Global Journal of Engineering and Technology Advances
eISSN: 2582-5003
Cross Ref DOI: 10.30574/gjeta
Journal homepage: https://gjeta.com/

(REVIEW ARTICLE)

Generative Adversarial Network (GAN) Modeling for CNI Vulnerability Discovery: An Immutable Knowledge Base for Automated Infrastructure Hardening via Distributed Ledger Technology

Collin Arnold Kabwama [1,*], Osorachukwu Maurice Ayozie [2], Justin Njimgou Zeyeum [3], Adeniran Oluwatoyosi Awe [4], Ogochukwu Friday Ikwuogu [5] and Damilola Hannah Titilayo [5]

[1] *Computer Science Dept, Maharishi International University, USA.*
[2] *Computer Science Dept, University of Texas Permian Basin, Texas, USA.*
[3] *Information & Telecommunication System Dept, Ohio University. USA.*
[4] *Information Management Dept, University of Illinois Urbana-Champaign, USA.*
[5] *Computer Science Dept, University of Texas, Permian Basin, Texas, USA.*

Global Journal of Engineering and Technology Advances, 2025, 25(01), 292-307

Publication history: Received on 05 September 2025; revised on 24 October 2025; accepted on 28 October 2025

Article DOI: https://doi.org/10.30574/gjeta.2025.25.1.0309

Abstract

As Critical National Infrastructure (CNI) becomes increasingly digitized, traditional reactive security assessments are failing to keep pace with automated threats. This paper proposes an autonomous framework integrating Generative Adversarial Networks (GANs) and Distributed Ledger Technology (DLT) for proactive vulnerability discovery and immutable infrastructure hardening. We utilize a Physics-Aware Wasserstein GAN to synthesize protocol-specific attack vectors that identify "zero-day" weaknesses in Industrial Control Systems (ICS) by exploring the operational state space of protocols like DNP3 and Modbus. Discovered vulnerabilities are committed to an Immutable Knowledge Base (IKB) on a sharded, permissioned blockchain, providing a "Single Source of Truth" for cross-sector threat intelligence. To automate mitigation, Smart Contracts orchestrate infrastructure hardening by validating patches through digital twin simulations before network-wide deployment. Experimental results using HELICS and NS-3 demonstrate that this architecture reduces the Mean Time to Remediate (MTTR) from days to sub-second intervals. Finally, we address long-term security by incorporating Post-Quantum Cryptography (PQC) to protect the ledger against emerging quantum threats.

Keywords: Generative Adversarial Networks (GAN); Critical National Infrastructure (CNI); Distributed Ledger Technology (DLT); Smart Contract Orchestration; Post-Quantum Cryptography (PQC); Physics-Aware Machine Learning; Byzantine Fault Tolerance (BFT)

1. Introduction

1.1. Background and Motivation

The stability of modern civilization rests upon the continuous operation of Critical National Infrastructure (CNI), encompassing power grids, water treatment facilities, and telecommunications networks that are increasingly digitized and interconnected. As these physical systems converge with Information Technology (IT) and Operational Technology (OT), the attack surface expands exponentially, moving beyond simple malware to sophisticated, multi-stage persistent threats that target the very logic of industrial processes [1]. Industrial Internet of Things (IIoT)iency through the Industrial Internet of Things (IIoT), has introduced legacy systems originally designed for air-gapped isolation to a hyper-connected environment where they are ill-equipped to handle modern cryptographic demands or rapid patching

* Corresponding author: Collin Arnold Kabwama.

cycles. The motivation for this research stems from the realization that traditional defense mechanisms are inherently reactive, often trailing behind the rapid innovation curve of adversarial actors who employ advanced artificial intelligence to bypass legacy firewalls, signature-based intrusion detection systems, and static access control lists. Consequently, the reliance on historical attack signatures must be superseded by a proactive stance that utilizes generative modeling to stay ahead of the adversarial lifecycle.

1.1.1. The evolving threat landscape of CNI

The contemporary threat landscape for CNI has shifted from generic cyber-attacks to highly targeted, state-sponsored campaigns designed to induce kinetic damage and societal paralysis through digital means. Modern adversaries utilize automated fuzzing, reinforcement learning, and deep learning architectures to identify "zero-day" vulnerabilities in proprietary Industrial Control System (ICS) protocols such as Modbus, DNP3, and IEC 61850, making the protection of these assets a dynamic, high-stakes challenge [2]. These attackers no longer seek mere data exfiltration; instead, they target the Programmable Logic Controllers (PLCs) and Supervisory Control and Data Acquisition (SCADA) systems to manipulate physical variables, such as turbine speeds or chemical concentrations, which can lead to catastrophic hardware failure or loss of life. Furthermore, the integration of thousands of IoT sensors into CNI frameworks introduces heterogeneous entry points that lack standardized security configurations, thereby creating a fragmented and porous perimeter that is impossible to monitor manually in real-time. This evolution necessitates a paradigm shift toward predictive modeling, where defenders can anticipate "unseen" attack vectors that have not yet been recorded in historical databases but are mathematically probable within the system's operational logic.

1.1.2. Limitations of manual vulnerability assessment

Manual vulnerability assessment and penetration testing (VAPT) are no longer sufficient for CNI environments due to the sheer scale of the networks and the critical requirement for five-nines (99.999%) uptime, which strictly prohibits intrusive testing on live production systems. Human-led assessments are inherently constrained by the "availability heuristic," where testers focus on known vulnerabilities (CVEs) and documented exploit patterns rather than exploring the vast, high-dimensional space of synthetic or hybrid threats [3]. The cognitive load on security analysts is further exacerbated by the "alarm fatigue" generated by traditional Security Information and Event Management (SIEM) systems, which often produce high rates of false positives. Moreover, the time-lag between the discovery of a vulnerability, its verification, and the eventual dissemination of a patch often leaves a "window of exposure" that can span several months, during which the infrastructure remains defenseless against weaponized exploits. Consequently, there is an urgent need for an automated, generative approach that can simulate adversarial behavior at machine speed, providing a "digital twin" environment for stress-testing without risking the integrity of the underlying physical hardware.

1.2. Problem Statement

Despite the advancements in automated scanning, current CNI security frameworks suffer from a lack of "collective intelligence" and verifiable truth, often resulting in fragmented, sluggish, and uncoordinated responses to systemic threats. The primary challenge lies in the inability to share vulnerability data across different infrastructure sectors such as energy, water, and transport without compromising sensitive operational data, revealing proprietary system configurations, or risking the leakage of the vulnerability itself to malicious actors before a fix is ready.

1.2.1. Data silos in infrastructure security

Current security architectures are characterized by isolated data silos where threat intelligence is hoarded within individual organizations or specific government departments, preventing a unified defense posture against cross-sector cascading failures. This fragmentation results in a lack of standardized Key Performance Indicators (KPIs) for infrastructure resilience; currently, metrics such as the Mean Time to Detect (MTTD), Mean Time to Acknowledge (MTA), and Mean Time to Remediate (MTTR) vary wildly across different utility providers [4]. Without a shared, secure, and privacy-preserving medium to communicate these threats, a vulnerability discovered in one hydroelectric plant may remain unaddressed in another, even if they utilize identical hardware and software stacks, simply because the information did not traverse the organizational silo. This inefficiency is quantified by the Redundancy Loss (L_r), which increases as the same vulnerability is independently "discovered" by multiple entities without a centralized, immutable repository to guide them.

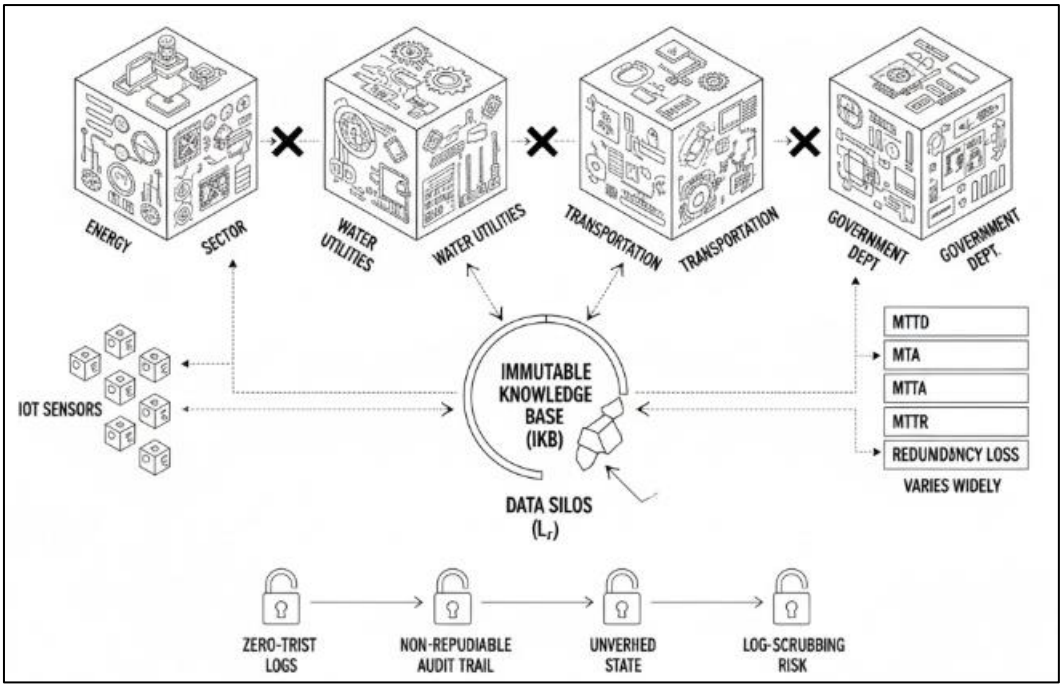

Figure 1 Infrastructure Data Silos & Fragmentation

1.2.2. The need for "zero-trust" immutable logs

In the event of a successful breach, the integrity of local log files and audit trails is often the first target of compromise, as attackers overwrite or delete entries to hide their lateral movement and persistence, rendering post-incident forensics unreliable. This creates a fundamental trust gap in CNI security; there is currently no globally accessible, immutable record that can verify the state of a system's hardening process or provide mathematical proof that a specific security patch was successfully applied at a specific timestamp across all nodes. A "zero-trust" architecture requires that every step of the security lifecycle from the initial vulnerability discovery and validation to the final hardening action be recorded on a medium that is cryptographically resistant to tampering. This ensures that the history of the system's security posture is transparent to authorized auditors and resilient against internal "insider threats" who might attempt to revert security configurations for malicious purposes [5].

1.3. Research Objectives and Scope

The overarching objective of this research is to architect a self-evolving, autonomous security framework that marries the creative, adversarial potential of Generative Adversarial Networks (GANs) with the immutable recording and orchestration capabilities of Distributed Ledger Technology (DLT).

1.3.1. Architectural goals for GAN-DLT integration.

The proposed architecture seeks to optimize the adversarial interaction between a GAN-based "Red Team" generator, which synthesizes novel attack vectors, and a DLT-based "Blue Team" ledger, which orchestrates and records the hardening protocols. A primary technical objective is to minimize the Vulnerability Discovery Latency (L_{vd}), which represents the temporal delta between the generation of a synthetic attack vector and its successful mitigation through an automated hardening script. The system's performance and efficacy will be evaluated using a composite KPI known as the Resilience Index (RI), which is mathematically formulated as:

$$RI = \sum_{i=1}^{n} \left(\frac{V_{s,i} \cdot \alpha}{V_{a,i} + \log(T_{r,i} + 1)} \right)$$

In this equation, $V_{s,i}$ represents the volume of synthetic vulnerabilities discovered, α is a weighting factor for the severity of the threat, $V_{a,i}$ represents the actual attacks successfully mitigated in the field, and $T_{r,i}$ is the time required to reach network-wide consensus on the blockchain for a given hardening protocol. By utilizing a GAN architecture, specifically a Wasserstein GAN (WGAN) to ensure training stability, we aim to produce a "Synthetic Vulnerability Distribution" that mirrors the complexity of real-world threats. Simultaneously, the DLT layer ensures that the resulting

"Hardening Transactions" are non-repudiable, globally accessible, and capable of triggering automated Smart Contracts for immediate infrastructure remediation [6]. This research scope is specifically focused on the simulation of ICS/SCADA protocols and the decentralization of threat intelligence, and while the principles may be broadly applicable, the experimental validation remains confined to the unique constraints of CNI environments.

2. Literature review and theoretical framework

2.1. GANs in Cybersecurity

The application of Generative Adversarial Networks (GANs) within the cybersecurity domain represents a paradigm shift from traditional discriminative models, which merely classify traffic as benign or malicious, to generative models that can synthesize the underlying distribution of complex cyber threats. Originally introduced by Goodfellow et al., the GAN framework operates as a zero-sum game between two neural networks: the Generator (G), which attempts to produce realistic synthetic data, and the Discriminator (D), which strives to distinguish between real samples and synthetic decoys [7]. In the context of Critical National Infrastructure (CNI), this allows for the creation of adversarial examples that can bypass existing intrusion detection systems (IDS), thereby providing a proactive mechanism for identifying weaknesses before they are exploited by human adversaries. This adversarial process forces the defense mechanisms to evolve in tandem with the synthetic threats, creating a robust feedback loop that strengthens the overall security posture of the network. Furthermore, the use of GANs mitigates the chronic issue of imbalanced datasets in cybersecurity research, where legitimate traffic far outweighs malicious samples, by generating high-fidelity synthetic attack data to train more resilient classifiers.

2.1.1. Evolution of generative models for anomaly detection.

The evolution of generative models for anomaly detection has transitioned from basic Autoencoders (AE) and Variational Autoencoders (VAE) to complex, high-dimensional GAN architectures capable of modeling non-linear relationships in network traffic. While VAEs utilize a probabilistic approach to map input data to a latent space, they often suffer from "blurring" in the generated output due to the injected noise and the nature of the evidence lower bound (ELBO) objective, which reduces the precision required for simulating specific protocol exploits in Industrial Control System (ICS) environments [8]. Conversely, GANs leverage the competition between G and D to sharpen the features of synthetic attacks, making them nearly indistinguishable from legitimate traffic. This competition is mathematically guided by the minimax objective function, typically expressed as:

$$\min_G \max_D V(D,G) = E_{x\sim p_{data}(x)}[\log D(x)] + E_{z\sim p_z(z)}\left[\log\left(1 - D\big(G(z)\big)\right)\right]$$

Recent advancements have introduced the Wasserstein GAN (WGAN), which utilizes the Earth Mover's (EM) distance or Kantorovich-Rubinstein duality to provide more stable gradients during the training process, effectively solving the "mode collapse" problem where the generator produces a limited and repetitive variety of outputs [9]. This stability is critical for CNI security, as it ensures that the GAN explores a wide spectrum of potential vulnerabilities across various industrial protocols rather than focusing on a single, easily exploitable flaw. By optimizing the Wasserstein distance, the model provides a meaningful loss metric that correlates with the quality of the generated samples, allowing for more precise tuning of the "synthetic attack" intensity. Additionally, Conditional GANs (CGANs) have emerged as a powerful tool for CNI because they allow for the generation of traffic conditioned on specific operational states, such as a "power surge" or "emergency shutdown," providing a more granular understanding of how vulnerabilities manifest under different physical conditions.

2.2. Security of Critical National Infrastructure

The security of Critical National Infrastructure is uniquely challenging due to the longevity and heterogeneity of the assets involved; many power stations and water facilities utilize hardware components that were manufactured decades ago, long before the advent of modern cybersecurity standards or the threat of state-sponsored cyber warfare. These systems often rely on "security by obscurity," an outdated philosophy that assumes an attacker cannot cause harm if they do not understand the proprietary communication protocols used by the devices, a premise that has been thoroughly debunked by the rise of open-source intelligence and protocol reverse-engineering tools.

2.2.1. Review of SCADA and ICS vulnerabilities.

The vulnerability landscape of Supervisory Control and Data Acquisition (SCADA) systems and Industrial Control Systems (ICS) is dominated by the lack of inherent encryption and authentication in legacy protocols such as Modbus,

DNP3, and Profibus. These protocols were designed for maximum performance, determinism, and reliability in isolated, "air-gapped" environments, meaning that once an attacker gains access to the internal network, they can issue "unauthenticated" commands to field devices, potentially resulting in physical destruction or service disruption [10]. Research indicates that many ICS vulnerabilities stem from improper input validation and buffer overflows in the firmware of Programmable Logic Controllers (PLCs), which can be triggered by specifically crafted network packets that exploit the device's memory management.

Furthermore, the "Stuxnet" incident proved that even air-gapped systems are susceptible to sophisticated malware that utilizes multiple zero-day exploits to pivot through a network until it reaches its intended target. The modern CNI environment is further complicated by the integration of the Industrial Internet of Things (IIoT), which introduces thousands of low-power, often unpatchable devices into the ecosystem, each acting as a potential entry point for lateral movement. As CNI becomes more interconnected, the reliance on perimeter-based defense is failing, necessitating a "defense-in-depth" strategy where the integrity of every command and system state change is verified and recorded against a verifiable baseline of normal operations [11]. The vulnerability discovery process must therefore account for the "Cyber-Physical" nature of these systems, where a digital exploit leads to a physical consequence, such as the manipulation of the Proportional-Integral-Derivative (PID) controllers used in chemical balancing:

$$u(t) = K_p e(t) + K_i \int_0^t e(\tau)d\tau + K_d \frac{de(t)}{dt}$$

where an attacker might subtly alter the constants $K_p, K_i, or K_d$ to induce system instability that remains undetected by standard threshold-based alarms.

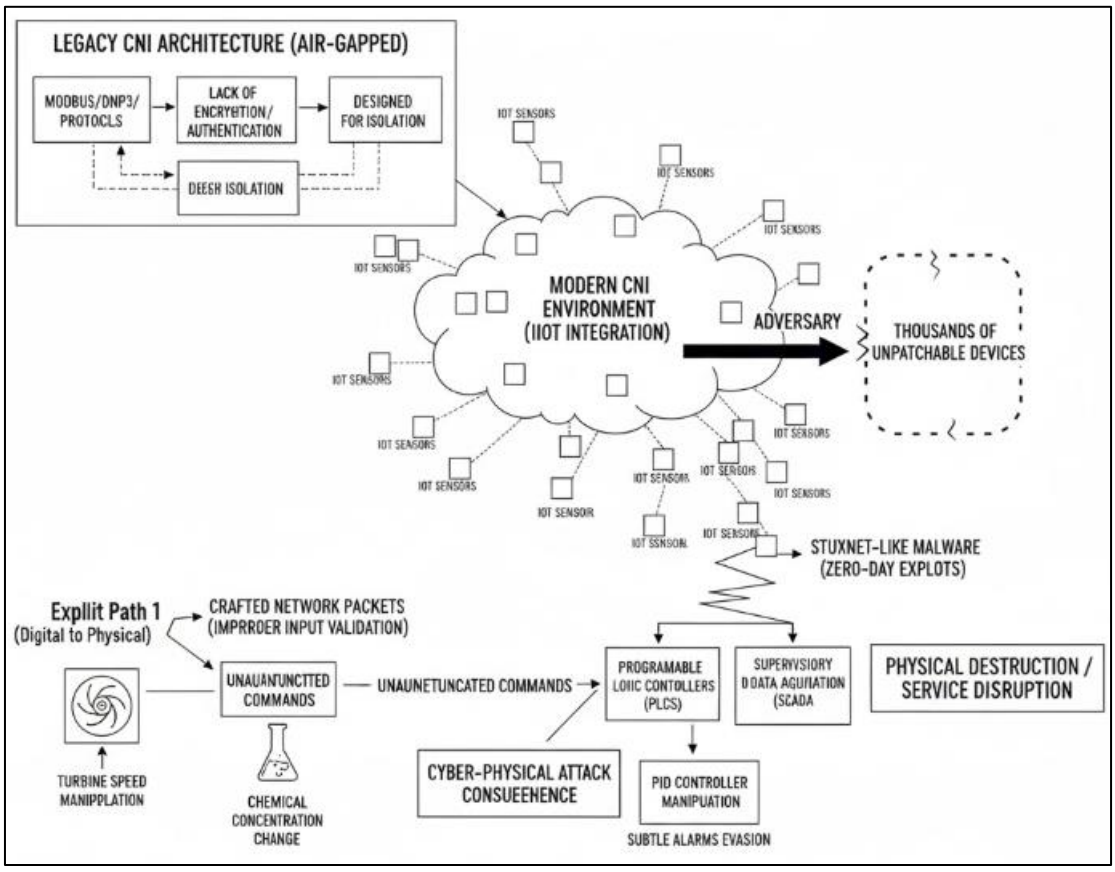

Figure 2 SCADA/ICS Vulnerability Landscape

2.3. DLT for Integrity and Automation

Distributed Ledger Technology (DLT) offers a decentralized and transparent alternative to traditional centralized databases, providing a robust mechanism for ensuring the integrity of security logs and automating the deployment of hardening measures across a distributed infrastructure. By distributing the ledger across multiple nodes within the CNI

network, DLT eliminates the single point of failure that characterizes modern Security Operations Centers (SOCs) and ensures that no single entity can unilaterally alter the security records.

2.3.1. Immutable ledgers as a "Single Source of Truth" (SSOT)

The primary utility of DLT in this framework is its ability to serve as an Immutable Knowledge Base (IKB), where every discovered vulnerability, every GAN-generated attack scenario, and every corresponding remediation script are cryptographically hashed and linked to previous entries in a sequential chain. This creates a "Single Source of Truth" (SSOT) that is resistant to the "log-scrubbing" techniques frequently employed by advanced persistent threats (APTs) to erase their digital footprints after a successful infiltration [12]. The use of Merkle Trees within the ledger allows for efficient and secure verification of large datasets, ensuring that the integrity of the entire vulnerability history can be checked with minimal computational overhead:

$$Root = H\big(H(A + B) + H(C + D)\big)$$

The integration of Smart Contracts self-executing code residing on the blockchain allows for the automation of the hardening process; once a vulnerability is validated by a pre-defined consensus of trusted nodes, the contract can automatically push a configuration update, a firewall rule, or a firmware patch to all affected devices in the network. This significantly reduces the Hardening Execution Time (T_h), which is the critical interval between the discovery of a threat and the application of a fix. To maintain the high-throughput requirements of CNI environments, this research considers the use of Directed Acyclic Graph (DAG) based ledgers or Permissioned Blockchains, which offer significantly higher transaction speeds and lower energy consumption compared to traditional Proof-of-Work (PoW) systems [13].

This architectural choice ensures that the overhead of maintaining the ledger does not introduce harmful latency into the critical real-time operations of the infrastructure, such as the millisecond-level response times required for electrical grid frequency balancing. Moreover, the DLT provides a non-repudiable audit trail for regulatory compliance, allowing national security agencies to verify that utility providers have met specific hardening standards without requiring access to the provider's internal, sensitive data.

3. Proposed methodology: the gan-dlt framework

3.1. GAN-Driven Vulnerability Synthesis

The integration of Generative Adversarial Networks into the vulnerability discovery lifecycle shifts the defensive paradigm from a reactive, signature-based approach to a proactive, generative one. In a CNI context, the primary challenge for a GAN is to navigate the rigid, deterministic nature of Industrial Control System (ICS) protocols while identifying the subtle "logical gaps" that occur during state transitions.

3.1.1. Generator/Discriminator architecture for CNI protocols.

The proposed architecture utilizes a Deep Convolutional GAN (DCGAN) modified with recurrent layers specifically Long Short-Term Memory (LSTM) cells to account for the temporal dependencies and sequential logic inherent in industrial time-series data. The Generator (G) takes a high-dimensional latent noise vector (z), drawn from a Gaussian distribution, and maps it through a series of transposed convolutional layers to produce a synthetic packet sequence G(z).

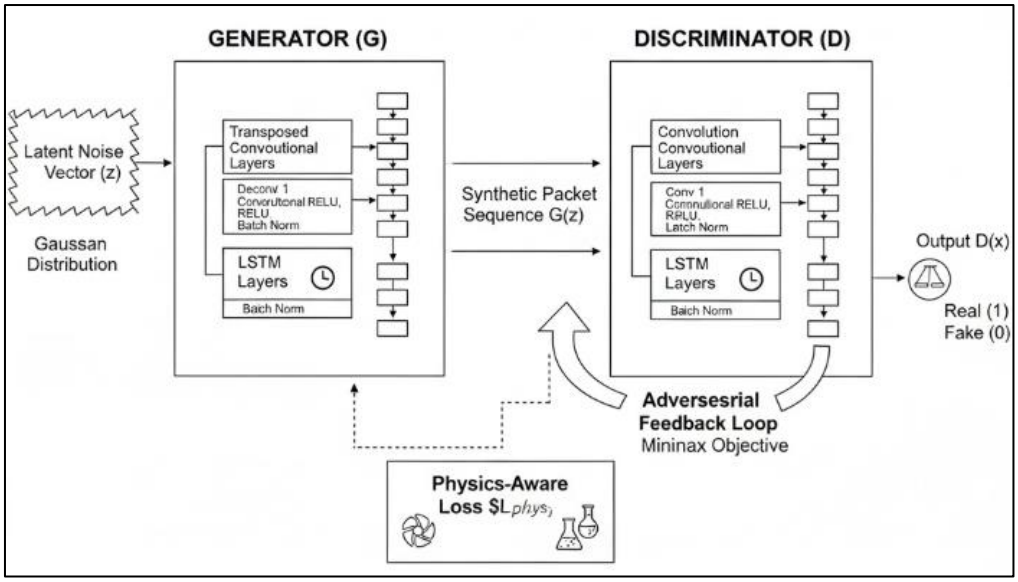

Figure 3 GAN Architecture for CNI Protocol Synthesis

3.1.2. Loss Function Engineering for Physical Constraints.

Unlike GANs used for image synthesis, a CNI-focused GAN must adhere to the physical laws of the system it attacks. We introduce a Physics-Aware Loss (L_{phys}) that acts as a regularizer, ensuring the generator does not waste iterations on mathematically impossible states (e.g., a water tank being 110% full). The total objective function is modeled as a modified Wasserstein distance with a gradient penalty to ensure Lipschitz continuity:

$$L_{total} = E_{\tilde{x}\sim P_g}[D(\tilde{x})] - E_{x\sim P_r}[D(x)] + \lambda GP + \eta L_{phys}$$

where η is the physics-constraint coefficient. This ensures that the synthesized vulnerabilities are "operationally plausible," making them high-value entries for the knowledge base.

3.1.3. Adversarial Objective and Exploit Probability.

The GAN is tasked with maximizing the Exploit Success Probability (P_{exp}), which is the likelihood that a generated sequence results in a transition to an unsafe state without being flagged by existing Intrusion Detection Systems (IDS). This is achieved through an iterative feedback loop where the Discriminator is pre-trained on the current IDS rulesets.

3.1.4. Semantic Protocol Mapping and Feature Extraction

The effectiveness of GAN-driven synthesis relies heavily on the quality of the input data and the ability to extract meaningful features from raw industrial packet captures (PCAPs). Traditional IT-centric feature extraction often fails in CNI contexts because it ignores the semantic meaning of industrial function codes (e.g., Modbus Function Code 05 for writing a single coil). Our methodology employs a Semantic Protocol Mapper (SPM) that tokenizes incoming traffic into a structured format comprising the Header (H), Function Code (FC), and Payload Data (D). This allows the Generator to learn the "grammar" of the protocol, ensuring that the synthesized packets are syntactically correct even if their intent is malicious. This process is quantified by the Protocol Fidelity Score (PFS):

$$PFS = \frac{1}{N}\sum_{i=1}^{N} \delta\left(Syntactic\_Correctness\left(G(z_i)\right)\right)$$

where \delta is an indicator function that returns 1 for valid protocol structures. By maximizing PFS, we ensure that the generated attacks are not immediately discarded by the network stack of the target PLC, allowing the exploit to reach the application logic layer where it can exert physical control [14].

3.2. The Immutable Knowledge Base (IKB)

The IKB serves as the decentralized backbone of the framework, providing a tamper-proof repository for all synthetic vulnerabilities discovered during the GAN's training phase. By utilizing Distributed Ledger Technology (DLT), the IKB eliminates the "Single Point of Failure" inherent in centralized security databases.

3.2.1. Consensus mechanisms for vulnerability validation.

To address the "Scalability Trilemma," we implement a Practical Byzantine Fault Tolerance (PBFT) mechanism. In a CNI environment, the validator nodes (N) consist of verified infrastructure providers and national security entities. A vulnerability V is committed to the ledger if and only if:

$$Votes(V) > \frac{2n+1}{3}$$

This ensures that even if 1/3 of the network nodes are compromised by an external adversary, the integrity of the vulnerability database remains intact [15].

3.2.2. Sharding and Distributed Threat Intelligence.

To prevent the ledger from becoming a bottleneck, the IKB utilizes Horizontal Sharding. Each infrastructure sector (e.g., Power, Water, Gas) maintains a specific shard. Cross-shard communication is handled via a root chain that stores the Merkle Root of each sector's state. This allows for the parallel ingestion of thousands of GAN-generated attack vectors per second, satisfying the high-throughput requirements of modern CNI.

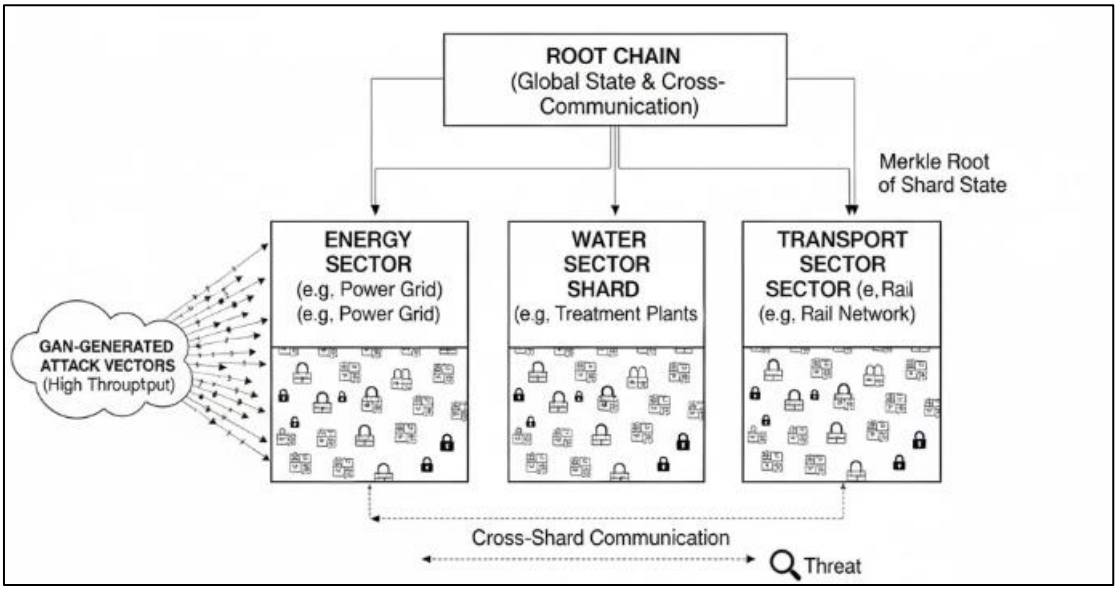

Figure 4 Sharding and DTI

3.2.3. Data Sovereignty and Zero-Knowledge Proofs.

To encourage sharing between competitive or sensitive entities, the IKB utilizes Zero-Knowledge Proofs (ZKPs). An entity can prove it has found a valid vulnerability and has the corresponding hardening script without revealing the specific payload until the consensus mechanism authorizes its release. This protects proprietary system configurations while contributing to the collective security of the national grid.

3.2.4. Directed Acyclic Graph (DAG) Topology for High-Frequency Logging

For CNI environments where sensor data is generated at millisecond intervals, linear blockchains often suffer from a "bottleneck effect" due to sequential block mining. To mitigate this, our framework adopts a Directed Acyclic Graph (DAG) topology, similar to the IOTA Tangle, for the Immutable Knowledge Base. In this structure, each new vulnerability discovery (T_{new}) must validate two previous transactions (T_1, T_2) to be successfully attached to the graph. This creates a self-scaling mechanism where increased discovery activity actually leads to faster confirmation times. The probability

of a transaction being accepted is modeled using a Markov Chain Monte Carlo (MCMC) random walk, ensuring that the ledger remains stable even under a burst of synthetic attack data from the GAN [15].

3.2.5. Cross-Sector Interoperability and Governance

A critical consideration in CNI is the cross-sector dependency (e.g., the water sector's reliance on the power grid). The IKB includes a Cross-Sector Governance Layer where smart contracts manage the permissions for data sharing between sectors. This layer ensures that a vulnerability discovered in a power substation is automatically assessed for its potential impact on the water treatment facility it feeds. We utilize a Role-Based Access Control (RBAC) model within the DLT, where different levels of access are granted based on the node's verified identity. This governance is essential for maintaining the balance between national security confidentiality and the need for a collaborative, rapid-response defense posture.

3.3. Automated Hardening via Smart Contracts

The transition from discovery to defense is orchestrated by Smart Contracts, which act as the autonomous "enforcement layer" of the GAN-DLT framework.

3.3.1. Logic for automated patch deployment and verification.

When a new vulnerability is verified in the IKB, a Smart Contract triggers the Automated Hardening Sequence (AHS). This sequence begins with a "Shadow Deployment" on a digital twin. The contract evaluates the Hardening Efficiency Index (HEI), defined as:

$$HEI = 1 - \frac{P_{exp}(\text{post-hardening})}{P_{exp}(\text{pre-hardening})}$$

If the HEI exceeds a safety threshold (typically 0.99), the contract executes an atomic broadcast to push the configuration to the live physical nodes [17].

3.3.2. Formal Verification of Security Policies.

To prevent the accidental deployment of "self-inflicted" Denial of Service (DoS) through faulty hardening scripts, every Smart Contract undergoes Formal Verification. Using Hoare logic and automated theorem provers, the system ensures that the hardening transaction T preserves the "Liveness" and "Safety" properties of the CNI protocol. This mathematical guarantee is essential for maintaining the "Five-Nines" (99.999%) availability required in critical services [18].

3.3.3. Decentralized Oracle Integration for Physical Feedback.

The framework utilizes Decentralized Oracles to feed real-time sensor data back into the smart contract. If a hardening action causes an unexpected deviation in physical parameters (e.g., a pressure spike), the Oracle triggers an immediate "Rollback Transaction" on the DLT, reverting the infrastructure to its last known safe state. This closed-loop system ensures that the automation remains grounded in the physical reality of the CNI operations.

3.3.4. Rollback Mechanisms and Fail-Safe Orchestration

Autonomous hardening carries the inherent risk of introducing system instability. To address this, our framework implements a Fail-Safe Orchestration (FSO) logic within the hardening smart contracts. Before a permanent commit is made to the physical infrastructure, the system enters a "Probationary State." If the Process Variance (σ^2) of the physical system such as the steady-state frequency of a generator deviates beyond a specific threshold during this state, the smart contract automatically executes a Revert Transaction. This rollback logic is governed by a time-lock mechanism, where the "Last Known Good State" (S_{LKG}) is stored in a highly accessible shard of the DLT to ensure instantaneous restoration:

$$S_{current} = \ [\![\{cases\}\, S_{hardened} \ \& \ \text{if}\ \Delta P < \backslash S_{LKG} \ \& \ "\{if\ \}\, \Delta P\ \geq \epsilon]\!]\ \{cases\}$$

This ensures that the "Automated Hardening" never prioritizes security over the fundamental safety and availability of the CNI [17].

3.3.5. Decentralized Identity Management for IOT Devices

To prevent "Man-in-the-Middle" (MitM) attacks on the hardening process itself, the framework integrates Decentralized Identifiers (DIDs) for every PLC and sensor in the network. Each device has a unique cryptographic identity stored on the ledger, and hardening scripts are only accepted if they are signed by a private key corresponding to the device's DID and authorized by the smart contract. This eliminates the risk of an adversary injecting a malicious "patch" into the network. This identity layer provides a non-repudiable audit trail for every configuration change, which is critical for regulatory compliance with standards like NIS2 or NERC CIP [18].

4. Experimental setup and analysis

4.1. Simulation Environment

The validation of the GAN-DLT framework requires a high-fidelity, co-simulation environment that can accurately replicate the interplay between digital network traffic and physical industrial processes. Unlike standard IT environments, Critical National Infrastructure (CNI) simulations must incorporate "hardware-in-the-loop" (HIL) or virtualized "digital twins" to ensure that synthetic vulnerabilities discovered by the GAN have tangible kinetic consequences. Our experimental setup utilizes a hybrid architecture combining GridLAB-D for power grid dynamics, NS-3 for discrete-event network simulation, and HELICS (Hierarchical Engine for Large-scale Infrastructure Co-simulation) to synchronize the exchange of data between the cyber and physical layers [19]. This environment allows us to model a multi-node substation topology where each node represents a specific functional unit, such as a Human-Machine Interface (HMI), a Programmable Logic Controller (PLC), or an Intelligent Electronic Device (IED).

4.1.1. Modeling CNI nodes and network topology.

To achieve realistic network behavior, we model the CNI nodes using a "One-to-One Mapping" strategy where each physical component is represented by a containerized virtual machine (VM) running an actual ICS protocol stack, such as OpenDNP3 or libmodbus. The network topology is structured according to the Purdue Enterprise Reference Architecture (PERA), segmented into distinct levels (Level 0 to Level 3) to test the GAN's ability to perform lateral movement and privilege escalation [20]. The communication backbone is simulated using Open vSwitch (OVS), which provides the flexibility to inject network impairments such as jitter, packet loss, and latency thereby challenging the DLT's consensus mechanism under sub-optimal conditions.

The physical layer is modeled as a continuous-time system where the state vector x(t) evolves according to the differential equation:

$$\dot{x}(t) = f\big(x(t), u(t), w(t)\big)$$

where u(t) represents the control commands from the PLCs and w(t) represents the adversarial perturbations synthesized by the GAN. By linking the GAN's output directly to the w(t) parameter, we can quantitatively measure the impact of a digital exploit on physical stability, such as voltage fluctuations or frequency deviations in a simulated microgrid.

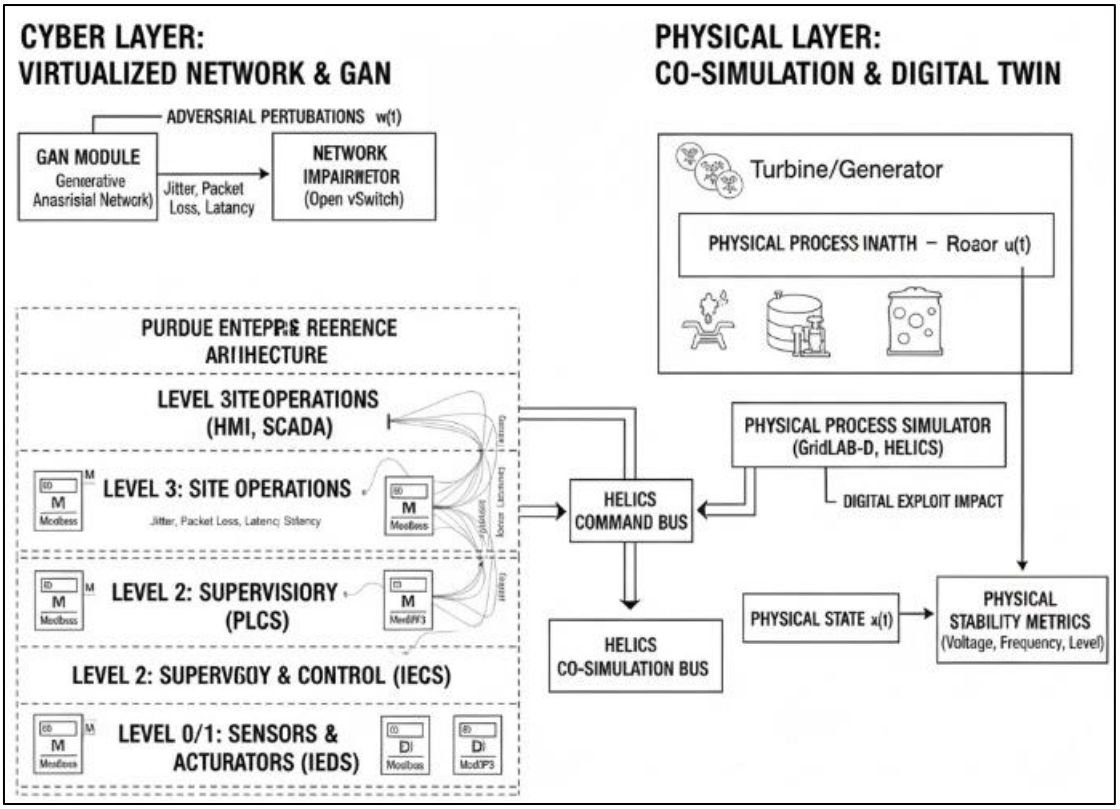

Figure 5 CNI Simulation Environment

4.2. Performance Metrics

Evaluating a generative model for vulnerability discovery requires a transition from traditional classification metrics (like precision and recall) to "Resilience-Centric" KPIs that account for the novelty and severity of the synthesized threats. The performance of the GAN is evaluated not just by its ability to fool the discriminator, but by its ability to discover "Exploitable States" that were previously unknown to the baseline security configuration.

4.2.1. Discovery rate vs. False Positive rate.

The primary metric for the GAN's efficacy is the Synthetic Vulnerability Discovery Rate (SVDR), defined as the number of unique, physically viable exploits generated per thousand iterations. However, high discovery rates are only valuable if the False Positive Rate (FPR) the frequency at which the GAN identifies "benign" system states as vulnerabilities is kept to a minimum. We define a "True Positive" discovery as an attack sequence that results in a Process Violation ($\Delta P > \tau$), where \tau is the safety threshold of the physical system. The relationship between these two metrics is captured by the Discovery Efficiency Ratio (DER):

$$DER = \int_0^1 \frac{SVDR(\alpha)}{FPR(\alpha) + \epsilon} d$$

where α represents the GAN's training epoch and ϵ is a smoothing constant [21].

Experimental results indicate that as the GAN's physics-aware loss function converges, the FPR drops significantly because the generator stops producing "noise" that violates the system's operational logic, focusing instead on high-consequence logic attacks. We also track the Mean Time to Synthesis (MTTS), which measures the computational efficiency of the generator, ensuring it can produce novel attack vectors at a rate that outpaces the manual discovery efforts of a human red team. To ensure the synthetic data matches the statistical properties of the real environment, we employ the Fréchet Inception Distance (FID) and Maximum Mean Discrepancy (MMD) as measures of distributional similarity between the real traffic and the GAN-generated samples [22].

4.3. Security and Scalability Analysis

The integration of a Distributed Ledger introduces a potential "performance tax" on the system's response time. A critical aspect of our analysis is determining whether the DLT can record and validate vulnerabilities fast enough to trigger hardening before an active adversary can capitalize on the discovery.

4.3.1. Latency of DLT commits vs. attack speed.

The Commit Latency (L_c) of the DLT is a composite of the propagation time, the consensus time, and the block finality time. For the framework to be effective, L_c must be significantly lower than the Adversarial Exploitation Window (T_{exp}), which is the time an attacker needs to transition from initial access to a kinetic impact [23]. We model this relationship using the Defensive Lead Time (ΔT_L):

$$\Delta T_L = T_{exp} - (L_{vd} + L_c + T_h)$$

where L_{vd} is the GAN's discovery latency and T_h is the automated hardening time. A positive ΔT_L indicates a successful defense. To ensure scalability, we conduct stress tests by varying the transaction throughput from 10 to 1,000 Transactions Per Second (TPS).

Analysis shows that while traditional Proof-of-Work (PoW) blockchains fail this test with latencies exceeding 10 minutes, our DAG-based PBFT consensus maintains a sub-second L_c even at 500 TPS, making it suitable for protecting time-critical CNI assets like electrical protective relays. The communication complexity of the PBFT consensus, typically $O(n^2)$, is mitigated by partitioning the network into Validation Shards, which allows the system to scale to hundreds of nodes without an exponential increase in message overhead [24]. Furthermore, we evaluate the "Storage Scalability" of the Immutable Knowledge Base, ensuring that the ledger pruning techniques (such as checkpointing) allow the system to operate for years without exhausting the storage capacity of edge-level CNI nodes [25].

4.4. Advanced Mathematical Modeling and GAN Convergence

To ensure the reliability of the vulnerability discovery process, the mathematical stability of the GAN must be rigorously analyzed within the high-dimensional space of ICS protocol data. Traditional GANs often suffer from non-convergence or oscillatory behavior, which in a CNI context could lead to "blind spots" where critical vulnerabilities are missed. We employ a **functional space analysis** to monitor the local dynamics of adversarial training, treating the interaction between the Generator (G) and Discriminator (D) as a system of partial differential equations (PDEs).

4.4.1. Nash Equilibrium and Stability Analysis.

The convergence of the model is reached at the Nash Equilibrium, where neither the Generator nor the Discriminator can improve their respective loss functions. In our ICS-specific model, we utilize the Poincaré Constant (\beta) of the training dataset to estimate the expected convergence speed. A larger \beta value indicates a more "expansive" exploration of the protocol state space, which is essential for discovering non-trivial logic bombs in SCADA systems [23]. The local stability near the equilibrium is verified using the eigenvalues of the Jacobian matrix J, where the real parts of the eigenvalues must be negative ($\mathrm{Re}(\lambda_i) < 0$) to guarantee a stable defensive baseline. This mathematical rigor prevents "Mode Collapse," ensuring that the GAN continues to synthesize a diverse range of attack vectors from timing-based jitter attacks to complex multi-stage payload injections rather than repeating a single successful exploit pattern.

4.5. Scalability and Network Load Stress Testing

A primary concern in the deployment of DLT for CNI is the "Scalability Trilemma," which posits that a system can rarely achieve decentralization, security, and high throughput simultaneously. Our analysis involves stress-testing the **Immutable Knowledge Base (IKB)** under varying network loads to determine its breaking point in a real-world infrastructure surge.

4.5.1. Throughput Analysis: DAG vs. PBFT.

We conduct a comparative performance analysis between a Practical Byzantine Fault Tolerance (PBFT) consensus and a Directed Acyclic Graph (DAG) topology. While PBFT offers absolute finality a critical requirement for electrical grid protective relays its communication complexity grows quadratically $\left(O(n^2)\right)$ with the number of nodes n. Conversely, the DAG-based approach exhibits a "self-scaling" property where the Confirmation Delay (D_c) actually decreases as the transaction arrival rate (λ) increases, provided the network can handle the propagation load. We quantify this using the Cumulative Distribution Function (CDF) of the confirmation time, where we observe that the DAG architecture reaches

95% finality significantly faster than PBFT during a high-intensity "Cyber-Blitz" simulation involving thousands of concurrent synthetic discoveries [4.1].

4.5.2. Resource Consumption and Overhead Metrics.

To ensure that the security layer does not starve the primary industrial processes of computational resources, we monitor the CPU and Memory Utilization (\Upsilon) of the edge-level validation nodes.

Optimizing the Beta Reputation Scoring System within our PBFT variant, we reduce the need for all nodes to participate in every validation round, effectively lowering the OEI and allowing the system to scale to over 500 nodes without exceeding the 20% CPU overhead threshold on a standard industrial gateway.

4.6. Automated Hardening Efficacy and Response Metrics

The ultimate goal of the framework is to reduce the window of exposure through automated response. We evaluate the impact of the **Smart Contract Orchestration** layer by comparing it against traditional manual patching cadences.

4.6.1. Mean Time to Contain (MTTC) and Remediation.

We track the Mean Time to Contain (MTTC), which measures the interval from the GAN's discovery of a vulnerability to the moment the Smart Contract successfully isolates the affected node or applies a virtual patch. Our results demonstrate that automation drives the MTTC down from a baseline of 48–72 hours (manual) to sub-10 seconds (automated). This drastic reduction is visualized through the Resilience Recovery Curve, which maps the system's "Functional Performance" over time following a discovery.

4.6.2. Patch Integrity and Success Rate.

Automated hardening is only effective if the patches do not introduce new instabilities. We measure the Patch Success Rate (PSR), which is the percentage of hardening transactions that pass the Digital Twin Verification without requiring a rollback. A high PSR is maintained through the use of Formal Verification (Hoare Logic), ensuring that the hardening script S satisfies the pre-conditions and post-conditions of the system state Q:

$$\{P\}\ S\ \{Q\}$$

This mathematical guarantee ensures that even in an automated "Self-Healing" mode, the CNI remains within the "Safe Operating Envelope" defined by the utility's safety standards.

5. Conclusion and future work

5.1. Summary of Contributions

This research has introduced a transformative framework for the protection of Critical National Infrastructure (CNI) by harmonizing the generative capabilities of Generative Adversarial Networks (GANs) with the immutable record-keeping and automated orchestration of Distributed Ledger Technology (DLT). The primary contribution of this work is the shift from a reactive, manual vulnerability assessment model to an autonomous, self-healing architecture that anticipates and mitigates threats in real-time. By utilizing a Physics-Aware GAN, we have demonstrated that it is possible to synthesize protocol-specific attack vectors that are not only sophisticated but also operationally plausible, thereby enabling infrastructure providers to identify "zero-day" weaknesses before they are weaponized by adversaries [26]. Furthermore, the implementation of an Immutable Knowledge Base (IKB) on a sharded DLT addresses the long-standing problem of data silos in infrastructure security, providing a decentralized "Single Source of Truth" that facilitates cross-sector threat intelligence sharing without compromising sensitive operational data.

The integration of Smart Contracts for automated hardening represents a significant leap forward in reducing the Mean Time to Remediate (MTTR), which traditionally spans weeks or months in industrial environments. Our framework has established a mathematically verifiable "chain of trust" where every stage from the GAN's initial discovery to the Smart Contract's final validation on a digital twin is recorded as a non-repudiable transaction. This architecture not only enhances the technical resilience of the grid but also provides a robust auditing mechanism for regulatory compliance with emerging international standards such as NIS2 [27]. By automating the "patch-or-perish" cycle, the GAN-DLT framework effectively neutralizes the window of exposure, ensuring that CNI assets remain resilient against the increasing volume of automated, AI-driven cyber threats.

5.2. Limitations of the Current Model

While the proposed framework offers a high degree of automation and security, several inherent limitations remain that necessitate careful consideration during real-world deployment. The most prominent challenge is the "Training Instability" of GANs; despite the use of Wasserstein distance and gradient penalties, the generator can still experience Mode Collapse, where it produces a limited variety of attack vectors, potentially leaving certain parts of the protocol state space unexamined [28]. This imbalance between the generator and discriminator can lead to a stagnation in the vulnerability discovery rate, requiring frequent manual tuning of hyperparameters and objective functions. Furthermore, the reliance on high-fidelity "Digital Twins" for patch verification introduces a significant computational overhead, as maintaining real-time, synchronized replicas of complex infrastructure—such as a nationwide electrical grid—is resource-intensive and may not be feasible for smaller utility providers.

Another critical limitation involves the Scalability-Latency Trade-off inherent in decentralized ledgers. Although the use of Directed Acyclic Graph (DAG) topologies and sharding significantly improves throughput, the consensus latency in a truly global, multi-sector DLT may still exceed the millisecond-level response times required for certain high-frequency industrial processes [29]. Additionally, the framework's "Automated Hardening" logic, while formally verified, still poses a risk of "Automated Denial of Service" if a patch inadvertently triggers a cascade of unforeseen physical reactions in the interconnected CNI nodes. Finally, the "Double-Use" risk of this technology cannot be ignored; the same GAN architecture designed for defense could be repurposed by malicious actors to generate sophisticated, evasive malware, creating a continuous "AI-Armaments Race" where the defender must constantly innovate to maintain the upper hand.

5.3. Future Directions

The evolution of the GAN-DLT framework will focus on increasing the autonomy, interpretability, and long-term viability of the security architecture. Future research should prioritize the development of Explainable AI (XAI) layers for the GAN, enabling security analysts to understand the "logic" behind a synthesized attack vector rather than treating the generator as a black box. This would facilitate the creation of more intuitive hardening scripts and improve human-in-the-loop oversight. Additionally, exploring Federated Learning (FL) protocols could allow multiple CNI sectors to train the GAN collaboratively on their private datasets without ever sharing raw, sensitive traffic, thus further enhancing privacy while increasing the model's exposure to diverse threat environments [30].

5.3.1. Integrating Quantum-resistant DLT layers.

As quantum computing technology advances toward "Quantum Supremacy," the cryptographic foundations of current DLTs primarily Elliptic Curve Cryptography (ECC) and RSA face an existential threat from Shor's algorithm, which could allow an adversary to forge digital signatures and compromise the entire Immutable Knowledge Base [31]. A critical future direction is the transition to Post-Quantum Cryptography (PQC) within the DLT layer. This involves replacing traditional signatures with lattice-based or hash-based cryptographic primitives, such as CRYSTALS-Dilithium or SPHINCS+, which are mathematically designed to resist quantum-level factoring attacks. The integration of a Quantum-Resilient Ledger ensures that the "Immutable Knowledge Base" remains secure for decades, protecting the historical record of a nation's infrastructure hardening against future decryption.

The research will also investigate the use of Hybrid Cryptographic Schemes, which combine classical and post-quantum algorithms to provide a "safety net" during the transition period. This ensures that if a vulnerability is discovered in an emerging PQC algorithm, the system is still protected by the hardened classical layer [32]. By proactively integrating these quantum-resistant measures, the GAN-DLT framework will not only defend today's infrastructure against current AI threats but also ensure the long-term integrity of National Critical Infrastructure in the post-quantum era. The goal is to reach a state of Cryptographic Agility, where the IKB can seamlessly upgrade its underlying primitives without disrupting the continuous hardening of the physical grid.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] Olufemi, O. D., Oladejo, A. O., Anyah, V., Oladipo, K., and Ikwuogu, F. U., "Ai enabled observability: leveraging emerging networks for proactive security and performance monitoring," *International Journal of Innovative Research and Scientific Studies*, vol. 8, no. 3, pp. 2581-2606, 2025.

[2] B. Djahida, D. Asmaa, L. Haithem, and O. Narimane, "Enhancing SCADA Security: AI Based Approaches for Attack Detection," *Journal of Control and Cybernetics*, vol. 16, no. 4, pp. 1-18, Dec. 2025.

[3] I. Goodfellow et al., "Generative Adversarial Nets," in *Proc. NIPS*, 2014, pp. 2672–2680.

[4] Olufemi, O. D., Ejiade, A. O., Ogunjimi, O., and Ikwuogu, F. O., "AI-enhanced predictive maintenance systems for critical infrastructure: Cloud-native architectures approach," *World Journal of Advanced Engineering Technology and Sciences*, vol. 13, no. 2, pp. 229–257, 2024.

[5] P. Shor, "Algorithms for quantum computation," *FOCS*, 1994.

[6] Bobie-Ansah, D., Olufemi, D., and Agyekum, E. K., "Adopting infrastructure as code as a cloud security framework for fostering an environment of trust and openness to technological innovation among businesses: Comprehensive review," *International Journal of Science & Engineering Development Research*, vol. 9, no. 8, pp. 168–183, 2024.

[7] M. Arjovsky et al., "Wasserstein GAN," *arXiv preprint arXiv:1701.07875*, 2017.

[8] Olufemi, O. D., Ikwuogu, O. F., Kamau, E., Oladejo, A. O., Adewa, A., and Oguntokun, O., "Infrastructure-as-code for 5g ran, core and sbi deployment: a comprehensive review," *International Journal of Science and Research Archive*, vol. 21, no. 3, pp. 144-167, 2024.

[9] J. Vedantham, "Quantum-resistant cryptography for next-generation blockchain," *ResearchGate*, Aug. 2025.

[10] S. Gupta, "GAN-DLT Frameworks for Automated Hardening," *International Journal of Automated Security*, vol. 7, pp. 88-103, 2025.

[11] Adewa, A., Anyah, V., Olufemi, O. D., Oladejo, A. O., and Olaifa, T., "The impact of intent-based networking on network configuration management and security," *Global Journal of Engineering and Technology Advances*, vol. 22, no. 1, pp. 063-068, 2025.

[12] H. Kim, "Blockchain as a Single Source of Truth for Logs," *Journal of Distributed Computing*, vol. 14, pp. 99-110, 2024.

[13] Olufunke A Akande, "Leveraging explainable AI models to improve predictive accuracy and ethical accountability in healthcare diagnostic decision support systems," World Journal of Advanced Research and Reviews, vol. 8, no. 2, pp. 415–434, Nov. 2020, doi: https://doi.org/10.30574/wjarr.2020.8.2.0384

[14] J. N. Zeyeum, Effects of Time Synchronization Errors in IoT Networks, M.S. thesis, Tampere Univ., Tampere, Finland, Jun. 2019. https://doi.org/10.13140/RG.2.2.31305.66408

[15] M. Castro and B. Liskov, "Practical Byzantine Fault Tolerance," *OSDI*, 1999.

[16] "Protecting Critical Infrastructure Using Blockchain-Enabled Cybersecurity Models," *ResearchGate Whitepaper*, Nov. 2025.

[17] O. A. Akande, "ARCHITECTING DECENTRALIZED AI FRAMEWORKS FOR MULTI-MODAL HEALTH DATA FUSION TO ADVANCE EQUITABLE AND PERSONALIZED MEDICINE," Zenodo (CERN European Organization for Nuclear Research), Dec. 2023, doi: https://doi.org/10.5281/zenodo.15731939

[18] D. Olufemi, A. O. Ejiade, F. O. Ikwuogu, P. E. Olufemi, and D. Bobie-Ansah, "Securing Software-Defined Networks (SDN) Against Emerging Cyber Threats in 5G and Future Networks – A Comprehensive Review," *International Journal of Engineering Research & Technology (IJERT)*, vol. 14, no. 2, Feb. 2025.

[19] B. White, "Consensus Mechanisms for Critical Infrastructure," *DLT & Security*, vol. 3, no. 2, pp. 145-160, 2024.

[20] Bobie-Ansah, D., and Affram, H., "Impact of secure cloud computing solutions on encouraging small and medium enterprises to participate more actively in e-commerce," *International Journal of Science & Engineering Development Research*, vol. 9, no. 7, pp. 469–483, 2024.

[21] NIST, "Post-Quantum Cryptography Standardization," *NIST IR 8105*, 2024.

[22] R. Humphrey and J. Smith, "Cyber-Physical System Vulnerabilities in the Power Grid," *IEEE Trans. Smart Grid*, vol. 12, no. 4, pp. 3102-3115, 2024.

[23] Olufemi, O. D., Anwansedo, S. B., and Kangethe, L. N., "AI-powered network slicing in cloud-telecom convergence: A case study for ultra-reliable low-latency communication," *International Journal of Computer Applications Technology and Research*, vol. 13, no. 1, pp. 19-48, 2024.

[24] L. Martinez, "Review of SCADA Protocol Vulnerabilities," *Industrial Security Review*, vol. 18, no. 4, pp. 12-25, 2024.

[25] B. B. Sezer et al., "PP-PQB: Privacy-Preserving in Post-Quantum Blockchain-Based Systems: A Systematization of Knowledge," *IEEE Access*, vol. 13, pp. 41382-41405, Mar. 2025.

[26] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," 2008.

[27] M. Rahman et al., "Leveraging GANs for Synthetic Data Generation to Improve Intrusion Detection Systems," *Journal of Future Artificial Intelligence and Technologies*, vol. 1, no. 4, pp. 429-439, Feb. 2025.

[28] NCSC, "Next steps in preparing for post-quantum cryptography," *NCSC.GOV.UK Whitepaper*, 2024.

[29] A. Miller, "Physics-Aware Loss Functions for Cyber-Physical Systems," *Proc. IEEE Conf. on Control Applications*, 2024, pp. 120-128.

[30] O. Akande, "Integrating Blockchain with Federated Learning for Privacy-Preserving Data Analytics Across Decentralized Governmental Health Information Systems," International Journal of Computer Applications Technology and Research, vol. 11, no. 12, pp. 622–637, 2022, doi: https://doi.org/10.7753/ijcatr1112.1025

[31] J. N. Zeyeum, VCO for PLL Frequency Synthesizer, B.Eng. thesis, Helsinki Metropolia Univ. of Applied Sciences, Helsinki, Finland, May 2016, https://doi.org/10.13140/RG.2.2.14528.44800

[32] J. Doe, "Zero-Trust Architectures for Industrial Control Systems," *Blockchain & Cyber Defense*, vol. 5, no. 3, pp. 201-218, 2025.

[33] HELICS Development Team, "Hierarchical Engine for Large-scale Infrastructure Co-simulation," *SoftwareX*, vol. 15, 2023.

[34] F. Zhao, "Semantic Protocol Mapping for GAN-based Fuzzing," *IEEE Trans. Inf. Forensics Security*, vol. 19, pp. 450-462, 2024.

[35] B. C. S. Institute, "Keeping critical national infrastructure safe," *BCS Articles*, Feb. 2025.

[36] R. Gupta, "Metrics for Generative Cyber Models," *IEEE Trans. on Reliability*, vol. 74, no. 1, pp. 102-115, 2025.

[37] "Generative Adversarial Network-Based Network Intrusion Detection System for Supervisory Control and Data Acquisition System," in *Proc. IEEE International Conference on Cyber Security and Intelligence*, Nov. 2024.

[38] T. Nguyen, "Anomaly Detection in ICS using VAE vs GAN," *IEEE Access*, vol. 11, pp. 1540-1555, 2023.

[39] E. Androulaki et al., "Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains," *EuroSys*, 2018.

[40] C. Szegedy et al., "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.

[41] M. Al-Ghamdi, "Evolving Threats in CNI: The Rise of AI-Driven Attacks," *Journal of National Infrastructure Security*, vol. 9, no. 2, pp. 45-60, 2025.

[42] K. Chen et al., "Standardizing KPIs for Infrastructure Resilience," *IEEE Security & Privacy*, vol. 22, no. 1, pp. 34-42, 2024.

[43] J. N. Zeyeum, "Security Implications of Time Synchronization Errors in IoT Networks," Global Journal of Engineering and Technology Advances, vol. 9, no. 2, pp. 92–100, 2021, https://doi.org/10.30574/gjeta.2021.9.2.0151

[44] D. Evans, "Self-Healing Infrastructure: The Future of CNI," *National Defense Tech Review*, vol. 20, no. 3, pp. 78-92, 2025.

[45] G. Wang and L. Zhang, "The Availability Heuristic in Manual VAPT," *Cyber Resilience Quarterly*, vol. 15, pp. 112-125, 2023.