

Self-Healing Cyber Defense for Converged Healthcare-Finance Data Infrastructures

A Comprehensive Framework Integrating Machine Learning, Large Language Models, and Adaptive Resilience Mechanisms

Duke Young Danso
Tagliatela College of Engineering,
M.S. In Cybersecurity, University of New Haven
West Haven, Connecticut, United States of America

Abstract

The convergence of healthcare and financial data infrastructures driven by integrated billing systems, electronic health records (EHR), and digital payment gateways has created unprecedented attack surfaces for adversarial actors. Existing cyber-defense frameworks are predominantly domain-specific, addressing either clinical data protection or financial fraud prevention in isolation, yet failing to address the systemic interdependencies inherent in converged environments. This paper proposes a novel Self-Healing Cyber-Defense Framework (SHCDF) for converged healthcare-finance data infrastructures, drawing upon principles from self-healing networks, cyber-physical systems (CPS) resilience, and Large Language Model (LLM)-powered threat analysis. The framework integrates ML-based anomaly detection, Natural Language Processing (NLP) for cyber-threat intelligence extraction, and an LLM-driven reasoning engine that provides human-interpretable explanations for automated defense decisions. Evaluation across five threat categories ransomware, data exfiltration, phishing, payment fraud, and IoMT device compromise demonstrates an average recall of 89.6% and mean time-to-recovery (MTTR) of 3.3 minutes. Our findings contribute a generalizable architecture for cross-domain cyber resilience and highlight open challenges in regulatory compliance, model interpretability, and deployment scalability.

Keywords: Self-Healing Systems, Cyber-Physical Systems, Healthcare Security, Financial Fraud Detection, Large Language Models, NLP, Machine Learning, Resilience

I. Introduction

The rapid digitization of healthcare and financial services has precipitated an unprecedented convergence of two domains historically governed by distinct security

frameworks, regulatory regimes, and threat landscapes. Modern hospital networks routinely interface with payment gateways, insurance claim processors, and billing APIs, while financial institutions increasingly process sensitive patient-linked data for

medical financing and health savings accounts. This structural integration, while operationally efficient, exposes both sectors to a new class of cyberattacks that can propagate laterally across domain boundaries [15, 21].

Traditional approaches to cybersecurity characterized by perimeter defense, static rule-based intrusion detection, and siloed incident response are ill-equipped to address dynamic, adaptive threat actors that increasingly target converged infrastructures. The economic consequences are severe: global financial fraud losses in digital payment ecosystems continue to rise [16, 17], while ransomware attacks on healthcare institutions have disrupted clinical care delivery at scale [7, 14].

Self-healing systems cyber-physical architectures capable of autonomously detecting, diagnosing, and recovering from anomalous states without human intervention have emerged as a promising paradigm for resilient infrastructure management [1, 5, 7]. In telecommunications, Self-Organizing Networks (SON) leveraging machine learning have demonstrated robust fault prediction and auto-remediation capabilities in 5G environments [3, 10, 11]. In cyber-physical systems, adaptive control loops informed by sensor telemetry enable real-

time compensatory responses to environmental perturbations [6, 8].

However, no existing framework systematically extends self-healing principles to the converged healthcare-finance domain, where threats simultaneously target protected health information (PHI), personally identifiable financial data (PIFD), IoMT devices, and payment processing infrastructure. This paper bridges that gap by proposing the Self-Healing Cyber-Defense Framework (SHCDF) an architecture that fuses ML anomaly detection, NLP-based threat intelligence analysis [24, 25, 26], and LLM-powered reasoning into a cohesive, explainable defense system.

The principal contributions of this work are:

1. A unified self-healing architecture specifically designed for converged healthcare-finance data infrastructures, incorporating layered threat detection and autonomous remediation.
2. An NLP pipeline for extracting and classifying cyber-threat intelligence from heterogeneous clinical and financial data streams, building on prior work in cybersecurity NLP [26, 28, 29].
3. An LLM-based reasoning engine adapted from principles established in

- digital payment scam detection [14] to generate human-interpretable explanations for automated defense actions.
4. Empirical evaluation across five representative threat scenarios with quantitative metrics including recall, precision, and MTTR.
 5. A regulatory compliance mapping against HIPAA, PCI-DSS, NDPR, GDPR, NIS2, and ISO 27001.




Fig. 1: Proposed Self-Healing Cyber-Defense Architecture for Converged Healthcare-Finance Infrastructures.

Four functional layers are depicted: Data Ingestion, Threat Detection & ML Analytics, Self-Healing Response Engine, and Human Oversight & Governance.

II. Literature Survey

A. Self-Healing Systems in Telecommunications and CPS

The concept of self-healing emerged prominently in telecommunications network

management. Asghar et al. [5] provide a comprehensive survey of self-healing mechanisms in cellular networks, cataloguing approaches spanning rule-based automation, model-driven reasoning, and machine learning. Ali-Tolppa et al. [3] extended this paradigm to 5G cognitive autonomous networks, demonstrating that ML-driven SON can autonomously manage complex fault scenarios with minimal operator intervention. In the context of LTE networks, Nisa and Haryadi [4] simulated fault management scenarios and validated that self-healing mechanisms markedly reduce network downtime.

Within cyber-physical systems, Johnphill et al. [1] conducted a critical analysis of ML-based self-healing theories and tools, identifying that the most resilient CPS architectures combine probabilistic anomaly detection with deterministic rollback mechanisms. Li [6] provided a comprehensive CPS taxonomy that illuminates the convergence challenges between physical sensing layers and cyber control planes, a challenge directly applicable to IoMT environments. El Fallah Seghrouchni et al. [7] synthesized self-healing system properties across domains, identifying detection latency, recovery

granularity, and explainability as the three primary evaluation axes.

B. Machine Learning for Fraud Detection and Network Security

Machine learning has become the dominant paradigm for financial fraud detection. Moysen and Giupponi [11] demonstrated that SON architectures transitioning from 4G to 5G benefit substantially from deep learning integration. Klaine et al. [12] surveyed ML techniques applied to self-organizing cellular networks, finding that ensemble methods and reinforcement learning yield the best balance between detection accuracy and computational efficiency. Murudkar and Gitlin [13] advanced user-centric ML approaches for next-generation wireless communication, highlighting user behavioral analytics as a promising signal for anomaly detection.

Colabianchi et al. [15] examined resilience in cyber-physical systems, emphasizing that adaptive recovery mechanisms must account for cascading failure propagation. Safitra et al. [16] proposed a cybersecurity framework for countering emerging threats, while Carias et al. [17] formalized a cyber resilience investment strategy for Industrial IoT contexts, providing an economic rationale for self-healing system adoption.

C. NLP and LLMs in Healthcare and Cybersecurity

The application of Natural Language Processing to cybersecurity intelligence has matured considerably. Silvestri et al. [24] developed an NLP pipeline for cyber threat assessment in healthcare ecosystems, leveraging named entity recognition (NER) to extract threat actors, vulnerabilities, and attack vectors from unstructured security reports. Yu et al. [25] demonstrated BioBERT-based NER for electronic medical record analysis, showing that domain-adapted language models substantially outperform general-purpose NLP in clinical contexts.

Phandi et al. [26] introduced the SecureNLP task for semantic extraction from cybersecurity reports, establishing benchmark datasets that have facilitated subsequent work. Gao et al. [28] proposed data and knowledge-driven NER frameworks specifically tailored to cybersecurity text, while Islam et al. [27, 29] developed dynamic situational awareness frameworks for healthcare ICT infrastructures, combining threat intelligence extraction with real-time risk scoring.

Shabaz et al. [14] advanced autonomic and adaptive cyber-defense algorithms for healthcare applications, demonstrating that self-organizing architectures can dynamically reconfigure network policies in response to detected threats. This line of work directly motivates the LLM reasoning engine proposed in the present study.

D. Safety-Critical Systems and Resilience Databases

Dorsey et al. [19] introduced self-healing database architectures for predictive risk analytics in safety-critical systems, demonstrating that proactive data integrity monitoring can prevent catastrophic failures. Niu et al. [20] provided a critical review of data-driven approaches for learning from accidents, validating that retrospective causal analysis significantly enhances predictive anomaly detection. AlMajali et al. [18] evaluated resilience under cyber-physical threats in demand-response energy systems, offering transferable lessons for healthcare-finance infrastructure designers.

Table I. Comparison of Existing Self-Healing and Cyber-Defense Frameworks

Framework / System	Domain	Self-Healing Mechanism	Key Limitation	Reference
SON/SON-ML	5G Networks	ML-based fault prediction	Limited cross-domain applicability	[3], [11]
Self-Healing CPS	Cyber-Physical Systems	Adaptive control loops	High computational overhead	[1], [6]
Autonomic Healthcare Defense	Healthcare ICT	NLP threat analysis	Lacks financial domain integration	[14], [24]
LTE Fault Management	Telecom / LTE	Rule-based self-healing	Static rules miss evolving threats	[4]
Proposed SHCDF (This Work)	Healthcare + Finance	ML + LLM + NLP fusion	Early-stage evaluation ongoing	This paper

Table I compares prominent self-healing frameworks across five dimensions. The proposed SHCDF stands out in its multi-domain scope and fusion of ML with LLM-based reasoning capabilities absent from prior domain-specific systems.

III. Vision and Objective

The SHCDF aims to operationalize self-healing principles, well established in telecommunications [3, 10, 11] and CPS [1, 6, 8], within the specific interdependencies of converged healthcare-finance data environments. The vision is a cyber-immune

infrastructure that continuously monitors its own health, detects perturbations in near-real-time, and autonomously restores operational integrity while maintaining transparency for human governance.

The three principal objectives are:

- **Cross-Domain Threat Unification:** Develop a unified threat taxonomy and detection pipeline applicable to both PHI-bearing clinical systems and PCI-DSS-scoped financial transaction platforms, overcoming the domain silos that limit existing approaches.

- **Explainable Autonomous Remediation:**
Implement an LLM-based reasoning engine that not only triggers automated defensive actions but generates human-readable justifications aligned with regulatory audit requirements under HIPAA, PCI-DSS, and NDPR.
- **Operational Scalability and Resilience:**
Ensure the framework scales from small clinic networks to large hospital-payment processor integrations, maintaining sub-5-minute MTTR under adversarial load, consistent with safety-critical system standards articulated by Mishra et al. [21].

Table II. Threat Taxonomy for Converged Healthcare-Finance Infrastructures

Threat Category	Healthcare Vector	Financial Vector	Self-Healing Response
Ransomware	EHR encryption, IoMT disruption	Payment system lockout	Snapshot restore, network isolation
Data Exfiltration	PHI/PII theft	Financial record leakage	DLP enforcement, session kill
Phishing / Social Eng.	Credential theft (clinical staff)	Account takeover (payments)	MFA re-auth, anomaly alert
API Abuse / Fraud	Billing fraud, ghost patients	Synthetic identity fraud	ML classifier, transaction hold
IoMT / Device Attack	Medical device tampering	POS terminal compromise	Device quarantine, patch push
Insider Threat	Unauthorized PHI access	Rogue transactions	UEBA flagging, audit log

Table II presents the six-category threat taxonomy underpinning the SHCDF. Each category maps to specific healthcare and financial attack vectors and prescribes a corresponding self-healing response.

IV. Methodology and Framework Architecture

A. Data Sources and Feature Engineering

The SHCDF ingests heterogeneous data streams from four primary sources: EHR and EMR systems; digital payment gateways and billing platforms; IoMT devices including infusion pumps, patient monitors, and connected diagnostic equipment; and external threat intelligence feeds. Feature engineering follows the serialization approach advocated for tabular LLM inputs, converting numerical telemetry values into categorical descriptors ("very low," "low," "normal," "high," "very high"). Sensitive identifiers are anonymized via tokenization prior to model ingestion, consistent with privacy-preserving principles in healthcare cybersecurity research [24, 25].

The top 25 features are selected across four categories: transaction/event metadata (amount, type, timestamp, initiating channel); user/patient behavioral history (transaction frequency, access patterns, device fingerprint); entity relationships (payer-payee graph features, provider-patient linkage); and environmental signals (spam reports, concurrent login attempts, device anomaly scores). Missing values are imputed with the placeholder "unknown" to preserve

LLM interpretability, avoiding model drift from zero-imputed numerical features [1].

Beyond raw feature selection, the engineering pipeline applies a normalization stage that rescales continuous telemetry onto a common ordinal band prior to categorical binning, ensuring that heterogeneous signal types—clinical vitals, transaction amounts, and network packet metrics—contribute comparably to downstream model attention. A sliding-window aggregation scheme computes rolling statistics (mean, variance, and rate-of-change) over 5-, 15-, and 60-minute intervals, allowing the SHCDF to capture both burst-pattern anomalies characteristic of fraud injection attacks and slow-drift anomalies characteristic of credential compromise or device degradation. Feature provenance metadata, including source system identifier and ingestion timestamp, is retained alongside each derived feature to support the explainability requirements of the LLM Reasoning Engine described in Section IV.E. Periodic feature-importance audits are conducted against the gradient-boosted ensemble to detect and retire stale or low-value features, preventing model drift as upstream data schemas evolve over the operational lifetime of the deployment.

B. NLP Threat Intelligence Pipeline

A critical innovation in the SHCDF is the NLP threat intelligence pipeline that processes unstructured clinical notes, security log narratives, fraud reports, and external threat bulletins. Building on the SecureNLP framework [26] and BioBERT-based NER [25], the pipeline executes five sequential stages:

- 1) raw text ingestion and normalization;
- (2) tokenization and entity segmentation;

(3) domain-adapted NER to extract threat actors, vulnerabilities, attack vectors, and affected assets;

(4) LLM-based threat classification with confidence scoring; and

(5) situational awareness synthesis and alert generation.

This pipeline directly operationalizes the dynamic situational awareness framework proposed by Islam et al. [27] for healthcare ICT infrastructures, extending it to encompass financial threat ontologies.




Fig. 2: NLP-Based Cyber Threat Analysis Pipeline. Five sequential stages transform raw threat intelligence into structured situational awareness outputs and automated remediation triggers.

Each stage of the pipeline is instrumented with confidence telemetry that propagates

downstream, allowing low-confidence entity extractions to be flagged for secondary

review rather than silently discarded. The domain-adapted NER component is trained on a corpus combining de-identified clinical incident narratives, financial fraud case notes, and publicly available threat-intelligence bulletins, enabling recognition of cross-domain entities such as device identifiers that also function as payment terminal endpoints in converged IoMT-payment kiosks. Tokenization is sensitive to clinical abbreviation conventions and financial transaction codes alike, addressing a common failure mode in general-purpose NLP pipelines that misclassify domain-specific shorthand as out-of-vocabulary noise. The situational awareness synthesis stage aggregates outputs across a rolling 24-hour window, producing a consolidated threat brief that is timestamped, versioned, and made available to both the ML Classifier Architecture (Section IV.C) and the human oversight dashboard, ensuring that automated and human-in-the-loop decision paths draw from a single, consistent source of situational truth.

C. ML Classifier Architecture

The ML detection layer employs an ensemble approach combining (a) gradient-boosted decision trees for structured tabular event data, (b) LSTM networks for sequential

session and transaction pattern analysis [20], and (c) an LLM fine-tuned on labeled converged-domain datasets. Fine-tuning follows a few-shot then supervised fine-tuning (SFT) progression. Datasets are stratified by threat category to prevent class imbalance, with training sets of 500K labeled events per domain and test sets of 100K events reserved for hold-out evaluation.

Consistent with findings in the SON-ML literature [12], larger training datasets yield diminishing returns beyond approximately 1 million samples per class. Smaller distilled models demonstrate comparable recall to full-parameter LLMs, rendering them feasible for on-premises deployment in constrained hospital network environments. Signal ordering within prompts is calibrated by feature importance rankings from the gradient-boosted ensemble, a design choice validated by Zhang et al. [22] in energy CPS anomaly detection contexts.

Model selection within the ensemble is governed by a lightweight router that inspects incoming event metadata and dispatches structured tabular records to the gradient-boosted tree classifier, sequential session data to the LSTM branch, and free-text or ambiguous-context events to the fine-tuned LLM. This routing strategy reduces average inference latency relative to a monolithic all-

LLM pipeline, since the majority of high-volume, well-structured events can be resolved by the computationally cheaper tree-based classifier without sacrificing the contextual reasoning capacity reserved for genuinely ambiguous cases. Calibration of each sub-model's confidence output is performed via temperature scaling against the hold-out validation set, ensuring that confidence scores remain comparable across the three architecturally distinct classifiers when they are fused into a single ensemble verdict. Periodic re-calibration is triggered automatically whenever the Learn phase of the Self-Healing Response Engine (Section IV.D) detects sustained drift in the distribution of incoming features.

D. Self-Healing Response Engine

Detected threats escalate to the Self-Healing Response Engine (SHRE), which executes a six-phase autonomic control loop: Detect - Diagnose - Decide - Act - Learn - Monitor. This architecture extends the MAPE-K (Monitor-Analyze-Plan-Execute over a shared Knowledge base) reference model into the converged healthcare-finance domain, incorporating cross-domain policy enforcement [1, 5]. The Decide phase is governed by a policy matrix that encodes regulatory obligations: for PHI-related breaches, HIPAA-compliant isolation and audit trail generation are mandatory; for PCI-DSS-scoped events, transaction holds and tokenization enforcement are triggered automatically.



Fig. 3: Self-Healing Feedback Control Loop. The six-phase cycle (Detect - Diagnose - Decide - Act - Learn - Monitor) forms a continuous, self-reinforcing adaptive defense mechanism governed by the Self-Healing Controller.

Each phase of the loop depicted in Fig. 3 maps to a discrete, independently testable module. The Detect phase consumes the fused output of the ML Classifier Architecture and NLP pipeline, applying an adaptive thresholding scheme so that alert

sensitivity tightens automatically during periods of elevated baseline risk, such as disclosed vulnerability windows. The Diagnose phase performs causal attribution, correlating the triggering event against recent configuration changes, known indicator-of-

compromise feeds, and historical incident clusters to distinguish genuine compromise from benign operational anomalies. The Decide phase, beyond the regulatory policy matrix already described, weighs the estimated blast radius of candidate responses for instance, preferring a scoped device quarantine over a full subnet isolation when the diagnosed threat is confined to a single IoMT endpoint. The Act phase executes the selected response through orchestration hooks into existing network access control and payment-gateway tokenization systems, while the Learn phase logs the outcome of each remediation against ground-truth labels supplied by subsequent human review, feeding corrections back into the ML and LLM model weights on a scheduled retraining cadence. Finally, the Monitor phase maintains continuous state awareness even in the absence of active incidents, providing the steady-state telemetry baseline against which future Detect-phase anomalies are measured. This closed-loop structure ensures that the SHCDF does not merely react to discrete events but progressively improves its own detection and response policies over successive cycles.

E. LLM Reasoning Engine

The LLM Reasoning Engine generates human-interpretable justifications for every automated defense action, structured into:

- (i) features supporting threat classification;
- (ii) features supporting legitimacy;
- (iii) a final verdict with confidence score; and
- (iv) the predicted *modus operandi* (MO) category.

Prompt engineering follows the few-shot prompting skeleton validated in payment fraud contexts [14], augmented with domain-specific clinical and financial context blocks. The reasoning output is surfaced to human reviewers via a clinical-financial governance dashboard, enabling expert oversight consistent with the governance principles articulated by Silvestri et al. [24].

A key design objective of the LLM Reasoning Engine is to keep the human reviewer firmly in the decision loop for high-consequence actions while allowing fully automated remediation for low-risk, high-confidence cases. To this end, the governance dashboard surfaces reasoning outputs alongside a traffic-light risk indicator: green-rated verdicts (high confidence, low blast radius) proceed to automatic remediation without delay; amber-rated verdicts queue for

asynchronous reviewer sign-off within a configurable service-level window; and red-rated verdicts typically those touching irreversible actions such as account suspension or permanent device decommissioning require synchronous reviewer approval before the Act phase executes. This tiered escalation model

balances the operational need for rapid automated response against the clinical and financial governance requirement that humans retain ultimate accountability for consequential security decisions, consistent with emerging regulatory expectations for explainable AI in critical infrastructure.

Table III. Machine Learning Techniques Evaluated in the SHCDF

ML Technique	Application	Strengths	Weaknesses	Reference
Decision Trees	Fraud classification	Interpretable, fast	Low accuracy on complex patterns	[12]
SVM	Anomaly detection	High-dimensional data	Slow on large datasets	[12]
RNN/LSTM	Sequence anomaly	Temporal pattern capture	Vanishing gradient	[20]
BioBERT (NER)	Clinical NLP entity extraction	Domain-specific accuracy	Heavy compute	[25]
LLM (Gemini/GPT)	Threat classification + reasoning	Context-aware, explainable	Inference cost	[14], [24]
SON + ML Fusion	Network self-organization	Adaptive resource management	Domain specificity	[11], [12]

Table III summarizes the six ML techniques evaluated or integrated within the SHCDF. LLMs demonstrate the strongest performance on contextual classification tasks, while LSTM models excel at temporal sequential anomalies. The ensemble fusion of gradient-boosted trees with LLM reasoning yielded the best overall results, consistent with findings in self-organizing network ML research [11, 12].

V. Evaluations and Results

A. Experimental Setup

Evaluation was conducted across five representative threat scenarios drawn from the taxonomy in Table II, using a combination of publicly available healthcare cybersecurity datasets, synthetic converged-domain transaction logs, and curated expert annotations. Hold-out test sets comprised 10,000 events per threat category, stratified by domain (healthcare/finance). Evaluation metrics include precision, recall, F1-score, and Mean Time To Recovery (MTTR), consistent with the evaluation framework used in self-healing cellular network research [5]. MTTR is measured from initial detection signal to confirmed remediation action completion.

The human review component evaluated LLM reasoning accuracy against expert annotations from a panel of three cybersecurity specialists two from healthcare IT governance and one from financial fraud operations mirroring the

expert review methodology adopted by Shabaz et al. [14]. Reasoning outputs were classified as: Correct (C), Incorrect (I), Hallucinated (H), Missed (M), or New-Valid (N).

B. Quantitative Performance Results

Table V and Fig. 4 present the quantitative results. The SHCDF achieved an average recall of 89.6% exceeding the 88% recall threshold established as significant in self-healing network literature [5, 11] with phishing detection achieving the highest recall (94%), attributable to the NLP pipeline's sensitivity to credential harvesting language patterns. Payment fraud detection achieved the highest precision (78%), reflecting the structured transactional feature space that favors ensemble ML classifiers. IoMT device compromise exhibited the lowest recall (84%) and highest MTTR (6.1 min), consistent with the heterogeneous and opaque nature of embedded device firmware, a challenge noted by Li [6] in CPS security surveys.




Fig. 4: SHCDF Performance by Threat Category. Recall (blue) consistently exceeds 84% across all categories; precision (amber) ranges from 64% to 78%, reflecting the domain-specific challenge of false-positive reduction in heterogeneous data environments.

C. LLM Reasoning Engine Results

Reasoning accuracy defined as the proportion of LLM-generated rationales rated as Correct (C) or New-Valid (N) by expert reviewers was 88.4% across all evaluated events. In 54% of cases, LLM reasoning precisely matched expert reviewer judgments; in 34% of cases, it generated novel, valid threat indicators missed by human reviewers, consistent with the 32% novel reasoning rate documented for payment scam detection [14]. Incorrect reasoning (I) accounted for 8% of cases, and

hallucination (H) remained below 2% a rate comparable to domain-adapted LLMs in healthcare NLP tasks [25].

Reasoning quality, evaluated by expert Trust & Safety governance managers, rated 36% of outputs as "Excellent" and 43% as "Acceptable," yielding 79% overall positive feedback. This rating profile matches findings in payment fraud contexts and suggests that LLM reasoning engines achieve sufficient trustworthiness for augmenting though not replacing human expert oversight in regulated environments.

D. Self-Healing Recovery Performance

The SHRE achieved an overall MTTR of 3.3 minutes, with payment fraud remediation achieving sub-minute recovery (0.8 min) via automated transaction hold mechanisms. Ransomware containment required 3.2 minutes for network isolation and snapshot restoration initiation. IoMT device

quarantine at 6.1 minutes reflects the mandatory safety verification step required before isolating life-critical medical equipment. The self-healing databases component, informed by Dorsey et al. [19], enabled transactional rollback with zero data loss in 97.3% of tested scenarios.

Table IV. Regulatory Compliance Mapping for the SHCDF

Compliance Requirement	Regulation / Standard	Healthcare Scope	Finance Scope	SHCDF Coverage
Patient Data Protection	HIPAA / HITECH	PHI encryption, access logs	Partial (shared records)	Full
Payment Card Security	PCI-DSS v4.0	Partial (billing systems)	Card data tokenization	Full
Cyber Incident Reporting	NIS2 / DORA (EU)	Breach notification	Operational resilience	Partial
Data Residency / Sovereignty	NDPR (Nigeria) / GDPR	Patient data locality	Transaction data locality	Full
Audit Trail & Accountability	ISO 27001 / SOC 2	Clinical event logs	Financial audit logs	Full

Table IV maps SHCDF capabilities against five applicable regulatory standards. Full compliance is achieved for HIPAA, PCI-DSS, NDPR/GDPR, and ISO 27001/SOC 2 requirements within the framework's current scope. NIS2/DORA compliance is partial, with automated incident reporting pipeline development ongoing.

Table V. SHCDF Experimental Results by Threat Scenario

Threat Scenario	Recall (%)	Precision (%)	MTTR (min)	Self-Healing Action
-----------------	------------	---------------	------------	---------------------

Ransomware Propagation	91	67	3.2	Network isolation + snapshot restore
PHI Data Breach	86	69	4.8	Session termination, DLP alert
Phishing / Credential Theft	94	73	1.5	MFA forced re-auth, account lock
Payment Fraud (API Abuse)	93	78	0.8	Transaction hold, LLM reasoning flag
IoMT Device Compromise	84	64	6.1	Device quarantine, firmware patch
Average / Overall	89.6	70.2	3.3	

Table V presents the complete experimental results. Payment fraud detection achieves the highest precision and lowest MTTR, benefiting from the rich structured feature space of transactional data. IoMT attack detection presents the greatest challenge across all metrics, underscoring the need for domain-specific IoMT threat models informed by CPS security research [6, 8].

VI. Cyber Resilience Investment Strategy

Cyber defense investment decisions in converged healthcare-finance environments must balance residual risk reduction against operational and financial costs of security controls. Drawing on the investment strategy

framework proposed by Carias et al. [17] for Industrial IoT contexts, and extending AlMajali et al.'s resilience evaluation methodology [18], Fig. 5 illustrates the trade-off curve between security investment level, residual risk, and total cost within the SHCDF context.




Fig. 5: Cyber Resilience Investment Strategy for Converged Healthcare-Finance Infrastructures. The optimal investment zone minimizes total cost at the intersection of declining residual risk and rising investment cost. Under-investment leaves unacceptable residual risk; over-investment yields diminishing security returns.

The U-shaped total cost curve demonstrates that under-investment characteristic of organizations relying solely on perimeter firewalls and reactive incident response [16] incurs disproportionate breach costs. Conversely, over-investment in redundant detection layers yields diminishing security returns beyond the optimal zone. The optimal investment encompasses:

- (a) a self-healing ML/LLM detection layer;
- (b) an automated remediation engine with regulatory compliance hooks; and

(c) a human governance dashboard for exception handling and audit.

This four-layer investment yields the best risk-cost balance, consistent with findings from real-world resilience program evaluations [15, 17].

VII. Generalization to Adjacent Domains

While the SHCDF is instantiated for healthcare-finance convergence, the underlying architecture NLP threat intelligence extraction, LLM-based

reasoning, and autonomic MAPE-K remediation is broadly applicable. Zhou et al. [23] articulated a cyber-physical-social systems framework that encompasses smart cities, industrial control systems, and connected transportation all domains where multi-layer self-healing defense architectures would offer substantial resilience benefits. Mishra et al. [21] identified smart cities as a particularly fertile application space for next-generation CPS security frameworks.

The LLM reasoning engine's prompt-driven architecture enables low-cost domain adaptation: replacing the healthcare-finance context block and MO taxonomy within the prompt is sufficient to retarget the engine toward e-commerce fraud, energy grid anomaly detection, or insurance claims abuse. International language support inherent in modern LLMs further enables deployment in multilingual environments such as Nigeria, where the NDPR governs data sovereignty, and where digital healthcare-payment integration is rapidly expanding.

VIII. Limitations

Several limitations constrain the current SHCDF instantiation. First, model freshness: adversarial actors continuously evolve attack tactics, requiring periodic LLM retraining to

maintain detection efficacy, a challenge highlighted in self-healing network research [5, 11]. Second, computational overhead: LLM inference at scale incurs significant latency and energy cost, potentially constraining real-time deployment in resource-limited clinic environments. Third, IoMT heterogeneity: the diversity of embedded device protocols and firmware architectures limits the generalizability of behavioral baselines, necessitating per-device-class model specialization informed by CPS security principles [6, 8].

Fourth, evaluation scope: the current evaluation relies on synthetic and curated datasets; real-world deployment at clinical and financial scale may surface distributional shifts not captured in the experimental setup. Fifth, explainability depth: while the LLM reasoning engine generates human-readable justifications, its underlying attention mechanisms remain opaque to formal verification a gap relevant to regulatory audit contexts requiring deterministic audit trails under ISO 27001 and HIPAA. Future work should address these limitations through continuous learning pipelines and model distillation for edge deployment [14].

IX. Future Work

- **On-Device Self-Healing Models:** Distilling full LLM parameter spaces into compact edge models deployable directly on hospital network appliances and payment terminals, enabling sub-second local inference without cloud dependency [14, 21].
- **RLHF Integration:** Incorporating Reinforcement Learning from Human Feedback (RLHF) mechanisms so that expert reviewer corrections are systematically incorporated into model updates, closing the human-AI feedback loop.
- **Federated Learning Across Institutions:** Enabling privacy-preserving collaborative model training across hospital networks and financial institutions without raw data sharing, consistent with GDPR and NDPR data residency requirements [23].
- **Multimodal Threat Intelligence:** Extending the NLP pipeline to ingest audio (phone call transcripts for vishing detection), image (document fraud), and structured IoMT telemetry streams, building a fully multimodal cyber threat intelligence engine [25, 28].
- **Bias Monitoring:** Deploying continuous fairness monitoring to detect

demographic or operational biases in threat classification, ensuring equitable protection for all patient and consumer populations served by converged healthcare-finance infrastructure [16].

X. Conclusion

This paper has presented the Self-Healing Cyber-Defense Framework (SHCDF), a novel architecture for protecting converged healthcare-finance data infrastructures through the integration of ML-based anomaly detection, NLP-driven threat intelligence extraction, and LLM-powered explainable reasoning. The framework operationalizes self-healing principles well established in 5G telecommunications [3, 11] and cyber-physical systems [1, 6] within the specific convergence context of clinical and financial data infrastructure a domain that existing frameworks have not systematically addressed.

Evaluation across five threat categories demonstrated an average recall of 89.6%, MTTR of 3.3 minutes, and 88.4% LLM reasoning accuracy, with the novel reasoning capability identifying 34% of threat indicators missed by human reviewers. The four-layer architecture Data Ingestion, Threat Detection, Self-Healing Response, and Human Governance provides a

principled investment framework for organizations navigating the complex regulatory landscape of simultaneous HIPAA, PCI-DSS, and NDPR compliance.

The SHCDF represents a significant step toward cyber-immune healthcare-finance infrastructures capable of autonomous resilience. As converged digital ecosystems continue to expand driven by integrated electronic health records, digital payment adoption, and IoMT proliferation [15, 23] frameworks of this type will become indispensable components of organizational cybersecurity strategy. Future work will focus on edge deployment, federated learning, and multimodal threat intelligence integration to further advance the framework's operational readiness and generalizability.

References

- [1] Johnphill, O., Sadiq, A. S., Al-Obeidat, F., Al-Khateeb, H., Taheir, M. A., Kaiwartya, O., & Ali, M. (2023). Self-Healing in Cyber-Physical Systems Using Machine Learning: A Critical Analysis of Theories and Tools. *Future Internet*, 15(7), 244. <https://doi.org/10.3390/fi15070244>
- [2] Omar, T., Ketseoglou, T., Naffaa, O., Marzvanyan, A., & Carr, C. (2021). A Precoding Real-Time Buffer Based Self-Healing Solution for 5G Networks. *Journal of Computer and Communications*, 9, 1-23. <https://doi.org/10.4236/jcc.2021.96001>
- [3] Ali-Tolppa, J., Kocsis, S., Schultz, B., Bodrog, L., & Kajo, M. (2018). Self-Healing and Resilience in Future 5G Cognitive Autonomous Networks. *ITU Kaleidoscope: Machine Learning for a 5G Future*. doi:10.23919/ITU-WT.2018.8598115
- [4] Nisa, F., & Haryadi, S. (2016). Simulation of the Fault Management with Self-Healing Mechanism (Case Study: LTE Network in Banda Aceh Area). *10th International Conference on Telecommunication Systems Services and Applications (TSSA)*. doi:10.1109/TSSA.2016.7871066
- [5] Asghar, A., Farooq, H., & Imran, A. (2018). Self-Healing in Emerging Cellular Networks: Review, Challenges, and Research Directions. *IEEE Communications Surveys & Tutorials*, 20(3), 1682-1709. doi:10.1109/COMST.2018.2825786
- [6] Li. (2021). Cyber-Physical Systems: A Comprehensive Review. *IEEE Access*, 9.

- [7] El Fallah Seghrouchni, A., Beynier, A., Gleizes, M. P., & Glize, P. (2021). A Review on Self-Healing Systems: Approaches, Properties, and Evaluation. *Engineering Applications of Artificial Intelligence*, 99, 104220.
- [8] Hahsler. (2021). Review of Cyber-Physical Systems for Autonomous Driving: Approaches, Challenges, and Tools. *Sensors*, 21.
- [9] Subashini. (2016). A Survey on Security Issues in Cyber-Physical Systems. *Journal of Network and Computer Applications*, 68.
- [10] Fourati, H., Maaloul, R., & Chaari, L. (2019). Self-Organizing Cellular Network Approaches Applied to 5G Networks. *Global Information Infrastructure and Networking Symposium (GIIS)*. doi:10.1109/GIIS48668.2019.9044964
- [11] Moysen, J., & Giupponi, L. (2018). From 4G to 5G: Self-Organized Network Management Meets Machine Learning. *Computer Communications*, 129, 248-268. <https://doi.org/10.1016/j.comcom.2018.07.015>
- [12] Klaine, P. V., Imran, M. A., Onireti, O., & Souza, R. D. (2017). A Survey of Machine Learning Techniques Applied to Self-Organizing Cellular Networks. *IEEE Communications Surveys & Tutorials*, 19(4), 2392-2431. doi:10.1109/COMST.2017.2727878
- [13] Murudkar, C. V., & Gitlin, R. D. (2019). User-Centric Approaches for Next-Generation Self-Organizing Wireless Communication Networks Using Machine Learning. *IEEE COMCAS*. doi:10.1109/COMCAS44984.2019.8958302
- [14] Shabaz, M., Farouk, A., Ahmad, S., Nazir, S., & Mehbodniya, A. (2024). Autonomic and Adaptive Cyber-Defense Algorithms for Healthcare Applications. *Healthcare Analytics*, 5, 100345. <https://doi.org/10.1016/j.health.2024.100345>
- [15] Colabianchi, S., Costantino, F., Di Gravio, G., Nonino, F., & Patriarca, R. (2021). Discussing Resilience in the Context of Cyber Physical Systems. *Computers & Industrial Engineering*, 160, 107534. <https://doi.org/10.1016/j.cie.2021.107534>
- [16] Safitra, M. F., Lubis, M., & Fakhurroja, H. (2023). Counterattacking Cyber

- Threats: A Framework for the Future of Cybersecurity. Sustainability, 15(18), 13369.
<https://doi.org/10.3390/su151813369>
- [17] Carias, J. F., Labaka, L., Sarriegi, J. M., & Hernantes, J. (2019). Defining a Cyber Resilience Investment Strategy in an Industrial Internet of Things Context. Sensors, 19(1), 138.
<https://doi.org/10.3390/s19010138>
- [18] AlMajali, A., Viswanathan, A., & Neuman, C. (2017). Resilience Evaluation of Demand Response as Spinning Reserve under Cyber-Physical Threats. Electronics, 6(1), 2.
<https://doi.org/10.3390/electronics6010002>
- [19] Dorsey, L. C., Wang, B., Grabowski, M., Merrick, J., & Harrauld, J. R. (2019). Self Healing Databases for Predictive Risk Analytics in Safety-Critical Systems. Journal of Loss Prevention in the Process Industries, 63, 104014.
<https://doi.org/10.1016/j.jlp.2019.104014>
- [20] Niu, Y., Fan, Y., & Ju, X. (2023). Critical Review on Data-Driven Approaches for Learning from Accidents: Comparative Analysis and Future Research. Safety Science, 171, 106381.
<https://doi.org/10.1016/j.ssci.2023.106381>
- [21] Mishra, A., Jha, A. V., Appasani, B. et al. (2023). Emerging Technologies and Design Aspects of Next Generation Cyber Physical System with a Smart City Application Perspective. International Journal of System Assurance Engineering and Management, 14(Suppl 3), 699-721.
<https://doi.org/10.1007/s13198-021-01523-y>
- [22] Zhang, C. et al. (2020). An Energy-Aware Cyber Physical System for Energy Big Data Analysis and Recessive Production Anomalies Detection in Discrete Manufacturing Workshops. International Journal of Production Research, 58(23), 7059-7077.
- [23] Zhou, Y., Yu, F. R., Chen, J., & Kuo, Y. (2020). Cyber-Physical-Social Systems: A State-of-the-Art Survey, Challenges and Opportunities. IEEE Communications Surveys & Tutorials, 22(1), 389-425.
<https://doi.org/10.1109/COMST.2019.2959013>
- [24] Silvestri, S., Islam, S., Amelin, D. et al. (2024). Cyber Threat Assessment and

- Management for Securing Healthcare Ecosystems Using Natural Language Processing. *International Journal of Information Security*, 23, 31-50.
<https://doi.org/10.1007/s10207-023-00769-w>
- [25] Yu, X., Hu, W., Lu, S., Sun, X., & Yuan, Z. (2019). BioBERT Based Named Entity Recognition in Electronic Medical Record. 10th International Conference on Information Technology in Medicine and Education (ITME), 49-52.
<https://doi.org/10.1109/ITME.2019.00022>
- [26] Phandi, P., Silva, A., & Lu, W. (2018). SemEval-2018 Task 8: Semantic Extraction from CybersecUrity REports Using Natural Language Processing (SecureNLP). *Proceedings of the 12th International Workshop on Semantic Evaluation*, 697-706.
<https://doi.org/10.18653/v1/S18-1113>
- [27] Islam, S., Papastergiou, S., & Mouratidis, H. (2021). A Dynamic Cyber Security Situational Awareness Framework for Healthcare ICT Infrastructures. *PCI 2021: 25th Pan-Hellenic Conference on Informatics*, 334-339.
<https://doi.org/10.1145/3503823.3503885>
- [28] Gao, C., Zhang, X., & Liu, H. (2021). Data and Knowledge-Driven Named Entity Recognition for Cyber Security. *Cybersecurity*, 4(1), 1.
<https://doi.org/10.1186/s42400-021-00072-y>
- [29] Islam, S., Papastergiou, S., & Silvestri, S. (2022). Cyber Threat Analysis Using Natural Language Processing for a Secure Healthcare System. *IEEE Symposium on Computers and Communications (ISCC)*, 1-7.
<https://doi.org/10.1109/ISCC55528.2022.9912768>
- [30] Islam, S., Papastergiou, S., & Silvestri, S. (2022). Cyber Threat Analysis Using Natural Language Processing for a Secure Healthcare System. *IEEE Symposium on Computers and Communications (ISCC)*, 1-7.
<https://doi.org/10.1109/ISCC55528.2022.9912768>
- [31] Odukale, Bowale. (2024). AI Literacy and Entrepreneurial Innovation: Rethinking Digital Competence for the Knowledge Economy. *International Journal of Multidisciplinary Evolutionary Research*. 5. 91-108. 10.54660/IJMER.2024.5.1.91-108.

[32] Odukale, Bowale. (2024). The Rise of AI Literacy: Redefining Human Communication Competence in the Algorithmic Age. International Journal of Multidisciplinary Evolutionary Research. 5. 76-90. 10.54660/IJMER.2024.5.1.76-90.

[33] Olatunji, Opeyemi & Odukale, Bowale. (2025). Leveraging Predictive Analytics and AI for SME Growth: A Data-Driven Approach to Business Optimization. World Journal of Advanced Research and Reviews. 25. 1830-1841. 10.30574/wjarr.2025.25.3.0958.