VOLT ACTIVE DATA

INDUSTRY BRIEF

Real-Time Rails: FedNow, SEPA Instant, RTP, Pix, UPI

Volt Active Data | BFSI & FinTech Solution Brief

Executive Summary

This brief examines the data infrastructure requirements created by real-time payment rails, including FedNow, SEPA Instant, RTP, Pix, and UPI. Today, FedNow is growing, SEPA instant adoption is accelerating under regulatory pressure across the eurozone, and UPI in India and Pix in Brazil have already exceeded traditional card networks in volume.

Who should read this:

Payments engineers, solutions architects, and technology leaders at banks, credit unions, and FinTechs building or modernizing instant payment infrastructure. Compliance and risk technology leaders responsible for AML, sanctions, and APP-fraud control frameworks will find the regulatory compliance sections directly relevant.

Related topics addressed in context:

Real-time fraud detection, AML rule execution, sanctions screening, and pre-funded liquidity management are discussed as integral parts of the rail processing workflow. Dedicated briefs on fraud decisioning architecture and payment gateway infrastructure are referenced at the end of this document.

The Infrastructure Requirement that Instant Rails Expose

Real-time payment rails process irrevocable instant credit transfers, typically with a 10-second end-to-end SLA and a bank-side processing budget that is often 2 to 3 seconds. Within that window, the bank's own systems need to complete a balance check, limit enforcement, fraud scoring, AML screening, sanctions and PEP lookup, ledger update, and customer notification. Once funds leave, they cannot be recalled.

UPI in India and Pix in Brazil already exceed traditional card networks in volume. FedNow is growing. SEPA Instant adoption is accelerating under regulatory pressure across the eurozone. The institutions on these rails made architectural decisions about their data plane years before the volume arrived, and many of those decisions are now under stress.

The irrevocability that makes instant rails attractive to consumers makes them attractive to fraudsters for the same reason. APP (authorized push payment) scam fraud is rising on every major rail. UK PSR now mandates sender-bank reimbursement for APP fraud, and analogous regulation is emerging elsewhere. Sanctions screening requirements have moved from post-event batch to sub-500ms per-transaction, which legacy AML platforms cannot meet without fundamental changes to how they are architected.

Where First-generation Rail Infrastructure Breaks Down

Most banks built their initial real-time rail capability on adapted batch infrastructure. That trade-off was acceptable at low volumes, but is not acceptable at the volumes these rails carry today, and certainly not at the volumes they are expected to handle in the near future.

The core architectural problem is that fraud, AML, balance, and limit checks are typically handled by separate services. On a batched or asynchronous flow, chaining those services together adds latency that gets absorbed into a multi-second processing window. On an instant rail with a 2-second bank-side budget, the same chain of network hops can exhaust the entire window before the decision is made. The response to this, in most first-generation implementations, is to reduce the scope of checks performed in real time rather than to fix the underlying architecture. That trade-off was acceptable at low volumes and is not acceptable at the volumes these rails now carry.

The solution most teams reach for next is caching. Stale cached state is not a minor accuracy issue in this context. A sanctions hit missed because the screening table has not been updated in 90 seconds, or an account balance read from a cache that does not reflect a concurrent transaction, produces outcomes that are regulatory failures or paid-out fraud events. The regulatory consequences of a missed sanctions hit now include fines measured in hundreds of millions of dollars in multiple jurisdictions.

Liquidity management adds another failure mode. Banks operating on pre-funded settlement models need to track their settlement position continuously. Institutions that track this in batch or near-real-time have experienced intraday outages when a position ran short in the gap between reads. The consequences range from reputationally damaging to materially disruptive depending on the volume of transactions affected.

ISO 20022 messages carry an order of magnitude more data than ISO 8583. Storage and indexing economics are not academic; they affect the cost of real-time AML lookups, the latency of sanctions screening, and the feasibility of in-database rule execution at the required throughput.

How Volt Addresses It

Volt runs fraud scoring, AML rule execution, sanctions lookup, balance check, and limit enforcement inside a single stored procedure. The entire set of checks completes in one database round trip rather than four or five network hops. At the throughputs real-time rails require, that architecture consistently delivers bank-side processing in under 10ms, leaving headroom within any rail's SLA for the wire calls to the rail operator and back.

The stored procedure model means that AML and sanctions screening tables live inside the same system as the transaction state and account balance. Updates to sanctions lists or fraud rules apply atomically. A rule that is deployed under live load either takes effect fully or does not take effect at all. There is no partial-deploy window in which some transactions are screened against old criteria and others against new.

Pre-funded settlement position tracking updates in the same transaction as the payment authorization. The institution's real-time view of its settlement position reflects every transaction that has been processed, including transactions in flight. Intraday liquidity shortfalls become visible before they cause an outage rather than after.

For resilience, Volt's active-active multi-region architecture allows every deployed region to accept writes, maintain authoritative state, and process transactions concurrently. When a region fails, traffic shifts automatically to healthy regions without manual promotion or reconfiguration. Regulators increasingly require sub-15-minute RTO for real-time rail operators; the active-everywhere architecture meets that requirement without a standby infrastructure that does nothing most of the time.

The ACID (Atomicity, Consistency, Isolation, and Durability) consistency guarantee carries a specific regulatory value that is worth stating plainly. Under UK PSR APP-fraud reimbursement rules and emerging equivalents elsewhere, institutions will need to demonstrate that their screening and fraud controls were operating correctly at the moment a specific transaction was processed. A system with eventual consistency cannot produce that demonstration reliably. Volt records each decision atomically, against a known and auditable state, at the point of transaction.

That record satisfies the evidentiary requirement in a way that post-hoc reconstruction from logs cannot.

Data Architecture

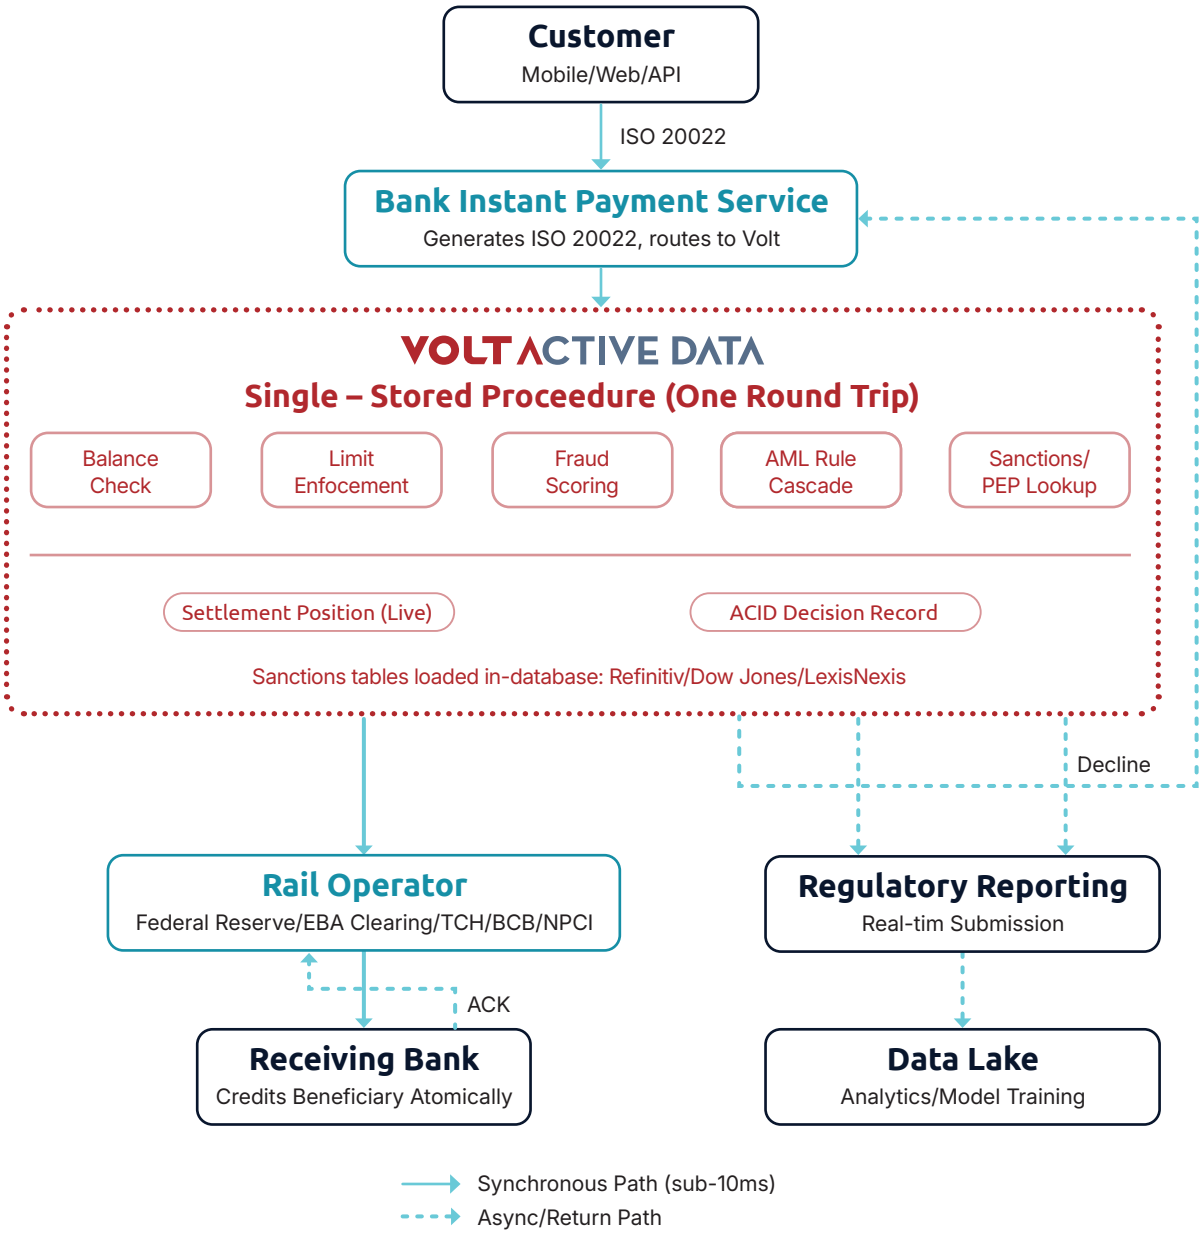

The bank's instant payment service receives a customer-initiated payment and generates an ISO 20022 message. The Volt stored procedure runs the full check set (balance, limits, fraud, AML, sanctions) in a single transaction and returns an approve or decline. On approval, the customer ledger is debited within the same transaction. The ISO 20022 message goes to the rail operator (Federal Reserve for FedNow, EBA Clearing for SEPA Instant, TCH for RTP, BCB for Pix, NPCI for UPI). Settlement position updates continuously in Volt. Regulatory reporting and data lake export run asynchronously from the Kafka stream, leaving the transaction path unaffected by downstream processing.

Sanctions and PEP screening tables from Refinitiv World-Check, Dow Jones, or LexisNexis load into Volt for in-database execution. Vendor APIs serve as the backup path for hits that require a live refresh.

Outcomes

P99 bank-side latency under 1 second, well inside the 10-second rail SLA. At a Tier-1 bank, a comparable consolidation of reporting infrastructure moved P99 from 150ms to under 12ms, demonstrating what single-system execution delivers relative to a multi-hop architecture.

50% or greater reduction in fraud losses on real-time rails versus first-generation implementations. The improvement comes from richer real-time state available at the point of the fraud decision and from the ability to deploy new rule logic within hours of a new attack pattern being identified.

30-50% reduction in pre-funding costs from continuous real-time settlement position tracking. Eliminating intraday blind spots removes the buffer institutions carry to protect against running short unexpectedly.

60-80% reduction in time to launch a new rail or rail feature. Hot-swappable rule logic means that adding a new rail connector, adjusting screening parameters, or responding to a regulatory change does not require a deployment window or an engineering sprint.

Business Value

The business case for modernizing real-time rail infrastructure is driven by four converging pressures, each with a direct financial consequence.

Regulatory compliance cost and exposure.

APP-fraud reimbursement mandates in the UK, and emerging equivalents across the EU and US, have created a new category of financial liability that did not exist three years ago. Under UK PSR rules effective from October 2024, sending banks are liable for reimbursing APP fraud losses up to £85,000 per claim, with liability shared with receiving banks. Institutions without the capability to demonstrate that their fraud controls were operating correctly at the moment of a specific transaction face both reimbursement liability and regulatory sanction. The architectural investment required to meet that evidentiary standard is the same investment that reduces the fraud losses themselves.

Revenue at risk from rail adoption.

Instant payment rails are growing faster than the fraud and compliance infrastructure most institutions originally built to support them. Institutions that constrain real-time screening to stay within their processing budget, or that disable checks under peak load, are accepting a fraud and compliance exposure that scales with their rail volume. The cost of that exposure is not linear: as volume grows, so does the attractiveness of the institution as a target for coordinated attacks.

Liquidity efficiency.

Banks operating pre-funded settlement positions on multiple rails carry buffer capital to protect against intraday blind spots. Real-time settlement position tracking reduces the required buffer. For institutions with large rail volumes, the capital released from reduced buffer requirements can be material.

Engineering capacity.

The fragmented, multi-service architecture that most first-generation rail implementations use is expensive to operate and slow to change. Regulatory updates, new rail connector requirements, and fraud rule changes all compete for the same deployment pipeline. Consolidating onto a single execution path reduces the operational overhead enough that institutions typically see meaningful reductions in both time-to-change and incident frequency within the first year of operation.

Agentic AI

Real-time rails are the use case where the cost of an agent acting on stale state is highest and most immediate. An autonomous APP-fraud agent that decides to hold a transaction must read the live state of the receiving account, the sender's behavioral history, and the bank's current fraud-loss exposure within the rail's processing window. A read that is 5 seconds old can be the difference between catching a fraud event and irretrievably releasing the funds.

The regulatory dimension matters here beyond the operational one. As APP-fraud reimbursement mandates require institutions to demonstrate that their controls were functioning correctly at a specific moment, the audit trail of what an agent queried, what data Volt returned, and what decision was recorded becomes a compliance artifact. Volt captures that chain completely, including every MCP tool call the agent made and the state values returned, which flows into both the regulatory record and the model retraining pipeline.

Related Reading

Real-Time Fraud Prevention Covers the fraud decisioning architecture in detail, including vendor signal integration, ML model orchestration, rule deployment under live load, and the per-transaction economics of in-house versus vendor fraud tooling. Available from voltactivedata.com/bfsi.

Payment Gateway Infrastructure Covers acquirer routing, idempotency, settlement state management, and the consistency requirements at gateway scale. The fraud and compliance controls discussed in this brief operate within the gateway's processing budget and are addressed in that context. Available from voltactivedata.com/bfsi.

Talk To Us

We work with payments engineering teams on FedNow, SEPA Instant, RTP, Pix, and UPI implementations. Reference architectures are available for each rail, and we are happy to walk through where Volt fits relative to your current data plane and what the migration path looks like.

voltactivedata.com/company/contact