Customer Case Study:

Real-Time Fraud Decisioning at Global Scale

How a Global ICT Provider Embedded Volt Active Data to Stop Card Fraud in Single-Digit Milliseconds

About the Customer

The provider is a leading global information and communications technology (ICT) solutions provider and consumer electronics leader, headquartered in APAC and operating in more than 170 countries and regions. Employee owned and founded in 1987, it employs roughly 208,000 people, around 113,000 of them in research and development, and reinvests more than 20 percent of annual revenue (approximately US$118 billion in 2024) into R&D, holding over 150,000 active patents. Its enterprise Big Data and analytics platform serves enterprises and retail financial services institutions across APAC.

The end customer is one of APAC's largest commercial retail banks and ranks among the world's top ten by Tier 1 capital. It holds total assets of approximately US$1.7 trillion, employs nearly 121,600 people, and operates one of the region's largest credit card portfolios, with close to 96 million active cards in circulation. That card base generates the high-volume payment traffic where every authorization must be evaluated for fraud within a strict latency budget. The provider embedded Volt Active Data into their Real time Fraud Prevention Solution to deliver fraud decisioning as a real-time system property: packaged, repeatable, and deployable across its banking customer base.

Note on confidentiality: *Customer identities have been anonymized at the request of the parties. Metrics and architecture reflect a live production deployment.*

Overview: Catching Fraud on the First Misused Card

Every time a customer swipes, inserts, taps, or scans a card, the bank has only milliseconds to decide whether to authorize, decline, or step up the transaction to two-factor authentication. Get it wrong by approving fraud, and someone has to absorb the loss. Get it wrong by declining a legitimate transaction, and the cardholder experience suffers and trust erodes.

The art is catching fraud on the very first misused payment method, before an account is flagged and added to a blacklist. That demands a decision made against accurate, current state, applied deterministically, and recorded as authoritative truth, all within the payment-authorization window. This is not just an analytics problem. It is a real-time decisioning problem.

Recognizing the value of combining fast data with big data, the provider sought a foundational engine to expand its analytics platform with real-time services. After evaluating many databases, it selected Volt Active Data as the high-velocity decisioning layer capable of running hundreds of fraud check rules per transaction and returning a decision in single-digit milliseconds, at a scale and cost no other system could match.

The Challenge: Deterministic Fraud Decisions Within a 200ms Budget

Real-time, in-event fraud analysis means evaluating each transaction before it is authorized, not detecting it after settlement. To do this, the platform needed a system that could perform hundreds to thousands of non-trivial queries per financial transaction, applying user-provided logic and rules against gigabytes to terabytes of recent history, blacklists, and other state-based data.

Because the capability would be packaged into the provider's analytics product and shipped to banking customers, it also needed deployment flexibility, manageable cost, and financial-grade security. Several constraints made this exceptionally hard:

- **A hard latency budget.** Banks allow up to 200 ms for a card-swipe decision. Within that window the system must run every fraud rule and model and return an answer, every time, not on average.
- **Extreme query density.** Each swipe requires multiple transactions and roughly 1,000 queries per transaction. At peak, the bank's traffic ranges from hundreds of thousands of transactions per second toward a million.
- **"Fast but inconsistent" is unacceptable.** A fraud decision made on stale or approximate state is a wrong decision. Behavior cannot degrade as load increases, and the same conditions must always produce the same outcome.
- **Rules that change intraday.** Fraud logic must be updatable as often as every few minutes, with the engine transitioning to entirely new logic transactionally, without impacting load or correctness.
- **Continuous availability is non-negotiable.** Card payments operate around the clock, and fraud controls cannot become a point of failure. The system must deliver deterministic decisions 24×7, surviving hardware failures, software upgrades, and traffic spikes without downtime, degraded accuracy, or interruption to transaction processing.

Traditional RDBMS platforms would have cost more, demanded more hardware, and required multiple architectural layers to achieve the same result. NoSQL and other NewSQL systems could not offer the query power needed to sustain that level of analysis at those rates. The fraud decision was emerging too late, from too many places, in an inconsistent state.

The Solution: Volt as the Real-Time Decisioning Layer

Volt Active Data was embedded as the Real-Time Database (RTD) component inside the provider's analytics platform, becoming the deterministic decisioning layer that turns a stream of card-swipe signals into authoritative allow, reject, or step-up decisions. Credit-card swipes arrive as a continuous stream of data over message queues. Volt maintains the authoritative operational state (recent history, balances, velocity, blacklists) and decides what should happen next, in the moment of execution.

User-provided rules and fraud detection logic are compiled into Volt stored procedures and loaded alongside dimension data such as blacklists. With the scoring procedures in place, the platform is ready to evaluate fraud against current, committed state rather than a delayed approximation.

1. In-Event Decisioning: Four Transactions Per Swipe

For every card swipe, Volt processes multiple transactions for fraud analysis: such as to log and record the event as authoritative truth, and reads as a part of fraud-detection rules and Anti-fraud check ML models. Each transaction invokes roughly 1,000 queries. Based on the combined rule outcomes and model score, the transaction is then allowed, rejected, or sent for two-factor authentication.

2. Deterministic Rules at Scale

Volt solves the core problem of running transactions through complex rule sets, hundreds of rules, at the lowest possible latency and at scale. Decisions are made atomically against current state and applied deterministically, so the same inputs always produce the same outcome. This is the difference between detection and decision authority: the platform does not merely flag risk, it makes and records the authoritative decision the payment system enforces.

3. Intraday Rule Dynamics

Fraud logic can be changed as often as every few minutes. Volt accepts and transactionally transitions to entirely new logic without impacting load time, so the bank can adapt to emerging fraud patterns without taking the decisioning layer offline.

4. Why Volt Active Data

The bank's decision came down to four properties that had to hold simultaneously:

- **Speed.** All four transactions, close to 4,000 queries per swipe, complete well inside the 200 ms budget, the Decision latency i.e. P99 latency at Volt is sub-10 ms and end-to-end responses are under 50 ms.
- **ACID.** Every decision is atomic, consistent, isolated, and durable, even under sustained load. Correctness does not degrade as scale increases.
- **Programmability.** Complex, evolving fraud logic is expressed as stored procedures that execute where the data lives, with no service hops and no external caches.
- **XDCR.** Active-Active Cross Data Center Replication keeps decision authority available and correct across regions, so a regional failure becomes capacity loss, not downtime.

Architecture Overview

Volt sits at the center as a fast-data layer in this architecture. Card-swipe transactions stream into the real-time decisioning tier, where a business rules engine evaluates each event against authoritative state and returns a decision in real time. A big-data tier (Spark / Hadoop) handles machine-learning model training, and updated rules are pushed back into the decisioning layer intraday, closing the loop between insight at rest and decisions in motion.

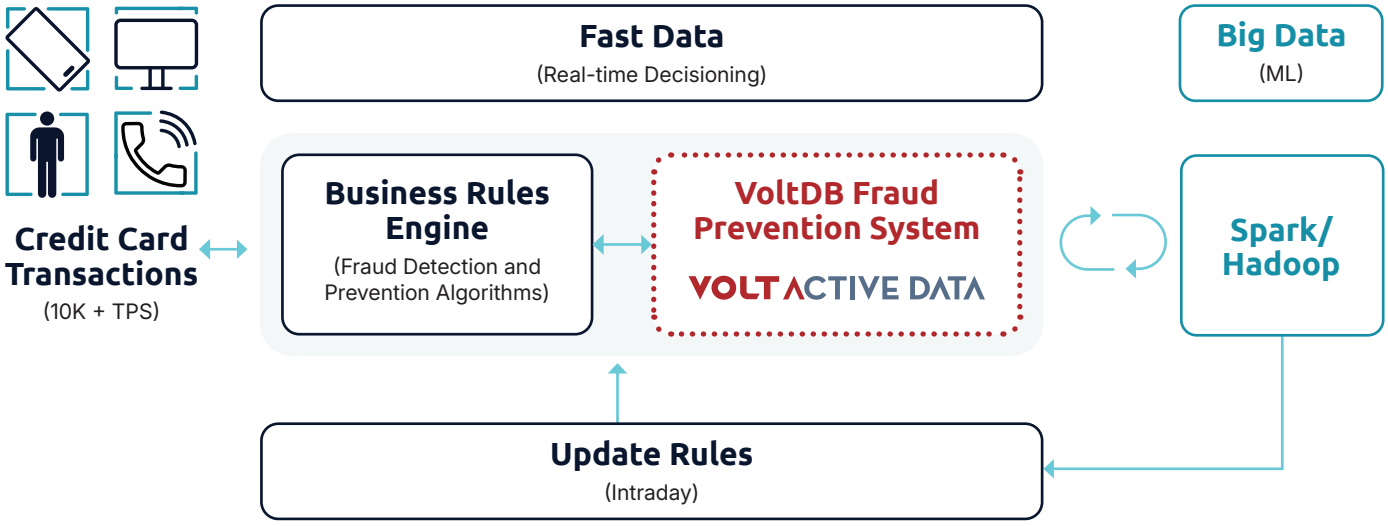

Figure 1: Real-time fraud decisioning architecture. Fast data for in-event decisions, big data for model training.

The production deployment spans the following footprint for VoltDB:

Deployment Metric	Production Value
Cluster Size	15 nodes per cluster
Total Clusters	10 clusters
Cross Data Center Replication (XDCR)	Active (Active-Active)
Sustained Throughput	70,000 TPS
Decision Latency	< 10 ms per transaction
Bank Latency Budget vs. Volt Response	200 ms budget, answered in under 50 ms

The Results: Fraud stopped before it completes

With Volt Active Data embedded as its decisioning layer, the provider delivered measurable, material improvements for its banking customer:

83%

REDUCTION IN FRAUDULENT TRANSACTIONS

By deciding in-event rather than reconciling after settlement.

10x

INCREASE IN PROCESSING CAPACITY

On a compact footprint, sustaining 70,000 TPS across the deployment.

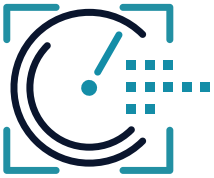

<50 ms

DECISIONS END-TO-END

against a 200 ms budget, sub-10 ms per transaction, collapsing the gap from ingestion to decision from minutes or hours to milliseconds.

HUNDREDS

OF RULES PER TRANSACTION

executed deterministically, with intraday rule updates applied transactionally and without downtime.

ACTIVE-ACTIVE

RESILIENCE

via XDCR across multiple clusters, keeping fraud prevention operational through regional failures.

> "The decisioning layer narrowed the gap from the point of data ingestion to the point of decision-making, from minutes, or even hours, to milliseconds."

Conclusion: A Decisioning layer built for risk and trust

Fraud prevention is a test of architecture, not detection. Systems that score events asynchronously and reconcile outcomes later can identify fraud, but they cannot prevent it within the authorization window. The gap between knowing and deciding is where losses occur and trust erodes.

By embedding Volt Active Data as the deterministic decisioning layer, the provider gave its banking customer the ability to make authoritative decisions at the moment of execution: correct, consistent, and recorded as authoritative truth, even as volume surges toward a million transactions per second. Speed, ACID guarantees, programmability, and Active-Active resilience hold together as hard constraints, exactly where the cost of a late or inconsistent decision is too high to accept.

Fraud prevention is no longer a function of after-the-fact reconciliation.
It is a function of architecture.

Ready to make fraud prevention a system property?

Volt Active Data is purpose-built for the latency, throughput, and determinism demands of real-time risk and fraud decisioning. Whether you are protecting payments, enforcing limits, or stopping fraud before it completes, Volt gives you the decisioning layer to be correct: provably, at scale, in real time..

Talk to a Volt Solutions Architect today.

Visit www.voltactivedata.com, email info@voltactivedata.com or contact your Volt account team to arrange a briefing and proof-of-concept scoping session.