Data platform playbook

How data platforms and governance must evolve for an agent-driven world

thoughtworks

Why governance is becoming architectural

“Scaling enterprise data platforms is no longer just about technology. It’s about enabling organizations to act on data with speed, security and trust.”

Bilal Jaffery

SVP, Head of Data & AI, Americas
Thoughtworks

Organizations have invested heavily in data infrastructure yet many struggle to enable AI at scale. Most data foundations were not designed for the demands of advanced analytics and AI agents. Fragmented, stale or ungoverned data slows AI adoption and compounds risk across automated workflows and business insights.

This playbook brings together perspectives from architects, engineers and data leaders working across governance, compliance, machine learning and real-time analytics who’ve seen these platform gaps firsthand. They share the decisions, tradeoffs and lessons that helped organizations move from experimentation into production.

Our goal is to give technology leaders practical guidance for building platforms where AI does more than generate recommendations and can act responsibly, reliably and at scale.

The agentic operating system

Why AI governance must become a runtime concern

“We’re trying to govern systems that compose tool calls with rules designed for systems that return a value.”

Shayan Mohanty

Chief Data & AI Officer
Thoughtworks

AI has shifted from reactive tools to systems that plan, select actions and act across the environments they run in. Modern agents call tools, reach sensitive data and chain decisions with little human mediation in between. Capability accelerated quickly. Governance did not.

Traditional governance models were built for users and static systems.

They struggle with three properties that define agentic software:

- **Non-deterministic behavior** – the same prompt can produce different action sequences.
- **Dynamic tool usage** – what an agent touches is decided at runtime, not declared in advance.
- **Delegated authority** – an agent may act on behalf of a person, across systems that never see that person directly.

A policy document cannot keep up with software that composes its own next step. By the time a review cycle runs, the action has already happened, sometimes thousands of times.

This is why governance must move into the runtime. Operating systems faced the same problem decades ago with many programs, shared resources and no ability to trust each one. They solved it not with policy but with primitives, which are foundational controls the system enforces on every process, whether or not the program cooperates. Agentic systems need the same architectural thinking. Governance becomes a property of the substrate an agent runs in, not a layer bolted on after the fact.

Validate the plan, not just the step

Most enforcement fires one action at a time: check the request, allow or deny and move on. That is necessary, but for systems that compose, it is not enough. An agent can run for an hour, mutate real state across several systems and only then reach the step that should never have been permitted. The check was correct. It was simply too late.

A runtime built for autonomy should be able to reason about the whole plan before any of it executes. It should prove that at least one compliant path through the entire workflow exists before the first action is taken. It is the difference between catching an error while the program runs and

The five primitives of an agentic runtime

Execution tracing

Full visibility into what was decided, what was called and why.

Identity

Every agent is verifiable and accountable, and the human behind it stays attached to the action it takes.

Capabilities

Permissions are explicit, scoped and time-boxed to specific tools and data. Not roles inherited once, but capabilities granted for a purpose and then surrendered.

Resource accounting

Cost and compute are first-class controls, not an afterthought discovered on an invoice.

Sandboxing

Agents run in isolated environments that contain blast radius, under policies that can tighten mid-execution; a sandbox that reads sensitive data, for example, can lose its right to write elsewhere.

rejecting an entire class of errors before it ever does, the way a type checker refuses to compile code that could fail.

Governing the agent is only half the problem

An agent is only as governed as the resources it can reach. The strongest controls on the agent itself still leave a gap if the data and tools behind it are governed separately or not at all. Governance must be held in two places at once: at the agent's execution boundary and in the substrate of data and tools it acts upon. Govern the driver and the roads.

That second half is where data governance and AI governance stop being separate disciplines.

Unifying AI and data governance

Enterprise AI started with a simple idea. Models accepted inputs, generated outputs and stopped there. AI agents' access to data was limited to predefined feature sets, a constrained context window and a specific API surface. Governance was manageable because organizations could control what entered the system and audit what came out.

That framework no longer reflects how modern agents operate.

Agentic AI does more than generate predictions. It executes actions across APIs and data systems, often making sequences of decisions that are difficult to fully predict. Without strong controls, scaling quickly becomes a risk.

Most organizations built separate layers for data and AI governance. Agents don't adhere to that boundary. An agent can access data, call tools, invoke models and act on behalf of users in a single workflow. When governance remains fragmented across systems, organizations lose visibility into what an agent can access, what actions it can take and how decisions are made.

The primitives already exist

The five primitives Shayan describes mirror the same governance challenges that led Databricks to build Unity Catalog.

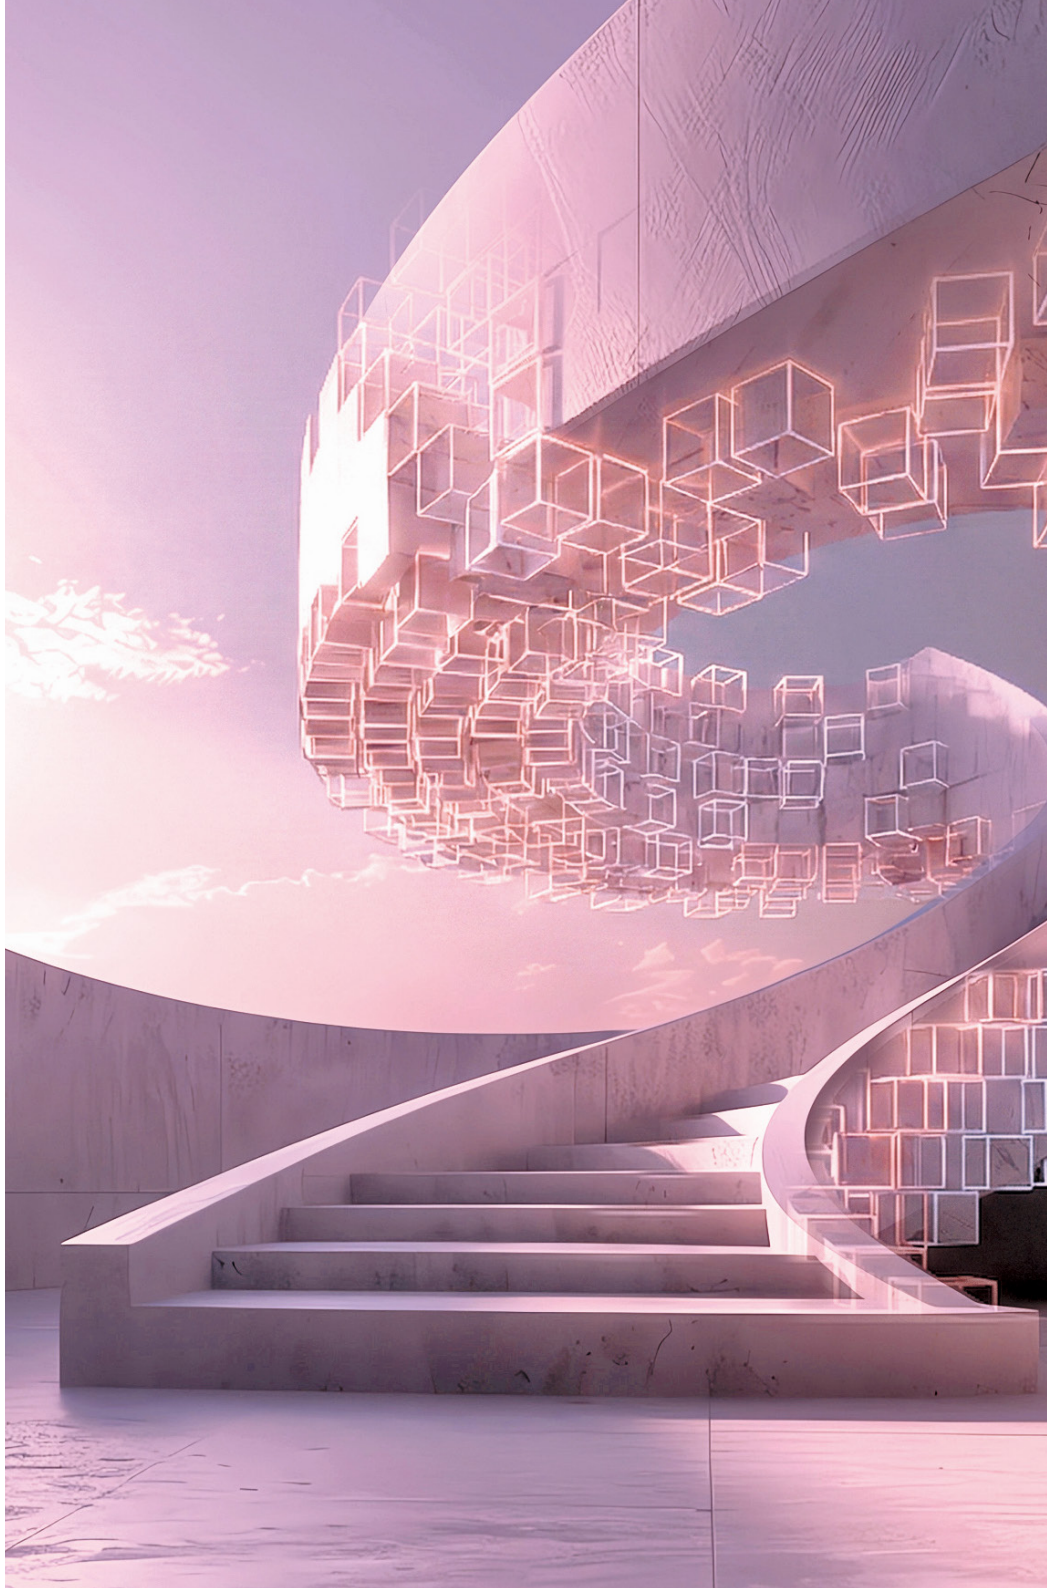

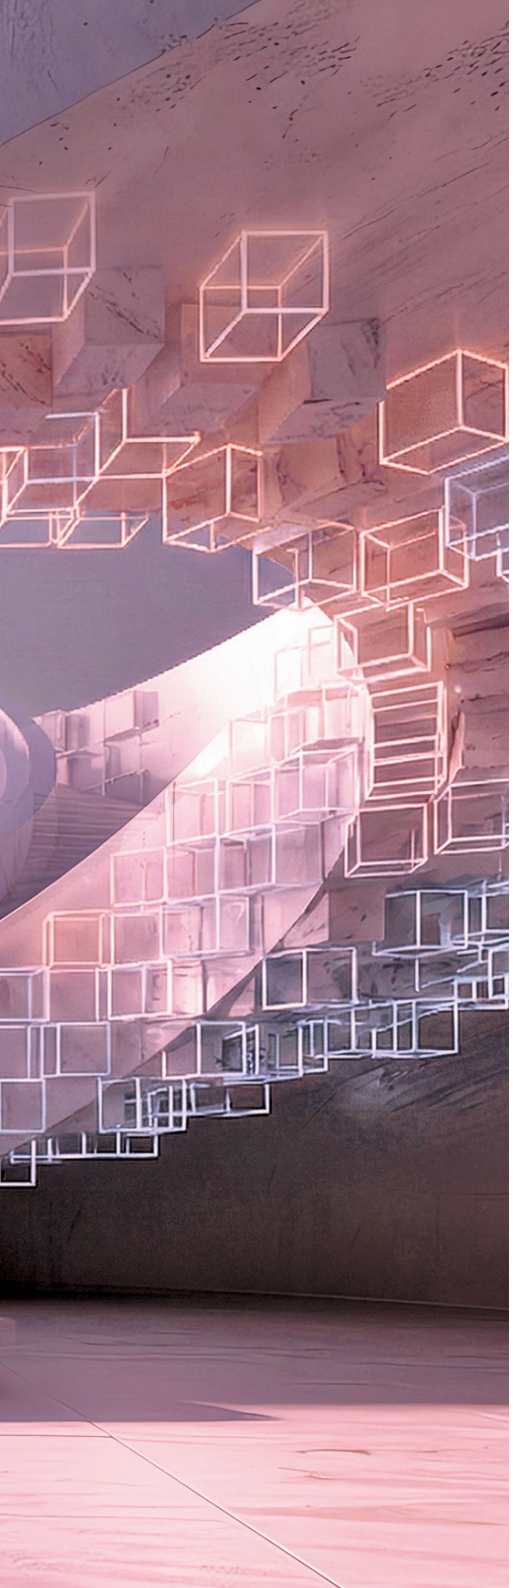

The catalog was built around a core principle that governance should live with the data, not inside every application that accesses it.

Enterprises already run hundreds of tools against shared datasets. Duplicating governance logic across every tool was never going to scale. That same model now naturally applies to agents. Whether an organization builds an agent with LangGraph or CrewAI or another orchestration framework, agents depend on the same governed resources. What an agent can do is largely shaped by the data and tools it can reach.

Governing the agent without governing the resources it accesses is like governing the driver but not the roads. This is why AI governance and data governance can no longer operate as separate systems.

When a unified catalog governs both, an agent automatically inherits access controls from the data it reads, policies from the tools it connects to and identity context from the user it represents. Governance is not something added onto the agent afterward. It already exists in the platform the agent operates within.

Extending governance into agent workflows

Agents reinforce the importance of a data-centric governance model. At Databricks, this led us to extend the same governance principles deeper into runtime execution through three core capabilities:

- Identity propagation across workflows.
- Policy enforcement at every interaction.
- Unified observability across data and AI systems.

Identity propagation

Identity must persist across agent workflows. When agents query data, call tools or invoke other AI agents, it becomes difficult to trace who originally initiated the action.

Traditional data systems assumed humans were the direct callers. Agents break that assumption. To preserve accountability, the platform carries the user's identity through the workflow so actions can always be tied back to the person behind the request, not just the agent or credential executing it.

Context-based policy enforcement

Permissions define what an agent can access but do not fully determine whether a specific action should happen in a specific context. That decision depends on the action being taken, the arguments involved, the resource being accessed and the user behind the request.

Modern agentic systems require policy evaluation at runtime. Databricks built service policies that evaluate interactions before agents request data or tools, using full contextual awareness of both the request and the identity that initiated it.

Unified observability

Observability must live alongside business data. Most monitoring systems isolate agent activity telemetry from the rest of the data platform. Teams can view traces and logs, but they cannot easily connect them to the data accessed, operational costs or interaction with business systems.

At Databricks, agent logs, traces, audit events and usage data are stored as standard Delta tables alongside enterprise data. That means agent behavior can be analyzed with the same tools, queries and governance model used everywhere else in the platform. This changes what becomes possible operationally.

This allows organizations to:

- Correlate agent behavior with data quality issues.
- Detect anomalous access patterns.
- Trace downstream business impact.
- Allocate costs across teams and workflows.
- Investigate incidents with full context of operational telemetry and business data.

Traditional observability systems were not designed for this level of integrated governance.

Unified governance is what makes autonomy possible

Extending the same governance model for enterprise data and agent workflows turns governance from a bottleneck into an enabler. The same controls that already govern enterprise data can now govern how agents access data, invoke tools and act. That makes it possible to scale agents safely across the organization.

David Nasi

Director, Product Management, AI and Agentic Platform, Databricks

Client spotlight

Real-world proof from the agentic enterprise

“OAG data sits inside the decisions airlines, airports and travel platforms make every minute of every day. With Thoughtworks, we’ve modernized the foundation so we can do more than tell our customers what happened. We can tell them what’s about to happen and give them the intelligence to act on it.

A governed, real-time data platform is the prerequisite for predictive and agentic aviation.

The foundation is built. Live predictions are in production. The intelligence layer that will power the next generation of airline, airport and travel operations is being built on top of it.“

OAG

James Cameron-Williams

Chief Product & Technology Officer, OAG

Organizations building for the agentic AI era understand that systems are only as effective as their governed data foundations.

This story shows how unified governance and architectural discipline help organizations move from experimentation to production.

OAG: Building the intelligence layer for global aviation

From foundational schedules data to predictive, real-time intelligence at scale.

OAG sits at the heart of the global aviation ecosystem. Every major airline, airport, online travel agency (OTA), global distribution system (GDS) and many emerging AI-driven travel systems depend on OAG schedules, status, fare and capacity data to make decisions.

As the industry shifts from reporting on what happened to predicting and acting on what happens next, the demands on OAG's data foundation have changed fundamentally.

Legacy architecture built for batch delivery of reference data could not provide the real-time context, governed access and machine-learning-ready pipelines required to power predictive models and emerging agentic travel systems.

A planned four-year migration timeline risked leaving OAG and the customers who built it a step behind the market.

The goal was to modernize the foundation, accelerate predictive intelligence capabilities, and position OAG as the trusted data and intelligence layer for an industry moving toward autonomous decision making.

The approach

OAG and Thoughtworks executed a strategic platform transformation that consolidated fragmented data ecosystems into a unified, governed foundation on Databricks with Unity Catalog at its core.

The work spanned three supporting layers:

- A modern data platform engineered for high-velocity aviation feeds.
- A governed, machine-learning-ready data product layer covering schedules, status and fares.
- An intelligence layer that turns data into predictive signals customers can build into their own operations.

One technical proof point is live status prediction that refines estimated departure and arrival times in real time using OAG's global status feed combined with other data sources. For airlines, that means earlier and more accurate visibility into delays, better turn planning and fewer downstream disruptions. For the platforms, OTAs and agents that depend on OAG, it means richer, predictive intelligence delivered through the same trusted data layer they already consume.

The outcome

Reduced time to onboard new data feeds by more than 50%, accelerating the pace at which OAG can improve data quality for customers and stand up more refined intelligence products.

Compressed a four-year legacy migration into a 12-month delivery window without disruption to mission-critical customer feeds.

Refined estimated time of departure (ETD) and estimated time of arrival (ETA) in real time with predictive status intelligence to give airlines and travel platforms earlier, more accurate operational signals.

Decreased issue remediation from five steps to two through unified governance, automated quality dashboards and end-to-end tracing across the platform. This improves our responsiveness to customer queries.

"When organizations guarantee data quality and orchestrate unified governance at scale, they empower AI to stop merely making predictions and start taking meaningful business action."

Bilal Jaffery

SVP, Head of Data & AI, Americas
Thoughtworks

Accountability at the moment of action

For decades, the boardroom conversations about technology focused on three questions: How fast can we ship? How much will it cost? What value will it create?

AI is changing that conversation. The new question CEOs and boards are beginning to ask is not what AI can do. It is what AI does on their behalf, and whether the organization is ready to be accountable for those actions.

That shift is more significant than it first appears. When a model generates a recommendation, a human still makes the final decision. When an agent acts autonomously, the organization has already decided. Every approval gate, audit trail and policy framework designed for traditional enterprise systems assumed a human was in control. Agentic systems challenge that assumption.

“When an agent acts, the organization has already decided.”

Asha Saxena

CEO, World Leaders in Data and AI (WLDA)

Chair, xnode.ai

Author, The AI Factor & Digital Human Advantage

The failure of static policy

In my work advising CEOs, CIOs and Chief Data and AI Officers, and through the AI steering committees and board readiness programs I lead, the conversation I hear most often is not about capability. Leaders know their agents can act. What they question is whether their organizations are ready to be accountable for those actions.

In most cases the honest answer is no. The problem is structural. Most enterprise governance still looks like a document that was written once, reviewed annually and signed off by the right committees. That model worked when the decisions it governed were human, slow enough for policy review and interpretation.

Autonomous systems do not consult policy. Agents act in real time across systems and datasets in ways policy authors could never fully anticipate. By the time an annual audit cycle runs the action has already happened, sometimes thousands of times.

That is not only a technical challenge. It is a leadership challenge.

What healthcare taught me about runtime governance

At Aculyst, the healthcare analytics firm I founded, we built data and AI capabilities for large healthcare providers. In healthcare, governance documentation alone is not a defense. Governance must be enforced at the moment of action, when a clinical recommendation is generated, when a record is accessed or when a downstream system acts on a result. Documentation comes second. Enforcement comes first.

That mindset is increasingly necessary for every industry in the agentic era. In “The AI Factor,” I introduced the concept of the “Power Zone,” the intersection of persistent business problems, rich operational data and high-value business potential.

That is where AI creates the greatest impact. The challenge is that agentic systems do not remain in an organization’s “power zone.” Agents compose, delegate and act across systems in ways static governance models cannot anticipate. The organizations succeeding are not adding a governance layer on top of their agents. They are embedding governance directly into every dataset, every tool call and every action the agent is authorized to perform.

Governance as architecture, not documentation

This is the transition enterprises need to make.

As Chair of xnode.ai, I have seen what it takes to make governance work at the infrastructure level where policy is enforced at the moment of request, role-based access is applied at each interaction and every tool call is logged with full context and timestamped before execution completes. It is not reviewed after the fact. It is enforced before the action occurs.

Shayan and David make the architectural case earlier in this playbook. From the boardroom perspective, the conclusion is straightforward. Governance is no longer a function inside the organization. It is a property of the platform itself. The organizations that understand this first will not only move faster with agentic AI but also be able to explain, trace and defend their agents when it matters most.

The leadership mandate

For leaders, the message is not to slow AI adoption. It is the opposite. Trust embedded into architecture is a competitive advantage.

Through WLDA, the peer network I founded for Fortune 1000 data and AI leaders, and through the board readiness programs I lead, I see executives proving this every day. They are moving faster with AI precisely because their governance model gives their boards, regulators and customers confidence in those systems they operate.

The question is no longer whether organizations can deploy AI. It is whether your governance can keep up with what systems are now authorized to do and whether it enforces that authorization at the moment it matters.

Asha Saxena

CEO, World Leaders in Data and AI (WLDA)

Chair, xnode.ai

Author, The AI Factor & Digital Human Advantage

What this means for you

Across industries, the same shift is happening:

- Data platforms are becoming AI platforms.
- Governance is becoming architectural.
- Speed and trust are becoming competitive differentiators.

Explore what's possible

1. Discover more real-world playbooks.
2. See technical demonstrations in action.
3. Connect with Thoughtworks experts.

Building data platforms that scale AI in production

Communications, Media and Entertainment Partner of the Year, Latin America, 2025

Innovation Partner of the Year, Australia and New Zealand, 2024

About Thoughtworks

We are a global technology consultancy that delivers extraordinary impact by blending design, engineering and AI expertise.

For over 30 years, our culture of innovation and technological excellence has helped clients strengthen their enterprise systems, scale with agility and create seamless digital experiences.

We're dedicated to solving our clients' most critical challenges, combining AI and human ingenuity to turn their ambitious ideas into reality.

thoughtworks.com

/thoughtworks

Design. Engineering. AI.

Last updated June 2026