

# BARC

## Engineering AI, Metadata & Context for Self-Service Agentic Analytics

Author: Kevin Petrie

Publication: February 2026

### Abstract

This Spotlight report defines what an agentic analytics fabric is, why it matters, and how it can drive self-service. We conclude by exploring use cases and an example solution from the vendor Promethium.

Research sponsored by:



# Table of Contents

- Introduction..... 3**
- Market Overview ..... 3**
- Architecture ..... 4**
  - Analytics agent..... 5
  - Context Engineering ..... 5
  - Data Access Layer ..... 6
  - Govern..... 6
- Use cases ..... 7**
  - Sales forecasting ..... 7
  - Supply chain management ..... 7
  - Price optimization..... 7
  - Retail product quality & returns ..... 7
- Solution Profile: Promethium..... 8**
- Getting Started ..... 9**
- About BARC.....10**
- About Promethium..... 11**

## Introduction

Analytics agents evoke the popular vision of self service. In this vision, business teams explore data and analyze opportunities or risks without burdening technical specialists. They make smarter, faster decisions at a lower cost for the enterprise.

But despite impressive gains in model intelligence, we are not there yet. Most agentic analytics tools still fail to generate production-level insights because they cannot access distributed data on a real-time basis. They also cannot organize metadata, which starves them of the context needed for accurate, personalized, and role-specific conversations. Such limitations prevent business users from achieving self-service. Instead, they must explore and analyze data in a slow, impersonal, and inefficient way, unsure which outputs to trust.

The time has come for a new architectural approach: the agentic analytics fabric. This emerging approach brings together agent engineering, context engineering, and real-time access to distributed datasets. Designed and implemented well, the agentic analytics fabric harnesses diverse metadata to enable accurate, AI-driven insights. It empowers humans to spend more time thinking and less time navigating data. And self-service becomes achievable.

While this market is nascent, we believe data leaders must scope requirements in this area and evaluate their options to meet them. This Spotlight report defines what an agentic analytics fabric is, why it matters, and how it can drive self-service. We conclude by exploring use cases and an example solution from the vendor Promethium.

***The agentic analytics fabric harnesses diverse metadata to enable self-service***

## Market Overview

We start by considering the state of adoption. While a recent BARC survey shows that 50% of organizations have agents in production, only 27% use them for BI or analytics. Complexity and lack of context and accuracy hold the rest back. Hybrid and multi-cloud environments contribute to the problem: nearly half (47%) of organizations have AI source data that sits in more than one physical location.

### Market evolution

Semantic layers address part of the problem. Traditional BI tools have long contained semantic layers that derive standard business metrics from database records. More recently, vendors offered standalone semantic layers that span multiple BI tools and data sources. Others now bundle semantic layer capabilities with catalogs or data platforms. But many of these approaches miss contextual factors such as user intent, history, and vocabulary—or restrict you to a specific type of environment.

The agentic analytics fabric (or simply “fabric”) goes further to codify such inputs across heterogeneous environments. It activates wide-ranging metadata on demand to help agents interpret data and analyze business scenarios. This gives analytics agents and humans the intelligence, context, access, and confidence they need to make smarter decisions faster. It enables scenarios such as the following.

- The sales VP asks her analytics agent about team performance. She receives a chart that applies her favorite KPIs to her territories, along with explanations of why certain reps might be underperforming.
- A supply-chain manager, meanwhile, evaluates warehouse inventories while the agent reconciles inconsistent product names in the background. To double check, he clicks to inspect and adjust this master data himself.

Human-agent conversations like these clarify user needs and incorporate tribal knowledge, making the fabric more intelligent over time. Vendors are starting to build such capabilities into their analytics, data management, and data fabric platforms.

## Architecture

Let's consider the architecture that makes this possible. We focus first on the analytics agent, then move down to the context engineering and data access layers.

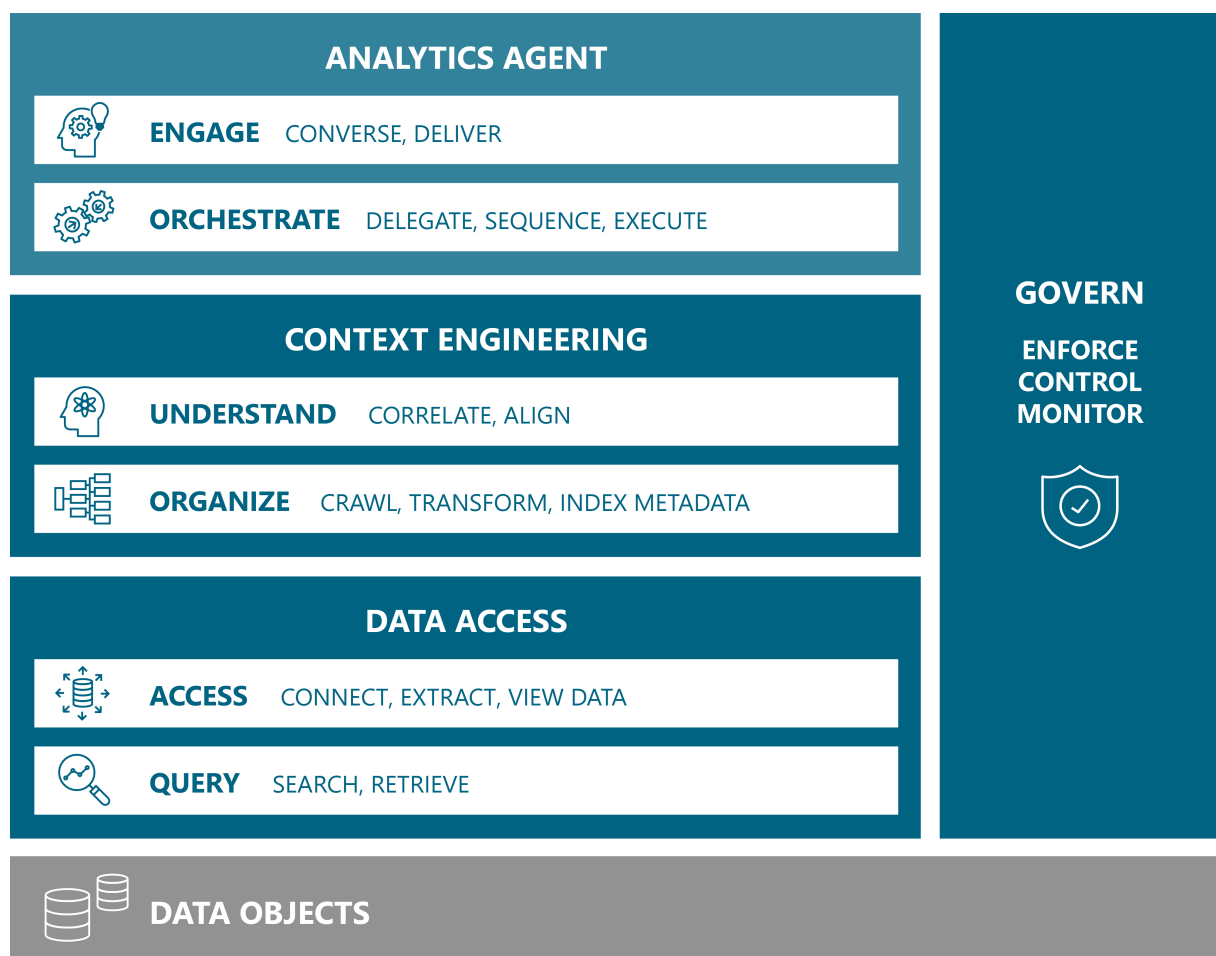


Figure 1: Architecture of the Agentic Analytics Fabric

## Analytics agent

The analytics agent engages with the business users and orchestrates necessary workflows with context and enterprise data to deliver insights.

### Engage

The analytics agent engages with the human business expert. It converses intelligently with the decision-maker in sales, logistics, finance and so on, delivering them personalized insights in context. These humans receive relevant data and nuanced explanations that fit their expectations, vocabulary, and usage histories. They also get transparent views into data lineage and model interpretations, along with the ability to adjust outputs.

### Orchestrate

The analytics agent orchestrates workflows using the context engineering layer (described below) as a foundation. It delegates actions to various models and tools: for example, the GenAI model converts text inquiries to SQL commands while an agentic tool populates visual templates with the query results. The agent creates sequences of such tasks, executes them, and monitors the status. All this orchestration depends on the context engineering layer to actively manipulate technical, business, operational, and governance metadata.

## Context Engineering

We can view the context engineering layer as the next generation of a semantic layer. It brings order to chaos by understanding business requests and organizing the metadata to resolve access and queries.

### Understand

The context engineering layer seeks to “understand” how the elements fit together. For example, it correlates user type and likely intention. If a marketing manager requests summaries of all digital interactions with their 25 largest customers, the semantic layer should tag this as a request for account-based marketing to guide similar future inquiries. This layer also correlates user type and terminology. If a finance manager asks the agent to “give an A/R update,” he needs an answer that aligns with his country’s definition of accounts receivable.

### Organize

Agents need well-organized data, which means they need well-organized metadata. The semantic layer crawls across data sources throughout the enterprise to identify metadata. It transforms this metadata, for example by filtering, merging, or reformatting it. The metadata ranges from business table names to technical schemas, user roles, and access rights. The semantic layer creates an over-arching index of all these elements to guide everything that follows.

The context engineering layer organizes the following types of metadata:

- **Technical** metadata defines the structure, format, version, storage, and processing requirements of data objects. It enables business users as well as data analysts, scientists, and engineers to prepare and consume data.

- **Business** metadata describes data products, terminology, definitions, ontologies, classifications, and KPIs. It makes data understandable and discoverable for business users, often through a shared glossary.
- **Operational** metadata provides context such as lineage, quality metrics, and consumption patterns as well as the intent of different roles and users. It helps technical teams orchestrate analytics and AI workflows.
- **Governance** metadata specifies data quality scores, access rights, privacy rules, and guardrails. It helps ensure data, models, and agents comply with governance policies and regulations.

## Data Access Layer

Now we come to the data access layer. As the name suggests, this layer accesses and queries the distributed data objects.

### Access

The data access layer connects to data sources such as databases, data warehouses, lakehouses, or telemetry data stores. It accesses the data, for example by extracting or replicating it to a central repository such as a lakehouse. It also might create virtual views that point to distributed data objects, including tables, documents, images, and so on, without moving them. These views bring together wide-ranging inputs that enrich analytics.

### Query

The data access layer performs various queries on behalf of the analytics agent. It searches tables or other objects and retrieves values, often using familiar SQL commands. Certain questions might require querying multiple data platforms simultaneously. As an example, revenue data might reside on a modern data platform such as Snowflake or Databricks, while customer data resides in a CRM such as Salesforce. In this case, the data access layer needs to support federated querying capabilities to query multiple databases simultaneously and combine the outputs. Building on our prior example, this would help a marketing VP get a global view of digital customer interactions across Asia, Europe, and North America.

## Govern

Enterprises must enforce governance policies, rules, and standards across all layers of the fabric. Agents might propagate inaccurate data, breach privacy rules, or recommend damaging business decisions. Organizations mitigate such risks by strengthening their governance policies, rules, and standards. And the fabric helps by enforcing policies with technical controls and monitoring compliance. For example, role-based access controls ensure that only authenticated users or agents perform only authorized tasks on permissible data sets. The data access layer uses governance metadata to identify and obfuscate personally identifiable information (PII) to satisfy privacy regulations. In addition, the context engineering layer documents all activities and shares that metadata with third-party catalogs and auditing tools.

## Use cases

To deepen our knowledge, here are four example use cases: sales forecasting, supply-chain management, price optimization, and retail return processes.

### Sales forecasting

A sales leader asks an analytics agent to forecast quarterly revenue. The fabric applies standardized pipeline stages, quota definitions, and regional hierarchies from business metadata to ensure consistency with prior reports. Operational metadata about data freshness and lineage helps the agent quantify the confidence level and highlight late or incomplete CRM updates. Governance metadata enforces territory-level access controls so that managers review only the pipelines they are authorized to see.

### Supply chain management

A supply-chain manager assesses inventory levels and costs across warehouses. The fabric uses technical and business metadata to reconcile inconsistent product IDs, supplier names, and units of measure. The agent combines distributed data from ERP, logistics, and demand-planning systems through queries and virtual views, guided by operational metadata. Governance metadata flag low-trust data sources and prevent the use of non-compliant supplier data in the recommendations.

### Price optimization

A pricing analyst asks the agent to simulate margin impacts under different discount scenarios, and the fabric aligns cost, revenue, and elasticity metrics using shared KPI definitions. Operational metadata captures prior pricing experiments and user intent, allowing the agent to explain why certain price changes are likely to outperform others in specific regions or channels. Governance metadata ensures sensitive cost inputs are masked and that recommendations comply with internal pricing policies.

### Retail product quality & returns

A retail manager investigates rising return rates, and the semantic layer links sales transactions, product attributes, customer feedback, and warranty data through common business definitions. The agent uses operational metadata on data lineage and quality to distinguish true defect signals from reporting noise. Governance metadata protects customer PII in return records while still allowing the agent to surface actionable quality insights by product line and supplier.

## Solution Profile: Promethium

Promethium.ai has launched a new AI Insights Fabric platform that enables business and data teams to “talk” to data for self-service agent analytics. Led by CEO Prat Moghe, Promethium is venture-backed by investors such as Insight Partners, .406 Ventures and Zetta Venture Partners. Promethium’s platform is in essence an agentic analytics fabric that accesses, organizes, understands, queries, and governs distributed data in data platforms and SaaS applications, wherever they reside, using context engineering to improve the accuracy of its real-time outputs.

The Promethium AI Insights Fabric has three primary components:

1. The Universal Query Engine serves as the data access layer, performing federated real-time querying of data across source platforms, SaaS applications and cloud sources. It governs these functions with fine-grained access controls.
2. 360 Context Hub ingests and curates technical, business, operational, and governance metadata to enable context engineering across enterprise environments.
3. Data Answer Agent Mantra™ provides agentic engineering to answer questions and engage business users while aligning with their individual role, usage history, and objectives.

Promethium offers two modes.

1. Business and data analysts can use the Mantra™ agent to discover data, ask business questions, develop data products, and gain insights.
2. Any enterprise model/agent/tool or application can connect with the AI Insights Fabric via MCP or APIs and ask questions and receive dynamically generated data answers and insights with explainability and context.

The Promethium’s agentic offering is ideal for large enterprises with a distributed and diverse set of data platforms and SaaS applications, and with a strong security, compliance and governance posture. Enterprises can deploy Promethium within their virtual private cloud to ensure that no data actually leaves the enterprise boundary.

## Getting Started

As analytics agents become the primary interface for decision-making, enterprises must break the metadata and data access bottleneck that prevents them from delivering trustworthy insights. A modern fabric unifies metadata and real-time data access with agent engineering to break this bottleneck and thereby realize the vision of self-service analytics. It enables analytics agents to deliver timely, reliable, and actionable intelligence as they converse with business users.

- **Establish a unified metadata and context foundation.** Data leaders should prioritize consolidating metadata across distributed platforms into a shared semantic and governance layer that AI agents can reason over. This foundation must encode business definitions, lineage, quality signals, and access policies, and user habits in machine-consumable form. Without it, conversational and autonomous analytics will remain brittle, opaque, and difficult to scale.
- **Operationalize analytics agents for high-value decisions.** Business and data teams should jointly identify a small set of repeatable, high-impact use cases—such as forecasting, pricing, or supply chain optimization—where agents can augment human decision-making. These deployments should emphasize explainability, traceability to source data, and alignment with business KPIs. Early wins will build confidence while informing broader rollout strategies.
- **Define requirements and start evaluating commercially packaged offerings that address part or all of the agentic analytics fabric.** As with many aspects of data management, commercial software offers lower cost and complexity than homegrown code. Define your must-have product criteria, including ease of use, governance features, performance, scale, and ecosystem support, and build a short list of commercial solutions that might fit the bill. Give higher priority to independent offerings that do not favor a specific data platform or analytics tool.

## About BARC

BARC is the leading analyst firm for data & analytics, AI, corporate performance management (CPM) and ESG with a reputation for unbiased and trusted advice. Our expert analysts deliver a wide range of research, events and consulting services for the data & analytics community. Our innovative research evaluates software, vendors and service providers rigorously and highlights market trends, delivering insights that enable our customers to innovate with data, analytics and AI. BARC's 25 years of experience with data strategy & culture, data architecture, organization and software selection helps clients transform into truly data-driven organizations.

### Research

BARC user surveys, software evaluations and analyst advisory services along with expert driven content such as research notes, trend analysis and blogs give organizations the confidence to make the right decisions. Our independent research gets to the heart of market developments, evaluates software, vendors and service providers thoroughly and gives valuable ideas on how to turn data, analytics and AI into added value and successfully transform businesses.

### Consulting

The BARC consulting practice is entirely focused on translating companies' requirements into future-proof decisions. The holistic advice we provide helps companies successfully implement their data & analytics strategy and culture as well as their architecture and technology.

BARC's research and experience-founded expert input sets organizations on the road to the successful use of data & analytics, from strategy to optimized data-driven business processes.

### Events

At BARC events, leading minds and industry experts come together to share insights and drive innovation. Our conferences, roundtables and online webinars attract over 10,000 participants annually, offering a unique blend of information, inspiration and interactivity. These events provide a platform to exchange ideas with peers, explore emerging trends and gain expert perspectives on market developments.

By engaging with thought leaders and industry practitioners, participants discover actionable strategies to enhance their business and stay ahead in the evolving world of data & analytics.



## About Promethium

Promethium enables enterprises to get accurate insights by “talking” to all their business data. With its AI Insights Fabric, all AI models, agents and business tools can get access to both distributed data and context seamlessly, dramatically reducing the cost and effort of context engineering and data access to enable production-level AI outcomes. Business and data teams can now become significantly more productive, by building data products on the fly or enabling self-service with AI-ready data, across all tools, models, and data platforms. A Gartner Cool Vendor, Promethium is backed by top investors like Insight Partners, .406 Ventures, and Zetta Venture Partners. Leading enterprises like National Grid trust Promethium to deliver faster and accurate insights at scale. Learn more at [promethium.ai](https://www.promethium.ai)

### Contact Info

**Promethium**  
101 Jefferson Dr  
Menlo Park, CA, 94025  
USA  
[www.promethium.ai](https://www.promethium.ai)



# BARC

Data Decisions. Built on BARC.

[www.barc.com](http://www.barc.com)

## Germany

BARC GmbH  
Berliner Platz 7  
D-97080 Würzburg  
+49 931 880651-0

## Austria

BARC GmbH  
Hirschstettner Straße 19 / I / IS314  
A-1220 Wien  
+43 660 6366870

## Switzerland

BARC Switzerland GmbH  
Buchhaldenstrasse 7  
CH-5442 Fislisbach  
+41 76 340 35 16

## United States

BARC US  
13463 Falls Drive  
Broomfield, CO 80020  
USA