

Cognitive BI

From Dashboards to Autonomous Analytics

AUTHOR:



Scott Watson

Vice President, Research



About Activant

Activant is a research-led global investment firm that partners with high-growth companies. Since 2015, we have invested in category-defining businesses during their most critical phases of growth, partnering with founders who have won their initial battles and are ready for the next challenge.

Our approach pairs deep, proprietary research with patient, flexible capital and hands-on operational partnership. We work alongside founders and leadership teams to refine strategy, strengthen operations, and accelerate sustainable growth.

Activant Research is dedicated to uncovering the most exciting emerging technologies, sectors, and companies we believe will shape the future. Our research-driven perspective informs everything we do, helping us invest at meaningful inflection points and support founders in building enduring, category-leading businesses.

You can find out more about Activant and our research at <https://activantcapital.com/>.

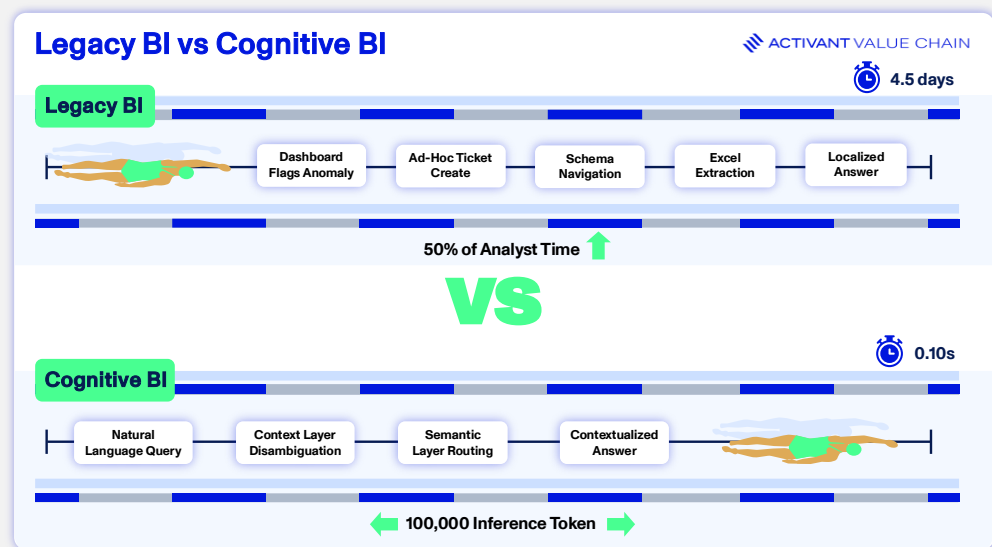
This is Part II of Activant's Cognitive BI series. [Part I \('Conversational Analytics: The Illusion of Chatting with Your Data'\)](#) diagnosed why Text-to-SQL fails in production and introduced the Semantic Layer and Context Layer architecture. This installment examines the commercial implications.

The Four-Day Answer

Picture a CFO spotting a sudden 4% margin drop in Europe during a quarterly close. The dashboard flags the number. It can't explain it.

She files a ticket.

That ticket lands with a junior analyst who, according to [dbt Labs'](#) 2025 State of Analytics Engineering report, already spends 57% of their time navigating messy databases before writing a single query.¹ Four days later, after exporting to Excel and crunching numbers manually, the analyst delivers an answer. By then, the decision window has closed. The company pays twice: once in wasted engineering time, and again in delayed action.



Now imagine the same CFO asking the platform directly: "Why are European margins down?"

Within seconds, the system scans recent Slack threads and detects that "Europe margin" currently excludes a newly acquired subsidiary. It locks the query to the company's GAAP-compliant gross margin definition via the Semantic Layer, generates the SQL, hits the database, runs the variance analysis, and surfaces the root cause: an unapproved discount from the German sales team.

The CFO has a verified, explainable answer before she finishes her coffee. No ticket. No analyst. No four-day lag.

What once took days now finishes in seconds, at a marginal compute cost measured in fractions of a cent per query. The AI isn't making a chart. It's doing the investigative work of a data analyst. The company effectively gains an unlimited, zero-latency data team.

That gap between the four-day answer and the four-second answer is where the entire enterprise BI market is about to fracture.

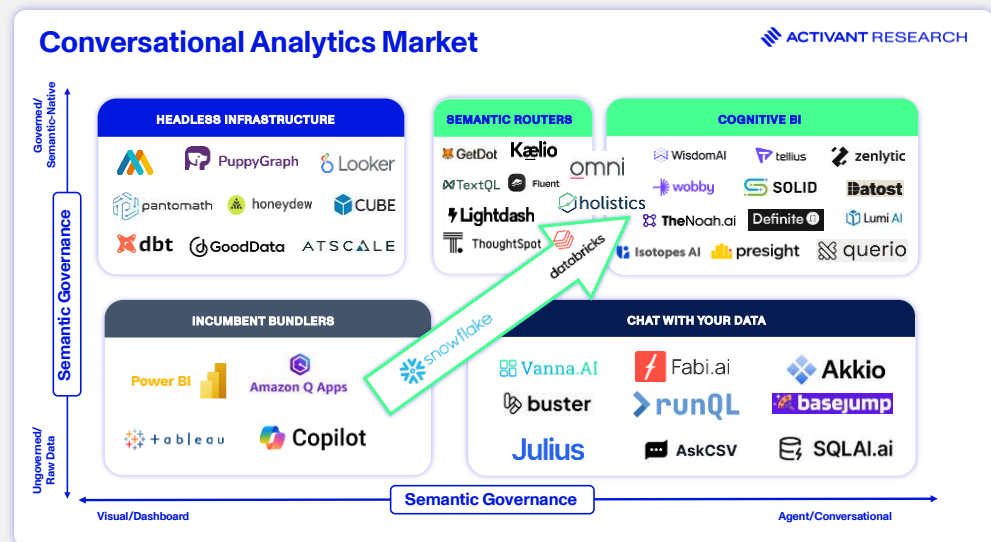
Why the Market is Splitting

In Part I, we showed that Text-to-SQL fails in production not because LLMs can't write code, but because they lack business context. Execution accuracy on real enterprise schemas can fall as low as 16.7%,² a number that explains why executives don't trust AI-generated reports.

It also reflects a broader pattern. MIT's 2025 State of AI in Business study found that 95% of enterprise AI pilots fail to deliver measurable P&L impact. The models aren't the problem; it's the systems around them that aren't architected for production.³

The fix we proposed in Part I is architectural: a Semantic Layer to enforce mathematical truth, and a Context Layer to capture institutional knowledge. But the harder question is commercial: who captures this value?

That question is splitting the vendor ecosystem into four camps, defined entirely by how they handle mathematical truth and operational context:



The Incumbents. Microsoft (Power BI) and Salesforce (Tableau) control the largest distribution footprints in enterprise software comprising hundreds of millions of seats and massive procurement relationships. But their architectures were built for visual dashboards, not conversational AI. They're now bolting generative text interfaces onto legacy engines to defend SaaS seat revenue.. We'll dig into how they're doing this, and why it's failing, in the next section.

The Chat-to-Data Startups. These solutions place a zero-shot language model on top of a cloud data warehouse with no semantic governance, business logic layer, or guardrails. The AI guesses at table joins and hopes for the best. It's effective if a skilled engineer is driving but risky if a non-technical executive is making decisions based on the output.

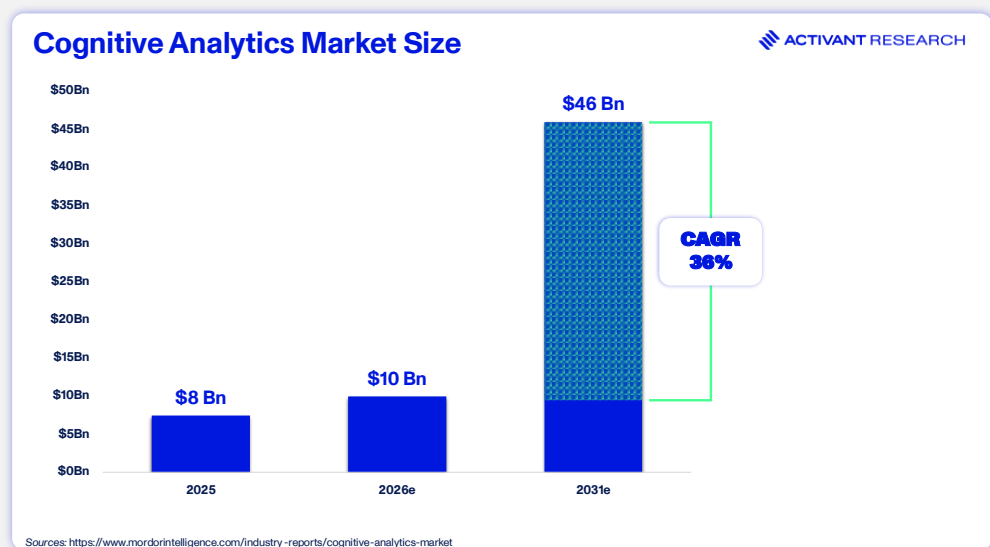
The Infrastructure Pure-Plays. Standalone semantic layers (e.g., dbt Labs or Cube) and context engines (vector databases and RAG platforms) solve the hallucination problem by enforcing absolute logic rules, but they remain isolated silos. Buying them means hiring data engineers just to stitch everything together.

The Unified Cognitive Platforms. Vendors like [Wisdom.ai](#), [Zenlytic](#) and [Holistics](#) collapse the stack into a single system, combining the semantic layer, the context engine, and the autonomous AI agent into one environment. They don't sell dashboard access, they sell verified analytical outcomes.

All four camps will coexist for a while. But enterprise budgets are shifting toward architectures that can guarantee accuracy without human intervention, and that favors the third and fourth camps.

The Market Behind the Split

The broader cognitive analytics market, encompassing natural language processing, machine learning platforms, and AI-native BI, is projected to grow from roughly \$10 billion in 2026 to \$46 billion by 2031, a 36% CAGR.⁴



Conversational analytics is a growing share of that market, driven by cheaper storage, better models, and enterprise demand for self-serve data access.

Industry estimates, admittedly rough and recycled so often they've become received wisdom, suggest that most enterprise data remains unstructured, buried in Slack threads, PDFs, and email chains. That's where the *why* behind the numbers lives, not in structured database tables. Legacy BI tools were never built to access it.

Meanwhile, underlying infrastructure economics are shifting. Data storage costs have fallen sharply, with some organizations reporting up to an 80% drop in data-lake operational costs.⁵ Cheaper storage, combined with faster analysis, is turning conversational analytics from an experiment into a hard enterprise requirement.

The Incumbent Tax

With active BI adoption still below 26%, legacy platforms have long struggled with usability.⁶ Part I made the case that these tools don't reduce cognitive load. They redirect it. Non-technical users still end up in IT ticket queues and Excel exports. Microsoft and Salesforce saw the threat coming from AI-native disruptors. Their response was defensive: bolt generative text interfaces onto existing rendering engines and hope the distribution moat holds.

It's not holding.



To their credit, both companies are moving toward semantic and context layers. Microsoft is embedding semantic models into Fabric and Salesforce is building out Data 360. The question isn't whether they'll get there but what it will cost customers along the way.

Microsoft: The Infrastructure Play

Microsoft reported 15 million paid M365 Copilot seats in January 2026, a 160% year-over-year increase.⁷ Impressive, until you consider that's just 3.3% of its 450 million commercial M365 base.⁸

Three percent. Despite \$37.5 billion in quarterly Capex and a full-court enterprise sales push, 97% of Microsoft's own customers haven't bought the AI add-on.

So where does the money come from? Not the interface, but the infrastructure underneath.

To unlock Copilot's full capabilities at scale, Microsoft steers buyers toward Fabric F64 capacity clusters. That's a minimum commitment of roughly \$60,000 to \$77,000 per year depending on region and contract terms.⁹ On paper, F64 bundles six workloads and eliminates per-user Power BI Pro licensing. In practice, it's an infrastructure floor that ties AI functionality directly to Azure compute consumption.

The Copilot interface is the hook. Fabric capacity is the revenue engine.

This isn't an organic software upgrade. It's a forced migration, with the AI interface acting as a lock-in mechanism that shifts costs downstream.¹⁰

Salesforce: The Churn Defender

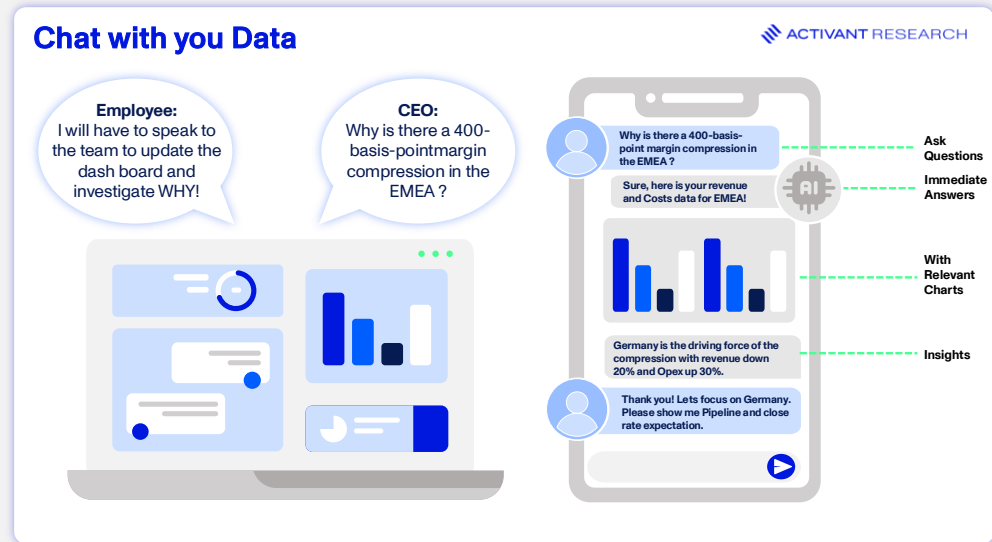
Salesforce is running a parallel strategy. Agentforce has reached \$800 million in ARR while Tableau's growth has faded in prominence.¹¹

The company deployed Tableau Pulse, which pushes automated metrics to users to combat dashboard fatigue and justify seat renewals. But unlocking Tableau's agentic features, requires moving to Tableau+ and higher-tier licenses. Enterprise-tier Creator licenses cost \$115 per user per month compared to the standard \$75 Creator license and the \$15 Viewer tier that most existing customers are on.¹² Adoption metrics reveal the gap: Salesforce has closed 29,000 cumulative Agentforce deals and reports \$800 million in ARR as well as strong growth with 50% quarter-over-quarter increases in both deal count and production accounts.¹³ But third-party analysis suggests the gap between signed contracts and active production deployment remains wide.¹⁴

That gap is where shelfware risk lives.

Both incumbents are investing in the right architectural direction. But retrofitting these capabilities onto legacy rendering engines pushes cost and complexity onto the buyer. The AI features only work once the underlying data is manually modelled. For enterprises weighing that cost against AI-native alternatives that work out of the box, the calculus is starting to shift.

Why 'Chat with Your Data' Isn't Enough



Chat-to-Data tools are powerful for experienced engineers. They eliminate grunt work and cut SQL drafting time from minutes to seconds. They also work perfectly in pristine sandboxes. Point them at a flat, simple dataset where complex table joins aren't required, and they perform reliably.

But these startups aren't selling to engineers. They're selling to non-technical business operators as autonomous tools. The accuracy data bears repeating: point an unconstrained model at a real enterprise schema and execution accuracy collapses to 16%.¹⁵

Presenting AI hallucinations as absolute mathematical facts destroys executive trust. An executive cannot make operational decisions on data they have to second-guess. If a CFO is forced to send an AI's output to the data engineering team just to verify the math, the entire purpose of a self-serve BI tool is defeated.

The Integration Trap

The market's first solution was modular: separate semantic layers to lock the math and context engines to capture institutional memory. That's the architectural split we unpacked in Part I.

Both work. But buying them separately creates its own mess.

Companies stack overlapping SaaS licenses while highly paid engineers burn through billable hours stitching silos together. The TCO arbitrage that made the modern cloud data stack attractive gets eaten alive by integration overhead.

The hallucination problem gets solved and a new one replaces it: integration hell, a fragmented stack nearly as expensive and labor-intensive as the old regime.

Unified Architecture

This is why unified cognitive platforms are gaining traction: they collapse the stack and make AI the execution engine rather than just a chat interface.

To illustrate what this looks like in practice, consider three startups taking different approaches to the same problem.

[Wisdom.ai](#) focuses on proactive context. Rather than building static dashboards, it deploys a context layer that bridges structured databases and unstructured company knowledge. Autonomous agents monitor KPIs continuously, detect anomalies, pull in relevant operational context from Slack threads and internal docs, and push an alert to the decision-maker before a ticket is ever filed.


The enterprise doesn't discover the problem during next week's review meeting. The platform catches it in real time. And it shows up with an explanation, not just a red number on a screen.

[Zenlytic](#) emphasizes explainability. Its engine locks generative AI to a rigid semantic model and exposes the underlying query logic and data lineage on screen. A non-technical executive can visually trace how the AI arrived at an answer: which tables it queried, which filters it applied, and which metric definition it used.

Once the executive can see the math on screen, they stop sending it back to engineering for verification. They can verify outputs directly, positively impacting the entire workflow.



The biggest misconception in AI analytics is that accuracy alone earns trust. It doesn't. If a CFO can't see exactly how an answer was derived, in a step-by-step way they can understand, they'll never act on it. That's why we built explainability into the core of the engine, where the math and the context are governed together from the start. In the AI era, the proof is the product."



Ryan Janssen
Co-Founder & CEO, Zenlytic

Holistics builds a semantic layer where metrics are first-class, composable assets. These are managed through a versioned, testable workflow so complex analytical questions can be answered reliably. It combines AI with BI interactivity, turning answers into interactive surfaces users can explore and refine in real time rather than static outputs.



Most semantic layers handle simple lookups. In the AI age, users ask more complex questions that push systems back to brittle text-to-SQL, where hallucination and distrust live. Conversational analytics doesn't break because of the model. It breaks because the semantics underneath can't answer the question."



Vincent Woon
CEO, Holistics

These platforms don't sell dashboard access. They sell the solution: verified, explainable, and ready to act on.

The Economics of Outcomes

The move from static dashboards to cognitive platforms rewrites the economics of enterprise software. Legacy B2B SaaS relied on a simple formula: more human headcount equals more seat licenses. Autonomous agentic workflows break that equation.

If one AI agent replaces the diagnostic labor of three junior data analysts, why would the enterprise keep paying for traditional visualization seats? Legacy vendors are staring down a deflationary trap.

The market is already reacting. Flexera reports 85% of SaaS leaders are actively pivoting toward hybrid or consumption-based pricing.¹⁶ Separately, Ibbaka's end-of-2025 analysis found that 65% of AI vendors are explicitly layering AI-specific metrics on top of their baseline subscriptions.¹⁷

The math behind enterprise data software is flipping. Legacy vendors charge per seat where the more humans you have, the more you pay. Cognitive platforms charge per outcome so when an agent replaces the work of multiple analysts, the vendor captures that value directly.

The enterprise doesn't need more seats. It needs fewer tickets.

Where We See This Going

Enterprises that buy verified outcomes instead of dashboard seats will have a structural cost advantage. Legacy BI won't disappear overnight because there's too much lock-in, too many procurement cycles already in motion. But its role will shrink. Unified platforms that own both the math and the context will control the execution layer. In a market shifting from software access to verified business outcomes, controlling the execution layer means controlling the margin.

At Activant, we believe the winners in cognitive BI will be the platforms that natively fuse semantic governance and operational context, delivering verified, explainable answers without requiring enterprises to assemble the architecture themselves. We're actively researching and investing in this space. If you're building here, we want to hear from you.

Disclaimer: The information contained herein is provided for informational purposes only and should not be construed as investment advice. The opinions, views, forecasts, performance, estimates, etc. expressed herein are subject to change without notice. Certain statements contained herein reflect the subjective views and opinions of Activant. Past performance is not indicative of future results. No representation is made that any investment will or is likely to achieve its objectives. All investments involve risk and may result in loss. This newsletter does not constitute an offer to sell or a solicitation of an offer to buy any security. Activant does not provide tax or legal advice and you are encouraged to seek the advice of a tax or legal professional regarding your individual circumstances.

This content may not under any circumstances be relied upon when making a decision to invest in any fund or investment, including those managed by Activant. Certain information contained in here has been obtained from third-party sources, including from portfolio companies of funds managed by Activant. While taken from sources believed to be reliable, Activant has not independently verified such information and makes no representations about the current or enduring accuracy of the information or its appropriateness for a given situation.

Activant does not solicit or make its services available to the public. The content provided herein may include information regarding past and/or present portfolio companies or investments managed by Activant, its affiliates and/or personnel. References to specific companies are for illustrative purposes only and do not necessarily reflect Activant investments. It should not be assumed that investments made in the future will have similar characteristics. Please see “full list of investments” at activantcapital.com/companies/ for a full list of investments. Any portfolio companies discussed herein should not be assumed to have been profitable. Certain information herein constitutes “forward-looking statements.” All forward-looking statements represent only the intent and belief of Activant as of the date such statements were made. None of Activant or any of its affiliates (i) assumes any responsibility for the accuracy and completeness of any forward-looking statements or (ii) undertakes any obligation to disseminate any updates or revisions to any forward-looking statement contained herein to reflect any change in their expectation with regard thereto or any change in events, conditions or circumstances on which any such statement is based. Due to various risks and uncertainties, actual events or results may differ materially from those reflected or contemplated in such forward-looking statements.

¹ dbt Labs, [State of Analytics Engineering](#), 2025.

² data.world, [A Benchmark to Understand the Role of Knowledge](#), 2023.

³ MIT NANDA, [The GenAI Divide: State of AI in Business](#), 2025.

⁴ Mordor Intelligence, [Cognitive Analytics Market Report](#), 2026.

⁵ Mordor Intelligence, [Cognitive Analytics Market Report](#), 2026.

⁶ Straits Research, [Business Intelligence Market Overview](#), 2026

⁷ Microsoft, [Q2 FY2026 Earnings Call](#), January 28, 2026.

⁸ Microsoft, [Directions on Microsoft analysis](#), January 2026.

⁹ Microsoft, [Azure Fabric Pricing](#), 2026.

¹⁰ Microsoft, [Q2 FY2026 Earnings Call](#), January 28, 2026.

¹¹ Salesforce, [Q4 FY2026 Earnings](#), February 2026.

¹² Tableau, [Pricing Page](#), 2026.

¹³ Salesforce, [Q4 FY2026 Earnings](#), February, 2026.

¹⁴ Finterra/FinancialContent, [The Agentic Pivot: Decoding Salesforce's Mixed Outlook](#), February 2026.

¹⁵ data.world, 2023 (ibid.).

¹⁶ Flexera, [From Seats to Consumption: Why SaaS Pricing Has Entered Its Hybrid Era](#), 2026.

¹⁷ Ibbaka, [The Evolution of AI Pricing Models: Update for the End of 2025](#), 2025.