D2A2
DIGITAL DATA ANALYTICS
ARTIFICIAL INTELLIGENCE

Expert Insight

Executive Guide: Bridging Enterprise Analytics & GenAI with Semantic Layers

  

Executive Guide: Bridging Enterprise Analytics and Generative AI with Semantic Layers

Dr. Arun Marar, PhD, MS
Dr. Prashanth Southekal, PhD, MBA
Tobias Zwingmann, MS

Generative AI (GenAI) is challenging the way businesses derive insights from data – from conversational data assistants to AI-generated reports. But its effectiveness hinges on the quality and consistency of the data it's fed.

One emerging solution is the semantic layer: a framework that translates raw data into consistent, business-friendly terms. By aligning technical data structures with business concepts, semantic layers can help GenAI systems generate more accurate, trustworthy, and relevant outputs.

This report explores the role of semantic layers in enterprise analytics: where they add value, where their limits lie, and how they fit into an AI-driven data strategy.

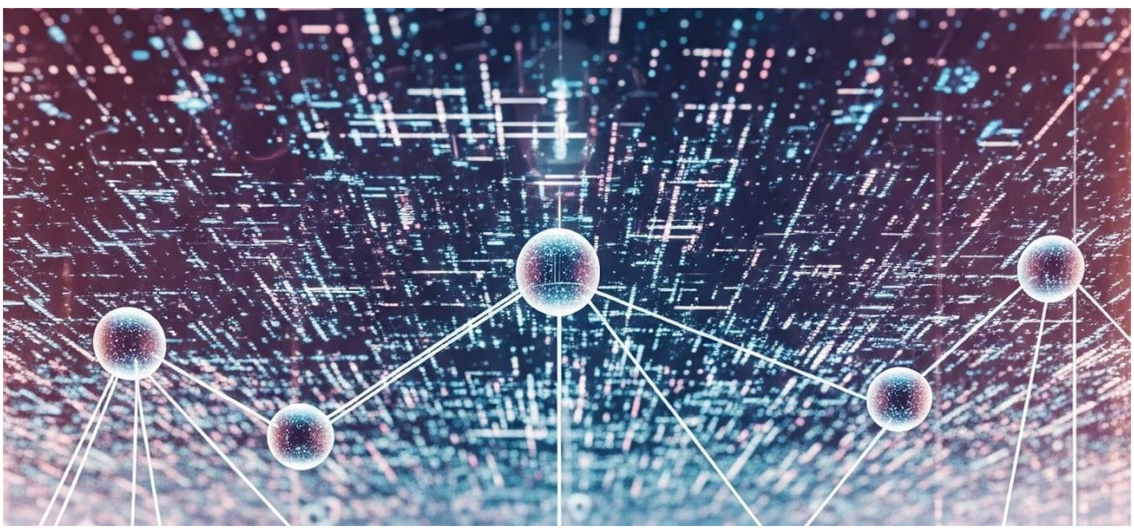

Understanding Semantic Layers

As enterprise data environments grow more complex, ensuring consistent definitions across teams has become a significant challenge. A semantic layer helps address this by translating raw data structures into standardized, business-friendly terms—establishing a shared vocabulary for metrics, dimensions, and logic.

At its core, the semantic layer abstracts technical data sources—like SQL tables or APIs—into familiar concepts such as “revenue,” “active customer,” or “churn rate.” This enables both people and AI systems to engage with data more intuitively and consistently, without needing deep expertise in underlying schemas. While semantic layers have long underpinned traditional BI platforms to support self-service analytics, their role is evolving. With the rise of GenAI tools that interact with data through natural language, the need for structured, governed semantics has become more pressing. Without a semantic foundation, AI systems are more prone to misinterpretation, ambiguity, and inconsistent outputs.

Closely related are knowledge graphs, which represent data as networks of entities and relationships, e.g., “Customer X purchased Product Y.” They add semantic richness by modeling real-world context and hierarchy. When paired with a semantic layer, knowledge graphs enhance GenAI’s ability to reason across data, providing both definitions and the relationships between them.

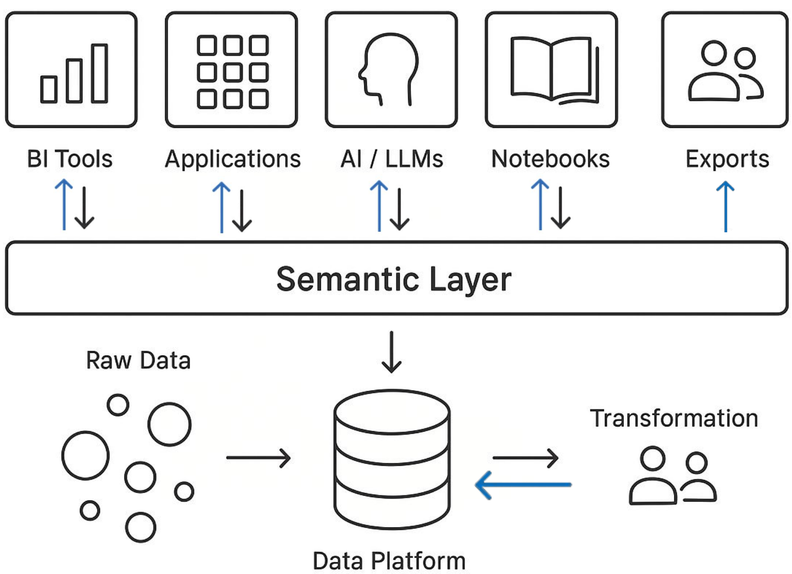

Simplified diagram of a semantic layer architecture

Fueling GenAI with Trustworthy Data

As organizations adopt generative AI to automate and augment analytics, a recurring obstacle emerges: inconsistent definitions and data quality issues often undermine the reliability of AI outputs. Generative models are only as good as the data they work with—a fact widely acknowledged across the industry.

A semantic layer addresses this challenge by enforcing standardized business definitions and quality controls upstream of AI usage. Concepts like “revenue” or “active customer” are provided with context and defined once, centrally, to be then applied consistently across dashboards, reports–and: GenAI queries. This ensures that regardless of the interface—be it a BI tool or a conversational assistant—the meaning remains stable.

According to Google, a robust semantic layer is “essential to deliver on the promise of trustworthy GenAI” for business intelligence by offering trust, business context, governance, and organizational alignment.[1]

In practice, this means that when an AI assistant is asked for last quarter’s sales, it doesn’t generate SQL against ad hoc or inconsistent source tables. Instead, it queries against centrally maintained, governed data models, ensuring the results align with established business definitions.

Grounding AI to Reduce Hallucinations and Errors

Large language models are known to produce plausible but inaccurate responses—a behavior known as “hallucination”—particularly when generating text outside their training data distribution. One effective way to mitigate this risk is by grounding AI outputs in well-defined, structured data. Semantic layers support this by serving as a source of truth, guiding AI interactions through pre-approved definitions, metrics, and timeframes.

[1] https://cloud.google.com/blog/products/business-intelligence/how-lookers-semantic-layer-enhances-gen-ai-trustworthiness

Experts increasingly see semantic layers as essential for trustworthy GenAI in business intelligence. For example, when asked to retrieve a “profit margin trend,” an AI model without guidance might misinterpret the metric or select the wrong data source. A semantic layer, however, directs the model to the correct calculation and scope, minimizing ambiguity.

David P. Mariani, CTO and Co-Founder of AtScale, a leading semantic layer platform provider, even goes as far as saying that "Generative AI can only deliver reliable answers when it's grounded in governed logic”.

Real-world benchmarks reflect this assumption. Research conducted by data.world found that GPT-4 answered enterprise data questions correctly just 16% of the time when using raw SQL, but accuracy jumped to 54% when the model was provided with a semantic knowledge graph—a more than threefold improvement.[2]

By filtering, contextualizing, and structuring enterprise data, semantic layers serve as both a quality gatekeeper and an interpretive guide. This enables GenAI systems to deliver more accurate, consistent, and reliable answers—grounded in business truth rather than inferred guesses.

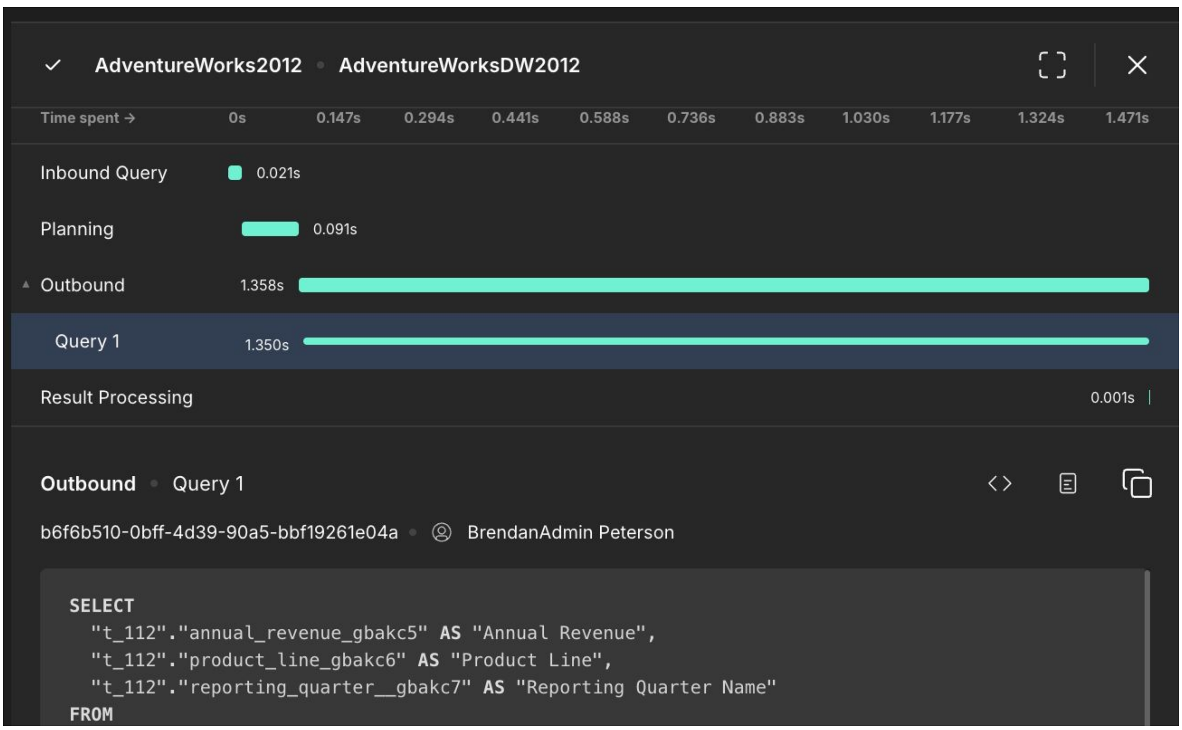

Automated SQL query planning in AtScale Semantic Layer

[2] https://arxiv.org/pdf/2311.07509

Generating AI-Powered Reports and Dashboards with a Semantic Layer

While semantic layers have traditionally focused on improving data accuracy and governance, their integration with generative AI introduces a new dimension: intuitive, accessible analytics. By anchoring AI tools to shared semantic definitions, organizations can move beyond conventional BI workflows and finally enable business users to ask questions in natural language—while still receiving structured, reliable outputs such as narratives, visualizations, or auto-generated dashboards.

Although semantic modeling is not new, its combination with GenAI elevates the usability and impact of analytics to a level that makes broader, enterprise-wide adoption far more feasible.

For example, a user might ask an AI assistant, "**What was our revenue in EMEA last quarter?**" Behind the scenes, the semantic layer ensures that:

- "**Revenue**" reflects the organization's approved definition (e.g., total net sales),
- "**EMEA**" maps to a specific regional grouping, and
- "**last quarter**" aligns with the company's fiscal calendar.

The resulting answer is consistent with what a data analyst or finance team would report—because it draws from centrally maintained, governed metrics.

A follow-up such as "How does that compare to the previous quarter, and can you show it as a chart?" builds on that same semantic logic, allowing the AI to generate a valid comparison and visual output using certified data.

Real-World Examples

Several platforms illustrate how this works in practice. dbt Labs, for instance, developed the “Ask dbt” assistant—an AI chatbot that interfaces with the dbt Semantic Layer. The assistant responds to business queries by drawing on semantic definitions maintained in dbt models, generating tabular or narrative answers that reflect consistent logic across teams.

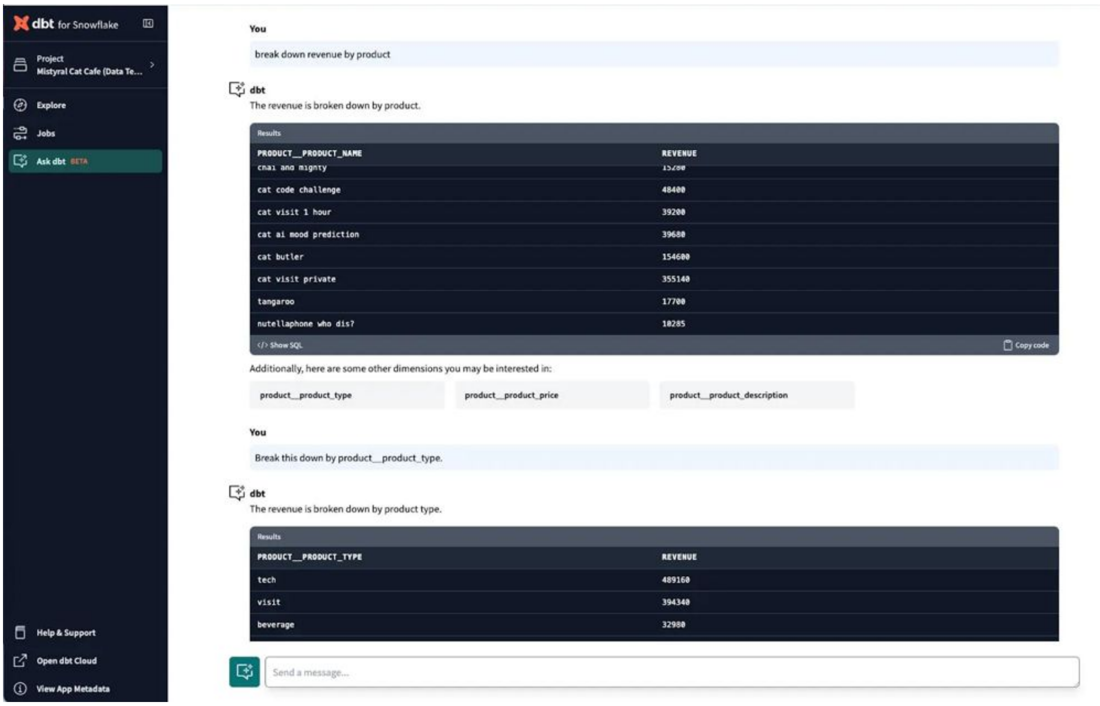

Example: dbt Labs’ Ask dbt in Snowflake

Similarly, AtScale provides a universal semantic layer that enables GenAI tools to interpret natural language questions and translate them into optimized, governed queries. When a user asks “**What is our total revenue by product line for last quarter?**”, the platform not only ensures that definitions like “Total Revenue” and “Product Line” are applied consistently across data sources such as Snowflake, BigQuery, or Databricks, but also optimizes the query for speed by leveraging pre-computed values in the layer, if they are available.

In this model, AI becomes a dynamic front-end for enterprise analytics and business users benefit from faster access to insights, while data teams gain confidence that AI-driven exploration remains within governed boundaries.

Role-based Data Access via AI Assistants

As generative AI becomes embedded in enterprise analytics, questions of who can access what data take on new urgency. AI assistants can quickly surface insights, but without proper controls, they risk exposing sensitive or unauthorized information.

Semantic layers help mitigate this by integrating with role-based access control (RBAC) systems and enforcing data governance policies at query time. This ensures that AI-generated outputs follow the same access rules that govern dashboards, reports, and SQL queries—down to the user's role, department, or region.

For example, a finance manager asking, "Show me last quarter's gross margin by region," may receive detailed results, while a sales rep with lower data privileges may only see aggregated or redacted figures for the same question. These permissions are not managed within the AI model itself, but at the semantic layer, ensuring consistency, auditability, and compliance.

In effect, the semantic layer acts as a policy-aware intermediary between GenAI tools and enterprise data systems. It filters not only what the AI can access, but also how it interprets and presents the data, based on user roles, data sensitivity, and contextual rules.

This approach has two key benefits:

- **Security and Compliance**: AI systems remain aligned with corporate governance, privacy, and regulatory requirements (e.g., GDPR, HIPAA, internal data classification).
- **User Trust and Adoption**: Business users gain confidence that AI assistants will not overstep access boundaries, making adoption safer and more scalable.

By combining AI's natural language capabilities with the semantic layer's control logic, organizations can safely democratize data access and deliver relevant insights to users without compromising security or compliance.

Example: Role-Based Responses from an AI Assistant

User Prompt:
“What was our revenue by product line for Q2?”

User Role	Data Access Level	AI Assistant Response (via Semantic Layer)
Sales Rep	Limited: Own region & summary metrics	*“Revenue in your region (EMEA) totaled **$12.5M** across 3 products. Detailed product-level breakdown is restricted.”*
Sales Manager	Moderate: All regions, summary + product	*“Q2 revenue by product: Product A: **$8.1M**, Product B: **$6.3M**, Product C: **$5.9M** (Global figures).”*
Finance Analyst	Full: Detailed access to all dimensions	*“Q2 global revenue by product and region: EMEA – Product A: **$3.2M**, Product B: **$2.5M**; NA – Product A: **$4.9M**, etc.”*
Executive	High-level KPIs, sensitive data hidden	*“Total Q2 revenue: **$20.3M**. Key drivers: Product A and B growth. Drill-down available on request.”*

Example responses generated using the same AI assistant, but governed by role-based access enforced through the semantic layer.

Not a Silver Bullet: When Raw Data is Still Needed

While semantic layers offer significant benefits for structured, governed analytics, especially when paired with generative AI, they are not a universal solution. As organizations explore more advanced AI use cases, it becomes clear that there are scenarios where raw data access remains essential.

Semantic layers are best suited for answering questions rooted in well-defined, curated business metrics: revenue, churn, customer segments, etc. But when an AI assistant—or an analyst—asks a novel or exploratory question that falls outside these established definitions, the semantic layer may not be equipped to respond. This could include combining data sources in new ways, analyzing unstructured content like customer support logs or social media data, or querying dimensions not yet modeled.

For example, a data scientist investigating behavioral trends in clickstream logs, or a machine learning engineer building models from semi-structured JSON data, will often need direct access to raw data in data lakes or event stores. These sources are rich in information but typically not represented in a semantic model. Similarly, certain investigative workflows—such as debugging data anomalies or tracking individual transaction histories—may require bypassing the abstraction layer to inspect the source records directly.

Another common edge case is AI model training. Fine-tuning a domain-specific large language model may involve ingesting raw documents, transcripts, or user-generated content that lie well outside the scope of a semantic layer. In these contexts, the semantic layer is not a substitute for ETL pipelines, data preparation tools, or flexible query frameworks.

Even within business intelligence, there are operational realities to consider. A newly introduced metric may not yet be modeled, or organizational changes may lag behind updates to the semantic definitions. In such cases, fallback access to raw data via SQL, notebooks, or APIs, ensures that insights aren't blocked by model coverage gaps.

A Complement, Not a Replacement

Recognizing these limitations does not diminish the value of semantic layers. Instead, it emphasizes their role as one component of a broader, modern data architecture. Forward-thinking organizations mitigate these gaps by:

- Continuously evolving and expanding their semantic models;
- Incorporating fallback mechanisms for free-form or low-context queries;
- Combining semantic layers with complementary tools, such as vector databases, context layers, or hybrid query systems.

In short, semantic layers work well for the "known knowns" of enterprise data. But AI systems still need a path to the raw and unstructured when operating in more dynamic, experimental, or edge-case scenarios. A balanced architecture that maps data requirements to different use cases ranging from classical BI workflows to more exploratory data science questions ensures full AI capability across the whole spectrum.

When to Use Semantic Layer vs. Raw Data Access

Scenario	Semantic Layer	Raw Data Access
Standard business KPIs (e.g., revenue, churn rate)	Ideal – pre-defined, governed metrics ensure consistency	Not required
Natural language queries by business users	Supports accurate AI responses using shared definitions	Risk of inconsistent results
Exploratory or ad hoc analysis	Limited – depends on model coverage	Necessary for flexibility
Model training and experimentation	Not designed for raw ingestion	Essential for access to full datasets
Debugging / root cause analysis	Limited granularity	Full visibility into raw records
Rapid schema changes / new metrics	May lag behind data updates	Immediate access possible

Comparison Table: Semantic Layer vs. Raw Data Access

Conclusion

The emergence of Generative AI in business analytics has redefined how organizations interact with data—enabling users to ask complex questions in plain language and receive immediate, insightful responses. But this new capability also surfaces familiar challenges: inconsistent metrics, questionable data lineage, and governance gaps that can erode trust in AI outputs.

Semantic layers offer a practical and increasingly essential foundation for addressing these issues. By embedding consistent definitions, contextual understanding, and access controls into the data workflow, they ensure that GenAI tools operate on a shared understanding of the business. Whether powering a chatbot, generating a report, or assembling a dashboard, AI systems grounded in a semantic layer can produce results that are accurate, reproducible, and aligned with organizational standards.

Organizations implementing this approach have reported tangible benefits, from improved data trust and reduced errors to faster decision-making and broader adoption of AI tools by non-technical users. Technologies from providers such as AtScale, dbt Labs, and Google's Looker demonstrate how semantic models can bridge traditional BI and modern AI-driven interfaces in a unified analytics architecture.

That said, semantic layers are not a complete solution in themselves. They require thoughtful design, ongoing maintenance, and must be supported by solid data engineering practices. They do not replace the need for direct access to raw data in exploratory or experimental workflows. But they do dramatically increase the reliability of AI in structured business contexts.

In a landscape where GenAI is poised to reshape how businesses use data, the semantic layer serves as a critical enabler. It ensures that AI-driven insights are not just fast and flexible, but also credible and aligned to the organization's data truths. Ultimately, it is this combination of innovation and integrity that allows generative AI to evolve from a promising tool into a trusted partner in enterprise decision-making.

Contributors: The Analyst Team

Arun Marar, Ph.D.

Arun Marar, Ph.D. is an expert at implementing AI/ML models utilizing a wide range of technologies such as supervised and unsupervised learning, text analytics and natural language processing and deep learning. Arun has consulted with clients all around the world across domains such as Banking and Financial Services, Retail, Healthcare and Transportation Logistics. Arun is also an expert in optimization, and this has helped him aid companies in moving up the value chain of analytics from predictive analytics to prescriptive analytics.

Prashanth Southekal, PhD, MBA

Prashanth Southekal, PhD, MBA, ICD.D is a Consultant, Author, and Educator. He has worked and consulted for over 80 organizations including P&G, GE, Shell, Apple, and SAP. He is the author of three books — “Data for Business Performance”, "Analytics Best Practices” and “Data Quality”. His second book, ANALYTICS BEST PRACTICES was ranked #1 analytics books of all time in May 2022 by BookAuthority. Apart from his consulting and advisory pursuits, he has trained over 4,500 professionals worldwide in Data and Analytics.

Tobias Zwingmann, MSc.

Tobias Zwingmann is Consultant, Author, and Educator and is on a mission to help companies implement machine learning and AI faster while delivering meaningful business value. He is the author of 2 books “AI-Powered Business Intelligence” and “Augmented Analytics”, He brings more than 15 years of professional experience working in a corporate setting where his responsibilities included building data science use cases and developing an enterprise-wide data strategy.

www.d2a2.ai

D2A2
DIGITAL DATA ANALYTICS
ARTIFICIAL INTELLIGENCE