CSAI Foundation | Cloud Security Alliance

NVD Triage Overhaul: AI Tooling's Governance Blind Spot

How NIST's Risk-Based Enrichment Model Leaves AI Infrastructure Vulnerabilities Unseen

2026-05-11

AI-assisted Rapid Research

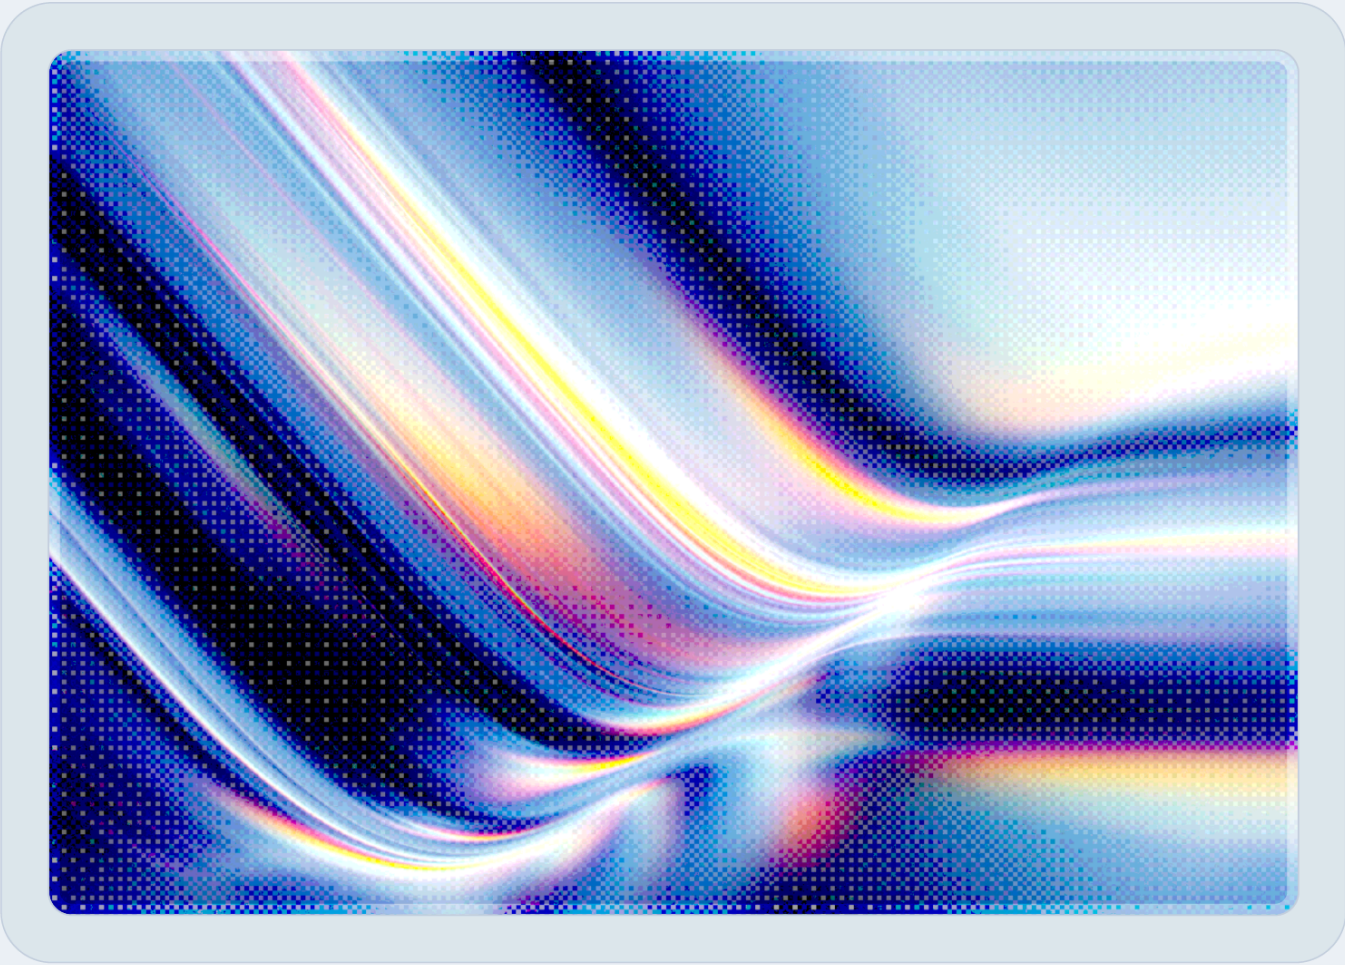

CSAI

CSA cloud security alliance®

Key Takeaways

- On April 15, 2026, NIST formally shifted the National Vulnerability Database to a risk-based triage model, limiting full enrichment to CVEs appearing in CISA's Known Exploited Vulnerabilities catalog, federal government software, and Executive Order 14028 critical software designations — an estimated 15–20% of total CVE volume [1][2][20].
- Most AI infrastructure libraries — including PyTorch, Hugging Face Transformers, LangChain, LangGraph, ONNX Runtime, and emerging Model Context Protocol (MCP) server implementations — fall outside all three NVD priority categories, meaning their CVEs will routinely land in the "Not Scheduled" state without CVSS severity scores, CPE product mappings, or CWE weakness classifications [3][4].
- Without enrichment metadata, the automated scanners, SIEM correlation rules, and patch-management workflows that enterprise security teams rely on cannot operationalize these vulnerabilities — creating a systematic blind spot precisely where AI adoption is accelerating attack surface [5][6].
- Recent critical vulnerabilities in AI tooling — including a CVSS 9.3 PyTorch deserialization flaw (CVE-2025-32434), a CVSS 9.3 LangChain serialization injection (CVE-2025-68664), and a CVSS 8.8 Azure MCP Server SSRF vulnerability (CVE-2026-26118) — demonstrate the severity of what will increasingly fall outside NVD's enrichment scope [7][8][9].
- Organizations operating AI infrastructure must now treat AI-specific vulnerability intelligence as a distinct, proactively managed discipline rather than an output of standard patch-management pipelines.

Background

The National Vulnerability Database has served as the foundational enrichment layer for enterprise vulnerability management for more than two decades. When a CVE is submitted to the system, NIST analysts add structured metadata: CVSS base scores that quantify severity, Common Platform Enumeration identifiers that link vulnerabilities to specific product versions, and Common Weakness

Enumeration classifications that describe root-cause weaknesses. This enrichment is what allows downstream tools — vulnerability scanners, asset management platforms, security orchestration workflows — to ingest raw CVE identifiers and surface actionable, prioritized findings for security teams.

Beginning in early 2024, the universal enrichment model showed measurable strain. Significant budget reductions at NIST triggered immediate staff constraints on the NVD program, which at its peak has only 21 analysts responsible for enriching every vulnerability submission from across the global software ecosystem [10][11]. By May 2024, reporting indicated that over 90% of newly submitted CVEs were awaiting analysis without enrichment [10]. NIST increased output substantially — enriching nearly 42,000 CVEs in 2025, a 45% year-over-year increase — yet still fell further behind, as CVE submissions surged 263% between 2020 and 2025 [1][12]. In the first quarter of 2026 alone, incoming submissions ran approximately one-third higher than the same period a year earlier [1].

Faced with a structural mismatch between analyst capacity and submission volume, NIST formally adopted a new operating model effective April 15, 2026. Rather than attempting universal enrichment, the agency moved to risk-based triage. CVEs appearing in CISA's Known Exploited Vulnerabilities catalog receive immediate priority. CVEs affecting software used within the federal government and software meeting Executive Order 14028's definition of critical software receive secondary priority. All other CVEs are categorized as "Lowest Priority — Not Scheduled for Immediate Enrichment." Additionally, all backlogged CVEs with an NVD publish date before March 1, 2026, that remained unenriched were formally moved into the "Not Scheduled" category rather than remaining in "Awaiting Analysis" [1][2][3].

This policy is operationally defensible given NIST's resource constraints and mission alignment. However, it embeds a structural assumption that is poorly suited to the current threat environment: that the software categories most likely to carry actively dangerous vulnerabilities are those already represented in federal procurement registries and CISA's exploitation tracking.

Security Analysis

The Criteria Mismatch for AI Infrastructure

The three NVD enrichment priority tiers are each anchored to existing federal governance instruments. CISA's KEV catalog reflects vulnerabilities with confirmed exploitation in the wild, though confirmation and listing typically lags initial disclosure by weeks or months. Federal software registries reflect procurement decisions made through formal acquisition processes, which favor established commercial

software with enterprise licensing and FedRAMP authorization. Executive Order 14028's critical software designations focus on software performing privileged functions in federal networks, such as operating systems, identity management platforms, and network security tools.

AI infrastructure libraries occupy a fundamentally different position in the software ecosystem. Frameworks such as PyTorch, Hugging Face Transformers, LangChain, ONNX Runtime, and the emerging category of MCP server implementations are primarily distributed through package managers rather than enterprise procurement channels. They are maintained by open-source communities and commercial AI companies whose customer base extends far beyond the federal government. They underpin training pipelines, inference infrastructure, autonomous agent runtimes, and AI-augmented developer tooling at scale — but they are not typically represented in federal software procurement inventories, and as of the time of writing, no CVEs affecting PyTorch, LangChain, LangGraph, ONNX Runtime, or Hugging Face Transformers appear in CISA's KEV catalog [4][3].

The consequence is structural: vulnerabilities in these libraries will, by default, receive no NVD enrichment. They will be assigned a CVE identifier, published in the NVD with only the submitting CVE Numbering Authority's base information, and receive no CVSS score from NIST, no CPE product mappings, and no CWE classification. Without CPE mappings in particular, the vulnerability cannot be linked to specific library versions in automated tooling, rendering it effectively invisible to scanners that depend on NVD as their data source.

The Scale of Recent AI Toolchain Vulnerabilities

The operational stakes of this blind spot are not hypothetical. The last eighteen months have produced a series of high-severity vulnerabilities across the core AI infrastructure stack that illustrate the exposure.

In April 2025, researchers disclosed CVE-2025-32434, a critical deserialization vulnerability in PyTorch with a CVSS base score of 9.3. The flaw undermines a foundational assumption in machine learning workflows — that model weights can be safely loaded in isolation — by allowing arbitrary code execution during deserialization without requiring network access or elevated privileges [7]. Given PyTorch's broad adoption across both research and production inference environments, the vulnerability's blast radius spans a significant proportion of enterprise AI deployments.

In December 2025, a serialization injection vulnerability in LangChain's core serialization functions was assigned CVE-2025-68664 and rated CVSS 9.3. Dubbed "LangGrinch" by researchers, the flaw allows attackers to extract environment variable secrets — including API keys and authentication tokens — during deserialization operations when `secrets_from_env=True` is set, a common configuration in agent frameworks [8]. A related vulnerability in LangChain.js, CVE-2025-68665, carries a CVSS score of 8.6.

Additional disclosures in early 2026 identified a path traversal vulnerability in LangChain's prompt-loading API (CVE-2026-34070) and an SQL injection flaw in LangGraph's SQLite checkpoint implementation (CVE-2025-67644) [13].

In March 2026, Microsoft patched CVE-2026-26118, a server-side request forgery vulnerability in Azure MCP Server Tools carrying a CVSS score of 8.8. MCP servers act as intermediaries between large language models and operational systems — code repositories, cloud APIs, enterprise data services — and the vulnerability enables an attacker to capture managed identity tokens by supplying malicious URLs as tool parameters, effectively granting access to any resource authorized to the MCP server's identity [9]. This class of vulnerability is particularly significant because the MCP ecosystem is growing rapidly and largely outside established software procurement frameworks.

None of these libraries are federal procurement standards. The probability that their CVEs will receive NVD enrichment under the new triage model is low.

Compounding Factors: Threat Volume and AI Toolchain Targeting

The NVD enrichment gap arrives at a moment of increasing adversarial focus on AI infrastructure. JFrog's 2025 Software Supply Chain Security report documented a 6.5-fold increase in malicious models uploaded to the Hugging Face model repository through 2024 [21]. A separate analysis by Protect AI, partnered with Hugging Face for model security scanning, identified 352,000 unsafe or suspicious issues across 51,700 Hugging Face models by April 2025 [22]. Open-source malware detections across major software repositories increased 73% in 2025 compared with 2024, with AI-adjacent components representing a growing share of that volume [14][24]. IBM's 2025 Cost of a Data Breach Report found that 13% of organizations had experienced breaches of AI models or applications, with 97% of those organizations lacking adequate AI access controls [14][23].

AI-assisted vulnerability discovery has accelerated CVE submissions across all software categories, contributing to the volume surge that forced NIST's triage model [15]. Large language models are being used to analyze code at scale and generate CVE submissions affecting the full breadth of the software ecosystem, not AI tooling specifically. The resulting triage policy nevertheless deprioritizes AI infrastructure libraries on the basis of federal procurement criteria — creating a coverage gap that is structural rather than a product of AI toolchain volume.

The downstream effect on scanner fidelity is concrete. Vulnerability scanners correlate detected software components against NVD's CPE dictionary to identify matches. When a vulnerable library version is present in an environment but its CVE lacks NVD enrichment and therefore carries no CPE mapping, the scanner reports no finding. Security teams operating under the assumption that their tooling provides comprehensive coverage will not be prompted to investigate.

Governance Infrastructure Has Not Kept Pace

Standard vulnerability management governance documentation — security policies, risk registers, compliance evidence packages — typically references NIST NVD enrichment as the authoritative basis for severity classification, reflecting its historical role as the primary public source of CVSS scores and CPE mappings. Under the previous universal enrichment model, this reference was largely unproblematic. Under the new triage model, it creates a documentation gap that auditors and regulators are likely to flag as programs mature their AI security posture.

SBOM practices, which are increasingly required under Executive Order 14028 and sector-specific regulations, help identify which components are present in an environment but do not resolve the enrichment gap. An SBOM entry for a specific version of LangChain can confirm the library's presence; without NVD enrichment, it provides no connection to known vulnerabilities affecting that version. The AI-BOM concept — extended bill of materials covering model weights, training datasets, fine-tuning provenance, and inference dependencies — is emerging in NIST's AI Risk Management Framework and CSA's AI Controls Matrix guidance, but tooling support for operationalizing AI-BOMs against vulnerability intelligence is still developing [16][17].

Recommendations

Immediate Actions

Security teams should audit their vulnerability management pipeline's data dependencies and explicitly document whether CVSS scores, CPE data, and CWE classifications originate from NVD or supplemental sources. This audit will surface the degree to which current scanner configurations are implicitly dependent on NVD enrichment for AI-adjacent components. Supplemental enrichment providers — including VulnCheck, Black Duck's BDSA service, and similar commercial sources — maintain independent enrichment pipelines that continue to cover software categories outside NVD's new priority scope and should be evaluated as additions to existing data feeds [5][18].

For AI infrastructure specifically, security teams should subscribe directly to the security advisories published by AI framework maintainers: PyTorch's GitHub Security Advisories, Hugging Face's security disclosure process, LangChain and LangGraph release notes, and the GitHub Advisory Database entries maintained by CNA-registered organizations. These sources provide enrichment data that does not depend on NVD and often predates NVD publication regardless of enrichment status.

Short-Term Mitigations

Organizations should establish a curated AI software inventory — a working AI-BOM that enumerates inference frameworks, orchestration libraries, model serving platforms, and MCP server components with their current versions. This inventory becomes the basis for a targeted vulnerability watch process: when a CVE is disclosed affecting any enumerated component, it can be manually triaged and routed to the appropriate remediation owner regardless of NVD enrichment status. The inventory should be integrated into change management processes so that updates to AI components trigger a vulnerability review before deployment.

Vulnerability management policy documentation should be updated to explicitly acknowledge the changed NVD enrichment model and specify the organization's alternative data sources, quality-assurance methodology for CNA-supplied CVSS scores, and the escalation path for CVEs where enrichment data is absent or conflicting. This documentation update is required for organizations subject to audit under frameworks that reference NVD as a normative source.

Container and package signing controls should be enforced for AI model artifacts pulled from registries such as Hugging Face, in recognition of the documented increase in malicious model uploads. Integrity verification at ingest time reduces the risk that a compromised model file introduces a threat vector that falls entirely outside the CVE and NVD systems.

Strategic Considerations

The NVD's new triage criteria embed an implicit policy judgment: that software not appearing in federal procurement or CISA's exploitation catalog is lower priority. For many software categories, this judgment is defensible. For AI infrastructure libraries that underpin enterprise AI deployments at scale, the judgment may not age well. As AI-assisted vulnerability discovery continues to accelerate submission rates and as adversarial targeting of AI toolchains continues to intensify, the gap between NVD's enrichment scope and the actual enterprise attack surface is likely to widen.

Industry engagement with NIST and CISA to advocate for AI infrastructure libraries to be considered within the scope of critical software designations under EO 14028 is a long-cycle effort but one that would address the structural mismatch at its source. In parallel, organizations should engage with CISA's CVE Numbering Authority program to understand how AI library maintainers can be supported in becoming CNAs, which would allow them to publish enriched CVSS data directly with their CVE submissions rather than waiting for NVD analysis.

At a strategic level, organizations integrating AI infrastructure into core business processes should treat AI-specific vulnerability intelligence as a distinct discipline requiring dedicated tooling, analyst attention, and governance documentation — not as a subset of general vulnerability management that can be served by existing pipelines without modification.

CSA Resource Alignment

Several CSA frameworks directly address the governance gap exposed by the NVD triage overhaul.

The CSA AI Controls Matrix (AICM) v1.0, released in July 2025, includes a dedicated AI Supply Chain Security domain covering 18 security domains and 243 control objectives across the AI lifecycle [16]. Controls within this domain address verification of AI-generated code artifacts, integrity of model outputs, and third-party AI component assessment — providing an organizational structure for the AI-BOM and targeted vulnerability watch processes recommended above. Organizations should map their AI infrastructure inventory against AICM's supply chain controls to identify coverage gaps.

The CSA MAESTRO framework for agentic AI threat modeling explicitly models supply chain threats at the deployment and runtime layers, covering the class of risks represented by LangChain serialization flaws and MCP server SSRF vulnerabilities [19]. Organizations deploying agentic AI systems should apply MAESTRO's threat modeling methodology to assess whether unpatched AI toolchain vulnerabilities create exploitable paths within agent execution environments, particularly where agents have access to external APIs, file systems, or privileged credentials.

The Cloud Controls Matrix (CCM) supply chain management domain provides additional governance structure for third-party software risk, which maps directly to the need for AI-specific vendor security assessment and software integrity verification controls. Organizations seeking audit-ready evidence packages for AI vulnerability governance should use CCM as the control mapping layer and AICM as the AI-specific supplementary framework.

CSA Labs published a related research note in April 2026 addressing how enterprise vulnerability programs should adapt their general processes to the NVD enrichment policy change [20]. The present note extends that guidance specifically to the AI infrastructure category, where the structural mismatch between NVD's triage criteria and the actual risk landscape is most acute.

References

[1] NIST. "NIST Updates NVD Operations to Address Record CVE Growth." NIST, April 2026.

[2] Cybersecurity Dive. "NIST limits vulnerability analysis as CVE backlog swells." Cybersecurity Dive, April 2026.

[3] Help Net Security. "NIST admits defeat on NVD backlog, will enrich only highest-risk CVEs going forward." Help Net Security, April 16, 2026.

[4] Black Duck. "NVD Changes 2026: NIST Vulnerability Database Shift & BDSA Solution." Black Duck Blog, 2026.

[5] Axonius. "NIST's NVD Enrichment Gap Is Now Permanent – What It Means." Axonius Blog, 2026.

[6] Socket. "NIST Officially Stops Enriching Most CVEs." Socket Blog, 2026.

[7] Medium / Egyda-AI. "PyTorch Security Crisis: How a Critical Vulnerability Threatens the AI Ecosystem." Medium, April 2025.

[8] The Hacker News. "Critical LangChain Core Vulnerability Exposes Secrets via Serialization Injection." The Hacker News, December 2025.

[9] PointGuard AI. "Microsoft MCP Server Vulnerability (CVE-2026-26118)." PointGuard AI, 2026.

[10] The Record. "Amid funding cuts, backlog of unanalyzed vulnerabilities in gov't database is growing." Recorded Future News, 2024.

[11] The Record. "NIST to limit work on CVE entries as submissions surge." Recorded Future News, 2026.

[12] Infosecurity Magazine. "NIST Drops NVD Enrichment for Pre-March 2026 Vulnerabilities." Infosecurity Magazine, 2026.

[13] The Hacker News. "LangChain, LangGraph Flaws Expose Files, Secrets, Databases in Widely Used AI Frameworks." The Hacker News, March 2026.

[14] Cloudsmith. "The 2026 Guide to Software Supply Chain Security." Cloudsmith Blog, 2026.

[15] Penligent AI. "NIST CVE Prioritization as AI Speeds Up Vulnerability Discovery." Penligent AI, 2026.

[16] Cloud Security Alliance. "AI Controls Matrix." CSA, 2025.

[17] NIST. "AI Risk Management Framework." NIST, January 2023.

[18] Tenable. "NIST Scales Back NVD CVE Enrichment: What to Know." Tenable Blog, 2026.

[19] Cloud Security Alliance. "Agentic AI Threat Modeling Framework: MAESTRO." CSA, February 2025.

[20] CSA Labs. "NVD Enrichment Triage: Enterprise Vulnerability Programs Must Adapt." CSA Lab Space, April 19, 2026.

[21] JFrog. "JFrog Enables Trusted AI – Uncovers Critical Security Threats Emerging from AI's Expansion in the Software Supply Chain." JFrog, April 2025.

[22] Protect AI / Hugging Face. "4M Models Scanned: Protect AI + Hugging Face 6 Months In." Hugging Face Blog, April 2025.

[23] IBM Security. "IBM Report: 13% of Organizations Reported Breaches of AI Models or Applications, 97% of Which Reported Lacking Proper AI Access Controls." IBM Newsroom, July 30, 2025.

[24] ReversingLabs. "2026 Software Supply Chain Security Report." ReversingLabs, January 2026.