

ATSCALE

Delivering Speed of Thought Analytics for Cloud Scale Data

The role of a semantic layer in scaling business intelligence and
enterprise AI

Table of Contents

3	What's the Speed of Thought?
4	What is Dimensional Analytics?
6	The Practical Limitations of SSAS and Traditional OLAP
8	A Modern Approach to Analytics and Business Intelligence Acceleration
11	How AtScale Accelerates Analytics Performance & Enables OLAP Modernization?
13	SSAS Replacement
14	How Betclic Delivers Cloud-First Data Analytics
16	AtScale for Speed of Thought Analytics
17	Take the Next Step

What's the Speed of Thought?

Speed of thought analytics — in other words, analytics performance that keeps pace with the cognitive reasoning of the analyst — requires near-instantaneous queries across large data sets with consistent results. When this is accomplished, users can follow their cognitive curiosity, unearth relevant insights quicker, and have the freedom to explore the data to generate better business insights.

Application user experience design shows that user attention drops substantially after one second UI response times and gets exponentially worse as the seconds tick by. When creating an interactive analytics experience, whether in a Tableau or Power BI dashboard, or an Excel Pivot Table, the time to return results is key.

While data visualization platforms can be a buffer, analytics experience is primarily driven by query performance.

As data sets grow and get more complicated, queries consume more compute cycles and take longer times to execute. While modern cloud data platforms have nearly unlimited computational capacity, compute cycles are never free. Delivering high-performance analytics experiences has become increasingly difficult as the size of modern cloud data sets and then number of users has increased. Even the most powerful query engines cannot execute full table scans

of large, complicated data sets to consistently deliver speed of thought performance for dozens, hundreds or even thousands of users.

Poor analytics performance drives data teams and business users to find workarounds that inevitably result in disconnected data pipelines, shadow analytics programs, and metrics sprawl.

This eBook will discuss a modern approach to delivering speed of thought analytics tied to live cloud data sources.



What is Dimensional Analytics?

Dimensional analysis is a paradigm to organize data for analytics that business teams have relied on for decades. At the heart of dimensional analysis is the definition and standardization of hierarchical **Dimensions** (aka drill-down categories) such as Time, Geography, or Product. **Measures** (aka metrics or KPIs) are quantitative values that can be aggregated at different levels along the dimensional **Hierarchies** to form a “cut” of data (e.g. total sales of a product during March in Massachusetts).

The first key benefit of dimensional analytics is that it supports consistency and self-service. Business users are able to easily ask questions of a multidimensional dataset and get a consistent response back. Analysts can slice and dice the data using business concepts like time, region, product, and price without the need for technical expertise. The next eBook in this series will delve more deeply into the topic of self-service and the potential of enterprise metrics hubs.

The second benefit of this analysis paradigm is the ability to leverage the predictability that comes with constraining user queries to predefined dimensions and metrics. This predictability can be used to architect data and structure queries in ways that optimize performance. The term Online Analytical Processing (OLAP) emerged in the 1990s to describe a data architecture and processing approach that leveraged dimensional analysis. Measures and dimensions combine to form a “data cube”, which expands upon two-dimensional tabular data to include pre-calculated aggregations of measures along the dimensional hierarchies. These cubes were cached in memory or faster tiers of storage to accelerate query response.



Dimensional analytics has historically been tightly coupled with traditional OLAP tools like **Microsoft SQL Server Analysis Services (SSAS)**. SSAS has been a workhorse analytics platform for decades, with thousands of organizations relying on SSAS for mission-critical analytics that are most commonly distributed to users within Power BI dashboards or Excel Pivot tables.



The Practical Limitations of SSAS and Traditional OLAP

Dimensional analysis provides a logical organization and data modeling paradigm that is leveraged by traditional OLAP tools like Microsoft SSAS to deliver high-speed analytics. SSAS analysts build logical data models and spin up cubes that they connect to using Excel Pivot Tables or Power BI dashboards. Users navigate through data using interactive tabular views or charts, enabling them to drill up and down, apply filters and slicers, and more. Since all data, including aggregations, are precalculated and cached, query response times are extremely fast.

Traditional OLAP approaches become more complicated, more expensive, and less practical as data sets scale. Precalculating cubes can take hours or even days to compute and build, which means they're no longer feasible for large data sets. Furthermore, any data analytics approach that requires physical data movement and transformation creates more work for data engineers and requires additional compute and storage resources, driving up the costs for data analytics. In practice, different user groups will build cubes based on smaller subsets of raw data. This creates disconnected data pipelines and often results in shadow analytics programs with no central governance.



These local OLAP cubes are physically disconnected from centralized data, meaning business users will always be out of sync with the most current version.

This may cause inconsistencies to emerge across different analyses. Metrics sprawl is most often a direct result of efforts to create performant analytics based on solutions that cannot scale with the size of source data sets. While cloud data platforms have democratized the ability to centrally manage large data sets, traditional approaches to delivering speed of thought analytics have not kept up.

High cost, complexity, lack of consistency, and physical limitations make traditional OLAP solutions impractical for the modern cloud-scale organization. That said, the concept of dimensional analysis can still be applied to create a modern, cloud-first approach to high-performance analytics.



A Modern Approach to Analytics and Business Intelligence Acceleration

Cloud data platforms have democratized enterprise access to limitless storage connected to high bandwidth networks, flexible architectures for both structured and unstructured data, data management, data transformation and movement utilities, and powerful query engines. Despite unlimited compute resources, it is impractical to deliver speed of thought analytics by querying raw cloud data sets directly. There is, however, an opportunity to combine the concept of dimensional analytics with cloud resource orchestration and AI/ML to deliver performant analytics without the limitations of traditional cube-based OLAP structures.





The approach can be summarized within three techniques

1

INTELLIGENT AGGREGATE MANAGEMENT

As discussed, OLAP cube structures contain pre-calculated aggregates for metrics at all levels in the hierarchy for relevant dimensions. Intelligent aggregate management techniques monitor user query patterns to identify what aggregates would be most beneficial to overall performance. This is a dynamic approach to creating aggregates on demand. Likewise, as aggregates are no longer used, they are released so as to not consume valuable storage resources. This approach dramatically reduces the storage required for aggregates, making it more feasible to provide performance acceleration for larger data sets.

2

DYNAMIC QUERY PUSHDOWN

The key to scaling performance across large datasets and large numbers of users is to leverage the architecture of the underlying cloud data platform. Dynamic query pushdown does exactly that by translating logical queries into physical queries, optimized for the cloud data platform. By performing resource intensive aggregations and transformations on the same data platform that stores the data, unnecessary data movement can be avoided and performance will scale along with the cloud data platform.

3

GRAPH-BASED QUERY PLANNING

As users create new requests for data, queries can be planned automatically using graph-based algorithms that find the most efficient approach. If an aggregate exists, the query may be re-written to access the pre-calculated value. If it doesn't exist, the query may execute against the raw data. If repeated queries against raw data occur, the planner may suggest the need for a new aggregate.




Leveraging dimensional analytics in architecting high-performance analytics does not necessarily mean conventional OLAP. Modern cloud data platforms bring a new set of capabilities that, when combined with a platform like AtScale's, enable speed of thought analytics on the most complicated cloud data sets.

How AtScale Accelerates Analytics Performance & Enables OLAP Modernization?

AtScale is a semantic layer platform that allows organizations to build a central hub of enterprise metrics tied to live cloud data to support speed of thought analytics. The AtScale platform combines six unique capabilities to deliver the most powerful analytics platforms in the market.





CONSUMPTION INTEGRATION

AtScale connects natively to common BI platforms including Microsoft Excel, Microsoft Power BI, Tableau, and Google Looker. Users connect to AtScale without custom client-side software and connect “live” to the data, without the need for data extracts or data loads



SEMANTIC MODELING

AtScale’s low-code modeling tools enable users to easily create multidimensional models tied to live cloud data platforms. Users define table relationships, create conformed dimensions, build new hierarchies, define new metrics, and create calculated columns. Users from different workgroups can leverage pre-existing models and definitions to compose new models.




QUERY VIRTUALIZATION

AtScale dynamically translates logical queries from consumer tools into physical cloud data platform queries using the dimensional semantic model. AtScale manages the various cloud data platform SQL dialects and creates a query plan using a graph-based optimizer. While data consumers interact with a simplified view of data based on the semantic model, the raw data stays in the cloud data platform.




PERFORMANCE OPTIMIZATION

AtScale leverages the dimensional analysis paradigm and the approach outlined above to dynamically optimize performance based on intelligent aggregates, dynamic query pushdown, and graph-based query planning. AtScale learns from user query behavior to constantly improve performance transparently and autonomously.



ANALYTICS GOVERNANCE

AtScale provides a single location to enforce governance policies, including both user identity management, data access control and consistency of business definitions.



DATA INTEGRATION

AtScale integrates with a variety of cloud data platforms including both data warehouses and data lakehouses and uses dynamic query pushdown to avoid data movement and scale performance.



SSAS Replacement

Many organizations are actively searching for alternatives to Microsoft SSAS as they work to modernize their data strategy. This can be difficult, particularly when SSAS is supporting mission-critical analytics. In this eBook, we've discussed an alternative approach to delivering speed of thought analytics without the limitations of traditional OLAP tools like SSAS. One important consideration is how data consumers will access analytics. SSAS users typically use Microsoft tools like Excel or Power BI, but they may also want to leverage other tools.

AtScale supports tools that speak OLAP like Excel and Power BI, tools that speak SQL like Tableau and Looker and tools that speak Python. With this universal approach, enterprises can leverage AtScale's semantic layer for business analysts, data scientists and application developers.

EXCEL

AtScale supports native MDX connections to Excel using PivotTables or CUBE functions for a seamless, high-performance analytics experience on live cloud data sets.

POWER BI

AtScale is currently the only solution to provide a live DAX interface for Power BI. This live connection enables business users to visualize AtScale's data models directly within the Power BI model canvas and execute high-performance "live" queries on cloud data without data extracts.

TABLEAU OR GOOGLE LOOKER

SSAS is a multidimensional analysis tool and, as such, is not supported by popular SQL-based data visualization tools like Google Looker. AtScale has the capability to translate tabular queries into multidimensional queries, opening up a new set of visualization options for high-performance, dimensional analytics.



How Betcllic Delivers Cloud-First Data Analytics

Betcllic — a European online gambling company — provides millions of players with services for sport and horse-racing betting, as well as poker and other casino games. Central to the company's success is leveraging massive amounts of data to improve the customer experience and manage the risks associated with the gambling industry.

As the Betcllic data team worked to keep pace with growth and modernize their analytics stack, they needed to overcome the limitations of their SQL Server Analysis Service (SSAS) implementation. By leveraging AtScale, Betcllic has been able to move away from SSAS while delivering a high-performance analytics experience to users and moving to a cloud-first data management approach.

Bitcllic uses AtScale to more easily connect Excel and Tableau to Snowflake data, enabling dimensional analysis against massive cloud data sets. Even sophisticated queries on large data sets that weren't possible with the legacy approach now take just a few seconds with AtScale. This high-performance cloud approach has enabled the company to decommission its legacy SSAS cubes.



As a result, Betclic has rolled out cloud-first analytics to over one hundred business users performing more than 7,000 queries per week. Speed of thought analytics has turned the company's business analysts into citizen data scientists, transforming the company into a truly data-driven business.



AtScale for Speed of Thought Analytics

Speed of thought analytics is critical for data-driven business decisions, but it's becoming harder to deliver as cloud data sets continue to grow and more users analyze data. A modern approach to dimensional analysis enables enterprises to ensure speed of thought analysis, no matter where data lives or how large it is.

With AtScale, enterprises can democratize access to business insights while reducing the data engineering effort and infrastructure costs required. Rather than creating physical cubes for specific team needs, AtScale's modern approach allows for more flexibility and eliminates additional infrastructure costs. Moreover, AtScale learns with each new query, creating aggregates that reduce query times and save on computing costs over the long run.

Perhaps most importantly, AtScale reduces reliance on data teams to create and maintain speed of thought analytics. That means enterprises can use dimensional analysis for a wider range of business use cases without heavy reliance on expensive and difficult-to-hire data engineers

This democratization of speed-of-thought analytics supports a more data-driven culture and sparks a flywheel of data productivity.



Take the Next Step

About AtScale

AtScale enables smarter decision-making by accelerating the flow of data-driven insights. The company's semantic layer platform simplifies, accelerates, and extends business intelligence and data science capabilities for enterprise customers across all industries. With AtScale, customers are empowered to democratize data, implement self-service BI and build a more agile analytics infrastructure for better, more impactful decision making. For more information, please visit www.atscale.com and follow us on LinkedIn, Twitter or Facebook.

[REQUEST A DEMO](#) →

[CONTACT US](#) →

[LEARN MORE](#) →

Read [The Practical Guide to Using a Semantic Layer](#)

Download [The Complete Buyers Guide to a Semantic Layer](#)

Review [How does using a semantic layer impact cloud data warehouse performance? \(benchmark reports\)](#)

Get [advice from fellow data & analytics leaders](#) on how to scale smarter data-driven decision-making

[Stay current](#) on analytics strategies with short [articles](#) and [webinars](#) on topics in analytics strategies

