

Semantic Layer & Conversational AI


Trusted, AI-ready analytics across the Snowflake + Power BI ecosystem

Leaders cannot trust AI-assisted decisions when the same KPI means different things across dashboards, SQL scripts, spreadsheets, and applications. Conversational AI does not hide fragmented business logic. It exposes inconsistency directly and at scale.

Persistent's point of view: Define business meaning once in a governed semantic foundation close to the data platform, then reuse it across dashboards, analysts, AI assistants, agents, and applications.

The offering

We design, build, and govern a single semantic layer in Snowflake enabling every dashboard, report, AI assistant, and application to operate from the same trusted business definitions.



One Governed definition

Approved entities, measures, relationships, hierarchies, synonyms, and access rules reused across every consumer.



Trusted Conversational Analytics

Natural-language answers grounded in approved business definitions rather than raw objects or duplicated logic.



Modernized BI estate

Reduced model sprawl, reconciliation effort, duplicated KPI logic and maintenance burden.



Governance by Design

Security, masking, lineage, and access boundaries applied consistently across analytics, and AI experiences.

Start here - Semantic Layer Readiness Assessment

A focused 2–3 weeks engagement to identify where fragmented KPI logic is slowing analytics and AI adoption, validate a priority business domain, and define a practical path to governed conversational analytics.



Discover

Identify where fragmented KPI definitions are creating trust, governance, and AI adoption challenges.



Validate

Prove one to two governed KPIs end-to-end across a dashboard and a conversational experience.



Plan

Define the target architecture, reusable standards, delivery backlog, operating model, and phased roadmap.

What you get:



Readiness and model-sprawl assessment



Reference architecture and consumption recommendations



Validated priority domain and one to two KPIs



Validation approach for trusted natural-language answers



Business glossary and semantic-definition starter set



Phased roadmap, backlog, operating model, and indicative cost view

How we scale it

Pilot 2 Months



Validate semantic layer for one priority domain; deliver 1–2 business KPIs; enable consumption through AI, Power BI, and SQL.

Productionize 2–3 Months



Productionize certified metrics with governance, monitoring, and CI/CD; expand KPI coverage within the same domain; establish operating model and standards.

Scale 3–6 Months

Onboard additional business domains; expand certified metric catalog; enable conversational analytics through Cortex Analyst and Copilot; drive self-service adoption.



Mature 6–12 Months



Establish enterprise semantic layer foundation; optimize governance, performance, and cost; enable enterprise-scale analytics, AI, and agentic experiences.



Outcomes: One Definition, Zero Divergence | Trustworthy AI | Lower BI Maintenance Cost | Onboard New Tools in Hours | Governance Enforced Across Every Consumer.

Why Persistent

Engineering-led depth in Data & AI, with hands-on expertise across the Snowflake and Microsoft ecosystems — we turn this offering into a governed, production-grade reality.

Ready to begin? Start with a 2–3 weeks Semantic Layer Readiness Assessment. Assessment. To start the conversation, contact your Persistent account team. [Add account contact / email]