

Article

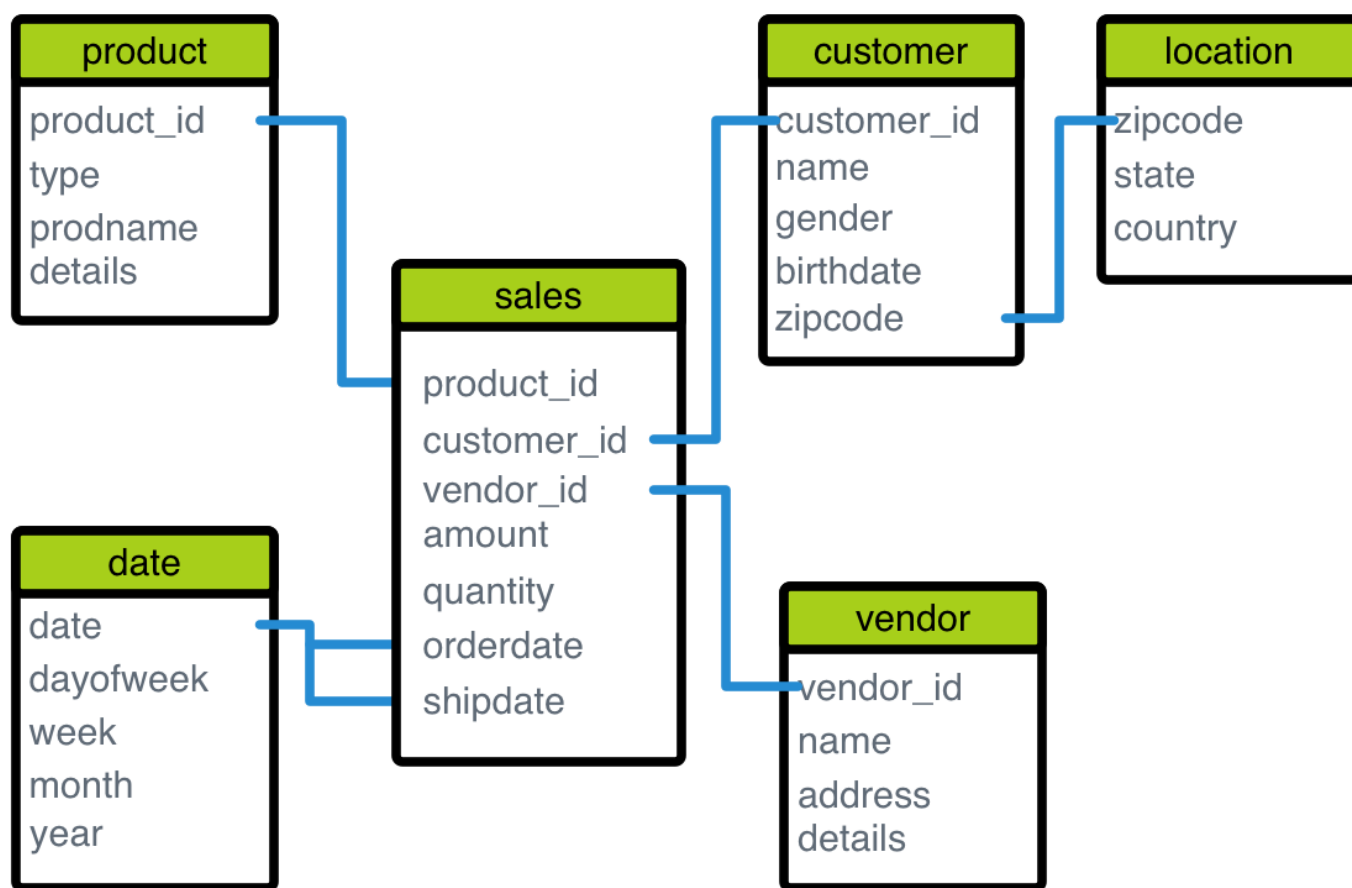
[Evgeniy Potapov](#) · Aug 4, 2022 6m read

## Benefits of Adaptive Analytics powered by AtScale + InterSystems Report Designer Powered by Logi Reports

AtScale pulls data from the IRIS database.

The AtScale product forms a virtual OLAP cube on the intermediate layer, which external applications can access using standard SQL and MDX (Multidimensional Expressions) languages. The solution includes three main components.

Design Center is used for designing OLAP cubes, forming links between metadata and dimensions of a virtual cube. Along with the task of designing a data schema, the issues of access policy to certain data and security are also solved here. Since Virtual Cube does not physically store Big Data, ensuring acceptable performance is a serious problem.

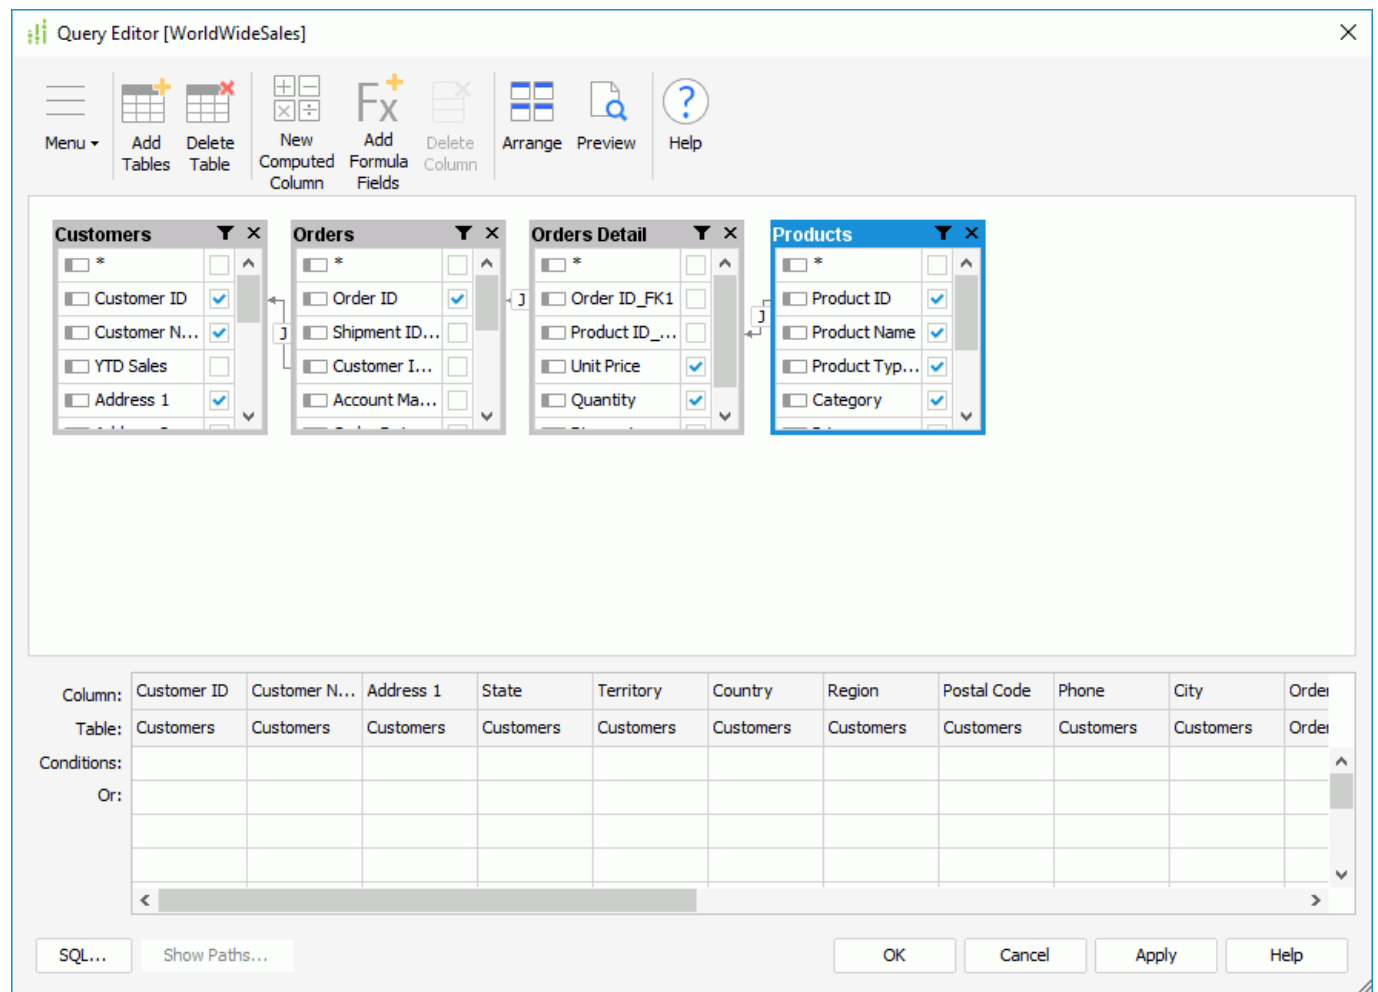


Adaptive Cache makes it possible not only to physically cache recent or frequently used data but also to predict what data will be needed soon so that it can be prefetched into the cache.

To connect Logi to AtScale, use the Hive2 JDBC connection type with the org.apache.hive.jdbc.HiveDriver driver.

The advantage of this type of connection is the addition of an unlimited number of cubes to one Catalog. This makes it possible to create complex reports with a large dataset. Also, it is possible to create a diagram inside Logi based on cubes and to make links between cubes from AtScale, which expands the possibilities of data generation. It is also possible to manually construct SQL queries for cubes. This feature is not used very often, but it can be

employed to manually establish relationships or design SQL formulas.

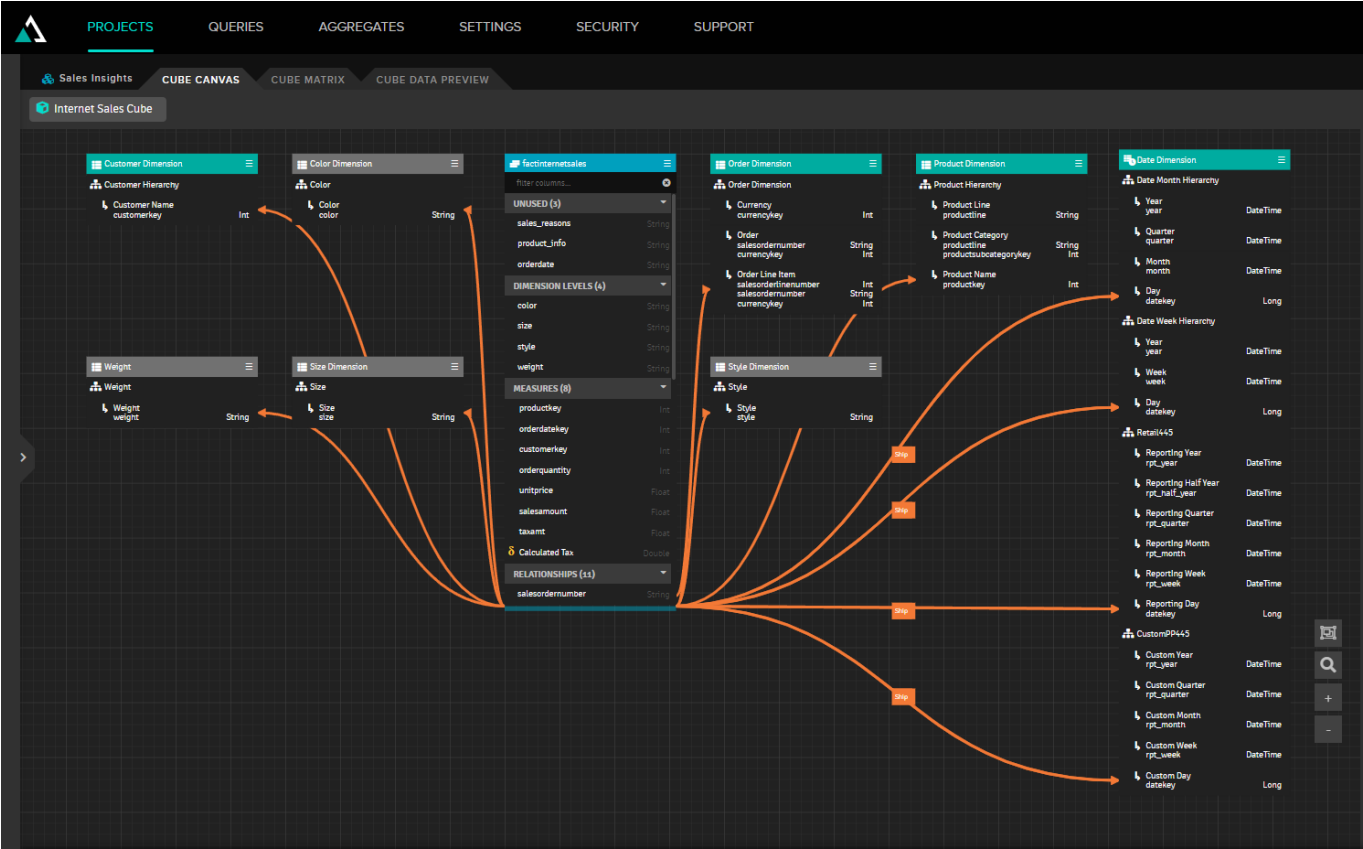


Most of the logic for generating a data set is performed on the AtScale side. Logi (InterSystems Report Designer) accepts the generated data sets and visualizes this data.

Sales Insights Application test dataset (AtScale versions >= 2019.2.x)


“ Gain insight into sales trends, product performance, and customer details using this sample data based on the AdventureWorks data set. ”


It can be downloaded from the link <https://downloads.atscale.com/>





This is what this data set looks like in InterSystems Reports Designer


## Catalog Manager [Unknown.cat]


  
New Data Source


  
Save Catalog


  
Paste


  
Edit


  
Copy


  
Delete


  
Rename

  
Configure Reference

  
Refresh Reference

  
Export Reference

  
Options

  
Data Source Driver

**Data Resource** A Z C

☒ Data Source 1

☒ jdbc:hive2:// :11111/Sales%20Insights

☐ Tables

☐ internet sales cube

☐ Average Customer Count per Order

☐ Average Est Customer Count per Order

☐ Average Last Product Unit Count per Order

☐ Average Max Tax Count per Order

☐ Average Sales Amount SD Count per Order

☐ Average Sales per Order

☐ Average Sold Product per Order

☐ calculatedtax1

☐ City

☐ Color

☐ CountryCity

☐ CountryZip

☐ currency

☐ Customer Count-Order CustomPP445-PrevYear

☐ Customer Count-Order CustomPP445-PrevYearGrowth

☐ Customer Count-Order CustomPP445-PrevYearGrowthPct

☐ Customer Count-Order Retail445-30PrdMvAvg

☐ Customer Count-Order Retail445-30PrdMvAvgPrevYear

☐ Customer Count-Order Retail445-30PrdMvAvgPrevYearGrowthPct

☐ Customer Count-Order Retail445-30PrdMvLowerBand

☐ Customer Count-Order Retail445-30PrdMvStdev

☐ Customer Count-Order Retail445-30PrdMvUpperBand

☐ Customer Count-Order Retail445-Next

☐ Customer Count-Order Retail445-Prev

☐ Customer Count-Order Retail445-PrevPeriodGrowth

☐ Customer Count-Order Retail445-PrevPeriodGrowthPct

The main advantages that we have received from using this bundle are:

1. Possibility of automatic generation of reports according to the schedule
2. An option of receiving PDF reports by email
3. An opportunity to use parameters automatically calculated for the report (dates)
4. Acceleration of report generation (UDAF, aggregates and caching)
5. Ability to clone, backup, and quickly propagate AtScale cube changes
6. Logi functionality extension: some of the comparative parameters were not previously calculated in Logi and DeepSee
7. Substitution of Logi functionality: calculation of comparative indicators depending on the parameter settings on the AtScale side

## Why UDAF is needed for Adaptive Analytics and how to configure it

InterSystems Adaptive Analytics (powered by AtScale) is a powerful addition to IRIS BI. It allows users to get their Dashboards built faster. Also, Developers and Maintainers can quickly switch between data sources. And last but not least are snapshots of Data Cubes versions. They give you the ability to backup data structure logic and rollback if needed.

Some of the functions are ready to use out-of-the-box. However, one important option needs to be configured. I mean an option that gives a boost to Dashboards and Reports speed.

It is called UDAF. If it is not configured, you can not see any optimisation or improvement of the abstraction layer mechanism.

You can get UDAF distribution from the same place you get Adaptive Analytics one.

UDAF gives Adaptive Analytics 2 main advantages:

- the ability to store query execution results (they call it Aggregate Tables), so that the next query, using aggregation on data, could take already pre-calculated results from the database.
- the ability to use additional functions (a.k.a User-Defined Aggregate Functions) and data processing algorithms that Adaptive Analytics is forced to store in the data source.

They are stored in the database in separate tables, and Adaptive Analytics can call them by name in auto-generated queries. When Adaptive Analytics uses these functions, the speed of queries increases.

The place where all these tables will be stored is specified when creating a connection to IRIS in the AGGREGATE SCHEMA item. If no such schema exists in the database, it will be assembled. It is a good idea to store that schema in another database. Later I will explain you why.

## Add Data Warehouse

### Type of Data Warehouse

InterSystems IRIS ▼

### General

NAME ?

EXTERNAL CONNECTION ID

Only change this if you need it to match imported project xml

☒ Override generated value?

NAMESPACE

AGGREGATE SCHEMA ?

AtScale ✓

Now let's imagine that IRIS, to which Adaptive Analytics is connected, is intended only for analytics and contains only a copy of data from a running database (so that requests from analytical systems do not load the resources of the main database). Once in a certain time, the data from the main system is copied to our IRIS, overwriting the old data.

In such a situation, at the time of rewriting, all the data recorded by Adaptive Analytics disappears, and we get the following result:

- queries that are trying to access aggregate tables do not find them and fall with an error.
- requests that use the functions stored in the database cannot access them and fall with an error.
- when deleting old paths to aggregate tables, Adaptive Analytics quietly creates new ones, and later the situation repeats itself.

The above described is only one of the possible cases of overwriting data in IRIS but by no means the most

common one.

The main problem is that Adaptive Analytics does not use fallback mechanisms of avoiding using UDAF options and making direct queries in case of such errors. It fails to report building or dashboard updates.

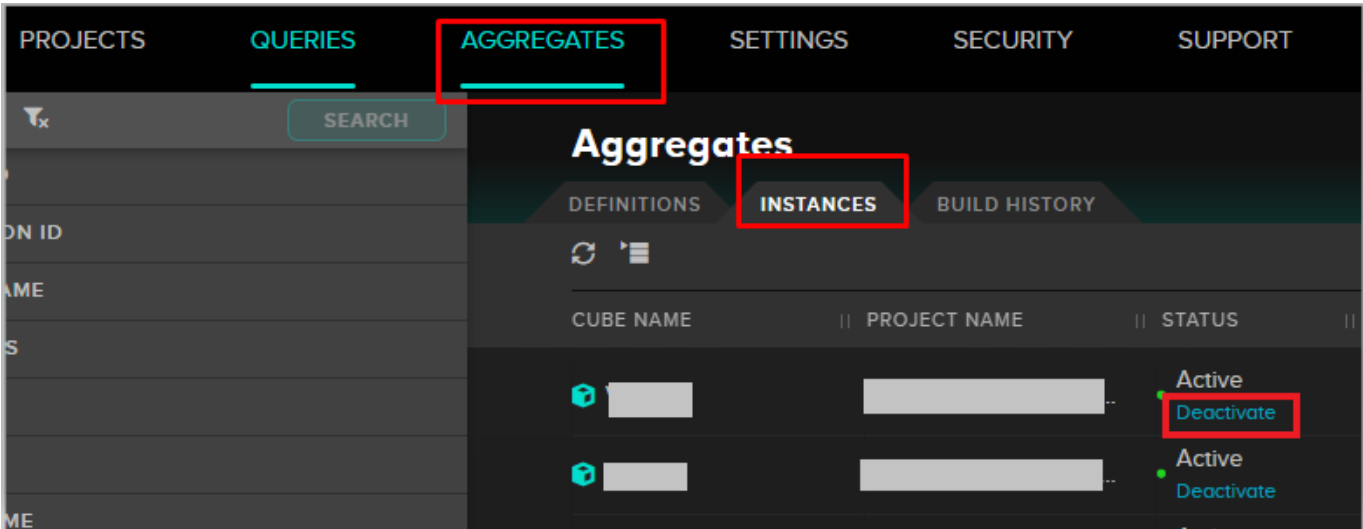
We can try to load the table with functions into the database along with the update manually. However, as mentioned above, the UDAF is not only the data processing functions but also the aggregate tables, and resaving them is also quite problematic.

The solution to the problems described above is the creation of a separate database where Adaptive Analytics will write its service tables. Such a solution is described in the documentation for connecting Adaptive Analytics to IRIS:

<https://docs.intersystems.com/irislatest/csp/docbook/DocBook.UI.Page.cls?KEY=AADAN#AADANconfig>

It's worth mentioning that the schema name you specify after "/<instancePath>/mgr/" must match the schema name you specify as AGGREGATE SCHEMA with Adaptive Analytics. Otherwise, a separate database will be created, but Adaptive Analytics will write data to the schema it created inside your database, ignoring the newly made separate database. The UDAF developers recommend using "AtScale" as the name of such a scheme. It also the name given as an example in the documentation.

In the event that this database becomes corrupted and the data will be impossible to recover, you can disable the aggregations in Adaptive Analytics to automatically recreate them on the next request.



To do this, in the Aggregates/instances tab, you must deactivate each aggregate. Aggregates are displayed 20 per page, so be sure to visit each page to verify they are all deactivated.

So if you are going to use Adaptive Analytics, please, pay attention to that small point. At the start of my work with Adaptive Analytics, this option was set up with a mistake, so I didn't get all the benefits it could give. Also, remember to check your database name. It is case-sensitive in setting and in the script of installation of UDAF.

I think there might be some options which I don't know about at the moment, as well as their influence on the performance of my analytics systems. If you find any, please write about them in the comments.

[#Adaptive Analytics](#) [#InterSystems Reports \(Logi\)](#) [#InterSystems IRIS](#)