



Unlocking trustworthy AI

A semantic layer guide for the AI era



Table of contents

01	Introduction	03	05	LookML delivers: meeting the Enterprise mandate	11
02	Why is a semantic layer important in the context of AI	04		LookML unites data and business teams	11
03	Introducing LookML for building a code- based semantic layer	05		LookML supports multiple data sources	12
04	The six important considerations for choosing a semantic layer	06		LookML integrates into broader AI platform strategy	12
	Unite data and business teams	07		LookML powers full BI capabilities that enforce its use	13
	Support multiple data sources	08		LookML is enterprise ready	14
	Easily integrate into broader AI platform strategy	08		LookML prevents lock-in with flexible design	14
	Be a part of the enterprise workflows	09	06	Can the semantic layer be built elsewhere?	15
	Be enterprise ready	10		Option 01: database "Semantic Views" or "Gold Layers"	16
	Avoid lock-in as technology rapidly changes	10		Option 02: embedded in the BI offering	17
				Option 03: a standalone semantic offering	17
				Semantic Layer Location Category Capability Matrix	18
			07	Your foundation for an intelligent enterprise	19



Introduction

The rise of Generative AI presents an opportunity for businesses to unlock insights from their data at unprecedented speed. However, that opportunity is met with a critical challenge: though immensely powerful, AI models lack business context. Without understanding your company's unique definitions for metrics, KPIs, and relationships, AI models are prone to hallucinations that can erode user trust and undermine the investments that you are making.

The solution is a semantic layer - a business representation of your data that acts as a universal translator between your business language and your data sources. It provides essential context required for both AI models and business intelligence (BI) tools to deliver consistent, reliable answers to questions.

However, not all semantic layers are created equal. The choice of where and how you build your semantic layer will have implications for your organization's agility, connectivity, and continuity.

This brief provides critical considerations and a clear framework for evaluating a semantic layer, empowering you to make the most strategic decision for your enterprise.



Why is a semantic layer important in the context of AI

For all of their transformative potential, large language models (LLMs) and various purpose-built agents are trained off generalized datasets. When questions are asked of an LLM using business terms but without providing any additional context, its probabilistic response is not optimized for the particular individual or organization asking it. For example, if a user asks, "What was our Monthly Active User (MAU) count last month?", the LLM won't know if your organization defines MAU as logged-in users or users who performed a specific in-app action. It also won't know if monthly refers to calendar month or a rolling 30-day window. Additionally, when that response requires the LLM to write SQL, it must properly write the SQL and optimize it for the particular database cost and performance structure, two things that are further challenges for an LLM.

A semantic layer is the solution to these problems. Since it already contains the business naming conventions, it is able to provide the additional context necessary for the LLM to understand the user's question correctly. In addition, since the semantic layer also knows how to translate user requests to SQL, it also aids the LLM in writing the correct SQL that is optimized for the particular database. In other words, when deployed correctly, a semantic layer shifts the burden of understanding a specific business question and compiling SQL away from the LLM to the semantic layer itself.

The impact of an LLM answering questions more accurately when used with a semantic layer is reflected in internal tests at Google Cloud Platform (GCP).



According to our internal benchmarks, LLMs can reduce errors by as much as two thirds when a semantic layer is used, when compared to when answering questions with raw SQL and no semantic layer.

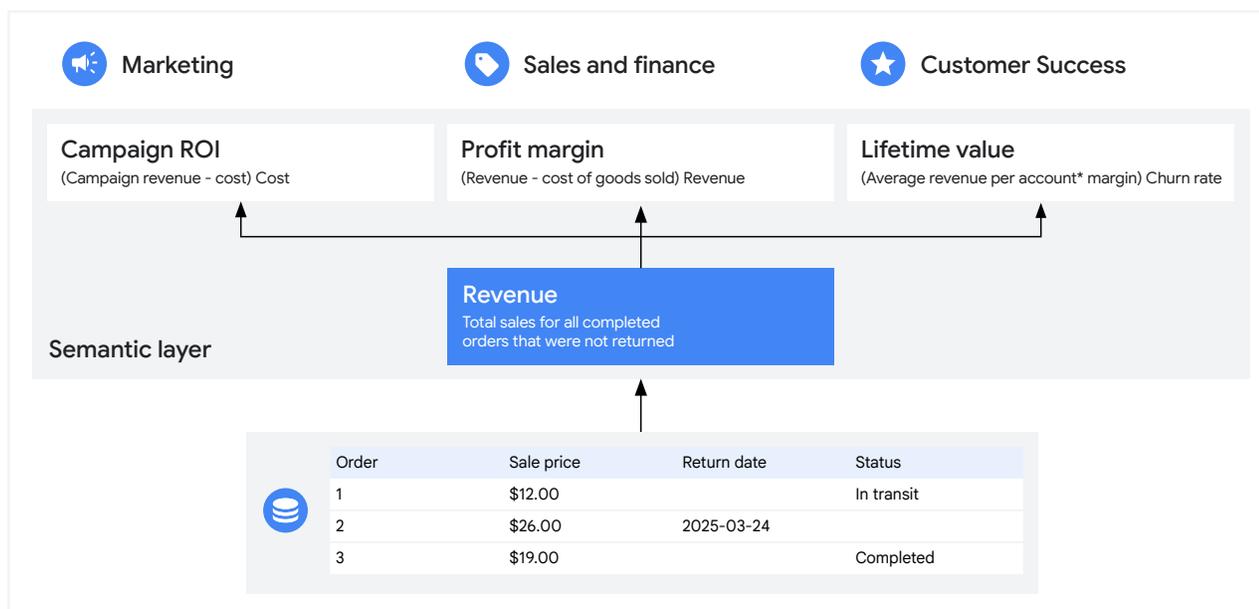


Introducing LookML for building a code-based semantic layer

The realization that deploying a semantic layer can significantly increase LLM accuracy is driving many organizations to prioritize its development. Investing the time to create a robust enterprise-wide semantic layer is now fundamental for many organizations looking to rise to the top in their adoption of transformative AI technologies.

[LookML](#) (Looker Modeling Language) is the declarative language that forms the core of the [Looker](#) platform's semantic layer. It defines your data structure and business logic in a human-readable, version-controlled context store that can [power AI agents](#) throughout an enterprise.

LookML abstracts the complexity of the raw database schema into consistent, reusable definitions that can help an AI agent translate a business-oriented prompt into a data request. By centralizing all business rules, metrics, and relationships in code, LookML ensures that every data consumer—whether running a simple report, generating a dashboard, or using natural language to query via an LLM—is using the exact same definitions for KPIs and dimensions. This enforced consistency is a critical first step toward building trust and accuracy in AI-driven insights.





The six important considerations for choosing a semantic layer

It is not sufficient to just deploy any semantic model using any technology; the model must be built in accordance with a broader Agentic Enterprise AI and BI strategy.

There are six critical considerations essential for any successful semantic layer implementation.

A semantic layer needs to:

01

Unite data and business teams

02

Support multiple data sources

03

Easily integrate into broader AI platform strategy

04

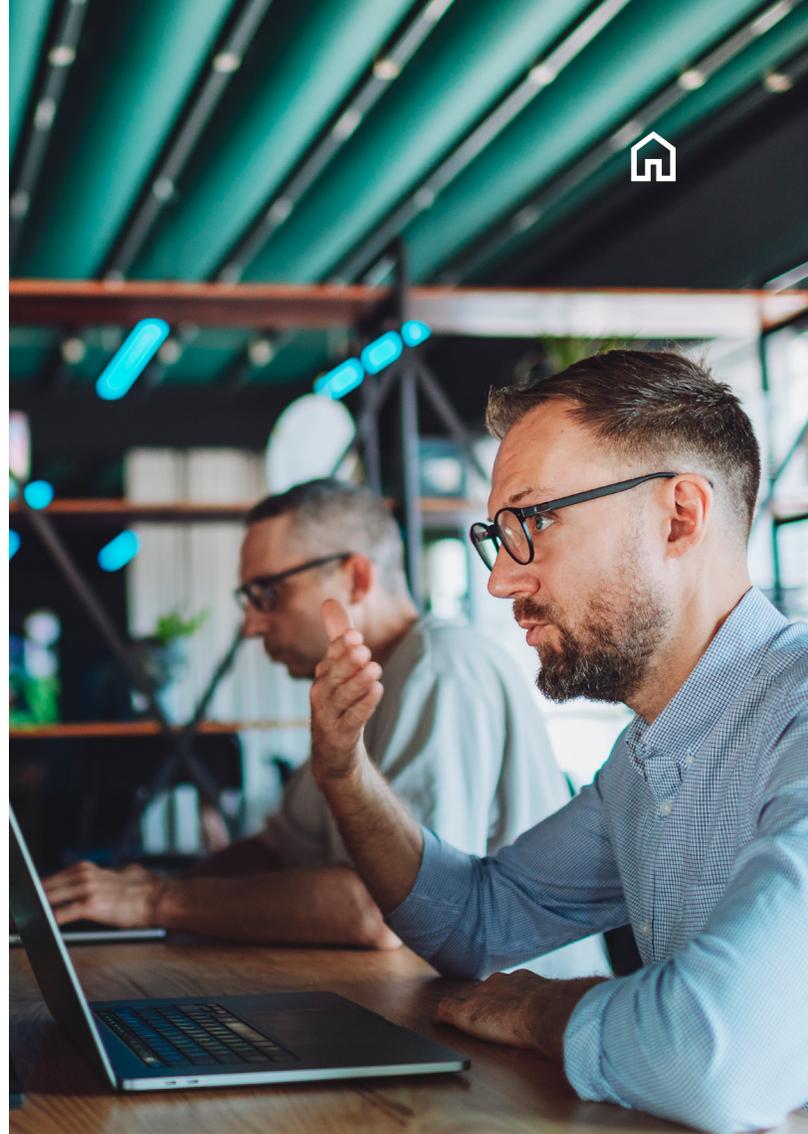
Be available to BI workflows that actually enforce its use

05

Be enterprise ready

06

Avoid lock-in as technology rapidly changes





Unite data and business teams

Beyond its technical impact, the most beneficial by-product of developing a semantic layer is that it acts as a forcing function to bring the domain experts, who understand line of business considerations best, together with the data experts, who best understand the underlying data to support that analysis. Successfully bringing these two groups together is something that many organizations have struggled to achieve, at large scale, for many years. When choosing a technology to develop a semantic layer on, it is critically important to consider whether that technology empowers both of these two core constituencies and brings them together. A semantic layer should not be a static repository, but should be a dynamic, collaborative workspace.

Look for specific capabilities that empower both teams:

A common, readable language: The core of the semantic layer should be defined in a way that is transparent and understandable to both business analysts and data engineers.

Modern development workflows: The platform should support collaborative processes. This includes features like version control, review history, and approval workflows to ensure definitions are properly vetted before being used in production.

Support for end user validation: Data experts need the ability to inspect the underlying queries generated by the semantic layer to debug and optimize them. Business analysts need the ability to explore the semantic outputs and make sure they are reflective of business needs.



It is important to build the semantic layer where these two constituencies interact, rather than building a semantic layer further upstream where the business teams will have minimal ability to participate.



Support multiple data sources

Centralizing and standardizing data as much as possible into a powerful data warehouse will unlock maximal value from Agentic AI experiences. But the reality is that no Enterprise organization can centralize 100% of their structured data in a single data source. To that end, it's important that the chosen semantic layer technology be able to span across multiple data sources, so when users ask a question, the LLM can respond appropriately regardless of where the structured data sits.

Easily integrate into broader AI platform strategy

While semantic layers are a valuable asset in driving improved conversational intelligence when using LLMs against structured data, the reality is that a semantic layer will never cover 100% of the usable enterprise information an AI agent might need access to. In practice, there will be a need for agents to leverage unstructured data that may not even be in the database or BI domains. As a result, it is critical that AI and Conversational Agents be able to resolve questions and perform actions that leverage both semantically modeled and non-semantically modeled data.

When choosing a semantic modeling technology, consider how seamlessly that technology can fit into your broader Agentic AI platform strategy.

For instance:

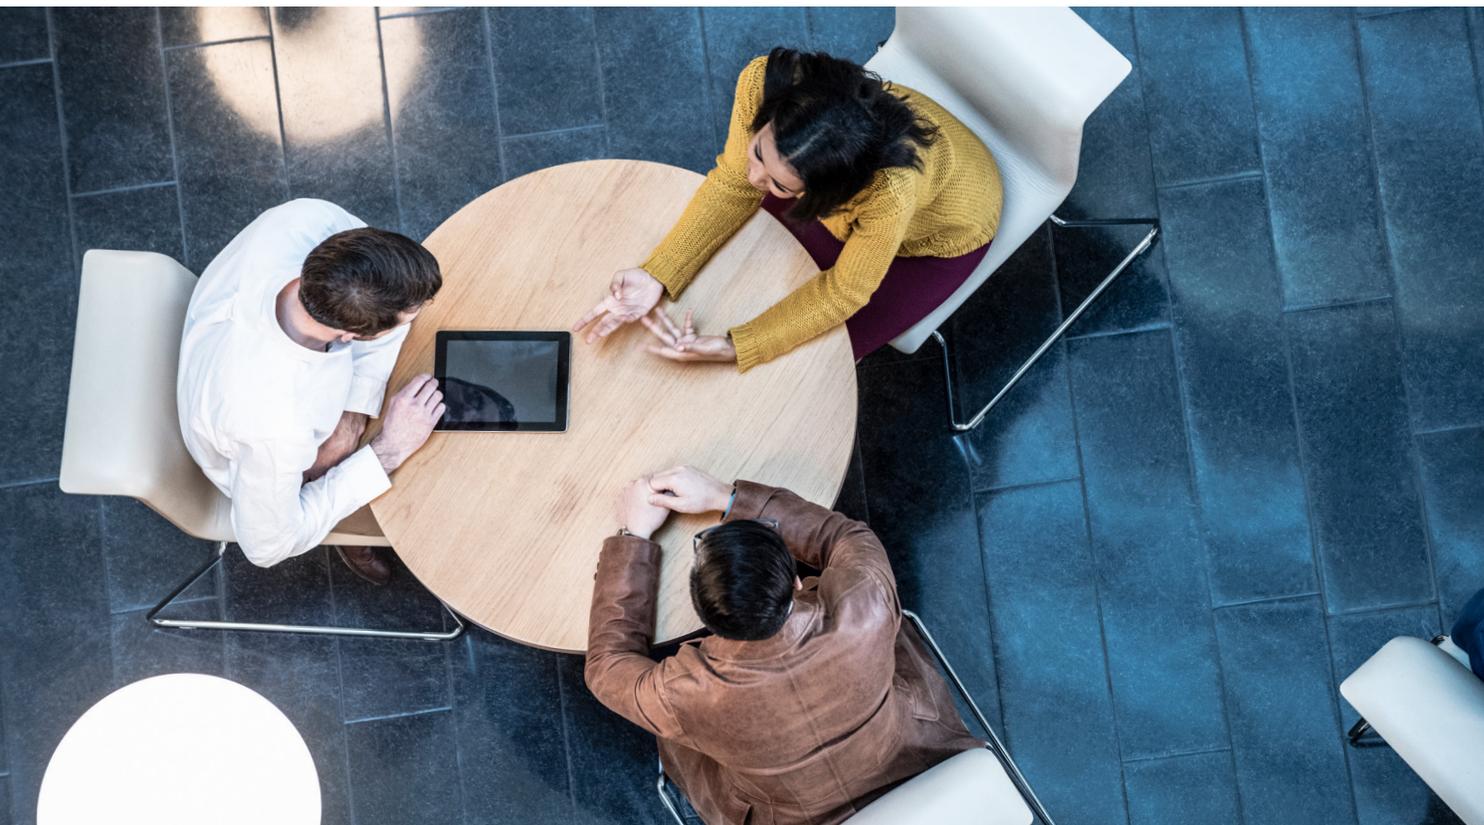
Can its querying capabilities be exposed via a Model Context Protocol (MCP) Server?

Can an agent easily access the semantic layer's APIs?

Can such an agent be published in the interfaces where your business users spend most of their time?

Conversational experiences that can be tightly integrated into a holistic AI platform strategy offer the most transformative potential.





Be a part of the enterprise workflows

The rise of conversationality will lead to a reduction in the importance of perfectly manicured reports and dashboards. But while the format of analysis might change, workflows that have always been critical in BI, such as saving, sharing, scheduling, alerting and activating, are even more crucial in an ecosystem where more analysis is occurring.

It is critical to ensure that your semantic layer's scope is not limited to a conversational experience that exists in isolation and that insights are not limited to the conversation itself. Your semantic layer offering should have its own native capabilities for the above workflows, rather than being dependent on 3rd party integrations. Corporate strategies can easily change and outright dependency on 3rd party integration is a potential risk.



Be enterprise ready

With AI models being new features in your environment, it is important that your supporting tech stack be established with a sturdy foundation.

Rather than building your semantic layer on top of technology that is still in its infancy and may not even be Generally Available (GA), choose a semantic modeling approach that has been proven to be successful at enterprise scale over a long period of time.

For enterprise adoption, this also extends to the compliance and security posture of the semantic layer. Enterprises require robust security competencies, including support for features like VPC Service Controls (VPC-SC) for perimeter defense, Customer-Managed Encryption Keys (CMEK) for enhanced control over sensitive data, Private IP connectivity to isolate traffic, and the ability to ensure Data Residency to meet regulatory requirements. A mature semantic layer should demonstrate established security and compliance capabilities to manage the risks inherent in using experimental LLMs and sensitive corporate data.

Avoid lock-in as technology rapidly changes

The need to maintain optionality in the semantic layer as technology evolves is even more pronounced in the age of GenAI due to the rapid pace of innovation.

The selected semantic layer should:

01

Decouple business logic from the underlying database

02

Ensure that all SQL written against the semantic layer is standard, non-proprietary SQL

03

Avoid embedding logic in specific workbooks and cubes.

With this approach, you'll be able to change the data warehouse, semantic layer, and BI interfaces as technology continues to innovate, best positioning you for success in the coming decade.

These six considerations form the blueprint for a successful, enterprise-grade semantic layer. With this framework established, the critical next question is a practical one: where should this semantic layer be built and managed? This is not merely a technical detail; it is a strategic architectural decision that will fundamentally enable—or inhibit—your ability to achieve the collaboration, governance, and flexibility your organization requires.



LookML delivers: meeting the Enterprise mandate

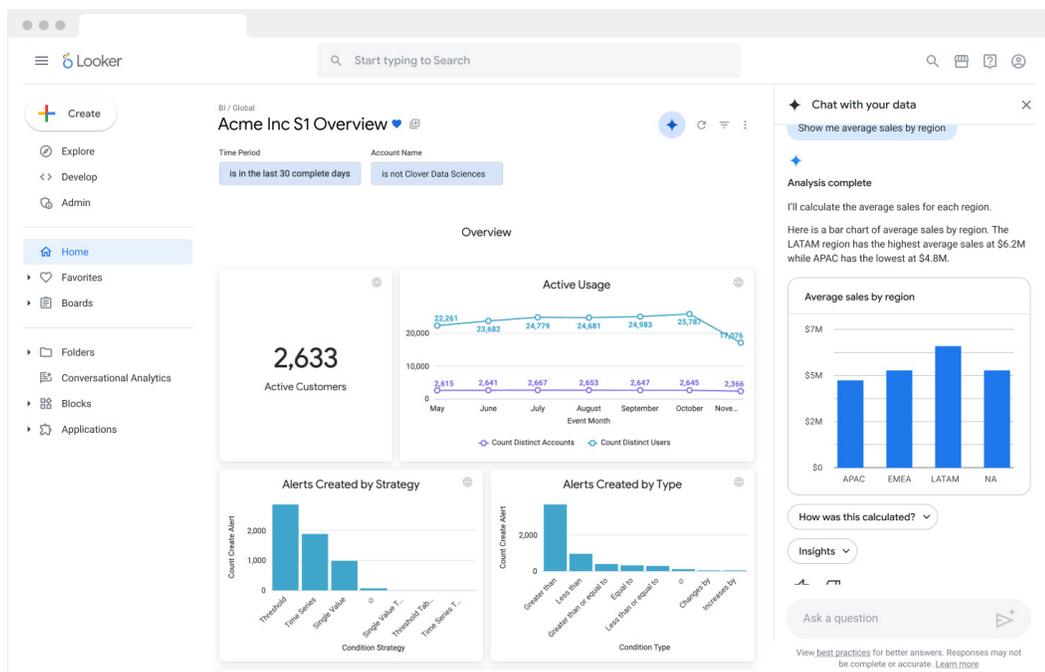
LookML, as part of a fully managed cloud data platform, is designed to deliver on all six enterprise readiness criteria, including security and compliance requirements.

LookML unites data and business teams:

LookML provides a common, human-readable language that serves as the contract between data engineers and business users. Data teams write LookML to encapsulate complex technical logic and transformations, ensuring data governance and efficiency. Simultaneously, the resulting semantic layer provides an intuitive, non-technical interface—the Explore—which allows business users to confidently self-serve accurate data without writing SQL.

This same, established semantic layer is precisely what Generative AI systems require: a trusted, well-defined vocabulary that allows the LLM to translate natural language questions into consistent, governed queries, thereby democratizing data access across the entire organization.

Within the Conversational Analytics UI, business users can simply and directly have a conversation directly with the data, knowing that the answers can be trusted thanks to the semantic model.



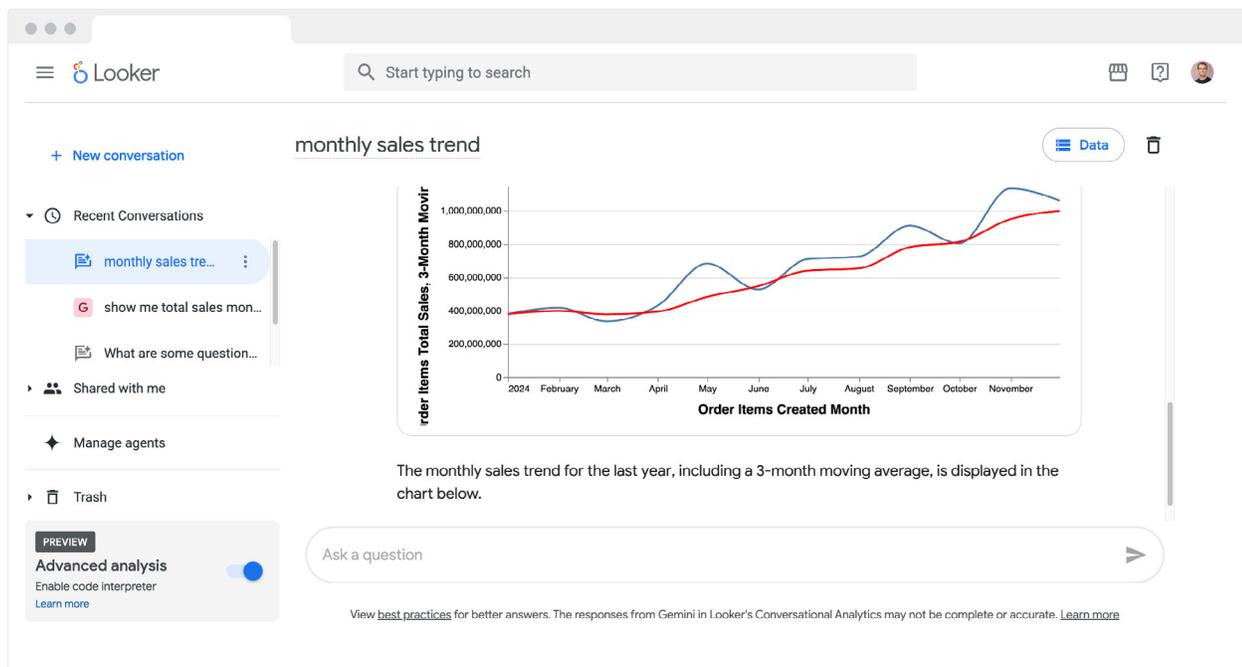


LookML supports multiple data sources:

LookML's strength lies in its database independence and ability to define a unified semantic layer across diverse data sources. It can connect to and model data from numerous platforms simultaneously, including all major cloud data warehouses (e.g., BigQuery, Snowflake) as well as traditional databases (e.g., Postgres, MySQL). Through the LookML structure, developers can create a single, holistic model for business analysis. This capability ensures that Generative AI applications can query a comprehensive view of the entire enterprise data landscape, even when the underlying data is fragmented.

LookML integrates into broader AI platform strategy:

The true power of using Looker's semantic layer is its ability to serve as the governed data foundation for the entire AI ecosystem. By centralizing trusted business logic and metadata, LookML provides a consistent, high-quality input for various AI interfaces. The resulting semantic model is inherently open; it can be exposed via APIs or shared metadata layers, allowing it to seamlessly integrate with and power third-party AI platforms and services. The Looker semantic model is also [exposed via MCP](#), making it available to any LLM that supports the MCP protocol. This interoperability ensures that your entire AI strategy, regardless of the tools used, is grounded in a single, authoritative definition of enterprise data.

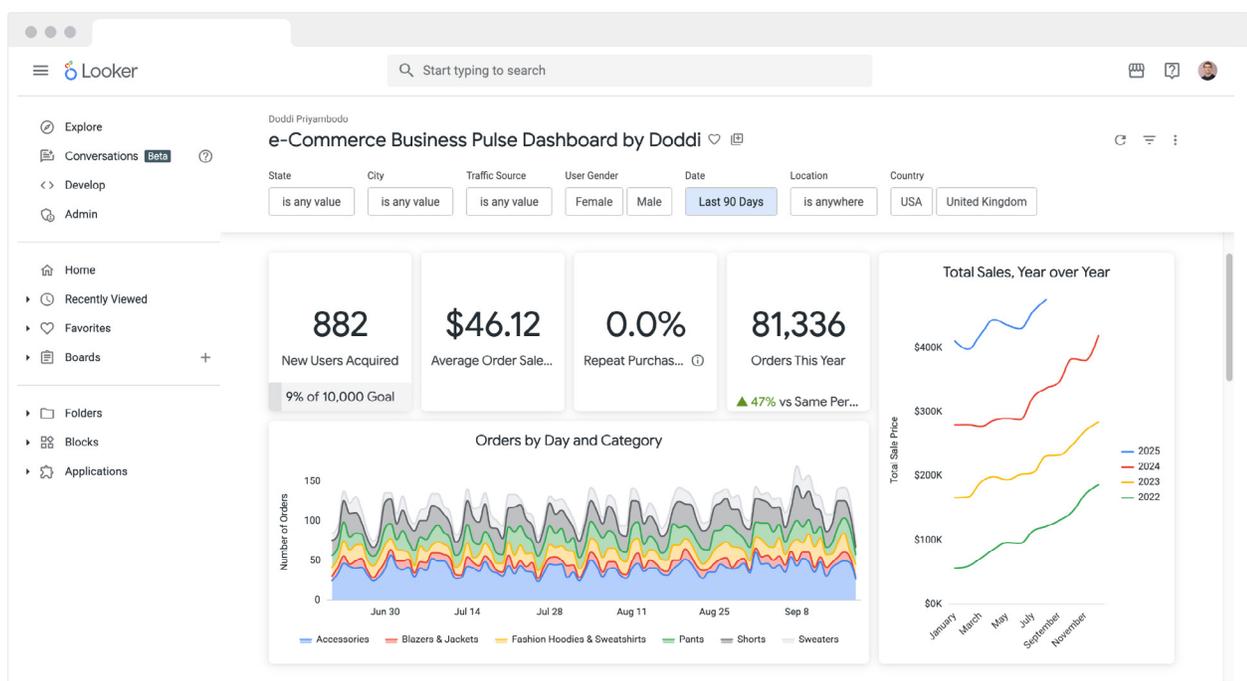




LookML powers full BI capabilities that enforce its use:

LookML is the foundation for a full-featured BI platform. As it powers the entire stack—from curated dashboards and visualizations to [embedded analytics](#)—users are inherently driven to use the LookML definitions. This enforced usage means that the same governed metrics used for the organization's standard BI reports are automatically the metrics used for Generative AI queries.

Since LookML is part of the enterprise-grade Looker AI and BI platform, any conversational insight it generates is immediately available for enterprise consumption.



This means that conversational insights aren't trapped in a one-off conversation, they can easily be saved, drilled into, added to dashboards, scheduled, distributed, or acted upon with just a few clicks.



LookML is enterprise ready:

True enterprise readiness for a semantic layer requires both operational maturity and a robust security and compliance posture. LookML is built on a proven platform, not an experimental technology, and has been deployed by thousands of enterprises over the past decade. This stability is paired with essential security competencies: the platform supports critical features like VPC Service Controls (VPC-SC) for network perimeter defense, Customer-Managed Encryption Keys (CMEK) for data sovereignty, and secure Private IP connectivity. By offering an established, proven framework with comprehensive, enterprise-grade security controls, LookML provides the stable and compliant foundation necessary for safely deploying Generative AI solutions with sensitive corporate data.

LookML prevents lock-in with flexible design:

A key concern for enterprise architecture is avoiding vendor lock-in. LookML addresses this by being an open, declarative language that defines the semantic layer on top of your existing data cloud architecture, rather than storing the data itself. Because LookML does not extract or move data, and instead executes queries in-database, the data remains sovereign and accessible in your chosen cloud data platform. Every query executed through LookML results in ANSI-compliant SQL, meaning that organizations can take any query generated and run it directly on their own in any other system. Any organization that builds its semantic model using LookML is retaining the flexibility to swap out any component of its data architecture.



Can the semantic layer be built elsewhere?

We've established that the six critical considerations—spanning collaboration, connectivity, governance, and agility—form the blueprint for a successful, enterprise-grade semantic layer. A common objection to the LookML approach is the belief that an equally robust semantic layer can be developed elsewhere.

To directly address this, the critical next question is a practical one:

Can a semantic layer be built and managed elsewhere—in the database or as a standalone service—and still comprehensively deliver on all six of these criteria?

This is not merely a technical detail; it is a strategic architectural decision that, upon rigorous evaluation, fundamentally enables—or inhibits—your ability to achieve the collaboration, governance, and flexibility your organization requires.

Broadly speaking, there are three options for where to develop and deploy a semantic model

01

In the database itself, either through purpose built “semantic views” or traditional database views (i.e. a medallion architecture where semantic layer is the “gold layer”)

02

Embedded within a BI offering

03

As a standalone solution, independent from the database or BI offerings





Option 01: Database “Semantic Views” or “Gold Layers”

A number of semantic layer offerings fit into this category, capitalizing on the perceived value proposition of institutionalizing data definitions at the source and then offering tool flexibility on top. These approaches generally do not add any security risk since the data is already present and has been secured & vetted.

But a deeper dive into database semantic views, using our six criteria above, demonstrates that these solutions leave quite a bit to be desired. By definition, they are limited to assets in that data warehouse/ lakehouse only, and cannot incorporate external assets. They lead to additional lock-in risk, as the data and the business logic are bundled together rather than separated, and they pose integration and enforcement risk to varying degrees as they are dependent on being correctly invoked by BI and AI interfaces. Lastly, due to their newness, they tend to be immature and not enterprise ready.

Importantly, they also create distance between those who are building the model and the experts that ultimately need to use the insights to drive decision making. While other categories of solutions bring together business experts and their direct data liaisons, the database-based semantic views often leave modeling to Data Engineers and DBAs, who are likely removed from the business teams, leading to increased risk of inaccurate definitions, lengthier release cycles, and increased friction.

Option 02: Embedded in the BI offering

This category of solutions offers a number of clear benefits. Generally, these solutions offer full BI capabilities, and also allow for data source independence, as almost all BI offerings can connect to different data sources.

However, a lot of discretion is needed as these offerings have different postures with respect to the criteria above. Some of these embedded semantic layers don't translate to SQL or are not necessarily enforced system wide, leading to lock-in risk. Most BI offerings have localized agentic strategies and are tough to fit into a broader agentic platform strategy. There also can be a wide range in establishment and enterprise-readiness of these semantic offerings.

While this category is an attractive place to develop a semantic model, discretion is needed to truly understand an offering's posture with respect to the criteria above.



Option 03: A standalone semantic offering

This category captures semantic layer offerings that are neither directly tied to a BI platform nor to a specific database. In theory, these solutions offer the most independence and flexibility in tool selection.

However, these solutions also suffer from a number of drawbacks. By definition, they are increasing complexity in the technical stack, and introducing bi-directional integration risk with every component of the stack. In an era of rapid pivoting of strategy from software companies, they are poorly positioned as a single neglected integration from a BI or database partner can cause massive disruption to workloads.

There is some appeal of solution agnosticism achieved by choosing a software framework in this category, but that appeal needs to be weighed against the compromises in the form of integration risk, security risk, awkward Agentic AI platform fits, and overall lock-in concerns.



Category capability matrix:

The below table directionally summarizes how each category performs against each of the six criteria we've established below:

	Unites Business and Data Teams	Supports Multiple Sources	Broader Agentic Platform Fit	Enforced and Activated Through BI	Enterprise Readiness	Lock In Risk
Database Semantic Views/Layer	Poor	Poor	Strong	Poor	Medium	Medium
Semantic Layer Embedded in BI	Strong	Strong	Medium	Strong	Medium	Medium
Separate Semantic Layer Solution	Medium	Strong	Poor	Poor	Poor/ Strong	Strong

This critical analysis demonstrates that the Semantic Layer embedded in BI architecture offers the optimal balance of governance, enforcement, and integration required for an Agentic Enterprise. While standalone solutions offer theoretical agnosticism, they fail on operational risk and enforcement. Looker's semantic layer is the definitive choice in the Embedded BI category, uniquely combining the strength of enforcement within the BI platform with the API-first interoperability and openness often sought in standalone solutions. This means you gain the governance of an integrated platform without sacrificing the flexibility to expose your semantic layer to third-party AI tools or services.





Your foundation for an intelligent enterprise

In the era of Generative AI, the question is not if you will use AI in your data strategy, but how you will ensure it is accurate, governed, and trustworthy. The promise of instant answers from complex data is immense, but entirely dependent on providing AI with the business context it fundamentally lacks.

The semantic layer provides this critical foundation. The architectural choice of where to build and manage the semantic layer is a defining strategic decision, impacting everything from data governance to technological agility. Based on the outlined rigorous evaluation framework, the Semantic Layer Embedded in BI architecture proves to be the most strategically sound choice. Looker, with its mature and proven LookML semantic layer, exemplifies this optimal approach, directly delivering on all six enterprise criteria—from uniting data and business teams to providing essential security competencies.

The path to becoming a data-driven organization, powered by reliable AI, begins with a single source of truth for your business metrics. By choosing Looker, you are making the strategic choice for a flexible, collaborative, and interoperable semantic layer that is enterprise ready, setting the stage for your enterprise's secure, trusted and intelligent future.

Get started with a Looker free trial.

[Contact us.](#)