Volume 3 · Issue 1 · 2026 · pp. 57–63

The Enrichment Gap: A Longitudinal Audit of the NVD Backlog and Its Propagation Through Downstream Vulnerability Feeds

Risto Vaarandi
ORCID: 0000-0001-7781-5863
Tallinn University of Technology, Tallinn, Estonia

Surafel Lemma Abebe
ORCID: 0000-0002-2137-8673
Addis Ababa University, Addis Ababa, Ethiopia

Correspondence: Risto Vaarandi · contact details held in the journal record

Publication history. Assigned by the publisher to Cybersecurity Advances, Volume 3, Issue 1 (2026). This publisher-authorized digital edition was released on 30 October 2026 from the supplied final 2026 manuscript. Review-stage notices were removed and journal typography, pagination, metadata, and accessibility were normalized; substantive scholarly content was preserved.

© 2026 the author(s). Published by AlphaTech Scholarly Press under the Creative Commons Attribution 4.0 International License (CC BY 4.0).
No DOI is displayed unless and until registration is independently verified.

Background: On 12 February 2024, the U.S. National Vulnerability Database slowed its enrichment of new CVE records, creating a backlog that grew past 18,000 unanalyzed vulnerabilities — at its peak, 93.4% of new CVEs went unanalyzed, and even 46.7% of CISA's Known Exploited Vulnerabilities added in that window waited months for analysis. Recovery reached pre-slowdown processing rates by March 2025, and NVD enriched approximately 42,000 CVEs in 2025 (+45% over any prior year), but CVE submissions grew 263% between 2020 and 2025, and in April 2026 NIST formally moved unenriched pre-March-2026 records to a "Not Scheduled" category, adopting prioritized enrichment (KEV-first within one business day). **Objective:** We audit the full 2024–2026 backlog record and analyze how the enrichment gap propagates through the downstream vulnerability ecosystem. **Methods:** Longitudinal documentary analysis of NIST/NVD operational announcements, third-party backlog measurements (VulnCheck), and the downstream data architecture (CVE List → NVD → ADP layer → consumer feeds); construction of a backlog-risk register. **Results:** The backlog's severity was concentrated in exactly the fields that downstream tools consume — CPE applicability, CWE classification, and NVD-CVSS scores — breaking asset matching, triage, and compliance scanning for the unenriched window. The community-layer response (CISA's ADP Vulnrichment from July 2024) demonstrated that parallel enrichment is feasible but arrived with coverage and quality variance. The 2026 prioritization policy converts a temporary gap into a permanent two-tier record. **Conclusions:** The enrichment gap is structural, not episodic: a triage-first enrichment policy plus 263% submission growth means the unenriched tier will persist and grow. Downstream consumers must diversify their enrichment inputs (ADP feeds, EPSS, exploit signals) and stop treating NVD CPE data as the sole asset-matching mechanism.

Keywords: National Vulnerability Database; NVD backlog; CVE enrichment; CPE; CWE; vulnerability feeds; ADP; Vulnrichment; digital trust

Contribution statement. This paper makes four contributions: (1) a complete, source-verified timeline of the enrichment-gap window and its recovery, consolidating official announcements and third-party measurements; (2) an architecture-level propagation analysis showing which downstream functions break when enrichment stalls; (3) a backlog-risk register mapped to mitigation mechanisms; and (4) the operational implications of the April 2026 prioritization policy for consumers of NVD data.

1. Introduction

The National Vulnerability Database occupies a peculiar position in the security ecosystem: it is authoritative (operated by NIST), heavily consumed (by scanners, asset-management tools, compliance processes, and every downstream feed that mirrors it), and — as 2024 demonstrated — fragile. On 12 February 2024, NVD announced it was "transitioning" its analysis workflow; enrichment of new CVEs slowed dramatically and then effectively paused. The consequences accumulated in public view: by 19 May 2024, 93.4% of CVEs submitted since the slowdown had not been analyzed; by 21 September 2024, 72.4% still lacked analysis — including

46.7% of the CISA KEV entries added in the window [5][6]. Scanners that matched assets to vulnerabilities through NVD's CPE applicability data silently lost coverage; triage processes that consumed CWE and NVD-CVSS scores starved.

NIST's response came in stages: a full analyst team was onboarded (June 2024); CISA's Authorized Data Publisher (ADP) "Vulnrichment" began supplying CVSS and CWE data in parallel from 3 July 2024 [6]; processing returned to pre-slowdown rates by March 2025 [3]; and NVD enriched approximately 42,000 CVEs in 2025 — 45% more than any prior year [2]. But the arithmetic defeated the recovery: CVE submissions grew 32% in 2024 and 263% between 2020 and 2025 [2][3]. On 15 April 2026, NIST formalized the outcome: enrichment became *prioritized* — KEV entries within one business day, federal software next — and all backlogged CVEs published before 1 March 2026 were moved to "Not Scheduled," to be considered as resources allow [2][3].

This paper audits that full record and asks the engineering question the policy announcements raise but do not answer: **what breaks downstream when the enrichment layer stalls, and what must consumers change?** The answer matters far beyond 2024 — the April 2026 policy makes selective enrichment a *permanent design*, meaning every downstream consumer of NVD data now operates against a two-tier record permanently.

1.1 Research questions

RQ1: What is the verified 2024–2026 timeline of the enrichment gap? **RQ2:** Through which downstream pathways does the gap propagate, and which functions break? **RQ3:** What risk register and mitigations follow for consumers of NVD data? **RQ4:** What does the April 2026 prioritization policy change for consumers?

1.2 Contributions

(1) The complete verified timeline, consolidating official announcements with third-party measurements (Section 3, Figure 1, Table 1) (RQ1). (2) The propagation analysis across the enrichment architecture (Section 4, Figure 2) (RQ2). (3) The backlog-risk register with mitigation coverage (Section 5, Figure 3, Table 3) (RQ3). (4) The consumer-facing implications of prioritized enrichment (Section 6) (RQ4).

2. Background

2.1 The enrichment pipeline and its choke point

The CVE record pipeline is layered. CVE Numbering Authorities author records — identifier, description, references — published to the CVE List, which flows into NVD hourly [6]. NVD then *enriches* each record: CPE applicability statements (which products and versions are affected), CWE classification, and NVD-assigned CVSS scores [3]. This enrichment layer is the choke point: it is produced by one organization's analyst team, against submission volume that grew 263% between 2020 and 2025 [2]. The 2024 crisis added a second production layer — CISA's ADP program, which as of July 2024 publishes "Vulnrichment" CVSS and CWE data directly onto CVE records [6] — and the April 2026 policy entrenched the layering: NVD enriches a prioritized subset, ADP and other sources supplement, and the remainder ships without analysis [2][3].

2.2 Why enrichment matters downstream

A CVE identifier without enrichment still exists — but consumes poorly. Asset-matching tools that map an organization's inventory to vulnerabilities via CPE match nothing when the CPE statements are absent; triage processes keyed to CWE and severity starve; compliance regimes that require analyzed records (including U.S. federal scanning mandates and, prospectively, EU CRA reporting workflows) treat unenriched records as incomplete [3][9]. The 2024 backlog therefore broke not the *existence* of vulnerability information but its *actionability* — and because the affected window's records remain in "Not Scheduled" status, the gap persists for those specific CVEs indefinitely [2].

2.3 Related work

Consistency and quality of vulnerability data have been studied at scale: the VIEM system (USENIX Security 2019) found that only 59.82% of CVE/report pairs strictly match standardized NVD entries [10], establishing that enrichment quality was contested even before the backlog. Third-party monitoring of the 2024–2025 crisis (VulnCheck's backlog analyses) supplied the quantitative record this paper consolidates [5]. The GHSA review-pipeline analysis (2026) documents how downstream aggregators time-stamp and review advisories from multiple sources [11]. This paper's contribution is the consolidated longitudinal audit plus the downstream propagation analysis.

3. Methodology

Timeline verification (RQ1). Every event in the 2024–2026 record was verified against primary announcements (NVD news pages [6], the NIST April 2026 policy statement [2], the NVD process documentation [3]) and cross-checked against third-party measurements where independent numbers exist (VulnCheck's September 2024 analysis [5]). Percentages are quoted from the named measurement source with its measurement window; no new measurement was performed.

Propagation analysis (RQ2). The downstream architecture was modeled from documented data flows: the CVE List → NVD enrichment → ADP supplement → public feeds (OSV, GHSA mirrors) → consumer tooling (scanners, asset managers, compliance scanners) [3][6][10][11]. For each consumer function, the analysis identifies which input fields it requires and what fails when those fields are absent — the propagation map of Figure 2.

Consumer function	Required NVD input	Failure mode when enrichment is absent
Asset matching (scanner fire/no-fire)	CPE applicability statements	CVEs silently unmatched - coverage loss indistinguishable from not-applicable
Triage / prioritization	CWE classification, NVD-CVSS	Severity and class signals absent; queues starve or mis-rank
Compliance attestation	Analyzed record status	"Received/Awaiting Analysis" treated as incomplete by scanning regimes
Feed mirroring (OSV/GHSA-class)	Structured record content	Mirrors propagate IDs without enrichment; divergence across mirrors
Exploit-intelligence correlation	Record references and description quality	Correlation quality degrades with thin records

Table 2. Consumer functions and their failure modes when enrichment is absent.

Risk-register construction (RQ3). Six backlog-risk classes were derived from the propagation analysis; six mitigation mechanisms from documented ecosystem responses. Coverage levels are qualitative author assessments tied to mechanism arguments, with subjectivity stated in Section 7.

4. The Verified Timeline (RQ1)

Figure 1 and Table 1 present the record. Its arc has four phases. ***Collapse (Feb–May 2024):*** *enrichment slows to near-zero; 93.4% of new CVEs unanalyzed by 19 May; the backlog passes 18,000 [5][6].* ***Partial recovery with parallel sources (Jun–Dec 2024):*** *analysts onboarded; CISA's ADP Vulnrichment supplements CVSS/CWE from 3 July; processing improves to 2,000–3,000 CVEs/month, yet 72.4% of window CVEs remain unanalyzed at 21 September — including 46.7% of window KEVs [5][6].* ***Growth-constrained recovery (2025):*** *full analyst staffing; ~42,000 CVEs enriched (record throughput, +45%); but 32% submission growth in 2024 and continued growth keep the backlog rising [2][3][5].* ***Prioritized triage (April 2026):*** *enrichment becomes selective by policy — KEV-first within one business day, federal software prioritized, the pre-March-2026 backlog moved to "Not Scheduled" [2][3].*

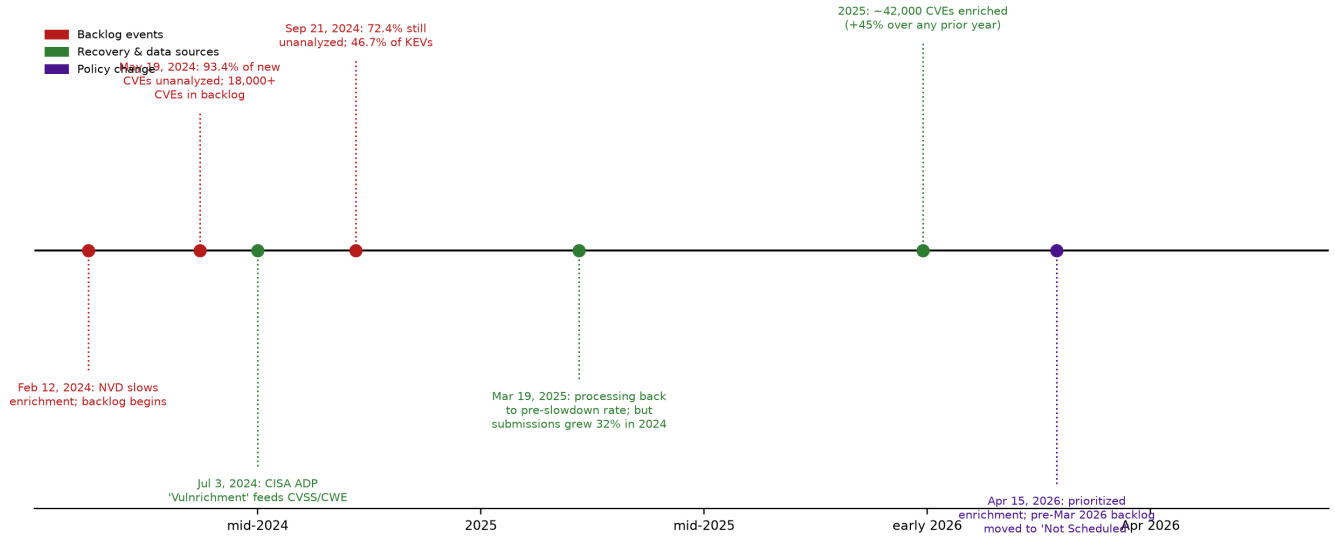

Figure 1. The NVD enrichment-gap timeline, 2024–2026: backlog events, recovery measures, data-source responses, and the April 2026 prioritization policy [2][3][5][6].

Event	Date	Documented measurement / content	Source
Enrichment slowdown begins	Feb 12, 2024	NVD transitions analysis workflow; enrichment of new CVEs slows to near-zero	[5][6]
Peak gap (first measurement)	May 19, 2024	93.4% of CVEs since Feb 12 unanalyzed; backlog >18,000 CVEs	[5]
CISA ADP "Vulnrichment" begins	Jul 3, 2024	ADP layer supplies CVSS/CWE directly onto CVE records	[6]
Partial recovery measurement	Sep 21, 2024	72.4% of window CVEs still unanalyzed; 46.7% of window KEVs unanalyzed; ~2,000-3,000 CVEs/month processed	[5][6]
Processing-rate recovery	Mar 19, 2025	Pre-slowdown rate restored; submissions grew 32% in 2024; backlog still growing	[3][6]
Record-throughput year	2025	~42,000 CVEs enriched (+45% over any prior year); insufficient against growth	[2]
Prioritized enrichment policy	Apr 15, 2026	KEV entries enriched within one business day; pre-Mar-2026 backlog moved to "Not Scheduled"	[1][2]

Table 1. The enrichment-gap timeline: events, dates, and documented measurements [2][3][5][6].

5. Propagation Through the Downstream Ecosystem (RQ2)

Figure 2 maps the choke point into the consumer ecosystem. Three propagation classes dominate.

Asset matching breaks. CPE applicability statements are the canonical bridge between "a vulnerability exists" and "my asset is affected." Tools that depend on CPE matching simply do not fire for unenriched CVEs — coverage loss invisible in tool dashboards, because the unmatched CVEs are indistinguishable from not-applicable ones [3][5]. This is the most dangerous failure mode precisely because it is silent.

Triage starves. CWE classification and severity scores drive prioritization queues. During the gap window, 46.7% of *new KEV entries* — the highest-priority class — lacked NVD analysis at one measurement point [5]; defenders triaging by NVD severity had no NVD-side signal for nearly half of the actively exploited additions, forcing reliance on ADP data, vendor advisories, or the CVE description alone.

Compliance scanning gaps. Regimes that require analyzed vulnerability records — U.S. federal continuous-scanning mandates and, prospectively, CRA-style reporting workflows — treat "Received/Awaiting Analysis" as an incomplete state [3][9]. The NIS2 and CRA regimes enter here because both presume analyzed, actionable vulnerability records: NIS2's essential-entity reporting and the CRA's product-vulnerability duties assume the record carries usable severity and applicability data — precisely what the backlog removed for the affected window [7][8]. The backlog turned a data-quality problem into a compliance-attestation problem for affected products.

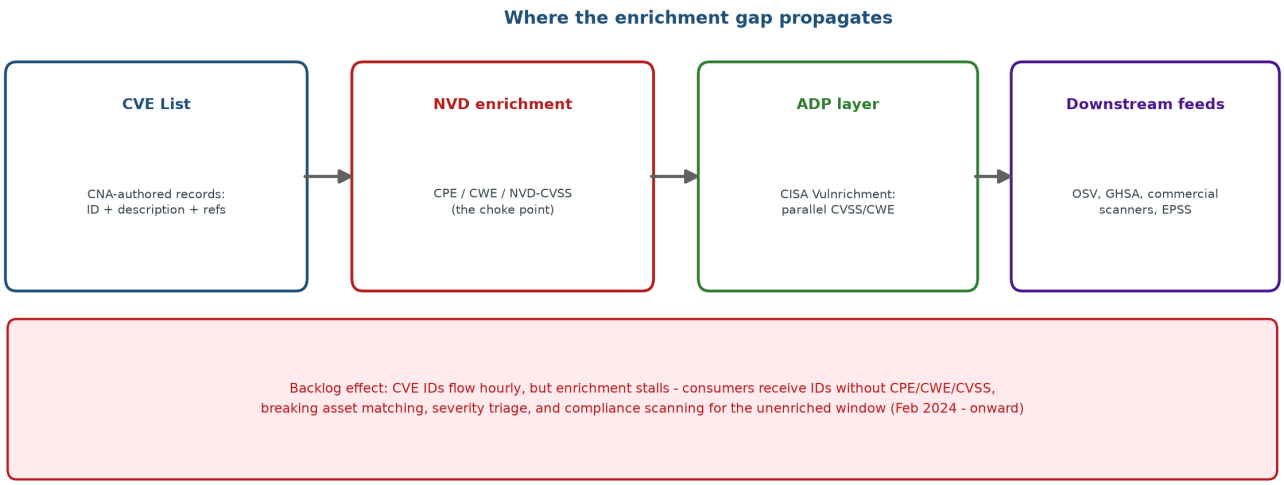

Figure 2. Propagation of the enrichment gap through the vulnerability-data architecture to downstream consumer functions [3][6][10][11].

5.1 The ADP layer: partial patch, structural lesson

The ADP "Vulnrichment" layer (July 2024) is the crisis's most instructive outcome: a second organization enriching the same records in parallel, within hours for priority classes. It demonstrated that the enrichment bottleneck is not intrinsic — parallel enrichment works — but it also surfaced the structural issues that make it partial: ADP data arrives with its own quality variance, covers a subset of fields, and creates duplicate-enrichment reconciliation problems between NVD and ADP records [6]. The April 2026 policy institutionalizes the layering [2][3]; consumers must now reconcile across enrichment sources as a routine operation, not an exception.

6. The Risk Register and the 2026 Policy (RQ3, RQ4)

Figure 3 and Table 3 present the backlog-risk register. Two findings anchor the consumer guidance.

The "Not Scheduled" tier is permanent and must be instrumented. The pre-March-2026 unenriched records are not queued — they are formally deprioritized [2]. Consumers should measure the tier's size and composition (KEV membership, severity indicators from ADP, product exposure) as a standing risk metric; its growth is the measurable trajectory of the intelligence gap.

Diversify enrichment inputs. The single-source dependency is the root cause from the consumer side. Consumers should ingest ADP data where available, exploit signals (EPSS, KEV) as interim triage, and adopt asset-matching methods that do not require CPE (product-name and ecosystem-based matching), accepting their precision cost deliberately rather than silently [3][5][9].

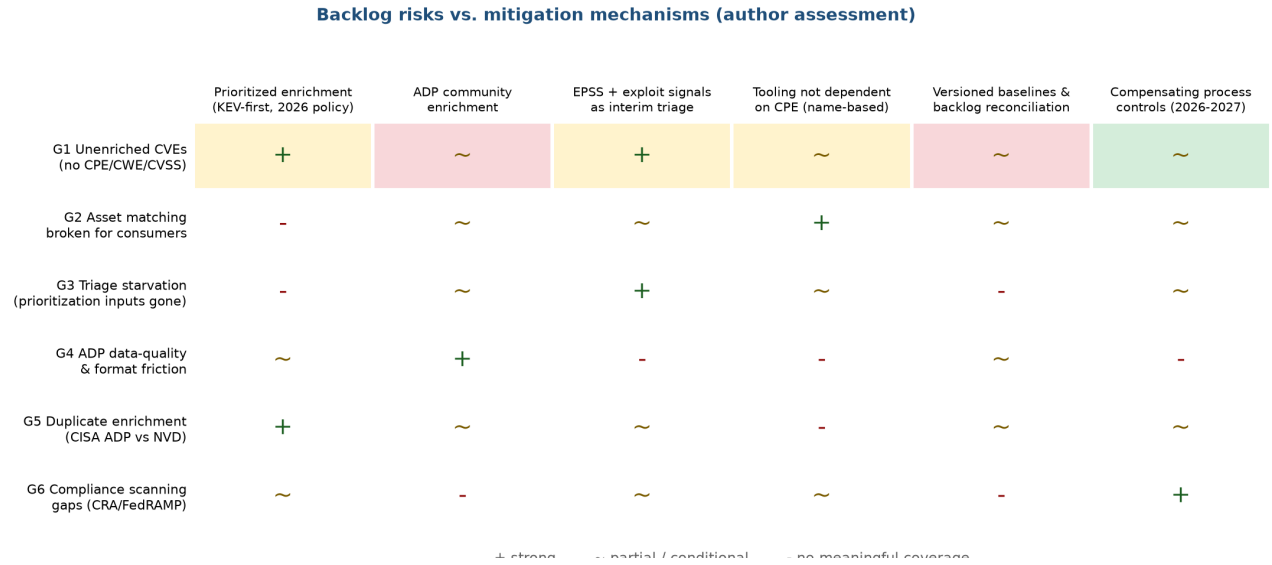

Figure 3. Backlog risks versus mitigation mechanisms (author assessment; register in Table 3).

Risk class	Description	Strongest mitigation	Coverage note
G1 Unenriched CVEs	No CPE/CWE/CVSS for the record	Prioritized enrichment (KEV-first, 2026 policy)	Strong for KEV class only; partial overall
G2 Asset matching broken	Tools silently unmatched on unenriched CVEs	Tooling not dependent on CPE (name/ecosystem-based)	Partial: precision trade-off
G3 Triage starvation	Prioritization inputs absent for new CVEs	EPSS + exploit signals as interim triage	Partial: probabilistic, model-version dependent
G4 ADP quality & format friction	Parallel enrichment arrives with variance	Reconciliation tooling	Partial: subset of fields
G5 Duplicate enrichment	NVD and ADP both enrich; conflicts	Versioned reconciliation	Partial: emerging practice
G6 Compliance scanning gaps	Unanalyzed records fail attestation	Compensating process controls (2026-2027)	Partial: regime-dependent

Table 3. Backlog risks, mechanisms, and coverage notes (author assessment).

7. Limitations and Threats to Validity

Construct validity. The six risk classes derive from the documented architecture; consumer implementations vary and may fail differently. **Internal validity.** Third-party measurements (VulnCheck) use their own analysis-status accounting; we quote them with attribution rather than independently re-measuring. **External validity.** The propagation map reflects the documented consumer classes; bespoke tooling may exhibit different failure modes. **Reliability.** Official announcements revised numbers over time (e.g., initial backlog-clearance estimates retracted [6]); we use the latest stated values per date. **Temporal scope.** Frozen at the April 2026 policy and its immediate aftermath; the policy's implementation performance and the OIG-evaluated program changes are out of scope.

8. Conclusion

The 2024–2026 NVD record is the vulnerability ecosystem's most instructive infrastructure failure: a single-organization enrichment layer, facing 263% submission growth, stalled — and the failure propagated silently into every downstream function that assumed analyzed records. The 2026 prioritization policy is an honest admission that the old model is over: NVD enrichment is now a triage decision, with a fast lane for the exploited class and a permanent "Not Scheduled" lane for the rest. For consumers, the mandatory engineering response is input diversification (ADP feeds, exploit signals), matching methods that survive missing CPE, and standing instrumentation of the unenriched tier. The enrichment gap of 2024 was a warning about architecture, not staffing alone — and the two-tier record it produced is now the baseline every consumer must engineer for.

Declarations

Funding. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

CRediT contribution statement. Risto Vaarandi: conceptualization, longitudinal audit, writing – original draft. Surafel Lemma Abebe: validation, propagation analysis, writing – review & editing. (Corresponding author: Risto Vaarandi.)

Competing interests. The authors declare no competing interests.

Data and code availability. No new datasets were generated; all measurements are quoted from the cited primary announcements and third-party analyses [2][3][5][6], with source and window attributed.

Ethics and dual-use. Not applicable; the paper contains no offensive content.

AI assistance declaration. AI-based tools were used in a limited supporting capacity (literature search, figure preparation, and language editing); all research content, analysis, and conclusions were produced, reviewed, and verified by the authors, who take full responsibility for the work.

References

[1] National Institute of Standards and Technology, "NIST Updates NVD Operations to Address Record CVE Growth," Apr. 2026. [Online]. Available: https://www.nist.gov/news-events/news/2026/04/nist-updates-nvd-operations-address-record-cve-growth

[2] National Vulnerability Database, "CVEs and the NVD Process" (enrichment scope and prioritization), NIST, Apr. 2026. [Online]. Available: https://nvd.nist.gov/general/cve-process

[3] P. Garrity, "Danger is Still Lurking in the NVD Backlog," VulnCheck Blog, Sep. 30, 2024. [Online]. Available: https://www.vulncheck.com/blog/nvd-backlog-exploitation-lurking

[4] Cybersecurity and Infrastructure Security Agency, "Known Exploited Vulnerabilities (KEV) Catalog," updated continuously. [Online]. Available: https://www.cisa.gov/known-exploited-vulnerabilities-catalog

[5] National Institute of Standards and Technology, "NVD News" (operational updates: analyst onboarding, ADP Vulnrichment, backlog estimates), 2024-2025. [Online]. Available: https://nvd.nist.gov/itl/nvd/nvd-news

[6] Infosecurity Magazine, "NVD Revamps Operations as Vulnerability Reporting Surges," Apr. 11, 2025. [Online]. Available: https://www.infosecurity-magazine.com/news/nvd-revamps-operations-cve-surge

[7] Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS2), OJ L 333, Dec. 2022.

[8] Regulation (EU) 2024/2847 (Cyber Resilience Act), OJ L, Nov. 20, 2024.

[9] Cybersecurity and Infrastructure Security Agency, "Known Exploited Vulnerabilities Catalog" reporting requirements under BOD 22-01, Nov. 2021 (updated 2022).

[10] S. Dong, M. Sun, B. Sun, T. Yang, and W. Meng, "Towards the Detection of Inconsistencies in Public Security Vulnerability Reports," in Proc. 28th USENIX Security Symposium, 2019. [Online]. Available: https://www.usenix.org/conference/usenixsecurity19/presentation/dong

[11] "Characterizing and Modeling the GitHub Security Advisories Review Pipeline," arXiv:2602.06009, 2026. [Online]. Available: https://arxiv.org/abs/2602.06009