The Two-Body Problem

Why AI Ambition Fails Without Data Foundations and Process Discipline

Jim Bartro, SVP of Consulting Solutions, Curate Partners

Sam Levine, Lead Data Strategy Consultant, Curate Partners

June 2026

Companies Do Not Have an AI Idea Problem

Many of the large enterprises we work with have AI ideas. Dozens of them. Sometimes hundreds. They come from department heads, from operations teams, from the C-suite, from vendor presentations, from board meetings. The ideas are rarely the bottleneck.

The bottleneck is what sits underneath them. Before any AI use case can move from concept to production, it must answer a set of foundational questions that most organizations are not equipped to answer consistently. Is the data it depends on governed, accessible, and trustworthy? Does the architecture support the integration pattern the use case requires? Are the access controls, lineage, and quality standards in place to make the output reliable and auditable? And even if the data is ready, is there a structured process to evaluate the use case, route it to the right build pattern, clear it through governance, and measure whether it actually delivered?

Most organizations treat these as separate problems. Data readiness lives with the platform team. AI use case intake lives with the innovation office or the line of business. Governance lives somewhere in between. The result is predictable: use cases that pass business evaluation stall at architecture review because the data is not ready. Use cases that have the data never make it through intake because there is no system to capture and prioritize them. Pilots launch without governance and get reworked mid-sprint. Investment is made without a measurement framework to determine whether it paid off.

This is not a technology problem. It is an execution problem. And solving it requires two things that most organizations are treating as separate workstreams when they are, in fact, inseparable: a well-governed, well-architected enterprise data foundation, and a structured, repeatable framework for managing AI use cases across the full lifecycle from ideation through measurement.

Two Pillars, One System

AI transformation at enterprise scale depends on the convergence of two foundational capabilities. We call this the two-body problem because neither pillar can succeed in isolation, and organizations that invest heavily in one while neglecting the other will find themselves rebuilding within 12 to 18 months.

Pillar 1: A Well-Governed, Well-Architected Enterprise Data Foundation

The first pillar is the data foundation. AI is only as reliable as the data it operates on. Organizations that skip this step, or treat it as a parallel initiative that will "catch up later," build on a foundation that cannot support the weight of what they are trying to do.

In practice, what we see is this: a use case makes it through intake, clears business value evaluation, and reaches the architecture and feasibility stage, only to stall because the data it depends on is fragmented across systems, ungoverned, or inaccessible to the tools that need it. The use case is not the problem. The data is.

A well-governed data foundation means consistent data products with clear ownership, lineage, and quality standards. It means a catalog that connects data assets to business meaning. It means access controls that are enforced by policy, not by tribal knowledge. And increasingly, it means a semantic layer that bridges raw data and the context that AI models need to produce reliable, auditable outputs.

Modern data platforms provide the technical infrastructure to achieve this. Lakehouse architectures unify structured and unstructured data. Catalog and governance tools enforce access, lineage, and classification at the platform level. Model lifecycle management tools track experiments, versions, and deployments. These capabilities exist today. But the platform is the enabler, not the strategy. Without a structured process to route AI use cases to the right data, through the right governance, and into the right architecture pattern, even a well-built data platform sits underutilized.

Governance cannot function as a documentation exercise. This means a unified catalog that spans structured data, unstructured data, model artifacts, and AI pipelines, with a single point of truth for who can access what, where data originated, and what transformations occurred before it reached a model. In regulated industries, the absence of this capability is not an operational inconvenience. An AI use case that cannot produce a complete lineage record from raw source to model output will not pass audit review regardless of how accurate its outputs are.

A frequently underestimated cost driver in enterprise AI programs is redundant feature computation. In most organizations, teams independently compute the same derived data constructs, including customer risk scores, utilization metrics, and transaction aggregations, with different logic, no version control, and no shared registry. Feature management infrastructure that centralizes how derived data constructs are defined, versioned, and served consistently across training and production workloads resolves this. A feature built for one use case becomes a validated, reusable building block for the next, which is one of the primary mechanisms by which programs scale to 100-plus use cases without proportional growth in data engineering cost.

Perhaps the most consequential gap in enterprise AI readiness is the absence of a semantic layer, a structured representation of what data means in business terms. An AI system querying data about customer churn needs to know what churn means in this organization, which data product holds the authoritative definition, what business rules govern the calculation, and which populations are in scope. Without this layer, every new AI use case must independently resolve the gap between raw data and business meaning, creating inconsistency and rework at scale. Organizations that invest in building a semantic layer are creating context infrastructure that every subsequent use case inherits. This is the difference between programs where business meaning is reconstructed by each team and programs where it compounds as a durable organizational asset.

Pillar 2: A Structured, Repeatable Process for AI Use Case Lifecycle Management

The second but equally important pillar is process discipline. This means a governed, end-to-end system for how AI use cases are captured, evaluated, prioritized, funded, built, and measured across the enterprise. Not a suggestion box. Not a spreadsheet. A repeatable operating model with defined phases, clear decision points, and accountability at every stage.

Without this, organizations face two compounding risks. The first is slow progress. Use cases stall in evaluation, get stuck waiting for architecture review, or sit in a backlog with no mechanism to sequence or decline them. Leadership loses visibility into what is actually in flight and why.

The second risk is less obvious but more dangerous: unmanaged cost. Every AI use case that reaches production consumes compute and token volume. Without structured intake and pattern routing, organizations accumulate AI workloads with no cost governance, no usage telemetry, and no mechanism to evaluate whether the output justifies the spend. Unmonitored token volume is the new shadow IT. The difference is that shadow IT was a security problem. Unmonitored AI consumption is a financial one, and it scales faster than most finance teams can track.

Process discipline means every use case enters through a standardized intake, is evaluated against consistent criteria (business value, technical feasibility, data readiness, risk, organizational impact), is routed to a defined solution pattern, clears governance gates before a dollar is spent on delivery, and is measured against its original hypothesis after deployment. This is not bureaucracy. It is the operating model that makes scale possible.

Why They Must Converge

A well-governed data foundation without a structured process to activate it produces expensive shelf-ware. You can invest in a modern data platform, build governed data products, catalog everything, enforce lineage and access controls, and still have no mechanism to connect that investment to the AI use cases that justify it. The platform is ready. The organization is not.

A structured process without data foundations produces well-organized failure. You can build a repeatable intake system, evaluate and prioritize use cases against consistent criteria, fund the right ones, and still watch them stall at architecture review because the data they depend on is ungoverned, fragmented, or inaccessible. The process works. The foundation does not exist.

The organizations that will demonstrate measurable AI value over the next 24 months are the ones that treat these two pillars as a single, integrated system. Not two workstreams. Not two roadmaps. One operating model.

What This Looks Like in Practice

A large integrated healthcare organization (payer and hospital system) faced exactly this challenge. Leadership had strong AI ambition and growing demand from across the enterprise. But requests were arriving through fragmented channels with no common intake, no evaluation framework, and no path from idea to delivery. At the same time, the state of data across the enterprise was uneven: some domains were well-governed, others were fragmented, and the connection between data readiness and AI feasibility was not being assessed until late in the process, when rework was expensive.

The engagement started with a structured assessment of where the organization stood across AI strategy, governance, data readiness, talent, adoption, and technology. That baseline became the evidence for what needed to change and where to start. It was not the campaign promise. It was the diagnostic that made the rest of the program credible.

Over the following months, a standardized intake process was established across the enterprise.

The heart of the intake process is intentionally human-centric. We begin by asking what the humans in a workflow are actually doing and why. We empathize with the challenges and opportunities they face in accomplishing their core objectives. Once we understand the workflow from a user/operator perspective, then we work to understand the appropriate and meaningful insertion of AI technology to aid the people (help them accomplish a task faster or with higher accuracy, for example).

What the structured process revealed was as valuable as what it produced:

Roughly 10% of submitted use cases were identified as not requiring AI at all. They were process problems, data problems, or training problems dressed up as AI opportunities. Identifying them early avoided wasted investment and redirected effort toward solutions that would actually work.

Approximately 30% of use cases could be addressed using enterprise tools the organization already owned. No custom build required. The structured evaluation surfaced capabilities that were already in place but not visible to the teams requesting AI solutions.

The remaining use cases were routed to defined solution patterns based on their complexity, data dependencies, and integration requirements. This routing mechanism is what made scale possible. Instead of treating every use case as a custom project, the program matched each one to one of four architecture patterns, from simple configurations requiring no new data connectors to full custom builds requiring multiple system integrations.

The economic leverage point was a Core AI Team that operated two sprints ahead of every delivery pod. This team built shared components once so that every downstream team could reuse them. A single capability served four divisions. A single data connector served more than ten use cases. This is what makes hundreds of use cases economically viable and what separates a coordinated program from a collection of separate projects.

The data foundation was not a separate initiative. It was embedded in the evaluation process from the start. Every use case was assessed for data readiness alongside business value and technical feasibility. Use cases that required data that was not yet governed, accessible, or trustworthy were flagged and either sequenced behind data readiness work or routed to vendor backlog. The process and the data foundation operated as one system.

That integrated approach spanned more than a dozen business areas, including Clinical, Compliance, Health Analytics, Integrated Network Delivery, Medicaid, Member Experience, Physician Services, Provider, Revenue Cycle, Sales, Marketing and Communications, HR, Contracting, and the organization's Core AI function. AI demand was not concentrated in a single department. It was enterprise-wide, which is precisely why a standardized intake and routing system was necessary.

Three specific examples illustrate how assumptions shifted through structured evaluation:

An HR team assumed that adding a copilot on top of benefits documents would let employees ask questions about their coverage. The structured intake revealed that real value required personalized answers tied to each employee's profile and elections, which meant integrating with HRIS systems and designing access controls. The MVP was scoped to guided benefits Q&A with policy grounding. The scale path required the data foundation work first.

A contracting team assumed that LLMs would dramatically reduce contract cycle time. The evaluation uncovered that human and legal review steps would still dominate total cycle time. The near-term win was accuracy and completeness: fewer back-and-forth cycles, fewer missing fields, better handoffs. The intake workflow was reduced from days to minutes through AI-assisted form completion, but the success metric was rework reduction, not cycle time elimination.

A marketing team assumed that LLMs could generate more content at scale. The real business objective was targeted, member-relevant messaging. But there was no customer data platform or golden record to support segmentation and personalization. The use case was valid, but the data foundation was not ready. The MVP was scoped to content generation for a small set of campaigns, with the scale path dependent on standing up the data assets first.

In every case, the structured process caught the gap between assumption and reality before significant investment was committed. And in every case, the data foundation was the determining factor in whether the use case could scale.

The organization is now projected to reach 500+ use cases by year-end, operating within a governed pipeline that connects intake to delivery to measurement.

Leveraging Databricks

The organization in this example operates on the Databricks Data Intelligence Platform. The technical capabilities required by this operating model map directly to platform components. Delta Lake provides the versioned, ACID-compliant storage foundation that makes training data reproducible and lineage traceable. Unity Catalog enforces the governance layer, including row-level security, automated lineage, and asset classification, that regulated industries require before AI outputs can be considered auditable. Databricks Feature Store addresses the training-serving consistency problem that causes programs to stall as they scale. MLflow closes the model accountability loop that structured intake opens at use case evaluation, registering every experiment and versioning every deployment decision. And Mosaic AI enables the retrieval-augmented generation pipelines and agent workflows that bring semantic context capabilities to production workloads.

These are not abstract capabilities. They are the platform-layer expression of the same requirements this operating model imposes at the process level. The platform does not substitute for the operating model, but an operating model without platform-level infrastructure to back it will find its practical ceiling far sooner than expected.

What This Means for CIOs and CDOs

If you do not have a structured intake process, you do not have an AI program. You have a backlog. The volume of AI ideas will only increase. Without a human-led system to evaluate, route, and govern them, the backlog grows faster than the team can respond, and the loudest request wins regardless of value.

If you have not resolved your data foundations, your AI use cases will stall at architecture review, not at deployment. The most common failure mode we see is not a technology limitation. It is a use case that passes business evaluation but cannot proceed because the data it depends on is ungoverned, fragmented, or inaccessible.

Cost governance is not a finance problem. It is an architecture problem. Every AI workload in production consumes resources. Without pattern routing and usage telemetry, organizations have no mechanism to evaluate whether the output justifies the spend. The organizations that approach AI with structure will exhibit managed cost profiles. Those that do not will discover the problem in their cloud invoices.

The organizations that will outperform are not the ones with the most pilots. They are the ones with a governed operating model that connects AI ambition to data readiness to measurable outcomes. Structure and rigor are competitive advantages.

For Transformation Leaders

If you are responsible for managing the intake flood, prioritizing competing requests from across the enterprise, and demonstrating measurable progress to executive leadership, this framework gives you something concrete to work with. It provides a repeatable, defensible process for evaluating AI opportunities against consistent criteria. It protects you from the political risk of picking the wrong use cases by grounding every decision in business value, technical feasibility, and data readiness. And it gives you a measurement framework that tracks what actually happened, not just what was predicted, so that when ROI shifts from the original hypothesis (and it will), you can explain why and what was learned.

About Curate Partners

Curate Partners is a boutique consulting firm specializing in practitioner-led engagement across data, analytics, AI, digital strategy, and platform-enabled operating models. The firm primarily serves large, regulated organizations in healthcare and financial services.

Curate has always been a 'People First' organization; our DNA is people: finding them, placing them, building teams around them. That heritage shapes how we approach every consulting engagement. We lead with the human element because AI transformation is not a technology problem. It is a people problem. Understanding who is doing the work, how they do it, and what would make them more effective is the starting point for everything we deliver. That instinct, to 'lead with humans and land with frameworks', is what separates our work from firms that start with architecture and work backward toward the business.

Curate lives at the junction of AI transformation and enterprise data. Our heritage is in building data products at enterprise scale in regulated industries, including helping organizations adopt and operationalize contemporary data platforms. Our AI practice extends that foundation into structured, repeatable programs that move organizations from experimentation to governed, measurable AI adoption.

The **Curate AI Accelerator**, powered by the **ATOM (AI Transformation Operating Model) Framework**, is our enterprise-scale offering for organizations ready to move from scattered pilots and uncertain ROI to a prioritized, governed, scalable AI operating model. The program starts by assessing where you are today across AI strategy, use cases, governance, data readiness, talent, adoption, technology, and value realization, then identifies the operating model gaps that are slowing ROI and builds a practical roadmap to accelerate execution.

Curate's AI portfolio spans three tiers to meet organizations where they are:

> **Tier 1: Activate** *(AI Activation QuickStart)*. A focused working session for a single team. Immediate wins in existing tools. Days, not months.
>
> **Tier 2: Accelerate** *(Curate AI Accelerator, powered by ATOM)*. 12 to 16 weeks. Enterprise-wide. Structured intake through measurement. The full operating model.

Tier 3: Automate *(AI Workflow Automation and Agent Activation).* Ongoing delivery sprints. System-executed workflows. Governance and lifecycle support at scale.

Start the conversation.

curatepartners.ai
