
Human-AI Collaborative Architecture for Enterprise Financial Platforms

Ravindra Rajasekhar Kavuru

Abstract: Co-branded credit card platforms combine high-volume consumer software with stringent financial regulation, creating architectural challenges that standard design approaches cannot adequately address. This paper presents a human-AI collaborative architecture built around five interlocking design commitments: an event-driven core that captures every state transition as an immutable, replayable domain event; regulation-aware caching that restricts sensitive data domains to narrow read surfaces; cryptographic boundaries with key isolation scoped to the service and regulatory domain; a Zero Trust posture that enforces continuous authentication on every inter-service request; and a tiered human-AI collaboration model that is policy-governed rather than autonomous. The central argument is that compliance is not an external control overlay but a first-class structural property of data models, service boundaries, and event schemas from the outset of design. The resulting platform demonstrates that regulatory requirements and platform innovation are structurally complementary when encoded from the beginning of the architecture.

Keywords: *Event-Driven Architecture, Human-AI Collaboration, Zero Trust Security, Compliance-Native Design, Co-Branded Credit Card Platforms*

1. Introduction

Co-branded credit card programs and their associated enterprise financial systems occupy a unique intersection between high-volume consumer-facing software and heavily regulated financial systems, making them both complex and mission-critical. Such systems handle the entire lifecycle of digital onboarding, data handoffs for underwriting decisions, rewards reconciliation, third-party integrations, and audit-ready reporting in a manner consistent with modern consumer expectations. This operational complexity is not purely functional but an architectural concern, one that generates independent data and governance requirements that must be solved in concert.

What distinguishes co-branded credit card platforms from general-purpose enterprise software is the structural prevalence of asynchronous, issuer-dependent decisions embedded at the center of their transaction flows. In a canonical case, a consumer applies for a credit card through a partner's digital channel. The issuer does not return a final decision but rather a pending response, acknowledging

receipt of the application while deferring a final underwriting determination. This response pattern forces downstream systems — partner rewards engines and customer-facing status APIs — to manage a prolonged period of uncertainty without violating data-minimization principles, exposing prohibited underwriting detail, or degrading the customer experience. Failure to do so results in various forms of operational debt: support queues for status inquiries, inconsistent partner integration reconciliations, and compliance artifacts scattered across ad hoc integration logs.

Timeout semantics also apply to pending states. Regulatory and operational policy often require that an application in a pending state transition to the appropriate terminal state if no final decision event is produced within a specified period. This requirement is difficult to satisfy in asynchronous distributed services where state machine modeling, durable event persistence, and replay of state transitions under supervisory review are not first-class platform concerns. Systems without these capabilities rely on brittle designs consisting of scheduled database sweeps, manual reconciliation processes, and fragile point-to-point integrations.

Independent Researcher, USA

A principled design philosophy encodes compliance as a structural property of the platform from the outset, rather than as an after-the-fact control overlay. Event-driven architecture enables precisely this: by making the pending receipt, the timeout evaluation, and the final terminal classification each an immutable, append-only domain event, the platform obtains a single source of truth that serves simultaneously as the integration backbone and the evidentiary record [1]. This yields replayable, auditable workflow semantics required in regulated environments, while the loosely coupled, event-driven nature enables horizontal scalability across large consumer populations.

Beyond the event-driven core, additional architecture must enable AI applications without imposing a governance overhead. This includes operational AI capabilities such as anomaly detection, event summarization, alert triage, and recommended actions. Unconstrained AI may fail to satisfy the explainability, model behavior, and regulatory accountability requirements of high-impact financial applications. AI should therefore be deployed as a bounded collaborator rather than an autonomous decision-maker, in line with established best practices in AI risk management and governance [2]. Recommendation outputs are logged with confidence metadata and policy-check results, human dispositions are captured, and no AI recommendation is translated into a controlled business action without explicit human authorization.

This paper describes a human-AI collaborative architecture for this operating environment that addresses platform requirements across five dimensions: asynchronous state management, regulation-aware data partitioning, cryptographic trust boundaries, Zero Trust access control, and multi-tier human oversight of AI recommendations. The resulting design treats compliance and innovation as mutually reinforcing forces when regulation is encoded as a first-class architectural citizen.

2. Event-Driven Architecture as the Compliance Foundation

2.1 Immutable Event Trails

The foundation of a compliant financial platform is an event-driven architecture. Each change of application state — application submission, issuer response, underwriting status update, rewards

balance adjustment, timeout event — is represented as an immutable domain event. These events, along with accompanying metadata, form an append-only event log that serves simultaneously as the integration backbone and evidentiary record, providing a single verifiable source of truth that cannot be modified, only appended to.

This design is particularly valuable in cases where issuer decisions are deferred, which is common in co-branded credit card flows where the issuer returns a pending acknowledgment rather than a definitive approval or decline. The system does not attempt to resolve this ambiguity through synchronous retry logic or polling chains. Instead, the pending acknowledgment is stored as a distinct domain event with a precise timestamp. To derive the final state, the platform executes an explicit policy-based state machine that measures elapsed time against a configurable timeout threshold. If no terminal decision event occurs within the allotted window, a timeout event is emitted and the application transitions automatically to the appropriate terminal state. This process is fully auditable and reproducible, and it does not depend on a live connection with the issuer.

This pattern addresses a common flaw in synchronous, tightly coupled integration designs, where timeout logic is typically embedded implicitly in request-retry counters, thread-level timeouts, or database-level expiry flags that are difficult to audit and become brittle under load. With the event-driven model, each state transition is a first-class, versioned artifact. Engineers and auditors alike can replay the event log in event order to reconstruct the complete history of any application instance and verify that every state transition conformed to the governing policy at the time [3]. This is not merely an operational convenience but a regulatory requirement in many financial systems, where supervisory examinations require evidence of consistent rule application and prohibit retrospective record modification — properties that an append-only event log satisfies structurally.

The evidential value of the event log is further strengthened through cryptographic hash chaining. Each record is hashed together with the hash of the preceding record, forming a mathematically verifiable chain of custody. Any attempt to backdate or modify a past record breaks the hash chain of all subsequent records, creating a detectable inconsistency that automated integrity

verification systems can identify. This property directly supports the auditability and non-repudiation requirements examined in security and compliance audits.

Beyond compliance, the immutability of the event log provides a strong basis for operational recovery. Because the current state of every entity — an application, a rewards balance, a partner

integration — is always derivable from the event history, systems can be rebuilt from scratch by replaying the log. This eliminates the class of data corruption failures that arise when mutable state stores become inconsistent under partial failures or concurrent writes [3].

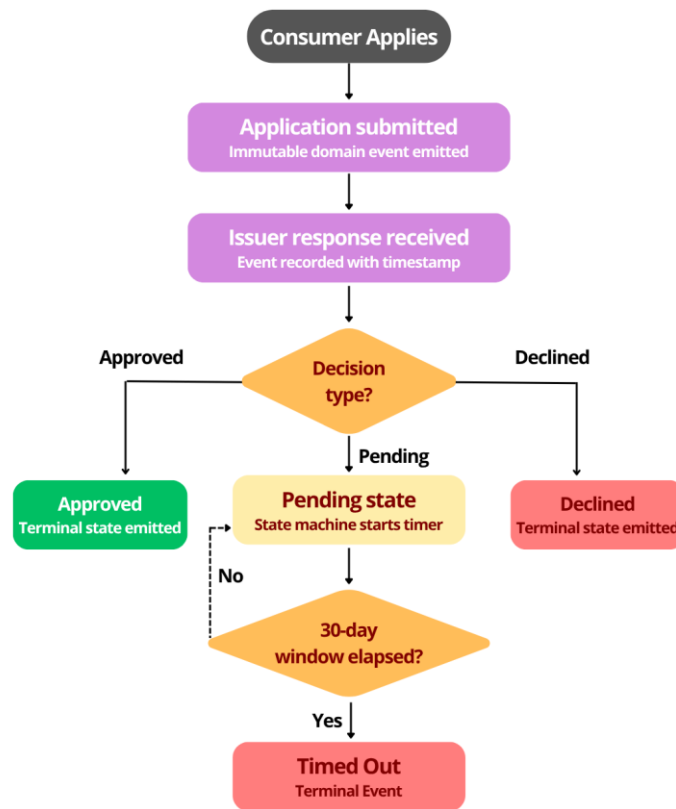


Figure 1: Asynchronous issuer decision and 30-day timeout state machine

2.2 Stateful Projections and Read Performance

While the append-only event log serves as the canonical data source, this structure is not suited to serving high-volume read requests. An application-status API that recalculates state by scanning an unbounded event log on every invocation would impose unacceptable latency in customer-facing workflows. The event-driven architecture addresses this through stateful projections: materialized read models that asynchronously consume the event stream and maintain a pre-computed representation of current entity state optimized for specific query patterns.

Projections are defined by subscribing to events of interest and applying deterministic projection logic to build a running state. An application-status projection subscribes to application-submitted,

issuer-pending, issuer-approved, and timed-out events, and projects the current state of each application indexed by customer and partner identifiers. This projection is exposed via a low-latency read API, providing the current state in a single indexed lookup. Because projections are derived views rather than authoritative state stores, they can be safely recomputed by replaying the event log, making them resilient to changes in query requirements.

A primary design tenant of projection design is the segregation of sensitive and non-sensitive data across projection boundaries. Non-sensitive projections — application status summaries, rewards eligibility flags, partner notification flags, and operational counts — can be aggressively cached at multiple levels in the read path. Because

these projections contain no personally identifiable information (PII), no payment account data, and no underwriting attributes, they present a small exposure surface even across cache boundaries. They can be deployed in distributed cache topologies close to API edge nodes to support the sub-200-millisecond response times required by consumer-facing flows under peak load [4].

Sensitive data domains — applicant PII, payment account data, and issuer underwriting attributes — are subject to distinct access and retention policies, stored within more secure service boundaries, never materialized into shared projection stores, never cached in distributed layers outside service boundaries, and never included in inter-service event payloads that cross trust boundaries. Access to sensitive projections is controlled through fine-grained policy and implemented through authenticated, authorized service-to-service calls that are logged for auditing. This transforms data minimization from a procedural policy into a structural property of the architecture: systems that do not expose sensitive data through non-sensitive read paths cannot do so inadvertently, even in the presence of application-layer defects [4].

The resulting regulation-aware caching model satisfies the dual requirements of co-branded credit card platforms: low-latency, high-throughput reads for frequently accessed operational data, and fine-grained isolation for data carrying the highest regulatory and privacy risk. This is not a tradeoff between performance and compliance but an architecture capable of satisfying both simultaneously by making data sensitivity a first-class parameter in the caching and projection design.

3. Cryptographic Boundaries and Zero Trust Posture

3.1 Cryptography as a Boundary-Setting Mechanism

Enterprise security models typically treat cryptography as a storage-layer control, protecting data at rest against physical or logical access to database and backup media. This framing is insufficient in a multiparty financial context, where data traverses service boundaries and integration adapters connecting to partner APIs and external card issuers—each with its own trust model, regulatory requirements, and information retention policies. In this architecture, cryptography functions as a boundary mechanism: a structural

control that defines the edges of a trust domain rather than simply encrypting content within one.

In practice, these requirements mean encryption in transit and at rest are non-negotiable defaults on every data path in the system. All inter-service communication, whether traversing external networks or an internal service mesh, uses current transport security protocols. There are no plaintext internal channels, no implicit-trust heuristics for traffic originating within private network segments, and no exceptions for high-throughput internal paths where overhead might be perceived as a performance cost. The performance argument against internal encryption has been rendered largely moot by hardware-accelerated cryptography and modern protocol optimization, while the risk of lateral movement following an initial compromise remains a well-documented attack surface [5].

Encryption at rest is applied with the additional architectural discipline of domain-based key isolation. Rather than encrypting data under a single organization-wide master key — which concentrates risk and complicates domain-specific key rotation — every logical service and regulatory boundary maintains its own encryption key. Payment account data is encrypted under keys scoped to payment security requirements, personal data under keys scoped to privacy law retention and deletion obligations, and issuer-specific fields under keys available only to the issuer integration service. A compromise of credentials in one service domain does not provide access to data in another — a property enforced architecturally rather than procedurally [5].

Audit log integrity is further protected through the cryptographic hash chaining described in Section 2.1. The resulting chain of custody is mathematically verifiable: any retroactive modification of a past log entry breaks the hash chain of all subsequent entries, creating a detectable inconsistency. In a regulatory audit where the integrity of archived logs must be demonstrably assured rather than merely asserted, hash-chained logs eliminate the category of dispute in which the integrity of the evidentiary record is challenged.

3.2 Zero Trust Architecture Across Multi-Party Boundaries

Zero Trust represents a fundamental departure from the perimeter-based security model that has dominated enterprise architecture for decades. In the perimeter model, network traffic originating

behind a defined boundary is implicitly trusted. This model has been repeatedly undermined by insider threats, supply chain attacks, credential theft, and the dissolution of the traditional network perimeter through cloud computing and partner business relationships. Zero Trust replaces implicit perimeter trust with a continuous, identity-based access model in which every resource request is authenticated and validated regardless of its network origin [6].

The architecture applies Zero Trust across three categories of service interaction, each with distinct trust properties. Internal platform services — application processors, rewards engines, and customer servicing APIs — are treated as mutually untrusted by default. All services authenticate using short-lived, cryptographically signed credentials issued by a centralized identity provider. Each access request is evaluated against a fine-grained authorization policy that specifies, per data type, which service identities may read or write which data and under which conditions. There are no broad internal trust grants, no service accounts with platform-wide data access, and no authorization decisions based solely on network location.

Partner API boundaries form the second trust domain, governed differently from internal services. These are the access points for external systems such as co-branded retail platforms, travel booking systems, and rewards redemption partners. Access is restricted strictly to the API surface each partner is contractually authorized to consume, enforced at the API gateway on ingress and egress. Authentication uses mutual transport-layer security and ephemeral token credentials scoped to a specific partner identity and immediately revoked upon expiration. Partner access is logged at the API gateway layer with sufficient fidelity to reconstruct the complete access history of any partner identity over any time range, supporting both contractual audit rights and regulatory third-party risk management expectations.

Issuer adapter boundaries carry the highest sensitivity in the architecture. Underwriting-relevant data exchanged with issuers is subject to the strictest data handling policies on the platform, accessible only to the issuer integration service, encrypted under issuer-domain keys as described above, and transported exclusively over mutually authenticated, encrypted channels with certificate pinning to prevent interception. Issuer adapter deployments are not network-accessible to other

platform services. The issuer integration boundary is not merely a logical construct defined by application-layer access control; it is a real security domain in which continuous authentication, fine-grained authorization, domain-isolated encryption, and per-boundary audit logging together form a Zero Trust posture enforced throughout the entire multi-party system. This posture remains effective even if individual credentials, tokens, or service instances are individually compromised [6].

4. Human-AI Collaboration Model

4.1 The Case for Bounded AI

AI in financial platforms delivers operational value when applied within a well-defined and circumscribed scope. Machine-learning classifiers can process event streams at a scale and velocity unachievable by human cognition; large language models can compress detailed incident histories into actionable summaries in seconds, and anomaly scoring can detect subtle patterns imperceptible to human reviewers. However, financial services is a high-consequence, heavily regulated operating environment in which any system output that affects a regulated business decision makes the institution, its partners, and its customers accountable for the outcome.

Uncontrolled AI introduces several interrelated risks. Explainability risk arises when credit decisions, timeout classifications, or partner notifications occur without traceable, interpretable reasoning. Supervisory and audit bodies require that decisions be explainable in terms of governing policy and underlying data. Model behavior risk arises from the statistical nature of machine learning inference: a model trained on historical event data will generalize from historical patterns—operational anomalies, seasonal distributional shifts, and issuer-specific effects—that should not be encoded as durable behavioral rules. Without model monitoring and governance in production, model drift can produce material operational failures [7]. Privacy risk arises when AI components are granted access to sensitive data domains—PII, underwriting attributes, and payment account data—beyond the scope strictly necessary for inference, unnecessarily expanding the data exposure surface and potentially violating data minimization principles.

The appropriate response is not to exclude AI from the platform but to deploy it as a bounded collaborator with explicit, documented policies and

assumptions, scoped to tasks within its demonstrated competence: aggregating event trails into incident narratives, scoring the severity of anomalies in alert streams, classifying clusters of support tickets, and presenting recommended next actions to human evaluators. The boundary between AI recommendation and human authorization is not a soft convention but an architecturally enforced policy checkpoint through which every AI output must pass before any controlled change to system state is made. This design is closely aligned with prevailing AI governance guidance, which holds that high-stakes AI applications require not only capable models but also constrained deployment contexts, an accountability framework, and continuous human oversight at critical decision points [7].

4.2 Tiered Decision Authority

The collaboration model operates across three tiers of decision authority, differentiated by the risk profile and reversibility of each action class. Rather than a binary distinction between automated and non-automated, the model allows AI to assist at every tier while calibrating the degree of human oversight to the consequence of the action.

Tier 1 — Low Risk. These are actions for which the model has high confidence, the governing policy is well-defined, and the action is fully reversible. Representative examples include attaching a model-generated runbook summary to an open incident record, automatically classifying a low-severity alert into a predefined category, and correlating an event stream anomaly with a known fault pattern. Actions at this tier cannot modify the authoritative business state, cannot produce customer- or partner-facing communications, and can be reversed immediately upon identification of a model error. The operational benefit at this tier is that AI can process high volumes of simple operational signals that would otherwise consume disproportionate engineering attention or go unaddressed.

Tier 2 — Medium Risk. These are actions where the AI-generated recommendation carries substantial operational impact but where execution entails nontrivial reversal costs. A representative example is the reconciliation of partner application status for a case that has exceeded its timeout threshold without receiving a terminal issuer decision. The AI module evaluates the event history to confirm that no terminal event was

received within the policy-defined window and that the case duration exceeds applicable thresholds, and outputs a structured, non-executed reconciliation recommendation with event citations and confidence metadata. This recommendation is presented to a human operations analyst, who reviews the full context and either approves or rejects the proposed state transition. The final disposition — including the AI recommendation, confidence score, policy-check results, and authorization timestamp — is captured as a complete, non-repudiable audit record [8].

Tier 3 — High Risk. At this tier, AI presents contextual information only and does not generate recommendations. Representative actions include policy-exception approvals, issuer integration rule changes, data retention schedule modifications, and cross-partner configuration changes. These changes have broad systemic effects, are expensive to reverse, and require contextual policy judgment and institutional accountability that no one can delegate. At this tier, AI functions as an information retrieval and summarization tool, surfacing prior exceptions, related configuration histories, and event histories of affected application populations to inform the human decision-maker without substituting for their judgment.

For each tier, AI outputs conform to a standard metadata structure capturing the model identifier and version, the features and event references used in generating the output, the confidence score, the results of automated policy checks performed prior to presenting the recommendation, the human disposition (accepted, modified, or rejected), and the authorized identity of the human decision-maker. Collectively, this audit trail supports posterior model measurement, supervisory review of human-AI decision patterns, and feedback data for model governance and improvement [8].

Decision Tier	Representative Action	Risk Profile	AI Role	Human Role	Audit Trail
Tier 1 – Low Risk	Alert classification; runbook attachment	Bounded; fully reversible	Autonomous execution	Post-hoc review if needed	Logged with model version and confidence
Tier 2 – Medium Risk	Timeout reconciliation; status confirmation	Non-trivial; partially reversible	Generates structured recommendation	Explicit authorization required	Logged with confidence, policy check, and disposition
Tier 3 – High Risk	Policy exceptions; issuer rule changes	Broad systemic; largely irreversible	Surfaces historical context only	Fully human-led decision	Logged with authorized identity and rationale

Table 1: AI Decision Tiers and Governance Characteristics

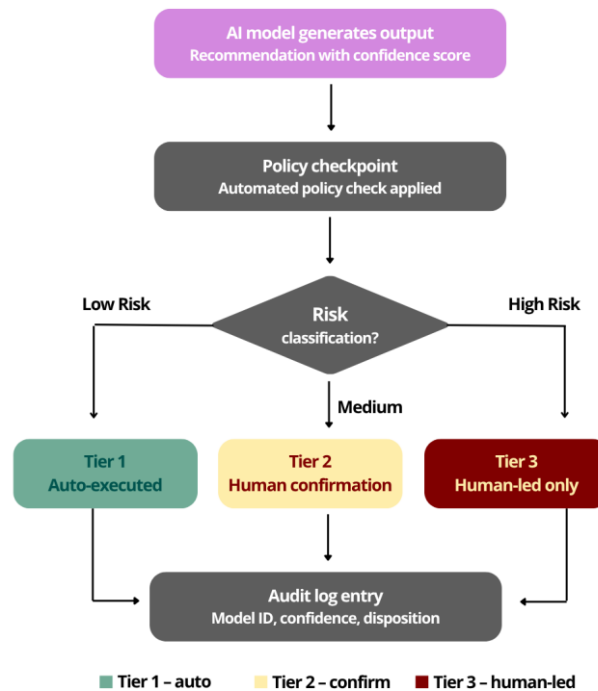


Figure 2: Tiered human-AI decision authority and audit trail flow

5. Compliance as an Architectural Property

5.1 From External Control Overlay to Structural Design Constraint

The most consequential architectural insight available to teams operating at the intersection of financial services and consumer technology is deceptively simple: regulatory compliance is least costly and most durable when it is native to the platform from the outset. Conventional enterprise compliance practice treats regulation as a follow-on concern—platforms are designed and deployed to satisfy functional objectives, then overlaid with access controls, audit logging, data masking, and

controlled environment configurations in subsequent iterations. This approach is intrinsically fragile: controls placed in an external overlay are, by design, not tightly coupled to the system behavior they govern, producing persistent semantic gaps between the control framework's model of the system and the system's actual runtime behavior.

An alternative is to treat compliance as a first-class architectural property. When compliance constraints are encoded in the initial design of data models, service boundaries, event schemas, caching topologies, and deployment configurations,

they become intrinsic behavioral properties of the system rather than procedural side effects of its operation. A system that cannot write PII into a shared projection cache — because that layer was architecturally scoped to exclude sensitive data domains from the outset — does not rely on application-layer logic, operational discipline, or developer attention to maintain that property. That boundary does not erode through configuration drift or the accumulation of technical debt [9].

The principle of least privilege provides a concrete illustration. Payment account data—credit card numbers, bank account identifiers, and authorization codes—is not protected solely by access control lists on a shared database. The rewards service has no API interface to the payment account service, no credentials to authenticate to it, and the payment account service's response schemas do not include payment account data in any form. A developer working in the rewards service cannot inadvertently read payment account data because the architectural design affords no interface through which to do so. The compliance control is an expression of a security engineering principle: security properties are most reliably achieved when security is a property of the design rather than a control applied to an insecure-by-default system [9].

5.2 Regulatory Frameworks as Architectural Dimensions

Each major regulatory framework applicable to a co-branded credit card platform contributes a first-class design dimension when its implications are treated as architectural inputs rather than compliance checkboxes.

Data privacy law, governing the collection, storage, and deletion of personal data, shapes the data lifecycle architecture of the platform in ways with direct technical implications. Data minimization is expressed architecturally through event schema design that excludes unnecessary personal attributes from integration events, projection design that derives aggregated status indicators rather than propagating raw personal data across service boundaries, and retention automation that treats data deletion as a durable platform behavior rather than an ad hoc operational task. Encoding data minimization as a schema constraint and service boundary property reduces the surface area of personal data to be secured and managed, producing a security benefit in addition to satisfying the privacy compliance requirement [10].

Security and availability control frameworks influence how evidence is created, stored, and presented to auditors. Rather than periodic, manual evidence collection — which consumes meaningful effort and is subject to gaps — compliance as an architectural property implies continuous, automated evidence generation. The immutable event log described in Section 2 produces a complete, timestamped record of every state transition managed by the platform. Access logs, policy-policy-verification results, human authorizations, and AI recommendation audit trails are all produced by the operational infrastructure of the platform without requiring separate evidence collection procedures. Evidence packages for an examiner are produced by querying structured log stores rather than manually reconstructing records from scattered logs or spreadsheets.

Operational resilience guidance for financial technology platforms requires that third-party and infrastructure dependencies be explicitly modeled, tested, and managed. Resilience is an architectural property in this design: risks arising from external dependencies are addressed through explicit integration contract design rather than implicit coupling. Issuer adapters, partner API integrations, and external data feeds are modeled as bounded integration domains with explicit failure semantics — circuit breakers to isolate downstream failures, event-driven decoupling to allow the platform to continue serving internal workflows when external integrations are unavailable, and defined degraded-state behaviors to sustain core customer-facing functions under partial outage conditions.

AI governance obligations — including model deployment monitoring and human oversight mechanisms — are similarly treated as architectural properties. Model boundaries are governed through the same policy checkpoint tooling that handles other controlled actions, avoiding the need for a parallel governance architecture maintained exclusively for AI components. Confidence thresholds, policy checks, and human authorization requirements are encoded as properties of the AI recommendation pipeline, ensuring continuous compliance with governance requirements rather than periodic model risk reviews that may not reflect actual production model behavior [10].

The result of treating compliance requirements as architectural inputs is a platform that is not only compliant but measurably more reliable, more

auditable, and more extensible than one for which compliance was an afterthought. Encoding compliance as a structural property eliminates entire classes of operational risk that procedural

compliance cannot address, producing the predictable, well-bounded system behavior that is the foundation of confident innovation and sustainable scale.

Architectural Mechanism	Corresponding Regulatory Framework	Structural Role in Platform
Immutable Event Trails	Audit and control frameworks; operational resilience guidance	Single source of truth; replayable evidentiary record for supervisory review
Regulation-Aware Caching	Payment security requirements; data privacy law	Separates sensitive and non-sensitive data across projection boundaries
Cryptographic Boundaries	Payment security requirements; data privacy law	Enforces trust domain edges; isolates keys per service and regulatory scope
Zero Trust Posture	Network security standards; third-party risk guidance	Continuous authentication and authorization across all inter-service interactions
Compliance-Native Design	All applicable regulatory frameworks	Encodes regulatory constraints as structural platform properties from outset

Table 2: Architectural Mechanisms and Their Regulatory Alignment

6. Conclusion

Co-branded credit card platforms represent some of the most architecturally demanding environments in enterprise software, combining high consumer expectations with opaque, asynchronous issuer workflows, multi-party integration requirements, and intensive regulatory scrutiny across payment security, data privacy, operational resilience, and AI governance domains. This paper has described an architecture that addresses these constraints through five interlocking design commitments: an event-driven core in which every state transition is immutable, replayable, and auditable; regulation-aware caching that enables low-latency reads while restricting sensitive data domains to narrow, controlled surfaces; strong cryptographic boundaries that isolate encryption keys by service and regulatory domain; a Zero Trust access posture that enforces continuous authentication and authorization across every inter-service interaction, regardless of network boundary; and a tiered human-AI collaboration model that deploys AI as a bounded, policy-governed collaborator rather than an autonomous decision-maker in regulated workflows.

The central finding is that compliance and innovation are not opposing forces. When regulatory requirements are encoded as first-class structural properties of the technology—rather than as procedural controls layered on afterward—they become enabling constraints that make the platform simultaneously more reliable, more auditable, and more extensible. An AI bounded by well-specified

policy guardrails, validated confidence scoring, auditable recommendation trails, and mandatory human authorization at consequential decision points is a genuine operational force multiplier, reducing engineering toil, accelerating incident triage, and improving explainability under supervisory scrutiny.

Several directions remain open for future investigation. First, the tiered human-AI collaboration model presented here relies on static confidence thresholds and policy-check rules; adaptive governance mechanisms that update tier boundaries in response to measured model drift or operational feedback represent a promising extension. Second, the architecture has been described in the context of co-branded credit card programs; applying the same compliance-native design principles to adjacent regulated domains—open banking, digital lending origination, and buy-now-pay-later platforms—would test the generalizability of the approach and surface domain-specific refinements. Third, while cryptographic hash chaining provides tamper evidence for the event log, integration with distributed ledger technologies or verifiable credential frameworks could further strengthen the non-repudiation properties of the evidentiary record in multi-party dispute resolution scenarios. Fourth, the human oversight model raises open questions about optimal interface design for AI-assisted decision support at Tier 2: how recommendation presentation format, confidence visualization, and audit trail transparency affect human decision

quality and throughput are empirical questions warranting controlled study.

For organizations operating at the intersection of consumer commerce and regulated finance, this architecture offers a principled, evidence-based, and practically deployable path toward platforms that are simultaneously scalable, secure, and trustworthy.

References

- [1] Michiel Overeem, et al., "An Empirical Characterization of Event Sourced Systems and Their Schema Changes: Lessons from Industry," *Journal of Systems and Software*, vol. 178, Aug. 2021, Art. no. 110970, doi: 10.1016/j.jss.2021.110970. [Online]. Available: <https://doi.org/10.1016/j.jss.2021.110970>
- [2] National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100-1, Jan. 2023. [Online]. Available: <https://doi.org/10.6028/NIST.AI.100-1>
- [3] Martin Fowler, "Event Sourcing," in martinfowler.com, Dec. 2005. [Online]. Available: <https://martinfowler.com/eaDev/EventSourcing.html>
- [4] Scott Rose, et al., "Zero Trust Architecture," NIST Special Publication 800-207, National Institute of Standards and Technology, Aug. 2020, doi: 10.6028/NIST.SP.800-207. [Online]. Available: <https://doi.org/10.6028/NIST.SP.800-207>
- [5] National Institute of Standards and Technology, "Guide to General Server Security," NIST Special Publication 800-123, Jul. 2008, doi: 10.6028/NIST.SP.800-123. [Online]. Available: <https://doi.org/10.6028/NIST.SP.800-123>
- [6] John Kindervag, "No More Chewy Centers: Introducing the Zero Trust Model of Information Security," Forrester Research, Cambridge, MA, USA, Tech. Rep., 2010. [Online]. Available: <https://media.paloaltonetworks.com/documents/Forrester-No-More-Chewy-Centers.pdf>
- [7] Ben Shneiderman, "Human-Centered AI: Reliable, Safe and Trustworthy," *International Journal of Human-Computer Interaction*, vol. 36, no. 6, pp. 495–504, 2020, doi: 10.1080/10447318.2020.1741118. [Online]. Available: <https://arxiv.org/pdf/2002.04087>
- [8] Percy Liang et al., "Holistic Evaluation of Language Models," *Annals of the New York Academy of Sciences*, vol. 1525, no. 1, pp. 140–146, Jul. 2023, doi: 10.1111/nyas.15007. [Online]. Available: <https://arxiv.org/pdf/2211.09110>
- [9] Ross Anderson, "Security Engineering: A Guide to Building Dependable Distributed Systems, 3rd ed." Indianapolis, IN, USA: Wiley, 2020, ch. 1, pp. 1–35. [Online]. Available: <https://www.cl.cam.ac.uk/archive/rja14/Papers/SEv3.pdf>
- [10] European Parliament and Council of the European Union, "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data (General Data Protection Regulation)," *Official Journal of the European Union*, L 119, pp. 1–88, May 2016. [Online]. Available: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679>