


Why a Late Decision Is a Wrong Decision

The Decisioning Layer in Financial Services

Fraud models are accurate. Payment rails are instant. AI investment is real. But between detection and enforcement, most FS architectures have a gap: no layer that turns a score into an authoritative, auditable decision while the transaction is still waiting.

FOR PAYMENTS, FRAUD, AND DATA LEADERS



TRANSACTION SCORED

The Fraud Losses Are Real. So Is the Cost of a Wrong Decision.



Detection has never been better. Losses have never been higher.

Financial Institutions and Their Architects Are Facing the Same Structural Gap From Different Sides.

FOR FINANCIAL INSTITUTIONS

- ⌚ A late score is an audit record, not prevention
- ⚖️ Wrong calls cost twice: fraud missed, customers declined
- ⚠️ Balances go negative under load; caught hours later
- 📁 Regulators require explainability the stack can't give


FOR ISV & SI ARCHITECTS

- 📁 One system holds the state, scoring, and enforcement
- 🔄 Concurrent reads race the counters fraudsters exploit
- 🔗 Four-component stacks leak at scale
- 📄 No complete trail of what was queried, scored, decided

The shared root cause: detection is strong; decision authority is missing.

THE GAP

The Rails Went Instant. The Mandates Followed. Most Architectures Have Not Caught Up.



THE REAL-TIME COMPLIANCE CHASM

Instant rails settle in seconds and are irreversible. Every enrichment, sanctions check, fraud score, routing decision, and balance verification must complete inside that window. Batch decisioning is now a compliance gap, not just a performance gap.

DEADLINE

What Authoritative Decisioning Looks Like in Production.

83%

A leading global technology vendor's fraud platform
Fewer fraud losses, in-line within the authorization window, sub-50ms scoring.

Published results →

12ms

A tier-1 UK bank
P99 down from 150ms to 12ms, MiFID II hub, 120,000+ events/sec, six-plus systems to one.

Published results →

3x

A global exchange trading company
Faster risk response; sub-5ms P99 trading latency at 50,000 TPS peak.

Published results →

Names anonymized at the customer's request. Each claim is verifiable against its published case study.

AI-Driven Finance Requires Three Layers. Most Architectures Have Two.

Cloud Analytics and AI Platforms

Model training, risk analytics, regulatory reporting

↓ SCORES IN

Real-Time Decisioning Layer

VOLT

Authoritative operational state with ACID guarantees. ML scores evaluated deterministically against current state via in-line inference. Decisions recorded atomically. Complete audit trail per decision.

Scores in. Authoritative decisions out.

Sub-10ms platform decision latency

↓ DECISIONS OUT

Data Foundation

Core banking, payment switches, market data, downstream execution

AI recommendations are only as safe as the state they are evaluated against; the decisioning layer is what makes agentic AI deployable in production finance.

Four Functions. One Execution Path. No Gap Between Score and Enforcement.



Ingest

Payments, auth, ML scores arrive continuously



Maintain

Balances, limits, sanctions state, held ACID-safe



Evaluate

Rules and scores evaluated where state lives



Record

Allow, block, route. Written atomically.

1M+ TPS | Sub-10ms under production load | ACID at scale | Active-active geo-resilience |

Zero-downtime upgrades

Read the Data. Explore the Architecture.

Two ways in: the executive brief behind the \$81B false-decline number, or the Dev Hub where the decisioning layer runs live. No form. No sales call.

Read the Executive Brief

Explore the Dev Hub

VOLT ACTIVE DATA

Sources: FTC testimony to the US Congress Joint Economic Committee (2026); PYMNTS Intelligence; FinCEN; Federal Reserve Financial Services; UK Finance; NPCI / RBI.