coalesce®

REPORT

# The Top Data Trends for 2026

Industry experts share their predictions for the year ahead

# Table of Contents


04 The AI-powered data core will be built on interoperable lakehouses

05 AI market hits hyperspeed, enters correction phase

07 Low-stakes tasks will drive early agentic AI adoption

09 Automation chains will redefine data engineering

11 The true end of the AI hype cycle

12 Data teams will own AI agent infrastructure in 2026

14 A year of post-bubble growth for AI

15 The rise of hybrid data architectures

17 Agents will replace traditional BI dashboards and workflows with actionable insights

18 People will need to learn to think like agents, not just users

19 AI will force a fundamental rethinking of data management

22 In 2026, AI will push the semantic layer into the mainstream

24 AI readiness becomes the new data quality standard

26 Data and platform readiness will define enterprise AI

27 The rise of citizen data practitioners

29 The 'Invisible Data Stack' will make data interaction effortless

31 Business intelligence becomes unified intelligence

32 Managing AI becomes a core skill

33 AI moves from promise to proof

34 Organizations finally get real about data readiness

35 Strong data foundations will quietly drive AI forward

# Contributors

**BHASKAR CHATURVEDI**
Enterprise Data Architect
Accenture Australia

**ARMON PETROSSIAN**
CEO
Coalesce

**SATISH JAYANTHI**
CTO
Coalesce

**GLEB MEZHANSKIY**
CEO
Datafold

**LEE DERKS**
CEO
DigBI Consulting

**DAVID YAFFE**
Co-Founder
Estuary

**ERIK DUFFIELD**
CEO
Hakkōda, an IBM Company

**MATT FLORIAN**
Practice Director - SAP Delivery
Hakkōda, an IBM Company

**TEJAS MANOHAR**
Co-CEO & Co-Founder
Hightouch

**FRANK BELL**
Founder
Data Thought Leader
ITS - Snowflake Solutions

**BARR MOSES**
CEO & Co-Founder
Monte Carlo

**COLIN ZIMA**
CEO & Co-Founder
Omni

**JAY GIMPLE**
Chief Data & Analytics Officer
PanTexAs Deterrence

**MAC NOLAND**
Chief Data Officer & Co-Founder
phData

**JOE REIS**
Author
Data architect
Host of The Joe Reis Show

**CHRIS RÜGE**
Head of Data & BI
RSG Group

**MIKE PALMER**
CEO
Sigma

**JENNIFER BELISSENT**
Principal Data Strategist
Snowflake

**NICHOLAS MANN**
Founder & CEO
Stratos Consulting

**KENT GRAZIANO**
The Data Warrior

**ANDREW CRISP**
VP, Director of Enterprise Data Services
United Community Banks

# AI moves from promise to proof

2026 will be the year of reckoning for AI initiatives. After several years of hype, experimentation, and inflated expectations, organizations will shift from proof-of-concept to proof-of-value. They will realize that AI readiness isn't about having more data, but about the right foundation for their data. The conversation will shift from "what can AI do for our business?" to "how can we actually make AI work?"

With this shift comes both clarity and complexity. AI agents are evolving from concept to reality, automating not just individual tasks but entire workflows. Natural language is becoming the primary interface for data interaction, opening up analytics capabilities to people who have never written a line of code. The semantic layer—once seen as a nice-to-have—is emerging as essential infrastructure to give AI the business context it needs to be truly useful.

Yet alongside this progress comes a renewed focus on the fundamentals: data quality, governance, and modeling, which are proving more critical than ever. The companies pulling ahead aren't those with perfect data—they're the ones building resilience into their systems, learning from production failures, and elevating data teams from pipeline builders to strategic enablers.

The insights shared in this year's report come from data leaders who are navigating this transformation firsthand, hard at work in the trenches of real AI implementation. Their predictions and observations reflect both the promise of what's possible as well as the pragmatism of what actually works. Here's to 2026 as a year of meaningful progress built on strong data foundations.

# The AI-powered data core will be built on interoperable lakehouses

**BHASKAR CHATURVEDI**
Enterprise Data Architect
Accenture Australia

## Open standards will unlock real-time, cross-cloud AI at scale

Let's face it—AI is here to stay, and it's starving for clean, unified data. Legacy silos are a death sentence for innovation, and that's why interoperable lakehouses are stealing the show. By 2026 and beyond, 60% of enterprises will lean on these platforms to merge structured and unstructured data, fueling AI with real-time, cross-cloud access. Open standards like Apache Iceberg and Delta Lake, paired with data fabric architectures, make data sharing effortless, slashing integration headaches and speeding up AI insights by up to 50% compared to 2024. This isn't just a trend—it's the foundation for staying competitive.

To capitalize on this shift, companies need to build lakehouses with open table formats and federated querying, ensuring their architecture is ready to feed AI apps with modular, low-latency data flows.

## The majority of businesses will query data in plain English

Imagine a world where anyone in your company can ask data questions and get instant, spot-on answers—no data scientist required. That's where conversational AI is taking us. By 2026, 70% of businesses will use NLP-driven platforms to let employees query lakehouses and vector databases with plain English, cutting analytics time by up to 60%. This shift is driven by the hunger for faster decisions in cutthroat markets—every second counts, and conversational data delivers. It's not just cool tech; it's a revolution in how enterprises unlock insights.

Success here requires rolling out semantic layers with AI models that turn natural language into real-time, context-rich insights, bridging SaaS and on-prem data seamlessly.

## Agentic AI emerges as the ETL and migration game-changer

Manual data migrations and clunky ETL pipelines are the stuff of nightmares, slowing down modernization and racking up costs. Agentic AI is here to save the day, automating everything from schema mapping to pipeline orchestration. By 2026, over 60% of enterprises will use these smart, collaborative AI agents to modernize legacy systems, slashing errors by 70% and halving migration timelines. With data sources growing more complex, this trend is critical for achieving zero-downtime transitions to cloud-native lakehouses. Ignore it, and you're stuck in the slow lane.

The winning approach is to deploy AI-orchestrated pipelines using AI-powered DataOps to automate migrations with real time monitoring.

# AI market hits hyperspeed, enters correction phase

**ARMON PETROSSIAN**
CEO
Coalesce

## Competition will outpace innovation in 2026

AI is overshadowing everything right now, and agentic models and their impact are dominating the conversation. There's definitely a bubble there, and while it's hard to say if it's going to pop in 2026, there's sure to be a correction. People are starting to realize that expectations were set too high. Think about Gartner's Hype Cycle for Artificial Intelligence report—this will be the "fall from hype."

What I'm really seeing is the rapid commoditization of AI. There are all these AI companies growing rapidly in the application layers. As soon as one of them catches fire, they have twelve new competitors within months, all doing the exact same things. It's basically a market on hyperspeed—one player emerges, quickly becomes the number one player, and then tons of other players emerge based on VC funding. Now they're all running in the same direction at the same time, which sounds a lot like what we first saw with the Modern Data Stack. I think there's a parallel there with how the AI market will unfold.

The question is whether there's a new stack emerging for AI. Everybody wants to believe that, but I think it will be made up of existing companies rather than new ones. I don't see any new data companies emerging that are leveraging AI in unique ways that the incumbents currently don't.

## The M&A fever in the data space will continue as AI pressure mounts

Large companies will continue to realize the importance of having a data strategy and scrambling to make moves into the data space. We saw clear indicators of this in 2025 with Salesforce acquiring Informatica, ServiceNow acquiring data. world, and now Fivetran and dbt merging. For a company like ServiceNow, I don't see the obvious fit, but there was enough market pressure for them to figure it out. AI is driving this to the next level, forcing the market along and accelerating these massive-scale mergers. Some type of existential threat is driving big companies to take their approach to data a lot more seriously than before.

> "The most exciting thing I see coming down the road is agentic workflows in data to help accelerate use cases. I don't think that's happened yet, but it's close and I think we'll get there in 2026. In the future, nontechnical people could call an agent to build out a data engineering use case. This wouldn't eliminate the need for humans—you'd still need a specialized data person to go in, refine the logic, validate it, and make sure it's of high quality."
>
> **ARMON PETROSSIAN**
> Coalesce

## AI will eventually break the stranglehold of legacy ETL tools

Right now, a lot of what AI is doing in data engineering is things like adding column descriptions or table-level descriptions. While this is helpful and shaves off a lot of time, it's not a radical shift or acceleration in development. But agentic workflows promise to help eliminate the massive pain that has always been caused by migrations. There are so many dissatisfied customers using legacy ETL tooling, but there's such a burden to migrate off of them that they've been historically very sticky. My hope is that AI can help people **take the pain out of data migrations** and get to a better framework faster than ever before, and we're already seeing that firsthand.

"One thing I'm looking forward to in 2026 is how AI will make data migrations significantly easier. Migrations have always been painful, so they're a perfect use case for leveraging AI—and it will only get better. As migrations become faster and easier, I expect we'll see data teams moving more decisively toward the best frameworks for scale, accelerating the shift away from legacy tooling to more modern solutions."

**ARMON PETROSSIAN**

Coalesce

"Right now, everybody's got this glitzy, starry-eyed approach to AI and how it's going to impact literally everything.

**But in 2026, after enough experimentation, people will start to realize where AI is effective and where it's not."**

**ARMON PETROSSIAN**

Coalesce

# Low-stakes tasks will drive early agentic AI adoption

SATISH JAYANTHI
CTO
Coalesce

## AI agents will start to handle the simple stuff

AI agents are getting better—they can now perform autonomously for up to an hour without human intervention. If you give it a task like "go debug my program," it will spend an hour figuring things out. That's a drastic improvement from where we started, but the accuracy is still only 50%. Is that something a business can live with? I don't think so. You can't spend an hour on an important task when the answer is only 50% likely to be correct.

In 2026, we'll start to see success with simple agentic AI use cases where the stakes are low, such as telling it, "Go set up an appointment with my barber." If it picks the wrong time, I can easily reschedule. These are the types of use cases where you don't lose anything. But most enterprise decisions require a high level of accuracy, and we're not there yet.

Technology adoption always follows this pattern: it happens in the application space first with simple, narrow tasks where the context is small enough to control and test. The more complex analytics use cases will come later, once AIs can handle the massive amount of context those require.

## We'll soon build data pipelines by just talking to them

The most exciting thing I see coming around the corner is the ability to talk to data tools and build pipelines in a conversational manner—just describe what I want it to build without having to drag and drop anything. I talk to ChatGPT every day, and it's amazing how many things it can do. I want to build pipelines the same way.

This is where everything is headed. We already know that at least 50% of this is possible. Now the question is, can we get to 100%? I'm excited to reach the next level and figure out how we get to 80% or higher.

In analytics, this approach is different from building an app. Data problems are unique: you're managing vast volumes of data, everyone wants something different from it, and you have to move fast. You don't always need 100% accuracy; you just need to know enough to make a decision. If I'm a business leader, getting an answer that's 80–90% accurate right away, versus waiting, is incredibly valuable because I can act on it immediately.

For example, I could say, "Go to my last two years of sales data, use these two sources, join whatever you need to join, and forecast this for the next 36 months with these assumptions." If it gives me a forecasting number that's 80% accurate, that's valuable. I'm building something and using it right away, without needing to be a data engineer. And if I want to dig deeper, I can look under the hood, see the lineage, and understand exactly what it did.

> "Discovery is where AI really helps a lot—you're simply gaining knowledge, quickly learning about something so you can do the real thing later. Healthcare is a perfect example: agents running experiments faster, trying to discover new drugs. That's a perfect use case for AI, and that's where it will shine."
>
> **SATISH JAYANTHI**
> Coalesce

"Data quality has always been important, but it will become even more critical in 2026 as agentic models take hold. When people rely on agents to operate autonomously and make decisions on their own, the quality of the data needs to be superior to anything we've had before.

**Quality, security, and governance may sound like the same priorities we've always had, but with agents in the mix, these fundamentals will matter more than ever."**

**SATISH JAYANTHI**

Coalesce

# Automation chains will redefine data engineering

GLEB MEZHANSKIY
Co-Founder & CEO
Datafold

## Data workflows become self-managing ecosystems

In 2026, we'll see AI agents start to orchestrate these automations together into what I call "automation of automation"—agents that don't just automate individual tasks, but chain together multiple automated processes to handle end-to-end workflows that previously required human decision-making at every junction. For the last decade, we've been automating individual data engineering tasks—write this SQL, run this test, check this quality metric. It's time to go further.

I've seen this firsthand with migrations at Datafold. We're not just using AI to convert SQL from one dialect to another—that's just the starting point. The breakthrough is when the agent can read and map out an entire legacy codebase spanning millions of lines of code, refactor and translate as needed, create a testing environment, run validation, identify discrepancies, adjust the transformation, and iterate until it's production-ready. All without human intervention between steps.

Migrations happened to be the killer use case because they're bounded enough to control, but complex enough to prove the concept. But the same pattern will expand across data engineering in 2026—any workflow that currently requires a data engineer to manually move between tools and make judgment calls at each step is fair game for agent orchestration. Data quality monitoring triggering automated remediation. Schema changes automatically propagating through dependent pipelines. Optimization agents identifying and fixing performance or cost issues.

The companies that figure out how to build and manage these "automation chains" will see 10-100x productivity gains. The companies that keep thinking about AI as a better autocomplete will wonder why they're falling behind.

## AI will elevate data engineers, not eliminate them

I don't think AI is replacing data engineers—I think the job is changing in the same way it changed when we moved from writing MapReduce jobs to using SQL, or from managing on-prem servers to using cloud warehouses. Data engineers will move to higher levels of abstraction. The tedious parts get automated, which frees you up to focus on the parts that actually require human judgment and collaboration.

Automation of automation doesn't reduce the need for data engineers—it actually increases the need for their expertise. When an agent can handle the mechanical work of a migration in weeks instead of months, you need your engineers making higher-level decisions more frequently: reviewing the agent's approach, validating business logic, setting quality standards, and working with stakeholders to understand what success actually looks like.

For many, the immediate skill shift is learning to work effectively with AI—and frankly, most still need to get better at it. But once you master that, the shift is getting comfortable spending time with the people who use your data. When you're not buried in tactical work, you can actually sit with the product team and understand what they're trying to measure. You can talk to executives about what questions they're asking of their dashboards and why the current data model makes those questions hard to answer. You can partner with analysts to understand which data quality issues actually impact decisions versus which ones are just noise.

Honestly, the data engineers who are most excited about what we're building at Datafold are the ones who never wanted to spend six months manually rewriting stored procedures anyway. They want to solve interesting problems and work on things that move the business forward. Automation just gets the grunt work out of the way faster.

## Intelligent automation will dismantle the legacy vendor lock-in model

In many cases, data teams see migration complexity as an unsolvable problem—so they optimize around the constraints of their stale tools. I'd like to see less of that in 2026.

Here's the pattern we need to change: A company runs critical processes on legacy ETL or warehousing solutions. Everyone knows it's holding them back. But when they consider migrating, they get quoted a multi-year, multi-million-dollar billable hours project with a high risk of failure.

So they don't migrate. They hire more legacy technology specialists, build workarounds, and create hybrid architectures that are somehow worse than the old stack. Meanwhile, individual teams spin up shadow IT—Databricks instances, Coalesce projects—because they literally can't do their jobs with the legacy platform. Now you're paying for both the expensive legacy stack nobody wants and dealing with governance issues.

The reality is that migration complexity has been used as a business moat by legacy vendors for years. It's not an accident that many old transformations are defined in barely readable XML, or that the logic is scattered across GUI mappings, shell scripts, and embedded SQL. That complexity keeps customers paying maintenance contracts year after year.

But here's what's changing: The same AI techniques that can understand and generate code can untangle these complex workflows, translate between paradigms, and validate that the new implementation matches the old. What was economically impossible becomes viable at a fraction of that cost. The barrier that kept companies trapped is dissolving. Let's lean into that and bring our industry forward.

"The best data engineers in 2026 won't be the ones who can write the most complex SQL—they'll be the ones who can understand what the business is trying to achieve, design the right data architecture to support it, and leverage AI agents to implement it quickly."

**GLEB MEZHANSKIY**

Datafold

# The true end of the AI hype cycle

**LEE DERKS**
CEO
DigBI Consulting

## AI will finally shift from hype to everyday workplace reality

I think 2026 will be the year when we finally see widespread implementation of AI in the workplace. Both 2023 and 2024 were the hype years about AI's potential, but few organizations had actually adopted it. By 2025, most mainstream SaaS platforms, including Coalesce, started making AI "copilots" available. Nowadays, if there's not an AI feature built into your product, it immediately feels old and outdated. But in the coming year, I believe the next iteration of AI will arrive in the workplace. Organizations will be required to adopt AI governance policies as more employees turn to AI for everyday tasks ranging from coding, process automation, data analysis, and various mundane tasks. That famous line from Star Trek now applies to AI: "We are the Borg. You will be assimilated. Resistance is futile!"

For companies striving to build an AI-ready data foundation, the first step is acceptance of AI. If organizations continue to resist AI due to fear, they will fall behind their peers. I'm not saying organizations should blindly allow AI everywhere, but they need to accept that their employees are eager to use AI to become more productive, and pivot to a strategy that's open to it. Collaborate with data and AI leaders and security experts to develop a comprehensive AI governance and implementation plan. Establish clear guidelines for permissible use of AI across the organization.

## The skills required to be a data pro will continue to evolve

Several years ago, having a "data pro mindset" was about wrangling data silos and unifying the data into centralized data warehouses and datamarts. Now, it's about finding quick solutions to analyzing and storing various types of data, both structured and unstructured. There should be no reason to say, "That type of data is unsupported." Any type of data can be extracted and stored, and having a true data pro mindset means knowing which tools are available to bring life to the data. For example, a user on your team who is not necessarily a pro could use something like Coalesce to build complex data pipelines without writing a single line of SQL. If it gives me a forecasting number that's 80% accurate, that's valuable. I'm building something and using it right away, without needing to be a data engineer. And if I want to dig deeper, I can look under the hood, see the lineage, and understand exactly what it did.

> "AI will not take away human jobs. Yes, AI will take over most simple tasks but it will be up to people to upskill themselves and learn how to partner effectively with AI."
>
> **LEE DERKS**
> DigBI Consulting

# Data teams will own AI agent infrastructure in 2026

**DAVID YAFFE**
Co-Founder
Estuary

## What do you believe will be the most significant data trend of 2026?

Data teams will become increasingly responsible for using AI in consumer-facing applications, such as enabling data discovery, RAG models, and more. This is the next step toward agentic, governed AI on enterprise data becoming the norm. Data teams will begin to be responsible for creating environments where AI agents—not just chat UIs—can plan and execute tasks over trusted, governed data to start powering real business workflows.

## Which AI trend do you anticipate will have the greatest impact on business?

AI agents embedded in applications and operations—from customer support to finance close—will turn "assist" into autonomous execution with human guardrails. Meanwhile, multimodal AI (text, image, audio, and video) will raise the quality of insights and automation across both customer and employee experiences.

## What 2025 trend surprised you most, and do you think it will continue into 2026?

The consolidation of the data stack. It seemed somewhat unnatural for unrelated tools to start seeing consolidation (for example, the EL and T coming together with Fivetran and dbt). This trend will inevitably continue, as the only option to build bigger businesses and extract more revenue from customers is by consolidating data use cases. The only thing that could change this would be a paradigm shift caused by AI.

## What's something you hope to see less of in 2026?

One-off AI demos with no data governance or ROI plan—it's time to ship fewer toys and more products.

## What's the biggest misconception about the data and AI industry you hope gets put to rest in 2026?

That "AI is just a model." In reality, value comes from data foundations, retrieval (RAG), orchestration, and controls around the model. The practitioner who creates the "model" will end up being crucial in the same way a data engineer is today—it may even become the same person as tools evolve.

## What do you see coming around the corner that's most exciting to you?

Composable agent teams capable of handling multi-step business processes end-to-end, auditable by default and with the proper human guardrails understood and maintained.

Also, richer "RAG 2.0" systems—knowledge graph-aware, multimodal, and event-driven—that safely tie real-time data into decisions-making.

## Which industry will see the highest disruption from AI in 2026, and why?

I think it's less about entire industries and more about specific professions. Industries will see new opportunities, while some professions may become less essential:

- Legal roles will change as paralegal work becomes more commoditized.
- Customer service and sales operations across sectors (horizontal functions) will likely shift the fastest as agentic workflows hit CRM, support, and field operations.

## What's the most critical element of building an AI-ready data foundation?

A trustworthy, highly structured retrieval layer built on governed data—lineage, quality, and PII controls—so AI can cite, explain, and stay current, all while keeping the correct data tokenized for a particular task.

**Biggest gap:** Lack of standardization in a quickly evolving industry prevents users from knowing what to do. For example, MCP (model context protocol) has been around for about six months, but Anthropic has already come out with a new standard. Staying relevant will be nearly impossible for a while.

## What does it mean to have a true "data pro mindset" in 2026?

Foundations are just as important as ever. Enable teams in a governance-first way—with reusable, policy-aware data and AI components—so every team can ship safely and quickly on a strong, accurate data foundation.

**Looking into 2026, how should organizations think about data modernization and legacy migrations?**

It's not just about "moving to the cloud." Data practices should be designed for your current and future use cases. Know where you're headed, and build the right foundations that can manage it. For example, you may need event-driven, real-time, and retrieval-friendly patterns that feed AI.

Design for regulation and efficiency from day one, with lineage, access control, and cost observability to stay compliant with AI and data localization rules in 2026 and beyond.

# A year of post-bubble growth for AI

**ERIK DUFFIELD**
CEO, Hakkōda an IBM Company

## AI market shakeout

I would not be surprised if, given how 'hot and frothy' the AI data space has become, there's a bit of a shakeout about the valuations of AI companies in 2026. While bubbles can be fantastic and exciting, a shakeout isn't necessarily a bad sign. Obviously, if you're concentrated in that space, it will hurt for a little bit. But it's the natural cycle of the market: things get hyper-inflated, then they ease up a bit. That means we're moving from one phase to the next.

Even though we're seeing our clients truly gain efficiency and productivity thanks to AI, we may still go through that public market cycle with AI software and foundational model companies—because at some point, the tea kettle has to blow off some steam.

As for the actual implementation of AI on the client side, we're moving into a new phase I think is really exciting: budgets and spending will increase, our clients will have better earnings calls, and their stock prices will go up. It feels like 2026 could be the year when we start to see this play out, and we shouldn't be afraid to ride through that. If you're willing to be part of the "bubble," you also have to be ready for the "post-bubble."

## Quantum computing on the horizon

Now that Hakkōda is part of IBM, I'm thinking more and more about quantum computing. I love the boldness IBM has shown around quantum computing—especially its plans to demonstrate fault-tolerant quantum as soon as next year. With a broader release in the 2029–2030 timeframe, we're only a few years away, and in 2026 IBM (and maybe others) will start showing proof points. It's already being tested in some fascinating case studies.

When I think about AI and data readiness, I love that engineers are innovating from the silicon up. By 2026, with all the capital invested—venture, private equity, and corporate—we'll see broader technology innovation: AI that builds on today's momentum and accelerates beyond it. Quantum is definitely something to keep an eye on.

> "Whenever the shakeout comes, the market will rebuild. AI isn't going away—getting your data ready for machine-driven consumption is the foundation of AI readiness, and that momentum will only accelerate."
>
> **ERIK DUFFIELD**
> Hakkōda, an IBM Company

# The rise of hybrid data architectures

**MATT FLORIAN**
Practice Director
SAP Delivery
Hakkōda, an IBM Company

## Hybrid architectures will serve as bridge to AI maturity

Data architecture still matters, but the dogma around it is slowing us down. I see it all the time in the field: people become emotionally invested in architecture. You've got the Kimball loyalists on one side, the Data Vault evangelists on the other, and a few one-big-table rebels in the corner. And here's the thing—they're both right and wrong.

The Modern Data Stack has matured to a point where we don't have to pick a single camp anymore. The real trend we'll see in 2026 will be the rise of hybrid data architectures—architectures that combine the best of each approach instead of forcing false choices. It's not about abandoning what came before; it's about integrating it intelligently.

Hybrid architectures are also the bridge to AI maturity. They connect our structured data to semantic meaning, which is the foundation for knowledge graphs and LLMs that don't just retrieve answers but reason over them. That's where the value lies: when your data can explain itself.

In short, the hybrid architecture elevates the medallion framework from being about data to being a hierarchy of understanding.

## The semantic layer becomes the AI layer

The AI trend that's going to have the biggest impact on business isn't the next big model—it's the rise of the semantic layer as the new AI layer.

For years, the semantic layer has been treated like a reporting convenience, something to keep KPIs consistent across dashboards. But now it's becoming the thing that gives data real meaning. As companies move deeper into AI, that shared business language is what lets analytics tools, search interfaces, and AI agents all "speak the same language." That's how you get answers that aren't just accurate, but explainable.

We're already seeing it happen. SAP's Business Data Cloud can share business context directly into Snowflake, and platforms like Coalesce can refine and extend that meaning before it's used by BI tools or AI systems. Suddenly, the same definition of "revenue" or "customer" carries through every layer—from ERP to the semantic layer to the AI agent that answers your question.

That's when the medallion framework stops being about data movement and starts being about understanding. The hybrid architecture evolves into a hierarchy of understanding.

## Interoperability will power the next wave of intelligence

What surprised me most in 2025 was how fast "interoperability" stopped being a buzzword and actually started showing up in production.

For years, we've talked about open ecosystems like they were a distant goal. But this year, we saw the walls finally come down. Snowflake launched OpenFlow and the Open Semantic Interchange to let data and meaning flow freely across tools. SAP opened up its Business Data Cloud to work seamlessly with Databricks and Google BigQuery. Suddenly, the data stack started acting like … well, a stack.

The other big surprise was how quickly agentic AI started to move from concept to reality. It's no longer just something being demoed at conferences—organizations are experimenting with agents that can actually perform tasks across systems, using real business data, not sample sets.

And what ties those two together is interoperability. Agents can only act intelligently if the data underneath them can move and be understood across systems. That's why this trend isn't slowing down. In 2026, we'll see interoperability mature from "integration" to true collaboration between tools, teams, and AI systems.

"Legacy and modern players alike are recognizing that openness is the key to unlocking enterprise-scale AI. Interoperability is the quiet revolution that will make AI real."

**MATT FLORIAN**

Hakkōda, an IBM Company

# Agents will replace traditional BI dashboards and workflows with actionable insights

**TEJAS MANOHAR**
Co-CEO & Co-Founder
Hightouch

"Agents that connect to your data warehouse and other sources will replace a lot of standard BI and reporting processes. Instead of just serving customers a dashboard that's refreshed on a schedule, agents can answer questions like 'What ads are performing worse this week and why?' and with the right context and guidance, start to answer questions like 'What are my top opportunities this week to increase sales?"

**TEJAS MANOHAR**
Hightouch

# People will need to learn to think like agents, not just users

**FRANK BELL**
Founder
Data Thought Leader
ITS - Snowflake Solutions

## 2026 will be the year of agentic thinking

I foresee developments in the next phase of "Automation of Automation" in 2026, where data sets will become increasingly self-discoverable and self-analyzing. That said, I think 2026 will mark an even greater data divide for many companies. While some companies are achieving amazing results with their AI initiatives, others are still not seeing the results they expected. This is often related to data preparation, data quality, and overall data strategy. But let's not forget that, even in the era of generative AI, true change takes time—it requires companies and individuals to think differently and not follow the standard practices that people have been trained to follow for years.

To succeed in the future, people will need to learn how to think agentically, a huge mindset shift that hasn't been addressed enough, especially when it comes to data preparation, cataloging, and near real-time quality assessment.

That near real-time quality assessment is also fundamental to building an AI-ready data foundation. Using near real-time data for AI analysis means finding any major quality issues immediately, before the data is fed into AI agents and operations. "Good enough" data is no longer acceptable as input to AI automations.

## AI is moving fast and getting real

It probably shouldn't surprise me anymore, but the frequency and quality of new generative AI capabilities—especially across major LLMs from OpenAI, Anthropic, Google Gemini, and Grok—have turned into an almost weekly competition to outdo one another. The pace of change is unlike anything we've seen before.

One thing I hope we see less of in 2026 is the endless wave of AI charlatans claiming to "change everything," yet offering merely a thin front end to an LLM without adding real additional value. But one exciting thing I think we can look forward to is more and more AI in robotics.

> "Data professionals must understand more than just data to be truly effective—especially with platforms like Snowflake and Databricks. These platforms go far beyond traditional data warehousing, embedding AI/ML capabilities directly alongside the data. With the introduction of Cortex AI, you could argue that Snowflake is evolving SQL itself into SQL-for-AI."
>
> **FRANK BELL**
> ITS

# AI will force a fundamental rethinking of data management

**BARR MOSES**
CEO & Co-Founder
Monte Carlo

## Data and AI will be managed as one unified system

The most significant trend in 2026 will be the death of data or AI as disparate functions, and the rise of unified data + AI observability.

Here's why: AI has fundamentally changed what reliability means. You can't just monitor your data pipelines anymore and call it a day. When an agent fails, the root cause could be anywhere: corrupted upstream data, stale embeddings in your vector database, model drift, or the agent itself. Traditional data observability illuminates part of that picture, but in the complex, interdependent environment of AI systems, it falls short.

The teams winning with AI in 2026 won't be the ones treating data and AI as separate disciplines. They'll be the ones with a single pane of glass across the entire stack, from source data to model output, with unified workflows that catch and fix problems fast. Because the reality is that AI requires data engineers to build and maintain its applications, and data engineers are increasingly responsible for how AI is layered into the enterprise. They're not two separate problems anymore. They're one system that needs to be managed as one.

## Context engineering will emerge as a critical AI discipline

The AI trend I anticipate will have the greatest impact on businesses will be context engineering becoming a core discipline. Here's why it matters: input costs for AI models are roughly 300–400x larger than outputs. If your context data is shackled with problems—incomplete metadata, unstripped HTML, empty vector arrays—you're facing massive cost overruns while processing at scale. And that's just the financial impact.

Confused or incomplete context is also a major reliability issue. Ambiguous product names and poor chunking confuse retrievers, while small changes to prompts or models lead to dramatically different outputs. We're seeing teams realize that the quality of what goes into their AI systems matters just as much as the models themselves.

The businesses that will win in 2026 are the ones mastering upstream context monitoring, ensuring a reliable corpus and embeddings before they hit expensive processing jobs. This means monitoring unstructured data (documents, emails, images) with the same rigor they apply to structured data. It's not sexy, but it's the difference between AI that delivers value and AI that burns budget.

## The AI enthusiasm gap widens into a competitive reckoning

What I found surprising in 2025 was the AI enthusiasm gap. MIT's report said it all: **95% of AI pilots are failing**, and the blame rests partly with executives who treat AI like magic. We're seeing this pattern on repeat: an executive who doesn't understand the technology sets the priority, the project fails to provide value, the pilot gets scrapped, rinse and repeat.

What surprised me was how widespread this became. Companies are spending billions on AI with no clear understanding of where or how it will drive impact. And yes, I think this continues into 2026 to an extent, but with a twist. The enthusiasm gap will widen into a reckoning. Companies that figure out how to get to value by empowering data + AI teams who understand both the technology and the business problems will pull dramatically ahead. Everyone else will be stuck in pilot purgatory.

## The next phase of AI maturity is resilience, not readiness

One thing I'd like to see less of in 2026 is perfectionism paralysis. Here's what I keep hearing from CDOs: "We're not AI-ready yet." They want perfect data, complete governance, the whole nine yards before deploying anything. But that readiness they're chasing? It doesn't exist.

I've watched teams spend budget cycle after budget cycle cataloging every asset and cleaning data sets across business units. Meanwhile, the data keeps changing. APIs get updated. Customer behaviors shift. By the time they declare themselves "ready," the landscape has already moved beneath their feet. And here's the kicker: they've missed out on the most valuable learning opportunity they'll ever get—actual production failures.

"Perfectionism is killing more AI initiatives than bad data ever will. The companies winning with AI aren't the ones with perfect data. They're the ones moving fast, learning from production, and building the observability and guardrails to respond quickly when things break. In 2026, I want to see less preparation theater and more production courage. Stop asking 'Are we AI-ready?' Start asking 'Do we have the right tooling and process to adapt?"

**BARR MOSES**
Monte Carlo

## Data engineers move from pipeline builders to AI enablers

Here's the biggest misconception I want to see corrected in 2026: that AI will replace data engineers. Everyone's talking about it, but I think the opposite is true—AI is going to make data engineering more essential, not less.

Here's why: AI will absolutely automate tasks. But automation doesn't eliminate the need for humans. It elevates what humans do. When our Troubleshooting Agent explores thousands of hypotheses to identify why an issue occurred, you'd think that mitigates the need for humans in the loop. On the contrary, data + AI folks are still crucial to actually investigate, triage, and execute on what the agent presents.

As AI takes over the tedious, time-intensive workflows, humans will spend more time on higher-order problems: reliability, governance, strategy, and designing systems that both humans and AI can trust. The skills that matter (abstract thinking, business understanding, contextual creation) become more valuable, not less. My advice to anyone in the field: learn how to work with AI, not compete with it. Understand how to validate outputs and design trustworthy systems. That's the future.

## Operational agility matters more than perfect data

To build an AI-ready data foundation, the most critical element isn't perfect data quality—it's continuous operational readiness. And that's where most organizations are falling short.

Teams are treating AI readiness like a checklist: clean the data, catalog the assets, build governance frameworks, then deploy. But your test environment will never replicate production. Those edge cases that break your model? They only show up when real users do real things. By the time you've checked all the boxes, the data has already changed.

The organizations at risk are the ones paralyzed by preparation. They're missing the most valuable learning opportunity they'll ever get: actual production failures. What works is shifting from "prepare then deploy" to "prepare and deploy"—building observability into the entire stack so you can detect, triage, resolve, and measure issues fast.

The real question isn't "Is our data AI-ready?" It's "Can we catch and fix problems quickly when they inevitably occur?" Build the response system, not just the perfect system. That's what production-ready actually means.

# In 2026, AI will push the semantic layer into the mainstream

**COLIN ZIMA**
CEO & Co-Founder
Omni

## Semantic layers will go from nice-to-have to non-negotiable

It's already begun, but over the next year, AI will push the semantic layer into the mainstream. For years, the data community debated whether defining consistent metrics and structure was worth the effort. Now, AI has made the answer obvious. You can't get reliable results by just pointing an LLM at your data. AI needs to understand your business context for it to be useful. That means your data needs to reflect how your organization actually talks and operates.

With traditional BI, messy or inconsistent models caused confusion, but there were buffers. Analysts could intervene. Dashboards limited scope. A bad result might be frustrating, but rarely damaging. In an AI workflow, you need to find new ways to build in guardrails. It's not just about building dashboards anymore; AI is querying your data in natural language and its results impact decisions. That only works when it's grounded in a consistent, shared understanding of your business.

The semantic layer is no longer a nice-to-have. It's the foundation that makes AI work.

## Business context matters more than ever for data quality

One of the most critical elements of an AI-ready data foundation is business context. You can have all the infrastructure in place, but if your AI doesn't understand how your business defines something like "active customer," it's going to guess. That's where most organizations fall short. They focus on exposing data to AI instead of structuring it for AI.

But quality doesn't come from just access, it comes from context. You need to teach AI how your business actually talks and operates, and that structure needs to be accessible to both the AI and the people using it, so they can trust and verify the results. Building an AI-ready foundation isn't just about connecting tools; it's about creating a shared understanding, so AI can answer questions the way a smart analyst who understands your business would.

"AI is powerful and can replace a lot of manual work, but it's not magic yet. It's fast, tireless, and great at pattern matching, but the logical limitations and nuance of data still leak into nearly every analysis. To get useful output, you need to invest in structure, teach it how your business talks, and refine it through real use.

**When you treat AI like a robot and design for its strengths, it stops being a novelty and starts becoming a force multiplier—but it won't be running our businesses in 2026."**

**COLIN ZIMA**
Omni

# AI readiness becomes the new data quality standard

**JAY GIMPLE**
Chief Data & Analytics Officer
PanTeXas Deterrence

## Data architectures evolve for the AI era

The most significant data trend of 2026 will be the shift to AI readiness as a core architectural mandate. This isn't just about launching a few AI pilots; it's the realization that data programs must pivot entirely to support the executive focus on AI. We'll see organizations aggressively re-architecting to ensure data is immediately governed, standardized, and available in real-time—because, simply put, AI models fail without it. The relevance of every data initiative hinges on building and demonstrating AI readiness, which has become the ultimate measure of data quality and speed.

## Causal AI ushers In a new era of enterprise intelligence

What I'm most excited about is the rise of causal AI, and the corresponding shift from predictive systems to prescriptive reasoning engines. Right now, almost all of our data and AI efforts are focused on correlation: figuring out what will happen (prediction). Causal AI, however, is about figuring out the why behind events, allowing systems to understand cause-and-effect relationships. This is a game-changer because it moves our data strategy from "what-if" modeling to "if we do X then Y will happen" reasoning.

This shift will enable true, trustworthy AI-driven decision-making. Imagine a system that doesn't just predict customer churn, but accurately identifies the specific sequence of actions that caused the churn, and then automatically recommends the best intervention to reverse the cause. It brings true intelligence—the ability to reason—to data at enterprise scale, which is hugely exciting.

## Leaving brittle, monolithic systems behind

Looking into 2026, when we think about data modernization and legacy migrations, it's absolutely about rethinking data pipelines from the ground up, not just shifting infrastructure. Legacy migrations must be viewed as a re-architecting opportunity. The goal shouldn't be to just move to the cloud for cost savings, but to build for agility, scale, and real-time processing. This means intentionally moving away from brittle, monolithic systems to adopting event-driven architectures and focusing on creating highly governed, self-service data products. This foundational shift is what's necessary to deliver the speed and quality that modern AI demands.

> "In 2026 I hope we see less focus on vanity metrics and more focus on data governance and lineage. Too many teams are still prioritizing the volume of dashboards or models over the trustworthiness of the data feeding them."
>
> **JAY GIMPLE**
> PanTeXas Deterrence

## AI will deliver the systemic fixes healthcare has long needed

I believe the healthcare industry will experience the most significant and immediate disruption from AI in 2026, driven by a pressing need to reduce waste and improve operational efficiency. The industry's deep-rooted administrative complexity—think of the ongoing struggle with claims, billing, and regulatory paperwork—is something AI is uniquely suited to handle. This disruption is accelerating because the underlying financial pressures are so intense that they demand a systemic fix. While executives are initially focused on efficiency gains, the real impact will be felt by consumers. My hope is that using AI to eliminate this entrenched waste will ultimately lead to lower overall healthcare costs, making care genuinely more affordable for everyone.

"One of the biggest misconceptions is that AI is a silver bullet that works with any data. The reality is that AI is an amplifier—it will only amplify the quality and bias of the data you feed it. As data leaders, we need to diligently retire the myth that the tool matters more than the foundation."

**JAY GIMPLE**

PanTeXas Deterrence

# Data and platform readiness will define enterprise AI

**MAC NOLAND**
Chief Data Officer &
Co-Founder
phData

## The AI edge starts with clean data

In 2026, we expect to see an acceleration in AI-ready data and platforms. Throughout 2025, many AI projects were hindered by the lack of available, usable data and unready platforms. Organizations have realized that AI's ROI has been limited by the need to ensure their data is both available and well organized. To fully maximize AI's impact within their operations, companies must have data curated to a usable state, along with platforms that can operate at enterprise scale.

The message is clear: the foundation of successful AI is having quality data and a strong enterprise platform. If you're not prioritizing those right now, you're at risk of falling significantly behind and being disrupted—but by focusing on these foundational pieces, you can position yourself for success and safeguard against disruption.

## Migration without modernization will stall progress

Looking into 2026, organizations need to understand that data modernization and legacy migrations aren't just about moving to the cloud—getting to the cloud is only the first step. Lifting and shifting a poorly designed system and declaring victory only gets you a subpar system running on a different platform. There must be an investment in the second step of modernizing the applications and data layers. This ensures you can capitalize on advances in technology and prepares you to support hundreds, thousands, and eventually millions of AI agents that will use your platform and data to automate business problem-solving.

Your migration plan should be step 1, followed by step 2. Once you've done that, take a breath and celebrate, but then get back to work and continue building step 3 and beyond.

> "The biggest misconception about data and AI these days is that this is all easy stuff, which I think gets said too cavalierly. Yes, some things are becoming easier, as they often do in technology, but building robust and scalable platforms, data layers, and AI solutions still requires highly skilled engineers and architects."
>
> **MAC NOLAND**
> phData

# The rise of citizen data practitioners

**JOE REIS**
Author, Data Architect
Host of
The Joe Reis Show

## What do you believe will be the most significant data trend of 2026?

The easy answer is all things agents, context, and memory. These things are obviously happening. What surprises me most is how the rise of "citizen data X" (X = analyst, data engineer, data scientist, etc.) is happening right under our noses. There's a whole generation of non-technical and non-data people using AI to get answers and build automations and apps. This is insanely cool … and slightly scary. And all of this is happening with or without proper data teams.

## Which AI trend do you anticipate will have the greatest impact on business?

A recent **MIT study found that 95% of generative AI efforts fail**, citing memory and context as the biggest roadblocks. I expect these issues to receive significant focus and improvement, which will unlock many new opportunities for businesses.

## What 2025 trend surprised you most, and do you think it will continue into 2026?

Vibe coding. End of story. I think it will continue in some fashion, largely because of shadow AI and its obvious utility.

## What's something you hope to see less of in 2026?

Vendors bolting AI onto products instead of reimagining what AI could look like from the ground up if they started over.

## What's the biggest misconception about the data and AI industry you hope gets put to rest in 2026?

The notion that "data needs to be AI-ready." Maybe so, but maybe not since agents don't necessarily need data to function.

## What do you see coming around the corner that's most exciting to you?

AI overlords helping strategize and build legit businesses and workflows in record time.

## Which industry will see the highest disruption from AI in 2026, and why?

It's more like, who do I see not getting disrupted as quickly—people like cattle ranchers, plumbers, and electricians. Everyone else will face big disruption in 2026 and beyond, even if they don't realize it yet.

## What's the most critical element of building an AI-ready data foundation?

Data modeling and semantics are still a massive challenge. Context won't happen until "the business" and "data" align for real. Same problem as ever.

## What does it mean to have a true "data pro mindset" in 2026?

A willingness to experiment with AI, make mistakes, and improve—this is how you win in 2026 and beyond.

**Looking into 2026, how should organizations think about data modernization and legacy migrations?**

The reset button has been hit. Everyone is back to the starting line and figuring out their next moves. The moves of yesterday might work, but definitely rethink **EVERYTHING**!

# The 'Invisible Data Stack' will make data interaction effortless

**CHRIS RÜGE**
Head of Data & BI
RSG Group

## Users will work via seamless, intuitive interfaces

The most significant data trend in 2026 will be the emergence of what I call the "Invisible Data Stack" that powers composable, intelligent experience layers. We're witnessing a fundamental shift where data interaction becomes deeply embedded and conversational, woven directly into the tools that business users use every day, rather than requiring them to log in to separate analytics platforms.

At RSG Group, where we operate fitness brands across multiple countries, I'm seeing this transformation accelerate beyond the tech industry. What was once considered a "tech company problem" has become table stakes for running gyms. We're no longer asking "should we invest in data infrastructure?" but rather, "how do we architect our member experience as a data product?"

The breakthrough is happening through what I observe as "unified UIs with interchangeable backends." Users will interact through seamless, intuitive interfaces while multiple specialized vendors handle specific workloads underneath. An operations manager won't log in to a business intelligence tool anymore; they'll ask our proprietary app, "Which classes have the lowest attendance on Tuesdays, and what member segments should we target with promotions?" In the background, a seamless orchestration of best-of-breed tools for ingestion, transformation, marketing automation, and AI will deliver the answer.

This matters because it solves two critical problems simultaneously: it eliminates vendor lock-in anxiety while dramatically reducing complexity for business users. Your marketing team doesn't need to learn five different platforms—they learn one interface that orchestrates your cloud data warehouse, marketing automation tools, visualization platforms, and whatever comes next. The data layer becomes the constant; the tools become variables.

This trend makes a "data first" culture finally attainable for traditionally non-tech industries because it meets users where they are, speaking their language. For fitness specifically, this means our frontline staff and members interact with seamless experiences while we maintain the flexibility to swap analytics engines, change CRM platforms, or adopt new AI capabilities without re-training our entire organization.

> "A simple 'lift and shift' of legacy logic to the cloud is a missed opportunity that often just moves problems to a more expensive environment. Real modernization means rethinking your entire data philosophy: how data is modeled, how teams access it, how governance is enforced, and how systems interoperate."
>
> **CHRIS RÜGE**
> RSG Group

## Operational and analytical data will finally converge

Looking ahead, what excites me most is the possibility of true data interoperability—the complete convergence of operational and analytical data architectures. I see the potential for real-time, unified data experiences where the boundary between transactional systems and analytical systems disappears entirely. The ultimate promise is a future where data is decoupled from compute, allowing us to leverage multiple query engines on a single, unified storage layer without vendor lock-in.

We're seeing glimpses with dynamic tables in cloud data warehouses, live tables in unified analytics platforms, and emerging streaming-first architectures. But the full vision, where every employee has real-time access to operational truth without impacting production systems, and where you can seamlessly switch between platforms to use the best tool for the job, is still a few years out. When we get there, the distinction between "data engineers" and "application developers" might finally blur in productive ways.

"Data professionals are no longer a service department—they are the central nervous system enabling the entire organization to think and act with data."

**CHRIS RÜGE**
RSG Group

"AI elevates the data team from builders of pipelines to architects of knowledge. As AI automates the creation of basic code and reports, the role of a data professional becomes more strategic. They become the ones who curate, model, and govern the trusted data foundation that the entire organization's AI ambitions rely on."

**CHRIS RÜGE**
RSG Group

# Business intelligence becomes unified intelligence

MIKE PALMER
CEO
Sigma

## AI will unify the data stack

The biggest trend in 2026 will be the collapse of boundaries between analytics, AI, and applications. The term "business intelligence" has reached the end of its useful life. AI is forcing convergence across every layer of the stack. What used to be separate tools are merging into unified, end-to-end workflows. The winners will be platforms that integrate schema, logic, and user experience into a single system.

## Seamless data will power the intelligent enterprise

One exciting thing I see coming around the corner is the rise of an open, interoperable data and AI ecosystem. Definitions, semantics, and logic have been locked inside tools for too long. Standards like the Open Semantic Interchange are breaking that open. When meaning becomes portable, collaboration scales. That will unlock connected applications and shared intelligence across the enterprise.

## Data modernization will be a way of working, not a project

Looking ahead to 2026, organizations shouldn't think of data modernization as something that ends when you migrate your legacy systems into the cloud. Rather, data modernization is about redesigning how data, applications, and people interact. While the cloud provides the foundation to do that work—after all, you need a cloud warehouse to move fast, operate at scale, and leverage the latest innovations—data modernization is really about unifying workflows around a cloud-centric system.

> "The idea that AI will replace data teams is dead wrong. AI will make them more critical. Prompting your way into an answer without understanding the data model is a recipe for failure. The need for expertise isn't going away—it's widening."
>
> **MIKE PALMER**
> Sigma

# Managing AI becomes a core skill

JENNIFER BELISSENT
Principal
Data Strategist
Snowflake

## AI agents for all

As agentic AI becomes more common, personal agents won't just be for actors or athletes anymore—everyone will use them. We're already comfortable with AI assistants, and over time many of these assistants will become more agentic, able to function autonomously and work collaboratively on complex tasks that we assign to them.

A recent Capgemini study on harnessing the value of gen AI showed a steady planned increase in the use of agents by large enterprises: 10% currently use them in some form, while more than 50% plan to use them in the next year, jumping to 82% in the next three years. There are high expectations for this emerging technology.

## Management in the agentic age

The rise of AI has sparked fears about what it means for our jobs. At the same time, we're also frequently being told that AI itself won't replace us, but someone who knows how to use it well just might. This includes people who understand how to effectively manage these emerging "agentic colleagues."

The science of managing human beings has been around for over a century. That gives us a head start, but the moment has come to sharpen our skills for an agentic future.

Social scientists have long studied task delegation through principal-agent theory. A principal delegates tasks to an agent; however, the agent may not act as expected or desired due to uneven information, misaligned incentives, or conflicting interests. There are management best practices to address these challenges, and fortunately, many also apply to agentic AI—with some new tools to guide the process. These best practices are an integral part of ensuring responsible AI.

In the year ahead, companies will need to train managers to become true principals to ensure autonomous AI agents deliver desired outcomes and subsequent business value.

### Guidelines for managing agentic agents

- **Define the role with an explicit "job description,"** including responsibilities, autonomy, and seniority.
- **Identify the right AI model candidate**, whether bought or built, based on capabilities and fit.
- **Onboard agents with appropriate training data**, ensuring access to contextual data and prioritization when requirements conflict.
- **Establish performance metrics** that clearly define what success looks like and how it will be measured.
- **Monitor performance and enforce rules**, applying mitigation strategies when anomalies arise.
- **Empower agents to govern themselves and peers**, using feedback loops and escalation when needed.

# AI moves from promise to proof

**NICHOLAS MANN**
Founder & CEO
Stratos

## In 2026, AI will have to prove its value

New AI projects in 2026 will require clear business cases and ROI before moving forward. Proof-of-concept failures will lead to skepticism, so it will be critical to show value at a low cost before committing to larger initiatives. Quality data will also be at the forefront for successful AI projects because of all of this year's AI failures and hype. People are starting to see the lack of results despite all the money being thrown at AI projects.

I think this will change in 2026 as companies tighten their budgets. Personally, I'd like to see less AI hype in the year ahead. Yes, AI is the future, but it isn't a solution for all business problems.

## Efficiency will redefine talent and disrupt industries

Leaner data and analytics teams will continue using a variety of modern technology (with and without AI) to do their jobs more effectively. Those people who can learn new tools quickly, communicate effectively with the business, and use tech to be more productive—working more efficiently as an individual than several people combined—will command the highest salaries.

I believe the manufacturing industry is poised to see the biggest disruption from AI in 2026. More and more robotics companies are starting to replace warehouse workers, so that could be a huge disruptor to the entire sector.

> "Generative AI solutions that accurately and quickly return data—work that previously took hours to cobble together in Excel or via other methods—will have the greatest impact on businesses in 2026."
>
> **NICHOLAS MANN**
> Stratos

# Organizations finally get real about data readiness

KENT GRAZIANO
The Data Warrior

## The AI era reveals that data discipline still matters

One big data trend we'll see in 2026 is organizations finally starting to realize that getting data "AI ready" is much harder than imagined. Folks think that AI and agentic AI are going to automatically clean up their data and immediately transform it into useful information and insights. But as organizations get more serious about rolling out AI-assisted programs and processes into production, they are finding out just how bad the data they want to use really is.

Back in the Hadoop era, many believed data modeling was no longer needed (e.g., schema-on-read) and data warehousing was dead. Now in 2026, organizations need to wake up to the reality that "garbage in, garbage out" is still true, that much of their data is not ready for use, and that it presents a major risk to the success of their AI initiatives. That means organizations will have to invest heavily in revamping their data ecosystem to include better data modeling, QA, and governance, and follow the basic data management practices that have been around for decades.

While some started to realize this in 2025, it will accelerate in 2026—there is no free lunch.

## Modernization becomes the key to surviving constant change

Organizations thinking about a data migration need to begin by asking two questions: Does this really need to be migrated and modernized, or can we just mothball it? What's the real value of the system or systems in question?

If the answer is simply "we need it" and there is proven business value, then they should think about rearchitecting and re-engineering the system and processes to use modern tools, platforms, and concepts. This will ensure the longevity and future extensibility of both the system and the data.

Things will continue to change at an even faster pace, so systems must be modernized to make it easier for them to evolve in the years ahead.

> "Data modeling, data quality, and data governance are fundamental to being able to use and trust the data consumed by AI. The risk is twofold—organizations aren't willing to truly invest in these areas, and there are fewer professionals available to guide and implement these programs due to an aging workforce and the lack of emphasis on these disciplines in most universities."
>
> **KENT GRAZIANO**
> The Data Warrior

# Strong data foundations will quietly drive AI forward

**ANDREW CRISP**
VP, Director of Enterprise Data Services
United Community Banks

## Efforts to clean, govern, and monitor data will outpace flashier AI trends

Data quality will remain a foundational focus in 2026, with more efforts to manage these platforms and programs. It's not the newest or most exciting trend, but it's the backbone of what's directly in the spotlight. There's no shortage of data today, but more data isn't what we need anymore. What users and developers leveraging modern tools and AI need is clean, reliable data that inspires confidence—otherwise your AI roadmap is just building the shiniest ship at the bottom of the ocean.

Recent studies and surveys, including one from the **Business Application Research Center** (BARC), report that data quality remains a top priority. Data leaders are clamoring for a more streamlined and manageable approach, and we're seeing product teams respond, integrating more functionality into catalogs and observability tools.

## Conversational analytics will redefine how we interact with data

Looking ahead, what excites me most is the rise of conversational analytics. It has a real opportunity to change how we've thought about reports, visualizations, and dashboards over the past decade or more. We're finding opportunities where analysts no longer have to anticipate the five or ten right questions, or spend countless hours in requirements meetings to ensure business data is adequately translated before an engineer spends hours building a dashboard with a few filters.

As we learn more about these semantic views and how to weave them into unstructured data, we're entering a moment when more comprehensive insights can be found by simply talking to the data directly from the platform itself.

> "I hope 2026 will be the year of effective, intentional AI. I'm seeing a dangerous trend where users are abdicating their personal responsibility to produce critical thought in favor of the convenience of a prompt. We've become incredibly eager to delegate deep thinking to large language models—and it's starting to show. People notice the lack of real effort to contribute to the output of a generative AI chatbot before something gets shared. As the hype fades and we settle into our new normal, I hope we defer less critical thinking to AI, and start leveraging it more as a tool to enhance the creativity and brilliance that we have in abundance in our industry."
>
> **ANDREW CRISP**
> United Community Banks

"It's finally time to retire the myth that more data equals better decisions. There's no shortage of data today, yet many still believe that with just a bit more, the answer they seek is going to jump out at them. I'd argue the opposite is true. If not well managed, the sheer volume of data that teams and companies now have access to can create confusion with dashboard bloat, nuanced differences in business definitions and rule sets, and a sprawl of data products that leave users overwhelmed and confused about where to find what they need.

**In 2026, I hope we'll see a shift in the idea that more is better, and instead turn our attention to the value, quality, structure, consistency, and insights within the data already at our fingertips."**

**ANDREW CRISP**
United Community Banks

coalesce®

**Coalesce** is the only data transformation and governance platform designed for the AI era. Built on a metadata-driven framework, Coalesce gives data teams the speed to build and deploy transformations 10× faster—while enforcing the standards, structure, and governance needed to scale sustainably. With Coalesce Catalog, transformation and metadata management come together in a single solution, enabling discovery, trust, and collaboration across the business. Whether accelerating AI-assisted migrations from legacy tools or future-proofing enterprise data architectures, Coalesce provides the guardrails and efficiency to keep data teams AI-ready.

Discover more at **Coalesce.io**.



Coalesce.io | sales@coalesce.io