



Unlocking Self-Service Data at AI Scale

A New Framework for Delivering
Real-Time, Explainable Data Answers
Across Your Enterprise



Introduction

AI has the potential to transform how organizations make decisions, innovate, and compete. The promise is that with the right data, AI can deliver real-time insights, automate complex processes, and unlock new opportunities. But despite this promise, most enterprises are struggling to make AI deliver meaningful results for now.

The problem isn't necessary with AI itself - it's with the data and the processes we are using to get it.

Enterprise data is notoriously complex. It's spread across multiple platforms, cloud environments, and operational systems. Connecting this data is time-consuming and resource-intensive. Traditional data processes rely on rigid pipelines designed for batch processing - too slow for the real-time insights that AI demands. Worse yet, data often lacks the essential context needed for AI to produce accurate and trusted results.

These challenges explain why so many organizations fail to realize the full potential of AI yet. Data teams are forced to spend the majority of their time wrangling data instead of delivering insights. Business users are left waiting for answers, delaying decisions and opportunities. And without the right context, even the most powerful AI models risk producing incomplete, inaccurate, or misleading outputs.

To unlock AI's full potential, organizations must rethink how they connect, prepare, and access data. What's needed is a new approach designed for speed, agility, and trust - one that enables AI to deliver insights in real time without relying on slow, manual data preparation processes.

In this paper, we'll explore:

- Why traditional data processes struggle to meet the demands of AI
- How Promethium's Data Answer approach overcomes these challenges with
- The unique value Promethium brings by enabling instant, trusted insights across all enterprise data

By addressing these core challenges, organizations can finally make AI work - not just in theory, but in practice - transforming data into actionable insights that drive business impact.



The Challenge:

Why AI on Enterprise Data Is Hard

AI holds enormous promise for enterprises — the ability to answer complex questions instantly, predict outcomes with precision, and automate decision-making. But for most organizations, these outcomes remain frustratingly out of reach. The reason? Data simply isn't ready yet.

Despite investing in modern data platforms and AI tools, enterprises still struggle to extract meaningful insights on the fly. Data is fragmented, slow to access, and often lacks the context AI needs to deliver trusted results.

Three major challenges stand in the way:

1. Data Is Distributed and Fragmented: Enterprise data is scattered across cloud platforms, on-prem systems, and operational applications. Data warehouses, data lakes, CRM platforms, ERP systems - the list goes on and on and increases every year. Connecting this data traditionally requires extensive engineering work, including data pipelines, ETL processes, and manual data preparation.

This fragmentation creates two major obstacles:

- **Data Access Is Slow:** Teams must build and maintain complex data pipelines that take weeks or months to complete.
- **Data Is Often Incomplete:** With data living in silos, key information may be overlooked, limiting the accuracy of AI outputs.

2. Current Data Processes Aren't Built for AI: The traditional modern data stack was designed to support static reporting and batch processing but not necessarily real-time, AI-driven insights. These processes rely heavily on predefined data pipelines and pre-built datasets that take significant time and effort to construct. While this approach is great for operational processes, delivering highly performant and highly reliable results, it is less ideal in an AI world.

As AI expands the art of the possible, users are encouraged to ask questions that have not been asked before and are trying to quickly experiment with data to surface additional business value. But what happens if your current dashboard is unable to answer these questions? It requires building new pipelines and new dashboards, processes that currently take weeks if not months in most enterprises. But by the time the business gets the answers they need, the opportunity might have already passed.


3. AI Lacks Enterprise Context: Even when AI has access to enterprise data, it often lacks the understanding necessary to produce accurate, meaningful results. Without the right context, AI models risk producing incomplete or misleading outputs.

For example:

- Different teams may define metrics like "revenue" or "customer churn" in different ways. Without that context, AI may apply the wrong logic to generate insights
- Technical metadata such as schemas, labels, and relationships often varies across platforms, further complicating data integration and reducing AI's accuracy

These challenges are why so many organizations find themselves investing heavily in AI initiatives - only to see slow adoption, limited impact, and mounting frustration.

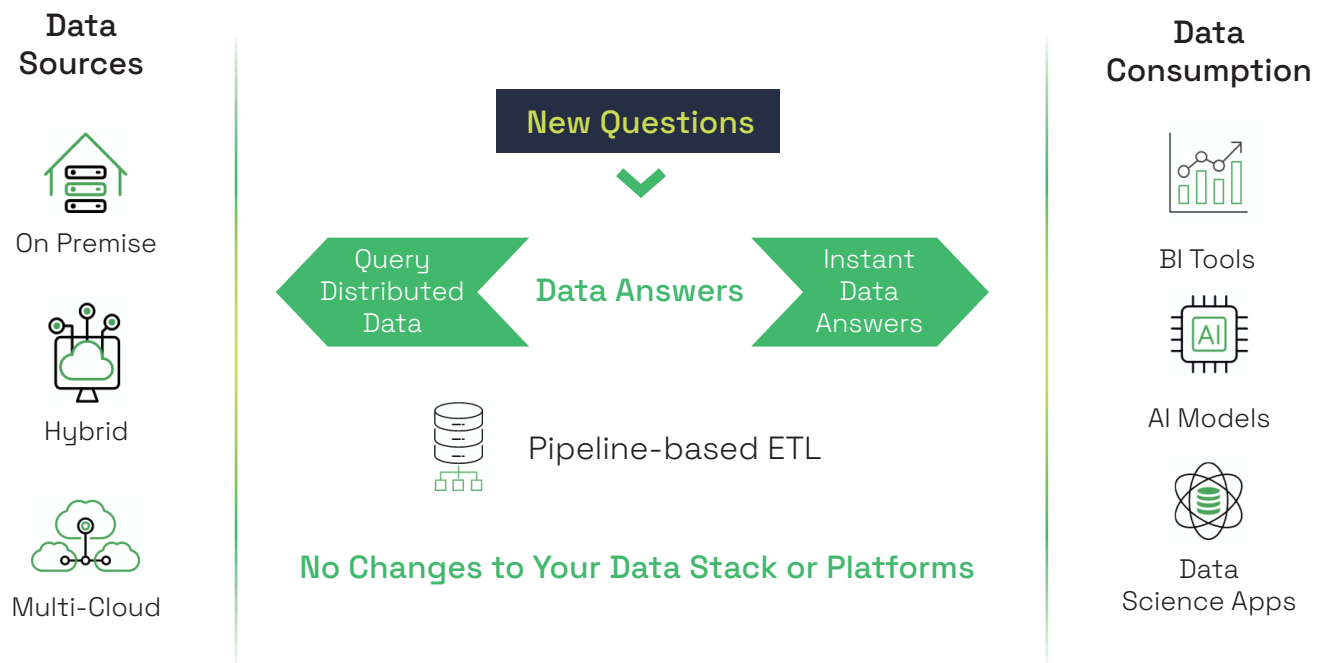
To make AI successful, organizations need a way to deliver trusted insights faster - without relying on always having to build new data pipelines or sacrificing data quality and governance.



Delivering Trusted,

Instant Data Answers With Promethium

To make AI successful in the enterprise, organizations need fast, trusted insights - yet fragmented data, slow processes, and incomplete context make this difficult. Promethium offers a new approach to get immediate insights for all your data no matter where it sits, and without having to build pipelines and dashboards up front.



Data Answers sit at the intersection of business and data & analytics teams. Instead of building new pipelines from scratch, data and analytics teams can route ad hoc and exploratory questions from the business to Promethium. Based on the question, Promethium queries distributed data across multiple platforms to dynamically build a Data Answer, while augmenting and creating the right business context to ensure that the Data Answers are relevant and accurate.

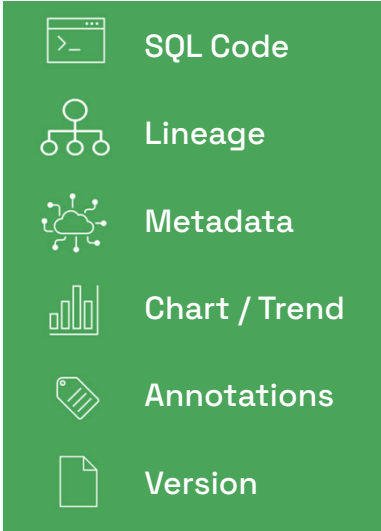
By design, Promethium does not require any changes to your data environment or the modern data stack. Data Answers are not a replacement for traditional data pipelines - they are a powerful complement. They provide a new way to enable fast, iterative analysis that would otherwise take weeks or months to do; and if results should be published or shared, they can be materialized into your existing data tools to become part of your production pipelines. It is a seamless addition and an accelerator to your current processes.

What Is a Data Answer?

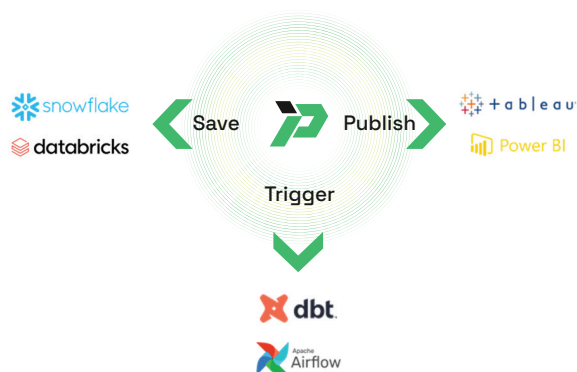
A **Data Answer** is an automatically generated response to a business question, derived from structured data across distributed enterprise sources. Unlike static reports or pre-built dashboards, Data Answers are created on demand with full context and reasoning, making them:

- **Accurate and Trusted:** Each Data Answer includes rich metadata, lineage, and annotations to ensure results are explainable and reliable.
- **Instantly Consumable:** Data Answers combine a SQL query (the dataset), relevant metadata (lineage, annotations), and visual outputs (charts or trends) in one complete deliverable.
- **Effortless to Generate:** Users can simply ask a question in natural language, and receive a structured response — with no need for manual pipeline development or extensive data engineering. They can be endorsed or downvoted, and that feedback is used to refine answers in the future.

Data Answer

- 
- SQL Code
 - Lineage
 - Metadata
 - Chart / Trend
 - Annotations
 - Version

Data answers can be directly integrated into the existing workflows of the data teams.



- They can be materialized into data platforms like Snowflake or Databricks (or any other platform of choice) for further analysis and consumption.
- Data engineers can use Data Answers as the starting point to build new pipelines to operationalize a hardened data product. For example, Data Answers can be directly saved as dbt models.
- Data answers just like any other dataset can be published into the BI tools of choice (e.g., PowerBI, Tableau, ThoughtSpot, etc.) or be accessed via APIs from agents or AI models.

By delivering fast, trusted insights in seconds, Data Answers give organizations the agility they need to unlock AI's true potential — without waiting months to build new pipelines or manually prepare data.

What Makes Data Answers Unique?

Data Answers are unique because they combine:


- **Speed:** Data Answers are generated instantly, turning questions into insights in seconds rather than weeks.
- **Accuracy:** Embedded metadata, lineage, and annotations ensure answers are trusted and explainable.
- **Ease of Use:** By leveraging natural language queries, Data Answers simplify data access for all users while still delivering the structure data teams require.

Everything within a Data Answer is purpose-built to deliver trusted insights quickly, helping data teams get business teams the answers they need without delays.

What AI Requires: The “Post-Modern” Data Stack

To move beyond the limitations of the traditional data stack and truly unlock AI’s potential, enterprises need more than pipelines and dashboards. They need a new foundation—one designed for agility, intelligence, and scale. This “post-modern” data stack is built on three core capabilities:

- **Access:** Instant access across all your distributed data platforms—cloud, on-prem, SaaS, and hybrid.
- **Context:** Built-in understanding of meaning, usage, and lineage—so every answer is accurate, explainable, and aligned with how your business actually works.
- **Self-Service:** Consumable data that can be used by anyone or any agent—delivered in a form that accelerates insights, not builds additional technical debt.

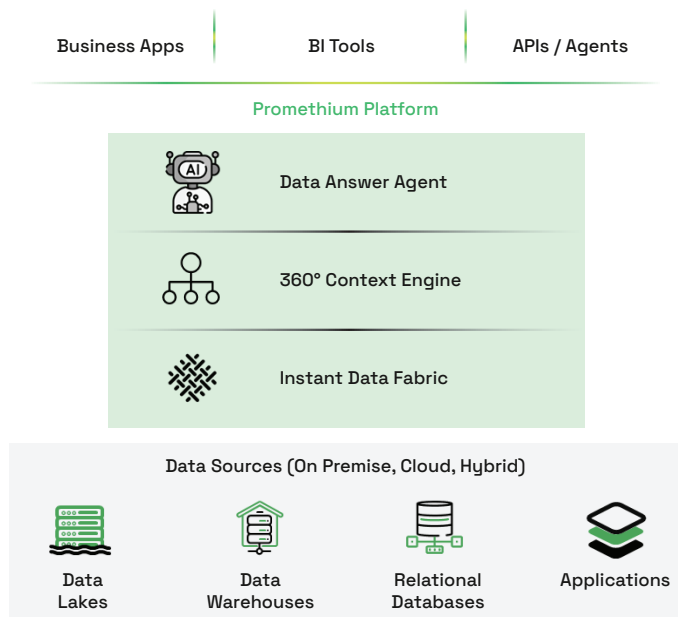


An Agentic Architecture Purpose-Built for Trusted, Instant Data Answers

Most tools deliver one or two of these pillars. Promethium is purpose-built to deliver all three in a unified, agentic platform - enabling organizations to scale AI with speed and trust. To do that, Promethium combines three distinct capabilities in a tightly integrated platform. Every element within the architecture is designed to fulfill this goal, ensuring fast, reliable, and explainable results in real time:

- **Instant Data Fabric** that provides real-time access to all the data platforms and data sources. The data fabric can discover enterprise data that is of value and query distributed data relevant to the analytics question, without requiring movement or duplication.
- **360° Context Engine** brings together business and technical context to ensure that the right data is selected for accuracy, enabling trusted results. Context is derived from various sources including your data platforms (e.g., technical metadata, query history), your business semantics (e.g., rules, dashboards, personas), and third-party sources (e.g., data catalogs).
- **Data Answer Agent** that orchestrates the building of an answer and enables data teams to iteratively test, validate, and share a data answer with business teams who can reinforce and endorse them.

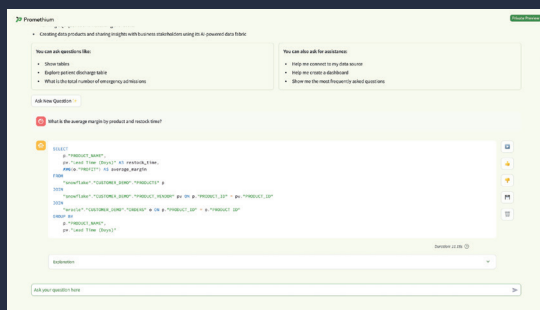
These three layers work in unison to dynamically assemble, explain, and deliver actionable insights - no matter how complex or fragmented the underlying data environment may be.



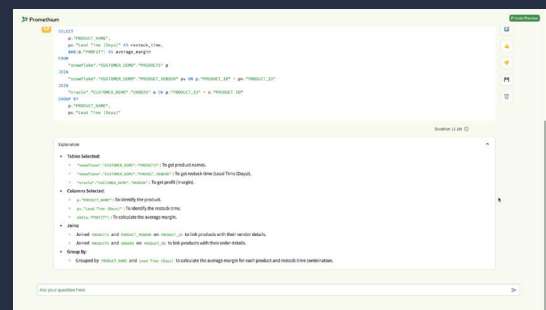
Bringing It All Together:

An Example Workflow

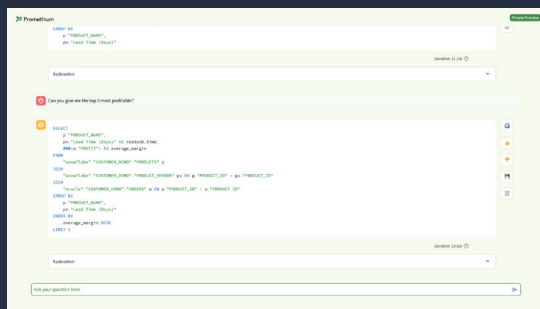
Promethium's architecture is designed for flexibility and iteration, allowing users to refine their questions without waiting weeks for new data pipelines. A typical Promethium workflow follows this path:



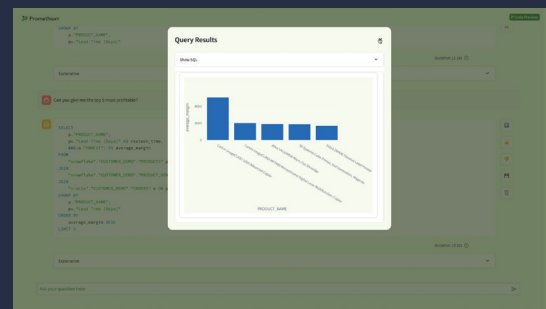
1. Ask a Question: A business user asks a question in natural language (e.g., "What is the average margin by product by restock time?")



2. See Results: Promethium generates a Data Answer with directly executable SQL query that delivers relevant insights and is visualized through charts, trends, and tabular data



3. Refine the Answer: If the result needs adjustment, the user can ask a follow-up question (e.g., "Which 5 products were the most profitable")



4. See Refined Results: The Data Answer updates in seconds, reflecting the refined query

5. Save the Answer: Once satisfied, the user can save the Data Answer for future reference, integrate it into a BI tool, or materialize it in a data platform like Snowflake or Databricks.

The Value Promethium

Brings to Organizations

Promethium is transforming how organizations access, analyze, and act on data by empowering both data teams and business users with faster insights, improved efficiency, and reduced complexity. As a **Gartner Cool Vendor**, Promethium has been recognized for its innovative approach to unlocking real-time insights - helping companies accelerate their AI journey, modernize their data infrastructure, and enable true self-service data.



Faster Insights for Better Decisions

Promethium dramatically reduces the time required to access and analyze data, enabling organizations to make informed decisions faster. By eliminating the need for extensive data preparation and pipeline building, teams can shift their focus from data wrangling to decision-making.

- Promethium empowers teams to **ask and answer questions instantly**, eliminating delays caused by traditional data preparation processes.
- Organizations can accelerate AI adoption without overhauling their data infrastructure, making it easier to integrate AI initiatives across the business.

Impact:

90%+
reduction

in time spent
by data teams
responding to
ad hoc business
questions.



Supercharging Data Teams

Promethium empowers data teams to handle more requests, faster - allowing them to focus on strategic initiatives instead of spending time on routine provisioning:

- Promethium enables data teams to rapidly answer unexpected or urgent questions without building complex pipelines.
- By streamlining workflows, Promethium helps data teams accomplish more with the same resources.

Impact:

5x
increase

in productivity for data analysts and data engineers.



Unlocking New Business Value

Promethium enables organizations to uncover insights that would otherwise remain hidden. By delivering trusted Data Answers in real-time, Promethium bridges the gap between technical data teams and business decision-makers:

- Promethium enables companies to identify trends, patterns, and insights that were previously inaccessible due to data silos or complex processes.
- These insights drive smarter business decisions and unlock new opportunities for innovation.

Impact:

10+
million of \$

in value generated by uncovering AI insights with new Data Answers.

Use Case Examples:

Promethium in Action



1

Empowering Non-Technical Users with True Self-Service Analytics

Business users can independently explore data in an intuitive and interactive way, gaining insights without requiring SQL knowledge or IT intervention and without being constrained to existing dashboards and datasets.



2

Enabling Data Teams to Handle Ad Hoc Questions Efficiently

Promethium allows data teams to respond to unexpected or urgent questions quickly, without the need for complex pipeline creation. As a result, teams can deliver more answers, faster.



3

Accelerating AI Adoption and Enabling Generative AI

By enabling real-time insights across distributed data sources, Promethium accelerates AI projects and expands access to Generative AI across teams, with your current level of data maturity.



4

Driving Strategic Impact by Delivering Hidden Insights

Promethium uncovers insights that would otherwise be missed, empowering both data teams and business leaders to make impactful decisions that were previously impossible due to inaccessible data.



5

Mitigating Data Migration Risks

Promethium's Instant Data Fabric acts as an abstraction layer, minimizing disruption during infrastructure changes. This ensures ongoing access to data regardless of whether systems are moved, upgraded, or replaced.

Delivering Value to Both

Data Teams and Business Users

Promethium's dual impact ensures that both technical and non-technical teams benefit:

- **For Data Teams:** Promethium streamlines data access and handling, improving productivity and freeing teams to focus on innovation rather than routine data tasks.
- **For Business Users:** Promethium removes barriers to data access, enabling decision-makers to explore insights independently, make faster decisions, and uncover previously hidden opportunities.

By combining these benefits, Promethium helps organizations become data-driven at scale, accelerating AI adoption, modernizing processes, and ensuring everyone in the organization can act on trusted, instant insights.

Promethium as the Connective

Tissue of Your Data Ecosystem

Promethium is designed to connect and enhance your existing data stack, acting as the connective tissue between data sources, BI tools, ETL platforms, and data catalogs.

By enabling organizations to query data where it lives, Promethium unlocks the full potential of existing investments - delivering fast, actionable insights without duplicating data or introducing new bottlenecks.

Why This Matters

and What Comes Next

The rapid acceleration of AI adoption is transforming industries, yet many enterprises risk falling behind because their data isn't AI-ready. While AI holds immense potential, fragmented data, slow processes, and a lack of trusted insights create major roadblocks.

Promethium addresses these challenges by delivering instant, trusted Data Answers enabling organizations to move faster, innovate sooner, and gain real-time insights in a governed, secure way.

By removing the bottlenecks that traditionally slow down data-driven decision-making, Promethium empowers organizations to:

- Respond faster to new questions, business needs, and opportunities.
- Equip business users with intuitive, self-service access to insights.
- Enable data teams to handle more requests efficiently and focus on high-impact projects.
- Accelerate AI adoption by ensuring enterprise data is accessible, trusted, and ready to power AI-driven innovation.

As companies race to harness the potential of AI, Promethium serves as the catalyst that accelerates this journey - connecting people to data faster, securely, and with full transparency.

Let's explore how Promethium can help you on your AI journey. Reach out to us today to learn more and schedule a demo.

About Promethium

Promethium is an agentic platform purpose-built for real-time AI insights, delivering Instant Data Answers directly from your distributed data ecosystem. Designed for enterprise data and analytics teams, Promethium enables organizations to go from business question to trusted answer in minutes instead of months. With dynamic agentic orchestration, a powerful context engine, and built on an Instant Data Fabric, Promethium eliminates existing bottlenecks to answer business questions and supercharges your data estate. Learn more at www.promethium.ai.

Promethium®, 2025. All rights reserved.

