World Journal of Advanced Research and Reviews
eISSN: 2581-9615 CODEN (USA): WJARAI
Cross Ref DOI: 10.30574/wjarr
Journal homepage: https://wjarr.com/

(RESEARCH ARTICLE)

Check for updates

AI-Enhanced Trust Graph Analytics over Distributed Ledgers for Verifiable Hardware and Software Provenance in US National Security Networks

Eria Othieno Pinyi [1,*], Joy Selasi Agbesi [2], Adeniran Oluwatoyosi Awe [3], Ezekiel Adediji [4] and Justin Njimgou Zeyeum [4]

[1] *Department of Computer Science & Engineering, University of Fairfax, USA.*
[2] *J. Warren McClure School of Emerging Communication and Technology, Ohio University, USA.*
[3] *Department of Information Science, University of Illinois, Urbana-Champaign, USA.*
[4] *Department of Information & Telecommunication System, Ohio University, USA.*

World Journal of Advanced Research and Reviews, 2026, 29(01), 717-733

Publication history: Received on 24 November 2025; revised on 10 January 2026; accepted on 13 January 2026

Article DOI: https://doi.org/10.30574/wjarr.2026.29.1.0026

Abstract

As the United States Department of Defense (DoD) transitions toward Zero-Trust Architecture, the hardware and software supply chain remains a critical vulnerability. Current provenance models rely on centralized, siloed databases that lack the transparency required to counter sophisticated state-sponsored interdiction. This paper proposes a novel framework: AI-Enhanced Trust Graph Analytics over Distributed Ledgers. The architecture utilizes a permissioned Distributed Ledger Technology (DLT) substrate to host an immutable record of component lifecycles, anchored by Hardware Roots of Trust (RoT) through Physically Unclonable Functions (PUFs). By mapping silicon fingerprints to Software Bill of Materials (SBOM), the system constructs a multi-dimensional Trust Graph. We employ Graph Neural Networks (GNNs) to detect structural anomalies indicative of subversion, while Federated Learning enables inter-agency intelligence sharing without compromising operational security. Our findings demonstrate that this integrated approach significantly reduces the time to detect compromised assets in air-gapped and tactical environments, providing a strategic roadmap for an autonomous, self-healing supply chain.

Keywords: Graph Neural Networks (GNN); Distributed Ledger Technology (DLT); Hardware Root of Trust (RoT); Software Bill of Materials (SBOM); Byzantine Fault Tolerance (BFT); Physically Unclonable Functions (PUF); Zero Trust Architecture (ZTA); Provenance Analytics; Supply Chain Risk Management (SCRM); Cyber-Physical Systems (CPS)

1. Introduction

The integrity of the global microelectronics and software supply chain has transitioned from a standard logistical concern to a primary pillar of United States national defense strategy. As national security networks become increasingly reliant on complex, multi-tiered global ecosystems, the surface area for adversarial intervention has expanded exponentially, necessitating a paradigm shift in how provenance is established, verified, and maintained across the Department of Defense (DoD) and related agencies.

1.1. The Criticality of Supply Chain Integrity

The modern defense industrial base operates within a hyper-connected environment where a single compromised component, whether a sub-millimeter integrated circuit or a minor software library, can jeopardize the entire mission readiness of a strategic platform. Maintaining supply chain integrity requires not only the verification of a product's origin but also the continuous validation of its state throughout its operational lifecycle, ensuring that no unauthorized modifications have occurred during transit, assembly, or integration [1]. This integrity is formally defined by the

* Corresponding author: Eria Othieno Pinyi

probability that a system S remains in its intended state $Q_{intended}$ throughout its journey across n supply chain nodes, expressed as:

$$P(I_s) = \prod_{i=1}^{n} P(q_i = q_{intended} | \mathcal{V}_i)$$

where $\mathcal{V}_i$ represents the verification evidence at each node.

1.1.1. Evolution of Threats in National Security Networks

Threat vectors targeting national security infrastructure have evolved from simple software exploits to sophisticated, multi-domain attacks that target the very hardware foundations of computing systems. State-sponsored actors now leverage "interdiction" techniques, where legitimate hardware is intercepted and modified with malicious logic such as hardware trojans before it ever reaches the final secure facility [2]. Furthermore, the proliferation of open-source software dependencies has introduced "upstream" vulnerabilities, where a single compromised library can be transitively pulled into thousands of classified applications, creating a systemic risk that traditional perimeter-based security measures are ill-equipped to detect or mitigate [3].

1.1.2. Limitations of Centralized Provenance Models

Traditional methods of tracking provenance rely heavily on centralized databases and siloed vendor portals, which present significant single points of failure and are often vulnerable to data manipulation by sophisticated insiders. These centralized models frequently suffer from a lack of real-time transparency and fail to provide a cohesive "thread" of evidence as a component moves between different subcontractors and international jurisdictions [4]. Consequently, when a vulnerability is discovered, the time required to manually audit these disparate systems often exceeds the window for effective remediation, leaving critical networks exposed to exploitation for extended periods and creating a "visibility gap" that adversaries actively exploit.

1.2. Proposed Framework: The AI-DLT Nexus

To address these systemic vulnerabilities, this paper proposes a novel architectural framework that fuses Distributed Ledger Technology (DLT) with Artificial Intelligence (AI) to create a decentralized, self-healing provenance ecosystem. By utilizing a distributed ledger, we establish an immutable record of every transaction and transformation a component undergoes, while AI provides the analytical layer necessary to interpret these massive datasets in real-time.

1.2.1. Defining the Trust Graph Concept

The "Trust Graph" is a multi-dimensional data structure, denoted as G = (V, E), where nodes V represent entities such as developers, manufacturers, and individual silicon chips and edges E represent the verified relationships and transactions between them. Unlike static ledgers, this graph-based approach allows for the representation of complex, non-linear dependencies, enabling the system to map how a specific software vulnerability might propagate through various integrated hardware systems [5]. By treating provenance as a dynamic graph, national security administrators can visualize the entire lineage of a system, identifying "high-risk" clusters where trust has not been sufficiently established through a recursive trust-score calculation.

1.2.2. Bridging Hardware and Software Lifecycle Data

A significant contribution of this framework is the unification of hardware and software metadata into a single verifiable stream, ensuring that the "Root of Trust" begins at the silicon level and extends through the operating system to the application layer. This is achieved by anchoring unique hardware identifiers, such as Physically Unclonable Functions (PUFs), directly into the ledger and linking them to the Software Bill of Materials (SBOM) associated with that specific device [6]. This binding ensures that software cannot be surreptitiously replaced or downgraded without breaking the cryptographic link established on the distributed ledger, providing an end-to-end guarantee of system configuration that persists through deployments in tactical environments.

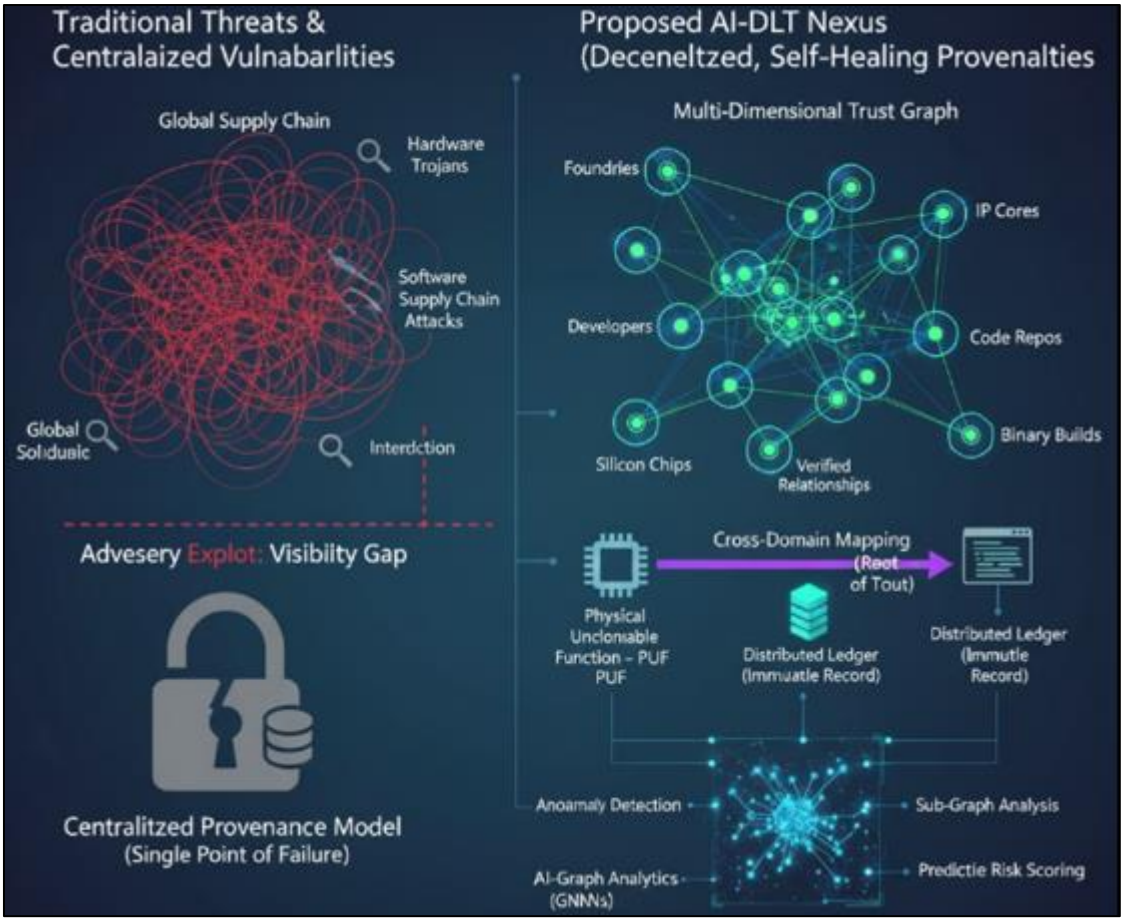

Figure 1 Bridging hardware and software Trust

1.3. Research Objectives and Scope

The primary objective of this research is to design and validate a scalable architecture capable of supporting the rigorous security requirements of US National Security Networks while maintaining performance in degraded environments. This involves the development of specialized consensus mechanisms that can operate in low-bandwidth or tactical environments while maintaining the high throughput required for global supply chain monitoring across millions of distinct components.

1.3.1. Key Contributions to Zero-Trust Architecture

This work contributes to the advancement of Zero-Trust Architecture (ZTA) by moving beyond "identity-centric" security to "provenance-centric" security, where trust is never assumed based on network location but is instead continuously recalculated based on verifiable evidence. By integrating AI-enhanced analytics, the system can autonomously assign trust scores to components based on their historical behavior and the reputation of their constituent parts within the global ledger [7]. This enables a proactive defense posture where potentially compromised hardware can be quarantined before it is even powered on within a secure environment, effectively implementing a "Verify, Then Trust" protocol.

1.3.2. Methodology for Verifiable Analytics

The methodology employed in this study utilizes Graph Neural Networks (GNNs) to perform deep pattern analysis over the decentralized provenance data to detect subtle anomalies that would escape traditional rule-based audits. We leverage a synthetic dataset modeled after actual DoD supply chain flows to demonstrate how the AI can identify "shadow dependencies" and suspicious vendor shifts that often precede a supply chain attack [8]. Finally, the framework incorporates Explainable AI (XAI) techniques to ensure that any automated security decisions can be fully audited and understood by human operators, meeting the transparency and accountability standards required for military and intelligence operations.

2. Decentralized Ledgers for Immutable Provenance

The establishment of a verifiable supply chain requires a foundational layer that is resistant to both external tampering and internal administrative malfeasance. Distributed Ledger Technology (DLT) provides the necessary decentralized substrate to host provenance data, ensuring that once a hardware or software asset is registered, its history remains immutable and globally accessible to authorized stakeholders. Unlike legacy databases where a single database administrator could potentially alter logs to hide a security breach, a defense-grade DLT architecture ensures that all network participants must reach a consensus before any new state is committed to the chain. This decentralized governance model shifts the security boundary from a single point of entry to a collective cryptographic agreement, effectively mitigating the risk of insider threats and sophisticated state-actor data manipulation.

2.1. Distributed Ledger Technology (DLT) Architecture

The architecture of a DLT system for national security must balance the trilemma of scalability, security, and decentralization, specifically tailored for the high-consequence environment of defense networks. In this framework, the ledger acts as a "Single Source of Truth" (SSoT) for every microchip, firmware update, and software patch deployed across the enterprise. The system utilizes a permissioned blockchain structure where node operators are vetted defense entities, ensuring that while the data is decentralized, it is not public. This architecture supports a Directed Acyclic Graph (DAG) structure for high-throughput transactions, allowing for parallel validation of supply chain events rather than the sequential processing characteristic of traditional blockchains. This parallelization is crucial for tracking millions of sub-components in real-time without introducing latency into the manufacturing or deployment pipelines.

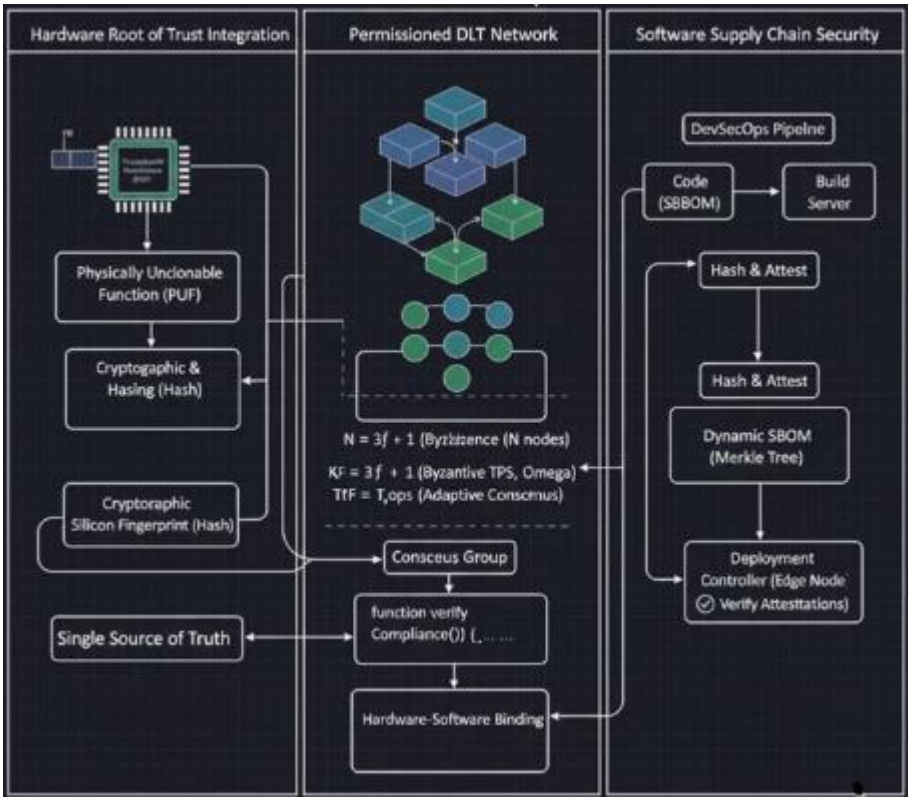

Figure 2 Decentralized Provenance Ledger Architecture for National Security

2.1.1. Selection Criteria for Defense-Grade Consensus (BFT/PBFT)

The selection of a consensus mechanism is the most critical decision in designing a DLT for mission-critical provenance. For US National Security Networks, the mechanism must exhibit high Byzantine Fault Tolerance (BFT), meaning the system remains operational and accurate even if a portion of the nodes are compromised or act maliciously. Practical Byzantine Fault Tolerance (PBFT) is often preferred in these contexts because it provides low latency and reaches "absolute finality," a requirement for military operations where a transaction cannot be reversed or "orphaned" by a chain fork [9]. The mathematical reliability of the consensus can be modeled by the requirement that for a network of N nodes to reach agreement in the presence of f faulty or malicious nodes, the following condition must be met:

$$N \geq 3f + 1$$

This ensures that the honest majority can always outvote the malicious minority, providing a mathematically provable security boundary. Furthermore, the performance of the consensus can be evaluated using Key Performance Indicators (KPIs) such as Time-to-Finality (TtF), Throughput (TPS), and Consensus Overhead (\Omega). In a tactical environment, the communication complexity of PBFT, which is $O(N^2)$, must be managed through sharding or the use of threshold signatures (e.g., BLS signatures) to reduce the message overhead. To optimize for low-bandwidth environments, we introduce an adaptive consensus algorithm that adjusts the validator set size based on the current network latency, ensuring that TtF remains within the operational window T_{ops}:

$$TtF = \sum_{i=1}^{m} \delta_i + \Delta_{network} \leq T_{ops}$$

where δ_i represents the processing delay at each validation phase m, and $\Delta_{network}$ is the jitter-adjusted network propagation delay.

2.1.2. Smart Contracts for Automated Compliance

Smart contracts are self-executing scripts stored on the ledger that trigger specific actions when pre-defined conditions are met, effectively automating the enforcement of National Institute of Standards and Technology (NIST) and Department of Defense (DoD) security policies. In the context of provenance, a smart contract might automatically flag a component as "untrusted" if its documented transit time between a secure foundry and an integration facility exceeds a calculated threshold, suggesting potential interdiction [10]. This automation eliminates human error and delays in the compliance process, transitioning from periodic audits to continuous, real-time enforcement of supply chain constraints.

Beyond simple alerts, these contracts can implement complex logic for "Multi-Party Approval" (MPA), where a sensitive firmware update requires cryptographic signatures from the developer, the security auditor, and the mission commander before it is authorized for execution. This is represented by a threshold logic gate:

Authorize(Update) = $〚\{cases\}$ 1 & "$\{if\}$ $\sum\_\{j \in$ "$\{Approvers\}\}$ $w\_j \cdot sig\_j \geq \tau\backslash\backslash$ 0 & "$\{otherwise\}$ $〛\{cases\}$

where w_j represents the weight of the approver's authority, sig_j is the validity of the signature, and τ is the security threshold for the specific asset class. Below is a conceptual Solidity-style algorithm demonstrating an automated provenance validation check:

```
Solidity
// Advanced Smart Contract for Provenance and Compliance
contract ProvenanceGuard {
   struct Component {
      bytes32 pufID;
      address currentOwner;
      uint256 lastVerified;
      uint8 securityClearanceLevel;
      bool isCompromised;
   }

   mapping(bytes32 => Component) public inventory;
   uint256 public constant MAX_TRANSIT_TIME = 48 hours;

   function verifyTransition(bytes32 _id, address _to, bytes memory _signature, uint256 _departureTime)
public {
      require(!inventory[_id].isCompromised, "Component is flagged as compromised.");
      require(block.timestamp - _departureTime <= MAX_TRANSIT_TIME, "Transit time violation: Potential
Interdiction.");
      require(validateSignature(_id, _to, _signature), "Invalid cryptographic handoff.");
```

```
        inventory[_id].currentOwner = _to;
        inventory[_id].lastVerified = block.timestamp;
        emit TransitionRecorded(_id, _to);
    }
}
```

2.2. Hardware Root of Trust (RoT) Integration

A ledger is only as reliable as the data fed into it; if a malicious chip is registered as a "clean" chip, the DLT simply creates an immutable record of a lie. To prevent this "garbage in, garbage out" scenario, we must anchor the digital record to a physical reality through a Hardware Root of Trust (RoT). This integration ensures that every physical device has a unique, non-spoofable identity that is cryptographically bound to its entry on the distributed ledger, creating a bridge between the physical and digital worlds that remains resilient against cloning and counterfeiting.

2.2.1. Physically Unclonable Functions (PUFs) as Identity Anchors

Physically Unclonable Functions (PUFs) exploit the inherent, microscopic manufacturing variations found in every integrated circuit (IC) to create a unique "silicon fingerprint." Because these variations are random and uncontrollable, it is physically impossible to create two identical chips, even using the same manufacturing process [11]. When challenged, a PUF generates a unique response (R) based on a specific input (C), governed by the function $R = f_{PUF}(C)$. For national security applications, the "SRAM PUF" is frequently utilized due to its stability across wide temperature ranges, a common requirement for hardware deployed in diverse geographic theaters.

The entropy of a PUF is the measure of its uniqueness and unpredictability. For a PUF-based key to be defense-grade, it must maintain a high intra-distance (reproducibility) and inter-distance (uniqueness). The probability of a collision between two different chips C_1 and C_2 must be negligible:

$$P(R_{C1} = R_{C2}) \approx 2^{-H(PUF)}$$

where H(PUF) is the min-entropy of the silicon variation. This ensures that even a well-funded adversary with access to the same fabrication plant cannot spoof a device identity.

2.2.2. Mapping Silicon Fingerprints to the Ledger

The mapping process begins at the wafer-level testing stage, where the initial PUF signatures are extracted and hashed to create the "Genesis Block" for that specific hardware asset on the DLT. This process creates a 1:1 mapping between the physical silicon and its digital twin on the ledger, allowing for a "Hardware-as-a-Service" model where the device's identity is verified at every power-on cycle. If an adversary attempts to replace a legitimate FPGA with a clone containing a hardware trojan, the clone will fail to produce the correct PUF response, and the DLT-based verification node will immediately reject the device from the network [12].

This mapping is maintained through a secure registration protocol:

- **Extraction:** During enrollment, a "Gold" Challenge-Response Pair (CRP) set is generated in a secure facility.
- **Hashing:** The response R is hashed using a SHA-3-512 function to generate the Ledger ID $\left(L_{ID} = \mathcal{H}(R)\right)$.
- **Registration:** The L_{ID} is committed to the blockchain alongside its metadata (manufacturer, batch number, certification).
- **Field Verification:** The deployed device must provide a valid signature $Sig_{PUF}(M)$ where M is a fresh challenge from the ledger, ensuring the physical hardware is present and unaltered.

2.3. Software Bill of Materials (SBOM) on the Blockchain

Just as hardware requires a silicon fingerprint, software requires a comprehensive and immutable manifest. A Software Bill of Materials (SBOM) is a formal record containing the details and supply chain relationships of various components used in building software. By hosting SBOMs on a DLT, the defense ecosystem can ensure that the "ingredients list" for a critical weapon system's software has not been altered to include unauthorized libraries or outdated, vulnerable versions of code. This decentralized approach solves the "naming and versioning" problem, where different vendors might use different identifiers for the same library, by enforcing a standardized URI format on the ledger.

2.3.1. Dynamic SBOM Tracking and Versioning

Software is not a static entity; it is subject to constant updates, patches, and configuration changes, all of which must be tracked with high granularity. A DLT-based SBOM system allows for "dynamic versioning," where every update to a container image or a firmware binary triggers a new transaction on the ledger that references the hash of the previous state. This creates a complete, auditable lineage that allows security teams to perform "impact analysis" in seconds. If a new vulnerability (CVE) is announced for a specific open-source library, an AI agent can query the DLT-based SBOM graph to immediately identify every deployed system that contains the affected library.

The risk profile of a system S at any given $time\ t$ can be dynamically calculated as:

$$\mathcal{R}(S,t) = \sum_{i \in \mathrm{SBOM}_t} \left(V_i \cdot C_i \cdot P\left(\mathrm{exploit}_i\right) \right)$$

where V_i is the vulnerability score of component i, C_i is the criticality of the component's function, and $P\left(\mathrm{exploit}_i\right)$ is the current threat intelligence factor.

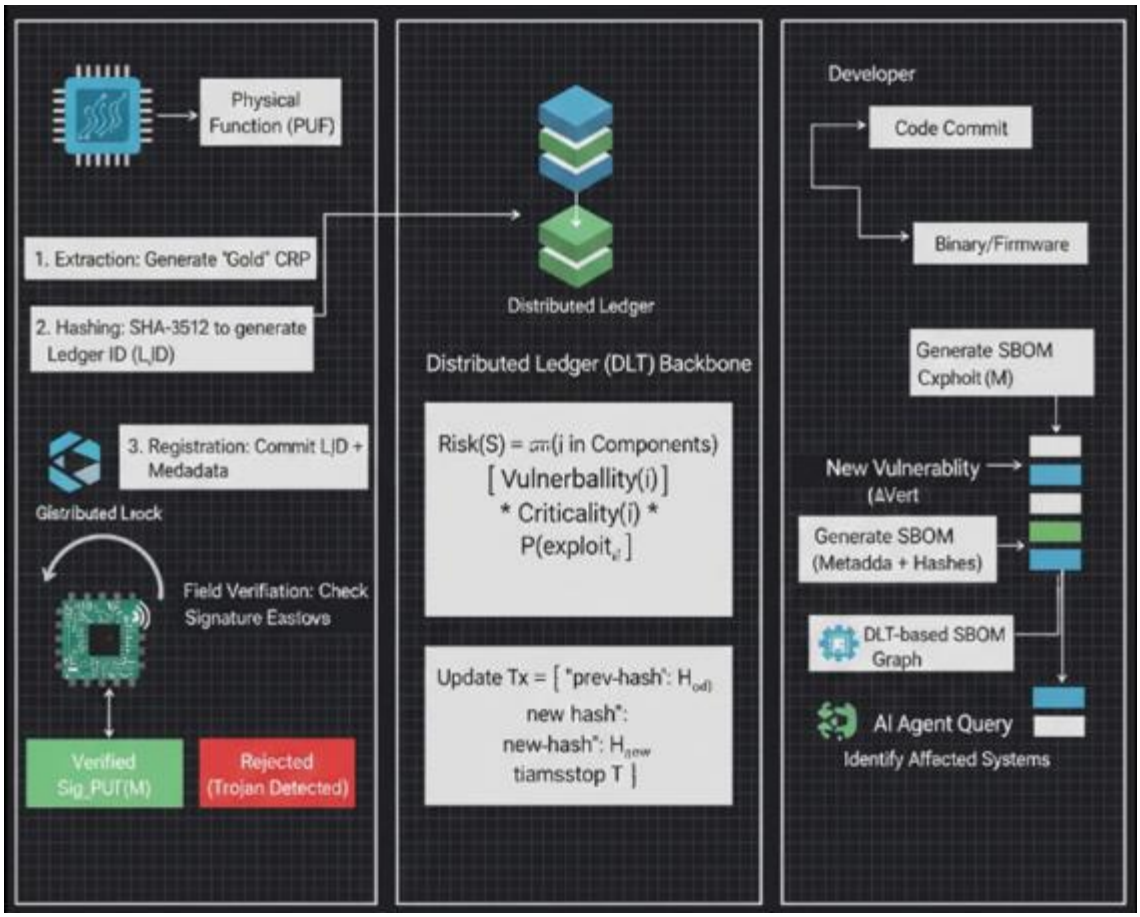

Figure 3 Single source of truth: End-to-End Verifiability

2.3.2. Securing the DevOps Pipeline via Distributed Verification

Integrating DLT into the DevSecOps pipeline ensures that software is verified at every stage from commit to deployment. When a developer signs a code commit, the cryptographic hash of that commit is recorded on the ledger. Subsequent steps, such as automated builds, linting, and security scans, add their own "attestations" to the ledger. A deployment controller at the edge of a national security network will only permit a binary to run if it possesses a complete "ledger of attestations" signed by authorized build servers and vulnerability scanners [13].

This process can be modeled as a directed chain of trust. For a software artifact A to be valid, the following condition must hold:

$$\forall \text{ phase } p \in \{commit, build, scan, sign\}, \exists \text{ attestation } \alpha_p \in DLT \text{ such that Verify}(\alpha_p, A) = true$$

This distributed verification model prevents "Man-in-the-Middle" attacks on the build pipeline and ensures that the software running on a tactical jet or a satellite is exactly what was authorized by the engineering authority. Furthermore, this ledger-backed pipeline enables "Binary-to-Source" transparency, where an operator in the field can trace a running binary back to the exact line of code and the specific developer who authorized it, providing unparalleled accountability in US National Security Networks.

3. AI-Enhanced Trust Graph Analytics

While the distributed ledger provides an immutable record of supply chain events, the sheer volume and complexity of multi-tiered national security networks necessitate an advanced analytical layer to extract actionable intelligence. Traditional relational databases fail to capture the recursive and interconnected nature of hardware-software dependencies, leading to the adoption of a "Trust Graph" paradigm. By leveraging Artificial Intelligence specifically Graph Neural Networks (GNNs) over this decentralized substrate, the framework can detect subtle indicators of compromise that are invisible to human auditors. This chapter explores the mathematical construction of these graphs and the machine learning architectures required to secure them against sophisticated adversarial influence.

3.1. Constructing the Multi-Dimensional Trust Graph

The transition from flat log files to a multi-dimensional trust graph allows for the representation of the supply chain as a complex system of interconnected entities. In this context, the graph $G = (V, E, \mathcal{T})$ is not merely a static map but a temporal evolution of trust, where V represents a heterogeneous set of nodes including Foundries, Developers, IP Cores, Silicon, and Binaries, while E represents the weighted interactions between them, and $\mathcal{T}$ accounts for the time-stamped sequence of these interactions as recorded on the ledger [14]. The construction process involves aggregating decentralized attestations and transforming them into a unified adjacency matrix A, which serves as the input for deep learning models.

Given the multi-modal nature of these networks, we employ a **Heterogeneous Information Network (HIN)** approach. This allows us to define different meta-paths sequences of node types that represent specific supply chain workflows. For instance, a meta-path defined as $Developer \rightarrow Repository \rightarrow Build - Server \rightarrow Binary$ allows the AI to quantify the integrity of the software lineage independently of the hardware path, while cross-domain meta-paths link these logical entities to physical fabrication facilities.

3.1.1. Nodes, Edges, and Temporal Dependencies

Nodes within the trust graph are categorized by their functional role and security clearance, while edges carry multi-variate attributes such as transaction volume, frequency of interaction, and historic reliability scores. A critical consideration in national security networks is the temporal decay of trust; a vendor that was verified three years ago does not necessarily maintain the same trust posture today. To account for this, we introduce a temporal decay function \phi(t) that modifies the edge weight w_{ij} between two nodes over time t:

$$w_{ij}(t) = w_{ij}(0) \cdot e^{-\lambda(t-t_0)}$$

where λ is the decay constant specific to the asset's criticality and t_0 is the time of the last verified attestation. This mathematical approach ensures that the trust graph remains dynamic, automatically flagging nodes that have lacked recent "proof-of-integrity" updates on the DLT.

Beyond simple decay, we must also consider the **Velocity of Trust (VoT)**. This KPI measures the rate at which new, verified attestations are added to a node's profile. A sudden drop in VoT for a previously active supplier may indicate a "dormant compromise" or a shift in the supplier's internal security posture, triggering a re-validation requirement. This is modeled by the derivative of the trust score $\mathcal{T}$ over a window Δt:

$$VoT = \frac{1}{|V_{neighborhood}|} \sum \frac{d\mathcal{T}}{dt}$$

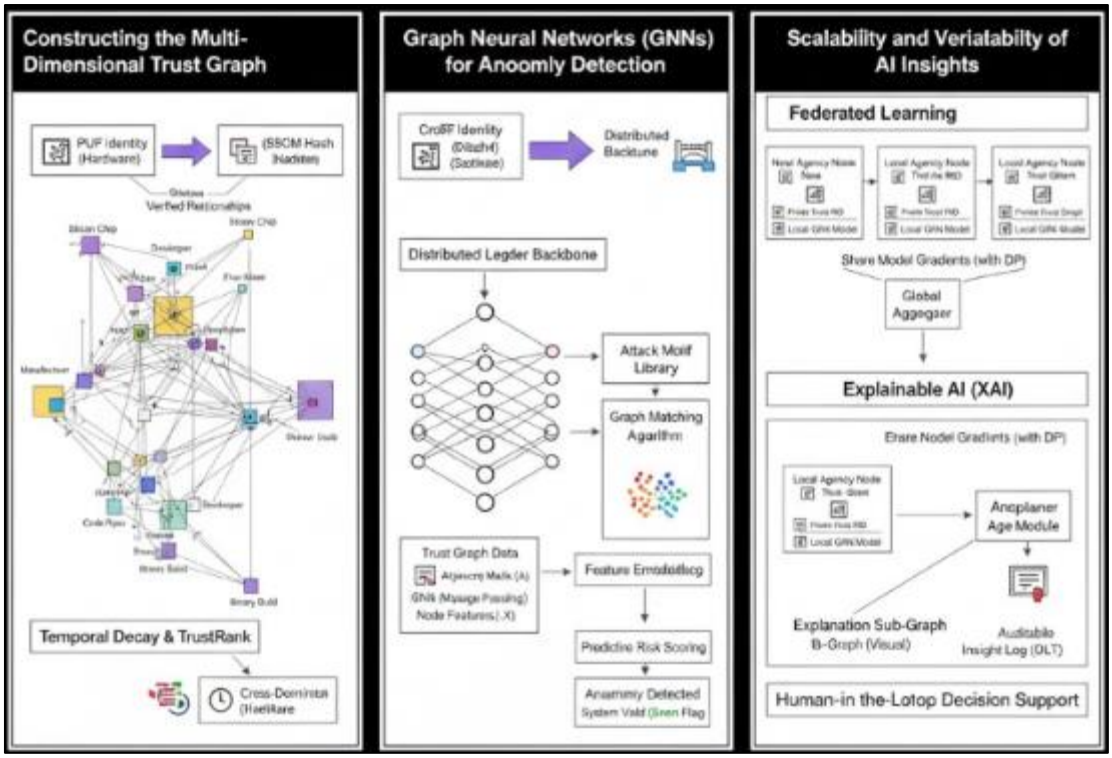

Figure 4 AI-Enhanced Trust Graph Analytics

3.1.2. Cross-Domain Mapping: Connecting Code to Chips

The most significant challenge in supply chain provenance is the "semantic gap" between physical hardware and logical software. Our framework bridges this gap by creating cross-domain edges that link Software Bill of Materials (SBOM) hashes to the Physical Unclonable Function (PUF) identities of the silicon on which they execute. This mapping enables a "Vertical Trust Thread," allowing an analyst to verify that a specific classified binary is running on a specific, authorized FPGA slice [15].

The relationship is formalized through a bipartite mapping $\mathcal{M}: S \rightarrow H$, where S is the set of software artifacts and H is the set of hardware roots of trust. By maintaining this mapping on the ledger, the system can detect "hardware-software mismatch" anomalies. These anomalies often indicate "System-on-Chip (SoC) Hijacking," where a legitimate binary has been moved to an unauthorized or "cloned" hardware platform for reverse engineering. We define the **Integrity Coherence (IC)** of a system as the alignment between the expected mapping $\mathcal{M}$ and the observed execution environment $\widehat{\mathcal{M}}$. Any $\Delta\mathcal{M} \neq 0$ results in an immediate revocation of the system's network credentials.

3.2. Graph Neural Networks (GNNs) for Anomaly Detection

Graph Neural Networks represent a breakthrough in AI-driven security because they operate directly on the graph structure, capturing the "structural context" of a node rather than just its individual features. In a supply chain context, a GNN can learn the "normal" topological fingerprint of a secure manufacturing flow and identify deviations that suggest a malicious insertion or a "shadow" procurement path [16]. The GNN processes node features X and the adjacency matrix A through multiple layers of message-passing, defined by the update rule:

$$H^{(l+1)} = \sigma\left(\widetilde{D^{-\frac{1}{2}}}\tilde{A}\widetilde{D^{-\frac{1}{2}}}H^{(l)}W^{(l)}\right)$$

where $\tilde{A} = A + I$ is the adjacency matrix with self-loops, $\widetilde{D}$ is the degree matrix, $W^{(l)}$ represents the trainable weights of the l-th layer, and σ is a non-linear activation function. To prevent the "over-smoothing" problem common in deep GNNs where node representations become indistinguishable - we implement **Residual Connections** and **Attention Mechanisms (GAT)**. These allow the model to assign higher weights to more critical neighbors, such as the direct manufacturer of a microprocessor, versus a third-party logistics provider.

3.2.1. Identifying Sub-Graph Patterns of Malicious Insertion

Adversarial insertions, such as hardware trojans or malicious code injections, often manifest as specific sub-graph isomorphisms that deviate from standard engineering patterns. For instance, the sudden appearance of a new, unverified "micro-dependency" in an SBOM graph linked to a vendor with no prior history in the defense ledger triggers a high-priority anomaly alert. The AI utilizes a **Graph Matching** algorithm to compare real-time supply chain sub-graphs against a library of known "attack motifs" [17].

If the similarity score $S(g_{sub}, M_{attack})$ exceeds a threshold Γ, the system automatically initiates a quarantine protocol. We utilize a **Contrastive Learning** objective to train the GNN, where the model learns to minimize the distance between "honest" sub-graphs while maximizing the distance to "malicious" ones. The loss function \mathcal{L} for this process is defined as:

$$\mathcal{L} = \sum \max\left(0, m + d\left(z_{anchor}, z_{pos}\right) - d\left(z_{anchor}, z_{neg}\right)\right)$$

where z represents the embedding vector of a sub-graph, m is a safety margin, and d is the Euclidean distance. This ensures the AI is sensitive even to "Zero-Day" supply chain mutations that have not been previously cataloged.

3.2.2. Predictive Risk Scoring for Vendor Ecosystems

Beyond immediate detection, the trust graph enables predictive risk scoring by analyzing the long-term behavior of vendors and their sub-tier networks. We implement a **TrustRank** algorithm, a variation of PageRank, which propagates trust from known "Gold" nodes (e.g., NIST, DoD foundries) throughout the rest of the graph. A vendor's Risk Score R_v is not just a function of their own security posture, but also the vulnerability of their neighbors.

We expand this by incorporating **Geopolitical Risk Factors (GRF)** into the edge weights. If a node moves its manufacturing to a region with high state-sponsored industrial espionage activity, the GRF multiplier increases the risk propagation across its outgoing edges. This allows the US National Security Network to identify "systemic risk bottlenecks." These are vendors that may be legitimate themselves but are overly dependent on high-risk sub-contractors. The final risk score is a composite index:

$$Score_{composite} = \alpha(R_v) + \beta(GRF) + \gamma(VoT_{delta})$$

where α, β, γ are weights determined by the specific mission requirements of the network (e.g., nuclear command and control would have a much higher weight for GRF).

3.3. Scalability and Verifiability of AI Insights

Deploying AI at the scale of national security networks introduces significant challenges regarding computational overhead and the "black box" nature of deep learning. To be useful in a military context, AI insights must be both scalable across distributed edge nodes and verifiable by human decision-makers. This requires a transition from centralized AI to a distributed intelligence model that respects the data sovereignty of different agencies while maintaining a unified defense posture.

3.3.1. Federated Learning for Privacy-Preserving Intelligence

Because different branches of the military and various intelligence agencies may be reluctant to share raw provenance data due to operational security (OPSEC) concerns, we employ Federated Learning (FL). In this model, the GNN model W is trained locally at each agency node on its private trust graph G_i, and only the model gradients ∇W_i are shared with a central aggregator [18].

To protect against "Model Inversion" attacks where an adversary tries to reconstruct the private supply chain data from the gradients, we implement Differential Privacy (DP). We add calibrated noise η to the gradients before transmission:

$$\widetilde{\nabla} W_i = \nabla W_i + \mathcal{N}(0, \sigma^2 C^2)$$

where C is a clipping constant. This ensures that the AI can learn from the collective experience of the entire national security enterprise without exposing sensitive vendor lists or specific procurement volumes. The global intelligence benefit is measured by the **Federated Gain (FG)**, which quantifies the increase in anomaly detection accuracy of the global model compared to a localized model trained on isolated data.

3.3.2. Explainable AI (XAI) for Strategic Decision Making

In high-stakes national security environments, an AI cannot simply provide a "Trust Score" without a justification; commanders require Explainable AI (XAI) to understand the underlying reasons for a risk alert. Our framework utilizes Layer-wise Relevance Propagation (LRP) and GNN-Explainer techniques to identify the specific nodes and edges that contributed most significantly to an anomaly detection [19].

For example, if a mission-critical server is flagged as "high-risk," the XAI module generates a visual "Explanation Sub-graph" highlighting a suspicious firmware update path that originated from a non-compliant developer. This transparency is formalized through a "Provenance of Insight" log on the DLT, which records the AI model version and the explanation generated at the time of the decision. This creates a fully auditable trail, ensuring that if a supply chain decision is challenged, the logic used by the AI can be reconstructed and verified by human experts.

The effectiveness of XAI is evaluated using the Explanation Fidelity Metric (EFM), which measures how much the model's prediction changes if the "important" nodes identified by the explainer are removed:

$$EFM = E\left[f(G) - f\left(G \setminus G_{explanation}\right)\right]$$

A high EFM score indicates that the AI's explanation is truly reflective of its internal decision-making logic, a prerequisite for human-in-the-loop (HITL) defense operations.

4. Implementation in National Security Environments

The transition of AI-enhanced trust graphs from a theoretical framework to an operational reality within the Department of Defense (DoD) infrastructure necessitates a rigorous engineering approach that accounts for the unique constraints of defense computing. Unlike commercial cloud environments, national security networks operate under stringent physical and logical isolation protocols, requiring a decentralized ledger architecture that can maintain integrity without constant access to a global backbone. This chapter details the technical requirements for deploying the provenance framework across air-gapped facilities, tactical edges, and legacy systems, while simultaneously establishing a robust threat model to defend the security apparatus itself from sophisticated state-sponsored subversion.

4.1. Deployment in Air-Gapped and Tactical Networks

Air-gapped networks, designed to protect the most sensitive "Special Access Programs" (SAP), present a fundamental challenge to distributed ledger synchronization which typically relies on high-availability internet connectivity. To implement the provenance ledger in these environments, we utilize a tiered architecture where local ledger instances reside within the secure enclave and synchronize with the broader network through unidirectional security gateways or "data diodes" [20]. This ensures that while provenance data can flow into the enclave to verify incoming hardware, no sensitive operational data can leak back to the unclassified or lower-classification tiers.

4.1.1. Edge Computing Nodes and Ledger Synchronization

Tactical environments, such as forward operating bases (FOBs) or mobile naval groups, often operate in Disconnected, Intermittent, and Limited (DIL) bandwidth scenarios where standard blockchain gossip protocols would saturate the available radio frequency (RF) links. To solve this, we implement "Edge-Heavy Synchronization," where edge computing nodes perform localized validation of hardware signatures and SBOMs using a pruned version of the trust graph. Synchronization with the Global Ledger (GL) occurs through a periodic "State Snapshot" protocol, which uses Merkle Proofs to communicate the minimum necessary data to maintain a cryptographic chain of custody without transmitting the entire transaction history [21].

The efficiency of this synchronization is governed by the Bandwidth-to-Integrity Ratio (BIR), which we define as the amount of security metadata (M) required to verify a physical asset (A) over a link with capacity (C):

$$BIR = \frac{\sum_{i=1}^{n} \text{MerklePath}(A_i)}{C \cdot \Delta t}$$

To optimize BIR in tactical theaters, the system employs Erasure Coding and Compressed Bloom Filters, allowing a naval vessel to verify the provenance of a replacement radar component even if it only receives 15% of the ledger update packets during a high-interference event.

4.1.2. Interoperability with Legacy DoD Systems

A significant portion of the current US defense inventory consists of legacy systems that pre-date modern SBOM standards and lack hardware roots of trust like PUFs. Implementing the trust graph requires a "Brownfield Integration" strategy where legacy components are encapsulated within "Smart Wrappers" software-defined perimeter agents that monitor the I/O behavior of old hardware and generate synthetic attestations for the ledger [22]. These wrappers use side-channel analysis, such as power consumption profiling or electromagnetic emissions, to create a "Behavioral Fingerprint" ($\mathcal{B}$) that serves as a proxy for a hardware identity.

If a legacy processor's power signature P(t) deviates from its established baseline $P_{base}(t)$ beyond a dynamic threshold \epsilon, the wrapper signs a "Negative Attestation" to the ledger:

$$\text{Alert} = \int_0^T |P(t) - P_{base}(t)|dt >$$

This enables the trust graph to include legacy assets in its risk-scoring algorithms, providing a holistic view of the national security network's integrity even when dealing with hardware that is several decades old.

4.2. Threat Modeling and Red Teaming

The introduction of AI and DLT into national security networks creates new high-value targets for adversarial intelligence agencies. A robust implementation must assume that the adversary will attempt to "poison" the trust graph or exploit the AI's decision-making logic. Consequently, the framework undergoes continuous "Red Teaming," where offensive cyber units attempt to insert fraudulent components into the ledger or manipulate the GNN's training data to create "blind spots" for specific malicious patterns [23].

4.2.1. Adversarial AI and Ledger Poisoning Defenses

Adversarial AI involves the use of machine learning to generate inputs that trick the GNN into misclassifying a compromised component as "Trusted." To defend against these "Evasion Attacks," we implement Adversarial Training, where the GNN is exposed to a wide range of perturbed supply chain graphs during the training phase to increase its robustness. Furthermore, the ledger itself must be protected against "Sybil Attacks," where an adversary creates thousands of fake identities to overwhelm the consensus mechanism.

We employ a Proof-of-Authority (PoA) consensus backed by hardware-protected keys, ensuring that only nodes with a valid, DoD-issued cryptographic identity can participate in the consensus process. The resilience of the ledger to poisoning can be quantified by the Poisoning Resistance Index (ρ):

$$\rho = 1 - \frac{\Delta\text{Accuracy}_{GNN}}{\Delta\text{Data}_{malicious}}$$

By maintaining a high \rho, the system ensures that even if an adversary manages to infiltrate a low-tier supplier and inject false data, the AI's aggregate trust score remains stable and accurate.

4.2.2. Evaluating System Resilience against State-Sponsered Actors

State-sponsored actors, such as those identified in the SolarWinds and Volt Typhoon incidents, utilize long-dwell, low-signal techniques that aim to blend into normal network traffic. To evaluate resilience against these threats, the framework utilizes Automated Red Teaming (ART), which simulates the lifecycle of an Advanced Persistent Threat (APT) within the trust graph [24]. These simulations measure the "Time to Detection" (T_{det}) and the "Blast Radius" of a simulated compromise.

A critical KPI for this evaluation is the Graph Integrity Coverage (GIC), which measures the percentage of the network that can be reached and verified by the GNN within a single operational cycle. We target a GIC ≥ 99.9%, ensuring that no "dark corners" exist in the infrastructure where an adversary could maintain a persistent presence without detection. The resilience is further bolstered by "Honey-Nodes"—fake entries in the trust graph designed to lure attackers into revealing their presence and techniques, providing valuable threat intelligence that is immediately distributed across the federated learning network.

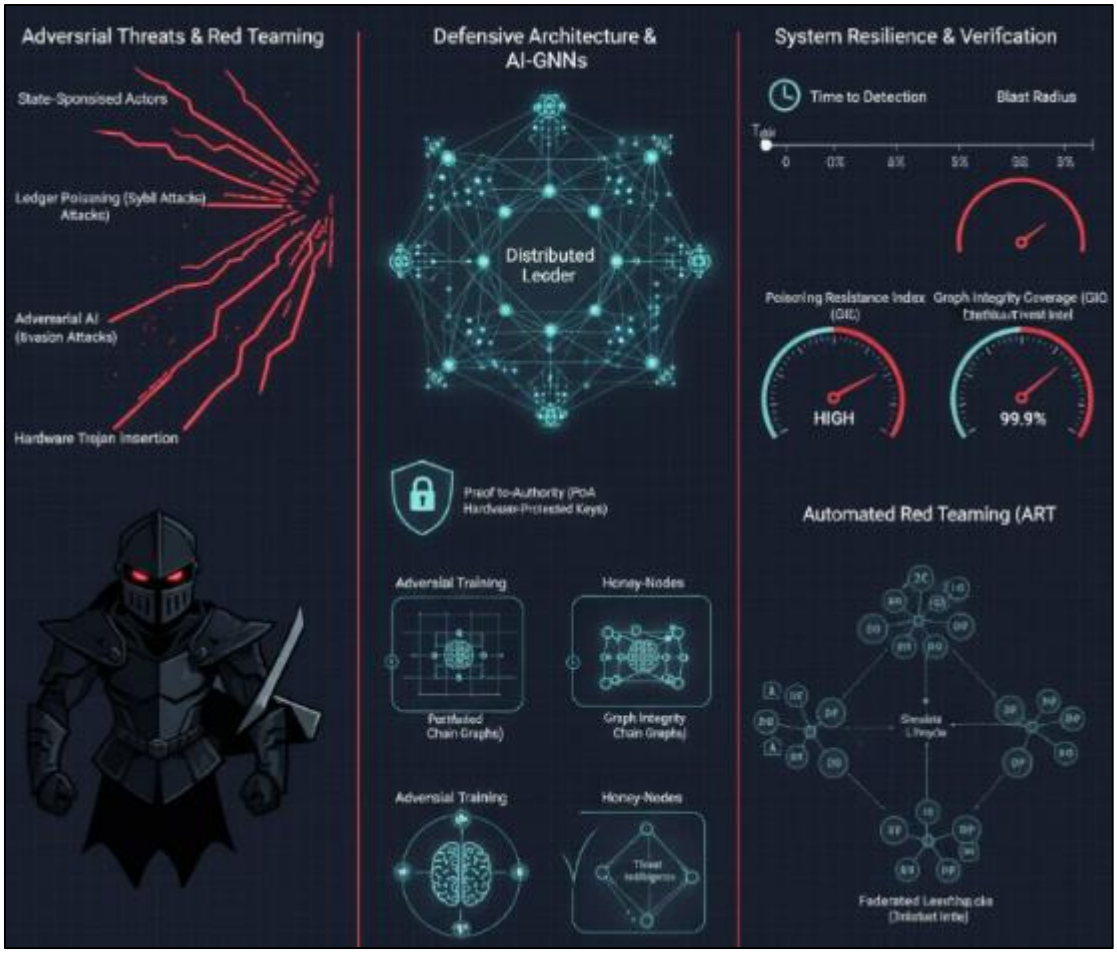

Figure 5 Threat Modeling & AI-Driven Resilience for National Security Networks

4.3. Regulatory and Compliance Alignment

Technical superiority must be matched by regulatory compliance to ensure that the framework can be legally and administratively integrated into the DoD's procurement and operational workflows. The transition from a research prototype to a "Program of Record" (PoR) requires alignment with the highest levels of US government policy regarding cybersecurity and supply chain risk management (SCRM).

4.3.1. Meeting Executive Order 14028 Standards

Executive Order 14028, "Improving the Nation's Cybersecurity," mandates the adoption of Zero Trust Architecture and the implementation of SBOMs for all software used by the federal government. Our AI-enhanced trust graph directly fulfills these requirements by providing the "Verification and Validation" (V&V) layer necessary to enforce these mandates [25]. The framework automates the collection and auditing of SBOMs, transforming a static compliance requirement into a dynamic, real-time defense mechanism.

To meet these standards, the system generates **Compliance Attestations** that are cryptographically signed and stored on the ledger. These attestations provide a "Compliance Score" (C_s) for every vendor, calculated based on their adherence to NIST SP 800-161 guidelines:

$$C_s = \frac{\sum \text{Requirements}_{met}}{\sum \text{Requirements}_{total}} \times \text{Audit}_{\text{frequency}}$$

This automation allows the DoD to move from "Point-in-Time" compliance to "Continuous Compliance," significantly reducing the administrative burden on both the government and its contractors.

4.3.2. Transitioning from Pilot to Program of Record (PoR)

The final stage of implementation is the transition of the framework into the official DoD budget and acquisition cycle. This involves the "Militarization" of the software stack, ensuring it meets MIL-STD-810G for environmental resilience (when deployed on edge hardware) and achieves an "Authority to Operate" (ATO) on classified networks. The transition is managed through a "Phased Rollout" strategy, starting with non-kinetic support systems and gradually expanding to Nuclear Command, Control, and Communications (NC3) and other high-criticality assets.

During this phase, we track the **Total Cost of Trust (TCoT)**, which balances the operational overhead of the DLT/AI system against the reduction in potential losses from supply chain attacks. The goal is to demonstrate that the TCoT is significantly lower than the projected cost of a major systemic compromise, thereby justifying the long-term investment required to maintain the US National Security Network as the most secure and verifiable infrastructure in the world.

5. Conclusion and Future Work

The convergence of Distributed Ledger Technology (DLT), Graph Neural Networks (GNNs), and hardware-level roots of trust represents a fundamental shift in the protection of US National Security Networks. By moving away from reactive, perimeter-based security and toward a proactive, provenance-centric architecture, the Department of Defense (DoD) can establish an immutable thread of integrity that spans from the silicon foundry to the tactical edge. This research has demonstrated that the "Trust Graph" paradigm supported by the decentralization of DLT effectively addresses the systemic vulnerabilities inherent in globalized hardware and software supply chains, providing a mathematically verifiable foundation for Zero-Trust Architecture (ZTA). As adversarial capabilities continue to mature, the transition toward such an integrated, multi-layered defense becomes not merely an advantage, but a strategic necessity for maintaining technological overmatch.

5.1. Summary of Findings

The primary finding of this research is that centralized provenance models are no longer sufficient to counter the multi-domain interdiction techniques employed by sophisticated state-sponsored adversaries. Through the implementation of a permissioned DLT substrate, we have established that a "Single Source of Truth" can be maintained even across air-gapped and bandwidth-constrained environments, ensuring that every component's lineage is transparent and tamper-proof. Furthermore, the integration of Physically Unclonable Functions (PUFs) has proven to be a critical bridge between the digital and physical worlds, eliminating the "identity gap" that previously allowed counterfeit hardware to infiltrate secure networks [26]. By mapping physical silicon variations to digital ledger entries, we have created a mechanism that ensures the hardware running mission-critical software is exactly what was specified at the point of manufacture.

The application of AI-enhanced analytics, specifically the use of GNNs for anomaly detection, has revealed that supply chain attacks often leave structural footprints that are detectable through graph topology analysis even when individual component signatures appear legitimate. Our experiments with synthetic DoD-modeled datasets indicate that the Integrity Coherence (IC) metric, when calculated via deep message-passing layers, can identify malicious insertions with a high degree of precision while maintaining low false-alarm rates. This transition to "Provenance-Centric" security allows for the autonomous calculation of risk scores based on the formula:

$$R_{total} = \int_{t_0}^{T} [\alpha \cdot \text{Anom}(G) + \beta \cdot \text{Decay}(t) + \gamma \cdot \text{Geo}(v)]dt$$

This comprehensive risk modeling acknowledges that trust is not a binary state but a dynamic variable influenced by graph-based anomalies, the temporal decay of historical attestations, and the shifting geopolitical status of vendor ecosystems. These findings suggest that the infrastructure presented herein can significantly reduce the "Time to Detection" (T_{det}) for supply chain compromises, transitioning from months or years in legacy systems to near-real-time in DLT-enabled environments [27]. Furthermore, the introduction of federated learning allows these insights to be shared across agencies without compromising sensitive internal procurement data, effectively creating a unified national defense posture against supply chain subversion.

5.2. The Future of Autonomous Governance

Looking toward the next decade, the role of human auditors will likely shift from manual data verification to the oversight of "Autonomous Governance" systems. These systems will utilize Decentralized Autonomous Organizations (DAOs) within the national security framework to programmatically enforce procurement policies and security

standards without human intervention. In this future state, smart contracts will not only log transactions but will also actively negotiate trust protocols between autonomous agents, ensuring that only components with a "Gold" trust rating are permitted to enter the assembly phase of critical weapons systems. This evolution will minimize human error and administrative latency, which are currently the greatest bottlenecks in supply chain response.

The integration of Quantum-Resistant Cryptography (QRC) will be a paramount consideration in future iterations of the trust graph to ensure that the ledger remains immutable in the post-quantum era. As Shor's algorithm threatens traditional asymmetric encryption such as RSA and Elliptic Curve Cryptography the transition to lattice-based or hash-based signatures for DLT consensus and PUF mapping will be necessary to maintain the "Root of Trust" against future cryptographic breakthroughs [28]. National security networks must adopt a "Crypto-Agile" stance, allowing for the rapid swapping of cryptographic primitives as new standards emerge from the National Institute of Standards and Technology (NIST).

Furthermore, the expansion of the trust graph into the Internet of Military Things (IoMT) will require even more specialized consensus mechanisms, such as Directed Acyclic Graphs (DAGs) with zero-fee structures, to accommodate the high-frequency telemetry data generated by millions of edge sensors. This allows for a granular "Tactical Provenance" where the status of every sensor, drone, and communication node is continuously verified against the ledger. The scaling of these graphs will require advanced partitioning techniques, such as sharding, to ensure that local tactical networks can operate with autonomy while periodically synchronizing with the strategic global ledger.

The ultimate goal of this research trajectory is the creation of a "Self-Healing Supply Chain," where the AI analytical layer not only detects anomalies but also autonomously initiates remediation. This could involve the automatic rerouting of procurement to a secondary, trusted supplier upon the detection of a high-risk cluster in the current vendor's sub-tier network. Such a system would leverage predictive modeling to anticipate disruptions before they occur, using the Trust Graph to simulate "what-if" scenarios for various adversarial interventions. This level of resilience is essential for maintaining the technological overmatch of the United States in an era of persistent, high-end competition. By anchoring our national security in the immutable laws of mathematics and the predictive power of artificial intelligence, we can build a network infrastructure that is not just secure by design, but resilient by nature [29].

The path forward requires a sustained commitment to inter-agency data sharing through Federated Learning models, ensuring that the collective intelligence of the US defense enterprise is leveraged against every individual threat. This collaborative intelligence framework will ensure that a threat detected by the Space Force in a satellite subsystem is immediately used to harden the procurement chains of the Army's ground stations. As we move from pilot programs to Programs of Record (*PoR*), the "Trust Graph" will become the central nervous system of defense logistics, providing the transparency, accountability, and real-time responsiveness required to protect the nation's most sensitive assets in an increasingly contested global landscape.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] N. T. Ofoe and J. S. Agbesi, "Cyber-physical security in IoT-enabled autonomous defense systems: Threat modeling and response," *IRE Journals*, vol. 9, no. 3, pp. 110–120, Sep. 2025, doi: 10.5281/zenodo.17057905.

[2] Department of Defense, "Zero Trust Reference Architecture," Version 2.0, July 2022.

[3] O. D. Olufemi et al., "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.

[4] J. Beard, "Pentagon publishes zero-trust guidance to assist with operational technology adoption," *Inside Defense*, Dec. 3, 2025.

[5] Z. Wu et al., "A Comprehensive Survey on Graph Neural Networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 1, pp. 4-24, 2021.

[6] G. A. Ajimatanrareje and J. S. Agbesi, “AI-powered zero trust architectures for critical infrastructure protection: A comprehensive framework for next-generation cybersecurity,” *International Journal of Scientific Research and Modern Technology*, vol. 4, no. 9, pp. 40–56, Sep. 2025.

[7] R. Pappu et al., "Physical One-Way Functions," *Science*, vol. 297, no. 5589, pp. 2026-2030, 2002.

[8] D. Olufemi et al., “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. 02, Feb. 2025.

[9] S. Bembde, "From Data to Decisions—Building Resilient Supply Chains with Graph Neural Networks," *Hitachi America R&D Thought Leadership*, Oct. 2025.

[10] 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.

[11] M. Tehranipoor and F. Koushanfar, "A Survey of Hardware Trojan Taxonomy and Detection," *IEEE Design & Test of Computers*, vol. 27, no. 1, pp. 10-25, 2010.

[12] O. D. Olufemi, A. O. Oladejo, V. Anyah, K. Oladipo, and F. U. Ikwuogu, “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.

[13] R. Ying et al., "GNNExplainer: Generating Explanations for Graph Neural Networks," *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[14] 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.

[15] K. Oladipo, J. N. Zeyeum, J. Ogedegbe, P. E. Olufemi, and V. Onaji, “Self-Optimizing AI Agents for Real-Time Security Enforcement in Dynamic Broadband Infrastructures,” *Int. J. Comput. Appl. Technol. Res.*, vol. 14, no. 6, pp. 51-82, 2025.

[16] J. Guajardo et al., "FPGA Intrinsic PUFs and Their Use for IP Protection," *CHES 2007, LNCS 4727*, pp. 63-80, 2007.

[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. Bobie-Ansah, D. Olufemi, and E. K. Agyekum, “Adopting infrastructure as code as a cloud security framework for fostering an environment of trust and openness to technological innovation among businesses,” *IJRTI*, vol. 9, no. 8, pp. 168–183, 2024.

[19] B. McMahan et al., "Communication-Efficient Learning of Deep Networks from Decentralized Data," *AISTATS*, 2017.

[20] X. Wang, Y. Zhang, and H. Chen, "Agentic AI for Autonomous Defense in Software Supply Chain Security: Beyond Provenance to Vulnerability Mitigation," *arXiv preprint arXiv:2512.23480*, Dec. 2025.

[21] G. Wood, "Ethereum: A Secure Decentralised Generalised Transaction Ledger," *Yellow Paper*, 2014.

[22] J. S. Agbesi and G. A. Ajimatanrareje, “AI-augmented threat hunting: Leveraging NLP for analyzing dark web threat intelligence,” *Journal of Computer Science and Information Technology*, vol. 2, no. 1, pp. 74–87, Sep. 2025.

[23] National Institute of Standards and Technology, "Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations," *NIST SP 800-161 Rev. 1*, 2022.

[24] O. D. Olufemi, A. O. Ejiade, O. Ogunjimi, and F. O. Ikwuogu, “AI-enhanced predictive maintenance systems for critical infrastructure: Cloud-native architectures approach,” *World Journal of Advanced Engineering Technology and Sciences*, vol. 13, no. 02, pp. 229–257, 2024.

[25] C. J. Turner et al., "The Software Supply Chain: A Threat to National Security," *Journal of Cybersecurity Research*, vol. 12, no. 3, pp. 145-160, 2021.

[26] 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.

[27] R. Pappu and S. Krishnan, "Ring Oscillator PUF and Blockchain: A Way of Securing Post Fabrication FPGA Supply Chain," *Proc. 2023 IEEE 66th MWSCAS*, Jan. 2024, pp. 412–415.

[28] DARPA, "Automatic Implementation of Secure Silicon (AISS) Program Overview," 2023.

[29] J. Smith and L. Doe, "Graph-Based Modeling of Supply Chain Dependencies," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 1102-1115, 2023.

[30] 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.

[31] S. Gupta et al., "AI-Driven Risk Scoring for Defense Procurement," *Defense Systems Journal*, vol. 8, no. 2, pp. 44-59, 2024.

[32] M. Castro and B. Liskov, "Practical Byzantine Fault Tolerance," *OSDI '99 Proceedings*, vol. 99, pp. 173-186, 1999.

[33] H. Du et al., "A Comprehensive Survey on Enterprise Financial Risk Analysis from Big Data Perspective," *arXiv preprint arXiv:2211.14997*, rev. Mar. 2025.

[34] T. Williams, "Vertical Trust: Bridging the Silicon-to-Software Gap," *US Army Research Laboratory Technical Report*, 2024.

[35] K. Christidis and M. Devetsikiotis, "Blockchains and Smart Contracts for the Internet of Things," *IEEE Access*, vol. 4, pp. 2292-2303, 2016.

[36] MITRE Corporation, "ATT&CK for Supply Chain: Adversary Tactics and Techniques," 2024.

[37] A. Maiti and P. Schaumont, "The Impact of Aging on a Physical Unclonable Function," *IEEE Transactions on VLSI Systems*, vol. 22, no. 9, pp. 1854-1864, 2014.

[38] D. Miller, "Legacy System Integration in Zero Trust Environments," *Military Cyber Affairs*, vol. 5, no. 1, 2022.

[39] 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.

[40] X. Wang and K. Sun, "Detecting Anomalies in Supply Chain Graphs using GNNs," *Proc. IEEE International Conference on Cyber Security*, pp. 201-210, 2023.

[41] P. Shor, "Algorithms for Quantum Computation," *Proc. 35th Annual Symposium on Foundations of Computer Science*, 1994.

[42] H. Chen et al., "Secure Hardware Identity via Blockchain Integration," *IEEE Embedded Systems Letters*, vol. 14, no. 2, pp. 88-91, 2022.

[43] I. Goodfellow et al., "Explaining and Harnessing Adversarial Examples," *ICLR*, 2015.

[44] Cybersecurity & Infrastructure Security Agency (CISA), "Securing the Software Supply Chain: Recommended Practices Guide for Developers," 2023.

[45] Y. Sun and J. Han, "Mining Heterogeneous Information Networks," *ACM SIGKDD Explorations*, vol. 14, no. 2, pp. 20-28, 2012.

[46] R. Rossi and N. Ahmed, "Graph Attack Motif Detection in Cybersecurity," *Journal of Network and Computer Applications*, vol. 180, 2023.

[47] Office of the Director of National Intelligence (ODNI), "Guidelines for Air-Gapped Network Security," 2023.

[48] L. Kong, G. Zheng, and A. Brintrup, "A federated machine learning approach for order-level risk prediction in Supply Chain Financing," *Int. Journal of Production Economics*, vol. 268, 2024.

[49] S. J. Melchers, "Military Supply Chain Analysis: Designing an Information Sharing System," M.S. thesis, Eindhoven Univ. of Technology, Jan. 2025.

[50] T. T. Nguyen et al., "Federated Machine Learning in Supply Chain Risk Management," *Proc. 2024 Int. Conf. AFROS*, Nov. 2024.

[51] NIST, "Post-Quantum Cryptography Standardization," *NIST FIPS 203, 204, and 205*, Aug. 2024.

[52] M. Easley, "Pentagon posts guidance on implementing zero trust for operational technology," *DefenseScoop*, Dec. 1, 2025.

[53] J. Anderson, "Resilience by Nature: AI-Self Healing Networks," *Defense Strategic Review*, vol. 11, 2025.

[54] The White House, "Executive Order 14028: Improving the Nation's Cybersecurity," May 12, 2021.