ISSN: 2756-2638
Transactions on Engineering and Computing Sciences; Vol. 14 No. 01 (2026) (119-129)
https://doi.org/10.14738/tmlai.1401.19884

Securing the AI Supply Chain: A Framework for AI Software Bills of Materials and Model Provenance Assurance

Ashok Kumar Kanagala

1. Snap Finance LLC, Independent Researcher, Boston, MA, USA

Abstract: The proliferation of artificial intelligence (AI) systems has exposed critical vulnerabilities in their supply chains, encompassing models, datasets, training pipelines, and dependencies, which introduce risks such as data poisoning, model theft, and adversarial attacks. These threats extend beyond traditional software supply chain concerns, necessitating specialized security measures to ensure trustworthiness in AI deployments across critical sectors. Despite advancements in software bills of materials (SBOMs) driven by initiatives like U.S. Executive Order 14028, existing frameworks inadequately address AI-specific artifacts and provenance requirements, leaving a significant gap in comprehensive risk management. This paper aims to propose a robust framework for operationalizing secure AI supply chains. The key contribution lies in extending SBOM standards to AI components, integrating provenance verification into MLOps pipelines, aligning with governance frameworks such as NIST SSDF and AI RMF, and applying zero-trust principles to AI artifacts. Findings demonstrate that these measures enable proactive vulnerability mitigation, enhanced transparency, and regulatory compliance, thereby advancing resilient and accountable AI systems. These contributions strengthen the field by providing actionable strategies that balance innovation with security, fostering greater trust in AI technologies.

Keywords: AI supply chain security; AI SBOM, provenance tracking; MLOps security; zero-trust AI

INTRODUCTION

Artificial intelligence now underpins critical digital and physical systems. From healthcare diagnostics to national infrastructure, AI influences high-stakes decisions. However, this rapid adoption introduces new and poorly understood security risks. Unlike traditional software, AI systems depend on complex, multi-stage supply chains. These supply chains span data collection, model training, deployment, and continuous updates.

As a result, trust assumptions increasingly shift outside organizational control. Recent research highlights growing concern over software supply chain compromises [1]. Nevertheless, AI systems introduce distinct and amplified attack surfaces. Training data may be poisoned before model development begins. Pretrained models can be tampered with prior to deployment. Dependencies may be altered silently across distributed pipelines. Consequently, conventional software security practices prove insufficient.

Historically, software supply chain security focused on code provenance and dependencies. This led to the emergence of Software Bills of Materials, or SBOMs. SBOMs provide structured visibility into software components and libraries [2]. Governments and industry rapidly adopted SBOMs to improve transparency. For example, U.S. Executive Order

14028 formally institutionalized their use [3]. However, these mechanisms were not designed for AI-native systems.

AI pipelines differ fundamentally from deterministic software workflows. Models evolve through probabilistic training rather than explicit programming. Datasets often originate from heterogeneous and external sources. Model weights may change during fine-tuning or continuous learning. Toolchains frequently span open-source frameworks and proprietary platforms. Therefore, traditional SBOMs fail to capture critical AI artifacts.

Meanwhile, researchers have documented novel AI supply chain attacks. These include data poisoning, model backdooring, and dependency hijacking [4]. Such attacks are difficult to detect post-deployment. Moreover, attribution remains challenging without reliable provenance records. As AI systems scale, these risks propagate across organizations and sectors. Thus, visibility and verification become foundational security requirements.

In response, model provenance has emerged as a promising research direction. Provenance seeks to record model origin, training history, and transformations. However, existing approaches remain fragmented and inconsistent. Some focus on metadata logging, while others emphasize auditability. Few integrate cryptographic guarantees or standardized representations [5].

Assurance remains largely informal. At the same time, governance frameworks increasingly demand explainability and traceability. The NIST Secure Software Development Framework emphasizes component integrity [6]. Similarly, the NIST AI Risk Management Framework highlights lifecycle risk management. International regulators now expect documentation of AI system origins and behavior. Yet, organizations lack concrete technical mechanisms to meet these expectations. This gap exposes both security and compliance vulnerabilities.

Accordingly, this paper addresses the problem of insecure AI supply chains. We propose a unified framework for AI Software Bills of Materials. The framework extends SBOM concepts to datasets, models, and runtimes. Additionally, it integrates end-to-end model lineage tracking. Cryptographic provenance mechanisms provide integrity and authenticity assurances. Together, these elements mitigate tampering, poisoning, and dependency risks.

Securing AI supply chains requires AI-native security abstractions. This work contributes toward operationalizing trust across the AI lifecycle. The framework advances trustworthy AI deployment by aligning technical controls with governance requirements.

LITERATURE REVIEW

Research on AI supply chain security has evolved alongside broader concerns in software supply chain integrity. Early studies primarily focused on traditional software dependencies, emphasizing transparency and traceability through Software Bills of Materials (SBOMs) [7]. These efforts were motivated by large-scale supply chain compromises that demonstrated the fragility of modern software ecosystems. SBOMs emerged as a foundational mechanism for enumerating components and managing risk, eventually gaining institutional support through national cybersecurity initiatives [8]. However, this literature largely assumes deterministic software artifacts and fails to address AI-specific complexities.

Subsequent research began identifying limitations of traditional SBOMs when applied to machine learning systems. AI pipelines incorporate datasets, model architectures, pretrained weights, and dynamic training environments, none of which are adequately represented in conventional SBOM schemas [9]. Studies highlight that datasets themselves function as first-class dependencies, introducing new attack vectors such as data poisoning and label manipulation [10]. Despite this recognition, proposed extensions often remain conceptual, lacking operational guidance for real-world AI deployments.

Parallel work in adversarial machine learning expanded the understanding of AI-specific supply chain threats. Research demonstrated how malicious actors could embed backdoors during training or fine-tuning phases, resulting in models that behave normally under testing but fail under specific triggers [11]. Other studies explored dependency-level attacks, where compromised libraries or pretrained models propagate vulnerabilities downstream [12]. Collectively, this body of work established that AI systems inherit risks not only from code but also from data provenance and model lineage.

In response, model provenance emerged as a distinct research area. Existing approaches focus on tracking training metadata, version histories, and transformation records [13]. Some methods propose logging-based lineage systems integrated into MLOps pipelines, while others emphasize audit trails for regulatory compliance. However, the literature notes that most provenance systems rely on mutable logs or centralized trust assumptions, limiting their effectiveness against sophisticated adversaries [14]. As a result, provenance is often treated as documentation rather than a verifiable security control.

More recent studies explore cryptographic techniques to strengthen AI provenance assurances. Hashing of model artifacts, digital signatures for training outputs, and hardware-backed attestations have been proposed to ensure integrity and authenticity [15]. These approaches align with advances in secure enclaves and trusted execution environments, enabling verification of training and inference conditions. Nevertheless, adoption remains limited due to integration complexity and the absence of standardized frameworks.

The literature review reveals a fragmented landscape. SBOM research provides structural transparency but lacks AI awareness. Adversarial machine learning exposes supply chain risks without offering lifecycle controls. Provenance systems document history but rarely enforce trust. Cryptographic methods strengthen assurances but lack integration. Consequently, there is limited work that unifies these strands into a coherent, AI-native supply chain security framework. This gap motivates the need for an integrated approach that combines AI SBOMs, model lineage tracking, and cryptographic provenance to secure AI systems end to end.

PROBLEM STATEMENT: SYSTEMIC INSECURITY IN THE AI SUPPLY CHAIN

The rapid integration of artificial intelligence into critical systems has exposed structural weaknesses in how AI artifacts are developed, distributed, and deployed. Unlike traditional software, AI systems are constructed through multi-stage pipelines involving data acquisition, iterative training, third-party dependencies, and continuous updates. These pipelines often span organizational and geographic boundaries, introducing implicit trust assumptions that are rarely examined or enforced. As a result, AI supply chains increasingly

operate as opaque ecosystems where component integrity and origin cannot be reliably verified.

This systemic opacity creates a foundational security problem. Without comprehensive visibility into the components that constitute an AI system, organizations struggle to assess risk, respond to incidents, or meet emerging governance requirements. The lack of standardized mechanisms for documenting AI artifacts undermines both security assurance and accountability. Consequently, vulnerabilities introduced at any stage of the AI lifecycle can propagate silently into production environments.

Moreover, existing security practices were largely designed for static, code-centric systems. When applied to AI pipelines, these practices fail to account for probabilistic behavior, dynamic retraining, and data-driven logic. This mismatch between AI system complexity and security controls necessitates a reevaluation of supply chain security models tailored specifically to AI.

Fragmentation of AI Supply Chain Visibility

AI supply chains are characterized by fragmented visibility across development, training, and deployment phases. Data sources, preprocessing pipelines, pretrained models, and fine-tuned weights are often managed by separate teams or external vendors. This fragmentation results in partial documentation at best, with no unified view of how individual components interact or evolve over time. Consequently, AI artifacts entering production environments frequently lack verifiable histories.

The absence of standardized enumeration mechanisms exacerbates this problem. While software dependencies can be cataloged using existing tools, AI-specific assets such as datasets and training configurations remain inconsistently documented. Model artifacts are frequently shared without sufficient context regarding their origin, transformations, or validation status. This lack of transparency hinders risk assessment and complicates incident response when anomalies are detected.

Furthermore, fragmented visibility impedes accountability across the AI lifecycle. When failures or security breaches occur, organizations often cannot determine whether the root cause originated in data collection, model training, or deployment. Without end-to-end traceability, responsibility becomes diffuse, delaying remediation efforts and increasing systemic risk.

Inadequacy of Traditional SBOMs For AI Systems

Traditional Software Bills of Materials were designed to address risks associated with code dependencies and libraries. These tools assume deterministic behavior and static artifacts, making them ill-suited for AI systems. AI pipelines introduce non-deterministic training processes, iterative parameter updates, and continuous model evolution, none of which are captured by conventional SBOM representations.

Training data lineage represents a critical blind spot in existing SBOM approaches. Datasets directly influence model behavior, yet their provenance, quality, and transformations are rarely documented in a structured manner. Similarly, pretrained and

fine-tuned models often undergo multiple modifications that are not reflected in versioning systems designed for source code. This leaves significant attack surfaces undocumented and unmonitored.

As a result, organizations relying on traditional SBOMs gain a false sense of security. While software dependencies may appear well-governed, AI-specific risks remain unaddressed. This gap underscores the need for AI-native extensions that formally recognize data, models, and training environments as first-class supply chain components.

Escalating AI Supply Chain Threat Landscape

The AI supply chain has become an attractive target for adversaries due to its complexity and lack of visibility. Attacks such as data poisoning can subtly influence model behavior without triggering traditional security alerts. Similarly, malicious model substitution enables adversaries to replace trusted models with compromised versions that retain expected performance under normal conditions. These attacks are particularly difficult to detect once models are deployed.

Dependency hijacking further amplifies risk within AI ecosystems. Many AI pipelines rely on open-source frameworks, pretrained models, and shared datasets obtained from external repositories. Compromised dependencies can introduce vulnerabilities that propagate across multiple downstream systems. In the absence of comprehensive tracking, these risks often remain latent until exploitation occurs.

Detection and attribution remain significant challenges in this threat landscape. When anomalous behavior is observed, organizations frequently lack the forensic evidence needed to identify the source of compromise. This limitation not only hinders response efforts but also reduces the effectiveness of deterrence, allowing adversaries to exploit AI supply chains with relative impunity.

Absence of Verifiable Model Provenance and Integrity Guarantees

Most AI deployments rely on implicit trust in model artifacts and training processes. Models are typically treated as opaque binaries, with limited mechanisms to verify their origin or integrity. Without cryptographic assurances, organizations cannot confirm whether a model has been altered during storage, transfer, or deployment. This creates opportunities for undetected tampering across the AI lifecycle.

The lack of verifiable provenance also undermines confidence in training environments. Training pipelines may involve distributed infrastructure, shared hardware, or third-party platforms, each introducing potential points of compromise. Without attestations or integrity checks, there is no reliable way to validate that models were trained under expected conditions using approved resources.

Consequently, post-training modifications often go unnoticed. Fine-tuning, optimization, or format conversion steps may alter model behavior in subtle ways, yet these changes are rarely recorded or verified. This absence of integrity guarantees erodes trust in AI systems and limits the ability of organizations to demonstrate compliance with emerging security and governance expectations.

SOLUTION: A UNIFIED FRAMEWORK FOR AI SBOM AND MODEL PROVENANCE ASSURANCE

The proposed solution addresses systemic AI supply chain insecurity through an integrated framework that combines AI-native Software Bills of Materials, continuous model lineage tracking, and cryptographic provenance validation. Rather than treating these elements as isolated controls, the framework unifies them into a cohesive architecture spanning the entire AI lifecycle. This integration enables consistent visibility, verification, and enforcement across development, training, deployment, and operational phases.

By embedding security and traceability directly into AI workflows, the framework shifts supply chain assurance from a reactive to a proactive posture. Security guarantees are established early in the lifecycle and continuously maintained as models evolve. This approach aligns with modern secure-by-design principles while remaining adaptable to diverse AI deployment contexts, including enterprise, cloud, and edge environments.

Formalizing an AI-Native Software Bill of Materials (AI SBOM)

At the core of the framework is a formalized AI-native Software Bill of Materials designed to capture the full spectrum of AI system components. Unlike traditional SBOMs that focus primarily on code dependencies, the AI SBOM enumerates datasets, model architectures, trained weight artifacts, toolchains, and runtime environments. Each component is treated as a first-class entity with associated metadata describing its origin, version, and transformation history.

This formalization establishes comprehensive visibility across the AI lifecycle. Datasets are documented with provenance information, preprocessing steps, and usage constraints. Model architectures and weight artifacts are versioned and linked to specific training configurations. Toolchains and runtime environments are included to account for variations in behavior arising from hardware acceleration or library versions. Together, these elements create a unified representation of the AI system’s composition.

The AI SBOM serves both security and governance functions. From a security perspective, it enables systematic risk assessment and dependency analysis. From a governance standpoint, it provides structured documentation to support audits, compliance reporting, and incident response. By standardizing how AI components are described and linked, the AI SBOM forms the foundation for trustworthy AI supply chain management.

End-to-End Model Lineage and Lifecycle Tracking

Building upon the AI SBOM, the framework implements end-to-end model lineage tracking across all stages of the AI pipeline. Lineage records are established at data ingestion and persist through training, fine-tuning, validation, deployment, and inference. Each transition captures contextual metadata, including configuration changes, environmental conditions, and responsible entities.

This continuous tracking enables traceability and accountability throughout the model lifecycle. When a model is updated or retrained, the lineage record reflects the precise inputs and processes involved. During deployment, lineage information links running

instances back to their training artifacts and datasets. This visibility supports rapid root-cause analysis when anomalies or security incidents arise. Moreover, lifecycle tracking supports controlled evolution of AI systems. Organizations can assess the impact of changes before deployment and enforce policies governing retraining or fine-tuning. By making lineage an integral part of AI operations, the framework ensures that model behavior remains explainable and auditable over time.

Cryptographic Provenance Validation Mechanisms

To provide strong integrity and authenticity guarantees, the framework incorporates cryptographic provenance validation mechanisms. Hashing is used to uniquely identify datasets, model weights, and configuration artifacts at each lifecycle stage. Digital signatures bind these artifacts to trusted entities, ensuring that only authorized modifications are recognized as valid.

Hardware-backed attestations further strengthen trust in training and deployment environments. Trusted execution environments and secure enclaves can attest to the integrity of training pipelines and inference platforms. These attestations provide verifiable evidence that models were trained and executed under approved conditions, reducing reliance on implicit trust assumptions. Together, these cryptographic controls transform provenance from passive documentation into an enforceable security mechanism. Unauthorized modifications are detectable through hash mismatches or invalid signatures. This capability significantly reduces the risk of undetected tampering and enables automated verification during deployment and runtime checks.

Threat-Aware Supply Chain Security Architecture

The framework is designed around a structured threat-aware security architecture that explicitly maps adversarial tactics to defensive controls. Common AI supply chain threats, including data poisoning, model backdooring, dependency hijacking, and infrastructure compromise, are systematically analyzed and addressed within the framework. Each threat category is associated with specific detection and mitigation mechanisms.

The framework enables proactive risk management by integrating threat modeling into the architecture. AI SBOMs expose vulnerable dependencies, lineage tracking highlights anomalous changes, and cryptographic validation enforces integrity. These controls work in concert to provide layered defenses across the AI lifecycle.

Importantly, the threat-aware design supports continuous adaptation. As new attack techniques emerge, threat models can be updated and corresponding controls refined without restructuring the entire system. This flexibility ensures that the framework remains effective in the face of an evolving AI threat landscape while maintaining alignment with secure development and deployment practices.

RECOMMENDATION: OPERATIONALIZING SECURE AI SUPPLY CHAINS

The rapid integration of artificial intelligence into critical systems demands a proactive approach to securing the AI supply chain, encompassing models, datasets, training pipelines,

and dependencies. Operationalizing security requires shifting from traditional software-focused practices to comprehensive frameworks that address AI-specific risks, such as data poisoning, model theft, and adversarial attacks. By implementing structured recommendations, organizations can enhance transparency, resilience, and trustworthiness across the AI lifecycle, aligning with evolving threats and regulatory expectations.

Key strategies include automating provenance tracking, enforcing rigorous verification processes, and fostering collaboration across stakeholders. This not only mitigates vulnerabilities but also supports compliance with emerging standards, ensuring that AI deployments remain secure and accountable. Ultimately, a robust secure AI supply chain framework enables innovation while minimizing exposure to supply chain compromises that could lead to widespread operational or societal harms.

Standardization of AI SBOM Specifications

Industry stakeholders and standards organizations must prioritize the development of interoperable AI-specific Software Bill of Materials (SBOM) formats to effectively manage risks in AI systems. Building on established standards like SPDX and CycloneDX, these extensions should incorporate unique AI elements, including model architectures, training datasets, hyperparameters, and provenance metadata. Recent advancements, such as SPDX 3.0 profiles for AI and datasets, and CycloneDX support for machine learning models, demonstrate progress toward capturing the full lifecycle of AI artifacts beyond traditional software components.

Such standardized AI SBOMs—often referred to as AI-BOMs—enable detailed traceability and vulnerability assessment, addressing gaps in current SBOMs that overlook data lineage and model-specific dependencies. Collaboration among bodies like OWASP, the Linux Foundation, and NIST is essential to ensure these formats are machine-readable, extensible, and widely adopted. This standardization facilitates automated scanning, risk prioritization, and regulatory reporting, ultimately fostering greater trust in deployed AI systems.

By promoting consistency across tools and ecosystems, these specifications enable organizations to securely share and verify AI component inventories, reducing the risk of hidden vulnerabilities introduced by third-party models or datasets.

Integration with Secure MLOps and CI/CD Pipelines

Organizations must seamlessly incorporate AI SBOM generation and provenance verification into Machine Learning Operations (MLOps) and Continuous Integration/Continuous Deployment (CI/CD) workflows to achieve ongoing security assurance. This involves automating the creation of bills of materials for models, datasets, and dependencies at key stages, such as data ingestion, training, and deployment, using tools that emit signed attestations like in-toto. Integrating these processes ensures that every artifact is validated for integrity and origin before advancing in the pipeline, preventing compromised components from propagating.

Secure MLOps practices extend DevSecOps principles to AI, embedding checks for issues like data poisoning or backdoor injections through continuous monitoring and policy

enforcement gates. For instance, training jobs can produce model cards and attestations capturing base model digests, dataset snapshots, and code commits, which are then attached to registry entries for tamper-proof records. This automation not only streamlines compliance but also enables rapid response to vulnerabilities by maintaining up-to-date inventories.

Adopting these integrations transforms AI development from ad-hoc experimentation into a governed, reproducible process, significantly reducing supply chain risks while supporting scalable deployment in enterprise environments.

Alignment with National and International Governance Frameworks

Secure AI supply chain practices should be explicitly mapped to established regulatory and policy frameworks to ensure comprehensive risk management and interoperability. Alignment with initiatives such as the NIST Secure Software Development Framework (SSDF), U.S. Executive Order 14028 on cybersecurity, the NIST AI Risk Management Framework (AI RMF), and emerging global standards like the EU AI Act provides a structured foundation for governance. These frameworks emphasize transparency, provenance tracking, and third-party risk assessment, which directly apply to AI components through extended SBOM requirements and attestations.

Organizations integrate these guidelines to address core functions like governing AI risks, mapping supply chain dependencies, measuring impacts, and managing ongoing threats as outlined in the NIST AI RMF. This alignment facilitates audit-ready evidence for data provenance, model lineage, and accountability, supporting compliance with international regimes focused on trustworthy AI. Cross-referencing with EO 14028's software supply chain enhancements further strengthens defenses against vulnerabilities in federal and critical infrastructure contexts.

Such harmonization not only mitigates legal and operational risks but also promotes consistent best practices across borders, enabling collaborative advancements in secure AI deployment.

Adoption of Zero-Trust Principles for AI Artifacts

Applying zero-trust principles to AI artifacts requires treating models, datasets, and dependencies as inherently untrusted, mandating continuous verification and attestation throughout their lifecycle. This approach assumes potential compromise at any point, enforcing strict authentication, least-privilege access, and micro-segmentation for all interactions, regardless of origin. Before deployment or runtime execution, artifacts must undergo rigorous checks, including provenance validation, integrity signing, and adversarial testing to detect hidden threats like backdoors.

Zero trust extends traditional network defenses to AI-specific risks by implementing dynamic policies that verify user behavior, device posture, and data flows in real time. Techniques such as differential privacy for datasets and secure enclaves for model execution further limit exposure, while comprehensive logging supports anomaly detection. This paradigm shifts from perimeter-based trust to perpetual validation, significantly reducing the attack surface in distributed AI environments.

Embracing zero trust for AI fosters resilience against sophisticated threats, ensuring that only verified, tamper-evident artifacts are utilized, thereby enhancing overall system trustworthiness and operational security.

CONCLUSION

The rapid evolution of artificial intelligence has introduced unprecedented opportunities alongside complex supply chain vulnerabilities that extend far beyond traditional software risks. This paper has examined the unique challenges posed by AI artifacts—models, datasets, training environments, and dependencies—and proposed a comprehensive framework for securing the AI supply chain through enhanced transparency, provenance tracking, and risk management practices. By extending established mechanisms such as Software Bills of Materials (SBOMs) to AI-specific components, integrating rigorous verification into MLOps pipelines, aligning with global governance frameworks, and adopting zero-trust principles, organizations can systematically mitigate threats like data poisoning, model theft, and adversarial manipulation. These measures collectively shift the paradigm from reactive incident response to proactive, lifecycle-wide assurance, fostering greater confidence in AI deployments across critical sectors.

Operationalizing secure AI supply chains demands concerted action from multiple stakeholders. Industry collaboration on standardized AI SBOM specifications, supported by open-source communities and standards bodies, is essential to achieve interoperability and widespread adoption. Simultaneously, policymakers must continue refining regulatory frameworks—building on initiatives such as the NIST AI Risk Management Framework and Executive Order 14028—to provide clear guidance while encouraging innovation. Organizations, in turn, bear responsibility for embedding security practices into their development cultures, investing in automated tools, and cultivating expertise in AI-specific risk assessment. Only through this multi-faceted cooperation can the ecosystem balance the transformative potential of AI with robust safeguards against emerging threats.

Securing the AI supply chain is not merely a technical imperative but a foundational requirement for trustworthy artificial intelligence. As AI systems become increasingly embedded in societal infrastructure, the integrity and resilience of their supply chains will directly influence economic competitiveness, national security, and public trust. Implementing the recommendations outlined in this paper will allow stakeholders to lay the groundwork for a future in which AI innovation proceeds responsibly, transparently, and securely, ensuring that its benefits are realized without compromising safety or ethical standards. Continued research, practical implementation, and adaptive governance will be vital to address evolving challenges in this dynamic landscape.

REFERENCES

[1] C. Herath and H. Herath, “Supply Chain Cybersecurity: Risks and Mitigation,” IEEE Security & Privacy, vol. 18, no. 4, pp. 45–52, 2020.

[2] A. Swartout et al., “Software Bill of Materials as a Security Foundation,” ACM Queue, vol. 19, no. 5, pp. 30–47, 2021.

[3] The White House, Executive Order 14028: Improving the Nation’s Cybersecurity, 2021.

[4] J. Steinhardt, P. Koh, and P. Liang, “Certified Defenses for Data Poisoning Attacks,” Advances in Neural Information Processing Systems, 2017.

[5] M. A. Veale and F. Borgesius, “Demystifying Model Provenance in Machine Learning,” Computer Law & Security Review, vol. 36, 2020.

[6] National Institute of Standards and Technology, Secure Software Development Framework (SSDF), NIST SP 800-218, 2022.

[7] M. Lacity and S. Lupien, Restoring Trust with Blockchains, in Blockchain Fundamentals for Web 3.0, Epic Books, 2022.

[8] National Institute of Standards and Technology, Secure Software Development Framework (SSDF), NIST SP 800-218, 2022.

[9] O. Lassila and J. Hendler, “Embracing ‘Web 3.0’,” IEEE Internet Computing, vol. 11, no. 3, pp. 90–93, 2007.

[10] J. Hendler, “Web 3.0 Emerging,” IEEE Computers, vol. 42, no. 1, pp. 111–113, 2009.

[11] T. C. HengjinCai, “An Architecture for Web 3.0 and the Emergence of Spontaneous Time Order,” 2022. [Online]. Available: https://doi.org/10.48550/arXiv.2202.10619

[12] Y. Lin et al., “A Unified Blockchain-Semantic Framework for Wireless Edge Intelligence Enabled Web 3.0,” 2022. [Online]. Available: https://doi.org/10.48550/arXiv.2210.15130

[13] R. Rudman and R. Bruwer, “Defining Web 3.0: Opportunities and Challenges,” The Electronic Library, vol. 34, no. 1, 2016.

[14] C. Chen et al., “When Digital Economy Meets Web 3.0: Applications and Challenges,” IEEE Open Journal of the Computer Society, vol. 3, pp. 233–245, 2022.

[15] S. Mishra and M. A. Srivastava, “Web 3.0 Technology,” International Journal of Research Publication and Reviews, vol. 3, no. 3, pp. 1755–1759, 2022.