

SQL Server Analysis Services (SSAS) to Modern Cloud OLAP Migration: A Practical Guide








TABLE OF CONTENTS

The Rise and Rise of Data Modeling	1
SSAS: The Most Pervasive OLAP Platform	2
Traditional SSAS Methodology Limitations	4
COLAP (Cloud OLAP) to the Rescue	5
Before Getting Started	7
SSAS to Modern Cloud OLAP Migration Checklist	8
Futureproof Your Data Infrastructure with AtScale	10





The Rise and Rise of Data Modeling

Growth and business continuity are the most important things that every organization has to worry about today. These can only be achieved by gaining a holistic view of the markets and understanding the target audience better. Enter Data Modeling, a game changer in the Business Intelligence (BI) space.

Before jumping into the technical aspects of this dynamic field, it's important to understand the basics. Data Modeling involves the mapping, digesting, and conceptualizing of information that comes from various data streams and warehouses. The business then generates actionable insights from this information.

There are three main aspects of Data Modeling that you should always factor into your final design - Conceptual (Enterprise), Logical, and Physical.

Data Modeling Design Components

Conceptual

Here the primary business needs are defined and the most significant entity relations are identified

Logical

This approach is more complex as it goes deeper to establish specific metrics and business functions

Physical

This is basically the blueprint of the Data Model design you will finally implement and work off repeatedly

Data Analysis just isn't enough anymore. The large amount of data available today, along with increased competition and fluctuating customer behaviour patterns, is pushing more and more businesses into the Data Modeling realm. Only a birds-eye view of your sector can help you achieve sustainable growth.



SSAS: The Most Pervasive OLAP Platform

In a nutshell, SQL Server Analysis Services (SSAS) is one of the most commonly used Online Analytical Processing (OLAP) and data mining tools today. It's inherited advantages and benefits have been extensively documented (and proven) over the years in multiple use-case scenarios.

There are many variations of OLAP out there right now, with ROLAP (Relational OLAP) and MOLAP (Multidimensional OLAP) being the prominent ones.

Here are some of the key SSAS characteristics:

- **Modes** - Native mode and SharePoint mode
- **Methodologies** - Tabular (better for specific analytics) and Multidimensional (better for corporate scale analytics)
- **Technologies** - A blend of server and client technologies
- **Technical** - SSAS is usually installed as a VM server instance or as an on-premise model to slice and dice huge chunks of data

Tabular and Multidimensional SSAS models differ in the way they operate, perform, and behave. Here are the biggest differences between the two.

Tabular SSAS Model

Uses tables for faster execution. It is significantly easier to use, develop, and maintain

Has an in-memory cache and reads directly from it

Can interpret DAX and MDX reports

A relational model that works best when data is simple

Offers support for two data access modes - DirectQuery and Cached

Multidimensional SSAS Model

Executes as a group of operations to fetch query results via cube or dimension coordinates

File-based storage system via OLAP cubes. Reads pre-aggregated data

Works with MDX reports only

A dimensional model that is good for complicated data in bulk

Offers support for two data access modes - MOLAP and ROLAP



There are some tried and tested techniques to get the most out of your SSAS models. One best practice is the defining efficient aggregations for larger data groups and maximizing the use of partitions. These aggregations should be created sparingly to minimize their effect on performance, which is key today.

Other “pro tips” include the pre-planned usage of the query engine cache, fine tuning of memory or processor usage, creating natural hierarchies to improve performance, and implementation of effective MDX. Furthermore, you must optimize the cube at all times and the group design should be measured too.

Unfortunately, SSAS is not getting the job done anymore due the exponential rise in data sources and the growing need for scaling up fast(er).

50% of companies are using data more often now than prior to the Covid-19 pandemic



Traditional SSAS Methodology Limitations

The problems with SSAS begin when you want to shift from one selection to another between Tabular and Multidimensional SSAS models. This will essentially require you to start over again, something that takes a lot of time and effort, not to mention the financial implications of such processes.

Furthermore, it's impossible to merge information between cubes of different types, which is a huge technical hurdle if you are looking to scale up fast.

There is also HOLAP (Hybrid OLAP), a fusion of the ROLAP and MOLAP methodologies. It was created to try and iron out their shortcomings by allowing the storage of data in both relational and multidimensional databases and a lot of caching to allow the use of data sets and complex queries on-demand.

But even HOLAP has its problems. Here are just a few of them.

- Data volumes often become too big to handle smoothly
- The extra storage requires additional investment in infrastructure
- “ROLAP-heavy” queries take a long time to process
- Every data update or deletion requires system processing
- The cache also need to be updated with every database change

All of these facts basically mean that traditional methodologies are basically becoming obsolete due to the current on-demand BI requirements.




COLAP (Cloud OLAP) to the Rescue

With more and more organizations realizing the potential of big data and BI initiatives, they are also finding it tough to meet their goals with traditional SSAS solutions. The huge investments required in computing power and storage databases are causing many problems.

This is why the world of BI is moving to the cloud to minimize investment in manpower and resources, maximize their performance metrics, and create as many actionable insights as possible. By using OLAP in the cloud, results are also much faster, with scaling up becoming smooth and error-free.

COLAP (Modern Cloud OLAP) is taking the world by storm. And rightly so.

This comprehensive solution basically allows you to take data virtualization to the next level by presenting a single data service view into multiple data warehouses, all with just a few clicks without worrying about engineering and operational requirements like with old SSAS technologies.



Here are just a few benefits of moving your OLAP to the cloud:

- Fast results even with complex multidimensional queries
- Unified and centralized solution for all crucial business metrics
- Simplified interface that can be used by multiple stakeholders
- Enhanced visibility for better cross-department collaboration

As mentioned earlier, COLAP gives you the ability to directly access any data warehouse or data lake, regardless of where they are located, on-prem or in the cloud. With loads of information piling up and the need to react faster to market developments, SSAS to Modern Cloud OLAP migration is trending up.



Furthermore, you can even use solutions like AtScale to elevate your COLAP performance by enhancing visibility, functionality, and ease-of-use.

If Amazon Redshift is your solution of choice, you can use AtScale to supplement two of Redshift's top data engineering based optimizations. These are Distribution Keys (for better parallelization) and Sort Keys (to skip irrelevant pages).

There is also a growing need to comply with Health Insurance Portability and Accountability Act (HIPAA), Sarbanes–Oxley Act of 2002 (SOX), General Data Protection Regulation (GDPR), and more. A centralized solution can hugely enhance your governance and monitoring capabilities.



Before Getting Started

It's also important to realize that SSAS to Modern Cloud OLAP migration is not just a technical hosting change. It's a massive shift on many levels, including the future-proofing of your on-premise data to work with new technologies that may emerge in the near future. More on this topic later.

Here are some crucial issues you need to factor into your migration process.

Security - Make sure your IT and Security teams are involved in the process. Get their inputs prior to the migration. They should be able to tell you if there are architectural flaws or limitations that may hamper your security or create issues with your upcoming GDPR, HIPAA, and SOX compliance audits.

You may also have to take care of applications and repoint them if they are generating data. Infrastructure reconfigurations may also be required.

Investment - You will probably be saving manpower and resources upon moving to the cloud. But you must act wisely in the cloud and pick the right options. It may be the better call to opt for Fine Grained data warehouse options because Coarse Grained data warehouses scale in larger units (costlier).

For example, scaling up Snowflake with larger data warehouses can eventually lead to excess unused capacity, that can exceed 25% of your billing cost.

Performance - Just like any other technological migration to the cloud, you may be looking at legacy applications or data that is still on-premise. Or you may be facing issues with migrating everything at once. It's extremely important to establish proper performance metrics and backups to be on the safe side.

It is also important to perform a thorough comparison between the various vendors to inspect the “performance for money” (bang for buck) aspect.



SSAS to Modern Cloud OLAP Migration Checklist

Now that we have covered the main reasons to make the move and touched on the prerequisites, let's dive into the SSAS to Modern Cloud OLAP migration checklist.

1. Load existing cube source file(s) to a cloud data warehouse of your choice. Keep in mind that there are many factors you need to consider before zeroing in on one that suits your budget and requirements.

<p>GOOGLE BIGQUERY</p> <p>The only entirely serverless product with best in class GIS support and integrated ML declarative language (BQML or BigQuery ML)</p>	<p>SNOWFLAKE</p> <p>This solution is extremely popular due to its ease-of-use and flexibility. Its unique architecture is optimized for seamless cloud compatibility</p>
<p>AMAZON REDSHIFT</p> <p>This proven and tested solution is considered by many to be the OG cloud data warehouse. Configuration and maintenance can be complicated</p>	<p>AZURE SYNAPSE ANALYTICS SQL</p> <p>This Microsoft solution is making its case with Synapse and Hyperscale offerings. It's based on the mature SQL Server IP (high compatibility)</p>

2. Migrate your SSAS models using a XMLA migration tool. You can also use a data mapping tool such as [Altova MapForce](#), which has a user-friendly interface, to gain added visibility into the entire process.

3. Port all existing reports and dashboards from your current installation to the new one. If you properly mapped your legacy cube XMLA to a Modern Cloud OLAP equivalent, all you need to do is change the existing connection strings and credentials.

4. Update and optimize data governance rules to reflect your latest business requirements and needs. With rules and regulations becoming stricter, you need to pay more attention to unstructured data.



5. Integrate the Single Sign-On (SSO) from your existing BI tool to your new cloud data warehouse. While you are at it, it may be a good idea to create new perspectives and enforce row level security rules.

6. Reconcile the old on-prem SSAS data with the new Modern Cloud OLAP data. You can do this by comparing key reports and dashboards. You can use JMeter to measure benchmarks and ensure optimal performance.

7. Decide on an appropriate schema for your cloud data warehouse. If you are opting for the Snowflake Schema, which is much more complex than the Star Schema, you need to make sure it's optimized for atomic queries.

Needless to say, you should also be aware of all cloud vendor options. Make a scaling up plan without sacrificing performance, while minimizing expenses.



Futureproof Your Data Infrastructure with AtScale

SSAS has probably served its purpose in the grand scheme of things. Modern Cloud OLAP is empowering businesses with the ability to shift their focus from cumbersome maintenance processes and tedious (re)programming efforts to what really matters - gaining the edge with actionable insights on demand.

The big data space is evolving fast. New requirements and technologies are being launched at a mind boggling pace, with regulatory requirements constantly changing. Having your data on the cloud can future-proof it and provide you with added peace of mind in the long run.

Also, once your data is in the cloud, you have the freedom to explore new technologies and elevate your performance with new solutions. As mentioned earlier, AtScale allows you to take your visualization, analysis, and reporting functionality to a whole new level, all in a user-friendly and intuitive way.

- **Performance** - In a market that requires fast time-to-market and optimal quality standards, you can now boost your query performance.
- **Security** - Complexity is reduced throughout the data lifecycle, allowing the maintenance of consistent data lineage for high security standards.
- **Compatibility** - A simplified and agile interface that can work with multiple BI and AI/ML tools can empower even the most inexperienced of users.

Harnessing the power of Modern Cloud OLAP has never been easier. Construct virtual OLAP cubes with AtScale over a data platform of your choice (Snowflake, Redshift, Azure SQL Server, etc.). Configure, publish, connect, and query away.

GOT Modern Cloud OLAP?

*Simplify and Accelerate Your Business Intelligence Ops to
Get a Better ROI on your Cloud Analytics Investment*

REQUEST A DEMO

ABOUT ATSCALE

AtScale powers the analysis used by the Global 2000 to make million dollar business decisions. The company's Intelligent Data Virtualization™ platform provides Cloud OLAP, Autonomous Data Engineering™ and a Universal Semantic Layer™ for fast, accurate data-driven business intelligence and machine learning analysis at scale.

For more information, visit www.atscale.com.