




Kyubit Semantic Modeling Guide: Designing and Maintaining Data Models for SQL and Excel/CSV Sources

©2026 Kyubit, all rights reserved

www.kyubit.com

Table of Contents

1. Introduction to Semantic Modeling & Kyubit BI	4
1.1 What is a Semantic Model?	4
1.2 How the Kyubit Semantic Engine Works	4
1.3 Core Structural Building Blocks	5
2. Data Source Connectivity & Modeling Modes	7
2.1 Supported Relational Data Sources	7
2.2 Model Approaches: Visual Diagram vs. Single-Dataset	7
3. Designing a Visual Diagram Semantic Model (SQL Sources).....	9
3.1 Navigating the Visual Model Designer	9
3.2 Canvas Navigation: Zoom and Pan	11
3.3 Defining the Central Fact Table	12
3.4 Adding Dimension Tables & Visual Table Joins	14
3.5 Auditing with Live 200-Row Sample Data	15
3.6 Applying SQL Table Filters	16
3.7 Mapping & Formatting Measures	17
3.8 Creating Dimension Entities & Mapping Attributes	18
3.9 Automated Time Dimensions & Date Hierarchies	20
3.10 Building Multi-Level Attribute Hierarchies.....	22
3.11 High-Cardinality Dimension Details (Drillthrough)	23
3.12 Value Bucketing with Dimension Attribute Ranges	25
4. Designing Single-Dataset Models (Excel, CSV, and Direct SQL Queries)	26
4.1 Excel and CSV File Uploads	26
4.2 SQL Query Datasets.....	26
4.3 Column-Based Data Definitions	27
4.4 Structure Organization & Sorting Keys.....	27
5. Model Processing, Lifecycle & Automation	28
5.1 Processing States & Compilation	28
5.1.1. Full Model Processing	29
5.1.2. Partial Processing (Rolling Time Frame Refresh)	30
5.1.3. Incremental Processing (Delta Ingestion)	30
5.2 Processing Logs & Troubleshooting	31

5.3 Automated Updates with Scheduled Jobs	33
5.4 Sizing Guidelines & Performance Best Practices.....	33
6. Data Analysis, Reporting, Dashboards & KPIs	35
6.1 Interactive Grid & Pivot Analysis.....	35
6.2 Advanced Date Period & Relative Filtering	35
6.3 Dynamic On-the-Fly Calculated Measures.....	37
6.4 Designing Interactive Dashboards & KPI Widgets.....	38
6.5 Permissions, Governance & Collaboration	40

1. Introduction to Semantic Modeling & Kyubit BI

1.1 What is a Semantic Model?

A **Semantic Model** serves as an intelligent business translation layer positioned between complex physical database structures (tables, foreign keys, views, raw files) and non-technical business end-users.

In raw transactional databases, data is frequently split across normalized tables with cryptic technical column names, complex join keys, and unformatted numeric values. A semantic model abstracts this complexity by:

- Mapping physical columns into intuitive business terminology (e.g., transforming `tbl_cust_addr_line1` into Customer Address).
- Defining numerical facts as aggregate **Measures** with consistent formatting and business logic.
- Organizing categorical attributes into structured **Dimensions** and drill-down **Hierarchies**.
- Predefining multi-table relationships so business users can slice, dice, and pivot data without writing SQL queries or understanding database schemas.

1.2 How the Kyubit Semantic Engine Works

Kyubit eliminates the requirement for heavy, specialized OLAP cube infrastructures (such as SSAS) while delivering the familiar flexibility of multidimensional pivot-table analytics.

- **Underlying Architecture:** When a semantic model is defined and processed, Kyubit structures the analytical metadata and storage inside the internal KyubitAnalyticModels database on Microsoft SQL Server.

- **ColumnStore Indexing & Caching:** Kyubit utilizes SQL Server **ColumnStore indexes** combined with smart in-memory caching mechanisms. This architecture provides near-instant aggregations across hundreds of thousands or millions of records.
- **Dynamic Query Generation:** When an end-user performs slice-and-dice, drill-down, or filtering operations in an analysis or dashboard, Kyubit dynamically translates those visual gestures into optimized SQL queries executed against the optimized model storage.

1.3 Core Structural Building Blocks

Every Kyubit semantic model is composed of distinct structural components:

Structure Type	Purpose	Behavior in Analysis
Measures	Quantitative, numerical values (e.g., Sales Amount, Units Sold, Margin).	Aggregated using mathematical functions (Sum, Avg, Count, Min, Max). Used on values/metrics axes.
Dimensions & Attributes	Categorical context (e.g., Customer Name, Product Category, Region).	Placed on the Categories (rows), Series (columns), or Slicers axes to slice and segment measures.
Date / Time Dimensions	Temporal data representing transaction dates or milestones.	Provides automated multi-level hierarchies (Year → Quarter → Month → Day) and flexible date-range filtering.
Hierarchies	Structured parent-child drill paths (e.g., Continent → Country → State → City).	Enables one-click drill-down and roll-up across related analytical levels.

Structure Type	Purpose	Behavior in Analysis
Dimension Details	High-cardinality metadata (e.g., Phone Numbers, Email Addresses, Street Names).	Excluded from top-level aggregations to protect performance; available immediately during drillthrough actions.
Attribute Ranges	Numeric bucketing intervals (e.g., Age brackets 20–29, 30–39 or Order Value tiers).	Transforms continuous numeric values into clean categorical segments for reporting.
Key Sorting Columns (this refers to older Kyubit semantic models)	Secondary numeric order identifiers (e.g., Month Number 1..12 for Month Name Jan..Dec).	Enforces natural chronological or business ordering instead of standard alphabetical sorting.

The screenshot displays the Kyubit Business Intelligence interface for the Adventure Works Semantic Model. The interface is divided into several sections:

- Database Tables:** A sidebar on the left lists various database tables, including 'AW SQL' and 'dbo.AdventureWorksDWBld...'. It also includes a search bar and navigation icons.
- Analytic Model:** The main area shows a diagram of the semantic model. It includes dimensions like 'dbo.DimGeography', 'dbo.DimCustomer', and 'dbo.DimPromotion', and a fact table 'dbo.FactInternetSales'. Arrows indicate relationships between these tables.
- Model Structure:** A panel on the right shows the hierarchy of the semantic model. It includes measures like 'TotalProductCost', 'SalesAmount', and 'TaxAmt', and dimensions like 'Geo', 'Customer', and 'OrderDate'.
- Processing and Schedule:** Tabs at the top right allow users to manage the processing and scheduling of the model.

2. Data Source Connectivity & Modeling Modes

2.1 Supported Relational Data Sources


Kyubit connects natively to major relational database management systems. Before building a visual semantic model, register a new **Data Source** inside the Kyubit administration panel pointing to your database:

- **Microsoft SQL Server**
- **PostgreSQL**
- **Oracle Database**
- **MySQL**
- **MariaDB**
- **SQLite**
- **ODBC-compliant Data Sources**

2.2 Model Approaches: Visual Diagram vs. Single-Dataset

Kyubit offers two distinct modeling workflows depending on the origin of your data:

1. **Diagram-Based Semantic Model (Visual Designer):** Available for native SQL data sources. Allows full multi-table modeling on an interactive canvas, including visual joins, multi-table attribute extraction, fact table configuration, and table-level SQL filtering.



2. **Single-Dataset Model (Query & File-Based):** Used for uploaded Excel/CSV files or custom single-statement SQL queries. Designed for flat tabular data where all measures, dimensions, and details reside in a single result set.


Measure	Measure	Measure	Date	Date	Dimension	Dimension	Dimension	Dimension	Dimension
SalesAmount	TaxAmt	Product Standard Cost	OrderDate	ShipDate	Color	Product Name	ModelName	Size	Product Subcategory Name
\$#,###	\$#,###	€#,###							
\$3,578.27	\$286.26	€2,171.29	29.12.2028. 0:00:00	5.1.2029. 0:00:00	Red	Road-150 Red, 62	Road-150	62	Road Bikes
\$3,399.99	\$272	€1,912.15	29.12.2028. 0:00:00	5.1.2029. 0:00:00	Silver	Mountain-100 Silver, 44	Mountain-100	44	Mountain Bikes
\$3,399.99	\$272	€1,912.15	29.12.2028. 0:00:00	5.1.2029. 0:00:00	Silver	Mountain-100 Silver, 44	Mountain-100	44	Mountain Bikes
\$699.1	\$55.93	€413.15	29.12.2028. 0:00:00	5.1.2029. 0:00:00	Black	Road-650 Black, 62	Road-650	62	Road Bikes
\$3,399.99	\$272	€1,912.15	29.12.2028. 0:00:00	5.1.2029. 0:00:00	Silver	Mountain-100 Silver, 44	Mountain-100	44	Mountain Bikes
\$3,578.27	\$286.26	€2,171.29	30.12.2028. 0:00:00	6.1.2029. 0:00:00	Red	Road-150 Red, 44	Road-150	44	Road Bikes
\$3,578.27	\$286.26	€2,171.29	30.12.2028. 0:00:00	6.1.2029. 0:00:00	Red	Road-150 Red, 62	Road-150	62	Road Bikes

3. Designing a Visual Diagram Semantic Model (SQL Sources)


3.1 Navigating the Visual Model Designer

When creating a new **Diagram-based Analytic Model**, select your target SQL data source and open the **Model Designer** tab. The workspace is divided into three functional zones:


1. **Left Panel (Database Tables):** Displays all tables and views discovered in the connected database, complete with a quick-search filter bar.



2. **Central Workspace (Interactive Canvas):** The visual modeling area where tables are arranged, joined via drag-and-drop connectors, filtered, and audited.




3. **Right Panel (Model Structure Tree):** Displays the final business semantic architecture—showing the configured Measures, Dimensions, Hierarchies, and Details as end-users will see them.



3.2 Canvas Navigation: Zoom and Pan

To manage schemas containing numerous entities:


- **Zoom:** Use the mouse scroll wheel forward or backward to smoothly scale the entire canvas in and out.
- **Pan:** Click and hold empty canvas space to drag and pan across large diagrams.
- **Resize & Move Entities:** Grab any table header to move it across the canvas; drag the bottom-right corner handles to expand or contract table containers.




3.3 Defining the Central Fact Table

A **Fact Table** contains the core quantitative, transactional metrics of your business process (e.g., Sales, Invoices, General Ledger).


1. Locate your central transactional table in the left panel.




2. Drag and drop the table onto the central canvas.



3. Click the context menu (...) in the table header and select **Set as Fact Table**.




4. The table header updates to indicate fact status, enabling metric selection checkboxes.



3.4 Adding Dimension Tables & Visual Table Joins

Dimension tables provide descriptive context (Customers, Products, Stores, Geography).

1. Drag relevant dimension tables from the left panel onto the canvas around the central fact table.
2. To create a SQL join relationship, click and drag the **Primary Key** column from the dimension table and drop it directly onto the corresponding **Foreign Key** column in the fact table.




3. A visual connector arrow connects the two entities, indicating an established relationship and color-coding the dimension table as active in the model.

3.5 Auditing with Live 200-Row Sample Data

To verify table contents and column formats without leaving the designer:

1. Click the context menu (...) on any table header on the canvas.
2. Select **Show Table Sample Data**.



3. A modal immediately displays the first 200 rows of live data from that table, allowing you to confirm column values and data quality.

Sample 200 rows.					
CustomerKey	GeographyKey	CustomerAlternateKey	Title	FirstName	MiddleName
11000	26	AW00011000	null	Jon	V
11001	37	AW00011001	null	Eugene	L
11002	31	AW00011002	null	Ruben	null
11003	11	AW00011003	null	Christy	null
11004	19	AW00011004	null	Elizabeth	null

3.6 Applying SQL Table Filters

To restrict processing scope or exclude obsolete historical data without modifying the source database views:

1. Open the table context menu on the canvas and choose **Table Filtering**.
2. Enter a standard SQL WHERE clause expression.




Table Filtering

Filters


[OrderDate] >= '2025-01-01' AND [Status] = 'Completed'

TEST OK CLOSE


3. **Requirement:** Always enclose column names in square brackets.
 - *Example:* [OrderDate] >= '2025-01-01' AND [Status] = 'Completed'
4. Click **OK** to save the boundary condition.

3.7 Mapping & Formatting Measures

1. In the Model Structure Panel, highlight the **Measures** area (indicated by a light-blue accent). This means that this structure is selected and ready to accept new items.




2. Check the box next to any numeric column you want to aggregate (e.g., ExtendedPrice, Quantity, DiscountAmount).



3. In the right-hand **Model Structure** panel, click on the newly added measure to open its properties:


- **Caption:** Provide a clean business title (e.g., "Total Revenue").


- **Description:** Add contextual tooltips explaining how the metric is calculated.
- **Format String:** Select a standard formatting pattern (e.g., Currency #,##0.00 \$, Percentage 0.0%, or Integer #,##0) or specify a custom mask.




3.8 Creating Dimension Entities & Mapping Attributes

1. In the right-hand **Model Structure** panel, open the menu and select **New Dimension** (e.g., "Customer").






2. Select the newly created dimension folder in the tree (highlighted in light blue).



3. On the canvas, navigate to the related dimension table and check the columns you wish to expose as analytical attributes (e.g., CustomerName, City, CustomerType).



- Click any mapped attribute in the structure tree to refine its **Caption**, Description or field length.

The image shows a configuration dialog for the 'MaritalStatus' attribute and a sidebar showing the model structure.

MaritalStatus

Table	dbo.DimCustomer
Primary Key	CustomerKey
Source Column	MaritalStatus
Related Table	dbo.FactInternetSales
Related Column	CustomerKey
Data Type	Text
Caption	Marital Status
Description	
Length	255
Usage	Dimension Attribute
Hierarchy	

Buttons: REMOVE, CLOSE, OK

Model Structure

- First Semantic Model
 - Measures
 - Customer
 - MaritalStatus
 - Gender
 - EmailAddress

3.9 Automated Time Dimensions & Date Hierarchies

Kyubit features automatic Time Intelligence for date columns:


- On any table containing a Date or DateTime column, check the date column.

The image shows a dialog for selecting columns from the 'dbo.FactInternetSales' table.


dbo.FactInternetSales

- ☐ Freight money
- ☐ CarrierTrackingNu...
- ☐ CustomerPONum...
- ☒ OrderDate dateti...

2. Kyubit automatically detects the date type and offers to create a complete and comprehensive **Time Dimension** or just to add a simple date field for data filtering.




3. The generated structure automatically provides:
 - Multi-level date hierarchies: **Year → Quarter → Month → Day**
 - Standard analytical attributes: *Quarter Name, Month Name, Day of Month, Day of Week, Week of Year.*




3.10 Building Multi-Level Attribute Hierarchies

Hierarchies enable drill-down pathways (such as *Product Category* → *Subcategory* → *Product Name* or *Country* → *State* → *City*).


1. In the **Model Structure** panel, right-click the parent dimension and select **New Hierarchy**.



2. Assign a business name to the hierarchy (e.g., "Geographic Location").



3. Drag and order the dimension attributes into the hierarchy levels to define the sequential drill path.



3.11 High-Cardinality Dimension Details (Drillthrough)

Columns containing mostly unique text values—such as CustomerEmail, StreetAddress, TrackingNumber, or PhoneNumber—introduce substantial indexing overhead and can noticeably degrade semantic model performance when configured as standard dimension attributes.

Because aggregating data by unique descriptors like street addresses or phone numbers offers little analytical value, Kyubit includes an automated safeguard: **by default, if a column exceeds 80% distinct values, model processing will stop with a "High-Cardinality" warning,** recorded in the processing log.

However, these granular fields are often essential when inspecting transactional line items. When handling high-cardinality columns, you can configure their **Usage** setting in one of two ways:

1. **Dimension Details (Recommended):** Routes the field into lightweight drillthrough storage. It is excluded from heavy aggregation indexing, keeping the model fast and compact while preserving the data for instant display when drilling through to leaf-level records.

2. **Dimension Attribute (Ignore High-Cardinality):** Overrides the safeguard and forces the field into the primary dimensional hierarchy, regardless of the associated performance and indexing overhead.

The screenshot displays the 'Email Address' dimension attribute configuration in SSDT. The 'Usage' field is set to 'Dimension Details', and a red arrow points to the 'Dimension Attribute (Ignore High Cardinality)' option in the dropdown menu. The 'Model Structure' pane on the right shows the hierarchy: First Semantic Model > Customer > Email Address.

Email Address	
Table	dbo.DimCustomer
Primary Key	CustomerKey
Source Column	EmailAddress
Related Table	dbo.FactInternetSales
Related Column	CustomerKey
Data Type	Text
Caption	Email Address
Description	
Length	255
Usage	Dimension Details
Hierarchy	Dimension Attribute Dimension Details Dimension Attribute (Ignore High Cardinality)
REMOVE	

Model Structure

- First Semantic Model
 - Measures
 - Customer
 - Marital Status
 - Gender
 - Email Address
 - OrderDate
 - Products

- Map these fields as **Dimension Details** instead of attributes.
- **Result:** The model processes faster, consumes significantly less storage, and makes these granular details instantly accessible when a user performs a **Drillthrough to Details** action on an aggregated report cell.
- To explicitly ignore high-cardinality warning, set the **Usage** field to **Dimension Attribute (Ignore High-Cardinality)**

3.12 Value Bucketing with Dimension Attribute Ranges

When analyzing continuous numeric variables (such as *Customer Age*, *Tenure Days*, or *Order Sizes*), distinct values clutter visual charts and pivot tables.

1. Click on the numeric dimension attribute in the structure tree and select **Attribute Ranges**.
2. Define custom value buckets:
 - *Example 1:* Range 18–29 (Min: 18, Max: 29)
 - *Example 2:* Range 30–49 (Min: 30, Max: 49)
 - *Example 3:* Range 50+ (Min: 50, Max: 120)
3. When processed, the dimension groups data into these distinct ranges, making demographic and bracket analysis clean and intuitive.

Ranges

Name	Min	Max
Single day	>= 0	< 2
Few days	>= 2	< 5
Many days	>= 5	< 100

ADD

OK

CLOSE

4. Designing Single-Dataset Models (Excel, CSV, and Direct SQL Queries)

For file-based uploads or custom single-statement SQL queries where a visual multi-table diagram is not applicable, Kyubit provides the streamlined **Single-Dataset** definition mode.

Measure	Measure	Measure	Date	Date	Dimension	Dimension	Dimension
SalesAmount \$#,##.##	TaxAmt \$#,##.##	Product Standard Cost €#,##.##	OrderDate	ShipDate	Color	Product Name	ModelName
\$3,374,99	\$270	€1.898,09	30.12.2028. 0:00:00	6.1.2029. 0:00:00	Black	Mountain-100 Black, 48	Mountain-100
\$3,399,99	\$272	€1.912,15	30.12.2028. 0:00:00	6.1.2029. 0:00:00	Silver	Mountain-100 Silver, 38	Mountain-100
\$3,578,27	\$286,26	€2.171,29	1.1.2029. 0:00:00	8.1.2029. 0:00:00	Red	Road-150 Red, 56	Road-150
\$3,578,27	\$286,26	€2.171,29	2.1.2025. 0:00:00	9.1.2025. 0:00:00	Red	Road-150 Red, 56	Road-150
\$3,578,27	\$286,26	€2.171,29	2.1.2029. 0:00:00	9.1.2029. 0:00:00	Red	Road-150 Red, 56	Road-150
\$3,578,27	\$286,26	€2.171,29	2.1.2025. 0:00:00	9.1.2025. 0:00:00	Red	Road-150 Red, 44	Road-150
\$3,578,27	\$286,26	€2.171,29	2.1.2029. 0:00:00	9.1.2029. 0:00:00	Red	Road-150 Red, 44	Road-150

4.1 Excel and CSV File Uploads

- **Prerequisites:** Data must be structured in rows and columns, with the first row containing column headers. Numeric and date formats must match the Kyubit server regional settings. CSV files should use standard semicolon (;) or comma separation.
- **Upload Procedure:** Create a new Analytic Model, select **File** input, save the model record, and upload the source file. The file is saved directly into the model configuration.

4.2 SQL Query Datasets

- Select **Query** input, pick an existing registered SQL/ODBC Data Source, and input your custom SQL query.
- Click **Run Query** or switch to the **Data** tab to verify the preview results.

4.3 Column-Based Data Definitions

On the **Data** tab, configure each column's role from the dropdown:

- **Measure:** Numeric values for aggregation. Click *Format Values* to set currency or numeric formatting masks.
- **Dimension:** Categorical columns for row and column series grouping.
- **Date / Date Dimension:** Check *Create Date Dimension* to auto-generate the complete Year-Quarter-Month-Date hierarchy.
- **Details:** Non-aggregated descriptive fields reserved for drillthrough inspection.

4.4 Structure Organization & Sorting Keys

- **Organize Dimensions:** On the **Structure** tab, move attributes from the "Default Dimension" into logically grouped parent containers and create custom multi-level hierarchies.
- **Key Sorting Column:** If an attribute requires non-alphabetical sorting (such as month names), assign a numeric order column as its **Key Column** (e.g., assign MonthNumber to sort MonthName). In analyses, selecting *Level Sorting* → *Member Key* will sort the values chronologically.

5. Model Processing, Lifecycle & Automation

5.1 Processing States & Compilation

Once the semantic model structure is configured, click **Process Analytic Model** on the **Processing** tab to build the data layer and make it ready for analysis.

**Full Process**

Process all rows from the Fact table into the Analytic Model structure.



**Partial Process**

Process rows within a defined relative date period from the Fact table.



**Incremental Process**

Process only new rows added since the last successful processing.



CLOSE

Processing Mode	Scope	Typical Execution Window	Database Impact	Best For
Full Process	Rebuilds all fact tables, dimension tables, ColumnStore indexes, and relational links from scratch.	Monthly, quarterly, yearly.	High (reads all source data).	Major schema updates, historical data revisions, or weekly maintenance.

Processing Mode	Scope	Typical Execution Window	Database Impact	Best For
Partial Process	Updates the semantic model by refreshing only fact table records within a designated rolling time frame (e.g., Last 30 Days).	Weekly, monthly	Low to Moderate	Refresh recent data with new entries.
Incremental Process	Ingests only new transactional records based on a high-watermark key or date filter.	High-frequency intervals (e.g., every 15–60 minutes).	Minimal (reads only delta records).	High-volume fact tables (e.g., daily sales, POS transactions, continuous telemetry data).

5.1.1. Full Model Processing

A **Full Process** operation completely drops and reconstructs the internal ColumnStore structures inside the KyubitAnalyticModels database.

- **When to Use:**
 - Initial deployment and compilation of a newly designed semantic model.
 - Structural model changes (e.g., adding/modifying table joins, adding new hierarchies, altering column data types).
 - Routine periodic maintenance (e.g., over the weekend) to rebuild indexes and capture retroactive historical data corrections.
- **Operational Behavior:**
 - Reads the entire historical range defined in the model's source queries.
 - Rebuilds all dimension tables, attribute ranges, and fact metrics simultaneously.
 - Updates the model state to Processed upon successful compilation.

5.1.2. Partial Processing (Rolling Time Frame Refresh)

In many business environments, historical transactional data (e.g., sales from previous years) remains static, while records within the current operating period (e.g., the last 30 days) may receive status updates, reconciliations, or adjustments.

Partial Processing updates the semantic model by selectively reloading fact data that falls within a designated rolling time frame, leaving older historical data intact.

- **How It Works:**

1. Navigate to **Processing → Process Settings** and set the **Anchor Date Column** to your fact table's primary transactional date field. After configuring the **Anchor Date Column** for the first time, you must run a Full Process once before Partial Processing becomes available for selection.
2. If the Kyubit host and database server clocks are not synchronized, specify the offset in **Data Source Time Difference**.
3. Under **Schedule → Jobs**, set the job action to **Partial Process** and define the rolling window (e.g., Last 30 Days).
4. The engine queries and refreshes only records matching the anchor date boundary, preserving older historical partitions without a costly full-table rebuild.

5.1.3. Incremental Processing (Delta Ingestion)


Incremental Processing provides high-speed, continuous synchronization for large datasets by capturing and loading only new transactional records generated since the previous execution cycle.

- **How It Works:**

1. Under **Processing → Process Settings**, designate the **Incremental Key Column** using a unique, auto-incrementing identity field. After configuring the **Incremental Key Column** for the first time, you must run a Full Process once before **Incremental Processing** becomes available for selection.
2. If an identity column is not available, leave the field blank and ensure the **Anchor Date Column** is selected.
3. The scheduler checks the watermark from the previous run and ingests only higher sequential keys or newer timestamps, merging them directly into the model's fact storage.

A model cycles through three operational states:

1. **Unprocessed:** The model definition is saved or modified, but data structures are not compiled. It cannot be used in analytics yet.
2. **Processing:** The engine is extracting data, building indexes, and establishing relationships. The model is temporarily locked.
3. **Processed:** The model is compiled, optimized, and ready for queries, reports, and dashboards.



5.2 Processing Logs & Troubleshooting

If processing fails:

1. Open the model and navigate to the **Processing** tab. Each processing session will be displayed in a separate row.

Analytic Model	Model Designer	Processing	Schedule
Status	Start Time		
✓ Processing Successful	20.8.2026. 11:17:25		
⚡ Processing Failed	20.8.2026. 11:05:06		
⚡ Processing Failed	20.8.2026. 11:01:05		
⚡ Processing Failed	19.8.2026. 15:59:12		
✓ Processing Successful	8.7.2026. 10:19:47		

2. Click the **processing session** (row) to view detailed execution logs, step timings, warnings, type mismatches, missing key relations, or query syntax issues.

Processing - Log			
📊	Level table created [dbo_DimProductSubcategory_EnglishProductSubcategoryName].	testuser1	20.8.2026. 11:17:47
✦	Level foreign values updated [dbo_DimProductCategory_EnglishProductCategoryName].	testuser1	20.8.2026. 11:17:47
📋	Level values inserted [Product Category].	testuser1	20.8.2026. 11:17:44
📊	Level table created [dbo_DimProductCategory_EnglishProductCategoryName].	testuser1	20.8.2026. 11:17:44
✦	Bulk Stream Copy Full successfully done.	testuser1	20.8.2026. 11:17:44
📊	Model table successfully created.	testuser1	20.8.2026. 11:17:25
	Process structure prepared.	testuser1	20.8.2026. 11:17:25
✦	Full data import process started (ModelID = 20).	testuser1	20.8.2026. 11:17:25
✓	Source data types and integrity checked successfully	testuser1	20.8.2026. 11:17:25
✓	Source values include all defined column names.	testuser1	20.8.2026. 11:17:25
📌	Processing started (ModelID = 20).	testuser1	20.8.2026. 11:17:25
COPY CLOSE			

5.3 Automated Updates with Scheduled Jobs

Keep semantic models synchronized with live transactional databases automatically:

Configuring Scheduled Jobs in Kyubit

1. Navigate to **Schedule** → **Jobs** in the top navigation bar.
2. Click **New Scheduled Job** and enter a descriptive job name (e.g., Daily Sales Partial Refresh - Last 30 Days).
3. Set the **Action Type** to **Process Analytic Model**.
4. Select your target **Analytic Model** and choose the desired **Processing Type** (*Full*, *Partial*, or *Incremental*).
5. If **Partial Process** is selected, define the target rolling window (e.g., 30 Days).
6. Define the **Execution Schedule**:
 - **Frequency**: Daily, Weekly, Hourly, or custom recurrence intervals.
 - **Start Time & Timezone**: Align execution times with production ETL completion and off-peak periods.
7. Configure **Alerts & Failure Notifications**:
 - Specify recipient email addresses to receive instant notifications with if a processing fail. Set “JobStatusEmailNotify” in the web.config file.

Analytic Model

Model Designer

Processing

Schedule

Title	Occurrence	Time
 Partial Processing, Last 7 days	Weekly	08:00

5.4 Sizing Guidelines & Performance Best Practices

- **Optimal Range**: Kyubit Semantic Models are optimized for datasets ranging from thousands of rows up to **10 million rows** in the Fact table (depending on the SQL Server hardware).

- **Large Datasets:** Performance with millions of records depends on the SQL server CPU and RAM. For massive multi-billion-row enterprise workloads, dedicated OLAP / SSAS cube connections should be considered.

- **Performance Best Practices:**
 - Map unique identifiers, emails, and notes as **Dimension Details** rather than Dimensions.


 - Use **Table Filtering** to exclude historical data beyond business reporting windows.

 - Limit the active series columns on pivot grids to what can be legibly rendered (Kyubit supports up to 128 active series columns per measure).

6. Data Analysis, Reporting, Dashboards & KPIs

6.1 Interactive Grid & Pivot Analysis

Once a model is processed, users can launch an **Analysis** session:



- **Drag-and-Drop Navigation:** Drag measures into the values area, and drag dimensions or hierarchies onto the **Categories (Rows)** and **Series (Columns)** axes.
- **Analytical Gestures:** Perform drill-down (expand levels), drill-by (switch dimension branches), sorting, top-N filtering, and isolating values.

6.2 Advanced Date Period & Relative Filtering

Kyubit Semantic Models provide flexible date slicing that simplifies date handling:

New Analysis

GRID CHART REPORT

Gender F, M OrderDate

Gender

F

Product Category Product Subcategory Sales Amount

Product Category	Product Subcategory	Sales Amount
Accessories		\$699,133,68
Bikes	Mountain Bikes	\$9,469,734,13
	Road Bikes	\$11,653,283,17
	Touring Bikes	\$3,875,490,66
	Total	\$24,998,507,96
Clothing		\$333,626,94
Total		\$26,031,268,58

Date Range
Select a fixed or relative date range

Date Range Year-To-Date

Selected: 2026-01-01 to 2026-08-26 (238 Days)

< JANUARY 2026 > < AUGUST 2026 >

SU	MO	TU	WE	TH	FR	SA	SU	MO	TU	WE	TH	FR	SA
				1	2	3							1
4	5	6	7	8	9	10	2	3	4	5	6	7	8
11	12	13	14	15	16	17	9	10	11	12	13	14	15
18	19	20	21	22	23	24	16	17	18	19	20	21	22
25	26	27	28	29	30	31	23	24	25	26	27	28	29
							30	31					

TD YD 3D 7D 10D WTD 1W 2W PW

MTD 1M 3M 6M PM QTD PQ YTD PY

1Y 2Y 4Y 8Y MAX

CUSTOM OK CLOSE

- **Absolute Date Filtering:** Select exact start and end calendar dates using the date picker dialog.
- **Relative Date Filtering:** Slices data dynamically relative to the current execution date.
 - Last 3 Years (without current year): Enter 3.
 - Last 3 Years (including current year to date): Enter 3+.
 - Similar relative expressions apply to days, weeks, and months.

6.3 Dynamic On-the-Fly Calculated Measures

Business users often require ad-hoc calculations (e.g., profit margin percentages, growth rates) that were not configured during initial data modeling.

The screenshot shows the 'User Calculated Measure' dialog box. On the left sidebar, under 'Measures', 'Net Sales' is selected, indicated by a red arrow. The main panel shows the measure name 'Net Sales', the available measure 'Sales Amount', and the expression `[dbo_FactInternetSales_SalesAmount] - [dbo_FactInternetSales_TaxAmt]`. The format is set to `##.##0`. Buttons for 'ADD MEASURE', 'TEST EXPRESSION', 'SAVE', 'DELETE', and 'CLOSE' are visible.

1. In the Analysis Grid view, click **Options** → **Create User Calculated Measure**.
2. Enter a measure name and write the arithmetic expression referencing existing measures.
3. **Preventing Division by Zero:** Always wrap divisor measures in NULLIF:

$$\text{Margin \%} = \frac{[\text{Revenue}] - [\text{Cost}]}{\text{NULLIF}([\text{Revenue}], 0)} \times 100$$

4. Supported operators and functions:

- **Arithmetic:** +, -, *, /, %
- **Math Functions:** ABS, ROUND, FLOOR, EXP, LOG, SIGN, SIN, COS, RAND, DEGREES

5. Click **Test Expression** to validate syntax, apply desired number formatting, and click **Save**.

6. The calculated measure appears instantly in the model tree for immediate use.

The screenshot displays the 'Adventure Works Semantic Model' interface. On the left, a 'Measures' list includes 'Sales Amount', 'Total Product Cost', 'Tax Amt', and 'Net Sales' (which is selected). Below this is a list of dimensions: 'Product', 'Customer', 'OrderDate', 'Promotion', and 'Currency'. On the right, a pivot table is shown with 'Gender' (F, M) and 'OrderDate Year-To-Date' as filters. The table has columns for 'Product Category', 'Product Subcategory', and 'Net Sales' (calculated as 'Net Sales'). The data is summarized by 'Bikes' (Mountain Bikes, Road Bikes) and a 'Total' row.


		Gender		
		F	M	Total
Product Category	Product Subcategory	Net Sales	Net Sales	Net Sales
Bikes	Mountain Bikes	578.016,77	520.907,09	Σ 1.098.923,86
	Road Bikes	1.116.788,83	1.072.150,13	Σ 2.188.938,96
	Total	Σ 1.694.805,60	Σ 1.593.057,22	Σ 3.287.862,82
Total		Σ 1.694.805,60	Σ 1.593.057,22	Σ 3.287.862,82

6.4 Designing Interactive Dashboards & KPI Widgets

Reports and analyses created from semantic models serve as building blocks for dashboards:

- **Dashboard Assembly:** Drag and drop saved analyses onto the responsive dashboard grid.
- **Visual Components:** Switch views between Pivot Tables, Bar Charts, Line Charts, Area Charts, Pie Charts, Bubble Visualizations, and KPI Indicators.

- **KPI Metrics:** Set target values, threshold bands (Red/Yellow/Green), and comparative periods (e.g., Actual vs. Budget) directly within KPI card elements.
- **Interactive Filtering:** Place global dashboard slicers (such as Date Pickers or Dimension Selectors) that automatically filter all connected dashboard widgets simultaneously.



6.5 Permissions, Governance & Collaboration

- **Role-Based Permissions:** Grant View or Manage access to individual models, analyses, and dashboards for specific users or security groups.
- **User Calculated Measure Governance:** Calculated measures created by an individual user remain private by default, but can be published and shared across teams via the **Shared Items** section.