ISSN: 2756-2638
Transactions on Engineering and Computing Sciences; Vol. 14 No. 02 (2026) (15-26)
https://doi.org/10.14738/tmlai.1402.20153

Managing Nth-Party Risk in AI Supply Chains: A Framework for Assessing Vendor, Model, and Dependency Risks in Multi-Layered AI Ecosystems

Ashok Kumar Kanagala

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

Abstract: AI systems are increasingly dependent on multi-layered supply chains, including foundation models, APIs, datasets, and tooling, creating complex Nth-party risk exposure. Traditional third-party risk management frameworks are inadequate for addressing dynamic dependencies, cascading vulnerabilities, and continuous model updates. This paper proposes a structured framework for assessing and governing Nth-party AI risks, combining supply chain mapping, dependency classification, risk propagation modeling, and multi-dimensional assessment metrics. Continuous monitoring and adaptive risk scoring provide real-time visibility into evolving vulnerabilities, while integration with enterprise risk management and regulatory standards ensures accountability and compliance. Operational recommendations emphasize vendor transparency, cross-functional governance, continuous auditing, and risk-based procurement strategies. By embedding these practices into AI lifecycles, organizations can proactively mitigate inherited risks, reduce systemic exposure, and maintain regulatory compliance. The framework provides a comprehensive approach to Nth-party AI risk, supporting resilient, secure, and auditable AI ecosystems capable of withstanding emerging threats and operational challenges.

Keywords: Nth-party AI risk management, AI supply chain security, Model dependency and vendor risk.

INTRODUCTION

Artificial intelligence systems increasingly rely on complex, interconnected supply chains that extend far beyond traditional vendors and contractual partners. Consequently, modern AI deployments inherit risk from multiple upstream sources, including foundation models, data providers, APIs, and tooling dependencies. As a result, AI risk no longer resides solely within organizational boundaries, but instead propagates across layered ecosystems that are often opaque and difficult to govern.

Historically, third-party risk management frameworks were designed to address direct vendor relationships. However, AI ecosystems introduce deeply layered dependencies and transitive relationships that remain largely invisible to downstream adopters. Therefore, security flaws, compliance gaps, or data integrity issues can silently cascade across interconnected models and services. Thus, organizations may unknowingly deploy AI systems that carry significant inherited risk.

Unlike conventional software, AI systems are rarely developed entirely in-house. Instead, they integrate pretrained models, external services, and shared datasets sourced from multiple providers. Moreover, continuous updates and fine-tuning further complicate dependency visibility over time. Each update may alter a system's risk posture without

explicit notification or contractual review. Consequently, risk assessment becomes dynamic, non-linear, and difficult to manage using static controls.

Nth-party risk emerges when indirect dependencies introduce hidden vulnerabilities beyond an organization's immediate oversight. For example, a foundation model may rely on undisclosed data sources or third-party preprocessing pipelines. Similarly, API providers may integrate additional analytics, monitoring, or logging services. Therefore, accountability becomes diffused across multiple entities. This diffusion significantly challenges existing governance, assurance, and liability models.

In regulated environments, these challenges intensify considerably. Regulatory frameworks increasingly emphasize transparency, traceability, and accountability for AI systems. However, organizations frequently lack visibility into the full dependency chains underlying deployed models. Thus, compliance efforts often become reactive rather than proactive. Moreover, enforcement actions increasingly target inherited risk exposures rather than isolated technical failures.

Existing risk management models remain insufficient for addressing these realities. Traditional approaches assume static vendors, stable software artifacts, and clearly defined service boundaries. In contrast, AI supply chains evolve continuously and autonomously. Therefore, risk cannot be meaningfully assessed at a single point in time. Instead, it requires continuous evaluation, contextual awareness, and adaptive governance mechanisms.

Recent research has begun to examine AI supply chain security, with particular attention to model provenance, data lineage, and dependency transparency. While these studies highlight critical vulnerabilities, many approaches remain fragmented or narrowly scoped. Few frameworks systematically address how risk propagates across multi-layered AI ecosystems. Consequently, a comprehensive model for Nth-party AI risk governance remains underdeveloped.

Additionally, organizations struggle to quantify AI dependency risk in a consistent manner. Current practices rely heavily on qualitative assessments and checklists. However, such methods lack comparability and decision-making precision. Therefore, leaders face difficulty prioritizing mitigation efforts across complex AI portfolios. Quantifiable metrics are essential to support risk-informed governance.

This paper addresses these gaps through a structured framework for managing Nth-party risk in AI supply chains. The proposed approach maps AI dependencies across vendors, models, and services while introducing metrics for risk attribution and propagation. It integrates governance and compliance considerations directly into AI lifecycle management. As a result, inherited risk becomes observable, measurable, and actionable. Ultimately, this work advances more resilient, transparent, and trustworthy AI ecosystems.

LITERATURE REVIEW

Artificial intelligence adoption has intensified the need for effective supply chain risk management. Early research in AI and digital supply chains focused on third-party relationships, emphasizing vendor accountability, contractual obligations, and risk mitigation strategies [1]. These studies provided foundational insights into dependency management, but primarily addressed direct vendors in static software ecosystems.

Consequently, they are insufficient for modern AI systems, which rely on layered, interdependent models and services.

As AI models became more complex, research began highlighting the risks associated with indirect dependencies, or Nth-party relationships. Studies revealed that AI systems inherit vulnerabilities from data providers, API integrations, pre-trained models, and software libraries [2]. These dependencies can introduce cascading failures, making it difficult for organizations to trace risk across multi-layered AI supply chains [3]. Traditional third-party risk frameworks fail to account for these indirect and often opaque connections, leaving organizations exposed to unanticipated operational and compliance challenges.

Methodological approaches to AI supply chain risk vary widely. Some studies employ dependency mapping to identify upstream vulnerabilities, including model provenance and data lineage [4]. Others propose probabilistic risk propagation models to evaluate how failures in one component may affect downstream systems [5]. These approaches collectively underscore that AI risk is dynamic and non-linear, as minor changes in upstream models or APIs can significantly alter overall system resilience.

Quantitative and qualitative assessments have been explored to evaluate AI dependency risk. Qualitative approaches, such as risk scoring frameworks and checklists, provide an initial overview but lack reproducibility and comparability across large portfolios [6]. Quantitative methods, including Bayesian risk modeling and Monte Carlo simulations, offer more precise insights into likelihood and impact, yet their adoption remains limited due to data and resource constraints [7]. A key gap identified in the literature is the absence of standardized metrics that integrate operational, security, and compliance considerations across all AI supply chain layers.

Regulatory and governance research emphasizes that inherited risk is increasingly a compliance concern. Emerging frameworks for AI governance require transparency, accountability, and traceability across all components of AI systems [8]. Organizations must demonstrate proactive risk mitigation not only for direct vendors but also for indirect dependencies. However, studies indicate that most organizations lack visibility into full dependency chains, making compliance challenging and reactive rather than proactive [9].

Despite progress, existing frameworks are fragmented and often narrowly focused. Research typically addresses discrete elements such as data integrity, model provenance, or API risk, without providing a holistic, multi-layered perspective [10]. Moreover, frequent updates in AI models, dynamic tool integrations, and versioning of dependencies necessitate continuous monitoring, which remains underdeveloped in current literature. Consequently, organizations struggle to implement scalable and operational Nth-party risk management practices.

The literature highlights the need for a comprehensive framework that maps, quantifies, and mitigates Nth-party AI risks. While foundational studies have explored dependency identification, risk propagation, and governance alignment, gaps remain in standardized metrics, continuous monitoring, and operational integration. This study seeks to address these gaps by proposing a structured framework for assessing vendor, model, and dependency risks in multi-layered AI ecosystems, enabling resilient and compliant AI operations.

PROBLEM STATEMENT: THE ESCALATING COMPLEXITY OF NTH-PARTY RISK IN AI SUPPLY CHAINS

Artificial intelligence systems are increasingly embedded within complex, multi-layered supply chains that extend far beyond direct vendor relationships. These ecosystems often include foundation models, pre-trained APIs, third-party data providers, and specialized tooling, each introducing its own risk vectors. As AI adoption accelerates, organizations face the challenge of assessing not only direct vendor risks but also the hidden vulnerabilities inherited from upstream and downstream dependencies. The interconnected nature of AI supply chains amplifies the potential impact of a single compromised component, making conventional risk management approaches insufficient. Understanding and mitigating Nth-party risk has therefore become a critical priority for organizations seeking to maintain operational resilience, security, and regulatory compliance.

The dynamic and adaptive nature of AI further complicates risk assessment. Continuous model updates, retraining, and fine-tuning introduce new dependencies and alter the system's behavior in ways that are often unpredictable. Moreover, automated decision-making pipelines rely on layered AI components, creating feedback loops in which errors or vulnerabilities in one component propagate downstream. This environment highlights the need for a systematic framework that can identify, quantify, and manage risk across multiple layers of AI dependencies, moving beyond the limitations of traditional vendor-focused risk assessment.

The stakes of unmanaged Nth-party risk are particularly high in regulated industries, where accountability, transparency, and traceability are mandatory. Without clear insight into AI dependency chains, organizations cannot reliably ensure compliance with frameworks such as GDPR, NIST AI RMF, ISO/IEC 42001, or sector-specific standards. As a result, failures in upstream components can lead to regulatory violations, operational disruptions, and reputational damage. The following sections elaborate on the core challenges associated with the escalating complexity of Nth-party AI risk Introduction.

Opaqueness of Multi-Layered AI Supply Chains

Modern AI systems rely on nested and often opaque dependencies, including foundation models, APIs, datasets, and third-party tooling. Many organizations lack visibility into the full supply chain of their AI components, making it difficult to assess the reliability, security, and provenance of upstream sources. This opaqueness increases uncertainty about potential vulnerabilities and amplifies the challenge of making informed risk-based decisions.

Limited transparency also complicates collaboration between internal teams and external partners. Without clear mapping of dependency relationships, it is challenging to understand how updates, patches, or security flaws in one component may affect other layers of the ecosystem. This lack of visibility makes organizations reliant on assumptions or incomplete information when evaluating AI system safety and resilience.

Furthermore, opaqueness hinders proactive risk mitigation. Security assessments, compliance checks, and monitoring activities are often reactive, triggered only after a vulnerability or incident becomes apparent. To address these challenges, organizations require frameworks capable of systematically mapping multi-layered dependencies and assessing risk propagation across opaque AI supply chains.

Risk Inheritance Across Model and Vendor Dependencies

AI systems inherently inherit risk from the vendors, models, and components they incorporate. Security vulnerabilities, privacy gaps, and compliance deficiencies in a single upstream component can propagate downstream, affecting the integrity and reliability of the entire AI ecosystem. This risk inheritance amplifies the potential impact of minor flaws, making small upstream issues capable of causing systemic failures.

Risk inheritance also complicates accountability. When multiple vendors and models contribute to an AI system, determining responsibility for a breach or failure becomes difficult. Organizations may unknowingly assume liability for vulnerabilities that originate outside their immediate control. This diffusion of responsibility highlights the need for comprehensive Nth-party risk assessment and monitoring strategies.

Moreover, inherited risks are dynamic. AI systems evolve over time through model updates, retraining, and dependency changes, which can introduce new vulnerabilities even after initial deployment. Without ongoing assessment, organizations remain exposed to cascading failures that compromise security, performance, and compliance objectives.

Inadequacy of Traditional Third-Party Risk Management Models

Conventional third-party risk frameworks were designed for static vendor relationships and predictable software artifacts. They often assume that risk is limited to directly contracted vendors, with clearly defined accountability and boundaries. However, AI supply chains are dynamic and multi-layered, involving frequent model updates, transitive dependencies, and automated decision-making processes, which traditional frameworks cannot adequately address.

These models also tend to emphasize compliance checklists and qualitative assessments rather than continuous, quantitative evaluation. As a result, they fail to detect emergent risks that arise from complex interactions between AI components or evolving threat vectors. This gap leaves organizations vulnerable to systemic weaknesses that may be invisible under traditional assessment methods.

Additionally, existing frameworks do not provide mechanisms for real-time monitoring or adaptive risk management. AI systems operate in rapidly changing environments, where new dependencies and vulnerabilities emerge continuously. The inadequacy of conventional models underscores the necessity for a novel framework specifically designed to manage Nth-party AI risks.

Regulatory Exposure and Compliance Gaps

Organizations deploying AI face increasing regulatory scrutiny due to the potential for inherited vulnerabilities to compromise privacy, safety, and fairness. Frameworks such as GDPR, NIST AI RMF, and ISO/IEC 42001 emphasize transparency, accountability, and traceability of AI components, but they often provide limited guidance on managing Nth-party dependencies. This regulatory gap exposes organizations to compliance risks that can result in legal, financial, and reputational consequences.

Unclear accountability for vulnerabilities in upstream vendors or models further complicates compliance. Organizations may be held responsible for incidents originating in components beyond their direct control, creating significant legal and operational challenges. These pressures necessitate a systematic approach to identifying and mitigating inherited risk across AI supply chains.

Moreover, sector-specific regulations, such as those in healthcare, finance, or critical infrastructure, increasingly require auditable records of AI system provenance and risk assessments. Failure to maintain compliance can trigger penalties, operational restrictions, or loss of trust among stakeholders. Addressing these gaps requires a framework that integrates regulatory expectations directly into AI dependency management processes.

Lack of Quantifiable Metrics for AI Dependency Risk

Current AI risk assessment practices often rely on qualitative methods, such as checklists and expert judgment, which lack precision and comparability. Metrics for evaluating model provenance, dependency trustworthiness, and risk amplification across supply chains are largely underdeveloped. Without quantifiable measures, organizations cannot systematically prioritize mitigation efforts or evaluate the effectiveness of interventions.

The absence of standardized metrics also limits transparency and decision-making. Stakeholders, including regulators, auditors, and internal risk teams, require objective evidence of AI system resilience and dependency risk. Qualitative assessments alone fail to provide the rigor needed for informed governance and compliance reporting.

Finally, without measurable metrics, organizations struggle to monitor evolving risks in dynamic AI supply chains. Dependencies, model updates, and emerging vulnerabilities continuously change the risk landscape. Establishing quantifiable indicators of Nth-party risk is therefore critical for continuous assessment, proactive mitigation, and long-term resilience of AI ecosystems.

SOLUTION: A STRUCTURED FRAMEWORK FOR NTH-PARTY AI RISK ASSESSMENT AND GOVERNANCE

The complexity and opacity of modern AI supply chains demand a structured framework to systematically identify, assess, and mitigate Nth-party risks. Traditional third-party risk management approaches are insufficient, as they fail to capture transitive dependencies, dynamic model updates, and cascading vulnerabilities. The proposed solution introduces a comprehensive methodology for mapping AI dependencies, quantifying risk propagation, defining measurable indicators, and integrating findings into enterprise governance structures. By embedding this framework into AI development and operational lifecycles, organizations can achieve proactive, continuous, and auditable management of Nth-party risk.

This framework emphasizes both technical and organizational components. On the technical side, it provides mechanisms to classify dependencies, model risk propagation, and continuously monitor changes in AI supply chains. Organizationally, it ensures that assessment outcomes inform procurement, deployment, and enterprise risk management

decisions. By aligning these elements, the framework enables organizations to bridge the gap between operational risk, compliance obligations, and strategic decision-making.

The solution aims to provide actionable insights and measurable metrics that extend beyond direct vendor risk. It enables organizations to anticipate vulnerabilities, prioritize mitigations, and respond to evolving threats across multi-layered AI ecosystems. The following sub-sections describe the key components of this framework in detail.

AI Supply Chain Mapping and Dependency Classification

The first step in Nth-party risk management involves systematically identifying all components contributing to an AI system. This includes foundation models, APIs, datasets, and tooling layers, as well as any indirect dependencies that may introduce hidden vulnerabilities. Mapping these relationships provides transparency into the AI supply chain, allowing organizations to visualize the full ecosystem of potential risk sources.

Dependency classification further organizes components based on criticality, sensitivity, and operational impact. By categorizing AI elements according to their function, trustworthiness, and update frequency, organizations can prioritize risk assessment efforts where vulnerabilities are likely to have the greatest systemic effect. This structured approach ensures that risk evaluations are both comprehensive and actionable.

Additionally, supply chain mapping supports traceability and auditability. Understanding which models, vendors, or datasets contribute to specific AI outputs enables organizations to pinpoint the origin of vulnerabilities or compliance gaps. This visibility is critical for regulatory reporting, incident response, and proactive mitigation of cascading risks in complex AI environments.

Risk Attribution and Propagation Modeling

Once dependencies are mapped, the framework applies formalized risk attribution to quantify the potential impact of each component. This process evaluates how vulnerabilities in a single model, vendor, or dataset may propagate throughout the AI system. Risk propagation modeling identifies paths through which failures or attacks could cascade, highlighting interdependencies that could amplify the consequences of a breach.

Propagation modeling uses both probabilistic and scenario-based approaches to simulate potential failure points and estimate their systemic impact. By considering factors such as model interconnections, update frequency, and data flow, organizations can identify critical nodes within the supply chain that require enhanced oversight or mitigation measures. This provides a predictive perspective on vulnerabilities, rather than a purely reactive view.

Furthermore, risk propagation insights inform decision-making across multiple levels of the organization. Understanding how risk travels across the AI ecosystem allows teams to prioritize interventions, allocate resources effectively, and implement safeguards that reduce the likelihood of cascading failures. This predictive capability strengthens both operational resilience and regulatory compliance.

Multi-Dimensional Risk Assessment Metrics

A central feature of the framework is the definition of measurable risk indicators across multiple dimensions. These include security, privacy, reliability, compliance, and operational resilience. Each dimension provides quantifiable metrics that can guide assessment, mitigation, and monitoring strategies. For example, security metrics may evaluate vulnerability exposure, while compliance metrics assess alignment with regulatory standards such as GDPR, NIST AI RMF, or ISO/IEC 42001.

Multi-dimensional metrics also enable a holistic view of risk. Organizations can assess not only technical vulnerabilities but also operational and strategic implications, such as the potential impact on service continuity, stakeholder trust, or regulatory obligations. This approach ensures that risk assessments capture both immediate and systemic consequences of Nth-party dependencies.

Additionally, standardized metrics facilitate comparability across AI systems and supply chain components. By applying consistent indicators, organizations can benchmark risk exposure over time, monitor trends, and evaluate the effectiveness of mitigation strategies. This data-driven approach enhances decision-making, transparency, and accountability in managing complex AI ecosystems.

Continuous Monitoring and Adaptive Risk Scoring

Nth-party risk is inherently dynamic, changing as AI models evolve, vendors update services, and regulatory requirements shift. To address this, the framework incorporates continuous monitoring of AI supply chain components and adaptive risk scoring mechanisms. Real-time data on model updates, dependency changes, and security incidents feed into risk evaluations, ensuring that assessments remain current and actionable.

Adaptive risk scoring allows organizations to dynamically adjust priorities based on emerging threats or newly discovered vulnerabilities. By integrating automated monitoring tools with analytics engines, the framework can generate alerts for high-risk components, quantify the potential impact of changes, and recommend appropriate mitigation actions. This continuous approach reduces the likelihood of exposure to cascading failures or regulatory non-compliance.

Moreover, continuous monitoring supports proactive decision-making. Organizations can respond quickly to updates or incidents, maintaining resilience across the AI ecosystem. By combining real-time visibility with adaptive scoring, the framework transforms Nth-party risk management from a static, periodic activity into an ongoing, strategic process that aligns with organizational objectives.

Governance Integration with Enterprise Risk and AI Lifecycle Management

Nth-party risk evaluation is embedded into enterprise governance structures and the AI lifecycle. Findings from risk assessments inform procurement decisions, model selection, deployment strategies, and operational controls, creating alignment between technical risk management and organizational decision-making. This integration enables organizations to manage AI risks proactively and consistently across all stages of the AI lifecycle.

Embedding Nth-party risk into enterprise risk management (ERM) structures ensures accountability at multiple levels. Governance teams can oversee mitigation actions, track residual risks, and report on compliance with internal policies and external regulations. This alignment enhances transparency, facilitates auditability, and provides stakeholders with confidence in the organization's ability to manage complex AI supply chains.

Finally, governance integration promotes continuous improvement. Feedback loops between risk assessments, operational outcomes, and strategic objectives allow organizations to refine dependency mapping, risk models, and mitigation strategies over time. By connecting Nth-party risk management with both AI lifecycle and ERM processes, the framework establishes a sustainable, proactive approach to securing multi-layered AI ecosystems.

RECOMMENDATIONS: OPERATIONALIZING NTH-PARTY AI RISK GOVERNANCE IN REGULATED ENVIRONMENTS

Effective management of Nth-party risk in AI supply chains requires operational strategies that go beyond traditional vendor oversight. Organizations must embed structured governance practices into AI lifecycles, ensuring that risk assessment, monitoring, and mitigation extend to multi-layered dependencies. Proactive operationalization enables organizations to maintain security, reliability, and regulatory compliance, while reducing exposure to cascading failures or legal liabilities. The following recommendations provide a roadmap for embedding Nth-party risk governance in regulated environments.

Operationalizing governance involves not only technical controls but also organizational processes that facilitate accountability, transparency, and cross-functional coordination. Integrating these practices into procurement, model development, and deployment workflows ensures that risk management becomes a continuous and measurable activity. By aligning operational procedures with industry standards and regulatory frameworks, organizations can transform reactive risk assessment into a systematic and auditable process.

These recommendations emphasize a holistic approach to AI risk management, combining disclosure, compliance alignment, continuous auditing, cross-functional oversight, and risk-based vendor management. Collectively, they provide actionable guidance for organizations seeking to manage Nth-party AI risk in complex and highly regulated ecosystems.

Mandate Nth-Party Risk Disclosure and Transparency Requirements

Organizations should require vendors to disclose comprehensive information on model lineage, dependency chains, and upstream risk controls during procurement and onboarding. Transparency ensures that hidden vulnerabilities are identified before deployment and allows risk teams to assess systemic exposure across multi-layered AI ecosystems. Clear disclosure policies help organizations evaluate not only the direct vendor but also transitive dependencies that could introduce cascading failures.

Mandatory risk disclosure facilitates informed decision-making and accountability. Understanding the sources, update frequencies, and security practices of each component

enables organizations to anticipate potential vulnerabilities and implement appropriate mitigation strategies. Without disclosure, Nth-party risk remains opaque, increasing operational and regulatory exposure.

Transparency also supports regulatory compliance. Documentation of model provenance and upstream risk controls provides auditable evidence that organizations have evaluated inherited risks and are taking measures to manage them. By making disclosure a formal requirement, organizations establish a foundation for proactive, accountable, and verifiable Nth-party risk governance.

Align AI Risk Controls with Regulatory and Standards Frameworks

Effective Nth-party risk governance must map directly to established regulatory standards and frameworks. Practices should align with guidance such as NIST AI RMF, ISO/IEC 27001, ISO/IEC 42001, and applicable sectoral regulations. This alignment ensures that risk assessment, mitigation, and monitoring meet recognized best practices and can withstand regulatory scrutiny.

Standards alignment also facilitates consistency across organizational units and AI projects. By applying uniform frameworks, organizations can compare risk exposures, assess mitigation effectiveness, and prioritize interventions in a structured, repeatable manner. This harmonization reduces ambiguity and ensures that risk management processes are both auditable and scalable.

Additionally, aligning controls with regulatory expectations helps organizations anticipate future compliance requirements. Proactive integration of standards-based risk management allows for adaptive responses to evolving regulations, minimizing the likelihood of enforcement actions or penalties. This approach strengthens both operational resilience and organizational credibility.

Implement Continuous Auditing and Assurance Mechanisms

Continuous auditing is critical for managing dynamic Nth-party risks in AI supply chains. Automated logging, validation, and monitoring tools can track changes in model dependencies, vendor updates, and emerging vulnerabilities. This real-time oversight ensures that inherited risks are identified and addressed promptly, rather than relying solely on periodic assessments.

Assurance mechanisms enhance accountability by providing auditable records of risk management activities. Logs, alerts, and validation reports document compliance with internal policies and regulatory standards. These mechanisms enable organizations to demonstrate that inherited risks are actively monitored, prioritized, and mitigated throughout the AI lifecycle.

Continuous auditing also supports predictive risk management. By analyzing historical trends and detecting deviations, organizations can identify high-risk components, anticipate cascading failures, and trigger proactive mitigation strategies. This approach transforms Nth-party risk governance from a reactive to a forward-looking discipline.

Establish Cross-Functional AI Risk Governance Bodies

Managing Nth-party risk effectively requires collaboration across multiple functional areas. Security, legal, compliance, data science, and procurement teams should collectively oversee AI supply chain risks, ensuring that technical, regulatory, and operational perspectives are considered in decision-making. Cross-functional governance enables comprehensive evaluation and prioritization of vulnerabilities.

Collaboration also promotes consistent communication and accountability. By defining clear roles and responsibilities, organizations can ensure that risk mitigation strategies are implemented effectively and that decisions are informed by both technical and business considerations. This integrated approach reduces blind spots and ensures holistic management of AI supply chain risks.

Cross-functional governance facilitates alignment between risk management and strategic objectives. Committees or councils can provide oversight, monitor emerging threats, and guide procurement and deployment decisions, embedding Nth-party risk considerations into the organization's broader AI governance framework. This ensures that risk management is both operationally effective and strategically aligned.

Prioritize Risk-Based Vendor Selection and Exit Strategies

Organizations should adopt risk-weighted scoring models to evaluate AI vendors and dependencies. Scoring criteria may include security posture, compliance history, update frequency, and operational impact. Prioritizing vendors based on quantified risk enables organizations to make informed procurement and retention decisions, reducing exposure to high-risk components.

Defining exit strategies is equally important. High-risk vendors or components should have contingency plans for replacement or mitigation in the event of vulnerability disclosure, service disruption, or regulatory non-compliance. Clear exit pathways ensure that organizations can maintain operational continuity while minimizing cascading risks.

Risk-based vendor management also supports continuous improvement. By monitoring performance, incident history, and dependency evolution, organizations can refine selection criteria, adjust mitigation strategies, and strengthen supply chain resilience over time. This proactive approach ensures that Nth-party risk is actively managed rather than passively tolerated.

CONCLUSION

The increasing complexity of AI supply chains has amplified the need for structured Nth-party risk management. Modern AI systems rely on multi-layered dependencies, including foundation models, APIs, datasets, and tooling, which introduce hidden vulnerabilities that propagate across the ecosystem. Traditional third-party risk frameworks are inadequate for this environment, as they fail to account for dynamic model updates, transitive dependencies, and continuous operational interactions. Organizations deploying AI without systematic evaluation of these risks face heightened exposure to security, privacy, and regulatory failures.

The proposed structured framework addresses these challenges through systematic supply chain mapping, dependency classification, risk propagation modeling, and multi-dimensional assessment metrics. Continuous monitoring and adaptive scoring provide real-time visibility into evolving vulnerabilities, while governance integration ensures alignment with enterprise risk management and regulatory standards. By embedding these processes into AI lifecycles, organizations can achieve proactive mitigation, reduce systemic exposure, and maintain compliance with frameworks such as NIST AI RMF, ISO/IEC 42001, and sector-specific regulations.

Operationalizing Nth-party risk governance further requires transparency, cross-functional collaboration, and risk-informed vendor management. Mandating disclosure, aligning controls with regulatory frameworks, implementing continuous auditing, establishing oversight bodies, and prioritizing risk-based vendor selection ensures comprehensive, actionable, and accountable risk management. Collectively, these practices enable organizations to manage AI supply chain risk holistically, fostering resilient, secure, and compliant AI ecosystems capable of withstanding evolving technological and regulatory challenges.

REFERENCES

[1] A. L. Choi, "Third-Party Risk Management in Digital Supply Chains," Journal of Supply Chain Management, vol. 57, no. 4, pp. 23-35, 2021.

[2] R. K. Gupta, "Understanding Indirect Dependencies in AI Systems," AI & Society, vol. 37, pp. 1221-1236, 2022.

[3] J. S. Lee, "Cascading Failures in Multi-Layered AI Supply Chains," Computers & Security, vol. 114, 102597, 2022.

[4] S. Patel, "Dependency Mapping for AI Model Risk Assessment," Journal of Risk Research, vol. 25, no. 7, pp. 837-854, 2022.

[5] M. Zhang, "Probabilistic Risk Propagation Models for AI Ecosystems," ACM Transactions on AI, vol. 3, no. 2, pp. 1-19, 2022.

[6] K. V. Rao, "Qualitative Approaches to AI Vendor Risk Assessment," International Journal of Information Management, vol. 63, 102453, 2022.

[7] L. Hernandez, "Quantitative Risk Assessment for AI Supply Chains," IEEE Access, vol. 10, pp. 45678-45691, 2022.

[8] OECD, AI Governance and Supply Chain Risk Management, OECD Publishing, 2023.

[9] ISO/IEC, ISO/IEC 42001: Artificial Intelligence Management Systems, International Organization for Standardization, 2023.

[10] D. S. Kumar, "Challenges in Multi-Layered AI Dependency Management," Computers in Industry, vol. 142, 103704, 2022.