

EXECUTIVE BRIEF

The \$81 Billion Decision

Why Real-Time Data Decides Whether Your AI Investment in Payments and Fraud Pays Off

Adapted from the Industry Perspective by Ken Ballou

Founder, NewEnding LLC · Presented by Volt Active Data, Fintech South 2026

SYNOPSIS

Enterprises spent an estimated \$1.5 trillion on AI in 2025, yet the leading barrier to success is not the models. It is the inability to put trusted, consistent data in front of an AI-driven decision at the precise commitment point. Payment processing and fraud detection are inextricably linked, and this failure has a measurable price: false declines alone cost North American e-commerce an estimated \$81 billion in permanently lost revenue every year, more than the cost of actual fraud.

This brief condenses the full Industry Perspective and accompanies the Volt Active Data session at Fintech South 2026. It sets out why the fix is a trusted, real-time operational data foundation, and the three things to do about it.

The Business Case

In execution, AI models surface a new challenge: accelerated decisions without equally accelerated data correlation and currency create a sequencing problem, and that sequencing problem gives new life to both fraud and false declines. In the quest for faster decisions and higher throughput, AI has removed the human buffer that used to offset fragmented data, disconnected systems, inconsistent records and siloed decision-making. Those problems constrained enterprises for years, but they were survivable provided humans stayed in the decision loop, because reconciliation, audits and manual intervention created the time needed to catch and correct inconsistencies. Autonomous decisioning at machine speed has eliminated that time advantage.

The scale of the gap is measurable. Gartner ranks data quality as the leading barrier to AI success, ahead of model accuracy, compute cost and talent, and further projects that by 2027 most organizations will fail to realize the value of their AI use cases because of a range of issues tied to fragmented, incohesive data and governance practices. The pattern across banking, payments and fintech is consistent: **successful AI depends less on acquiring new AI capability and more on resolving long-standing data and decisioning challenges.**

A payment authorization, fraud assessment or compliance check may depend on information from dozens of 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 its execution model and the data available to it in that moment, yet financially and operationally wrong in fact. That sequence misalignment is exactly where opportunity and profit are lost.

Lost Opportunity and Fraud: The Two Costs of a Bad Decision

Cost one: the fraud you miss. US consumers lost \$12.5 billion to fraud in 2024. In card and payment fraud, the window between submission and settlement is the only moment a block is meaningful. Once the transaction clears, the money has moved, and on instant rails the settlement is irreversible. A fraud decision made after the transaction clears is the same as a wrong decision, and this exposure is growing at roughly 25% a year.

Cost two: the customers you decline. The larger and quieter cost is false declines: legitimate transactions wrongly rejected because the model scored against stale or incomplete context. At an estimated \$81 billion a year, false declines cost more than the fraud itself. A decision that is technically defensible can still be operationally expensive.

The third cost compounds both: governance. Regulators increasingly expect institutions to demonstrate which data influenced a decision, whether that data was current at the time,

which version of the model was used, and whether the decision can be reproduced under examination. When decision velocity outpaces governance visibility, an incorrect attribute or outdated risk score can propagate through thousands of automated decisions before anyone notices.

Across the three environments the full Industry Perspective examines, payment routing, fraud and AML, and governance, risk and compliance, the diagnosis is the same. The model is rarely the problem. Stale, fragmented state is.

The Fix: Decisions Made Inside the Window

Every consequential financial decision has a window, and the window is defined by the payment rail or the risk event, not by the institution: the card authorization hold, the instant payment clearing cycle, the moment before a trade executes. The requirement that follows is precise.

The state a decision reads, the logic that produces it, the enforcement that acts on it, and the auditable record that proves it must execute together, in a single transaction, at the precise commitment point.

Anything less recreates the gap. When state lives in one database, scoring arrives over an API and enforcement travels through a queue, the decision assembles across systems while the transaction settles underneath it. Connecting systems faster does not close this; the data must be continuously consistent, validated and available at the point of decision.

Where agentic AI belongs: the 90/10 Rule. Roughly 90% of operational financial decisions are deterministic: the rules are known, the data requirements are known, and the winning attribute is speed with proof. Deterministic decisioning should own that 90% inside the window. Agentic AI belongs in the ambiguous 10%, the novel patterns and judgement calls, operating under human governance and reading the same live, consistent state. Institutions that invert this, putting probabilistic agents in the deterministic path, add latency and unpredictability exactly where neither can be afforded.

What this buys. Fraud is blocked before the money moves, and false declines fall because scoring runs on current context rather than a snapshot. Payments route on current truth: authorization success, cost and network conditions as they are now, not as they were at the last batch run. And every decision is recorded with its full context in the same transaction, so it can be explained, reproduced and defended without a reconciliation exercise.

What to Do Next

Three actions, drawn from the Volt session at Fintech South:

- **Map your decision windows.** List the decisions in your stack that have a hard window, fraud blocks, payment authorizations, risk limits, and identify which rail or obligation defines each window.
- **Audit the gap.** For each windowed decision, trace where state, logic and enforcement live today and count the hops between them. Every hop is latency, and every gap between systems is where the costly mistakes get created.
- **Put AI where it belongs.** Apply the 90/10 Rule before scaling agents: deterministic decisioning inside the window for the 90%, agentic AI on the ambiguous 10% under governance, both reading the same live state.

Volt Active Data: The Decisioning Layer

- Decisioning layer latency: sub-10ms. Throughput: 1M+ TPS with linear scaling.
- Accuracy: full ACID, with the decision and its context written atomically as the outcome record.
- Availability: 99.999%, zero-downtime failover.

In production: a Tier-1 bank moved P99 validation latency on MiFID II transaction reporting from 150ms to under 12ms, making compliance a real-time system property rather than a post-trade reconstruction exercise and cutting audit reconciliation effort by more than 70%; a Tier-1 global bank cut fraud 83% executing 2,000+ rules atomically inside the authorization window; a global crypto exchange runs ACID-guaranteed decisioning across billions of dollars in daily volume.

Go deeper. This brief condenses Ken Ballou's full Industry Perspective, *AI: What Could Possibly Go Wrong with Fragmented Data at Machine Speed?* Ask the Volt team for a copy, or visit voltactivedata.com/fintech-south.

Sources: Gartner, enterprise AI spend and data-quality barrier, 2025, and fragmented data governance prediction for 2027. PYMNTS Intelligence, false declines in North American e-commerce. FTC Consumer Sentinel Network Data Book 2024. Volt Active Data published case studies and benchmark results.