Journal of Science, Technology and Social Transformation Volume 1, Issue 01, 2025

Data Engineering Architectures for AI-Based Digital Repair Twins in Financial Cybersecurity

Chum Derin*

Independent Researcher, Hong Kong SAR China

Abstract

Financial institutions face a growing volume of cyber threats — credential-stuffing attacks, account-takeover attempts, insider misuse, and infrastructure intrusions — that must be detected and contained within tight operational windows to prevent financial and reputational harm. This paper examines the data engineering architectures required to support AI-based Digital Repair Twins (DRTs): live, diagnostic, and action-capable virtual models that detect security anomalies, localize root causes, and initiate bounded remediation within financial infrastructure. We present a five-layer reference architecture spanning ingestion and streaming, feature and twin-state storage, AI/ML model serving, reasoning and policy, and action and remediation, and we describe the data pipeline design decisions specific to financial cybersecurity, including entity-graph modeling of accounts and sessions, lineage tracking for regulatory audit, and strict data governance for sensitive financial and personally identifiable information. A table-based comparison of pipeline layers, technologies, and data artifacts is provided, alongside a second table of evaluation metrics for operational assessment. Two figures illustrate the layered architecture and the closed-loop detection-to-repair workflow. A simulated case study involving a credential-stuffing attack against an online banking portal demonstrates how the architecture supports rapid detection, root-cause correlation, and automated session-level containment. The paper closes with a discussion of governance, explainability, and open data engineering challenges for AI-driven self-healing in financial cybersecurity.

Keywords: Data engineering; digital repair twin; financial cybersecurity; machine learning; anomaly detection; self-healing systems; SIEM; feature store; incident response.

Introduction

Financial institutions operate under a persistent and evolving threat landscape that includes account-takeover fraud, credential-stuffing campaigns, insider data exfiltration, and targeted intrusions against core banking and payment infrastructure. The cost of delayed detection in this domain is measured not only in direct financial loss but in regulatory exposure and customer trust. Security operations centers have traditionally relied on SIEM platforms and human analysts to triage alerts, a model that struggles to keep pace with both the volume of telemetry generated by modern financial infrastructure and the speed at which automated attacks unfold.

The Digital Repair Twin (DRT) concept extends the digital twin — a continuously updated virtual model of a system — into a diagnostic and action-capable form specifically suited to this challenge. Rather than simply visualizing security posture, a DRT for financial cybersecurity ingests telemetry, maintains a live model of accounts, sessions, and infrastructure state, applies machine learning to detect and localize anomalies, and executes or recommends bounded remediation actions such as session termination, credential invalidation, or network segmentation. The feasibility and reliability of such a system depends critically on the underlying data engineering architecture, which is the focus of this paper.

This paper makes three contributions. First, it presents a five-layer reference architecture for data engineering in AI-based DRTs applied to financial cybersecurity, detailing the technologies and data artifacts appropriate to each layer. Second, it provides a structured comparison, in tabular form, of pipeline layers and of evaluation metrics suited to operational assessment of such systems. Third, it illustrates the architecture through two figures and a simulated case

Corresponding Author: Chum Derin, Independent Researcher, Hong Kong SAR China

How to cite this article: Derin, C. (2025). Data Engineering Architectures for AI-Based Digital Repair Twins in Financial Cybersecurity *Journal of Science, Technology and Social Transformation* 1(1), 64-68.

Source of support: Nil

Conflict of interest: None

study of a credential-stuffing incident, showing how the components interact end to end.

Background

Digital Twins and Repair Twins in Security Operations

Digital twin techniques have been applied to IT and security operations to provide a live, structured view of infrastructure state, complementing traditional log-based monitoring. Extending this to a repair twin adds an explicit diagnostic and action layer, aligning with the broader movement toward AIOps and security orchestration, automation, and response (SOAR) platforms, which apply automation to reduce analyst workload for well-understood incident classes.

Data Engineering Challenges Unique to Financial Cybersecurity

Financial cybersecurity telemetry presents distinctive data engineering challenges relative to generic IT security monitoring. Data volumes are large and highly bursty during attack events; data sensitivity is high, since telemetry often includes account identifiers, session tokens, and behavioral data subject to financial privacy regulation; and correlation must often span multiple systems — authentication services, core banking platforms, payment gateways, and fraud engines — that were not originally designed to share a common data model. A DRT's data engineering layer must therefore unify heterogeneous telemetry into a consistent representation while enforcing strict access controls and retention policies.

Reference Architecture

Figure 1 presents the proposed five-layer reference architecture. Telemetry flows upward from ingestion through feature computation and model inference to reasoning and, where approved, automated action, while outcome data flows back into the ingestion layer to support continuous model improvement.

Table 1 summarizes the function, representative technologies, and key data artifacts associated with each layer of the architecture.

Data Pipeline Design for Financial Cybersecurity

Ingestion and Normalization

The ingestion layer must normalize telemetry from authentication services, core banking systems, payment gateways, endpoint detection agents, and network sensors into a common schema before downstream processing. Given the bursty nature of attack traffic, ingestion pipelines are typically designed with elastic throughput and backpressure handling to avoid data loss during high-volume events such as a credential-stuffing campaign, which can generate orders of magnitude more authentication attempts than baseline traffic within minutes.

Entity-Graph Modeling for Twin State

Unlike generic infrastructure monitoring, financial cybersecurity DRTs benefit from an explicit entity graph linking accounts, devices, sessions, and IP ranges, since many attack patterns manifest as unusual relationships

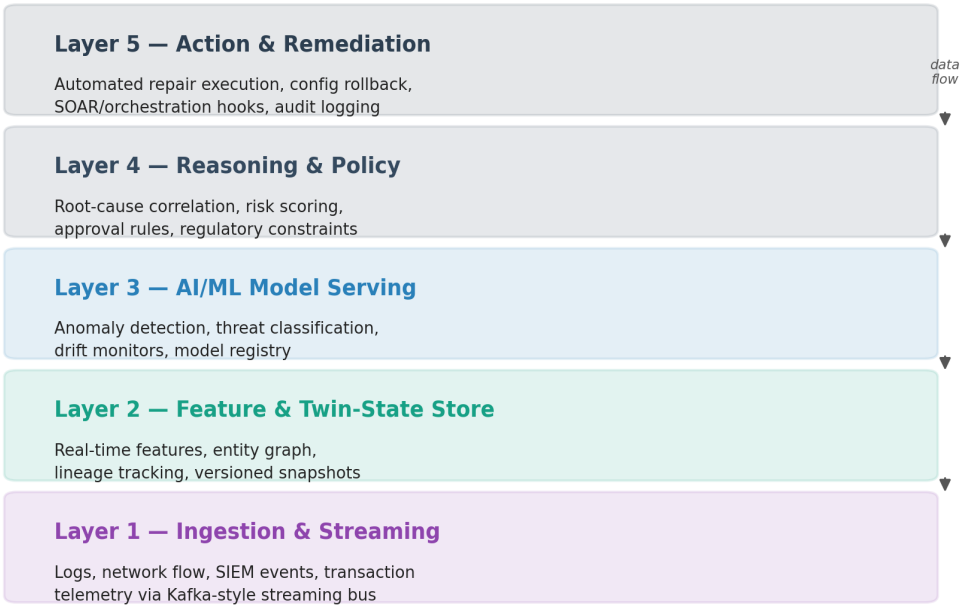

Figure 1: Five-layer data engineering architecture for AI-based Digital Repair Twins in financial cybersecurity.

Table 1: Data engineering layers, functions, and representative artifacts.

Layer	*Primary Function*	*Representative Technologies*	*Key Data Artifacts*
Ingestion & Streaming	Real-time collection of security and transaction telemetry	Kafka-style event bus, log shippers, network flow collectors	Raw events, SIEM alerts, flow records
Feature & Twin-State Store	Consistent feature computation and live entity-graph state	Feature store, graph database, time-series store	Rolling features, entity graph snapshots, lineage metadata
AI/ML Model Serving	Anomaly detection, threat classification, drift monitoring	Low-latency inference service, model registry	Anomaly scores, classification labels, drift alerts
Reasoning & Policy	Root-cause correlation and remediation decisioning	Rule engine, causal inference module, risk scorer	Diagnosis reports, risk scores, policy decisions
Action & Remediation	Execution of approved repair actions with full audit trail	Orchestration/SOAR tooling, config management APIs	Action logs, audit records, outcome telemetry

between entities — for example, a single device attempting authentication against many distinct accounts, or a single account being accessed from geographically improbable session origins in rapid succession. Maintaining this graph as a continuously updated twin-state store allows root-cause and pattern-detection models to reason over relationships rather than isolated events.

Governance, Lineage, and Data Minimization

Because financial cybersecurity telemetry frequently includes regulated personal and account data, the data engineering layer must enforce field-level access controls, tokenization of sensitive identifiers, and data minimization at the point of ingestion wherever feasible. Full lineage tracking, recording which data and model version produced a given diagnosis or action, is essential both for regulatory audit and for post-incident forensic review, particularly where an automated action affected a customer-facing account or session.

AI/ML Model Layer

The model serving layer hosts several complementary model classes operating on the data produced by the layers below. Unsupervised anomaly detection models, including isolation forests and autoencoders, flag statistically unusual authentication and transaction patterns without requiring labeled examples of every attack type. Supervised classifiers, trained on historical labeled incidents, provide finer-grained classification of detected anomalies into categories such as credential stuffing, account takeover, or insider misuse.

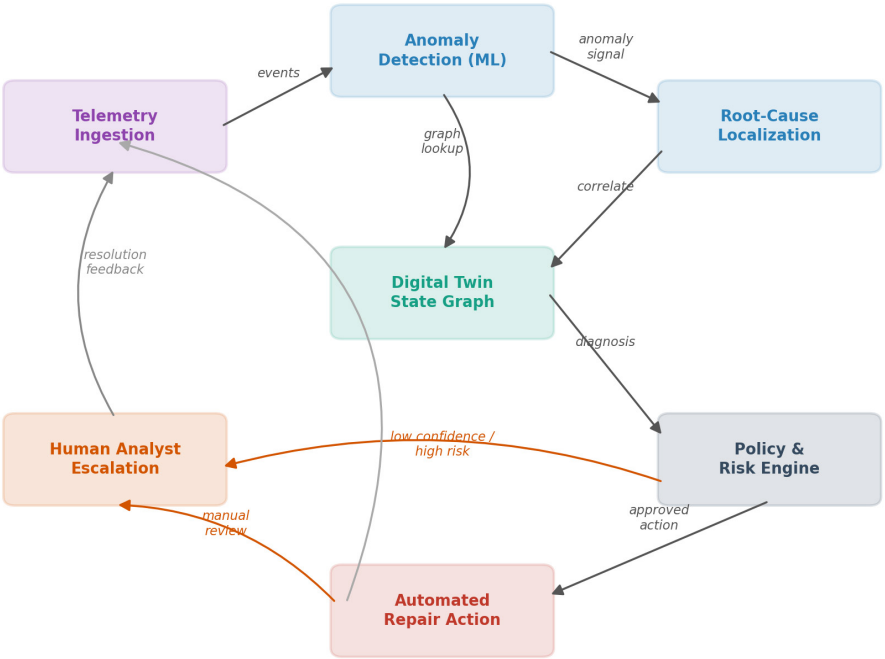

Figure 2 : Closed-loop detection-to-repair workflow within the Digital Repari Twin.

Table 2: Proposed evaluation metrics for AI-based Digital Repair Twins in financial cybersecurity.

Metric	*Definition*	*Illustrative Target Range*
Mean Time to Detect (MTTD)	Interval between fault onset and confirmed anomaly flag	Seconds to low minutes
Mean Time to Repair (MTTR)	Interval between detection and successful remediation	Under five minutes for pre-approved actions
False-Positive Remediation Rate	Share of automated actions later found unnecessary on review	As low as operationally achievable, tracked continuously
Root-Cause Localization Accuracy	Proportion of incidents where the flagged cause matches confirmed cause	High accuracy with confidence thresholds
Escalation Rate	Share of incidents routed to human analysts rather than auto-resolved	Tracked to gauge automation coverage over time

Graph-based models operating on the twin-state entity graph support root-cause localization by identifying which entities and relationships are most implicated in an observed anomaly. Drift-monitoring models track the statistical stability of upstream features and downstream model performance over time, triggering retraining or human review when behavior shifts beyond expected bounds, which is particularly important in a domain where attacker behavior adapts continuously to evade static defenses.

Simulated Case Study: Credential-Stuffing Detection and Containment

Consider a simulated scenario in which an attacker launches a credential-stuffing campaign against an online banking portal, using a large set of previously breached username-password pairs sourced from unrelated services. The ingestion layer captures a sharp rise in authentication attempts across a distinct pool of source IP addresses. The anomaly detection model flags the aggregate failed-login rate as a statistically significant deviation from baseline within the observation window. The root-cause localization model, operating over the twin-state entity graph, identifies that the anomalous traffic is concentrated among a set of source IP ranges associated with known proxy infrastructure and that affected accounts show no prior association with those ranges, distinguishing the event from a legitimate regional traffic shift.

Figure 2 illustrates the resulting closed-loop workflow: the anomaly signal and root-cause diagnosis are passed to the policy and risk engine, which evaluates the incident against pre-approved response tiers. Because the identified action — temporary rate-limiting and step-up authentication for the affected IP ranges — is reversible and narrowly scoped, it is executed automatically by the action layer, while accounts showing signs of successful unauthorized access are flagged for mandatory credential reset and escalated to a human analyst for review rather than being handled automatically, reflecting the higher risk associated with actions that directly affect customer account access.

Evaluation Considerations

Assessing the operational effectiveness of a financial cybersecurity DRT requires metrics that capture both detection speed and the safety of automated action. Table 2 summarizes five metrics proposed for this purpose, spanning detection latency, repair latency, false-positive remediation, root-cause accuracy, and the rate at which incidents are escalated to human analysts rather than resolved automatically. Tracking escalation rate alongside false-positive remediation rate is particularly important, since an architecture tuned purely to minimize escalation risks increasing unsafe automated action, while an architecture tuned purely to minimize false positives risks reverting to a largely manual, slower incident response process.

Governance and Compliance

Data engineering architectures supporting automated security remediation in financial services must align with applicable regulatory frameworks governing data protection, financial recordkeeping, and operational resilience. In practice, this shapes several architectural decisions: sensitive fields are tokenized or encrypted at ingestion rather than relying on downstream masking; retention policies are enforced within the feature and twin-state store rather than left to ad hoc cleanup; and every automated action is logged with sufficient detail — triggering signal, model version, policy rule, and outcome — to support both internal audit and external regulatory examination.

Open Challenges and Future Directions

Several data engineering challenges remain open. Unifying telemetry schemas across legacy core banking systems and modern cloud-native services remains labor-intensive and is a common source of blind spots in twin-state accuracy. Maintaining entity-graph freshness at scale, particularly for institutions with very large customer bases, requires careful attention to storage and query latency trade-offs. Cross-institution threat intelligence sharing could improve root-cause accuracy for coordinated attacks but remains

constrained by competitive and regulatory barriers, motivating continued interest in privacy-preserving collaborative learning techniques. Finally, establishing shared, realistic benchmark datasets for financial cybersecurity DRTs would allow more rigorous comparison of architectural choices across institutions and vendors, an area that remains underdeveloped relative to the maturity of the underlying AI/ML techniques.

Conclusion

This paper has presented a data engineering reference architecture for AI-based Digital Repair Twins in financial cybersecurity, spanning ingestion, feature and twin-state storage, model serving, reasoning and policy, and automated action layers. Through a tabular comparison of architectural layers and evaluation metrics, together with a simulated credential-stuffing case study, we have shown how disciplined data engineering — particularly entity-graph modeling, lineage tracking, and governance-aware pipeline design — underpins the ability of a DRT to detect, diagnose, and safely contain security incidents in financial infrastructure. As attack automation continues to increase in speed and sophistication, the resilience of financial institutions will depend increasingly on data engineering architectures capable of supporting this kind of AI-driven, closed-loop self-healing.

References

[1] Satish Kumar Nalluri, Venkata Krishna Bharadwaj Parasaram & Varun Teja Bathini, "Machine Learning-Based Management Models for Scalable and Resilient Industrial Platforms", International Journal of Engineering Research and Modern Education, Volume 1, Issue 1, Page Number 760-785, 2016. https://doi.org/10.5281/zenodo.19634396

[2] Nalluri, S. K., & Parasaram, V. K. B. (2016). Early Approaches to Robotic Process Automation in Enterprise Systems. International Journal of Humanities and Information Technology, 1(01), 12-28. https://doi.org/10.21590/ijhit.01.01.06

[3] 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.

[4] 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.

[5] 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.

[6] Venkata Krishna Bharadwaj Parasaram, Satish Kumar Nalluri & Varun Teja Bathini, "Artificial Intelligence Driven Management Systems for Optimizing Efficiency in Smart Industrial Environments", International Journal of Multidisciplinary Research and Modern Education, Volume 1, Issue 2, Page Number 489-514, 2015. https://doi.org/10.5281/zenodo.19634549

[7] Parasa, M. (2021). Encryption-aware data integrity and quality controls in SAP SuccessFactors integrations using machine learning and cryptographic hash chains for tamper detection. International Journal of Computer Technology and Electronics Communication, 4(6), 4304–4316. https://doi.org/10.15680/IJCTECE.2021.0406014

[8] Manne, V. T. (2025, December). AI-Powered Fraud Detection in Payments Using Long-Term Behavior Sequence Modeling. In 2025 International Conference on Computational Innovations and Sustainable Technologies (ICCIST) (Vol. 1, pp. 1-7). IEEE.

[9] Barua, S. (2023). Hybrid Electro-membrane Reactors for Decentralized Removal of Forever Chemicals From Industrial Wastewater. SAMRIDDHI: A Journal of Physical Sciences, Engineering and Technology, 15(04), 461-468.

[10] Parasa, M. (2024). Mapping the Latent Risk Layer in Enterprise Platforms: A Practical Model for Workforce Data Integrity, Access Behavior, and Cyber Threat Detection. Advanced International Journal of Multidisciplinary Research, 2(6). https://doi.org/10.62127/aijmr.2024.v02i06.1366

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