Enabling Conversations with Data: From Engineering to Real-World Use

Time Series Highway / VAIA / 2026

2026 © delaware

delaware

Rémi Tossut

Consultant Data Engineering
+2.5 years @ delaware

MSc Business & Info. Sys. Engineering

Wouter Labeeuw

Senior Manager
+10 years @ delaware

PhD Engineering Science
MSc Computer Science option AI

Outline

1. **Situation**
   Wouter
2. **Case**
   Wouter
3. **Build**
   Rémi
4. **Consume**
   Rémi
5. **Conclusion**
   Wouter

Situation

intelligent agent

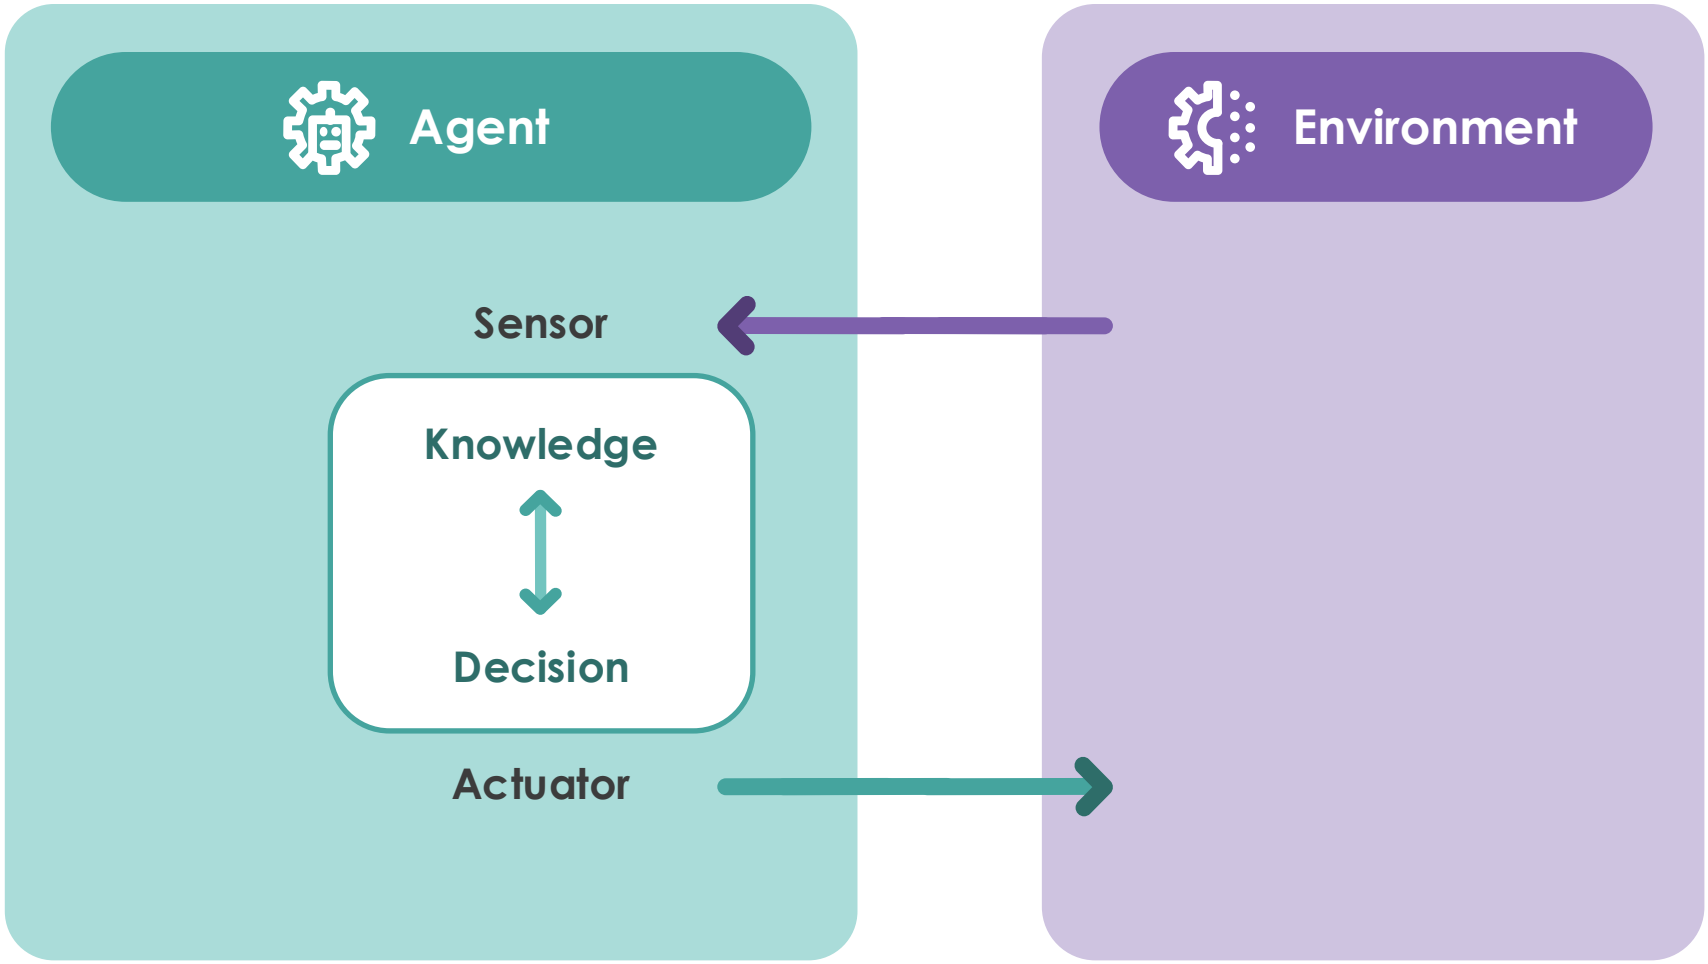

Minimal requirements

Sense
Reason
Act
Goal

Practical examples

Thermostat (Nest)
Spam Filters
Trading bots

Russell, S., & Norvig, P. (1995). Artificial Intelligence: A Modern Approach.

agents for data

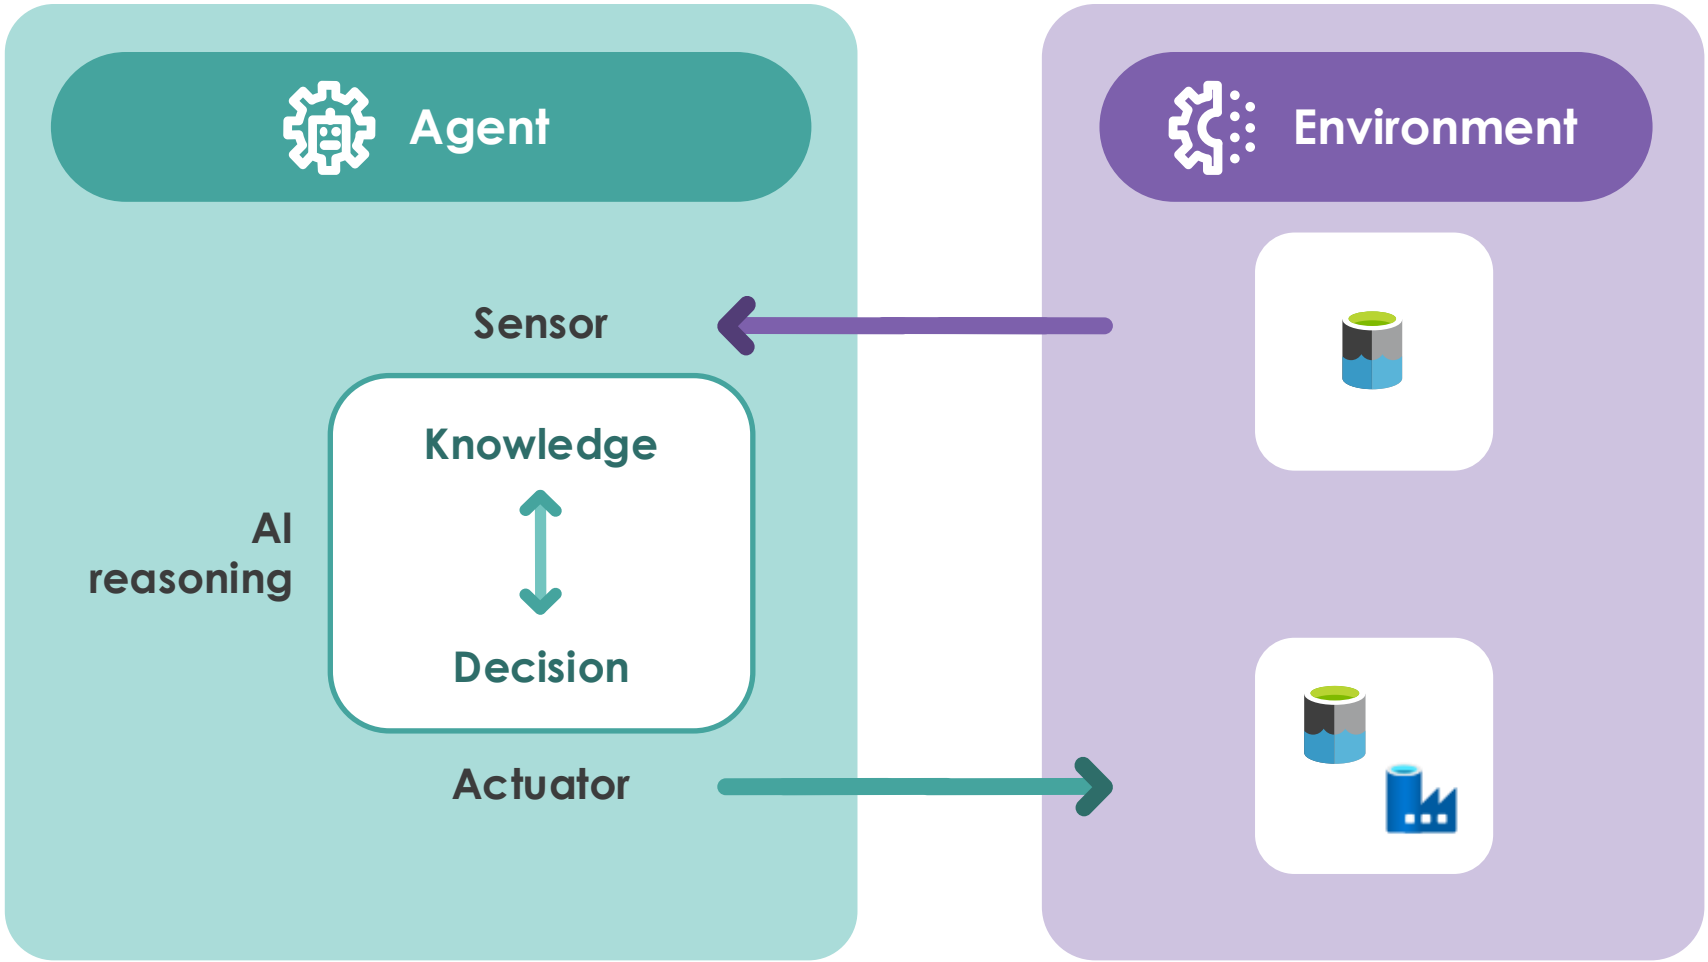

Sensing

Reading source data
Trigger

AI reasoning

Transformations
Detect anomalies

Act

Generate and run transformation

The **task-completion time horizon** is the task duration (measured by human expert completion time) at which an AI agent is predicted to succeed with a given level of reliability. For example, the 50%-time horizon is the duration at which an agent is predicted to succeed half the time. The graph below shows the 50%- and 80%-time horizons for frontier AI agents, calculated using their performance on over a hundred diverse software tasks.

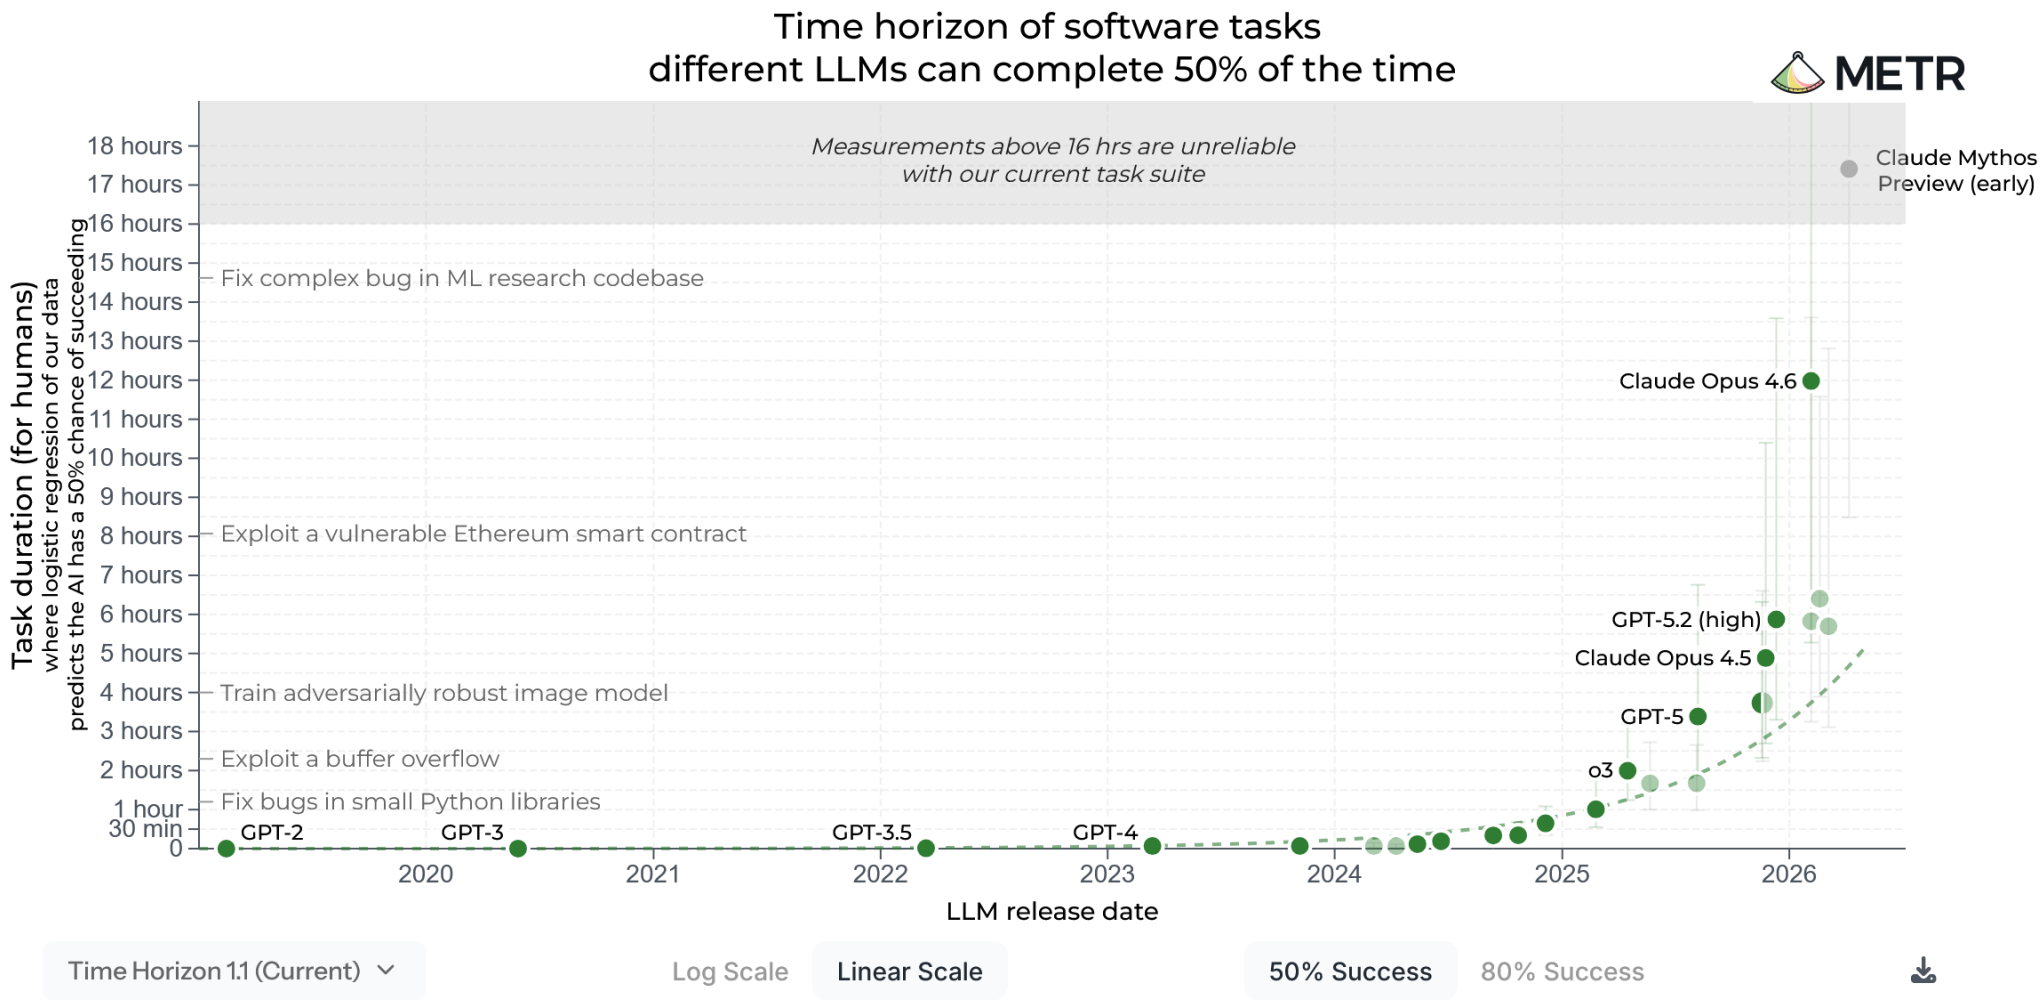

Analysis code is available on GitHub. Raw data available here

The **task-completion time horizon** is the task duration (measured by human expert completion time) at which an AI agent is predicted to succeed with a given level of reliability. For example, the 50%-time horizon is the duration at which an agent is predicted to succeed half the time. The graph below shows the 50%- and 80%-time horizons for frontier AI agents, calculated using their performance on over a hundred diverse software tasks.

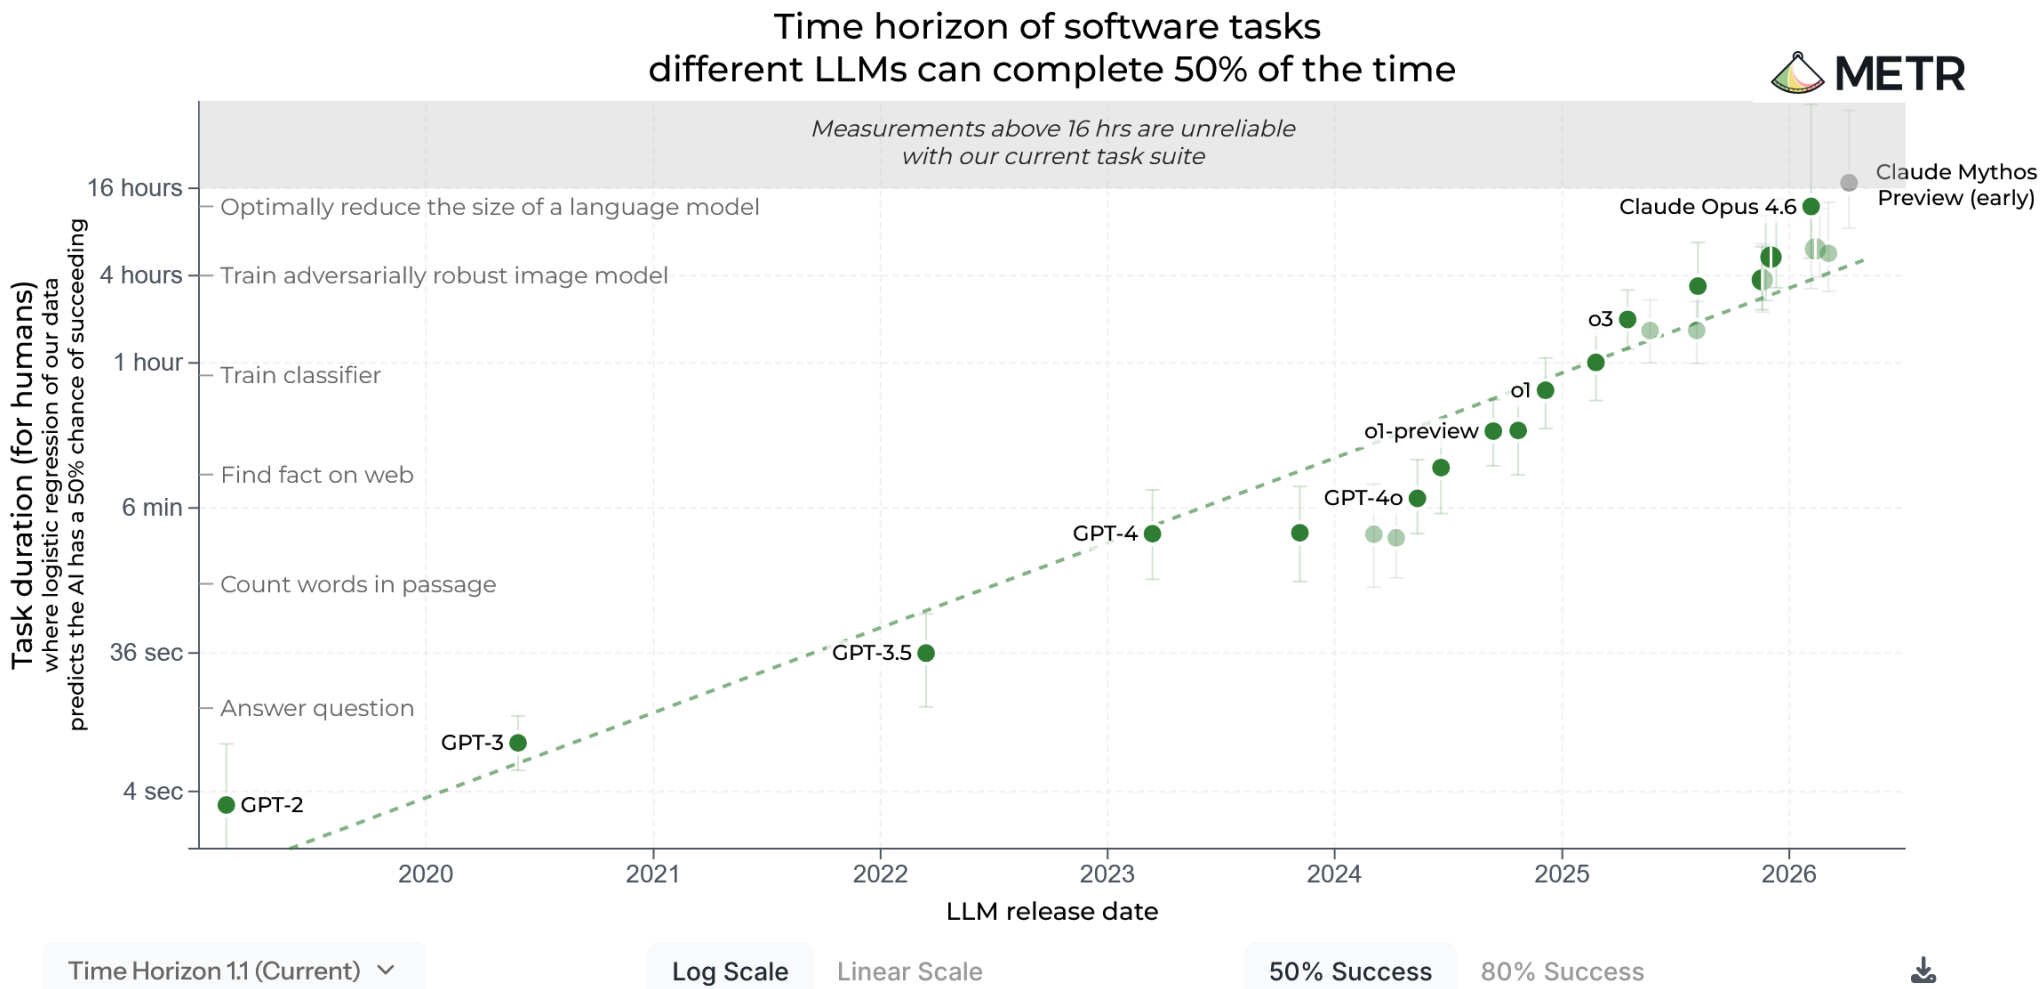

Analysis code is available on GitHub. Raw data available here

Commentary Breakingviews

AI labs' consulting push defies corporate history

Karen Kwok

May 11, 2026 4:50 PM GMT+2 · Updated May 12, 2026

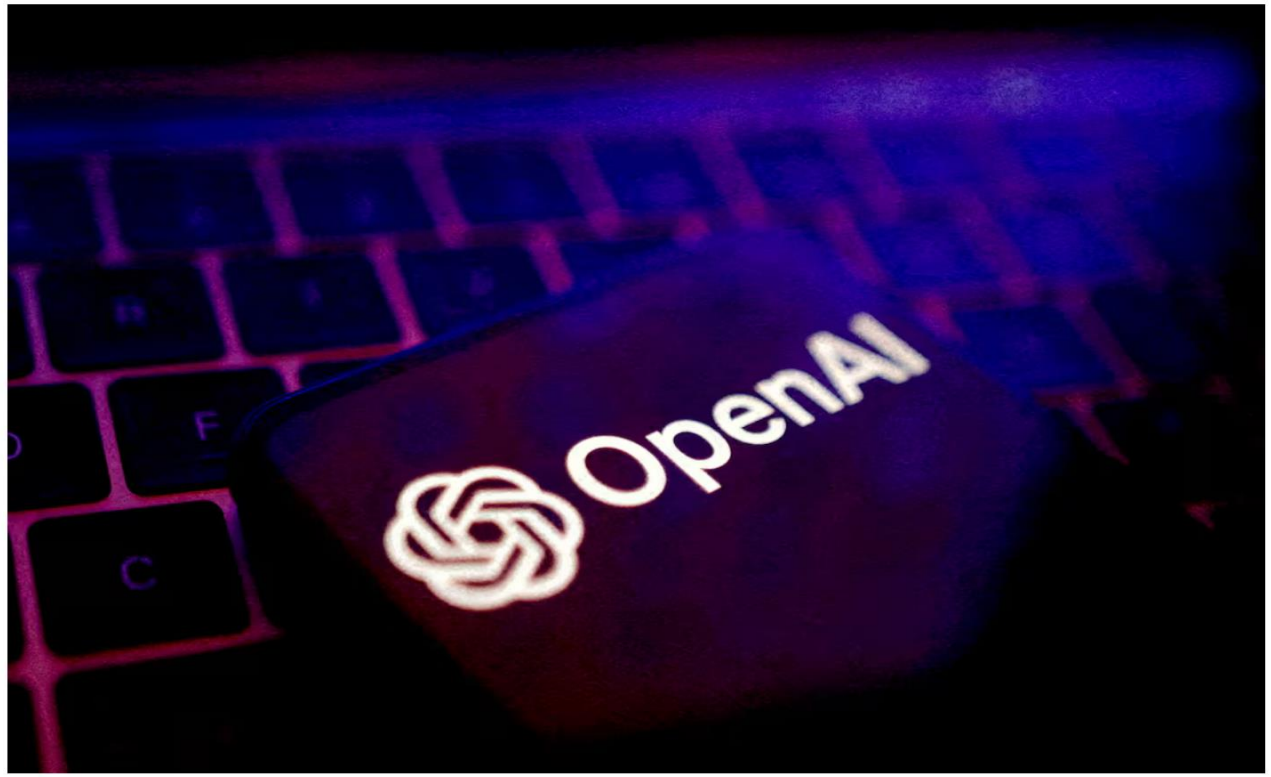

OpenAI logo is seen in this illustration taken May 20, 2024. REUTERS/Dado Ruvic/Illustration/File Photo Purchase Licensing Rights

AI Agents for Data

Build

Pipeline & Transformation

AI generates logic
AI generates pipelines
Human verifies Build

Operate

Quality & Observability

Data Quality
Monitoring
Optimisation

Consume

Analytics & Decisions

Natural language querying
Prescriptive insights
Ad-hoc questions

AI Agents for Data

Build

Pipeline & Transformation

AI generates logic
AI generates pipelines
Human verifies Build

Operate

Quality & Observability

Data Quality
Monitoring
Optimisation

Consume

Analytics & Decisions

Natural language querying
Prescriptive insights
Ad-hoc questions

Case

Detect anomalies and trends in IoT data

Any problems in the water network?

Input

Water Meter Readings

1 reading / device / day

-

Functional/non-functional

-

One or more per location

Build

Process Data for Anomalies

Device linking

-

Rolling statistics

-

Anomaly detection (reverse flow, z-score)

Consume

Interpret Anomalies

Seasonal patterns

-

Upward trends

-

Water Network Health

Build

The way we build data platforms is changing

writing code

developers spend most of their time writing, debugging, and refactoring data flows manually

conducting agents

like a conductor leading an orchestra, the developer directs agents and reviews their output

The way we build data platforms is changing

Traditional delivery

AI-driven delivery

	Traditional delivery	AI-driven delivery
Workflow	**Linear content creation** Design, code, test, fix, refactor, document	**Iterative prompting, spec-driven** Plan, instruct AI to build, review and refine
Human role	**Implementing & debugging** Deep familiarity with every line of code	**Architecting & directing** Focus on patterns and best-practices

Our team of data & analytics agents

BI analyst

Captures BI requirements through a guided Q&A conversation

Phases: Analysis Design

Data generator

Generates realistic sample data for a given target data model

Phases: Build Test

Power BI modeler

Builds semantic models, DAX measures & relationships in Power BI

Phases: Build Test

Databricks data engineer

Builds and optimizes data pipelines in Azure Databricks.

Phases: Build Test

MS Fabric data engineer

Builds and optimizes data pipelines in Microsoft Fabric.

Phases: Build Test

Power BI report creator

Builds interactive dashboards and reports in Power BI.

Phases: Build Test

Demo

1 Create Skill

- Define a reusable skill capturing best practices for the data platform.

Databricks · Genie

2 Build Data Layer

- Use Genie to explore data and shape the bronze → silver → gold layers.

Databricks · Genie

3 Enhance Semantic Layer

- Refine measures, relationships and metadata on the Power BI model.

Databricks · Genie

Claude + PBI MCP

4 Consume the Data

- Query the semantic model in natural language to get answers fast.

Claude + PBI MCP

Outcome: **an AI-ready data platform where business users get trusted answers in natural language.**

Demo 1 — Create Skill

1. Create Skill

- Define a reusable skill in Claude Code capturing best practices for the data platform.

Databricks · Genie

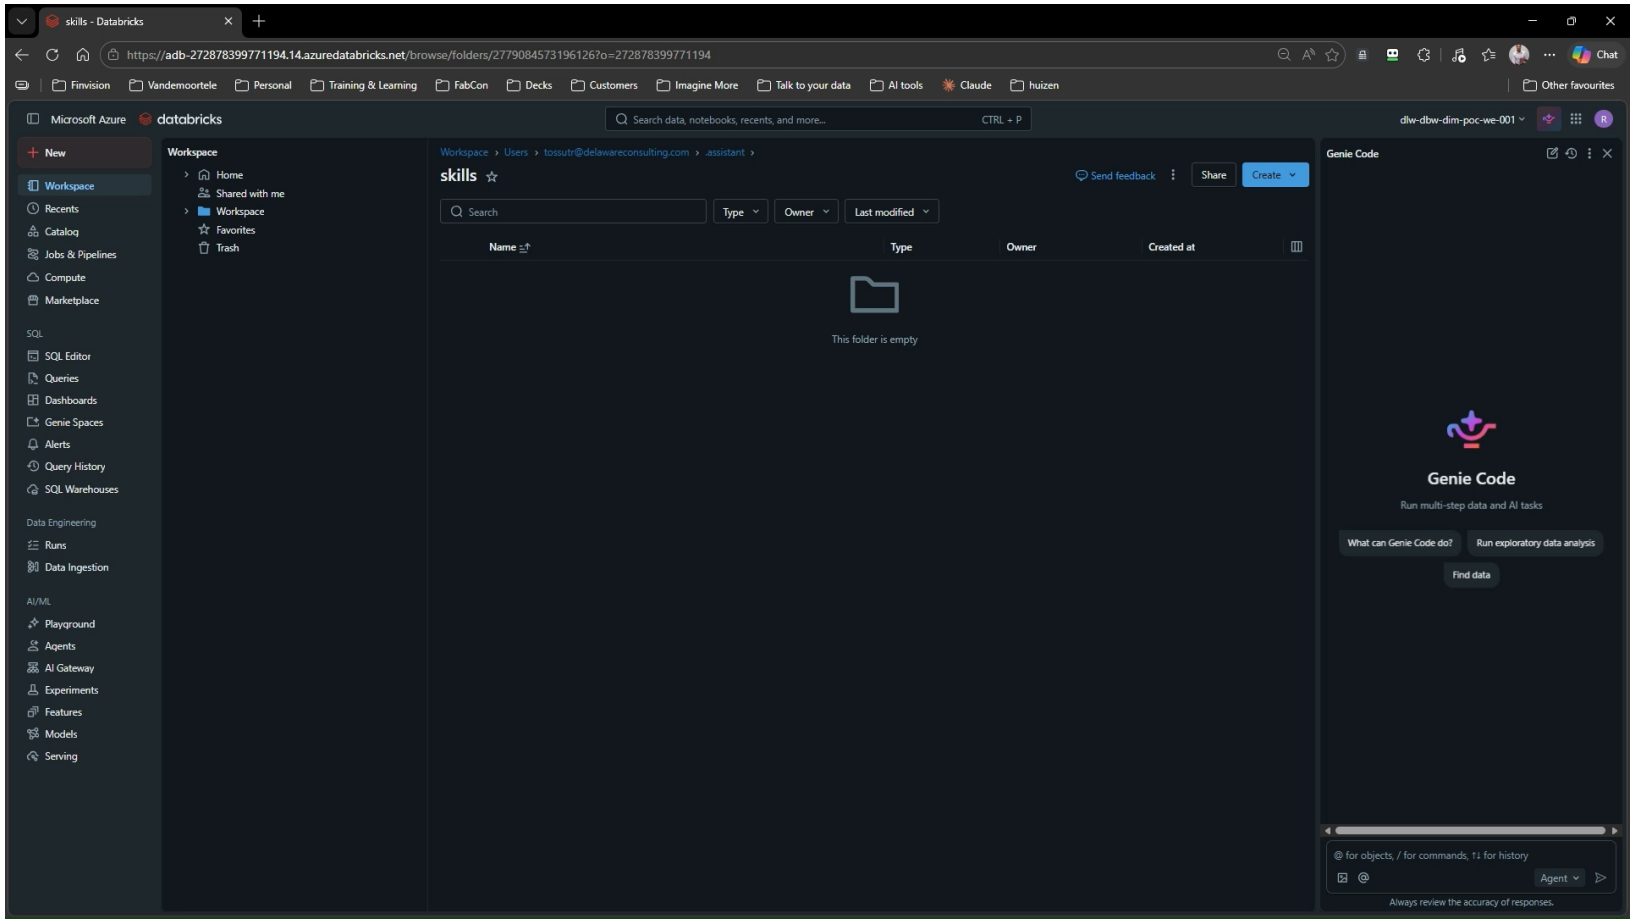

Demo 2 — Build Data Layer

2. Build Data Layer

- Use Genie to explore data and shape the bronze → silver → gold layers.

Databricks · Genie

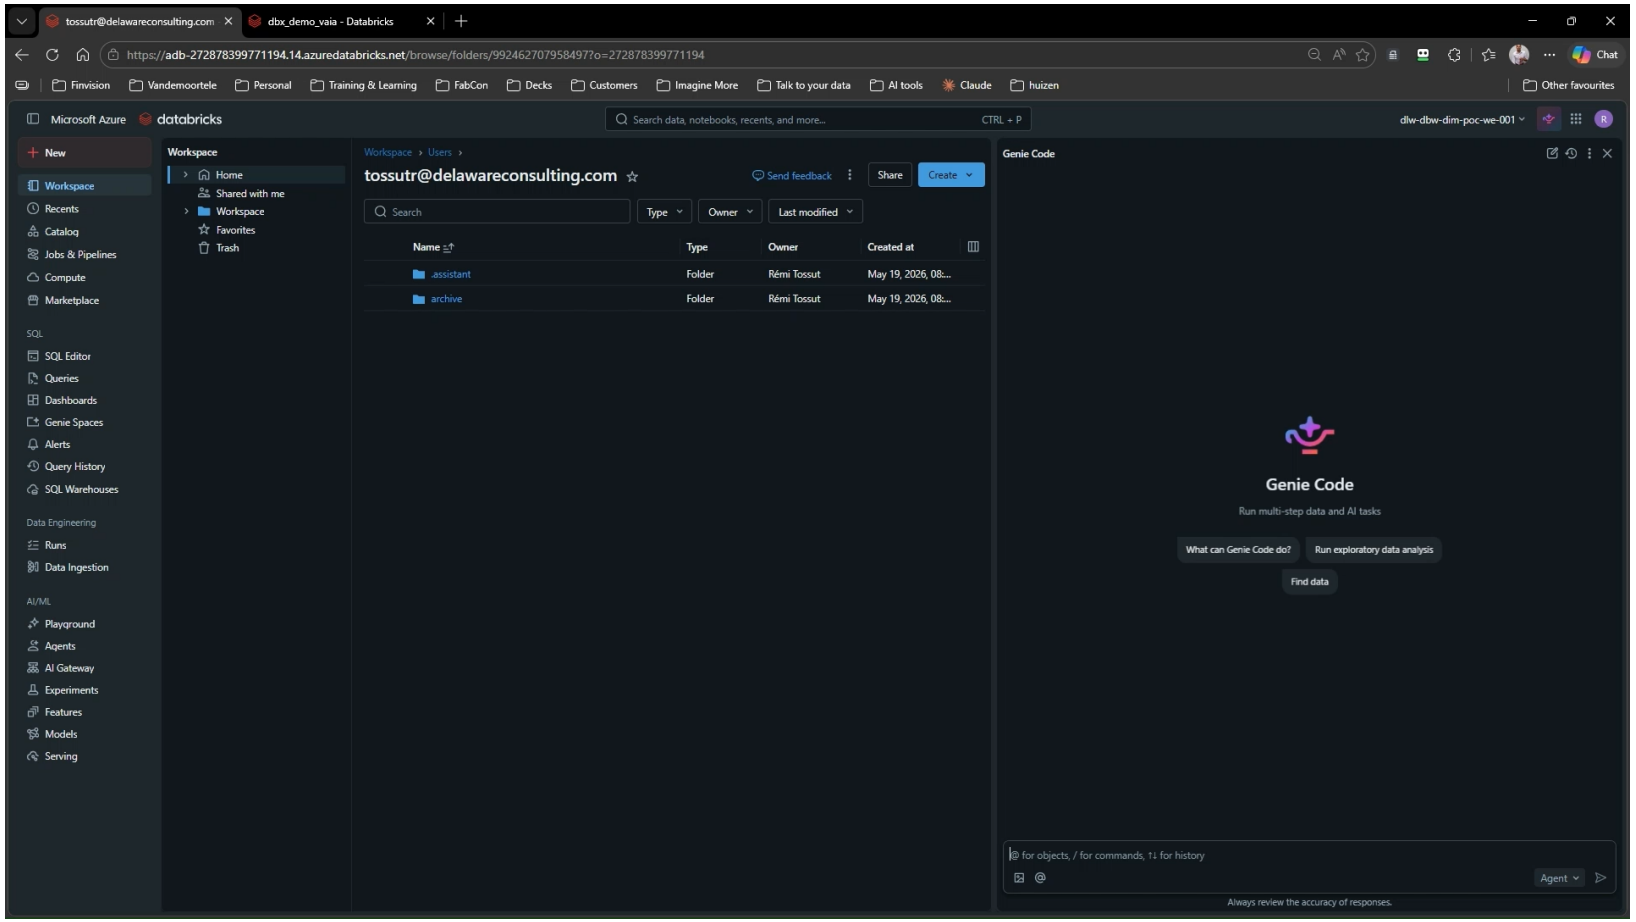

Demo 3 — Enhance Semantic Layer

3. Enhance Semantic Layer

- Refine measures, relationships and metadata on the Power BI model.

Databricks ·Genie

Claude + PBI MCP

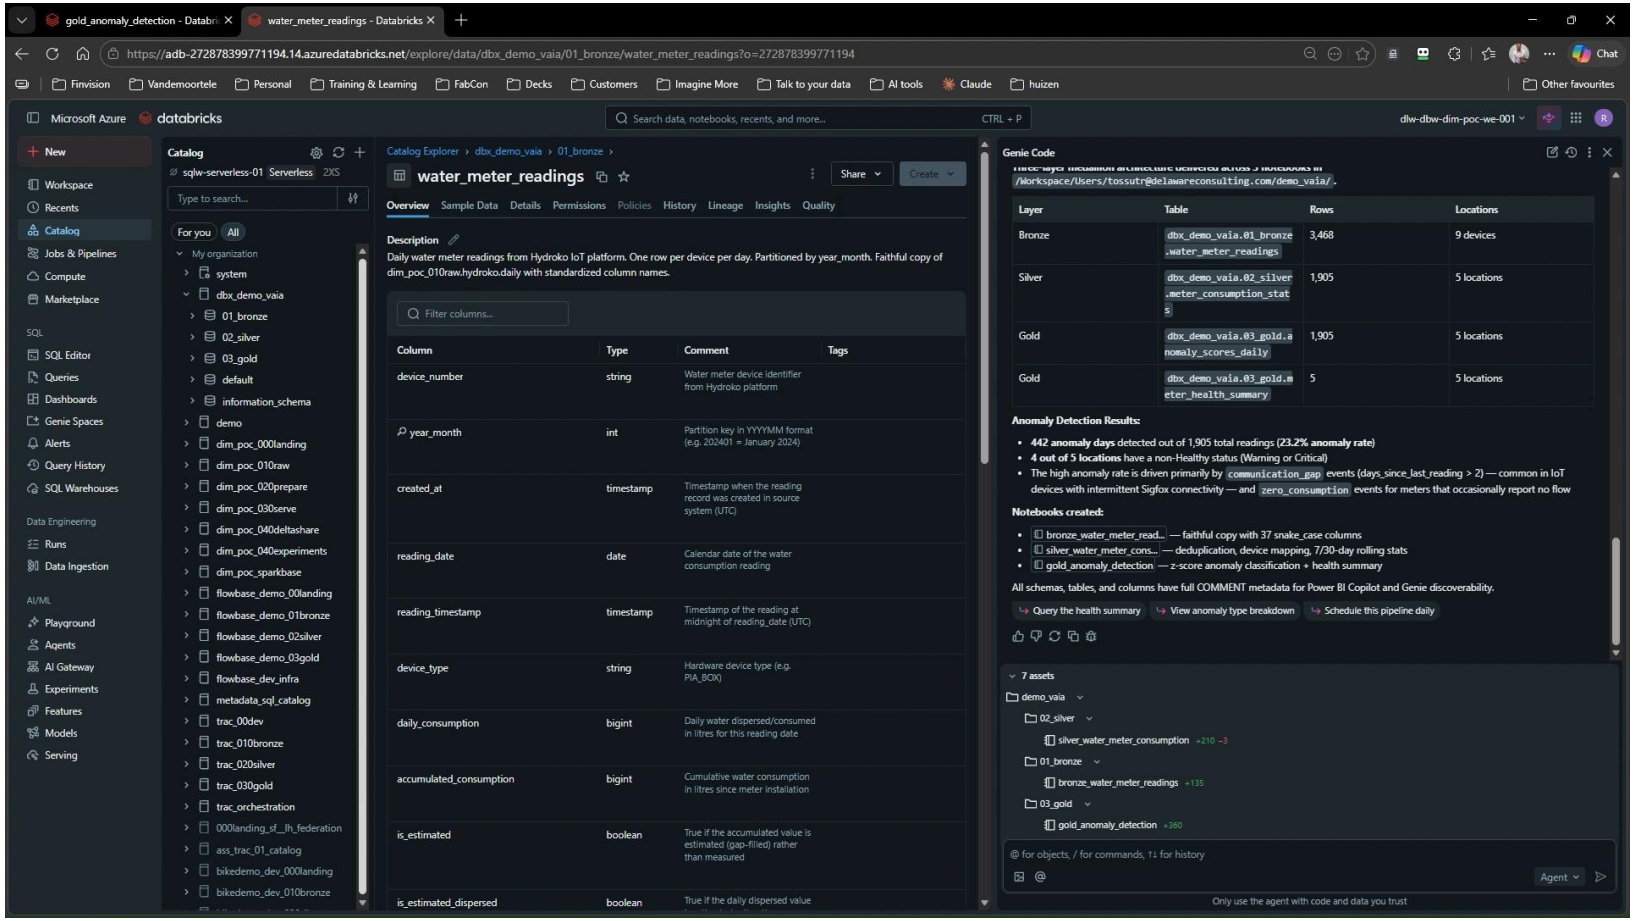

Consume

5 waves of ‘BI is dead’

Each wave was driven by a new technology paradigm, each time dashboards survived

The real question was never whether to replace dashboards, but **who was underserved by them**

The BI adoption gap

20% POWER USERS

Analysts, data professionals — comfortable with BI tools, building and consuming dashboards.

Recurring questions | Known KPIs | Visual exploration

80% EVERYONE ELSE

Business managers, operations, executives — need answers without knowing how to query data.

Ad-hoc questions | No BI skills needed | Conversational access

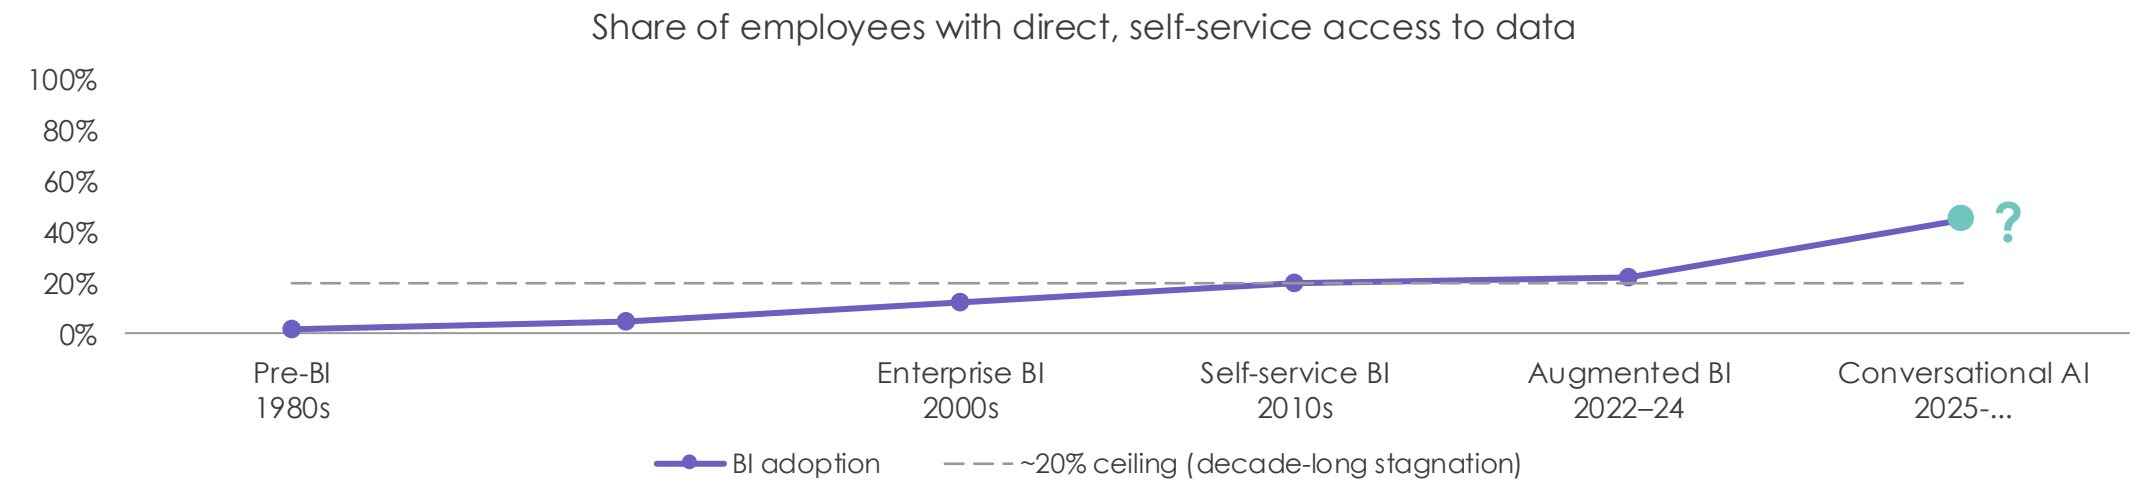

Confirmed data: Forrester, Boris Evelson — "Bring Data to the Other 80% of BI Users" (2024); BARC & Eckerson Group BI Adoption Study (2022); Gartner BI Maturity Survey (2018)

Self-service BI never broke the 20% ceiling, no matter how many 'user-friendly' features vendors added. Conversational BI is the first paradigm shift with a credible case for reaching further.

two fundamentally different approaches to data access

"Why did churn increase last quarter in the EMEA segment?"

TRADITIONAL BI

1. User files ticket with analytics team
2. Analyst checks if EMEA churn is in the model
3. Data isn't segmented — new ETL needed
4. Sprint planned, built, deployed
5. Answer arrives — weeks later

Time to answer: days to weeks

CONVERSATIONAL BI

1. User types question in natural language
2. Orchestrator identifies relevant sources
3. Joins CRM + billing + ERP dynamically
4. Answer with context, in seconds

Time to answer: seconds to minutes

Decision Matrix

DECISION FACTOR	USE DASHBOARDS WHEN…	USE CONVERSATIONAL BI WHEN…
Question type	Recurring, stable, known KPIs	Ad-hoc, unexpected, one-off
User profile	Analytics-trained, BI-literate	Non-technical, time-pressured
Primary goal	Monitor & align on shared numbers	Explore, investigate, find causes
Governance need	Audit trail, single version of truth	Speed & flexibility over formality
Analyst capacity	Team available, backlog manageable	Multi-day waits for ad-hoc answers

Demo 4 — Consume the Data

4. Consume the Data

- Query the semantic model in natural language to get trusted answers fast.

Claude + PBI MCP

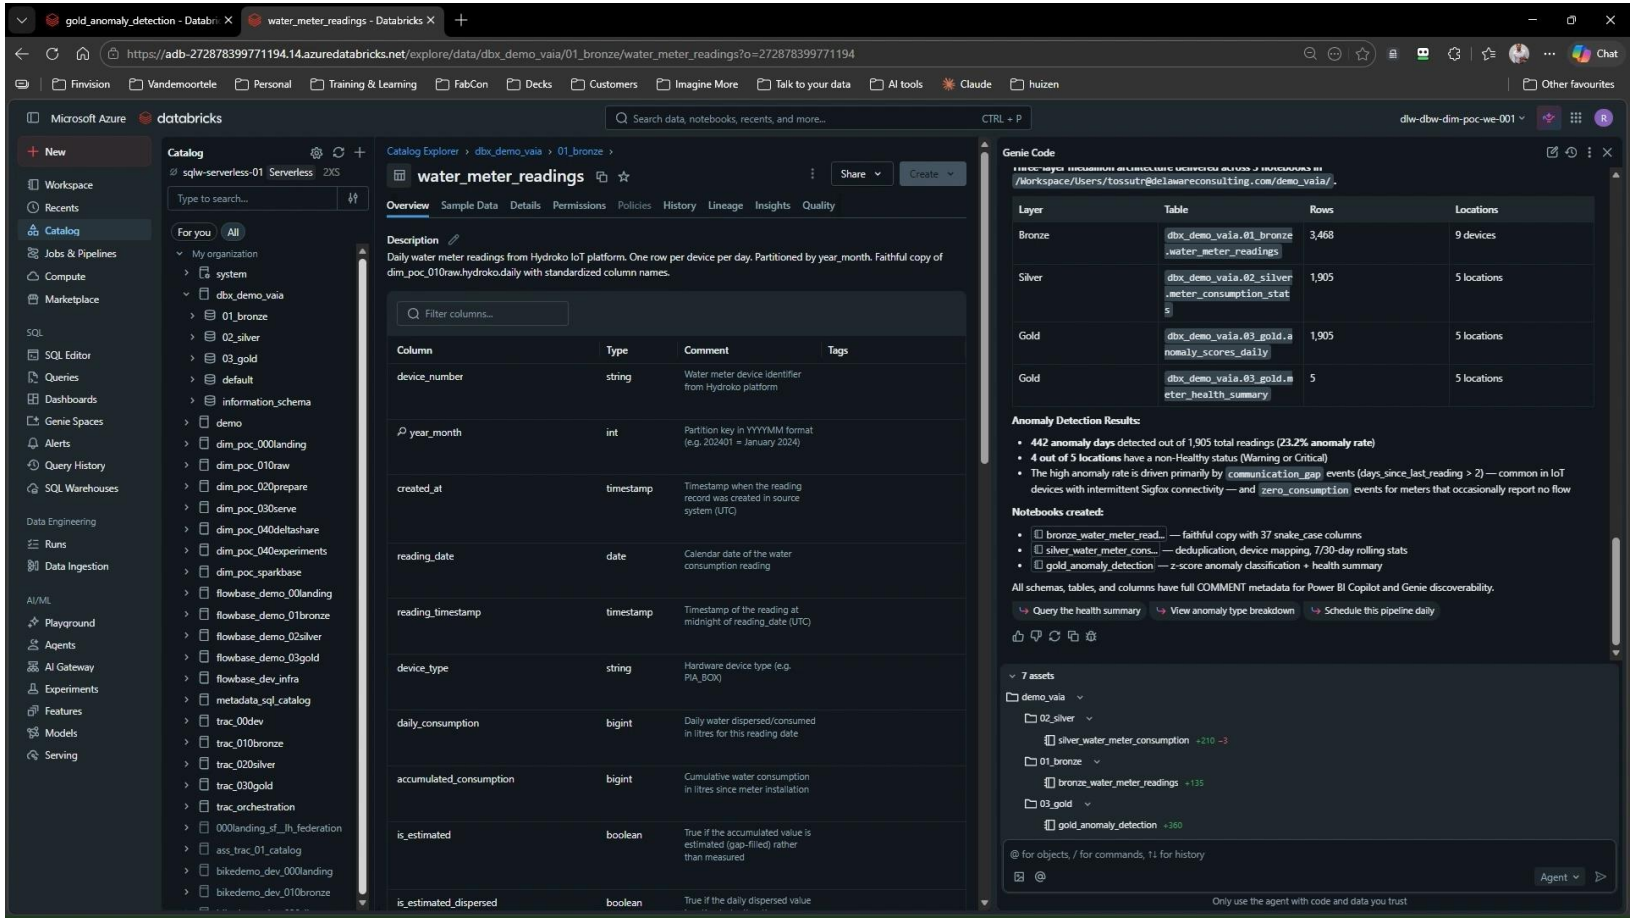

we commit. we deliver.

Conclusions

AI Agents for Data

Build

Pipeline & Transformation

AI generates logic
AI generates pipelines
Human verifies Build

Operate

Quality & Observability

Data Quality
Monitoring
Optimisation

Consume

Analytics & Decisions

Natural language querying
Prescriptive insights
Ad-hoc questions

thank you for your time

questions welcome

The materials, including any images in this presentation, are protected under international copyright law.

we commit. we deliver.

delaware