



The Semantic Graph: MicroStrategy's Unique Strength and Your Competitive Advantage

WHITEPAPER

Table of Contents

- Executive Summary3**
 - Semantic Layer Approaches 3
 - The MicroStrategy Advantage..... 4
- Introduction5**
- The Promise and Limitations of a Universal Semantic Layer5**
 - Lack of a Calculation Engine and Its Implications 5
 - Proximity of Semantic Layer Impacts How It’s Utilized 6
- MicroStrategy's Semantic Graph Approach..... 7**
 - An Integrated Engines Framework 7**
 - Dynamic SQL Engine 8
 - Cache and Cube Engine 8
 - Dynamic Sourcing Engine..... 8
 - Federated Multi-Source Engine 9
 - Superior AI + BI Integration..... 11**
 - The Role of Semantic Graph in AI..... 11
 - Depth and Breadth of Semantic Graph Matters 11
 - Portability and Data Source Agnostic14**
 - Open Architecture15**
- Conclusion 16**

Executive Summary

MicroStrategy was founded on a very simple idea – that analytics programs should exist to drive better decision making at every level of the organization and that the availability of these insights should be pervasive. This is why the vision for our platform is – and always has been - ‘Intelligence Everywhere’.

To deliver on this mission, MicroStrategy offers the broadest array of delivery modes for analytics. This is why MicroStrategy is trusted by thousands of the largest most well-known brands in the world. Another key reason is the MicroStrategy Semantic Graph which weaves together data, business rules, relationships, and metadata so that components are **reusable, inheritable, and privacy aware**.

This unique object-oriented design within the MicroStrategy ONE platform makes it easy to maintain the full suite of data experiences without the burdensome effort required by other platforms. Unlike other solutions, MicroStrategy enables data teams to focus on creating foundational reusable assets that serve as the basis for all upstream analysis. Changes to this foundational layer are immediately propagated throughout the system – eliminating rework, driving alignment, and increasing user trust in the data they use for critical everyday decisions.

Semantic Layer Approaches

MicroStrategy has been the leader in providing a fully integrated semantic layer that allows organizations to connect to nearly any data source, model business definitions such as attributes (aka dimensions) and metrics (aka measures), govern data and reuse these definitions for downstream analytics purposes. Over the years, many competitive semantic layers have come and gone, while the MicroStrategy approach has stood the test of time. The market today broadly has three semantic layer approaches:

1. **Hyperscaler-tied Semantic Layer:** Tied to hyperscaler data infrastructure like LookML and Microsoft Fabric. These solutions often offer deep integration with specific hyperscaler platforms, catering to organizations heavily invested in those ecosystems.
2. **Universal Semantic Layer:** Primarily acts as a presentation layer for data models, reliant on underlying sources for processing (e.g., dbt, Cube, Kyvos, Kyligence, AtScale, Denodo, etc.). While providing modeling flexibility, these tools often lack a calculation engine, requiring additional tools for advanced and ad-hoc calculation at the analytics layer.
3. **BI Semantic Layer:** Integrates schema modeling with analytics and data visualization layers. It is often portable across hyperscaler infrastructures and includes an integrated calculation engine for advanced multi-dimensional analysis. MicroStrategy, Business Objects, and Cognos fall into this category, though the latter two have lagged in modernization and innovations.

The MicroStrategy Advantage

MicroStrategy stands out in the semantic layer landscape through three key differentiators that enhance its effectiveness:

1. **An Integrated Engines Framework:** The MicroStrategy Semantic Layer is not just a data mapping tool; it actively transforms data into multi-dimensional representations suitable for sophisticated analysis. By integrating various components such as a Dynamic SQL Engine, Dynamic Sourcing Engine, Cube and Caching Engines, and a Federated Multi-Source Engine, MicroStrategy provides a cohesive and scalable solution. This integrated framework ensures that data modeling is both powerful and efficient, enhancing interaction with underlying data systems.
2. **Superior AI + BI Integration:** MicroStrategy AI empowers users with features like Auto Answers, advanced Q&A, and AI visualizations, delivering in-depth data analysis and insights. By combining AI with MicroStrategy's robust Semantic Graph, it ensures reliable data interpretation and computation while maintaining data governance, security, and integrity.
3. **Portability and Data Source Agnostic:** Unlike many semantic layers that are tightly bound to specific data sources, the MicroStrategy Semantic Layer offers unparalleled flexibility. Its design allows organizations to seamlessly switch between different data sources or hyperscalers such as Google Cloud to Azure without needing to reconstruct data models. This portability and agnostic approach free businesses from vendor lock-in, enabling them to choose the most suitable technologies without compatibility concerns.
4. **Open Architecture:** The MicroStrategy Semantic Layer supports a broad spectrum of analytics tools through its open architecture. Whether it's enhancing BI functionalities with low-code/no-code options, integrating with popular visualization tools like Qlik, Tableau and Power BI, or empowering data scientists with optimized datasets for Python models, MicroStrategy ensures that data is not only accessible but also usable in diverse ways. This openness extends the utility of MicroStrategy's platform, accommodating a wide range of business needs and user preferences.

Explore these differentiations in more detail in the subsequent sections to understand how MicroStrategy's unique approach can benefit your organization.

Introduction

In the evolving analytics landscape, the semantic layer is essential for translating complex data into accessible business insights. While past discussions have focused on comparing semantic layer capabilities across various BI vendors, recent dialogue has expanded to include new vendors and technologies advocating for a universal semantic layer, independent from traditional analytics and BI functions. This section will examine the strengths and weaknesses of these semantic layer solutions, with a particular focus on MicroStrategy's approach within this changing landscape.

The Promise and Limitations of a Universal Semantic Layer

In recent years, a range of technologies and vendors have sought to establish a universal semantic layer, intended as a translation mechanism between data and business insights, accessible via robust APIs. This approach has become popular with the emergence of tools like dbt, Kyvos, AtScale, Denodo and Cube, reflecting a shift towards uniform data definitions. These platforms support data modeling across varied data sources and encourage a standardized approach to data management. However, they typically lack a robust calculation engine necessary for advanced data operations. This shortcoming, along with the complexities of managing multiple semantic systems within organizations, presents significant challenges. This discussion will delve into these issues and contrast them with MicroStrategy's integrated approach, which illustrates how a well-designed semantic layer can significantly enhance data utility and analytical capabilities.

Lack of a Calculation Engine and Its Implications

Despite their growing popularity, universal semantic layer solutions often lack a crucial component: a robust calculation engine necessary for complex data aggregations and multi-dimensional analysis. For instance, accurately calculating average prices at various aggregation levels requires the semantic layer to effectively manage pre-aggregations of revenue and units sold. The inability to handle intermediate metric calculations such as smart totals, nested aggregations, and case-specific computations means that these complex calculations must often be performed elsewhere. Consequently, without a comprehensive calculation engine, the semantic layer essentially serves as a limited presentation layer, heavily reliant on underlying data sources for data processing.

Moreover, this limitation indicates that universal semantic layers mainly address the data presentation layer needs of other BI and downstream tools. This arrangement does not eliminate the necessity for BI tools to develop their own semantic layers, equipped to manage and leverage their internal aggregation engine logic. As a result, organizations must manage multiple tools and semantic layer configurations, an endeavor that requires significant operational efforts and incurs considerable costs.

Proximity of Semantic Layer Impacts How It's Utilized

The positioning of a semantic layer within the data stack significantly influences its application and best practices. Universal semantic layers, typically situated close to the data layer, are designed to provide a consistent and stable interface for data access and integration. This proximity to the data source makes them ideal for operational reporting and scenarios where data structures and queries are relatively static. They offer a stable environment for predefined queries and standard reports, ensuring reliability and consistency. However, this static nature can limit their flexibility and responsiveness to dynamic analytical needs.

In contrast, MicroStrategy's semantic graph resides within the analytics and visualization layer, making it more dynamic and adaptable. This integration allows for real-time interaction with data, supporting complex, ad-hoc analyses and evolving business questions. The semantic layer's proximity to the analytics tools means that it can quickly adjust to new queries and data models, providing users with the flexibility to explore data in a more iterative and interactive manner.

However, it is not always an either-or scenario. There are instances where both approaches can complement each other effectively. For example, dbt provides tremendous value in ELT data modeling by transforming and materializing data to a database. In a number of cases, customers use dbt to prepare and structure their data, and then leverage MicroStrategy to build a semantic layer on top of this data. This combination allows organizations to benefit from the robust data processing capabilities of dbt while also taking advantage of MicroStrategy's dynamic and flexible semantic layer for advanced analytics and visualization.

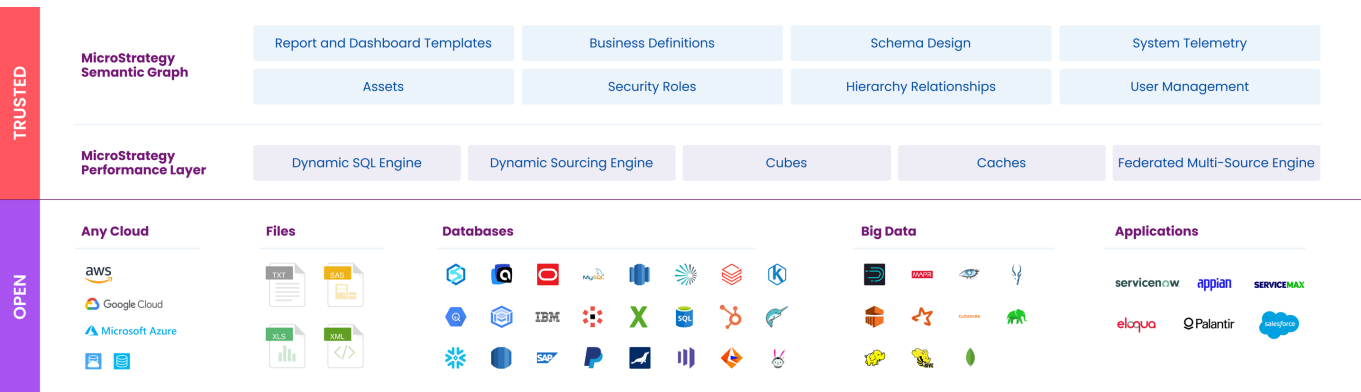
MicroStrategy's Semantic Graph Approach

In this section, we will explore three key aspects of MicroStrategy's semantic graph approach that set it apart from others. An integrated framework that ensures seamless functionality across diverse systems, portability that facilitates vendor-agnostic flexibility, and an open architecture that supports a wide range of downstream analytics tools. Each component is designed to reinforce the others, providing a robust platform that adapts to evolving business needs.

An Integrated Engines Framework

MicroStrategy's semantic graph does more than just map data structures; it actively transforms data into multi-dimensional representations suitable for advanced analysis. MicroStrategy stands out by providing robust data modeling capabilities across diverse data sources and deploying a suite of components that form an integrated framework to support the semantic graph:

- **Dynamic SQL Engine:** Generates on-demand, source-native, and optimized multi-pass query. This includes the ability to push down advanced calculations and data processing to the source for scalable performance.
- **Cache and Cube Engine:** Enhances data retrieval speed and reduces the load on underlying data sources by storing pre-calculated, summarized data for quick access.
- **Dynamic Sourcing Engine:** Automatically selects the optimal data retrieval path for best performance—either querying an existing in-memory data cube or generating a source-native, optimized query directly from the backend source.
- **Federated Multi-Source Engine:** Allows seamless connection to multiple data sources within a single project.



Dynamic SQL Engine

The Dynamic SQL Engine generates on-demand, source-native, and optimized multi-pass queries. This includes the ability to push down advanced calculations and data processing to the source for scalable performance. However, a robust SQL engine must also address fundamental data modeling challenges such as multi-fact and multi-granularity issues, non-aggregatable facts, slowly changing dimensions (SCD), level metrics, nested aggregations, set qualifications, distinct counts and many other complex data modeling scenarios.

Current solutions often fail to adequately tackle these real-world challenges due to either a lack of a dynamic SQL engine or one that is not robust enough to handle them, requiring users to solve these issues manually through SQL, data pipelines, or other methods. In contrast, MicroStrategy's Dynamic SQL Engine is designed to handle these complexities, providing a comprehensive solution that not only generates efficient queries but also simplifies data integration and management. This ensures better performance and scalability without compromising on the integrity and usability of the data.

Incorporating these capabilities into the Dynamic SQL Engine will not only streamline the process but also reinforce our commitment to addressing the foundational aspects of data management that are often overlooked in modern solutions.

Cache and Cube Engine

In-memory cubes and caching strategies are becoming increasingly important with the rise of cost-per-query databases like Snowflake and Databricks. MicroStrategy's approach includes the use of in-memory caching and shared in-memory dataset cubes.

A cache is a stored result set that improves response times for future requests by retrieving results from MicroStrategy's server instead of re-executing queries against a database. Our platform supports various caches, including result caches, data element caches, and semantic layer object caches.

Additionally, MicroStrategy's Intelligent Cubes allow multiple reports to access data from a single shared in-memory dataset, significantly reducing query execution times. This scalable solution minimizes data consumption and redundancy by creating only the necessary data sets. Intelligent Cubes can also use partitioning to load more data into memory and perform parallel queries, leveraging multi-core hardware architectures for optimal performance.

By combining these caching strategies, MicroStrategy not only boosts performance but also reduces overall database query costs, ensuring data is synchronized across platforms and tools.

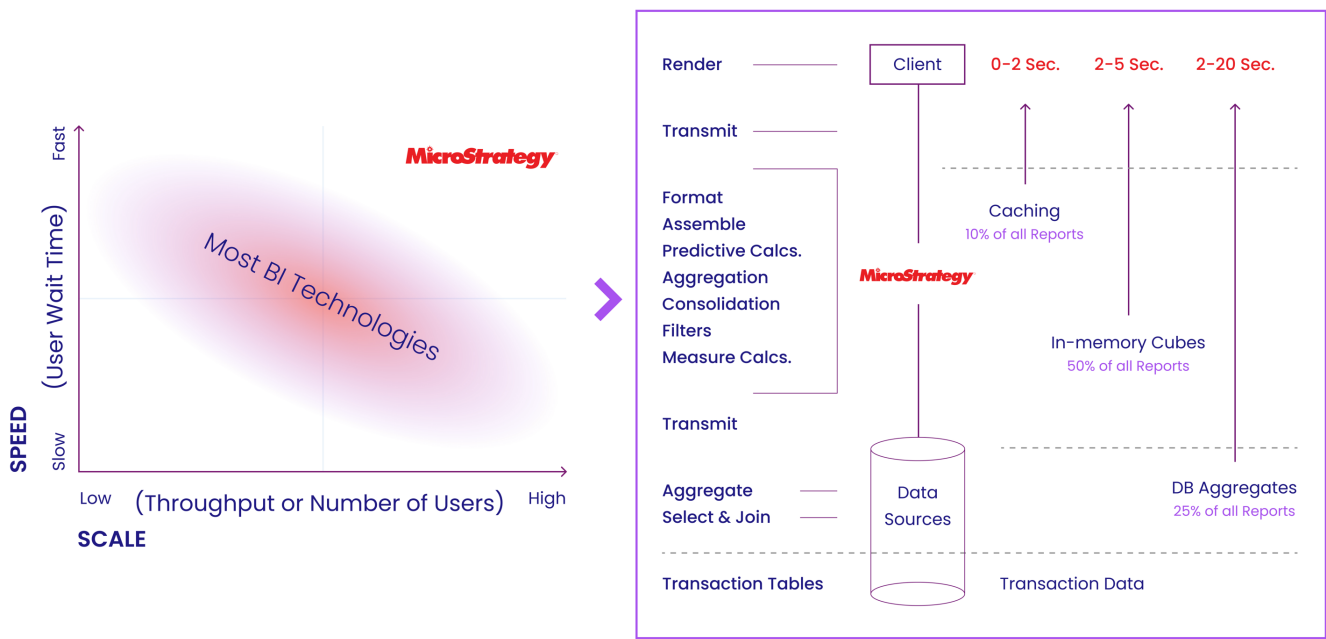
Dynamic Sourcing Engine

Unlike most existing semantic layers where the source of a report request is statically mapped and pre-determined, MicroStrategy's approach is more dynamic. Dynamic Sourcing is a powerful feature of MicroStrategy that optimizes report requests by determining at run-time whether the required data can be obtained from an in-memory data cube, an aggregate table, or a base table from the warehouse.

If the data needed for a report is available in a published cube, the cube satisfies the request, eliminating the need to execute a separate query against the data warehouse. The most suitable cube is automatically selected from the available options.

If the data cannot be found in an existing cube, MicroStrategy generates a query directly against the backend data warehouse. It first checks the semantic graph for available aggregate tables, querying the smallest aggregate table that meets the request's requirements. If no aggregate table is available, MicroStrategy will generate a source-native, optimized query against the base table(s) to retrieve the necessary data.

This dynamic sourcing approach is a significant advantage and enables MicroStrategy to achieve optimal performance at scale.

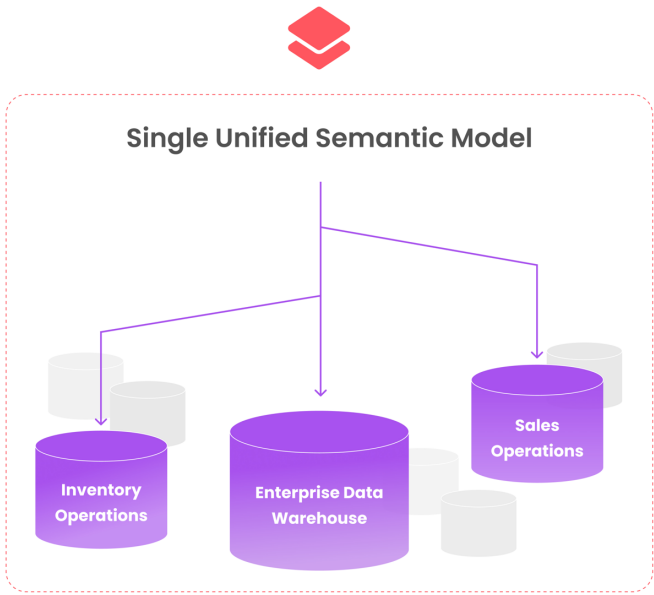


Federated Multi-Source Engine

Organizations generally maintain silos of data across different data sources. This disparate model makes it very difficult for users to gain a unified view of the business. Unfortunately, universal semantic layer solutions like dbt, Cube, AtScale and others lack native support for connecting to multiple databases within the same project, often leading to vendor lock-in as they require migrating all data to a dedicated source.

In contrast, MicroStrategy’s Federated Multi-Source feature allows seamless connection to multiple data sources within a single project. This integration enables you to consolidate information from various databases into a unified MicroStrategy project, leveraging a single semantic layer schema. The ability to access multiple data sources in one project offers significant advantages. For example, a sales manager wants to include forecast data available in Salesforce alongside actual sales data from Snowflake. This

capability to connect, query, and join data across sources allows for flexible, comprehensive, and integrated reporting solutions without the need to centralize all your data into one dedicated source.



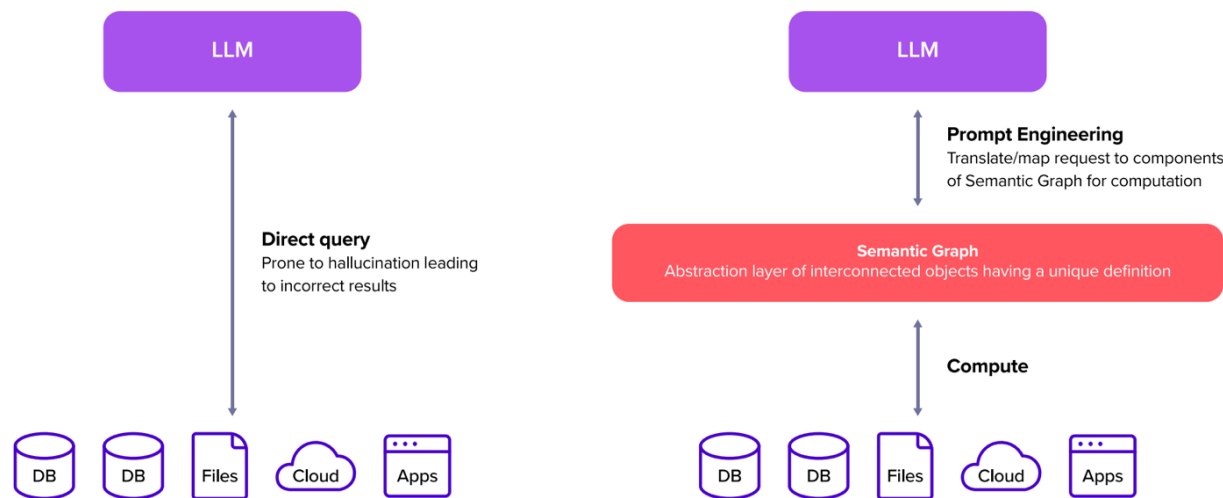
The integration of these components not only facilitates sophisticated data modeling but also significantly improves how these models interact with underlying data systems. This cohesive functionality provides flexibility in design and scalability in performance. MicroStrategy's approach ensures that organizations can fully leverage their data assets, achieving deeper insights and maintaining efficiency at scale, without the complexities of managing multiple semantic layers or compromising on analytical capabilities.

Superior AI + BI Integration

MicroStrategy AI provides an array of AI features to empower users with varying skill levels and different roles within an organization. Business users and analysts can take advantage of Auto Answers, a chatbot experience that allows them to dig deeper into their dashboard by providing insights and in-depth data analysis. This includes advanced Q&A and AI visualizations that leverage machine learning capabilities to generate key driver analysis, forecasts, and trends. They can also use bots that are focused on a specific use case or persona and allow for additional customizations leveraging Knowledge Asset and Custom Instructions fields to provide more business context. Other features available in MicroStrategy AI include Auto Dashboard, which enables users to design dashboards more efficiently, and Auto SQL, which helps administrators and architects to expedite data modeling by generating SQL.

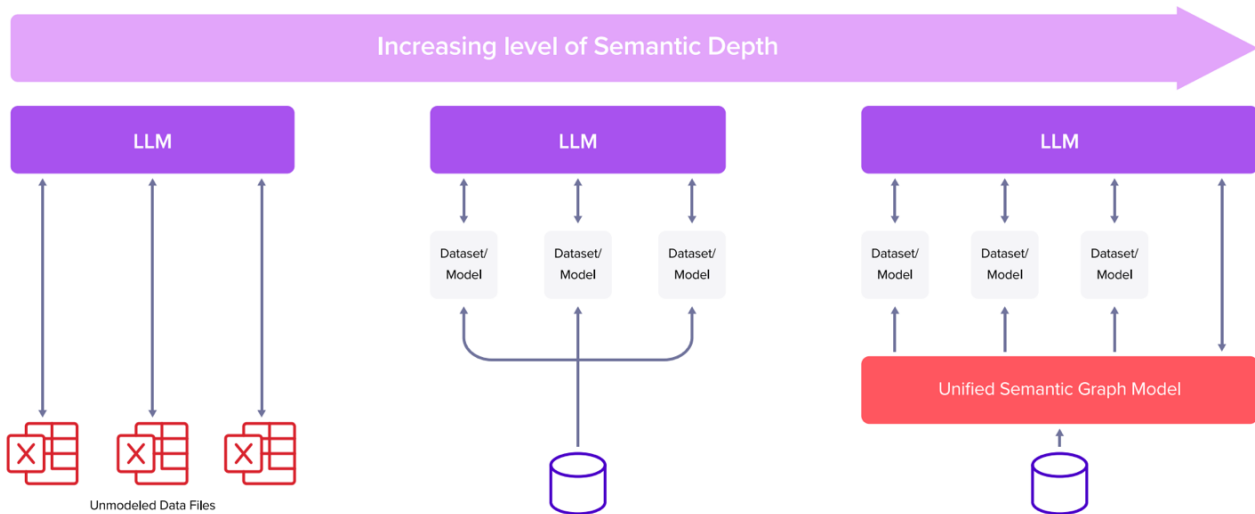
The Role of Semantic Graph in AI

Using AI to directly interpret and aggregate data can be unreliable. However, employing AI prompt engineering with a Semantic Graph resolves this issue. Prompt engineering fine-tunes the AI's language abilities to translate specific queries into actionable commands for the Semantic Graph. In this setup, the AI excels at translation rather than computation. It interprets the user's request through prompt engineering and maps it to components the Semantic Graph comprehends. The MicroStrategy Semantic Graph, with its robust grasp of data relationships and standardized business logic, handles the computational task for accurate and reliable results. The illustration below details this intricate process, showing how both prompt engineering and the Semantic Graph act together as a crucial bridge in AI-BI integration:



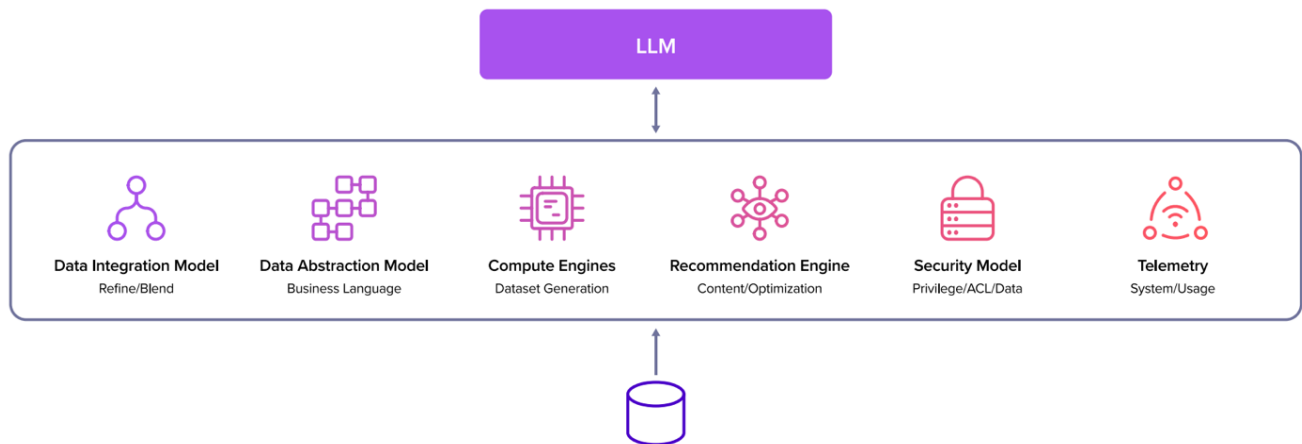
Depth and Breadth of Semantic Graph Matters

The Semantic Graph in BI is often expressed in the context of various downstream data tools and applications. The design and sophistication of the Semantic Graph can vary widely, reflecting the specific needs and capabilities of the tools and platforms in use, as shown below:



From a data abstraction perspective, basic productivity tools like Microsoft Excel often lack a semantic layer/model altogether, whereas point-solution BI tools usually confine semantic definitions within individual datasets. In contrast, MicroStrategy utilizes a foundational Semantic Graph that resides below the dataset level, serving as a base to create various upstream objects, including datasets.

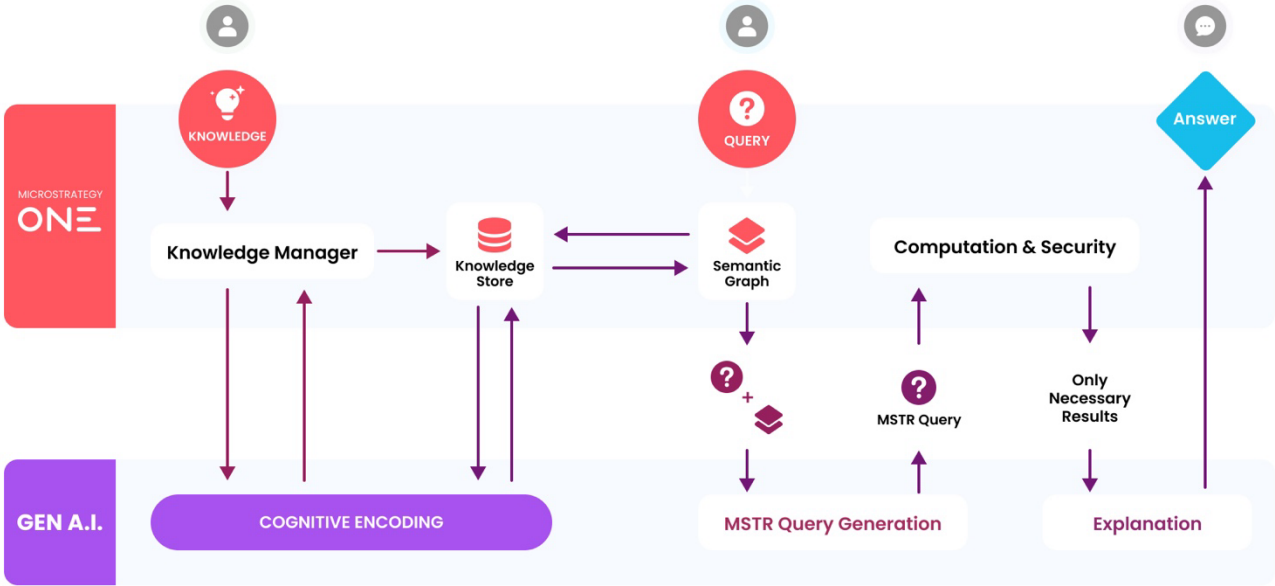
A Semantic Graph that's embedded deeper into a platform can offer more possibilities for more comprehensive AI integration. These capabilities can include, but are not limited to, those listed in the illustration below:



Together, a Semantic Graph and an AI model enable prompt engineering to tap into additional elements of the graph. For example, the security model in the MicroStrategy Semantic Graph can actively evaluate access permissions or data restrictions against a dataset when a user makes an AI request. In this way, it can solve data privacy and governance challenges. Usage and system telemetry can further refine AI interactions, optimizing requests or tailoring recommendations for queries. The more robust a Semantic Graph, the more flexibility it offers for deeper AI integration in analytics.

Furthermore, MicroStrategy enables users to fine-tune their AI deployment by providing additional semantic layer context through a file using the Knowledge Assets functionality. The Knowledge Manager processes all the information uploaded through an Excel file to augment the knowledge of MicroStrategy AI. The embedding model then processes this information, transforming it into definitions saved on the Knowledge Store. The Knowledge Store acts as a secure vault for storing domain knowledge that has been encoded using cognitive processing techniques. This encoding not only preserves the integrity of the data but also enhances its accessibility for cognitive search operations.

When interacting with the Gen AI module, the Knowledge Store plays a critical role by supplying contextually relevant information. This ensures that the Gen AI can formulate precise and accurate MicroStrategy queries.

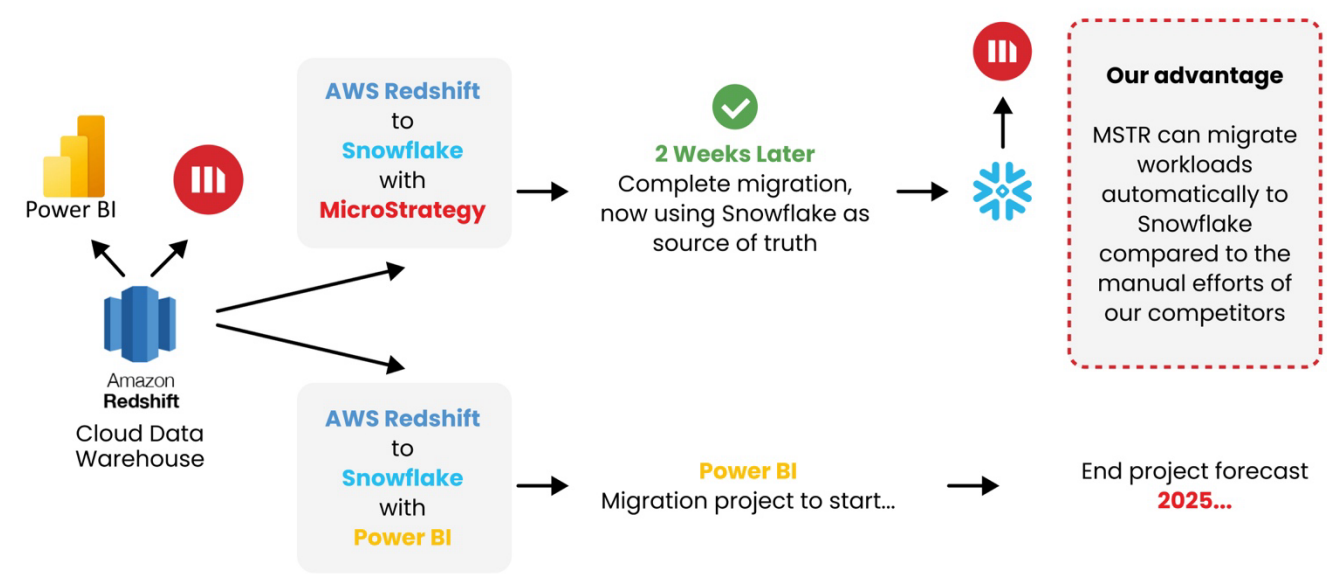


Portability and Data Source Agnostic

Unlike many semantic layer models that are intricately linked to a specific data source's syntax, MicroStrategy's approach stands out for its flexibility and vendor-agnostic design. Typically, semantic models are constructed using syntax unique to their initial data source, posing substantial challenges when transitioning between technologies. For instance, migrating data from an Oracle system to Snowflake often requires a complete redesign or reconstruction of the existing model due to its deep integration with the original data source. Over time, these challenges can result in an organization becoming locked into a specific data source vendor, even if they wish to switch.

This issue is exacerbated when semantics are embedded in a platform-specific tool like Looker on Google Cloud. Transitioning those semantic models to Azure can be particularly challenging due to the tool's deep integration with the platform's native features and services. This scenario highlights the broader issue of hyperscaler lock-in, where switching cloud providers becomes a complex and resource-intensive endeavor.

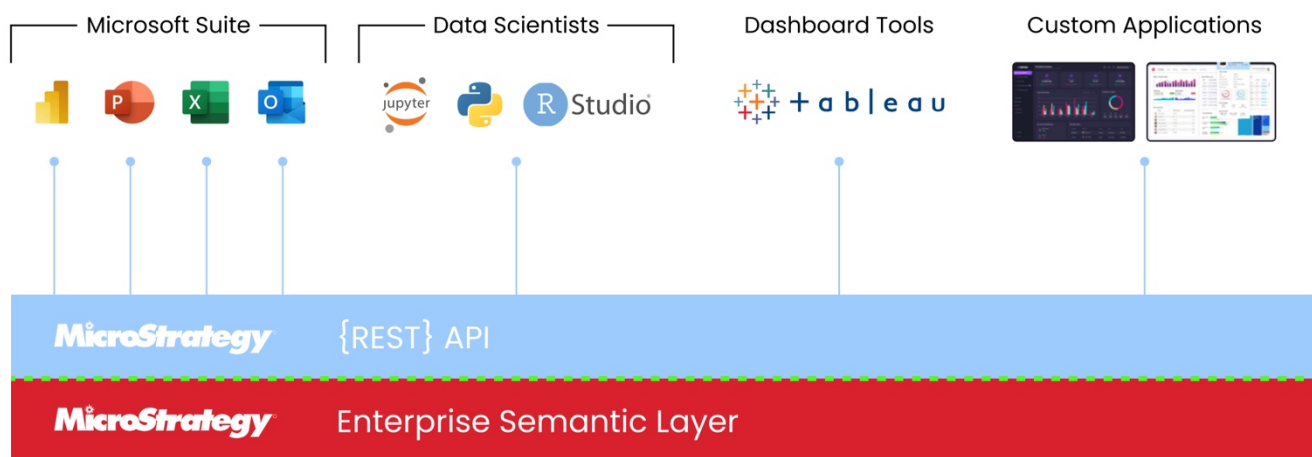
In contrast, MicroStrategy's semantic graph is designed to be independent of the underlying data source. This independence allows organizations to seamlessly migrate their data assets from one vendor to another without the need to rebuild or significantly alter their existing models. This aspect of MicroStrategy's design is particularly beneficial as it conserves resources that would otherwise be spent redesigning models that are too closely tied to specific data environments. Below is a use case involving a large retailer migrating their cloud data warehouse from Redshift to Snowflake. As you can see, the migration effort was substantially less with MicroStrategy compared to other semantic layer platforms, in this case, Power BI:



In an era where data strategies must be adaptable to changing technological landscapes, the ability to maintain a consistent semantic model without frequent redesigns is invaluable. MicroStrategy's approach ensures that organizations can focus on leveraging their data effectively, rather than being constrained by vendor-specific dependencies.

Open Architecture

MicroStrategy's semantic graph offers a sophisticated and open architecture, establishing it as a central source of trusted data that can be utilized across various platforms. This approach ensures that data from MicroStrategy's semantic graph can be maximized through its REST API, enhancing the impact and accessibility of its capabilities.



Key extensions enabled by MicroStrategy's semantic graph include:

- **Office Productivity Tools:** Supports real-time updates and direct importation of visualizations into Excel and PowerPoint, ensuring presentations and reports reflect the most current data.
- **Data Science Support:** Offers data scientists curated datasets optimized for running complex models in Python, promoting efficiency and innovation in data-driven projects.
- **Dashboard Integration:** Facilitates seamless connectivity with popular dashboard tools like Power BI, Tableau, and Qlik, harnessing the power of a centralized data model for improved data visualization and decision-making.
- **Low-Code/No-Code Customization:** Empowers users to easily customize, integrate, and extend BI and AI functionalities into applications and websites.

Through these integrated capabilities, MicroStrategy not only promotes data consistency and trust across business units but also enhances organizational agility in deploying advanced data solutions.

Conclusion

The goal of this paper was to provide a framework for understanding the various semantic layer approaches and to illustrate where MicroStrategy fits into that framework. In recent years, there's been a proliferation of semantic layer options, each offering differing approaches. This speaks to the growing maturity of the marketplace and the increasing need for robust semantic layers. Although each solution has its own unique approaches, many come with challenges. Hyperscaler-tied semantic layers offer deep integration with specific platforms but can limit flexibility and lead to vendor lock-in. Universal semantic layers may provide modeling flexibility but often lack a robust calculation engine, resulting in reliance on additional tools for complex analytics. In contrast, the MicroStrategy Semantic Graph approach addresses these challenges head-on with several key differentiations:

- **Proximity to Analytics Layer:** Positioned within the analytics and visualization layer, it allows real-time interaction with data, supporting complex, ad-hoc analyses and ever evolving business questions.
- **Integrated Engines Framework:** Combines dynamic SQL engine, advanced caching, and federated multi-source capabilities for a cohesive and scalable solution.
- **Superior AI + BI Integration:** Enhances AI capabilities while ensuring governance, data security, and integrity.
- **Portability and Data Source Agnostic:** Allows seamless transitions between different data sources without needing to reconstruct data models, avoiding vendor lock-in.
- **Open Architecture:** Supports a wide range of analytics tools, ensuring data is accessible and usable across various platforms.

In short, our goal is to set the industry standard for semantic layer modeling and analytics solutions. At MicroStrategy, our vision is to deliver "Intelligence Everywhere", empowering organizations to achieve deeper analytical insights. And it all starts with a robust and powerful semantic graph foundation.



MicroStrategy Incorporated, 1850 Towers Crescent Plaza, Tysons Corner, VA 22182
Copyright © 2024. All Rights Reserved.
microstrategy.com