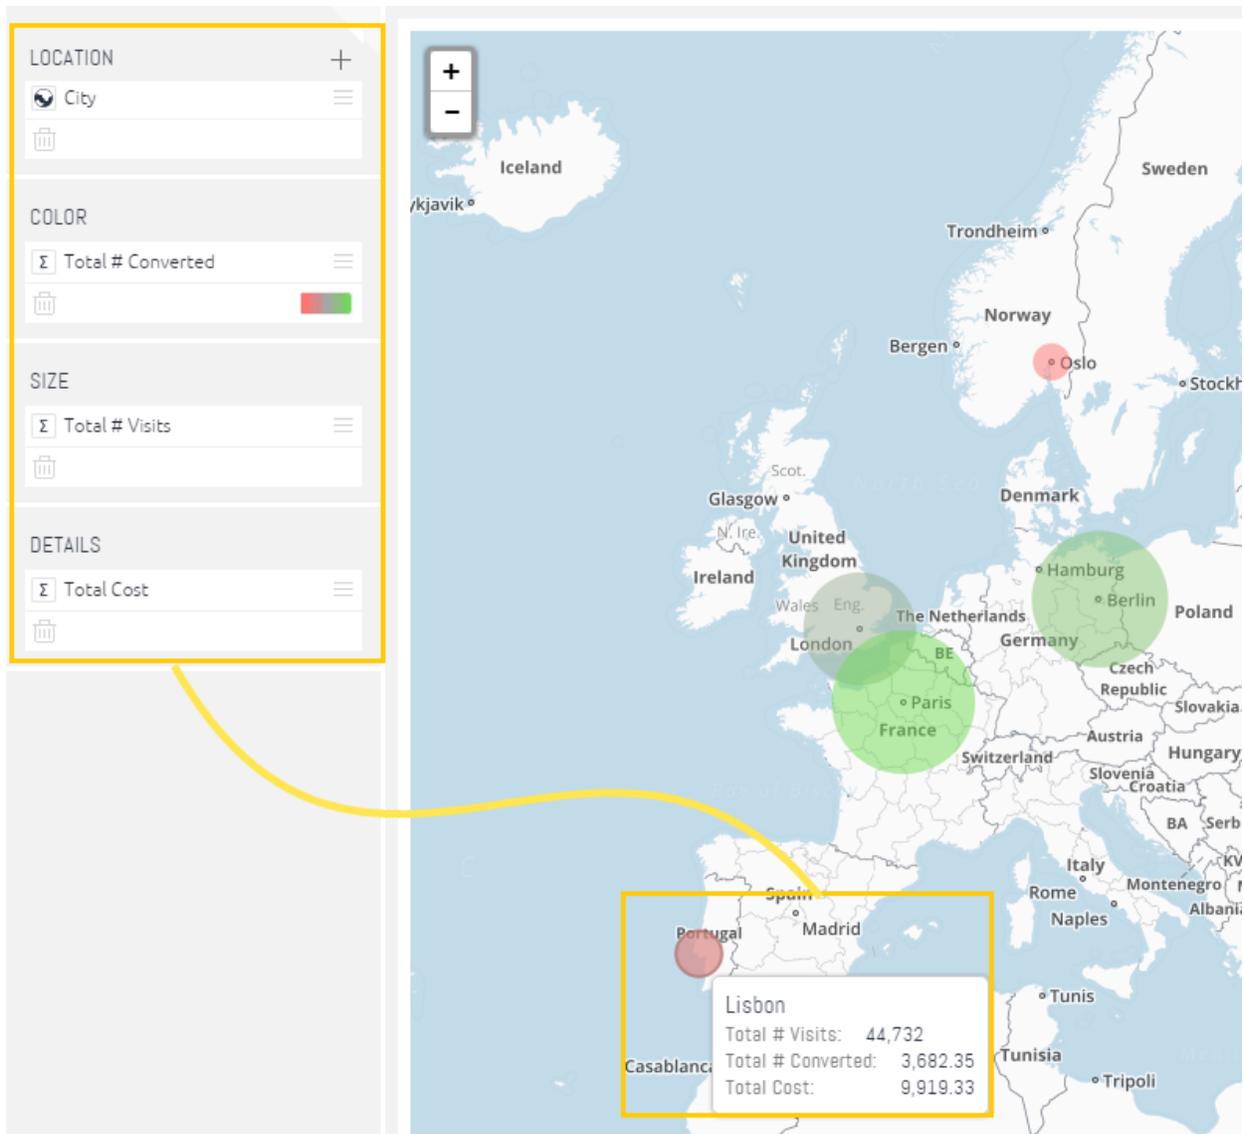
Hover over the point to see it.

### **Size**

Select a numeric field here to make the size of the scatter points on the map relative to their value in the field's value range.

### **Details**

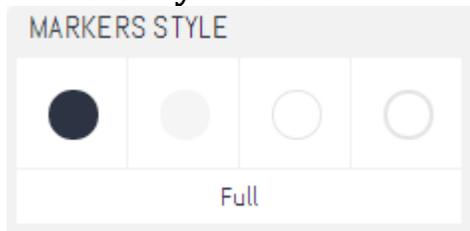
You can select additional data (numeric or descriptive) that will appear when hovering over the data point in the map.



## Designing the Scatter Map

Fine-tune the appearance of the Scatter Map Widget, using the following tools.

- **Marker Style:** Choose between various marker styles.



- **Marker Size:** Select the relative size for the markers.

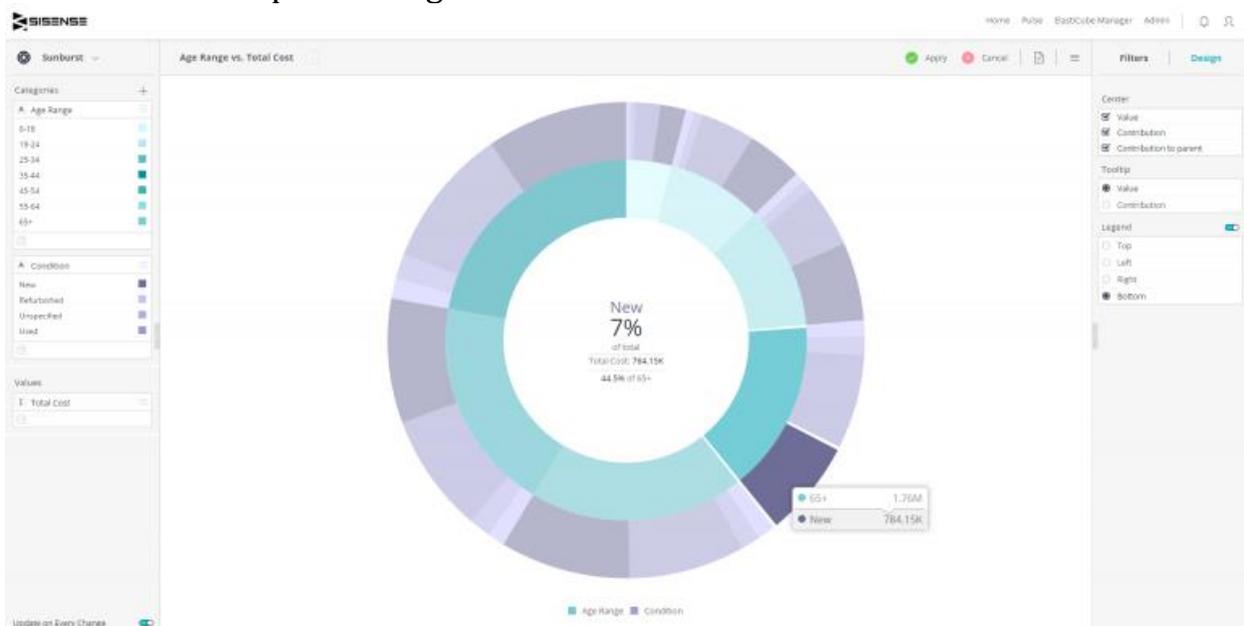


When the Size has a field in it, the slider has two levers controlling the minimum and maximum size of the points.



# Sunburst Widget

The Sunburst widget is similar to a pie chart but is multi-dimensional. Whereas a pie chart combines one field and one numeric value, the Sunburst widget can display multiple rings, one for each field. Each ring in the Sunburst shows a breakdown of its parent ring slice.



## Adding Data

1. In the **Categories** panel, select the fields whose values will be represented as slices of the ring.  
Typically, the **Categories** of a Sunburst widget contain descriptive data. Multiple fields can be added, so that each field is represented as a ring whose values are broken up by its parent ring slice. You can reorder the fields using drag and drop.
2. In the **Values** panel, click **Add +** to select the field whose values determine the size of the ring. Values are typically used to represent numeric data. You can only add one field.

# Designing the Sunburst Widget

Fine-tune the appearance of the Sunburst widget using the following tools.

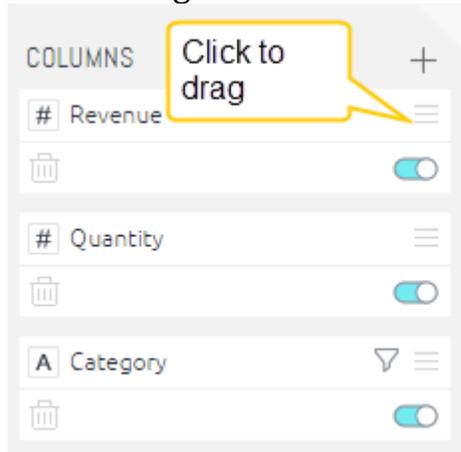
- **Center:** Select which info is presented in the center of the Sunburst widget when hovering over a ring slice:
  - Value – Show the actual numeric value associated with the selected slice.
  - Contribution – Show the contribution (percentage) of the selected slice to the entire field.
  - Contribution to parent – Show the contribution (percentage) of the selected slice to its parent ring slice.
- **Tooltip:** Determine whether to show the actual value or contribution in the slice tooltip when hovering over it.
- **Legend:** Specify whether to show or hide the Legend, and select its position in the widget.

# Table

The Table widget displays a broader view of your data, presenting raw and non-aggregated data in columns, with as much fields and metrics as needed.

## Adding Data

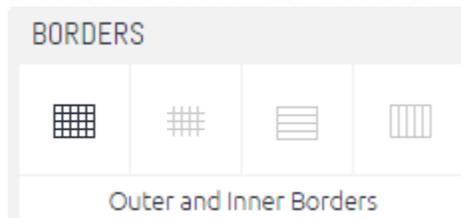
- Click **Add +** to select the field(s) that you want to add to the table.  
To rearrange the order of the columns, click and drag a column up or down.



## Designing the Table Panel

Fine-tune the appearance of the Table widget, using the following tools:

- **Borders:** Select the borders to display in the table.



- **Column Width:** You can resize the column widths as follows:
  - **Fit to Window Size:** Columns are resized automatically to fit the contents in the window.

- **Fit to Content:** You can manually resize the column widths by dragging the column separator.



- **Colors:** Select whether to color (highlight) alternating columns and rows, and the table header.
- **Word Wrap:** Select whether to wrap texts in headers and rows.

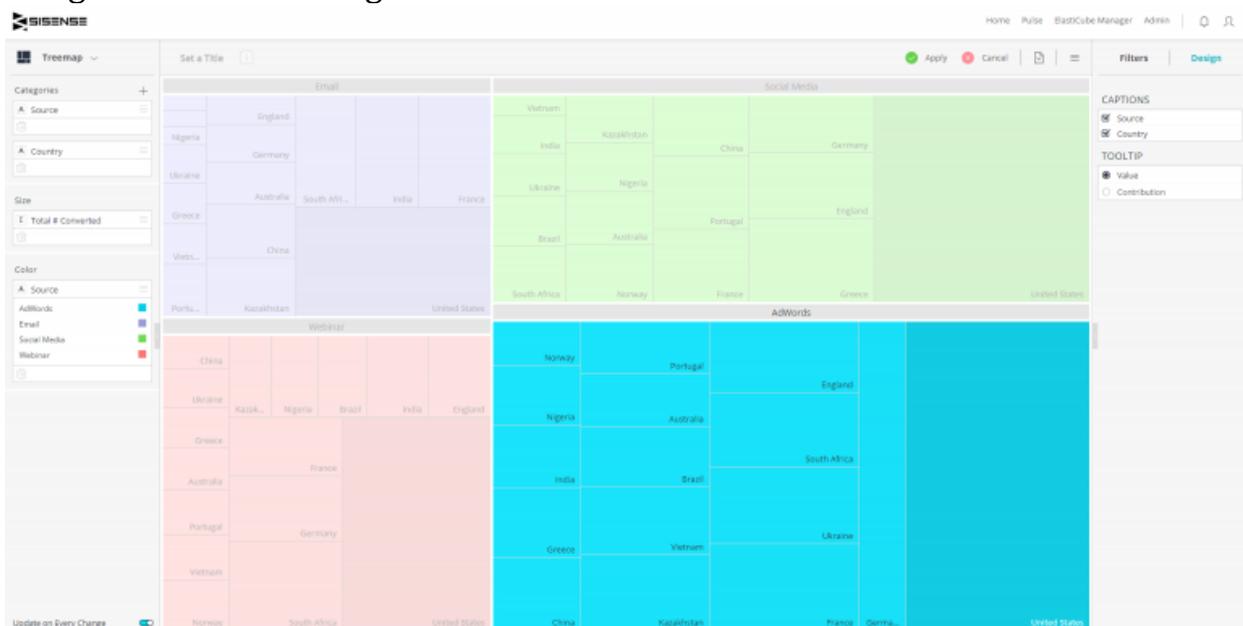
Ca...	Age Range
Monit...	25-34
Monit...	35-44
PDA's	35-44
Proce...	25-34
Table...	25-34
Table...	35-44

Category	Age Range
Monitors	25-34
Monitors	35-44
PDA's	35-44
Processor Upgrades	25-34
Tablet PCs	25-34

- **Infinite Scroll:** Toggle the Infinite Scroll option On or Off:
  - On: Include all items in a single page with scrolling. Additional data loads as you scroll down in the page.
  - Off: Specify how many rows you want in a single page. Above the specified number, you will have paging options.
- Infinite Scroll in the Off position.

# Treemap

The Treemap is a multi-dimensional widget that displays hierarchical data in the form of nested rectangles. This type of chart can be used in different scenarios, for example, instead of a column chart if you have to compare too many categories and sub-categories.



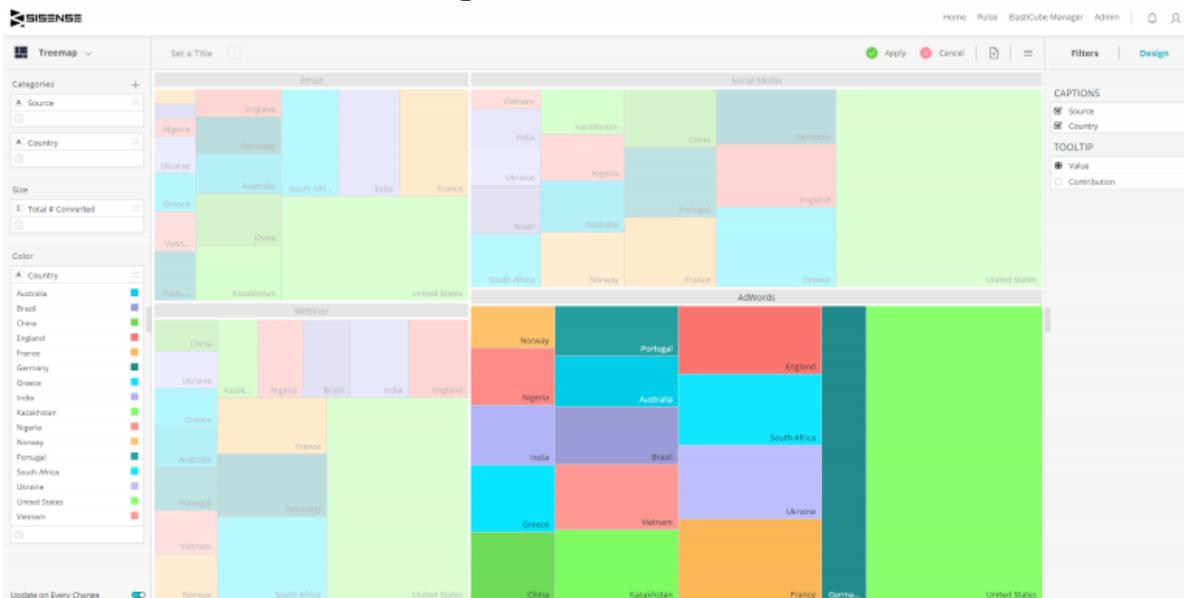
## Adding Data

1. In the **Categories** panel, click **Add+** to select the fields whose values will be represented as rectangles in the chart.  
Typically, the Categories of a Treemap widget contain descriptive data. You can add up to three fields, so that a rectangle hierarchy is created. You can reorder the fields using drag and drop.
2. In the **Size** panel, select the field whose values determine the size of the rectangles by clicking **Add +**. Values are typically used to represent numeric data.  
You can add only one field.

3. In the Color panel, click **Add +** to select the field to use as the color basis for your rectangles.

The following coloring scenarios are supported:

- **Color by a numeric value:** If you choose a numeric value, then each rectangle will be colored according to its value in the field's range of values. You can manipulate the coloring rules in the color formatting menu.
- **Color by field members:** If you choose a field, then all rectangles that are members of that field will be colored the same. For example, if you choose to color by the field that is also used as the first category, then the entire Treemap will be split into groups of colors, like in the image at the top of this page. If you choose to color by the field that is used as the second category, then all rectangles that share the same value will also share the same color, as in the image below:



**Note:** You can only color by a field that is also used in the Categories panel. If you choose a field that is not there, it will be automatically added.

## Designing the Treemap

Fine-tune the appearance of the Treemap widget, using the following tools.

- **Captions:** Control which info is presented in the treemap. A checkbox will be visible for each field that is used in the **Categories** panel. Turning a field on/off will determine if the field member name will be visible inside the

treemap. Note that this info will always be available in the tooltip when hovering over a rectangle, regardless of the on/off setting.

- **Tooltip:** Select whether to show the actual value or contribution in the rectangle tooltip when hovering over it.

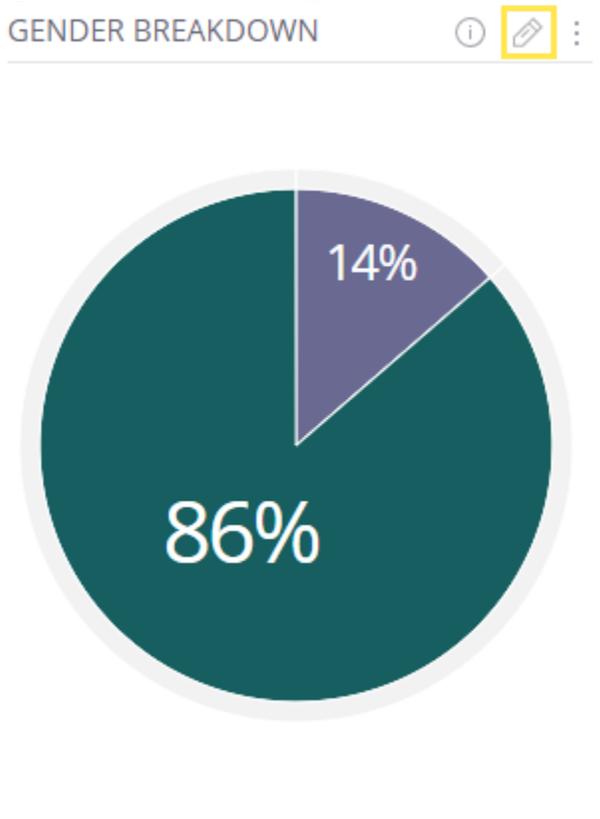
# Widget Designer

The Widget Designer lets you fine-tune a widget's appearance and behavior.

## Opening the Widget Designer

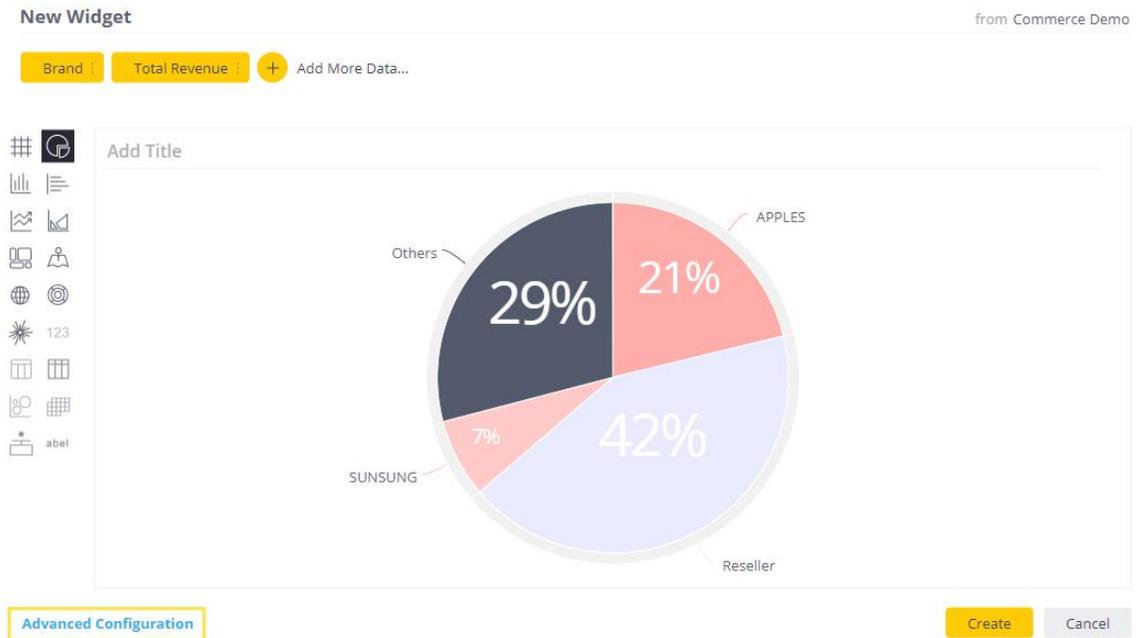
**To display the Widget Designer:**

- On the dashboard, click the **Pencil (Edit)**  button that appears in the top-right corner of a Widget.



- Alternatively, while creating a new widget, in the Widget Wizard, click the **Advanced Configuration** option that appears at the bottom-left of each window. The Widget Designer is then displayed showing the widget in the

same state as in the Widget Wizard preview.



## Navigating the Widget Designer

The Widget Designer enables you to select the data to be included in a widget in addition to providing a variety of options for customizing the visualization used to show the data.



- 1. Selecting the Widget Visualization:** Enables you to change the visualization of the widget.
- 2. Adding Data to the Widget (Data Panel):** Used for selecting the values that you want to appear in the widget and those needed for grouping the data. The options differ depending on the visualization (chart type).
- 3. Previewing the Widget:** The center of the window displays the current design and content of the widget, which automatically changes each time you select a different option in this window.
- 4. Adding a title to the dashboard:** Click Set a Title and type in a new title for the widget. You can also add a widget title directly from the dashboard without entering the Widget Designer. [Click here](#) to learn more.
- 5. Accessing More Options:** Displays a menu of additional options for the widget.
- 6. Filtering the Widget (Filter Panel):** Lets you manage the filters that affect this widget. Learn more in the [filtering page](#).
- 7. Designing the Widget (Design Panel):** Provides a variety of options for fine-tuning the appearance of the Widget, including labels, legends, line types and more. Some of these options need to be turned on for you to configure them. The options differ depending on the visualization (chart type).
- 8. Update on every change:** Selected – Data changes are automatically and dynamically updated in the display. Not selected – you have to click **UPDATE** to display your latest changes in the widget. Clearing this checkbox may be useful when you are working with very large datasets where query times might be slower.

# Fine-tuning a Widget

- To get an overview of the different widgets available, [click here](#).
- To learn more about adding data and fine-tuning the design of a specific widget, click on a widget from the list below.
  - Indicator
  - Column Chart
  - Line Chart
  - Area Chart
  - Bar Chart
  - Pie Chart
  - Polar Chart
  - Scatter Chart
  - Pivot
  - Sunburst
  - Treemap
  - Calendar Heatmap
  - Table

See also [Additional Widget Design Options](#) and [Extending Dashboard Functionality with JavaScript](#).

# Changing the Dashboard's Color Palette

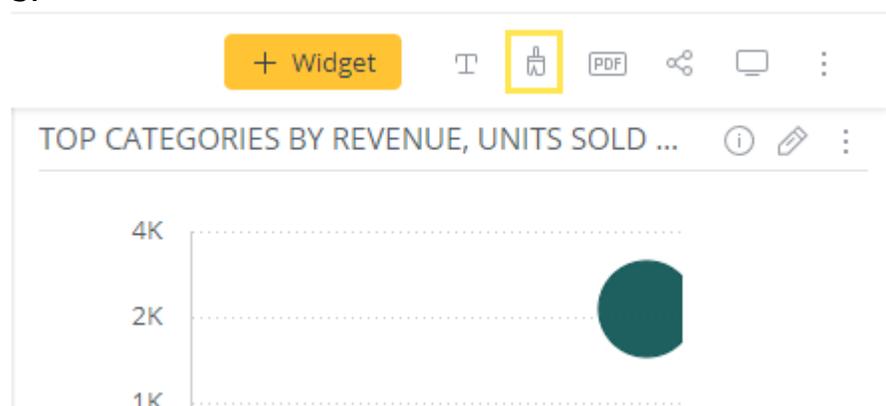
You can easily change the color scheme of the dashboard to suit your professional or brand preferences.

**Note:** The following procedure explains how to change the existing palette by way of selecting a different predefined palette. If you want to create a custom palette, it is possible using the REST API. [Click here](#) to learn more about using the REST API. To learn more about creating a custom palette using the REST API, see this [support article](#).

## To change the color palette:

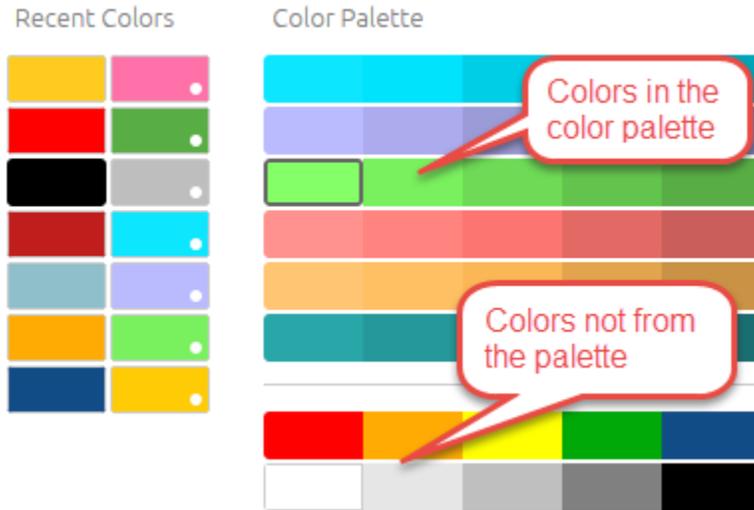
- Click  in the top menu, and select the preferred color palette. The default palette is Vivid.

**Note:** Three of the palettes are suitable for color blindness: ColorBlind 1, 2, and 3.

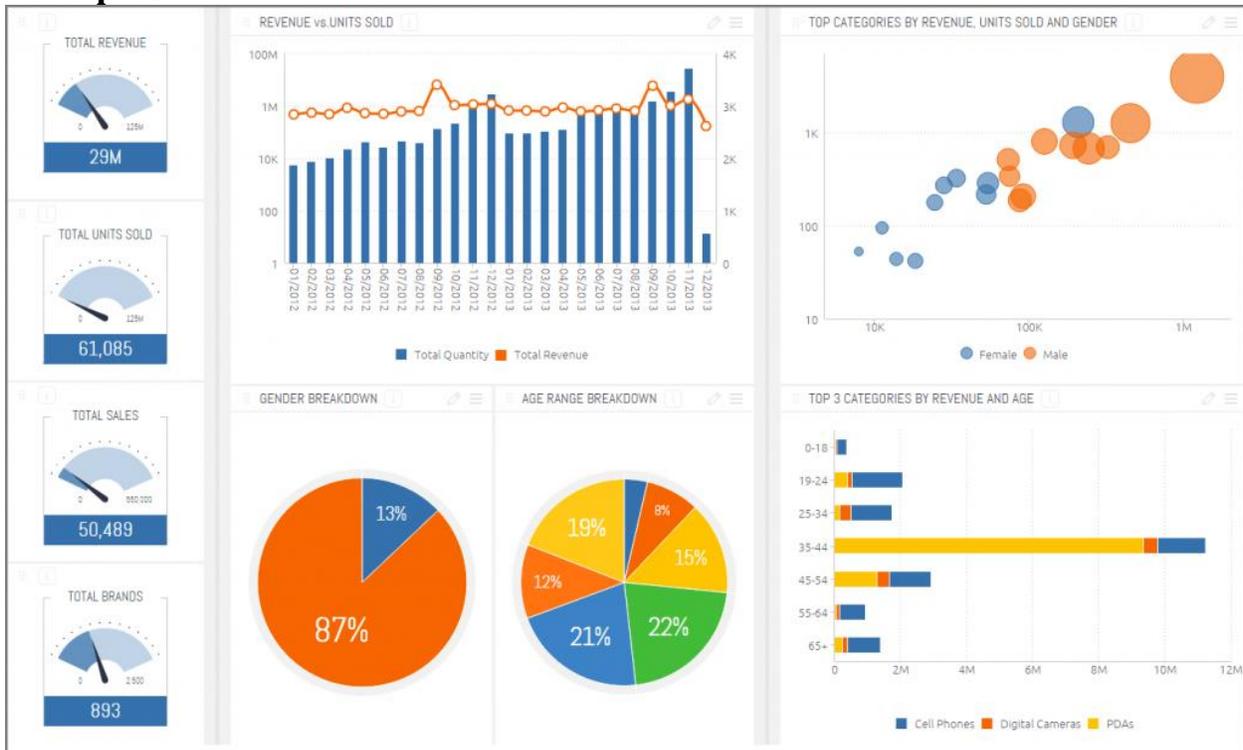


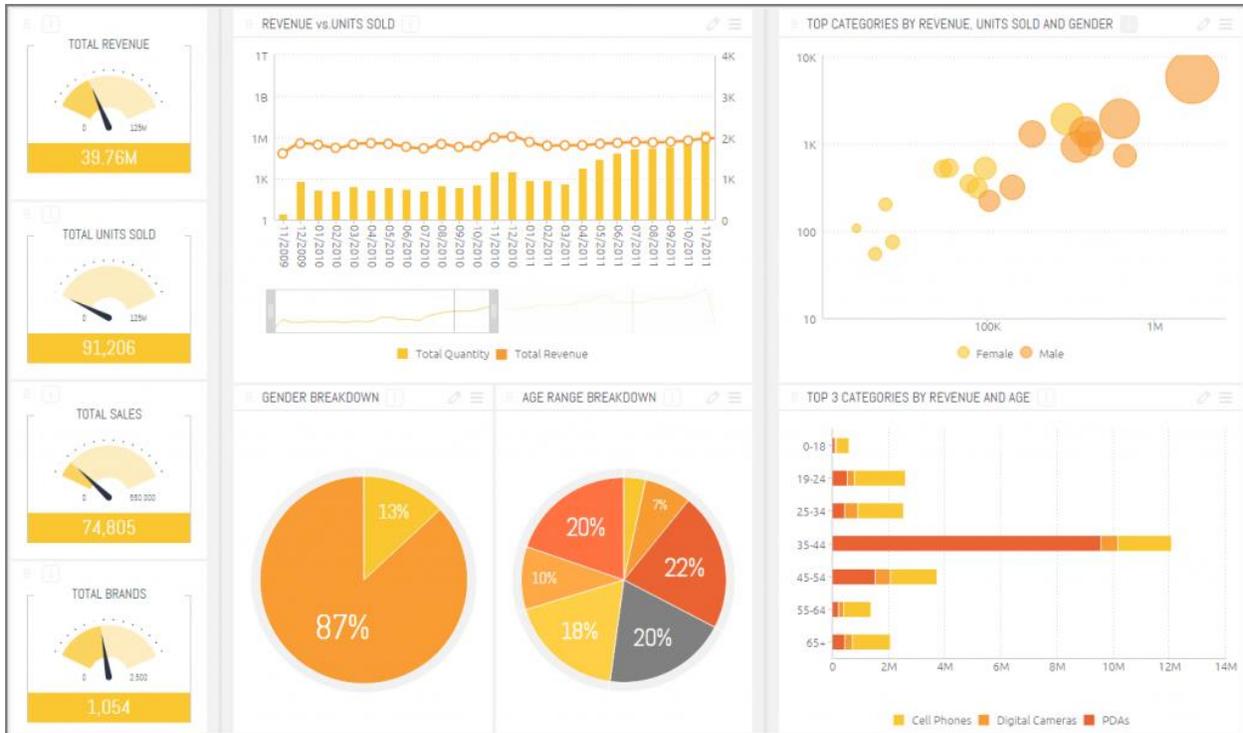
**Note:** When you change the color of an individual widget to a new color that is not one of the palette colors, the new color will be applied to the widget regardless of the applied palette. If you select a new color from one of the palette

colors, then when changing the dashboard's palette, the color will change according to the new palette.



## Examples



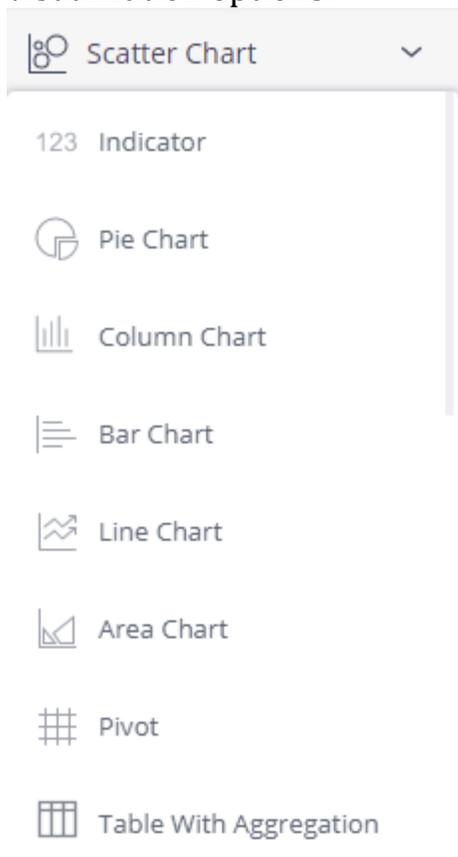


# Changing a Widget's Visualization

You can change the visualization of a widget (for example, from a Line Chart to a Pie Chart). When you change the type, all relevant values and definitions are transferred to the new visualization.

## To change a widget's visualization:

1. Open the Widget Designer, as described in [Opening the Widget Designer](#).
2. Click the Visualization selector, shown below, to display a menu of visualization options.

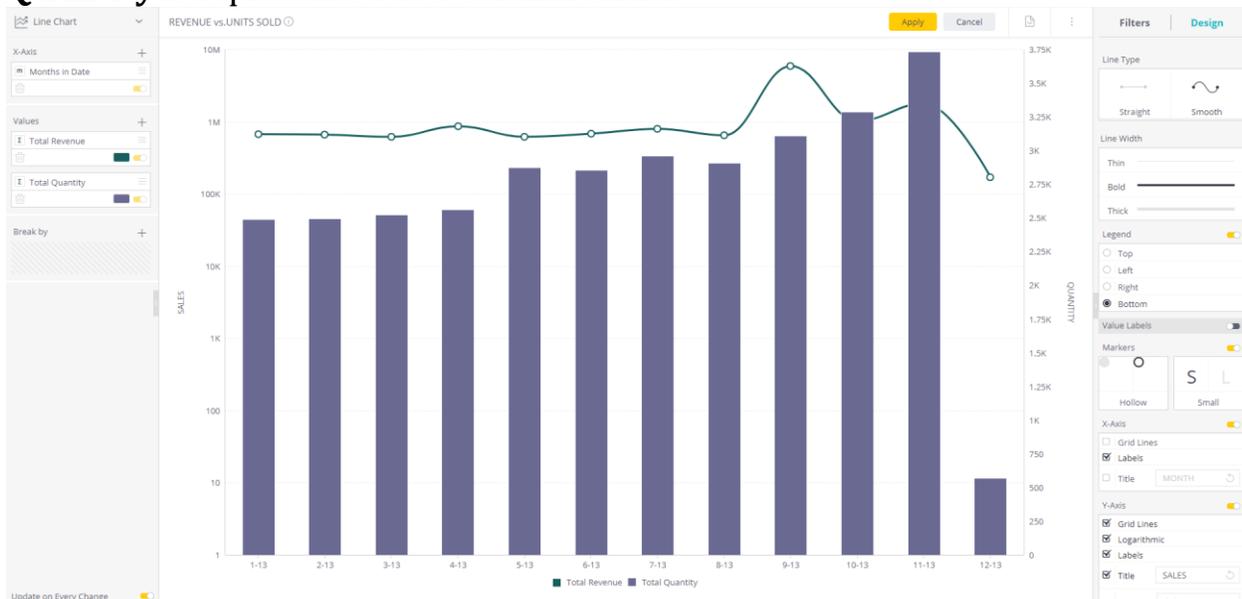


3. Select a different visualization, and click **Apply**.

# Combining Two Types of Visualizations in a Widget

In chart widgets (such as Line Charts, Area Charts, Column Charts and Bar Charts), you can represent one or more selected field(s) using a different visualization than the visualization of the original widget.

For example, the following example shows a Line Chart in which the **Total Revenue** is still represented as a Line Chart (the default), but the **Total Quantity** is represented as a Column Chart.



## To select a different visualization for a field in the widget:

1. In the Widget Designer, in the **Values** area, right-click the field that you want to change.
2. Select **Series Type** from the menu. A menu of alternative visualization types is displayed.
3. Select the visualization to use for this field only.

**Tip:** If necessary, drag the field up or down in the Widget Designer to have a visualization brought to back or front.

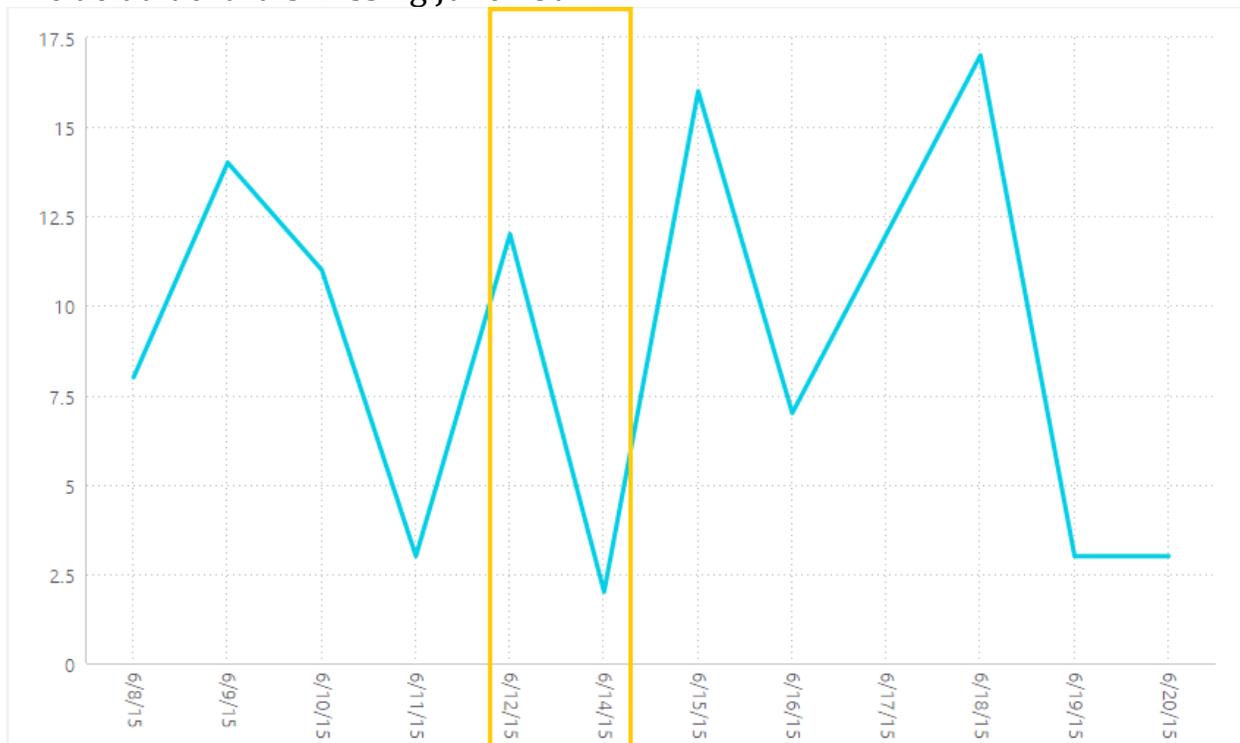
**Note:** The chart keeps the same visualization type as was originally created. The Design panel still displays only the design options of the widget's original visualization type.

# Creating a Continuous Chart with Missing Date Values

If your data is missing date values, you can add the missing dates to your chart to create a continuous data flow. In addition, you can select whether to display missing dates as gaps in your chart, or include null values as zeros, to avoid gaps, and maintain a continuous data flow.

For example, if your chart begins June 8th and ends June 20th, but there is no value on June 13th, you can choose to include the missing date on the x-axis. You can also choose whether to continue the graph (create zero values) or cut the graph where there are missing values.

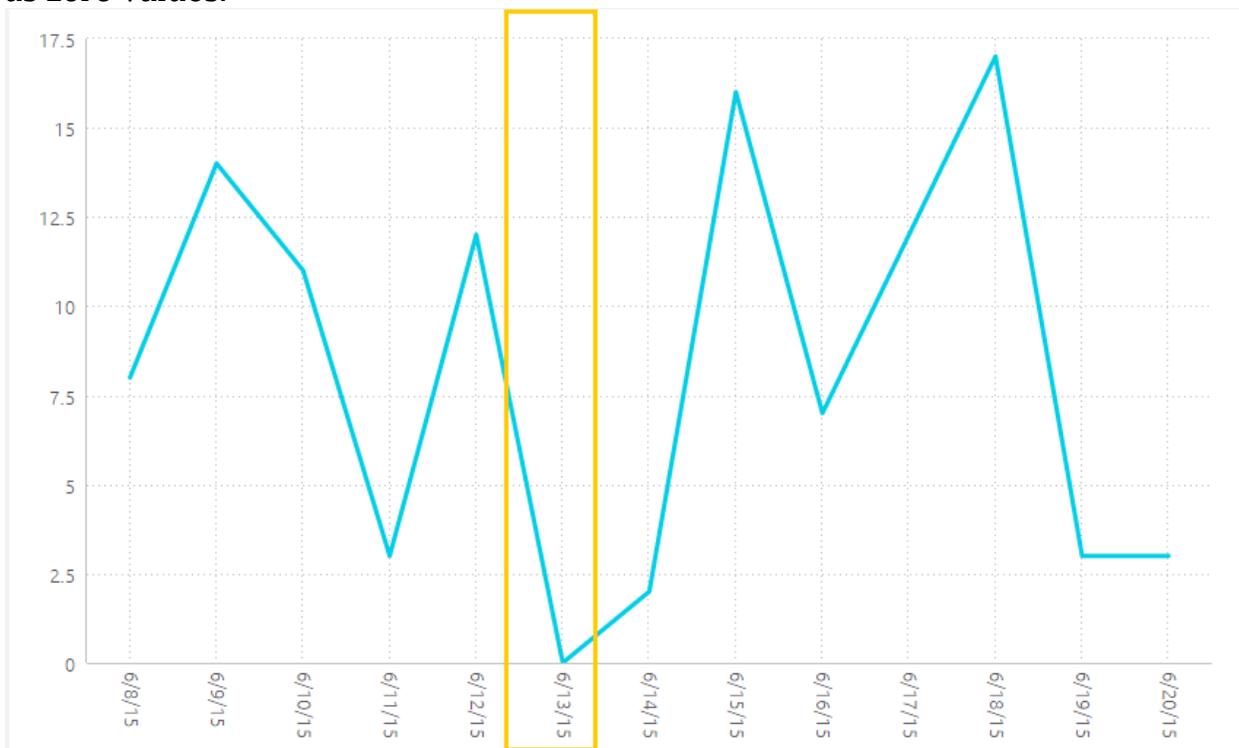
The default chart is missing June 13th.



After selecting the Continuous option, the chart includes the missing date(s).



In addition, you can select to continue the line, by treating null (missing) values as zero values.

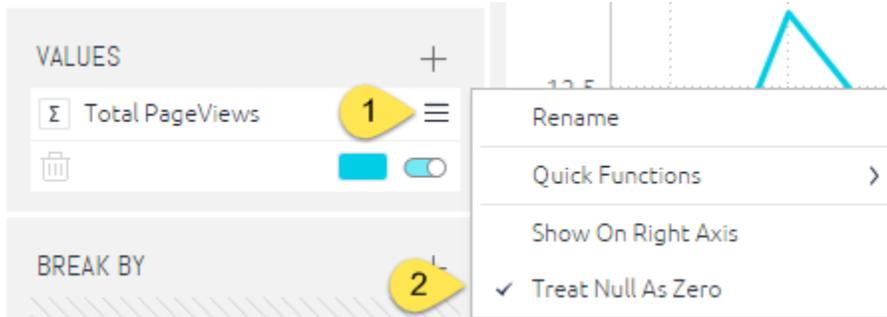


**To create a continuous chart:**

1. In the X-Axis panel, click on the menu, and select **Continuous**.



2. To treat null values as zero values, click on the menu of the Values panel, and select **Treat Null as Zero**.



# Adding Drill Hierarchies to Widgets

Dashboard Viewers have the option to [drill down in a widget](#), and get an in-depth view of a selected value.

While Viewers can select a drill-down path from the complete list of fields, it's easier to select a commonly needed drill hierarchy from a short list. This is especially true when there is a lot of data, and the Viewer needs to remember specific fields, and select them each time.

As a Designer, you can add predefined drill hierarchies to widgets.

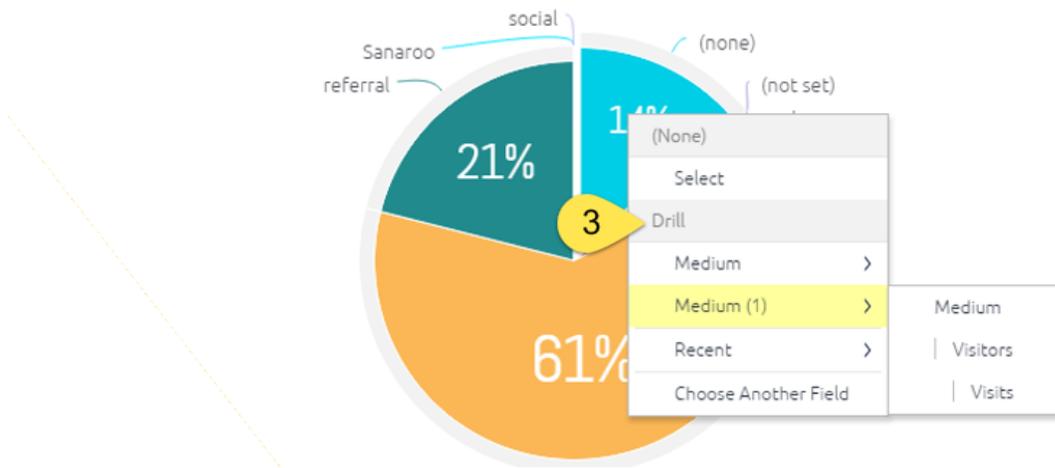
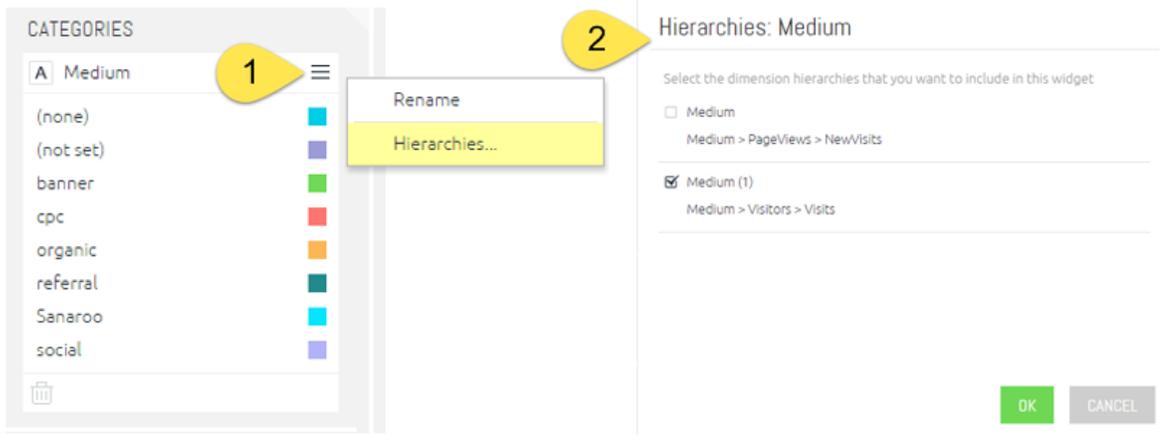
**Note:** Your Sisense Administrator must first create the drill hierarchies before you can select and add them in the Widget Designer. [Click here to learn more.](#)

## **To enable drill hierarchies in a widget:**

1. In the Widget Designer, click on the menu of the value (dimension), and select Hierarchies...
2. Select the checkboxes next to the drill hierarchies that you want to make available to dashboard viewers.

**Note:** Hierarchies with a lock icon will always appear in the widget. To

unlock this option, please contact your Sisense Administrator.



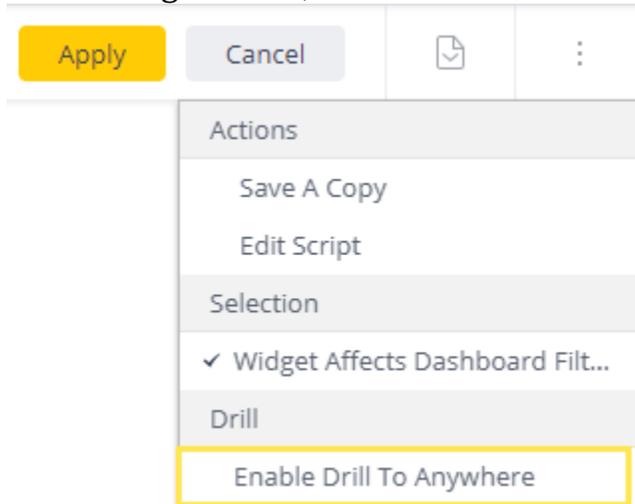
## Disabling the Drill to Anywhere Option

As a Designer, you can restrict viewers to predefined hierarchies only, by disabling the option to drill down freely to any field.

The drill-to-anywhere option is enabled by default for all users and widgets. You can disable the drill-to-anywhere option in the Sisense web app per widget (see below), or by user role in the REST API.

**To disable the drill-to-anywhere option:**

- In the widget menu, deselect the Enable Drill to Anywhere option.



As a result, the Choose Another... option is removed from the widget's Drill options.

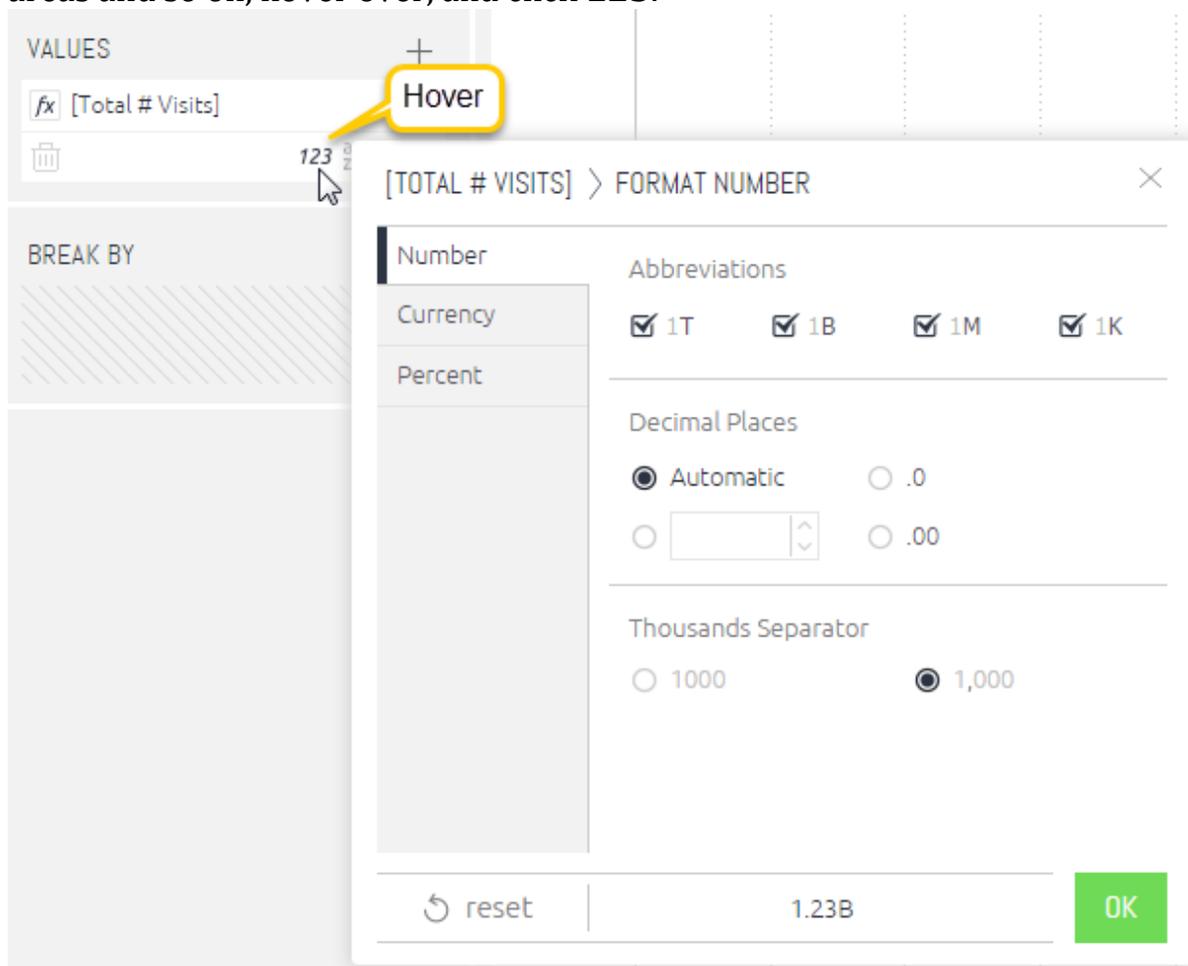
The left image below displays the widget with the option to drill to anywhere. The image on the right displays the widget with only predefined drill hierarchies.

# Formatting Numbers in Widgets

You can change the way numbers are displayed and formatted in widgets. For example, you can display a number as a percentage, and you can define the number of decimal places.

## To format numbers in a widget:

1. In the Widget Designer, for example, in the X-Axis, Y-Axis, Values, Categories areas and so on, hover over, and click **123**.



2. Select how the number, currency or percentage appears. The preview area at the bottom of this window shows a preview of how each option will appear in the widget.

## Locale Settings

The formats for dates, times and numbers in your dashboards are based on your computer's operating system or browser's locale settings (depending on which browser you use).

The following locales are supported:

- Unites States
- United Kingdom
- Israel
- Canada
- South Africa
- Australia
- Netherlands
- Germany
- Ireland
- Mexico
- France
- China
- Brazil

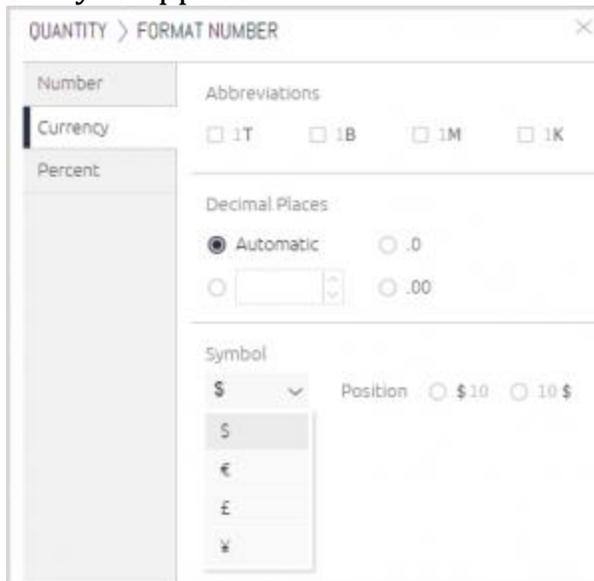
**Notes** You can override the default locale settings for all users, by updating a parameter via the Rest API. To use a locale file that is not included in the above list of countries, you can manually add a locale file in the following location: `C:\Program Files\Sisense\PrismWeb\client\resources\base\localization`. Click [here](#) to access Locale files and view their codes.

## Currency Settings

The default currency symbol used in the dashboard is the US dollar (\$), however you can change the symbol by either selecting another symbol from the drop-down list, or by typing in a different symbol. The list of currency symbols will

include common currency symbols, the currency defined in your browser's locale settings, and the symbol that you added (if you did so).

After you select a different symbol, additional viewers will see the same symbol that you applied.

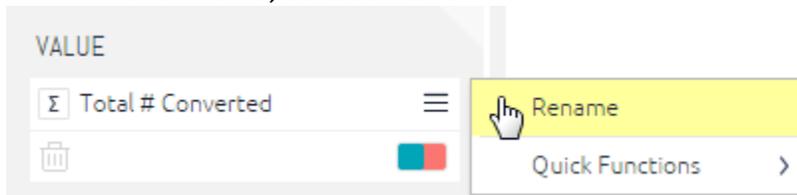


# Renaming Fields in Widgets

The names of fields in widgets are taken from the raw data. To make these names more comprehensible, you can rename the fields in your widget.

## To rename a field in a widget:

1. In the Widget Designer, right-click on a field, for example, in the X-AXIS, Y-AXIS, VALUES, CATEGORIES areas and so on.
2. In the menu, select **Rename**.



**Note:** Renaming a field only affects how it is displayed in the current widget. This field's name is not affected in the Data Browser or in other widgets in the dashboard.

# Selecting Colors in Widgets

Different field values are represented by different colors. The Widget Designer automatically assigns these colors and displays them in the widget legend that is displayed in the Widget Designer and in the dashboard.

The Widget Designer provides the following options for manipulating these colors:

- **Single Color:** For selecting a specific color for all values
- **Range:** For selecting a range of colors
- **Conditional:** For defining colors based on conditional data.

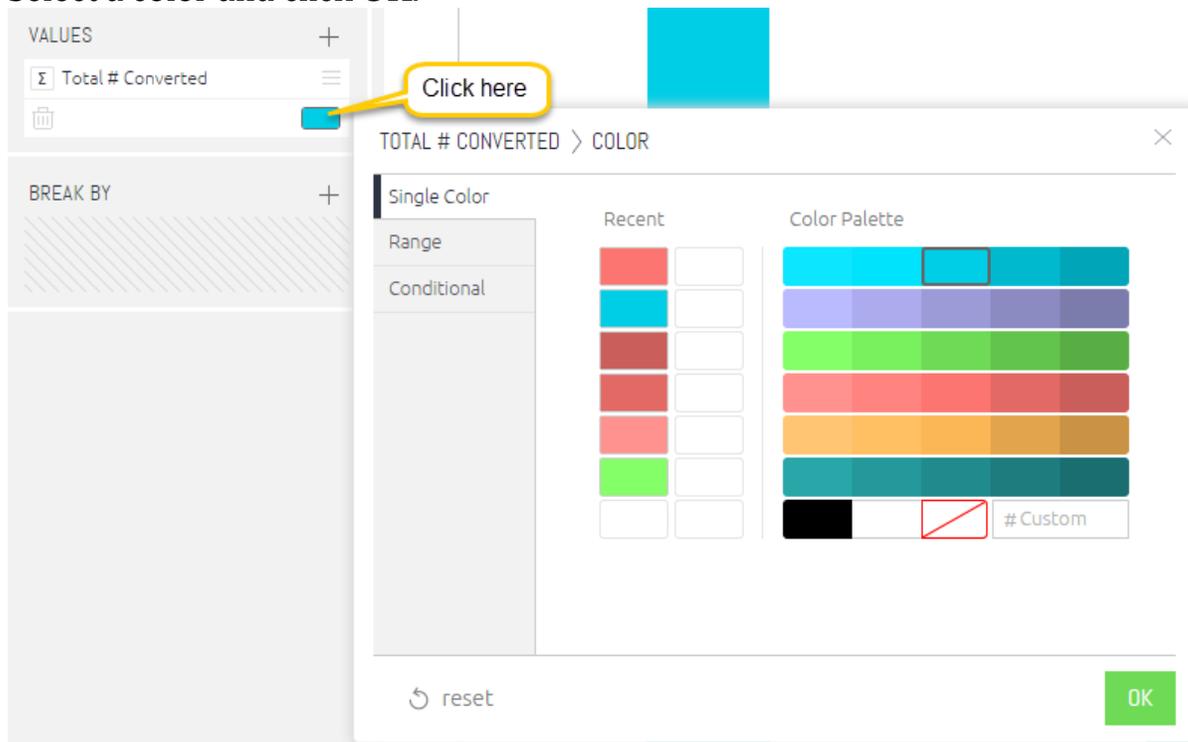
## Selecting a Specific Color – Color Tab

The following procedure describes how to assign a specific color to a data value that does not change. This color only applies to the widget in which it is defined.

**To select a specific color:**

1. Click on a color in the Widget Designer Data Panel to display a color picker window.
2. Select the **Single Color** tab.

3. Select a color and click **OK**.



**Note:** To select a custom color that does not exist in the palette, enter a hexadecimal value into the # Custom field.

## Selecting a Range of Colors – Range Tab

The following procedure describes how to assign a range of colors to a field according to its value.

### To select a color range:

1. Click on a color in the Widget Designer Data panel to display a color picker window.
2. Select the **Range** tab.
3. By default, the range configuration is set to 'Auto', which automatically assigns a color range.

Select **Manual** to manually configure the color range, as follows:

- **Change Left/Right Colors:** Click the color rectangles at the right and left to change the start and end color of the range.
- **Customize Range:** By default, the color range is set to match the minimum and maximum values of the data. Instead, you can manually set a minimum value. All values that are below this value will be assigned to the color on

the left. Accordingly, selecting a Max value means that all values larger than it will be assigned the color on the right. By default, the 'Middle' color is grey. All values larger than this will start gradually getting the right color. All values smaller than this will start gradually getting the left color. Changing the 'Middle' value controls which value the color gray is assigned to.

- **Coloring Method:** Selecting **Min/Max** will display the full color range. Selecting **Min** will display all the colors between the minimum and middle values, and higher than the middle value will be gray. Selecting **Max** will display all values between the middle value and the maximum value, while all other values will be displayed in gray.
- **Transition:** Select Steps to define a distinct number of colors in the range. Select **Gradient** for a linear color range.

TOTAL # VISITS > COLOR ×

Single Color

Configuration:  Auto  Manual

Range



Conditional

Coloring method:  Min  Max  Min / Max

Customize range:

Transition:

Steps   Gradient

↻ reset
OK

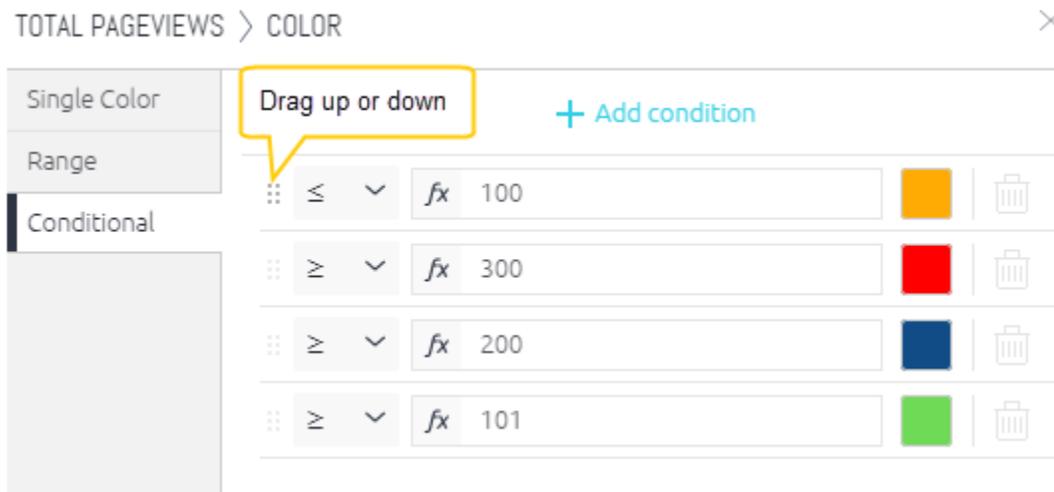
## Defining Conditional Coloring – Conditional Tab

Define colors based on the value of the field as defined by a logical expression.

In addition, for Indicator Widgets, you can also define whether an icon appears in the widget and select the icon's color. For example, you can define a red arrow pointing downwards for a negative value.

**To set conditional coloring:**

1. Click on a color in the Widget Designer Data panel to display a color picker window.
2. Select the **Conditional** tab. Each row in this window represents a condition to which you can apply a value. If the condition is true, then that value is shown in the color that you select in this row. For example, one row can define that a negative value is displayed in red and another row can define that a positive value is green.
3. In each row, define an expression to be applied to the value of this field by selecting an operator and specifying a value, as shown below.
4. Select a color for the field.
5. If required, add and define additional condition rows by clicking **+ Add condition**.
6. In the dashboard, these conditions are evaluated in the order in which they appear in this window – from top to bottom. The first condition that is true, determines the color. Reorder the condition rows as necessary by dragging them up or down.
7. Click **OK**.



# Configuring How Filters Affect the Dashboard and Widgets

This topic explains how dashboard filters work together with widget filters, and how you can configure different filtering behaviors and their effects on the dashboard.

The following procedures are explained below:

- [Configuring how Dashboard Filters Affect a Widget](#)
- [Defining how a Widget Behaves when Filtered](#)
- [Defining how a Selection on the Widget Affects the Dashboard](#)

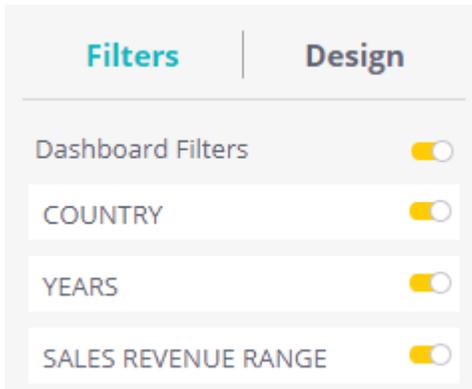
## Configuring how Dashboard Filters Affect a Widget

Dashboard filters affect all the widgets in the dashboard, except widgets that you have defined to be independent. An independent widget is not affected by dashboard filters.

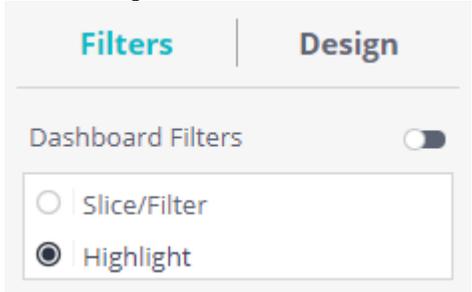
### **To make a widget filter-independent:**

1. Open (edit) this Widget in the Widget Designer, as described in [Adding Widgets to a Dashboard](#). The top part of the Widget Filters panel on the right of the Widget Designer lists the Dashboard Filters that are defined for

this Dashboard.

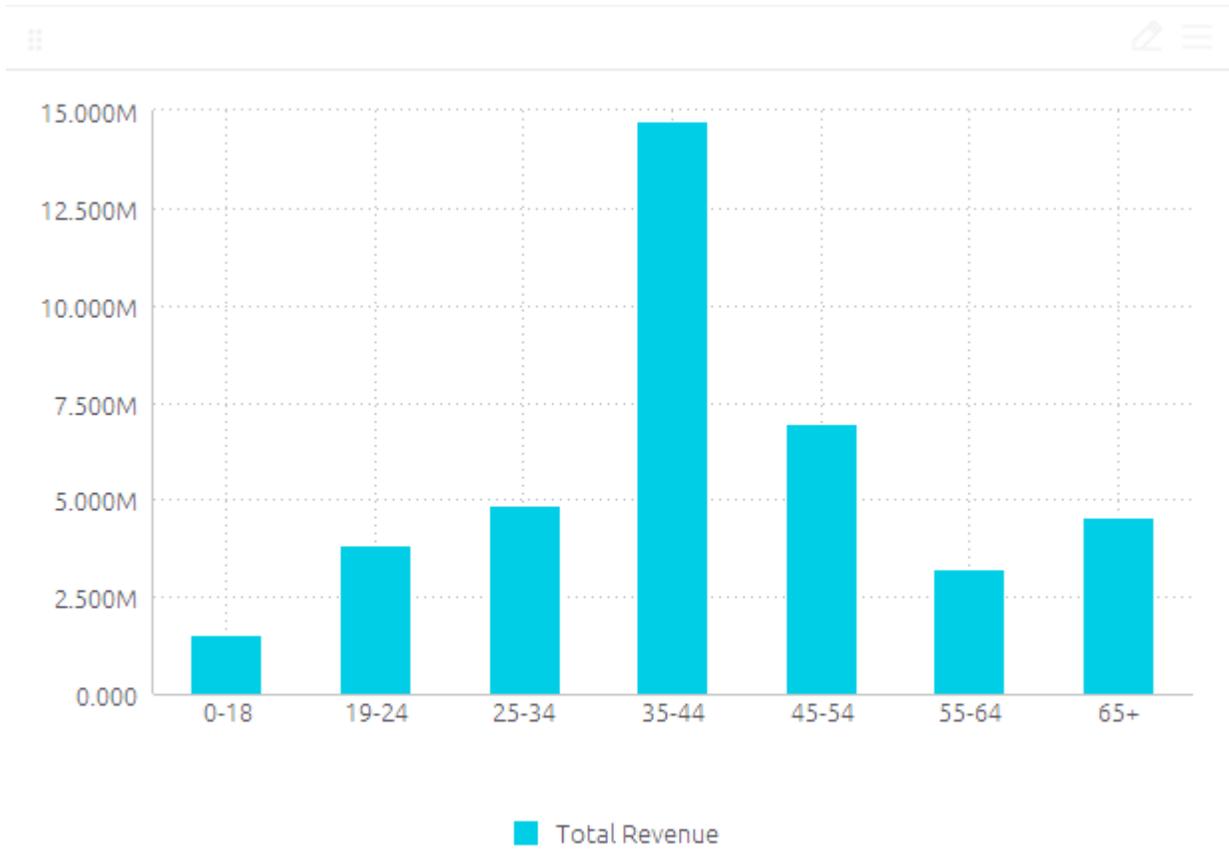


2. Turn off the Dashboard filters that you do not want to affect this Widget: The toggle on icon  (default) indicates that the Widget is affected by this Dashboard filter. The toggle off  icon indicates that the widget is not affected by this Dashboard filter. You also have the option to make this Widget independent of all current and future Dashboard Widgets by turning the Dashboard Filters option off, as shown below:



## Defining how a Widget Behaves when Filtered

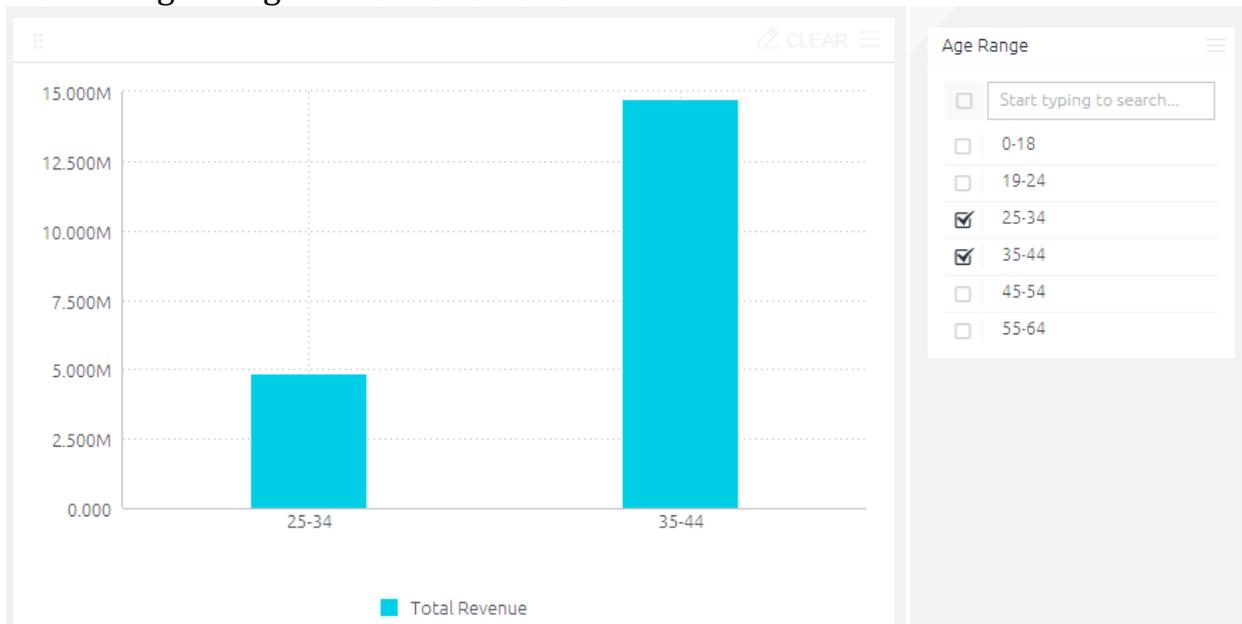
When fields that are visible in a widget are filtered, it can be displayed in two ways – Slice/Filter or Highlight. For example, this column chart shows revenue for different age categories:



If the chart is set to Highlight mode, filtering a subset of age categories will highlight them in the chart:

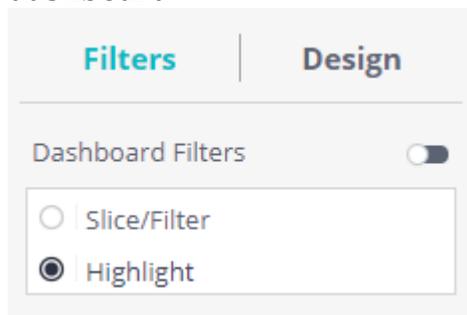


If the chart is set to **Slice/Filter** mode, the same filter will remove all non-selected age categories from the chart:



**To define the filter behavior mode:**

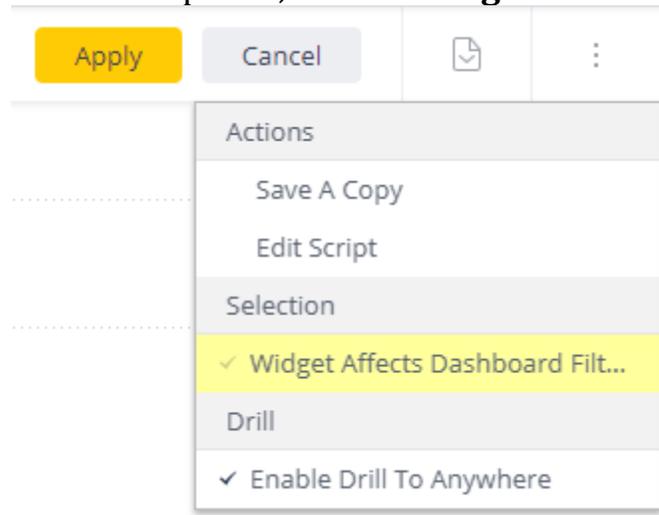
- Select **Slice/Filter** or **Highlight** in the Filters panel on the right side of the dashboard.



## Defining how a Selection on the Widget Affects the Dashboard

When making a selection on a widget, a filter is added to the dashboard with the selected part of the chart as the applied filter.

To disable this option for a widget, open the widget in editing mode, and from the menu options, disable **Widget Affects Dashboard Filters**.



# Creating Dashboard Filters

This topic explains how to create dashboard filters using the different filtering options.

In this section you will learn how to:

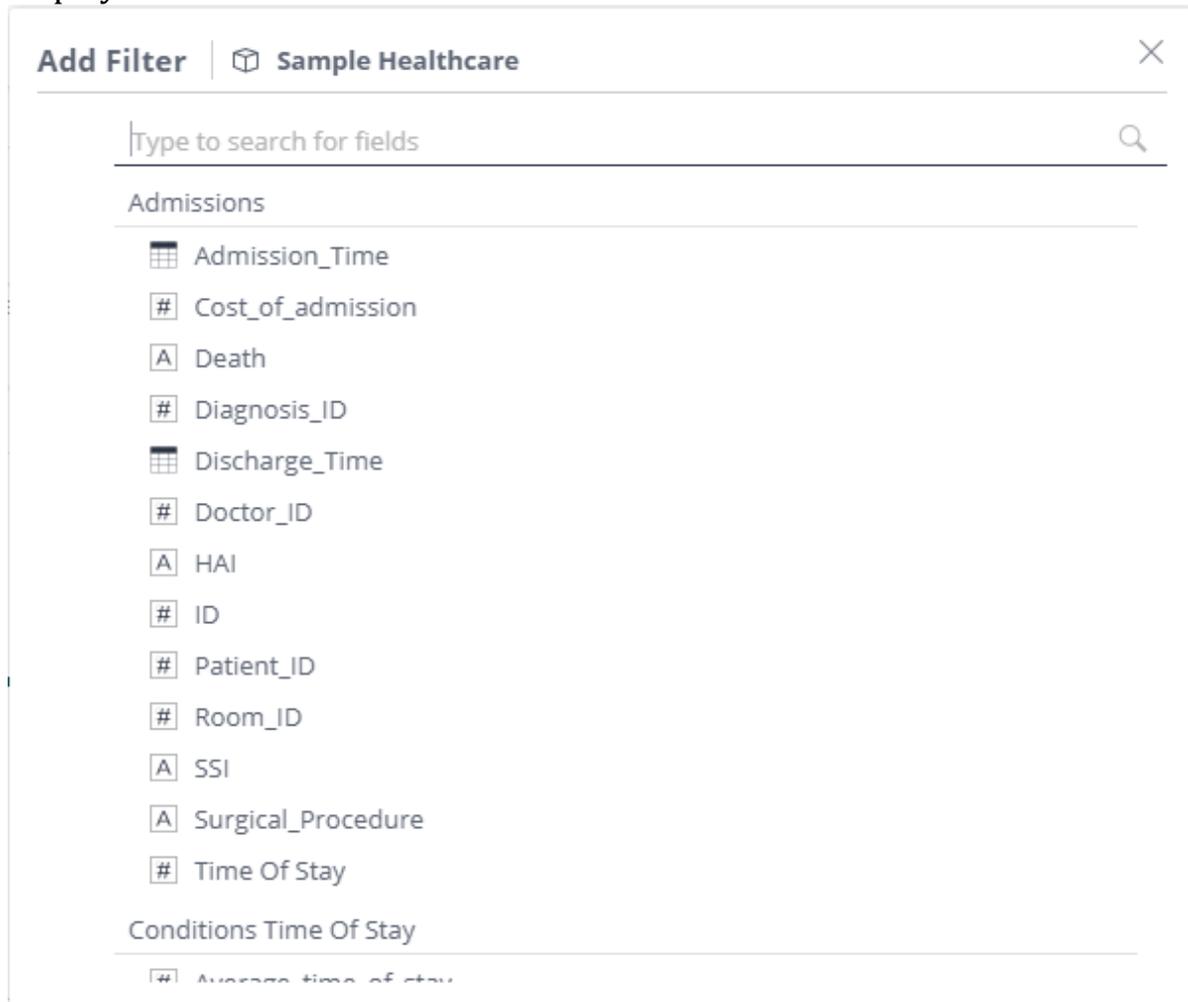
- [Create a Dashboard Filter](#)
- [Create a List Filter](#)
- [Create a Text Filter](#)
- [Create a Numeric Value Filter](#)
- [Create a Top/Bottom Ranking Filter](#)
- [Create a Time Filter](#)
- [Wildcards](#)

## Creating a Dashboard Filter

**To create a dashboard filter:**

1. Click **Filter Your Dashboard** in the dashboard's right panel (if it's your first filter), or **+** if you are adding another filter. The Add Filter dialog box is

displayed.



2. (Optional) If you have multiple data sources, select the data source that contains the fields you want to filter. Filters applied to fields from one data source do not affect fields from another data source even on the same dashboard.
3. Select the field by which to filter.
4. Select the type of filter from the pane on the left, such as **List**, **Criteria**, **Calendar** or **Ranking**. The options provided in this window may vary according to the selected field type (Date, Numeric or Descriptive).
5. Select the required filtering details in the panel on the right to create one of the following types of filters:
  - Creating an Include/Exclude Filter
  - Creating a Text/Numeric/Date Filter
  - Creating a Top/Bottom Ranking Filter
  - Creating a Calendar FilterThese are explained in depth below.

6. Click the **Star Favorite** button to save this filter. You can give your starred filter a unique name by clicking in the text field and editing the given name. This will be useful when you need to quickly add a new filter (with same field type) or update a filter to reflect the unique parameters that you defined in your starred filter.  
Note: Your favorite filters are visible to shared users.
7. Click **OK** to filter the data in the Dashboard and add this filter to the FILTER panel on the right of the Dashboard.  
After a filter has been added to the Dashboard, it stays displayed on the Dashboard's FILTER panel. Any changes that you make to the filter are immediately reflected in the Dashboard.  
Tip: You can rearrange the filters in the Filters panel by dragging a filter above or below another filter.

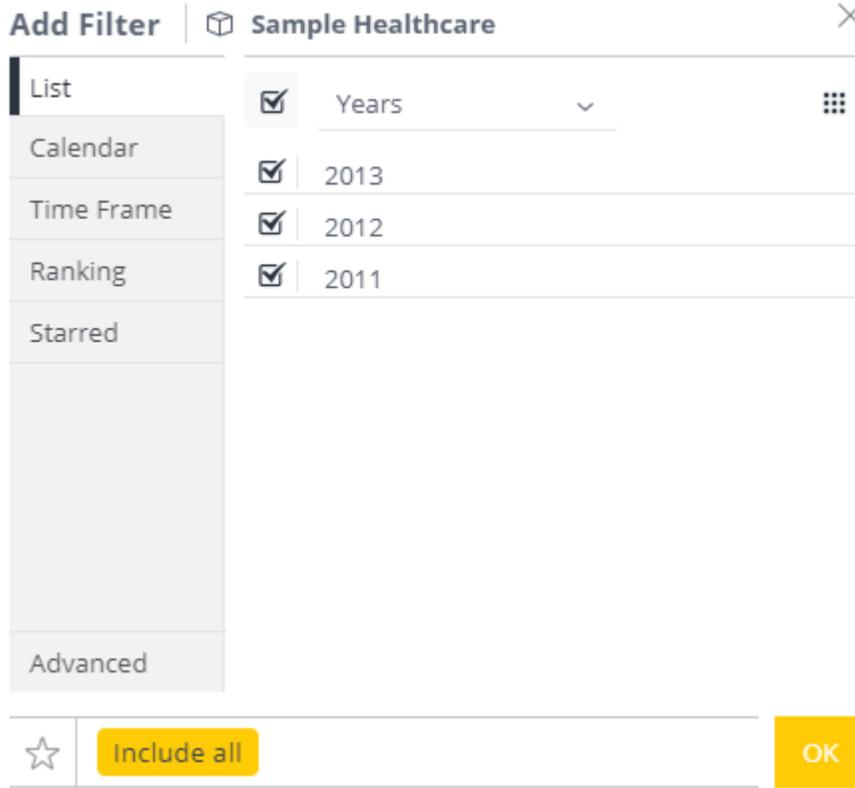
## Creating an Include/Exclude Filter (List Filter)

This type of filter enables you to define whether selected field values are included or excluded.

### To define a List filter:

1. Follow Steps 1 – 3 of [Creating Dashboard Filters](#).

2. Select the **List** option in the left panel.



3. Select/clear the field values to be included/excluded in the Dashboard.
4. Click **OK** to filter the data in the Dashboard and add a List filter to the Filter panel of the Dashboard.

## Selecting Single Selection or Multi Selection

By default list filters are set to multi selection. To toggle between single and multi-selection click the button on the upper right of the list filter:

## Multi Selection:

**Add Filter** |  **Sample Healthcare** ×

List	<input type="checkbox"/>	
Values	<input checked="" type="checkbox"/>   1	
Ranking	<input type="checkbox"/>   2	
Starred	<input type="checkbox"/>   3	
	<input type="checkbox"/>   4	
	<input type="checkbox"/>   5	
	<input type="checkbox"/>   6	
	<input type="checkbox"/>   7	
	<input type="checkbox"/>   8	
	<input type="checkbox"/>   9	
	<input type="checkbox"/>   10	
Advanced	<input type="checkbox"/>   11	

 **1** OK

## Single Selection:

The screenshot shows a dialog box titled "Add Filter" for "Sample Healthcare". On the left, there is a vertical list of filter types: "List", "Values", "Ranking", "Starred", and "Advanced". The "Values" filter type is selected. To the right of this list, there is a vertical list of numbers from 1 to 11. The number "1" is selected, indicated by a filled radio button. At the bottom of the dialog, there is a star icon, a yellow box containing the number "1", and an "OK" button.

## Creating a Text Filter

This type of filter enables you to filter according to text matching.

### To define a Text filter:

1. Select the Text tab.
2. Follow steps 1 – 3 of [Creating Dashboard Filters](#). Select a Textual field.
3. Select the operator, for example Starts with or Containing.
4. Type in the value by which to filter.
5. If you require additional conditions, click **+ Add condition**. Select **AND** or **OR** to define how filter operators are combined.

## Creating a Numeric Value Filter

This type of filter enables you to limit your dashboards to specific value ranges.

For example:

- Keep only sales above 100\$.
- Keep only product IDs between 1000020 and 1000030.

**To define a Numeric filter:**

1. Select the **Values** tab.
2. Follow steps 1 – 3 of [Creating Dashboard Filters](#).
3. Select the operator, for example Equals or Greater than.
4. Type in the value by which to filter.
5. If you require additional conditions, click **+ Add condition**. Select **AND** or **OR** to define how filter operators are combined.

## Creating a Top/Bottom Ranking Filter

This type of filter enables you to select whether to include only the top/bottom ranking fields.

**To define a Ranking filter:**

1. Follow steps 1 – 3 of [Creating Dashboard Filters](#).
2. Select the **Ranking** filter tab.
3. Select **Top** or **Bottom** and the number of items to include.
4. In the **By** field, select a field by which to rank the fields. A list of the fields that match these definitions is displayed so that you can verify that the results are as you expect.
5. Click **OK** to filter the data in the Dashboard and add a Ranking filter to the FILTER panel of the Dashboard.

## Creating a Time Filter

There are two types of time filtering options, Calendar and Dynamic Time filters:

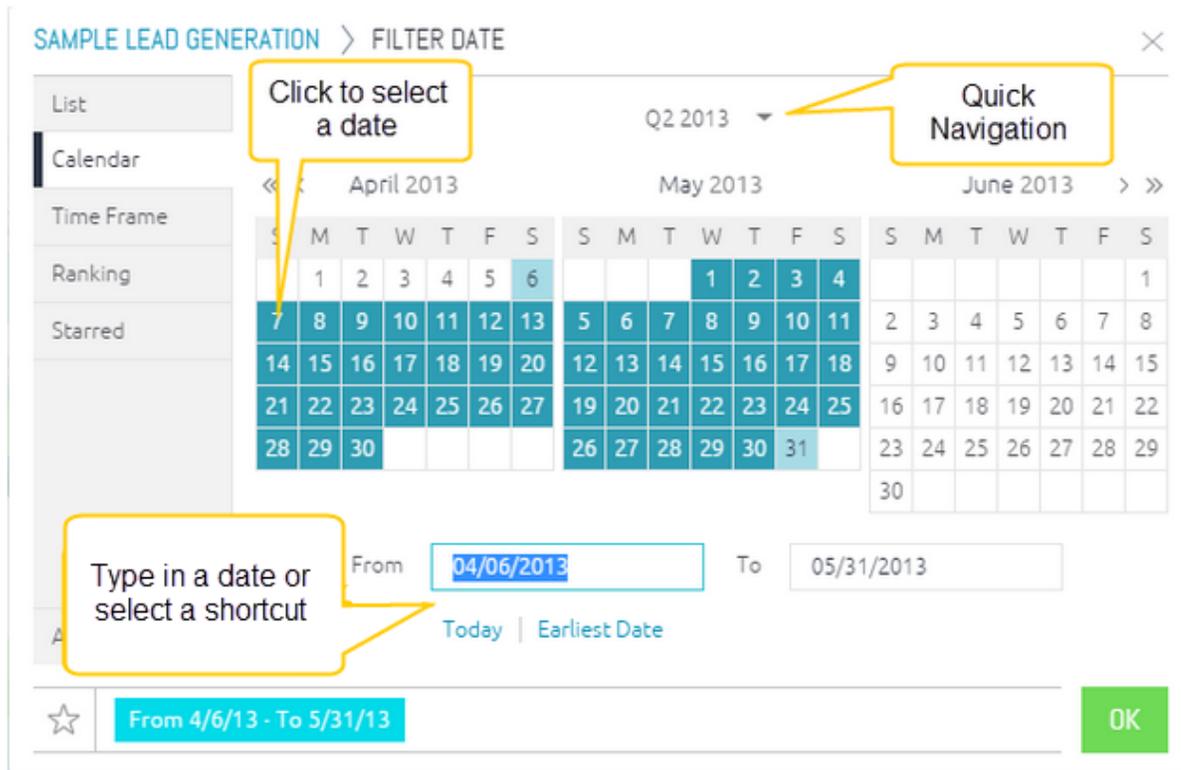
### Calendar Filter

This filter enables you to select custom date ranges from a calendar.

**To define a Calendar filter:**

1. Follow steps 1 – 3 of [Creating Dashboard Filters](#). (Select a date field)
2. Select the **Calendar** filter tab.
3. Use the quick navigation menu to jump to a point in time.

4. Select a date range using any of the following methods:
    - Click on a start and end date.
    - Type in the dates, or use a shortcut such as Earliest Date or Today.
- Note:** Earliest Date and Latest Date refer to the earliest and latest dates with data.



**Note:** Unlike other filters that you define, viewers will be able to change the type of date filter, providing more flexibility for viewing date ranges. For example, from Time Frame to Calendar, etc.

## Dynamic Time Filter

This filter enables you filter dashboards to preset dynamic time frames, such as 'Last Year', 'Last 2 Years', '2 Year ago', etc.

### To define a dynamic time filter:

1. Follow steps 1 – 3 of [Creating Dashboard Filters](#). (Select a date field)
2. Select the **Time** filter tab.
3. Use the calendar to select a date range to filter by. Use the left calendar to select the starting date of the range, and the right calendar to select the end date.

4. Click **OK** to filter the data in the Dashboard and add a date filter to the Filter panel of the Dashboard.

## Advanced Filtering Criteria

If you require advanced filtering options that are not included in the provided settings (for dashboard or widget), you can add or edit filters in the **Advanced** section.

### To add or edit filtering criteria:

1. Click **Advanced** in the left panel.
2. Add or edit filtering criteria. Click **Test** to see the results of your new/edited filter.
3. Click **OK** to add the filter.

### Example 1 – Adding an OR statement

The example below shows an OR statement being used to display multiple names. After clicking **Test**, the matches are displayed to the right of the script.

The screenshot shows a window titled "FILTER DIRECTOR1" with a close button (X) in the top right corner. On the left side, there is a vertical menu with options: "List", "Text", "Ranking", "Starred", and "Advanced". The "Advanced" option is highlighted with a yellow border. Below the menu is a "Test" button. The main area of the window is split into two panels. The left panel contains a JSON script: 

```
{
  "or": [
    {
      "contains": "martin"
    },
    {
      "contains": "alfred"
    }
  ]
}
```

 The right panel displays the results: "Showing 2 matches: Alfred Hitchcock (dir.)" and "Martin Scorsese (dir.)". At the bottom of the window, there is a star icon on the left and a green "OK" button on the right.

### Examples using JAQL Queries

## Example 2 – Filtering by Measures

This option lets you filter dimensions by measures, returning only the filtered members and associated values.

In this example, all products with a total price equal or greater than 1 million will be returned.

The query:

```
"datasource": "laptopsales",
  "metadata": [
    {
      "dim": "product"
    },
    {
      "dim": "price",
      "agg": "sum",
      "filter": {
        ">=": 1000000
      }
    }
  ]
}
```

Results include the relevant product names and price, for example:

- Product A, 1000000
- Product B, 1551246

## Example 3 – Filtering by Dimension Attributes

This option lets you filter dimensions by filtered attributes of other dimensions, returning only the filtered values.

In this example, all customers that purchased Product A and Product B will be returned.

```

{
  "datasource": "laptopsales",
  "metadata": [
    {
      "dim": "customers",
      "filter": {
        "attributes": [
          {
            "dim": "products",
            "filter": {
              "members": ["Product A",
                "Product B" ]
            }
          }
        ]
      }
    }
  ]
}

```

Note that you can add as many nested filters as needed.

## Updating on Every Change

When this option is on, any changes you make to the dashboard filters are updated immediately, and the dashboard is refreshed to reflect the changes. If you are working with large or complex data sets that result in longer refresh times, then switch off this option. Make all your changes, and click the Update button when you are ready to update your dashboard.

**Note:** This option is available for Designers only. Sisense viewers can make changes that will update automatically on every change.

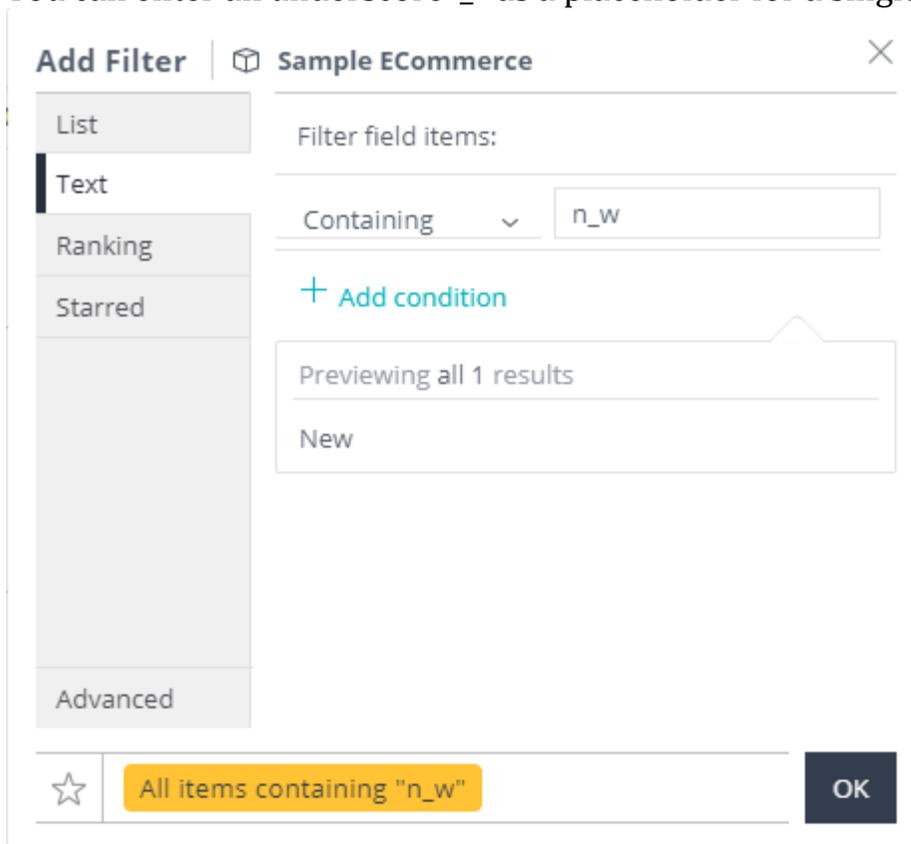
## Wildcards

To filter lists or text quickly, you can use wildcards such as '\_' or '%', which will be interpreted as literal characters or an empty string.

Sisense supports the following wildcards:

## Single Characters

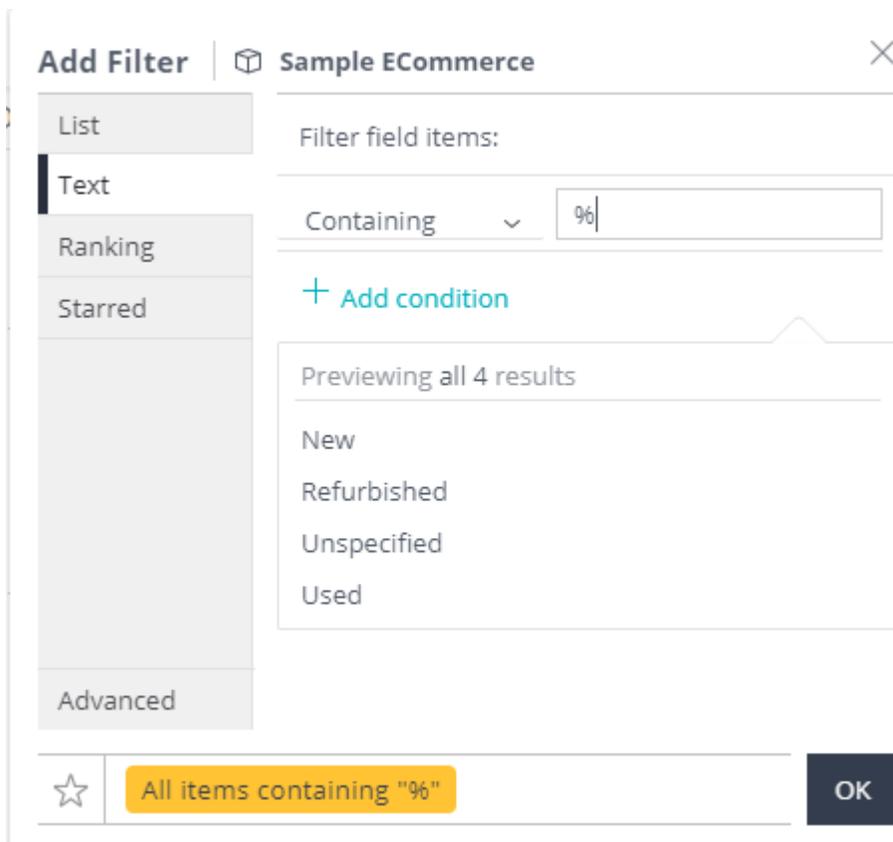
You can enter an underscore '\_' as a placeholder for a single unknown character.



The screenshot shows a dialog box titled "Add Filter" for a "Sample ECommerce" dataset. On the left, there is a vertical menu with filter types: "List", "Text" (selected), "Ranking", "Starred", and "Advanced". The main area is titled "Filter field items:" and contains a dropdown menu set to "Containing" and a text input field containing "n\_w". Below this is a "+ Add condition" button. A results preview section shows "Previewing all 1 results" and a single item named "New". At the bottom, there is a star icon, a yellow summary bar that reads "All items containing 'n\_w'", and an "OK" button.

## Multiple Characters

You can use a percent % symbol to search for an undefined length of string.



## Escape Characters

If you are searching for a string with an underscore or percent symbol, add a backslash (\) in front of the special character. Backslashes are treated as an escape character by Sisense.

For example, to locate results for the string “\_te”, use a backslash before the string:

**Edit A Filter** |  **escape2** ✕

List

**Text**

Ranking

Starred

Advanced

---

Filter field items:

Containing v

+ Add condition

Previewing all 1 results

---

my\_text

---

 All items containing "\\_te" OK

# Creating Widget Filters

This topic describes how you can use the Widget Designer to create widget filters and add widget filters to an existing field.

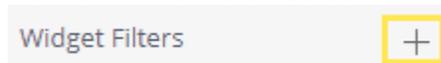
In the Widget Designer, you can define a filter for the data in a widget. or example, you can create a filter that only shows the fields of relevant countries. The filters are not visible in the dashboard, nor can they be edited directly from the dashboard.

Note about duplicate filtering: When creating a widget filter for a field that is already filtered in a dashboard filter, the widget filter will override the dashboard filter. Example: You have a dashboard filter, filtering the months field, with the values January and February selected. At the same time you create a filter for a widget, also for months, but this time March and April are selected. When both filters are applied in the dashboard, you will see data for March and April in the widget, according to the widget's filter. If your widget is filtered using measured values, then the measured value will override any other widget or dashboard filters you have for the same fields.

## Creating a Widget Filter

**To add a widget filter:**

1. Open (edit) this Widget in the Widget Designer, as described in [Adding Widgets to a Dashboard](#).
2. Click the **Filters** tab in the Widget Designer.
3. The Filters panel is displayed showing the filters defined for this Widget. Click on the **Add Filter** button to display the Data Browser, which lists the fields in the ElastiCube.



- Click on a field in the Data Browser to add it as a filter. You can define multiple filters, each for a different field. These filters have an **AND** relationship between them. Each filter that you define appears on the **Widget Filters** panel on the right enabling you to easily edit it if needed. The procedure for defining a Widget filter is similar to the procedure for defining a Dashboard filter, as described in [Creating Dashboard Filters](#). The difference is that the widget filter only affects its hosting widget, and is accessible only via the widget designer, not directly from the dashboard.

## Adding a Widget Filter to an Existing Field

Adding a new widget filter, as described above, lets you use any field as a widget filter. Sometimes you want to directly filter a field that is already part of the widget and visible in the data panel on the left.

### To add a Widget Filter to an existing field:

- Hover over the field you want to filter, and click the **FILTER** icon.

The screenshot illustrates the process of adding a widget filter to an existing field. On the left, the 'VALUES' panel shows two fields: 'Total Revenue' and 'Total Cost'. The 'Total Revenue' field has a filter icon (a downward-pointing triangle) highlighted with a yellow box. On the right, the 'FILTER TOTAL REVENUE' dialog box is open. It features a 'Values' tab selected, with a 'Keep Items' dropdown set to 'greater than' and a value of '142062'. Below this, there is a 'Ranking' section with a progress bar and a 'Starred' section. A preview window shows 'Previewing first 5 results' with a list of ranges: 19-24, 25-34, 35-44, 45-54, and 55-64. At the bottom, there is a star icon, a summary text 'All items greater than 1420629', and an 'OK' button.

# Creating Dependent Filters

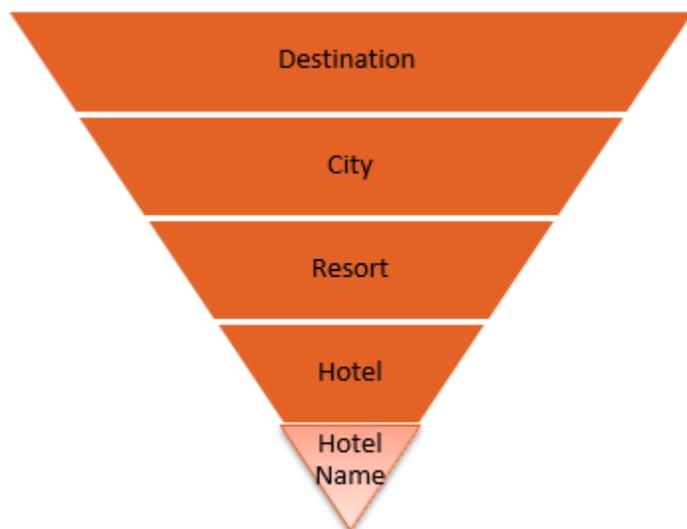
This topic describes dependent filters.

Dependent Filters combine multiple dashboard filters into a hierarchy of dependent filters that affect each other from parent, to child, to grandchild and so on. When you filter the parent filter, fewer options will be available in the child filter.

## Example Use Case

The following example shows how using a dependent filter helps you easily focus information in a dashboard.

This example describes an online travel agency with 250,000 hotels in its database. An agent wants to find a specific hotel, but does not remember its name. To narrow the results, the agent can define a dependent filter. This dependent filter is comprised of four sub-dependent filters defined with the following hierarchy:



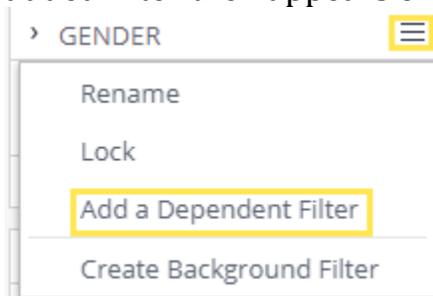
- **Destination:** Displays a list of States and Countries. In this example, the agent will select the Destination Florida.

- **City:** Because Florida was selected, only Destinations in Florida are displayed in the child filter, such as: Miami, Orlando and Tampa. In this example, the agent will select the City Orlando.
- **Resort:** Because Orlando was selected, only resorts in Orlando are displayed in the grandchild filter, such as Floridays, Buena Vista, Walt Disney World. In this example, the agent will select the Walt Disney World resort.
- **Hotel:** Because Walt Disney World was selected, only the Hotels in Walt Disney World are displayed in the great-grandchild filter.

**Note:** By default, filters are not dependent. Dependent Filters can be defined only from selection type filters.

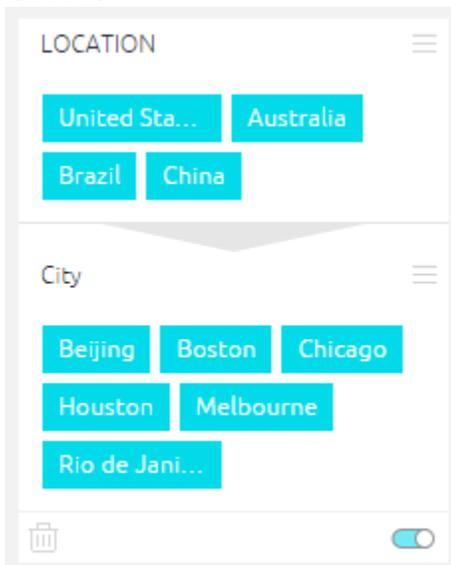
### To define a Dashboard Dependent Filter:

1. The top level (most inclusive) filter of the Dependent Filter must first be defined in the usual manner (see [Creating Dashboard Filters](#)). This newly added filter then appears on the dashboard.



2. On the parent filter, click + or click → **Add a dependent filter**.
3. Define the child filter in the usual manner (see [Creating Dashboard Filters](#)). The number of available options for selection in the child filter are less because of the selections in the parent filter. The child filter appears on the dashboard with a down arrow indicating the dependent filter and its parent

above.



4. Repeat this step for each new child filter to create as many sub-levels of Dependent Filters as you need.

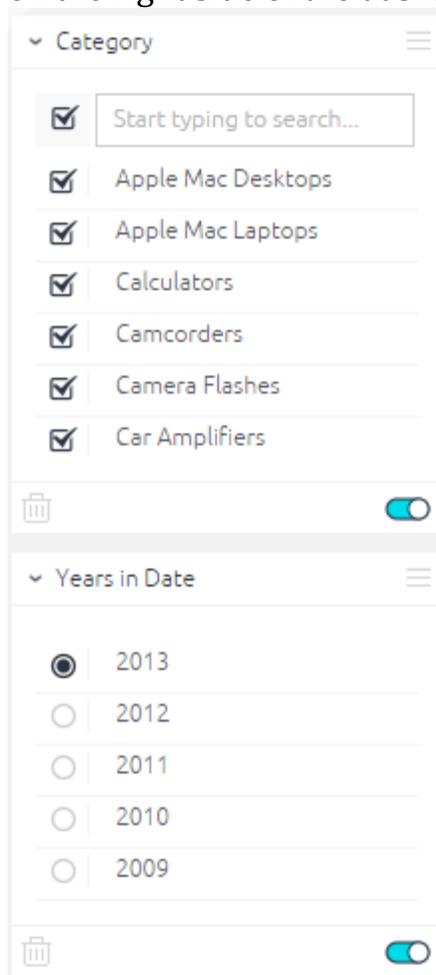
# Editing and Deleting a Dashboard Filter

This topic explains how to edit a dashboard or delete a dashboard filter.

## Editing a Dashboard Filter

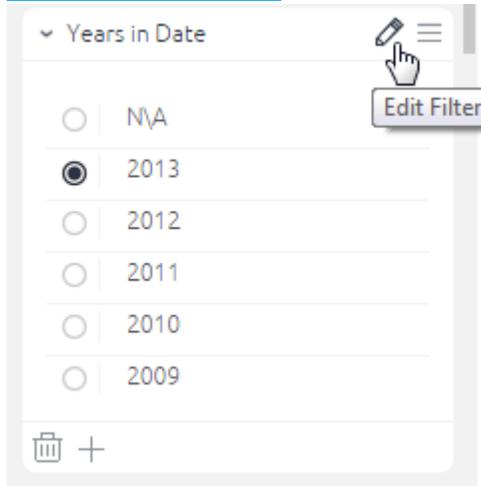
**To edit a dashboard filter:**

- Make a different selection in the filter controls displayed in the Filter panel on the right side of the dashboard, as shown below:



OR

Click on the pencil icon next to the filter name (shown above) in the dashboard to display the Filter Definition window, as described in [Creating Dashboard Filters](#).



## Deleting a Dashboard Filter

### To delete a filter:

- Click the bin  icon. To temporarily disable a widget's filter, you can toggle the filter on or off using the toggle switch. Delete the widget only if you do not need it any longer.

# Interacting with Filters as a Designer

If you are a Dashboard Designer, you can define both dashboard filters and widget filters, as follows:

- **Dashboard Filters:** This type of filter affects all the widgets in the dashboard (except the widgets that you configured as independent). Designers can create dashboard filters and Viewers can view and interact with them.
- **Widget Filters:** This type of filter affects only the specific widget it was applied to. A Designer can define which subset of data values is displayed in a widget. Viewers do not see these filters and cannot change them.

**Note:** Dashboard editing rights are assigned by Dashboard Designers and Designers.

The following topics include procedures for:

- [Creating Dashboard Filters](#)
- [Editing and Deleting a Dashboard Filter](#)
- [Creating Widget Filters](#)
- [Creating Dependent Filters](#)
- [Configuring how Filters Affect the Dashboard and Widgets](#)
- [Limiting Filters to Specific Values with Background Filters](#)
- [Locking Filters](#)
- [Restoring Filters](#)
- [Switching Filters On and Off](#)
- [Saving your Default Filters View](#)

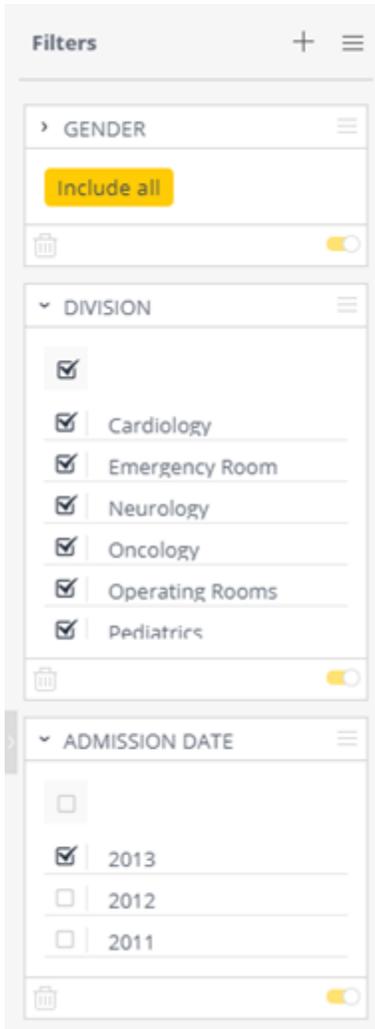
# Interacting with Filters as a Viewer

Viewers can interact with filters for analyzing data. Through the Filters pane, you can modify filters in a dashboard and change the way values are sorted in a visualization. This is useful for focusing on data displayed in dashboards that have been shared by the dashboard's owner.

As a Viewer, you can select and deselect filter visuals on a dashboard. Each time you interact with a filter, for example by selecting or entering a value, the filter is immediately applied to your dashboard. After you have configured the relevant filters, you can save it as your default filters view.

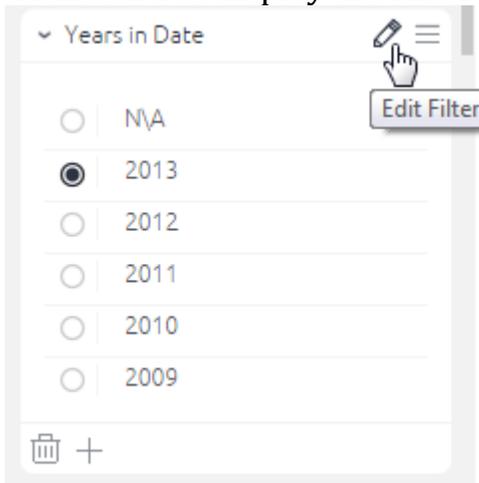
Viewers can interact with filters as follows:

- Make a different selection in the filter controls displayed in the Filter pane on the right side of the dashboard, as shown below:



OR

Click on the pencil icon next to the filter name (shown above) in the dashboard to display the Filter Definition window.



In addition, Viewers can perform the following procedures:

- [Switching Filters On and Off](#)
- [Saving your Default Filters View](#)

To create a filter you must be a [Designer](#).

[Click here](#) to learn more about filters.

# Limiting Filters to Specific Values with Background Filters

This topic describes how you can leverage background filters to limit what data is exposed to your Viewers or to reduce the amount of filtering a Viewer needs to do.

As a Designer, you can define which of the fields in a filter will be visible to Viewers. After defining which fields will be included in the filter, viewers will be able to use the updated and restricted filter to further filter the dashboard or widget using the available fields.

This can be useful in cases where Viewers only require selected data, and then they do not have to search through a long list of fields, or for security purposes in cases where you do not want to expose all the data.

**Note:** You must have access to edit a dashboard to view this option.

## **Example 1 – Simple List Filter**

A filter includes a list of 10 countries. As a designer, you can limit the filter to include only five specified countries. Viewers will see the filter as having only five countries to include or exclude in the filter.

## **Example 2 – Conditional Filter**

Another example requires having only the top 5 countries by population included in the available list of countries. In this case, Viewers will only see five countries, however, the list of countries will depend on the population sizes and any changes that may affect which countries are in the top 5 at the given time.

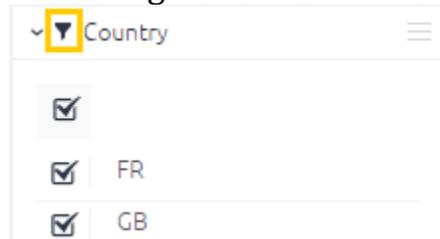
### To create a background filter:

- Create a background filter using one of the following methods:
  - In a filter without selected values, open the filter's menu and click **Create background filter**. Define your filter settings, and click **OK**.

OR

If you made changes to your filter (either by editing the filter or by selecting values directly in the filter), you can open the filter's menu, and click **Set as background filter**.

The background filter is indicated by a filter  icon.



**Note:** The icon is not visible to viewers.

### To edit a background filter:

1. In the filter's menu, click **Edit background filter**.
2. Update the filter settings, and click **OK**.

### To remove a background filter:

- In the filter's menu, click **Remove background filter**.

# Locking Filters

As a Designer, you can lock a filter, preventing Viewers from making any changes. The Viewer will be able to see the defined filter settings, but the filter will be disabled for editing.

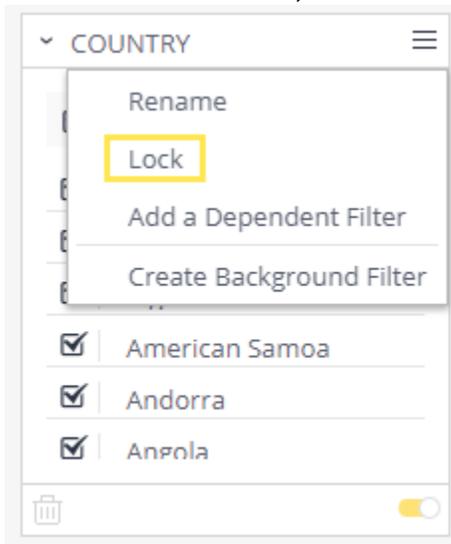
**Note:** You must have access to edit a dashboard to view this option.

## To lock a filter:

1. Click the Filter menu.
2. Click **Lock**.

A lock icon appears to indicate that the filter is locked.

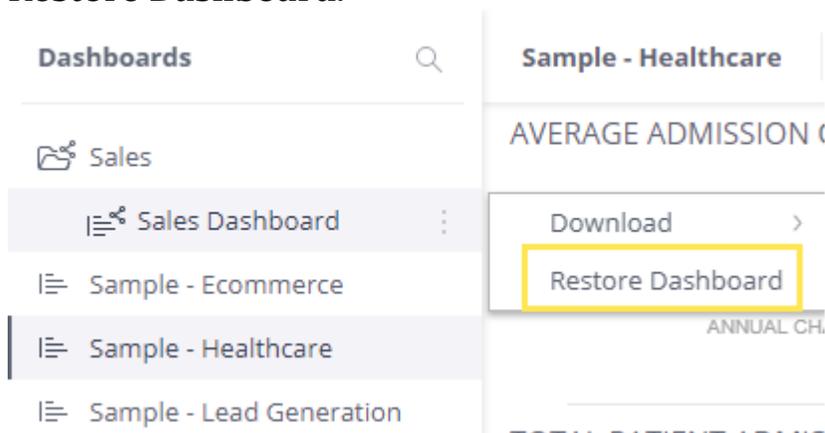
To unlock the filter, click on the menu, and select **Unlock**.



# Restoring Filters

As a Designer, if you changed the dashboard's filters, you can revert to the latest copy shared by the dashboard's owner at any time. When you restore a dashboard, local changes to a dashboard by a Viewer are overridden and the default dashboard is restored for all users.

- Under the Dashboards list, click on the menu of the dashboard and select **Restore Dashboard**.



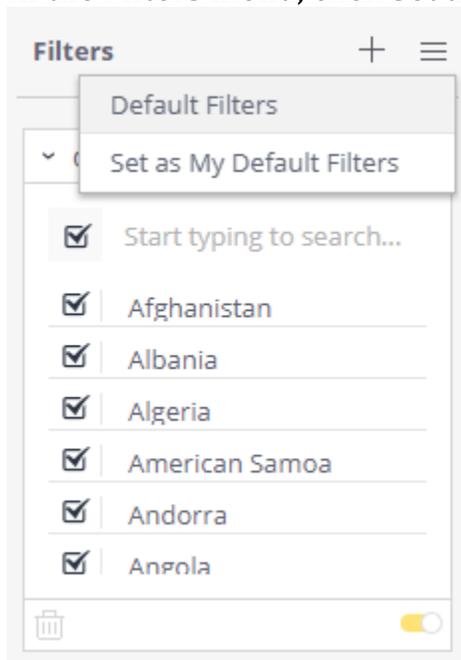
# Saving Your Default Filters View

You can save the current state of your filters and their settings at any time. The current state includes the existing filters, their configuration, and the order in which they appear in the filters panel.

After making changes to any of the above settings, you will be able to restore your filters to their previously saved state.

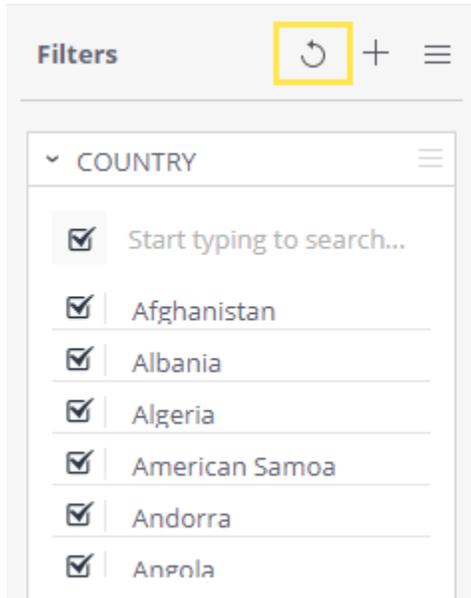
## To save your current set of filters:

- In the Filters menu, click Set as **My Default Filters**.



## To restore a saved filters set:

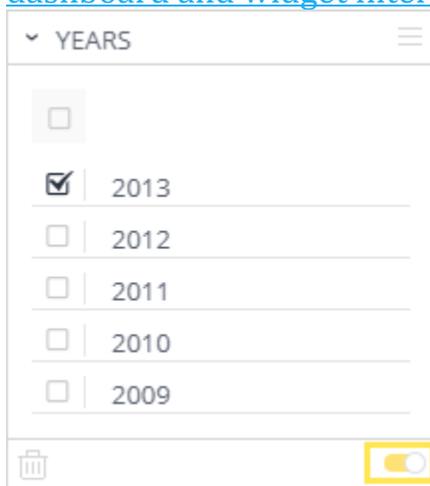
- Click the restore icon next to the Filters menu.



# Switching Filters On and Off

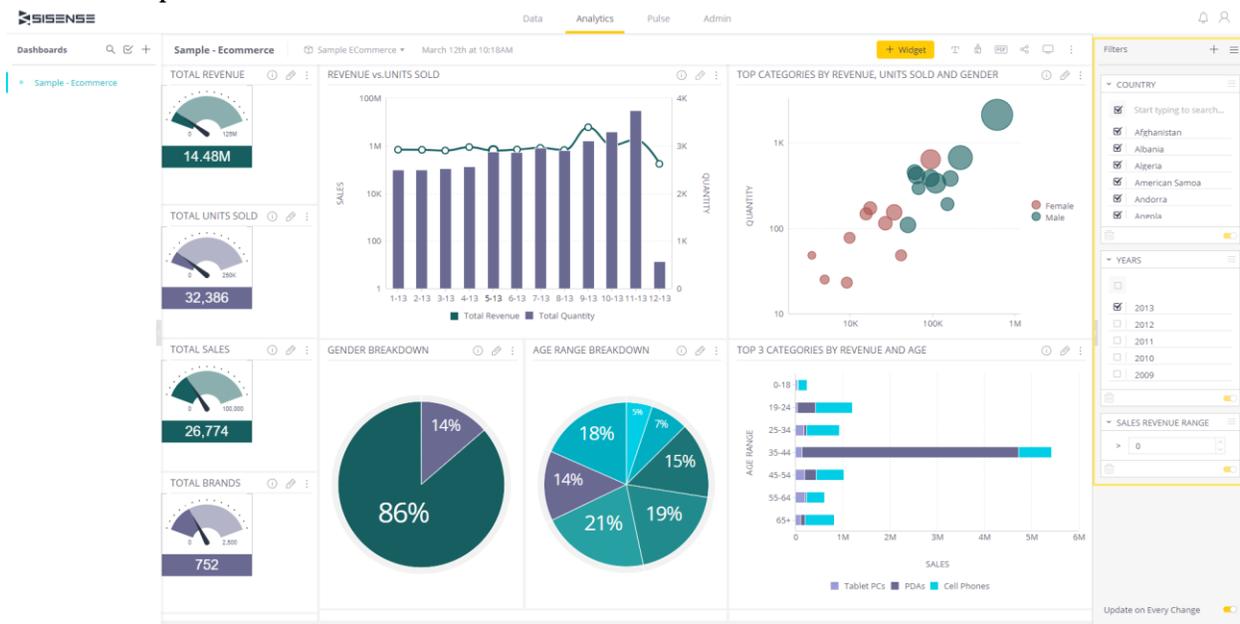
You can easily toggle filters on and off using the toggle switch. Use this option to compare states (with and without the filter), or to temporarily disable a filter, rather than deleting it.

Toggling a filter on or off in the dashboard affects the data in all the widgets, unless a widget is defined not to be affected by dashboard filters (see [how dashboard and widget filters work together](#)).



# Introduction to Filters

The Filters panel on the right of the dashboard provides data filtering options that affect the data displayed in the widgets. You can change these filters to focus on more specific data.



Your Sisense user role determines how you can interact with filters. As a Designer you can create filters that are distributed to Viewers who interact with your dashboards. See [Interacting with Filters as a Designer](#).

Viewers can interact with filters and save their changes locally. See [Interacting with Filters as a Viewer](#) for more information.

# Building Formulas with Functions

Functions are operations that perform common types of calculations, and can be used to build formulas. In this topic, you can read about four types of functions. The functions' syntax is explained and examples are provided.

## Combine Data: Aggregate Functions

Aggregations are used to perform mathematical calculations on data. Although this is an essential function Sisense offers an advantage in the ability to run multiple aggregations on several fields simultaneously – this makes it easy to summarize data based on multiple factors.

### *Aggregate Syntax:*

Function(Numeric Field)

### *Parameters*

Function: Sum(), Avg(), Count(), DupCount(), Max(), Min()

Numeric Field: A numeric field or formula.

An example **simple aggregation** is a sales manager who wants to calculate the average sales revenue for each sales rep. They can create a pivot table which shows the sales rep and their average sales revenue.

A more complex example is a **multi-pass aggregation** (or grouping) that is an aggregation that performs multiple calculations simultaneously. Following our first example let's assume the sales manager wants to also see average sales per day for each sales rep. Instead of having to add an additional column for day in the pivot table the manager can create a multi-pass aggregation that first

performs a sum of sales per day and then averages the results for each rep. This requires two fields – a day from a date field and the revenue field, as well as two aggregations, sum of sales and average. This result is the sales manager does not need to add a column for days in the pivot.

Data Set			Results Example 1: Avg(Revenue)	
Day	Sales Rep	Revenue	Sales Rep	Avg Revenue
1 Jan	John	10	John	$(10+30+10)/3 = 16.67$
1 Jan	Jane	20	Jane	$(20+40+5)/3 = 21.67$
1 Jan	Jane	40		
2 Feb	John	30		
2 Feb	John	10		
2 Feb	Jane	5		

Results Example 2: Avg(Days, Sum(Revenue))	
Sales Rep	Avg Daily Revenue
John	$(10+(30+10))/2 = 25$
Jane	$((20+40)+5) / 2 = 32.5$

The above example as shown in the formula editor.

Sales Rep	Average Revenue	Avg(([Days in Date]),([Total Revenue]))
Jane	21.67	32.5
John	16.67	25

COUNTRIES > EDIT FORMULA ×

AVG(([Days in Date]),([Total Revenue])) ★

## Summarize Data: Statistical Functions

Descriptive statistics provide meaningful summaries of data and help make more informed decisions. This is particularly important for large data sets where descriptive statistics can help to focus analysis.

*Statistical Syntax:* Function(Numeric Field)

*Parameters*

Function: Median(), Mode(), Largest(), Stdev(), Stdevp(), Varp(), Var(), Quartile(), Percentile()  
Numeric Field: A numeric field or formula.

An example of statistical functions is a marketing team that has a large data set on leads generated from various channels and want to understand where to focus their budget. Descriptive statistics can be used to summarize valuable insight about each channel such as the central tendency or median leads generated along with standard deviations to assess typical lead volume.

*Statistical Syntax:* Function(Numeric Field)

*Parameters*

Function: Median(), Mode(), Largest(), Stdev(), Stdevp(), Varp(), Var(), Quartile(), Percentile()  
Numeric Field: A numeric field or formula.

## Accumulate Data: Running Total and Average

Often to measure performance data must be viewed in a continuous and accumulative format over extended periods such as years, quarters or months. Sisense provides functions to create running totals and averages over standard or custom time periods.

*Accumulate Syntax:* Function(Numeric Field)

*Parameters*

Function: YTD Avg(), QTD Avg(), MTD Avg(), YTD Sum(), QTD Sum(), MTD Sum(), RPSum()/ RPAvg()  
Numeric Field: A numeric field or formula.

For example, a support team has a goal to reduce the average monthly cost to resolve open issues. A Year to Date Average can be used to track progress towards reducing the average cost of support.

<i>Data</i>		<i>Results Example 1: YTD Avg(Cost)</i>	
<i>Month</i>	<i>Cost</i>	<i>Month</i>	<i>YTD Avg Cost</i>
Jan	30	Jan	30
Feb	35	Feb	32.5
Mar	40	Mar	35
Apr	20	Apr	31.25
May	25	May	30
Jun	10	Jun	26.67

## Change over Time: Time Functions

Time is critical for business decisions, Sisense time functions make it easy to compare outcomes at different points in time, determine growth rates and calculate the time differences. Functions can be set for common time periods such as year, month or day as well as for custom periods.

*Time Syntax: Function(Numeric Field) \**

*Parameters*

Function:

PastYear(), PastQuarter(), PastMonth(), Growth(), GrowthRate(), Contribution(), YDiff(), QDiff(), MDiff(), DDiff(), HDiff(), MnDiff(), SDiff(), Next(), Prev(), DiffPastYear(), DiffPastMonth(), DiffPastQuarter(), DiffPastDay()

Numeric Field: A numeric field or formula.

*\* To calculate time functions require a time filter applied on the dashboard or widget.*

For example, an executive team wants to compare growth in revenue to the same period in the previous year. A Difference in Past Year function can be used to compare past values based on the current month to the same month in the previous year.

*Data*

<i>Date</i>	<i>Revenue</i>
01-2013	5
02-2013	10
03-2013	15
01-2014	25
02-2014	5
03-2014	10

*Results Example 1: DiffPastYear(Revenue)*

<i>Month</i>	<i>Difference to Past Year</i>
01-2013	
02-2013	
03-2013	
01-2014	$25-5 = 20$
02-2014	$5-10 = -5$
03-2014	$10-15 = -5$

# Conditional Statements

Conditional statements are formulas that enable you to create additional categories in your data set. When the condition is met, the new category is added to your visualization without modifying the data source itself.

For example, if you want to calculate only certain values according to conditions you define, you can use a conditional statement to include only the values of interest.

Sisense provides several conditional statements that enable you to evaluate criteria and display only the relevant results in your dashboard.

**Note:** Conditional states work on measures and aggregations only and return numeric values.

## IF

```
IF (Boolean_expression, true value, false value)
```

Evaluates a boolean expression and returns the first value when evaluated as true, and returns the second value when evaluated as false.

For example:

```
(if(Sum(Sales) < 10, 0, Sum(Sales) )
```

Nested IF statements are also supported.

For example:

```
if (Sum(Sales) < 100 , 1 , (if (Sum(Sales) < 1000 , 2, 3) )
```

# Case

```
CASE
WHEN Boolean_expression THEN result_expression
[ ... ]
[ ELSE else_result_expression ]
END
```

The CASE function evaluates boolean expressions, and when evaluated as true, returns the corresponding result\_expression. If no match is found, the else\_result\_expression is returned. If there is no default returned and no values match, then Null is returned.

For example:

```
CASE
WHEN Sum(Sales) < 100 THEN 1
WHEN Sum(Sales) < 1000 THEN 2
ELSE 3
END
```

# Creating and Editing a Formula

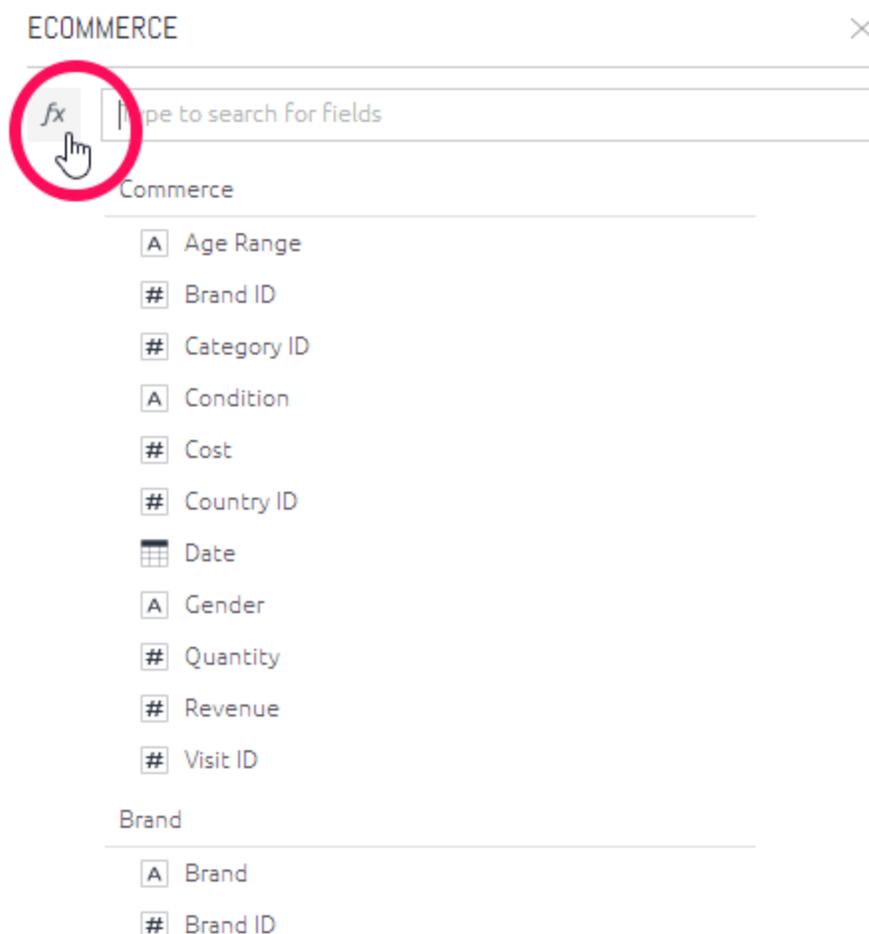
The Data Browser enables you to define formulas (freeform expressions) that define the values and filters of a widget.

A rich variety of functions are provided for you to use in the formula that you define.

## To define a formula:

1. Open the formula editor in the Data Browser:
  - For a new widget, click Select Data, and then  .

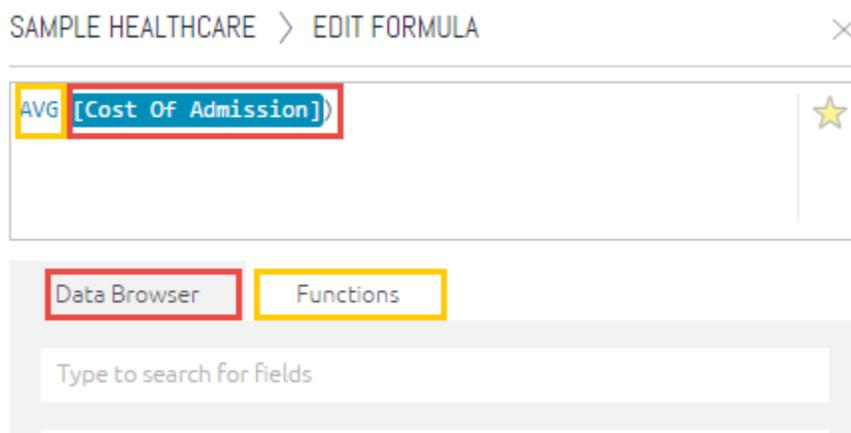
- For an existing widget, click on the edit formula button .



The Data Browser then changes to display the Formula Editor, which has two tabs: Data Browser and Functions.

- The Data Browser tab provides fields to choose from.
  - The Functions tab lists the functions that you can include in your formula by selecting them. You can read a description of each function in a tooltip by hovering over it.
2. Define the formula as follows:
- From the **Data Browser** tab, select one or more fields.

- From the **Functions** tab, select the required functions.



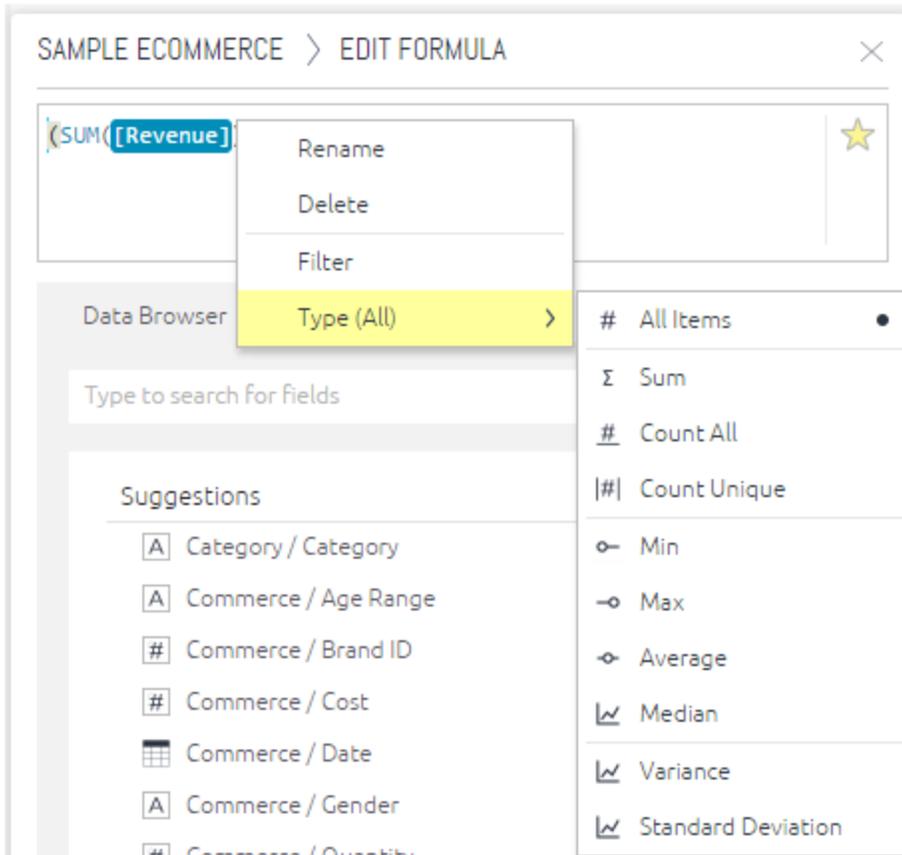
3. Type in the required parts of the formula. To see examples, see [Formulas Based on Criteria and Conditions](#), and [Functions to Build Formulas](#).
4. Click **OK**.

### To edit a formula:

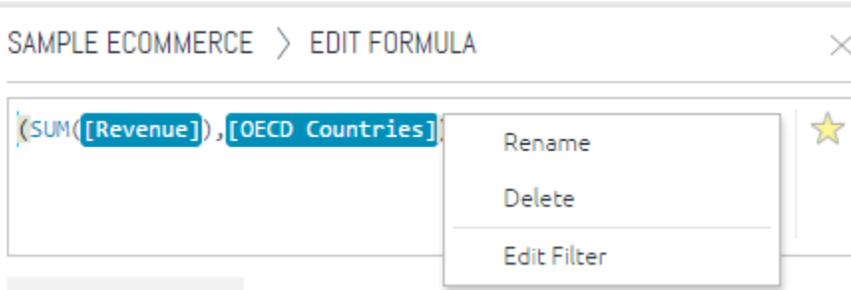
You can easily edit formulas using the right-click options. They include:

- **Rename:** Rename the formula, for example, give a name that represents a real-life task or expected result from the formula, or include in the name filters that you have added to the formula.
- **Filter:** Add filters to the formula.
- **Type:** Change the default aggregation method, for example, from Sum to Average.

The following image shows the right-click options.



The following image shows right-click options for a filtered formula. Read more [below](#).



# Creating Formulas Based on Criteria and Conditions (Filters)

Often formulas must take into account specific criteria. To do this Sisense provides a feature called **Measured Value**, which like the SUMIF function in Excel, only performs a calculation when the values meet a set of criteria. Criteria for Measured Values may be based on any logical operators in a filter.

*Measure Value Syntax:*

(Measure, Scope1 , Scope2...ScopeN)

*Parameters*

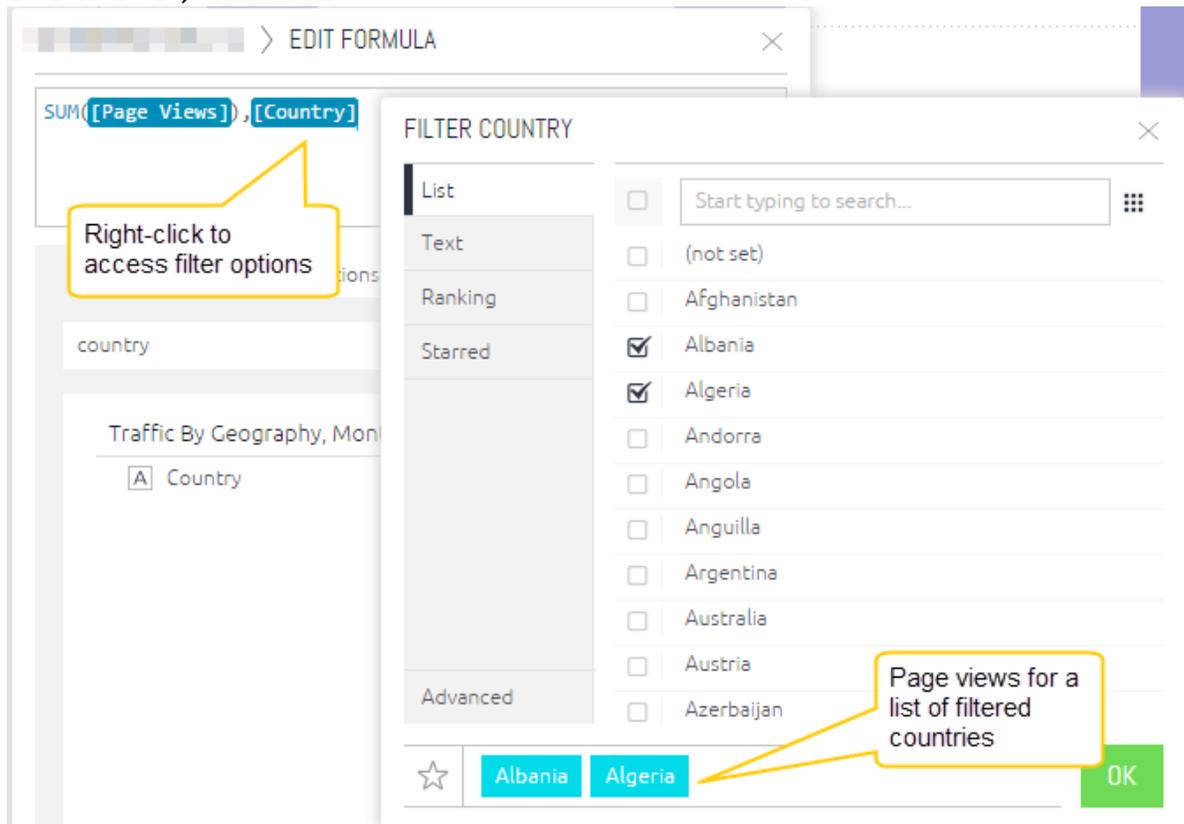
Measure: A field measure or formula.

Scope: A filter including *Value, Text, List, Ranking and Time filters*.

## To filter the formula:

1. In the Data Browser, create your formula from the Data Browser and Functions, as explained in [Formula Editor](#).
2. Add the field (criteria) by which you want to filter the formula. Right-click the field and select **Filter**.

3. You can then filter the formula by listed items, text options, ranking, etc. When done, click **OK**.



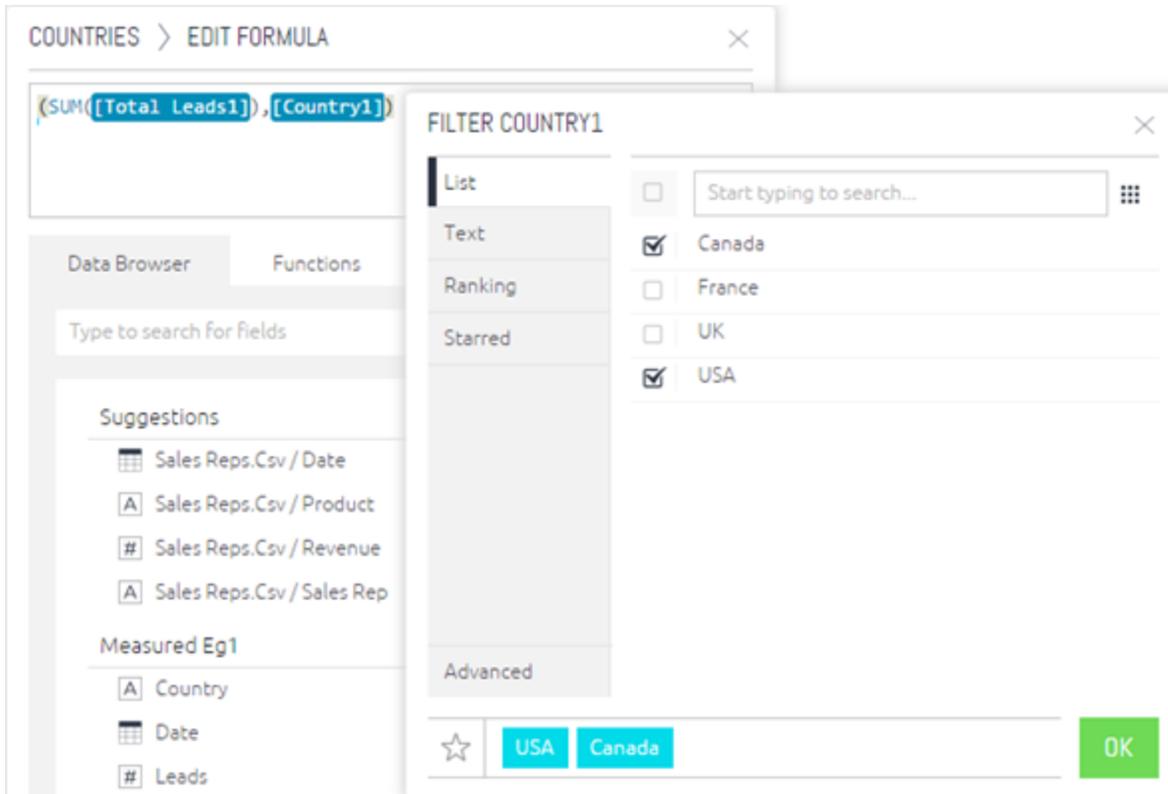
A simple example of Measured Value is the use of a list filter. A marketing team may need to count leads generated for a specific region such as North America. Even if leads come from many different countries, the measured value calculates leads generated only when the lead originates from the United States or Canada.

Country	# Leads
USA	15
France	20
Canada	5
UK	30

*Example 1: Measured Value with List Filter*  
 (Sum(Leads), (List Filter: Country = United States, Canada))

*Result*  
 15+5 = 20

The above example as defined in the Formula Editor.



A more sophisticated case is the **use of a ranking filter**, for example a sales team may want to track the contribution of best-selling products to total revenue. However, what constitutes a popular product may change over time. A measured value can be created for sales which includes a condition that only shows sales for the top products for any month. This simultaneously filters the data but also takes into account changes in what classifies as a top product over time.

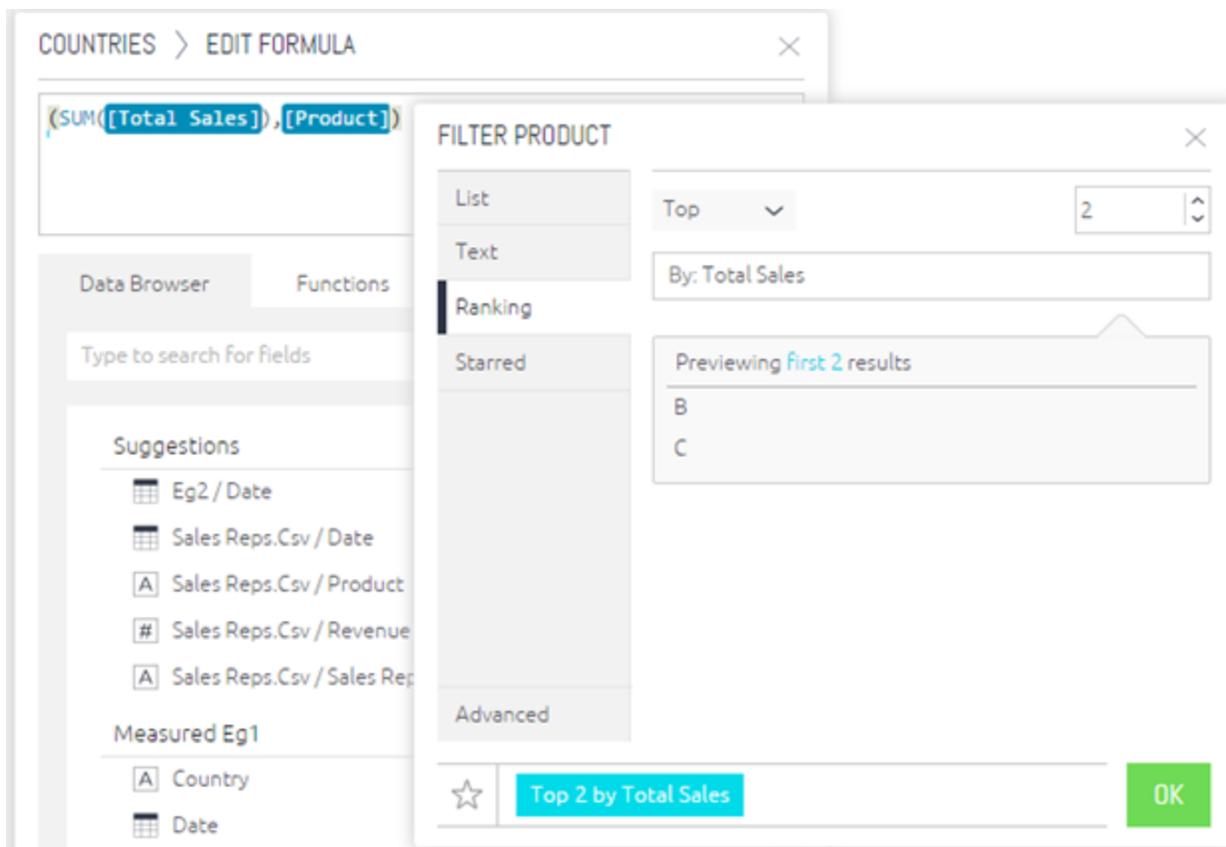
Month	Product	Sales
Jan	A	10
Jan	B	5
Jan	C	15
Feb	A	1
Feb	B	25
Feb	C	5

*Example 2: Measured Value with Top Ranking Filter  
(Sum(Sales), (Top Ranking Filter: Top 2 Products by Sales))*

*Results*

Month	Total Sales Top 2 Products
Jan	(A/C) 10 + 15 = 25
Feb	(B/C) 25 + 5 = 30

The above example as defined in the Formula Editor.



Measured Values are a powerful feature to take into account business logic and quickly perform calculations only when a specific set of criteria is met.

**Note:** If your widget is filtered using measured values, then the measured value will override any other widget or dashboard filters you have for the same fields.

## Calculating Contributions Using the ALL Function

The All() function returns the total amount for a dimension, and can be used for various use cases. In the following example, we will use the All function to calculate how much each country contributed towards the total cost of a campaign.

Our final widget includes the following information:

Country	Total Cost	Total Cost per Countries	Contribution
Australia	9,643.09	283,755.53	3.40%
Brazil	9,638.80	283,755.53	3.40%
China	9,607.4	283,755.53	3.39%
England	19,492.8	283,755.53	6.87%
France	19,440.59	283,755.53	6.85%
Germany	19,377.61	283,755.53	6.83%
Greece	9,842.84	283,755.53	3.47%
India	9,603.16	283,755.53	3.38%
Kazakhstan	9,899.33	283,755.53	3.49%
Nigeria	9,764.44	283,755.53	3.44%
Norway	9,681.69	283,755.53	3.41%
Portugal	9,919.33	283,755.53	3.50%
South Africa	19,211.15	283,755.53	6.77%
Ukraine	9,693.03	283,755.53	3.42%
United States	99,241.60	283,755.53	34.97%
Vietnam	9,698.67	283,755.53	3.42%
Grand Total	283,755.53	283,755.53	100%

**Step 1:** The second column above represents a formula that sums up the total cost for all countries and does not represent the breakdown per country. The formula includes the calculation (total cost) followed by the all function (filter), followed by the dimension (country) in parenthesis. It looks like this:

```
sum([Total Cost], all([Country]))
```

We can save (star) the above formula and call it Total cost for Countries, which will be used in the next step.

**Step 2:** We can now use the above formula in another formula to calculate the contribution, like this:

```
sum([Cost])/[Total Cost for Countries]
```

The result is the third column above (plus formatting the results as percentages).

# Function Reference

This page contains a list of all the functions you can use in Sisense's formula editor.

## Statistical Functions

### Average

`Avg(<numeric Field>)`

Calculates the mean average of the given values.

For example – `AVG(Score)` will calculate the mean average of the given scores.

`Avg(<group by field>, <aggregation>)`

Calculates the average of the given aggregation grouped by another field.

For example – `Avg( Product, Total Sales)` will calculate the average of the total sales per product.

### Contribution

`Contribution(<numeric field>)`

Calculates the percentage of total.

For example – `Contribution( Total Sales )` will calculate the percentage of total sales per group (for example per day or per product) out of total sales (for all days or all products).

### Correlation

`CORREL(<Numeric Field a>, <Numeric Field b>)`

Returns the correlation coefficient of two numeric fields.

For example – CORREL(Revenue, Cost) will return the correlation between revenue and cost.

```
CORREL(<group by field>, <aggregation a>, <aggregation b>)
```

Returns the correlation coefficient of two fields aggregations grouped by another field.

For example – CORREL(Products, AVG(Revenue), AVG(Cost)) will return the correlation between the average of revenue and cost per product.

## Count

```
Count(<Numeric Field>)
```

Counts the number of unique values withing the given values.

## Count All

```
DupCount(<Numeric Field>)
```

Returns the actual item count of the given list of items, including duplicates.

## Covariance (Population)

```
COVARP(<Numeric Field a>, <Numeric Field b>)
```

Returns the population covariance of <Numeric Field a> and <Numeric Field b>.

For example – COVARP(Revenue, Cost) will return the population covariance of revenue and cost.

```
COVARP(<group by field>, <aggregation a>, <aggregation b>)
```

Returns the population covariance of two fields aggregations grouped by another field.

For example – COVARP(Products, AVG(Revenue), AVG(Cost)) will return the population covariance of the average revenue and the average cost per product.

## Covariance (Sample)

```
COVAR(<Numeric Field a>, <Numeric Field b>)
```

Returns the sample covariance of <Numeric Field a> and <Numeric Field b>.

For example – COVAR(Revenue, Cost) will return the sample covariance of revenue and cost.

```
COVAR(<group by field>, <aggregation a>, <aggregation b>)
```

Returns the sample covariance of two fields aggregations grouped by another field.

For example – COVAR(Products, AVG(Revenue), AVG(Cost)) will return the sample covariance of the average revenue and the average cost per product.

## Exponential Distribution

```
EXPONDIST(<numeric value>, <lambda>, <Cumulative  
(true/false)>)
```

Returns the exponential distribution for a given value and a supplied distribution parameter lambda. Cumulative: TRUE = Cumulative distribution function, FALSE = Probability density function. For example – EXPONDIST( Count(Leads), 2, False ) will return the exponential distribution density of the number of leads per country where lambda is 2.

## Intercept

```
INTERCEPT(<field>, <numeric value>)
```

Returns the intercept of the linear regression line through a supplied series of x- and y- values.

For example – INTERCEPT(Date.Quarter, Total Sales) will return the intercept of the regression line that represents the trend over quarter of the sum of sales.

## Largest

LARGEST(<Numeric Field>, <k>)

Returns the k-th largest value in a field.

## Maximum

Max(<Numeric Field>)

Returns the maximum value among the given values.

## Median

MEDIAN( <Numeric Field> )

Calculates the median of the given values. The median of a set of data is the middlemost number in the set. The median is also the number that is halfway into the set.

## Minimum

Min(<Numeric Field>)

Returns the minimum value among the given values.

## Mode

MODE(<Numeric Field>)

Returns the most frequently occurring value from the column.

## Normal Distribution

`NORMDIST(<Numeric Field>, <Mean>, <Standard Deviation>, <Cumulative (true/false)>)`

Returns the standard normal distribution for a given value, a supplied distribution mean and standard deviation. Cumulative: TRUE = Cumulative Normal Distribution Function, FALSE = Normal Probability Density Function. For example – `NORMDIST(Score, ( Mean(Score), All(Score)), ( STDEV(Score), All(Score) ), False )` will return the normal probability density of a given score.

## Percentile

`PERCENTILE(<Numeric Field>, <k>)`

Returns the k-th percentile value from the given field.  
k is any number between 0..1 (inclusive).

## Possion Distribution

`POISSONDIST( <numeric value>, <mean>, <Cumulative (true/false)>)`

Returns the poisson distribution for a given value and a supplied distribution mean. Cumulative: TRUE = Cumulative distribution function, FALSE = Probability mass function.

For example – `POISSONDIST( Score, ( Mean(Score), All(Score) ), ( STDEV(Score), All(Score) ), False )` will return the poisson probability density of a given number of sales

## Quartile

`QUARTILE(<Numeric Field>, <k>)`

Returns the k-th quartile for the given field.

- k = 0 returns the Minimum value
- k = 1 returns the first quartile (25th percentile)
- k = 2 returns the Median value (50th percentile)
- k = 3 returns the third quartile (75th percentile)
- k = 4 returns the Maximum value

## Rank

`RANK(<numeric value>, [DESC/ASC], [Rank Type], [<group by field 1>, ... , <group by field n>])`

Returns the rank of a value in a list of values.

[DESC/ASC] – Optional. By default sort order is descending.

[Rank Type] – Optional. By default the type is standard competition ranking (“1224” ranking). Support also modified competition ranking (“1334” ranking), dense ranking (“1223” ranking) and ordinal ranking (“1234” ranking).

[<Group by field 1>, ... , <Group by field n>] – Optional. Rank partitions fields.

For example – `RANK(Total Cost, “ASC”, “1224”, Product, Years)` will return the rank of the total annual cost per each product were sorted in ascending order.

## Running Sum (RSUM)

`RSUM ( <numeric value> ), RSUM ( <numeric value> , <continuous> )`

Returns the running total of the measure by the defined dimension according to the current sorting order in the widget.

By default, RSUM accumulates a measure by the sorting order of the dimension.

To accumulate by another order, the relevant measure should be added as an additional column and sorted.

<continuous> is a boolean value that that accumulates the sum continuously when there are two or more dimensions. The default value is False.

**Note:** Filtering the RSUM column by Values, filters the dimensions and recalculates the RSUM from the first filtered value.

## Skewness (Population)

`SKEWP(<numeric value>)`

Returns the skewness of the distribution of a given value in the population.

For example – `SKEWP(Revenue)` will return the skewness of the distribution of revenue in the population.

## Skewness (Sample)

`SKEW(<numeric value>)`

Returns the skewness of the distribution of a given value.

For example – `SKEW(Revenue)` will return the skewness of the distribution of revenue.

## Slope

`SLOPE(<field>, <numeric value>)`

Returns the slope of the linear regression line through a supplied series of x- and y- values.

For example – `SLOPE(Date.Quarter, Total Sales)` will return the slope of the regression line that represent the trend over quarter of the sum of sales.

## Standard Deviation (Population)

`STDEVP( <Numeric Value> )`

Returns the Standard Deviation of the given values (Population). Standard deviation is the square root of the average squared deviation from the mean. The

standard deviation of a population gives researchers the amount of dispersion of data for an entire population of survey respondents.

## Standard Deviation (Sample)

```
STDEV( <Numeric Value> )
```

Returns the Standard Deviation of the given values (Sample). Standard deviation is the square root of the average squared deviation from the mean. A standard deviation of a sample estimates the amount of dispersion in a given data set, based on a random sample.

## T Distribution

```
TDIST( <numeric value x>, <degrees_freedom>, <Cumulative  
(true/false)> )
```

Returns the student's T-distribution for a given value and a supplied number of degrees of freedom (must be  $\geq 1$ ). Cumulative: TRUE = Cumulative Distribution Function, FALSE = Probability Density Function.

For example – TDIST( Score, 3, TRUE ) will return the student's T-distribution of a given score, with 3 degrees of freedom.

## Variance (Population)

```
VARP( <Numeric Value> )
```

Returns the Variance of the given values (Population). Variance (Sample) is the average squared deviation from the mean, based on an entire population of survey respondents.

## Variance (Sample)

```
VAR( <Numeric Value> )
```

Returns the Variance of the given values (Sample). Variance (Sample) is the average squared deviation from the mean, based on a random sample of the population.

## Mathematical Functions

### Absolute

`Abs(<Numeric value>)`

Returns the absolute value of the given value.

For example – `ABS(Cost)`, where the absolute result for the value '2' or '-2' is '2'.

### Acos

`ACOS(<numeric value>)`

Returns the angle, in radians, whose cosine is the given numeric expression. Also referred to as arccosine.

For example – `ACOS(Total Revenue)` will return the angle, in radians, whose cosine is the given total revenue.

### Asin

`ASIN(<numeric value>)`

Returns the angle, in radians, whose sine is the given numeric expression. Also referred to as arcsine.

For example – `ASIN(Total Revenue)` will return the angle, in radians, whose sine is the given total revenue.

### Atan

`ATAN(<numeric value>)`

Returns the angle in radians whose tangent is the given numeric expression. Also referred to as arctangent.

For example – ATAN(Total Revenue) will return the angle in radians whose tangent is the given total revenue.

## Ceiling

CEILING(<numeric value>)

Returns number rounded up, away from zero, to the nearest multiple of significance.

For example – CEILING(Cost), where the result of '83.2' rounded up is '84'.

## Cos

COS(<numeric value>)

Returns the trigonometric cosine of the given angle (in radians).

For example – COS(Average Angle) will return the trigonometric cosine of the average angle.

## Cosh

COSH(<numeric value>)

Returns the hyperbolic cosine of the given value.

For example – COSH(Total Revenue) will return the hyperbolic cosine of the total revenue.

## Cot

COT(<numeric value>)

Returns the trigonometric cotangent of the given angle (in radians).

For example – COT(Average Angle) will return the trigonometric cotangent of the average angle.

## Exp

EXP(<numeric value>)

Returns the exponential value of the given value.

For example – EXP(Sales) will return the exponential value of sales.

## Floor

FLOOR(<numeric value>)

Returns number rounded down, toward zero, to the nearest multiple of '1'.

For example – FLOOR(Revenue), where the result of '88.6' rounded down is '88'.

## Ln

LN(<numeric value>)

Returns the base-e logarithm of the given value.

For example – LN(Cost) will return the base-e logarithm of cost.

## Log10

LOG10(<numeric value>)

Returns the base-10 logarithm of the given value.

For example – LOG10(Revenue) will return the base-10 logarithm of revenue.

## Mod

MOD(<numeric value>, divisor)

Returns the remainder after a number is divided by a divisor.

For example – MOD(Cost, 10), where the remainder of '255' divided by '10' is '5'.

## Power

`Power(value, power)`

Returns the results of the given value raised to a supplied power.

For example – `POWER(Revenue, 2)` will return revenue raised by the power of 2.

## Quotient

`QUOTIENT(<numeric value>, divisor)`

Returns the integer portion of a division.

For example – `QUOTIENT(Cost, 2)`, where the integer portion of '5' divided by '2' is '2'.

## Round

`ROUND(<numeric value>, num_digits)`

Returns number rounded to a specified number of digits.

For example – `ROUND(Revenue, 2)` will return the revenue rounded to two decimal places.

## Sin

`SIN(<numeric value>)`

Returns the trigonometric sine of the given angle (in radians).

For example – `SIN(Average Angle)` will return the trigonometric sine of the average angle.

## Sinh

`SINH(<numeric value>)`

Returns the hyperbolic sine of the given value.

For example –  $\text{SINH}(\text{Total Revenue})$  will return the hyperbolic sine of the total revenue.

## Square root

$\text{SQRT}(\langle \text{Numeric value} \rangle)$

Returns the square root of the given value.

For example –  $\text{SQRT}(\text{Cost})$  will return the square root of cost.

## Sum

$\text{Sum}(\langle \text{Numeric Field} \rangle)$

Calculates the total of the given values.

## Tan

$\text{TAN}(\langle \text{numeric value} \rangle)$

Returns the trigonometric tangent of the given angle (in radians).

For example –  $\text{TAN}(\text{Average Angle})$  will return the trigonometric tangent of the average angle.

## Tanh

$\text{TANH}(\langle \text{numeric value} \rangle)$

Returns the hyperbolic tangent of the given value.

For example –  $\text{TANH}(\text{Total Revenue})$  will return the hyperbolic tangent of the total revenue.

# Time Related Functions

## Day Difference

`DDiff( <Start Time>, <End Time> )`

Returns the difference between <Start Time> and <End Time> in days.

## Growth

`Growth( <Numeric Value> )`

Calculates growth over time. The time dimension to be used is determined by the time resolution in the widget/dashboard.

Formula: (current value - compared value) / compared value.

For example:

- If this month your value is 12, and last month it was 10, your Growth for this month is 20% (0.2).  
Calculation:  $(12 - 10) / 10 = 0.2$
- If this year your value is 80, and last year it was 100, your Growth for this year is -20% (-0.2).  
Calculation:  $(80 - 100) / 100 = -0.2$

## Growth Rate

`GrowthRate( <Numeric Value> )`

Calculates growth rate over time. The time dimension to be used is determined by the time resolution in the widget/dashboard.

For example:

- If this month your value is 12, and last month it was 10, your Growth Rate for this month is  $12/10 = 120\%$  (1.2).  
Calculation:  $12 / 10 = 1.2$
- If this year your value is 80, and last year it was 100, your Growth for this year is  $80/100 = 80\%$  (0.8).  
Calculation:  $80 / 100 = 0.8$

## Hour Difference

HDiff( <Start Time>, <End Time> )

Returns the difference between <Start Time> and <End Time> in hours.

## Prev

Prev( <Time Field> [, <N>] )

Returns the Time period Member in <Time Field> which is N periods back from the current Member. This function only works as a scope function and not by itself.

For example – This formula will return the numeric value 2 months ago:

(<Numeric Value>, Prev(<Month Field>, 2))

## Minute Difference

MnDiff( <Start Time>, <End Time> )

Returns the difference between <Start Time> and <End Time> in minutes.

## Month Difference

MDiff( <Start Time>, <End Time> )

Returns the difference between <Start Time> and <End Time> in months.

Returns whole numbers.

## Month to Date Average

MTDAvg( <Numeric Value> )

Returns the running average starting from the beginning of the month up to the current time period member.

The time dimension to be used is determined by the time resolution in the widget/dashboard.

Returns 0 if the active time resolution is quarters or years.

## Month to Date Sum

```
MTDSum( <Numeric Value> )
```

Returns the running total starting from the beginning of the month up to the current time period member.

The time dimension to be used is determined by the time resolution in the widget/dashboard.

Returns 0 if the active time resolution is quarters or years.

## Next

```
Next( <Time Field> [, <N>] )
```

Returns the Time period Member in <Time Field> which is N periods after the current Member. This function only works as a scope function and not by itself.

For example – This formula will return the numeric value 2 months ahead of now:

```
(<Numeric Value>, Next(<Month Field>, 2)
```

## Now

```
Now(<Day from Date field>)
```

Returns the value for the current time period. Supports day, month, quarter or year. The Now function receives a date dimension and its level and returns all the members in that dimension which match the current query execution time.

**Note:** This function only works as a scope function and not by itself.

The following example will return the value for the current day.

```
([Total Sales], Now([Days in Datefield]))
```

## Past Year

```
PastYear( <Numeric Value> )
```

Calculates the value for the same period in the past (previous) year.

For example:

- If you're looking at a specific day, you will see the value of the same day one year back.
- If you're looking at a specific month, you will see the value of the same month one year back.

## Past Quarter

```
PastQuarter( <Numeric Value> )
```

Calculates the value for the same period in the past (previous) quarter.

For example:

- If you're looking at a specific day, you will see the value of the same day one quarter back.
- If you're looking at a specific month, you will see the value of the same month one quarter back.

## Past Month

```
PastMonth( <Numeric Value> )
```

Calculates the value for the same period in the past (previous) month.

For example:

- If you're looking at a specific day, you will see the value of the same day one month back.

## Quarter Difference

```
QDiff( <Start Time>, <End Time> )
```

Returns the difference between <Start Time> and <End Time> in quarters.  
Returns whole numbers.

## Quarter to Date Average

`QTDAvg( <Numeric Value> )`

Returns the running average starting from the beginning of the quarter up to the current time period member.

The time dimension to be used is determined by the time resolution in the widget/dashboard.

Returns 0 if the active time resolution is years.

## Quarter to Date Sum

`QTDSum( <Numeric Value> )`

Returns the running total starting from the beginning of the quarter up to the current time period member.

The time dimension to be used is determined by the time resolution in the widget/dashboard.

Returns 0 if the active time resolution is years.

## Range

`range( <Field1>, <Field2> )`

Returns a graphical range selector for a data set where two members of the same dimension and level define the minimum and maximum values of the range.

## Second Difference

`SDiff( <Start Time>, <End Time> )`

Returns the difference between <Start Time> and <End Time> in seconds.

## Year Difference

YDiff( <Start Time>, <End Time> )

Returns the difference between <Start Time> and <End Time> in years. Returns whole numbers.

## Year to Date Average

YTDAvg( <Numeric Value> )

Returns the running average starting from the beginning of the year up to the current time period member.

The time dimension to be used is determined by the time resolution in the widget/dashboard.

## Year to Date Sum

YTDSum( <Numeric Value> )

Returns the running total starting from the beginning of the year up to the current time period member.

The time dimension to be used is determined by the time resolution in the widget/dashboard.

## Other

### All

All(<Field>)

Ignores the scope set on the dimension.

## Ordering

`ORDERING(<expression1>, <expression2>)`

Returns the numeric order position of rows sorted into ascending or descending order, breaking ties with further arguments.

The expressions must be aggregated by applying the MIN/MAX functions as in the example below:

```
ORDERING(MIN([Sales Person Name]), MIN([Days in Transaction_Date]), -  
1*Sum([Sales]))
```

## Rdouble

`RDOUBLE(<R expression>, <numeric value 1>, [<numeric value 2>, ..., <numeric value n>] )`

Returns a numeric result for a given R expression and a list of numeric values (use 'args[[i]]' in the R expression to reference numeric values parameters).

The R expression is passed to the running Rserve.

The optional ordering expression determines the order in which the rows are sent to R. The argument of the Ordering parameter can be an index in your data source or you can use the ORDERING() function to determine the order of the field. For more information about the ORDERING() function, [click here](#).

For example – `RDOUBLE("m <- log(matrix(unlist(args), ncol=2));`

`kmeans(m,3)$cluster", [Total Cost], [Total Revenue])` will return the k-means cluster (R expression) of the args: [Total Cost] and [Total Revenue].

`RDOUBLE(<recycle>, <R expression>, <numeric value 1>, [<numeric value 2>, ..., <numeric value n>] )`

`recycle = TRUE` (default) – Results will be cached for unchanged functions and data.

recycle = FALSE – Results will not be cached. Use this option if your R code contains randomness.

## Rint

```
RINT(<R expression>, <numeric value 1>, [<numeric value 2>, ..., <numeric value n>] )
```

Returns an integer result for a given R expression and a list of numeric values (use 'args[[i]]' in the R expression to reference numeric values parameters).

The R expression is passed to the running Rserve.

The optional ordering expression determines the order in which the rows are sent to R. The argument of the Ordering parameter can be an index in your data source or you can use the ORDERING() function to determine the order of the field. For more information about the ORDERING() function, [click here](#).

For example – RINT("m <- log(matrix(unlist(args), ncol=2));

kmeans(m,3)\$cluster", [Total Cost], [Total Revenue]) will return the k-means cluster (R expression) of the args: [Total Cost] and [Total Revenue].

```
RINT(<recycle>, <R expression>, <numeric value 1>, [<numeric value 2>, ..., <numeric value n>] )
```

recycle = TRUE (default) – Results will be cached for unchanged functions and data.

recycle = FALSE – Results will not be cached. Use this option if your R code contains randomness.

# Introduction to Formulas

This topic provides an overview to working with formulas in Sisense along with important tips and examples.

Formulas are custom calculations performed on one or more fields in the data.

They offer an important way to analyze results and express business logic. .

Sisense formula capabilities are designed around several principles;

- Create complex business calculations without IT or technical knowledge.
- Easily combine fields from different data sources together.
- Customize formulas to reflect specific criteria and conditions.
- Work with raw data without the need to summarize data sets before creating formulas.
- Instantly recalculate formulas based on any filter, variable or level of granularity.

The table below provides a reference to the main formula functions available in Sisense. [Click here](#) to see a complete reference.

Goal	Function	Types and Syntax
Perform calculation based on criteria	Measured value	Value Filters: $\neq$ , $=$ , $>$ , $<$ , between Text Filters: Contains, Doesn't Contain, Doesn't End With, Doesn't Start With, Ends With, Start With, Equals, Not Equal List Filter: Include, Exclude Ranking Filters: Top, Bottom Ranking Time Filter: Date and Calendar
Combine data/apply simple mathematics	Aggregate functions	Operator: $+$ , $-$ , $*$ , $/$ Aggregate: Sum() Average: Avg() Count: Count(), DupCount() Range: Max()/Min()

Goal	Function	Types and Syntax
Summarize data	Statistical function	Central Tendency: Median(), Model, Largest() Std Deviation and Variance: Stdev(), Stdevp(), Varp(), Var() Quartile and Percentile: Quartile(), Percentile()
Accumulate data	Rolling sum/average	Sum to Date: YTDSum(), QTDSum, MTDSum() Avg to Date: YTDAvg, QTDAvg, MTDAvg()
Compare Time or Trends	Time functions	Past Periods: PastYear(), PastQuarter(), PastMonth(), Next(), Prev() Growth Trend: Growth(), GrowthRate() Time Difference: YDiff(), QDiff(), MDiff(), DDiff(), HDiff(), MnDiff(), SDiff()

The following topics explain how to create formulas, and describe the features of the formula editor.

- [Using the Formula Editor](#)
- [Creating and Editing a Formula](#)
- [Reusing Formulas](#)
- [Using Quick Functions](#)
- [Creating Formulas Based on Criteria and Conditions \(Filters\)](#)
- [Building Formulas with Functions](#)
- [Function Reference](#)

# Quick Functions

Just like [reusing formulas](#), quick functions is another feature to make working with formulas easier. The Widget Designer provides a variety of predefined commonly used functions that you can easily apply in the Data Browser when selecting a Value to be included in a widget.

Quick Functions instantly add a time dimension to any existing value and formula. These functions include calculations for past values, change over time, contribution and running totals. Quick Functions include all the Time Functions previously discussed but they can only be accessed by clicking on a formula that is already present in a widget.

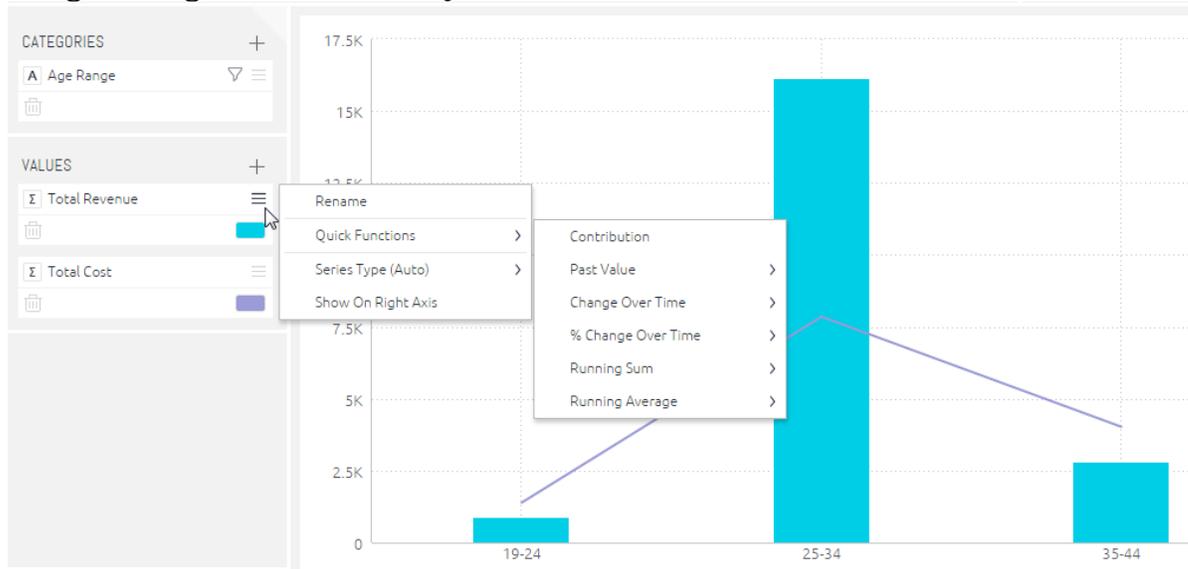
A simple example of a Quick Function is a finance manager who is reviewing total costs per month but needs to track the accumulated annual costs. They can simply use a quick function to calculate the year to date total for costs.

Months in Date	Total Cost
01/2013	12,349.68
02/2013	11,247.97
03/2013	12,504.97
04/2013	11,900.94
05/2013	12,329.42
06/2013	12,000.75

The screenshot shows the 'ROWS' section with 'Months in Date' and the 'VALUES' section with 'Total Cost'. A context menu is open over the 'Total Cost' value, with 'Quick Functions' highlighted. The sub-menu includes: Contribution, Past Value, Change Over Time, % Change Over Time, Running Sum, and Running Average.

**To use a quick function:**

1. Hover and click on the menu icon of a numeric field in the data panel of the widget designer, and select Quick Functions from the menu:

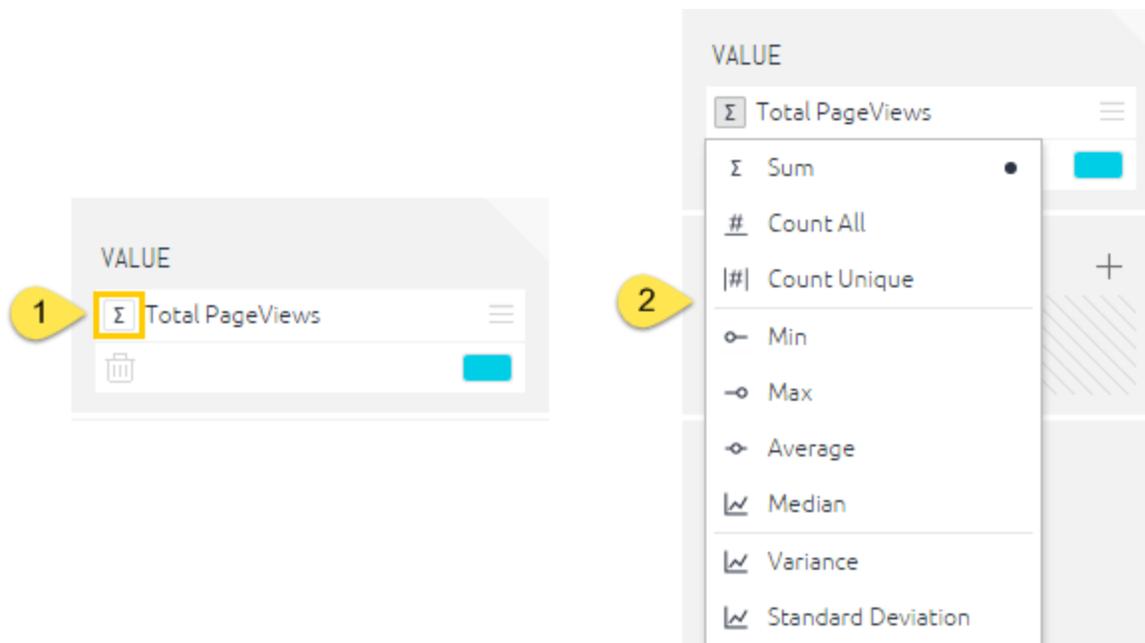


- A list of commonly used functions is displayed.
2. Select a function. The widget will be updated immediately.

## Adding Aggregate Functions

You can also add aggregate functions to your formula without opening the formula editor.

Click on the Value icon to open a list of aggregate functions, and select the function to apply to your formula.



## Starring Formulas with Quick functions

A more complex example uses starring with multiple quick functions. Let's say a finance manager also wants to compare year to date costs to the same period for the previous year. They can first apply the year to date function to total cost and then save it to the formula repository. They can then add the saved year to date formula but apply another quick formula for past values which will perform the same calculation but on data from the previous year.

**Note:** Starred (favorite) formulas will be shared with other users.

# Using R in Formulas

[R](#) is a software environment for statistical computing and graphics. Sisense supports the integration of R functions in your formulas.

You can write R code directly in the formula editor, and send fields as parameters. You can also combine R functionality with the regular functions to create advanced and tailored formulas. (For general help with the formula editor and functions, [click here](#).)

Before you can integrate R into your formulas, you must have an R server set up in your organization.

This document will cover the following:

- Understanding How R works in Sisense
- Connecting Sisense to Your R Server
- Using R functions in Formulas
- Guidelines for using R code in Sisense
- Example 1: Simple R Function – Kmeans Clustering
- Example 2: R Logistic Regression via Sisense
- Example 3: Loading an Existing R Model
- Example 4: Building and Saving a R Model using Data from Sisense

## Understanding How R works in Sisense

Within a widget, R code is used and assigned fields as parameters, the following happens:

1. The selected fields are sent fully to R as a list of arrays.
2. The R code is executed on the R server, referencing the fields as needed.
3. The numeric results from R are sent back to the widget for visualization.

Each request from the R server is done in a self-contained namespace, thus R variables and results cannot be reused across requests unless they are loaded from a saved location.

R calculations are computed in the R server, and not in the ElastiCube, this means that for larger data sets results will not return as fast as regular widgets. Sisense recommends [Revolution Analytics](#) (Revolution R Open) for enhanced R performance.

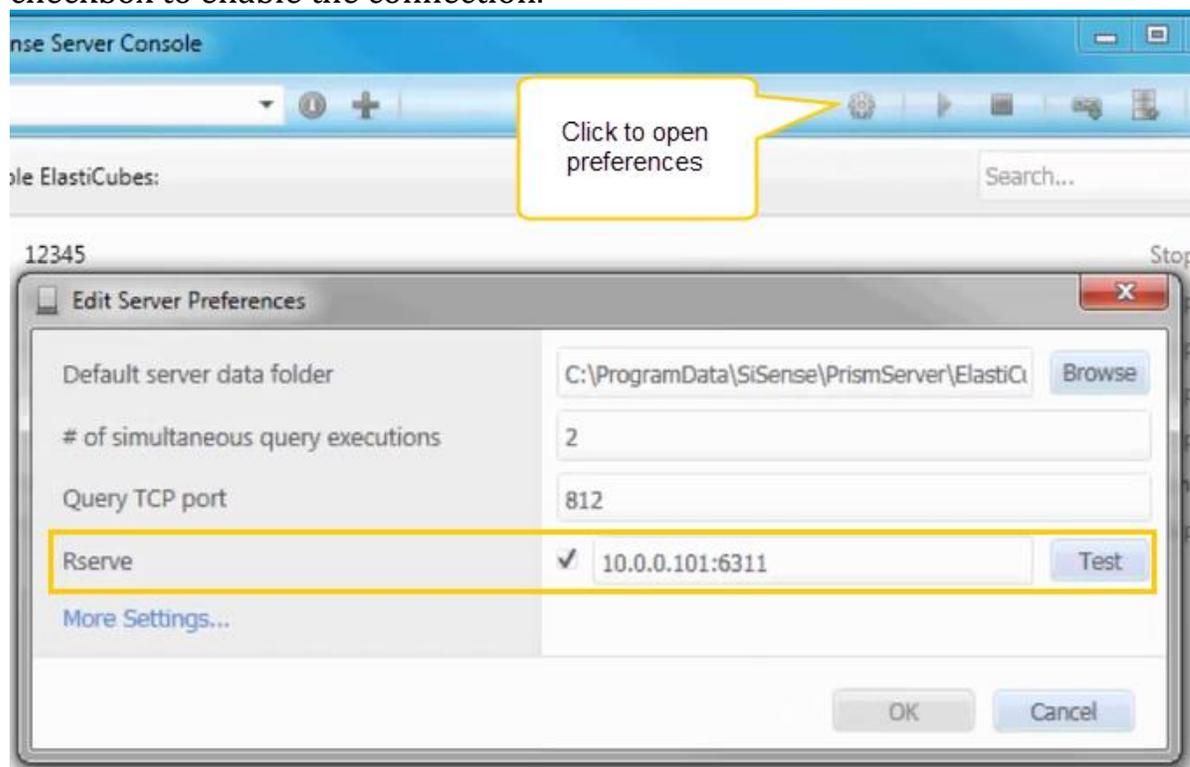
## Running R on Windows

If your R server is running on Windows you will be bound to the concurrency limitations the R distribution has on Windows. This means that concurrent requests to R may return wrong results. Sisense recommends using R on a Linux machine for production environment with concurrent user activity.

## Connecting Sisense to Your R Server

**To connect your Sisense instance to your R server:**

1. Open the Sisense Server Console from the start menu.
2. Click on the Server Preferences icon to open the server settings.
3. In the **Rserve** field, enter the IP address of your R server, and select the checkbox to enable the connection.



You can test the connection by clicking on the **Test** button.

## Installing Rserve

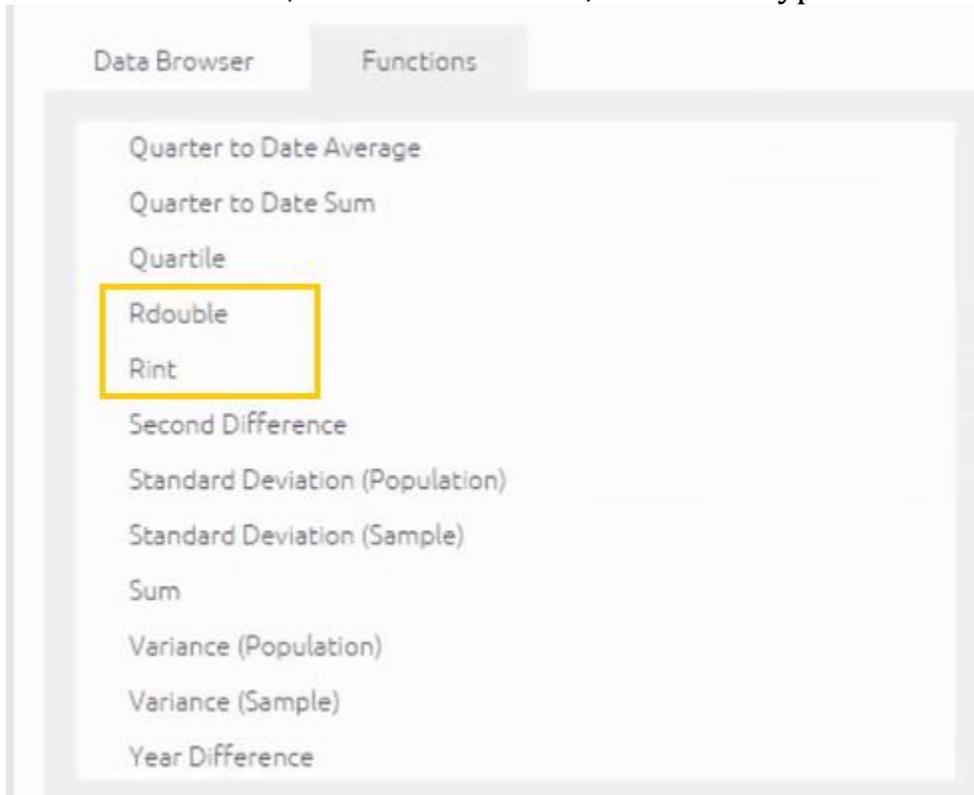
Sisense can only connect to a running Rserve instance. For more information about installing and running Rserve [follow this link](#).

Please note that if you're using RStudio, Rserve still has to be installed.

## Using R functions in Formulas

Using R code in Sisense is achieved by using two dedicated formulas – RINT & RDOUBLE. Depending on the result type of your R code, use the appropriate function within the [formula editor](#).

For either formula, RINT or RDOUBLE, the return type has to be an exact match.



Multiple numeric field values can be passed as parameters to the R functions, which within your R code can be accessed via the 'args' argument (See K-means example below).

## Syntax:

```
RINT([recycle (true)], [<Ordering>], <R expression>,
[<numeric value 1>, ..., <numeric value n>] )
```

```
RDOUBLE([recycle (true)], [<Ordering>], <R expression>,
[<numeric value 1>, ..., <numeric value n>] )
```

Sisense does not support the use of a final ';' in R statements. If you add a semi-colon at the end of your R statement, the code will fail.

## Parameters:

- **Recycle true/false (default = true)** – This is an optional parameter that controls whether the results from R should be recycled (cached), so that consequent queries will not have to be recalculated unless they or the data have changed. Generally, this behavior is automatically managed by the ElastiCube automatically for all functions. However, since R code might have non-deterministic components to it (such as randomness functions or date-specific functions), the ElastiCube cannot rely on a data-set and function that hasn't changed not to return a different result in multiple executions.

By default, the Recycle value is set to true. Use 'off' if your R code contains randomness or other non-deterministic content.

- **Ordering** – This is an optional parameter that defines the sort order in which numeric data is sent to R. The argument of the Ordering parameter can be an index in your data source or you can use the ORDERING() function to determine the order of your fields. This function arranges the values of the arguments into ascending or descending order, breaking ties by further arguments.

For example:

```
ORDERING([Total Sales], -1*[COUNT Salesman],
MIN(<Office Name>))
```

For more information about the ORDERING() function, [click here](#).

- **R expression** – Your R code needs to be passed here, wrapped in double quotes. R expects the return type to be an array with the same size as widget's row count. Nulls will be used to make up for shorter arrays, and longer arrays will be trimmed.

Use single quotes to wrap strings within your R code when using the Rint/Rdouble functions, so that there will be no double-quote collision with the quotes wrapping your R code within the Rint/Rdouble function.

- **Numeric Value Arguments** – Numeric values can be passed as arguments to your R code.
- All arguments are passed to R as a 1-based list named “args”. Each item in the list contains an array that represents the field.

For example:

- `args[[1]]` will return an array which represents the first field that was used as an argument.
- `args[[2]][3]` will return the 3rd data value within the 2nd field that was used as an argument.

## Guidelines for using R code in Sisense

In order for your R code to work well in Sisense you need to make sure of the following:

1. The resultset returned from R into Sisense has to be equal in size to the dataset that is sent to R as parameters.
2. The order of the resultset return from R into Sisense has to match the incoming order. No assumptions can be made about the incoming order – This means that you cannot assume that the incoming order matches the order you visually see in the widget.

## Example 1: Simple R Function – Kmeans Clustering

**Description:** In the example below the Total Cost and Total Revenue fields from Sisense are being used to cluster data via a Kmeans function. The result will cluster the data based on the kmeans settings, in this case 4 clusters.

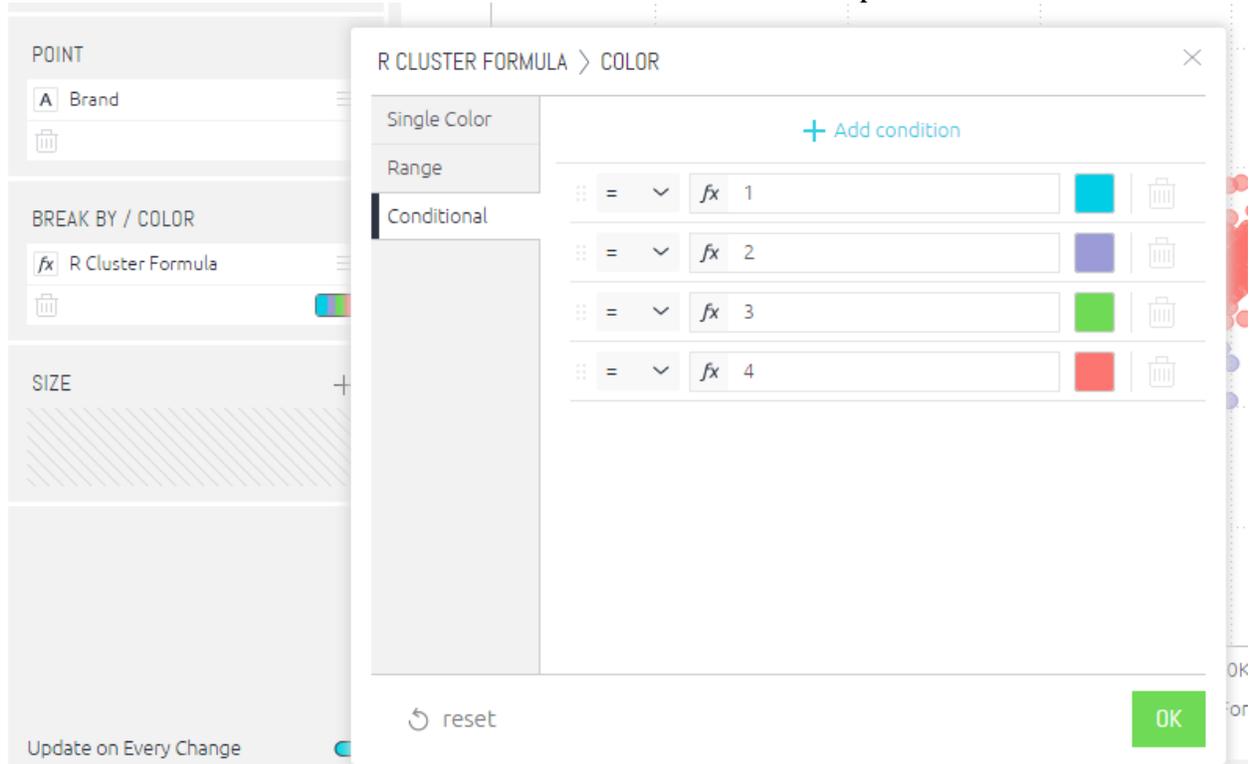
**Sisense Syntax:**

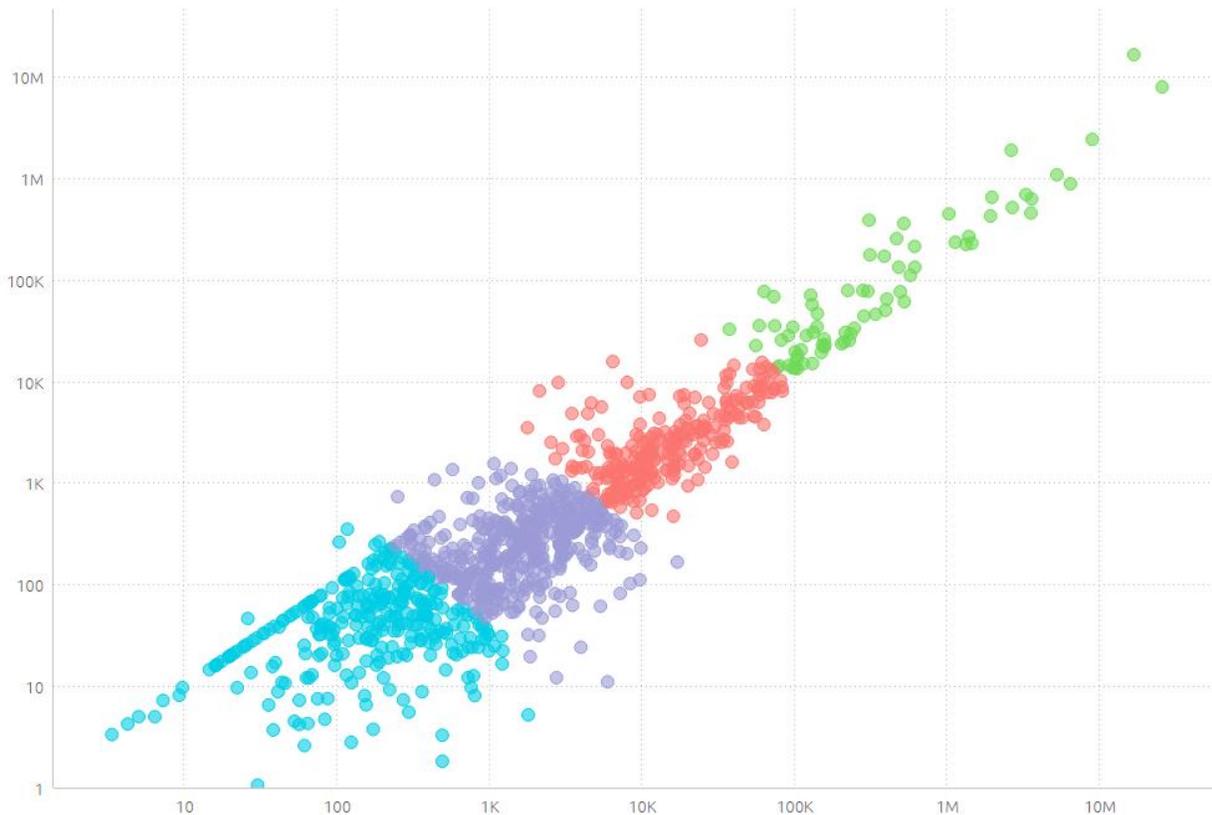
```
RINT(TRUE, "m<- log(matrix(unlist(args), ncol = 2));
kmeans(m,4)$cluster" , [Total Cost], [Total Revenue])
```

**Result:** The R expression will return a result from 1 to 4. The widgets color settings can be configured to color corresponding values based on if they equal 1, 2, 3 or 4.

**Tip:** You can [save \(star\)](#) the complex formula above, and use it again in additional formulas and charts.

In the chart's BREAK BY/COLOR panel, you can change the color of the clusters as well as define the number of clusters and their break points.





## Example 2: R Logistic Regression via Sisense

**Description:** In the example below website traffic data is used to predict if a conversion is likely to occur based on the source of web traffic. The following fields contained in the ElastiCube are passed to the R logistic model: Average Outcome (1=conversion, 0= no conversion), Average Time on Page, Average Pages (viewed) and Average Bounce Rate. This data is passed to a model to predict the outcome and likelihood the traffic will convert.

[For the sake of illustration both the training and test data set are the same in the example]

### Sisense Syntax:

```
RDOUBLE(TRUE, "mydata<-data.frame(convert=args[[1]],
time=args[[2]], pages=args[[3]],
bounce=args[[4]]);lrmodel <- glm(convert ~ time + pages +
```

```
bounce, data = mydata, family = 'binomial');prob <-
predict(lrmodel, newdata = mydata, type = 'response')",
[Average Outcome], [Average TimeonPage], [Average Pages],
[Average BounceRate])
```

**Result:** The R expression will return a result from 0 to 1 showing the likelihood of a conversion occurring. Note a similar process can be applied to create a linear regression.

Ref	Predicted Outcome
(direct) / (none)	0.60
t.co / referral	0.57
social / twitter	0.56
bing / organic	0.54
amddatawiz.com / referral	0.54
bing.com / referral	0.53
answers.yahoo.com / referral	0.50
bing.com / referral	0.45
bingsandbox.com / referral	0.43
buy-cheap-online.info / referral	0.41
checkpagerank.net / referral	0.40
getpocket.com / referral	0.39
criticue.com / referral	0.39

Predicted outcome of website conversion based on logistic regression.

## Example 3: Loading an Existing R Model

**Description:** In the example below a saved logistic regression built on previous training data is called and used with new website traffic data to predict if a conversion occurs. The following fields contained in the ElastiCube are passed to the saved R logistic model: Average Outcome (1=conversion, 0= no conversion), Average Time on Page, Average Pages (viewed) and Average Bounce Rate. This data is passed to a model to predict the outcome and likelihood the traffic will convert.

**Sisense Syntax:**

```
RDOUBLE(TRUE, "mydata<-data.frame(convert=args[[1]],
time=args[[2]], pages=args[[3]],
bounce=args[[4]]);load('C:\rdata');prob <-
predict(lrmodel, newdata = mydata, type = 'response')",
[Average Outcome], [Average TimeonPage], [Average Pages],
[Average BounceRate])
```

**Result:** The R expression will return a result from 0 to 1 showing the likelihood of a conversion occurring based on the loaded R model in this case stored in 'rdata'.

## Example 4: Building and Saving a R Model using Data from Sisense

**Description:** In the example below a logistic regression to predict if a website conversion will occur is built using data from Sisense and then saved as an R model. The following fields contained in the ElastiCube are passed to create the R logistic model: Average Outcome (1=conversion, 0= no conversion), Average Time on Page, Average Pages (viewed) and Average Bounce Rate. This data is saved as an rdata file and can be used on other data sets to predict outcomes.

**Sisense Syntax:**

```
RDOUBLE(TRUE, "mydata<-data.frame(convert=args[[1]],
time=args[[2]], pages=args[[3]],
bounce=args[[4]]);lrmodel <- glm(convert ~ time + pages +
bounce, data = mydata, family = 'binomial');
save(lrmodel, file = 'C:\rdata') ;1", [Average Outcome],
[Average TimeonPage], [Average Pages], [Average
BounceRate])
```

**Result:** The R expression will use Sisense data to build and save a logistic model in the specified file in this case 'rdata'.

# Reusing Formulas

This topic describes how you can reuse formulas that you have previously marked as a favorite (starred).

Formulas are saved per ElastiCube.

## To mark a formula as a Favorite:

1. While defining a formula, click the Favorite (Star) button.
2. Enter a name for this Formula.
3. Click **OK**.

Filter Age Range / Filter a Field

List	<input checked="" type="checkbox"/>	Start typing to search...	☰
Text	<input checked="" type="checkbox"/>	0-18	
Ranking	<input checked="" type="checkbox"/>	19-24	
Starred	<input checked="" type="checkbox"/>	25-34	
	<input checked="" type="checkbox"/>	35-44	
	<input checked="" type="checkbox"/>	45-54	
	<input checked="" type="checkbox"/>	55-64	
	<input checked="" type="checkbox"/>	65+	

Advanced

Include all **OK**

## To reuse a favorite formula:

- Favorite Formulas appear in the Data Browser under the title Formulas, as shown below. Simply select it to use it.

*fx* Type to search for fields

Suggestions

- brand / Brand
- category / Category
- Commerce / Age Range
- Commerce / Condition
- # Commerce / Cost
- Commerce / Date
- Commerce / Gender
- # Commerce / Quantity
- # Commerce / Revenue
- # Commerce / Visit ID

Formulas

MyRevenueFormula

Commerce

- Age Range

**Note:** To see the underlying building blocks of your favorite formula, right-click the favorite formula, and select **Use Flattened Formula**.

[Past Year Cost]

- Rename
- Delete
- Use Flattened Formula

# Formula Editor

The Formula Editor is where Designers can define formulas for a dashboard's widgets.

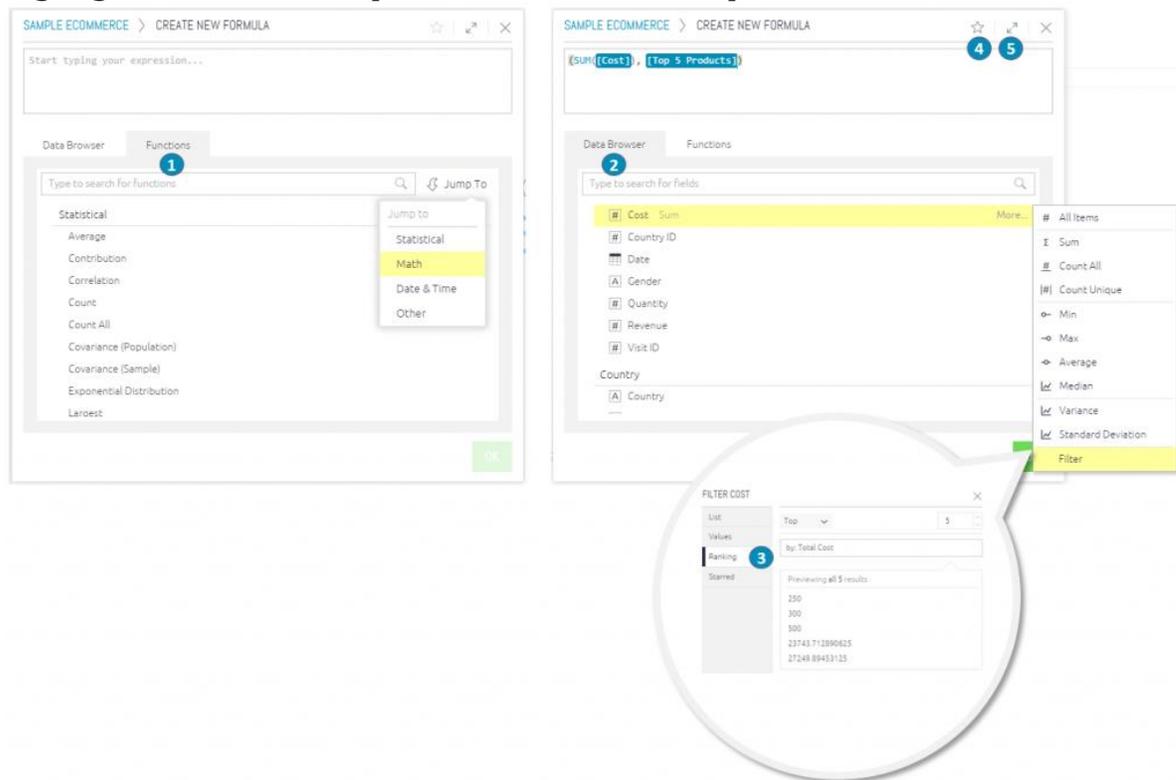
If you are familiar with the Formula Editor, see the [Quick Reference Guide](#). For a detailed guide with examples, continue reading below.

## To open the Formula Editor:

- Do one of the following:
  - For a new widget, click **Select Data**, and then .
  - For an existing widget, click on the edit formula button .

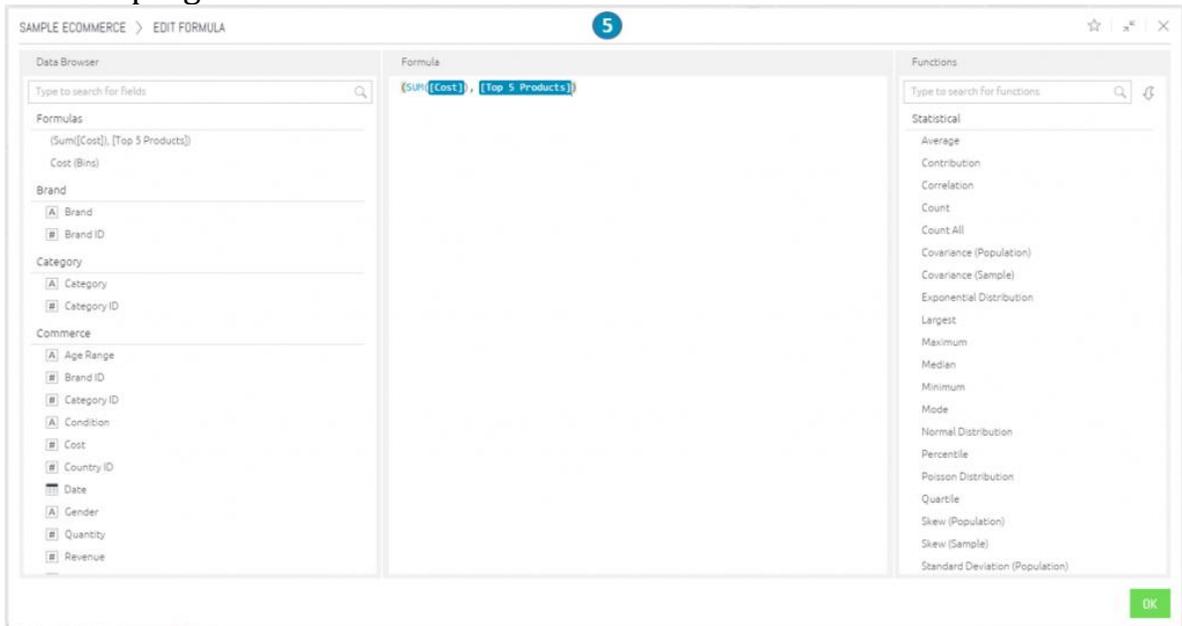
The Formula Editor has two tabs, the Data Browser to select fields and the Functions tab to select formula operations. You can create a formula combining one or more function, field and filters. The diagram below

highlights the main components of the formula panel.



1. **Functions** are operations which perform different calculations, for example a sum. Use the 'Jump To' menu or the search box to quickly find the formula you need.
2. Fields in the **Data Browser** are variables contained in the data set (ElastiCube). Clicking on a field in the data browser will include it as part of the formula.
3. [Filters](#) can be applied to restrict formulas based on criteria.
4. [Starring](#) is a way to save a formula for later use.

5. The Formula Editor window can be expanded by clicking the expand button at the top right.

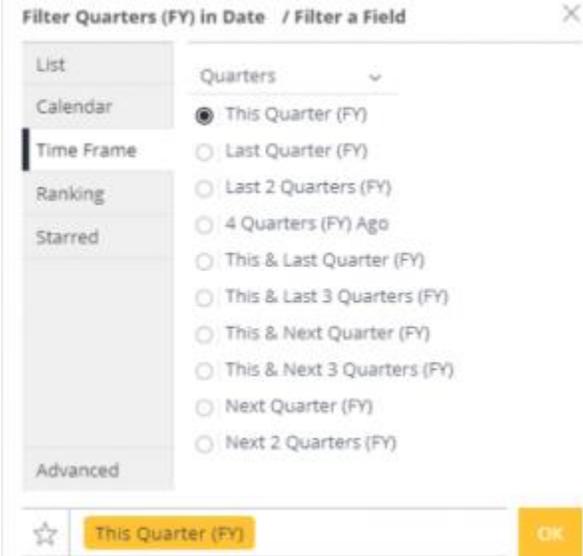


# Fiscal Years

By default, Sisense calculates the beginning of the fiscal year as January 1st. In some cases, you might need to express your company’s date fields according to another fiscal month, such as April 1st. Sisense allows you to easily shift your fiscal calendar to start at a different quarter, April, July, and October.

This can be defined in the system level, by the Administrator, and also at the ElastiCube level, by the Data Designer. This means that you can support one fiscal calendar month for your system, such as January 1st, while a specific ElastiCube could be set according to the US tax year which begins October 1st.

When you change the fiscal calendar, this affects how time is calculated in your dashboard and you are also getting an indication, for example when filtering your widgets and dashboards by time, through the FY tag.



Date Level	Effects
Year	Changes the calendar year to fiscal. By default, the name reflects the year in which it ends. For example, if the fiscal year

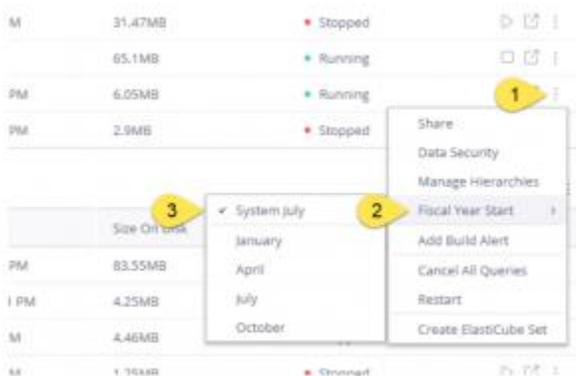
Date Level	Effects
	starts in April, the year for the date June 1, 2014 would be shown as FY 2015.
Quarter	Changes the calendar quarter to that of the quarter in the fiscal year. For example, if the fiscal year starts in April, the quarter for the date June 1, 2014 would be Q1.
Month	Not affected by fiscal year settings. The calendar month is the same as the fiscal month.
Week	Changes the calendar week to the week of the fiscal year. For example, if your fiscal year begins April 1st, April 3rd would fall under Week 1.
Day/Minute/Hour	Not affected by fiscal year settings.

To set your fiscal year on the system level, see [System Settings](#).

Administrators and Data Designers can set the fiscal year on the ElastiCube level.

**To set your fiscal year on the ElastiCube level:**

1. In the Sisense Web Application, select **Admin > Data Sources**.
2. For the relevant ElastiCube, select the ElastiCube menu > **Fiscal Year Start**, and then select the month.



# Formatting Fiscal Dates

When you apply fiscal years, the default format for your dates is 2018 FY for years or Q1 2018 FY for quarters.

You can customize the format of the date through placeholders such as “y” and “p” where “y” represents the year and “p” the previous year.

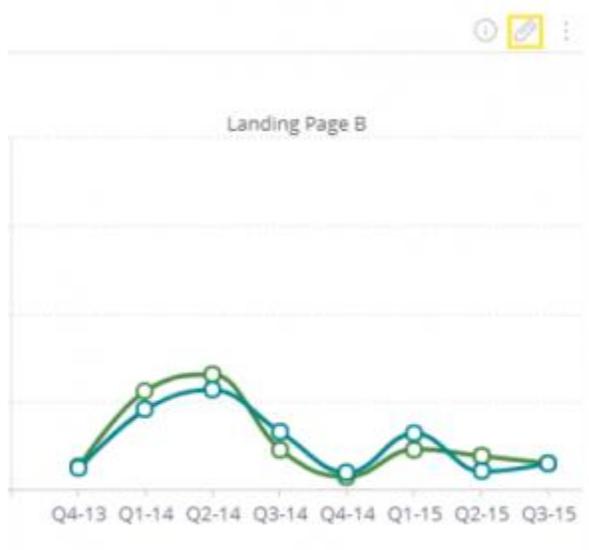
The examples below illustrate how you can format your dates:

- 2017/18: yyyypp/yy
- 17/18: yp/yy
- FY17/18: FY yp/yy
- Q1 2018 FY: Q yyyy FY

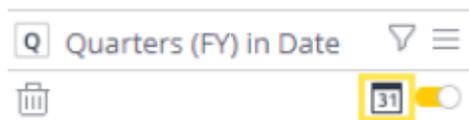
By default, Sisense applies FY at the end of the year when fiscal years are applied. However, if you make any changes, Sisense will no longer add the FY automatically, so you must define this with the placeholder FY in the location you want as shown above.

## To format your fiscal dates:

1. On the Dashboard, click the **Pencil (Edit)** button  that appears in the top-right corner of a widget with a date dimension.



2. Select the Calendar icon to define your formatting.



3. Select the relevant time period and enter the relevant placeholders in the format you want to be displayed in your widget. A preview of the custom format is displayed below.

Quarters (FY) in Date > Format Date

Years

Quarters

Months

Weeks

Days

Time

Format:

Q1 1970

1970 Q1

Q-yy-FY

Q1-70-FY

reset OK

4. Click OK. The format of your dates is updated in the widget.

# Date and Time Fields

Many widgets include data that is relative to date and time fields. Sisense enables you to apply date and time filters so you can easily identify trends in your data. When you create a widget in your dashboard that includes a date field, Sisense automatically breaks the data down by year, however, you can change this to a shorter resolution. For example, a Date field included in a widget can be separated into Years, Quarters, Months, Weeks, Days, and Time from your widget's filter.

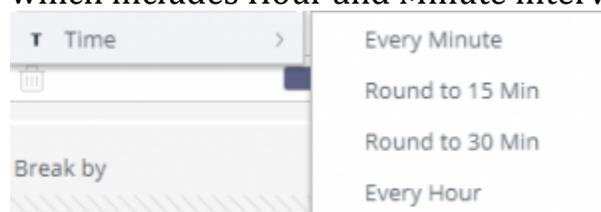


Dashboard Viewers have the option to drill down the widget according to the Date level as shown in the following image:



Viewers can drill down into the lowest date or time resolution available in your data.

The highest resolution available is by Year and the lowest resolution is Time, which includes Hour and Minute intervals.



For more information about filters, see [Creating Dashboards Filters](#).

# Time in Sisense

Delete this text and replace it with your own content.

# Changing a Dashboard's Data Source

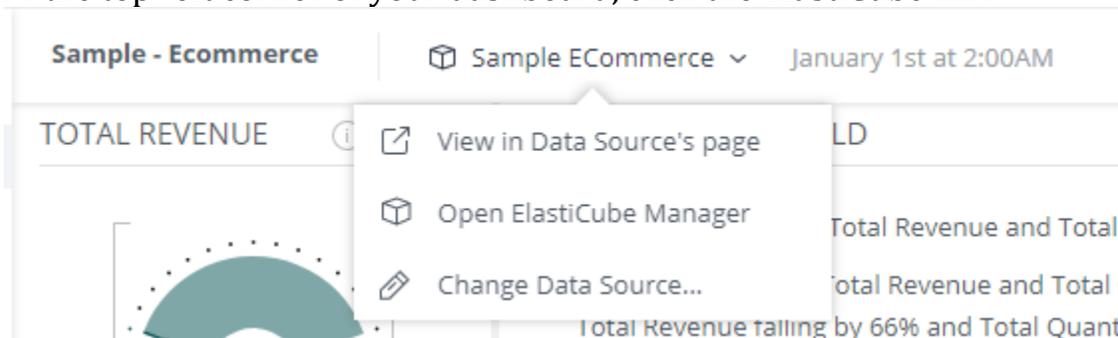
This topic describes how you can change your dashboard's data source.

Changing a dashboard's data source is useful when you have recently changed your server and you need to reassign your dashboard's ElastiCube or in cases where you have deleted an ElastiCube Set and you want to reassign the dashboards in that set to your remaining ElastiCubes.

Dashboards can also support multiple data sources, which means you can have widgets built on several data sources within a single dashboard. For example, if you have widgets from an ElastiCube on a dashboard, you can add a data set to the dashboard and create widgets built on that data set's live data source. The widgets from both data sources continue to function independently of each other in the same dashboard. In addition, filters applied to fields from each data source, do not affect the fields from another data source. For more information about filtering, click [here](#).

## To change your dashboard's ElastiCube:

1. In the top-left corner of your dashboard, click the ElastiCube link.



2. Hover over Change Data Source and select the new data source.  
The source is added to your list. You now add new widgets from any of your data sources to the dashboard.

# Creating Dashboards

This topic describes how to create a new dashboard.

There are two ways you can create dashboards, first, through the Sisense Web Application, and the second, through the Sisense REST API.

The difference between these two methods is the Sisense Web Application provides an interface where you populate the dashboard with widgets.

To create a dashboard through the REST API, you define the dashboard and then populate it with widgets that already exist through their widget IDs.

## Creating a Dashboard

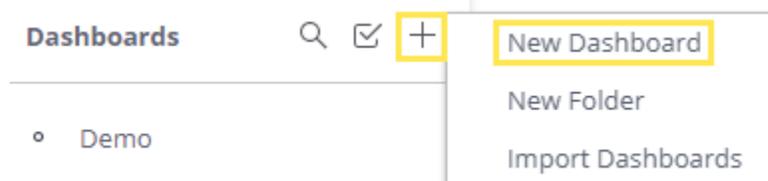
Each Dashboard is based on the data selected from one or more ElastiCubes or data sets. The instructions below assume that the required ElastiCube(s) and data set(s) have already been created. You can also use the sample ElastiCubes that are installed with Sisense if you have not yet created your own.

### To create a dashboard:

1. In the **Analytics** page, click **+** above the Dashboards list.

OR

Right-click on the folder menu and select **New Dashboard**.



The following window appears:

New Dashboard / Please select a Data Set to connect to

Data Set: Sample ECommerce

Title: Sample ECommerce

Create

Cancel

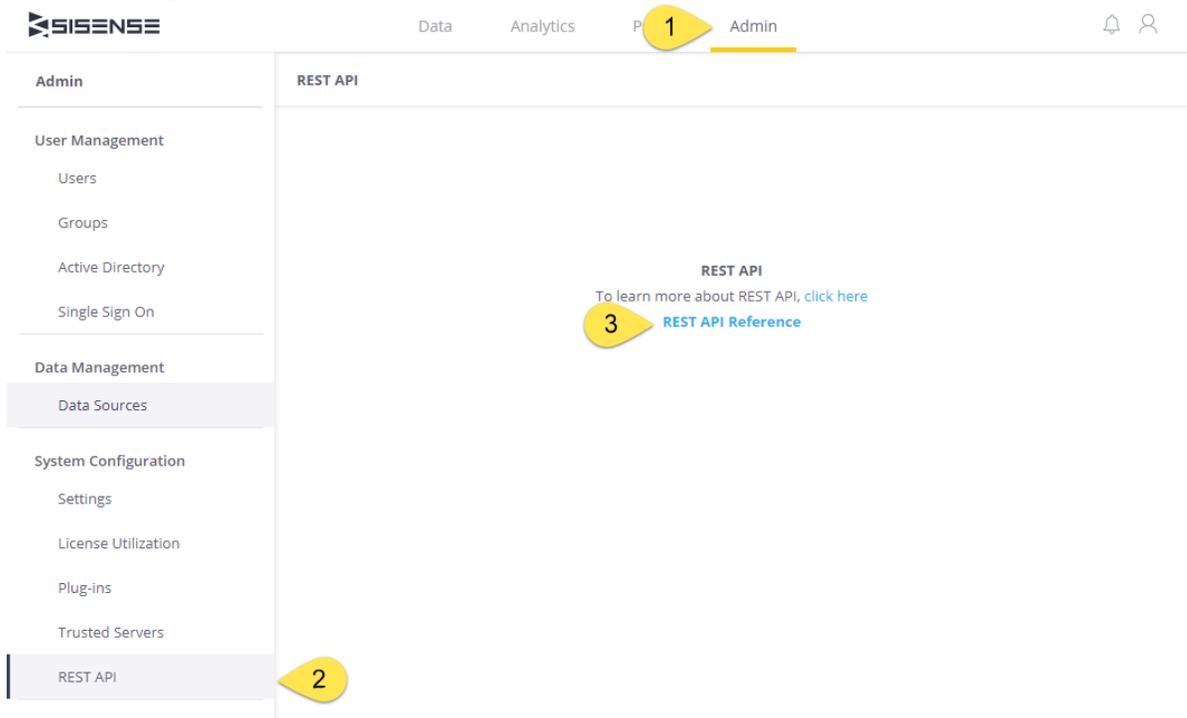
2. Click the name of the **Data Set** displayed to select the ElastiCube or data set you want to work with. ElastiCubes have the  next to their name while data sets have the  icon next to their name.
3. Click the name of the **Title**, and enter a name for the dashboard. This name will appear at the top of the dashboard, and in the **Dashboards** list.
4. Click **Create**.  
You are then automatically guided through the process of creating your first widget in the Widget Wizard, as described in [Adding Widgets to a Dashboard](#).

## Creating Dashboards through the REST API

The Sisense REST API enables you to programmatically create dashboards by defining the basic properties of a dashboard such as Title, ElastiCube (datasource), and owner. After defining the properties of the dashboard, you define its visualizations by including the widget IDs of pre-existing widgets.

**To create a dashboard through the REST API:**

1. In the Sisense Web Application, click **Admin** at the top, and then **REST API**, followed by REST API Reference.



2. Select **Settings**, then **POST /dashboards**.
3. In the body of your call, define the values of the following keys:

Key	Value
Title	String that defines the name of the dashboard as displayed in the Sisense Web Application.
desc	String that represents a short description of the dashboard.
datasource	Object that defines the data source where the dashboard pulls its data. This object contains the following elements: <b>title:</b> The name of the ElastiCube. <b>id:</b> The address and name of the ElastiCube where the dashboard queries its data from. <b>address:</b> The address of the dashboard. <b>database:</b> The name of the ElastiCube where the data is pulled from by the dashboard.

Key	Value
	<b>fullname:</b> The name of the ElastiCube.
filters	<p>This object defines the dashboard filters that the user add goes to the "filters" object in the .dash file.</p> <p>This object contains the following elements:</p> <p><b>datatype:</b> The type of data affected by the filter, for example, text.</p> <p><b>title:</b> The name of the filter displayed in the Filters pane.</p>
layout	<p>Object that defines what widgets are displayed in the dashboard and the layout.</p> <p>This object contains the following elements:</p> <p><b>type:</b> String that defines how the layout of the dashboard. This value should be "columnar".</p> <p><b>columns:</b> Array objects that defines the width of each column in the dashboard and how widgets are displayed in each column.</p> <p><b>cells:</b> Object that contains the subcells.</p> <p><b>subcells:</b> Object that contains the elements objects where widgets are defined.</p> <p><b>elements:</b> This objects contains the widget and defines its dimensions within the ElastiCube.</p>
defaultFilters	This object defines the default filters of a dashboard. When a user selects "Set as my default filters" from the Filters menu in the Sisense Web Application, the filters are defined in this object.

4. Click **POST**. The Dashboard is added to the Dashboards list in the Sisense Web Application.

# Customizing the Dashboard Layout

This topic describes how to rearrange your dashboard layout by moving your widgets around and resizing them.

By default, widgets are arranged in a dashboard one underneath the other. When you add a new widget, it is added at the bottom of the dashboard.

You can drag widgets one on top of another to split them horizontally by selecting the widget's title and moving the widget to the relevant location.

The dashboard can also be organized in columns. By default, a dashboard is created with a single column, but you can add up to four columns. Each column can contain multiple widgets, but a widget cannot span more than one column.

You can rearrange the order of the widgets, reorganize the widgets in columns, add/remove widgets from columns, and resize columns/widgets.

## Layout Mode and View Mode

A dashboard can be viewed in either Layout mode or in View mode.

### Layout Mode

Layout mode is the default mode in which you can add new widgets and rearrange them on the dashboard, as described throughout this section.

 is displayed in the top-right corner of the screen when the dashboard is in Layout mode. Click this button to display View mode.

### View Mode

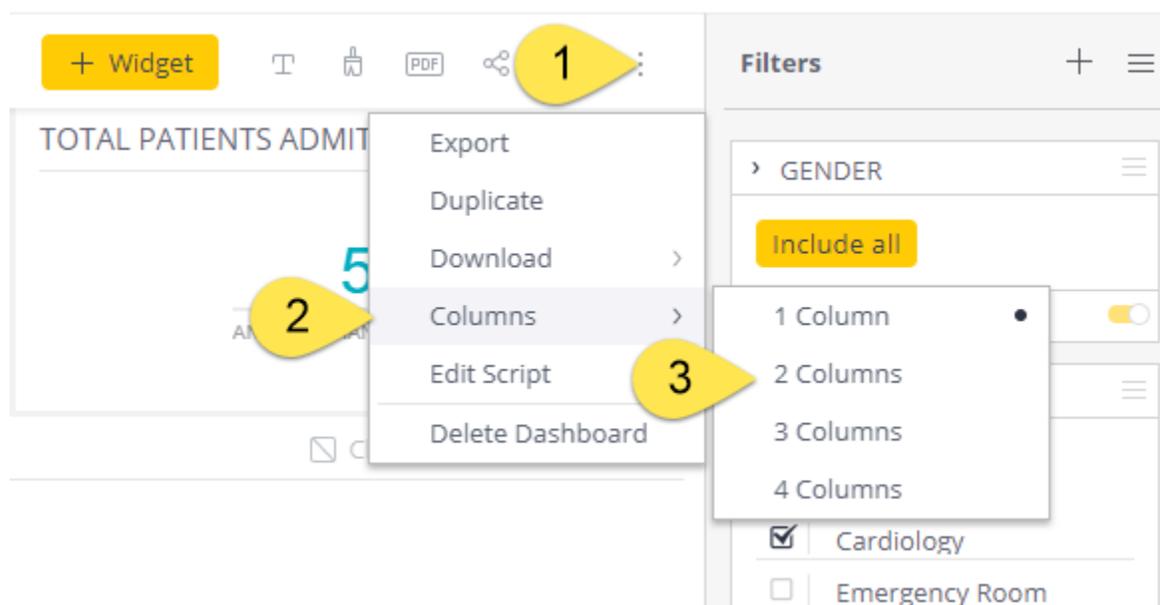
View mode enables you to see what the dashboard looks like when a user is only viewing the dashboard. In View mode, the dashboard cannot be rearranged or edited.

 is displayed when the dashboard is in View mode. Click this button to display Layout mode.

## Adding Dashboard Columns

**To change the number of columns:**

- Click on the dashboard's menu, and select Columns and the number of columns.

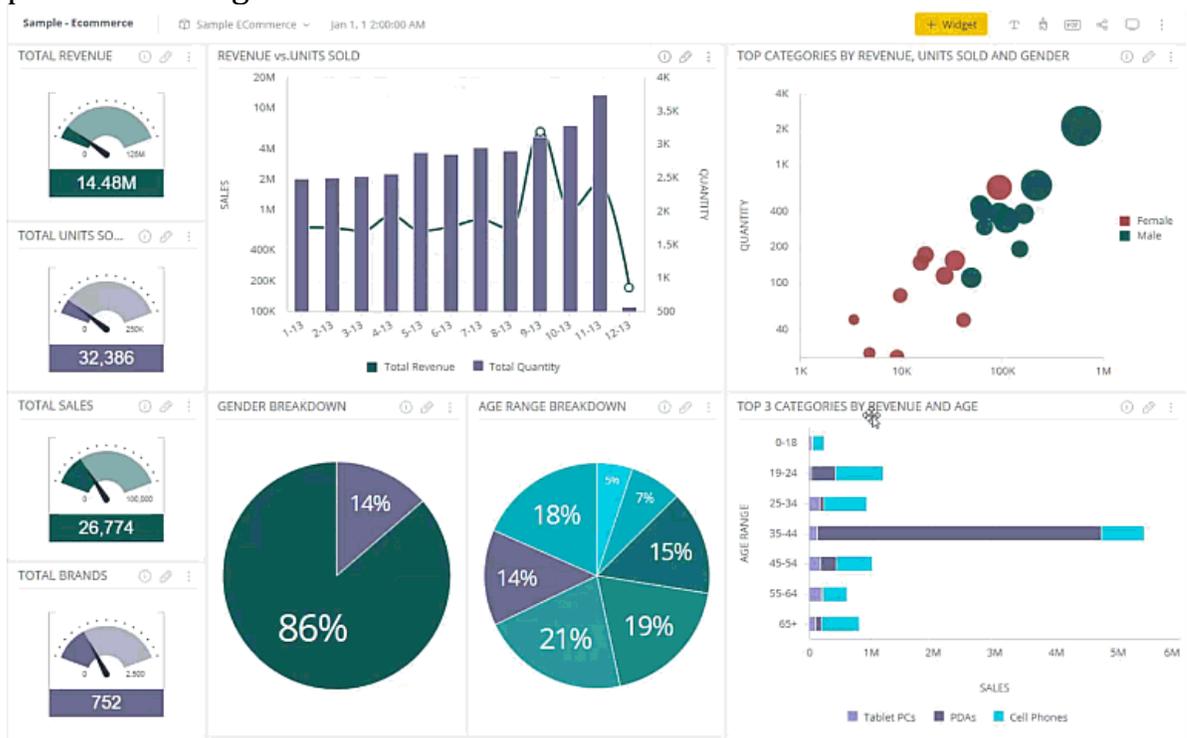


## Moving a Widget on the Dashboard

**To move a widget to a different column or position:**

1. Make sure that you are in Layout mode (default mode). The Layout mode indicator  is displayed in the top-right corner of the dashboard. If not, click  to go from View Mode to Layout mode.

2. Drag and drop the widget into the desired position by dragging it by the top part of the widget.



The position of a widget can be split to contain multiple widgets within a single column by dragging one widget on top of another.

## Placing Widgets Side-by-Side

You can place widgets side-by-side using one of the following methods:

- By putting each widget in a different column.

OR

By dragging one widget on top of the other. Both these widgets will then be in the same column.

## Resizing Widgets

Resizing the column height or width automatically resizes the widgets in the best possible way.

**To resize a widget:**

- Resize the dashboard column in which the widget is located by dragging the edge of the column right or left. All the widgets in this same column are resized accordingly in an optimal manner.
- Resize the widget by dragging its edges right, left, up or down.



Refer to Ticker widgets below for a description of how an Indicator Widget turns into a ticker type widget when you reduce its height.

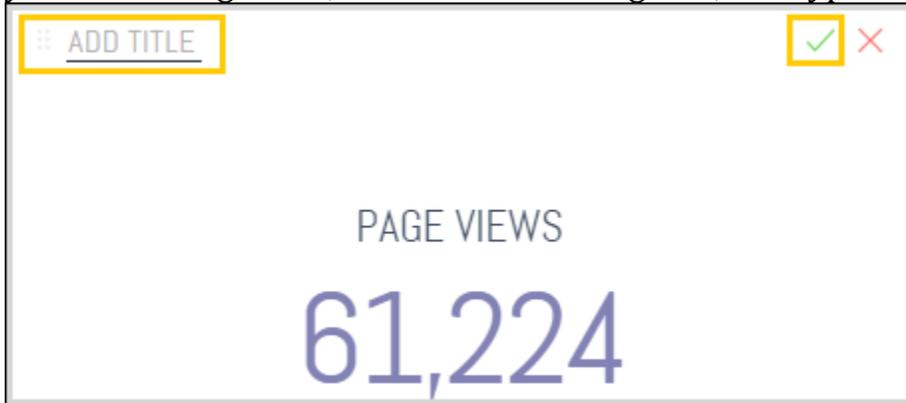
## Adding Widget Titles

You can add titles to your widgets directly in the dashboard view.

If you already added a title when you created the widget in the Widget Designer, you can edit the title directly in the dashboard.

**To add a widget title from the dashboard:**

1. To add a title, click **Add Title** at the top of the widget, and type in the title. If you are editing a title, click on the existing title, and type in the new title.



2. Click on the green ✓ icon to save the title.

## Working with Ticker Widgets

If you resize the height of an Indicator widget, it automatically turns into a Ticker-type widget. An Indicator widget that was a numeric Indicator appears differently to an Indicator widget that was a gauge.

### Regular Indicators

Full size Indicator widget	Full size Gauge widget

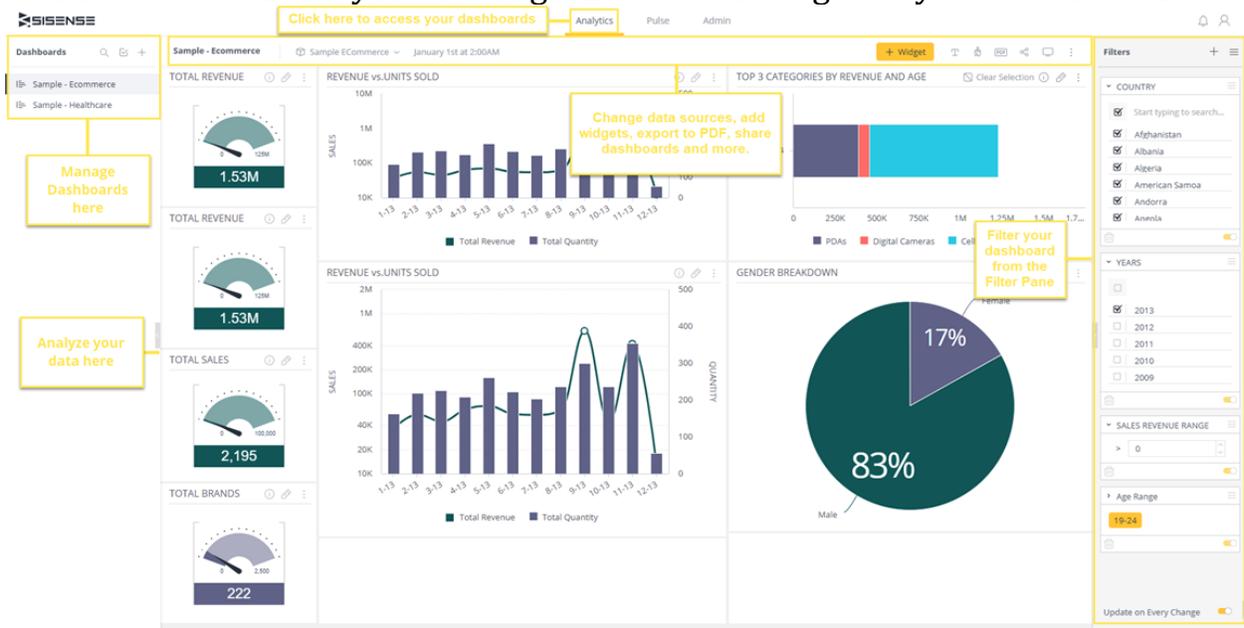
### Ticker Indicators

GROWTH OF VISITORS: <b>13.2%</b>   # of Visitors: 220K	   REVENUE: <b>2M</b>   Growth: 0.2%
Indicator widget after resizing down	Gauge widget after resizing down

# Dashboards

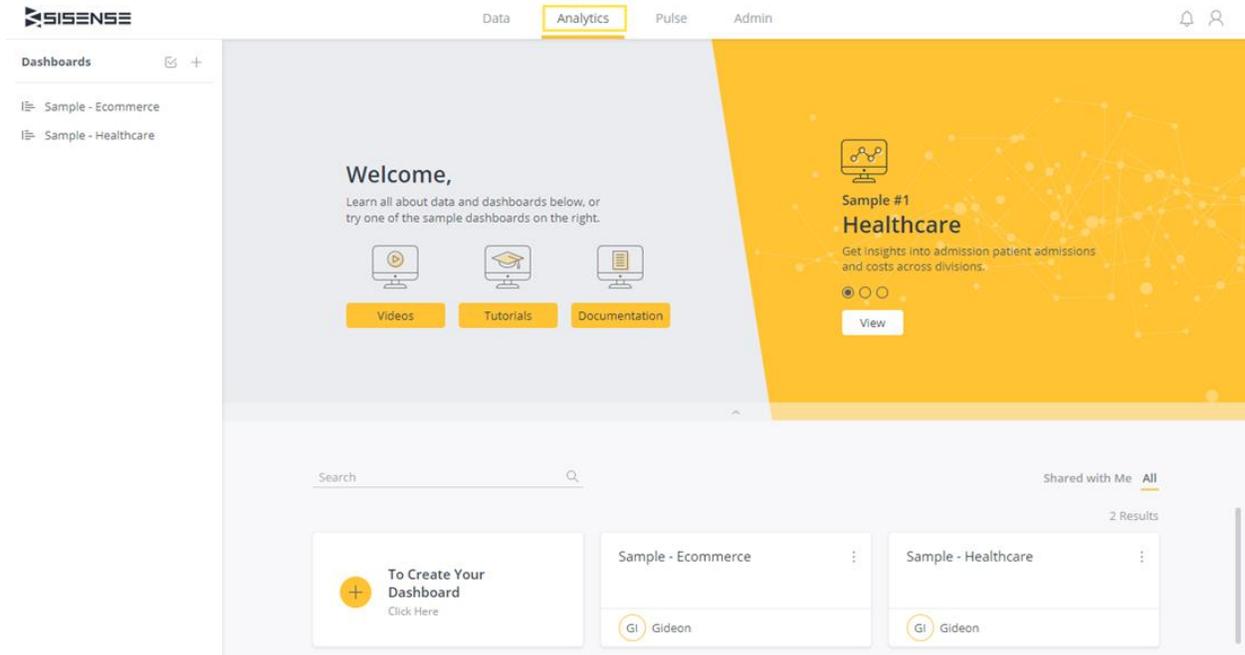
After importing data into Sisense, the next step is to create dashboards that Viewers can interact with.

This section will show you how to get started working with your dashboards.



## Creating Dashboards

Designers and Dashboard Designers can create dashboards from the **Analytics** page where all your existing dashboards are displayed.



When you create a dashboard, your dashboard is initially empty. The next step is to add widgets, or visualizations, to your dashboard to make it meaningful. The topics below provide information about creating dashboards, adding widgets to it, and customizing the data displayed in the dashboard.

- [Creating Dashboards](#)
- [Adding Widgets to a Dashboard](#)
- [Managing Widgets](#)
- [Filtering Dashboards](#)
- [Formulas](#)
- [Working with Time](#)
- [Customizing the Dashboard Layout](#)
- [Changing the Color Palette](#)
- [Changing a Dashboard's Datasource](#)

## Managing and Sharing Dashboards

As a Designer, you can manage dashboards that you have created, or that have been shared with you; these topics explain how.

- [Organizing Dashboards](#)
- [Locating Dashboards](#)
- [Deleting Dashboards](#)

- [Copying Dashboards](#)
- [Sharing Dashboards](#)

## Interacting with Dashboards

Now that you've created your dashboards, it's time to start interacting with them.

The topics below will help you learn how to get the most out of your dashboards.

- [Viewing Dashboards](#)
- [Interacting with Filters as a Viewer](#)
- [Changing Sisense's Language](#)
- [Drilling Down in a Widget](#)
- [Making Selections in a Widget](#)
- [Exporting Dashboards](#)

# Reporting in Sisense

After you have viewed your dashboards, you may have discovered some insights you want to share. Sisense supports a variety of ways you can share your data with other Sisense users.

Downloading widgets and dashboards as various file types is useful for sharing your data with users or including in your reports and presentations, however, Sisense makes it easy to share dashboards directly with other Sisense users by allowing you to share the dashboard's URL from the Sisense Web Application or sending your insights via email.

In addition, you can schedule automatic reports to be sent to yourself or other Sisense users. See the [Sending Email Dashboard](#) reports for more information.

The topics below describe how you can share widgets and dashboards with other Sisense users:

- [Exporting Widgets](#)
- [Exporting and Sharing Dashboards](#)
- [Creating PDF Reports](#)
- [Customizing PDF Reports](#)

# Managing Dashboards

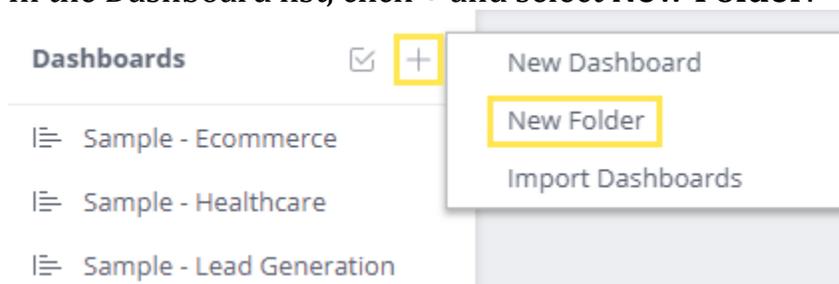
This page is only for Sisense users who have dashboard editing rights (Designers).

## Organizing Dashboards in Folders

You can organize your dashboards into folders and sub-folders within any of your folders to improve how your viewers navigate the Dashboard list.

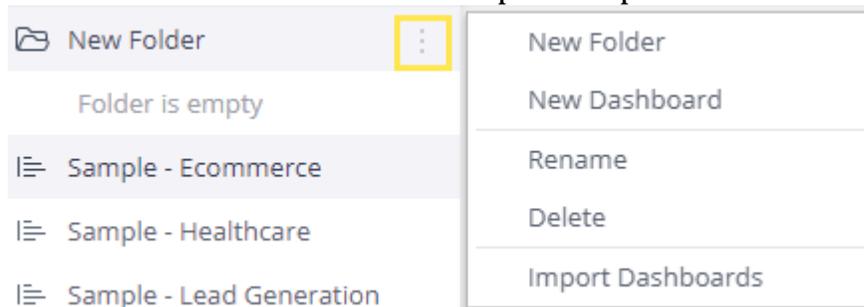
**To create a folder:**

1. In the Dashboard list, click **+** and select **New Folder**.



The new folder is added to the dashboards list.

2. Hover over the new folder and click **>** **Rename** to name the folder. To create additional sub-folders within your folders, hover over the folder and select **> New Folder** and repeat Step 2.

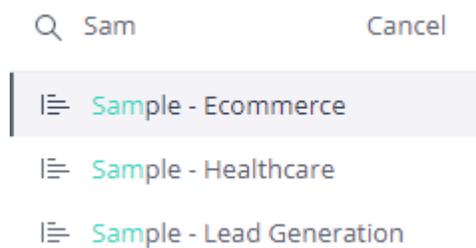


# Locating a Dashboard

You can quickly and easily locate dashboards without having to navigate through your entire list of dashboards. When you enter any characters in the **Search** field, the list of dashboards is filtered according to the characters found within your dashboard titles.

## To locate a dashboard:

1. In the Dashboard list, click  to display the **Search** field.
2. In the **Search** field, enter the title of your dashboard. Any dashboards that contain the characters you entered are displayed.



# Copying a Dashboard

Copying a dashboard (even one that was shared with you) creates a new copy of the dashboard with a different name. This new dashboard is totally separate from the original and you are its owner, meaning that you have full editing rights. When you copy a dashboard, the new dashboard is added below the copied dashboard in your Dashboard list with the same name plus a number in the order it was created.

For example, Sample Ecommerce(1) is a copy of Sample Ecommerce, while Sample Ecommerce(1)(1) is a copy of Sample Ecommerce(1):

Q Sam Cancel

☰ Sample - Ecommerce

☰ Sample - Healthcare

☰ Sample - Lead Generation

**To copy a Dashboard:**

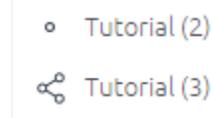
- Right-click the dashboard's name in the Dashboard list in the Home page and select **Duplicate**.

# Sharing Dashboards

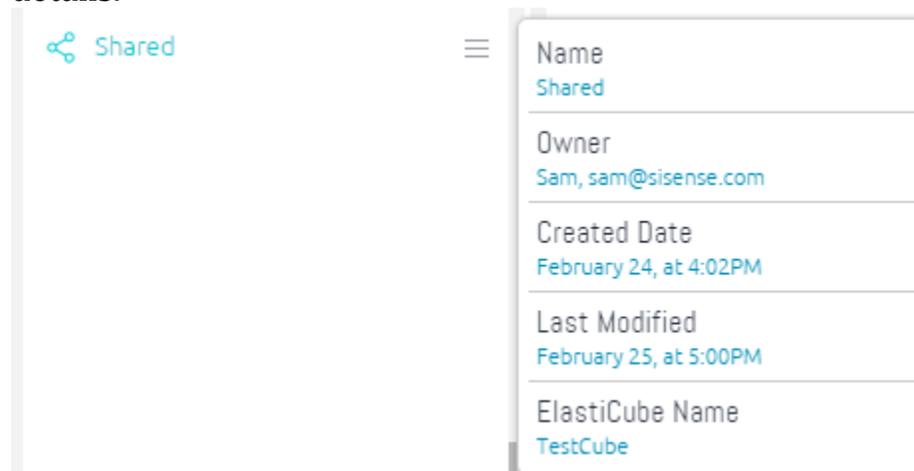
Owners can share dashboards that they have created with other Sisense users (Viewers and Designers). An owner is any Designer who has created a dashboard. When you share a dashboard, you can define whether the recipient has editing rights:

- Designers have editing rights.
- Viewers have viewing/filtering rights.

After you share a dashboard, or for any dashboard shared with you, in the Dashboards list the  icon appears next to the dashboard.



You can hover over the dashboard in the Dashboards list to view ownership details.



When you share your dashboard, you can configure Sisense to send periodic email notifications containing the dashboard inside the email's body to your dashboard's recipients.

# Working with Shared Dashboards

A dashboard recipient can edit the dashboard if you grant them permission, **but the changes they make in the dashboard do not affect your copy of the dashboard**. The recipient's changes are automatically performed on a separate copy of the dashboard (to which only they have access).

Dashboard recipients can revert to your latest shared copy at any time by right-clicking the dashboard's name in the Dashboards list on the left of the Sisense environment and selecting **Restore Dashboard**.

When you share a dashboard that's contained within a folder, the dashboard is shared using the same structure (under the same folder).

When a Designer shares a dashboard with you, you become a recipient and all of the above applies to you.

## To share a Dashboard:

1. In the heading of your dashboard, click the Share  button.  
OR  
In the Dashboards list, select **> Share** for the dashboard you want to share.
2. Enter the emails or groups of users with whom to share the dashboard.  
Selecting **Everyone** specifies that this dashboard is shared with all Sisense users and user groups.  
The contact information of the people you add is listed in the center of the window and is displayed the next time you open this window.  
If a person does not have a Sisense account, they receive a link that guides them through the activation process.
3. In the dropdown menu next to each user/user group select whether the user is a designer with editing rights (**Can edit**) or a viewer (**Can view**).
4. Toggle the Email Reports  button next to each user/user group on/off to define whether the user receives scheduled email reports of the dashboard.  
Read more on email reports below.
5. The **Dashboard URL** field at the bottom displays a direct link to this dashboard that you can copy/paste.
6. Click **Save**.

## Publishing Dashboards to Recipients

When an owner modifies a dashboard, for example, by adding a new widget, the dashboard is not automatically updated for recipients. The owner of a dashboard must publish their updated dashboard to update the dashboards of their recipients and the dashboard recipients must refresh the dashboard before any of the owner's changes are applied.

### To publish a dashboard to your recipients:

- In the heading of your dashboard, click **Republish**.

OR

In the Dashboards list, right-click the dashboard's name and select : > **Republish**.

## Updating a Shared Dashboard

When a recipient modifies a shared dashboard, those changes remain local unless the recipient updates their dashboard to match the dashboard of the owner. If the recipient updates their shared dashboard to match the dashboard of the owner, any of their local changes to the dashboard are overwritten. Sisense recommends that recipients save any local versions of a shared dashboard as a new dashboard before updating a dashboard.

### To update a shared dashboard:

1. In the heading of your dashboard, click : > **Restore Dashboard**.

OR

In the Dashboards list, right-click the dashboard's name and select : > **Restore Dashboard**.

2. Click **YES** to confirm that you want to update the shared dashboard.

## Changing Ownership of a Dashboard

The owner of a dashboard can transfer ownership of the dashboard to an Admin or another Designer. Transferring ownership of dashboards can be useful when

the original owner leaves your company and you do not want to lose their dashboards.

When you transfer ownership, the dashboard remains as a shared dashboard within your Dashboards list, and you become a dashboard recipient. As a dashboard recipient, you are limited to how you can modify the dashboard as described in **Working with Shared Dashboards**.

**To transfer ownership of a dashboard:**

1. In the heading of the dashboard you want to transfer ownership of, click the **Share** button.

OR

In the Dashboards list, select **> Share** for the relevant dashboard.

The Share window is displayed.

Share: Sample - Ecommerce

Enter names, email addresses or groups

**G** **Gideon** gideon@sisense.com Owner

Owner is not selected to receive scheduled reports. [Select](#)

Dashboard URL <http://10.20.4.15:8081/app/main#/dashboards/5aa66b91e67ded5f2424e0bf>

Save Cancel

2. Enter the name of the Administrator or Designer you want to transfer ownership to.

3. In the User Privilege list of the user you want to make an owner, select **Make Owner**.

Share: Sample - Ecommerce

Gideon  
gideon@sisense.com Owner

Gidi B  
gidib@sisense.com

- Can View
- Can View
- Can Design
- Make Owner

No users are selected to receive scheduled reports. [Select All](#)

Dashboard URL <http://10.20.4.15:8081/app/main#/dashboards/5aa66b91e67ded5f2424e0bf>

Save Cancel

4. Click **Save**. A Warning message is displayed.
5. Click **OK** to confirm that you want to transfer ownership to the user.

# Changing Sisense's Language

This topic discusses how Sisense Administrators can change the language of the Sisense Web Application for all of their users.

Administrators can set the default language for a Sisense server from the Admin section. If you are an Administrator and want to set the language for your system, users groups, or users, see [Changing the Sisense Web Application's Language](#).

If your users would like to set another language for the Sisense Web Application other than the language defined by the Administrator, they can select a language from their User Profiles or from the Login page.

Sisense supports the following languages:

- English
- Chinese (Simple)
- Dutch
- French (France)
- German
- Japanese
- Portuguese (BR)
- Russian
- Spanish (ES)
- Spanish (LA)
- Italian

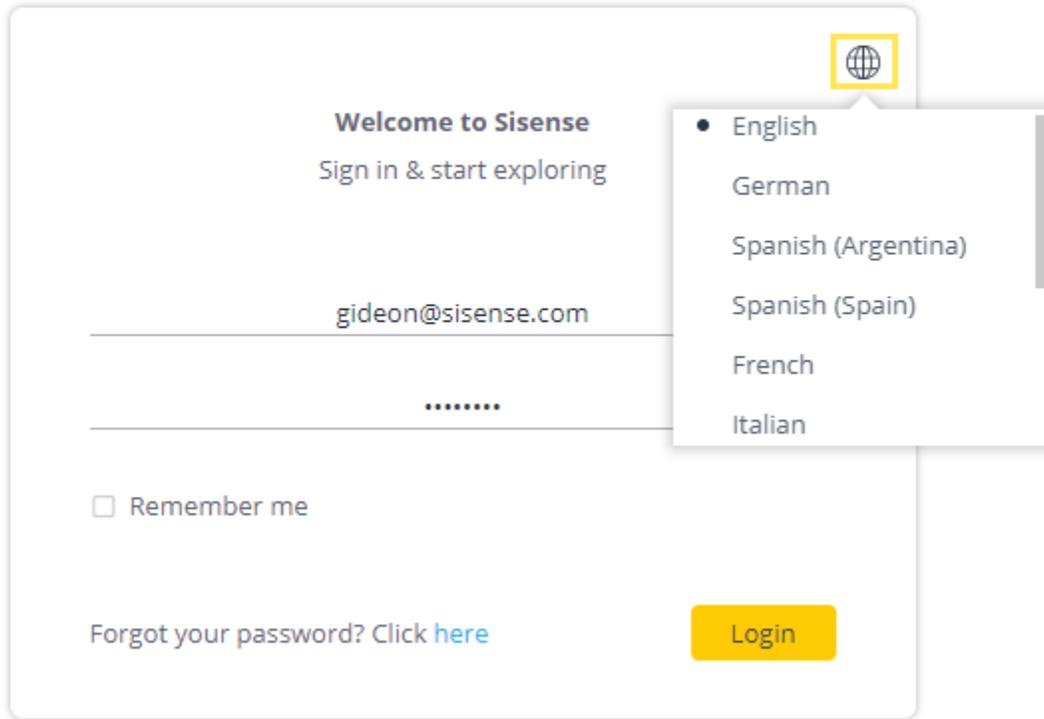
In addition, any languages that your company have translated will appear in the list according to the name defined in the translated system files.

Note, the following content is not being translated:

- Built-in Javascript editor
- ElastiCube Manager and Server Console
- Sisense documentation and online help
- REST API and documentation
- Sisense automated emails

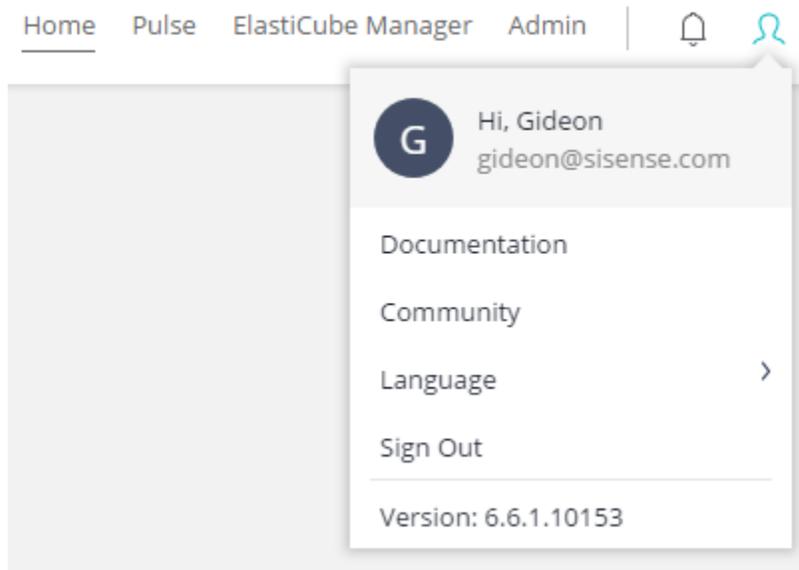
**To change your system's language:**

- Select the language from the Login page of the Sisense Web Application.



OR

In the Sisense Web Application User Options, open the languages list and select the language you want to display.



After setting the user's language, the Sisense Web Application will automatically be displayed in the selected language in future sessions.

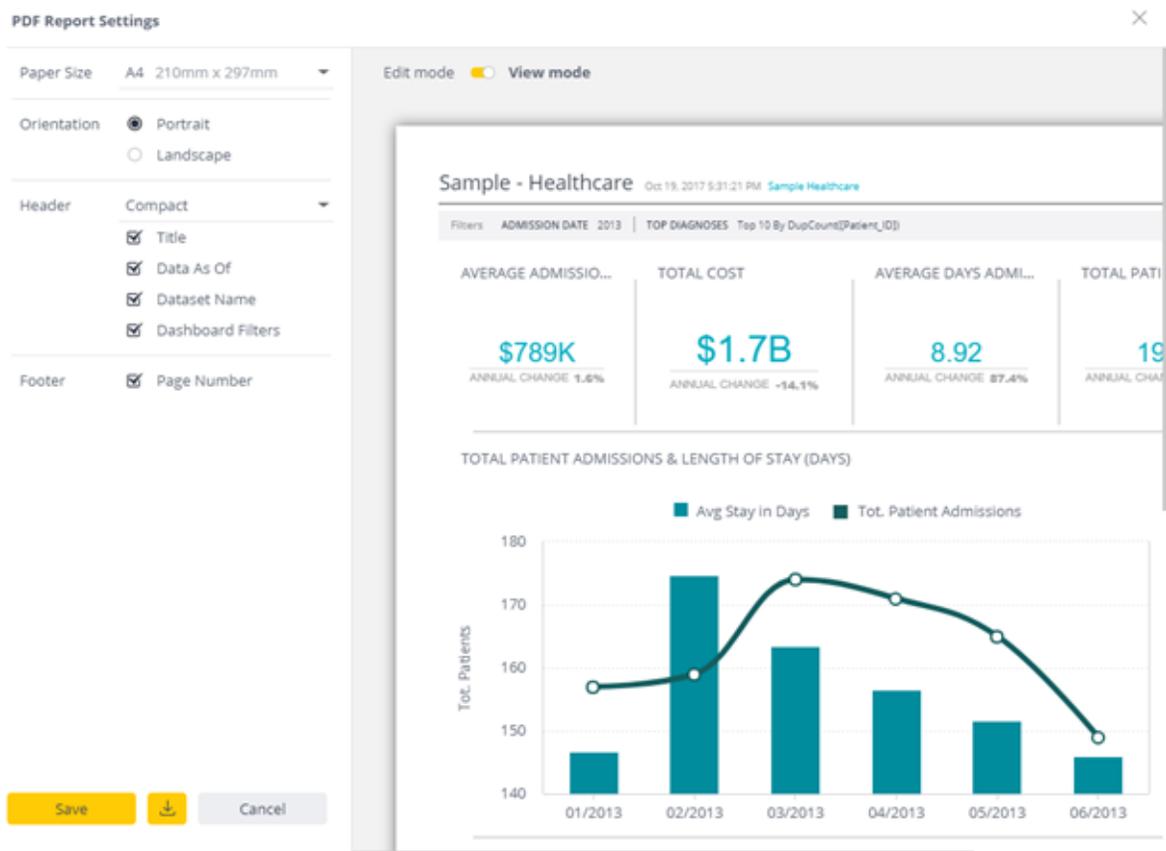
# Creating PDF Reports

When you need to take copies of your dashboards with you for meetings or sharing with others, you can generate a PDF report of your dashboard.

By clicking the  icon on your dashboard, you can customize your dashboard's appearance through the PDF Report Settings page. After you have defined how your dashboard is to be displayed, you can download your dashboard locally as a PDF.

## To create a PDF report:

1. From your dashboard menu, click the PDF icon . The PDF Report Settings page is displayed.
2. Customize how your dashboard is displayed in the PDF. For more information, see [Customizing PDF Reports](#).



**PDF Report Settings** ✕

Paper Size: A4 210mm x 297mm

Orientation:  Portrait  Landscape

Header: Compact

- Title
- Data As Of
- Dataset Name
- Dashboard Filters

Footer:  Page Number

Edit mode  View mode

**Sample - Healthcare** Oct 19, 2017 5:31:21 PM [Sample Healthcare](#)

Filters: ADMISSION DATE 2013 | TOP DIAGNOSES Top 10 By DupCount(Patient\_ID)

AVERAGE ADMISSIONS	TOTAL COST	AVERAGE DAYS ADMITTED	TOTAL PATIENTS
\$789K ANNUAL CHANGE 1.6%	\$1.7B ANNUAL CHANGE -14.1%	8.92 ANNUAL CHANGE 87.4%	19 ANNUAL CHANGE

TOTAL PATIENT ADMISSIONS & LENGTH OF STAY (DAYS)

■ Avg Stay in Days ■ Tot. Patient Admissions

Month	Tot. Patient Admissions	Avg Stay in Days
01/2013	148	158
02/2013	175	159
03/2013	165	175
04/2013	158	172
05/2013	152	165
06/2013	145	148

Save  Cancel

3. Click . The dashboard is downloaded locally as a PDF file.

# Customizing PDF Reports

If you need to share dashboards with other users, or create a hardcopy for yourself, Sisense allows you to send an email report to your users or export your dashboard to PDF. Dashboards in email reports or PDFs though look different compared to online dashboard as they serve different purposes. Through the Sisense PDF Report Settings, you can customize and create formatted email reports and PDFs according to your requirements. The Sisense PDF Report Settings allow you to define both the content and design of your report quickly and easily.

**PDF Report Settings** [Close]

Paper Size: A4 210mm x 297mm

Orientation:  Portrait  Landscape

Header: Compact

- Title
- Data As Of
- Dataset Name
- Dashboard Filters

Footer:  Page Number

Edit mode  View mode

**Sample - Healthcare** Oct 19, 2017 5:31:21 PM Sample Healthcare

Filters: ADMISSION DATE 2013 | TOP DIAGNOSES Top 10 By DupCount(Patient\_ID)

AVERAGE ADMISSIO...	TOTAL COST	AVERAGE DAYS ADML...	TOTAL PATI
<b>\$789K</b> ANNUAL CHANGE 1.6%	<b>\$1.7B</b> ANNUAL CHANGE -14.1%	<b>8.92</b> ANNUAL CHANGE 87.4%	<b>19</b> ANNUAL CHA...

TOTAL PATIENT ADMISSIONS & LENGTH OF STAY (DAYS)

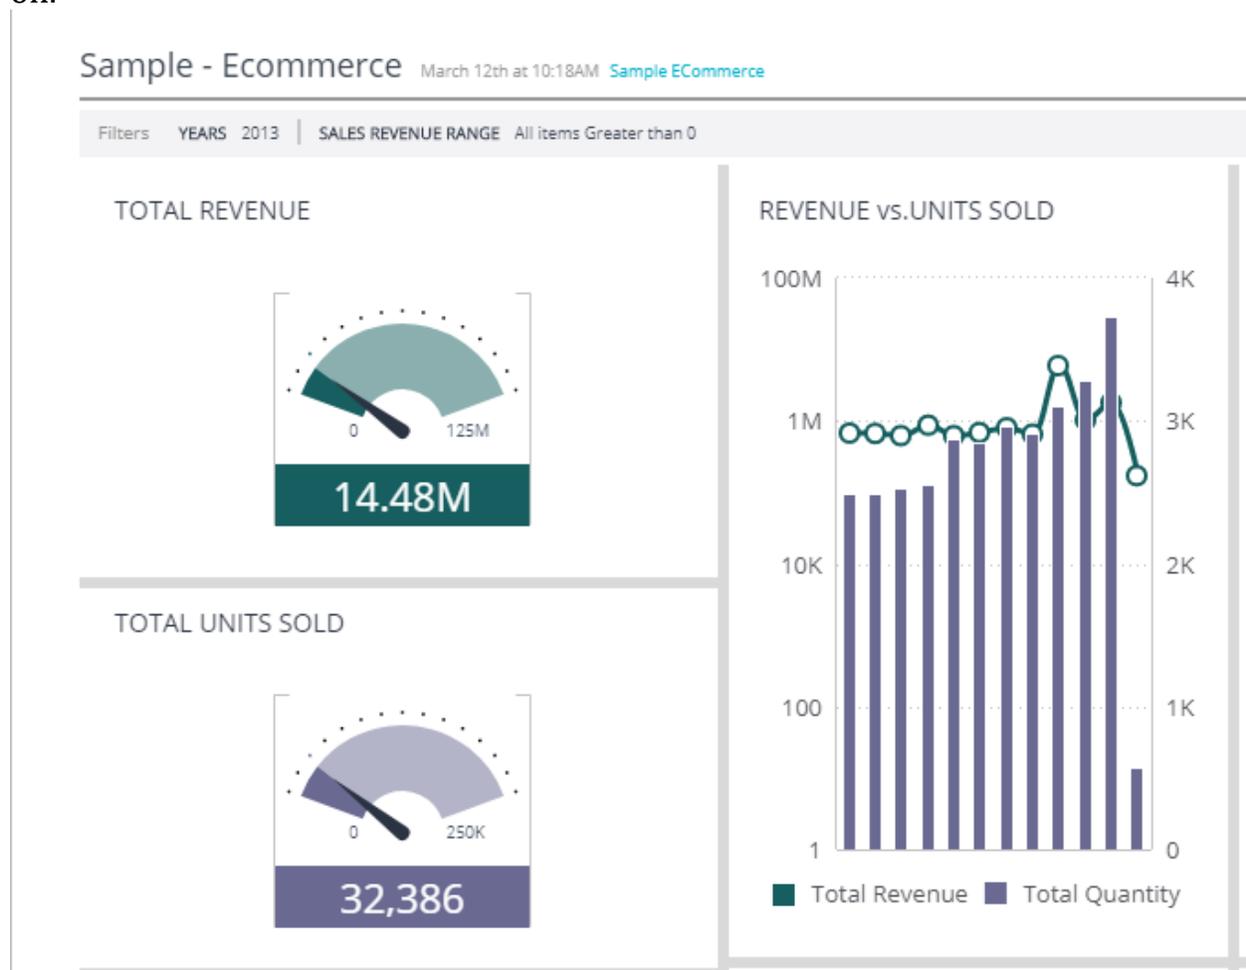
■ Avg Stay in Days ■ Tot. Patient Admissions

Month	Tot. Patient Admissions	Avg Stay in Days
01/2013	~148	~158
02/2013	~175	~160
03/2013	~165	~175
04/2013	~158	~172
05/2013	~152	~168
06/2013	~145	~150

Save [Download] Cancel

From the PDF Report Settings page, you have two modes, Edit mode, for customizing your PDF, and View mode, for seeing how the dashboard will be displayed in the PDF. In Edit Mode, all the settings you can apply to your report are displayed in the left menu.

In Edit mode, you can set the size of your widgets by selecting and dragging the borders of each widget. These borders are displayed when Edit mode is toggled on.



In View and Edit mode, you can define your dashboard's orientation, its layout, and what is displayed in the header and footer on your report.

After you have customized your report, you can save the settings and use them when sharing reports, or download the report as a PDF.

**To customize a PDF report:**

1. In your dashboard's menu, click . The PDF Report Settings page is displayed.
2. Edit the dashboard by selecting any of the following options:

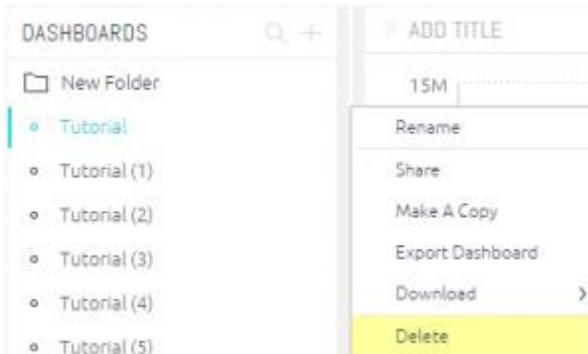
Setting	Description
Paper Size	The size of the email report in the PDF.
Orientation	The orientation of the dashboard in the PDF, landscape or portrait. Landscape orientations display the dashboard horizontally while portrait orientations display the dashboard vertically.
Header	<p>The design and content of your report's header. The header is displayed on the top of each page or your PDF report.</p> <p>From the Header list, you can define how your header is to be displayed: Compact, Medium, Large, or No Header.</p> <p>After you have set the design, you can determine what information is included in the header.</p> <p><b>Title:</b> Select to display the dashboard name at the top of the PDF.</p> <p><b>Dashboard As Of:</b> Select to display the ElastiCube name and last build time at the top of the PDF.</p> <p><b>Dataset Name:</b> Select to display the name of the ElastiCube that contains the dashboard's data.</p> <p><b>Dashboard Filters:</b> Select to display dashboard filter selections.</p>
Footer	<p>The content displayed in the footer section of each page of your report. You can display the following information:</p> <p><b>Page Number:</b> Select to display the page number on each page of the report.</p>

3. Click **Save** to save your settings when sharing reports or  to download your report.

# Deleting Dashboards

## To delete a dashboard:

- Click on the menu ☰ of the Dashboard to be deleted in the list in the Home



page, and select **Delete**.

OR

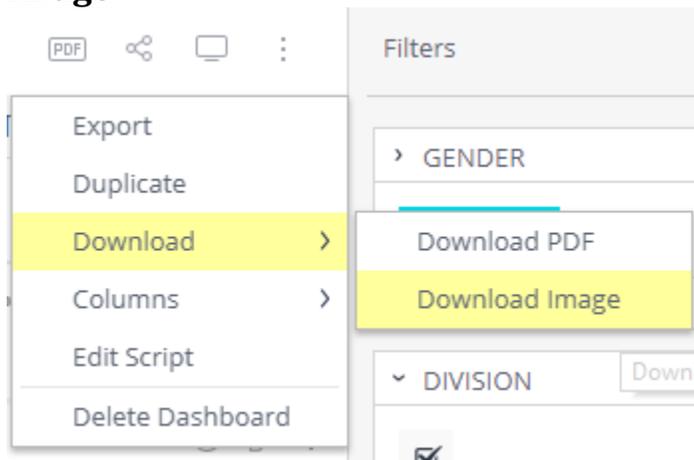
Click the Dashboard Menu ☰ and select **Delete Dashboard**.

# Downloading a Dashboard as an Image

You can download the dashboard as an image. This can be useful for including the dashboard in a presentation, web site, etc.

## To download the image:

1. Click  at the top right of the dashboard, and select **Download > Download Image**.



2. You can also define the following additional settings:
  - **Show Dashboard Title:** Turning this on will display the dashboard name at the top of the PDF.
  - **Show Dashboard Filters:** Turning this on will display a summary of the dashboard filter selections at the top of the PDF.
  - **Show ElastiCube Info:** Turning this on will display the ElastiCube name and last build time at the top of the PDF.

## Download Image

Size (Width)

Current (1439 Pixels) ▾

Data

- Show Dashboard Title
- Show Dashboard Filters
- Show Data Set Info

Download

Cancel

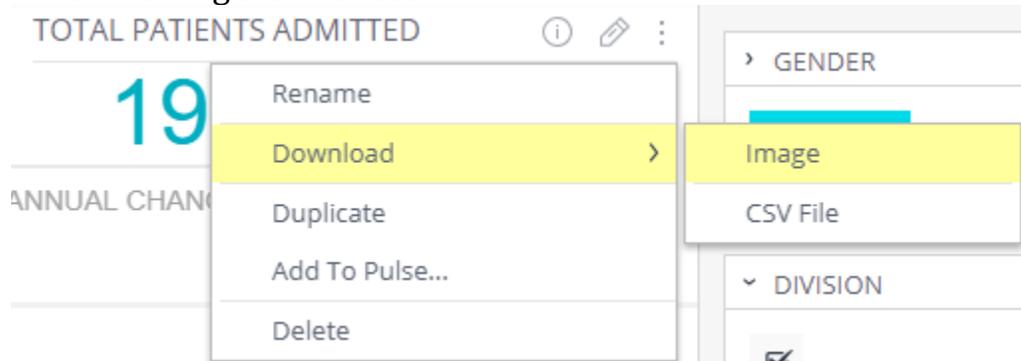
# Downloading a Widget as an Image

You can download a widget as an image (PNG format).

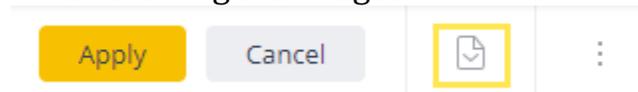
## To download the widget as an image:

- In dashboard mode, click on the widget's menu, and select **Download > Image**.
- In edit widget mode, click on the download icon, and select **Download Image**.

When viewing the dashboard:



When viewing the widget:

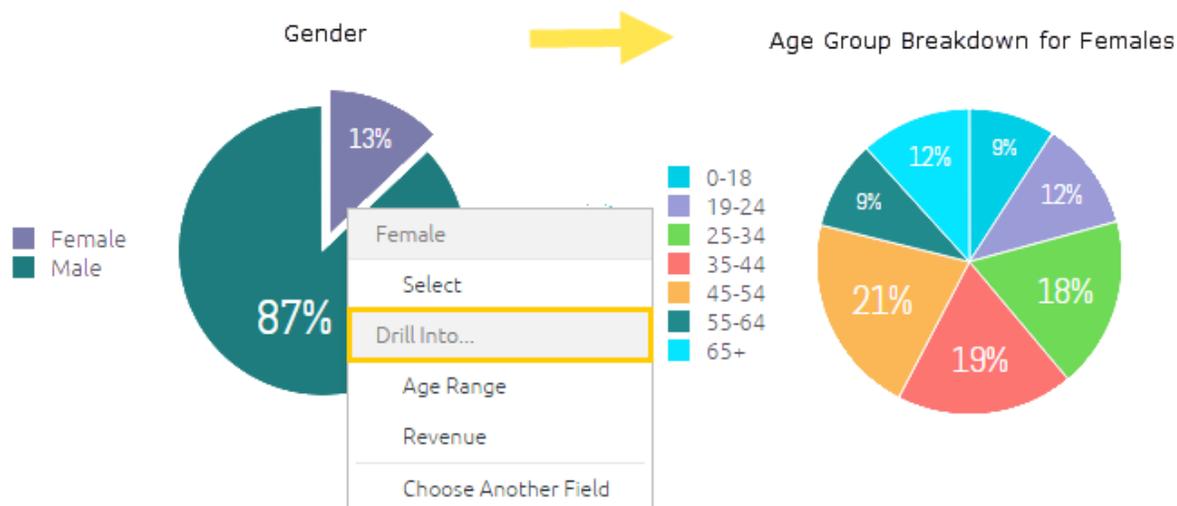


**Tip:** The image size will reflect the size of the widget on the screen when you download it. To create a larger image, we recommend enlarging the widget space in the dashboard, or to open the widget in edit view, to get the bigger picture.

# Drilling Down in a Widget

This topic describes how you can drill down into your data in most widgets to get an in-depth view of a selected value.

In the following example, the pie chart on the left shows a breakdown by gender. This is the original chart as created by the dashboard's designer. On the right side, the chart shows a breakdown or drill down showing age groups of the 'female' segment from the original pie chart.

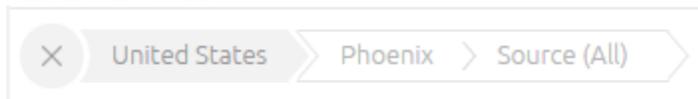


With Sisense you can drill down from any field to any field, unless disabled by the dashboard's Designer.

## To drill down into a chart:

1. Right-click on the item in the widget into which you want to drill down.
2. To manually select the drill hierarchy path, select **Drill Into**, and then in the Data Browser, select the field into which to drill down. If you have already drilled into this chart, then you will have shortcuts to previously selected fields, or select **Choose Another Field** to select a different field for the first time.  
or  
Select a predefined drill hierarchy (if available).

To drill up to a higher level, click on a breadcrumb. To drill all the way up, click on the X icon.



# Email Troubleshooting

Users of Sisense receive a variety of automated emails from Sisense. Sisense emails fall into two categories of emails:

**System Emails:** These emails include activation emails, password recovery emails, and updates regarding published dashboards.

**Email Reports:** These emails are dashboard reports in PDF/PNG format that are sent manually or according to a schedule you define.

The following steps will help you troubleshoot email-related issues such as emails not being sent/received.

Step 1 – [Verify that Emails are Configured Correctly](#)

Step 2 – [Send a Test Email](#)

Step 3 – [Export the Dashboard to PDF](#)

Step 4 – [Verify System Load](#)

Step 5 – [Verify Port Bindings](#)

Step 6 – [Custom Email Server](#)

Step 7 – [Verify the default.yaml Settings](#)

Step 8 – [Check Rebranded Email Settings](#)

Step 9 – [Verify Spam Protection](#)

## Step 1 – Verify that Emails are Configured Correctly

You may not receive emails from Sisense due to configuration issues. There are four potential configuration issues that can occur.

1. Activate Sisense Emails: (All email types) Make sure that Sisense emails have been enabled within the System Configuration page of the Admin

menu.

The screenshot shows the Sisense Admin interface. At the top left is the Sisense logo. At the top right are navigation links: Home, ElastiCube Manager, Admin, and a search icon. A left sidebar menu contains the following items: ADMIN, USERS, GROUPS, ELASTICUBES, ACTIVE DIRECTORY, SINGLE SIGN ON, SYSTEM CONFIGURATION (highlighted with a dark bar), and REST API. The main content area is titled 'SYSTEM CONFIGURATION' and contains the 'SISENSE WEB SERVER' settings. These settings include: 'Alias' (text input with placeholder 'Server public IP/Domain (ex: www.company.com)'), 'FIRST DAY OF WEEK' (dropdown menu set to 'Monday'), 'DEFAULT LANGUAGE' (dropdown menu set to 'English'), 'SSL' (toggle switch turned off), and 'SEND EMAILS' (toggle switch turned on, highlighted with a yellow box). Each setting has an information icon (i) to its right. An 'EDIT' button is located at the bottom right of the settings area.

2. Verify that the user you want to share the report with is selected: (Email reports only) You can verify which users are selected by checking the  icon associated with their name. If the  icon is not highlighted, the user will not receive an email report. You can select the user by clicking the icon.

## Share: Sample - Ecommerce



G

Gideon  
gideon@sisense.com

Owner



S

Sam Smith  
sam@sisense.com

Resend invitation

Can View



Some users are selected to receive scheduled reports.

[Select All](#)

Dashboard URL

<http://192.168.5.222:8081/app/main#/dashboards/57c2e5f75f19bc58>

SAVE

CANCEL

If you are sharing your reports with a group, verify that all the members are included. An Administrator can check the list of users in a group from the Admin page by **Selecting Admin > Groups > Edit**.

### ADMIN

- USERS
- GROUPS **2****
- ELASTICUBES
- ACTIVE DIRECTORY
- SINGLE SIGN ON
- SYSTEM CONFIGURATION
- REST API

### GROUPS

ADD GROUP **+**

<input type="checkbox"/>	Group Name	Number of Users	Origin	
<input type="checkbox"/>	Admins	1	Sisense	
<input type="checkbox"/>	Everyone	4	Sisense	
<input type="checkbox"/>	Group 1	2	Sisense	<b>3</b>

Page 1 of 1      << < 1 > >>      3 Groups Total

## Modify Group

Group Name

Group 1|

Group Users

+ Add Group Users



gideon@sisense.com



sam@sisense.com



SAVE

CANCEL

3. **Data Security:** (Email reports) Verify with your Sisense Administrator that the user you want to share the report with has access rights to the dashboard's ElastiCube.
4. **Verify Build Success:** (Email reports only) If you send a report following each build, verify that dashboard's ElastiCube was built successfully.

## Step 2 – Send a Test Email

If you do not receive any emails from Sisense, email reports or system emails, you can verify whether or not you can receive emails from Sisense by sending a report to yourself. To send an email report, click the  icon of a dashboard and in the bottom-left corner of the window, click **Send me a report** now to test the email's functionality. If you receive an email, but do not receive your scheduled reports, see Contacting Technical Support.



Email Report Type

Email report only

PDF Settings	<input checked="" type="checkbox"/> Dashboard Title	Layout	Paper Size	Orientation
	<input checked="" type="checkbox"/> Dashboard Filters	As Is	A4	Landscape
	<input checked="" type="checkbox"/> ElastiCube Info			

Schedule

 Every ElastiCube update, no more than  emails per day Daily:  at  [Send me a report now](#)

SAVE

CANCEL

## Step 3 – Export the Dashboard to PDF

This step is relevant for email reports. If your issue is related to system emails, goto Step 5.

If you cannot receive any email reports, this may indicate a conflict with your plugins. By exporting your dashboard to PDF, you can verify if email reports are being sent and if not, what is the source of the problem.

1. **Check Plugins:** You can check if your plugins are interfering with your email reports by renaming the plugin folder located at C:\Program Files\Sisense\PrismWeb\plugins to another name, for example, to “plugin1”. If the email succeeds, this indicates that one of your plugins is interfering with your emails and needs to be corrected.

2. **Check SSL:** After checking your plugins, if you still do not receive dashboard reports, this may indicate a problem with SSL. If you are using SSL, follow the instructions in this [forum post](#).

## Step 4 – Verify System Load

This step is relevant for email reports. If your issue is related to system emails, goto Step 5.

Verify that the system is not overloading by sending too many reports at the same time. For example, if your email reports are all scheduled for 5:00pm, this may overload the email server. You can verify if this is causing problems by scheduling email loads throughout the day.

## Step 5 – Verify Port Bindings

Verify that your Port bindings are configured properly by using an HTTP port that is not used by other web sites on your IIS.

If you have a DNS configured make sure you have the required certificate for it.

## Step 6 – Custom Email Server

This step is relevant for all email types.

By default, Sisense uses Mandrill for emailing services. Verify that you can reach mandrill.com from the server. If you can't, this might indicate a network security restriction. Verify that the IP addresses listed here are whitelisted. If your network security restrictions prevent you from communicating with Mandrill email services, you can implement your own custom email server, see [Setting Up a Custom Email Server](#).

## Step 7 – Verify the default.yaml Settings

The default.yaml settings file is located in: C:\Program Files\Sisense\PrismWeb\vnext\config. These settings define which path Sisense uses for exporting emails.

'webServer' -> iisPort should be the same port as the binding in the IIS.

```
58 webServer:
59   ver: '5.0'
60   port: 3000
61   iisPort: 8081
```

'exporting' -> host = "localhost", port = "default".

Note that the 'exporting' address should be accessible from inside the server, i.e. if you can't access localhost from your browser, but you can access it using a DNS, change the export settings to use the DNS.

```
101 exporting:
102   host: "localhost"
103   port: "default"
104   protocol: "http"
```

## Step 8 – Check Rebranded Email Settings

This step is relevant for all email types.

If you have changed the default Sisense configuration, for example, when rebranding Sisense, run the following requests through REST API admin page

- GET /branding
- GET /settings/system

Verify the 'senderEmail' is configured correctly and not listed as 'string' or another erroneous sender. This can happen when someone uses POST /branding and does not enter a valid email address or leaves the value as 'string'.

If you are whitelabeling Sisense and rebranding your emails and experience email problems, review the following steps as described here.

## Step 9 – Verify Spam Protection

This step is relevant for all email types.

Make sure that emails sent from Sisense are whitelisted and not filtered by any spam protection. If you do not have access to your company's Spam folder or Sisense emails may be filtered before arriving to your inbox, consult with your company's IT administrator to verify that Sisense emails are not being filtered by your company's Spam filters.

If none of the above steps have resolved the issue please contact Sisense Technical Support. When contacting Sisense Technical Support, please provide the following information to reduce the amount of time it takes to resolve your issue:

- **Time:** When the issue occurred.
- **Logs:** When encountering an error while trying to send email reports, Sisense maintains error logs that you can review or send to Sisense Technical Support. The logs are located at ...Program Files\Sisense\PrismWeb\vnext\iisnode.
- **Steps Taken:** List each of the steps you have taken to when trying to resolve the issue yourself.

# Exporting and Sharing Dashboards

You can share your Sisense dashboards and widgets with other Sisense users to distribute your Sisense insights.

The topics below describe how to share your dashboards in a variety of formats:

- [Exporting Pivot Tables to PDF](#)
- [Exporting Pivot Tables to Excel](#)
- [Export Tables to PDF](#)
- [Exporting CSVs](#)
- [Downloading Dashboards](#)
- [Downloading Widgets as Images](#)
- [Sharing Dashboards](#)
- [Sending Email Dashboard Reports](#)
- [Email Troubleshooting](#)

# Exporting Pivot Tables to Excel

You can export your pivot tables to Excel (XLSX). When you export a pivot table to Excel, the data that is exported is the data that is displayed in your widget. This means that any active filters, layout structure, and masks that you have applied in Sisense to your data are exported as well. For example, if you have modified how currency, percentages, or numbers are formatted in the widget in Sisense, these will be reflected in your exported Excel file.

The example below shows a pivot table exported to Excel versus CSV:

## Excel

Years in Quarters		AdWords				Email				Social Media				Webinar							
		Diff	Past	Period	Total #	Visi	Total #	Co	Total #	Co	Diff	Past	Period	Total #	Visi	Total #	Co	Total #	Co	Total #	Co
2013	2013 Q1	61150.00	61150.00	282.04	6421.30	490.00	490.00	4.41	23.70	1306.00	1306.00	785.28	1943.15	1675.00	1675.00	120.38	314.50				
	2013 Q2	-11644.00	49506.00	1017.49	5934.20	1985.00	2475.00	4.42	358.55	181713.00	183019.00	184.01	11467.75	11477.00	13152.00	10.02	1327.30				
	2013 Q3	50706.00	100212.00	188.08	4275.15	17212.00	19687.00	3.09	1340.40	-139697.00	43322.00	4.60	1996.65	236170.00	249322.00	15.37	10360.20				
	2013 Q4	-76158.00	24054.00	199.50	1583.95	-2086.00	17601.00	1.65	1182.75	-17549.00	25773.00	10.06	2412.45	-152716.00	96606.00	4.89	4131.00				
Total 2013		234922.00	234922.00	1687.11	18214.60	40253.00	40253.00	13.57	2905.40	253420.00	253420.00	983.94	17820.00	360755.00	360755.00	150.65	16133.00				
2014	2014 Q1	11470.00	35524.00	795.96	3022.35	-17575.00	26.00	0.19	2.40	-25773.00				-92826.00	3780.00	14.12	163.20				
	2014 Q2	48696.00	84220.00	284.37	6432.80	59137.00	59163.00	5.49	7588.35	14012.00	14012.00	333.25	6425.45	5482.00	9262.00	185.72	1201.35				
	2014 Q3	-25576.00	58644.00	190.74	3559.80	-28778.00	30385.00	3.46	3756.75	-13828.00	184.00	1.50	11.10	-708.00	8554.00	18.00	891.60				
	2014 Q4	8206.00	66850.00	181.44	4131.75	-30362.00	23.00	0.16	0.75	2701.00	2885.00	5.00	65.55	-2173.00	6381.00	14.30	227.70				
Total 2014		10316.00	245238.00	1452.51	17146.70	49344.00	89597.00	9.30	11348.25	-236339.00	17081.00	339.75	6502.10	-332778.00	27977.00	232.14	2483.85				

## CSV

	A	B	C	D	E	F	G	H	I	J
1	Years in Date	Quarters	Source	DiffPastPe	Total # Vis	Total # Co	Total # Converted			
2	2013	2013 Q1	AdWords	61150	61150	282.0375	6421.3			
3	2013	2013 Q1	Email	490	490	4.405235	23.7			
4	2013	2013 Q1	Social Me	1306	1306	785.275	1943.15			
5	2013	2013 Q1	Webinar	1675	1675	120.3757	314.5			
6	2013	2013 Q2	AdWords	-11644	49506	1017.49	5934.2			
7	2013	2013 Q2	Email	1985	2475	4.42218	358.55			
8	2013	2013 Q2	Social Me	181713	183019	184.0085	11467.75			
9	2013	2013 Q2	Webinar	11477	13152	10.0205	1327.3			
10	2013	2013 Q3	AdWords	50706	100212	188.0823	4275.15			
11	2013	2013 Q3	Email	17212	19687	3.094383	1340.4			
12	2013	2013 Q3	Social Me	-139697	43322	4.601848	1996.65			
13	2013	2013 Q3	Webinar	236170	249322	15.36633	10360.2			
14	2013	2013 Q4	AdWords	-76158	24054	199.5022	1583.95			
15	2013	2013 Q4	Email	-2086	17601	1.646973	1182.75			
16	2013	2013 Q4	Social Me	-17549	25773	10.05553	2412.45			

Exporting to Excel maintains your pivot’s layout and structure, and values such as sub-totals, which are not maintained when you export a pivot table to CSV.

**To export a widget to Excel:**

- In Dashboard mode, click on the widget’s menu, and select **Download > Excel File**.
- In Edit mode, click on the download icon, and select **Download Excel File**.

## Exporting to Excel via the REST API

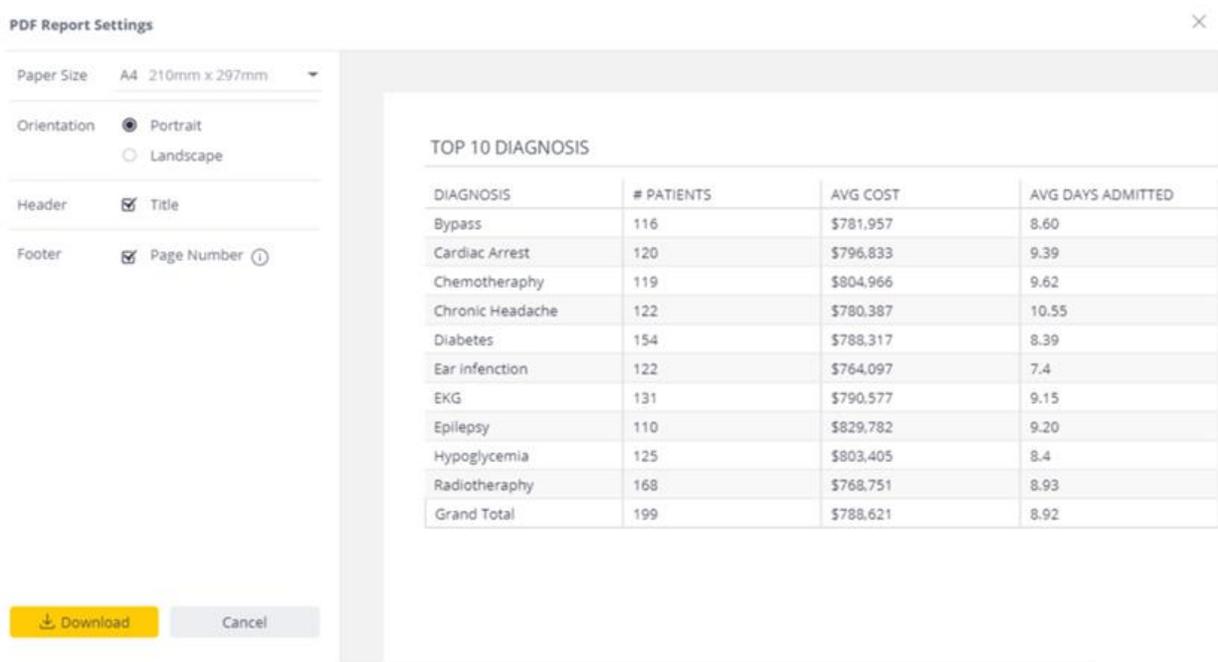
In addition to exporting to Excel via the Sisense Web Application, you can export binary content via the Sisense REST API through the POST /engine/excelExport endpoint. You can then convert this content to an Excel file (.xlsx).

**Note:** This endpoint is not available from Sisense's interactive REST API. For more information about the endpoint, see the POST /engine/excelExport in [Version 1.0 of the REST API](#).

# Exporting Table Widgets to PDF

You can export table widgets in PDF format. When you export your table to PDF, the entire table is included.

Through the PDF Report Settings page, you can customize how your table is displayed in your PDF report including the table's orientation and page size.



PDF Report Settings

Paper Size: A4 210mm x 297mm

Orientation:  Portrait,  Landscape

Header:  Title

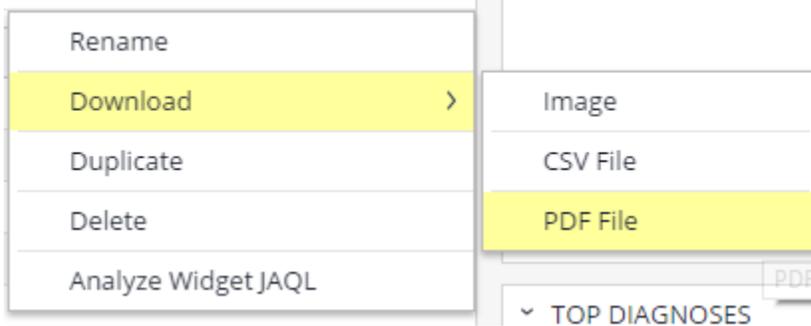
Footer:  Page Number

Download Cancel

DIAGNOSIS	# PATIENTS	AVG COST	AVG DAYS ADMITTED
Bypass	116	\$781,957	8.60
Cardiac Arrest	120	\$796,833	9.39
Chemotherapy	119	\$804,966	9.62
Chronic Headache	122	\$780,387	10.55
Diabetes	154	\$788,317	8.39
Ear Infection	122	\$764,097	7.4
EKG	131	\$790,577	9.15
Epilepsy	110	\$829,782	9.20
Hypoglycemia	125	\$803,405	8.4
Radiotherapy	168	\$768,751	8.93
Grand Total	199	\$788,621	8.92

**To export a Table widget to PDF:**

1. In the Widget menu, select **Download > PDF File**.

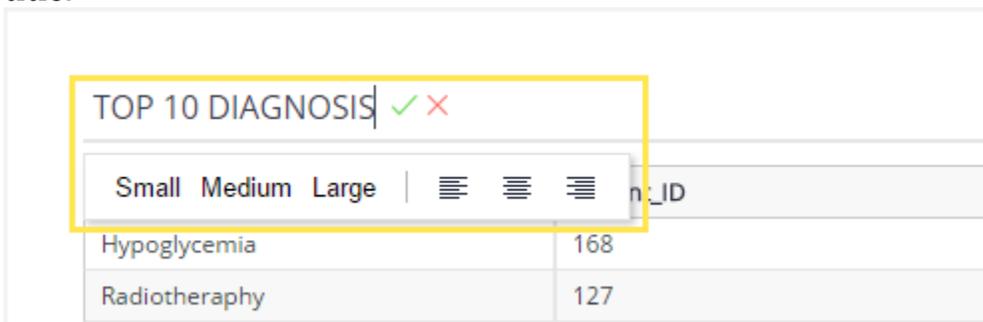


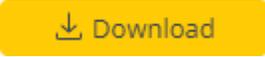
The PDF Report Settings page is displayed.

2. You can apply the following settings:

Setting	Description
Paper Size	The page size of the PDF report.
Orientation	The orientation of the dashboard in the PDF, landscape or portrait. Landscape orientations display the dashboard horizontally while portrait orientations display the dashboard vertically.
Header	Select Title to display the widget's title in the header of the PDF report.
Footer	Select Page Number to display page numbers for each page of the PDF report.

3. In the Preview area, click the title to edit the text, size, and alignment of the title.



4. Click the  to download your report.

# Exporting Widgets to CSV/Excel

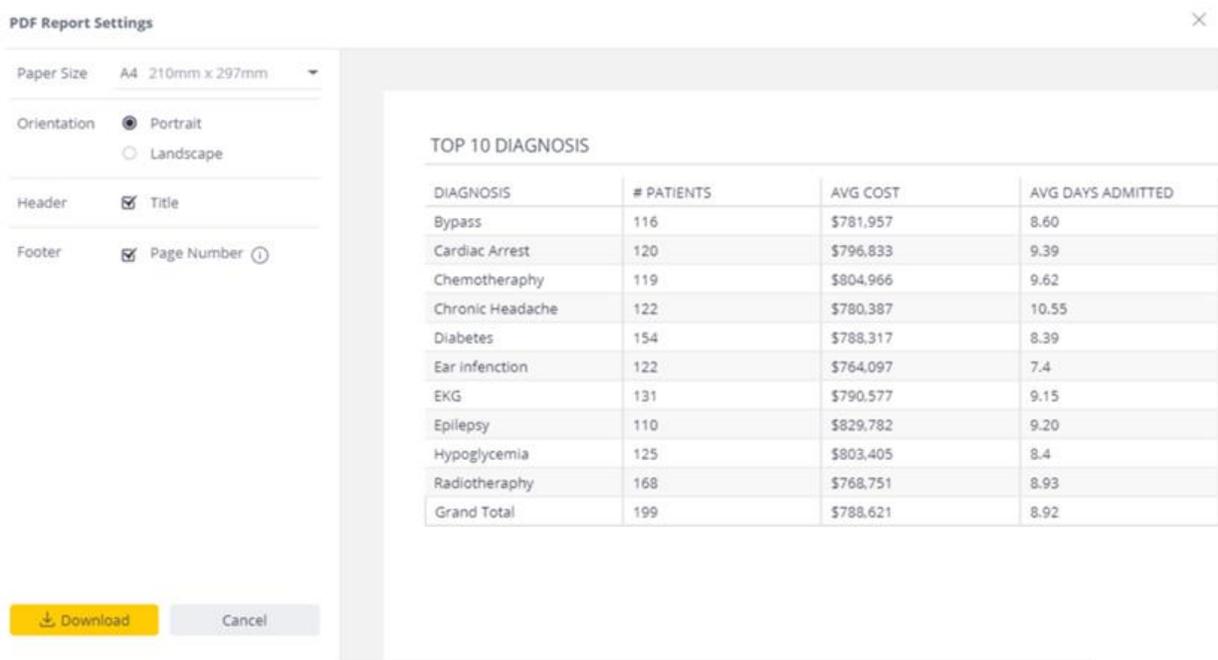
You can export widgets in the CSV format. This is also useful for using with Excel. Only the displayed (filtered) data is exported according to the currently selected filters.

## To export a widget to CSV:

- In dashboard mode, click on the widget's menu, and select **Download > CSV File**.
- In edit widget mode, click on the download icon, and select **Download CSV**.

# Exporting Pivot Tables to PDF

You can export pivot table widgets in PDF format. Through the PDF Report Settings page, you can customize how your table is displayed in your PDF report including the table's orientation and page size. In the PDF Report Settings, the first 14 pages of your Pivot table are displayed in the preview window, however, when you export your Pivot table to PDF, the entire table is included, up to 10,000 rows on multiple pages.



PDF Report Settings

Paper Size: A4 210mm x 297mm

Orientation:  Portrait  Landscape

Header:  Title

Footer:  Page Number ⓘ

Download Cancel

DIAGNOSIS	# PATIENTS	AVG COST	AVG DAYS ADMITTED
Bypass	116	\$781,957	8.60
Cardiac Arrest	120	\$796,833	9.39
Chemotherapy	119	\$804,966	9.62
Chronic Headache	122	\$780,387	10.55
Diabetes	154	\$788,317	8.39
Ear infection	122	\$764,097	7.4
EKG	131	\$790,577	9.15
Epilepsy	110	\$829,782	9.20
Hypoglycemia	125	\$803,405	8.4
Radiotherapy	168	\$768,751	8.93
Grand Total	199	\$788,621	8.92

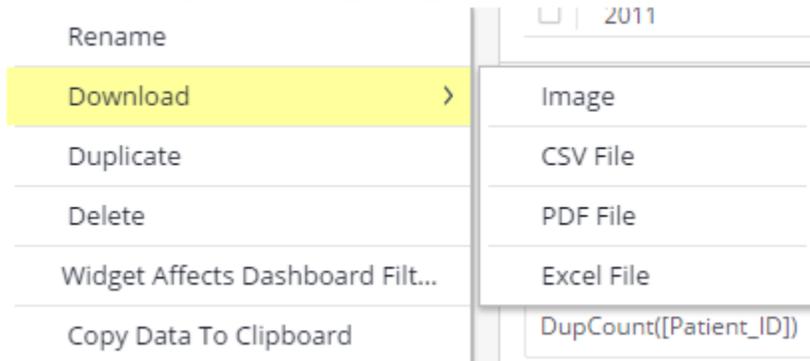
**Note:** The following plug-ins are not supported when exporting Pivot tables into PDF reports:

- Embed images to Pivot tables
- Conditionally format Pivot foreground
- Embed a sparkline chart into a Pivot

**To export a Pivot table widget to PDF:**

1. In the Widget menu, select **Download > PDF File**.

2. The PDF Report Settings page is displayed.



3. You can apply the following settings:

Setting	Description
Paper Size	The size of the email report in the PDF.
Orientation	The orientation of the dashboard in the PDF, landscape or portrait. Landscape orientations display the dashboard horizontally while portrait orientations display the dashboard vertically.
Header	<p>The design and content of your report's header. The header is displayed on the top of each page or your PDF report.</p> <p>From the Header list, you can define how your header is to be displayed: Compact, Medium, Large, or No Header.</p> <p>After you have set the design, you can determine what information is included in the header.</p> <p><b>Title:</b> Select to display the dashboard name at the top of the PDF.</p> <p><b>Dashboard As Of:</b> Select to display the ElastiCube name and last build time at the top of the PDF.</p> <p><b>Dataset Name:</b> Select to display the name of the ElastiCube that contains the dashboard's data.</p> <p><b>Dashboard Filters:</b> Select to display dashboard filter selections.</p>
Footer	The content displayed in the footer section of each page of your report. You can display the following information:

Setting	Description
	Page Number: Select to display the page number on each page of the report.

- In the Preview area, click the title to edit the text, size, and alignment of the title.

TOP 10 DIAGNOSIS ✓ ✗

Small Medium Large | 
 ☰ ☷ ☷

Bypass	8
Cardiac Arrest	7

- Click **Download** to download your report.

# Exporting Widgets

You can download widgets in one of the following formats depending on the widget type:

- **CSV:** Download the data described by the widget in a CSV file. If you have applied any filters, the filtered data is the data included in the file.
- **Excel:** Download the data described by the widget in an Excel file. When you export a pivot table to Excel, the data that is exported is the data that is displayed in your widget. This means that any active filters, layout structure, and masks that you have applied in Sisense to your data are exported as well. For example, if you have modified how currency, percentages, or numbers are formatted in the widget in Sisense, these will be reflected in your exported Excel file. Exporting to Excel maintains your pivot's layout and structure, and values such as sub-totals, which are not maintained when you export a pivot table to CSV.
- **Image:** You can download a widget as an image in PNG format. The image size will reflect the size of the widget on the screen when you download it. To create a larger image, enlarge the widget space in the dashboard, or open the widget in edit view, to get a bigger picture.
- **PDF:** Download the widget as a PDF. For Table widgets, you can customize how your table is displayed in your PDF including the table's orientation and page size. In the PDF Report Settings, the first 14 pages of your Pivot table are displayed in the preview window, however, when you export your Pivot table to PDF, the entire table is included, up to 10,000 rows on multiple pages. For more information, see [Customizing PDF Reports](#).

## To download a widget:

- In dashboard view, click on the widget's menu, and select **Download** and select the relevant file type.
- In edit widget view, click on the download icon, and select **Download** and select the relevant file type.

# Interacting with Filters as a Viewer

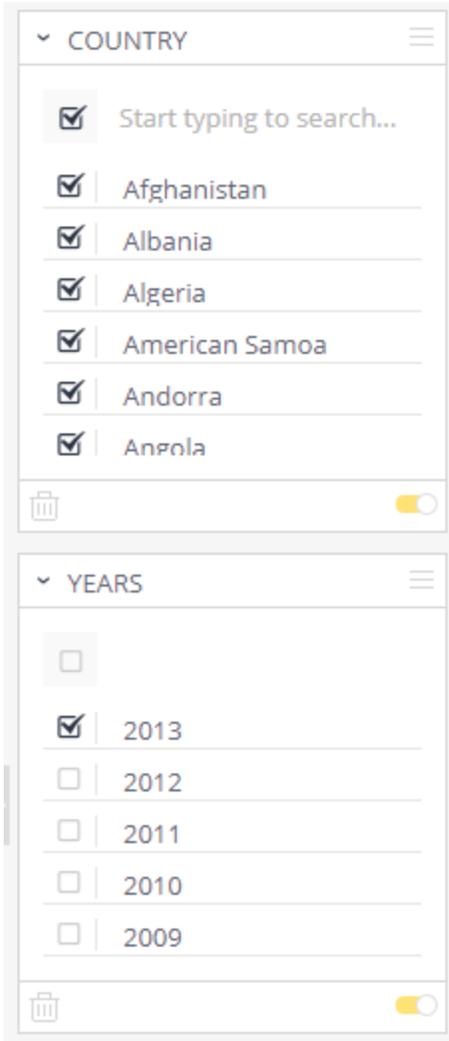
## Filtering Dashboards

This topic describes how Viewers can interact with filters for analyzing data. Through the Filters pane, you can modify filters in a dashboard and change the way values are sorted in a visualization. This is useful for focusing on data displayed in dashboards that have been shared by the dashboard's owner.

As a Viewer, you can select and deselect filter visuals on a dashboard. Each time you interact with a filter, for example by selecting or entering a value, the filter is immediately applied to your dashboard. After you have configured the relevant filters, you can save it as your default filters view.

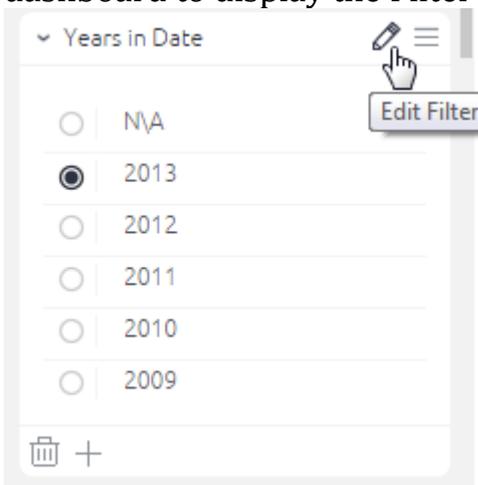
As a Viewer, you can interact with filters as follows:

- Make a different selection in the filter controls displayed in the Filters pane on the right side of the dashboard, as shown below:



OR

- Click on the pencil icon next to the filter name (shown above) in the dashboard to display the Filter Definition window.



In addition, Viewers can perform the following procedures:

- [Switching Filters On and Off](#)
  - [Saving your Default Filters View](#)
- To create a filter, you must be a [Designer](#).  
[Click here](#) to learn more about filters.

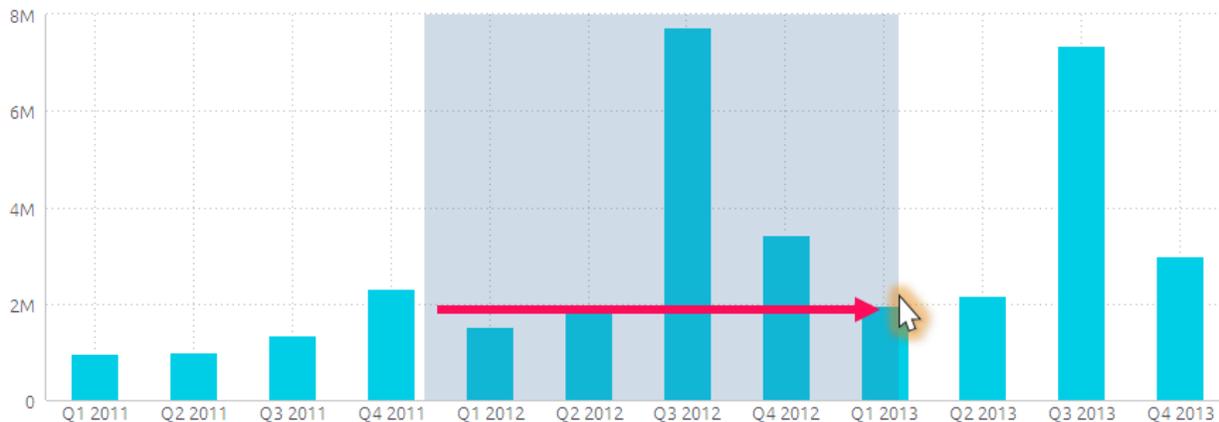
# Making Selections in a Widget

You can click on a specific portion of a widget to select it. This filters the dashboard according to the selected data by adding a filter to the dashboard's Filters panel.

You can also use the selection to drill down in the widget (see [Drilling Down in a Widget](#)).

## To select an item in a widget for filtering:

- Left-click an item in a widget to automatically select it and add a dashboard filter according to the selection.
- For multiple selection, use the Ctrl key. When released, a menu will appear. Click **Select**.
- In chart widgets, you can drag and draw the area to be selected, as shown below:



# Exporting Dashboards

You can export dashboards as images (PNG format) or PDF files. This is useful for storing dashboards or sharing them with other users.

For more information, see the following topics:

- [Creating PDF Reports](#)
- [Downloading Dashboards as Images](#)

# Working with Dashboards

## User Guide

# Sending Email Dashboard Reports

After you share your dashboard, you can configure the Sisense Scheduler to send periodic email reports containing a static version of the dashboard that can be clicked to access the dashboard or PDFs containing images of your dashboard. From your dashboard, you can configure how many emails the Scheduler sends and when.

To test how the report is to be displayed to your shared users, in the Shared Dashboard window, click the **Send me a report now** link, which generates an email report sent to your email address immediately.

## To schedule email notifications:

1. In your dashboard, click the Share icon to display the Share window.



- In the Shared Dashboard window, click the  tab.

Share: Sample - Ecommerce  

Owner is not selected to receive scheduled reports. [Select](#) [Dismiss](#)

Email Report Content

Email Report

PDF Attachment

Schedule

Every ElastiCube update, no more than  emails per day

Frequency  at

[Send me a report now](#)

- Select the type of report you want to send. There are two types of email reports you can send:
  - Email Report:** The email report is attached as an image file.
  - PDF Attachment:** The email report is attached as a PDF. You can define how the dashboard is to be displayed in the PDF through the PDF Report Settings page by clicking Edit/View PDF Report. For more information, see [Customizing PDF Reports](#).
- In the Schedule area, select the frequency for sending email notifications. You can select one of the following options:

Option	Description
Every ElastiCube update, no more than X emails per day	Select to send a report each time you update the ElastiCube up to the amount of emails you define in the Number box.  In the Number box, type or select the amount of emails that can be sent a day, regardless of the amount of builds.

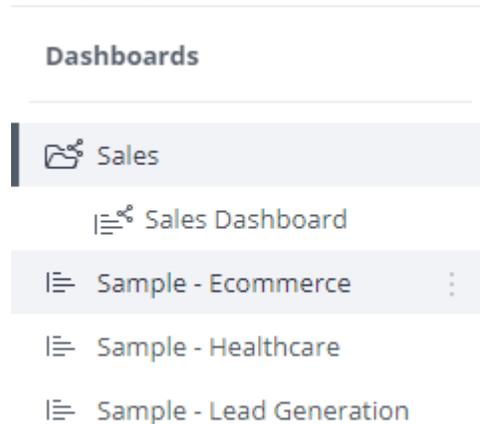
Option	Description
Daily	<p>Select to send one email report according to the time you define.</p> <p>In the Everyday box, select the days the scheduled report is to be sent. You can select multiple days by clicking the day.</p> <p>Select the Hour and the Time Zone that the email is to be sent on the days you selected.</p>

5. Click **Save**.

# Viewing Dashboards

This topic describes how you can view dashboards that you created or dashboards that were shared with you.

A dashboard that was shared with you appears with this  icon in the Dashboards list.



You may receive an email with a link or you may simply see a new dashboard in your Dashboards list on the Home page.

## To view a dashboard:

- Click on the link in the email. If you do not have a Sisense account, you will be automatically guided through the activation process.

OR

Browse to your Sisense Web page, log in and click the name of the dashboard in the Dashboards list on the Analytics page.

# Advanced Settings

By default, when an alert is triggered, you are notified via your email address registered with Sisense.

You can also send alerts using more advanced configurations that let you notify additional users when an alert is triggered, notify them across multiple channels, and customize the message displayed in the notification.

The following pages describe the functionality available in the Advanced Settings options:

- [Notifying Additional Users](#)
- [Setting Notification Channels](#)
- [Customizing Notification Messages](#)
- [Configuring Additional Options](#)

# Creating Build Alerts

Administrators can set build alerts to notify themselves or other users when a build fails or succeeds. Build alerts are created from the Sisense Web Application and are sent following an ElastiCube build.

Build alerts allow you to monitor the status of ElastiCube builds.

Administrators can create build alerts from the Admin section of the Sisense Web Application. After a build alert is created, it is added to Sisense Pulse from where you can manage it.

## To create a build alert:

1. In the Sisense Web Application, select **Admin > Data Sources > Add Build Alert**.

OR

From the ElastiCube menu, **Add Build Alert**.

Server	ElastiCube	Editors	Last Model Change	Last Build	Size On Disk	Status
LOCALHOST	Facebook Ads Demo	gideon@senses.com=1	5/5/17 1:09 PM	5/5/17 10:45 AM	4.89MB	Stopped
All Local Machines						
ElastiCube	Editors	Last Model Change	Last Build	Size On Disk	Status	
Commerce Demo	gideon@senses.com	10/28/15 12:03 PM	8/9/17 2:29 PM	566.95MB	Running	
Facebook Ads (1)	gideon@senses.com=1	5/3/17 1:09 PM	5/3/17 2:12 PM	5.38MB	Stopped	3
Facebook Ads Demo	gideon@senses.com=1	5/3/17 1:09 PM	5/3/17 10:45 AM	4.89MB	Stopped	
Facebook New 20175430	gideon@senses.com=1	8/3/17 12:45 PM	8/3/17 12:45 PM	13.07MB	Stopped	
Facebook Page	gideon@senses.com=1	4/25/17 2:13 PM	8/16/17 12:49 PM	18.09MB	Stopped	
new	gideon@senses.com	10/10/16 5:45 PM	6/6/17 11:56 AM	39.48MB	Stopped	
Sample Ecommerce	gideon@senses.com	2/10/14 2:21 PM	3/12/14 10:18 AM	59.06MB	Running	
Sample Healthcare	gideon@senses.com	3/2/13 1:07 AM	3/12/14 2:01 PM	1.33MB	Running	
Sample Lead Generation	gideon@senses.com	3/3/14 7:40 PM	3/12/14 11:35 AM	2.2MB	Running	
Training	gideon@senses.com	10/16/14 2:36 PM	10/16/14 2:28 PM	100.07KB	Stopped	

The Add Build Alert window is displayed.

2. In the **Name** field, enter a name for the alert.

**Add Build Alert**

Name

Server  ▾

ElastiCube  ▾

Notify when

- Build failed
- Build succeeded
- Build succeeded after failure

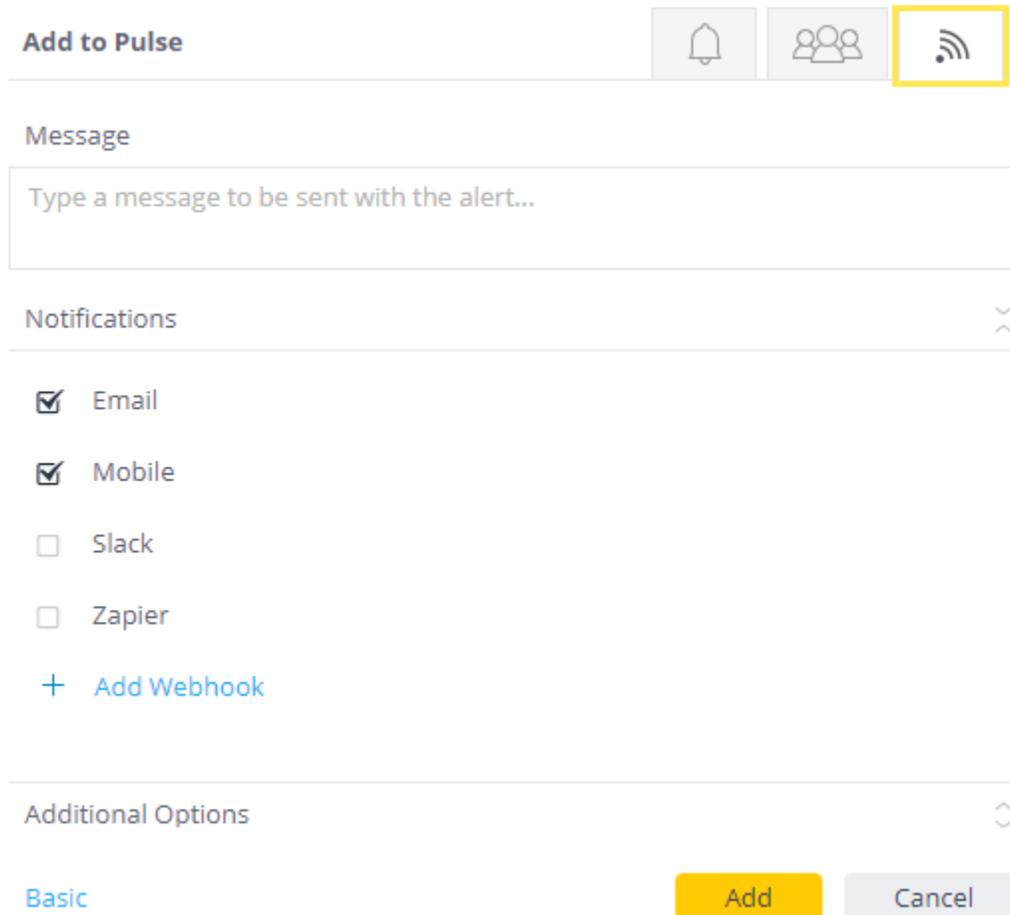
3. From the **Server** list, select the server where the ElastiCube is hosted.
4. From the ElastiCube list, select the ElastiCube(s) that triggers an alert following a build.
5. Select when to notify you following a build. There are three conditions that when met, can trigger an alert:
  - Build failed:** The alert is sent after the ElastiCube build fails.
  - Build succeeded:** The alert is sent after the ElastiCube is built successfully.
  - Build succeeded after failure:** The alert is sent following a successful build after the ElastiCube previously failed to build.
6. Click **Save**. The Alert is added to your Sisense Pulse page.  
For additional advanced configuration options, see [Advanced Settings](#).

# Configuring Additional Options

You can configure how often and when a notification is sent after an event is triggered. For example, you can limit the amount of notification sent to one when a data alert is triggered. This is useful if you build your ElastiCubes frequently, but do not need multiple notifications.

**To configure additional options:**

1. When adding or editing a widget to Pulse, in the Advanced Options window, select  > **Additional Options**.



**Add to Pulse**

Message

Type a message to be sent with the alert...

Notifications

- Email
- Mobile
- Slack
- Zapier
- + Add Webhook

Additional Options

Basic

Add Cancel

2. Select any of the following options:
  - For data alerts:**
    - Select **Notify only once after a condition is met** to limit how many notifications are sent after an event occurs.
    - Select **Notify when a condition is no longer met** to notify users after the event is no longer triggered.
  - For build alerts:**
    - Select **Notify once when there is a status change** to notify users only once when the latest build status is different than the previous build status.
3. Click **Add** to save the alert.

# Creating Data Alerts

This topic describes data alerts and how you can add them to Sisense Pulse. Data alerts are triggered by conditions you define for your widgets. For example, if you have an Indicator widget that represents your company's revenue, you can define an alert that notifies you when this revenue reaches or falls below a certain threshold.

Sisense supports alerts for the following widgets:

- Indicator Widgets
- Area Chart Widgets
- Bar Chart Widgets
- Column Chart
- Line Chart

For each type of widget you can add to Pulse, you can monitor a single value. For example, for Column charts you can monitor a single section (value) of a column. To monitor multiple values, you can add each value you want to monitor to Pulse separately.

When creating data alerts, you define the condition that triggers the alert.

Sisense provides three alert conditions that can trigger an alert:

 Threshold

Sisense notifies you when a threshold is reached following a build. This condition is useful for monitoring metrics that follow a regular, cyclical schedule.

 Automatic

Sisense automatically notifies you when the Sisense alert engine identifies an anomaly in your data based on machine learning algorithms. The Sisense algorithm actively learns from your data and eliminates noise to identify anomalies in your data. Sisense's anomaly detection is based on a smoothing

algorithm that takes the latest values into account. The more times an ElastiCube is built, the more accurate Sisense becomes at identifying any anomalies and notifying you.

 Always

Sisense notifies you after each build, regardless if there is a change in your data or not.

In addition, you can configure advanced settings that determine who receives the alert and what channel the alert is sent through. For more information see [Advanced Settings](#).

### To create a Data alert:

1. From a indicator widget's menu, select **Add To Pulse**.
2. The Add to Pulse window is displayed.

**Add to Pulse**

---

Name

Alert Condition

Threshold      Automatic      Always

---

The system will monitor this value and notify you when anomalies are detected.

[Advanced](#)      **Add**      Cancel

3. In the **Name** field, enter a meaningful name that represents your alert.
4. In the Alert Condition area, select the condition that triggers the alerts.  
**Threshold:** Sisense defines the threshold value as the current value displayed in the widget. You can increase or decrease this value and define the equality and relational operators that determine when the threshold is met.

Alert Condition



Threshold



Automatic



Always

**Automatic:** Sisense automatically identifies anomalies in your data and alerts you.

Alert Condition



Threshold



Automatic



Always

**Always:** Sisense notifies you every time the data changes.

Alert Condition



Threshold



Automatic



Always

You will be notified about this measure  
whenever it updates.

5. Click **Add** to create the alert. The alert is added to Sisense Pulse. For additional advanced configuration options, see [Advanced Settings](#).

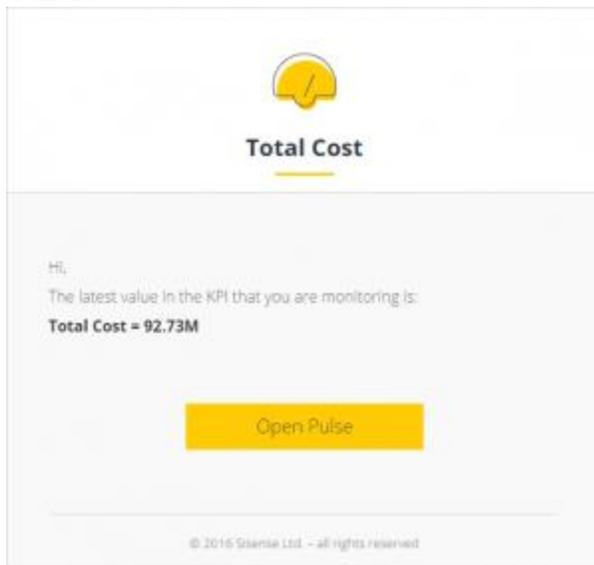
# Customizing Notification Messages

By default, the following message is provided in your email alert notifications:

*Hi,*

*The latest value in the KPI that you are monitoring is:*

**Value**



You can provide customized messages for your notifications when configuring your alert through the alert's Advanced Settings.

The message you provide replaces the string: The latest value in the KPI that you are monitoring is:

For example, if your message is "The last value of my KPI is:" the message in your alert's notification will be:

Hi,

The last value of my KPI is:

**Value**

## To customize your notification's text:

1. For Build Alerts, select the  tab.

**Add to Pulse**

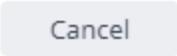
Message

Type a message to be sent with the alert...

Notifications 

- Email
- Mobile
- Slack
- Zapier
- [+ Add Webhook](#)

Additional Options 

[Basic](#)  

2. For Data Alerts, select **Advanced** > .

#### Add to Pulse

Name

Total Quantity

Alert Condition



Threshold



Automatic



Always

You will be notified when the value is:

Greater than ▾

32386

Advanced

Add

Cancel

3. In the Message area, enter your custom text.
4. Click **Save**. The next time a notification is sent, your custom message is displayed in the notification.
5. You can modify your alert notifications further including customizing the text and images displayed. For more information, see [Rebranding Sisense Automated Emails](#).

# Managing Tiles in Sisense Pulse

This topic describes how to manage your Pulse tiles and notifications when an alert is triggered.

After you have created alerts or added tiles to Sisense Pulse, you can manage those tiles from Sisense Pulse. Through the Sisense Pulse page, you can duplicate alerts and modify them, disable notifications, and delete alerts.

## Duplicating Tiles in Sisense Pulse

You can create copies of alerts and modify them to quickly generate new alerts with different notification messages or notify various groups across separate channels.

For example, if you want to notify Group 1 via email and Group 2 via Slack when the same event occurs, you can duplicate an alert and modify the notification channel.

### To duplicate a tile:

1. In the Sisense Web Application, select **Pulse**.
2. In the relevant alert's menu, select **Duplicate**. A copy of the alert is created and the Add to Pulse window is displayed.

**Add to Pulse**

Name: Copy of TOTAL      Value: TOTAL

Filters: 5 active filters [Show](#)

Alert Condition - Threshold ▾

Greater than ▾      90000

**Add**      **Cancel**

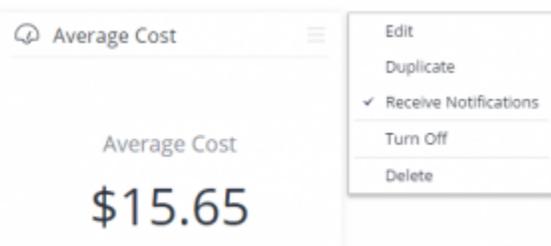
You can modify the alert's settings from the **Add to Pulse** window. For more information about these settings, click [here](#).

## Deactivating Your Notifications

Each time you add a tile to Sisense Pulse or create an alert, by default, you will receive notifications when the event is triggered. If you want to just monitor KPIs from Sisense Pulse, or you do not want to receive any notifications, you can disable notifications from being sent by toggling the notifications. When you toggle the notifications, you no longer receive notifications when an alert is triggered, however, other users with whom the alert was shared will continue to be notified. See [Deactivating Notifications for All Users](#) to turn off notifications for everyone associated with an alert.

### To deactivate notifications:

1. In the Sisense Web Application, select **Pulse**.
2. From the Tile menu of the relevant alert, toggle **Receive Notifications**.



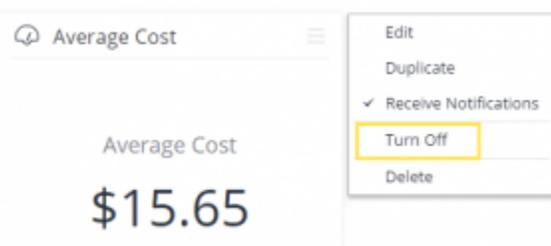
## Deactivating Notifications for All Users

If you have added an alert that is no longer relevant or you want to deactivate the alert for all users, you can turn off the alert.

When you turn off an alert, no users are notified when the event that triggers the alert occurs.

### To deactivate notifications for all users:

1. In the Sisense Web Application, select **Pulse**.
2. From the Tile menu of the relevant alert, select **Turn Off**. A confirmation dialog box is displayed.



3. Click **Yes** to confirm that you want to turn off all notifications for the alert. The alert is turned off.  
Alerts that have been turned off appear in the Pulse page as shaded tiles. You can reactivate the alert by selecting **Turn On** from the Tile menu.

## Deleting Tiles in Sisense Pulse

When a tile is no longer needed, you can delete the tile through the Sisense Pulse page.

Deleting an alert removes it from the Sisense Pulse page and users are no longer notified.

### To delete an alert:

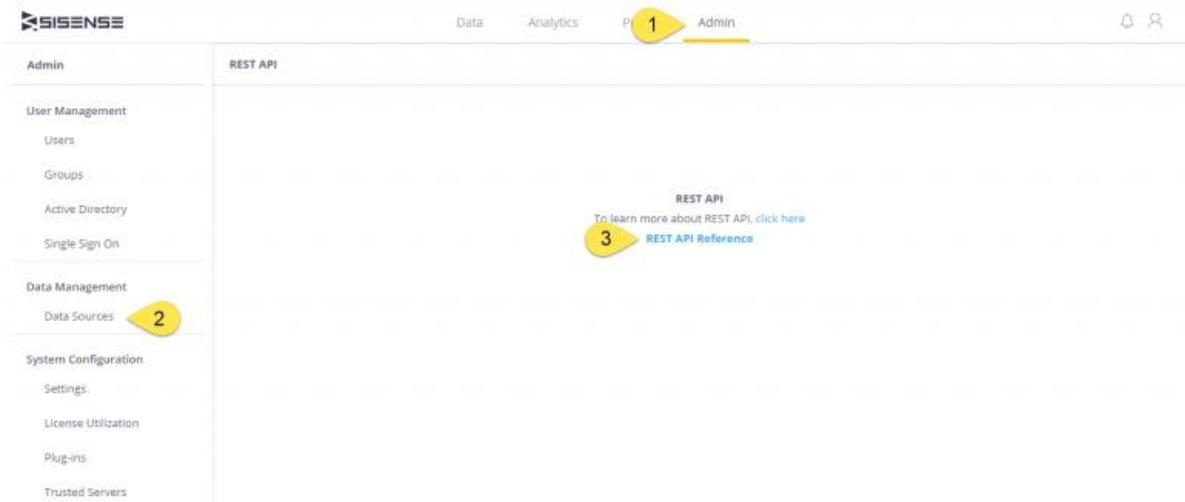
1. In the Sisense Web Application, select **Pulse**.
2. In the relevant alert's menu, select **Delete**.

## Removing Sisense Pulse

You can remove the Sisense Pulse through the Sisense REST API. Removing Sisense Pulse removes the Pulse page from the Sisense Web Application and deactivates all the alerts defined in Sisense Pulse.

### To remove Sisense Pulse:

1. Access the Sisense REST API.



2. In version 1.0 of the REST API, select the **PATCH Dashboard** endpoint, **/features/{key}/toggle**.
3. In the **Authorization** field, enter your Sisense authorization key. The value of this field is the user's API token, preceded by the keyword Bearer (with a space between it and the token). For information about how to retrieve this key, see [Using the REST API](#).
4. In the **Key** field, enter the name of the feature to be enabled or disabled. In the case of Sisense Pulse, the key is alerting. You can retrieve the available features that you can toggle from the GET /features endpoint.
5. In the **toggleFeature** field, set the active key value as true or false to enable (true) or disable (false) a feature.

6. Click **RUN** to toggle the alert.

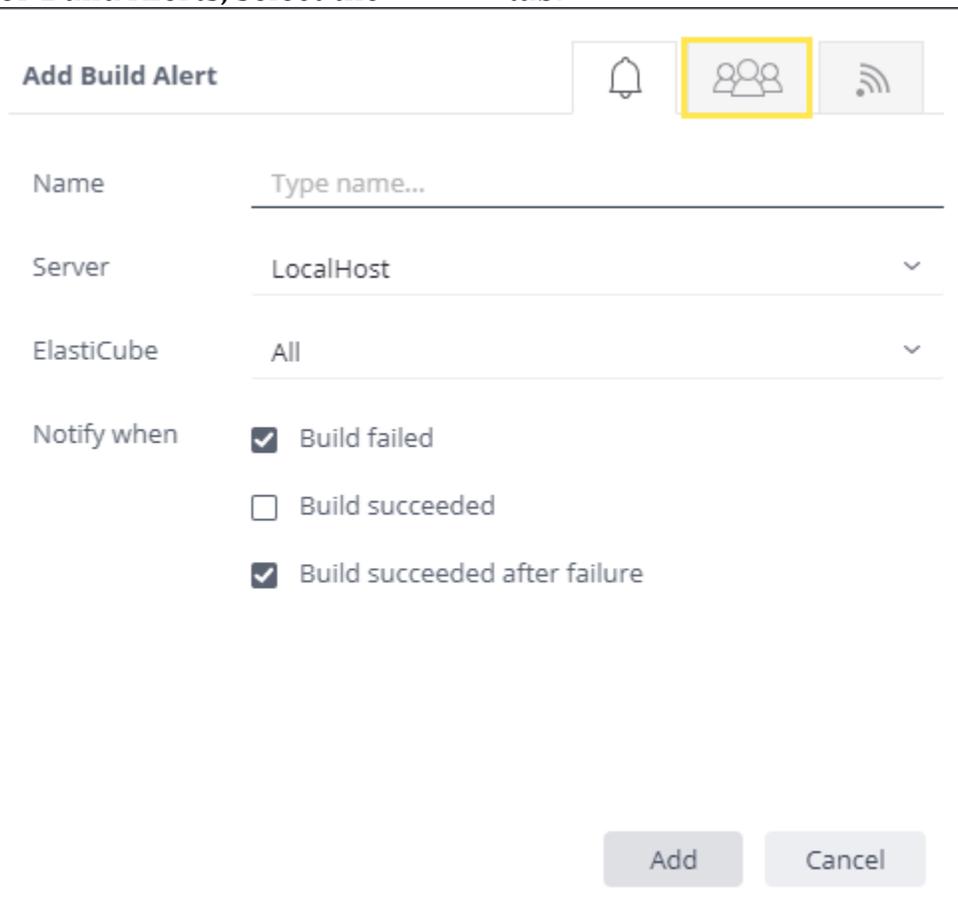
# Notifying Additional Users

When an alert is triggered, Administrators and Designers can notify additional Sisense users and groups by configuring the alert's advanced settings. Viewers cannot share alerts with additional users.

Users who are notified receive a notification through Sisense Pulse and any other additional channels you define.

## To notify additional users:

1. For Build Alerts, select the  tab.



**Add Build Alert**

Name

Server  

ElastiCube  

Notify when

- Build failed
- Build succeeded
- Build succeeded after failure

For Data Alerts, select **Advanced Options** > 

**Add to Pulse**   

---

Name	Value
Total Revenue	Total Revenue

Filters

3 active filters [Show](#)

Alert Condition - Threshold ▾

Greater than ▾  

---

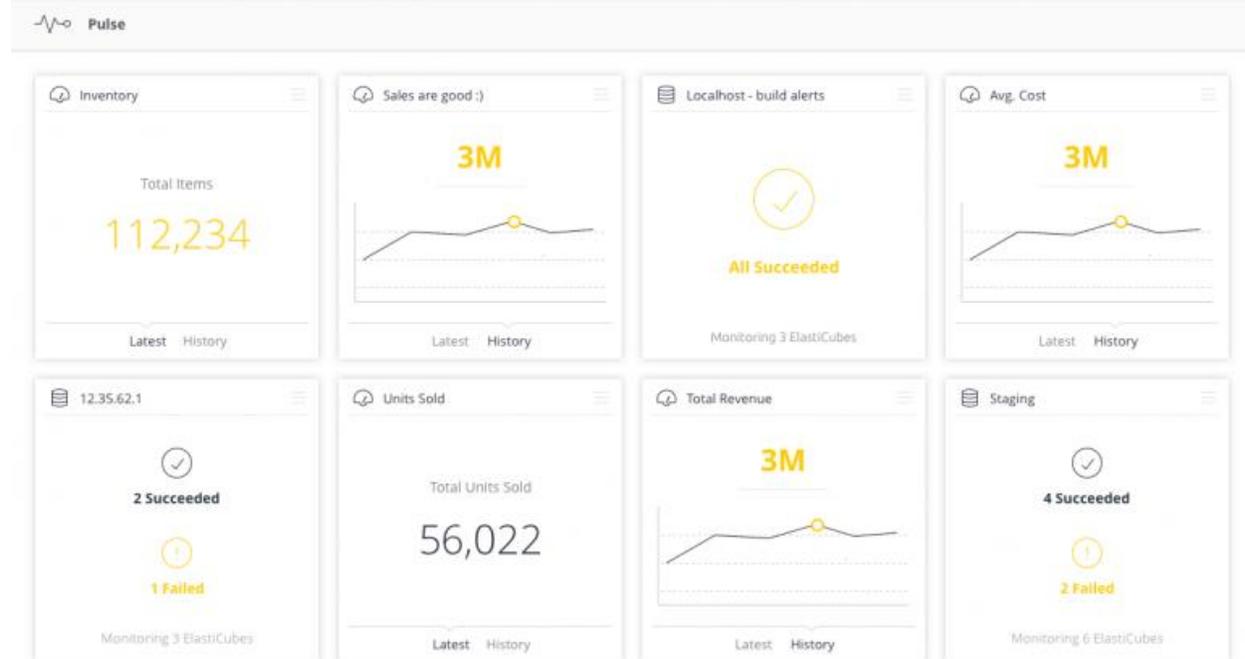
[Basic](#) [Add](#) [Cancel](#)

2. Enter the email address, Sisense username, or group name of the users you want to notify when an alert is triggered. Entering **Everyone** specifies that all Sisense users and user groups associated with your account are notified. The contact information of the people you add is listed in the center of the window and is displayed the next time you open this window.
3. Click **Add** to add the alert to Pulse.

# Introduction Sisense Pulse

Sisense Pulse is a centralized location where you can stay on top of your most important KPIs across multiple dashboards or manage your data and build alerts. For example, you can consolidate and monitor important aspects of your business by adding Indicator widgets from various dashboards to the Sisense Pulse page.

Then, you can create alerts to notify you when certain thresholds are met or anomalies in your data are detected. Sisense Pulse provides you with access to your most important data and notifies you when to take action.



The following sections describe Sisense alerts and managing alerts and KPIs in Sisense Pulse.

## Sisense Pulse

- [Overview](#)
- [Managing Pulse Tiles](#)

## Alerts

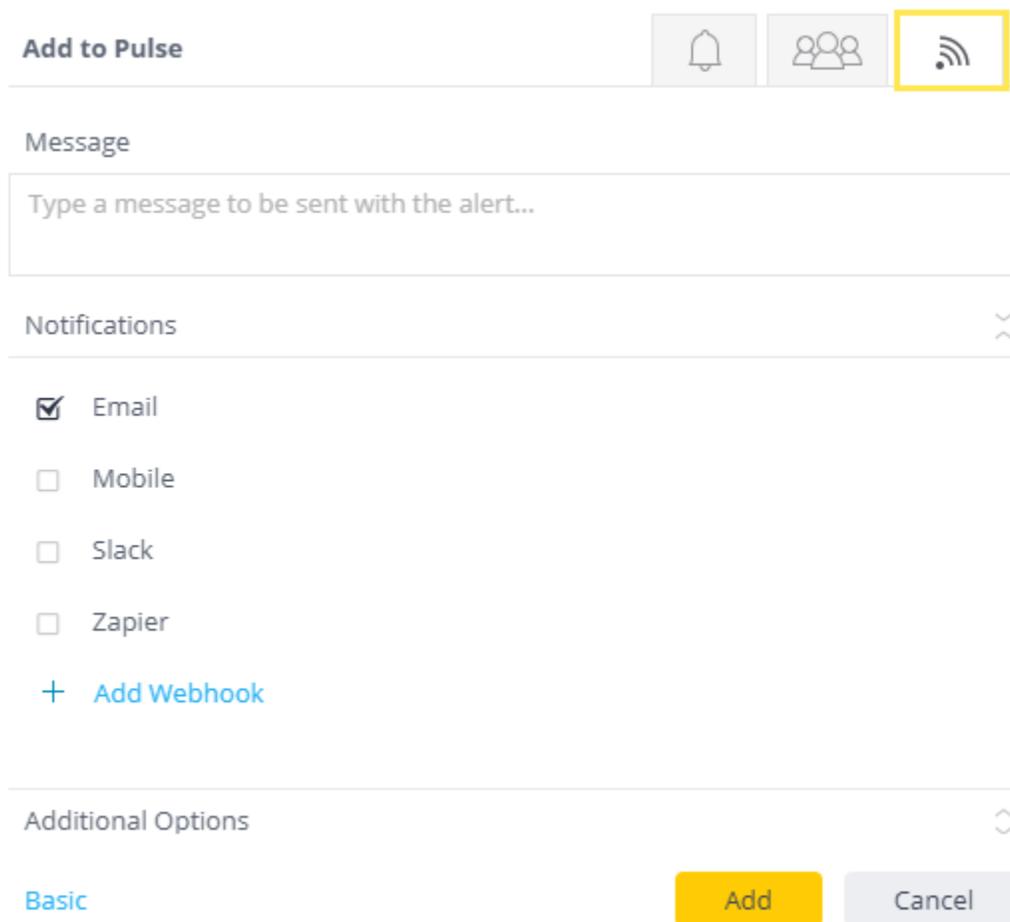
- [Overview](#)
- [Creating Data Alerts](#)
- [Creating Build Alerts](#)
- [Advanced Alert Settings](#)

# Setting Notification Channels

By default, your users are notified via their email addresses registered with Sisense. In addition, Sisense supports a variety of workflows by sending notifications across multiple applications and services. For example, your users can receive notifications through the Sisense mobile application, Zapier, Slack, and other 3rd services via webhooks.

## To define how users are notified:

1. For Build Alerts, select the  tab.
2. For Data Alerts, select **Advanced Options** > .



**Add to Pulse**

Message

Type a message to be sent with the alert...

Notifications

- Email
- Mobile
- Slack
- Zapier

+ [Add Webhook](#)

Additional Options

Basic Add Cancel

3. In the Notifications area, select the relevant channel for your notification. The default channel, Email, is already selected.

You can select one or more of the following channels:

- Email
- Mobile
- Slack
- Zapier
- Webhook
- mobile

4. Click **Save**.

## Sisense Mobile

Sisense Mobile enables users to view Sisense dashboards through their mobile devices. You or your users can receive notifications through Sisense Mobile. After selecting the users you want to notify, any users with a Sisense user account and Sisense Mobile will receive a notification when an alert is triggered.



## Slack

Slack is a real-time team collaboration messaging platform that allows you to receive incoming Webhooks from external sources such as Sisense. To send

notifications through Slack, you need to enter the URL and optionally, your Slack channel or username, where the notification is to be displayed.

The image shows a user interface for configuring alerts. At the top, there's a section titled "Add to Pulse" with three icons: a bell, a group of people, and a Wi-Fi symbol. Below this is a "Message" section with a text input field containing the placeholder "Type a message to be sent with the alert...". Underneath is a "Notifications" section with three checked checkboxes: "Email", "Mobile", and "Slack". A modal dialog is open over the "Slack" checkbox, containing three input fields: "URL" with the placeholder "Type in the URL...", "#channel/@username", and "Optional". At the bottom of the modal are green and red checkmarks. To the right of the modal are two buttons: a yellow "Add" button and a grey "Cancel" button.

You can obtain these details by adding the Incoming WebHooks app into your Slack configuration.

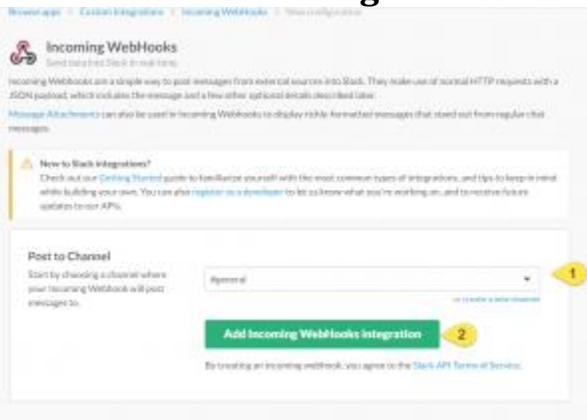
**To obtain a Webhook URL and channel from Slack:**

1. Add the Incoming WebHooks app to Slack. You can find this app in Slack's App Directory.

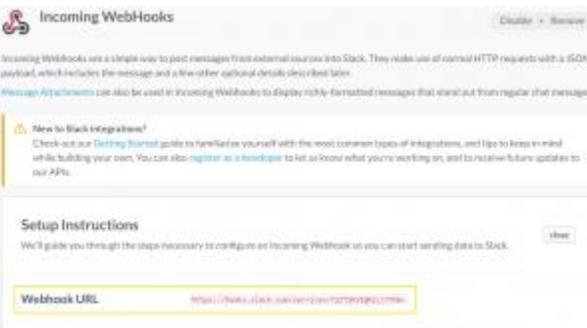
2. In the Incoming WebHooks app page, click **Add Configuration**.



3. After you add the Incoming WebHooks app, you can select which channel to post the notifications to. The channel name is optional and displays the notification in the channel associated with your WebHook URL. In the Post to Channel area, select the channel where the notification is to be displayed and click **Add Incoming WebHooks integration**.



4. In the Setup Instructions the Webhook URL is displayed. Copy the Webhook URL.



5. In Sisense when sending notifications through Slack, in the **URL** field, enter the URL copied in the previous step and in the **#channel/@username** field, enter the channel name if you have defined one.

# Zapier

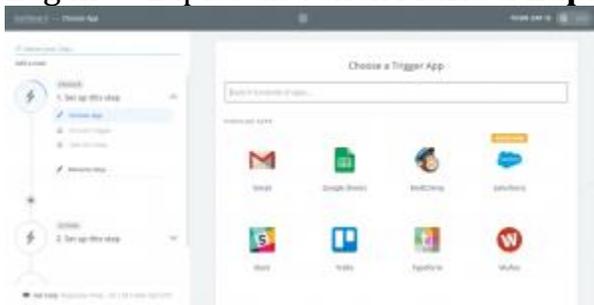
Zapier is an online service that lets you create automated actions connecting disparate business apps and services such as Sisense. To receive notifications through Zapier, you must create a Zap and provide Sisense with the relevant URL when configuring your alerts' notifications.

The screenshot shows a configuration window for an alert. At the top, there is a section titled "Add to Pulse" with three icons: a bell, a group of people, and a Wi-Fi signal. Below this is a "Message" section with a text input field containing the placeholder text "Type a message to be sent with the alert...". Underneath is a "Notifications" section with a list of options: "Email" (checked), "Mobile" (unchecked), "Slack" (unchecked), and "Zapier" (checked). A small edit icon is next to the "Zapier" option. A modal dialog box is open over the "Zapier" option, titled "URL" and containing a text input field with the placeholder "Type in the URL...". At the bottom of the dialog are green and red checkmarks. Below the dialog, there is a "Basic" tab, a yellow "Add" button, and a grey "Cancel" button.

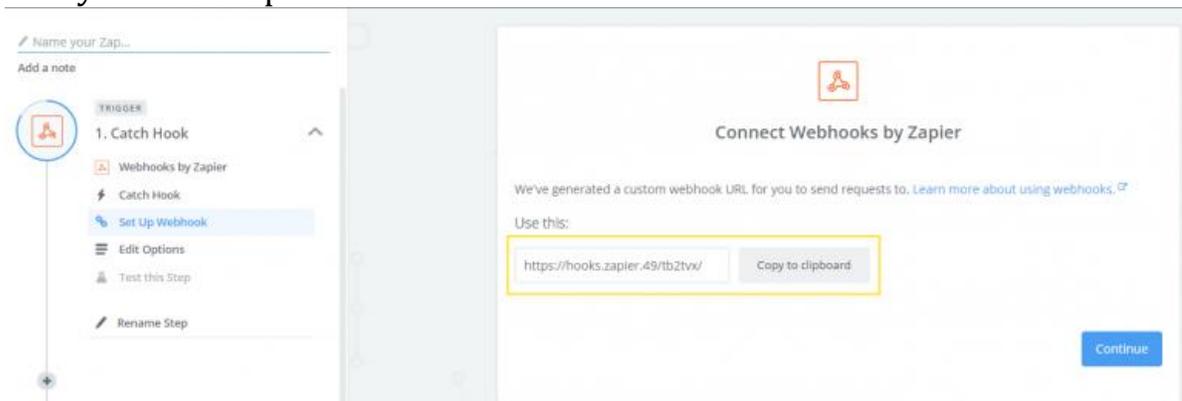
When an alert is triggered, Sisense sends the notification to the Webhook you define in Zapier.

**To obtain a Webhook URL from Zapier:**

1. Log in to Zapier and click **Make a Zap**. The Trigger App page is displayed.



2. From the Trigger App page select Webhooks. The Connect to Webhooks by Zapier page is displayed.
3. In Step 1 > Set Up Webhook, click **Copy to clipboard** to copy the Webhook URL you need to provide to Sisense.

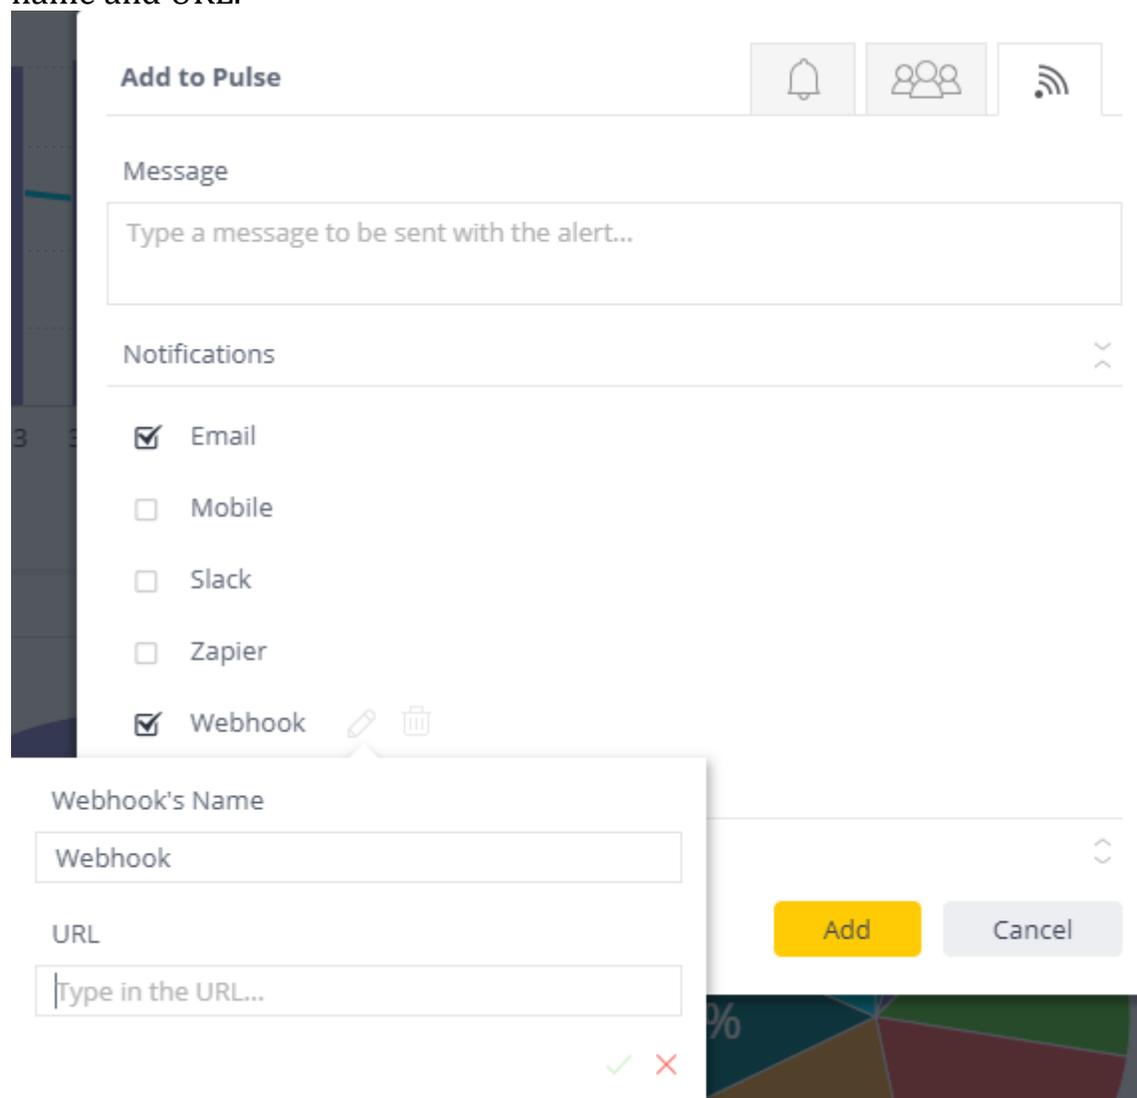


4. In Sisense, when sending notifications through Zapier, in the **URL** field, enter the URL copied in the previous step.

## Webhooks

Webhooks are a system of automated notifications that indicate that an event occurred, in this case an alert being triggered. To send notifications through additional 3rd party channels, you can select **Webhook** and enter the Webhook

name and URL.



Your Webhook server listens for incoming messages from Sisense across the URL specified in the URL field.

When Sisense sends notifications to your custom Webhook, they are sent within a JSON payload. Your Webhook provider must interpret the payload and display the notification to your users. The following are examples of a notifications sent in a JSON payload and a description of the contents of the payload:

```
{ type: 'kpi', name: 'Total Sales', message: 'Hello World', originDashboardLink: '58128745677f52dc0f000019', measure: 'Revenue', value: 39759625.51502721, condition: 'none', filters: [ { asDimension: false, title: 'Category', value: 'Include
```

```
all', jaql: [Object], panel: 'scope' }, { asDimension: false, title: 'Gender', value:
'Include all', jaql: [Object], panel: 'scope' } ]
```

Data Events	Build Events
<pre>{ type: 'kpi',   name: 'Total Sales',   message: 'Hello World',   originDashboardLink: '58128745677f52dc0f000019',   measure: 'Revenue', value: 39759625.51502721,   condition: 'none', filters:   [ { asDimension: false,       title: 'Category', value: 'Include all',       jaql: [Object],       panel: 'scope' },     { asDimension: false, title: 'Gender', value: 'Include all',       jaql: [Object],       panel: 'scope' } ]</pre>	<pre>{ type: 'build', name: 'Success Event', measure: null, value: 1, condition: 'none', filters: null }</pre>

--	--

Name	Type	Description
Type	String	Type of alert. 'kpi' for data events and 'build' for build events.
Name	String	Name of the alert.
Message	String	Notification message.
originDashboardLink	String	Dashboard ID provided by Sisense. You can find the dashboard ID from the URL of the dashboard. For example, the ID of the dashboard below if: <code>http://localhost:8081/app/main#/dashboards/<b>573b4322048e5e382a000001</b></code>
value	Double	For data events when a threshold condition is applied, this is the value of the threshold. For build events, the value is 0 for failed builds and 1 for successful builds.
condition	String	Type of condition that generated the alert.
filters	Object	Object that defines filters applied to the dashboard.

# Sisense Alerts

This topic describes the two types of alerts you can define, data alerts and system alerts in Sisense.

Data alerts are triggered by changes to your data. Sisense provides several options for defining the type of changes that trigger data alerts including threshold based alerts, automatic anomaly detection, and Always which is triggered by any change to your data. For example, if you have an Indicator widget that represents total sales, Sisense can notify you when a threshold you defined is met or exceeded.

Build alerts can be configured to be sent following build attempts.

When creating an alert, you can customize your notifications, the channel through which they are sent, and how often a notification is sent following a data or build.

After you create an alert, the alert is added to the Sisense Pulse page where you can manage the alert and its notifications. For example, you can delete alerts or enable and disable notifications.

If you want to make a similar alerts, for example, alerts with different notifications for different groups of users, you can duplicate the alert and modify its settings.

When an alert is triggered, you and your users are notified through the channel you defined when creating the alert and through the Pulse Feed. The Pulse Feed is displayed in the Sisense Web Application and shows a list of your most recent triggered alerts. From the Pulse Feed you can open the dashboard from where

the alert originated or open the Sisense Pulse page.

 Build Alert - Success	The ElastiCube <b>Test</b> on <b>LocalHost</b> was built successfully.	
 Conv		11/1/16 3:30 PM
 Converted (%)	Converted (%) = 7.29%	
		10/31/16 9:35 AM
 Average Cost	Average Cost = \$15.65	
		10/31/16 9:35 AM
 Converted (%)	Converted (%) = 7.29%	
		10/25/16 3:55 PM
 Average Cost	Average Cost = \$15.65	
		10/25/16 3:55 PM
 Build Alert - Success	The ElastiCube <b>Test</b> on <b>LocalHost</b> was built successfully.	
		10/25/16 10:12 AM
 Build Alert - Success	The ElastiCube <b>Test</b> on <b>LocalHost</b> was built successfully.	
		10/13/16 9:20 AM
<a href="#">Clear all</a>		

## **Related Topics**

[Creating Data Alerts](#)

[Creating Build Alerts](#)

# Sisense Pulse

Sisense Pulse enables you can stay on top of your most important KPIs across all your dashboards and your latest ElastiCube builds. By adding important KPIs from your dashboards to Sisense Pulse, you can get a comprehensive picture of your data from a single location.

Sisense Pulse contains tiles that display information from your dashboards and the status of your ElastiCube builds.

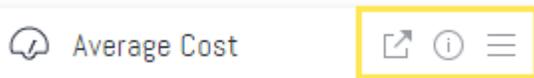
Within Sisense Pulse, there are two types of tiles, data tiles that display your data's current status and recent history and build tiles that display the latest status of your ElastiCube builds.

You can access Sisense Pulse from the menu bar of the Sisense Web

Data      Analytics      **Pulse**      Admin

Application.

By hovering over the header of any tile in Sisense Pulse, you can view additional options for your tiles.



The  icon opens a new tab in your browser. For data tiles the dashboard where the widget originated is displayed and for build tiles the Data Sources page in the Admin area is displayed.

The  icon is a tooltip that describes the tile in more detail. For data tiles the tooltip describes who added the tile to Pulse and when it was added. Build tile tooltips describe when the tile was created, which ElastiCube is being monitored, and under what conditions an alert is triggered when alerts are enabled.

The ☰ icon displays a menu for managing tiles in Sisense Pulse. For more information regarding managing tiles in Sisense Pulse, see [Managing Tiles in Sisense Pulse](#).

## Adding Tiles to Sisense Pulse

To keep track of your important KPIs and ElastiCube builds, you add the relevant data and build tiles to Sisense Pulse.

The flow for adding tiles to Sisense Pulse is the same for creating alerts for Indicator widgets and ElastiCube builds. For more information about these flows, see the following topics:

For data tiles, see [Creating Data Alerts](#).

For build tiles, see [Creating Build Alerts](#).

Each of these flows adds tiles to Sisense Pulse and creates an alert. If you want to monitor your data and ElastiCube builds without alerts, see [Deactivating Notifications](#).

# Accessing the Sisense Application Database

When Sisense is installed, Sisense installs a MongoDB database as the application database with three MongoDB users:

- **AppUser:** The user used by the Sisense application to access data stored in the Sisense MongoDB.
- **ReadUser:** A user with read-only access to the Sisense MongoDB data.
- **WriteUser:** A user with read-write access to the Sisense MongoDB data.

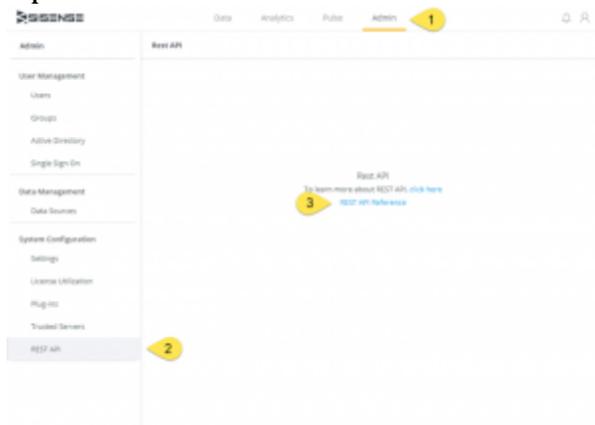
The AppUser is created to allow the Sisense Application to interact with the application database. This password is unique to your Sisense installation and should not be changed. The credentials of the AppUser are shared between each node of the replica set. If you want to create a Sisense MongoDB instance on server where Sisense is not installed, see [Hosting an Application Server](#).

## Creating Application Database Users

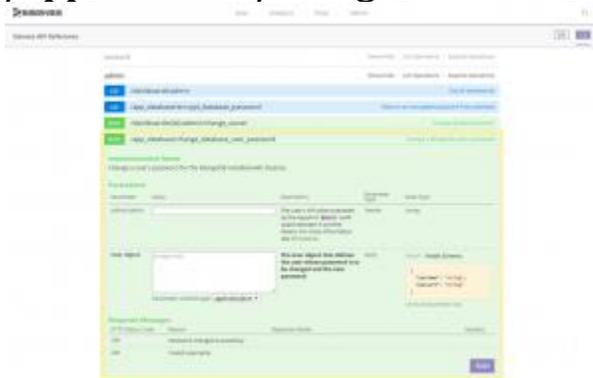
The ReadUser and WriteUser can be used to access your Sisense application database. These users have the relevant read-only and write-only MongoDB roles assigned to them for accessing the application database.

**To create an application database user:**

1. Open the Sisense REST API.



2. In the REST API, select **admin POST /app\_database/change\_database\_user\_password**.



3. In the field **Authorization**, enter your API token. See [Using the REST API](#) for more information.
4. In the User object, provide the username (ReadUser or WriteUser depending on which user's password you want to change) and the new password:

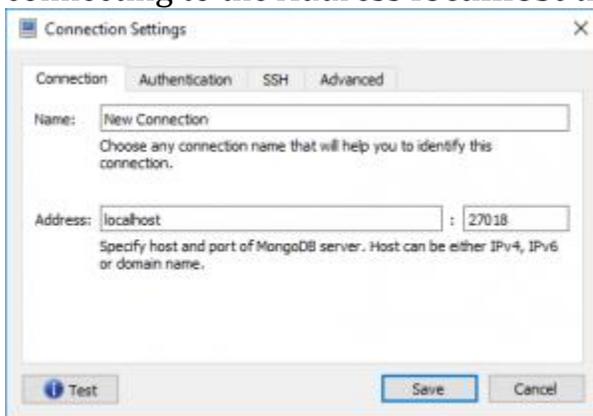
```
{
  "userName": "string", (ReadUser or WriteUser)
  "password": "string" (The password you enter here is
the password you will provide when connecting to the
database)
}
```
5. Click **RUN**.

This endpoint encrypts your password and saves it in Sisense application database. When you connect to the application database, use the plaintext password you entered as the value of the password parameter, not the encrypted password returned by the API.

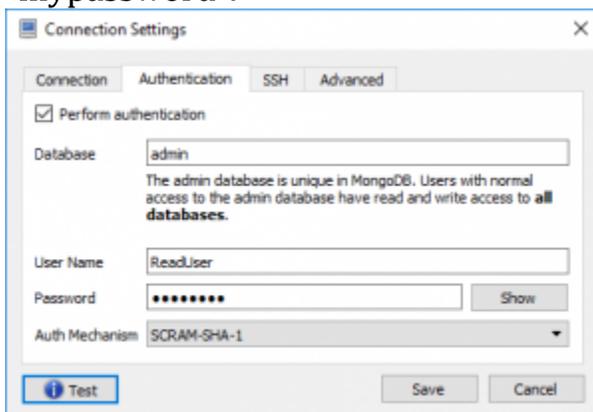
# Connecting to the Application Database

Once you have defined the username and passwords for your application database users, these users can connect to the application database through applications such as RoboMongo or when importing data into an ElastiCube from your Sisense MongoDB as described in [Connecting to MongoDB](#).

For example, in RoboMongo, you can connect to your local Sisense MongoDB by connecting to the Address **localhost** through port **27018**.



In the Authentication tab, enter the name of the Sisense MongoDB, admin, and your username and password. The username is the user whose password you changed, in this example, ReadUser. The password is the password you provided through the **POST admin/app\_database/change\_database\_user\_password** endpoint. This password is not the encrypted password that was returned, but the password you defined as the value of the password parameter, for example “mypassword”.



Click **Test** to verify that your user can connect. If your credentials are authenticated successfully, click **Save** and then connect to the database. Your Sisense MongoDB tables will be displayed.

# Enabling SSL for the Application Database

If the application database is hosted on another server, Sisense supports the SSL (Secure Sockets Layer) protocol to ensure that your connection to the application database is secure.

To establish an SSL connection with an application database hosted on another server, you must first obtain a certificate by creating your own or through a 3rd party provider and then import the certificate into both your Sisense and non-Sisense servers. Once the certificates are on both servers, you need to update Sisense configuration files that define the location of those files.

This procedure explains how to modify these configuration files to enable SSL.

## To enable SSL for the Sisense application database:

1. Open the MongoDB configuration file, `mongodbconfig.conf`, located at:  
`...\Program Files\Sisense\Infra\MongoDB\mongodb.conf`
2. In the configuration file, define the locations of your PEM file and the CRT file if you are using signed certificates.  
net:  
ssl:  
mode: requireSSL  
**PEMKeyFile:** "path to mongodb.pem file"  
**CAFile:** "path to certificate.crt file" (only if using signed certificate)
3. Save your changes to the `mongodbconfig.conf` file and open the `default.yaml` located at:  
`...\Program Files\Sisense\PrismWeb\vnext\config\`
4. In the `default.yaml` file, define the values for the following parameters under `db_security`:  
**sslEnabled:** A boolean value. The default value is false. This value should be true when using SSL.  
**sslValidate:** A boolean value. The default value is false. This value should be true when using self-signed certificates.

**certFilePath:** A string value that represents the location of your CRT file.

**pemKeyFilePath:** A string value that represents the location of your mongodb.pem file.

**caFilePath:** A string value that represents the location of your CA file when using self-signed certificates.

db\_security:

mongoAppUser: AppUser

mongoAppPassword: o9VacJmhgygqmnJSq4Q==

sslEnabled: true

sslValidate: false (only if using self-signed certificate)

certFilePath: 'path to certificate.crt file'

pemKeyFilePath: 'path to mongodb.pem file'

caFilePath: 'path to CA file' (only if using signed certificate)

5. Save your changes and open the Sisense.WebQueryService.exe.config file located at:

...\Program Files\Sisense\PrismWeb\Bin\

6. Edit the connection string and add the SSL flag:

```
<connectionStrings>
```

```
<add name="Mongo"
```

```
connectionString="mongodb://AppUser:we6jBsd0K6lvfaTmA==@localhost:27018/admin?ssl=true" />
```

```
</connectionStrings>
```

7. Under <apSettings>, add two new fields:

```
<add key="SslPfxFilePath" value="path to the PFX file"/><add
```

```
key="SslPfxFilePassword" value="PFX file password"/>
```

**Note:** If you are using self-signed certificates, you can add it to the trusted root certificate authorities of your operating system by running the PFX file.

8. Save your changes. SSL is now enabled. If you need to connect directly to your MongoDB with an application such as RoboMongo, you can enter your

## SSL configuration when creating or editing a connection.

Connection Settings ✕

Connection Authentication SSH **SSL** Advanced

Use SSL protocol

Authentication Method: Use CA Certificate ▾

CA Certificate:  ...

Use PEM Cert./Key: Enable this option to connect to a MongoDB that requires CA-signed client certificates/key file.

Advanced Options

ⓘ Test Save Cancel

# Hosting the Application Database

If you wish to host the Sisense application database on another server, you need to create an AppUser that Sisense can use to communicate with your application database. The AppUser you create must have two application database roles assigned to it, readWriteAnyDatabase and dbAdminAnyDatabase. These roles provide the Sisense AppUser read/write privileges to the application database. In addition, the password for the AppUser should be encrypted using the same encryption mechanism Sisense uses. You can encrypt a plaintext password using this encryption mechanism in the Sisense REST API through the endpoint:

```
GET /app_database/encrypt_database_password
```

This encrypted password should be included in all the relevant configuration files as described in [Connecting Sisense to the Replica Set](#).

# Creating a KeyFile

To ensure that your connection between your application database servers is secure, Sisense uses KeyFiles that you generate.

The contents of the keyfile serves as a shared password for each of the nodes in your replica set. The content of the keyfile must be the same for all members of the replica set.

You can generate a keyfile using any method you choose. The contents of the keyfile must be between 6 and 1024 characters long.

The following operation uses OpenSSL to generate a random character string to use for a keyfile.

## **To create a KeyFile:**

1. Download [OpenSSL](#) and extract the files to your machine.
2. Open the Command Prompt and navigate to your bin folder of your OpenSSL directory, for example, .../Openssl/bin.
3. Run the command:

```
openssl rand -base64 756 > <path-to-keyfile>
```

For example: `openssl rand -base64 756 > c:\openSSL\key`

The file 'Key' is created in the path you defined in the command. The location of this file should be referenced in the relevant configuration files.

The file should be copied to each node in your replica set.

# Limiting Access to the Application Database

For security reasons, by default, access to application database from outside of localhost (127.0.0.1) is prevented.

If you need to access the application database from an IP address other than localhost, you can add the IP address of that server to the Sisense MongoDB configuration file, which contains a list of all servers allowed to access the application database. The Sisense MongoDB configuration file, `mongodb.conf`, is located in the following directory

```
..\Program Files\Sisense\Infra\MongoDB\
```

To allow external IP addresses to access your application database, add the IP addresses in this format `[127.0.0.1, XX.XX.XX.X]` as the value of `bindIp`, or change it to `0.0.0.0` to enable connections from any IP.

```
net:
```

```
bindIp: [127.0.0.1, XX.XX.XX.XX]
```

**Note:** The value of `bindIp` has been changed from an array to a string. As a string, the IP addresses should be separated by a comma with no space in between. In addition, the only IP addresses to be included in the string as `127.0.0.1` and the public IP of your local server.

# Sisense Application Database

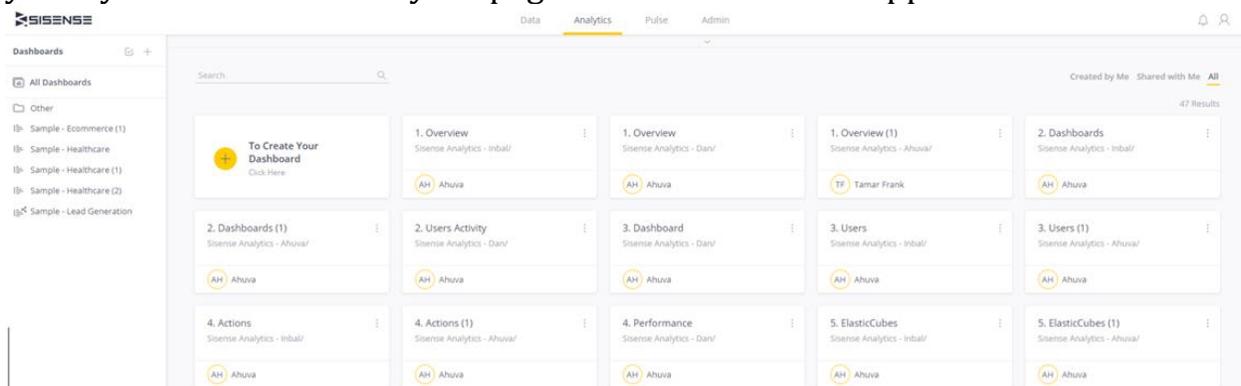
When you install Sisense, an application database is installed locally on your machine that contains metadata for ElastiCubes, dashboards, and users. The application database is used by Sisense to display the correct dashboards and provide access to the relevant user.

The following topics describe how you can host the application database remotely, back it up, secure it, and maintain it in a high availability configuration:

- [Hosting the Application Database](#)
- [Accessing the Sisense Application Database](#)
- [Replicating the Application Database](#)
- [Limiting Access to the Application Database](#)
- [Enabling SSL for the Application Database](#)
- [Creating a KeyFile](#)

# Managing Your Users Dashboards

As a Sisense Administrator, you can manage and interact with all dashboards in the system, including your own dashboards and dashboards of other users across your system from the Analytics page of the Sisense Web Application.



For Administrators, the Analytics page becomes a centralized location for managing all the dashboards in your system and performing administrative actions on them.

This option is available from the “All Dashboards” item above the Dashboard list.

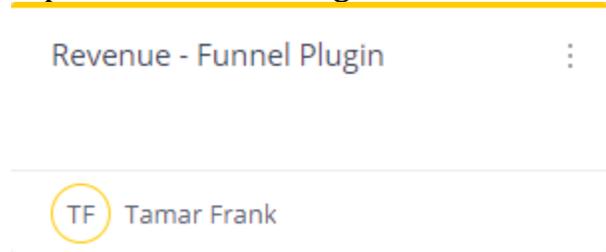
As an Administrator, you can perform the following operations:

- Sharing the dashboard with other users in the system, or change the dashboard’s owner
- Accessing a user’s dashboard as a Viewer
- Copying the dashboard to another server
- Downloading a PDF of the dashboard
- Deleting the dashboard

This is in addition to the functionality available to you for your own dashboards.

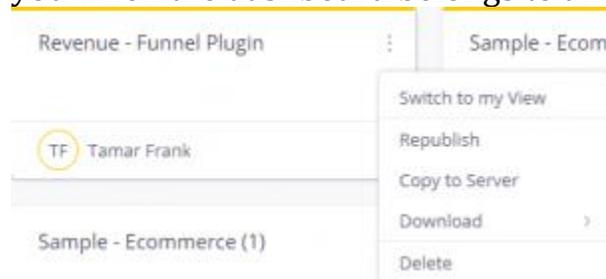
When you select the **All Dashboards** button on the Analytics page, you can see all your dashboards and the dashboards of your users. A yellow border displayed on the top of the dashboard tile indicates that you are looking at the dashboard

version of one of your users, and not your own version of the dashboard. This can be useful to view a dashboard that the owner modified, but has not yet republished the changes.



You can click the dashboard tile to open the dashboard. When opening a dashboard not shared with you, you will see the dashboard as a Viewer without editing rights. If you need to modify the dashboard, you must first make yourself the owner of the dashboard. You will then have full edit permissions for the dashboard.

The menu on the right of the dashboard tile is from where you can manage your or your users' dashboards. The example below includes the options available to you when the dashboard belongs to another user.



Note that the All Dashboards button is displayed only if the grid-view area of the Analytics page is available to you. If you have rebranded the whole analytics landing page, the All Dashboards button will not be displayed. If you rebrand just the top area of the analytics landing page, the All Dashboards button will be available.

# Assigning Security Rights

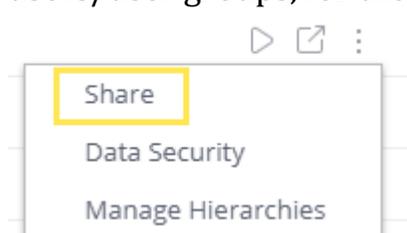
## Assigning Access Rights to an ElastiCube Server

By default, when you add an ElastiCube server (as described in Manage ElastiCube), all users are allowed to access it.

A user that has access to an ElastiCube Server can access all the ElastiCubes on it. Once you start assigning users/user groups access rights to the ElastiCube Server, then only the users/user groups that you assign can access the data of the ElastiCube Server.

### To assign access rights to access an ElastiCube Server:

1. Click Admin and select the Data Sources tab in the left menu.
2. To define that an ElastiCube Server can only be accessed by specific users/user groups, for the relevant ElastiCube, select  and click **Share**.



3. Enter the users/user groups that you want to give access to this ElastiCube Server.
4. Click Save.

## Assigning Access Rights to an ElastiCube

See [Sharing ElastiCubes](#) for more information.

# Data Security

Data security in Sisense can be divided into two types, data encryption and data access. Data communication is related to how data is secured by Sisense while being imported into Sisense and written on your server's disk.

When data is imported into Sisense, the protocol used depends on the protocols supported by the data source. Sisense supports importing data over SSL, if the source supports it. Sisense supports SSL for data Retrieval, for example, when viewing data in dashboards.

Configuration data, such as account credentials and authorization profiles, are encrypted prior to being written to the disk. The encryption technology used by Sisense includes:

1. SHA-256
2. TripleDES
3. AES-256

For data at rest, Sisense supports OS based disk encryption, Windows file system encryption - Transparent Data Encryption (TDE). When using Windows transparent encryption, the key pair (private/public) is bound to the user identity. For more information, [click here](#).

The second type of data security is data access. This type of data security refers to who can access your data after it's imported into Sisense and displayed in a dashboard.

## **What is Data Access Security?**

Sisense enables you to define Data Security Rules that control which users can access which portions of the raw data in an ElastiCube, down to row granularity.

For example:

- Each widget only shows the data permitted by the Data Security Rules that apply, including totals, averages and so on.
- The data browsers used while building dashboards and widgets only show the data permitted by the Data Security Rules that apply.
- You can define a single dashboard that automatically displays different results for each user (or user group), based on the rows that user is permitted to see.

### Use Case Example

- A Sales Order table has a column representing the salesperson that closed a deal.
- You created a quarterly performance dashboard for your salespeople, but want each of them to see only their own data.
- You do not want any of them to be exposed to data that represents the performance of others.

### How Does Data Access Security Work?

Each ElastiCube contains tables and each table contains fields. A Data Security Rule defines that a specific user can only see any data of an entire row of a table, if a specific field in that row has a specific value(s).

For example, in a Sales widget a salesperson (for example, Dan) will only see the sales amounts from the rows of a Sales ElastiCube whose Salesperson field contains the value Dan (rows 1 and 4).

### Sales Table

#	Salesperson	Product	Amount
1	Dan	HD-TV	\$100
2	Matthew	TV	\$300
3	Amber	Media Center	\$700
4	Dan	Player	\$200
5	Matthew	Air Conditioner	\$600

Dan will not see any part of a row in the ElastiCube that does not contain the value Dan in the Salesperson field, nor will any amounts from this row be included in totals.

**Note:** The entire row of data is not seen by the relevant user even when the field to which the rule applies does not appear in the widget.

If a widget that shows the amount spent per product is shared with Dan , then he will only see HD-TV and Player and the sales total will be 300\$.

## Defining Data Access Security for an ElastiCube

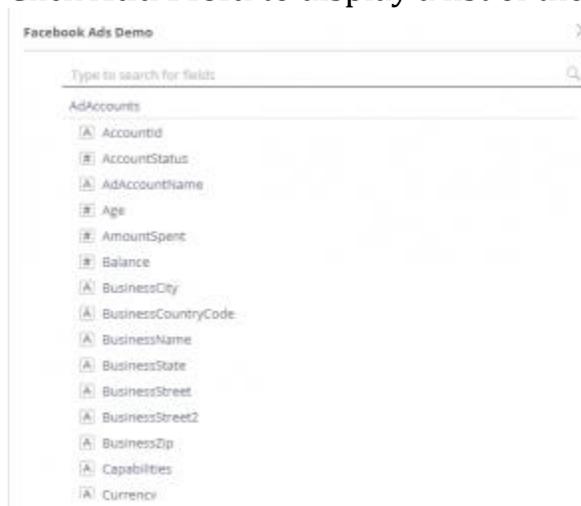
Each Data Security Rule applies to a specific field in an ElastiCube and to specific user(s)/user group(s). It enables you to define the values that must be contained in a specific field to enable that entire row of data to be available to a user.

### To access Data Security:

1. Click Admin and select the ElastiCube tab in the menu.
2. For the relevant ElastiCube, select and click Data Security.  
If no data security rules have yet been defined for this ElastiCube, then the following message is displayed:



3. Click **Add Field** to display a list of the fields in this ElastiCube.



4. Select a field. For example, **Brand**. The following window is then displayed in which you can define rules.

**Note:** You cannot select date type fields.



The left side of this table enables you to define which users/user groups can access this data. Click **+ Add Restriction** and start typing into the **Restricted User/Groups** field to get a drop-down list.

Add as many users/user groups, as necessary.

The right side of this window enables you to define which values the specified users/user groups are permitted to see.

Start typing into the **Values** field to get a drop-down list.



### Multiple Values can be Selected.

The value of numeric type fields must be typed into this field, as no auto-complete option appears for numeric type fields.

Alternatively, you can select:

- **Everything:** To specify that the selected users/user groups can see this data no matter what the value is in this field.
- **Nothing:** To specify that the selected users/user groups cannot see this data no matter what the value is in this field.

For example, you can define that the following Users/User Groups must have the following values in the Product Category column to enable them to see their data row in a widget.

#	User/User Group	Product Category
1	Management	Everything
2	Bob	Apple Mac Desktops
3	Don	Calculators, Camera Flashes
4	Everyone else	Nothing

This means that management can see the data of all Product Categories, Don can only see the data of Calculators and Camera Flashes, Bob can only see the data of Apple Mac Desktops, and Everyone else won't see anything.

### **How Does Data Level Security Work for Tables with Relationships?**

Tables in an ElastiCube may have a relationship between them.

As described above, each widget only shows any data of an entire row of a table, if a specific field in that row has a specific value.

In addition, a widget may further restrict the data shown to a specific user when a rule is defined for a table that has a relationship to a table that has a field in the widget.

This means that a widget only shows the data permitted by the **combined** Data Security Rules assigned to all the tables that have **any field in the widget**.

As described above, the entire data row is restricted even when the field to which the rule applies does not appear in the widget. The entire row of data is also restricted even when the field of the relationship between the two tables does not appear in the widget.

### **Use Case Example - Expanding Upon the Example Above**

- The Sales table has a column that has a relationship with a Deal Contacts table that holds the contacts that were involved in each deal.
- You created a Deal Contacts widget for your salespersons.
- As described in the example above, the Sales table has a Data Security Rule that maps each user to his/her matching field value, so that each sales person only sees their own data.
- Even though the Deal Contacts table does not have any Data Security Rules defined for it, the Deal Contacts widget only enables each sales person to see the contacts associated with their own sales, because of the Data Security Rule assigned to the Sales table.

# Security Bulletin

## **Update regarding Meltdown (CVE-2017-5754) and Spectre (CVE-2017-5753 and CVE-2017-5715)**

The Meltdown/Spectre vulnerabilities are the recent vulnerabilities found in Intel's processors. These are vulnerabilities in all of the Windows operating systems. Sisense provides an application and does not provide the server hardware or the Windows OS. These are provided by customers. As such, it's the customers' responsibility to secure their OS.

We recommend customers follow all security recommendations of the vendors who provided the customers OS. The current Windows recommendation is to install the latest available security updates for the OS. Sisense has run extensive performance tests on the OS patch (Windows patch) and found a minimal impact on performance.

For more information about these vulnerabilities and how to negate them, see:

[Meltdown and Spectre](#): A high-level description of the vulnerabilities.

[Project Zero](#): Describes the theory behind the vulnerabilities by those who discovered them.

[Intel Official Announcement](#): Describes Intel's response to the vulnerabilities.

[AMD Processor Security](#): Describes AMD's response to the vulnerabilities.

# Optional Security Hardening for Sisense Web Pages

This page provides additional security options that can be applied for hardening the security of Sisense web pages for your needs.

## Cookie Security

You can make your Sisense web page cookies more secure by adding a `secureFlag` to your web pages. This flag instructs the browser that the cookie should only be returned to the Sisense Web Application over encrypted connections (HTTPS).

To add a `secureFlag`, a new property, “`secureCookie`”, was added to the `POST /settings/security` endpoint in Sisense REST V0.9. This boolean adds a “`secureFlag`” to cookies of users who connect to the Sisense Web Application. This flag instructs the browser that the cookie should only be returned to the Sisense Web Application over encrypted connections (HTTPS).

## Strict Transport Security

HTTP Strict Transport Security (HSTS) is a method for preventing any communications from being sent over HTTP to the specified domain and allows only communication over HTTPS. This is useful for preventing man-in-the-middle attacks or users with invalid certificates from accessing your dashboards.

In the Sisense web.config file, you can add a custom header that informs the browser that it should never load a site using HTTP and should automatically convert all attempts to access the site using HTTP to HTTPS requests instead. Sisense has added the following line in the <customHeaders> section that is commented out by default.

```
<!-- <add name="Strict-Transport-Security" value="max-age=31536000" /> -->
```

You can remove the comments to restrict communication to HTTPS requests only.

For a more secure option, you can add **includeSubDomains** which restricts communication to HTTPS from all domains and their subdomains.

```
<add name="Strict-Transport-Security" value="max-age=31536000;
includeSubDomains/">
```

#### **To implement HTTP Strict Transport Security:**

1. Open the web.config file located at: C:\Program Files\Sisense\PrismWeb
2. Remove the comments (<!-- and -->) from the following lines under <customHeaders>  

```
<!-- <add name="Strict-Transport-Security"
value="max-age=31536000" /> -->
```
3. Save the web.config file.

## **X-Frame-Options and Content-Security-Policy Headers**

If you have embedded your Sisense dashboard into your site, you can configure an X-Frame-Options header to defend against clickjacking attacks. This will prevent other web pages from framing your dashboard by indicating whether or not a browser should be allowed to render a page in a <frame>, <iframe> or <object>.

#### **X-Frame-Options Header Types**

There are three possible values for the X-Frame-Options header:

- **DENY:** Prevents any domain from framing the content. If you do not embed your dashboard with iframes, use this option to prevent any domain from framing your dashboard.
- **SAMEORIGIN:** Allows only the current site to frame the content.
- **ALLOW-FROM uri:** Permits only the specified 'uri' to frame this page and prevents all other domains. (e.g., ALLOW-FROM http://www.example.com)

## Browser Support

The following browsers support X-Frame-Options headers.

Browser	DENY/SAMEORIGIN Support Introduced	ALLOW-FROM Support Introduced
Chrome	<a href="#">4.1.249.1042</a>	<a href="#">Doesn't support</a>
Firefox	<a href="#">3.6.9 (1.9.2.9)</a>	<a href="#">18.0</a>
Internet Explorer	<a href="#">8.0</a>	<a href="#">9.0</a>
Opera	<a href="#">10.50</a>	Not supported in Opera 26 or below
Safari	<a href="#">4.0</a>	<a href="#">Doesn't support – Supports CSP frame-ancestors instead</a>

## Setting X-Frame-Options

You can set the X-Frame-Options for your dashboards in the Sisense web.config file. This file contains a section where custom headers are defined for the webpages where your dashboards are hosted.

### To set the X-Frame-Options header:

1. Open the web.config file located at: C:\Program Files\Sisense\PrismWeb
2. Remove the comments (<!-- and -->) from the following lines under <customHeaders>
 

```
<!-- <add name="X-Frame-Options" value="ALLOW-FROM" /> -->
<!-- <add name="Content-Security-Policy" value="frame-ancestors :host" />
-->
```

For IE:

```
<add name="X-Frame-Options" value="ALLOW-FROM  
https://dashboardurl.com" />
```

For other browsers:

```
<add name="Content-Security-Policy" value="frame-ancestors  
https://dashboardurl.com"/>
```

For both:

```
<!-- <add name="X-Frame-Options" value="ALLOW-FROM  
https://dashboardurl.com" /> -->
```

```
<!-- <add name="Content-Security-Policy" value="frame-ancestors  
https://dashboardurl.com" /> -->
```

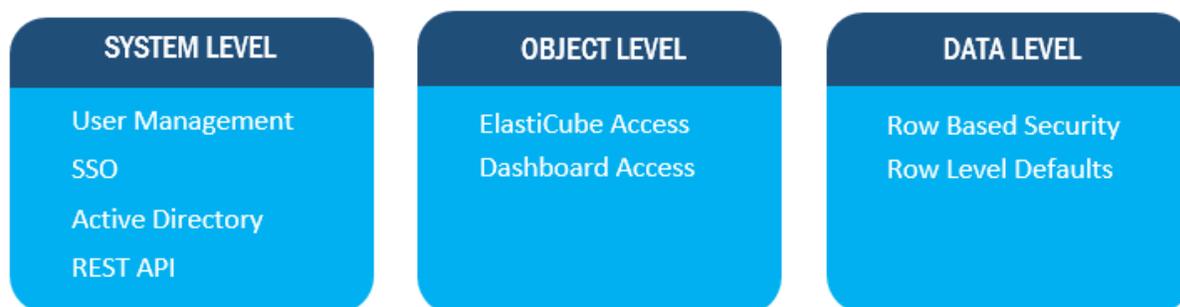
3. Save the web.config file.

# Sisense Security Overview

## Sisense Security Architecture

Sisense is built around a robust and flexible security architecture that is both comprehensive and intuitive. This architecture has been designed to ensure security processes are enforced while scaling to enterprise deployments of Sisense. This includes the ability to secure dashboards and data as well as implement custom security requirements that suit your organization. This section provides a general overview of the main security features.

Security is based around three levels associated with sets of security features. The diagram below maps this security architecture on a system, data and object level.



## System Level Security

System-level security encompasses security features for role-based settings and integration options. This includes user and server management, connection to an active directory, Single Sign-On (SSO) implementation, and use of the security REST API.

### User Management

You can assign five primary roles to Sisense users:

- Administrators
- Data Admin
- Data Designer
- Designer
- Viewer

These roles can be defined on either a user or group level to determine sharing, access and security. To learn more about user roles, [click here](#).

### **ElastiCube Server Access Rights**

You can assign access rights to different ElastiCube servers for individual users, groups or to everyone. These settings allow the management of different environments such as a testing and production server, or servers for specific projects or departments. See also ElastiCube Server and ElastiCube Security.

**Note:** Following improvements to ElastiCube security in Sisense V7.0 and later, ElastiCubes created prior to Sisense V7.0 are accessible to everyone by default, unless you have defined the ElastiCube's access rights. For ElastiCubes created after Sisense V7.0, the default access is only for the ElastiCube owner. When the ElastiCube is ready, it should be shared with the relevant people, or with "Everyone" if that is the desired access policy.

### **Active Directory**

Connect existing users and groups from your organization's Active Directory to define security and sharing properties and reduce deployment time. This removes 'password fatigue' as users can rely on existing credentials while organizational policies around security credentials such as updates can be enforced. See also [Integrating Active Directory](#).

### **Single Sign-On (SSO)**

SSO facilitates seamless integration between Sisense and other systems in your organization while offering standardization of authentication policies across your organization. This can improve user productivity by avoiding password fatigue and reduce support overhead. See also [Configuring SSO](#).

## REST API

The Security REST API provide access to parameters to integrate and automate restrictions and access control based on existing settings and standards. Specify access rights and security to dashboards, ElastiCubes and data. Manage users via the API to create, edit and assign new users or groups. [Click here](#) to visit the API documentation site.

## Object Level Security

Object security defines access rights for different users and groups to various components within Sisense. The two main objects are dashboards and ElastiCubes.

### Dashboards

You can share dashboards on either a user or group level. The sharing options include the configuration of access rights for all users as well as whether users defined as designers may edit the dashboard. The sharing options also include subscription settings that define which users and groups will receive email reports. See also [Sharing Dashboards](#).

### ElastiCubes

You can define access rights to different ElastiCubes on a user or group level. This enables flexibility to create ElastiCubes for specific user or group needs while offering strict access control. See also [Assigning Access Rights to an ElastiCube](#).

## Data Level Security

Data access must provide data to people only to the extent that they need to complete their jobs. Data Level Security provides the necessary control to enforce varying degrees of data visibility and access to support the separation of

duties. A single dashboard can be shared with many users, but each viewer sees only data relevant to their needs. This reduces both development time and provides for security.

### **Security on the Row Level**

You can grant user and group permissions to specific rows in the data. For each ElastiCube, you can apply multiple rules to enforce granular access control.

### **Row Level Defaults**

Control which data is accessible for users or groups that do not have explicit security rules. For example, enable new employees to access a restricted data set until they are added to relevant groups. You can set defaults to include everything, nothing or view based on a security rule.

See more in Data Security below.

## **Security Levels**

Sisense provides three levels of security:

- [ElastiCube Server Security](#)
- [ElastiCube Security](#)
- [Data Security](#)

## **ElastiCube Server and ElastiCube Security**

### **What is ElastiCube Server and ElastiCube Security?**

Sisense enables you to define access rights to control which users can access which ElastiCubes, whether they are creating new dashboards or trying to access shared dashboards.

### **ElastiCube Server Security – Use Case Example**

You may already have access to an ElastiCube Server named Production. As an administrator, you want to create a Staging environment with its own ElastiCube

Server; and you do not want other users to be exposed to it. In this case, you can set the ElastiCube Server rights to allow only you to access it.

**Note:** Following improvements to ElastiCube security in Sisense V7.0 and later, ElastiCubes created prior to Sisense V7.0 are accessible to everyone by default, unless you have defined the ElastiCube's access rights. For ElastiCubes created after Sisense V7.0, the default access is only for the ElastiCube owner. When the ElastiCube is ready, it should be shared with the relevant people, or with "Everyone" if that is the desired access policy.

### **ElastiCube Security – Use Case Example**

You may have an ElastiCube named Marketing and only want the CEO and Marketing team to have access to it. You can grant rights only to them using ElastiCube Security, thus denying anyone else access.

### **How Does ElastiCube Security Work?**

- You can define that users/user groups have access to either:
  - An entire ElastiCube Server, including all the ElastiCubes on the ElastiCube Server.
  - and/or
  - Only to specific ElastiCubes.
- By default, all users are allowed to access the data of all ElastiCube Servers and ElastiCubes. Once you start assigning users/user groups access rights to an ElastiCube Server/ElastiCube, then only the users/user groups that you assign can access that ElastiCube Server/ElastiCube.
- Access to ElastiCube Servers/ElastiCubes is by Grant rules. There are no Deny rule options.
- When no access rights have been defined for a specific ElastiCube, then that ElastiCube inherits its access rights from the ElastiCube Server settings.
- When a user attempts to access a Dashboard using a direct link and that dashboard is based on an ElastiCube to which that user does not have access rights, a security message is displayed.
- When access rights are defined for a specific ElastiCube, then those access rights override the ElastiCube Server access rights. This means that a user can access an ElastiCube to which they are assigned access, even if that user has not been assigned access rights to the ElastiCube server on which that ElastiCube is stored.

# Security at Sisense

Sisense provides a variety of ways in which you can control which data is exposed to users and secure your connections to Sisense.

The following topics describe how you can configure and maintain the security of your data in Sisense:

- [Security Overview](#)
- [Security Bulletin](#)
- [Assign Security Rights](#)
- [Data Security](#)
- [SSL](#)
- [Optional Security Features](#)

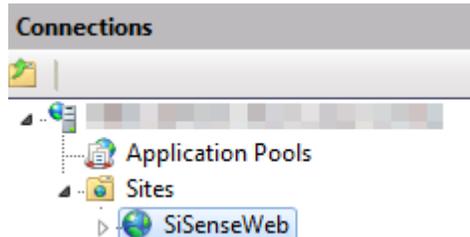
# Setting Up SSL

SSL (Secure Sockets Layer) is a security protocol that encrypts the server-client data channel. In Sisense, SSL secures the link between the Sisense server, specifically the IIS web server included in the Sisense installation, and the clients' browsers (where dashboards are created and viewed). With SSL in place, you can access Sisense via an HTTPS secure connection, which is password protected. This is recommended when you want users to be able to access dashboards from outside the organization.

**Note:** Configuring SSL on your Sisense server and in your IT environment should be performed by an IT Specialist or Web Administrator.

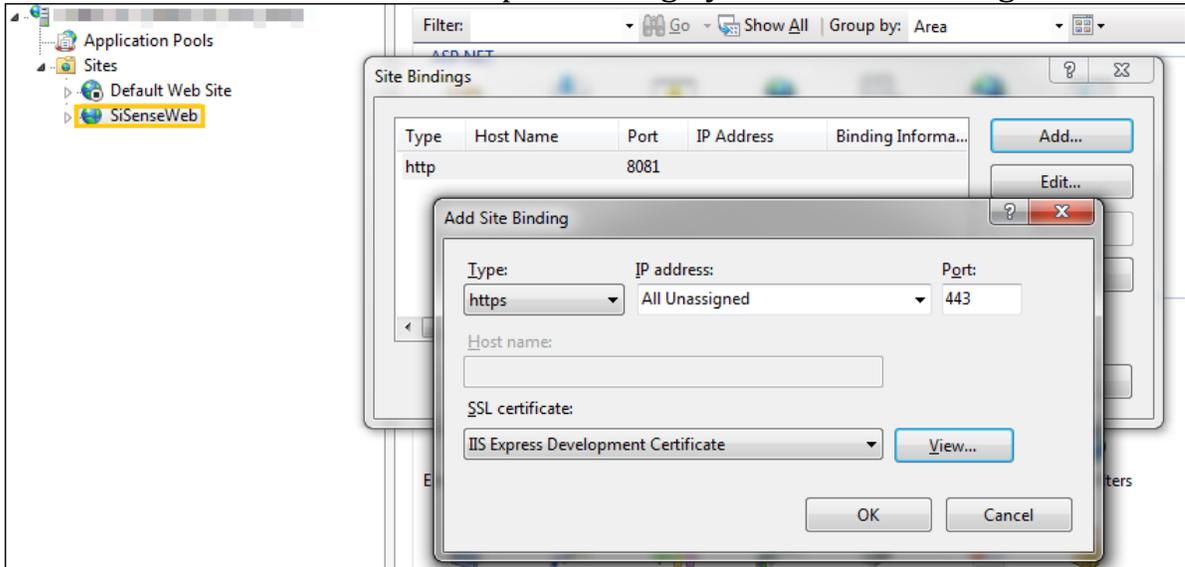
The following prerequisites and steps are required for setting up SSL in Sisense:

1. Make sure your Sisense instance is using the IIS web server and not the default IIS Express server. You can check this by verifying the following settings:
  1. Open your IIS Manager from the Windows Start menu. (The name of the IIS Manager may differ between different versions of Windows.)
  2. Expand the server, and under Sites, you should see SisenseWeb or PrismWeb sites.



3. Open your Control Panel, and open **Administrative Tools > Services**. Verify that the SisenseWeb.Express (Prism.WebExpress for earlier versions) does not appear in the list of services. If you need help moving over from IIS Express to IIS, click [here](#).
2. Install a valid SSL certificate on your IIS server.
3. In your IIS Manager, add HTTPS to the Site Bindings settings. Right-click your web site and select Edit Bindings... Click Add, and select HTTPS, and the IP Address. The default port for HTTPS is 443.

**Note:** You must maintain at least one non-HTTPS port bound to the website, which Sisense uses for internal communication. You can still restrict external traffic over this port through your firewall settings.



4. If your organization is rerouting HTTP requests to HTTPS, directly to your web server, please consult your IT specialist for the correct binding and routing configuration.
5. After completing the above steps, the Sisense Web site will answer for both secure and non-secure requests (HTTP & HTTPS). If you want your users to be automatically redirected to the HTTPS session, please read the instructions in this [support post](#).
6. If you use a reverse proxy or load balancer:
  1. Navigate to the following folder: "C:\Program Files\Sisense\PrismWeb\AnalyticalEngine\dist\config" and open the production.json file with a text editor.
  2. Add the line "baseUrl": "XXXXXXXXXXXX", before the line "forkEnable", after the bracket where XXXXXXXXXXXXX should be replaced with the URL of the Sisense server, followed by a comma. Note: If your baseUrl is not on port 80, the URL should include the port number, for example, "http://localhost:8081" or "http://yourdashboardsite.com:8081". Do not replace the value of Port in the config file with the baseUrl's port number. In addition, do not modify the values of "forkEnable" or "Port" when attempting to configure secure connections for exporting tables to Excel. Example:

```
{
  "baseUrl": "http://yourserver.com:8081",
```

"forkEnable":true,

"Port":40003

}

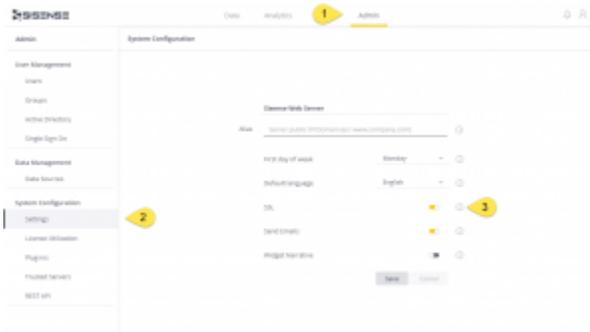
3. Restart the IIS service.

7. To enable HTTPS in the URL links for your shared dashboard reports and email invites, you must enable SSL in your System Configuration settings.

1. In your Sisense Web Application, click **Admin** and select **Settings**.

2. Toggle the SSL switch to on.

3. Click **Save**.



# Data Management

The topics below describe how you can share ElastiCubes with other users and you can create drill hierarchies that you want to make available for Designers:

- [Sharing ElastiCubes](#)
- [Adding and Removing ElastiCube Servers](#)
- [Managing Drill Hierarchies](#)

# Managing Drill Hierarchies

Dashboard viewers can drill down in most widgets to get an in-depth view of a selected value. To make it easier for viewers to select drill-down paths, dashboard designers can add predefined drill hierarchies to their widgets. As an Administrator, you can create a pool of drill hierarchies, from which Dashboard Designers can select the drill hierarchies they want to include in the widgets they build.



This section describes how to create and manage the drill hierarchies that you want to make available for designers.

## To create a drill-down hierarchy:

1. Click **Admin** in the Sisense Web Application. In the left menu, select **Data Sources**.
2. For the relevant ElastiCube , click and select **Manage Hierarchies**.
3. If this is the first hierarchy you are adding, select **Add Hierarchy**. If you already have hierarchies, click **Add Hierarchy** in the top-right corner.
4. Select the field. This will be the first field in the drill hierarchy, and the same field used in the widget.
5. Click **Add Field** to add the next level in the drill hierarchy. Repeat this step for as many levels as you need.

Source

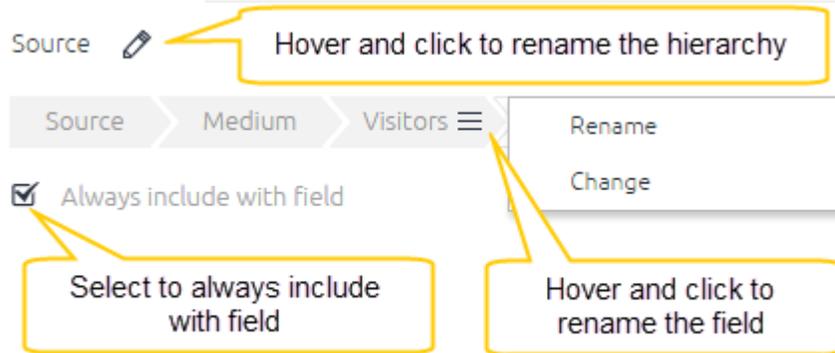


Additional options include:

**Always include with field:** Select this checkbox to always include the drill hierarchy in the widget for this field.

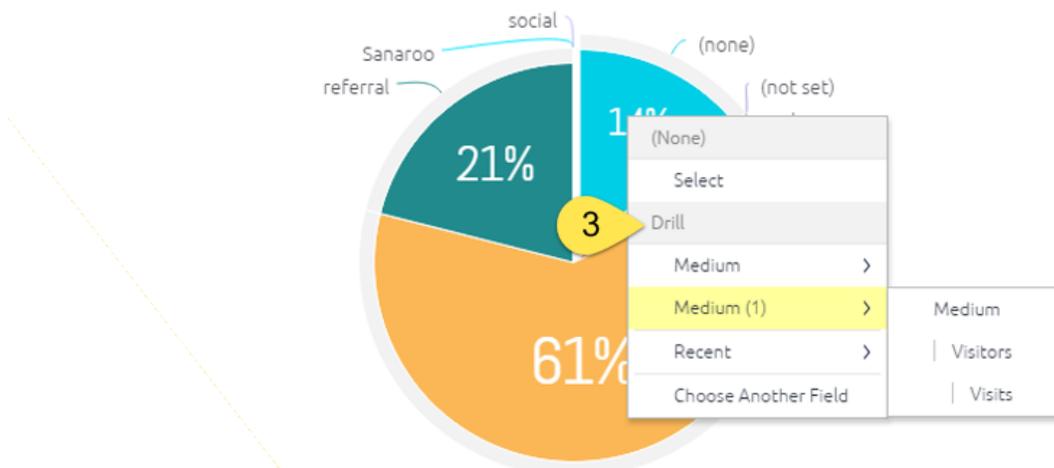
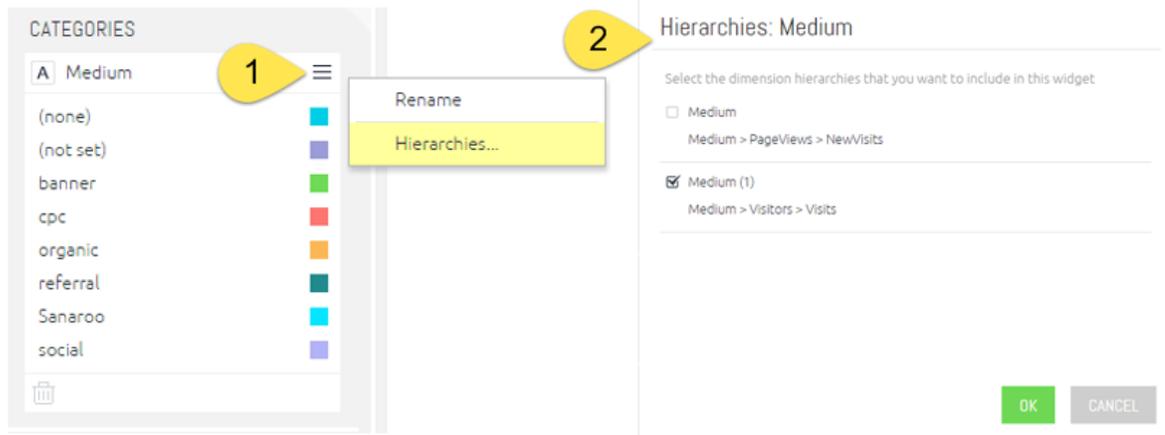
**Editing the Hierarchy name:** To edit the name of the drill hierarchy, click on the pencil icon, and type in the name. Click on the tick icon to confirm the change.

**Editing the field name:** To edit the name of a field in the hierarchy, hover over the field name and click on the menu > Rename. Type in the name, and click on the tick icon to confirm the change.



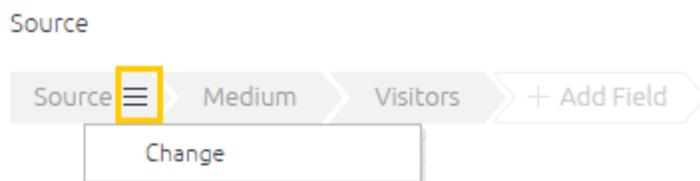
6. Add as many hierarchies as you need for the ElastiCube. When done, dashboard designers will have access to the hierarchies in the widget

designer, and will be able to select and add them to the widget.



### To change a field in a hierarchy:

- Hover over the field in the hierarchy, and click on the menu that appears. Click **Change** and select a different field.



### To delete a hierarchy:

- Click on the bin to the right of the hierarchy.

# Adding and Removing ElastiCube Servers

This topic describes how to manage your ElastiCubes.

By default, the Sisense Web Application is connected to the ElastiCube server that is installed with it on the same machine.

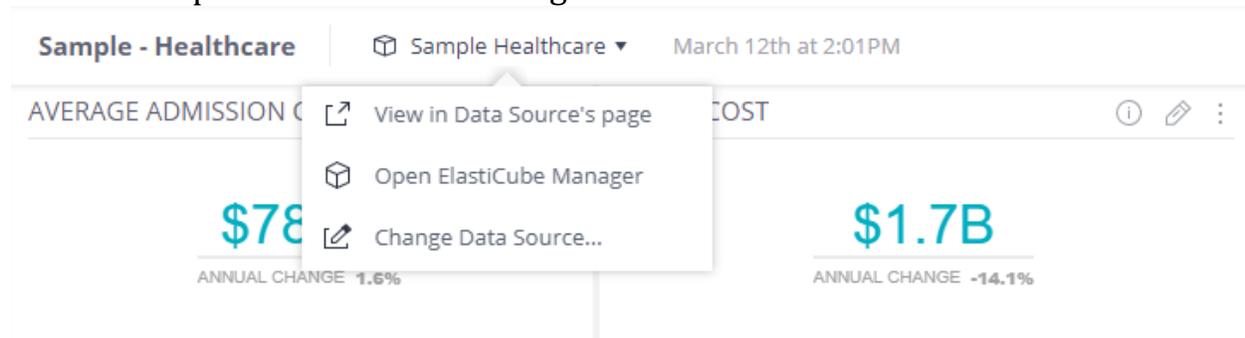
By default, all ElastiCubes on that machine are accessible to all dashboard users, and are visible under LOCALHOST.

## Opening an ElastiCube in the ElastiCube Manager

You have a number of ways to open your ElastiCube in the ElastiCube Manager:

From the Sisense Web Application:

- If you have an open dashboard, click on the ElastiCube name next to the dashboard name. Select View in ElastiCube's Page to open the ElastiCube management page in the web app. In the ElastiCubes page, select the ElastiCube and click Open. To open the ElastiCube of the open dashboard, click Open in ElastiCube Manager .



- From Windows: From the Start menu, open Sisense ElastiCube Manager.

# Adding Another ElastiCube Server

## To add another ElastiCube Server to Sisense:

1. Click **Admin** and select the **Data Sources** tab on the left.
2. Click **Add Server**.

**Add Server**

IP or Server Name

Who can access this server?  Admins Only  Everyone ⓘ

Add Server Cancel

3. Select whether this ElastiCube Server is intended for Admins only. This selection also applies to all new ElastiCubes that will be added to this ElastiCube Server. A user that has access to an ElastiCube Server can access all the ElastiCubes on it.
4. Add the IP of the ElastiCube Server to which you want to connect and click Add Server. The new ElastiCube Server is added to the list with all its ElastiCubes.
5. Click Save. Once a new ElastiCube Server has been added, its ElastiCubes immediately become available to all designers (by default) when they create a new Dashboard.
6. To define which users are allowed to access the data of an ElastiCube Server, see [Assigning Rights to an ElastiCube Server](#).
7. To define which users are allowed to access each specific ElastiCube, see [Assigning Rights to an ElastiCube](#).

# Removing an ElastiCube Server

## To remove an ElastiCube server from the list:

- In the Data Sources page, for the relevant server, click the Trash  button.

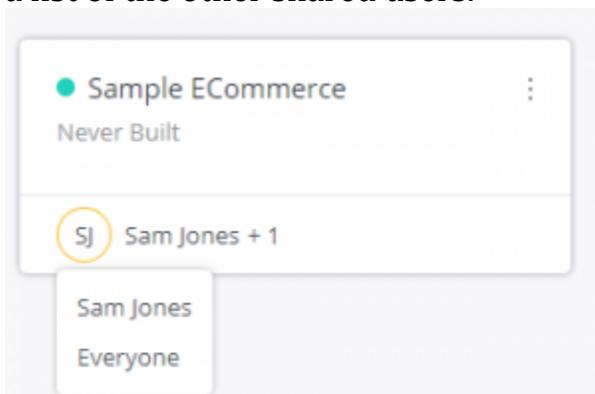
**Note:** The LOCAL MACHINE server cannot be removed.

# Sharing ElastiCubes

By default, all users are allowed to access each ElastiCube.

You can share ElastiCubes with other Sisense users in your account through the Admin section of the Sisense Web Application.

ElastiCubes that are shared with other users are displayed in the Data page with a list of the other shared users.



Once you start assigning specific users/user groups access rights to this ElastiCube, then only the users/user groups that you assign can access the ElastiCube.

When you share an ElastiCube, the recipient's user role determines what access you can assign to them.

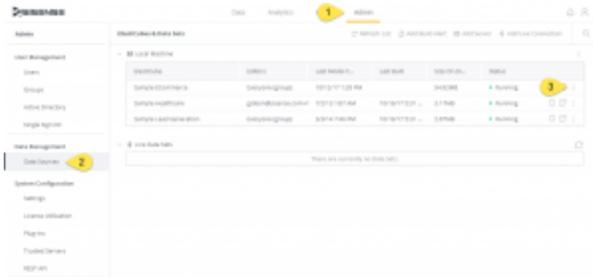
If the recipient is a Data Designer, you can choose to give them Query or Edit rights. Query rights allow the user to query or create a dashboard based on the ElastiCube while Edit rights allow the user to modify the ElastiCube.

For Administrators and Data Administrator, these users automatically receive Edit rights and Viewers and Dashboard Designers automatically receive Query rights.

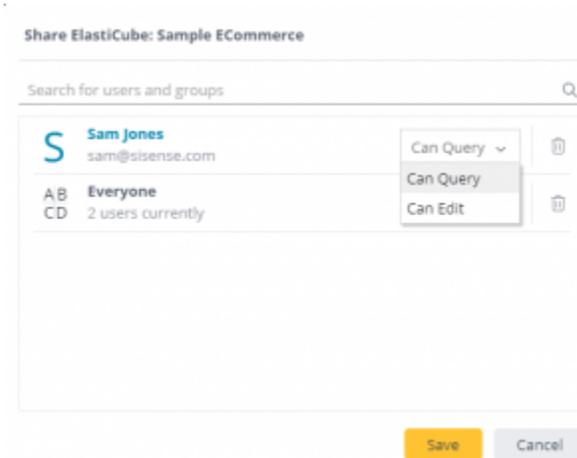
Sharing an ElastiCube does not affect its ownership.

**To share an ElastiCube:**

1. In the Sisense Web Application, select **Admin > Data Sources** and then the menu button for the relevant ElastiCube.



2. Click **Share**.



3. Enter the emails or groups of users with whom to share the dashboard. Selecting Everyone specifies that this dashboard is shared with all Sisense users and user groups.
4. The contact information of the people you add is listed in the center of the window and is displayed the next time you open this window.
5. If a person does not have a Sisense account, they receive a link that guides them through the activation process.
6. In the dropdown menu next to each user/user group select what rights the user has. You can only select the rights for a Data Designer as all other roles are predefined.
7. Click **Save**.

# Embedding Dashboards and Widgets

This article is about getting started with embedded Sisense dashboards. If you're new to Sisense, check out our [embedded analytics solutions](#).

You can embed Sisense dashboards and widgets in non-Sisense environments, such as in your own websites or applications by concatenating parameters to the URL of your dashboard or widget.

In addition, you can customize what aspects of the dashboard and widget are displayed in your environment.

For example, if you want to remove the Sisense header from the dashboard, you can concatenate the “?embed=true&h=false” parameter in the URL of the dashboard to hide the header.

`mysite.com/app/main#/dashboards/5541dc7a80a4e2181e00011a?embed=true&h=false`

In this example, the “embed=true” embeds the dashboard into your environment. When you embed a dashboard or widget, viewers can only view the dashboard or widget and do not have any editing privileges. The “h=false” parameter hides the Sisense header allowing you to whitelabel the dashboard.

This article provides a full list of the parameters you can concatenate, their functionality, and examples.

# Embedding Dashboards

To embed a dashboard, add “?embed=true” to the URL of the dashboard, and refresh the page. If your dashboard is within a folder, then first delete “?folder=” and everything that follows, before adding “?embed=true”.

This will open the dashboard without the surrounding Sisense environment.

For example:

Dashboard URL:

`mysite.com/app/main#/dashboards/5541dc7a80a4e2181e00011a`

Embedded URL:

`mysite.com/app/main#/dashboards/5541dc7a80a4e2181e00011a?embed=true`

Dashboard URL for dashboard within folder:

`mysite.com/app/main#/dashboards/550952417404b2981a000029?folder=550955a27404b2981a00003b`

Embedded URL for dashboard within folder:

`http://localhost:8081/app/main#/dashboards/550952417404b2981a000029?folder=550955a27404b2981a00003b`

You can use this URL to embed the dashboard in an iframe, for example:

```
<iframe id="ifm" name="ifm" width="100%" height="100%"
frameborder="0"
src="http://mysite.com/app/main#/dashboards/536f2b70d093e
26c280000d5?embed=true" scrolling="auto"></iframe>
```

## Dashboard Embedding Configuration Options

You can control which aspects of the Sisense environment are available when embedding by concatenating the following parameters to the embedded URL:

**Show/Hide the left navigation panel**

Parameter: "l" (left) or "i" to hide it completely

Default: false

Example: ?embed=true&l=true

### Show/Hide the right filter panel

Parameter: "r" (right)

Default: true

Example: ?embed=true&r=false

### Show/Hide the toolbar

Parameter: "t" (toolbar)

Default: false

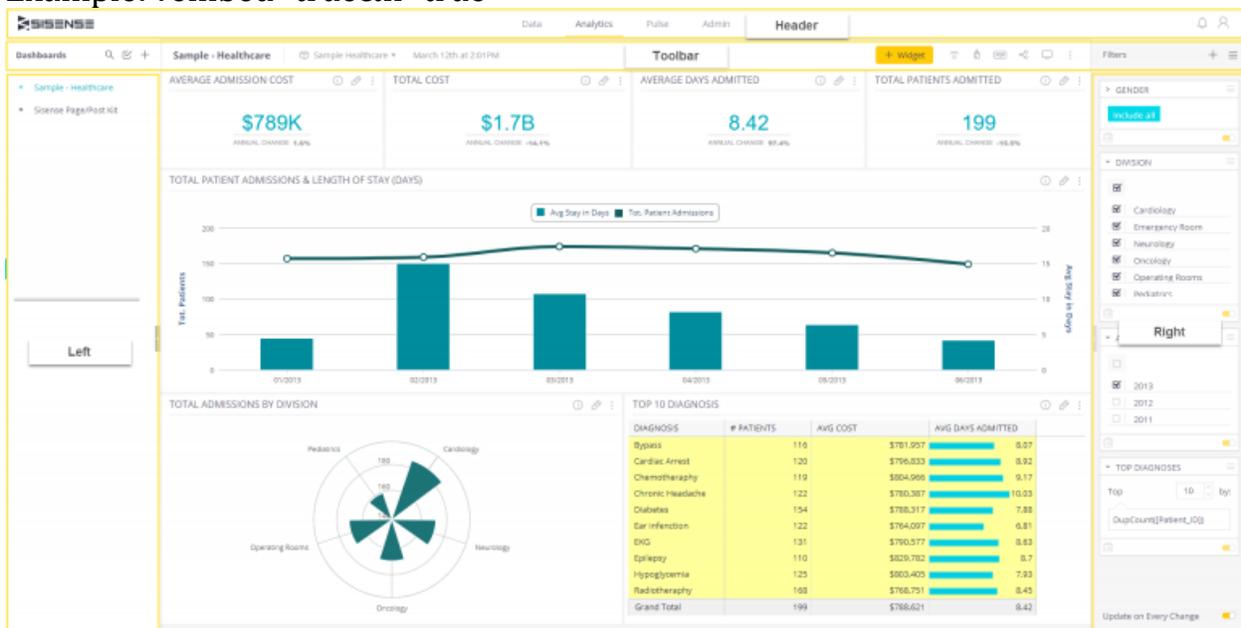
Example: ?embed=true&t=true

### Show/Hide the environment header

Parameter: "h" (header)

Default: false

Example: ?embed=true&h=true



# Embedding Widgets

To embed a single widget, add “?embed=true” to the URL of the widget when it edit mode.

This will open the widget without the surrounding Sisense environment.

For example:

## **Dashboard URL:**

<http://mysite.com/app/main#/dashboards/536f3a54d093e26c28000114/widgets/536f3a54d093e26c2800011b>

## **Embedded URL:**

<http://mysite.com/app/main#/dashboards/536f3a54d093e26c28000114/widgets/536f3a54d093e26c2800011b?embed=true>

You can use this URL to embed the widget in and iframe, for example:

```
<iframe id="ifm" name="ifm" width="100%" height="100%" frameborder="0"
src="http://mysite.com/app/main#/dashboards/536f3a54d093e26c28000114
/widgets/536f3a54d093e26c2800011b?embed=true"
scrolling="auto"></iframe>
```

## Widget Embedding Configuration Options

You can control which aspects of the Sisense environment are available when embedding by concatenating the following parameters to the embedded URL:

### **Show/Hide the left data panel**

Parameter: “l” (left)

Default: false

Example: ?embed=true&l=true

### **Show/Hide the right design panel**

Parameter: “r” (right)

Default: false

Example: ?embed=true&r=false

### **Show/Hide the toolbar**

Parameter: “t” (toolbar)

Default: false

Example: ?embed=true&t=true

### **Show/Hide the environment header**

Parameter: “h” (header)

Default: false

Example: ?embed=true&h=true

## Authentication in Embedded Dashboards and Widgets

All Sisense dashboards and widgets require authentication. When embedding dashboards and widgets outside of the Sisense environment, usually the use case is to not enforce the Sisense environment authentication, but rather use SSO (Single Sign On) with existing corporate authentication. To read more on configuring Single Sign On in Sisense, [click here](#).

# Displaying Custom URLs

You can display your company's URL in the address of your Sisense dashboards by defining an alias URL from the Sisense Admin page.

This is useful when you are providing OEM services and you want to mask Sisense's URLs and branding.

## To implement an alias URL in Sisense:

1. Click **Admin** in the top menu, and then **Settings** in the left menu.
2. In the **Alias** field, enter the web server's alias or IP address. This is especially important when you have configured your IIS with a domain name (alias), but when you share a dashboard, the IP address and/or the incorrect port is sent instead. Reports will be sent from the alias entered in this field.
3. Click **Save**.

# Rebranding Sisense Automated Emails

Sisense automated emails are sent to users according to predefined scenarios.

When a predefined scenario such as a build alert is triggered, an automated email is sent to the relevant Sisense users.

You can customize these emails to match your company's branding.

This page describes the Sisense automated emails and how you can replace them with your own so each time a predefined scenario is triggered, Sisense sends your customized emails to your Sisense users.

## Sisense Automated Emails

You can rebrand the following automatic emails:

- **Build Alert:** Email sent each time a build alert is triggered.
- **Dashboard Errors Reports:** Email sent when a dashboard report fails.
- **Dashboard Errors Reports Bulk Fail:** Email sent when a large amount of email reports fail.
- **Dashboard Reports:** Email sent to the recipient of a shared dashboard when you share a dashboard.
- **Empty Dashboard Report:** Email sent to an administrator when an empty dashboard report is sent.
- **KPI Alert:** Email sent each time a KPI alert is triggered.
- **New User Invitation:** Email sent to a new user invited to Sisense.
- **Password Recovery:** Email sent to a user who requested a new password.
- **Share ElastiCube:** Email sent to the recipient of a shared ElastiCube.
- **Share With:** Email sent to the recipient of a shared dashboard.
- **Transfer Ownership:** Email sent to a user who received ownership of a dashboard from another owner.
- **User Created:** Email sent to a user when a user is created.

- **User Created AD:** Email sent when a user is created in Sisense after being added from Active Directory.

For each automated email, Sisense provides an email template in the format of EJS and LESS files. These files are located within template folders inside the Template Parent folder located at:

```
...Program  
Files\Sisense\PrismWeb\vnext\src\features\emails\templates
```

## LESS Files

All templates contain the style.less file that refers to a stylesheet that defines the branding used in Sisense's automated emails. The content of every style.less is as follows:

```
@import '../styles.less';
```

The stylesheet referred to by all the email templates is called styles.less located at:

```
...\Program  
Files\Sisense\PrismWeb\vnext\src\features\emails\templates
```

By replacing the content of the styles.less file with your own CSS, you can redefine the branding used in all of the email templates to your custom branding.

## EJS Files

Each template folder contains an EJS file, which contains HTML and JavaScript that describe the content of the email.

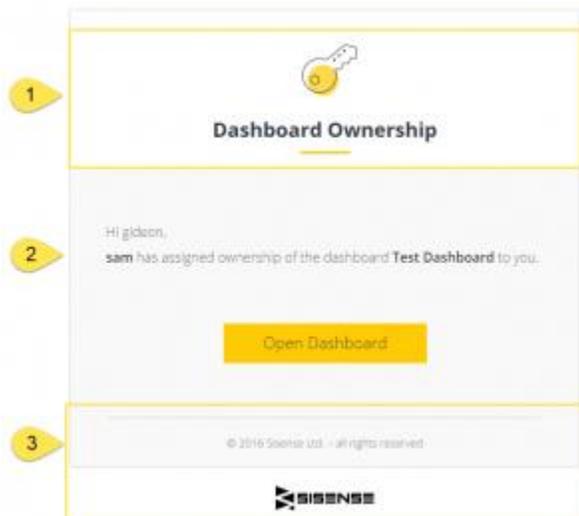
You can use any text/HTML editor to modify the HTML.ejs files according to your needs.

The following is an example of an HTML.ejs file.

```

<% include ../header %>
  <tr>
    <td>
      <div class="contentText">
        <p class="userText" style="margin: 0;"><%=
i18nContent.hi %> <%= newOwnerUserName %>,</p>
        <p class="generalInfo" style="margin:
0;"><%= i18nContent.generalInfo %></p>
      </div>
    </td>
  </tr>
<% include ../footer %>

```



Each EJS file has three sections, a header, body, and footer; each of which describes various areas in the email your users receive.

Section 1 is the Header section defined in the EJS file as `<% include ../header %>`.

This is a reference to the header.ejs file located at:

...\Program

Files\Sisense\PrismWeb\Resources\branding\emails

The header.ejs file defines the title and image used in the header section of each email template. You can modify the header by replacing the header.ejs file with or replace the `<% include ../header %>` tag from each template with your own HTML and JavaScript.

Section 2 is the Body section, which contains the content displayed in the email your users receive. There are two types of content displayed in automated emails, tokens and strings. Tokens refer to Sisense variables such as the Sisense username or dashboard owner's name. For example, the `<%=newOwnerUserName %>` token displays the new owner of a dashboard's Sisense username in the automated email.

```
<tr>
  <td>
    <div class="contentText">
      <p class="userText" style="margin: 0;"><%=
i18nContent.hi %> <%= newOwnerUserName %>,</p>
      <p class="generalInfo" style="margin: 0;"><%=
i18nContent.generalInfo %></p>
    </div>
  </td>
</tr>
```

When you replace the HTML.ejs file with your own files, you can still use the Sisense tokens. For more information about tokens you can add to your automated emails, see [Applying Sisense Tokens](#).

In addition, each template contains an object such as `i18nContent.hi` that displays strings predefined by Sisense. These strings cannot be modified, however, you can remove and replace the object with your own content.

The final section, Section 3 contains the footer. Like the header, the footer, `<% include ../footer %>`, is a reference to a footer.ejs file located at:

```
...\Program Files\Sisense\PrismWeb\Resources\branding\emails
```

The footer.ejs file defines the copyright text and image used in the footer section of each email template. You can modify the footer by replacing the footer.ejs file with or replace the `<% include ../footer %>` tag from each template with your own HTML and JavaScript.

Images that are displayed in the Header and Footer are stored in the following location:

```
...\Program  
Files\Sisense\PrismWeb\vnext\src\features\emails\templates\images
```

You can replace these with your own images by maintaining the same file names used, however, it is recommended to replace the content of the Header and Footer sections altogether with your own content.

## Replacing Sisense Automated Emails

If you prefer to implement your own email templates for Sisense's predefined scenarios, you can replace Sisense's emails with your own through the Sisense REST API.

Note: Rebranding Sisense emails must be enabled for your license. Contact your Account Manager for more information.

For each automated email, Sisense provides an email template in the format of EJS and LESS files. These files are located within template folders inside the Template Parent folder located at:

```
...Program  
Files\Sisense\PrismWeb\vnext\src\features\emails\templates
```

Some template folders may contain an html-org.ejs file. These files are legacy templates and are maintained for backwards compatibility. Newer developments should only modify or replace the html.ejs files.

**Note:** While it is possible to modify the content of your automated emails through the language.js file located in the templates folder, the recommended method is modify the email templates as described in the procedure below.



6. In the email object define the following information:

**Note:** If you have already rebranded your Sisense Web Application, copy the current settings through the GET /branding endpoint and paste them into the POST /branding endpoint modifying or adding the emails object. If you only modify the emails object, the remaining objects and keys are set to their default Sisense and will overwrite any existing settings.

Key	Description
senderEmail	The email address from which the automated email is to be sent. When you define the senderEmail value, you must verify your domain in Sisense's Mandrill email service. Click <a href="#">here</a> for more information.
senderName	The name to be displayed as the Sender in the automated email.
templates_directory	The directory location of your template folders. Sisense accepts /resources/branding/emails as a custom location for automated email templates. This directory is located on your Sisense server at ...Program Files\Sisense\PrismWeb\Resources\branding\emails
passwordRecoverySubject	The subject line of the password recovery email.
newUserInviteSubject	The subject line of the new user email.
createdUserSubject	The subject line of the new user invite email.
shareWithNewUserSubject	The subject line of the shared dashboard with a new user email.
shareWithExistingUserSubject	The subject line of the shared dashboard email.

Key	Description
transferOwnership	The subject line of the dashboard transfer ownership email.

For example:

```
"emails": {
  "senderEmail": "sys-admin@company.com",
  "senderName": "sys-admin",
  "templates_directory": "/resources/branding/emails",
  "passwordRecoverySubject": "Recover Password",
  "newUserInviteSubject": "new User Account Activation",
  "shareWithNewUserSubject": "A Dashboard has been
shared with you:",
  "shareWithExistingUserSubject": "A Dashboard has been
shared with you:"
},
```

7. Click **Run**. The next time an automated email is sent, your templates are sent instead of Sisense's predefined emails.

## Applying Sisense Tokens

Most Sisense automated emails include tokens that you can apply to your automated emails. When the email is sent, Sisense populates the token with the relevant value. For example, the `<%= user %>` token represents the Sisense's username.

Each template has tokens defined by Sisense that you can apply listed below. Tokens from one automated email cannot be applied to another. Applying tokens to other automated emails, for example adding `<%= user %>` to the KPI alert email results

**build\_alert:** Sent when a build alert is triggered.

## No Template Tokens Available

**Dashboard\_errors\_report:** Sent when an Dashboard email report fails.

## Template Tokens Available

<%= user %>

<%= errors[i].user %>

<%= errors[i].error %>

<%= errors[i].widgetFailed %>

<%= errors[i].widgetCount %>

**Dashboard\_errors\_report\_bulk\_fail:** Sent when a Dashboard email report sent to a large amount of users fails.

## Template Tokens Available

<%= user %>

<%= errors[i].user %>

<%= errors.groups[i] %>

**Dashboard\_report:** An automated email that contains a dashboard report sent to users.

## Template Tokens Available

<%= url %>

<%= images[i] %>

**Empty\_dashboard\_report:** An automated email that contains an empty dashboard report sent to users.

## Template Tokens Available

<%= user %>

**Kpi\_alert:** Sent to shared users when a KPI alert is triggered.

## Template Tokens Available

<%= measure %>

<%= value %>

<%= message %>

**New\_user\_invitation:** Sent to a new user that you have invited to join.

Template Tokens Available

<%= owner %>

\* The folder for this template contains a file called HTML-ORG.ejs. This file is for older versions of Sisense and is maintained for backward compatibility.

**Password\_recovery:** Sent when a user requests their password from the Forgot Password page.

Template Tokens Available

<%= user %>

\* The folder for this template contains a file called HTML-ORG.ejs. This file is for older versions of Sisense and is maintained for backward compatibility.

**Share\_with:** Sent to a new recipient of a dashboard when the dashboard is shared with them.

Template Tokens Available

<%= user %>

\* The folder for this template contains a file called HTML-ORG.ejs. This file is for older versions of Sisense and is maintained for backward compatibility.

**transfer\_ownership:** Sent to the new owner of a dashboard when ownership is transferred.

Template Tokens Available

<%= newOwnerUserName %>

<%= oldOwnerUserName %>

<%= sharingDashboardName %>

**User\_created:** Sent to the recipients of new Sisense accounts.

Template Tokens Available

<%= user %>

**User\_created\_ad:** Sent when a user is created in Sisense after being added from Active Directory

## Template Tokens Available

<%= user %>

# White Labeling Sisense

Sisense enables you to rebrand (white label) the Sisense Web Application into your native application or site by using the Sisense REST API to replace existing icons and images with your own.

In addition, you can rebrand the automated emails Sisense sends to match your company's branding.

White labeling must be enabled in your license. To enable this feature, please contact your Sisense representative or open a request through our Help Center.

## Rebranding Sisense

You can rebrand the Sisense Web Application and the automated emails sent by Sisense through Sisense's REST API.

There are two steps to rebranding the Sisense Web Application. The first step is to place your branded content into the Sisense branding directory where Sisense is installed. For a list of the files you can rebrand and their specific dimensions, see [White Labeling Configuration Options](#).

The second step is to access the Sisense REST API, where you can embed your content directly into Sisense and modify text displayed to customers.

### **To rebrand the Sisense Web Application:**

1. Create a subfolder with your brand name in this directory (on the machine where Sisense is installed): "C:\Program Files\Sisense\PrismWeb\Resources\". Put your branded files in this directory. The following files are supported:
  - Desktop Logo (for example: Logo-Desktop.png) – This is the main icon that is visible at the top left of the Sisense Web Application.
  - Tablet Logo (for example: Logo-Tablet.png)
  - Mobile Logo (for example: Logo-Mobile.png)

- Favicon (for example: favicon.ico)  
**Note:** Optimal dimensions for logo images: 92×26
  - In the Sisense Web Application, click **Admin**, and then **Rest API**, followed by **REST API Reference**.
2. In the REST API Reference section, select version 0.9 in the top right corner.
  3. Click /branding and then click the POST operation /branding.
  4. To get the required JSON format for the request, click the sample model schema to place the sample text in the body parameter.

POST /branding Adds new branding to your Sisense dashboards

**Implementation Notes**  
 Branding includes customized texts, logos and emails.

**Response Class (Status 200)**  
 Model | Model Schema

```

{
  "logo": {
    "desktop": {
      "small": "string",
      "large": "string"
    },
    "tablet": {
      "small": "string",
      "large": "string"
    },
    "phone": {

```

Response Content Type: application/json

**Parameters**

Parameter	Value	Description	Parameter Type	Data Type
body	<pre> {   "logo": {     "desktop": {       "small": "string",       "large": "string"     },     "tablet": { </pre>	branding object that needs to be added to the collection	body	Model   Model Schema

Parameter content type: application/json

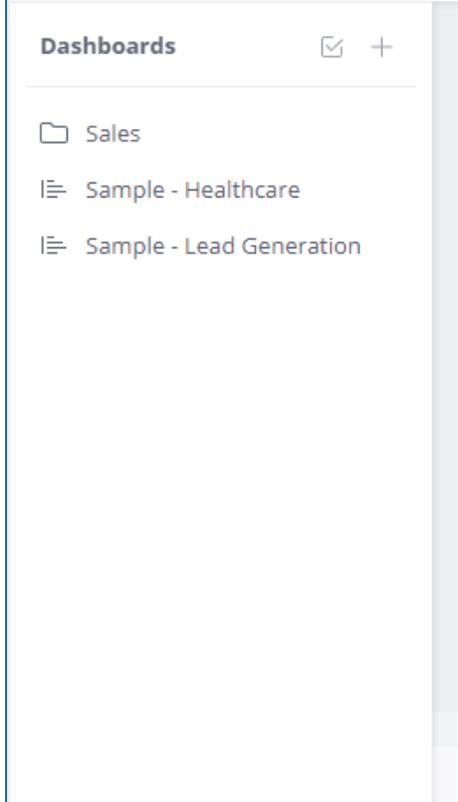
Click to set as parameter value

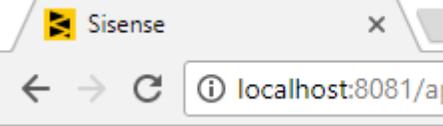
Replace the sample values with your required values. For an explanation of each parameter and how it modifies the Sisense Web Application, see [White Labeling Configuration Options](#).

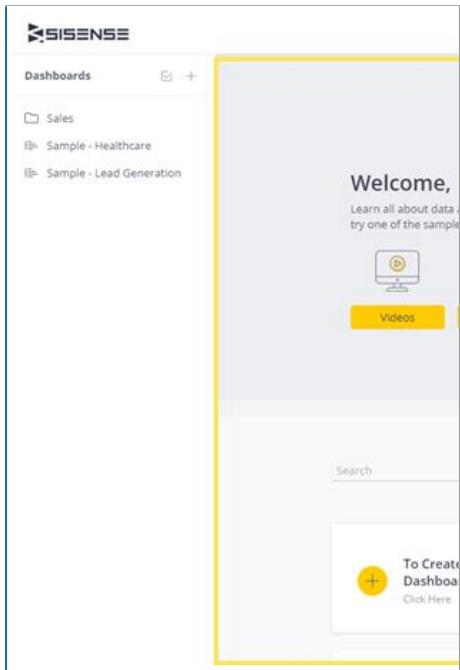
5. Click **RUN**. Refresh the Sisense Web Applicationpage to view your changes.  
**Note:** You can change just some of the parameters by leaving out those you do not want to update.

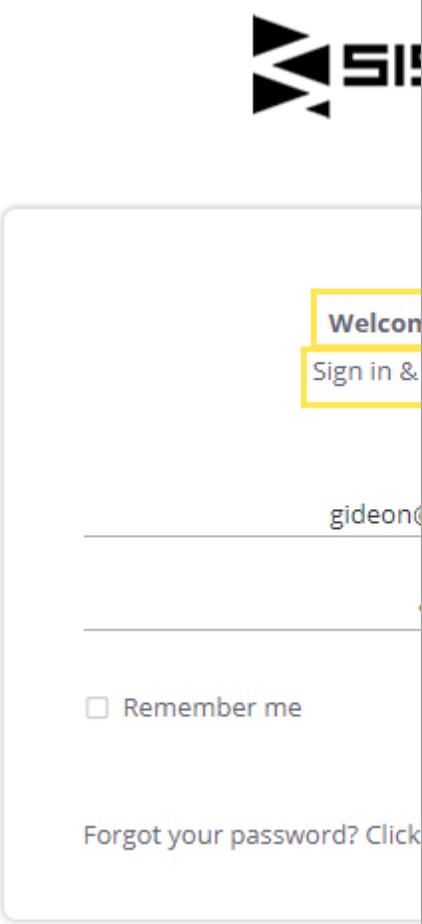
# White Labeling Configuration Options

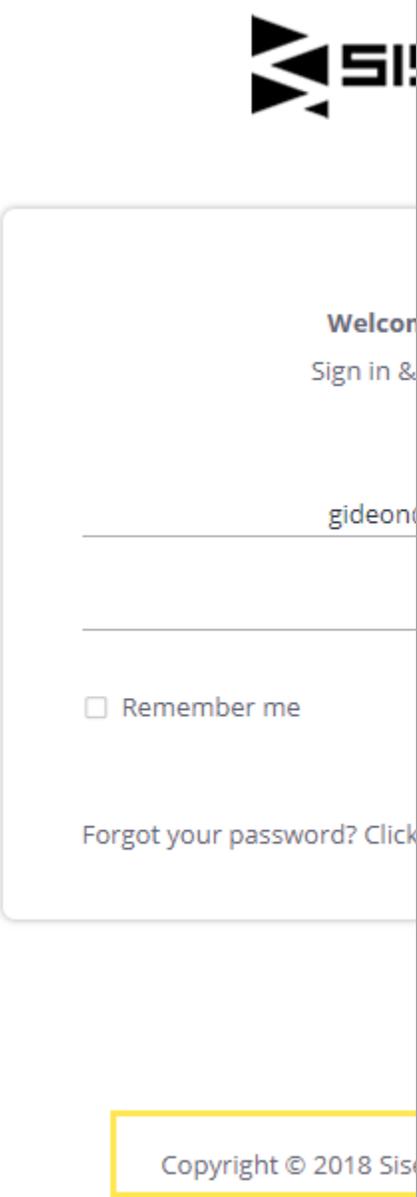
The following table provides a list of the features you can rebrand within the Sisense Web Application and the JAQL code that must be modified to rebrand it. Click the thumbnail image of the feature to see the default branding provided by Sisense and where your branding will appear.

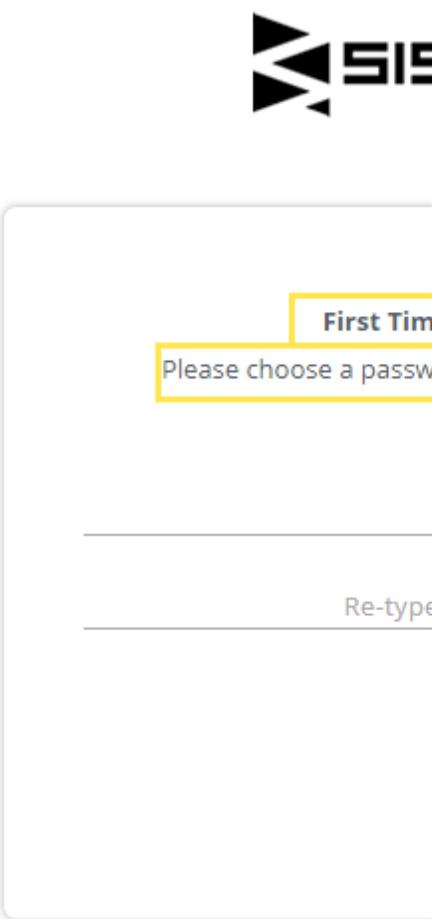
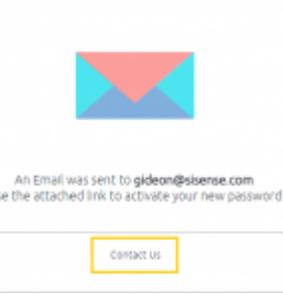
Name	Description												
<p>Sisense Logo</p>  	<p>The Sisense logo can be replaced with your own. You need to supply two sizes of the logo, for each device (Desktop / Tablet / Mobile). These logos affect the following screens:</p> <ul style="list-style-type: none"> <li>• Login Screen</li> <li>• Account Activation Screen</li> <li>• 'Forgot Password' Screen</li> </ul> <p>Below is a table with the ideal logo dimensions (pixels):</p> <table border="1" data-bbox="639 1010 1589 1304"> <thead> <tr> <th></th> <th>Small</th> <th>Large</th> </tr> </thead> <tbody> <tr> <td>Desktop</td> <td>92×26</td> <td>204</td> </tr> <tr> <td>Tablet</td> <td>92×24</td> <td>160</td> </tr> <tr> <td>Phone</td> <td>92×24</td> <td>192</td> </tr> </tbody> </table> <p><b>Important Notes:</b></p> <ol style="list-style-type: none"> <li>1. Images must be in URI format (web address)</li> <li>2. No spaces are allowed in the address of the images</li> <li>3. Transparent .png format is preferred</li> </ol> <pre data-bbox="639 1581 1589 1862"> "logo":{ "desktop":{ "small":"/resources/&lt;YourBrand&gt;/logoDesktopSmall.png", "large":"/resources/&lt;YourBrand&gt;/logoDesktopLarge.png" }, </pre>		Small	Large	Desktop	92×26	204	Tablet	92×24	160	Phone	92×24	192
	Small	Large											
Desktop	92×26	204											
Tablet	92×24	160											
Phone	92×24	192											

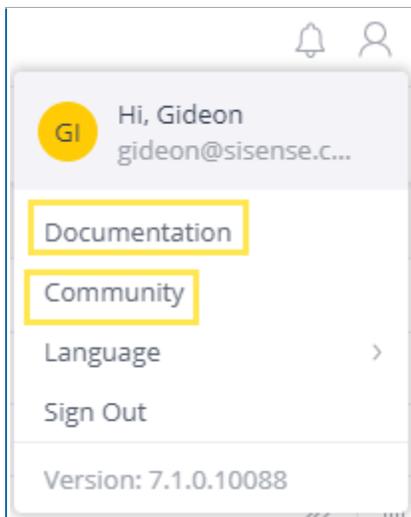
	<pre>"phone":{ "small":"/resources/&lt;YourBrand&gt;/logoPhoneSmall.png", "large":"/resources/&lt;YourBrand&gt;/logoPhoneLarge.png" }, "tablet":{ "small":"/resources/branding/logoTabletSmall.png", "large":"/resources/branding/logoTabletLarge.png" } },</pre>
<p>Page Title and Favicon</p> 	<p>The page title that appears in the browser tab/page and the Sisense favicon.</p> <p><b>Note:</b></p> <ol style="list-style-type: none"> <li>1. Image must be in URI format (web address)</li> <li>2. .ico format is preferred</li> </ol> <p>You can modify these images by updating the values of the following JAQL parameters in the REST API:</p> <pre>"favicon" : "/resources/&lt;YourBrand&gt;/fav.ico" "pageTitle": "Sisense",</pre>
<p>Homepage</p>	<p>The home screen when your users log into Sisense.</p> <p><b>Note:</b> The prefix http:// or https:// must be included.</p> <pre>"homePage" : "http://www.mysite.com",</pre>

	
<p>Homepage Config</p> 	<p>The Welcome banner of the homepage displayed when users log into Sisense.</p> <p>You define the following properties in the homePageConfig object:</p> <p><b>hideSearchInNavver:</b> Indicate if you want to display the Search bar below or hide it. By default, the Search bar is displayed.</p> <p><b>marketingPart:</b> The Videos button. Enter a URL that opens your Video or marketing site.</p> <p><b>tutorialsLink:</b> The Tutorials button. Enter a URL that opens your site.</p> <p><b>documentationLink:</b> The Documentation button. Enter a URL that opens your documentation site.</p>
<p>Login Page</p>	<p>The Login page title and subtitles.</p> <p>“loginTitle” : “WELCOME TO SISENSE”, “loginSubtitle” : “Sign in &amp; start exploring”</p>

	
<p>Copyright Text</p>	<p>The copyright text that appears in Sisense.  “copyrightText”:”Copyright © 2018 Sisense inc. All rights reserved”,</p>

	
<p>Activation Titles</p>	<p>The titles on the activation page for new users signing in to Sisense.</p> <p>“activationTitle”: “First Time Activation”, “activationSubtitle”: “Please choose a password to begin using Sisense”,</p>

	
<p><b>Contact Us</b></p> 	<p>The Contact Us text is displayed when you request a password and enter your email on the Recover Password page. This text displays a message and provides a link to your Support services.</p> <p>“contactUsText” : “For further assistance, please contact our success management team!: &lt;a href='mailto:support@sisense.com'&gt;support@sisense.com&lt;/a&gt;”,</p>
<p><b>User Support</b></p>	<p>The User Support text is displayed in the top right corner of the Sisense Web Application.</p> <p>Be default, the links redirect customers to the Sisense user forums and documentation site. You can modify these URLs</p>



to redirect users to your user forums or documentation. This is useful for example when providing OEM services or if you have special forums or instructions to display to your users.

“documentationUrl”: “string”,  
 “forumUrl”: “string”

Replace the string with the URL to the relevant locations to redirect your users to your support sites.

## Emails

The sender information and subject line text displayed in automated emails.

For further customization, see [Rebranding Sisense Automated Emails](#).

Note: Branded emails (for example, using the senderEmail and senderName parameters) will be used for new user/group invitation emails and password recovery emails. When sharing a dashboard report, or for subscription emails, the email will be sent without the original name and email of the dashboard’s creator.

```

“emails” : {“senderEmail”:“sys-admin@company.com”,“senderName”:“sys-admin”,“templates_directory” :
“/resources/<YourBrand>/emails”,“passwordRecoverySubject” :
  “Sisense password recovery”,
“newUserInviteSubject” :
  “Sisense account activation”,
“shareWithNewUserSubject” :

```

```
"A Dashboard has been shared with you:",  
"shareWithExistingUserSubject" : "A  
Dashboard has been shared with you:"  
"transferOwnership": "string" }
```

# Embedded Analytics

Sisense's embedded analytics solution is an end-to-end, single stack BI solution embedded within your application that lets your customers easily prepare, analyze and visualize complex data. Sisense's embedded analytic solution offers built-in features and exposes a variety of functionality through several interfaces, such the Sisense JavaScript and REST APIs for customizing your integration.

You can accomplish the following through the Sisense embedded analytics solution:

- Seamlessly integrate Sisense with your current interface and branding
- Manage data, reporting and visualization
- Integrate SSO and Active Directory
- Govern user and data security
- Support high availability

This documentation is intended for two audiences, product managers and developers. For product managers, those who want to enhance their products and applications with Sisense embedded analytics, this section describes the features and functionality Sisense offers as part of an embedded analytics solution. For developers, those who embed Sisense in the application, this section describes how to embed Sisense analytics through the relevant APIs.

## Overview

When you embed Sisense analytics, you can provide your users with Sisense's analytical capabilities such as data management, reporting and visualization, built into your business applications and solutions.

Each OEM has their own requirements when embedding Sisense's analytic solution. This section covers the multiple ways in which Sisense's OEM features

can be embedded to provide BI capabilities in your applications. Sisense exposes functionality through several interfaces. For example, Sisense’s REST API allows you to white label Sisense and Sisense’s command shell allows you to couple Sisense with your on-premise installation to modify connection settings. You determine which features you want to embed according to the type of OEM services you provide to your customers.

The following table lists the various OEM-related tasks you can perform with Sisense and the description with links to the relevant documentation for completing each tasks:

Task	Description
Rebranding Sisense	<p>Sisense enables you to white label the Sisense Web Application to display your company's logos and colors and mask the URL to display your company's domain.</p> <p>For documentation regarding how you can customize and rebrand your Sisense solution, see the following topics:</p> <p><a href="#">White Labeling Sisense</a></p> <p><a href="#">Rebranding Sisense Automated Emails</a></p> <p><a href="#">Displaying Custom URLs</a></p>
Embedding Sisense	<p>Sisense also allows you to embed widgets and dashboards directly into your site or application through the use of iFrames or SisenseJS, Sisense's custom JavaScript library for embedding widgets.</p> <p>For documentation regarding how you can embed Sisense, see the following topics:</p> <p><a href="#">Embedding Dashboards and Widgets via iFrames</a></p> <p><a href="#">Embedding Dashboards and Widgets via JavaScript</a></p>

# Considerations

Before embedding Sisense analytics, you should keep in mind several considerations that may affect your design and deployment:

- **Users:** Who will be using your application? This determines what interfaces need to be white labeled and which API calls you should leverage.
- **Performance:** As with any deployment, you must provide sufficient CPU and memory on the client and the server for your anticipated user load.
- **Availability:** What is your desired up-time and how many resources can you dedicate to maintaining system availability.

# Querying ElastiCubes during Builds

Sisense ElastiCube Sets are collections of ElastiCubes with identical schemas that allow you to query running ElastiCubes within the ElastiCube Set while other ElastiCubes are building.

**Note:** From Sisense V6.5 onwards, you can query ElastiCubes during accumulative builds without ElastiCube Sets. However, to ensure that your users can query ElastiCubes after a failed build, you can implement ElastiCube Sets to handle your queries while Sisense rebuilds the failed ElastiCube.

Grouping ElastiCubes into ElastiCube sets has several benefits:

- Allowing viewers to query the most up-to-date cubes within an ElastiCube Set.
- Reducing server load during builds by using multiple nodes in which only the non-building node is queried.
- Achieving data redundancy by running builds interchangeably across multiple nodes.

To maximize availability and reduce build and dashboard response times, administrators can distribute resource allocation across multiple Sisense nodes. By adding ElastiCubes across multiple servers to a single ElastiCube Set, administrators can configure ElastiCube Sets to run builds interchangeably via the Sisense Orchestrator Service so the most up-to-date and complete ElastiCubes are queried.

The Sisense Orchestrator Service is an automated service that can be configured to synchronize builds across the ElastiCube Set. If you change your schemas within any of the ElastiCubes, you must manually change the schemas in all of the ElastiCubes and rebuild the cube entirely.

This page provides the following information:

1. Instructions on [Creating](#) and [Deleting ElastiCube Sets](#).
2. A walkthrough for how you can create a set and configure the Sisense Orchestrator Service. See [Working with ElastiCube Sets](#).
3. (Optional) Instructions on how to implement and configure the Sisense Orchestrator Service. See [Activating the Sisense Orchestrator Service](#).
4. How to update an ElastiCube within an ElastiCube Set.

## Creating ElastiCube Sets

From the Admin page, Administrators can create ElastiCube Sets.

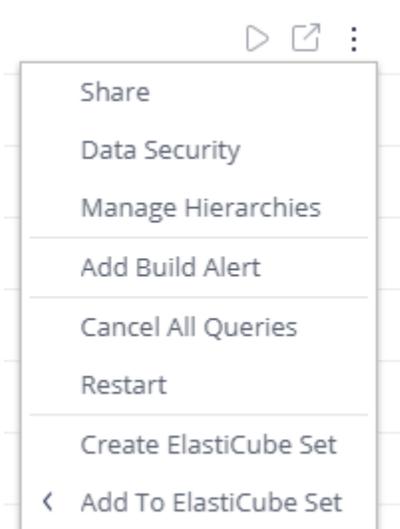
ElastiCube Sets require at least two ElastiCubes. The Administrator can add multiple ElastiCubes to a set, however, ElastiCubes can only be part of a single ElastiCube Set at a time. After an Administrator creates an ElastiCube Set or adds an ElastiCube to an existing set, the ElastiCube is disabled in the original server and no further actions can be performed outside of the set.

Sisense recommends that you develop ElastiCubes in a Dev environment or separate ElastiCube to ensure they work before adding them to an ElastiCube Set.

### To create an ElastiCube Set:

1. Click **Admin** and select the **Data Sources** tab on the left.

2. For the relevant ElastiCube you want to add to the ElastiCube Set, select : and click **Create ElastiCube Set**.



3. In the Create ElastiCube Set window, enter a name for the set and select the ElastiCubes you want to include in the set.
4. From the Routing Strategy list, select a strategy for how you want to route traffic to your ElastiCube Set. See [Routing Queries in ElastiCube Sets](#) for more information.
5. From the Failover ElastiCube list, select which ElastiCube you prefer to use if the other ElastiCubes in the set are not available. Failover ElastiCubes only handle queries when an ElastiCube that is part of the ElastiCube Set

fails.(Note: this feature is available in Sisense V6.7 onwards).

### Create ElastiCube Set

ElastiCube Set Name

Type in a descriptive name for the new ElastiCube Set.

Routing Strategy

Single ElastiCube

ElastiCubes

Search ElastiCubes...

<input type="checkbox"/>	Facebook Ads Demo	4.01MB
<input type="checkbox"/>	new	39.46MB
<input type="checkbox"/>	Sample ECommerce	59.09MB
<input type="checkbox"/>	Sample Healthcare	1.56MB
<input checked="" type="checkbox"/>	Sample Lead Generation	2.25MB
<input type="checkbox"/>	Training	814.28KB

1 ElastiCube Selected

Failover ElastiCube

None

Cancel

6. Click **Create**.
7. In the Existing Dashboards popup message, indicate if you want to associate your ElastiCube's dashboards with the ElastiCube Set or maintain the association with the original ElastiCube.

Existing Dashboards in added ElastiCubes

- Connect existing dashboards to new ElastiCube Set
- Leave existing dashboards with their current ElastiCube

OK

Cancel

8. Click **OK**. The ElastiCube Set is created and displayed in the Admin page. After you have created an ElastiCube Set, you can configure the ElastiCube Set's Hierarchies, Access Rights, and Data Security.

## Adding ElastiCubes to ElastiCube Sets

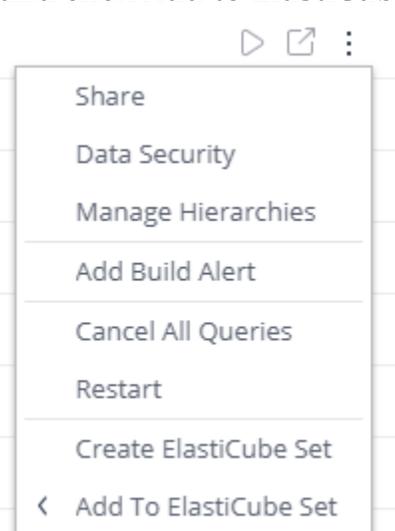
After you have created an ElastiCube Set, you can add more ElastiCubes to the set.

When you add an ElastiCube to an existing ElastiCube Set, any dashboards associated with that ElastiCube are passed on to the ElastiCube Set.

**Note:** You cannot add 'localhost' to an ElastiCube Set. Instead, you can add your localhost as a new server with its IP address as the server name and then add ElastiCubes from it to an ElastiCube Set.

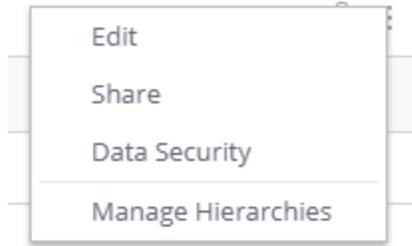
### To add an ElastiCube to the ElastiCube Set:

1. Click **Admin** in the upper right corner and select the **ElastiCubeS** tab on the left.
2. For the relevant ElastiCube you want to add to the ElastiCube Set select : and click Add to ElastiCube Set.



OR

In the ElastiCube Set table, click : and select **Edit** and select the ElastiCubes to be added to the set. Click Save after you have selected the relevant ElastiCubes.



The ElastiCube is added to the ElastiCube Set and displayed in the ElastiCube Set table on the Admin page.

## Deleting ElastiCube Sets

Administrators can delete ElastiCube Sets. Deleting an ElastiCube Set removes the set from the Admin page and any dashboards associated with that set.

### To delete an ElastiCube Set:

1. Hover over its title and click the Trash  icon.
2. Click **Delete** to confirm that you want to delete the ElastiCube Set.

## Working with ElastiCube Sets

This section describes how to create your first ElastiCube Set and configure it for high availability.

In this example, let's assume that you have an ElastiCube, ElastiCube A, and you want to add it to a set to ensure that the data can always be queried, even during a build.

The first step is to create a copy of ElastiCube A. The ElastiCubes can have different names, but the schemas must be identical.

**Note:** If your ElastiCubes are located on a single node, the ElastiCubes cannot have the same name.

In the ElastiCube Manager, we will save a new version of ElastiCube A as ElastiCube B. In Sisense, you can see both ElastiCubes are identical (Select **Admin > ELASTICUBES**)

Local Machine

ElasticCube	Status	Created On	Last Built	Size On Disk	
Elastic Cube A	Running	3/3/16 9:11 AM	3/3/16 9:11 AM	39.42MB	Open  
Elastic Cube B	Running	3/3/16 9:11 AM	3/3/16 9:14 AM	39.42MB	Open  

Now, let's create an ElasticCube Set that contains both ElasticCubes. Click **>** **Create ElasticCube Set**. For more information, see [Creating ElasticCube Sets](#). Select the ElasticCubes to be added to the set. For this set, we will select ElasticCube A and B, which have different names, but identical schemas.

## Create ElasticCube Set

ElasticCube Set Name

ElasticCube Set Example

ElasticCubes

Search ElasticCubes... 

<input checked="" type="checkbox"/>	Elastic Cube A	39.42MB	
<input checked="" type="checkbox"/>	Elastic Cube B	39.42MB	
			
			
			
			
			

2 ElasticCubes Selected

CREATE

CANCEL

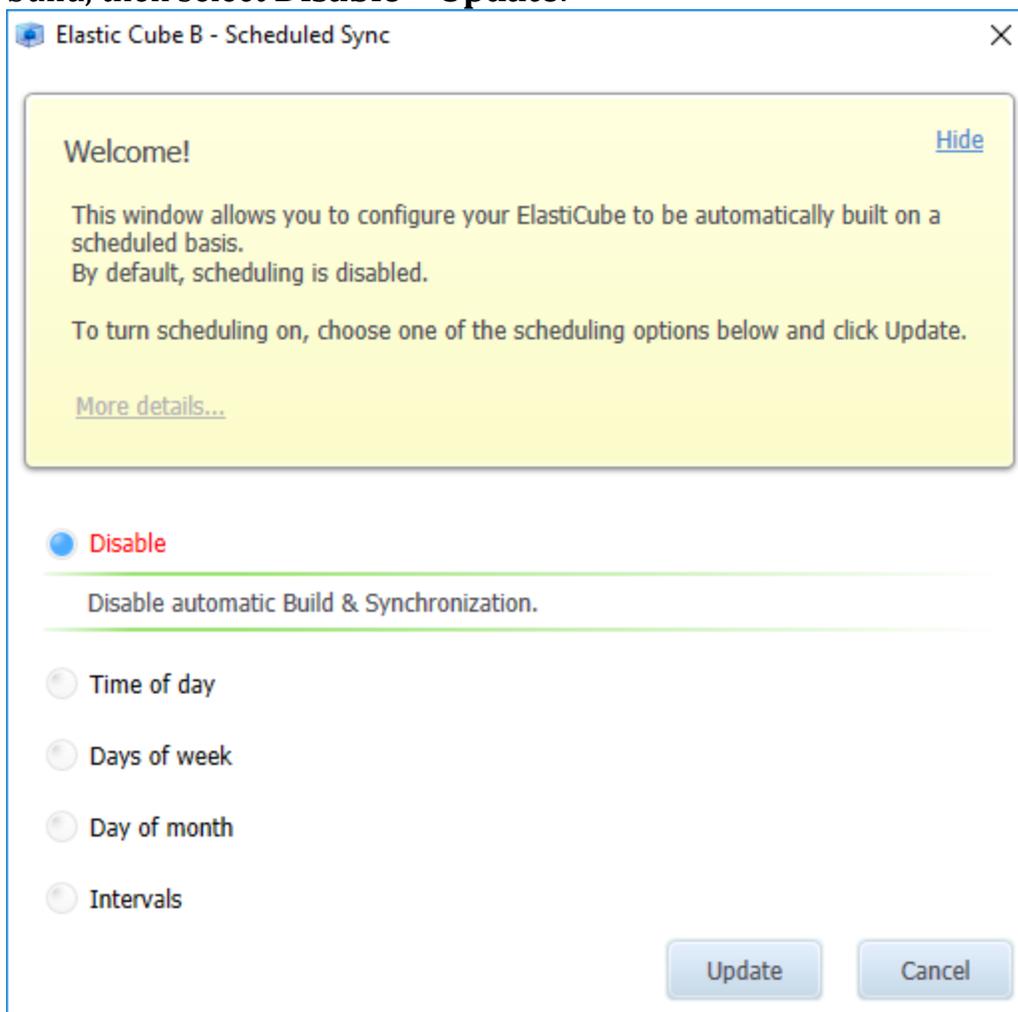
Click **Create** and confirm that you want to create a new set with these ElasticCubes. The set and its ElasticCubes are now displayed at the top of the Admin page.

 ElasticCube Example [Manage Hierarchies](#) | [Access Rights](#) | [Data Security](#) |  

Server	ElasticCube	Status	Created On	Last Built	Size On Disk	
LocalHost	Elastic Cube A	Running	3/3/16 9:11 AM	3/3/16 9:11 AM	39.42MB	Open 
LocalHost	Elastic Cube B	Running	3/3/16 9:11 AM	3/3/16 9:14 AM	39.42MB	Open 

Now that the set is built, it's time to schedule automatic builds with the Sisense Orchestrator Service. This service automatically builds the ElastiCubes according to your preferences. For example, you can schedule the builds to run every 10 minutes or at midnight when queries are at their fewest.

To schedule automatic builds, first you should disable any scheduled builds you have for the ElastiCubes. In the ElastiCube Manager in the top menu, select **ElastiCube > Schedule Build Settings** for each ElastiCube with a scheduled build, then select **Disable > Update**.



Now that all the ElastiCubes' scheduled builds are disabled, it's time to configure the Sisense Orchestrator Service to build the ElastiCubes within the set. To configure the Sisense Orchestrator Service, you must modify the Config.JSON file

that was installed when you installed the latest version of Sisense. This file defines how the ElastiCube Set is being built, the order of the build, and how frequently a build occurs.

This file is located in the Sisense.Orchestration Config folder (/Sisense.Orchestration/Config/) of your Sisense Installation folder.

For this set, let's assume that you want to schedule an accumulative build for midnight. Within the **elastiCubes** object, we listed both ElastiCubes to be built by the Sisense Orchestrator Service. The order of the ElastiCube builds is in the order of oldest to newest. Next, in the **timer** object, we listed the hour when the build is to take place. The value is defined in Cron format, so midnight equals 0 minutes 0 hours "0 0 \* \* \* \*".

Finally, we define the flow of the build. We want an accumulative build, but if for some reason this fails, we want to run an entire build as a back-up. Remember, when an accumulative build fails, the ElastiCube needs to be completely rebuilt to ensure that the data is intact. In the build object, we define the type of builds to be initiated and their order. The order of the values determines what build type is initiated first. In this example, the value is "FullUpdateExisting", "Full", which means run an accumulative build first and if that fails, run a full build. For a complete explanation on how to define the rest of the objects list in the Config.JSON file, see [Activating the Sisense Orchestrator Service](#).

Now, save the file and Sisense automatically begins to run the builds according to the schedule defined in the timer object.

```
{
  "servers": {
    "localhost": {
      "ip": "localhost",
      "port": "8081",

      "apiKey": "eyJ0eXAiHDHDHsXKSSZJdnUzI1NiJ9.eyJ1c2VyIjoiaNTY1N332RW2eygL0uVLufGYC-8vY"
```

```

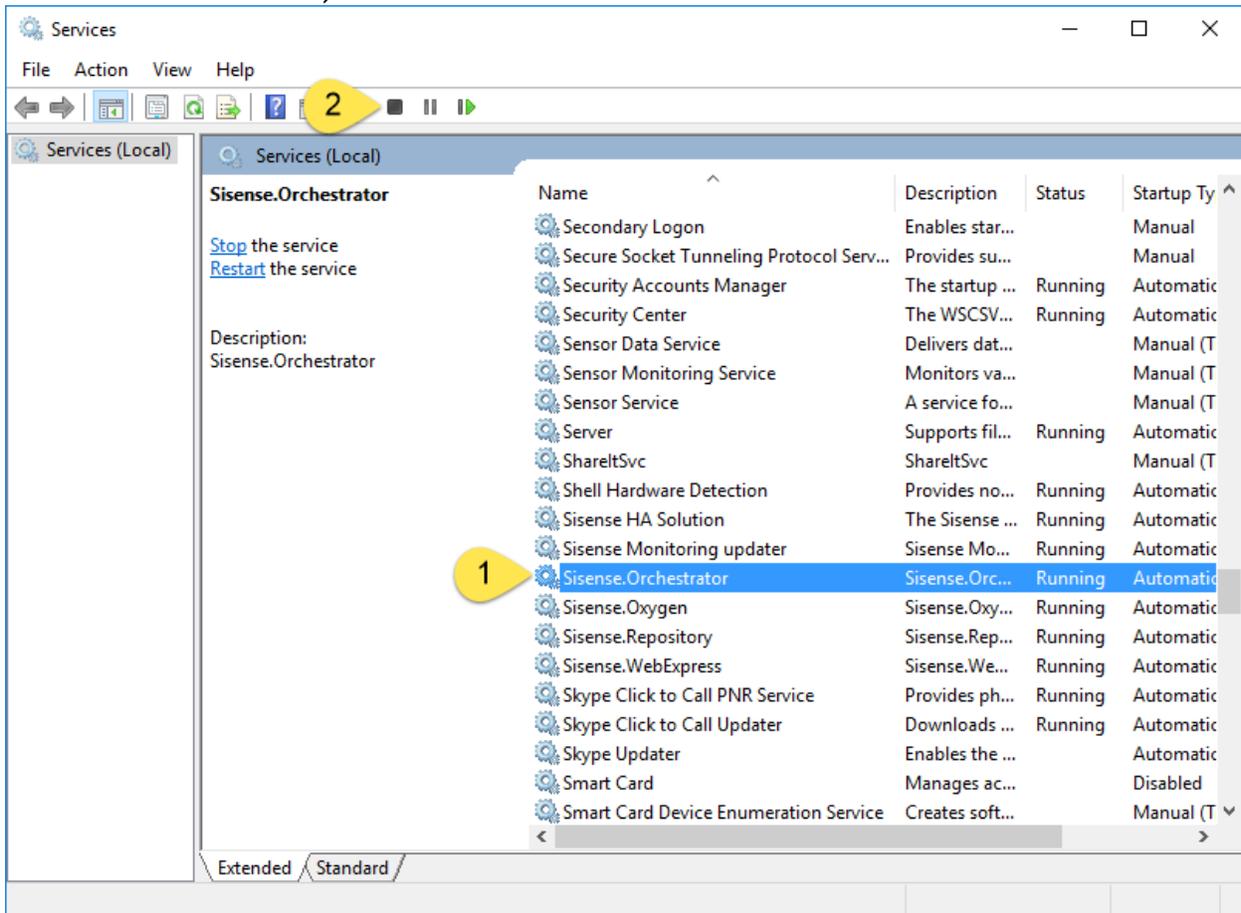
    }
  },
  "tasks": [
    {
      "name": "taskName",
      "elasticubes": [
        {
          "server": "localhost",
          "cube": "ElastiCube A"
        },
        {
          "server": "localhost",
          "cube": "ElastiCube B"
        }
      ],
      "timer": "0 0 * * *",
      "build": {
        "queue": [
          "FullUpdateExisting",
          "Full"
        ],
        "flow": "sync"
      },
      "email": {
        "when": "Fail"
      }
    }
  ],
  "emails": [
    {
      "email": "JohnDoe@sisense.com",
      "name": "John Doe"
    }
  ],
  "logLevel": "info"
}

```

The Sisense Orchestrator Service automatically builds the ElastiCubes within the set, but if the schema changes, you must manually rebuild each ElastiCube within

the set so the schemas match across all the ElastiCubes. While working with the schemas of an ElastiCube that is part of an ElastiCube Set, you should stop the Sisense Orchestrator Service.

To stop the Sisense Orchestrator Service, open Windows Services, select Sisense.Orchestrator, and click ■.



After the ElastiCubes are rebuilt, restart the Sisense Orchestrator Service by selecting Sisense.Orchestrator and click ▶.



The ElastiCube Set is now configured and running ElastiCubes can be queried while the other ElastiCubes are being built.

# Activating the Sisense Orchestrator Service

The Sisense Orchestrator Service automates the build process for ElastiCube Sets and determines how frequently the build process occurs.

To activate the Sisense Orchestrator Service, you must create and save a JSON file called `config.json` in the `Sisense.Orchestration Config` folder (`/Sisense.Orchestration/Config/`) of your Sisense Installation folder. This file defines how the ElastiCube Set is being built, the order of the build, and how frequently a build occurs.

After you save the file, the Sisense Orchestrator Service automatically begins to build ElastiCubes in your ElastiCube Set as defined by you in the `config.json` file. You can review the progress of the build in the `/Sisense.Orchestration/Logs` folder in a log file whose file name you define in the JSON file.

Note: During this phase it is recommended to stop the Sisense Orchestrator Service while importing data and reactivate it after the importing process is completed.

The objects that you must include in the `config.json` file and their descriptions are provided in the table below:

Key	Value
Servers	The servers block contains a server object that defines each server that hosts an ElastiCube within your ElastiCube Set. Each server must include a server name, IP, port, and a unique API key. For example, if you have two different servers, each server must be represented as a separate object with its own name, IP, and port.
ServerName	Your server's name. The value of this key is used when associating the ElastiCube Set with a server in the value of the Tasks key.
IP	The IP address of your server.

Key	Value
Port	The port of your web server.
API	The access token used to identify your requests to the Sisense web server. For more information about the access token, see Retrieving a Token.
Tasks	The Tasks object contains an array of ElastiCube objects. This object can contain multiple ElastiCube objects, but a minimum of two are required for a set.
Name	The name of the plan. The value of this parameter is the filename of the log file for the Sisense Orchestrator Service.
ElastiCubes	The ElastiCube object that holds an array of ElastiCubes. Each ElastiCube must include the name of its server and the name of the cube as defined in the Elastic Manager.
Server	The name of the server that hosts the ElastiCube.
Cube	The name of the cube to be included in the set.
Timer	The frequency of builds in Cron format. Some examples: -To run a build each night at midnight, enter the value "0 0 * * *"-To run a build every hour, enter the value "0 * * * *" See the full Cron format reference <a href="#">here</a> .
Build	The Build object contains two keys, Queue and Flow, which determine the order in which the build is completed and the type of build.
Queue	The value of this key is the type of build Sisense should attempt and in what order. There are two possible values delimited by a comma: FullUpdateExisting: Attempts an accumulative build. Full: Attempts an entire build. Sisense recommends the following

Key	Value
	value:"FullUpdateExisting","Full" In this example, Sisense attempts to do an accumulative build first, and if it fails, Sisense attempts an entire build. If the build fails for all builds, Sisense retries again according to the next scheduled build.
Flow	Determines the order of the build. A build occurs according to the value of Timer. For example, if your builds are scheduled for midnight, then the Orchestrator begins to build the ElastiCubes at midnight according to the order you define below. Single: Builds a single cube, then waits for the next time a build is scheduled according to the value of Timer. The cube to be built is a cube that previously failed or the cube with the oldest build. Sync: Builds the cubes with the oldest build first and then the remaining cubes in the order of oldest to newest until all ElastiCubes are updated. All: Builds the cube with the oldest build first and then the remaining cubes asynchronously.
Email	This object determines what triggers activate Sisense to send an email to users you define in the Emails block
When	Indicates when an email alert is triggered. Email alerts can be triggered for the following scenarios: None: No email is ever sent. Build: An email is sent when a build is complete. Fail: An email is sent when a build fails.
Emails	This object defines who receives an email and their email address when an email is triggered according to the events you define in the Email object.
Email	The email address where emails are to be sent depending on the value of the When key.

Key	Value
Name	The name of the recipient of the email.
LogLevel	The type of info returned in the log. The possible values you can enter are Info and Error.

### Example:

The following is an example of a config.json file. In this example, the ElastiCube Set contains two ElastiCubes, ElastiCube1 and ElastiCube2, hosted on the two different servers, Server1 and Server2. The ElastiCubes, even though they are located on separate servers, are connected through the Tasks object. The Sisense Orchestrator Service attempts to build the ElastiCube Set every 10 minutes. If the build fails, Sisense sends an email to JohnDoe@Sisense.com.

```
{
  "servers":{
    "Server1":{
      "ip":"localhost",
      "port":"80",

      "apiKey":"eyJ0eXAiHDHDHsXKSSZJdnUzI1NiJ9.eyJ1c2VyIjojoiNT
Y1N332RW2eygL0uVLufGYC-8vY"
    }
    "Server2":{
      "ip":"192.168.5.134",
      "port":"8081",

      "apiKey":"dgfhdgSSdSSZJdnUzI1NiJ9.eyJ1c2SDFSDfSDFSDfSDFSDfSDF
GSSDFVfasddgdeygL0uVLufDFeY"
    }
  },
  "tasks":[
    {
      "name":"taskName",
      "elasticubes":[
        {
          "server":"Server1",
          "cube":"ElastiCube1"
        }
      ]
    }
  ]
}
```

```

        },
        {
            "server": "Server2",
            "cube": "ElastiCube2"
        }
    ],
    "timer": "* / 10 * * * *",
    "build": {
        "queue": [
            "FullUpdateExisting",
            "Full"
        ],
        "flow": "sync"
    },
    "email": {
        "when": "Fail"
    }
}
],
"emails": [
    {
        "email": "JohnDoe@sisense.com",
        "name": "John Doe"
    }
],
"logLevel": "info"
}

```

## Overcoming Build Failures in ElastiCube Sets

In Sisense, if an accumulative build fails, you must rebuild the ElastiCube that failed entirely to ensure the validity of the ElastiCube.

If you use the Sisense Orchestrator Service, Sisense recommends that you define the value of the **Queue** key as “FullUpdateExisting, Full”

```
"queue": ["FullUpdateExisting", "Full"]
```

In this configuration, the Sisense Orchestrator Service attempts to build the set using accumulative builds and in the event of a failure, a full build is initiated.

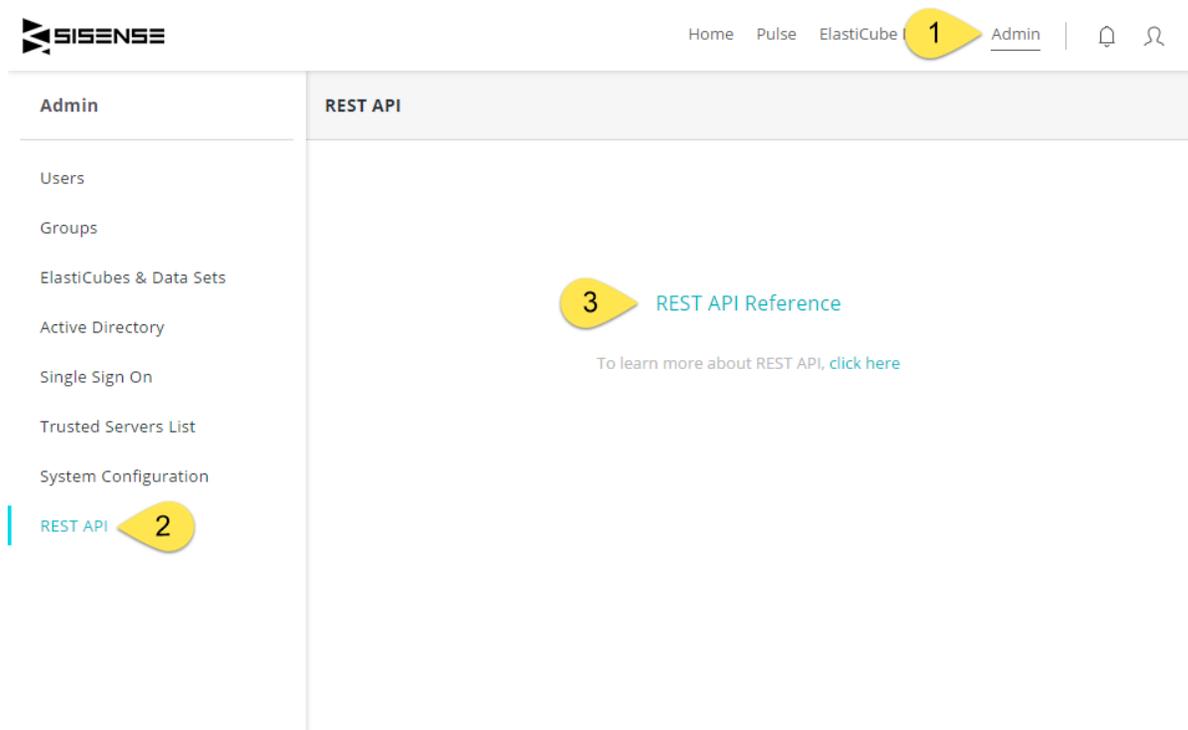
## Retrieving a Token

As part of the config.json file, you must provide a token that Sisense uses for authenticating your ElastiCube Sets.

You can retrieve this token through Sisense's API documentation. After you retrieve the token, save it as the value of apiKey in the config.json file.

### To retrieve a token:

1. In Sisense, access the API documentation, select **Admin**> **REST API** > **REST API Reference**.



2. In the REST API Reference site, select Version 1.0 in the top-right corner of the page.
3. Open the authentication method.

4. Enter your login credentials and click **Run**. The access\_token is displayed in the Response Body below.

authentication Show/Hide | List Operations | Expand Operations

**POST** /authentication/login Authenticate and receive token

**Implementation Notes**  
The login endpoint validates passed credentials and returns an API token for subsequent requests to the API.

**Response Class (Status 200)**  
Model | Model Schema

```
{
  "success": true,
  "message": "string",
  "access_token": "string",
  "profile": {}
}
```

Response Content Type:

**Parameters**

Parameter	Value	Description	Parameter Type	Data Type
x-device-id	<input type="text"/>	The device id.	header	string
username	<input type="text" value="(required)"/>	User's username or email	formData	string
password	<input type="text" value="(required)"/>	User's password in plain text	formData	string
localeId	<input type="text"/>	User's locale, as ISO language code (See <a href="#">table</a> )	formData	string

**Response Messages**

HTTP Status Code	Reason	Response Model	Headers
401	Unauthorized: Invalid credentials. Wrong username or password.		

**RUN**

5. Copy the token and paste it into the value of apiKey in the config.json file.

[Hide Response](#) RUN

**Curl**

```
curl -X POST --header "Content-Type: application/x-www-form-urlencoded" --header "Accept: application/json" -d "username=gideon%40"
```

**Request URL**

```
http://localhost:8081/api/v1/authentication/login
```

**Response Body**

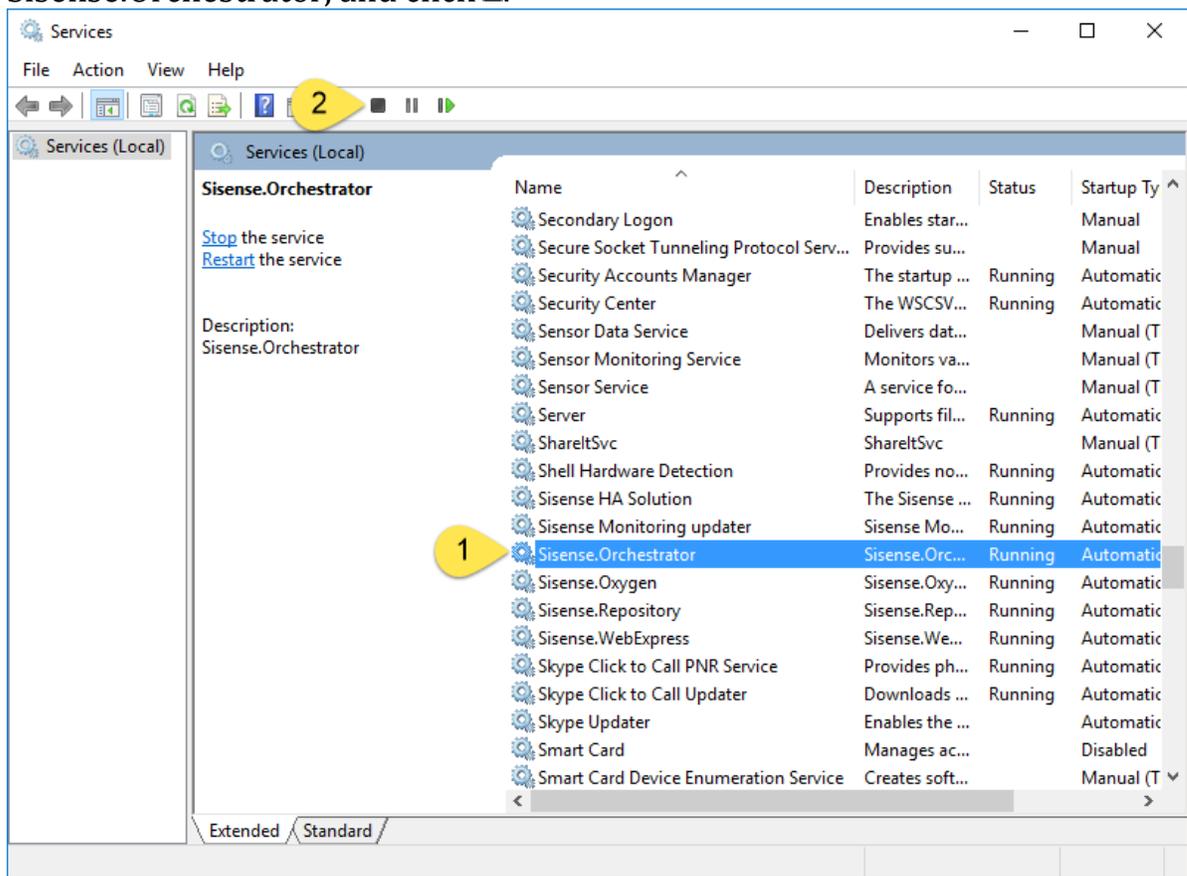
```
{
  "success": true,
  "message": "",
  "access_token": "eyJ0eXAiOiJKV1QiLCJhbGciOiJIUzI1NiJ9.eyJ1c2VyIjoiaWNTZi0wM3YTFiZGQ2YmVlNDY1MDNjZWNlIiwiaWF0IjoiNTk1MDU3N",
  "profile": {}
}
```

## Updating ElastiCubes within an ElastiCube Set

The Sisense Orchestrator Service automatically updates your data across ElastiCubes within an ElastiCube Set. If you want to modify the schema, for example, by adding new tables, you must manually rebuild the ElastiCubes and then import that ElastiCube to the other machines hosting your ElastiCube Set. While rebuilding your ElastiCube, you should deactivate the Sisense Orchestrator Service to prevent it from updating the ElastiCubes with different schemas. After the schemas are identical for all the ElastiCubes across the ElastiCube Set, you can reactivate the Sisense Orchestrator Service to automatically update your ElastiCubes.

### To update an ElastiCube within ElastiCube Set:

1. Stop the Sisense Orchestrator Service. Open Windows Services, select Sisense.Orchestrator, and click .



2. Apply your schema changes to the ElastiCube and run a full build.

3. Export the ElastiCube and import it on your other machines. For more information, see [Importing and Exporting ElastiCube Data](#).
4. In Windows Services, restart the Sisense Orchestrator Service by selecting Sisense.Orchestrator and clicking ▶.

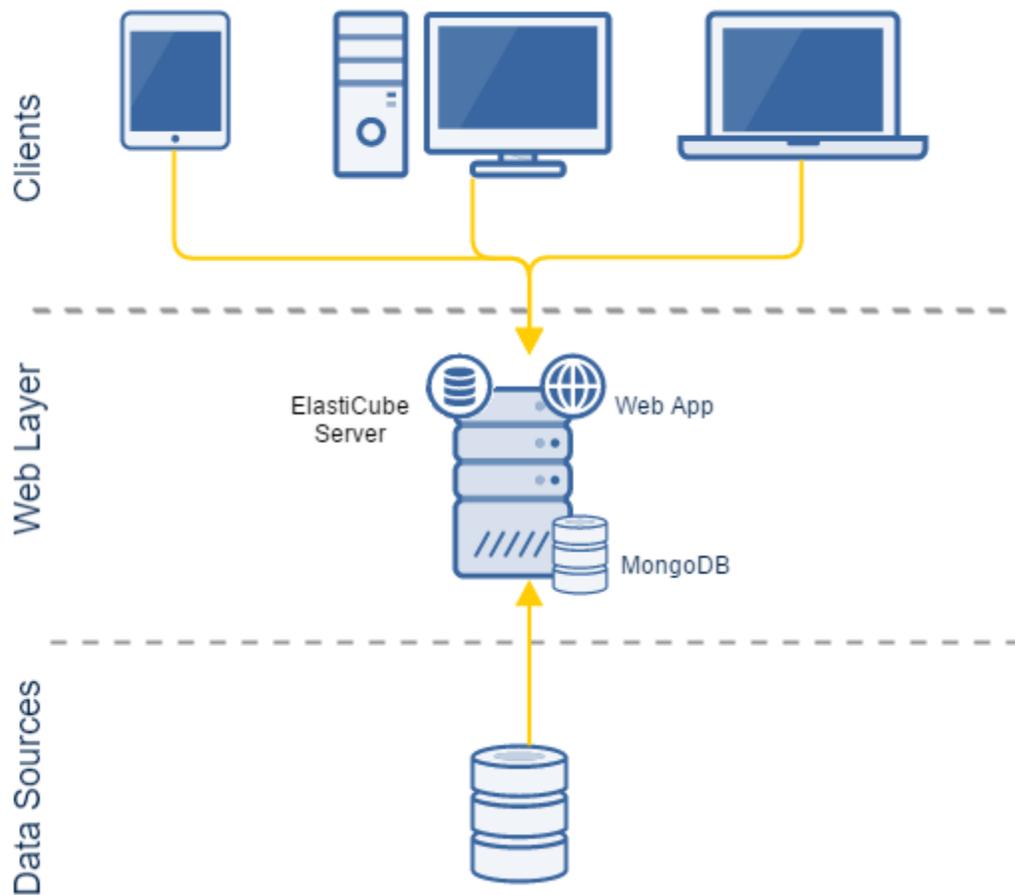
# High Availability in Sisense

## Overview

Sisense provides flexible design choices for supporting high availability (HA) and scalability for your Sisense configuration. To configure Sisense for high availability or optimized performance by scaling out Sisense servers, you must build in redundancy, thus reducing potential downtime or bottlenecks.

**Note:** High availability must be enabled in your license. For more information, please contact your Sisense representative or submit a request to Sisense Support through our community website.

In a simple configuration, you interface with the Sisense Web Application that interacts with the ElastiCube.



This configuration represents the Sisense full stack solution. At least one instance of each of the following components must be active to enable you update your data and allow your users to query that data from the dashboard:

### **Sisense Web Server**

The Sisense Web Server hosts the Sisense Web Application that provides the user interface and hosts the API endpoints. If the Sisense Web Server fails, your users cannot access the Sisense Web Application to view dashboards or use the Sisense APIs.

### **Application Database**

The application database is installed with Sisense and supports the Sisense Web Application. The application database contains dashboard, filters, and user information necessary for ensuring data consistency across all web servers. In

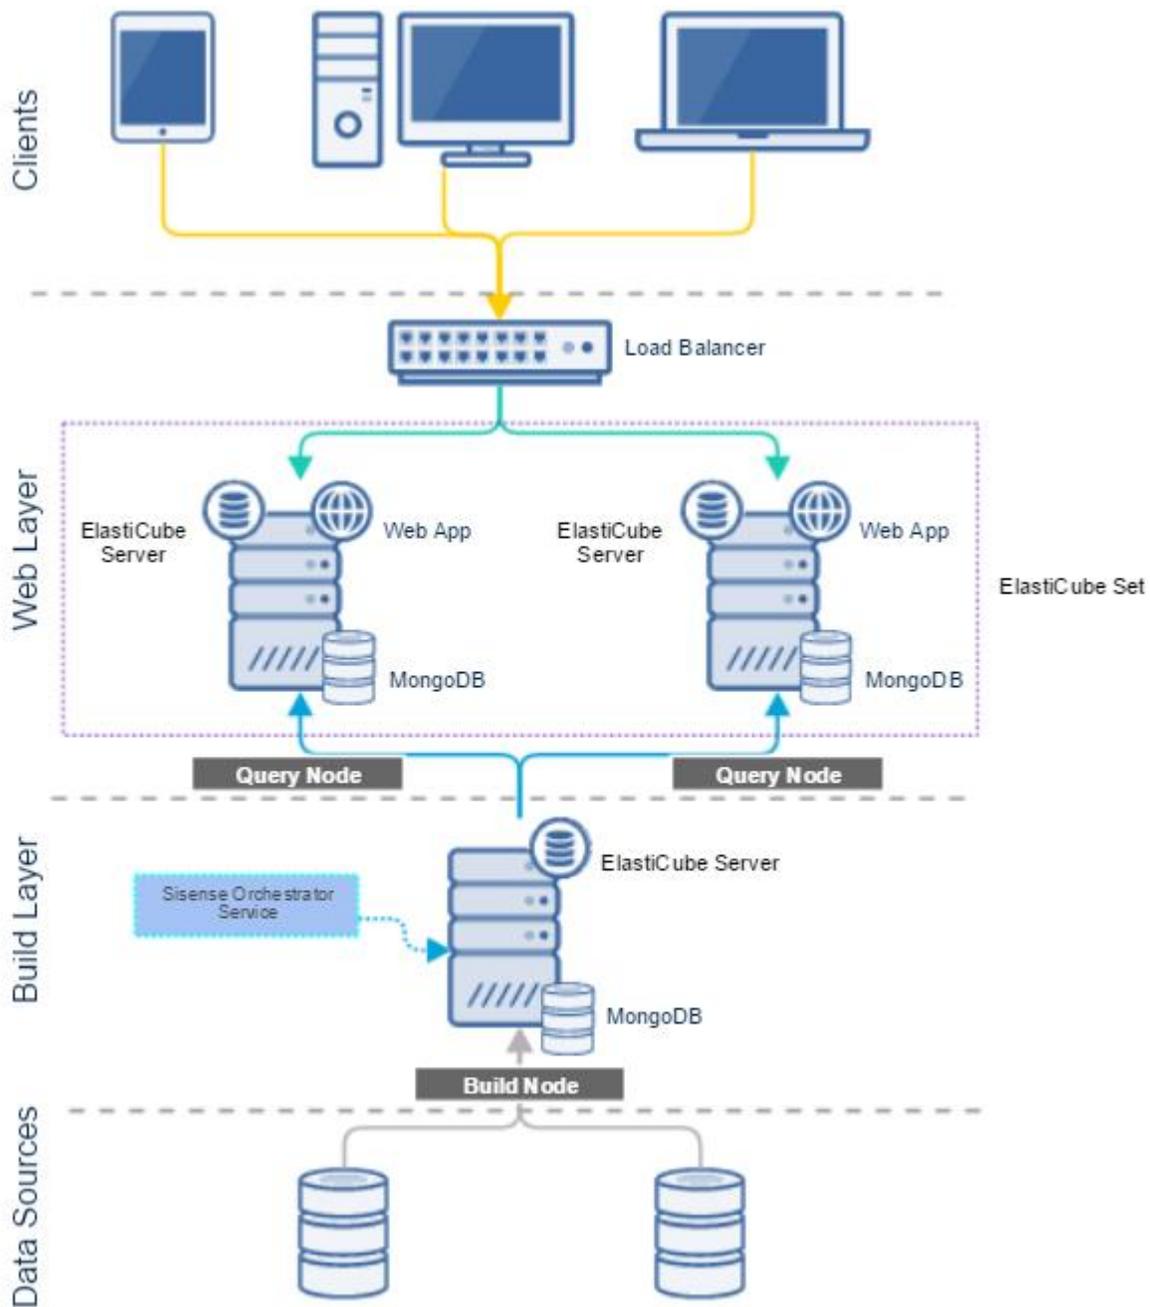
addition, the application database is used for authentication when you make calls to Sisense's APIs. If the application database fails the Sisense Web Application will fail.

### **ElastiCube Server**

The ElastiCube Server is installed locally on your computer and provides access to ElastiCubes. If access to the ElastiCube Server fails, queries from the Sisense Web Application will fail.

## **Understanding High Availability in Sisense**

While Sisense is fully-functional in an environment without high availability, a multi-node configuration is necessary for scalability to support large amounts of concurrent users and redundancy in case of a failure of one of the Sisense components.



In a multi-node configuration, Sisense components are replicated. Replicating each of these components provides redundancy and fault tolerance against the failure of any single component. The replicated components are combined into nodes. There are two types of nodes, a build node and query nodes. The query nodes, which handle user queries from the Sisense Web Application, are replicated to support high availability. The build node is typically not replicated

as its failure only prevents building new ElastiCubes not issuing queries from the Sisense Web Application.

### **Build Node**

The build node is responsible for building ElastiCubes and distributing the build to query nodes via the Sisense Orchestrator Service. The Sisense Orchestrator Service is an automated service that you configure on the build node to synchronize and distribute builds to the query nodes. For more information, see [Distributing ElastiCube Builds to Query Nodes](#).

### **Query Nodes**

Query nodes are responsible for supporting queries from Sisense users. These nodes contain a web server, MongoDB, and an ElastiCube server. ElastiCubes are distributed by the build node to the query node. The query nodes' ElastiCubes are combined into ElastiCube Sets to support high availability by separating the web and ElastiCube servers across multiple query nodes. If a build node is distributing a build to one ElastiCube server, Sisense automatically directs any queries to the other ElastiCubes in the ElastiCube Set.

The query node can be configured as a single application stack where each node hosts a Sisense Web Application, ElastiCube Manager, and a MongoDB. In this configuration, if the machine hosting the components fails, the whole query node will fail. Queries will then be redirected to the next available query node. Another option is to host each component of the query node separately in a distributed application stack. In this configuration, if a component of the query node fails, the rest of the query node is not affected. For an example of a single application stack, see [Scenario 1](#) and for an example of a distributed application stack, see [Scenario 2](#).

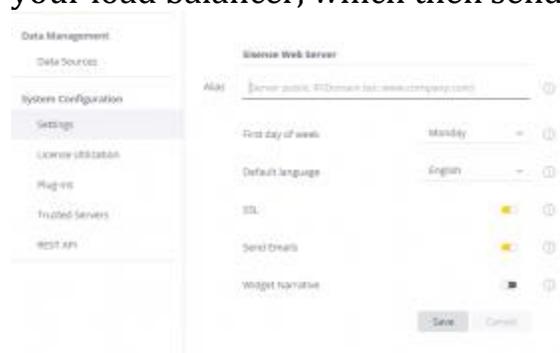
### **Load Balancer**

In addition to query and build nodes, to support a high availability configuration, you must handle load balancing on your side prior to directing traffic to one of

your Sisense nodes. Load balancing spreads requests across multiple query nodes according to an algorithm you define and the current status of the query node.

When implementing ElastiCube Sets, Sisense's query nodes operate in active-active mode. This means that each of the query nodes is active and can handle requests when the node is not building and its components are available. For example, traffic could be spread 50-50 across two web servers and if a component fails, a load balancer should redirect traffic to the other available web server.

The URL of your load balancer should be provided as an Alias in **Admin > Settings** in the Sisense Web Application. This directs Sisense to send traffic to your load balancer, which then sends the traffic to the relevant server.



Within an ElastiCube set, you can define how queries are routed to the available ElastiCubes. For more information, see [Routing Queries in ElastiCube Sets](#).

## Creating a High Availability Environment

There are many possible configurations you can implement to support high availability. What determines the type of configuration you choose based on three parameters:

1. **Redundancy:** Which components require redundancy?
2. **Concurrency:** How many users do you need to support?
3. **Resources:** How many systems can you add and maintain to your configuration?

Regardless of which configuration you choose, there are three steps to achieving high availability in Sisense:

1. Install Sisense on each machine in the build and query nodes. The Sisense installation includes the Sisense Web Application, ElastiCube server, and the MongoDB.
2. Replicate the MongoDB instance on the build node and distribute it to each of the query nodes. For more information, see [Replicating MongoDB](#).
3. Configure the Sisense Orchestrator Service on the build node to distribute the latest ElastiCube builds to each of the query nodes. For more information, see [Distributing ElastiCube Builds to Query Nodes](#) below.

## Distributing ElastiCube Builds to Query Nodes

The Sisense Orchestrator Service is an automated service that distributes builds and balances queries across an ElastiCube Set. The Sisense Orchestrator Service is installed during the default installation of Sisense. To distribute ElastiCube builds across all query nodes, you must configure the Sisense Orchestrator Service. When an ElastiCube is building, or receiving a build from the build node, the Sisense Orchestrator Service redirects requests to another available ElastiCube in the set.

**Note:** After configuring the Sisense Orchestrator Service, the directory `C:\ProgramData\SiSense\PrismServer\ElastiCubeData\` and the relevant ElastiCube folders within it must be shared with permissions for 'Everyone' so Sisense can build or update the ElastiCubes in those folders.

The Sisense Orchestrator Service is defined in a JSON file called `config.json` in the `Sisense.Orchestration Config` folder (`Sisense/Sisense.Orchestration/Config/`) of your Sisense Installation folder of your build node.

The `config.json` file has three main objects that you must define. The first object is the ElastiCube object that defines which ElastiCube is to be built on the build node and which ElastiCubes are located on query nodes. The second object is the Tasks object that determines how and when the build ElastiCube is to be

distributed to the query nodes. The final object that you must define is the Schedule object that defines when an ElastiCube is to be built and distributed. To distribute your ElastiCube builds through the Sisense Orchestrator Service, you define and save the config.json file. After saving the file, the Sisense Orchestrator Service automatically begins to build ElastiCubes in your ElastiCube Set. You can review the progress of the build in the /Sisense.Orchestration/Logs folder in a log file whose file name you define in the JSON file.

The objects that you must include in the config.json file and their descriptions are provided in the table below:

Key	Value
General	The General object contains two objects, logLevel and emails that define what types of logs Sisense generates and where to send them to.
LogLevel	The type of info returned in the log. The possible values you can enter are Info, debug, and Error.
Emails	This object defines who receives an email and their email address when an email is triggered according to the events you define in the mail object.
Email	The email address where emails are to be sent depending on the value of the When key.
Name	The name of the recipient of the email.
Cubes	The cubes object contains all your ElastiCubes in your configuration and their location. As the Sisense Orchestrator Service is installed on the build node, the build cubes are local, so you must define the name of the ElastiCube. For ElastiCubes on a query node, you must define the ElastiCube name, URL of the remote server, and its directory.

Key	Value
Build ElastiCubes	The Build ElastiCube object defines the ElastiCubes that will be used as Build ElastiCubes nodes.
ecube	The name of the ElastiCube.
Query ElastiCubes	The Query ElastiCube object defines the name of the remote ElastiCube, its URL address, and directory.
ecube	The name of the ElastiCube.
URL	The URL address of the ElastiCube server. If the URL is to a secure address, the value should include the username and password credentials. For example, "ssh://username:password@10.50.1.128:/C/ecubes". See Scenario 1 for an example.
localPath	Directory of the ElastiCubes.
Tasks	The Tasks object contains a task array that defines which ElastiCube should be built, the type of build, and to which ElastiCubes the build should be distributed to. Currently, Sisense only supports one Task object.
Task	The name of the task array. Currently, you can define only one task. Defining multiple tasks in the config.json file may cause builds to fail.
Build	The build object defines the ElastiCube to be built and distributed.
Cube	The name of the cube to be built.
Queue	The value of this key is the type of build Sisense should attempt and in what order. There are two possible values delimited by a comma:

Key	Value
	<p><b>accumulate:</b> Attempts an accumulative build.</p> <p><b>entire:</b> Attempts an entire build.</p> <p><b>schemaChanges:</b> Attempts to update the build only if changes were made to the schema since the previous build.</p> <p>Sisense recommends the following value: "accumulate","entire"</p> <p>In this example, Sisense attempts to do an accumulative build first, and if it fails, Sisense attempts an entire build. If the build fails for all builds, Sisense retries again according to the next scheduled build.</p>
Distribute	<p>Determines which cubes the latest build should be distributed to. The value should be the ElastiCube name for your ElastiCubes on your query nodes. For example, ["cube1, cube2, cube3"].</p> <p>You can add multiple ElastiCubes delimited by a comma separator.</p>
Reattach	<p>Reattach is an object that contains two objects, DeleteOldDbfarm and Cube.</p>
DeleteOldDbfarm	<p>A boolean value that determines if an old ElastiCube is to be deleted after it has been updated.</p>
Cube	<p>The name of the ElastiCube to be reattached or not.</p>
Mail	<p>Indicates when an email alert is triggered. Email alerts can be triggered for the following scenarios:</p> <p><b>None:</b> No email is ever sent.</p> <p><b>Build:</b> An email is sent when a build is complete regardless if the build was successful or not.</p> <p><b>Fail:</b> An email is sent when a build fails.</p>
Scheduler	<p>The Scheduler object contains an array of tasks that define when a task is to be task is to be initiated.</p>
Task	<p>Defines which task to perform and the order. Currently, Sisense</p>

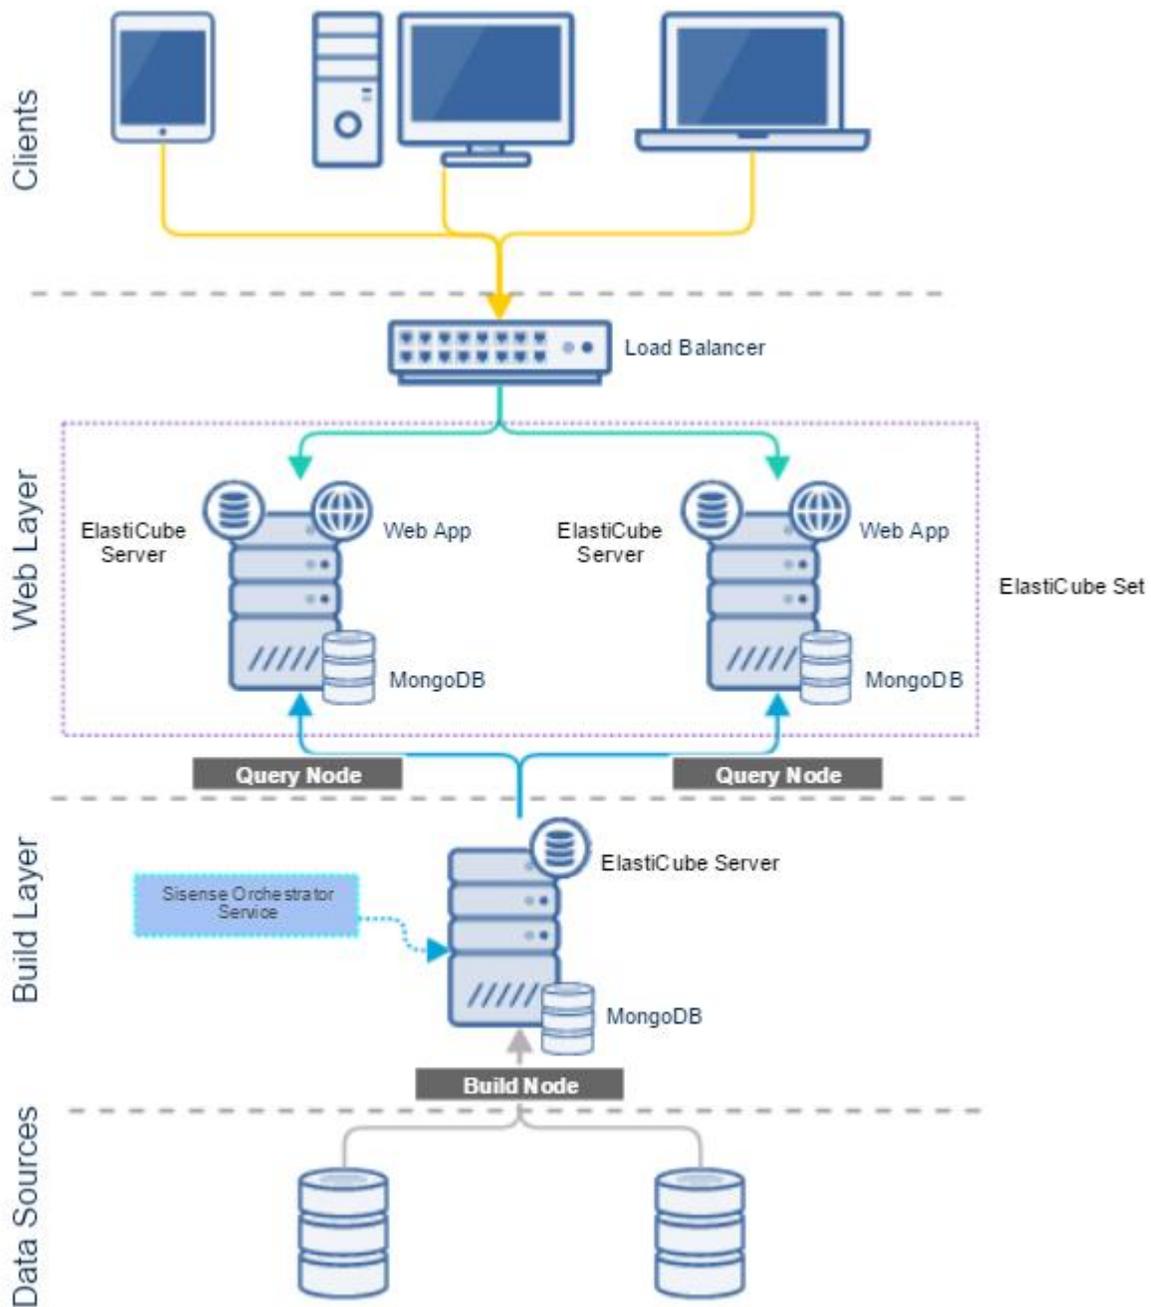
Key	Value
	only supports one task.
Schedule	The time that a task is to be initiated in Cron format. Some examples: -To run a build each night at midnight, enter the value "0 0 * * * *" -To run a build every hour, enter the value "0 * * * * *" See the full Cron format reference <a href="#">here</a> .
Enabled	A boolean value that indicates if the task is to be executed or not.

## Use Cases

The following section presents two use cases and an example config.json file that demonstrates how to support these use cases. For assistance with additional HA configurations, contact your Customer Service Manager.

### Scenario 1 – Single Stack

In this example, there is one build node and two query nodes. Each query node is hosted on a separate machine while the components that make up the query node are hosted on the same machine as a single application stack.



In this example, the build cube is called “BuildTest1”. This Elasticube is distributed to two Elasticube servers defined in the cube1 and cube2 objects. The build cube is distributed first to cube1 and then to cube2 as defined in the tasks object. The first time the Sisense Orchestrator Service tries to build the BuildTest1 Elasticube, it attempts an accumulative build and if that fails, an entire build. After each build is replicated, the previous version of the Elasticube

is deleted as specified in the reattach object. This task takes place hourly. If the distribution fails, Sisense sends an email to johndoe@Sisense.com with a debug log attached.

```
{
  "general":{
    "logLevel":"INFO",
    "emails":[
      {
        "email":"johndoe@sisense.com",
        "name":"John Doe"
      }
    ]
  },
  "cubes":{
    "localCubeBuildTest1":{
      "ecube":"BuildTest1"
    },
    "cube1":{
      "ecube":"BuildTest1",
      "url":"\\\\10.49.1.68\\ecubes",
      "localPath":"d:\\ecubes"
    },
    "cube2":{
      "ecube":"BuildTest1",
      "url":"\\\\10.70.0.32\\ecubes",
      "localPath":"e:\\ecubes"
    },
    "cube3":{
      "ecube":"BuildTest1",
      "url":"ssh://administrator:Password@10.70.1.133:/C/ecubes",
      "localPath":"e:\\ecubes"
    }
  },
  "tasks":{
    "task1":[
      {
```

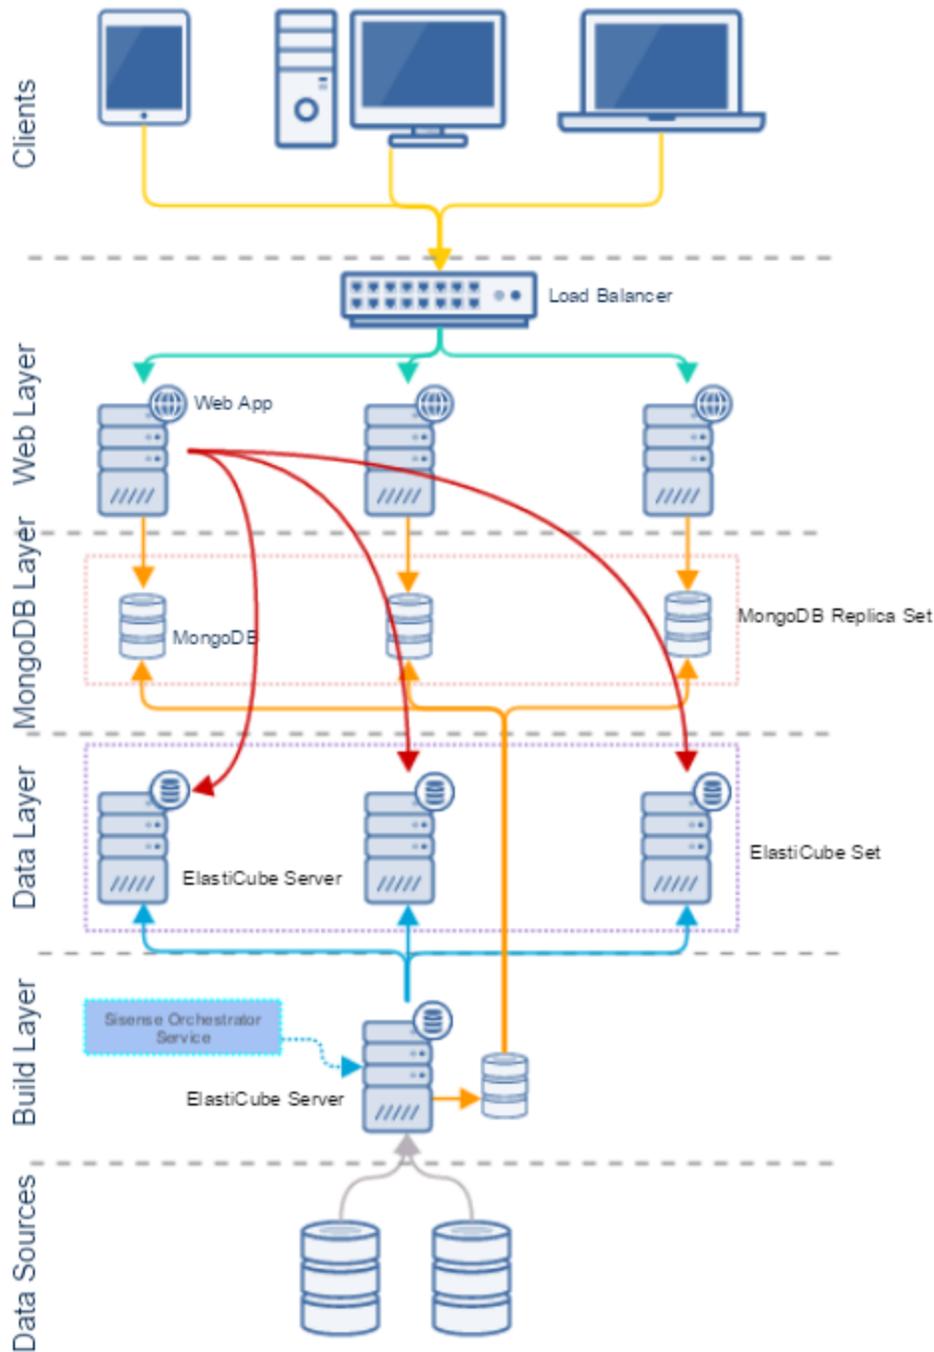
```

        "build":{
            "cube":[
                "localCubeBuildTest1"
            ],
            "queue":[
                "accumulate",
                "entire"
            ]
        }
    },
    {
        "distribute":[
            "cube1",
            "cube2",
            "cube3"
        ]
    },
    {
        "reattach":{
            "deleteOldDbfarm":true,
            "cube":[
                "cube1",
                "cube2",
                "cube3"
            ]
        }
    },
    {
        "mail":"build"
    }
]
},
"scheduler":[
    {
        "task":"task1",
        "schedule":"0 * * * *",
        "enabled":true
    }
]
}

```

## Scenario 2 – Distributed Stack

In this example, there are three query nodes in which each component is located in a different remote location. This scenario, while costly to implement, demonstrates a distributed application stack configuration where each component is replicated and located on a separate server independent of the other components. If any component fails, the rest of the architecture remains intact. For example, the web server is on one server, the ElastiCubes are stored on another server, and the MongoDB is located on another server.



The build cube in this example is called “BuildTest1”. This ElastiCube is distributed to three remote ElastiCube servers defined in the cube1, cube2, and cube3 objects. The build cube is distributed first to cube1, then cube2, and then cube3 as defined in the tasks object. The URLs are standard URLs for cube1 and

cube2 while the URL for cube3 is secured through SSH. When the URL is secured, you must provide the URL and the required credentials to access it.

The first time the Sisense Orchestrator Service tries to build the BuildTest1 ElastiCube, it attempts an accumulative build and if that fails, an entire build. If both builds fail, Sisense attempts to update the schema only. After each build is replicated, the previous version of the ElastiCube is deleted as specified in the reattach object. This task takes place At 14:15 on the 1st of every month. After a build is successful, Sisense sends an email to johndoe@Sisense.com.

```
{
    "general":{
        "logLevel":"DEBUG",
        "emails":[
            {
                "email":"johndoe@sisense.com",
                "name":"John Doe"
            }
        ],
        "cubes":{
            "localCubeBuildTest1":{
                "ecube":"BuildTest1"
            },
            "cube1":{
                "ecube":"BuildTest1",
                "url":"\\\\10.49.1.68\\ecubes",
                "localPath":          "d:\\ecubes"
            },
            "cube2":{
                "ecube":"BuildTest1",
                "url":"\\\\10.70.0.32\\ecubes",
                "localPath":          "e:\\ecubes"
            },
            "cube3":{
                "ecube":"BuildTest1",
                "url":          "ssh:"
            }
        }
    }
}
//administrator:
```

```

Password@10.70.1.133:/C/ecubes",
  "localPath": "e:\\ecubes"
}
},
"tasks":{
"task1":[
{
"build":{
"cube":[
"localCubeBuildTest1"
],
"queue":[
"accumulate",
"entire",
"schemachanges"
]
}
},
{
"distribute":[
"cube1",
"cube2",
"cube3"
]
},
{
"reattach":{
"deleteOldDbfarm":true,
"cube":[
"cube1",
"cube2",
"cube3"
]
}
},
{
"mail":"fail"
}
]
},

```

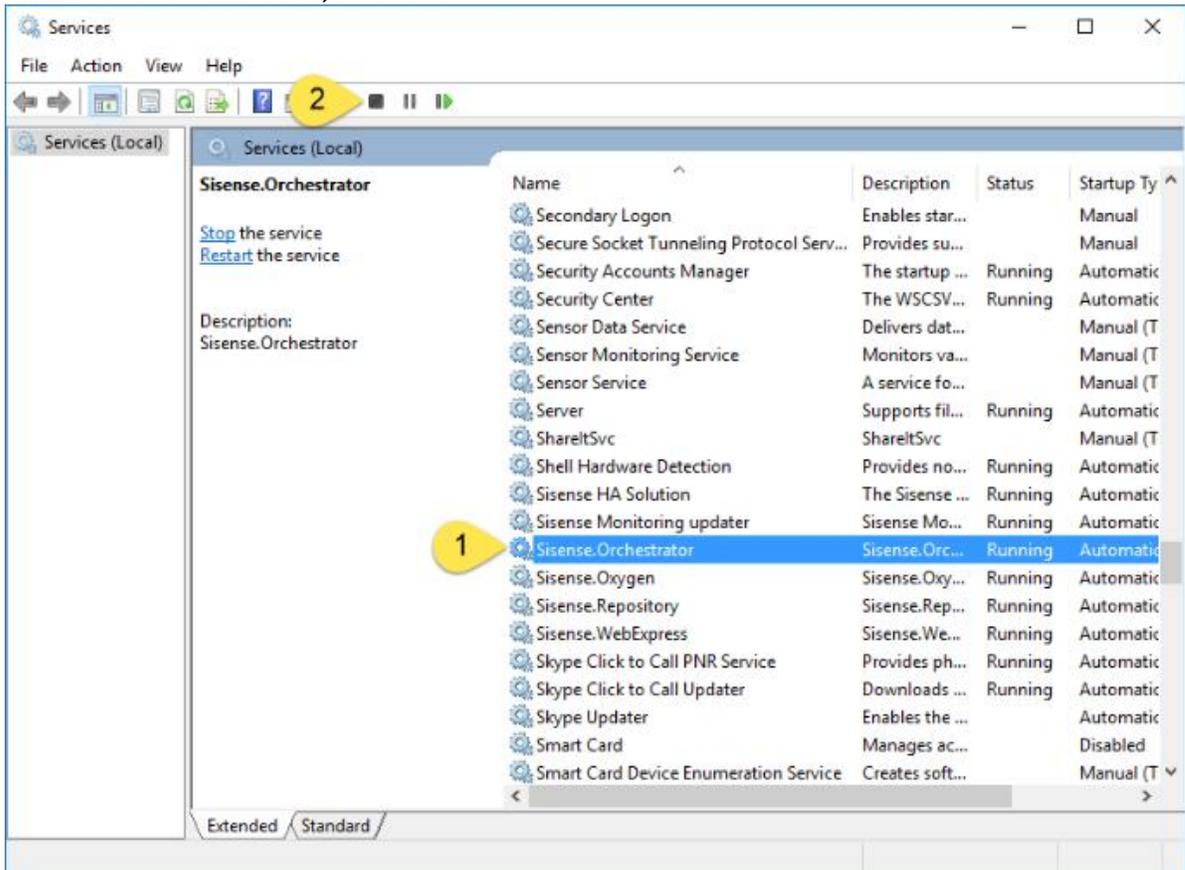
```
    "scheduler": [
      {
        "task": "task1",
        "schedule": "15 14 1 * *",
        "enabled": true
      }
    ]
  }
```

## Add Query Nodes

In some configurations, replicating components can improve performance by scaling out and preventing any potential bottlenecks. Each component you replicate must be added to the ElastiCube Set and the config.json file. You can continue to add query nodes according to your requirement so long as the nodes are included in the config.json file.

**To add a query node:**

1. Stop the Sisense Orchestrator Service. Open Windows Services, select Sisense.Orchestrator, and click .



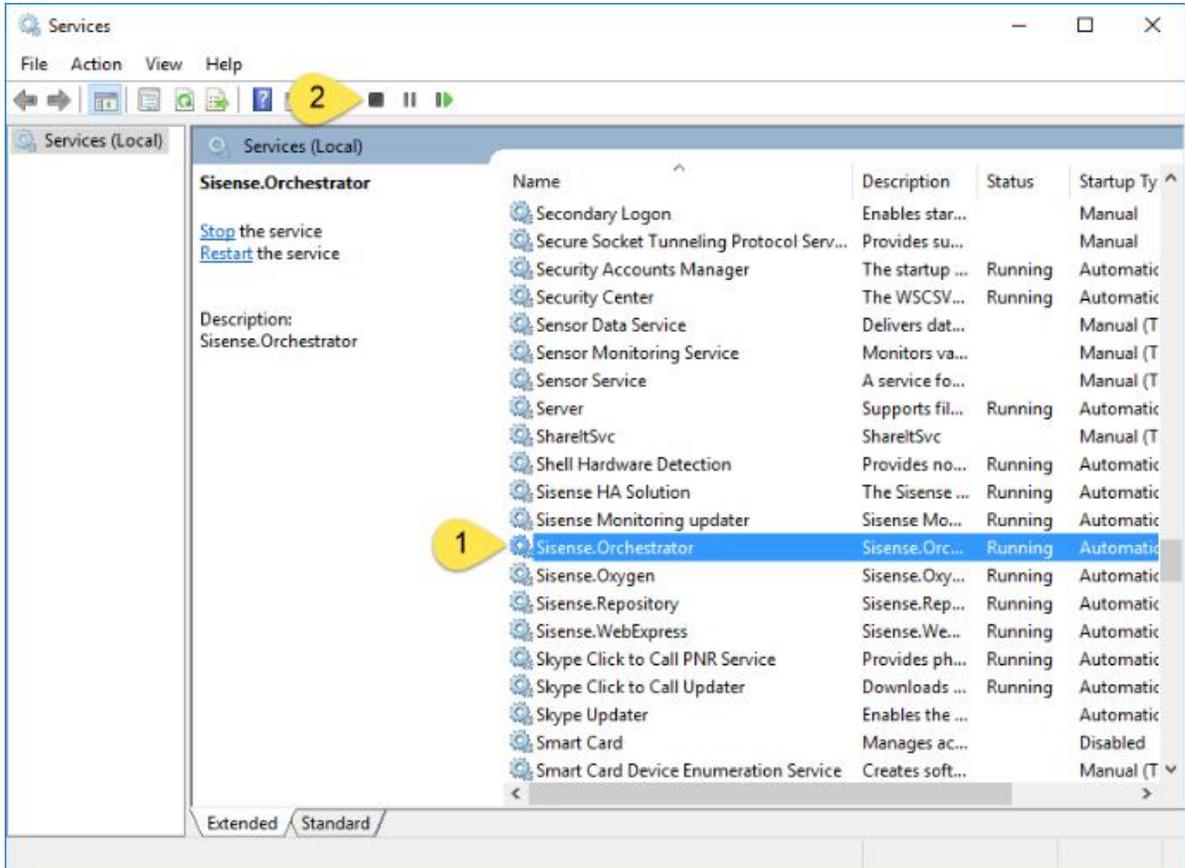
2. In the Sisense Web Application, click **Admin** and select the **Data Sources** tab on the left.
3. Hover over the ElastiCube you want to add to the ElastiCube Set and click on the menu that appears. Click **Add to ElastiCube Set** and select the set that you want to add the ElastiCube to.
4. Edit the config.json located at Sisense/Sisense.Orchestration/Config/ and add the new ElastiCube to the Cubes object.
5. Save the config.json file.
6. Restart the Sisense Orchestrator Service.

## Remove Query Nodes

You can remove query nodes by removing them from the ElastiCube Set and the config.json file.

## To remove a query node:

1. Stop the Sisense Orchestrator Service. Open Windows Services, select Sisense.Orchestrator, and click .



2. In the Sisense Web Application, click **Admin** and select the **Data Sources** tab on the left.
3. In the ElastiCube Set table, click and clear the checkboxes of the ElastiCubes to be removed. Click **Save** after you have selected the relevant ElastiCubes.
4. Edit the config.json located at Sisense/Sisense.Orchestrator/Config/ and delete the relevant ElastiCubes from the Cubes object.
5. Save the config.json file.
6. Restart the Sisense Orchestrator Service.

## Limitations

If your user queries a query node (Query Node 1), then the build node begins to replace a build on that node (Query Node 1), there are two scenarios in which the data may be inconsistent.

1. When generating PDFs or Images, the widgets may not match the dashboard depending.
2. In a Pivot widget, the subtotal value may not match the expected value.

This occurs because the widgets exported to PDF or subtotal values inside Pivot widgets run new queries against the node. If the data was originally taken from query node 1, then subsequently queries were run against another query node (Query Node 2), the data may have changed leading to inconsistent results.

# Replicating the Application Database

This information is relevant for versions of Sisense V6.7 and later. For earlier versions, click [here](#).

When you install Sisense, an application database is installed locally on your machine that contains metadata for ElastiCubes, dashboards, and users. If you have several application databases across multiple servers, you can combine them to form a replica set. The goal of a replica set is to ensure that each application database has the same metadata about ElastiCubes, dashboards, and filters. If one database fails on one query server, the next available database can replace it. To maintain consistency throughout your configuration and support high availability, each database must be identical with the other databases in your configuration.

There are two types of nodes within a replica set. The primary node, which supports write operations and the secondary nodes, which replicate the primary database's operations log and syncs the data to reflect that of the primary. The primary and secondary nodes can be hosted on the build or query nodes in a high availability configuration on servers with or without Sisense installed.

There are four steps to support high availability in Sisense:

1. **Setting up replica nodes:** Each database must be configured to connect to other instances in the Replica Set and share a KeyFile, that is used for authentication.
2. **Creating a replica set:** After setting up your nodes, you need to set one as the primary, which has write access to the other nodes in the replica set.
3. **Connecting to a replica set:** Once you have created the replica set, now it's time to connect it to Sisense.

4. **Connecting Sisense Web Application to the replica set:** Your replica set can support your Sisense Web Application. To do so, you need to connect it to the replica set instead of the localhost that it is connected to by default. This way, if your server's database fails, the Sisense Web Application is supported by another database within the replica set.

The Sisense uses MongoDB as its application database. This page describes how to set up your databases in a replica and connect it to Sisense.

### Prerequisites

1. Servers with port 27018 open.
2. RoboMongo installed on a machine that has access to those servers on port 27018 (any node can be on one of the servers).
3. MongoDB installed on the machine. By default, this is installed when Sisense is installed. If Sisense is not installed, download MongoDB version 3.4.5 zip on the secondary servers from [here](#).

## Setting up Replica Nodes

**Note:** For each machine that you have installed Sisense, you must stop the Sisense-related services from running to prepare the primary database that will be replicated to the other secondary nodes.

Each node in your replica set must have a MongoDB configuration file. The configuration file contains information about the database such as its open port and the name of the replica set the application database belongs to. When Sisense is installed on the node, this configuration file is created for you. For nodes in which Sisense is not installed, you must create this configuration file as described in the procedure below.

In addition, each database must share a KeyFile used for authentication. This KeyFile must be the same on each instance, whether Sisense is installed or not.

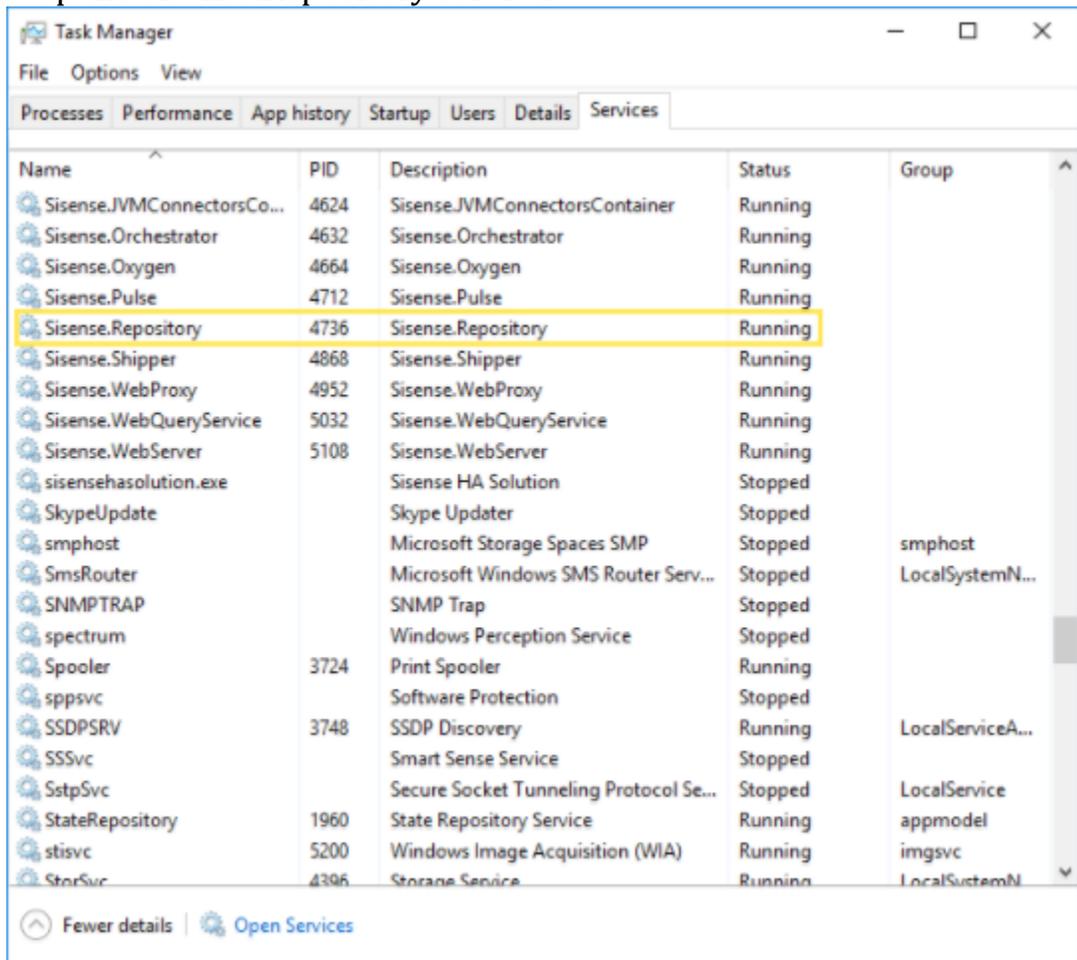
This procedure describes how to set up a replica node on a Sisense

### To set up a replica node:

1. Create a KeyFile. See [Creating KeyFiles](#) for more information.
2. Create a MongoDB user with write access. See [Accessing the Application Database](#) for more information.

### 3. When Sisense is installed:

1. Stop the Sisense.Repository service.



2. Copy the Key File you prepared above as described in Step 1 to the directory ... \Program Files \Sisense \Infra \MongoDB.
3. In the directory, '... \Program Files \Sisense \Infra \MongoDB \mongodbconfig.conf', backup the MongoDB configuration and save the backup copy in another directory.
4. In the MongoDB directory, open the mongodbconfig.conf file and uncomment the commented lines by removing # from the beginning of the row.  
storage:  
dbPath: "C:/ProgramData/Sisense/PrismWeb/Repository/DB"  
journal:  
enabled: true  
systemLog:  
destination: file

```

path:
"C:/ProgramData/Sisense/PrismWeb/Repository/Logs/sisenseRepository.log"
net:
bindIp: 127.0.0.1
port: 27018
# ssl:
# mode: requireSSL
# PEMKeyFile: keyFile path
# PEMKeyPassword: keyFile password
security:
authorization: "enabled"
#keyFile: "C:/Program Files/Sisense/Infra/MongoDB/keyfile"
# replication:
# replSetName: replica set name

```

5. Bind each of your servers in your high availability configuration together by adding their IP addresses and separating them with a comma:

**For Sisense V6.7 and earlier:**

```

net:
bindIp: [127.0.0.1, XX.XX.XX.XX, XX.XX.XX.XX]

```

**For Sisense V6.7.1 and later**

```

net:
bindIp: 127.0.0.1,XX.XX.XX.XX

```

**Note:** The value of bindIp has been changed from an array to a string. As a string, the IP addresses should be separated by a comma with no space in between. In addition, the only IP addresses to be included in the string as 127.0.0.1 and the public IP of your local server.

6. Under **security**, define the location of your KeyFile created in Step 1.

```

security:
authorization: "enabled"
keyFile: "C:/Program Files/Sisense/Infra/MongoDB/keyfile"

```
7. Set a name for the replica set.

```

replication:
replSetName: replicasetname

```

Save and close the mongodconfig.conf file.
8. Start the Sisense.Repository service.

**4. When Sisense is not installed:**

1. Copy the Key File you prepared above as described in Step 1 to your MongoDB's configuration file folder.

2. In the MongoDB configuration folder, open the mongodconfig.conf file.
3. Bind each of your servers in your high availability configuration together by adding their IP addresses and separating them with a comma:  
net:  
bindIp: [127.0.0.1, XX.XX.XX.XX, XX.XX.XX.XX]  
**Note:** You can add your PC's IP address as well to allow you connect to the MongoDB with an application such as RoboMongo.
4. Under **security**, define the location of your KeyFile created in Step 1.  
security:  
authorization: "enabled"  
keyFile: "C:/MongoDB/keyfile"
5. Set a name for the replica set.  
replication:  
replSetName: <replicasetname>
6. Save and close the mongodconfig.conf file.

## Creating a Replica Set

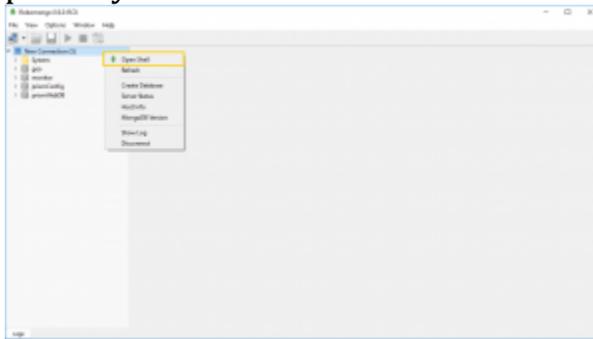
After you have prepared the MongoDB configuration file for each of your nodes, you must create a replica set with these nodes. With a replica set, each member has the same Sisense data but is otherwise independent. If the primary becomes unavailable, the replica set holds an election to select a new primary from one of the secondary nodes.

You should perform this procedure on your primary node. It sets the server as the primary node, and connects the secondary nodes to it.

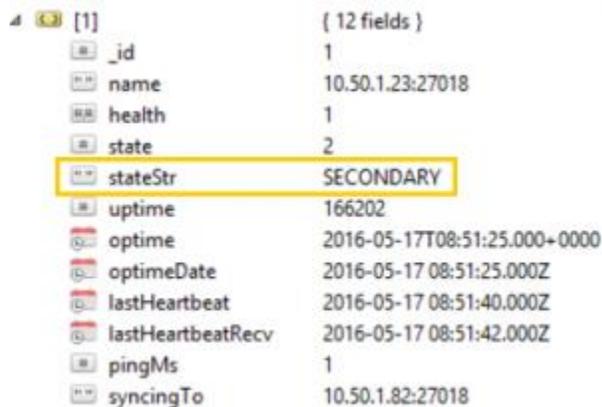
### To create replica set:

1. Open Robomongo with a WriteUser account as described in [Accessing the Sisense Application Database](#) and select Open Shell to connect to the

primary node.



2. In the new shell, run the command `rs.initiate()` to create a replica set.
3. Run the following commands:  
`cfg = rs.conf()`  
`cfg.members[0].host = "XX.XX.XX.XX:27018"` (Enter your IP address of the primary node and the MongoDB port, by default 27018).  
`rs.reconfig(cfg)`
4. Run the command `rs.add("XX.XX.XX.XX:27018")`, where the XX is the IP address of your secondary node. Run this command to add the unique IP address for each secondary node. For example:  
`rs.initiate()`  
`cfg = rs.conf()`  
`cfg.members[0].host = "10.50.1.199:27018"`  
`rs.reconfig(cfg)`  
`rs.add("10.20.1.196:27018")`  
`rs.add("10.20.1.197:27018")`
5. Run the command **`rs.status()`** to verify that the replica set was configured properly. Run this command every couple of minutes to check the progress until `stateStr` of the secondaries service displays **SECONDARY**.



Field	Value
_id	1
name	10.50.1.23:27018
health	1
state	2
stateStr	SECONDARY
uptime	166202
optime	2016-05-17T08:51:25.000+0000
optimeDate	2016-05-17 08:51:25.000Z
lastHeartbeat	2016-05-17 08:51:40.000Z
lastHeartbeatRecv	2016-05-17 08:51:42.000Z
pingMs	1
syncingTo	10.50.1.82:27018

# Connecting Sisense to the Replica Set

After you have created the replica set, you must connect the primary web application server to your replica set to support your high availability configuration.

To connect your replica set with your primary web server, you need to configure the db.config and default.yaml files with the name of your replica set and the IP addresses of each of the servers in the set.

## To modify the primary server's configuration files:

1. In the directory ...\\Program Files\\Sisense\\PrismWeb\\App\_Data\\Configurations backup the file **db.config**.
2. In the original db.config file of the primary server, modify the value of the parameter **connectionString** to:  
"mongodb://AppUser:we6jBUYGoOrh0K6l+XpTmA==@XX.XX.XX.XX:27018,XX.XX.XX.XX:27018,XX.XX.XX.XX:27018/admin?replicaset=<replicasetname>&readPreference=primaryPreferred"

replacing '@localhost:27018/admin' with the IP addresses of each of your nodes in the replica set separated by a comma and replacing replicasetname with the name of your replica set. After you have added the IP addresses and your replica set name, your connection string should look like this:

```
<add name="MongoProvider"
connectionString="mongodb://AppUser:we6jBUYGoOrh0K6l+XpTmA==@
XX.XX.XX.XX:27018,XX.XX.XX.XX:2
018,XX.XX.XX.XX:27018/admin?
replicaset=<replicasetname>&readPreference=primaryPreferred"
mongoAppIV="ePtH5UDfff6YvSR0ihwkcQ=="
mongoAppKey="dDRRM1p3SSS85N0JlQ3g4Mg=="SslPfxFilePath=""
SslPfxFilePassword=""/></connectionStrings>
```

## For example:



```
1 <add name="MongoProvider" connectionString="mongodb://localhost:27018/" />
2 <add name="MongoProvider" connectionString="mongodb://localhost:27018/" />
3 <add name="MongoProvider" connectionString="mongodb://localhost:27018/" />
4 <add name="MongoProvider" connectionString="mongodb://localhost:27018/" />
5 <add name="MongoProvider" connectionString="mongodb://localhost:27018/" />
6 <add name="MongoProvider" connectionString="mongodb://localhost:27018/" />
7 <add name="MongoProvider" connectionString="mongodb://localhost:27018/" />
8 <add name="MongoProvider" connectionString="mongodb://localhost:27018/" />
9 <add name="MongoProvider" connectionString="mongodb://localhost:27018/" />
10 <add name="MongoProvider" connectionString="mongodb://AppUser:we6jBUYGoOrh0K6l+XpTmA==@10.10.10.10:27018,10.10.10.10:27018,10.10.10.10:27018/admin?replicaset=rs0&readPreference=primaryPreferred" mongoAppIV="ePtH5UDfff6YvSR0ihwkcQ==" mongoAppKey="dDRRM1p3SSS85N0JlQ3g4Mg==" SslPfxFilePath="" SslPfxFilePassword="" /></connectionStrings>
```

**Note:**The only values that should be changed are the IP addresses in the connectionString and replicasetname.

3. In this directory "... \Program Files \Sisense \PrismWeb \vnext \config", backup the file default.yaml
4. In the original default.yaml file, uncomment the section replica\_set replacing the value of the parameter name with the name of your replica set and XX.XX.XX.XX with the IP addresses of each of the nodes in your replica set:

replica\_set:

name: <replicasetname>

servers:

- host: XX.XX.X.XX

port: 27018

- host: XX.XX.X.XX

port: 27018

- host: XX.XX.X.XX

port: 27018

**For example:**

```
db_server:
  host: 127.0.0.1
  port: 27018
  process: \mongodb\mongod.exe
  dbpath: '/sisense/PrismWeb/UpgradedDB'
  ttlNotFound: 604800

db_security:
  mongoAppUser: AppUser
  mongoAppPassword: 2eXY2+I15fd5PnFbYMH==
  mongoAppKey: EDB3jh5Q2IIH13pQm9-Q==
  mongoAppIV: ePtH*FDHF83KVSR0ihwkoQ==
  sslEnabled: true
  sslValidate: false
  certFilePath: 'C:/ProgramData/Sisense/PrismWeb/ssl/my-cert.pem'
  pemKeyFilePath: 'C:/ProgramData/Sisense/PrismWeb/ssl/pem-key.pem'
  caFilePath: 'C:/ProgramData/Sisense/PrismWeb/ssl/my-cert.pem'
  pemKeyPassword: password

replica_set:
  name: 'highavailability'
  servers:
    - host: 10.70.18.62
      port: 27018
    - host: 10.70.18.54
      port: 27018
    - host: 10.70.18.61
      port: 27018
```

5. Depending on your Sisense installation,  
**Full IIS:** Open the IIS Manager and start the website SisenseWeb.  
**IIS Express:** Open Windows Services and start the service Sisense.WebExpress.

# Connecting Additional Web Servers to your Set

To add additional query nodes to support your primary web server, each additional server must be configured with the IP address of the primary and its AppUser credentials. The AppUser defines who can access the application database. For more information about AppUsers, see [Accessing the Sisense Application Database](#).

This procedure describes how to configure your query nodes for high availability.

## To configure your query nodes:

1. In the directory `..\Program Files\Sisense\PrismWeb\App_Data\Configurations` backup the file **db.config**.
2. In the original **db.config** file for each of your query nodes, replace the following line with the lines from the same file in the **Primary Node**.

```
<add name="MongoProvider"
connectionString="mongodb://AppUser:we6jBUYGoOrh0K6l
+XpTmA==@XX.XX.XX.XX:27018,XX.XX.XX.XX:27018,XX.XX.X
X.XX:27018/admin?replicaset=<replicasetname>&readPreference=primaryPreferred"
mongoAppIV="ePtH5UDfff6YvSR0ihwkcQ=="
mongoAppKey="dDRRM1p3SSS85N0JlQ3g4Mg=="
SslPfxFilePath=""
SslPfxFilePassword="" /></connectionStrings>
```

Where `XX.XX.XX.XX:27018` is the IP address and port of your primary node and `<replicasetname>` with your replica set name. All the content from the parameter `connectionString` should be included including `mongoAppIV` and `mongoAppKey` as this information is essential to enabling a connection between the primary and secondary nodes.
3. Open `default.yaml` file located in `C:\Program Files\Sisense\PrismWeb\vnex\config\` and edit the file and replace the `db_security`, `replica_set` and `db_users` sections with your parameters from the `default.yaml` file stored on the Primary Node.  
`db_security:`

```
mongoAppUser: AppUser
mongoAppPassword: we6jBUsdRh0K6l+XpTmA==
mongoAppKey: dDRsdm85N0JlQ3g4Mg==
mongoAppIV: ePtH5ds6YvSR0ihwkcQ==
replica_set:
name: 'replicaSet name'
servers:
- host: hostName
port: port
- host: hostName
port: port
- host: hostName
port: port
```

4. Depending on your Sisense installation:

**Full IIS:** Open the IIS Manager and start the website SisenseWeb.

**IIS Express:** Open Windows Services and start the service Sisense.WebExpress.

## Defining the Application Database's Read Preferences

Read preferences specify where each member of a replica set should direct its read operations.

If you have implemented high availability with replica sets to support your Sisense implementation, and your application databases are geographically spread out, you may want to define different read preferences to different nodes within the set.

Defining your own read preferences is useful in cases such as preventing a secondary node which has fallen behind a primary node to giving old data.

You can define the following read preferences for the application database:

- **primary:** This preference routes all reads to the replica set primary. If the primary is unavailable for some reason a read operation would produce an error or exception.
- This preference is useful to never return stale data.

- (Sisense Default) primaryPreferred: Reads are sent to the primary, but if it is unavailable, the reads come from secondary members instead.
- secondary: Reads are only allowed on secondary members of the replica set. When no secondaries are available, an error or exception is returned.
- secondaryPreferred: Reads are routed to a secondary, but if no secondary is available read operations are sent to the primary.
- Nearest: Reads are performed on the nearest available node, disregarding if it is a primary or secondary member. The nearest node is determined by sending pings to all members and measuring the response time.

For additional information about read preferences, [click here](#).

To modify the application database's read preferences, you must update two files located in your Sisense server, the db.config file and the default.yaml file.

The db.config file is located in the following directory:

C:\Program Files\Sisense\PrismWeb\App\_Data\Configurations

Open the db.config in a text editor and add the property

readPreference="yourPreferenceType" to the MongoProvider connection string.

```
<add name="MongoProvider"
connectionString="mongodb://AppUser:mzfsb/DqUHAu+fsdAg==@
10.60.22.54:27018,10.60.22.64:27018,10.60.22.74:27018/adm
in?replicaset=sisense;readPreference=primaryPreferred"
mongoAppIV="ePtHfsdvvSR0ihwkcQ=="
mongoAppKey="ZzJBOFhfsdTW02aw==" SslPfxFilePath=""
SslPfxFilePassword=""/></connectionStrings>
```

The default.yaml file is located in the following directory:

C:\Program Files\Sisense\PrismWeb\vnext\config

Open the default.yaml file in a text editor and add the property readPreference:

'yourPreferenceType' to the replica set object.

```
replica_set:
name: 'replicaSet name'
readPreference: 'primaryPreferred'
servers:
```

```
- host: hostName
```

```
port: port
```

```
- host: hostName
```

```
port: port
```

```
- host: hostName
```

```
port: port
```

**Note:** Remember, the comments '#' must be removed from each row of the replica\_set object.

# Routing Queries in ElastiCube Sets

When a user queries an ElastiCube that is part of an ElastiCube Set, Sisense routes the query to a single ElastiCube with the latest ElastiCube build. By default, ElastiCube Sets operate in active-passive mode. In active-passive mode, one ElastiCube, by default the ElastiCube with the latest build, handles incoming queries. This ElastiCube is active while the remaining ElastiCubes are on standby in passive mode.

Sisense also supports active-active mode where queries are spread across available ElastiCubes within an ElastiCube Set. In active-active mode, each ElastiCube that is not currently building can handle queries.

You can modify how queries are routed to ElastiCubes within ElastiCube Sets by selecting one of the following routing strategies:

**Single:** (Default). ElastiCubes operate in active-passive mode where the ElastiCube with the freshest data receives all the queries while the remaining ElastiCubes are on standby.

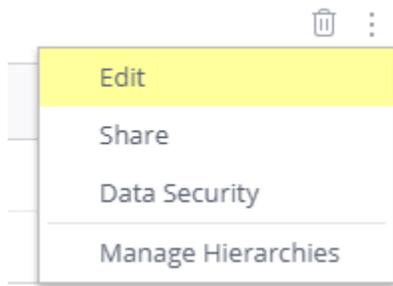
**Multiple ElastiCubes – Fresh Data:** ElastiCubes operate in active-active mode where queries are spread across all ElastiCubes with the latest data.

**Multiple ElastiCubes – Best Spread:** ElastiCubes operate in active-active mode where queries are spread across the largest group of ElastiCubes with identical data.

**To set a query routing strategy:**

1. Click **Admin** and select the **Data Sources** tab on the left.

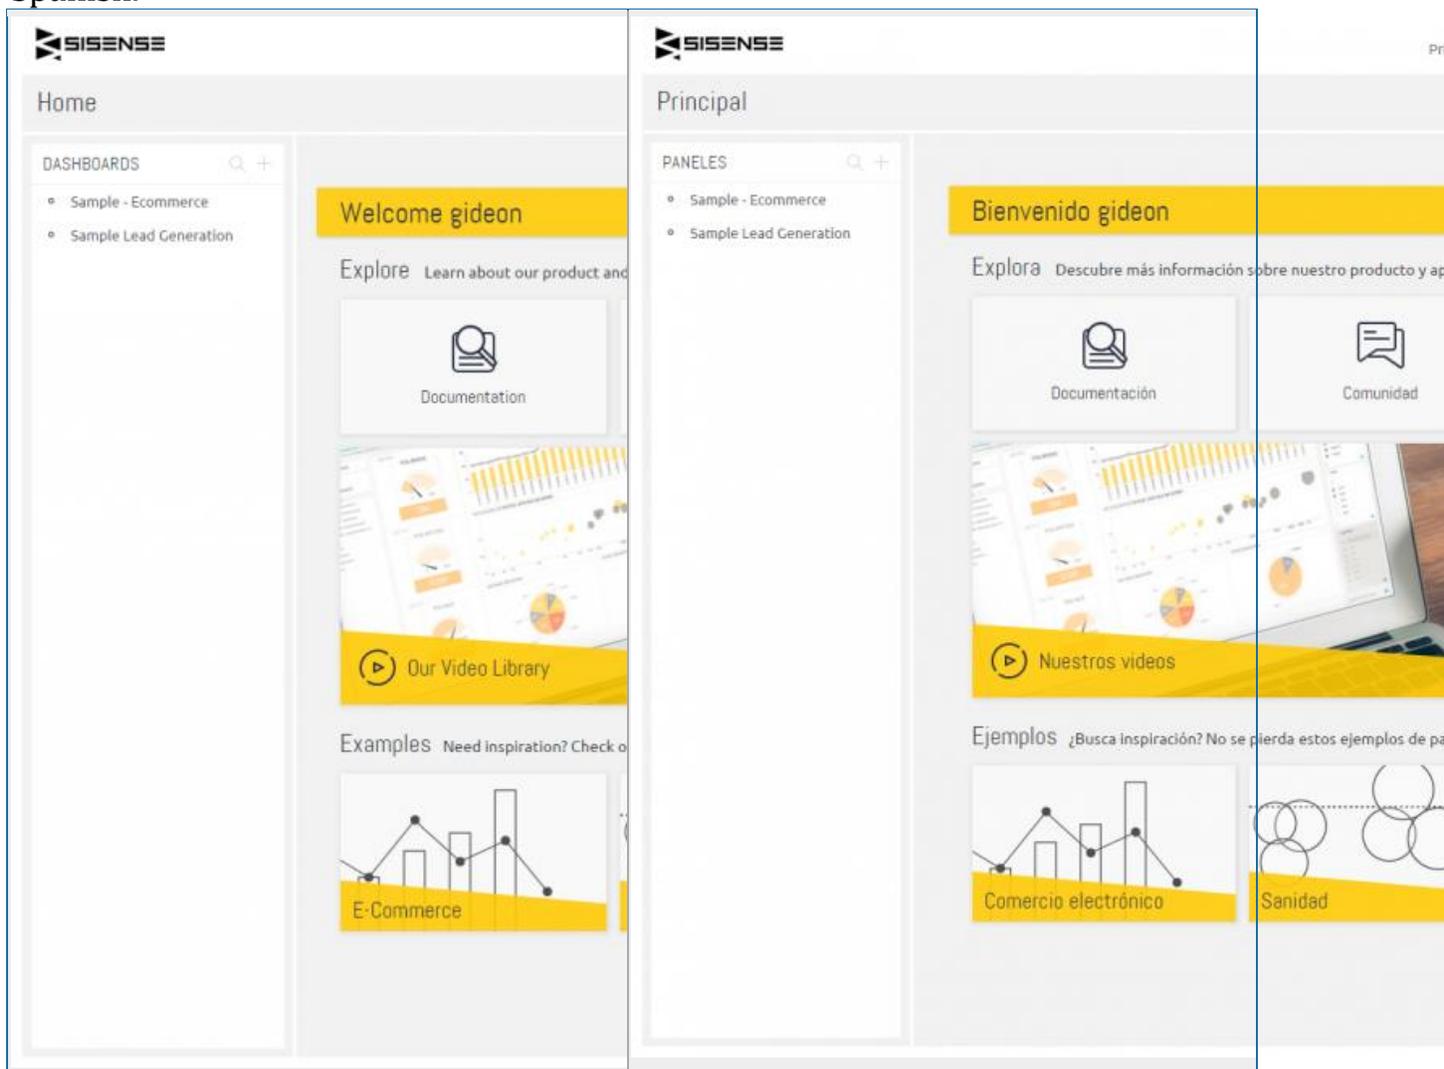
2. For the relevant ElastiCube Set, select **Edit**.



3. From the Routing Strategy list, select the relevant strategy.
4. Click **Save**.

# Changing the Sisense Web Application's Language

Sisense supports multiple languages in the Sisense Web Application making it easier to extract meaningful insights from your data in your users' local language. The following images display two interfaces, one in English and the other in Spanish.



English	Spanish
---------	---------

## What has been Translated?

Sisense has translated text that your users might see including menu names, buttons, messages, and other elements in the Sisense Web Application.

Translating the Sisense Web Application does not affect your data or how it is displayed.

The Sisense Web Application has been localized to the following languages:

- Chinese (Simple)
- French (France)
- German
- Italian
- Japanese
- Portuguese (BR)
- Russian
- Spanish (ES)
- Spanish (LA)
- Dutch

Note, the following content is not being localized:

- Built-in Javascript editor
- ElastiCube Manager and Server Console
- Sisense documentation and online help
- REST API and documentation
- Sisense automated emails

**Note:** You can use your own translation of Sisense automated emails by defining your own. [Click here](#) for more information.

## Changing Sisense's Default Language

Administrators can select the default language for the Sisense Web Application. Setting the default language affects all users and groups who access the Sisense

dashboard. Additionally, Administrators can then set the language for specific groups and users when adding new users or by editing a group or user account. After selecting a new language, each string with a translation is displayed in that language across all your Sisense users' interfaces.

If you have defined a custom translation, the name of the language as defined in the settings.json is displayed in this list. See [Embedding Custom Languages](#) below for more information.

Sisense's default language is English. After changing the system language, you can revert back to the English or any other language by selecting it from the list.

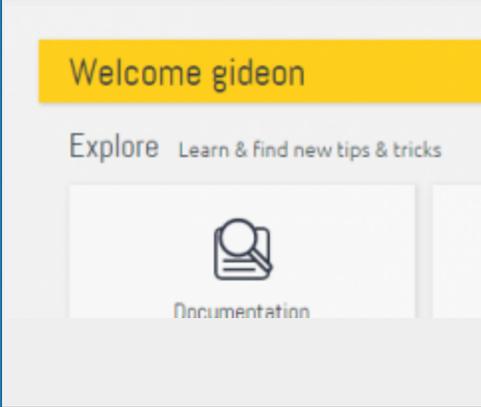
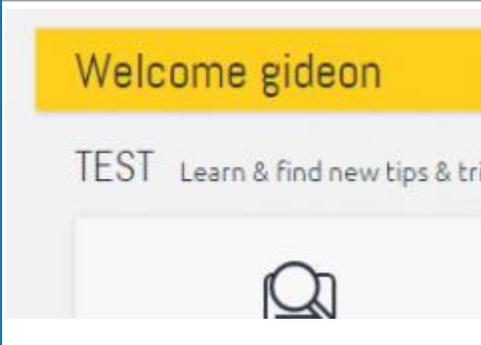
#### **To modify your default language:**

1. In the Sisense Web Application, click **Admin** in the top menu, and then **Settings** in the left menu.
2. Select the relevant language from the Language list. After selecting the language, the UI is immediately translated.
3. Click **Save**. The Sisense Web Application is translated in the selected language.

## **Embedding Customized Languages**

Sisense enables you to display your own translations in the Sisense Web Application by replacing the values of strings in any language you choose. You can embed customized languages by modifying configuration and resource files replacing the string values according to your requirements.

For example, in the homepage of the Sisense Web Application, the word "Explore" is displayed. In the resource file home.js, you can modify the value of this string to display your own text as shown below.

	<pre>"welcome_user": "Welcome", "explore": { "explore": "Explore", "learn": "Learn &amp; find new tips &amp; tricks", "links": {</pre>
	<pre>"welcome_user": "Welcome", "explore": { "explore": "TEST", "learn": "Learn &amp; find new tips &amp; tricks", "links": {</pre>

In addition to modifying the resource files, you must modify the settings.json file and specify your language's name and update the version. In this example, the language has been changed to "test" and the version increased from 1.0.0 to 1.0.1.

```
{ "displayName": "test", "version": "1.0.1" }
```

The value of "displayName" is displayed as your language's name in Default Language list on the System Configuration page. In the example below, the displayed was set to Test.

Upgrading Sisense may replace the bundled translation files. If you wish to modify a bundled language, copy of all of the translation files to a new directory, and change the copied versions to prevent your translation from being overwritten.

### To customize the Sisense Web Application's language:

1. Navigate to the directory "...\\Sisense\\PrismWeb\\translations\\"
2. In the \\translations directory, copy the **en-us(default)** directory, paste it in the \\translation directory and rename it.
3. In the new directory, open each resource file and translate the strings into the desired language saving each file after your changes are made.

4. Open the settings.json file and replace the following key values:  
**displayName:** Enter the name of the language.  
**version:** Enter the name of the new version. Currently, the version is 1.0.0. For each version, increase this value, for example, 1.0.1.
5. Save the changes in a file called settings.json in your new language folder.
6. In the Sisense Web Application, click Admin in the top menu, and then Settings in the left menu.
7. Select your customized translation from the Language list. The Sisense Web Application is displayed in your customized translation.
8. Click **Save**.

# Internationalization and Localization

Sisense supports localization and internationalization by displaying your users' native language in the Sisense Web Application. You can select the language displayed in the Sisense Web Application from a list of languages provided by Sisense or display your own translation. Each time you select a language, the Sisense Web Application interface is displayed in that language.

When you change the default language displayed in the Sisense Web Application, words that are part of the Sisense framework such as menu items, buttons, and messages are displayed in the selected language. Words specific to your dashboard, such as table names, field names, and filters, are not automatically translated when you select a new language to display. These words are called metadata and you can provide your own translation of these terms and display those translations to your users through a Dashboard script.

For instructions on how to change the language displayed in the Sisense Web Application, see [Changing the Sisense Web Application's Language](#).

For instructions on how to translate your metadata, see [Translating Sisense Metadata](#).

# Translating Sisense Metadata

Sisense enables you to translate metadata in your users' native language. Sisense metadata is displayed in your users' dashboards and includes the following:

- ElastiCube fields
- ElastiCube tables
- Widget Titles
- Dashboard Names
- Folder Names
- Saved formulas
- Default and custom hierarchies
- Customized renamed field titles

The following images display two interfaces with metadata, one in English and the other in Russian.

	
English	Russian

As fields such as ElastiCube names and tables are defined by you, these fields are not translated when changing the default language of the Sisense Web Application.

You can translate your metadata by modifying the default values using a JavaScript plugin. Through the plugin, you define the new values to be displayed. You may also include additional logic to determine which users or user groups the translated values should be displayed to.

When you create a JavaScript plugin, you should create a folder for the plugin and place it inside the plugins folder in the following location:

```
...\Program Files\Sisense\PrismWeb\plugins\
```

If the `\plugins\` folder does not exist, create it and then create a new folder within the `\plugins\` folder with the relevant files.

You can download an example of a Metadata Translation plugin [here](#).

Extract the enclosed folder into the plugins folder, `...\Program Files\Sisense\PrismWeb\plugins\`.

This file contains a folder called Metadata and two files:

**plugin.json:** The `plugin.json` file contains the name of the plugin and a link to the file that describes the plugin's functionality, in this example, `run.js`.

**run.js:** The `run.js` file contains the plugin's functionality and describes which metadata is to be translated and how it is to be translated.

This example, and any metadata translation JavaScript file you write should include three parts:

- **Provider Function:** These functions provide the mapping context contained in the Mapping object.
- **Event Subscription:** This event is triggered when a datasource changes, and calls the Provider function.
- **Mapping Object:** Contains key/value pairs with your default and translated metadata.

## Provider Functions

The code sample below is an example of two Provider functions, the first is a global provider function for folder and dashboard aliasing. The second is the datasource provider function for translating ElastiCube and formula names. You can replace the parameters in this example to provide the mapping context you define in the Mapping object.

### Global Provider

```
function (resolve, reject) {  
  //an ajax request or any other logic.  
  setTimeout(function () {  
    resolve(globalAliasing);  
  }, 0);  
};
```

The following table describe the Provider function and its arguments:

Parameter	Description
resolve	A callback function which should be called with the provided aliasing context object or without any parameter if aliases shouldn't be provided for the datasource.
reject	A callback function to be called in case of error, so the system can properly handle it.

### Datasource Provider

```
function (datasourceId, resolve, reject) {  
  //an ajax request or any other logic.  
  setTimeout(function () {  
    resolve(datasourceAliasing);  
  }, 0);  
};
```

The following table describe the Provider function and its arguments:

Parameter	Description
datasourceId	An object that represents the ID of the datasource with the following properties: <ul style="list-style-type: none"><li>• address</li><li>• database</li><li>• fullname</li><li>• ID</li><li>• title</li></ul>
resolve	A callback function which should be called with the provided aliasing context object or without any parameter if aliases shouldn't be provided for the datasource.
reject	A callback function to be called in case of error, so the system can properly handle it.

## Event Subscription

To trigger the Provider functions, your plugin should include a handler for a global event 'beforealiascontextinit'.

In the example below, datasourceProviderFunction and globalAliasingProvider are the Provider functions. The timeout parameter is optional and configures the amount of time in milliseconds which the system waits for the aliasing context to be resolved by the provider function. By default the system uses a timeout of 2000 ms

```
prism.on("beforealiascontextinit", function (ev, args) {  
  var timeout = 2000;  
  args.register(datasourceProviderFunction,  
    globalAliasingProvider, timeout);  
});
```

## Mapping Object

The translation context is a simple JavaScript object which contains key value pairs for tables, columns, hierarchies, formulas and custom item titles, all case-insensitive.

The sample below is a skeleton of a Mapping object that defines what metadata is to be translated and its new value.

```
/**
 * tables - aliases for table names
 * "tables": {
 * "table_name_1": "alias_1",
 * "table_name_2": "alias_2",
 * }
 * columns - aliases for column names in scope of table
 * "columns": {
 * "table_name_1": {
 * "column_name_1": "alias_for_column_name_1"
 * "column_name_2": "alias_for_column_name_2"
 * "column_name_3": "alias_for_column_name_3"
 * },
 * "table_name_2": {
 * "column_name_1": "alias_for_column_name_1"
 * "column_name_2": "alias_for_column_name_2"
 * "column_name_3": "alias_for_column_name_3"
 * }
 * }
 * titles - aliases for custom dimension titles, etc
 * "titles": {
 * "Custom title": "Alias",
 * "Sample Revenue": "Sample Profit"
 * },
 * formulas - aliases for custom formula titles
 * "formulas": {
 * "Custom title": "Alias"
 * },
 * hierarchies - aliases for hierarchy titles
 * "hierarchies": {
```

```

* "Custom title": "Alias"
* }
*/
var datasourceAliasing = {
  "tables": {
    "brand": "бренд",
    "category": "категория",
    "country": "страна",
    "commerce": "коммерция"
  },
  "formulas": {
    "Revenue": "Чистый доход"
  },
  "hierarchies": {
    "Category by Brand and Age Range": "Категория через
Бренд и Год"
  },
  "titles": {
    "CUSTOM REVENUE": "ДОХОДЫ"
  },
  "columns": {
    "Brand": {
      "Brand": "Бренд",
      "Brand ID": "Бренд Id"
    },
    "Category": {
      "Category": "Категория",
      "Category ID": "Id Категории"
    },
    "Widgets": {
      "Revenue": "Чистый доход",
    },
    "Commerce": {
      "Age Range": "Возрастные диапазоны",
      "Brand ID": "Id Бренда",
      "Category ID": "Id Категории",
      "Condition": "Состояние",
      "Cost": "Затраты",
      "Country ID": "Id Страны",
      "Date": "Дата",

```

```

"Gender": "Пол",
"Quantity": "Количество",
"Revenue": "Доходы",
"Visit ID": "Id Посетителя"
},
"Country": {
"Country ID": "Id Страны",
"Country": "Страна"
}
};
/**
 * Global object example
 * folder - aliases for folder names
 * "folder": {
 *   "folder_name_1": "alias_1",
 *   "folder_name_2": "alias_2",
 * }
 * dashboards- aliases for dashboard names
 * "dashboards- ": {
 *   "dashboard_name_1": "alias_1",
 *   "dashboard_name_2": "alias_2",
 * }*
 */
var globalAliasing = {
"folders": {
"New Folder": "Folder renamed",
"my folder": "my folder renamed",
"Any name": "любое название"
},
"dashboards": {
"Sample - Ecommerce (1)": "Коммерческий Dashboard",
"Sample - Healthcare": "Sample Healthcare renamed",
"Sample - Healthcare (1)": "Healthcare renamed",
"my dashboard": "my dashboard renamed"
},
};

```

# Server Settings

Sisense Administrators have access to the Data Settings page located in the Data Management section of the Admin page.

From the Data Settings page, Administrators can modify memory allocation, the number of simultaneous queries that can be handled, and query caching executed by Sisense.

**Note:** Sisense recommends consulting with Technical Support before changing the advanced settings.

You can modify the following settings:

Settings	Description
Default ElastiCube Data Location	The location of your ElastiCubes on your server. Sisense recommends storing ElastiCube data on a separate non-system drive.
Data Import Chunk Size (Rows)	The number of rows in every chunk imported when building an ElastiCube. This can be increased or decreased based on your build performance and network throughput.
R Server Address	The address of your R Server. By default, this feature is not activated. Click the toggle button to enable R server integration in Sisense and enter the IP address and port. Click Test to verify the connection to the R server. For more information about using R in Sisense, see <a href="#">Using R in Formulas</a> .
Max ElastiCube Memory Allocation	The percentage of RAM out of the server's total RAM on the machine that can be allocated to each ElastiCube. By default, the value is set to Dynamic, which automatically allocates and frees memory in each ElastiCube to ensure stability and performance.

Settings	Description
	Setting a static percentage will disable this mechanism.
Simultaneous Query Executions	The number of simultaneous queries that can be processed on a server across all ElastiCubes. This value is determined by half of the machine's vCPU cores.
Query Timeout (Seconds)	The maximum amount of seconds that must pass before a query times out.
ElastiCube Query Recycler	The ElastiCube Query Recycler caches similar and identical queries to the ElastiCube to reduce the amount of time to return a response. This reduces the amount of time it takes to process queries. This should be turned off for debugging purposes only.

# Setting Up a Custom Email Server

Sisense uses Mandrill by MailChimp to send automated emails such as dashboard reports or password reset emails. Custom email servers allow all our email transactions to be completely managed on your side.

Through the Sisense REST API, you can send an HTTP POST request to issue Sisense emails through your own custom email server instead of the default Sisense email server. The REST API exposes properties that you can configure for the custom email servers such as security, timeout, and logging properties.

## To set up a custom email server:

1. In the Sisense Web Application, click Admin, and then Rest API, followed by REST API Reference.
2. In version 1.0 of the REST API, select Settings, then POST `/settings/email_server`.
3. In the body of your call, define the values of the following keys:

Key	Value
enable	If True, allows access to the custom email server instead of the default Sisense email server. <b>Note:</b> This property is different from the enable property found in version .9 of the REST API POST <code>/settings/system</code> request. In .9, enable activates Sisense emails. In version 1.0, enable sends Sisense emails through your custom server.
host	The hostname or IP address to connect to. By default, this is 'localhost'.
port	The port to connect to (defaults to 25 or 465).

Key	Value
secure	<p>When you send an email, some email servers check if secure transport (TLS) is available for that address or domain.</p> <p>This value determines if the connection should use TLS (if True) or not (if False). The default value is false. The connection may be upgraded to TLS if the email server requires it.</p>
ignoreTLS	<p>If this is True and secure is False, TLS will not be used even if the email server attempts to upgrade the connection to TLS-secure.</p> <p>This field should be set to True when connecting to a non-secure SMTP server.</p>
requireTLS	<p>If this is True and secure is False, it forces the connection to be upgraded to TLS-secure even if the email server does not advertise support for TLS.</p>
connectionTimeout	<p>How many milliseconds to wait for the connection to establish.</p>
greetingTimeout	<p>How many milliseconds to wait for the greeting after connection is established.</p>
socketTimeout	<p>How many milliseconds of inactivity to allow.</p>
user	<p>The user name of the user.</p> <p>As non-secure or restricted SMTP servers do not require authentication, do not include this field when connecting to a non-secure SMTP server.</p>
pass	<p>The user's password for the email server.</p> <p>As non-secure or restricted SMTP servers do not require authentication, do not include this field when connecting to a non-secure SMTP server.</p>

Key	Value
logger	If set to True, Sisense logs events such as which server you connected to.If the value is not defined or is False, then nothing is logged.
debug	If set to True, then Sisense logs SMTP traffic. When False, Sisense logs only transaction events.Debug logs record everything and may increase in size quickly and impact performance. You should only set activate debug logging when troubleshooting a specific issue and then reset the logging level back to info.

4. Click **Run**.
5. After you have set up the email server, you need to change the sender email address to the domain of your custom email server so they are consistent. You can update the sender email through the POST 0.9 /branding endpoint in the SisenseREST API.

```
{
  "emails": {
    "senderEmail": "username@newdomain.com"
  }
}
```

6. Click **Run**.

## Limitations

- When authenticating an HTTP request, Sisense only supports the basic authentication through the Authorization method using plain text.
- No Errors are displayed in the UI.

# Troubleshooting a Custom Email Server

Sisense does not provide a UI for troubleshooting an issues with your custom email server. To troubleshoot any issue with your custom email server, activate logging through the **logger** and **debug** keys in the REST API. When **logger** and **debug** are set to True, Sisense generates logs that are saved in the following directory:

```
c:\Program Files\Sisense\PrismWeb\vnext\iisnode
```

After troubleshooting the issue, the debug logger should be disabled as the logs can grow in size rapidly and this may impact performance.

# System Configuration

The System Configuration settings let you define the web server's alias, the first day of the week, SSL security, and email settings.

## To update one or more System Configuration settings:

1. Click **Admin** in the top menu, and then **Settings** in the left menu.
2. Update one or more of the following settings:
  - **Alias:** Define the web server's alias or IP address. This is especially important when you have configured your IIS with a domain name (alias), but when you share a dashboard, the IP address and/or the incorrect port is sent instead. Reports will be sent from the alias entered in this field.
  - **First Day of the Week:** Select the first day of the week so all of your Sisense users' queries will be calculated according to the day you define. By default, the first day of the week is Monday when querying an ElastiCube. Sisense defines Monday as the first day according to ISO 8601. This means that when selecting filters such as Week 1, Week 2, etc., the first day of the week is not that defined by first day of that year, but rather the first Monday of the year. Each week starts on a Monday and ends on a Sunday for 52-53 full weeks in a year.  
**Note:** Changing the first day of week after using a dashboard or widget filter/drill with Weeks filter in use might lead to bad results. To resolve the issue, the Weeks filter should be removed and added again.
  - **First Month of Fiscal Calendar:** Select the first month of your company's fiscal calendar. By default, the first month of the fiscal calendar is January.  
**Note:** Changing the first month of the fiscal calendar after using a dashboard or widget filter/drill with the related date fields is in use might lead to bad results. To resolve the issue the date filter should be removed and added again.
  - **Default Language:** Select the language to be displayed in the Sisense Web Application. If you have defined a custom translation, the name of the language as defined in the settings.json is displayed in this list. Select it to display your custom translation. For more information about custom translations, [click here](#).

- **SSL:** Enable SSL if you have configured your web environment with SSL security (https) and certificates. This will enable https in the URLs of shared dashboards. [Click here](#) to learn more about setting up SSL.
  - **Send Emails:** You can disable sending emails. This will disable emails sent when sharing dashboards as well as any scheduled email reports.
  - **Widget Narrative:** You can add narrated text to your dashboard's widgets. For more information about this feature, see [Sisense Narratives](#).
  - **Field Suggestions:** You can toggle on field suggestions. Field suggestions are displayed to Dashboard Designers to help them add the relevant fields to widgets quickly and easily. For more information, see [Field Suggestions](#).
3. Click **Save** to update your system settings.

# System Settings

Sisense Administrators can perform a variety of tasks related to implementing and maintaining an implementation of Sisense.

The following topics describe system-related tasks that Sisense Administrators are responsible for handling:

- [System Configuration](#)
- [Internationalization and Localization](#)
- [Managing Plug-ins](#)
- [Setting Up a Custom Email Server](#)
- [Sisense Narratives](#)

# Introduction to SSO

Single Sign-On (SSO) is a mechanism that allows a system to authenticate users and subsequently tell Sisense that the user has been authenticated. The user is then allowed to access Sisense without being prompted to enter separate login credentials.

The SSO security mechanism allows Sisense to trust the login requests it gets from your corporate authentication system, and grant access to the users that have been authenticated by it. An SSO session begins when the authenticated user requests a secured resource from Sisense while logged into your site or application. The user's browser sends an HTTP request to Sisense that includes a cookie which contains session and authentication information. This information is then used for session validation.

Users who already have Sisense accounts can continue to access Sisense through the Sisense Login page with their current accounts. To prevent users from directly logging in to Sisense instead of your login page, your Sisense administrator can change the passwords of your current users forcing them to log in with your company's credentials in your company's login page.

Sisense recommends that administrators always keep a Sisense password, so that the administrator can access Sisense in case the SSO server is not available.

Sisense SSO supports two SSO protocols for securing the exchange of user authentication data, JSON Web Token (JWT) and SSO via SAML 2.0 (Security Assertion Markup Language). For more information, see [SSO via JWT](#) or [SSO via SAML 2.0](#).

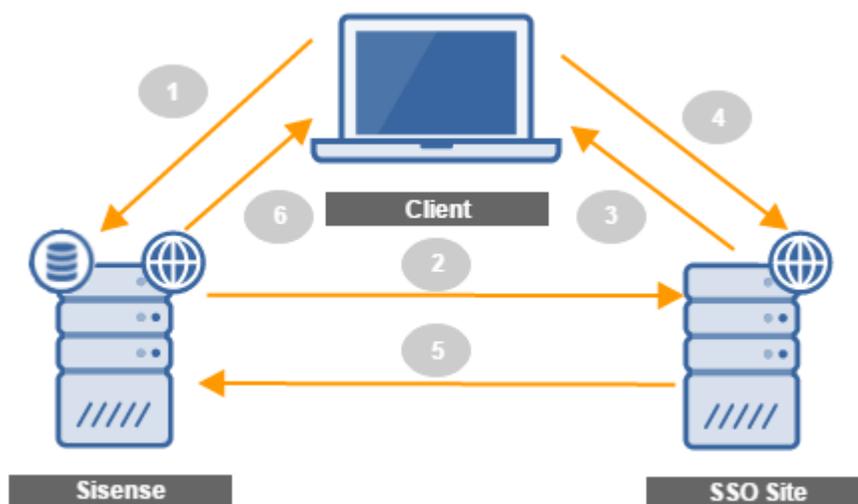
# SSO via JWT

JWT is a technique that can be used for single sign-on (SSO) between your site and Sisense. JWT is a token that represents your users credentials wrapped in a single query string. In addition, Sisense uses the `jti` parameter (see below), which adds a unique ID to the token that prevents the token from being used more than once, thus preventing attacks on the system (also known as replay attacks). The Sisense SSO via JWT authentication flow is explained in the following diagram.

**Note:** See the [Developers Community](#) for a tutorial that shows how to implement SSO via JWT.

## SSO Authentication Flow

The following is a diagram of the SSO authentication flow from your site or application to Sisense.



1. Your user requests a resource from Sisense, typically a dashboard.

2. Sisense recognizes that no authenticated cookie is present. If you have enabled SSO in Sisense, the SSO handler redirects the user to your Remote Login URL defined in the Sisense Web Application.

**Single Sign On Configuration**

Method  JWT  SAML 2.0 (i)

Remote Login URL  (i)

Remote Logout URL  (i)

Shared Secret  (i)

3. Your user is challenged to authenticate their account.
4. Your Remote Login application authenticates your user and generates a JWT (JSON Web Token).
5. You redirect the user back to Sisense with the encoded JWT in a query string. Sisense sets a cookie that authenticates the user's session until they end it or you log them out via the Sisense REST API. For more information see Logging Users Out.
6. Sisense provides the authenticated user with the request resource.

A common scenario that illustrates SSO is when an unauthenticated user navigates to your site in which Sisense is embedded via an iFrame. Sisense redirects this user to your SSO script. Your script authenticates the user through your login process and builds a JWT request with all the relevant credentials wrapped together. You then redirect the customer back to Sisense with the JWT payload. Sisense then decodes the user details from the JWT payload and then grants the user a session.

## Configuring SSO in Sisense

While SSO is highly customizable, there are generally four steps you should complete when configuring SSO:

**Note:** Configuring SSO requires technical expertise and should be conducted by an administrator or developer with SSO experience.

1. **Enabling SSO in Sisense:** Through the Sisense Web Application, an administrator can enable SSO in Sisense and define the relevant Login and Logout URLs.
2. **Creating a JWT:** After you authenticate a user, you generate a JWT with the user's credentials to Sisense, so Sisense knows this user is allowed to access resources from Sisense through your site.
3. **Configure Sisense as a sub-domain:** When authenticating users, you should configure SSO as a sub-domain.
4. **Logging Users Out:** A user can access Sisense so long as a session is maintained. To end a session, the user's cookie that Sisense provides must be deleted. To delete this cookie, you can use the Sisense REST API.

## Enabling SSO in Sisense

For Sisense to recognize that your users should be authenticated through SSO, you must enable SSO in the Sisense Web Application. In the SSO menu of the Admin page of the Sisense Web Application, you define the URL where Sisense redirects users to authenticate on your side and where Sisense redirects users after they log out from Sisense.

When you access the SSO menu of the Admin page, Sisense displays the Shared Secret key. The Shared Key is a JWT encryption public key used to encrypt the JWT payload. It is generated once when the SSO configuration is saved. You include this key in the JWT payload when redirecting the user back to Sisense after authenticating them on your side.

### To access and set up SSO:

1. Log into Sisense.
2. Select Admin screen and click **Single Sign On** in the left menu.
3. Fill in the following SSO configuration fields:
  - **Remote Login URL:** This is the URL that Sisense will invoke to attempt remote authentication. In that endpoint the participating application user authentication script is triggered and the JWT payload is generated.
  - **Remote Logout URL:** This is the URL that users will be redirected to after they log out from Sisense (i.e. the participating application's home page).
4. Click **Save**.

# Creating a JWT

Your script builds a JWT request that contains the user data.

The table below provides a list and descriptions of the attributes your JWT should contain.

In addition, several samples are provided below in various languages.

Attribute	Mandatory	Description
iat	Yes	Issued at the time the token was generated. This is used to help ensure that a given token gets used shortly after it is generated. The value must be the number of seconds since UNIX epoch. Sisense allows up to five minutes clock skew. <b>Note:</b> The date must be an integer and not a float.
sub	Yes	Email of the user being signed in, used to uniquely identify the user in Sisense. If the user does not exist in Sisense, it will be created with default viewer privileges.
jti	Yes*	A unique string added to the token that is used to prevent replay attacks, by making sure the token is used only once.

\* You can set this attribute as optional in the Sisense REST API v1.0 through the POST settings/SSO endpoint.

SSO Code Samples

[C#](#)

[Java](#)

[Javascript](#)

[PHP](#)

[Ruby](#)

[Python](#)

## return\_to URL

When Sisense redirects a user to your login script, Sisense passes a `return_to` parameter in the URL. This parameter contains the page that Sisense will return the user to after the authentication succeeds. For example:

A customer visits your site opens a dashboard embedded through an iFrame.

Sisense recognizes that the user is not authenticated.

Sisense redirects the user to:

```
https://yourcompany.com/sisense/sso?return_to=https://yourcompany.sisense.com/dashboards/
```

All your script needs to do, is take the `return_to` value from the invoked URL and pass it back to Sisense when submitting the JWT token. In other words, upon authentication on your side, your script redirects the user to:

```
https://yourcompany.com/access/jwt?jwt=payload&return_to=https://yourcompany.sisense.com/dashboards/
```

## Configuring Sisense as a Sub-Domain with SSO

To authenticate your users locally and allow them to access Sisense, your first step should be to configure Sisense as a sub-domain of your web application and embed Sisense into your web application with SSO.

**Note:** Sisense also works when embedded in cross-domain iFrames.

**To configure Sisense as a Sub-Domain:**

1. In IIS Management Console, add your website to IIS.



2. In the **Add Website** window, in the field **Site Name**, enter the site name.  
**Note:** You can not leverage SSO when you enter a DNS address with an underscore “\_” with Internet Explorer.

A screenshot of the 'Add Website' dialog box. The dialog has a title bar with a question mark and a close button. It contains several sections: 'Site name:' with an empty text box; 'Application pool:' with a dropdown menu showing 'DefaultAppPool' and a 'Select...' button; 'Content Directory' section with 'Physical path:' and an empty text box followed by a browse button (...); 'Pass-through authentication' section with 'Connect as...' and 'Test Settings...' buttons; 'Binding' section with 'Type:' dropdown set to 'http', 'IP address:' dropdown set to 'All Unassigned', and 'Port:' text box set to '80'; 'Host name:' with an empty text box and an example: 'Example: www.contoso.com or marketing.contoso.com'; and a checkbox labeled 'Start Website immediately' which is checked. At the bottom right are 'OK' and 'Cancel' buttons.

3. In **Physical path**, enter the subdomain directory.

4. Under the Binding area, in **Host name**, change the existing SisenseWeb site binding to use **sisense.example.website.com** as host name on **port 80**.
5. Open the file **C:\Windows\System32\drivers\etc\hosts** and add mapping for the sites:  
 192.168.5.148 sisense.examplewebsite.com  
 192.168.5.148 example.website.com
6. Sign in to the Sisense Web Application at **sisense.examplewebsite.com** and configure the SSO as pictured below:
7. Place the following SSO script in the server location corresponding to the Remote Login URL in the server's root directory. The SSO script can be implemented in any server-side language. This example uses Python. Example code has been attached for C# and Python.  
[Python script example](#)  
[C# script example](#)
8. In index.html from examplewebsite.com, the IFrame source is the Sisense dashboard URL.
9. index.html from examplewebsite.com:

```
<html>
<head>
<title>Example Website</title>
</head>
<body>
<p><b>examplewebsite.com</b> - <b>SSO</b> login with
embedded dashboard from
<b>sisense.examplewebsite.com</b></p>
<iframe width="100%" height="100%"
src='http://sisense.examplewebsite.com/app/main#/dashboards/53b29843751b655443000018?embed=true' />
</body>
</html>
```

Navigate to **examplewebsite.com** and you should see the specific dashboard you embedded.

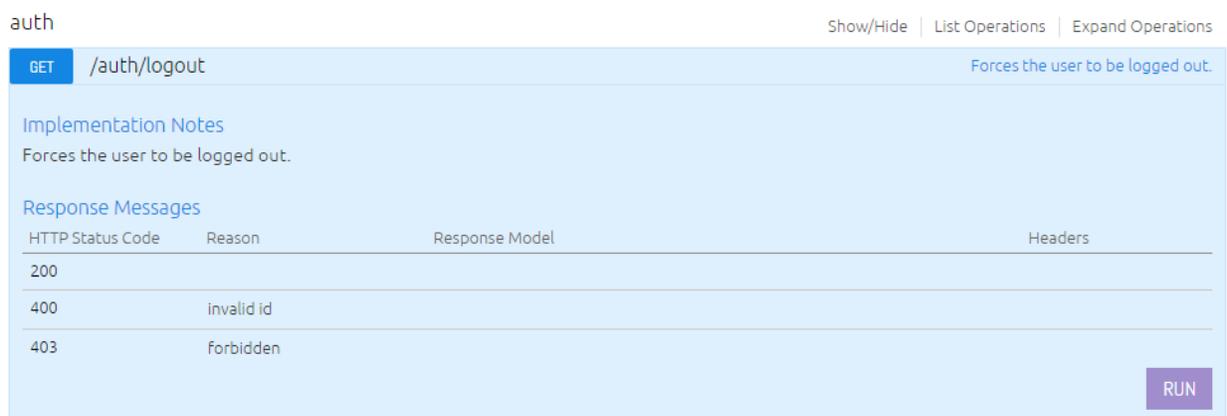
# Logging Users Out

When a user is logged in, anyone using that browser can access the session, or users may encounter an issue where they remain logged in until the Sisense cookie is cleared.

Users are logged out when the session ends. A session ends when the user closes their browser or according to the value of the attribute exp you send in the JWT payload.

You can log the user out through the Sisense REST API.

To manually log a user out, access version .9 of the REST API. Through the Auth method, you can issue a get request to log out specific users.



The screenshot shows the REST API interface for the endpoint `auth` with the method `GET` and path `/auth/logout`. The description is "Forces the user to be logged out." Below this, there is a table of response messages:

HTTP Status Code	Reason	Response Model	Headers
200			
400	invalid id		
403	Forbidden		

A `RUN` button is visible in the bottom right corner of the interface.

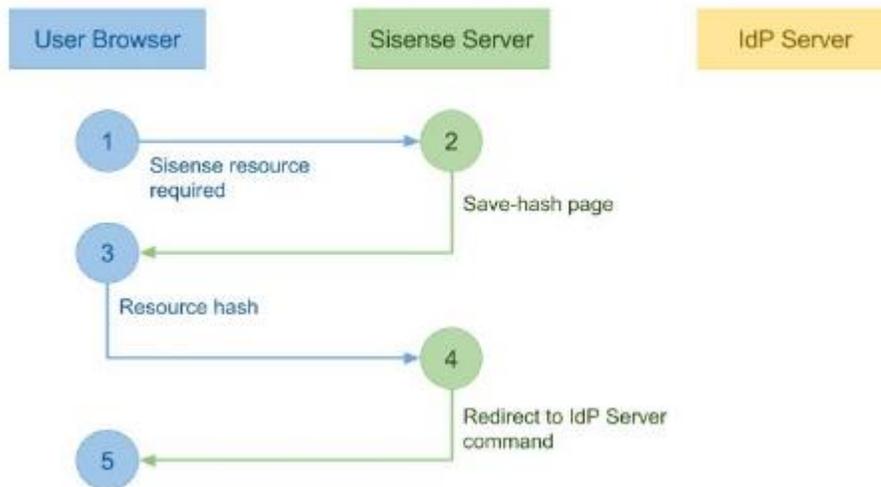
While the logout REST API can delete the SSO authentication cookie, it can only delete it when the call is made from within the Sisense domain. Scripts on different pages can access each other only if the pages that executed them are at locations with the same protocol.

If you have embedded Sisense in an iFrame and you want to log out the user from your application and Sisense, you can use the `window.postMessage` method to call the logout when the users asks to logout from your application. This method overcomes any cross-origin communication limitations. Sisense has created a plugin that implements a listener, which calls the Logout API when the `postMessage` method() is called. For more information, click here.

# SSO via SAML 2.0

The Sisense SAML authentication process is based on the SAML 2.0 protocol, and is explained in the following diagrams:

## Sisense SAML auth workflow (1st phase)



## Phase 1

The first phase of this process begins when a user requests a resource from Sisense via their browser (1). The browser generates a resource request to the Sisense server. The server processes this request (2), and for unauthenticated users, returns a special save-hash page to save the requested hash data.

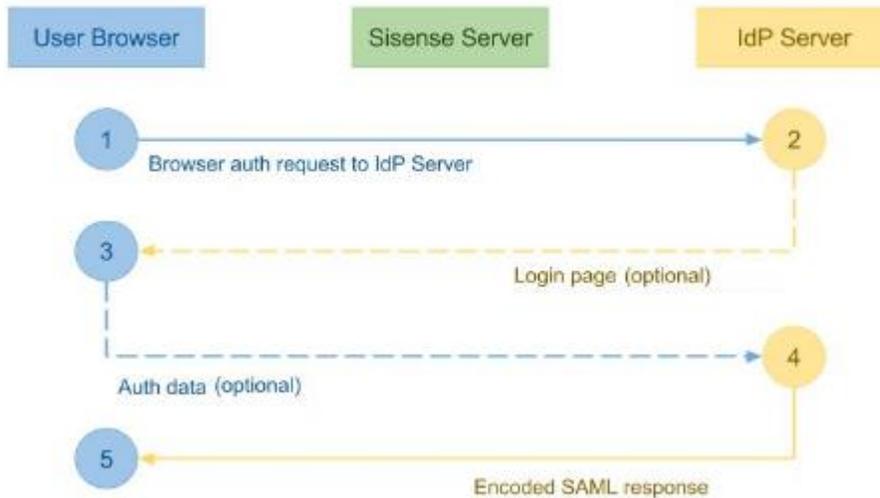
After the browser loads the save-hash page from the Sisense server, it runs a script, which loads the requested resource hash from the URL and sends it to the Sisense Server (3).

At the end of this phase, the Sisense Server converts the requested resource URL and sends it to the browser as a part of the redirect to IdP Server command (4).

## Phase 2

The second phase of the authentication process starts after the browser retrieves the redirect command from the Sisense Server, and sends the authentication request to the IdP Server (1) as described in the diagram below.

### Sisense SAML auth workflow (2nd phase)



The base URL for this request is taken from the Remote Login URL field of the Single Sign On Admin page in the Sisense Web Application.

The requested resource address is passed to this request as a RelayState parameter. All other data is provided as a SAMLRequest parameter.

The next step of this phase depends on the user authentication state and the IdP implementation.

If the current user isn't logged in as an IdP user, the IdP server redirects the browser to your Login page (2), where the user enters their IdP credentials. After logging in, the IdP Server sends the browser the Auto-Sign-In page with encoded data about the currently logged-in user (4).

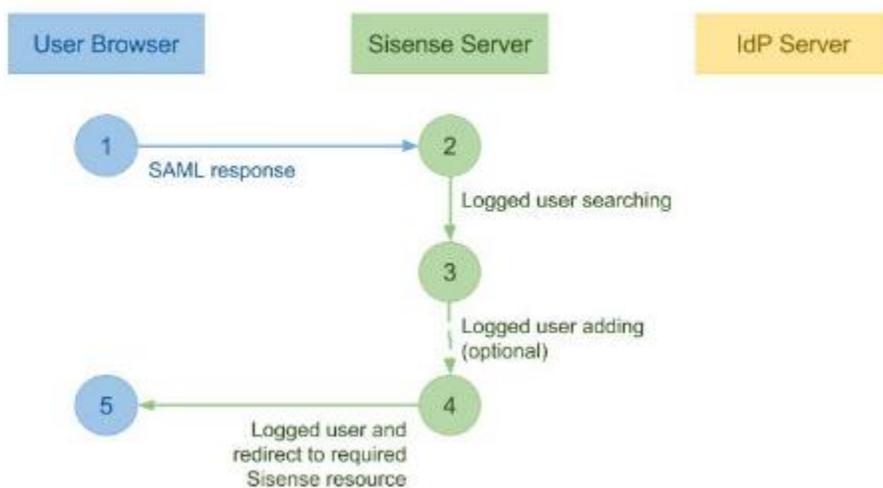
**Note:** Two-factor authentication for Sisense is supported for SSO providers that support two-factor authentication.

After this document is loaded in the user's browser, it runs a script which creates the POST-query to the Sisense Server API, and passes the SAML response to this query (1). At this time, the third phase of the authentication process begins.

## Phase 3

The Sisense Server handles the POST-query (2), decrypts it with the certificate specified in the Public X.509 Certificate field defined in the Single Sign On Admin page. The Sisense Server uses the decrypted User ID as a key to locate the Sisense user in the internal database.

### Sisense SAML auth workflow (3rd phase)



If a user is not found in the system, Sisense creates a new user (3).

When Sisense creates a new user, Sisense analyzes the "memberOf" field to locate one or more groups related to the logged-in user.

If the "memberOf" field is empty, the user is assigned the role: "Viewer". If the "memberOf" field contains one or more groups, and the groups were previously defined in Sisense, the newly created users' Role is taken from the groups' default role. When multiple Sisense Groups are found, the user is assigned the role with the maximum privileges.

After a user is created in the system, an administrator can modify the user role, if needed.

Below, is an example of an SAML XML where the "Test" group is specified:

```
<saml:Attribute
  NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-
format:basic
  Name="memberOf">
  <saml:AttributeValue
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-
instance"
    xsi:type="xs:string">
    Test
  </saml:AttributeValue>
</saml:Attribute>
```

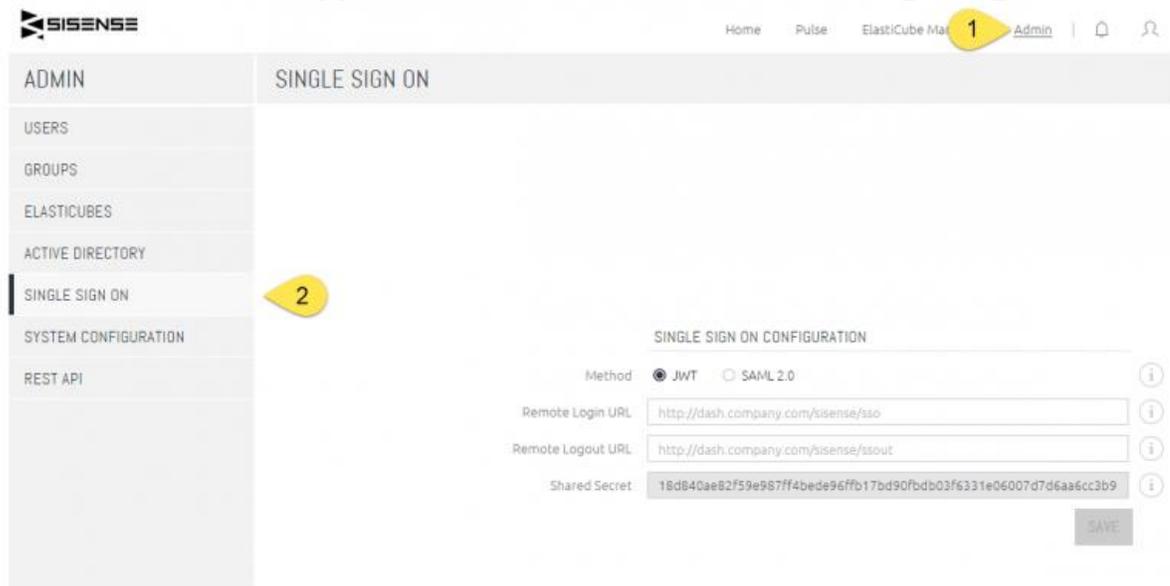
After the user is located or created, a new session is initialized, and the Sisense Server redirects the user browser to the requested resource (4).

## Enabling SSO via SAML in Sisense

After you have configured your SAML server, sign in to Sisense as an Administrator and follow the instructions below.

**To enable SAML in Sisense:**

1. In the Sisense Web Application, click **Admin** and select **Single Sign On**.



2. In the Single Sign On page, select **SAML 2.0**.
3. In the **Remote Login URL** field, enter the SAML Login endpoint. Sisense redirects the user to this field when they sign in. This value should be provided by the IdP Service.
4. In the **Remote Logout URL** field, enter the SAML Logout endpoint. Sisense redirects the user to this field when they sign out. This value should be provided by the IdP Service.
5. In the **Public X.509 Certificate** field, enter your public key for your SAML configuration. This value should be provided by the IdP Service.
6. Click **Save**. SSO via SAML 2.0 is configured.

## Sisense Default Role Set-Up

When an authenticated user is not found in the Sisense database, a new account is created. The user role is specified based on the user group/groups default role.

**To define a group's role:**

1. In the Sisense Web Application, click **Admin** and select **Groups**.

The screenshot shows the Sisense Admin interface. The top navigation bar includes links for Home, Pulse, and ElastiCube M. The 'Admin' link is highlighted with a yellow callout '1'. The left sidebar contains a menu with options: ADMIN, USERS, GROUPS (highlighted with a yellow callout '2'), ELASTICUBES, ACTIVE DIRECTORY, SINGLE SIGN ON, SYSTEM CONFIGURATION, and REST API. The main content area is titled 'GROUPS' and features an 'ADD GROUP' button with a plus sign. Below the button is a table with the following data:

<input type="checkbox"/>	Group Name	Number of Users	Origin	
<input type="checkbox"/>	Admins	2	Sisense	
<input type="checkbox"/>	Everyone	6	Sisense	
<input type="checkbox"/>	Sales	2	Sisense	

At the bottom of the page, it says 'Page 1 of 1' and '3 Groups Total'.

2. Click **Add Group**. The Create a New Group window is displayed.

3. In the **Create a New Group** window, select the default role of the group.

Create a New Group

Group Name

|Type in a descriptive name for the new group.

Group Users

+ Add Group Users

No members added yet

Choose default members role:  Viewer  Designer

CANCEL

4. Click **Save**.

**Note:** Changes to the group's default role are applied when new users are created, and do not affect existing users. After a user is created in the system, an administrator can modify the user role, if needed.

# Customizing User Roles

This section describes how to customize the permissions given to the Sisense user roles (designer and viewer), using Sisense's [REST API](#). At the bottom of the page, you can find a video tutorial demonstrating the following procedure.

[Click here](#) to learn more about each role type.

**Note:** The API uses a previous terminology for Designers and Viewers. When entering the role name into the operation parameters, use Contributor (for Designer), and Consumer (for Viewer).

## Customizing Roles

You can customize the three default roles with Sisense's REST API, by defining which permissions to enable or disable for each role. For example, you can prevent a user from drilling into widgets, or exporting data to a CSV file. On the other hand, you may want to let viewers change the widget type and format the style.

The following table lists all the actions that you can customize using the API. True and False indicate values that you can change. Never indicates values that you cannot change.

## Dashboards

Dashboards	Admin	Designer	Viewer	Description
create	True	True	Never	Create a new dashboard.
duplicate	True	True	False	Duplicate an existing dashboard.
toggle_edit_mode	True	True	False	Select between Layout and View modes in the dashboard. Layout mode is the default mode in which you can add new

Dashboards	Admin	Designer	Viewer	Description
				widgets and rearrange them on the dashboard. In View mode, the dashboard cannot be rearranged or edited.
edit_script	True	True	Never	Edit a dashboard using the Sisense JavaScript API.
export_dash	True	True	Never	Export a dashboard as a .dash file that can be imported back into Sisense.
export_jpg	True	True	True	Download a widget as an image in PNG format. The downloaded image represents the current view.
export_pdf	True	True	True	Export the dashboard to PDF format. The exported file includes the current view.
restore	True	True	True	Revert changes made in the dashboard to the latest dashboard that was shared.
import	True	True	Never	Import a .dash file into the Sisense Web Application.
select_palette	True	True	Never	Change the selected color palette.
<b>Filters</b>				
create	True	True	False	Create a new filter.
on_off	True	True	True	Switch filters on or off.
toggle_expansion	True	True	True	Enable the expansion of filter settings to see more filter parameters.
modify_type	True	True	False	Include in the filter options 'Ranking' and

Dashboards	Admin	Designer	Viewer	Description
				'Starred', in addition to 'List' and 'Text'.

## Widgets

Widgets	Admin	Designer	Viewer	Description
create	True	True	Never	Create a new widget.
delete	True	True	Never	Delete a widget.
rename	True	True	False	Rename a widget.
duplicate	True	True	Never	Duplicate a widget.
edit_script	True	True	Never	Edit a widget using the JavaScript API.
change_type	True	True	False	Change the chart type.
export_csv	True	True	True	Export a widget's data to a CSV file.
export_png	True	True	True	Download a widget as an image in the PNG format. The downloaded image represents the current view.
modify_selection_mode	True	True	False	Enable the option to change the filter behavior between 'Slice' (only filtered items are included in the chart) and 'Highlight' (filters items are highlighted).
drill_to_anywhere	True	True	True	Enable the option to drill down into a field to get an in-depth view of a selected value. If False, the user can only drill down with predefined drill hierarchies (if defined).
<b>Items</b>				
reorder	True	True	False	Reorder the fields and values in the data

Widgets	Admin	Designer	Viewer	Description
				panel and their representation in the chart.
modify_type	True	True	Never	Enables the option to edit values and categories in the data panel.
on_off	True	True	True	Enable the option to turn fields on or off in the widget.
Filters				
on_off	True	True	False	Switch filters on or off.
toggle_expansion	True	True	True	Enable the expansion of filter settings to display more filter parameters.
modify_type	True	True	False	Include in the filter options 'Ranking' and 'Starred', in addition to 'List' and 'Text'.
use_starred	True	True	False	Enables the option to star (bookmark as favorite) a filter selection.

## ElastiCubes

This object is not displayed in the Sisense REST API, however, you can still toggle these features through the **POST** /roles/{idOrName}/manifest/{path} endpoint. The value of path should be in the format of manage/elasticubes/[value] where [value] is the feature listed in the table below, add or schedule\_build.

Parameter	Value	Description	Parameter Type	Data Type
idOrName	dataAdmin	The name or ID of the role.	path	string
path	manage/elasticubes/add	The path to the manifest. Enter the path to the setting that you want to change as it appears in the hierarchy. Do not include the setting in the path. Example: To change the 'Create' setting for dashboard filters, enter dashboard/filters/. For creating dashboards, enter dashboards. See the reference table above for more settings.	path	string

ElastiCubes	Admin	Data Admin	Data Designer	Designer	Viewer

<b>ElastiCubes</b>	<b>Admin</b>	<b>Data Admin</b>	<b>Data Designer</b>	<b>Designer</b>	<b>Viewer</b>
<b>Add</b> Allows the user to add new ElastiCubes.	True	True	True	False	False
<b>schedule_build</b> Allows the user to schedule builds for an ElastiCube.	True	True	True	False	False
<b>sql_api</b> Allows the user to access the SQL API ( <a href="#">SQL Runner</a> ) feature.	True	False	False	False	False

**To change permissions for a user role:**

1. In the Sisense Web Application, click **Admin** at the top right of the screen, and then **REST API** in the left menu.
2. Click **REST API Reference** to view the API documentation.
3. In the API Documentation, select version **0.9**.
4. Click on **/roles** to access the different API operations.
5. Find the operation that you require (see examples below), and type in the required parameters in the operation. You can click on **Model Schema** to see and add example code for the parameters.

When changing a permission for a role, you must use the **Path** parameter, which defines the permission you want to update and its relative path in the hierarchy that appears in the table above. For example, to change the 'Create' setting for dashboard filters, enter `dashboards/filters/`. For creating dashboards, enter `/dashboards`. See the reference table above for more

settings.

<b>Dashboards</b>	<b>Admin</b>	<b>Designer</b>
create	true	true
duplicate	true	true
toggle_edit_mode	true	true
edit_script	true	true
export_dash	true	true
export_jpeg	true	true
export_pdf	true	true
restore	true	true
import	true	true
select_palette	true	true
<b>Filters</b>		
create	true	true
on_off	true	true
toggle_expansion	true	true
path	dashboards/filters/	
visibility_type	true	true

6. Click **RUN** to apply the changes.

Here are some examples of what you can do.

### See all permissions for all roles

Method: **Get /roles**

#### Parameters

- **includeManifest** – Use the default setting (true)
- **compiledRoles** – Use the default setting (true)

#### Example API Call



## Restore a specific setting for a specific role

Method: **DELETE** /roles/{idOrName}/manifest/{path}

### Parameters

- **idOrName** – The role name can be one of the following: super, dataDesigner, dataAdmin, admin, contributor, consumer.
- **Path** – Enter the full path of the setting you want to change, as it appears in the hierarchy. Example: dashboards/filters/modify\_type. See the reference table above for more settings.

### Example API Call

```
/api/roles/viewer/manifest/dashboards/filters/modify_type
```

## Restore all settings for a specific role

Method: **DELETE** /roles/{idOrName}/manifest/{path}

### Parameters

- **idOrName** – The role name can be one of the following: super, dataDesigner, dataAdmin, admin, contributor, consumer.
- **Path** – / Enter just a slash to restore all settings.

### Example API Call

```
/api/roles/viewer/manifest/
```

# Integrating Active Directory

Sisense supports integrating Active Directory users and groups with your current Sisense users, so that you can share dashboards and email reports with any of your users. This works in parallel with Sisense's existing User Management, so you can continue to add users and user groups that are not part of the Active Directory domain.

To enable Sisense to identify your Active Directory users, you import your Active Directory users and groups into Sisense. When you import an Active Directory user, Sisense automatically creates a Sisense user. When you import an Active Directory group, Sisense does not automatically create a Sisense user for each member. The first time a member of an Active Directory group logs into Sisense, then Sisense automatically creates a Sisense user. Active Directory users that do not log in, are not assigned a Sisense user until their first login.

This section describes how to configure Active Directory in Sisense, and how to import users and user groups from Active Directory.

## Adding an Active Directory Domain to Sisense

From the Active Directory page, you can add Active Directory domains to your Sisense configuration. Each Active Directory domain you add to your Sisense configuration is added to the Active Directory list displayed in the Active Directory page. The first time you open the Active Directory page, the Add Domain link is displayed, which you can click to begin adding Active Directory domains.

After you have added your first Active Directory domain, it is displayed in the Active Directory list.

From the Active Directory list, you can add additional Active Directory domains. Each Active Directory domain you add to Sisense enables you to share your dashboards and reports with users from any of the domains as well as Sisense users not part an Active Directory.

You can filter the list of the Active Directory domains by entering the name of the relevant domain in the Search field.

### To add an Active Directory:

1. Click **Admin** and select the **Active Directory** tab on the left.
2. In the Active Directory page, click **Add Domain**.
3. In the **Add Domain** window, fill in the Active Directory configuration fields to point to your Active Directory instance:
  - **Name:** The name of your Active Directory domain. This name is used to locate your Active Directory domain in Sisense when you have multiple domains.
  - **Connection String:** The full path to your Active Directory domain. Example: **ldap://dc.domain.com**
  - **Base DN:** The Active Directory location from where a server searches for users. Enter the domain by its distinguished name (DN) with its domain components (DC) connected by commas. For example: For the domain domain.sisense.test, use **DC=domain,DC=sisense,DC=test**.  
**Note:** If you have a hierarchy of user groups in your organization, you can define the hierarchy in your string as follows:
    - If your domain is domain.sisense.com/OU1/OU2, then use **OU=OU2, OU=OU1,DC=domain,DC=sisense,DC=com**. Only users and user groups from OU2 will be available. Users/groups under OU1 will not be available.
  - **Username:** Enter the username with domain. Example: **domain\username**
  - **Password:** Enter the password for your Active Directory instance.
4. Click **Test** to verify that your configuration successfully connects to your Active Directory domain.

5. In the **Sync Data** field, select the frequency that Sisense automatically synchronizes with the users and groups in this Active Directory domain.
6. Click **Save**.

## Editing Active Directory Domain Settings

### To edit an Active Directory domain:

1. From the Active Directory list, select  for the Active Directory settings you want to edit. The Edit Active Directory window is displayed.
2. In the Edit Active Directory window, edit the relevant fields. See **Adding an Active Directory in Sisense** for more information about each field.
3. Click **Save**.

## Disabling an Active Directory Domain

Supports allows you to disable Active Directory domains in Sisense.

An Active Directory domain that has been disabled is still displayed in the Active Directory list, however, users from this domain will not be able to log in to Sisense, and you will not be able to share dashboards or reports with them.

Disabled Active Directory domains have the status Disabled in the Status column in the Active Directory list.

### To disable an Active Directory domain:

- From the Active Directory list, select  > **Disable** for the Active Directory domain you want to disable. The Active Directory domain is disabled.  
You can enable the Active Directory domain by selecting  > **Enable**.

## Deleting an Active Directory

You can delete an Active Directory domain from your configuration. When you delete an Active Directory domain, the users included in that domain are no longer displayed when sharing dashboards or reports. Deleted Active Directory domains cannot be restored, however you can add the Active Directory domain again from the Active Directory page.

### To delete an Active Directory domain:

1. From the Active Directory list, select  for the Active Directory settings you want to delete. The Delete Active Directory window is displayed.
2. Click **Delete**. The Active Directory domain settings are removed from Sisense.

## Adding a User from Active Directory

Once a connection with Active Directory has been established, an additional **Import from External Server** option is added to the **Add Users** button in the User tab.

The user's username and email will be added from Active Directory.

**Note:** You can add a user from Active Directory, even if you have not imported that user's user group from Active Directory.

See [Managing Dashboard Users](#) for general information about working with users/user groups in Sisense.

## Adding a User Group from Active Directory

Once a connection with Active Directory has been established, an additional **IMPORT FROM ACTIVE DIRECTORY** option is added to the **ADD GROUP** button in the **GROUPS** tab.

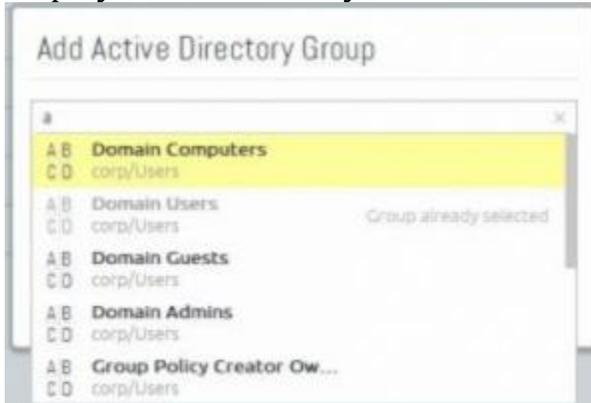
The following procedure describes how to import a list of users from a group in Active Directory to Sisense.

Users that are added in this way only appear in the Sisense users list after each user logs in to Sisense. In this way, your Sisense license slots are only consumed by users who actually log in to Sisense.

### To add a user group from Active Directory:

1. Click **Admin** and select the **Groups** tab on the left.
2. Click the **Add Group** button and then select the **From Active Directory** option from the menu.

3. In the **Search Active Directory Group** field start typing in the name of the desired Active Directory Group. An auto-complete dropdown menu is displayed from which you can select a user group.



If an Active Directory user group has already been imported into Sisense, the words **Group already selected** appear next to it, as shown above. You can add multiple Active Directory groups, by selecting one group after another. The groups appear in this window as follows:

4. Select the default Role for the members of this user group: Viewer or Designer. It will be assigned to new users from this group, when the users are created. Subsequently, user's role can be changed in the Sisense Users list.

Changes to a group's default role will affect users created after that time, but will not affect existing users.

5. Click **Add**.

# License Utilization

Sisense provides your licensing information in the Admin page of the Sisense Web Application.

In the License Utilization section, you can see when your license expires, how many user licenses are in use for your account, and additional features that are supported in your account.

The License Utilization section includes the following information:

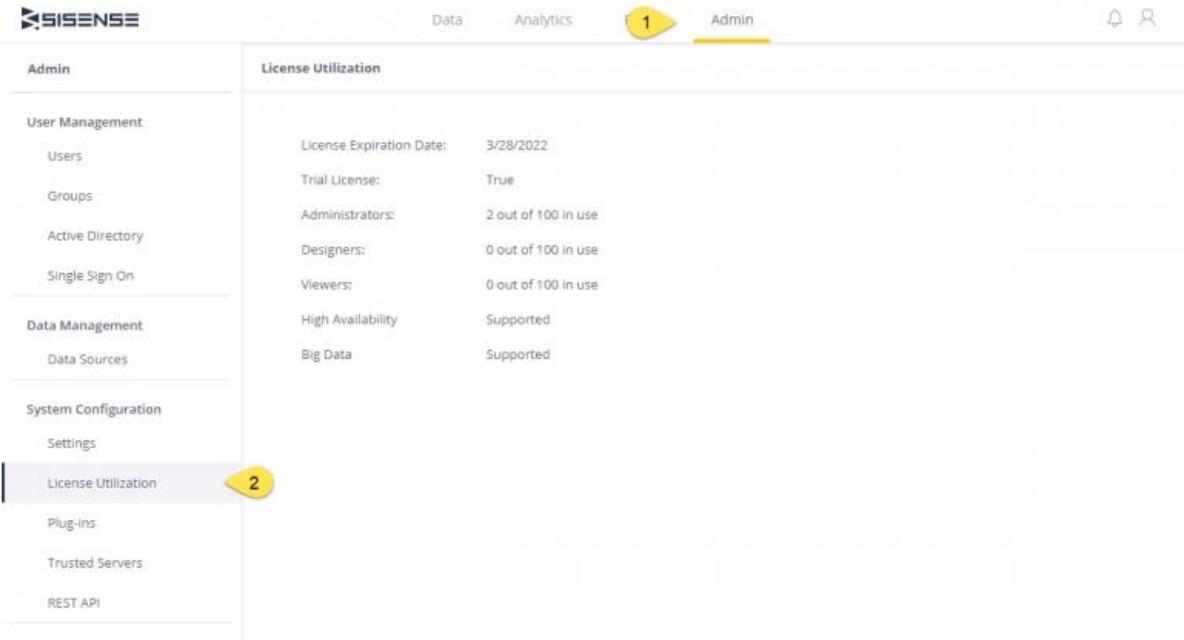
Details	Description
License Expiration Date	The date when your Sisense license expires in the format of Month/Day/Year.
Trial License	Indicates if your license is for a trial or full version of Sisense. True indicates that you are using a trial license.
Administrators	The number of Administrator licenses in use out of the number of available licenses in your account. Administrator licenses includes Admin, Sys.Admin, and Data Admin.
Designers	The number of Designer licenses in use out of the number of available licenses in your account. Designer licenses include Designers and Data Designers.
Viewers	The number of Viewer licenses in use out of the number of available licenses in your account.
High Availability	Indicates if your account supports high availability.
Big Data	Indicates if your account supports big data. Big data as defined by Sisense is typically considered to be data sets above ~1 billion

Details	Description
	rows.

**To access the License Utilization section:**

**Note:** This section is accessible to System Administrators only.

- In the Sisense Web Application, click **Admin** and select License Utilization in the left menu.



# Managing Dashboard Users

This section describes how to add and manage Sisense users that can access Sisense dashboards.

You can add users using the following methods:

- Add new users by entering them one by one or in bulk
- Add new users from your Active Directory database. [Click here](#) to learn how to configure your Active Directory in Sisense

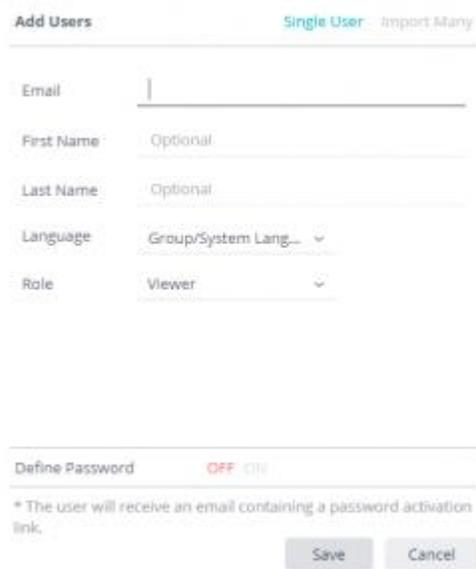
**To access the user management settings:**

- Click **Admin** and select the **Users** tab on the left. By default, the only available user after installing Sisense is the user that installed Sisense. By default, this user is an System Administrator.

## Adding a New User

**To add a new user:**

1. Click **+ Users**. If a connection with Active Directory also exists, an additional Import Many option is available. Click it to select a user from Active Directory.



The screenshot shows the 'Add Users' form in the Sisense user management interface. At the top, there are three tabs: 'Add Users' (selected), 'Single User', and 'Import Many'. Below the tabs, there are several input fields: 'Email' (a text input field), 'First Name' (a text input field with 'Optional' text below it), 'Last Name' (a text input field with 'Optional' text below it), 'Language' (a dropdown menu with 'Group/System Lang...' selected), and 'Role' (a dropdown menu with 'Viewer' selected). At the bottom of the form, there is a 'Define Password' section with a toggle switch set to 'OFF'. Below this, there is a note: '\* The user will receive an email containing a password activation link.' At the very bottom, there are two buttons: 'Save' and 'Cancel'.

2. Fill in the user's details in the window.
3. (Optional) The **Language** for your account is the language set by an Administrator for your group or server. From the Language list, you can override those settings and select the language you prefer to be displayed in the Sisense Web Application.
4. Select the **Role** of the user:
  - **Admin:** Admins can access the Admin section of the Sisense Web Application from where they can manage ElastiCubes, users, user groups and more. Sisense supports multiple Admin users.
  - **Data Admin:** Data Admins can manage all data sources in the system. They have access to the Admin section of the Sisense Web Application from where they can manage ElastiCubes and Live connections. Data Admins can also add and manage servers. This is useful for migrating Sisense across environments. Data Admins do not have access to User Management and System Configuration.
  - **Data Designers:** Data Designers can create and edit ElastiCubes they created or were shared with them. Data Designers can access the Admin section of the Sisense Web Application from where they can manage the ElastiCubes and Live Connections they have access to in the Data Sources page. In addition, Data Designers can see servers, but can not add new servers. Data Designers do not have access to User Management and System Configuration.
  - **Designers:** Designers can create, design, edit and share Dashboards. Sisense Designers determine whether the user with who they share a Dashboard has editing rights (is a Designer) or only viewing rights (Viewer).
  - **Viewers:** Viewers can view, explore, drill down, make selections and filter the Dashboards that Designers share with them. Viewers only need a standard web browser; no plugins or downloads. They cannot create new Dashboards or edit existing ones.
5. Click **Save**. The user then receives an email from Sisense that redirects the user to activate their account and set a password.

**To manually set a password for the new user:**

- Set the **Define Password** option to **ON** and enter a password before saving.

Define Password	OFF <b>ON</b>
Set Password	<input type="password"/>
Confirm Password	<input type="password"/>

In this case, the user receives an invitation email. You must personally notify the user of the password that you set for them.

## Adding Multiple Users

### To add multiple users:

1. Click the **Import Many** option at the top right of the Add users window.
2. In the displayed text box, paste a comma separated list of emails and click **Next**.
3. The next screen enables you to assign User rights to the users that you entered and to verify that the user count and permissions match your license.
4. After you click **Add**, each user receives an activation email.

## Editing an Existing User

### To edit an existing user:

1. Click **Admin** and select the **Users** tab on the left.
2. Click  on the right of the user's entry in the list.

## Deleting a User

### To delete an existing user:

1. Click **Admin** in the upper right corner and select the **Users** tab on the left.
2. Click  at the right of the user's entry in the list.

# Managing User Groups

This section describes how to add and use user groups in Sisense.

## How Do User Groups Work?

1. When building and sharing dashboards, wherever you can select a user, you can also select a user group. For example, you can share a dashboard with an entire user group at once.
2. When you add an additional user to an existing user group, then that user is automatically awarded the permissions of that user group.
3. When users are added in Active Directory to a previously imported Active Directory user group, then those users are automatically added to the corresponding Sisense user group.
4. User groups can only contain users. They do not support nested user groups.
5. Additional users cannot be added to a user group that was imported from Active Directory.

You can add users to a user group using the following methods:

- Add existing Sisense users to a user group
- Add a user group from Active Directory

## Accessing User Group Management

You can add, edit and delete user groups, and assign users to them.

### To access the user group settings:

- Click **Admin** and select the **Groups** tab on the left. A list of the previously defined user groups is displayed. The **Origin** column displays **Sisense** or **Active Directory** to indicate whether the user group was defined in Sisense or imported from Active Directory.

If no user groups have yet been defined, the following message is displayed: No Groups Found

# Adding a New User Group

The following procedure describes how to add a new user group by selecting users from existing users in Sisense.

See [Adding a User Group from Active Directory](#) for a description of how to add a user group from Active Directory.

## To add a new user group:

1. Click **Admin** and select the **Groups** tab on the left.
2. Click **Add Group**.

The screenshot shows the 'Create a New Group' form. It has a 'Group Name' input field with a placeholder text 'Type in a descriptive name for the new group.'. Below that is the 'Group Users' section, which includes a '+ Add Group Users' button and a search icon. A large empty box below this section contains the text 'No members added yet'. At the bottom of the form, there are two dropdown menus: 'Default Language' is set to 'System Language' and 'Default members role' is set to 'Viewer'. At the very bottom are 'Create' and 'Cancel' buttons.

3. In a **Group Name** field, enter the name for this user group.
4. In the **Group Users** field select each of the users that belong to this group. To add each user, start typing in the user's name. An auto-complete list is displayed showing the users that already exist in Sisense. Select a username to add to the group.  
**Note:** If you want, you can leave the group empty and add members to it later.
5. (Optional) The **Default Language** for a group is the language set by an Administrator for that server. From the Default Language list, you can select another default language for a group. For example, if your server supports groups located in different countries, you can set the language for each group.
6. (Optional) Select the default user role. Each member of the group is assigned this role when they are added to the group

7. Click **Create**.

## Adding Users to an Existing User Group

The following procedure describes how to add one or more users to an existing user group.

### To add multiple users to a user group:

1. Click **Admin** and select the **Users** tab on the left.
2. Select the relevant users.
3. Click **+ Add to Group**.
4. In the **Group Users** field, you can type additional usernames to add more users. An auto-complete list is displayed showing the users that already exist in Sisense. Select a username to add to the group.
5. Click **Add**.

## Editing an Existing User Group

### To edit an existing user group:

1. Click **Admin** and select the **Groups** tab on the left.
2. Click  on the right of the user group's entry in the list.  
**Note:** To remove a user from the Group, hover over the user name in the list and click .

## Deleting a User Group

When a user group that was entered in Sisense is deleted, its users are not deleted. They remain active users in Sisense.

When a user group that was imported from Active Directory is deleted, a prompt is displayed asking whether you want to also delete all the users in this user group from Sisense.

### To delete an existing user group:

1. Click **Admin** and select the **Groups** tab on the left.
2. Click  at the right of the user group's entry in the list.

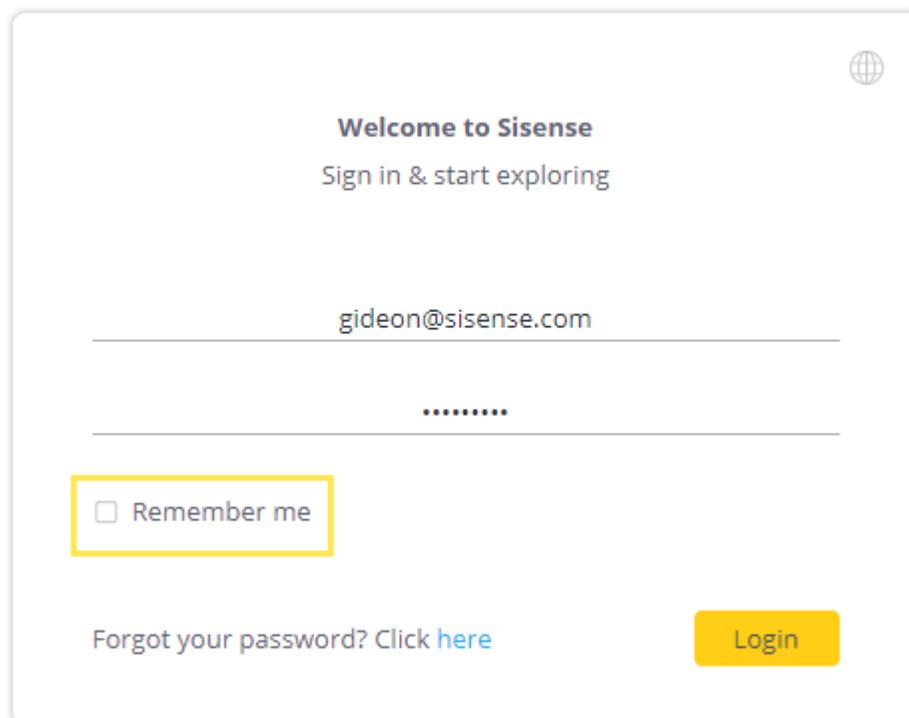


# Managing User Sessions

When a user signs into the Sisense Web Application, a session cookie is stored in their browser. The session cookie allows the user to remain logged in and authenticated even after ending a session by closing the browser.

For security reasons, you may want to logout your users or end their session when they close the browser. Through the Sisense REST API, you can add an additional property, `sessionOnlyCookie`, to the POST, `/api/settings/security` of REST API v.9.

When the Boolean value of `sessionOnlyCookie` is true, the login page does not display the “Remember Me” checkbox in the user’s browser and logging in to Sisense generates a session cookie.



gideon@sisense.com

.....

Remember me

Forgot your password? [Click here](#) Login

When the user logs out or closes their browser, the session ends and they must log in again to access Sisense.

The following is an example of a cURL command in which the Remember Me button is no longer displayed in the Login page.

```
<pre>curl -X POST -header "Content-Type:
application/json" -header "Accept: application/json" -d
"{
\"sessionOnlyCookie\" : \"true\"
}" "http://localhost:8081/api/settings/security"</pre>
```

# Managing Sisense Users

Sisense Administrators have access to the User Management section of the Admin Console where users, groups, and advanced user settings can be configured.

The topics below describe how to create users and groups, and allow users you authenticate to access your implementation of the Sisense Web Application.

- [Managing Dashboard Users](#)
- [Managing User Groups](#)
- [Managing User Sessions](#)
- [Customizing User Roles](#)
- [License Utilization](#)
- [Integrating Active Directory](#)
- [Introduction to SSO](#)

# Administration

Sisense supports several [user roles](#) with Administrative privileges for customizing Sisense for your company. The following roles: System Administrator, Data Administrator, and Administrator have varying levels of access to your system including servers, users, ElastiCubes and more. The Admin page is your dedicated location in your web application for managing your system.

This section describes how your Sisense Administrator can configure and manage users, groups, data sources, and system settings from the Admin page.

The screenshot shows the Sisense Admin console interface. The top navigation bar includes 'Data', 'Analytics', 'Pulse', and 'Admin'. The 'Admin' tab is selected. The left sidebar contains a menu with 'User Management', 'Data Management', and 'System Configuration'. The main content area displays a table of users with columns for 'User Name', 'Email', 'Name', 'Groups', 'Roles', and 'Origin'. Three callout boxes provide instructions: 'Manage Users, Groups, and User Authentication here' points to the 'User Management' menu item; 'Access the Admin Console here' points to the 'Admin' tab; and 'Admins and Data Admins can manage data sources here' points to the 'Data Sources' menu item. Another callout 'Customize your Sisense configuration from here' points to the 'Settings' menu item. The table lists three users: gideon@sisense.com (Sys. Admin), sam@sisense.com (Data Admin), and John@sisense.com (Viewer). The footer shows 'Page 1 of 1' and '3 Users found'.

**Note:** The Admin page is only displayed to [users](#) with the relevant privileges. In addition, certain settings are only available to certain user roles except for Administrations who have access to all settings.

- [User Management](#)
- [Dashboard Management](#)
- [Data Source Management](#)
- [System Settings](#)
- [Embedding and Rebranding](#)

- [Security](#)
- [Migration](#)
- [High Availability](#)

# ElastiCube Manager Management

Data Administrators can manage ElastiCube servers, create ElastiCube Sets for querying during builds and share ElastiCubes with other users.

The following topics describe how you can manage ElastiCubes in Sisense:

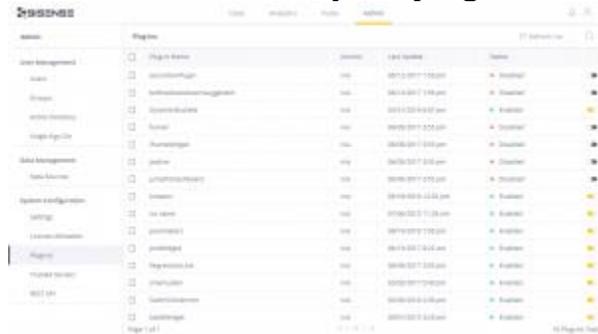
- [Managing ElastiCubes](#)
- [Sharing ElastiCubes](#)
- [Managing Drill Hierarchies](#)
- [Managing Live Data Sources](#)
- [Managing Row-Based Data Security](#)

# Managing Plug-ins

Sisense plug-ins enable you to extend the functionality and capabilities of your Sisense dashboard through standard web development languages such as JavaScript and HTML.

For more information about developing Sisense Plug-ins, see [JavaScript API & Plugins](#).

In the Sisense Admin page, you can view a list of your current plug-ins, their version as defined in your plug-in's script, and their status, enabled or disabled.



From the Plug-ins page, you can enable or disable your plug-ins. This is useful for testing new plug-ins or troubleshooting issues with your dashboard. When a plug-in is disabled, its functionality is no longer available in the Sisense Web Application.

## To enable or disable a plug-in:

1. In the Sisense Web Application, select **Admin > Plugin-Ins**.
2. Toggle the Enable/Disable switch to activate or deactivate a plug-in.

OR

To activate/deactivate multiple plug-ins, select the checkbox of the relevant plug-ins. When you select any of the checkboxes, an Enable/Disable switch is displayed in the menu bar. Toggle the switch to enable or disable all the selected plug-ins.



After enabling or disabling a plug-in, Sisense rebuilds your Plug-ins' list. The

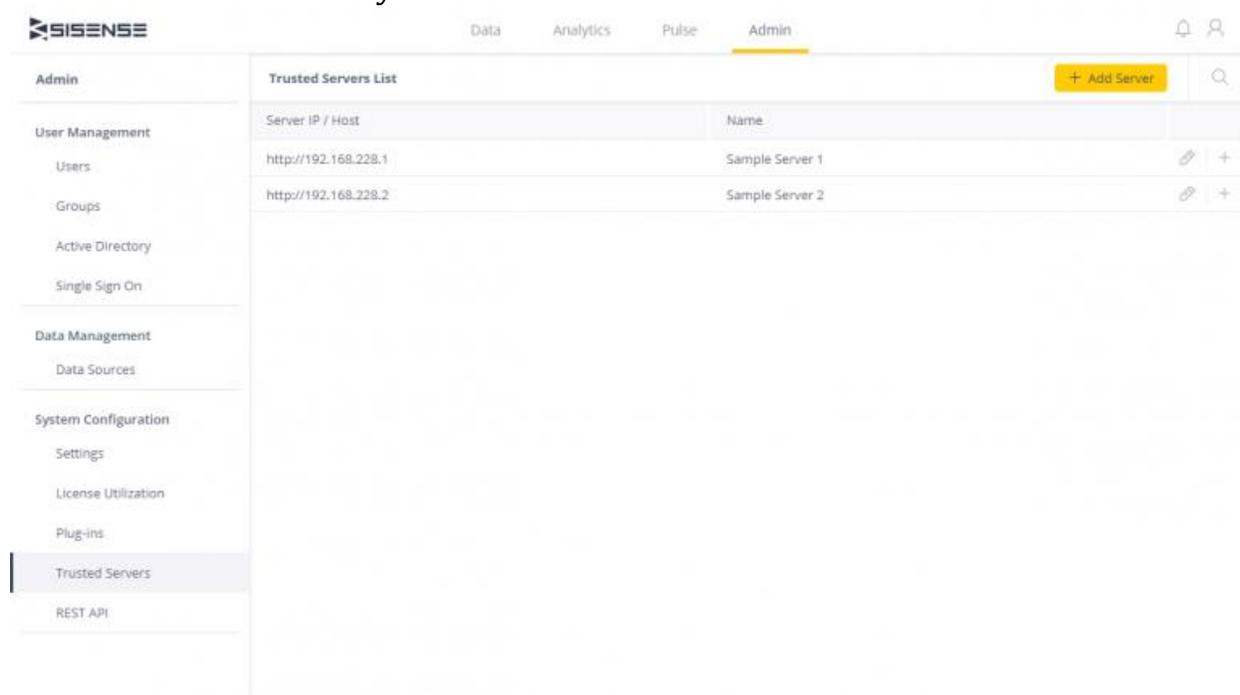
next time you refresh your dashboard, the plug-ins will be enabled or disabled as defined in the Plug-ins page.

# Migrating Sisense across Environments

Many users that implement Sisense maintain copies of their dashboards locally or across multiple environments for a variety of reasons. With multiple instances of Sisense, you can test and verify your data's accuracy before publishing dashboards to a production environment where accuracy is critical. You can save external copies of dashboards for disaster recovery. OEMs can support a master server and separate servers for each of their customers, copying the relevant dashboards from the master server to the appropriate customer server. Sisense enables Administrators and Designers to maintain multiple copies of their dashboards and migrate them quickly and easily across multiple environments. There are two options for migrating dashboards from the Sisense Web Application. The first option is to export your dashboards into a .dash file locally, and then import the dashboards into the relevant environments. This option is useful if you do not have access to the server, for example, if you are an OEM providing dashboards to your customers. For more information about exporting dashboards, see [Exporting Sisense Dashboards](#).

The second option is to copy the dashboard to another server from the Sisense Web Application. In this scenario, you can copy dashboards from one environment directly to another environment after the Administrator of the target server provides access to the source server. The target server can only accept files from servers available on the Trusted Server List. For example, if you have two machines with Sisense installed, on the server where you want to copy

dashboards to, you can provide access to the source server by adding the source server's IP address to my list of allowed servers.



After providing access, an Administrator or Designer can copy the dashboard directly to the target server. If the dashboard already exists on the target server, the Administrator or Designer will be notified and can overwrite the existing dashboard or create a duplicate dashboard on the target server.

The option you choose depends on which methodology is convenient for you and what you need to achieve. The end result is the same.

There are several things you should remember when copying a dashboard to another server. First, copied dashboards and their widgets keep their IDs. This ensures that any users with bookmarks of the dashboard or its widgets can continue using the same bookmarks. If the dashboards already exist on the recipient server, Sisense prompts you to select how you want to handle each dashboard that already exists. In the Copy Dashboards to Server window, a list of all the dashboards to be copied to the recipient server is displayed. The list includes the dashboards name, the destination path where the dashboard is to be

copied, and several options for how you wish to handle the dashboard in case it already exists.

In the example below, you can see that the Sample – Healthcare dashboard has been renamed to “Renamed Dashboard”. Sisense recognizes that these dashboards are the same even though the name on the recipient server has been changed.

For the Sample – Ecommerce dashboard, the dashboard has been saved in a folder called “test” on the recipient server. The dashboard’s full destination path on the recipient server is displayed. When the dashboard is located in the main navigation tree, only the dashboard name is displayed.

### Copy Dashboards to Server

The following dashboards already exist. Select action:

Sample - Healthcare	Renamed Dashboard	1	Skip ▾
Sample - Lead Generation	Sample - Lead Generation		Skip ▾
Sample - Ecommerce	test\Sample - Ecommerce	2	Skip ▾

Skip All | Duplicate All | Overwrite All

OK Cancel

Second, when copying a dashboard, the dashboard’s scripts on the target server are overwritten. The scripts from the source dashboard are used instead. If your data source is not the same in each environment, you will need to manually adjust the data source after copying the dashboards. If plug-ins are installed on the original server, these plug-ins will not work on the dashboard unless they are also installed on the target server.

Third, changes to your dashboards are not reflected on the same dashboard in another environment. To modify a remote dashboard you can either modify it on the remote server, or modify it in your source server, and repeat the migration process. Once you have completed your changes, you will need to republish the dashboard.

## Copying Dashboards between Sisense Servers

Administrators and Designers with access to a remote Sisense server can copy dashboards from one Sisense server to another from the Sisense Web Application. When you copy a dashboard or a folder to a remote server, all the dashboards, including those in the selected folders are duplicated into the main Dashboard list in the target server. Folders are not duplicated on the target server.

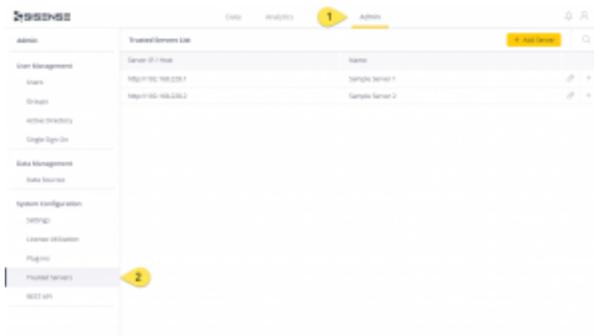
Before copying a dashboard to a target server, the Administrator of the target server must provide access to target server. The procedure below explains how to provide access to another server and then how to copy a dashboard from the source server to the target server through the Sisense Web Application.

**Note:** When migrating a dashboard to another server, you must be logged in to your local server with the IP address as the URL. You cannot copy dashboards from localhost.

### **To copy your dashboard to another Sisense server:**

**Note:** Some of the dialog boxes may be considered as pop-ups by your ad blockers. Make sure that Sisense is whitelisted or that your ad blocker is disabled when performing this procedure.

1. On the target server, the Administrator needs to provide access for the source server to the target server. On the target Sisense server, select **Admin > Trusted Server List**.



The Trusted Server List page is displayed.

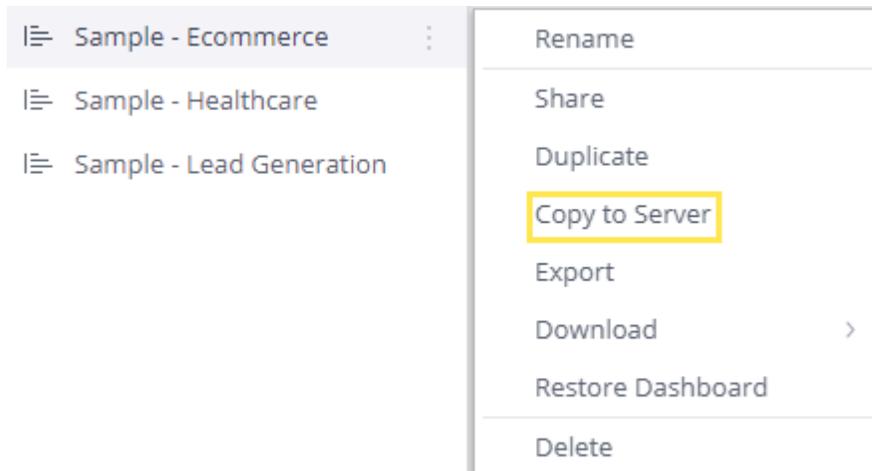
2. In the Trusted Server List page, click **Add Server**. The Create a New Server Access dialog box is displayed.

### Add Server

Server	<input type="text" value="Type URL or Hostname (Including Port)"/>
Name	<input type="text" value="Optional"/>

Cancel

3. In the **Server** field, enter the IP address or hostname of the source server that contains the original dashboard.
4. In the **Name** field, enter a name for the source server. This is an optional identifier to help you keep track of your servers.
5. Click **OK**. The server is added to the list and now has copy access to the target server.
6. On the source server, from the Dashboard list, click the dashboard's menu and click **Copy to Server**.



OR

To copy multiple dashboards, in the Dashboards list, click the Select Dashboards  button to display checkboxes next to the dashboard and folders, which allows you to select multiple dashboard to be exported. Then, select the relevant dashboards and folders to be exported and click . The Copy Dashboards to Server dialog box is displayed.

### Copy Dashboards to Server

Server

Type URL (including port number)

Republish dashboards after copying

Cancel

7. In the Copy Dashboards to Server dialog box, enter the following information:

**Server:** The URL or Hostname of the remote server. Sisense supports SSL connections so dashboards can be transferred securely when SSL is activated.

**Republish dashboards after copying:** Select to automatically republish

the dashboard on the target server after the selected dashboards are copied. Republishing a dashboard refreshes the dashboard definitions (widget types, titles, scripts etc.) for each of the users who share the dashboard.

8. Click **OK**.
9. If the dashboards already exist on the recipient server, Sisense prompts you to select how you want to handle each dashboard that already exists. You can select one of the following options for each dashboard:
  - **Skip:** No changes are made to the dashboard on the recipient server.
  - **Overwrite:** The new dashboard writes the existing dashboard on the recipient server.
  - **Duplicate:** A new dashboard is created on the recipient server with a number added to the end of the dashboard's title, for example, "Sample Dashboard (1)".After you have defined how the dashboards are to be handled, the selected dashboards are copied to the remote server and displayed in the Dashboard's list.

## Limitations

- You cannot copy a dashboard to another server when SSL is enabled on the origin server

# Viewing Dashboards on Mobile Devices

Sisense Mobile connects you to your dashboard anytime, anywhere on your Android or iPhone mobile devices. Sisense Mobile decreases time to insight with an intuitive UI, search bar for locating dashboards, and responsive visualizations; You can explore your dashboards with simple gestures optimized to let you dive deep into your data or apply filters to view the data you want. Check out a demo video below to see how you can tap, zoom, and swipe to interact with your data.

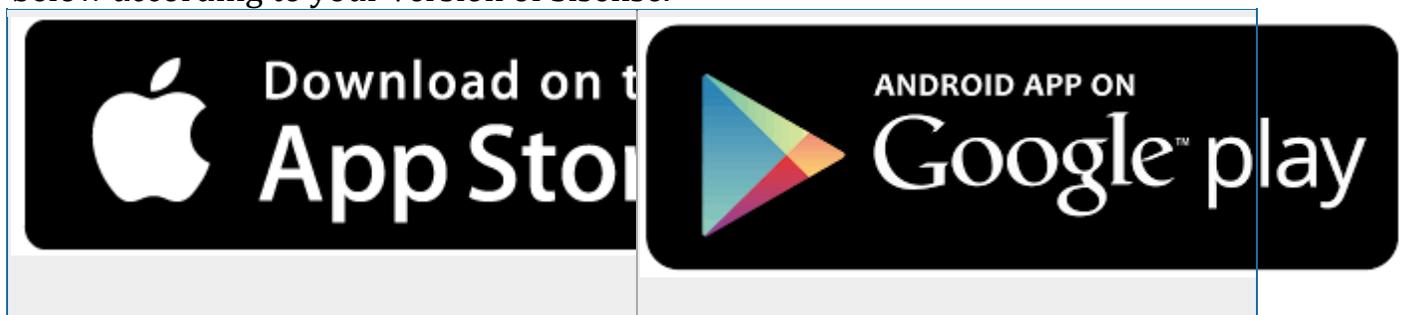
## Getting Started

Supported Phones and Operating Systems

- iPhone – Requires iOS 8 or later, Supported iPhone 5 device or higher.
- Android – Requires Android versions 4.3 or later.

## Get Sisense Mobile

To download the Sisense Mobile App for your device, click the relevant link below according to your version of Sisense.



--	--

**Note:** Sisense Mobile is supported by Sisense V6.5 and later.

## Logging in to Sisense Mobile

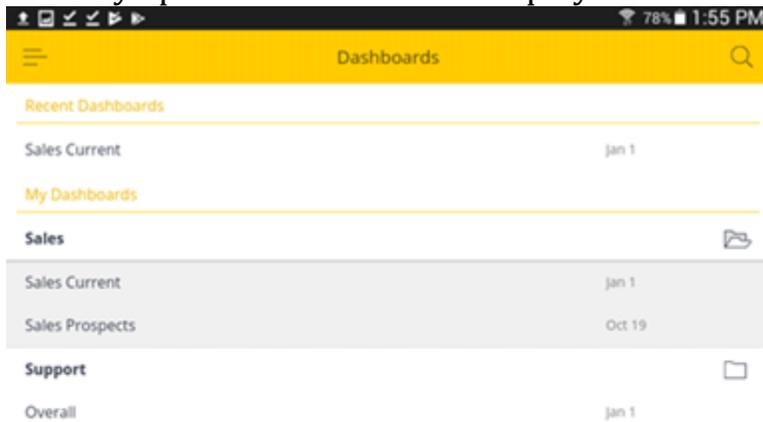
After you have downloaded Sisense Mobile, open the app to launch the login screen and enter the server address and your Sisense user credentials.

The server address should include the server IP or domain name (alias). If you are not using the default port (80) you will need to add the port as well, for example: `http://xx.xxx.xx.xx:8081`. When working with SSL security protocol, use the 'https' prefix.

You can sign out by opening the navigation menu and tapping Sign Out.

### ***Navigating Sisense Mobile***

When you first log in to Sisense Mobile, a list of your folders, dashboards, and recently opened dashboards is displayed.

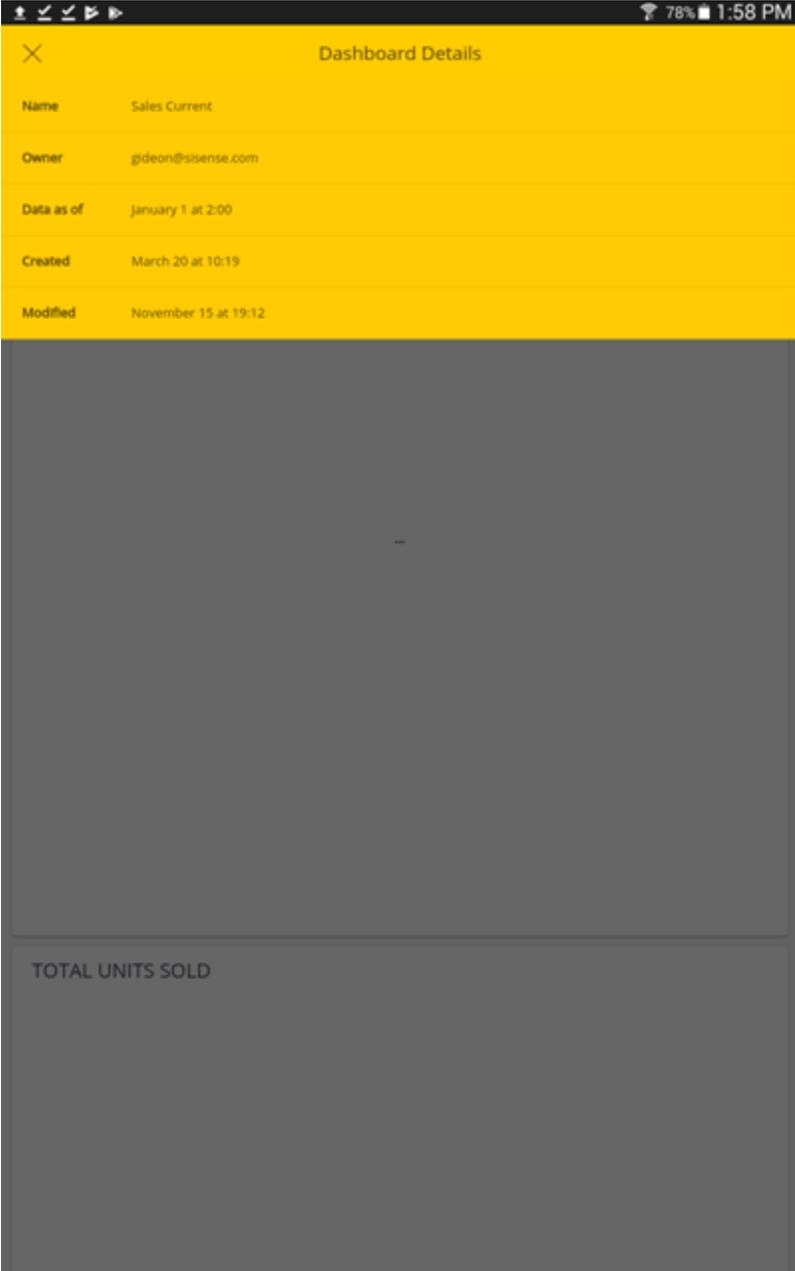


By tapping a dashboard, you can open the dashboard or if your dashboard is located in a folder, tap the folder then tap the relevant dashboard to open it. You can also locate a dashboard quickly by tapping the search icon and typing the name of the dashboard. As you begin to type, any dashboards that match the

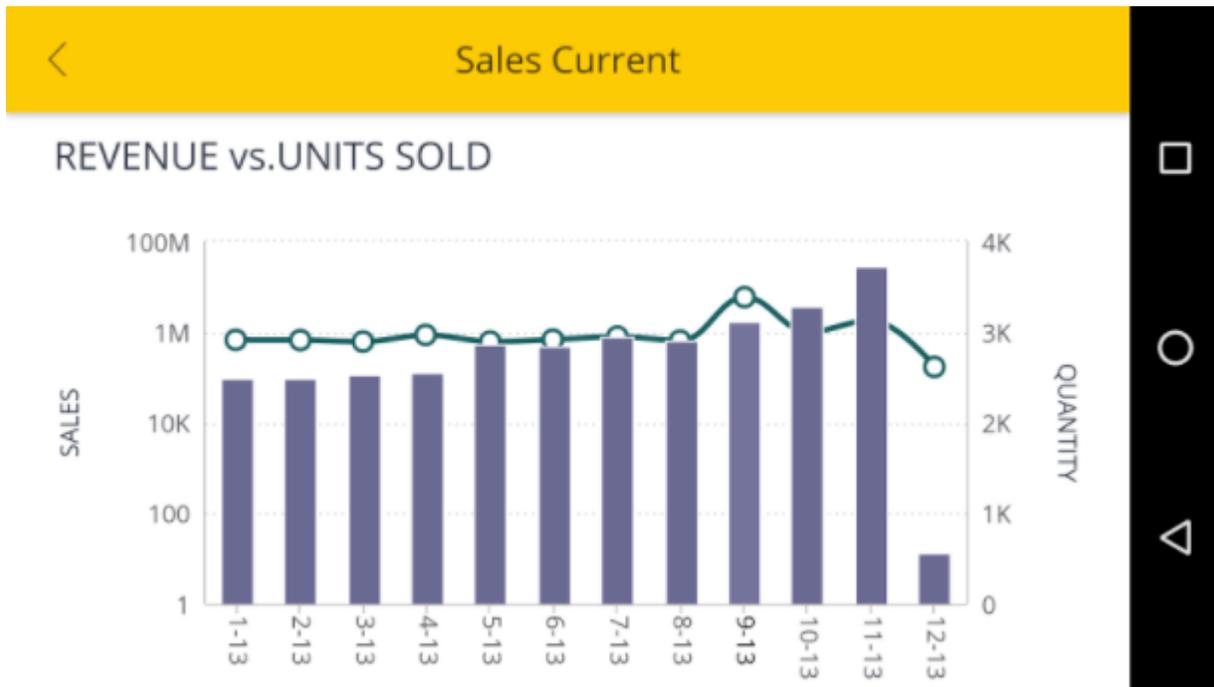
characters you entered are displayed.



Once the dashboard is displayed, you can tap the dashboard's title to display more details.



By rotating your device horizontally, you can view your dashboard in landscape view.

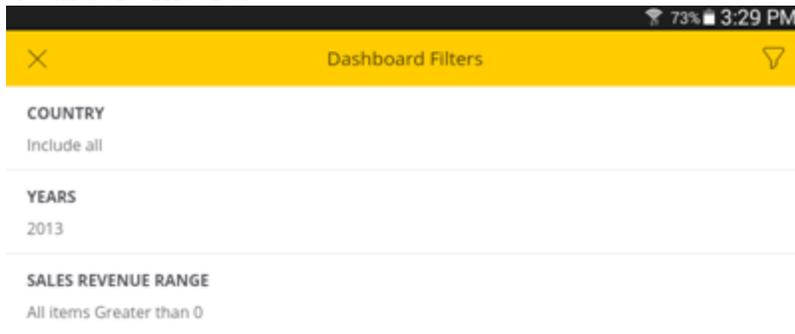


### *Exploring Data*

In Sisense Mobile, you can filter and drill down into your data much like you do from your PC.

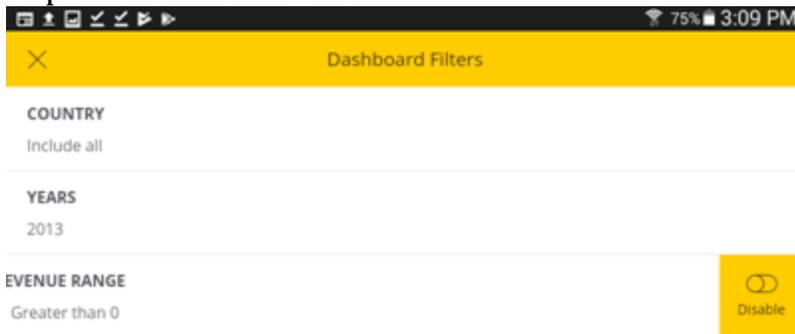
## Filtering Widgets

To filter a dashboard, you can tap the filter icon . This displays the list of available filters.

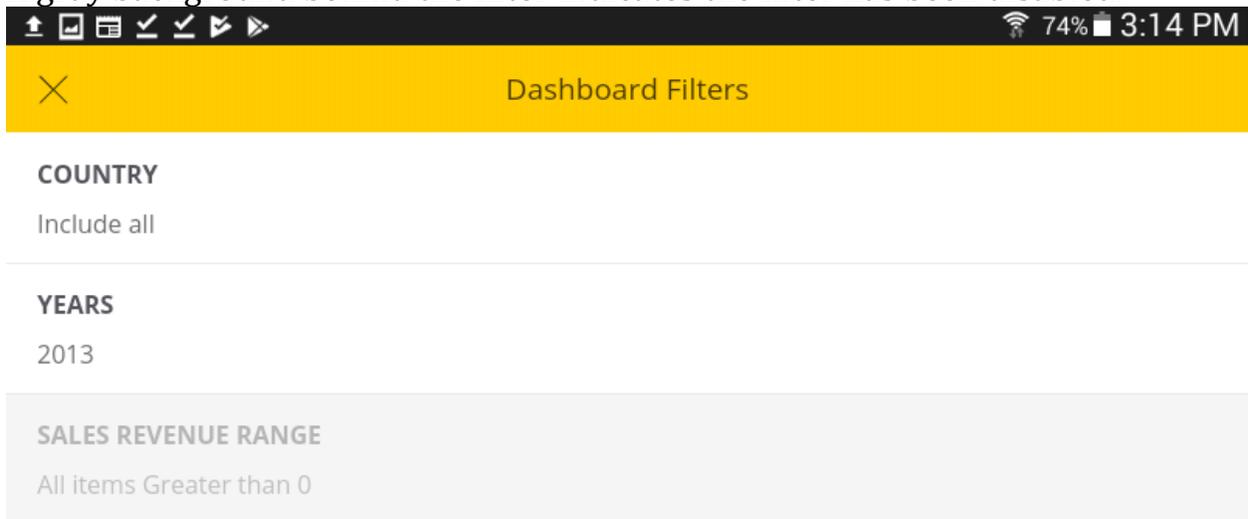


You can disable your filters by swiping left on the filter to display a toggle button.

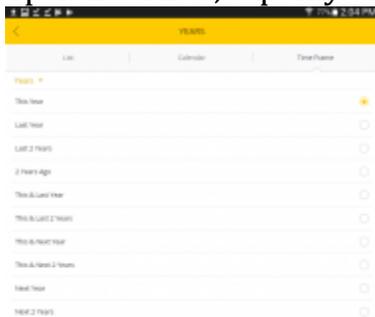
Tap it to disable a filter.



A gray background behind the filter indicates the filter has been disabled.



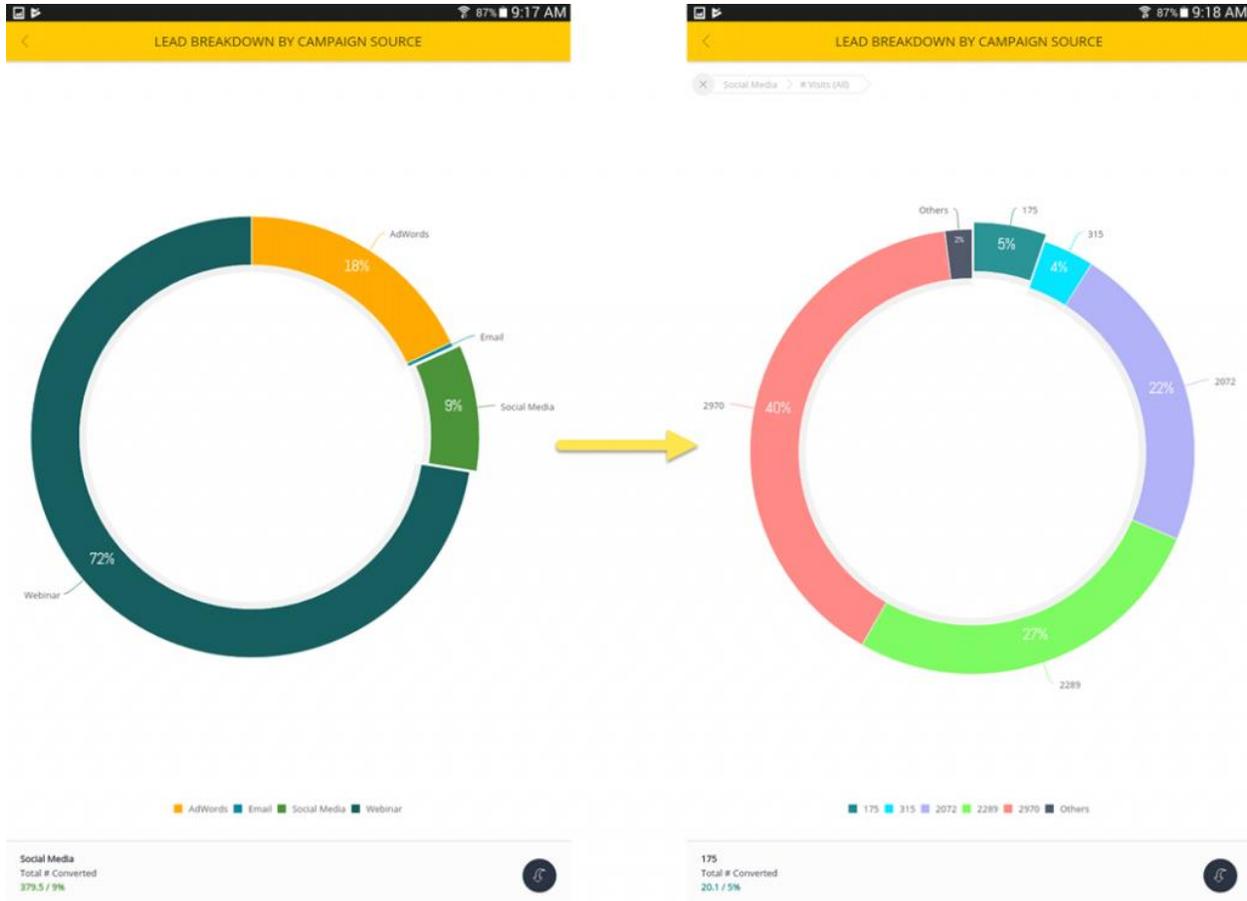
To dive deeper into your filters, tap the relevant filter to open all of its available options. Then, tap any of the filter options to toggle the filters to your data.



## Drilling into your Data

You can drill down in most widgets to get an in-depth view of a selected value. For example, tap on a slice in a pie chart to see the absolute value in addition to the percentage.

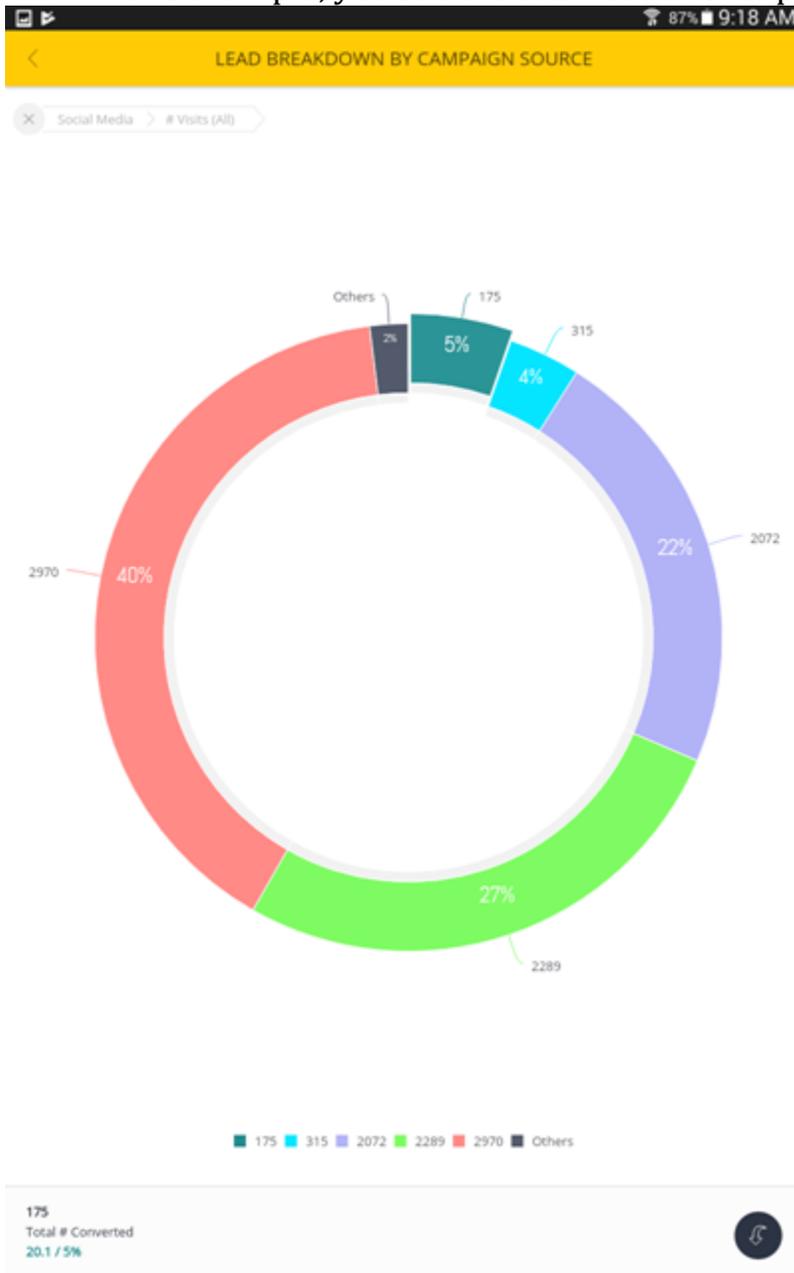
In the following example, the left chart shows the lead breakdown by campaign source, and on the right side, the chart shows a breakdown of visits from various social media sources.



Any time you select a value, a tooltip is displayed at the bottom of your screen, which provides a description of the selected value and the drill down

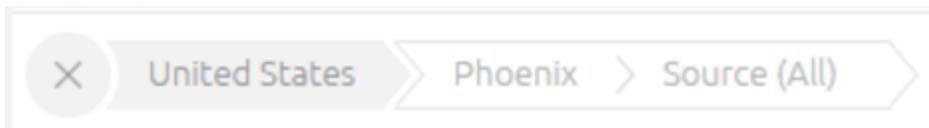


icon. In this example, you see the total lead and its percentage.



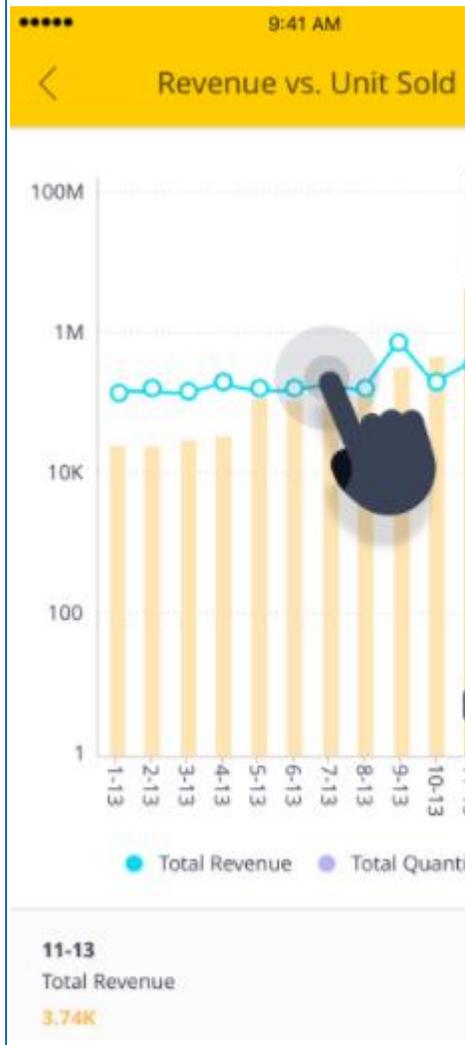
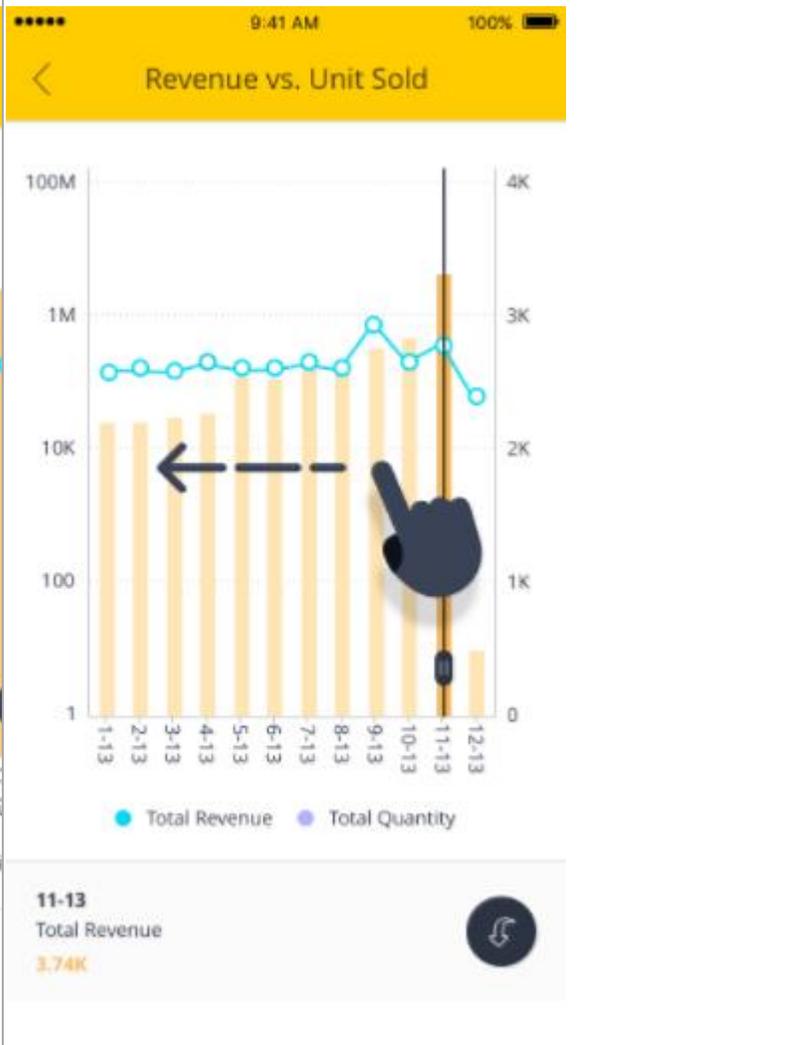
**Note:** You cannot drill down while in landscape mode.

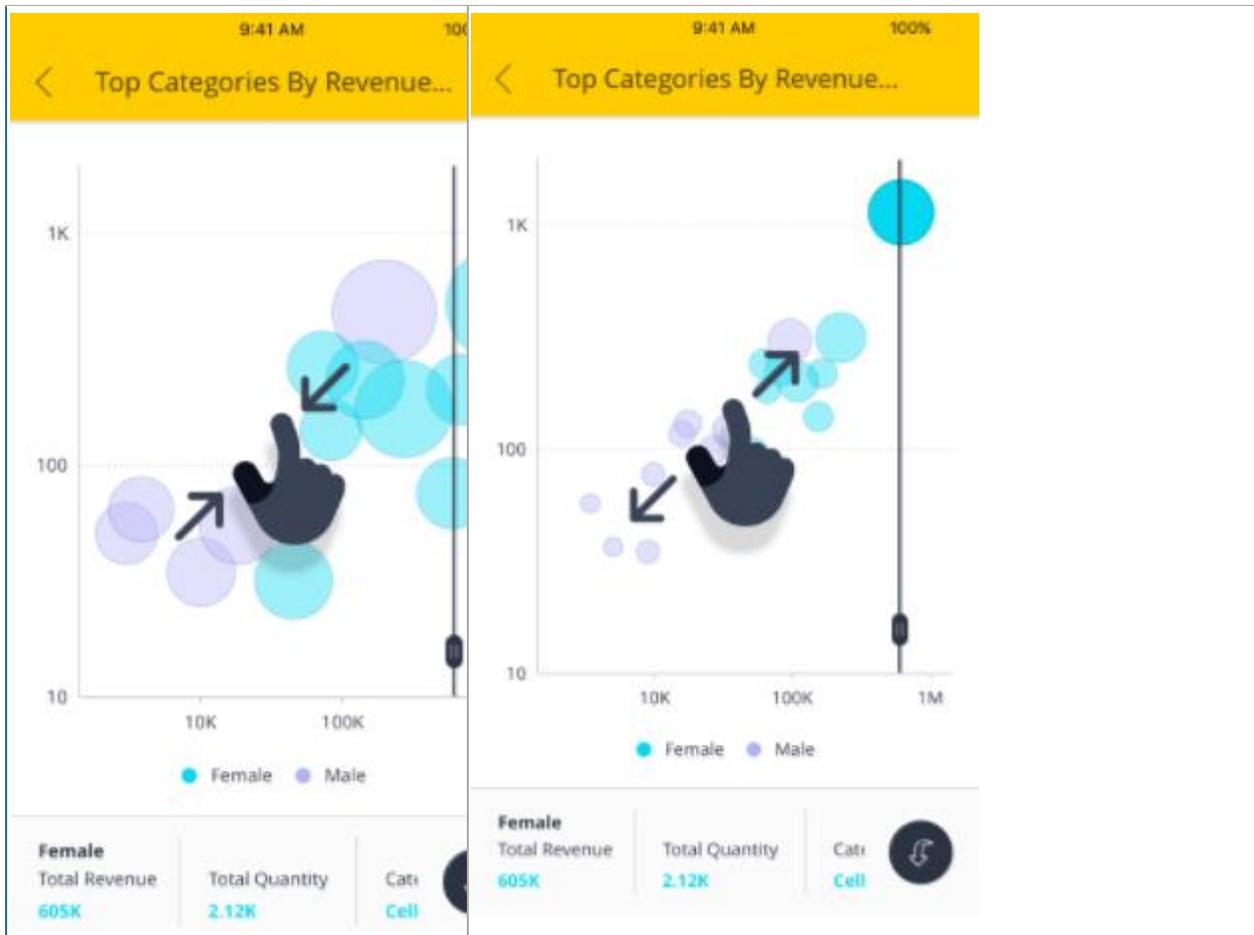
To return to a higher level, click on a breadcrumb. To drill all the way up, click on the X icon.



## Viewing More Details

Sisense Mobile supports a variety of gestures that let you dive deeper in your data and see more details.

Highlight a Value	Swipe across your Visualizations
 <p><b>11-13</b> Total Revenue 3.74K</p>	 <p><b>11-13</b> Total Revenue 3.74K</p>
Pinch in to Zoom	Pinch out to Expand

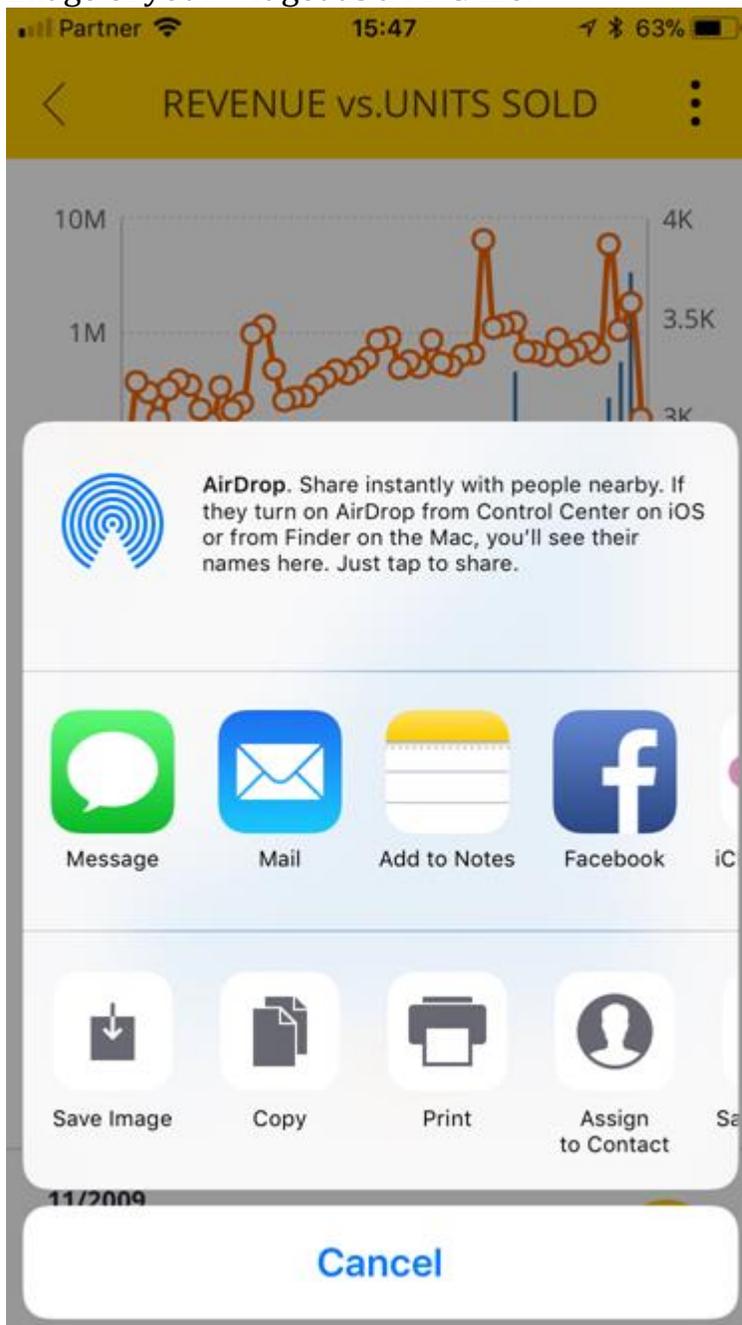


## Sharing Widgets

You can share widgets as PNG images from Sisense Mobile across various apps supported on your device.

In your dashboard, tap a widget to open Widget Mode, then tap the menu icon  and tap Share. Your device will display a list of apps that you can use to send the

image of your widget as a PNG file.

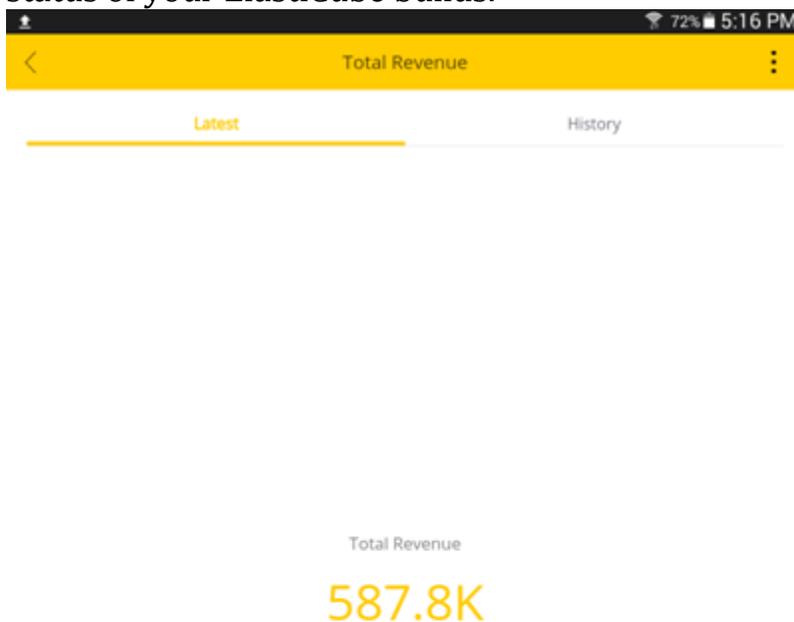


## ***Sisense Pulse***

Sisense Pulse is a centralized location where you can stay on top of your most important KPIs across multiple dashboards or manage your data and build alerts. By adding important KPIs from your dashboards to Sisense Pulse, you can get a comprehensive picture of your data from a single location.

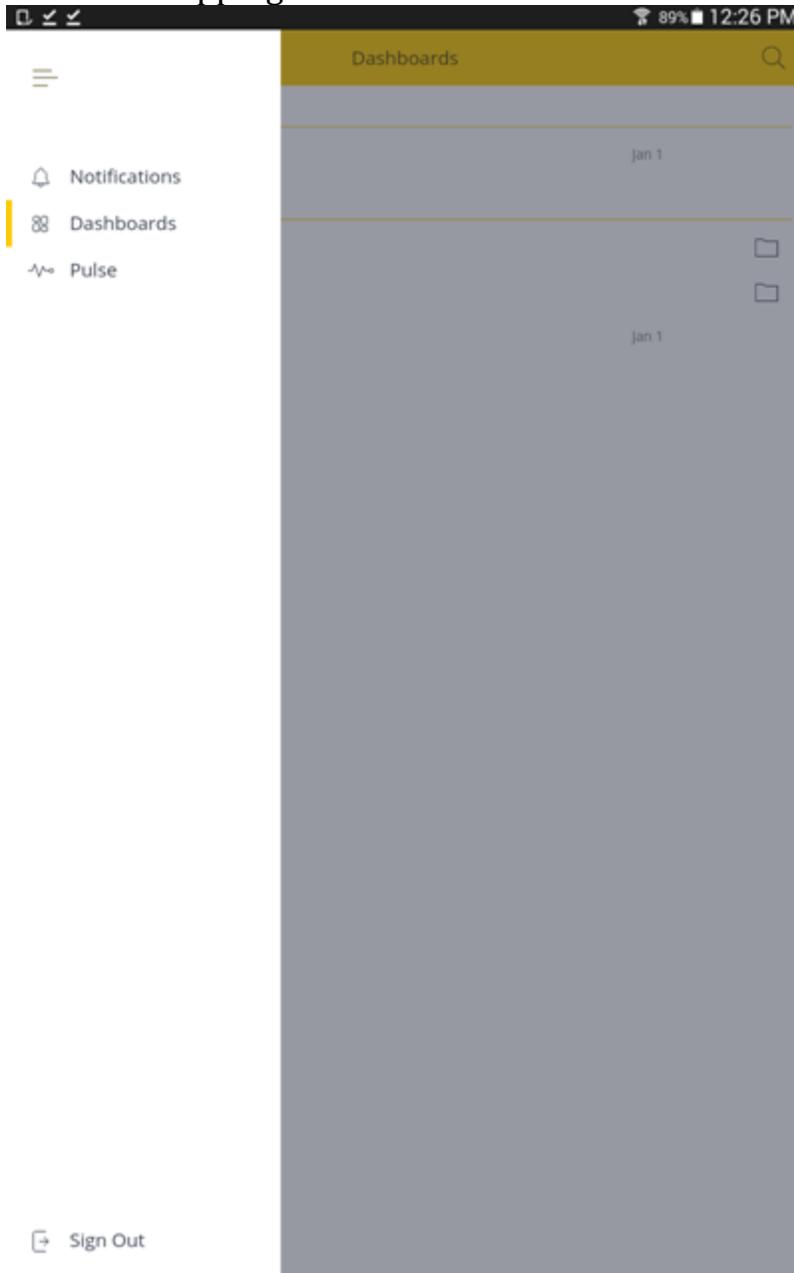
Sisense Pulse contains tiles that display information from your dashboards and the status of your ElastiCube builds.

Within Sisense Pulse, there are two types of tiles, data tiles that display your widgets' current status and recent history and build tiles that display the latest status of your ElastiCube builds.



In Sisense Mobile you can manage your alerts from Sisense Pulse or create data alerts based on widgets in your dashboards.

**Note:** Build alerts can only be created from the Sisense Web Application. You can access Sisense Pulse from Sisense Mobile by opening the navigation menu and tapping Pulse.



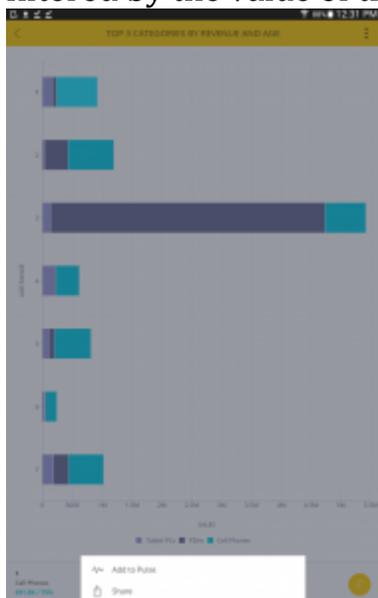
## Creating Data Alerts

Data alerts are triggered by conditions you define for your widgets. For example, if you have an Indicator widget that represents your company's revenue, you can

define an alert that notifies you when this revenue reaches or falls below a certain threshold.

For each type of widget you can add to Pulse, you can monitor a single value. For example, for Column charts, you can monitor a single section (value) of a column. To monitor multiple values, you can add each value you want to monitor to Pulse separately.

When you create an alert for a specific value, Sisense applies any filters that were already applied, for example, if you create a data alert for the total revenue of a given age range, the pulse alert created will be called Total Revenue and will be filtered by the value of the age range you selected.



To create a data alert, from your dashboard, tap the relevant widget to open it in Widget Mode. In Widget Mode, tap the value you want to monitor, then tap the menu icon  and tap **Add to Pulse**. This opens the Alert Conditions menu where you select the type of data alert you want to create. For example, if you create a

threshold data alert, the following settings are displayed.

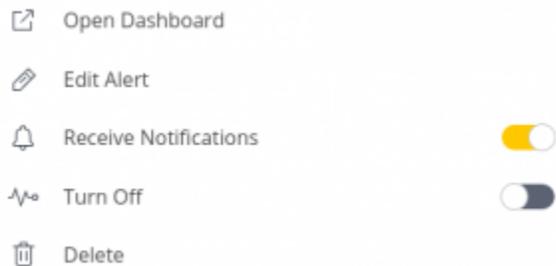
The screenshot shows a mobile application interface for configuring an alert. At the top, there is a yellow header bar with a back arrow on the left, the text "Add to Pulse" in the center, and a "Save" button on the right. Below the header, the settings are organized into sections:

- Alert Name:** Total Revenue
- Alert Condition:** Threshold > 691797
- Notification Methods:**
  - Mobile:
  - Email:
- Notification Limitations:**
  - Notify once:
  - Notify when no longer met:

After you have set the conditions, tap **Save** and the new alert will be added to Sisense Pulse.

## Managing Alerts

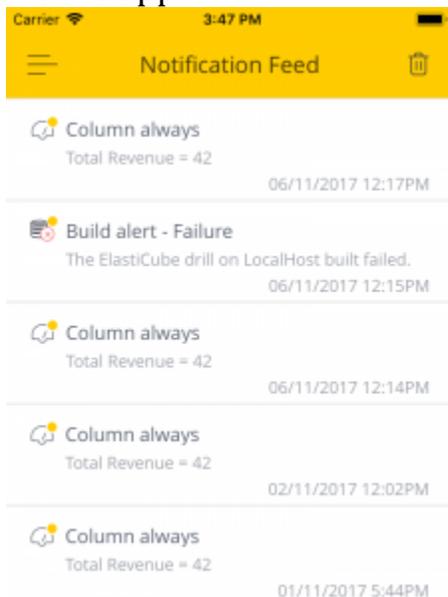
You can manage your alerts from Sisense Pulse. In Sisense Pulse, tap the relevant tile to open Tile Mode. Once you're in Tile Mode, the menu icon  is displayed in the top right corner. Tap it to display a list of alert management options.



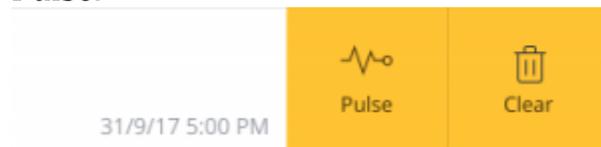
From the options area, you can edit alerts, delete alerts, deactivate alerts, and toggle notifications.

## Pulse Notifications

If your Sisense user has been configured to receive Pulse alerts to your mobile device, when an alert is triggered, you will receive a notification from the Sisense Mobile app.



You can swipe left on an alert to clear the alert or to view the alert in Sisense Pulse.



To clear all alerts, click the trash can icon  at the top of the notification feed.  
**Note:** You must have Sisense Mobile installed to receive Pulse notifications.

## ***Single Sign Open***

If you have implemented Single Sign On (SSO), your users can log in to the Sisense Mobile app.

This is useful for maintaining one set of login credentials across multiple apps.

See [Single Sign On](#) (SSO) for more information.

## ***Release Notes***

### **March 3, 2018**

- Support for Hours and Minutes time resolutions in date filters (Sisense 7.0.1 and above)
- Ability to collapse/expand Recent Dashboards
- Improved support for SSO connectivity
- Bug fixes

# Connecting to Sisense via ODBC

Sisense has its own ODBC driver that allows you to connect your tool of choice directly to the Sisense ElastiCube.

Some use case examples:

- Connect a 3rd party visualization tool directly to the ElastiCube.
- Build pixel-perfect reports using your existing reporting tools by connecting directly to the ElastiCube
- Import data directly into Excel from ElastiCube tables
- Importing Sisense ElastiCube tables from one ElastiCube to another.  
Contact us through our [Help Center](#) for more information for implementing this use case.

**Note:** Once installed with your Administrator's user credentials, any user can access the ElastiCube over ODBC. Connecting to Sisense via ODBC is not recommended for use cases that require data security.

This page will cover the following:

- Installing the Sisense ODBC driver
- Creating a new Sisense ODBC connection
- Use case example 1: Importing data from an ElastiCube into Excel
- Use case example 2: Importing data from an ElastiCube into Tableau

The Sisense ODBC Driver is licensed for use by the number of licensed Users (Administrator, Designer or Viewer) of the Software licensed hereunder. Only the number of licensed Users of the Software are entitled to use or view data or query results produced via the Sisense ODBC Driver in third party tools.

# Installing the Sisense ODBC driver

Depending on whether you will be connecting to Sisense through a 32 or 64 bit application, download and install the appropriate driver:

- Download Sisense ODBC for 32-bit applications
- Download Sisense ODBC for 64-bit applications

## Creating a New Sisense ODBC Connection

Sisense ODBC is a licensed feature.

Connections will only be accepted by the Sisense ElastiCube if the ODBC functionality has been enabled in your license and if you are using Sisense Version 6.0 or later. Please contact your customer support manager or open a request through our Help Center if your license currently does not include this.

To upgrade to the latest version of Sisense, [click here](#).

The following procedure explains how to create a new Sisense ODBC connection.

Creating such a connection has to be done on every machine from which you want to connect to an ElastiCube via ODBC. Each connection is defined for a specific ElastiCube, so repeat this step for multiple ElastiCubes if necessary.

### To create a new Sisense ODBC connection:

1. From the **Start menu**, open the Windows ODBC Data Sources application. Notice that there are two different instances of it. One for 64 bit and another for 32 bit.  
Tip for Windows 7 users: To quickly access the 32-bit ODBC window type "c:\Windows\SysWOW64\odbcad32.exe" from the Run... menu.
2. Under the 'System DSN' tab click Add..., and select "Sisense ODBC Driver". Click **Finish**.
3. In the "New Sisense Datasource" window, fill in the following details:  
In the **Data Source Name** field, enter a descriptive name for your connection.  
(Optional) In the **Description** field, you can enter a description.  
In the **Server** field, enter the Sisense server address.

localhost:8081/app/settings#/users

4. Click on the **ElastiCube** drop-down list. You will be prompted to enter your username and password.
5. From the **ElastiCube** drop-down list, select the ElastiCube.
6. If you are using a secured connection (SSL), select Use secure connection (SSL).
7. Click **Test** to verify the connection to the ElastiCube.
8. Click **OK** to save the connection.

## Connecting to Sisense in Excel

In Excel, you can use the ElastiCube ODBC connector to import data from your ElastiCube into your Excel worksheet or pivot, and also through Microsoft Query. After downloading and adding the ODBC connector, you can use any of the following methods to import the data from the ElastiCube.

### Method 1: Import data directly into a pivot table

1. In Excel, create a Pivot Table by clicking **Pivot Table** under the **Insert** ribbon.
2. Select **Use an external data source**. Click **Choose Connection...** to browse and locate the source table.  
If the connection does not appear, click **Browser for More...**, and then the **Connect to New Data Source** option.  
Select **ODBC DSN**, and then the data source. Enter your username and password in the appropriate fields, and click **OK**.
3. Click **OK** to connect to the data.  
Now you can drag and drop the fields into the relevant columns and rows, as you would for any pivot table in Excel.

## Method 2: Import data using the Microsoft Query option

1. Under the **Data** ribbon, select **From Other Sources > From Microsoft Query**.
2. Select the database, and click **OK**.
3. You can either click on the tables you want to add, or click **Cancel** to enter a visual environment and add the relevant tables.
4. You can now drag and drop connectors between fields from the different tables to create relationships.
5. Click on a field to add the data to Excel. Repeat for each field that you want to add.
6. You can now import the data into the worksheet. Click on the import button, and select the import options, such as where you want to import the data in the worksheet.

## Using Sisense as a Live Connection in Tableau

In Tableau, you can use the ElastiCube ODBC connector to run your queries directly in the Sisense ElastiCube instead of in Tableau's in memory data engine.

After downloading and adding the ODBC connector, follow these steps:

1. In Tableau, in the **Data** tab, click **Connect to Data > Other Databases (ODBC)**.
2. From the **DSN** dropdown menu, select the ElastiCube, and click **Connect**.
3. You can now begin working with the data and define relationships. Add tables as you would for any other data source in Tableau, and click Update Now.
4. Create a new worksheet in Tableau, and add the tables from the ElastiCube as needed to run queries and create visualizations.

## Accumulative Build Support

The Sisense ODBC connector supports accumulative builds for certain data sources. Some of these are supported by default, while others, you must activate through the ODBC configuration file. The list of supported data sources is below.

If you want to activate accumulative build support for a data source, open the file `odbcConfig.json` located in the directory:

`C:\Program`

`Files\Sisense\DataConnectors\DotNetContainer\Connectors\GenericODBC\`

In the configuration file are JSON objects that represent each data source. Within each data source object is a boolean called "IsAccumulativeSupported". When this value is true, the option to perform an accumulated build is available, and the value false removes this option. Below is an example object for MsSQL.

```
{  
  "Provider": "Mssql",  
  "DriverNames": [  
    "SQLSRV32.DLL",  
    "SQLNCLI11.DLL"  
  ],  
  "IsAccumulativeSupported": true  
},
```

After changing the value of `IsAccumulativeSupported`, save the file to update the supported build options in Sisense.

Data Source	IsAccumulativeSupported
MySQL	true
MsSQL	true
MongoDB	true
BigQuery	false
HiveODBC	false
Snowflake	false
MapR Drill	false