

# Anzo<sup>®</sup> 5.4 User Guide

Last Updated: 1/29/2024

Online documentation is available at docs.cambridgesemantics.com

© 2024 Cambridge Semantics, Inc.

# **Table of Contents**

About This Doc	
Onboard & Virtualization	
Onboard Structured Data	21
Adding Data Sources	
Connecting to a Database	23
Defining a Database Schema	
Adding an HTTP or SPARQL Data Source	
Adding a CSV Data Source	60
Adding a JSON Data Source	68
Adding an XML Data Source	75
Adding a SAS Data Source	81
Adding a Parquet Data Source	
Configuring a CSV or Parquet Source for Incremental Processing	92
Assigning Primary Keys in a Schema	
Adding Foreign Keys	
Onboarding Data with the Automated Workflow	105
Creating a Graphmart from a Data Source	
Adding a Data Source to an Existing Graphmart	
Direct Load Advanced Settings Reference	112
Onboarding or Virtualizing Data with SPARQL Queries	117
Introduction to the GDI	
GDI Concepts and Basic Usage	120
Options for Data Types, Data Linking, and Models	
Advanced Usage by Data Source Type	

GDI Property Reference	
Onboard Unstructured Data	
Unstructured Onboarding Process Overview	
Creating an Unstructured Pipeline	
Running an Unstructured Pipeline	
Pipeline Settings Reference	
Annotator Settings Reference	
External Service Annotator	
Keyword and Phrase Annotator	
Knowledgebase Annotator	
Regex Annotator	
Model	
Model Concepts and Vocabulary	
Managed Model Concepts	
Model Requirements	
Uploading a Model	
Creating a Model	
Editing a Custom Model	
Opening Models in the Editor	
Changing Model Components	
Class Editor Reference	
Property Editor Reference	
Editing a Managed Model	
Downloading a Model	
Defining Resource Templates	

Blend	
Working with Datasets	
Adding an Empty Dataset for an Export Step	
Importing an Existing Dataset (FLDS)	
Creating a Dataset from RDF Files	451
Managing Dataset Editions	
Introduction to Editions	
Creating an Edition	
Modifying an Edition	
Deleting a Saved Edition	
Limiting the Number of Editions in a Dataset	
Creating a Graphmart from a Dataset	
Adding a Dataset to a Graphmart	
Dataset FAQ	
Working with Graphmarts	
Creating a Graphmart	
Copying a Graphmart	
Graphmart Settings Reference	
Creating an Elasticsearch Index from a Graphmart	
Adding Data Layers to Graphmarts	
Creating a New Layer	
Cloning a Layer	
Using Query Contexts	
Defining Execution Conditions	
Advanced Data Access Settings	
Adding Steps to Layers	

Directly Load a Data Source (Direct Load Step)	
Create an Elasticsearch Index (Elasticsearch Indexing Step)	
Take a Snapshot of an Index (Elasticsearch Snapshot Step)	
Export Data to an FLDS (Export Step)	
Load a Dataset from the Catalog (Load Dataset Step)	531
Pre-Compile a Query (Pre-Compile Query Step)	
Create a Reusable Query Template	
Run a Transformation Query (Query Step)	
Infer New Data (RDFS+ Inference Step)	
Validate the Data (Validation Step)	
Construct a View of the Data (View Step)	
Creating Data on Demand Endpoints	
Creating an Auto-Generated Endpoint	
Creating a Custom Endpoint	
Sharing Access to Graphmarts	
Sharing Concepts	
Graphmart FAQ	
Profiling Datasets and Graphmarts	
Generating a Dataset Data Profile	
Generating a Graphmart Data Profile	
Data Profiling Metrics	
Access & Analyze	
Access Data with Hi-Res Analytics Dashboards	
Introduction to Hi-Res Analytics	
Getting Started: Explore and Visualize Your Data	

Create a New Dashboard	637
Explore the Data	
Create Visualizations of the Data	
Working with Dashboards	
Creating a Graphmart Dashboard	
Creating a Network Navigator Dashboard	651
Configuring a Dashboard to Update in Batch Reporting vs. Interactive Mode	
Capturing User-Defined Values in Dashboards	
Working with Lenses	
Creating a Lens	
Cloning a Lens	
Exporting a Lens	
Deleting a Lens	
Working with Filters	
Adding a Cloud Filter	
Adding a Date Range Filter	
Adding a Hierarchy Filter	
Adding a Limit Filter	
Adding a List Filter	
Adding a Numeric Range Filter	
Adding a Presence Filter	
Adding a Quartile Filter	
Adding a Range Slider Filter	
Adding a Relative Time Filter	
Adding a Search Filter	
Adding a Single Select List Filter	

Adding a Types Filter	
Calculating Values in Lenses and Filters	
Combining Data from Multiple Classes	
Searching for Text in Unstructured Documents	
Supported Search Syntax	
Sharing Access to Dashboards and Lenses	
Configure Dashboard or Lens Permissions	
Access Data with the Query Builder	
Running SPARQL Queries in the Query Builder	
Searching for Quads in the Query Builder	
Access Data on Demand Endpoints	
Accessing an Endpoint Programmatically	
Accessing an Endpoint from an Application	
Accessing Data via the OData API	
Downloading the Anzo ODBC and JDBC Drivers	
JDBC Driver Documentation	
OData Reference	
OData URL Conventions	
Supported Query Operators	
Access the SPARQL Endpoint	
Access the HTTP Client Interface	
Share Access to Artifacts	
Version and Migrate Artifacts	
Creating and Restoring Versions of Artifacts	
Create a Backup Version	

Restore a Backup Version	
Exporting an Artifact	
Making Properties Replaceable on Export	
Importing Exported Versions of Artifacts	
SPARQL Best Practices and Query Templates	
SPARQL Best Practices	
SPARQL Query Templates	
Function and Formula Reference	
String Functions	
BUSINESS_ENTITY_EXCLUDER	
CONCATENATE	
CONCATURL	
CONTAINS	
ENCODE_FOR_URI	
ESCAPEHTML	
FIND	
FINDREVERSE	
GROUP_CONCAT	
GROUPCONCAT	
LANG	
LANGMATCHES	
LCASE	
LEFT	
LEN	
LEVENSHTEIN_DIST	
LOWER	

MD5	
MID	
REGEX	
REGEXP_SUBSTR	
REPLACE	
RIGHT	
SEARCH	
SHA1	
SHA224	
SHA256	
SHA384	
SHA512	
STRAFTER	
STRBEFORE	
STRBEFORE	
STRBEFORE	
STRBEFORE STRDT STRENDS STRLANG	
STRBEFORE	
STRBEFORE STRDT STRENDS STRLANG STRLEN STRSTARTS	
STRBEFORE STRDT STRENDS STRLANG STRLEN STRSTARTS STRUUID	
STRBEFORE STRDT STRENDS STRLANG STRLEN STRSTARTS STRUUID SUBSTITUTE	
STRBEFORE STRDT STRENDS STRLANG STRLEN STRSTARTS STRUUID SUBSTITUTE SUBSTR	
STRBEFORE STRDT STRENDS STRLANG STRLEN STRSTARTS STRUUID SUBSTITUTE SUBSTR TOURI	
STRBEFORE STRDT STRENDS STRLANG STRLANG STRLEN STRSTARTS STRUUID SUBSTITUTE SUBSTR TOURI TRIM	
STRBEFORE STRDT STRENDS STRLANG STRLANG STRLEN STRSTARTS STRUUID SUBSTITUTE SUBSTR TOURI TRIM UCASE	

Math Functions	
ABS	
ADD	
AVERAGEIF	
AVERAGEIFS	
AVG	
CEILING	
COS	
DIVIDE	
EQUAL	
EXP	
FACT	
FLOOR	
GE	
GT	
HAMMING_DIST	
HAVERSINE_DIST	
LE	
LN	
LOG	
LOG2	
LT	
MAXVAL	
MINVAL	
MOD	
MULTIPLY	

NOT_EQUAL	
NPV	
NUMERIC-ADD	
NUMERIC-SUBTRACT	
PI	
POWER	
PRODUCT	
QUOTIENT	
RAD	
RAND	
RANDBETWEEN	
ROUND	
ROUNDDOWN	
ROUNDUP	
SIN	
SQRT	
SUBTRACT	
SUM	
SUMIF	
SUMIFS	
SUMPRODUCT	
SUMSQ	
TAN	
Aggregate Functions	
AVERAGEIF	
AVERAGEIFS	

AVG	
CHOOSE_BY_MAX	
CHOOSE_BY_MIN	
COUNT	
COUNT_DISTINCT	
COUNTIF	
COUNTIFS	
GROUP_CONCAT	
GROUPCONCAT	
MAX	
MEDIAN	
MIN	
MODE	
MODEPERCENT	
PERCENTILE_CONT	
MODEPERCENT PERCENTILE_CONT PERCENTILE_DISC	
MODEPERCENT PERCENTILE_CONT PERCENTILE_DISC PRODUCT	
MODEPERCENT PERCENTILE_CONT PERCENTILE_DISC PRODUCT SAMPLE	
MODEPERCENT PERCENTILE_CONT PERCENTILE_DISC PRODUCT SAMPLE STDEV	
MODEPERCENT PERCENTILE_CONT PERCENTILE_DISC PRODUCT SAMPLE STDEV STDEVP	
MODEPERCENT PERCENTILE_CONT PERCENTILE_DISC PRODUCT SAMPLE STDEV STDEV SUM	1082 1083 1084 1084 1084 1085 1085 1085 1086 1086
MODEPERCENT PERCENTILE_CONT PERCENTILE_DISC PRODUCT SAMPLE STDEV STDEVP SUM SUMIF	1082 1083 1083 1084 1084 1085 1085 1085 1086 1086 1087
MODEPERCENT PERCENTILE_CONT PERCENTILE_DISC PRODUCT SAMPLE STDEV STDEV SUM SUMIF SUMIFS	1082 1083 1084 1084 1084 1085 1085 1085 1086 1086 1087 1088
MODEPERCENT PERCENTILE_CONT PERCENTILE_DISC PRODUCT SAMPLE STDEV STDEV SUM SUMIF SUMIFS SUMPRODUCT	1082 1083 1083 1084 1084 1085 1085 1085 1086 1086 1087 1088 1088
MODEPERCENT PERCENTILE_CONT PERCENTILE_DISC PRODUCT SAMPLE STDEV STDEV SUM SUMIF SUMIFS SUMPRODUCT SUMSQ	1082 1083 1083 1084 1084 1085 1085 1085 1086 1086 1087 1088 1088 1088

VARP	
WEIGHTEDAVERAGE	
Date and Time Functions	
DATE	
DATEPART	
DATETIME (or xsd:dateTime)	
DAY	
DAYSFROMDURATION	
DUR_TO_MILLIS	
DURATION	
DURATIONFORMAT	
DURATIONPERIODFORMAT	
FORMATDATE	
HOUR	
MASKEDDATETIME	1101
MILLIS	
MINUTE	
MONTH	
NOW	
NOWMILLIS	
PARSEDATETIME	
SECOND	
TIME	
TIMEPART	
TIMEVALUE	
TIMEZONE	

TODAY	
ΤΖ	
WEEKDAY	
WEEKNUM	
xsd:date	
YEAR	
YEARMONTH	
Casting Functions	
BNODE	
BOOLEAN	
BYTE	
CONCATURL	
DATATYPE	
DATETIME (or xsd:dateTime)	
DATEVALUE	
DECIMAL	
DOUBLE	
DURATION	
DURATIONFORMAT	
ENCODE_FOR_URI	
FLOAT	
FORMATDATE	
FORMATFRACTION	
FORMATNUMBER	
INT	
INTEGER	

LONG	
PARSEDATETIME	
RAD	
SERIALIZE	
SHORT	
STR	
TEXT	
TIMEVALUE	
TOURI	
UUID	
xsd:date	
Logical Functions	
AND	
BOUND	
CASE	
COALESCE	
EQUAL	
IF	
IFERROR	
IN	
NOT	
NOT_EQUAL	
NOT_IN	
OR	
PARTITIONINDEX	
SAMETERM	

UNBOUND	
Informational or Testing Functions	
CONTAINS	
GE	
GT	
ISBLANK	
ISDATATYPE	
ISERROR	
ISIRI	1147
ISLITERAL	1147
ISNUMERIC	
ISURI	
LANG	
LANGMATCHES	
LE	
LOCALNAME	
LT	1151
METADATAGRAPHURI	
NAMESPACE	
SAMETERM	
Hash Functions	
MD5	
SHA1	
SHA224	
SHA256	
SHA384	

SHA512	
Window Aggregate and Ranking Functions	
WINDOW_AVG	
WINDOW_COUNT	
WINDOW_MAX	
WINDOW_MIN	
WINDOW_NTILE	
WINDOW_PERCENTILE	
WINDOW_PERCENTILE_CONT	
WINDOW_PERCENTILE_DISC	
WINDOW_PRODUCT	
WINDOW_QUARTILE	
WINDOW_SUM	
Develop	
Develop Anzo Rest API	<b>1175</b> 1176
Develop Anzo Rest API Introduction to the API	<b>1175</b> 
Develop Anzo Rest API Introduction to the API Viewing the API Documentation	<b>1175</b> 1176 
Develop Anzo Rest API Introduction to the API Viewing the API Documentation Enabling Cross-Origin Resource Sharing	<b>1175</b> 1176 1177 1177 1180 1181
Develop Anzo Rest API Introduction to the API Viewing the API Documentation Enabling Cross-Origin Resource Sharing Step Type Schemas	<b>1175</b> 1176 1177 1180 1181 1183
Develop Anzo Rest API Introduction to the API Viewing the API Documentation Enabling Cross-Origin Resource Sharing Step Type Schemas Direct Load Step	<b>1175</b> 1176 1177 1180 1181 1183 1183
Develop Anzo Rest API Introduction to the API Viewing the API Documentation Enabling Cross-Origin Resource Sharing Step Type Schemas Direct Load Step Elasticsearch Indexing Step	<b>1175</b> 1176 1177 1180 1181 1183 1183 1183
Develop Anzo Rest API Introduction to the API Viewing the API Documentation Enabling Cross-Origin Resource Sharing Step Type Schemas Direct Load Step Elasticsearch Indexing Step Elasticsearch Snapshot Step	<b>1175</b> 1176 1177 1177 1180 1181 1183 1183 1183 1187 1191
Develop Anzo Rest API Introduction to the API Viewing the API Documentation Enabling Cross-Origin Resource Sharing Step Type Schemas Direct Load Step Elasticsearch Indexing Step Elasticsearch Snapshot Step Export Step	<b>1175</b> 1176 1177 1177 1180 1181 1183 1183 1183 1187 1191
Develop Anzo Rest API Introduction to the API Viewing the API Documentation Enabling Cross-Origin Resource Sharing Step Type Schemas Direct Load Step Elasticsearch Indexing Step Elasticsearch Snapshot Step Export Step Load Dataset Step	<b>1175</b> 1176 1177 1180 1181 1183 1183 1187 1191 1202
Develop Anzo Rest API Introduction to the API Viewing the API Documentation Enabling Cross-Origin Resource Sharing Step Type Schemas Direct Load Step Elasticsearch Indexing Step Elasticsearch Snapshot Step Export Step Load Dataset Step Pre-Compile Query Step	<b>1175</b> 1176 1177 1177 1180 1181 1183 1183 1183 1187 1191 1195 1202 1207

Q	Query Step	1215
R	RDFS+ Inference Step	1219
Т	emplated Step	1222
V	/alidation Step	1227
Anzo J	Java SDK	1233

# About This Doc

This document provides guidance on onboarding, modeling, blending, and accessing data in Anzo.

## Тір

You can view the contents of this guide as well as release notes, getting started, deployment, and administration documentation online at <u>docs.cambridgesemantics.com</u>. You can also find PDF versions of the getting started, deployment, and administration documentation <u>here</u>.

The following list introduces the sections in this guide.

- Onboard & Virtualization: Provides instructions for connecting to data sources and onboarding or virtualizing structured and semi-structured data as well as details about creating pipelines to onboard unstructured data.
- Model: Provides conceptual information about data models, describes model requirements, and includes instructions for creating, editing, uploading, and downloading models.
- Blend: Includes information about the Datasets catalog and instructions on creating and managing datasets and editions. Also provides instructions on creating graphmarts, data layers, steps, Data on Demand endpoints, and managing graphmart permissions, as well as creating data profiles for datasets and graphmarts.
- Access & Analyze: Discusses the options for accessing and analyzing your data, including creating Hi-Res Analytics dashboards, using the Query Builder to run SPARQL queries, and accessing the SPARQL endpoint or Data on Demand endpoints. Also includes information about sharing, versioning, and migrating artifacts.
- Develop: Includes information for developers about the Anzo REST API and Java SDK.

# **Onboard & Virtualization**

The topics in this section provide instructions for onboarding unstructured data and onboarding or virtualizing structured or semi-structured data. For instructions on importing files that are in RDF format (Turtle or N-Triple), see Creating a Dataset from RDF Files.

In this section:	
Onboard Structured Data	21
Onboard Unstructured Data	

# **Onboard Structured Data**

For structured and semi-structured data sources—databases, HTTP REST endpoints, CSV, JSON, XML, Parquet, and SAS files, and raw data—there are two ways to onboard and/or virtualize the data:

# **Automated Direct Data Load**

Data from databases and CSV, JSON, XML, Parquet, and SAS files can be onboarded to Anzo via the automated direct data load workflow. This workflow follows an extract, load, and transform (ELT) process to ingest data. In the ELT workflow, data sources are extracted and loaded to graphmarts. Data layers and Direct Load Steps with SPARQL queries are automatically generated to transform and blend the data to analytics-ready knowledge graphs. The AnzoGraph Graph Data Interface (GDI) Java plugin is used to connect to the sources, create a model, and generate the data layer queries. For more information about this workflow, see Onboarding Data with the Automated Workflow.

# Load or Virtualization with SPARQL Queries

Raw data and data from databases, HTTP REST endpoints, and CSV, JSON, XML, Parquet, and SAS files can be onboarded or virtualized by invoking the Graph Data Interface (GDI) with manually written SPARQL queries. The GDI is extremely flexible, allowing you to connect directly to sources and control all aspects of the extract, load, and transform process. You can onboard data into Anzo by adding a Direct Load Step query. Or you can create a virtual graph by adding a View step query. By creating views, you can access the source data exactly when it is needed, without requiring you to ingest all of the data into Anzo up front. For more information about this workflow, see Onboarding or Virtualizing Data with SPARQL Queries.

## Note

Whether you plan to use the automated workflow or manually invoke the GDI, the first step in the process is to connect your data sources to Anzo and onboard the schemas. Start with Adding Data Sources.

# **Adding Data Sources**

This topics in this section provide instructions for connecting to structured and semi-structured data sources and working with schemas.

## Tip

For instructions on onboarding RDF files (Turtle or N-Triple files) to Anzo, see Creating a Dataset from RDF Files.

In this section:	
Connecting to a Database	23
Defining a Database Schema	
Adding an HTTP or SPARQL Data Source	
Adding a CSV Data Source	60
Adding a JSON Data Source	
Adding an XML Data Source	
Adding a SAS Data Source	81
Adding a Parquet Data Source	88
Configuring a CSV or Parquet Source for Incremental Processing	
Assigning Primary Keys in a Schema	
Adding Foreign Keys	100

## **Connecting to a Database**

The topics in this section provide instructions on configuring connections to the database types that Anzo and AnzoGraph support by default.

## Tip

To connect to other types of databases, first add JDBC drivers to Anzo (see <u>Uploading a</u> <u>Plugin</u> in the Administration Guide) and AnzoGraph (see <u>Deploy Drivers for Custom Database</u> <u>Sources</u> in the Deployment Guide). Then you can follow any of the instructions in this section to configure the connection. All database connections are similar, but settings may vary based on the JDBC driver.

- Connecting to a Databricks Source
- Connecting to a DB2 Source
- Connecting to an MSSQL Source
- Connecting to an Oracle Source
- Connecting to a PostgreSQL Source
- Connecting to a Snowflake Source
- Connecting to a Sybase Source

**Connecting to a Databricks Source** 

Follow the instructions below to connect to a Databricks database.

	Data Sources S	chemas						
<b>Y</b>	<b>L</b> Search	= • <b>ii</b>	≡			Ad	d Data Sou	irce 👻
•	Title $\uparrow$ Description	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
	Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
	🗇 Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	:
	∃ DB-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
	E Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
	E Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
	Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
					Rows per page:	25 ¥ 1-6 of	6 <	>

Click the Add Data Source button, select Database, and then select Databricks Database
 Data Source. The Create screen is displayed:

Create Databricks Database Data Source		
Title *		
Description		
User*		
Password *		0
Server		
Database		
HTTP Path *		
	CANCEL	SAVE

- 3. Complete the required fields and configure any optional settings. The list below describes each setting.
  - Title: The name to give to this data source connection.
  - **Description**: An optional description of the connection.

- **User**: The user name to use for logging in to the database.
- **Password**: The password for the user name.
- Server: The host name or IP address for the source.
- **Database**: An optional partition that contains the data to onboard.
- HTTP Path: The Databricks compute resources URL.

Create Databricks Database Data Sourc	e	
Title *		
Databricks		
Burn de la companya d		
Description		
llear *		
token		
Password *		
		O
Server		
dbc-e7fc434a-24aa.cloud.databricks.com		
Database		
HTTP Path *		
/sql/1.0/warehouses/45e050b8db136b0e		
		0.015
	CANCEL	SAVE

4. Click **Save** to save the data source connection. The **Tables** tab is displayed.

Databricks			Create	Graphmart
Overview	Tables	Versions	Discussion	Sharing >
		•		
		•		
	Please import a schema, o	or create a schema fro	m a custom query.	
	Create Schemas fr	om Query	rt Schemas	
	oreate denemas in		- Contentino	

 Before proceeding to select or create a schema, you may want to test connectivity to the database. To do so, click the **Overview** tab and then click the **Test Connection** button at the bottom of the screen.

2	Databricks			(	Create Graphmart	D
	Overview	Tables	Versions	Discussion	Sharing	>
Description None User token Password				General Type Creator Updated Released	Database System Administrator 5 minutes ago 5 minutes ago	Î
dbc-e7fc Database None	434a-24aa.cloud.dat	abricks.com		http://cambridge	semantics.com/Data	
HTTP Path /sql/1.0/ Test Co	warehouses/45e050	b8db136b0e				
				Tags None		Ŧ

If the connection fails, adjust the data source details until the connection is successful.

## Note

Only the connection between Anzo and the database is tested. Connectivity is not tested between AnzoGraph and the database. Cambridge Semantics recommends that you perform a manual test to ensure that AnzoGraph can also connect to the database.

After connecting to the source, the next step is to add one or more schemas. See Defining a Database Schema for instructions.

#### **Connecting to a DB2 Source**

Follow the instructions below to connect to a DB2 database.

Data Sources	Schemas						
Y Q Search	<b>₹</b> • <b>П</b>				Ad	d Data So	urce 👻
Title 🛧 Description	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
E Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	:
∃ DB-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
E Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
E Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
				Rows per page:	25 👻 1-6 of	6 <	>

 Click the Add Data Source button, select Database, and then select DB2 Database Data Source. The Create screen is displayed:

Create DB2 Database Data Source		
Title *		
Description		
User*		
Password *		ø
Server		
Database		
	CANCEL	SAVE

- 3. Complete the required fields and configure any optional settings. The list below describes each setting.
  - Title: The name to give to this data source connection.
  - **Description**: An optional description of the connection.
  - **User**: The user name to use for logging in to the database.

- **Password**: The password for the user name.
- Server: The host name or IP address for the source.
- **Database**: An optional partition that contains the data to onboard.

Create DB2 Database Data Source		
Yale 4		
DB2		
Description		
User*		
jscott		
Password *		
		0
Server		
10.22.20.19		
Database		
	CANCEL	SAVE

4. Click Save to save the data source connection. The Tables tab is displayed.

DB2			Create	Graphmart
Overview	Tables	Versions	Discussion	Sharing >
	•			
	Please import a schema,	or create a schema fror	n a custom query.	
	Create Schemas	from Query Impor	t Schemas	

5. Before proceeding to select or create a schema, you may want to test connectivity to the database. To do so, click the **Overview** tab and then click the **Test Connection** button at the bottom of the screen.

1	DB2 Not Versioned 🗹			Creat	e Graphmart
	Overview	Tables	Versions	Discussion	Sharing >
Description None User jscott Password ****** Server	n			General Type Datab Creator Syster Updated a few Released a few	ase n Administrator seconds ago seconds ago
10.22.20 Database None Test C	0.19 connection			http://cambridgesemantie Tags None	s.com/DB2

If the connection fails, adjust the data source details until the connection is successful.

## Note

Only the connection between Anzo and the database is tested. Connectivity is not tested between AnzoGraph and the database. Cambridge Semantics recommends that you perform a manual test to ensure that AnzoGraph can also connect to the database.

After connecting to the source, the next step is to add one or more schemas. See Defining a Database Schema for instructions.

**Connecting to an MSSQL Source** 

Follow the instructions below to connect to a SQL Server or MSSQL database.

	Data Sources	Schemas						
<b>Y</b> Q	Search	= •				Ad	d Data Sou	urce 👻
•	Title	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
	Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
	🗇 Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	:
	∃ 08-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
	Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
	Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
	Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
					Rows per page:	25 ¥ 1-6 of	6 <	>

 Click the Add Data Source button, select Database, and then select MSSQL Database Data Source. The Create screen is displayed:

Create MSSQL Database Data Source		
Title *		
Description		
User *		
Password *		ø
Server		
Database		
Domain		
	CANCEL	SAVE

- 3. Complete the required fields and configure any optional settings. The list below describes each setting.
  - Title: The name to give to this data source connection.
  - **Description**: An optional description of the connection.

- **User**: The user name to use for logging in to the database.
- **Password**: The password for the user name.
- Server: The host name or IP address for the source.
- **Database**: An optional partition that contains the data to onboard.
- **Domain**: When using NTLM authentication, this is the domain to authenticate against.

Create MSSQL Database Data Source		
Title *		
MSSQL		
Description		
User*		
\$8		
Password *		
		0
Server		
ms-sql-db.cambridesemantics.com		
Database		
NORTHWIND		
Domain		
	CANCEL	CAVE
	GANCEL	SAVE

4. Click Save to save the data source connection. The Tables tab is displayed.



 Before proceeding to select or create a schema, you may want to test connectivity to the database. To do so, click the **Overview** tab and then click the **Test Connection** button at the bottom of the screen.

					Create Graphmart	Д
	Overview	Tables	Versions	Discussion	Sharing	>
Description None User Sa Password	on 1			General Type Creator Updated Released	Database System Administrator a minute ago a minute ago	
Server ms-sql- Database NORTH	db.cambridesemantic	s.com		http://cambridgese	mantics.com/MSS	-
Domain None Test (	Connection			Tags None		

If the connection fails, adjust the data source details until the connection is successful.

## Note

Only the connection between Anzo and the database is tested. Connectivity is not tested between AnzoGraph and the database. Cambridge Semantics recommends that you perform a manual test to ensure that AnzoGraph can also connect to the database.

After connecting to the source, the next step is to add one or more schemas. See Defining a Database Schema for instructions.

#### **Connecting to an Oracle Source**

Follow the instructions below to connect to an Oracle database.

Data Sources S	chemas						
Y Q Search	<b>F</b> • <b>II</b>	≡			Ad	ld Data Sou	irce 👻
Title $\uparrow$ Description	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
E Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	П	:
∃ DB-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
E Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
E Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
				Rows per page:	25 👻 1-6 of	6 <	>

 Click the Add Data Source button, select Database, and then select Oracle Database Data Source. The Create screen is displayed:

Create Oracle Database Data Source		
Title *		
Description		
User *		
Password *		•
Server		
	CANCEL	SAVE

- 3. Complete the required fields and configure any optional settings. The list below describes each setting.
  - Title: The name to give to this data source connection.
  - Description: An optional description of the connection.
  - **User**: The user name to use for logging in to the database.

- **Password**: The password for the user name.
- Server: The host name or IP address for the source.

Create Oracle Database Data Source		
Title *		
Oracle		
Description		
User*		
admin		
Password *		
		0
Server		
db-oracle.cambridgesemantics.com		
	CANCEL	SAVE

4. Click Save to save the data source connection. The Tables tab is displayed.



 Before proceeding to select or create a schema, you may want to test connectivity to the database. To do so, click the **Overview** tab and then click the **Test Connection** button at the bottom of the screen.

	Oracle			(	Create Graphmar	t 🔲
	Overview	Tables	Versions	Discussio	on Sh	aring >
Description None User admin Password ****** Server	1			General Type Creator Updated Released	Database System Administrat a minute ago a minute ago	or
db-oracle Test C	e.cambridgesemanti	cs.com		http://cambridge Tags None	semantics.com/Oracl	

If the connection fails, adjust the data source details until the connection is successful.

## Note

Only the connection between Anzo and the database is tested. Connectivity is not tested between AnzoGraph and the database. Cambridge Semantics recommends that you perform a manual test to ensure that AnzoGraph can also connect to the database.

After connecting to the source, the next step is to add one or more schemas. See Defining a Database Schema for instructions.

Connecting to a PostgreSQL Source

Follow the instructions below to connect to a PostgreSQL database.

Dat	a Sources S	chemas						
T Q Sea	arch	F •	≡			Ad	d Data Sou	irce 👻
🗌 🕶 Title	↑ Description	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
8.4	Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
	Datafox	JSON Deta S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	:
80	AD-80	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
E 1	Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
8 1	Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
8 1	Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
					Rows per page:	25 🛩 1-6 of	6 <	>

Click the Add Data Source button, select Database, and then select PostgreSQL Database
 Data Source. The Create screen is displayed:

Create PostgreSQL Database Data Source		
Title *		
Description		
User*		
Password *		0
Server		
Database		
	CANCEL	SAVE

- 3. Complete the required fields and configure any optional settings. The list below describes each setting.
  - Title: The name to give to this data source connection.
  - **Description**: An optional description of the connection.
  - **User**: The user name to use for logging in to the database.
- **Password**: The password for the user name.
- Server: The host name or IP address for the source.
- **Database**: An optional partition that contains the data to onboard.

The image below shows an example of a completed configuration.

Create PostgreSQL Database Data Source		
Trile*		
PostgreSQL		
Description		
admin		
Password *		
		0
Server		
db-postgres.cambridgesemantics.com		
Database		
	CANCEL	CAVE
	GANGEL	SAVE

4. Click Save to save the data source connection. The Tables tab is displayed.

			Create	Graphmart 🔲
Overview	Tables	Versions	Discussion	Sharing >
	Please import a schema	, or create a schema fror	n a custom query.	
	Create Schemas	from Query Impor	t Schemas	

 Before proceeding to select or create a schema, you may want to test connectivity to the database. To do so, click the **Overview** tab and then click the **Test Connection** button at the bottom of the screen.

				Cre	ate Graphmart
	Overview	Tables	Versions	Discussion	Sharing >
Description None User admin Password				General Type Data Creator Syst Updated a mi Released a mi	abase tem Administrator inute ago inute ago
db-postg	res.cambridgeseman	tics.com		http://cambridgesemar	ntics.com/Post
None Test Co	onnection				
				Tags None	÷

If the connection fails, adjust the data source details until the connection is successful.

### Note

Only the connection between Anzo and the database is tested. Connectivity is not tested between AnzoGraph and the database. Cambridge Semantics recommends that you perform a manual test to ensure that AnzoGraph can also connect to the database.

After connecting to the source, the next step is to add one or more schemas. See Defining a Database Schema for instructions.

**Connecting to a Snowflake Source** 

Follow the instructions below to connect to a Snowflake database.

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

Data S	ources	Schemas						
Y Q Search	1	F •	≡			A	ld Data Sou	irce 👻
🗌 👻 Title	↑ Description	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
🗄 Airlin	es	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
🗇 Data	fox	JSON Deta S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	:
∃ 084	A.	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
E Mov	es	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
E Tick	ets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
🗄 Тор	Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
					Rows per page:	25 👻 1-6 of	6 <	>

Click the Add Data Source button, select Database, and then select Snowflake Database
 Data Source. The Create screen is displayed:

Create Snowflake Database Data Source		
Title *		
Description		
User*		
Password *		ø
Server		
Database		
	CANCEL	SAVE

- 3. Complete the required fields and configure any optional settings. The list below describes each setting.
  - Title: The name to give to this data source connection.
  - Description: An optional description of the connection.
  - **User**: The user name to use for logging in to the database.

- **Password**: The password for the user name.
- Server: The host name or IP address for the source.
- Database: An optional partition that contains the data to onboard.

The image below shows an example of a completed configuration.

Create Snowflake Database Data Source		
Teles		
Snowflake		
Description		
User*		
admin		
Password *		
		0
Server		
https://zdstbou-sp69514.snowflakecomputing.com		
Database		
SNOWFLAKE_SAMPLE_DATA		
	0.1105	0.015
	CANCEL	SAVE

4. Click Save to save the data source connection. The Tables tab is displayed.



5. Before proceeding to select or create a schema, you may want to test connectivity to the database. To do so, click the **Overview** tab and then click the **Test Connection** button at the bottom of the screen.

	Snowflake				Create Graphm	art	D
	Overview	Tables	Versions	Discussi	on	Sharing	>
Description None User admin Password				General Type Creator Updated Released	Database System Administ 2 minutes ago 2 minutes ago	rator	Î
https://z Database SNOWFL Test C	dstbou-sp69514.snov AKE_SAMPLE_DATA	wflakecomputing.com		http://cambridg Tags None	esemantics.com/Sno	0	

If the connection fails, adjust the data source details until the connection is successful.

## Note

Only the connection between Anzo and the database is tested. Connectivity is not tested between AnzoGraph and the database. Cambridge Semantics recommends that you perform a manual test to ensure that AnzoGraph can also connect to the database.

After connecting to the source, the next step is to add one or more schemas. See Defining a Database Schema for instructions.

**Connecting to a Sybase Source** 

Follow the instructions below to connect to a Sybase database.

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

Data Sources	Schemas						
Y Q Search	<b>₹</b> • <b>П</b>				Ad	d Data So	urce 👻
Title 🛧 Description	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
E Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	:
∃ DB-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
E Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
E Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
				Rows per page:	25 👻 1-6 of	6 <	>

 Click the Add Data Source button, select Database, and then select Sybase Database Data Source. The Create screen is displayed:

Create Sybase Database Data Source		
Title *		
Description		
User*		
Password *		0
Server		
Database		
	CANCEL	SAVE

- 3. Complete the required fields and configure any optional settings. The list below describes each setting.
  - Title: The name to give to this data source connection.
  - **Description**: An optional description of the connection.
  - **User**: The user name to use for logging in to the database.

- **Password**: The password for the user name.
- Server: The host name or IP address for the source.
- **Database**: An optional partition that contains the data to onboard.

The image below shows an example of a completed configuration.

Create Sybase Database Data Source		
Title *		
Sybase		
Description		
User*		
admin		
Password *		
		0
Server		
sybase-db.cambridgesemantics.com		
Database		
	CANCEL	SAVE
	or another	Sec. 1

4. Click Save to save the data source connection. The Tables tab is displayed.

1	Sybase			Create	Graphmart 🔲
	Overview	Tables	Versions	Discussion	Sharing >
			•		
		Please import a schema	, or create a schema fro	m a custom query.	
		Create Schemas	from Query Impor	t Schemas	

 Before proceeding to select or create a schema, you may want to test connectivity to the database. To do so, click the **Overview** tab and then click the **Test Connection** button at the bottom of the screen.

	Sybase				Create Graphmart	Д
_	Overview	Tables	Versions	Discussi	on Sharing	, <b>)</b>
Description None User admin Password ******* Server sybase-C Database None Test C	lb.cambridgesemant	ics.com		General Type Creator Updated Released http://cambridg	Database System Administrator a few seconds ago a few seconds ago gesemantics.com/Syba	
				Tags None		-

If the connection fails, adjust the data source connection details until the connection is successful.

### Note

Only the connection between Anzo and the database is tested. Connectivity is not tested between AnzoGraph and the database. Cambridge Semantics recommends that you perform a manual test to ensure that AnzoGraph can also connect to the database.

After connecting to the source, the next step is to add one or more schemas. See Defining a Database Schema for instructions.

# **Defining a Database Schema**

The schema defines the source data to onboard. There are multiple options available for defining a database schema. You can import a predefined schema from the database, you can write a static SQL query that defines the data to onboard, or, if you want to import data incrementally, you can write an incremental SQL query that ingests a subset of the data.

### Note

You can import or create up to 5 schemas per database data source. To include more than 5 schemas, create another data source for the additional schemas.

# Tip

By default, Anzo is configured to exclude Views from the list of available Schemas to import. For information about including Views as tables that can be imported, see <u>Including Views as</u> <u>Schemas for Database Data Sources</u> in the Administration Guide.

- Import a Predefined Schema
- Create a Schema from an SQL Query
- Create an Incremental Schema

### Import a Predefined Schema

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

	Data Sources S	chemas						
<b>Y</b>	<b>X</b> Search	<b>F</b> •	≡			A	dd Data Sou	irce 👻
•	Title	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
	Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
	🗇 Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	:
	≡ DB-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
	Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	8
	E Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
	Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	8
					Rows per page:	25 👻 1-6 of	6 <	>

2. Click the database for which you want to import a schema. Anzo displays the Tables tab for the source. For example, the image below shows the Tables tab before any schemas have been added:

DB-QA				Create Graphmart	D
Overview	Tables	Versions	Discussion	Sharing	
	Please import a sche	ma, or create a schem	a from a custom query.		
	Create Schem	has from Query	mport Schemas		

3. Click the **Import Schemas** button. Anzo displays the Import Schemas dialog box, which lists any predefined schemas in the database. For example:



# Note

If you do not see a schema that you expect to see, make sure that you have the appropriate permissions to access to the data source.

4. To import entire schemas, select the checkbox next to each schema name that you want to import. If you want to import a subset of the tables in a schema, expand the schema and select the checkbox next to each table that you want to include. For example:



5. When you have finished selecting schemas, click **OK**. Anzo imports the selected schemas and lists them on the Tables tab. You can expand a schema to view its tables. Selecting a row in the schema displays the sample data on the right side of the screen.

## Important

The automated data load workflow ignores all changes that are made to the schema on the Tables screen—except for changes to primary and foreign keys. For example, if you edit a column heading to change its semantic type, that change is disregarded when the graphmart is created. Only the original type from the data source is considered. If you add or change primary and foreign keys on the Tables screen, however, the automated data load workflow will retain those changes.

	MySQL DB										Create Grap	hmart 🔲
	Overview	Tables	Ver	sions	Discust	sion	Sharing					
Scher	nas		*	:: ←	Alphabo	etical list of p	roducts					
Cre	ate Schemas Fro	m Query	Import S	chemas	Creator	System Administr	ator Last Mod	ified Date 03	28/2023 Colum	Count 11	Row Count	Not Calculated
<b>Q</b> Se	arch	<b>- -</b>			Sample	e Data	Foreign Keys	Schem	a Sharing			
•	Schema Name	Table Name 🋧 Type	Actions	1	Producti0	ProductNa 7' String	SupplierID % Int	Category(D	QuantityPer 7' String	UnitPrice	UnitsinStock	UnitsOnOrd
	northwind		:	~ I	70	Outback Lager	7	1	24-355 milbottk	150000	15	10
		Alphabetical L.	:		1	Chai	1	1	10 boxes x 20 ba	18,0000	39	0
		Categories	:		38	Cte de Blaye	18	1	12-75 cl bottles	263.5000	17	0
		Category Sal	:		2	Chang	1	1	24-12 oz botties	19.0000	17	40
		Current Prod	:		67	Laughing Lumbe	16	1	24-12 oc bottles	140000	52	0
-				∎ → <sup>*</sup>	43	Ipph Coffee	20	1	16-500 g tins	46/0000	17	10

For information about creating or changing primary keys and foreign keys, see Assigning Primary Keys in a Schema and Adding Foreign Keys.

#### Create a Schema from an SQL Query

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

Data Sources	Schemas						
Y Q Search	= • 🖬 :				-	dd Data Sou	rce 👻
Title $\uparrow$ Description	Type Sc	hema	Creator	Created Date	Updated Date	Actions	÷
Airlines	CSV Data Sou Ai	rlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
🗇 Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	8
E DB-QA	Database Dat Er	mployees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Movies	CSV Data Sou M	lovie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	8
E Tickets	CSV Data Sou Ti	ckets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Top Movies	CSV Data Sou M	lovies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
			F	tows per page:	25 <del>v</del> 1-6 o	f6 <	>

2. Click the database for which you want to create a schema. Anzo displays the Tables tab for the source. For example, the image below shows the Tables tab before any schemas have been created:

DB-QA				Create Graphmart	Д
Overview	Tables	Versions	Discussion	Sharing	
	Please import a sche	ema, or create a schem	a from a custom query.		
	Create Schen	has from Query	mport Schemas		

3. Click the **Create Schemas From Query** button. Anzo displays the Create Schemas dialog box:

Create Schemas		20
		Á
Schema Name *		- 1
Table Name *		
The name of the schema table		
1 Database Query *		
	CANCEL	SAVE

- 4. In the Create Schemas dialog box, specify a name for this schema in the Schema Name field.
- 5. In the Table Name field, specify a name for the table in the schema that the query will create.
- 6. Type the SQL statement in the text box. The statement can include any functionality that the source database supports. Anzo does not validate the SQL.

#### Note

If the SQL query requires quotes around values, such as '2010-01-01' or 'TestValue', make sure that you use single quotes ('). For example:

SELECT \* FROM Movies WHERE production day='2021-08-01'

Including double quotes (") in a schema query results in an error when the query is run.

The following example creates a schema named employees. A table named all\_employees will be created in the schema, and the table is created from the following SQL query:

```
SELECT EmployeeID, FirstName, LastName, Title, Salary, BirthDate, HireDate,
Region, Country
FROM northwind.Employees
WHERE EmployeeID
```

Create Schemas	20	
Schema Name * employees		ĺ
Table Name* all_employees		
The name of the schema table Database Query* 1 SELECT EmployeeID, FirstName, LastName, Title, Salary, BirthDate, HireDate, Region, Country 2 FAOH northwind.Employees 3 WHERE EmployeeID 4	r	
CANCEL SJ	AVE	

7. Click **Save** to save the query. Anzo creates the new schema and adds it to the list of schemas on the Tables screen. Selecting the schema displays sample data on the right side of the screen.

### Important

The automated data load workflow ignores all changes that are made to the schema on the Tables screen—except for changes to primary and foreign keys. For example, if you edit a column heading to change its semantic type, that change is disregarded when the graphmart is created. Only the original type from the data source is considered. If you add or change primary and foreign keys on the Tables screen, however, the automated data load workflow will retain those changes.

	DB-QA						Pierce de		Photos				Create Gra	phmort []
	Overview		ables	Ver	sions		Discussion	1	sharing					
Schen	nas				* C (	÷	all_empk	oyees						
Crea	te Schemas Fro	m Query		Import	Schemas		Creator S	ystem Adminis	trator Last Mod	dified Date 10/	19/2022	Column Count 9	Row Count	Not Calculated
Q, Sea	irch	F	٣				Sample	Data	Foreign Keys	Schema	Sharing			
•	Sohema Name	Table Name ↑	Type	Actions	1		EmployeeID	FirstName	LastNorme	Title	Salary	BirthDate	HireDate	Region
	employees			:	~		"Nate	T String	T String	7 String	Float	ED Date time	ED Date time	T String
							6	Michael	Suyama	Sales Representa	2004.07	1965-07-02.00.0	1993-10-17 00	0
		alLempioyees	Manual				5	Steven	Buchanan	Sales Manager	1744.21	1955-03-04 00.0	1993-10-17 00	01
	northwind			:	<		10	Martin	Mark:	Sales Manager	2500.06	1957-04-05-00.0	1995-07-06-00	00
							1	Nancy	Develo	Sales Representa	2954.55	1948-12-08 00 0	1992-05-01-00	OC VIA
							7	Robert	King	Sales Representa	1991.55	1960-05-29 00 0	1994-01-02-00	00
4						×	2	Andrew	Puller	Vice President, S	2254,49	1955-02-19:00:0	1992-08-14-00	00 WA
							1							

- 8. If you want to create additional tables in the schema, follow these steps:
  - a. Click the menu icon (i) in the Actions column for the schema name and select Add
    - Table. For example:

 Schema Name	Table Name ↑	Туре	Actions		:
employees			:	<	
northwind		Add Table			
			Reload Schema		
			Delete		

The Create New Table dialog box is displayed.

Create New Table		20
Table Name *		
The name of the schema table		
Enter Query Here		
	CANCEL	SAVE

- b. In the Create New Table dialog box, specify a name for the new table in the **Table Name** field.
- c. In the **Schema Query** field, write the SQL query that defines the data for the new table.
- d. Click **Save** to add the table to the schema and return to the Tables screen.

For information about creating or changing primary keys and foreign keys, see Assigning Primary Keys in a Schema and Adding Foreign Keys.

#### **Create an Incremental Schema**

Follow the instructions below to create a schema by writing an SQL query that defines a subset of the data to onboard in increments.

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

Data Sources	Schemas						
Y Q Search	= •	=			Ad	d Data So	urce 👻
Title 🛧 Description	n Type	Schema	Creator	Created Date	Updated Date	Actions	÷
Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
🗇 Datafox	JSON Deta S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	:
E DB-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
				Rows per page:	25 👻 1-6 of	б <	>

2. Click the database for which you want to create an incremental schema. Anzo displays the Tables tab for the source. For example, the image below shows the Tables tab before any schemas have been created:

DB-QA				Create Graphmart
Overview	Tables	Versions	Discussion	Sharing
	Please import a sch	ema, or create a schem	a from a custom query.	
	Create Scher	mas from Query	mport Schemas	

3. Click the **Create Schemas From Query** button. Anzo displays the Create Schemas dialog box:

Create Schemas		20
		A
Schema Name *		- 1
Table Name *		
The name of the schema table		
1 Database Query *		
		- 1
	CANCEL	SAVE

- 4. In the Create Schemas dialog box, specify a name for this schema in the Schema Name field.
- 5. In the **Table Name** field, specify a name for the table in the schema that the query will create.
- 6. At the bottom of the screen, enable the **Include increment data** option by sliding the slider to the right. Anzo displays additional settings:

Create Schemas	- 22	
		^
Include increment data		
Incremental Column Name *		I
Value *		
Comparator *	~	
CANCEL	SAVE	

- 7. Populate the following fields so that you can use the values as a guide for writing the schema query:
  - Incremental Column Name: The source column whose value will be used to increment the data.

• **Value**: The value in the column to use as the stopping point for the first import process and the starting point for the next import.

### Note

Do not include quote characters in the **Value** field. If the SQL query requires quotes around values, such as '2010-01-01' or 'TestValue', include the quotes around the {INCREMENTVALUE} parameter in the query and not in the Value field. For example, if the value to increment on is '2010-01-01', specify **2010-01-01** in the Value field and add the quotes to the query like the following example:

SELECT \* FROM Orders WHERE OrderData > '{INCREMENTVALUE}'

In addition, make sure that you use single quotes (') in schema queries. Including double quotes (") in a schema query results in an error when the query is run.

- Comparator: The operator to use for comparing source values against the value above.
- 8. In the query text field, type the SQL statement that will target the appropriate source data. The WHERE clause must include the incremental column name, the comparison operator, and an INCREMENTVALUE parameter. This parameter is substituted with the Value at runtime. For example, in the query below, the incremental column name is EmployeeID, the comparator is > (greater than), and the {INCREMENTVALUE} parameter is specified after the comparator. {INCREMENTVALUE} is replaced with the value in the Value field at runtime:

```
SELECT EmployeeID, FirstName, LastName, Title, Salary, BirthDate, HireDate,
Region, Country
FROM northwind.Employees
WHERE EmployeeID > {INCREMENTVALUE}
```

Make sure that the query includes the INCREMENTVALUE parameter and uses the same Incremental Column Name and Comparator values as the fields below the query. For example:

Create Schemas	- 20
Table Name *	
new_employees	
The name of the schema table	_
Database Guery* 1 SELECT EmployeeID, FirstName, LastName, Title, Salary, BirthDate, HireDate, Region, Country 2 FROM monthwind, Employees 3 WHERE EmployeeID > (INCREMENTVALUE)	
Include increment data	
Internet Advance Manua A	
EmployeeID	
vaue* 5	
-	
Comparator	
Greater Than X	$\sim$
CANCEL	SAVE

9. Click **Save** to save the query. Anzo creates the new schema and adds it to the list of schemas on the Tables screen. Selecting the schema displays sample data on the right side of the screen.

### Important

The automated data load workflow ignores all changes that are made to the schema on the Tables screen—except for changes to primary and foreign keys. For example, if you edit a column heading to change its semantic type, that change is disregarded when the graphmart is created. Only the original type from the data source is considered. If you add or change primary and foreign keys on the Tables screen, however, the automated data load workflow will retain those changes.

	DB-QA											Create Graph	mart 🔲
	Overview	т	ibles	Vera	ions	Discussi	on	Sharing					
Schem	185			Ŧ	;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;	new_em	ployees						
Crea	te Schemas Fror	n Query		Import Sc	hemas	Creator S	ystem Adminis	trator Last Mod	fied Date 10/1	9/2022	Column Count 9 R	ow Count N	iot Calculated
Q Sea	rch	7	٣			Sample	Data	Foreign Keys	Schema	Sharing			
•	Schema Name	Table Name ↑	Туре	Actions	:	EmployeeID	FirstName	LastName	Title	Salary	BirthDate	HireDate	Region
	employees			1	<	13	FIRST NAME	7 string Test	2 string	on Most	1955-05-07 00 00	1993-10-17:00	or wa
	new_hites			:	~	4	Mithael	Suama	Sales Recreterts	2004/07	1043-07-02-00-00	1003-10-17-00	"
		new_employ_	Manual	1		-							
	northwind				<	10	Martin	Mark	Sales Manager	2500.06	1957-04-05 00:00	1995-07-06-00	α
						7	Robert	King	Sales Representa	1991.55	1960-05-29 00:00	1994-01-02:00	01
						11	sebastian	ken.	Inside Sales Coor	2100.5	1958-01-09-00.00	1994-01-02:00	or wa
4					•	14	First Test	Test2	HR2		2022-07-28 00:00	2022-05-28-00	ox
						4							,

- 10. If you want to create additional tables in the schema, follow these steps:
  - a. Click the menu icon (1) in the Actions column for the schema name and select Add
     Table. For example:



The Create New Table dialog box is displayed.

Create New Table		22	
			Â
Table Name *			I
The name of the schema table			I
			I
Enter Query Here			I
			I
			I
			I
			I
			Ē
	CANCEL	CAVE	

b. In the Create New Table dialog box, specify a name for the new table in the **Table Name** field.

- c. In the **Schema Query** field, write the SQL query that defines the data for the new table.
- d. Click **Save** to add the table to the schema and return to the Tables screen.

For information about creating or changing primary keys and foreign keys, see Assigning Primary Keys in a Schema and Adding Foreign Keys.

When you are ready to onboard the data to Anzo, see Onboarding Data with the Automated Workflow for next steps. Or, if you want to onboard or virtualize the source by manually writing SPARQL queries against the Graph Data Interface service, see Onboarding or Virtualizing Data with SPARQL Queries.

# Adding an HTTP or SPARQL Data Source

Follow the instructions below to add a connection to a SPARQL or HTTP endpoint. Though you cannot use the automated workflow to ingest data to Anzo from an endpoint, as you can from filebased and database sources, configuring the connection to endpoints adds those sources as Context Providers so that you can reference context keys in Graph Data Interface (GDI) queries against the endpoints.

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

Data Sources So	chemas						
Y Q Search	F •	≡				dd Data So	urce 👻
🗌 🕶 Title 🛧 Description	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	:
∃ 08-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
E Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
				Rows per page:	25 <del>v</del> 1-6 o	f6 <	>

 Click the Add Data Source button, select HTTP, and then choose the type of endpoint to connect to. Note that HTTP Datasource and Sparql Datasource can be used interchangeably at this time. There are no differences in the configuration options. Anzo opens the Create Data Source screen. For example:

Create HTTP Datasource		
Title *		
Description		
Url *		
User		
Password		0
Confirm Password		0
	CANCEL	SAVE

- 3. At the top of the screen, specify a **Title** for the source and enter an optional **Description**.
- 4. Enter the endpoint connection details:
  - **Url**: The URL for the endpoint. When connecting to an Anzo SPARQL endpoint, make sure that any URIs, such as a graphmart URI, are URL-encoded. For example:

https://10.0.10.10/sparql/graphmart/http%3A%2F%2Fcambridgesemantics.com%2F Graphmart%2F1686168b-3eaf-4fdc-9730-1903717b9e62

- User: The user name to use for authentication.
- **Password**: The password for the user name.
- Confirm Password: Confirm the password for the user.
- 5. Click **Save** to save the data source connection. Anzo displays the Overview tab where you can view the connection details.

Now that the connection to the endpoint is configured, you can write Graph Data Interface (GDI) queries to virtualize or ingest data from the source. For information about using the GDI, see Onboarding or Virtualizing Data with SPARQL Queries.

# Adding a CSV Data Source

Follow the instructions below to add a CSV data source and import the data from the files.

## Tip

If your CSV data source is consistently updated with new or changed files, you can configure the source to process the data incrementally. For details, see Configuring a CSV or Parquet Source for Incremental Processing.

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

	Data Sources So	hemas						
<b>Y</b>	<b>L</b> Search	F •	≡			Ad	d Data Sou	urce 👻
•	Title	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
	Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	÷
	🗇 Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	÷
	E DB-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
	Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	÷
	Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
	Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
					Rows per page:	25 👻 1-6 of	6 <	>

 Click the Add Data Source button and select File > CSV Data Source. Anzo opens the Create CSV Data Source screen.

Create CSV Data Source		
Title *		
THUC -		
Description		
	CANCEL	SAVE

3. Specify a name for the Data Source in the **Title** field, and type an optional description in the **Description** field. Then click **Save**. Anzo saves the source and displays the Tables tab.

	Movies				Create Graphmart					
	Overview	Tables	Versions	Discussion	Sharing					
	CSV									
Please add new files to get started.										
	Add New File									

4. Click the Add New File button. Anzo displays the Add New File dialog box.

Add New File		
Source		
From Your Computer O From File Store		
Upload to		
sysadmin User Folder	Change	
The location where your uploaded files will be put on the Anzo Server. Defaults to your user folder		
Drag and drop files or browse from comp Please select the .csv files to import Selected: None	puter.	LEAR ALL
	CANCEL	NEXT
	OVINOUL.	iners.1

5. Follow the appropriate steps below depending on whether the CSV files are on your computer or the shared File Store:

If the files are on your computer:

# Note

The **From Your Computer** option is a convenient way to do a one-time ingestion so you can quickly get started with your data. It should not be relied upon as part of a regular onboarding workflow unless the server is configured to store uploaded files on the shared file store as described in <u>Setting the Default Base File Store Path for File</u> <u>Uploads</u> in the Administration Guide. Data source files that are routinely updated and re-ingested should be hosted on a shared file store.

a. As a best practice, check the upload location that is listed in the Upload To field by hovering your pointer over the value to view the tooltip. Make sure the upload location is a directory on the shared file store and not in the server installation path. If the file is not uploaded to the shared file store it is not accessible by applications like AnzoGraph. In addition, other users cannot create graphmarts from the data source because they typically do not have access to the file location.

For example, viewing the Upload To location for the screen above shows that the file will be uploaded to the server installation path, /opt/Anzo/Server/data...



If your Upload To location is configured to upload the file to the server installation path, click **Change** and select an upload location that is on the shared file store. For example, the image below shows the Upload Folder Location dialog box that is presented after clicking **Change**. A folder called **fileUploads** is selected on the shared store.

Upload Folder Location							
Current Folder			GO				
Selected: /fileUploads/ 🛞				CLEAR ALL			
Server Shared Filesystem		docs					
sysadmin User Folder		fileUploads					
		flights					
		fmcsa					
	<	chib10		¥			
CREATE NEW FOLDER			CANCE	EL OK			

 b. Drag and drop the files onto the screen or click browse to navigate to the files and select them. Anzo attaches the files and the Next button becomes active. For example:



c. Click **Next**. Anzo lists the uploaded files on the left side of the screen with a status of Pending. For example:

	Movies Not Versioned 🖸 Overview	Tables	Versions		Create Graphmart Discussion Sharing
Schen Add N	nas Iew File Process Pe Q. Search	nding Files	<b>∓</b> ⊡	¢	
•	Title 个	Status	Actions	÷	csv
	MovieComposers	Pending     Pending	:	ł	
	MovieDirectors	··· Pending	:	l	Schema Details
	MovieEditors	··· Pending	:		These process are new to new anothe
	MovieProducers	Pending	:	÷	
	Rows per pa	ge: 👻 1-6 of 6	<	>	

If the files are on the File Store:

a. Click the From File Store radio button. Anzo displays the file selection dialog box. For example:



b. On the left side of the screen, select the file store that hosts the CSV files. On the right side of the screen, navigate to the directory that contains the files to import. The screen displays the list of files in the directory.

c. Select each file that you want to import. If you have multiple files with the same schema— the files contain the same columns listed in the same order—and you want the files to be imported as if they are a single file, you can select the Insert Wildcard option. Then type a string using asterisks as wildcard characters to find the files with similar names. Files that match the specified string will be imported as one file. After typing a string, click Apply to include that string in the Selected list.

## Example

The image below shows a directory with several CSV files. For this example, **part.csv** and **partsupp.csv** have the same schema and can be imported as one file. The **Insert Wildcard** option is selected, and **part\*.csv** is specified to identify the two files.



d. When you have finished selecting files, click **Next** to close the dialog box. Anzo lists the uploaded files on the left side of the screen with a status of Pending. For example:

	Movies Not Versioned 🖄 Overview	Tables	Versions		Create Graphmart Discussion Sharing
Add N	nas ew File Process Pe Q Search	nding Files	* D	÷	
•	Title 个	Status	Actions	E .	csv
	MovieComposers	Pending	:	ł	
	MovieDirectors	··· Pending	:	l	Schema Details
	MovieEditors	· Pending			
	MovieProducers	- Pending	÷	¥	
	Rows per pag	ge: 👻 1-6 of 6	< >	>	

6. If you do not need to change CSV file options, click the Process Pending Files button to import all of the pending files. Anzo imports the data and updates the status to Processed. If you do need to change CSV file options, click the menu icon (:) in the Actions column for that file and select Edit. To change the options for multiple files, select the checkbox next to each file, and then click the Edit button at the bottom of the table. Anzo displays the Edit CSV File screen. For example, the image below shows the Edit screen for a single file:

Edit CSV File	
Title	
MovieActors	
File Path	
/movies/MovieActors.csv	BROWSE
☑ Contains Header 🔲 Use Extended Sample 📄 Infer All Columns as Strings	
Delimiter	
Row Delimiter	
Csv Escape Character	
Char Set	
Quote Character	
CANCEL	SAVE & IMPORT

Change the options as needed and then click **Save & Import** to import the file or files. Anzo imports the data and updates the status to Processed.

7. Once the files are processed, you can click a table row on the left side of the screen to display the schema on the right side of the screen.

#### Important

The automated data load workflow ignores all changes that are made to the schema on the Tables screen—except for changes to primary and foreign keys. For example, if you edit a column heading to change its semantic type, that change is disregarded when the graphmart is created. Only the original type from the data source is considered. If you add or change primary and foreign keys on the Tables screen, however, the automated data load workflow will retain those changes.

	Movies						Create Graphmart
	Overview	Tables	Versions	Discussion	Sharing		
Schemas 🔻 🗅 🗧		* ∷ ←	MovieActors				
Add New File Process Pending Files		Pending Files		Creator System Ad	ministrator Last Modified Date 1	0/20/2022 Column Count	4 Row Count Not Calculated
T	<b>Q</b> Search	- *		Sample Data	Foreign Keys Schen	na Sharing	
•	Title	↑ Status	Actions 🗄	MovielD	MovieTitle	ActorID	ActorName
	MovieActors	Processed	- E - ^	''s Int	T String	**• int	T String
	MovieComposers	⊘ Processed		15400287	Mission: Impossible (film series)	689567	Ving Rhames
	the lationary	0		30271424	Alvin and the Chipmunks (film serie	e 921466	Tony Hale
	Movieurectors	Ø+tocessed	-	30271424	Alvin and the Chipmunks (film serie	e 24202290	Jenny Slate
	MovieEditors	Processed	1	30271424	Alvin and the Chipmunks (film series	e 708892	Jason Lee (actor)
	MovieProducers	Processed	÷	15400087	Station Impossible (Streaghter)	51.640	Tem Coulea
	Rows per	page: - 1-6 of 6	$\langle \rangle$	10400207	master impassive (initiaenes)	31400	ion crase
				30271424	Alvin and the Chipmunks (film seri-	e 2647835	Zachery Levi

For information about creating or changing primary keys and foreign keys, see Assigning Primary Keys in a Schema and Adding Foreign Keys.

When you are ready to onboard the data to Anzo, see Onboarding Data with the Automated Workflow for next steps. Or, if you want to onboard or virtualize the source by manually writing SPARQL queries against the Graph Data Interface service, see Onboarding or Virtualizing Data with SPARQL Queries.

# Adding a JSON Data Source

This topic provides instructions for adding a JSON data source.

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

Data Sources So	chemas						
Y Q Search	<b>-</b> •	≡			A	id Data Sou	irce 👻
□ ▼ Title ↑ Description	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	П	:
∃ DB-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	П	:
E Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
E Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	П	:
				Rows per page:	25 👻 1-6 of	6 <	>

 Click the Add Data Source button and select File > JSON Data Source. Anzo opens the Create JSON Data Source screen.

Create JSON Data Source		
Title 4		
nue -		
Description		
File Location *		BROWSE
	CANCEL	SAVE

3. Specify a name for the data source in the **Title** field, and type an optional description in the **Description** field.

4. Click the **JSON File Location** field to open the File Location dialog box.



5. Follow the appropriate steps below depending on whether the file is on your computer or the shared File Store:

If the file is on your computer:

### Note

The **From Your Computer** option is a convenient way to do a one-time ingestion so you can quickly get started with your data. It should not be relied upon as part of a regular onboarding workflow unless the server is configured to store uploaded files on the shared file store as described in <u>Setting the Default Base File Store Path for File</u> <u>Uploads</u> in the Administration Guide. Data source files that are routinely updated and re-ingested should be hosted on a shared file store.

a. As a best practice, check the upload location that is listed in the Upload To field by hovering your pointer over the value to view the tooltip. Make sure the upload location is a directory on the shared file store and not in the server installation path. If the file is not uploaded to the shared file store it is not accessible by applications like AnzoGraph. In addition, other users cannot create graphmarts from the data source because they typically do not have access to the file location.

For example, viewing the Upload To location for the screen above shows that the file will be uploaded to the server installation path, /opt/Anzo/Server/data...



If your Upload To location is configured to upload the file to the server installation path, click **Change** and select an upload location that is on the shared file store. For example, the image below shows the Upload Folder Location dialog box that is presented after clicking **Change**. A folder called **fileUploads** is selected on the shared store.

Jpload Folder Locatio	n			
Current Folder			GO	
Selected: /fileUploads/ @			cı	EAR ALL
Server Shared				•
Filesystem		docs		- 1
sysadmin User Folder		fileUploads		
		flights		
		fmcsa		
		nhih10		•
CREATE NEW FOLDER			CANCEL	OK

b. Drag and drop the file onto the screen or click browse to navigate to the file and select it.
 Anzo attaches the file and the OK button becomes active. For example:

File Location			
Source  From Your Computer  From File Store Upload to Server Shared Filesystem fileUploads/ Cha The location where your uploaded files will be put on the Anzo Server. Defaults to your user folder.	nge		
Drag and drop files or browse from computer. Please select the .json file to import			
Selected: datafox-api-signals.json 82.4		CLEAR	ALL
CREATE NEW FOLDER	CANCE	L	ок

c. Click **OK**. Anzo lists the path to the file in the JSON File Location field.

If the file is on the File Store:

- a. Click the From File Store radio button.
- In the File Location dialog box, on the left side of the screen, select the shared file store that hosts the file. On the right side of the screen, navigate to the directory that contains the file to import. The screen displays the list of files in the directory. For example:

File Location					
Source ○ From Your Computer ( Current Folder /json/ □ Insert Wildcard Ex: * json, *north* json	From File Store  APPLY	GO Show All File Types			
Selected: None		CLEAR ALL			
Server Shared Filesystem	split split				
sysadmin User Folder	add-api-signals.json				
	aircheck_carlsontucker_asset_summary.json				
	datafox-api-signals.json				
	oncogrid-modded.json				
	VF_Style_Vol3_Ep1_Summa	ary-ONE.json			
CREATE NEW FOLDER		CANCEL OK			

c. Select the file that you want to import and then click **OK** to close the dialog box. If you have multiple files with the same schema— the files contain the same arrays in the same order—you can select the **Insert Wildcard** option. Then type a string using asterisks as wildcard characters to find the files with similar names. Files that match the specified string will be imported as **one file**. You can specify up to 16,000 files using a wildcard. After typing a string, click **Apply** to include that string in the Selected list.

# Example

The image below shows a directory with multiple JSON files. For this example, add-api-signals.json and datafox-api-signals.json have the same schema and can be imported as one file. The Insert Wildcard option is selected, and \*signals.json is specified to identify the two files.
Source From Your Computer (	From File Store	
Current Folder ∕json/		GO Show All File
<ul> <li>Insert Wildcard</li> <li>Ex: *.json, *north*.json</li> </ul>	APPLY	
Selected: *signals.json 🛇		CLEAR ALL
Server Shared Filesystem	n split	
sysadmin User Folder	add-api-signals.jsor	n
	aircheck_carlsontud	cker_asset_summary.json
	datafox-api-signals.	json

6. Click **Save** to create the data source. Anzo adds the source and displays the Overview screen. For example:

	Datafox	Varaiana	Dissussion	Cheving	Create Graphmart
Title Datafox Descriptio None JSON File /json/d: Root Elem None	n Location atafox-api-signals.json ent Name	verai015	DIOLUSSIUI	General Type Creator Updated Released	Graph System Administrator 38 minutes ago 38 minutes ago

7. By default, when the data from this source is ingested, the entire root node is captured; the node at the root of the hierarchy is loaded to AnzoGraph in its entirety. Building the hierarchical record for a large file is extremely memory intensive. To increase load performance and decrease memory usage when onboarding a large file with many repeating elements, Cambridge Semantics recommends that you configure the **Root Element Name** field on the Overview tab. This field designates the element in the hierarchy that should be treated as the root node. Specifying the desired root node tells the Graph Data Interface to scan into memory only the data that you are interested in and not the entire file. To set the root element, follow these steps:

- a. Click in the **Root Element Name** field to make it editable.
- b. Add the name of the element to designate as the root element. Type the name the same way it appears in the file. Data is captured from whichever node in the hierarchy matches the Root Element Name value in its entirety.

#### Tip

It is not necessary to express the path to an element if it is low in the hierarchy. For specificity, however, you can use dot notation to supply the path. For example, specifying "city" captures all city elements anywhere in the file. But specifying "country.state.city" captures only the city elements that are under state and city in the hierarchy. You can also include the dollar sign (\$) character to anchor the selector at the root of the file. For example, "data" captures all data elements anywhere in the file. But "\$.data" captures only the data elements that are at the root of the hierarchy.

As an example, for a file that contains weather data in daily, hourly, minutely, and currently hierarchies, "hourly" is specified to target only the data under the hourly hierarchy:

	Weather Not Versioned				Create Graphmart	Д
	Overview	Versions	Discussion	1	Sharing	
Title Weather Description None JSON File I /json/we	n Location eather.json			General Type Creator Updated Released	JSON System Administrator a few seconds ago a few seconds ago	
Root Eler	nent Name		✓ X	http://cambrid Tags None	igesemantics.com/JSON	

c. Click the checkmark icon ( $\checkmark$ ) to save the change.

For information about creating or changing primary keys and foreign keys, see Assigning Primary Keys in a Schema and Adding Foreign Keys.

When you are ready to onboard the data to Anzo, see Onboarding Data with the Automated Workflow for next steps. Or, if you want to onboard or virtualize the source by manually writing SPARQL queries against the Graph Data Interface service, see Onboarding or Virtualizing Data with SPARQL Queries.

# Adding an XML Data Source

This topic provides instructions for adding an XML data source.

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

Data Sources	Schemas						
Y Q Search	F • 🖬	≡			Ad	ld Data Sou	urce 👻
□ ▼ Title ↑ Description	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
🗇 Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	П	:
E DE-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	П	:
E Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
E Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
				Rows per page:	25 👻 1-6 of	6 <	>

 Click the Add Data Source button and select File > XML Data Source. Anzo opens the Create XML Data Source screen.

Create XML Data Source		
Title *		
Description		
File Location *		BROWSE
	CANCEL	SAVE

- 3. Specify a name for the data source in the **Title** field, and type an optional description in the **Description** field.
- 4. Click the File Location field to open the File Location dialog box.



5. Follow the appropriate steps below depending on whether the file is on your computer or the shared File Store:

If the file is on your computer:

## Note

The **From Your Computer** option is a convenient way to do a one-time ingestion so you can quickly get started with your data. It should not be relied upon as part of a regular onboarding workflow unless the server is configured to store uploaded files on the shared file store as described in <u>Setting the Default Base File Store Path for File</u> <u>Uploads</u> in the Administration Guide. Data source files that are routinely updated and re-ingested should be hosted on a shared file store.

a. As a best practice, check the upload location that is listed in the Upload To field by hovering your pointer over the value to view the tooltip. Make sure the upload location is a directory on the shared file store and not in the server installation path. If the file is not uploaded to the shared file store it is not accessible by applications like AnzoGraph. In addition, other users cannot create graphmarts from the data source because they typically do not have access to the file location.

For example, viewing the Upload To location for the screen above shows that the file will be uploaded to the server installation path, /opt/Anzo/Server/data...



If your Upload To location is configured to upload the file to the server installation path, click **Change** and select an upload location that is on the shared file store. For example, the image below shows the Upload Folder Location dialog box that is presented after clicking **Change**. A folder called **fileUploads** is selected on the shared store.

Upload Folder Location	n			
Current Folder			GO	
Selected: /fileUploads/ ③				CLEAR ALL
Server Shared Filesystem		docs		· · ·
sysadmin User Folder		fileUploads		
		flights		
		fmcsa		
	<	chih10	_	• •
CREATE NEW FOLDER			CANCE	EL OK

- b. Drag and drop the file onto the screen or click browse to navigate to the file and select it.
   Anzo attaches the file and the OK button becomes active.
- c. Click **OK**. Anzo lists the path to the file in the XML File Location field.

If the file is on the File Store:

- a. Click the From File Store radio button.
- b. In the File Location dialog box, on the left side of the screen, select the appropriate File Store. On the right side of the screen, navigate to the directory that contains the file to import. The screen displays the list of files in the directory. For example:

File Location		
Source From Your Computer	From File Store	
Current Folder /xml/		GO Show All File
Insert Wildcard Ex: *.json, *north*.json	APPLY	.,,
Selected: None		CLEAR ALL
Server Shared Filesystem	people.xml	
sysadmin User Folder	sample.xml	
	test.xml	
CREATE NEW FOLDER		CANCEL OK

c. Select the file that you want to import and then click **OK** to close the dialog box. Anzo lists the path to the file in the XML File Location field.

#### Note

If you have multiple files with the same schema— the files contain the same elements in the same order—and you want the files to be imported as if they are a **single file**, you can select the **Insert Wildcard** option. Then type a string using asterisks as wildcard characters to find the files with similar names. Files that match the specified string will be imported as **one file** and will result in one job being created in the pipeline to ingest all of the files that are selected by the specified string. After typing a string, click **Apply** to include that string in the Selected list.

6. Click **Save** to create the data source. Anzo adds the source and displays the Overview screen. For example:

<>	People Not Versioned 🖉 Overview	Versions	Discussion	Sharin	Create Graphmart
Title People Descriptio None XML File L /xml/pe	n ocation cople.xml			General Type Creator Updated Released	Graph System Administrator a few seconds ago a few seconds ago
Repeating None	Element Paths			http://cambridgese	emantics.com/XMLDataSource/

7. By default, when the data from this source is ingested, the entire root node is captured; the node at the root of the hierarchy is loaded to AnzoGraph in its entirety. Building the hierarchical record for a large file is extremely memory intensive. To increase load performance and decrease memory usage when onboarding a large file with many repeating elements, Cambridge Semantics recommends that you configure the **Root Element Name** field on the Overview tab. This field designates the element in the hierarchy that should be

treated as the root node. Specifying the desired root node tells the Graph Data Interface to scan into memory only the data that you are interested in and not the entire file. To set the root element, follow these steps:

- a. Click in the Root Element Name field to make it editable.
- b. Add the name of the element to designate as the root element. Type the name the same way it appears in the file. Data is captured from whichever node in the hierarchy matches the Root Element Name value in its entirety.

#### Tip

It is not necessary to express the path to an element if it is low in the hierarchy. For specificity, however, you can use dot notation to supply the path. For example, specifying "city" captures all city elements anywhere in the file. But specifying "country.state.city" captures only the city elements that are under state and city in the hierarchy. You can also include the dollar sign (\$) character to anchor the selector at the root of the file. For example, "data" captures all data elements anywhere in the file. But "\$.data" captures only the data elements that are at the root of the hierarchy.

As an example, for a file that contains weather data in daily, hourly, minutely, and currently hierarchies, "hourly" is specified to target only the data under the hourly hierarchy:

	Weather				Create Graphmart	Д
	Overview	Versions	Discussion	:	Sharing	
Title Weather Description None XML File Lo /xml/we	cation ather.xml			General Type Creator Updated Released	XML System Administrator a few seconds ago a few seconds ago	
Root Elen hourly	nent Name			http://cambrid	gesemantics.com/XMLD	
			✓×	Tags None		

c. Click the checkmark icon ( $\checkmark$ ) to save the change.

For information about creating or changing primary keys and foreign keys, see Assigning Primary Keys in a Schema and Adding Foreign Keys.

When you are ready to onboard the data to Anzo, see Onboarding Data with the Automated Workflow for next steps. Or, if you want to onboard or virtualize the source by manually writing SPARQL queries against the Graph Data Interface service, see Onboarding or Virtualizing Data with SPARQL Queries.

## Adding a SAS Data Source

Follow the instructions below to add a SAS data source and import data from SAS7BDAT files.

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

Data Sources S	chemas						
Y Q Search	F •	≡			Ad	d Data Sou	irce 👻
Title 🛧 Description	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
🗇 Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	:
≡ 08-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
E Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
				Rows per page:	25 ¥ 1-6 of	6 <	>

 Click the Add Data Source button and select File > SAS Data Source. Anzo opens the Create SAS Data Source screen.

Create SAS Data Source		
Title *		
Description		
C/	ANCEL	SAVE

3. Specify a name for the source in the **Title** field, and type an optional description in the **Description** field. Then click **Save**. Anzo saves the source and displays the Tables tab.



4. Click the Add New File button. Anzo displays the Add New File dialog box.

Add New File
Source <ul> <li>From Your Computer</li> <li>From File Store</li> </ul>
Upload to sysadmin User Folder
The location where your uploaded files will be put on the Anzo Server. Defaults to your user folder. Change
<b>%</b>
Drag and drop files or browse from computer.
Please select the .sas7bdat files to import
Selected: None CLEAR ALL
CANCEL NEXT

5. Follow the appropriate steps below depending on whether the SAS files are on your computer or the shared File Store:

If the files are on your computer:

#### Note

The **From Your Computer** option is a convenient way to do a one-time ingestion so you can quickly get started with your data. It should not be relied upon as part of a regular onboarding workflow unless the server is configured to store uploaded files on the shared file store as described in <u>Setting the Default Base File Store Path for File</u> <u>Uploads</u> in the Administration Guide. Data source files that are routinely updated and re-ingested should be hosted on a shared file store.

a. As a best practice, check the upload location that is listed in the **Upload To** field by hovering your pointer over the value to view the tooltip. Make sure the upload location is

a directory on the shared file store and not in the server installation path. If the file is not uploaded to the shared file store it is not accessible by applications like AnzoGraph. In addition, other users cannot create graphmarts from the data source because they typically do not have access to the file location.

For example, viewing the Upload To location for the screen above shows that the file will be uploaded to the server installation path, /opt/Anzo/Server/data...



If your Upload To location is configured to upload the file to the server installation path, click **Change** and select an upload location that is on the shared file store. For example, the image below shows the Upload Folder Location dialog box that is presented after clicking **Change**. A folder called **fileUploads** is selected on the shared store.

Current Folder			GO	
Selected: /fileUploads/ @	3			CLEAR AL
Server Shared				
Filesystem	<u> </u>	docs		
sysadmin User Folder		fileUploads		
		flights		
		fmcsa		
		chib10		

- b. Drag and drop the files onto the screen or click **browse** to navigate to the files and select them. Anzo attaches the files and the Next button becomes active.
- c. Click **Next**. Anzo lists the uploaded files on the left side of the screen with a status of Pending. For example:

	Overview		Tables		Ve	sions	Discussion	Sharing
Schem	as			Ŧ C	←	1		
Add N	New File	Pro	cess Pending Files					
<b>Q</b> Sear	ch		= •					
	Title	$\uparrow$	Status	Action	s :		SAS	
	airline		··· Pending	:			S	
	cars		··· Pending	:				
	demo		· Pending	:			Schema I Please process the f	Details lles to view them.

If the files are on the File Store:

- a. Click the From File Store radio button.
- b. In the File Location dialog box, on the left side of the screen, select the appropriate File Store. On the right side of the screen, navigate to the directory that contains the file to import. The screen displays the list of files in the directory. For example:



c. Select each file that you want to import. When you finish selecting files, click **Next** to close the dialog box. Anzo lists the uploaded files on the left side of the screen with a status of Pending. For example:

S	Demo	Tablas	,	losoiono	Discussion	Create Graphmart	Д
	Overview	Tables	`	rsions	Discussion	Sharing	
Scher	mas		<b>∓ :: ∢</b>	-			
Add	d New File Pro	cess Pending Files					
Q Se	arch						
•	Title 🛧	Status	Actions		SAS		
	airline	··· Pending	:		S		
	cars	··· Pending	:				
	demo	··· Pending	:		Schema Please process the f	Details files to view them.	
	Rows per page:	▼ 1-3 of 3	$\langle \rangle$				

6. If you do not need to change SAS file options, click the **Process Pending Files** button to import all of the pending files. Anzo imports the data and updates the status to Processed. If you do need to change SAS file options, click the menu icon (:) for that file and select **Edit**. To change the options for multiple files, select the checkbox next to each of the files, and then click the **Edit** button at the bottom of the table. Anzo displays the Edit SAS File screen. For example, the image below shows the Edit screen for a single file:

Edit SAS Files	
Title	
airline	
File Path	
/sas/airline.sas7bdat	BROWSE
Force column names to lower case Infer numeric colummn as decimal	
Scale of infered decimal	
📋 Infer numeric colummn as Float 📋 Infer numeric colummn as Int 📋 Infer numeric colummn as Long	
Infer numeric colummn as Short  Extract column labels	
Number of seconds to allow reading of file metadata	
Minimum byte length of input splits	
Maximum byte length of input splits	
CANCEL SAVE	& IMPORT

Change the options as needed and then click **Save & Import** to import the SAS file or files. Anzo imports the data and updates the status to Processed.

7. Once the files are processed, you can click a table row on the left side of the screen to display the schema on the right side of the screen.

#### Important

The automated data load workflow ignores all changes that are made to the schema on the Tables screen—except for changes to primary and foreign keys. For example, if you edit a column heading to change its semantic type, that change is disregarded when the graphmart is created. Only the original type from the data source is considered. If you add or change primary and foreign keys on the Tables screen, however, the automated data load workflow will retain those changes.

Demo Not Versioned						Create	e Graphmart
Overview	Tables	Versions	Discussion	Sharing			
Schemas	∓ ¤ <	airline					
Add New File Proces	ss Pending Files	Creator System	n Administrator Las	st Modified Date 12	/02/2022 Column	Count 6 Row Co	unt Not Calculated
<b>Q</b> Search	<b>- -</b>	Sample Data	Foreign K	eys Schem	a Sharing		
🗌 🕶 Title 🗠 Sta	atus Actions :	YEAR	Y	W	R	L	К
airine Q	Devenues :	100 Double	128 Double	ue Double	100 Double	ua Double	128 Double
		1951	1.94799995422363	0.296999999940094	0.393999993801117	1.54999995231628	0.564000010490417
cars 🧭	Processed	1957	4 42000007629395	0.379000008106232	0.304500013589859	2.70700001716614	0.921000003814697
demo 🧭	Processed						
		1956	3.97900009155273	0.361000001430511	0.382200002670288	2.43499994277954	0.800000011920929
		1954	3.02500009536743	0.3350000834465	0.402500003576279	1.96399998664856	0.776000022888184
Rows per page: 👻	1-3 of 3 < >	1948	1.21399998664856	0.243000000715256	0.145400002598763	1.41499996185303	0.611999988555908
							Ŧ

For information about creating or changing primary keys and foreign keys, see Assigning Primary Keys in a Schema and Adding Foreign Keys.

When you are ready to onboard the data to Anzo, see Onboarding Data with the Automated Workflow for next steps. Or, if you want to onboard or virtualize the source by manually writing SPARQL queries against the Graph Data Interface service, see Onboarding or Virtualizing Data with SPARQL Queries.

# Adding a Parquet Data Source

Follow the instructions below to create a Parquet data source. You can onboard one file or multiple files with the identical format (schema) per data source.

#### Tip

If your Parquet source is consistently updated with new or changed files, you can configure the source to process the data incrementally. For details, see Configuring a CSV or Parquet Source for Incremental Processing.

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

	Data Sources So	chemas						
T C	<b>L</b> Search	= • <b>ii</b>	≡			Ad	d Data Sou	irce 👻
•	Title	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
	E Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
	🗇 Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	:
	E DB-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
	E Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
	Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
	Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
					Rows per page:	25 👻 1-6 of	6 <	>

 Click the Add Data Source button and select File > Parquet Data Source. Anzo opens the Create Parquet Data Source screen.

Create Parquet Data Source		
Title *		
Description		
	CANCEL	SAVE

 Specify a name for the source in the Title field, and type an optional description in the Description field. Then click Save. Anzo saves the source and displays the Overview tab. For example:

E E	Events				Create Graphmart
	Overview	Versions	Discussion	Shar	ing
Description None Parquet File None				General Type Creator Updated Released http://cambridg Tags None	Parquet System Administrator Dec 2, 2022 9:43 AM Dec 2, 2022 9:43 AM

- On the Overview tab, click in the Parquet File field to make the value editable. Then click
   Browse to open the File Location dialog box and select the file to import.
- 5. In the File Location dialog box on the left side of the screen, select the file store for the Parquet file. On the right side of the screen, navigate to the directory that contains the file to import. The screen displays the list of files in the directory. For example:



6. Select the file that you want to import. If you have multiple files with the identical format you can select the **Insert Wildcard** option. Then type a string using asterisks as wildcard characters to find the files with similar names. Files that match the specified string will be imported as **one file** and will result in one job being created in the pipeline to ingest all of the files that are selected by the specified string. You can specify up to 16,000 files using a wildcard. After typing a string, click **Apply** to include that string in the Selected list.

#### **Example**

The image below shows a directory with multiple parquet files. The **events.parquet** and **events-2.parquet** file have the identical format and can be imported as one file. The **Insert Wildcard** option is selected, and **event**\* is specified to identify the two files.

Current Folder	GO D Show All File	è
Insert Wildcard Ex: *.json, *north*.json	Use Incremental Processing	
Selected: event* 😒	CLEAR	ALL
Server Shared Filesystem	events-2.parquet	•
sysadmin User Folder	events.parquet	
	parts-10-21.parquet	
	parts-2.parquet	

 After selecting the file, click OK to close the File Location dialog box. Then click the checkmark icon (
 to save the change to the Parquet File field. Anzo imports the file and generates a data model.

	Events lot Versioned 🖄				Create Graphmart
	Overview	Versions	Discussion	Shar	ing
Description None Parquet File /parquet/e	vent*			General Type Creator Updated Released	Parquet System Administrator Dec 2, 2022 9:49 AM Dec 2, 2022 9:43 AM
				http://cambrid	gesemantics.com/ParquetData 🔲
				Tags None	

For information about creating or changing primary keys and foreign keys, see Assigning Primary Keys in a Schema and Adding Foreign Keys.

When you are ready to onboard the data to Anzo, see Onboarding Data with the Automated Workflow for next steps. Or, if you want to onboard or virtualize the source by manually writing SPARQL queries against the Graph Data Interface service, see Onboarding or Virtualizing Data with SPARQL Queries.

# Configuring a CSV or Parquet Source for Incremental Processing

If you have a CSV or parquet data source that is consistently updated with new or changed files, you can configure the source to process the data incrementally. When a source is set up to onboard data incrementally and new files are added to the data source directory, those new files are ingested and added to the existing dataset. This topic describes how incremental processing works, gives recommendations for organizing the source files, and provides instructions for configuring incremental processing.

- Incremental Processing Overview
- Organizing the Data Source Directory
- Configuring Incremental Processing

#### **Incremental Processing Overview**

When you configure file-based incremental processing, there are two methods to choose from:

#### File Name Strategy

The first method is called the **File Name Strategy**. When this strategy is chosen, Anzo saves the names of the files that are onboarded. Subsequent workflow runs ingest any files whose names are not saved. This method is useful when you know that new files will be added to the data source directory but the existing files will not be edited. If the contents of a previously ingested file changes but the file name does not, that file will not be reprocessed during the next run.

#### Last Modified Strategy

The second method is called the **Last Modified Strategy**. When this strategy is chosen, Anzo saves the last modified timestamps for the files that are ingested when the workflow is run. Subsequent runs ingest any files whose last modified timestamps are greater than the saved timestamps. Each time an incremental workflow run, the job-specific last modified date will be updated to the latest modified date of the files processed. This method is useful when you know that the contents of previously ingested files may change in addition to adding new files. Since the last modified date is updated when a file is changed, the changes will be processed during the next run.

#### **Organizing the Data Source Directory**

Regardless of the strategy you use to process data incrementally, it is important to consider the schema when organizing the data source files on the file store. In order to enable incremental processing, the source files to import must be specified using wildcard characters (\*). That means the list of files that are targeted by the wildcard need to have the same schema.

#### **CSV Data Sources**

For CSV data sources, you have two options. You can create one subdirectory per schema and then add files multiple times, once for each schema. With this structure, you would import the files in each directory separately and could specify the wildcard like \*.csv to import all of the

files in a directory. You can also place all of the files into a single directory and use more detailed text when specifying wildcards. For example, when you add files you apply multiple wildcard values such as patients \*.csv and medication \*.csv.

#### **Parquet Data Sources**

For Parquet data sources, you can only choose one schema per source. You must create a separate data source for each schema type. You may want to create one directory per schema. Then each Parquet source can target one directory and specify a wildcard value such as \*.parquet.

#### **Configuring Incremental Processing**

When adding or modifying a CSV or Parquet data source, you configure incremental processing when you are adding files to the source from the file store. This section provides instructions for adding files to a source and configuring incremental processing. For instructions on adding a new data source, see Adding a CSV Data Source or Adding a Parquet Data Source.

 When selecting the source files on the file store, select the Insert Wildcard checkbox. Enabling the wildcard option activates the Use Incremental Processing option.

File Location	
Current Folder     /nfs/data/parquet/	Show All File
Insert Wildcard Ex: *.json, *north*.json	Use Incremental Processing     APPLY
Selected: None	CLEAR ALL
Server Filesystem	parts-5-21.parquet
sysadmin User Folder	parts-6-21.parquet
	parts-7-21.parquet
	parts-8-21.parquet
	parts-9-21.parquet
	parts.parquet
	CANCEL OK

2. Below the checkbox, type a string using asterisks (\*) as wildcard characters to find the files to be processed. Then click **Apply** to apply the string. If you are configuring a CSV data source,

you can apply multiple wildcard strings to target files with different schemas. The image below shows an example for a Parquet source. The string parts\* is applied to select all of the files with names that start with "parts."

File Location	
Current Folder /nfs/data/parquet/	GO Show All File
Ex: *.json, *north*.json	Use Incremental Processing
Selected: parts* 😒	CLEAR ALL
Server Filesystem	parts-5-21.parquet
sysadmin User Folder	parts-6-21.parquet
	parts-7-21.parquet
	parts-8-21.parquet
	parts-9-21.parquet
	parts.parquet
	CANCEL OK

3. Next, select the **Use Incremental Processing** checkbox. The Configure Incremental Processing dialog box is displayed:

Configure Incremental Processing							
Charles							
Strategy		•					
	CANCEL	APPLY					

- Click the Strategy drop-down list and select the strategy to use for incremental processing. The following list describes the options. For more details about the strategies, see Incremental Processing Overview above.
  - File Name Strategy: Select this option if file names should be used to target the new source data to process each time the graphmart that contains this source is reloaded or refreshed.

- Last Modified Strategy: Select this option if the last modified date should be used to target the new source data to process each time the graphmart that contains this source is reloaded or refreshed.
- If you chose File Name Strategy, click Apply and then click OK or Next to proceed. The source is now configured to process data incrementally. If you chose Last Modified Strategy, proceed to the next step.
- 6. If you chose **Last Modified Strategy**, the Baseline options are displayed. The Baseline determines when the last modified date begins.

Configure Incremental Processing		
_ Strategy Last Modified Strategy		•
Baseline		
<ul> <li>All matching files</li> </ul>		
O All matching files since 11/16/2021 3:29pm		
	CANCEL	APPLY

- 7. By default All matching files is selected as the Baseline. This means all files in the directory that are matched by the wildcard string will be ingested when the graphmart is reloaded or refreshed. If you have older files that you do not want to be ingested, you can select All matching files since. Then click the timestamp field and specify the date and time to use as the Baseline.
- 8. When you have finished configuring the Strategy, click **Apply**. Then click **OK** or **Next** to proceed and finish configuring the data source if necessary.

When the graphmart is created, all of the data that matches the wildcard string and meets the baseline requirements will be onboarded. When new files are added to the data source directory, reloading or refreshing the graphmart will update the data.

For information about creating or changing primary keys and foreign keys, see Assigning Primary Keys in a Schema and Adding Foreign Keys.

When you are ready to onboard the data to Anzo, see Onboarding Data with the Automated Workflow for next steps. Or, if you want to onboard or virtualize the source by manually writing SPARQL queries against the Graph Data Interface service, see Onboarding or Virtualizing Data with SPARQL Queries.

# Assigning Primary Keys in a Schema

If you have a data source without primary keys and you want to be able to create relationships, you can assign primary keys in the schema. Follow the instructions below to edit a schema and assign primary keys. For instructions on creating foreign keys, see Adding Foreign Keys.

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

Data Sources Si	hemas						
Y Q Search	F *	≡			Ad	ld Data So	urce 👻
□   Title	Туре	Schema	Creator	Created Date	Updated Date	Actions	÷
Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	П	:
🗇 Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	П	:
∃ D8-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	П	:
Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
				Rows per page:	25 👻 1-6 of	6 <	>

2. Click the **Schemas** tab to view the list of all schemas. (You can also access a schema by selecting the data source that contains the schema.) The image below shows a view of the Schemas tab:

Data Sources	Schemas				
T Q Search	F • 🖬 🔳			Import Scl	nemas
Title	↑ Data Source	Updated Date	Tags	Actions	:
Airlines	Airlines	Oct 18, 2022		Д	:
i employees	DB-QA	Oct 19, 2022		Д	:
Movies	Top Movies	Sep 30, 2022		Д	:
Movies	Movies	Oct 20, 2022		Д	:
new_hires	DB-QA	Oct 19, 2022		Д	:
i northwind	DB-QA	Oct 19, 2022		Д	:
Tickets	Tickets	Oct 1, 2022		Д	:
		Rows per page: 25	5 👻 1-7 of 1	7 <	>

3. Click the schema for which you want to assign primary keys. Anzo displays the Tables tab for the data source that contains the schema. For example:

	Movies						Create Graphmart
	Overview	Tables	Versions	Discussion	Sharing		
Schem	nas		∓ :: ←	MovieActors			
Add N	ew File Process P	ending Files		Creator System A	dministrator Last Modified Date	10/20/2022 Column Count	4 Row Count Not Calculated
<b>Y</b>	Q Search	= •		Sample Data	Foreign Keys Sche	ma Sharing	
•	Title	↑ Status	Actions	MovieID	MovieTitle	ActorID	ActorName
	MovieActors	Processed	:	15400197	T String	24 Int 680547	T String
	MovieComposers	Processed	:	10400207	mission, il guosolore (il il selles)	0.5007	They Relative
	MovieDirectors	Processed	:	30271424	Alvin and the Chipmunks (film se	ne 921466	Tony Hale
		0		30271424	Alvin and the Chipmunks (film as	rie: 24202290	Jenny Slate
	MovieEditors	Processed	i	30271424	Alvin and the Chipmunks (film se	rie: 708892	Jason Lee (actor)
	MovieProducers	Processed	:	15400287	Mission: Impossible (film series)	31460	Tom Cruise
	Movies	Processed	:				
				30271424	Alvin and the Chipmunks (film se	xie 2647835	Zachary Levi
	Rows p	erpage: 👻 1-6	of6 < >	30271424	Alvin and the Chipmunks (film se	rie 1582246	Kimberly Williams-Paisley

- 4. On the left side of the screen select the table for which you want to assign a primary key.
- On the right side of the screen, find the column that you want to label as the primary key. Hover your pointer over the column name to display additional icons. Edit and Delete icons replace the data type under the column name. For example:

Sample Data	Foreign Keys So	chema Sharing
MovieID	MovieTitle T String	ActorID /
15400287	Mission: Impossible (film ser	ies) 689567
30271424	Alvin and the Chipmunks (film	n serie: 921466
30271424	Alvin and the Chipmunks (film	n serie: 24202290

6. Click the edit icon ( $\checkmark$ ). The Edit dialog box is displayed. For example:

Edit	
Name *	
ActorID	
Name of the column	
Semantic Type	
Int	× × ×
Type of the column	
Primary Key	
	CANCEL SAVE

7. On the Edit screen, select the **Primary Key** checkbox. Then click **Save** to save the change. The column is now the primary key for the table, and a key icon is displayed next to the column name. For example:

Sample Data	Foreign Keys	Schema Sharing
MovieID <sup>1</sup> 2 <sub>9</sub> Int	MovieTitle T String	ActorID
30271424	Alvin and the Chipmunks (film	: 2647835
30271424	Alvin and the Chipmunks (film	: 876916
30271424	Alvin and the Chipmunks (film	: 24202290

Repeat the steps above to assign primary keys for additional tables. For instructions on creating foreign keys, see Adding Foreign Keys.

# Adding Foreign Keys

Follow the instructions below to edit a schema to add foreign keys. For instructions on designating primary keys, see Assigning Primary Keys in a Schema.

1. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing data sources. For example:

	Data Sources	Schemas						
Ŧ	Q Search	= •	≡			Ad	d Data So	urce 👻
	▼ Title ↑ Descriptio	on Type	Schema	Creator	Created Date	Updated Date	Actions	8
	Airlines	CSV Data Sou	Airlines	System Admi	Oct 18, 2022	Oct 18, 2022	Д	:
	Datafox	JSON Data S		System Admi	Oct 19, 2022	Oct 19, 2022	Д	:
	E 08-QA	Database Dat	Employees	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
	Movies	CSV Data Sou	Movie FKs	System Admi	Oct 1, 2022	Oct 19, 2022	Д	:
	E Tickets	CSV Data Sou	Tickets	System Admi	Oct 1, 2022	Oct 1, 2022	Д	:
	Top Movies	CSV Data Sou	Movies	System Admi	Sep 30, 2022	Oct 18, 2022	Д	:
					Rows per page:	25 👻 1-6 of	6 <	>

2. Click the **Schemas** tab to view the list of all schemas. (You can also access a schema by selecting the data source that contains the schema.) The image below shows a view of the Schemas tab:

Data Sources	Schemas				
T Q Search	F • 🖬 🔳			Import Scl	nemas
Title	↑ Data Source	Updated Date	Tags	Actions	:
Airlines	Airlines	Oct 18, 2022		Д	:
employees	DB-QA	Oct 19, 2022		Д	:
Movies	Top Movies	Sep 30, 2022		Д	:
Movies	Movies	Oct 20, 2022		Д	:
new_hires	DB-QA	Oct 19, 2022		Д	:
i northwind	DB-QA	Oct 19, 2022		Д	:
Tickets	Tickets	Oct 1, 2022		Д	:
		Rows per page: 25	5 👻 1-7 of 1	7 <	>

3. Click the schema for which you want to assign foreign keys. Anzo displays the Tables tab for the data source that contains the schema. For example:

	Movies Not Versioned 🖄 Overvlew	Tables		Vers	sions Discussion	Sharing		Create Graphmart
Schen	nas		* C	÷	MovieActors			
Add N	ew File Proces	ss Pending Files	•		Creator System Administ	rator Last Modified Date 10. Foreign Keys Schema	20/2022 Column Count 4	Row Count Not Calculated
•	Title 🔶	Status	Actions	÷	MovieID	MovieTitle	ActoriD	ActorName T. Strike
	MovieActors	Processed	1	Î	30271424	Alvin and the Chipmunks (film serier	876916	Kaley Cuoco
	MovieComposers	⊘Processed	:		30271424	Alvin and the Chipmunks (film serie:	1778716	Jesse McCartney
	MovieDirectors	Processed	:	J.	30271424	Alvin and the Chipmunks (film serier	1582246	Kimberly Williams-Paisley
	MovieEditors	⊘Processed	:		15400287	Mission Impossible (film series)	31460	Tom Cruise
	MovieProducers	Processed	:	÷	30271424	Alvin and the Chipmunks (film serier	24202290	Jenny Slate
	Rows per page:	▼ 1-6 of 6	<	>	30271424	Alvin and the Chipmunks (film serie:	4995125	Matthew Gray Gubler

4. Click the **Foreign Keys** tab on the right side of the screen. The tab lists any existing keys. For example, the image below shows a schema that does not have keys defined:

Sample Data	Foreign Keys	Schema Sharing	
Q Search  No foreign keys found		<b>m =</b>	Create

5. To change an existing key, click a row in the table to open the Edit Foreign Key dialog box. To create a new key, click the **Create** button to open the Create Foreign Key dialog box.

Your key will not be saved unless you select colun	nns from two different tabl	es.	
Key Type: Manual			
Name			
Name of the foreign key			
Source Table		Target Table	
Select source table	~	Select target table	~
Select source columns	×	Select target columns	~
Only View Selected Columns 🛛 💼		Only View Selected Columns	

- 6. On the Create Foreign Key screen, specify a name for the key in the **Name** field.
- 7. Then specify the source and target tables for this key.
  - **Source Table**: The source table is the table where the new foreign key is created. This table refers to the primary key from the Target Table. Click the **Source Table** drop-down list and select the schema table where the foreign key should be created.
  - **Target Table**: The target table is the table that contains the primary key to be referenced by the Source Table. Click the **Target Table** drop-down list and select the table that will pass values to the source table.
- 8. Next, specify the source and target columns for this key:
  - Source Columns: The source column is the column that becomes the foreign key to the target table's primary key. Click the Select Source Columns drop-down list and select the source column. To create a composite key by selecting an additional column, click the Select Source Columns drop-down list again and select a column.

#### Tip

By default the screen shows sample values from the selected source column. If you want to view sample values from all columns in the source table, you can disable the **Only View Selected Columns** option by sliding the slider to the left.

• **Target Columns**: The target column is the primary key column in the target table. Click the **Select Target Columns** drop-down list and select the target column. To create a composite key by selecting an additional column, click the Select Target Columns drop-down list again and select a column.

#### Tip

By default the screen shows sample values from the selected target column. If you want to view sample values from all columns in the target table, you can disable the **Only View Selected Columns** option by sliding the slider to the left.

For example, the image below creates a relationship called **actedIn** where the **MovieID** column in the **MovieActors** table becomes the foreign key and references the values from the primary key column, **MovieID**, in the **Movies** table.

Name actedIn			
Name of the foreign key Source Table		Target Table	
MovieActors X	(   ~	Movies	$x \mid  \checkmark$
MovielD ×	<   ~	MovielD ×	x   ~
Only View Selected Columns		Only View Selected Columns 🛑	
MovielD		MovieID	
30271424		3837	
30271424		3746	

9. When you have finished supplying values, click **Save** to create the new key and return to the Foreign Key list. Repeat this process to create additional keys.

When you onboard the data using this schema, the foreign keys become RDF OWL object properties in the data model. For example, the image below shows a portion of the model that was generated after ingesting the schema that has the foreign key in the example above. In the model, **Acted In** is an object property in the **Movie Actors** class:



# **Onboarding Data with the Automated Workflow**

There are two ways to load a data source with the automated direct data load workflow. You can build a graphmart from a selected data source or you can add a data source to an existing, activated graphmart. Both procedures automatically generate data layers to extract, load, and transform the data to a knowledge graph.

#### Tip

When you build a new graphmart from a data source, advanced graphmart options are made available that are not presented when you add a data source to an existing graphmart. See Graphmart Options for information about the additional options.

- Creating a Graphmart from a Data Source
- Adding a Data Source to an Existing Graphmart
- Direct Load Advanced Settings Reference

## Creating a Graphmart from a Data Source

Follow the steps below if you want to create a new graphmart from a data source.

#### Note

AnzoGraph uses the Graph Data Interface (GDI) Java plugin to connect directly to sources. For file-based sources, make sure the source files are available to AnzoGraph on the shared file store. For databases, if you have configured custom drivers to access those sources in Anzo, the same drivers need to be added to AnzoGraph. For instructions, see <u>Deploy Optional</u> <u>Drivers for Accessing Custom Database Sources</u> in the Deployment Guide.

- 1. If necessary, add the data source. See Adding Data Sources for instructions.
- 2. In the Anzo application, expand the **Onboard** menu and click **Structured Data**. Anzo displays the Data Sources screen, which lists any existing sources. For example:

Data Sources So	hemas						
Y Search	<b>=</b> •	≡			4	Add Data Sou	irce 👻
🗌 🕶 Title 🗠 Description	Туре	Schema	Creator	Created Date	Updated Date	Actions	:
≡ DB-QA	Database Dat	books, north	System Admi	Nov 1, 2022	Nov 1, 2022	Д	:
E Flights	CSV Data Sou	Flights	System Admi	Nov 1, 2022	Nov 1, 2022	Д	:
B Movies	CSV Data Sou	Movies	System Admi	Nov 1, 2022	Nov 1, 2022	Д	:
E Tickets	CSV Data Sou	Tickets	System Admi	Nov 2, 2022	Nov 2, 2022	Д	:
				Rows per page:	25 👻 1-4 0	f4 <	>

3. Select the checkbox next to the data source that you want to ingest. Options are enabled at the bottom of the screen. For example:

	Data Sources	s S	chemas						
Ŧ	Q Search		₹ •	≡			A	dd Data Sou	urce 👻
•	Title 🛧	Description	Туре	Schema	Creator	Created Date	Updated Date	Actions	:
	≡ 08-QA		Database Dat	books, north	System Admi	Nov 1, 2022	Nov 1, 2022	П	:
	E Flights		CSV Data Sou	Flights	System Admi	Nov 1, 2022	Nov 1, 2022	П	:
	Movies		CSV Data Sou	Movies	System Admi	Nov 1, 2022	Nov 1, 2022	П	:
	Tickets		CSV Data Sou	Tickets	System Admi	Nov 2, 2022	Nov 2, 2022	Д	:
						Rows per page:	25 <del>v</del> 1-4 of	4 <	>
<u>1 lte</u>	ms Selected	0		Delete	Bookmark	Create Graph	mart Add To	Package	•

#### Тір

Users with the **Batch Direct Data Loading** permission can select multiple data sources to ingest.

4. Click the Create Graphmart button. The Create Graphmart dialog box is displayed. For

example:

Create Graphmart		
This will create a graphmart selected data sources. Each of each data source.	containing da layer will cont	ta layers for each of your ain a graph representation
Graphmart Name		
Movies		
Description		
Advanced		~
	CANCEL	CREATE GRAPHMART

## Note

If the selected data source includes more than one schema, the number of schema

instances is shown at the top of the screen. For example:



Clicking Show displays the selected schemas. For example:



If you would like to exclude one or more schemas, clear the checkbox for each schema that you want to exclude.

5. On the Create Graphmart dialog box, the Graphmart Name is populated with the name of the selected source. If multiple sources were selected, the Graphmart Name is blank. Edit the Graphmart Name if necessary and add an optional Description for the new graphmart.

- If you would like to configure any of the advanced settings, click Advanced to view the options. For details about the each of the Advanced settings, see Direct Load Advanced Settings Reference.
- 7. When you have finished configuring the workflow, click Create Graphmart. The new graphmart is created and activated and the data layers and steps are generated according to the chosen strategy. A managed model is also generated. See Managed Model Concepts for information. If you chose to export a dataset, the new dataset is also added to the Datasets catalog.

Once the graphmart is online, the data can be analyzed. See Access & Analyze for next steps. Or see Working with Graphmarts for information about managing graphmarts.

# Adding a Data Source to an Existing Graphmart

Follow the steps below if you want to add a data source to an existing graphmart.

#### Note

AnzoGraph uses the Graph Data Interface (GDI) Java plugin to connect directly to sources. For file-based sources, make sure the source files are available to AnzoGraph on the shared file store. For databases, if you have configured custom drivers to access those sources in Anzo, the same drivers need to be added to AnzoGraph. For instructions, see <u>Deploy Optional</u> Drivers for Accessing Custom Database Sources in the Deployment Guide.

- 1. If necessary, add the data source. See Adding Data Sources for instructions.
- 2. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:
| <b>Y</b> | <b>X</b> Search  | =           |                |            |                         | Add Graphn | nart 👻 |
|----------|------------------|-------------|----------------|------------|-------------------------|------------|--------|
| •        | Title $\uparrow$ | Description | Status         | Statements | Creator                 | Actions    | ÷      |
|          | 😻 Books          |             | C Ready to use | 156        | System Administrator    | Д          | :      |
|          | Flights          |             | C Ready to use | 265,339    | System Administrator    | Д          | :      |
|          | ≋ Movies         |             | Ready to use   | 104,166    | System Administrator    | Д          | :      |
|          | © Northwind      |             | C Ready to use | 131,386    | System Administrator    | Д          | :      |
|          |                  |             | C Ready to use | 3,696,548  | System Administrator    | Д          | :      |
|          |                  |             |                |            |                         |            |        |
|          |                  |             |                |            |                         |            |        |
|          |                  |             |                | Row        | rs per page: 25 👻 1-5 o | 15 <       | >      |

3. Click the name of the graphmart that you want to add the data source to. The Overview is displayed. For example:

*	Movies Not Versioned 🖄 INACTIVE	ACTIVE 🔗 Ready to use	AZG Static		Save a Copy	Create Dashboard
_	Overview	Explore	Datasets	Data Layers	Dashboards	Data on Demand 💙
escription						General
one						Type : Granhmart
ad Priorit	ly .					Creator : System Administrat
one					La	st Accessed : 6 hours ago
					Struct	ast Updated: 6 nours ago are Modified: 7 hours ago
etrics						Released: 7 hours ago
ofile data	upon activation 🛛 🗩					
isable L	oad Counts				http://ca	mbridgesemantics.com/Graph
sable pop	ulating counts during load	d operations 🖉 🗯			Dynamio	Resource Log
ata Loa	ding Settings				А	nzoGraph Server Details
ad layers	that do not fail 🛛 💻					AnzoGraph : AZG
inore Sc	ource Errors					Status : Ready to use
nore any s	step or view's source layer	if they failed to load			La	st Accessed : 2 minutes ago emory Lised : 281 25 MB/ 38 )
					M	emory Total: 11.7 GB
eave Gr	aphmart Online Dur	ing Refresh				
hen refres	shing layers, leave graphm	art and layer online.			Inactivit	y Deactivate Timeout
lanual F	Refresh Graphmart				None	
nce loaded	d, changes only pushed to	AnzoGraph manually.(Only affect	s Journal Based Data) 🦳 💻			

- 4. If necessary, activate the graphmart. Graphmarts must be online to be able to add data sources to them.
- 5. Click the **Data Layers** tab. On the Data Layers tab, click **Add** on the right side of the screen and select **Data Source**. The Select Data Source dialog box is displayed. For example:

Select Data Source					
Select data sources and sample size to gene	rate layers and lo	ad steps.			
Title ↑ Description	Туре	Schema	Creator	Created Date	Updated Date
≡ DB-QA	Database Data Sou	northwind	System Administra	Nov 16, 2022	Nov 16, 2022
E Flights	CSV Data Source	Flights	System Administra	Nov 16, 2022	Nov 16, 2022
Movies	CSV Data Source	Movies	System Administra	Nov 16, 2022	Nov 16, 2022
Tickets	CSV Data Source	Tickets	System Administra	Nov 16, 2022	Nov 16, 2022
			Rows per page:	25 <del>→</del> 1-4 of 6	4 < >
Advanced					~
				CANCEL CREAT	E LAYER & STEP

6. Select the checkbox next to the data source that you want to add to the graphmart. (Users with the **Batch Direct Data Loading** permission can select multiple sources.) The selected schema or schemas are shown at the bottom of the screen. For example, in the image a below, a source with two schema instances is selected.

<b>T</b>	<b>Q</b> Search		= •	1 =					
•	Title	Description	Туре	Schema	Version	Creator	Created Date	Updated Date	11
	∃ DB-QA		Database Dat	emrdbsmall,		System Admi	Nov 16, 2022	Nov 16, 2022	ŕ
	B Flights		CSV Data So	Flights		System Admi	Nov 16, 2022	Nov 16, 2022	
	Movies		CSV Data So	Movies		System Admi	Nov 16, 2022	Nov 16, 2022	
-	🖬 Tinkata		CCI/ Data Co	Tinkata		Ovetam Islani	Nov 16 1000	Nov 16, 2022	Ψ - F
						Rows per page	e: 25 🛩	1-4 of 4 <	$\rightarrow$
<u>1 Item</u>	n <u>s Selected</u> (	3							
2 of 2 Sci	hema Instances	selected. SH	w						
Advar	nced								~
							CANCEL	CREATE LAYER 8	STEP

Clicking **Show** displays the selected schemas. For example:



If you would like to exclude one or more schemas, clear the checkbox for each schema that you want to exclude.

- If you want to configure any of the advanced settings, click Advanced to view the options. For details about the each of the Advanced settings, see Direct Load Advanced Settings Reference.
- 8. When you have finished configuring the workflow, click **Create Layer & Step**. The new layer is created and the steps are generated according to the chosen strategy. A managed model is also generated for the layer. See Managed Model Concepts for information.

Once the layer is online, the data can be analyzed. See Access & Analyze for next steps. Or see Working with Graphmarts for information about managing graphmarts.

## **Direct Load Advanced Settings Reference**

This topic describes the Advanced options that are available when you create or configure the Direct Data Load workflow to load a data source via auto-generated data layers.

Advanced		^
Graphmart Options		
Export to Dataset	Connections 🗌 Profile Data	
Lawar Constation Strategies		
cayer Generation Strategies		
Each strategy results in the same graph created for each strategy are stored in the stored in the strategy are stored in the stored in the strategy are stored in the stored in the stored in the stored in the stored in the stored in the stored in the	data but uses a different method to pe teir own layer. Only one of these layers	form the load. The queries should be enabled at a time.
Single Step 🗌 Multiple Ste	eps	
Enabled Layer		
Single Step		*
Ontology URI optional		
Used as the base for any class or property U	Ris in the generated model.	
Load Options	Sampling Limit	
Load Options Enable Partitioning	Sampling Limit	

- Graphmart Options
  - Export to Dataset
  - Find Connections
  - Profile Data
- Layer Generation Strategies
  - Single Step
  - Multiple Steps
  - Enabled Layer
- Ontology URI
- Enable Partitioning
- Sampling Limit

## **Graphmart Options**

The Graphmart Options are available when you create a graphmart from a data source (as described in Creating a Graphmart from a Data Source). They are not available when you add a data source to an existing graphmart (as described in Adding a Data Source to an Existing Graphmart).

### **Export to Dataset**

This setting controls whether the automated workflow generates a dataset by exporting the graphmart. This option does not create an Export Step in the graphmart but it does generate a dataset in the selected data store and adds the dataset to the Datasets catalog. If you leave Export to Dataset disabled, a dataset is not automatically generated. For information about creating an Export Step to export a graphmart to a dataset at a later time, see Export Data to an FLDS (Export Step). If you enable Export to Dataset, the following settings are displayed. The list below the image describes the settings.

Export to Dataset	
Advanced	^
Dataset Anzo Data Store *	
Store	*
Post Dataset Export Action *	
<ul> <li>Leave Graphmart Activated</li> </ul>	
<ul> <li>Deactivate Graphmart</li> </ul>	
O Delete Graphmart	
Dataset Format	
ttl.gz	
Set the dataset's default data format, such as ttl or ttl.gz.	

- Dataset Anzo Data Store: This required setting specifies the Anzo Data Store to export the dataset to. The data store must be a location on the shared file store that AnzoGraph has access to.
- Post Dataset Export Action: This required setting specifies how to treat the graphmart after the Export Step is processed. To leave the graphmart online, select Leave Graphmart Activated (the default value). To retain the graphmart but disable it and remove the data from

AnzoGraph, select **Deactivate Graphmart**. And to designate the graphmart as temporary and remove it after the dataset is exported, select **Delete Graphmart**.

• **Dataset Format**: This setting specifies the file format for the RDF TTL files that are generated, i.e., whether they are compressed or not. The valid options are **ttl** for uncompressed and **ttl.gz** for compressed.

### **Find Connections**

This optional setting specifies whether to find relationships between tables in the schema (or between data sources if multiple sources are selected). Finding connections is useful if the schema does not define primary and foreign key relationships and you want the Graph Data Interface to create the connections.

## Note

When **Find Connections** is enabled, two models are created, one that contains the classes and properties and one that contains only the connections. To view the complete model, both models must be added to the Working Set in the Model viewer.

#### **Profile Data**

This optional setting specifies whether to generate a Data Profile after the graphmart is activated. For information about the metrics that are run when a profile is generated, see Data Profiling Metrics.

#### **Layer Generation Strategies**

These settings control the strategy to use for auto-generating the data layer queries in the graphmart. Each strategy produces the same graph data but uses a different method for structuring the queries that produce the data. You can select both options if you want and review the resulting layers and steps. Each option results in a separate layer. However, only one of the resulting layers can be enabled by default.

### Single Step

This is the default strategy. When **Single Step** is selected, a layer with a single Direct Load Step is created. The single query loads all tables from the selected schema or schemas. The generated query is an RDF and Ontology Generator query. See Onboarding Data with a Direct Load Step to learn more about the GDI RDF and Ontology Generator.

#### **Multiple Steps**

When Multiple Steps is selected, the layer has a separate Direct Load Step for each table in the selected schema or schemas. The generated query in each step is also an RDF and Ontology Generator query. With this strategy, you can enable and disable certain steps to control which tables are included in the graphmart.

#### **Enabled Layer**

This required setting specifies the data layer that should be enabled by default when the graphmart is activated.

#### **Ontology URI**

This optional setting specifies the custom URI to use for the model that is automatically generated. The value must be a valid URI without a hash (#) or slash (\) character at the end. If you do not specify a custom URI, the Graph Data Interface generates a URI in the following format:

http://cambridgesemantics.com/Layer/<layer\_ID>/Model

#### **Enable Partitioning**

This option specifies whether to enable file partitions for file-based data sources. When file partitions are enabled, files will be partitioned and ingested in parallel for increased performance.

#### Note

Multiline CSV files may fail to load when this option is enabled. If the following type of error message is returned when onboarding files, disable the Enable File Partitions setting:

File uses multiline records and cannot be segmented. Please disable segmenting for this file.

## Sampling Limit

This optional setting specifies the number of rows to scan before inferring the data types for each column.

# **Onboarding or Virtualizing Data with SPARQL Queries**

The topics in this section provide information about exploring, analyzing, virtualizing, and ingesting data by writing federated SPARQL queries that invoke the Graph Data Interface (GDI), an AnzoGraph plugin that enables you to connect directly to sources and control all aspects of the extract, load, and transform process. Depending on the type of query you write, you can ingest data into Anzo or create a virtual graph that accesses the source only when it is needed without ingesting the data.

## Tip

This section focuses on writing your own SPARQL queries to read, ingest, or virtualize data from various sources. The GDI can also be used to automatically generate models, data layers, and steps with ingestion queries that you can edit. For more information about the automated workflow, see Onboarding Data with the Automated Workflow.

In this section:	
Introduction to the GDI	
GDI Concepts and Basic Usage	
Options for Data Types, Data Linking, and Models	
Advanced Usage by Data Source Type	
GDI Property Reference	

## Introduction to the GDI

The Graph Data Interface (GDI) (sometimes called the Data Toolkit) is an extremely flexible and configurable AnzoGraph plugin that enables users to access a variety of data sources via federated SPARQL queries. Depending on the type of query you write, i.e., whether it is an INSERT query against the GDI service or a CONSTRUCT query against the view or virtualized service, you can ingest source data into Anzo or create a virtual graph that accesses the source only when it is needed without ingesting the data into Anzo.

The GDI has built-in, native support for various file format types, HTTP/REST endpoints, and common database types. Internally, the GDI API has a records-oriented view of data. This view enables the GDI to bridge graph operations to operations for data in other formats. Though the GDI views the source as rows in a table, ultimately it has the capability to convert the records to graph format, enabling the data to be incorporated into data layers to augment existing data.

## Tip

When you query a source such as a database, the GDI service leverages that source to retrieve only the data that it needs for the query. Unlike a JDBC driver, the GDI service does not need to retrieve all values and then complete an often time-consuming step to filter the results.

## **Supported Data Sources**

Source	Description
HTTP/REST Endpoints	The GDI natively supports reading or ingesting data from HTTP/REST endpoints.
Databases	<ul> <li>Cambridge Semantics supplies JDBC drivers for the following databases:</li> <li>Databricks</li> <li>H2</li> <li>IBM DB2</li> <li>Microsoft SQL Server</li> <li>MariaDB</li> <li>Oracle</li> <li>PostgreSQL</li> <li>SAP Sybase (jTDS)</li> </ul>

This table below lists the data sources, file systems, and applications that the GDI supports.

Source	Description
	• Snowflake To extend the service to access other databases, additional JDBC drivers can be added to AnzoGraph. For information about acquiring additional JDBC drivers, contact your Cambridge Semantics Customer Success manager. For instructions on deploying drivers, see <u>Deploy Optional Drivers</u> for Accessing Custom Database Sources in the Deployment Guide.
File Formats	<ul> <li>The following file types are supported:</li> <li>CSV and TSV</li> <li>JSON and NDJSON</li> <li>Parquet</li> <li>SAS (SAS Transport XPT and SAS7BDAT formats)</li> <li>XML</li> <li>Raw text format</li> </ul>
File Systems	<ul> <li>The following types of file storage systems are supported:</li> <li>Amazon S3</li> <li>FTP &amp; FTPS</li> <li>Google Cloud Storage</li> <li>HDFS (Kerberized HDFS is not supported at this time.)</li> <li>NFS</li> <li>SFTP</li> <li>WebDAV</li> </ul>
Applications	Queries against Elasticsearch and Kafka applications are supported.

## **GDI Concepts and Basic Usage**

The topics in this section help you get to know the Graph Data Interface (GDI) by introducing you to the main concepts and giving a general overview of the query syntax, available properties, and functionality that is applicable across query and data source types.

- Getting Started with GDI Queries
- Onboarding Data with a Direct Load Step
- Reading a Data Source's Metadata
- Paginating Requests
- Binding and Hierarchy Concepts
- Incremental Onboarding Concepts

#### **Getting Started with GDI Queries**

This topic provides details about the structure to use when writing GDI queries. It focuses on the properties that are common to all types of data sources. It also includes example queries that demonstrate the data integration capabilities for different types of sources.

#### Tip

Rather than manually writing complex queries, you can use the GDI to automatically generate graphs and ontologies by including a few key statements in a relatively simple query. For information, see Onboarding Data with a Direct Load Step.

- GDI Query Syntax
- GDI Query Examples

#### **GDI Query Syntax**

The following query syntax shows the structure of a GDI query. The clauses, patterns, and placeholders that are links are described below.

```
# PREFIX Clause
PREFIX s: <a href="http://cambridgesemantics.com/ontologies/DataToolkit#">http://cambridgesemantics.com/ontologies/DataToolkit#</a>>
```

```
PREFIX rdf:
             <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd:
             <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
PREFIX dc:
             <http://purl.org/dc/elements/1.1/>
# Result Clause
{
  [ GRAPH ${targetGraph} { ]
  triple patterns
 [ } ]
}
[ ${usingSources} ]
WHERE
{
   # SERVICE Clause: Include the following service call when reading or inserting data.
   SERVICE [ TOPDOWN ] <http://cambridgesemantics.com/services/DataToolkit>
   # View SERVICE Clause: Or use the service call below when constructing a view.
    SERVICE <http://cambridgesemantics.com/services/DataToolkitView>(${targetGraph})
    {
      ?data a s:source type ;
# Based on the source type, additional connection and input parameters are
# available. The options below are valid for all sources. For source-related
# options, see GDI Property Reference.
        s:url "string" ;
        [ s:username "string" ; ]
        [ s:password "string" ; ]
        [ s:timeout int ; ]
        [ s:batching boolean | int ; ]
        [ s:paging [ pagination options ; ]
        [ s:concurrency int | [ list of properties ] ; ]
        [ s:rate int | "string" ; ]
        [ s:locale "string" ; ]
        [ s:sampling int ; ]
        [ s:selector "string" | [ list ] ; ]
        [ s:model "string" ; ]
        [ s:key ("string") ; ]
        [ s:reference [ s:model "string" ; s:using ("string") ]
```

```
[ s:formats [ datatype_formatting_options ] ; ]
[ s:normalize boolean | [ normalization_rules ] ; ]
[ s:count ?variable ; ]
[ s:offset int ; ]
[ s:limit int ; ]
# Mapping variables
?mapping_variable ( [ "binding" ] [ datatype ] [ "datetime_format" ] ) ;
... ;
.
# Additional clauses such as BIND, VALUES, FILTER
}
```

## Note

}

#### For readability, the parameters below exclude the base URI

<http://cambridgesemantics.com/ontologies/DataToolkit#> as well as the s: prefix. As shown in the examples, however, the s: prefix or full property URI does need to be included in queries.

Option	Туре	Description
PREFIX Clause	N/A	The PREFIX clause declares the standard and custom prefixes for GDI service queries. Generally, queries include the prefixes from the query template (or a subset of them) plus any data-specific declarations.
Result Clause	N/A	The result clause defines the type of SPARQL query to run and the set of results to return, i.e., whether you want to read (SELECT or CONSTRUCT) from the source or ingest the data into Anzo (INSERT).
GRAPH \${targetGraph}	N/A	Include the GRAPH keyword and target graph parameter \${targetGraph} when writing an INSERT query to ingest data into a graphmart. Anzo automatically populates the query with the appropriate target URIs when the query runs.

Option	Туре	Description
\${usingSource s}	N/A	Include the source graph parameter \$ {usingSources} when writing a "topdown" query that passes values from the data that is in the graphmart to the data source. Anzo automatically populates the query with the appropriate FROM clauses when the query runs. When passing literal values to the remote source, you do not need to include the source graph parameter. The SERVICE Clause description below includes more information about passing input to data sources.
SERVICE Clause	N/A	<pre>Include the SERVICE call SERVICE [ TOPDOWN ] <http: cambridgesemantics.com="" datatool="" kit="" services=""> to invoke the GDI service when you are running a SELECT, INSERT, or CONSTRUCT query that is not creating a view. When writing a CONSTRUCT query in a View Step, use the DataToolkitView service call, as described below in View SERVICE Clause.</http:></pre>
		Include the optional TOPDOWN keyword when you want to pass input values from the graphmart to the data source. When you include TOPDOWN in the service call, it indicates that the rest of the query produces values to send to the source. In this case, the GDI makes repeated calls to pass in each of the specified values and retrieve the data that is based on those values.
View SERVICE Clause	N/A	When writing a CONSTRUCT query that creates a view of the data (usually in a View Step), include the following SERVICE call: SERVICE <http: cambridgesemantics.com="" datatool<br="" services="">kitView&gt;(\${targetGraph}). Using the DataToolkitView call optimizes query execution because it tells the GDI to inspect the query and determine which filters to push to the data source. It also limits the result set and retrieves only the data that is needed,</http:>

Option	Туре	Description
		i.e., the source data is fully mapped but all of the mapped data is not necessarily returned.
source_type	object	The ?data a s:source_type triple pattern specifies the type of data source that the query will run against. For example, ?data a s:DbSource, specifies that the source type is a database. The list below describes the available types:
		DbSource to connect to any type of database.
		<ul> <li>FileSource for flat files. The supported file types are CSV and TSV, JSON, NDJSON, XML, Parquet, and SAS (SAS Transport XPT and SAS7BDAT formats). The GDI automatically determines the file type from the file extensions. When querying file sources, make sure that the files are accessible to both Anzo and AnzoGraph.</li> <li>HttpSource to connect to HTTP endpoints.</li> <li>ElasticSource to connect to Elasticsearch indexes on an Elasticsearch server.</li> <li>KafkaSource to connect to Kafka streaming sources.</li> <li>MetadataSource for metadata discovery.</li> </ul>
		<b>Tip</b> Certain connection and input parameters are available based on the specified source type. For details about the options for your source, see GDI Property Reference.
url	string	This property specifies the URL for the data source, such as the database URL, Elasticsearch URL, or HTTP endpoint URL. For file-based sources, the url property specifies the file system

Туре

#### Description

location of the source file or directory of files. When specifying a directory (such as s:url "/opt/shared-files/loads/"), the GDI loads all of the file formats it recognizes. To specify a directory but limit the number or type of files that are read, you can include the pattern and/or maxDepth properties described in FileSource Properties.

#### Important

For security, it is a best practice to reference connection information (such as the url, username, and password) from a Query Context so that the sensitive details are abstracted from any requests. In addition, using a Query Context makes connection details reusable across queries. See Using Query Contexts for more information. For example, the triple patterns below reference keys from a Query Context:

```
?data a s:DbSource ;
s:url "{{@db.eca4bfa8...ff9a.url}}" ;
s:username "{{@db.eca4bfa8...ff9a.user}}"
;
s:password "
{{@db.eca4bfa8...ff9a.password}}" ;
```

username	string	If authentication is required to access the source, include this property to specify the user name.
password	string	This property lists the password for the given username.
timeout	int	This property can be used to specify the timeout (in milliseconds) to use for requests against the source. For example, s:timeout

Option	Туре	Description			
		5000 configures a 5 second timeout.			
batching	boolean or int	This property can be used to disable batching, or it can be used to change the default the batch size. By default, batching is set to 5000 (s:batching 5000). To disable batching, you can include s:batching false in the query. Typically users do not change the batching size. However, it can be useful to control the batch size when performing updates. To configure the size, include s:batching int in the query. For example, s:batching 3000.			
paging	RDF list	This property can be used to configure paging so that the GDI can access large amounts of data across a number of smaller requests. For details about the paging property, see Paginating Requests.			
concurrency	int or RDF list	This property can be included to configure the maximum level of concurrency for the query. The value can be an integer, such as s:concurrency 8. If the value is an integer, it configures a maximum limit on the number of slices that can execute the query. For finer-grained control over the number of nodes and slices to use, concurrency can also be included as an object with limit, nodes, and/or executorsPerNode properties. For example, the following object configures a concurrency model that allows a maximum of 24 executors distributed across 4 nodes with 8 executors per node: s:concurrency [ s:limit 24 ; s:nodes 4 ; s:executorsPerNode 8 ; ] ;			

Option	Туре	Description			
rate	int or string	This property can be included to control the frequency with which a request is sent to the source. The limit applies to the number of requests a single slice can make. If you specify an integer for the rate, then the value is treated as the maximum number of requests to issue per minute. If you specify a string, you have more flexibility in configuring the rate. The sample values below show the types of values that are supported:			
		s:rate "90/minute" ;			
		<pre>s.fate "30 per minute", s:rate "200000 every week" ; s:rate "10000 every 6 hours" ;</pre>			
		To enforce the rate limit, the GDI introduces a sleep between requests that is equal to the rate delay. The more executing slices, the longer the rate delay needs to be to enforce the limit in aggregate.			
		Given the example of s:rate "90/minute", the GDI would optimize the concurrency and only use 1 slice for execution with a rate delay of 666ms between requests. If s:rate "240/minute", the GDI would use 3 executors with a rate delay of 750ms between requests.			
locale	string	This property can be used to specify the locale to use when parsing locale-dependent data such as numbers, dates, and times.			
sampling	int	This property can be used to configure the number of records in the source to examine for data type inferencing.			
selector	string or RDF list	This property can be used as a binding component to identify the path to the source objects. For example, s:selector "Sales.SalesOrderHeader" targets the SalesOrderHeader			

Option	Туре	Description				
		table in the Sales schema. For more information about binding components and the selector property, see Using Binding Trees and Selector Paths.				
model	string	This property defines the class (or table) name for the type of data that is generated from the specified data source. For example, s:model "employees". Model is optional when querying a single source. If your query targets multiple sources, however, and you want to define resource templates (primary keys) and object properties (foreign keys), you must specify the model value for each source.				
key	string	This property can be used to define the primary key column for the source file or table. This column is leveraged in a resource template for the instances that are created from the source. For example, s:key ("EMPLOYEE_ID"). For more information about key, see Data Linking Options.				
reference	RDF list	This property can be used to specify a foreign key column. The reference property is an RDF list that includes the model property to list the target table and a using property that defines the foreign key column. For more information about reference, see Data Linking Options.				
formats	RDF list	To give users control over the data types that are used when coercing strings to other types, this property can be included in GDI queries to define the desired types. In addition, it can be used to describe the formats of date and time values in the source to ensure that they are recognized and parsed to the appropriate date, time, and/or dateTime values. For details about the formats property, see Data Type Formatting Options.				

Option	Туре	Description					
normalize	boolean and/or RDF list	To give users control over the labels and URIs that are generated, the GDI offers several options for normalizing the model and/or the fields that are created from the specified data source(s). For details about the normalize property, see Model Normalization Options.					
count	variable	If you want to turn the query into a COUNT query, you can include this property with a <code>?variable</code> to perform a count. For example, <code>s:count</code> ?count.					
offset	int	This property can be used to offset the data that is returned by a number of rows.					
limit	int	You can include this property to limit the number of results that are returned. s:limit maps to the SPARQL LIMIT clause.					
mapping_ variable	variable	The mapping variables, in <code>?mapping_variable</code> (["binding"] [datatype] ["datetime_format"]) format, define the triple patterns to output. When the specified <code>?variable</code> matches the source column name, the GDI uses the variable as the source data selector. If you specify an alternate variable name, a binding needs to be specified to map the new variable to the source. You also have the option to transform the data using the datatype and datetime_format options.					
		<b>Note</b> The parentheses around the binding, data type, and format specifications are not required but are included in this document for readability.					
binding	string	The binding is a literal value that binds a ?mapping_variable to a					

#### Туре

#### Description

source column. If you specify a ?variable that matches the source column name, then that variable name is the data selector and it is not necessary to specify a binding. If you specify an alternate variable name or there is a hierarchical path to the source column, then the binding is needed to map the new variable to that source column.

For example for a flat source like CSV, the following pattern simply binds the source column AIRLINE to the lowercase variable ?airline: ?airline ("AIRLINE"). For a database source, this example binds the ?subject variable by navigating to the SUBJECT column in the FILM table in the dbo schema: ?subject ("dbo.FILM.SUBJECT"). And for an HTTP source, this example binds the ?time variable to the time object under the minutely data path: ?time

("minutely.data.time").

#### Note

For **FileSource** and **HttpSource**, periods (.), forward slashes (/), and brackets ([]) are parsed as path notation. Therefore, if a source column name includes any of those characters they must be escaped in the binding. Use two backslashes (\\) as an escape character. For example, if a column name is **average/day**, the variable and binding pattern could be written as ?averagePerDay

("average\\/day").

For **DbSource**, database, schema, and table names in bindings are parsed according to the specific rules for that database type. You do not need to escape characters in database names. However, database names with

Option	Туре	Description			
		<pre>characters that do not match (_ A-Z a-z) (_ A-Z a- z 0-9) * should be quoted, such as ("'Adventure.Works'.Sales.'Daily.Totals'").</pre>			
datatype	URI	The datatype is the data type to convert the column to. If you do not specify a data type, the GDI infers the type. The GDI supports the following types: xsd:int, xsd:long, xsd:float, xsd:double, xsd:boolean, xsd:time, xsd:dateTime, xsd:date, xsd:duration, xsd:timeDuration, xsd:yearMonthDuration, xsd:gMonthDay, xsd:gMonth, xsd:gYearMonth, xsd:anyURI			
datetime_ format	string	This option is used to specify the format to use for date and time data types. The GDI supports Java date and time formats. Specify days as "d," months as "M," and years as "y." For the time, specify "H" for hours, "m" for minutes, and "s" for seconds. For example, " $yyyyMMdd HH:mm:ss"$ or "ddMMMyy" to display date values such as "01JAN19." Note The GDI's default base year is 2000. If the source data has years with only two digits, such as $02-04-99$ , the GDI prepends 20 to the digits. The value $02-04-99$ is parsed to $02-04-2099$ . To specify an alternate base year to use for two-digit values, you can include the notation $^nnnn$			
		(e.g., ^1900) in the format value. For example, to set the base year to 1900 instead of 2000, use a format value such as xsd:date "dd-MMM-yy^1900" or xsd:date			

Option	Туре	Description
		"dd-MMM-yy^1990". When one of those values is specified, 02-04-99 is parsed to 02-04-1999.

### **GDI Query Examples**

The query below reads data from a sample HTTP source that compiles worldwide weather statistics. The source has several models available for retrieving data that is current, daily, historical, etc. To target current data, the query includes <code>s:selector "currently"</code> as an input parameter. In addition, the query demonstrates the use of the "topdown" functionality, where the query sends values to the source to narrow the results. The query includes the TOPDOWN keyword in the GDI service call, and the VALUES clause specifies the latitude and longitude values for the cities to return data for. In addition, since this sample source requires parameters to be specified in the connection URL, the <code>s:url</code> value includes <code>?lat</code> and <code>?long</code> as parameters as part of the value.

```
PREFIX s:
               <http://cambridgesemantics.com/ontologies/DataToolkit#>
               <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdf:
PREFIX rdfs:
               <http://www.w3.org/2000/01/rdf-schema#>
               <http://www.w3.org/2001/XMLSchema#>
PREFIX xsd:
               <http://www.w3.org/2002/07/owl#>
PREFIX owl:
               <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX anzo:
               <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
PREFIX zowl:
PREFIX dc:
               <http://purl.org/dc/elements/1.1/>
               <http://example.org/ontologies/City#>
PREFIX ex:
SELECT
       ?city ?state ?summary ?temp ?rainChance
       ?humidity ?pressure ?windSpeed
WHERE
{
       SERVICE TOPDOWN <http://cambridgesemantics.com/services/DataToolkit>
    {
         ?data a s:HttpSource ;
         s:url "https://sampleEndpoint.com/forecast/{{?lat}},{{?long}}";
         s:selector "currently" ;
         ?lat ("latitude") ;
         ?long ("longitude") ;
```

```
?temp ("temperature") ;
         ?rainChance ( "precipProbability" ) ;
         ?humidity () ;
         ?pressure () ;
         ?windSpeed () .
    }
    VALUES( ?city ?state ?lat ?long )
    {
         ( "Lakeway" "TX" 30.374563 -97.975892 )
         ( "Boston" "MA" 42.358043 -71.060415 )
         ( "Seattle" "WA" 47.590720 -122.307053 )
         ( "Chicago" "IL" 41.837741 -87.823296 )
         ( "Hilo" "HI" 19.702040 -155.090312 )
    }
}
ORDER BY ?city
```

The query returns the following results:

city		state		summary		temp		rainChance		humidity		pressure		windSpeed
	•+•		+-		-+-		-+-		•+•		+ -		+-	
Boston		MA		Overcast		79.81		0		0.6		1018.7		7.71
Chicago		IL		Clear	Ι	81.7		0		0.52		1021.1		5.13
Hilo		HI		Partly Cloudy	I	72.6		0.13		0.79		1018.6		4.86
Lakeway		TX		Partly Cloudy	Ι	92.43		0		0.48		1013.3		10.85
Seattle		WA		Mostly Cloudy	Ι	61.82		0		0.76		1018.2		4.57
5 rows														

The example below ingests data into a data layer from a database source using an INSERT query in

## a Direct Load Step.

PREFIX	s:	<http: cambridgesemantics.com="" datatoolkit#="" ontologies=""></http:>
PREFIX	rdf:	<http: 02="" 1999="" 22-rdf-syntax-ns#="" www.w3.org=""></http:>
PREFIX	rdfs:	<http: 01="" 2000="" rdf-schema#="" www.w3.org=""></http:>
PREFIX	xsd:	<http: 2001="" www.w3.org="" xmlschema#=""></http:>
PREFIX	owl:	<http: 07="" 2002="" owl#="" www.w3.org=""></http:>
PREFIX	anzo:	<http: 07="" 2008="" anzo#="" ontologies="" openanzo.org=""></http:>
PREFIX	zowl:	<http: 05="" 2009="" anzoowl#="" ontologies="" openanzo.org=""></http:>
PREFIX	dc:	<http: 1.1="" dc="" elements="" purl.org=""></http:>
PREFIX	:	<http: example.com="" kl_hosp#="" ontologies=""></http:>
INSERT		
ſ		

```
GRAPH ${targetGraph}
    {
       ?InputEvent cv a :InputEvent cv ;
           :row_id ?row_id ;
           :subject id ?subject id ;
           :hadm id ?hadm id ;
           :icustay id ?icustay id ;
           :charttime ?charttime ;
           :itemid ?itemid ;
           :amount ?amount ;
           :amountuom ?amountuom ;
           :rate ?rate ;
           :rateuom ?rateuom ;
           :storetime ?storetime ;
           :cgid ?cgid ;
           :orderid ?orderid ;
           :linkorderid ?linkorderid ;
           :stopped ?stopped ;
           :newbottle ?newbottle ;
           :originalamount ?originalamount ;
           :originalamountuom ?originalamountuom ;
           :originalroute ?originalroute ;
           :originalrate ?originalrate ;
           :originalrateuom ?originalrateuom ;
           :originalsite ?originalsite .
    }
}
WHERE
{
       SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
       ?data a s:DbSource ;
           s:url "{{@db.eca4bfa83481f3638b93ab5fdf93ff9a.url}}" ;
            s:username "{{@db.eca4bfa83481f3638b93ab5fdf93ff9a.user}}"
            s:password "{{@db.eca4bfa83481f3638b93ab5fdf93ff9a.password}}"
           s:selector "kl hosp schema.inputevents cv" ;
           ?row id (xsd:int) ;
           ?subject id (xsd:int) ;
           ?hadm id (xsd:int) ;
           ?icustay id (xsd:int) ;
           ?charttime (xsd:dateTime) ;
           ?itemid (xsd:int) ;
           ?amount (xsd:float) ;
```

```
?amountuom (xsd:string) ;
           ?rate (xsd:float) ;
           ?rateuom (xsd:string) ;
           ?storetime (xsd:dateTime) ;
           ?cgid (xsd:int) ;
           ?orderid (xsd:int) ;
           ?linkorderid (xsd:int) ;
           ?stopped (xsd:string) ;
           ?newbottle (xsd:int) ;
           ?originalamount (xsd:float) ;
           ?originalamountuom (xsd:string) ;
           ?originalroute (xsd:string) ;
           ?originalrate (xsd:float) ;
           ?originalrateuom (xsd:string) ;
           ?originalsite (xsd:string) ;
      BIND(IRI("http://example.com/inputevent cv/{{?row id}}") AS ?InputEvent cv)
      BIND(IRI("http://example.com/patients/{{?subject id}}") AS ?patient)
      BIND(IRI("http://example.com/admissions/{{?hadm id}}") AS ?admission)
    }
}
```

#### The following query ingests airport-related data from a CSV file.

```
PREFIX s:
                <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX rdf:
                <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs:
                <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd:
                <http://www.w3.org/2001/XMLSchema#>
PREFIX owl:
                <http://www.w3.org/2002/07/owl#>
PREFIX anzo:
                <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl:
                <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
PREFIX dc:
                <http://purl.org/dc/elements/1.1/>
INSERT
{
  GRAPH ${targetGraph}
  {
   ?code a <http://anzograph.com/airport> ;
      <http://anzograph.com/airport/name> ?name ;
      <http://anzograph.com/airport/city> ?city ;
      <http://anzograph.com/airport/state> ?state ;
      <http://anzograph.com/airport/latitude> ?lat;
      <http://anzograph.com/airport/longitude> ?long.
  }
```

```
}
WHERE
{
       SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
       ?data a s:FileSource ;
          s:url "/opt/shared-files/airports.csv" ;
          ?iata code ("IATA CODE" xsd:string) ;
          ?name ("AIRPORT" xsd:string) ;
          ?city ("CITY" xsd:string) ;
          ?state ("STATE" xsd:string) ;
          ?lat ("LATITUDE" xsd:double) ;
          ?long ("LONGITUDE" xsd:double).
       BIND(IRI("http://anzograph.com/airport/{{?iata code}}") as ?code)
    }
}
```

#### The query below creates a view of a database source.

```
PREFIX s:
              <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
PREFIX dc: <http://purl.org/dc/elements/1.1/>
PREFIX ont: <a href="http://cambridgesemantics.com/ont/autogen/Rh/MIMIC-III-Data Source/mimic">http://cambridgesemantics.com/ont/autogen/Rh/MIMIC-III-Data Source/mimic</a>
iii schema#>
CONSTRUCT
{
   ?caregiversURI a ont:caregivers ;
      ont:caregivers cgid ?cgid ;
      ont:caregivers description ?description ;
      ont:caregivers label ?label .
}
WHERE
{
   GRAPH ?q
   {
      SERVICE <http://cambridgesemantics.com/services/DataToolkitView>(${targetGraph})
```

```
{
    ?data a s:DbSource ;
    s:url "{{@db.eca4bfa83481f3638b93ab5fdf93dd9a.url}}" ;
    s:username "{{@db.eca4bfa83481f3638b93ab5fdf93dd9a.user}}"
    s:password "{{@db.eca4bfa83481f3638b93ab5fdf93dd9a.password}}"
    s:selector "mimic_iii_schema.caregivers" ;
    ?row_id (xsd:int) ;
    ?cgid (xsd:int) ;
    ?label (xsd:string) ;
    ?description (xsd:string) .
    BIND(IRI("http://cambridgesemantics.com/class/caregivers/{{?row_id}}") AS
?caregiversURI) }
  }
}
```

## **Onboarding Data with a Direct Load Step**

With no mapping required, Anzo's Direct Load Step functionality automatically generates a graph and an ontology (model) for a data source. Using a relatively simple SPARQL query, the direct load option invokes the Graph Data Interface (GDI) RDF and Ontology Generators. The GDI Generators recognize the structure of a data source and automatically generate the necessary statements.

Invoking the Generators is preferable to producing a hand-written query, especially when the structure of the data is very complex, such as a JSON data source with many inner repeating structures or a database with many tables and keys. When the source contains complex structures, the GDI will generate only the required statements and avoid cross-products, optimizing query execution and memory usage. In addition, the GDI Generator parallelizes the load across the AnzoGraph cluster so that a data source (such as a database) can be ingested with a single query.

This topic provides details about invoking the GDI RDF and Ontology Generators. The Generators can be used with all of the supported data source types.

- How to Use the GDI Generator
- GDI Generator Query Syntax
- GDI Generator Example Queries

## How to Use the GDI Generator

To invoke the GDI Generator in a data layer, you add a **Direct Load Step** to the layer. In the Direct Load Step, you compose a SPARQL query that incorporates the GDI Generator parameters as detailed below in GDI Generator Query Syntax.

## Tip

For instructions on adding steps to layers, see Adding Steps to Layers.

# Why Use a Direct Load Step?

It is important to use a Direct Load Step with the RDF and Ontology Generators because it is the only type of step with the ability to manage the generated ontologies (models). An ontology that is generated in a Direct Load Step is automatically registered in Anzo. The registered model is linked to and managed by the data layer that contains the step. If an Ontology Generator query is changed, additional Direct Load Steps are added to the same layer, or the underlying source schema changes, the managed model is automatically updated when the graphmart is reloaded or refreshed. See Managed Model Details below for important details about layer-managed models.

## **Managed Model Details**

Though an ontology that is generated in a Direct Load Step is registered in Anzo and is available for viewing in the Model editor, **the model is owned and managed by the data layer that contains the Direct Load Step**. That means any manual changes made to the model outside of the step, such as from the Model editor, will be overwritten any time the graphmart or layer is refreshed or reloaded. **Do not modify generated managed models except by editing (or adding) Direct Load Step queries**. For information on updating managed models, see Editing a Managed Model.

**There is only one managed model per layer**. If you include multiple Direct Load Steps in the same layer, they will all update the same ontology. This functionality can be useful if you want to align the data and generated model across multiple steps. If you have multiple sources that are not intended to align or update the same model, create separate layers.

**If you delete a layer that includes a managed model, the model is also deleted**. Use caution when referencing a managed model outside of a graphmart. For example, if you create a dataset and reference a managed model when you select the ontology, the reference will break if the data layer that manages the model is deleted.

#### **GDI Generator Query Syntax**

The following query syntax shows the structure of a GDI Generator query. The clauses, patterns, and placeholders that are links are described below.

```
# PREFIX Clause
PREFIX s: <a href="http://cambridgesemantics.com/ontologies/DataToolkit#">http://cambridgesemantics.com/ontologies/DataToolkit#</a>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
#Result Clause
INSERT {
  GRAPH ${targetGraph} {
    ?s ?p ?o .
  }
}
WHERE
{
       SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
      ?data a s:source type ;
# Based on the source type, additional connection and input parameters are
# available. The options below are valid for all sources. For source-related
# options, see GDI Property Reference.
        s:url "string" ;
        [ s:model "class name for this source" ; ]
        [ s:username "string" ; ]
        [ s:password "string" ; ]
        [ s:timeout int ; ]
        [ s:maxConnections int ; ]
        [ s:batching boolean | int ; ]
        [ s:paging [ pagination options ; ]
        [ s:concurrency int | [ list of properties ] ; ]
        [ s:rate int | "string" ; ]
        [ s:locale "string" ; ]
```

```
[ s:sampling int ; ]
     [ s:selector "string" | [ list ] ; ]
     [ s:key ("string") ; ]
     [ s:reference [ s:model "string" ; s:using ("string") ]
     [ s:formats [ datatype formatting options ] ; ]
     [ s:normalize boolean | [ source normalization rules ] ; ]
     [ s:count ?variable ; ]
     [ s:offset int ; ]
     [ s:limit int ; ] .
   # Multiple data sources can be merged if they project a similar set
   # of output variables. Make sure each source has a unique subject variable.
     [ ?unique variable a s:source type ;
       . . .
     . ]
   ?rdf a s:RdfGenerator, s:OntologyGenerator ;
     s:as (?s ?p ?o);
     s:ontology ontology uri ;
     s:base base uri ;
     [ s:normalize boolean | [ global_normalization_rules ] ; ]
   # Additional clauses such as BIND, FILTER
}
```

## Note

}

For readability, the parameters below exclude the base URI

<http://cambridgesemantics.com/ontologies/DataToolkit#> as well as the s: prefix. As shown in the examples, however, the s: prefix or full property URI does need to be included in queries.

Option	Туре	Description
PREFIX Clause	N/A	The PREFIX clause declares the standard and custom prefixes for GDI service queries. Generally, queries include the prefixes from the query template (or a subset of them) plus any data-specific declarations.

Option	Туре	Description	
Result Clause	N/A	The result clause for Direct Load Steps is typically an INSERT query with the graph pattern in the template above.	
		Note It is important to include the <b>GRAPH</b> keyword and target graph parameter <b>\${targetGraph}</b> when you are writing an INSERT query. Anzo automatically replaces the \${targetGraph) parameter with the appropriate target URI(s) when the query runs.	
source_type	object	The ?data a s:source_type triple pattern specifies the type of data source that the query will run against. For example, ?data a s:DbSource, specifies that the source type is a database. The list below describes the available types:	
		• <b>DbSource</b> to connect to any type of database.	
			• FileSource for flat files. The supported file types are CSV and TSV, JSON, NDJSON, XML, Parquet, and SAS (SAS Transport XPT and SAS7BDAT formats). The GDI automatically determines the file type from the file extensions. When querying file sources, make sure that the files are accessible to both Anzo and AnzoGraph.
		HttpSource to connect to HTTP endpoints.	
		• ElasticSource to connect to Elasticsearch indexes on an Elasticsearch server.	

Option	Туре	Description
		<ul> <li>KafkaSource to connect to Kafka streaming sources.</li> </ul>
		<b>Tip</b> Certain connection and input parameters are available based on the specified source type. For details about the options for your source, see GDI Property Reference.
url	string	This property specifies the URL for the data source, such as the database URL, Elasticsearch URL, or HTTP endpoint URL. For file-based sources, the url property specifies the file system location of the source file or directory of files. When specifying a directory (such as s:url "/opt/shared-files/loads/"), the GDI loads all of the file formats it recognizes. To specify a directory but limit the number or type of files that are read, you can include the pattern and/or maxDepth properties described in FileSource Properties.
		Important For security, it is a best practice to reference connection information (such as the url, username, and password) from a Query Context so that the sensitive details are abstracted from any requests. In addition, using a Query Context makes connection details reusable across queries. See Using Query Contexts for more information. For example, the triple patterns below reference keys from a Query Context:

Option	Туре	Description
		<pre>?data a s:DbSource ; s:url "{{@db.eca4bfa8ff9a.url}}" ; s:username " {{@db.eca4bfa8ff9a.user}}" ; s:password " {{@db.eca4bfa8ff9a.password}}" ;</pre>
model	string	This property defines the class (or table) name for the type of data that is generated from the specified data source. For example, s:model "employees". Model is optional when querying a single source. If your query targets multiple sources, however, and you want to define resource templates (primary keys) and object properties (foreign keys), you must specify the model value for each source.
username	string	If authentication is required to access the source, include this property to specify the user name.
password	string	This property lists the password for the given username.
timeout	int	This property can be used to specify the timeout (in milliseconds) to use for requests against the source. For example, s:timeout 5000 configures a 5 second timeout.
maxConnections	int	For database sources, this property can be used to set a limit on the maximum number of active connections to the source. For example, s:maxConnections 16 sets the

Option	Туре	Description				
		limit to 16 connections. The default value is 10.				
batching	boolean or int	This property can be used to disable batching, or it can be used to change the default the batch size. By default, batching is set to 5000 (s:batching 5000). To disable batching, you can include s:batching false in the query. Typically users do not change the batching size. However, it can be useful to control the batch size when performing updates. To configure the size, include s:batching int in the query. For example, s:batching 3000.				
paging	RDF list	This property can be used to configure paging so that the GDI can access large amounts of data across a number of smaller requests. For details about the paging property, see Paginating Requests.				
concurrency	int or RDF list	This property can be included to configure the maximum level of concurrency for the query. The value can be an integer, such as s:concurrency 8. If the value is an integer, it configures a maximum limit on the number of slices that can execute the query. For finer-grained control over the number of nodes and slices to use, concurrency can also be included as an object with limit, nodes, and/or executorsPerNode properties. For example, the following object configures a concurrency model that allows a maximum of 24 executors distributed across 4 nodes with 8 executors per node: s:concurrency [ s:limit 24 ; s:nodes 4 ;				
Option	Туре	Description				
---------------	---------------	---	--	--	--	--
		<pre>s:executorsPerNode 8 ; ] ;</pre>				
rate	int or string	This property can be included to control the frequency with which a request is sent to the source. The limit applies to the number of requests a single slice can make. If you specify an integer for the rate, then the value is treated as the maximum number of requests to issue per minute. If you specify a string, you have more flexibility in configuring the rate. The sample values below show the types of values that are supported: s:rate "90/minute";				
		s:rate "90 per minute" ;				
		s:rate "200000 every week" ;				
		s:rate "10000 every 6 hours" ;				
		To enforce the rate limit, the GDI introduces a sleep between requests that is equal to the rate delay. The more executing slices, the longer the rate delay needs to be to enforce the limit in aggregate.				
		Given the example of s:rate "90/minute", the GDI would optimize the concurrency and only use 1 slice for execution with a rate delay of 666ms between requests. If s:rate "240/minute", the GDI would use 3 executors with a rate delay of 750ms between requests.				
locale string		This property can be used to specify the locale to use when parsing locale-dependent data such as numbers, dates, and times.				

Option	Туре	Description
sampling	int	This property can be used to configure the number of records in the source to examine for data type inferencing.
selector	string or RDF list	This property can be used as a binding component to identify the path to the source objects. For example, s:selector "Sales.SalesOrderHeader" targets the SalesOrderHeader table in the Sales schema. For more information about binding components and the selector property, see Using Binding Trees and Selector Paths.
key	string	This property can be used to define the primary key column for the source file or table. This column is leveraged in a resource template for the instances that are created from the source. For example, s:key ("EMPLOYEE_ID"). For more information about key, see Data Linking Options.
reference	RDF list	This property can be used to specify a foreign key column. The reference property is an RDF list that includes the model property to list the target table and a using property that defines the foreign key column. For more information about reference, see Data Linking Options.
formats	RDF list	To give users control over the data types that are used when coercing strings to other types, this property can be included in GDI queries to define the desired types. In addition, it can be used to describe the formats of date and time values in the source to ensure that they are recognized and parsed to the appropriate date, time,

Option	Туре	Description
		and/or dateTime values. For details about the formats property, see Data Type Formatting Options.
normalize	RDF list	To give users control over the labels and URIs that are generated, the GDI offers several options for normalizing the model and/or the fields that are created from the specified data source(s). For details about the normalize property, see Model Normalization Options.
count	variable	If you want to turn the query into a COUNT query, you can include this property with a <code>?variable</code> to perform a count. For example, <code>s:count</code> ?count.
offset	int	This property can be used to offset the data that is returned by a number of rows.
limit	int	You can include this property to limit the number of results that are returned. s:limit maps to the SPARQL LIMIT clause.
RdfGenerator	object	Include this property to invoke the RDF Generator. If you only want to generate a model without RDF, you can exclude RdfGenerator.
OntologyGenerator	object	Include this property to invoke the Ontology Generator. If you only want to generate RDF without a model, you can exclude OntologyGenerator.
as	N/A	This property provides the variable bindings for the RDF Generator's projection to RDF. Typically the value is s:as (?s ?p ?o) to match the variables in the result clause.

Option	Туре	Description				
ontology	URI	This property specifies the URI to use as the base URI for any generated ontology artifacts. For example, s:ontology <http: abc.com="" myontology="" ontologies="">.</http:>				
		Note In the graphmart, the data layer ID is appended to the ontology URI that is generated. The complete URI is based on the layer and cannot be customized.				
base	URI	This property specifies the base URI for instance data. The base value should NOT end in #. The Generator will add a trailing slash (/) if one does not exist. For example, s:base <http: abc.com=""></http:> .				

**GDI Generator Example Queries** 

This section includes sample queries that may be useful as a starting point for writing your own RDF and Ontology Generator queries.

- Basic Query that Generates RDF and Ontology for a JSON File
- Basic Query that Generates an Ontology for a Directory of CSV Files
- Query that Normalizes and Generates RDF and Ontology for a Database
- Query with Query Context that Normalizes and Generates RDF and Ontology for a Database
- Query for Multiple Sources that Generates RDF and Ontology with Resource Templates and Object Properties

## Basic Query that Generates RDF and Ontology for a JSON File

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
INSERT {
  GRAPH ${targetGraph} {
    ?s ?p ?o .
 }
}
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
      ?data a s:FileSource ;
         s:model "People" ;
         s:url "/opt/shared-files/json/people.json" .
      ?rdf a s:RdfGenerator , s:OntologyGenerator ;
         s:as (?s ?p ?o) ;
         s:ontology <http://cambridgesemantics.com/ontologies/People> ;
         s:base <http://cambridgesemantics.com/data/> .
  }
}
```

### Basic Query that Generates an Ontology for a Directory of CSV Files

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
INSERT {
    GRAPH ${targetGraph} {
        ?s ?p ?o .
    }
}
WHERE {
    SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
        ?data a s:FileSource ;
        s:model "Sales" ;
        s:url "/opt/shared-files/csv/sales" ;
        s:format [
            s:delimiter "," ;
            s:headers true ;
            s:comment "#" ;
```

```
s:quote "\"";
s:maxColumns 22;
].
?rdf a s:OntologyGenerator;
s:as (?s ?p ?o);
s:ontology <http://cambridgesemantics.com/ontologies/Sales>;
s:base <http://cambridgesemantics.com/data/>.
}
```

# Query that Normalizes and Generates RDF and Ontology for a Database

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
INSERT {
  GRAPH ${targetGraph} {
     ?s ?p ?o .
 }
}
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
      ?data a s:DbSource ;
      s:url "jdbc:mysql://10.11.12.9/emrdbbig" ;
      s:username "root" ;
     s:password "sql1@#" ;
     s:normalize [
         s:model [
            s:removeStart "emr " ;
            s:words "activity 'patient complaint' medication observation patient
specialty study" ;
        ];
         s:field [
            s:removePartialPrefix true ;
            s:words "provider description start end drug complaint date medication
normal code
                     observation product active dose generic route admin strength
collection
                     activity home first last status first year birth death directed
complex
```

```
period age flag gender language";
];
].
?rdf a s:RdfGenerator , s:OntologyGenerator ;
s:as (?s ?p ?o) ;
s:ontology <a href="http://cambridgesemantics.com/ontologies/EMR">http://cambridgesemantics.com/ontologies/EMR</a>;
s:base <a href="http://cambridgesemantics.com/EMR">http://cambridgesemantics.com/ontologies/EMR</a>;
}
```

# Query with Query Context that Normalizes and Generates RDF and Ontology for a Database

The query below references a Query Context to supply the username and password for the database connection.

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
INSERT {
  GRAPH ${targetGraph} {
     ?s ?p ?o .
  }
}
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
      ?data a s:DbSource ;
         s:url "jdbc:sqlserver://localhost;databaseName=AdventureWorks2012" ;
         s:username "{{@db.username}}";
         s:password "{{@db.password}}";
         s:schema "Production", "HumanResources", "Person", "Sales", "Purchasing";
         s:normalize [
            s:model [
               s:localNamePrefix "C " ;
               s:localNameSeparator " " ;
               s:match [ s:pattern "(.+)Enlarged" ; s:replace "$1" ] ;
            ];
            s:field [
               s:localNamePrefix "P ";
               s:localNameSeparator " " ;
               s:ignore "rowguid ModifiedDate" ;
```

```
s:match (
      [ s:pattern "(.+)GUID$" ; s:replace "$1" ]
      [ s:pattern "(.+)ID$" ; s:replace "$1" ]
      ) ;
      ] ;
      ] ;
      ] .
    ?rdf a s:RdfGenerator, s:OntologyGenerator ;
      s:as (?s ?p ?o) ;
      s:ontology <a href="http://cambridgesemantics.com/ontologies/AdventureWorks">http://cambridgesemantics.com/ontologies/AdventureWorks</a> ;
    s:base <a href="http://cambridgesemantics.com/AdventureWorks">http://cambridgesemantics.com/AdventureWorks</a> ;
}
```

# Query for Multiple Sources that Generates RDF and Ontology with Resource Templates and Object Properties

This query also includes global normalization rules for normalizing the data across all Data Sources.

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
INSERT {
  GRAPH ${targetGraph} {
     ?s ?p ?o .
 }
}
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
      ?event a s:FileSource ;
         s:model "event" ;
         s:url "/opt/shared-files/csv/events.csv";
         s:key ("EVENT ID") .
      ?listing a s:FileSource ;
         s:model "listing" ;
         s:url " /opt/shared-files/csv/listings.csv" ;
         s:key ("LIST ID") ;
         s:reference [ s:model "event" ; s:using ("EVENT ID") ; s:key ("EVENT ID") ] .
      ?date a s:FileSource ;
         s:model "date" ;
```

```
s:url "/opt/shared-files/csv/event dates.csv";
   s:key ("DATE ID") ;
   s:reference [ s:model "event" ; s:using ("EVENT ID") ; s:key ("EVENT ID") ] .
?venue a s:FileSource ;
   s:model "venue" ;
   s:url " /opt/shared-files/csv/venues.csv" ;
   s:key ("VENUE ID") ;
   s:reference [ s:model "event" ; s:using ("EVENT ID") ; s:key ("EVENT ID") ] .
?sale a s:FileSource ;
   s:model "sale" ;
   s:url " /opt/shared-files/csv/sales.csv" ;
   s:key ("SALE ID") ;
   s:reference [ s:model "event" ; s:using ("EVENT ID") ; s:key ("EVENT ID") ] ;
   s:reference [ s:model "listing" ; s:using ("LIST ID") ; s:key ("LIST ID") ] .
?rdf a s:RdfGenerator, s:OntologyGenerator ;
   s:as (?s ?p ?o) ;
   s:ontology <http://cambridgesemantics.com/tickets> ;
   s:base <http://cambridgesemantics.com/data> ;
   s:normalize [
     s:all [
        s:casing s:UPPER ;
        s:localNameSeparator " " ;
     ];
   ].
```

#### Reading a Data Source's Metadata

}

If you want to retrieve instance data from a source but are unsure about the data model, schema, or the exact names of columns and their data types, you can use the Graph Data Interface (GDI) to explore the source's metadata. The GDI can be used to return a list of the catalogs (schemas), models, columns, data types, and other data source information. This topic describes the metadata query syntax and provides several example queries.

- Metadata Query Syntax
- Metadata Query Examples

#### Metadata Query Syntax

The following query syntax shows the structure of a metadata query. The clauses, patterns, and placeholders in blue are described below.

```
# PREFIX Clause
PREFIX s:
               <http://cambridgesemantics.com/ontologies/DataToolkit#>
               <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdf:
              <http://www.w3.org/2000/01/rdf-schema#>
PREFIX rdfs:
               <http://www.w3.org/2001/XMLSchema#>
PREFIX xsd:
              <http://www.w3.org/2002/07/owl#>
PREFIX owl:
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl:
              <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
              <http://purl.org/dc/elements/1.1/>
PREFIX dc:
# Result Clause
SELECT *
WHERE
{
      # SERVICE Clause: Include the following service call
       SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
        [] s:select ?metadata .
        ?data a s:source type ;
            s:url "string" ;
            [ s:username "string" ; ]
            [ s:password "string" ; ]
        ?metadata a s:MetadataSource ;
            s:from ?data ;
      # The metadata selector below specifies the type of metadata to return.
        ?catalogs | ?fields | ?models [
            ?metadata type datatype ;
            ...;
       ].
    }
}
```

Option	Туре	Description
PREFIX Clause	N/A	The PREFIX clause declares the standard and custom prefixes for GDI service queries. Generally, queries include the prefixes from the query template (or a subset of them) plus any data-specific declarations.
Result Clause	N/A	The result clause defines the results to return. For metadata queries, the result clause is typically $SELECT *$ .
SERVICE Clause		Include the required GDI SERVICE call in the WHERE clause. The rest of the WHERE clause defines the patterns to look for in the source.
[] s:select ?metadata	N/A	Include this required triple pattern in metadata queries. The select property specifies the source that should be used to return data.
source_type	object	The ?data a s:source_type triple pattern specifies the type of data source that the query will run against. For example, ?data a s:DbSource, specifies that the source type is a database. The list below describes the available types:
		DbSource to connect to any type of database.
		• FileSource for flat files. The supported file types are CSV and TSV, JSON, NDJSON, XML, Parquet, and SAS (SAS Transport XPT and SAS7BDAT formats). The GDI automatically determines the file type from the file extensions. When querying file sources, make sure that the files are accessible to both Anzo and AnzoGraph.
		HttpSource to connect to HTTP endpoints.
		ElasticSource to connect to Elasticsearch indexes on an Elasticsearch server.

Option	Туре	Description
		<ul> <li>KafkaSource to connect to Kafka streaming sources.</li> <li>Tip         Certain connection and input parameters are available         based on the specified source type. For details about the         options for your source, see GDI Property Reference.     </li> </ul>
url	string	This property specifies the URL for the data source, such as the database URL, Elasticsearch URL, or HTTP endpoint URL. For file-based sources, the url property specifies the file system location of the source file or directory of files.
		<pre>Important For security, it is a best practice to reference connection information (such as the url, username, and password) from a Query Context so that the sensitive details are abstracted from any requests. In addition, using a Query Context makes connection details reusable across queries. See Using Query Contexts for more information. For example, the triple patterns below reference keys from a Query Context: ?data a s:DbSource ; s:url "{{@db.eca493ff9a.url}}"; s:username "{{@db.eca493ff9a.user}}";</pre>
username	string	<pre>s:password " {{@db.eca493ff9a.password}}"; If authentication is required to access the source, include this</pre>
		property to specify the user name.

Option	Туре	Description
password	string	This property lists the password for the given username.
catalogs	variable	This selector narrows the results to schema-related metadata such as the schema names. Even when additional metadata types (metadata_type datatype) are specified as objects, only catalog (schema) information is returned.
fields	variable	This selector is the broadest and most flexible option. Using the fields selector enables users to return any and all of the source metadata information, depending on the specified metadata types (metadata_type datatype).
models	variable	This selector narrows the results to model-related metadata such as the model names. Even when additional metadata types (metadata_type datatype) are specified as objects, only model information is returned.
metadata_ type datatype	N/A	The triple patterns in the array for the metadata selector specify the type of metadata to return as well as the data type for the return value. The following list shows all of the valid options. You can include any combination of properties. The results that are returned depend on the type of data source and whether the information exists in the source. The parentheses around the data type are not required but are included in this document for readability.
		<ul> <li>format. For file sources, this property returns file names.</li> <li>?field (xsd:string): Returns column names</li> </ul>
		<ul> <li>?catalog (xsd:string): Returns schema names.</li> </ul>
		Adatatype (owl:Thing): Returns the data types of the columns.

Option	Туре	Description
		<ul> <li>?keys (xsd:string): Returns primary and foreign key columns. For compound keys, the GDI returns a commaseparated list of columns comprising the key.</li> <li>?format (xsd:string): Returns the format of the source.</li> <li>?cardinality (xsd:string): Returns the cardinality of relationships between tables: optional, many, or required.</li> <li>?count (xsd:int): Returns the number of times the field appears in the source.</li> <li>?order (xsd:int): Returns the order in which the field was encountered.</li> </ul>

#### **Metadata Query Examples**

This section includes sample metadata queries that run against different types of data sources.

- List Database Schemas
- Explore a Database Schema
- Explore a Directory of SAS Files
- Explore an HTTP Endpoint
- Explore a Directory of CSV Files

### **List Database Schemas**

The query below sends a metadata query to a MySQL database to return a list of the schemas that are available:

PREFIX	s:	<http: cambridgesemantics.com="" datatoolkit#="" ontologies=""></http:>
PREFIX	rdf:	<http: 02="" 1999="" 22-rdf-syntax-ns#="" www.w3.org=""></http:>
PREFIX	rdfs:	<http: 01="" 2000="" rdf-schema#="" www.w3.org=""></http:>
PREFIX	xsd:	<http: 2001="" www.w3.org="" xmlschema#=""></http:>
PREFIX	owl:	<http: 07="" 2002="" owl#="" www.w3.org=""></http:>

```
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
PREFIX dc: <a href="http://purl.org/dc/elements/1.1/">http://purl.org/dc/elements/1.1/></a>
SELECT *
WHERE
{
       SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
           [] s:select ?metadata .
           ?data a s:DbSource ;
                 s:url "jdbc:mysql://10.100.2.9:5555/?user=root&password=Mysql1@#" .
           ?metadata a s:MetadataSource ;
                 s:from ?data ;
           ?catalogs [
                 ?catalog (xsd:string) ;
                 ?order (xsd:int) ;
           ] .
   }
}
ORDER BY ?catalog
```

#### The query returns the following results:

catalog	Ι	order
	-+-	
BANKTEST_DB	Ι	1
EMR	Ι	4
GOLFCLUB_DB	Ι	8
NORTHWIND	Ι	10
SPORTDB	Ι	13
SQLPOCKET_DB	Ι	14
WORDPRESS_DB	Ι	16
classicmodels	Ι	2
crm_national_patients	Ι	3
emrdbbig	Ι	5
emrdbsmall	Ι	6
emrnational_schema	Ι	7
mysql	Ι	9
optum	Ι	11
performance_schema	Ι	12

## **Explore a Database Schema**

Using the list of schemas that were returned in the example above (List Database Schemas), the query below returns metadata about the columns in one of the schemas. To narrow the results to a schema, the schema name (NORTHWIND) is added to the connection URL.

```
PREFIX s:
              <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX rdf:
              <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
PREFIX dc:
             <http://purl.org/dc/elements/1.1/>
SELECT *
WHERE
{
       SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
       [] s:select ?metadata .
       ?data a s:DbSource ;
               s:url "jdbc:mysql://10.100.2.9:5555/NORTHWIND?user=root&password=Mysql1@#" .
       ?metadata a s:MetadataSource ;
               s:from ?data ;
       ?fields [
               ?model (xsd:string) ;
               ?field (xsd:string) ;
               ?datatype (owl:Thing) ;
       ].
   }
}
ORDER BY ?model
```

The query returns the following results:

model

| data

\_\_\_\_

Alphabetical list of products	CategoryID
http://www.w3.org/2001/XMLSchema#	int
Alphabetical list of products	Discontinued
http://www.w3.org/2001/XMLSchema#	boolean
Alphabetical list of products	SupplierID
http://www.w3.org/2001/XMLSchema#	int
Alphabetical list of products	UnitPrice
http://www.w3.org/2001/XMLSchema#	decimal
Alphabetical list of products	ProductName
http://www.w3.org/2001/XMLSchema#	string
Alphabetical list of products	QuantityPerUnit
http://www.w3.org/2001/XMLSchema#	string
Alphabetical list of products	UnitsOnOrder
http://www.w3.org/2001/XMLSchema#	short
Alphabetical list of products	CategoryName
http://www.w3.org/2001/XMLSchema#	string
Alphabetical list of products	ProductID
http://www.w3.org/2001/XMLSchema#	int
Alphabetical list of products	ReorderLevel
http://www.w3.org/2001/XMLSchema#	short
Alphabetical list of products	UnitsInStock
http://www.w3.org/2001/XMLSchema#	short
Categories	CategoryID
http://www.w3.org/2001/XMLSchema#	int
Categories	Description
http://www.w3.org/2001/XMLSchema#	string
Categories	Picture
http://www.w3.org/2001/XMLSchema#	base64Binary
Categories	CategoryName
http://www.w3.org/2001/XMLSchema#	string
Categories	categoryid
Category Sales for 1997	CategoryName
http://www.w3.org/2001/XMLSchema#	string
Category Sales for 1997	CategorySales
http://www.w3.org/2001/XMLSchema#	double
Current Product List	ProductName
http://www.w3.org/2001/XMLSchema#	string
Current Product List	ProductID
http://www.w3.org/2001/XMLSchema#	int

### **Explore a Directory of SAS Files**

The query below explores a directory of SAS files to return the model, catalog (schema), field, data type, and cardinality information. The query also orders the results by model name, which is the file name for file sources of a data model does not exist.

```
<http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX s:
             <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdf:
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
PREFIX dc:
             <http://purl.org/dc/elements/1.1/>
SELECT *
WHERE
{
       SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
       [] s:select ?metadata .
       ?data a s:FileSource ;
               s:url "/opt/shared-files/sas" .
       ?metadata a s:MetadataSource ;
               s:from ?data ;
       ?fields [
               ?model (xsd:string) ;
               ?field (xsd:string) ;
               ?catalog (xsd:string) ;
               ?datatype (owl:Thing) ;
               ?cardinality (xsd:string) ;
       ].
   }
}
```

ORDER BY ?model

The query returns the following results:

model		field	I	catalog		datatype		cardinality
_	•+•		-+-		+-		+-	
demand		P1		les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
demand		P2		les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
demand		P3		les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
demand		Y		les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
demand		Q1	Ι	les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
demand		Q2		les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
demand		Q3		les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
demo		YEAR		les/sas		http://www.w3.org/2001/XMLSchema#long		REQUIRED
demo		QTR		les/sas		http://www.w3.org/2001/XMLSchema#long		REQUIRED
demo		GDP		les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
demo		PR		les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
demo		M1	Ι	les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
demo		RS	Ι	les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
airline		YEAR		les/sas		http://www.w3.org/2001/XMLSchema#long		REQUIRED
airline		Y		les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
airline		W		les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
airline		R	Ι	les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
airline		L	Ι	les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
airline		K		les/sas		http://www.w3.org/2001/XMLSchema#double		REQUIRED
cars		MPG		les/sas		http://www.w3.org/2001/XMLSchema#long		REQUIRED
cars		CYL		les/sas		http://www.w3.org/2001/XMLSchema#long		REQUIRED
50 rows								

# **Explore an HTTP Endpoint**

The query below explores the metadata for a sample HTTP source that compiles worldwide weather statistics.

PREFIX	s:	<http: cambridgesemantics.com="" datatoolkit#="" ontologies=""></http:>
PREFIX	rdf:	<http: 02="" 1999="" 22-rdf-syntax-ns#="" www.w3.org=""></http:>
PREFIX	rdfs:	<http: 01="" 2000="" rdf-schema#="" www.w3.org=""></http:>
PREFIX	xsd:	<http: 2001="" www.w3.org="" xmlschema#=""></http:>
PREFIX	owl:	<http: 07="" 2002="" owl#="" www.w3.org=""></http:>
PREFIX	anzo:	<http: 07="" 2008="" anzo#="" ontologies="" openanzo.org=""></http:>
PREFIX	zowl:	<http: 05="" 2009="" anzoowl#="" ontologies="" openanzo.org=""></http:>
PREFIX	dc:	<http: 1.1="" dc="" elements="" purl.org=""></http:>
SELECT	*	

```
WHERE
```

```
{
       SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
       [] s:select ?metadata .
       ?data a s:HttpSource ;
               s:url "https://sampleEndpoint.com/forecast/30.374563,-97.975892" .
       ?metadata a s:MetadataSource ;
              s:from ?data ;
       ?fields [
               ?model (xsd:string) ;
               ?field (xsd:string) ;
               ?datatype (owl:Thing) ;
               ?cardinality (xsd:string) ;
               ?order (xsd:int) ;
       ].
   }
}
ORDER BY ?model ?order
```

### The query returns the following results:

model   cardinality	field   order	datatype	I
+		-+	+
currently	time	http://www.w3.org/2001/XMLSchema#int	I
REOUIRED			1
currently	summary	http://www.w3.org/2001/XMLSchema#string	
REQUIRED	7		
currently	icon	http://www.w3.org/2001/XMLSchema#string	
REQUIRED	8		
currently	nearestStormDistance	http://www.w3.org/2001/XMLSchema#int	
REQUIRED	9		
currently	nearestStormBearing	http://www.w3.org/2001/XMLSchema#int	
REQUIRED	10		
currently	precipIntensity	http://www.w3.org/2001/XMLSchema#int	
REQUIRED	11		
currently	precipProbability	http://www.w3.org/2001/XMLSchema#int	
REQUIRED	12		

currently | temperature http://www.w3.org/2001/XMLSchema#float REQUIRED 13 | http://www.w3.org/2001/XMLSchema#float currently | apparentTemperature REQUIRED | 14 currently | dewPoint http://www.w3.org/2001/XMLSchema#float REOUIRED 15 | http://www.w3.org/2001/XMLSchema#float currently | humidity REQUIRED 1 16 currently | pressure http://www.w3.org/2001/XMLSchema#float REQUIRED 1 17 currently | windSpeed http://www.w3.org/2001/XMLSchema#float REQUIRED | 18 currently | windGust http://www.w3.org/2001/XMLSchema#float REQUIRED 19 currently | windBearing | http://www.w3.org/2001/XMLSchema#int REQUIRED | 20 currently | cloudCover http://www.w3.org/2001/XMLSchema#float REQUIRED | 21 | http://www.w3.org/2001/XMLSchema#int currently | uvIndex REOUIRED 22 currently | visibility http://www.w3.org/2001/XMLSchema#int REQUIRED | 23 currently | ozone http://www.w3.org/2001/XMLSchema#float REQUIRED | 24 | http://www.w3.org/2001/XMLSchema#string | daily summary REOUIRED 75 | icon daily http://www.w3.org/2001/XMLSchema#string | REQUIRED | 76 daily | data I MANY 77 data | time http://www.w3.org/2001/XMLSchema#int 29 REQUIRED data | precipIntensity http://www.w3.org/2001/XMLSchema#float REQUIRED 30 data | precipProbability http://www.w3.org/2001/XMLSchema#float REQUIRED 31 data | summary | http://www.w3.org/2001/XMLSchema#string | OPTIONAL 32 . . . 81 rows

The following query retrieves the model, field, and data type metadata for the United States from the publicly available <u>Data API Covid Tracking Project</u>.

```
PREFIX s:
             <http://cambridgesemantics.com/ontologies/DataToolkit#>
             <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdf:
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
PREFIX dc: <http://purl.org/dc/elements/1.1/>
SELECT *
WHERE
{
      SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
       [] s:select ?metadata .
       ?data a s:HttpSource ;
               s:url "https://covidtracking.com/api/v1/us/current.csv" .
       ?metadata a s:MetadataSource ;
              s:from ?data ;
       ?fields [
              ?model (xsd:string) ;
              ?field (xsd:string) ;
              ?datatype (owl:Thing) ;
       ].
    }
}
```

#### The query returns the following results:

model	field		datatype
+		-+-	
us	date	Ι	http://www.w3.org/2001/XMLSchema#string
us	states	Ι	http://www.w3.org/2001/XMLSchema#string
us	positive	Ι	http://www.w3.org/2001/XMLSchema#string
us	negative		http://www.w3.org/2001/XMLSchema#string
us	pending		http://www.w3.org/2001/XMLSchema#string
us	hospitalizedCurrently	Ι	http://www.w3.org/2001/XMLSchema#string
us	hospitalizedCumulative	Ι	http://www.w3.org/2001/XMLSchema#string
us	inIcuCurrently	Ι	http://www.w3.org/2001/XMLSchema#string
us	inIcuCumulative	I	http://www.w3.org/2001/XMLSchema#string

us		onVentilatorCurrently		http://www.w3.org/2001/XMLSchema#string
us		onVentilatorCumulative		http://www.w3.org/2001/XMLSchema#string
us		recovered		http://www.w3.org/2001/XMLSchema#string
us		dateChecked		http://www.w3.org/2001/XMLSchema#string
us		death		http://www.w3.org/2001/XMLSchema#string
us		hospitalized		http://www.w3.org/2001/XMLSchema#string
us		lastModified		http://www.w3.org/2001/XMLSchema#string
us		total		http://www.w3.org/2001/XMLSchema#string
us		totalTestResults		http://www.w3.org/2001/XMLSchema#string
us		posNeg		http://www.w3.org/2001/XMLSchema#string
us		deathIncrease		http://www.w3.org/2001/XMLSchema#string
us		hospitalizedIncrease		http://www.w3.org/2001/XMLSchema#string
us		negativeIncrease		http://www.w3.org/2001/XMLSchema#string
us		positiveIncrease	I	http://www.w3.org/2001/XMLSchema#string
us		totalTestResultsIncrease	I	http://www.w3.org/2001/XMLSchema#string
us		hash		http://www.w3.org/2001/XMLSchema#string
25	rows			

### **Explore a Directory of CSV Files**

The query below explores a directory of CSV files to return the model, field, and data type. The query also orders the results by model name, which is the file name for file sources of a data model does not exist. In addition, the query includes s:sampling true, which means the GDI will scan the entire file or files before returning results.

```
PREFIX s:
             <http://cambridgesemantics.com/ontologies/DataToolkit#>
             <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdf:
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
PREFIX dc:
             <http://purl.org/dc/elements/1.1/>
SELECT *
WHERE
{
       SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
       [] s:select ?metadata .
       ?data a s:FileSource ;
```

```
s:url "/opt/shared-files/movie-csv" .
       ?metadata a s:MetadataSource ;
              s:from ?data ;
       # Sample the whole file
       s:sampling true ;
       # Sample the first N records #
       # s:sampling 1000 ;
       ?fields [
               ?model (xsd:string) ;
               ?field (xsd:string) ;
               ?datatype (owl:Thing) ;
      ].
  }
ORDER BY ?model
```

#### The query returns the following results:

}

model	field	datatype				
	-+	+				
MovieActors1	MovieID	I				
http://www.w3.org/2001,	/XMLSchema#int					
MovieActors1	MovieTitle	1				
http://www.w3.org/2001,	/XMLSchema#string					
MovieActors1	ActorID	1				
http://www.w3.org/2001,	/XMLSchema#int					
MovieActors1	ActorName	1				
http://www.w3.org/2001,	/XMLSchema#string					
MovieActors2	MovieID	1				
http://www.w3.org/2001,	/XMLSchema#int					
MovieActors2	MovieTitle	1				
http://www.w3.org/2001,	http://www.w3.org/2001/XMLSchema#string					
MovieActors2	ActorID	1				
http://www.w3.org/2001/XMLSchema#int						
MovieActors2	ActorName	1				
http://www.w3.org/2001/XMLSchema#string						
MovieActors2	ActorCategory	I				
http://www.w3.org/2001/XMLSchema#string						

MovieCategory	MovieID	
http://www.w3.org/2001/X	MLSchema#int	
MovieCategory	MovieTitle	
http://www.w3.org/2001/X	MLSchema#string	
MovieCategory	MoveCategoryID	
http://www.w3.org/2001/X	MLSchema#int	
MovieCategory	MovieCategory	
http://www.w3.org/2001/X	MLSchema#string	
MovieCinematographers	MovieID	
http://www.w3.org/2001/X	MLSchema#int	
MovieCinematographers	MovieTitle	
http://www.w3.org/2001/X	MLSchema#string	
MovieCinematographers	MovieCinematographerID	
http://www.w3.org/2001/X	MLSchema#int	
MovieCinematographers	MovieCinematographerName	
http://www.w3.org/2001/X	MLSchema#string	
MovieComposers	MovieID	
http://www.w3.org/2001/X	MLSchema#int	
MovieComposers	MovieTitle	
http://www.w3.org/2001/X	MLSchema#string	
MovieComposers	MovieComposerID	
http://www.w3.org/2001/X	MLSchema#int	
MovieComposers	MovieComposerName	
http://www.w3.org/2001/X	MLSchema#string	
MovieDirectors	MovieID	
http://www.w3.org/2001/X	MLSchema#int	
MovieDirectors	MovieTitle	
http://www.w3.org/2001/X	MLSchema#string	
79 rows		

#### The following example shows a query that returns metadata for an Elasticsearch source.

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
PREFIX dc: <http://purl.org/dc/elements/1.1/>
PREFIX ex: <http://example.org/ontologies/City#>
PREFIX es: <http://elastic.co/search/>
```

```
PREFIX :
             <http://example.org/cities/>
SELECT *
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit>
{
      [] s:select ? fields .
      ?data a es:ElasticSource ;
         es:url "http://localhost:9200/" ;
         es:index "account" ;
         ?account number xsd:long ;
         ?age (xsd:long) ;
         ?balance (xsd:long) ;
         ?address (xsd:string) ;
         ?city (xsd:string) ;
         ?state (xsd:string) ;
         ?email (xsd:string) ;
         ?employer (xsd:string) ;
         ?firstname (xsd:string) ;
         ?lastname (xsd:string) ;
         ?gender (xsd:string) .
      ? fields a s:MetadataSource ;
         s:from ?data ;
         ?fields [
            ?catalog () ;
            ?model () ;
            ?field () ;
            ?cardinality () ;
            ?datatype () ;
            ?type () ;
            ?object () ;
         ].
  }
}
ORDER BY ?catalog ?model ?field
```

For instructions on querying the instance data based on the data source metadata, see Getting Started with GDI Queries.

### **Paginating Requests**

The GDI exposes paging models that enable you to access large amounts of data across a number of smaller requests. Paging is configured by including the **paging** property in a query and configuring a combination of the pagination options described below. The GDI supports keyset-based, page-based, cursor-based, and offset-based pagination. Paging is supported for all data source types.

- Paging Syntax
- Paging Examples

### Paging Syntax

```
s:paging [
  s:key (?variable) ;
  s:page ?variable ;
  s:cursor ?variable ;
  s:offset ?variable ;
  s:size int ;
  s:limit ?variable ;
] ;
```

Option	Туре	Description
key	variable	Include this property if you want to configure keyset-based pagination where a key is specified to act as a delimiter of the page. The $s: key$ value is a variable that is bound to an expression that defines how to delimit the data. It is usually calculated by an aggregate expression and/or filter that can be pushed to the source. The aggregate expression is typically MAX, but MIN can also be used to page through data in reverse order, such as when working with temporal data. See Key-Based Examples below for examples that configure paging using the $s: key$ property.
page	variable	Include this property if you want to configure page-based pagination where the set is divided into pages. The s:page property value is a variable that the GDI can use to track the current page across requests.

Option	Туре	Description			
		See Page-Based Example below for an example that configures paging using the s:page property.			
cursor	variable	Include this property if you want to configure cursor-based pagination. The s:cursor property value is a variable that is bound against the source to capture the "cursor" value. The GDI uses this value as input to the source to deliver the next page of data. See Cursor-Based Example below for an example that configures paging with the s:cursor property.			
offset	variable	Include this property along with the limit property if you want to configure offset-based pagination. The $s:offset$ property value is a variable that the GDI can use to track the current offset across requests. See Offset-Based Example below for an example that configures paging using the $s:offset$ property.			
size	int	This property can be included with any of the paging models to configure the maximum size of each page. For example, s:size 5000 limits the page size to 5,000 rows.			
limit	variable	This property can be included to define the variable that the GDI should use to push the page size back to the source.			

### Paging Examples

- Key-Based Examples
- Page-Based Example
- Cursor-Based Example
- Offset-Based Example

# **Key-Based Examples**

The example SERVICE clause below pages data based on the <code>?LastID</code> key, which is calculated by finding the maximum value of <code>SalesOrderID</code> and binding it to <code>?LastID</code>. A FILTER is used to filter for data where the <code>SalesOrderID</code> is greater than <code>?LastID</code>.

```
SERVICE <http://cambridgesemantics.com/services/DataToolkit>
{
  BIND(MAX(?SalesOrderID) AS ?LastID)
  FILTER(?SalesOrderID > ?LastID)
  ?SalesOrderHeaderEnlarged a s:DbSource ;
    s:url "jdbc:sqlserver://...";
    s:table "Sales.SalesOrderHeaderEnlarged" ;
    s:paging [
     s:key (?LastID) ;
     s:size 5000 ;
   1;
    ?SalesOrderID (xsd:int) ;
    ?RevisionNumber (xsd:int) ;
    ?OrderDate ("OrderDate" xsd:dateTime) ;
    ?DueDate (xsd:dateTime) .
}
```

The SERVICE clause below shows an example where key-based paging is configured to page through temporal data in reverse order. The s:limit property is configured on the s:HttpSource to limit the overall number of results returned across all pages. This query retrieves at most 1000 records (s:limit 1000), 100 rows (s:size 100) at a time.

```
SERVICE <http://cambridgesemantics.com/services/DataToolkit>
{
  BIND(MIN(?Timestamp) AS ?LastTimestamp)
  ?api a s:HttpSource ;
    s:url "http://slack.com/api/messages/latest" ;
    s:parameter [ s:name "before" ; s:value ?LastTimestamp ] ;
    s:parameter [ s:name "limit" ; s:value ?limit ] ;
    s:limit 1000 ;
    s:paging [
        s:key (?LastTimestamp) ;
        s:limit ?limit ;
    }
}
```

```
s:size 100 ;
] ;
?Message (xsd:string) ;
?Author (xsd:string) ;
?Timestamp (xsd:dateTime) .
```

## Page-Based Example

}

The SERVICE clause below shows an example that uses the s:page property to configure pagebased paging where the page size is 100 rows. This query retrieves at most 1000 records (s:limit 1000), 100 rows (s:size 100) at a time.

```
SERVICE <http://cambridgesemantics.com/services/DataToolkit>
{
    ?api a s:HttpSource ;
    s:url "http://slack.com/api/messages" ;
    s:parameter [ s:name "page" ; s:value ?page ] ;
    s:parameter [ s:name "size" ; s:value ?limit ] ;
    s:limit 1000 ;
    s:paging [
        s:page ?page ;
        s:limit ?limit ;
        s:size 100 ;
    ] ;
    ?Message (xsd:string) ;
    ?Author (xsd:string) ;
    ?Timestamp (xsd:dateTime) .
}
```

### **Cursor-Based Example**

The SERVICE clause below shows an example that uses the s:cursor property to configure cursor-based paging.

```
SERVICE <http://cambridgesemantics.com/services/DataToolkit>
{
     ?api a s:HttpSource ;
     s:url "http://slack.com/api/messages" ;
     s:parameter [ s:name "cursor" ; s:value ?cursor ] ;
     s:parameter [ s:name "limit" ; s:value ?limit ] ;
```

```
s:limit 1000 ;
s:paging [
s:cursor ?cursor ;
s:limit ?limit ;
s:size 100 ;
] ;
?Message (xsd:string) ;
?Author (xsd:string) ;
?Timestamp (xsd:dateTime) ;
?cursor ("next_cursor" xsd:string) .
}
```

### **Offset-Based Example**

The SERVICE clause below shows an example that uses the s:offset property to configure offset-based paging.

```
SERVICE <http://cambridgesemantics.com/services/DataToolkit>
{
  ?api a s:HttpSource ;
   s:url "http://slack.com/api/messages" ;
   s:parameter [ s:name "offset" ; s:value ?offset ] ;
    s:parameter [ s:name "limit" ; s:value ?limit ] ;
   s:limit 1000 ;
   s:paging [
     s:offset ?offset ;
     s:limit ?limit ;
     s:size 100 ;
   ];
    ?Message (xsd:string) ;
    ?Author (xsd:string) ;
    ?Timestamp (xsd:dateTime) .
}
```

#### **Binding and Hierarchy Concepts**

As part of the Graph Data Interface's (GDI) flexibility, there are multiple ways to express binding hierarchies in queries. This topic describes the options for expressing hierarchies.

- Using Binding Trees and Selector Paths
- Unpacking JSON with Bindings and Arrays

Returning Hierarchies as JSON Strings

#### **Using Binding Trees and Selector Paths**

One way to express hierarchies in queries is to use brackets ([]) to group objects into binding trees. For example, the WHERE clause snippet below organizes mapping variable objects into an hourly/data hierarchy by nesting the ?data patterns inside the ?hourly [] tree:

```
WHERE
{
  SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
      ?data a s:HttpSource;
        s:url "https://sampleEndpoint.com/forecast/" ;
        ?latitude (xsd:double) ;
        ?longitude (xsd:double) ;
        ?timezone (xsd:string) ;
        ?hourly
[
          ?data
I
            ?time (xsd:long) ;
            ?summary (xsd:string) ;
            ?rainIntensity ("precipIntensity" xsd:double) ;
            ?rainProbability ("precipProbability" xsd:double) ;
            ?temperature (xsd:double) ;
            ?feelsLike ("apparentTemperature" xsd:double) ;
            ?humidity (xsd:double) ;
            ?pressure (xsd:double) ;
            ?windSpeed (xsd:double) ;
        ];
    ] .
  }
}
```

When constructing object binding trees, if you choose to introduce the hierarchy with a variable name that is not an exact match to the source label, include a **selector** property to list the value from the source. For example, in the WHERE clause snippet below, s:selector is included to select eventHeader in the source as ?event in the query and statLocation as ?location.

WHERE

{

```
SERVICE <http://cambridgesemantics.com/services/DataToolkit>
  {
      ?data a s:FileSource ;
      s:url "/mnt/data/json/part 1.json" ;
      ?event
Γ
         s:selector "eventHeader" ;
           ?eventId (xsd:string) ;
           ?eventName (xsd:string) ;
           ?eventVersion (xsd:string) ;
           ?eventTime (xsd:dateTime) ;
      ];
      ?location
Γ
         s:selector "statLocation" ;
           ?locationId (xsd:string) ;
           ?lineNo (xsd:int) ;
           ?statNo (xsd:int) ;
           ?statId (xsd:int) ;
      1.
 }
}
```

As an alternative to grouping objects in binding trees, the **selector** property also supports using dot notation to specify paths. For example, the WHERE clause snippet below rewrites the first example query to express the same <code>hourly/data</code> hierarchy as a path in the <code>s:selector</code> value:

```
WHERE
{
  SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
      ?data a s:HttpSource;
         s:url "https://sampleEndpoint.com/forecast/" ;
         ?latitude (xsd:double) ;
         ?longitude (xsd:double) ;
         ?timezone (xsd:string) ;
         s:selector: "hourly.data" ;
         ?time (xsd:long) ;
         ?summary (xsd:string) ;
         ?rainIntensity ("precipIntensity" xsd:double) ;
         ?rainProbability ("precipProbability" xsd:double) ;
         ?temperature (xsd:double) ;
         ?feelsLike ("apparentTemperature" xsd:double) ;
```

```
?humidity (xsd:double) ;
?pressure (xsd:double) ;
?windSpeed (xsd:double) .
}
```

You can also include the \$ character to anchor the selector at the root of the file. For example, s:selector "data" captures all data elements anywhere in the file. But s:selector "\$.data" captures only the data elements that are at the root of the hierarchy.

**Unpacking JSON with Bindings and Arrays** 

In addition to object binding trees and selectors, the GDI offers additional syntax for reading or ingesting JSON sources with nested objects and arrays. For example, following the JSON sample file below is a query that captures each value in the arrays:

```
{
    "payload" :
    {
        "IBP_IndEvent_MSR" :
        {
            "unit" : "ms",
            "value" : [ 0, 1 ]
        },
        "IBP_IndEvent_RMF" :
        {
            "unit" : "-",
            "value" : [ 0.012, 1.398, 3.1415 ]
        }
    }
}
```

To read the JSON file above, the following query uses an object binding (?values []) to drill down to the value arrays in the source. An @ selector is specified in the ?value variable binding (?value ("@" xsd:double)) to retrieve each of the array values. For an array of primitive values, the @ selector captures each value in the array. If the source value was an array of objects, the @ selector would retrieve a JSON representation for each object in the array. In addition to creating a new binding context for the primitive array values, the ?values object binding also includes ?index ("!array::index") to capture the index array with the primitive value.

```
PREFIX s: <a href="http://cambridgesemantics.com/ontologies/DataToolkit#">http://cambridgesemantics.com/ontologies/DataToolkit#</a>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
SELECT *
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
       ?data a s:FileSource ;
       s:url "/mnt/data/json/array-index.json" ;
       s:selector "payload.*" ;
       ?unit (xsd:string) ;
       ?values [
          s:selector "value" ;
          ?value ("@" xsd:double) ;
          ?index ("!array::index") ;
       ] .
  }
}
```

The results of the query are shown below:

unit		value		index
	+ -		-+-	
ms		0		0
ms		1		1
-		0.012		0
-		1.398		1
-		3.1415		2

If you do not want to retrieve all of the values in an array, you can include the specific index number to retrieve instead of using the @ symbol. In the variable binding, the index number is appended in brackets ([]) to the binding column name. For example, the following variable binding retrieves the second index value (the third value in the array) from a "projects" array: <code>?project ("projects [2]")</code>. The next example uses the following JSON file:

```
{
   "field1" : "value1",
   "arrayfield" : [
       "arrayvalue1",
       "arrayvalue2"
  ]
}
```

To retrieve only the second value in the array, the following query appends the index value 1 to the

array column name, arrayfield:

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
SELECT *
WHERE {
    SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
        ?json a s:FileSource ;
        s:url "/mnt/data/json/array-index-2.json" ;
        ?field1 (xsd:string) ;
        ?arrayval ("arrayfield[1]" xsd:string) .
    }
}
```

The results of the query are shown below:

### **Returning Hierarchies as JSON Strings**

When working with schema-less sources, you can also capture a tree of data as a JSON string. For example, the query snippet below targets an HTTP endpoint. In this case, the properties under the **hourly** class of data are unknown. So the query binds all of the data below hourly to the **?hourly** variable by including empty parentheses. As a result, the GDI returns a JSON string representation of all of the properties and instance data under hourly:

```
WHERE
{
   SERVICE <http://cambridgesemantics.com/services/DataToolkit>
   {
        ?data a s:HttpSource;
        s:url "https://sampleEndpoint.com/forecast/" ;
        ?latitude (xsd:double) ;
        ?longitude (xsd:double) ;
        ?longitude (xsd:double) ;
        ?timezone (xsd:string) ;
        ?hourly () .
   }
}
```

For example, the results look like this:
```
latitude | longitude | timezone | hourly
```

30.374563 | -97.975892 | America/Chicago | {"summary":"\"Humid and partly cloudy throughout the day.\"","icon":"\"partly-cloudy-day\"","data":[{"time":"1595559600", summary":"\"Clear\"","icon":"\"clear-night\"","precipIntensity":"0", "precipProbability":"0","temperature":"88.39","apparentTemperature":"91.72", "dewPoint":"67.42","humidity":"0.5","pressure":"1011.7","windSpeed":"7.48", "windGust":"16.71","windBearing":"109","cloudCover":"0.06","uvIndex":"0", "visibility":"10","ozone":"285.2"},{"time":"1595563200","summary":"\"Clear\"", "icon":"\"clear-night\"","precipIntensity":"2.0E-4","precipProbability":"0.01", "precipType":"\"rain\"","temperature":"86.69","apparentTemperature":"90.1", "dewPoint":"67.84","humidity":"0.54","pressure":"1012","windSpeed":"7.05", "windGust":"17.56","windBearing":"110","cloudCover":"0.12","uvIndex":"0", "visibility":"10","ozone":"284.9"},...

Similar to the example above, you can write a query that specifically captures some of the properties in a hierarchy and then returns the rest of the properties and their values as a JSON string representation. To do so, use "@" as the binding path. For example:

```
WHERE
{
  SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
      ?data a s:HttpSource;
      s:url
"https://api.darksky.net/forecast/bdbe3f638eb908c9b94919537dad5945/30.374563,-
97.975892";
      ?latitude (xsd:double) ;
      ?longitude (xsd:double) ;
      ?timezone (xsd:string) ;
      ?hourly [
        s:selector "hourly.data" ;
        ?time (xsd:long) ;
        ?summary (xsd:string) ;
        ?hourly data ("@") ;
      ].
   }
}
```

#### Sample results are shown below:

latitude | longitude | timezone | time | summary | hourly\_data

```
30.374563 | -97.975892 | America/Chicago | 1595559600 | Clear
                                                                            {"time":"1595559600","summary":"\"Clear\"",
"icon":"\"clear-
night\"","precipIntensity":"0","precipProbability":"0","temperature":"88.39",
"apparentTemperature":"91.72","dewPoint":"67.42","humidity":"0.5","pressure":"1011.7","
windSpeed":"7.48",
"windGust":"16.71", "windBearing":"109", "cloudCover":"0.06", "uvIndex":"0", "visibility":"
10", "ozone": "285.2"}
 30.374563 | -97.975892 | America/Chicago | 1595563200 | Clear
{"time":"1595563200","summary":"\"Clear\"",
"icon":"\"clear-night\"","precipIntensity":"2.0E-
4", "precipProbability":"0.01", "precipType":"\"rain\"", "temperature":"86.69",
"apparentTemperature":"90.1","dewPoint":"67.84","humidity":"0.54","pressure":"1012","wi
ndSpeed":"7.05", "windGust":"17.56",
"windBearing":"110","cloudCover":"0.12","uvIndex":"0","visibility":"10","ozone":"284.
9"}
 30.374563 | -97.975892 | America/Chicago | 1595566800 | Partly Cloudy
                                                                          {"time":"1595566800","summary":"\"Partly Cloudy\"",
"icon":"\"partly-cloudy-night\"","precipIntensity":"3.0E-4","precipProbability":"0.01",
"precipType":"\"rain"","temperature":"85.63","apparentTemperature":"89.21",
"dewPoint":"68.33", "humidity":"0.56", "pressure":"1012.6", "windSpeed":"6.48", "windGust":
"17.92", "windBearing": "110",
"cloudCover":"0.34","uvIndex":"0","visibility":"10","ozone":"284.5"}
. . .
```

#### **Incremental Onboarding Concepts**

When loading data from a database or file-based data source with a Graph Data Interface (GDI) query, you can add a few statements to the query to load a portion of the data incrementally rather than all of the data at once. As data is added or changed in the source, new data can be ingested without having to reload all of the previously ingested data. Because incremental ingestion is

configured as a filter in a SPARQL query, it is extremely flexible, allowing for various conditions to be defined for diverse data sources. When the data is ingested, the GDI evaluates the current state of the data and then loads only the data that meets the conditions defined in the query.

This topic provides example incremental queries to get you started and includes instructions for configuring a Data Layer as an incremental ingestion workflow.

- Incremental DbSource Example
- Incremental FileSource Example
- Setting Up a Data Layer to Ingest Data Incrementally

## Incremental DbSource Example

The following query from a Direct Load Step ingests data from a database. All of the values for the requested columns in the ORDER\_DETAILS table will be loaded.

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX s: <a href="http://cambridgesemantics.com/ontologies/DataToolkit#">http://cambridgesemantics.com/ontologies/DataToolkit#</a>>
INSERT {
   GRAPH ${targetGraph} {
      ?s ?p ?o .
  }
}
${usingSources}
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
      ?data a s:DbSource ;
          s:url "jdbc:oracle:thin:@10.10.10.10:1111/XE" ;
          s:username "northwind" ;
          s:password "NORTHWIND123" ;
          s:schema "NORTHWIND" ;
          s:table "ORDER DETAILS" ;
          ?database ("!") ;
          ?schema ("!") ;
          ?table ("!") ;
          ?OrderID (xsd:int) ;
```

```
?ProductID (xsd:int) ;
?UnitPrice (xsd:double) ;
?Quantity (xsd:short) ;
?Discount xsd:double .
BIND(IRI("http://cambridgesemantics.com/orders/{{?OrderID}}") AS ?resource)
?rdf a s:RdfGenerator, s:OntologyGenerator ;
s:as (?s ?p ?o) ;
s:ontology <http://cambridgesemantics.com/ontologies/northwind> ;
s:base <http://cambridgesemantics.com/data> .
}
```

The query below adds statements that configure the same Direct Load Step to ingest data incrementally. It captures the maximum order ID as the incremental value. When the source is updated with records that increase the order ID, only the records with larger order IDs than the previous maximum value will be ingested when the Graphmart is refreshed or reloaded. In the query:

- A ?MaxID variable is bound to the result of MAX(?OrderID): BIND (MAX(?OrderID) AS ?MaxID).
- The ?MaxID variable is defined as the incremental value: ?MaxID a s:IncrementalValue.
- A filter clause is added to create a condition that ingests only the records where the order ID is greater than the previously ingested maximum ID: FILTER (?OrderID > ?MaxID).

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
INSERT {
    GRAPH ${targetGraph} {
        ?s ?p ?o .
    }
} ${usingSources}
```

```
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
      ?MaxID a s:IncrementalValue .
      FILTER (?OrderID > ?MaxID)
      BIND (MAX(?OrderID) AS ?MaxID)
      ?data a s:DbSource ;
         s:url "jdbc:oracle:thin:@10.10.10.10:1111/XE" ;
         s:username "northwind" ;
         s:password "NORTHWIND123" ;
         s:schema "NORTHWIND" ;
         s:table "ORDER DETAILS" ;
         ?database ("!") ;
         ?schema ("!") ;
         ?table ("!") ;
         ?OrderID (xsd:int) ;
         ?ProductID (xsd:int) ;
         ?UnitPrice (xsd:double) ;
         ?Quantity (xsd:short) ;
         ?Discount xsd:double .
      BIND(IRI("http://cambridgesemantics.com/orders/{{?OrderID}}") AS ?resource)
      ?rdf a s:RdfGenerator, s:OntologyGenerator ;
      s:as (?s ?p ?o);
      s:ontology <http://cambridgesemantics.com/ontologies/northwind> ;
      s:base <http://cambridgesemantics.com/data> .
  }
}
```

#### Incremental FileSource Example

#### The following query from a Direct Load Step ingests data from all of the CSV files in the

/nfs/data/fmcsa directory:

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
DELETE {
```

```
GRAPH ${targetGraph} {
 }
}
INSERT {
  GRAPH ${targetGraph} {
     ?s ?p ?o .
  }
}
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
      ?data a s:FileSource ;
         s:model "fmcsa" ;
         s:url "/nfs/data/fmcsa" ;
         s:pattern "*.csv" .
      ?rdf a s:RdfGenerator , s:OntologyGenerator ;
         s:as (?s ?p ?o) ;
         s:ontology <http://cambridgesemantics.com/ontologies/fmcsa> ;
         s:base <http://cambridgesemantics.com/data/> .
  }
}
```

The query below adds statements that configure the same Direct Load Step to ingest data incrementally. It uses a "last modified" strategy to determine what files are new or modified and should be ingested the next time the Graphmart is refreshed or reloaded. In the query:

- The modified timestamp metadata on the files is captured with ?Modified ("!").
- The ?LastRun variable is bound to the result of the NOW() function: BIND (NOW() AS ?LastRun).
- A filter clause is added to check whether the modified timestamp is later than the timestamp from the last time the query was run: FILTER (?Modified > ?LastRun).
- **?LastRun is defined as the incremental value: ?LastRun a s:IncrementalValue.**

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
```

```
DELETE {
  GRAPH ${targetGraph} {
 }
}
INSERT {
  GRAPH ${targetGraph} {
     ?s ?p ?o .
  }
}
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
      ?LastRun a s:IncrementalValue .
      FILTER (?Modified > ?LastRun)
      BIND (NOW() AS ?LastRun)
      ?data a s:FileSource ;
         s:model "fmcsa" ;
         s:url "/nfs/data/fmcsa" ;
         s:pattern "*.csv" ;
         ?Modified ("!") .
      ?rdf a s:RdfGenerator , s:OntologyGenerator ;
         s:as (?s ?p ?o) ;
         s:ontology <http://cambridgesemantics.com/ontologies/fmcsa> ;
         s:base <http://cambridgesemantics.com/data/> .
  }
}
```

## Setting Up a Data Layer to Ingest Data Incrementally

- Create a new empty dataset in the Anzo Data Store. For instructions, see Adding an Empty Dataset for an Export Step.
- 2. In the graphmart where you want to add a GDI query that ingests data incrementally, add a new layer.
- 3. In the new layer, add a Load Dataset Step as the first step. The Linked Dataset for this step should be the empty dataset that you created in the first step.
- 4. Now, add a Direct Load Step as the next step in the layer. Edit the query template in the step to compose the incremental query.

5. As the last step in the layer, add an Export Step. The Target FLDS for the step should also be the empty dataset that you created in the first step. For example, the image below shows a graphmart with a layer that is set up to ingest data incrementally.



6. Activate the graphmart to ingest the data and export it to an FLDS. Once the graphmart is activated, enable the Load Dataset Step.

# **Options for Data Types, Data Linking, and Models**

The topics in this section describe the options that are available across data source types for controlling the way that strings are coerced to other data types, the relationships that define connections across multiple sources, and the way label and URI values are generated in the data model.

- Data Type Formatting Options
- Data Linking Options
- Model Normalization Options

#### **Data Type Formatting Options**

To give you control over the data types that are used when coercing strings to other types, the **formats** property can be included in GDI queries to define the desired types. In addition, formats can be used to describe the formats of date and time values in the source to ensure that they are recognized and parsed to the appropriate date, time, and/or dateTime values. You can also use the formats property to suppress the conversion so that the generated values are typed the same way as the source.

## Tip

The GDI takes locale into account when formatting the generated date and time values.

For sources that do not include data type specifications and natively treat values as strings, the GDI Generator automatically converts the values to the appropriate type. For example, if a CSV file includes the value "Feb-18-2022," the GDI parses the string to an xsd:date with the format "2022-02-18". A column with numbers is converted to an xsd:int type and a column with a decimal value is converted to xsd:float. The formats property usage is described below.

- Formats Syntax
- Formats Examples

## Formats Syntax

```
s:formats [
    s:strict boolean ; [
        xsd:data_type "format"
        | xsd:data_type boolean ;
        [ ... ; ]
    ]
];
```

Option	Туре	Description
strict	boolean	This property enables or disables the automatic type conversion feature. By default, strict is set to false (s:strict false), meaning the GDI's automatic type conversion feature is enabled. When strict is false, any formats specified in s:formats [] augment the GDI's built-in date and time formats. You can selectively disable certain type conversions, however, by including xsd:data_type false. For example, xsd:dateTime false disables the parsing of strings to dateTime. When strict is true (s:strict true), the auto conversion logic is essentially disabled and the generated data will be represented the same way it is in the source. When strict is true, you can selectively enable certain conversions by including xsd:data_ type true or by defining xsd:data_type "format". In this case, values that do not match any of the formats provided will be
		typed as xsd:string.
xsd:data_ type "format"	N/A	Include xsd:data_type "format" when you want to describe the formats of date and time values in the source. The GDI supports Java date and time format notation. For example, if dates in the source are formatted like "yyyy-MM-dd," include the statement xsd:date "yyyy-MM-dd". If the source uses multiple formats for dates, e.g., 18-MAR-1978 and 03/18/1978, you can list multiple formats for xsd:date, such as xsd:date "dd-MMM-yyyy",

Option	Туре	Description	
		"MM/dd/yyyy". Note The GDI's default base year is 2000. If the source data has years with only two digits, such as 02-04-99, the GDI prepends 20 to the digits. The value 02-04-99 is parsed to 02-04-2099. To specify an alternate base year to use for two-digit values, you can include the notation ^nnnn (e.g., ^1900) in the format value. For example, to set the base year to 1900 instead of 2000, use a format value such as xsd:date "dd-MMM-yy^1900" or xsd:date "dd-MMM-yy^1990". When one of those values is specified, 02-04-99 is parsed to 02-04-1999.	
xsd:data_ type boolean	N/A	When strict is false or not set, you can disable specific type conversions by listing data types and setting their values to false.For example, if you want the GDI to convert strings to integers or floats when possible but you want the dates in the source to be preserved as strings, you can include xsd:date false to disable the conversion of strings to dates. When strict is true, you can enable specific type conversions by listing data types and setting their values to true. For example, if you want the GDI to preserve the strings in the source except for when the string is a number, you can include xsd:int true to enable the conversion of strings to integers.	

## **Formats Examples**

The example below sets strict to true and forces the GDI to parse values only to the data types that are enabled with true. It also defines the format to look for when converting strings to dateTime:

```
s:formats [
   s:strict true ;
   xsd:int true ;
   xsd:dateTime true ;
   xsd:dateTime "yyyy-MM-dd-HH-mm-ss" ;
];
```

The example below does not set strict, so the default value of false is used. The data type definitions specify the formats of the values to parse as date, time, and dateTime values. The example also disables the conversion from string to long:

```
s:formats [
    xsd:date "MM/dd/yyyy", "MMM dd", "MMM dd yyyy";
    xsd:time "HH[:mm][:ss][]a";
    xsd:dateTime "M/d/yyyy HH:mm:ss a", "yyyy-MM-dd-HH-mm-ss";
    xsd:long false;
];
```

#### **Data Linking Options**

When a data source does not define keys (such as a CSV or JSON source), the GDI provides properties that enable you to create a connected knowledge graph by defining relationships, resource templates (primary keys) and object properties (foreign keys), when you are loading data from multiple sources. The properties that are available are described below.

- Data Linking Syntax
- Data Linking Examples

## Data Linking Syntax

```
s:key ("column_name") ;
s:reference [
    s:model "table_to_reference" ;
    s:using ("foreign_key_column")
]
```

Option	Туре	Description
key	string	Include this property when you want to define the primary key column for the source file or table. This column is leveraged in a resource

Option	Туре	Description	
		<pre>template for the instances that are created from the source. For example, s:key ("EMPLOYEE_ID"). If none of the columns contain unique values, you can specify a combination of columns that would create a unique value. For example, s:key ("FlightNumber", "TailNumber").</pre>	
reference	RDF list	<pre>Include this property when you want to specify a foreign key column. The reference property is an RDF list that includes the model property to list the target table and a using property that defines the foreign key column in the source table. s:reference [    s:model "table_to_reference" ;    s:using ("foreign_key_column") ]</pre>	
		<pre>Tip You can also include an optional key property within the s:reference list that defines the key column in the target table and can be used as a way to expose additional metadata that helps inform the GDI how to name the object property. For example: s:reference [ s:model "Employees" ; s:using ("EMPLOYEE_ID") ; s:key ("EMPLOYEE_ID") ]</pre>	

#### **Data Linking Examples**

For example, the query snippet below defines two data sources. The s:model property defines the table/class for each source, and the s:key defines the primary key for each table/class. The s:reference property for the "venue" table defines a foreign key relationship from venue.EVENT\_ID to event.EVENT\_ID.

```
?event a s:FileSource ;
   s:model "event" ;
   s:url "/opt/shared-files/csv/events.csv" ;
   s:key ("EVENT_ID") .
?venue a s:FileSource ;
   s:model "venue" ;
   s:url " /opt/shared-files/csv/venues.csv" ;
   s:key ("VENUE_ID") ;
   s:reference [ s:model "event" ; s:using ("EVENT_ID") ] .
```

The following query for multiple file sources generates RDF and an ontology with resource templates and object properties. The query also includes global normalization rules for normalizing the data across all sources (see Model Normalization Options for information about normalization).

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
INSERT {
  GRAPH ${targetGraph} {
     ?s ?p ?o .
 }
}
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
      ?event a s:FileSource ;
         s:model "event" ;
         s:url "/opt/shared-files/csv/events.csv" ;
         s:key ("EVENT ID") .
      ?listing a s:FileSource ;
         s:model "listing" ;
         s:url " /opt/shared-files/csv/listings.csv" ;
         s:key ("LIST ID") ;
         s:reference [ s:model "event" ; s:using ("EVENT ID") ; s:key ("EVENT ID") ] .
      ?date a s:FileSource ;
         s:model "date" ;
         s:url "/opt/shared-files/csv/event dates.csv" ;
         s:key ("DATE ID") ;
         s:reference [ s:model "event" ; s:using ("EVENT ID") ; s:key ("EVENT ID") ] .
      ?venue a s:FileSource ;
```

```
s:model "venue" ;
   s:url " /opt/shared-files/csv/venues.csv" ;
   s:key ("VENUE ID") ;
   s:reference [ s:model "event" ; s:using ("EVENT ID") ; s:key ("EVENT ID") ] .
?sale a s:FileSource ;
   s:model "sale" ;
   s:url " /opt/shared-files/csv/sales.csv" ;
   s:key ("SALE ID") ;
   s:reference [ s:model "event" ; s:using ("EVENT ID") ; s:key ("EVENT ID") ] ;
   s:reference [ s:model "listing" ; s:using ("LIST ID") ; s:key ("LIST ID") ] .
?rdf a s:RdfGenerator, s:OntologyGenerator ;
   s:as (?s ?p ?o) ;
   s:ontology <http://cambridgesemantics.com/tickets> ;
   s:base <http://cambridgesemantics.com/data> ;
   s:normalize [
     s:all [
         s:casing s:UPPER ;
        s:localNameSeparator " " ;
     ];
   ].
```

## **Model Normalization Options**

To give users control over the labels and URIs that are generated in the data model, the GDI offers several options for normalizing the class and property fields that are created from the specified data source(s). Normalization rules can be specified at the source level to normalize the data from each source independently, or they can be used at the RDF Generator level to apply global rules across all specified data sources.

#### Note

}

Normalization rules are applied only at the model level. The rules to do affect the instance data values that are ingested.

Including the **normalize** parameter is optional. If you include it, you can specify any combination of rules. See Default Normalization Behavior below for details about the Generator's default behavior when normalization rules are not specified in your query.

- Default Normalization Behavior
- Normalize Syntax
- Normalize Examples

#### **Default Normalization Behavior**

The GDI Generator normalizes data according to the following rules by default. If you do not include the s:normalize parameter in your query, these are the rules that are applied:

```
s:normalize [
s:all [
s:removePrefix true ;
s:removePartialPrefix false ;
s:allowWhiteSpace false ;
s:allowPunctuation false ;
s:allowSymbols false ;
s:separator " " ;
s:singularize false ;
s:casing s:UpperCamel ;
s:localNameSeparator "." ;
]
```

## **Normalize Syntax**

```
s:normalize boolean | [
   s:model | s:field | s:all
   [
    s:removeStart "string";
    s:removeEnd "string";
    s:removePrefix boolean;
    s:removePartialPrefix boolean;
    s:match [ s:pattern "java_regex"; s:replace "java_regex"];
    s:disambiguationLevel int;
    s:ignore "string";
    s:words "string";
    s:preserve "string";
```

```
s:split "string";
s:allowWhiteSpace boolean;
s:allowPunctuation boolean;
s:allowSymbols boolean;
s:singularize boolean;
s:casing property;
s:separator "string";
s:localNamePrefix "string";
s:localNameSeparator "string";
];
```

```
Description
Property
                             Type
                             N/A
boolean
                                           Normalize is enabled by default for all GDI Generator
                                           queries. If you want to disable normalization, you can
                                           include s:normalize false. If normalization is
                                           disabled, the names in the source will be used verbatim
                                           both for labeling and in generating the local names for
                                           property and class URIs. However, when normalization
                                           is disabled, the labels in the data source are used
                                           verbatim. In addition, the Generator creates hard-to-
                                           read, URL-encoded local names for property and class
                                           URIs.
s:model | s:field | s:all
                             N/A
                                           This property defines whether the specified
                                           normalization rules should be applied across the model
                                           or only to the classes or properties. The list below
                                           describes each option:
                                              s:model: Applies the rules to the file/table/class
                                             •
                                                names only.
                                              s:fields: Applies the rules to the
                                             •
                                                column/property/field names only.
                                              s:all: (Default) Applies the rules to both the
                                                class and property names. This is the default
```

Property	Туре	Description
		value if not specified.
removeStart	string	If you want to remove text from the beginning of identifiers, include the <b>removeStart</b> rule to specify the string to remove. For example, s:removeStart "temp_".
removeEnd	string	If you want to remove text from the end of identifiers, include the <b>removeEnd</b> rule to specify the string to remove. For example, s:removeEnd "NEW".
removePrefix	boolean	If there are property identifiers that share a prefix with the class, the RDF Generator automatically removes the shared prefix from the property name; the <b>removePrefix</b> rule is set to true by default. For example, if there is an EMPLOYEE class with an EMPLOYEE_ID column, the shared prefix "EMPLOYEE" is removed from the generated property so that it becomes "ID." If you do not want the Generator to remove prefixes, you can include s:removePrefix false.
removePartialPrefix	boolean	If there are property identifiers that share a partial prefix with the class, you can enable <b>removePartialPrefix</b> to remove the partial prefix from the property name. The <b>removePartialPrefix</b> rule is set to false by default. If you want the Generator to remove partial prefixes, you can include s:removePrefix true.
match	RDF list	This rule provides a way to use regular expressions (REGEX) to match a <b>pattern</b> against source identifiers and <b>replace</b> the matched text in the normalized name.

Property	Туре	Description
		<pre>The s:pattern property defines the Java REGEX pattern to match against, and s:replace defines the Java REGEX replacement pattern. As shown in the example below, the match rule can also be configured with an rdf:List of objects to perform match evaluation in a certain order: s:match ( [ s:pattern "(.+)GUID\$" ; s:replace "\$1" ; ] [ s:pattern "(.+)ID\$" ; s:replace "\$1" ; ] )</pre>
disambiguationLevel	int	This rule specifies the number of levels to use to resolve ambiguities between similarly named elements in a hierarchical source. For example, an element named "Data" appears in two contexts: "Currently" and "Hourly." By default, the Generator retains all levels, meaning two classes are generated: "Currently Data" and "Hourly Data." If s:disambiguationLevel is set to 0, a single class named "Data" is generated and both the Currently and Hourly classes have a "Data" property. The disambiguationLevel value is also used to determine the number of hierarchy levels to use when encoding the local name of the generated URI.
ignore	string	This rule can be used to list identifiers that should be ignored. Properties and classes will not be generated for identifiers that match the specified string(s). The ignore rule is a multi-valued property. For simplicity,

Property	Туре	Description
		you can enter a list by separating words with a space, rather than quoting each term and separating them with a comma. For multi-word identifiers, use single quotes. For example, s:ignore "sample example 'test column' old".
words	string	Since many sources do not encode word boundaries very well, the <b>words</b> rule can be used to list the set of words that should be separate identifiers. This rule tells the Generator which words may be encountered. The words rule is a multi-valued property. For simplicity, you can enter a list by separating words with a space, rather than quoting each term and separating them with a comma. For multi-word identifiers, use single quotes. For example: s:words "activity 'patient complaint' medication observation patient signal specialty study" ;
preserve	string	This rule can be used to identify any words whose casing should be preserved in the input identifiers. For example, if casing is set to lower but you want preserve the original upper casing of certain words, you can specify the words to preserve. The preserve rule is a multi-valued property. For simplicity, you can enter a list by separating words with a space, rather than quoting each term and separating them with a comma. For multi-word identifiers, use single quotes. For example: s:preserve "ABC 'Laundry List' TriG". The preserve rule is case-insensitive. You do not have to match the casing of the words to preserve.

Property	Туре	Description
split	string	This rule specifies the string that should be used to split source identifiers into individual terms. If neither <code>split</code> nor words is specified, input identifiers are split on casing changes and character class changes.
allowWhiteSpace	boolean	This rule specifies whether or not white space should be preserved in identifiers after they have been split into individual terms. This rule is set to false by default, meaning white space is not preserved. You can specify s:allowWhiteSpace true to preserve spaces.
allowPunctuation	boolean	This rule specifies whether or not punctuation should be preserved in identifiers after they have been split into individual terms. This rule is set to false by default, meaning punctuation is not preserved. You can specify s:allowPunctuation true to preserve punctuation.
allowSymbols	boolean	This rule specifies whether or not symbols should be preserved in identifiers after they have been split into individual terms. This rule is set to false by default, meaning symbols are not preserved. You can specify s:allowSymbols true to preserve symbols.
singularize	boolean	This rule specifies whether or not to change any plural identifiers to singular. This rule is set to false by default, meaning plural identifiers are preserved. You can specify s:singularize true to change plural terms to the singular version of the term.
casing	object	This rule specifies how the generated labels should be

Property	Туре	Description
		cased. By default, the Generator outputs labels in upper camel case (s:casing s:UpperCamel). To use a different casing, specify any of the following properties:
		<ul> <li>default: This object preserves the casing from the source. Labels will not be converted.</li> </ul>
		<ul> <li>UPPER: This object converts all characters to uppercase. For example, "uppercase" becomes "UPPERCASE."</li> </ul>
		<ul> <li>lower: This object converts all characters to lowercase. For example, "Lower Case" becomes "lower case".</li> </ul>
		• UpperCameI: This is the default casing value and converts labels to upper cameI case, where terms are concatenated and the first letter of each word is capitalized. For example, "upper cameI case" becomes "UpperCameICase."
		• <b>IowerCamel</b> : This object converts labels to lower camel case, where terms are concatenated and the first letter of the first word is lower case. The first letter of subsequent terms is capitalized. For example, "lower camel case" becomes "lowerCamelCase."
separator	string	This rule specifies the character or characters to use to separate terms in the generated label. The default separator is a space (s:separator " ").

Property	Туре	Description
localNamePrefix	string	This rule specifies a string to use as the prefix for local names when generating a URI.
IocalNameSeparator	string	This rule specifies the string to use for separating local names when encoding hierarchies according to the specified disambiguationLevel. By default, localNameSeparator is a period (s:localNameSeparator "."). If localNameSeparator is empty, hierarchical context will not be encoded into the local name of any properties or child classes. The result would be an ontology where only the class or property name is used to determine the local name. For example, a property URI would look like ont:employeeID rather than ont:Employee.employeeID. The result could lead to "conflicts" in the generated ontology, but those "conflicts" may be desired as properties with same name are reused across the generated ontology.

## Tip

You can specify normalization rules at both the source and global level in the same query. If you include multi-valued rules (such as ignore, words, or preserve) at both levels, the Generator combines all values from both instances of the rule. If you specify single value rules at both levels and the values are conflicting, the Generator applies the value at the source level.

#### **Normalize Examples**

The example below uses the normalize property to normalize data at both the model and field level.

```
s:normalize [
    s:model [
        s:localNamePrefix "C_";
```

```
s:localNameSeparator "_";
s:match [ s:pattern "(.+)Enlarged"; s:replace "$1" ];
];
s:field [
s:localNamePrefix "P_";
s:localNameSeparator "_";
s:ignore "rowguid ModifiedDate";
s:match (
[ s:pattern "(.+)GUID$"; s:replace "$1" ]
[ s:pattern "(.+)ID$"; s:replace "$1" ]
);
];
```

# Advanced Usage by Data Source Type

The topics in this section provide more advanced GDI usage information by including descriptions for all of the query options for each type of supported data source.

## Note

Although connections to data sources can be made directly in GDI queries. Creating the connection in Anzo prior to writing queries enables your organization to get an overview of the sources that are in use or available to use. Connecting to a source also gives a sample view of the data to be onboarded (except for JSON and XML sources). And, for sources that require input of sensitive connection and authorization information, connecting the sources to Anzo allows you to later hide that sensitive information (through the use of a Query Context) from any queries that are generated for that source. See Adding Data Sources for instructions on connecting to sources.

## Tip

Rather than manually writing complex queries, you can use the GDI to automatically generate graphs and ontologies by including a few key statements in a relatively simple query. For information, see Onboarding Data with a Direct Load Step.

- Querying a Database Source
- Querying an HTTP Source
- Querying an Elasticsearch Source
- Querying a File Source

## **Querying a Database Source**

This topic provides details about the structure to use when writing GDI queries to read or ingest data from database data sources. It also includes example queries that may be useful as a starting point for writing your own GDI queries.

- Supported Databases
- Query Syntax
- Query Examples

#### Supported Databases

The GDI supports querying any database through a JDBC connection. AnzoGraph installations include JDBC drivers for the following databases:

- Databricks
- H2
- IBM DB2
- Microsoft SQL Server
- MariaDB
- Oracle
- PostgreSQL
- SAP Sybase (jTDS)
- Snowflake

To extend the service to access other databases, additional JDBC drivers can be added to AnzoGraph. For information about acquiring additional JDBC drivers, contact your Cambridge Semantics Customer Success manager. For instructions on deploying drivers, see <u>Deploy Optional</u> Drivers for Accessing Custom Database Sources in the Deployment Guide.

#### **Query Syntax**

The following query syntax shows the structure of a GDI query for database sources. The clauses, patterns, and placeholders that are links are described below.

```
# PREFIX Clause
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
```

```
PREFIX xsd:
             <http://www.w3.org/2001/XMLSchema#>
PREFIX owl:
             <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
            <http://purl.org/dc/elements/1.1/>
PREFIX dc:
# Result Clause
{
  [ GRAPH ${targetGraph} { ]
  triple patterns
 [ } ]
}
[ ${usingSources} ]
WHERE
{
   # SERVICE Clause: Include the following service call when reading or inserting data.
    SERVICE [ TOPDOWN ] <http://cambridgesemantics.com/services/DataToolkit>
   # View SERVICE Clause: Or use the service call below when constructing a view.
    SERVICE <http://cambridgesemantics.com/services/DataToolkitView>(${targetGraph})
    {
      ?data a s:DbSource ;
        s:url "string" ;
        s:username "string" ;
        s:password "string" ;
        [ s:token "string" ; ]
        [ s:driver "string" ; ]
        [ s:property [ s:name "string" ; s:value "string" ] ; ]
        [ s:timeout int ; ]
        [ s:maxConnections int ; ]
        [ s:batching boolean | int ; ]
        [ s:concurrency int | [ list of properties ] ; ]
        [ s:rate int | "string" ; ]
        [ s:partitionBy "string" | ?variable | s:auto ; ]
        [ s:paging [ pagination options ; ]
        [ s:locale "string" ; ]
        [ s:sampling int ; ]
        [ s:selector "string" | [ list ] ; ]
        [ s:model "string" ; ]
        [ s:key ("string") ; ]
        [ s:reference [ s:model "string" ; s:using ("string") ]
```

```
[ s:formats [ datatype_formatting_options ] ; ]
        [ s:normalize boolean | [ normalization rules ] ; ]
        [ s:query "string" ; ]
        [ s:database "string" ; ]
        [ s:schema "string" ; ]
        [ s:table "string" ; ]
        [ s:count ?variable ; ]
        [ s:offset int ; ]
        [ s:orderBy "string" | ?variable ; ]
        [ s:limit int ; ]
        # Mapping variables
       ?mapping variable ( [ "binding" ] [ datatype ] [ "datetime format" ] ) ;
        ....;
    # Additional clauses such as BIND, VALUES, FILTER
  }
}
```

## Note

For readability, the parameters below exclude the base URI

<http://cambridgesemantics.com/ontologies/DataToolkit#> as well as the s: prefix. As shown in the examples, however, the s: prefix or full property URI does need to be included in queries.

Option	Туре	Description
PREFIX Clause	N/A	The PREFIX clause declares the standard and custom prefixes for GDI service queries. Generally, queries include the prefixes from the query template (or a subset of them) plus any data-specific declarations.
Result Clause	N/A	The result clause defines the type of SPARQL query to run and the set of results to return, i.e., whether you want to read (SELECT or CONSTRUCT) from the source or ingest the data into Anzo (INSERT).
GRAPH	N/A	Include the GRAPH keyword and target graph parameter

Option	Туре	Description
\${targetGraph}		<pre>\$ {targetGraph} when writing an INSERT query to ingest data into a graphmart. Anzo automatically populates the query with the appropriate target URIs when the query runs.</pre>
\${usingSource s}	N/A	Include the source graph parameter $\{usingSources\}\$ when writing a "topdown" query that passes values from the data that is in the graphmart to the data source. Anzo automatically populates the query with the appropriate FROM clauses when the query runs. When passing literal values to the remote source, you do not need to include the source graph parameter. The SERVICE Clause description below includes more information about passing input to data sources.
SERVICE N/A Clause		<pre>Include the SERVICE call SERVICE [ TOPDOWN ] <http: cambridgesemantics.com="" datatool="" kit="" services=""> to invoke the GDI service when you are running a SELECT, INSERT, or CONSTRUCT query that is not creating a view. When writing a CONSTRUCT query in a View Step, use the DataToolkitView service call, as described below in View SERVICE Clause.</http:></pre>
		Include the optional TOPDOWN keyword when you want to pass input values from the graphmart to the data source. When you include TOPDOWN in the service call, it indicates that the rest of the query produces values to send to the source. In this case, the GDI makes repeated calls to pass in each of the specified values and retrieve the data that is based on those values.
View SERVICE Clause	N/A	When writing a CONSTRUCT query that creates a view of the data (usually in a View Step), include the following SERVICE call: SERVICE

Option	Туре	Description
		<http: cambridgesemantics.com="" datatool<br="" services="">kitView&gt;(\${targetGraph}). Using the DataToolkitView call optimizes query execution because it tells the GDI to inspect the query and determine which filters to push to the data source. It also limits the result set and retrieves only the data that is needed, i.e., the source data is fully mapped but all of the mapped data is not necessarily returned.</http:>
url	string	<pre>This property specifies the URL to use to access the database. Important For security, it is a best practice to reference connection information (such as the url, username, and password) from a Query Context so that the sensitive details are abstracted from any requests. In addition, using a Query Context makes connection details reusable across queries. See Using Query Contexts for more information. For example, the triple patterns below reference a Query Context and add a JDBC driver level connection property: ?data a s:DbSource ; s:url "{{@db.eca4bfff9a.url}}" ; s:password " {{@db.eca4bfff9a.password}}" ; s:property [ s:name "access" ; s:value "all" ] ;</pre>
username	string	This property lists the user name to use for the connection to the database.

Option	Туре	Description
		<pre>Fip If you want to group the username and password properties, you can wrap them with s:credentials [    </pre>
password	string	This property lists the password for the given username.
token	string	For connections that require a bearer token, this property can be included to specify the token.
driver	string	This property can be included to specify the JDBC driver to use.
property	RDF list	<pre>This property can be included to list any JDBC driver-specific connection properties. To incorporate property, use the following syntax: s:property [ s:name "custom_driver_property_name" ; s:value "custom_value" ]</pre>
timeout	int	This property can be used to specify the timeout (in milliseconds) to use for requests against the source. For example, s:timeout 5000 configures a 5 second timeout.
maxConnection	int	This property can be used to set a limit on the maximum number

Option	Туре	Description
S		of active connections to the source. For example, s:maxConnections 16 sets the limit to 16 connections. The default value is 10.
batching	boolean or int	This property can be used to disable batching, or it can be used to change the default the batch size. By default, batching is set to 5000 (s:batching 5000). To disable batching, you can include s:batching false in the query. Typically users do not change the batching size. However, it can be useful to control the batch size when performing updates. To configure the size, include s:batching int in the query. For example, s:batching 3000.
concurrency	int or RDF list	<pre>This property can be included to configure the maximum level of concurrency for the query. The value can be an integer, such as s:concurrency 8. If the value is an integer, it configures a maximum limit on the number of slices that can execute the query. For finer-grained control over the number of nodes and slices to use, concurrency can also be included as an object with limit, nodes, and/or executorsPerNode properties. For example, the following object configures a concurrency model that allows a maximum of 24 executors distributed across 4 nodes with 8 executors per node: s:concurrency [ s:limit 24 ; s:nodes 4 ; s:executorsPerNode 8 ; ] ;</pre>
rate	int or string	This property can be included to control the frequency with which a request is sent to the source. The limit applies to the number of

Option	Туре	Description
		requests a single slice can make. If you specify an integer for the rate, then the value is treated as the maximum number of requests to issue per minute. If you specify a string, you have more flexibility in configuring the rate. The sample values below show the types of values that are supported:
		<pre>s:rate "90/minute" ; s:rate "90 per minute" ; s:rate "200000 every week" ; s:rate "10000 every 6 hours" ;</pre>
		To enforce the rate limit, the GDI introduces a sleep between requests that is equal to the rate delay. The more executing slices, the longer the rate delay needs to be to enforce the limit in aggregate.
		Given the example of s:rate "90/minute", the GDI would optimize the concurrency and only use 1 slice for execution with a rate delay of 666ms between requests. If s:rate "240/minute", the GDI would use 3 executors with a rate delay of 750ms between requests.
partitionBy	string, variable, object	The GDI attempts to partition queries automatically across the available cores (slices) in AnzoGraph. To determine how to partition the query, the GDI uses metadata from the source database. It looks for any column in an index, preferring the primary key column if it is interpolable. However, it only considers the first column in any index on the table. After determining the partition column, the GDI does a MIN/MAX on the column as well as a basic sizing query. To specify which column or columns the GDI should partition on, you can include the partitionBy property in the query. The property supports a list of source field names, bound variables, or the object s:auto, which forces the

Option	Туре	Description
		GDI to partition the data when the source does not define partitioning metadata.
paging	RDF list	This property can be used to configure paging so that the GDI can access large amounts of data across a number of smaller requests. For details about the paging property, see Paginating Requests.
locale	string	This property can be used to specify the locale to use when parsing locale-dependent data such as numbers, dates, and times.
sampling	int	This property can be used to configure the number of records in the source to examine for data type inferencing.
selector	string or RDF list	This property can be used as a binding component to identify the path to the source objects. For example, s:selector "Sales.SalesOrderHeader" targets the SalesOrderHeader table in the Sales schema. As an alternative to including the selector property for identifying the target data, you could use the database, schema, and/or table properties.
model	string	This property defines the class (or table) name for the type of data that is generated from the specified data source. For example, s:model "employees". Model is optional when querying a single source. If your query targets multiple sources, however, and you want to define resource templates (primary keys) and object properties (foreign keys), you must specify the model value for each source.
key	string	This property can be used to define the primary key column for the source file or table. This column is leveraged in a resource

Option	Туре	Description
		template for the instances that are created from the source. For example, s:key ("EMPLOYEE_ID"). For more information about key, see Data Linking Options.
reference	RDF list	This property can be used to specify a foreign key column. The reference property is an RDF list that includes the model property to list the target table and a using property that defines the foreign key column. For more information about reference, see Data Linking Options.
formats	RDF list	To give users control over the data types that are used when coercing strings to other types, this property can be included in GDI queries to define the desired types. In addition, it can be used to describe the formats of date and time values in the source to ensure that they are recognized and parsed to the appropriate date, time, and/or dateTime values. For details about the formats property, see Data Type Formatting Options.
normalize	RDF list	To give users control over the labels and URIs that are generated, the GDI offers several options for normalizing the model and/or the fields that are created from the specified data source(s). For details about the normalize property, see Model Normalization Options.
query	string	If you want to access the source data by running an SQL query, you can include this property to specify the query string to run. The language does not have to be SQL if the source supports another language. However, some GDI features where the query is dynamically altered may not work with a non-SQL language. Including {{?variable}} substitutions is supported within s:query strings.

Option	Туре	Description
		<b>Important</b> If you include s:query, you must also specify table and partitionBy. Specify the table name in s:table and the column to partition the table on in s:partitionBy. If the table and partition column are not specified, the GDI will not partition the query and query execution may fail or perform very poorly.
database	string	This property can be used to specify the database to target in the source if the database is not listed in the s:url or s:selector strings.
schema	string	This property can be included to specify the target schema to query. If you include s:schema "schema_name" without specifying s:table (described below) or s:query, all tables in the schema are queried.
table	string	This property can be included to specify the target table or tables for the query.
count	variable	If you want to turn the query into a COUNT query, you can include this property with a <code>?variable</code> to perform a count. For example, <code>s:count</code> ?count.
offset	int	This property can be used to offset the data that is returned by a number of rows.
orderBy	string, variable, list	You can include this property to order the result set by a field name, a bound variable, or a list of names or bound variables.
Option	Туре	Description
----------------------	----------	--
limit	int	You can include this property to limit the number of results that are returned. s:limit maps to the SPARQL LIMIT clause.
mapping_ variable	variable	The mapping variables, in ?variable (["binding"] [datatype] ["datetime_format"]) format, define the triple patterns to output. When the specified ?variable matches the source column name, the GDI uses the variable as the source data selector. If you specify an alternate variable name, a binding needs to be specified to map the new variable to the source. You also have the option to transform the data using the datatype and datetime_format options. Note The parentheses around the binding, data type, and format specifications are not required but are included in this document for readability.
binding	string	The <b>binding</b> is a literal value that binds a ?variable to a source column. If you specify a ?variable that matches the source column name, then that variable name is the data selector and it is not necessary to specify a binding. If you specify an alternate variable name or there is a hierarchical path to the source column that is not already identified by the selector, database, schema, table, or query properties, then the binding is needed to map the new variable to that source column. For example, ?subject ("dbo.FILM.SUBJECT") binds the ?subject variable by navigating to the SUBJECT column in the FILM table in the dbo schema.
		<b>Note</b> Database, schema, and table names in bindings are

Option	Туре	Description
		parsed according to the specific rules for that database type. You do not need to escape characters in database names. However, database names with characters that do not match (_ A-Z a-z) (_ A-Z a-z 0-9) * should be quoted, such as ("'Adventure.Works'.Sales.'Daily.Totals'").
datatype	URI	The datatype is the data type to convert the column to. If you do not specify a data type, the GDI infers the type. The GDI supports the following types: xsd:int, xsd:long, xsd:float, xsd:double, xsd:boolean, xsd:time, xsd:dateTime, xsd:date, xsd:duration, xsd:dayTimeDuration, xsd:yearMonthDuration, xsd:gMonthDay, xsd:gMonth, xsd:gYearMonth, xsd:anyURI
datetime_ stri format	string	This option is used to specify the format to use for date and time data types. The GDI supports Java date and time formats. Specify days as "d," months as "M," and years as "y." For the time, specify "H" for hours, "m" for minutes, and "s" for seconds. For example, "yyyyMMdd HH:mm:ss" or "ddMMMyy" to display date values such as "01JAN19."
		Note The GDI's default base year is 2000. If the source data has years with only two digits, such as 02-04-99, the GDI prepends 20 to the digits. The value 02-04-99 is parsed to 02-04-2099. To specify an alternate base year to use for two-digit values, you can include the notation

Option	Туре	Description
		^nnnn (e.g., ^1900) in the format value. For
		example, to set the base year to 1900 instead of 2000,
		use a format value such as xsd:date "dd-MMM-
		yy^1900" or xsd:date "dd-MMM-yy^1990". When
		one of those values is specified, 02-04-99 is parsed to
		02-04-1999.

#### **Query Examples**

The example below selects data from the AdventureWorks2012 database. The s:selector property is used to specify the table (salesOrderHeader in the Sales schema) to target.

```
<http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX s:
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
SELECT (COUNT(*) as ?count)
WHERE
{
   SERVICE <http://cambridgesemantics.com/services/DataToolkit>
{
      ?SalesOrderHeader a s:DbSource ;
         s:url "{{@db.eca4bfa83481f3638b93ab5fdf93ff9a.url}}" ;
         s:username "{{@db.eca4bfa83481f3638b93ab5fdf93ff9a.user}}";
         s:password "{{@db.eca4bfa83481f3638b93ab5fdf93ff9a.password}}" ;
         s:selector "Sales.SalesOrderHeader" ;
         ?SalesOrderID (xsd:int) ;
         ?RevisionNumber (xsd:int) ;
         ?OrderDate (xsd:dateTime) ;
         ?DueDate (xsd:dateTime) ;
         ?TerritoryID (xsd:int) ;
         ?TotalDue (xsd:decimal) .
      FILTER(?TerritoryID IN (1, 2, 3))
      FILTER(?TotalDue < 11.0 || ?TotalDue > 250)
  }
}
```

The example below ingests data from a database. To define the data to target, the query includes the s:query property to run an SQL query. The s:table and partitionBy properties are also included to aid the GDI in partitioning the query.

```
<http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX s:
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX ont:
<http://cambridgesemantics.com/Layer/2f1e926b130a402db6fc10fa54199d49/Model#>
INSERT {
   GRAPH ${targetGraph} {
      ?resource a ont:EmrPatient ;
      ont:EmrPatient.patientid ?PATIENTID ;
      ont:EmrPatient.gender ?GENDER ;
      ont:EmrPatient.language ?LANGUAGE ;
      ont:EmrPatient.patientfirstdocactivitydate ?PATIENTFIRSTDOCACTIVITYDATE .
  }
}
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
      ?data a s:DbSource ;
      s:url "{{@db.eca4bfa83481f3638b93ab5fdf93ff9a.url}}" ;
      s:username "{{@db.eca4bfa83481f3638b93ab5fdf93ff9a.user}}";
      s:password "{{@db.eca4bfa83481f3638b93ab5fdf93ff9a.password}}";
      s:query "select * from emrdbsmall.emr_patient where emr_patient.PATIENTID < 500"
;
      s:partitionBy "PATIENTID" ;
      s:table "emrdbsmall.emr patient" ;
      ?PATIENTID (xsd:int) ;
      ?GENDER (xsd:string) ;
      ?LANGUAGE (xsd:string) ;
      ?PATIENTFIRSTDOCACTIVITYDATE (xsd:dateTime "M/d/yyyy HH:mm:ss") .
   BIND(IRI("http://cambridgesemantics.com/Layer/2f1e926b130a402db6fc10fa54199d49/
{{?PATIENTID}}") AS ?resource)
  }
}
```

#### **Querying an HTTP Source**

This topic provides details about the structure to use when writing GDI queries to read or ingest data from HTTP data sources. It also includes example queries that may be useful as a starting point for writing your own GDI queries.

- Query Syntax
- Mapping the Content Property to JSON
- Query Examples

#### **Query Syntax**

The following query syntax shows the structure of a GDI query for HTTP sources. The clauses, patterns, and placeholders that are links are described below.

```
# PREFIX Clause
PREFIX s: <a href="http://cambridgesemantics.com/ontologies/DataToolkit#">http://cambridgesemantics.com/ontologies/DataToolkit#</a>
PREFIX rdf:
             <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
PREFIX dc: <http://purl.org/dc/elements/1.1/>
# Result Clause
{
   [ GRAPH ${targetGraph} { ]
   triple patterns
 [ } ]
}
[ ${usingSources} ]
WHERE
{
   # SERVICE Clause: Include the following service call when reading or inserting data.
    SERVICE [ TOPDOWN ] <http://cambridgesemantics.com/services/DataToolkit>
   # View SERVICE Clause: Or use the service call below when constructing a view.
    SERVICE <http://cambridgesemantics.com/services/DataToolkitView>(${targetGraph})
```

```
?data a s:HttpSource ;
  s:url "string" ;
  [ s:authorization [
      a s:BearerToken ; s:token "string" ;
    | a s:AWSSignature ; s:accessKey "string" ; s:region "string" ;
             s:secretKey "string" ; s:serviceName "string" ;
             s:sessionToken "string" ;
   | a s:BasicAuth ; s:username "string" ; s:password "string" ;
  1;1
  [ s:trust "string" ; ]
  [ s:proxy "string" | [ s:host "string" ; s:port int ] ]
  [ s:header [ s:name: "string" ; s:value "string" ] ; ]
  [ s:mimetype "string" ; ]
  [ s:contentType "string" ; ]
  [ s:content """string"""; ]
  [ s:parameter [ s:name "string" ; s:value "string" ] ; ]
  [ s:method "string" ; ]
  [ s:encoding "string" ; ]
  [ s:form [ s:name: "string" ; s:value "string" ] ; ]
  [ s:format [ source format options ; ] ; ]
  [ s:timeout int ; ]
  [ s:batching boolean | int ; ]
  [ s:paging [ pagination options ; ]
  [ s:concurrency int | [ list of properties ] ; ]
  [ s:rate int | "string" ; ]
  [ s:partitionBy "string" | ?variable ; ]
  [ s:locale "string" ; ]
  [ s:sampling int ; ]
  [ s:selector "string" | [ list ] ; ]
  [ s:model "string" ; ]
  [ s:key ("string") ; ]
  [ s:reference [ s:model "string" ; s:using ("string") ]
  [ s:formats [ datatype formatting options ] ; ]
  [ s:normalize boolean | [ normalization rules ] ; ]
  [ s:count ?variable ; ]
  [ s:offset int ; ]
  [ s:orderBy "string" | ?variable ; ]
  [ s:limit int ; ]
  # Mapping variables
  ?mapping_variable ( [ "binding" ] [ datatype ] [ "datetime_format" ] ) ;
  . . . . ;
```

```
# Additional clauses such as BIND, VALUES, FILTER
}
```

### Note

### For readability, the parameters below exclude the base URI

<http://cambridgesemantics.com/ontologies/DataToolkit#> as well as the s: prefix. As shown in the examples, however, the s: prefix or full property URI does need to be included in queries.

Option	Туре	Description
PREFIX Clause	N/A	The PREFIX clause declares the standard and custom prefixes for GDI service queries. Generally, queries include the prefixes from the query template (or a subset of them) plus any data-specific declarations.
Result Clause	N/A	The result clause defines the type of SPARQL query to run and the set of results to return, i.e., whether you want to read (SELECT or CONSTRUCT) from the source or ingest the data into Anzo (INSERT).
GRAPH \${targetGraph}	N/A	Include the GRAPH keyword and target graph parameter \${targetGraph} when writing an INSERT query to ingest data into a graphmart. Anzo automatically populates the query with the appropriate target URIs when the query runs.
\${usingSource s}	N/A	Include the source graph parameter $\{usingSources\}\$ when writing a "topdown" query that passes values from the data that is in the graphmart to the data source. Anzo automatically populates the query with the appropriate FROM clauses when the query runs. When passing literal values to the remote source, you do not need to include the source graph parameter. The SERVICE Clause description below includes more information about passing

Option	Туре	Description
		input to data sources.
SERVICE Clause	N/A	<pre>Include the SERVICE call SERVICE [ TOPDOWN ] <http: cambridgesemantics.com="" datatool="" kit="" services=""> to invoke the GDI service when you are running a SELECT, INSERT, or CONSTRUCT query that is not creating a view. When writing a CONSTRUCT query in a View Step, use the DataToolkitView service call, as described below in View SERVICE Clause.</http:></pre>
		Include the optional TOPDOWN keyword when you want to pass input values from the graphmart to the data source. When you include TOPDOWN in the service call, it indicates that the rest of the query produces values to send to the source. In this case, the GDI makes repeated calls to pass in each of the specified values and retrieve the data that is based on those values.
View SERVICE Clause	N/A	When writing a CONSTRUCT query that creates a view of the data (usually in a View Step), include the following SERVICE call: SERVICE <http: cambridgesemantics.com="" datatool<br="" services="">kitView&gt;(\${targetGraph}). Using the DataToolkitView call optimizes query execution because it tells the GDI to inspect the query and determine which filters to push to the data source. It also limits the result set and retrieves only the data that is needed, i.e., the source data is fully mapped but all of the mapped data is not necessarily returned.</http:>
url	string	This property specifies the URL to use to access the source. Query binding variables can be inserted into the url string by surrounding the variable name with double curly braces. For

Option	Туре	Description
		<pre>example, "{{?name}}".</pre>
		<pre>Important For security, it is a best practice to reference connection information (such as the url, username, and password) from a Query Context so that the sensitive details are abstracted from any requests. In addition, using a Query Context makes connection details reusable across queries. See Using Query Contexts for more information. For example:</pre>
authorization	RDF list	This property specifies the type of authorization to use and the values for authentication. The options are BearerToken, AWSSignature, or BasicAuth. s:authorization [ a s:BearerToken   s:AWSSignature   s:BasicAuth ]
BearerToken	string	<pre>Specify this property when a bearer token is used for authentication, and include the token property. s:authorization [ a s:BearerToken ; s:token "string"</pre>

Option	Туре	Description
		]
AWSSignature	RDF list	<pre>For authorization to AWS service endpoints, specify this property and include the appropriate authentication properties from the list below:     accessKey: Include this property to specify the AWS     access key.     region: Include this property to specify the AWS region.     secretKey: Include this property to specify the AWS     secret key.     serviceName: Include this property to specify the AWS     service name.     sessionToken: Include this property to specify the AWS     session token. s:authorization [     a s:AWSSignature ; s:accessKey "string";     s:region "string" ; s:sessionToken "string" ; ]</pre>
BasicAuth	RDF list	<pre>Specify this property when basic authentication is used, and include the username and password properties. s:authorization [ a s:BasicAuth ; s:username "string" ; s:password "string" ; ]</pre>

Option	Туре	Description
trust	string	Include this property to set the level of trust for the source's SSL certificate. The value can be either "system" or "all".
proxy	string or RDF list	<pre>Include this property to specify proxy information if a proxy is used. The value can be a string, such as s:proxy "host_url:port_ number", or an RDF list that includes host and port properties, such as s:proxy [ s:host "host_url" ; s:port port_ number ].</pre>
header	RDF list	You can use this property to specify name-value pairs to include as headers in the request. For example:
		<pre>s:header [ s:name "Accept" ; s:value "application/json" ]</pre>
		If you are creating a view, you can include variables in the s:header list. When another query is run against a view with variables, that query can map the variables through the view by including predicates in the CONSTRUCT clause.
mimetype	string	You can include this property to specify the MIME type of the source. For example, s:mimetype "text/html".
contentType	string	<pre>Include this property to specify the content type of the body of the request. For example, s:contentType "application/sparql-query" or s:contentType "application/json".</pre>
content	string or RDF list	This property can be included to send content to the source in the body of the request. For example, content can be a SPARQL query, JSON arrays, or a list of key-value pairs. Content can also be configured with an inline object (blank node) that gets translated to JSON. For more information, see Mapping the

Option	Туре	Description
		Content Property to JSON below.
parameter	RDF list	You can include this property to list any URL parameters as name- value pairs. For example, the s:parameter property below adds format to return results in CSV format and the named-graph- uri parameter to target a specific layer in a graphmart. s:parameter [ s:name "format" ; s:value "csv" ] , [ s:name "named-graph-uri" ; s:value "http://cambridgesemantics.com/Layer/d541" ] If you are creating a view, you can include variables in the s:parameter list. When another query is run against a view with variables, that query can map the variables through the
		view by including predicates in the CONSTRUCT clause.
method	string	You can include this property to specify the HTTP method. For example, s:method "GET" or s:method "POST".
encoding	string	When targeting a file, you can include this property to specify the character encoding used by the file. The default value is s:encoding "utf8".
form	RDF list	To send data to the HTTP endpoint, you can use this property to post the data. Form is a list of name-value pairs. When including s:form, you must also include s:contentType "multipart/form-data". The GDI sends the form object as an application/x-www-form-urlencoded string that contains the specified parameters. See Example form Parameter Usage below for sample usage.

Option	Туре	Description
format	RDF list	If the data is file-based, you can include the format property to add parameters that describe the source. See File Source Format Options for details about the supported parameters.
timeout	int	This property can be used to specify the timeout (in milliseconds) to use for requests against the source. For example, s:timeout 5000 configures a 5 second timeout.
batching	boolean or int	This property can be used to disable batching, or it can be used to change the default the batch size. By default, batching is set to 5000 (s:batching 5000). To disable batching, you can include s:batching false in the query. Typically users do not change the batching size. However, it can be useful to control the batch size when performing updates. To configure the size, include s:batching int in the query. For example, s:batching 3000.
paging	RDF list	This property can be used to configure paging so that the GDI can access large amounts of data across a number of smaller requests. For details about the paging property, see Paginating Requests.
concurrency	int or RDF list	This property can be included to configure the maximum level of concurrency for the query. The value can be an integer, such as s:concurrency 8. If the value is an integer, it configures a maximum limit on the number of slices that can execute the query. For finer-grained control over the number of nodes and slices to use, concurrency can also be included as an object with limit, nodes, and/or executorsPerNode properties. For example, the following object configures a concurrency model that allows a maximum of 24 executors distributed across 4 nodes with 8 executors per node:

Option	Туре	Description
		<pre>s:concurrency [    s:limit 24 ;    s:nodes 4 ;    s:executorsPerNode 8 ; ] ;</pre>
rate	int or string	This property can be included to control the frequency with which a request is sent to the source. The limit applies to the number of requests a single slice can make. If you specify an integer for the rate, then the value is treated as the maximum number of requests to issue per minute. If you specify a string, you have more flexibility in configuring the rate. The sample values below show the types of values that are supported:
		<pre>s:rate "90/minute" ; s:rate "90 per minute" ; s:rate "200000 every week" ; s:rate "10000 every 6 hours" ;</pre>
		To enforce the rate limit, the GDI introduces a sleep between requests that is equal to the rate delay. The more executing slices, the longer the rate delay needs to be to enforce the limit in aggregate.
		Given the example of s:rate "90/minute", the GDI would optimize the concurrency and only use 1 slice for execution with a rate delay of 666ms between requests. If s:rate "240/minute", the GDI would use 3 executors with a rate delay of 750ms between requests.
partitionBy	string, variable,	The GDI attempts to partition queries automatically across the available cores (slices) in AnzoGraph. To determine how to

Option	Туре	Description
	list	partition the query, the GDI uses metadata from the source. It looks for any column in an index, preferring the primary key column if it is interpolable. However, it only considers the first column in any index on the table. After determining the partition column, the GDI does a MIN/MAX on the column as well as a basic sizing query. To specify which column or columns the GDI should partition on, you can include the partitionBy property in the query. The property supports a list of source field names, bound variables, or the object s:auto, which forces the GDI to partition the data when the source does not define partitioning metadata.
locale	string	This property can be used to specify the locale to use when parsing locale-dependent data such as numbers, dates, and times.
sampling	int	This property can be used to configure the number of records in the source to examine for data type inferencing.
selector	string or RDF list	This property can be used as a binding component to identify the path to the source objects. For example, s:selector "Sales.SalesOrderHeader" targets the SalesOrderHeader table in the Sales schema. For more information about binding components and the selector property, see Using Binding Trees and Selector Paths.
model	string	This property defines the class (or table) name for the type of data that is generated from the specified data source. For example, s:model "employees". Model is optional when querying a single source. If your query targets multiple sources, however, and you want to define resource templates (primary keys) and object properties (foreign keys), you must specify the model value for

Option	Туре	Description
		each source.
key	string	This property can be used to define the primary key column for the source file or table. This column is leveraged in a resource template for the instances that are created from the source. For example, s:key ("EMPLOYEE_ID"). For more information about key, see Data Linking Options.
reference	RDF list	This property can be used to specify a foreign key column. The reference property is an RDF list that includes the model property to list the target table and a using property that defines the foreign key column. For more information about reference, see Data Linking Options.
formats	RDF list	To give users control over the data types that are used when coercing strings to other types, this property can be included in GDI queries to define the desired types. In addition, it can be used to describe the formats of date and time values in the source to ensure that they are recognized and parsed to the appropriate date, time, and/or dateTime values. For details about the formats property, see Data Type Formatting Options.
normalize	RDF list	To give users control over the labels and URIs that are generated, the GDI offers several options for normalizing the model and/or the fields that are created from the specified data source(s). For details about the normalize property, see Model Normalization Options.
count	variable	If you want to turn the query into a COUNT query, you can include this property with a <code>?variable</code> to perform a count. For example, s:count <code>?count</code> .

Option	Туре	Description
offset	int	This property can be used to offset the data that is returned by a number of rows.
orderBy	string, variable, list	You can include this property to order the result set by a field name, a bound variable, or a list of names or bound variables.
limit	int	You can include this property to limit the number of results that are returned. s:limit maps to the SPARQL LIMIT clause.
mapping_ variable	variable	The mapping variables, in <code>?mapping_variable</code> (["binding"] [datatype] ["datetime_format"]) format, define the triple patterns to output. When the specified ?variable matches the source column name, the GDI uses the variable as the source data selector. If you specify an alternate variable name, a binding needs to be specified to map the new variable to the source. You also have the option to transform the data using the datatype and datetime_format options.
		Note The parentheses around the binding, data type, and format specifications are not required but are included in this document for readability.
binding	string	The binding is a literal value that binds a ?mapping_variable to a source column. If you specify a ?variable that matches the source column name, then that variable name is the data selector and it is not necessary to specify a binding. If you specify an alternate variable name or there is a hierarchical path to the source column, then the binding is needed to map the new variable to that source column.

Option	Туре	Description
		For example for CSV, the following pattern simply binds the source column AIRLINE to the lowercase variable ?airline: ?airline ("AIRLINE").
		Note For FileSource, periods (.), forward slashes (/), and brackets ([]) are parsed as path notation. Therefore, if a source column name includes any of those characters they must be escaped in the binding. Use two backslashes (\\) as an escape character. For example, if a column name is <b>average/day</b> , the variable and binding pattern could be written as ?averagePerDay ("average\\/day").
datatype URI	The datatype is the data type to convert the column to. If you do not specify a data type, the GDI infers the type. The GDI supports the following types:	
		<pre>xsd:int, xsd:long, xsd:float, xsd:double, xsd:boolean, xsd:time, xsd:dateTime, xsd:date, xsd:duration, xsd:dayTimeDuration, xsd:yearMonthDuration, xsd:gMonthDay,</pre>
		xsd:gMonth, xsd:gYearMonth, xsd:anyURI
datetime_ format	string	This option is used to specify the format to use for date and time data types. The GDI supports Java date and time formats. Specify days as "d," months as "M," and years as "y." For the time, specify "H" for hours, "m" for minutes, and "s" for seconds. For example, "yyyyMMdd HH:mm:ss" or "ddMMMyy" to display date values such as "01JAN19."

Option	Туре	Description
		Note The GDI's default base year is 2000. If the source data has years with only two digits, such as 02-04-99, the GDI prepends 20 to the digits. The value 02-04-99 is parsed to 02-04-2099. To specify an alternate base year to use for two-digit values, you can include the notation ^nnnn (e.g., ^1900) in the format value. For example, to set the base year to 1900 instead of 2000, use a format value such as xsd:date "dd-MMM-yy^1900" or xsd:date "dd-MMM-yy^1990". When one of those values is specified, 02-04-99 is parsed to 02-04-1999.

### Mapping the Content Property to JSON

The s:content property can be configured with an inline object (blank node) that gets translated to JSON in the request body. This mapping allows for creation of embedded objects and arrays as well as a mechanism for iterating over all available input so that HTTP endpoints that support batching can be used more effectively.

### **Using Blank Nodes**

Blank nodes are used to create an object in the output JSON. The local name of any predicate used within content becomes a key in the generated JSON object. Blank nodes can be embedded within each other, allowing the hierarchical nature of JSON to be represented. For example:

```
s:content [ ex:firstName "Mary" ; ex:lastName "Barry" ] ;
```

Or

```
s:content [ ex:person [ ex:firstName "Mary" ] ] ;
```

# **Using Variables**

Variables can be also used in the object position to construct a request from input at runtime. For example:

s:content [ ex:firstName ?firstName ; ex:lastName ?lastName ] ;

The values for the variables can come from a TOPDOWN variable, a VALUES clause in the SERVICE block, or another data source. Any unbound variables in the input will not be added to the generated JSON object.

# **Using RDF Lists**

An RDF list can also be used to create an array in the output JSON. For example:

s:content [ ex:allKnownNames ( ?firstName ?lastName ?nickName ) ]

An RDF list can also be embedded inside another list to create an array in the output JSON and populate it with items evaluated against a repeating pattern across all available input rows for a slice. That pattern can be a variable, which generates an array of primitive values, or a blank node, which generates an array of mapped JSON objects. For example:

s:content [ ex:documents ((?id)) ] ;

Or

s:content [ ex:documents (([ ex:id ?id ; ex:title ?title ])) ;

# Example

The following example query demonstrates the use of s:content to generate JSON. The query also includes the s:concurrency property to restrict execution to a single slice. Without limiting execution when there are a small number of inputs (as in the VALUES clause), each input row gets executed on its own. As the inputs increase, each slice operates over a larger number of inputs until the default s:batching 5000 is applied.

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX api: <http://contoso.com/api/>
```

```
SELECT *
WHERE {
   SERVICE TOPDOWN <http://cambridgesemantics.com/services/DataToolkit>
   {
     VALUES (?firstName ?lastName ?dob ?email)
      {
        ("Gray" "Hay" "1978-03-18"^^xsd:date "gray@abc.com")
        ("Ana" "Bana" "1974-10-20"^^xsd:date "ana@abc.com")
        ("George" "Forge" "1975-08-13"^^xsd:date "george@abc.com")
        ("Miles" "Giles" "1977-04-12"^^xsd:date "miles@abc.com")
      }
      ?data a s:HttpSource ;
        s:url "https://postman-echo.com/post" ;
        s:header [ s:name "Accept" ; s:value "application/json" ] ;
        s:concurrency 1 ;
        s:content
        (([
           api:dateOfBirth ?dob ;
           api:email ?email ;
           api:year 2020 ;
           api:person [ api:firstName ?firstName ; api:lastName ?lastName ] ;
        ]));
        s:selector "data" ;
        ?firstName ("person.firstName" xsd:string) ;
        ?lastName ("person.lastName" xsd:string) ;
        ?dob ("dateOfBirth" xsd:date) ;
        ?email ("email" xsd:string) ;
        ?year ("year" xsd:int) .
  }
}
```

The content portion of the request that the query generates is shown below:

```
[{
    "firstName": "Gray",
    "lastName": "Hay",
    "dateOfBirth": "1978-03-18",
    "email": "gray@abc.com",
    "year": 2020
},
{
```

```
"firstName": "Ana" ,
    "lastName": "Bana" ,
    "dateOfBirth": "1974-10-20" ,
    "email": "ana@abc.com" ,
    "year": 2020
    },
    {
        ...
}]
```

### **Query Examples**

- Topdown Query with URL Parameters
- Generator Query against an Anzo SPARQL Endpoint
- API Queries
- Example form Parameter Usage

### **Topdown Query with URL Parameters**

The query below reads data from a sample HTTP source that compiles worldwide weather statistics. The source has several models available for retrieving data that is current, daily, historical, etc. To target current data, the query includes <code>s:selector "currently"</code>. In addition, the query demonstrates the use of the "topdown" functionality, where the query sends values to the source to narrow the results. The VALUES clause specifies the latitude and longitude values for the cities to return data for. In addition, since this sample source requires parameters to be specified in the connection URL, the <code>s:url</code> value includes <code>?lat</code> and <code>?long</code> as parameters as part of the value.

```
SERVICE TOPDOWN <http://cambridgesemantics.com/services/DataToolkit>
    {
         ?data a s:HttpSource ;
         s:url "https://sampleEndpoint.com/forecast/{{?lat}},{{?long}}";
         s:selector "currently" ;
         ?lat ("latitude") ;
         ?long ("longitude") ;
         ?temp ("temperature") ;
         ?rainChance ( "precipProbability" ) ;
         ?humidity () ;
         ?pressure () ;
         ?windSpeed () .
    }
   VALUES ( ?city ?state ?lat ?long )
    {
         ("Lakeway" "TX" 30.374563 -97.975892)
         ( "Boston" "MA" 42.358043 -71.060415 )
         ("Seattle" "WA" 47.590720 -122.307053)
         ( "Chicago" "IL" 41.837741 -87.823296 )
         ("Hilo" "HI" 19.702040 -155.090312)
    }
}
ORDER BY ?city
```

# Generator Query against an Anzo SPARQL Endpoint

The example below is a GDI Generator query that retrieves graphmart data from a remote Anzo SPARQL endpoint.

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
INSERT {
    GRAPH ${targetGraph} {
        ?s ?p ?o }
}
${usingSources}
WHERE {
    SERVICE <http://cambridgesemantics.com/services/DataToolkit>
```

```
?data a s:HttpSource ;
            s:url
"https://10.10.0.10/sparql/graphmart/http%3A%2F%2Fcambridgesemantics.com%2FGraphmart%2F
1686168b-3eaf-4fdc-9730-1903717b9e62";
            s:trust "all" ;
            s:username "user";
            s:password "pass";
            s:contentType "application/sparql-query" ;
            s:header [ s:name "Accept" ; s:value "text/csv" ] ;
        s:content """
          PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
          SELECT ?s ?p ?o
          WHERE {
            ?s ?p ?o .
           FILTER(ISLITERAL(?o))
          }
        .....
        ?rdf a s:RdfGenerator, s:OntologyGenerator ;
            s:as (?s ?p ?o) ;
            s:ontology <http://cambridgesemantics.com/ontologies/TopMovies> ;
            s:base <http://cambridgesemantics.com/data> .
  }
}
```

# **API Queries**

The following example queries the Google Recognize API to request transcriptions for voice recordings that are stored in a Google bucket.

```
?record <http://google.com/confidence> ?confidence .
    ?record <http://google.com/file> ?file .
 }
}
${usingSources}
WHERE {
   BIND(<gs://csi-se/demo/emergency-test.mp3> as ?file)
    BIND(UUID() as ?record)
{
        SERVICE <http://cambridgesemantics.com/services/DataToolkit>
        {
            ?data a s:HttpSource ;
              s:selector "results.alternatives" ;
              s:url "https://speech.googleapis.com/v1p1beta1/speech:recognize" ;
              s:authorization [ a s:BearerToken ; s:token """ya29...""" ] ;
              s:content """
    {
      "config": {
        "encoding": "MP3",
        "sampleRateHertz": 16000,
        "languageCode": "en-US",
        "enableWordTimeOffsets": false
    },
      "audio": {
        "uri":"gs://csi-se/demo/emergency-test.mp3"}
    }
    ?confidence ("confidence") ;
            ?transcript ("transcript") .
      }
   }
```

The example below includes the header and content properties to send a request that contains small text snippets for sentiment analysis.

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX ont: <http://cambridgesemantics.com/ontologies/Sentiment_Analysis#>
INSERT {
```

```
GRAPH ${targetGraph} {
    ?requirement a ont:Sentiment ;
      ont:p Sentiment Type ?sentiment ;
      ont:p Sentiment Score ?polarity .
  }
}
${usingSources}
WHERE {
  ?requirement a
<http://cambridgesemantics.com/Layer/3b4163e7f53149d5a815627be5d409bd/Model#Requirement
s> ;
<http://cambridgesemantics.com/Layer/3b4163e7f53149d5a815627be5d409bd/Model#Requirement
s.reqText> ?requirement text .
  SERVICE TOPDOWN <a href="http://cambridgesemantics.com/services/DataToolkit">http://cambridgesemantics.com/services/DataToolkit</a> {
      ?data a s:HttpSource ;
         s:url "https://text-analysis12.p.rapidapi.com/sentiment-analysis/api/v1.1" ;
        s:method "POST" ;
        s:header [ s:name "Accept" ; s:value "application/json" ] ,
                  [ s:name "X-RapidAPI-Key" ; s:value "key" ] ,
                  [ s:name "X-RapidAPI-Host" ; s:value "text-analysis12.p.rapidapi.com"
];
         s:contentType "application/json" ;
         s:content """{ "text": "{{?requirement text}}" , "language": "english" }""" ;
         ?polarity ("aggregate sentiment/compound" xsd:double);
         ?sentiment () .
  }
}
```

### **Example form Parameter Usage**

The query snippet below shows an example that incorporates the s:form parameter.

```
SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
   VALUES (?reviews_text) {
    ("Horrible, terrible, will never use again.")
    ("Wonderful, magnificent, will recommend to everyone!")
    ("Simply OK. I might purchase this product again.") }
   ?data a s:HttpSource ;
    s:url "https://api.meaningcloud.com/sentiment-2.1" ;
    s:contentType "multipart/form-data" ;
```

```
s:form [ s:name "key" ; s:value "9eab751142..." ],
      [ s:name "lang" ; s:value "auto" ],
      [ s:name "txt" ; s:value ?reviews_text ] ;
?confidence ();
?score_tag ();
?subjectivity () ;
?irony () ;
?agreement () .
```

#### **Querying an Elasticsearch Source**

This topic provides details about the structure to use when writing GDI queries to read or ingest data from Elasticsearch data sources. It also includes example queries that may be useful as a starting point for writing your own GDI queries.

• Query Syntax

}

- Query DSL and Filter Mapping
- Query Examples

#### **Query Syntax**

The following query syntax shows the structure of a GDI query for Elasticsearch sources. The clauses, patterns, and placeholders that are links are described below.

```
# PREFIX Clause
PREFIX s: <a href="http://cambridgesemantics.com/ontologies/DataToolkit#">http://cambridgesemantics.com/ontologies/DataToolkit#</a>>
             <http://elastic.co/search/>
PREFIX es:
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
             <http://purl.org/dc/elements/1.1/>
PREFIX dc:
# Result Clause
{
   [ GRAPH ${targetGraph} { ]
   triple patterns
 [ } ]
```

```
[ ${usingSources} ]
```

#### WHERE

{

```
# SERVICE Clause: Include the following service call when reading or inserting data.
SERVICE [ TOPDOWN ] <http://cambridgesemantics.com/services/DataToolkit>
```

# View SERVICE Clause: Or use the service call below when constructing a view. SERVICE <http://cambridgesemantics.com/services/DataToolkitView>(\${targetGraph})

```
{
 ?data a s:ElasticSource ;
   s:url "string" ;
   [ s:username "string" ; ]
   [ s:password "string" ; ]
   [ s:property [ s:name "string" ; s:value "string" ; ]
   [ es:aggregations [ rdf list ] ; ]
   [ es:config "string" ; ]
   [ es:document "string" ; ]
   [ es:field "string" | ?variable ; ]
   [ es:highlight [ rdf list ] ; ]
   [ es:html boolean ; ]
   [ es:index "string" ; ]
   [ es:minScore float ; ]
   [ es:query "string" | [ rdf_list ] ; ]
   [ es:routing "string" ; ]
   [ es:searchAfter [ rdf list ] ; ]
   [ es:size int ; ]
   [ es:source boolean | [ rdf_list ] ; ]
   [ s:timeout int ; ]
   [ s:batching boolean | int ; ]
   [ s:paging [ pagination options ; ]
   [ s:concurrency int | [ list of properties ] ; ]
   [ s:rate int | "string" ; ]
   [ s:partitionBy "string" | ?variable ; ]
   [ s:locale "string" ; ]
   [ s:sampling int ; ]
   [ s:selector "string" | [ list ] ; ]
   [ s:model "string" ; ]
   [ s:key ("string") ; ]
   [ s:reference [ s:model "string" ; s:using ("string") ]
    [ s:formats [ datatype formatting options ] ; ]
```

```
[ s:normalize boolean | [ normalization_rules ] ; ]
[ s:count ?variable ; ]
[ s:offset int ; ]
[ s:orderBy "string" | ?variable ; ]
[ s:limit int ; ]
# Mapping variables
?mapping_variable ( [ "binding" ] [ datatype ] [ "datetime_format" ] ) ;
...;
# Additional clauses such as BIND, VALUES, FILTER
}
```

### Note

}

#### For readability, the parameters below exclude the base URIs

<http://cambridgesemantics.com/ontologies/DataToolkit#> and <http://elastic.co/search/> as well as the s: and es: prefixes. As shown in the examples, however, the prefixes or full property URIs do need to be included in queries.

Option	Туре	Description
PREFIX Clause	N/A	The PREFIX clause declares the standard and custom prefixes for GDI service queries. Generally, queries include the prefixes from the query template (or a subset of them) plus any data-specific declarations.
Result Clause	N/A	The result clause defines the type of SPARQL query to run and the set of results to return, i.e., whether you want to read (SELECT or CONSTRUCT) from the source or ingest the data into Anzo (INSERT).
GRAPH \${targetGraph}	N/A	Include the GRAPH keyword and target graph parameter \${targetGraph} when writing an INSERT query to ingest data into a graphmart. Anzo automatically populates the query with the appropriate target URIs when the query runs.

Option	Туре	Description
\${usingSource s}	N/A	Include the source graph parameter $\{usingSources\}\$ when writing a "topdown" query that passes values from the data that is in the graphmart to the data source. Anzo automatically populates the query with the appropriate FROM clauses when the query runs. When passing literal values to the remote source, you do not need to include the source graph parameter. The SERVICE Clause description below includes more information about passing input to data sources.
SERVICE Clause	N/A	<pre>Include the SERVICE call SERVICE [ TOPDOWN ] <http: cambridgesemantics.com="" datatool="" kit="" services=""> to invoke the GDI service when you are running a SELECT, INSERT, or CONSTRUCT query that is not creating a view. When writing a CONSTRUCT query in a View Step, use the DataToolkitView service call, as described below in View SERVICE Clause.</http:></pre>
		Include the optional TOPDOWN keyword when you want to pass input values from the graphmart to the data source. When you include TOPDOWN in the service call, it indicates that the rest of the query produces values to send to the source. In this case, the GDI makes repeated calls to pass in each of the specified values and retrieve the data that is based on those values.
View SERVICE Clause	N/A	When writing a CONSTRUCT query that creates a view of the data (usually in a View Step), include the following SERVICE call: SERVICE <http: cambridgesemantics.com="" datatool<br="" services="">kitView&gt;(\${targetGraph}). Using the DataToolkitView call optimizes query execution because it tells the GDI to inspect the query and determine which filters to push to the data source. It</http:>

Option	Туре	Description
		also limits the result set and retrieves only the data that is needed, i.e., the source data is fully mapped but all of the mapped data is not necessarily returned.
url	string	This property specifies the URL to use to access the source.
		<pre>Important For security, it is a best practice to reference connection information (such as the url, username, and password) from a Query Context so that the sensitive details are abstracted from any requests. In addition, using a Query Context makes connection details reusable across queries. See Using Query Contexts for more information. For example:     ?data a s:ElasticSource ;     s:url "{{@es.hostname}}:{{@es.port}}" ;     s:username "{{@es.username}}" ; </pre>
username	string	<pre>This property lists the user name to use for the connection to the Elasticsearch.  Tip If you want to group the username and password properties, you can wrap them with s:credentials [     J.For example:     s:credentials [     s:username "username" ;     s:password "password" ;</pre>

Option	Туре	Description
		]
password	string	This property lists the password for the given username.
property	RDF list	<pre>This property can be included to list any source-specific configuration values. s:property [ s:name "custom_property_name" ; s:value "custom_value" ]</pre>
aggregations	RDF list	You can include this property to calculate aggregations over the specified bindings. For information about aggregations, see <u>Aggregations</u> in the Elasticsearch documentation.
config	string	To enable you to use explicit mappings, you can include this property to specify the URL to the index configuration file to employ. For example, es:config "/opt/shared/elastic/mapping.json".
document	string	This property lists the document(s) to search.
field	string or variable	This property defines the field to operate on. The value can be a string or bound variable.
highlight	RDF list	You can include this property to define how results are highlighted. For information about the available properties, see Highlighting Elasticsearch Results.
html	boolean	This property controls whether to output HTML for highlighted results. Defaults to true.

Option	Туре	Description
index	string	This property can be included to specify the indexes to search. Specify multiple indexes in a comma-separated list. For example, es:index "projectA_mar", "projectA_apr" ;.
minScore	float	This property defines the minimum score for matching documents. Documents with a lower score are not included in the search results.
query	string or RDF list	This property defines the query to execute. The value can be a string or a query object that maps to the Elasticsearch Query DSL. To generate the final query, the GDI combines es:query with any filters it can push to the Elasticsearch DSL. For more information about the query property and mapping Elasticsearch filters to SPARQL FILTER clauses, see Query DSL and Filter Mapping below.
routing	string	This property can be included to route a document to a specific shard or to limit the search to a particular shard.
searchAfter	RDF list	You can include this property to define the key values to start searching from.
size	int	This property maps to the size parameter in the Elasticsearch Search API and configures the batch size or maximum number of hits to return in a single call. Defaults to 10 and typically does not need to be changed.
source	boolean or RDF list	This property can be included to specify the source data to include in results. The value can be a boolean, list of fields, or a list of variable bindings. When true, all source data is returned. When false, no source data is returned.

Option	Туре	Description
timeout	int	This property can be used to specify the timeout (in milliseconds) to use for requests against the source. For example, s:timeout 5000 configures a 5 second timeout.
batching	boolean or int	This property can be used to disable batching, or it can be used to change the default the batch size. By default, batching is set to 5000 (s:batching 5000). To disable batching, you can include s:batching false in the query. Typically users do not change the batching size. However, it can be useful to control the batch size when performing updates. To configure the size, include s:batching int in the query. For example, s:batching 3000.
paging	RDF list	This property can be used to configure paging so that the GDI can access large amounts of data across a number of smaller requests. For details about the paging property, see Paginating Requests.
concurrency	int or RDF list	This property can be included to configure the maximum level of concurrency for the query. The value can be an integer, such as s:concurrency 8. If the value is an integer, it configures a maximum limit on the number of slices that can execute the query. For finer-grained control over the number of nodes and slices to use, concurrency can also be included as an object with limit, nodes, and/or executorsPerNode properties. For example, the following object configures a concurrency model that allows a maximum of 24 executors distributed across 4 nodes with 8 executors per node: s:concurrency [ s:limit 24 ; s:nodes 4 ;

Option	Туре	Description
		<pre>s:executorsPerNode 8 ; ] ;</pre>
rate	int or string	This property can be included to control the frequency with which a request is sent to the source. The limit applies to the number of requests a single slice can make. If you specify an integer for the rate, then the value is treated as the maximum number of requests to issue per minute. If you specify a string, you have more flexibility in configuring the rate. The sample values below show the types of values that are supported:
		s:rate "90/minute" ;
		<pre>s:rate "90 per minute" ; s:rate "200000 every week" ;</pre>
		s:rate "10000 every 6 hours" ;
		To enforce the rate limit, the GDI introduces a sleep between requests that is equal to the rate delay. The more executing slices, the longer the rate delay needs to be to enforce the limit in aggregate.
		Given the example of s:rate "90/minute", the GDI would optimize the concurrency and only use 1 slice for execution with a rate delay of 666ms between requests. If s:rate "240/minute", the GDI would use 3 executors with a rate delay of 750ms between requests.
partitionBy	string, variable, list	The GDI attempts to partition queries automatically across the available cores (slices) in AnzoGraph. To determine how to partition the query, the GDI uses metadata from the source database. It looks for any column in an index, preferring the primary key column if it is interpolable. However, it only considers

Option	Туре	Description	
		the first column in any index on the table. After determining the partition column, the GDI does a MIN/MAX on the column as well as a basic sizing query. To specify which column or columns the GDI should partition on, you can include the partitionBy property in the query. The property supports a list of source field names, bound variables, or the object s:auto, which forces the GDI to partition the data when the source does not define partitioning metadata.	
locale	string	This property can be used to specify the locale to use when parsing locale-dependent data such as numbers, dates, and times.	
sampling	int	This property can be used to configure the number of records in the source to examine for data type inferencing.	
selector	string or RDF list	This property can be used as a binding component to identify the path to the source objects. For example, s:selector "Sales.SalesOrderHeader" targets the SalesOrderHeader table in the Sales schema. For more information about binding components and the selector property, see Using Binding Trees and Selector Paths.	
model	string	This property defines the class (or table) name for the type of data that is generated from the specified data source. For example, s:model "employees". Model is optional when querying a single source. If your query targets multiple sources, however, and you want to define resource templates (primary keys) and object properties (foreign keys), you must specify the model value for each source.	
key	string	This property can be used to define the primary key column for the	
Option	Туре	Description	
-----------	----------	--	--
		source file or table. This column is leveraged in a resource template for the instances that are created from the source. For example, s:key ("EMPLOYEE_ID"). For more information about key, see Data Linking Options.	
reference	RDF list	This property can be used to specify a foreign key column. The reference property is an RDF list that includes the model property to list the target table and a using property that defines the foreign key column. For more information about reference, see Data Linking Options.	
formats	RDF list	To give users control over the data types that are used when coercing strings to other types, this property can be included in GDI queries to define the desired types. In addition, it can be used to describe the formats of date and time values in the source to ensure that they are recognized and parsed to the appropriate date, time, and/or dateTime values. For details about the formats property, see Data Type Formatting Options.	
normalize	RDF list	To give users control over the labels and URIs that are generated, the GDI offers several options for normalizing the model and/or the fields that are created from the specified data source(s). For details about the normalize property, see Model Normalization Options.	
count	variable	If you want to turn the query into a COUNT query, you can include this property with a <code>?variable</code> to perform a count. For example, <code>s:count ?count</code> . The GDI runs an Elasticsearch <u>value count</u> aggregation.	
offset	int	This property can be used to offset the data that is returned by a number of rows.	

Option	Туре	Description		
orderBy	string, variable, list	You can include this property to order the result set by a field name, a bound variable, or a list of names or bound variables.		
limit	int	You can include this property to limit the number of results that are returned. s:limit maps to the SPARQL LIMIT clause.		
mapping_ variable	variable	The mapping variables, in <code>?mapping_variable</code> (["binding"] [datatype] ["datetime_format"]) format, define the triple patterns to output. When the specified <code>?variable</code> matches the source column name, the GDI uses the variable as the source data selector. If you specify an alternate variable name, a binding needs to be specified to map the new variable to the source. You also have the option to transform the data using the datatype and datetime_format options.  Note The parentheses around the binding, data type, and format specifications are not required but are included in this document for readability.		
binding	string	The binding is a literal value that binds a ?mapping_variable to a source column. If you specify a ?variable that matches the source column name, then that variable name is the data selector and it is not necessary to specify a binding. If you specify an alternate variable name or there is a hierarchical path to the source column, then the binding is needed to map the new variable to that source column.		
datatype	URI	The datatype is the data type to convert the column to. If you do not specify a data type, the GDI infers the type. The GDI supports		

Option	Туре	Description		
		<pre>the following types: xsd:int, xsd:long, xsd:float, xsd:double, xsd:boolean, xsd:time, xsd:dateTime, xsd:date, xsd:duration, xsd:dayTimeDuration, xsd:yearMonthDuration, xsd:gMonthDay, xsd:gMonth, xsd:gYearMonth, xsd:anyURI</pre>		
datetime_ string format	This option is used to specify the format to use for date and time data types. The GDI supports Java date and time formats. Specify days as "d," months as "M," and years as "y." For the time, specify "H" for hours, "m" for minutes, and "s" for seconds. For example, "yyyyMMdd HH:mm:ss" or "ddMMMyy" to display date values such as "01JAN19."			
		Note The GDI's default base year is 2000. If the source data has years with only two digits, such as 02-04-99, the GDI prepends 20 to the digits. The value 02-04-99 is parsed to 02-04-2099. To specify an alternate base year to use for two-digit values, you can include the notation ^nnnn (e.g., ^1900) in the format value. For example, to set the base year to 1900 instead of 2000, use a format value such as xsd:date "dd-MMM-yy^1900" or xsd:date "dd-MMM-yy^1990". When one of those values is specified, 02-04-99 is parsed to 02-04-1999.		

**Query DSL and Filter Mapping** 

The vocabulary used in GDI queries against an ElasticSource closely mimics the Elasticsearch <u>Query DSL</u>. The table below shows a side-by-side view of a DSL query that is mapped to SPARQL using the es:query property:

DSL	SPARQL		
{	es:query [		
"query": {	a es:BoolQuery ;		
"bool" : {	es:must [		
"must" : {	a es:TermQuery ;		
"term" : { "user.id" :	es:field "user.id" ;		
"kimchy" }	es:value "kimchy" ;		
},	];		
"filter": {	es:filter [		
"term" : { "tags" :	a es:TermQuery ;		
"production" }	es:field "tags" ;		
},	es:value "production" ;		
"must_not" : {	] ;		
"range" : {	es:mustNot [		
"age" : { "gte" : 10,	a es:RangeQuery ;		
"lte" : 20 }	es:field "age" ;		
}	es:gte 10 ;		
},	es:lte 20 ;		
"should" : [	] ;		
{ "term" : { "tags" :	es:should (		
"env1" } },	<pre>[ a es:TermQuery ; es:field "tags" ;</pre>		
{ "term" : { "tags" :	es:value "env1" ]		
"deployed" } }	<pre>[ a es:TermQuery ; es:field "tags" ;</pre>		
],	es:value "deployed" ]		
"minimum_should_match" : 1,	) ;		
"boost" : 1.0	es:minimumShouldMatch 1 ;		
}	es:boost 1.0 ;		
}	]		

DSL	SPARQL
}	

The following example SERVICE clause with comments provides details about how the GDI es: query property can be mapped to DSL:

```
SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
  ?data a es:ElasticSource ;
    s:url "http://localhost:9200/" ;
# When the value of es:query is a simple literal,
# it is mapped to an Elastic query string query.
    es:query "literal"
# When the value of es:query is an RDF list,
# you can specify other query types,
# such as a match query:
   es:query [
      a es:MatchQuery ;
      es:field "title" | ?title ; # field can be a literal or bound variable
      es:query "moby dick" ;
   ];
# or a boolean query:
   es:query [
     a es:BoolQuery ;
      es:should ([
       a es:RangeQuery ;
       es:field ?amount ;
       es:gt 500 ;
       es:lt 1000 ;
     ] [
       a es:TermQuery ;
       es:field ?status ;
        es:value 'late' ;
     ]);
    1;
}
```

### **Filter Mapping**

Filtering can be performed inside the es:query list or you can add a FILTER clause to the query. For example, the table below shows the SPARQL snippet above expressed as a FILTER clause.

SPARQL Query	FILTER Clause
es:query [	FILTER(?user_id = "kimchy" &&
a es:BoolQuery ;	?tags = "production" &&
es:must [	!(?age >= 10 && ?age <=
a es:TermQuery ;	20) &&
es:field "user.id" ;	(?tags == "env1"    ?tags
es:value "kimchy" ;	== "deployed"))
];	
es:filter [	
a es:TermQuery ;	
es:field "tags" ;	
es:value "production" ;	
] ;	
es:mustNot [	
a es:RangeQuery ;	
es:field "age" ;	
es:gte 10 ;	
es:lte 20 ;	
] ;	
es:should (	
<pre>[ a es:TermQuery ; es:field "tags" ;</pre>	
es:value "env1" ]	
<pre>[ a es:TermQuery ; es:field "tags" ;</pre>	
es:value "deployed" ]	
) ;	
es:minimumShouldMatch 1 ;	
es:boost 1.0 ;	
]	

The table below shows each of the supported ElasticSource FILTER translations. Only expressions matching the list below will be translated by the GDI. If the expression is of the form value <= ?field, the inequality is flipped to ?field > value before translating.

es:query Expression	FILTER Clause Expression
es:query [ a es:BoolQuery ; es:mustNot expr ]	!expr
<pre>es:query [ a es:BoolQuery ; es:must (left right) ]</pre>	left && right
es:query [ a es:BoolQuery ; es:should (left right) ]	left    right
<pre>es:query [ a es:RangeQuery ; es:field ?field ; es:lt value ]</pre>	?field < value
es:query [ a es:RangeQuery ; es:field ?field ; es:lte value ]	?field <= value
es:query [ a es:TermQuery ; es:field ?field ; es:value value ]	?field = value
es:query [ a es:BoolQuery ; es:mustNot [ a es:TermQuery ; es:field ?field ; es:value value ] ]	?field != value
es:query [ a es:RangeQuery ; es:field ?field ; es:gte value ]	?field >= value
<pre>es:query [ a es:RangeQuery ; es:field ?field ; es:gt value ]</pre>	?field > value
<pre>es:query [ a es:QueryStringQuery ; es:field ?field ; es:query pattern ; es:defaultOperator "AND" ]</pre>	REGEX(?field, pattern, "q")
es:query [ a es:TermsQuery ; es:field ?field ; es:value value, ]	IN(?field, value,)
es:query [ a es:BoolQuery ; es:mustNot [ a es:TermsQuery ; es:field ?field ; es:value value, ] ]	NOT IN(?field, value,)

es:query Expression	FILTER Clause Expression
es:query [ a es:MatchQuery ; es:field ?field ; es:query value ; es:lenient true ]	CONTAINS (?field, value)
es:query [ a es:PrefixQuery ; es:field ?field ; es:value value ]	STRSTARTS (?field, value)
es:query [ a es:ExistsQuery ; es:field ?field ]	BOUND(?field)

#### **Query Examples**

- General Query
- Aggregations
- Highlighting

### **General Query**

The following example queries any Elasticsearch indexes that are loaded in the graphmart for which you run the query. No configuration is needed because Anzo manages the indexes that it loads and uses predictable naming conventions and aliases.

```
PREFIX docm: <http://cambridgesemantics.com/ontologies/2011/07/DocumentMetadata#>
PREFIX es: <http://elastic.co/search/>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX services: <http://cambridgesemantics.com/services/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
SELECT ?title
WHERE {
    SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
    ?data a es:ElasticSource;
    s:selector "unstructuredfile";
    es:query "string"; # input the text search string
    es:fields "fullText";
    ?title (xsd:string) .
```

}

}

### Aggregations

The following example query performs terms aggregations.

```
PREFIX es: <http://elastic.co/search/>
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
SELECT DISTINCT *
WHERE {
  SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
    ?data a es:ElasticSource;
    s:url "https://{{@es.hostname}}:{{@es.port}}/" ;
    s:username "{{@es.username}}" ;
    s:password "{{@es.password}}" ;
    es:index "templated consumption es" ;
    es:query "*ELM*" ;
    ?instance () ;
    es:aggregations [
      ?artifactTypes [
        a es:TermsAggregation ;
        es:field ?artifactType ;
        es:meta [
          ?label "artifactType" ;
        ];
        ?value () ;
        ?count () ;
      ];
      ?fileTypes [
        a es:TermsAggregation ;
        es:field ?fileType ;
        es:meta [
          ?label "fileType" ;
        ];
        ?value () ;
        ?count () ;
      ];
      ?managedBys [
        a es:TermsAggregation ;
        es:field ?managedBy ;
```

### Highlighting

The following example configures highlighting for fragments from the actor field.

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX es: <http://elastic.co/search/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
SELECT *
WHERE {
  SERVICE TOPDOWN <http://cambridgesemantics.com/services/DataToolkit>
  {
    ?data a es:ElasticSource ;
      es:url "http://localhost:9200/" ;
      es:index "films" ;
      es:html false ;
      es:query "Clint" ;
      es:field ?actor, ?director ;
      es:highlight [
       es:field ?actor ;
       es:type "plain" ;
        es:fragmentSize 200 ;
        es:numberOfFragments 10 ;
        es:preTags "<mark hit='true'>" ;
        es:postTags "</mark>" ;
      ];
      s:selector "film" ;
      ?actor (xsd:string) ;
      ?awards (xsd:string) ;
      ?director (xsd:string) ;
      ?image (xsd:string) ;
      ?length (xsd:long) ;
      ?popularity (xsd:long) ;
      ?subject (xsd:string) ;
```

```
?title (xsd:string) ;
?year (xsd:long) ;
?score () ;
?id () ;
?highlights [
    ?field () ;
    ?fragment () ;
    ] .
FILTER(?year = 1990 || ?length > 103)
FILTER(REGEX(?title, "Manhattan", "q") || REGEX(?subject, "Comedy", "q") || REGEX
(?subject, "Drama", "q"))
}
```

#### **Highlighting Elasticsearch Results**

By including the **highlight** property in ElasticSource GDI queries, you can configure the response to include highlights for search results. For general information about highlighting Elasticsearch responses, see <u>Highlighting</u> in the Elasticsearch documentation. Highlight property usage is described below.

- Highlight Syntax
- Highlight Examples

### **Highlight Syntax**

```
es:highlight [
    es:boundaryChars "string" ;
    es:boundaryMaxScan int ;
    es:boundaryScannerLocale "string" ;
    es:boundaryScannerType "string" ;
    es:field "string" ;
    es:forceSource boolean ;
    es:fragmentSize int ;
    es:fragmenter "string" ;
    es:highlightFilter boolean ;
    es:highlightQuery "string" | [rdf_list ] ;
    es:highlighterType "string" ;
    es:noMatchSize int ;
    es:numberOfFragments int ;
```

```
es:order "string" ;
  es:phraseLimit int ;
  es:postTags "string" ;
  es:preTags "string" ;
  es:requireFieldMatch boolean ;
];
```

Option	Туре	Description
boundaryChars	string	This property can be used to define the boundary characters to look for. Defaults to ., $!? \t n$ .
boundaryMaxScan	int	This property can be used to place a limit on the number of characters to scan when looking for boundary characters. Defaults to 20.
boundaryScannerLocale	string	This property defines the language tag (such as "en-US" or "fr-FR") to apply when searching for sentence and word boundaries.
boundaryScannerType	string	If highlighterType is unified or fvh, this property can be used to specify how to break the highlighted fragments. This property is ignored when the highlighter type is plain. The list below describes the valid values:
		<ul> <li>chars: Valid when the highlighter type is fast vector highlighter (fvh) (es:highlighterType "fvh").</li> <li>Specifies that the highlighting boundaries are the characters specified by boundaryChars. The boundaryMaxScan value controls how far to scan for boundary characters. This is the default value for fvh.</li> </ul>

Option	Туре	Description
		<ul> <li>sentence: This is the default value for the unified highlighter. It configures highlighted fragments to break at the next sentence boundary. You can specify the locale to use with boundaryScannerLocale. When used with the unified highlighter, the sentence scanner splits sentences bigger than fragmentSize at the first word boundary next to fragmentSize. You can set fragmentSize to 0 to avoid splitting sentences.</li> <li>word: Configures highlighted fragments to break at the next word boundary. You can specify the locale to use with boundaryScannerLocale.</li> </ul>
field	string or variable	<pre>This property specifies the field to retrieve highlights for. It can include a ?variable (which the GDI maps to the full path of the field in the Elasticsearch document), a field name, or a field name pattern. For example: es:highlight [    es:field ?actor ;    es:field ?actor ;    es:field "film.actor" ;    es:field "film.*" ;    es:field "*" ;</pre>
forceSource	boolean	This property controls whether to highlight based

Option	Туре	Description
		on the source even if the field is stored separately. Defaults to false.
fragmentSize	int	This property specifies the number of characters to include in highlighted fragments. Defaults to 100.
fragmenter	string	<ul> <li>If highlighterType is plain, this property can be used to specify how to break up text in highlight snippets. The list below describes the valid values:</li> <li>simple: Breaks text into fragments that are the same size (as specified by fragmentSize).</li> <li>span: The default value. Breaks text into fragments that are the same size but tries to avoid breaking text between highlighted terms.</li> </ul>
highlightFilter	boolean	This property controls whether to highlight filter results.
highlightQuery	string or object	This property specifies the highlight query. The value can be a string or a query object that maps to the Elasticsearch query DSL.
highlighterType	string	This property defines the type of highlighter to use, "plain", "unified", or "fvh".
noMatchSize	int	This property specifies the number of characters to return from the beginning of the field if there are no matching fragments to highlight. Defaults to 0

Option	Туре	Description
		(nothing is returned).
numberOfFragments	int	This property can be used to set the maximum number of fragments to generate. If this property is set to 0, no fragments are returned. Instead, the entire field contents are highlighted and returned, which can be useful if you want to highlight short text (such as a title or address) for which fragmentation is not required. Defaults to 5. If the number of fragments is 0, fragmentSize is ignored.
order	string	This property can be included to sort highlighted fragments by score. When es:order "score", the most relevant fragments are output first. Defaults to "none"; fragments are output in the order they appear in the field.
phraseLimit	int	If highlighterType is fvh, this property can be used to limit the number of matching phrases to consider. Limiting the number of phrases prevents the fvh highlighter from analyzing too many phrases and consuming too much memory. Defaults to 256.
postTags	string	This property is used in conjunction with preTags to define the HTML tags to use for the highlighted elements. This property defines the closing tag to use after the highlighted text. Defaults to .
preTags	string	This property is used in conjunction with postTags to define the HTML tags to use for the highlighted elements. This property defines the opening tag to use before the highlighted text. Defaults to <em>.</em>

Option	Туре	Description
requireFieldMatch	boolean	This property controls whether to highlight only the fields that contain a query match. Defaults to true. If false, all fields are highlighted.

## **Highlight Examples**

The following example configures highlighting for fragments from the actor field.

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX es: <http://elastic.co/search/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
SELECT *
WHERE {
  SERVICE TOPDOWN <http://cambridgesemantics.com/services/DataToolkit>
  {
    ?data a es:ElasticSource ;
      es:url "http://localhost:9200/" ;
      es:index "films" ;
      es:html false ;
      es:query "Clint" ;
      es:field ?actor, ?director ;
      es:highlight [
        es:field ?actor ;
        es:type "plain" ;
        es:fragmentSize 200 ;
        es:numberOfFragments 10 ;
        es:preTags "<mark hit='true'>" ;
        es:postTags "</mark>" ;
      ];
      s:selector "film" ;
      ?actor (xsd:string) ;
      ?awards (xsd:string) ;
      ?director (xsd:string) ;
      ?image (xsd:string) ;
      ?length (xsd:long) ;
      ?popularity (xsd:long) ;
      ?subject (xsd:string) ;
      ?title (xsd:string) ;
```

```
?year (xsd:long) ;
?score () ;
?id () ;
?highlights [
?field () ;
?fragment () ;
] .
FILTER(?year = 1990 || ?length > 103)
FILTER(REGEX(?title, "Manhattan", "q") || REGEX(?subject, "Comedy", "q") || REGEX
(?subject, "Drama", "q"))
}
```

### **Querying a File Source**

The Graph Data Interface (GDI) uses the <u>Apache Commons VFS</u> library to work with file systems. Many of the properties specified in queries against file sources reflect the requirements of the VFS library for that source. The topics in this section provide guidance on writing GDI queries for each of the supported file types.

- Querying CSV and TSV Files
- Querying JSON and NDJSON Files
- Querying XML Files
- Querying Parquet and SAS Files
- File Source Format Options
- File Storage Connection Options

### Querying CSV and TSV Files

This topic provides details about the structure to use when writing GDI queries to read or ingest data from CSV or TSV files. It also includes example queries that may be useful as a starting point for writing your own GDI queries.

- Query Syntax
- Query Examples

# **Query Syntax**

The following query syntax shows the structure of a GDI query for CSV and TSV sources. The clauses, patterns, and placeholders that are links are described below.

```
# PREFIX Clause
PREFIX s:
              <http://cambridgesemantics.com/ontologies/DataToolkit#>
             <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdf:
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
            <http://purl.org/dc/elements/1.1/>
PREFIX dc:
# Result Clause
{
   [ GRAPH ${targetGraph} { ]
  triple_patterns
 [ } ]
}
[ ${usingSources} ]
WHERE
{
   # SERVICE Clause: Include the following service call when reading or inserting data.
    SERVICE [ TOPDOWN ] <http://cambridgesemantics.com/services/DataToolkit>
   # View SERVICE Clause: Or use the service call below when constructing a view.
    SERVICE <http://cambridgesemantics.com/services/DataToolkitView>(${targetGraph})
    {
      ?data a s:FileSource ;
        s:url "string" ;
        [ s:options [ file storage connection options ] ; ]
        [ s:pattern "string" ; ]
        [ s:maxDepth int ; ]
        [ s:format [ source_format_options ; ] ; ]
        [ s:mimetype "string" ; ]
        [ s:username "string" ; ]
```

[ s:password "string" ; ]

```
[ s:timeout int ; ]
   [ s:batching boolean | int ; ]
   [ s:paging [ pagination options ; ]
   [ s:concurrency int | [ list_of_properties ] ; ]
   [ s:rate int | "string" ; ]
   [ s:locale "string" ; ]
   [ s:sampling int ; ]
   [ s:selector "string" | [ list ] ; ]
   [ s:model "string" ; ]
   [ s:key ("string") ; ]
   [ s:reference [ s:model "string" ; s:using ("string") ]
   [ s:formats [ datatype formatting options ] ; ]
   [ s:normalize boolean | [ normalization rules ] ; ]
   [ s:count ?variable ; ]
   [ s:offset int ; ]
   [ s:limit int ; ]
   # Mapping variables
   ?mapping variable ( [ "binding" ] [ datatype ] [ "datetime format" ] ) ;
   ....;
# Additional clauses such as BIND, VALUES, FILTER
```

### Note

}

}

#### For readability, the parameters below exclude the base URI

<http://cambridgesemantics.com/ontologies/DataToolkit#> as well as the s: prefix. As shown in the examples, however, the s: prefix or full property URI does need to be included in queries.

Option	Туре	Description
PREFIX Clause	N/A	The PREFIX clause declares the standard and custom prefixes for GDI service queries. Generally, queries include the prefixes from the query template (or a subset of them) plus any data-specific declarations.
Result Clause	N/A	The result clause defines the type of SPARQL query to run and the

Option	Туре	Description
		set of results to return, i.e., whether you want to read (SELECT or CONSTRUCT) from the source or ingest the data into Anzo (INSERT).
GRAPH \${targetGraph}	N/A	Include the GRAPH keyword and target graph parameter \${targetGraph} when writing an INSERT query to ingest data into a graphmart. Anzo automatically populates the query with the appropriate target URIs when the query runs.
\${usingSource s}	N/A	Include the source graph parameter $\{usingSources\}\$ when writing a "topdown" query that passes values from the data that is in the graphmart to the data source. Anzo automatically populates the query with the appropriate FROM clauses when the query runs. When passing literal values to the remote source, you do not need to include the source graph parameter. The SERVICE Clause description below includes more information about passing input to data sources.
SERVICE Clause	N/A	<pre>Include the SERVICE call SERVICE [ TOPDOWN ] <http: cambridgesemantics.com="" datatool="" kit="" services=""> to invoke the GDI service when you are running a SELECT, INSERT, or CONSTRUCT query that is not creating a view. When writing a CONSTRUCT query in a View Step, use the DataToolkitView service call, as described below in View SERVICE Clause.</http:></pre>
		Include the optional TOPDOWN keyword when you want to pass input values from the graphmart to the data source. When you include TOPDOWN in the service call, it indicates that the rest of the query produces values to send to the source. In this case, the GDI makes repeated calls to pass in each of the specified values and retrieve the data that is based on those values.

Option	Туре	Description
View SERVICE Clause	N/A	When writing a CONSTRUCT query that creates a view of the data (usually in a View Step), include the following SERVICE call: SERVICE <http: cambridgesemantics.com="" datatool<br="" services="">kitView&gt;(\${targetGraph}). Using the DataToolkitView call optimizes query execution because it tells the GDI to inspect the query and determine which filters to push to the data source. It also limits the result set and retrieves only the data that is needed, i.e., the source data is fully mapped but all of the mapped data is not necessarily returned.</http:>
url	string	This property specifies the file system location of the source file or directory of files. When specifying a directory (such as s:url "/opt/shared-files/loads/"), the GDI loads all of the file formats it recognizes. To specify a directory but limit the number or type of files that are read, you can include the pattern and/or maxDepth properties.
options	RDF list	If additional connection information needs to be provided to access the file storage system, include the options property to list any storage-specific connection parameters. See File Storage Connection Options for information about the supported properties for each storage type.
pattern	string	This property can be used to specify a wildcard pattern for matching file names. For example, s:pattern "common_ prefix*.csv". You can include one s:pattern property per FileSource. The GDI supports Unix file globbing syntax outside of parentheses. Within parentheses, full Java regular expression language is supported. For example, including s:pattern "data/**/customer_*.csv" tells the GDI to load all files that match the pattern "customer_*.csv" from any number of

Option	Туре	Description
		<pre>subdirectories under the data directory. Similarly s:pattern " (\d+) /transaction_*.csv" tells the GDI to load all files that match the pattern "transaction_*.csv" in all subdirectories.</pre>
maxDepth	int	This property can be used to limit the directory traversal depth. By default, when $s:url$ specifies a directory (and $as:pattern$ that limits that traversal depth is not specified), all subdirectories are processed. To process only the files in the top level directory, set maxDepth to 0 ( $s:maxDepth$ 0). To process the files in the top level directory plus the first-level subdirectories, set maxDepth to 1 ( $s:maxDepth$ 1), and so on.
format	RDF list	You can include the format property to add parameters that describe the source files. See File Source Format Options for details about the supported parameters.
mimetype	string	This property can be included to specify the MIME type of the data. If you are querying TSV files that do not have a .tsv file extension, include the mimetype property with a value of text/tsv (s:mimetype "text/tsv").
username	string	If authentication is required to access the source, include this property to specify the user name.
password	string	This property lists the password for the given username.
timeout	int	This property can be used to specify the timeout (in milliseconds) to use for requests against the source. For example, s:timeout 5000 configures a 5 second timeout.
batching	boolean or int	This property can be used to disable batching, or it can be used to change the default the batch size. By default, batching is set to

Option	Туре	Description
		5000 (s:batching 5000). To disable batching, you can include s:batching false in the query. Typically users do not change the batching size. However, it can be useful to control the batch size when performing updates. To configure the size, include s:batching int in the query. For example, s:batching 3000.
paging	RDF list	This property can be used to configure paging so that the GDI can access large amounts of data across a number of smaller requests. For details about the paging property, see Paginating Requests.
concurrency	int or RDF list	<pre>This property can be included to configure the maximum level of concurrency for the query. The value can be an integer, such as s:concurrency 8. If the value is an integer, it configures a maximum limit on the number of slices that can execute the query. For finer-grained control over the number of nodes and slices to use, concurrency can also be included as an object with limit, nodes, and/or executorsPerNode properties. For example, the following object configures a concurrency model that allows a maximum of 24 executors distributed across 4 nodes with 8 executors per node: s:concurrency [ s:limit 24 ; s:nodes 4 ; s:executorsPerNode 8 ; ] ;</pre>
rate	int or string	This property can be included to control the frequency with which a request is sent to the source. The limit applies to the number of requests a single slice can make. If you specify an integer for the

Option	Туре	Description
		rate, then the value is treated as the maximum number of requests to issue per minute. If you specify a string, you have more flexibility in configuring the rate. The sample values below show the types of values that are supported:
		<pre>s:rate "90/minute" ; s:rate "90 per minute" ; s:rate "200000 every week" ; s:rate "10000 every 6 hours" ;</pre>
		To enforce the rate limit, the GDI introduces a sleep between requests that is equal to the rate delay. The more executing slices, the longer the rate delay needs to be to enforce the limit in aggregate.
		Given the example of s:rate "90/minute", the GDI would optimize the concurrency and only use 1 slice for execution with a rate delay of 666ms between requests. If s:rate "240/minute", the GDI would use 3 executors with a rate delay of 750ms between requests.
locale	string	This property can be used to specify the locale to use when parsing locale-dependent data such as numbers, dates, and times.
sampling	int	This property can be used to configure the number of records in the source to examine for data type inferencing.
selector	string or RDF list	This property can be used as a binding component to identify the path to the source objects. For example, s:selector "Sales.SalesOrderHeader" targets the SalesOrderHeader table in the Sales schema. For more information about binding components and the selector property, see Using Binding Trees and Selector Paths.

Option	Туре	Description
model	string	This property defines the class (or table) name for the type of data that is generated from the specified data source. For example, s:model "employees". Model is optional when querying a single source. If your query targets multiple sources, however, and you want to define resource templates (primary keys) and object properties (foreign keys), you must specify the model value for each source.
key	string	This property can be used to define the primary key column for the source file or table. This column is leveraged in a resource template for the instances that are created from the source. For example, s:key ("EMPLOYEE_ID"). For more information about key, see Data Linking Options.
reference	RDF list	This property can be used to specify a foreign key column. The reference property is an RDF list that includes the model property to list the target table and a using property that defines the foreign key column. For more information about reference, see Data Linking Options.
formats	RDF list	To give users control over the data types that are used when coercing strings to other types, this property can be included in GDI queries to define the desired types. In addition, it can be used to describe the formats of date and time values in the source to ensure that they are recognized and parsed to the appropriate date, time, and/or dateTime values. For details about the formats property, see Data Type Formatting Options.
normalize	RDF list	To give users control over the labels and URIs that are generated, the GDI offers several options for normalizing the model and/or the fields that are created from the specified data source(s). For details about the normalize property, see Model Normalization

Option	Туре	Description
		Options.
count	variable	If you want to turn the query into a COUNT query, you can include this property with a <code>?variable</code> to perform a count. For example, <code>s:count ?count</code> .
offset	int	This property can be used to offset the data that is returned by a number of rows.
limit	int	You can include this property to limit the number of results that are returned. s:limit maps to the SPARQL LIMIT clause.
mapping_ vari variable	variable	The mapping variables, in <code>?mapping_variable</code> (["binding"] [datatype] ["datetime_format"]) format, define the triple patterns to output. When the specified <code>?variable</code> matches the source column name, the GDI uses the variable as the source data selector. If you specify an alternate variable name, a binding needs to be specified to map the new variable to the source. You also have the option to transform the data using the datatype and datetime_format options.
		Note The parentheses around the binding, data type, and format specifications are not required but are included in this document for readability.
binding	string	The binding is a literal value that binds a ?mapping_variable to a source column. If you specify a ?variable that matches the source column name, then that variable name is the data selector and it is not necessary to specify a binding. If you specify an alternate variable name or there is a hierarchical path to the source column,

Option	Туре	Description
		<pre>then the binding is needed to map the new variable to that source column. For example for CSV, the following pattern simply binds the source column AIRLINE to the lowercase variable ?airline: ?airline ("AIRLINE").</pre>
		Note For FileSource, periods (.), forward slashes (/), and brackets ([]) are parsed as path notation. Therefore, if a source column name includes any of those characters they must be escaped in the binding. Use two backslashes (\\) as an escape character. For example, if a column name is <b>average/day</b> , the variable and binding pattern could be written as ?averagePerDay ("average\\/day").
datatype	URI	The datatype is the data type to convert the column to. If you do not specify a data type, the GDI infers the type. The GDI supports the following types: xsd:int, xsd:long, xsd:float, xsd:double, xsd:boolean, xsd:time, xsd:dateTime, xsd:date, xsd:duration, xsd:timeDuration, xsd:yearMonthDuration, xsd:gMonthDay, xsd:gMonth, xsd:gYearMonth, xsd:anyURI
datetime_ format	string	This option is used to specify the format to use for date and time data types. The GDI supports Java date and time formats. Specify days as "d," months as "M," and years as "y." For the time, specify "H" for hours, "m" for minutes, and "s" for seconds. For example,

Option	Туре	Description
		"yyyyMMdd HH:mm:ss" or "ddMMMyy" to display date values such as "01JAN19."
		Note The GDI's default base year is 2000. If the source data has years with only two digits, such as 02-04-99, the GDI prepends 20 to the digits. The value 02-04-99 is parsed to 02-04-2099. To specify an alternate base year to use for two-digit values, you can include the notation ^nnnn (e.g., ^1900) in the format value. For example, to set the base year to 1900 instead of 2000, use a format value such as xsd:date "dd-MMM-yy^1900" or xsd:date "dd-MMM-yy^1990". When one of those values is specified, 02-04-99 is parsed to 02-04-1999.

## **Query Examples**

The example below ingests a directory of CSV files into a Graphmart. The pattern property (s:pattern "post\_[0-9]\*\_[0-9]\*.csv") is used to narrow down the set of files to load. Since the files use a pipe (|) as the delimiter rather than a comma (,), the delimiter property is also included to specify the | character.

```
snvoc:id ?id ;
      snvoc:imageFile ?imageFile ;
      snvoc:locationIP ?locationIP ;
      snvoc:browserUsed ?browserUsed ;
      snvoc:language ?language ;
      snvoc:content ?content ;
      snvoc:length ?length ;
  }
}
WHERE {
  SERVICE <http://cambridgesemantics.com/services/DataToolkit>
  {
      ?data a s:FileSource ;
      s:url "/opt/shared-files/data/csv/post 6 0/" ;
      s:pattern "post [0-9]* [0-9]*.csv" ;
      s:format [ s:delimiter "|" ] ;
      ?creationDate (xsd:dateTime) ;
      ?id (xsd:string) ;
      ?imageFile (xsd:string) ;
      ?locationIP (xsd:string) ;
      ?browserUsed (xsd:string) ;
      ?language (xsd:string) ;
      ?content (xsd:string) ;
      ?length(xsd:string) .
  BIND(IRI("http://www.ldbc.eu/ldbc socialnet/1.0/data/Post/{{?id}}") as ?postIRI)
  }
}
```

## The example below is similar to the query above but it specifies the formats to use for the

xsd:date values.

```
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
PREFIX dc: <http://purl.org/dc/elements/1.1/>
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX kd: <http://cambridgesemantics.com/ont/autogen/Rh/Kaggle_Diabetes#>
INSERT {
    GRAPH ${targetGraph}
```

```
?URI a kd:Diagnosis ;
      kd:Diagnosis DiagnosisGuid ?diagnosis guid ;
      kd:Diagnosis PatientGuid ?patient guid ;
      kd:Diagnosis ICD9Code ?icd9Code ;
      kd:Diagnosis DiagnosisDescription ?diagnosisDescription ;
      kd:Diagnosis StartDate ?cus start date ;
      kd:Diagnosis EndDate ?Date End ;
      kd:Diagnosis Acute ?acute ;
      kd:Diagnosis UserGuid ?UserGuid .
  }
}
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit>
  {
      ?csv a s:FileSource ;
      s:url "/opt/shared-files/source data/kaggle diabetes/" ;
      s:pattern "Diagnosis.csv" ;
      s:format [ s:delimiter "," ] ;
      ?diagnosis guid ("DiagnosisGuid" xsd:string) ;
      ?patient guid ("PatientGuid" xsd:string) ;
      ?icd9Code ("ICD9Code" xsd:string) ;
      ?diagnosisDescription ("DiagnosisDescription" xsd:string) ;
      ?acute ("Acute" xsd:int ) ;
      ?UserGuid ("UserGuid" xsd:string) ;
      ?cus_start_date ("CUSTOMER_START_DATE" xsd:date "yyyy-MM-dd") ;
      ?Date_End ("Date End" xsd:date "MM/dd/yy") .
  }
  BIND (IRI (CONCAT ("urn://cambridgesemantics.com/kaggle diabetes/patient/",ENCODE FOR
URI(?diagnosis guid))) as ?URI)
```

### **Querying JSON and NDJSON Files**

This topic provides details about the structure to use when writing GDI queries to read or ingest data from JSON or NDJSON files. It also includes example queries that may be useful as a starting point for writing your own GDI queries.

- Query Syntax
- Hierarchical Bindings and Arrays

- Capturing Property Keys
- Query Examples

## **Query Syntax**

The following query syntax shows the structure of a GDI query for JSON sources. The clauses, patterns, and placeholders that are links are described below.

```
# PREFIX Clause
            <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX s:
            <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdf:
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
            <http://purl.org/dc/elements/1.1/>
PREFIX dc:
# Result Clause
{
   [ GRAPH ${targetGraph} { ]
   triple patterns
 [ } ]
}
[ ${usingSources} ]
WHERE
{
   # SERVICE Clause: Include the following service call when reading or inserting data.
    SERVICE [ TOPDOWN ] <http://cambridgesemantics.com/services/DataToolkit>
   # View SERVICE Clause: Or use the service call below when constructing a view.
    SERVICE <http://cambridgesemantics.com/services/DataToolkitView>(${targetGraph})
    {
      ?data a s:FileSource ;
        s:url "string" ;
        [ s:options [ file storage connection options ] ; ]
        [ s:pattern "string" ; ]
        [ s:maxDepth int ; ]
```

```
[ s:format [ source_format_options ; ] ; ]
   [ s:mimetype "string" ; ]
   [ s:username "string" ; ]
   [ s:password "string" ; ]
   [ s:timeout int ; ]
   [ s:batching boolean | int ; ]
   [ s:paging [ pagination options ; ]
   [ s:concurrency int | [ list_of_properties ] ; ]
   [ s:rate int | "string" ; ]
   [ s:locale "string" ; ]
   [ s:sampling int ; ]
   [ s:selector "string" | [ list ] ; ]
   [ s:model "string" ; ]
   [ s:key ("string") ; ]
   [ s:reference [ s:model "string" ; s:using ("string") ]
   [ s:formats [ datatype formatting options ] ; ]
   [ s:normalize boolean | [ normalization rules ] ; ]
   [ s:count ?variable ; ]
   [ s:offset int ; ]
   [ s:limit int ; ]
   # Mapping variables and hierarchical bindings
   ?mapping_variable ( [ "binding" ] [ datatype ] [ "datetime_format" ] ) ;
   ...;
# Additional clauses such as BIND, VALUES, FILTER
```

### Note

}

#### For readability, the parameters below exclude the base URI

<http://cambridgesemantics.com/ontologies/DataToolkit#> as well as the s: prefix. As shown in the examples, however, the s: prefix or full property URI does need to be included in queries.

Option	Туре	Description
PREFIX Clause	N/A	The PREFIX clause declares the standard and custom prefixes for GDI service queries. Generally, queries include the prefixes from the query template (or a subset of them) plus any data-specific

Option	Туре	Description
		declarations.
Result Clause	N/A	The result clause defines the type of SPARQL query to run and the set of results to return, i.e., whether you want to read (SELECT or CONSTRUCT) from the source or ingest the data into Anzo (INSERT).
GRAPH \${targetGraph}	N/A	Include the GRAPH keyword and target graph parameter \${targetGraph} when writing an INSERT query to ingest data into a graphmart. Anzo automatically populates the query with the appropriate target URIs when the query runs.
\${usingSource s}	N/A	Include the source graph parameter $\{usingSources\}$ when writing a "topdown" query that passes values from the data that is in the graphmart to the data source. Anzo automatically populates the query with the appropriate FROM clauses when the query runs. When passing literal values to the remote source, you do not need to include the source graph parameter. The SERVICE Clause description below includes more information about passing input to data sources.
SERVICE Clause	N/A	<pre>Include the SERVICE call SERVICE [ TOPDOWN ] <http: cambridgesemantics.com="" datatool="" kit="" services=""> to invoke the GDI service when you are running a SELECT, INSERT, or CONSTRUCT query that is not creating a view. When writing a CONSTRUCT query in a View Step, use the DataToolkitView service call, as described below in View SERVICE Clause. Include the optional TOPDOWN keyword when you want to pass</http:></pre>
		input values from the graphmart to the data source. When you include TOPDOWN in the service call, it indicates that the rest of

Option	Туре	Description
		the query produces values to send to the source. In this case, the GDI makes repeated calls to pass in each of the specified values and retrieve the data that is based on those values.
View SERVICE Clause	N/A	When writing a CONSTRUCT query that creates a view of the data (usually in a View Step), include the following SERVICE call: SERVICE <http: cambridgesemantics.com="" datatool<br="" services="">kitView&gt;(\${targetGraph}). Using the DataToolkitView call optimizes query execution because it tells the GDI to inspect the query and determine which filters to push to the data source. It also limits the result set and retrieves only the data that is needed, i.e., the source data is fully mapped but all of the mapped data is not necessarily returned.</http:>
url	string	This property specifies the file system location of the source file or directory of files. When specifying a directory (such as s:url "/opt/shared-files/loads/"), the GDI loads all of the file formats it recognizes. To specify a directory but limit the number or type of files that are read, you can include the pattern and/or maxDepth properties.
options	RDF list	If additional connection information needs to be provided to access the file storage system, include the options property to list any storage-specific connection parameters. See File Storage Connection Options for information about the supported properties for each storage type.
pattern	string	This property is used to specify a wildcard pattern for matching file names. For example, s:pattern "common_prefix*.json". You can include one s:pattern property per FileSource. The GDI supports Unix file globbing syntax outside of parentheses.

Option	Туре	Description
		Within parentheses, full Java regular expression language is supported. For example, including s:pattern "data/**/customer_*.json" tells the GDI to load all files that match the pattern "customer_*.json" from any number of subdirectories under the data directory. Similarly s:pattern " (\d+)/transaction_*.json" tells the GDI to load all files that match the pattern "transaction_*.json" in all subdirectories.
maxDepth	int	This property can be used to limit the directory traversal depth. By default, when s:url specifies a directory (and a s:pattern that limits that traversal depth is not specified), all subdirectories are processed. To process only the files in the top level directory, set maxDepth to 0 (s:maxDepth 0). To process the files in the top level directory plus the first-level subdirectories, set maxDepth to 1 (s:maxDepth 1), and so on.
format	RDF list	You can include the format property to add parameters that describe the source files. See File Source Format Options for details about the supported parameters.
mimetype	string	This property can be included to specify the MIME type of the data.
username	string	If authentication is required to access the source, include this property to specify the user name.
password	string	This property lists the password for the given username.
timeout	int	This property can be used to specify the timeout (in milliseconds) to use for requests against the source. For example, s:timeout 5000 configures a 5 second timeout.
batching	boolean	This property can be used to disable batching, or it can be used to

Option	Туре	Description
	or int	change the default the batch size. By default, batching is set to 5000 (s:batching 5000). To disable batching, you can include s:batching false in the query. Typically users do not change the batching size. However, it can be useful to control the batch size when performing updates. To configure the size, include s:batching int in the query. For example, s:batching 3000.
paging	RDF list	This property can be used to configure paging so that the GDI can access large amounts of data across a number of smaller requests. For details about the paging property, see Paginating Requests.
concurrency	int or RDF list	<pre>This property can be included to configure the maximum level of concurrency for the query. The value can be an integer, such as s:concurrency 8. If the value is an integer, it configures a maximum limit on the number of slices that can execute the query. For finer-grained control over the number of nodes and slices to use, concurrency can also be included as an object with limit, nodes, and/or executorsPerNode properties. For example, the following object configures a concurrency model that allows a maximum of 24 executors distributed across 4 nodes with 8 executors per node: s:concurrency [ s:limit 24 ; s:nodes 4 ; s:executorsPerNode 8 ; ] ;</pre>
rate	int or string	This property can be included to control the frequency with which a request is sent to the source. The limit applies to the number of
Option	Туре	Description
----------	-----------------------	--
		requests a single slice can make. If you specify an integer for the rate, then the value is treated as the maximum number of requests to issue per minute. If you specify a string, you have more flexibility in configuring the rate. The sample values below show the types of values that are supported:
		s:rate "90/minute" ;
		<pre>s:rate "90 per minute" ; s:rate "200000 every week" ;</pre>
		s:rate "10000 every 6 hours" ;
		To enforce the rate limit, the GDI introduces a sleep between requests that is equal to the rate delay. The more executing slices, the longer the rate delay needs to be to enforce the limit in aggregate.
		Given the example of s:rate "90/minute", the GDI would optimize the concurrency and only use 1 slice for execution with a rate delay of 666ms between requests. If s:rate "240/minute", the GDI would use 3 executors with a rate delay of 750ms between requests.
locale	string	This property can be used to specify the locale to use when parsing locale-dependent data such as numbers, dates, and times.
sampling	int	This property can be used to configure the number of records in the source to examine for data type inferencing.
selector	string or RDF list	This property can be used for JSON path extraction to traverse nested structures and target specific data. For example, s:selector "projects" targets the projects class of data. To express a hierarchy, use dot notation. For example, s:selector "region.state.city" navigates a hierarchy to

Option	Туре	Description
		target $city$ data. For more information about binding components and the selector property, see Using Binding Trees and Selector Paths.
model	string	This property defines the class (or table) name for the type of data that is generated from the specified data source. For example, s:model "employees". Model is optional when querying a single source. If your query targets multiple sources, however, and you want to define resource templates (primary keys) and object properties (foreign keys), you must specify the model value for each source.
key	string	This property can be used to define the primary key column for the source file or table. This column is leveraged in a resource template for the instances that are created from the source. For example, s:key ("EMPLOYEE_ID"). For more information about key, see Data Linking Options.
reference	RDF list	This property can be used to specify a foreign key column. The reference property is an RDF list that includes the model property to list the target table and a using property that defines the foreign key column. For more information about reference, see Data Linking Options.
formats	RDF list	To give users control over the data types that are used when coercing strings to other types, this property can be included in GDI queries to define the desired types. In addition, it can be used to describe the formats of date and time values in the source to ensure that they are recognized and parsed to the appropriate date, time, and/or dateTime values. For details about the formats property, see Data Type Formatting Options.

Option	Туре	Description
normalize	RDF list	To give users control over the labels and URIs that are generated, the GDI offers several options for normalizing the model and/or the fields that are created from the specified data source(s). For details about the normalize property, see Model Normalization Options.
count	variable	If you want to turn the query into a COUNT query, you can include this property with a <code>?variable</code> to perform a count. For example, <code>s:count</code> ?count.
offset	int	This property can be used to offset the data that is returned by a number of rows.
limit	int	You can include this property to limit the number of results that are returned. s:limit maps to the SPARQL LIMIT clause.
mapping_ variable	variable	The mapping variables, in <code>?mapping_variable</code> (["binding"] [datatype] ["datetime_format"]) format, define the triple patterns to output. When the specified <code>?variable</code> matches the source column name, the GDI uses the variable as the source data selector. If you specify an alternate variable name, a binding needs to be specified to map the new variable to the source. You also have the option to transform the data using the datatype and datetime_format options.
		<b>Tip</b> See Hierarchical Bindings and Arrays below for more information about configuring mapping variables and unpacking JSON files with nested objects and arrays.
binding	string	The binding is a literal value that binds a ?mapping_variable to a

Option	Туре	Description
		<pre>source column. If you specify a ?variable that matches the source column name, then that variable name is the data selector and it is not necessary to specify a binding. If you specify an alternate variable name or there is a hierarchical path to the source column, then the binding is needed to map the new variable to that source column.</pre> For example for CSV, the following pattern simply binds the source column AIRLINE to the lowercase variable ?airline: ?airline ("AIRLINE").
		Note For FileSource, periods (.), forward slashes (/), and brackets ([]) are parsed as path notation. Therefore, if a source column name includes any of those characters they must be escaped in the binding. Use two backslashes (\\) as an escape character. For example, if a column name is <b>average/day</b> , the variable and binding pattern could be written as ?averagePerDay ("average\\/day").
datatype	URI	The datatype is the data type to convert the column to. If you do not specify a data type, the GDI infers the type. The GDI supports the following types: xsd:int, xsd:long, xsd:float, xsd:double, xsd:boolean, xsd:time, xsd:dateTime, xsd:date, xsd:duration, xsd:dayTimeDuration, xsd:yearMonthDuration, xsd:gMonthDay, xsd:gMonth, xsd:gYearMonth, xsd:anyURI

Option	Туре	Description
datetime_ format	string	This option is used to specify the format to use for date and time data types. The GDI supports Java date and time formats. Specify days as "d," months as "M," and years as "y." For the time, specify "H" for hours, "m" for minutes, and "s" for seconds. For example, "yyyyMMdd HH:mm:ss" or "ddMMMyy" to display date values such as "01JAN19."
		Note The GDI's default base year is 2000. If the source data has years with only two digits, such as 02-04-99, the GDI prepends 20 to the digits. The value 02-04-99 is parsed to 02-04-2099. To specify an alternate base year to use for two-digit values, you can include the notation ^nnnn (e.g., ^1900) in the format value. For example, to set the base year to 1900 instead of 2000, use a format value such as xsd:date "dd-MMM-yy^1900" or xsd:date "dd-MMM-yy^1990". When one of those values is specified, 02-04-99 is parsed to 02-04-1999.

## **Hierarchical Bindings and Arrays**

When configuring the mapping variables in a query, the GDI provides syntax for unpacking JSON files with nested objects and arrays. One way to express hierarchies in queries is to use brackets ([]) to group objects into binding trees. For example, the WHERE clause snippet below organizes mapping variable objects into an hourly/data hierarchy by nesting the ?data patterns inside the ?hourly [] tree:

```
WHERE
{
   SERVICE <http://cambridgesemantics.com/services/DataToolkit>
   {
```

```
?data a s:FileSource;
        s:url "/mnt/data/json/weather.json" ;
        ?latitude (xsd:double) ;
        ?longitude (xsd:double) ;
        ?timezone (xsd:string) ;
        ?hourly
Γ
          ?data
Γ
            ?time (xsd:long) ;
            ?rainProbability ("precipProbability" xsd:double) ;
            ?temperature (xsd:double) ;
            ?feelsLike ("apparentTemperature" xsd:double) ;
            ?windSpeed (xsd:double) ;
        ];
   ] .
 }
}
```

When constructing object binding trees, if you choose to introduce the hierarchy with a variable name that is not an exact match to the source label, include a **selector** property to list the value from the source. For example, in the WHERE clause snippet below, s:selector is included to select eventHeader in the source as ?event in the query and statLocation as ?location.

```
WHERE
{
   SERVICE <http://cambridgesemantics.com/services/DataToolkit>
  {
      ?data a s:FileSource ;
      s:url "/mnt/data/json/part 1.json" ;
      ?event
[
         s:selector "eventHeader" ;
           ?eventId (xsd:string) ;
           ?eventName (xsd:string) ;
           ?eventVersion (xsd:string) ;
           ?eventTime (xsd:dateTime) ;
      1;
      ?location
Γ
         s:selector "statLocation" ;
           ?locationId (xsd:string) ;
```

```
?lineNo (xsd:int) ;
    ?statNo (xsd:int) ;
    ?statId (xsd:int) ;
    ].
}
```

As an alternative to grouping objects in binding trees, the **selector** property also supports using dot notation to specify paths. For example, the WHERE clause snippet below rewrites the first example query to express the same <code>hourly/data</code> hierarchy as a path in the <code>s:selector</code> value:

```
WHERE
{
  SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
      ?data a s:FileSource;
        s:url "/mnt/data/json/weather.json" ;
        ?latitude (xsd:double) ;
        ?longitude (xsd:double) ;
        ?timezone (xsd:string) ;
        s:selector: "hourly.data" ;
        ?time (xsd:long) ;
        ?rainProbability ("precipProbability" xsd:double) ;
        ?temperature (xsd:double) ;
        ?feelsLike ("apparentTemperature" xsd:double) ;
        ?windSpeed (xsd:double) .
  }
}
```

In addition to object binding trees and selectors, the GDI offers additional syntax for reading or ingesting JSON sources with nested objects and arrays. For example, following the JSON sample file below is a query that captures each value in the arrays:

```
{
    "payload" :
    {
        "IBP_IndEvent_MSR" :
        {
            "unit" : "ms",
            "value" : [ 0, 1 ]
        },
        "IBP_IndEvent_RMF" :
```

```
{
    "unit" : "-",
    "value" : [ 0.012, 1.398, 3.1415 ]
    }
}
```

To read the JSON file above, the following query uses an object binding (?values []) to drill down to the value arrays in the source. An @ selector is specified in the ?value variable binding (?value ("@" xsd:double)) to retrieve each of the array values. For an array of primitive values, the @ selector captures each value in the array. If the source value was an array of objects, the @ selector would retrieve a JSON representation for each object in the array. In addition to creating a new binding context for the primitive array values, the ?values object binding also includes ?index ("!array::index") to capture the index array with the primitive value.

```
<http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX s:
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
SELECT *
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
      ?data a s:FileSource ;
      s:url "/mnt/data/json/array-index.json" ;
      s:selector "payload.*" ;
      ?unit (xsd:string) ;
      ?values [
         s:selector "value" ;
         ?value ("@" xsd:double) ;
        ?index ("!array::index") ;
      ] .
  }
}
```

The results of the query are shown below:

unit		value		index
	+-		-+-	
ms		0		0
ms		1		1
-		0.012		0
-		1.398		1
-		3.1415		2

If you do not want to retrieve all of the values in an array, you can include the specific index number to retrieve instead of using the @ symbol. In the variable binding, the index number is appended in brackets ([]) to the binding column name. For example, the following variable binding retrieves the second index value (the third value in the array) from a "projects" array: <code>?project ("projects [2]")</code>. The next example uses the following JSON file:

```
{
   "field1" : "value1",
   "arrayfield" : [
        "arrayvalue1",
        "arrayvalue2"
]
}
```

To retrieve only the second value in the array, the following query appends the index value 1 to the array column name, arrayfield:

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
SELECT *
WHERE {
    SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
        ?json a s:FileSource ;
        s:url "/mnt/data/json/array-index-2.json" ;
        ?field1 (xsd:string) ;
        ?arrayval ("arrayfield[1]" xsd:string) .
    }
}
```

The results of the query are shown below:

## **Capturing Property Keys**

In GDI Generator queries, the names of property keys can be captured from files by including a variable as the s:selector and using the same variable as the s:key. For example, the GDI query below ingests the following simple JSON file.

```
# company.json
{
    "AAPL": {
        "name": "Apple Corp"
    },
    "MSFT": {
        "name": "Microsoft"
    },
    "IBM": {
        "name": "IBM"
    }
}
```

In the query, the keys "AAPL," "MSFT," and "IBM" are selected as the ?TickerSymbol variable and the key is set to the same value.

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
INSERT ${targetGraph} {
 ?s ?p ?o .
 }
}
WHERE {
   SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
      ?data a s:FileSource ;
        s:url "/opt/shared/data/company.json" ;
        s:selector "?TickerSymbol" ;
        s:key (?TickerSymbol) ;
        s:model "Company" ;
        ?TickerSymbol (xsd:string) ;
        ?name (xsd:string) .
      ?rdf a s:RdfGenerator, s:OntologyGenerator;
        s:as (?s ?p ?o) ;
        s:ontology <http://cambridgesemantics.com/ontologies/company> ;
        s:base ${targetGraph} .
  }
}
```

Selecting the predicates and objects from the graph shows the tickerSymbol predicate and value.

```
SELECT ?p ?o
${usingSources}
WHERE { ?s ?p ?o . }
ORDER BY desc(?o)
```

p	0
http://anzograph.com/ontologies/company#Company.name	Microsoft
<pre>http://anzograph.com/ontologies/company#Company.tickerSymbol  </pre>	MSFT
http://anzograph.com/ontologies/company#Company.name	IBM
<pre>http://anzograph.com/ontologies/company#Company.tickerSymbol  </pre>	IBM
http://anzograph.com/ontologies/company#Company.name	Apple Corp
<pre>http://anzograph.com/ontologies/company#Company.tickerSymbol  </pre>	AAPL

## **Query Examples**

The example query below reads a JSON file that contains data about weather. Since the file is hierarchical, the s:selector property is included to specify the path to data in the hourly/data hierarchy:

```
<http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX s:
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
SELECT *
WHERE
{
  SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
      ?data a s:FileSource;
        s:url "/mnt/data/json/weather.json" ;
        ?latitude (xsd:double) ;
        ?longitude (xsd:double) ;
        ?timezone (xsd:string) ;
        s:selector: "hourly.data" ;
        ?time (xsd:long) ;
        ?rainProbability ("precipProbability" xsd:double) ;
        ?temperature (xsd:double) ;
        ?feelsLike ("apparentTemperature" xsd:double) ;
        ?windSpeed (xsd:double) .
  }
}
```

The following example query ingests data from a JSON file that contains data about the New York Times best selling books.

```
<http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX s:
PREFIX books: <http://cambridgesemantics.com/ontologies/NYT Bestsellers Ontology#>
INSERT {
    GRAPH ${targetGraph}{
      ?book a books:Book ;
      books:p Title ?title ;
      books:p Description ?description ;
      books:p Author ?author ;
      books:p Publisher ?publisher ;
     books:p Date ?rawdate .
  }
}
WHERE {
    SERVICE <http://cambridgesemantics.com/services/DataToolkit>
  {
      ?data a s:FileSource ;
      s:url "/mnt/data/json/nyt best sellers.json" ;
      ?title () ;
      ?author () ;
      ?description () ;
      ?publisher () ;
      ?price() ;
      ?rawdate ("bestsellers date.$date.$numberLong").
  }
  BIND (IRI (CONCAT ("http://cambridgesemantics.com/ontologies/NYT Bestsellers Ontology/",
ENCODE FOR URI(?title))) AS ?book) .
}
```

A snippet of the file's contents is shown below:

```
{
    "_id": {
        "$oid": "5b4aa4ead3089013507db18b"
    },
    "bestsellers_date": {
        "$date": {
            "$numberLong": "1211587200000"
        }
    },
    "published_date": {
    }
}
```

```
"$date": {
      "$numberLong": "1212883200000"
    }
  },
  "amazon_product_url": "http://www.amazon.com/Odd-Hours-Dean-
Koontz/dp/0553807056?tag=NYTBS-20",
  "author": "Dean R Koontz",
  "description": "Odd Thomas, who can communicate with the dead, confronts evil forces
in a California coastal town.",
  "price": {
    "$numberInt": "27"
    },
  "publisher": "Bantam",
  "title": "ODD HOURS",
  "rank": {
    "$numberInt": "1"
    },
  "rank last week": {
    "$numberInt": "0"
    },
  "weeks on list": {
    "$numberInt": "1"
    }
}
```

#### **Querying XML Files**

This topic provides details about the structure to use when writing GDI queries to read or ingest data from XML files. It also includes example queries that may be useful as a starting point for writing your own GDI queries.

- Query Syntax
- Hierarchical Bindings and Arrays
- Query Examples

## **Query Syntax**

The following query syntax shows the structure of a GDI query for XML sources. The clauses, patterns, and placeholders that are links are described below.

```
# PREFIX Clause
PREFIX s:
             <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX rdf:
             <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
            <http://purl.org/dc/elements/1.1/>
PREFIX dc:
# Result Clause
{
  [ GRAPH ${targetGraph} { ]
  triple patterns
 [ } ]
}
[ ${usingSources} ]
WHERE
{
   # SERVICE Clause: Include the following service call when reading or inserting data.
    SERVICE [ TOPDOWN ] <http://cambridgesemantics.com/services/DataToolkit>
   # View SERVICE Clause: Or use the service call below when constructing a view.
    SERVICE <http://cambridgesemantics.com/services/DataToolkitView>(${targetGraph})
    {
      ?data a s:FileSource ;
       s:url "string" ;
        [ s:options [ file storage connection options ] ; ]
        [ s:pattern "string" ; ]
        [ s:maxDepth int ; ]
        [ s:format [ source format options ; ] ; ]
        [ s:mimetype "string" ; ]
        [ s:username "string" ; ]
        [ s:password "string" ; ]
        [ s:timeout int ; ]
        [ s:batching boolean | int ; ]
        [ s:paging [ pagination options ; ]
        [ s:concurrency int | [ list of properties ] ; ]
        [ s:rate int | "string" ; ]
        [ s:locale "string" ; ]
        [ s:sampling int ; ]
```

```
[ s:selector "string" | [ list ] ; ]
[ s:model "string" ; ]
[ s:key ("string") ; ]
[ s:reference [ s:model "string" ; s:using ("string") ]
[ s:formats [ datatype_formatting_options ] ; ]
[ s:normalize boolean | [ normalization_rules ] ; ]
[ s:count ?variable ; ]
[ s:count ?variable ; ]
[ s:limit int ; ]
# Mapping variables and hierarchical bindings
?mapping_variable ( [ "binding" ] [ datatype ] [ "datetime_format" ] ) ;
...;
# Additional clauses such as BIND, VALUES, FILTER
}
```

### Note

}

#### For readability, the parameters below exclude the base URI

<http://cambridgesemantics.com/ontologies/DataToolkit#> as well as the s: prefix. As shown in the examples, however, the s: prefix or full property URI does need to be included in queries.

Option	Туре	Description
PREFIX Clause	N/A	The PREFIX clause declares the standard and custom prefixes for GDI service queries. Generally, queries include the prefixes from the query template (or a subset of them) plus any data-specific declarations.
Result Clause	N/A	The result clause defines the type of SPARQL query to run and the set of results to return, i.e., whether you want to read (SELECT or CONSTRUCT) from the source or ingest the data into Anzo (INSERT).
GRAPH	N/A	Include the GRAPH keyword and target graph parameter

Option	Туре	Description
\${targetGraph}		<pre>\$ {targetGraph} when writing an INSERT query to ingest data into a graphmart. Anzo automatically populates the query with the appropriate target URIs when the query runs.</pre>
\${usingSource s}	N/A	Include the source graph parameter $\{usingSources\}\$ when writing a "topdown" query that passes values from the data that is in the graphmart to the data source. Anzo automatically populates the query with the appropriate FROM clauses when the query runs. When passing literal values to the remote source, you do not need to include the source graph parameter. The SERVICE Clause description below includes more information about passing input to data sources.
SERVICE Clause	N/A	<pre>Include the SERVICE call SERVICE [ TOPDOWN ] <http: cambridgesemantics.com="" datatool="" kit="" services=""> to invoke the GDI service when you are running a SELECT, INSERT, or CONSTRUCT query that is not creating a view. When writing a CONSTRUCT query in a View Step, use the DataToolkitView service call, as described below in View SERVICE Clause.</http:></pre>
		Include the optional TOPDOWN keyword when you want to pass input values from the graphmart to the data source. When you include TOPDOWN in the service call, it indicates that the rest of the query produces values to send to the source. In this case, the GDI makes repeated calls to pass in each of the specified values and retrieve the data that is based on those values.
View SERVICE Clause	N/A	When writing a CONSTRUCT query that creates a view of the data (usually in a View Step), include the following SERVICE call: SERVICE <http: cambridgesemantics.com="" datatool<="" services="" td=""></http:>

Option	Туре	Description
		<pre>kitView&gt;(\${targetGraph}). Using the DataToolkitView call optimizes query execution because it tells the GDI to inspect the query and determine which filters to push to the data source. It also limits the result set and retrieves only the data that is needed, i.e., the source data is fully mapped but all of the mapped data is not necessarily returned.</pre>
url	string	This property specifies the file system location of the source file or directory of files. When specifying a directory (such as s:url "/opt/shared-files/loads/"), the GDI loads all of the file formats it recognizes. To specify a directory but limit the number or type of files that are read, you can include the pattern and/or maxDepth properties.
options	RDF list	If additional connection information needs to be provided to access the file storage system, include the options property to list any storage-specific connection parameters. See File Storage Connection Options for information about the supported properties for each storage type.
pattern	string	This property is used to specify a wildcard pattern for matching file names. For example, s:pattern "common_prefix*.xml". You can include one s:pattern property per FileSource. The GDI supports Unix file globbing syntax outside of parentheses. Within parentheses, full Java regular expression language is supported. For example, including s:pattern "data/**/customer_*.xml" tells the GDI to load all files that match the pattern "customer_*.xml" from any number of subdirectories under the data directory. Similarly s:pattern " (\d+)/transaction_*.xml" tells the GDI to load all files that match the pattern "transaction_*.xml" in all subdirectories.

Option	Туре	Description
maxDepth	int	This property can be used to limit the directory traversal depth. By default, when $s:url$ specifies a directory (and $as:pattern$ that limits that traversal depth is not specified), all subdirectories are processed. To process only the files in the top level directory, set maxDepth to 0 ( $s:maxDepth$ 0). To process the files in the top level directory plus the first-level subdirectories, set maxDepth to 1 ( $s:maxDepth$ 1), and so on.
format	RDF list	You can include the format property to add parameters that describe the source files. See File Source Format Options for details about the supported parameters.
mimetype	string	This property can be included to specify the MIME type of the data.
username	string	If authentication is required to access the source, include this property to specify the user name.
password	string	This property lists the password for the given username.
timeout	int	This property can be used to specify the timeout (in milliseconds) to use for requests against the source. For example, s:timeout 5000 configures a 5 second timeout.
batching	boolean or int	This property can be used to disable batching, or it can be used to change the default the batch size. By default, batching is set to 5000 (s:batching 5000). To disable batching, you can include s:batching false in the query. Typically users do not change the batching size. However, it can be useful to control the batch size when performing updates. To configure the size, include s:batching int in the query. For example, s:batching 3000.

Option	Туре	Description
paging	RDF list	This property can be used to configure paging so that the GDI can access large amounts of data across a number of smaller requests. For details about the paging property, see Paginating Requests.
concurrency	int or RDF list	<pre>This property can be included to configure the maximum level of concurrency for the query. The value can be an integer, such as s:concurrency 8. If the value is an integer, it configures a maximum limit on the number of slices that can execute the query. For finer-grained control over the number of nodes and slices to use, concurrency can also be included as an object with limit, nodes, and/or executorsPerNode properties. For example, the following object configures a concurrency model that allows a maximum of 24 executors distributed across 4 nodes with 8 executors per node: s:concurrency [ s:limit 24 ; s:nodes 4 ; s:executorsPerNode 8 ; ] ;</pre>
rate	int or string	This property can be included to control the frequency with which a request is sent to the source. The limit applies to the number of requests a single slice can make. If you specify an integer for the rate, then the value is treated as the maximum number of requests to issue per minute. If you specify a string, you have more flexibility in configuring the rate. The sample values below show the types of values that are supported: <pre>s:rate "90/minute" ; s:rate "90 per minute" ; s:rate "200000 every week" ;</pre>

Option	Туре	Description
		s:rate "10000 every 6 hours" ;
		To enforce the rate limit, the GDI introduces a sleep between requests that is equal to the rate delay. The more executing slices, the longer the rate delay needs to be to enforce the limit in aggregate.
		Given the example of s:rate "90/minute", the GDI would optimize the concurrency and only use 1 slice for execution with a rate delay of 666ms between requests. If s:rate "240/minute", the GDI would use 3 executors with a rate delay of 750ms between requests.
locale	string	This property can be used to specify the locale to use when parsing locale-dependent data such as numbers, dates, and times.
sampling	int	This property can be used to configure the number of records in the source to examine for data type inferencing.
selector	string or RDF list	This property can be used for XML path extraction to traverse nested structures and target specific data. For example, s:selector "projects" targets the projects class of data. To express a hierarchy, use dot notation. For example, s:selector "region.state.city" navigates a hierarchy to target city data. For more information about binding components and the selector property, see Hierarchical Bindings and Arrays below.
model	string	This property defines the class (or table) name for the type of data that is generated from the specified data source. For example, s:model "employees". Model is optional when querying a single source. If your query targets multiple sources, however, and

Option	Туре	Description
		you want to define resource templates (primary keys) and object properties (foreign keys), you must specify the model value for each source.
key	string	This property can be used to define the primary key column for the source file or table. This column is leveraged in a resource template for the instances that are created from the source. For example, s:key ("EMPLOYEE_ID"). For more information about key, see Data Linking Options.
reference	RDF list	This property can be used to specify a foreign key column. The reference property is an RDF list that includes the model property to list the target table and a using property that defines the foreign key column. For more information about reference, see Data Linking Options.
formats	RDF list	To give users control over the data types that are used when coercing strings to other types, this property can be included in GDI queries to define the desired types. In addition, it can be used to describe the formats of date and time values in the source to ensure that they are recognized and parsed to the appropriate date, time, and/or dateTime values. For details about the formats property, see Data Type Formatting Options.
normalize	RDF list	To give users control over the labels and URIs that are generated, the GDI offers several options for normalizing the model and/or the fields that are created from the specified data source(s). For details about the normalize property, see Model Normalization Options.
count	variable	If you want to turn the query into a COUNT query, you can include this property with a <code>?variable</code> to perform a count. For example,

Option	Туре	Description
		s:count ?count.
offset	int	This property can be used to offset the data that is returned by a number of rows.
limit	int	You can include this property to limit the number of results that are returned. s:limit maps to the SPARQL LIMIT clause.
mapping_ variable	variable	The mapping variables, in <code>?mapping_variable</code> (["binding"] [datatype] ["datetime_format"]) format, define the triple patterns to output. When the specified <code>?variable</code> matches the source column name, the GDI uses the variable as the source data selector. If you specify an alternate variable name, a binding needs to be specified to map the new variable to the source. You also have the option to transform the data using the datatype and datetime_format options. Tip See Hierarchical Bindings and Arrays below for more information about configuring mapping variables and
		unpacking files with nested objects and arrays.
binding	string	The binding is a literal value that binds a ?mapping_variable to a source column. If you specify a ?variable that matches the source column name, then that variable name is the data selector and it is not necessary to specify a binding. If you specify an alternate variable name or there is a hierarchical path to the source column, then the binding is needed to map the new variable to that source column. For example, the following pattern simply binds the source column AIRLINE to the lowercase variable ?airline: ?airline

Option	Туре	Description
		("AIRLINE").
		Note For FileSource, periods (.), forward slashes (/), and brackets ([]) are parsed as path notation. Therefore, if a source column name includes any of those characters they must be escaped in the binding. Use two backslashes (\\) as an escape character. For example, if a column name is <b>average/day</b> , the variable and binding pattern could be written as ?averagePerDay ("average\\/day").
datatype	URI	The datatype is the data type to convert the column to. If you do not specify a data type, the GDI infers the type. The GDI supports the following types: xsd:int, xsd:long, xsd:float, xsd:double, xsd:boolean, xsd:time, xsd:dateTime, xsd:date, xsd:duration, xsd:timeDuration, xsd:yearMonthDuration, xsd:gMonthDay, xsd:gMonth, xsd:gYearMonth, xsd:anyURI
datetime_ format	string	This option is used to specify the format to use for date and time data types. The GDI supports Java date and time formats. Specify days as "d," months as "M," and years as "y." For the time, specify "H" for hours, "m" for minutes, and "s" for seconds. For example, "yyyyMMdd HH:mm:ss" or "ddMMMyy" to display date values such as "01JAN19."

Option	Туре	Description
		The GDI's default base year is 2000. If the source data has years with only two digits, such as 02–04–99, the GDI prepends 20 to the digits. The value 02–04–99 is parsed to 02–04–2099. To specify an alternate base year to use for two-digit values, you can include the notation ^nnnn (e.g., ^1900) in the format value. For example, to set the base year to 1900 instead of 2000, use a format value such as xsd:date "dd-MMM-yy^1900" or xsd:date "dd-MMM-yy^1990". When one of those values is specified, 02–04–99 is parsed to 02–04–1999.

## **Hierarchical Bindings and Arrays**

When configuring the mapping variables in a query, the GDI provides syntax for unpacking XML files with nested objects and arrays. One way to express hierarchies in queries is to use brackets ([]) to group objects into binding trees. For example, the WHERE clause snippet below organizes mapping variable objects into an hourly/data hierarchy by nesting the ?data patterns inside the

```
?hourly [ ] tree:
```

```
WHERE
{
   SERVICE <http://cambridgesemantics.com/services/DataToolkit>
   {
        ?data a s:FileSource;
        s:url "/mnt/data/xml/weather.xml";
        ?latitude (xsd:double);
        ?latitude (xsd:double);
        ?longitude (xsd:double);
        ?timezone (xsd:string);
        ?hourly
[
            ?data
[
```

```
?time (xsd:long) ;
    ?rainProbability ("precipProbability" xsd:double) ;
    ?temperature (xsd:double) ;
    ?feelsLike ("apparentTemperature" xsd:double) ;
    ?windSpeed (xsd:double) ;
    ];
  ].
}
```

When constructing object binding trees, if you choose to introduce the hierarchy with a variable name that is not an exact match to the source label, include a **selector** property to list the value from the source. For example, in the WHERE clause snippet below, s:selector is included to select eventHeader in the source as ?event in the query and statLocation as ?location.

```
WHERE
{
  SERVICE <http://cambridgesemantics.com/services/DataToolkit>
  {
      ?data a s:FileSource ;
      s:url "/mnt/data/xml/part 1.xml" ;
      ?event
Γ
         s:selector "eventHeader" ;
           ?eventId (xsd:string) ;
           ?eventName (xsd:string) ;
           ?eventVersion (xsd:string) ;
           ?eventTime (xsd:dateTime) ;
      1;
      ?location
[
         s:selector "statLocation" ;
           ?locationId (xsd:string) ;
           ?lineNo (xsd:int) ;
           ?statNo (xsd:int) ;
           ?statId (xsd:int) ;
      ].
  }
}
```

As an alternative to grouping objects in binding trees, the **selector** property also supports using dot notation to specify paths. For example, the WHERE clause snippet below rewrites the first example query to express the same <code>hourly/data</code> hierarchy as a path in the <code>s:selector</code> value:

```
WHERE
{
  SERVICE <http://cambridgesemantics.com/services/DataToolkit>
    {
      ?data a s:FileSource;
        s:url "/mnt/data/xml/weather.xml" ;
        ?latitude (xsd:double) ;
        ?longitude (xsd:double) ;
        ?timezone (xsd:string) ;
        s:selector: "hourly.data" ;
        ?time (xsd:long) ;
        ?rainProbability ("precipProbability" xsd:double) ;
        ?temperature (xsd:double) ;
        ?feelsLike ("apparentTemperature" xsd:double) ;
        ?windSpeed (xsd:double) .
  }
}
```

# **Query Examples**

The following example query ingests data from an XML file that contains hierarchies.

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX fmcsa: <http://census.gov/ontologies/FMCSA#>
INSERT {
    GRAPH ${targetGraph} {
    }
}
${usingSources}
WHERE {
    SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
        ?data a s:FileSource ;
        s:url "file:///opt/shared/data/xml/define.xml" ;
```

```
?ItemGroupDef [
        ?OID (xsd:string) ;
        ?Name (xsd:string) ;
        ?Repeating (xsd:string) ;
        ?IsReferenceData (xsd:string) ;
        ?Purpose (xsd:string) ;
        ?Label (xsd:string) ;
        ?Structure (xsd:string) ;
        ?DomainKeys (xsd:string) ;
        ?Class (xsd:string) ;
        ?ArchiveLocationID (xsd:string) ;
        ?Comment (xsd:string) ;
        ?ItemRef [
           ?ItemOID (xsd:string) ;
           ?OrderNumber (xsd:int) ;
           ?Mandatory (xsd:string) ;
        ];
   ].
}
```

### **Querying Parquet and SAS Files**

This topic provides details about the structure to use when writing GDI queries to read or ingest data from Parquet and SAS files. It also includes example queries that may be useful as a starting point for writing your own GDI queries.

#### Note

}

Loading parquet files from Amazon S3 is not supported. Amazon does not provide a VFS driver that supports random access reads, which are necessary for reading parquet files.

- Query Syntax
- Query Examples

# **Query Syntax**

The following query syntax shows the structure of a GDI query for Parquet and SAS sources. The clauses, patterns, and placeholders that are links are described below.

```
# PREFIX Clause
PREFIX s:
              <http://cambridgesemantics.com/ontologies/DataToolkit#>
             <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdf:
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
            <http://purl.org/dc/elements/1.1/>
PREFIX dc:
# Result Clause
{
   [ GRAPH ${targetGraph} { ]
  triple_patterns
 [ } ]
}
[ ${usingSources} ]
WHERE
{
   # SERVICE Clause: Include the following service call when reading or inserting data.
    SERVICE [ TOPDOWN ] <http://cambridgesemantics.com/services/DataToolkit>
   # View SERVICE Clause: Or use the service call below when constructing a view.
    SERVICE <http://cambridgesemantics.com/services/DataToolkitView>(${targetGraph})
    {
      ?data a s:FileSource ;
        s:url "string" ;
        [ s:options [ file storage connection options ] ; ]
        [ s:pattern "string" ; ]
        [ s:maxDepth int ; ]
        [ s:format [ source_format_options ; ] ; ]
        [ s:mimetype "string" ; ]
        [ s:username "string" ; ]
```

[ s:password "string" ; ]

```
[ s:timeout int ; ]
   [ s:batching boolean | int ; ]
   [ s:paging [ pagination options ; ]
   [ s:concurrency int | [ list_of_properties ] ; ]
   [ s:rate int | "string" ; ]
   [ s:locale "string" ; ]
   [ s:sampling int ; ]
   [ s:selector "string" | [ list ] ; ]
   [ s:model "string" ; ]
   [ s:key ("string") ; ]
   [ s:reference [ s:model "string" ; s:using ("string") ]
   [ s:formats [ datatype formatting options ] ; ]
   [ s:normalize boolean | [ normalization_rules ] ; ]
   [ s:count ?variable ; ]
   [ s:offset int ; ]
   [ s:limit int ; ]
   # Mapping variables and hierarchical bindings
   ?mapping variable ( [ "binding" ] [ datatype ] [ "datetime format" ] ) ;
   ...;
# Additional clauses such as BIND, VALUES, FILTER
```

```
}
```

}

Option	Туре	Description
PREFIX Clause	N/A	The PREFIX clause declares the standard and custom prefixes for GDI service queries. Generally, queries include the prefixes from the query template (or a subset of them) plus any data-specific declarations.
Result Clause	N/A	The result clause defines the type of SPARQL query to run and the set of results to return, i.e., whether you want to read (SELECT or CONSTRUCT) from the source or ingest the data into Anzo (INSERT).
GRAPH \${targetGraph}	N/A	Include the GRAPH keyword and target graph parameter \${targetGraph} when writing an INSERT query to ingest data into a graphmart. Anzo automatically populates the query with the

Option	Туре	Description
		appropriate target URIs when the query runs.
\${usingSource s}	N/A	Include the source graph parameter \${usingSources} when writing a "topdown" query that passes values from the data that is in the graphmart to the data source. Anzo automatically populates the query with the appropriate FROM clauses when the query runs. When passing literal values to the remote source, you do not need to include the source graph parameter. The SERVICE Clause description below includes more information about passing input to data sources.
SERVICE Clause	N/A	<pre>Include the SERVICE call SERVICE [ TOPDOWN ] <http: cambridgesemantics.com="" datatool="" kit="" services=""> to invoke the GDI service when you are running a SELECT, INSERT, or CONSTRUCT query that is not creating a view. When writing a CONSTRUCT query in a View Step, use the DataToolkitView service call, as described below in View SERVICE Clause.</http:></pre>
		Include the optional TOPDOWN keyword when you want to pass input values from the graphmart to the data source. When you include TOPDOWN in the service call, it indicates that the rest of the query produces values to send to the source. In this case, the GDI makes repeated calls to pass in each of the specified values and retrieve the data that is based on those values.
View SERVICE Clause	N/A	When writing a CONSTRUCT query that creates a view of the data (usually in a View Step), include the following SERVICE call: SERVICE <http: cambridgesemantics.com="" datatool<br="" services="">kitView&gt;(\${targetGraph}). Using the DataToolkitView call optimizes query execution because it tells the GDI to inspect</http:>

Option	Туре	Description
		the query and determine which filters to push to the data source. It also limits the result set and retrieves only the data that is needed, i.e., the source data is fully mapped but all of the mapped data is not necessarily returned.
url	string	This property specifies the file system location of the source file or directory of files. When specifying a directory (such as s:url "/opt/shared-files/loads/" or s:url "gs://shared-files/my-files/"), the GDI loads all of the file formats it recognizes. To specify a directory but limit the number or type of files that are read, you can include the pattern and/or maxDepth properties.
options	RDF list	If additional connection information needs to be provided to access the file storage system, include the options property to list any storage-specific connection parameters. See File Storage Connection Options for information about the supported properties for each storage type.
pattern	string	This property is used to specify a wildcard pattern for matching file names. For example, s:pattern "common_ prefix*.parquet". You can include one s:pattern property per FileSource. The GDI supports Unix file globbing syntax outside of parentheses. Within parentheses, full Java regular expression language is supported. For example, including s:pattern "data/**/customer_*.parquet" tells the GDI to load all files that match the pattern "customer_*.parquet" from any number of subdirectories under the data directory. Similarly s:pattern " (\d+)/transaction_*.xpt" tells the GDI to load all files that match the pattern "transaction_*.xpt" in all subdirectories.
maxDepth	int	This property can be used to limit the directory traversal depth. By

Option	Туре	Description
		default, when s:url specifies a directory (and a s:pattern that limits that traversal depth is not specified), all subdirectories are processed. To process only the files in the top level directory, set maxDepth to 0 (s:maxDepth 0). To process the files in the top level directory plus the first-level subdirectories, set maxDepth to 1 (s:maxDepth 1), and so on.
format	RDF list	You can include the format property to add parameters that describe the source files. See File Source Format Options for details about the supported parameters.
mimetype	string	This property can be included to specify the MIME type of the data.
username	string	If authentication is required to access the source, include this property to specify the user name.
password	string	This property lists the password for the given username.
timeout	int	This property can be used to specify the timeout (in milliseconds) to use for requests against the source. For example, s:timeout 5000 configures a 5 second timeout.
batching	boolean or int	This property can be used to disable batching, or it can be used to change the default the batch size. By default, batching is set to 5000 (s:batching 5000). To disable batching, you can include s:batching false in the query. Typically users do not change the batching size. However, it can be useful to control the batch size when performing updates. To configure the size, include s:batching int in the query. For example, s:batching 3000.
paging	RDF list	This property can be used to configure paging so that the GDI can

Option	Туре	Description
		access large amounts of data across a number of smaller requests. For details about the paging property, see Paginating Requests.
concurrency	int or RDF list	<pre>This property can be included to configure the maximum level of concurrency for the query. The value can be an integer, such as s:concurrency 8. If the value is an integer, it configures a maximum limit on the number of slices that can execute the query. For finer-grained control over the number of nodes and slices to use, concurrency can also be included as an object with limit, nodes, and/or executorsPerNode properties. For example, the following object configures a concurrency model that allows a maximum of 24 executors distributed across 4 nodes with 8 executors per node: s:concurrency [ s:limit 24 ; s:nodes 4 ; s:executorsPerNode 8 ; ] ;</pre>
rate	int or string	This property can be included to control the frequency with which a request is sent to the source. The limit applies to the number of requests a single slice can make. If you specify an integer for the rate, then the value is treated as the maximum number of requests to issue per minute. If you specify a string, you have more flexibility in configuring the rate. The sample values below show the types of values that are supported: <pre>s:rate "90/minute" ; s:rate "90 per minute" ; s:rate "200000 every week" ; s:rate "10000 every 6 hours" ;</pre>

Option	Туре	Description
		To enforce the rate limit, the GDI introduces a sleep between requests that is equal to the rate delay. The more executing slices, the longer the rate delay needs to be to enforce the limit in aggregate.
		Given the example of s:rate "90/minute", the GDI would optimize the concurrency and only use 1 slice for execution with a rate delay of 666ms between requests. If s:rate "240/minute", the GDI would use 3 executors with a rate delay of 750ms between requests.
locale	string	This property can be used to specify the locale to use when parsing locale-dependent data such as numbers, dates, and times.
sampling	int	This property can be used to configure the number of records in the source to examine for data type inferencing.
selector	string or RDF list	This property can be used as a binding component to identify the path to the source objects. For example, s:selector "Sales.SalesOrderHeader" targets the SalesOrderHeader table in the Sales schema. For more information about binding components and the selector property, see Using Binding Trees and Selector Paths.
model	string	This property defines the class (or table) name for the type of data that is generated from the specified data source. For example, s:model "employees". Model is optional when querying a single source. If your query targets multiple sources, however, and you want to define resource templates (primary keys) and object properties (foreign keys), you must specify the model value for each source.

Option	Туре	Description
key	string	This property can be used to define the primary key column for the source file or table. This column is leveraged in a resource template for the instances that are created from the source. For example, s:key ("EMPLOYEE_ID"). For more information about key, see Data Linking Options.
reference	RDF list	This property can be used to specify a foreign key column. The reference property is an RDF list that includes the model property to list the target table and a using property that defines the foreign key column. For more information about reference, see Data Linking Options.
formats	RDF list	To give users control over the data types that are used when coercing strings to other types, this property can be included in GDI queries to define the desired types. In addition, it can be used to describe the formats of date and time values in the source to ensure that they are recognized and parsed to the appropriate date, time, and/or dateTime values. For details about the formats property, see Data Type Formatting Options.
normalize	RDF list	To give users control over the labels and URIs that are generated, the GDI offers several options for normalizing the model and/or the fields that are created from the specified data source(s). For details about the normalize property, see Model Normalization Options.
count	variable	If you want to turn the query into a COUNT query, you can include this property with a <code>?variable</code> to perform a count. For example, <code>s:count</code> ?count.
offset	int	This property can be used to offset the data that is returned by a number of rows.

Option	Туре	Description	
limit	int	You can include this property to limit the number of results that are returned. s:limit maps to the SPARQL LIMIT clause.	
mapping_ variable	variable	The mapping variables, in <code>?mapping_variable</code> (["binding"] [datatype] ["datetime_format"]) format, define the triple patterns to output. When the specified <code>?variable</code> matches the source column name, the GDI uses the variable as the source data selector. If you specify an alternate variable name, a binding needs to be specified to map the new variable to the source. You also have the option to transform the data using the datatype and datetime_format options.	
binding	string	The binding is a literal value that binds a ?mapping_variable to a source column. If you specify a ?variable that matches the source column name, then that variable name is the data selector and it is not necessary to specify a binding. If you specify an alternate variable name or there is a hierarchical path to the source column, then the binding is needed to map the new variable to that source column. For example, the following pattern simply binds the source column AIRLINE to the lowercase variable ?airline: ?airline ("AIRLINE").	
		Note For FileSource, periods (.), forward slashes (/), and brackets ([]) are parsed as path notation. Therefore, if a source column name includes any of those characters they must be escaped in the binding. Use two backslashes (\\) as an escape character. For example, if a column name is <b>average/day</b> , the variable and binding pattern	
Option	Туре	Description	
------------------	--	---	--
		<b>could be written as</b> ?averagePerDay ("average\\/day").	
datatype	URI	The datatype is the data type to convert the column to. If you do not specify a data type, the GDI infers the type. The GDI supports the following types: xsd:int, xsd:long, xsd:float, xsd:double, xsd:boolean, xsd:time, xsd:dateTime, xsd:date, xsd:duration, xsd:timeDuration, xsd:yearMonthDuration, xsd:gMonthDay, xsd:gMonth, xsd:gYearMonth, xsd:anyURI	
datetime_ string	This option is used to specify the format to use for date and time data types. The GDI supports Java date and time formats. Specify days as "d," months as "M," and years as "y." For the time, specify "H" for hours, "m" for minutes, and "s" for seconds. For example, "yyyyMMdd HH:mm:ss" or "ddMMMyy" to display date values such as "01JAN19."		
		Note The GDI's default base year is 2000. If the source data has years with only two digits, such as 02-04-99, the GDI prepends 20 to the digits. The value 02-04-99 is parsed to 02-04-2099. To specify an alternate base year to use for two-digit values, you can include the notation ^nnnn (e.g., ^1900) in the format value. For example, to set the base year to 1900 instead of 2000, use a format value such as xsd:date "dd-MMM-yy^1900" or xsd:date "dd-MMM-yy^1990". When one of those values is	

Option	Туре	Description
		<b>specified,</b> 02-04-99 <b>is parsed to</b> 02-04-1999.

# **Query Examples**

The following query inserts data from a Parquet file on the shared file system and uses a binding tree to express hierarchical data.

```
PREFIX rdf:
            <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
             <http://purl.org/dc/elements/1.1/>
PREFIX dc:
             <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX s:
INSERT {
   GRAPH ${targetGraph} {
     ?s ?p ?o
  }
}
WHERE {
  SERVICE <http://cambridgesemantics.com/services/DataToolkit>
  {
    ?data a s:FileSource ;
      s:url "file:///opt/shared/data/parquet/part-c000.snappy.parquet" ;
      s:model "BID" ;
      ?bidid (xsd:string) ;
      ?cur (xsd:string) ;
      ?id (xsd:string) ;
      ?nbr (xsd:long) ;
      ?seatbid [
        ?seat (xsd:string) ;
        ?group (xsd:long) ;
        ?bid [
         ?adid (xsd:string) ;
          ?cid (xsd:string) ;
```

```
?crid (xsd:string) ;
  ?dealid (xsd:string) ;
  ?h (xsd:long) ;
  ?w xsd:long ;
  ?bundle (xsd:string) ;
  ?price (xsd:double) ;
  ];
 ].
}
```

The query below reads data from a Parquet file and filters out data by zip code.

```
PREFIX rdf:
             <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX anzo: <http://openanzo.org/ontologies/2008/07/Anzo#>
PREFIX zowl: <http://openanzo.org/ontologies/2009/05/AnzoOwl#>
            <http://purl.org/dc/elements/1.1/>
PREFIX dc:
            <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX s:
SELECT *
WHERE {
  SERVICE <http://cambridgesemantics.com/services/DataToolkit>
  {
    ?data a s:FileSource ;
      s:url "file:///opt/shared/data.gov/FMCSA CENSUS1 2023Nov.parquet" ;
      ?DOT NUMBER (xsd:string) ;
      ?LEGAL NAME (xsd:string) ;
      ?CARRIER OPERATION (xsd:string) ;
      ?HM FLAG (xsd:string) ;
      ?PC FLAG (xsd:string) ;
      ?PHY STREET (xsd:string) ;
      ?PHY CITY (xsd:string) ;
      ?PHY STATE (xsd:string) ;
      ?PHY ZIP (xsd:string) ;
      ?PHY COUNTRY (xsd:string) .
    FILTER(?PHY ZIP = "78735" || ?PHY ZIP = "77040" || ?PHY ZIP = "78705")
  }
}
```

#### **File Source Format Options**

For file sources, you can include the format property to list additional parameters that describe the source. The supported format parameters are described below.

```
s:format [
  s:delimiter "string" ;
  s:headers boolean ;
  s:columns "string" ;
  s:start int ;
  s:skip int ;
  s:comment "string" ;
  s:quote "string" ;
  s:escape "string" ;
  s:maxColumns int ;
  s:segment boolean ;
```

```
];
```

Option	Туре	Description
delimiter	string	This property specifies the string that is used to delimit columns in the file(s). For example, s:delimiter " ".
headers	boolean	This property indicates whether or not the file(s) include headers. By default the headers value is true (s:headers true). For files that do not have headers, specify s:headers false.
columns	string	If you want the GDI to target only certain columns in the source file(s), you can include the <b>columns</b> property to list the names of columns to include. The value is a single string that is a comma-separated list. For example, s:columns "employee_id, name, address, start date, title".
start	int	If the file includes headers that take up more than one row, include the start property to specify the row number where the

Option	Туре	Description
		data starts to exclude headers. For example, s:start 8.
skip	int	This property can be used to specify the number of rows/records to skip before reading or ingesting the file(s). By default, skip is set to 0 (s:skip 0).
comment	string	This property specifies the string that is used as the comment character in the file(s). The comment value is set to # by default (s:comment "#").
quote	string	This property is used to specify the string that is used as the quote character.
escape	string	This property is used to specify the escape string that is used in the file(s). For example, s:escape "\".
maxColumns	int	This property can be used to set a limit on the maximum number of columns to read or ingest. The maxColumns property is set to -1 (unlimited) by default (s:maxColumns -1).
segment	boolean	This property indicates whether or not the file(s) can be segmented (partitioned). The default value is s:segment true. If you have CSV files that contain embedded new lines, include s:segment false in the query as those files cannot be segmented.

#### **File Storage Connection Options**

If you are querying a **FileSource** and additional connection information needs to be provided to access the file storage system, include the **options** property in the query and define the necessary

storage-specific connection parameters. The parameters that the GDI supports for each type of storage system are pulled directly from the Java API for that system. The supported properties for each storage type are listed below.

- Amazon S3
- FTP & FTPS
- Google Cloud Storage
- HDFS
- SFTP
- WebDAV

## Amazon S3

#### Note

Loading parquet files from S3 is not supported. Amazon does not provide a VFS driver that supports random access reads, which are necessary for reading parquet files.

```
s:options [
   s:accessKey "string";
   s:region "string";
   s:secretKey "string";
   s:serviceName "string";
   s:sessionToken "string";
   s:createBucket boolean;
   s:disableChunkedEncoding boolean;
   s:serverSideEncryption boolean;
   s:useHttps boolean;
```

];

Option	Туре	Description
accessKey	string	The accessKey property can be included to specify the access key.

Option	Туре	Description
region	string	The region property can be included to specify the region.
secretKey	string	The secretKey property can be included to specify the secret key.
serviceName	string	For connections to AWS service endpoints, the serviceName property can be included to specify the service name.
sessionToken	string	The sessionToken property can be included to specify the session token.
createBucket	boolean	Refer to the S3 API documentation.
disableChunkedEncoding	boolean	For increased performance, Amazon S3 requests use chunked encoding by default. To disable chunked encoding, you can include s:disableChunkedEncoding true in the query.
serverSideEncryption	boolean	Refer to the S3 API documentation.
useHttps	boolean	Refer to the S3 API documentation.

## **FTP & FTPS**

```
s:options [
    s:autodetectUtf8 boolean ;
    s:connectTimeout int ;
    s:controlEncoding "string" ;
    s:dataTimeout int ;
```

```
s:defaultDateFormat "string" ;
s:entryParser "string" ;
s:fileType "string" ;
s:passiveMode boolean ;
s:proxy "string" ;
s:recentDateFormat "string" ;
s:remoteVerification boolean ;
s:serverLanguageCode "string" ;
s:serverTimeZoneId "string" ;
s:shortMonthNames "string" ;
s:socketTimeout int ;
s:userDirIsRoot boolean ;
s:dataChannelProtectionLevel "string" ;
s:ftpsMode "string" ;
s:keyManager "string" ;
s:trustManager "string" ;
```

];

Option	Туре	Description
autodetectUtf8	boolean	For FTP connections, the autodetectUtf8 property can be included to indicate whether the FTP server is set to UTF-8 mode or Auto- detect encoding.
connectTimeout	int	For FTP connections, you can include the connectTimeout property to specify the maximum number of seconds to hold a connection before timing out.
controlEncoding	string	Refer to the FTP API documentation.
dataTimeout	int	For FTP connections, you can include the dataTimeout property to specify the maximum number of seconds to transfer data before timing out.

Option	Туре	Description
defaultDateFormat	string	Refer to the FTP API documentation.
entryParser	string	Refer to the FTP API documentation.
fileType	string	Refer to the FTP API documentation.
passiveMode	boolean	For FTP connections, the passiveMode property can be included to indicate whether the data transfer mode is passive or active. If you use passive mode, set passiveMode to true (s:passiveMode true).
proxy	string	If you are using an FTP proxy, include the proxy property to specify the proxy connection details.
recentDateFormat	string	Refer to the FTP API documentation.
remoteVerification	boolean	For FTP connections, the remoteVerification property can be included to indicate whether remote authentication is enabled. If you use remote authentication, set remoteVerification to true (s:remoteVerification true).
serverLanguageCode	string	If the FTP server language is not set to English, include the serverLanguageCode property to specify the language code for the server. For example, s:serverLanguageCode "ES".
serverTimeZoneld	string	For FTP connections, the serverTimeZoneId property can be included

Option	Туре	Description
		to specify the timezone ID for the server.
shortMonthNames	string	Refer to the FTP API documentation.
socketTimeout	int	For FTP connections, you can include the socketTimeout property to specify the maximum number of seconds to transfer data before timing out.
userDirlsRoot	boolean	Refer to the FTP API documentation.
dataChannelProtectionLevel	string	For FTPS connections, the dataChannelProtectionLevel property specifies the Data Channel Protection Level for the server.
ftpsMode	string	For FTPS connections, the ftpsMode property specifies whether the FTPS is in implicit or explicit mode.
keyManager	string	For FTPS connections, the keyManager property specifies the KeyManager value for making an SSL connection to the server.
trustManager	string	For FTPS connections, the trustManager property specifies the TrustManager value for the SSL connection to the server.

# **Google Cloud Storage**

```
s:options [
    s:serviceAccountKey "string" ;
```

];

Option	Туре	Description
serviceAccountKey	string	For connections to GCS, the serviceAccountKey property can be included to specify the key for the service account.

# HDFS

```
s:options [
    s:configName "string";
    s:configPath "string";
    s:configURL "string";
];
```

Option	Туре	Description
configName	string	For connections to HDFS, the configName property can be included to specify the name of the configuration file to read.
configPath	string	For connections to HDFS, the configPath property can be included to list the path to the specified configuration file.
configURL	string	Refer to the HDFS API documentation.

## SFTP

```
s:options [
    s:compression "string";
    s:configRepository "string";
    s:fileNameEncoding "string";
    s:identityProvider "string";
    s:identityRepositoryFactory "string";
    s:keyExchangeAlgorithm "string";
    s:knownHosts "string";
    s:loadOpenSSHConfig boolean;
    s:preferredAuthentications "string";
    s:sessionTimeout int;
    s:strictHostKeyChecking "string";
    s:userInfo "string";
}
```

];

Option	Туре	Description
compression	string	Refer to the SFTP API documentation.
configRepository	string	Refer to the SFTP API documentation.
fileNameEncoding	string	Refer to the SFTP API documentation.
identityProvider	string	Refer to the SFTP API documentation.
identityRepositoryFactory	string	Refer to the SFTP API documentation.
keyExchangeAlgorithm	string	For SFTP connections, you can include the keyExchangeAlgorithm property to specify the key exchange algorithm to use.
knownHosts	string	Refer to the SFTP API documentation.

Option	Туре	Description
IoadOpenSSHConfig	boolean	For SFTP connections, the <pre>loadOpenSSHConfig property indicates whether to read the ~/.ssh/config file.</pre>
preferredAuthentications	string	For SFTP connections, the preferredAuthentications property can be included to specify the authentication order to use.
sessionTimeout	int	For SFTP connections, you can include the sessionTimeout property to specify the maximum number of seconds to leave the session open before timing out.
strictHostKeyChecking	string	For SFTP connections, you can include the <pre>strictHostKeyChecking property to specify</pre> how host keys are checked.
userInfo	string	Refer to the SFTP API documentation.

# WebDAV

```
s:options [
   s:creatorName "string" ;
   s:versioning boolean ;
] ;
Option Type Description
```

Option	Туре	Description
creatorName	string	For WebDAV connections, the creatorName property can be included to add a description of the creator of the resource.

Option	Туре	Description
versioning	boolean	Refer to the WebDAV API documentation.

### **GDI Property Reference**

This topic describes the Graph Data Interface (GDI) properties that are available to use in queries. The first section describes the options that are available regardless of data source type, and the second section describes the source-specific options.

- Universal Properties
- DbSource Properties
- FileSource Properties
- HttpSource Properties
- ElasticSource Properties

#### **Universal Properties**

The table below lists the properties that are valid in queries against all data source types.

Option	Туре	Description
batching	boolean or int	This property can be used to disable batching, or it can be used to change the default the batch size. By default, batching is set to 5000 (s:batching 5000). To disable batching, you can include s:batching false in the query. Typically users do not change the batching size. However, it can be useful to control the batch size when performing updates. To configure the size, include s:batching int in the query. For example, s:batching 3000.
concurrency	int or RDF list	This property can be included to configure the maximum level of concurrency for the query. The value can be an integer, such as s:concurrency 8. If the value is an integer, it configures a maximum limit on the number of slices that can execute the query. For finer-grained control over the number of nodes and slices to use, concurrency can also be included as an object with limit, nodes, and/or executorsPerNode properties. For

Option	Туре	Description
		<pre>example, the following object configures a concurrency model that allows a maximum of 24 executors distributed across 4 nodes with 8 executors per node: s:concurrency [    s:limit 24 ;    s:nodes 4 ;    s:executorsPerNode 8 ; ] ;</pre>
count	variable	If you want to turn the query into a COUNT query, you can include this property with a <code>?variable</code> to perform a count. For example, <code>s:count</code> ?count.
errors	boolean	Controls whether the GDI ignores errors (such as query or file errors) or stops processing the query when an error is encountered. This property is set to true by default (s:errors true). Processing stops when an error is encountered. To ignore errors, you can include s:errors false.
formats	RDF list	To give users control over the data types that are used when coercing strings to other types, this property can be included in GDI queries to define the desired types. In addition, it can be used to describe the formats of date and time values in the source to ensure that they are recognized and parsed to the appropriate date, time, and/or dateTime values. For details about the formats property, see Data Type Formatting Options.
key	string	This property can be used to define the primary key column for the source file or table. This column is leveraged in a resource template for the instances that are created from the source. For

Option	Туре	Description
		example, s:key ("EMPLOYEE_ID"). For more information about key, see Data Linking Options.
limit	int	You can include this property to limit the number of results that are returned. s:limit maps to the SPARQL LIMIT clause.
locale	string	This property can be used to specify the locale to use when parsing locale-dependent data such as numbers, dates, and times.
model	string	This property defines the class (or table) name for the type of data that is generated from the specified data source. For example, s:model "employees". Model is optional when querying a single source. If your query targets multiple sources, however, and you want to define resource templates (primary keys) and object properties (foreign keys), you must specify the model value for each source.
normalize	boolean and/or RDF list	To give users control over the labels and URIs that are generated, the GDI offers several options for normalizing the model and/or the fields that are created from the specified data source(s). For details about the normalize property, see Model Normalization Options.
offset	int	This property can be used to offset the data that is returned by a number of rows.
paging	RDF list	This property can be used to configure paging so that the GDI can access large amounts of data across a number of smaller requests. For details about the paging property, see Paginating Requests.

Option	Туре	Description
password	string	This property lists the password for the given username.
rate	int or string	This property can be included to control the frequency with which a request is sent to the source. The limit applies to the number of requests a single slice can make. If you specify an integer for the rate, then the value is treated as the maximum number of requests to issue per minute. If you specify a string, you have more flexibility in configuring the rate. The sample values below show the types of values that are supported: s:rate "90/minute";
		s:rate "90 per minute" ;
		s:rate "200000 every week" ;
		To enforce the rate limit, the GDI introduces a sleep between requests that is equal to the rate delay. The more executing slices, the longer the rate delay needs to be to enforce the limit in aggregate.
		Given the example of s:rate "90/minute", the GDI would optimize the concurrency and only use 1 slice for execution with a rate delay of 666ms between requests. If s:rate "240/minute", the GDI would use 3 executors with a rate delay of 750ms between requests.
reference	RDF list	This property can be used to specify a foreign key column. The reference property is an RDF list that includes the model property to list the target table and a using property that defines the foreign key column. For more information about reference, see Data Linking Options.
sampling	int	This property can be used to configure the number of records in

Option	Туре	Description
		the source to examine for data type inferencing.
selector	string or RDF list	This property can be used as a binding component to identify the path to the source objects. For example, s:selector "Sales.SalesOrderHeader" targets the SalesOrderHeader table in the Sales schema. For more information about binding components and the selector property, see Using Binding Trees and Selector Paths.
strict	boolean	This property can be used to force the GDI to limit the data to strictly what is stated in the query. For example, when ingesting data from a CSV file, you can include <code>s:strict true</code> on the <code>s:FileSource</code> to ensure that the GDI only ingests columns for which a variable binding exists in the query. In addition, this property can be included in <code>s:formats</code> to control the automatic data type conversion feature (as described in Data Type Formatting Options). The default value is <code>false</code> .
timeout	int	This property can be used to specify the timeout (in milliseconds) to use for requests against the source. For example, s:timeout 5000 configures a 5 second timeout.
url	string	This property specifies the URL for the data source, such as the database URL, Elasticsearch URL, or HTTP endpoint URL. For file-based sources, the url property specifies the file system location of the source file or directory of files. When specifying a directory (such as s:url "/opt/shared-files/loads/"), the GDI loads all of the file formats it recognizes. To specify a directory but limit the number or type of files that are read, you can include the pattern and/or maxDepth properties.

Option	Туре	Description		
		<pre>For security, it is a best practice to reference connection information (such as the url, username, and password) from a Query Context so that the sensitive details are abstracted from any requests. In addition, using a Query Context makes connection details reusable across queries. See Using Query Contexts for more information. For example, the triple patterns below reference keys from a Query Context: ?data a s:DbSource ; s:url "{{@db.eca4bf3ff9a.url}}"; s:password " {{@db.eca4bf3ff9a.password}}";</pre>		
username	string	If authentication is required to access the source, include this property to specify the user name.		

#### **DbSource Properties**

The table below lists the properties that are available for queries against database data sources. For more information about database sources, see Querying a Database Source.

Option	Туре	Description
database	string	This property can be used to specify the database to target in the source if the database is not listed in the s:url or s:selector strings.
driver	string	This property can be included to specify the JDBC driver to

Option	Туре	Description
		use.
orderBy	string, variable, list	You can include this property to order the result set by a field name, a bound variable, or a list of names or bound variables.
maxConnections	int	This property can be used to set a limit on the maximum number of active connections to the source. For example, s:maxConnections 16 sets the limit to 16 connections. The default value is 10.
partitionBy	string, variable, list	The GDI attempts to partition queries automatically across the available cores (slices) in AnzoGraph. To determine how to partition the query, the GDI uses metadata from the source database. It looks for any column in an index, preferring the primary key column if it is interpolable. However, it only considers the first column in any index on the table. After determining the partition column, the GDI does a MIN/MAX on the column as well as a basic sizing query. To specify which column or columns the GDI should partition on, you can include the partitionBy property in the query. The property supports a list of source field names, bound variables, or the object s:auto, which forces the GDI to partition the data when the source does not define partitioning metadata.
property	RDF list	<pre>This property can be included to list any JDBC driver- specific connection properties. To incorporate property, use the following syntax: s:property [ s:name "custom_driver_property_name" ;</pre>

Option	Туре	Description
		s:value "custom_value" ]
query string	If you want to access the source data by running an SQL query, you can include this property to specify the query string to run. The language does not have to be SQL if the source supports another language. However, some GDI features where the query is dynamically altered may not work with a non-SQL language. Including {{?variable}} substitutions is supported within s:query strings.	
		<pre>Important If you include s:query, you must also specify table and partitionBy. Specify the table name in s:table and the column to partition the table on in s:partitionBy. If the table and partition column are not specified, the GDI will not partition the query and query execution may fail or perform very poorly.</pre>
schema	string	This property can be included to specify the target schema to query. If you include s:schema "schema_name" without specifying s:table (described below) or s:query, all tables in the schema are queried.
table	string	This property can be included to specify the target table or tables for the query.

#### **FileSource Properties**

The table below lists the properties that are available for queries against file-based data sources. For more information about file sources, see Querying a File Source.

Option	Туре	Description
format	RDF list	You can include the format property to add parameters that describe the source files. See File Source Format Options for details about the supported parameters.
maxDepth	int	This property can be used to limit the directory traversal depth. By default, when s:url specifies a directory (and a s:pattern that limits that traversal depth is not specified), all subdirectories are processed. To process only the files in the top level directory, set maxDepth to 0 (s:maxDepth 0). To process the files in the top level directory plus the first-level subdirectories, set maxDepth to 1 (s:maxDepth 1), and so on.
mimetype	string	This property can be included to specify the MIME type of the data. If you are querying TSV files that do not have a .tsv file extension, include the mimetype property with a value of text/tsv (s:mimetype "text/tsv").
options	RDF list	If additional connection information needs to be provided to access the file storage system, include the options property to list any storage-specific connection parameters. See File Storage Connection Options for information about the supported properties for each storage type.
pattern	string	This property can be used to specify a wildcard pattern for matching file names. For example, s:pattern "common_prefix*.csv". You can include one s:pattern property per FileSource. The GDI supports Unix file globbing syntax outside of parentheses. Within parentheses, full Java regular expression language is supported. For

Option	Туре	Description
		example, including s:pattern "data/**/customer_*.csv" tells the GDI to load all files that match the pattern "customer_*.csv" from any number of subdirectories under the data directory. Similarly s:pattern "(\d+)/transaction_*.csv" tells the GDI to load all files that match the pattern "transaction_*.csv" in all subdirectories.

#### **HttpSource Properties**

The table below lists the properties that are available for queries against HTTP data sources. For more information about HTTP sources, see Querying an HTTP Source.

Option	Туре	Description
authorization	RDF list	This property specifies the type of authorization to use and the values for authentication. The options are BearerToken, AWSSignature, or BasicAuth. s:authorization [ a s:BearerToken   s:AWSSignature   s:BasicAuth ]
AWSSignature	RDF list	<ul> <li>For authorization to AWS service endpoints, specify this property and include the appropriate authentication properties from the list below:</li> <li>accessKey: Include this property to specify the AWS access key.</li> <li>region: Include this property to specify the AWS region.</li> <li>secretKey: Include this property to specify the AWS secret key.</li> <li>serviceName: Include this property to specify the AWS service name.</li> </ul>

Option	Туре	Description
		<pre>• sessionToken: Include this property to specify the AWS session token. s:authorization [ a s:AWSSignature ; s:accessKey "string" ; s:region "string" ; s:secretKey "string" ; s:serviceName "string" ; s:sessionToken "string" ; ]</pre>
BasicAuth	RDF list	<pre>Specify this property when basic authentication is used, and include the username and password properties. s:authorization [ a s:BasicAuth ; s:username "string" ; s:password "string" ; ]</pre>
BearerToken	string	<pre>Specify this property when a bearer token is used for authentication, and include the token property. s:authorization [ a s:BearerToken ; s:token "string" ]</pre>
content	string or RDF list	This property can be included to send content to the source in the body of the request. For example, content can be a SPARQL query, JSON arrays, or a list of key-value pairs. Content can also be configured with an inline object (blank node) that gets translated to JSON. For more information, see Mapping the Content Property to JSON.

Option	Туре	Description
contentType	string	<pre>Include this property to specify the content type of the body of the request. For example, s:contentType "application/sparql-query" or s:contentType "application/json".</pre>
encoding	string	When targeting a file, you can include this property to specify the character encoding used by the file. The default value is s:encoding "utf8".
form	RDF list	To send data to the HTTP endpoint, you can use this property to post the data. Form is a list of name-value pairs. When including s:form, you must also include s:contentType "multipart/form-data". The GDI sends the form object as an application/x-www-form-urlencoded string that contains the specified parameters. See Example form Parameter Usage in Querying an HTTP Source for sample usage.
format	RDF list	If the data is file-based, you can include the format property to add parameters that describe the source. See File Source Format Options for details about the supported parameters.
header	RDF list	You can use this property to specify name-value pairs to include as headers in the request. For example: <pre>s:header [ s:name "Accept" ; s:value "application/json" ]</pre> If you are creating a view, you can include variables in the <pre>s:header list. When another query is run against a view with variables, that query can map the variables through the view by including predicates in the CONSTRUCT clause.</pre>

Option	Туре	Description
method	string	You can include this property to specify the HTTP method. For example, s:method "GET" or s:method "POST".
mimetype	string	You can include this property to specify the MIME type of the source. For example, s:mimetype "text/html".
orderBy	string, variable, list	You can include this property to order the result set by a field name, a bound variable, or a list of names or bound variables.
parameter	RDF list	You can include this property to list any URL parameters as name-value pairs. For example, the s:parameter property below adds format to return results in CSV format and the named-graph-uri parameter to target a specific layer in a graphmart. s:parameter [ s:name "format" ; s:value "csv" ] , [ s:name "named-graph-uri" ; s:value "http://cambridgesemantics.com/Layer/d541" ]
		If you are creating a view, you can include variables in the s:parameter list. When another query is run against a view with variables, that query can map the variables through the view by including predicates in the CONSTRUCT clause.
partitionBy	string, variable, list	The GDI attempts to partition queries automatically across the available cores (slices) in AnzoGraph. To determine how to partition the query, the GDI uses metadata from the source. It looks for any column in an index, preferring the primary key

Option	Туре	Description
		column if it is interpolable. However, it only considers the first column in any index on the table. After determining the partition column, the GDI does a MIN/MAX on the column as well as a basic sizing query. To specify which column or columns the GDI should partition on, you can include the partitionBy property in the query. The property supports a list of source field names, bound variables, or the object s:auto, which forces the GDI to partition the data when the source does not define partitioning metadata.
proxy	string or RDF list	<pre>Include this property to specify proxy information if a proxy is used. The value can be a string, such as s:proxy "host_ url:port_number", or an RDF list that includes host and port properties, such as s:proxy [ s:host "host_url" ; s:port port_number ].</pre>
trust	string	Include this property to set the level of trust for the source's SSL certificate. The value can be either "system" or "all".

#### **ElasticSource Properties**

The table below lists the properties that are available for queries against Elasticsearch data sources. For more information about Elasticsearch sources, see Querying an Elasticsearch Source.

Option	Туре	Description
aggregations	object	You can include this property to calculate aggregations over the specified bindings. For information about aggregations, see <u>Aggregations</u> in the Elasticsearch documentation.
config	string	To enable you to use explicit mappings, you can include this property to specify the URL to the index configuration file to employ. For example, es:config

Option	Туре	Description
		"/opt/shared/elastic/mapping.json".
document	string	This property lists the document(s) to search.
field	string or variable	This property defines the field to operate on. The value can be a string or bound variable.
highlight	RDF list	You can include this property to define how results are highlighted. For information about the available properties, see Highlighting Elasticsearch Results.
html	boolean	This property controls whether to output HTML for highlighted results. Defaults to true.
index	string	This property can be included to specify the indexes to search. Specify multiple indexes in a comma-separated list. For example, es:index "projectA_mar", "projectA_apr" ;.
minScore	float	This property defines the minimum score for matching documents. Documents with a lower score are not included in the search results.
query	string or RDF list	This property defines the query to execute. The value can be a string or a query object that maps to the Elasticsearch Query DSL. To generate the final query, the GDI combines es:query with any filters it can push to the Elasticsearch DSL. For more information about the query property and mapping Elasticsearch filters to SPARQL FILTER clauses, see Query DSL and Filter Mapping.
routing	string	This property can be included to route a document to a specific

Option	Туре	Description
		shard or to limit the search to a particular shard.
searchAfter	RDF list	You can include this property to define the key values to start searching from.
size	int	This property maps to the size parameter in the Elasticsearch Search API and configures the batch size or maximum number of hits to return in a single call. Defaults to 10 and typically does not need to be changed.
source	boolean or RDF list	This property can be included to specify the source data to include in results. The value can be a boolean, list of fields, or a list of variable bindings. When true, all source data is returned. When false, no source data is returned.
url	string	The Elasticsearch endpoint URL.

# **Onboard Unstructured Data**

Anzo processes unstructured data using an Anzo Distributed Unstructured cluster and Elasticsearch. Configurable text analytics and natural language processing (NLP) pipelines find and extract data and convert it to the graph data model. Anzo can process all common file types such as Office documents, PDFs, web pages, and email messages, and can analyze text within Excel, databases, and knowledgebases, or XML columns, properties, and fields. Anzo finds, analyzes, extracts, and ingests concepts, entities, sentiment, topics, classifications, events, facts, and thousands of types of relationships.

In this section:	
Unstructured Onboarding Process Overview	
Creating an Unstructured Pipeline	
Running an Unstructured Pipeline	
Pipeline Settings Reference	
Annotator Settings Reference	

### **Unstructured Onboarding Process Overview**

Anzo onboards unstructured data through pipelines that run in a distributed environment where a cluster of worker nodes process the incoming documents and generate output artifacts. This topic provides an overview of the Anzo Distributed Unstructured (DU) pipeline process and infrastructure.

The diagram below provides a high level overview of the Anzo platform architecture with integration of DU and Elasticsearch. The description below the diagram describes the unstructured data onboarding process and resulting artifacts.



When an unstructured pipeline is run, a crawler service streams data to a pipeline service. The pipeline service reads the stream of files and constructs the appropriate request payloads—one request per document to process. Anzo sends the requests to the DU leader instance, and the leader queues the requests and distributes them to the worker instances to process in parallel. When each worker processes a document, it creates a temporary output artifact on the shared file system. The artifact includes the following items:

- An RDF file that describes the text annotations and general metadata about the processed document.
- A binary store artifact for Anzo.
- A JSON artifact that contains a reference to the extracted text of the document. Elasticsearch uses this artifact to generate the document index.

When the DU workers have processed all of the documents, Anzo completes the following postprocessing steps:

- Consolidate the RDF artifacts from the workers and create a file-based linked data set (FLDS) for loading to AnzoGraph.
- Read the JSON artifacts and instruct the Elasticsearch server to build an index with the text extracted from the documents. A snaphsot of the index is saved on the file system with the FLDS. Any time a graphmart that includes that FLDS is loaded to an AnzoGraph instance, Anzo loads the corresponding snapshot into the Elasticsearch server that is associated with the AnzoGraph connection.

When the post-processing is finished, the pipeline service finalizes the FLDS metadata to store in its catalog. The new unstructured data set becomes available in the Datasets catalog, and it can be added to a graphmart and loaded to AnzoGraph for use in Hi-Res Analytics dashboards.

### **Creating an Unstructured Pipeline**

Follow the instructions below to create and run a new unstructured pipeline.

- 1. Create the Pipeline
- 2. Add Crawlers to the Pipeline
- 3. Add Annotators to the Pipeline
- 4. Run the Pipeline

### **Create the Pipeline**

1. In the Anzo application, expand the **Onboard** menu and click **Unstructured Data**. Anzo displays the Pipelines screen, which lists any existing unstructured pipelines. For example:

Components	Pipelines				
Y Q Search	= • 🔳 🔳			Add Unstructured Pip	eline 👻
🗌 🕶 Title	↑ Description	Updated Date	Tags	Actions	÷
A: Docs		Nov 14, 2022		Д	:
fic: Guide		Nov 14, 2022		Д	:
			Rows per page: 25	5 ¥ 1-2 of 2 <	>

2. Click the **Add Unstructured Pipeline** button and select **Distributed Unstructured Pipeline**. Anzo opens the Create Distributed Unstructured Pipeline dialog box. For example:

Create Distributed Unstructured Pipeline
Title *
Description
Target Anzo Data Store *
The datasource to use for autocreating linked datasets from this pipeline Deploy Unstructured Infrastructure Dynamically
Static Elastic Search Config 🛛 🗸
A static elastic search config to use in post processing. If none is provided, then user will be prompted to pick a cloud location for dynamic ES spinup while triggering pipeline.
CANCEL SAVE

3. In the **Title** field, type a name for the pipeline.

#### Note

The title serves as a key to identify this pipeline and its corpus in multiple contexts. Specify a title that is unique and stable. The pipeline's corpus dataset name is derived from this value.

- 4. Type an optional description for the pipeline in the **Description** field.
- 5. If necessary, click the **Target Anzo Data Store** field and select the Anzo Data Store for this pipeline.
- If the environment is configured for dynamic Kubernetes-based deployments of the infrastructure, select the Deploy Unstructured Infrastructure Dynamically checkbox and leave the Static Elasticsearch Config field blank.
- 7. If necessary, click the **Static Elasticsearch Config** field and select the Elasticsearch connection to use for this pipeline. If you use dynamic deployments to deploy Elasticsearch instances on-demand, leave this field blank. Anzo prompts the user to choose a Cloud Location when the pipeline is run.

8. Click **Save** to create the pipeline. Anzo displays the pipeline Overview screen. For example:

Anzo Docs Net Vesioned 🖄	Д	
Overview Crawlers Annotators History Versions Discussion	>	
Description         General           PDF versions of the documentation         Type         Unstructured Pipeline           Target Arao Data Store         Creator         System Administrator           Store         Updated         a few seconds ago           Deploy Unstructured Infrastructure Dynamically         Released         an hour ago		
Static Elastic Search Config ES http://cambridgesemantics.com/Distribute	http://cambridgesemantics.com/DistributedU	
+ Advanced Dynamic Resource Log	Dynamic Resource Log	
None		

#### Note

A pipeline saves automatically and constantly undergoes validation to make sure that it is valid based on the current configuration. Anzo displays validation issues in red on the top of the screen. The warnings will disappear as you add components to the pipeline.

- 9. If necessary, click **Advanced** to configure the advanced pipeline settings. For details about the advanced settings, see Pipeline Settings Reference.
- 10. Next, follow the instructions in Add Crawlers to the Pipeline to add one or more crawlers to the pipeline.

### Add Crawlers to the Pipeline

After creating a pipeline, the next step is to add one or more crawlers. Crawlers determine what text to process.
1. In the pipeline, click the **Crawlers** tab.



 Next, click the Add Input button. Anzo displays the Add Component dialog box. The New tab is selected and lists all available crawlers. The Existing Components tab lists crawlers that have been previously configured for other pipelines.

Add Con	nponent
N	iew Existing Components
<b>Y</b> Q	Search Fr *
	File Based Dataset Crawler
	A crawler for the RDF of a File Based Dataset.
	Filesystem Crawler
	A crawler of unstructured documents on File Systems
	Graphmart RDF Crawler
	A crawler for RDF in an Online Graphmart.
	Local Volume Dataset Crawler
	A crawler for the RDF of a Local Volume Dataset.
	•
	Rows per page: 25 + 1-4 of 4 < >
	CANCEL OK

- 3. To add a new crawler, select the crawler. To add an existing crawler, click the **Existing Components** tab and select a crawler. The list below describes each of the crawlers:
  - File Based Dataset Crawler: Include this crawler to process data from a file-based linked data set (FLDS) on a file store.
  - Filesystem Crawler: Include this crawler to process documents, such as email messages, PDF, XML, PowerPoint, Excel, OneNote, or Word files, and images, that are available on a file store.

- **Graphmart RDF Crawler**: Include this crawler to process RDF in an online graphmart or specific data layer.
- Local Volume Dataset Crawler: Include this crawler to process RDF data that is stored as a linked data set (LDS) in an Anzo journal.
- 4. After selecting a crawler, click **OK**. Anzo opens the Create dialog box for that crawler so that you can configure it. Click a crawler name in the list below to view the details for that component:
  - File Based Dataset Crawler
  - Filesystem Crawler
  - Graphmart RDF Crawler
  - Local Volume Dataset Crawler

#### File Based Dataset Crawler

Create File Based Dataset Crawler		
		Î
Title *		.
Description		
Backing Dataset *	× ~	
The backing dataset		
Backing Ontology *	× ~	
The backing ontology		
RDF Resource Type	~	
RDF Class		
Link Property	~	
Property(s) to match on for filepath reference		
	CANCEL SAV	E

- Title: Required field that specifies the unique name for this crawler.
- **Description**: Optional field that provides a description of this crawler.

- Backing Dataset: Required field that specifies the Anzo dataset to crawl.
- Backing Ontology: Required field that specifies the model for the dataset.
- **RDF Resource Type**: Required field that specifies the resource type or class of data to target with this crawler.
- Link Property: Optional field that specifies any link properties to crawl. A link property is a property whose value identifies the location of a linked document. When linked properties are specified, the crawler will crawl the linked documents. For example, in the triples below, fileLocation is a link property:

```
<urn://someUnstructuredDocument> <urn://someProperty> "file metadata" ;
<urn://fileLocation> "/path/to/file.pdf" .
```

### Note

In typical use cases, this crawler is configured to define either a Link Property or a Content Property but not both.

Content Property: Optional field that identifies any content properties to crawl. A
content property is a property whose value is a string literal and you want the crawler to
crawl and annotate those strings. For example, in the triples below, longDescription is
a content property:

```
<urn://someUnstructuredDocument> <urn://someProperty> "file metadata" ;
<urn://longDescription> "this is some interesting, likely long,
unstructured text
with a lot of information, and I want it to be annotated" .
```

- **Base Path Connection**: Required field whose value depends on whether you specified a Link Property or a Content Property:
  - If a Link Property was specified, the Base Path Connection is the base path to use for resolving relative file paths in the Link Property values. For example, using the example triples:

```
<urn://someUnstructuredDocument> <urn://someProperty> "file metadata"
;
<urn://fileLocation> "/path/to/file.pdf" .
```

The <urn://fileLocation> value of /path/to/file.pdf could be a relative path to a location like s3://location/bucket/path/to/file.pdf or /opt/anzoshare/data/path/to/file.pdf. Therefore, the Base Path needs to be specified to resolve any relative paths and locate the linked documents.

 If a Content Property was specified, the Base Path Connection is a directory on the file store where the crawler can save a copy of the Content Property strings for the Anzo Unstructured worker instances. Saving the content to a shared file location avoids the overhead of sending the strings to the workers over the network.

### **Filesystem Crawler**

Create Filesystem Crawler		
Title *		
Description		
File Crawl Location *		BROWSE
The file location to crawl		
Crawl subfolders		
	CANCEL	SAVE

- Title: Required field that specifies the unique name for this crawler.
- Description: Optional field that provides a description of this crawler.
- File Crawl Location: Required field that specifies the file system crawl location. Click the field to open the File Location dialog box:

File Location		
Current Folder		GO
Selected: None		CLEAR ALL
Server Shared Filesystem	noot	
sysadmin User Folder	books books	
	crawl	
	csv	
	datafox	
CREATE NEW FOLDER		CANCEL OK

On the left side of the screen, select the storage location for the files to crawl. On the right side of the screen, navigate to the directory that contains the files. Select a directory, and then click **OK**.

 Crawl subfolders: Optional field that specifies whether to crawl the subdirectories under the VFS Crawl Location. To crawl the subdirectories, select the Crawl subfolders checkbox. To ignore subdirectories, clear the Crawl subfolders checkbox.

#### **Graphmart RDF Crawler**

Create Graphmart RDF Crawler		
		í
Title *		
Description		
Backing Graphmart	[ ~	
A backing graphmart		
Backing Layer	~	
A backing layer		
Backing Ontology *	X ¥	
The backing ontology		
RDF Resource Type		
RDF Class	CANCEL SAVE	

- Title: Required field that specifies the unique name for this crawler.
- Description: Optional field that provides a description of this crawler.
- Backing Graphmart: Optional field that specifies the graphmart to crawl. To configure the grawler to crawl at the graphmart level, select one or more graphmarts in the Backing Graphmart field and leave the Backing Layer field blank.
- Backing Layer: Optional field that specifies the data layer or layers that you want the pipeline to crawl. To crawl specific layers and not an entire graphmart, make sure that you leave the Backing Graphmart field blank and select the layers to crawl in the Backing Layer field. If you specify both a Backing Graphmart and a Backing Layer, the Backing Graphmart value supersedes the Backing Layer value, resulting in the entire graphmart being crawled.
- **Backing Ontology**: Required field that specifies the model for the Backing Graphmart or Data Layer.
- **RDF Resource Type**: Required field that specifies the resource type or class of data to target with this crawler.
- Link Property: Optional field that specifies any link properties to crawl. A link property is a property whose value identifies the location of a linked document. When linked properties are specified, the crawler will crawl the linked documents. For example, in the triples below, fileLocation is a link property:

```
<urn://someUnstructuredDocument> <urn://someProperty> "file metadata" ;
<urn://fileLocation> "/path/to/file.pdf" .
```

#### Note

In typical use cases, this crawler is configured to define either a Link Property or a Content Property but not both.

• **Content Property**: Optional field that identifies any content properties to crawl. A content property is a property whose value is a string literal and you want the crawler to

crawl and annotate those strings. For example, in the triples below, **longDescription** is a content property:

```
<urn://someUnstructuredDocument> <urn://someProperty> "file metadata" ;
<urn://longDescription> "this is some interesting, likely long,
unstructured text
with a lot of information, and I want it to be annotated" .
```

- Base Path Connection: Required field whose value depends on whether you specified a Link Property or a Content Property:
  - If a Link Property was specified, the Base Path Connection is the base path to use for resolving relative file paths in the Link Property values. For example, using the example triples:

```
<urn://someUnstructuredDocument> <urn://someProperty> "file metadata"
;
<urn://fileLocation> "/path/to/file.pdf" .
```

The <urn://fileLocation> value of /path/to/file.pdf could be a relative path to a location like s3://location/bucket/path/to/file.pdf or /opt/anzoshare/data/path/to/file.pdf. Therefore, the Base Path needs to be specified to resolve any relative paths and locate the linked documents.

 If a Content Property was specified, the Base Path Connection is a directory on the file store where the crawler can save a copy of the Content Property strings for the Anzo Unstructured worker instances. Saving the content to a shared file location avoids the overhead of sending the strings to the workers over the network.

## Local Volume Dataset Crawler

Create Local Volume Dataset Crawler		
		î
Title *		
Description		
Backing Dataset *	× ×	
The backing dataset		
Backing Ontology *	x ~	
The backing ontology		
RDF Resource Type		
RDF Class		
Link Property	~	
Property(s) to match on for filepath reference		•
	CANCEL SAVE	

- Title: Required field that specifies the unique name for this crawler.
- Description: Optional field that provides a description of this crawler.
- Backing Dataset: Required field that specifies the Anzo dataset to crawl.
- Backing Ontology: Required field that specifies the model for the dataset.
- **RDF Resource Type**: Required field that specifies the resource type or class of data to target with this crawler.
- Link Property: Optional field that specifies any link properties to crawl. A link property is a property whose value identifies the location of a linked document. When linked properties are specified, the crawler will crawl the linked documents. For example, in the triples below, fileLocation is a link property:

```
<urn://someUnstructuredDocument> <urn://someProperty> "file metadata" ;
<urn://fileLocation> "/path/to/file.pdf" .
```

#### Note

In typical use cases, this crawler is configured to define either a Link Property or a Content Property but not both.

Content Property: Optional field that identifies any content properties to crawl. A
content property is a property whose value is a string literal and you want the crawler to
crawl and annotate those strings. For example, in the triples below, longDescription is
a content property:

```
<urn://someUnstructuredDocument> <urn://someProperty> "file metadata" ;
<urn://longDescription> "this is some interesting, likely long,
unstructured text
with a lot of information, and I want it to be annotated" .
```

- **Base Path Connection**: Required field whose value depends on whether you specified a Link Property or a Content Property:
  - If a Link Property was specified, the Base Path Connection is the base path to use for resolving relative file paths in the Link Property values. For example, using the example triples:

```
<urn://someUnstructuredDocument> <urn://someProperty> "file metadata"
;
<urn://fileLocation> "/path/to/file.pdf" .
```

The <urn://fileLocation> value of /path/to/file.pdf could be a relative path to a location like s3://location/bucket/path/to/file.pdf or /opt/anzoshare/data/path/to/file.pdf. Therefore, the Base Path needs to be specified to resolve any relative paths and locate the linked documents.

- If a Content Property was specified, the Base Path Connection is a directory on the file store where the crawler can save a copy of the Content Property strings for the Anzo Unstructured worker instances. Saving the content to a shared file location avoids the overhead of sending the strings to the workers over the network.
- 5. When you have finished configuring the crawler, click **Save**. Anzo adds the crawler to the pipeline and returns to the Crawlers screen. For example:

🔺 At le	ast one Annotator is rec	ommended for a Pipelin					
	Anzo Docs					Run Pipeline	Д
	Overview	Crawlers	Annotators	History	Versions	Discussion	>
Add In Crawl	put Input data sour	rce options			/	Enter a single, or multipl sources and keywords o extensions for the data are looking for and how deliver it in the fields to left.	le you to the

- 6. If you want to change the crawler configuration, click the Edit icon (✓) for the crawler and modify the settings as needed. If you want to add another crawler to the pipeline, repeat the steps above.
- 7. When you have finished adding crawlers, follow the instructions in Add Annotators to the Pipeline to add one or more annotators to the pipeline.

## Add Annotators to the Pipeline

After adding crawlers, the next step is to add one or more annotators. Annotators extract facts or references in the text as annotations.

1. In the pipeline, click the **Annotators** tab.



2. Next, click the **Add Output** button. Anzo opens the Add Component dialog box. The **New** tab is selected and lists the available annotators and the **Existing Components** tab lists

annotators that have been previously configured for other pipelines.

Add Con	nponent
ħ	ew Existing Components
<b>T</b> Q	Search 🖉 🖛 👻
	Custom Relationship Annotator
	Annotates based on custom relationships between concepts found via the other annotators in the pipeline.
	External Panelas Associates
	External Service Annotator
	Keyword and Phrase Annotator
	Annotates based on fixed keywords and phrases
	Knowledgebase Annotator
	Annotates based on existing datasets within Anzo
	Rows per page: 25 👻 1.5 of 5 < >
	CANCEL OK

- 3. To add a new annotator to the pipeline, click the annotator name to select it. To add an existing annotator to the pipeline, click the Existing Components tab, and then select an annotator. The list below describes each of the default annotators:
  - **Custom Relationship Annotator**: Include this annotator to map relationships between annotations based on the number of characters between the annotations.
  - External Service Annotator: Include this annotator to hit an HTTP endpoint that provides annotations.
  - Keyword and Phrase Annotator: Include this annotator to create annotations based on the phrases that you specify.
  - Knowledgebase Annotator: Include this annotator to link structured and unstructured data by finding instances in data layers, graphmarts, or Anzo linked datasets. Based on the names and aliases of entities present or patterns that are indicative of the entities, this annotator marks up the documents with the structured entities linked.
  - **Regex Annotator**: Include this annotator to use regular expression rules to identify entities such as email addresses, URLs, phone numbers, or any other entity that can be matched using a regular expression.

- 4. After selecting an annotator, click **OK**. Anzo opens the Create dialog box for the component. Complete the fields to configure the annotator. The list below provides details about the settings for the annotators that are typically used in pipelines. Click an annotator name to view the details for that component:
  - External Service Annotator
  - Keyword and Phrase Annotator
  - Knowledgebase Annotator
  - Regex Annotator

#### **External Service Annotator**

Create External Service Annotator		
Title *		
Description		
UTTO Desured Canfe #		
Config for connecting and sending a request to the external NLP server		1 ~
Document ID Response Path *		
the Document ID path for entities returned in the service response		
Entity Name Path *		
The entity Name path for entities returned in the service response		
Entity Class Path *		
The entity Class path for entities returned in the service response		
	CANCEL	SAVE

## Тір

For information about the options that are presented when you edit an External Service Annotator, see Annotator Settings Reference.

- Title: Required field that specifies the unique name for this annotator.
- Description: Optional field that provides a description of this annotator.

- **HTTP Request Config**: Required field that specifies the HTTP source object that contains the URL and method to use when sending data for annotations.
- **Document ID Response Path**: Required field that specifies where to find the document ID in the response.
- Entity Name Path: Required field that specifies the annotation object name path.
- Entity Class Path: Required field that specifies the class URI for an annotation.

#### **Keyword and Phrase Annotator**

Create Keyword and Phrase Annotator		
Title *		
Description		
		ADD
Phrase *		
Phrase to look for		
	CANCEL	SAVE

## Tip

For information about the options that are presented when you edit a Keyword and Phrase Annotator, see Annotator Settings Reference.

- Title: Required field that specifies the unique name for this annotator.
- **Description**: Optional field that provides a description of this annotator.
- **Phrase**: Required field that specifies the terms or phrases to annotate. Type a word or phrase in the field and then click **Add** to add the phrase. You can add any number of phrases.

## **Knowledgebase Annotator**

Create Knowledgebase Annotator		
		Î
Title *		
Description		
Backing Graphmart		
A backing graphmart		
Backing Layer	×	
A backing layer		
Backing Ontology *	× ×	1
The backing ontology		
Term Class	~	
The owl:Class of the knowledge base terms		*
	CANCEL SAVE	

## Тір

For information about the options that are presented when you edit a Knowledgebase Annotator, see Annotator Settings Reference.

- Title: Required field that specifies the unique name for this annotator.
- **Description**: Optional field that provides a description of this annotator.
- **Backing Graphmart**: Optional field that specifies the graphmart or graphmarts to annotate.

### Note

If you want the annotator to run against a linked dataset or Anzo knowledgebase instead of a data layer or graphmart, leave the Backed Layer and Backed Graphmart fields blank. After saving the pipeline, you can edit the pipeline and specify a **Backed Dataset** at that time. • Backing Layer: Optional field that specifies the data layer or layers to annotate.

#### Note

The Backing Layer and Backing Graphmart fields are treated independently. Layers that you select do not have to be part of the graphmart that you specify in **Backing Graphmart**. And specifying a layer does not mean that you must select a Backing Graphmart. However, any layers or graphmarts that you select must contain classes and properties from the **Backing Ontology** or the data will not be annotated.

- **Backing Ontology**: Required field that specifies the model for the backing data layers and/or graphmart. Click the field and select a model from the drop-down list.
- Term Class: Required field that specifies the class of data for the annotations.
- **Term Label Property**: Required field that lists the primary name or label property of the resources.
- **Term Identifying Properties**: Required field that specifies the properties that contain names, aliases, or other identifiers to use for identifying the resources.

### **Regex Annotator**

Create Regex Annotator		
Title *		
Description		
Regular Expression Rule *		~
nine regional expression rune(a) onar ons attitio(ator will use	CANCEL	SAVE

## Tip

For information about the options that are presented when you edit a Regex Annotator, see Annotator Settings Reference.

- Title: Required field that specifies the unique name for this annotator.
- Description: Optional field that provides a description of this annotator.
- Regular Expression Rule: Required field that lists the regular expression rules for this annotator. To add a rule, click drop-down field and select Create New. Anzo opens the Create Regular Expression Rule dialog box where you can define the rule:

Create Regular Expression Rule	
Title *	
Class Structure *	
Description	
Description	 
Description Regular Expression *	 
Description Regular Expression * The regular expression to look for in the text	

- <sup>o</sup> **Title**: Required field that specifies the name of the rule.
- Class Structure: Required field that specifies the class in the model that should be created for this rule. The value should be in the format group\_number:class\_name, where group\_number corresponds to a group in the regex capture. Each rule should start with group 0. Include groups 1 and higher if needed to represent parts of the expression that are contained in parentheses. The class\_name is a label that describes the type of data the rule will find. For example, for a rule that finds hyphenated words 0:Hyphens.
- <sup>o</sup> **Description**: Optional field that describes the rule.
- Regular Expression: Required field that specifies the regular expression to use for finding matching entities.

5. When you have finished configuring the annotator, click **Save**. Anzo adds the annotator to the pipeline and returns to the Annotators screen. For example:

	Anzo Docs					Run Pipeline
	Overview	Crawlers	Annotators	History	Versions	Discussion >
Add Outpu	ut Output data se	ource options				0
Find phra	ises				/ 1	Enter a single, or multiple sources and keywords or extensions for the data you are looking for and how to deliver it in the fields to the left.

- 6. If you want to change the annotator configuration, click the Edit icon (✓) for the annotator and modify the settings as needed (see Annotator Settings Reference for information about settings). If you want to add another annotator to the pipeline, repeat the steps above.
- 7. When you have finished adding annotators to the pipeline, proceed to Run the Pipeline below.

## **Run the Pipeline**

When you are ready to run the pipeline, click the **Run Pipeline** button on the top right of the screen. For example:

l.	Anzo Docs				I	Run Pipeline	Д
	Overview	Crawlers	Annotators	History	Versions	Discussion	•
Descriptio PDF ver Target And Store	n sions of the documer 20 Data Store eploy Unstructured In	ntation frastructure Dynamic	cally		General Type Uni Creator Sys Updated a fe Released an I	structured Pipeline tem Administrator w seconds ago hour ago	
Static Elar ES	stic Search Config				http://cambridgesema	ntics.com/Distributed	0
+ Advar	nced				Dynamic Resource	Log	
					Tags None		

The process can take several minutes to complete. You can click the **Progress** tab to view details such as the pipeline status, runtime, number of documents processed, and errors. For example:

	Crawlers	Annotators	History	Progress	Versions	Discussion
0 Errored	0 Cancelled	0/2 Completed	6min 52s Elapsed Time	09:30: Start	14 AM	N/A Finish Time
0% Finished Processing Pending Writ Indexing	a Dets per Minuce	4 5 2 3 0 09:31:00	083200 093300	09.34:00 Timestamp	09:35:00	09:36:00 09:37
Cancelled Error			+	Processed		
All (2) PendingWrite (1 C Search	) Processing (1)	<b>III =</b>	•	Processed		
Concelled     Error	) Processing (1)	Document Size	Elapsed Time(in sec)	Processed Start Time	Finish Time	Actions
Cancelled     Error	) Processing (1) Status Processing	Document Size	Elapsed Time(in sec) 394.651	Processed Start Time 09:0025 AM	Finish Time	Actions

When the pipeline finishes, a new dataset becomes available in the Datasets catalog. From the catalog, you can create a graphmart from the dataset so that you can explore and analyze the data. For instructions, see Creating a Graphmart from a Dataset. You can also add the dataset to an existing graphmart by following the steps in Adding a Dataset to a Graphmart.

# **Running an Unstructured Pipeline**

This page provides instructions for running an unstructured pipeline.

1. In the Anzo application, expand the **Onboard** menu and click **Unstructured Data**. Anzo

displays the Pipelines screen, which lists any existing unstructured pipelines. For example:

Components	Pipelines					
▼ Q Search	= •			Add Unstruct	tured Pipe	eline 👻
Title	↑ Description	Updated Date	Tags		Actions	÷
R: Docs		Nov 14, 2022			Д	:
R: Quide		Nov 14, 2022			Д	:
			Rows per page: 2	25 👻 1-2 of 2	<	>

2. Click the name of the pipeline that you want to run. Anzo displays the pipeline Overview screen. For example:

le:	Anzo Docs					Run Pipeline
	Overview	Crawlers	Annotators	History	Versions	Discussion >
Description PDF ver Target And Store	n sions of the docume 20 Data Store eploy Unstructured I	ntation nfrastructure Dynamica	slly		General Type Uns Creator Sys Updated a fe Released an t	tructured Pipeline tem Administrator w seconds ago lour ago
Static Elec ES	stic Search Config				http://cambridgesemar	ntics.com/DistributedU 📋
+ Advar	nced				Dynamic Resource Tags None	Log

3. Click **Run Pipeline** to run the pipeline.

The process can take several minutes to complete. You can click the **Progress** tab to view details such as the pipeline status, runtime, number of documents processed, and errors. For example:

	Crawlers	Annotators	History	Progress	Versions	Discussion >
0 Errored	0 Cancelled	0/2 Completed	6min 52s Elapsed Time	09:30: Start	14 AM	N/A Finish Time
0% Finished Processing Pending Wri Indexing	te Dors per Minue	0 09:31:00	0832.00 0933.00	09:34:00 Timestamp	09:35:00	09:36:00 09:37:00
<ul> <li>Queued</li> <li>Cancelled</li> <li>Error</li> </ul>			+	Processed		
All (2) PendingWrite (1	1) Processing (1)	m =	+	Processed		
All (2) PendingWrite ( Queved Error All (2) PendingWrite ( C Search Itle	1) Processing (1)	Document Size	Elapsed Time(in sec)	Processed	Finish Time	Actions
Charceled     Canceled     Canceled     Error	1) Processing (1) Transformed and the second secon	Document Size	Elapsed Time(in sec) 394.651	Processed Start Time 09:0025 AM	Finish Time	Actions :

If this is the first time the pipeline was run, a new dataset becomes available in the Datasets catalog. From the catalog, you can create a graphmart from the dataset so that you can explore and analyze the data. For instructions, see Creating a Graphmart from a Dataset. If the pipeline was run previously, the existing dataset was updated and you can refresh or reload the graphmart that contains the dataset to make the new data available for analytics.

# **Pipeline Settings Reference**

The table below defines the Advanced settings that are available on the Overview tab when viewing an unstructured pipeline.

Setting	Description
Append Timestamp	Controls whether to add a timestamp to unstructured document URIs. This setting is enabled by default.
Diagnostic Logging	Controls whether verbose diagnostic logging is enabled for the pipeline. This setting is disabled by default. When enabled, debug-level logging is performed for the duration of the pipeline.
Current Pipeline Run	This setting is a pointer to the pipeline run object that tracks the ongoing execution of the pipeline.
Pipeline Network Connection	This setting specifies the network connection configuration to be used by the pipeline's worker nodes to connect to the Anzo server. If not specified, this setting defaults to the Unstructured Cluster connection configuration.
Persist Extracted Text	Controls whether to persist the extracted text from documents. This setting is enabled by default.
Persist HTML	Controls whether to persist the extracted highlighted/annotated HTML from documents. This setting is enabled by default.
Persist Original Binary	Controls whether to persist the binary from the original documents. This setting is enabled by default.
Persist Hit Spans	Controls whether to persist the hit spans for the annotations of unstructured documents. This setting is disabled by default.
Persist Nothing	Controls whether RDF data about the documents or annotations are saved or

Setting	Description
	persisted. This setting is disabled by default.
Skip Elastic Search Indexing	Controls whether to skip Elasticsearch indexing. This setting is disabled by default.
Skip Elastic Search JSON creation	Controls whether to skip creating Elasticsearch JSON. This setting is disabled by default.
ls Corpus Cumulative	Controls whether to add the components of each pipeline run to the working edition of the dataset. This setting is disabled by default.
Skip Text Extraction	Controls whether to skip text extraction. This setting is disabled by default.
Delete Elastic Search JSON files	Controls whether to delete the Elasticsearch JSON files after they are indexed. This setting is enabled by default.
Allow Empty Documents	Controls whether to allow documents that have no text to proceed through the pipeline. This setting is disabled by default.
Archive and Host Content	Controls whether to download, cleanse, encapsulate, archive, and host complete document content with inline artifacts. This setting is enabled by default.
HTTP Fetch in Archive	Controls whether the archiving process should resolve and download HTTP URLs that are specified in documents. This setting is disabled by default.
Corpus Linked Dataset	Specifies the FLDS used for documents and annotations from this pipeline. This setting defaults to the name of the pipeline.

Setting	Description
Corpus Name	Specifies the name of the corpus (collection of documents) for the pipeline.
Phase Status Persistence	Specifies how phase status metadata is persisted for each document in the pipeline.
Write Status Updates to Jnl	Controls whether status updates for pipeline runs are written to the journal. This setting is enabled by default.
Write Status Updates to FLDS	Controls whether status updates for pipeline runs are written to an FLDS. This setting is disabled by default.
Write Original Binary On Timeout	Controls whether the original binary is written if the pipeline times out or errors. This setting is disabled by default.
RamDisk Directory Location	Specifies an optional RamDisk base directory to create temporary files under. Using a RamDisk may speed up the pipeline.
Use File Name as Document Title	Controls whether to use the file's name on disk as the document title. This setting is disabled by default.
RDF Statement Buffer Size	Specifies the maximum number of statements to buffer before writing. The default value is 10,000.
RDF File Statement Count	Specifies the maximum number of statements to include in each RDF output file.
Batch Size	Specifies the number of documents to include in one batch.
Maximum	Specifies the maximum number of issues that can be encountered in a run of

Setting	Description
Allowed Session Issues	this pipeline before failing the pipeline.
UI Update Interval (in milliseconds)	The interval of time to wait between running queries to update the data on the pipeline Progress screen. The default value is 30,000 milliseconds (30 seconds).
Document Processing Timeout	Specifies the timeout in milliseconds for each document batch to be processed. Leave this value unset (or set it to 0) to use the microservice cluster's default timeout value.
Error On No Documents Found	Controls whether to fail the pipeline if no documents are found. This setting is enabled by default.
Maximum Pipeline Run Status Journals	Specifies the maximum number of pipeline run status journals to keep before aging them off to an FLDS. By default, only the status of the most recent run of a pipeline remains stored in a status journal. All previous reports are automatically converted to an FLDS and the original status journal is deleted.
Elastic Search Bulk Actions	Specifies the maximum number of indexing actions to queue during Elasticsearch indexing. The default value is 2,000.
Elastic Search Bulk Size	Specifies the maximum size of the document queue during Elasticsearch indexing. The default value is 5.
Elastic Search Bulk Concurrent Requests	This setting specifies the maximum number of concurrent bulk requests to allow during Elasticsearch indexing. The default value is 1.
Elastic Search Bulk Max	This setting specifies the maximum number of threads to use for Elasticsearch indexing. The default value is 1.

Setting	Description
Threads	
Elastic Search Mapping	This setting specifies (in JSON format) the mapping to use when indexing unstructured documents in Elasticsearch.
Elastic Search Pipeline Configuration	This setting specifies (in JSON format) the Elasticsearch pipeline configuration to use when indexing unstructured documents.
Elastic Search Directory Write- all	Controls whether to give write-all permission to the $\tt esi$ directory in the output corpus FLDS.
Elasticsearch Index Settings	This setting specifies (in JSON format) the index settings to use when indexing unstructured documents in Elasticsearch.
Skip Teardown Of Dynamic Resources	Controls whether dynamic K8s-based resources associated with the pipeline are left running after the pipeline is complete. This setting is disabled by default. Enabling it can result in increased cloud resource usage.
Default Finish Pending Writes On Pipeline Cancellation	Controls whether to finish any pending writes for documents during pipeline cancellation. There is a flag in the cancellation request that can be used to override this setting. This setting is enabled by default.
Post-persist Postprocessor	Specifies any post-persist semantic postprocessors in the pipeline.
Rich Text Extractor	Lists the HTML extractors to use in the pipeline.
Post Worker	Specifies a service to invoke on documents after they are successfully

Setting	Description
Service	processed by the pipeline worker processes.
Pre-persist Postprocessor	Specifies any pre-persist semantic postprocessors in the pipeline.
Status Journal Base Path	Specifies the base path for storage of the status journal. By default, status journals are written to a status_journals subdirectory in the Anzo Data Store that is specified for the pipeline.
Content Transformer	Specifies any content transformation and metadata extraction components to use in the pipeline.
Document Crawler Thread Count	Specifies the number of threads to use for document crawling. The default value is 4.
Worker Service ID	Specifies the worker service ID to send requests to. If not specified, the default is pipelineWorkerService.

## **Annotator Settings Reference**

When you edit an existing annotator, additional options become available for refining the annotation criteria or customizing the generated model or instance data. This topic describes the advanced settings that are available when editing each type of annotator.

- External Service Annotator
- Keyword and Phrase Annotator
- Knowledgebase Annotator
- Regex Annotator

## **External Service Annotator**

The table below defines the settings that are displayed when an External Service Annotator is edited.

Setting	Description
Title	Required field that specifies the unique name for the annotator.
Description	Optional field that provides a description of the annotator.
HTTP Request Config	Required field that specifies the HTTP source object that contains the URL and method to use when sending data for annotations.
Document ID Response Path	Required field that specifies where to find the document ID in the response.
Entity Name Path	Required field that specifies the annotation object name path.
Entity Class Path	Required field that specifies the base class URI for an annotation.
Result Path Root	The path to the object that contains the annotation results.

Setting	Description
Store NLP Service Response	Controls whether the service's response is stored in the binary store.
Result Field Path	The external NLP-specific result configuration for returned entities.
Socket Timeout	Specifies the socket timeout (in milliseconds) to use for requests against the source.
Entity Snippet Path	The snippet path for entities returned in the service response.
Entity End Offset Path	The end text offset location in the document for entities returned in the service response.
Entity Begin Offset Path	The start text offset location in the document for entities returned in the service response.
Entity Span Path	The text offset location in the document for entities returned in the service response.
Entity Text Path	The text path for entities returned in the service response.
Entity ID Path	The ID path for entities returned in the service response.
Document ID Request Field	The Document ID parameter for the external service.
Class Name Property	Specifies an annotation property whose value you want to map to the name of the class. For example, if a Category property has the value Disease and you want the name of the class to be "Disease," add Category to this field. When Class Name Property is not defined, the class name is auto-generated.

Setting	Description
Unintended Property Names	A list of any property names to filter out. Type a name in the field and then click <b>Add</b> to add the value.
Unintended Classes	A list of any classes to filter out. Type a class in the field and then click <b>Add</b> to add the value.
Unintended Instances	A list of any entities or instances of the class to filter out. Type an instance in the field and then click <b>Add</b> to add the value.
Is Combine Annotation Instances	Controls whether to combine multiple instances of an extraction into one annotation.
Create Additional General Annotation Type	Controls whether the annotator creates a general shared annotation type in addition to the specific annotation types that are created.
Explicit Property Datatypes/Objecttypes	A list of keys that map property names to a particular object property type.
Output the Detections	Controls whether to include specific detections as an annotation property.
Unintended Property Values	A list of any property values to filter out. Type a value in the field and then click <b>Add</b> to add the value.
Use General Annotator Name in Ontology URI	Controls whether to use ExternalServiceAnnotator or the specific annotator name in the ontology URI.
Entity URI Property	Specifies an annotation property whose value you want to map to the URIs for instances of the class. For example, if a Disease_ID property has the value http://example.com/Asthma and you want to use http://example.com/Asthma as the base URI for instances of the class, add Disease_ID to this field. When Entity

Setting	Description
	URI Property is not defined, the URI is auto-generated based on the name.
Entity Name Property	Specifies an annotation property whose value you want to map to the names for instances of the class. For example, if a Preferred_ Label property includes disease names and you want to use those label values as the names for instances of the Disease class, add Preferred_Label to this field. When Entity Name Property is not defined, the name is auto-generated.
Domain Object Base Class URI	If creating a general annotation type (Create Additional General Annotation Type is enabled), this setting specifies the class to use as the base type for the annotator's domain objects.
Class URI Property	Specifies an annotation property whose value you want to map to the class URI in the model. For example, if a Category_ID property has the value http://example.com/Disease and you want to use http://example.com/Disease as the base class URI, add Category_ID to this field. When Class URI Property is not defined, the URI is auto-generated based on the name.
Is Error Fatal	Controls whether to fail the pipeline if this annotator fails to create annotations.

## Keyword and Phrase Annotator

The table below defines the settings that are displayed when a Keyword and Phrase Annotator is edited.

Setting	Description
Only Consider Text	Controls whether to find phrases via a simplified format of the

Setting	Description
	document. Enabling this setting can be beneficial for a document such as a rich HTML file. Enabling this option is less ideal for documents with multibyte characters.
Require Nonstandard Word Boundaries	Indicates whether the specified phrase can be present with or without surrounding character breaks (e.g. for Chinese) or with regex-nonstandard word boundaries (e.g. for Tagalog).
Title	Required field that specifies the unique name for the annotator.
Description	Optional field that provides a description of the annotator.
Phrase	Required field that specifies the terms or phrases to annotate. Type a word or phrase in the field and then click <b>Add</b> to add the phrase. You can add any number of phrases.
Unintended Property Names	A list of any property names to filter out. Type a name in the field and then click <b>Add</b> to add the value.
Create Additional General Annotation Type	Controls whether the annotator creates a general shared annotation type in addition to the specific annotation types that are created.
Entity URI Property	Specifies an annotation property whose value you want to map to the URIs for instances of the class. For example, if a Disease_ID property has the value http://example.com/Asthma and you want to use http://example.com/Asthma as the base URI for instances of the class, add Disease_ID to this field. When Entity URI Property is not defined, the URI is auto-generated based on the name.
Explicit Property Datatypes/Objecttypes	A list of keys that map property names to a particular object property type.

Setting	Description
Entity Name Property	Specifies an annotation property whose value you want to map to the names for instances of the class. For example, if a Preferred_ Label property includes disease names and you want to use those label values as the names for instances of the Disease class, add Preferred_Label to this field. When Entity Name Property is not defined, the name is auto-generated.
Domain Object Base Class URI	If creating a general annotation type (Create Additional General Annotation Type is enabled), this setting specifies the class to use as the base type for the annotator's domain objects.
Unintended Classes	A list of any classes to filter out. Type a class in the field and then click <b>Add</b> to add the value.
Class Name Property	Specifies an annotation property whose value you want to map to the name of the class. For example, if a Category property has the value Disease and you want the name of the class to be "Disease," add Category to this field. When Class Name Property is not defined, the class name is auto-generated.
Unintended Instances	A list of any entities or instances of the class to filter out. Type an instance in the field and then click <b>Add</b> to add the value.
Is Combine Annotation Instances	Controls whether to combine multiple instances of an extraction into one annotation.
Class URI Property	Specifies an annotation property whose value you want to map to the class URI in the model. For example, if a Category_ID property has the value http://example.com/Disease and you want to use http://example.com/Disease as the base class URI, add Category_ID to this field. When Class URI Property is not defined, the URI is auto-generated based on the name.

Setting	Description
Output the Detections	Controls whether to include specific detections as an annotation property.
Unintended Property Values	A list of any property values to filter out. Type a value in the field and then click <b>Add</b> to add the value.
Is Error Fatal	Controls whether to fail the pipeline if this annotator fails to create annotations.

## Knowledgebase Annotator

The table below defines the settings that are displayed when a Knowledgebase Annotator is edited.

Setting	Description
Title	Required field that specifies the unique name for the annotator.
Description	Optional field that provides a description of the annotator.
Backing Graphmart	Optional field that specifies the graphmart or graphmarts to annotate.
Backing Layer	Optional field that specifies the data layer or layers to annotate.
	Note The Backing Layer and Backing Graphmart fields are treated independently. Layers that you select do not have to be part of the graphmart that you specify in Backing Graphmart. And specifying a layer does not mean that you must select a Backing Graphmart. However, any layers or graphmarts that you select must contain classes and

Setting	Description
	properties from the Backing Ontology or the data will not be annotated.
Backing Ontology	Required field that specifies the model for the backing data layers and/or graphmart.
Term Class	Required field that specifies the class of data for the annotations.
Term Label Property	Required field that lists the primary name or label property of the resources.
Term Identifying Properties	Required field that specifies the properties that contain names, aliases, or other identifiers to use for identifying the resources.
Backing Dataset	Optional field that specifies the dataset or datasets to annotate.
Case Sensitive	Controls whether matches must be case-sensitive.
Invalidating Properties	A list of any properties for which you do not want to find matching resources.
Discard Matches Of Common Words	Controls whether to discard matches of the most common words.
Discard Matches of Substrings	A list of the substrings for which you want matches to be discarded. Type a string in the field and then click <b>Add</b> to add the value.
Text Search Query Pattern Precedence	When text search query properties are specified, this setting controls whether resource names or aliases are included as matches. When enabled, resource names and aliases will not be matched.

Setting	Description
Lucene Pattern Properties	A list of properties that contain Lucene query syntax for document categorization.
Approximate Label Properties	A list of properties that contain phrases that may be matched only approximately, i.e., fault-tolerantly, via slightly alternate spellings or misspellings.
Simplified Regex Pattern Properties	A list of properties that contain simplified regular expressions.
Regex Pattern Properties	A list of properties that contain regular expressions.
Strip Characters for Match	Characters to strip out before determining if there is a match.
Clear Caches	Controls whether to clear any existing caches when the pipeline is run.
Rows Per Query	The maximum number of rows to query at a time when paging through the knowledgebase.
Minimum Hit Length	The minimum span length that can count as a match.
Domain Object Base Class URI	If creating a general annotation type (Create Additional General Annotation Type is enabled), this setting specifies the class to use as the base type for the annotator's domain objects.
Class URI Property	Specifies an annotation property whose value you want to map to the class URI in the model. For example, if a Category_ID property has the value http://example.com/Disease and you want to use http://example.com/Disease as the base class URI, add Category_ID to this field. When Class URI Property is

Setting	Description
	not defined, the URI is auto-generated based on the name.
Entity URI Property	Specifies an annotation property whose value you want to map to the URIs for instances of the class. For example, if a Disease_ID property has the value http://example.com/Asthma and you want to use http://example.com/Asthma as the base URI for instances of the class, add Disease_ID to this field. When Entity URI Property is not defined, the URI is auto-generated based on the name.
Is Combine Annotation Instances	Controls whether to combine multiple instances of an extraction into one annotation.
Unintended Instances	A list of any entities or instances of the class to filter out. Type an instance in the field and then click <b>Add</b> to add the value.
Explicit Property Datatypes/Objecttypes	A list of keys that map property names to a particular object property type.
Unintended Classes	A list of any classes to filter out. Type a class in the field and then click <b>Add</b> to add the value.
Unintended Property Names	A list of any property names to filter out. Type a name in the field and then click <b>Add</b> to add the value.
Create Additional General Annotation Type	Controls whether the annotator creates a general shared annotation type in addition to the specific annotation types that are created.
Output the Detections	Controls whether to include specific detections as an annotation property.
Unintended Property	A list of any property values to filter out. Type a value in the field and
Setting	Description
----------------------	--
Values	then click <b>Add</b> to add the value.
Entity Name Property	Specifies an annotation property whose value you want to map to the names for instances of the class. For example, if a Preferred_ Label property includes disease names and you want to use those label values as the names for instances of the Disease class, add Preferred_Label to this field. When Entity Name Property is not defined, the name is auto-generated.
Class Name Property	Specifies an annotation property whose value you want to map to the name of the class. For example, if a Category property has the value Disease and you want the name of the class to be "Disease," add Category to this field. When Class Name Property is not defined, the class name is auto-generated.
Is Error Fatal	Controls whether to fail the pipeline if this annotator fails to create annotations.

### **Regex Annotator**

The table below defines the settings that are displayed when a Regex Annotator is edited.

Setting	Description
Title	Required field that specifies the unique name for the annotator.
Description	Optional field that provides a description of the annotator.
Regular Expression Rule	Required field that lists the regular expression rules for this annotator.
Case-Insensitive	Enables or disables case-insensitive matching. By default, case-

Setting	Description
	insensitive matching assumes that only characters in the US-ASCII character set are being matched. Unicode-aware case-insensitive matching can be enabled by enabling Unicode Case Folding in conjunction with this option.
Multiline Mode	Enables or disable multiline mode. When multiline mode is enabled, the expressions ^ and \$ match immediately after or before a line terminator or the end of the input sequence. When multiline mode is disabled, these expressions only match at the beginning and end of the entire input sequence.
Allow Comments	Controls whether whitespace and comments are allowed in a pattern. When enabled, whitespace and embedded comments starting with # are ignored until the end of a line.
Canonical Equivalence	Controls whether canonical equivalence is taken into account when finding matches. When enabled, characters are considered a match if and only if their full canonical decompositions match. For example, the expression $a \u030A$ will match the string $\u00E5$ .
Enable Dotall	Controls whether dotall mode is used. When enabled, the expression . matches any character, including a line terminator. When disabled, . does not match line terminators.
Literal Parsing	Controls whether literal parsing is employed. When enabled, the input string that specifies the pattern is treated as a sequence of literal characters and metacharacters and escape sequences have no special meaning.
Unicode Case Folding	Controls whether case-insensitive matching is done in a manner that is consistent with the Unicode Standard. By default, Case- Insensitive matching assumes that only characters in the US-ASCII

Setting	Description
	set are being matched.
Unix Lines	Enables or disables Unix line mode. When enabled, only the $\n$ line terminator is recognized in the behavior of ., ^, and $\$$ .
Is Combine Annotation Instances	Controls whether to combine multiple instances of an extraction into one annotation.
Class URI Property	Specifies an annotation property whose value you want to map to the class URI in the model. For example, if a Category_ID property has the value http://example.com/Disease and you want to use http://example.com/Disease as the base class URI, add Category_ID to this field. When Class URI Property is not defined, the URI is auto-generated based on the name.
Entity URI Property	Specifies an annotation property whose value you want to map to the URIs for instances of the class. For example, if a Disease_ID property has the value http://example.com/Asthma and you want to use http://example.com/Asthma as the base URI for instances of the class, add Disease_ID to this field. When Entity URI Property is not defined, the URI is auto-generated based on the name.
Class Name Property	Specifies an annotation property whose value you want to map to the name of the class. For example, if a Category property has the value Disease and you want the name of the class to be "Disease," add Category to this field. When Class Name Property is not defined, the class name is auto-generated.
Unintended Classes	A list of any classes to filter out. Type a class in the field and then click <b>Add</b> to add the value.

Setting	Description
Entity Name Property	Specifies an annotation property whose value you want to map to the names for instances of the class. For example, if a Preferred_ Label property includes disease names and you want to use those label values as the names for instances of the Disease class, add Preferred_Label to this field. When Entity Name Property is not defined, the name is auto-generated.
Unintended Instances	A list of any entities or instances of the class to filter out. Type an instance in the field and then click <b>Add</b> to add the value.
Unintended Property Values	A list of any property values to filter out. Type a value in the field and then click <b>Add</b> to add the value.
Create Additional General Annotation Type	Controls whether the annotator creates a general shared annotation type in addition to the specific annotation types that are created.
Output the Detections	Controls whether to include specific detections as an annotation property.
Explicit Property Datatypes/Objecttypes	A list of keys that map property names to a particular object property type.
Unintended Property Names	A list of any property names to filter out. Type a name in the field and then click <b>Add</b> to add the value.
Domain Object Base Class URI	If creating a general annotation type (Create Additional General Annotation Type is enabled), this setting specifies the class to use as the base type for the annotator's domain objects.
Is Error Fatal	Controls whether to fail the pipeline if this annotator fails to create annotations.

# Model

The topics in this section provide introductory information about data models, describe model requirements, and include instructions for working with models.

In this section:	
Model Concepts and Vocabulary	
Managed Model Concepts	
Model Requirements	
Uploading a Model	
Creating a Model	
Editing a Custom Model	
Editing a Managed Model	
Downloading a Model	
Defining Resource Templates	

# **Model Concepts and Vocabulary**

Models define the business meaning of the source data. They describe the concepts, attributes, and relationships in or across data sets. Instead of reflecting the format or schema of the source data, models reflect the desired structure of the data after it is onboarded to Anzo. Anzo links data to models to provide flexibility for capturing data coming from various sources and structures and to enable users to search for and visualize data in Hi-Res Analytics dashboards or other applications. Users can create models in the Models editor or import existing or external models (OWL ontologies) to Anzo. Models can be shared and reused.

The following image shows a portion of the model for a dataset that captures sales activity for a fictional website where people buy and sell tickets for sporting events, shows, and concerts.



## **Model Vocabulary**

The table below defines key terms to know when working with models.

Term	Description
Class	Models are made up of classes. Classes describe a concept or a group of related objects. For example, the model above contains Tickit Events, Tickit Dates, Tickit Categories, Tickit Sales, and Tickit Listings classes.
Base Class	A base class is a more general version of a class. For example, a film model might

Term	Description
	have a Person base class with subclasses that categorize types of people.
Subclass	A subclass is a more specific version of another class. For example, in a film model a Person class might have subclasses of Actors, Directors, and Producers. Subclasses share the properties from the base class and can include additional properties.
Property	Properties are attributes that describe the data in a class. For example the Tickit Events class has properties such as eventName, dateID, and startTime. There are two types of properties:
	<ul> <li>Data property: Relates a class to a simple or literal value. For example, in the events class, the eventName and startTime properties relate to simple values.</li> </ul>
	Object property: Relates a class to another class. For example, the listID property relates to the Tickit Sales and Tickit Events classes.
Property Type	A property type is the data type of the values for the property, also known as the Property Range.
Instance	Instances are concrete occurrences of a class. For example, an event's name is an instance of the Tickit Events class.
Simple Value	A simple value is also known as a literal value. The list below describes literal values:
	• Numbers (for example, 15, -9, 10.35)
	<ul> <li>Text strings (for example, "Jane Doe" or "a long description")</li> </ul>
	<ul> <li>Dates and times (for example, "13-Dec-2021", or "April, 2022")</li> </ul>
	Boolean (true or false)

Term	Description
Туре	Either a class or a simple value.

For conceptual information about Anzo-generated models, called managed models, see Managed Model Concepts.

# Managed Model Concepts

This topic provides an overview of managed models and includes important concepts to know when working with them.

## What is a Managed Model?

When you onboard a data source with the automated direct load workflow (as described in Onboarding Data with the Automated Workflow) or manually onboard a source with a Direct Load Step (as described in Onboarding Data with a Direct Load Step), a *Managed Model* is produced. A managed model is generated, owned, and managed by the data layer that contains the autogenerated or user-created Direct Load Step. If a Direct Load Step query is changed, additional Direct Load Steps are added to the same layer, or the underlying source schema changes, the managed model is automatically updated when the graphmart or layer is reloaded or refreshed.

## Managed Models Cannot be Edited Outside of Direct Load Steps

Though a model that is generated in a Direct Load Step is registered in Anzo and is available for viewing in the Model editor, the model is owned and managed by the layer that contains the Direct Load Step. That means any manual changes made to the model outside of the step, such as from the Model editor, will be overwritten any time the graphmart or layer is refreshed or reloaded. **Do not modify managed models except by editing (or adding) Direct Load Step queries**. For guidance on editing managed models, see Editing a Managed Model.

## There is One Managed Model Per Data Layer

If you include multiple Direct Load Steps in the same layer, they will all update the same model. This functionality can be useful if you want to align the data and generated model across multiple steps. If you have multiple sources that are not intended to align or update the same model, create separate data layers.

## **Deleting a Layer Deletes the Model**

If you delete a layer that includes a managed model, the model is also deleted. Use caution when referencing managed models outside of graphmarts. For example, if you create a dataset and reference a managed model when you select the ontology, the reference will break if the data layer that manages the model is deleted.

# **Model Requirements**

To ensure that data structures are properly defined, Anzo requires that data models include certain information and avoid unsupported information. This topic provides details about the requirements and guidelines to follow when uploading or creating models.

### **Requirements**

This section lists the requirements or rules to follow when uploading or creating a data model. Models that are generated by Anzo during the auto-ingest process conform to these rules.

- Define each model as an owl:Ontology
- Define the model name with rdfs:label
- The named graph URI must match the ontology URI
- Define classes and concepts with owl:Class
- Define taxonomy with rdfs:subClassOf
- Define properties as owl:DatatypeProperty or owl:ObjectProperty
- Include rdfs:domain and rdfs:range for all properties
- Reference only Anzo-stored models

#### Define each model as an owl:Ontology

Define each data model as an **owl:Ontology**. To do so, include the following triple in the model:

<myOntology> a owl:Ontology

Where myOntology is the URI that names the model. The URI must be unique. To avoid unexpected results when saving a model, do not include a hash (#) character at the end of the model URI.

#### Define the model name with rdfs:label

Use an rdfs:label property to define name of the model as a string. Include the following triple:

For example, you can use the following statement as a template for inserting owl:Ontology and rdfs:label into the model:

```
<myOntology> a owl:Ontology ;
rdfs:label "My ontology"^^xsd:string .
```

#### The named graph URI must match the ontology URI

Make sure that the named graph URI for the model matches the ontology URI. For example:

<myOntology> { <myOntology> a owl:Ontology . }

Like a linked data set, an ontology is a core component that is used throughout the system. The registries that store and track the graphs for core components, such as the ontology registry, expect that each graph contains a resource that matches the graph URI and specifies the type of graph. Having a mismatched graph and ontology URI can break core Anzo functionality.

#### Define classes and concepts with owl:Class

Use **owl:Class** for class or concept definitions. Do NOT include skos:Concept or rdfs:Class. For example, the following statement requires modification to make it valid in an Anzo model:

```
<myConcept> a skos:Concept
```

Changing the statement as follows correctly uses owl: Class instead of skos: Concept:

```
<myConcept> a owl:Class ;
rdfs:label <businessFacingClassLabel> .
```

#### Define taxonomy with rdfs:subClassOf

Use **rdfs:subClassOf** for taxonomy. Do NOT use skos:broader. For example, the following statement requires modification to make it valid in an Anzo model:

<childSkosConcept> skos:broader <parentSkosConcept> .

Changing the statement as follows correctly uses rdfs:subClassOf instead of skos:broader:

```
<childOwlClass> rdfs:subClassOf <parentOwlClass> .
```

#### Define properties as owl:DatatypeProperty or owl:ObjectProperty

Define properties using owl:DatatypeProperty or owl:ObjectProperty. For example:

<myObjectProperty> a owl:ObjectProperty .

Or

```
<myDataTypeProperty> a owl:DatatypeProperty .
```

### Include rdfs:domain and rdfs:range for all properties

Define **rdfs:domain** and **rdfs:range** for all properties. For example, the following property definition is incomplete:

```
<myObjectProperty> a owl:ObjectProperty .
```

The statement below completes the definition by adding rdfs:label, rdfs:domain, and rdfs:range:

```
<myObjectProperty> a owl:ObjectProperty ;
rdfs:label <businessFacingPropertyLabel> ;
rdfs:domain <myClass> ;
rdfs:range <myOtherClass> .
```

The example below shows a valid data type definition:

```
<myDataTypeProperty> a owl:DatatypeProperty ;
rdfs:label <businessFacingPropertyLabel> ;
rdfs:domain <myClass> ;
<myDataTypeProperty> rdfs:range <literal> .
```

### Note

When defining the property range for integer values, use xsd:int instead of xsd:integer.

### **Reference only Anzo-stored models**

Models must be self-contained or include references only to models that are stored in Anzo.

## Guidelines

This section lists additional guidelines and important information to know when working with data models in Anzo.

- Property Range Guidelines
- TriG is the preferred format for models to upload
- Load RDFS and OWL vocabularies as graphs
- Axiomatically defined classes and property hierarchies are not processed

### **Property Range Guidelines**

When creating or editing properties in the model editor, Anzo offers several RDF property ranges or data types to choose from. Certain types are preferred over others, however, because they are treated consistently and predictably across systems. Cambridge Semantics recommends that you specify one of the following preferred property range values:

- Boolean: For true or false values.
- Byte: For 1-byte integers from -128 to 127.
- Date: For date values that follow a format such as YYYY-MM-DD.
- **Date time**: For 8-byte date and time values that follow a format such as YYYY-MM-DDThh:mm:ss. Note that dateTime values are normalized to GMT.
- Double: For up to 8-byte double floating point values.
- **Duration**: For a duration of time expressed as a number of years, months, days, hours, minutes, and seconds in a format such as PnYnMnDTnHnMnS.
- Float: For up to 4-byte floating point values with potential decimal places.
- Int: For up to 4-byte integers from -2,147,483,648 to 2,147,483,647.
- Long: For up to 8-byte integers from -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807.

- Short: For up to 2-byte integers from -32,768 to 32,767.
- String: For character values of varying length.
- **Time**: For time values that follow a format such as hh:mm:ss.

### TriG is the preferred format for models to upload

Anzo accepts model files in OWL (.owl), RDF (.rdf), TriG (.trig), TTL (.ttl), and XML (.xml) format. The preferred format for models that will be uploaded to Anzo is **TriG** (.trig) format.

### Load RDFS and OWL vocabularies as graphs

Anzo loads but does not process additional vocabulary data (such as rdf:subPropertyOf, owl:sameAs, and owl:intersectionOf, etc.) if they are encoded in models. Models that contain vocabularies rather than structural information should be loaded as RDF graphs instead. Anzo can load any valid RDF data. Since RDFS, SKOS, and OWL are valid RDF formats, the vocabulary information can be loaded as a graph, and the data can be interpreted with SPARQL in data layers and Hi-Res Analytics.

### Axiomatically defined classes and property hierarchies are not processed

When models include axiomatically defined classes or property hierarchies, Anzo loads the information but does not process the data. For example, Anzo does not infer information from axiomatically defined classes.

# Uploading a Model

This topic provides instructions for uploading an existing model to Anzo. Follow these instructions if you have a model that was created outside of Anzo or was downloaded from Anzo as described in Downloading a Model. Anzo accepts model files in OWL (.owl), RDF (.rdf), TriG (.trig), TTL (.ttl), and XML (.xml) format.

#### Important

When uploading a model to Anzo, follow the requirements and guidelines defined in Model Requirements.

If you want to import a version of a model that was exported from Anzo (as described in Exporting an Artifact), follow the instructions in Importing Exported Versions of Artifacts to import the model.

#### Note

One of the following outcomes will occur if two users upload the same data model:

- If the second user does not have permission to modify the model that the first user uploaded, the second user receives an access denied error and cannot upload the model.
- If the second user does have permission to modify the model that the first user uploaded, Anzo overwrites the existing model with the version from user two.
- 1. In the Anzo application, click **Model**. Anzo displays the Manage Data Model Working Set screen. For example:

Manage Data Model Working	Set					
Y Q Search	<b>≓</b> * <b>≣</b>	≡ ~		4	Add Moo	lel 👻
Title	Class #	Description	Updated Date 🛧	Actions		÷
K Movies Layer Model	4	The managed model for 'Movies Layer' in the 'Movies' graphmart.	Nov 1, 2022	Д	:	
K Flights Layer Model	1	The managed model for 'Flights Layer' in the 'Flights' graphmart.	Nov 1, 2022	Д	:	
🖧 DB-QA Layer Model	3	The managed model for 'DB-QA Layer' in the 'Books' graphmart.	Nov 2, 2022	Д	:	
유석 DB-QA Layer Model	27	The managed model for 'DB-QA Layer' in the 'Northwind' graphmart.	Nov 2, 2022	Д	:	
🖧 Tickets Layer Model	6	The managed model for Tickets Layer in the Tickets' graphmart.	Nov 2, 2022	Д	:	
🕂 Data Toolkit	32	Defines the interface to the Data Toolkit service.		Д	:	
K SKOS Vocabulary	4	An RDF vocabulary for describing the basic structure and content of concept schem		Д	:	
		Rows per	page: 25 🔻	1-7 of 7	<	>
UPLOAD MODELS				CANO	EL	ок

2. On the bottom left corner of the screen, click **Upload Models**. The Upload Data Models dialog box opens.



3. To upload a model, drag and drop the file onto the dialog box or click the text to browse and select the file on your computer. Anzo uploads the model that you selected and displays the file name and size. For example:

Upload Data Models	
×	
Drag and drop OWL (.owl, .rdf, .ttl, .t or click to browse.	trig, .xml,) files
model_v0.1_version.trig	4.04 kB
DON	E ADD TAGS

If you want to upload additional models, you can repeat the process and drag and drop or select files on the Upload Data Models dialog box.

- If you want to add a tag or edit the tag that was specified in the uploaded model, you can click
   Add Tags and specify the tag in the dialog box. Then click OK.
- 5. Click **Done** when you finish uploading models. The new models become available on the Manage Data Model Working Set screen.

For information about editing models using the model editor, see Editing a Custom Model.

# **Creating a Model**

This topic provides instructions for creating a new data model in the Anzo application. For instructions on uploading an existing model to Anzo, see Uploading a Model.

1. In the Anzo application, click **Model**. Anzo displays the Manage Data Model Working Set screen. For example:

Manage Data Model Working	Set					
T Q Search	<b>=</b> •			4	dd Moo	del 👻
Title	Class #	Description	Updated Date 🛧	Actions		:
K Movies Layer Model	4	The managed model for 'Movies Layer' in the 'Movies' graphmart.	Nov 1, 2022	Д	8	
상 Flights Layer Model	1	The managed model for 'Flights Layer' in the 'Flights' graphmart.	Nov 1, 2022	П	:	
성 DB-QA Layer Model	3	The managed model for 'DB-QA Layer' in the 'Books' graphmart.	Nov 2, 2022	Д	8	
승석 DB-QA Layer Model	27	The managed model for 'DB-QA Layer' in the 'Northwind' graphmart.	Nov 2, 2022	Д	:	
<b>β</b> <sup>4</sup> Tickets Layer Model	6	The managed model for Tickets Layer in the Tickets' graphmart.	Nov 2, 2022	Д	:	
🖧 Data Toolkit	32	Defines the interface to the Data Toolkit service.		Д	:	
H SKOS Vocabulary	4	An RDF vocabulary for describing the basic structure and content of concept schem		Д	:	
		Rows per	page: 25 <del>v</del>	1-7 of 7	<	>
UPLOAD MODELS				CANC	EL	ок

2. Click the **Add Model** button on the top right of the screen and select **New Model**. Anzo displays the Model editor.

Modify Set Create New Model Search Ontology Items	Untitled     Remove From Set     Delete       Not Versioned ID     Explore     Details
Cuntified Model	
	Label *
	The label of the entity
	V 🖸
	Description
	A brief description of the entity
	Serialization Prefix
	The prefix that should be used when serializing elements defined in this ontology
	Class Prefix
Save All	The base URI that should be used when creating classes defined in this ontology

- 3. In the **Label** field, type a unique name for the model.
- 4. Provide the following optional information as needed:
  - Description: A brief description of the model.
  - Serialization Prefix: The prefix to use for this model when Anzo serializes it. For example, the prefix for the Friend of a Friend (FOAF) model is "foaf," and the prefix for Dublin Core is "dc."

#### Tip

The Prefix value is also used to provide hints when typing queries in the Query Builder. When writing a query against a data source that has this model in scope, typing in the PREFIX clause presents this Prefix value as a suggestion.

- Class Prefix: The custom URI template to follow for classes in this model. The value must be a valid URI. When the Class Prefix is set, the URIs for the classes in this model will follow the specified scheme. For example, if Class Prefix is set to http://cambridgesemantics.com/class/ and a class called Employees is created in the model, the URI that is generated for the Employees class will be http://cambridgesemantics.com/class/Employees.
- Property Prefix: The custom URI template to follow for properties in this model. The value must be a valid URI. When the Property Prefix is set, the URIs for the properties in this model will follow the specified scheme. For example, if Property Prefix is set to http://cambridgesemantics.com/property/ and a property called LastName is created in the model, the URI that is generated for the LastName property will be http://cambridgesemantics.com/property/LastName.
- Imports: Lists any definitions that you want to import from another model into this model.
   To select models to import, click in the Imports field and select a model from the dropdown list. Select the field again to select additional models.
- System Model: Indicates that the data model is a system model only and not related to

business data.

- Hidden Model: Hides the data model so that it is not associated with business data.
- 5. Click **Save** to save the model.

For information about adding classes and properties to the new model, see Editing a Custom Model.

# **Editing a Custom Model**

This topic provides information about using the Anzo model editor to open a data model and modify it to add, edit, or remove classes, properties, data ranges, and annotations.

#### Important

Do not modify auto-generated, layer-managed models. Changes will be overwritten whenever the host graphmart or layer is reloaded or refreshed. For more information, see Managed Model Concepts.

- Opening Models in the Editor
- Changing Model Components
- Class Editor Reference
- Property Editor Reference

#### Tip

Before editing a model, you have the option to create a backup of the current version. For more information, see Creating and Restoring Versions of Artifacts.

## **Opening Models in the Editor**

1. In the Anzo application, click **Model**. Anzo displays the Manage Data Model Working Set screen. For example:

Manage Data Model Working Set						
Y Q Search	= • <b>ii</b>			A	dd Mod	el 👻
🗌 🕶 Title	Class #	Description	Updated Date 🛧	Actions		÷
Hovies Layer Model	4	The managed model for 'Movies Layer' in the 'Movies' graphmart.	Nov 1, 2022	Д	:	
Hights Layer Model	1	The managed model for 'Flights Layer' in the 'Flights' graphmart.	Nov 1, 2022	Д	:	
🖧 DB-QA Layer Model	3	The managed model for 'DB-QA Layer' in the 'Books' graphmart.	Nov 2, 2022	Д	:	
<b>∂석</b> DB-QA Layer Model	27	The managed model for 'DB-QA Layer' in the 'Northwind' graphmart.	Nov 2, 2022	Д	:	
🖧 Tickets Layer Model	6	The managed model for Tickets Layer in the Tickets' graphmart.	Nov 2, 2022	Д	:	
🖧 Data Toolkit	32	Defines the interface to the Data Toolkit service.		Д	:	
K SKOS Vocabulary	4	An RDF vocabulary for describing the basic structure and content of concept schem		Д	:	
		Rows per	page: 25 <del>√</del>	1-7 of 7	<	>
UPLOAD MODELS				CANC	EL	OK

2. On the Manage Working Set screen, select the checkbox next to the model (or models) that you want to add to the working set and edit. Then click **OK**. Anzo opens the selected model in the editor. For example:

Modify Set Create New Mo	del	DB-QA Layer Mode	I		Remove From Set	Delete
Search Ontology Items		Explore	Details	Comments	Sharing	>
🗸 🕂 DB-QA Layer Model	1 ^	Label *		-		
> O Alphabetical List Of Products	+	DB-QA Layer Model				
> O Categories	1	The label of the entity				
> O Category Sales For 1997	1					. 6
> O Current Product List	3	http://cambridgesemantics	.com/Layer/acfdea8ci	e55c42ca979fa90eb	c6a43e9/Model	¢ 'U
> O Customer And Suppliers By City	1	Description				
> O Customers		The model for the 'Northwi	nd' graphmart.			
> O Employee Territories						
> O Employees	+					
> O Invoices	+	A brief description of the entity				
> O Order Details						
> O Order Details Extended	+	Serialization Prefix				
> O Order Subtotals		The prefix that should be used whe	en serializing elements defin	ed in this ontology		
> O Orders	Ξ					
Save	All	Class Prefix				
		The base URI that should be used i	when creating classes define	ed in this ontology		

3. You can edit the following model-level settings or view the Changing Model Components section below for information about working with classes, properties, annotations, and data

ranges.

- **Description**: A brief description of the model.
- Serialization Prefix: The prefix to use for this model when Anzo serializes it. For example, the prefix for the Friend of a Friend (FOAF) model is "foaf," and the prefix for Dublin Core is "dc."

#### Tip

The Prefix value is also used to provide hints when typing queries in the Query Builder. When writing a query against a data source that has this model in scope, typing in the PREFIX clause presents this Prefix value as a suggestion.

- Class Prefix: The custom URI template to follow for classes in this model. The value must be a valid URI. When the Class Prefix is set, the URIs for the classes in this model will follow the specified scheme. For example, if Class Prefix is set to http://cambridgesemantics.com/class/ and a class called Employees is created in the model, the URI that is generated for the Employees class will be http://cambridgesemantics.com/class/Employees.
- Property Prefix: The custom URI template to follow for properties in this model. The value must be a valid URI. When the Property Prefix is set, the URIs for the properties in this model will follow the specified scheme. For example, if Property Prefix is set to http://cambridgesemantics.com/property/ and a property called LastName is created in the model, the URI that is generated for the LastName property will be http://cambridgesemantics.com/property/LastName.
- Imports: Lists any definitions that you want to import from another model into this model. To select models to import, click in the Imports field and select a model from the drop-down list. Select the field again to select additional models.
- System Model: Indicates that the data model is a system model only and not related to business data.
- Hidden Model: Hides the data model so that it is not associated with business data.

## **Changing Model Components**

The sections below provides instructions for working with model components. When modifying models, make sure that you click **Save** periodically to save your changes.

- Creating a New Class
- Creating a New Property
- Adding an Existing Property to a Class
- Editing a Class
- Deleting a Property
- Deleting a Class
- Adding a Data Range
- Adding an Annotation

#### **Creating a New Class**

Open the model menu by clicking the menu icon (1) to the right of the model name. Then select **Add Class**.



Anzo opens the class editor so that you can configure the new class. See Class Editor Reference below for information about class settings.

#### **Creating a New Property**

Open the class menu by clicking the menu icon (:) to the right of the class name. Then select **Add Property**.



Anzo opens the property editor so you can configure the new property. See Property Editor Reference below for information about property settings.

#### Adding an Existing Property to a Class

To add an existing property to a class, click the class in the left pane to display the class details in the editor. In the editor, click in the **Properties** field and select the property that you want to add from the drop-down list. For example:



#### **Editing a Class**

To change an existing class, select the class in the left pane. Anzo expands the class to show its properties and displays the details for that class in the editor. You can make changes in the editor. See Class Editor Reference below for information about class settings.

#### **Deleting a Property**

In the left pane of the working set, select the property that you want to delete. Anzo opens that property in the editor. To remove the property, click the **Delete** button on the top right of the screen. Then click **Delete** in the dialog box to confirm that you want to delete the property.

#### **Deleting a Class**

Click the menu icon (:) to the right of the class that you want to remove from the model.



Click **Delete**. Anzo displays a dialog box that asks if you want to delete only the class or all of the subclasses and properties in the class. Select the appropriate option and then click **Delete** to confirm that you want to delete the class. **This action cannot be undone.** Anzo removes the class and saves the model.

#### Adding a Data Range

Click the menu icon (:) to the right of the model name. Then select **Add Data Range**. Anzo opens the data range editor so that you can configure the new range.

#### Adding an Annotation

Click the menu icon (:) to the right of the model name. Then select **Add Annotation**. Anzo opens the editor so that you can configure the annotation.

## **Class Editor Reference**

This section describes each of the fields that are available for configuring classes.

Field	Description
Label	The name of the class.
Description	A brief description of the class.
Parent Classes	Lists any parent classes under which this class becomes a child or subclass. Click in the field to select parent classes from the drop-down list. Or click the X to the left of a class name to remove that parent class from the list.
Properties	Lists the properties under this class. Click in the field to a property from the drop- down list. Or click the X to the right of a property name to remove that property from the list.
Inherited Properties	Properties that the class has inherited from a super class or the model.
Preview Property	Defines a property from the class to use as the "name" or entity on default displays. For example, if there is a reference to entity X, and entity X has Name, Title, and Label properties, you could specify that you want Title to display by default instead of "X."
Resource Template	Defines the Uniform Resource Identifier (URI) template to use for instances of the class. You can construct URI templates by typing a value and pressing <b>Enter</b> or by choosing an available property from the drop-down list. For more information, see Defining Resource Templates.
Graph Template	Defines the graph URI template to use for instances of the class. You can construct graph URI templates by typing a value and pressing <b>Enter</b> or by choosing an available property from the drop-down list. You can concatenate the

Field	Description
	<pre>specified graph template value with values of properties in the class. For example, http://cambridgesemantics.com/graph/ and Title</pre>

# **Property Editor Reference**

This section describes each of the fields that are available for configuring properties.

Field	Description
Label	The name of the property.
Description	A brief description of the property.
Required	Indicates whether a value is required for this property.
Multi Value	Indicates whether more than one value can exist for this property.
	<b>Note</b> Some business intelligence (BI) applications have limitations on the retrieval of multi-value properties. If you use the Anzo Data on Demand service to query data from BI tools, consider whether your application supports multi-value properties before creating them.
Has Data Range	Indicates whether the property has a single data type or a data range. Selecting this checkbox displays the Data Range field so that you can choose the data range.
Property Range	The data type for the property. See Property Range Guidelines for recommendations on choosing property ranges.
Domain	Lists the class or classes that the property belongs to.
Min Cardinality	The minimum number of distinct values a property can have. When Min Cardinality is blank, the number of values is unrestricted.
Мах	The maximum number of distinct values a property can have. When Max

Field	Description
Cardinality	Cardinality is blank, the number of values is unrestricted.
Value Restriction	Indicates whether to restrict the property's values to certain data types or specific values in a list.

# **Editing a Managed Model**

Managed models that are generated by a Direct Load Step are owned and managed by the data layer that contains the Direct Load Step. Any changes made to the model outside of the step, such as from the Model editor, are overwritten any time the graphmart or layer is refreshed or reloaded. This topic provides guidance on updating the properties in a managed model by editing the Direct Load Step query.

#### Tip

Before starting the procedure, you might want to start another session of Anzo. In the new session, open in the Model editor the model that you will be updating. That way you can copy model, class, and property URIs from the model while you edit the Direct Load Step query.

- 1. In the graphmart for which you want to update a model, click the **Data Layers** tab. Expand the layer that was generated when you created the graphmart and find the Direct Load Step that inserts the data.
- 2. Open the Direct Load Step for editing and click the **Query** tab. For example, the image below shows the query that was generated to onboard data about books from a CSV file:



3. Make sure the following prefixes are declared in the PREFIX clause at the top of the query. If any are missing, add them to the query:

```
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
```

4. Next, add the following GRAPH clause statements under the SERVICE clause:

```
GRAPH <model_URI> {
    <property_URI> a owl:DatatypeProperty;
    rdfs:comment "comment";
    rdfs:label "property_label";
    rdfs:domain <class_URI>;
    rdfs:range <dataype> .
[ <property2_URI> a owl:DatatypeProperty;
    rdfs:label "property2_label";
    rdfs:label "property2_label";
    rdfs:domain <class_URI>;
    rdfs:range <datatype> .
]
```

```
[ ... ]
```

• model\_URI: The URI of the model that was generated by the Direct Load Step. You can copy the URI from the Model editor from the Details tab for the model. For example:

<http://cambridgesemantics.com/Layer/698d17.../Model>

• **property\_URI**: The URI to use for the new property or the URI of the existing property that you want to update. When adding a property, you can look at the URIs for other properties in the model and follow the same scheme. For example:

```
<http://cambridgesemantics.com/Layer/6.../Model#GoodReads100kBooks.FirstPr int>
```

- comment: A string that describes the property. For example, "Date the book first went to print".
- property\_label: The label to give the property. For example, "First Print Date".
- class\_URI: The URI of the class that the existing property belongs to or the new property should be added to. You can copy the URI from the Model editor from the Details tab for the class. For example:

<http://cambridgesemantics.com/Layer/698d1794.../Model#GoodReads100kBooks>

• datatype: The URI for the datatype of the property. For example, xsd:date.

For example, the following query adds one new property to the model that is generated by the Direct Load Step query shown above:

```
WHERE {
  SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
    GRAPH
<http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Model> {
<http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Model#Goo</pre>
dReads100kBooks.FirstPrint> a owl:DatatypeProperty;
      rdfs:comment "Date the book first went to print" ;
      rdfs:label "First Print Date" ;
      rdfs:domain
<http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Model#Goo</pre>
dReads100kBooks> ;
     rdfs:range xsd:date .
  }
    []
      s:segment "false"^^xsd:boolean .
    ?generator a s:RdfGenerator, s:OntologyGenerator ;
      s:as (?subject ?predicate ?object) ;
      s:base ${targetGraph} .
    good reads books files:GoodReads 100k books a s:FileSource ;
      s:url "/nfs/data/csv/GoodReads 100k books.csv" ;
      s:model "GoodReads 100k books" .
  }
}
```

- 5. Next, if you added properties and the class for any of the new properties does not have a primary key defined, you must create a key for that class. If all of the classes you referenced in the query have primary keys, you can continue to the next step. If one or more of the classes do not have primary keys, follow the instructions below:
  - a. Locate in the query the statement block for each class that needs a key definition. For example, in the query above, there is only one class, s:model "GoodReads\_100k\_Books". If you have multiple classes, the query has several blocks, such as this example:

```
...
emrdbsmall:emr_complaint a s:DbSource ;
   s:using mysql_db:MySQL_DB ;
```

```
s:table "emrdbsmall.emr_complaint";
s:model "emr_complaint" .
emrdbsmall:emr_patient a s:DbSource;
s:using mysql_db:MySQL_DB;
s:table "emrdbsmall.emr_patient";
s:model "emr_patient" .
emrdbsmall:emr_complaintdescription a s:DbSource;
s:using mysql_db:MySQL_DB;
s:table "emrdbsmall.emr_complaintdescription";
s:model "emr_complaintdescription" .
...
```

- b. At the end of the block for the class you want to add a key to, change the period (.) after s:model to a semicolon (;).
- c. Next, add the following line below s:model:

s:key ("key\_property" [, "key\_property2" ] [, ... ]) .

Where key\_property is the label of the property to use as a key for the class. The property that you choose must have unique values. If there is not a property in the class with unique values, you can specify a combination of properties that would create a unique value. Make sure that the value of key\_property matches the label for that property in the model. For example, for the query in step 4, the Isbn property can be used as a unique key for the GoodReads\_100k\_Books class:

```
good_reads_books_files:GoodReads_100k_books a s:FileSource ;
   s:url "/nfs/data/csv/GoodReads_100k_books.csv" ;
   s:model "GoodReads_100k_books" ;
   s:key ("Isbn") .
```

The final, completed query is shown below:

```
PREFIX good_reads_books_files:
<http://cambridgesemantics.com/DataSource/0e5ab8601249264a318b4a4bcfa81700
/Good_Reads_Books/>
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
```
```
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
INSERT {
 GRAPH ${targetGraph} {
    ?subject ?predicate ?object .
 }
}
WHERE {
 SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
   GRAPH
<http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Mode
1> {
<http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Mode
l#GoodReads100kBooks.FirstPrint> a owl:DatatypeProperty;
      rdfs:comment "Date the book first went to print" ;
      rdfs:label "First Print Date" ;
      rdfs:domain
<http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Mode
l#GoodReads100kBooks> ;
     rdfs:range xsd:date .
  }
    []
     s:segment "false"^^xsd:boolean .
    ?generator a s:RdfGenerator, s:OntologyGenerator ;
      s:as (?subject ?predicate ?object) ;
      s:base ${targetGraph} .
    good reads books files:GoodReads 100k books a s:FileSource ;
      s:url "/nfs/data/csv/GoodReads 100k books.csv" ;
      s:model "GoodReads 100k books" ;
      s:key ("Isbn") .
  }
}
```

6. Save the step and then refresh or reload the graphmart or layer to update the model with the new properties.

For more information about Direct Load Steps, see Directly Load a Data Source (Direct Load Step).

# **Downloading a Model**

This topic provides instructions for downloading a data model from Anzo to your computer.

1. In the Anzo application, click **Model**. Anzo displays the Manage Data Model Working Set screen. For example:

Manage Data Model Working	Set					
Y Q Search	₹ •				Add Moo	iel 👻
🗌 🕶 Title	Class #	Description	Updated Date 🛧	Actions		÷
K Movies Layer Model	4	The managed model for 'Movies Layer' in the 'Movies' graphmart.	Nov 1, 2022	Д	:	
K Flights Layer Model	1	The managed model for 'Fights Layer' in the 'Flights' graphmart.	Nov 1, 2022	Д	:	
K DB-QA Layer Model	3	The managed model for 'DB-QA Layer' in the 'Books' graphmart.	Nov 2, 2022	Д	:	
CB-QA Layer Model	27	The managed model for 'DB-QA Layer' in the 'Northwind' graphmart.	Nov 2, 2022	Д	:	
K Tickets Layer Model	6	The managed model for Tickets Layer' in the Tickets' graphmart.	Nov 2, 2022	Д	:	
🕂 Data Toolkit	32	Defines the interface to the Data Toolkit service.		Д	:	
K SKOS Vocabulary	4	An RDF vocabulary for describing the basic structure and content of concept schem		D	:	
		Rows pa	arpage: 25 👻	1-7 of 7	<	>
UPLOAD MODELS				CANO	EL	ок

2. On the Manage Working Set screen, select the checkbox next to the model that you want to download, and then click **OK**. Anzo opens the selected model in the editor. For example:

Modify Set Create New M	odel E		l		Remove From Set	Delete
Search Ontology Items		Explore	Details	Comments	Sharing	>
∽ 🚜 DB-QA-Layer Model	1 1	ahol *				
> O Alphabetical List Of Products	i D	B-QA Layer Model				
> O Categories	: 1	he label of the entity				
> O Category Sales For 1997	1.1					6
> O Current Product List	: <u>h</u>	ttp://cambridgesemantics.	com/Layer/acfdea8c	e55c42ca979fa90eb	c6a43e9/Model	
> O Customer And Suppliers By Cit	y i 🛛 D	escription				
> O Customers	- T	he model for the 'Northwin	d' graphmart.			
> O Employee Territories						
> O Employees	÷ .					
> O Invoices	- · ·	brief description of the entity				
> O Order Details						
> O Order Details Extended	s	erialization Prefix				
> O Order Subtotals	: 1	he prefix that should be used when	n serializing elements defin	ed in this ontology		
> O Orders	÷					
	C All	lass Prefix				
	T	he base URI that should be used w	hen creating classes define	ed in this ontology		

Open the model menu by clicking the menu icon (1) to the right of the model name. Then select
 Download Model.

Search Ontology Items		Explo	re
V 😽 DB-QA Layer Model	t,	Include Imports	-
> O Alphabetical List Of Products	۵	Download Model	~
> O Categories	Θ	Remove from Set	n
> O Category Sales For 1997	+	Add Class	
> O Current Product List	+	Add Property	1
> O Customer And Suppliers By City	+	Add Annotation	
> O Customers	+	Add Data Range	1
> O Employee Territories	Î	Delete	
> O Employees	:		

Anzo displays the Download Model dialog box:



 In the Download Model dialog box, select the format to save the model in. By default Anzo saves models in TRIG format. If you want to save the file in OWL format, select the OWL radio button. Then click Download.

Anzo downloads the model to your computer in the selected format.

#### Note

When a model is downloaded from Anzo, the resulting TriG or OWL file size can be significantly larger than the file size of the original data model file that was uploaded. The original model likely includes prefix specifications and abbreviated URIs. When a model is exported, however, Anzo replaces the prefixes with full URIs. In addition, the downloaded model includes the Anzo-generated metadata for the model.

For example, the following simple example TTL content shows part of a data model that uses prefixes:

```
@prefix csi: <http://cambridgesemantics.com/2017/02/ont#> .
csi:testModel a owl:Ontology ;
  rdfs:label "Test Model"^^xsd:string .
csi:DOB a owl:Class;
  rdfs:domain csi:Demographics ;
  rdfs:label "DOB" ;
  rdfs:range xsd:string .
csi:HEIGHT a owl:Class;
  rdfs:domain csi:Demographics ;
  rdfs:label "HEIGHT" ;
  rdfs:range xsd:decimal .
```

# After uploading the TTL file and then downloading the model in TriG format, the resulting file includes full URIs as well as the model's metadata:

```
<http://cambridgesemantics.com/2017/02/ont#testModel> {
  <http://cambridgesemantics.com/2017/02/ont#DOB> a
<http://www.w3.org/2002/07/owl#Class> ;
    <http://www.w3.org/2000/01/rdf-schema#domain>
<http://cambridgesemantics.com/2017/02/ont#Demographics> ;
    <http://www.w3.org/2000/01/rdf-schema#label> "DOB" ;
    <http://www.w3.org/2000/01/rdf-schema#range>
<http://www.w3.org/2001/XMLSchema#string> .
  <http://cambridgesemantics.com/2017/02/ont#HEIGHT> a
<http://www.w3.org/2002/07/owl#Class> ;
    <http://www.w3.org/2000/01/rdf-schema#domain>
<http://cambridgesemantics.com/2017/02/ont#Demographics> ;
    <http://www.w3.org/2000/01/rdf-schema#label> "HEIGHT" ;
    <http://www.w3.org/2000/01/rdf-schema#range>
<http://www.w3.org/2001/XMLSchema#decimal> .
  <http://cambridgesemantics.com/2017/02/ont#testModel> a
<http://www.w3.org/2002/07/owl#Ontology> ;
    <http://www.w3.org/2000/01/rdf-schema#label> "Test Model" .
}
<http://cambridgesemantics.com/registries/Ontologies> {
  <http://cambridgesemantics.com/registries/Ontologies>
    <http://openanzo.org/ontologies/2008/07/Anzo#defaultNamedGraph>
    <http://cambridgesemantics.com/2017/02/ont#testModel> ;
    a <http://openanzo.org/ontologies/2008/07/Anzo#Dataset> .
<http://openanzo.org/metadataGraphs
(http%3A%2F%2Fcambridgesemantics.com%2F2017%2F02%2Font%23testModel)> {
```

}

# **Defining Resource Templates**

When you open a data model in the Model editor, there is a **Resource Template** setting for each of the classes in the model. A Resource Template defines the Uniform Resource Identifier (URI) pattern that Anzo should follow when ingesting data and generating the URIs for the instances of each class. Defining a Resource Template for the classes in your models helps link and relate data by using URI patterns that express the meaning of the data and combine similar concepts. Additionally, simpler and more meaningful URIs are easier to read and therefore easier to write in queries.

#### Important

Keep the following points in mind when defining class instance URI patterns:

- Do not modify auto-generated, layer-managed models. Changes will be overwritten whenever the host graphmart or layer is reloaded or refreshed. For more information, see Managed Model Concepts.
- Avoid joining data that should not be joined. For example, using a property such as YearProduced in a movies Resource Template would group all movies from a given year as a single instance.
- Resource Templates with multiple components must have all components present. If a component is missing, Anzo generates random strings for missing Resource Template components.
- Resource templates do not work across different classes. You must define resource templates on individual classes.

#### Tip

For property URIs, the default URI prefix is **http://cambridgesemantics.com**/. The value is controlled by the URI Prefix option in server settings. See <u>Configure URI Prefix and SPARQL</u> <u>Options</u> in the Administration Guide for more information. 1. In the Anzo application, click **Model**. Anzo displays the Manage Data Model Working Set screen. For example:

Q Search				ľ	Add Mor	jel 👻
Title	Class #	Description	Updated Date ↑	Actions		:
K Movies Layer Model	4	The managed model for 'Movies Layer' in the 'Movies' graphmart.	Nov 1, 2022	Д	:	
🙀 Flights Layer Model	1	The managed model for 'Flights Layer' in the 'Flights' graphmart.	Nov 1, 2022	Д	:	
🕂 DB-QA Layer Model	з	The managed model for 'DB-QA Layer' in the 'Books' graphmart.	Nov 2, 2022	Д	:	
K DB-QA Layer Model	27	The managed model for 'DB-QA Layer' in the 'Northwind' graphmart.	Nov 2, 2022	Д	:	
K Tickets Layer Model	6	The managed model for Tickets Layer' in the Tickets' graphmart.	Nov 2, 2022	Д	:	
🖧 Data Toolkit	32	Defines the interface to the Data Toolkit service.		Д	:	
K SKOS Vocabulary	4	An RDF vocabulary for describing the basic structure and content of concept	t schem	Д	:	
		R	lows per page: 25 👻	1-7 of 7	<	>
UPLOAD MODELS				CANO	EL	ок

2. On the Manage Working Set screen, select the checkbox next to the model (or models) that you want to add to the working set for editing. Then click **OK**. Anzo opens the selected model in the editor. For example:



3. Click the **Details** tab.

 Select a class in the model to display the settings for that class. Then scroll down to the Resource Template field. For example, the image below shows the Resource Template field for the selected Movie Actors class.

Search Ontology Items		Properties	
→ 🔊 Movies Laver Model		Acted In × Movie Title × Actor ID × Actor Name ×	$\times \mid \vee$
<ul> <li>O Movie Actors</li> </ul>	:	The properties which belong to this class	
Acted In		Inherited Properties	
12, Actor ID		type ×	~
T Actor Name		The properties which belong to super classes of this class or to the model itself	
$oldsymbol{T}$ Movie Title			
> O Movie Composers	:	Storage	~
> O Movie Directors	:		
> O Movie Producers	:	A star Name	
> O Movies	:	Actor Name The property to use as a preview or label for the given class	~
		The property to use as a preview of tablet of the gmen class	
		Resource Template	v
		A template describing the URI for a newly created resource of this Class	
		Graph Template	
Sa	ve All	A template describing the URI of the graph container for a newly created resource of this Class	

5. Click the **Resource Template** field and specify the URI pattern to use for instances of this class. First, type a base value in the field and press **Enter** to add the value to the field. For example, for Movie Actors in the step above:

http://cambridgesemantics.com/Movie/Actor.



Then click the field again and select a property in the class that defines the class, i.e.,

contains unique values. For example, in the Movie Actors class, Actor ID provides unique values.



 Click Save to save the change, and then select another class for which to set a Resource Template. Repeat the step above for each class in the model.

# Blend

This section provides information about working with datasets and blending data in graphmarts.

In this section:	
Working with Datasets	
Working with Graphmarts	
Profiling Datasets and Graphmarts	

# **Working with Datasets**

The topics in this section provide guidance on working with datasets in the Datasets catalog.

In this section:	
Adding an Empty Dataset for an Export Step	443
Importing an Existing Dataset (FLDS)	
Creating a Dataset from RDF Files	
Managing Dataset Editions	456
Creating a Graphmart from a Dataset	466
Adding a Dataset to a Graphmart	471
Dataset FAQ	475

## Adding an Empty Dataset for an Export Step

Follow the steps below to create an empty dataset that can be used to create an FLDS from the data that is output from an Export Step.

1. In the Anzo application, expand the **Blend** menu and click **Datasets**. Anzo displays the Datasets screen, which lists the catalog of datasets. For example:

<b>T</b>	<b>\</b> Search		₹.					Add	Dataset	
	Title	$\uparrow$	Description		Updated	Date	Actions		:	
	1)- Books		Exported Books dataset				Д	:		
	1). Flights		Exported Flights dataset				Д	:		
	1]. Movies		Exported Movies graphmart				Д	:		
	$\imath\}_*$ Northwind		Exported Northwind graphmart				Д	:		
				Rows per	r page:	25 👻	1-4 of 4	<	>	

2. On the Datasets screen, click Add Dataset. Anzo opens the Create Dataset dialog box.

Create Dataset			
<ul> <li>Empty Dataset</li> </ul>	From Existing RDF	From Existing D	ataset
Create a new dataset bas empty dataset.	ed on RDF data files from an exte	ernal source, or create a	new
To create new RDF data fr	rom an external source, go to the	Data Sources Tab.	
Title *			
Description			
RDF File Location *			BROWSE
NOTE: Only ttl is supported. Data e.g. /path/to/files/rdf.ttl/data.ttl	should be in a folder named either "rdf.tt . To upload, click BROWSE and select the	tl" or "rdf.ttl.gz" depending on t folder that contains your datas	ne file type, iet.
Ontologies			v
Ontologies associated with the f	ile based linked data set		
Include System Data			

3. Select the Empty Dataset radio button.

Create Dataset		
<ul> <li>Empty Dataset</li> </ul>	O From Existing RDF	O From Existing Dataset
Create a new dataset bas empty dataset.	sed on RDF data files from an ext	ernal source, or create a new
To create new RDF data	from an external source, go to the	e Data Sources Tab.
Title *		
Description		
RDF File Location *		BROWSE
Click BROWSE and select the fo	older where you want to create your datase	
Click BROWSE and select the fo	older where you want to create your datase	et.
Ontologies associated with the	older where you want to create your datase file based linked data set	et.
Ontologies Ontologies associated with the Rdf Format	older where you want to create your datase file based linked data set	et.   ~
Ontologies Ontologies associated with the Rdf Format ttl.gz	older where you want to create your datase file based linked data set	et.
Ontologies Ontologies associated with the Rdf Format ttl.gz Set the dataset's default data for	older where you want to create your datase file based linked data set ormat, such as ttl or ttl.gz.	et.

4. Type a name for the new dataset in the **Title** field and an optional description in the **Description** field.

5. Click the **RDF File Location** field to open the File Location dialog box.

File Location			
Current Folder			GO
Selected: None			CLEAR ALL
Server Shared Filesystem		root	
sysadmin User Folder		books	
		csv	
		datafox	
		docs	
		fileUploads	
	4		
CREATE NEW FOLDER			CANCEL OK

- Find and select the directory where the FLDS for this dataset should be created. If needed, you can click Create New Folder to create a new directory. Then click OK to close the dialog box and return to the Create Dataset screen.
- 7. Next, if you know which model or models are associated with the data that you plan to export, select the models from the **Ontologies** drop-down list. Leave **Ontologies** blank if you do not know which models apply or do not want to select one at this time. When the Export steps runs, Anzo automatically exports any models that are related to the exported data.
- 8. Lastly, the RDF Format setting defaults to **ttl.gz**, meaning files that are output to this dataset will be in compressed Turtle format. This is the ideal format for preserving space on the file store. If you do not want the files to be compressed, you can change the value to **ttl**.
- 9. Click **Save** to save the empty dataset and return to the Datasets screen. The new dataset becomes available as a selection when choosing the Target FLDS for an Export Step.

For more information about Export Steps, see Export Data to an FLDS (Export Step).

### Importing an Existing Dataset (FLDS)

Follow the instructions below to add an existing FLDS to the Datasets catalog. Make sure that the FLDS meets the following File Requirements.

#### **File Requirements**

To add an FLDS to the Datasets catalog, the location of the files, the file format, and the directory structure must meet the following requirements:

- **Supported File Locations**: Files can be staged on a configured file store or they can be uploaded from your computer as a .zip file.
- Supported Directory Structure: FLDS directories should contain an flds.trig file, an onts directory that includes the model .trig file, and an rdf.ttl or rdf.ttl.gz directory that contains the data files. For example:

#### Note

Models must be in TriG format, regardless of the file type of the data files.

#### **Importing an FLDS**

Follow the steps below to import an FLDS from the file store or from a .zip file on your computer.

1. In the Anzo application, expand the **Blend** menu and click **Datasets**. Anzo displays the Datasets screen, which lists the catalog of datasets. For example:

Y Q Search		₹ • 🖬 ≣				Add D	ataset
🗌 👻 Title	$\uparrow$	Description	Uţ	odated Date	Actions		:
1). Books		Exported Books dataset			Д	:	
1). Flights		Exported Flights dataset			Д	:	
V. Movies		Exported Movies graphmart			Д	:	
나. Northwind		Exported Northwind graphmart			Д	:	
			Rows per pa	age: 25 👻	1-4 of 4	<	>

2. On the Datasets screen, click Add Dataset. Anzo opens the Create Dataset dialog box.

Create Dataset			
<ul> <li>Empty Dataset</li> </ul>	From Existing RDF	From Existing Data	set
Create a new dataset bas	ed on RDF data files from an exte	ernal source, or create a new	1
To create new RDF data fi	rom an external source, go to the	Data Sources Tab.	
Title *			
Description			
RDF File Location *		BR	OWSE
NOTE: Only ttl is supported. Data e.g. /path/to/files/rdf.ttl/data.ttl	a should be in a folder named either "rdf.tt I. To upload, click BROWSE and select the	tl" or "rdf.ttl.gz" depending on the fil folder that contains your dataset.	e type,
Ontologies			I
Ontologies associated with the f	file based linked data set		
Include System Data			

3. Select the From Existing Dataset radio button.



4. Click the **RDF File Location** field to open the File Location dialog box.



5. Follow the appropriate steps below depending on whether the FLDS is on your computer or the shared File Store:

If the file is on your computer:

a. As a best practice, check the upload location that is listed in the **Upload To** field by hovering your pointer over the value to view the tooltip. Make sure the upload location is a directory on the shared file store and not in the server installation path. If the file is not

uploaded to the shared file store it is not accessible by applications like AnzoGraph. In addition, other users cannot create graphmarts from the data source because they typically do not have access to the file location.

For example, viewing the Upload To location for the screen above shows that the file will be uploaded to the server installation path, /opt/Anzo/Server/data...



If your Upload To location is configured to upload the file to the server installation path, click **Change** and select an upload location that is on the shared file store. For example, the image below shows the Upload Folder Location dialog box that is presented after clicking **Change**. A folder called **fileUploads** is selected on the shared store.

Upload Folder Locatic	n			
℃ Current Folder			GO	
Selected: /fileUploads/ @				CLEAR ALL
Server Shared Filesystem		docs		
sysadmin User Folder		fileUploads		
		flights		
		fmcsa		
	<	chib10		• •
CREATE NEW FOLDER			CANCE	EL OK

- b. Drag and drop the file onto the screen or click browse to navigate to the file and select it.
   Anzo attaches the file and the OK button becomes active.
- c. Click **OK**. The file is added to the RDF File Location field on the Create Dataset screen.

If the files are on the File Store:

a. Click the **From File Store** radio button. Anzo displays the file selection dialog box. For example:

File Location		
Source	From File Store	<u>م</u>
Current Folder		GO Show All File Types
Selected: None		CLEAR ALL
Server Shared Filesystem	noot	Â
sysadmin User Folder	books	
	crawl	
	csv	
	datafox	
CREATE NEW FOLDER		CANCEL OK

- b. On the left side of the screen, select the file store that hosts the FLDS. On the right side of the screen, navigate to the root directory for the dataset. This is the directory that contains the **flds.trig** file, the **onts** directory, and the **rdf.ttl** or **rdf.ttl.gz** directory.
- c. Select the root directory for the FLDS and then **click OK**. The location is added to the RDF File Location field on the Create Dataset screen.
- 6. Click **Save** to import the FLDS and return to the Datasets screen.

You can now select the dataset in the catalog and create a new graphmart or add the dataset to an existing graphmart. See Creating a Graphmart from a Dataset or Adding a Dataset to a Graphmart for instructions.

### **Creating a Dataset from RDF Files**

Source data that is not in RDF format is onboarded through the automated direct data load workflow or unstructured pipelines, where the data is converted to RDF format. If you have data that is already in RDF format in Turtle or N-Triple files, those files can be added to the Datasets catalog directly, making the data available to add to a graphmart for loading and analyzing in AnzoGraph.

#### Note

To import data from CSV, JSON, XML, Parquet, or SAS files, follow the processes described in Adding Data Sources.

Follow the instructions below to create a dataset from a directory of Turtle or N-Triple files. Make sure that the files and directory meet the requirements in File Requirements.

#### **File Requirements**

To add data to the Dataset catalog, the location of the files, the file format, and the directory structure must meet the following requirements:

- **Supported File Locations**: Files can be staged on a configured file store, or they can be uploaded from your computer as a .zip file.
- Supported File Formats: Files must be in one of the following formats:
  - Turtle (.ttl file type)
  - N-Triple (.n3 and .nt file types)

Either of the file types listed above can be compressed in GZIP format and named as **<filename>.<filetype>.gz** files.

Supported Directory Structure: When importing RDF files that are not part of an FLDS, the files must be placed in a directory named rdf.<filetype> or rdf.<filetype>.gz. Stage uncompressed TTL files in a directory called rdf.ttl, and stage compressed TTL files in a directory called rdf.ttl, and stage compressed TTL files in a directory called rdf.ttl, and stage compressed TTL files in a directory called rdf.ttl, and stage compressed TTL files in a directory called rdf.ttl.gz. Stage uncompressed N-Triple files in a directory called rdf.nt or rdf.n3, depending on the file type extension. Place compressed files in an rdf.nt.gz or rdf.n3.gz directory. For example:

#### Important

All files inside an rdf.<filetype> or rdf.<filetype>.gz directory must be the same format and end in the same extension. Data in mixed formats will not load successfully. If you plan to import multiple file types, organize files into separate directories by file extension type, and then import each directory separately.

#### Note

To upload files from your computer, use the same directory structure as shown above. Zip the top-level directory so that the upload file is <code>External-RDF-Top-Level-Directory.zip</code> and contains the <code>rdf.ttl.gz</code> directory.

#### **Importing RDF Files**

Follow the steps below to create a dataset from RDF files.

#### Tip

Anzo provides the option to link the files to an existing data model during the import. If the model is not yet available in Anzo, consider uploading it before importing the RDF files. See Uploading a Model for instructions. You are not required to include a model at import time; a model can be associated with a data set at any time. How do I associate a model with a dataset?

1. In the Anzo application, expand the **Blend** menu and click **Datasets**. Anzo displays the Datasets screen, which lists the catalog of datasets. For example:

Y Q Search	₹ • 🖬 🔳		l	Add Da	itaset
Title 1	Description	Updated Date	Actions		:
1]- Books	Exported Books dataset		Д	:	
1]- Flights	Exported Flights dataset		Д	:	
V. Movies	Exported Movies graphmart		Д	:	
1]. Northwind	Exported Northwind graphmart		Д	:	
	Row	s per page: 25 👻	1-4 of 4	<	>

2. On the Datasets screen, click Add Dataset. Anzo opens the Create Dataset dialog box.

Create Dataset			
<ul> <li>Empty Dataset</li> </ul>	From Existing RDF	From Existing E	)ataset
Create a new dataset bas	ed on RDF data files from an ext	ernal source, or create a	new
To create new RDF data f	rom an external source, go to the	Data Sources Tab.	
Title *			
Description			
RDF File Location *			BROWS
Description RDF File Location * NOTE: Only ttl is supported. Dat e.g. /path/to/files/rdf.ttl/data.tt	a should be in a folder named either "rdf.t I. To upload, click BROWSE and select the	I" or "rdf.ttl.gz" depending on t folder that contains your data	BROW: he file typ set.
Description RDF File Location * NOTE: Only til is supported. Dat e.g. /path/to/files/rdf.til/data.tt Ontologies	a should be in a folder named either "rdf.t I. To upload, click BROWSE and select the	tl" or "rdf.ttl.gz" depending on t folder that contains your data:	BROW he file typ set.
Description RDF File Location * NOTE: Only this supported. Dat e.g. /path/to/files/rdf.ttl/data.tt Ontologies Ontologies associated with the t	a should be in a folder named either "rdf.t I. To upload, click BROWSE and select the file based linked data set	II" or "rdf.ttl.gz" depending on t folder that contains your data	BROWS he file typ set.
Description  RDF File Location *  NOTE: Only ttl is supported. Dat e.g. /path/to/files/rdf.ttl/data.tt  Ontologies  Ontologies  Include System Data	a should be in a folder named either "rdf.t I. To upload, click BROWSE and select the file based linked data set	II" or "rdf.ttl.gz" depending on t folder that contains your data:	BROWS he file typ set.
Description  RDF File Location *  NOTE: Only ttl is supported. Dat e.g. /path/to/files/rdf.ttl/data.tt  Ontologies  Ontologies Include System Data	a should be in a folder named either "rdf.t I. To upload, click BROWSE and select the file based linked data set	tl" or "rdf.ttl.gz" depending on t folder that contains your data	BROW: he file typ set.

- 3. The **From Existing RDF** radio button is selected by default. Type a name for the new dataset in the **Title** field and an optional description in the **Description** field.
- 4. Click the **RDF File Location** field to open the File Location dialog box and follow the appropriate steps below depending on the location of the files.

If you are uploading a .zip file from your computer:

a. As a best practice, check the upload location that is listed in the Upload To field by hovering your pointer over the value to view the tooltip. Make sure the upload location is a directory on the shared file store and not in the server installation path. If the file is not uploaded to the shared file store it is not accessible by applications like AnzoGraph. In addition, other users cannot create graphmarts from the data source because they typically do not have access to the file location.

For example, viewing the Upload To location for the screen above shows that the file will be uploaded to the server installation path, /opt/Anzo/Server/data...



If your Upload To location is configured to upload the file to the server installation path, click **Change** and select an upload location that is on the shared file store. For example, the image below shows the Upload Folder Location dialog box that is presented after clicking **Change**. A folder called **fileUploads** is selected on the shared store.

Upload Folder Locatio	n			
Current Folder			GO	
Selected: /fileUploads/ @				CLEAR ALL
Server Shared Filesystem		docs		i i
sysadmin User Folder		fileUploads		
		flights		
		fmcsa		
	4	chih10		¥
CREATE NEW FOLDER			CANC	EL OK

- b. Drag and drop the .zip file with the RDF files onto the screen or click Browse to navigate to the file on your computer and select it.
- c. Click **OK** to close the dialog box and return to the Create Dataset screen.

If the files are on the File Store

- a. Select the From File Store radio button.
- b. Find and select the rdf.<filetype> directory that you want to import, and then click OK to close the dialog box and return to the Create Dataset screen.
- 5. If you want to associate a model with this dataset, click the **Ontologies** drop-down list and select the model. To include a system model, select the **Include System Data** checkbox. If you do not want to associate a model with the data at this time, leave the Ontologies field blank.

#### Note

Datasets without a model cannot be viewed in Hi-Res Analytics dashboards, but the imported data can still be queried. A model can be associated with the data set at a later time. How do I associate a model with a dataset?

6. Click **Save**. Anzo creates the FLDS and adds the new dataset to the Datasets catalog, and return to the Datasets screen.

#### Note

Anzo generates an flds.trig file at the same level as the rdf.<filetype> directory. The file contains metadata about the load files.

You can now select the dataset in the catalog and create a new graphmart or add the dataset to an existing graphmart. See Creating a Graphmart from a Dataset or Adding a Dataset to a Graphmart for instructions.

### **Managing Dataset Editions**

The topics in this section introduce the concepts to know when working with editions and provide instructions for creating, deleting, and modifying editions.

- Introduction to Editions
- Creating an Edition
- Modifying an Edition
- Deleting a Saved Edition
- Limiting the Number of Editions in a Dataset

#### Introduction to Editions

Editions are collections of the data components that are published by an unstructured pipeline or generated by an Export Step in a graphmart. Editions can be assembled by users and can include any subset of components. This topic introduces you to the concepts that are helpful to know when working with editions.

- What is a Data Component?
- What is the Managed Edition?
- What is a Saved Edition?

#### What is a Data Component?

A **Data Component** is the data that is generated by a successful run of an unstructured pipeline or by the processing of an Export Step in a graphmart. Each time a pipeline or step runs to completion, a new data component is created that contains the version of the data that was generated by that run.

All data components are automatically included in the **Managed Edition** (see What is the Managed Edition?) and any of the components can be added to a **Saved Edition** (see What is a Saved Edition?).

#### What is the Managed Edition?

When an unstructured pipeline runs or an Export Step is processed, the resulting data is added to the **Managed Edition**. This Edition is managed by Anzo and always contains all of the published data components.

The Managed Edition cannot be changed, but it can be cloned (via the **Actions** menu) and saved as a **Saved Edition**. Saved Editions can be modified.

What is a Saved Edition?

A **Saved Edition** is a user-assembled collection of data components. A Saved Edition can contain any combination of and any version of data components. Saved Editions can be created from scratch or can be cloned from the Managed Edition or another Saved Edition.

The Managed Edition or any Saved Edition can be added to a graphmart for analysis.

#### **Creating an Edition**

Follow the instructions below to create a new edition.

1. In the Anzo application, expand the **Blend** menu and click **Datasets**. Anzo displays the Datasets screen, which lists the catalog of datasets. For example:

<b>Y</b>	<b>L</b> Search		₹• 🖬 🔳				Add D	ataset
•	Title	$\uparrow$	Description	U	pdated Date	Actions		:
	1]. Books		Exported Books dataset			Д	:	
	1). Flights		Exported Flights dataset			Д	:	
	i). Movies		Exported Movies graphmart			Д	:	
	$\iota\}_*$ Northwind		Exported Northwind graphmart			Д	:	
			F	Rows per p	age: 25 👻	1-4 of 4	<	>

2. Click the dataset for which you want to create an edition. Anzo displays the Explore tab for the dataset. Click the **Overview** tab, which lists the existing editions. For example:

Description None Advanced   Advanced    Managed Editions Title	EMR Not Version	ied 🖉	ore Graphmarts	+ Dashboards	Add to Cart Cres	ate Graphmart
Title               Description            Most Recent Published Date           Actions           Statistics        EMR - Edition Cumulative Export Ste           :           Statistics           Statistics        aved Editions           Q. Search           Create New Edition           Tags         None	escription ione dvanced - fanaged Editions	3			Ty Crea Structure Modifi Releas	General pe; Catalog tor: System Administrat ed: a minute ago ed: a minute ago
Aved Editions  Q. Search  Q. Search  Create New Edition	Title EMR - Edition Cumula	↑ Description	Most Recent Po	ablished Date Actions	Statistics not avai	Statistics Jable
Q Search					http://cambridge	semantics.com/FileBa 🛅 e Log
No editions found	Q Search	= *		Create New Edition	Tags None	

3. To create a new edition from scratch, click the **Create New Edition** button at the bottom of the screen.

#### Tip

If you want to create an edition by cloning the Managed Edition, click the menu icon (

in the Actions column for the Managed Edition and select Clone Edition. To clone a

Saved Edition, click the menu icon for the Saved Edition and select Clone Edition.

The Create New Edition (or Clone Edition) screen is displayed. For example:

Create New	/ Edition						
Title *		Description					
Data Compone	nt for selected job included in th	is Edition					
🗌 👻 Title			Created Date			Actions	:
Crea	ated : 11/21/2022 02:30PM		11/21/2022 02:30PM			:	
Crea	ated : 11/21/2022 02:25PM		11/21/2022 02:25PM			:	
			Rows per page:	25 🔻	1-2 of 2	< :	>
					CANCE	L SAV	/E

- 4. Specify a name for the edition in the **Title** field and include an optional description in the **Description** field.
- 5. In the Data Components list, select the checkbox next to each component that you want to add to this edition. For example:

Create New Edition		
Title * Last Run Only Data Component for selected job include	Description	
Title	Created Date	Actions 🚦
Created : 11/21/2022 02:30PM	11/21/2022 02:30PM	:
Created : 11/21/2022 02:25PM	11/21/2022 02:25PM	:
	Rows per page: 25 ▼ 1-2 of 2	< >
1 Items Selected		
	CANCEL	SAVE

6. When you are finished selecting data components, click **Save** to save the edition. The new edition is added to the list of **Saved Editions** on the Overview screen. For example:

Saved Editions				
<b>Q</b> Search				Create New Edition
Title	$\uparrow$	Description	Last Modified Date	Most Recent Published Da Actions
Last Run Only			11/21/2022 02:35PM	:

From the **Actions** menu for an edition, you can create a graphmart, or you can browse, clone, or delete the edition.

#### **Modifying an Edition**

You cannot change an existing edition, but you can clone and edit a copy of an edition. Follow the steps below to create a new edition based on an existing version.

1. In the Anzo application, expand the **Blend** menu and click **Datasets**. Anzo displays the Datasets screen, which lists the catalog of datasets. For example:

<b>T</b>	<b>L</b> Search		₹• 🖬					Add	Dataset
	Title	$\uparrow$	Description	l	Updated	Date	Actions		:
	나- Books		Exported Books dataset				Д	:	
	1). Flights		Exported Flights dataset				Д	:	
	1}. Movies		Exported Movies graphmart				Д	:	
	$\iota\}_*$ Northwind		Exported Northwind graphmart				Д	:	
				Rows per	page:	25 👻	1-4 of 4	<	>

2. Click the dataset for which you want to modify an ddition. Anzo displays the Explore tab for the dataset. Click the **Overview** tab, which lists the existing editions. For example:

4.	EMR Not Versioned 🗹						+ Add to Cart	Create Graphmart
	Overview		Explore	Graphma	rts Dashb	oards	Versions	Discussion >
Description None Advanced Managed	d Editions						Structu	General Type: Catalog Creator: System Administrat reModfied: 7 minutes ago Released: 13 minutes ago
Title EMR - Edi	ition Cumulative Exp	↑ oort Sta	Description	Most R	ecent Published Date	Actions	Statistics	Statistics s not available
							http://ca	mbridgesemantics.com/FileBa 📋
							Dynamic	Resource Log
Saved Ed	litions ch		- •		Creat	e New Edit	Tags None	
Title	<b>Υ</b>	Descrip	tion	Last Modified Date	Most Recent Published	Dal Actions		
Last Run (	Only			11/21/2022 02:35PM		:		

3. Click the menu icon in the Actions column for the edition that you want to copy and select Clone Edition. Or select Browse Edition if you want to review the edition before making a copy. When you are ready to make a copy, click the Edit a Copy button. Anzo opens the Edition for editing. For example:

Clone Edition									
Title *		Description							
Data Cor	nponent for selected job included in th	his Edition							
-	Title		Created Date			Actions	÷		
	Created : 11/21/2022 03:58PM		11/21/2022 03:58PM			÷			
	Created : 11/21/2022 02:30PM		11/21/2022 02:30PM			÷			
	Created : 11/21/2022 02:25PM		11/21/2022 02:25PM	:					
			Rows per page:	25 🕶	1-3 of 3	<	>		
3.Item	s Selected								
					CANCE	L SA	VE		

4. Specify a name for the edition in the **Title** field and include an optional description in the **Description** field.

5. To make changes to the edition, select or clear the Data Component checkboxes on the left side of the screen to include or exclude components.

#### Note

When you make changes to an edition while creating or changing a graphmart, Anzo creates a copy of the edition with the changes and uses the copy as a dataset in the graphmart. The original published edition remains unchanged.

6. When you have finished modifying the edition, click **Save**. Anzo creates the edition and adds it to the list of Saved Editions on the Overview screen.

The new edition is now available to add to a new or existing graphmart. To quickly create a new graphmart, you can click the menu icon in the Actions column for the new edition and select **Create Graphmart with this Edition**.

#### **Deleting a Saved Edition**

Follow the steps below to delete a Saved Edition.

#### Note

Before deleting an edition, ensure that there are no graphmarts that require that edition.

#### Tip

You cannot delete the Managed Edition, but users with administrative privileges can clear out the existing components so that the edition is recreated from scratch the next time the pipeline is published. For instructions, see How do I clear the components from the Managed Edition?

1. In the Anzo application, expand the **Blend** menu and click **Datasets**. Anzo displays the Datasets screen, which lists the catalog of datasets. For example:

Y Q Search	₹ • 🖬 🔳		I	Add Da	ataset
Title	▶ Description	Updated Date	Actions		:
1)- Books	Exported Books dataset			:	
1]. Flights	Exported Flights dataset		Д	:	
1). Movies	Exported Movies graphmart		Д	:	
1)+ Northwind	Exported Northwind graphmart		Д	:	
	R	Rows per page: 25 👻	1-4 of 4	<	>

2. Click the dataset for which you want to delete an edition. Anzo displays the Explore tab for the dataset. Click the **Overview** tab, which lists the existing editions. For example:

EMR Not Versioned	Ø				+ Add to Cart	Create Graphmart
Overview	N	Explore	Graphmarts	Dashboards	Versions	Discussion >
Vescription None Advanced V Managed Editions		Receiving	Mart Ground	n-blinked Parks Astronomy	Structure I F	General Type: Catalog Creator: System Administrat todified: 6 minutes ago teleased: 2 hours ago
The Strike Grandship	-T*	beschpuon	Most Recent	Published Date Activ	Statistics no	Statistics available
	- Enport du				http://camb	ridgesemantics.com/FileBa 📋
					Dynamic Re	source Log
aved Editions				Create New D	Tags None	
Q Search		- *		Create New E	aition	
Title 个	Descriptio	on Last N	iodified Date Mos	t Recent Published Dal Actio	ons	
Last Run Only		11/21,	2022 02:35PM	:		
Original Component		11/21/	2022 04:05PM			

In the list of Saved Editions, click the menu icon in the Actions column for the edition that you
want to delete and select **Delete**. Anzo displays a confirmation message. Click **OK** to confirm
the delete operation and remove the edition.

#### Limiting the Number of Editions in a Dataset

By default, there is no limit on the number of editions that can be created and preserved in a dataset. Each time an unstructured pipeline or Export Step is run, a new edition is generated and the TTL for that edition is written to disk. You can modify a dataset, however, to set a limit on the number of editions to archive. When a limit is set, Anzo to ages off and removes older editions from disk. Follow the steps below to limit the number of editions for a dataset.

1. In the Anzo application, expand the **Blend** menu and click **Datasets**. Anzo displays the Datasets screen, which lists the catalog of datasets. For example:

<b>Y</b>	<b>L</b> Search		₹• 🖬 🛢					Add	Dataset
•	Title	$\uparrow$	Description		Updated	i Date	Actions		:
	ν}- Books		Exported Books dataset				Д	:	
	ν}. Flights		Exported Flights dataset				Д	:	
	1). Movies		Exported Movies graphmart				Д	:	
	$\iota\}_*$ Northwind		Exported Northwind graphmart				Д	:	
				Rows per	r page:	25 👻	1-4 of 4	<	>

- 2. Click the dataset for which you want to configure the maximum number of editions. The Explore screen is displayed.
- 3. Click the **Overview** tab. Then click **Advanced** to expand the screen and show the advanced settings. The image below shows the settings.



4. Click in the **Maximum Number of Archives** field to make it editable. Then type the number of archived editions that you want to keep for the dataset. For example:

\$
$\checkmark$ ×

5. Click the checkmark icon ( $\checkmark$ ) to save the change.

The dataset is now configured to limit the number of editions that are retained. When the maximum number is reached, Anzo ages off and removes the oldest edition. For unstructured datasets, Anzo also removes the corresponding Elasticsearch Index JSON backups and pipeline runs.

## Creating a Graphmart from a Dataset

Follow the steps below to create a new graphmart from a dataset in the Datasets catalog.

1. In the Anzo application, expand the **Blend** menu and click **Datasets**. Anzo displays the Datasets screen, which lists the catalog of datasets. For example:

▼ C	<b>L</b> Search		₹ • 🖬 ≣				Add D	ataset
•	Title	$\uparrow$	Description	Upda	ted Date	Actions		:
	1]. Books		Exported Books dataset			Д	:	
	ν}. Flights		Exported Flights dataset			Д	:	
	1). Movies		Exported Movies graphmart			Д	:	
	$\iota_{I^*}^*$ Northwind		Exported Northwind graphmart			Д	:	
			I	Rows per page	e: 25 👻	1-4 of 4	<	>

2. In the catalog, click the checkbox next to each dataset that you want to add to the new graphmart. Hover the pointer over an item to display the checkbox in the left column. Anzo adds the datasets to the shopping cart and additional icons become available at the bottom of the screen. For example:

<b>T</b>	<b>Q</b> Search	₹ • 🔳 🗮		Add D	ataset
•	Title	↑ Description	Updated Date	Actions	:
	4. Books	Exported Books dataset		Д	:
	4). Flights	Exported Flights dataset		Д	:
	1]+ Movies	Exported Movies graphmart	Nov 2, 2022	Д	:
	4)- Northwind	Exported Northwind graphmart		Д	:
	1]+ Tickets	Exported Tickets graphmart		Д	:
			Rows per page: 25 👻 1-5 of	5 <	>
<u>1 Item</u>	ns Selected 🛞	Bookmark 🗍 Delete	Create Graphmart	Package	•

3. Click the **Create Graphmart** button. Anzo displays the Create Graphmart screen and populates the Title field by appending "Graphmart" to the name of the dataset.

Create Graph	hmart							
<sup>Title*</sup> Movies Graphn	nart							
The title of the grap	phmart							
Description								
A brief description	of the graphmart							
Title	$\uparrow$	Edition	Description		Last Published	Date	Actions	:
Movies		Movies - Edition Cumulative Export					/	
				Rows p	erpage: 25 🔻	- 1-1 of 1	<	>
				C/	ANCEL CRE	ATE & ACTIVATE	CREA	TE

- 4. On the Create Graphmart screen, you can edit the **Title** and add an optional **Description**.
- 5. By default the current working edition (Managed Edition) of the dataset is selected. If you want to select a different edition, follow these steps:
  - a. click the Modify Edition (\*) icon in the Actions column. The Modify Edition dialog box is displayed. For example:

Modify	Edition						
Title Movies Data Cor	Descript	lion					
<b>v</b> •	Title	Created Date			Actions	÷	
	Created : 11/02/2022 03:07PM	11/02/2022 03:07PM			:	^	
	Created : 11/02/2022 02:13PM	11/02/2022 02:13PM			:		
	Created : 11/02/2022 12:28PM	11/02/2022 12:28PM			÷		
	Created : 11/02/2022 10:27AM	11/02/2022 10:27AM			:	-	
			Rows per page: 2	5 <del>v</del> 1-4 of 4	<	>	
<u>4 Item</u>	s Selected						
SAVE A	SAVE AS NEW EDITION CANCEL SAVE CUSTOR						

- b. To choose a different edition, click the drop-down list at the top of the screen and select the edition to use.
- c. If you want to make changes to the selected edition, select or clear the Data Component checkboxes on the left side of the screen to include or exclude components.

#### Tip

You can also rename a component by clicking the menu icon in the Actions column and selecting **Rename Component**. In addition, you can remove a component from a dataset by clicking the Actions menu and selecting **Remove Component from Dataset**. However, note that choosing this option deletes the component from the edition and **deletes that component's data from the file system**.

- d. When you are finished making changes, choose one of the following options for saving the changes:
  - To save the changes as a new edition, click Save As New Edition on the left side of the screen. The Create New Edition dialog box is displayed. Specify a Title and optional Description and click Save.
| Create New Edition |        |      |
|--------------------|--------|------|
| Title *            |        |      |
| Description        |        |      |
|                    | CANCEL | SAVE |

- To save the changes as a copy of the existing edition, click Save Customizations. Anzo clones the edition and adds the copy to the list on the Create Graphmart screen.
- To create the graphmart without activating it, click Create. If you want to create the graphmart and activate it, click Create & Activate. Anzo creates the graphmart and displays the Overview screen. For example:

#### Note

If you activate the graphmart and have more than one static AnzoGraph engine configured or you have a Cloud Location configured for dynamic AnzoGraph deployments, Anzo displays a **Select an AnzoGraph Query Engine** dialog box. Click the drop-down list to select the engine to load the graphmart to, or select **Spin up new AnzoGraph** (if available) to deploy a new instance. Then click **OK**.



# Tip

If you want to cancel graphmart activation while data is loading, open the Activity Log by clicking the Activity Log icon (()) in the main menu bar. Then click **Cancel** for the

Provisioning...graphmart activity. For example:



When graphmart creation is complete and the graphmart is activated, the data is available to access and analyze. For more information, see Access & Analyze. For more information about graphmarts, see Working with Graphmarts.

# Adding a Dataset to a Graphmart

This topic provides instructions for adding a dataset from the Datasets catalog to an existing graphmart.

1. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

<b>Y</b> (	<b>Q</b> Search	=	· • 🖬 \Xi			Add Gra	phmart	•
•	Title $\uparrow$	Description	Status	Statements	Creator	Actio	ns	÷
	🕸 Books		C Ready to use	156	System Administrator	Д	:	
	Flights		C Ready to use	265,339	System Administrator	Д	:	
	≋ Movies		Ready to use	104,166	System Administrator	Д	:	
	© Northwind		C Ready to use	131,386	System Administrator	Д	:	
	© Tickets		C Ready to use	3,696,548	System Administrator	Д	:	
				Rows	s per page: 25 👻 1	I-5 of 5	<	>

- 2. On the Graphmarts screen, click the name of the graphmart that you want to add data to. Anzo displays the Overview.
- Click the **Datasets** tab. The screen lists the datasets in the graphmart. For example, the image below shows a graphmart without any datasets because this is a graphmart that was auto-generated from a data source:

	Movies Derived from: Movies_3 INACTIVE	20221118182145 <b>COVI</b> 2 ACTIVE ⊘ Ready to u	se 🛆 AZG 🛛 Static		Save a Copy	Create Dashboard	D
	Overview	Explore	Datasets	Data Layers	Dashboards	Data on Demand	>
Member Q Seal No editio	Datasets rch	<b></b>				Add Data	aset

4. Click the Add Dataset button. The Select Dataset dialog box is displayed.

Select Dataset						
Y Search	F • 🖬 🗉					
🗌 🕶 Title 🔶 🛧	Description	Updated Date	Tags	A	tions	÷
i]- Northwind		Nov 18, 2022			Д	:
V). Tickets		Nov 18, 2022			Д	:
		0	05	10-40	,	
		ROWS	s per page: 25 ♥	1-2 01 2	<	2
				CAN	CEL	ADD

 In the dialog box, select the checkbox next to the dataset that you want to add, then click Add. Anzo adds the dataset to the graphmart. (A new data layer is generated with a Load Dataset Step that will load the dataset.)

$\phi$ This	graphmart is modified a	and needs refresh.	REFRESH			
**	Movies Derived from: Movies_2022 INACTIVE A	ACTIVE Ø Ready to u	ise 🛆 AZG 🛭 St	atic	Save a Copy Cre	ate Dashboard 📃
	Overview	Explore	Datasets	Data Layers	Dashboards	Data on Demand 🔸
Member	Datasets					
Q Sear	rch	= -				Add Dataset
Title	$\uparrow$	Edition		Description	Last Published Date	Actions :
Tickets		Tickets - Edition Cumulativ	e Export Starting Wed			/ 🖂

- 6. By default the current working edition (Managed Edition) of the dataset is selected. If you want to select a different edition, follow these steps:
  - a. Click the Edit icon (
     in the Actions column. The Modify Edition dialog box is displayed.
     For example:

Modify	Edition						
Title Tickets Data Cer	Description						
<b>v</b> •	Title	Created Date				Actions	÷
	Created : 11/21/2022 09:13AM	11/21/2022 09:13AM				:	Â
	Created : 11/18/2022 04:46PM	11/18/2022 04:46PM					I.
	Created : 11/18/2022 04:26PM	11/18/2022 04:26PM				:	
	Created : 11/18/2022 02:25PM	11/18/2022 02:25PM					
	Created : 11/18/2022 09:18AM	11/18/2022 09:18AM					v
			Rows per page:	25 👻	1-8 of 8	<	>
8.item	is Selected						
SAVE A	AS NEW EDITION		CAI	NCEL	SAVE CUS	TOMIZATIO	INS

- b. To choose a different edition, click the drop-down list at the top of the screen and select the edition to use.
- c. If you want to make changes to the selected edition, select (to include) or clear (to exclude) the data component checkboxes on the left side of the screen.

#### Note

When you make changes to an edition while creating or changing a graphmart, Anzo creates a copy of the edition (with the changes) and uses the copy as a dataset in the graphmart. The original published edition remains unchanged.

- d. When you are finished making changes, choose one of the following options for saving the changes:
  - If you want to save the changes as a new Saved Edition, click Save As New Edition. Anzo displays the Create New Edition dialog box. Specify a Title and optional Description for the edition, and click Save. Then click Save Customizations on the Modify Edition screen.

Create New Edition		
Title *		
Description		
	CANCEL	SAVE

- If you want to save the changes as a copy of the existing edition, click Save Customizations. Anzo clones the edition and adds the copy to the list on the screen.
- 7. To reload the graphmart and add the new dataset to AnzoGraph, click the **Data Layers** tab, and then click the **Reload** button (C).

### Tip

If you want to cancel graphmart activation while data is loading, open the Activity Log by clicking the Activity Log icon (()) in the main menu bar. Then click **Cancel** for the

Provisioning...graphmart activity. For example:



Once the graphmart is loaded into AnzoGraph, the data is available to access and analyze. For more information, see Access & Analyze.

# **Dataset FAQ**

This topic provides answers to frequently asked questions about datasets.

- How do I find the URI for a dataset?
- How do I find the catalog entry URI for a dataset?
- How do I associate a model with a dataset?
- How do I clear the components from the Managed Edition?

# How do I find the URI for a dataset?

Anzo displays dataset details on the Overview screen for the dataset. To view and copy a dataset URI, go to the Overview tab for the dataset. The URI is under General information on the right side of the screen. You can click the clipboard icon ( $\Box$ ) to copy the URI.



# How do I find the catalog entry URI for a dataset?

To query from a remote client (such as over the SPARQL endpoint) a linked data set (LDS) that is stored in a local volume, you need to specify the catalog entry URI for that LDS as the target data set. The catalog entry URI uniquely identifies an LDS because it encodes both the LDS and its data source (local volume) in the URI. Follow the steps below to find the catalog entry for an LDS.

1. First, retrieve the URI for the dataset whose catalog entry URI you want to find. For instructions, see How do I find the URI for a dataset? above.

 Next, open the Find tab in the Query Builder. In the Anzo application, expand the Access menu and click Query Builder. Then click the Find tab. The Find screen opens and the System Datasource is selected as the target data source.

Query	Find		
Select a datasource:			
System Datasource			x   ~
Subject	Predicate	Object	Graph
			Clear Add Statements Find

- 3. If the LDS is in a different volume, click the **Source** drop-down list and select the appropriate volume. Typically, linked data sets are stored in the system volume.
- 4. Paste in the **Object** field the LDS URI that you copied in the first step. Then click **Find**. Anzo returns the set of quads for which the LDS URI is the object. For example:

Result	s (6)	Show native 🗹 Subject 🗹 Pr	redicate 🧧 Object 📋 Named Graph	0
EDIT	DELETE			
	Subject $\downarrow$	Predicate	Object	
	* http://cambridgesemantics.com/FileBackedLinkedDataSet/de9fa6f84bc74d06 a0d643917fd5132d +	http://cambridgesemantics.com/ ontologies/2009/05/LinkedData#d ataset >	http://cambridgesemantics.com/Fil eBackedLinkedDataSet/de9fa6f84bc 74d06a0d643917fd6132d >	:
	http://cambridgesemantics.com/registries/LinkedDataSets >	< http://openanzo.org/ontologies/ 2008/07/Anzo#defaultNamedGrap h >	http://cambridgesemantics.com/Fil eBackedLinkedDataSet/de9fa6f84bc 74d06a0d643917fd6132d >	:
	http://openanzo.org/catEntry[%58http%3A%2F%2Fcambridgesemantics.com% 2FFileBackedLinkedDataSet%2Fde91a6/B4bc74d06a0d643917td6132d%50%4 0%58http%3A%2F%2Fopenanzo.org%2Fdatasource%2FsystemDatasource%50) >	http://cambridgesemantics.com/ ontologies/2009/05/LinkedData/IIi nkedDataset *	http://cambridgesemantics.com/Fil eBackedLinkedDataSet/de9fa6f84bc 74d06a0d643917fd6132d >	:
	<ul> <li>http://openanzo.org/catEntry(%58http%3A%2F%2Fcambridgesemantics.com% 2FFileBackedLinkedDataSet%2Fde91a6/84bc74d06a0d643917td6132d%50%4 0%58http%3A%2F%2Fopenanzo.org%2Fdatasource%2FsystemDatasource%50) &gt;</li> </ul>	< http://cambridgesemantics.com/ ontologies/2009/05/LinkedData#d ataset >	http://cambridgesemantics.com/Fil eBackedLinkedDataSet/de9fa6f84bc 74d06a0d643917fd6132d >	:
	http://openanzo.org/catEntry[\\5Bhttp\\3A\\21\\2Fcambridgesemantics.com\\ 2FFileBackedLinkedDataSet\\2Fde91a6/84bc74d06a0d643917td6132d\\5D\\4 0\\5Bhttp\\3A\\2f\\2Fopenanzo.org\\2Fdatasource\\2FsystemDatasource\\5D) >	« http://openanzo.org/ontologies/ 2008/07/Anzo#inheritsFrom »	http://cambridgesemantics.com/Fil eBackedLinkedDataSet/de9fa6f84bc 74d06a0d643917fd6132d >	÷
	http://openanzo.org/datasets#NamedGraphs +	« http://openanzo.org/ontologies/ 2008/07/Anzo#namedGraph »	« http://cambridgesemantics.com/Fil eBackedLinkedDatsSet/de9fa6f84bc 74d06a0d643917fd6132d »	:

5. In the Subject field in the results, look for a URI that begins with

http://openanzo.org/catEntry. The value is the catalog entry URI for the LDS. For example:



 Copy the entire URI. This is the URI to use as the target data source for SPARQL endpoint queries against the LDS. For more information about the SPARQL endpoint, see Access the SPARQL Endpoint.

### How do I associate a model with a dataset?

Follow the instructions below to associate a model that is in Anzo with an onboarded dataset.

- 1. In the Anzo application, expand the **Blend** menu and click **Datasets**. Anzo displays the Datasets catalog, which lists the existing datasets.
- 2. Click the dataset that you want to add a model to. Anzo displays the Explore screen for the dataset.
- 3. Click the **Overview** tab. Under the Description field, click **Advanced** to display the advanced options. For example:



- 4. Click in the **Models** field to make it editable, then click the Models drop-down and select the model to add to this dataset. To include a system model, select the **Include System Data** checkbox. To select multiple models, click the drop-down list again and select another model.
- 5. When you have finished selecting models, click the checkmark icon ( $\checkmark$ ) to save the change and associate the model or models with the dataset.

# How do I clear the components from the Managed Edition?

Follow the instructions below if you want to clear out all of the existing components from the Managed Edition so that the edition is recreated from scratch the next time the pipeline is published or the dataset is exported from a graphmart.

#### Note

Permission to Manage Semantic Services is required to complete this task.

- First, copy the URI of the dataset for which you want to clear the Managed Edition. How do I find the URI for a dataset?
- 2. Next, In the Administration application, expand the **Monitoring & Diagnostics** menu and select **Semantic Services**.
- 3. Search for the LinkedDataService and view its details. Then click the Service Builder tab in Semantic Service Details.
- 4. Click the **Please Select an Operation** field and select **clearWorkingEdition** from the dropdown list. The Request Statements for the service call are populated:

mantic Service Details				0	0
Overview	Operations	Service Builder			
clearWorkingEdition					
tp://cambridgesemantics.co	m/semanticServices/Linked	Data#clearWorkingEdition 📋			
Request Statements					
1 @prefix xsd: <http 2 @prefix foaf: <htt 3 @prefix rdfs: <htt 4 @prefix rdfs: <http: 5 @prefix dcterms: &lt; 6 @prefix rdf: <http 7 @prefix owl: <http: 8 @prefix system: <ht 10 @prefix graphmar: 12 13 <http: servicereq<br="">14 <http: servicereq<br="">15 <http: servicereq<br="">16 <http: servicereq<br="">17 <http: servicereq<br="">18 <http: servicereq<br="">19 <http: servicereq<br="">10 <http: servicereq<br="">10 <http: servicereq<br="">11 <http: servicereq<br="">12 <http: servicereq<br="">13 <http: servicereq<br="">14 <http: servicereq<br="">15 <http: servicereq<br="">16 <http: servicereq<br="">17 <http: servicereq<br="">18 <http: servicereq<br="">18 <http: servicereq<br="">18 <http: servicereq<br="">19 <http: servicereq<br="">10 <http: servicereq<br="">10 <http: servicereq<br="">10 <http: servicereq<br="">10 <http: servicereq<br="">11 <http: servicereq<br="">12 <http: servicereq<br="">13 <http: servicereq<br="">14 <http: servicereq<br="">15 <http: servicereq<br="">16 <http: servicereq<br="">17 <http: servicereq<br="">18 <http: servicereq<br="">19 <http: servicereq<br="">10 <http: servicereq<br="">11 <http: servicereq<br="">12 <http: servicereq<br="">13 <http: servicereq<br="">14 <http: servicereq<br="">15 <http: servicereq<br="">16 <http: servicereq<br="">16 <http: servicereq<br="">17 <http: servicereq<br="">17 <http: servicereq<br="">18 <http: servicereq<br="">19 <http: servicereq<br="">19 <http: servicereq<br="">10 <http: servicereq<="" td=""><td>://www.w3.org/2001/XM p://xmlns.com/fodf/0. p://wmlw3.org/2000/0 //purl.org/dc/element http://purl.org/dc/element http://purl.org/dc/element http://openanzo.org/ong/07 ttp://openanzo.org/onto p://openanzo.org/onto //cambridgesematics. <http: cambridgesematics.<br=""><http: cambridgesematics.<br=""><http: cambridgesematics.<br=""><http: cambridgesematics.<br=""><http: cambridgesematics.<br=""><http: cambridgesematics.<br=""><http: www.scambridgesematics.<br=""><http: wwww.scambridgesematics.<br=""><http: www.scambridgesematics.<br=""><http: td="" wwwwwwwwwwwwwwwwwwwwwwwwwwwwwwwwww<=""><td>LSchema#&gt; . 1/&gt; . 1/&gt; . s/l.1/&gt; . rms/&gt; . //22-rdf-symtax-ns#&gt; . /valles/2008/07/Asstem#&gt; . logies/2008/07/Asstem#&gt; . logies/2008/07/Asstem#&gt; . com/ontologies/2009/05/Li antics.com/ontologies/6rs 6-94a7-lef9dfble65c&gt; { -42f6-94a7-lef9dfble65c&gt; { -42f6-94a7-lef9dfble65c&gt; { -fillin]; o_fillin];</td><td>nkedData#&gt; . phmarts#&gt; . a ld:ClearWorkingEditionRequest ;</td><td></td><td>•</td></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></td></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></ht </http: </http </http: </htt </htt </http 	://www.w3.org/2001/XM p://xmlns.com/fodf/0. p://wmlw3.org/2000/0 //purl.org/dc/element http://purl.org/dc/element http://purl.org/dc/element http://openanzo.org/ong/07 ttp://openanzo.org/onto p://openanzo.org/onto //cambridgesematics. <http: cambridgesematics.<br=""><http: cambridgesematics.<br=""><http: cambridgesematics.<br=""><http: cambridgesematics.<br=""><http: cambridgesematics.<br=""><http: cambridgesematics.<br=""><http: www.scambridgesematics.<br=""><http: wwww.scambridgesematics.<br=""><http: www.scambridgesematics.<br=""><http: td="" wwwwwwwwwwwwwwwwwwwwwwwwwwwwwwwwww<=""><td>LSchema#&gt; . 1/&gt; . 1/&gt; . s/l.1/&gt; . rms/&gt; . //22-rdf-symtax-ns#&gt; . /valles/2008/07/Asstem#&gt; . logies/2008/07/Asstem#&gt; . logies/2008/07/Asstem#&gt; . com/ontologies/2009/05/Li antics.com/ontologies/6rs 6-94a7-lef9dfble65c&gt; { -42f6-94a7-lef9dfble65c&gt; { -42f6-94a7-lef9dfble65c&gt; { -fillin]; o_fillin];</td><td>nkedData#&gt; . phmarts#&gt; . a ld:ClearWorkingEditionRequest ;</td><td></td><td>•</td></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:></http:>	LSchema#> . 1/> . 1/> . s/l.1/> . rms/> . //22-rdf-symtax-ns#> . /valles/2008/07/Asstem#> . logies/2008/07/Asstem#> . logies/2008/07/Asstem#> . com/ontologies/2009/05/Li antics.com/ontologies/6rs 6-94a7-lef9dfble65c> { -42f6-94a7-lef9dfble65c> { -42f6-94a7-lef9dfble65c> { -fillin]; o_fillin];	nkedData#> . phmarts#> . a ld:ClearWorkingEditionRequest ;		•

 Toward the bottom of the request, replace the <temp://value\_to\_fill\_in\_1> placeholder URI with the URI for the dataset.

```
<http://serviceRequesta9e72cca-a79b-42f6-94a7-1ef9dfb1e65c> {
    <http://serviceRequesta9e72cca-a79b-42f6-94a7-1ef9dfb1e65c> a
ld:ClearWorkingEditionRequest ;
    ld:fldsToClear <temp://value_to_fill_in_1> ;
    ld:typesToClear <temp://value_to_fill_in_0> .
```

```
<temp://value_to_fill_in_1> a ld:FileBasedLinkedDataSet .
```

#### }

#### For example:

```
<http://serviceRequesta9e72cca-a79b-42f6-94a7-1ef9dfb1e65c> {
    <http://serviceRequesta9e72cca-a79b-42f6-94a7-1ef9dfb1e65c> a
ld:ClearWorkingEditionRequest ;
    ld:fldsToClear
<http://csi.com/FileBasedLinkedDataSet/ee8d3d5792fd218a03b70fdf850b6a4c> ;
    ld:typesToClear <temp://value_to_fill_in_0> .
<http://csi.com/FileBasedLinkedDataSet/ee8d3d5792fd218a03b70fdf850b6a4c> a
ld:FileBasedLinkedDataSet .
}
```

6. Comment out the ld:typesToClear <temp://value\_to\_fill\_in\_0> line. For
 example:

7. Click the **Run Service** button to clear the edition. Anzo returns a response such as the following example when the request is processed:

```
@prefix n-1060687345: <http://openanzo.org/ClearWorkingEditionResponse/> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix ld: <http://cambridgesemantics.com/ontologies/2009/05/LinkedData#> .
@prefix ss: <http://openanzo.org/ontologies/2008/07/SemanticService#> .
```

}

```
n-1060687345:effbda5b-ce9c-4190-abb1-6f6efd07f5d8 {
    n-1060687345:effbda5b-ce9c-4190-abb1-6f6efd07f5d8 ld:wasWorkingEditionCleared
    "true"^^<http://www.w3.org/2001/XMLSchema#boolean> .
    n-1060687345:effbda5b-ce9c-4190-abb1-6f6efd07f5d8 rdf:type
ld:ClearWorkingEditionResponse .
    n-1060687345:effbda5b-ce9c-4190-abb1-6f6efd07f5d8 rdf:type ss:ServiceResponse
}
```

Now, if you browse the Managed Edition for the dataset, you will see that the edition does not contain any data components.

# **Working with Graphmarts**

The topics in this section provide guidance on working with graphmarts.

In this section:	
Creating a Graphmart	
Copying a Graphmart	
Graphmart Settings Reference	
Creating an Elasticsearch Index from a Graphmart	
Adding Data Layers to Graphmarts	
Adding Steps to Layers	
Creating Data on Demand Endpoints	
Sharing Access to Graphmarts	
Graphmart FAQ	601

# **Creating a Graphmart**

Besides importing an exported graphmart, there are four ways to create a new graphmart:

- 1. You can create a graphmart directly from a connected data source. For instructions, see Creating a Graphmart from a Data Source.
- 2. You can create a graphmart from a dataset in the Datasets catalog. For instructions, see Creating a Graphmart from a Dataset.
- 3. You can create a graphmart by copying an existing one. For instructions, see Copying a Graphmart.
- 4. Or you can create an empty graphmart and build it from scratch by adding data layers and steps that load and transform data. To create an empty graphmart, go to the Graphmarts screen in the Anzo application, click Add Graphmart, and select Create New Graphmart. See Graphmart Settings Reference for information about graphmart settings, and see Adding Data Layers to Graphmarts for information about adding layers and steps.

# **Copying a Graphmart**

Follow the instructions below if you want to create a graphmart by making a copy of an existing graphmart and its data layers and steps. For instructions on creating a new graphmart from scratch, see Creating a Graphmart.

1. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

<b>Y</b>	<b>Q</b> Search	=				Add Graphr	mart 👻
•	Title 个	Description	Status	Statements	Creator	Actions	÷
	🕸 Books		C Ready to use	156	System Administrator	Д	:
	Flights		C Ready to use	265,339	System Administrator	Д	:
	≋ Movies		C Ready to use	104,166	System Administrator	Д	:
	© Northwind		C Ready to use	131,386	System Administrator	Д	:
	Tickets		C Ready to use	3,696,548	System Administrator	Д	:
				Roy	ws per page: 25 👻 1-5 o	nf 5 <	>

2. On the Graphmarts screen, click the name of the graphmart that you want to copy. Anzo displays the graphmart Overview. At the top of the screen, click the **Save a Copy** button.



The Clone dialog box is displayed. For example:

Clone		
Clone Title		
Flights (Cloned 11/02/2022 10:1	2 AM)	
Clone Description		
(Cloned 11/02/2022 10:12 AM)		
	CANCEL	SAVE A COPY

- 3. In the Clone dialog box you have the option to edit the new graphmart name in the **Clone Title** field and modify the description in the **Clone Description** field.
- 4. When you are ready to save the copy, click **Save a Copy**.

Anzo creates and displays the new graphmart. You can proceed with modifying the graphmart settings, changing and adding datasets, data layers, or Data on Demand endpoints, or activating the new graphmart.

# **Graphmart Settings Reference**

This topic describes the graphmart configuration settings that are available on the graphmart

#### Overview tab.



Setting	Description
Load Priority	If you want Anzo to prioritize the order in which graphmarts are activated when reconnecting to AnzoGraph or resetting and reloading AnzoGraph, you can designate a Load Priority for each graphmart. When reloading AnzoGraph, Anzo activates the graphmarts in sequence, starting with the lowest Load Priority number. The default value is <b>100</b> .
Metrics	<b>Profile Data Upon Activation</b> This setting is disabled by default and controls whether a data profile is automatically generated each time the graphmart is activated. For information about data profiles, see Generating a Graphmart Data Profile.
Disable Load	Disable Populating Counts During Load Operations

Setting	Description
Counts	This setting is disabled by default controls whether Anzo periodically sends select (count(*) as ?count) queries to AnzoGraph to count the total number of statements that are being loaded in each data layer. Disabling the load counts may increase load performance as it decreases the number of queries that run during graphmart activation.
Data Loading Settings	Load Layers that Do Not Fail This setting is enabled by default and controls what to do if a data layer fails during graphmart activation. When enabled (the default setting), the graphmart is configured to load all layers that succeed and skip any layers that fail. When disabled, the entire graphmart activation is aborted if any layer fails.
Ignore Source Errors	<b>Ignore any Step or View's Source Layer if they Failed to Load</b> This setting is enabled by default and controls what to do if a source that is referenced by a step or view fails to load. For example, if the source for a Query Step is set to "All Previous Layers Within Graphmart" and one of the previous layers fails to load, this setting controls whether to run the Query Step but ignore the failed layer or fail the step since one of the sources failed.
Ignore Source Errors	If Ignore Source Errors is enabled (the default setting), Anzo ignores the failed source and runs the step against the sources that did not fail. For example, if \${usingSources} in a Query Step translates to USING <layer1> USING <layer2> USING <layer3></layer3></layer2></layer1>
	And layer1 failed to load, Anzo runs the Query Step but ignores layer1 and automatically changes the query to

Setting	Description
	USING <layer2> USING <layer3></layer3></layer2>
	If <b>Ignore Source Errors</b> is disabled and a source layer fails, any steps or views that use that source will also fail since the source is not available.
Leave Graphmart Online During Refresh	When Refreshing Layers, Leave Graphmart and Layer Online This setting is disabled by default and controls whether a graphmart remains online while it is being refreshed in AnzoGraph. When this option is enabled, if a user clicks the <b>Refresh</b> button to refresh a graphmart (or the Refresh icon on a data layer), Anzo copies the existing layers into temporary graphs so that the data remains online while the original graphs are refreshed. When the refresh is complete, the temporary graphs are deleted.
	<b>Note</b> This setting applies only to <b>Refresh</b> operations. If <b>Leave Graphmart</b> <b>Online During Refresh</b> is enabled and a user clicks <b>Reload</b> , the data layers will not remain online. During reloads all of the data is dropped and then loaded again.
Manual Refresh Graphmart	Once Loaded, Changes only Pushed to AnzoGraph Manually (Only Affects Journal Based Data)
	in this graphmart are automatically deployed to AnzoGraph without requiring a manual refresh or reload of the graphmart. This setting only applies to graphmarts with Load Dataset Steps that load a journal-based data set, such as a system metadata graph. When this option is enabled, changes to the journal-based data set are only deployed to AnzoGraph when the graphmart is manually reloaded or refreshed. When this option is

Setting	Description
	disabled, changes to the dataset are automatically loaded to AnzoGraph without requiring a manual refresh.
Generate Statistics After Each Layer	Call Generate Statistics after Loading Each Layer, Instead of Only at the End Typically the AnzoGraph connection is configured to automatically initiate AnzoGraph's internal statistics gathering queries after loading a graphmart. However, if a user refreshes individual layers rather than the entire graphmart, those queries are not triggered. Enabling this setting initiates the statistics gathering queries each time a layer is loaded. This helps the AnzoGraph query planner generate ideal query execution plans for queries that are run against the refreshed data layers.
Require Elasticsearch	This Graphmart Requires Elasticsearch This setting is disabled by default. If you plan to include unstructured datasets in this graphmart or configure a data layer that creates an Elasticsearch index, you can enable this option to ensure that Anzo validates the connection to Elasticsearch whenever this graphmart is activated, reloaded, or refreshed.

# **Creating an Elasticsearch Index from a Graphmart**

By associating an Elasticsearch index with a data layer, you can load data from a graphmart to an Elasticsearch index, enabling you to perform free-text and pattern searches on your knowledge graphs. This topic provides instructions for configuring a workflow that generates an Elasticsearch index and snapshot from a graphmart.

- Prerequisites
- Add an Export Step to the Graphmart
- Add a Layer to Manage the Index
- Add a Step to Create the Index
- Add a Step to Take a Snapshot of the Index

### **Prerequisites**

Before configuring a graphmart to create an Elasticsearch index, make sure that the following requirements are met:

- 1. A supported version of Elasticsearch is installed and configured. For more information, see Elasticsearch Requirements in the Deployment Guide.
- 2. The Elasticsearch instance is connected to Anzo. See <u>Connecting to Elasticsearch</u> in the Administration Guide for more information.
- The AnzoGraph instance that you will load the graphmart to is also connected to Elasticsearch. Configure the connection by selecting the Elasticsearch instance in the Elasticsearch Configuration field in the AnzoGraph configuration. For more information, see <u>Connecting to AnzoGraph</u> in the Administration Guide.

# Add an Export Step to the Graphmart

First, if a file-based linked data set (FLDS) has not been generated for this graphmart, add an **Export Step** to the **last layer** in the graphmart so that all of the graph data is exported to an FLDS on the file store. For instructions on adding an Export Step, see Export Data to an FLDS (Export

#### Step).

If the graphmart has an Export Step, note the Target FLDS for the step and proceed to Add a Layer to Manage the Index below.

#### Add a Layer to Manage the Index

Follow the steps below to add a layer to a graphmart and configure it to manage an Elasticsearch index.

- Add a new data layer to the graphmart. This new layer will be associated with the Elasticsearch index and contain steps to generate the index and snapshot. For instructions on adding a layer, see Creating a New Layer.
- In the new layer, click the Elasticsearch Indexing tab. Then select the Associate an Elasticsearch index with this layer checkbox. When this setting is enabled, Anzo creates an index on the Elasticsearch server and links this layer to that index.



#### Note

Though the index will be populated with data from other layers in the graphmart, you do not need to modify those layers to associate them with the index. **Associate an Elasticsearch index with this layer** should be disabled on all layers that do not contain Elasticsearch Indexing Steps.

3. The Advanced settings on this screen relate to managing the index's life cycle. It is not necessary to modify the settings to successfully generate an index for the graphmart. If you have an advanced use case, such as a case that requires linking this layer to an existing index or needing to add custom Elasticsearch-specific index or mapping configurations, expand

Advanced to access the options. The list below describes the available settings.

	Details	Execution Condition	Data Access	Query Context	Elasticsearch Indexing
Asso	ciate an Elastic	search index with this layer	r Changao may signif	icently offect Electicocours	ek bebeuler
Advanced A	ese are advanced	i setungs. Proceed with cauto	n. Changes may signi	icantiy arrect clasticsear	cir benavior.
Elasticsean Name wi	ch Index Name II be system ge	enerated if left blank			
Name will I	e system generate	d if left blank	ecutes 🔽 Delete I	Elasticsearch index on	layer unload
Elasticse	arch Index Set	tings			
Settings (J	SON format) for Ela	asticsearch index			
Elasticse	arch Index Ma	pping			
Mapping (J	SON format) for El	asticsearch index			

- Elasticsearch Index Name: By default, the new index is given a system-generated name. If you are linking this layer to an existing index, add the existing index name to this field. Or, if you plan to reference this index elsewhere and want to give it a human-readable name, you can specify a custom name.
- Clear Elasticsearch index before this layer executes: This option is enabled by default and configures the layer so that the index is cleared and recreated each time this layer is run.
- Delete Elasticsearch index on layer unload: This option is enabled by default and configures the layer so that the index is deleted any time the layer is deactivated or reloaded. The index is recreated during the reload or when the graphmart is activated again. If you do not want the index to be deleted when the layer is offline, clear the checkbox.
- Elasticsearch Index Settings: This field can be used to add any Elasticsearch-specific index settings that you want to apply. Add the settings in the following JSON format:

{			
	"index":	{	

```
"<settings_and_values>"
}
```

#### For example:

• Elasticsearch Index Mapping: This field can be used to add any Elasticsearch-specific mapping properties that you want to apply. Add the properties in the following JSON format:

```
{
   "properties": {
    "<properties_and_values>"
   }
}
```

For example:

```
{
  "properties": {
   "movie_Abstract": { "type": "keyword" },
   "movie_ID": { "type": "long" },
   "movie_Title": { "type": "text" }
  }
}
```

4. When you have finished configuring the layer, click **Save** to add it to the graphmart and return to the Data Layers screen. Then proceed to Add a Step to Create the Index below.

# Add a Step to Create the Index

Follow the steps below to add an Elasticsearch Indexing Step to the new layer and configure it to generate an index.

- On the Data Layers screen, click the menu icon (1) for the new layer and select Add Step/View.
- 2. On the Add Step/View screen, select **Elasticsearch Indexing Step** and click **OK**. The Create dialog box is displayed.
- 3. On the Details tab, add a name for the step in the **Title** field, and configure any optional settings. For details about the settings, see Create an Elasticsearch Index (Elasticsearch Indexing Step).

Details	Query	Execution Condi	tion >
Title *			
The title of the step			
Description			
A brief description of the step			
Enabled			
Source			
Self ×			x   ~
Source data to act upon			
Data models			÷ ~
Associated data models			
		CANCEL	SAVE

4. When you have finished configuring the Details tab, click the Query tab. The tab includes a template for writing a SPARQL SELECT query that incorporates the Graph Data Interface (GDI) service to generate an index on the Elasticsearch server.



Unlike queries for other steps, this step does not run an INSERT query because the data is not being inserted into AnzoGraph. Edit the template as needed. You can click the **Preview in Query Builder** button to open the query in the Query Builder, where you can perform practice runs to see results without having to refresh the layer. For more information about writing GDI queries against an Elasticsearch source, see Querying an Elasticsearch Source.

#### Important

Do not include Elasticsearch connection or index parameters in the query. Anzo automatically populates that information from the AnzoGraph configuration when the step is run.

- 5. When you have completed the indexing query, click **Save** to save the configuration and add the step to the layer.
- 6. Next, users typically add an Elasticsearch Snapshot Step to the same layer. The step takes a snapshot of the index and saves it to the dataset (FLDS) on disk. Storing the snapshot with

the FLDS ensures that the index is included if the dataset is added to another graphmart. If you want to create a snapshot, continue to Add a Step to Take a Snapshot of the Index below.

# Add a Step to Take a Snapshot of the Index

Follow the steps below to add an Elasticsearch Snapshot Step to the new layer and configure it to save a snapshot of the index to the FLDS.

- On the Data Layers screen, click the menu icon (1) for the new layer and select Add Step/View.
- On the Add Step/View screen, select Elasticsearch Snapshot Step and click OK. The Create dialog box is displayed.

Details	Execution Condition	Query Cont	text
Title *			
The title of the step			
Description			
A brief description of the step			
Enabled			
Source			
Self ×			
This option cannot be changed.			
Target FLDS *			~
FLDS only.			
		CANCEL	SAVE

- 3. On the Details tab, add a name for the step in the **Title** field.
- 4. Click the **Target FLDS** field and select the FLDS that is created by the Export Step in the graphmart.
- 5. Configure any other optional settings on the Details tab. For information about the settings, see Take a Snapshot of an Index (Elasticsearch Snapshot Step).
- 6. Click **Save** to save the configuration and add the step to the layer.

Once you have finished configuring the new layer and steps, reload the graphmart to generate (or update) the FLDS, create the Elasticsearch index, and save a snapshot of the index to the FLDS.

# **Adding Data Layers to Graphmarts**

The topics in this section provide instructions for creating and configuring data layers. For conceptual information about graphmarts, layers, and steps, see <u>Graphmart Concepts</u> in the Getting Started Guide.

- Creating a New Layer
- Cloning a Layer
- Using Query Contexts
- Defining Execution Conditions
- Advanced Data Access Settings

#### **Creating a New Layer**

Follow the steps below to create a new data layer in a graphmart. For instructions on creating a new layer by copying an existing one, see Cloning a Layer.

1. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

<b>Y</b> (	<b>Q</b> Search	₹ • 🖬 🗏		I	Add Graphn	nart 👻
•	Title	Status	Statements	Creator	Actions	÷
	🕸 Books	Ready to use	156	System Administrator	Д	:
	☺ Flights	Ready to use	265,339	System Administrator	Д	:
	🕸 Movies	Ready to use	104,166	System Administrator	Д	:
	➢ Northwind	Ready to use	131,386	System Administrator	Д	:
	⊯ Tickets	Ready to use	3,696,548	System Administrator	Д	:
			Rows	per page: 25 👻 1-5	of 5 <	>

2. Click the name of the graphmart that you want to add a layer to. The Overview tab is displayed. For example:



3. Click the Data Layers tab. Anzo displays the existing data layers. For example:

*	Not Vers	thwind	ACTIVE @ Ready to us	se 🛆 AZG State	Ð	Save a Copy	Create Dashboard	Ω
	Ove	rview	Explore	Datasets	Data Layers	Dashboards	Data on Demand	>
C RE	LOAD	ø HIDE :	SIMPLE LOAD LAYERS	EXPAND ALL			ADD	-
: >	0	1. DB-QA La Load data from	iyer 🤗 DB-QA				-	:
: >		2. Create Da	ataset					:

4. Click Add and select New Layer. Anzo displays the Create dialog box.

Create						20
	Details	Execution Condition	Data Access	Query Context	Elasticsearch	>
Title t						
The title of t	he layer					
Descriptio	on					
A brief desc	ription of the layer					
Aut	o Deploy Ontology	y Changes 🔲 Disable Lo	oading Counts			
					CANCEL S	AVE

- 5. Specify a name for the layer in the **Title** field and an optional description in the **Description** field.
- 6. Determine how you want to control changes to this layer's dependent data models:
  - If you want Anzo to automatically deploy to AnzoGraph any changes to the related models without having to manually refresh the layer or graphmart, select the Auto Deploy Ontology Changes checkbox.

#### Note

The **Manual Refresh Graphmart** setting on the graphmart must be **disabled** for automatic deployment of models to work. See Graphmart Settings Reference for information about graphmart settings.

- If you want model changes to be deployed to AnzoGraph only when this layer (or entire graphmart) is refreshed or reloaded, leave the Auto Deploy Ontology Changes checkbox empty (disabled).
- 7. Determine whether to Disable Loading Counts for this layer. This setting controls whether Anzo periodically sends COUNT queries to AnzoGraph while this layer is reloading or refreshing. Disabling the load counts may increase load performance as it decreases the number of queries that are executed during loads.

8. Click **Save** to add the new layer to the graphmart and return to the Data Layers screen.

The new layer becomes the last layer in the graphmart. If you want to change the order of the layers, you can click the black bar on the left side of a layer and drag the layer up or down. Data layers are processed from top to bottom.

### Tip

The Refresh icon ( $\mathcal{O}$ ) on the new layer indicates that the layer is out of sync with the data that is in AnzoGraph. Once you configure the new layer and add data processing steps, you can click the **Reload** button ( $\mathcal{O}$ ) at the top of the screen to reload the entire graphmart, or you can click the **Refresh** icon ( $\mathcal{O}$ ) on the layer to reload only that layer.

See Adding Steps to Layers for instructions on creating steps.

# **Cloning a Layer**

Follow the steps below to create a new data layer by copying an existing one from any graphmart. For instructions on creating a layer from scratch, see Creating a New Layer.

1. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

<b>Y</b>	<b>Q</b> Search	=	<b>iii</b> =			Add Graphn	nart 👻
•	Title 个	Description	Status	Statements	Creator	Actions	÷
	🕸 Books		C Ready to use	156	System Administrator	Д	:
	Flights		C Ready to use	265,339	System Administrator	Д	:
	≋ Movies		Ready to use	104,166	System Administrator	Д	:
	© Northwind		C Ready to use	131,386	System Administrator	Д	:
	© Tickets		C Ready to use	3,696,548	System Administrator	Д	:
				Rov	ws per page: 25 👻 1-5	i of 5 <	>

2. Click the name of the graphmart that you want to add a layer to. The Overview tab is displayed. For example:

**	Northwind	ACTIVE Ø Ready to use	△ AZG Static		Save a Copy	Create Dashboard		
	Overview	Explore	Datasets	Data Layers	Dashboards	Data on Demand 🔉		
escription		_						
lone						General		
ad Driarit	ha					Type : Graphmart		
one	y.				1.00	Creator : System Administrat		
one					Las	st Updated : Shours ago		
					Structur	e Modified : 6 hours ago		
etrics						Released : 6 hours ago		
ofile data	upon activation 🛛 💻							
isable L	oad Counts				http://car	nbridgesemantics.com/Graph 📋		
sable pop	ulating counts during loa	ad operations			Dynamic	Dynamic Resource Log		
ata Loa	ding Settings				An	AnzoGraph Server Details		
ad layers	that do not fail 🛛 💻					AnzoGraph: AZG		
	-					Status : Ready to use		
nore So	ource Errors				Las	Accessed : 3 hours ago		
nore any s	step or view's source laye	r if they failed to load 🛛 💻			Me	mory Used: 280.5 MB(3%)		
eave Gr	aphmart Online Du	ring Refresh			Me	mory Total: 11.7 GB		
hen refres	thing layers, leave graphs	nart and laver online						
1011101100	und where, icone diabuti				Inactivity	Deactivate Timeout		
anual F	Refresh Graphmart				None			
nce loade	d, changes only pushed to	o AnzoGraph manually.(Only affects .	lournal Based Data) 🦳 💻					

3. Click the **Data Layers** tab. Anzo displays the existing data layers. For example:

*	Northwind	ACTIVE Ø Ready to us	e 🛆 AZG Stati	0	Save a Copy	Create Dashboard	D
	Overview	Explore	Datasets	Data Layers	Dashboards	Data on Demand	>
C RE	ELOAD 🛷 HIDE S	IMPLE LOAD LAYERS	EXPAND ALL			ADD	•
: >	I. DB-QA Lay	rer 🤗 B-QA				-	:
: >	2. Create Dat	laset					:

4. click **Add** and select **Add Existing**. Anzo opens the Select a value to add dialog box, which lists the existing layers for all graphmarts. For example:

Select a value to add					
T Q Search	₹ • 🔳		l	Hide simple lo	ad layers
Title 个	Description	Updated Date	Tags	Actions	÷
🃚 Create Dataset		Nov 16, 2022		Д	:
📚 Create Dataset		Nov 16, 2022		Д	:
📚 DB-QA Layer	Load data from DB-QA	Nov 16, 2022	Auto-Gen	Д	:
📚 Flights Løyer	Load data from Flights	Nov 16, 2022	Auto-Gen	Д	:
S Movies Layer	Load data from Movies	Nov 16, 2022	Auto-Gen	Д	:
S Movies Layer	Load data from Movies	Nov 16, 2022	Auto-Gen	Д	:
📚 Tickets Layer	Load data from Tickets	Nov 16, 2022	Auto-Gen	Д	:
			Rows per page: 25 👻	1-8 of 8	$\rightarrow$
				CANCEL	ОК

5. Select the layer that you want to copy and click **OK**. The Clone dialog box is displayed. For example:

Clone		
Clone Title		
Flights Layer (Cloned 11/16/2022	05:31 PM)	
Clone Description		
(Cloned 11/16/2022 05:31 PM)		
	CANCEL	SAVE AS

- In the Clone dialog box, you have the option to edit the Clone Title and/or Clone Description.
- 7. Click Save As to add the cloned layer and any steps that the layer contains to the graphmart.

The new layer becomes the last layer in the graphmart. If you want to change the order of the layers, you can click the black bar on the left side of a layer and drag the layer up or down. Data layers are processed from top to bottom.

# Тір

The Refresh icon ( $\mathcal{O}$ ) on the new layer indicates that the layer is out of sync with the data that is in AnzoGraph. Once you configure the new layer and add data processing steps, you can click the **Reload** button ( $\mathcal{O}$ ) at the top of the screen to reload the entire graphmart, or you can click the **Refresh** icon ( $\mathcal{O}$ ) on the layer to reload only that layer.

See Adding Steps to Layers for instructions on creating steps.

# **Using Query Contexts**

When you use the Graph Data Interface (GDI) for querying data sources, you may connect to sources that require input of sensitive connection and authorization information such as keys, tokens, and user credentials. When configuring a step that runs a GDI query, Cambridge Semantics recommends that you refer to a Query Context whenever possible. Contexts store sensitive information as key-value pairs. Queries reference only the keys from the context and the sensitive values are abstracted from the requests that are sent to the data source and AnzoGraph. This topic provides information on configuring Query Contexts and referring to Context Variables in a query.

- Overview of Query Contexts
- Referencing Context Variables in a Query

#### **Overview of Query Contexts**

Query contexts are accessible from the **Query Context** tab that is available when creating or editing a data layer or step. The image below shows the Query Context tab for a layer.

	Details	Execution Condition	Data Access	Query Context	Elasticsearch Indexing		
Select the Context Providers that will be made available to queries via Context Variables.							
Context Pr	oviders						
Reference ex	xisting database	connections to provide conne	ction information for s	ervices accessed during	the query.		
Search to	add Contaxt D	rovidoro					
Search to	add Context P	roviders			· · ·		
Context Va	ariahlee						
Access Cont	text Variables for	each Context Provider added	above.				
Context Br	ovider						
Context Pr	ovider	· ·					
Custom Co	ontext				Add Key		
Name			Value				
None							
					CANCEL SAVE		

#### **Context Providers**

Connections in Anzo implement the **Context Provider** interface. For example, file store connections and data source connections provide contexts (in the form of JSON objects) that contain key-value pairs. The contexts contain the source connection details such as URLs, database names, user names, passwords, and tokens. A context is passed to the data source when a request is made against that source. To use one of the Anzo-generated Context Providers that was created for a pre-existing connection, select that provider from the drop-down list. When you select a provider, the variables from that context are displayed under Context Variables, as shown in the image below:
	Details	Execution Condition	Data Access	Query Context	Elasticsearch Indexing	
Select the C	ontext Provid	lers that will be made avail	able to queries via C	Context Variables.		
Context Pro Reference exis	<b>viders</b> sting database	connections to provide conne	ction information for s	ervices accessed during 1	he query.	
Context Provide	rs					
DB-QA ×					×	~
Dontext Provide	r •	<pre>#DB-QA db.718cc8ab3e9fb2b22345a db.718cc8ab3e9fb2b22345a db.718cc8ab3e9fb2b22345a db.718cc8ab3e9fb2b22345a db.718cc8ab3e9fb2b22345a db.http://cambridgeseman db.http://cambridgeseman db.http://cambridgeseman</pre>	d4dc31e685c.label d4dc31e685c.password d4dc31e685c.wrver d4dc31e685c.wrl d4dc31e685c.wrl d4dc31e685c.wser tics.com/MySQLDatabase tics.com/MySQLDatabase	:DataSource/30eb8be931f94 DataSource/30eb8be931f94 DataSource/30eb8be931f94	140c8e570625709d1377.label 140c8e570625709d1377.passwo	rd (
		db.http://cambridgeseman db.http://cambridgeseman	tics.com/MySQLDatabase tics.com/MySQLDatabase	DataSource/30eb8be931f94 DataSource/30eb8be931f94 DataSource/30eb8be931f94	40c8e570625709d1377.url 40c8e570625709d1377.user	
Custom Con	itext				Ad	d Key
					CANCEL	SAV

## **Context Variables**

When you select a Context Provider, the variables from the selected context are displayed under Context Variables. You can copy the text and then use any key in the list as a variable in any query that connects to this data source. Typically there are two versions of each variable. Either version can be used in a query. The short versions are generated from the long versions for ease of use and readability in queries. For information about referencing variables in a query, see Referencing Context Variables in a Query below.

## **Custom Context**

Custom Contexts are user-defined key-value pairs that are not associated with a particular Context Provider. To add a key and define its value, click the **Add Key** button. Then specify the **Key Name** and **Key Value** in the Create Context Key dialog box. Click **Create** to add the keyvalue pair to the context.

Create Context Key		
Key Name *		
Key Value *		
	CANCEL	CREATE

## Tip

When defining custom keys at the layer level, make sure that the key names are unique. A layer may have different steps that connect to different sources but reference the same custom context. The image below, for example, creates URL, username, and password Context Keys for a MySQL database.

ontext Keys		Add Key
Name	Value	
mysql-password	admin123	Z 1
mysql-username	sysadmin	2 T
mysql-northwind-url	jdbc:mysql://10.111.4.9:3306/NORTHWIND	21

**Referencing Context Variables in a Query** 

The format that you use for referencing a Context Variable in a query depends on the type of AnzoGraph plugin or extension that is being called by the query. Generally, contexts are only used in steps that contain GDI queries. When referencing context keys in GDI queries, use the following format:

{{@context\_key\_name}}

For example, the following GDI query references three variables from the Context Provider shown above: url, user, and password.

```
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
SELECT DISTINCT *
WHERE
{
SERVICE <http://cambridgesemantics.com/services/DataToolkit>
```

```
?data a s:DbSource ;
          s:url "{{@db.718cc8ab3e9fb2b22345ad4dc31e685c.url}}";
          s:username "{{@db.718cc8ab3e9fb2b22345ad4dc31e685c.user}}";
          s:password "{{@db.718cc8ab3e9fb2b22345ad4dc31e685c.password}}";
          s:selector "[dbo].[FILM]" ;
          ?year ("[YEAR]" xsd:int);
          ?length (xsd:int) ;
          ?title (xsd:string) ;
          ?subject ("[dbo].[FILM].[SUBJECT]" xsd:string) ;
          ?actor ("[ACTOR]" xsd:string) ;
          ?actress (xsd:string) ;
          ?director (xsd:string) ;
          ?popularity (xsd:int) ;
          ?awards (xsd:string) ;
          ?image (xsd:string) .
      FILTER(?year >= 1990 && ?year < 2000)
      FILTER(?subject = "Drama" || ?subject = "Action")
      FILTER(?length <= 90)</pre>
}
```

## **Defining Execution Conditions**

Execution conditions can be defined at the data layer or step level and are used to conditionalize the execution of the layer or step based on the result of a specified Validation Condition.

#### Note

}

Validation Conditions are defined in Validation Steps. In order to set up an Execution Condition, the Graphmart needs to have at least one **Validation Step** that defines a **Condition Variable**. Validation Conditions can be used across all Data Layers in the same Graphmart. For guidance on configuring a Validation Step, see Validate the Data (Validation Step).

This topic focuses on configuring an execution condition at the data layer level.

## **Configuring an Execution Condition**

Execution Conditions are configured from the **Execution Condition** tab that is available when creating or editing a layer or step. The image below shows the Execution Condition tab for a layer.

🗌 Ena	ble layer based	on boolean condition	
Select c	ondition variable	e	
onditional	Variable		
Execute	only if returned	value is	~
	Variable If Decult		
onditional	variable il Result		
onditional	vanable il Result		
onditional	vanable il Result		
conditional	vanadie il Result		

## **Enable Layer Based on Boolean Condition**

This setting indicates whether to enable this data layer only if the returned value from the Validation Condition is either **true** or **false**. You specify true or false in the **Conditional Variable If Result** field. If the Validation Condition fails, the layer is disabled.

## **Conditional Variable**

This field specifies the variable that you want to base this execution condition on. If the list is empty, that means either there are no Validation Steps in the graphmart or a Validation Step exists but it does not include a Condition Variable (defined on the Options tab). Refer to Validate the Data (Validation Step) for more information.

## **Conditional Variable If Result**

If you enabled the Enable Layer Based on Boolean Condition setting, select true or false from the drop-down list. The data layer will be enabled only if the result of the Validation Step Query matches the value that you specified. If Enable Layer Based on Boolean Condition is disabled, leave this field blank.

## **Advanced Data Access Settings**

This topic provides reference information about the Advanced Data Layer Hi-Res Analytics settings that control how a layer is exposed to and affects Hi-Res Analytic dashboards.

#### Important

Changing these settings can have unexpected consequences.

The Advanced Hi-Res Analytics settings are available on the **Data Access** tab when you create or edit a data layer:



The sections below describe each of the available settings:

- Display data outputted by the layer in Hi-Res Analytics dashboards
- Display Layer in Data Layers in Hi-Res Analytics
- User can both view and hide the data outputted by the layer
- By default, the data outputted by the layer is visible
- Hide from Other Layers

Display data outputted by the layer in Hi-Res Analytics dashboards

This setting controls whether the data generated by the steps is available to query and display in Hi-Res Analytics dashboards:

- When the setting is **enabled** (the default value), the layer's data is available to dashboards.
- When the setting is **disabled**, other data layers in the graphmart can use the layer's data, but the data is not available to dashboards.

**Display Layer in Data Layers in Hi-Res Analytics** 

This setting controls whether Anzo displays the layer in the Data Layers panel on Hi-Res Analytics dashboards. The image below shows an example Data Layers panel:

∧ Data Layers	
Tickets Layer	
View	
Export	

- When the setting is **enabled** (the default value), the layer is listed in the Data Layers panel in dashboards.
- When the setting is **disabled**, the layer's data is always used in dashboards but users do not see the layer listed in the Data Layers panel.

User can both view and hide the data outputted by the layer

This setting controls whether users have the option to show and hide the layer in the Data Layers panel on dashboards:

- When the setting is **enabled** (the default value), the layer is listed in the Data Layers panel in dashboards and users have the option to show and hide the layer.
- When the setting is **disabled**, whether the layer shows up in the Data Layers panel depends on the **By default**, **the data outputted by the layer is visible** setting. If the layer is visible in the Data Layers panel ("By default, the data outputted by the layer is visible" is enabled), users cannot toggle it on and off.

By default, the data outputted by the layer is visible

This setting controls whether the data that is generated by the steps in the layer is visible in queries and Hi-Res Analytics dashboards:

- When the setting is **enabled** (the default value), the layer is listed in the Data Layers panel in dashboards and is selected by default. The layer's data is also included by default when queries are run against the graphmart.
- When the setting is disabled, the layer shows up in the Data Layers panel but is not selected. To include the layer's data in Hi-Res Analytic queries, users must select the layer. In addition, the layer's data is automatically excluded from queries that are run against the graphmart. To include data from the layer in results, the queries must explicitly list the layer's URI.

### Hide from Other Layers

This setting controls whether the other layers in the graphmart can act upon the source data in this layer.

- When the setting is **disabled** (the default value), this layer is available as a choice in the **Source data to act upon** drop-down list when a step's source is configured.
- When the setting is **enabled**, this layer is not listed as a choice in the **Source data to act upon** list when the a step's source configured.

# **Adding Steps to Layers**

The steps in a data layer perform the data operations, such as loading, creating, deleting, changing, or exporting data. You can add any number of steps to a layer. The topics in this section provide information about configuring each type of step.

- Directly Load a Data Source (Direct Load Step)
- Create an Elasticsearch Index (Elasticsearch Indexing Step)
- Take a Snapshot of an Index (Elasticsearch Snapshot Step)
- Export Data to an FLDS (Export Step)
- Load a Dataset from the Catalog (Load Dataset Step)
- Pre-Compile a Query (Pre-Compile Query Step)
- Create a Reusable Query Template
- Run a Transformation Query (Query Step)
- Infer New Data (RDFS+ Inference Step)
- Validate the Data (Validation Step)
- Construct a View of the Data (View Step)

## Directly Load a Data Source (Direct Load Step)

With no mapping required, a Direct Load Step can be used to automatically generate a graph and model for a data source. The Direct Load Step is the only type of step with the ability to manage generated models. A model that is generated by a Direct Load Step is automatically registered in Anzo is linked to and managed by the layer that contains the step. If a query is changed, additional Direct Load Steps are added to the same layer, or the underlying source schema changes, the managed model is automatically updated when the graphmart is reloaded or refreshed. Follow the steps below to create a Direct Load Step.

- 1. Go to the graphmart for which you want to add a step and then click the **Data Layers** tab.
- On the Data Layers tab, find the layer that you want to add the step to. Click the menu icon (1) for that layer and select Add Step/View. The Add Step/View dialog box is displayed with the New tab selected.

Add Step	o/View						
N	lew	Existing Steps	Existing Views				
<b>T</b> Q	Search	= -					
	Direct Load	Step					^
Ð	Load data direc	ctly from an external source					-11
							-1
	Elasticseard	ch Indexing Step					- 11
la	Create a single	Elasticsearch index to associa	te with a layer.				
	Elasticseard	ch Snapshot Step					
<b>O</b>	Create an Elast	ticsearch snapshot of the index	associated with a layer.				
	Export Step	•					
Ē.	This allows you	u to export contents of graphma	art to an fids.				
			Rows per page:	25 👻	1-12 of 12	<	>
					CANO	EL	ОК

3. To create a new Direct Load step, select **Direct Load Step** and then click **OK**. If you want to clone an existing step, click the **Existing Steps** tab, select the step that you want to clone, and then click **OK**. Anzo creates or clones the step and displays the Details tab:

iitle *		
e title of the step		
scription		
rief description of the step		
Enabled		
Self ¥		×   ~
urce data to act upon		
ita models		

- 4. On the Details tab, configure the following options as needed:
  - Title: The required name of the step.
  - **Description**: An optional short description of the step.
  - **Enabled**: When creating a new step, the Enabled option is selected by default, indicating that the step is enabled and will run when the layer is loaded or refreshed. If you want to disable the step so that it is not processed, clear the **Enabled** checkbox.
  - **Source**: The source data that this step should act upon. Steps can build upon the data generated by steps in other layers or can be self-contained, applying changes that relate only to the data defined in the layer that contains this step. You can select any number of the following options:
    - **Self**: This option is selected by default and means that the step runs against the data that is in the parent layer.
    - All Previous Layers Within Graphmart: This option means that the step runs against the data in all of the successful layers that precede the layer this step is in. Any failed layers are ignored.
    - **Previous Layer Within Graphmart**: This option means that the query runs against the data that is in the one layer that precedes the layer this step is in.
    - *Layer Name*: The Source drop-down list also includes options for layer names.
       You can choose a specific layer to act upon the data in that layer.
  - Data models: This optional field specifies the model or models to associate with this step. By default, Managed is selected. If you are onboarding a source that does not have a model, make sure Managed remains specified so that the step generates a model. See Important Notes on Managed Models below for more information about managed models.

The Data Models list displays all of the available models. By default, the field is set to **Exclude System Data** (\*). If you want to choose a system model, click the toggle button on the right side of the field to change it to **Include System Data** (\*). When system data is included, the drop-down list displays the system models in addition to the user-generated models.

- **Pre-Run Generate Statistics**: This option controls whether to initiate AnzoGraph's internal statistics gathering queries before running the query to pre-compile. The statistics gathering helps ensure that the AnzoGraph query planner generates ideal query execution plans for queries that are run against the graphmart.
- 5. When you have finished configuring the Details tab, click the **Query** tab. This tab defines the query that this step should run.



6. Typically Direct Load Step queries are GDI RDF and Ontology Generator queries. Using a relatively simple SPARQL query, the GDI Generators recognize the structure of a data source and automatically generate the necessary statements. Invoking the Generators is preferable when the structure of the data is very complex, such as a JSON data source with many inner repeating structures or a database with many tables and keys. When the source contains complex structures, only the required statements are generated, avoiding cross-products and optimizing query execution and memory usage. For details about writing GDI Generator queries, see GDI Generator Query Syntax.

### Note

If your query connects to a source that requires input of connection and authorization information, Cambridge Semantics recommends that you do not include the connection and authorization values directly in the query. Instead, replace those values with Context Variables from a Query Context. You can access Context Providers for each data source from the step's Query Context tab. For detailed information about query contexts and referencing variables in a query, see Using Query Contexts.

7. When you have finished writing the query, click **Save** to save the step configuration.

Once the Details tab is configured and the query is written, the step can be run. For information about running this step conditionally by setting up an execution condition, see Defining Execution Conditions.

#### **Important Notes on Managed Models**

Though an ontology that is generated in a Direct Load Step is registered in Anzo and is available for viewing in the Model editor, **the model is owned and managed by the data layer that contains the Direct Load Step**. That means any manual changes made to the model outside of the step, such as from the Model editor, will be overwritten any time the graphmart or layer is refreshed or reloaded. **Do not modify generated managed models except by editing (or adding) Direct Load Step queries**. For information on updating managed models, see Editing a Managed Model.

**There is only one managed model per layer**. If you include multiple Direct Load Steps in the same layer, they will all update the same ontology. This functionality can be useful if you want to align the data and generated model across multiple steps. If you have multiple sources that are not intended to align or update the same model, create separate layers.

**If you delete a layer that includes a managed model, the model is also deleted**. Use caution when referencing a managed model outside of a graphmart. For example, if you create a dataset and reference a managed model when you select the ontology, the reference will break if the data layer that manages the model is deleted.

## Create an Elasticsearch Index (Elasticsearch Indexing Step)

This topic provides guidance on configuring an Elasticsearch Indexing Step that generates an Elasticsearch index from data in a graphmart. Follow the steps below to create an Elasticsearch Indexing Step.

## Tip

For an overview on creating an Elasticsearch index and snapshot from a graphmart, see Creating an Elasticsearch Index from a Graphmart.

- 1. Go to the graphmart for which you want to add a step and then click the **Data Layers** tab.
- 2. On the Data Layers tab, find the layer that you want to add the step to. Click the menu icon (I) for that layer and select Add Step/View. The Add Step/View dialog box is displayed with the New tab selected.

ħ	Vew Existing Steps Existing Views	
Q	Search	
-	Direct Load Step	
Ð	Load data directly from an external source	
	Flashanan kadada Ana	
	Elasticsearch indexing Step	
"IQ	Create a single classicisearch index to associate with a layer.	
-	Elasticsearch Snapshot Step	
0	Create an Elasticsearch snapshot of the index associated with a layer.	
	Export Step	
	This allows you to export contents of graphmart to an flds.	
	Rows per page: 25 ▾ 1-12 of 12 <	

3. To create a new Elasticsearch Indexing step, select **Elasticsearch Indexing Step** and then click **OK**. If you want to clone an existing step, click the **Existing Steps** tab, select the step that you want to clone, and then click **OK**. Anzo creates or clones the step.

If this is the first Elasticsearch-related step in the layer, the following warning is displayed:



The warning is to inform you that the Elasticsearch index-related configuration settings at the layer level are automatically adjusted to the default settings for layers that contain Elasticsearch Indexing or Snapshot Steps. Those settings are on the layer's Elasticsearch Indexing tab. The image below shows the tab with the default settings:



4. If necessary, click **OK** to close the warning. The Details tab is displayed:

Details	Query	Execution Conditi	on >
Title *			
The title of the step			
Description			
A brief description of the step			
Enabled			
Source			
Self ×		>	$\langle   \rangle$
Source data to act upon			
Data models			÷ ~
Associated data models			
		CANCEL	SAVE

- 5. On the Details tab, configure the following options as needed:
  - Title: The required name of the step.
  - Description: An optional short description of the step.
  - **Enabled**: When creating a new step, the Enabled option is selected by default, indicating that the step is enabled and will run when the layer is loaded or refreshed. If you want to disable the step so that it is not processed, clear the **Enabled** checkbox.
  - **Source**: The source data that this step should act upon. Steps can build upon the data generated by steps in other layers or can be self-contained, applying changes that relate only to the data defined in the layer that contains this step. You can select any number of the following options:
    - **Self**: This option is selected by default and means that the step runs against the data that is in the parent layer.
    - All Previous Layers Within Graphmart: This option means that the step runs against the data in all of the successful layers that precede the layer this step is in. Any failed layers are ignored.

- **Previous Layer Within Graphmart**: This option means that the query runs against the data that is in the one layer that precedes the layer this step is in.
- *Layer Name*: The Source drop-down list also includes options for layer names.
   You can choose a specific layer to act upon the data in that layer.
- Data models: This optional field specifies the model or models to associate with this step. The list displays all of the available models. By default, the field is set to Exclude System Data (\*). If you want to choose a system model, click the toggle button on the right side of the field to change it to Include System Data (\*). When system data is included, the drop-down list displays the system models in addition to the user-generated models.
- 6. When you have finished configuring the Details tab, click the Query tab. This tab contains the Elasticsearch indexing query to run. The template includes the syntax for writing a SPARQL SELECT query that uses the Graph Data Interface (GDI) to generate the index on the Elasticsearch server. Unlike queries for other steps, this step does not run an INSERT query because the data is not being inserted into AnzoGraph. For general information about writing GDI queries against Elasticsearch sources, see Querying an Elasticsearch Source.

#### Note

Do not include Elasticsearch connection or index parameters in the query. Anzo automatically populates that information from the AnzoGraph configuration when the step is run.



7. Click **Save** to save the step configuration.

Once the Details tab is configured and the query is written, the step can be run. For information about running this step conditionally by setting up an execution condition, see Defining Execution Conditions.

## Take a Snapshot of an Index (Elasticsearch Snapshot Step)

This topic provides guidance on configuring an Elasticsearch Snapshot Step that generates a snapshot of an Elasticsearch index that was created from a graphmart. Follow the steps below to create an Elasticsearch Snapshot Step.

#### Tip

For an overview on creating an Elasticsearch index and snapshot from a graphmart, see Creating an Elasticsearch Index from a Graphmart.

- 1. Go to the graphmart for which you want to add a step and then click the **Data Layers** tab.
- On the Data Layers tab, find the layer that you want to add the step to. Click the menu icon (I) for that layer and select Add Step/View. The Add Step/View dialog box is displayed with the New tab selected.

Add Step	o/View					
N	ew Existing Steps Existing Views					
<b>T</b> Q	Search					
	Direct Load Step					
	Load data directly from an external source					
	Elasticsearch Indexing Step					
<u>a</u>	Create a single Elasticsearch index to associate with a layer.					
-						
	Elasticsearch Snapshot Step					
	Create an Elasticsearch snapshot of the index associated with a layer.					
	Export Step					
	This allows you to export contents of graphmart to an flds.					
	Rows per page: 25 🛩 1-12 of 12 < >					
	CANCEL OK					

3. To create a new Elasticsearch Snapshot step, select **Elasticsearch Snapshot Step** and then click **OK**. If you want to clone an existing step, click the **Existing Steps** tab, select the step that you want to clone, and then click **OK**. Anzo creates or clones the step.

If this is the first Elasticsearch-related step in the layer, the following warning is displayed:

Warning	
Adding this step to this layer will enable this Layer's ES Indexing with a default configuration. You can change that configuration at the la	yer level.
	ок

The warning is to inform you that the Elasticsearch index-related configuration settings at the layer level are automatically adjusted to the default settings for layers that contain Elasticsearch Indexing or Snapshot Steps. Those settings are on the layer's Elasticsearch Indexing tab. The image below shows the tab with the default settings:

	Details	Execution Condition	Data Access	Query Context	Elasticsearch Indexing
🗹 Assoc	iate an Elasti	icsearch index with this laye	r		
Warning: The Advanced 🔺	se are advance	ed settings. Proceed with cautio	on. Changes may signi	ficantly affect Elasticsear	rch behavior.
Elasticsearch Name will	h Index Name   be system ç	generated if left blank			
Name will be	e system genera	ited if left blank			
🗹 Clea	ar Elasticsear	rch index before this layer ex	ecutes 🗹 Delete	Elasticsearch index on	layer unload
Elasticsea	arch Index Se	ettings			
Settings (JS	ON format) for I	Elasticsearch index			
Elasticsea	arch Index M	apping			
Mapping (JS	SON format) for	Elasticsearch index			

4. If necessary, click **OK** to close the warning. The Details tab is displayed:

Details	Execution Condition	Query Con	text
Title *			
The title of the step			
Description			
A brief description of the step			
Enabled			
Source			
Self ×			$\sim$
This option cannot be changed.			
Target FLDS *			~
FLDS only.			
		CANCEL	SAVE

- 5. On the Details tab, configure the following options as needed:
  - Title: The required name of the step.
  - **Description**: An optional short description of the step.

- **Enabled**: When creating a new step, the Enabled option is selected by default, indicating that the step is enabled and will run when the layer is loaded or refreshed. If you want to disable the step so that it is not processed, clear the **Enabled** checkbox.
- **Source**: The source data that this step should act upon. Steps can build upon the data generated by steps in other layers or can be self-contained, applying changes that relate only to the data defined in the layer that contains this step. You can select any number of the following options:
  - **Self**: This option is selected by default and means that the step runs against the data that is in the parent layer.
  - All Previous Layers Within Graphmart: This option means that the step runs against the data in all of the successful layers that precede the layer this step is in. Any failed layers are ignored.
  - **Previous Layer Within Graphmart**: This option means that the query runs against the data that is in the one layer that precedes the layer this step is in.
  - *Layer Name*: The Source drop-down list also includes options for layer names.
     You can choose a specific layer to act upon the data in that layer.
- Target FLDS: This is the target FLDS to save the snapshot to. Typically this is the FLDS that was created by an Export Step in a previous layer. If an FLDS does not exist, you can select -Create New- to create an empty dataset. See Adding an Empty Dataset for an Export Step for instructions.
- 6. Click **Save** to save the step configuration.

Once the Details tab is configured, the step can be run. For information about running this step conditionally by setting up an execution condition, see Defining Execution Conditions.

## Export Data to an FLDS (Export Step)

This topic provides guidance on configuring an Export Step to use for exporting the knowledge graphs in memory to a file-based linked data set (FLDS) on the shared file store. Follow the steps below to create an Export Step.

#### Important

If you add an Export Step to a graphmart that has been activated, you must reload the entire graphmart after adding the step. Simply refreshing the layer or graphmart after adding the step does not create the ontology graph that the Export Step requires.

- 1. Go to the graphmart for which you want to add a step and then click the **Data Layers** tab.
- 2. On the Data Layers tab, find the layer that you want to add the step to. Click the menu icon (I) for that layer and select Add Step/View. The Add Step/View dialog box is displayed with the New tab selected.

Add Step	o/View						
N	lew	Existing Steps	Existing Views				
<b>T</b> Q	Search	- *					
	Direct Load	Step					*
Ð	Load data direc	ctly from an external source					- 1
	Elasticsear	ch Indexing Step					
<u>م</u>	Create a single	Elasticsearch index to associat	e with a layer.				
	Elasticsear	ch Snapshot Step					
•	Create an Elast	ticsearch snapshot of the index	associated with a layer.				
	Export Step	•					
Ē.	This allows you	u to export contents of graphma	rt to an flds.				
			Rows per page: 25	<b>▼</b> 1	-12 of 12	<	>
					CAN	EL	ок

3. To create a new Export step, select **Export Step** and then click **OK**. If you want to clone an existing step, click the **Existing Steps** tab, select the step that you want to clone, and then click **OK**. Anzo creates or clones the step and displays the Details tab:

Details	Elastic Search	Execution Condition	Query Context	
Title *				
The title of the step				
Description				
A brief description of the step				
Source				
Self K				×   ~
Source data to act upon				
Data models				~
Associated data models				
				CANCEL SAV

- 4. On the Details tab, configure the following options as needed:
  - Title: The required name of the step.
  - Description: An optional short description of the step.
  - Enabled: When creating a new step, the Enabled option is selected by default, indicating that the step is enabled and will run when the layer is loaded or refreshed. If you want to disable the step so that it is not processed, clear the Enabled checkbox.
  - **Source**: The source data that this step should act upon. Steps can build upon the data generated by steps in other layers or can be self-contained, applying changes that relate only to the data defined in the layer that contains this step. You can select any number of the following options:
    - **Self**: This option is selected by default and means that the step runs against the data that is in the parent layer.
    - All Previous Layers Within Graphmart: This option means that the step runs against the data in all of the successful layers that precede the layer this step is in. Any failed layers are ignored.
    - **Previous Layer Within Graphmart**: This option means that the query runs against the data that is in the one layer that precedes the layer this step is in.
    - *Layer Name*: The Source drop-down list also includes options for layer names.
       You can choose a specific layer to act upon the data in that layer.

- Data models: This optional field specifies the model or models to associate with this step. The list displays all of the available models. By default, the field is set to Exclude System Data (?). If you want to choose a system model, click the toggle button on the right side of the field to change it to Include System Data (X). When system data is included, the drop-down list displays the system models in addition to the user-generated models.
- Target FLDS: This is the target FLDS for this export. If an FLDS does not exist, you can select -Create New- to create an empty dataset. See Adding an Empty Dataset for an Export Step for instructions. If you select an existing target FLDS, you also have the option to specify whether or not to overwrite the existing dataset.
- Overwrite FLDS: This setting controls whether the existing FLDS is replaced with the exported files each time the Export Step runs or whether the exported files are added to the existing FLDS:
  - If you want Anzo to replace the current edition of the dataset, select the Overwrite FLDS checkbox. When Overwrite FLDS is enabled, Anzo archives the existing files in a new timestamped export subdirectory directory under the Target FLDS directory. Each time the Export step runs, Anzo archives the current edition, and creates a new export directory. If you add this dataset to a graphmart, only the latest version of the exported data will be loaded to AnzoGraph.
  - If you want Anzo to add the exported files to the existing FLDS, leave the Overwrite FLDS checkbox unchecked. When Overwrite FLDS is disabled, Anzo adds all of the exported components to a cumulative export directory under the Target FLDS directory. The dataset will contain the original files as well as all cumulative working editions. If you subsequently add this dataset to a graphmart, all of the data from all of the subdirectories will be loaded into AnzoGraph.
- **Export Binary Store Contents**: This option applies to exports of unstructured graphmarts and controls whether the binary store is exported along with the data.
- Always Move Binary Store: This option also applies to exports of unstructured graphmarts and controls whether the binary store is moved or copied during the export.

Since the binary store can be large and have a nested structure, copying the data can take a very long time. Since moving the binary store is almost instantaneous, however, enabling **Always Move Binary Store** can reduce the time it takes to complete the export.

 Export Elasticsearch Contents: This option is enabled by default and controls whether Elasticsearch contents, such as associated indexes, are copied to disk in the Target FLDS. If you do not want Elasticsearch contents to be exported to the FLDS, clear the checkbox.

#### Note

When exporting data from an AnzoGraph instance that does not have an associated Elasticsearch connection, clear the **Export Elasticsearch Contents** checkbox.

• Keep Elasticsearch Index Online: This option controls whether the Elasticsearch index that is associated with the dataset remains stored in Elasticsearch or is removed from Elasticsearch once it is exported.

#### Tip

For information about advanced Elasticsearch options that are available for Export Steps, see Elastic Search Tab below.

- Generate Metrics: This option controls whether a data profile is generated before the data is exported. Enabling this option increases the time it takes to run the step, but enabling it ensures that the information on the Dataset Explore is complete if the dataset is viewed. If you load the exported files in the future, the data profile is also loaded.
- Do Not Create New Edition in Dataset on Export: This option controls whether or not a new edition is created for the dataset each time the Export Step is run. By default this option is disabled, which means each export results in a new edition. If you do not want a

new edition to be created on each export, select the **Do Not Create New Edition in Dataset on Export** checkbox.

- Maximum Number of Components in Edition: This option controls the maximum number of components to retain in an edition. The default value is blank, which means unlimited. If you specify a number in this field and the limit is reached, Anzo ages off the oldest components as new ones are created.
- 5. Click Save to save the step configuration.

Once the Details tab is configured, the step can be run. If you added the step to an active graphmart, make sure that you deactivate and then reactivate the graphmart. Simply refreshing the layer or graphmart after adding the step does not create the ontology graph that the Export Step requires. For information about running this step conditionally by setting up an execution condition, see Defining Execution Conditions.

Once the graphmart is reactivated, the new dataset becomes available in the Datasets catalog.

#### Note

Unless you enable **Generate Metrics**, when you view in the Datasets catalog a dataset that was created by an Export Step, the counts in the Ontology panel on the Explore tab remain **0** until a Data Profile is generated for the dataset. For information about generating a profile, see **Generating a Dataset Data Profile**. For example, the image below shows the Explore tab for a dataset that was created by an Export Step and has not been profiled.



### **Elastic Search Tab**

The **Elastic Search** tab contains optional settings that you can use to set any desired limits on Elasticsearch indexing processes.

Details	Elastic Search	Execution Condition	Query Context
Elasticsearch Index Setting	S		
Elasticsearch Index Settings			
Elasticsearch Bulk Size			
Elasticsearch Bulk Size (max batcl	h size in MB)		
Elasticsearch Bulk Actions			
Elasticsearch Bulk Actions (max n	umber of documents per ba	itch)	
Elasticsearch Bulk Concurr	ent Requests		
Elasticsearch Bulk Concurrent Reg	uests (max number of conc	current bulk requests)	
Elasticsearch Bulk Max Thr	eads Per FLDS		
Elasticsearch Bulk Max Threads P	er FLDS (max number threa	ds to use for indexing per FLDS	1)
Elasticsearch Bulk Max FLI	OS Threads		
Elasticsearch Bulk Max FLDS Thre	ads (max number FLDSes to	o index concurrently)	
			CANCEL SA

Setting	Description
Elasticsearch Index Settings	A custom JSON list of any Elasticsearch index settings that you want to apply to the export.
Elasticsearch Bulk Size	The maximum batch size in MB.
Elasticsearch Bulk Actions	The maximum number of documents to include in each batch.
Elasticsearch Bulk Concurrent Requests	The maximum number of bulk requests that can run concurrently.
Elasticsearch Bulk Max Threads Per FLDS	The maximum number of threads to use for indexing per file- backed linked data set (FLDS).
Elasticsearch Bulk Max FLDS Threads	The maximum number of FLDSes to index concurrently.

## Load a Dataset from the Catalog (Load Dataset Step)

This topic provides guidance on configuring a Load Dataset Step to use for adding a dataset from the Datasets catalog to a graphmart. Follow the steps below to create a Load Dataset Step.

- 1. Go to the graphmart for which you want to add a step and then click the **Data Layers** tab.
- On the Data Layers tab, find the layer that you want to add the step to. Click the menu icon (1) for that layer and select Add Step/View. The Add Step/View dialog box is displayed with the New tab selected.

	New Existing Steps Existing Views
<b>q</b>	Search 📃 🗐 💌
-	Direct Load Step
Ð	Load data directly from an external source
	Elasticsearch Indexing Step
Q	Create a single Elasticsearch index to associate with a layer.
	Elasticsearch Snapshot Step
0	Create an Elasticsearch snapshot of the index associated with a layer.
_	Export Step
	This allows you to export contents of graphmart to an fids.
	Rows per page: 25 🛩 1-12 of 12 <

3. Select Load Dataset Step and then click OK. Anzo creates the step and displays the Details tab:

Create					20
	Details	Execution Condition	Filter	Query Context	
					Chabled
Title *					
The title of t	the step				
Descripti	on				
A brief desc	ription of the step				
Linked (	Dataset *				o . ~
Unked Data	net to load				
U Wa	tch Flds Directory				
🗌 Ign	ore Missing File or	Directory			
🗌 Skij	p Elastic Search Sr	apshot Restoration if Inde	x Already Exists		
					CANCEL SAVE

- 4. On the Details tab, configure the following options as needed:
  - Title: The required name of the step.
  - Description: An optional short description of the step.
  - Enabled: When creating a new step, the Enabled option is selected by default, indicating that the step is enabled and will run when the layer is loaded or refreshed. If you want to disable the step so that it is not processed, slide the Enabled slider to the left.
  - Linked Dataset: This field specifies the dataset to load. The list displays all of the datasets in the Dataset catalog. By default, the field is set to Exclude System Data (\*). If you want to choose a system dataset, click the toggle button on the right side of the field to change it to Include System Data (\*). When you select a dataset, the current working edition (Managed Edition) of the dataset is selected as the data to load. If you want to change the edition, you can click Modify Edition and follow the steps in Modifying an Edition.
  - Watch FLDS Directory: This option controls whether the FLDS directory is monitored for changes. If Watch FLDS Directory is enabled and changes to the files in the FLDS directory are detected, Anzo will mark this step (and layer) as needing a refresh.
  - Ignore Missing File or Directory: This option controls whether to ignore missing files or subdirectories in the FLDS directory and proceed with the load or fail the step if files or directories are missing.

- Skip Elastic Search Snapshot Restoration if Index Already Exists: This option applies to graphmarts with Elasticsearch indexes and controls whether Anzo first checks to see if an index with the alias for the dataset already exists in Elasticsearch. If this setting is enabled and the index does exist, Anzo will not reload the index snaphsot into Elasticsearch.
- 5. Typically when users add a dataset to a graphmart, they want to load the entire dataset. However, if you are familiar with the data and want to exclude certain predicates from the dataset or write an INSERT query that filters the data, you can configure filtering options on the Filter tab. For information, see Filter Tab below.
- 6. Click **Save** to save the step configuration.

Once the Details tab is configured, the step can be run. For information about running this step conditionally by setting up an execution condition, see Defining Execution Conditions.

#### **Filter Tab**

The **Filter** tab includes options for filtering out some of the data in the dataset. If you want to load all of the statements in the dataset, do not configure Filter options. If you want to exclude some statements, configure the Filter options.

	Details	Execution Condition	Filter	Query Context	
Filter Ty	ype Itiple Select				
O Qu	ery				
Maskee	d Predicate				~

## **Multiple Select**

This option enables you to exclude certain triples from the load by selecting the predicates to filter out. These are known as **Masked Predicates**. To exclude predicates, select the **Multiple Select** radio button, then click the **Masked Predicate** drop-down list and select a predicate to add it to the Masked Predicate field. Click the field again to select additional predicates. You can remove a property from the masked list by clicking the X next to the predicate name.

## Query

If you want to hand-pick the data to load, you can use this option to write a SPARQL query that inserts specific values or filters out certain values. To write a query, select the **Query** radio button, and then type an INSERT query in the text box. For example, you can use the following format to filter out properties from the files:

```
INSERT {
   GRAPH ${targetGraph}{
      ?s ?p ?o.
   }
}
${usingSources}
WHERE {
   ?s ?p ?o .
   FILTER EXISTS { ?s a ?type . }
   FILTER(?type = <URI>)
}
```

#### Note

Including the **\${targetGraph}** and **\${usingSources}** parameters are required. Anzo automatically populates the query with the appropriate graph URIs when the step is run.

#### Important

In load filter queries, URIs are not supported in the object position. To specify a URI as an object, include the standard ?s ?p ?o triple pattern in the WHERE clause and then apply FILTER statements with URIs as needed. URIs are supported in the subject or predicate position.

For example, the following query filters the data in a sample dataset that includes information about people and the events they buy tickets for. The WHERE clause filters the data to load only the triples that are related to person1 (personid=1):

```
INSERT { GRAPH ${targetGraph} {
    ?s ?p ?o
  }
```

```
}
${usingSources}
WHERE {
   ?s ?p ?o;
   <http://cambridgesemantics.com/ont/autogen/c89d/Tickets#tickit_users_personid> ?id
   FILTER (?id=1)
}
```

## Pre-Compile a Query (Pre-Compile Query Step)

The first time a user runs an analytic query against AnzoGraph, AnzoGraph performs a code compilation process to generate the code for running that query. It then executes the query using that compiled code, and the same code is reused for subsequent runs of the query. If you determine that a particular query has a long code compilation time, you can add that query to a Pre-Compile Query Step. That way the query is run during the graphmart load and the compiled code is available before an end-user runs that query. Follow the steps below to create a Pre-Compile Query Step.

- 1. Go to the graphmart for which you want to add a step and then click the **Data Layers** tab.
- On the Data Layers tab, find the layer that you want to add the step to. Click the menu icon (I) for that layer and select Add Step/View. The Add Step/View dialog box is displayed with the New tab selected.

Add Step	o/View						
N	lew	Existing Steps	Existing Views				
<b>T</b> Q	Search	= *					
	Direct Load	Step					^
	Load data direct	tly from an external source					- 1
	Elasticsearc	h Indexing Step					- 11
) ସି	Create a single	Elasticsearch index to associa	te with a layer.				
	Elasticsearc	h Snapshot Step					
•	Create an Elasti	csearch snapshot of the index	associated with a layer.				
	Export Step						
E.	This allows you	to export contents of graphm	art to an flds.				
			Rows per page:	25 👻	1-12 of 12	<	>
					CAN	CEL	ОК

To create a new Pre-Compile Query Step, select Pre-Compile Query Step and then click
 OK. If you want to clone an existing step, click the Existing Steps tab, select the step that you want to clone, and then click OK. Anzo creates or clones the step and displays the Details tab:

Details	Query	Execution Condition	Query Context	
Title *				
The title of the step				
Description				
brief description of the step	compile query fails	the layer will be marked a	sfailed	
☐ If the precompile query	fails the whole or	and mart will be marked as	failed	
iource	rono, the more gr			
Self x			×	~
Source data to act upon				
Pre-Run Generate Statis	tics			
			CANCEL.	SAVE

- 4. On the Details tab, configure the following options as needed:
  - Title: The required name of the step.
  - Description: An optional short description of the step.
  - **Enabled**: When creating a new step, the Enabled option is selected by default, indicating that the step is enabled and will run when the layer is loaded or refreshed. If you want to disable the step so that it is not processed, clear the **Enabled** checkbox.
  - If the precompile query fails, the layer will be marked as failed: This option controls whether loading the data layer that contains this step is completed or aborted if this step fails. Select this option if you want Anzo to fail the layer if this step fails.
  - If the precompile query fails, the whole graphmart will be marked as failed: This option controls whether loading the graphmart is completed or aborted if this step fails.
     Select this option if you want Anzo to fail the entire graphmart if this step fails.
  - **Source**: The source data that this step should act upon. Steps can build upon the data generated by steps in other layers or can be self-contained, applying changes that relate only to the data defined in the layer that contains this step. You can select any number of the following options:

- **Self**: This option is selected by default and means that the step runs against the data that is in the parent layer.
- All Previous Layers Within Graphmart: This option means that the step runs against the data in all of the successful layers that precede the layer this step is in. Any failed layers are ignored.
- **Previous Layer Within Graphmart**: This option means that the query runs against the data that is in the one layer that precedes the layer this step is in.
- *Layer Name*: The Source drop-down list also includes options for layer names.
   You can choose a specific layer to act upon the data in that layer.
- **Pre-Run Generate Statistics**: This option controls whether to initiate AnzoGraph's internal statistics gathering queries before running the query to pre-compile. The statistics gathering helps ensure that the AnzoGraph query planner generates ideal query execution plans for queries that are run against the graphmart.
- 5. When you have finished configuring the Details tab, click the **Query** tab. This tab contains the query to pre-compile. The tab provides a template for writing a SPARQL SELECT query. You can edit the template to write the query, or you can paste in query contents that you copied from a log file, dashboard, the Query Builder, etc.

#### Note

Make sure that you include the \${fromSources} parameter in the query. Anzo automatically populates the query with the appropriate source graph URIs according to the Source configured from the Details tab.

	Details	Query	Execution Condition	Query Context		
Precon	ple query *					
1						
2	#fromSources is repla	ced with the URIs	of the Layer's Sources at	t runtime		
3	SELECT *					
4	\${fromSources}					
5.	WHERE{					
6	1					
Query	used to pre-compile					
					CANCEL	SAVE

## Note

If your query connects to a source that requires input of connection and authorization information, Cambridge Semantics recommends that you do not include the connection and authorization values directly in the query. Instead, replace those values with Context Variables from a Query Context. You can access Context Providers for each data source from the step's Query Context tab. For detailed information about query contexts and referencing variables in a query, see Using Query Contexts.

6. Click **Save** to save the step configuration.

Once the Details tab is configured and the query is written, the step can be run. For information about running this step conditionally by setting up an execution condition, see Defining Execution Conditions.

## Create a Reusable Query Template

There are two types of steps that enable you to create a query template that is reusable across data layers and graphmarts:

## **Templated Step**

Templated Steps use user-defined key-value pairs. The keys are represented by parameters in the query. Creating the key-value pairs requires familiarity with the data and properties defined in the model. When this step is reused, users do not need to rewrite the query; they modify the values for

the keys.

### **Query-Driven Templated Step**

Query-Driven Templated Steps are similar to Templated Steps in that they provide a way to create query templates that use parameters to represent key-value pairs. The queries are reusable across datasets because the existing parameters can be substituted for alternate key-value pairs. The difference between the two types of steps is that the key-value pairs for Templated Steps are user-defined. In Query-Driven Templated Steps, a parameter query is run that automatically generates the key-value pairs. Then the defined template query is run for each key-value solution from the parameter query.

This section includes instructions for creating both types of templated steps.

- Create a Templated Step
- Create a Query-Driven Template

### **Create a Templated Step**

Templated Steps use user-defined key-value pairs. The keys are represented by parameters in the query. When this step is reused, users do not need to rewrite the query; they modify the values for the keys. Follow the steps below to create a Templated Step.

## Tip

This type of template step uses key-value pairs that are user-defined. Creating the key-value pairs requires familiarity with the data and properties defined in the model. To create a query template that enables you to run a query and automatically generate the key-value pairs, see Create a Query-Driven Template.

- 1. Go to the graphmart for which you want to add a step and then click the **Data Layers** tab.
- On the Data Layers tab, find the layer that you want to add the step to. Click the menu icon (1) for that layer and select Add Step/View. The Add Step/View dialog box is displayed with the New tab selected.

New Existing Steps Existing Views	
▼ Q Search = ▼	
Direct Load Step	A
Load data directly from an external source	
Elasticsearch Indexing Step	
Create a single Elasticsearch index to associate with a layer.	
Elasticsearch Snapshot Step	
Create an Elasticsearch snapshot of the index associated with a layer.	
-	
Export Step	
This allows you to export contents of graphmart to an fids.	
-	
Rows per page: 25 ¥ 1-12 of 1	2 < >
c	CANCEL OK

3. To create a new Templated step, select **Templated Step** and then click **OK**. If you want to clone an existing step, click the **Existing Steps** tab, select the step that you want to clone, and then click **OK**. Anzo creates or clones the step and displays the Details tab:

Details	Template	Execution Condition	Query Context
Title *			
The title of the step			
Description			
A brief description of the step			
Enabled			
Source			
Self ×			×   ~
Source data to act upon			
Data models *			0 V
Associated data models			
			CANCEL SAV
			CHARGE CAN

- 4. On the Details tab, configure the following options as needed:
  - Title: The required name of the step.
  - **Description**: An optional short description of the step.
- **Enabled**: When creating a new step, the Enabled option is selected by default, indicating that the step is enabled and will run when the layer is loaded or refreshed. If you want to disable the step so that it is not processed, clear the **Enabled** checkbox.
- **Source**: The source data that this step should act upon. Steps can build upon the data generated by steps in other layers or can be self-contained, applying changes that relate only to the data defined in the layer that contains this step. You can select any number of the following options:
  - **Self**: This option is selected by default and means that the step runs against the data that is in the parent layer.
  - All Previous Layers Within Graphmart: This option means that the step runs against the data in all of the successful layers that precede the layer this step is in. Any failed layers are ignored.
  - **Previous Layer Within Graphmart**: This option means that the query runs against the data that is in the one layer that precedes the layer this step is in.
  - *Layer Name*: The Source drop-down list also includes options for layer names.
     You can choose a specific layer to act upon the data in that layer.
- Data models: This required field specifies the model or models to associate with this step. The list displays all of the available models. By default, the field is set to Exclude System Data (\*). If you want to choose a system model, click the toggle button on the right side of the field to change it to Include System Data (\*). When system data is included, the drop-down list displays the system models in addition to the user-generated models.
- **Pre-Run Generate Statistics**: This option controls whether to initiate AnzoGraph's internal statistics gathering queries before running the query to pre-compile. The statistics gathering helps ensure that the AnzoGraph query planner generates ideal query execution plans for queries that are run against the graphmart.
- 5. When you have finished configuring the Details tab, click the **Template** tab. This tab defines the template query and the key-value pairs.



6. On the left side of the screen create the query template. The default template includes the syntax for writing SPARQL INSERT and DELETE queries and includes source and target graph parameters that Anzo replaces at runtime. In the query, include the parameters in the format \${key\_name}. Each parameter will become a key in the Key/Value Table will you click the **Refresh Keys** button. For example, the following INSERT query includes several parameters that represent properties and functions:

```
INSERT {
   GRAPH ${targetGraph}{
        ?lsubj ${linkProperty} ?rsubj
   }
}
${usingSources}
WHERE {
        ?lsubj ${sourceProperty} ?lobj .
        ?rsubj ${targetProperty} ?robj .
        FILTER (${lFunction}(?lobj) ${operator} ${rFunction}(?robj))
}
```

#### Note

If your query connects to a source that requires input of connection and authorization information, Cambridge Semantics recommends that you do not include the connection and authorization values directly in the query. Instead, replace those values with Context Variables from a Query Context. You can access Context Providers for each data source from the step's Query Context tab. For detailed information about query contexts and referencing variables in a query, see Using Query Contexts.

7. Once the template query has been defined, populate the Key/Value Table with the keys from the query by clicking the **Refresh Keys** button. For example, using the example query above, the Key/Value table is populated with the following keys:

,		
Search keys Q	Search values	0
linkProperty	Value	8
Function	Value	î
targetProperty	Value	
rFunction	Value	8
operator	Value	8
sourceProperty	Value	î

8. In each row, specify the desired **Value** for the key. For example, in the image below, the property URI

http://cambridgesemantics.com/ont/autogen/c89d/Tickets#tickit

events\_eventid is specified as the Value for the linkProperty key.

Search keys (	Q	Search values C	λ	_
IFunction		Value		8
rFunction		Value		8
linkProperty		http://cambridgesemantics.com/ont/autogen/c89d/Tickets#tickit_events_eventi	id	î
operator		Value		a,

9. Click **Save** to save the step configuration.

Once the Details tab is configured and the template and key-value pairs are defined, the step can be run. For information about running this step conditionally by setting up an execution condition, see Defining Execution Conditions.

### **Create a Query-Driven Template**

Query-Driven Templated Steps are similar to Templated Steps in that they provide a way to create query templates that use parameters to represent key-value pairs. The queries are reusable across datasets because the existing parameters can be substituted for alternate key-value pairs. The difference between the two types of steps is that the key-value pairs for Templated Steps are user-defined. In Query-Driven Templated Steps, a parameter query is run that automatically generates the key-value pairs. Then the defined template query is run for each key-value solution from the parameter query. Follow the steps below to create a Query-Driven Templated Step.

- 1. Go to the graphmart for which you want to add a step and then click the **Data Layers** tab.
- On the Data Layers tab, find the layer that you want to add the step to. Click the menu icon (I) for that layer and select Add Step/View. The Add Step/View dialog box is displayed with the New tab selected.

		Estados Otros	Foliation Minute				
	lew	Existing Steps	Existing Views				
٩	Search	= •					
-	Direct Load	d Step					
Ð	Load data dire	ctly from an external source					
_	Elasticsear	ch Indexing Step					
<u>D</u>	Create a single	e Elasticsearch index to asso	ciate with a layer.				
_	Elasticsear	ch Snapshot Step					
Ο	Create an Elas	sticsearch snapshot of the in	dex associated with a layer.				
	Export Step	p					
	This allows yo	u to export contents of grap	hmart to an flds.				
			Rows per page: 2	25 👻	1-12 of 12	<	>

3. To create a new Query-Driven Templated step, select Query Driven Templated Step and then click OK. If you want to clone an existing step, click the Existing Steps tab, select the step that you want to clone, and then click OK. Anzo creates or clones the step and displays the Details tab:

Details	Parameters Query	Template	Execution Condition	Query Context	
Title *					
The title of the step					
Description					
brief description of the step					
lource					
Self ×					x   ~
source data to act upon					
Data models *					0.~
Associated data models					
				CAN	CEL SAVE

- 4. On the Details tab, configure the following options as needed:
  - Title: The required name of the step.
  - Description: An optional short description of the step.
  - **Enabled**: When creating a new step, the Enabled option is selected by default, indicating that the step is enabled and will run when the layer is loaded or refreshed. If you want to disable the step so that it is not processed, clear the **Enabled** checkbox.
  - **Source**: The source data that this step should act upon. Steps can build upon the data generated by steps in other layers or can be self-contained, applying changes that relate only to the data defined in the layer that contains this step. You can select any number of the following options:
    - **Self**: This option is selected by default and means that the step runs against the data that is in the parent layer.
    - All Previous Layers Within Graphmart: This option means that the step runs against the data in all of the successful layers that precede the layer this step is in. Any failed layers are ignored.
    - **Previous Layer Within Graphmart**: This option means that the query runs against the data that is in the one layer that precedes the layer this step is in.
    - *Layer Name*: The Source drop-down list also includes options for layer names.
       You can choose a specific layer to act upon the data in that layer.

- Data models: This required field specifies the model or models to associate with this step. The list displays all of the available models. By default, the field is set to Exclude System Data (?). If you want to choose a system model, click the toggle button on the right side of the field to change it to Include System Data (X). When system data is included, the drop-down list displays the system models in addition to the user-generated models.
- **Pre-Run Generate Statistics**: This option controls whether to initiate AnzoGraph's internal statistics gathering queries before running the query to pre-compile. The statistics gathering helps ensure that the AnzoGraph query planner generates ideal query execution plans for queries that are run against the graphmart.
- 5. When you have finished configuring the Details tab, click the **Parameters Query** tab. This tab defines the query to use for determining the key-value pairs for the Source that was selected on the Details tab. The tab provides a template for writing a SPARQL SELECT DISTINCT query. Edit the template as needed.

#### Note

Make sure that you include the \${fromSources} parameter in the query. Anzo automatically populates the query with the appropriate source graph URIs according to the Source configured on the Details tab.

	Details	Parameters Query	Template	Execution Condition	Query Context	
Templat	e Parameter *					
1	#fromSources is n	eplaced with the URIs of t	he Layer's Sourc	es at runtime		
2	#For each result	row, the template query is	run with the se	lected variables as the	input bindings	
3	SELECT DISTINCT ?	param1 ?param2 ?param3				
4	\$(fromSources)					
5.*	WHERE{					
6	?param1 ?pa	ran2 ?paran3.				
7	3					
					CANCEL	SAVE

### Note

If your query connects to a source that requires input of connection and authorization information, Cambridge Semantics recommends that you do not include the connection and authorization values directly in the query. Instead, replace those values with Context Variables from a Query Context. You can access Context Providers for each data source from the step's Query Context tab. For detailed information about query contexts and referencing variables in a query, see Using Query Contexts.

6. When you have finished configuring the Parameters Query tab, click the **Template** tab. This tab contains the query to run for each of the key-value pairs identified by the Parameters Query. The template includes the syntax for writing SPARQL DELETE and INSERT queries and includes source and target graph parameters as well as the placeholder parameters from the Parameters Query.

#### Note

By default, Anzo uses RDF encoding for parameters, meaning a parameter specified as  $f{param}$  is translated as  $f{rdf.param}$ . If you do not want to use RDF encoding, you can specify plain text by adding text. before the parameter name, for example,  $f{text.param}$ .

	Details	Parameters Query	Template	Execution Condition	Query Context	
Transfo	mation Template *					
1	#tangetGraph is re	eplaced with the Layers UK	I at runtime			
2	#usingSources is a	replaced with the URIs of	the Layen's Sources	at runtime		
3	#Query is run once	e per result from the para	meters query			
4 *	DELETE{					
5 *	GRAPH \$(target	tGraph)-{				
6						
7	}					
8	)					
9 *	INSERT(					
10 *	GRAPH \${target	tGraph)-{				
11						
12	)					
13	}					
14	\${usingSources}					
15 *	WHERE (					
16	\${paran1} \$	<pre>(param2) \${param3} .</pre>				
17	}					
Query to	emplate used to perform	the transformation				
					CANCEL SA	VE

7. Click Save to save the step configuration.

Anzo replaces the parameters at runtime. The query is executed *n* times, where *n* is the number of rows returned by the Parameters Query.

Once the Details tab is configured, the parameters query is in place, and the template is completed, the step can be run. Anzo replaces the parameters at runtime, and the query is executed *n* times, where *n* is the number of rows returned by the parameters query. For information about running this step conditionally by setting up an execution condition, see Defining Execution Conditions.

# Run a Transformation Query (Query Step)

This topic provides guidance on configuring a Query Step that you can use for creating, cleaning, conforming, or transforming the data in a layer. Follow the steps below to create a Query Step.

- 1. Go to the graphmart for which you want to add a step and then click the **Data Layers** tab.
- 2. On the Data Layers tab, find the layer that you want to add the step to. Click the menu icon (I) for that layer and select Add Step/View. The Add Step/View dialog box is displayed with the New tab selected.

Add Ste	p/View						
N	lew	Existing Steps	Existing Views				
<b>T</b> Q	Search	- *					
	Direct Load	Step					^
	Load data direc	ctly from an external source					- 11
	Elasticseard	ch Indexing Step					- 11
<u></u>	Create a single	Elasticsearch index to assoc	iate with a layer.				
	Flasticsear	ch Snanshot Sten					
6	Create an Elast	ticsearch snapshot of the inde	ex associated with a layer.				
•							
	Export Step						
E.	This allows you	u to export contents of graphr	nart to an fids.				
			Rows per page:	25 👻	1-12 of 12	<	>
					CANO	EL	ок

3. To create a new Query step, select **Query Step** and then click **OK**. If you want to clone an existing step, click the **Existing Steps** tab, select the step that you want to clone, and then click **OK**. Anzo creates or clones the step and displays the Details tab:

Details	Query	Execution Condition	Query Context	
Title *				
the title of the step				
Description				
brief description of the step				
outce				
Self ×				x   ~
Source data to act upon				
Data models *				0.4
Associated data models				

- 4. On the Details tab, configure the following options as needed:
  - Title: The required name of the step.
  - Description: An optional short description of the step.
  - **Enabled**: When creating a new step, the Enabled option is selected by default, indicating that the step is enabled and will run when the layer is loaded or refreshed. If you want to disable the step so that it is not processed, clear the **Enabled** checkbox.
  - **Source**: The source data that this step should act upon. Steps can build upon the data generated by steps in other layers or can be self-contained, applying changes that relate only to the data defined in the layer that contains this step. You can select any number of the following options:
    - **Self**: This option is selected by default and means that the step runs against the data that is in the parent layer.
    - All Previous Layers Within Graphmart: This option means that the step runs against the data in all of the successful layers that precede the layer this step is in. Any failed layers are ignored.
    - **Previous Layer Within Graphmart**: This option means that the query runs against the data that is in the one layer that precedes the layer this step is in.
    - *Layer Name*: The Source drop-down list also includes options for layer names.
       You can choose a specific layer to act upon the data in that layer.

- Data models: This required field specifies the model or models to associate with this step. The list displays all of the available models. By default, the field is set to Exclude System Data (?). If you want to choose a system model, click the toggle button on the right side of the field to change it to Include System Data (X). When system data is included, the drop-down list displays the system models in addition to the user-generated models.
- **Pre-Run Generate Statistics**: This option controls whether to initiate AnzoGraph's internal statistics gathering queries before running the query to pre-compile. The statistics gathering helps ensure that the AnzoGraph query planner generates ideal query execution plans for queries that are run against the graphmart.
- 5. When you have finished configuring the Details tab, click the Query tab. This tab defines the query that this step should run. The template includes the syntax for writing SPARQL INSERT and DELETE queries and includes the target and source graph parameters (\${targetGraph} and \${usingSources}). Anzo replaces the parameters with the appropriate URIs when the step runs. Edit the template as needed. You can click the Preview in Query Builder button to open the query in the Query Builder, where you can perform practice runs to see results without having to refresh the graphmart or layer.

#### Note

If your query connects to a source that requires input of connection and authorization information, Cambridge Semantics recommends that you do not include the connection and authorization values directly in the query. Instead, replace those values with Context Variables from a Query Context. You can access Context Providers for each data source from the step's Query Context tab. For detailed information about query contexts and referencing variables in a query, see Using Query Contexts.



6. Click **Save** to save the step configuration.

Once the Details tab is configured and the query is defined, the step can be run. For information about running this step conditionally by setting up an execution condition, see Defining Execution Conditions.

## Infer New Data (RDFS+ Inference Step)

This topic provides guidance on configuring an RDFS+ Inference Step that uses RDFS+ and OWL rules to create new relationships based on the vocabularies in the existing data. Follow the steps below to create an RDFS+ Inference Step.

- 1. Go to the graphmart for which you want to add a step and then click the **Data Layers** tab.
- 2. On the Data Layers tab, find the layer that you want to add the step to. Click the menu icon (I) for that layer and select Add Step/View. The Add Step/View dialog box is displayed with the New tab selected.

Add Step	o/View						
N	lew	Existing Steps	Existing Views				
<b>T</b> Q	Search	=	*				
	Direct Load S	tep					
	Load data direct	y from an external sou	irce				- 11
	Elasticsearch	Indexing Step					
ā	Create a single E	lasticsearch index to a	associate with a layer.				
-							
	Elasticsearch	Snapshot Step					
6	Create an Elastic	search snapshot of th	e index associated with a layer.				
	Export Step						
	This allows you t	o export contents of g	raphmart to an flds.				
			Rows per page	25 👻	1-12 of 12	<	>
					CAN	EL	ок

3. To create a new RDFS+ Inference step, select **RDFS+ Inference Step** and then click **OK**. If you want to clone an existing step, click the **Existing Steps** tab, select the step that you want to clone, and then click **OK**. Anzo creates or clones the step and displays the Details tab:

	Details	Execution Condition	Query Context
Title *			
The title of the ste	ep		
Description			
A brief description	n of the step		
Chabled			
Inference Rule	s To Run		
Comma separateo exclude a rule	d rules: 'all'- for all	set of rules, 'rdfsplus' - for RDF	$FS*$ set of rules, {rulenames} - from owl2rl specific spec rules, use prefix $\forall$ to
Source			
Self $\mathbf{x}$			x   ~
			CANCEL SAVE

- 4. On the Details tab, configure the following options as needed:
  - Title: The required name of the step.
  - Description: An optional short description of the step.

- **Enabled**: When creating a new step, the Enabled option is selected by default, indicating that the step is enabled and will run when the layer is loaded or refreshed. If you want to disable the step so that it is not processed, clear the **Enabled** checkbox.
- Inference Rules to Run: By default the step runs all of the RDFS-plus inference rules and a subset of the OWL 2 RL rules (see Inference Rule Reference below for specifics). If you want to customize the step to include or exclude certain rules, specify any combination of the following options in the Inference Rules To Run field. Specify multiple options in a comma-separated list:
  - ° all: Run all rules.
  - ° **rdfsplus**: Run only the RDFS-plus rules.
  - *rule\_names*: List specific rules to run only those rules. See Inference Rule Reference.
  - -rule\_name: Specify a hyphen (-) in front of a rule name to exclude that rule. For example, -scm-svf2 excludes the scm-svf2 rule.

For example, the following value runs all of the inference rules except prp-fp and prp-ifp:

all,-prp-fp,-prp-ifp

#### Note

Certain inference rules are coupled. Specifying either of the rules in the pair automatically runs the coupled rule. The list below describes the paired rules:

- scm-dom1 and scm-rng1
- scm-dom2 and scm-rng2
- prp-inv1 and prp-inv2

In addition, running scm-eqc1and cax-sco also runs cax-eqc1 and cax-eqc2. And running scm-eqp1 and prp-spo1 also runs prp-eqp1 and prp-eqp2.

- **Source**: The source data that this step should act upon. Steps can build upon the data generated by steps in other layers or can be self-contained, applying changes that relate only to the data defined in the layer that contains this step. You can select any number of the following options:
  - **Self**: This option is selected by default and means that the step runs against the data that is in the parent layer.
  - All Previous Layers Within Graphmart: This option means that the step runs against the data in all of the successful layers that precede the layer this step is in. Any failed layers are ignored.
  - **Previous Layer Within Graphmart**: This option means that the query runs against the data that is in the one layer that precedes the layer this step is in.
  - *Layer Name*: The Source drop-down list also includes options for layer names.
     You can choose a specific layer to act upon the data in that layer.
- Data models: This required field specifies the model or models to associate with this step. The list displays all of the available models. By default, the field is set to Exclude System Data (\*). If you want to choose a system model, click the toggle button on the right side of the field to change it to Include System Data (\*). When system data is included, the drop-down list displays the system models in addition to the user-generated models.
- **Pre-Run Generate Statistics**: This option controls whether to initiate AnzoGraph's internal statistics gathering queries before running this step. The statistics gathering helps ensure that the AnzoGraph query planner generates ideal query execution plans for queries that are run against the graphmart.
- 5. Click **Save** to save the step configuration.

Once the Details tab is configured, the step can be run. For information about running this step conditionally by setting up an execution condition, see Defining Execution Conditions.

#### **Inference Rule Reference**

This section provides reference information for the RDFS-plus rules and the subset of OWL 2 RL rules that inference steps run.

- RDFS+ Rules
- OWL 2 RL Rules

#### **RDFS+** Rules

The tables below define the RDFS-plus inference rules.

- Semantics of Class Axioms
- Semantics of Axioms about Properties
- Semantics of Schema Vocabulary

#### **Semantics of Class Axioms**

#### Note

Because **cax-eqc1** and **cax-eqc2** (described in the table below) are implied rules that are coupled with scm-eqc1 and cax-sco, including **cax-eqc1** or **cax-eqc2** in the Inference Rules to Run field will result in an invalid inference rule name error. To run the cax-eqc1 and cax-eqc2 rules, specify **scm-eqc1** and **cax-sco** (scm-eqc1,cax-sco).

Rule	Description	IF	THEN
cax- eqc1	Two classes are synonymous.	T(?c1, owl:equivalentClass, ?c2) T(?x, rdf:type, ?c1)	T(?x, rdf:type, ?c2)
cax- eqc2	Two classes are synonymous.	T(?c1, owl:equivalentClass, ?c2) T(?x, rdf:type, ?c2)	T(?x, rdf:type, ?c1)

Rule	Description	IF	THEN
cax-sco	Members of a subclass are also members of the superclass.	T(?c1, rdfs:subClassOf, ?c2) T(?x, rdf:type, ?c1)	T(?x, rdf:type, ?c2)

#### **Semantics of Axioms about Properties**

#### Note

Because **prp-eqp1** and **prp-eqp2** (described in the table below) are implied rules that are coupled with scm-eqp1 and prp-spo1, including **prp-eqp1** or **prp-eqp2** in the Inference Rules to Run field will result in an invalid inference rule name error. To run the prp-eqp1 and prp-eqp2 rules, specify **scm-eqp1** and **prp-spo1** (scm-eqp1,prp-spo1).

Rule	Description	IF	THEN
prp-dom	Infer the subject's type from the predicate's domain.	T(?p, rdfs:domain, ?c) T(?x, ?p, ?y)	T(?x, rdf:type, ?c)
prp- eqp1	Two properties are synonymous.	T(?p1, owl:equivalentProperty, ?p2) T(?x, ?p1, ?y)	T(?x, ?p2, ?y)
prp- eqp2	Two properties are synonymous.	T(?p1, owl:equivalentProperty, ?p2) T(?x, ?p2, ?y)	T(?x, ?p1, ?y)
prp-fp	If predicate p is a functional property, then a subject can be related to only one specific object by p.	T(?p, rdf:type, owl:FunctionalProperty) T(?x, ?p, ?y1) T(?x, ?p, ?y2)	T(?y1, owl:sameAs, ?y2)
prp-ifp	If predicate p is an inverse	T(?p, rdf:type,	T(?x1,

Rule	Description	IF	THEN
	functional property, then a specific object can be related to only one subject by p.	owl:InverseFunctionalProperty) T(?x1, ?p, ?y) T(?x2, ?p, ?y)	owl:sameAs, ?x2)
prp-inv1	Two properties are the inverse of each other.	T(?p1, owl:inverseOf, ?p2) T(?x, ?p1, ?y)	T(?y, ?p2, ?x)
prp-inv2	Two properties are the inverse of each other.	T(?p1, owl:inverseOf, ?p2) T(?x, ?p2, ?y)	T(?y, ?p1, ?x)
prp-rng	Infer the object's type from the predicate's range.	T(?p, rdfs:range, ?c) T(?x, ?p, ?y)	T(?y, rdf:type, ?c)
prp- spo1	Relationships that are described by a subproperty also hold for the superproperty.	T(?p1, rdfs:subPropertyOf, ?p2) T(?x, ?p1, ?y)	T(?x, ?p2, ?y)
prp- symp	The inverse is true for a property.	T(?p, rdf:type, owl:SymmetricProperty) T(?x, ?p, ?y)	T(?y, ?p, ?x)
prp-trp	Chains of relationships collapse into a single relationship.	T(?p, rdf:type, owl:TransitiveProperty) T(?x, ?p, ?y) T(?y, ?p, ?z)	T(?x, ?p, ?z)

# Semantics of Schema Vocabulary

Rule	Description	IF	THEN
scm-cls	Every class is its own	T(?c, rdf:type, owl:Class)	T(?c, rdfs:subClassOf, ?c)

Rule	Description	IF	THEN
	subclass and equivalent class, and it is a subclass of owl:Thing.		T(?c, owl:equivalentClass, ?c) T(?c, rdfs:subClassOf, owl:Thing) T(owl:Nothing, rdfs:subClassOf, ?c)
scm- dom1	A property with domain c also has domain c's superclasses.	T(?p, rdfs:domain, ?c1) T(?c1, rdfs:subClassOf, ?c2)	T(?p, rdfs:domain, ?c2)
scm- dom2	A subproperty inherits the domains of the superproperties.	T(?p2, rdfs:domain, ?c) T(?p1, rdfs:subPropertyOf, ?p2)	T(?p1, rdfs:domain, ?c)
scm- eqc1	Equivalent classes are subclasses of each other.	T(?c1, owl:equivalentClass, ?c2)	T(?c1, rdfs:subClassOf, ?c2) T(?c2, rdfs:subClassOf, ?c1)
scm- eqc2	If two classes are subclasses, they are also equivalent classes.	T(?c1, rdfs:subClassOf, ?c2) T(?c2, rdfs:subClassOf, ?c1)	T(?c1, owl:equivalentClass, ?c2)
scm- eqp1	Equivalent properties are subproperties of each other.	T(?p1, owl:equivalentProperty, ?p2)	T(?p1, rdfs:subPropertyOf, ?p2) T(?p2, rdfs:subPropertyOf, ?p1)
scm- eqp2	If two properties are subproperties, they are	T(?p1, rdfs:subPropertyOf, ?p2)	T(?p1, owl:equivalentProperty,

Rule	Description	IF	THEN
	also equivalent properties.	T(?p2, rdfs:subPropertyOf, ?p1)	?p2)
scm- rng1	A property with range c also has range c's superclasses.	T(?p, rdfs:range, ?c1) T(?c1, rdfs:subClassOf, ?c2)	T(?p, rdfs:range, ?c2)
scm- rng2	A subproperty inherits the ranges of its superproperties.	T(?p2, rdfs:range, ?c) T(?p1, rdfs:subPropertyOf, ?p2)	T(?p1, rdfs:range, ?c)
scm-sco	owl:subClassOf relationships are transitive	T(?c1, rdfs:subClassOf, ?c2) T(?c2, rdfs:subClassOf, ?c3)	T(?c1, rdfs:subClassOf, ?c3)
scm-spo	owl:subPropertyOf relationships are transitive.	T(?p1, rdfs:subPropertyOf, ?p2) T(?p2, rdfs:subPropertyOf, ?p3)	T(?p1, rdfs:subPropertyOf, ?p3)

### Note

The scm-dp and scm-op schema vocabulary rules are not run. Those rules add significant compute overhead but do not result in meaningful inference results.

### OWL 2 RL Rules

The tables below define the subset of OWL 2 RL inference rules that inference steps run.

- Semantics of Equality
- Semantics of Schema Vocabulary
- Semantics of Classes

# Semantics of Equality

Rule	Description	IF	THEN
eq-rep- o	Describes the replacement property of the owl:sameAs axiom.	T(?o, owl:sameAs, ?o') T(?s, ?p, ?o)	T(?s, ?p, ?o')
eq-rep- p	Describes the replacement property of the owl:sameAs axiom.	T(?p, owl:sameAs, ?p') T(?s, ?p, ?o)	T(?s, ?p', ?o)
eq-rep- s	Describes the replacement property of the owl:sameAs axiom.	T(?s, owl:sameAs, ?s') T(?s, ?p, ?o)	T(?s', ?p, ?o)
eq-sym	Describes the symmetric property of the owl:sameAs axiom.	T(?x, owl:sameAs, ?y)	T(?y, owl:sameAs, ?x)
eq- trans	Describes the transitive property of the owl:sameAs axiom.	T(?x, owl:sameAs, ?y) T(?y, owl:sameAs, ?z)	T(?x, owl:sameAs, ?z)

# Semantics of Schema Vocabulary

Rule	Description	IF	THEN
scm- svf1	A property restriction c1 is a subclass of c2 if they are both someValuesFrom restrictions on the same property and c1's target class is a subclass of c2's target class.	T(?c1, owl:someValuesFrom, ?y1) T(?c1, owl:onProperty, ?p) T(?c2,	T(?c1, rdfs:subClassOf, ?c2)

Rule	Description	IF	THEN
		owl:someValuesFrom, ?y2) T(?c2, owl:onProperty, ?p) T(?y1, rdfs:subClassOf, ?y2)	
scm- svf2	A property restriction c1 is a subclass of c2 if they are both someValuesFrom restrictions on the same class where c1's target property is a subproperty of c2's target property.	T(?c1, owl:someValuesFrom, ?y) T(?c1, owl:onProperty, ?p1) T(?c2, owl:someValuesFrom, ?y) T(?c2, owl:onProperty, ?p2) T(?p1, rdfs:subPropertyOf, ?p2)	T(?c1, rdfs:subClassOf, ?c2)
scm-int		T(?c, owl:intersectionOf, ?x) LIST[?x, ?c1,, ?cn]	T(?c, rdfs:subClassOf, ?c1) T(?c, rdfs:subClassOf, ?c2)  T(?c, rdfs:subClassOf, ?cn)

#### **Semantics of Classes**

Rule	Description	IF	THEN
cls-svf1	At least one object of a property is a member of the specified class.	T(?x, owl:someValuesFrom, ?y) T(?x, owl:onProperty, ?p) T(?u, ?p, ?v) T(?v, rdf:type, ?y)	T(?u, rdf:type, ?x)
cls-int1	An instance belongs to every one of the specified classes.	T(?c, owl:intersectionOf, ?x) LIST[?x, ?c1,, ?cn] T(?y, rdf:type, ?c1) T(?y, rdf:type, ?c2)  T(?y, rdf:type, ?cn)	T(?y, rdf:type, ?c)

# Validate the Data (Validation Step)

This topic provides guidance on configuring a Validation Step to use for validating the data in a layer and optionally setting up execution conditions. Follow the steps below to create a Validation Step.

- 1. Go to the graphmart for which you want to add a step and then click the **Data Layers** tab.
- 2. On the Data Layers tab, find the layer that you want to add the step to. Click the menu icon (I) for that layer and select Add Step/View. The Add Step/View dialog box is displayed with the New tab selected.

Add Step	p/View
N	ew Existing Steps Existing Views
<b>T</b> Q	Search 📃 🖛 👻
	Direct Load Step
Ð	Load data directly from an external source
	Elasticsearch Indexing Step
ā	Create a single Elasticsearch index to associate with a layer.
	Elasticsearch Snapshot Step
•	Create an Elasticsearch snapshot of the index associated with a layer.
	Export Step
	This allows you to export contents of graphmart to an flds.
	Rows per page: 25 👻 1-12 of 12 < >
	CANCEL OK

3. To create a new Validation step, select **Validation Step** and then click **OK**. If you want to clone an existing step, click the **Existing Steps** tab, select the step that you want to clone, and then click **OK**. Anzo creates or clones the step and displays the Details tab:

Details	Options	Query	Query Context	
Title *				
The title of the step				
Description				
brief description of the step				
Validation Source Datasour	ce			~
ptional Datasource used to perform	m the validation query			
ource				
Self ×				×   ~
ource data to act upon				
				CANCEL SAVE

- 4. On the Details tab, configure the following options as needed:
  - Title: The required name of the step.
  - **Description**: An optional short description of the step.

- **Enabled**: When creating a new step, the Enabled option is selected by default, indicating that the step is enabled and will run when the layer is loaded or refreshed. If you want to disable the step so that it is not processed, clear the **Enabled** checkbox.
- Validation Source Datasource: This optional field enables you to select a data source (such as a system data source) to perform the validation against if you do not want the query to run against the graphmart that the step is in. When Validation Source Datasource is unset, the validation is performed against the graphmart.
- **Source**: The source data that this step should act upon. Steps can build upon the data generated by steps in other layers or can be self-contained, applying changes that relate only to the data defined in the layer that contains this step. You can select any number of the following options:
  - **Self**: This option is selected by default and means that the step runs against the data that is in the parent layer.
  - All Previous Layers Within Graphmart: This option means that the step runs against the data in all of the successful layers that precede the layer this step is in. Any failed layers are ignored.
  - **Previous Layer Within Graphmart**: This option means that the query runs against the data that is in the one layer that precedes the layer this step is in.
  - *Layer Name*: The Source drop-down list also includes options for layer names.
     You can choose a specific layer to act upon the data in that layer.
- **Pre-Run Generate Statistics**: This option controls whether to initiate AnzoGraph's internal statistics gathering queries before running the query to pre-compile. The statistics gathering helps ensure that the AnzoGraph query planner generates ideal query execution plans for queries that are run against the graphmart.
- 5. When you have finished configuring the Details tab, click the **Options** tab. This tab includes the settings that specify the type of check to perform on the data as well as instructions for what to do if the validation fails.

	Details	Options	Query	Query Context		
Check Type						
<ul> <li>Validati</li> </ul>	on					
O Conditi	on					
If the	validation query f	fails, the layer will be mar	ked as failed.			
If the	validation query f	fails, the whole graphman	t will be marked as	failed.		
Optional var	iable name to w	hich results are placed.				
Validation Res	alt Variable Name					
					CANCEL	SAVE

- 6. On the Options tab, determine which type of check to perform and select the appropriate radio button. There are two check types:
  - Validation: A Validation check validates the data according to the defined query (on the Query tab) and can be configured to take action depending on whether the validation passes or fails.
  - **Condition**: A Condition check takes the results of the query and associates it with the specified variable. That variable can then be used for setting up an execution condition at the layer or step level.
- 7. If you selected the Validation check type, you have the option to configure what to do if the validation query fails:
  - If the validation query fails, the layer will be marked as failed: Select this option if you want Anzo to abort the load of the layer if this step fails.
  - If the validation query fails, the whole graphmart will be marked as failed: Select this option if you want Anzo to abort the load of the entire graphmart if this step fails.
- 8. If you selected the Condition check type, you are required to specify the variable name that you want to use to store the result from the query. This variable becomes available as a choice when configuring an execution condition.
- 9. When you have finished configuring the Options tab, click the **Query** tab and compose the validation query that the step should run. The tab includes the syntax for writing a SPARQL

ASK query, which is useful for determining whether a certain pattern exists in the data. ASK queries return "true" or "false" to indicate whether a solution exists. The template includes a source graph parameter (\${fromSources}). Using the configured Source options from the Details tab, Anzo automatically populates the query with the appropriate source graph URIs when the query runs.

### Note

If your query connects to a source that requires input of connection and authorization information, Cambridge Semantics recommends that you do not include the connection and authorization values directly in the query. Instead, replace those values with Context Variables from a Query Context. You can access Context Providers for each data source from the step's Query Context tab. For detailed information about query contexts and referencing variables in a query, see Using Query Contexts.



10. Click **Save** to save the step configuration.

Once the Details tab is configured and the validation options and query are defined, the step can be run. For information about setting up an execution condition that uses this step, see Defining Execution Conditions.

# Construct a View of the Data (View Step)

This topic provides guidance on configuring a View Step to create a custom view of the data that does not change the graphmart or necessarily materialize any new data. View steps contain SPARQL CONSTRUCT queries to create a view definition in AnzoGraph. Follow the steps below to create a View.

- 1. Go to the graphmart for which you want to add a step and then click the **Data Layers** tab.
- 2. On the Data Layers tab, find the layer that you want to add the step to. Click the menu icon (1) for that layer and select Add Step/View. The Add Step/View dialog box is displayed with the New tab selected.

Add Step	p/View						
N	lew	Existing Steps	Existing Views				
<b>T</b> Q	Search	- *					
	Direct Load	i Step					^
	Load data dire	ectly from an external source					- 1
							- 1
	Elasticsear	ch Indexing Step					
<u></u>	Create a single	e Elasticsearch index to assoc	iate with a layer.				
	Elasticsear	ch Snapshot Step					
	Create an Elas	sticsearch snapshot of the inde	ex associated with a layer.				
	Export Step	p					
	This allows yo	u to export contents of graphr	nart to an flds.				
			Rows per page:	25 👻	1-12 of 12	<	>
					CANO	CEL	ОК

To create a new View step, select View and then click OK. If you want to clone an existing view, click the Existing Views tab, select the view that you want to clone, and then click OK. Anzo creates or clones the view and displays the Details tab:

Details	Query	Execution Condition	HI-Res Analytics	>
Title *				
The title of the view				
Description				
A brief description of the view Materialize the view wh	en activated, other	wise at runtime. 🗹 Enab	led	
Source				
Source data to act upon				
Data models *				¢ - V
Associated data models				
			CANCEL	SAVE

- 4. On the Details tab, configure the following options as needed:
  - Title: The required name of the step.
  - Description: An optional short description of the step.
  - Materialize the view when activated, otherwise at runtime: This option controls
    whether a copy of the data the view creates is saved in the graphmart (materialized) or
    whether this is a virtual view where the data is recreated each time the view runs. If you
    are creating a view against an extremely large data source or a source that changes
    often, typically the view should not be materialized. If you want to store a copy of the
    data that the view creates, select the Materialize the view when activated... check
    box. When this option is disabled Anzo creates a virtual view where only the view
    definition is stored in memory and not a copy of the data.
  - **Enabled**: When creating a new step, the Enabled option is selected by default, indicating that the step is enabled and will run when the layer is loaded or refreshed. If you want to disable the step so that it is not processed, clear the **Enabled** checkbox.
  - **Source**: The source data that this step should act upon. Steps can build upon the data generated by steps in other layers or can be self-contained, applying changes that relate only to the data defined in the layer that contains this step. You can select any number of the following options:

- **Self**: This option is selected by default and means that the step runs against the data that is in the parent layer.
- **All Previous Views Within Layer**: This options means that the step runs against the data that is generated by all of the previous views in the same layer.
- **Previous View Within Layer**: This options means that the step runs against the data that is generated by the previous view in the same layer.
- All Previous Layers Within Graphmart: This option means that the step runs against the data that is generated by all of the successful layers that precede the layer this step is in. Any failed layers are ignored.
- Previous Layer Within Graphmart: This option means that the query runs against only the data that is generated by the one layer that precedes the layer this step is in.
- Data models: This required field specifies the model or models that you want to create this view against. The list displays all of the available models. By default, the field is set to Exclude System Data (\*). If you want to choose a system model, click the toggle button on the right side of the field to change it to Include System Data (\*). When system data is included, the drop-down list displays the system models in addition to the user-generated models.
- 5. When you have finished configuring the Details tab, click the **Query** tab. This tab contains the query to use to create the view.



 Edit the provided template to compose the CONSTRUCT query the step should run. For information about CONSTRUCT queries, see <u>CONSTRUCT</u> in the W3C SPARQL 1.1 Query Language specification.

### Note

Do not include a GRAPH keyword in the CONSTRUCT clause. Anzo uses the view's URI as the graph URI for the constructed triples. In addition, Anzo uses the configured Source options from the Details tab to automatically replace the \${fromSources} parameter with the appropriate FROM clauses when the query runs.

You can click the **Open in Query Builder** button to open the query in the Query Builder, where you can perform practice runs to see results without having to refresh the graphmart or layer.

#### Note

If your query connects to a source that requires input of connection and authorization information, Cambridge Semantics recommends that you do not include the connection and authorization values directly in the query. Instead, replace those values with Context Variables from a Query Context. You can access Context Providers for each data source from the step's Query Context tab. For detailed information about query contexts and referencing variables in a query, see Using Query Contexts.

### Tip

If your view query employs the Graph Data Interface, be sure to use the following DataToolkitView service call in the query:

```
SERVICE <http://cambridgesemantics.com/services/DataToolkitView>
(${targetGraph})
```

For more information, see GDI Query Syntax.

7. Click **Save** to save the step configuration.

Once the Details tab is configured and the query is written, the step can be run. For information about running this step conditionally by setting up an execution condition, see Defining Execution Conditions. For details about the advanced Hi-Res Analytics settings that control how the view affects dashboards, see Hi-Res Analytics Tab below.

#### **Hi-Res Analytics Tab**

The **Hi-Res Analytics** tab contains advanced settings that control how the layer is exposed to and affects Hi-Res Analytic dashboards.

### Note

Changing these settings can have unexpected consequences, and Cambridge Semantics recommends that you do not modify them unless you understand the repercussions.



# Display data outputted by the view in Hi-Res Analytics dashboards

This setting controls whether the data accessed by the view is available to query and display in dashboards:

- When the setting is enabled (the default value), the view data is available to dashboards.
- When the setting is **disabled**, other data layers in the graphmart can use the view's data, but the data is not available to use Hi-Res Analytics dashboards.

# **Display View in Data Layers Facet in Hi-res Analytics**

This setting controls whether the view name is displayed in the Data Layers panel on dashboards. The image below shows an example Data Layers panel:

∧ Data Layers	0
Tickets Layer	
View	
Export	

- When the setting is **enabled** (the default value), the view is listed in the Data Layers panel.
- When the setting is **disabled**, the view's data is always used in dashboards for this graphmart but users do not see the view listed in the Data Layers panel.

### User can both view and hide the data outputted by the view

This setting controls whether users have the option to show and hide the view in the Data Layers panel on dashboards:

- When the setting is **enabled** (the default value), the view is listed in the Data Layers panel and users have the option to show and hide the layer.
- When the setting is disabled, whether the view shows up in the Data Layers panel depends on the By default, the data outputted by the view is visible setting. If the view is visible in the Data Layers panel ("By default, the data outputted by the view is visible" is enabled), users cannot toggle it on and off.

## By default, the data outputted by the view is visible

This setting controls whether the data generated by the view is visible in dashboards:

- When the setting is **enabled** (the default value), the view is listed in the Data Layers panel in dashboards and is selected by default.
- When the setting is disabled, the view shows up in the Data Layers panel but is not selected.
   To include the view's data in a dashboard, the user must select the view.

# Is Dynamic

Typically this option is used only for Graph Data Interface (GDI) connections where a remote data source is accessed and that source data changes dynamically. If the source is dynamic and you want Anzo to automatically refresh the view of the data at certain internals, select the **Is Dynamic** checkbox. Then set the Dynamic Refresh Interval (described below).

# **Dynamic Refresh Interval**

If the **Is Dynamic** option is enabled, this setting configures the interval at which Anzo queries the data source to retrieve any updated view data. Specify the number of milliseconds to wait between refreshes of the data.

# Hide from Other Layers/Views

This setting controls whether the other layers in the graphmart can act upon the data in this view.

- When the setting is **disabled** (the default value), this view is available as a choice in the **Source** drop-down list when a step is configured.
- When the setting is enabled, this view is not listed as a choice in the Source list.

# **Creating Data on Demand Endpoints**

With the Anzo Data on Demand service you can generate Open Data Protocol (OData)-based feeds that can be used to access Graphmarts programmatically via a RESTful API or from third-party business intelligence applications such as TIBCO Spotfire, Tableau, and Microsoft Power BI. The OData protocol enables web clients to use simple HTTP messages to access resources that are identified using URLs. OData shares some similarities with JDBC and ODBC. Like ODBC, OData is not limited to relational databases. The Anzo Data on Demand service follows the OData Version 4.0 specification, which defines the standard URL conventions, query options, and metadata schema.

Anzo supports two types of Data on Demand endpoints. The first type is called an **Auto-Generated** endpoint. Auto-Generated endpoints are the quickest type to create. They simply make available as-is all of the data in the selected Data Layers. Any joins, filters, and other operations must be performed by the consumer of the data outside of Anzo and AnzoGraph.

The second type of endpoint is called a **Custom** endpoint (sometimes called a **Table** endpoint). Since queries that join data often perform very poorly when run in BI applications with a JDBC driver, Custom endpoints let you assemble custom queries that join classes and apply filters and formulas. The endpoint becomes a view in AnzoGraph and AnzoGraph executes the custom queries in memory. Results can then be viewed from the endpoint without having to run the complex analytic queries over JDBC.

The topics in this section provide instructions for creating both types of endpoints.

- Creating an Auto-Generated Endpoint
- Creating a Custom Endpoint

## **Creating an Auto-Generated Endpoint**

Follow the instructions below to create an Auto-Generated Data on Demand endpoint. Auto-Generated endpoints can quickly be created to make available all of the data in the selected Data Layers. The data cannot be customized to exclude certain classes, join data across classes, or apply functions and formulas to properties. For instructions on creating an endpoint that can be customized, see Creating a Custom Endpoint.

1. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

▼ Q Search							
•	Title 🔶 🔶	Description	Status	Statements	Creator	Actions	÷
	🕸 Books		C Ready to use	156	System Administrator	Д	:
			C Ready to use	265,339	System Administrator	Д	:
	🕸 Movies		Ready to use	104,166	System Administrator	Д	:
	🕸 Northwind		C Ready to use	131,386	System Administrator	Д	:
	⊯ Tickets		C Ready to use	3,696,548	System Administrator	Д	:
				Rows	per page: 25 👻 1-5	of 5 <	>

- 2. On the Graphmarts screen, click the name of the graphmart for which you want to create an endpoint.
- 3. Click the **Data on Demand** tab. Anzo displays the Data on Demand screen, which lists any existing endpoints. For example, the image below shows a graphmart that does not have any endpoints configured:

*	EMR DB	ACTIVE @ Ready to	ouse c	S AZG Stat	0		Save a Copy	Create Dashboard	D
<	Overview	Explore	Datas	sets	Data Layers	t	Dashboards	Data on Demand	,
Creat	API Query End	points ≭ ←		It looks lii End	se you dont have point by clicking t	any Endpi the + Crea	oints yet. Yo ate New End	u can create a new point button.	

4. Click the **Create New Endpoint** button on the left side of the screen. Anzo displays the Create REST API Query Endpoint screen.

Create REST API Query En	dpoint	
Endpoint Name *		
Name of Endpoint		
Endpoint Description		
Description of Endpoint		
Endpoint Creation		
Auto-Generated		
Custom		
Denormalize Results		
Some BI tools work best with de	normalized resul	ts.
	CANCEL	SAVE

- 5. Configure the endpoint options on the screen as needed. The list below describes each setting:
  - Endpoint Name: Specify a name for the endpoint in this field. The endpoint name must be unique.
  - Endpoint Description: You can add an optional description for the endpoint in this field.
  - Endpoint Creation: This field specifies the type of endpoint to create. By default, the type is set to Auto-Generated. Leave the Auto-Generated radio button selected.
  - Denormalize Results: By default (when Denormalize Results is not selected), OData returns multi-valued properties as arrays. Certain BI tools, however, do not support arrays or multi-valued properties. If your data includes multi-valued properties and you plan to view the endpoint using a BI tool that does not support them, you can select the Denormalize Results setting to denormalize all multi-valued properties that are exposed in the endpoint. For JSON, XML, and CSV output formats, denormalization expands the properties into new rows so that they can be viewed in BI tools.
#### Tip

The following image shows an example of CSV output of multi-valued properties when **Denormalize Results** is disabled:

TeamId, TeamName, TeamToPlayer\_PlayerName, league\_tbal44\_key
1,Al Thomas, "[Fred Wynn, Steve Jones, James Smith]", aHR0cDovL2NzaS5jb20vVGVhbXMvMQ
2,Black Sox, "[Tim Hooper, Jared Bonds, Matt Butler]", aHR0cDovL2NzaS5jb20vVGVhbXMvMg
3,Braves, "[Billy Roper, Alex Granderson, Ted Sale]", aHR0cDovL2NzaS5jb20vVGVhbXMvMw
4,Somerville," [Chris Underwood, Mike Magazine, Ted James]", aHR0cDovL2NzaS5jb20vVGVhbXMvNNA

And the example below shows the output of multi-valued properties when

#### Denormalize Results is enabled:

TeamId, TeamName, PlayerName, table51_key
1,Al Thomas, James Smith, aHR0cDovL2NzaS5jb20vVGVhbXMvMQ
1,Al Thomas, Fred Wynn, aHR0cDovL2NzaS5jb20vVGVhbXMvMQ
1,Al Thomas, Steve Jones, aHR0cDovL2NzaS5jb20vVGVhbXMvMQ
2, Black Sox, Matt Butler, aHR0cDovL2NzaS5jb20vVGVhbXMvMg
2,Black Sox,Jared Bonds,aHR0cDovL2NzaS5jb20vVGVhbXMvMg
2,Black Sox,Tim Hooper,aHR0cDovL2NzaS5jb20vVGVhbXMvMg
3, Braves, Billy Roper, aHR0cDovL2NzaS5jb20vVGVhbXMvMw
3, Braves, Alex Granderson, aHR0cDovL2NzaS5jb20vVGVhbXMvMw
3, Braves, Ted Sale, aHR0cDovL2NzaS5jb20vVGVhbXMvMw
4, Somerville, Mike Magazine, aHR0cDovL2NzaS5jb20vVGVhbXMvNA
4. Somerville, Chris Underwood, aHR0cDovL2NzaS5jb20vVGVhbXMvNA
4. Somerville, Ted James, aHR0cDovL2NzaS5ib20vVGVbbXMvNA
.,,,,,,,,,,,,,

#### Note

Keep in mind that denormalizing the results means that more rows are created and transferred. Depending on the number of multi-valued properties and how the data is set up, you may see slower performance when creating and querying denormalized endpoints.

6. Click **Save** to create the endpoint and view the configuration details. For example:

Overview	Explore		Datasets	Data Layers	Dashboards	Data on Demand
REST API Query Endpoints Create New Endpoint	×	÷	Endpoint Name EMR data Endpoint Description Nome			
EMR data		- II	Enabled			
			Included Data layers <ul> <li>All</li> <li>Selected</li> </ul>			
			DB-QA Layer Cre	ate EMR dataset	and the set of the left	lana a Danadi sa sa
			Predicate used to retrieve v	ter or not to look up r	rame using endPointN	amePredicate.
			Cache Enabled	i		
			Denormalize R	esults		
			Connection String			
			Service root URLs used to	o retrieve class data from th	lis graphmart.	
			ODBC (SQL)	JDBC (SQL)		

- 7. You can make changes to any of the following additional configuration options that become available after the endpoint is created:
  - **Enabled**: By default the endpoint is set to **Enabled**, indicating that the endpoint is active. If you want to disable the endpoint, slide the Enabled slider to the left.

#### Note

If a request is sent to a disabled endpoint, Anzo displays a 503: Service Unavailable error with a message indicating that the endpoint is disabled. For example, "Unable to process request. The endpoint '<name>' is DISABLED."

• Included Data Layers: By default the Included Data Layers option is set to All, indicating that all of the layers in the graphmart are available from the endpoint. The included layers are listed below the radio buttons.

If you do not want to include all layers, click the **Selected** radio button. An **Edit Selections** link becomes available under the list of layers. Click **Edit Selections** to open the Select Data Layers dialog box. For example:

Select	Data Layers							
Q, Sear	ch	= • <b>1</b>	=					
	Title	$\uparrow$	Description	Updated Date	Tags			÷
	📚 Create EMR dataset			Nov 21, 2022				
	🕏 DB-OA Layer		Load data from DB-QA	Nov 21, 2022	Auto-Gen			
				Rows per page	25 🛩	1-2 of 2	<	>
2.ltem	s Selected ⊗							
					CANCE	L SAVE	SELECT	IONS

Clear the checkbox for any layer that you want to exclude from the endpoint, and then click **Save Selections** to save the change and return to the configuration screen.

- Controls whether or not to look up name using endPointNamePredicate: This setting controls which predicate value from the related model is used for the class and property display names in the endpoint. By default, the setting is enabled and the Predicate used to retrieve value for name from class or property is blank. That means Anzo uses the rdfs:label (http://www.w3.org/2000/01/rdf-schema#label) value for each class and property name.
  - If you want the endpoint to use a different value for class and property names, you can edit Predicate used to retrieve value for name from class or property to specify the URI for another predicate from the model. For example, specifying <a href="http://purl.org/dc/elements/1.1/description.would-use-each-entity">http://purl.org/dc/elements/1.1/description.would-use-each-entity</a>'s Description value.
  - If you disable the Controls whether or not to look up name using endPointNamePredicate setting, each entity's local name is used.
- Cache Enabled: When the endpoint is accessed, Anzo translates the OData query to a SPARQL query and sends it to AnzoGraph for execution. The Cache Enabled setting controls whether the results of that AnzoGraph query are cached in Anzo so that subsequent endpoint requests can run against the cache in Anzo. When Cache Enabled is disabled (the default setting), Anzo does not store the cache, and endpoint requests are sent to AnzoGraph. When Cache Enabled is selected, Anzo stores the

cached results and AnzoGraph only gets queried if the cached results are invalidated and need to be refreshed.

Once you are satisfied with the configuration, this Data on Demand endpoint is ready for access via OData/ODBC or JDBC. At the bottom of the screen, retrieve the ODBC or JDBC service URL to use to access the endpoint. For example:

Connection String	
Service root URLs used to retri	ieve class data from this graphmart.
ODBC (SQL)	JDBC (SQL)
https://10.102.0.17/dataonder	mand/EMR-DB/EMR-data 📋

To test whether the endpoint is active, you can copy the ODBC service URL and paste it into a web browser. If the endpoint is active, the browser shows an XML feed of the schema. For example:

This XML file does not appear to have any style information associated with it. The document tree is shown below.
<pre>vapp:service xmlns:atom="http://www.w3.org/2005/Atom" xmlns:app="http://www.w3.org/2007/app" xmlns:metadata="http://docs.oasis-open.org/edata/ns/metadata"</pre>
<pre>metadata:context="https://10.102.0.17/dataondemand/EPR-DB/EPR-data/\$metadata"&gt;</pre>
# capp:workspace>
<aton:title>Feeds.Default</aton:title>
<pre>v cappicollection bref="fm_Dbservationdescription" metadatione="fm_Dbservationdescription"&gt; cationalitileStructureS</pre>
<a )<="" for="" href="https://www.european.com/action/actio&lt;/td&gt;&lt;/tr&gt;&lt;tr&gt;&lt;td&gt;wighter collection bref. " matadate(name="For Observation" observation"="" td=""></a>
/ togethis a Serie Champeline ( stratting a series of the
<a estadata-name="Fee Medication" fee="" href="https://www.com/com/com/com/com/com/com/com/com/com/&lt;/td&gt;&lt;/tr&gt;&lt;tr&gt;&lt;td&gt;wrannenllection broke" medication"=""></a>
(atom itle) Em_Medication(atom tile)
<pre>vcapp:collection href="Emr_Specialty" metadata:name="Emr_Specialty"&gt;</pre>
<aton:title>Emr_Specialty</aton:title>
<pre>v<app:collection href="tmr_Complaint" metadata:name="tmr_Complaint"></app:collection></pre>
<atom:title>Emr_Complaint</atom:title>
<pre>vap:collection href~Tmr_Complaintdescription" metadata:name~Tmr_Complaintdescription"&gt;</pre>
<pre>vcap:collection href="Em_Medicationdescription" metadata:name="Em_Medicationdescription"&gt;</pre>
<pre>v.cpp:collection hrsf="Em_Activity" metadata:name="Em_Activity"&gt;</pre>
<pre>v<april of="" state="" td="" th<="" the="" to=""></april></pre>
<aton:title>Enr_Study</aton:title>
▼ <app:collection href="Emr_Patient" metadata:name="Emr_Patient"></app:collection>
<atom:title>Em_Patient</atom:title>
<pre>v<apr:collection href="Em_Signal" metadata:name="Em_Signal"></apr:collection></pre>

#### Note

The endpoint is accessible only when it is **Enabled** and the associated graphmart is **Active**.

For information about accessing endpoints programmatically, see Accessing an Endpoint Programmatically. For information about accessing endpoints with third-party analytics tools, see Accessing an Endpoint from an Application. For information about the supported OData operators, output format, and query examples, see OData Reference.

## **Creating a Custom Endpoint**

Follow the instructions below to create a Custom Data on Demand endpoint (sometimes called a Table endpoint). Creating a custom endpoint is similar to creating a dashboard Table lens in that you build a table with the columns that you want to see. You can traverse the relationships and join classes, add filters, and apply functions to properties. The tables are translated to SPARQL queries that create views in AnzoGraph, allowing you to interact with the graph for complex analytics but generate results in the tabular format that BI tools expect.

1. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

<b>Y</b>	<b>Q</b> Search	=	<b>ii</b> =			Add Graphn	hart 👻
•	Title $\uparrow$	Description	Status	Statements	Creator	Actions	÷
	🕸 Books		C Ready to use	156	System Administrator	Д	:
			C Ready to use	265,339	System Administrator	Д	:
	🕸 Movies		Ready to use	104,166	System Administrator	Д	:
	a Northwind		C Ready to use	131,386	System Administrator	Д	:
	© Tickets		C Ready to use	3,696,548	System Administrator	Д	:
				Rows	per page: 25 👻 1-5 o	f5 <	>

- 2. On the Graphmarts screen, click the name of the graphmart for which you want to create an endpoint.
- 3. Click the **Data on Demand** tab. Anzo displays the Data on Demand screen, which lists any existing endpoints. For example, the image below shows a graphmart without any endpoints:



4. Click the **Create New Endpoint** button on the left side of the screen. Anzo displays the Create REST API Query Endpoint screen:

Create REST API Query Endpoin	it	
Endpoint Name *		
Name of Endpoint		
Endpoint Description		
Description of Endpoint		
Endpoint Creation		
Auto-Generated		
O Custom		
-		
Dopormaliza Poquita		
Some BI tools work best with denormal	lized resul	ts.
	CANCEL	SAVE

- 5. Configure the endpoint options on the screen as needed. The list below describes each setting:
  - Endpoint Name: Specify a name for the endpoint in this field. The endpoint name must be unique.
  - Endpoint Description: You can add an optional description for the endpoint in this field.
  - Endpoint Creation: This field specifies the type of endpoint to create. By default, the type is set to Auto-Generated. Select the Custom radio button.

Denormalize Results: By default (when Denormalize Results is not selected), OData returns multi-valued properties as arrays. Certain BI tools, however, do not support arrays or multi-valued properties. If your data includes multi-valued properties and you plan to view the endpoint using a BI tool that does not support them, you can select the Denormalize Results setting to denormalize all multi-valued properties that are exposed in the endpoint. For JSON, XML, and CSV output formats, denormalization expands the properties into new rows so that they can be viewed in BI tools.

#### Tip

The following image shows an example of CSV output of multi-valued properties when **Denormalize Results** is disabled:

TeamId,TeamName,TeamToPlayer\_PlayerName,league\_tbal44\_key
1,Al Thomas,"[Fred Wynn,Steve Jones,James Smith]",aHROcDovL2NzaS5jb20vVGVhbXMvMQ
2,Black Sox,"[Tim Hooper,Jared Bonds,Matt Butler]",aHROcDovL2NzaS5jb20vVGVhbXMvMg
3,Braves,"[Billy Roper,Alex Granderson,Ted Sale]",aHROcDovL2NzaS5jb20vVGVhbXMvMw
4,Somerville,"[Chris Underwood,Mike Magazine,Ted James]",aHROcDovL2NzaS5jb20vVGVhbXMvNNA

And the example below shows the output of multi-valued properties when **Denormalize Results** is enabled:

TeamId, TeamName, PlayerName, table51_key
1,Al Thomas,James Smith,aHR0cDovL2NzaS5jb20vVGVhbXMvMQ
1,Al Thomas,Fred Wynn,aHR0cDovL2NzaS5jb20vVGVhbXMvMQ
1,Al Thomas,Steve Jones,aHR0cDovL2NzaS5jb20vVGVhbXMvMQ
2,Black Sox,Matt Butler,aHR0cDovL2NzaS5jb20vVGVhbXMvMg
2,Black Sox,Jared Bonds,aHR0cDovL2NzaS5jb20vVGVhbXMvMg
2,Black Sox,Tim Hooper,aHR0cDovL2NzaS5jb20vVGVhbXMvMg
3,Braves,Billy Roper,aHR0cDovL2NzaS5jb20vVGVhbXMvMw
3,Braves,Alex Granderson,aHR0cDovL2NzaS5jb20vVGVhbXMvMw
3, Braves, Ted Sale, aHR0cDovL2NzaS5jb20vVGVhbXMvMw
4,Somerville,Mike Magazine,aHR0cDovL2NzaS5jb20vVGVhbXMvNA
4, Somerville, Chris Underwood, aHR0cDovL2NzaS5jb20vVGVhbXMvNA
4.Somerville.Ted James.aHR0cDovL2NzaS5ib20vVGVhbXMvNA

#### Note

For Custom endpoints, you also have the option to denormalize data on a percolumn basis. If you do not want to denormalize all multi-valued properties, you can leave **Denormalize Results** disabled and then enable denormalization for specific columns when you build the endpoint views. Also note that denormalizing all results means that more rows are created and transferred. Depending on the number of multi-valued properties and how the data is set up, you may see slower performance when creating and querying denormalized endpoints.

6. When you have finished configuring the endpoint options, click **Save**. The endpoint is created and is empty until you create queries. For example, the image below shows a new endpoint:

Flights     Save a Copy     Create Dashboard       Derived from: Flights_20221123174517     Control     Control       INACTIVE     ACTIVE Control     ACG								
Overview	Explore	Datasets	Data Layers	Dashboards Data	on Demand >			
REST API Query Endpoints	₹ ←	Flights-Airlines	Queries	Connection Strings	Data Acce >			
Flights-Airlines	Φ 🔳	♥ Q Search No tables found	7	m	+ Add Query			

By default the endpoint is set to **Enable**, indicating that the endpoint is active. If you want to disable the endpoint, slide the **Enable** slider to the left.

#### Note

If a request is sent to a disabled endpoint, Anzo displays a 503: Service Unavailable error with a message indicating that the endpoint is disabled. For example, "Unable to process request. The endpoint '<name>' is DISABLED."

7. (Optional) You can click the **Details** tab to make changes to any of the following additional configuration options that become available after the endpoint is initially created:

Details	Queries
Endpoint Description flights by airline	
EndPoint Query Priority	
Extraction	
🗹 Cache Enabled	
Denormalize Results	

- Endpoint Query Priority: This setting controls where in Anzo's query queue the generated SPARQL queries for this endpoint are placed for processing, i.e., the priority for executing queries against this endpoint versus other types of queued queries like data layer and dashboard queries. By default, Endpoint Query Priority is set to
   Extraction, which is priority level 3 out of 8. The options are Interactive = 1, Extraction = 3, and Batch Process = 8.
- Cache Enabled: When the endpoint is accessed, Anzo translates the OData query to a SPARQL query and sends it to AnzoGraph for execution. The Cache Enabled setting controls whether the results of that AnzoGraph query are cached in Anzo so that subsequent endpoint requests can run against the cache in Anzo. When Cache Enabled is selected (the default setting), Anzo stores the cached results and AnzoGraph only gets queried if the cached results are invalidated and need to be refreshed. When Cache Enabled is disabled, Anzo does not store the cache, and endpoint requests are sent to AnzoGraph.
- 8. To start building a view, click the **Queries** tab (if necessary), and then click the **Add Query** button. The New Query dialog box is displayed.

New Query			- 20
Details	Query		
Title *			
The title of the view			
Description			
A brief description of the view			
Materialize			
Include Data Layers			
All layers ×			$\times \mid  \sim$
Source data to act upon			
		CANCEL	SAVE

- 9. On the **Details** tab, configure the following options as needed:
  - **Title**: Type a name for the table in the Title field. The name must be unique for the endpoint.
  - **Description**: You can add an optional description for the view in this field.
  - **Materialize**: If you want to store a copy of the data that this view creates (materialize the data), select the **Materialize** check box. When this option is disabled AnzoGraph creates a virtual view where only the view definition is stored in memory and not a copy of the data. If a request is made against this view, AnzoGraph temporarily materializes the data in memory, performs the query operations, and then drops the temporary data.
  - Include Data Layers: By default, the Include Data Layers option is set to include All Layers. If you do not want the query to target the source data in all layers, select the field and choose alternate layers from the drop-down list.

10. Click the **Query** tab to start building a table.

New Que	v				- 22
	Details	Query			
Starting point:	Select starting	oint	~		
	Set a	starting point to begin	defining your query		
				CANCEL	SAVE

- On the Query screen, get started by selecting a class to use as the starting point. Click the Select Starting Point drop-down list and select a class. Once you select the class, the Add Column button is displayed.
- 12. Click Add Column to create a column in the table. The Add Column dialog box is displayed:

Edit Column	
Formula Value	Sample Data No expression specified
Column Header Label	
Denormalize column data	
	CANCEL SAVE

- 13. On the Add Column screen, configure the following options as needed:
  - Formula Value: Click this field to choose the column header. Like building columns in a dashboard table, you can navigate the relationships to join data from different classes, and you can apply functions to the values. When you click the field and start to select paths and properties, sample data is shown on the right side of the screen. For example:



- **Column Header Label**: This is the label that you want to use for displaying the Formula Value.
- **Denormalize column data**: If you enabled **Denormalize Results** at the endpoint level, leave this setting disabled. If Denormalize Results is disabled at the endpoint level, you can enable this setting to denormalize the values for this column only.

#### Tip

If you do not denormalize the data, you can use the **Format** tab to specify the character to use for separating the values in the arrays that are returned for multi-valued properties. By default, the Value Separator is set to comma (,). The Format tab becomes available after saving the column.

- 14. Click **Save** to add the column to the table.
- 15. Add new columns to the table by clicking **Add Column**. You can change the order of columns by dragging a column up or down, and you can enable or disable columns by sliding the slider for the column. Click the menu icon (1) for a column to edit or delete that column.

vew Quer	y .		
	Details	Query	
arting point.	Flights10k	x   ~	
Airline	Flight Number	Distance	Elapsed Time
AS	08	1445	194
AA	2336	2830	279
us	840	2295	293
AA	258	2342	281
AS	135	1448	215
2	804	1485	
0	er page: 25 -	1-25 of 9999	$\langle \rangle$

- 16. If you would like to group by clause, you can click the Group By tab and select the column or columns to group.
- 17. If you would like to filter out some data, you can add one or more filters to the overall query. To add a filter, click the **Filters** tab and then click **Add Filter**. In the **Formula Value** field, specify the formula to use to determine which values should be included in the results. Then specify a name for the column in the **Column Header Label**. For example, the filter below is configured to include only the results where the Player ID is less than or equal to 5.

Ado	l Filter
Forr	nula Value /{Players}PlayerId{int} , 5)
Colu	nn Header Label
LE(	PlayerId, 5)

18. Once you are satisfied with the configuration, click Save to create the endpoint.

This endpoint is ready for access via OData/ODBC or JDBC. On the **Connection Strings** tab for the endpoint, retrieve the ODBC or JDBC service URL to use to access the endpoint. For example:



To test whether the endpoint is active, you can copy the ODBC service URL and paste it into a web browser.

#### Important

After pasting the URL into the browser, add the table name to the end of the string. The URL for a Custom endpoint is <ODBC\_or\_JDBC\_URL>/<table\_name>. For example, if the ODBC URL is https://10.10.0.11/dataondemand/Flights/Flights-Airlines, and the table name is Airlines, the connection string is <a href="https://10.10.0.11/dataondemand/Flights/Flights-Airlines/Airlines">https://10.10.0.11/dataondemand/Flights/Flights-Airlines</a>.

If the endpoint is active, the browser shows an XML feed of the data. For example:



## Note

The endpoint is accessible only when it is **Enabled** and the associated graphmart is **Active**.

For information about accessing endpoints programmatically, see Accessing an Endpoint Programmatically. For information about accessing endpoints with third-party analytics tools, see Accessing an Endpoint from an Application. For information about the supported OData operators, output format, and query examples, see OData Reference.

# **Sharing Access to Graphmarts**

This topic introduces the concepts to know when working with graphmart and data layer access control and provides instructions for configuring permissions.

- Sharing Concepts
- Changing Configuration-Level Access
- Changing Data-Level Access

## **Sharing Concepts**

This section describes the concepts that are helpful to know when working with graphmart and data layer permissions. It also gives of overview of the graphmart sharing settings and the predefined permission sets and associated privileges.

- Default Access Configuration
- Configuration vs. Data Access Control
- Permission Inheritance
- Configuration Permissions

#### **Default Access Configuration**

When a new graphmart is created, the access control configuration of that Graphmart is defined by the **Graphmarts Registry** Default Access Policy that is configured by your administrator (see <u>Managing Default Access Policies</u> in the Administration Guide for information). The graphmart also inherits permissions from other artifacts in the onboarding workflow. For example, when a graphmart is created from a data source, the graphmart inherits permissions from the source schema (which inherits permissions from the data source). Users who have permission to modify graphmart access can share that graphmart with other users and groups.

#### **Configuration vs. Data Access Control**

Graphmart and data layer sharing is managed on two levels: **Configuration** and **Data Access**. When managing access at the **Configuration** level, you are controlling who can view or modify the *configuration* of the graphmart, such as who can edit the graphmart settings on the Overview tab, who can enable, disable, modify, or add layers, and who can view or modify the graphmart permissions. The **Data Access** configuration controls who can view the data that is contained within the graphmart.

#### **Permission Inheritance**

When assigning Configuration and Data Access permissions at the graphmart level, you can configure the graphmart to inherit the permissions from another artifact and/or pass on its permissions to additional artifacts. For example, you can configure one graphmart to pass its permissions to other graphmarts. Inheritance transmits all of the artifact's permissions for all users and groups.

#### Note

Since data layers are created in graphmarts, they inherit their permissions from the graphmart by default—with one exception: Layers with Load Dataset Steps inherit their Data Access permissions from the dataset. Data on Demand endpoints also inherit their permissions from the parent graphmart by default.

The following inheritance settings are displayed at the top of the Configuration and Data Access tabs on the graphmart Sharing screen.

#### **Configuration Inheritance**

The image below shows a graphmart Configuration tab with the default inheritance settings. The **Inherit permissions from** field shows that the graphmart inherits permissions from the schema instance that the graphmart was created from.

	Configuration	Data Access	Warnings
Set ac within	ccess control on the Graphmart.	the Graphmart a	nd Data Layer configurations. These permissions do not control data access
Inherit	ance		
Inherit	permissions from		
Mov	ies-instance ×		x   ~

#### **Data Access Inheritance**

The image below shows the Data Access tab for the same graphmart. The **Graphmart Level View Permissions** are set to **Inherit from Graphmart** by default. And **Default Layer View Permissions** (for new Layers) is also set to **Inherit from Graphmart**.

Below the inheritance settings, the **Permissions Overview** provides a detailed view of the permission inheritance for each layer, view, and Data on Demand endpoint in the graphmart.

Configuration Dat	ta Access Warning	s				
Set access control to the d included Datasets. Inherita Access permissions are ch permissions to edit the cor	ata within the Graphr nce and custom pern anged from the origin figuration of the Grap	nart. Data Acc nissions may l nal permissior phmart or Dat	cess settings are inherited by be set for all Data Layers. Wi is inherited from the Dataset a Layers.	r default from either the Graphr amings are displayed whenever t. These permissions do not co	nart or the r the Data ntrol the	
Viewing Data in Layers						
Graphmart-Level View Permissi	ons		Default Layer View Permi:	ssions (for new Layers)		
Inherit from Graphmart	CHANGE		Inherit from Graphma	rt CHANGE		
Permissions Overview	<u>-</u> •					
Title	Туре	Warnings	Inherits ACLs from	Roles that can view	Actions	:
⊜ Movies	Graphmart		Graphmart Configuration	Everyone, sysadmin,	1	
S Movies Layer	Layer		Layer Configuration	Everyone, sysadmin,	/	
📚 Create dataset	Layer		Layer Configuration	Everyone, sysadmin,	1	

#### **Configuration Permissions**

Graphmart Configuration permissions control who can view or modify the graphmart settings, who can enable, disable, modify, or add data layers, and who can view or modify the graphmart permissions. There are three predefined permission sets that can be applied to a user or group. The permission sets include a combination of six permissions. You also have the option to customize the set of permissions that are applied to a user or group.

The tables below list the predefined permission sets and describe the privileges that are granted for each permission that is part of the set:

#### View

The following table describes the permissions in the View set.

Permission	Allows a user to:
View	<ul> <li>See the graphmart in the Anzo application.</li> <li>Copy the graphmart URI from the Overview tab.</li> <li>Copy data layer URIs from the data layers tab.</li> <li>See the existing Data on Demand endpoints on the Data on Demand tab.</li> <li>View and clone the dataset editions that are included in the graphmart.</li> <li>Reload and refresh the graphmart.</li> </ul>
	Create and import graphmart versions.
Meta View	This permission relates only to the graphmart Sharing tab. A user with this permission can see the Sharing tab, but they cannot modify, add, or remove permissions.

#### Modify

In addition to the **View** and **Meta View** permissions described above, the **Modify** set includes the **Add/Edit** and **Delete** permissions described below.

Permission	Allows a user to:
Add/Edit	<ul> <li>Rename the graphmart and edit the description.</li> </ul>
	Create Data on Demand endpoints.
	<ul> <li>Add datasets and data sources to the</li> </ul>

Permission	Allows a user to:
	<ul> <li>graphmart.</li> <li>Enable, disable, add, or edit layers and steps.</li> <li>Activate and deactivate the graphmart.</li> </ul>
Delete	<ul> <li>Remove datasets from the graphmart.</li> <li>Delete data layers and steps from the graphmart.</li> <li>Cannot delete the graphmart.</li> </ul>

#### Admin

In addition to the View, Meta View, Add/Edit, and Delete permissions described above, the Admin set includes the Meta Add/Edit and Meta Delete permissions described below.

Permission	Allows a user to:
Meta Add/Edit	This permission relates only to the graphmart Sharing tab. A user with this permission can modify the sharing settings by adding permissions to a user or group.
Meta Delete	<ul> <li>Modify the sharing settings by removing permissions from a user or group.</li> <li>Delete the graphmart.</li> </ul>

## **Changing Configuration-Level Access**

Follow the steps below if you want to modify the configuration-level access for a graphmart.

1. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

<b>Y</b> (	<b>Q</b> Search	=	· · 🖬 🗏			Add Graphr	nart 👻
•	Title $\uparrow$	Description	Status	Statements	Creator	Actions	÷
	🕸 Books		C Ready to use	156	System Administrator	Д	:
	Flights		C Ready to use	265,339	System Administrator	Д	:
	🕸 Movies		Ready to use	104,166	System Administrator	Д	:
	© Northwind		C Ready to use	131,386	System Administrator	Д	:
	Tickets		C Ready to use	3,696,548	System Administrator	Д	:
				Rows	per page: 25 👻 1-5 c	nf5 <	>

 Click the name of the graphmart for which you want to configure permissions. Then click the Sharing tab. The Sharing screen is displayed and the Configuration tab is selected. For example:

								T
*	Northwine	d			Save a Copy	Create Dashboard	Д	1
~	Derived from: North	hwind_20221123174519	2					
	INACTIVE (	■ ACTIVE ⊘ Ready	to use 💫 🛆 AZG 🤇 Static	2				
<	Data Layers	Dashboards	Data on Demand	Versions	Discussion	Sharing	>	
							-	
	Configuration	Data Access Warning	38					
Set acc	ess control on t	he Graphmart and Data	Layer configurations. These	e permissions do n	not control data access	vithin the Graphmart.		
			, ,					
Inheritar	nce							
Inherit p	ermissions from							
north	wind-instance ×		x V					
User and	d Group Permission	ns						
Permiss	ions defined for sp	ecific users or groups are ap	plied in addition to any inherite	d permissions.				
Q S	earch users, role	es or groups						
🚢 Ev	eryone	View 📗		1	0			
				/	9			
👗 Sy	stem Administra	ator Admin						
				Define the permi	ssions of this artifact.			1
			s	tart by selecting a	user or group on the let	t.		
								-

 If you want to change how the Configuration permissions are inherited, use the Inherit permissions from field at the top of the screen. To apply all of the permissions from another artifact to this one, select the artifact to inherit from in the Inherit permissions from field.

#### Tip

For more information about permission inheritance at the graphmart level, see Permission Inheritance.

4. To modify Configuration access to this graphmart with a particular user or group, type a value in the Search users, roles or groups field to find and display the user or group. The resulting list shows the current permission level that is set for each user or group in the search results. For example, the image below shows the current permissions for the IT group (None):



5. Select the user or group for which you want to configure permissions. The permissions settings are displayed on the right side of the screen. For example:

ermissions defined for s	pecific users or groups	are applied in addition	to any inherited	permissions.		
् info		Set permission for Info	nology (IT) armation Technol	907 (IT)		
2. Information Tech	nology (IT) None	Permissions	O View	O Modify	O Admin	O Custom
		Add/Edit		~	~	
		View	~	~	~	
		Delete		~	~	
		Meta Add/Edit			~	
		Meta View	~	~	~	
		Meta Delete			~	

6. To assign a predefined set of permissions, click the View, Modify, or Admin radio button to assign that level of access to the selected user or group. Refer to Configuration Permissions for details about the permission sets. For example, the image below gives Admin permissions to users in the IT group:

User and Group Permissions Permissions defined for specific users or groups	are applied in addition	to any inherited	permissions.				
info X Information Technology (IT) Set permission for information Technology (IT)							
Information Technology Admin     (IT)	Permissions	O View	O Modify	Admin	O Custom		
	Add/Edit		~	~			
	View	~	~	~			
	Delete		~	~			
	Meta Add/Edit			~			
	Meta View	~	$\checkmark$	~			
	Meta Delete			~			

If you want to customize the permissions, click the **Custom** radio button and then select or deselect the permissions checkboxes. To clear permissions for a user or group, click the trashcan icon (**I**) next to the name.

## **Changing Data-Level Access**

Follow the steps below if you want to modify permissions at the Data Access level for a Graphmart.

1. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

<b>Y</b>	<b>X</b> Search	=	· • 🖬 =					A	dd Graphr	nart 👻
•	Title 🛧 🛙	Description	Status	Statements		Creator			Actions	÷
	😸 Books		C Ready to use	156		System Ad	ministrator		Д	:
	Flights		C Ready to use	265,339		System Ac	ministrator		Д	:
	© Movies		C Ready to use	104,166		System Ac	ministrator		Д	:
	© Northwind		C Ready to use	131,386		System Ad	ministrator		Д	:
			C Ready to use	3,696,548		System Ad	iministrator		Д	:
					Rows pe	er page:	25 👻	1-5 of 5	5 <	>

 Click the name of the graphmart for which you want to configure permissions. Then click the Sharing tab. The Sharing screen is displayed and the Configuration tab is selected. Click the Data Access tab. For example:

Configuration Data Ac	cess Warnings				
Set access control to the data included Datasets. Inheritance Access permissions are chang permissions to edit the configu	within the Graphma and custom permis ed from the original ration of the Graph	rt. Data Acc sions may b permission mart or Data	ess settings are inherited by be set for all Data Layers. We is inherited from the Dataset a Layers.	default from either the Graphn Irnings are displayed whenever . These permissions do not co	nart or the the Data ntrol the
Viewing Data in Layers					
Graphmart-Level View Permissions			Default Layer View Permis	sions (for new Layers)	
Inherit from Graphmart CHA	NGE		Inherit from Graphma	rt CHANGE	
Permissions Overview					
Y Q Search	≞. ▼				
Title	Туре	Warnings	Inherits ACLs from	Roles that can view	Actions :
	Graphmart		Graphmart Configuration	Everyone, sysadmin,	1
S Movies Layer	Layer		Layer Configuration	Everyone, sysadmin,	1
📚 Create dataset	Layer		Layer Configuration	Everyone, sysadmin,	1

- 3. If you want to change how the Data Access permissions are inherited, use the fields at the top of the screen:
  - **Graphmart-Level View Permissions** controls who can view the data within the entire graphmart.
  - Default Layer View Permissions (for new Layers) controls who can view the data within the data layers.

#### Tip

For more information about permission inheritance at the graphmart level, see Permission Inheritance.

4. To change the permissions for an individual layer, Data on Demand endpoint, or another graphmart component that is listed in the Permissions Overview, click the Edit icon (✓) in the Actions column in the in the row for that component.

Changes to graphmart and layer permissions take effect immediately. Users do not need to log out and log back in, and affected graphmarts do not need to be reloaded or refreshed.

# **Graphmart FAQ**

This topic provides answers to frequently asked questions about graphmarts.

- What is the difference between graphmart reload vs. refresh?
- How do I find the URI for a graphmart?
- How do I find the URI for a layer?
- How do I see the models in a graphmart?

### What is the difference between graphmart reload vs. refresh?

When you make modifications to data layers in a graphmart, Anzo displays **Reload** and **Refresh** buttons on the top of the screen. For example:

Overview	Explore Datasets		Data Layers
C RELOAD C REFRESH		LAYERS 🗘	EXPAND ALL
Cost data from Movies			

The Refresh option becomes available when changes have been made to one or more data layers. Clicking **Refresh** resets (deletes from AnzoGraph) and reloads only the data layers that have changed. Clicking **Reload** resets and reloads the entire graphmart to AnzoGraph, including the data layers that have not changed.

## How do I find the URI for a graphmart?

Anzo displays graphmart details on the Overview screen for the graphmart. To view and copy a graphmart URI, go to the Overview tab for the graphmart. The URI is under General information on the right side of the screen. You can click the clipboard icon (<sup>[]</sup>) to copy the URI.

Gen	eral
Type :	Graphmart
Creator :	System Administrat
Last Accessed :	an hour ago
Last Updated :	an hour ago
Structure Modified :	Feb 23, 2023 10:58
Released :	Jan 23, 2023 5:19 P
http://cambridgesema	intics.com/Gra
http://cambridgesema	ntics.com/Gra
http://cambridgesema Dynamic Resource Lo	intlics.com/Gra
http://cambridgeseme Dynamic Resource Log AnzoGraph S	intics.com/Gra [] 9 ierver Details
http://cambridgeseme Dynamic Resource Log AnzoGraph S AnzoGraph :	intics.com/Gra [
http://combridgesense Dynamic Resource Log AnzoGraph S AnzoGraph : Status :	erver Details A2G Ready to use
http://cambridgesens Dynamic Resource Lo AnzoGraph S AnzoGraph : Status : Last Accessed :	AZG Ready to use 32 minutes ago
Dynamic Resource Lo AnzoGraph S AnzoGraph S Satus : Last Accessed : Memory Used :	Inticacionn/fera G ierver Details AZG Ready to use 32 minutes ago 187.05 MB(2%)
Dynamic Resource Los AnzoGraph S AnzoGraph S Status : Last Accessed : Memory Used : Memory Total :	erver Details AZG Ready to use 32 minutes ago 187.05 MB(2%) 11.7 GB

You can also copy a URL-encoded version of the graphmart URI from the address bar in the browser when viewing the graphmart. For example:



## How do I find the URI for a layer?

You can retrieve a data layer URI on the Data Layers screen for a graphmart. To view and copy a layer URI, go to the Data Layers tab for the graphmart. Anzo displays the layers. Each layer is a graph. For example:



Click the menu icon (I) for the layer whose URI you want to copy and click Copy URI:

Add Step/View
 Edit Layer
 Copy Layer
 Copy URI
 Delete Layer

## How do I see the models in a graphmart?

Anzo displays the list of included models on the Explore screen for a graphmart. To view the list of models:

- 1. Go to the Explore tab for the graphmart for which you want to view models.
- 2. In the top right corner of the graph view in the center of the screen, there are three icons:



- 3. To view the associated models, click the contents icon ( $\equiv$ ) on the right.
- 4. For example, the image below shows a graphmart with one model:



You can click the arrow icon  $(\rightarrow)$  to close the model panel.

# **Profiling Datasets and Graphmarts**

To help you explore your data and assess its quality, Anzo provides the option to generate a data profile for datasets and graphmarts. Creating a profile runs several metrics against the data and reports statistics at the class, property, and instance levels. Data profile metrics measure data quality, perform data discovery, and can help you decide on the types of analytics to run. The topics in this section provide instructions for generating data profiles and describe each of the metrics that Anzo runs.

In this section:	
Generating a Dataset Data Profile	
Generating a Graphmart Data Profile	610
Data Profiling Metrics	

## **Generating a Dataset Data Profile**

Similar to generating a profile for a graphmart, you can generate a data profile for a dataset in the Datasets catalog. Generating a dataset profile helps users perform data discovery, assess the quality of the onboarded data, and decide whether to use the dataset in a graphmart. The reports can also assist users in determining the types of data layer steps to create and writing the queries to include in the steps.

#### Important

To generate a dataset data profile, AnzoGraph must be online. If you have dynamic AnzoGraph deployments enabled, AnzoGraph will be provisioned automatically when the profile is generated.

1. In the Anzo application, expand the **Blend** menu and click **Datasets**. Anzo displays the Datasets screen, which lists the catalog of datasets. For example:

Y Q Search	₹ • 🖬 🔳			Add Da	ataset
□ - Title ↑	Description	Updated Date	Actions		:
1]- Books	Exported Books dataset		Д	:	
1). Flights	Exported Flights dataset		Д	:	
1). Movies	Exported Movies graphmart		Д	:	
1]+ Northwind	Exported Northwind graphmart		Д	:	
	Row	s per page: 25 👻	1-4 of 4	<	>

 On the Datasets screen, click the name of the dataset that you want to generate a profile for. Anzo displays the Explore tab for the dataset. For example, the image below shows a dataset for which a data profile has not been generated and the class and property counts are 0:



3. Click the **Profile Data** button on the left side of the screen. The Data Profile dialog box is displayed::



4. On the Data Profile screen, choose the Profiling Method to use. The **Primary** method focuses on data quality type analysis. The **Advanced** method includes the Primary data quality analytics plus several advanced metrics for deeper data discovery and analysis. For details about the Primary and Advanced metrics, see Data Profiling Metrics.

#### Note

To run the Advanced metrics, the optional AnzoGraph C++ extensions and dependencies must be installed. If you use dynamic, K8s-based deployments of AnzoGraph, the extensions are included. If you installed a static AnzoGraph instance with the installer, the C++ extensions are optional and are only installed if **yes** was **specified for the** Do you want to install C++ UDXs packaged with AnzoGraph DB? prompt. For information about the C++ dependencies, see Install the Optional C++ Extension Dependencies in the Deployment Guide.

5. Click **Run Profiling** to start generating the profile.

#### Note

If the dataset is used in a graphmart that is active, Anzo displays the Select Graphmart dialog box (shown below), which prompts you to choose whether the online dataset can be used for running the profiling queries or whether to provision another temporary graphmart for the dataset.

Select Graphmart		
Select Graphmart		~
	CANCEL	ОК

To use the dataset in the graphmart that is online, click the **Select Graphmart** dropdown list and select the graphmart name. If you want Anzo to provision a temporary graphmart instead, select **Don't reuse the Graphmart**.

The profiling process may take several minutes, especially for large volumes of data. You can check the status of the process in the Activity Log. The Activity Log also presents the option to stop the profiling process by clicking **Cancel** under the progress bar for the task. For example:

<b>Y</b> Q	Search		
*	Provisioning Graphmart - Movies graphmart. Provisioning Graphmart - Movies graphmart. Successfully activated the graphmart on AnzoGraph View Logs Run Time: 2 secs   End time: Nov 2, 2022 - 2:25:26 pm	Д	
	Profiling Metrics LinkedDataset - Movies Profiling Metrics LinkedDataset - Movies Profiling Metrics View Logs Elapsed Time: 7 secs Updating Dataset metrics for:Movies Cancel	Д	
	Rows per page: 25 ▾ 1-25 of 64 <	>	•
View All Histo	эгу		

6. Once the profiling is complete, the Profile Metrics panel is expanded on the Explore tab. To populate the panel, click a class or property in the ontology or a class in the graph view in the middle of the screen. For example:

Movies				+ Add to Cart	Create Graphmart
Overview Explore	Graphmarts	Pipelines	Dashboards	Versions	Discussion >
Ontology     Profile Date       Q. Search Ontology Items       ~ A Movies Layer Model		> : ≡	Profile Metrics Last profiled 2 minutes ag Movies Layer Model J Metrics	o / Movie Directors (2,77 Sample Data	∓ C →
O Movie Actors (4,891)     O Movie Directors (2,775)			Director Name		T String
or <sup>6</sup> Directed T Director Name <sup>1</sup> 2 <sub>3</sub> Movie Director ID T Movie Title > O Movie Producers (2,348) > O Movies (8,157)	Hore Active (4.97)	neer Upon Producers (1,500)	100% of instances have values	ner s	.5 0.75 1

You can click the Expand icon  $(\Box)$  on the right side of the screen to collapse the graph view and expand the metrics view.

Select any class or property to view its metrics. For details about each of the metrics that are run, see Data Profiling Metrics.

# Generating a Graphmart Data Profile

Similar to generating a data profile for a dataset, you can profile a graphmart, which may include multiple datasets. When metrics are generated for graphmarts, Anzo profiles the data that results from all of the enabled layers and reports metrics for the classes and properties in the model as well as statistics about the values for the properties. Generating a graphmart profile helps users perform data discovery, assess the quality of the data, and decide on the types of analytics to perform.

#### Important

To generate a graphmart data profile, AnzoGraph must be online. If you have dynamic AnzoGraph deployments enabled, AnzoGraph will be provisioned automatically when the profile is generated.

1. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

<b>Y</b>	<b>Q</b> Search	=			l	Add Graphr	mart 👻
•	Title 🔶 🛧	Description	Status	Statements	Creator	Actions	÷
	🕸 Books		C Ready to use	156	System Administrator	Д	:
	8 Flights		C Ready to use	265,339	System Administrator	Д	:
	≋ Movies		C Ready to use	104,166	System Administrator	Д	:
	© Northwind		C Ready to use	131,386	System Administrator	Д	:
	≆ Tickets		C Ready to use	3,696,548	System Administrator	Д	:
				Rows	s per page: 25 👻 1-5	of 5 <	>

2. On the Graphmarts screen, click the name of the graphmart for which you want to generate metrics. Anzo displays the Overview for that graphmart. For example:

*	Tickets Not Versioned 🖄 INACTIVE 💶	ACTIVE	🛆 AZG (Static)		Save a Copy	Create Dashboard
	Overview	Explore	Datasets	Data Layers	Dashboards	Data on Demand
Description						General
None						General
oad Priorit	tv					Type: Graphmart Creator: Sustem Administrat
lone					Last /	Accessed : Sminutes ado
					Last	Updated : 5 minutes ago
					Structure	Modified : 6 minutes ago
vietrics						Released : 10 minutes ago
Profile data	upon activation				http://caml	bridgesemantics.com/Graph 📋
Disable L Disable pop	LOAD COUNTS sulating counts during loa	d operations 📃 💻			Dynamic R	esource Log
Data Loa	ding Settings				Anz	oGraph Server Details
oad layers	that do not fail 🛛 🗩				A	nzoGraph: AZG
	_					Status : Ready to use
gnore Si	ource Errors				Last /	Accessed : 5 minutes ago
gnore any :	step or view's source laye	r if they failed to load			Mem	tory Used: 282.92 MB( 3%)
.eave Gr	aphmart Online Du	ring Refresh			Mem	tory Total := 11.7 GB
when refree	shing lavers, leave graphn	nart and laver online.				
					Inactivity D	eactivate Timeout
Manual F	Refresh Graphmart				None	
Once loade	d, changes only pushed to	o AnzoGraph manually.(Only affects J	ournal Based Data) 🔵 🗯			

#### 3. Click the **Explore** tab.



4. If necessary, activate the graphmart, and then click the **Profile Data** button on the left side of the screen. The Data Profile dialog box is displayed:

Profile Data		
Generate metrics for data quality and discovery at the class, property, and instance level. Data Profiling can be a time-consuming operation for larger volumes of data.		
Profiling method		
Primary     Executes a subset of metrics focused on data quality.     Advanced     Executes all available metrics (including Primary ones) for further da	ta discovery an	d analytics.
LEARN MORE		
	CANCEL	RUN PROFILING

5. On the Data Profile screen, choose the Profiling Method to use. The **Primary** method focuses on data quality type analysis. The **Advanced** method includes the Primary data quality analytics plus several advanced metrics for deeper data discovery and analysis. For details about the Primary and Advanced metrics, see Data Profiling Metrics.

#### Note

To run the Advanced metrics, the optional AnzoGraph C++ extensions and dependencies must be installed. If you use dynamic, K8s-based deployments of AnzoGraph, the extensions are included. If you installed a static AnzoGraph instance with the installer, the C++ extensions are optional and are only installed if **yes** was **specified for the** Do you want to install C++ UDXs packaged with AnzoGraph DB? prompt. For information about the C++ dependencies, see Install the Optional C++ Extension Dependencies in the Deployment Guide.

6. Click **Run Profiling** to start generating the profile. The process may take several minutes, especially for large volumes of data. You can check the status of the process in the Activity Log. The Activity Log also presents the option to stop the profiling process by clicking **Cancel** under the progress bar for the task. For example:


7. Once the profiling is complete, the Profile Metrics panel is expanded on the Explore tab. To populate the panel, click a class or property in the Ontology or a class in the graph view in the middle of the screen. For example:



You can click the Expand icon (C) on the right side of the screen to collapse the graph view and expand the metrics view.

Select any class or property to view its metrics. For details about each of the metrics that are run, see Data Profiling Metrics.

# **Data Profiling Metrics**

When a data profile is generated for a dataset or graphmart, Anzo runs several metrics that can help users measure data quality and perform data discovery at the class, property, and instance level. The metrics are grouped into two categories: a **Primary** category that focuses on data quality type analysis, and an **Advanced** category that includes the Primary data quality analytics plus several advanced metrics for deeper data discovery and analysis.

#### Note

To run the Advanced metrics, the optional AnzoGraph C++ extensions and dependencies must be installed.

The lists below give a summary of the Primary and Advanced metrics. For more information and sample images of the visualizations that are generated, click a metric name.

## **Primary**

- Row Count: Reports the total row (instance) count per class.
- Absent Property: For each property, reports the total number of instances that do not have a value.
- Empty Property: For each string property, reports the total number of empty strings.
- Sample Values: Returns sample values for each property.
- Average: For each numeric property, computes the average of all values.
- Sum: For each numeric property, computes the sum of all values.
- Average String Length: For each string property, computes the average length of the strings.
- Value Present: For each property, computes the percentage of instances that have at least one value.
- Unique Pattern Count: For each property, counts the total number of unique value patterns.
- Unique Values Count: For each property, counts the total number of unique values.

- Median: For each numeric property, computes the median of all values.
- Standard Deviation: For each **numeric** property, computes the standard deviation of all values.
- Mode: For each numeric property, computes the mode of all of the values.
- Presence: For each property in a class, reports the percentage of instances that have values vs. do not have values.
- Top Value Counts: For each property, computes the top *N* most occurring values.
- Bottom Value Counts: For each property, computes the *N* least occurring values.
- Top Pattern Counts: For each property, computes the top *N* most common value patterns.
- Bottom Pattern Counts: For each property, computes the *N* least common value patterns.
- Range: For numeric properties, reports the total range of values.
- Value Types: For each property, returns the data types for the instances.
- DateTime Distribution By Year/Month/Day: For **dateTime** properties, computes a histogram that shows the distribution of values by year, month, and day.

#### Advanced

- Pearson Skewness: For each numeric property, computes the Pearson coefficient of skewness.
- Geometric Mean: For each numeric property, computes the geometric mean of all values.
- Variance: For each numeric property, computes the variance of all values.
- Discrete Entropy: For each property, computes the discrete entropy of all values.
- Discrete Probability: For each property, computes the discrete probability of all values.
- String Length Range: For each string property, reports the range of string lengths.
- Unique Values: For each property, computes the percentage of unique values.

- Lower Case Strings: For each **string** property, computes the percentage of values with all lower case characters.
- Upper Case Strings: For each **string** property, computes the percentage of values with all upper case characters.
- Trivial Values: For each string property, computes the percentage of instances that have one of the following values: NA, N/A, NONE, or NULL.

#### **Row Count**

This metric computes the total row count per class. The row count is in parenthesis at the top of the screen next to the class name, as shown in the image below.



#### **Absent Property**

For each property, this metric reports on the total number of instances that do not have a value for that property.



#### **Empty Property**

For string properties, this metric reports the total number of empty strings.

northwind - Auto / Employees (17) / Country					
Primary N	Primary Metrics				
<b>64.71%</b>	Absent Property: 6	Average String Leng	th: <b>2.55</b>		
Present	Empty Property: 0	Value Present:	64.71%		
	Unique Pattern Count: 2	Unique Values Coun	t: <b>2</b>		

#### Sample Values

This metric returns sample values for each property.

Sample Values	
Beverages	Condiments
Confections	Dairy Products
Grains/Cereals	Meat/Poultry
Produce	Seafood

#### Average

For each numeric property, this metric computes the average of all values.



#### Sum

For each numeric property, this metric computes the sum of all values.

northwind	- Auto / Products	(77) / UnitPrice
Primary N	Metrics	
	Absent Property:	0
$\checkmark$	Range:	3 - 264
100%	Median:	19.5
instances	Value Present:	100%
have values	Sum:	2222.71
	Unique Values Cour	it: 62

For each string property, this metric computes the average length of the values.

northwind - Auto / Territories $(53)$ / TerritoryDescription				
Primary N	Metrics			
~	Absent Property:	0	Average String Leng	th: 50
100%	Empty Property:	0	Value Present:	100%
instances have values	Unique Pattern Cou	nt: <b>17</b>	Unique Values Coun	t: <b>52</b>

#### **Value Present**

For each property, this metric computes the percentage of instances that have at least one value.

northwind - Auto / Current Product List (69) / ProductID				
Primary N	<b>Netrics</b>			
_	Absent Property:	0	Empty Property:	0
$\langle \rangle$	Range:	1 - 77	Average:	40.52
100%	Median:	41	Mode:	16
or instances	Value Present:	100%	Standard Deviation:	22.59
have values	Sum:	2796	Unique Pattern Count	t: <b>2</b>
$\bigcirc$	Unique Values Cour	nt: 69		

#### **Unique Pattern Count**

For each property, this metric counts the total number of unique value patterns.



#### **Unique Values Count**

For each property, this metric counts the total number of unique values.



#### Median

For each numeric property, this metric computes the median of all values.

northwind - Auto / Invoices (2,155) / Quantity				
Primary Metrics				
	Absent Property:	0		
$\sim$	Range:	1 - 130		
100%	Median:	20		
instances	Value Present:	100%		
have values	Sum:	51317		
	Unique Values Coun	t: 55		

#### **Standard Deviation**

For each numeric property, this metric computes the standard deviation of all values.



#### Mode

For each numeric property, this metric computes the mode of all of the values.



#### Presence

The metric is available when viewing a class. For each property in the class, this metric reports on the percentage of instances that have values and the percentage of instances that do not have values.



#### **Top Value Counts**

For each property, this metric computes the top *N* most occurring values.



#### **Bottom Value Counts**

For each property, this metric computes the *N* least occurring values.



### **Top Pattern Counts**

For each property, this metric computes the top *N* most common value patterns.



#### **Bottom Pattern Counts**

For each property, this metric computes the *N* least common value patterns.



#### Range

For numeric properties, this metric computes the range of all values.

northwind	- Auto / Product	s (77) / UnitsO	nOrder	
Primary I	Metrics			
	Absent Property:	0	Empty Property:	0
<ul> <li></li> <li></li> </ul>	Range:	0 - 100	Average:	10.13
100%	Median:	0	Mode:	0
or instances	Value Present:	100%	Standard Deviation:	23.14
have values	Sum:	780	Unique Pattern Count	t <b>3</b>
10.000	Unique Values Cou	int: 10		

#### Value Types

For each property, this metric returns the data types for the instances.



#### DateTime Distribution By Year/Month/Day

For each dateTime property, this metric computes a histogram that shows the distribution of values by year, month, and day.



#### **Pearson Skewness**

For each numeric property, this metric computes the Pearson coefficient of skewness to show the distribution of values. A value of 0 indicates no skew, a positive number indicates positive skew, and a negative number indicates negative skew.

Advanced Metrics	
Mean Absolute Deviation:	2.15
Unique Values:	0%
Geometric Mean:	3.59
Interquartile Range:	5
Negative Numbers:	0%
Pearson Skewness:	0.48
Zero Numbers:	0%

#### **Geometric Mean**

For each numeric property, this metric computes the geometric mean of all of the values.

	Advanced Metrics	
	Mean Absolute Deviation:	2.15
	Unique Values:	0%
<	Geometric Mean:	3.59
	Interquartile Range:	5
	Negative Numbers:	0%
	Pearson Skewness:	0.48
	Zero Numbers:	0%

#### Variance

For each numeric property, this metric computes the variance of all values.

Advanced Metrics	
Mean Absolute Deviation:	16.15
Unique Values:	0%
Geometric Mean:	0
Interquartile Range:	20
Negative Numbers:	0%
Positive Numbers:	99.67%
Variance:	891.91

#### **Discrete Entropy**

For each property, this metric computes the discrete entropy of all values.

orthwind - Auto / Ord	er Details Extend	ed (2,155) / UnitPrice	
Advanced Metrics			
Mean Absolute Deviation:	16.15	Discrete Entropy:	6.48
Unique Values:	0%	Galton Skewness:	0.36
Geometric Mean:	0	Interquartile Mean:	19.75
Interquartile Range:	20	Kurtosis:	33.98
Negative Numbers:	0%	Outliers Count:	97
Positive Numbers:	99.67%	Pearson Skewness:	0.78
Variance:	891.91	Zero Numbers:	0.33%

#### **Discrete Probability**

**Discrete Probability** 18.0 10.0 14.0 12.5 19.0 38.0 14.4 15.0 21.0 34.0 0 0.005 0.01 0.015 0.02 0.025 0.03 0.035 0.04 0.045 0.05 Values

For each property, this metric computes the discrete probability of all values.

#### **String Length Range**

For each string property, this metric reports the range of string value lengths.



#### **Unique Values**

For each property, this metric computes the percentage of unique values.

er Details Extend	led (2,155) / ProductName	
6.1	Unique Values:	0%
0%	Lower Case Strings:	0%
0%	Numeric Strings:	0%
4 - 32	Spaces in values:	1
0%	Upper Case Strings:	0%
	er Details Extend 6.1 0% 0% 4 - 32 0%	er Details Extended (2,155) / ProductName 6.1 Unique Values: 0% Lower Case Strings: 0% Numeric Strings: 4 - 32 Spaces in values: 0% Upper Case Strings:

#### Lower Case Strings

For each string property, this metric computes the percentage of values that have all lower case characters.

rthwind - Auto / Ord	ers (830) / ShipA	ddress	
Advanced Metrics			
Discrete Entropy:	6.23	Unique Values:	0%
HTML Tags:	0%	Lower Case Strings:	2.29%
Non Printable Characters:	0%	Numeric Strings:	0%
String Length Range:	10 - 45	Spaces in values:	2
Trivial Values:	0%	Upper Case Strings:	0%

#### **Upper Case Strings**

For each string property, this metric computes the percentage of values that have all upper case characters.

northwind - Auto / Orde	ers (830) / Ship	Address	
Advanced Metrics			
Discrete Entropy:	6.23	Unique Values:	0%
HTML Tags:	0%	Lower Case Strings:	2.29%
Non Printable Characters:	0%	Numeric Strings:	0%
String Length Range:	10 - 45	Spaces in values:	2
Trivial Values:	0%	Upper Case Strings:	0%

#### **Trivial Values**

For each string property, this metric computes the percentage of instances that have one of the following values: **NA**, **N/A**, **NONE**, or **NULL**.

n	northwind - Auto / Customers (93) / ContactName						
	Advanced Metrics						
	Discrete Entropy:	6.54	Unique Values:	100%			
	HTML Tags:	0%	Lower Case Strings:	0%			
	Non Printable Characters:	0%	Numeric Strings:	0%			
	String Length Range:	4 - 23	Spaces in values:	1			
	Trivial Values:	0%	Upper Case Strings:	0%			

For additional metrics based on the type of data quality checks needed, contact Cambridge Semantics.

# Access & Analyze

Once data has been onboarded, users have several options for accessing and analyzing the data. Anzo includes the Hi-Res Analytics application where users can create dashboards for exploring and visualizing the data without needing to have specialized query knowledge. The Query Builder enables users to find specific statements or run SPARQL queries. Users can also access data from the SPARQL endpoint or by using the Data on Demand service to generate data feeds for thirdparty business intelligence tools. The topics in this section provide information about the ways to access data in Anzo as well as information about sharing, versioning, and migrating artifacts.

In this sectio	n:
----------------	----

Access Data with Hi-Res Analytics Dashboards	628
Access Data with the Query Builder	
Access Data on Demand Endpoints	921
Access the SPARQL Endpoint	
Access the HTTP Client Interface	
Share Access to Artifacts	
Version and Migrate Artifacts	
SPARQL Best Practices and Query Templates	
Function and Formula Reference	

# Access Data with Hi-Res Analytics Dashboards

The Anzo Hi-Res Analytics application enables users to answer both ad-hoc and pre-determined questions using custom dashboards. Automated query generation eliminates the need to have specialized query knowledge, and users can traverse complex, multi-dimensional data by building exploratory charts, filters, tables, and network views.

The topics in this section provide guidance on getting started with Hi-Res Analytics dashboards and include instructions for creating and modifying dashboards and dashboard components.

In this section:	
Introduction to Hi-Res Analytics	629
Getting Started: Explore and Visualize Your Data	637
Working with Dashboards	648
Working with Lenses	697
Working with Filters	827
Calculating Values in Lenses and Filters	881
Combining Data from Multiple Classes	887
Searching for Text in Unstructured Documents	892
Sharing Access to Dashboards and Lenses	899

# **Introduction to Hi-Res Analytics**

Hi-Res Analytics dashboards enable you to create visual representations of your data using the latest in powerful web technologies. This introduction defines the fundamental concepts of working with dashboards and provides an overview of the Hi-Res Analytics user interface.

#### Tip

To fully leverage the advanced capabilities of Hi-Res Analytics, it helps to have skills working with Excel functions and formulas, SPARQL, and JavaScript and HTML. You can create dashboards without these skills but may not be able to take advantage of all functions.

- Concepts and Vocabulary
- Application Overview

## **Concepts and Vocabulary**

Term	Description
Dashboard	Dashboards are containers for the elements that enable you to visualize, analyze, and share your data. Data is viewed through lenses, such as tables and charts, that format the data for display. You can also apply filters to the data to refine the results.
	There are two types of dashboards that you can create: a <b>Graphmart</b> <b>Dashboard</b> that offers several choices of lens and filter types, and a <b>Network</b> <b>Navigator Dashboard</b> , which is an interactive graph visualization tool for exploring the relationships in a graph.
Data Layer	Since graphmarts typically have multiple data layers, users can include or exclude the data from certain layers when creating or viewing dashboards.
Lens	Lenses are the structures that display your data. You must have at least one lens in a dashboard. You can reuse existing lenses or create new ones. For more information, see Creating a Lens.

Term	Description
Filter	Filters narrow and further define the data to display. Dashboard-level filters apply globally to all lenses in a dashboard. Lens-level filters apply only to a specific lens and are not displayed on the lens. They are shown only in the lens designer. For more information, see Working with Filters.
Property	A property is a predicate that contains the instance data to display. The data type of a property determines the functional aspects within a dashboard. For example, certain filters act only on dates or numbers.
Path	Paths are relationships. They are transitional elements that allow you to connect data across classes.
Functions and formulas	Functions and formulas can be applied to properties to modify the data that is presented. Available functions depend on the property's data type. For more information, see Calculating Values in Lenses and Filters.

# **Application Overview**

This section gives an overview of the user interface. The images below show an administrator view. Some options are not available to users with lower permission levels.



#### Main Toolbar

The toolbar at the top of the screen provides the following options:

- **Dashboard**: This menu includes options to save or reset the dashboard and create a new dashboard or open an existing one. This menu also provides access to the permission settings, properties, and the option to delete the dashboard.
- Lenses: This menu includes options to create a new lens or open an existing one.
- Filters: This menu includes options to create a new filter, clear selected filters, and show all active filters.
- Refresh: This menu includes an Automatic option that controls whether the dashboard is automatically updated when the underlying graphmart data changes. In addition, the menu includes a Show Update Controls option that, when enabled, adds pause icons to lenses and filters so that users can pause the refresh of individual components and then manually update them when desired.
- **Designer**: This menu opens the dashboard designer, which includes settings that control dashboard layout and design as well as the update method to use for the overall dashboard.

- **Help**: This menu includes options to open the Progress window as well as the Query Manager and the documentation.
- User: This menu presents the option to log out of the application.

#### **Dashboard Tabs**

The dashboard tabs under the main toolbar display the open dashboards and enable you to navigate between dashboards. When you change a dashboard, an asterisk appears on the dashboard tab. Save the dashboard to preserve the changes.



#### **Filters Panel**

The left panel (shown below) contains the dashboard filters. Each of the panels are described below.



#### Manage Lenses

Clicking **Manage Lenses** opens the Manage Lenses dialog box, which enables you to show and hide the items in the panel. For example:

Manage Lenses	×
Filters hidden from view will still impact displayed results are currently applying filters.	
Graphmart	🔽 ON
Data Types	<b>Ο</b> Ν
Data Layers	IN ON
Venue City	ON 💟
Categories	🔽 ON
Event Date	V ON
Event Name	🗹 ON
Venue State	ON 🔽
Restore Defaults	
CANCEL	ОК

Clearing an **ON** checkbox hides that item in the panel. You can select a checkbox to show an item or click **Restore Defaults** to show all items.

## Tip

When the dashboard layout is a **Vertical List Container**, there is also a **Manage Lenses** button available for showing and hiding lenses. For example:

Manage Lenses													
∧ All Events												:	×
eventid		event	name					star	ttime				
521		West	Side St	ory				1/2/2	2008 10:0	00 AM			
			<	<	>	>1	1 - 20	) of 115	Page	1	of 6	\$	•
∧ Venues												:	×
Venue Venue													
wenue/	Charles Playho	use		ve	Fenv enue/v	vay Parl Venuen	c ame		TD B	anknor	th Gard	en	

#### Graphmart

The Graphmart panel displays the selected graphmart for the dashboard:

∧ Graphmart	+
Northwind	

### Data Types

The Data Types panel displays the selected data type from the graphmart.

∧ Data Types	+	:
Orders (48)		

#### **Data Layers**

The Data Layers panel displays the layers in the graphmart.

∧ Data Layers	:
DB - northwind	

#### Filters

By default, filters that you create appear in the left column of the dashboard.

<ul> <li>ShipCity (5)</li> </ul>	:	$  \times$
Search		↓A
(+) show filters		
Aachen (6)		
Albuquerque (18)		
Anchorage (10)		
Barcelona (5)		

#### **Object Toolbar and Tabs**

The object toolbar and tabs enable you to manage the lenses and filters in the selected dashboard. The tabs display the open lenses, and the toolbar enables you work with the lens configuration.

🔅 Designer	Properties		Manage Columns	
What can I	do next? ×	Novie Table $ imes$	Popularity Chart $\times$	

# **Getting Started: Explore and Visualize Your Data**

When you start to build a new dashboard, you might not know what data exists in the knowledge graph, which values in graph you ultimately want to display, and the most pertinent way to visualize the results. This topic introduces the available lenses and filters and provides guidance on getting started by using the Hi-Res Analytics tools to perform data discovery. By experimenting with simple objects, you can explore the data, determine which questions you want to answer, and start to visualize the end result.

To get started:

- 1. Create a New Dashboard
- 2. Explore the Data
- 3. Create Visualizations of the Data

#### **Create a New Dashboard**

 In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

<b>Y</b>	<b>Q</b> Search	Ŧ	- <b>.</b> =			Add Graphn	nart 👻
•	Title 🔶	Description	Status	Statements	Creator	Actions	÷
	🕸 Books		C Ready to use	156	System Administrator	Д	:
	Flights		C Ready to use	265,339	System Administrator	Д	:
	🕸 Movies		C Ready to use	104,166	System Administrator	Д	:
	🕸 Northwind		C Ready to use	131,386	System Administrator	Д	:
	🕸 Tickets		C Ready to use	3,696,548	System Administrator	Д	:
				Row	rs per page: 25 <del>↓</del> 1-5	iof5 <	>

2. On the Graphmarts screen, click the name of the graphmart for which you want to create a dashboard. The Overview screen is displayed. For example:

*	Northwind	ACTIVE	AZG Static	Sav	e a Copy Create Dashboard
	Overview	Explore	Datasets	Data Layers	Dashboards Data on >
escription					General
one ad Priorit one	ty				Type : Graphmart Creator : System Administrat Last Accessed : a few seconds ago Last Updated : a few seconds ago Structure Motified : More 2 0023 1472 PM
etrics	upon activation				Released : Nov 2, 2022 1:00 A
sable L	oad Counts				http://cambridgesemantics.com/Graph
sable pop	ulating counts during loa	ad operations			Dynamic Resource Log
ata Loa ad layers	that do not fail				AnzoGraph Server Details AnzoGraph: AZG
nore So	ource Errors step or view's source laye	r if they failed to load			Status: Ready to use Last Accessed: a few seconds ago Memory Used: 266.78 MB(3%)
ave Gr	aphmart Online Du shing layers, leave graphr	ring Refresh nart and layer online.			Memory Total: 11.7 GB
Aanual F	Refresh Graphmart	Ŭ			None

- 3. Click the **Create Dashboard** button. The Hi-Res Analytics application opens and displays the Create Dashboard dialog box. Leave **Graphmart Dashboard** selected and click **Next**.
- 4. Next, type a name for the dashboard in the **Title** field and enter an optional **Description**.

Create Dashboard			×
General Information			
Title:*			
Description:			
Graphmart:*			
Northwind			
	CANCEL	PREVIOUS	FINISH

5. Click **Finish** to create the dashboard. The new dashboard appears as a new tab and contains a sub-tab titled **What can I do next?**. This tab acts as a wizard to guide you through the initial dashboard creation. For example:

	ard 🗸	✿ Lenses - ♡ Filters - C Refresh - ✿ Designer   � Help -
•Northwind Orders ×  • Ready to use <b>Factors Amage Lenses</b> • Graphmart Northwind <b>A Data Types</b> + No items of the selected type available	< + :	What can I do next? ×         Dashboard         This dashboard will allow you to visualize linked data         Image: Specify the type of data you would like to see
Data Layers     DB - northwind to store	Ξ	

6. In the main toolbar, click the **Dashboard** button and select **Save**. Proceed to Explore the Data below for guidance on next steps.

# **Explore the Data**

Once you create a new dashboard, you can experiment with Hi-Res Analytics tools to get to know the data and decide the best way to display it.

Decide What Type (Class) of Data You Want to See

 First, review the types of data or classes that exist in the data. On the What can I do next? tab, click Specify the types of data you would like to see. The Select Data Types dialog box displays the available data types. The value in parentheses shows the total number of instances of that type exist in the data set:



2. Select one data type. The property that you choose determines the fields that become available to filter on.

#### Tip

Though you must choose one base data type for a dashboard, you can leverage the relationships in the graph to access and integrate data from additional classes. See Combining Data from Multiple Classes for more information.

Click **OK** to close the Select Data Types dialog box. The data type is added to the Data Types panel on the left side of the screen. Proceed to Create Filters to See the Values for Properties below for next steps.

#### Тір

You might want to create multiple dashboards so that you can click between dashboards and view multiple classes of data at the same time.

#### **Create Filters to See the Values for Properties**

To dive deeper into the data and quickly determine what values exist for the class of properties you selected, you can start adding filters to the dashboard. Filters reveal the values associated with fields and help you learn the data set specifics such as whether data exists for certain properties

and whether the data includes many duplicate or unique values. Learning more about the details enables you to start making decisions about what properties to group on, what properties have relationships, and what results you want to visualize on the dashboard.

1. To create a filter, click **Add facets to filter the data** on the What can I do next tab. The Create Filter dialog box opens.

Create Filter	×
Fields:*	
Click to edit	
Filter:*	

- 2. In the Create Filter dialog box, click the **Fields** field and select the property or property path to filter on.
- 3. Then click the **Filter** field to select the filter type. The list of available choices depends on the data type of the property you selected in Fields. The table below describes each filter type.

Filter	Description
Cloud	Cloud filters display values in term clouds where each term is written in a font size that represents the number of results for that value. Unlike list filters, which enable you to select and filter on multiple values at once, cloud filters allow you to filter on one value at a time. The Cloud filter is available for all data types.
List	List filters display values in a list and allow you to filter on multiple values at the same time. List filters are available for all data types.
Single Select List	Single Select List filters are similar to List filters but only allow you to filter one value from the list at a time. This type of filter is available for properties of all data types but is not available for paths.

#### **Filter Types**

Filter	Description
Limit	Limit filters are used to limit the results on the dashboard to a specified number of either the largest or smallest values. The Limit filter is available for any data type. For strings, results are ordered alphabetically. "Largest" orders by the last letters in the alphabet and "Smallest" orders by the first letters in the alphabet.
Date Range	Date Range filters are used to limit the results on a dashboard to data that falls in (or outside of) certain date and time groupings. Date Range filters are available for properties with date, dateTime, and time data types.
Numeric Range	Numeric Range filters are used to limit the results on a dashboard to data that falls in (or outside of) certain numeric groupings. Numeric Range filters are available for properties with integer and double data types.
Range Slider	Range Slider filters display a slider control that enables you to filter dashboard results by setting one range that you can adjust as needed. This type of filter is available for properties with integer, double, date, time, and dateTime data types. It is not available for paths.
Relative Time	Relative Time filters are used to filter for records that fall into a specified increment of time relative to the current time. This type of filter is available for date, time, and dateTime data types.
Search	Search filters are used to search for values of a property that contain a partial match, exact match, or do not equal the text that you specify. The search is case-insensitive. This type of filter is available for all data types. It is not available for use with paths.
Presence	Presence filters group results based on whether the value exists or does not exist. This type of filter is useful for testing whether there are records that are missing a particular value. Presence filters are available for paths and properties of all data types.

Filter	Description
Quartile	Quartile filters group and rank the values for a property into four equal ranges. This filter is available for properties with integer, double, date, time, and dateTime data types. It is not available for paths.
Hierarchy	If hierarchies exist in your knowledge graph, you can create a Hierarchy filter to explore the parent and child relationships and filter the dashboard based on the relationships. Unlike the majority of dashboard filters, where you select a property to filter on, Hierarchy filters operate on relationships and are only available as a filter type when you select a path to filter on.
Types	Types filters are used to filter data according to the types of data (classes) that are connected by a specified path. This type of filter is available only for paths and not properties.

For information about configuring each type of filter, see Working with Filters.

# **Create Visualizations of the Data**

Once you have a good understanding of the values and relationships that exist in the data, you can experiment with the Hi-Res Analytics lenses and decide on the most appropriate way to display the data. Creating a Table lens is a quick way to view the data that you filtered. This section provides instructions for creating a table lens and describes each of the lenses available in Hi-Res Analytics.

1. To create a Table lens, click **Select or create visualizations of your data** in the What can I do next tab. Anzo displays the Lens Selection dialog box.

Lens Selection	×
Q Search for lenses	
+ Path to lens values Results: 0 No suitable lenses found	
Selection: 0	
Create a new lens	ок

2. In the dialog box, click **Create a new lens**. Anzo displays the Create Lens dialog box.

Cre	ate Lens		×
Selec	ct the type of lens to create		
0	Search for lenses		
	Solicit for Istibus		
+ Path Result	to lens values s: 10		-
	Table Tabular visualization of data Properties		
₽/	Form View and edit resource attributes. Properties		
սի	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

3. Select the Table lens and click Next.

4. Type a **Title** for the lens, and then click **Finish**. Anzo opens the Table Designer:

Designer		×
Columns + = d A	Column Details     Set properties of the selected column.	
Untitled Disable   View   Clone Remove	Column Header Label:	
	Column Value Expression:	
	Click to edit	
		Format   Hyperlink   Editor
	Column Footer Expression:	
	Click to edit	Second Hard A
	Column Width:	Format   Hyperlink
	Content Filters	
	Group rows by:	
Add a column I Display Properties	Click to edit	
Preview changes		CANCEL SAVE

5. In the Designer, click the **Auto-generate columns** icon (\*) to add all available columns to the table. Then click **Save**.

The new lens displays as a new sub-tab on the dashboard and displays the data according to the data type and filter or filters that you created. Now that you can view a summary of the data in a table, it can help you determine how to further narrow or expand the results by adding, changing, or removing filters. In addition, you can experiment by adding other lenses to the dashboard to find the ideal way to display the data to answer the questions that you have. The table below describes each type of lens. For more information about each lens, see Creating a Lens.

Lens	Description
AnzoKO Web Page	AnzoKO Web Page lenses include the <u>Knockout JavaScript</u> framework and enable you to display data on a web page that you create using HTML, CSS, and JavaScript.
Chart	Anzo offers several types of chart lenses. These lenses are useful for displaying large amounts of complex data and have the widest format range of any lens type. The ability to add an axis enables you to compare data, such as for comparing monthly sales data for multiple stores.
Drill Down	Drill down lenses combine other lenses into a hierarchical interface. Clicking an

Lens	Description	
	object in one lens opens the next lens in the hierarchy and can present the data in a different view.	
Form	Form lenses enable you to create an editable or read-only form on the dashboard. Creating forms can be useful for displaying many details about each record instead of using a table where the large number of columns makes the data hard to read.	
	<b>Note</b> By default, only the sysadmin user has access to create Form lenses. In addition, Form lenses are valid in only in Linked Dataset dashboards. For more information, see Creating a Form Lens.	
List	List lenses display results in a list layout, similar to the Microsoft Windows® Explorer interface. The lens enables you to add icons for each data value, and results are grouped onto pages according to the given page size value.	
Query	Query lenses enable you to retrieve data using a custom SPARQL query and display the results by writing basic HTML and CSS. You can use a Query lens to access data from external sources. Query lenses do not bind directly to the linked data set, data type, or filters defined on the dashboard.	
Resource Tree Navigator	Resource Tree Navigator lenses display results in a hierarchical tree view. You can click parent data points to open the successive child data points. This lens is useful for presenting small amounts of data; each discrete group appears on a separate page in the dashboard. You can also click certain objects to view the object's data properties in the left panel.	
	<b>Note</b> By default, the only user who has permission to create a Resource Tree Navigator lens is the <b>sysadmin</b> user.	

Lens	Description
Table	Table lenses present results in a basic table grid consisting of rows and columns. Table lenses are useful for presenting data aggregates or summaries.
Web Page	Web Page lenses enable you to display results on a web page that you create using HTML, CSS, and JavaScript.

# Working with Dashboards

The topics in this section provide instructions on creating and configuring dashboards.

- Creating a Graphmart Dashboard
- Creating a Network Navigator Dashboard
- Configuring a Dashboard to Update in Batch Reporting vs. Interactive Mode
- Capturing User-Defined Values in Dashboards

## **Creating a Graphmart Dashboard**

Follow the instructions below to create a new dashboard for a graphmart. For instructions on creating a Network Navigator dashboard, see Creating a Network Navigator Dashboard.

1. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

▼ Q Search =				Add Graphmart 👻	
🗌 👻 Title 🛛 🛧	Description Status	Statements	Creator	Actions	÷
🕸 Books	C Ready to use	156	System Administrator	Д	:
😂 Flights	Ready to use	265,339	System Administrator	Д	:
≅ Movies	Ready to use	104,166	System Administrator	Д	:
© Northwind	C Ready to use	131,386	System Administrator	Д	:
© Tickets	C Ready to use	3,696,548	System Administrator	Д	:
		Ro	wys per page: 25 ┯ 1-5 c	of 5 <	>

2. On the Graphmarts screen, click the name of the graphmart for which you want to create a dashboard. The Overview screen is displayed. For example:
| *  | Northwind                                  | ACTIVE Ø Ready to use        | AZG Static | Save  | e a Copy Create Dsshboard  |
|--|--|------------------------------|------------|---|--|
| _  | Overview                                   | Explore                      | Datasets   | Data Layers                                 | Dashboards Data on   |
| Description  | 1  |                              |            |   | General  |
| None<br>Load Priorit<br>None   | ty   |                              |            |   | Type : Graphmart<br>Creator : System Administrat.<br>Last Accessed : a few seconds ago<br>Last Updated : a few seconds ago<br>Structure Modified : Nov 2, 2022 1:57 PM |
| <b>letrics</b><br>rofile data  | upon activation                            |                              |            |   | Released : Nov 2, 2022 10:06 A.  |
| isable l   | oad Counts                                 |                              |            |   | http://cambridgesemantics.com/Graph  |
| isable pop   | pulating counts during loa                 | d operations                 |            |   | Dynamic Resource Log   |
| Data Loading Settings AnzoGraph Server De Load layers that do not fail AnzoGraph Server De AnzoGraph : A20 |  |                              |            | AnzoGraph Server Details<br>AnzoGraph : AZG |  |
| gnore So<br>gnore any :  | ource Errors<br>step or view's source laye | r if they failed to load 🛛 💻 |            |   | Status: Ready to use<br>Last Accessed: a few seconds ago<br>Memory Used: 266.78 MB(3%)   |
| eave Gr  | aphmart Online Du                          | ring Refresh                 |            |   | Memory Total: 11.7 GB  |
| /hen refres  | shing layers, leave graphn                 | nart and layer online.       |            |   | Inactivity Deactivate Timeout  |
| Manual F   | Refresh Granhmart                          |                              |            |   | None   |

- 3. Click the **Create Dashboard** button. The Hi-Res Analytics application opens and displays the Create Dashboard dialog box. Leave **Graphmart Dashboard** selected and click **Next**.
- 4. Next, type a name for the dashboard in the **Title** field and enter an optional **Description**.

General Information		
Titler#		
1 ILNE,		
Description:		
Graphmart.*		

5. Click **Finish** to create the dashboard. The new dashboard appears as a new tab and contains a sub-tab titled **What can I do next?**. This tab acts as a wizard to guide you through the initial dashboard creation. For example:

	d <del>-</del>	✿ Lenses - ♡ Filters - C Refresh - ✿ Designer   ❶ Help -
*Northwind Orders ×		
PLTONS  Manage Lenses  Graphmart	< +	What can I do next? ×           Dashboard           This dashboard will allow you to visualize linked data
Northwind	i	<ul> <li>B Add data you are interested in exploring</li> <li>If Specify the type of data you would like to see</li> <li>Select or create visualizations of your data</li> <li>Add facets to filter the data</li> </ul>
No items of the selected type available	I	
Data Layers     DB - northwind to store		

6. On the What can I do next? tab, click Specify the types of data you would like to see. The Select Data Types dialog box displays the available data types. The value in parentheses shows the total number of instances of that type exist in the graphmart:



7. In the Select Data Types dialog box, select the data type or class of data that you want to display on the dashboard. Anzo uses the type, along with any filters, to populate the visualizations (lenses) that you add to the dashboard.

### Tip

Though you must choose one base data type for a dashboard, you can leverage the relationships in the graph to access and integrate data from additional classes. See Combining Data from Multiple Classes for more information.

- 8. Click **OK** to close the Select Data Types dialog box. The data type is added to the Data Types panel on the left side of the dashboard.
- 9. In the main toolbar, click the **Dashboard** button and select **Save** to save your progress.

Now that the dashboard basics are defined, see Creating a Lens and Working with Filters for instructions on adding lenses and filters to the dashboard.

## **Creating a Network Navigator Dashboard**

The Network Navigator Dashboard is an interactive graph visualization tool that enables you to find relationships in the knowledge graph and explore the paths to build out a Network View. This topic helps you get started building a Network Navigator Dashboard by covering the basic steps and functionality that is presented by default. Additional topics describe the functionality in further detail and cover more advanced dashboard configuration options. For instructions on creating a Graphmart dashboard, see Creating a Graphmart Dashboard.

#### Note

By default, only the sysadmin user has permission to create a Network Navigator Dashboard. However, the sysadmin user can share created dashboards with other users and groups (see Sharing Access to Dashboards and Lenses).

1. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

<b>Y</b>	<b>Q</b> Search	=				Add Graphn	nart 👻
•	Title 🔶 🔶	Description	Status	Statements	Creator	Actions	÷
	🕸 Books		C Ready to use	156	System Administrator	Д	:
	Flights		C Ready to use	265,339	System Administrator	Д	:
	⇒ Movies		C Ready to use	104,166	System Administrator	Д	:
	© Northwind		C Ready to use	131,386	System Administrator	Д	:
			C Ready to use	3,696,548	System Administrator	Д	:
				Rows	s per page: 25 👻 1-5 o	15 <	>

2. On the Graphmarts screen, click the name of the graphmart for which you want to create a dashboard. The Overview screen is displayed. For example:

*		n Graphmart		Sav	e a Copy Create Da	shboard 🔲 🔲
	INACTIVE 🛑	ACTIVE	AZG Static			
	Overview	Explore	Datasets	Data Layers	Dashboards	Data on 🔸
escription					Gener	al
one					Type: G	raphmart
ad Priorit	ty				Creator : S	ystem Administrat
00					Last Accessed : 9	minutes ago
					Structure Modified : 9	minutes ago
<b>Netrics</b>					Released: 9	minutes ago
rofile data	upon activation 🛛 💻					-
isable L	oad Counts				http://cambridgesemant	ics.com/Graph
isable pop	ulating counts during loa	d operations			Dynamic Resource Log	
ata Loa	ding Settings				AnzoGraph Ser	ver Details
oad layers	that do not fail 🛛 🔲				AnzoGraph : A	ZG
anoro Si	ource Errore				Status : R	eady to use
ynore arv s	sten or view's source lave	r if they failed to load			Last Accessed : 8	minutes ago
pare only a	way or men a availed laye				Memory Used: 2 Memory Total: 1	1.7 GB
eave Gr	aphmart Online Du	ring Refresh				
/hen refres	shing layers, leave graphn	nart and layer online. 🛛 💻			Inactivity Deactivate Tim	neout
Manual Pafrash Granhmart None						

3. Click the **Create Dashboard** button. The Hi-Res Analytics application opens and displays the Create Dashboard dialog box. Select **Network Navigator Dashboard** and click **Next**.

Remember that only the **sysadmin** user has the option to create this type of dashboard. If you do not see the option and think you should, check with your system administrator.



4. Next, type a name for the dashboard in the **Title** field and enter an optional **Description**. For example:

Create Dashboard Specify details about the new dashboard			×
General Information			
Explore Supply Chain			
Explore suppliers for raw materials Graphmart*			
[VW] Supply Chain Graphmart			
	CANCEL	PREVIOUS	FINISH

5. Click **Finish** to create the dashboard. The dashboard appears as a new tab in the application. For example:

	📰 Dashboard 👻   💠 Designer   \Theta Help 👻	<b>.</b>
Explore Supply Chain ×		
	Find Data   Open View Save View   Download Image If Clear   Image Image   Image   Image Image   Image Imag	> Find Paths Select two nodes to find paths between them. Degrees of Separation G Cayout Tightness Re-layout C d Cayout C C Cayout C C C C C C C C C C C C C C C C C C C

6. The first step in rendering the Network View is to choose a node (or group of nodes) as the starting point. A node is an object in the data. In order to be able to traverse the network, start with an object for a property that has one or more relationships with other properties. To select a value, click the **Find Data** button at the top of the dashboard. The Find Data dialog box is displayed and lists all of the values for each of the properties in the data model. Select one or more objects in the list.

Fin	d Data		×
	d Data Q. Search Results: 171 ADAPTER_BLOWBY-ROHR_N57_US type: Item AESC type: Tier-1 Altainnano type: Tier-1 AUSCHUESC DESCRIPTOR NATTO		×
	ANSCHLUSS_ZKH_BBR_US_N47_O1_L7 type: Item Selection: 0	CANCEL	ок

To narrow down the list of values by class or type of data, you can click **Filter** to open the Filter panel on the left side of the screen. The Filter panel lists the classes and subclasses from the model. For example, the image below shows the classes for a knowledge graph that contains supply chain data for an automobile manufacturer:

FILTERS	<
Q Find types	
<ul> <li>Supply Chain</li> </ul>	
+ Company Country Item Message Rohmaterial	

You can expand or collapse the classes to display or hide subclasses. And you can start typing text in the **Find Types** field to search for specific classes. Select the checkbox next to any of the classes from which you want to display data. The list in the main part of the screen is refreshed to display only the values for the properties in the selected classes.

7. Once you have selected the starting node in the Find Data dialog box, click **OK** to add the node to the Network View. For example, in the image below the value **Vale mines** was selected as the starting point, and the Vale mines node is added to the View:

*Explore Supply Chain ×		
Ready to use		
<ul> <li>Ready to use</li> <li>Select a single node to view its properties.</li> </ul>	Find Dats   Open View Save View Download Image If: Clear Vale mines	> Find Paths Select two nodes to find paths between them. Degrees of Reparation 4 €0 Layout Tightness 3 ♂ ★ Be-layout €dge Labels ● Node Labels ●

### Tip

You can configure the dashboard to **Auto-Expand** by a certain number of degrees or hops so that adding a node to a View automatically adds the specified number of related nodes and paths. For information, see <u>Auto-Expanding a Network View</u>.

8. To start building out the network, select a node to view its connections and properties on the left side of the screen. For example, selecting the **Vale mines** node shows its incoming and outgoing connections as well as properties like the label.

<	Find Data	🗀 Open View	Save View	🕹 Download Image	≣⊧ Clear
Vale mines					
Connections <pre></pre>				Vale mines	
Properties label Vale mines					

The image shows that Vale mines has an incoming connection (or backward path o«) from **Mining Company** via **hasMine** and an outgoing connection (or forward path o») to **Country** via **locatedIn**.

9. There are two options for exploring the connections:

## See a list of connected nodes without adding them to the View

To see a list of nodes that are connected by a path without adding the nodes and path to the View, you can click the path name link.



Details are shown in the lower panel, but the nodes are hidden from the network View by default. For example, clicking the **locatedIn** path from the image above, shows that Vale mines is located in Canada.

Vale mines / locatedIn / Country				
Q Search Vale mines locatedIn				
Displaying 1 Results				
🐼 Canada				

If you decide you want a node and path to be added to the View, you can click the hidden node to make it visible. For example, clicking **Canada** adds the path and node to the View:

	Find Data Dopen View 🗟 Save View 🛃 Download Image 📦 Cle	ear
Vale mines		
Mine		
	Ð	
Connections	÷	
O≪ MiningCompany		
hasMine (1)	E E	
O>> Country	$\sim$	
locatedin (1)		
	Vale mines Ca	anad
Vale mines / locatedIn / Country		
Q Search Vale mines locatedI		
Displaying 1 Results		
O Canada		

## Add a path and node directly to the View

To add a path and node directly to the network View, you can select the checkbox next to the path.

For example, selecting the **hasMine** checkbox adds the path and the MiningCompany value (**Vale**) to the View:



### Tip

Any time you want to clear the View and start again, you can click **Clear** at the top of the screen. To toggle between showing and hiding a path or node, you can click the path or node in the panel on the left side of the screen to select or deselect the item.

You can save the View by clicking **Save View** and specifying a Title and optional Description in the dialog box. When you save the dashboard, the View that is visible at the time it is saved is presented when the dashboard is reopened.

10. To continue to explore and expand the Network View, you can select another node and repeat the previous step to identify additional paths to follow. Or you also have the option to select two nodes and do a search for paths between them. To find paths, find and add to the view, if necessary, the two nodes whose relationships you want to explore. Then hold the **Ctrl** button and select the two nodes.

For example, in the image below, two nodes, an OEM called VW Group, and a raw material, Nickel, were added to the View, and both nodes are selected:



11. On the right side of the screen under **Find Paths**, adjust the **Degrees of Separation** value as needed (the default is 4) and then click **Go**. Anzo runs an All Paths graph algorithm to find all of the paths that exist between the two nodes—where the maximum number of hops is the value in the **Degrees of Separation** field. If a "No paths found" message is returned, that means there is no path between the two nodes that is *N* or fewer hops away, where *N* is the value in Degrees of Separation. You can change the value and try again.

In the example for VW Group and Nickel, finding paths with the default value of 4 returned "No paths found." However, increasing Degrees of Separation to 5 renders several nodes and relationships.



### Тір

You can remove a single node or a hide a group of nodes by right-clicking a node and selecting an option from the pop-up menu (shown below). **Remove Selected Node** deletes the node and its incoming and outgoing connections. **Hide Selected Tree** hides the selected node and all of its connected nodes and paths.



12. You can continue to explore the network using the methods described above. In addition, you can adjust the layout of the View using the options under **Layout** on the right side of the screen:

# Tightness

The **Tightness** setting controls how close or far apart the nodes are. Valid values are 1 - 9. The lower the Tightness value, the looser the nodes are. Increasing Tightness brings the nodes closer together. Type a number in the field and then press **Enter** to apply it. In the simple example below, the default Tightness value of **3** shows nodes that are far apart.



Increasing Tightness to 9 attracts the nodes closer together:



## Standard vs. Hierarchical Layout

The layout buttons (shown below) enable you to switch between a **Standard** ( $\varkappa$ ) layout (the default setting) and a **Hierarchical** ( $\clubsuit$ ) layout.



To change to a Hierarchical layout from a Standard layout, select the node in the View that is the root node and then click the **Hierarchy** (♣) button. For example, in the image below,

**VW Group** was selected as the root node. Changing to the Hierarchical layout changes the View to a hierarchy:



## **Re-Layout**

The **Re-Layout** setting controls whether the entire View is refreshed and rearranged when a node is added. When Re-Layout is enabled, the View is rearranged when a node is added. If Re-Layout is disabled, new nodes are added to the View and existing nodes remain in place.

## Show or Hide Edge and Node Labels

The Eye (●) icons next to **Edge Labels** and **Node Labels** enable you to toggle between showing or hiding the edge and node labels in the View.

You can save the View any time by clicking **Save View**. You are prompted to specify a **Title** and optional **Description**. To avoid creating multiple Views with the same name, make sure that you specify a unique Title. When you save the dashboard, the View that is visible at the time it is saved is presented when the dashboard is reopened. Other Views can be opened by clicking **Open View** and selecting a View. Only one View is displayed at time per dashboard. You also have the option to export a PNG version of the Network View by clicking **Download Image**.

#### Adding Icons to a Network Navigator Dashboard

In Network Navigator Dashboards, you can use custom icons to represent different types (classes) of data in the knowledge graph. The steps below guide you through identifying the types in the Network View, if necessary, so that you can choose icons, uploading the icons the library, and applying the icons to specific classes.

- Identifying Types and Preparing Icons
- Uploading and Applying Icons to a Dashboard

#### Note

To be able to upload and apply icons, a Network Navigator Dashboard must exist and have a Network View that displays at least one node. For instructions on creating a dashboard, see Creating a Network Navigator Dashboard.

#### **Identifying Types and Preparing Icons**

When configuring a dashboard to display icons, the icons are applied at the class (type) level. That way each property of the same type is represented by the same icon. If you are unfamiliar with the data that is displayed in a View, there are a couple of ways you can identify the classes:

You can select a node in the View and see the type below the node label on the left side of the screen. For example, the image below shows that the AxionPowerInternational node is in the Tier-1 class.



• You can click **Find Data** and open the **Filter** panel. Expand the contents to see the classes and subclasses. For example:



Once you know the types of the nodes in the View, you can determine the appropriate icon to use for each type. Download the icons to your computer, if necessary. When uploading icons to the library, you will browse your computer to select them. The application supports standard image formats like PNG, JPEG, and SVG.

The examples from the list above show a portion of a View and a list of the classes and subclasses for an automobile supply chain knowledge graph. The entire View (shown below) shows companies that are potential suppliers of Cobalt to the VW Group. In the example, adding icons would hep users distinguish between the first tier companies, mining companies, refineries, mines, etc.



Uploading and Applying Icons to a Dashboard

Follow the instructions below to upload images to the icon library and then apply them to the classes in your Network View.

#### Note

When you apply icons to the classes (types) in a dashboard, those icons will apply to all Network Views that include those types.

- 1. Open in the Hi-Res Analytics application the Network Navigator Dashboard that you want to add icons to.
- 2. In the main toolbar, click the **Designer** button. The **Properties by Type** tab is displayed in the Designer. For example:



3. Expand the view of the model to display each of the types (classes). If the model contains subclasses, you can view them by clicking the **View Sub Types** link next to the parent class (shown below).

- Supply Chain						
Company (view su	b types)					
Label Path: http://www.w3.org/2000/01/rdf-schema#label						
+ Node Styling						
+ Edge Styling						
+ Advanced						

When viewing the subtypes, you can click **Back** to return to the parent level. For example:

Properties by Type	Options
Back	
Mine	
Label Path:	http://www.w3.org/2000/01/rdf-schema#label
+ Node Styling	
+ Edge Styling	
+ Advanced	
MiningCompany	
MiningCompany	
Label Dath:	http://www.w3.org/2000/01/rdf-schema#label
Laber Path.	
+ Node Styling	
+ Node Styling + Edge Styling	

 Under the class that you want to apply an icon to, expand Node Styling. Then expand Default Node Style. For example, the image below shows the default node style configuration for the Mine type.

Label Path:	http://www.w3.org/2000/01/rdf-schema#label	
Node Styling		
Default Node Sty	yle	
Node Options		
Type: Shape 🛩	Shape: Circle 🗸	Preview
Background Color:	#8FD0DD Border Color: #FFFFFF	
Scale Factor: 1		
Label Options		Preview
Font Family: Robot	to v Bold	
Font Color: #00	0000 Font Size: 14	
Background Color:	#FFFFF	
Label Position: Bot	tom 🛩	
- Conditional Node	e Styles	

 Under Node Options, click the Type drop-down list and select Icon. An Icon field with a Browse button is displayed below Type:

Type: Icon	~	Shape:	Circle	~	
Icon:		Browse			

6. Click **Browse** to open the Icon Browser. The Browser lists the icons that have been uploaded to the application and gives you the option to upload new icons. For example, the image below shows that two icons have been uploaded previously and are in the library.

lco	n Browser							×
> 10	Sort by: Title Direction: Ascending							
CHS	(i) info	oem 🛞						
	1							
			1< <	>	×	1-2 of 2 Page	1of 1	¢ -
				CAN	CEL	UPLOAD NEW	USE SE	LECTED

To apply an existing icon, select the icon and click Use Selected. To add new icons, click
 Upload New. The Icon Uploader dialog box is displayed:

Icon Uploader	×	
Properties Title: Collection: Keywords: Type: Icon	Instances NEW Browse x h Upload	
	CANCEL OK	

- 8. In the Icon Uploader dialog box, complete the following fields as needed:
  - Title: Required field that specifies the name of the icon.
  - Instances: Required field that specifies the icon to upload. Only one icon can be selected at a time. Click Browse to navigate to the icon on your computer and select it. When an icon is selected, the width and height fields are populated to show the size of the selected image. You can adjust the size but it is not required. The images are automatically downsized when displayed in the View.
  - **Collection**: Optional field that specifies the collection to add the icon to. If you want to be able to filter icons by collection, you can select an existing collection from the drop-down list or click **New** and type a new name to add to the list.
  - **Keywords**: Optional field that specifies any keywords to tag the icon with. Keywords also enable users to filter icons.
  - **Type**: Required field that specifies the type of image you are adding to the library. **Icon** is the default value and should not be changed.

The image below shows an example of the properties set for an icon that will be used to represent the nodes for mines.

Properties		Instances NEW				
Title:	mine	mine.jpeg	Browse	1490	× 1609	
Collection:	New					Upload
Keywords:	Type: Icon					

9. When you have finished configuring the icon properties, click Upload to upload the icon. The dialog box shows a preview of the new icon. Click OK to close the dialog box and add the icon to the Icon Browser.

### Тір

It is helpful to repeat Steps 7 - 9 until you have uploaded all of the icons that you want to use for the dashboard. That way you can more efficiently apply the icons to the types in the Designer.

 When you have finished uploading icons to the Icon Browser, the next step is to apply them to the classes. First, select the icon that you want to apply to the class you selected in Step 4.
 For example, in the image below, the mine icon is selected for the Mine class.

lcon	Browser				×
> FIC	Sort by: Title _ Direction: Ascending				
TERS	info info	A mine	Ĩ	o mining company	
	oem	refinery		tier1	
		1< <	> >! 1-6of6	Page 1 of 1	¢ -
			CANCEL UPL	OAD NEW USE	SELECTED

11. Click **Use Selected** to apply the icon and return to the Designer where you can preview the icon and configure additional node options.

Node Options	
Type: Icon V Shape: Circle V	Preview
Icon: mine A 32 Browse	
Background Color: #8FD0DD Border Color: #FFFFFF	
Auto-trim Images	A
Scale Factor: 1	Preview

- 12. If you want to change the shape, size, and background or border color for the icon, you can adjust the following properties:
  - Shape: Enables you to set the icon shape to Circle or Square.
  - **Background Color**: Controls the color of the background if the icon is transparent or is smaller than the chosen shape. You can click the colored square to use the color picker or type the hex value for the desired color in the text box.
  - **Border Color**: Controls the color of the border around the icon. You can click the colored square to use the color picker or type the hex value for the desired color in the text box.
  - **Auto-Trim Images**: This setting trims the image to fit in the chosen shape if it is larger than the shape.
  - Scale Factor: Enables you to scale the icon size up or down.
- 13. When you have finished configuring the icon for the selected class, navigate to the other classes and repeat the steps above to apply and configure additional icons. When you have finished applying icons, click **Save** to apply the changes and close the Designer. The new icons will be applied to the nodes in each Network View for the dashboard. For example, the image below shows the revised View of the sample automobile supply chain knowledge graph.



#### Auto-Expanding a Network View

When you add nodes to a Network View, the Network Navigator Dashboard is configured by default to add only the selected nodes and paths to the View without automatically expanding the paths to show related nodes. You have the option to enable the **Auto-Expand** feature in the Dashboard Designer, however, so that the network is automatically expanded by the specified number of degrees any time a new node or path is added to the View. To configure a dashboard to auto-expand, follow the steps below.

- 1. Open in the Hi-Res Analytics application the Network Navigator Dashboard that you want to configure.
- In the main toolbar, click the **Designer** button. The Designer opens and the **Properties** by Type tab is displayed. Click the **Options** tab:

Designer				×
Properties by Ty	options			
Defaults	Path Finding	Expansion		
Hide system	and hidden resources in	visualization.		
			0.00051	0.01/5
			CANCEL	SAVE

3. Click the **Expansion** tab to view the expansion options. For example, the image below shows the default configuration.

Properties by Type Options
Defaults Path Finding Expansion
Auto expand 3 degrees by default.
Limit number of outgoing adjacent nodes to show to 4 nodes per node.
Limit number of incoming adjacent nodes to show to 4 nodes per node.
Show all connections:
Property Inclusion List
Use Inclusion List Select the properties to auto expand. You will still be able to manually expand properties that are not selected. Selecting none is the same as selecting all.
CANCEL SAVE

- 4. On the Expansion tab, configure the options as needed:
  - Number of Degrees to Expand: At the top of the screen, select the default number of degrees (hops) to auto-expand. The options are 0 6. When this option is set to a value greater than 0 and a node is added to a View in the dashboard, Anzo automatically populates the View with the nodes and paths that are *N* or fewer hops from the added node, where *N* is the specified value. If the View does not auto-expand when a node is added, that means there are no paths to or from that node that are *N* or fewer hops.
  - **Show All Connections**: This option controls whether or not all connections are shown for all nodes.
  - Show or Hide Inverse Edges: When auto-expanding the network multiple levels, this option controls whether to show or hide the inverse edges, i.e. whether or not to show both the incoming and outgoing connections to the same nodes. When this option is set to don't show, only one edge (path) will be shown between nodes. When this option is set to show, both the incoming and outgoing paths will be displayed.
  - Limit the Number of Adjacent Nodes: This option controls the number of adjacent nodes to show per node.
  - **Property Inclusion List**: If you would like auto-expansion to apply only to particular properties, you can select the **Use Inclusion List** checkbox. When Use Inclusion List is enabled, the dashboard's properties are listed at the bottom of the screen. Select each of the properties that you want to include.
- 5. When you have finished configuring the expansion options, click **Save** to apply the changes and close the Designer.

Once Auto-Expand is enabled, adding a node or path to a Network View in this dashboard will also automatically add to the View the related nodes and paths that meet the specified criteria.

Swapping the Graphmart in a Network Navigator Dashboard

Follow the instructions below if you want to change a Network Navigator dashboard to access a different graphmart.

- 1. Open in the Hi-Res Analytics application the Network Navigator Dashboard that you want to configure.
- 2. In the main toolbar, click the **Dashboard** menu and select **Properties**. The Properties dialog box for the dashboard is displayed. For example:

General	Security	Advanced				
General In	formation					
Explore Su	pply Chain			 	 	
Description:						
Explore	suppliers	; for raw m	aterials			
raphmart:*						
iraphmart:* Automobil	le Supply Cha	ain Risk				
Graphmart:* Automobil	le Supply Cha	ain Risk				
Graphmart:* Automobil	le Supply Cha	ain Risk				
Graphmart:* Automobil	le Supply Cha	ain Risk				

3. On the General tab, click the Graphmart drop-down list and select the graphmart that you want to use in place of the current selection.

General Security Advanced
General Information
Title:*
Explore Supply Chain
Description:
Explore suppliers for raw materials.
Graphmart:*
Automobile Supply Chain Risk
Q Search
Automobile Supply Chain Risk - no description -
Movies - no description -
Northwind - no description -
Tickets Graphmart     - no description -

4. Click **Save** to save the selection and close the Properties dialog box.

# Configuring a Dashboard to Update in Batch Reporting vs. Interactive Mode

Dashboards are configured to run in interactive mode by default. In interactive mode, Anzo updates dashboards automatically each time a user adds, selects, or modifies a filter or lens. Each change generates one or more SPARQL queries that are run against the graphmart data in AnzoGraph. Depending on conditions such as the number of concurrent users, the size of the data, the number of filters, and/or the complexity of the calculations, running in interactive mode can result in significant server workloads, as many complex queries are executed simultaneously. If dashboard performance becomes unsatisfactory, there are two options for changing how and when a dashboard is updated:

You can configure a dashboard to run in "Reporting" mode. Dashboards in reporting mode are
not refreshed immediately with every change. Instead, update queries are paused until a user
completes their changes and manually initiates the update process. The update refreshes all
of the lenses and filters on the dashboard.

 For finer-grained control over which lenses and filters are updated automatically and which are not, you can enable update controls on each lens and filter. Adding update controls does not change the update method for the entire dashboard, but certain components can be paused and then updated individually.

#### Tip

By default, dashboards are also automatically refreshed when the data in the backing graphmart changes. The **Automatic** setting in the Refresh menu controls this behavior. When Automatic is disabled and the graphmart is updated, the user is informed that the data changed and they can manually refresh the dashboard to get the updates.

This topic provides instructions for changing the update method at the dashboard level and enabling update controls so that lenses can be updated individually.

- Changing the Update Method for a Dashboard
- Adding Manual Update Controls to Lenses and Filters

Changing the Update Method for a Dashboard

- 1. Open in the Hi-Res Analytics application the dashboard that you want to change.
- 2. In the main toolbar, click **Designer** to open the dashboard designer. For example:

Dashboard designer Configure the properties of the dashboard	
Layout Hidden Lenses Toolbar Update Method	
Please select a region to configure	
Required region	
Configuring the center region	
Tab Container	
Tab position:	
Тор	
	CLOSE

3. In the designer, click the **Update Method** tab. The image below shows the Update Method tab for a dashboard that is in Interactive mode.



- 4. Click the radio button for the update mode that you want the dashboard to use. Then click **Close** to close the designer.
- 5. Save the dashboard to save the configuration change.

When a dashboard is in Reporting mode, an Update button is displayed on the top right corner of the dashboard. The button is disabled when the dashboard is up to date. The image below shows the disabled button.



When a user makes a change to a component, the Update button is enabled (as shown in the image below) and the dashboard can be refreshed as needed. When the update is initiated, all lenses are refreshed.



Adding Manual Update Controls to Lenses and Filters

- 1. Open in the Hi-Res Analytics application the dashboard that you want to change.
- 2. In the main toolbar, click the Refresh menu and select **Show Update Controls**. A pause icon and status icon are added to each filter and lens on the dashboard, as shown in the image below.



The pause icon (III) is used to pause or re-enable automatic updates of the component, and the status icon indicates whether the component is up-to-date (O) or out of sync (O).

- 3. Save the dashboard to save the configuration change.
- 4. In the dashboard, click the **pause** icon for any filters and lenses that you do not want to be updated automatically.

When a filter or lens is paused and a user makes an update that affects that component, an update message is displayed. For example, the image below shows a lens and filter that are out of sync. Clicking **UPDATE** in the message refreshes that component.



## **Capturing User-Defined Values in Dashboards**

You can configure a Table lens in a dashboard to allow users to input values and save those values to a linked dataset (LDS) in a graphmart. This topic provides instructions for configuring a graphmart, LDS, and dashboard to enable and save user-defined input. There are several steps required to configure the environment:

### Note

The procedures below must be performed by the sysadmin user or a user with Anzo Administrator permissions. Once the setup is complete, the components can be shared with other users and groups.

- 1. Create a Volume for Storing User Input
- 2. Configure the Graphmart
- 3. Add User Input Properties to the Model
- 4. Create and Configure the Dashboard
- 5. Share the LDS and Add it to the Graphmart

#### Create a Volume for Storing User Input

When you configure a Table lens to save user-entered values to an LDS, you must choose the volume in which to save the LDS. Cambridge Semantics strongly recommends that you do not write to the system journal (<install\_path>/Server/data/journal/anzo.jnl by default). Instead, use a separate volume that is dedicated to storing user-entered values. If necessary, create a new volume before proceeding. See <u>Creating a New Volume</u> in the Administration Guide for instructions.

#### Note

Cambridge Semantics also strongly recommends that you create a separate volume for each graphmart that will include user-defined values.

#### **Configure the Graphmart**

The first step is to deactivate the graphmart that will contain the user-entered values and disable manual refresh. When manual refresh is disabled, the graphmart is refreshed automatically any time a user updates the LDS by adding values in the dashboard. Follow the steps below to configure the graphmart.

- 1. In the graphmart for which you want to add user-defined text, go to the Overview tab.
- From the Overview tab, deactivate the graphmart if it is online, and then find the Manual Refresh Graphmart setting at the bottom of the screen:

**	Good Reads	Books			Save a Copy	Save a Copy Create Dashboard		
	Overview	Explore	Datasets	Data Layers	Dashboards	Data on Demand	>	
Description None Load Priorit 100 Metrics	у				Las La Structu	General Type: Graphmart Creator: System Admini t Accessed: N/A ist Updated: N/A re Modified: 37 minutes ago Released: 43 minutes ago	strat	
Profile data	oad Counts				http://car	mbridgesemantics.com/Gra	Ō	
Disable pop	ulating counts during load	d operations 🛛 💻			Dynamic	Resource Log		
Data Loa Load layers	ding Settings that do not fail				Inactivity	Deactivate Timeout		
Ignore So Ignore any s	ource Errors itep or view's source layer	r if they failed to load 🛛 🗕			Tags			
Leave Gr When refres	aphmart Online Du	ring Refresh nart and layer online.			Auto-Gen	1		
Manual F Once loaded	Refresh Graphmart d, changes only pushed to	AnzoGraph manually.(Only affe	cts Journal Based Data) 🔹	•				
Generate	Statistics After Ea	ch Layer						

- 3. If Manual Refresh Graphmart is enabled, slide the slider to the left to disable it.
- 4. Activate the graphmart and proceed to the next step to update the model.

Add User Input Properties to the Model

The next step is to update the model that was generated when the data was onboarded. Any new properties that you want to add as columns for text entry in a Table lens need to be added to the model. Follow the steps below to update the Direct Load Step query to add new properties to the model.

1. Click the **Data Layers** tab. Expand the layer that was generated when you created the graphmart and find the Direct Load Step that inserts the data.

#### Tip

You might want to start another session of Anzo so that you can view the layer's model in the Model editor and the graphmart query at the same time.

2. Open the Direct Load Step for editing and click the **Query** tab. For example, the image below shows the query that was generated to onboard data about books from a CSV file:


3. Make sure the following prefixes are declared in the PREFIX clause at the top of the query. If any are missing, add them to the query:

```
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
```

4. Next, add the following GRAPH clause statements under the SERVICE clause:

```
GRAPH <model_URI> {
    <property_URI> a owl:DatatypeProperty;
    rdfs:comment "comment";
    rdfs:label "property_label";
    rdfs:domain <class_URI>;
    rdfs:range <dataype>.
[ <property2_URI> a owl:DatatypeProperty;
    rdfs:label "property2_label";
    rdfs:label "property2_label";
    rdfs:domain <class_URI>;
    rdfs:range <datatype>.
]
```

```
[ ... ]
}
```

Value	Data T ype	Description
mode I_URI	URI	The URI of the model that was generated by the Direct Load Step. You can copy the URI from the Model editor from the Details tab for the model. For example, <http: 698="" cambridgesemantics.com="" layer="" mode<br="">l&gt;.</http:>
proper ty_URI	URI	The URI to use for the new property. You can look at the URIs for other properties in the model and follow the same scheme. For example, <http: 698="" cambridgesemantics.com="" layer="" model<br="">#GoodReads100kBooks.Notes&gt;.</http:>
comm ent	string	A comment that describes the property. For example, "Notes from users".
proper ty_ label	string	The label to give the property. For example, "Notes".
class_ URI	URI	The URI of the class that the property should be added to. You can copy the URI from the Model editor from the Details tab for the class. For example, <http: 698="" cambridgesemantics.com="" layer="" model<br="">#GoodReads100kBooks&gt;.</http:>
datayp e	URI	The datatype of the new property. For example, xsd:string.

For example, the following query adds one new property to the model that is generated by the Direct Load Step query shown above:

```
PREFIX good reads books files:
<http://cambridgesemantics.com/DataSource/0e5ab8601249264a318b4a4bcfa81700/Goo
d Reads Books/>
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
INSERT {
 GRAPH ${targetGraph} {
   ?subject ?predicate ?object .
 }
}
WHERE {
 SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
   GRAPH
<http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Model> {
<http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Model#Goo</pre>
dReads100kBooks.Notes> a owl:DatatypeProperty;
      rdfs:comment "Notes from users" ;
      rdfs:label "Notes" ;
      rdfs:domain
<http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Model#Goo</pre>
dReads100kBooks> ;
     rdfs:range xsd:string .
 }
   []
      s:segment "false"^^xsd:boolean .
    ?generator a s:RdfGenerator, s:OntologyGenerator ;
      s:as (?subject ?predicate ?object) ;
      s:base ${targetGraph} .
   good reads books files:GoodReads 100k books a s:FileSource ;
      s:url "/nfs/data/csv/GoodReads 100k books.csv" ;
      s:model "GoodReads 100k books" .
 }
}
```

- 5. Next, **if the class for any of the new properties does not have a primary key defined**, **you must create a key for that class**. If all of the classes you referenced in the query have primary keys, you can continue to the next step. If one or more of the classes do not have primary keys, follow the instructions below:
  - a. Locate in the query the statement block for each class that needs a key definition. For example, in the query above, there is only one class, s:model "GoodReads\_100k\_Books". If you have multiple classes, the query has several blocks, such as this example:

```
emrdbsmall:emr_complaint a s:DbSource ;
s:using mysql_db:MySQL_DB ;
s:table "emrdbsmall.emr_complaint" ;
s:model "emr_complaint" .
emrdbsmall:emr_patient a s:DbSource ;
s:using mysql_db:MySQL_DB ;
s:table "emrdbsmall.emr_patient" ;
s:model "emr_patient" .
emrdbsmall:emr_complaintdescription a s:DbSource ;
s:using mysql_db:MySQL_DB ;
s:table "emrdbsmall.emr_complaintdescription" ;
s:model "emr_complaintdescription" .
...
```

- b. At the end of the block for the class you want to add a key to, change the period (.) after s:model to a semicolon (;).
- c. Next, add the following line below s:model:

```
s:key ("key_property" [, "key_property2" ] [, ... ]) .
```

Where  $key\_property$  is the label of the property to use as a key for the class. The property that you choose must have unique values. If there is not a property in the class with unique values, you can specify a combination of properties that would create a

unique value. Make sure that the value of key\_property matches the label for that property in the model. For example, for the query in step 4, the Isbn property can be used as a unique key for the GoodReads\_100k\_Books class:

```
good_reads_books_files:GoodReads_100k_books a s:FileSource ;
  s:url "/nfs/data/csv/GoodReads_100k_books.csv" ;
  s:model "GoodReads_100k_books" ;
  s:key ("Isbn") .
```

The final, completed query is shown below:

```
PREFIX good reads books files:
<http://cambridgesemantics.com/DataSource/0e5ab8601249264a318b4a4bcfa81700
/Good Reads Books/>
PREFIX s: <http://cambridgesemantics.com/ontologies/DataToolkit#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
INSERT {
 GRAPH ${targetGraph} {
   ?subject ?predicate ?object .
 }
}
WHERE {
 SERVICE <http://cambridgesemantics.com/services/DataToolkit> {
   GRAPH
<http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Mode
1> {
<http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Mode
l#GoodReads100kBooks.Notes> a owl:DatatypeProperty;
      rdfs:comment "Notes from users" ;
     rdfs:label "Notes" ;
      rdfs:domain
<http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Mode
l#GoodReads100kBooks> ;
     rdfs:range xsd:string .
 }
    []
      s:segment "false"^^xsd:boolean .
    ?generator a s:RdfGenerator, s:OntologyGenerator ;
```

```
s:as (?subject ?predicate ?object) ;
s:base ${targetGraph} .

good_reads_books_files:GoodReads_100k_books a s:FileSource ;
s:url "/nfs/data/csv/GoodReads_100k_books.csv" ;
s:model "GoodReads_100k_books" ;
s:key ("Isbn") .
}
```

6. Save the step and then refresh or reload the graphmart to update the model with the new properties.

**Create and Configure the Dashboard** 

Now that the graphmart is updated, the next step is to create and configure the dashboard that will enable input for the new properties.

- 1. Create a dashboard for the graphmart that was updated in the previous task.
- 2. Add a Table lens to the dashboard. In the lens, add any columns that you would like to see, including columns for the user input properties that you added to the model.
- 3. In the lens Designer, select a column that will allow input and click the **Editor** link under **Column Value Expression** (shown in the image below).

Designer	×
Columns + 🗊 🦑 🦄 List of column configurations.	Column Details Set properties of the selected column.
Img	Column Header Label:
1 Genre	Column Value Expression:
Desc	Notes
1 Author	Column Footer Expression:
1 Link	Click to edit Format   Hyperlink
III Pages	Column Width:
Rating	
Reviews	Content Filters
Notes	Group route but
Add a column   Display Properties	Click to edit
Preview changes	CANCEL SAVE

The Column Editor Configuration dialog box is displayed and populated with the name of the chosen property:

Column Editor Configuration	×
Editable	
1. Select Linked Dataset	
2. Graph to use for newly created data Match entity URI	
3. Select property to edit	
➡ Notes	
4. Set editor properties	
	CLOSE

4. In the dialog box, select the **Editable** checkbox. Then click the **Select Linked Dataset** field and select **Create a linked dataset**. The New Linked Data Set dialog box is displayed:

New Linked Data Set	×
Title:*	
Description:	
[+] more	
CANCEL	OK

5. Specify a title for the new LDS and add an optional description. Then click the **[+] more** link to expose the following additional settings.

New Linked Data Set	×
Title:*	
Book Notes	
Description:	
user-defined text about the books	
[-] less	
Datasource:*	
System Datasource	
Storage:*	
http://openanzo.org/creation/9cb6ba22-6747-432e-8d3f-e1cd5c7ae44c	
Ontology:*	add value
http://openanzo.org/creation/9cb6ba22-6747-432e-8d3f-e1cd5c7ae44c_ontolo	ogy
	add value
CANCEL	ок

- 6. Complete the required fields:
  - **Datasource**: Click this field and select the volume to save the LDS in. This is the volume that you created in Create a Volume for Storing User Input.

- **Storage**: This field lists the system-generated URI for the LDS and dataset storage graph. Cambridge Semantics recommends that you do not change this value.
- Ontology: Replace this value with the URI of the model that contains the writable properties. For the example shown above, the value is http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1 f611/Model.

The image below shows the completed screen:

New Linked Data Set	×
Title:*	
Book Notes	
Description:	
user-defined text about the books	
[-] less	
Datasource:*	
User-defined Text	
Storage:*	
http://openanzo.org/creation/9cb6ba22-6747-432e-8d3f-e1cd5c7ae44c	]
add valu	Je Je
http://cambridgesemantics.com/Layer/698d179498f945a393f8a91ee9f1f611/Model	
add valu	Je
CANCEL	ж

7. Click OK to save the LDS configuration and create the LDS. The LDS is added to the Datasets catalog and it is shown on the Column Editor Configuration screen. Additional editor properties are made available on the screen. For example:

Co	olumn Editor	Configuration	×
Edita 1. Se Bo 2. G	able 🔽 elect Linked Datas ook Notes raph to use for ne	set wly created data	
Ma	atch entity URI		
3. S(	elect property to e Notes	dit	
4. Ses	et editor propertie arch	S	0
-	General		~
	control	Single Line Text	•
	max length		
	place holder		
			CLOSE

- 8. Under **Set editor properties**, you have the option to modify the constraints for the editable property.
- 9. When you are finished setting any constraints, click Close and then click Save in the Designer. The lens is now configured to display the editable field. Hovering over a row in the table shows an edit icon (✓) in the field. For example:

🗘 Designer 🛛 🛃 Pro	Designer 🛃 Properties 🔁 Export 🏢 Manage Columns							
Table ×								
Title 🛧	Img	Genre	Desc	Author	Pages	Rating	Notes	
1984, Spring: A Choice of Futures		Nonfiction,Scien		Arthur C. Clarke	254	3.68	1	
200 Years to Christmas/Rebel of the Red Planet			For almost two centuries the huge spaceship had speared its way through the stars, bound for another	J.T. McIntosh,James Murdoch MacGregor,Charl L. Fontenay	81	3.75		

- 10. Repeat steps 3 9 for any additional properties that you added to the model and want to make editable.
- 11. To add text in the lens, double-click an editable field to open the text editor, as shown in the image below.

Table ×						
						Notes
1984, Spring: A Choice of Futures	Nonfiction,Scien		Arthur C. Clarke	254	3.68	1
200 Years to Christmas/Rebel of the Red Planet		For almost two centuries the huge spaceship had speared its way through the stars, bound for another 200 hundre	J.T. McIntosh,James Murdoch MacGregor,Charl L. Fontenay		CA	add value

Add text in the provided field, and then click **OK** to add the text to the dashboard.



The text is displayed in the table:

ſ	Table $\times$								
	Title		Img	Genre	Desc	Author	Pages	Rating	Notes
	1984, Sp A Choice Futures	ring: e of		Nonfiction,Scien		Arthur C. Clarke	254	3.68	Sample text about this book.
	200 Year Christma of the Re Planet	rs to as/Rebel ed	200		For almost two centuries the huge spaceship had speared its way through the stars, bound	J.T. McIntosh,James Murdoch MacGregor,Charl L. Fontenay	81	3.75	

You can edit the text by double-clicking the field. You can also add another text entry by opening the text editor and clicking **add value**. Click **OK** after making changes or adding values.

12. Finally, save the dashboard and proceed to the next task.

### Share the LDS and Add it to the Graphmart

In order for other users to be able to add values to the dashboard, they need to have permission to modify the LDS. Follow the steps below to share the LDS and then add it to the graphmart.

1. In the Datasets catalog, open the LDS that is associated with the dashboard you configured in the previous task. Then click the **Sharing** tab to configure the permissions.

¥.	Book Notes				+ Add to Cart	Create Graphmart	Д
	Overview	Explore	Graphmarts	Dashboards	Versions	Discussion	>
Inheritance	•						
Inherit pern	nissions from			Pass permissions to			
Inherit p	ermissions from		~	Pass permissions to	)		$ $ $\vee$
Q Sea	rch users, roles or	groups					
💄 Syst	em Administrator	Admin 📋		6	•		
🚢 Ever	yone	View 📋		4			
				Define the permission	ons of this artifact.		
				Start by selecting a use	er or group on the left.		

- 2. On the Sharing tab, share the dataset with the appropriate users and/or groups. For information about configuring permissions, see Share Access to Artifacts.
- 3. Next go to the graphmart and click the **Datasets** tab.
- On the Datasets tab, click Add Dataset. Then select the dataset that you shared and click Add to add the dataset to the graphmart. For more information about adding a dataset to a graphmart, see Adding a Dataset to a Graphmart.
- 5. Refresh the graphmart to load the dataset to AnzoGraph.

The graphmart, linked dataset, and dashboard are now configured to allow and save user-defined values.

### Working with Lenses

The topics in this section provide instructions on creating, exporting, and deleting lenses.

- Creating a Lens
- Cloning a Lens
- Exporting a Lens
- Deleting a Lens

### **Creating a Lens**

Lenses define the data's visual presentation. Each type of lens represents a unique method for displaying data. For instance, in a column chart, you can present multiple data series for comparison. You can also apply custom formats such as fonts and colors to any lens. The topics in this section provide instructions for creating lenses.

- Creating a Chart Lens
- Creating a Drill Down Lens
- Creating a Form Lens
- Creating a List Lens
- Creating a Table Lens
- Advanced Lenses

#### **Creating a Chart Lens**

Anzo Hi-Res Analytics employs the <u>Highcharts</u> API to provide interactive chart lenses. The topics in this section provide information about creating each type of chart.

- Create an Area Chart
- Create a Bar Chart
- Create a Bubble Chart
- Create a Column Chart

- Create a Funnel Chart
- Create a Heat Map
- Create a Line Chart
- Create a Polar Chart
- Create a Scatter Chart

### Create an Area Chart

Area charts are useful for emphasizing trends in your data. They are similar to line charts but include options for displaying stacked data series. This topic provides instructions for creating an area chart with minimal configuration. Descriptions of all of the available configuration options are included below the steps.

- Complete the Minimum Configuration
- Optional Configuration Settings

# **Complete the Minimum Configuration**

Follow the instructions below to create an area chart. The instructions guide you through completing the minimum configuration needed to display your data in the chart. Additional, optional configuration settings are described in Optional Configuration Settings.

 In the dashboard that you want to add a lens to, click Lenses in the main toolbar and select New. The Create Lens dialog box is displayed:

Cre	ate Lens		×
Selec	t the type of lens to create		
۹.	Search for lenses		
+ Path Result	to lens values s: 10		
	Table Tabular visualization of data Properties		
₽/	Form View and edit resource attributes. Properties		
սհ	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

2. On the Create Lens dialog box, select **Chart**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens		×
General Information		
Title.*		
Description:		
(+) more		

3. Type a **Title** and optional **Description** for the lens.

4. Click Finish. The lens Designer dialog box is displayed:

Designer	×
Chart Configuration Please select options below to begin configuring your chart	
Chart Type Column	>
Chart Data Series not specified	>
Chart Theme Chart theme selection	>
Chart Details Title: unspecified, Subtitle: unspecified, Legend is enabled, Tooltip is enabled	>
O Preview changes CANCEL PREVIEW	SAVE

5. Click Chart Type to open the Chart Types dialog box:

Designer			×
Chart Types Please select a chart ty	rpe		<del>©</del> 🖧
Column Bar Line Gubble Funnel Heat Map Pie Polar Scatter	Column Clustered Column Bar Clustered Bar Clustered Bar	1 100% Stacked Column	terfat
Area     Preview changes	Lina	CANCEL PREVIEW	SAVE

- 6. On the left side of the screen, select **Area**. On the right side of the screen, select the icon for the type of area chart that you want to create. There are five types:
  - Area: Connects value points on the chart with straight lines and shades the area below the lines.
  - **Step Area**: Connects value points on the chart with short horizontal steps and shades below the lines. This chart emphasizes the extent of value change by expanding the data points across the X axis.

- Area Spline: Connects value points on the chart with curved lines and shades the area below the lines.
- **Stacked Area**: Connects value points on the chart with straight lines and shades the area below the lines. Includes the option to configure multiple groups within a series to distinguish between groups of values inside the total value.
- 100% Stacked Area: Compares each value as a percentage of the total and shades the area below each series. Includes the option to configure multiple groups within a series to distinguish between groups of values inside the total value.

Once the type is selected, you are returned to the Chart Configuration screen, and the Chart Type value is set to the type that you chose.



7. Next, select **Chart Data** to add the data that will populate the chart.

The **Series Configurations** section of the screen contains settings to manage each series. One series is created by default. A series is a set of data to display on the chart, for example a line on a line or area chart or one set of columns on a column or bar chart. The details for the selected series appear in **Series Details**, which contains settings to define the details of the selected series, such as the data to display as well as formatting and labels. From Series Configurations, you can clone, rename, disable (remove the series from the chart without deleting it), or remove a series. You can add a series by clicking the plus icon (+) or delete all series by clicking the trashcan icon (**I**). Removing a series cannot be undone.

- 8. Under Series Configurations, click **Rename** under Series 1 and give the series a meaningful name that describes the data that will populate the chart.
- 9. Under Series Details, click the Group field and select the property to use for grouping the data in the series. These are the values for the X axis. You can use functions to derive the values for the groups. The list below describes the icons and options that are available when choosing a property:
  - The Root Value (A) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IF) or back (IF) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.

- 10. Once you select the Group, the **Group Label** field is populated with the same value. This setting configures the property whose values should serve as the group label. You can edit the value if necessary.
- 11. Next, click the **Value** field and select the property to use for the Y axis values. You can include functions to derive the values.
- 12. If you want to add a group for the series, you can click **Show formulas used for creating calculated series** and add values for the following fields:
  - Series Group: Specifies a property to use for grouping data in addition to the Group value.
  - Series Label: Specifies the property whose values should serve as the Series Group label. It is typically set to the same value as Series Group.
  - Series Sort Direction: Specifies the sort direction for the series groups.
- 13. You can click Save any time to render on the dashboard the data that you have configured so far. You can add another series and repeat the steps above. You can also configure several more options for the chart by changing the optional parameters. The additional tabs and options are described below in Optional Configuration Settings.

# **Optional Configuration Settings**

This section describes each of the tabs and additional settings that are available for customizing an area chart.

- Chart Data: Plot Tab
- Chart Data: X Axis Tab
- Chart Data: Y Axis Tab
- Chart Data: Filters Tab
- Chart Theme
- Chart Details

# **Chart Data: Plot Tab**

This tab contains settings that control series formatting, such as data labels, legends, and other

#### display options.

Series Details Set properties of the selected series configuration	
Series Chart Type Stacked Area	>
Series Chart Style Plot style information	>
Series Chart Data Labels Data labels are set to automatic enablement	>
Series Chart Markers Markers are enabled	>
Show:	
Show in legend:	
Connect missing points:	
Threshold:	
0	
Data Plot X Axis Y Axis Filters	

Field	Description
Series Chart Type	This option opens the Chart Types selection screen where you can select a different chart type for the selected series. The selection must be compatible with the types selected for any other series in the chart, however, or the incompatible series will fail validation and not be displayed on the chart until it is corrected.
Series Chart Style	This option contains the plot style settings. These settings control the thickness and color of lines for the series as well as the color and opacity of the fill area below the lines.
Series Chart Data Labels	This option contains the data label settings, which control the font, alignment, and font effects of the labels for the series.
Series Chart	This option contains the data marker settings. These settings control the

Field	Description
Markers	placement and format of the data points for a series that appear when a user hovers the pointer over a point in the line. For example, the image below shows a data marker in an area chart:
Show	This setting can be used to limit the data included in the series. The limit can be defined by the largest or smallest Group Label or Value.
Show in legend	This setting controls whether to show the name of the series in the legend.
Connect missing points	Selecting this option connects the series line across any missing points. For example, if this option was enabled for the chart below, the two dark blue areas would be connected.
Threshold	Defines the Y axis value to use as a base (starting point) for the shaded area. For example, a threshold of 0 begins all shading at the value 0. A threshold of 10 begins the shading at 10 and draws the area chart above or below the threshold as required. The image below shows an area chart with a threshold of 30.

Field	Description

## **Chart Data: X Axis Tab**

This tab defines the formats and labels for the X axis values in the series.

Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Talo	
The.	
The second	
Type:	
Linear	
Display values as evenly distributed categories	
Display axis on the opposite side	
Axis Title Details	>
Axis title is Year Of Release	
Axis Labels	>
Axis labels are enabled	
Axis Style	>
Avia etula information	

Field	Description
Axis	The Axis and Title settings are populated with the values from the Group and Group Label values from the Data tab. If multiple axes exist, you can

Field	Description
	selected a different value to use for the X axis.
Title	This setting defines the title for the X axis.
Туре	This setting controls the scale for the X axis, linear or logarithmic.
Display values as evenly distributed categories	This setting controls whether to evenly distribute the X axis values.
Display axis on the opposite side	This setting can be enabled to move the X axis to the opposite side of the chart.
Axis Title Details	This option contains the X axis title settings, which control the font, alignment, and font effects of the axis title.
Axis Labels	This option contains the X axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option contains the X axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

### Chart Data: Y Axis Tab

This tab defines the formats and labels for the Y axis values in the series.

Series Details Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Title:	_
Type:	
Linear	_
☑ Display axis on the opposite side	
Axis Title Details Axis title is Rating	>
Axis Labels Axis labels are enabled	>
Axis Style Axis style information	>
Data Plot X Axis Y Axis Filters	

Field	Description
Axis	The Axis and Title settings are populated with the Value from the Data tab. If multiple axes exist, you can selected a different value to use for the Y axis.
Title	This setting defines the title for the Y axis.
Туре	This setting controls the scale for the Y axis, linear or logarithmic.
Display axis on the opposite side	This setting can be enabled to move the Y axis to the opposite side of the chart.
Axis Title Details	This option contains the Y axis title settings, which control the font, alignment, and font effects of the axis title.
Axis Labels	This option contains the Y axis label settings, which control the font, alignment, and font effects of the labels for the axis.

Field	Description
Axis Style	This option contains the Y axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

# **Chart Data: Filters Tab**

This tab defines any filters to apply to the series.

Series Details Set properties of the selected series configuration	
Series Filters No filters specified	>
Group Filters No filters specified	
Value Filters No filters specified	
Data Plot X Axis Y Axis Filters	

Field	Description
Series Filters	This option can be used to define filters that apply to the entire series.
Group Filters	This option can be used to define filters that apply only to the Group values for the series.
Value Filters	This option can be used to define filters that apply only to the <code>Value</code> values for the series.

### **Chart Theme**

This setting presents options for configuring the theme or color scheme for the chart.



# **Chart Details**

This option offers finer-grained customization settings than the Chart Theme. You can further customize the chart design by adding details such as a chart title and subtitle. You can also modify chart-level styles and fonts as well as legend and tooltip formats.

Chart Details Please specify details about your chart	3
Title:	
Subtitle:	
Chart Style Chart style information	>
Title Details Title: unspecified	>
Subtitle Details Subtitle: unspecified	>
Legend Details Legend is enabled	>
Tooltip Details Tooltip is enabled	>

#### **Create a Bar Chart**

Bar charts are useful for comparing different sets of data or emphasizing drastic changes in a data set over time. This topic provides instructions for creating a bar chart with minimal configuration. Descriptions of all of the available configuration options are included below the steps.

- Complete the Minimum Configuration
- Optional Configuration Settings

# **Complete the Minimum Configuration**

Follow the instructions below to create a bar chart. The instructions guide you through completing the minimum configuration needed to display your data in the chart. Additional, optional configuration settings are described in Optional Configuration Settings.

 In the dashboard that you want to add a lens to, click Lenses in the main toolbar and select New. The Create Lens dialog box is displayed:



2. On the Create Lens dialog box, select **Chart**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens				×
General Information				
Title.*				
Description:				
(+) more				
	CAL	NCEL PRE	IVIOUS	FINISH

3. Type a Title and optional Description for the lens.

4. Click Finish. The lens Designer dialog box is displayed:

Designer	×
Chart Configuration Please select options below to begin configuring your chart	
Chart Type Column	>
Chart Data Series not specified	>
Chart Theme Chart theme selection	>
Chart Details Title: unspecified, Subtitle: unspecified, Legend is enabled, Tooltip is enabled	>
Preview changes     CANCEL     PREVIEW	SAVE

5. Click Chart Type to open the Chart Types dialog box:

Designer		×
Chart Types Please select a chart ty	rpe	G &
Column Bar Line Bubble Funnel Heat Map Pie Polar Scatter	Column	Image: Stacked Column     Image: Stacked Column       Image: Stacked Education     Image: Stacked Education
Area	Line	CANCEL PREVIEW SAVE

- 6. On the left side of the screen, select **Bar**. On the right side of the screen, select the icon for the type of bar chart that you want to create. There are three types:
  - Clustered Bar: Compares values across categories.
  - **Stacked Bar**: Compares the contribution of each value to a total across categories. Includes the option to configure multiple groups within a series to distinguish between groups of values inside the total value.

100% Stacked Bar: Compares each value as a percentage of the total. Includes the
option to configure multiple groups within a series to distinguish between groups of
values inside the total value.

Once the type is selected, you are returned to the Chart Configuration screen, and the Chart Type value is set to the type that you chose.

7. Next, select Chart Data to add the data that will populate the chart.

Designer		×
Chart Data Please specify what you would like to se	e in the chart	0 🖨
Series Configurations + ins Select a configuration to edit	Series Details Set properties of the selected series configuration *Group:	
Clone   Rename   Disable   Remove	Click to set the formula used to group by Group label:	0
	Click to set formula for the Group label value *Value:	0
	Click to set formula for the Value value + Show formulas used for creating calculated series	0
+ Add a series	Data Plot Y Axis X Axis Filters	
O Preview changes	CANCEL PREVIEW	SAVE

The **Series Configurations** section of the screen contains settings to manage each series. One series is created by default. A series is a set of data to display on the chart, for example a line on a line or area chart or one set of columns on a column or bar chart. The details for the selected series appear in **Series Details**, which contains settings to define the details of the selected series, such as the data to display as well as formatting and labels. From Series Configurations, you can clone, rename, disable (remove the series from the chart without deleting it), or remove a series. You can add a series by clicking the plus icon (+) or delete all series by clicking the trashcan icon (**I:**). Removing a series cannot be undone.

- 8. Under Series Configurations, click **Rename** under Series 1 and give the series a meaningful name that describes the data that will populate the chart.
- 9. Under Series Details, click the **Group** field and select the property to use for grouping the data in the series. These are the values for the X axis. You can use functions to derive the values for the groups. The list below describes the icons and options that are available when

choosing a property:

 The Root Value (A) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IP) or back (IP) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 10. Once you select the Group, the **Group Label** field is populated with the same value. This setting configures the property whose values should serve as the group label. You can edit the value if necessary.
- 11. Next, click the **Value** field and select the property to use for the Y axis values. You can include functions to derive the values.
- 12. If you want to add a group for the series, you can click **Show formulas used for creating calculated series** and add values for the following fields:

- Series Group: Specifies a property to use for grouping data in addition to the Group value.
- Series Label: Specifies the property whose values should serve as the Series Group label. It is typically set to the same value as Series Group.
- Series Sort Direction: Specifies the sort direction for the series groups.
- 13. You can click Save any time to render on the dashboard the data that you have configured so far. You can add another series and repeat the steps above. You can also configure several more options for the chart by changing the optional parameters. The additional tabs and options are described below in Optional Configuration Settings.

# **Optional Configuration Settings**

This section describes each of the tabs and additional settings that are available for customizing a bar chart.

- Chart Data: Plot Tab
- Chart Data: Y Axis Tab
- Chart Data: X Axis Tab
- Chart Data: Filters Tab
- Chart Theme
- Chart Details

### **Chart Data: Plot Tab**

This tab contains settings that control series formatting, such as data labels, legends, and other display options.

Series Details	
Set properties of the selected series configuration	
Series Chart Type Clustered Bar	>
Series Chart Style Plot style information	>
Series Chart Data Labels Data labels are set to automatic enablement	>
Show:	
Largest 70 -Automatic-	
Show in legend:	
Data Plot Y Axis X Axis Filters	

Field	Description
Series Chart Type	This option opens the Chart Types selection screen where you can select a different chart type for the selected series. The selection must be compatible with the types selected for any other series in the chart, however, or the incompatible series will fail validation and not be displayed on the chart until it is corrected.
Series Chart Style	This option contains the plot style settings, which control the thickness and color of the bars for the series.
Series Chart Data Labels	This option contains the data label settings, which control the font, alignment, and font effects of the labels for the series.
Show	This setting can be used to limit the data included in the series. The limit can be defined by the largest or smallest Group Label or Value.
Show in legend	This setting controls whether to show the name of the series in the legend.

# Chart Data: Y Axis Tab

This tab defines the formats and labels for the Y axis values in the series.

Series Details Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Title:	
Display axis on the opposite side	
Axis Title Details Axis title is Title	>
Axis Labels Axis labels are enabled	>
Axis Style Axis style information	>
Data Plot YAxis XAxis Filters	

Field	Description
Axis	The Axis and Title settings are populated with the Value from the Data tab. If multiple axes exist, you can selected a different value to use for the Y axis.
Title	This setting defines the title for the Y axis.
Туре	This setting controls the scale for the Y axis, linear or logarithmic.
Display axis on the opposite side	This setting can be enabled to move the Y axis to the opposite side of the chart.
Axis Title Details	This option opens the Y axis title settings, which control the font, alignment, and font effects of the axis title.

Field	Description
Axis Labels	This option opens the Y axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option contains the Y axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

# Chart Data: X Axis Tab

This tab defines the formats and labels for the X axis values in the series.

Series Details Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Title:	
Туре:	
Linear	
Display values as evenly distributed categories	
Display axis on the opposite side	
Axis Title Details	5
Axis title is Year Of Release	
Axis Labels Axis labels are enabled	>
Axis Style Axis style information	>
Data Plot Y Axis X Axis Filters	

Field	Description
Axis	The Axis and Title settings are populated with the values from the Group and Group Label values from the Data tab. If multiple axes exist, you can selected a different value to use for the X axis.

Field	Description
Title	This setting defines the title for the X axis.
Туре	This setting controls the scale for the X axis, linear or logarithmic.
Display values as evenly distributed categories	This setting controls whether to evenly distribute X axis values.
Display axis on the opposite side	This setting can be enabled to move the X axis to the opposite side of the chart.
Axis Title Details	This option accesses the X axis title settings, which control the font, alignment, and font effects of the axis title.
Axis Labels	This option accesses the X axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option accesses the X axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

### **Chart Data: Filters Tab**

This tab defines any filters to apply to the series.

Series Details Set properties of the selected series configuration		
Series Filters No filters specified	>	
Group Filters No filters specified	>	
Value Filters No filters specified		
Data Plot Y Axis X Axis Filters		

Field	Description
Series Filters	This option can be used to define filters that apply to the entire series.
Group Filters	This option can be used to define filters that apply only to the Group values for the series.
Value Filters	This option can be used to define filters that apply only to the Value values for the series.

## **Chart Theme**

This setting presents options for configuring the theme or color scheme for the chart.


## **Chart Details**

This option offers finer-grained customization settings than the Chart Theme. You can further customize the chart design by adding details such as a chart title and subtitle. You can also modify chart-level styles and fonts as well as legend and tooltip formats.

Chart Details Please specify details about your chart	G 🖨
Title:	
Subtitle:	
Chart Style Chart style information	>
Title Details Title: unspecified	>
Subtitle Details Subtitle: unspecified	>
Legend Details Legend is enabled	>
Tooltip Details Tooltip is enabled	>

#### **Create a Bubble Chart**

A bubble chart is similar to a scatter chart but has a third dimension. In addition to the X and Y axes, there is a size axis that determines the size of the bubbles in the chart. This topic provides instructions for creating a bubble chart with minimal configuration. Descriptions of all of the available configuration options are included below the steps.

- Complete the Minimum Configuration
- Optional Configuration Settings

## **Complete the Minimum Configuration**

Follow the instructions below to create a bubble chart. The instructions guide you through completing the minimum configuration needed to display your data in the chart. Additional, optional configuration settings are described in Optional Configuration Settings.

1. In the dashboard that you want to add a lens to, click **Lenses** in the main toolbar and select

New. The Create Lens dialog box is displayed:

Crea	ate Lens t the type of lens to create		×
٩.5	Search for lenses		
+ Path Results	to lens values s: 10		
	Table Tabular visualization of data Properties		
₽⁄	Form View and edit resource attributes. Properties		
dı	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

2. On the Create Lens dialog box, select **Chart**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens		×
General Information		
Title.*		
Description:		
(+) more		

3. Type a **Title** and optional **Description** for the lens.

4. Click Finish. The lens Designer dialog box is displayed:

Designer	×
Chart Configuration Please select options below to begin configuring your chart	
Chart Type Column	>
Chart Data Series not specified	>
Chart Theme Chart theme selection	>
Chart Details Title: unspecified, Subtitle: unspecified, Legend is enabled, Tooltip is enabled	>
O Preview changes CANCEL PREVIEW	I SAVE

5. Click Chart Type to open the Chart Types dialog box:

Designer			>
Chart Types Nease select a chart t	/pe		<b>G</b> (
Column C	Column	100% Stacked Colum	winefat
<ul> <li>Pie</li> <li>Polar</li> <li>Scatter</li> <li>Area</li> </ul>	Custered Bar Stacked Bar	100% Stacked Bar	
Preview changes		CANCEL	PREVIEW SAVE

6. On the left side of the screen, select **Bubble**, and then select the Bubble icon in the main part of the screen. Once the type is selected, you are returned to the Chart Configuration screen, and the Chart Type value is set to the type that you chose. 7. Next, select **Chart Data** to add the data that will populate the chart.

Designer		×	
Chart Data Please specify what you would like to se	te in the chart	0 🕁	
Series Configurations +	Series Details Set properties of the selected series configuration		
Series 1 Clone   Rename   Disable   Remove	*Group: Click to set the formula used to group by	0	
	**:		
	*X:		
	Click to set formula for the X value *Size:	0	
	Click to set formula for the Size value	0	
	Data Plot X Axis Y Axis Filters		
+ Add a series			
O Preview changes	CANCEL PREVIEW	SAVE	

The **Series Configurations** section of the screen contains settings to manage each series. One series is created by default. A series is a set of data to display on the chart, for example a line on a line or area chart or one set of columns on a column or bar chart. The details for the selected series appear in **Series Details**, which contains settings to define the details of the selected series, such as the data to display as well as formatting and labels. From Series Configurations, you can clone, rename, disable (remove the series from the chart without deleting it), or remove a series. You can add a series by clicking the plus icon (+) or delete all series by clicking the trashcan icon (IF). Removing a series cannot be undone.

- 8. Under Series Configurations, click **Rename** under Series 1 and give the series a meaningful name that describes the data that will populate the chart.
- 9. Under Series Details, click the **Group** field and select the property to use for grouping the data in the series. You can use functions to derive the values for the groups. The list below describes the icons and options that are available when choosing a property:
  - The Root Value (A) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IF) or back (IF) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 10. Next, click the **Y** field and select the property to use for the Y axis values in the series. You can include functions to derive the values.
- 11. Then select the property to use for the X axis values by clicking the **X** field and choosing a property. You can also include functions.
- 12. Next, click the **Size** field and select the property to use for the size axis, which determines the sizes of the bubbles for the series.
- If you want to add a group for the series, you can click Show formulas used for creating calculated series and add values for the following fields:
  - Series Group: Specifies a property to use for grouping data in addition to the Group value.

- Series Label: Specifies the property whose values should serve as the Series Group label. It is typically set to the same value as Series Group.
- Series Sort Direction: Specifies the sort direction for the series groups.
- 14. You can click **Save** any time to render on the dashboard the data that you have configured so far. You can add another series and repeat the steps above. You can also configure several more options for the chart by changing the optional parameters. The additional tabs and options are described below in Optional Configuration Settings.

# **Optional Configuration Settings**

This section describes each of the tabs and additional settings that are available for customizing a bubble chart.

- Chart Data: Plot Tab
- Chart Data: X Axis Tab
- Chart Data: Y Axis Tab
- Chart Data: Filters Tab
- Chart Theme
- Chart Details

#### **Chart Data: Plot Tab**

This tab contains settings that control series formatting, such as data labels, legends, and other display options.

Series Details	
Set properties of the selected series configuration	
Series Chart Type Bubble	>
Series Chart Style Plot style information	>
Series Chart Data Labels Data labels are set to automatic enablement	>
Series Chart Markers Markers are enabled	>
Show:	
Largest 70 -Automatic-	
Show in legend:	
Data Plot X Axis Y Axis Filters	

Field	Description
Series Chart Type	This option opens the Chart Types selection screen where you can select a different chart type for the selected series. The selection must be compatible with the types selected for any other series in the chart, however, or the incompatible series will fail validation and not be displayed on the chart until it is corrected.
Series Chart Style	This option contains the plot style settings that control the color of the bubbles.
Series Chart Data Labels	This option accesses the data label settings, which control the font, alignment, and font effects of the labels for the series.
Series Chart Markers	This option accesses the data marker settings. These settings control the placement and format of the data points for a series that appear when a user hovers the pointer over a bubble. For example, the image below shows a data marker:

Field	Description
	Reviews X: Parasite Y: 758K Value: 8.5
Show	This setting can be used to limit the data included in the series. The limit can be defined by the largest or smallest Group Label or Value.
Show in legend	This setting controls whether to show the name of the series in the legend.

### Chart Data: X Axis Tab

This tab defines the formats and labels for the X axis values in the series.

Series Details Set properties of the select	ted seri	es configura	tion	
Create a new axis   Dele	te curre	ent axis		
Axis:				
Title:				
Sort by:				
Soft by.				
Y Values				Ascending
Display axis on the op	posite s	ide		
Avia Titla Dataila				
Axis title is Title				>
Axis Labels				>
Axis labels are enabled				F
Axis Style Axis style information				>
Data Plot X	Axis	Y Axis	Filters	

Field	Description
Axis	The Axis and Title settings are populated with the X axis value from the Data tab. If multiple axes exist, you can selected a different value to use for the X axis.
Title	This setting defines the title for the X axis.
Sort by	This setting controls how to sort the bubbles, by X, Y, or Size axis.
Display axis on the opposite side	This setting can be enabled to move the X axis to the opposite side of the chart.
Axis Title Details	This option contains the X axis title settings, which control the font, alignment, and font effects of the axis title.

Field	Description
Axis Labels	This option contains the X axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option accesses the X axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

### Chart Data: Y Axis Tab

This tab defines the formats and labels for the Y axis values in the series.

Series Details Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Title:	_
Display axis on the opposite side	
Axis Title Details Axis title is Number Of Reviews	>
Axis Labels Axis labels are enabled	>
Axis Style Axis style information	>
Data Plot X Axis Y Axis Filters	

Field	Description
Axis	The Axis and Title settings are populated with the Y axis value from the Data tab. If multiple axes exist, you can selected a different value to use for the Y axis.
Title	This setting defines the title for the Y axis.

Field	Description
Display axis on the opposite side	This setting can be enabled to move the Y axis to the opposite side of the chart.
Axis Title Details	This option contains the Y axis title settings, which control the font, alignment, and font effects of the axis title.
Axis Labels	This option accesses the Y axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option contains the Y axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

# **Chart Data: Filters Tab**

This tab defines any filters to apply to the series.

Series Details	
Set properties of the selected series configuration	
Series Filters No filters specified	>
Y Filters No filters specified	>
X Filters No filters specified	>
Size Filters No filters specified	>
Data Plot X Axis Y Axis Filters	

Field	Description
Series Filters	This option can be used to define filters that apply to the entire series.

Field	Description
Y Filters	This option can be used to define filters that apply only to the ${\tt Y}$ axis values for the series.
X Filters	This option can be used to define filters that apply only to the ${\bf X}$ axis values for the series.
Size Filters	This option can be used to define filters that apply only to the $Size$ axis values for the series.

#### **Chart Theme**

This setting presents options for configuring the theme or color scheme for the chart.



## **Chart Details**

This option offers finer-grained customization settings than the Chart Theme. You can further customize the chart design by adding details such as a chart title and subtitle. You can also modify chart-level styles and fonts as well as legend and tooltip formats.

Chart Details Please specify details about your chart	G 🖾
Title:	
Subine:	
Chart Style Chart style information	>
Title Details Title: unspecified	>
Subtitle Details Subtitle: unspecified	>
Legend Details Legend is enabled	>
Tooltip Details Tooltip is enabled	>

#### **Create a Column Chart**

Like bar charts, column charts are also useful for comparing different sets of data or emphasizing drastic changes in a data set over time. This topic provides instructions for creating a column chart with minimal configuration. Descriptions of all of the available configuration options are included below the steps.

- Complete the Minimum Configuration
- Optional Configuration Settings

# **Complete the Minimum Configuration**

Follow the instructions below to create a column chart. The instructions guide you through completing the minimum configuration needed to display your data in the chart. Additional, optional configuration settings are described in Optional Configuration Settings.

 In the dashboard that you want to add a lens to, click Lenses in the main toolbar and select New. The Create Lens dialog box is displayed:

Cre	ate Lens		×
Joint	t the type of lens to create		
۹.	Search for lenses		
+ Path Result	to lens values s: 10		
	Table Tabular visualization of data Properties		
₽/	Form View and edit resource attributes. Properties		
սի	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

2. On the Create Lens dialog box, select **Chart**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens			×
General Information			
Description:			
(+) more			
	CANCEL	PREVIOUS	FINISH

3. Type a **Title** and optional **Description** for the lens.

4. Click Finish. The lens Designer dialog box is displayed:

Designer	×
Chart Configuration Please select options below to begin configuring your chart	
Chart Type Column	>
Chart Data Series not specified	>
Chart Theme Chart theme selection	>
Chart Details Title: unspecified, Subtitle: unspecified, Legend is enabled, Tooltip is enabled	>
O Preview changes CANCEL PREVIEW	SAVE

5. Click Chart Type to open the Chart Types dialog box:

Designer				×
Chart Types Please select a chart t	/pe			3
III Column Bar ✓ Line ✓ Bubble ✓ Funnel III Heat Map ✓ Pie ✓ Polar ✓ Scatter ▲ Area	Column Clustered Column Bar Clustered Bar Clustered Bar Line	100% Stacked Co	Auron Waterfal	1
Preview changes		CANCEL	PREVIEW	SAVE

- 6. On the left side of the screen, select **Column**. On the right side of the screen, select the icon for the type of bar chart that you want to create. There are three types:
  - Clustered Column: Compares values across categories.
  - **Stacked Column**: Compares the contribution of each value to a total across categories. Includes the option to configure multiple groups within a series to distinguish between groups of values inside the total value.

 100% Stacked Column: Compares each value as a percentage of the total. Includes the option to configure multiple groups within a series to distinguish between groups of values inside the total value.

Once the type is selected, you are returned to the Chart Configuration screen, and the Chart Type value is set to the type that you chose.

7. Next, select Chart Data to add the data that will populate the chart.

Designer		×
Chart Data Please specify what you would like to se	e in the chart	6
Series Configurations +	Series Details Set properties of the selected series configuration	
Clone   Rename   Disable   Remove	~Group: Click to set the formula used to group by	0
	Group label: Click to set formula for the Group label value	0
	*Value:	
	Click to set formula for the Value value + Show formulas used for creating calculated series	0
+ Add a series	Data Plot X Axis Y Axis Filters	
Preview changes	CANCEL PREVIEW	SAVE

The **Series Configurations** section of the screen contains settings to manage each series. One series is created by default. A series is a set of data to display on the chart, for example a line on a line or area chart or one set of columns on a column or bar chart. The details for the selected series appear in **Series Details**, which contains settings to define the details of the selected series, such as the data to display as well as formatting and labels. From Series Configurations, you can clone, rename, disable (remove the series from the chart without deleting it), or remove a series. You can add a series by clicking the plus icon (+) or delete all series by clicking the trashcan icon (I). Removing a series cannot be undone.

- 8. Under Series Configurations, click **Rename** under Series 1 and give the series a meaningful name that describes the data that will populate the chart.
- 9. Under Series Details, click the Group field and select the property to use for grouping the data in the series. These are the values for the X axis. You can use functions to derive the values for the groups. The list below describes the icons and options that are available when

choosing a property:

 The Root Value (A) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IP) or back (IP) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 10. Once you select the Group, the **Group Label** field is populated with the same value. This setting configures the property whose values should serve as the group label. You can edit the value if necessary.
- 11. Next, click the **Value** field and select the property to use for the Y axis values. You can include functions to derive the values.
- 12. If you want to add a group for the series, you can click **Show formulas used for creating calculated series** and add values for the following fields:

- Series Group: Specifies a property to use for grouping data in addition to the Group value.
- Series Label: Specifies the property whose values should serve as the Series Group label. It is typically set to the same value as Series Group.
- Series Sort Direction: Specifies the sort direction for the series groups.
- 13. You can click Save any time to render on the dashboard the data that you have configured so far. You can add another series and repeat the steps above. You can also configure several more options for the chart by changing the optional parameters. The additional tabs and options are described below in Optional Configuration Settings.

# **Optional Configuration Settings**

This section describes each of the tabs and additional settings that are available for customizing a column chart.

- Chart Data: Plot Tab
- Chart Data: X Axis Tab
- Chart Data: Y Axis Tab
- Chart Data: Filters Tab
- Chart Theme
- Chart Details

#### **Chart Data: Plot Tab**

This tab contains settings that control series formatting, such as data labels, legends, and other display options.

Series Details Set properties of the	selected seri	es configura	tion	
Series Chart T Clustered Column	уре			>
Series Chart S Plot style informati	<b>tyle</b> on			>
Series Chart D Data labels are set	ata Labels to automatic	s enablement		>
Show:				
Largest 70	-Automati	c-		
Show in legend:				
$\checkmark$				
Data Plot	X Axis	Y Axis	Filters	

Field	Description
Series Chart Type	This option opens the Chart Types selection screen where you can select a different chart type for the selected series. The selection must be compatible with the types selected for any other series in the chart, however, or the incompatible series will fail validation and not be displayed on the chart until it is corrected.
Series Chart Style	This option contains the plot style settings that control the thickness and color of the columns for the series.
Series Chart Data Labels	This option accesses the data label settings, which control the font, alignment, and font effects of the labels for the series.
Show	This setting can be used to limit the data included in the series. The limit can be defined by the largest or smallest Group Label or Value.
Show in legend	This setting controls whether to show the name of the series in the legend.

## Chart Data: X Axis Tab

This tab defines the formats and labels for the X axis values in the series.

Series Details Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Title:	
Sort by:	
Group label Values As	scending
Display axis on the opposite side	
Axis Title Details Axis title is Director	>
Axis Labels Axis labels are enabled	>
Axis Style Axis style information	>

Field	Description
Axis	The Axis and Title settings are populated with the values from the Group and Group Label values from the Data tab. If multiple axes exist, you can selected a different value to use for the X axis.
Title	This setting defines the title for the X axis.
Sort by	This setting controls how to sort the data for the X axis, by the Group Label values or the Value values.
Display axis on the opposite	This setting can be enabled to move the X axis to the opposite side of the chart.

Field	Description
side	
Axis Title Details	This option contains the X axis title settings, which control the font, alignment, and font effects of the axis title.
Axis Labels	This option accesses the X axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option contains the X axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

#### Chart Data: Y Axis Tab

This tab defines the formats and labels for the Y axis values in the series.

Series Details Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Title:	
Type:	
Linear	
Display axis on the opposite side	
Axis Title Details Axis title is Rating	>
Axis Labels Axis labels are enabled	>
Axis Style Axis style information	>
Data Plot X Axis Y Axis Filters	

Field	Description
Axis	The Axis and Title settings are populated with the Value from the Data tab. If multiple axes exist, you can selected a different value to use for the Y axis.
Title	This setting defines the title for the Y axis.
Туре	This setting controls the scale for the Y axis, linear or logarithmic.
Display axis on the opposite side	This setting can be enabled to move the Y axis to the opposite side of the chart.
Axis Title Details	This option contains the Y axis title settings, which control the font, alignment, and font effects of the axis title.
Axis Labels	This option contains the Y axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option accesses the Y axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

# **Chart Data: Filters Tab**

This tab defines any filters to apply to the series.

Series Details Set properties of the selected series configuration	
Series Filters No filters specified	>
Group Filters No filters specified	>
Value Filters No filters specified	>
Data Plot X Axis Y Axis Filters	

Field	Description
Series Filters	This option can be used to define filters that apply to the entire series.
Group Filters	This option can be used to define filters that apply only to the Group values for the series.
Value Filters	This option can be used to define filters that apply only to the Value values for the series.

# **Chart Theme**

This setting presents options for configuring the theme or color scheme for the chart.



## **Chart Details**

This option offers finer-grained customization settings than the Chart Theme. You can further customize the chart design by adding details such as a chart title and subtitle. You can also modify chart-level styles and fonts as well as legend and tooltip formats.

Chart Details Please specify details about your chart	3 🖨
Title:	
Subtitle:	
Chart Style Chart style information	>
Title Details Title: unspecified	>
Subtitle Details Subtitle: unspecified	>
Legend Details Legend is enabled	>
Tooltip Details Tooltip is enabled	>

#### **Create a Funnel Chart**

A funnel chart is useful for showing progress through stages in which the data values typically decrease. This topic provides instructions for creating a funnel chart with minimal configuration. Descriptions of all of the available configuration options are included below the steps.

- Complete the Minimum Configuration
- Optional Configuration Settings

## **Complete the Minimum Configuration**

Follow the instructions below to create a funnel chart. The instructions guide you through completing the minimum configuration needed to display your data in the chart. Additional, optional configuration settings are described in Optional Configuration Settings.

1. In the dashboard that you want to add a lens to, click **Lenses** in the main toolbar and select

New. The Create Lens dialog box is displayed:

Crea	ate Lens		×
<b>Q</b> 5	Search for lenses		
+ Path Results	to lens values s: 10		
	Table Tabular visualization of data Properties		
₽,	Form View and edit resource attributes. Properties		
ılı	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

2. On the Create Lens dialog box, select **Chart**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens		×
General Information		
Title.*		
Description:		
(+) more		

3. Type a **Title** and optional **Description** for the lens.

4. Click Finish. The lens Designer dialog box is displayed:

Designer	×
Chart Configuration Please select options below to begin configuring your chart	
Chart Type Column	>
Chart Data Series not specified	>
Chart Theme Chart theme selection	>
Chart Details Title: unspecified, Subtitle: unspecified, Legend is enabled, Toolitip is enabled	>
O Preview changes CANCEL PREVIEW	SAVE

5. Click Chart Type to open the Chart Types dialog box:

Designer		×
Chart Types Please select a chart t	/pe	0 6
III Column Bar Bar Cline Gubble Funnel III Heat Map Pie Polar Scatter ▲ Area	Column	100%. Stacked Bar
Preview changes	CA	ANCEL PREVIEW SAVE

6. On the left side of the screen, select **Funnel**, and then select the Funnel icon in the main part of the screen. Once the type is selected, you are returned to the Chart Configuration screen, and the Chart Type value is set to the type that you chose. 7. Next, select **Chart Data** to add the data that will populate the chart.

Designer		×
Chart Data Please specify what you would like to se	e in the chart	0 🖨
Series Configurations + == Select a configuration to edit	Series Details Set properties of the selected series configuration	
Series 1	*Group:	
Clone   Rename   Disable   Remove	Click to set the formula used to group by	0
	Group label:	
	Click to set formula for the Group label value	0
	*Value:	
	Click to set formula for the Value value	θ
	+ Show formulas used for creating calculated series	
+ Add a series	Data Plot Filters	
O Preview changes	CANCEL PREVIEW	SAVE

The **Series Configurations** section of the screen contains settings to manage each series. One series is created by default. A series is a set of data to display on the chart, for example a line on a line or area chart or one set of columns on a column or bar chart. The details for the selected series appear in **Series Details**, which contains settings to define the details of the selected series, such as the data to display as well as formatting and labels. From Series Configurations, you can clone, rename, disable (remove the series from the chart without deleting it), or remove a series. You can add a series by clicking the plus icon (+) or delete all series by clicking the trashcan icon (**I**). Removing a series cannot be undone.

- 8. Under Series Configurations, click **Rename** under Series 1 and give the series a meaningful name that describes the data that will populate the chart.
- 9. Under Series Details, click the Group field and select the property to use for grouping the data in the series. You can use functions to derive the values for the groups. The list below describes the icons and options that are available when choosing a property:
  - The **Root Value** (**•**) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IP) or back (IP) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 10. Once you select the Group, the **Group Label** field is populated with the same value. This setting configures the property whose values should serve as the group label. You can edit the value if necessary.
- 11. Next, click the **Value** field and select the property to use for the values in the groups. You can include functions to derive the values.
- 12. If you want to add a group for the series, you can click **Show formulas used for creating calculated series** and add values for the following fields:
  - Series Group: Specifies a property to use for grouping data in addition to the Group value.

- Series Label: Specifies the property whose values should serve as the Series Group label. It is typically set to the same value as Series Group.
- Series Sort Direction: Specifies the sort direction for the series groups.
- 13. You can click **Save** any time to render on the dashboard the data that you have configured so far. You can add another series and repeat the steps above. You can also configure several more options for the chart by changing the optional parameters. The additional tabs and options are described below in Optional Configuration Settings.

# **Optional Configuration Settings**

This section describes each of the tabs and additional settings that are available for customizing a funnel chart.

- Chart Data: Plot Tab
- Chart Data: Filters Tab
- Chart Theme
- Chart Details

#### **Chart Data: Plot Tab**

This tab contains settings that control series formatting, such as data labels, legends, and other display options.

Series Details	
Set properties of the selected series configuration	
Series Chart Type Funnel	>
Series Chart Style Plot style information	>
Series Chart Data Labels Data labels are set to automatic enablement	>
Show:	
Largest 70 -Automatic-	
Show in legend:	
Data Plot Filters	

Field	Description
Series Chart Type	This option opens the Chart Types selection screen where you can select a different chart type for the selected series. The selection must be compatible with the types selected for any other series in the chart, however, or the incompatible series will fail validation and not be displayed on the chart until it is corrected.
Series Chart Style	This option contains the plot style settings that control the color of the groups in the series.
Series Chart Data Labels	This option accesses the data label settings, which control the font, alignment, and font effects of the labels for the series.
Show	This setting can be used to limit the data included in the series. The limit can be defined by the largest or smallest Group Label or Value.
Show in legend	This setting controls whether to show the name of the series in the legend.

### **Chart Data: Filters Tab**

This tab defines any filters to apply to the series.

Series Details Set properties of the selected series configuration	
Series Filters No filters specified	>
Value Filters No filters specified	>
Data Plot Filters	

Field	Description
Series Filters	This option can be used to define filters that apply to the entire series.
Value Filters	This option can be used to define filters that apply only to the Value values for the series.

#### **Chart Theme**

This setting presents options for configuring the theme or color scheme for the chart.



## **Chart Details**

This option offers finer-grained customization settings than the Chart Theme. You can further customize the chart design by adding details such as a chart title and subtitle. You can also modify chart-level styles and fonts as well as legend and tooltip formats.

Chart Details Please specify details about your chart	G 🕁
Title:	
Subtitle:	
Chart Style Chart style information	>
Title Details Title: unspecified	>
Subtitle Details Subtitle: unspecified	>
Legend Details Legend is enabled	>
Tooltip Details Tooltip is enabled	>

#### **Create a Heat Map**

A heat map depicts values of interest across two axes. The axis variables are divided into ranges like a bar chart, and each cell's color indicates the value in the corresponding cell range. This topic provides instructions for creating a heat map with minimal configuration. Descriptions of all of the available configuration options are included below the steps.

- Complete the Minimum Configuration
- Optional Configuration Settings

## **Complete the Minimum Configuration**

Follow the instructions below to create a heat map. The instructions guide you through completing the minimum configuration needed to display your data in the chart. Additional, optional configuration settings are described in Optional Configuration Settings.

1. In the dashboard that you want to add a lens to, click **Lenses** in the main toolbar and select

New. The Create Lens dialog box is displayed:

Crea	ate Lens		×
<b>Q</b> 5	Search for lenses		
+ Path Results	to lens values s: 10		
	Table Tabular visualization of data Properties		
₽,	Form View and edit resource attributes. Properties		
ılı	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

2. On the Create Lens dialog box, select **Chart**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens		×
General Information		
Title.*		
Description:		
(+) more		

3. Type a **Title** and optional **Description** for the lens.

4. Click Finish. The lens Designer dialog box is displayed:

Designer	×
Chart Configuration Please select options below to begin configuring your chart	
Chart Type Column	>
Chart Data Series not specified	>
Chart Theme Chart theme selection	>
Chart Details Title: unspecified, Subtitle: unspecified, Legend is enabled, Toolitip is enabled	>
O Preview changes CANCEL PREVIEW	SAVE

5. Click Chart Type to open the Chart Types dialog box:

Designer		×
Chart Types Please select a chart t	/pe	0 6
III Column Bar Bar Cline Gubble Funnel III Heat Map Pie Polar Scatter ▲ Area	Column	100%. Stacked Bar
Preview changes	CA	ANCEL PREVIEW SAVE

6. On the left side of the screen, select **Heat Map**, and then select the Heat Map icon in the main part of the screen. Once the type is selected, you are returned to the Chart Configuration screen, and the Chart Type value is set to the type that you chose.

7. Next, select **Chart Data** to add the data that will populate the chart.

Designer		×
Chart Data Please specify what you would like to	see in the chart	<mark>6</mark> 🔏
Series Configurations +	Series Details Set properties of the selected series configuration	
Series 1 Clone   Rename   Disable   Remove	*Group:	
	Click to set the formula used to group by	0
	-x. Click to set formula for the X value	0
	*Å:	
	Click to set formula for the Y value	0
	*Value:	•
	- Click to set formula for the Value value     - Show formulas used for creating calculated series	•
	Data Plot X Axis Y Axis Color Axis Filters	5
+ Add a series		
Preview changes	CANCEL PREVIEW S	AVE

The **Series Configurations** section of the screen contains settings to manage each series. One series is created by default. A series is a set of data to display on the chart, for example a line on a line or area chart or one set of columns on a column or bar chart. The details for the selected series appear in **Series Details**, which contains settings to define the details of the selected series, such as the data to display as well as formatting and labels. From Series Configurations, you can clone, rename, disable (remove the series from the chart without deleting it), or remove a series. You can add a series by clicking the plus icon (+) or delete all series by clicking the trashcan icon (**I:**). Removing a series cannot be undone.

- 8. Under Series Configurations, click **Rename** under Series 1 and give the series a meaningful name that describes the data that will populate the chart.
- 9. Under Series Details, click the Group field and select the property to use for grouping the data in the series. You can use functions to derive the values for the groups. The list below describes the icons and options that are available when choosing a property:
  - The Root Value (A) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IP) or back (IP) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 10. Next, click the **X** field and select the property to use for the X axis values in the series. You can include functions to derive the values.
- 11. Then select the property to use for the Y axis values by clicking the Y field and choosing a property. You can also include functions.
- 12. Next, click the **Value** field and select the property to use for the value range in the series.
- 13. If you want to add a group for the series, you can click **Show formulas used for creating calculated series** and add values for the following fields:
  - Series Group: Specifies a property to use for grouping data in addition to the Group value.
- Series Label: Specifies the property whose values should serve as the Series Group label. It is typically set to the same value as Series Group.
- Series Sort Direction: Specifies the sort direction for the series groups.
- 14. You can click **Save** any time to render on the dashboard the data that you have configured so far. You can add another series and repeat the steps above. You can also configure several more options for the chart by changing the optional parameters. The additional tabs and options are described below in Optional Configuration Settings.

# **Optional Configuration Settings**

This section describes each of the tabs and additional settings that are available for customizing a heat map.

- Chart Data: Plot Tab
- Chart Data: X Axis Tab
- Chart Data: Y Axis Tab
- Chart Data: Color Axis Tab
- Chart Data: Filters Tab
- Chart Theme
- Chart Details

### Chart Data: Plot Tab

This tab contains settings that control series formatting, such as data labels, legends, and other display options.

Series Details Set properties of the selected series configuration	
Series Chart Type Heat Map	>
Series Chart Style Plot style information	>
Series Chart Data Labels Data labels are set to automatic enablement	>
Show: Largest 70 -Automatic-	
Show in legend:	
Data         Plot         X Axis         Y Axis         Color Axis         Filters	

Field	Description
Series Chart Type	This option opens the Chart Types selection screen where you can select a different chart type for the selected series. The selection must be compatible with the types selected for any other series in the chart, however, or the incompatible series will fail validation and not be displayed on the chart until it is corrected.
Series Chart Style	This option contains the plot style settings that control the color of the ranges.
Series Chart Data Labels	This option contains the data label settings, which control the font, alignment, and font effects of the labels for the series.
Show	This setting can be used to limit the data included in the series. The limit is defined based on the largest or smallest x axis, y axis, or Value values.
Show in legend	This setting controls whether to show the name of the series in the legend.

# Chart Data: X Axis Tab

This tab defines the formats and labels for the X axis values in the series.

Series Details	
Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Title:	
Display values as evenly distributed categories	
Display axis on the opposite side	
Axis Title Details	
Axis title is Listime	>
Axis Labels	, ,
Axis labels are enabled	
Axis Style	\$
Axis style information	
Data Diat VAvie VAvie ColorAvie Eiltere	
Data Flot A Axis T Axis Color Axis Fliters	

Field	Description
Axis	The Axis and Title settings are populated with the X axis value from the Data tab. If multiple axes exist, you can selected a different value to use for the X axis.
Title	This setting defines the title for the X axis.
Display values as evenly distributed categories	This setting controls whether to evenly distribute X axis values.
Display axis on the opposite side	This setting can be enabled to move the X axis to the opposite side of the chart.
Axis Title Details	This option contains the X axis title settings, which control the font,

Field	Description
	alignment, and font effects of the axis title.
Axis Labels	This option accesses the X axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option contains the X axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

### Chart Data: Y Axis Tab

This tab defines the formats and labels for the Y axis values in the series.

Series Details Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Title:	
Туре:	
Linear	
Display axis on the opposite side	
Axis Title Details Axis title is Numtickets	>
Axis Labels Axis labels are enabled	>
Axis Style Axis style information	>
Data         Plot         X Axis         Y Axis         Color Axis         Filters	

Field	Description
Axis	The Axis and Title settings are populated with the Y axis value from the Data tab. If multiple axes exist, you can selected a different value to use for the Y axis.

Field	Description
Title	This setting defines the title for the Y axis.
Туре	This setting controls the scale for the Y axis, linear or logarithmic.
Display axis on the opposite side	This setting can be enabled to move the Y axis to the opposite side of the chart.
Axis Title Details	This option contains the Y axis title settings, which control the font, alignment, and font effects of the axis title.
Axis Labels	This option accesses the Y axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option contains the Y axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

### **Chart Data: Color Axis Tab**

This tab defines the value range block colors and axis labels and styles.

Series Details Set properties of the selected series configuration	
Minimum color:	
Maximum color:	
Axis Labels Axis labels are enabled	>
Axis Style Axis style information	>
Data Plot X Axis Y Axis Color Axis Filters	

Field	Description
Minimum color	This option defines the color for the minimum range of values.
Maximum color	This option defines the color for the maximum range of values.
Axis Labels	This option accesses the color axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option contains the color axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

## **Chart Data: Filters Tab**

This tab defines any filters to apply to the series.

Series Details Set properties of the selected series configuration	
Series Filters No filters specified	>
X Filters No filters specified	>
Y Filters No filters specified	>
Value Filters No filters specified	>
Data Plot X Axis Y Axis Color Axis Fi	Iters

Field	Description
Series Filters	This option can be used to define filters that apply to the entire series.

Field	Description
X Filters	This option can be used to define filters that apply only to the ${\rm X}$ axis values for the series.
Y Filters	This option can be used to define filters that apply only to the ${\tt Y}$ axis values for the series.
Value Filters	This option can be used to define filters that apply only to the <code>Value</code> values for the series.

### **Chart Theme**

This setting presents options for configuring the theme or color scheme for the chart.



## **Chart Details**

This option offers finer-grained customization settings than the Chart Theme. You can further customize the chart design by adding details such as a chart title and subtitle. You can also modify chart-level styles and fonts as well as legend and tooltip formats.

Chart Details Please specify details about your chart	G 🖾
Title:	
Subtitle:	
Chart Style Chart style information	>
Title Details Title: unspecified	>
Subtitle Details Subtitle: unspecified	>
Legend Details Legend is enabled	>
Tooltip Details Tooltip is enabled	>

#### **Create a Line Chart**

Line charts are useful for emphasizing trends in your data. This topic provides instructions for creating a line chart with minimal configuration. Descriptions of all of the available configuration options are included below the steps.

- Complete the Minimum Configuration
- Optional Configuration Settings

# **Complete the Minimum Configuration**

Follow the instructions below to create a line chart. The instructions guide you through completing the minimum configuration needed to display your data in the chart. Additional, optional configuration settings are described in Optional Configuration Settings.

 In the dashboard that you want to add a lens to, click Lenses in the main toolbar and select New. The Create Lens dialog box is displayed:

Cre	ate Lens		×
Joint	t the type of lens to create		
۹.	Search for lenses		
+ Path Result	to lens values s: 10		
	Table Tabular visualization of data Properties		
₽/	Form View and edit resource attributes. Properties		
սի	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

2. On the Create Lens dialog box, select **Chart**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens		×
General Information		
Title.*		
Description:		
(+) more		

3. Type a **Title** and optional **Description** for the lens.

4. Click Finish. The lens Designer dialog box is displayed:

Designer	×
Chart Configuration Please select options below to begin configuring your chart	
Chart Type Column	>
Chart Data Series not specified	>
Chart Theme Chart theme selection	>
Chart Details Title: unspecified, Subtitle: unspecified, Legend is enabled, Tooltip is enabled	>
O Preview changes CANCEL PREVI	IEW SAVE

5. Click Chart Type to open the Chart Types dialog box:

Designer		×
Chart Types Please select a chart ty	rpe	0 6
Column Column Sar Sar Column Sar	Column Clustered Column Bar	100% Stacked Column
<ul> <li>Pie</li> <li>Polar</li> <li>Scatter</li> <li>Area</li> </ul>	Custered Bar	100% Stacked Bar
Preview changes		CANCEL PREVIEW SAVE

- 6. On the left side of the screen, select **Line**. On the right side of the screen, select the icon for the type of line chart that you want to create. There are three types:
  - Line: Connects value points on the chart with straight lines.
  - **Spline**: Connects value points on the chart with curved lines.
  - Step Line: Connects value points on the chart with horizontal lines.

Once the type is selected, you are returned to the Chart Configuration screen, and the Chart Type value is set to the type that you chose.

7. Next, select **Chart Data** to add the data that will populate the chart.

Designer		×
Chart Data Please specify what you would like to se	e in the chart	<mark>6</mark> 🕁
Series Configurations + IIF Select a configuration to edit	Series Details Set properties of the selected series configuration *Group:	
Clone   Rename   Disable   Remove	Click to set the formula used to group by	θ
	Group label: Click to set formula for the Group label value	0
	*Value: Click to set formula for the Value value	0
	+ Show formulas used for creating calculated series	-
+ Add a series	Data Plot X Axis Y Axis Filters	
O Preview changes	CANCEL PREVIEW	SAVE

The **Series Configurations** section of the screen contains settings to manage each series. One series is created by default. A series is a set of data to display on the chart, for example a line on a line or area chart or one set of columns on a column or bar chart. The details for the selected series appear in **Series Details**, which contains settings to define the details of the selected series, such as the data to display as well as formatting and labels. From Series Configurations, you can clone, rename, disable (remove the series from the chart without deleting it), or remove a series. You can add a series by clicking the plus icon (+) or delete all series by clicking the trashcan icon (IF). Removing a series cannot be undone.

- 8. Under Series Configurations, click **Rename** under Series 1 and give the series a meaningful name that describes the data that will populate the chart.
- 9. Under Series Details, click the Group field and select the property to use for grouping the data in the series. These are the values for the X axis. You can use functions to derive the values for the groups. The list below describes the icons and options that are available when choosing a property:
  - The Root Value (A) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IF) or back (IF) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 10. Once you select the Group, the **Group Label** field is populated with the same value. This setting configures the property whose values should serve as the group label. You can edit the value if necessary.
- 11. Next, click the **Value** field and select the property to use for the Y axis values. You can include functions to derive the values.
- 12. If you want to add a group for the series, you can click **Show formulas used for creating calculated series** and add values for the following fields:
  - Series Group: Specifies a property to use for grouping data in addition to the Group value.

- Series Label: Specifies the property whose values should serve as the Series Group label. It is typically set to the same value as Series Group.
- Series Sort Direction: Specifies the sort direction for the series groups.
- 13. You can click **Save** any time to render on the dashboard the data that you have configured so far. You can add another series and repeat the steps above. You can also configure several more options for the chart by changing the optional parameters. The additional tabs and options are described below in Optional Configuration Settings.

# **Optional Configuration Settings**

This section describes each of the tabs and additional settings that are available for customizing a line chart.

- Chart Data: Plot Tab
- Chart Data: X Axis Tab
- Chart Data: Y Axis Tab
- Chart Data: Filters Tab
- Chart Theme
- Chart Details

### **Chart Data: Plot Tab**

This tab contains settings that control series formatting, such as data labels, legends, and other display options.

Series Details	selected seri	es configura	ation	
Series Chart T	ype	es comigure		>
Series Chart S Plot style informati	<b>tyle</b> on			>
Series Chart D Data labels are set	ata Labels	enablement		>
Series Chart Markers Markers are enabled		>		
Show:				
Largest 70	-Automati	C-		
Show in legend:				
$\checkmark$				
Connect missing poi	nts:			
Data Plot	X Axis	Y Axis	Filters	

Field	Description
Series Chart Type	This option opens the Chart Types selection screen where you can select a different chart type for the selected series. The selection must be compatible with the types selected for any other series in the chart, however, or the incompatible series will fail validation and not be displayed on the chart until it is corrected.
Series Chart Style	This option contains the plot style settings that control the thickness and color of lines for the series.
Series Chart Data Labels	This option contains the data label settings, which control the font, alignment, and font effects of the labels for the series.
Series Chart Markers	This option contains the data marker settings. These settings control the placement and format of the data points for a series that appear when a user hovers the pointer over a point in the line. For example, the image below shows a data marker:

Field	Description
	Series 1 X: 7 Y: 20,000
Show	This setting can be used to limit the data included in the series. The limit can be defined by the largest or smallest Group Label or Value.
Show in legend	This setting controls whether to show the name of the series in the legend.
Connect missing points	Selecting this option connects the series line across any missing points.

# Chart Data: X Axis Tab

This tab defines the formats and labels for the X axis values in the series.

Series Details	
Create a new axis   Delete current axis	
Axis:	
Title:	
Туре:	
Linear	
Display values as evenly distributed categories	
Display axis on the opposite side	
Avis Title Details	
Axis title is hasTickit Events/Catid	>
Avis Lahole	
Axis labels are enabled	>
Axis Style	
Axis style information	>
Data Plot X Axis Y Axis Filters	

Field	Description
Axis	The Axis and Title settings are populated with the Group and Group Label values from the Data tab. If multiple axes exist, you can selected a different value to use for the X axis.
Title	This setting defines the title for the X axis.
Туре	This setting controls the scale for the X axis, linear or logarithmic.
Display values as evenly distributed categories	This setting controls whether to evenly distribute X axis values.
Display axis on the opposite side	This setting can be enabled to move the X axis to the opposite side of the chart.
Axis Title Details	This option contains the X axis title settings, which control the font,

Field	Description
	alignment, and font effects of the axis title.
Axis Labels	This option accesses the X axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option contains the X axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

### Chart Data: Y Axis Tab

This tab defines the formats and labels for the Y axis values in the series.

Series Details Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Title:	
Туре:	
Linear	
Display axis on the opposite side	
Axis Title Details	
Axis title is Totalprice	
Axis Labels Axis labels are enabled	>
Axis Style	
Axis style information	/
Data Plot X Axis Y Axis Filters	

Field	Description
Axis	The Axis and Title settings are populated with the ${\tt Value}$ from the Data tab. If

Field	Description
	multiple axes exist, you can selected a different value to use for the Y axis.
Title	This setting defines the title for the Y axis.
Туре	This setting controls the scale for the Y axis, linear or logarithmic.
Display axis on the opposite side	This setting can be enabled to move the Y axis to the opposite side of the chart.
Axis Title Details	This option contains the Y axis title settings, which control the font, alignment, and font effects of the axis title.
Axis Labels	This option accesses the Y axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option contains the Y axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

# **Chart Data: Filters Tab**

This tab defines any filters to apply to the series.

Series Details Set properties of the selected series configuration	
Series Filters No filters specified	>
Group Filters No filters specified	>
Value Filters No filters specified	>
Data Plot X Axis Y Axis Filters	

Field	Description
Series Filters	This option can be used to define filters that apply to the entire series.
Group Filters	This option can be used to define filters that apply only to the Group values for the series.
Value Filters	This option can be used to define filters that apply only to the Value values for the series.

### **Chart Theme**

This setting presents options for configuring the theme or color scheme for the chart.



# **Chart Details**

This option offers finer-grained customization settings than the Chart Theme. You can further customize the chart design by adding details such as a chart title and subtitle. You can also modify chart-level styles and fonts as well as legend and tooltip formats.

Chart Details Please specify details about your chart	G 🖾
Title:	
Subine:	
Chart Style Chart style information	>
Title Details Title: unspecified	>
Subtitle Details Subtitle: unspecified	>
Legend Details Legend is enabled	>
Tooltip Details Tooltip is enabled	>

#### **Create a Polar Chart**

Polar or spider charts can be useful for showing similarities, differences, and outliers in your data. These charts plot one or more groups of values over common variables. The variables are axes that are arranged radially around a central point. This topic provides instructions for creating a polar chart with minimal configuration. Descriptions of all of the available configuration options are included below the steps.

- Complete the Minimum Configuration
- Optional Configuration Settings

## **Complete the Minimum Configuration**

Follow the instructions below to create a polar chart. The instructions guide you through completing the minimum configuration needed to display your data in the chart. Additional, optional configuration settings are described in Optional Configuration Settings.

 In the dashboard that you want to add a lens to, click Lenses in the main toolbar and select New. The Create Lens dialog box is displayed:

Cre Selec	ate Lens t the type of lens to create		×
Q :	Search for lenses		
+ Path Result	to lens values s: 10		
	Table Tabular visualization of data Properties		
₽/	Form View and edit resource attributes. Properties		
dı	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

2. On the Create Lens dialog box, select **Chart**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens			×
General Information			
Description:			
(+) more			
	CANCEL	PREVIOUS	FINISH

3. Type a **Title** and optional **Description** for the lens.

4. Click Finish. The lens Designer dialog box is displayed:

Designer	×
Chart Configuration Please select options below to begin configuring your chart	
Chart Type Column	>
Chart Data Series not specified	>
Chart Theme Chart theme selection	>
Chart Details Title: unspecified, Subtitle: unspecified, Legend is enabled, Tooltip is enabled	>
O Preview changes CANCEL PREVI	IEW SAVE

5. Click Chart Type to open the Chart Types dialog box:

	×
ype	<b>G</b> (
Column	d Column
Clustered Bar Stacked Bar 100% Stack	ked Bar
	ype Column Coutered Column Bar Coutered Bar Coutered Bar

6. On the left side of the screen, select **Polar**, and then select the Spider icon in the main part of the screen. Once the type is selected, you are returned to the Chart Configuration screen, and the Chart Type value is set to the type that you chose.

7. Next, select **Chart Data** to add the data that will populate the chart.

Designer		×
Chart Data Please specify what you would like to se	e in the chart	3
Series Configurations + #F Select a configuration to edit Series 1 Clone   Rename   Disable   Remove	Series Details Set properties of the selected series configuration *Group: Click to set the formula used to group by Group label: Click to set formula for the Group label value	9
	*Value: Click to set formula for the Value value * Show formulas used for creating calculated series	0
+ Add a series	Data Plot X Axis Y Axis Filters CANCEL PREVIEW	SAVE

The **Series Configurations** section of the screen contains settings to manage each series. One series is created by default. A series is a set of data to display on the chart, for example a line on a line or area chart or one set of columns on a column or bar chart. The details for the selected series appear in **Series Details**, which contains settings to define the details of the selected series, such as the data to display as well as formatting and labels. From Series Configurations, you can clone, rename, disable (remove the series from the chart without deleting it), or remove a series. You can add a series by clicking the plus icon (+) or delete all series by clicking the trashcan icon (IF). Removing a series cannot be undone.

- 8. Under Series Configurations, click **Rename** under Series 1 and give the series a meaningful name that describes the data that will populate the chart.
- 9. Under Series Details, click the Group field and select the property to use for grouping the data in the series. These are the values for the X axis. You can use functions to derive the values for the groups. The list below describes the icons and options that are available when choosing a property:
  - The Root Value (A) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IP) or back (IP) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 10. Once you select the Group, the **Group Label** field is populated with the same value. This setting configures the property whose values should serve as the group label. You can edit the value if necessary.
- 11. Next, click the **Value** field and select the property to use for the Y axis values. You can include functions to derive the values.
- 12. If you want to add a group for the series, you can click **Show formulas used for creating calculated series** and add values for the following fields:
  - Series Group: Specifies a property to use for grouping data in addition to the Group value.

- Series Label: Specifies the property whose values should serve as the Series Group label. It is typically set to the same value as Series Group.
- Series Sort Direction: Specifies the sort direction for the series groups.
- 13. You can click **Save** any time to render on the dashboard the data that you have configured so far. You can add another series and repeat the steps above. You can also configure several more options for the chart by changing the optional parameters. The additional tabs and options are described below in Optional Configuration Settings.

# **Optional Configuration Settings**

This section describes each of the tabs and additional settings that are available for customizing a polar chart.

- Chart Data: Plot Tab
- Chart Data: X Axis Tab
- Chart Data: Y Axis Tab
- Chart Data: Filters Tab
- Chart Theme
- Chart Details

### **Chart Data: Plot Tab**

This tab contains settings that control series formatting, such as data labels, legends, and other display options.

Series Details				
Set properties of the	selected seri	es configura	tion	
Series Chart T Spider	уре			>
Series Chart S Plot style informati	<b>tyle</b> on			>
Series Chart D Data labels are set	ata Labels	enablement		>
Show:				
Largest 70	-Automati	C-		
Show in legend:				
$\checkmark$				
Data Plot	X Axis	Y Axis	Filters	

Field	Description
Series Chart Type	This option opens the Chart Types selection screen where you can select a different chart type for the selected series. The selection must be compatible with the types selected for any other series in the chart, however, or the incompatible series will fail validation and not be displayed on the chart until it is corrected.
Series Chart Style	This option contains the plot style setting that controls whether the plot lines have shadows.
Series Chart Data Labels	This option accesses the data label settings, which control the font, alignment, and font effects of the labels for the series.
Show	This setting can be used to limit the data included in the series. The limit can be defined by the largest or smallest Group Label or Value.
Show in legend	This setting controls whether to show the name of the series in the legend.

## Chart Data: X Axis Tab

This tab defines the formats and labels for the X axis values in the series.

Series Details Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Title:	
Display values as evenly distributed categories	
Display axis on the opposite side	
Axis Title Details Axis title is Listime	>
Axis Labels Axis labels are enabled	>
Axis Style Axis style information	>
Data Plot X Axis Y Axis Filters	

Field	Description
Axis	The Axis and Title settings are populated with the Group and Group label values from the Data tab. If multiple axes exist, you can selected a different value to use for the X axis.
Title	This setting defines the title for the X axis.
Display values as evenly distributed categories	This setting controls whether to evenly distribute X axis values.
Display axis on the opposite side	This setting can be enabled to move the X axis to the opposite side of the chart.

Field	Description
Axis Title Details	This option contains the X axis title settings, which control the font, alignment, and font effects of the axis title.
Axis Labels	This option contains the X axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option accesses the X axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

### **Chart Data: Y Axis Tab**

This tab defines the formats and labels for the Y axis values in the series.

Series Details Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Title:	
Туре:	
Linear	
Display axis on the opposite side	
Axis Title Details Axis title is hasTickit Users/hasTickit Sales/Qtysold	>
Axis Labels Axis labels are enabled	>
Axis Style Axis style information	>
Data Plot X Axis Y Axis Filters	

Field	Description
Axis	The Axis and Title settings are populated with the Value from the Data tab. If multiple axes exist, you can selected a different value to use for the Y axis.
Title	This setting defines the title for the Y axis.
Туре	This setting controls the scale for the Y axes, linear or logarithmic.
Display axis on the opposite side	This setting can be enabled to move the Y axis to the opposite side of the chart.
Axis Title Details	This option contains the Y axis title settings, which control the font, alignment, and font effects of the axis title.
Axis Labels	This option contains the Y axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option accesses the Y axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

# **Chart Data: Filters Tab**

This tab defines any filters to apply to the series.

Series Details Set properties of the selected series configuration	
Series Filters No filters specified	>
Group Filters No filters specified	>
Value Filters No filters specified	>
Data Plot X Axis Y Axis Filters	

Field	Description
Series Filters	This option can be used to define filters that apply to the entire series.
Group Filters	This option can be used to define filters that apply only to the Group values for the series.
Value Filters	This option can be used to define filters that apply only to the ${\tt Value}$ values for the series.

### **Chart Theme**

This setting presents options for configuring the theme or color scheme for the chart.



## **Chart Details**

This option offers finer-grained customization settings than the Chart Theme. You can further customize the chart design by adding details such as a chart title and subtitle. You can also modify chart-level styles and fonts as well as legend and tooltip formats.

Chart Details Please specify details about your chart	G 🕻
Title:	
Subtitle:	
Chart Style Chart style information	>
Title Details Title: unspecified	>
Subtitle Details Subtitle: unspecified	>
Legend Details Legend is enabled	>
Tooltip Details	>

#### **Create a Scatter Chart**

Scatter charts are also useful for observing relationships between variables, patterns, or outliers in your data. The position of dots on the horizontal and vertical axes represent values for individual data points. This topic provides instructions for creating a scatter chart with minimal configuration. Descriptions of all of the available configuration options are included below the steps.

- Complete the Minimum Configuration
- Optional Configuration Settings

## **Complete the Minimum Configuration**

Follow the instructions below to create a scatter chart. The instructions guide you through completing the minimum configuration needed to display your data in the chart. Additional, optional configuration settings are described in Optional Configuration Settings.

1. In the dashboard that you want to add a lens to, click **Lenses** in the main toolbar and select

New. The Create Lens dialog box is displayed:

Crea	ate Lens		×
۹.	Search for lenses		
+ Path Results	to lens values s: 10		
	Table Tabular visualization of data Properties		
₽⁄	Form View and edit resource attributes. Properties		
dı	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

2. On the Create Lens dialog box, select **Chart**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens			×
General Information			
Title.*			
Description:			
(+) more			
	CANCEL	PREVIOUS	FINISH

3. Type a **Title** and optional **Description** for the lens.

4. Click Finish. The lens Designer dialog box is displayed:

Designer	×	<
Chart Configuration Please select options below to begin configuring your chart		
Chart Type Column	>	
Chart Data Series not specified	>	
Chart Theme Chart theme selection	>	
Chart Details Title: unspecified, Subtitle: unspecified, Legend is enabled, Tooltip is enabled	>	_
O Preview changes CANCEL Pr	REVIEW SAVE	

5. Click Chart Type to open the Chart Types dialog box:

Designer		×
Chart Types Please select a chart ty	pe	0 (
Column Co	Column Chustered Column Bar Custered Bar Custered Bar	100% Stacked Column
Area	Line	CANCEL PREVIEW SAVE

6. On the left side of the screen, select **Scatter**, and then select the Scatter icon in the main part of the screen. Once the type is selected, you are returned to the Chart Configuration screen, and the Chart Type value is set to the type that you chose. 7. Next, select **Chart Data** to add the data that will populate the chart.

Designer		×
Chart Data Please specify what you would like to se	e in the chart	6
Series Configurations + in Select a configuration to edit	Series Details Set properties of the selected series configuration	
Series 1 Clone   Rename   Disable   Remove	*Group:	•
	*X:	•
	Click to set formula for the X value	0
	*Y:	
	Click to set formula for the Y value	•
	+ Show formulas used for creating calculated series	
+ Add a series	Data Plot X Axis Y Axis Filters	
Preview changes	CANCEL PREVIEW	SAVE

The **Series Configurations** section of the screen contains settings to manage each series. One series is created by default. A series is a set of data to display on the chart, for example a line on a line or area chart or one set of columns on a column or bar chart. The details for the selected series appear in **Series Details**, which contains settings to define the details of the selected series, such as the data to display as well as formatting and labels. From Series Configurations, you can clone, rename, disable (remove the series from the chart without deleting it), or remove a series. You can add a series by clicking the plus icon (+) or delete all series by clicking the trashcan icon (IF). Removing a series cannot be undone.

- 8. Under Series Configurations, click **Rename** under Series 1 and give the series a meaningful name that describes the data that will populate the chart.
- 9. Under Series Details, click the Group field and select the property to use for grouping the data in the series. You can use functions to derive the values for the groups. The list below describes the icons and options that are available when choosing a property:
  - The **Root Value** (**•**) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IP) or back (IP) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 10. Next, click the **X** field and select the property to use for the X axis values in the series. You can include functions to derive the values.
- 11. Then select the property to use for the Y axis values by clicking the Y field and choosing a property. You can also include functions.
- 12. If you want to add a group for the series, you can click **Show formulas used for creating calculated series** and add values for the following fields:
  - Series Group: Specifies a property to use for grouping data in addition to the Group value.
  - Series Label: Specifies the property whose values should serve as the Series Group

label. It is typically set to the same value as Series Group.

- Series Sort Direction: Specifies the sort direction for the series groups.
- 13. You can click **Save** any time to render on the dashboard the data that you have configured so far. You can add another series and repeat the steps above. You can also configure several more options for the chart by changing the optional parameters. The additional tabs and options are described below in Optional Configuration Settings.

# **Optional Configuration Settings**

This section describes each of the tabs and additional settings that are available for customizing a scatter chart.

- Chart Data: Plot Tab
- Chart Data: X Axis Tab
- Chart Data: Y Axis Tab
- Chart Data: Filters Tab
- Chart Theme
- Chart Details

### Chart Data: Plot Tab

This tab contains settings that control series formatting, such as data labels, legends, and other display options.
Series De	tails		
Set propertie	es of the	selected series configuration	
Series C Scatter	hart T	уре	>
Series C Plot style i	hart St	tyle on	>
Series C Data labels	Series Chart Data Labels Data labels are set to automatic enablement		>
Series C Markers ar	Series Chart Markers Markers are enabled		>
Show:			
Largest	70	-Automatic-	
Show in lege	end:		
$\checkmark$			
Trend line:			
None			
Data	Plot	X Axis Y Axis Filters	

Field	Description
Series Chart Type	This option opens the Chart Types selection screen where you can select a different chart type for the selected series. The selection must be compatible with the types selected for any other series in the chart, however, or the incompatible series will fail validation and not be displayed on the chart until it is corrected.
Series Chart Style	This option contains the plot style settings that control the color of the dots for the series.
Series Chart Data Labels	This option contains the data label settings, which control the font, alignment, and font effects of the labels for the series.
Series Chart Markers	This option accesses the data marker settings. These settings control the placement and format of the data points for a series that appear when a user hovers the pointer over a dot in the chart. For example, the image below shows a data marker:

Field	Description		
	Series 1 X: 10.06 Y: 9.96		
Show	This setting can be used to limit the data included in the series. The limit can be defined by the largest or smallest $x$ or $y$ axis values.		
Show in legend	This setting controls whether to show the name of the series in the legend.		
Trend line	This setting controls whether to display a linear or exponential trend line for the series.		

## Chart Data: X Axis Tab

This tab defines the formats and labels for the X axis values in the series.

Series Details Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
	_
Title:	
Туре:	_
Linear	
Display values as evenly distributed categories	
Display axis on the opposite side	
Axis Title Details	5
Axis title is Price	_
Axis Labels Axis labels are enabled	>
Axis Style Axis style information	>
Data Plot X Axis Y Axis Filters	

Field	Description
Axis	The Axis and Title settings are populated with the $x$ axis value from the Data tab. If multiple axes exist, you can selected a different value to use for the X axis.
Title	This setting defines the title for the X axis.
Туре	This setting controls the scale for the X axis, linear or logarithmic.
Display values as evenly distributed categories	This setting controls whether to evenly distribute X axis values.
Display axis on the opposite side	This setting can be enabled to move the X axis to the opposite side of the chart.

Field	Description
Axis Title Details	This option contains the X axis title settings, which control the font, alignment, and font effects of the axis title.
Axis Labels	This option contains the X axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option accesses the X axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

## **Chart Data: Y Axis Tab**

This tab defines the formats and labels for the Y axis values in the series.

Series Details Set properties of the selected series configuration	
Create a new axis   Delete current axis	
Axis:	
Title:	_
Туре:	_
Linear	
Display axis on the opposite side	_
Axis Title Details Axis title is Price Earnings	>
Axis Labels Axis labels are enabled	>
Axis Style Axis style information	>
Data Plot X Axis Y Axis Filters	

Field	Description
Axis	The Axis and Title settings are populated with the ${\tt Y}$ axis value from the Data tab. If multiple axes exist, you can selected a different value to use for the Y axis.
Title	This setting defines the title for the Y axis.
Туре	This setting controls the scale for the Y axis, linear or logarithmic.
Display axis on the opposite side	This setting can be enabled to move the Y axis to the opposite side of the chart.
Axis Title Details	This option contains the Y axis title settings, which control the font, alignment, and font effects of the axis title.
Axis Labels	This option contains the Y axis label settings, which control the font, alignment, and font effects of the labels for the axis.
Axis Style	This option accesses the Y axis style settings. These settings control the minimum and maximum values for the axis and the style and position of the grid lines and tick marks.

## **Chart Data: Filters Tab**

This tab defines any filters to apply to the series.

Series Details Set properties of the selected series configuration			
Series Filters No filters specified	>		
X Filters No filters specified	>		
Y Filters No filters specified	>		
Data Plot X Axis Y Axis Filters			

Field	Description
Series Filters	This option can be used to define filters that apply to the entire series.
X Filters	This option can be used to define filters that apply only to the ${\bf X}$ axis values for the series.
Y Filters	This option can be used to define filters that apply only to the ${\tt Y}$ axis values for the series.

# **Chart Theme**

This setting presents options for configuring the theme or color scheme for the chart.



# **Chart Details**

This option offers finer-grained customization settings than the Chart Theme. You can further customize the chart design by adding details such as a chart title and subtitle. You can also modify chart-level styles and fonts as well as legend and tooltip formats.

Chart Details Please specify details about your chart	G 🖄
Title:	
Subtitle:	
Chart Style Chart style information	>
Title Details Title: unspecified	>
Subtitle Details Subtitle: unspecified	>
Legend Details Legend is enabled	>
Tooltip Details Tooltip is enabled	>

### **Creating a Drill Down Lens**

Drill down lenses combine other lenses into a hierarchical interface. Clicking on an object in one lens opens the next lens in successive order. Follow the steps below to create a drill down lens.

 In the dashboard that you want to add a lens to, click Lenses in the main toolbar and select New. The Create Lens dialog box is displayed:

Cre	ate Lens		×
Joint	t the type of lens to create		
۹.	Search for lenses		
+ Path Result	to lens values s: 10		
	Table Tabular visualization of data Properties		
₽/	Form View and edit resource attributes. Properties		
սի	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

2. On the Create Lens dialog box, select **Drill Down**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens			×
General Information			
Description:			
(+) more			
	CANCEL	PREVIOUS	FINISH

3. Type a **Title** and optional **Description** for the lens.

4. Click **Finish**. The lens Designer dialog box is displayed so that you can configure the lens.

Designer	×
Lenses Drag to reorder views.	+ 🖬
Add lanses	Act on filter
CANC	EL SAVE

5. Drill down lenses do not require any property selections or format configurations. Click the plus icon (+) at the top or bottom of the Designer to add the lenses to include in the drill down functionality. The lens listed first becomes the first lens in the hierarchy. Clicking a drill down icon takes you to the next lens. You can drag the lenses in the Designer to change the display order.

Designer		×
Lenses Drag to reorder views.		+ 🗊
Table Disable   Remove		
chart		
👍 Add lenses	Ac	t as filter 📃
	CANCEL	SAVE

6. When you finish adding lenses, click **Save**. Anzo adds the drill down functionality to the dashboard, and you can configure each lens in the hierarchy using the Designer for that lens.

The image below shows a dashboard with drill down functionality. Clicking the drill down icon ( $\nabla$ ) in the left column displays a chart lens, which shows details about the venue for that event:

events		
	Eventid	Eventname
Ŷ	1	Gotterdammerung
Y	2	Boris Godunov
Ŷ	3	Salome
Ŷ	4	La Cenerentola (Cinderella)
Y	5	II Trovatore
Υ	6	L Elisir d Amore
Υ	7	Doctor Atomic
Ŷ	8	The Magic Flute
Y	9	The Fly
<i>∀</i>	10	Rigoletto
Y	11	Doctor Atomic
$\nabla$	12	Ring Cycle
Υ	13	Lucia di Lammermoor
Υ	14	La Rondine
Υ	15	Die Walkure
Y	16	La Gioconda
Y	17	La Gioconda
Ÿ	18	Gotterdammerung

### **Creating a Form Lens**

Form lenses enable you to create an editable or read-only form on the dashboard. Creating forms can be useful for displaying many details about each record instead of using a table where the large number of columns makes the data hard to read. Follow the steps below to create a Form lens.

### Important

By default, only the sysadmin user has access to create Form lenses. In addition, Form lenses are valid in only in Linked Dataset dashboards. They do not display data for graphmart dashboards. To create a Linked Dataset dashboard, select the **Show advanced dashboards** checkbox when creating a dashboard, and then select **Linked Dataset Dashboard**.

 In the dashboard that you want to add a lens to, click Lenses in the main toolbar and select New. The Create Lens dialog box is displayed:

Cre	ate Lens		×
٩.	Search for lenses		
+ Path Result	to lens values s: 10		
	Table Tabular visualization of data Properties		
₽/	Form View and edit resource attributes. Properties		
dı	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

2. On the Create Lens dialog box, select **Form**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens			×
General Information			
Description:			
[+] more			
	CANCEL	PREVIOUS	FINISH

3. Type a **Title** and optional **Description** for the lens.

4. Click **Finish**. The lens Designer dialog box is displayed so that you can configure the lens.

Designer		×
form		
Content Field Settings Form Settings		
∧ Data		
Please select a property		
Q Search for properties		
Results: 7 O>> Movies <sup>1</sup> 2 <sub>3</sub> Movie ID - no description - T Movie Title - no description - O Type The subject is an instance of a class.		
O≪Movie Actors		
✓ Layout		
Show advanced C	ANCEL	SAVE

5. On the Content tab, drag onto the dashboard each property or relative path that you want to appear as a field on the form. After adding objects, you can rearrange the form layout and use the Field Settings tab to further configure each field.

💠 Designer 🛛 😝 Properties	Des	igner		×
Maximize   Snap to original	form	1		
*Movie ID	Co	ontent Field S	Settings Form Settings	
	Sea	arch		Q
	-	General		
*Movie Title		control	Label - Read Only	•
		label	*Acted In/Actor Name	
Acted III/Actor Name		Separator	,	
	-	Layout		
		height		рх
		width	100	%
	-	Advanced		
		class		
		id		
	⊖ Sho	ow advanced	CANCEL	SAVE

If you want to arrange the fields in a different layout, such as a two-column layout, click
 Layout below the list of properties on the Content tab. The Designer displays the available
 layout containers.



- 7. Drag a container onto the form to create the layout template. You can then drag properties into the template.
- 8. Click **Save** to save the configuration and add the lens to the dashboard.

### **Creating a List Lens**

List lenses display the values for the selected property in a list layout with icons, similar to a directory explorer view. Follow the steps below to create a List lens.

 In the dashboard that you want to add a lens to, click Lenses in the main toolbar and select New. The Create Lens dialog box is displayed:

Cre Selec	ate Lens t the type of lens to create		×
٩.	Search for lenses		
+ Path Result	to lens values s: 10		
	Table Tabular visualization of data Properties		
₽/	Form View and edit resource attributes. Properties		
dı	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

2. On the Create Lens dialog box, select **List**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens			×
General Information			
Description:			
(+) more			
	CANCEL	PREVIOUS	FINISH

3. Type a **Title** and optional **Description** for the lens.

4. Click **Finish**. The lens Designer dialog box is displayed so that you can configure the lens.

Design	er		×
Title —			
Path*	Click to select a path		0
Default			
Subtitle			
Path	Click to select a path		
Default			
Icon —			
Path	Click to select a path		
Default			
Hyperlink			
Preview	changes	CANCEL	SAVE

- 5. In the **Path** field under Title, select the property that contains the values you want to display in the list.
- 6. If you want to include a subtitle for each value in the list, click the **Path** field under Subtitle and select the property to supply the subtitle values. For example, in the image below, Actor Name is chosen as the Title with Movie Name as the subtitle.

🌣 Designer 🛛 📄 Properties		
What can I do next? $\times$ list $\times$		
View: Medium Icons Sort by: Title	Direction: Ascending	
Aasha Davis, Kim Wayans,	S Abbie Cornish, Jena Malon	J Abigail Breslin, Max Thieri
🚰 Abigail Breslin, Michelle Pf	Series Adam Baldwin, Patrick Sw	🚰 Adam Beach, Nicolas Cage
🚰 Adam Brody, Kerry Washin	Series Adam Campbell (actor)	🚰 Adam Carolla, David Alan
F Adam Rich, Bill Cosby, Juli	Jadam Sandler, CeeLo Gree	🚰 Adam Wade (singer), Roxie
F Adetokumboh M'Cormack,	🚰 Adolfo Celi, Genevi�ve Bu	🚰 Adriana Barraza, Radha Mi
🚰 Adriano Giannini, Madonna	Adrienne Shelly, Arye Gros	Series Aisha Tyler, Constance Zim
🚰 AJ Michalka, James Dento	🚰 Ajay Naidu, Michael Keato	
	I< < > >I	1 - 20 of 8157 Page 1 of 408 🌣 👻

- 7. If you have a property that contains an icon to use in place of the default question mark icons, click the Path field under lcon and select the property.
- 8. If you want the Title to be formatted as a hyperlink to another lens, you can click **Hyperlink** and select the lens to link to.
- 9. When you have finished configuring the lens, click **Save** to save the configuration and add the lens to the dashboard.

### **Creating a Table Lens**

Table lenses display data in a standard column and row grid layout. Follow the steps below to create a table lens.

### Note

The list below describes the default display formats for date and numeric values in tables.

• **Date**: By default Anzo displays date values in "short" date format. The order of the month, day, and year depends on the location of your browser. For example, in the

United States the default date format is MM/DD/YYYY. In Australia, the default date format is DD/MM/YYYY. Note that this is not dependent on the Anzo server location but on the location auto-detected by the browser.

- Numeric: Anzo displays the complete value without a limit on precision. Numeric formats are also dependent on the location of the browser. For example, in the United States the default format for a large number is 4,294,967,295.00 and in Canada the default format is 4,294,967,295,000.
- In the dashboard that you want to add a lens to, click Lenses in the main toolbar and select New. The Create Lens dialog box is displayed:

Cre	ate Lens		×
Selec	trane type of tens to create		
٩.5	Search for lenses		
+ Path Results	to lens values s: 10		
	Table Tabular visualization of data Properties		
₽⁄	Form View and edit resource attributes. Properties		
ılı	Chart Bar, scatter, pie, etc Properties		
	AnzoKO Web Page A user-defined web page. Properties		
		CANCEL	NEXT

2. On the Create Lens dialog box, select **Table**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens			×
General Information			
Description:			
(*) more			
	CANCEL	PREVIOUS	FINISH

- 3. Type a **Title** and optional **Description** for the lens.
- 4. Click Finish. The lens Designer dialog box is displayed so that you can configure the lens.

olumns + 🗊 🛷 🌂 st of column configurations.	Column Details     Set properties of the selected column.
Untitled Disable   View   Clone Remove	Column Header Label:
	Column Value Expression:
	Click to edit
	Format   Hyperlink   Edito
	Column Footer Expression:
	Column Width:
	Content Filters
	Group raws by:
	Children and M

5. To add a column to the table, click the **Column Value Expression** field and select the property name or calculation to use to populate the values in the column. For information about calculating values, see Calculating Values in Lenses and Filters.

Designer		×
Columns + 🗊 🦑 🦄 List of column configurations.	Column Details Set properties of the selected column.	
Untitled Disable   View   Clone Remove	Column Header Label:	
	Column Value Expression: Please select values	
	fx •	¥ ≣≣
	Q Search for properties	ik
	Results: 8  Root Value  - no description -	
	O» Movies	
	G - no description -	-
Add a column   Display Properties	T Movie Title - no description - Close	

Tip

As a shortcut to browse the available columns, you can click the Auto-generate columns icon () to add all properties (for the data type selected on the dashboard) as columns. Then you can hover the pointer over columns to remove the ones you do not want to keep. For example:

st of column configurations.	Set properties of the selected column.	
CustomerID Disable   View   Clone Remove	Column Header Label:	
Employee ID	Column Value Expression: CustomerID	
🔢 Freight	Forma	at   Hyperlink   Edito
🔢 OrderDate	Column Footer Expression: Click to edit	
OrderID Disable   View   Clone Remove	Column Width:	Format   Hyperlin
RequiredDate		
ShipAddress	Content Filters	
ShipCity	Group rows by:	
A solumn   Display Properties	Click to edit	
Preview changes	CANC	CEL SAV

- 6. You have the option to create a filter or configure any of the following optional fields:
  - **Column Header Label**: The column name to display. Overrides the Column Value Expression property name.
  - Column Footer Expression: The property to use to supply the table footer.
  - Column Width: The width of the column in pixels.
  - Group rows by: The property to group data by.
  - Filters Tab: You can click **Create filter** to add a filter on the column. For more information, see Working with Filters.
- 7. Click **Save** to save the configuration and add the lens to the dashboard.

### Advanced Lenses

The topics in this section provide information about creating advanced dashboard lenses. Creating these types of lenses require users to advanced SPARQL query language skills or coding skills in the areas of HTML, CSS, and JavaScript. For information about Network Navigator lenses, see Creating a Network Navigator Dashboard.

- Creating an AnzoKO Web Page Lens
- Creating a Query Lens
- Creating a Resource Tree Navigator Lens
- Creating a Web Page Lens

### Creating an AnzoKO Web Page Lens

The custom AnzoKO Web Page lens includes the <u>Knockout JavaScript</u> framework and enables you to create visualizations of RDF resources and metadata using knockout.js-like syntax without needing to write additional JavaScript to declare which parts of the data to render in which sections of the HTML. Follow the steps below to create an AnzoKO Web Page Lens.

 In the dashboard that you want to add a lens to, click Lenses in the main toolbar and select New. The Create Lens dialog box is displayed:



2. On the Create Lens dialog box, select **AnzoKO Web Page**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens			×
General Information			
Description:			
[*] more			
	CANCEL	PREVIOUS	FINISH

- 3. Type a **Title** and optional **Description** for the lens.
- 4. Click **Finish**. The lens Designer dialog box is displayed so that you can configure the lens.

Designer	×
View HTML CSS Javascript	
Font ■ Z U ∓ E Ξ Ξ   Font ■ Size x-small ■ A =	
Binding to input <u>uti</u> : Name:	
Preview changes     CANCEL     PREVIEW     SA	VE

5. Configure the lens and then click **Save** to save the configuration and add the lens to the dashboard.

### **Creating a Query Lens**

The query lens allows you to retrieve and display data using custom a SPARQL query. You format the query results using HTML and CSS. This lens can access external SPARQL-compatible data sources. See SPARQL Best Practices and Query Templates for guidance on writing SPARQL queries.

# **Query Lens Configuration**

The Query lens Designer has three tabs:

Designer		>
Query HTML CSS		
<pre>1 PREFIX dc: <http: 1.1="" dc="" elements="" purl.org=""></http:></pre>		
2 PREFIX 1d: <http: 05<="" 2009="" cambridgesemantics.com="" ontologies="" p=""></http:>	/LinkedData#:	>
3 PREFIX owl: <a href="http://www.w3.org/2002/07/owl#">http://www.w3.org/2002/07/owl#&gt;</a>		
<pre>4 PREFIX rdf: <http: 02="" 1999="" 22-rdf-syntax-ns#="" www.w3.org=""></http:></pre>	*	
5 PREFIX rdfs: <a href="http://www.w3.org/2000/01/rdf-schema#">http://www.w3.org/2000/01/rdf-schema#</a> >		
6 PREFIX anzo: <a href="http://openanzo.org/ontologies/2008/07/Anzo#">http://openanzo.org/ontologies/2008/07/Anzo#</a>	Þ	
7		
8 SELECT DISTINCT ?value ?p ?o		
9 WHERE {		
10 ?value ?p ?o .		
11 <%mixins%>		
12 }		
13 ORDER BY (%orderD1r%>(<%orderBy%>)		
14 OFFSET (AOTTSETA)		
16 CINII CALINICA		
10		
Preview changes	CANCEL S	AVE

- **Query**: This tab displays a SPARQL query template that you can use to write the query. Note the default code that reflects inherent functionality:
  - <%mixins%>: Incorporates a filter function.
  - **ORDER BY**: Incorporates a sort function.
- HTML: This tab includes default HTML and basic JavaScript code with sample values. You can edit the content to design the results that the query returns. The default HTML code automatically adds returned query data to a table and organizes it so that new rows are created for each record. Make sure that the <option> elements correspond to the elements in your query.
- CSS: This tab enables you to create a cascading style sheet to format the HTML and define the look and feel of the lens. Cambridge Semantics recommends that you define all CSS classes as namespaces to avoid global format changes.

### **Creating a Resource Tree Navigator Lens**

The Resource Tree Navigator lens displays data in a tree format with points that you can click to open successive child data points. Follow the steps below to create a lens.

### Note

By default, the only user who has permission to create a Resource Tree Navigator lens is the **sysadmin** user. However, the sysadmin user can share created lenses with other users and groups (see Sharing Access to Dashboards and Lenses).

 In the dashboard that you want to add a lens to, click Lenses in the main toolbar and select New. The Create Lens dialog box is displayed:



2. On the Create Lens dialog box, select **Resource Tree Navigator**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens		×
General Information Title.*		
Description:		
el more		
( ) market		

- 3. Type a **Title** and optional **Description** for the lens.
- 4. Click **Finish**. The lens Designer dialog box is displayed so that you can configure the lens. The Designer displays all of the properties for the linked classes. Select each property that

you want the resource tree to include.

Designer		×
Q Search for properties		
Results: 99  Categories  - no description -		
• Suppliers - no description -		
• Type The subject is an instance of a class.		
O>> Products By Category		
O Untitled Selection: 1		
Categories 🕲		
Preview changes	CANCEL	SAVE

5. Click **Save** to save the configuration and add the lens to the dashboard. The image below shows the information that is displayed.



- 1. **Related data**: Displays the class data related to the selected data property. Data changes when another class is selected.
- 2. **Class property**: Displays the label property of the target class as the initial (start) point of the resource tree. Expand the tree to view child properties by clicking the plus icon for a data point.

- 3. **Selected linked property**: Displays the initial selected property that links to other classes.
- 4. **Navigation tools**: Use the arrows to navigate to other pages. The Showing text box displays the current page number and total number of pages.

### **Creating a Web Page Lens**

Web Page lenses enable you to display data by creating a web page using HTML, CSS, and JavaScript. This lens is for advanced users with coding skills in these areas. A powerful feature of the this lens is the ability to bind data to Anzo graphs so that updates are reflected in real time. Follow the steps below to create a Web Page lens.

 In the dashboard that you want to add a lens to, click Lenses in the main toolbar and select New. The Create Lens dialog box is displayed:



2. On the Create Lens dialog box, select **Web Page**, and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens			×
General Information			
Description:			
[4] more			
	CANCEL	PREVIOUS	FINISH

- 3. Type a **Title** and optional **Description** for the lens.
- 4. Click **Finish**. The lens Designer dialog box is displayed so that you can configure the lens.

Designer	×
View HTML CSS Javascript	
▷ ~   B I 및 주   톤 푼 Ξ	
Font   Size x-small	<u>A</u> -
A • B • 60 •	
Preview changes     CANCEL	SAVE

5. Configure the lens and then click **Save** to save the configuration and add the lens to the dashboard. See Web Page Lens Configuration below for information about the configuration.

# Web Page Lens Configuration

The Web Page Designer has four tabs:

View     HTML     CSS     Javascript       Image: Size     Image: Size     Image: Size     Image: Size       Font     Image: Size     Image: Size     Image: Size       Image: Size     Image: Size     Image: Size     Image: Size       Image: Size     Image: Size     Image: Size     Image: Size	Designer				×
Image: Size     Image: Size       Image: Size     Image: Size       Image: Size     Image: Size       Image: Size     Image: Size	View HTML	CSS Javascript	t		
	Font B	·⊻ ∓   ≣ •	≣ ≣   Size x-small	-	<u> </u>
		63 •			
Preview changes     CANCEL SAVE	Preview changes			CANCEL	SAVE

- View: Provides a rich text interface for viewing the page (WYSIWYG). Changes made to this page are reflected in the HTML code.
- **HTML**: This tab enables HTML coding and data binding. The example HTML image below shows code that defines text format as well as data binding using the anzowbind:innerhtml command.

Designer	×
View HTML CSS Javascript	
<pre>1     <div style="text-align: center;"><b>         <i><font face="Times New Roman" size="4">My excellent lens</font></i></b></div>         <div xmlns:film="http://cambridgesemantics.com/ontologies/2009/08/Film\$">         <gpan anzowtbind:innerhtml="film:mpName"> </gpan></div></pre>	
Preview changes     CANCEL	SAVE

For more information about data binding, see the Data Binding Example section below.

- CSS: This tab enables you to create a cascading style sheet to format the HTML and define the look and feel of the web page. Cambridge Semantics recommends that you define all CSS classes as namespaces to avoid global format changes.
- Javascript: This tab enables you to write JavaScript code to implement functions such as if statements, animations, or event notifications.

## Data Binding Example

When data is bound to a web page lens using HTML code, the web page lens behaves as follows:

- The lens will reflects data changes in real time.
- If the lens is oriented to the left-hand column (using the Orientation drop-down), selecting data in an active lens prompts the web page lens to display the related data.

In the example below, the active table lens row is selected, prompting the web page lens on the left ("My excellent lens") to display the corresponding data.



## **Cloning a Lens**

Cloning a lens makes a copy of the lens that can be changed without affecting the original lens or other dashboards. Follow the steps below to clone a lens.

### Note

You can only clone lenses from dashboards that you have permission to modify. If you open a dashboard with read-only access, the clone options are not available. To clone a lens from a read-only dashboard, save a copy of the dashboard so that you become the owner of the copy. To save a copy, click the **Dashboard** button in the main toolbar and select **Save As**. Then follow the procedure below to clone a lens into the dashboard that you own.

 Open a dashboard in the Hi-Res Analytics application, then click Lenses in the main toolbar and select Open. Anzo opens the Lens Selection dialog box, which lists the lenses that are available to open. For example:

Lens Selection		×
Q Search for lenses		
Path to lens values Results: 4     Airlines     - no description -     Properties   Delete   Clone     Airports     - no description -     .		
Properties   Delete   Clone  Properties   Delete   Clone  Selection: 0		
Create a new lens	CANCEL	ок

2. Click the **Clone** link for the lens that you want to clone. Anzo displays the Clone lens dialog box and populates the Title field with the existing lens name and "(clone)." For example:

×

- 3. Modify the **Title** to name the new copy of the lens, and add or change the **Description** if necessary. Then click **OK**.
- 4. Anzo adds the new copy of the lens to the Lens Selection dialog box and selects it. Click **OK** to add the lens to the dashboard.

### **Exporting a Lens**

If you have dashboards with Table, Chart, and/or Query lenses, you can export those lenses from the Hi-Res Analytics application. Charts can be exported as images in JPEG, PNG, SVG, or PDF format. Tables can be exported to CSV or JSON files. And Query lenses can be exported to CSV files. Follow the instructions below to export a lens.

- 1. Open the dashboard that contains the lens that you want to export.
- 2. If necessary click the tab for the lens to make it active. For example, the image below shows a chart lens.



3. In the object toolbar for the lens, click the **Export** button.



For Query lenses, adjust the Export options as needed, and then click **Export Results**. For example:



4. If the lens is a chart, select the one of the image types from the drop-down list. Anzo creates the image as that type and downloads the file to your computer.

5. If the lens is a table, the Export Options dialog box is displayed:

Export Options	×
File name: Format: CSV Multi Valued Column Separator: [pipe] Export Headers Column to Export	
Eoscription	I ON
I Director	ON ON
≣ Movie Cast	Ø ON
I Number Of Reviews	🗹 ON
Popularity Index	M ON
∏ Rating	Ø ON
≣ Title	🗹 ON
CA	ANCEL OK

- 6. In the Export Options dialog box, specify the following file options:
  - File name: Specify a name for the file. Do not specify the file type extension.
  - Format: Click the Format field and select CSV to create a .csv file or JSON to create a .json file.
  - Multi Valued Column Separator: For CSV files, click this column to select the character to use as a separator in the file. This option does not apply to JSON files.
  - **Export Headers**: Indicates whether to include column headers in the file. Clear the checkbox to exclude headers from the file. This option does not apply to JSON files.
  - **Columns to Export**: By default, all columns are selected for export. You can clear the checkboxes next to columns that you do not want to include. In addition, you can reorder columns by dragging and dropping the rows.
- 7. Click **OK** to download the file to your computer.

## **Deleting a Lens**

This topic provides information about the permissions that are required for deleting lenses as well as instructions for deleting a lens.

### Note

By default, only the **sysadmin** user and **lens creator** have permission to delete a lens. To delete a lens, a user must have the **Manage** permission assigned for that lens. The Manage permission is included in the **Admin** predefined lens permission set. If lens permissions have not been changed since the lens was created, the sysadmin user and the lens creator are the only users who have permission to delete that lens. The Manage permission is also required to change lens security settings and grant privileges to other users. Users who have read access to a lens (granted through the View, Modify, or Admin lens permission sets) can view the lens security settings to identify which non-sysadmin users have permission to delete the lens. For more information, see Sharing Access to Dashboards and Lenses.

Follow the steps below to delete a lens.

- 1. Open a dashboard that contains the lens that you want to delete.
- In the Hi-Res Analytics application, click the Lenses menu in the main toolbar and select
   Open. The Lens Selection dialog box is displayed.
- 3. In the Lens Selection dialog box, find the lens that you want to delete and then click the **Delete** link for that lens. For example:

Q Search for lenses	
+ Path to lens values Results: 3	
E Actors - no description - Properties   <u>Delete</u>   Clone	
Actors           - no description -           Properties   Delete   Clone	
All Actors - no description - Properties   Delete   Clone	
Selection: 0	
Create a new lens	
CANCEL OK	

4. The application presents a confirmation message. Click **Yes** to delete the lens. The lens is removed from the Lens Selection dialog box, and you can repeat this process to delete additional lenses for which you have the required privileges.

## **Working with Filters**

Filters narrow the data presented in a dashboard. You can define filter criteria using Microsoft Excel-like functions such as AVG, SUM, or UPPER, or groupings such as a date range or aggregation. Though you can filter data in some lens types, such as tables and charts, when you add a filter to a dashboard, all lenses on the dashboard update simultaneously based on your filter selection. Unlike lenses, filters cannot be shared by other users or dashboards and must be created for each dashboard.

The topics in this section provide instructions for creating and configuring each type of filter.

- Adding a Cloud Filter
- Adding a Date Range Filter
- Adding a Hierarchy Filter
- Adding a Limit Filter
- Adding a List Filter
- Adding a Numeric Range Filter
- Adding a Presence Filter
- Adding a Quartile Filter
- Adding a Range Slider Filter
- Adding a Relative Time Filter
- Adding a Search Filter
- Adding a Single Select List Filter
- Adding a Types Filter

### **Adding a Cloud Filter**

Cloud filters display values in term clouds where each term is written in a font size that represents the number of results for that value. Unlike list filters, which enable you to select and filter on multiple values at once, cloud filters allow you to filter on one value at a time. The cloud filter is available for all data types. Follow the instructions below to create a cloud filter.

- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.

Create Filter	×
Fields:*	
Click to edit	
Filter:*	
	ANCEL OK

3. Click in the **Fields** field to open the property drop-down list and determine the property to filter on. The values for this property will be the terms that are displayed in the cloud. The list of available properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about tickets sold for various types of events. The type for the dashboard is tickit\_events:

Create Filter	×
Fields:*	
Please select values	
fe -	<b>1</b> 7
Q Bearch for properties	
Results: 13	
Root Value     no description -	
O>> tickit_events	- 11
Category     - no description -	
<sup>1</sup> 2 <sub>3</sub> Catid - no description -	₽
<sup>1</sup> 2 <sub>3</sub> Dateid - no description -	Þ
Close	11

The list below describes the icons and options that are available when choosing a property:

 The Root Value (n) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.
### Tip

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IP) or back (IP) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 4. After you have selected the property to filter on, click Close to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **Cloud** from the drop-down list. The dialog box is refreshed to show the Filter Properties and other options that are available for the filter type:

Create Filter	>	<
Fields:*		
venue/venuecity		
Filter.*		
Cloud		
Filter Properties		
Title:		
Label field:		
Click to edit		
	Exclude:	
	Show counts: 🛃 Respond to other filters: 🛃	
Subfilters + Create Filter		
	CANCEL OK	

- 6. Configure any of the following properties. All of the fields are optional:
  - **Title**: Defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the Fields value is used as the title.
  - Label Field: If you want to populate the cloud with values from a property other than the one specified in Fields, you can select an alternate property in this field.
  - **Exclude**: This setting controls whether selecting a term in the filter narrows the results to show only the records that *include* that term or whether selecting a term *excludes* the records that include that term. When Exclude is disabled, selecting a term narrows the dashboard results to show only the records that include that term. When Exclude is enabled, selecting a term filters out all of the records that include that term.
  - **Show Counts**: This setting controls whether the number of results for each term are displayed when you hover the pointer over a term.
  - **Respond to Other Filters**: This setting controls whether the results of this filter change based on the selections made in other filters on the dashboard.
- 7. If you would like to be able to further constrain the data that appears in the filter, you can add one or more subfilters. To add a subfilter, click **Create Filter** under Subfilters. The process of creating a subfilter is the same as the process for the parent filter. However, the subfilter is not

displayed on the dashboard. It is visible only when editing the parent filter, and the subfilter's configuration affects only the parent filter and any sibling subfilters.

8. When you have finished configuring the filter, click **OK** to add it to the dashboard. The new filter appears in the left pane of the dashboard and displays the values that are available for filtering the data.

For example, the filter in the image below shows cities with event venues. The size of the terms represent the number of events that were held in that city.



Depending on whether the Exclude option is enabled or disabled, clicking a term in the cloud refreshes the dashboard to show only the data that either contains or excludes the selected term.

When working with the filter on the dashboard, the following options are available for sorting and configuration:

- **show/hide sort details**: Shows or hides the following options for sorting the results in the filter:
  - Sort by: Select Count to order results according to the total number of results for each value, or select Value to sort string values alphabetically.
  - Direction: Select Ascending to order results in alphabetical order. Or select
     Descending to order results in reverse order.
- **show/hide filters**: This option is displayed when a term is selected in the cloud. It shows or hides the selection.

- Menu (1): The menu contains the following options:
  - Select All Visible: This option does not work for cloud filters.
  - Clear: This option becomes available when a term is selected. Clicking Clear removes the selection.
  - Designer: Selecting this option opens the designer so that you can view or change the filter configuration.
- Close (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

# Adding a Date Range Filter

Date range filters are used to limit the results on a dashboard to data that falls in (or outside of) certain date and time groupings. Date range filters are available for properties with date, dateTime, and time data types. Follow the instructions below to create a date range filter.

- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.

Create Filter	×
Fields:*	
Click to edit	
Filter:*	
CANCEL	OK

3. Click in the **Fields** field to open the property drop-down list and determine the date, dateTime, or time type property to filter on. The values for this property will be used to determine the date ranges for the filter. The list of available properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about tickets sold for various types of events. The type for the dashboard is tickit events:

Create Filter	×
Fields:*	
Please select values	
fe •	<b>1</b> 7
Q Search for properties	
Results: 13	
no description -	- In
O>> tickit_events	
Category     - no description -	
<sup>1</sup> 2 <sub>3</sub> Catid - no description -	÷
<sup>1</sup> 2 <sub>3</sub> Dateid - no description -	Þ
Close	1

The list below describes the icons and options that are available when choosing a property:

• The **Root Value** (**h**) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

### Tip

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons. The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties. For this example filter, selecting EventDate, navigates to the tickit\_dates class where a date type property can be selected.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IF) or back (IF) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on

the data type of the selected property. To choose a more advanced function or type a formula, click **Advanced**. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.

- 4. After you have selected the property to filter on, click **Close** to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **Date Range** from the drop-down list. The dialog box is refreshed to show the Filter Properties, Format, and other options that are available for the filter type:

Create Filter		×
Fields:*		
eventDate/caldate		
Filter:*		_
Date Range		
Filter Properties		
Title:		
Interval Unit:*		
MILLENNIUM		
Interval:*		
	Exclude: Show Bars: Always Show Checkboxes:	~
	Show counts: 🛃 Respond to other filters:	~
Format		
Format Type:		
Automatic		
Subfilters		
+ Create Filter		
	CANCEL	ж

- 6. Configure any of the following properties. **Interval Unit** and **Interval** are required fields, and the rest of the fields are optional:
  - **Title**: Defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the Fields value is used as the title.
  - Label Field: If you want to populate the ranges with a label other than the one specified in Fields, you can select an alternate property in this field.

- Interval Unit: Click this field to choose the unit of time for the Interval value. Depending on the data type of the selected property, a subset of the following values are available to choose from: Millennium, Century, Decade, Year, Month, Week, Day, Hour, Minute, or Second.
- Interval: This setting specifies a number that defines the length of time in each grouping.
   For example, if the Interval Unit is "Decade," an Interval value of 2 creates groups of twodecade increments.
- **Exclude**: This setting controls whether selecting a range in the filter narrows the results to show only the records that are *included* in that range or whether selecting a range *excludes* the records that fall in that range. When Exclude is disabled, selecting a range narrows the dashboard results to show only the records that fall in that range. When Exclude is enabled, selecting a range filters out all of the records that fall in that range.
- **Show Bars**: This setting controls whether the total values for the selected property appear as a bar graphic in the background of the filter.
- Always Show Checkboxes: This setting controls whether checkboxes are shown next to the items in the filter.
- **Show Counts**: This setting controls whether the number of results for each range are displayed in parentheses next to the range.
- **Respond to Other Filters**: This setting controls whether the results of this filter change based on the selections made in other filters on the dashboard.
- If you would like to specify the format to for date values that are displayed in the filter, click the Format Time field and select a format from the drop-down list.
- 8. If you would like to be able to further constrain the data that appears in the filter, you can add one or more subfilters. To add a subfilter, click **Create Filter** under Subfilters. The process of creating a subfilter is the same as the process for the parent filter. However, the subfilter is not displayed on the dashboard. It is visible only when editing the parent filter, and the subfilter's configuration affects only the parent filter and any sibling subfilters.

9. When you have finished configuring the filter, click **OK** to add it to the dashboard. The new filter appears in the left pane and displays the values that are available for filtering the data.

For example, the filter in the image below shows date ranges for events. The Interval Unit is Month, and the Interval is 3 months.

∧ Event Date	:	$  \times$
[+] show sort details		
1/1/2008 - 1/31/2008 (14)		
1/31/2008 - 3/1/2008 (11)		
<b>3/1/2008 - 3/31/2008</b> (14)		
<b>3/31/2008 - 4/30/2008</b> (12)		
<b>4/30/2008 - 5/30/2008</b> (11)		
<b>5/30/2008 - 6/29/2008</b> (12)		
<b>6/29/2008 - 7/29/2008</b> (15)		
7/29/2008 - 8/28/2008 (12)		
8/28/2008 - 9/27/2008 (11)		
<b>9/27/2008 - 10/27/2008</b> (11)		
13 choices   more		

Depending on whether the Exclude option is enabled or disabled, clicking a range in the filter refreshes the dashboard to show only the data that is in the selected range or only the data that is outside of the range.

When working with the filter on the dashboard, the following options are available for sorting and configuration:

- show/hide sort details: Shows or hides the following option for sorting the results in the filter:
  - Direction: This option controls how you want to order the date ranges in the filter, depending on the Format specified for the values in the filter. For number values, Count Ascending orders results from the earliest to latest date and Count Descending orders results from the latest to earliest date. For character values, Name Ascending orders results in alphabetical order and Name Descending orders results in reverse alphabetical order.
- show/hide filters: This option is displayed when a range is selected. It shows or hides the selection.

- Menu (1): The menu contains the following options:
  - <sup>o</sup> **Select All Visible**: This option selects all of the ranges that are listed in the filter.
  - Clear: This option is available when one or more ranges are selected. Clicking Clear removes the selection..
  - Designer: Selecting this option opens the designer so that you can view or change the filter configuration.
- Close (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

# Adding a Hierarchy Filter

If hierarchies exist in your knowledge graph, you can create a hierarchy filter to explore the parent and child relationships and filter the dashboard based on the relationships. Unlike the majority of dashboard filters, where you select a property to filter on, hierarchy filters operate on relationships and are only available as a filter type when you select a path to filter on.

In order to produce hierarchies in the data, you typically need a self-referential data model, where properties have relationships to themselves. The SKOS ontology is a good example of a self-referential model. As shown in the image below, many of the properties in the Concept class refer to themselves. These paths define a hierarchy.



Follow the instructions below to create a hierarchy filter.

- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.

Create Filter		×
Fields:*		
Click to edit		
Filter:*		
	CANCEL	OK

 Click in the Fields field to open the property drop-down list and determine the path to filter on. The list of available classes, paths, and properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about protein classification. The type for the dashboard is Classification:

Create Filter	×
Fields.*	
Please select values	
fe 🔹 \leftrightarrow	Ŵ?
Q Search for properties	-
Root Value - no description -	
O» Classification	'
<ul> <li>Child         <ul> <li>no description -</li> </ul> </li> </ul>	
T ClassificationName	
<ul> <li>Type The subject is an instance of a class.</li> </ul>	
O≪Cluster	
<ul> <li>Classification</li> <li>- no description -</li> </ul>	- 1
Close	

The list below describes the icons and options that are available when choosing a path:

• The **Root Value** (**•**) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard. Typically **Root Value** is

the chosen Field for Hierarchy filters.

### Tip

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IP) or back (IP) arrows to go forward or backward one or more levels at a time.
- After you have selected a path, you can apply a function or formula to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that become available depend on the data type of the selected path. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 4. After you have selected the path to filter on, click **Close** to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **Hierarchy** from the drop-down list. The dialog box is refreshed to show the Filter Properties and other options that are available for the filter type:

Create Filter		>
Fields:*		
ROOT		
Filter:*		
Hierarchy		
Filter Properties		
Title:		
	Show counts: 🛃 Respond to other filters:	~
Label Field:*		
Click to select a path		
Children Field:*		
Click to select a path		
Subfilters		_
+ Create Filter		
	CANCEL	)K

- 6. Configure any of the following properties. **Label Field** and **Children Field** are required fields, and the rest of the fields are optional:
  - **Title**: Defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the Fields value is used as the title.
  - **Show Counts**: This setting controls whether the number of results for each filter value are displayed in parentheses next to the value.
  - **Respond to Other Filters**: This setting controls whether the results of this filter change based on the selections made in other filters on the dashboard.
  - Label Field: Species the property that supplies the label for the parent values in the hierarchy.
  - Children Field: Specifies the child value in the relationship.
- 7. If you would like to be able to further constrain the data that appears in the filter, you can add one or more subfilters. To add a subfilter, click **Create Filter** under Subfilters. The process of creating a subfilter is the same as the process for the parent filter. However, the subfilter is not displayed on the dashboard. It is visible only when editing the parent filter, and the subfilter's configuration affects only the parent filter and any sibling subfilters.

8. When you have finished configuring the filter, click **OK** to add it to the dashboard. The new filter appears in the left pane of the dashboard and displays the values that are available for filtering the displayed data.

For example, the filter in the image below shows a hierarchy of classifications.



Selecting checkboxes in the filter refreshes the dashboard to show only the data that includes the selected values.

When working with the filter on the dashboard, the following options are available:

- Menu (1): The menu contains the following options:
  - Clear: This option becomes available when an item is selected. Clicking Clear removes the selection.
  - **Designer**: Selecting this option opens the designer so that you can view or change the filter configuration.
- Close (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

# Adding a Limit Filter

Limit filters are used to limit the results on the dashboard to a specified number of either the largest or smallest values. The limit filter is available for any data type. For strings, results are ordered alphabetically. "Largest" orders by the last letters in the alphabet and "Smallest" orders by the first letters in the alphabet. Follow the instructions below to create a limit filter.

- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.

Create Filter	×
Fields:*	
Click to edit	
Filter.*	
C/	ANCEL OK

3. Click in the **Fields** field to open the property drop-down list and determine the property to filter on. The list of available properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about tickets sold for various types of events. The type for the dashboard is tickit\_events:

Create Filter	×
Fields:*	
Please select values	
fe -	<b>1</b> 7
Q [Search for properties	
Results: 13	
Root Value - no description -	- It
O>> tickit_events	
Category     -no description -	
<sup>1</sup> 2 <sub>3</sub> Catid - no description -	₽
<sup>1</sup> 2 <sub>3</sub> Dateid - no description -	Þ
Close	1

The list below describes the icons and options that are available when choosing a property:

• The **Root Value** (**h**) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons. The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties. For this example filter, selecting the Venue path navigates to the tickit\_venues class where an integer type property, Venueseats, is selected.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IF) or back (IF) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 4. After you have selected the property to filter on, click **Close** to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **Limit** from the drop-down list. The dialog box is refreshed to show the Filter Properties and other options that are available for the filter type:

Create Filter		×
Fields:*		
venue/venueseats		
Filter:*		
Limit		
Filter Properties		
	Limit by res	ource: 🗸
Subfilters + Create Filter		
	CANCEL	ок

- 6. Configure any of the following properties. All of the fields are optional:
  - **Title**: Defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the Fields value is used as the title.
  - Limit by Resource: This setting controls whether the limit is also applied to the resource that is the data type of the dashboard. When Limit by Resource is enabled, the specified limit applies to the resource as well as the property. When Limit by Resource is disabled, the limit applies only to the specified property. Using the example above for the property of Venueseats and the resource (data type) tickit\_events, when Limit by Resource is enabled, filtering for the 5 largest values of Venueseats returns the 5 events with the largest venues. When Limit by Resource is disabled, filtering for the 5 largest venues returns all of the events that were held in one of the 5 largest venues.
- 7. If you would like to be able to further constrain the data that appears in the filter, you can add one or more subfilters. To add a subfilter, click **Create Filter** under Subfilters. The process of creating a subfilter is the same as the process for the parent filter. However, the subfilter is not displayed on the dashboard. It is visible only when editing the parent filter, and the subfilter's configuration affects only the parent filter and any sibling subfilters.
- 8. When you have finished configuring the filter, click **OK** to add it to the dashboard. The new filter appears in the left pane of the dashboard and displays the options for configuring the limit.
- 9. To configure the limit for filtering data, specify a number in the **Include the** field. And click the drop-down list on the right to choose **Largest** or **Smallest**.

For example, the filter in the image below excludes all but the 10 venues with the most number of seats.

∧ Venue S	eats	:	$\mid$ ×
Include the	10	Largest	

Changing the values in the filter refreshes the dashboard according to the new limit.

When working with the filter on the dashboard, the following options are available:

- Menu (1): The menu contains the following options:
  - Clear: Clicking Clear removes the value in the Include the field.
  - **Designer**: Selecting this option opens the designer so that you can view or change the filter configuration.
- Close (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

## **Adding a List Filter**

List filters display values in a list and allow you to select and filter on multiple values at the same time. List filters are available for all data types. Follow the instructions below to create a list filter.

- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.

Create Filter		×
Fields:*		
Click to edit		
Filter:*		
	CANCEL	ОК

3. Click in the **Fields** field to open the property drop-down list and determine the property to filter on. The list of available properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about tickets sold for various types of events. The type for the dashboard is tickit\_events:

Create Filter	×
Fields:*	
Please select values	
fe •	<b>1</b> 7
Q Search for properties	
Results: 13	
no description -	- In
O>> tickit_events	
Category     - no description -	
<sup>1</sup> 2 <sub>3</sub> Catid - no description -	÷
<sup>1</sup> 2 <sub>3</sub> Dateid - no description -	Þ
Close	1

The list below describes the icons and options that are available when choosing a property:

• The **Root Value** (**h**) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

### Tip

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (=) to clear the path and start again or you can click the forward (>) or back (>) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to

choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.

- 4. After you have selected the property to filter on, click Close to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **List** from the drop-down list. The dialog box is refreshed to show the Filter Properties and other options that are available for the filter type:

Create Filter		
Fields:*		
category/catname		
Filter:*		
List		
Filter Properties		
Title:		
Label field:		
Click to edit		
Exclude: S	Show Bars: 🗌 Show Blanks: 🗌 Always Show Che	ckboxes: 🗸
	Show counts: 🛃 Respond to oth	er filters: 🗸
Subfilters		
+ Create Filter		
	CANCEL	ок

- 6. Configure any of the following properties. All of the fields are optional:
  - **Title**: Defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the Fields value is used as the title.
  - Label Field: If you want to populate the list with values from a property other than the one specified in Fields, you can select an alternate property in this field.
  - Exclude: This setting controls whether selecting a value in the List Filter narrows the results to show only the records that *include* that value or whether selecting a value *excludes* the records that include that value. When Exclude is disabled, selecting a value in the list narrows the Dashboard results to show only the records that include that value. When Exclude is enabled, selecting a value filters out all of the records that include that value.

- **Show Bars**: This setting controls whether the total values for the selected property appear as a bar graphic in the background of the filter.
- Show Blanks: This setting controls whether to include null values for the selected property by listing them as **Blank** in the filter.
- Always Show Checkboxes: This setting controls whether checkboxes are shown next to the items in the filter.
- **Show Counts**: This setting controls whether the number of results for each item in the filter is displayed in parentheses next to the item.
- **Respond to Other Filters**: This setting controls whether the results of this filter change based on the selections made in other filters on the dashboard.
- 7. If you would like to be able to further constrain the data that appears in the filter, you can add one or more subfilters. To add a subfilter, click **Create Filter** under Subfilters. The process of creating a subfilter is the same as the process for the parent filter. However, the subfilter is not displayed on the dashboard. It is visible only when editing the parent filter, and the subfilter's configuration affects only the parent filter and any sibling subfilters.
- 8. When you have finished configuring the filter, click **OK** to add it to the dashboard. The new filter appears in the left pane of the dashboard and displays the values that are available for filtering the displayed data.

For example, the filter in the image below shows a filter that lists the categories of events. One or more of the categories can be selected to filter the data on the dashboard.

∧ Categories	:	$  \times$
Search		↓å
Musicals (21)     Plays (46)     Pop (77)		
3 choices		

Depending on whether the Exclude option is enabled or disabled, selecting an item in the filter refreshes the dashboard to show only the data includes the selected value or only the data that does not include the selected value.

When working with the filter on the dashboard, the following options are available for sorting and configuration:

- Search: Enables you to search for a value in the list. The search is case-insensitive.
- **Sort**  $(\downarrow_2^{*})$ : Shows the following options for sorting the results in the filter:
  - <sup>o</sup> **Count Ascending**: Orders results from the smallest count to the largest.
  - ° **Count Descending**: Orders results from the largest count to the smallest.
  - Name Ascending: Orders results in alphabetical order.
  - Name Descending: Orders results in reverse alphabetical order.
- show/hide filters: This option is displayed when a value is selected. It shows or hides the selection.
- Menu (1): The menu contains the following options:
  - ° Select All Visible: This option selects all of the items that are listed in the filter.
  - Clear: This option becomes available when an item is selected. Clicking Clear removes the selection.
  - Designer: Selecting this option opens the designer so that you can view or change the filter configuration.
- Close (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

## Adding a Numeric Range Filter

Numeric range filters are used to limit the results on a dashboard to data that falls in (or outside of) certain numeric groupings. Numeric range filters are available for properties with integer and double data types. Follow the instructions below to create a numeric range filter.

- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.

Create Filter		×
Fields:*		
Click to edit		
Filter:*		
	CANCEL	OK

3. Click in the **Fields** field to open the property drop-down list and determine the integer or double type property to filter on. The values for this property will be used to determine the numeric ranges for the filter. The list of available properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about tickets sold for various types of events. The type for the dashboard is tickit listings:

Create Filter	×
Fields.*	
Please select values	
fe •	Ĥ7
Q Search for properties	
Results: 13	
Root Value	In
O>> tickit_listings	
<sup>1</sup> 2 <sub>3</sub> Dateid - no description -	÷
<sup>1</sup> 2 <sub>3</sub> Eventid - no description -	÷
<sup>1</sup> 2 <sub>3</sub> Listid - no description -	÷
Listime - no description -	9-
Close	

The list below describes the icons and options that are available when choosing a property:

• The **Root Value** (**n**) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IF) or back (IF) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 4. After you have selected the property to filter on, click Close to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **Numeric Range** from the drop-down list. The dialog box is refreshed to show the Filter Properties, Format, and other options that are available for the filter type:

Create Filter		×
Fields:*		
totalprice		
Filter:*		
Numeric Range		
Filter Properties -		
Title:		
Interval:*		
	Exclude: Show Bars: Always Show Checkboxes	
	Show counts: 🔽 Respond to other filters	: 🗸
Format		
Type:		
Automatic		
Subfilters		
+ Create Filter		
	011051	
	CANCEL	UK

- 6. Configure any of the following properties. **Interval** is a required field, and the rest of the fields are optional:
  - **Title**: Defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the Fields value is used as the title.
  - Interval: This setting specifies a number that defines the size of the groupings.
  - Exclude: This setting controls whether selecting a range in the filter narrows the results to show only the records that are *included* in that range or whether selecting a range *excludes* the records that fall in that range. When Exclude is disabled, selecting a range narrows the dashboard results to show only the records that fall in that range. When Exclude is enabled, selecting a range filters out all of the records that fall in that range.
  - **Show Bars**: This setting controls whether the total values for the selected property appear as a bar graphic in the background of the filter.
  - Always Show Checkboxes: This setting controls whether checkboxes are shown next to the items in the filter.

- **Show Counts**: This setting controls whether the number of results for each range are displayed in parentheses next to the range.
- **Respond to Other Filters**: This setting controls whether the results of this filter change based on the selections made in other filters on the dashboard.
- 7. If you would like to specify the format for the numeric values that are displayed in the filter, click the **Type** field and select a format from the drop-down list.
- 8. If you would like to be able to further constrain the data that appears in the filter, you can add one or more subfilters. To add a subfilter, click **Create Filter** under Subfilters. The process of creating a subfilter is the same as the process for the parent filter. However, the subfilter is not displayed on the dashboard. It is visible only when editing the parent filter, and the subfilter's configuration affects only the parent filter and any sibling subfilters.
- 9. When you have finished configuring the filter, click **OK** to add it to the dashboard. The new filter appears in the left pane of the dashboard and displays the values that are available for filtering the displayed data.

For example, the filter in the image below shows numeric ranges for the total price paid for tickets to events. The Interval is 500, and format is Money, resulting in ranges for each group of \$500.

∧ Total Price Paid	:	$  \times$
[+] show sort details		1
\$20.00 - \$520.00 (31,394)		
\$520.00 - \$1,020.00 (26,807)		
\$1,020.00 - \$1,520.00 (21,069)		
\$1,520.00 - \$2,020.00 (17,317)		
\$2,020.00 - \$2,520.00 (14,274)		
\$2,520.00 - \$3,020.00 (11,869)		
S3,020.00 - \$3,520.00 (10,044)		
\$3,520.00 - \$4,020.00 (8,462)		
\$4,020.00 - \$4,520.00 (7,424)		
\$4,520.00 - \$5,020.00 (6,253)		
40 choices   more		

Depending on whether the Exclude option is enabled or disabled, clicking a range in the filter refreshes the dashboard to show only the data that is in the selected range or only the data that is outside of the range.

When working with the filter on the dashboard, the following options are available for sorting and configuration:

- **show/hide sort details**: Shows or hides the following option for sorting the results in the filter:
  - Direction: This option controls how you want to order the ranges in the filter, depending on the Format specified for the values in the filter. Count Ascending and Name Ascending order results from the smallest range to the largest and Count Descending and Name Descending order results from the largest range to the smallest.
- show/hide filters: This option is displayed when a range is selected. It shows or hides the selection.
- Menu (1): The menu contains the following options:
  - <sup>o</sup> **Select All Visible**: This option selects all of the ranges that are listed in the filter.
  - Clear: This option becomes available when a range is selected. Clicking Clear removes the selection.
  - **Designer**: Selecting this option opens the designer so that you can view or change the filter configuration.
- Close (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

## Adding a Presence Filter

Presence filters group results based on whether the value exists or does not exist. This type of filter is useful for testing whether there are records that are missing a particular value. Presence filters are available for paths and properties of all data types. Follow the instructions below to create a presence filter.

- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.

Create Filter		×
Fields:*		
Click to edit		
Filter.*		
	CANCEL	ок

3. Click in the **Fields** field to open the property drop-down list and determine the path or property to filter on. The list of available properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about tickets sold for various types of events. The type for the dashboard is tickit events:

Create Filter	×
Fields:*	
Please select values	
fe •	87
Q Bearch for properties	
Results: 13	
Root Value - no description -	- II
O>> tickit_events	
Category     - no description -	Ĩ
<sup>1</sup> 2 <sub>3</sub> Catid - no description -	₽
<sup>1</sup> 2 <sub>3</sub> Dateid - no description -	Þ
Close	LI II

The list below describes the icons and options that are available when choosing a property:

 The Root Value (A) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IP) or back (IP) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 4. After you have selected the property to filter on, click **Close** to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **Presence** from the drop-down list. The dialog box is refreshed to show the Filter Properties and other options that are available for the filter type:

Create Filter		×
Fields:*		
eventid		
Filter:*		
Presence		
Filter Properties		
Title:		
	Show counts: 🛃 Respond to other filte	ers: 🗸
Subfilters		
+ Create Filter		
	CANCEL	ок

- 6. Configure any of the following properties. All of the fields are optional:
  - **Title**: Defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the Fields value is used as the title.
  - **Show Counts**: This setting controls whether the number of results for each item in the filter is displayed in parentheses next to the item.
  - **Respond to Other Filters**: This setting controls whether the results of this filter change based on the selections made in other filters on the dashboard.
- 7. If you would like to be able to further constrain the data that appears in the filter, you can add one or more subfilters. To add a subfilter, click **Create Filter** under Subfilters. The process of creating a subfilter is the same as the process for the parent filter. However, the subfilter is not displayed on the dashboard. It is visible only when editing the parent filter, and the subfilter's configuration affects only the parent filter and any sibling subfilters.
- 8. When you have finished configuring the filter, click **OK** to add it to the Dashboard. The new filter appears in the left pane of the dashboard and displays the values that are available for filtering the displayed data.

For example, the filter in the image below shows a Presence Filter that tests whether there are records that are missing the Category ID value. Selecting **Exists** or **Does not exist** filters the dashboard data to show only the records that fall into that category.



When working with the filter on the dashboard, the following options are available for sorting and configuration:

- Menu (1): The menu contains the following options:
  - Clear: This option becomes available when an item is selected. Clicking Clear removes the selection.

- **Designer**: Selecting this option opens the designer so that you can view or change the filter configuration.
- **Close** (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

# Adding a Quartile Filter

Quartile filters group and rank the values for a property into four equal ranges. This filter is available for properties with integer, double, date, time, and dateTime data types. It is not available for paths. Follow the instructions below to create a quartile filter.

- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.

Create Filter		×
Fields:*		
Click to edit		
Filter:*		
	CANCEL	OK

3. Click in the **Fields** field to open the property drop-down list and determine the property to filter on. The list of available properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about tickets sold for various types of events. The type for the dashboard is tickit\_events:

Create Filter	×
Fields:*	
Please select values	
fe •	¥۲
Q Search for properties	
Results: 13	
Root Value     no description -	- II
O>> tickit_events	- 11
Category     - no description -	
<sup>1</sup> 2 <sub>3</sub> Catid - no description -	÷
<sup>1</sup> 2 <sub>3</sub> Dateid - no description -	Þ
Close	11

The list below describes the icons and options that are available when choosing a property:

• The **Root Value** (**h**) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

### Tip

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (=) to clear the path and start again or you can click the forward (>) or back (>) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to

choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.

- 4. After you have selected the property to filter on, click **Close** to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **Quartile** from the drop-down list. The dialog box is refreshed to show the Filter Properties that are available for the filter type:

Create Filter		×
Fields:*		
priceperticket		
Filter:*		
Quartile		
Filter Properties		
Filter Properties		
Filter Properties		
Title:		
Title:		

- 6. The only property that is configurable for Quartile filters is the **Title**. This field is optional and defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the value from the Fields field is used as the title.
- 7. When you have finished configuring the filter, click **OK** to add it to the dashboard. The new filter appears in the left pane of the dashboard and displays the values that are available for filtering the displayed data.

For example, the filter in the image below shows a filter that groups the price paid for tickets to events into four equal ranges.



One or more ranges can be selected to filter the dashboard data to show only the records from the selected ranges.

When working with the filter on the dashboard, the following options are available for sorting and configuration:

- show/hide filters: This option is displayed when a value is selected. It shows or hides the selection.
- Menu (1): The menu contains the following options:
  - ° Select All Visible: This option selects all of the items that are listed in the filter.
  - Clear: This option becomes available when an item is selected. Clicking Clear removes the selection.
  - **Designer**: Selecting this option opens the designer so that you can view or change the filter configuration.
- Close (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

# Adding a Range Slider Filter

Range slider filters display a slider control that enables you to filter dashboard results by setting one range that you can adjust as needed. This type of filter is available for properties with integer, double, date, time, and dateTime data types. It is not available for paths. Follow the instructions below to create a range slider filter.

- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.



3. Click in the **Fields** field to open the property drop-down list and determine the property to filter on. The list of available properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about tickets sold for various types of events. The type for the dashboard is tickit events:

Create Filter	×
Fields:*	
Please select values	
fe •	₩7
Q [Search for properties	
Results: 13	
Root Value - no description -	Ir
O>> tickit_events	
<ul> <li>Category         <ul> <li>no description -</li> </ul> </li> </ul>	
<sup>1</sup> 2 <sub>3</sub> Catid - no description -	₽
<sup>1</sup> 23 Dateid - no description -	₽
Close	

The list below describes the icons and options that are available when choosing a property:

 The Root Value (A) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IF) or back (IF) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function

button (**fx**) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click **Advanced**. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.

- 4. After you have selected the property to filter on, click **Close** to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **Range Slider** from the drop-down list. The dialog box is refreshed to show the Filter Properties and other options that are available for the filter type:

Create Filter		×
Fields:*		
eventDate/caldate		
Filter:*		
Range Slider		
Filter Properties		
Title:		
Label field:		
Click to edit		
Subfilters		
+ Create Filter		
	CANCEL	ок

- 6. Configure any of the following properties. All of the fields are optional:
  - **Title**: Defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the Fields value is used as the title.
  - Label Field: If you want to populate the list with values from a property other than the one specified in Fields, you can select an alternate property in this field.
- 7. If you would like to be able to further constrain the data that appears in the filter, you can add one or more subfilters. To add a subfilter, click **Create Filter** under Subfilters. The process of creating a subfilter is the same as the process for the parent filter. However, the subfilter is not displayed on the dashboard. It is visible only when editing the parent filter, and the subfilter's configuration affects only the parent filter and any sibling subfilters.

8. When you have finished configuring the filter, click **OK** to add it to the dashboard. The new filter appears in the left pane of the dashboard and displays the values that are available for filtering the displayed data.

For example, the filter in the image below shows a Range Slider filter that uses a date property to define the range. The dashboard can be filtered by adjusting the Min and Max values to decrease or increase the range.



9. To refresh the results on the dashboard you can click and drag the Min or Max end of the slider or click the plus and minus buttons to adjust the range in small increments.

When working with the filter on the dashboard, the following options are available for sorting and configuration:

- Menu (1): The menu contains the following options:
  - Clear: This option is available once the slider has been adjusted. Clicking Clear resets the range to the default configuration.
  - **Designer**: Selecting this option opens the designer so that you can view or change the filter configuration.
- Close (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

## Adding a Relative Time Filter

Relative time filters are used to filter for records that fall into a specified increment of time relative to the current time. This type of filter is available for date, time, and dateTime data types. Follow the instructions below to create a relative time filter.
- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.

Create Filter		×
Fields:*		
Click to edit		
Filter:*		
	CANCEL	OK

3. Click in the **Fields** field to open the property drop-down list and determine the date, dateTime, or time type property to filter on. The list of available properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about tickets sold for various types of events. The type for the dashboard is tickit\_events:

Create Filter	×
Fields:*	
Please select values	
fe -	87
Q Bearch for properties	
Results: 13	
Root Value     -no description -	
O>> tickit_events	
Category     -no description -	Ĩ
<sup>1</sup> 2 <sub>3</sub> Catid - no description -	₽
<sup>1</sup> 2 <sub>3</sub> Dateid - no description -	₽
Close	

The list below describes the icons and options that are available when choosing a property:

• The **Root Value** (**•**) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IP) or back (IP) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 4. After you have selected the property to filter on, click Close to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **Relative Time** from the drop-down list. The dialog box is refreshed to show the Filter Properties and other options that are available for the filter type:

Create Filter		×
Fields:*		
starttime		
Filter:*		
Relative Time		
Filter Properties		
Title:		
Label field:		
Click to edit		
	Respond to other filter	rs: 🗸
Subfilters		
+ Create Filter		
	CANCEL	ок

- 6. Configure any of the following properties. All of the fields are optional:
  - **Title**: Defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the Fields value is used as the title.
  - Label Field: If you want to populate the filter with a label other than the one specified in Fields, you can select an alternate property in this field.
  - **Respond to Other Filters**: This setting controls whether the results of this filter change based on the selections made in other filters on the dashboard.
- 7. If you would like to be able to further constrain the data that appears in the filter, you can add one or more subfilters. To add a subfilter, click **Create Filter** under Subfilters. The process of creating a subfilter is the same as the process for the parent filter. However, the subfilter is not displayed on the dashboard. It is visible only when editing the parent filter, and the subfilter's configuration affects only the parent filter and any sibling subfilters.
- 8. When you have finished configuring the filter, click OK to add it to the dashboard. The new filter appears in the left pane of the dashboard and displays the values that are available for filtering the displayed data.

For example, the filter in the image below filters data by the Sale Date for tickets.

∧ Sale Date		:	$\mid$ ×
Last	years		

- 9. To configure the requirements for the data to display on the Dashboard, specify the following options:
  - On the left, select Last or Next to configure the relative time direction.
  - In the middle, specify a number to represent the amount of time.
  - On the right, select the time increment to use.

When working with the filter on the dashboard, the following options are also available:

- Menu (1): The menu contains the following options:
  - Clear: This option is available when once the filter options are selected. Clicking Clear resets the filter to the default values.
  - **Designer**: Selecting this option opens the designer so that you can view or change the filter configuration.
- Close (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

## Adding a Search Filter

Search filters are used to search for values of a property that contain a partial match, exact match, or do not equal the text that you specify. The search is case-insensitive. This type of filter is available for all data types. It is not available for use with paths. Follow the instructions below to create a search filter.

- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.

Create Filter	×
Fields:*	
Click to edit	
Filter.*	
C/	ANCEL OK

3. Click in the **Fields** field to open the property drop-down list and determine the property to filter on. The list of available properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about tickets sold for various types of events. The type for the dashboard is tickit\_events:

Create Filter	×
Fields:*	
Please select values	
fe -	<b>1</b> 7
Q [Search for properties	
Results: 13	
Root Value - no description -	- It
O>> tickit_events	
Category     -no description -	
<sup>1</sup> 2 <sub>3</sub> Catid - no description -	₽
<sup>1</sup> 2 <sub>3</sub> Dateid - no description -	Þ
Close	1

The list below describes the icons and options that are available when choosing a property:

• The **Root Value** (**h**) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

#### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IP) or back (IP) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.
- 4. After you have selected the property to filter on, click **Close** to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **Search Filter** from the drop-down list. The dialog box is refreshed to show the Filter Properties and other options that are available for the filter type:

Create Filter		×
Fields:*		
eventname		
Filter:*		
Search Filter		
Filter Properties		
Subfilters + Create Filter		
	CANCEL	ок

- 6. The only property that is configurable for Search Filters is the **Title**. This field is optional and defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the value from the Fields field is used as the title.
- 7. If you would like to be able to further constrain the data that appears in the filter, you can add one or more subfilters. To add a subfilter, click **Create Filter** under Subfilters. The process of creating a subfilter is the same as the process for the parent filter. However, the subfilter is not displayed on the dashboard. It is visible only when editing the parent filter, and the subfilter's configuration affects only the parent filter and any sibling subfilters.
- 8. When you have finished configuring the filter, click **OK** to add it to the dashboard. The new filter appears in the left pane of the dashboard and displays the values that are available for filtering the displayed data.

For example, the filter in the image below defines Event Name as the property to search on.

∧ Event Name		:   X
	Event Name	Filter

- 9. To perform a search, configure the following options:
  - On the left, select the type of match to perform. **Matches** includes partial matches and **Equals** is an exact match. Note that matches are case-insensitive.

Q	Matches
-	Equals
÷	Does not equal

- In the middle, specify the value to search for. The search is case-insensitive.
- 10. Press Enter or click Filter to perform the search and refresh the data on the dashboard.

When working with the filter on the dashboard, the following options are also available:

- Menu (1): The menu contains the following options:
  - Clear: This option becomes available when a search term is entered. Clicking Clear removes the term.

- **Designer**: Selecting this option opens the designer so that you can view or change the filter configuration.
- **Close** (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

## Adding a Single Select List Filter

Single select list filters are similar to list filters but only allow you to filter one value from the list at a time. This type of filter is available for properties of all data types but is not available for paths. Follow the instructions below to create a single select list filter.

- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.

Create Filter	×
Fields:*	
Click to edit	
Filter:*	
c	ANCEL OK

3. Click in the **Fields** field to open the property drop-down list and determine the property to filter on. The list of available properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about tickets sold for various types of events. The type for the dashboard is tickit\_events:

Create Filter	×
Fields:*	
Please select values	
fe •	<b>1</b> 7
Q Search for properties	
Results: 13	
no description -	- In
O>> tickit_events	
Category     - no description -	
<sup>1</sup> 2 <sub>3</sub> Catid - no description -	÷
<sup>1</sup> 2 <sub>3</sub> Dateid - no description -	Þ
Close	1

The list below describes the icons and options that are available when choosing a property:

• The **Root Value** (**h**) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IF) or back (IF) arrows to go forward or backward one or more levels at a time.
- After you have selected a property, you can apply a function or formula to that property to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that are available depend on the data type of the selected property. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to

choose additional properties and functions. For more information, see Calculating Values in Lenses and Filters.

- 4. After you have selected the property to filter on, click **Close** to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **Single Select List** from the drop-down list. The dialog box is refreshed to show the Filter Properties and other options that are available for the filter type:

Create Filter	×
Fields:*	
category/catname	
Filter:*	
Single Select List	
Filter Properties	
Title:	
Label field:	
Click to edit	
	Exclude: Show Bars: Show Blanks:
	Show counts: 🛃 Respond to other filters: 🛃
Subfilters + Create Filter	
	CANCEL OK

- 6. Configure any of the following properties. All of the fields are optional:
  - **Title**: Defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the Fields value is used as the title.
  - Label Field: If you want to populate the list with values from a property other than the one specified in Fields, you can select an alternate property in this field.
  - Exclude: This setting controls whether selecting a value in the filter narrows the results to show only the records that *include* that value or whether selecting a value *excludes* the records that include that value. When Exclude is disabled, selecting a value in the list narrows the dashboard results to show only the records that include that value. When Exclude is enabled, selecting a value filters out all of the records that include that value.

- **Show Bars**: This setting controls whether the total values for the selected property appear as a bar graphic in the background of the filter.
- Show Blanks: This setting controls whether to include null values for the selected property by listing them as **Blank** in the filter.
- Show Counts: This setting controls whether the number of results for each item in the filter is displayed in parentheses next to the item.
- **Respond to Other Filters**: This setting controls whether the results of this filter change based on the selections made in other filters on the dashboard.
- 7. If you would like to be able to further constrain the data that appears in the filter, you can add one or more subfilters. To add a subfilter, click **Create Filter** under Subfilters. The process of creating a subfilter is the same as the process for the parent filter. However, the subfilter is not displayed on the dashboard. It is visible only when editing the parent filter, and the subfilter's configuration affects only the parent filter and any sibling subfilters.
- 8. When you have finished configuring the filter, click **OK** to add it to the dashboard. The new filter appears in the left pane of the dashboard and displays the values that are available for filtering the displayed data.

For example, the filter in the image below shows a filter that lists each of the states that have event venues. One state at a time can be selected to filter the data on the dashboard.

∧ Venue State	:   X
AB (52)	
AZ(128)	
BC(31)	
CA(911)	
CO(140)	
DC(114)	
FL (205)	
GA (81)	
IL (196)	
IN (43)	
33 choices   more	

Depending on whether the Exclude option is enabled or disabled, selecting an item in the filter refreshes the Dashboard to show only the data includes the selected value or only the data that does not include the selected value.

When working with the filter on the dashboard, the following options are available for sorting and configuration:

- Menu (1): The menu contains the following options:
  - Clear: This option is available when an item is selected. Clicking Clear removes the selection.
  - **Designer**: Selecting this option opens the designer so that you can view or change the filter configuration.
- Close (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

## **Adding a Types Filter**

Types filters are used to filter data according to the types of data (classes) that are connected by a specified path. This type of filter is available only for paths and not properties. Follow the instructions below to create a types filter.

- 1. Open the dashboard that you want to add the filter to.
- 2. In the Hi-Res Analytics main toolbar, click **Filters** and select **Create a Filter**. The Create Filter dialog box is displayed.



 Click in the Fields field to open the property drop-down list and determine the path to filter on. The list of available classes, paths, and properties depends on the selected data type for the dashboard.

For example, the following image shows the list of properties that are available for a dashboard whose source is a graphmart that contains data about tickets sold for various types of events. The type for the dashboard is tickit events:

Create Filter	×
Fields:*	
Please select values	
fe •	<b>1</b> 7
Q Search for properties	
Results: 13	
no description -	- In
O>> tickit_events	
Category     - no description -	
<sup>1</sup> 2 <sub>3</sub> Catid - no description -	÷
<sup>1</sup> 2 <sub>3</sub> Dateid - no description -	Þ
Close	1

The list below describes the icons and options that are available when choosing a path:

• The **Root Value** (**h**) is the instance URI for the root resource—the URI for the instances of the class that was chosen as the data type for the dashboard.

### Tip

To view the root values, you can use the STR function to show a string representation of the URIs.

- Linked classes are represented by incoming (o«) and outgoing (o») connection icons.
   The properties in those classes with a path to another class are denoted with a circle icon (o). Selecting a linked property navigates to that class and displays its properties.
- When a property or path is selected, the breadcrumbs at the top of the dialog box show you the property path. You can click the clear icon (IF) to clear the path and start again or you can click the forward (IF) or back (IF) arrows to go forward or backward one or more levels at a time.
- After you have selected a path, you can apply a function or formula to calculate the values that are displayed in the filter. To add a function, click the function button (fx) at the top of the drop-down list. The functions that become available depend on the data type of the selected path. To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose

additional properties and functions. For more information, see Calculating Values in Lenses and Filters.

- 4. After you have selected the path to filter on, click **Close** to close the Fields drop-down list.
- 5. Next, click the **Filter** field and select **Types** from the drop-down list. The dialog box is refreshed to show the Filter Properties and other options that are available for the filter type:

Create Fi	lter	;
Fields:*		
ROOT		
Filter:*		
Types		
	Exclude: Show Bars: Show Blanks: Always Show Checkboxe	s: 🔽
	Exclude: Show Bars: Show Blanks: Always Show Checkboxe Show counts: Respond to other filter	s: 🔽
Subfilters + Create Filt	er	
	CANCEL	ок

- 6. Configure any of the following properties. All of the fields are optional:
  - **Title**: Defines the title that appears at the top of the filter when it is added to the dashboard. If Title is blank, the Fields value is used as the title.
  - Exclude: This setting controls whether selecting a value in the filter narrows the results to show only the records that *include* that value or whether selecting a value *excludes* the records that include that value. When Exclude is disabled, selecting a value in the filter narrows the dashboard results to show only the records that include that value. When Exclude is enabled, selecting a value filters out all of the records that include that value.
  - **Show Bars**: This setting controls whether the total values for the selected property appear as a bar graphic in the background of the filter.
  - **Show Blanks**: This setting controls whether to include null values for the selected property by listing them as **Blank** in the filter.

- Always Show Checkboxes: This setting controls whether checkboxes are shown next to the items in the filter.
- **Show Counts**: This setting controls whether the number of results for each filter value are displayed in parentheses next to the value.
- **Respond to Other Filters**: This setting controls whether the results of this filter change based on the selections made in other filters on the dashboard.
- 7. If you would like to be able to further constrain the data that appears in the filter, you can add one or more subfilters. To add a subfilter, click **Create Filter** under Subfilters. The process of creating a subfilter is the same as the process for the parent filter. However, the subfilter is not displayed on the dashboard. It is visible only when editing the parent filter, and the subfilter's configuration affects only the parent filter and any sibling subfilters.
- 8. When you have finished configuring the filter, click **OK** to add it to the dashboard. The new filter appears in the left pane of the dashboard and displays the values that are available for filtering the displayed data.

For example, the filter in the image below shows that there is one type of class of data that is connected by the chosen path.

<ul> <li>Types for 'listing</li> </ul>	:	$  \times$	
Search		↓å	
tickit_listings (640)			
1 choices			

Depending on whether the Exclude option is enabled or disabled, selecting an item in the filter refreshes the dashboard to show only the data includes the selected value or only the data that does not include the selected value.

When working with the filter on the dashboard, the following options are available for sorting and configuration:

- Search: Enables you to search for a value in the filter. The search is case-insensitive.
- Sort  $(\downarrow_{z}^{A})$ : Shows the following options for sorting the results in the filter:
  - <sup>o</sup> **Count Ascending**: Orders results from the smallest count to the largest.
  - ° **Count Descending**: Orders results from the largest count to the smallest.
  - Name Ascending: Orders results in alphabetical order.
  - Name Descending: Orders results in reverse alphabetical order.
- **show/hide filters**: This option is displayed when a value is selected. It shows or hides the selection.
- Menu (1): The menu contains the following options:
  - <sup>o</sup> **Select All Visible**: This option selects all of the items that are listed in the filter.
  - Clear: This option is available when one or more items are selected. Clicking Clear removes the selections.
  - Designer: Selecting this option opens the designer so that you can view or change the filter configuration.
- Close (X): Clicking X deletes the filter from the dashboard. This action cannot be undone.

# **Calculating Values in Lenses and Filters**

Anzo provides many standard and advanced functions that you can use to compute the values that are displayed on a dashboard. When selecting properties and paths for lenses and filters, you can add calculations by selecting functions from a list or by writing your own formula. The Hi-Res Analytics application enables you to save your formulas as computed properties that can be reused on other dashboards, lenses, and filters. This topic provides instructions for using functions and formulas to calculate displayed values, saving formulas as computed properties, and reusing computed properties.

- Computations in Filters and Lenses
- Applying Functions and Formulas to Properties
- Saving Formulas for Reuse
- Reusing Computed Properties

### **Computations in Filters and Lenses**

When you apply formulas to properties in filters, Anzo performs the calculation across all of the values that exist for the selected property and then groups the results into the list of values that the calculations return. For multiple value properties, all values for that property are included in the calculations.

When you apply formulas to properties in lenses, the calculation results depend on the data type of the dashboard or lens. If the property belongs to a class that allows multiple values, Anzo performs the calculation on each set of multiple values and returns the results as one record in the lens. If the class includes single value properties, the calculation is performed separately for each single value.

## **Applying Functions and Formulas to Properties**

Follow the instructions below to use a function or formula to compute the values in a lens or filter.

- 1. Create a new lens or filter or open the Designer for an existing lens or filter.
- 2. In the drop-down list for selecting properties or fields, select the property or path for which you want to compute the values. For example:

∫ ← priceperticket	Î
Q Search for properties	
Results: 13	
1.23 Priceperticket - no description -	۵
Seller     - no description -	
<sup>1</sup> 2 <sub>3</sub> Sellerid - no description -	€
Close	1

3. Click the function button (**fx**) to display the list of standard functions. The list varies depending on the data type of the selected property.

f <sub>x</sub> •	priceperticket
AVG	1
SUM	1
MIN	
MA	ĸ
COL	INT
COL	INT_DISTINCT
STR	
Adv	anced
60 7.20 MI	VI

- 4. Click a function to apply it to the property that you chose. For information about each of the available functions, see Function and Formula Reference.
- To choose a more advanced function or type a formula, click Advanced. The Calculated Value dialog box opens and enables you to choose additional properties and functions as well as type your own calculation.

Calculated Value	X
Please select values	
Formula [{tickit_listings}priceperticket]	
Fields	Save as computed property   Clear Functions
Q Search for properties	Q Search for functions
Results: 13      Root Value     . no description - O≫ tickit_listings      123 Dateid      - no description -	Results: 182 $f_X$ ABS(number) Calculate absolute value of term. $f_X$ ADD(value1, value2) Calculate addition between 2 terms.
	CLOSE

## Тір

To create an advanced formula, it might help to get started by viewing the functions listed in the Functions column on the right side of the screen. Each function includes the syntax to use for creating a formula that uses that function.

- 6. In the Functions column, double-click a function to add it to the Formula field at the top of the dialog box.
- 7. Place your cursor in the Formula where you want to insert the property to perform the calculation on (for example, inside the parentheses) and then double-click the property in the Fields column. If the syntax for the function includes characters such as commas, type the characters in the appropriate location in the formula. You can click the **Clear** link on the bottom right of the Formula field any time to clear that field and start over.

For example, the formula below calculates the total price paid by multiplying the values in the price per ticket column with the value in the number of tickets column:

Calculated Value	×
Please select values	
Formula MULTIPLY([{tickit_listings}numtickets	],[{tickit_listings}priceperticket])
	Save as computed property   Clear
Fields	Functions
Q Search for properties	Q multi I 🗆
Results: 13	Results: 1
<sup>1</sup> 2 <sub>3</sub> Numtickets → no description -	f <sub>x</sub> MULTIPLY(v1, v2) Calculate numeric multiply
1.23 Priceperticket ⇒ - no description -	
Insert Field   Reset	
	CLOSE

- 8. When you are finished writing a formula, you have two options:
  - If you want to use the formula now without saving it for later use, click **Close** to close the Calculated Value dialog box. Then complete the lens or filter configuration.
  - If you want save the formula for reuse, click the **Save as computed property** link and follow the instructions below in **Saving Formulas** for Reuse.

## **Saving Formulas for Reuse**

Follow these instructions to save a formula as a computed property that you can use in other lenses and filters that target the same class of data.

 When you have finished writing a formula in the Calculated Value dialog box, click the Save as computed property link below the Formula field. The Save formula as computed property dialog box opens. 2. In the **Title** field, type a name for the new computed property.

Save formula as a computed property			
Title:*			
Description:			
Ontology:			
Tickets - Auto			
	CANCI	EL SAV	

- 3. Type a description of the new property in the **Description** field.
- 4. If necessary, click in the **Ontology** field to choose another ontology to save the property in. If you want to save this property in multiple ontologies, you can click the **Save as computed property** link again after saving the property in the current ontology.
- 5. Click Save. Anzo saves the new property and labels it with an asterisk (\*). The property becomes available in the Fields column in the Calculated Value dialog box. For example, the image below shows the Computed Total Price property that saves the formula from the Applying Functions and Formulas to Properties section above.



6. Click **Close** to close the dialog box.

# **Reusing Computed Properties**

When an ontology contains computed properties, any other dashboards, lenses, and filters that use that ontology can also use the computed properties as long as they also use the same data type or class of data that the computed property is saved in. Follow the instructions below to use a computed property.

- 1. Open the Designer for the filter or lens where you want to apply the computed property.
- Click in the Fields or Column Value Expression field to open the property selection dropdown list. The drop-down list includes any computed properties that are available for use with the selected data type. Computed properties are labeled with an asterisk (\*).

Fields	Functions
Q Search for properties	Q Search for functions
Results: 13  Root Value - no description -	Results: 182 $f_X$ ABS( number ) Calculate absolute value of term.
O≫ tickit_listings	f <sub>X</sub> ADD( value1, value2 ) Calculate addition between 2 terms. f <sub>X</sub> AND( logical1, logical2 )
<sup>1</sup> 2 <sub>3</sub> Dateid → no description -	Calculate logical AND of params. f <sub>x</sub> AVERAGEIF( values_to_test, criteria, valu Averages only those elements that match the
Insert Field   Reset	

- 3. To use the property as-is, select the property and then close the drop-down list.
- 4. If you want to make changes to the formula and save it as a different computed property, select the property and then click the function button (fx) to open the Calculated Value dialog box. Follow the instructions in Saving Formulas for Reuse above to edit the formula and then save a new computed property.

# **Combining Data from Multiple Classes**

Though you must choose one base data type (or class) for each dashboard, selecting a data type with connections to other classes enables you to configure lenses and filters that combine the data from those classes. This powerful capability can help surface the semantic relationships in your data and enable you to leverage those relationships to access and integrate all of the data in the knowledge graph. When choosing the base data type for a dashboard, it helps to consider all of the desired filters.



For example, consider the following data model for a movie data set:

By creating a dashboard that specifies **Movies** as the base data type, the lenses and filters in the dashboard can navigate the paths to properties in the other classes, Movie Actors, Movie Directors, Movie Producers, and so on. This topic provides guidance on accessing data from multiple classes in filters and lenses.

- Combining Classes in a Lens
- Filtering on Multiple Classes

## **Combining Classes in a Lens**

The image below shows a dashboard that accesses the graph for the above model. The specified data type is **Movies**, and a table lens displays all of the columns/properties in the Movies class:

Deside the second		
ready to use		
LTERS	<	🗘 Designer 🔄 Properties 🔄 Export 🔠 Manage Columns
E Manage Lenses		What can I do next? × All Movies ×
∧ Graphmart	+	Movie ID 🔶 Movie Title
Movies Graphmart		3217 Army of Darkness
		3746 Blade Runner
		3837 Blazing Saddles
		3947 Blue Velvet (film)
A Data Types	+ :	4227 Barry Lyndon
		4231 Buffy the Vampire Stayer (film)
Movies ( 8,157 )	1	4560 Braveheart
		4726 Batman (1989 film)
		4727 Batman (1966 film)
		4728 Batman Returns
<ul> <li>Data Lawers</li> </ul>	:	4729 Batman & Robin (film)
		4730 Batman Forever
Movies		5313 Crouching Tiger, Hidden Dragon
		7906 Destry Rides Again
		8481 Dressed to Kill (1980 film)

Lenses and filters can be configured to leverage the relationships from the base class to the connected classes. For example, adding a column that navigates the **ActedIn** path to access the **Movie Actors** class could be used to display values such as the names of the actors who starred in the movies. To navigate the relationship in the lens Designer, the **ActedIn** path from the **MovieActors** class is selected for the new column:

Designer		×
Columns + in de Van	Column Details Set properties of the selected column.	
II Movie Title	Column Value Expression:	
Untitled Disable   View   Clone Remove	Please select values	yr ≣≓
	Q Search for properties	k
	Results: 8 The subject is an instance of a class.	
	Acted In     - no description -	
Add a column   Display Properties	O  O  O  O  O  O  O  O  O  O  O  O  O	
Thereit changes	Close	-

Once the path is chosen, all of the properties from the Movie Actor class are displayed:



Selecting **ActorName** adds the column to the dashboard. The actors from each movie are now integrated into the lens even though the actor name values are not in the base class.

🗘 Designer	Properties	∋ Export	III Manage Columns	
All Movies	×			
Movie ID		$\uparrow$	Movie Title	Actor Name
4727			Batman (1966 film)	Lee Meriwether, Frank Gorshin, Cesar Romero, Burt Ward, Burgess Meredith, Adam West
4728			Batman Returns	Pat Hingle, Michelle Pfeiffer, Michael Murphy (actor), Michael Keaton, Michael Gough, Danny DeVito, Christopher Walken
4729			Batman & Robin (film)	Uma Thurman, Pat Hingle, Michael Gough, John Glover (actor), George Clooney, Elle Macpherson, Chris O'Donnell, Arnold Schwarzenegger, Alicia Silverstone
4730			Batman Forever	Val Kilmer, Tommy Lee Jones, Pat Hingle, Nicole Kidman, Michael Gough, Jim Carrey, Chris O'Donnell
5313			Crouching Tiger, Hidden Dragon	Zhang Ziyi, Michelle Yeoh, Chow Yun-fat, Chang Chen
7906			Destry Rides Again	Mischa Auer, Marlene Dietrich, James Stewart, Brian Donlevy
8481			Dressed to Kill (1980 film)	Nancy Allen (actress), Michael Caine, Keith Gordon, Angie Dickinson
			I< < > >I 1-20 of	8157 Page 1 of 408 🍄 👻

# **Filtering on Multiple Classes**

In addition to combining classes in lenses, you can also apply filters across classes. Like the example above, the image below shows a dashboard where the specified data type is **Movies**, and a Table lens displays all of the columns/properties in the Movies class:

Ready to use				
LTERS	<	🗘 Designer 🛛 🕞 Properties 🕞 Exp	ort 🛛 🗰 Manage Colu	imns
Manage Lenses		What can I do next? × All Movies	×	
∧ Graphmart	+	Movie ID	Ŷ	Movie Title
Movies Graphmart		3217		Army of Darkness
		3746		Blade Runner
		3837		Blazing Saddles
		3947		Blue Velvet (film)
<ul> <li>Data Types</li> </ul>	+ :	4227		Barry Lyndon
		4231		Buffy the Vampire Slayer (film)
Movies ( 8,157 )		4560		Braveheart
		4726		Batman (1989 film)
		4727		Batman (1966 film)
		4728		Batman Returns
∧ Data Lavers	1	4729		Batman & Robin (film)
,		4730		Batman Forever
Movies		5313		Crouching Tiger, Hidden Dragon
		7906		Destry Rides Again
		8481		Dressed to Kill (1980 film)

A filter can be configured to leverage the relationships from the base class to the connected classes. For example, adding a filter that navigates the **Directed** path to access the **Movie Directors** class could be used to display, and filter on, the name of the director for each movie. To navigate the relationship in the Create Filter dialog box, the **Directed** path from the **MovieDirectors** class is selected for the Field to filter on:

Create Filter	×
Fields:*	
Please select values	
fr. •	ŧ÷
Q Search for properties	
Results: 8	
Acted In     - no description -	
O« Movie Directors	
Directed     no description -	
O « Movie Editors	
O Edited Close	
oroutining rigel, maach orugen	

Once the path is chosen, the properties from the **Movie Directors** class are displayed:



Selecting the **Director Name** property and choosing **List** as the type of filter adds a filter to the dashboard that lists all of the directors in the graph. Users can select particular director names to filter the lens so that it only shows the movies that include one or more of the selected directors.

*Movies ×						
Ready to use						
FILTERS	<	🌣 Designer	Properties	Export	III Manage Columns	
Manage Lenses		All Movies	~			
✓ Graphmart	+	Movie ID	^	Ŷ	Movie Title	Actor Name
✓ Data Types		1308304			Sidekicks (1992 film)	
✓ Data Layers		2683330			The Invisible Woman (1940 film)	
Directed/Director Name (2)	:   X	5638447			The Hitman	Saved by the Bell: The New Clas Michael Parks, Chuck Norris, Al Waxman
Search [+] show filters	ţş	6653265			Delta Force 2: The Colombian Connection	Richard Jaeckel, Paul Perri, Marl Margolis, John P. Ryan (actor), Chuck Norris, Billy Drago
A. Edward Sutherland (4)		8001525			Follow the Boys	Vera Zorina, George Raft
Aaron Natz (nimmaker) (1)		9058143			Top Dog (1995 film)	Chuck Norris
Aaron Schneider (1)		9095434			Dixie (film)	
Aaron Scott Moorhead (1)		32129077			The Navy Comes Through	
Abel Ferrara (9)						

Using the same path traversal, filters could also be created to narrow the data to certain actors or producers by following the relationships to the Movie Actors or Movie Producers classes. For more information about creating or editing lenses and filters, see Creating a Lens and Working with Filters.

# **Searching for Text in Unstructured Documents**

The Hi-Res Analytics application incorporates the Elasticsearch search engine to enable you to perform full text searches on unstructured documents. This topic provides instructions for creating a dashboard with text search capability and running a search across unstructured documents.

1. In the Anzo application, expand the **Blend** menu and click **Graphmarts**. Anzo displays a list of the existing graphmarts. For example:

<b>Y</b>	<b>Q</b> Search	=			l	Add Graphr	nart 👻
•	Title 🔶 🔶	Description	Status	Statements	Creator	Actions	1
	🕸 Books		C Ready to use	156	System Administrator	Д	:
	Flights		C Ready to use	265,339	System Administrator	Д	:
	⇒ Movies		Ready to use	104,166	System Administrator	Д	:
	© Northwind		C Ready to use	131,386	System Administrator	Д	:
	⊜ Tickets		C Ready to use	3,696,548	System Administrator	Д	:
				Row	s per page: 25 👻 1-5	of 5 <	>

- 2. On the Graphmarts screen, click the name of the graphmart that contains the unstructured data. Anzo displays the graphmart overview screen.
- 3. Click the **Create Dashboard** button. The Hi-Res Analytics application opens and displays the Create Dashboard dialog box. Leave **Graphmart Dashboard** selected and click **Next**.
- Type a name for the dashboard in the Title field and enter an optional Description. Then click OK. Anzo creates the dashboard and populates the Graphmart and Data Layers panels. For example:

	📕 Dashboard	- @ Lenses - ♥ Filters - C Refresh - ♪ Designer ? Help - 🚨
*DU_Cancer_Dashboard (g) $\times$		
Ready to use		
Graphmart     PL_Cancer Graphmart     A Data Types     No types selected	+ / +	What can I do next? ×         Dashboard         This dashboard will allow you to visualize linked data         State         Add data you are interested in exploring         *s Specify the type of data you would like to see
∽ Data Layers PL_Cancer	Y	

5. In the Data Types panel, click the plus icon (+) to open the Select Data Types dialog box. In the dialog box, select **Unstructured Document**.



- 6. Click **OK**. Anzo adds the data type to the Data Types panel.
- Next, click the Lenses button in the main toolbar and select New from the drop-down list. Anzo opens the Create Lens dialog box.

Crea Selec	ate Lens t the type of lens to create		>
<b>c,</b> s	earch for lenses		
Path	to lens values 11		
	Web Page A user-defined web page. Properties		
6	Document Search Document Search lens Properties		
⊞	List List visualization of data Properties		
		CANCEL	NEXT

8. In the Create Lens dialog box, select **Document Search** and then click **Next**. Anzo displays the General Information dialog box.

Create Lens Specify details about the new lens			×
General Information			
Description:			
[+] more			
	CANCEL	PREVIOUS	FINISH

 Type a name for the lens in the Title field and include an optional Description. Then click Finish. Anzo opens the Document Search Designer where you can configure the search settings or customize the style sheet, query, and HTML, if necessary. For example:

Designer					×
Search Settings	Query	HTML	CSS		
General Show No Results On Er	mpty Search:				
Synonym Expansion Dictionary:					
Dataset:  Ontology:					
Predicates:					
Preview changes	]			CANCEL	SAVE

- 10. In the Designer, change the optional search settings as needed. The list below describes each option:
  - Show No Results on Empty Search: Determines whether documents are listed in the search results before a search is run. When enabled, the Document Search lens remains blank until a search is run.
  - Allow Multi Select: Determines whether a user can select multiple documents at a time in the results. When enabled, multiple documents can be selected by holding the Shift key and clicking documents in the results.
  - Synonym Expansion Dictionary: Determines whether to display an option for including synonyms in text searches. When enabled, the lens displays an Include Synonyms checkbox next to the Search field.
  - **Knowledge Base Dataset**: Enables you to include a knowledge base in the search if one exists. Click the field to select an available knowledge base.
  - Ontology: Enables you to select a data model to use for the search.
  - Predicates: Enables you to select specific predicates from the model.

11. Click **Save**. Anzo add the lens to the dashboard. Depending on the search settings, the lens displays the list of documents. For example:

	_		
∧ Graphmart	+	🖌 Designer 🛛 El Properties	
PL_Cancer Graphmart		What can I do next? $\times$ Annotated Document $\times$ Original Document $\times$ Table of Documents $\times$ 26 $\times$	-
	_	Beards	۹
Unstructured Document (27)		Sort by Last Modified V Bescending V	
		27 Results	
A Data Lawar		Growth factor receptors_promising drug largets in cancer.pdf file //iglobalhfs/data/du/docs/cancer/Growth%20factor%20receptors_%20premising%20dug%20fargets%20in%20caUnknown	Mon Dec 02 2
PL.Cancer	-	Upregulation of miR-375 inhibits human liver cancer cell growth by modulating cell proliferation and apoptosis via tar	1
		the Anglobal misioana burbles rean can be pregulation to 2007 and the 2007 and 10 10 10 10 10 10 10 10 10 10 10	Mon Dec 02 2
	-	Emergence of ERBB2 Mutation as a Biomarker and an Actionable Target in Solid Cancers pdf file //iglobalhs/tata/da/doca/cancer/Emergence%20e%/20eRB82%20Mutation%20ar%20e%20Biomarker%20and% Unknown	Mon Dec 02.2
		Recognition of HER2 expression in hepatocellular carcinoma and its significance in postoperative tumor recurrence	Has Dec 42.2
		sexwjeosants oata/suvocis canceri kacegine en %/ustrx/unick/tx/ueppression%/Ulin%/Uhepstocellutar%/Ucarcer Unix/own	won uvec viz 2
		Expression of ERBB2 in cancer - Summary - The Human Protein Atlas, off file //iglobalinfsidata/du/docs/cancerExpression%200%20ERBB2%20in%20cancer%20%20Summary%20%20The Unknown	Mon Dec 02 2
		Increased Expression of c.erbB.2 in Liver is Associated with Hepatitis B x Antisen and Shorter Survival in Patients w file //globalits/data/du/ocs/cancerhereased%20Expression%20FW2Dc-erbB-2%20in%20ExerM2Dis%20Associat Unknown	Mon Dec 02 2

12. To run a search, type the text to find in the **Search** field and press **Enter**. See the Supported Search Syntax section below for information about supported search syntax. Anzo finds documents that include the search value and displays the documents, snippets of text to show the context of where the matches were found, and the Elasticsearch relevance score for the match. For information about how the relevance score is calculated, see <u>What Is Relevance?</u> in the Elasticsearch documentation. For example:



Clicking **Show More** expands the result to display additional matches. For example:



13. To refine the search, alter the text in the **Search** field and press **Enter** again. You can also click highlighted terms in the search results to open a dialog box that shows the full annotated document where the match was found. For example:

Annotated Document	×
DOCUMENT       LAST MODIFIED         12/2/2019       Epidermal @rowth Factor Receptor (EGFR)-targeted Therapies in Esophac         Anticancer Research antianjournals.org       Anticancer Research October 2013 vol. 33 no. 10 4139–4155         Epidermal Growth Factor Receptor (EGFR)- targeted Therapies in Esophagogastric Cancer         SABARISH AYYAPPAN <sup>1</sup> , DHIVYA PRABHAKAR <sup>2</sup> and NEELESH SHARMA <sup>1</sup> 9         Author Affiliations         Correspondence Rot: Neelesh Sharma, MD, Ph. D, Assistant Professor Thoracic Oncology. University Hospitals/ Case Comprehensive cancer Center, Case Western	> Find: Growth Find Concepts: ☑(KB: Biomarker) (1)
2 ✔ 🔉 < > 🚳 C	CANCEL SAVE

## **Supported Search Syntax**

This section describes the keyword search syntax that Anzo supports.

#### Wildcard Characters: ? and \*

• **?**: Use a question mark (?) to represent a single wildcard character. For example, in the search**co?I**, the resulting documents will include terms like "cool" or "coal."

 \*: Use an asterisk (\*) to represent multiple wildcard characters. For example, in the searchcol\*, the resulting documents will include terms like "collect" or "color."

#### Boolean Operators: +, -, OR, AND, NOT

- +: Use a plus (+) character to indicate mandatory matches. For example, in the search flight
   +New York, the resulting documents can include "flight" as an optional match and must include "New York."
- -: Use a minus (-) character to indicate a term that must not match. For example, in the search flight +New York -Los Angeles, the resulting documents can include "flight" as an optional match, must include "New York," and must not include "Los Angeles."
- **OR**: In the search **New York OR Los Angeles**, the resulting documents will include a match for either "New York" or "Los Angeles."
- AND: In the search New York AND Los Angeles, the resulting documents must include matches for both "New York" and "Los Angeles."
- NOT: In the search New York NOT Los Angeles, the resulting documents must include "New York" and cannot contain "Los Angeles."
- Grouping operators: In the search(flight AND New York) OR Los Angeles, the resulting documents will include "flight" and "New York" and optionally include "Los Angeles."

### Fuzzy Matches: ~n

To search for a fuzzy match, use a tilde (~) character followed by a number to represent the number of fuzzy or incorrect characters. For example, in the search **Figth~3**, the resulting documents could include the term "Flight."

### **Regular Expressions**

For example, the following search expression matches email addresses: /([a-zA-Z0-9\_\-\.]+)@([a-zA-Z0-9\_\-\.]+)\.([a-zA-Z]{2,5})/.

For more information about the regular expression syntax that Elasticsearch supports, see <u>Regular</u> <u>expression syntax</u> in the Elasticsearch documentation.

# **Sharing Access to Dashboards and Lenses**

This topic includes reference information about dashboard and lens permissions and provides instructions for configuring permissions.

- Dashboard Level Permissions
- Lens Level Permissions
- Configure Dashboard or Lens Permissions

## **Dashboard Level Permissions**

Dashboard level permissions affect a user's ability to view, modify, delete, design, or configure dashboards and dashboard permissions. There are three predefined permission sets that can be assigned to an Anzo user or group. You also have the option to customize the set of permissions that are applied to a user or group.

The table below lists the predefined permission sets and describes the privileges that are granted for each permission that is part of the predefined set:

Set	Permission	Allows a user to:
View	Read	<ul> <li>Search for and open accessible dashboards.</li> <li>Save As a new dashboard.</li> <li>Share the dashboard.</li> <li>View dashboard Properties.</li> <li>View lens Properties.</li> <li>Export lenses</li> </ul>
Modify	In addition to the <b>Read</b> permission described above, the <b>Modify</b> set include <b>Write</b> and <b>Delete</b> permissions described below.	
	Write	Use the dashboard Designer to change the

Set	Permission	Allows a user to:
		<ul><li>dashboard.</li><li>Clone lenses.</li></ul>
	Delete	Delete the dashboard.
Admin	In addition to the <b>Read</b> , <b>Write</b> , and <b>Delete</b> permissions described above, the <b>Admin</b> set includes the <b>Manage</b> permission described below.	
	Manage	<ul> <li>The Manage permission relates only to the Security tab. If a user has this permission, they can modify dashboard access by changing permissions for a user, group, or role.</li> </ul>

#### **Default Dashboard Permissions**

The table below lists the predefined permission sets that are applied by default when a new dashboard is created. Besides the sysadmin user, the dashboard creator is granted **Admin** privileges by default. The Everyone role is granted **View** privileges by default. No other users, groups, or roles have dashboard permissions assigned by default.

User or Role	Applied Permission Set
Sysadmin	Admin
Dashboard Creator	Admin
Everyone Role	View
# **Lens Level Permissions**

Lens level permissions affect a user's ability to view, modify, delete, design, or configure lenses and lens permissions. There are three predefined lens permission sets that can be assigned to an Anzo user or group. You also have the option to customize the set of permissions that are applied. While dashboard level permissions can affect a user's ability to clone a lens, the appropriate lens level permissions are required to be able to perform functions such as deleting or redesigning a lens.

The table below lists the predefined permission sets and describes the privileges that are granted for each permission that is part of the predefined set:

Set	Permission	Allows a user to:		
View	Read	<ul><li>Search for and open accessible lenses.</li><li>View lens Properties.</li><li>Export lenses.</li></ul>		
Modify	In addition to the <b>Read</b> permission described above, the <b>Modify</b> set includes the <b>Write</b> and <b>Delete</b> permissions described below.			
	Write	<ul><li>Use the lens Designer to change the lens.</li><li>Rename the lens.</li><li>Clone the lens.</li></ul>		
	Delete	Delete the lens.		
Admin	In addition to the <b>Read</b> , <b>Write</b> , and <b>Delete</b> permissions described above, the <b>Admin</b> set includes the <b>Manage</b> permission described below.			
	Manage	• The Manage permission relates only to the Security tab. If a user has this permission, they can modify		

Set	Permission	Allows a user to:
		lens access by changing permissions for a user, group, or role.

#### **Default Lens Permissions**

The table below lists the predefined permission sets that are applied by default when a new lens is created. Besides the sysadmin user, the lens creator is granted **Admin** privileges by default. The Everyone role is granted **View** privileges by default. No other users, groups, or roles have lens permissions assigned by default.

User or Role	Applied Permission Set
Sysadmin	Admin
Lens Creator	Admin
Everyone Role	View

## **Configure Dashboard or Lens Permissions**

This section provides instructions for modifying dashboard or lens properties to grant or restrict access to your dashboards and lenses.

### Note

Data can be restricted at a higher level than a dashboard. Though users might have access to view your dashboards and lenses, graphmart permissions determine whether they can view the data that the dashboard displays.

- 1. Open the dashboard for which you want to modify access.
- 2. Open the Properties dialog box for the either dashboard or for a specific lens:
  - To change access at the dashboard level, click **Dashboard** in the main toolbar and select **Properties**.

### Note

Sharing a dashboard automatically shares the lenses in that dashboard.



• To change access for a lens in the dashboard, click the lens to display it and then click the **Properties** button in the object toolbar.



3. In the Properties dialog box, click the **Security** tab. This tab lists the available groups and users who can view this dashboard or lens. The image below shows the Security tab at the dashboard level:

Listings		×
General Security Advanced		
User and group access - Set user and group privileges.	+ Access Details Permissions for Everyone	
Severyone everyone View	Access	
	Other	
+ Add a user or group		
	CANCEL	SAVE

- 4. Select a user or user group to manage, and then modify any of the following options:
  - **Remove a user or group**: Click the delete icon (X) next to the user or group.
  - Add a user or group: Click Add a user or group. On the User/Group Browser dialog box, select the users or groups that you want to add. Then click OK.



Access Details: Select the access level for the selected user or group. Refer to
 Dashboard Level Permissions or Lens Level Permissions above for details about each of

the access options.

- ° None: No permissions set for the selected dashboard or lens.
- View: Grants the View predefined permission set for the selected dashboard or lens.
- Modify: Grants the Modify predefined permission set for the selected dashboard or lens.
- Admin: Grants the Admin predefined permission set for the selected dashboard or lens.
- Other: Enables you to set custom access levels for the selected dashboard or lens.
   Select the checkboxes to enable any combination of the following permissions:
   Read, Write, Delete, or Manage (administrator permissions).
- 5. Click **Save** to save the changes.

To get a URL to your dashboard that you can send to users, click **Dashboard** in the main toolbar and select **Share**. The Share Dashboard dialog box opens and displays a URL for the dashboard. You can copy the link and send it to users.

# Access Data with the Query Builder

The Query Builder in the user interface provides options for accessing data in various data sources. The Query Builder includes a **Find** option that enables users to search for quads by specifying a single subject, object, predicate, or graph name. It also includes a **Query** option that enables users to write, run, and save SPARQL queries. The topics in this section provide information about accessing data using the Query Builder.

In this section:	
Running SPARQL Queries in the Query Builder	. 907
Searching for Quads in the Query Builder	

# **Running SPARQL Queries in the Query Builder**

The Query Builder includes a Query tab for writing and running SPARQL queries. The query editor provides syntax assistance, type-ahead suggestions for model entity names, and automated prefix creation and query formatting for readability. It also includes the option to save queries for later use.

The Query tab supports running queries against the following data sources:

- Graphmarts and specific data layers within graphmarts
- Linked Datasets
- Data sources: Anzo System Data Source, AnzoGraph, Anzo System Tables, Data Profiling Metrics, LDAP Primary Data Source

### Note

To ensure that queries perform well and do not consume too many resources on the system, keep the following guidelines in mind when developing and testing queries:

- Set a limit on the number of results to return.
- Avoid cross-product joins
- Consider using VALUES clauses instead of FILTER clauses.
- When retrieving a large number of values, use subqueries instead of OPTIONAL clauses.

For query templates and additional details about best practices, see SPARQL Best Practices and Query Templates.

Follow the instructions below to write and run SPARQL queries against any of the supported data sources.

1. In the Anzo application, expand the **Access** menu and click **Query Builder**. Anzo displays the query editor.

Query	Find	
Queries	* ←	Save
Add 👻		Target No target has been selected yet.
Q Search		1
	ф 🗆 🗖	
<ul> <li>My Queries</li> </ul>	:	
		Run Clear Format  More Skip server Cache S00

- 2. At the top of the screen, click the edit ( $\checkmark$ ) icon next to the Target data source:
  - To query data that is in a graphmart, select the Graphmart radio button.

Graphmart O Linked Dataset O Datasource	
Select Graphmart	~
Select Layers	~

Click the **Select Graphmart** drop-down list and select the graphmart to query. If you want to narrow the scope of the query by selecting one or more data layers in the graphmart, click the **Select Layers** drop-down list and select the layer or layers to target.

• To query data that is in the Datasets catalog, select the Linked Dataset radio button.

C oropinitati C canto	leu balaser () balasource	
Select linked dataset		~

Click the **Select linked dataset** drop-down list and select the linked dataset to query.

• To run queries against the system data source, data profiling metrics, Anzo system tables, LDAP server, or AnzoGraph, select the **Datasource** radio button.

○ Graphmart ○ Linked Dataset    ● Datasource	
System Datasource	x   ~
Online	
Named Datasets	
Add a new Named Dataset here	ADD
Named Graphs	
http://openanzo.org/namedGraphs/reserved/graphs/ALL	×
Add a new Named Graph here	ADD
Default Named Graphs	
http://openanzo.org/namedGraphs/reserved/graphs/ALL	×
Add a new Default Named Graph here	ADD

Click the **Datasource** drop-down list and select the target source:

- <sup>o</sup> Select **System Datasource** to search the local Anzo system volume.
- Select the name of an AnzoGraph instance to search for data in graphmarts that are loaded to that instance.
- <sup>o</sup> Select **Data Profiling Metrics** to search the data metrics volume.
- <sup>o</sup> Select LDAP Primary Datasource to search the directory server.
- <sup>o</sup> Select **System Tables** to search Anzo system table data.

By default, the Named Graphs and Default Named Graphs values are set to all named graphs (http://openanzo.org/namedGraphs/reserved/graphs/ALL). If you want to narrow the scope of the query, you can replace the values with specific graph URIs. To list multiple graphs, separate URIs with a space.

- 3. When you have finished configuring the target, click outside of the dialog box to return to the Query screen.
- In the text box below the target, compose the SPARQL query. See SPARQL Best Practices and Query Templates for tips on writing queries. For information about the supported SPARQL functions, see Function and Formula Reference.

## Tip

When adding PREFIX statements, once you type **prefix** followed by a space Anzo displays a tooltip that lists all of the global prefixes that are defined for your system. Clicking a prefix in the list inserts a PREFIX statement into the query. For example:

1 v prefix dc: <http://purl.org/dc/elements/1.1/</pre> dcterms: <http://purl.org/dc/terms/> foaf: <http://xmlns.com/foaf/0.1/> owl: <http://www.w3.org/2002/07/owl#> rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> rdfs: <http://www.w3.org/2000/01/rdf-schema#> xsd: <http://www.w3.org/2001/XMLSchema#>

In addition, typing the abbreviation for a global prefix followed by a colon (:) automatically inserts the PREFIX statement into the query without opening the tooltip. When typing entity URIs in the WHERE clause, the query builder also offers suggestions by listing the properties in the data source. You can click an item in the list to insert that entity. For example:



When a red exclamation mark icon (①) is displayed next to a line number, you can hover the pointer over the icon to view guidance on how to resolve the issue. For example:



- 5. If you want to format the query for readability, click the **Format** button and select from the following options:
  - Format: Auto-creates prefixes, inserts URI abbreviations, and restructures the query for readability.
  - Format with simplified variable names: Auto-creates prefixes, inserts URI abbreviations, simplifies variable names by changing them to ?\_var1, ?\_var2, ?\_varN, and restructures the query for readability.

For example, the image below shows a query before it is formatted.

1 56	elect Porigin	?destina	tion ?flight		
2 * wł	here {				
3	?s <http: c<="" td=""><td>ambridges</td><td>semantics.com/Layer/f9df7ce18bc1478</td><td>199d3085f0bc7f84e/Model#Flights10</td><td>0k.originAirport&gt; ?origin .</td></http:>	ambridges	semantics.com/Layer/f9df7ce18bc1478	199d3085f0bc7f84e/Model#Flights10	0k.originAirport> ?origin .
4	<pre>?s <http: c<="" pre=""></http:></pre>	ambridges	emantics.com/Layer/f9df7ce18bc1478	i99d3085f0bc7f84e/Model#Flights10	<pre>3k.destinationAirport&gt; ?destination .</pre>
5	?s <http: c<="" td=""><td>ambridges</td><td>emantics.com/Layer/f9df7ce18bc1478</td><td>i99d3085f0bc7f84e/Model#Flights10</td><td>3k.flightNumber&gt; ?flight .</td></http:>	ambridges	emantics.com/Layer/f9df7ce18bc1478	i99d3085f0bc7f84e/Model#Flights10	3k.flightNumber> ?flight .
6 }					
7					
Dry Rur	Run	Clear	Format	ip server cache	Limit results to: 500 🛩
Dry Rur	Run	Clear	Format	ip server cache	Limit results to: 500 🛩
Dry Rur	Run	Clear	Format Format with simplified variable names	ip server cache	Limit results to: 500 ~

After the query is formatted, prefixes and URI abbreviations are added. For example:



- 6. If the query is an INSERT or DELETE query, the Dry Run button becomes active. You can click **Dry Run** to do a test run of the update. In a test run, Anzo runs a version of the query where INSERT or DELETE is replaced with CONSTRUCT, and the results report the number of statements that the query affects, i.e., the number of additions or removals per graph. If the results are unexpected, you can adjust the query before clicking running the query and committing the updates.
- If necessary, change the query limit. By default, query results are limited to 500. To adjust the limit, click the Limit results to drop-down list below the query editor and select a value. For example:



8.	To run the query, click <b>Run</b> .	The results appear at the bottom of the screen. For	or example:
----	--------------------------------------	---	-------------

Results (100) Executed in 0.367 s	<b>V</b> :	se prefixes 🔽 S	how native Clear	Download
origin	destination	flight	distance	airline
LAX	PBI	2336	2330	AA
SEA	ANC	135	1448	AS
PDX	MSP	1670	1426	DL
LAS	DFW	214	1055	NK
LAS	IAH	298	1222	NK
ANC	SEA	134	1448	AS
SJU	JFK	304	1598	B6
IAG	FLL	647	1176	NK
BQN	EWR	1162	1585	AU
BQN	MCO	1030	1129	B6
SJU	MCO	2134	1189	B6
ORD	РНХ	602	1440	US
GEG	SEA	695	224	AS
HNL	ITO	102	216	HA

## Тір

You can click any value in the result list to copy that value to the clipboard.

9. To save the query for later use, click the **Save** button at the top of the screen. Anzo displays the New Query dialog box.

Title *	
Description	
occomption -	
Folder	
My Queries	×   ~

10. In the dialog box, specify a name for the query in the **Title** field and an optional description in the **Description** field.

11. You can click the **Folder** drop-down list to select a different folder or create a new one to save the query in. Then click **Save**. The query is saved and added to the **Queries** panel on the left side of the screen.

# Searching for Quads in the Query Builder

The Query Builder includes a Find tab for searching for data by specifying a single subject, object, predicate, graph name or any combination of those elements. Statements that match the search criteria are returned in quads, and the screen includes quick filters that enable users to toggle filters on and off to show or hide any of the quad elements. The Find tab supports searches against the following data sources:

- Anzo System Data Source
- AnzoGraph
- Anzo System Tables
- Data Profiling Metrics
- LDAP Primary Data Source

When finding data in the system data source, users have the option to modify or delete statements directly in the user interface. Follow the instructions below to find data in any of the supported data sources.

 In the Anzo application, expand the Access menu and click Query Builder. Anzo displays the query editor.

Query	Find		
Queries	* ←		Save
Add 👻		arget No target has been selected yet. 🥜	
Q Search		1	0
	000		
<ul> <li>My Queries</li> </ul>	:		
		Dry Run Clear Format + More +	Skip server Limit results to: cache 500 ~

#### 2. Click the Find tab.

Find			
		×	[ •
Predicate	Object	Graph	
		Clear Add Statements F	Find
	Find Predicate	Find Predicate Object	Find

- 3. Click the **Select a Datasource** drop-down list and select the data source that you want to search.
  - Select System Datasource to search the local Anzo system volume.
  - Select the name of an AnzoGraph instance to search for data in graphmarts that are loaded to that instance.
  - Select Data Profiling Metrics to search the data profiling metrics volume.
  - Select LDAP Primary Datasource to search the directory server.
  - Select System Tables to search Anzo system table data.

- 4. Follow the guidelines below to specify the data to find in the data source:
  - Specify any subject, predicate, object, or graph name in the appropriate field. You can specify a value for one field in the quad or any combination of fields.
  - Any URIs and/or literal values that you specify must match the value in the data. Partial values, wildcard characters, and regular expressions are not supported.
  - If you want to get a list of all of the statements in the data source, you can leave all of the fields blank.
- 5. Click **Find** to search for the statements that match the search criteria. Anzo displays the matching statements. For example:

Subject	Predicate		Object BOS		Grap	h	
					Clear	Add Statements	Find
Results	; (50)	Show native	🖌 Subject 🛛	Predicate	🗹 Object	🗹 Named Graph	0
EDIT	DELETE						
	Subject $\downarrow$	Predicate		Object	Named	Graph	
	« http://cambridgesemantics.com/La yer/f9df7ce18bc1478d99d3085f0bc7 f84e/Flights10k/0dWi2IJAKiRsnW1sY 190ie/0000_00000/0008610 »	« http://cambridgese m/Layer/f9df7ce18t 85f0bc7f84e/Model estinationAirport »	emantics.co oc1478d99d30 #Flights10k.d	« BOS »	http://ca tics.com e18bc14 bc7f84e	mbridgeseman n/Layer/f9df7c 178d99d3085f0	:
	« http://cambridgesemantics.com/La yer/f9df7ce18bc1478d99d3085f0bc7 f84e/Flights10k/0dWi2IJAKiRsnW1sY 190ie/0000_00000/0008656 »	« http://cambridgese m/Layer/f9df7ce18t 85f0bc7f84e/Model iginAirport »	emantics.co oc1478d99d30 #Flights10k.or	« BOS »	http://ca tics.com e18bc14 bc7f84e	ambridgeseman n/Layer/f9df7c 178d99d3085f0	:
	« http://cambridgesemantics.com/La yer/f9df7ce18bc1478d99d3085f0bc7 f84e/Flights10k/0dWi2IJAKiRsnW1sY 190le/0000_00000/0008681 »	« http://cambridgese m/Layer/f9df7ce18t 85f0bc7f84e/Model iginAirport »	emantics.co oc1478d99d30 #Flights10k.or	« BOS »	http://ca tics.com e18bc14 bc7f84e	mbridgeseman n/Layer/f9df7c 178d99d3085f0	:
	« http://cambridgesemantics.com/La yer/f9df7ce18bc1478d99d3085f0bc7 f84e/Flights10k/0dWi2IJAKiRsnW1sY 190ie/0000_00000/0008682 »	« http://cambridgese m/Layer/f9df7ce18t 85f0bc7f84e/Model iginAirport »	emantics.co oc1478d99d30 #Flights10k.or	« BOS »	http://ca tics.com e18bc14 bc7f84e	mbridgeseman n/Layer/f9df7c 178d99d3085f0	:
	« http://cambridgesemantics.com/La	« http://cambridgese	emantics.co		http://ca	ambridgeseman	

- 6. The following options are available for working with the results:
  - To filter results by showing or hiding parts of the quads in the statements, you can select or clear the following checkboxes above the results.



Clearing a checkbox hides that part of the quad in the result list. You can display the item again by selecting the checkbox.

- To modify the search parameters, you can click any of the graph, subject, predicate, or object values in the results. The search is automatically run again using only the value that you clicked.
- If the source that you searched is the **System Datasource**, you can edit, delete, or add statements directly. See System Datasource Options below for details.

### **System Datasource Options**

This section provides information about editing, deleting, and adding statements on the Find screen.

### Note

Though the options described below are available for all data sources, adding, deleting, or editing statements is only successful when the data source is **System Datasource**.

- Editing a Statement
- Deleting a Statement
- Adding a Statement

#### **Editing a Statement**

To edit a statement, click the menu icon (:) to the right of the statement and select Edit.

Result	is (50)	Show native Subject 🗹	Predicate	🗹 Object 🗹 Named Graph	0
EDIT	DELETE				
	Subject $\downarrow$	Predicate	Object	Named Graph	
	« http://cambridgesemantics.com/La yer/f9df7ce18bc1478d99d3085f0bc7 f84e/Flights10k/0dWi2IJAKiRsnW1sY 190ie/0000_00000/0008610 »	« http://cambridgesemantics.co m/Layer/f9df7ce18bc1478d99d30 85f0bc7f84e/Model#Flights10k.d estinationAirport »	« BOS »	http://cambridgeseman tics.com/Layer/f9df7c e18bc1478d99d3085f0 bc7f84e	Edit Delete
	« http://cambridgesemantics.com/La yer/f9df7ce18bc1478d99d3085f0bc7 f84e/Flights10k/0dWi2IJAKiRsnW1sY 190ie/0000_00000/0008656 »	« http://cambridgesemantics.co m/Layer/f9df7ce18bc1478d99d30 85f0bc7f84e/Model#Flights10k.or iginAirport »	« BOS »	http://cambridgeseman tics.com/Layer/19df7c e18bc1478d99d3085f0 bc7f84e	:

Anzo displays the Edit Statement dialog box. For example:

subject* <http: cam<="" td=""><td>bridgesemantics.com</td><td>n/Layer/f9df7ce</td><td>18bc1478d99d3</td><td>085f0bc7f84e/Fligl</td><td>hts1</td></http:>	bridgesemantics.com	n/Layer/f9df7ce	18bc1478d99d3	085f0bc7f84e/Fligl	hts1
Predicate * <http: cam<="" td=""><td>bridgesemantics.com</td><td>n/Layer/f9df7ce</td><td>18bc1478d99d3</td><td>085f0bc7f84e/Moc</td><td>lel#</td></http:>	bridgesemantics.com	n/Layer/f9df7ce	18bc1478d99d3	085f0bc7f84e/Moc	lel#
Object * " <b>BOS"</b>					
Named Graph I <http: cam<="" td=""><td><sub>RI *</sub> bridgesemantics.com</td><td>n/Layer/f9df7ce</td><td>18bc1478d99d3</td><td>085f0bc7f84e&gt;</td><td></td></http:>	<sub>RI *</sub> bridgesemantics.com	n/Layer/f9df7ce	18bc1478d99d3	085f0bc7f84e>	

Change any of the quad values, and then click Save.

### Important

If you edit URI values, make sure that the modified value is a valid URI.

#### **Deleting a Statement**

To delete a statement, click the menu icon (:) to the right of the statement and select **Delete**. Anzo displays the statement in a confirmation dialog box. For example:



Click **Delete** to remove the statement from the system data source.

### Adding a Statement

To add a quad to the data source, click the Add Statements button at the top of the result list.

Subject	Predicate	Object BOS	Add a quad containing a subject, object, predicate and graph name to the data source to search for.
			Clear Add Statements Find
Results (50)	Show native	🗹 Subject 🔽 Predicate	🗹 Object 🗹 Named Graph []

Anzo displays the Add Statements dialog box.

Add S	Add Statements					
×	Subject *					
	Predicate *					
	Object*					
	Named Graph URI *					
Add						
	CANCEL ADD					

Specify the new quad by adding the subject, predicate, object, and named graph URI in the appropriate fields. Each field is required. URIs must be valid, and the Named Graph URI that you specify must be present in the data source. You cannot add a new named graph. Click **Save** to add the new quad to the data source.

# **Access Data on Demand Endpoints**

The topics in this section provide information about accessing Data on Demand endpoints and using the OData API as well as the Anzo ODBC and JDBC drivers. For information about creating endpoints, see Creating Data on Demand Endpoints.

In this section:	
Accessing an Endpoint Programmatically	.922
Accessing an Endpoint from an Application	.926
OData Reference	944

# Accessing an Endpoint Programmatically

This topic provides guidance on accessing Data on Demand endpoints programmatically by showing some example implementations using R and Python.

- Authentication and Data Access
- Access an Endpoint with R (RStudio)
- Access an Endpoint with Python (Linux Terminal)

## **Authentication and Data Access**

Connections to Data on Demand endpoints must be authenticated. Users can submit their Anzo username and password when accessing data. Ultimately the data that is available to users from OData endpoints is subject to the security and composition of the graphmart as configured in Anzo.

## Access an Endpoint with R (RStudio)

The following example shows how to connect to an OData endpoint from RStudio. The example uses the R programming language to access a Data on Demand endpoint and pull in data via a standard dataframe. New or existing R scripts can then be used with the data.

The first step in accessing data from RStudio is to prepare the R script that will construct the target URL and retrieve the resulting information via HTTP. The example script below accesses a preconfigured "Sample Data" endpoint. The script has sections for filtering the results as well as expanding the selection to include information from multiple classes:

```
require("httr")
require("jsonlite")
require("rstudioapi")
user <- rstudioapi::showPrompt("Username", "Enter Anzo username", "sysadmin")
pw <- rstudioapi::askForPassword(paste("Enter password for",user,sep=" "))
## Data on Demand endpoint
odata <- "https://10.100.0.10/dataondemand/Sample-Graphmart/Sample-Data"
## Start from Probe class</pre>
```

```
startClass <- "Probe?"</pre>
## Filter results for Homo sapiens species
filterKw <- "$filter="</pre>
filterVal <- "Species eq 'Hs'"
urlify <- URLencode(filterVal)</pre>
filterStr <- paste(filterKw,urlify,sep="")</pre>
## Select properties of interest (FeatureID) from base class
selectKw <- "&$select="</pre>
selectVal <- "FeatureID"</pre>
selectStr <- paste(selectKw,selectVal,sep="")</pre>
## Select properties of interest (symbol) from Gene class
## via corresponds to property on base Probe class
expandKw <- "&$expand="
expandClass <- "corresponds to"
expandProps <- "symbol"
expSelStr <- "$select="
expandStr <- paste(expandKw,expandClass,"(",expSelStr,expandProps,")",sep="")</pre>
## Specify format
format <- "&$format=json"</pre>
## Generate OData URL using fragments above
url <- paste(odata,startClass,filterStr,selectStr,expandStr,format,sep="")</pre>
## Access OData endpoint
resultRaw <- GET(url, (authenticate(user,pw, type = "basic")))</pre>
resultTxt <- content(resultRaw, "text")</pre>
resultJson <- fromJSON(resultTxt, flatten = TRUE)</pre>
print(url)
## Read results into dataframe
resultDataFrame <- as.data.frame(resultJson)</pre>
View(resultDataFrame)
```

Executing the above R script from RStudio results in a dataframe that represents columns from the **Probe** and **Gene** classes.

# Access an Endpoint with Python (Linux Terminal)

Many users have existing Python scripts to use with data in Anzo or a familiarity with Python that would make exploring, retrieving, and leveraging the data easier. The following example shows how to connect to an OData endpoint by executing a Python script from a Linux terminal.

The first step in accessing data using Python is to prepare the Python script that will construct the target URL and retrieve the resulting information via HTTP. The example script below accesses a pre-configured "Sample Data" endpoint. The script has sections for filtering the results as well as expanding the selection to include information from multiple classes (the same filter and class properties that were used in the R example above).

```
import requests
import getpass
from urllib.parse import urlparse
un = getpass.getpass(prompt='Username: ')
pw = getpass.getpass(prompt='Password: ')
## OData endpoint
# Data on Demand URL
odata = 'https://10.100.0.10/dataondemand/Sample-Graphmart/Sample-Data/'
## Start from Lease class
startClass = "Probe?"
## Filter results
filterKw = "$filter="
filterVal = "Species eq 'Hs'"
urlify = urlparse(filterVal)
filterStr = filterKw + urlify.geturl()
## Select properties of interest (start date, missed payments, lease status) from base
class
selectKw = "&$select="
selectVal = "FeatureID"
selectStr = selectKw + selectVal
## Select properties of interest (name, social security number, credit score) from
Individual class
expandKw = "&$expand="
```

```
expandClass = "corresponds_to"
expandProps = "symbol"
expSelStr = "$select="
expandStr = expandKw + expandClass + "(" + expSelStr + expandProps + ")"
## Specify format
format = "&$format=text/csv"
## Generate OData URL using fragments above
url = odata + startClass + filterStr + selectStr + expandStr + format
## Access OData endpoint
r = requests.get(url, auth=(un, pw), verify=False)
print("URL")
print(url)
print("CONTENT")
print(r.content.decode('unicode escape'))
print(type(r))
print(type(r.content))
```

In this example, the output is returned in CSV format (rather than JSON like the R example).

# Accessing an Endpoint from an Application

Since Anzo's Data on Demand service conforms to the OData standard, any tool that supports the OData V4 REST API can access a Data on Demand endpoint to leverage data in Anzo. In addition, applications that support ODBC or JDBC APIs can use the Anzo CData ODBC or JDBC drivers to interact with Data on Demand endpoints. This capability enables users to leverage the benefits of Anzo's semantic layer, data model, and data blending capabilities in their favorite analytics tools.

This topic provides information about accessing Data on Demand endpoints from third-party applications.

- JDBC Driver Considerations
- Authentication and Data Access
- Accessing Data via the OData API
- Downloading the Anzo ODBC and JDBC Drivers
- JDBC Driver Documentation

### **JDBC Driver Considerations**

This section describes important items to consider when using JDBC clients for accessing Data on Demand endpoints:

- Join Performance
- Querying Multi-Valued Properties
- Working with Long Column Names

#### Join Performance

To join results from multiple classes, Cambridge Semantics strongly recommends using OData or SPARQL. Hi-Res Analytics and SPARQL are designed to quickly return large results from multiple classes and should be strongly considered for these use cases. Joins on large data sets are also well-supported with OData when best practices around paging are applied.

You can also join tables upstream in AnzoGraph by creating data layers. For example, you can create a view that joins the data using a CONSTRUCT query. The view becomes available as an OData table. For information about view steps, see Construct a View of the Data (View Step).

In addition, Custom Data on Demand endpoints (sometimes called Table endpoints) enable you to join classes, add filters, and apply functions to properties during endpoint creation. The tables that you create are automatically translated to SPARQL queries that create views in AnzoGraph, allowing you to perform complex analytics on the graph yet generate results in the tabular format that BI tools expect.

Because the JDBC driver generates multiple OData queries and joins the results in memory, SQL queries that include JOINs on large data sets may take a very long time to complete. When using the JDBC driver, Cambridge Semantics recommends that you query one class at a time and then use the BI tool to do analytics on the returned data. For more information, see JDBC Performance Details below.

### **Querying Multi-Valued Properties**

Some applications do not directly support Anzo's RDF graph data structures. For example, sometimes the JDBC driver presents multi-valued properties as arrays, which are not supported by some BI tools. When creating a Data on Demand endpoint for a graphmart that includes multi-valued properties, consider denormalizing the results to expand the properties into new rows so that they can be viewed in BI tools. For more information, see Creating Data on Demand Endpoints.

### Working with Long Column Names

By default, the JDBC driver creates column names based on the property labels in the data model. The property labels can be too long for some clients. For example, Informatica is limited to 128 characters. When ingesting data from a tabular source, the label is a concatenation of the table and column name. Users may need to shorten the property labels to work with JDBC clients. If the label is missing, Anzo uses the localName of the IRI. For information about configuring the column names to be used for a Data on Demand endpoint, see Creating Data on Demand Endpoints.

## **Authentication and Data Access**

Connections to Data on Demand endpoints must be authenticated. Users can submit their Anzo username and password when accessing data. If your applications use single sign-on (SSO) authentication, you can also use SSO with Anzo. When using SSO, the client authenticates the user against the SSO provider and then passes the credentials to Anzo. All data is secured according to the user's SSO profile. For information about the supported SSO providers and instructions on configuring SSO access, see <u>Connecting to an SSO Provider</u> in the Administration Guide.

### Note

Ultimately the data that is available to users from Data on Demand endpoints is subject to the access control configuration of the graphmart in Anzo.

## Accessing Data via the OData API

This section provides guidance on accessing a Data on Demand endpoint from an application that supports the OData REST API. It includes an example that configures an OData connection in TIBCO Spotfire. The example steps can also be applied to OData connections in other similar business intelligence tools.

The first step is to connect to the OData endpoint using the Spotfire Data sources user interface. When setting up the OData connection, the Service URL is the OData/ODBC URL from the Data on Demand endpoint configuration details in Anzo. The OData connection uses the user's Anzo credentials for authentication.

TIBCO	Spotfire	
	Q Files and data	OData Connection ×
	<ul> <li>Recommended</li> <li>Connect to</li> <li>Browse local file</li> <li>Other</li> </ul>	Service URL: Inttps://10.102.0.17/dateondemand/Semple-Movie-Date-Graphmart/Movie/ Authentication method: Username: Sysadmin Password: Password:
	LIBRARY Home Cambridge Semantics Samples	Account key: Help Connect Cancel Data
	Shared with me Shared by me	IBM Netezza     IBM Netezza     Search for more
		Microsoft SQL Server
		Microsoft SQL Server Analy
		OData

Once the connection is established, Sportfire prompts the user to select the classes and properties to work with. In this example, the **FeatureID** property from the **Probe** class and the **symbol** property from the **Gene** class are selected:

varable fables in dafabase.	Views in connection:	Columns in selected view.
ype to search tables	Define Prompting	gene key
Relations + Custom Query + Edit Tables	Probe	i symbol EastroiD
Exposure_Dataset Exposure_Dataset Exposure_Dataset1 Exposure1 Fedure Gene_Rearrangements Leboratory_Dataset1 Laboratory_Dataset1 Laboratory_Dataset1 Laboratory_Dataset1 Laboratory_Dataset1 Medical_History_Dataset1 Medical_History_Dataset1 MedicalHistory1 MedicalHi	Add > < Remove	FeatureID

Once the properties are chosen, the data is loaded in Spotfire and can be used to inform existing analytics and data visualizations or create new ones.

# Downloading the Anzo ODBC and JDBC Drivers

This section provides guidance on accessing Data on Demand endpoints from applications that support ODBC or JDBC APIs. Your Anzo deployment includes CData ODBC and JDBC drivers to use with applications. The first step is to retrieve the appropriate driver for your client. To download a driver, open a web browser and go to the following URL:

```
https://<Anzo_server>/installs/anzodataaccess
```

Where <Anzo\_server> is the Anzo server DNS name or IP address. The Anzo Data Access Software Installation page provides links to download each driver. For example:



Download the appropriate driver to the client server:

- The **CData JDBC Driver for Anzo** is used to connect to Anzo from most Java applications and database management tools.
- The CData ODBC Driver for Anzo for Windows or Mac is for use with applications and database management tools that support open database connectivity, such as Microsoft Excel or Tableau.

Configuring the Driver and Connecting to the Endpoint

This section provides guidance configuring an ODBC or JDBC driver by showing examples of configuring DbVisualizer and Tableau to access a Data on Demand endpoint using Anzo's JDBC driver and configuring Power BI to access an endpoint using the ODBC driver.

### Example JDBC Setup with DbVisualizer

- 1. In DbVisualizer, go to **Tools**  $\rightarrow$  **Driver Manager**.
- 2. In the Driver Manager, click the green plus icon to create a new driver.
- 3. Specify a name for the driver. For example, **Anzo JDBC Driver**.
- 4. In the URL Format field, specify the format jdbc:anzo.
- 5. In the Driver File Paths or Driver jar Files section of the screen, click the folder icon and then browse to and select the directory where you saved the CData JDBC Driver for Anzo cdata.jdbc.anzo.jar file that you downloaded to the server. DbVisualizer reads the jar and sets the Driver Class to cdata.jdbc.anzo.AnzoDriver. For example:

• • •	Driver Manager	
+ + -   10 12   ^ V   🖿		
Driver Name Cache Cache Gache Gache Gache Gache DB2 (DataDirect) DB2 (DataDirect) DB2 (Series (ITOpen) DB2 (Series (ITOpen) Series (ITOpen) Seri	Driver Settings         Name:       Anzo JDBC Driver         URL Format:       jdbc:anzo:         Driver Class:       III cdata.jdbc.anzo.AnzoDriver         Image: Image	►
McKoi	Show Full Path	
MySQL     MySQL     MySQL     DataDirect)     Netezza     NunDR     g = Custom Driver	Driver Meta Data Version: 0.0 Name: CData JDBC Driver for Anzo 2019 Version String: 19.0	

- To connect to the endpoint in DbVisualizer, go to Database → Create Database
   Connection. Click No Wizard when prompted.
- 7. Specify a name for the connection in the **Name** field.
- 8. In the Driver (JDBC) field, select the Anzo JDBC driver connection.
- 9. In the **Database URL** field, specify the JDBC URL from the Anzo Data on Demand endpoint configuration. For example:

jdbc:anzo:URL=https://10.10.0.11/dataondemand/Northwind/Northwind

Overview	Explore		Datasets	Data Layers	Dashboards	Data on Demand	>
REST API Query Endpoints Create New Endpoint	1	₹ ←	Predicate used to retri None Cache Enal	ieve value for name from o bled	lass or property.		*
Northwind		Ĩ	Connection Stri	te Results			
			Service root URLs ut	sed to retrieve class data f	rom this graphmart.		
			ODBC (SQ)	L) JDBC (5	sqL)	Ō	•

10. Under **Authentication**, enter your Anzo user ID and password. You should now be able to connect to the endpoint and view the available schemas. For example:

Connections					
V B Movie-DOD	Connection 🛠 Prope	rties 📙 Database Info 🗧 Data Types			
V Anzo (Default)					
🔻 🚽 Movie-DOD	Connection				
TABLE					
ison 🔛					
MovieActors1	Database				
MovieActors2		🚥 Database URL			
MovieComposers					
MovieDirectors		🗸 Anzo JDBC Driver			
MovieEditors					
MovieProducers	Authentication				
Movies					
VIEW					
Procedures	V Options				
		₹			

Example JDBC Setup with Tableau

- 1. After downloading the CData JDBC Driver for Anzo cdata.jdbc.anzo.jar file, place the .jar in the appropriate directory depending on your operating system:
  - Windows: C:\Program Files\Tableau\Drivers
  - MacOS: ~/Library/Tableau/Drivers
  - Linux: /var/opt/tableau/tableau\_server/data/tabsvc/vizqlserver/Datasources/
- 2. Restart Tableau and then go to Add a Connection  $\rightarrow$  To a Server.
- 3. Click Other Databases (JDBC).
- 4. In the **Database URL** field, specify the JDBC URL from the Anzo Data on Demand endpoint configuration. For example:

jdbc:anzo:URL=https://10.10.0.11/dataondemand/Northwind/Northwind

Overview	Explore		Datasets	Data Layers	Dashboards	Data on Demand	>
REST API Query Endpoints Create New Endpoint	Ŧ	÷	Predicate used to retri None Cache Enal	ieve value for name from o bled	lass or property.		•
Northwind		Ĩ	Denormalia Connection Stri Service root URLs us ODBC (SQ)	ze Results ing sed to retrieve class data f L) JDBC (1 1//10.102.0.17/deteorside	irom this graphmart. SQL)	Ō	-

5. Enter your Anzo username and password and click **Sign In**. You should now be able to connect to the endpoint and view the available schemas.

Example ODBC Setup with Microsoft Power BI

- 1. After downloading the Windows CData ODBC Driver for Anzo executable file, run the executable to start the installation wizard. The wizard guides you through installing the driver.
- 2. At the end of the installation, make sure the **Configure ODBC Data Source** checkbox is selected and click **Finish**. The wizard opens the driver's DNS Configuration screen. For example:

DSN Configuration Data Source Name: CData Anzo Source	Iest Connection	set Connection
Connection Properties		
C Authentication		^
User		
Password		
Caching		
Auto Cache	False	
Cache Connection		
Cache Location	%APPDATA%\CData\Anzo Data Provider	
Cache Metadata	False	
Cache Provider		
Cache Tolerance	600	
Offline	False	
E Firewall		
Firewall Type	NONE	
Firewall Server		
Firewall User		
Firewall Password		
Firewall Port		
Logging		
Logfile		
Max Log File Count		~
URL to the Anzo endpoint.		

3. Under Authentication in Connection Properties, specify the URL, User, and Password to use for connecting to the Data on Demand endpoint. The User and Password are the Anzo username and password to use for authentication, and URL is the OData/ODBC service root URL for the endpoint. You can retrieve the URL from the Data on Demand screen for the endpoint. For example:

Overview	Explore	Datasets	Data Layers	Dashboards	Data on Demand	>
REST API Query Endpoin Create New Endpoint	nts ∓ ←	Predicate used t None	o retrieve value for name from-	class or property.		*
Northwind		Denor	malize Results			
		Connection Service root U	1 String RLs used to retrieve class data	from this graphmart.		ł
		ODBC	(SQL) JDBC (	(SQL)		
		https://10.102	0.17/dataondemand/Northwir	nd/Northwind		

- 4. Click **OK** to save the configuration changes and close the dialog box.
- 5. Next, connect to the ODBC data source from Power BI. Open Power BI and click the **Get Data** button in the tool bar. In the Get Data dialog box, search for "ODBC." For example:

Get Data	
oded ×	All
All	ODBC
Other	
Certified Connectors	Connect Cancel

 The search opens the wizard for creating an ODBC connection to a specified data source. Select CData Anzo Source from the drop-down list. You do not need to configure the advanced options.

From ODBC	
Data source name (DSN)	
CData Anzo Source	•
Advanced options	

7. Click **OK** to create the connection. Power BI opens the Navigator screen. For example:

	Status_Entry	1
y Options *	a status_entry_key Basis	Cat
ODBC (dsn=CData Anzo Source) [1]	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDAyZTBINmItYTI3OS00MD Cumu	lative Phy
Anna [1]	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDA0NzkSODEtZTk52500Mj Period	i Phy
Anzo [1]	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDA1OTIhYWYtNmJhNC000 Cumul	lative Co
InitialSample [1]	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDBhOWI1MWQt0GJjYy00 Cumu	lative Phy
Status_Entry	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDEwZmVhNzAtYWNj2S00F Cumul	lative Co
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDEzY2RhMzQtZTVjOS00YT Period	i Co
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDEzZDlyYjMtYTI0MS00YW Cumul	lative Phy
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDE2MGE0MzEtNWY5Mi00 Period	1 Phv
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDE2MzRJN2ltMJNiY500ND' Period	f Co
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDFkZGIwMGltZWQwZC00f Cumul	lative Phy
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDIwNDZm0GUtNGQ5Mi0( Cumul	lative Phy
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDIyYzcyOGItY2YzNy00ZTkz Period	i Phy
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDI0MGJI0TgtYjcyMi00YW1 Cumul	lative Phy
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDI3ZjgwZWQtNGE2MC002 Cumul	lative Phy
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDJkYTcyZDMtYWEyMy00% Cumul	lative Co
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDJmODdkZDEtZWZINy00N Cumul	lative Phy
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDMwOTk3NjUtMTMxZi00I Period	i Phy
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDMxMmUzN2MtYzBjY500l Period	f Phy
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDMyZDRkMzMtZmVjY500f Cumul	lative Phy
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDM0YzQ3YmMtMDA2Zi00 Period	i Phy
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDM4MjUwNDUtMWE2MC Cumul	lative Co
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDNjM2JIZjItMGFkNS00Y2h Period	1 Co
	aHR0cDovL2NzaS5jb20vU3RhdHVzRW50cnkvMDNkNGRhODEtZTJhYy00M Cumul	lative Phy
	<	>

Under Display Options, the top level container in the view represents the ODBC driver, the **Anzo** item represents the database, and the **InitialSample** item represents the schema. Each table is represented as a table entry under the schema. In the example above there is one table. If you select a table, sample data from that table is displayed on the right side of the screen. To load table(s), select the checkbox for each table and click the **Load** button. You can also use the advanced features of Power BI to transform the data as you load it into the tool.

## **JDBC Driver Documentation**

This section provides a quick reference for JDBC driver support.

- For the complete JDBC driver documentation, see CData JDBC Driver for Anzo.
- For the complete ODBC driver documentation, see <u>CData ODBC Driver for Anzo</u>.

### **SQL** Compliance

The JDBC driver supports most of the standard operations for querying data. The exceptions are listed below.
- The driver does not currently support transactions.
- The driver does not support batching of SQL statements.
- The driver has support for inserting, updating, and deleting records. However, performing updates via the driver can have unexpected consequences.

For more information about SQL compliance, see the <u>SQL Compliance</u> section in the CData JDBC Driver documentation.

### **JDBC Performance Details**

By default, the JDBC driver offloads to Anzo as much of the SELECT statement processing as possible and then processes the rest of the query locally in memory.

- For joins, the driver generates multiple OData queries and joins the results in memory. As a result, SQL queries that include JOINs can take up to several minutes to complete.
- For aggregates, the driver retrieves all rows necessary to process in memory.
- For predicates, the driver determines which clauses Anzo supports and sends them to Anzo to retrieve the smallest possible superset of rows that would satisfy the query. It then filters the rest of the rows client-side.
- The driver's SupportEnhancedSQL setting can be disabled to limit SQL execution to only what the Anzo API supports. For more information, see the <u>Support Enhanced SQL</u> section in the CData JDBC Driver documentation.

### Tip

To determine which query capabilities the driver can offload to the Anzo API, you can query the **sys\_sqlinfo** system table. The table contains information about the functionality that is supported by the connected source. For example:

```
SELECT * FROM sys_sqlinfo WHERE name='AGGREGATE_FUNCTIONS'
or name = 'COUNT' or name = 'SUPPORTED_OPERATORS' or name = 'GROUP_BY'
or name = 'OUTER_JOINS' or name = 'OJ_CAPABILITIES' or name = 'SUBQUERIES'
or name = 'STRING_FUNCTIONS' or name = 'NUMERIC_FUNCTIONS'
or name = 'TIMEDATE_FUNCTIONS';
```

For more information, see the <u>sys\_sqlinfo</u> section in the CData JDBC Driver documentation.

### **Data Caching**

Due to the client-side in-memory processing of aggregates and joins, the performance of queries against extremely large data sets may suffer. If this is a common use case, consider leveraging caching in the JDBC driver. If the driver maintains a local copy of the data, it reduces the number of API calls and can increase performance for long-running queries. For more information, see the <u>Caching Data</u> section in the CData JDBC Driver documentation.

### **Supported SELECT Statement Clauses**

The following list shows the supported SELECT statement clauses. For more information, see the <u>SELECT Statement</u> section in the CData JDBC Driver documentation.

- SELECT
- INTO
- FROM
- JOIN
- WHERE
- GROUP BY
- HAVING
- UNION
- ORDER BY
- LIMIT

### Supported Aggregate Functions

The following list shows the supported aggregate functions. For more information, see the <u>Aggregate Functions</u> section in the CData JDBC Driver documentation.

- COUNT
- COUNT\_DISTINCT
- AVG
- MIN
- MAX
- SUM

### **Supported Joins**

The following list shows the supported JOIN types. For more information, see the <u>JOIN Queries</u> section in the CData JDBC Driver documentation.

- Inner Join: Selects only the rows from both tables that match the join condition.
- Left Join: Selects all of the rows in the FROM table and only matching rows in the JOIN table.

### **SQL** Function Reference

The JDBC driver provides implementations of the following common SQL functions. For more information, see the <u>SQL Functions</u> section in the CData JDBC Driver documentation.

### Note

The driver interprets all function input as either column names or strings. Therefore, all string literals must be escaped with single quotes. For example, SELECT DATENAME

('yy',GETDATE())).

### **String Functions**

- ASCII(character\_expression)
- CHAR(integer\_expression)
- CHARINDEX(expressionToFind ,expressionToSearch [,start\_location ])
- CONCAT(string\_value1, string\_value2 [, string\_valueN])
- CONTAINS(expressionToSearch, expressionToFind)

- ENDSWITH(character\_expression, character\_suffix)
- FORMAT(value, format)
- FROM\_UNIXTIME(time, format, issecond)
- INDEXOF(expressionToSearch, expressionToFind [,start\_location ])
- ISNULL(check\_expression, replacement\_value)
- JSON\_AVG(json, jsonpath)
- JSON\_COUNT(json, jsonpath)
- JSON\_EXTRACT(json, jsonpath)
- JSON\_MAX(json, jsonpath)
- JSON\_MIN(json, jsonpath)
- JSON\_SUM(json, jsonpath)
- LEFT(character\_expression, integer\_expression)
- LEN(string\_expression)
- LOWER(character\_expression)
- LTRIM(character\_expression)
- NCHAR(integer\_expression)
- PATINDEX(pattern, expression)
- QUOTENAME(character\_string [, quote\_character])
- REPLACE(string\_expression, string\_pattern, string\_replacement)
- REPLICATE(string\_expression, integer\_expression)
- REVERSE(string\_expression)
- RIGHT(character\_expression, integer\_expression)
- RTRIM(character\_expression)

- SOUNDEX(character\_expression)
- SPACE(repeatcount)
- STARTSWITH(character\_expression, character\_prefix)
- STR(float\_expression [, integer\_length [, integer\_decimal]])
- STUFF(character\_expression, integer\_start, integer\_length, replaceWith\_expression)
- SUBSTRING(expression, integer\_start, integer\_length)
- TOSTRING(string\_value1)
- TRIM(character\_expression)
- UNICODE(ncharacter\_expression)
- UPPER(character\_expression)
- XML\_EXTRACT(xml, xpath [, separator])

### **Date Functions**

- CURRENT\_DATE()
- CURRENT\_TIMESTAMP()
- DATEADD(datepart, integer\_number, date [, dateformat])
- DATEDIFF(datepart, startdate, enddate)
- DATEFROMPARTS(integer\_year, integer\_month, integer\_day)
- DATENAME(datepart, date)
- DATEPART(datepart, date [,integer\_datefirst])
- DATETIME2FROMPARTS(integer\_year, integer\_month, integer\_day, integer\_hour, integer\_ minute, integer\_seconds, integer\_fractions, integer\_precision)
- DATETIMEFROMPARTS(integer\_year, integer\_month, integer\_day, integer\_hour, integer\_ minute, integer\_seconds, integer\_milliseconds)
- EOMONTH(start\_date [, integer\_month\_to\_add ])

- GETDATE()
- GETUTCDATE()
- ISDATE(date, [date\_format])
- SMALLDATETIMEFROMPARTS(integer\_year, integer\_month, integer\_day, integer\_hour, integer\_minute)
- SYSDATETIME()
- SYSUTCDATETIME()
- TIMEFROMPARTS(integer\_hour, integer\_minute, integer\_seconds, integer\_fractions, integer\_precision)
- YEAR(date)

### Math Functions

- ABS(numeric\_expression)
- ACOS(float\_expression)
- ASIN(float\_expression)
- ATAN(float\_expression)
- ATN2(float\_expression1, float\_expression2)
- CEILING(numeric\_expression)
- COS(float\_expression)
- COT(float\_expression)
- DEGREES(numeric\_expression)
- EXP(float\_expression)
- EXPR(expression)
- FLOOR(numeric\_expression)
- LOG(float\_expression [, base ])

- LOG10(float\_expression)
- PI()
- POWER(float\_expression, y)
- RADIANS(float\_expression)
- RAND([integer\_seed])
- ROUND(numeric\_expression , integer\_length [ ,function ])
- SIGN(numeric\_expression)
- SIN(float\_expression)
- SQRT(float\_expression)
- SQUARE(float\_expression)
- TAN(float\_expression)

# **OData Reference**

The Anzo Data on Demand service follows the <u>OData Version 4.0 specification</u>, which defines the standard URL conventions and query options. This topic provides a quick reference for learning OData basics and viewing the supported string operators and output formats. It also provides some example queries.

- OData URL Conventions
- Supported Query Operators
- Example OData Requests

## **OData URL Conventions**

An OData service URL has three main parts:

- 1. The **Service Root URL** that Anzo provides. The service root URL is the metadata that describes all of the available feeds (tables).
- 2. The optional **Resource Path** that narrows the scope of the available data to the individual table (class) level, property level, or the schema.
- 3. The Query Options for analyzing the data.

For example, the following OData URL shows the service root from the Data on Demand screen in Anzo, a resource path that narrows the scope of the data to the Employees table (class), and query options that filter the result set to show data for the NA region only:

https://10.100.0.10/dataondemand/Northwind-Graphmart/Northwind/Employees?\$filter=contains(Region,			'NA')	
\	/ \	/ \		/
Service Root URL	Resource	Path	Query Options	

OData requests need to be URL-encoded. Typically you can configure programs to encode requests automatically. And browsers encode URLs that are pasted into the address bar.

# **Supported Query Operators**

OData query options are used to dynamically query data via the endpoint and control the amount and order of the data returned. The Data on Demand service supports the following OData query operators. See Example OData Requests below for example queries that employ the operators.

Operator	Description
\$count	Used to count the number of matching resources in the result set.
\$expand	Used to retrieve related data and include it in the results. When you query data via OData, the default response does not include related entities. The \$expand option provides flexibility for exploring data across the data model. It allows the related information to be embedded in the response.
\$filter	Used to filter a result set. The expression specified with \$filter is evaluated for each resource identified by the resource path, and only items where the expression evaluates to true are included in the response.
\$format	Used to specify the output format for the results. The supported formats are text/CSV, JSON, and XML. For example: \$format=json
\$metadata	Used to return the schema, entity set, and property metadata.
\$orderby	Used to return results in ascending (asc) or descending (desc) order. If asc or desc is not specified, solutions are returned in ascending order.
\$select	Used to specify the subset of properties to include in the result set.
\$skip	Used to specify the number of solutions to exclude in the results. The \$top and \$skip OData query options are similar to the LIMIT and OFFSET clauses in SPARQL queries.
\$top	Used to limit the number of solutions that are returned.

# **Example OData Requests**

This section demonstrates the use of OData query operators by providing examples of common types of OData requests.

The examples below are run against a sample graphmart, called **LeagueGM**, that contains data about the teams and players in a small local baseball league. The Data on Demand endpoint is named **LeagueData**. The following service root URL was created by Anzo:

https://10.100.0.10/dataondemand/LeagueGM/LeagueData

For readability, the examples below abbreviate "https://10.100.0.10/dataondemand" to **dataondemand**. In addition, the examples are not URL-encoded.

The data has Leagues, Teams, Players, and Positions classes (or entities in OData). And the image below shows a graph view of the data model. To view the TriG version of the model, click <u>here</u>.



To view the instance data for each class, you can click a link below to view the data for that class. The data is in JSON format.



### **Retrieving Metadata**

The request below retrieves the schema, entity set, and property metadata for the endpoint.

dataondemand/LeagueGM/LeagueData/\$metadata

The results are in XML format. A snippet of the results is shown below. To view the complete response, click <u>here</u>.

```
<?xml version="1.0" encoding="UTF-8"?>
<edmx:Edmx Version="4.0" xmlns:edmx="http://docs.oasis-open.org/odata/ns/edmx">
  <edmx:DataServices>
    <Schema xmlns="http://docs.oasis-open.org/odata/ns/edm" Namespace="Feeds">
      <EntityContainer Name="Default">
        <EntitySet Name="Leagues"
EntityType="com.cambridgesemantics.ont.autogen.LeagueDict.LeagueData.Leagues">
          <NavigationPropertyBinding Path="LeagueToTeam" Target="Teams"/>
        </EntitySet>
        <EntitySet Name="Teams"
EntityType="com.cambridgesemantics.ont.autogen.LeagueDict.LeagueData.Teams">
          <NavigationPropertyBinding Path="TeamToLeague" Target="Leagues"/>
          <NavigationPropertyBinding Path="TeamToPlayer" Target="Players"/>
        </EntitySet>
        <EntitySet Name="Positions"
EntityType="com.cambridgesemantics.ont.autogen.LeagueDict.LeagueData.Positions">
          <NavigationPropertyBinding Path="PositionToPlayer" Target="Players"/>
        </EntitySet>
        <EntitySet Name="Players"
EntityType="com.cambridgesemantics.ont.autogen.LeagueDict.LeagueData.Players">
          <NavigationPropertyBinding Path="PlayerToPosition" Target="Positions"/>
          <NavigationPropertyBinding Path="PlayerToTeam" Target="Teams"/>
        </EntitySet>
      </EntityContainer>
    </Schema>
```

### **Counting an Entity**

The request below returns the number of teams in the graphmart. Adding the resource path **Teams** to the request narrows the scope to the Teams entity (or class in Anzo).

 $\tt data on demand/LeagueGM/LeagueData/Teams/\$count$ 

#### Result

4

### This request returns the number of players:

dataondemand/LeagueGM/LeagueData/Players/\$count

### Result

12

### Counting a Property of an Entity

The request below counts the number of players on the Al Thomas team. The request uses the team\_key to identify the team and the TeamToPlayer to identify each player.

```
dataondemand/LeagueGM/LeagueData/Teams
('aHR0cDovL2NzaS5jb20vVGVhbXMvMQ')/TeamToPlayer/$count
```

### Result

3

This request counts the number of positions played by James Smith:

```
dataondemand/LeagueGM/LeagueData/Players
('aHR0cDovL2NzaS5jb20vUGxheWVycy8y')/PlayerToPosition/$count
```

#### Result

2

### Filtering Data via Text Search

The request below filters the results to show data for the TeamName that equals "Black Sox." The request also returns results in JSON format:

dataondemand/LeagueGM/LeagueData/Teams?\$filter=TeamName eq 'Black Sox'&\$format=json

### Result

```
"TeamName": "Black Sox",
    "teamtoplayer_key": [
        "aHR0cDovL2NzaS5jb20vUGxheWVycy80",
        "aHR0cDovL2NzaS5jb20vUGxheWVycy81",
        "aHR0cDovL2NzaS5jb20vUGxheWVycy82"
    ]
}
]
```

### This request filters the data to find the players whose name contains "Ted."

```
dataondemand/LeagueGM/LeagueData/Players?$filter=contains(PlayerName,'Ted')
```

The request can also use "startswith" in place of contains to filter specifically for player names that start with "Ted."

```
dataondemand/LeagueGM/LeagueData/Players?$filter=startswith(PlayerName,'Ted')
```

### Result

```
{
  "@odata.context":
"https://10.100.0.10/dataondemand/LeagueGM/LeagueData/$metadata#Players",
  "value": [
    {
      "players key": "aHR0cDovL2NzaS5jb20vUGxheWVycy8xMA",
      "playertoposition key": [
          "aHR0cDovL2NzaS5jb20vUG9zaXRpb25zLzM",
          "aHR0cDovL2NzaS5jb20vUG9zaXRpb25zLzI"
      ],
      "PlayerId": 10,
      "playertoteam key": [
          "aHR0cDovL2NzaS5jb20vVGVhbXMvNA"
      ],
      "PlayerName": "Ted James",
      "DefensiveRating": 92.55
    },
      "players key": "aHR0cDovL2NzaS5jb20vUGxheWVycy84",
      "playertoposition key": [
          "aHR0cDovL2NzaS5jb20vUG9zaXRpb25zLzI",
          "aHR0cDovL2NzaS5jb20vUG9zaXRpb25zLzEw"
```

```
],
   "PlayerId": 8,
   "playertoteam_key": [
        "aHR0cDovL2NzaS5jb20vVGVhbXMvMw"
   ],
   "PlayerName": "Ted Sale",
   "DefensiveRating": 77.33
   }
]
```

### **Selecting Properties and Ordering Results**

The request below selects player names and their defensive ratings. The results are ordered by defensive rating in descending order so that the player with the highest defensive rating is listed first. The request also formats the results in text/csv.

```
dataondemand/LeagueGM/LeagueData/Players?$select=PlayerName,DefensiveRating&$orderby=De
fensiveRating desc&$format=text/csv
```

#### Result

```
PlayerName, DefensiveRating
James Smith, 98.33
Alex Granderson, 98.22
Matt Butler, 95.66
Tim Hooper, 93.43
Steve Jones, 93.28
Ted James, 92.55
Fred Wynn, 88.68
Jared Bonds, 86.34
Billy Roper, 83.44
Mike Magazine, 78.33
Ted Sale, 77.33
Chris Underwood, 66.22
```

### Expanding the Results to Include Related Entities

The request below uses the \$expand operator to retrieve data from the Players entity and include the related Positions data for each player. For this example, the request limits the number of results returned to 2 players by adding \$top=2:

dataondemand/LeagueGM/LeagueData/Players?\$expand=PlayerToPosition&\$top=2

### Result

```
{
 "@odata.context":
"https://10.100.0.10/dataondemand/LeagueGM/LeagueData/$metadata#Players",
  "value": [
   {
     "players key": "aHR0cDovL2NzaS5jb20vUGxheWVycy8x",
     "playertoposition key": [
          "aHR0cDovL2NzaS5jb20vUG9zaXRpb25zLzg"
     ],
     "PlayerId": 1,
     "playertoteam key": [
          "aHR0cDovL2NzaS5jb20vVGVhbXMvMQ"
     ],
     "PlayerName": "Steve Jones",
      "DefensiveRating": 93.28,
      "PlayerToPosition": [
          {
            "positions key": "aHR0cDovL2NzaS5jb20vUG9zaXRpb25zLzg",
            "PositionId": 8,
            "ShortName": "CF",
            "positiontoplayer key": [
                "aHR0cDovL2NzaS5jb20vUGxheWVycy8xMg",
                "aHR0cDovL2NzaS5jb20vUGxheWVycy8x"
            ],
            "Description": "Centerfield"
          }
     1
    },
    ł
     "players key": "aHR0cDovL2NzaS5jb20vUGxheWVycy8xMA",
      "playertoposition key": [
          "aHR0cDovL2NzaS5jb20vUG9zaXRpb25zLzI",
          "aHR0cDovL2NzaS5jb20vUG9zaXRpb25zLzM"
     ],
     "PlayerId": 10,
     "playertoteam key": [
          "aHR0cDovL2NzaS5jb20vVGVhbXMvNA"
     ],
     "PlayerName": "Ted James",
      "DefensiveRating": 92.55,
      "PlayerToPosition": [
```

```
{
            "positions_key": "aHR0cDovL2NzaS5jb20vUG9zaXRpb25zLzI",
            "PositionId": 2,
            "ShortName": "C",
            "positiontoplayer key": [
                "aHR0cDovL2NzaS5jb20vUGxheWVycy84",
                "aHR0cDovL2NzaS5jb20vUGxheWVycy8xMA"
            ],
            "Description": "Catcher"
          },
          {
            "positions key": "aHR0cDovL2NzaS5jb20vUG9zaXRpb25zLzM",
            "PositionId": 3,
            "ShortName": "1B",
            "positiontoplayer_key": [
                "aHR0cDovL2NzaS5jb20vUGxheWVycy83",
                "aHR0cDovL2NzaS5jb20vUGxheWVycy8xMA"
            ],
            "Description": "First Base"
          }
      ]
    }
  ]
}
```

# Access the SPARQL Endpoint

Anzo offers a standard HTTP(S) SPARQL endpoint for sending SPARQL requests between client applications and Anzo. The endpoint is enabled by default. This topic provides the base endpoint URL and describes the supported HTTP methods and parameters.

### Tip

You can generate a cURL request against the SPARQL endpoint from a query in the Query Builder. Once you have a valid query, click the **More** button under the query and select **Copy CURL Command**. For more information about writing queries in the Query Builder, see Running SPARQL Queries in the Query Builder.

# Authentication

The Anzo SPARQL endpoint supports Basic Authentication. The endpoint can be configured to enable other Anzo-supported authentication methods. However, implementing alternate authentication mechanisms can have unexpected results. For more information, contact Cambridge Semantics Support.

### Note

Ultimately the data that is available to users from SPARQL endpoints depends on the access control configuration of the graphmart or linked data set as configured in Anzo.

# **HTTP Methods and Options**

The Anzo SPARQL endpoint accepts HTTP GET and POST methods. GET is used to retrieve data from the endpoint, and POST is used to send data to the endpoint. Update queries must use the POST method, and read queries can be submitted using GET or POST.

## **Endpoint Base URL**

Use the following base URL to access data in Anzo via the SPARQL endpoint. The table below describes each base URL component:

Option	Description
protocol	The protocol to use for the connection: <b>http</b> for HTTP protocol or <b>https</b> for SSL protocol.
hostname	The DNS name or IP address of the Anzo server.
port	The port for the endpoint. The port that you specify depends on the protocol that you choose. By default, the HTTP port is <b>80</b> and the HTTPS port is <b>443</b> . To view the ports that are configured for your Anzo instance, see <b>Server Settings</b> in the <b>Administration</b> menu.
sparql	Required keyword for the SPARQL endpoint.
store_type	The type of RDF store for the data. Typically users specify <b>graphmart</b> to query data that is in a graphmart. It is also possible to query the metadata for a linked data set (LDS) in the Dataset catalog. To query an LDS that is stored in a local volume, specify <b>Ids</b> as the store type.
url- encoded_ dataset_uri	The URI for the graphmart or the catalog entry for the LDS. The URI must be URL- encoded using upper case hexadecimal digits. Lower case hexadecimal digits are not supported at this time. How do I find the URI for a graphmart? How do I find the catalog entry URI for a dataset?

#### For example, the following base endpoint URL targets the data in a graphmart:

https://10.100.10.20:8443/sparql/graphmart/http%3A%2F%2Fcambridgesemantics.com%2FGraphm art%2F1ad0ee911b834097ad7f71ee0ae1c0ff

The example below shows a base endpoint URL that targets a dataset catalog entry:

https://10.100.10.20:8443/sparql/lds/http%3A%2F%2Fopenanzo.org%2FcatEntry
(%255Bhttp%253A%252F%252Fcsi.com%252FFileBasedLinkedDataSet%252F001e517db4f0eaea9f27942
7e4e2a828%255D%2540%255Bhttp%253A%252F%252Fopenanzo.org%252Fdatasource%252FsystemDataso
urce%255D)

# **HTTP Header Options**

The HTTP header provides information related to the transfer of data between the requesting client and the SPARQL endpoint. The table below describes the supported HTTP header options. Both of the fields are optional.

Option	Description
Content- Type	The Content-Type specifies the type of request that is being sent by the client. Anzo supports the following Content-Type values:
	<ul> <li>application/x-www-form-urlencoded: Including this value specifies that the query string will be passed as the value of a "query" or "update" HTTP parameter. This is the default value. When Content-Type is not specified, the endpoint behaves as if Content-Type: application/x-www-form-urlencoded is specified.</li> <li>application/sparql-query: Including this value specifies that the HTTP request body includes a SPARQL read (non-update) query.</li> <li>application/sparql-update: Including this value specifies that the HTTP request body includes a SPARQL update query.</li> </ul>
Accept	The Accept field specifies the response formats that are acceptable for the server to send back to the client. You can use this field to specify the output serialization format for query results in place of the format HTTP parameter. For details about the supported formats, see Format Options below.

# **HTTP Body Parameters**

The HTTP parameters in the body of the request provide the rest of the information about the request. Certain parameters are appropriate for read-only queries, SELECT and CONSTRUCT, and others are appropriate for updates, INSERT and DELETE. The tables below describe the supported parameters for query and update requests.

### **Query Parameters**

Parameter	Description
query	Specifies the full read-only query string to run. If you do not specify a url-encoded_dataset_uri, default-graph-uri or named-graph-uri in the request, the query string should contain the appropriate FROM clauses. To run an update query (INSERT or DELETE), use the update parameter.
default-graph-uri	Specifies a default graph URI to query. You can include this parameter multiple times in a request. When the base URL specifies a graphmart URI, you can specify a data layer URI to narrow the scope of the query to a specific data layer in the graphmart.
named-graph-uri	Specifies a named graph URI to query. You can include this parameter multiple times in a request. When the base URL specifies a graphmart URI, you can specify a data layer URI to narrow the scope of the query to a specific data layer in the graphmart.
format	Specifies the serialization format to use for the results of the query. For details about the supported formats, see Format Options below.

Parameter	Description
includeMetadataGraphs	A boolean value that specifies whether to query the metadata graphs. Only valid for queries that target a linked data set (LDS) that is stored in a local volume. The default value is <b>includeMetadataGraphs=false</b> .
delim	Specifies a custom delimiter character to use in CSV output results. Valid only for SELECT queries where the output format is <b>text/csv</b> . This field accepts any character. When delim is not specified the default value is a , (comma).
dedup	A boolean value that specifies whether to deduplicate CONSTRUCT results on the client side. When dedup is not specified, the default value is <b>dedup=true</b> .
serverDedup	A boolean value that specifies whether to deduplicate CONSTRUCT results on the server side. When serverDedup is not specified, the default value is <b>serverDedup=true</b> .
skipCache	A boolean value that specifies whether to skip the reuse of any query cache that exists from a previous run of the query. When skipCache is not specified, the default value is <b>skipCache=false</b> . Anzo server's query cache should be forcibly skipped/ignored
hasHeader	A boolean value that specifies whether to include headers in CSV results. Valid only for SELECT queries where the output format is <b>text/csv</b> . When hasHeader is not specified, the default value is <b>hasHeader=false</b> .
attachResult	A boolean value that specifies whether to provide the query response as a file "attachment," i.e. the HTTP response will include the Content-Disposition of <b>attachment</b> . When attachResult is not specified, the default value is <b>attachResult=false</b> . When returning

Parameter	Description
	results as an attachment, you can specify a file name in filename the parameter.
filename	If attachResult is true, this parameter specifies the file name to use for the attachment, excluding the file extension. If attachResult is true and filename is not specified, the default file name is <b>QueryResult</b> .

### **Format Options**

The table below describes the options for specifying the serialization format of the results that the server sends back to the client. These format options, i.e., MIME types or file extensions, can be specified in the format parameter in the body of the request or in the Accept header.

### Note

When the request does not include the format parameter or Accept header, the default result format for SELECT queries is <u>SPARQL XML</u> (application/sparql-results+xml). For CONSTRUCT queries, the default format depends on whether the query includes GRAPH clauses. If no GRAPH clause is present, the default format for CONSTRUCT results is <u>RDF</u> <u>Turtle</u>. If GRAPH clauses are present, the default format is <u>RDF TriG</u>.

Format	Accepted Values	Query Type	Description
XML	application/sparql- results+xml application/xml xml xml2 srx	SELECT only	Returns results in <u>SPARQL</u> <u>Query Results XML Format</u> .
	application/rdf+xml rdf	CONSTRUCT only	Returns results in <u>RDF 1.1 XML</u> format.

Format	Accepted Values	Query Type	Description
	owl rdfs		
JSON	application/json json	SELECT and CONSTRUCT	For SELECT queries, results are returned in <u>SPARQL Query</u> <u>Results JSON Format</u> . For CONSTRUCT queries, results are returned in Anzo's native JSON RDF serialization format. See Anzo JSON RDF Serialization for details.
	application/sparql- results+json	SELECT only	Returns results in <u>SPARQL</u> Query Results JSON Format.
CSV	text/csv csv	SELECT only	Returns results in <u>SPARQL</u> Query Results CSV Format.
TriG and Gzipped TriG	application/x-trig trig application/x-trigz trigz gz trig.gz	CONSTRUCT only	CONSTRUCT queries with a GRAPH clause return <u>RDF 1.1</u> <u>TriG</u> by default if no format is specified.
Turtle and Gzipped Turtle	application/x-turtle ttl application/x-turtlez ttlz ttl.gz	CONSTRUCT only	Returns <u>RDF 1.1 Turtle</u> . CONSTRUCT queries without a GRAPH clause return Turtle by default if no format is specified.

Format	Accepted Values	Query Type	Description
N-Triples	text/plain nt	CONSTRUCT only	Returns results in <u>RDF 1.1 N-</u> <u>Triples</u> format.
Notation3 and Gzipped Notation3	text/rdf+n3 n3 text/rdf+n3z n3z n3z.gz	CONSTRUCT only	Returns results in <u>RDF Notation3</u> format.
N-Quads	text/x-nquads nq nquad nquads	CONSTRUCT only	Returns results in <u>RDF 1.1 N-</u> <u>Quads</u> format.
TriX	application/trix trix	CONSTRUCT only	Returns results in <u>RDF Triples in</u> <u>XML</u> format.

# **Update Parameters**

Parameter	Description
update	Specifies the full update string to run. If you do not specify a url- encoded_dataset_uri, using-graph-uri or using-named-graph-uri in the request, the update query should contain the appropriate USING clauses.
	To run a non-update query (SELECT or CONSTRUCT), use the query parameter.
using-graph-uri	Specifies a default graph URI to update. You can include this parameter multiple times in a request. When the base URL specifies a graphmart URI, you can specify a data layer URI to narrow the scope of the update to a specific data layer in the

Parameter	Description
	graphmart.
using-named-graph-uri	Specifies a named graph URI to update. You can include this parameter multiple times in a request. When the base URL specifies a graphmart URI, you can specify a data layer URI to narrow the scope of the update to a specific data layer in the graphmart.
includeMetadataGraphs	A boolean value that specifies whether to query the metadata graphs. Only valid for queries that target a linked data set (LDS) that is stored in a local volume. The default value is <b>false</b> .

# **Examples**

The following example uses cURL to send a request that runs a SELECT query against a graphmart. Since the request does not include an Accept header or format parameter, results will be returned in SPARQL XML format.

```
curl --user sysadmin:@nz0 -c cookiejar.txt -L -v -k
http://10.100.10.20/sparql/graphmart/http%3A%2F%2Fcambridgesemantics.com%2FGraphmart%2F
2dc579b101654ae29eb91b0c7d046ca1
--data-urlencode "query=SELECT * WHERE{ ?s ?p ?o . } LIMIT 100"
```

The following example sends a GET request that runs a SELECT query against a graphmart. The format parameter is included to format the results in text/csv serialization.

GET	https://10.102.0.17:443/sparql/graphmart/http%3/	A%2F%2Fcambridgesemantics.com%2FGraphmart%2Fbbf2d4b4c	138403bab1c671eb6d9763 Send V Save V		
Para	ns Authorization Headers (9) Body Pre-rec	quest Script Tests Settings	Cookies Code		
Quei	y Params				
	KEY	VALUE	DESCRIPTION ••• Bulk Edit		
$\checkmark$	format	text/csv			
$\checkmark$	hasHeaders	yes			
~	query	select * where {?s ?p ?o} limit 100			
	Key	Value	Description		
Body	Body Cookles (1) Headers (14) Test Results Status: 200 OK Time: 511ms Size: 18.06 KB Save Response 🔻				
Pre	Pretty Raw Preview Visualize <sup>BETA</sup> Text 🖛 🚍				
<pre>1 "s","p","o" 2 "http://csi.com/Shippers/1","http://www.w3.org/1999/02/22-rdf-syntax-ns#type","http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Shippers" 3 "http://csi.com/Territories/10406","http://www.w3.org/1999/02/22-rdf-syntax-ns#type","http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Territories" 4 "http://csi.com/Territories/10406","http://www.w3.org/1999/02/22-rdf-syntax-ns#type","http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Territories" 5 "http://csi.com/Territories/10406","http://www.w3.org/1999/02/22-rdf-syntax-ns#type","http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Territories" 6 "http://csi.com/Territories/04082","http://www.w3.org/1999/02/22-rdf-syntax-ns#type","http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Territories" 7 "http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Suppliers_Region","http://www.w3.org/1999/02/22-rdf-syntax-ns#type","http://www.w3.org/1099/02/22-rdf-syntax-ns#type","http://www.w3.org/1099/02/22-rdf-syntax-ns#type","http://www.w3.org/2002/07/oul#Datatype 8 "http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Suppliers_Region","http://www.w3.org/1999/02/22-rdf-syntax-ns#type","http://www.w3.org/2002/07/oul#Datatype 9 "http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Customers_Fax","http://www.w3.org/1999/02/22-rdf-syntax-ns#type","http://www.w3.org/2002/07/oul#Datatype 10 "http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Customers_Fax","http://www.w3.org/1999/02/22-rdf-syntax-ns#type","http://www.w3.org/2002/07/oul#Datatype 11 "http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Suppliers_HomePage","http://www.w3.org/1999/02/22-rdf-syntax-ns#type","http://www.w3.org/2002/07/oul#Datat 12 "http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Suppliers_HomePage","http://www.w3.org/1999/02/22-rdf-syntax-ns#type","http://www.w3.org/2002/07/oul#Datat 12 "http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Suppliers_HomePage","http://www.w3.org/1999/02/22-rdf-synta</pre>					

For reference, below is the URL-encoded version of the request string shown in the image above.

When sending a request from a client that does not automatically encode requests, you must

convert the string. Line breaks are added for readability:

```
http://10.100.10.20/sparql/graphmart/http%3A%2F%2Fcambridgesemantics.com%2F
Graphmart%2F646861d1bab54d67bc79dea94e02f3e6
?query=select%20*%20where%20%7B%3Fs%20%3Fp%20%3Fo%7D%20limit%20100
```

The example below sends a POST request that runs a SELECT query. In this example, the query is included in the body of the request and the response format is XML.

POST 🔻 htt	ps://10.102.0.17:443/sparql/graphmart	/http%3A%2F%2Fcambridgesemantics.com%2FGraphmart%	2Fbbf2d4b4c138403 Send	▼ Save ▼	
Params  Authoriza	tion  Headers (11) Body	Pre-request Script Tests Settings		Cookies Code	
🔵 none 🔵 form-dat	none of form-data of x-www-form-urlencoded raw binary GraphQL BETA				
KEY		VALUE	DESCRIPTION	••• Bulk Edit	
v query		select * where {?s ?p ?o} limit 20			
🔽 format		xml			
Кеу		Value	Description		
Body Cookies (1) Headers (14) Test Results Status: 200 OK Time: 153ms Size: 34.87 KB Save Response 🔻					
Pretty Raw Preview Visualize BETA XML -					
<pre>8 <result> 9</result></pre>					

The example below sends a GET request that runs a CONSTRUCT query. The response format is set to JSON, and the results are formatted in Anzo JSON RDF Serialization.

GEI	<ul> <li>https://10.102.0.17:443/sparql/graphmart/http%</li> </ul>	3A%2F%2Fcambridgesemantics.com%2FGraphmart%2Fbbf2d4b4c	138403bab1c671eb6d9763 Send	▼ Save ▼
Para	as ● Authorization ● Headers (9) Body Pre-re	equest Script Tests Settings		Cookies Code
Quei	/ Params			
	KEY	VALUE	DESCRIPTION	••• Bulk Edit
$\checkmark$	query	construct { graph <http: csi.com="" test=""> {?s ?p ?o}} where {?s ?p</http:>		
$\checkmark$	format	json		
	Key	Value	Description	
3ody Pre	Cookies (1) Headers (14) Test Results	् न	Status; 200 OK Time: 536ms Size: 6.89 KB	Save Response 🔻
<pre>1 2 4 3 3     "namedGraphUri": "http://csi.com/test", 3 4     "subject": { 5         "objectType": "uri", 6         "value": "http://cambridgesemantics.com/ont/autogen/Fu/DB/northwind#Customers" 7         }, 8         "predicate": "http://www.w3.org/1999/02/22-rdf-syntax-ns#type", 9         "object": { 10             "objectType": "uri", 11             "value": "http://www.w3.org/2000/01/rdf-schema#Resource" 12         } 13      }, 14      { </pre>				

The example below uses a Python script to send a request that runs a SPARQL query.

```
import requests
import urllib
server = 'https://company.anzo.com:'
port = 443
graphmart = 'http://cambridgesemantics.com/Graphmart/be4bd080c5654628b6fff90ca1b647d6'
url = server + str(port) + '/sparql/graphmart/' + urllib.quote_plus(graphmart)
#urllib.parse.quote_plus(graphmart) in Python 3
queryText = 'SELECT * WHERE {?instance a ?type .} LIMIT 10'
payload = {'query':queryText, 'format':'text/csv'}
r = requests.post(url, data = payload, auth = ('sysadmin','<pw>'))
print r.text
```

# Access the HTTP Client Interface

In addition to the SPARQL HTTP(S) endpoint that enables users to send SPARQL queries to Anzo over HTTP, Anzo provides an HTTP(S) servlet that can be used to invoke Anzo client operations over HTTP. The client servlet enables external systems to interact with Anzo semantic services as well as custom services. It also enables remote servers to interact with Anzo without needing the Anzo command line interface.

# **HTTP Methods and Options**

The Anzo client servlet accepts HTTP GET and POST methods. GET is used for operations that retrieve data, and POST is used for update operations that add or remove data. Update operations must use the POST method, and read operations can be submitted using GET or POST.

# **Client Servlet Base URL**

Use the following URL to access Anzo services via the HTTP client servlet. The table below describes each URL component:

<protocol>://<hostname>:<port>/anzoclient/<client\_operation>

### For example:

```
https://10.100.10.20:8443/anzoclient/call
```

Option	Description
protocol	The protocol to use for the connection: <b>http</b> for HTTP protocol or <b>https</b> for SSL protocol.
hostname	The DNS name or IP address of the Anzo server.
port	The port for the endpoint. The port that you specify depends on the protocol that you choose. By default, the HTTP port is <b>80</b> and the HTTPS port is <b>443</b> . To view the ports that are configured for your Anzo instance, see <b>Server Settings</b> in the

Option	Description
	Administration menu.
anzoclient	Required keyword for the client servlet.
client_ operation	The type of Anzo client operation to invoke. The list below provides an overview of the supported operation types. For more information about the operations, see Client Operations below.
	<ul> <li>call: Invokes the semantic service operation identified by the URI provided in the request (analogous to the anzo call CLI command)</li> </ul>
	<ul> <li>add: Imports the specified statements to Anzo (analogous to the anzo import CLI command)</li> </ul>
	<ul> <li>remove: Removes the specified statements from Anzo (analogous to the anzo update -r CLI command)</li> </ul>
	<ul> <li>get: Gets the specified named graph from Anzo (analogous to the anzo get CLI command)</li> </ul>
	• find: Finds the statements in Anzo that match the specified pattern (analogous to the anzo find CLI command)

# **Client Operations**

This section provides usage information and examples for each of the Anzo client operations.

- Call
- Add
- Remove
- Get
- Find

# Call

The call operation invokes a semantic service. Identify the service to call by providing the URI for the service in the request header. The call operation is supported with HTTP GET and POST methods. When including RDF data as input to the service, the request must use the POST method.

### **Call Header Options**

Call operations support the following header parameters. Only the uri parameter is required:

- uri: Required parameter that specifies the URI of the semantic service to invoke.
- mimeType, Content-Type, or format: Include one of these optional parameters to specify the MIME type for the RDF serialization used in the request body as well as the response from the service. The default type is application/json if the header does not specify the format. For more information about the supported RDF serialization types, see Format Options.

### **Call Body Options**

If the call operation supplies data as input to the service, include the data in the request body. The data must be serialized as specified in the request header, or **application/json** if the header does not specify a serialization type.

### **Call Examples**

The following cURL example uses a GET call to invoke a health check service.

```
curl https://10.100.10.20:8443/anzoclient/call \
    --user sysadmin:123 \
    --header 'uri: http://www.csi.com/service/genericIngestManager#healthCheck'
```

The example below uses a POST call to invoke a service operation. The call passes in a request data set that is serialized as RDF JSON.

```
"object" : {"objectType": "uri" ,"value" : "urn://object"},
"namedGraphUri" : "urn://ng"}'
```

The example below uses a POST call to invoke a service operation. The call passes in a request data set that is serialized as TriG.

```
curl https://10.100.10.20:8443/anzoclient/call \
    --header 'Content-Type: application/x-trig' \
    --user sysadmin:123 \
    --header 'uri: http://www.csi.com/service/genericIngestManager#healthCheck'
    --data '<urn://ng> { <urn://test> <urn://predicate> <urn://object> .}'
```

# Add

The add operation adds statements to the Anzo graphstore. Add is supported with the HTTP POST method. The header can include **mimeType**, **Content-Type**, or **format** to specify the MIME type for the RDF serialization used in the request body as well as the response from the service. The default type is **application/json** if the header does not specify the format. For more information about the supported RDF serialization types, see Format Options. The request body includes the statements to add.

### Add Examples

The following example add operation uses cURL to issue a POST call to add a statement to the graphstore. The statement is specified in Anzo JSON RDF serialization format.

```
curl https://10.100.10.20:8443/anzoclient/add \
   --user sysadmin:123 \
   --data '{"subject" : {"objectType": "uri" ,"value" : "urn://test"},
        "predicate" : "urn://predicate",
        "object" : {"objectType": "uri" ,"value" : "urn://object"},
        "namedGraphUri" : "urn://ng"}'
```

## Remove

The remove operation deletes statements from the Anzo graphstore. Remove is supported with the HTTP POST method. The header can include **mimeType**, **Content-Type**, or **format** to specify the MIME type for the RDF serialization used in the request body as well as the response from the

service. The default type is **application/json** if the header does not specify the format. For more information about the supported RDF serialization types, see Format Options. The request body specifies the statements to remove.

### **Remove Examples**

The following example remove operation uses cURL to issue a POST call to remove a statement from the graphstore. The statement is specified in Anzo JSON RDF serialization format.

```
curl https://10.100.10.20:8443/anzoclient/remove \
   --user sysadmin:123 \
   --data '{"subject" : {"objectType": "uri" ,"value" : "urn://test"},
        "predicate" : "urn://predicate",
        "object" : {"objectType": "uri" ,"value" : "urn://object"},
        "namedGraphUri" : "urn://ng"}'
```

# Get

The get operation retrieves a named graph from the Anzo graphstore. The get operation is supported with HTTP GET and POST methods. The named graph URI that contains the contents to retrieve can be included as a query parameter or as a **uri** parameter in the request body. The get operation also returns the metadata graph, which is equivalent to running anzo get -m <named\_graph\_uri> with the Anzo admin CLI. The header can include **mimeType**, **Content-Type**, or **format** to specify the MIME type for the RDF serialization used to format the response from the service. The default type is **application/json** if the header does not specify the format.

### **Get Examples**

The following example get operation uses cURL to retrieve the contents of a named graph.

```
curl -k -XPOST https://10.100.10.20:8443/anzoclient/get --user sysadmin:123
    --data-urlencode
    "uri=http://cambridgesemantics.com/Graphmart/9da211618a15476daa10cead2292d8e7"
```

### This example uses Python with requests:

```
import requests
url = "https://10.100.10.20:8443/anzoclient/get"
data = {"uri":
```

```
"http://cambridgesemantics.com/Graphmart/9da211618a15476daa10cead2292d8e7"}
username = "sysadmin"
password = "123"
r = requests.post(url, data=data, auth=(username, password), verify=False)
print (r.text
```

# Find

The find operation finds the statements in the graphstore that match the pattern that is specified in the request. The find operation is supported with HTTP GET and POST methods. Header options are not applicable. The list below describes each of the supported parameters. These parameters can be included as query parameters in the URL or as parameters in the request body:

- graph: The named graph URI for the find pattern.
- **sub**: The subject of the find pattern.
- pred: The predicate of the find pattern.
- lit: The object of the find pattern if that object is a literal value.
- uri: The object URI of the find pattern if that object is a URI.
- **type**: If the object is a literal, this parameter can be used to specify the data type of the literal value.
- **lang**: If the object is a literal, this parameter can be used to specify the language of the literal value.

Results returned by the find operation are in Anzo JSON RDF serialization format. See Anzo JSON RDF Serialization below for details.

### **Find Examples**

The following example find operation (using the GET HTTP method) finds all of the statements in the graphstore with predicate http://w3.org/1999/02/22-rdf-syntax-ns#type and an object URI of

http://cambridgesemantics.com/ontologies/2009/05/LinkedData#LinkedDataSet . The parameters are specified as query parameters in the URL.

```
curl https://10.100.10.20:8443/anzoclient/find?pred=http://www.w3.org/1999/02/22-rdf-
syntax-
ns%23type&uri=http://cambridgesemantics.com/ontologies/2009/05/LinkedData%23LinkedD
ataSet' \
    --user sysadmin:123
```

The example below finds the same statements but issues a POST call. The URL-encoded parameters are specified in the request body.

```
curl https://10.100.10.20:8443/anzoclient/find \
    --user sysadmin:123
    --data 'pred=http%3A%2F%2Fwww.w3.org%2F1999%2F02%2F22-rdf-syntax-
ns%23type&uri=http%3A%2F%2Fcambridgesemantics.com%2Fontologies%2F2009%2F05%2FLinkedData
%23LinkedDataSet'
```

# **Anzo JSON RDF Serialization**

Anzo's JSON RDF serialization standard is straightforward but differs from the common public JSON RDF serialization standards. In Anzo JSON serialization format, a set of statements (quads) are represented as an array of JSON objects. Each JSON object (statement) is defined as a key/value pair, where the key specifies the component of the statement, i.e., the subject, predicate, object, or namedGraphUri. Depending on the component, properties such as the component's value and data type are specified in nested objects.

The following example array shows Anzo's JSON serialization. The list below the example describes the structure.

```
[
    {
        "subject" : {
            "objectType": "uri",
            "value" : "urn://test"
        },
        "predicate" : "urn://predicate",
        "object" : {
            "objectType": "uri",
            "value" : "urn://object"
        },
        "namedGraphUri" : "urn://ng"
        },
```

```
{
   "subject" : {
      "objectType": "uri",
      "value" : "urn://test"
   },
   "predicate" : "urn://predicate2",
   "object" : {
      "objectType": "literal",
      "value" : "test literal",
      "dataType" : "http://www.w3.org/2001/XMLSchema#string"
   },
   "namedGraphUri" : "urn://ng"
}
]
```

- subject is a JSON object with two properties:
  - objectType: The resource type of the subject value. This is either a "uri" or "bnode" (blank node).
  - <sup>o</sup> value: The blank node value or a string literal that specifies the URI.
- predicate is a string literal that specifies the predicate URI.
- object is a JSON object with two required properties and two optional properties:
  - objectType: Required property that specifies whether the object is a "uri," "literal," or "bnode."
  - <sup>o</sup> value: Required property that specifies the string representation of the object value.
  - dataType: Optional property for use if the objectType is "literal." This property describes the data type of the literal value. It is a string literal of the XSD data type URI. For example: "http://www.w3.org/2001/XMLSchema#string"
  - language:Optional property for use if the objectType is "literal." This property describes the language of the literal value.
- namedGraphUri is a string literal that specifies the named graph URI.
## **Share Access to Artifacts**

All Anzo artifacts—data sources, schemas, models, graphmarts, etc.—that you create can be shared with other groups (or users) from the **Sharing** tab in the Anzo application. This topic provides an overview of the Sharing tab and basic instructions for configuring artifact permissions.

#### Note

For specifics about sharing multifaceted artifacts like graphmarts that include layers and dashboards that include multiple lenses, see Sharing Access to Graphmarts and Sharing Access to Dashboards and Lenses.

- Sharing Tab Overview
- Permission Settings
- Sharing an Artifact

## **Sharing Tab Overview**

Access the Sharing tab by navigating to an artifact and clicking **Sharing**. For example, the image below shows the Sharing tab for a data source.

Movies					Create Graphmart	
Overview		Tables	Versions	Discussion	Sharing	
nheritance						_
nherit permissions from			Pass per	missions to		
Inherit permissions from			✓ Movie	s-instance × Movies	×	×
Jser and Group Permissions 'ermissions defined for specific	users or (	groups are applie	ed in addition to any inherit	ed permissions.		
User and Group Permissions Permissions defined for specific Q Search users, roles or	users or ( groups	groups are applie	ed in addition to any inherit	ed permissions.		
Jser and Group Permissions Permissions defined for specific Search users, roles or System Administrator	users or g groups Admin	groups are applie	ed in addition to any inherit	ed permissions.		
Jser and Group Permissions Permissions defined for specific Search users, roles or System Administrator Everyone	users or g groups Admin View	groups are applie	ed in addition to any inherit	ed permissions.		
Jser and Group Permissions Permissions defined for specific Q Search users, roles or System Administrator Everyone	users or g groups Admin View	groups are applie	ed in addition to any inherit	ed permissions.	this artifact.	

## Inheritance

The top of the screen displays the permission inheritance settings:

>
${\bf x}$ Movies ${\bf x}$ $\qquad \times$ $\mid$ $\checkmark$

To facilitate common workflows, the Anzo application applies logic so that artifacts in the same workflow inherit the same permissions. You can alter the inheritance configuration by choosing the artifact or artifacts that this artifact should inherit from by choosing the artifacts to **Inherit permissions from**. Since the example above is a data source and no artifacts precede the source in Anzo, the Inherit Permissions From setting is empty.

You can also configure an artifact to **Pass permissions to** other artifacts. In the example above, the Movies data source passes permissions to the Movies schema instance and the Movies graphmart.

ss permissions to	)		
Movies-instance	×	Movies ×	×   ~

## Note

Permissions are additive. Copying permissions from multiple artifacts with differing permission levels results in the super set being acquired by the artifact that is inheriting the permissions. In addition, any permissions that are configured in the table at the bottom of the screen are also added to the set. For more conceptual information about permission inheritance, see <u>Permission Inheritance</u> in the Administration Guide.

### **User and Group Permissions**

In addition to the inheritance settings, the bottom of the screen lists the users and groups that this artifact has been shared with. For example:



The current level of access is listed next to each name: **View**, **Modify**, or **Admin**. View, Modify, and Admin are predefined permission sets. Each predefined set selects a certain combination of six permissions. You also have the option to create a **Custom** set of permissions. Selecting a user or group from the list displays the following permissions table on the right side of the screen:

Permissions	View	O Modify	O Admin	O Custom
Add/Edit		~	~	
View	~	~	~	
Delete		~	~	
Meta Add/Edit			~	
Meta View	~	~	~	
Meta Delete			~	

Typing a value in the **Search users, roles or groups** field finds and displays the users or groups that you can add to the list. The Permission Settings section below defines each of the permission sets.

## **Permission Settings**

The tables below list the predefined permission sets and describe the privileges that are granted for each permission that is part of the set:

## View

The following table describes the permissions in the View set.

Permission	Allows a user to:
View	<ul><li>See the artifact in the Anzo application.</li><li>Create versions of the artifact.</li></ul>
Meta View	Relates only to an artifact's permissions. A user with Meta View can see the permissions on the Sharing tab but they cannot modify, add, or remove permissions.

## Modify

In addition to the **View** and **Meta View** permissions described above, the **Modify** set includes the **Add/Edit** and **Delete** permissions described below.

Permission	Allows a user to:
Add/Edit	<ul> <li>Change an artifact, such as to rename it or edit its description.</li> <li>Add an entity to an artifact. For example, add a schema to a data source or a layer to a graphmart.</li> </ul>
Delete	<ul> <li>Remove an entity from the artifact. For example, delete a layer from a graphmart or a schema from a data source.</li> <li>Does not give permission to remove the parent artifact. For example, a user can remove a schema from a source but cannot delete the data source.</li> </ul>

## Admin

In addition to the View, Meta View, Add/Edit, and Delete permissions described above, the Admin set includes the Meta Add/Edit and Meta Delete permissions described below.

Permission	Allows a user to:
Meta Add/Edit	Relates only to an artifact's permissions. A user with Meta Add/Edit can add permissions to a user or group. They cannot remove permissions from any user or group.
Meta Delete	<ul><li>Remove permissions from a user or group.</li><li>Delete the parent artifact and its related entities.</li></ul>

## **Sharing an Artifact**

Follow the instructions below to share access to an artifact.

- 1. In the Anzo application, navigate to the artifact that you want to share access to. Then click the **Sharing** tab.
- 2. If you want to change the inheritance for the artifact, use the fields at the top of the screen:
  - To apply all of the permissions from another artifact to this one, select the artifact to inherit from in the **Inherit permissions from** field.
  - To pass this artifact's permissions to other artifacts, select the artifacts to pass permissions to in the **Pass permissions to** field.

#### Note

Permissions are additive. Copying permissions from multiple artifacts with differing permission levels results in the super set being acquired by the artifact that is inheriting the permissions. In addition, any permissions that are configured in the table at the bottom of the screen are also added to the set. For more conceptual information about permission inheritance, see Permission Inheritance in the Administration Guide.

3. To share access to this artifact with a particular user or group, type a value in the Search users, roles or groups field to find and display the user or group. The resulting list shows the current permission level that is set for each user or group in the search results. For example,

the image below shows the current permissions for the IT group (None):



4. Select the user or group for which you want to configure permissions. The permissions settings are displayed on the right side of the screen. For example:

User and Group Permissions	5					
Permissions defined for spe	cific users or groups a	ire applied in addition	to any innerned	permissions.		
○ info		Information Tech Set permission for Info	nology (IT) armation Technol	ogy (IT)		
2. Information Technol	ology (IT) None	Permissions	O View	O Modify	O Admin	O Custom
		Add/Edit		~	~	
		View	$\sim$	$\sim$	~	
		Delete		~	$\checkmark$	
		Meta Add/Edit			~	×
		Meta View	$\checkmark$	~	$\checkmark$	
		Meta Delete			$\checkmark$	×.

5. To assign a predefined set of permissions, click the View, Modify, or Admin radio button to assign that level of access to the selected user or group. Refer to Permission Settings above for details about the permissions sets. For example, the image below gives Admin permissions to users in the IT group:

User and Group Permissions Permissions defined for specific users or groups	are applied in addition	to any inherited	permissions.		
$\bigcirc$ info $\qquad \times$	Information Tech Set permission for Info	hnology (IT) ormation Technol	ogy (IT)		
Information Technology Admin     (IT)	Permissions	O View	O Modify	<ul> <li>Admin</li> </ul>	O Custom
	Add/Edit		~	~	
	View	~	~	~	
	Delete		~	~	
	Meta Add/Edit			~	
	Meta View	~	~	~	
	Meta Delete			~	

If you want to customize the permissions, click the **Custom** radio button and then select or deselect the permissions checkboxes. To clear permissions for a user or group, click the trashcan icon (**I**) next to the name.

Repeat the steps above to share the artifact with additional groups. Changes to permissions take effect immediately. Users do not need to log out of the application and log back in.

# **Version and Migrate Artifacts**

The topics in this section provide guidance on creating backup versions of artifacts and migrating artifacts by exporting them from one server and importing them to another server.

### Tip

For information about migrating artifacts and their related entities in bulk, see <u>Migration</u> <u>Packages</u> in the Administration Guide.

In this section:	
Creating and Restoring Versions of Artifacts	
Exporting an Artifact	
Making Properties Replaceable on Export	
Importing Exported Versions of Artifacts	

## **Creating and Restoring Versions of Artifacts**

Before making changes to data sources, schemas, unstructured pipelines, models, graphmarts, etc., users can take a snapshot of the current version of that artifact. When a backup is created, Anzo automatically creates a backup version of each entity that is related to that artifact. For example, backing up a graphmart backs up the same version of any related models, datasets, and dashboards. In addition, Anzo backs up the metadata graphs for all of the entities. Metadata graphs store information about the artifacts such as the creator and creation date and the permissions configuration. Changed artifacts can be reverted at any time to any of the saved versions. If an artifact is restored to a previous version, Anzo automatically saves a version of the current state of the artifact and its related entities and metadata. Follow the appropriate instructions to create or restore a backup version of an artifact.

- Create a Backup Version
- Restore a Backup Version

## **Create a Backup Version**

Follow the instructions below to save a snapshot of an artifact.

 In the Anzo application, navigate to the artifact that you want to back up, and then click the Versions tab. For example, the image below shows the Versions tab for a graphmart. In this example, the graphmart does not have any backup versions:

**		E ⊘ Ready to use	AZG Static			Save a Copy	Create Dashboard	D
	Overview	Explore	Datasets	Data Layers	Dashboards	Data on Demand	Versions	>
Crea	te Import							

#### Note

For data models, add the model that you want to back up to the Working Set and then open it in the Model editor. Then click the **Versions** tab.

2. Click the Create button. Anzo displays the Create New Version screen.

Create New Version		
Name for New Version *		
Comment for New Version		
	CANCEL	SAVE

- 3. In the **Name for New Version** field, type a name for the backup version. Then type details about the version in the optional **Comment for New Version** field.
- 4. Click **Save**. Anzo takes a snapshot of the artifact as well as its related entities and adds the version to the list on the left side of the screen. Depending on the size and number of related entities, the backup operation can take a few minutes to complete. For example:

*	Tickets Derived from: Tickets Backup ( INACTIVE	TIVE 🔗 Ready to us	ie 🛆 AZG 📑	Static		Save a Copy Create	Dashboard
	Overview	Explore	Datasets	Data Layers	Dashboards	Data on Demand	Versions >
Crea Tickets I Derived fro	te Import Backup corre	Tickets Backu Description None	q			Z Expor	t Restore
		Derived from Created Na Creator Sys	rv 10, 2022 stem Administrator		Version IDs Current Env : Cacva		
		Related Entitie Type	es (15)	Title	Version ID	Actions	
		Config	gurable lens	all	DEV 1		
		Config	gurable lens	doc	DEV 1		
		Config	gurable lens	events	DEV 1		
		Config	gurable lens	Events	DEV 1		
		Config	gurable lens	listings	DEV 1		
		Config	gurable lens	venues	DEV 1		

5. If necessary, select the new version in the list to view details on the right side of the screen. The screen displays details such as the version creator and created date and lists each of the related entities that were also backed up. In the list of related entities, the **Actions** column displays a compare icon next to each entity that has changed since the previous version.

### Example

In the image below, the compare icon in the Graphmart row indicates that this version of the graphmart includes changes that were not in the previous version:

Related Entities (16)	Title	Version ID	Actions
La. Anzo Data Store	Store	DEV 1	
CSV Data Sour	Flights	DEV 1	
Dataset Project	Load Flights	DEV 1	
File Backed Lin	Flights	DEV 1	
<>> File Based Dat	Flights	DEV 1	
File Connection	Server Filesystem	DEV 1	
File Connection	sysadmin User Folder	DEV 1	
💥 Graphmart	Flights Graphmart	DEV 2	Pe
📚 Layer	Flights	DEV 1	
S Layer	Queries	DEV 1	

Clicking the icon in the Actions column opens the Compare Versions dialog box, which shows a side-by-side comparison of the TriG files for the two versions:



Users can now make changes to the current version of the backed up artifacts, and the new changes can be reverted to a backup version at any time.

## **Restore a Backup Version**

Follow the instructions below to restore an artifact and its related entities to a previous backup version.

1. In the Anzo application, go to **Versions** tab for the artifact that you want to restore.

#### Note

For data models, add the model that you want to restore to the Working Set and then open it in the Model editor. Then click the **Versions** tab.

2. On the Versions screen, select the backup version that you want to restore. For example:

**	Tickets Derived from: Tickets Backup (1993) INACTIVE (1994) ACTIV	■ Ø E ⊘ Ready to use	AZG Static			Save a Copy Create D	ashboard
	Overview	Explore	Datasets	Data Layers	Dashboards	Data on Demand	Versions >
Crea Tickets Derived flo Nov 10, 30	te Import Backup CCCB In: 221 System Administrator	Tickets Backup Description None Derived from Orested New T0, 2 Orester System A Related Entities (1	022 dministrator		Version IDs Gurrent Env : (1971)	Z Export	Restore
		Туре	Tr	tle	Version ID	Actions	
		Configurabl	e lens all	I	DEV 1		
		Configurabl	e lens do	oe	DEV 1		
		Configurabl	ie lens ev	vents	DEV 1		
		Configurabl	ie lens Ev	rents	DEV 1		
		Configurab	e lens lis	stings	DEV 1		
		Configurab	e lens ve	nues	DEV 1		

3. Click the **Restore** button to restore the artifact to the version that you selected. Since Anzo automatically creates a snapshot of the current version before you restore an artifact, Anzo displays the Restore Version dialog box so that you can specify a label for the new version.

Restore Version: DEV 1 (Tickets Backup)						
Before restoring version "DEV 1 (Tickets Backup)", the current state will be saved as a new version.						
Name for New Version *						
Comment for New Version						
Restore metadata graphs						
Create version of current state before restoring						
CANCEL	SAVE					

- 4. In the Restore Version dialog box, type a name for the version in the **Name for New Version** field. You can also add an optional comment in the **Comment for New Version** field.
- 5. Specify whether you want to revert to the backed up version's metadata graphs for this artifact and its related entities:
  - If you want the restored version to use the metadata, such as permission configuration and last created date, that was saved at the time of the backup, select the **Restore metadata graphs** checkbox. Anzo will revert the metadata to the saved version.
  - If changes were made to the metadata for the current version of the artifact and you want to preserve those changes, such as if the permissions were modified to further restrict or allow access, leave the **Restore metadata graphs** checkbox blank. Anzo will preserve the current metadata graphs instead of reverting the metadata to the saved version.
- 6. Specify whether you would like Anzo to create a backup version of the current version of the artifact before restoring it:
  - If you want Anzo to make a backup copy of the current version before restoring it, leave the **Create version of current state before restoring** checkbox selected.
  - If you want do not want Anzo to create a backup copy of the current version before restoring it, clear the **Create version of current state before restoring** checkbox.
- 7. Click **Save**. Anzo saves the current version and restores the current files to the backup version. The new version is added to the list of available backups.

## **Exporting an Artifact**

You can export the current version of an artifact and its related entities or any backup version. Follow the instructions below to export an artifact. For instructions on creating a backup version, see Creating and Restoring Versions of Artifacts.

1. In the Anzo application, navigate to the artifact that you want to export.

#### Note

For data models, add the model that you want to export to the Working Set and then open it in the Model editor.

- 2. Follow one of the options below, depending on whether you want to export the current version of the artifact or a backup version:
  - If you want to export the current, working version of the artifact, click the Export icon (☑) under the artifact name.



### Tip

The image below shows an example of the Export icon for an artifact that has a backup version. Clicking the Export icon ( $\square$ ) exports the current version of the

artifact, not the backup version that is listed.

Derived from: Tickets Backup	
INACTIVE C ACTIVE OReady to use	AZG Static

 If you want to export a backup version of the artifact, click the Versions tab, which lists the backups that exist for the artifact. Select the version that you want to export and then click the Export button. For example, the image below shows the Versions tab for a

#### graphmart:

*	Movies Derived from: Transform Layer	avar ⊘ /E⊘ Ready to use)	AZG Static	Save a Cos	y Create Dashboard
	Overview	Explore Datas	sets Data	Layers Dashboar	ds Data on Demand >
Creat	te Import	Transform Layer			Export Restore
Transform Derived from Nov 10, 202	m Layer CSV2 m : new 22   System Administrator	Description None			
new Derived from	The second secon	Derived from new Created Nov 10, 2022 Creator System Administra	tor	Version IDs Current Env :	
		Related Entities (11)	Tiste	Version ID	Antina
		Configurable I	Movies	DEV 1	Actions
		Direct Load St	Load Movies	DEV 1	
		Export Step	Export Movies	DEV 1	
		<>> File Backed Li	Movies	DEV 1	
		File Based Dat	Movies	DEV 1	

The Export dialog box is displayed. For example:

Export Current State		
🧭 Include Related Entities 🐷 Include Registry Statements 👿 Include Metadata 👿 Include Dataset Editions And Compone	nts	
Advanced		~
	CANCEL	EXPORT

- 3. On the Export dialog box, configure the following export options as needed. The list below describes each option:
  - Include Related Entities: Indicates whether to export the artifact's related entities. Since most artifacts have dependencies with other artifacts, Cambridge Semantics recommends that you enable Include Related Entities (selected by default) and export all related entities. The number and type of related entities that are included varies by the type of artifact that is being exported.
  - Include Registry Statements: Indicates whether to export the registry statements for the artifact and each of its related entities.

- Include Metadata: Indicates whether to export the metadata graph for the artifact and its related entities, such as the permissions configuration and last modified date. If you exclude the metadata, the artifacts in this export will follow the metadata configuration on the destination server when they are imported. Select Include Metadata if you want to migrate the existing metadata to the destination server. Enabling this setting also gives you the option to change the permissions configuration for the exported entities.
- Include Dataset Editions and Components: If a dataset is included in the export package, this setting controls whether to include all of the editions and components with that dataset.
- 4. If you want to change permissions or replace the values for certain properties in the exported version of an entity, such as the user name and password for a data source, the base folder location for a File Store connection, or the file path for an Anzo Data Store, expand the **Advanced** option to view the Included Entities list. For example:

vanced			^
Included	Entities Sharing		
~	Туре	Title	
~	Configurable lens	Movies	
~	Direct Load Step	Load Movies	
~	Export Step	Export Movies	
~	C File Backed Linked Data Set	Movies	^
~	<> File Based Dataset	Movies	
~	File Connection	Server Shared Filesystem	^

The entities with replaceable values are expandable. Click the ^ character to the right of an entity name to expand the options and view the editable properties. Replace any of the existing values with the new values that you want to define for the exported version of the artifact.

## Tip

For information about configuring additional properties so that their values are replaceable on export, see Making Properties Replaceable on Export.

If you specified **Include Metadata** and want modify permissions for the exported entities, click the **Sharing** tab. For information about changing permissions on the Sharing tab, see **Share** Access to Artifacts.

## Tip

When a row in the Related Entities list includes the compare versions icon ( $\mathbb{F}^{\mathbb{F}}$ ) in the Actions column, you can click the icon to open a side-by-side comparison of the TriG files for the two versions. For example:

er: Added Airports	Ver: Original Backup	DEV 1
erved from Original Backup Jun 24, 2020   System Administrator	Derived from:	Jun 24, 2020   System Administrator
<pre>A control of the second control of the</pre>		Alber 77173 ver7231 ver3743 ve

5. Click **Export** to export the artifacts.

Anzo packages the files into a .zip file and downloads it to your computer. You do not need to extract the files in order to import the artifacts to another Anzo server. See Exported ZIP File Contents below for a description of the files that are included in the .zip file.

## **Exported ZIP File Contents**

Depending on the options configured for the export, the resulting .zip file contains one or more of the following files:

- <artifact\_name>\_export.trig contains statements about the type of artifact that was exported and the export settings that were configured.
- <artifact\_name>\_graph.trig contains the artifact definitions.
- <artifact\_name>\_metadata.trig contains metadata statements such as the access control configuration and last modified date for the exported entities.
- <artifact\_name>\_registry.trig contains registry statements such as the named graph information for the artifacts.

## **Making Properties Replaceable on Export**

When exporting artifacts, users can replace the existing values for certain properties like the user name and password for database data sources, the base folder location for file connections, and the file path for Anzo Data Stores. This topic provides instructions for configuring additional properties so that their values can be modified in the exported version of an artifact.

To configure a property so that its value is replaceable on export, add the following statement to the **http://cambridgesemantics.com/annotations/replaceStatements** graph:

```
<class_URI> http://cambridgesemantics.com/ontologies/2018/06/Export#replaceStatement <property_URI>
```

Where **<class\_URI>** is the URI for the class that defines the property whose value should be replaceable. And **<property\_URI>** is the URI of the property.

#### Important

The specified property must be a datatype property that contains a literal value.

You can use the following TriG contents as a template for defining properties with replaceable values. The contents show the default replaceable properties. You can add your statements to the ann:replaceStatements list and then import the file.

```
@prefix ds: <http://cambridgesemantics.com/ontologies/DataSources#> .
@prefix exp: <http://cambridgesemantics.com/ontologies/2018/06/Export#> .
@prefix ann: <http://cambridgesemantics.com/annotations/> .
#Mode:ADD
ann:replaceStatements {
   ds:PathConnection exp:replaceStatement ds:filePath .
   ds:FileConnection exp:replaceStatement ds:fileConnectionBaseFolder .
   ds:DbDataSource exp:replaceStatement ds:dbUser , ds:dbDatabase, ds:dbPassword .
}
```

## **Importing Exported Versions of Artifacts**

Follow the instructions below to import an exported version of an artifact (as described in Exporting an Artifact) and its related entities.

#### Note

To import a model that was created outside of Anzo or was downloaded from Anzo (as described in Downloading a Model), see Uploading a Model.

1. In the Anzo application, go to the resource selection screen for the type of artifact that you want to import. For example, the image below shows the Graphmarts screen:

Y Q Search	= ·	≡				Add Graphn	nart 👻
🗌 🕶 Title 🕥 Description	Status	Statements	Version	Creator	Tags	Actions	÷
😂 Books	$\otimes$			System Admi	Auto-Gen	Д	:
🕸 Flights	8			System Admi	Auto-Gen	Д	:
IS Movies	A Ready to us	104,166	Derived from: Transform Layer	System Admi	Auto-Gen	Д	÷
* Northwind	🛆 Ready to us	131,386		System Admi	Auto-Gen	Д	:
🕸 Supply Chain	8			System Admi		Д	:
S Tickets	A Ready to us	3,696,548	Derived from: Tickets Backup	System Admi	Auto-Gen	Д	:
				Rows per page: 2	25 <del>v</del> 1-6 o	f6 <	>

2. Click the **Add** ... button on the top of the screen and select **Import** .... Anzo opens the Import dialog box.

Import		
Version Location *		BROWSE
Please select the zip file to import		
☑ Include Related Entities ☑ Include Registry Statements 🗌 Include Me	etadata 🔽	Include Dataset Editions And Components
	CANCEL	PROCEED WITH ADVANCED PROCEED

3. On the Import screen, click the Version Location field to open the File Location dialog box.



- If the exported .zip file is on your computer, drag and drop the file onto the screen or click browse to navigate to the file and select it. If the .zip file is on a File Store, select the From File Store radio button and select the file on the File Store.
- 5. Click **OK** to save the file location value and close the File Location dialog box.
- 6. Enable or disable the following options as needed, depending on the data that the import file contains:
  - Include Related Entities: Indicates whether to import the artifact's related entities. Since most artifacts have dependencies with other artifacts, the Include Related Entities option is selected by default when an artifact is exported. Capturing all related entities on export ensures that all of an artifact's dependencies are included when that artifact is migrated. For example, an exported pipeline has the data source, schema, and model artifacts that it relies on when the pipeline is run. If the exported package includes related entities, Cambridge Semantics recommends that you enable Include Related Entities on import.
  - Include Registry Statements: Indicates whether to import the registry statements for the artifact and its related entities. This option is selected by default. Registry statements

should be included in imports. When registry statements are not included, the imported artifacts are not displayed in Anzo. For example, if a data source artifact is imported without registry statements, it would not be added to the Data Sources Registry and therefore not be displayed in the list of data sources in the Anzo application.

- Include Metadata: Indicates whether to import the metadata graph for the artifact and its related entities, such as the access control configuration and last modified date. If you select Include Metadata, you have the option to edit the permission configuration before importing the artifact.
- Include Dataset Editions And Components: This option specifies whether to import all of the editions and components for each dataset included in the import package.
- 7. Choose one of the following options to proceed with the import:
  - If you want to import the files as alternate versions of artifacts and not as the current, working version, and you do not want to replace any values or change permissions, click
     Proceed. Anzo imports the data and the imported files become available as versions on the relevant Version screens for the imported artifacts.
  - If you want to import these files as the current working version, and/or you want to change values or modify the permissions, click **Proceed With Advanced**. Anzo opens the Import Advanced Options dialog box. For example:

Import Ad	lvanced Options	
$\checkmark$	Configurable lens	doc
$\checkmark$	Configurable lens	events
$\checkmark$	Configurable lens	Events
$\checkmark$	Configurable lens	listings
$\checkmark$	Configurable lens	venues
$\checkmark$	Direct Load Step	Load Tickets
$\checkmark$	Export Step	Export Tickets
$\checkmark$	K File Backed Linked Data Set	Tickets
$\checkmark$	File Based Dataset	Tickets
$\checkmark$	File Connection	Server Shared Filesystem
$\checkmark$	😂 Graphmart	Tickets
		CANCEL IMPORT IMPORT& APPLY

Click the ^ character to the right of an entity name to expand the options and view the editable properties. Replace any of the existing values with the new values that you want to define for the imported version of the entity. If you specified **Include Metadata** and want modify permission settings for the import, click the **Sharing** tab and edit or add permissions for users and groups. For details about the Sharing tab, see <u>Share Access</u> to Artifacts.

When you are ready to import the entities, choose one of the following options:

- If you want to import the files as alternate versions and not as the current, working version, click Import. Anzo imports the files and the entities become available as versions on the relevant Version screens.
- If you want to import the files so that they become the current, working versions of the artifacts, click Import & Apply. Anzo creates a backup version of the existing working versions and then imports the artifacts as the new working versions.

# **SPARQL Best Practices and Query Templates**

To provide guidance on developing performant SPARQL queries and avoiding unexpected results, this section offers SPARQL best practices and query templates that you can use as a starting point for writing SPARQL queries in Anzo, such as in data layers, dashboard query lenses, and the Query Builder.

In this section:	
SPARQL Best Practices	
SPARQL Query Templates	

## **SPARQL Best Practices**

This topic provides some background information on query execution and gives general guidelines to follow to help ensure that your SPARQL queries are optimized and do not exhaust Anzo or AnzoGraph resources.

- Query Execution Overview
- Clause Execution Order and Optimization
- General Guidelines

### **Query Execution Overview**

When you run a new query against a graphmart, AnzoGraph parses the text and then plans the query execution strategy. The planner determines the steps that are required to compute the results in an optimal way. The plan is based on the statistics about your data that is gathered during graphmart activation. After the plan is determined, code for the plan is generated, compiled, and then distributed throughout the cluster to be executed. The plan XML file and compiled code are both saved to disk so that they can be reused when the query is run again. That is why the first run of a query takes longer than subsequent runs.

In addition to the caching that AnzoGraph performs, Anzo caches the query results. If you enable the **Skip server cache** option when running a query in the Query Builder, it means that Anzo will not access the cache that is available from a previous run of the query.

## **Clause Execution Order and Optimization**

The following image shows the execution order of clauses in a read (SELECT or CONSTRUCT) query. The list below describes the order.



- 1. First, any subselects or subqueries are performed. Note that the ordering presented in the image also applies to the clauses included in subqueries.
- 2. Then the graph patterns that are not included in a subquery are scanned.
- 3. Next, filtering is applied if a FILTER clause is included.
- 4. Once the WHERE clause has been processed, grouping is performed if a GROUP BY clause is included.
- 5. Next, the result clause is processed, including keywords like DISTINCT and any aggregation or other functions.
- Lastly, the results are ordered when ORDER BY is specified. The ORDER BY clause is more complex when GROUP BY and LIMIT are included: Nlog (M) where N is the number of grouped solutions and M is the LIMIT.

The GROUP BY and result clauses (steps 4 and 5) add linear complexity on top of the complexity in the WHERE clause. Query execution time increases proportionally to the total number of solutions found by the WHERE clause. To reduce complexity, trim the WHERE clause to have as few solutions as possible so that the expensive clauses are executed against as few solutions as possible.

## **General Guidelines**

This section offers guidance on developing performant queries and avoiding unexpected results.

- Limit the Results
- Replace FILTER with VALUES or Triple Patterns
- Avoid Cross-Product Joins
- Use Subqueries for Large Amounts of Data

### Limit the Results

The easiest way to reduce query execution time in some cases is to apply a LIMIT statement to limit the result set to a specific number of solutions. Limiting the number of results improves performance for cases where query results are calculated and returned in a streaming fashion. Limiting results is particularly useful when results need to be ordered so that the first group of results are the only ones of interest.

### Example: Find the sample ID and binding density for the top 10 most dense samples

```
PREFIX biom: <http://identifiers.csi.com/pharmakg/def/biomarker#>
SELECT ?sampleId ?bindingDensity
WHERE {
    ?sample a biom:Sample ;
        biom:sampleId ?sampleId ;
        biom:bindingDensity ?bindingDensity .
}
ORDER BY DESC(?bindingDensity)
LIMIT 10
```

#### **Replace FILTER with VALUES or Triple Patterns**

While a FILTER clause is useful for narrowing down selected data per a set of requirements, only use FILTER when the logic does not lend to other operations. In many cases, replacing FILTER with a VALUES clause or a well-organized set of triple patterns increases query performance. When

processing a FILTER statement, all non-filtered data must be retrieved before the FILTER can be applied. Using a VALUES clause or triple pattern, however, reduces the amount of data that is retrieved and processed after the retrieval.

#### Example 1: Inappropriate use of FILTER for a value-driven SELECT query

```
PREFIX uriRoot: <http://example.com/rootOfUris#>
SELECT ?var1 ?var2
WHERE {
    ?instanceOfClass a uriRoot:Class ;
        uriRoot:varName1 ?var1 ;
        uriRoot:varName2 ?var2 ;
        uriRoot:filteredVar ?filteredVar .
FILTER(?filteredVar = 'COMPAREDDATA1' || ?filteredVar = `COMPAREDDATA2' ||
?filteredVar = `COMPAREDDATA3')
    # filteredVar is first retrieved and then run through several comparisons
}
```

#### Solution 1: Use VALUES to select certain values

```
PREFIX uriRoot: <http://example.com/rootOfUris#>
SELECT ?var1 ?var2
WHERE {
    ?instanceOfClass a uriRoot:Class ;
        uriRoot:varName1 ?var1 ;
        uriRoot:varName2 ?var2 ;
        uriRoot:filteredVar ?valueVar .

VALUES (?valueVar) {
        ('COMPAREDDATA1')
        ('COMPAREDDATA2')
        ('COMPAREDDATA3')
    }
    # selection is performed once for each entry in the VALUES clause,
    # retrieving no more data than necessary
}
```

#### Example 2: Inappropriate use of FILTER for a value-driven SELECT query

```
PREFIX uriRoot: <http://example.com/rootOfUris#>
```

```
SELECT ?var1 ?filteredVar
WHERE {
    ?instanceOfClass a uriRoot:ClassName ;
        uriRoot:varName1 ?var1 ;
        uriRoot:varName2 ?var2 ;
        uriRoot:filteredVar ?filteredVar .
FILTER(?filteredVar = 'COMPAREDDATA1')
    # filteredVar is first retrieved and then compared
}
```

#### Solution 2: Use a triple literal to select certain values

```
PREFIX uriRoot: <http://example.com/rootOfUris#>
SELECT ?var1 ?filteredVar
WHERE {
    ?instanceOfClass a uriRoot:Class ;
        uriRoot:varName1 ?var1 ;
        uriRoot:filteredVar 'COMPAREDDATA' .
        # data is only retrieved if filteredVar matches desired compared data upon
initial retrieval
}
```

#### **Avoid Cross-Product Joins**

When trying to gather data from multiple classes at once, it is possible to accidentally create a cross-product join, a selection that combines the selected data in a hyper-linear way rather than simply assembling the data and returning an unprocessed set.

#### Example: Accidental cross-product query

```
PREFIX uriRoot: <http://example.com/rootOfUris#>
SELECT ?var1 ?var2
WHERE {
    ?instanceOfClass1 a uriRoot:Class1 ;
        uriRoot:varName1 ?var1 .
    ?instanceOfClass2 a uriRoot:Class2 ;
        uriRoot:varName2 ?var2 .
}
```

In the above example, the goal may have been to retrieve IDs from all instances of Class1 and all instances of Class2, for example, all of the participants and all of the subjects. However, the result of the query would be every combination of participant and subject. If there are 10 participants and 5 subjects, there would be 50 results rather than 15. In large data sets, this severely affects performance and puts the system under unnecessary strain.

There are two straightforward ways to separate or parameterize data to write a more performant query:

Solution 1: Use UNION to replace the cross-product

```
PREFIX uriRoot: <http://example.com/rootOfUris#>
SELECT ?commonVar
WHERE {
  {
    ?instanceOfClass1 a uriRoot:Class1 ;
       uriRoot:varName1 ?var1 .
    BIND(?var1 as ?commonVar)
  }
  UNION
{
    ?instanceOfClass2 a uriRoot:Class2 ;
       uriRoot:varName2 ?var2 .
    BIND (?var2 as ?commonVar)
  }
  # this creates a temporary graph that is a union of two graphs
  # in each of the two graphs, the desired data is saved under the same name
}
```

#### Solution 2: Use VALUES to replace the cross-product

```
PREFIX uriRoot: <http://example.com/rootOfUris#>
SELECT ?commonVar
WHERE {
    ?instanceOfClass a ?classURI ;
        ?propertyURI ?commonVar .
VALUES (?classURI ?propertyURI) {
        (uriRoot:Class1 uriRoot:varName1)
```

```
(uriRoot:Class2 uriRoot:varName2)
}
```

### Use Subqueries for Large Amounts of Data

When analyzing data, there may be a need to aggregate data and then perform a selection or derivation on the resulting aggregate. In this case, it is advisable to use one or more subselects or subqueries, where a SELECT query is included inside the WHERE clause and the remainder of the WHERE clause operates on the results of that SELECT as though the data were immediately available in the graph.

#### Example: Aggregate a variable and then process the aggregation

```
PREFIX uriRoot: <http://example.com/rootOfUris#>
SELECT ?var1 ?var2Aggregation
WHERE {
  {
    SELECT ?var1 (GROUP CONCAT(?var2) as ?var2Aggregation)
   WHERE {
        ?instanceOfClass1 a uriRoot:Class1 ;
            uriRoot:varName1 ?var1 .
       ?instanceOfClass2 a uriRoot:Class2 ;
            uriRoot:varName2 ?var2 .
    }
   GROUP BY ?var1
  }
  # var1 and var2Aggregation are now available for the usual processing
  # while var2 is no longer available as it only existed within the subselect
  FILTER(regex(?var2Aggregation, `DESIREDVAR2VAL'))
  # FILTER is used for illustrative purposes, but any processing would work
}
```

## **SPARQL Query Templates**

This topic provides templates that you can use as a starting point for writing SPARQL queries. The templates are based on the best practices described in SPARQL Best Practices.

- Basic Data Selection
- Graph Traversal Data Selection
- Text Cleanup with REGEX
- Data Aggregation
- Apply a Filter to Selected Data
- Create or Derive New Variables

### **Basic Data Selection**

The most fundamental use case for writing SPARQL queries is to select data from properties from a collection of instances. The following template and example query illustrate how to access a class in a model and return the properties on that class using their URIs.

#### **Abstracted Query Template**

Replace the bold text to modify the query.

```
PREFIX uriRoot: <http://example.com/rootOfUris#>
# select the variables that are populated in the WHERE clause
SELECT ?var1 ?var2
WHERE {
    ?instanceOfClass a uriRoot:ClassName ;
        uriRoot:varName1 ?var1 ;
        # use a prefix to abbreviate a property URI as shown above
        # or use the full URI as shown below
        <http://example.com/rootOfUris#varName2> ?var2 .
}
```

#### Example: Get the sample ID and anatomical location for each sample

PREFIX bm: <http://identifiers.csi.com/pharmakg/def/biomarker#>

## **Graph Traversal Data Selection**

The graph model enables the flexibility to combine data from different classes. The following template illustrates how to traverse between classes in the data model and access data from properties on multiple classes.

#### **Abstracted Query Template**

Replace the bold text to modify the query.

```
PREFIX uriRoot: <http://example.com/rootOfUris#>
# select the variables that are populated in the WHERE clause
SELECT ?var1 ?var2 ?varFromOtherClass
WHERE {
    ?instanceOfClass a uriRoot:ClassName ;
        uriRoot:varName1 ?var1 ;
        # use a prefix to abbreviate a property URI as shown above
        # or use the full URI as shown below
        <http://example.com/rootOfUris#varName2> ?var2 ;
        # getting data from other classes requires traversing per the model
        uriRoot:pointerToOtherClass ?instanceOfOtherClass .
    ?instanceOfOtherClass a uriRoot:OtherClass .
}
```

### **Text Cleanup with REGEX**

Once data is onboarded to Anzo, it is common to encounter strings that include issues such as unintended characters, missing spaces, and inconsistent formatting. You can use regular expressions in a data layer query to manipulate those values so that they are consistent and readable in analytics against the graphmart. The BIND clause in the query below trims any white space from before and after the string, converts the characters to upper case, and removes all non-alphanumeric characters and non-spaces. Replace the bold text as needed:

```
PREFIX : <http://csi.com/>
DELETE {
  GRAPH ${targetGraph}{
    ?s ?pred ?old val
 }
}
INSERT {
 GRAPH ${targetGraph}{
  ?s ?pred ?new val
 }
}
${usingSources}
WHERE {
  ?s a :Class ;
    ?pred ?old val .
 VALUES (?pred) {
    (:property)
  }
BIND (TRIM (UPPER (REPLACE (?val, "[^a-zA-Z0-9[[:space:]]", ""))) as ?new_val)
}
```

## **Data Aggregation**

Grouping data selections around a central property yields a more complete representation or summary of the data available. The following template illustrates how to use one property to act as a pivot point for collecting all the data from another property.

#### **Abstracted Query Template**

Replace the bold text to modify the query

```
# apply selection/filtering logic to narrow the aggregation
# or get summaries of total data by applying only simple restrictions
?instance a pref:Class ;
    pref:instanceId ?instanceId ;
    pref:instanceDetail ?instanceDetail .
}
GROUP BY ?instanceId
# all non-aggregated variables must be grouped in GROUP BY
```

## Apply a Filter to Selected Data

Filtering the results for a query gives the ability to focus on specific aspects of the data. The following template illustrates how to restrict the total selected result set by including a filter on a variable.

#### **Abstracted Query Template**

Replace the bold text to modify the query.

```
PREFIX pref1: <http://example.com/rootOfUris1#>
PREFIX pref2: <http://example.com/rootOfUris2#>
SELECT ?varFromClass1 ?varFromClass2 ?varFromClass3 ?filteredVar
WHERE {
  ?instance1 a pref1:Class1 ;
     pref1:varName1 ?varFromClass1 ;
      # the path on the model points from Class1 to Class2
      pref1:pointerToClass2 ?instance2 .
  ?instance2 a pref1:Class2 ;
      pref1:varName2 ?varFromClass2 .
  # models with different prefixes can still be joined
  ?instance3 a pref2:Class3 ;
      # the path on the model points from Class3 to Class2
      pref2:pointerToClass2 ?instance2 ;
      pref2:filteredVarName ?filteredVar .
  # filters use comparisons to scope the selected data
  # they can use existence checks or other boolean expressions as well
  FILTER(?filteredVar = 'COMPAREDDATA')
}
```

Tip

For optimal query performance, replace FILTER clauses. See Replace FILTER with VALUES or Triple Patterns for more information.

### **Create or Derive New Variables**

Storing intermediate or derived data within a query enables a single query to answer more complex questions. The following template illustrates how to bind a derived value to a variable. That variable is then available for selection or further manipulation.

#### **Abstracted Query Template**

Replace the bold text to modify the query.

```
PREFIX pref1: <http://example.com/rootOfUris1#>
PREFIX pref2: <http://example.com/rootOfUris2#>
PREFIX pref3: <http://example.com/rootOfUris3#>
SELECT ?var1 ?filterVar ?var2AndVar3
WHERE {
  ?instance1 a pref1:Class1 ;
     pref1:varName1 ?var1 .
  ?filterInstance a pref2:MedicalHistory ;
     pref2:filterVarName ?filterVar ;
      # multiple traversals between classes may be necessary to link appropriate data
      pref2:pointerToIntermediateClass ?intermediateInstance .
  ?intermediateInstance a pref2:IntermediateClass ;
      pref2:pointerToClass1 ?instance1 .
  ?instance2 a pref3:Class2 ;
      # forwards traversals tend to be more performant
      # it is still possible to identify a latter class and do a backwards traversal
      pref3:pointerToClass1 ?instance1 ;
      pref3:varName2 ?var2 .
  ?instance3 a pref3:Class3 ;
     pref3:pointerToClass2 ?instance2 ;
     pref3:varName3 ?var3 .
```
# filters can be executed on various data types
FILTER(?filterVar < "filterData"^^xsd:filterDataType)</pre>

# binding allows population of new/derived variables
BIND(CONCAT(?var2, "--", ?var3) as ?var2AndVar3)

}

# **Function and Formula Reference**

This section describes the standard and advanced built-in functions that are available when working with Hi-Res Analytics dashboards or writing SPARQL queries in data layers or the Query Builder.

In this section:	
String Functions	1011
Math Functions	1040
Aggregate Functions	1071
Date and Time Functions	1092
Casting Functions	1113
Logical Functions	1131
Informational or Testing Functions	1142
Hash Functions	1154
Window Aggregate and Ranking Functions	1158

# **String Functions**

This topic describes the Anzo functions that operate on string data types.

**Typographical Conventions** 

This documentation uses the following conventions in function syntax:

- CAPS: Although SPARQL is case-insensitive, function names and other keywords are written in uppercase for readability.
- [ argument ]: Brackets are used to indicate optional arguments. Arguments without brackets are required.

#### Functions

- BUSINESS\_ENTITY\_EXCLUDER: Removes suffixes that represent business entities.
- CONCATENATE: Concatenates two or more strings and returns the result as a string.
- CONCATURL: Concatenates two or more strings and returns the result as a URI.
- CONTAINS: Evaluates whether the specified string contains the given pattern.
- ENCODE\_FOR\_URI: Encodes the specified string as a URI.
- ESCAPEHTML: Escapes the specified string for use in HTML.
- FIND: Returns the position—from left to right—of a string within another string.
- FINDREVERSE: Returns the position—from right to left—of a string within another string.
- **GROUP\_CONCAT**: Concatenates a group of strings into a single string.
- GROUPCONCAT: Concatenates a group of strings into a single string. This function is a customizable version of GROUP\_CONCAT.
- LANG: Returns any language tags that are included with strings.
- LANGMATCHES: Evaluates whether a string includes a language tag that matches the specified language range.
- LCASE: Converts the letters in a string literal to lower case.

- LEFT: Returns the specified number of characters starting from the beginning (left side) of the string.
- LEN: Calculates the length (number of characters) in a string.
- LEVENSHTEIN\_DIST: Calculates the Levenshtein distance or measure of similarity between two strings.
- LOWER: Converts all letters in a string to lower case.
- MD5: Returns the MD5 checksum of a string as a hexadecimal string.
- MID: Returns the specified number of characters from a string, starting from a given position in the string.
- **REGEX**: Evaluates whether a string matches the specified regular expression pattern.
- **REGEXP\_SUBSTR**: Searches a string for the specified regular expression pattern and returns the substring that matches the pattern.
- REPLACE: Extends the REGEX function to provide the ability to find a pattern in a string and replace it with another pattern.
- RIGHT: Returns the specified number of characters starting from the end (right side) of the string.
- SEARCH: Uses text search semantics to evaluate whether the specified string matches the given pattern.
- SHA1: Calculates the SHA-1 digest of a string value.
- SHA224: Calculates the SHA-224 digest of a string value.
- SHA256: Calculates the SHA-256 digest of a string value.
- SHA384: Calculates the SHA-384 digest of a string value.
- SHA512: Calculates the SHA-512 digest of a string value.
- STRAFTER: Returns the portion of a string that comes after the specified substring.
- STRBEFORE: Returns the portion of a string that comes before the specified substring.

- STRDT: Constructs a literal value with the specified data type.
- STRENDS: Evaluates whether the specified string ends with the specified substring.
- STRLANG: Constructs a literal value with the specified language tag.
- STRLEN: Calculates the length of a string.
- STRSTARTS: Evaluates whether the specified string starts with the specified substring.
- STRUUID: Returns a string that is the result of generating a Universally Unique Identifier (UUID).
- SUBSTITUTE: Substitutes the existing text for the specified new text.
- SUBSTR: Returns a substring from a string value.
- TOURI: Casts a string to a URI.
- TRIM: Removes all spaces from a string except for any single spaces between words.
- UCASE: Converts all letters in a string to upper case.
- UPPER: Converts the letters in a string literal to upper case.

# BUSINESS\_ENTITY\_EXCLUDER

This function removes from strings the suffixes that represent business entities.

#### **Syntax**

BUSINESS_ENTITY_EXCLUDER(text)		
Argument	Туре	Description
text	string	The string from which you want to remove business entities.

Data Type	Description	
string	The string without the business entity suffix.	

# CONCATENATE

This function concatenates two or more strings and returns the result as a string.

## Syntax

CONCATENATE (te	ext1, text2	[, textN ])
Argument	Туре	Description
text1–N	string The strings that you want to concatenate to form a single string	
Returns		

Туре	Description	
string	The concatenated string.	

# CONCATURL

This function concatenates two or more strings and returns the result as a URI.

## Syntax

CONCATURL(text1, text2 [, textN ])		
Argument	Туре	Description
text1–N	string	The strings that you want to concatenate to form a URI.

Туре	Description	
URI	The concatenated string as a URI.	

# CONTAINS

This function evaluates whether the specified strings contain the given pattern. Results are grouped under "true" or "false."

# Syntax

CONTAINS(text, pattern)			
Argument	Туре	Description	
text	string	The string value that you want to check against the specified pattern.	
pattern	string	The string pattern that you want to look for in the supplied text.	

#### Returns

Туре	Description	
boolean	${\tt True}\ {\it if}\ {\it the}\ {\it strings}\ {\it contain}\ {\it the}\ {\it pattern}\ {\it and}\ {\tt false}\ {\it if}\ {\it they}\ {\it do}\ {\it not}.$	

# ENCODE\_FOR\_URI

This function encodes the specified string as a URI and returns a string in URI format.

## Syntax

ENCODE\_FOR\_URI(text)

Argument	Туре	Description
text	string	The string value to encode as a URI.

Туре	Description
string	The string as a URI.

# **ESCAPEHTML**

This function escapes the specified string for use in HTML.

## Syntax

ESCAPEHTML (text	<b>z</b> )	
Argument	Туре	Description
text	string	The string value to escape for HTML.

#### Returns

Туре	Description
string	The string escaped for HTML.

# FIND

This function returns the position—from left to right—of a string within another string.

# Тір

You can use FINDREVERSE to find the character or substring position from right to left.

```
FIND (find text, within text, start num)
```

Argument	Туре	Description
find_text	string	The string to look for in the within_text.
within_text	string	The string to search within.
start_num	int	An integer that indicates the position to start from when looking for the find_text. The starting position is at the beginning of the within_

Argument	Туре	Description
		text value and characters are counted from left to right.
Returns		

# Type Description int The character position (from left to right) where the substring starts.

# FINDREVERSE

Similar to FIND, this function returns the position—from right to left—of a string within another string.

#### **Syntax**

FINDREVERSE(find\_text, within\_text, start\_num)

Argument	Туре	Description
find_text	string	The string to look for in the within_text value.
within_text	string	The string to search within.
start_num	int	An integer that indicates the position to start from when looking for the find_text. The starting position is the end of the within_text value and characters are counted from right to left.

Туре	Description	
int	The character position (from right to left) where the substring starts.	

# GROUP\_CONCAT

This function concatenates a group of strings into a single string. It is a simplified version of GROUPCONCAT as it takes only one argument.

## Syntax

GROUP_CONCAT(text)		
Argument	Туре	Description
text	string	The string property whose values to concatenate into a single string.

#### Returns

Туре	Description
string	The concatenated string.

# GROUPCONCAT

This function concatenates a group of strings into a single string. Unlike GROUP\_CONCAT, this function allows for customization of the separator to use as well as the configuration of limits and options like prefixes and suffixes.

Argument	Туре	Description
group1–N	string	The group(s) of strings to concatenate.
group_value_ separator	string	The separator string to use between the groups of strings if you specified more than one group.

Argument	Туре	Description
separator	string	The separator string to use between the values in a concatenated group of strings.
serialize	boolean	A boolean value that indicates whether returned values should be serialized with the value's data type.
row_limit	int	An integer that puts a maximum limit on the number of rows to retrieve for a group.
value_limit	int	An integer that puts a maximum limit on the number of values to retrieve from a group of rows.
delimit_blanks	boolean	A boolean value that indicates whether to delimit blanks with the separator value.
prefix	string	Optional string to add as a prefix to the resulting string.
suffix	string	Optional string to add as a suffix to the resulting string.
max_length	int	Optional integer that puts a maximum limit on the number of characters the resulting string can have.

Туре	Description
string	The concatenated string.

# LANG

This function returns any language tags that are included in the string. The results are grouped by each language tag or by "blank" if a value does not have a language tag.

## Syntax

LANG(text)		
Argument	t Type	Description
text	string	The string to search for language tags.
Returns		
Type	Description	

51	
string	The found language tags.

# LANGMATCHES

This function tests whether a string includes a language tag that matches the specified language range.

## Syntax

```
LANGMATCHES(text, language_range)
```

Argument	Туре	Description
text	string	The string to evaluate.
language_range	string	The language tag to match in the text.

## Example

LANGMATCHES(LANG(?prop), "en")

Туре	Description
boolean	${\tt True}$ if strings include a language tag that matches the range and ${\tt false}$ if they do

Туре	Description
	not.

# LCASE

This function converts the letters in a string literal to lower case.

## Syntax

LCASE (text)

Argument	Туре	Description
text	string	The string literal to convert to lower case.

## Returns

Туре	Description	
string	The string with lower case letters.	

# LEFT

This function returns the specified number of characters starting from the beginning (left side) of the string.

## Syntax

LEFT(text, num\_chars)

Argument	Туре	Description
text	string	The string from which to return the specified number of characters.
num_chars	int	An integer that specifies the number of characters to return, starting from the left side of the $text$ .

Туре	Description
string	The specified number of characters from the string.

# LEN

This function calculates the length (number of characters) in a string.

#### Syntax

LEN(text)		
Argument	Туре	Description
text	string	The string for which to calculate the length.
Returns		
Type	Description	

Туре	Description
int	The number of characters in the string.

# LEVENSHTEIN\_DIST

This function calculates the Levenshtein distance or measure of similarity between two strings. The distance is the number of edits required to transform the first string into the second string.

LEVENSHTEIN_DIST(text1, text2)		
Argument	Туре	Description
text1	string	The string that would be transformed into text2.
text2	string	The string to measure text1 against.

Туре	Description
int	The Levenshtein distance between the strings.

# LOWER

This function converts all letters in a string to lower case.

## Syntax

LOWER(text)

Argument	Туре	Description
text	string	The string to convert to lower case.

#### Returns

Туре	Description
string	The string with lower case letters.

# MD5

This function returns the MD5 checksum of a string as a hexadecimal string.

MD5(text)		
Argument	Туре	Description
text	string	The string for which to return the MD5 checksum.

Туре	Description
string	The hexadecimal string.

# MID

This function returns the specified number of characters from a string, starting from a given position in the string.

#### Syntax

```
MID(text, start_num, num_chars)
```

Argument	Туре	Description
text	string	The string from which to return the specified characters.
start_num	int	An integer that indicates the starting position in the string.
num_chars	int	An integer that specifies the number of characters to return, starting with the character indicated by <pre>start_num</pre> .

## Returns

Туре	Description
string	The specified number of characters from the string.

# REGEX

This function tests whether a string matches the specified regular expression pattern.

```
REGEX(text, pattern [, flags ])
```

Argument	Туре	Description
text	string	The string to test against the pattern.
pattern	string	The regular expression pattern to look for in the text. For information about the supported regular expression syntax, see the <u>Regular</u> <u>Expression Syntax</u> section of the W3C XQuery 1.0 and XPath 2.0 Functions and Operators specification.
flags	string	You can include one or more optional modifier flags that further define the pattern. For information about flags, see the <u>Flags</u> section of the W3C Functions and Operators specification.

Туре	Description
boolean	${\tt True}$ if the string matches the regular expression pattern and ${\tt false}$ if it does not.

# **REGEXP\_SUBSTR**

This function searches a string for the specified regular expression pattern and returns the substring that matches the pattern.

#### Syntax

REGEXP\_SUBSTR(text, pattern [, start\_position ] [, nth\_appearance ])

Argument	Туре	Description
text	string	The string to test against the pattern.
pattern	string	The regular expression pattern to look for in the text. For information about the supported regular expression syntax, see the <u>Regular Expression Syntax</u> section of the W3C XQuery 1.0 and XPath 2.0 Functions and Operators specification.

Argument	Туре	Description
start_position	int	An optional integer that specifies the number of characters from the beginning of the string to start searching for matches (the default value is 1).
nth_ appearance	int	An optional integer that specifies which occurrence of the pattern to match (the default value is 1).

Туре	Description
string	The substring that matches the regular expression pattern.

# REPLACE

This function extends the REGEX function to provide the ability to find a pattern in a string and replace it with another pattern. The function returns the replaced string.

#### Syntax

REPLACE(text, pattern, replacement\_pattern [, flags ])

Argument	Туре	Description
text	string	The string to test against the pattern.
pattern	string	The regular expression pattern to look for in the text. For information about the supported regular expression syntax, see the <u>Regular Expression Syntax</u> section of the W3C XQuery 1.0 and XPath 2.0 Functions and Operators specification.
replacement_ pattern	string	The pattern to replace the pattern with.

Argument	Туре	Description
flags	string	You can include one or more optional modifier flags that further define the pattern. For information about flags, see the <u>Flags</u> section of the W3C Functions and Operators specification.

Туре	Description
string	The string that contains the replacement pattern.

# **RIGHT**

This function returns the specified number of characters starting from the end (right side) of the string.

## Syntax

RIGHT(text, num\_chars)

Argument	Туре	Description
text	string	The string from which to return the specified number of characters.
num_chars	int	An integer that specifies the number of characters to return, starting from the right side of the $text$ .

Туре	Description
string	The specified characters from the string.

# SEARCH

This function uses text search semantics to evaluate whether the specified string matches the given pattern.

## Syntax

SEARCH(text, pattern [, required ] [, wildcard ] [, escape ])

Argument	Туре	Description
text	string	The string to search.
pattern	string	The search string to look for in the $text$ . Anzo automatically converts the value to a regular expression pattern that uses text search semantics.
required	boolean	An optional boolean value that indicates whether the $text$ must include all elements of the search pattern to qualify as a match or whether matching just part of the pattern qualifies as a match.
wildcard	boolean	An optional boolean value that indicates whether or not to add the wildcard character $\star$ to the end of the search pattern.
escape	boolean	An optional boolean value that indicates whether or not escape all of the special characters (such as +, -, or $ $ ) in the text.

#### Returns

Туре	Description
boolean	True if strings match the pattern and false if they do not.

# SHA1

This function calculates the SHA-1 digest of a string.

## Syntax

SHA1(text)		
Argument	Туре	Description
text	string	The string for which to calculate the SHA-1 digest.

#### Returns

Туре	Description
string	The SHA-1 digest.

## **SHA224**

This function calculates the SHA-224 digest of a string.

# Syntax

SHA224 (text)

Argument	Туре	Description
text	string	The string for which to calculate the SHA-224 digest.

#### Returns

Туре	Description
string	The SHA-224 digest.

## **SHA256**

This function calculates the SHA-256 digest of a string.

## Syntax

SHA256(text)

Argument	Туре	Description
text	string	The string for which to calculate the SHA-256 digest.

Туре	Description
string	The SHA-256 digest.

# **SHA384**

This function calculates the SHA-384 digest of a string.

# Syntax

SHA384 (t	text)
-----------	-------

Argument	Туре	Description
text	string	The string for which to calculate the SHA-384 digest.

#### Returns

Туре	Description
string	The SHA-384 digest.

## **SHA512**

This function calculates the SHA-512 digest of a string.

## Syntax

SHA512(text)

Argument	Туре	Description
text	string	The string for which to calculate the SHA-512 digest.

Туре	Description
string	The SHA-512 digest.

# STRAFTER

This function returns the portion of a string that comes after the specified substring.

## Syntax

STRAFTER(text, substring)		
Argument	Туре	Description
text	string	The string from which to return the characters that follow the substring.
substring	string	The string to match in the $text$ . The function will return the part of the text that comes after this substring.

#### Returns

Туре	Description
string	The part of the string that comes after the substring.

# STRBEFORE

This function returns the portion of a string that comes before the specified substring.

# Syntax

STRAFTER(text, substring)

Argument	Туре	Description
text	string	The string from which to return the characters that precede the substring.
substring	string	The string to match in the $text$ . The function will return the part of the text that comes before this substring.

## Returns

Туре	Description
string	The part of the string that comes before the substring.

# STRDT

This function constructs a literal value with the specified data type.

STRDT(text, datatype)		
Argument	Туре	Description
text	string	The string to add a data type specification to.
datatype	URI	The data type URI to add to the text. For example, xsd:integer or <a href="http://www.w3.org/2001/XMLSchema#integer">http://www.w3.org/2001/XMLSchema#integer</a> >.

Туре	Description
string	The typed literal value.

# **STRENDS**

This function evaluates whether the specified string ends with the specified substring.

# Syntax

STRENDS(text, substring)		
Argument	Туре	Description
text	string	The string to search for the substring.
substring	string	The string to match at the end of $text$ . The function returns true if the text ends in the specified substring and false if it does not.

#### Returns

Туре	Description
boolean	True if strings end with the specified substring and false if they do not.

# STRLANG

This function constructs a literal value with the specified language tag.

#### Syntax

STRLANG(text, language\_tag)

Argument	Туре	Description
text	string	The string to add the language tag to.
language_tag	string	The language tag to add to the text.

Туре	Description
string	The literal value with the language tag.

# STRLEN

This function calculates the length (in characters) of a string value.

#### Syntax

STRLEN(text)

Argument	Туре	Description
text	string	The string for which to return the length.

#### Returns

Туре	Description
long	The number of characters in the string.

# **STRSTARTS**

This function evaluates whether the specified string starts with the specified substring.

## Syntax

STRENDS(text, substring)

Argument	Туре	Description
text	string	The string to search for the substring.
substring	string	The string to match at the beginning of $text$ . The function returns true if the text starts with the specified substring and false if it does not.

Туре	Description	
boolean	True if strings begin with the specified substring and false if they do not.	

# STRUUID

This function returns a string that is the result of generating a Universally Unique Identifier (UUID).

## Syntax

STRUUID()

#### Returns

Туре	Description
string	The UUID.

## SUBSTITUTE

This function substitutes the existing text for the specified new text.

#### Syntax

SUBSTITUTE(text, old\_text, new\_text [, instance\_num ])

Argument	Туре	Description
text	string	The string to substitute text in.

Argument	Туре	Description
old_text	string	The string within the text to replace.
new_text	string	The string to replace the old_text with.
instance_ num	int	An optional integer that specifies the number of $old\_text$ instances to replace.

Туре	Description
string	The string with the new text.

# SUBSTR

This function returns a substring from a string value.

<pre>SUBSTR(text, start [, length ])</pre>		
Argument	Туре	Description
text	string	The string to find the substring in.
start	int	An integer that specifies the number of the character in the $text$ that should be the start of the substring.
length	int	An optional integer that specifies the total number of characters to include in the substring. If not specified, the substring will end at the end of the $text$ value.

Туре	Description
string	The substring.

# TOURI

This function casts a string literal value to a URI.

# Syntax

TOURI (text)

Argument Type		Description
text	string	The string literal to cast to a URI.

#### Returns

Туре	Description	
URI	The literal value as a URI.	

# TRIM

This function removes all spaces from a string except for any single spaces between words.

TRIM(text)		
Argument	Туре	Description
text	string	The string to trim.

Туре	Description	
string	The string with spaces removed.	

# UCASE

This function converts all letters in a string to upper case.

# Syntax

UPPER(text)		
Argument	Туре	Description
text	string	The string value to convert to upper case.

#### Returns

Туре	Description	
string	The string with upper case characters.	

# UPPER

This function converts all letters in a string literal to upper case.

UPPER(text)		
Argument	Туре	Description
text	string	The string literal to convert to upper case.

Туре	Description	
string	The string with upper case characters.	

# **Math Functions**

This topic describes the mathematical functions in Anzo.

#### **Typographical Conventions**

This documentation uses the following conventions in function syntax:

- CAPS: Although SPARQL is case-insensitive, function names and other keywords are written in uppercase for readability.
- [ argument ]: Brackets are used to indicate optional arguments. Arguments without brackets are required.

#### **Functions**

- ABS: Calculates the absolute value of the specified number.
- ADD: Adds two numeric values.
- AVERAGEIF: Calculates the average of the range of values that meet the specified criterion.
- AVERAGEIFS: Calculates the averages of the ranges of values that meet the specified criteria.
- AVG: Calculates the average (arithmetic mean) value for a group of numbers.
- CEILING: Rounds up a numeric value to the nearest integer.
- COS: Calculates the cosine of an angle.
- DIVIDE: Divides a number by another number.
- EQUAL: Evaluates whether two values are equal.
- EXP: Raises e to the specified power.
- FACT: Calculates the factorial of the specified number.
- FLOOR: Rounds down a numeric value to the nearest integer.
- GE: Evaluates whether one value is greater than or equal to (>=) another value.
- GT: Evaluates whether one value is greater than (>) another value.

- HAMMING\_DIST: Calculates the hamming distance between two values.
- HAVERSINE\_DIST: Computes the haversine distance between two latitude and longitude values.
- LE: Evaluates whether one value is less than or equal to (<=) another value.
- LN: Calculates the natural logarithm of a double value.
- LOG: Calculates the specified base logarithm of a double value.
- LOG2: Calculates the base two logarithm of a double value.
- LT: Evaluates whether one value is less than (<) another value.
- MAXVAL: Determines the maximum value from the given literal values.
- MINVAL: Determines the minimum value from the given literal values.
- MOD: Calculates the modulo of the division between two numbers.
- MULTIPLY: Multiplies two number values.
- NOT\_EQUAL: Evaluates whether two values are not equal.
- NPV: Calculates the net present value of an investment.
- NUMERIC-ADD: Adds two numeric values.
- NUMERIC-SUBTRACT: Subtracts one numeric value from another numeric value.
- PI: Returns the value for PI.
- POWER: Raises the specified number to the specified power.
- **PRODUCT**: Calculates the product of a group of numbers.
- QUOTIENT: Calculates the quotient between two numbers.
- RAD: Converts to radians an angle value that is in degrees.
- RAND: Returns a random double value between 0 and 1.
- RANDBETWEEN: Returns a random integer that falls between two specified integers.
- ROUND: Rounds a numeric value to the nearest integer.

- ROUNDDOWN: Rounds a numeric value down to the specified number of digits.
- ROUNDUP: Rounds a numeric value up to the specified number of digits.
- SIN: Calculates the sine of an angle.
- SQRT: Calculates the square root of a number.
- SUBTRACT: Subtracts one RDF term from another RDF term type value.
- SUM: Calculates the sum of the numbers within a group.
- SUMIF: Calculates the sum of the range of values that meet the specified criterion.
- SUMIFS: Calculates the sums of the ranges of values that meet the specified criteria.
- SUMPRODUCT: Multiplies the numbers in a group and adds the results.
- SUMSQ: Calculates the square root of each number in a group and adds the results.
- TAN: Calculates the tangent of an angle.

## ABS

This function calculates the absolute value of the specified number.

#### Syntax

ABS(number)

Argument	Туре	Description
number	numeric	The numeric value for which to calculate the absolute value.

Туре	Description
number	The absolute value.

# ADD

This function adds two numeric values.

## Syntax

ADD(value1, value2)

Argument	Туре	Description
value1	numeric	The first numeric value to add.
value2	numeric	The second numeric value to add.

#### Returns

Туре	Description	
number	The result of the addition operation.	

# AVERAGEIF

This function calculates the average of the range of values that meet the specified criterion.

## Тір

You can use AVERAGEIFS to specify multiple value ranges and conditions.

```
AVERAGEIF(values_to_test, criterion [, range_of_values ])
```

Argument	Туре	Description
values_to_ test	RDF term	The literal, URI, or blank node value that defines the values to test against the criteria.
criterion	RDF term	The literal, URI, or blank node value that defines the condition to

Argument	Туре	Description
		test values against.
range_of_ values	numeric	An optional number that defines the range of values to average. When omitted, values_to_test is used.

Туре	Description
number	The average value from the range of values that meet the criterion.

# **AVERAGEIFS**

This function calculates the averages of the ranges of values that meet the specified criteria. Unlike AVERAGEIF, this function enables you to specify multiple ranges and multiple conditions.

Argument	Туре	Description
values_to_ average	numeric	The numeric value that defines the overall range of values to evaluate.
value_ range1–N	RDF term	The literal, URI, or blank node value that defines the range of values to test against the corresponding criteria.
criteria1–N	RDF term	The literal, URI, or blank node value that defines the condition to test the corresponding value_range against.
Туре	Description	
--------	--	
number	The average values from the ranges of values that meet the criteria.	

## AVG

This function calculates the average (arithmetic mean) value for a group of numbers.

#### **Syntax**

AVG(number)		
Argument	Туре	Description
number	numeric	The numeric value for which to calculate the average.
Returns		

Туре	Description
number	The arithmetic mean of the input values.

# CEILING

This function rounds up a numeric value to the nearest integer if the value has a fractional part. CEILING returns the value itself if it is a whole number.



Туре	Description
number	The rounded up value.

# COS

This function calculates the cosine of the specified angle.

### Syntax

COS(angle)

Argument	Туре	Description
angle	double	The angle in radians (double data type) to calculate the cosine for. If you have angle values in degrees, you can use RAD to convert the degrees to radians.

### Returns

Туре	Description
double	The cosine of the angle.

# DIVIDE

This function divides one number by another number.

```
  DIVIDE (value1, value2)

  Argument
  Type
  Description

  value1
  numeric
  The number that is the dividend in the equation.
```

Argument	Туре	Description
value2	numeric	The number to divide value1 by.

Туре	Description
number	The result of the division operation.

## EQUAL

This function evaluates whether value1 is equal to value2.

### Syntax

EQUAL(value1, value2)

Argument	Туре	Description
value1	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value2.
value2	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value1.

#### Returns

Туре	Description
boolean	True if the values are equal and false if they are not.

### EXP

This function raises the base of the natural logarithms, e, to the specified power.

### Syntax

Returns		
power	numeric	The number to raise e to.
Argument	Туре	Description
EXP(power)		

Туре	Description
number	E raised to the specified power.

# FACT

This function calculates the factorial of the specified number.

### Syntax

FACT(number)

number int The number for which to calculate the fa	
	ctorial.

#### Returns

Туре	Description
int	The factorial of the input values.

# FLOOR

This function rounds down a numeric value to the nearest integer if the value has a fractional part. FLOOR returns the value itself if it is a whole number.

## Syntax

FLOOR (number	)	
Argument	Туре	Description
number	numeric	The numeric value to round down.
Returns		
Туре	Description	
number	The rounded do	wn value.

## GE

This function evaluates whether value1 is greater than or equal to (>=) value2.

## Syntax

GE(value1, value2)

Argument	Туре	Description
value1	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value2. This is the value that will be checked to see if it is greater than or equal to value2.
value2	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value1.

Туре	Description
boolean	True if value1 >= value2. False if not.

# GT

This function evaluates whether value1 is greater than (>) value2.

## Syntax

GE(value1, value2)

Argument	Туре	Description
value1	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value2. This is the value that will be checked to see if it is greater than value2.
value2	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value1.

#### Returns

Туре	Description
boolean	True if value1 > value2. False if not.

# HAMMING\_DIST

This function calculates the hamming distance between two values.

## Syntax

HAMMING	DIST	(number1,	number2)

Argument	Туре	Description
number1	long	The first number.
number2	long	The second number.

#### Returns

Туре	Description
int	The hamming distance.

## HAVERSINE\_DIST

This function computes the haversine distance between two latitude and longitude values and returns the distance in kilometers. The function uses the Haversine formula, which is accurate for most purposes but assumes a spherical Earth. Since the Earth is elliptical, distances involving points near the poles will be more inaccurate than other points.

```
HAVERSINE_DIST(latitude1, longitude1, latitude2, longitude2)
```

Argument	Туре	Description
latitude1	double	The first latitude value.
longitude1	double	The first longitude value.
latitude2	double	The second latitude value.
longitude2	double	The second longitude value.

Туре	Description
double	The distance in kilometers.

# LE

This function evaluates whether value1 is less than or equal to (<=) value2.

## Syntax

LE(value1, value2)

Argument	Туре	Description
value1	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value2. This is the value that will be evaluated to see if it is less than or equal to value2.
value2	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value1.

### Returns

Туре	Description
boolean	True if value1 <= value2. False if not.

# LN

This function calculates the natural logarithm of a double value.

## **Syntax**

LN(number)		
Argument	Туре	Description
number	double	The double value for which to calculate the natural logarithm.
Daturna		

#### Returns

Туре	Description
double	The natural logarithm of the input value.

## LOG

This function calculates the specified base logarithm of a double value.

## Syntax

LOG(number [, base ])

Argument	Туре	Description
number	double	The double value for which to calculate the base logarithm.
base	double	An optional double value that specifies the base for the logarithm. If omitted, base e is used.

#### **Returns**

Туре	Description
double	The base logarithm of the input value.

## LOG2

This function calculates the base two logarithm of a double value.

## Syntax

LOG2(number)		
Argument	Туре	Description
number	double	The double value for which to calculate the base 2 logarithm.

#### Returns

Туре	Description
double	The base two logarithm of the input value.

# LT

This function evaluates whether value1 is less than (<) value2.

# Syntax

LT(value1, value2)

Argument	Туре	Description
value1	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, or RDF term type value to compare to value2. This is the value that will be evaluated to see if it is less than value2.
value2	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, or RDF term type value to compare to value1.

Туре	Description
boolean	True if value1 < value2. False if not.

## MAXVAL

This function determines the maximum value from the given literal values.

## Syntax

MAXVAL(value1 [, value2 ] [, valueN ])			
Argument	Туре	Description	
value1–N	literal	A literal value from which you want to find the maximum value.	

#### Returns

Туре	Description
literal	The maximum value.

## MINVAL

This function determines the minimum value from the given literal values.

MINVAL(value1 [, value2 ] [, valueN ])			
Argument	Туре	Description	
value1–N	literal	A literal value from which you want to find the minimum value.	

Туре	Description
literal	The minimum value.

## MOD

This function calculates the modulo or remainder of the division between two numbers.

### Syntax

MOD(number, divisor)

Argument	Туре	Description
number	numeric	The number that is the dividend in the equation.
divisor	numeric	The number to divide the dividend by.

#### Returns

Туре	Description	
number	The modulo between the input numbers.	

# MULTIPLY

This function multiplies two numbers.

```
MULTIPLY(value1, value2)
```

Argument	Туре	Description
value1	numeric	The first number in the multiplication equation.

Argument	Туре	Description
value2	numeric	The number to multiply value1 by.

Туре	Description
number	The result of the multiplication operation.

# NOT\_EQUAL

This function evaluates whether value1 is not equal to value2.

#### Syntax

NOT\_EQUAL(value1, value2)

Argument	Туре	Description
value1	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value2.
value2	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value1.

#### Returns

Туре	Description
boolean	True if value1 is not equal to value2. False if they are equal.

#### NPV

This function calculates the net present value of an investment by using a discount rate and a series of future payments (negative values) and income (positive values).

## Syntax

NPV(rate, year, value)			
Argument	Туре	Description	
rate	double	A double value that defines the discount rate to use in the calculation.	
year	double	A double value that indicates which year the value is in.	
value	double	The double values that represent payments and income.	

#### Returns

Туре	Description
double	The net present value.

# NUMERIC-ADD

This function adds two numeric values.

## Syntax

NUMERIC-ADD(value1, value2)

Argument	Туре	Description
value1	numeric	The first number in the addition equation.
value2	numeric	The number to add to value1.

Туре	Description
number	The result of the addition operation.

## NUMERIC-SUBTRACT

This function subtracts one numeric value from another numeric value.

#### Syntax

```
NUMERIC-SUBTRACT(value1, value2)
```

Argument	Туре	Description
value1	numeric	The first number in the subtraction equation.
value2	numeric	The number to subtract from value1.

#### Returns

Туре	Description
number	The result of the subtraction operation.

## ΡΙ

This function returns the value for PI.

### Syntax

PI()

Туре	Description
double	The PI value.

## POWER

This function raises the specified number to the specified power.

### Syntax

POWER(value, power)

Argument	Туре	Description
value	numeric	The number to raise by the power.
power	numeric	The number to raise value by.

#### Returns

Туре	Description
number	The result of value raised to the specified power.

# PRODUCT

This function calculates the product of a group of numbers.

PRODUCT (number)	)	
Argument	Туре	Description
number	numeric	The group of numbers to multiply.

Туре	Description
number	The product of the group.

## QUOTIENT

This function calculates the quotient between two numbers.

### Syntax

```
QUOTIENT (numerator, denominator)
```

Argument	Туре	Description
numerator	numeric	The number to divide by the denominator.
denominator	numeric	The number to divide the numerator by.

#### Returns

Туре	Description
long	The quotient between the input values.

# RAD

This function converts to radians an angle value that is in degrees.

### Syntax

RAD(angle)

Argument	Туре	Description
angle	double	The angle value to convert to radians.

Туре	Description
double	The angle in radians.

## RAND

This function returns a random double value between 0 and 1.

### Syntax

RAND()

### Returns

Туре	Description
double	The random value between 0 and 1.

## RANDBETWEEN

This function returns a random integer that falls between the two specified integers. The two integers are included as options to be returned.

RANDBETWEEN (low	_number, h:	igh_number)
Argument	Туре	Description
low_number	int	The lowest integer in the range of values.
high_number	int	The highest integer in the range of values.

### Note

If the arguments are decimal values, Anzo returns a random integer between CEIL(low\_ number) and FLOOR(high\_number).

#### Returns

Туре	Description
int	The random value between the given low and high numbers.

## ROUND

This function rounds a numeric value to the nearest integer.

### Syntax

ROUND(number)

Argument	Туре	Description
number	numeric	The number to round to the nearest integer.

### Returns

Туре	Description
long	The rounded value.

# ROUNDDOWN

This function rounds a numeric value down to the specified number of digits.

### Syntax

ROUNDDOWN(number, num\_digits)

Argument	Туре	Description
number	numeric	The number to round down.
num_digits	int	An integer that specifies the number of digits to round down to.

Туре	Description
number	The rounded down value.

# ROUNDUP

This function rounds a numeric value up to the specified number of digits.

#### Syntax

ROUNDUP(number, num\_digits)

Argument	Туре	Description
number	numeric	The number to round up.
num_digits	int	An integer that specifies the number of digits to round up to.

#### Returns

Туре	Description
number	The rounded up value.

## SIN

This function calculates the sine of the specified angle.

## Syntax

SIN(angle)		
Argument	Туре	Description
angle	double	The angle in radians to calculate the sine for. If you have angle values in degrees, you can use RAD to convert the degrees to radians.

#### Returns

Туре	Description
double	The sine of the angle.

# SQRT

This function calculates the square root of the specified number.

# Syntax

SQRT(number)		
Argument	Туре	Description
number	numeric	The number for which to calculate the square root.

### Returns

Туре	Description
double	The square root of the input value.

# SUBTRACT

This function subtracts one RDF term type (a literal value, URI, or blank node) value from another RDF term type value.

## Syntax

SUBTRACT(term1,	term2)	
Argument	Туре	Description
term1	RDF term	The literal, URI, or blank node value that term2 will be subtracted from.
term2	RDF term	The literal, URI, or blank node value to subtract from term1.

#### Returns

Туре	Description
RDF term	The result of the subtraction operation.

## SUM

This function calculates the sum of the numbers within a group.

### Syntax

SUM(number)

Argument	Туре	Description
number	numeric	The group of numbers to sum.

Туре	Description
number	The sum of the values in the group.

### SUMIF

This function calculates the sum of the range of values that meet the specified criterion.

### Tip

You can use SUMIFS to specify multiple value ranges and conditions.

### Syntax

```
SUMIF(values_to_test, criterion [, range_of_values ])
```

Argument	Туре	Description
values_to_ test	RDF term	The literal, URI, or blank node value that defines the values to test against the criterion.
criterion	RDF term	The literal, URI, or blank node value that defines the condition to test values against.
range_of_ values	numeric	An optional number that defines the range of values to sum. When omitted, values_to_test is used.

#### Returns

Туре	Description
number	The sum of the range of values.

## SUMIFS

This function calculates the sums of the ranges of values that meet the specified criteria. Unlike SUMIF, this function enables you to specify multiple ranges and multiple conditions.

### Syntax

```
SUMIFS(values_to_sum, value_range1, criteria1,
value_range2, criteria2
[, value_rangeN, criteriaN ])
```

Argument	Туре	Description
values_to_ sum	numeric	The numeric value that defines the overall range of values to evaluate.
value_ range1–N	RDF term	The literal, URI, or blank node value that defines the range of values to test against the corresponding criteria.
criteria1–N	RDF term	The literal, URI, or blank node value that defines the condition to test the corresponding value_range against.

#### Returns

Туре	Description
number	The sums of the ranges of values.

## SUMPRODUCT

This function multiplies the numbers in a group and adds the results.

### Syntax

SUMPRODUCT (number)

Argument	Туре	Description
number	numeric	The group of numbers to multiply and then sum the results.

Туре	Description
number	The sum of the product of the numbers in the group.

## SUMSQ

This function calculates the square root of each number in a group and adds the results.

### Syntax

SUMSQ(number)		
Argument	Туре	Description
number	numeric	The group of numbers for which to calculate the square root and then sum the results.

#### Returns

Туре	Description
number	The sum of the square root of the numbers in the group.

## TAN

This function calculates the tangent of the specified angle.

#### Syntax

TAN(angle)

Argument	Туре	Description
angle	double	The angle in radians to calculate the tangent for. If you have angle values in degrees, you can use RAD to convert the degrees to radians.

Туре	Description
double	The tangent of the angle.

# **Aggregate Functions**

This topic describes the aggregate functions in Anzo.

**Typographical Conventions** 

This documentation uses the following conventions in function syntax:

- CAPS: Although SPARQL is case-insensitive, function names and other keywords are written in uppercase for readability.
- [ argument ]: Brackets are used to indicate optional arguments. Arguments without brackets are required.

#### Functions

- AVERAGEIF: Calculates the average of the range of values that meet the specified criterion.
- AVERAGEIFS: Calculates the averages of the ranges of values that meet the specified criteria.
- AVG: Calculates the average (arithmetic mean) value for a group of numbers.
- CHOOSE\_BY\_MAX: Returns the value from a group that corresponds to the maximum value from another group.
- CHOOSE\_BY\_MIN: Returns the value from a group that corresponds to the minimum value from another group.
- COUNT: Counts the number of values that exist for a group.
- COUNT\_DISTINCT: Counts the number of unique values that exist for a group.
- COUNTIF: Counts the number of values that meet the specified criterion.
- COUNTIFS: Counts the number of values that meet the specified criteria.
- **GROUP\_CONCAT**: Concatenates a group of strings into a single string.
- GROUPCONCAT: Concatenates a group of strings into a single string. This function is a customizable version of GROUP\_CONCAT.
- MAX: Returns the maximum value from each group of values.

- MEDIAN: Returns the median number out of a group of numbers.
- MIN: Returns the minimum value from each group of values.
- MODE: Returns the mode (the value that occurs most frequently) from a group of values.
- MODEPERCENT: Calculates the percentage of values in a group that belong to the mode.
- PERCENTILE\_CONT: Calculates a percentile based on the continuous distribution of the specified group of values.
- PERCENTILE\_DISC: Calculates a percentile based on the discrete distribution of the specified group of values.
- **PRODUCT**: Calculates the product of a group of numbers.
- SAMPLE: Returns an arbitrary value from the specified group of values.
- STDEV: Measures the standard deviation in a group of numbers.
- STDEVP: Calculates the product of the standard deviation for a group of numbers.
- SUM: Calculates the sum of the numbers within a group.
- SUMIF: Calculates the sum of the range of values that meet the specified criterion.
- SUMIFS: Calculates the sums of the ranges of values that meet the specified criteria.
- SUMPRODUCT: Multiplies the numbers in a group and adds the results.
- SUMSQ: Calculates the square root of each number in a group and adds the results.
- VAR: Calculates the unbiased (sample) variance of a group of numbers.
- VARP: Calculates the biased (population) variance of a group of numbers.
- WEIGHTEDAVERAGE: Calculates the weighted average of a group of values.

### AVERAGEIF

This function calculates the average of the range of values that meet the specified criterion.

### Tip

You can use AVERAGEIFS to specify multiple value ranges and conditions.

#### Syntax

AVERAGEIF(values\_to\_test, criterion [, range\_of\_values ])

Argument	Туре	Description
values_to_ test	RDF term	The literal, URI, or blank node value that defines the values to test against the criteria.
criterion	RDF term	The literal, URI, or blank node value that defines the condition to test values against.
range_of_ values	numeric	An optional number that defines the range of values to average. When omitted, values_to_test is used.

#### Returns

Туре	Description
number	The average value from the range of values that meet the criterion.

### **AVERAGEIFS**

This function calculates the averages of the ranges of values that meet the specified criteria. Unlike AVERAGEIF, this function enables you to specify multiple ranges and multiple conditions.

```
AVERAGEIFS(values_to_average, value_range1, criteria1,
        value_range2, criteria2
    [, value_rangeN, criteriaN ])
```

Argument	Туре	Description
values_to_ average	numeric	The numeric value that defines the overall range of values to evaluate.
value_ range1–N	RDF term	The literal, URI, or blank node value that defines the range of values to test against the corresponding criteria.
criteria1–N	RDF term	The literal, URI, or blank node value that defines the condition to test the corresponding value_range against.

Туре	Description
number	The average values from the ranges of values that meet the criteria.

# AVG

This function calculates the average (arithmetic mean) value for a group of numbers.

## Syntax

AVG(number)

Argument	Туре	Description
number	numeric	The numeric value for which to calculate the average.

#### Returns

Туре	Description
number	The arithmetic mean of the input values.

# CHOOSE\_BY\_MAX

This function calculates the maximum value for one group and returns the value from another group that corresponds to the maximum from the first group.

### Syntax

CHOOSE_BY_MAX(test, value)		
Argument	Туре	Description
test	RDF term	The group of literal, URI, or blank node values from which to find the maximum value.
value	RDF term	The group of literal, URI, or blank node values from which to return the value that corresponds to the maximum value of test.

#### Example

In a fictional ticket sales data set, the following statement returns the ID of the buyer who paid the most:

CHOOSE BY MAX(?totalPaid, ?buyerID)

#### Returns

Туре	Description
RDF	The term from the ${\tt value}$ group that corresponds to the maximum value from the ${\tt test}$
term	group.

### CHOOSE\_BY\_MIN

This function calculates the minimum value for one group and returns the value from another group that corresponds to the minimum from the first group.

#### Syntax

CHOOSE\_BY\_MIN(test, value)

Argument	Туре	Description
test	RDF term	The group of literal, URI, or blank node values from which to find the minimum value.
value	RDF term	The group of literal, URI, or blank node values from which to return the value that corresponds to the minimum value of test.

#### Example

In a fictional ticket sales data set, the following statement returns the ID of the seller who sold the least number of tickets:

CHOOSE\_BY\_MIN(?totalTickets, ?sellerID)

#### Returns

Туре	Description
RDF term	The term from the ${\tt value}$ group that corresponds to the minimum value from the ${\tt test}$ group.

## COUNT

This function counts the number of values that exist for a group.

#### Syntax

COUNT(value)

Argument	Туре	Description
value	RDF term	The group of literal, URI, or blank node values to count.

Туре	Description
long	The number of values in the group.

## COUNT\_DISTINCT

This function counts the number of unique values that exist for a group.

#### Syntax

COUNT_DISTINCT	(value)	
Argument	Туре	Description
value	RDF term	The group of literal, URI, or blank node values for which to count the number of distinct values.

#### Returns

Туре	Description
long	The number of unique values in the group.

# COUNTIF

This function counts the number of values that meet the specified criterion.

Tip

You can use COUNTIFS to specify multiple conditions.

### Syntax

COUNTIF(values\_to\_test, criterion)

Argument	Туре	Description
values_to_ test	RDF term	The literal, URI, or blank node value that defines the values to test against the criterion.
criterion	RDF term	The literal, URI, or blank node value that defines the condition to test values against.

Туре	Description
long	The number of values that meet the criterion.

# COUNTIFS

This function counts the number of values that meet the specified criteria. Unlike COUNTIF, this function enables you to specify multiple conditions.

### Syntax

```
COUNTIFS (values_to_count, criteria1 [, criteria2 ] [, criteriaN ])
```

Argument	Туре	Description
values_to_ count	RDF term	The literal, URI, or blank node value to compare against the criteria.
criteria1–N	RDF term	A literal, URI, or blank node value that defines a condition to test the values_to_count against.

### Returns

Туре	Description
long	The number of values that meet the specified conditions.

# GROUP\_CONCAT

This function concatenates a group of strings into a single string. It is a simplified version of GROUPCONCAT as it takes only one argument.

### Syntax

GROUP_CONCAT(te	ext)	
Argument	Туре	Description
text	string	The string property whose values to concatenate into a single string.

#### Returns

Туре	Description
string	The concatenated string.

## GROUPCONCAT

This function concatenates a group of strings into a single string. Unlike GROUP\_CONCAT, this function allows for customization of the separator to use as well as the configuration of limits and options like prefixes and suffixes.

Argument	Туре	Description
group1–N	string	The group(s) of strings to concatenate.
group_value_ separator	string	The separator string to use between the groups of strings if you specified more than one group.

Argument	Туре	Description
separator	string	The separator string to use between the values in a concatenated group of strings.
serialize	boolean	A boolean value that indicates whether returned values should be serialized with the value's data type.
row_limit	int	An integer that puts a maximum limit on the number of rows to retrieve for a group.
value_limit	int	An integer that puts a maximum limit on the number of values to retrieve from a group of rows.
delimit_blanks	boolean	A boolean value that indicates whether to delimit blanks with the separator value.
prefix	string	Optional string to add as a prefix to the resulting string.
suffix	string	Optional string to add as a suffix to the resulting string.
max_length	int	Optional integer that puts a maximum limit on the number of characters the resulting string can have.

Туре	Description
string	The concatenated string.

# MAX

This function returns the maximum value from each group of values.
MAX(value1 [, v	value2 ] [,	valueN ])
Argument	Туре	Description
value1–N	RDF term	The group(s) of literal, URI, or blank node values for which to return the maximum value.

#### Returns

Туре	Description
RDF term	The maximum value from each group.

# **MEDIAN**

This function returns the median value from a group of numbers. The median is the number in the group where half of the numbers are greater than the number and half are less than the number.

### Syntax

MEDIAN(number)		
Argument	Туре	Description
number	numeric	The group of numeric values for which to calculate the median.

#### Returns

Туре	Description
number	The median for the group.

### MIN

This function returns the minimum value from each group of values.

MIN(value1 [, v	value2 ] [,	valueN ])
Argument	Туре	Description
value1–N	RDF term	The group(s) of literal, URI, or blank node values for which to return the minimum value.

#### Returns

Туре	Description
RDF term	The minimum value from each group.

### MODE

This function returns the mode from a group of values. The mode is the value that occurs most frequently in the group.

### Syntax

MODE(value)

Argument	Туре	Description
value	RDF term	The group of literal, URI, or blank node values for which to return the mode.

#### Returns

Туре	Description
RDF term	The mode from the group.

## MODEPERCENT

This function calculates the percentage of values in a group that belong to the mode.

MODEPERCENT (va	lue)	
Argument	Туре	Description
value	RDF term	The group of literal, URI, or blank node values for which to calculate the modepercent.

#### Returns

Туре	Description
double	The percentage of values that belong to the mode.

# PERCENTILE\_CONT

This function calculates a percentile based on the continuous distribution of the specified group of values. The returned value is interpolated and may not be equal to any of the values in the group.

### **Syntax**

```
PERCENTILE CONT (percentile, value)
```

Argument	Туре	Description
percentile	float	A float value that specifies the percentile to compute.
value	RDF term	The group of literal, URI, or blank node values for which to calculate the percentile.

Туре	Description
RDF term	The interpolated percentile.

# PERCENTILE\_DISC

This function calculates a percentile based on the discrete distribution of the specified group of values.

### Syntax

PERCENTILE\_DISC(percentile, value)

Argument	Туре	Description
percentile	float	A float value that specifies the percentile to compute.
value	RDF term	The group of literal, URI, or blank node values for which to calculate the percentile.

### Returns

Туре	Description	
RDF term	The percentile based on the discrete distribution of the group.	

# PRODUCT

This function calculates the product of a group of numbers.

### Syntax

PRODUCT (number)

Argument	Туре	Description
number	numeric	The group of numbers to multiply.

Туре	Description
number	The product of the group.

# SAMPLE

This function returns an arbitrary value from the specified group of values.

## Syntax

SAMPLE(value)		
Argument	Туре	Description
value	RDF term	The group of literal, URI, or blank node values from which to choose a sample value.

### Returns

Туре	Description
RDF term	The arbitrary value from the group.

# STDEV

This function measures the standard deviation (amount of dispersion) of a group of numbers.

STDEV(value)		
Argument	Туре	Description
value	numeric	The numeric value that defines the set of numbers for which to measure the standard deviation.

Туре	Description
number	The standard deviation of the group.

# **STDEVP**

This function calculates the product of the standard deviation for a group of numbers.

# Syntax

STDEVP(value)		
Argument	Туре	Description
value	numeric	The numeric value that defines the set of numbers for which to measure the standard deviation and compute the product.

### Returns

Туре	Description
number	The product of the standard deviation for the group.

# SUM

This function calculates the sum of the numbers within a group.

# Syntax

SUM(number)

Argument	Туре	Description
number	numeric	The group of numbers to sum.

Туре	Description
number	The sum of the values in the group.

# SUMIF

This function calculates the sum of the range of values that meet the specified criterion.

## Tip

You can use SUMIFS to specify multiple value ranges and conditions.

### Syntax

```
SUMIF(values_to_test, criterion [, range_of_values ])
```

Argument	Туре	Description
values_to_ test	RDF term	The literal, URI, or blank node value that defines the values to test against the criterion.
criterion	RDF term	The literal, URI, or blank node value that defines the condition to test values against.
range_of_ values	numeric	An optional number that defines the range of values to sum. When omitted, values_to_test is used.

Туре	Description
number	The sum of the range of values.

# SUMIFS

This function calculates the sums of the ranges of values that meet the specified criteria. Unlike SUMIF, this function enables you to specify multiple ranges and multiple conditions.

### Syntax

```
SUMIFS(values_to_sum, value_range1, criteria1,
    value_range2, criteria2
    [, value_rangeN, criteriaN ])
```

Argument	Туре	Description
values_to_ sum	numeric	The numeric value that defines the overall range of values to evaluate.
value_ range1–N	RDF term	The literal, URI, or blank node value that defines the range of values to test against the corresponding criteria.
criteria1–N	RDF term	The literal, URI, or blank node value that defines the condition to test the corresponding value_range against.

### Returns

Туре	Description
number	The sums of the ranges of values.

# SUMPRODUCT

This function multiplies the numbers in a group and adds the results.

### Syntax

SUMPRODUCT (number)

Argument	Туре	Description
number	numeric	The group of numbers to multiply and then sum the results.

Туре	Description
number	The sum of the product of the numbers in the group.

# SUMSQ

This function calculates the square root of each number in a group and adds the results.

### Syntax

SUMSQ(number)		
Argument	Туре	Description
number	numeric	The group of numbers for which to calculate the square root and then sum the results.

#### Returns

Туре	Description
number	The sum of the square root of the numbers in the group.

# VAR

This function calculates the unbiased (sample) variance for a group of numbers.

### Syntax

VAR(value)

Argument	Туре	Description
value	numeric	The numeric value that defines the set of numbers for which to measure the variance.

Туре	Description
number	The unbiased variance for the group.

# VARP

This function calculates the biased (population) variance for a group of numbers.

### **Syntax**

VARP(value)		
Argument	Туре	Description
value	number	The value that defines the set of numbers for which to measure the population variance.

# **Returns**

Туре	Description
decimal	The biased variance for the group.

# WEIGHTEDAVERAGE

This function calculates the weighted average of a group of values.

# **Syntax**

WEIGHTEDAVERAGE (value, weight)

Argument	Туре	Description
value	decimal	The decimal value that defines the group of values for which to calculate the weighted average.
weight	decimal	The decimal value that defines the weight to use in the calculation.

Туре	Description
decimal	The weighted average for the group.

# **Date and Time Functions**

This topic describes the date, time, and duration functions in Anzo.

**Typographical Conventions** 

This documentation uses the following conventions in function syntax:

- CAPS: Although SPARQL is case-insensitive, function names and other keywords are written in uppercase for readability.
- [ argument ]: Brackets are used to indicate optional arguments. Arguments without brackets are required.

### **Functions**

- DATE: Returns an xsd:date value based on the specified year, month, and day.
- DATEPART: Returns the date part of a literal string, date, long, or dateTime value.
- DATETIME (or xsd:dateTime): Returns a dateTime value from the given string, long, or dateTime.
- DAY: Returns the day of the month from the specified date.
- DAYSFROMDURATION: Returns the days portion of a duration value.
- DUR\_TO\_MILLIS: Calculates the time in milliseconds from a duration or numeric value.
- DURATION: Returns an xsd:duration value from the given numeric or duration value.
- DURATIONFORMAT: Converts a duration or numeric value to a string in the specified duration format.
- DURATIONPERIODFORMAT: Calculates the duration between the given start and end values and returns a string in the specified duration format.
- FORMATDATE: Converts a numeric or date value into a string with the specified date format.
- HOUR: Returns the hour portion of the given dateTime value.

- MASKEDDATETIME: Replaces the year, month, day, hour, minute, second, and millisecond values for the given date or dateTime value with the new date and time values that you specify.
- MILLIS: Calculates the number of milliseconds in the given date or dateTime value.
- MINUTE: Returns the minutes portion of the given dateTime value.
- MONTH: Returns the month portion of the given dateTime value.
- NOW: Returns the current server date and time.
- NOWMILLIS: Returns the current server date and time in epoch milliseconds.
- PARSEDATETIME: Attempts to convert the given string or plain literal to a date, time, or dateTime value.
- SECOND: Returns the seconds portion of the given dateTime value.
- TIME: Returns an xsd:time value based on the specified hour, minute, and second values.
- TIMEPART: Returns the time part of a time or dateTime value.
- TIMEVALUE: Converts the specified RDF term type value to an xsd:time value.
- TIMEZONE: Returns as a duration the timezone part of a dateTime value.
- TODAY: Returns today's date based on the server date.
- TZ: Returns as a string the timezone from a dateTime value.
- WEEKDAY: Returns the day of the week from a date or dateTime value.
- WEEKNUM: Returns the week of the year in which the given date or dateTime occurs.
- xsd:date: Converts the specified string, date, or dateTime value to an xsd:date.
- YEAR: Returns the year portion of the given dateTime value.
- YEARMONTH: Returns the year and month (in the format "year-month") from the specified date or dateTime value.

# DATE

This function returns an xsd:date value based on the specified year, month, and day values.

# Syntax

```
DATE(year, month, day)
```

Argument	Туре	Description
year	int	An integer that represents the year.
month	int	An integer that represents the month.
day	int	An integer that represents the day.

## Returns

Туре	Description
date	The date according to the input values.

# DATEPART

This function returns the date part of a literal string, date, long, or dateTime value.

# Syntax

DATEPART (value)

Argument	Туре	Description
value	literal string, date, long, or dateTime	The literal string, date, long, or dateTime value from which to return the date.

Туре	Description
date	The date part of the input values.

# DATETIME (or xsd:dateTime)

This function returns a dateTime value from the given string, long, or dateTime.

### Syntax

DATETIME(value)		
Argument	Туре	Description
value	string, long, dateTime	The string, long, or dateTime value from which to return a dateTime.

### Returns

Туре	Description
dateTime	The dateTime value.

# DAY

This function returns the day of the month from the specified date value.

DAY(value)		
Argument	Туре	Description
value	date	The date value from which to return the day of the month.

Туре	Description
int	The day of the month.

# DAYSFROMDURATION

This function returns the days portion of a duration value.

### Syntax

DAYSFROMDURATIO	ON(value)	
Argument	Туре	Description
value	duration	The duration value from which to return the days.

#### Returns

Туре	Description
long	The number of days in the duration.

# DUR\_TO\_MILLIS

This function calculates the time in milliseconds from a duration or numeric value.

DUR_TO_MILLIS (	value)	
Argument	Туре	Description
value	duration, numeric	The duration or numeric value from which to calculate the time in milliseconds.

Туре	Description
long	The number of milliseconds.

# DURATION

This function returns an xsd:duration value from the given numeric or duration value.

# Syntax

DURATION(value)	)	
Argument	Туре	Description
value	duration, numeric	The duration or numeric value from which to return an xsd:duration.

### Returns

Туре	Description
duration	The duration value.

# DURATIONFORMAT

This function converts a duration or numeric value to a string in the specified duration format.

DURATIONFORMAT(value [, format ])		
Argument	Туре	Description
value	duration, numeric	The duration or numeric value to format.

Argument	Туре	Description
format	string	An optional value that specifies the format to use for the resulting duration string. Anzo supports Pattern Tokens for defining the format:
		• <b>y</b> for year digits
		• <b>M</b> for months
		• d for days
		H for hours
		• <b>m</b> for minutes
		• s for seconds
		S for milliseconds
		'text' for arbitrary text content
		<pre>If format is not specified, H:mm:ss.SSS is used.</pre>

Туре	Description
string	The duration as a string.

# DURATIONPERIODFORMAT

This function calculates the duration between the given start and end dateTime or numeric values and returns a string in the specified duration format.

### Syntax

DURATIONPERIODFORMAT(start, end [, format ])

Argument	Туре	Description
start	dateTime, numeric	The dateTime or numeric value that is the start of the duration period.
end	dateTime, numeric	The dateTime or numeric value that is the end of the duration period.
format	string	An optional value that specifies the format to use for the resulting duration string. Anzo supports Pattern Tokens for defining the format:
		• <b>y</b> for year digits
		• <b>M</b> for months
		• d for days
		H for hours
		• <b>m</b> for minutes
		• s for seconds
		S for milliseconds
		text' for arbitrary text content
		If format is not specified, the default is 'P'yyyy'Y'M'M'd'DT'H'H'm'M's.SSS'S'. The default value uses 'text' patterns with Pattern Tokens, which results in a string such as P1Y3M4DT1H4M44.000S.

Туре	Description
string	The duration as a string.

# FORMATDATE

This function converts a numeric or date value into a string with the specified date format.

### Syntax

FORMATDATE(value, format)

value o	date, numeric	The date or numeric value to convert to a string in the specified date format.
format	string	<ul> <li>The format to use for the resulting date string. Anzo supports</li> <li>Pattern Tokens for defining the format: <ul> <li>y for year digits</li> <li>M for months</li> <li>d for days</li> <li>'text' for arbitrary text content</li> </ul> </li> </ul>

#### Returns

Туре	Description
string	The date as a string.

# HOUR

This function returns the hour portion of the given dateTime value.

### Syntax

HOUR(value [, timezone ])

Argument	Туре	Description
value	dateTime	The dateTime value from which to return the hour portion.
timezone	string	An optional value that specifies the timezone for the value.

Туре	Description
int	The hour.

# MASKEDDATETIME

This function replaces the year, month, day, hour, minute, second, and millisecond values for the given date or dateTime value with the new date and time values that you specify.

### Syntax

MASKEDDATETIME (value, year, month, day, hour, minute, second, milliseconds)

Argument	Туре	Description
value	date, dateTime	The date or dateTime for which to replace the year, month, date, hour, minute, second, and milliseconds values.
year	int	The year to include in the resulting dateTime value.
month	int	The month to include in the resulting dateTime value.
day	int	The day to include in the resulting dateTime value.
hour	int	The hour to include in the resulting dateTime value.
minute	int	The minutes value to include in the resulting dateTime value.

Argument	Туре	Description
second	int	The seconds value to include in the resulting dateTime value.
milliseconds	int	The milliseconds value to include in the resulting dateTime value.

Туре	Description	
dateTime	The dateTime value with the specified input values.	

# MILLIS

This function calculates the number of milliseconds in the given date or dateTime value.

# Syntax

MILLIS(value)

Argument	Туре	Description
value	date, dateTime	The date or dateTime value for which to calculate the number of milliseconds.

### Returns

Туре	Description
long	The number of milliseconds.

# MINUTE

This function returns the minutes portion of the given dateTime value.

MINUTE (value)		
Argument	Туре	Description
value	dateTime	The dateTime value from which to return the minutes portion.
Returns		

Туре	Description
int	The minutes portion of the input value.

## MONTH

This function returns the month portion of the given dateTime value.

# Syntax

MONTH(value)

Argument	Туре	Description
value	dateTime	The dateTime value from which to return the month portion.

### Returns

Туре	Description
int	The month number.

# NOW

This function returns the current server date and time.

### Syntax

NOW([ timezone ])

Argument	Туре	Description
timezone	string	An optional value that specifies the timezone for which to return the current dateTime.

Туре	Description
dateTime	The current server date and time.

## NOWMILLIS

This function returns the current server date and time in epoch milliseconds.

#### **Syntax**

NOWMILLIS()

#### Returns

Туре	Description
long	The current server date and time in milliseconds.

# PARSEDATETIME

This function attempts to convert the given string or plain literal to a date, time, or dateTime value. For values that do not include a timezone, Anzo stores them in GMT. Values that include a timezone are stored as the appropriate value in GMT for that timezone.

```
PARSEDATETIME(value [, output_type ] [, format ])
```

Argument	Туре	Description
value	string,	The string or plain literal value to convert to a date, time, or dateTime.

Argument	Туре	Description	
	literal		
output_type	URI	An optional URI (xsd:date, xsd:time, or xsd:dateTime) that specifies the type of value to return. If output_type is not specified, dateTime is returned.	
format	string	An optional string that species the format to use for the resulting date, time, or dateTime value. Anzo supports Pattern Tokens for defining the format:	
		• <b>y</b> for year digits	
		• <b>M</b> for months	
		• d for days	
		• H for hours	
		• <b>m</b> for minutes	
		• s for seconds	
		S for milliseconds	
		'text' for arbitrary text content	
		For example, "yyyy.MM.dd HH:mm" or "dd/MM/yyyy HH:mm:ss".	

Туре	Description
date, time, or dateTime	The conversion of the string to the desired type.

# SECOND

This function returns the seconds portion of the given dateTime value.

SECOND(val	ue)	
Argument	t Type	Description
value	dateTime	The dateTime value from which to return the seconds portion.
Returns		
Туре	Description	

int	The seconds portion of the input value.

# TIME

This function returns an xsd:time value based on the specified hour, minute, and second values.

# Syntax

TIME (hour, minute, second)

Argument	Туре	Description
hour	int	An integer that represents the hour.
minute	int	An integer that represents the minute.
second	int	An integer that represents the seconds.

Туре	Description
time	The time according to the input values.

# TIMEPART

This function returns the time part of a time or dateTime value.

# Syntax

TIMEPART(value)		
Argument	Туре	Description
value	time, dateTime	The time or dateTime value from which to return the time portion.

### Returns

Туре	Description
time	The time portion of the input value.

# TIMEVALUE

This function converts the specified RDF term type value to an xsd:time value.

# Syntax

TIMEVALUE(value	=)	
Argument	Туре	Description
value	RDF term	The literal, URI, or blank node value to convert to a time value.

Туре	Description
time	The conversion of the term to a time value.

# TIMEZONE

This function returns the timezone part of a dateTime value as a duration. An error is returned if the input value does not include the timezone.

### Syntax

TIMEZONE (value)		
Argument	Туре	Description
value	dateTime	The dateTime value to return the timezone from.

### Returns

Туре	Description
duration	The timezone in duration format.

# TODAY

This function returns today's date based on the server date.

### Syntax

TODAY()

### Returns

Туре	Description
date	Today's date according to the server.

# ΤZ

This function returns the timezone part of a dateTime value as a string.

### Syntax

TZ(value)

Argument	Туре	Description
value	dateTime	The dateTime value to return the timezone from.

Туре	Description
string	The timezone as a string.

# WEEKDAY

This function returns the day of the week from a date or dateTime value.

# Syntax

WEEKDAY(value [, day\_number\_start ])

Argument	Туре	Description
value	date, dateTime	The date or dateTime value from which to return the day of the week.
day_ number_ start	int	<ul> <li>An optional value of 1, 2, or 3 that defines how the days of the week are represented as numbers.</li> <li>1 means Sunday is day 1. Saturday is day 7.</li> <li>2 means Monday is day 1. Sunday is day 7.</li> <li>3 means Monday is day 0. Sunday is day 6.</li> <li>If day_number_start is not specified, the default value is 1.</li> </ul>

Туре	Description
int	The day of the week from the input values.

# WEEKNUM

This function returns the week of the year in which the given date or dateTime occurs.

### Syntax

```
WEEKNUM(value [, day_week_begins ])
```

Argument	Туре	Description
value	date, dateTime	The date or dateTime value from which to return the week number.
day_week_ begins	int	An optional value of <b>1</b> or <b>2</b> that defines which day the weeks start on.
		<ul> <li>1 means a new week starts on Sunday.</li> </ul>
		• 2 means a new week starts on Monday.
		<pre>If day_week_begins is not specified, the default value is 1.</pre>

#### Returns

Туре	Description	
int	The week of the year the input value falls in.	

# xsd:date

This function converts the specified string, date, or dateTime value to an xsd:date.

### Syntax

xsd:date(value)

Argument	Туре	Description
value	string, date, dateTime	The string, date, or dateTime value to convert to an xsd:date.

Туре	Description
date	The input values converted to dates.

# YEAR

This function returns the year portion of the given dateTime value.

The year portion of the input values.

### Syntax

YEAR(value	2)	
Argumen	t Туре	Description
value	dateTime	The dateTime value to return the year from.
Returns		
Туре	Description	

# YEARMONTH

This function returns the year and month (in the format "year-month") from the specified date or dateTime value.

### Syntax

int

YEARMONTH(value)

Argument	Туре	Description
value	literal date or dateTime	The value to return the year-month from.

Туре	Description
gYearMonth	The year-month from the input values.

# **Casting Functions**

This topic describes the functions that are available for coercing data types in Anzo.

**Typographical Conventions** 

This documentation uses the following conventions in function syntax:

- CAPS: Although SPARQL is case-insensitive, function names and other keywords are written in uppercase for readability.
- [ argument ]: Brackets are used to indicate optional arguments. Arguments without brackets are required.

### Functions

- BNODE: Creates a blank node.
- BOOLEAN: Casts a literal value to a boolean data type.
- BYTE: Casts a literal value to a byte data type.
- CONCATURL: Concatenates two or more strings and returns the result as a URI.
- DATATYPE: Returns the data type of the given value.
- DATETIME (or xsd:dateTime): Returns a dateTime value from the given string, long, or dateTime.
- DATEVALUE: Casts a string to a date.
- DECIMAL: Casts a literal value to a decimal data type.
- DOUBLE: Casts a literal value to a double data type.
- DURATION: Returns an xsd:duration value from the given numeric or duration value.
- DURATIONFORMAT: Converts a duration or numeric value to a string in the specified duration format.
- ENCODE\_FOR\_URI: Encodes the specified string as a URI.
- FLOAT: Casts a literal value to a float data type.

- FORMATDATE: Casts a numeric or date value to a string in the specified date format.
- FORMATFRACTION: Converts a numeric value into a fraction string.
- FORMATNUMBER: Casts a numeric value to a string in the specified format.
- INT: Casts a literal value to an int data type.
- INTEGER: Casts a literal value to an integer data type.
- LONG: Casts a literal value to a long data type.
- PARSEDATETIME: Attempts to convert the given string or plain literal to a date, time, or dateTime value.
- RAD: Converts to radians an angle value that is in degrees.
- SERIALIZE: Creates a literal value from the string representation of the specified RDF term.
- SHORT: Casts a literal value to a short data type.
- STR: Casts an RDF term type value to a string.
- TEXT: Casts a numeric or dateTime value to a string in the specified format.
- TIMEVALUE: Converts the specified RDF term type value to an xsd:time value.
- TOURI: Casts a string literal value to a URI.
- UUID: Generates a Universally Unique Identifier (UUID).
- xsd:date: Converts the specified string, date, or dateTime value to an xsd:date.

# BNODE

This function creates a blank node.

BNODE([ value ])	BNODE ([	value	])
------------------	----------	-------	----

Argument	Туре	Description
value	string	An optional string value from which to create the blank node.

Туре	Description
blank node	The generated blank node.

# BOOLEAN

This function casts a literal value to a boolean data type.

# Syntax

BOOLEAN(value)

Argument	Туре	Description
value	literal	The literal value to cast to a boolean type.

### Returns

Туре	Description
boolean	The input cast to boolean.

# BYTE

This function casts a literal value to a byte data type.

BYTE(value)		
Argument	Туре	Description
value	literal	The literal value to cast to a byte type.

Туре	Description
byte	The input value cast to a byte.

# CONCATURL

This function concatenates two or more strings and returns the result as a URI.

# Syntax

CONCATURL (text)	1, text2 [,	textN ])
Argument	Туре	Description
text1–N	string	The strings that you want to concatenate to form a URI.

#### Returns

Туре	Description
URI	The concatenated string as a URI.

# DATATYPE

This function returns the data type of the given literal value.

DATATYPE(value)		
Argument	Туре	Description
value	literal	The literal value for which to return the data type.
Туре	Description	
------	----------------	
URI	The data type.	

# DATETIME (or xsd:dateTime)

This function returns a dateTime value from the given string, long, or dateTime.

# Syntax

DATETIME (value)		
Argument	Туре	Description
value	string, long, dateTime	The string, long, or dateTime value from which to return a dateTime.

#### Returns

Туре	Description
dateTime	The dateTime value.

# DATEVALUE

This function casts the given string value to a date.

DATEVALUE(value	e)	
Argument	Туре	Description
value	string	The string value from which to return the date.

Туре	Description
date	The string cast to a date.

# DECIMAL

This function casts a literal value to a decimal data type.

# Syntax

DECIMAL(value)		
Argument	Туре	Description
value	literal	The literal value to convert to a decimal type.

### Returns

Туре	Description
decimal	The literal value cast to a decimal type.

# DOUBLE

This function casts a literal value to a double data type.

DOUBLE(value)			
Argument	Туре	Description	
value	literal	The literal value to convert to a double type.	

Туре	Description
double	The literal value cast to a double type.

# DURATION

This function returns an xsd:duration value from the given numeric or duration value.

### Syntax

DURATION (value)	)	
Argument	Туре	Description
value	duration, numeric	The duration or numeric value from which to return an xsd:duration.

#### Returns

Туре	Description
duration	The duration value.

# DURATIONFORMAT

This function converts a duration or numeric value to a string in the specified duration format.

DURATIONFORMAT(value [, format ])		
Argument	Туре	Description
value	duration, numeric	The duration or numeric value to format.

Argument	Туре	Description
format	string	An optional value that specifies the format to use for the resulting duration string. Anzo supports Pattern Tokens for defining the format:
		• <b>y</b> for year digits
		• <b>M</b> for months
		• d for days
		H for hours
		• <b>m</b> for minutes
		• s for seconds
		S for milliseconds
		'text' for arbitrary text content
		<pre>If format is not specified, H:mm:ss.SSS is used.</pre>

Туре	Description
string	The duration as a string.

# ENCODE\_FOR\_URI

This function encodes the specified string as a URI and returns a string in URI format.

ENCODE_FOR_URI(text)			
Argument	Туре	Description	
text	string	The string value to encode as a URI.	

Туре	Description
string	The string as a URI.

# FLOAT

This function casts a literal value to a float data type.

# Syntax

FLOAT(value)		
Argument	Туре	Description
value	literal	The literal value to convert to a float type.

#### Returns

Туре	Description
float	The literal value cast to a float type.

# FORMATDATE

This function converts a numeric or date value into a string with the specified date format.

FORMATDATE(value, format)		
Argument	Туре	Description
value	date, numeric	The date or numeric value to convert to a string in the specified date format.
format	string	The format to use for the resulting date string. Anzo supports

Argument	Туре	Description
		Pattern Tokens for defining the format:
		• <b>y</b> for year digits
		• <b>M</b> for months
		• <b>d</b> for days
		'text' for arbitrary text content
		For example, "yyyy.MM.dd" or "dd/MM/yyyy".

Туре	Description
string	The date as a string.

# FORMATFRACTION

This function converts a numeric value into a fraction string.

# Syntax

FORMATFRACTION(value [, tolerance ] [, whole\_number ])

Argument	Туре	Description
value	numeric	The numeric value to convert to fraction text.
tolerance	double	An optional double value that specifies the precision of the fraction. The default value is 0.0001. The resulting fractional representation is the original value + or - the tolerance. The smaller the tolerance, the more precise the fraction is. For example, 399/800 vs. 1/2.

Argument	Туре	Description
whole_ number	boolean	An optional boolean value that specifies whether to include whole numbers in the result. For example, if true, the result would be formatted like 1 2/3 instead of 5/3.

Туре	Description
string	The fraction string.

# FORMATNUMBER

This function casts a numeric value to a string in the specified format.

### Syntax

TEXT(value, format)			
Argument	Туре	Description	
value	numeric	The numeric value to convert to a string.	
format	string	A text string that specifies the format to follow when converting the	

value to a string. Anzo supports Java Decimal Format.

#### Returns

Туре	Description	
string	The numeric value as a string.	

# INT

This function casts a literal value to an int data type.

### Syntax

INT(value)		
Argument	Туре	Description
value	literal	The literal value to convert to an int type.

#### Returns

Туре	Description	
int	The literal value cast to an int type.	

### INTEGER

This function casts a literal value to an integer data type.

# Syntax

INTEGER(value)

Argument	Туре	Description
value	literal	The literal value to convert to an integer type.

#### Returns

Туре	Description	
integer	The literal value cast to an integer type.	

# LONG

This function casts a literal value to a long data type.

### Syntax

LONG(value)

Argument	Туре	Description
value	literal	The literal value to convert to a long type.

Туре	Description
long	The literal value cast to a long type.

# PARSEDATETIME

This function attempts to convert the given string or plain literal to a date, time, or dateTime value. For values that do not include a timezone, Anzo stores them in GMT. Values that include a timezone are stored as the appropriate value in GMT for that timezone.

#### Syntax

PARSEDATETIME(value [, output\_type ] [, format ])

Argument	Туре	Description	
value	string, literal	The string or plain literal value to convert to a date, time, or dateTime.	
output_type	URI	An optional URI (xsd:date, xsd:time, or xsd:dateTime) that specifies the type of value to return. If output_type is not specified, dateTime is returned.	
format	string	An optional string that species the format to use for the resulting date, time, or dateTime value. Anzo supports Pattern Tokens for defining the format: • <b>y</b> for year digits	
		• <b>M</b> for months	

Argument	Туре	Description	
		<ul> <li>d for days</li> <li>H for hours</li> <li>m for minutes</li> <li>s for seconds</li> <li>S for milliseconds</li> </ul>	
		• 'text' for arbitrary text content For example, "yyyy.MM.dd HH:mm" or "dd/MM/yyyy HH:mm:ss".	

Туре	Description
date, time, or dateTime	The conversion of the string to the desired type.

# RAD

This function converts to radians an angle value that is in degrees.

# Syntax

RAD(angle)

Argument	Туре	Description
angle	double	The angle value to convert to radians.

Туре	Description
double	The angle in radians.

# SERIALIZE

This function creates a literal value from the string representation of the specified RDF term type value.

### Syntax

SERIALIZE (value	e)	
Argument	Туре	Description
value	RDF term	The literal, URI, or blank node value for which to generate a plain literal.

### Returns

Туре	Description
string	The string representation of the input term.

# SHORT

This function casts a literal value to a short data type.

# Syntax

SHORT(value)		
Argument	Туре	Description
value	literal	The literal value to convert to a short type.

Туре	Description
short	The literal value cast to a short type.

# STR

This function casts the specified RDF term type value to a string.

# Syntax

STR(value)		
Argument	Туре	Description
value	RDF term	The literal, URI, or blank node value to convert to a string.

### Returns

Туре	Description
string	The term cast to a string type.

# TEXT

This function casts a numeric or dateTime value to a string in the specified format.

TEXT (value, for	rmat)	
Argument	Туре	Description
value	numeric, dateTime	The numeric or datetime value to convert to a string.
format	string	A text string that specifies the format to follow when converting the value to a string. For numeric values, Anzo supports <u>Java</u> <u>Decimal Format</u> . For dateTime values, <u>Simple Date Format</u> is supported.

Туре	Description
string	The string in the specified format.

# TIMEVALUE

This function converts the specified RDF term type value to an xsd:time value.

# Syntax

TIMEVALUE(value	e)	
Argument	Туре	Description
value	RDF term	The literal, URI, or blank node value to convert to a time value.

#### Returns

Туре	Description
time	The conversion of the term to a time value.

# TOURI

This function casts a string literal value to a URI.

TOURI	(text)		
Argu	ument	Туре	Description
text		string	The string literal to cast to a URI.

Туре	Description
URI	The literal value as a URI.

# UUID

This function generates a Universally Unique Identifier (UUID).

#### Syntax

UUID()

#### Returns

Туре	Description
URI	The UUID.

# xsd:date

This function converts the specified string, date, or dateTime value to an xsd:date.

### Syntax

xsd:date(value)

Argument	Туре	Description
value	string, date, dateTime	The string, date, or dateTime value to convert to an xsd:date.

Туре	Description	
date	The input values converted to dates.	

# **Logical Functions**

This topic describes the logical functions in Anzo.

**Typographical Conventions** 

This documentation uses the following conventions in function syntax:

- CAPS: Although SPARQL is case-insensitive, function names and other keywords are written in uppercase for readability.
- [ argument ]: Brackets are used to indicate optional arguments. Arguments without brackets are required.

#### Functions

- AND: Evaluates two logical expressions and returns true if both expressions are true.
- BOUND: Evaluates whether an RDF term type is bound.
- CASE: Evaluates a series of conditions for the specified value and returns the matching result.
- COALESCE: Evaluates a number of expressions and returns the results for the first expression that is bound and does not raise an error.
- EQUAL: Evaluates whether two values are equal.
- IF: Evaluates a condition and returns the specified result depending on the outcome of the test.
- IFERROR: Synonym for COALESCE.
- IN: Evaluates whether the specified RDF term is found in any of the given test values.
- NOT: Evaluates whether the specified logical expression is not true.
- NOT\_EQUAL: Evaluates whether two values are not equal.
- NOT\_IN: Evaluates whether the specified RDF term is not found in any of the given test values.

- OR: Evaluates two logical expressions and returns true if at least one of the expressions is true.
- PARTITIONINDEX: Returns the zero-based index of the bucket in which the specified value falls.
- SAMETERM: Evaluates whether two RDF term type values are the same.
- UNBOUND: Extends the SPARQL UNDEF functionality to enable users to include an undefined value as a function argument.

### AND

This function evaluates two logical expressions. If both expressions are true, the function returns true. If one or both arguments are false, the function returns false.

#### Syntax

AND(logical\_expression1, logical\_expression2)

Argument	Туре	Description
logical_expression1	evaluates to boolean	The first logical expression to evaluate.
logical_expression2	evaluates to boolean	The second logical expression to evaluate.

#### Returns

Туре	Description
boolean	True if both conditions are true and false if either condition is false.

# BOUND

This function evaluates whether the specified RDF term has a value bound to it.

### Syntax

BOUND (term)

Argument	Туре	Description
term	RDF term	The literal, URI, or blank node value to evaluate.

Туре	Description
boolean	True if the term is bound and false if it is not.

# CASE

This function evaluates a series of conditions for the specified value and returns the matching result. CASE acts like an IF-THEN-ELSE statement.

Argument	Туре	Description
test_value	RDF term	The literal, URI, or blank node value to compare to the list of conditions (condition1-N).
condition1–N	RDF term	The conditions to be evaluated in the order that they are specified. Once a condition evaluates to true, the corresponding result is returned and the remaining conditions are not evaluated.
result1–N	RDF term	The results to return for the specified conditions.
default	RDF term	An optional value to be returned if none of the specified conditions pass.

Туре	Description
RDF term	The specified result according to the evaluation of the conditions.

# COALESCE

This function evaluates a number of expressions and returns the results for the first expression that is bound and does not raise an error.

### Syntax

COALESCE(expression1 [, expression2 ] [, expressionN ] )		
Argument	Туре	Description
expression1–N	RDF term	The literal, URI, or blank node expressions to evaluate.

#### Returns

Туре	Description
RDF term	The result of the first expression that is bound and does not error.

# EQUAL

This function evaluates whether value1 is equal to value2.

```
EQUAL(value1, value2)
```

Argument	Туре	Description
value1	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value2.

Argument	Туре	Description
value2	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value1.

Туре	Description
boolean	True if the values are equal and false if they are not.

### IF

This function evaluates a condition and returns the specified result depending on the outcome of the test. If the condition evaluates to true, the first result is returned. If the condition evaluates to false, the second result is returned. And if the condition results in an error, the third result is returned.

#### **Syntax**

IF(logical\_expression, true\_result, false\_result [, error\_result ])

Argument	Туре	Description
logical_ expression	evaluates to boolean	The condition that evaluates to true or false.
true_result	RDF term	The value that defines the result to return if the condition evaluates to true.
false_result	RDF term	The value that defines the result to return if the condition evaluates to false.
error_result	RDF term	An optional value that defines the result to return if the condition evaluates to an error. If the condition results in an error and error_result is not specified, logical_expression (error) is returned.

Туре	Description
RDF term	The result based on the evaluation of the condition.

# IFERROR

This function is a synonym for COALESCE.

### IN

This function evaluates whether the specified RDF term type value is found in any of the given test values.

#### **Syntax**

IN(term, test\_value1 [, test\_value2 ] [, test\_valueN])

Argument	Туре	Description
term	RDF term	The literal, URI, or blank node value to look for in the test values.
test_value1– N	RDF term	The literal, URI, or blank node values to look for the specified term in.

#### Returns

Туре	Description
boolean	True if the given term is found in the test values and false if it is not.

# NOT

This function evaluates whether the specified logical expression is not true.

### Syntax

NOT(logical\_expression)

Argument	Туре	Description
logical_expression	evaluates to boolean	The condition to evaluate.

#### Returns

Туре	Description
boolean	True if the condition is false and false if it is true.

# NOT\_EQUAL

This function evaluates whether value1 is not equal to value2.

# Syntax

NOT\_EQUAL(value1, value2)

Argument	Туре	Description
value1	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value2.
value2	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value1.

Туре	Description
boolean	True if value1 is not equal to value2. False if they are equal.

# NOT\_IN

This function evaluates whether the specified RDF term type value is not found in any of the given test values.

### Syntax

IN(term, test\_value1 [, test\_value2 ] [, test\_valueN])

Argument	Туре	Description
term	RDF term	The literal, URI, or blank node value to look for in the test values.
test_value1– N	RDF term	The literal, URI, or blank node values in which to look for the specified term.

#### Returns

Туре	Description
boolean	True if the given term is not in the test values and false if it is found in the test values.

# OR

This function evaluates two logical expressions. If at least one expression is true, the function returns true. If both expressions are false, the function returns false.



Argument	Туре	Description
logical_expression2	evaluates to boolean	The second logical expression to evaluate.

Туре	Description
boolean	True if one or both conditions are true and false if both conditions are false.

# PARTITIONINDEX

This function returns the zero-based index of the bucket in which the specified value falls. Buckets start at the specified start value and are sized according to the specified interval. The first bucket is [start, start+interval). That means it is closed on the low end and open on the high end. PARTITIONINDEX returns less than 0 if the value does not fall into any bucket, such as when the given value is less than start or if the comparison is indeterminate for date and time data types.

#### **Syntax**

PARTITIONINDEX(value, start, interval)

Argument	Туре	Description
value	literal	The literal value for which to determine the zero-based index.
start	literal	The literal value that indicates the start of the first bucket.
interval	literal	The literal value that specifies the size of the bucket.

Туре	Description
long	The zero-based index of the bucket in which the specified value exists.

# SAMETERM

This function evaluates whether two RDF term type values are the same.

#### Syntax

 SAMETERM (term1, term2)

 Argument
 Type
 Description

 term1
 RDF term
 The first literal, URI, or blank node value to compare.

 term2
 RDF term
 The literal, URI, or blank node value to compare to term1.

#### Returns

Туре	Description
boolean	True if the terms are the same and false if they are not.

# UNBOUND

This function is like the SPARQL UNDEF keyword but extends that functionality to enable users to include an undefined value as a function argument, as UNDEF is only supported in VALUES clauses.

#### Syntax

UNBOUND()

#### Example

The following example statement incorporates UNBOUND to return null if the specified condition (2x > 5) fails:

BIND(IF(?x > 5 , "Win", UNBOUND())) as ?testResult)

In this case, ?testResult is bound if ?x is greater than 5. If ?x is not greater than 5, ?testResult is not bound.

Туре	Description
RDF term	The specified result according to the evaluation of the condition.

# **Informational or Testing Functions**

This topic describes the functions in Anzo that retrieve information from your values and let you ask questions about them or test whether the values match expectations.

#### **Typographical Conventions**

This documentation uses the following conventions in function syntax:

- CAPS: Although SPARQL is case-insensitive, function names and other keywords are written in uppercase for readability.
- [ argument ]: Brackets are used to indicate optional arguments. Arguments without brackets are required.

#### Functions

- CONTAINS: Evaluates whether the specified string contains the given pattern.
- GE: Evaluates whether one value is greater than or equal to (>=) another value.
- GT: Evaluates whether one value is greater than (>) another value.
- ISBLANK: Evaluates whether the given RDF term is a blank node.
- ISDATATYPE: Evaluates whether the given literal value is typed as the specified data type.
- ISERROR: Tests whether the given RDF term evaluates to an error. Only valid in queries against the Anzo System Datasource or other volume.
- ISIRI: Evaluates whether the given RDF term is an IRI.
- ISLITERAL: Evaluates whether the given RDF term is a literal value.
- ISNUMERIC: Evaluates whether the given RDF term is a numeric literal value.
- ISURI: Evaluates whether the given RDF term is a URI.
- LANG: Returns any language tags that are included with strings.
- LANGMATCHES: Evaluates whether a string includes a language tag that matches the specified language range.

- LE: Evaluates whether one value is less than or equal to (<=) another value.
- LOCALNAME: Retrieves the local name from the given URI.
- LT: Evaluates whether one value is less than (<) another value.
- METADATAGRAPHURI: Returns the metadata graph URI for the given URI.
- NAMESPACE: Retrieves the namespace for the specified URI.
- SAMETERM: Evaluates whether two RDF term type values are the same.

### CONTAINS

This function evaluates whether the specified strings contain the given pattern. Results are grouped under "true" or "false."

#### **Syntax**

CONTAINS(text, pattern)			
Argument	Туре	Description	
text	string	The string value that you want to check against the specified pattern.	
pattern	string	The string pattern that you want to look for in the supplied text.	

#### Returns

Туре	Description
boolean	True if the strings contain the pattern and false if they do not.

#### GE

This function evaluates whether value1 is greater than or equal to (>=) value2.

#### Syntax

GE(value1, value2)

Argument	Туре	Description
value1	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value2. This is the value that will be checked to see if it is greater than or equal to value2.
value2	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value1.

Туре	Description
boolean	True if value1 >= value2. False if not.

# GT

This function evaluates whether value1 is greater than (>) value2.

# Syntax

GE(value1, value2)

Argument	Туре	Description
value1	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value2. This is the value that will be checked to see if it is greater than value2.
value2	numeric, boolean, dateTime,	The number, boolean, datetime, literal, URI, or blank node value to compare to value1.

Argument	Туре	Description
	RDF term	

Туре	Description
boolean	True if value1 > value2. False if not.

#### **ISBLANK**

This function evaluates whether the given RDF term value is a blank node. It returns true if it is a blank node or false if it is not.

#### **Syntax**

ISBLANK (value)		
Argument	Туре	Description
value	RDF term	The literal, URI, or blank node value to test and determine if it is a blank node.

#### Returns

Туре	Description
boolean	True if the term is a blank node and false if it is not.

# **ISDATATYPE**

This function evaluates whether the given literal value is typed as the specified data type. It returns true if the value is typed as the provided type or false if it is not.

#### Syntax

ISDATATYPE(value, datatype\_uri)

Argument	Туре	Description
value	literal	The literal value that you want to test against the datatype_uri.
datatype_uri	URI	The URI for the data type that you want to test the value against.

Туре	Description
boolean	True if the literal is typed as specified data type and false if it is not.

# **ISERROR**

This function tests whether the given RDF term evaluates to an error. It returns true if it does evaluate to an error or false if it does not.

#### Note

The ISERROR function is only for use in queries that are run against the Anzo System Datasource or other volume. It is invalid for graphmart queries.

### Syntax

ISERROR(term)

Argument	Туре	Description
term	RDF term	The literal, URI, or blank node value to evaluate for an error.

Туре	Description
boolean	True if the term evaluates to an error and false if it does not.

# ISIRI

This function evaluates whether the given RDF term type value is an IRI. It returns true if the value is an IRI or false if it is not.

### Syntax

ISIRI(term)		
Argument	Туре	Description
term	RDF term	The literal, URI, or blank node value to evaluate whether it is an IRI.

#### Returns

Туре	Description
boolean	True if the term is an IRI and false if it is not.

# **ISLITERAL**

This function evaluates whether the given RDF term type value is a literal value. It returns true if the value is a literal or false if it is not.

### Syntax

ISLITERAL(term)

Argument	Туре	Description
term	RDF term	The literal, URI, or blank node value to evaluate whether it is a literal.

Туре	Description
boolean	True if the term is a literal value and false if it is not.

# ISNUMERIC

This function evaluates whether the given RDF term type value is a numeric literal. It returns true if the value is a numeric literal or false if it is not.

# Syntax

ISNUMERIC(term)		
Argument	Туре	Description
term	RDF term	The literal, URI, or blank node value to evaluate whether it is a numeric literal.

#### Returns

Туре	Description
boolean	True if the term is a numeric literal and false if it is not.

# ISURI

This function evaluates whether the given RDF term type value is a URI. It returns true if the value is a URI or false if it is not.

### Syntax

ISURI(term)		
Argument	Туре	Description
term	RDF term	The literal, URI, or blank node value to evaluate whether it is a URI.

Туре	Description
boolean	True if the term is a URI and false if it is not.

# LANG

This function returns any language tags that are included in the string. The results are grouped by each language tag or by "blank" if a value does not have a language tag.

### Syntax

LANG(text)		
Argument	Туре	Description
text	string	The string to search for language tags.

#### Returns

Туре	Description
string	The found language tags.

# LANGMATCHES

This function tests whether a string includes a language tag that matches the specified language range.

### Syntax

```
LANGMATCHES(text, language_range)
```

Argument	Туре	Description
text	string	The string to evaluate.
language_range	string	The language tag to match in the text.

### Example

LANGMATCHES (LANG(?prop), "en")

Туре	Description
boolean	True if strings include a language tag that matches the range and false if they do not.

# LE

This function evaluates whether value1 is less than or equal to (<=) value2.

#### Syntax

```
LE(value1, value2)
```

Argument	Туре	Description
value1	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value2. This is the value that will be evaluated to see if it is less than or equal to value2.
value2	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, literal, URI, or blank node value to compare to value1.

#### Returns

Туре	Description	
boolean	True if value1 <= value2. False if not.	

# LOCALNAME

This function retrieves the local name from the given URI.

# Syntax

LOCALNAME (uri)		
Argument	Туре	Description
uri	URI	The URI from which to retrieve the local name.

#### Returns

Туре	Description
string	The local name.

# LT

This function evaluates whether value1 is less than (<) value2.

# Syntax

LT(value1, value2)

Argument	Туре	Description
value1	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, or RDF term type value to compare to value2. This is the value that will be evaluated to see if it is less than value2.
value2	numeric, boolean, dateTime, RDF term	The number, boolean, datetime, or RDF term type value to compare to value1.

Туре	Description
boolean	True if value1 < value2. False if not.

# **METADATAGRAPHURI**

This function returns the metadata graph URI for the specified URI.

# Syntax

METADATAGRAPHU	RI(uri)	
Argument	Туре	Description
uri	URI	The URI for which you want to return the corresponding metadata graph URI.

#### Returns

Туре	Description
URI	The metadata graph URI.

# NAMESPACE

This function retrieves the namespace for the given URI.

NAMESPACE(uri)		
Argument	Туре	Description
uri	URI	The URI from which to retrieve the namespace.
Туре	Description	
--------	----------------	
string	The namespace.	

## SAMETERM

This function evaluates whether two RDF term type values are the same.

## Syntax

SAMETERM	(term1,	term2)
----------	---------	--------

Argument	Туре	Description
term1	RDF term	The first literal, URI, or blank node value to compare.
term2	RDF term	The literal, URI, or blank node value to compare to term1.

#### Returns

Туре	Description
boolean	True if the terms are the same and false if they are not.

## **Hash Functions**

This topic describes the hash functions in Anzo.

**Typographical Conventions** 

This documentation uses the following conventions in function syntax:

- CAPS: Although SPARQL is case-insensitive, function names and other keywords are written in uppercase for readability.
- [ argument ]: Brackets are used to indicate optional arguments. Arguments without brackets are required.

#### Functions

- MD5: Returns the MD5 checksum of a string as a hexadecimal string.
- SHA1: Calculates the SHA-1 digest of a string value.
- SHA224: Calculates the SHA-224 digest of a string value.
- SHA256: Calculates the SHA-256 digest of a string value.
- SHA384: Calculates the SHA-384 digest of a string value.
- SHA512: Calculates the SHA-512 digest of a string value.

### MD5

This function returns the MD5 checksum of a string as a hexadecimal string.

MD5(text)		
Argument	Туре	Description
text	string	The string for which to return the MD5 checksum.

Туре	Description
string	The hexadecimal string.

## SHA1

This function calculates the SHA-1 digest of a string.

## Syntax

SHA1(text)

Argument	Туре	Description
text	string	The string for which to calculate the SHA-1 digest.

#### Returns

Туре	Description
string	The SHA-1 digest.

## SHA224

This function calculates the SHA-224 digest of a string.

SHA224 (text)		
Argument	Туре	Description
text	string	The string for which to calculate the SHA-224 digest.

Туре	Description
string	The SHA-224 digest.

## **SHA256**

This function calculates the SHA-256 digest of a string.

## Syntax

SHA256(text)

Argument	Туре	Description
text	string	The string for which to calculate the SHA-256 digest.

#### Returns

Туре	Description
string	The SHA-256 digest.

## SHA384

This function calculates the SHA-384 digest of a string.

SHA384(text)		
Argument	Туре	Description
text	string	The string for which to calculate the SHA-384 digest.

Туре	Description
string	The SHA-384 digest.

## SHA512

This function calculates the SHA-512 digest of a string.

## Syntax

SHA512(text)

Argument	Туре	Description
text	string	The string for which to calculate the SHA-512 digest.

#### Returns

Туре	Description
string	The SHA-512 digest.

## Window Aggregate and Ranking Functions

Window aggregates operate on a particular partition or window of the result set. Unlike grouped aggregate functions that group the result set and return a single row, window aggregates retain the resulting rows and return a value for each row. For example, using the grouped aggregate SUM function to add the total number of tickets sold in a year returns one value: the total number of tickets sold for the year. By using WINDOW\_SUM, the results could be partitioned by month so that the query returns 12 values: the sum of the number of tickets sold in each month of the year. This topic describes the window aggregate functions in Anzo.

#### **Typographical Conventions**

This documentation uses the following conventions in function syntax:

- CAPS: Although SPARQL is case-insensitive, function names and other keywords are written in uppercase for readability.
- [ argument ]: Brackets are used to indicate optional arguments. Arguments without brackets are required.

#### Functions

- WINDOW\_AVG: Calculates the average value of each group of values.
- WINDOW\_COUNT: Counts the number of values in each group of values.
- WINDOW\_MAX: Calculates the maximum value of each group of values.
- WINDOW\_MIN: Calculates the minimum value of each group of values.
- WINDOW\_NTILE: Divides the rows in the partition into the specified number of ranked groups and returns the group that each value belongs to.
- WINDOW\_PERCENTILE: Divides the rows in the partition into 100 ranked groups and returns the group that each value belongs to.
- WINDOW\_PERCENTILE\_CONT: Calculates a percentile based on the continuous distribution of the specified groups of values.

- WINDOW\_PERCENTILE\_DISC: Calculates a percentile based on the discrete distribution of the specified groups of values.
- WINDOW\_PRODUCT: Calculates the product of each group of values.
- WINDOW\_QUARTILE: Divides the rows in the partition into four ranked groups and returns the group that each value belongs to.
- WINDOW\_SUM: Calculates the sum of each group of values.

## WINDOW\_AVG

This function calculates the average value of each group of values.

```
WINDOW_AVG(value [, partition_over ] [, order_by ] [, order ]
      [, start_frame_type ] [, start_frame_value ]
      [, end_frame_type ] [, end_frame_value ])
```

Argument	Туре	Description
value	numeric	Required argument that defines the groups of values to operate on.
partition_ over	variable	Optional argument that partitions the results into groups of rows. If you do not include <code>partition_over</code> , the partition becomes the entire set identified by <code>value</code> .
order_by	variable	Optional argument that defines the order or sequence of rows within each partition.
order	boolean	Optional argument that controls whether the order is ascending or descending. When true, the order is ascending. When false, the order is descending.
start_frame_ type	string	When order_by is specified, the optional start_frame_type, start_frame_value, end_frame_type, and end_frame_

Argument	Туре	Description
		value arguments can be included to refine the set of rows to include in the partitions or groups.
		The start_frame_type argument defines the starting row of the partition and can be one of the following values:
		CURRENT ROW
		UNBOUNDED PRECEDING
		PRECEDING start_frame_value
		<ul> <li>FOLLOWING start_frame_value</li> </ul>
start_frame_ value	int	<b>Optional argument that specifies the starting row based on the</b> <pre>start_frame_type value.</pre>
end_frame_ type	string	This argument defines the ending row of the partition and can be one of the following values:
		CURRENT ROW
		UNBOUNDED FOLLOWING
		<ul> <li>PRECEDING end_frame_value</li> </ul>
		<ul> <li>FOLLOWING end_frame_value</li> </ul>
end_frame_ value	int	<b>Optional argument that specifies the ending row based on the</b> end_frame_type value.

\_

Туре	Description
int	The average values.

## WINDOW\_COUNT

This function counts the number of values in each group of values.

```
WINDOW_COUNT(value [, partition_over ] [, order_by ] [, order ]
      [, start_frame_type ] [, start_frame_value ]
      [, end_frame_type ] [, end_frame_value ])
```

Argument	Туре	Description
value	numeric	Required argument that defines the groups of values to operate on.
partition_ over	variable	Optional argument that partitions the results into groups of rows. If you do not include <code>partition_over</code> , the partition becomes the entire set identified by <code>value</code> .
order_by	variable	Optional argument that defines the order or sequence of rows within each partition.
order	boolean	Optional argument that controls whether the order is ascending or descending. When true, the order is ascending. When false, the order is descending.
start_frame_ type	string	When order_by is specified, the optional start_frame_type, start_frame_value, end_frame_type, and end_frame_value arguments can be included to refine the set of rows to include in the partitions or groups.
		The start_frame_type argument defines the starting row of the partition and can be one of the following values:
		CURRENT ROW
		UNBOUNDED PRECEDING
		<ul> <li>PRECEDING start_frame_value</li> </ul>

Argument	Туре	Description
		FOLLOWING start_frame_value
start_frame_ value	int	<b>Optional argument that specifies the starting row based on the</b> <pre>start_frame_type value.</pre>
end_frame_ type	string	<ul> <li>This argument defines the ending row of the partition and can be one of the following values:</li> <li>CURRENT ROW</li> <li>UNBOUNDED FOLLOWING</li> <li>PRECEDING end_frame_value</li> <li>FOLLOWING end_frame_value</li> </ul>
end_frame_ value	int	<b>Optional argument that specifies the ending row based on the</b> end_frame_type <b>value</b> .

Туре	Description
int	The counts of values.

### WINDOW\_MAX

This function calculates the maximum value of each group of values.

```
WINDOW_MAX(value [, partition_over ] [, order_by ] [, order ]
        [, start_frame ] [, start_frame_type ] [, start_frame_value ]
        [, end_frame_type ] [, end_frame_value ])
```

Argument	Туре	Description
value	numeric	Required argument that defines the groups of values to operate on.
partition_ over	variable	Optional argument that partitions the results into groups of rows. If you do not include <code>partition_over</code> , the partition becomes the entire set identified by <code>value</code> .
order_by	variable	Optional argument that defines the order or sequence of rows within each partition.
order	boolean	Optional argument that controls whether the order is ascending or descending. When true, the order is ascending. When false, the order is descending.
start_frame_ type	string	When order_by is specified, the optional start_frame_type, start_frame_value, end_frame_type, and end_frame_value arguments can be included to refine the set of rows to include in the partitions or groups.
		The start_frame_type argument defines the starting row of the partition and can be one of the following values:
		CURRENT ROW
		UNBOUNDED PRECEDING
		PRECEDING start_frame_value
		<ul> <li>FOLLOWING start_frame_value</li> </ul>
start_frame_ value	int	<b>Optional argument that specifies the starting row based on the</b> <pre>start_frame_type value.</pre>
end_frame_ type	string	This argument defines the ending row of the partition and can be one of the following values:

Argument	Туре	Description
		<ul> <li>CURRENT ROW</li> <li>UNBOUNDED FOLLOWING</li> <li>PRECEDING end_frame_value</li> <li>FOLLOWING end_frame_value</li> </ul>
end_frame_ value	int	<b>Optional argument that specifies the ending row based on the</b> end_frame_type value.

Туре	Description
int	The maximum values.

## WINDOW\_MIN

This function calculates the minimum value of each group of values.

```
WINDOW_MIN(value [, partition_over ] [, order_by ] [, order ]
      [, start_frame_type ] [, start_frame_value ]
      [, end_frame_type ] [, end_frame_value ])
```

Argument	Туре	Description
value	numeric	Required argument that defines the groups of values to operate on.
partition_ over	variable	Optional argument that partitions the results into groups of rows. If you do not include <code>partition_over</code> , the partition becomes the entire set identified by <code>value</code> .
order_by	variable	Optional argument that defines the order or sequence of rows

Argument	Туре	Description
		within each partition.
order	boolean	Optional argument that controls whether the order is ascending or descending. When true, the order is ascending. When false, the order is descending.
start_frame_ type	string	When order_by is specified, the optional start_frame_type, start_frame_value, end_frame_type, and end_frame_value arguments can be included to refine the set of rows to include in the partitions or groups.
		The start_frame_type argument defines the starting row of the partition and can be one of the following values:
		CURRENT ROW
		UNBOUNDED PRECEDING
		PRECEDING start_frame_value
		<ul> <li>FOLLOWING start_frame_value</li> </ul>
start_frame_ value	int	<b>Optional argument that specifies the starting row based on the</b> start_frame_type value.
end_frame_ type	string	This argument defines the ending row of the partition and can be one of the following values:
		CURRENT ROW
		UNBOUNDED FOLLOWING
		<ul> <li>PRECEDING end_frame_value</li> </ul>
		<ul> <li>FOLLOWING end_frame_value</li> </ul>

Argument	Туре	Description
end_frame_ value	int	<b>Optional argument that specifies the ending row based on the</b> end_frame_type value.

Туре	Description
int	The minimum values.

## WINDOW\_NTILE

This function divides the rows in the partition into the specified number of ranked groups and returns the group that each value belongs to.

#### Syntax

WINDOW\_NTILE(ntile, value, order\_by [, partition\_over ])

Argument	Туре	Description
ntile	int	Required argument that specifies the number of ranking groups.
value	numeric	Required argument that defines the groups of values to operate on.
order_by	variable	Required argument that defines the order or sequence of rows within each partition.
partition_ over	variable	Optional argument that partitions the results into groups of rows. If you do not include <code>partition_over</code> , the partition becomes the entire set identified by <code>value</code> .

Туре	Description
int	The group that the values belong to.

## WINDOW\_PERCENTILE

This function divides the rows in the partition into 100 ranked groups and returns the group that each value belongs to.

#### **Syntax**

```
WINDOW PERCENTILE (value, order by [, partition over ])
```

Argument	Туре	Description
value	numeric	Required argument that defines the groups of values to operate on.
order_by	variable	Required argument that defines the order or sequence of rows within each partition.
partition_ over	variable	Optional argument that partitions the results into groups of rows. If you do not include partition_over, the partition becomes the entire set identified by value.

#### Returns

Туре	Description
int	The group that the values belong to.

## WINDOW\_PERCENTILE\_CONT

This function calculates a percentile based on the continuous distribution of the specified groups of values. The returned value is interpolated and may not be equal to any of the values in the group.

#### Syntax

WINDOW\_PERCENTILE\_CONT(percentile, value, order\_by [, partition\_over ])

Argument	Туре	Description
percentile	int	Required argument that specifies the percentile for the calculation.
value	numeric	Required argument that defines the groups of values to operate on.
order_by	variable	Required argument that defines the order or sequence of rows within each partition.
partition_ over	variable	Optional argument that partitions the results into groups of rows. If you do not include <code>partition_over</code> , the partition becomes the entire set identified by <code>value</code> .

#### Returns

Туре	Description
int	The interpolated percentiles.

## WINDOW\_PERCENTILE\_DISC

This function calculates a percentile based on the discrete distribution of the specified groups of values.

```
WINDOW_PERCENTILE_DISC(percentile, value, order_by [, partition_over ])
```

Argument	Туре	Description
percentile	int	Required argument that specifies the percentile for the calculation.

Argument	Туре	Description
value	numeric	Required argument that defines the groups of values to operate on.
order_by	variable	Required argument that defines the order or sequence of rows within each partition.
partition_ over	variable	Optional argument that partitions the results into groups of rows. If you do not include <code>partition_over</code> , the partition becomes the entire set identified by <code>value</code> .

Туре	Description	
int	The percentiles based on the discrete distribution of the groups.	

## WINDOW\_PRODUCT

This function calculates the product of each group of values.

```
WINDOW_PRODUCT(value [, partition_over ] [, order_by ] [, order ]
       [, start_frame_type ] [, start_frame_value ]
       [, end_frame_type ] [, end_frame_value ])
```

Argument	Туре	Description
value	numeric	Required argument that defines the groups of values to operate on.
partition_ over	variable	Optional argument that partitions the results into groups of rows. If you do not include <code>partition_over</code> , the partition becomes the entire set identified by <code>value</code> .
order_by	variable	Optional argument that defines the order or sequence of rows

Argument	Туре	Description
		within each partition.
order	boolean	Optional argument that controls whether the order is ascending or descending. When true, the order is ascending. When false, the order is descending.
start_frame_ type	string	When order_by is specified, the optional start_frame_type, start_frame_value, end_frame_type, and end_frame_value arguments can be included to refine the set of rows to include in the partitions or groups.
		The start_frame_type argument defines the starting row of the partition and can be one of the following values:
		CURRENT ROW
		UNBOUNDED PRECEDING
		PRECEDING start_frame_value
		<ul> <li>FOLLOWING start_frame_value</li> </ul>
start_frame_ value	int	<b>Optional argument that specifies the starting row based on the</b> start_frame_type value.
end_frame_ type	string	This argument defines the ending row of the partition and can be one of the following values:
		CURRENT ROW
		UNBOUNDED FOLLOWING
		<ul> <li>PRECEDING end_frame_value</li> </ul>
		<ul> <li>FOLLOWING end_frame_value</li> </ul>

Argument	Туре	Description
end_frame_ value	int	Optional argument that specifies the ending row based on the end_frame_type value.

Туре	Description
int	The products of the groups.

## WINDOW\_QUARTILE

This function divides the rows in the partition into four ranked groups and returns the group that each value belongs to.

#### **Syntax**

WINDOW\_QUARTILE(value, order\_by [, partition\_over ])

Argument	Туре	Description
value	numeric	Required argument that defines the groups of values to operate on.
order_by	variable	Required argument that defines the order or sequence of rows within each partition.
partition_ over	variable	Optional argument that partitions the results into groups of rows. If you do not include partition_over, the partition becomes the entire set identified by value.

#### Returns

Туре	Description
int	The group that the values belong to.

## WINDOW\_SUM

This function calculates the sum of each group of values.

```
WINDOW_SUM(value [, partition_over ] [, order_by ] [, order ]
      [, start_frame_type ] [, start_frame_value ]
      [, end_frame_type ] [, end_frame_value ])
```

Argument	Туре	Description
value	numeric	Required argument that defines the groups of values to operate on.
partition_ over	variable	Optional argument that partitions the results into groups of rows. If you do not include <code>partition_over</code> , the partition becomes the entire set identified by <code>value</code> .
order_by	variable	Optional argument that defines the order or sequence of rows within each partition.
order	boolean	Optional argument that controls whether the order is ascending or descending. When true, the order is ascending. When false, the order is descending.
start_frame_ type	string	When order_by is specified, the optional start_frame_type, start_frame_value, end_frame_type, and end_frame_value arguments can be included to refine the set of rows to include in the partitions or groups.
		The start_frame_type argument defines the starting row of the partition and can be one of the following values:
		CURRENT ROW
		UNBOUNDED PRECEDING
		PRECEDING start_frame_value

Argument	Туре	Description
		<ul> <li>FOLLOWING start_frame_value</li> </ul>
start_frame_ value	int	<b>Optional argument that specifies the starting row based on the</b> <pre>start_frame_type value.</pre>
end_frame_ type	string	<ul> <li>This argument defines the ending row of the partition and can be one of the following values:</li> <li>CURRENT ROW</li> <li>UNBOUNDED FOLLOWING</li> <li>PRECEDING end_frame_value</li> <li>FOLLOWING end_frame_value</li> </ul>
end_frame_ value	int	<b>Optional argument that specifies the ending row based on the</b> end_frame_type <b>value</b> .

Туре	Description
int	The sum of each group.

#### Example

The example below first creates data by running the following INSERT DATA query in a graphmart Query Step:

```
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX ex: <http://example.com/csi#>
INSERT DATA {
    GRAPH ${targetGraph} {
    ex:sale1 ex:date "2014-05-23T10:20:13"^^xsd:dateTime ; ex:volume 15 .
    ex:sale2 ex:date "2014-06-23T10:20:13"^^xsd:dateTime ; ex:volume 3 .
    ex:sale3 ex:date "2014-06-23T10:22:13"^^xsd:dateTime ; ex:volume 35 .
```

```
ex:sale4 ex:date "2014-07-23T10:20:13"^^xsd:dateTime ; ex:volume 66 .
ex:sale5 ex:date "2014-09-23T10:20:13"^^xsd:dateTime ; ex:volume 19 .
ex:sale6 ex:date "2014-11-23T10:20:13"^^xsd:dateTime ; ex:volume 33 .
ex:sale7 ex:date "2014-12-23T10:20:13"^^xsd:dateTime ; ex:volume 12 .
}
```

The following query against the new data uses the WINDOW\_SUM function to return the total volume of sales for each month:

```
PREFIX ex: <http://example.com/csi#>
SELECT DISTINCT ?Month ?Total_Volume
WHERE {
    {
        SELECT
        ?Month (WINDOW_SUM(?o, ?Month) AS ?Total_Volume)
        WHERE {
            ?s ex:volume ?o .
            ?s ex:date ?date .
            BIND (MONTH(?date) AS ?Month)
        }
    }
}
```

The query returns the following results:

Month | Total\_Volume -----| 15 5 6 38 7 66 19 9 11 33 12 12

# Develop

The topics in this section provide information for developers.

In this section:	
Anzo Rest API	1176
Anzo Java SDK	1233

# Anzo Rest API

Anzo includes a REST API that you can use when developing applications or automation tasks. The API supports create, read, update, and delete operations on datasets, graphmarts, data layers, and steps. The interface also supports read and update operations for artifact access control lists, upload, download, and delete operations for models, run and cancel operations for unstructured pipelines, import and export operations for migration packages, and read operations for dynamic launch configurations and static AnzoGraph instances.

The topics in this section give an overview of the API, list the options for viewing the detailed documentation, provide instructions for enabling CORS, and describe the schemas for the data layer step types.

In this section:	
Introduction to the API	1177
Viewing the API Documentation	1180
Enabling Cross-Origin Resource Sharing	1181
Step Type Schemas	1183

## Introduction to the API

The API supports create, read, update, and delete operations on datasets, graphmarts, data layers, and steps. The interface also supports read and update operations for artifact access control lists, upload, download, and delete operations for models, run and cancel operations for unstructured pipelines, import and export operations for migration packages, and read operations for dynamic launch configurations and static AnzoGraph instances.

This topic gives a summary of the API request URL, endpoints and methods, URI encoding requirements, and error handling.

- Request URL
- Endpoints and Methods
- URI Encoding Requirements
- Error Handling

#### **Request URL**

The path in the URL that you use to access the REST API endpoints (as well as the documentation on your Anzo server) differs depending on the version of Anzo that is installed:

- In 5.4.1 releases, the URL is https://<hostname>/api/v1/<endpoint>.
- In 5.4.2+ releases, the URL is https://<hostname>/api/<endpoint>.

The v1 in the path was removed in 5.4.2. The URL to access the documentation on your Anzo server was also changed from https://<hostname>/api/v1/docs/index.html in 5.4.1 to https://<hostname>/api/docs/index.html in 5.4.2.

#### **Endpoints and Methods**

There is an endpoint for each type of Anzo artifact or object (dataset, edition, model, graphmart, layer, step, etc.). Each endpoint supports a subset or all of the following methods:

- POST or PUT for create
- GET for reads
- PATCH for updates
- DELETE for delete

## **URI Encoding Requirements**

When including a URI (e.g., a graphmart, layer, or step URI) in a request URL, the URI must be URL-encoded. The following example shows a layer URI,

http://cambridgesemantics.com/Layer/858c521bc7d84364a5a2112e38dc0b52, that has been URL-encoded:

http%3A%2F%2Fcambridgesemantics.com%2FLayer%2F858c521bc7d84364a5a2112e38dc0b52

The encoded value should be used in the request URL. For example, the following request retrieves the status of the layer:

```
http://10.100.0.10:8080/api/layers/http%3A%2F%2Fcambridgesemantics.com%2FLayer%2F858c52 lbc7d84364a5a2112e38dc0b52/status
```

#### Note

URIs that appear in the body of a request do not need to be URL-encoded.

An error message such as the one below indicates that a request URL included a URI that was not URL-encoded:

```
{
   "summary": "Rest API Error: No handler found for GET
/api/graphmarts/http://cambridgesemantics.com/Graphmart/811ece67d61e436cb128a929797b68d
f",
   "detail": "No handler found for GET
/api/graphmarts/http://cambridgesemantics.com/Graphmart/811ece67d61e436cb128a929797b68d
f"
}
```

To resolve the error, encode the URI as a URL and then resend the request.

## **Error Handling**

Errors returned from the API contain a summary and detailed message. Stack traces from the server are not included in API responses. Stack traces can be obtained via the server logs. The documentation describes the error codes and response details.

## Viewing the API Documentation

An Open API definition is used to generate the API documentation. There are two ways to view the document:

• Access the document on your Anzo server at the following URL:

https://<hostname>/api/docs/index.html

#### Note

In releases prior to 5.4.2, the URL to view the documentation on your server is https://<hostname>/api/v1/docs/index.html.

Where <hostname> is the Anzo server DNS name or IP address. When you access the API documentation on your server, a **Try It Out** button is available for each request. Clicking Try It Out opens a request body that you can edit and execute against the server.

 Access the document on the Cambridge Semantics documentation website at the following URL:

https://docs.cambridgesemantics.com/anzo/v5.4/api/

This version of the documentation is static and the Try It Out functionality is disabled. However, it remains available if you are unable to access the Anzo server.

## **Enabling Cross-Origin Resource Sharing**

By default, cross-origin resource sharing (CORS) is disabled for the Anzo REST API service. If you plan to access the API from a web client and need to enable cross-origin requests, follow the steps below.

- In the Administration application, expand the Servers menu and click Advanced Configuration. Click I understand and accept the risk.
- 2. Search for the Anzo REST API bundle and view its details.
- 3. Click the Services tab and expand Anzo REST API.
- 4. Locate the **org.openanzo.servlet.allowCrossOriginAccess** property (shown in the image below).

A	nzo REST API	Anzo Swagger J	SONLD API	STARTED		^
0	sgi://14BF-24AB-7402-46E0-B936-397F/c	onfiguration/instar	nce/config/com.cambridg	jesemantics.anzo.api.rest	ADD PROF	PERTY
	org.openanzo.services.enab	led				
	org.openanzo.servlet.authorizationType BASIC				$\checkmark X$	
	org.openanzo.servlet.allowCrossOriginAccess false					Î
	org.openanzo.servlet.pathSpec /api-anon/*					Î
	org.openanzo.servlet.protectedPathRoles http://openanzo.org/Role/authentic	atedUsers				Î
	org.openanzo.servlet.protectedPathSpec /api/*					Î

5. Click the property to make it editable, and then change false to true.

org.openanzo.servlet.allowCrossOriginAccess
true

- 6. Click the checkmark icon ( $\checkmark$ ) for that property to save the change.
- 7. Restart Anzo to apply the configuration change.

The Anzo API service is now configured to allow cross-origin requests. As an example of a CORS request, the following snippet from a React application reloads a graphmart via an API call. The request retrieves an authorization token and configures an Axios request object with the token and CORS-related headers:

```
export default async function reloadAnzo() {
  const authorization = getAnzoAuth(); //basicAuth
  console.log('Authorization', authorization);
  const config: AxiosRequestConfig = {
   headers: {
      Authorization: authorization,
      'Access-Control-Allow-Origin': '*',
      'Access-Control-Allow-Credentials': true,
      'Access-Control-Allow-Methods': 'POST, OPTIONS',
    },
  };
  const url =
    'https://{{AnzoURL}}/api/v1/graphmarts/{{GraphmartIRI}}/reload';
  try {
    const success = await axios.post(url, config);
   console.log('return', success);
   return success;
  } catch (error) {
    console.log(error);
   return {error};
  }
```

 $\checkmark \times$ 

## Step Type Schemas

This section provides reference information for each type of step that can be created or updated via the Anzo REST API. For each step type in the list below, there is a JSON request that includes all of the step's body parameters, excluding the read-only options. Below the request is a description of the step's schema, including the read-only parameters.

## **Direct Load Step**

This type of step loads data directly from an external source.

#### **JSON Request**

The following template shows the body of a JSON request that could be used in a Direct Load Step PUT or PATCH request. It lists all of the step's required and optional body parameters but excludes the read-only options. The default values for each parameter are shown. Below the request (in Schema Details) is a table that describes the complete schema, including the read-only parameters. Clicking a link in the template takes you to the schema details for that parameter.

```
{
  "title" : "string",
  "transformQuery" : "string",
  "incrementalData" : [ "string" ],
  "type" : "DirectLoadStep",
  "enabled" : true,
  "contextProvider" : [ "string" ],
  "description" : "string",
  "ontology" : [ "string" ],
  "source" : [ "string" ],
  "ignoreLoadErrors" : true,
  "disableLoadCounts" : true,
  "preGenerateStatistics" : true,
  "tags" : [ {
                "description" : "string",
                "title" : "string"
            }],
  "tagTitle" : [ "string" ]
}
```

#### **Schema Details**

The table below describes the Direct Load Step schema.

#### Tip

You can also see the Direct Load Step schema by expanding **Schemas** at the bottom of the Anzo REST API document and viewing **DirectLoadStep**.

Property	Format	Require- d?	Description
<b>uri</b> (read-only)	"uri"	Auto- generat ed	The URI of the step.
<b>creator</b> (read- only)	"uri"	Auto- generat ed	The creator of the step.

Property	Format	Require- d?	Description
<b>created</b> (read- only)	"dateTi me"	Auto- generat ed	The timestamp when the step was created.
<b>modifier</b> (read- only)	"uri"	Auto- generat ed	The user who modified the step.
<b>alltypes</b> (read- only)	Array of strings	Optional	A list of the types related to the step, such as DirectLoadStep, Step, LayerChild, etc.
<b>contextAttribut</b> <b>e</b> (read-only)	Array of strings	Optional	A list of any context attributes that are used.
title	"string"	Require d	The name of the step.
transformQuer y	"string"	Require d	The SPARQL query to run.
incrementalDat a	"string"	Optional	Incremental load data associated with the step.
type	"string"	Require d	The type of step: "DirectLoadStep".
enabled	boolean	Optional	Controls whether the step is enabled or disabled.
contextProvide r	[ "uri", " ." ]	Optional	A list of any referenced context providers (the data source URI). You can retrieve data for the parent layer to get a list of providers for that layer.

Property	Format	Require- d?	Description
description	"string"	Optional	A brief description of the step.
ontology	[ "uri", " ." ]	Optional	This type of step automatically generates a managed model that is owned and managed by the data layer that contains the Direct Load Step. If you would like to associate additional models with this step, you can include the URIs for those models. For more information about managed models, see Managed Model Concepts.
SOURCE	["uri", " ."]	Require d	<ul> <li>The source data for the step. Options are any combination of the following values:</li> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#Self": The source is the data that is in this step's layer.</li> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#AllPrevious": The source is the data from all of the successful layers that precede this step's layer. Failed layers are ignored.</li> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#Previous": The source is the data that is in the one layer that precedes this step's layer.</li> <li>"layer_uri": The source is a specific layer in the graphmart.</li> </ul>
ignoreLoadErr ors	boolean	Optional	Controls whether to ignore errors and proceed with the load or fail the step if there is an error.

Property	Format	Require- d?	Description
disableLoadCo unts	boolean	Optional	Controls whether Anzo periodically queries AnzoGraph to count the total number of statements that are processed. Disabling the load count decreases the number of queries that run during activation.
preGenerateSt atistics	boolean	Optional	Controls whether AnzoGraph generates statistics on the data before the step is run.
tags	Array of objects	Optional	Any tags on the step.
tagTitle	["string"]	Optional	A virtual property that is available for all objects. It lists the tags associated with the step or can be used to add a tag to the step without including a description.

## **Elasticsearch Indexing Step**

This type of step creates an Elasticsearch index to associate with a layer.

#### **JSON Request**

The following template shows the body of a JSON request that could be used in an Elasticsearch Indexing Step PUT or PATCH request. It lists all of the step's required and optional body parameters but excludes the read-only options. The default values for each parameter are shown. Below the request (in Schema Details) is a table that describes the complete schema, including the read-only parameters. Clicking a link in the template takes you to the schema details for that parameter.

```
{
    "elasticsearchIndexingQuery" : "string",
    "title" : "string",
    "incrementalData" : [ "string" ],
    "type" : "ElasticsearchIndexingStep",
```

#### **Schema Details**

The table below describes the Elasticsearch Indexing Step schema.

#### Tip

You can also see the Elasticsearch Indexing Step schema by expanding **Schemas** at the bottom of the Anzo REST API document and viewing **ElasticsearchIndexingStep**.

Property	Format	Requir- ed?	Description
<b>uri</b> (read-only)	"uri"	Auto- generat ed	The URI of the step.
<b>creator</b> (read-only)	"uri"	Auto- generat ed	The creator of the step.
created (read-only)	"dateTi me"	Auto- generat ed	The timestamp when the step was created.
Property	Format	Requir- ed?	Description
--	------------------	------------------------	---
modifier (read-only)	"uri"	Auto- generat ed	The user who modified the step.
alltypes (read-only)	Array of strings	Optiona I	A list of the types related to the step, such as ElasticsearchIndexingStep, Step, LayerChild, etc.
<b>contextAttribute</b> (read-only)	Array of strings	Optiona I	A list of any context attributes that are used.
elasticsearchIndex ingQuery	"string"	Require d	The SPARQL query to run for creating the Elasticsearch index.
title	"string"	Require d	The name of the step.
incrementalData	"string"	Optiona I	Incremental load data associated with the step.
type	"string"	Require d	The type of step: "ElasticsearchIndexingStep".
enabled	boolean	Optiona I	Controls whether the step is enabled or disabled.
contextProvider	["uri", ". "]	Optiona I	A list of any referenced context providers (the data source URI). You can retrieve data for the parent layer to get a list of providers for that layer.
description	"string"	Optiona I	A brief description of the step.

Property	Format	Requir- ed?	Description
ontology	[ "uri", ". " ]	Optiona I	A list of any models to associate with this step.
source	[ "uri", ". " ]	Require d	The source data for the step. Options are any combination of the following values:
			<ul> <li>"http://cambridgesemantics.com/ontologies/ Graphmarts#Self": The source is the data that is in this step's layer.</li> </ul>
			<ul> <li>"http://cambridgesemantics.com/ontologies/ Graphmarts#AllPrevious": The source is the data from all of the successful layers that precede this step's layer. Failed layers are ignored.</li> </ul>
			<ul> <li>"http://cambridgesemantics.com/ontologies/ Graphmarts#Previous": The source is the data that is in the one layer that precedes this step's layer.</li> </ul>
			<ul> <li>"layer_uri": The source is a specific layer in the graphmart.</li> </ul>
ignoreLoadErrors	boolean	Optiona I	Controls whether to ignore errors and proceed with the load or fail the step if there is an error.
disableLoadCount s	boolean	Optiona I	Controls whether Anzo periodically queries AnzoGraph to count the total number of statements that are processed. Disabling the load count decreases the number of queries that run during activation.

Property	Format	Requir- ed?	Description
preGenerateStatist ics	boolean	Optiona I	Controls whether AnzoGraph generates statistics on the data before the step is run.
tags	Array of objects	Optiona I	Any tags on the step.
tagTitle	["strin g"]	Optiona I	A virtual property that is available for all objects. It lists the tags associated with the step or can be used to add a tag to the step without including a description.

# **Elasticsearch Snapshot Step**

This type of step creates an Elasticsearch snapshot of the index associated with a layer.

#### **JSON Request**

The following template shows the body of a JSON request that could be used in an Elasticsearch Snapshot Step PUT or PATCH request. It lists all of the step's required and optional body parameters but excludes the read-only options. The default values for each parameter are shown. Below the request (in Schema Details) is a table that describes the complete schema, including the read-only parameters. Clicking a link in the template takes you to the schema details for that parameter.

```
{
   "gmLinkedDataset" : "string",
   "title" : "string",
   "incrementalData" : [ "string" ],
   "type" : "ElasticsearchSnapshotStep",
   "enabled" : true,
   "contextProvider" : [ "string" ],
   "description" : "string",
   "ontology" : [ "string" ],
   "source" : [ "string" ],
```

```
"ignoreLoadErrors" : true,
"disableLoadCounts" : true,
"preGenerateStatistics" : true,
"tags" : [ {
            "description" : "string",
            "title" : "string"
        } ],
"tagTitle" : [ "string" ]
}
```

The table below describes the Elasticsearch Snapshot Step schema.

## Tip

You can also see the Elasticsearch Snapshot Step schema by expanding **Schemas** at the bottom of the Anzo REST API document and viewing **ElasticsearchSnapshotStep**.

Property	Format	Require- d?	Description
<b>uri</b> (read-only)	"uri"	Auto- generat ed	The URI of the step.
<b>creator</b> (read- only)	"uri"	Auto- generat ed	The creator of the step.
<b>created</b> (read- only)	"dateTi me"	Auto- generat ed	The timestamp when the step was created.
<b>modifier</b> (read- only)	"uri"	Auto- generat ed	The user who modified the step.

Property	Format	Require- d?	Description
<b>alltypes</b> (read- only)	Array of strings	Optional	A list of the types related to the step, such as ElasticsearchSnapshotStep, Step, LayerChild, etc.
<b>contextAttribut</b> <b>e</b> (read-only)	Array of strings	Optional	A list of any context attributes that are used.
gmLinkedData set	"uri"	Require d	The URI of the Linked Dataset Catalog entry that represents the target FLDS for the export. To get the catalog entry, you can retrieve data about the dataset and use the catalogEntry value.
title	"string"	Require d	The name of the step.
incrementalDat a	"string"	Optional	Incremental load data associated with the step.
type	"string"	Require d	The type of step: "ElasticsearchSnapshotStep".
enabled	boolean	Optional	Controls whether the step is enabled or disabled.
contextProvide r	[ "uri", " ." ]	Optional	A list of any referenced context providers (the data source URI). You can retrieve data for the parent layer to get a list of providers for that layer.
description	"string"	Optional	A brief description of the step.
ontology	[ "uri", " ." ]	Optional	A list of any models to associate with this step.

Property	Format	Require- d?	Description
source	[ "uri", " ." ]	Require d	The source data for the step. Options are any combination of the following values:
			<ul> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#Self": The source is the data that is in this step's layer.</li> </ul>
			<ul> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#AllPrevious": The source is the data from all of the successful layers that precede this step's layer. Failed layers are ignored.</li> </ul>
			<ul> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#Previous": The source is the data that is in the one layer that precedes this step's layer.</li> </ul>
	<ul> <li>"layer_uri": The source is a specific layer in the graphmart.</li> </ul>		
ignoreLoadErr ors	boolean	Optional	Controls whether to ignore errors and proceed with the load or fail the step if there is an error.
disableLoadCo unts	boolean	Optional	Controls whether Anzo periodically queries AnzoGraph to count the total number of statements that are processed. Disabling the load count decreases the number of queries that run during activation.
preGenerateSt atistics	boolean	Optional	Controls whether AnzoGraph generates statistics on the data before the step is run.

Property	Format	Require- d?	Description
tags	Array of objects	Optional	Any tags on the step.
tagTitle	["string"]	Optional	A virtual property that is available for all objects. It lists the tags associated with the step or can be used to add a tag to the step without including a description.

# **Export Step**

This type of step exports the contents of a graphmart to an FLDS on disk.

### **JSON Request**

The following template shows the body of a JSON request that could be used in an Export Step PUT or PATCH request. It lists all of the step's required and optional body parameters but excludes the read-only options. The default values for each parameter are shown. Below the request (in Schema Details) is a table that describes the complete schema, including the read-only parameters. Clicking a link in the template takes you to the schema details for that parameter.

```
{
  "doNotCreateEditionsOnExport" : true,
  "generateMetrics" : true,
  "exportBinaryStoreContents" : true,
  "overwriteFlds" : true,
  "alwaysMoveBinaryStore" : true,
  "gmLinkedDataset" : "string",
  "edition" : "string",
  "elasticsearchBulkSize" : 0,
  "elasticsearchBulkActions" : 0,
  "elasticsearchBulkMaxThreadsPerFlds" : 0,
  "elasticsearchBulkConcurrentRequests" : 0,
  "keepEsIndexOnline" : true,
  "elasticsearchBulkMaxFldsThreads" : 0,
  "maxComponentsInEdition" : 0,
  "elasticsearchIndexSettings" : "string",
  "exportElasticSearchContents" : true,
```

```
"title" : "string",
  "incrementalData" : [ "string" ],
  "type" : "ExportStep",
  "enabled" : true,
  "contextProvider" : [ "string" ],
  "description" : "string",
  "ontology" : [ "string" ],
  "source" : [ "string" ],
  "ignoreLoadErrors" : true,
  "disableLoadCounts" : true,
  "preGenerateStatistics" : true,
  "tags" : [ {
                "description" : "string",
                "title" : "string"
            }],
  "tagTitle" : [ "string" ]
}
```

The table below describes the Export Step schema.

## Tip

You can also see the Export Step schema by expanding **Schemas** at the bottom of the Anzo REST API document and viewing **ExportStep**.

Property	Forma- t	Requir- ed?	Description
<b>uri</b> (read-only)	"uri"	Auto- genera ted	The URI of the step.
<b>creator</b> (read-only)	"uri"	Auto- genera ted	The creator of the step.
created (read-only)	"dateTi	Auto-	The timestamp when the step was created.

Property	Forma- t	Requir- ed?	Description
	me"	genera ted	
modifier (read-only)	"uri"	Auto- genera ted	The user who modified the step.
alltypes (read-only)	Array of strings	Option al	A list of the types related to the step, such as ExportStep, Step, LayerChild, etc.
<b>contextAttribute</b> (read- only)	Array of strings	Option al	A list of any context attributes that are used.
doNotCreateEditionsOn Export	boolea n	Option al	Controls whether a new edition is created for the dataset each time the step is run.
generateMetrics	boolea n	Option al	Controls whether a data profile is generated before the data is exported. If you load the exported files in the future, the data profile is also loaded.
exportBinaryStoreCont ents	boolea n	Option al	Applies to exports of unstructured graphmarts and controls whether the binary store is exported along with the data.
overwriteFlds	boolea n	Option al	Controls whether the existing FLDS is replaced with the exported files whenever the step is run or whether the exported files are added to the existing FLDS.

Property	Forma- t	Requir- ed?	Description
			When overwriteFlds is true, Anzo archives the existing files in a new timestamped export subdirectory under the FLDS directory each time the step runs. If you add the exported dataset to a graphmart, only the latest version of the data will be loaded. When overwriteFlds is false, Anzo adds all of the exported datasets to a cumulative export directory under the FLDS directory. The dataset will contain the original files as well as all cumulative working editions. If you add this dataset to a graphmart, all of the data from all of the subdirectories will be loaded.
alwaysMoveBinaryStor e	boolea n	Option al	This option also applies to exports of unstructured graphmarts and controls whether the binary store is moved or copied during the export. Since the binary store can be large and have a nested structure, copying the data can take a very long time. Since moving the binary store is almost instantaneous, however, leaving this option set to true can reduce the time it takes to complete the export.
gmLinkedDataset	"uri"	Requir ed	The URI of the Linked Dataset Catalog entry that represents the target FLDS for the export. To get the catalog entry, you can retrieve data about the dataset and use the catalogEntry value.

Property	Forma- t	Requir- ed?	Description
edition	"uri"	Option al	The URI of the edition that will be created on export.
elasticsearchBulkSize	long	Option al	The maximum batch size in MB.
elasticsearchBulkActio ns	int	Option al	The maximum number of documents to include in each batch.
elasticsearchBulkMaxT hreadsPerFlds	int	Option al	The maximum number of threads to use for indexing per FLDS.
elasticsearchBulkConc urrentRequests	int	Option al	The maximum number of bulk requests that can run concurrently.
keepEsIndexOnline	boolea n	Option al	Controls whether the Elasticsearch index remains stored in Elasticsearch or is removed from Elasticsearch once it is exported.
elasticsearchBulkMaxFl dsThreads	int	Option al	The maximum number of FLDSes to index concurrently.
maxComponentsInEditi on	int	Option al	Controls the maximum number of components to retain in an edition. The default value is 0, which means unlimited. If you specify a number in this field and the limit is reached, Anzo ages off the oldest components as new ones are created.
elasticsearchIndexSetti ngs	"string"	Option al	A JSON-formatted list of any Elasticsearch- specific index settings to apply.

Property	Forma- t	Requir- ed?	Description
exportElasticSearchCo ntents	boolea n	Option al	Indicates whether to export any Elasticsearch contents that are included in the graphmart.
title	"string"	Requir ed	The name of the step.
incrementalData	"string"	Option al	Incremental load data associated with the step.
type	"string"	Requir ed	The type of step: "ExportStep".
enabled	boolea n	Option al	Controls whether the step is enabled or disabled.
contextProvider	[ "uri", " " ]	Option al	A list of any referenced context providers (the data source URI). You can retrieve data for the parent layer to get a list of providers for that layer.
description	"string"	Option al	A brief description of the step.
ontology	[ "uri", " " ]	Option al	A list of any models to associate with this step.
source	[ "uri", " " ]	Requir ed	<ul> <li>The source data for the step. Options are any combination of the following values:</li> <li>"http://cambridgesemantics.com/ontolog ies/Graphmarts#Self": The source is the</li> </ul>

Property	Forma- t	Requir- ed?	Description
			<ul> <li>data that is in this step's layer.</li> <li>"http://cambridgesemantics.com/ontolog ies/Graphmarts#AllPrevious": The source is the data from all of the successful layers that precede this step's layer. Failed layers are ignored.</li> <li>"http://cambridgesemantics.com/ontolog ies/Graphmarts#Previous": The source is the data that is in the one layer that precedes this step's layer.</li> <li>"layer_uri": The source is a specific layer in the graphmart.</li> </ul>
ignoreLoadErrors	boolea n	Option al	Controls whether to ignore errors and proceed with the load or fail the step if there is an error.
disableLoadCounts	boolea n	Option al	Controls whether Anzo periodically queries AnzoGraph to count the total number of statements that are processed. Disabling the load count decreases the number of queries that run during activation.
preGenerateStatistics	boolea n	Option al	Controls whether AnzoGraph generates statistics on the data before the step is run.
tags	Array of objects	Option al	Any tags on the step.

Property	Forma- t	Requir- ed?	Description
tagTitle	["strin g"]	Option al	A virtual property that is available for all objects. It lists the tags associated with the step or can be used to add a tag to the step without including a description.

## Load Dataset Step

This type of step loads a system dataset or a dataset that is in the Datasets catalog.

#### **JSON Request**

The following template shows the body of a JSON request that could be used in a Load Dataset Step PUT or PATCH request. It lists all of the step's required and optional body parameters but excludes the read-only options. The default values for each parameter are shown. Below the request (in Schema Details) is a table that describes the complete schema, including the read-only parameters. Clicking a link in the template takes you to the schema details for that parameter.

```
{
  "skipEsLoadIfIndexExists" : true,
  "edition" : "string",
  "gmLinkedDataset" : "string",
  "watchFldsDirectory" : true,
  "maskedPredicate" : [ "string" ],
  "fldsComponent" : [ "string" ],
  "transformQuery" : "string",
  "title" : "string",
  "type" : "LoadDatasetStep",
  "incrementalData" : [ "string" ],
  "enabled" : true,
  "contextProvider" : [ "string" ],
  "description" : "string",
  "ontology" : [ "string" ],
  "ignoreLoadErrors" : true,
  "disableLoadCounts" : true,
  "preGenerateStatistics" : true,
  "tags" : [ {
                "description" : "string",
```

```
"title" : "string"

} ],

"tagTitle" : [ "string" ]

}
```

The table below describes the Load Dataset Step schema.

## Tip

You can also see the Load Dataset Step schema by expanding **Schemas** at the bottom of the Anzo REST API document and viewing **LoadDatasetStep**.

Property	Format	Required?	Description
<b>uri</b> (read-only)	"uri"	Auto- generated	The URI of the step.
<b>creator</b> (read-only)	"uri"	Auto- generated	The creator of the step.
created (read-only)	"dateTime"	Auto- generated	The timestamp when the step was created.
modifier (read-only)	"uri"	Auto- generated	The user who modified the step.
alltypes (read-only)	Array of strings	Optional	A list of the types related to the step, such as LoadDatasetStep, Step, LayerChild, etc.
contextAttribute (read-only)	Array of strings	Optional	A list of any context attributes that are used.

Property	Format	Required?	Description
skipEsLoadIfIndexExists	boolean	Optional	This option applies to graphmarts with Elasticsearch indexes. It controls whether or not Anzo first checks to see if an index with the alias for the dataset already exists in Elasticsearch. If this setting is true and the index does exist, Anzo will not reload the index snaphsot into Elasticsearch.
edition	"uri"	Optional	If you want to load a certain edition in the dataset, you can include this property to specify the URI of the edition.
gmLinkedDataset	"uri"	Optional	The URI of the Linked Dataset Catalog entry that represents the target FLDS for the load. To get the catalog entry, you can retrieve data about the dataset and use the catalogEntry value.
watchFldsDirectory	boolean	Optional	Controls whether the FLDS directory is monitored for changes. If true and files change, Anzo marks the step (and layer) as needing a refresh.

Property	Format	Required?	Description
maskedPredicate	[ "uri", "" ]	Optional	To exclude certain triples from the load, you can specify a list of predicate URIs to filter out.
fldsComponent	[ "uri", "" ]	Optional	If you want to load specific components in the dataset, you can include this property to list the component URIs.
transformQuery	"string"	Optional	If you want to hand-pick the data to load, you can include this property to run a SPARQL query that inserts specific values or filters out certain values.
title	"string"	Required	The name of the step.
type	"string"	Required	The type of step: "LoadDatasetStep".
incrementalData	"string"	Optional	Incremental load data associated with the step.
enabled	boolean	Optional	Controls whether the step is enabled or disabled.
contextProvider	[ "uri", "" ]	Optional	A list of any referenced context providers (the data source URI). You can retrieve data for the parent layer to get a list of providers for that

Property	Format	Required?	Description
			layer.
description	"string"	Optional	A brief description of the step.
ontology	[ "uri", "" ]	Optional	A list of any models to associate with this step.
ignoreLoadErrors	boolean	Optional	Controls whether to ignore errors and proceed with the load or fail the step if there is an error.
disableLoadCounts	boolean	Optional	Controls whether Anzo periodically queries AnzoGraph to count the total number of statements that are processed. Disabling the load count decreases the number of queries that run during activation.
preGenerateStatistics	boolean	Optional	Controls whether AnzoGraph generates statistics on the data before the step is run.
tags	Array of objects	Optional	Any tags on the step.
tagTitle	["string"]	Optional	A virtual property that is available for all objects. It lists the tags associated with the step or can be used to add a

Property	Format	Required?	Description
			tag to the step without including a description.

# **Pre-Compile Query Step**

This type of step runs the specified query immediately after the graphmart load so that the query is pre-compiled by AnzoGraph.

#### **JSON Request**

The following template shows the body of a JSON request that could be used in a Pre-Compile Query Step PUT or PATCH request. It lists all of the step's required and optional body parameters but excludes the read-only options. The default values for each parameter are shown. Below the request (in Schema Details) is a table that describes the complete schema, including the read-only parameters. Clicking a link in the template takes you to the schema details for that parameter.

```
{
  "validationFailsGraphmart" : true,
  "validationFailsLayer" : true,
  "validationQuery" : "string",
  "title" : "string",
  "incrementalData" : [ "string" ],
  "type" : "ValidationStep",
  "enabled" : true,
  "contextProvider" : [ "string" ],
  "description" : "string",
  "ontology" : [ "string" ],
  "source" : [ "string" ],
  "ignoreLoadErrors" : true,
  "disableLoadCounts" : true,
  "preGenerateStatistics" : true,
  "tags" : [ {
                "description" : "string",
                "title" : "string"
            }],
  "tagTitle" : [ "string" ]
}
```

The table below describes the Pre-Compile Query Step schema.

# Тір

You can also see the Pre-Compile Query Step schema by expanding **Schemas** at the bottom of the Anzo REST API document and viewing **PreCompileQueryStep**.

Property	Format	Requir- ed?	Description
<b>uri</b> (read-only)	"uri"	Auto- generat ed	The URI of the step.
<b>creator</b> (read- only)	"uri"	Auto- generat ed	The creator of the step.
<b>created</b> (read- only)	"dateTi me"	Auto- generat ed	The timestamp when the step was created.
<b>modifier</b> (read- only)	"uri"	Auto- generat ed	The user who modified the step.
<b>alltypes</b> (read- only)	Array of strings	Optiona I	A list of the types related to the step, such as PreCompileQueryStep, Step, LayerChild, etc.
<b>contextAttribute</b> (read-only)	Array of strings	Optiona I	A list of any context attributes that are used.
validationFailsGr	boolean	Optiona	If the query fails, this value controls whether the

Property	Format	Requir- ed?	Description
aphmart		I	entire graphmart load should fail.
validationFailsLa yer	boolean	Optiona I	If the query fails, this value controls whether the layer should fail.
validationQuery	"string"	Require d	The SPARQL query that the step should run.
title	"string"	Require d	The name of the step.
incrementalData	"string"	Optiona I	Incremental load data associated with the step.
type	"string"	Require d	The type of step: "PreCompileQueryStep".
enabled	boolean	Optiona I	Controls whether the step is enabled or disabled.
contextProvider	[ "uri", ". " ]	Optiona I	A list of any referenced context providers (the data source URI). You can retrieve data for the parent layer to get a list of providers for that layer.
description	"string"	Optiona I	A brief description of the step.
ontology	[ "uri", ". " ]	Optiona I	A list of any models to associate with this step.

Property	Format	Requir- ed?	Description
SOURCE	["uri", ". "]	Require d	<ul> <li>The source data for the step. Options are any combination of the following values:</li> <li>"http://cambridgesemantics.com/ontologies/Graphmarts#Self": The source is the data that is in this step's layer.</li> <li>"http://cambridgesemantics.com/ontologies/Graphmarts#AllPrevious": The source is the data from all of the successful layers that precede this step's layer. Failed layers are ignored.</li> <li>"http://cambridgesemantics.com/ontologies/Graphmarts#Previous": The source is the data that is in the one layer that precedes this step's layer.</li> <li>"layer_uri": The source is a specific layer in the graphmart.</li> </ul>
ignoreLoadError s	boolean	Optiona I	Controls whether to ignore errors and proceed with the load or fail the step if there is an error.
disableLoadCou nts	boolean	Optiona I	Controls whether Anzo periodically queries AnzoGraph to count the total number of statements that are processed. Disabling the load count decreases the number of queries that run during activation.
preGenerateStati stics	boolean	Optiona I	Controls whether AnzoGraph generates statistics on the data before the step is run.

Property	Format	Requir- ed?	Description
tags	Array of objects	Optiona I	Any tags on the step.
tagTitle	["strin g"]	Optiona I	A virtual property that is available for all objects. It lists the tags associated with the step or can be used to add a tag to the step without including a description.

# **Query-Driven Templated Step**

This type of step creates a reusable template for creating additional query steps.

## **JSON Request**

The following template shows the body of a JSON request that could be used in a Query-Driven Templated Step PUT or PATCH request. It lists all of the step's required and optional body parameters but excludes the read-only options. The default values for each parameter are shown. Below the request (in Schema Details) is a table that describes the complete schema, including the read-only parameters. Clicking a link in the template takes you to the schema details for that parameter.

```
{
    "parametersTemplate" : [ "string" ],
    "template" : "string",
    "title" : "string",
    "incrementalData" : [ "string" ],
    "type" : "QueryDrivenTemplatedStep",
    "enabled" : true,
    "contextProvider" : [ "string" ],
    "description" : "string",
    "ontology" : [ "string" ],
    "source" : [ "string" ],
    "ignoreLoadErrors" : true,
    "disableLoadCounts" : true,
    "preGenerateStatistics" : true,
    "
```

```
"tags" : [ {
        "description" : "string",
        "title" : "string"
        } ],
    "tagTitle" : [ "string" ]
}
```

The table below describes the Query-Driven Templated Step schema.

## Тір

You can also see the Query-Driven Templated Step schema by expanding **Schemas** at the bottom of the Anzo REST API document and viewing **QueryDrivenTemplatedStep**.

Property	Format	Require- d?	Description
<b>uri</b> (read-only)	"uri"	Auto- generat ed	The URI of the step.
<b>creator</b> (read- only)	"uri"	Auto- generat ed	The creator of the step.
<b>created</b> (read- only)	"dateTi me"	Auto- generat ed	The timestamp when the step was created.
<b>modifier</b> (read- only)	"uri"	Auto- generat ed	The user who modified the step.
<b>alltypes</b> (read- only)	Array of strings	Optional	A list of the types related to the step, such as QueryDrivenTemplatedStep, Step, LayerChild, etc.

Property	Format	Require- d?	Description
<b>contextAttribut</b> <b>e</b> (read-only)	Array of strings	Optional	A list of any context attributes that are used.
parametersTe mplate	"string"	Require d	The query to run for determining the key-value pairs.
template	"string"	Require d	The query template. In the query, include keys as parameters in the format $\{key_name\}$ . The keys are replaced at runtime with the values defined for the key.
title	"string"	Require d	The name of the step.
incrementalDat a	"string"	Optional	Incremental load data associated with the step.
type	"string"	Require d	The type of step: "QueryDrivenTemplatedStep".
enabled	boolean	Optional	Controls whether the step is enabled or disabled.
contextProvide r	[ "uri", " ." ]	Optional	A list of any referenced context providers (the data source URI). You can retrieve data for the parent layer to get a list of providers for that layer.
description	"string"	Optional	A brief description of the step.
ontology	[ "uri", " ." ]	Optional	A list of any models to associate with this step.

Property	Format	Require- d?	Description
SOURCE	[ "uri", " ." ]	Require d	<ul> <li>The source data for the step. Options are any combination of the following values:</li> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#Self": The source is the data that is in this step's layer.</li> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#AllPrevious": The source is the data from all of the successful layers that precede this step's layer. Failed layers are ignored.</li> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#Previous": The source is the data that is in the one layer that precedes this step's layer.</li> <li>"layer_uri": The source is a specific layer in the graphmart.</li> </ul>
ignoreLoadErr ors	boolean	Optional	Controls whether to ignore errors and proceed with the load or fail the step if there is an error.
disableLoadCo unts	boolean	Optional	Controls whether Anzo periodically queries AnzoGraph to count the total number of statements that are processed. Disabling the load count decreases the number of queries that run during activation.
preGenerateSt atistics	boolean	Optional	Controls whether AnzoGraph generates statistics on the data before the step is run.

Property	Format	Require- d?	Description
tags	Array of objects	Optional	Any tags on the step.
tagTitle	["string"]	Optional	A virtual property that is available for all objects. It lists the tags associated with the step or can be used to add a tag to the step without including a description.

# **Query Step**

This type of step runs a SPARQL query that creates, cleans, conforms, or transforms data.

### **JSON Request**

The following template shows the body of a JSON request that could be used in a Query Step PUT or PATCH request. It lists all of the step's required and optional body parameters but excludes the read-only options. The default values for each parameter are shown. Below the request (in Schema Details) is a table that describes the complete schema, including the read-only parameters. Clicking a link in the template takes you to the schema details for that parameter.

```
{
  "transformQuery" : "string",
  "title" : "string",
  "incrementalData" : [ "string" ],
  "type" : "QueryStep",
  "enabled" : true,
  "contextProvider" : [ "string" ],
  "description" : "string",
  "ontology" : [ "string" ],
  "source" : [ "string" ],
  "ignoreLoadErrors" : true,
  "disableLoadCounts" : true,
  "preGenerateStatistics" : true,
  "tags" : [ {
                "description" : "string",
                "title" : "string"
            }],
```

```
"tagTitle" : [ "string" ]
}
```

The table below describes the Query Step schema.

## Tip

You can also see the Query Step schema by expanding **Schemas** at the bottom of the Anzo REST API document and viewing **QueryStep**.

Property	Format	Require- d?	Description
<b>uri</b> (read-only)	"uri"	Auto- generat ed	The URI of the step.
<b>creator</b> (read- only)	"uri"	Auto- generat ed	The creator of the step.
<b>created</b> (read- only)	"dateTi me"	Auto- generat ed	The timestamp when the step was created.
<b>modifier</b> (read- only)	"uri"	Auto- generat ed	The user who modified the step.
<b>alltypes</b> (read- only)	Array of strings	Optional	A list of the types related to the step, such as QueryStep, Step, LayerChild, etc.
<b>contextAttribut</b> <b>e</b> (read-only)	Array of strings	Optional	A list of any context attributes that are used.

Property	Format	Require- d?	Description
transformQuer y	"string"	Require d	The SPARQL query to run.
title	"string"	Require d	The name of the step.
incrementalDat a	"string"	Optional	Incremental load data associated with the step.
type	"string"	Require d	The type of step: "QueryStep".
enabled	boolean	Optional	Controls whether the step is enabled or disabled.
contextProvide r	[ "uri", " ." ]	Optional	A list of any referenced context providers (the data source URI). You can retrieve data for the parent layer to get a list of providers for that layer.
description	"string"	Optional	A brief description of the step.
ontology	[ "uri", " ." ]	Optional	A list of any models to associate with this step.
source	[ "uri", " ." ]	Require d	The source data for the step. Options are any combination of the following values:
			aphmarts#Self": The source is the data that is in this step's layer.
			<ul> <li>"http://cambridgesemantics.com/ontologies/Gr</li> </ul>

Property	Format	Require- d?	Description
			<ul> <li>aphmarts#AllPrevious": The source is the data from all of the successful layers that precede this step's layer. Failed layers are ignored.</li> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#Previous": The source is the data that is in the one layer that precedes this step's layer.</li> <li>"layer_uri": The source is a specific layer in the graphmart.</li> </ul>
ignoreLoadErr ors	boolean	Optional	Controls whether to ignore errors and proceed with the load or fail the step if there is an error.
disableLoadCo unts	boolean	Optional	Controls whether Anzo periodically queries AnzoGraph to count the total number of statements that are processed. Disabling the load count decreases the number of queries that run during activation.
preGenerateSt atistics	boolean	Optional	Controls whether AnzoGraph generates statistics on the data before the step is run.
tags	Array of objects	Optional	Any tags on the step.
tagTitle	["string"]	Optional	A virtual property that is available for all objects. It lists the tags associated with the step or can be used to add a tag to the step without including a description.

# **RDFS+ Inference Step**

This type of step uses RDFS and OWL rules to infer new information about your data based on the vocabularies in the existing data.

#### **JSON Request**

The following template shows the body of a JSON request that could be used in an RDFS+ Inference Step PUT or PATCH request. It lists all of the step's required and optional body parameters but excludes the read-only options. The default values for each parameter are shown. Below the request (in Schema Details) is a table that describes the complete schema, including the read-only parameters. Clicking a link in the template takes you to the schema details for that parameter.

```
{
  "inferenceRules" : "string",
  "title" : "string",
  "incrementalData" : [ "string" ],
  "type" : "RDFSInferenceStep",
  "enabled" : true,
  "contextProvider" : [ "string" ],
  "description" : "string",
  "ontology" : [ "string" ],
  "source" : [ "string" ],
  "ignoreLoadErrors" : true,
  "disableLoadCounts" : true,
  "preGenerateStatistics" : true,
  "tags" : [ {
                "description" : "string",
                "title" : "string"
            }],
  "tagTitle" : [ "string" ]
}
```

#### **Schema Details**

The table below describes the RDFS+ Inference Step schema.

## Тір

You can also see the RDFS+ Inference Step schema by expanding **Schemas** at the bottom of the Anzo REST API document and viewing **RDFSInferenceStep**.

Property	Format	Require- d?	Description
<b>uri</b> (read-only)	"uri"	Auto- generat ed	The URI of the step.
<b>creator</b> (read- only)	"uri"	Auto- generat ed	The creator of the step.
<b>created</b> (read- only)	"dateTi me"	Auto- generat ed	The timestamp when the step was created.
<b>modifier</b> (read- only)	"uri"	Auto- generat ed	The user who modified the step.
<b>alltypes</b> (read- only)	Array of strings	Optional	A list of the types related to the step, such as RDFSInferenceStep, Step, LayerChild, etc.
<b>contextAttribut</b> e (read-only)	Array of strings	Optional	A list of any context attributes that are used.
inferenceRules	"string"	Optional	By default all inference rules are run. This property can be used to list a subset of rules to run or specific rules to exclude. For more information about the inference rules, see Infer New Data (RDFS+ Inference Step).

Property	Format	Require- d?	Description
title	"string"	Require d	The name of the step.
incrementalDat a	"string"	Optional	Incremental load data associated with the step.
type	"string"	Require d	The type of step: "RDFSInferenceStep".
enabled	boolean	Optional	Controls whether the step is enabled or disabled.
contextProvide r	["uri", " ." ]	Optional	A list of any referenced context providers (the data source URI). You can retrieve data for the parent layer to get a list of providers for that layer.
description	"string"	Optional	A brief description of the step.
ontology	[ "uri", " ." ]	Optional	A list of any models to associate with this step.
source	["uri", " ."]	Require d	<ul> <li>The source data for the step. Options are any combination of the following values:</li> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#Self": The source is the data that is in this step's layer.</li> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#AllPrevious": The source is the data from all of the successful layers that precede this step's layer. Failed layers are ignored.</li> </ul>

Property	Format	Require- d?	Description
			<ul> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#Previous": The source is the data that is in the one layer that precedes this step's layer.</li> <li>"layer_uri": The source is a specific layer in the graphmart.</li> </ul>
ignoreLoadErr ors	boolean	Optional	Controls whether to ignore errors and proceed with the load or fail the step if there is an error.
disableLoadCo unts	boolean	Optional	Controls whether Anzo periodically queries AnzoGraph to count the total number of statements that are processed. Disabling the load count decreases the number of queries that run during activation.
preGenerateSt atistics	boolean	Optional	Controls whether AnzoGraph generates statistics on the data before the step is run.
tags	Array of objects	Optional	Any tags on the step.
tagTitle	["string"]	Optional	A virtual property that is available for all objects. It lists the tags associated with the step or can be used to add a tag to the step without including a description.

# **Templated Step**

This type of step creates a reusable query template based on key-value pairs.

#### **JSON Request**

The following template shows the body of a JSON request that could be used in a Templated Step PUT or PATCH request. It lists all of the step's required and optional body parameters but excludes the read-only options. The default values for each parameter are shown. Below the request (in Schema Details) is a table that describes the complete schema, including the read-only parameters. Clicking a link in the template takes you to the schema details for that parameter.

```
{
   "template" : "string",
   "templateParameters" : [
     {
       "type": "TemplateParameterSet",
       "templateParameter": [
         {
           "type": "TemplateParameter",
           "parameterName": "key",
           "parameterValue": [ "uri" ],
         }
       1
    }
  ],
   "title" : "string",
   "incrementalData" : [ "string" ],
   "type" : "TemplatedStep",
  "enabled" : true,
   "contextProvider" : [ "string" ],
  "description" : "string",
  "ontology" : [ "string" ],
   "source" : [ "string" ],
  "ignoreLoadErrors" : true,
  "disableLoadCounts" : true,
   "preGenerateStatistics" : true,
   "tags" : [ {
                "description" : "string",
                "title" : "string"
            }],
   "tagTitle" : [ "string" ]
}
```

#### **Schema Details**

The table below describes the Templated Step schema.

# Тір

You can also see the Templated Step schema by expanding **Schemas** at the bottom of the Anzo REST API document and viewing **TemplatedStep**.

Property	Format	Require- d?	Description
<b>uri</b> (read-only)	"uri"	Auto- generat ed	The URI of the step.
<b>creator</b> (read- only)	"uri"	Auto- generat ed	The creator of the step.
<b>created</b> (read- only)	"dateTi me"	Auto- generat ed	The timestamp when the step was created.
<b>modifier</b> (read- only)	"uri"	Auto- generat ed	The user who modified the step.
<b>alltypes</b> (read- only)	Array of strings	Optional	A list of the types related to the step, such as TemplatedStep, Step, LayerChild, etc.
<b>contextAttribut</b> <b>e</b> (read-only)	Array of strings	Optional	A list of any context attributes that are used.
template	"string"	Require d	The query template. In the query, include keys as parameters in the format $\{key_name\}$ . The keys are replaced at runtime with the values defined for the key.
Property	Format	Require- d?	Description
------------------------	---------------------	----------------	--
templateParam eters	Array of objects	Require d	The list of objects that define the key-value pairs.
templateParam eter	Array of objects	Require d	The list of key-value pairs that will be substituted for the \${key_name} parameters in the query template. Each templateParameter must include type, parameterName, and parameterValue. The following snippet shows the format of template parameters.
			<pre>"templateParameter": [     {         "type": "TemplateParameter",         "parameterName": "key",         "parameterValue": [ "uri" ],     } ], "templateParameter": [     {      } ]</pre>
parameterNam e	"string"	Require d	The name of the key that is used in the query template.
parameterValue	[ "uri" ]	Require d	The value for the corresponding key (parameterName).
title	"string"	Require d	The name of the step.
incrementalDat a	"string"	Optional	Incremental load data associated with the step.

Property	Format	Require- d?	Description
type	"string"	Require d	The type of step: "TemplatedStep".
enabled	boolean	Optional	Controls whether the step is enabled or disabled.
contextProvide r	["uri", " ." ]	Optional	A list of any referenced context providers (the data source URI). You can retrieve data for the parent layer to get a list of providers for that layer.
description	"string"	Optional	A brief description of the step.
ontology	[ "uri", " ." ]	Optional	A list of any models to associate with this step.
source	[ "uri", " ." ]	Require d	The source data for the step. Options are any combination of the following values:
			<ul> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#Self": The source is the data that is in this step's layer.</li> </ul>
			<ul> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#AllPrevious": The source is the data from all of the successful layers that precede this step's layer. Failed layers are ignored.</li> </ul>
			<ul> <li>"http://cambridgesemantics.com/ontologies/Gr aphmarts#Previous": The source is the data that is in the one layer that precedes this step's layer.</li> </ul>
			<ul> <li>"layer_uri": The source is a specific layer in the</li> </ul>

Property	Format	Require- d?	Description
			graphmart.
ignoreLoadErr ors	boolean	Optional	Controls whether to ignore errors and proceed with the load or fail the step if there is an error.
disableLoadCo unts	boolean	Optional	Controls whether Anzo periodically queries AnzoGraph to count the total number of statements that are processed. Disabling the load count decreases the number of queries that run during activation.
preGenerateSt atistics	boolean	Optional	Controls whether AnzoGraph generates statistics on the data before the step is run.
tags	Array of objects	Optional	Any tags on the step.
tagTitle	["string"]	Optional	A virtual property that is available for all objects. It lists the tags associated with the step or can be used to add a tag to the step without including a description.

### **Validation Step**

This type of step runs a validation query to ensure that data conforms to expectations.

#### **JSON Request**

The following template shows the body of a JSON request that could be used in a Validation Step PUT or PATCH request. It lists all of the step's required and optional body parameters but excludes the read-only options. The default values for each parameter are shown. Below the request (in Schema Details) is a table that describes the complete schema, including the read-only parameters. Clicking a link in the template takes you to the schema details for that parameter.

```
{
```

```
"validationFailsLayer" : true,
"validationDatasourceUri" : "string",
"validationQuery" : "string",
"validationFailsGraphmart" : true,
"resultVar" : "string",
"validationOrConditional" : "string",
"title" : "string",
"incrementalData" : [ "string" ],
"type" : "ValidationStep",
"enabled" : true,
"contextProvider" : [ "string" ],
"description" : "string",
"ontology" : [ "string" ],
"source" : [ "string" ],
"ignoreLoadErrors" : true,
"disableLoadCounts" : true,
"preGenerateStatistics" : true,
"tags" : [ {
             "description" : "string",
             "title" : "string"
         }],
"tagTitle" : [ "string" ]
```

#### **Schema Details**

The table below describes the Validation Step schema.

#### Tip

}

You can also see the Validation Step schema by expanding **Schemas** at the bottom of the Anzo REST API document and viewing **ValidationStep**.

Property	Format	Requir- ed?	Description
uri (read-only)	"uri"	Auto- generat ed	The URI of the step.

Property	Format	Requir- ed?	Description
<b>creator</b> (read- only)	"uri"	Auto- generat ed	The creator of the step.
<b>created</b> (read- only)	"dateTi me"	Auto- generat ed	The timestamp when the step was created.
<b>modifier</b> (read- only)	"uri"	Auto- generat ed	The user who modified the step.
<b>alltypes</b> (read- only)	Array of strings	Optiona I	A list of the types related to the step, such as ValidationStep, Step, LayerChild, etc.
<b>contextAttribute</b> (read-only)	Array of strings	Optiona I	A list of any context attributes that are used.
validationFailsLa yer	boolean	Optiona I	If validationOrConditional is validation, and the validation fails, this value controls whether the layer should fail.
validationDataso urceUri	"uri"	Optiona I	The source to perform the validation on.
validationQuery	"string"	Require d	The SPARQL query that the step should run.
validationFailsGr aphmart	boolean	Optiona I	If validationOrConditional is validation, and the validation fails, this value controls whether the entire graphmart load should fail.

Property	Format	Requir- ed?	Description
resultVar	"string"	Optiona I	The variable name to use to store the result from the query. This variable becomes available when configuring an execution condition for a layer or step.
validationOrCon ditional	"string"	Optiona I	The type of check to perform: validation or condition. A validation check validates the data according to the defined query and takes the action configured with validationFailsLayer and validationFailsGraphmart. A condition check takes the results of the query and associates it with the specified resultVar.
title	"string"	Require d	The name of the step.
incrementalData	"string"	Optiona I	Incremental load data associated with the step.
type	"string"	Require d	The type of step: "ValidationStep".
enabled	boolean	Optiona I	Controls whether the step is enabled or disabled.
contextProvider	["uri", ". "]	Optiona I	A list of any referenced context providers (the data source URI). You can retrieve data for the parent layer to get a list of providers for that layer.
description	"string"	Optiona I	A brief description of the step.

Property	Format	Requir- ed?	Description
ontology	[ "uri", ". " ]	Optiona I	A list of any models to associate with this step.
SOURCE	["uri", ". "]	Require d	<ul> <li>The source data for the step. Options are any combination of the following values:</li> <li>"http://cambridgesemantics.com/ontologies/Graphmarts#Self": The source is the data that is in this step's layer.</li> <li>"http://cambridgesemantics.com/ontologies/Graphmarts#AllPrevious": The source is the data from all of the successful layers that precede this step's layer. Failed layers are ignored.</li> <li>"http://cambridgesemantics.com/ontologies/Graphmarts#Previous": The source is the data that is in the one layer that precedes this step's layer.</li> <li>"layer_uri": The source is a specific layer in the graphmart.</li> </ul>
ignoreLoadError s	boolean	Optiona I	Controls whether to ignore errors and proceed with the load or fail the step if there is an error.
disableLoadCou nts	boolean	Optiona I	Controls whether Anzo periodically queries AnzoGraph to count the total number of statements that are processed. Disabling the load count decreases the number of queries that run during activation.

Property	Format	Requir- ed?	Description
preGenerateStati stics	boolean	Optiona I	Controls whether AnzoGraph generates statistics on the data before the step is run.
tags	Array of objects	Optiona I	Any tags on the step.
tagTitle	["strin g"]	Optiona I	A virtual property that is available for all objects. It lists the tags associated with the step or can be used to add a tag to the step without including a description.

# Anzo Java SDK

This topic provides instructions for setting up an Anzo development environment using the Anzo software development kit (SDK) and Eclipse integrated development environment (IDE). The sample instructions below deploy the Anzo SDK in a Windows environment with Eclipse IDE for Java Developers Version 4.12.0. Anzo SDK and Eclipse can also be deployed on Linux and Mac operating systems.

### Requirements

Make sure that the Anzo development server meets the requirements in <u>Anzo Requirements</u> in the Deployment Guide. In addition, install the following programs for working with the Anzo Java SDK:

- Eclipse for Java Developers Version 4.7.3+: Install the Eclipse IDE for Java Developers or Eclipse IDE for Enterprise Java Developers.
- Java Runtime Environment Version 8: Eclipse and the Anzo SDK require JDK version 8.
   Cambridge Semantics tests with jdk1.8.0\_181.

## **Deploying the Anzo SDK with Eclipse**

Follow the instructions below to import the Anzo Java SDK to Eclipse and configure and test the environment.

- 1. Download the Anzo SDK .zip file to the host server. Do not unpack the file.
- 2. In Eclipse, click the **File** menu and select **Import**. Eclipse opens the Import dialog box. For example:

(	🖨 Im	nport				$\times$			
:	Select Create new projects from an archive file or directory.								
	Select	t an import wizard:							
	type	filter text							
	× 6	→ General				^			
		. Archive File							
		😂 Existing Projects into Workspace							
		😂 File System							
		Preferences							
		🛱 Projects from Folder or Archive				$\sim$			
	?	< Back Next > Fin	ish		Cance	l			
			1311		Curice				

3. In the Import dialog box, expand the **General** folder and select **Existing Projects into Workspace** and click **Next**. Eclipse opens the Import Projects dialog box. For example:

Import		
Import Projects Select a directory to search for existing Eclipse projects.		
Select root directory:	~	Browse
○ Select archive file:	~	Browse
Projects:		
		Select All
		Deselect All
		Refresh
Options Search for nested projects Copy projects into workspace Close newly imported projects upon completion Hide projects that already exist in the workspace		
Working sets		
Add project to working sets		New
Working sets:	~	Select
? < Back Next > Finish		Cancel

4. Select the **Select archive file** radio button and then browse to and select the Anzo SDK .zip file. Eclipse loads the .zip file and lists the contents in the Projects field. For example:

Import				$\times$
Import Projects Select a directory to sea	rch for existing Eclipse projects.			
$\bigcirc$ Select root directory:		~	Brows	e
• Select archive file:	\Anzo_SDK_4.3.0	×	Brows	e
Projects:				
com.cambridgese	mantics.anzo.sdk (com.cambridgesemantics.anzo	.sdk,	Select	All
✓ com.cambridgese	mantics.anzo.sdk.api (com.cambridgesemantics.a mantics anzo.sdk.server (com.cambridgesemantic	nzo.:	Deselec	t All
€ coni.cumbridgese	manaes.anzo.saleserver (conneumonagesernande	5.011	Refree	sh
<		>		
Options				
Search for nested pro	ojects			
Copy projects into w	orkspace			
Hide projects that all	a projects upon completion eady exist in the workspace			
Working sets				
Add project to work	ing sets		New	
Working sets:	5	~	Select	
5				
?	< Back Next > Finish		Cance	I

The Anzo SDK contains three projects:

- **com.cambridgesemantics.anzo.sdk**: This core project is required for creating solutions. It contains the Anzo libraries that provide the Anzo APIs and extension points as well as the libraries that enable Anzo to run in the development environment.
- **com.cambridgesemantics.anzo.sdk.server** This core project is required for creating solutions. It contains configuration files for running Anzo as well as a launcher for starting the Anzo server.
- **com.cambridgesemantics.anzo.sdk.api**: This is an example project that contains sample Java programs that illustrate several aspects of the Anzo client APIs. Each program is a simple example that demonstrates how to communicate with the Anzo server to read, write, and query data. See the comments in each example for an explanation of what each one demonstrates.
- 5. Click **Finish** to import the Anzo SDK .jar files. The process may take a few minutes. When the import is complete, Eclipse opens the workspace. At this point in the process, expect to see several errors in the workspace. For example:



- 6. Import your Anzo license:
  - a. Make sure that you have a copy of the Anzo license on the server. If necessary, you can view and download a copy from the <u>Cambridge Semantics Support Center</u>.
  - b. Rename the license file so its file extension is .lic. For example, license.lic.
  - c. In the Eclipse Package Explorer, right-click com.cambridgesemantics.anzo.sdk.server and select Import.
  - In the Import dialog box, expand the General folder and select File System. Then click
     Next. Eclipse opens the File System Import dialog box. For example:

Comport			×
File system			
Source must not be empty.			-
From directory:	~	Browse	
Filter Types Select All Deselect All			
Into folder: com.cambridgesemantics.anzo.sdk.server		Browse	
Options			
✓ Overwrite existing resources without warning			
Advanced >>			
? < Back Next > Finish		Cancel	

e. Click the **Browse** button next to the From directory field and select the directory that contains the license file. Eclipse displays the directory and its contents.

lmport	_	
File system Import resources from the local file system.		
From directory: C:\Users\Erin Martin\Downlo	oads\license ~	Browse
Filter Types Select All D	Deselect All	
Into folder: com.cambridgesemantics.anzo.s	dk.server	Browse
Overwrite existing resources without warn     Create top-level folder	ing	
Advanced >>		
? < Back Next >	Finish	Cancel

f. Select the license file in the right pane, and then click **Finish**.

- 7. Install the Eclipse Plugin Development Tools:
  - Click the Help menu and select Install New Software. Eclipse opens the Install dialog box.
  - In the Install dialog box, click the Work with drop-down list and select All Available
     Sites. In the search field below the Work with field, type "PDE" and wait for Eclipse to
     find the plugin tools. Select the checkbox next to Eclipse Plugin Development Tools,
     including Eclipse PDE Plug-in Developer Resources. For example:

Install			×
<b>Available Software</b> Check the items that you wish to install.			
Work with:All Available Sites		Manag	je
PDE	0_	Select	All
Name	eady i	Deselec	t All
⑦ < Back Next > Finish		Cancel	

- c. Click **Next** and accept the license agreement, then click **Finish**. Eclipse installs the software and then prompts you to restart the application.
- 8. After restarting Eclipse, load the Anzo SDK Target Platform:
  - a. Click the **Window** menu and select **Preferences**.
  - b. In the Preferences dialog box, expand **Plug-in Development** and select **Target Platform**.

c. In the Target Platform definitions, select the **Anzo SDK Devel Target** checkbox. For example:

Preferences	_	
type filter text	Target Platform	⇔ ▼ ⇔ ▼ ▼
<ul> <li>&gt; General</li> <li>&gt; Ant</li> <li>Gradle</li> <li>&gt; Help</li> </ul>	Add, edit and remove target definitions. The active target definition will be used as the which workspace plug-ins will be compiled and tested against. New definitions are stor can be moved to a project in the workspace and shared with others. Target definitions:	target platform ed locally, but they
Install/Update	Anzo SDK Devel Target (Active) - /com.cambridgesemantics.anzo.sdk/devel.ta	Reload
> Java > Maven > Mylyn	□ <sup>®</sup> Running Platform	Add
> Oomph		Edit
✓ Plug-in Development		Remove
API Baselines API Errors/Warnings API Use Scans Compilers DS Annotations Editors		Share
OSGi Frameworks	< >	
Target Platform	Locations:	
<ul> <li>&gt; Run/Debug</li> <li>&gt; Team</li> <li>Validation</li> <li>&gt; XML</li> </ul>	<ul> <li>\${workspace_loc:/com.cambridgesemantics.anzo.sdk}/plugins/dependencies - C:\</li> <li>\${workspace_loc:/com.cambridgesemantics.anzo.sdk}/plugins/anzo - C:\Users\Erii</li> <li>\${workspace_loc:/com.cambridgesemantics.anzo.sdk}/plugins/openanzo - C:\User</li> </ul>	Users\Erin Martin\e n Martin\eclipse-wc rs\Erin Martin\eclip: >
	Restore Defaults	Apply
? Þ 🖌 🖲	Apply and Close	Cancel

- d. Click Apply and Close. Eclipse loads the Anzo SDK Target Platform.
- 9. Test the environment:
  - a. In the Eclipse workspace, click the **Run** menu and select **Run Configurations**. Eclipse opens the Run Configurations dialog box.
  - b. On the left side of the dialog box, expand the OSGi Framework folder and select Anzo (SDK). For example:

Run Configurations	_	
Create, manage, and run co Create a configuration to launch	nfigurations n the OSGi framework.	
Image: Contract of the second seco	Name:       Anzo (SDK)         Bundles       Arguments         Launch with:       bundles selected below         Framework:       Equinox         Uppe filter text         Bundles       Start L         Auto-S         Workspace         Include optional dependencies when computing required Bundles         Add new workspace Bundles to this launch configuration automatically         Validate Bundles automatically prior to launching	Default /   A
Filter matched 11 of 11 items	Re <u>v</u> ert	Apply
0	Run	Close

 c. Click Run to run the Anzo SDK target platform. A Console tab opens in Eclipse and shows the status messages. When Anzo starts, the console displays the message "All Currently Registered Services started." For example:



If Anzo fails to start, one of the common reasons for the failure is that one or more of the Anzo ports are in use by other software. See <u>Firewall Requirements</u> in the Deployment Guide for information about the ports that Anzo uses.

To explore the sample Java programs that are included in the Anzo SDK, expand the **com.cambridgesemantics.anzo.sdk.api** package in the Package Explorer. In the package, expand the **src** directory and then the **com.cambridgesemantics.anzo.sdk.api** directory to see the list of sample programs. For example:



To run a program, right-click the .java file and select **Run As > Java Application**. For more information about using the Anzo SDK, see the **Anzo Java SDK Guide.pdf** that is distributed in the SDK .zip file.