

GenAI-Ready Semantic Layer Checklist

A simple yes/no assessment to evaluate
your semantic layer readiness for GenAI.

Instructions: Check each box that
applies to your current semantic
layer. Count your total checkmarks
to see your GenAI readiness score.

1. Semantic Modeling & Definitions

Your semantic layer needs standardized business definitions that work across all tools.

- Business metrics have consistent definitions across all systems
- Dimensions and hierarchies are standardized (time, geography, products, etc.)
- Business rules and calculations are centrally defined
- Metadata includes business context, not just technical definitions

2. Query & Integration

GenAI applications need programmatic access to your semantic models.

- APIs available for programmatic access to semantic models
- Standard SQL access to semantic models
- Semantic models can be exported in standard formats (JSON Schema, OpenAPI)
- Integration with existing GenAI tools and frameworks





3. Performance & Optimization

GenAI applications generate high query volumes that need fast, cost-effective responses.

- Automated creation and management of aggregates/cubes
- Query result caching and reuse capabilities
- Cost-aware query optimization and execution planning
- Horizontal scaling to handle high concurrent query volumes





4. Security & Governance

GenAI applications need enterprise-grade security and data governance.

-  Role-based access controls for semantic models
-  Row and column-level security controls
-  Complete audit logs and query lineage tracking
-  Data masking and PII protection capabilities





5. GenAI-Specific Features

Advanced capabilities that directly enable GenAI applications to work safely with your data.

-  Natural language to SQL query translation
-  Semantic model ontology APIs for LLM consumption
-  Query guardrails and safety checks for GenAI queries
-  Quality monitoring and hallucination detection





6. Developer & Operations Experience

Your team needs modern development workflows and operational visibility.

-  Version control and CI/CD workflows for semantic models
-  Infrastructure-as-Code deployment capabilities
-  Comprehensive observability: performance metrics, query costs, usage analytics
-  Migration tools for existing semantic models (LookML, dbt, etc.)

Your GenAI Readiness Score

Count your checkmarks and find your readiness level:

Score	Status	What This Means
20-24	 GenAI Ready	You have a strong foundation for scaling GenAI initiatives
16-19	 Mostly Ready	Address a few key gaps before enterprise GenAI deployment
10-15	 Needs Work	Significant investment required for GenAI success
Below 10	 Not Ready	Focus on foundational semantic layer capabilities first

What to Do Next

Based on your score:

- Review the unchecked items - these are your priority gaps
- Focus first on Semantic Modeling & Definitions (Section 1) and Security & Governance (Section 4)
- Evaluate whether your current platform can be enhanced or if you need a new solution
- Create an implementation timeline that aligns with your GenAI deployment goals

Need help getting GenAI-ready? AtScale's Universal Semantic Layer is designed specifically for GenAI workloads.

Schedule a Demo

See exactly how AtScale's semantic layer would work with your retail data and tools

Download the Buyers Guide

Get the detailed playbook for successful retail semantic layer deployment

AtScale's Universal Semantic Layer empowers organizations to deliver trusted, consistent insights for every user—from dashboards to AI agents.

By aligning business definitions across tools and teams, AtScale eliminates conflicting metrics, improves GenAI accuracy, and enables AI agents to act on governed, reliable data. Built on open semantics, AtScale ensures interoperability across platforms—so you can scale without vendor lock-in. The result: faster decisions, greater trust, and a single source of truth that works across your entire data ecosystem.