thinkia

BUSINESS INTELLIGENCE · JUNE 2026

Ask, don't *filter*.

The new business intelligence is a conversation with data, not a visualization of it.

EDITORIAL

For years I've walked into boardrooms with more dashboards than ever and watched decisions get made the way they were a decade ago: late, and on instinct.

I've stopped believing the problem is training or design. The dashboard does one thing well: it shows. We ask it to do another: reason. That mismatch isn't fixed by a better panel.

At Thinkia we start from an uncomfortable idea: the asset was never the chart. It was the shared semantics and the answer you can trace. On that base, conversing with data stops being a demo promise and becomes a reliable way to decide.

This document doesn't sell a product. It describes a shift: stop filtering, start asking. Anyone looking for a catalog won't find one here; anyone looking for the why and the how of that shift will.

If you finish it rethinking how your organization talks to its data, it will have done its job.

Rafael Conde
Chief Executive Officer · Thinkia · Madrid, June 2026
thinkia.com

EXECUTIVE SUMMARY

Large organizations have never had more dashboards, and yet they decide with the same slowness as a decade ago. The problem isn't a shortage of panels; it's that the panel only answers the question you already arrived with, and most decisions begin before you know what to ask.

This document defends a simple, demanding idea: the unit of the new business intelligence isn't the chart, it's the conversation with the data. The system stops waiting to be interrogated and starts pointing out what to look at, why it matters, and on what evidence.

Here are the problem, the anatomy of what causes it, the inversion of the flow that corrects it, the architecture that sustains it, and the errors that break the demos once they reach production.

23×

DECISIONS

greater effectiveness in organizations that decide with data.

$4.4T

ANNUAL IMPACT

in potential value from generative AI for the global economy.

75%

AI ADOPTION

of knowledge workers already use it in their work.

CONTENTS

01

Why Nobody Opens the *Dashboard*.

The past decade was sold as the era of data. Companies bought Tableau, Power BI, and Qlik licenses, built analytics teams, and filled their intranets with panels. The promise was that, with enough visualization, decisions would become obvious. They didn't. Spending on tools grew; the pace at which decisions get made did not.

The reason is structural, not one of adoption. A dashboard is a query device: it waits, idle, for someone to interrogate it. It answers the question the analyst already knew to ask, and stays silent on the one that truly mattered, the one no one posed because no one knew it had to be posed. The panel shows what you asked to see. It rarely shows what you should be looking at.

The result is an uncomfortable paradox. The more visual surface an organization produces, the fewer decisions it generates per unit of effort. Real adoption of BI platforms has been stuck near 30%[1] for years, a figure that barely moves no matter how many licenses are renewed. The other 70% is investment no one consults.

More dashboards than ever, and most of them unopened for weeks. The metric that best describes the state of corporate BI isn't one of performance; it's one of abandonment.

70%

of corporate dashboards go unused: adoption of analytics platforms stays near 30%.

Source: Gartner, BI adoption benchmarks.

It's worth naming the cost precisely, because the license bill is the smallest of its line items. Every panel no one opens is a decision made on instinct when evidence was available, a data team that spends its week maintaining reports no one reads, and an organization that believes it is data-governed while operating out of habit.

The usual diagnostic error is to treat this as a training or design problem: if people knew how to read the panels, or if the panels were prettier, they'd use them. It's a comfortable explanation and almost always false. The tool works exactly as it was conceived. The problem is the conception.

A dashboard shifts the entire cognitive load of analysis onto the human. It asks them to know in advance which metric to watch, which filter to apply, and which comparison makes sense. That load is invisible in a demo and exhausting day to day. The consequence is predictable: the panel gets consulted for the first two weeks and then abandoned, not for lack of value but for excess of friction.

The question that organizes the rest of this document is direct. If the classic dashboard fails by design, what replaces it? The answer isn't a better panel. It's a change in who does the work: the system stops waiting for the question and starts asking it.

> "
>
> **More visualization has produced *less* decision, not more.**

02

The Three Cracks in *Traditional* BI.

To replace something, you have to understand why it breaks. The classic dashboard doesn't fail from one bad implementation; it fails from three limitations shared by every platform, however good. These aren't product defects. They're properties of the model.

01

It requires knowing what to ask before you look.

A panel is a frozen answer to a question someone decided in advance. To be useful, the user has to arrive with the hypothesis already formed: which metric, which dimension, which period. The most valuable knowledge, the question you hadn't thought to ask, falls outside by construction, because no one designed a view for it.

02

The friction of filtering is a tax.

Every finding is several clicks away: open the panel, choose the filter, cross the dimension, adjust the range. Each click is small; the sum, across hundreds of queries a month, is the real reason the panel gets abandoned. Friction doesn't show in the demo and decides adoption in production.

03

Semantics are fragmented by department.

"Revenue" doesn't mean the same thing in finance as in sales; "active customer" shifts between marketing and operations. Without a shared definition, two correct panels produce two different figures, and the meeting is spent arguing which one is right instead of what to do about it.

The three cracks share a common origin. The dashboard is a medium of presentation, asked to do the work of a medium of reasoning. It presents figures with elegance; it doesn't understand the question, doesn't detect the anomaly, and doesn't reconcile two conflicting definitions. Those tasks are delegated to the person looking at the screen.

That's why "more BI" rarely solves the symptom. Adding panels multiplies the three cracks: more views that demand a prior hypothesis, more filters to cross, more loose definitions competing with one another. Saturation isn't a governance accident; it's the natural destiny of the tool when it's used for something it isn't.

The way out isn't polishing the surface. It's moving the reasoning work from the person to the system, and letting the screen go back to what it was always good at: showing, not deciding.

THE DIAGNOSIS

A medium of presentation doing the work of reasoning

Classic BI excels at showing. It's also asked to understand the question, watch for the unexpected, and unify semantics. Three things its design never included.

03

Inverting the Flow: From *Asking* to Being Told.

The dashboard lives in a pull model. The initiative belongs to the human: you ask, the panel answers. The system is passive by design and activates only when someone comes to it with a hypothesis. Anything not consulted does not exist.

AI-native business intelligence inverts the direction. It moves to a push model: the system observes the data continuously, recognizes what deviates from the expected, and comes to you with the finding, the explanation, and the next reasonable question. The initiative changes hands. You no longer ask what happened; you're told what to look at and why.

This isn't a product nuance. It's who carries the work. In the pull model, the cost of discovering something falls on someone happening to look for it. In the push model, discovery is the system's responsibility, and the person spends their attention judging and deciding, which is what they're hired for.

PULL MODEL

You ask

The panel waits, idle. What isn't consulted doesn't exist.

PUSH MODEL

It tells

The system watches, detects the deviation, and brings the question.

04

Three Layers of an *AI-Native* Dashboard.

Inverting the flow stacks into three capabilities. Each closes one crack in the classic model; together they turn consulting data into conversing with it.

01

Natural Conversation

You ask in language; you don't build a query. The system translates the sentence into SQL or MDX against the data model. The friction of filtering disappears: dimension, period, and comparison are inferred from the question.

02

Proactive Anomaly

The system doesn't wait for the question. It watches the series continuously, learns what's normal, and flags when something breaks the pattern (a margin slipping, a cohort drifting) before anyone thought to look. It closes the crack of having to know what to ask.

03

Generated Narrative

The output is the explanation, not the chart: what changed, against which reference, what explains it, and what decision to make. The figure arrives interpreted, with its context, ready to act on rather than to decode.

The three layers hold each other up. Natural conversation without shared semantics returns debatable figures; proactive anomaly without narrative generates noise instead of judgment; narrative without traceability is a well-written opinion. Value appears when all three operate over the same base of knowledge about the data.

It's the difference between a tool that shows and a collaborator that reasons. The chart is still there when needed; it stops being the starting point.

"

The figure stops being displayed and starts to *answer*.

05

The Architecture of a *Conversational* Dashboard.

Conversing with data sounds like magic and is engineering. What separates a flawless demo from a system that holds up in production is four pieces that work in order. A question doesn't go straight to the data: it passes through a controlled path that ends in an answer you can audit.

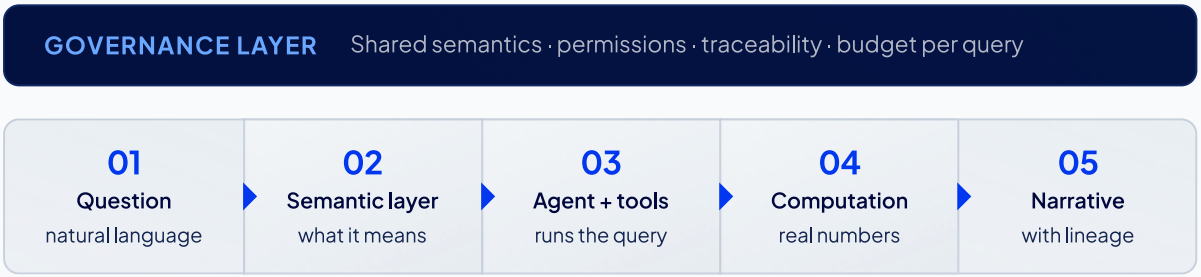

The path matters because each stage holds a failure the demos don't warn about. The semantic layer fixes what the terms mean; the agent executes instead of inventing; the computation runs on real data; the narrative arrives with its source attached. The following pages break it down piece by piece.

The Semantic Layer Is the Contract

Before translating any question, you have to fix what things mean. The semantic layer defines the metrics, the dimensions, and the relationships once, and every answer is computed against it. It's what closes the third crack: "revenue" has one definition, not one per department.

RAG over the Catalog, Not the Data

The model doesn't retrieve raw rows; it retrieves the knowledge about the data: metric definitions, table descriptions, business rules, and lineage, to understand the question and plan the computation. Retrieval feeds the reasoning; it never replaces the figure.

An Agent with Access to Tools

The central piece is an agent with real tools: it runs the query against the warehouse, calls the computation function, and composes the answer from what the engine returns. The rule is strict: the agent orchestrates and explains; the numbers always come from a query, never from the model.

KEY TERM

Data traceability

Every answer carries its full lineage: the source table, the query executed, the transformations applied, and a confidence score. Anyone can reproduce the computation and audit where each figure comes from.

It's the same discipline as in regulated systems: an answer is worth something not because it sounds right, but because it can be traced back to the data that supports it.

If you had to choose one piece as the backbone, it would be traceability. A conversational dashboard without lineage is more dangerous than useful: it produces convincing sentences no one can verify, and a single wrong number with good wording is enough for leadership to lose confidence in the whole system.

That's why the architecture is designed the opposite of how it's usually presented. First you guarantee that every answer is reproducible and that every figure points back to its source; only then do you make the conversation fluent. Fluency impresses in the demo; traceability is what survives the first audit.

> **An answer without lineage isn't a fact; it's a well-written *opinion*.**

06

What Changes in the *C-Suite*.

The value of inverting the flow shows better by function than in the abstract. Three distinct roles, three ways to stop filtering and start asking. These are patterns already running in production, described by what they change in the work, not as success stories.

THREE FUNCTIONS, ONE SHIFT

From interrogating panels to *talking with the business*.

CFO · Conversational close

Asks "why did the margin in Iberia fall this quarter?" and gets the breakdown by product and cost, with the lineage of each figure, without waiting for finance to prepare the report.

COO · Early deviation

The system flags a spike in returns at a plant on Tuesday, not in Friday's report, and proposes the next question: which batch, which shift, which supplier.

CMO · On-demand segmentation

Builds and compares segments by asking, without opening a ticket with the data team. The hypothesis is tested in the conversation, not in next week's request queue.

07

What Breaks the *Demos*.

A conversational demo is easy to impress with and hard to sustain. Four failures separate the launch video from the system that holds up through a quarterly close. They're worth naming, because each has a concrete containment and none is solved by a bigger model.

FAILURE	WHY IT HAPPENS	HOW IT'S CONTAINED
Numeric hallucination	*The model "fills in" a plausible figure when it doesn't have one.*	The agent only runs queries; it never generates numbers on its own.
Semantic ambiguity	*"Revenue" or "active" mean different things depending on who asks.*	The semantic layer is the single source of definitions.
No auditability	*No one can reproduce how the answer was obtained.*	Every answer carries its lineage: table, query, transformation, score.
Runaway cost	*Every question launches computation with no limit or measure.*	Budget per query, result caching, and cost-based routing.

The pattern is common to all four: the language model isn't the source of truth, it's the orchestrator. Truth lives in the data and in the semantics; the agent converses, retrieves context, and composes, but doesn't decide what the margin was. Confuse the two and you have a memorable demo and a system that doesn't survive the first committee meeting.

08

Measuring What Truly *Matters*.

If the classic dashboard was measured by number of panels published, conversational intelligence asks for different indicators. The question stops being how many views exist and becomes how many decisions get made, how fast, and with what verifiable accuracy. Four metrics are enough to tell whether the change is working.

METRIC	WHAT IT CAPTURES	SIGNAL OF SUCCESS
Decisions per user per month	*Real use of the system, not open sessions.*	Rises and holds after the first few weeks.
Time-to-insight	*Time from question to a grounded answer.*	Goes from hours or days to seconds.
Real adoption vs. licenses	*The gap between what's bought and what's used.*	The gap closes instead of widening.
Audited accuracy	*Answers verifiable against their source.*	Trends toward 100% reproducible.

The metric that costs the most and says the most is the first. Counting decisions, not logins, forces you to define what counts as a decision and to follow it to its effect. It's uncomfortable to instrument and it's the only one that tells a tool that's used from a tool that's bought.

The dashboard isn't read. It's *asked*.

Conclusion

Data Stops Being Displayed and Starts to *Answer*.

The dashboard isn't going to disappear as a screen. What changes is its job. It stops being the surface you interrogate and becomes the system that interrogates the data for you, brings you what's relevant, and explains it. The chart survives as illustration; it loses the lead role.

Seen this way, the asset was never the visualization. It was the shared semantics and the traceable answer: one definition of each metric and a figure you can trace back to its origin. On that base, conversation is the natural interface; without it, it's a trick that doesn't survive the first audit.

The organizations that win the next decade won't be the ones with the most panels. They'll be the ones that get their data to answer in sentences, with lineage, to the person who has to decide. The new business intelligence is a conversation with data, not a visualization of it.

> **The asset was never the chart.**
> **It was the *answer* you can trace.**

Ask, don't filter. Talk with the data.

THINKIA · JUNE 2026

REFERENCES

Sources

Figures and references cited throughout this document.

01 Gartner, *Magic Quadrant & Market Survey: Analytics and Business Intelligence Platforms*. BI adoption stays near 30% of the target headcount.

02 McKinsey & Company, *Five Facts: How Customer Analytics Boosts Corporate Performance*. Data-intensive organizations are up to 23× more effective at acquisition and decision-making.

03 McKinsey & Company, *The Economic Potential of Generative AI: The Next Productivity Frontier*, 2023. Annual potential impact of $2.6 to $4.4 trillion.

04 Microsoft and LinkedIn, *Work Trend Index Annual Report*, 2024. 75% of knowledge workers already use AI at work.

HOW TO CITE THIS PAPER

Conde, R. (2026, June). *Ask, Don't Filter: From Visualization to Conversation with Data*. Thinkia. thinkia.com

thinkia

thinkia.com

© 2026