



Talk to **ALL** Your Business Data

Powered by the AI Insights Fabric

Gartner

COOL
VENDOR
2024

Disclaimer

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved. GARTNER and the GARTNER COOL VENDOR badge are trademarks of Gartner, Inc., and/or its affiliates. All rights reserved. Gartner does not endorse any vendor, product or service depicted in its research publications and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's Research & Advisory organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

Gartner, Cool Vendors in Data Management: GenAI Disrupts Traditional Technologies, Nina Showell, Ehtisham Zaidi, Aaron Rosenbaum, 30 May 2024.



Today's Presenter

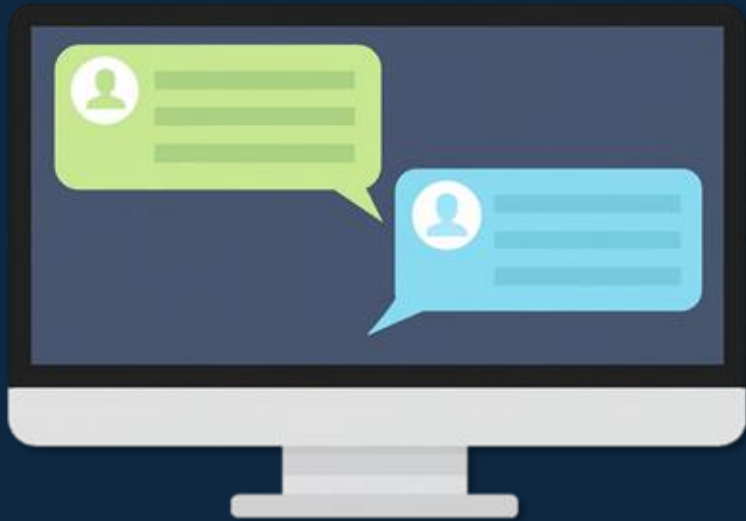


Prat Moghe
CEO, Promethium

Prat Moghe is a thought leader in data and analytics and serves as the CEO of Promethium.ai - the instant data fabric platform. He is a successful serial founder and CEO/GM with experience scaling multiple startups and public companies. Prior to Promethium, his roles include EVP at Cloudera, Founder/CEO at Cazena, and SVP at IBM Netezza. Prat has over 20 patents and holds a Ph.D. from UCLA in Electrical Engineering.



The AI Promise: Ask Questions, Get Answers



“Why did churn spike in EMEA last quarter?”

“Which products are driving the most margin in retail vs online?”

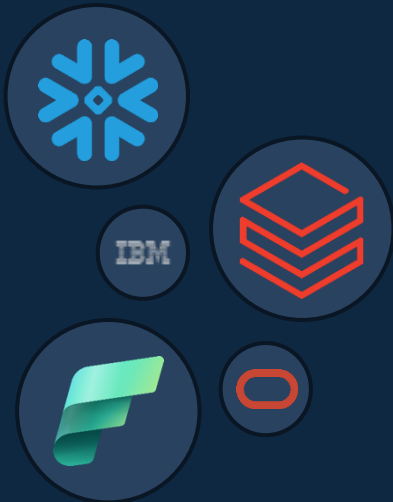
“What’s our pipeline coverage for Q4 across regions?”

The business expects agents to just work!



A Cambrian Explosion of “Talk to Your Data” Agents

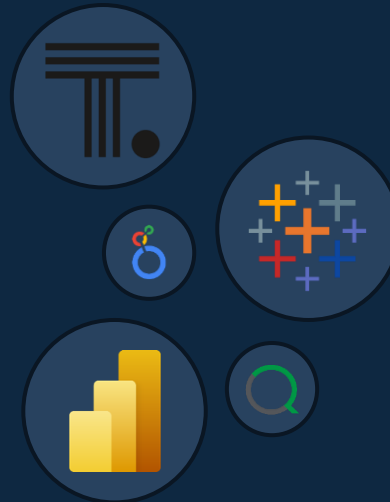
Data Platforms



Catalogs



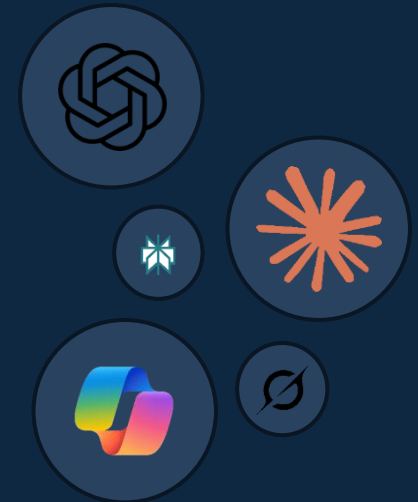
BI Tools



Applications



LLMs



In reality, each agent has access to only a subset of data and context!



The AI Disconnect: POC's vs Production

16%

Share of AI-generated answers to open ended questions in enterprises settings that are good enough for decision making.

95%

Share of enterprise AI pilots that fail to generate meaningful ROI.

60%

Share of AI Projects that will fail due to missing AI-ready data management practices.

70%

Share of AI projects where missing or bad data is identified as the root cause of failure.

Sources: BIRD Interactive Framework, MIT Project Nanda, Gartner, Nortal.



© 2026. All rights reserved.

The Enterprise Architecture Disconnect



Data is
Distributed

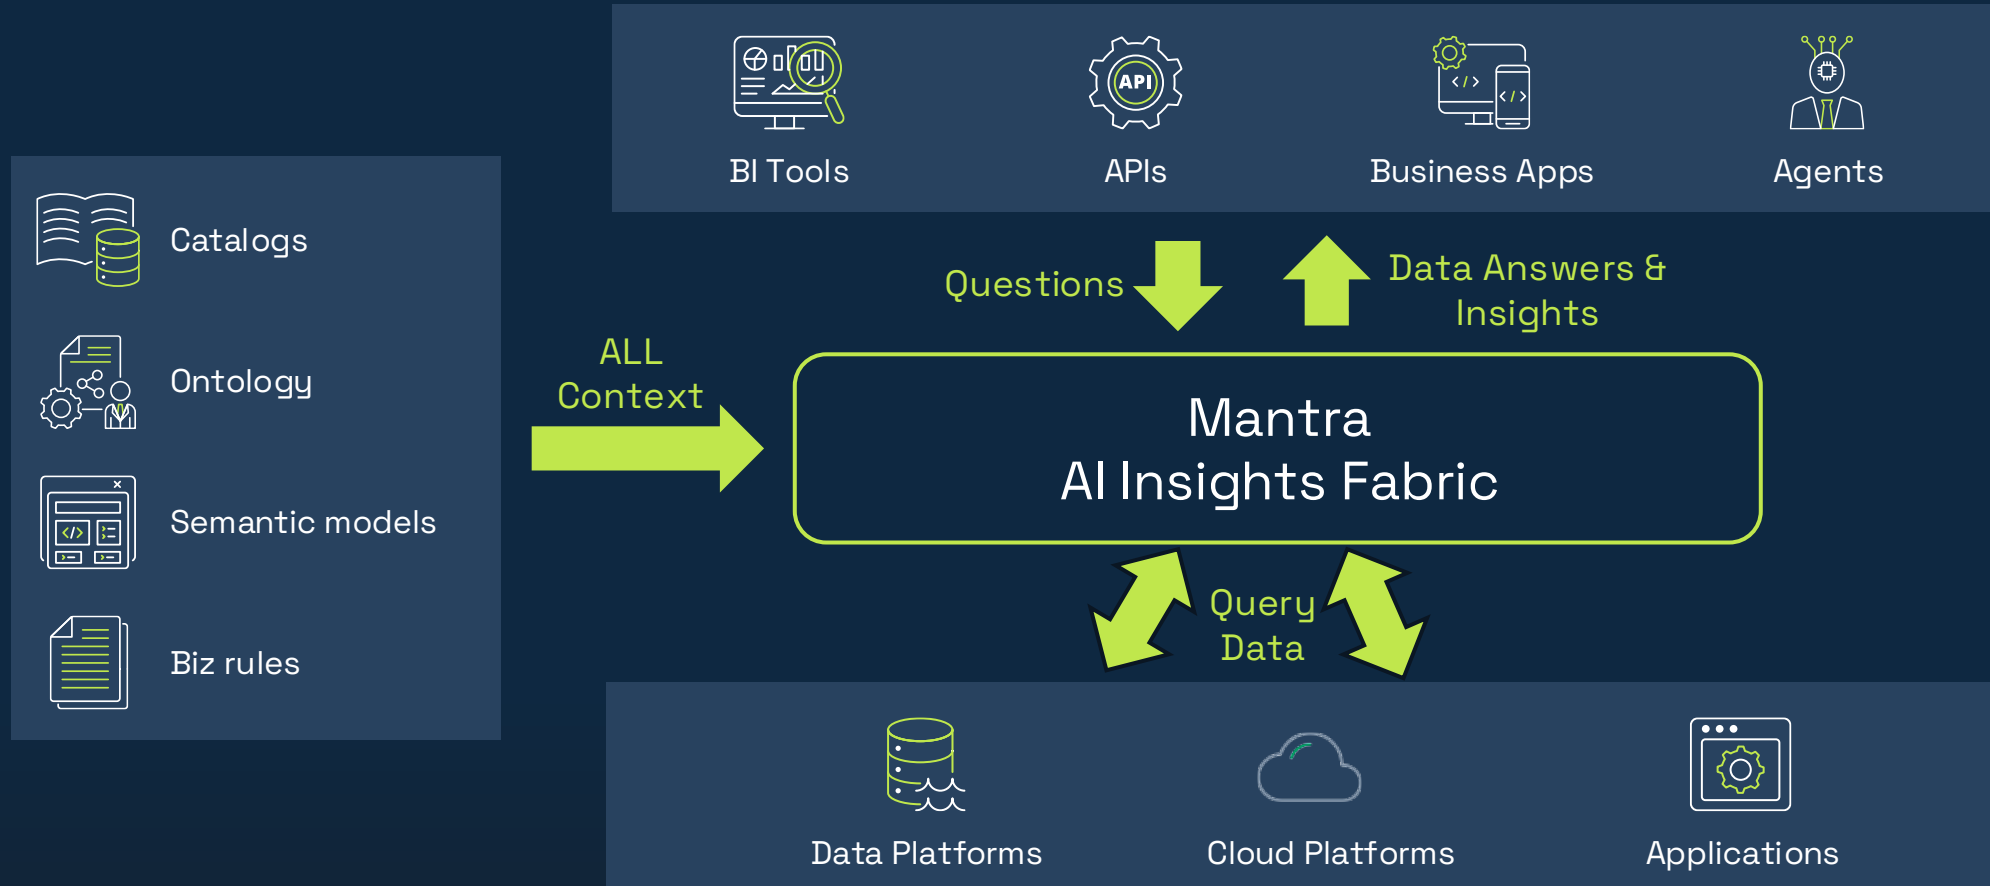


Context is
Fragmented

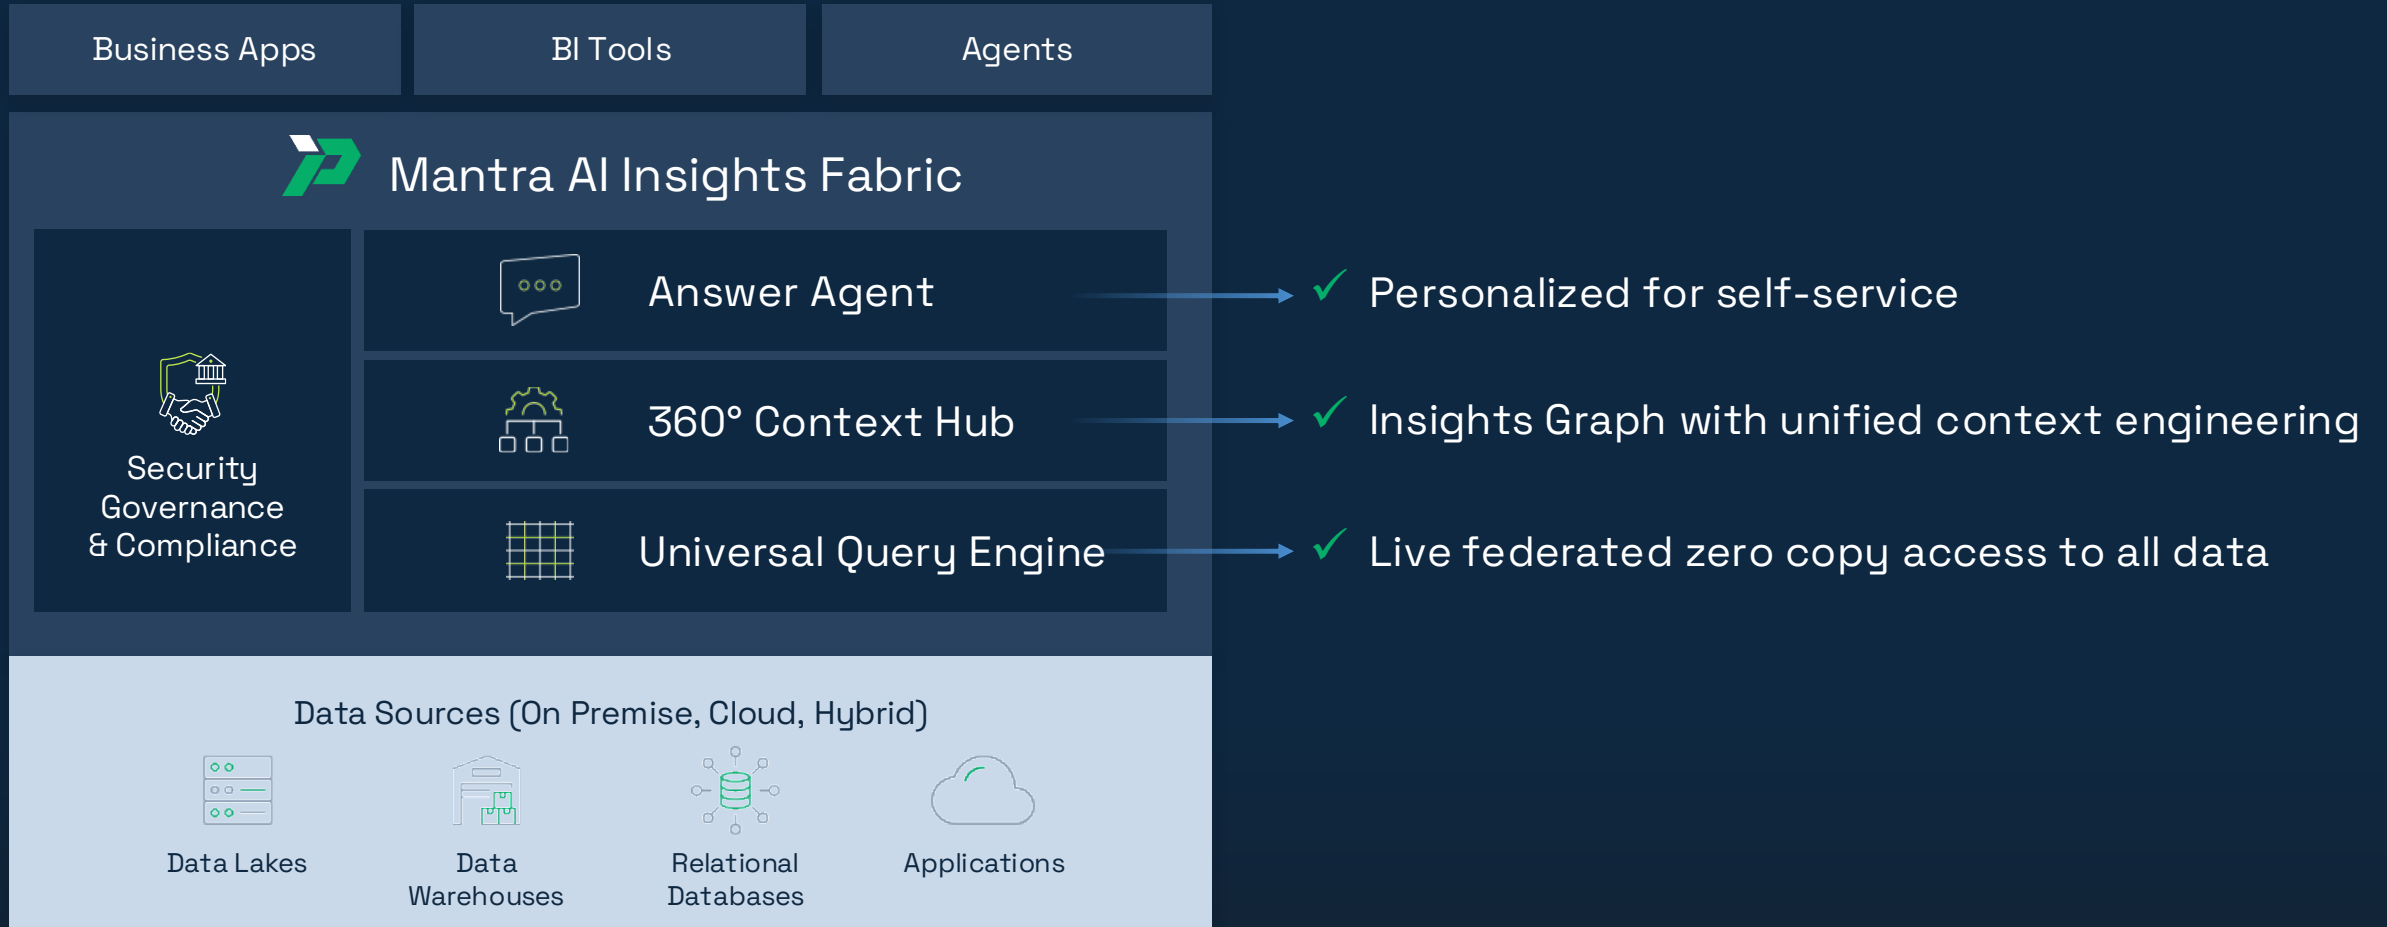


Cannot Scale
Personalization

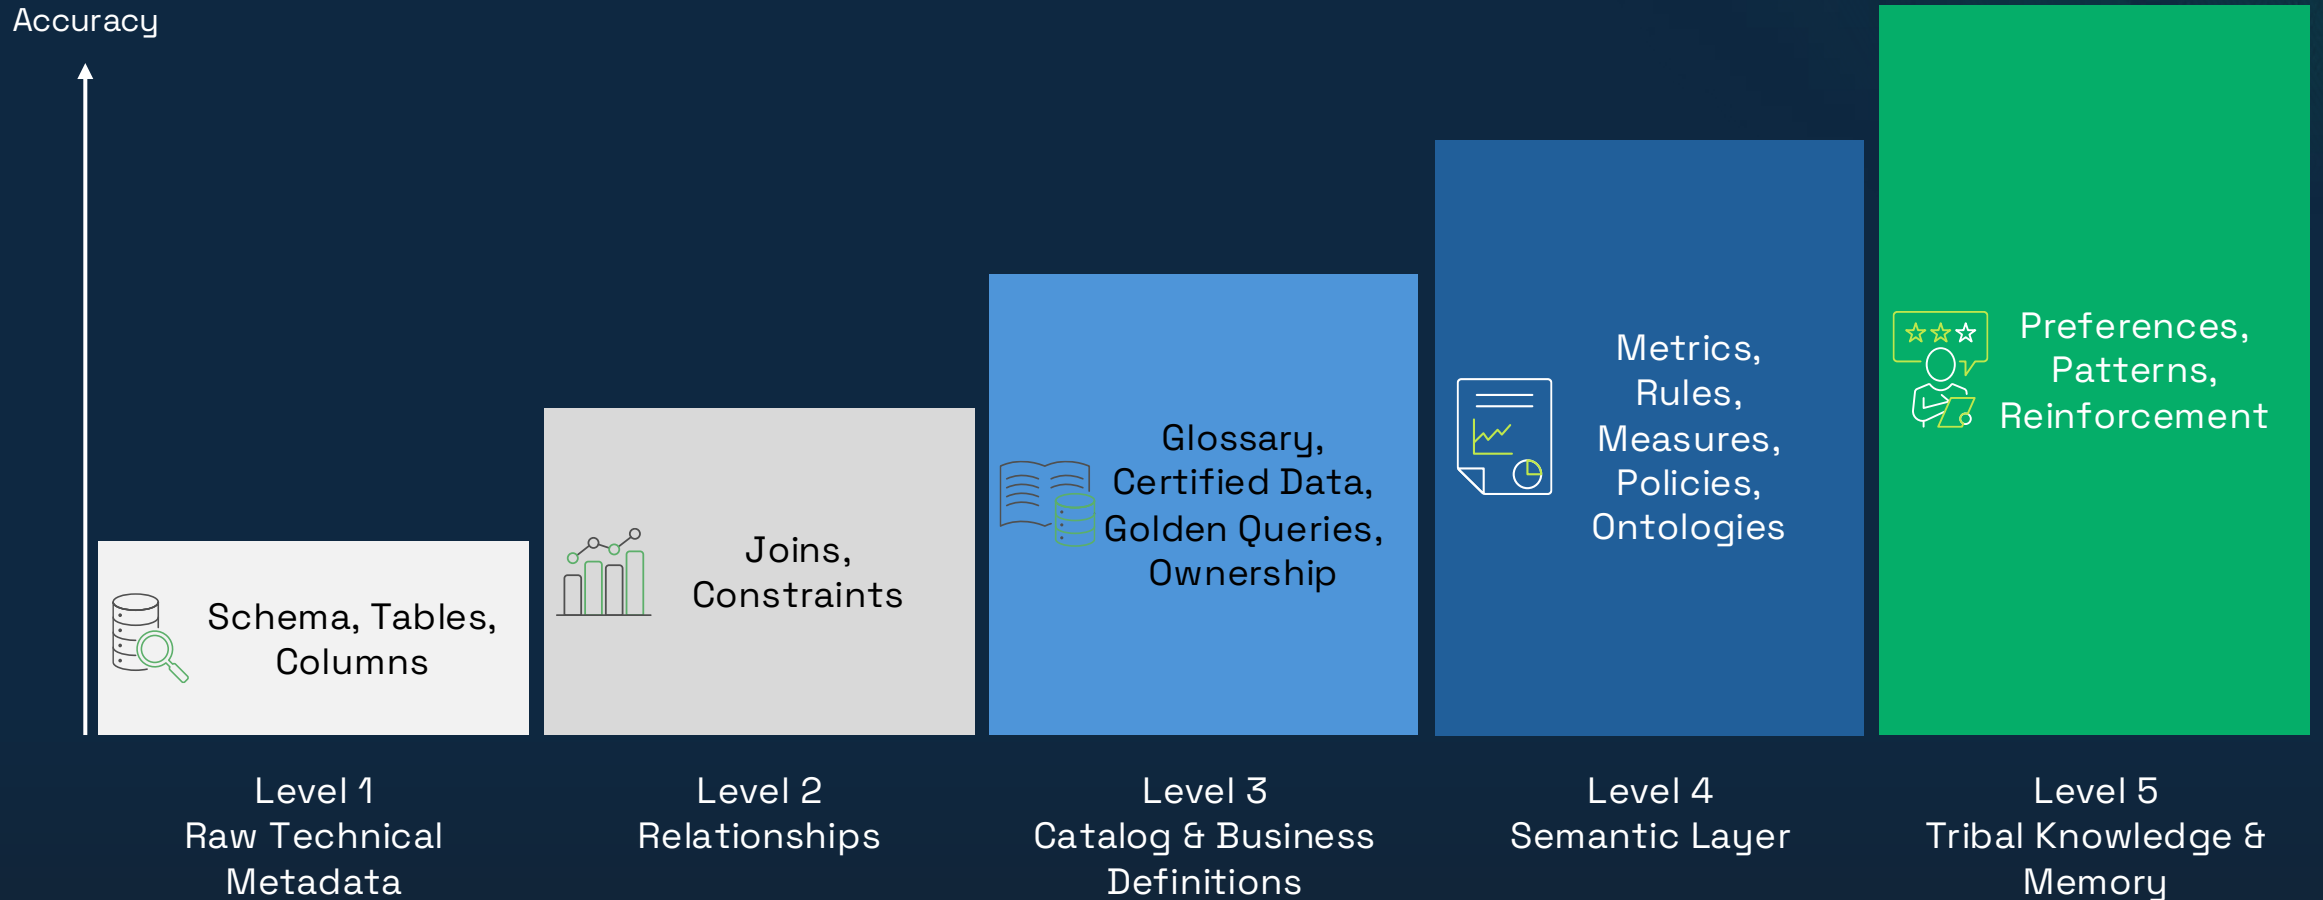
The AI Insights Fabric: Open Agentic Analytics



AI Insights Fabric Architecture

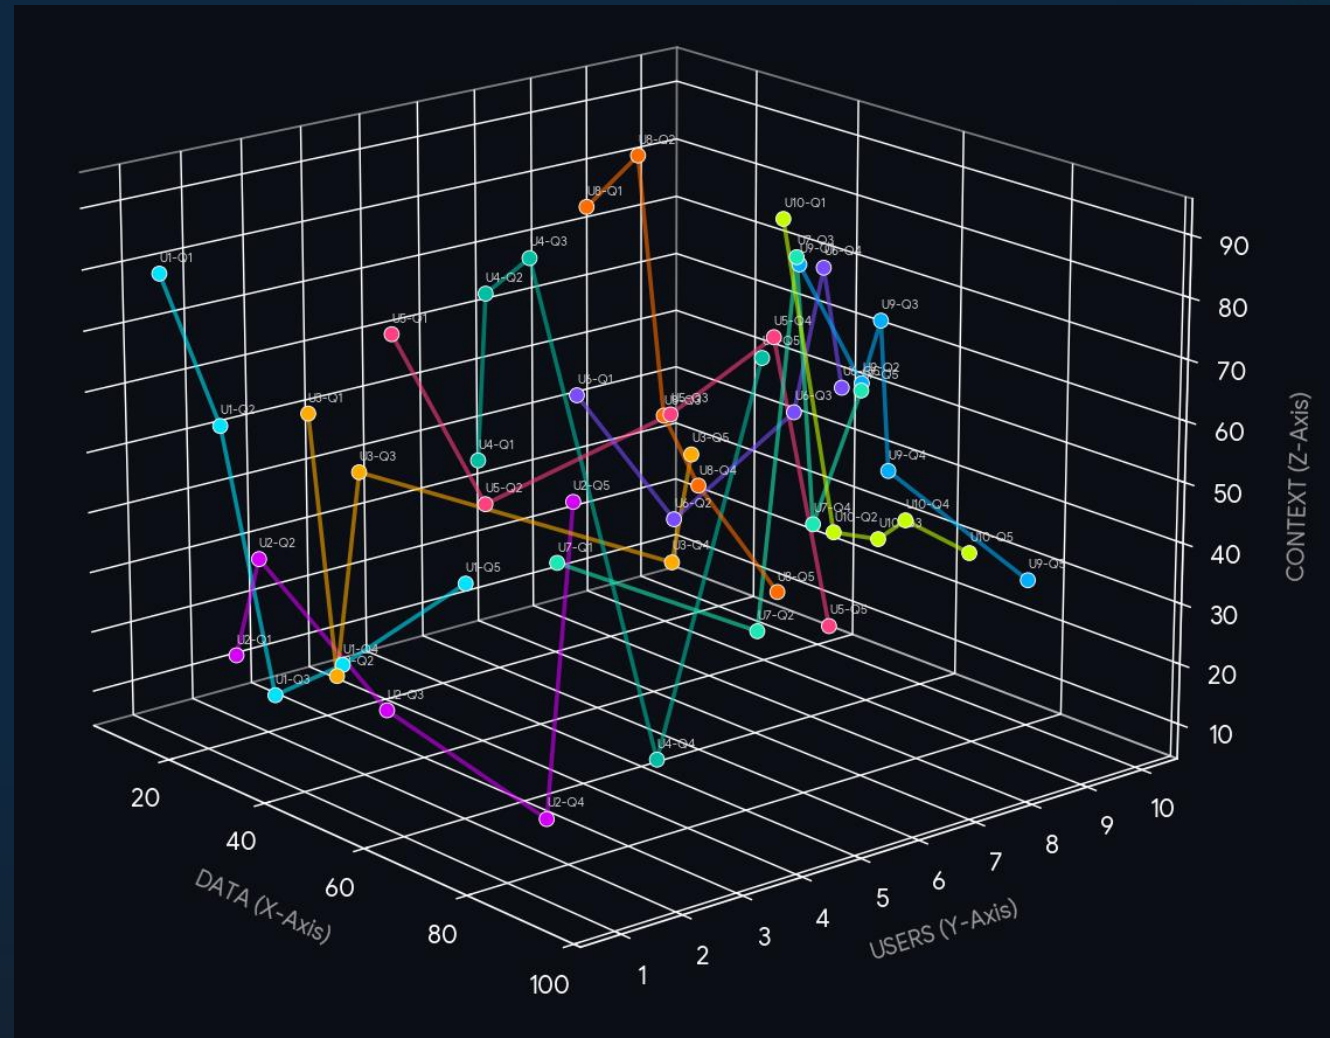


Multi-dimensional Context Graph for Insight Decisioning

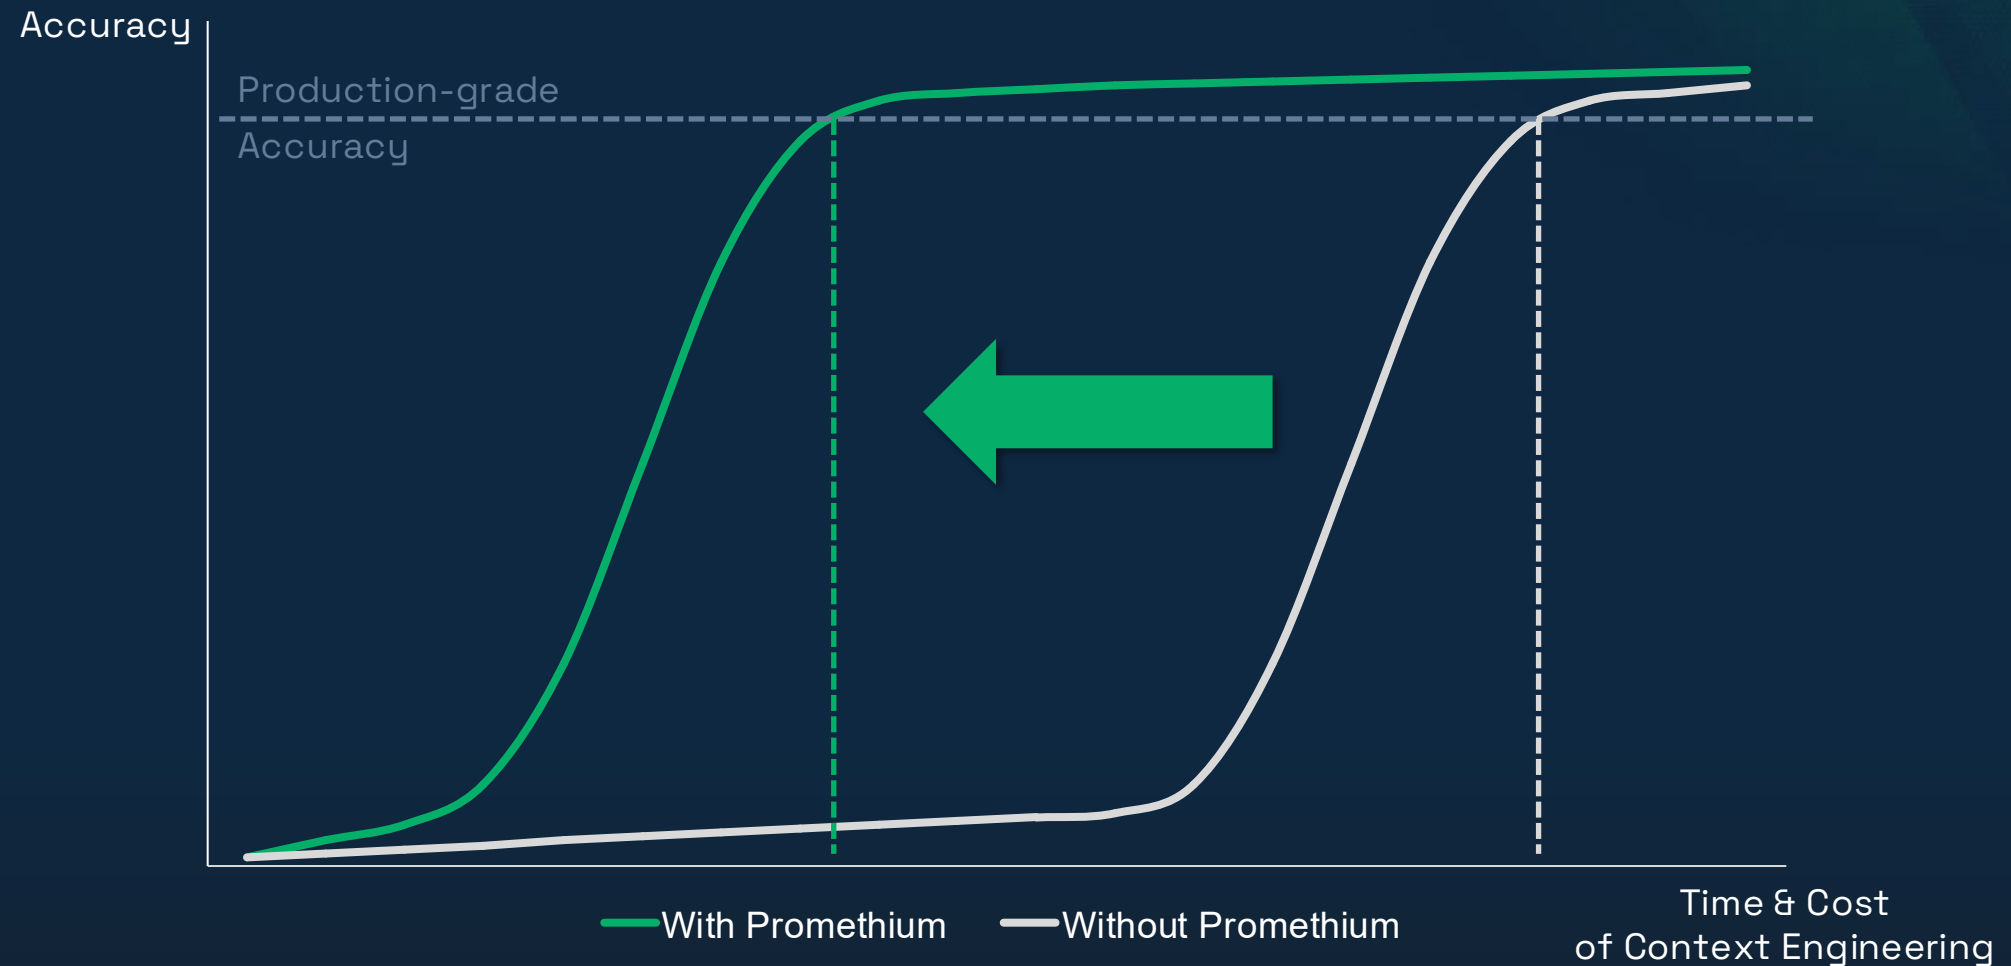


Mantra Insights Graph

User-Intent-Data-Context Driven



Mantra: Accurate Insights at Scale



Claude Asks Mantra to Generate Claims Report

ask mantra to compare the average claim amounts by coverage type and show the results in a table and chart

I'll start by loading the Mantra tool to make this query.

Searched available tools >

Now let me ask Mantra to compare average claim amounts by coverage type.

Ask mantra >

Good, Mantra has generated the SQL. Let me execute it to get the data.

Execute query >

I have the results. Let me now create a table and chart for you.

Average Claim Amounts by Coverage Type

Interactive artifact



Here are the results from Mantra. A few key takeaways:

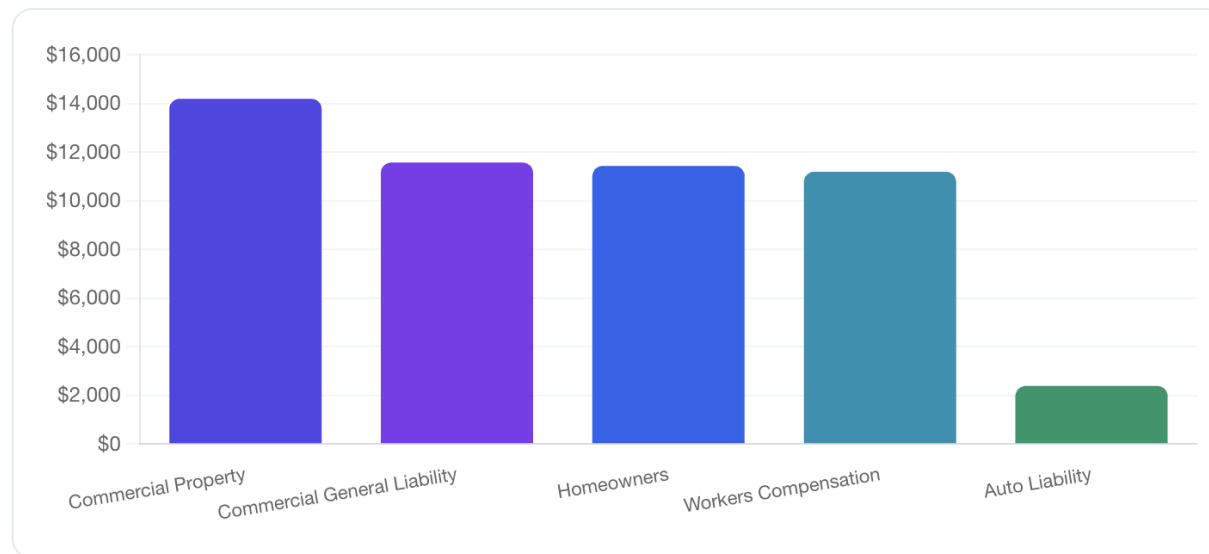
Commercial Property leads with the highest average claim amount at **\$14,200**, followed closely by **Commercial General Liability** (\$11,575), **Homeowners** (\$11,437.50), and **Workers Compensation** (\$11,200) — all clustered in a similar range.

Auto Liability stands out as significantly lower at **\$2,386.67**, roughly one-sixth of the top category. This suggests auto claims tend to be much smaller in severity compared to property and liability claims.

The chart and table above give you both a visual comparison and the exact figures. Would you like to drill deeper into any particular coverage type?

Average Claim Amounts by Coverage Type

Comparison of mean claim amounts across all coverage categories



Coverage Type

Avg Claim Amount

Distribution

Commercial Property

\$14,200.00

Commercial General Liability

\$11,575.00

Homeowners

\$11,437.50

Workers Compensation

\$11,200.00

Auto Liability

\$2,386.67

Drill down: Explainable and Validated Answers

Profitable Policies

Lineage

The screenshot displays a data analytics workspace with the following components:

- Query:** A SQL query to identify the most profitable insurance policies, including customer names and coverage types. The query is as follows:

```
1 SELECT
2   p.POLICY_ID,
3   COALESCE(c.COMPANY_NAME, CONCAT(c.FIRST_NAME, ' ', c.LAST_NAME)) AS CUSTOMER_NAME,
4   ct.COVERAGE_NAME,
5   MIN(f.combined_ratio) AS best_combined_ratio
6 FROM ins_sql.ins_fin.policy_monthly_financials f
7 JOIN sf_ins.INSURANCE_OPS.POLICIES p
8   ON f.policy_id = p.POLICY_ID
9 JOIN sf_ins.INSURANCE_OPS.CUSTOMERS c
10  ON p.CUSTOMER_ID = c.CUSTOMER_ID
11 JOIN sf_ins.INSURANCE_OPS.COVERAGE_TYPES ct
12  ON p.COVERAGE_TYPE_ID = ct.COVERAGE_TYPE_ID
13 WHERE f.combined_ratio < 1.0
14 GROUP BY p.POLICY_ID, c.COMPANY_NAME, c.FIRST_NAME, c.LAST_NAME, ct.COVERAGE_NAME
15 ORDER BY best_combined_ratio ASC
16 LIMIT 100
```
- Results:** A section titled "Data Answer" containing the question: "What are the most profitable policies? including customer name and coverage type".
- Lineage:** A diagram showing the data sources used in the query: sf_ins CUSTOMERS, sf_ins COVERAGE_TYPES, ins_sql policy_monthly_financials, and sf_ins POLICIES. Arrows indicate the flow of data from these sources into the query.
- Explanation:** A text box explaining the query: "I've created a query to identify the most profitable insurance policies, including customer names and coverage types. How it works: The query joins policy financials with policy details, customer information, and coverage types."

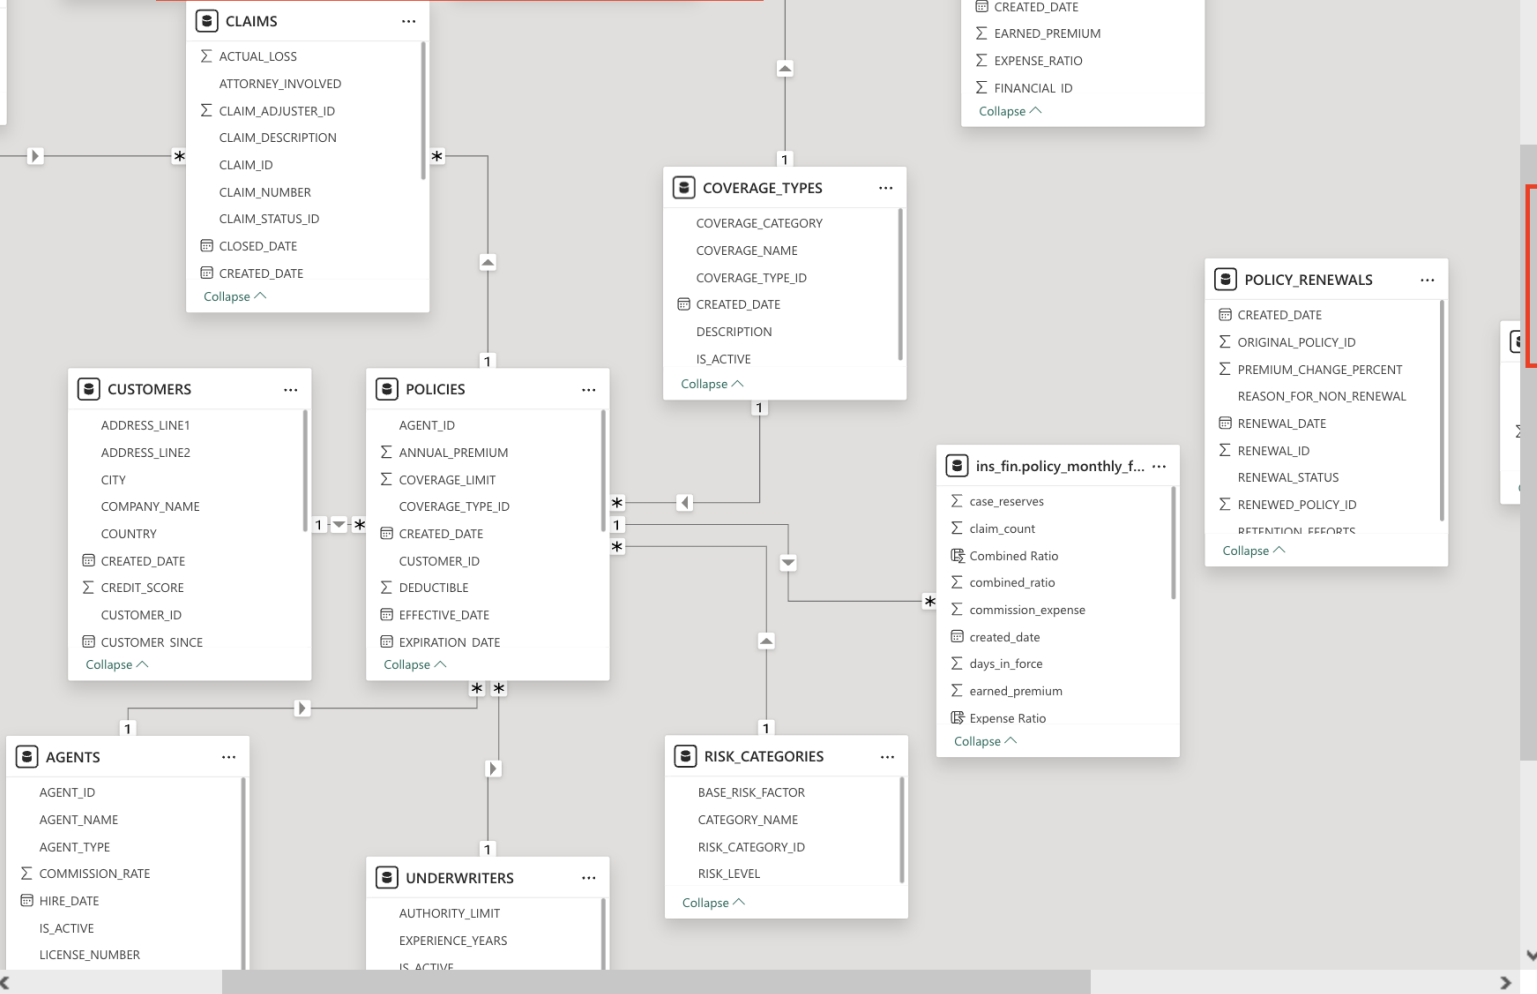


Business Rules & Metrics from PBI

File Home Help

Get data OneLake Transform Refresh data New measure New column New table Calculation group New parameter Manage roles Manage relationships Best practice analyzer Memory analyzer Community notebooks Model health

```
Profitable Policy = [Combined Ratio] < 1.0
```



Properties

General

Name: Profitable Policy

Description: Profitable policy is defined as combined_ratio less than 1.0 from policy_monthly_financials table

Display folder: Enter the display folder

Is hidden: No

Formatting

Data type: True/false

Format: True/false

Data

Tables Model

Search

- ins_fin.policy_monthly_financials
 - case_reserves
 - claim_count
 - Combined Ratio
 - commission_expense
 - created_date
 - days_in_force
 - earned_premium
 - Expense Ratio
 - expense_ratio
 - Loss Ratio
 - loss_ratio
 - paid_claims
 - policy_financial_id
 - policy_id
 - policy_status
 - Profitable Policy**
 - reporting_month
 - underwriting_expense
 - written_premium

Sign up for the “Talk to Data” Production Challenge!

Top 5 Business Questions
Week 1



Deploy Mantra
Week 2-3



Domain 2



Domain 3



**Quickly scale in
production across
additional domains**

Deliver Insights
Week 4



AI Insights Fabric Self-Service Use-cases



AI Analyst

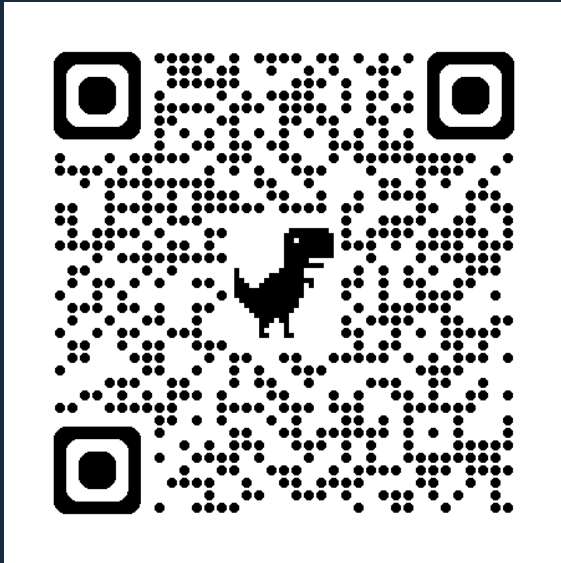
Mantra empowers your teams to instantly build and share data answers and insights



AI / BI + AI Ready Data

Plug your agents into the AI Insights Fabric to get access to data and context





Scan to schedule your
personalized demo

**Find us at
booth 101
to see live
demos & more**