CSAI Foundation

Cloud Security Alliance AI Safety Initiative

NVD Enrichment Triage: Guidance for AI Security Programs

Adapting Vulnerability Management After NIST's April 2026 Shift to Selective CVE Enrichment

Unofficial AI-assisted Research

2026-04-22

Key Takeaways

- On April 15, 2026, NIST formally transitioned the National Vulnerability Database to a risk-based enrichment model, concentrating analyst resources on CVEs in the CISA Known Exploited Vulnerabilities (KEV) catalog, federal government software, and critical software defined by Executive Order 14028 [1].
- All backlogged CVEs with an NVD publish date earlier than March 1, 2026—approximately 29,000 entries—have been reclassified as "Not Scheduled" with no committed enrichment timeline, creating an immediate blind spot for scanner configurations that depend on NVD-derived metadata [3].
- Based on the scope of NIST's three priority categories relative to total CVE volume, an estimated 80–85% of new CVE submissions will proceed without NIST-provided CVSS scores, CPE data, or CWE classifications—a proportion the authors derive from the policy's own criteria definitions.
- AI infrastructure dependencies—ML serving frameworks, model registries, inference libraries, and AI agent tooling—fall almost entirely outside the three priority categories and are therefore disproportionately likely to remain unenriched, exposing AI-native organizations to coverage gaps the vulnerability management tooling will not automatically surface [3].
- Research on CVSS scoring reliability has found that more than 40% of CVEs receive materially different scores when re-evaluated by the same analyst over time, illustrating the inherent subjectivity in severity assessment; documented examples include score shifts of 6.5 CVSS points—enough to cross severity-tier boundaries that drive patching urgency [4].
- AI security programs must immediately audit their tooling for NVD dependency, diversify enrichment sources to include GHSA, OSV, CISA Vulnrichment, and commercial threat-intelligence feeds, and rebuild prioritization logic around exploitability signals rather than CVSS score alone [5][6].

Background

The NVD's Central Role in Vulnerability Management

The National Vulnerability Database has served as the authoritative enrichment layer for the CVE identifier system since 2005. When a CVE Numbering Authority publishes a new CVE record, the raw entry typically includes only a brief description and a list of references. NIST's NVD team adds the structured metadata that makes the record operationally useful: Common Platform Enumeration (CPE) identifiers that allow scanners to match a CVE to affected product versions, a CVSS base score and vector string that encodes severity, and a Common Weakness Enumeration (CWE) classification that categorizes the underlying flaw type. Most commercial vulnerability scanners, application security posture management (ASPM) platforms, and software composition analysis (SCA) tools ingest NVD data as a primary enrichment source. This architecture placed substantial operational weight on a single government-maintained database, and that weight has now exceeded what NIST can sustain.

A System Under Structural Overload

The volume of CVE submissions has grown far beyond the capacity of any centralized team to manually enrich at the rate of publication. Between 2020 and 2025, CVE submissions increased by 263%, reaching 44,509 disclosures in 2025 alone, of which 14,593 had publicly available exploits at the time of publication [6]. NIST responded with substantially increased throughput: the agency enriched nearly 42,000 CVEs in 2025, approximately 45% more than any prior year [1]. Despite that record-setting output, the volume of new submissions kept pace. In the first quarter of 2026, submissions were running approximately one-third higher than the same period in 2025 [1], with the full year projected to reach 50,000–70,000 CVEs if first-quarter growth rates persist (authors' estimate based on [1]).

Part of what may be driving the acceleration is the emergence of AI-assisted vulnerability discovery. Tools capable of autonomous code analysis can identify classes of bugs that previously required days of expert researcher time, and they can operate in parallel across large codebases at a fraction of the cost. The consequence is a research productivity multiplier that flows directly into CVE volume: a bug-class that once yielded a handful of individual disclosures per year may now yield dozens as AI tooling systematically sweeps an ecosystem. This dynamic—AI enabling faster discovery, which floods the disclosure pipeline, which overwhelms a human-staffed enrichment team—is structural in nature and unlikely to be resolved through staffing increases alone [7].

NIST's April 2026 Decision

Effective April 15, 2026, NIST announced a formal transition from universal coverage to a risk-based triage model [1]. Under the new framework, the agency will concentrate enrichment resources on three categories of vulnerabilities: CVEs listed in CISA's Known Exploited Vulnerabilities catalog (with a target enrichment time of one business day), CVEs affecting software procured or operated by the U.S. federal government, and CVEs in software meeting the definition of "critical software" established by Executive Order 14028—a category that includes operating systems, web browsers, identity and access management systems, network security tools, and hypervisors. CVEs not qualifying under any of these criteria are categorized as "Lowest Priority – not scheduled for immediate enrichment." They remain published in the NVD, but without CPE strings, NIST CVSS scores, or CWE classifications. Organizations may request enrichment of specific CVEs by contacting NVD directly, though response timelines are not guaranteed [1]. All previously backlogged CVEs with NVD publish dates before March 1, 2026, have been moved to a "Not Scheduled" status without a committed resolution date.

Security Analysis

The Scope and Shape of the New Gap

The practical effect of the prioritization threshold is that a large majority of new CVE volume will proceed without NIST enrichment—an outcome that is not evenly distributed across the software landscape. Vendors who hold CNA status—Microsoft, Google, Adobe, Oracle, and other major commercial software publishers—will continue to provide their own CVSS scores and supporting metadata when they publish CVEs for their own products. Enterprises that rely primarily on commercial off-the-shelf software from these large vendors will experience less operational disruption. The organizations most exposed are those whose software supply chains extend into the open-source ecosystem, particularly in categories that fall outside federal procurement patterns and EO 14028's critical software definitions.

The scanner blind-spot risk is more acute than it may appear at first glance. Most vulnerability scanners perform asset-to-CVE matching through CPE identifiers: the scanner queries an asset for its software inventory, generates CPE strings from that inventory, and looks up matching CVE records in NVD. When a CVE lacks CPE data, the scanner cannot complete the match and produces no alert—even if the vulnerable software is present on the asset. Historical analysis of NVD enrichment coverage following the 2024 enrichment slowdown documented substantial gaps across the CVE corpus [8], and the new policy

will extend that dynamic indefinitely for non-priority submissions. Organizations that interpret a clean scanner result as confirmation of a clean environment are now at greater risk of that assumption being incorrect.

AI Infrastructure in the Blind Spot

AI security programs face a compounded version of this problem. For most organizations, the software stack underlying AI-native applications—ML serving frameworks such as SGLang, vLLM, and Triton Inference Server; model packaging and distribution tooling such as Hugging Face Hub clients and ONNX Runtime; agentic workflow libraries; vector database clients; and the Python package ecosystem that underlies all of these—falls outside NIST's three priority criteria. These libraries are typically not classified as federal government software in the sense relevant to the NVD enrichment model, nor as EO 14028 critical software, and while some CVEs in this category may eventually make CISA's KEV catalog following active exploitation, that designation arrives after exploitation has already occurred in the wild.

This creates an asymmetry that is particularly dangerous for AI security programs. The teams most likely to be operating novel software stacks with thin commercial support—and therefore the most likely to face unenriched CVEs in their dependency chains—are also the teams whose security posture is least likely to be captured by existing scanner configurations optimized for traditional enterprise environments. An organization running a production AI inference tier on SGLang, for example, will need to track CVEs in that framework through a combination of GitHub Security Advisories, direct monitoring of the project's security advisories, and manual enrichment review, because NVD-based scanning alone will not reliably surface those vulnerabilities going forward.

The AI tooling acceleration loop adds a further wrinkle: organizations that have deployed AI-assisted vulnerability discovery capabilities are themselves contributing to the CVE volume that has overwhelmed NIST's enrichment capacity. A security team using frontier AI models to conduct continuous code auditing of their own or their vendors' software may generate CVE submissions that, under the new model, enter the "Not Scheduled" pool unless they can be positioned as KEV-qualifying or federal-scope. Teams in this position should plan for the possibility that CVEs they help discover and disclose may require them to advocate for enrichment or provide their own enrichment data in coordination with the CNA.

CNA-Assigned Scores: A Reliability Caveat

Non-priority CVEs published after April 15, 2026, will carry only the CVSS score provided by the originating CNA—without a separate NIST review. This matters because CNA-assigned scores are not equivalent in quality to NVD-reviewed scores, a fact that is often overlooked in vulnerability

management workflows that treat any CVSS vector as a reliable severity signal. Research on CVSS scoring consistency has documented that more than 40% of CVEs receive materially different scores when re-evaluated by the same analyst over time, illustrating the inherent subjectivity built into the scoring methodology [4]. Documented examples include a curl vulnerability initially scored 9.8 that was later corrected to 3.3—a 6.5-point shift that moved the CVE from Critical to Low severity and would have produced materially different patching timelines had the original score stood [4]. The incentive structures for CNAs vary considerably: vendors with their own CNA status may—consciously or not—produce lower scores for commercial reasons, while independent researchers without deep CVSS expertise may miscalibrate scores in either direction.

For AI security programs that use CVSS as an input to automated prioritization pipelines—including AI-assisted patch prioritization, risk scoring in ASPM platforms, or SIEM alert triage—the practical recommendation is to treat CNA-only CVSS scores as starting estimates requiring supplemental verification rather than authoritative inputs. This is especially true for CVEs in AI infrastructure software where the affected product's CNA sophistication varies widely. Prioritization logic should be augmented with exploitability signals from sources that do not depend on NVD enrichment, most importantly the CISA KEV catalog and commercial exploit-intelligence feeds.

Recommendations

Immediate Actions

AI security programs should begin by auditing their vulnerability management tooling for NVD dependency. The relevant question is not only whether scanners are configured to pull from NVD, but whether the scanner's matching logic silently drops CVEs that lack CPE data. Some SCA and ASPM tools may suppress unenriched CVEs by default, and those CVEs will not appear in result sets or alert queues. Teams should request documentation from their tool vendors on how the product handles CVEs in "Not Scheduled" or unenriched status, and whether the product supplements NVD with alternative enrichment sources.

Organizations should immediately add the CISA KEV catalog as an independent prioritization input at the highest tier of urgency. KEV entries are the one category of CVE that NIST has committed to enriching within one business day [1], so they represent the most reliable overlap between NIST's new model and organizational prioritization. Any CVE on the KEV list that affects a component in the organization's environment should trigger immediate response regardless of CVSS score, since the presence of active exploitation is a stronger severity signal than any scoring-based proxy.

For AI infrastructure dependencies specifically, teams should establish direct monitoring of GitHub Security Advisories (GHSA) and the Open Source Vulnerabilities (OSV) database for every major framework in their stack. GHSA and OSV are maintained closer to the source of open-source software development and are often populated before NVD enrichment occurs even under normal conditions. Both databases publish machine-readable advisory feeds that can be integrated into CI/CD pipelines and ASPM tooling. This approach is particularly effective for ML frameworks such as PyTorch, Hugging Face Transformers, and ONNX Runtime, where the upstream maintainers and GitHub's security team are more likely to provide timely enrichment than NVD is, given that these libraries fall outside NVD's new priority criteria.

Short-Term Mitigations

Within 30 to 60 days, AI security programs should establish a diversified enrichment stack. A practical configuration integrates NVD for high-priority CVEs, GHSA and OSV for open-source dependencies, CISA Vulnrichment for CVEs the government team has independently enriched, and at least one commercial threat-intelligence feed for exploit status and exploitation trend data. Commercial feeds such as those offered by VulnCheck, Flashpoint, and Recorded Future provide enrichment that is independent of NVD and is not subject to the same resource constraints. Several of these providers maintain enrichment pipelines that aim to score and classify CVEs within hours of publication, faster than NVD's current operational tempo [9].

Security teams should also recalibrate their risk-scoring models to weight exploitability signals more heavily relative to raw CVSS base scores. The canonical exploitability signal is KEV membership, but several vulnerability intelligence providers publish exploit-prediction scores and exploitation-in-the-wild signals that can supplement KEV. Organizations using AI to assist in vulnerability triage should ensure the model's training data and prompt context accounts for the reduced reliability of CVSS scores from CNAs, and that the model is instructed to consider exploit evidence as a primary severity dimension rather than a secondary modifier.

For CVEs affecting AI infrastructure that appear in the "Not Scheduled" category and represent meaningful organizational risk, security teams can request enrichment from NIST via nvd@nist.gov with a justification describing the affected software's criticality to the organization's environment. While response timelines are uncertain, for a small number of high-priority unscheduled CVEs this pathway is worth pursuing. Teams should also monitor whether their AI infrastructure vendors hold CNA status, since CNA-published CVEs will at least include vendor-assigned severity data, even if the quality of that data is variable.

Strategic Considerations

The NVD enrichment transition reflects a structural limit of the centralized, human-staffed enrichment model given the current and projected CVE volume trajectory. AI security programs should anticipate that this model will continue to narrow, not expand, and plan their vulnerability management architectures accordingly. Long-term resilience requires treating NVD as one input among several rather than a source of record, and investing in tooling that can aggregate and reconcile enrichment from multiple upstream databases.

CSA's AI Safety Initiative is exploring AI-powered enrichment as a potential alternative—using frontier models to generate draft CPE strings, CVSS vectors, and CWE classifications for "Not Scheduled" CVEs, with human expert review. The AI Safety Initiative believes this approach represents a credible alternative enrichment model at a cost and volunteer capacity achievable within existing community structures. Organizations interested in contributing to or consuming from such a pipeline should engage with CSA's Vulnerability Data Working Group and the AICM working group.

At a governance level, CISOs and security architects should formally document the limitation of NVD-dependent controls in their risk registers and vendor assessments. STAR attestations and SOC 2 controls that reference NVD-sourced CVSS data should note the evolving nature of NVD coverage and identify the supplemental data sources the organization has adopted. This documentation matters not only for internal risk management but also for customer-facing assurance obligations where vulnerability management program completeness is audited.

CSA Resource Alignment

The NIST NVD transition reinforces priorities across several active CSA frameworks. The AI Controls Matrix (AICM) addresses vulnerability management requirements for AI-native systems, and the NVD enrichment gap is a concrete implementation challenge for organizations seeking to satisfy AICM's Threat and Vulnerability Management domain. AICM-governed programs should evaluate their CVE data sources explicitly against the new NVD priority criteria, since the framework's intent—visibility into AI system vulnerabilities—may not fully be satisfied by NVD-only configurations given the current enrichment policy.

CSA's Cloud Controls Matrix (CCM) v4.1 Threat and Vulnerability Management domain requires organizations to maintain continuous and risk-based vulnerability management programs. The NVD enrichment gap does not relax those obligations; it raises the operational bar for satisfying them. Organizations implementing CCM TVM controls must now demonstrate that their scanner configurations

and enrichment sources cover their actual software inventory, including components that NVD will not actively analyze. Compensating controls that integrate GHSA, OSV, Vulnrichment, and commercial exploit-intelligence feeds should be documented as part of CCM compliance evidence.

The MAESTRO threat modeling framework, which addresses the seven layers of agentic AI system risk, has direct application here. Layer 4 of MAESTRO covers tool and integration security, and an AI agent that queries a vulnerability database to assist with security triage must account for the possibility that the database it queries returns incomplete data. AI agents used for vulnerability management tasks—automatic patch prioritization, SBOM analysis, security advisory summarization—should be designed to recognize and surface unenriched CVEs rather than treating the absence of CVSS data as a signal of low severity. The MAESTRO principle of trust boundary management—applied here to external data sources, not only agent-to-agent communications—is directly relevant to how these systems should be designed.

CSA's STAR continuous-monitoring program provides a vehicle for organizations to attest to the quality of their vulnerability management programs. Given the NVD changes, STAR assessors and the organizations they evaluate should revisit how vulnerability management controls are evidenced, particularly for cloud-native and AI-native software stacks whose dependencies are now systematically outside NIST's enrichment scope.

References

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

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

[3] Socket.dev. "NIST Officially Stops Enriching Most CVEs as Vulnerability Volume Surges." Socket.dev, April 2026.

[4] The Register. "Vulnerability scores, huh, what are they good for? Almost nothing." The Register, October 16, 2025.

[5] CyberScoop. "NIST narrows scope of CVE analysis to keep up with rising tide of vulnerabilities." CyberScoop, April 2026.

[6] Flashpoint. "National Vulnerability Database (NVD) Shifts to Selective Enrichment as CVE Volume Su rges." Flashpoint, April 2026.

[7] The Hacker News. "NIST Limits CVE Enrichment After 263% Surge in Vulnerability Submissions." The Hacker News, April 2026.

[8] VulnCheck. "Danger is Still Lurking in the NVD Backlog." VulnCheck, 2025.

[9] SecureWorld. "The NVD Course Correction: Navigating NIST's Strategic Pivot for 2026." SecureWorld, April 2026.

[10] Aikido Security. "Reliable CVE sources in the age of NIST NVD cutbacks." Aikido Security, April 2026.

