INDUSTRY PERSPECTIVE

AI: What Could Possibly Go Wrong with Fragmented Data at Machine Speed?

Why a Trusted, Real-Time Data Foundation is the Key Prerequisite for AI Success

By Ken Ballou
Founder, NewEnding LLC

SYNOPSIS

Enterprises spent an estimated $1.5 trillion on AI in 2025, yet the leading barrier to success is not model sophistication or compute capacity. It is data. This perspective argues that AI has not changed the fundamental challenges facing enterprise technology organizations, including fragmented data, disconnected systems, and siloed decision-making. It has simply removed the human buffer that once absorbed them.

Drawing on three decades in enterprise software and financial technology, Ken Ballou examines how autonomous decision-making at machine speed turns longstanding data-quality weaknesses into strategic enterprise risk. Through three financial services use cases, payment routing, fraud and AML, and governance, risk and compliance, he makes the case that a trusted, real-time operational data foundation is now a prerequisite, not a supporting initiative, for AI to deliver competitive advantage rather than accelerate existing weaknesses.

Introduction

I recognize that the title of this paper may seem provocative at a time when nearly every technology conversation begins and ends with Artificial Intelligence.

Yet after more than three decades in enterprise software, financial technology, and data-driven business transformation, I keep arriving at the same conclusion: AI has not fundamentally changed the challenges facing enterprise technology organizations. It has accelerated them.

The issues that have constrained enterprises for years, including fragmented data, disconnected systems, inconsistent records, duplicate processes, and siloed decision-making, remain firmly in place. What has changed is speed. AI now consumes, analyses, and acts on information fast enough that the impact of those longstanding weaknesses is magnified rather than absorbed.

As organizations invest heavily in AI, many are discovering that the primary barrier to success is not model sophistication, compute capacity, or algorithm design. The barrier is data. More precisely, it is the inability to put trusted, consistent data in front of a decision at the moment that decision is made.

The AI Readiness Gap

In conversations with technology leaders across banking, payments, insurance, and fintech, a few themes recur:

- AI has become the single most important technology investment category for many organizations.
- At the same time, decades of application-specific buying decisions have produced highly fragmented environments. What was once prudent vendor diversification has become an ecosystem of disconnected systems, redundant data stores, and competing versions of operational truth.

The financial scale of that gap is now measurable. Enterprises spent an estimated $1.5 trillion on AI in 2025, yet **data quality is consistently ranked the leading barrier to AI success, ahead of model accuracy, compute cost, and talent (Gartner, 2025).**

The pattern is consistent: successful AI depends less on acquiring new AI capability and more on resolving long-standing data and decisioning challenges. **In many cases AI is simply exposing problems that were survivable while humans still made most operational decisions.** As those decisions move to autonomous systems, data quality, consistency, and availability stop being technical details and become strategic business concerns.

Connecting Systems Are the Key, But That Alone is Not Enough

In the past I've talked to the idea that we are now increasingly living in the "Connected Systems Economy" where enterprise integration focuses on not just moving information between systems but correlating it to fit the demands of application architectures and use cases. This objective is still the key, but success in the AI paradigm imposes a more demanding requirement.

Data must now be not only connected but continuously consistent, validated, and available in real time at the point of decision. A payment authorization, fraud assessment, compliance check, or underwriting recommendation may depend on information from dozens of internal and external systems at once. When those systems hold conflicting information, delayed updates, or incomplete records, an AI system can reach a decision that is technically correct given the data available to it, yet operationally wrong in practice.

The challenge, then, is not simply moving data faster. It is establishing a trusted operational data foundation capable of supporting correct decisions at machine speed.

Use Case 1: Payment Processing and Intelligent Transaction Routing

For payment processors, payment service providers, acquiring banks, fintech platforms, and embedded-finance providers, transaction routing has moved well beyond infrastructure plumbing. It is now a direct driver of revenue, margin, customer experience, and competitive differentiation.

As volumes grow, the next generation of payment platforms will increasingly rely on AI agents that continuously optimize acquiring-bank selection, authorization success rates, processing costs, network utilization, fraud-review workflows, and customer outcomes. These decisions often occur within milliseconds, and their quality depends entirely on the data available at the instant the decision is made.

When transaction histories, merchant data, customer profiles, and risk indicators are fragmented across systems, routing decisions become less reliable. The result is lower optimization accuracy, higher operational risk, and lost revenue. Organizations that maintain a continuously consistent view of transaction activity gain a real advantage as payment ecosystems become more autonomous.

Use Case 2: Fraud Detection and Anti-Money Laundering

Fraud prevention is now one of the most demanding real-time decision environments in financial services. Modern threats, including synthetic identity fraud, account takeover, authorized push-payment scams, mule networks, and cross-channel patterns, operate at machine speed.

Effective detection requires current customer information, transaction history, behavioural signals, device intelligence, sanctions data, and external risk indicators, all available at once. In many institutions these elements remain spread across platforms, so fraud models operate with incomplete context.

The consequences cut both ways. A missed fraud event produces immediate loss. A false positive creates friction, erodes trust, and drives customers away. **The scale of the second problem is often underestimated: false declines, legitimate transactions wrongly rejected, cost North American e-commerce an estimated $81 billion in permanently lost revenue each year, a larger figure than the cost of actual fraud** (PYMNTS Intelligence). A decision that is technically defensible can still be operationally expensive.

AML operations face the same structural issue. Suspicious-activity detection depends on correlating information across systems, business units, and channels. Without a unified, trusted view of customer and transaction activity, institutions carry higher compliance risk and higher investigative cost. As AI takes on more of this work, the quality of the underlying data increasingly determines the quality of the decision.

Use Case 3: Governance, Risk, and Compliance

The most overlooked AI challenge may sit inside Governance, Risk, and Compliance. Traditional GRC programs were built around periodic reviews, manual controls, and human oversight. AI changes that operating model.

As AI participates in lending, onboarding, payment authorization, fraud investigation, sanctions screening, and regulatory reporting, decision velocity accelerates sharply. What once took hours or days now takes milliseconds. **The difficulty is that AI inherits the same data-quality and integration limits enterprises have lived with for years.**

When customer, transaction, and compliance information is fragmented, AI can act on incomplete or conflicting data, creating a new category of risk in which decision velocity outpaces governance visibility.

This is not a theoretical concern. Gartner projects that by 2027, most organizations will fail to realize the value of their AI use cases because of fragmented, incohesive data governance. Regulators, meanwhile, increasingly expect institutions to demonstrate how a decision was made, which data informed it, and whether it can be reproduced and explained. Without consistent data and visibility across systems, those expectations become progressively harder to meet.

In many respects AI is not creating new governance problems. It is exposing existing ones at unprecedented speed.

Consider several examples:

Regulatory Reporting Risk

Financial institutions routinely generate regulatory reports by aggregating information from dozens or hundreds of internal systems.

When underlying data definitions differ across those systems, organizations often rely on reconciliation processes to establish consistency before reporting.

AI-driven operational systems may consume and act upon these same inconsistent data sources long before reconciliation occurs, potentially creating situations where operational decisions and regulatory reporting are based upon different versions of the truth.

Model Governance Risk

Regulators increasingly expect organizations to demonstrate explainability, lineage, and accountability for AI-driven decisions.

When training data, transaction data, and operational data reside in disconnected repositories, establishing complete lineage becomes difficult.

Organizations may struggle to answer fundamental governance questions:

- Which data sources influenced a decision?
- Was the data current at the time of the decision?
- Which version of the model was used?
- Can the decision be reproduced during an audit or regulatory examination?

Without integrated and correlated enterprise data, explainability becomes increasingly difficult as AI adoption expands.

Operational Risk Amplification

Historically, data inconsistencies often surfaced during batch processing cycles, reconciliations, or audits.

AI systems dramatically reduce the time available for detection.

An incorrect customer attribute, outdated risk score, or missing compliance flag can now propagate through thousands of automated decisions before human operators become aware of the issue.

The underlying problem is not the AI model itself. The problem is the speed at which existing data-quality issues are amplified.

Third-Party and Ecosystem Risk

Modern enterprises increasingly depend upon external partners, fintech providers, cloud services, payment processors, data vendors, and embedded-finance ecosystems.

Each participant introduces additional data sources, data models, and integration points.

As AI systems begin making decisions across these extended ecosystems, maintaining a consistent and trusted view of customers, transactions, counterparties, and risk indicators becomes increasingly difficult.

Organizations that cannot correlate data across internal and external environments face growing exposure to compliance failures, operational disruptions, and regulatory scrutiny.

The Governance Imperative

The emerging challenge for enterprise leaders is not simply deploying AI responsibly. It is establishing the foundational data infrastructure required to govern AI responsibly.

Before autonomous systems can be trusted to make increasingly consequential decisions, organizations must first establish the ability to continuously correlate, synchronize, and validate data across the enterprise.

Only then can governance frameworks provide the transparency, auditability, and control necessary to support AI operating at enterprise scale.

The Emerging Architecture Requirement

A clear pattern is emerging across industries. Organizations that successfully scale AI are investing in architectures that can:

- Correlate data across disparate systems
- Maintain operational consistency across environments
- Synchronize information in real time
- Support high-volume transaction processing at low latency
- Preserve data lineage and auditability
- Provide a trusted operational view of enterprise activity

These are becoming foundational requirements, not optional enhancements. **The goal is not simply faster data movement. It is a continuously trusted data environment capable of supporting correct, autonomous decision-making at scale.**

Conclusion

The AI revolution is not exposing new weaknesses in enterprise architecture. It is exposing weaknesses that have existed for decades.

For years, fragmented data, disconnected applications, inconsistent records, and siloed systems remained manageable because humans stayed in the decision loop. Reconciliation, audits, and manual intervention provided enough time to catch and correct inconsistencies. AI removes that buffer.

As autonomous systems take on payments, fraud prevention, compliance monitoring, and underwriting, execution speeds begin to rival the volumes once associated with financial trading platforms. At those speeds, data-quality issues stop being operational inefficiencies and become enterprise risks.

Organizations pursuing aggressive AI strategies should treat a trusted, real-time operational data foundation not as a supporting initiative but as a prerequisite. **The ability to put correct, consistent data in front of a decision, at the moment it is made, will increasingly determine whether AI delivers competitive advantage or simply accelerates existing weaknesses.**

In that sense, nothing has changed. The challenge remains what it has always been: ensuring that the right decision is made at the moment it matters, and that it can be recorded, explained, and defended. The only difference is that now everything is moving much faster.

For technology leaders, the practical question is no longer whether to invest in AI. It is whether the data foundation beneath it can make decisions correctly, and fast enough to matter.

Sources

- Gartner: enterprise AI spend reached an estimated $1.5 trillion in 2025; data quality ranked the leading barrier to AI success (2025).
- PYMNTS Intelligence: false declines estimated at $81 billion in permanently lost North American e-commerce revenue annually, exceeding losses from actual fraud.
- Gartner: prediction that 60% of organizations will fail to realize AI value by 2027 due to fragmented data governance.

ABOUT THE AUTHOR

Ken Ballou

Founder, NewEnding LLC

Ken Ballou is a senior technology executive, entrepreneur, and investor with more than three decades of experience in enterprise software, financial technology, and data-driven business transformation.

He is the founder of NewEnding LLC, a business development and advisory firm providing growth strategy and management guidance to organizations ranging from early-stage ventures to Fortune 500 enterprises.

Across his career, Ken has held senior sales, business development, and leadership roles at a range of leading technology companies, including Palantir Technologies, Sensedia, Camunda, Alfresco, Hewlett Packard Enterprise, and CA Technologies, with deep specialization in banking, financial services, and insurance. His work consistently focuses on helping enterprises connect, correlate, and operationalize their data to unlock measurable business outcomes.

He writes and speaks regularly on enterprise data strategy, the practical realities of AI adoption, and the architectural foundations required to support decision-making at machine speed.

https://www.linkedin.com/in/kenballou/