


5 REASONS

YOUR BUSINESS NEEDS A UNIVERSAL SEMANTIC LAYER TODAY

The demand for data continues to grow from inside and outside every organization and from machines now, too. However, storing massive amounts of data is not enough. **Organizations must solve the complex challenges** of organizing, managing, and accessing data despite the complexity of their data stacks, the proliferation of siloed data, and

inconsistent business definitions. This is why organizations need a universal semantic layer, an emerging—and sometimes poorly understood—concept. **A universal semantic layer unifies data to create a single source of truth** with consistent metrics, centralizes and enforces fine-grained data access controls, and helps organizations achieve faster, more cost-efficient results from data.

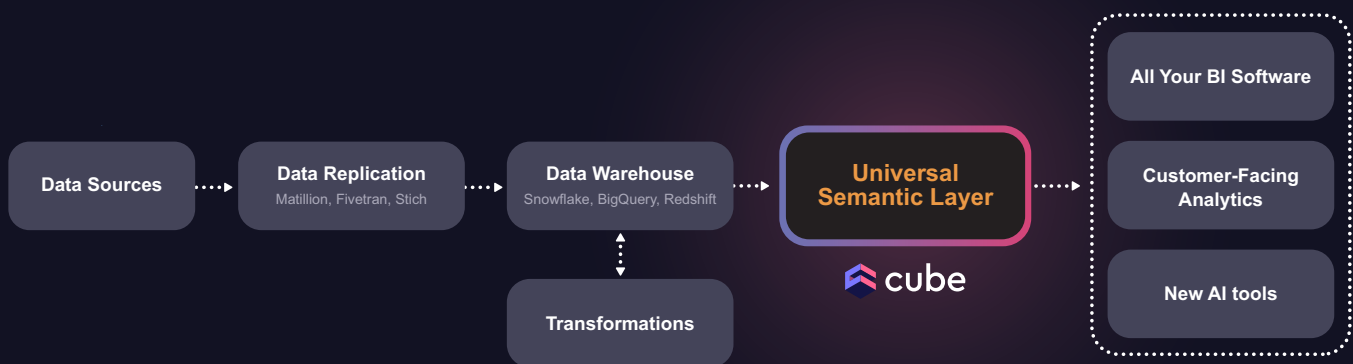
The Demand for Data Continues to Grow



Why a Universal Semantic Layer?



While the concept of a semantic layer has existed since the 1990s, self-service analytics in the 2000s led to the fragmentation of business definitions spread across data consumers. To solve the problem, a few vendors have introduced the concept of a **universal semantic layer**.



A universal semantic layer is an independent yet interoperable part of the modern data stack that sits between data sources and consumers.

It allows every data consumer, whether AI, BI, spreadsheets, or embedded analytics, to work with the same semantics and underlying data, increasing consistency and trust. The result? Consistent data drives confident business decisions across the organization.

Although a universal semantic layer can't solve all business problems, it can help organizations use complete, consistent data insights to navigate uncertainty. Having massive amounts of data alone doesn't create value, but consistent, complete, trusted data for informed decision-making does.



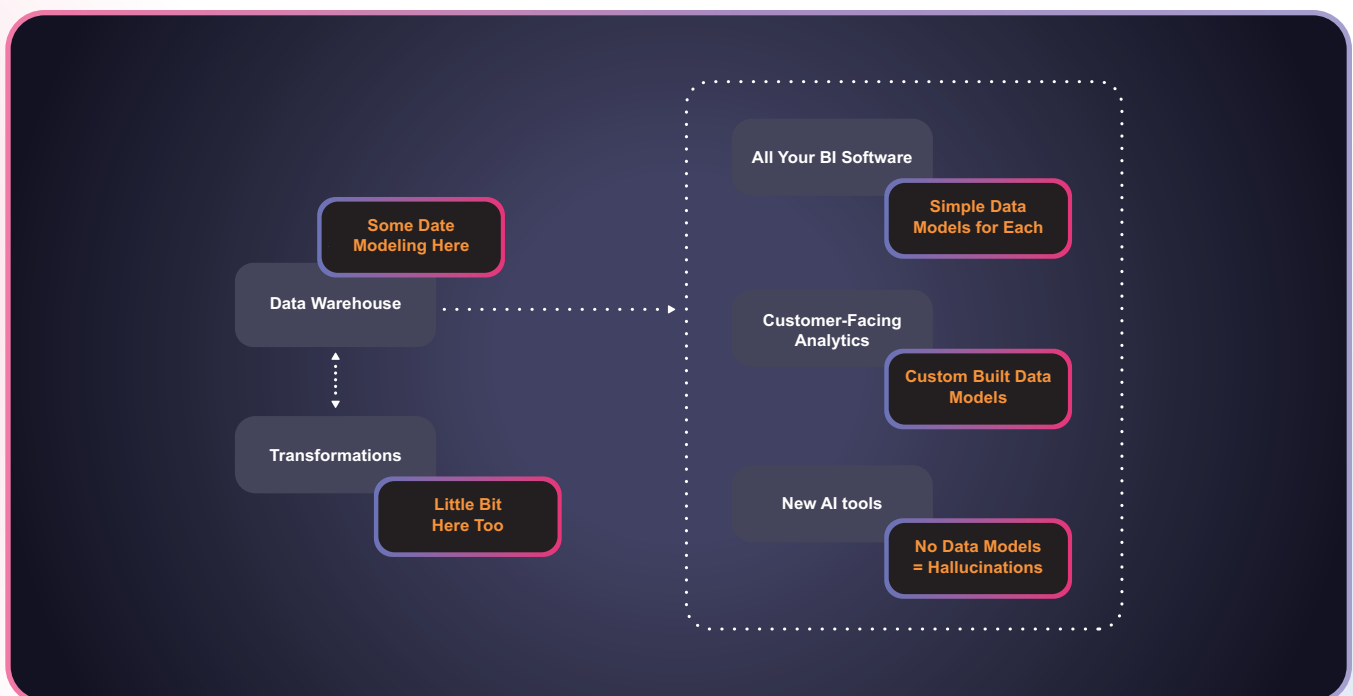
Here are five ways a universal semantic layer can help organizations solve problems associated with organizing, managing, and accessing data.

1 Unify Fragmented Business Logic

Modern organizations struggle with inconsistent metrics and duplicated business logic across siloed systems. Cube Cloud lets you model your data once and deliver it anywhere, creating a consistent, single source of truth for every data consumer. This unified view not only eliminates redundancies but also forms the foundation of modern OLAP and AI workloads. Multidimensional queries run reliably on Cube's cloud-native platform, enabling legacy OLAP replacement and ending the hassle of reconciling results from disparate data sources. At the same time, Cube Cloud prepares your data to be AI- and BI-ready so that you can share and reuse context-enriched data models across data consumers today and tomorrow for consistent results.

2 Govern Data Access Centrally

As regulatory demands and data security concerns grow, maintaining centralized control over data access is critical. Cube Cloud centralizes data access controls to make it easy to manage data access for users inside and outside of your organization. It can automatically rewrite queries and inject the appropriate security context, including user or role details, to ensure that only the proper users have access. With centralized and enforced fine-grained access controls, you can ensure that sensitive data is governed consistently across the organization. This unified approach not only helps meet compliance requirements but also increases trust in data so that AI outputs are accurate and BI content is actionable.



Your Business Needs a Semantic Layer



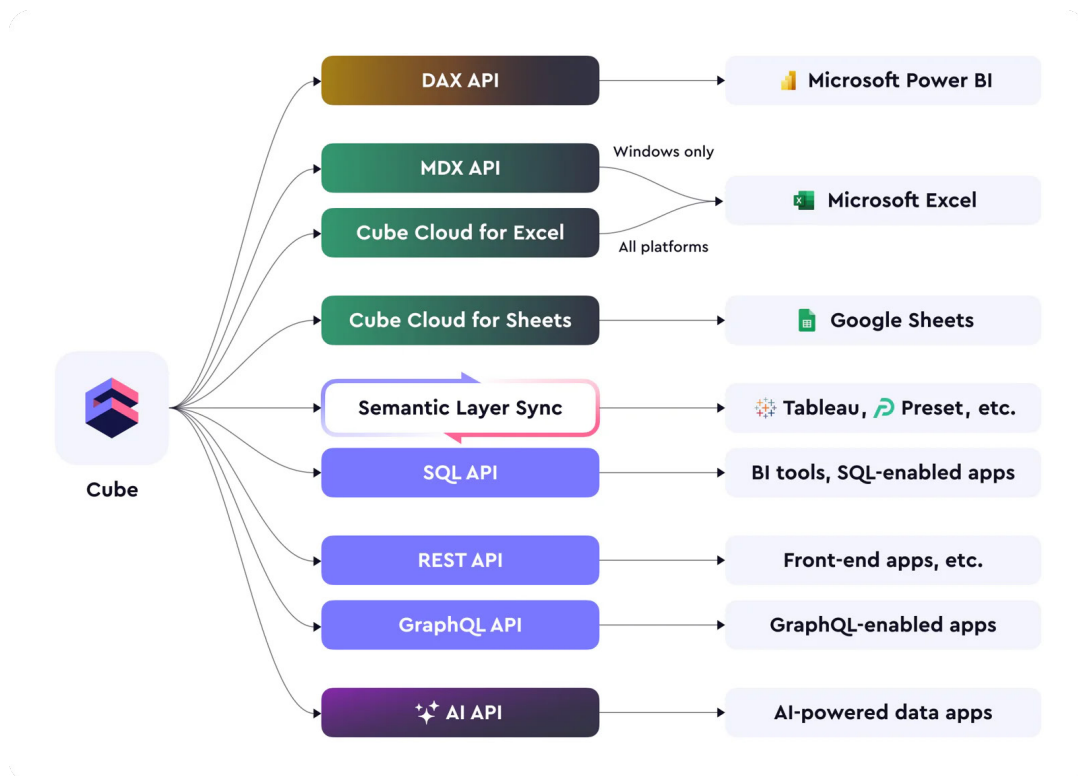
3 Optimize for Speed and Savings

Performance and cost-efficiency go hand-in-hand for cloud workloads. Cube Cloud's intelligent caching capabilities dramatically reduce query latency and lower cloud compute costs, making it ideal for both real-time modern OLAP and high-performance AI workloads. Through performance insights in the universal semantic layer, organizations can easily find redundant queries or other opportunities for caching and pre-aggregating query results, reducing the amount of data warehouse compute required. This optimization means your teams can derive actionable insights faster while enjoying significant operational savings, ensuring that your data stack is both agile and cost-effective.



4 Integrate Using API Endpoints

Cube Cloud is committed to interoperability, allowing you to integrate with your existing technology investments and adopt new platforms more easily. With an API-based approach, your unified data models and governed metrics can be delivered wherever it's needed, whether it be AI, BI, spreadsheets, or embedded analytics applications. This flexibility ensures that your organization can easily execute on its data strategy across the enterprise, maximizing the value of both modern OLAP and AI analytics without disrupting existing workflows.

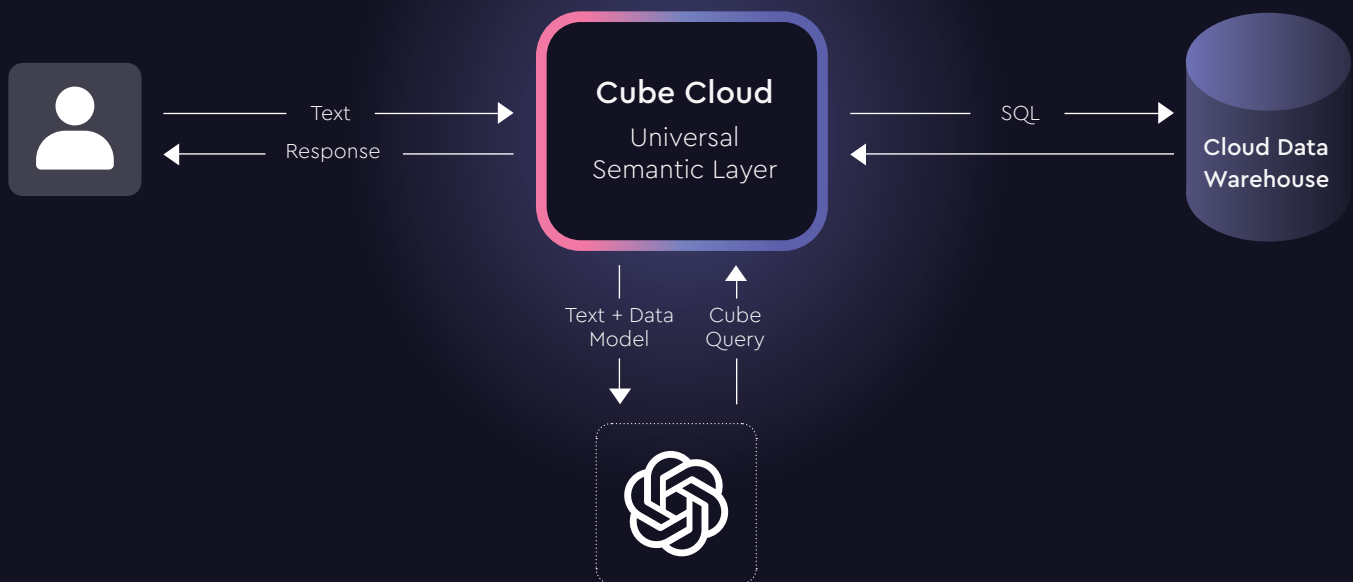




5 Provide Semantic Context to AI

Every data team is being asked what their AI plan is. A quick AI application for any data team is to enable internal and external users the ability to ask natural language questions and get charts, commentary, and tables of data as the results. But you already know, GenAI needs more context to deliver accurate results. Cube Cloud provides this business logic to LLMs and ensures accurate, consistent, and relevant responses from AI. A universal semantic makes it easy to deliver modeled data to LLMs to create AI chatbots that can answer questions from your proprietary data. Cube Cloud's AI API delivers trusted answers, integrating with Anthropic Claude models as a turnkey solution or BYO-LLM from top providers.

Text-to-Semantic Layer Scenario Improves Accuracy of Results



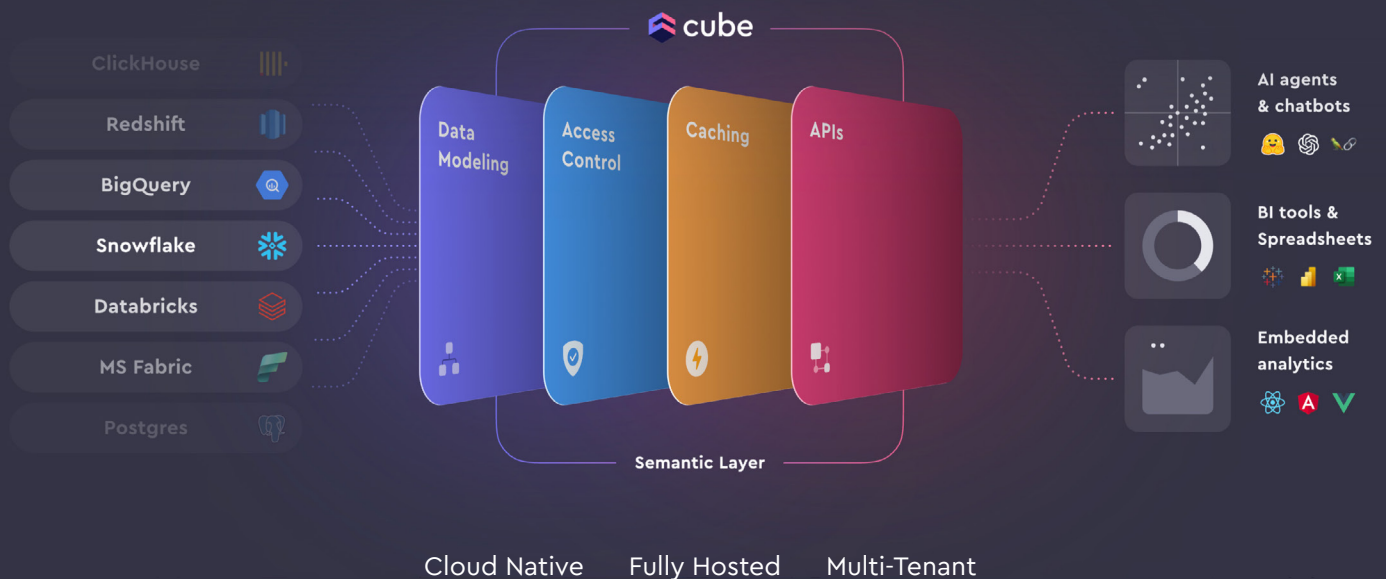
A universal semantic layer makes it easy to deliver modeled data

Your Business Needs a Semantic Layer



A Final Word

Effective organizations need uniform, accurate data to make actionable decisions. The organization's connective tissue, a universal semantic layer, links critical knowledge, information, and data assets. Regardless of the industry, a universal semantic layer allows employees to work together based on shared data definitions, ensure data governance and security, unite data silos, optimize cloud expenses, and harness the power of AI. A universal semantic layer translates to meaningful competitive advantage — and for any organization, there's no substitute.



About Cube

Cube brings consistency, context, and trust to the next generation of data experiences. Cube Cloud is a leading, AI-powered universal semantic layer platform, helping companies of any size manage and deliver trusted data with a single source of truth. Any data source can be unified, governed, optimized, and integrated with any data application: AI, BI, spreadsheets, and embedded analytics.

Cube is installed on 90,000 servers and used by over 5 million users. Customers include 20% of the Fortune 1000. Based in San Francisco, Cube is backed by Decibel, Bain Capital Ventures, Eniac Ventures, 645 Ventures, Databricks Ventures, and Betaworks. To learn more visit cube.dev.

Visit cube.dev
for more information

TRY CUBE FOR FREE