

What CIOs, CMOs, and CDOs Are Getting Right/Wrong About Enterprise AI — and the Patterns That Actually Drive Adoption

Cross-Industry Enterprise AI Adoption Brief

Vijay Kukreja, Chief Strategy Officer, Curate Partners

Key Contributions from Curate Data & AI Experts:

- Sam Levine
- Aseem Rastogi
- Jonathan Pulliza

Executive Summary

This brief distills patterns observed across active AI adoption programs in healthcare, financial services, nonprofits, and SaaS platforms during Q4 2025 and Q1 2026. These are practitioner insights drawn from the rooms where architecture decisions are made; budgets are debated, and teams confront the gap between AI ambition and operational reality.

Key Findings:

- **The foundation crisis is universal.** Every organization had discovered their data infrastructure was not ready for the AI strategy leadership had approved. Most are operating at Tier 1 maturity with promised Tier 3 delivery milestones.
- **Cost governance is the invisible blocker.** AI operates as a variable cost that scales non-linearly. Fewer than 20% of organizations had appropriate cost controls in place before activating their first production use case.
- **Productivity tool deployment is not an AI strategy.** Organizations that conflate deploying a single enterprise AI tool with having an effective AI strategy delay meaningful progress by 6–12 months.
- **Model and framework proliferation creates paralysis.** Teams simultaneously evaluating 4–5+ LLMs and agent frameworks are accumulating technical debt before shipping anything meaningful.
- **Change management is the actual adoption bottleneck.** Organizational resistance and fear of displacement are what stall adoption. Trust-building beyond technical training is the unlock.
- **The agentic frontier demands human-in-the-loop design.** The competitive advantage has shifted from "we deployed agents" to "our agents know when to ask for human judgment."

The bottom line: The organizations making progress treated AI as a capability to be earned through foundational work — not a product to be purchased.

The State of Enterprise AI Adoption: What We're Seeing


Over the past quarter, Curate Partners has been embedded in active AI adoption programs across healthcare payers, financial services institutions, nonprofit organizations, and SaaS platforms. This brief distills the patterns we observed firsthand — not from survey data or vendor marketing, but from the rooms where

architecture decisions are made, budgets are debated, and teams confront the gap between AI ambition and operational reality.

The central finding is this: the organizations struggling most with AI adoption are not struggling with the technology. They are struggling with the prerequisites — data readiness, governance clarity, organizational change management, and the discipline to sequence investments properly. Meanwhile, the organizations making progress share a common trait: they treated AI as a capability to be earned through foundational work, not a product to be purchased.

The Core Insight

Most enterprises want Tier 3 AI capabilities (generative, agentic) but have not completed Tier 1 foundations (data governance, canonical models, API layers). The gap between aspiration and readiness is where value — and risk — concentrates.



Pattern 1: The Foundation Crisis Is Universal

Every organization we worked with this quarter — regardless of size, vertical, or AI maturity — discovered the same uncomfortable truth: their data infrastructure was not ready for the AI strategy their leadership had approved.

What We Observed

At a healthcare organization, leadership purchased a new data platform and mandated a specific architecture to enable AI use cases. However, implementation revealed major legacy challenges: core systems required modernization, governance frameworks were missing, and the ingestion layer was handling preprocessing that newer systems could manage natively. The tools were modern. The thinking was not.

At a financial institution, multiple data product initiatives were running in parallel across different business units, each unaware of the others. No unified ontology. No shared object model. Teams were building prototypes for problems they had not yet researched.

At another organization, leadership saw AI support chatbots on their homepage and wanted their own specialized version launched immediately — before the backend data, compliance frameworks, or business cases were in place. Early enthusiasm collided with a six-month realistic development timeline.


The Data Maturity Pyramid

We use a three-tier framework to assess organizational readiness:

Tier	Capability	What It Looks Like
Tier 1: Foundation	Data ingestion, governance, cloud enablement, canonical models	Clean pipelines, unified schemas, role-based access, documented lineage
Tier 2: Enablement	Data products, customer 360, operational insights, integration APIs	Business-consumable data, cross-system joins, self-service analytics
Tier 3: Intelligence	Advanced analytics, ML models, generative AI,	AI-powered decisioning, autonomous agents,

	agentic workflows	predictive operations
--	-------------------	-----------------------

Most organizations we assessed in the last two quarters are operating at Tier 1 with Tier 3 ambitions. The gap is not a technology problem — it is a sequencing problem.



Pattern 2: Cost Governance Is the Invisible Blocker

One of the most consistent and least discussed patterns we observed: AI cost management is poorly understood and rarely governed before deployment.

What We Observed

An organization experienced large expenses in GPU and token costs within three to four days of internal AI usage. There was no consumption monitoring, no role-based limits, and no cost attribution by department. The pilot was paused indefinitely.

Multiple clients discovered that large tech SaaS AI enabled features operate on consumption-based pricing with variable and often unpredictable cost profiles. Without proactive governance, a successful pilot can become a budget crisis at scale.

Across our engagements, we found that fewer than 20% of organizations had implemented AI-specific cost controls before activating their first production use case.

For CFOs and CIOs

AI is not a capital expense. It is a variable operating cost that scales with usage, often non-linearly. Every AI initiative should have three things before activation: a cost ceiling, a consumption dashboard, and a defined escalation path when spending exceeds forecast.

Pattern 3: The Productivity Tool Conflation

The most common misconception we encountered: organizations that have deployed a widely adopted enterprise AI productivity tool believe they have an AI strategy.

What We Observed

We heard variations of the same statement in across business contexts from leadership: "We have our enterprise AI productivity tool rolled out, so we're covered on AI." This conflates tool deployment with strategic capability. An enterprise AI productivity tool is an individual productivity layer. It does not address data integration, workflow automation, customer-facing intelligence, or operational decision-making.

In one engagement we saw the AI productivity tool deployment had been treated as the entirety of the AI initiative. When we facilitated a workshop to identify high-friction business tasks suitable for AI, the team discovered dozens of use cases that the chosen tool could not address.

The pattern repeated across verticals. These tools address the "individual knowledge worker" problem. They do not solve the "enterprise intelligence" problem. Organizations need both, and confusing one for the other creates a false sense of progress.

The Distinction That Matters

An enterprise AI productivity tool is a personal productivity layer. Enterprise AI strategy requires data products, integration architectures, governance frameworks, and use-case-specific implementations that a single productivity tool cannot provide. Conflating the two delays meaningful progress by 6–12 months.

Pattern 4: Framework Proliferation Is Creating Organizational Paralysis

Organizations are experimenting with too many AI models and frameworks simultaneously, creating technical debt before they have shipped anything meaningful.

What We Observed

Across our engagements, teams were simultaneously evaluating 4 or more LLM's and multiple agent orchestration frameworks — often with different teams evaluating different tools for overlapping problems. The result was not innovation. It was analysis paralysis.

Some clients had multiple separate proof-of-concept projects using multiple agent frameworks, none of which had reached production. Each POC team was learning independently, building incompatible infrastructure, and competing for the same budget.

The organizations making progress picked one or two frameworks and built deep expertise in a specific stack. For instance, one enterprise who had one stack for customer-facing agents and another distinct stack for internal use cases saw clear progress. It reduced the decision surface, concentrated learning, and allowed teams to move from POC to production.

The Framework Selection Discipline

We recommend organizations make three explicit choices early:

- 1. Where do your agents interact with customers, members, or external stakeholders?**
- 2. Where do your employees interact with AI for productivity and operations?**
- 3. Where do your data teams build models, pipelines, and intelligence layers?**

Making these three choices — and committing to them for 12 months — eliminates 80% of the framework proliferation problem.

Pattern 5: Change Management Is the Adoption Bottleneck

Technical implementation is rarely what stalls AI adoption. Organizational resistance, role confusion, and fear of displacement are the real blockers.

What We Observed


At a healthcare organization, widespread employee anxiety about AI replacing jobs surfaced in team standups, 1:1s, and performance conversations. Leadership recognized this was not a training problem — it was a trust and framing problem. Their breakthrough came when they reframed AI literacy as a career skill equivalent to standard office productivity suite proficiency twenty years ago: “If you can’t use these tools, you will fall behind. If you can, you will expand what you’re capable of.”

At a nonprofit, staff feared automation would eliminate positions. Leadership’s reframe: “We are not replacing people. We are freeing them from paperwork to serve more participants.” Same jobs, expanded scope, reduced administrative burden.

The most effective change management approach we observed follows a specific sequence: leadership uses the tools first and shares results publicly; small groups experience “aha moments” with personal productivity improvements; early skeptics are specifically invited to pilot programs (converting dissenters into advocates is more powerful than preaching to believers); and results are shared weekly in short, informal formats.

For CHROs and CIOs

AI change management is not a training program. It is a trust-building exercise. The single most effective tactic: have senior leaders publicly share how they use AI in their own work, including mistakes and learning curves. Vulnerability from leadership creates permission for adoption.



Pattern 6: The Agentic Frontier and the Human-in-the-Loop Imperative

The most forward-looking organizations are beginning to build agentic AI systems — autonomous workflows that execute multi-step processes without human intervention. But the most sophisticated thinkers recognize that pure autonomy is neither achievable nor desirable in regulated industries.

What We Observed

In healthcare claims processing, agentic systems can handle straightforward adjudication. But complex cases — rare disease conditions, multi-comorbidity patients, empathy-required member interactions — require human judgment. The emerging architecture is not “agent replaces human” but “agent retains ownership and requests targeted human input at specific decision points.”

This is fundamentally different from today’s chatbot escalation model, where ownership transfers entirely to the human. In the human-in-the-loop agentic model, the agent sends context and a specific request to a qualified expert, receives the judgment, incorporates it, and continues executing. The human provides judgment. The agent provides throughput.

This pattern requires three capabilities most organizations have not built: natural language routing to qualified experts based on case complexity; configuration of human expertise profiles (availability, specialization, authority levels); and audit trails that satisfy regulatory requirements for every human-agent interaction.

The Next Wave

Organizations that automate well will face a second challenge: knowing when not to automate. The competitive advantage shifts from “we deployed agents” to “our agents know when to ask for human judgment.” This is the design problem of the next 18 months.

Industry-Specific Observations

Healthcare Payers

- Regulatory and compliance requirements (BAAs, HIPAA) add months to any AI deployment timeline. Organizations that build compliance into their architecture from day one avoid costly retrofits.
- Provider data aggregation could be an ideal early use case for agentic AI — high volume, rule-based, with clear validation criteria.
- Contact center audio analytics POCs are emerging as the bridge between individual productivity tool capabilities and full agentic deployment.
- Annual Enrollment Period (AEP) content production and plan design submission are high-friction, repetitive workflows where AI delivers immediate measurable ROI.

Financial Services

- Board-level AI mandates are driving aggressive timelines that outpace organizational readiness. The gap between executive ambition and engineering capacity is widest in this vertical.
- Data commercialization (packaging internal data as products for clients or partners) is emerging as a strategic initiative alongside traditional AI use cases.
- Competition for specialized talent (GenAI architects, cloud data platform engineers, agentic workflow designers) is intense, with major technology companies and cloud providers offering different packages.

- Multi-vendor complexity across major enterprise platforms requires explicit integration architecture decisions early — organizations that defer these decisions accumulate technical debt exponentially.

Nonprofits

- Budget constraints force harder ROI conversations but also create innovation pressure. Nonprofits that adopt AI successfully tend to focus on a single high-impact use case (grant writing, donor communication, program matching) rather than broad deployment.
- AI literacy programs have an outsized cultural impact in mission-driven organizations where staff are motivated by purpose, not efficiency metrics.
- Website modernization and AI adoption are increasingly bundled because both improve participant/donor experience — creating natural entry points for strategic consulting.

SaaS and Technology Platforms

- The transition from point-solution analytics to a unified cloud data platform is a prerequisite for AI-powered product decisioning — and most SaaS companies have not made this transition.

What CIOs Get Wrong / Right: Observations

#	The Misconception	The Reality
1	"We deployed an enterprise AI productivity tool, so we have an AI strategy."	An enterprise AI productivity tool is a personal productivity layer. Enterprise AI requires data products, integration architectures, and use-case-specific implementations.
2	"We need to experiment with every framework to find the best	Framework proliferation creates paralysis. Pick 1–2 platforms per layer

	one.”	(customer-facing, internal, data) and build depth over breadth.
3	“Our data team can handle AI on top of their current workload.”	AI implementation requires dedicated capacity. Teams running BAU data operations cannot simultaneously build AI infrastructure without quality collapse.
4	“We’ll figure out governance after we launch the pilot.”	Ungoverned pilots become ungoverned production systems. Cost controls, access policies, and audit trails must be designed before activation, not after.
5	“The technology is the hard part.”	Change management, organizational trust, and executive alignment are the actual bottlenecks. The technology works. Getting humans to work with it is the challenge.

About Curate Partners

Curate Partners is a strategic consulting and talent firm specializing in AI adoption, data governance, and digital transformation for regulated industries. We operate as fractional C-suite leaders embedded in our clients' organizations, providing cross-industry pattern recognition that single-company executives cannot access.

Ready to Close the Gap Between AI Ambition and Execution?

The patterns in this brief are not theoretical — they are the challenges we help organizations solve every week. Whether you are standing up your first AI use case or trying to move from pilot to production, Curate Partners can help you:

- **Assess your AI readiness** with our Data Maturity Pyramid framework
- **Identify high-impact, low-risk use cases** tailored to your industry and organizational maturity
- **Build governance, cost controls, and change management** before you scale — not after
- **Accelerate from POC to production** with embedded, fractional leadership

Start the conversation.

Visit curatepartners.ai to explore our AI services or reach out directly to schedule a complimentary AI Readiness Assessment with our team.

Disclaimer: All product names, logos, brands, trademarks, and registered trademarks referenced in this document, including but not limited to Microsoft, Microsoft Office, Salesforce, and other third-party technologies, are the property of their respective owners and are used solely for identification and informational purposes; no affiliation with, endorsement by, sponsorship from, or partnership with the respective trademark holders is implied, and no claim of ownership or rights is asserted. This document is provided for informational and educational purposes only and reflects a technical and governance-oriented synthesis of publicly available regulatory texts, standards, and industry guidance; it does not constitute legal advice, regulatory advice, or compliance certification. Organizations remain solely responsible for assessing applicability to their specific circumstances and for engaging qualified legal, compliance, and regulatory professionals before making any deployment or operational decisions involving AI systems.

© 2026 Curate Partners. All rights reserved. This document is the proprietary intellectual property of Curate Partners and may not be reproduced, distributed, transmitted, displayed, published, or broadcast in whole or in part without the prior written permission of Curate Partners. Unauthorized use, copying, or dissemination of this material is strictly prohibited and may result in legal action.