

STATE OF DATA PRODUCTS

QUARTERLY RELEASE: COMMUNITY CURATION

Stepping towards an AI Activation Layer

Rethinking Data Platforms for AI

The data platform direction to become AI ready with next-gen AI systems that are context and memory enriched.



A TECH-REPORTING INITIATIVE
RITWIKI CHOWDHURY X MODERN DATA 101


WHAT'S IN STORE



The Perils of
Data Platforms



Platform Paradigms for
AI Readiness



The Fundamentals for
AI success



AI Ready Data to Action
Ready Data



Context & Memory
Layer

...and so much more!

Editor's Note: 2025 Q3

When Platform Efficiency Doesn't Translate to AI Readiness

Most data platforms run efficiently; pipelines hum, dashboards refresh, and storage scales. Yet AI stumbles. Why? Because efficiency isn't the same as readiness. These systems were built to integrate, not to interpret. Data moves, but meaning doesn't. Context stays trapped in silos, semantics get lost in transit, and teams spend more time wrangling data than training models.

The irony: even the most advanced platforms can slow AI down. They're optimised for analytics, not intelligence. To unlock real AI outcomes, organisations need a new setup, one that treats data as connected, evolving products built for activation, not just access.


AI systems don't thrive on data volume or compute power alone — they thrive on context, continuity, and connectivity. When platforms fail to carry meaning across domains, AI operates in isolation: powerful engines with no shared understanding. The result is intelligence without memory, automation without awareness, and endless reinvention of logic.

What worked for dashboards doesn't work for models. In analytics, accuracy was enough; in AI, understanding is everything. Models can't self-correct or collaborate without data that acts like a living product, one that preserves semantics, evolves with feedback, and stays discoverable.

The shift to AI-ready platforms isn't optional. Enterprises don't fail at AI because of weak algorithms but because their foundations were never built for reasoning.



From The Editor's Desk,
Modern Data 101



Failure Points of Platforms: A Hurdle for AI Success

After years of development, teams realise they optimised for capability, not usability. Employees are told to use it, not drawn to it. That's where platforms stall, feature-rich, yet lifeless. A mature platform is one where product evolution and user adoption grow in lockstep. Without that pairing, even the most advanced stack becomes an expensive ghost town.

Older platforms often didn't anticipate scale of complexity: more data sources, more domains, more real-time needs. When such platform teams are small/foundational, everything sounds feasible; when demands multiply, the old structure breaks down.

Reactive vs intent-driven

Traditional platforms tend to react: build dashboards when asked, answer queries, and perform BI. They don't bake in intent feedback loops. There's no built-in mechanism for agents to propose changes, suggest improvements to data products, or evolve with business logic.

Often prioritize technical capability over usability, leaving users disconnected. Silos trap context and semantics, making data hard to activate. Governance and lineage often break under scale, causing inconsistent outputs. Integration is brittle, and pipelines are fragile. Most critically, platforms react to requests instead of evolving with business intent, leaving AI initiatives starved of meaningful, connected data.

When Architecture Imitates Yesterday

As Peter Baumann, Principal Management Consultant Data & AI Strategy at INFOMOTION notes, most data architectures still mirror the systems they were meant to replace. We’ve simply moved operational silos to the cloud; each domain with its own schema, logic, and governance, then tried to stitch them together after the fact.

The result? Endless integration code, brittle pipelines, and architectures that stabilise instead of evolve. Every new use case means rebuilding connections instead of reusing context.

Baumann’s point is sharp: we can’t fix this with more cloud, mesh, or metadata tooling. The flaw isn’t distribution, it’s design.

Until platforms treat semantics and relationships as first-class citizens, they’ll stay stuck managing data, not activating it.

Data Architecture Pattern Evolve Over Time and Grow Stronger into an Organizational Context




Image: Positioning different data architecture pattern | [Source](#)

What industry experts say about data platform challenges




John Wernfeldt  • Following
I share insights about Data & AI | Data Governance Consultant | ex-Gartn...
[View my newsletter](#)
2w • 

Most data platforms don't collapse because of technology.
They collapse because of what's missing: governance.

Here's the pattern I keep seeing:

- New data source? "Just hook it up."
- New metric request? "Hardcode it for now."
- Lineage questions? "We'll document later."
- Ownership issue? "Ask around."

On the surface, everything looks fine. Dashboards load. Queries run. Leadership claps.
But underneath, the wall is crooked. Every new brick makes it worse.



Paolo Platter  • 1st
CTO - Co-founder at Agile Lab | Product Manager on Witboost
[Visit my website](#)
4mo • 

Most data platforms scale infrastructure. Very few scale decision-making.

As the number of data teams grows, the challenge isn't just provisioning environments or automating pipelines.

It's ensuring that:

- Everyone knows what good looks like
- Practices evolve without breaking trust
- Change is codified and visible, not just tribal
- Governance scales without centralizing everything

This is the real platform maturity curve:


- ➡ From automating tasks
- ➡ to shaping behaviors
- ➡ to enable change at scale.



Rob Strechay  • 2nd
Analyst | Advisor | Speaker | Product | Marketing | Sales | M...
1mo • Edited • 

Data platforms aren't just plumbing anymore; they're becoming the control plane for enterprise intelligence. Don't take it from me, hear it from the voices shaping the space:

◆ "Not all data is equally valuable... features feed decision processes that must occur in milliseconds." — Bill Waid, FICO




Call for Humans. The Value of Culture and Communication

We've automated so much, yet the biggest gaps aren't in systems—they're in people. Culture and communication aren't perks; they're the scaffolding that makes work meaningful and outcomes real. It's time to put humans back at the center.

Thomas Nys, Enterprise Data Architect & Delivery Lead, in one of his posts, pointed out a crucial aspect around Human APIs. At the core, platforms enable, but people are the ones who integrate and operationalise.

Tools and architecture alone can't enforce intent or clarity; human communication is what keeps data platforms coherent and actionable.



Thomas Nys  · 2nd

Freelance Enterprise Data Architect & Delivery Lead | Turni...

1w · 

Your real integration layer isn't code. It's people.

Every data platform diagram illustrates systems communicating with each other. In real life, people do the talking.

They translate business goals into data needs.

They translate tradeoffs into choices that leadership can act on.

They align intent when the documentation doesn't.

These are your human APIs — the connectors who keep things moving when tools don't.

Ignore them, and your architecture will drift.

Empower them, and your architecture will scale.

Platforms don't integrate organizations.

People do.


When human APIs disappear, does your platform still work?

Is AI Adoption All About a Race?

Successful organisations anchor AI on trust, context, and accessibility: data that's accurate, well-governed, semantically rich, and easy to use.

Those struggling with AI didn't fail at ambition. They skipped the foundation. Governance, stewardship, and alignment aren't afterthoughts; they are the AI strategy.

Because in the end, your data maturity decides your AI reality. As Sol Rashidi mentions in one of her articles, AI projects fail not because of technology, but because of misaligned expectations and lack of readiness.



“You can have an ambitious vision about AI's transformational potential (the art of the possible). But you have to ground that vision in what you can actually execute (the art of the practical),” mentions Sol Rashidi, Chief Strategy Officer (CSO), AI & Data at Cyera.

Rather than chasing large, abstract transformation goals, companies should start small, execute practical use cases, and build credibility through early wins. This phased approach creates the momentum and confidence needed for true AI transformation.



Rethinking Platform Strategies to Drive the AI Momentum

As AI keeps transforming how businesses operate, a lot of focus rests on the strategies we pick for data infrastructures and platforms. The same article mentioned above pinpoints the approach of unified access to data as an approach for working around the AI drill.

As Andi Gutmans, VP/GM, Data Cloud at Google, pointed out,

“For companies to succeed with agentic AI, they must shift from an incremental approach to a comprehensive data strategy that is ready for AI's needs.”

New data platforms unify all data; legacy, structured, or unstructured, into AI-ready, semantically connected systems. This shift lets teams build intelligent, automated workflows instead of wrestling with integration, turning accessibility into a real business advantage.

An article published by The World Bank similarly draws attention to the need for working on the foundational pillars, like the platforms for AI readiness.

What makes data “AI-ready?”

AI-ready development data is systematically organized and thoroughly documented to ensure its meaning and context are clear not only for subject matter experts, but also for general users and AI systems.

Three core pillars define AI-ready development data:

1. **AI-Ready Data Systems:** The foundational infrastructure—encompassing discovery platforms, APIs, and technical standards—ensures that data is not only stored but also readily discoverable, interoperable, and accessible.
2. **High-Quality Data and Metadata:** Reliable, up-to-date, and thoroughly documented data, accompanied by comprehensive and structured metadata. For AI applications, this entails datasets that are systematically organized and described with sufficient specificity to ensure both machine and human analysts can accurately interpret the information.
3. **Robust Governance and Strategic Partnerships:** The implementation of comprehensive policies, standardized procedures, and collaborative efforts across sectors is essential to ensuring data integrity, enhancing transparency, and advancing responsible utilization. These measures are fundamental to cultivating public trust among both human and AI stakeholders.

Embedding Product Thinking to Accelerate AI Success



“Data doesn’t carry any intrinsic value on its own. Its value relies on you having a purpose — and a process — for it.”

Emily Gorcenski
Head of Data & AI, Thoughtworks Europe

Emily Gorcenski highlights the importance of product thinking that helps integrate different data resources and ensure they can be accessed and applied, offering a foundation for a comprehensive digital strategy. Once this is in place, a data platform allows for gathering data insights and building reliable AI systems.

Some key strategies to consider:

- Treat Data as a Product, Not a Byproduct

Data isn’t exhaust from operations; it’s an asset. Each domain should own data products built for clear business outcomes, not pipelines. When data is designed with purpose and quality, it becomes AI-ready and directly tied to impact.

- Decentralise to Scale

Central control slows intelligence. Distribute ownership so domains manage their own governed, high-quality data products. Decentralisation turns scale and agility into default settings for AI deployment.

With New Platform Paradigms in Place, It's Time to Rethink the AI Fundamentals

As data platforms evolve and the idea of productised data takes root, one truth stands out, AI doesn't thrive on data alone, it thrives on context.

For all the talk of models and compute, the real breakthroughs will come from how we structure meaning, memory, and continuity across systems. Before any enterprise can expect intelligent outcomes, it must solve for three fundamentals: how context travels through data pipelines, how memory persists across interactions, and how control frameworks orchestrate this intelligence reliably.

In one of her posts, [Jessica Talisman](#), an Information Architect mentioned, as AI projects surge, critical foundational work is under-resourced, forcing employees to stretch beyond their roles.

She rightly quoted [Association of College and Research Libraries \(ACRL\)](#),

“No Human, No AI”




Image: The crucial aspects defining AI success parameters

Shifting Priority to Context Engineering




Rakesh Gohel 
@rakeshgohel01

Prompt Engineering VS Context Engineering


Prompt Engineering	Context Engineering
Involves crafting and refining inputs to guide AI models towards the desired responses.	Building systems to provide the right information and tools in the right format
Works with static, predefined inputs – the user is responsible for getting the right output.	Written for dynamic, multi-source inputs that automatically pulls from different sources.

What is Context Engineering?



Context Engineering Implementation


Prompt Engineering Implementation



Initial prompt
Craft the optimal prompt to elicit the desired response from the LLM

System Prompt
Design the prompt to assign a particular role the system

Context Engineering Implementation



Tools
slack

Tool Specification
• Tool schemas that are semantically clear to the LLM
• Data normalization for consistent processing

Build the Switch Logic
Design logic that enables the LLM to adapt based on input

Large Language Model × **Large Reasoning Model**

Context Filtering
• Applies business logic to filter noise from signal
drant, neo4j Database

State Management
• Maintaining conversation context across interactions
Memory



Ethan Mollick  • 3rd+
Associate Professor at The Wharton School. Author of Co-I...
3mo • 

The hottest discussion in AI is about "context engineering" - how you give AI the data and information it needs to make decisions. The most important thing to consider about context engineering is that cannot (and must) not be a solely technical function.

"Context" is actually how your company operates; the ideal versions of your reports, documents & processes that the AI can use as a model; the tone & voice of your organization. It is a cross-functional problem..


Don't offload how your company operates to some sort of RAG search of every document in your shared storage. Make choices about what kinds of context you need, build the ideal version of a report or document or instruction manual, and take context seriously.



Markus Hoerr • 1st
Driving x-functional excellence and lateral value creation w/ end-to-e...
3mo • 

The breakdown into 'real' engineering disciplines like "context engineering" is the way to go, as it cuts through many noise (e.g. there is more than GenAI) levels and helps creating real, sustained value.

A Semantic Shift for Data Platforms: The Context Needed



Modern enterprises are realising that being AI-ready isn't enough; data must be action-ready. The missing link? A semantic layer that carries business meaning across changing systems, keeping data understandable, contextual, and consistent.

An [article](#) by [Sagar Paul](#), SVP Global Growth & Solutions at The Modern Data Company, highlights this point perfectly.

In legacy models, system changes break AI logic because raw data loses meaning. The semantic layer fixes this by preserving intent. It translates technical fields into a business context, enabling AI to act, not just predict.


Image: How Semantic Layer Enables rich Supply of Context to AI Agents, Apps, and Workflows in the Application Layer | Source: [Animesh Kumar](#), CTO at The Modern Data Company

A Semantic Shift in Data Platforms: The Context Needed

To scale this across domains, organisations need comprehensive data management platforms built on data products; standardised, semantic-rich units that represent core entities like customers, transactions, or inventory. These platforms handle the creation, governance, and evolution of semantics across distributed systems, so AI applications can plug into reliable, reusable context instead of rebuilding logic from scratch.

Data products and semantic layers turn fragmented data stacks into cohesive intelligence systems, reducing AI implementation time from months to weeks, maintaining continuity during platform changes, and ensuring that every AI action is context-aware and business-aligned.

Image: The Integrated System which Connects Data in siloed sources and analytics dashboards to Action in operational systems | Source: [Sagar Paul](#), SVP Global Growth & Solutions at The Modern Data Company




Understanding (context) Without Memory can be Short-Lived

Embedding AI meaningfully goes beyond dashboards or chatbots. It's about enabling AI-native applications where data isn't just accessed, but actively consumed, orchestrated, and translated into action. Context-rich, AI-ready data is critical, but it's only half the story.

Once context is engineered, meaning alone isn't enough; memory must follow. Context helps AI understand the “what”; memory ensures it remembers the “why.” This is where platforms evolve from being data-intelligent to experience-intelligent, systems that learn from interactions, recall outcomes, and refine performance over time.

To make your platform the connective tissue for AI agents and analytics tools alike. You're not replacing visualisation; there are a few fundamental layers, of which the presence of the memory layer is non-negotiable today.

Image: The elements of context engineering | [Source](#)



AI agents today can't just react; they have to remember. A memory layer that lets models retain meaningful context across sessions: user preferences, past decisions, and ongoing dependencies. It extracts the most salient facts, updates or prunes memory entries, and retrieves what truly matters for the next interaction.

In enterprise settings, a memory layer preserves business continuity even when backend systems change. AI agents built with memory systems show 26% higher accuracy, 90% fewer token costs, and 91% lower latency compared to raw context-window approaches.

A Network of 10k Subscribers

The 10K Signal marked an important milestone for our community – a reflection of how far the Modern Data 101 movement has come. Built as a shared asset for learning and conversation, it continues to shape how we think about data, AI, and product thinking.

Crossing 10,000 readers, it reaffirmed the value of community-built assets like The Data Product Playbook, The Modern Data Survey Report 2024-25; resources that have become reference points for teams rethinking data, AI, and product thinking.

The Data Product Playbook has been a community highlight with 3k+ downloads, along with a strong LinkedIn community of more than 12k members.

This isn't just growth in numbers but a signal of a movement that's defining the next generation of data and AI thinking across the space.



MODERN DATA MASTERCLASS

KNOWLEDGE IN MOTION


Sessions for the Community by the Community

The Modern Data 101 Community has launched its first ever 'The Modern Data Masterclass,' this is a new community-led series

Here data practitioners, thought leaders, and product builders share practical insights that go beyond theory. Each session dives into real-world frameworks, challenges, and strategies defining the future of data.

True to the ethos, this is for the community, by the community – a space to learn, engage, and contribute. Nominations are now open for upcoming sessions led by you or your favourite voices in Modern Data.

Explore more [here](#) .



Join The Modern Data 101 Crusade!

If this playbook was helpful for your implementation or even as a means to organise your thoughts, we’d highly recommend you join us and help us build the Data Product Expertise in the data industry.

We are ALL THINGS DATA PRODUCTS!

Find a truly community-driven voice with experts from all walks of the data industry voicing their insights in the Data Product space on our channels. We’d love for you to come forward and share your insights with us so we can improve our delivery and truly create for YOU.

Subscribe

15000+

Data peers

200+


Resources

60+

Contributors

130+

Countries reached



Your Community, Your Benefits

Explore a world of reliable information, connect with experienced peers, share insights, boost visibility, receive valuable feedback, and propel your career forward in the ever-evolving data industry.



Access To Resources

In an age of infinite data perspectives, have trustworthy information to make a judgment.



Opportunity to Network

When learning about data, surround yourself with people more experienced than you.



Data Product Expertise

Find all things data products, be it strategy, implementation, or a directory of top data product experts & their insights to learn from.



Visibility & Recognition

Showcase your expertise, contribute to discussions, and build your reputation within the community.




Career Growth

Stay updated on the latest trends and job opportunities in the data industry, and gain valuable career advice.

Data Platform Challenges, AI Readiness

....and much more in store

Access Now!



moderndata101.com/state-of-data-products

