

CSA Whitepaper — NVD Infrastructure Crisis & AI Vulnerability Discovery

Contents

1 CSA Whitepaper — NVD Infrastructure Crisis & AI Vulnerability Discovery	1
1.1 Boardroom Citation Context	1
1.2 Source Location	1
1.3 Captured Content	2

1 CSA Whitepaper — NVD Infrastructure Crisis & AI Vulnerability Discovery

Field	Value
Reference key	csa-nvd
Availability	AVAILABLE
Capture method	primary
Source type	remote_url
URL / path	https://labs.cloudsecurityalliance.org/wp-content/uploads/2026/05/CSA_whitepaper_NVD_infrastructure_crisis_AI_vuln_csa-styled.pdf
Captured (UTC)	2026-06-23T03:31:04Z
Content type	application/pdf
HTTP status	200

1.1 Boardroom Citation Context

GOV-10 NVD backlog; P7-08 redundant feeds.

1.2 Source Location

https://labs.cloudsecurityalliance.org/wp-content/uploads/2026/05/CSA_whitepaper_NVD_infrastructure_crisis_csa-styled.pdf

1.3 Captured Content

CSAI Foundation | Cloud Security Alliance

The NVD Infrastructure Crisis: AI
Discovery Overwhelms Tracking
How Autonomous Vulnerability Research Is Straining Global
Tracking Infrastructure and What Security Programs Must Do Now

2026-05-04

2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright information remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with appropriate attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval.

2026 Cloud Security Alliance. AI-assisted Rapid Research.

Table of Contents	
Executive Summary	4
Introduction: A Cornerstone Under Pressure	5
The Long Build: How the CVE Surge Developed	6
The Accelerant: AI-Powered Vulnerability Discovery	7
The April 2026 Inflection: What Changed and What It Means	8
Enterprise Consequences: Blind Spots in Vulnerability Programs	9
Fragmentation and Opportunity: The Emerging Alternative Ecosystem	11
The Global Governance Dimension	12
Conclusions and Recommendations	13
Conclusions	
Recommendations	
CSA Resource Alignment	16
References	18

Executive Summary

For more than twenty years, the National Vulnerability Database has been the silent backbone of vulnerability management. Every time a vulnerability scanner flags a finding, every time a compliance platform assigns a severity tier, every time a patch management tool prioritizes its queue---the identifiers, CVSS scores, and CWE classifications produced by NIST's small NVD team are doing the underlying work. That infrastructure is now in crisis.

Between 2020 and 2025, CVE submissions to the NVD grew by 263 percent, with 48,185 vulnerabilities published in 2025 alone [1]. The first quarter of 2026 arrived running approximately one-third higher than the same period a year before [1]. NIST's response---enriching nearly 42,000 CVEs in 2025, 45 percent more than any prior year---was the most ambitious enrichment effort in the database's history [1]. It was not enough. On April 15, 2026, NIST formally changed course, announcing a risk-based triage policy that would route full enrichment only to vulnerabilities in CISA's Known Exploited Vulnerabilities catalog, those identified by the federal government, and critical software defined under Executive Order 14028 [2]. NIST estimates that that filter covers 15 to 20 percent of anticipated CVE volume, leaving the remaining 80 to 85 percent as unenriched shells [2] [20].

Artificial intelligence is not the sole cause of the surge, but it is the accelerant that has converted a manageable backlog into a structural failure. Autonomous research systems are discovering vulnerabilities at throughputs that no manual enrichment program can match. Google's Big Sleep agent found 20 previously unknown vulnerabilities in widely used open-source projects by August 2025 [3]. AISLE, the autonomous cyber reasoning system described in CSA's January 2026 research coverage, replicated the full set of 15 CVEs in the OpenSSL coordinated release---including bugs dating back decades---against one of the most heavily audited codebases on the internet [4]. Anthropic's Claude Mythos and Project Glasswing, discussed in CSA's April 2026 emergency strategy briefing co-produced with SANS Institute, OWASP, and MITRE, [un]prompted, demonstrated the operational scale of AI-driven exploitation: Claude Mythos generated 100 working exploits targeting Firefox vulnerabilities, while Claude Opus 4.6 independently identified 100 high-severity vulnerabilities across open-source software projects, including flaws that had previously gone undetected for years [5]. The discovery side of the vulnerability lifecycle has been fundamentally transformed, while the tracking and enrichment infrastructure has not.

The consequences arrive across the full spectrum of security operations. A CVE without a CPE identifier is effectively invisible to most vulnerability scanners, which match software inventory against NVD identifiers and CPE strings to generate alerts. A CVE without a CVSS score cannot be placed in a risk-tiered remediation queue. When only 32 percent of CVE-2025 entries had been fully enriched by early 2026 [11], the surface that security programs thought they were measuring was already materially incomplete. The coverage problem is about to become structurally permanent under the new triage policy.

This whitepaper traces the forces behind the NVD's crisis, maps the enterprise security consequences, describes the emerging multi-source ecosystem that is developing in the gap, and provides actionable recommendations for security organizations navigating this transition.

Introduction: A Cornerstone Under Pressure

The National Vulnerability Database was established in 2005 to serve as the authoritative reference for the Common Vulnerabilities and Exposures program that MITRE had been running since 1999. While MITRE's CVE program produces identifiers and descriptions, NIST's NVD adds the structured metadata that makes those identifiers operationally useful: CVSS scores to quantify severity, Common Platform Enumeration identifiers to specify which software versions are affected, and Common Weakness Enumeration classifications to categorize the underlying flaw. Together, these data points allow security tools to translate a CVE identifier into an actionable security finding--matching it against internal policies, assigning a priority, and routing it to the appropriate remediation workflow.

The dependency on this infrastructure runs deep. Vulnerability scanners from every major vendor--Tenable, Rapid7, Wiz, and their competitors---use CPE data to determine whether a given CVE affects software present in the scanned environment. Compliance frameworks including the Payment Card Industry Data Security Standard, HIPAA Security Rule guidance, and FedRAMP continuous monitoring requirements reference CVSS scores for severity thresholds and remediation timelines. Software composition analysis tools depend on NVD enrichment to flag vulnerable dependencies in application builds. Security incident and event management platforms ingest NVD data to contextualize threat feeds and calculate risk scores. The database is so deeply embedded in security tooling that many organizations do not realize they are dependent on it until it stops working.

The crisis building inside NVD since early 2024 represents the collision of three converging pressures: a secular increase in vulnerability discovery volume driven partly by the growth of software complexity, partly by the democratization of security research tools, a sudden acceleration of that discovery attributable directly to AI-powered research capabilities, and a funding and staffing model for NVD designed for a world with far fewer CVEs to enrich. Those pressures have now produced a system that formally acknowledges it cannot do the job it was built to do.

2026 Cloud Security Alliance. AI-assisted Rapid Research.

The Long Build: How the CVE Surge Developed

The growth in CVE volume did not begin with AI. The Common Vulnerabilities and Exposures program expanded its network of CVE Numbering Authorities throughout the 2010s, extending the ability to generate CVE identifiers to hundreds of organizations including major software vendors, security research firms, and government agencies. That expansion was deliberate and appropriate: it distributed the work of vulnerability identification and reduced bottlenecks at MITRE while improving coverage for software outside major enterprise platforms. It also meant that the inbound volume to NVD grew with each new CNA added to the program.

Between 2017 and 2023, CVE publication rates increased steadily but within a range that NIST could manage with effort, track. The 2024 inflection was different in character. CVE submissions began outpacing e-

capacity in the first quarter of that year, and NIST acknowledged a developing backlog. By mid-2025, the unprocessed queue had grown large enough to generate industry concern. Security researchers began documenting CVEs that had sat without CVSS scores or CPE data for weeks or months, rendering them invisible to automated scanning tools for the duration of that delay. NIST worked to address the backlog by hiring contractors and adjusting workflows, but the submissions kept arriving faster than the existing team could process them.

The 2025 data tells the story in numbers. NIST published approximately 48,185 CVEs in 2025, representing a 263 percent increase over 2020 [1]. To put that growth in operational terms: in 2020, an NVD team of 100 professionals needed to enrich on the order of 18,000 entries per year, roughly 350 per week. In 2025, maintaining pace with submissions would have required enriching nearly 1,000 entries per week---a workload with the same team size [1]. NIST responded by expanding contractor support and optimizing enrichment processes, achieving what it described as its most productive enrichment year on record, but the gap was not close. As 2026 opened, submissions arrived at a pace approximately one-third above the already high rate of 2025 [1].

The staffing and funding environment compounded the operational problem. NIST experienced roughly a ten percent reduction in its federal funding in 2024, prompting voluntary departures and limiting its enrichment capacity [8]. Requests for supplemental appropriations to scale NVD operations received partial approval, but that did not close the gap. The agency found itself in the position of running faster to fall behind, and, slowly, an unsustainable posture that ultimately produced the April 2026 policy change.

The broader CVE program faced a parallel crisis in the spring of 2025. The contract under which Lockheed Martin Corporation managed the CVE program---funded by the Department of Homeland Security---expired on April 15, 2025, without immediate renewal, briefly raising the prospect of the entire CVE identification program ceasing operations [9]. CISA exercised an emergency contractual option to extend the program for approximately eleven additional months while longer-term arrangements were negotiated, averting an immediate collapse. The incident revealed how much of the world's vulnerability tracking infrastructure

2026 Cloud Security Alliance. AI-assisted Rapid Research.

The program had rested on a single contracting relationship between a US government agency and a nonprofit research organization. In direct response, the CVE Board announced the formation of the CVE Foundation to ensure governance independence from any single national sponsor [9]. The structural fragility exposed in the first forty-eight hours has not been fully resolved.

The Accelerant: AI-Powered Vulnerability Discovery

The secular growth in CVE volume provides the background condition for the NVD crisis. Artificial intelligence provides the accelerant that has transformed a serious operational problem into a systemic failure.

Prior to the widespread availability of AI-assisted research tools, large-scale vulnerability discovery was the domain of skilled human researchers, typically working on specific targets, applying specialized knowledge through years of experience. The economics constrained throughput. A talented security researcher could find a handful of novel vulnerabilities in a complex codebase in a quarter of sustained effort. Many programs and coordinated disclosure frameworks were calibrated to this human-speed discovery rate. Implicitly, was NIST's NVD enrichment capacity.

AI systems have broken that calibration. Google's Big Sleep project, developed jointly by DeepMind and Project Zero, had identified 20 previously unknown vulnerabilities in widely used open-source software in August 2025, including an exploitable stack buffer underflow in SQLite that represented, in Google's description, the first time an AI agent had been used to foil active exploitation of a vulnerability [3]. The system found and reproduced each vulnerability without human intervention, with human effort brought in only for the disclosure phase [3]. Parallel work on FFmpeg, ImageMagick, and other foundational open-source libraries demonstrated that AI-powered discovery was not limited to narrow domains of code analysis.

The AISLE autonomous cyber reasoning system, subject of CSA's January 2026 research coverage, demonstrated a capability of a different order: reproducing the full set of 12 CVEs disclosed in the 2026 OpenSSL coordinated release, including historical vulnerabilities that had persisted undetected for years in one of the most heavily reviewed codebases in existence [4]. The significance of that result is what it implies about the discovery backlog. If AI systems can find previously unknown vulnerabilities in heavily audited code at this rate, the volume of CVEs yet to be discovered across the much larger, less carefully reviewed universe of software is difficult to bound.

Anthropic's Claude Mythos and Project Glasswing, described in CSA's April 2026 emergency strategy briefing co-produced with SANS Institute, OWASP, and [un]prompted, demonstrated the operational reality of AI-driven exploitation: Claude Mythos generated 181 working exploits targeting Firefox vulnerabilities.

2026 Cloud Security Alliance. AI-assisted Rapid Research.

Claude Opus 4.6 independently identified over 500 high-severity vulnerabilities across open-source software projects, including flaws that had persisted undetected for years [5]. CSA's companion report on AI vulnerability weaponization quantified the economic differential underlying these results, a 156-times cost reduction relative to equivalent human effort and execution speeds 3,600 times faster.

These are not isolated proof-of-concept demonstrations. They represent the early, publicly visible signs of a technology trend that researchers across the academic, government, and commercial sectors are actively developing. The AI tools available in 2026 are not the ceiling; they are the floor from which more capable successors will be built. Each generation of more capable research tooling will find more vulnerabilities, submit more CVEs, and deepen the mismatch between discovery velocity and enrichment capacity.

The feedback dynamic is worth examining explicitly. More powerful AI discovery tools find more vulnerabilities. More CVE submissions overwhelm NVD's enrichment capacity. More unenriched CVEs

that vulnerability management programs operate with incomplete data. Incomplete data means that vulnerabilities are never remediated not because defenders chose to deprioritize them, but because tracking infrastructure never surfaced them as findings. Meanwhile, offensive actors---state-sponsored groups and criminal enterprises---are applying equivalent AI tooling to find and weaponize vulnerabilities before they are even known to the defender community. The asymmetry is compounding.

The April 2026 Inflection: What Changed and What It Means

NIST's April 15, 2026 announcement described the new policy as a "risk-based approach to CVE enrichment" [2]. The language was accurate, but the operational reality it describes is a triage defaults to non-enrichment for most of the CVE population. Under the new policy, NIST will provide enrichment---CVSS scores, CPE identifiers, CWE classifications---only for three categories of vulnerabilities: those appearing in CISA's Known Exploited Vulnerabilities catalog, those affecting software in use by the federal government, and those affecting critical software as defined under Executive Order 14028. NIST estimates that these three categories will capture approximately 15 to 20 percent of anticipated CVE volume going forward [2][20].

The remaining 80 to 85 percent of CVEs will enter the database as shells: identifiers and descriptions without the structured metadata that operational tools require. All backlogged CVEs with an NVD publication date earlier than March 1, 2026---approximately 29,000 entries---have been reclassified as "Not Reviewed" meaning NIST does not currently intend to enrich them [11]. The backlog is not cleared; it is formally acknowledged as beyond current capacity to address.

2026 Cloud Security Alliance. AI-assisted Rapid Research.

Additional changes compound the operational impact. NIST announced that it will no longer routinely produce independent CVSS scores for CVEs that have already been scored by the issuing CVE Number Authority, instead accepting the CNA-assigned score as the database record [2]. This change shifts responsibility for severity scoring to an ecosystem of hundreds of CNAs with varying methodologies, incentives, and expertise. VulnCheck's analysis of dual-scored CVEs found disagreement between NIST and CNA CVSS scores in more than half of cases examined, with individual divergences sometimes large enough to move a vulnerability across severity tiers---from Medium to Critical, for example---depending on the score the security tool consumed [6]. Under the new policy, that inconsistency will not be audited or corrected by NIST.

The NVD will also no longer routinely re-analyze modified CVEs unless NIST is specifically informed of a modification that materially affects the enrichment data [2]. CVE records are frequently revised upon publication as affected version ranges are refined, additional affected products are identified, and assessments are updated. Under the previous model, NIST would review those modifications and update the NVD record accordingly. That quality-assurance function is now substantially curtailed.

The practical effect on security operations is immediate. A CVE without CPE identifiers is, from the perspective of most vulnerability scanners, invisible. These tools work by matching the software they observe in a scanned environment against the CPE strings in NVD records; where that mapping is absent, a vulnerable component generates no scanner alert for organizations relying exclusively on NVD enrichment [12]. A CVE without a CVSS score cannot be placed in a severity-tiered remediation queue. A combination of missing CPE and missing CVSS means that the vulnerable software is present in the environment, the CVE identifier exists in the database, and the security program's tooling nonetheless fails to surface the exposure.

Enterprise Consequences: Blind Spots in Vulnerability Management Programs

The enterprise security consequences of the NVD's triage policy extend well beyond the operational inconvenience of incomplete database records. They touch the fundamental validity of the vulnerability management programs that organizations have built on the assumption of NVD completeness.

Vulnerability management programs typically operate on a severity-tiered patching cadence: critical vulnerabilities within a defined window (commonly seven to thirty days), high-severity vulnerabilities in a longer window, medium and lower severity vulnerabilities on routine patch cycles. That cadence depends on having severity scores for the vulnerabilities under management. When a significant fraction of vulnerabilities are without CVSS scores---or with CNA-assigned scores of uncertain consistency---the tiering logic breaks down.

2026 Cloud Security Alliance. AI-assisted Rapid Research.

Organizations face a choice between treating all unenriched CVEs as unknown risks (which would overwhelm remediation capacity) or ignoring them until enrichment arrives (which creates unacknowledged exposure windows).

The compliance dimension adds further urgency. Many regulatory and contractual frameworks specify thresholds for remediation timing. PCI DSS 4.0 requires that vulnerabilities with CVSS scores of 7.0 and above be addressed within defined timeframes. FedRAMP continuous monitoring requires that High severity findings be remediated within thirty days, a determination that depends on CVSS scores. When CVEs lack NVD-sourced CVSS scores, the compliance mapping breaks. Organizations may either be unaware of qualifying vulnerabilities or may consume inconsistent CNA-assigned scores that produce different compliance outcomes depending on the data source their tools happen to use.

The software composition analysis use case is particularly exposed. Organizations managing open-source dependencies at scale use SCA tools to flag components with known CVEs; those tools depend heavily on NVD CPE data to match package names and version ranges to CVE records. The open-source ecosystem generates a large proportion of CVEs that will not qualify for NIST enrichment under the new triage policy: they will not appear in CISA's KEV catalog, they will not be federal government software, and they will not meet the EO 14028 critical software threshold---meaning that many CVEs affecting widely deployed

source packages will arrive in the NVD as identifier-only records that SCA tools cannot match to components.

The timing environment makes all of this more dangerous, not less. As CSA documented in "The Collapsing Exploit Window," the mean time to exploit a disclosed vulnerability has compressed dramatically. The 2026 Global Threat Landscape Report found that exploited high- and critical-severity vulnerabilities increased 105 percent year-over-year in 2025, from 71 instances to 146, while the median time between vulnerability publication and confirmed exploitation fell from 8.5 days to 5.0 days [13]. The mean time to exploit across the same population dropped from 61.0 days to 28.5 days [13]. In 32.1 percent of tracked exploit cases, exploitation occurred on or before the CVE's public disclosure date---meaning that in nearly a third of exploits, the public vulnerability record did not exist when the attack began.

Contrasted with these exploitation timelines is the enterprise patching reality. Qualys' 2026 enterprise patch benchmark found that the mean time to remediation for complex enterprise applications reached five months and ten days [21]. The gap between a median exploitation timeline of five days and a mean remediation timeline of five months is not a gap that can be managed through operational efficiencies; it requires that vulnerability intelligence be actionable within hours of CVE publication. An enterprise system that takes weeks or months to assign CVSS scores and CPE identifiers---or, under the new policy, may never do so---cannot support that operational tempo.

IoT, operational technology, and embedded systems environments face a particular dimension of this challenge. CVEs in these categories rarely appear in CISA's KEV catalog, whose additions are driven primarily by evidence of exploitation in enterprise and federal environments. They do not affect

2026 Cloud Security Alliance. AI-assisted Rapid Research.

Industrial control software, government software in the conventional sense, and they frequently do not qualify as EO 14028 covered software. Under the new NVD triage policy, a vulnerability in industrial control software, medical device firmware, or building automation systems may go indefinitely without CVSS or CPE enrichment, invisible to the scanning tools that would otherwise surface it for remediation. Security teams responsible for these environments were already operating with limited NVD coverage; that limitation is now formally pronounced [14].

Fragmentation and Opportunity: The Emerging Alternative Ecosystem

NIST's triage announcement confirmed what commercial vulnerability intelligence vendors had been anticipating for more than a year: the era of a single authoritative enrichment source for the entire population was over. The market had already begun fragmenting before April 15, 2026; the policy change accelerated the process.

VulnCheck had been operating a parallel enrichment pipeline for years, documented in its analysis.

that NVD provided CPE identifiers for only 41.35 percent of published vulnerabilities, compared to 100 percent coverage in VulnCheck's own enrichment service [6]. VulnCheck, which had already operated NVD++ as a free API-compatible alternative to the NVD API feed since early 2024, expanded promotion of the platform following NIST's April 2026 triage announcement [15]. VulnDB, operated by Risk Based Security (now part of Flashpoint), offers a commercially supported alternative with more detailed exploit, and patch information than NVD typically provides. GitHub Security Advisories provided an advisory publication for open-source projects whose maintainers participate. Google's OSV (Open Source Vulnerabilities) database offers version-range focused enrichment well-suited to software composition analysis use cases.

ENISA launched the European Union Vulnerability Database (EUVD) on May 13, 2025, accelerating its availability in direct response to the instability visible in the US CVE infrastructure [16]. The EUVD is designed as a downstream repository---Europe's counterpart to NVD---rather than a replacement for the NVD identifier program, but its existence changes the governance landscape meaningfully. ENISA was formally designated as a CVE Program Root in November 2025, giving the EU agency the authority to assign CVE identifiers directly and to serve as a coordination hub for European national CSIRTs participating in coordinated disclosure [17]. The EUVD and ENISA's CVE Root status together create an independent enrichment path for vulnerabilities of particular relevance to European organizations and regulatory compliance, separate from the NIST pipeline.

2026 Cloud Security Alliance. AI-assisted Rapid Research.

The proliferation of enrichment sources resolves some of the operational problems created by NVD's reduced scope, but it introduces new ones. Organizations previously consuming a single NVD API feed now decide which sources to integrate, how to handle conflicts between sources, and which source to use for remediation decisions when scores disagree. VulnCheck's own research has documented CVSS score disagreements between NVD and CNAs in more than half of dual-scored CVEs [6], and the fragmentation of the enrichment ecosystem will only increase the frequency and magnitude of such disagreements. A vulnerability flagged as Critical by one enrichment source and Medium by another places the remediation decision in the hands of whoever wrote the organization's vulnerability management policy---which is presumably not written with inter-source arbitration in mind.

Commercial security vendors including Qualys, Tenable, Rapid7, Anchore, Mend, and Snyk have responded to the enrichment gap by building independent enrichment pipelines that supplement NVD with their own research, threat intelligence feeds, and CNA relationships [15]. For organizations using these products, operational disruption may be less severe than for those relying on raw NVD API feeds. The distribution matters: large enterprises with mature vulnerability management programs and tier-one security product contracts are better positioned to absorb the transition than mid-market organizations or those relying on open-source scanner tooling against the NVD API directly.

The NVD crisis is not occurring in isolation; it is one episode in a broader reconfiguration of infrastructure for vulnerability identification, enrichment, and coordination that has been under way at least since 2024.

The CVE program's near-collapse in April 2025 and the formation of the CVE Foundation in its aftermath illustrated the tension between the program's practical status as global infrastructure and its legal structure as a US government contractor arrangement [9]. The CVE Foundation's mandate is to develop governance structures that insulate the CVE program from the funding and political decisions of its national government sponsor---a reasonable ambition given the scale of global dependence on CVE identifiers, but one that requires coordinating the interests of hundreds of CNAs, national CSIRs, and commercial stakeholders across many jurisdictions.

ENISA's emergence as a parallel enrichment authority adds a constructive competitive pressure but also an interoperability challenge. The EUVD and NVD will not always produce identical scores, CPE mappings, or CWE classifications for the same CVE. European organizations subject to NIS2 compliance requirements will increasingly reference EUVD as their authoritative source; US federal contractors will reference

2026 Cloud Security Alliance. AI-assisted Rapid Research.

enterprises operating across jurisdictions will need to manage both, along with the divergences between them. Security tooling vendors will need to decide which enrichment sources to integrate and how to represent disagreements to end users.

At the identification layer, the CNA ecosystem's continued expansion changes the distribution of data quality. CNAs range from major software vendors with dedicated product security teams and rigorous disclosure processes to small organizations with part-time security programs. The CVSS scores, advisory version ranges, and CWE classifications produced by these organizations vary substantially in accuracy and completeness. Under NIST's previous enrichment model, NVD review provided a quality-assurance layer that caught and corrected many CNA errors. Under the new model, CNA-assigned scores pass directly into the database. Organizations consuming CVE data will increasingly be consuming it at a quality level determined by the issuing CNA rather than a centralized review standard.

CISA's KEV catalog has taken on increased importance in this environment, functioning as a curated confidence signal within a noisier database landscape. The catalog's 20 percent growth in 2025 to 1,000 confirmed exploitation entries, with 20.5 percent linked to ransomware groups [18], provides a default zero priority queue for organizations that need an unambiguous starting point for remediation efforts. The catalog's limitation is its coverage: it documents exploitation that has been observed and reported, but not exploitation that is occurring in less visible environments or that will occur. A vulnerability that is not in the KEV catalog is not safe; it is merely not yet confirmed as exploited. With 85 percent of CVEs not in NIST's routine enrichment scope, the KEV catalog becomes a necessary-but-insufficient substitute for comprehensive vulnerability intelligence.

Conclusions and Recommendations

Conclusions

The NVD crisis is not a temporary operational setback that will resolve when NIST secures additional appropriations. It is a structural mismatch between the pace of AI-accelerated vulnerability discovery and the capacity of any centralized enrichment program operating at current staffing and funding levels. Even if NIST received a threefold increase in its NVD budget tomorrow, the trajectory of AI discovery tools and capabilities suggests that vulnerability submission rates will continue to grow faster than any enrichment staff can address. The April 2026 triage policy is the system's acknowledgment of a new equilibrium---one in which comprehensive enrichment of the full CVE population by a single authoritative source is no longer operationally feasible.

2026 Cloud Security Alliance. AI-assisted Rapid Research.

For security organizations, the combination of an enrichment gap affecting 80 to 85 percent of new CVEs and an exploitation timeline that has compressed to a median of five days for high-severity vulnerabilities represents a structural vulnerability in vulnerability management programs that were designed for a more predictable information environment. Programs that remain dependent on NVD as their sole or primary enrichment source will have decreasing visibility into the exposure surface they are responsible for managing.

The fragmentation of the enrichment ecosystem into multiple commercial and governmental sources is a problem to be solved; it is the new normal to be managed. Multi-source enrichment, explicit source arbitration policies, and threat-intelligence-driven triage are not optional enhancements for mature vulnerability management programs---they are baseline requirements for programs that need to maintain meaningful coverage of the exposure surface.

Recommendations

Immediate Actions

Organizations should audit their vulnerability management tooling to understand where each product source is getting its CVE enrichment data. Tools relying exclusively on the NVD API feed will experience increased coverage gaps under the new triage policy. This audit should document each scanner, SCA tool, commercial platform, and SIEM that ingests NVD data and identify which products have independent enrichment capabilities.

Organizations should integrate CISA's Known Exploited Vulnerabilities catalog as a first-tier priority signal for remediation triage, independent of CVSS score availability. The KEV catalog represents the highest confidence signal available for CVEs requiring immediate attention. All CVE-identified vulnerabilities appearing in the KEV catalog should enter expedited remediation queues regardless of whether NVD enrichment data is available.

Security teams should establish written procedures for CVEs that arrive without CVSS scores or C identifiers. Those procedures should define a default severity assumption for unenriched CVEs (conservative programs may choose to treat unknown severity as High by default), an escalation path for unenriched CVEs in software known to be externally accessible, and a review trigger for when enrichment arrives that may change the initial triage decision.

Short-Term Adaptations

Organizations should adopt at least one alternative enrichment source to supplement NVD, selected based on their primary technology environments. VulnCheck NVD++ provides API-compatible enrichment at no cost for community tiers. VulnDB offers more detailed timeline and exploit context at commercial rates. Organizations with significant open-source software dependencies should additionally integrate GitHub

2026 Cloud Security Alliance. AI-assisted Rapid Research.

OSV database, which is particularly well-suited to package ecosystem vulnerability tracking. European organizations subject to NIS2 should monitor ENISA's EUVD as a complementary reference for EU-related disclosures.

Vulnerability management policy documentation should be revised to acknowledge the changed NVD enrichment model explicitly. Updated policies should specify the organization's chosen primary and secondary enrichment sources, the methodology for resolving score disagreements between sources, the acceptable delay between CVE publication and remediation triage in the absence of enrichment. Without explicit policy, ad-hoc decisions by individual team members will produce inconsistent outcomes across the vulnerability population.

Software composition analysis tooling configurations should be reviewed to ensure that alternative enrichment sources are enabled and that CVEs in open-source packages are surfaced even when NVD coverage is absent. Organizations using open-source scanner tooling against the raw NVD API should evaluate migration to commercial platforms with independent enrichment pipelines, or supplement with VulnCheck NVD++ as an interim measure.

Strategic Considerations

Organizations with sufficient technical capacity should establish internal threat-intelligence-enabled vulnerability prioritization workflows that go beyond CVSS scoring to incorporate exploitation evidence, public exploit availability, asset criticality, and network exposure context. CVSS scores were an imperfect proxy for exploitability; the NVD crisis makes the limitations of CVSS-only triage impossible to ignore. Frameworks such as EPSS (Exploit Prediction Scoring System) and the combination of KEV scores with threat-feed-backed exploitation evidence provide more operationally relevant priority signals.

The NVD crisis should prompt a broader reconsideration of vulnerability management program maturity. Programs that were built on the assumption of timely, complete, centralized enrichment---essentially

them---need a maturity upgrade that assumes distributed, inconsistent, and incomplete enrichment builds compensating controls around that assumption. This is not a matter of incremental improvement, but a reorientation of the program's operating model.

Organizations that operate software platforms large enough to qualify as CVE Numbering Authorities should evaluate CNA status as a way to control the quality and timeliness of enrichment for vulnerabilities in their own products. As NIST's quality-assurance function for CNA-submitted scores diminishes, CNAs' ability to produce accurate, complete, and timely enrichment data become relatively more trustworthy data sources for the ecosystem.

At the industry and policy level, the CSA joins other security organizations in advocating for sustained investment in vulnerability tracking infrastructure proportionate to the scale of the problem. The CSA's operating budget and staffing should reflect its status as critical infrastructure for the global security ecosystem, not as a line item in a federal research agency's information technology budget. The CSA's

2026 Cloud Security Alliance. AI-assisted Rapid Research.

The CSA's governance work deserves support from the international security community as a mechanism for insulating the CVE identifier program from single-point-of-failure funding risks. ENISA's EU-wide effort represents a valuable diversification of global enrichment capacity that should be encouraged through technical interoperability standards with NVD and other databases.

CSA Resource Alignment

This whitepaper addresses structural challenges in vulnerability management infrastructure that align with multiple areas of CSA's framework ecosystem. Security practitioners applying CSA guidance to address the challenges described here should consider the following resources.

The CSA AI Controls Matrix (AICM), released in July 2025, provides a comprehensive 18-domain control framework for AI systems, including domains directly relevant to the AI-accelerated discovery domains described here. The Threat and Vulnerability Management domain within AICM addresses proactive identification and mitigation of AI-related risks, including risks arising from AI systems used in security operations. Organizations deploying AI-powered vulnerability scanners or research tools should adopt controls governing model security, output validation, and human-in-the-loop review before disclosure to ensure that AI-discovered vulnerabilities are handled responsibly and accurately [19].

The MAESTRO threat modeling framework for agentic AI systems provides structure for assessing risks from autonomous vulnerability discovery and exploitation tools. The seven-layer threat model is applicable to AI systems used defensively (automated vulnerability research, AI-assisted triage) and to the risks posed by adversaries using equivalent tools. MAESTRO's analysis of how AI agents can be misused to compress exploitation timelines directly informs the urgency of the patching and triage acceleration recommendations in this paper.

The CSA Cloud Controls Matrix (CCM) and its AI superset, the AICM, provide control mappings for vulnerability management practices applicable to cloud-hosted software. Organizations using CCM compliance self-assessment should note that CCM control VM-06 and related vulnerability management controls assume timely access to severity-scored CVE data; the NVD triage change creates a compliance risk for organizations that have not updated their enrichment sources.

CSA's STAR for AI program, currently under development, will provide third-party certification against controls. Organizations that use AI-powered vulnerability discovery tools or AI-assisted security, should consider STAR for AI certification as a mechanism for demonstrating that their AI systems operate under appropriate governance and quality controls, including those relevant to vulnerability discovery and coordinated response.

2026 Cloud Security Alliance. AI-assisted Rapid Research.

Security practitioners seeking further context on the AI-acceleration dimension of the vulnerability management challenge should consult CSA's April 2026 whitepaper "The Collapsing Exploit Window: AI-Speed Vulnerability Weaponization" and the companion "AI Vulnerability Storm" emergency strategy briefing produced jointly with SANS Institute, OWASP, and [un]prompted [5][10]. The CSA research note "NVD Enrichment Triage: Enterprise Vulnerability Programs Must Adapt," published April 19, 2026, provides operational guidance for the immediate transition period [20].

2026 Cloud Security Alliance. AI-assisted Rapid Research.

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 forward." Help Net Security, April 16, 2026.
- [3] TechCrunch. "Google says its AI-based bug hunter found 20 security vulnerabilities." TechCrunch, August 4, 2025.
- [4] AISLE. "AISLE Discovered 12 out of 12 OpenSSL Vulnerabilities." AISLE Blog, January 2026. See also GlobeNewswire. "AISLE Researchers Identify 12 New Security Vulnerabilities in OpenSSL Using AI-Driven Discovery." January 27, 2026.
- [5] SANS Institute, Cloud Security Alliance, [un]prompted, and OWASP GenAI Security Project. "The Vulnerability Storm: Building a Mythos-Ready Security Program." April 14, 2026.
- [6] VulnCheck. "VulnCheck's Commitment to Expanding Access to Vulnerability Enrichment." VulnCheck Blog, 2026.

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

[8] Just Security. "NIST Can't Keep Up. The Whole Digital Ecosystem Will Soon Feel It.." Just Security, 2026.

[9] Krebs on Security. "Funding Expires for Key Cyber Vulnerability Database." Krebs on Security, 2026.

[10] Cloud Security Alliance. "The Collapsing Exploit Window: AI-Speed Vulnerability Weaponization at Cloud Security Labs, April 2026.

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

[12] VulnCheck. "Outpacing NIST NVD with VulnCheck NVD++." VulnCheck Blog, 2024.

[13] Rapid7. "2026 Global Threat Landscape Report Shows Exploited High and Critical-Severity Vulnerabilities Surged 105% as Attack Timelines Collapsed." Rapid7, March 2026.

[14] SecureIoT.house. "The Database That Tells You If Your Devices Are Safe Just Stopped Working Due to Vulnerabilities." SecureIoT, 2026.

2026 Cloud Security Alliance. AI-assisted Rapid Research.

[15] VulnCheck. "VulnCheck introduces VulnCheck NVD++ as a Reliable, High-Performance Alternative to NIST NVD 2.0 API." VulnCheck Press Release, 2024.

[16] InfoQ. "Goodbye CVE? European Vulnerability Database EUVD Now Live." InfoQ, June 2025.

[17] ENISA. "Stepping up our role in Vulnerability Management: ENISA Becomes CVE Root." ENISA, November 2025.

[18] SecurityWeek. "CISA KEV Catalog Expanded 20% in 2025, Topping 1,480 Entries." SecurityWeek, 2025.

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

[20] Cloud Security Alliance AI Safety Initiative. "NVD Enrichment Triage: Enterprise Vulnerability Management Must Adapt." CSA Labs, April 19, 2026.

[21] Qualys. "Enterprise Patch & Remediation Benchmark 2026." Qualys Blog, April 20, 2026.

2026 Cloud Security Alliance. AI-assisted Rapid Research.

Archived reference article for AICSR-STUDY-2026-001 footnote corpus.