UNIVERSITY OF POTSDAM · 2 NOVEMBER 2026

The All Hands Intelligence Model

Rethinking enterprise AI beyond dashboards

Olalekan Elesin · Data & AI Executive · AWS ML Community Hero

Agenda

1 **Why dashboards stalled**
The question nobody can answer quickly, and the assumption behind it

2 **The old way versus the new way**
Two process maps of the same question

3 **The All Hands Intelligence Model**
Four layers, a new organisation, three changing roles

4 **Making it real**
An open-source stack, ownership, and how to measure value

5 **Where to start**
A maturity check, Monday actions, and open research questions

MONDAY, 9:00

“Why did conversion drop last week?”

One question from the CEO. Count the people, tools and days it takes to answer it.

Hands up if you have waited more than a week for a number like this.

25%

of employees actively use BI and analytics tools

The average in a BARC and Eckerson Group survey of 214 companies, with minimal growth over the seven years they had tracked it.

Two decades of dashboards, and three in four people still ask someone else.

Source: BARC & Eckerson Group, 2022 · 4

The constraint was never access to SQL

What we assumed

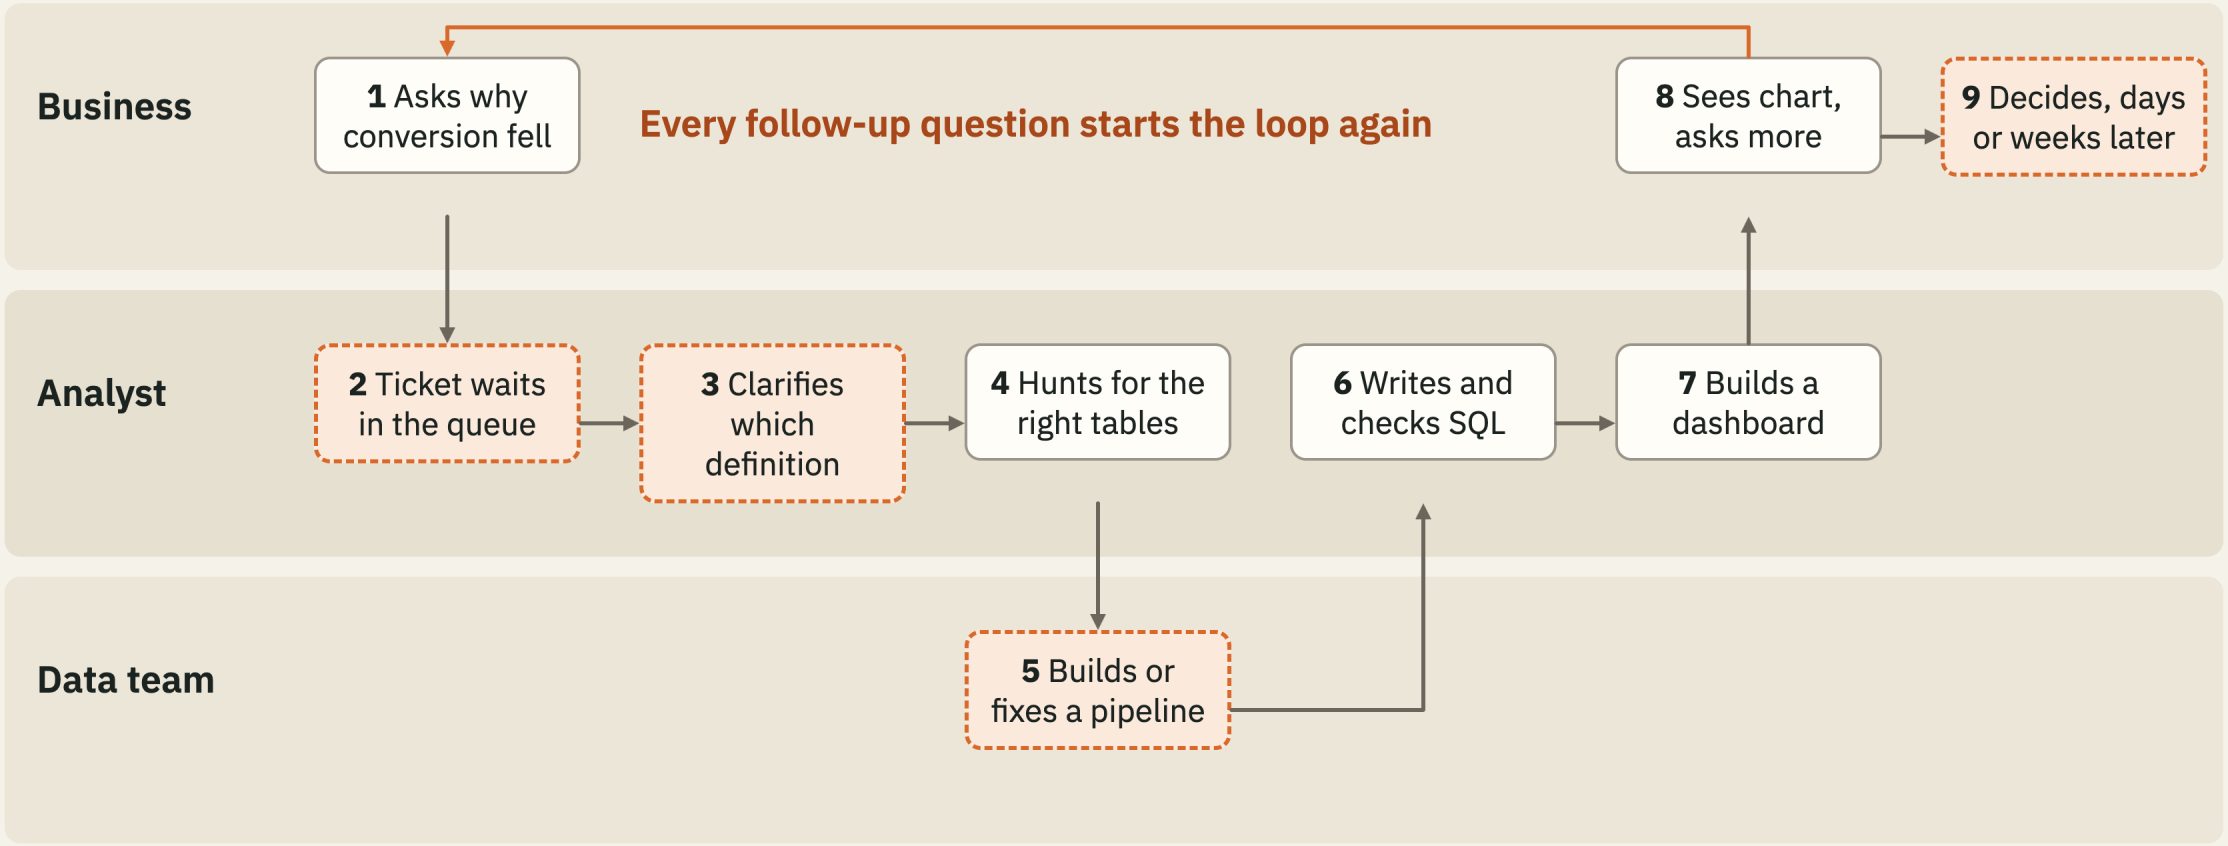
Insight is scarce because few people can write queries. So we hired analysts, bought BI licences and built dashboards. Capacity grew in a straight line with headcount.

What actually limits us

Meaning. What “active customer” or “conversion” means lives in a few people’s heads. An agent can write SQL in seconds; it cannot guess a definition nobody wrote down.

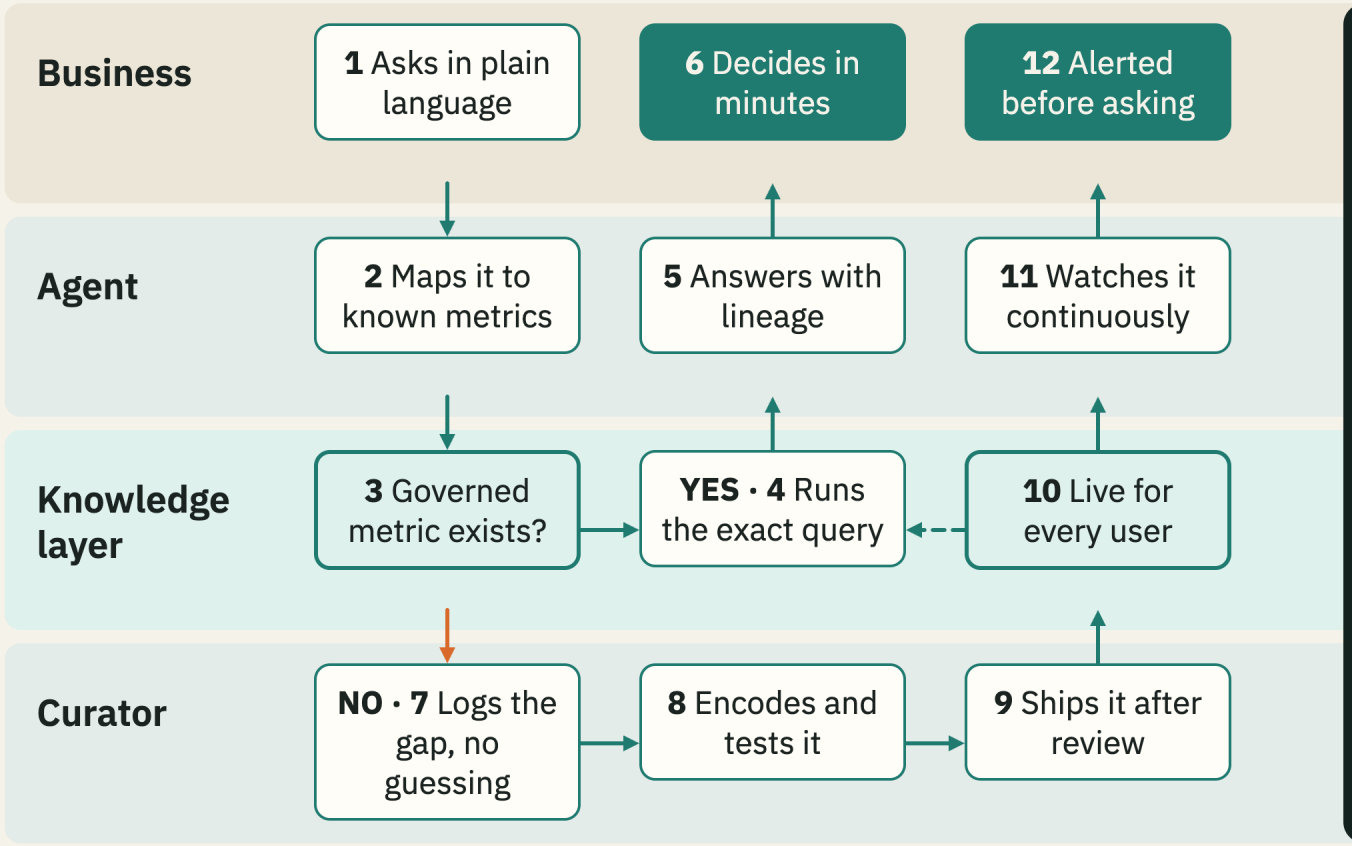
The scarce asset is the knowledge layer between raw data and decisions.

The old way: every question becomes a project



Process map, request-driven analytics. Orange dashed: waiting, not work · 6

The new way: questions become reusable knowledge



WHAT CHANGED

Minutes, not weeks

No queue between the question and the governed answer.

Gaps become assets

A missing definition is logged, encoded once and served to everyone.

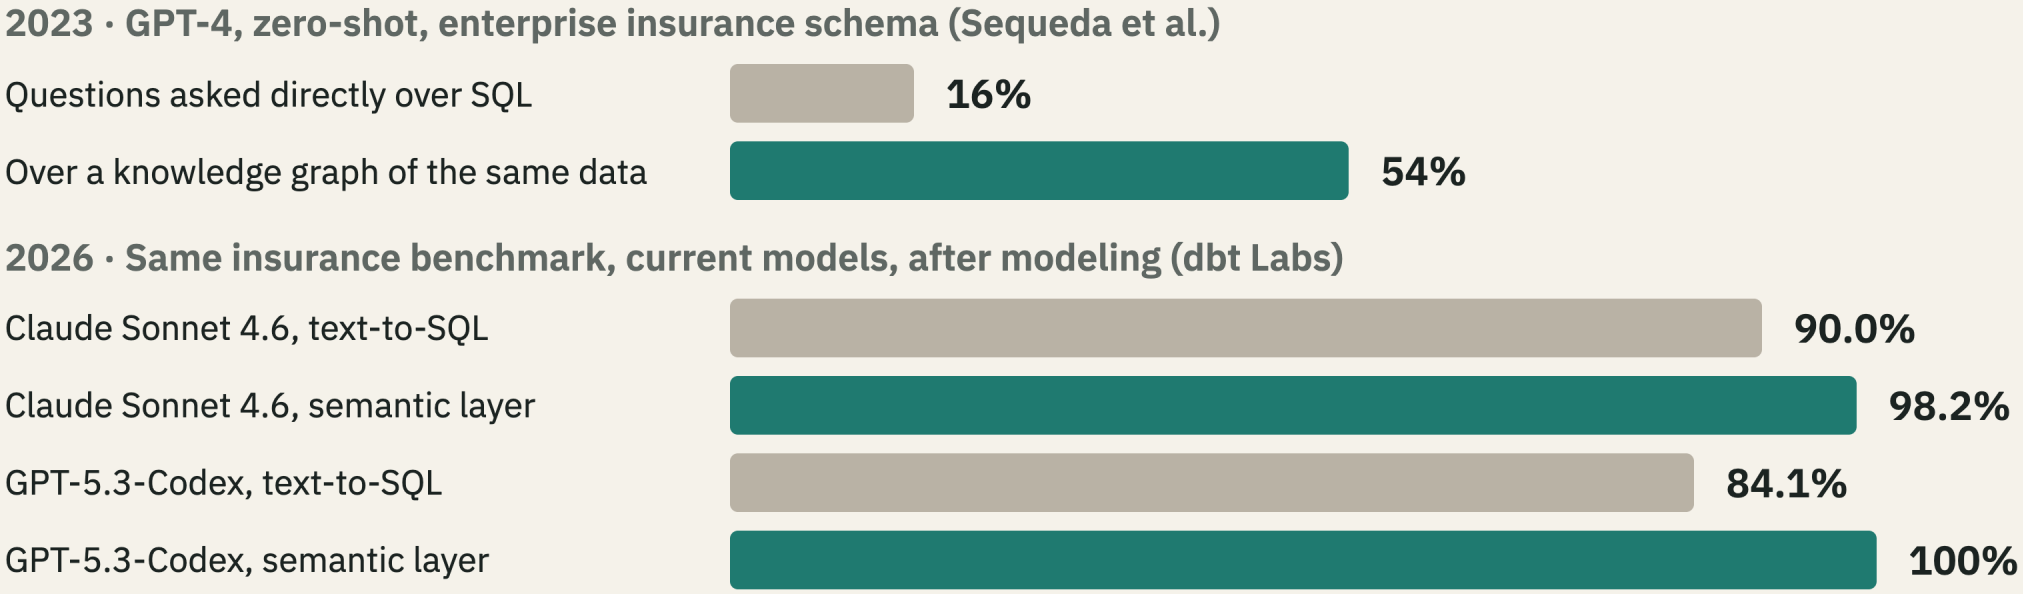
Agents go first

Step 11 watches the metric and alerts before anyone asks.

Side by side: what actually moved

	Old way	New way
Unit of work	A query for one person	A governed definition for everyone
Where knowledge lives	In analysts' heads and chat threads	Versioned, tested, owned
How it scales	One more analyst per queue	One more definition per gap
Timing	Reactive: after someone asks	Proactive: agents monitor and alert
When it does not know	Someone improvises	It says so and logs the gap

Grounding in business meaning lifts accuracy



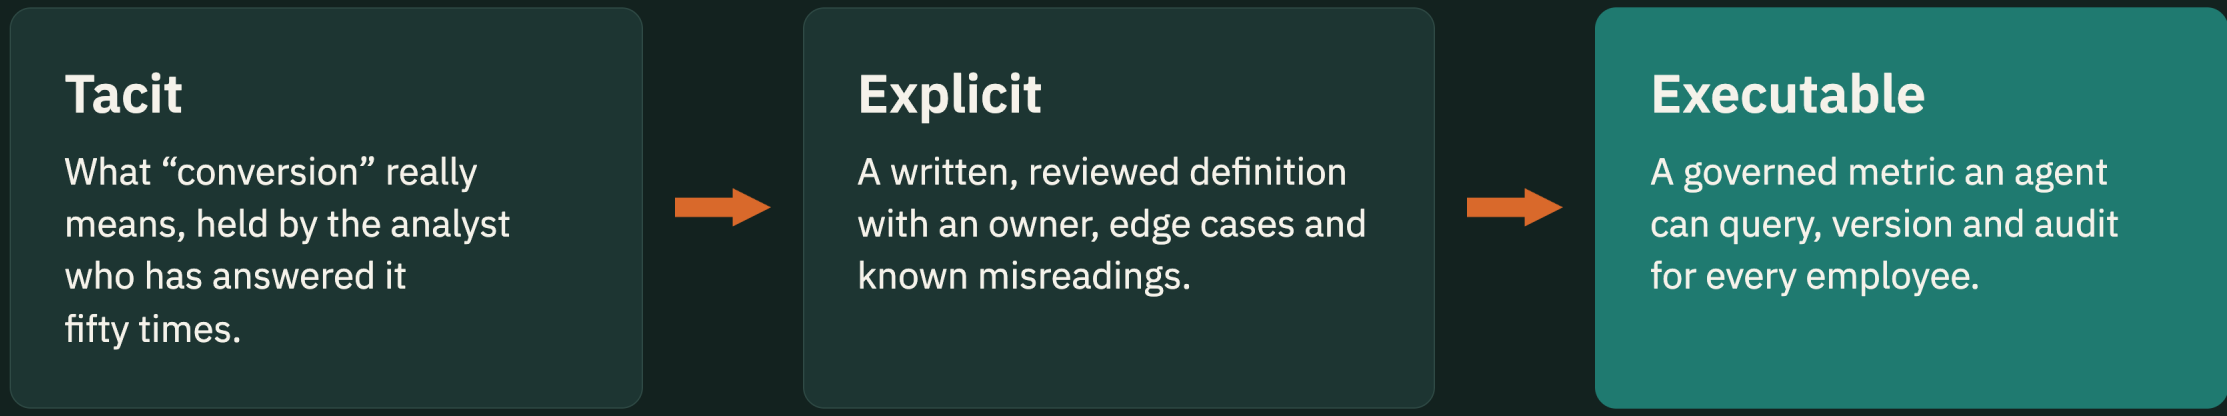
> **Caveat:** both benchmarks were run by vendors with a stake in the result (data.world, dbt Labs), on one insurance schema with small question sets. Directional, not conclusive.

Better models narrowed the gap. Governed definitions closed it on this benchmark.

Sources: Sequeda, Allemang & Jacob, arXiv 2023; dbt Labs, 7 April 2026 · 9

Standing on KMDL's shoulders

KMDL, the Knowledge Modeling and Description Language developed by Professor Gronau and colleagues, represents the development, use, offer and demand of knowledge along business processes, tacit knowledge included.



The semantic layer is knowledge conversion made machine-readable.

The All Hands Intelligence Model: four layers

People	A quarter of staff use BI; the rest ask	→	**Everyone asks in plain language**
Agents	Chatbots bolted onto dashboards	→	**Agents that answer, monitor and act with approval**
Knowledge layer	Definitions in heads and threads	→	**Governed, versioned definitions with named owners**
Data platform	Pipelines built per request	→	**Same platform, plus lineage and quality signals**

Accuracy is decided in the third layer. Most budgets go to the second.

From service desk to platform plus domains

Data & AI platform team

Pipelines, semantic engine, agent runtime, catalogue, governance. Enables the domains; owns no business logic.

Commercial & revenue

Analytics lead, semantic curators for pricing and pipeline, analytics engineer

Product & growth

Analytics lead, curators for engagement and acquisition, analytics engineer

Operations & finance

Analytics lead, curators for unit economics and cost, analytics engineer

Cross-domain metrics

Revenue, blended acquisition cost, churn: ratified federally, never owned by one domain.

Similar headcount. Different accountability: knowledge belongs to versioned data products, not people.

Three roles change the most

Analyst

WAS

Answers questions from a queue

BECOMES

Semantic curator: encodes and tests definitions

Analytics engineer

WAS

Builds tables that analysts query

BECOMES

Architect of the layer every agent queries

Analytics lead

WAS

Coordinates queues and stakeholders

BECOMES

Owns a domain’s semantic layer as a product

“I answer your data questions” is vulnerable. “I govern the system that answers them reliably” is not.

One open-source stack: Vanna 2.0

Step in the new way	Vanna 2.0 capability
Ask in plain language	A ready-made chat web component for React, Vue or plain HTML
Know who is asking	Identity from your existing cookies, JWTs or sessions
Ground in business meaning	Context enrichers add your documentation and definitions
Run the query safely	Permission-checked tools; row-level security per user
Trace every answer	Every query logged per user for audit
Choose model and database	Bedrock, Anthropic, OpenAI, Ollama; Postgres, Redshift, DuckDB

MIT licensed, so anyone in this room can run it. It does not write your definitions: the knowledge layer is still yours to build.

Sources: Vanna 2.0.2 on PyPI, 2 Feb 2026; vanna.ai documentation · 14

When the agent is wrong, who answers for it?

Not the model. Not the vendor. The owner of the definition.

Metric registry

Every metric has an owner, a definition, a sign-off and a last-validated date.

Review cadence

High-usage definitions are re-checked against how the business works today.

Audit trail

Every answer traces back to the exact definition that produced it.

Industry signals, not settled evidence

95%

of enterprise generative AI pilots delivered little to no measurable P&L impact

MIT NANDA, The GenAI Divide, 2025. An industry report, not peer-reviewed, with success measured narrowly as P&L impact. Its diagnosis: a learning gap and flawed integration.

>40%

of agentic AI projects will be cancelled by the end of 2027, Gartner predicts

Gartner, 25 June 2025. An analyst forecast, not an observed rate. Reasons given: escalating costs, unclear business value, inadequate risk controls.

Treat them as signals. Both point at the same gap: context, integration and ownership.

Sources: MIT NANDA via Fortune, 18 Aug 2025; Gartner, 25 June 2025 · 16

How would we measure the value?

1 · Analyst capacity reclaimed

Measure: share of recurring questions answered without an analyst, and where the freed hours go.

2 · Decision latency

Measure: time from question to decision, before and after, on a few high-value decisions.

3 · Reach

Measure: monthly active data users and cost per answered question.

WHAT MAKES THIS HARD TO MEASURE

Baselines

Few organisations log today how long a question takes to answer.

Attribution

Separating the agent's effect from everything else that changed at the same time.

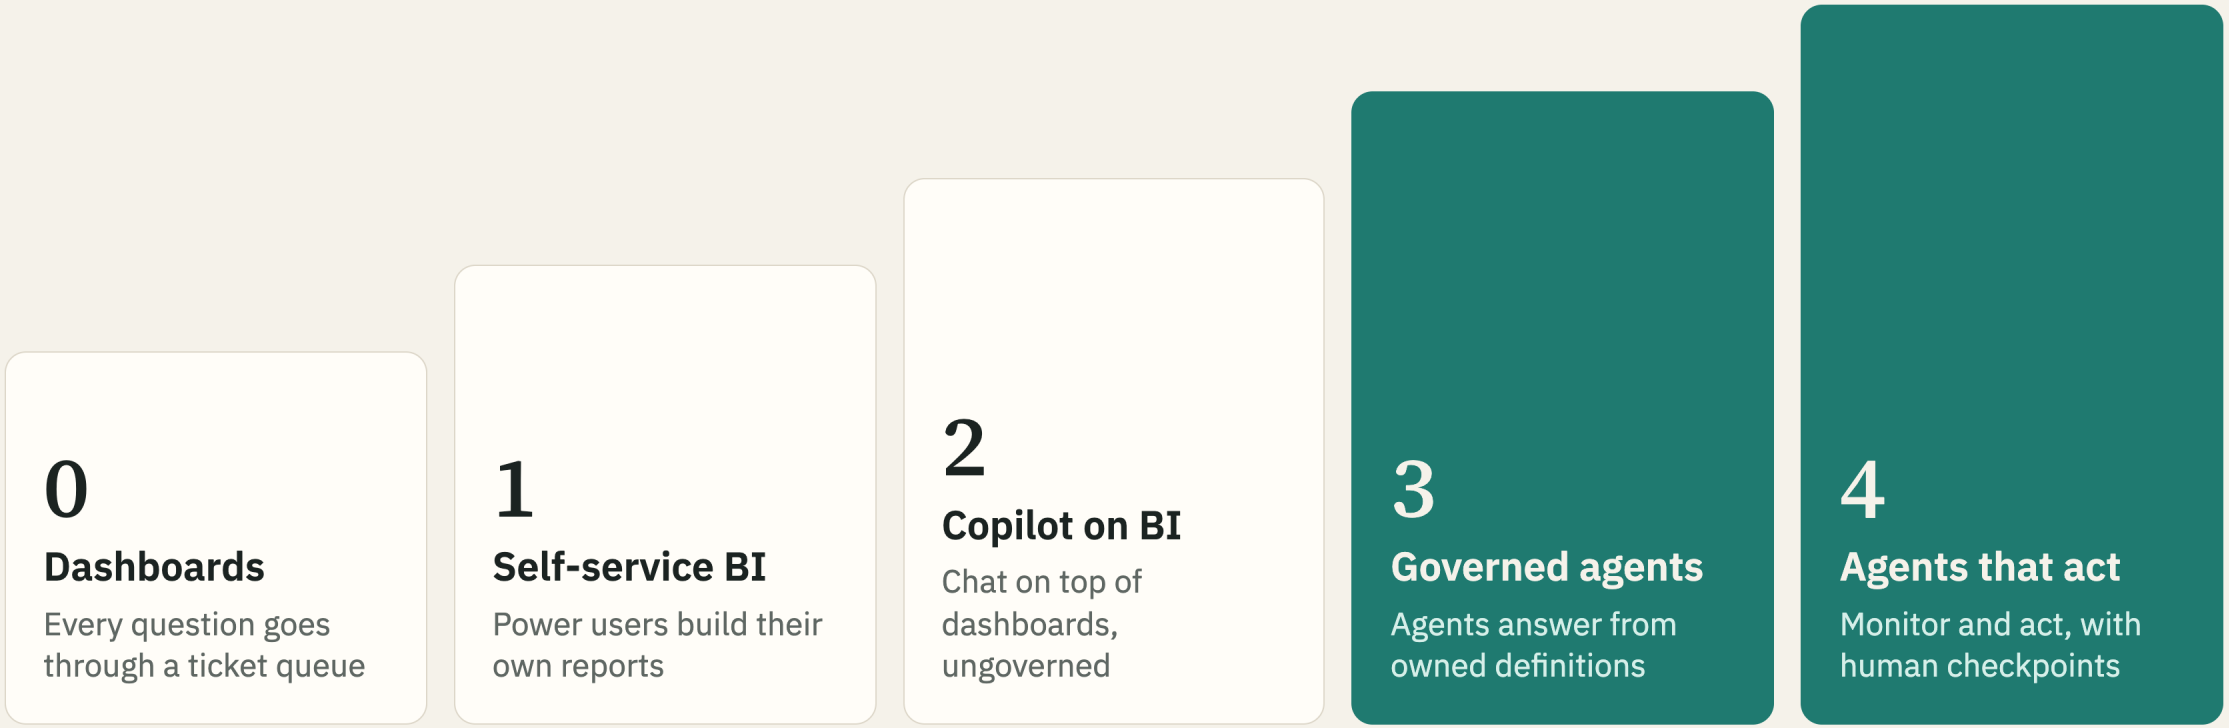
Answer quality

A fast wrong answer is a cost, not a gain, so accuracy has to be measured alongside speed.

Good candidates for a joint study.

A proposed maturity scale for agentic analytics

Think of an organisation you know well, from work, an internship or a case study. Where does it sit?



What to do on Monday morning

If you lead a data function

1. Move analyst time from report delivery to building the knowledge layer.

2. Name an owner for each of your most-used metrics.

3. Measure active data users and cost per insight, not dashboards shipped.

If you are an analyst or student

1. Pull your last 90 days of queries and write down what each metric means.

2. Rewrite model and column docs for business meaning, not data types.

3. Point an agent at your definitions and log every place it fails.

The analyst's output used to be a query. Now it is a governed definition that everyone can ask.

Certainty is not coming. Build the knowledge layer first.

Olalekan Elesin · elesin-olalekan.medium.com · Questions welcome

Sources: research and analysis

BARC & Eckerson Group, Strategies for Driving Adoption and Usage with BI and Analytics, 5 April 2022
barc.com/news/new-study-identifies-drivers-of-bi-and-analytics-adoption

Sequeda, Allemang & Jacob, A Benchmark to Understand the Role of Knowledge Graphs on LLM Accuracy on Enterprise SQL Databases, Nov 2023
arxiv.org/abs/2311.07509

Ganz & Perigaud, dbt Labs, Semantic Layer vs. Text-to-SQL: 2026 Benchmark Update, 7 April 2026
docs.getdbt.com/blog/semantic-layer-vs-text-to-sql-2026

Gronau, Müller & Korf, KMDL: Capturing, Analysing and Improving Knowledge-Intensive Business Processes, J.UCS 11(4), 2005
doi.org/10.3217/jucs-011-04-0452

MIT NANDA, The GenAI Divide: State of AI in Business 2025, reported by Fortune, 18 August 2025
fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots

Gartner, Over 40% of Agentic AI Projects Will Be Canceled by End of 2027, press release, 25 June 2025
gartner.com/en/newsroom/press-releases/2025-06-25

Elesin, Rethinking Analytics Teams in the Age of Agentic AI, Medium, 6 May 2026
elesin-olalekan.medium.com/rethinking-analytics-teams

Sources: software and standards

Vanna 2.0.2 on PyPI: architecture, web component, identity, row-level security, audit, models and databases, 2 February 2026
pypi.org/project/vanna

Vanna, Why We Built Vanna 2.0: MIT licence and identity-first design
vanna.ai/docs/why-we-built-this

Snowflake, Open Semantic Interchange: first specification published, 27 January 2026
snowflake.com/blog/open-semantic-interchanges-specs-finalized

AWS, Amazon Quick Suite general availability, 9 October 2025 (commercial alternative, speaker notes)
aws.amazon.com/about-aws/whats-new/2025/10/amazon-quick-suite

Discussion Points

1. How do we model the conversion of tacit analyst knowledge into executable definitions, and measure what is lost?
2. How do we audit decisions that an agent shaped, at the scale of every employee?
3. How should a federated organisation govern metrics that span domains?
4. How do we price decision latency, so the business case rests on evidence?