databricks

Whitepaper

How to Deploy AI Agents in Financial Services

Most financial services institutions aren't scaling AI agents — not because models aren't smart enough, but because their architecture cannot provide the governed, contextual execution agents need. The problem is not more compute, but providing AI agents with the right enterprise context and controls, at the right time.

By: Kim Hatton, Junta Nakai, Jennifer Miller, Andrea DeSosa, Marcela Granados, Ashraf Safdar, Antoine Amend, and Raj Suresh

databricks

Table of Contents

Executive Summary: A Decision You Cannot Defer ... 3

Chapter 1: The Strategic Decision in Front of You ... 4
- The stakes ... 4
- Why most financial institutions will stall ... 4
- What separates leaders from laggards ... 5

Chapter 2: Your Three Paths Forward ... 6
- Why fragmented architectures fail in the agentic era ... 7

Chapter 3: Four Executive Capabilities, One Platform ... 8
- How this translates to board language ... 8

Chapter 4: Where the Complexity Tax Shows Up (And How to Remove It) ... 9
- Failure 1: Duplicated data multiplies risk ... 9
- Failure 2: Handoffs kill production velocity ... 9
- Failure 3: Inconsistent definitions erode executive confidence ... 9

Chapter 5: Real Institutions, Measurable Outcomes ... 11
- National Australia Bank — Foundation at Scale ... 11
- Banco Bradesco — Real-Time Personalization as an Operating Capability ... 11
- RBC Brewin Dolphin — Agentic Automation with Auditability ... 12

Chapter 6: The Leadership Playbook ... 13
- Action 1: Declare data a governed asset — with board-level ownership ... 13
- Action 2: Launch lighthouse programs — with P&L accountability ... 14
- Action 3: Treat governance as a growth lever, not a compliance cost ... 14

Conclusion: The Architecture Decision Is the Business Decision ... 15

Next Steps ... 16

EXECUTIVE SUMMARY: A Decision You Cannot Defer

CHAPTER 1: The Strategic Decision in Front of You

CHAPTER 2: Your Three Paths Forward

CHAPTER 3: Four Executive Capabilities, One Platform

CHAPTER 4: Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5: Real Institutions, Measurable Outcomes

CHAPTER 6: The Leadership Playbook

CONCLUSION: The Architecture Decision is the Business Decision

NEXT STEPS

Executive Summary: A Decision You Cannot Defer

Financial services leaders face a structural fork in the road. AI agents are moving from experimentation into production workflows — multi-agent orchestration grew 327% in under four months[1] — and the institutions that reach scale first will capture a durable cost, risk, and growth advantage. McKinsey estimates AI could reduce banks' aggregate cost base by 15–20%, with gross reductions of up to 70% in specific cost categories.[2] That is not a marginal efficiency gain. It is a cost-income ratio inflection.

But 95% of enterprise AI pilots fail to reach sustained production.[3] Forrester finds only 10–15% of AI projects scale beyond controlled environments.[4] The failure mode is not the model, it is the architecture underneath it.

The three questions every financial services executive must answer now:

1. **Cost base:** Can your current architecture deliver a 15–20% cost reduction, or does it compound the complexity tax every time you add a new AI capability?
2. **Risk exposure:** Can your governance and lineage framework scale to cover AI agents operating across workflows, or will your next regulatory examination expose gaps?
3. **Growth ceiling:** Are you building a platform that gets cheaper and faster with every use case, or one that requires rebuilding controls from scratch each time?

This guide is not a product overview. It is a framework for making the right architectural choice, and understanding the cost of getting it wrong.

> "The frontier is no longer model intelligence. It's whether you can make the AI you already have actually work inside your enterprise — by giving agents secure access to governed, contextual data and a reliable execution path."
>
> MCKINSEY

EXECUTIVE SUMMARY:
A Decision You Cannot Defer

CHAPTER 1:
The Strategic Decision in Front of You

CHAPTER 2:
Your Three Paths Forward

CHAPTER 3:
Four Executive Capabilities, One Platform

CHAPTER 4:
Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5:
Real Institutions, Measurable Outcomes

CHAPTER 6:
The Leadership Playbook

CONCLUSION:
The Architecture Decision is the Business Decision

NEXT STEPS

Chapter 1: The Strategic Decision in Front of You

The stakes

The complexity tax is not abstract. It shows up on the P&L, in audit findings, and in the gap between what your technology teams promise and what the business receives. IDC estimates companies lose 20–30% of annual revenue to data silo inefficiencies.[5]

Accenture's *2026 Banking Trend Report* finds that nearly 70% of bank IT budgets are consumed by maintaining legacy systems, with technology costs growing four times faster than banking revenues over the past 15 years.[6] Capgemini finds that more than 80% of banking executives are failing to see anticipated gains from AI investments, with only 29% of IT budgets allocated to transformative technologies.[7]

Meanwhile, McKinsey warns that customer adoption of AI agents could erode up to $170 billion in bank profits if institutions do not adapt.[8] The competitive risk of inaction is now measurable in profit pool terms.

Why most financial services institutions will stall

KPMG finds that while 99% of companies plan to put autonomous agents into production, only 11% have actually done so.[9] The top barriers are not talent or budget, they are architectural:

BARRIER	% OF RESPONDENTS CITING IT
Security and risk gaps	63%
Lack of interoperability across technology ecosystem	55%
Poor data governance	48%

Source: Forrester / AWS, cited in Neurons Lab Agentic AI in Financial Services (2026)[10]

Each of these barriers is a symptom of the same root cause: In a pre-agentic world, you could afford to plug these gaps with human handoffs. In an agentic world, those handoffs are exactly where programs die. The models are already capable; what's missing is a unified, governed data and execution layer that gives agents the context they need without multiplying risk.

EXECUTIVE SUMMARY:
A Decision You Cannot Defer

CHAPTER 1:
The Strategic Decision in Front of You

CHAPTER 2:
Your Three Paths Forward

CHAPTER 3:
Four Executive Capabilities, One Platform

CHAPTER 4:
Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5:
Real Institutions, Measurable Outcomes

CHAPTER 6:
The Leadership Playbook

CONCLUSION:
The Architecture Decision is the Business Decision

NEXT STEPS

What separates leaders from laggards

The Databricks 2026 State of AI Agents report, drawing on telemetry across thousands of enterprise deployments, is unambiguous on what determines production success:

- Companies actively using evaluation tools get nearly **6x more** AI projects into production
- Companies using AI governance put over **12x more** AI projects into production

This is not a compliance story. It is a competitive advantage story. Governance and evaluation are the production multipliers that convert pilot investment into P&L impact.

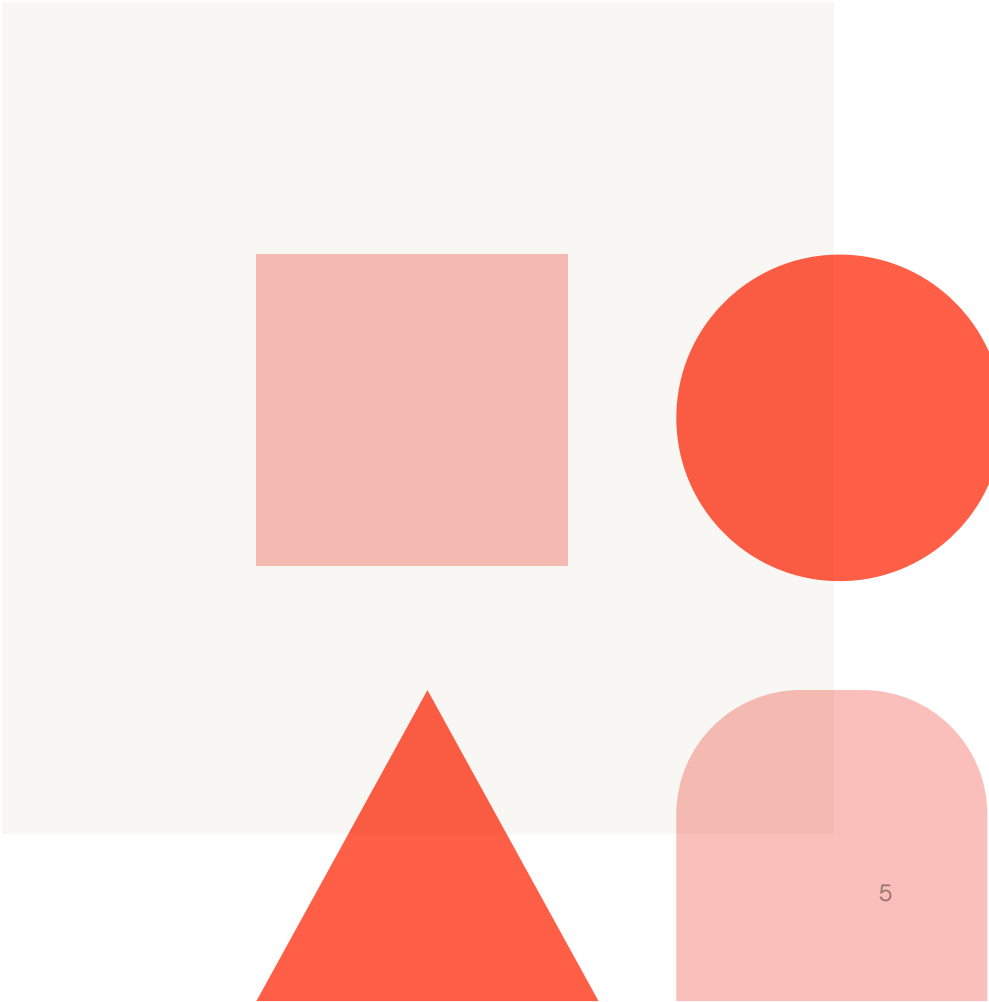

EXECUTIVE SUMMARY:
A Decision You Cannot Defer

CHAPTER 1:
The Strategic Decision in Front of You

CHAPTER 2:
Your Three Paths Forward

CHAPTER 3:
Four Executive Capabilities, One Platform

CHAPTER 4:
Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5:
Real Institutions, Measurable Outcomes

CHAPTER 6:
The Leadership Playbook

CONCLUSION:
The Architecture Decision is the Business Decision

NEXT STEPS

Chapter 2: Your Three Paths Forward

Before evaluating any platform, executives should be clear on the strategic options, and the real cost of each.

PATH	DESCRIPTION	SHORT-TERM DISRUPTION	LONG-TERM COST	PRODUCTION AI READINESS
Extend the Franken-stack	Add new AI tools on top of existing fragmented architecture	Low	Highest — complexity tax compounds with every addition	Low — governance gaps multiply as agents scale
Replatform fully	Rip and replace the entire stack with a new architecture	Highest	Potentially low — but rarely achieved	Uncertain — high execution risk, long time to value
Unify with a governance layer	Add a governed foundation that connects and extends existing investments without wholesale replacement	Low-to-medium	Lowest — eliminates redundancy without disruption	Highest — governance and lineage are built in from day one

The rational choice for most institutions is the third path. Databricks is designed to be that governance layer — the enterprise intelligence foundation that centralizes data in the Lakehouse, governs it with Unity Catalog and exposes it to AI agents through open formats and native agent frameworks like Agent Bricks.

EXECUTIVE SUMMARY:
A Decision You Cannot Defer

CHAPTER 1:
The Strategic Decision in Front of You

CHAPTER 2:
Your Three Paths Forward

CHAPTER 3:
Four Executive Capabilities, One Platform

CHAPTER 4:
Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5:
Real Institutions, Measurable Outcomes

CHAPTER 6:
The Leadership Playbook

CONCLUSION:
The Architecture Decision is the Business Decision

NEXT STEPS

In practice, this is an organizational knowledge store on the Lakehouse where schemas, queries, logs, dashboards and documentation are pre-computed into a governed memory layer that agents can query in real time.

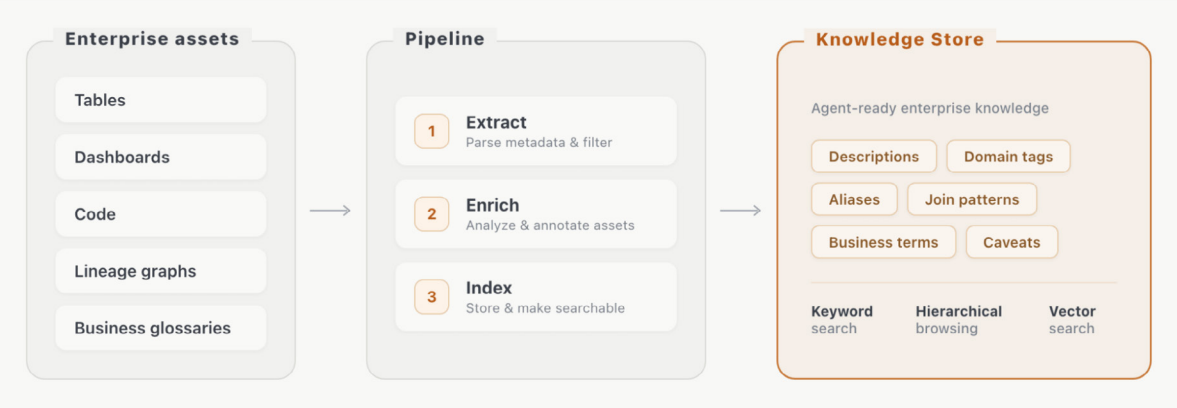

Building an organizational knowledge store from a collection of enterprise assets

Why fragmented architectures fail in the agentic era

Warehouse-centric stacks, hyperscaler-native stacks, and best-of-breed architectures all share the same structural weakness in an agentic world: they optimize one layer and create handoffs everywhere else. In practice, this means:

- **Duplication**
  Every new AI use case requires a new copy of data, a new pipeline, and a new permission model, multiplying cost and risk at each step.
- **Latency**
  Handoffs between storage, compute, model serving, and governance layers add operational friction that makes real-time agentic execution unreliable.
- **Governance gaps**
  When lineage and access control live in separate systems, audit and regulatory response becomes a manual, error-prone reconciliation exercise.
- **Fragile models**
  AI agents that draw from multiple uncoordinated data sources produce outputs that cannot be traced, validated, or defended in a model risk review.

The Databricks 2026 State of AI Agents report found that 80%+ of databases deployed in production are now being built or managed by AI agents, meaning the governance of those agents is inseparable from the governance of the data they touch. Architectures that treat these as separate concerns will accumulate risk faster than they can manage it.

EXECUTIVE SUMMARY: A Decision You Cannot Defer

CHAPTER 1: The Strategic Decision in Front of You

CHAPTER 2: Your Three Paths Forward

CHAPTER 3: Four Executive Capabilities, One Platform

CHAPTER 4: Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5: Real Institutions, Measurable Outcomes

CHAPTER 6: The Leadership Playbook

CONCLUSION: The Architecture Decision is the Business Decision

NEXT STEPS

Chapter 3: Four Executive Capabilities, One Platform

Rather than a component inventory, here is what Databricks delivers in terms executives fund:

EXECUTIVE CAPABILITY	WHAT IT MEANS	P&L / RISK IMPACT
One governed data plane	All data — structured, unstructured, real-time, historical — is accessible, governed, and AI-ready in one place	Eliminates duplicate pipelines; Unity Catalog reduces audit response from weeks to hours
One production path from model to workflow	The same platform that builds models also serves them, governs them, and embeds them in operational processes	Eliminates 6-month "productionization" handoffs; Forrester finds this is the primary reason 60%+ of pilots never reach production[11]
One control layer for audit and risk	Lineage, access control, and policy enforcement travel with every data asset and AI agent automatically	Reduces regulatory exposure; satisfies SR 11-7 model risk requirements without manual reconciliation
One foundation that compounds in value	Each new use case reuses governed data assets and agent components rather than rebuilding from scratch	NAB: use cases 2–10 are cheaper and faster than the first; Bradesco: 95% reduction in processing time

How this translates to board language

METRIC	INDUSTRY BENCHMARK	DATABRICKS-ENABLED OUTCOME
Cost-income ratio	McKinsey: AI can reduce aggregate cost base 15–20%[2]	A 15% cost base reduction on a 60% cost-income ratio bank translates to ~9 percentage points of C/I improvement
Time-to-production	Forrester: 60%+ of pilots never scale; average "productionization" takes 6+ months[11]	Unified platform eliminates the handoff; RBC BD deployed client review automation in weeks
Operational leverage	Capgemini: 80%+ of banking executives not seeing AI ROI[7]	Shared data foundation means marginal cost of each new use case declines, not compounds
Regulatory risk	SR 11-7 requires model lineage, validation, and audit trail for every AI model in production	Unity Catalog provides automated lineage and audit trail across all data, models, and agents

EXECUTIVE SUMMARY:
A Decision You Cannot Defer

CHAPTER 1:
The Strategic Decision in Front of You

CHAPTER 2:
Your Three Paths Forward

CHAPTER 3:
Four Executive Capabilities, One Platform

CHAPTER 4:
Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5:
Real Institutions, Measurable Outcomes

CHAPTER 6:
The Leadership Playbook

CONCLUSION:
The Architecture Decision is the Business Decision

NEXT STEPS

Chapter 4: Where the Complexity Tax Shows Up (And How to Remove It)

Three structural failures account for the majority of stalled AI programs in financial services.

Failure 1: Duplicated data multiplies risk

The symptom: Marketing has one customer view, risk has another, compliance has a third. When regulators ask for lineage, teams scramble to reconcile. Data governance failures cost financial services firms an average of 77% in financial losses and 55% in reputational damage per AI incident.[12]

> THE FIX: Unity Catalog provides one permission model, one audit trail, and one lineage graph, whether a human analyst queries a table or an AI agent accesses it for retrieval-augmented generation.

Failure 2: Handoffs kill production velocity

The symptom: A model works in the lab. Six months later it still isn't in production. Accenture's 2026 report identifies layered complexity — separate tools for experimentation, serving, operations, and governance — as the primary driver of technology costs growing four times faster than banking revenues.[6]

> THE FIX: The same platform that trains models also serves them, builds apps, and powers operational workloads. Governance travels with the pipeline. Existing tools connect through open formats rather than being displaced.

RBC Brewin Dolphin: 80–90% automation of annual client review meeting packs, 4,700 hours saved annually, and ~50% reduction in administrative costs in some offices, deployed in weeks, not quarters, with full Unity Catalog auditability.

Across these institutions, the common pattern is consistent. Success wasn't achieved by chasing the newest model, but by standardizing on a governed data and intelligence layer where agents, applications and analytics all operate from the same trusted context.

EXECUTIVE SUMMARY:
A Decision You Cannot Defer

CHAPTER 1:
The Strategic Decision in Front of You

CHAPTER 2:
Your Three Paths Forward

CHAPTER 3:
Four Executive Capabilities, One Platform

CHAPTER 4:
Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5:
Real Institutions, Measurable Outcomes

CHAPTER 6:
The Leadership Playbook

CONCLUSION:
The Architecture Decision is the Business Decision

NEXT STEPS

Failure 3: Inconsistent definitions erode executive confidence

The symptom: Marketing says conversion is up 15%. Finance says it is flat. Both are right, by their own definitions. Gartner and Forrester project that by 2026, autonomous governance modules will become a prerequisite embedded in core enterprise platforms, not managed as separate overlays.[13]

> THE FIX: Unity Catalog enforces consistent definitions across SQL, Python, AI agents, and BI dashboards. When a business term changes, it propagates everywhere automatically.

EXECUTIVE SUMMARY: A Decision You Cannot Defer

CHAPTER 1: The Strategic Decision in Front of You

CHAPTER 2: Your Three Paths Forward

CHAPTER 3: Four Executive Capabilities, One Platform

CHAPTER 4: Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5: Real Institutions, Measurable Outcomes

CHAPTER 6: The Leadership Playbook

CONCLUSION: The Architecture Decision is the Business Decision

NEXT STEPS

Chapter 5: Real Institutions, Measurable Outcomes

National Australia Bank — Foundation at Scale

NAB used Databricks to build a shared data foundation across the organization, one that scales use cases across teams without rebuilding pipelines, governance, or permissions for each one. The platform now supports use cases from marketing to financial crime detection to digital decisioning.

> **EXECUTIVE TAKEAWAY:** When the data foundation is shared, the next ten use cases get cheaper and faster than the first. This is operational leverage, the kind that improves cost-income ratios over time, not just in a single program.

Banco Bradesco — Real-Time Personalization as an Operating Capability

Bradesco built a real-time customer data platform using Spark Declarative Pipelines on the Databricks platform:

METRIC	OUTCOME
Processing time	**95% reduction**
CRM data maintenance	**87% reduction**
Contextual offer conversion	**2x improvement**

> **EXECUTIVE TAKEAWAY:** Real-time data on an open, governed foundation turns personalization from a marketing ambition into a measurable revenue lever.

EXECUTIVE SUMMARY: A Decision You Cannot Defer

CHAPTER 1: The Strategic Decision in Front of You

CHAPTER 2: Your Three Paths Forward

CHAPTER 3: Four Executive Capabilities, One Platform

CHAPTER 4: Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5: Real Institutions, Measurable Outcomes

CHAPTER 6: The Leadership Playbook

CONCLUSION: The Architecture Decision is the Business Decision

NEXT STEPS

RBC Brewin Dolphin — Agentic Automation with Auditability

RBC BD deployed GenAI workflows for client review preparation, automating the most time-intensive part of the advisory process while maintaining full traceability:

METRIC	OUTCOME
Client review pack automation	80–90%
Hours saved annually	4,700 (projected)
Administrative cost reduction	~50% in some offices
Auditability	Full lineage via Unity Catalog

> **EXECUTIVE TAKEAWAY:** This is how agentic systems earn institutional trust — automation plus governance, not automation instead of controls.

EXECUTIVE SUMMARY: A Decision You Cannot Defer

CHAPTER 1: The Strategic Decision in Front of You

CHAPTER 2: Your Three Paths Forward

CHAPTER 3: Four Executive Capabilities, One Platform

CHAPTER 4: Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5: Real Institutions, Measurable Outcomes

CHAPTER 6: The Leadership Playbook

CONCLUSION: The Architecture Decision is the Business Decision

NEXT STEPS

Chapter 6: The Leadership Playbook

Based on research from McKinsey, Accenture, and EY, leading institutions are converging on three concrete actions.

ACTION 1:

Declare data a governed asset — with board-level ownership

Form a cross-functional team to architect data with clear ownership, return expectations, and a 12–18 month roadmap for unified, real-time, and interoperable data. Accenture finds that banks redirecting legacy maintenance budget toward modern architecture are better positioned to absorb AI-native capabilities without rebuilding from scratch.[14] This is a capital allocation decision, not a technology decision.

The target state is an architecture where an agent's identity lives in its memory, not its model weights: skills, enterprise knowledge and episodic/semantic memories stored in a governed database like Lakebase, with LLMs acting as swappable reasoning engines on top.

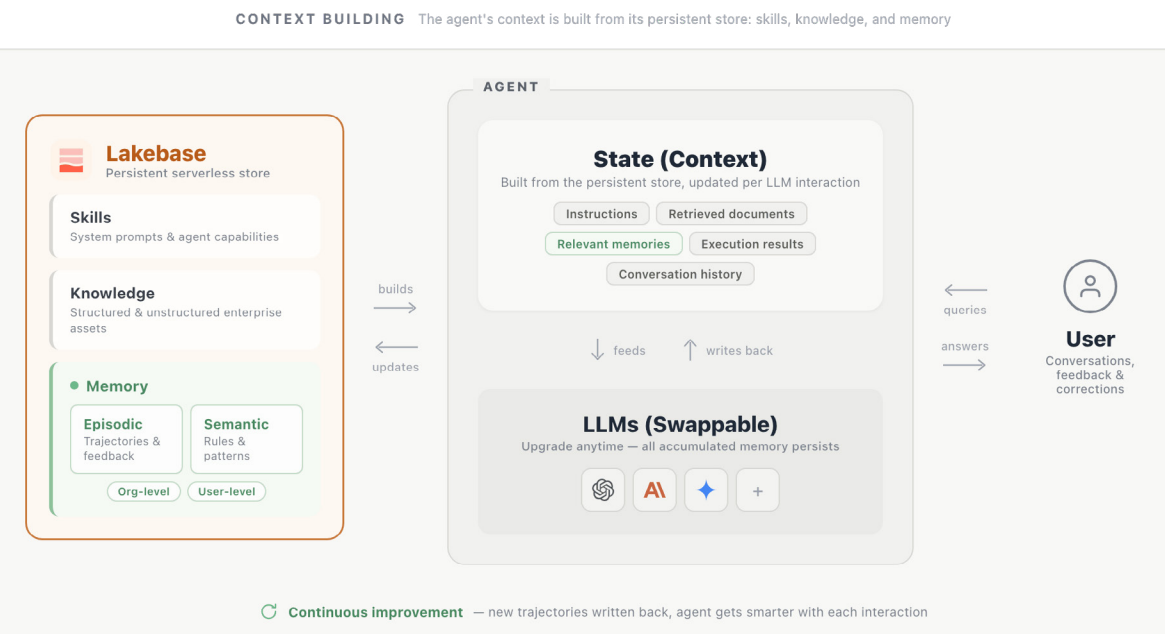

A memory-powered agent framework built on Lakebase.

ACTION 2:

Launch lighthouse programs — with P&L accountability

Establish an AI Center of Excellence to own agent standards, risk guardrails, and human-in-the-loop design. Pick 2–3 high-impact workflows (fraud triage, client reviews, claims FNOL). Baseline current cost and cycle time. Track value with the same rigor as capital investments. Gartner and Forrester confirm 2026 is the year of workflow automation breakthroughs in finance, the window for first-mover advantage is open but narrowing.[15]

ACTION 3:

Treat governance as a growth lever, not a compliance cost

A 2025 EY study reveals financial services firms are underinvesting in responsible AI by approximately 30%, even as 83% believe future AI regulations will support adoption.[16] *The Databricks 2026 State of AI Agents* data is definitive: organizations using governance put 12x more projects into production. That is not a risk management argument. That is a return on investment argument.

EXECUTIVE SUMMARY:
A Decision You Cannot Defer

CHAPTER 1:
The Strategic Decision in Front of You

CHAPTER 2:
Your Three Paths Forward

CHAPTER 3:
Four Executive Capabilities, One Platform

CHAPTER 4:
Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5:
Real Institutions, Measurable Outcomes

CHAPTER 6:
The Leadership Playbook

CONCLUSION:
The Architecture Decision is the Business Decision

NEXT STEPS

Conclusion: The Architecture Decision Is the Business Decision

The financial institutions that lead in 2027 will not be those with the biggest AI budgets. They will be the ones that made the right architectural choice in 2026, and built a foundation that compounds in value rather than complexity. BCG confirms that simplifying the technology stack is the primary lever available to banks for reducing run-the-bank costs and expanding capacity for innovation.[17]

The choice is clear:

- Extend the Franken-stack — and pay the complexity tax indefinitely
- Replatform fully — and absorb years of execution risk
- Unify with a governance layer — and build a foundation where every new use case gets faster and cheaper than the last

Databricks is that governance layer that connects existing investments through open formats, governing all data and AI assets in one place, and providing the production path that turns pilots into P&L impact. It is where your enterprise AI-powered models, agents, and analytics are actually built, governed and run on the data that already drives your boardroom decision.

The Bottom Line for the Board

What does this mean for my cost base?	McKinsey: 15–20% aggregate cost reduction potential; only achievable on a unified, governed architecture
What does this mean for my risk exposure?	Fragmented governance is the #1 reason AI programs fail regulatory scrutiny; Unity Catalog closes that gap structurally
What does this mean for my growth ceiling?	A compounding data foundation, where use cases 2–10 are cheaper than the first, is the only architecture that scales agentic AI profitably

EXECUTIVE SUMMARY: A Decision You Cannot Defer

CHAPTER 1: The Strategic Decision in Front of You

CHAPTER 2: Your Three Paths Forward

CHAPTER 3: Four Executive Capabilities, One Platform

CHAPTER 4: Where the Complexity Tax Shows Up (And How to Remove It)

CHAPTER 5: Real Institutions, Measurable Outcomes

CHAPTER 6: The Leadership Playbook

CONCLUSION: The Architecture Decision is the Business Decision

NEXT STEPS

Next Steps

FOR BANKING AND PAYMENTS LEADERS: Explore how unified customer data powers real-time fraud detection, personalized offers, and next-best-action recommendations that drive conversion and reduce loss ratios.

FOR CAPITAL MARKETS LEADERS: See how multi-agent systems accelerate investment research, automate surveillance, and deliver real-time risk analytics at scale, with full auditability.

FOR INSURANCE LEADERS: Learn how intelligent claims processing, smart underwriting, and predictive analytics transform customer experience while improving combined ratios.

Ready to reduce the complexity tax?

Contact your Databricks account executive or visit **databricks.com/financial-services** to explore solution accelerators, customer stories, and proven implementations.

Footnotes

[1] Databricks, *2026 State of AI Agents*

[2] McKinsey & Company, *Global Banking Annual Review 2025* (October 2025)

[3] NTT Data, *The $100 Billion Stack Gap* (September 2025)

[4] Forrester, *The State of AI 2025*, cited in Economic Times (January 2026)

[5] IDC Market Research, *The Cost of Data Silos*, cited in Defease with Ease (2022)

[6] Accenture, *Banking Trends 2026* (March 2026)

[7] Capgemini Research Institute, cited in CIO Dive / Banking Dive (March 2026)

[8] McKinsey & Company, *Agentic AI Will Shake Up Banking, Shrinking Global Profit Pools* (November 2025)

[9] KPMG, cited in Neurons Lab, *Agentic AI in Financial Services* (2026)

[10] Forrester / AWS, cited in Neurons Lab, *Agentic AI in Financial Services* (2026)

[11] Forrester, *The State of AI 2025*, cited in Economic Times (January 2026)

[12] Neurons Lab, *Agentic AI in Financial Services*, citing AI incident cost data (2026)

[13] Gartner and Forrester, *AI Adoption Predictions for Finance 2026*, cited on LinkedIn (January 2026)

[14] Accenture, *Banking Trends* 2026 (March 2026)

[15] Gartner and Forrester, cited on LinkedIn (January 2026)

[16] EY, cited in Neurons Lab, *Agentic AI in Financial Services* (2026)

[17] BCG, *Tech in Banking 2025: Smarter Tech Investment* (April 2025)