Journal of Contemporary Science and Technology Management

Volume 02, Issue 02, December 2026

Digital Repair Twins for Self-Healing Contactless Payment Gateways Using Artificial Intelligence and Machine Learning

Vipin Malik
Department of Computor Science, CCS University Meerut, UP India

Abstract

Contactless payment gateways — the software and infrastructure layers that process NFC tap-to-pay, QR-based, and wearable-device transactions — must sustain sub-second response times at very high transaction volumes while remaining resilient to hardware faults, network jitter, certificate expiry, and fraud pattern shifts. Failures in this layer are especially costly because contactless transactions rely on an implicit assumption of instant approval; even brief degradation is immediately visible to consumers at the point of sale. This paper proposes a Digital Repair Twin (DRT) framework tailored specifically to contactless payment gateways, combining a continuously updated virtual model of gateway topology and state with artificial intelligence (AI) and machine learning (ML) models for real-time anomaly detection, root-cause diagnosis, and autonomous or semi-autonomous self-healing. We detail the distinctive technical characteristics of contactless transaction flows — including terminal-to-gateway handshakes, tokenization, EMV contactless kernels, and issuer authorization round trips — and show how these characteristics shape the design of a repair twin's data pipeline and model set. A reference architecture is presented alongside a simulated operational scenario in which the DRT detects a terminal-firmware-induced spike in failed NFC handshakes, isolates the affected terminal fleet, and triggers an automated fallback to chip-and-PIN processing while alerting field engineering teams. We further examine security, latency, and regulatory constraints specific to contactless payments, including EMVCo certification and PCI requirements, and we discuss evaluation metrics such as mean time to detect (MTTD), mean time to repair (MTTR), and false-positive remediation rate. The paper concludes with a discussion of open challenges, including edge-device heterogeneity, offline transaction handling, and the limits of automated repair in safety-critical financial infrastructure.

Keywords: Contactless payments; digital repair twin; self-healing systems; machine learning; NFC; EMV; payment gateway; anomaly detection; fraud detection; fintech infrastructure.

1. Introduction

Contactless payments — tap-to-pay cards, NFC-enabled smartphones, and wearable devices — have become the dominant point-of-sale interaction pattern in many retail markets, prized for their speed and convenience. This convenience, however, is underpinned by a payment gateway stack that must complete terminal authentication, tokenization, risk scoring, and issuer authorization within a strict latency budget, typically under one second end to end. Because the entire user experience is built around near-instant approval, even transient

degradation in gateway performance is immediately felt by merchants and consumers, unlike card-present chip transactions where a slightly longer wait is tolerated.

Operating contactless gateways at scale therefore demands infrastructure that not only monitors itself but actively repairs itself before human operators can intervene. Traditional monitoring dashboards and static alerting thresholds are poorly suited to the fast-moving, high-cardinality failure modes typical of contactless processing, which can originate from terminal firmware bugs, NFC antenna interference, tokenization service latency, or sudden shifts in fraud-model behavior triggered by a new device type entering the field. This paper extends the general Digital Repair Twin (DRT) concept — a live, diagnostic, and action-capable virtual model of a system — into a domain-specific framework for contactless payment gateways.

The remainder of this paper proceeds as follows. Section 2 reviews the technical architecture of contactless payment processing and prior work on digital twins and AIOps in financial infrastructure. Section 3 defines the contactless-specific DRT and its required data and model components. Section 4 presents a reference architecture. Section 5 details the AI/ML techniques used for detection, diagnosis, and repair. Section 6 walks through a simulated operational scenario. Section 7 discusses security, compliance, and evaluation considerations, and Section 8 concludes with open research directions.

2. Background

2.1 Anatomy of a Contactless Transaction

A typical contactless transaction begins with a terminal-to-card or terminal-to-device NFC handshake governed by an EMV contactless kernel, which negotiates supported application identifiers and exchanges cryptographic data. The terminal forwards an authorization request to the acquiring gateway, which performs tokenization or de-tokenization, applies risk and fraud scoring, and routes the request to the appropriate card network and issuing bank. The issuer returns an approval or decline, which propagates back through the gateway to the terminal, typically within a few hundred milliseconds. Each of these hops is a potential point of failure, and because contactless flows are optimized for speed, there is little slack in the latency budget to absorb degradation gracefully.

2.2 Digital Twins and AIOps in Payments

Digital twin techniques, originally developed for physical asset monitoring, have been adapted to model software and infrastructure state in domains ranging from cloud operations

to telecommunications. In parallel, AIOps platforms apply machine learning to operational telemetry to reduce alert fatigue and accelerate incident response. Prior work applying these ideas to payments has generally focused on account-to-account or card-present transaction pipelines; the specific latency, protocol, and hardware constraints of contactless processing — including terminal fleets, NFC-specific error codes, and EMV kernel negotiation failures — have received comparatively little attention, motivating the domain-specific framework proposed here.

2.3 Why Contactless Gateways Need a Specialized Repair Twin

Three characteristics distinguish contactless gateways from generic payment infrastructure. First, the transaction path includes physical, field-deployed hardware — terminals and readers — whose firmware and antenna condition directly affect success rates, introducing a hardware-telemetry dimension absent from purely server-side systems. Second, the latency budget is extremely tight, leaving little room for retries or fallback logic that would be acceptable elsewhere. Third, contactless fraud patterns evolve quickly as new device types, wallets, and wearables enter the ecosystem, requiring models that adapt without long retraining cycles. A DRT for this domain must therefore ingest terminal-level telemetry alongside server-side metrics and must favor low-latency, incremental model updates over infrequent batch retraining.

3. Defining the Contactless Payment Gateway DRT

We define a Contactless Payment Gateway Digital Repair Twin as a continuously updated model of the terminal fleet, gateway services, tokenization infrastructure, and issuer connections involved in contactless transaction processing, coupled with AI/ML components that detect anomalies in handshake success rate, authorization latency, and decline patterns; localize the anomaly to a specific terminal cohort, gateway service, or network path; and execute or recommend a bounded remediation action such as terminal firmware rollback, traffic rerouting, or temporary fallback to an alternate transaction mode. Unlike a general-purpose payment DRT, this model must represent the physical terminal layer explicitly, since a meaningful share of contactless failures originate below the software stack, in antenna hardware, firmware versions, or terminal configuration drift.

4. Reference Architecture

The proposed architecture consists of five cooperating layers. The telemetry layer aggregates terminal logs, NFC handshake outcomes, gateway service metrics, tokenization latency, and issuer response codes through a combination of terminal-embedded agents and server-side

stream processors. The twin-state layer maintains a live graph representation linking terminal cohorts (grouped by model, firmware version, and location) to the gateway services and issuer connections they depend upon, updated continuously as new telemetry arrives. The modeling layer hosts anomaly detection models operating at both the terminal-cohort and service level, together with root-cause localization models that traverse the twin-state graph. The policy layer evaluates candidate remediations against a pre-approved action catalog — for example, firmware rollback, terminal cohort quarantine, traffic rerouting, or fallback-mode activation — filtered by risk thresholds and regulatory constraints. The action layer executes approved remediations through terminal management systems and gateway configuration APIs, logging every decision for audit and feeding outcomes back into the telemetry layer to support continuous model refinement.

A key architectural decision is the separation of terminal-cohort-level anomaly detection from transaction-level fraud scoring. The former operates on aggregated success-rate and latency statistics per cohort and is optimized for fast detection of hardware- or firmware-driven degradation; the latter operates on individual transaction features and is optimized for fraud risk. Keeping these concerns separate, while allowing the root-cause layer to correlate across both, avoids conflating operational failures with fraud events, which require very different remediation paths.

5. AI and ML Techniques

5.1 Terminal-Level Anomaly Detection

Because terminal telemetry is inherently grouped by cohort (device model, firmware version, geographic region), anomaly detection is applied at the cohort level using streaming statistical methods and lightweight sequence models to flag deviations in NFC handshake success rate or average tap-to-approval time. Cohort-level modeling allows the system to distinguish an isolated terminal fault from a systemic firmware issue affecting an entire device generation, which requires a markedly different remediation.

5.2 Root-Cause Localization Across Hardware and Software Layers

Root-cause localization in this domain must span both hardware and software layers. Graph-based propagation techniques applied to the twin-state graph allow the system to test whether an anomaly is better explained by a terminal-cohort factor (for example, a specific firmware version), a gateway-service factor (for example, a tokenization service instance), or a network-path factor (for example, a specific issuer connection), by comparing anomaly concentration across each candidate grouping. This cross-layer reasoning is essential, since a

naive service-level anomaly detector would misattribute a firmware-driven handshake failure to the gateway service simply because that is where the resulting errors surface.

5.3 Adaptive Fraud and Risk Scoring

Contactless fraud-scoring models must adapt quickly as new device types and wallet providers enter the market, since a purely static model risks either blocking legitimate new device cohorts or failing to catch fraud patterns specific to a new channel. Online learning and incremental model updates, combined with human-in-the-loop review of newly emerging device cohorts, help balance adaptability against the risk of destabilizing a production fraud model.

5.4 Repair Action Selection

Given a localized root cause, the DRT selects a remediation from a bounded, pre-approved action catalog rather than an open-ended action space, reflecting the safety-critical nature of payment infrastructure. Candidate actions for contactless-specific faults include quarantining a terminal cohort pending a firmware update, temporarily disabling a specific EMV kernel version known to be problematic, rerouting traffic away from a degraded issuer connection, or activating a fallback transaction mode such as chip-and-PIN or magnetic-stripe processing for affected terminals. Each action is associated with an estimated blast radius and reversibility score, which the policy layer uses to decide whether the action can be executed automatically or must be routed to a human operator for confirmation.

6. Simulated Operational Scenario

To illustrate the framework, consider a simulated scenario in which a retail chain deploys a firmware update to a specific terminal model across several thousand locations. Within minutes, the telemetry layer observes a rise in NFC handshake failures concentrated among terminals running the new firmware version, while terminals on the previous firmware version remain unaffected. The cohort-level anomaly detector flags the handshake failure rate for the new-firmware cohort as a statistically significant deviation from baseline. The root-cause model, traversing the twin-state graph, confirms that the anomaly is concentrated by firmware version rather than by gateway service, geographic region, or issuer connection, ruling out a server-side cause.

The policy layer consults the action catalog and identifies a pre-approved, low-risk remediation: activating an automatic fallback to chip-and-PIN processing for terminals running the affected firmware version, while flagging the firmware version itself for rollback

by field engineering. Because this action is reversible, scoped narrowly to the affected cohort, and does not alter settled transaction records, it is executed automatically by the action layer. Consumers at affected terminals experience a brief prompt to insert their card rather than tap, and authorization success rates for the affected cohort recover immediately. A complete audit record — anomaly signal, root-cause confidence, policy rule invoked, and action taken — is logged for later review, and the incident is escalated to field engineering for firmware remediation, illustrating how the DRT reduces both mean time to detect and mean time to repair relative to a manual, ticket-driven response process.

7. Security, Compliance, and Evaluation Considerations

7.1 EMVCo and PCI Requirements

Any automated action affecting contactless transaction processing must remain compliant with EMVCo contactless specifications and PCI DSS requirements governing cardholder data handling. In practice, this means that fallback and rerouting actions must be limited to pre-certified transaction modes and configurations, and that telemetry used to train DRT models must exclude or tokenize sensitive cardholder data at the point of collection rather than relying on downstream anonymization.

7.2 Latency Budget for Detection and Repair

Given the sub-second latency budget of contactless transactions, the DRT's detection and repair loop must itself operate on a timescale of seconds to minutes for cohort-level issues, since individual transaction-level intervention is not feasible within the authorization window. This distinguishes the appropriate scope of automated repair in this domain: the DRT protects the health of the terminal and gateway population over short time horizons rather than intervening within any single transaction.

7.3 Evaluation Metrics

We propose evaluating contactless payment DRTs along three primary metrics: mean time to detect (MTTD), the interval between the onset of a real degradation and its flagging by the anomaly detection layer; mean time to repair (MTTR), the interval between detection and successful remediation; and false-positive remediation rate, the proportion of automated actions taken in response to anomalies that, on later review, did not correspond to a genuine fault. This last metric is particularly important in a safety-critical domain, since an unnecessary fallback action, while reversible, still degrades the customer experience the system is meant to protect.

8. Conclusion and Future Work

This paper has presented a Digital Repair Twin framework specifically tailored to contactless payment gateways, addressing the domain's distinctive combination of field-deployed hardware, extremely tight latency budgets, and rapidly evolving fraud patterns. By extending the twin-state representation to include terminal cohorts alongside gateway services, and by pairing cohort-level anomaly detection with cross-layer root-cause localization and a bounded, auditable action catalog, the proposed framework offers a path toward contactless infrastructure that detects and repairs many classes of degradation before they are visible to consumers. Future work should address edge-device heterogeneity across manufacturers and firmware ecosystems, robust handling of offline or intermittently connected terminals, and the development of shared, privacy-preserving benchmarks that would allow different institutions and vendors to compare DRT performance on a common basis. As contactless adoption continues to grow globally, the resilience of the underlying gateway infrastructure will increasingly depend on exactly this kind of AI- and ML-driven self-healing capability.

References

Manne, V. T. (2025, October). Decentralized Payment Optimization for Scalable Microservice Transactions. In 2025 IEEE International Conference on Blockchain and Distributed Systems Security (ICBDS) (pp. 1-6). IEEE.

MARASANI, Y. (2023). Machine Learning Models for Predicting Patient Treatment Switching Using Claims Data. Frontiers in Computer Science and Artificial Intelligence, 2(1), 59-66.

Barua, S. . (2023). Biochar-Based Treatment Technologies for PFAS Removal from Industrial Storamwater and Wastewater: Mechanisms, Field Applications, and Future Regulatory Implications. International Journal of Technology, Management and Humanities, 9(04), 257-275.

Venkata, S. B. (2026, January). Runbook Mesh: MCP-Orchestrated Terraform and Ansible Co-Execution on Azure. In 2026 Second International Conference on Intelligent Systems for Communication, IoT and Security (ICISCoIS) (pp. 1-7). IEEE.

Venkata, S. B. (2025, December). Predictive infrastructure orchestration in azure using terraform and dynatrace for medical systems. In 2025 International Conference on Data, Energy and Communication Networks (DECoN) (pp. 1-6). IEEE.

MARASANI, Y. (2024). Enterprise Readiness for Generative AI: The Critical Role of Data Engineering. Frontiers in Computer Science and Artificial Intelligence, 3(2), 59-71.

Manne, V. T. (2026). A Lightweight Application-Layer Defense Against Relay Attacks in Contactless Transactions.

Marasani, Y. (2025). Explainable AI Frameworks for Patient-Level Claims Data Analytics. J Artif Intell Mach Learn & Data Sci, 8(1), 3382-3390.

Venkata, S. B. (2026, March). HERA-QI: Vision Language Quality Inspection for Hearing Aid Hardware and Software. In 2026 9th International Conference on Intelligent Computing and Control Systems (ICICCS) (pp. 2000-2006). IEEE.

Barua, S. (2026). AI-Based Early Warning Systems for Industrial Stormwater Exceedances: A Data-Driven Approach to Regulatory Compliance and Environmental Protection. *Journal of Science Technology and Social Transformation*, *2*(02), 9-21.