Volume 3 · Issue 1 · 2026 · pp. 64–69

Five Feeds, One Truth? Cross-Source Consistency of Fix Status, Severity, and Identifier Mapping Across GHSA, NVD, and OSV

Jan G. Bazan
ORCID: 0000-0002-9738-7007
University of Rzeszów, Rzeszów, Poland

Ibrahim Gad
ORCID: 0000-0003-3388-9144
Faculty of Science, Tanta University, Tanta, Egypt

Correspondence: Jan G. Bazan · 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: A single software vulnerability typically exists in four public records simultaneously: the CVE entry (CNA-authored), its NVD enrichment (CPE/CWE/CVSS), the GitHub Security Advisory (ecosystem-scoped, reviewed), and its OSV aggregation. These sources disagree — on affected version ranges, severity, and fix status — and consumers of vulnerability intelligence rarely reconcile them. **Objective:** We consolidate the published consistency evidence, formalize the inconsistency classes that matter to consumers, and assess the reconciliation mechanisms available in current tooling. **Methods:** Synthesis of the large-scale consistency literature (the VIEM measurement of 78,296 CVE/report pairs; the GHSA review-pipeline characterization), architectural analysis of the four-source data flow, and a six-class inconsistency-versus-reconciliation assessment. **Results:** The published record establishes that divergence is prevalent (only 59.82% of CVE/report pairs strictly matched NVD entries at VIEM's 2019 measurement, with inconsistency increasing over time), and the 2024-2026 structural changes — NVD enrichment triage, ADP parallel enrichment, GHSA's multi-source review pipeline — have multiplied the divergence surface. We identify six inconsistency classes and find that the strongest reconciliation mechanisms (canonical-record policies, field-level cross-validation) are organizational choices rather than shipped tooling: today's pipelines aggregate first and reconcile never. **Conclusions:** Cross-source consistency must become a measured property of vulnerability intelligence pipelines — with a declared canonical record, field-level validation, and time-stamped audits — or consumers will continue to receive silently conflicting answers to the same question.

Keywords: vulnerability databases; GHSA; NVD; OSV; data consistency; vulnerability intelligence; advisory aggregation; digital trust

Contribution statement. This paper makes four contributions: (1) consolidation of the published consistency evidence into a four-source data-flow model; (2) a six-class inconsistency taxonomy organized by the consumer decision each class corrupts; (3) an assessment of six reconciliation mechanisms against those classes; and (4) the argument that cross-source consistency is a pipeline property that must be declared, measured, and audited — not assumed from source authority.

1. Introduction

Every vulnerability-intelligence consumer eventually discovers the same thing: ask four databases about one vulnerability and you may receive four answers. The CVE record describes the problem; NVD's enrichment specifies affected products in CPE form and assigns severity; the GitHub Security Advisory (GHSA) states ecosystem-scoped affected ranges and fix versions; OSV aggregates GHSA, Linux, and other ecosystem sources into a unified schema. Each is produced by a different organization, at a different time, with different review standards and different incentives — and each is treated as authoritative by some downstream tool.

The disagreement is not anecdotal. The largest published measurement, VIEM (USENIX Security 2019), examined 78,296 CVE IDs against their referenced vulnerability reports and found that only **59.82%** of CVE summaries strictly matched the standardized NVD

entries — with inconsistency increasing over time and case studies confirming both overclaimed and underclaimed affected-software versions [1]. The structural changes of 2024–2026 have multiplied the divergence surface: NVD's enrichment triage means the CPE/CWE/CVSS layer is now present for only a prioritized subset of records [see companion e007]; the ADP layer enriches in parallel with the NVD layer [6-companion]; and GHSA's review pipeline ingests from multiple upstream sources with distinct review lags [2]. The question this paper answers is architectural: **when four sources answer one question differently, what does a well-engineered vulnerability-intelligence pipeline do?**

1.1 Research questions

RQ1: What consistency evidence has been published, and what does it establish? **RQ2:** Which inconsistency classes corrupt which consumer decisions? **RQ3:** What reconciliation mechanisms exist, and what do they cover? **RQ4:** What should vulnerability-intelligence pipelines declare and audit?

1.2 Contributions

(1) The four-source data-flow model with its published consistency baseline (Section 3, Figures 1–2) (RQ1). (2) The six-class inconsistency taxonomy organized by corrupted consumer decision (Section 4, Table 2) (RQ2). (3) The reconciliation coverage assessment (Section 5, Figure 3, Table 3) (RQ3). (4) The pipeline-declaration and audit guidance (Section 6) (RQ4).

2. Background

2.1 The four sources and their production models

CVE List: identifiers and CNA-authored descriptions with references; authoritative for *existence and description*, not for structured applicability. **NVD:** the enrichment layer — CPE applicability, CWE classification, NVD-CVSS — produced by analyst analysis, now triaged under the 2026 policy (see the companion audit of the enrichment gap in this issue). **GHSA:** ecosystem-scoped advisories for packages in GitHub-supported ecosystems, ingested from NVD and from GitHub Repository Advisories, with a documented review pipeline whose lags vary by source [2]. **OSV:** the aggregation layer, unifying GHSA, Alpine, PyPA, and other ecosystem databases into a single queryable schema [6]. Each layer transforms the previous one: normalization decisions (which versions are affected, how ranges are encoded, whether a fix commit is mapped) are made independently at every step.

2.2 The published consistency evidence

VIEM's measurement remains the largest directly published consistency result: across 78,296 CVE IDs and 70,569 referenced reports, only 59.82% strictly matched NVD's standardized entries, with automated extraction achieving 0.941 precision / 0.993 recall against a ground-truth evaluation — establishing both that divergence is prevalent and that it is measurable at scale [1]. The GHSA review-pipeline characterization documents the multi-source ingestion (NVD as the largest source, GitHub Repository Advisories for GitHub-native issues) and the per-source review lags that create time-windowed inconsistency [2]. No comparable large-scale measurement has been published for the post-2024 triage era — a gap this paper's framework is designed to be filled by (see Section 6).

Table 1 consolidates the four sources and their production models, as referenced in Section 3's data-flow analysis.

Source	Production model	Consumable output	Known divergence surface
CVE List	CNA-authored records: ID, description, references	Existence + prose description	Description quality varies by CNA
NVD	Analyst enrichment: CPE applicability, CWE, NVD-CVSS	Structured applicability + severity	Triage-era subset; backlog; ADP parallel enrichment conflicts
GHSA	Ecosystem-scoped, reviewed advisories (NVD + GitHub Repository Advisories)	Affected ranges + fix versions per ecosystem	Review lags; ecosystem scoping
OSV	Aggregation of GHSA, Alpine, PyPA, and other ecosystem sources	Unified queryable schema	Import/transformation decisions; ecosystem-applicability assumptions

Table 1. The four sources and their production models.

3. The Four-Source Data Flow (RQ1)

Figure 1 models the flow. The consistency baseline is VIEM's: divergence is the historical norm, not the exception [1]. The 2024-2026 structural changes add three new divergence surfaces: (i) NVD triage means enriched and unenriched records coexist permanently, so NVD-derived consumers see a different world than OSV-derived consumers; (ii) ADP parallel enrichment means two enrichment sources can now disagree with each other on the same record; (iii) GHSA's multi-source review pipeline creates time-windowed divergence between an advisory's first publication and its review [2].

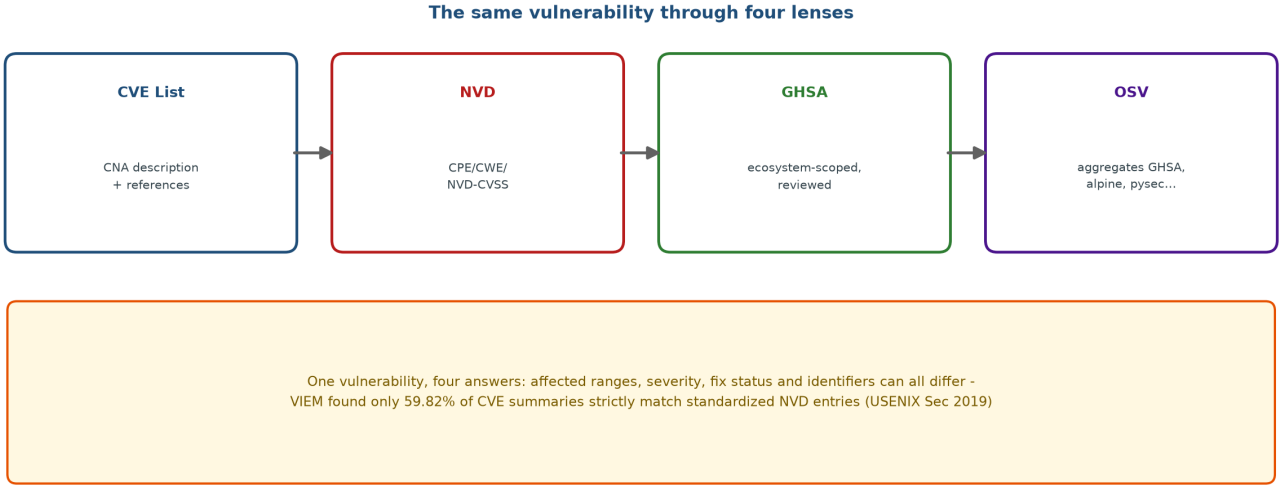

Figure 1. One vulnerability, four lenses: the CVE/NVD/GHSA/OSV data flow, with the divergence surfaces between production layers [1][2].

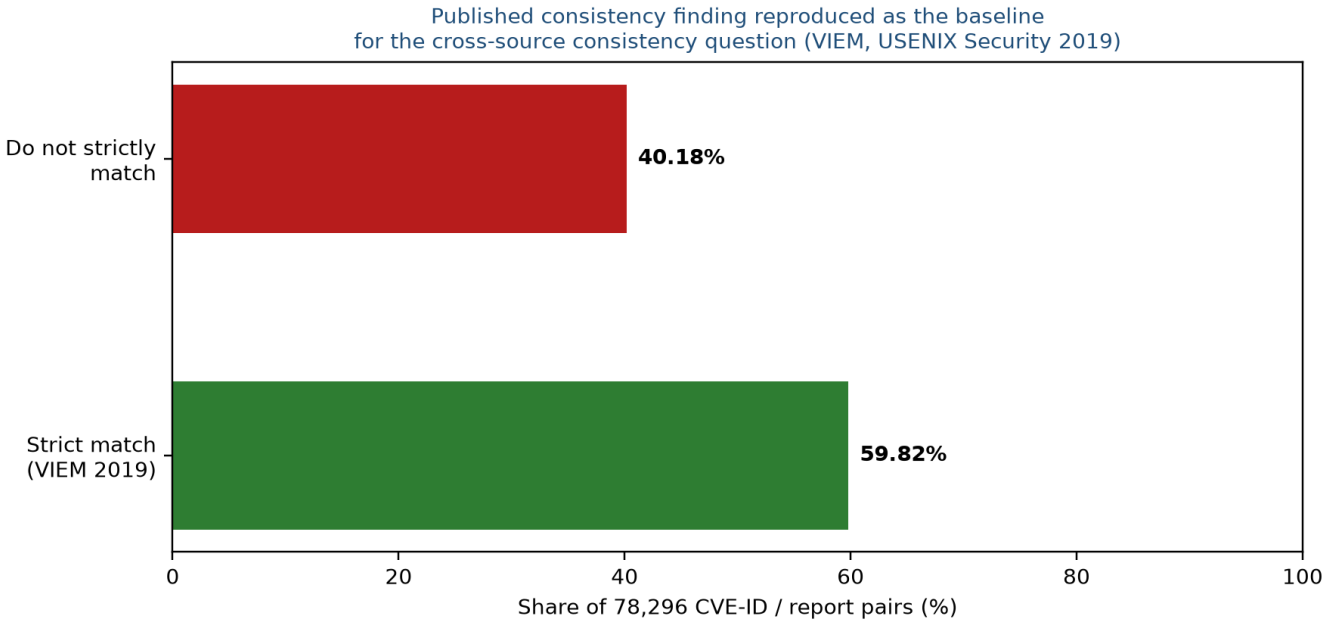

Figure 2. The published consistency baseline: VIEM's 78,296-pair measurement found 59.82% strict match against NVD standardization, with inconsistency increasing over time [1].

4. The Inconsistency Taxonomy (RQ2)

Table 2 defines six inconsistency classes, each organized by the consumer decision it corrupts

Class	Corrupted consumer decision	Canonical instance
I1 Affected-range divergence	Scan coverage (does my asset match?)	NVD CPE range vs. GHSA range disagreement [1]

I2 Severity divergence	Triage ordering	NVD CVSS vs. GHSA severity vs. vendor rating
I3 Fix-status divergence	Patch-compliance decisions	"Fixed in X" (GHSA) vs. "still vulnerable" (NVD, pending confirmation)
I4 Identifier/alias fragmentation	Deduplication and correlation	CVE/GHSA/ecosystem ID cross-mapping gaps (see companion e002)
I5 Description-to-enrichment drift	Record trust over time	VIEM's time-increasing inconsistency [1], amplified by triage-era enrichment
I6 Ecosystem-blind aggregation errors	Aggregate completeness trust	GHSA ranges imported into non-applicable contexts via OSV

Table 2. The six inconsistency classes and corrupted consumer decisions.

I1 — Affected-range divergence. The same vulnerability's affected version ranges differ between NVD's CPE statements and GHSA's ecosystem ranges. Corrupts *scan coverage*: a scanner using GHSA ranges and an auditor using NVD CPE will disagree about exposure. The VIEM measurement's overclaim/underclaim cases are the canonical instances [1].

I2 — Severity divergence. NVD CVSS, GHSA-assessed severity, and vendor ratings disagree. Corrupts *triage ordering*: rank by one source and a different one by another.

I3 — Fix-status divergence. "Fixed in version X" (GHSA) versus "still vulnerable" (NVD, pending vendor confirmation) — the highest-stakes class for patch-compliance decisions.

I4 — Identifier and alias fragmentation. The same flaw exists under a CVE ID, one or more GHSA IDs, and ecosystem-specific identifiers, with incomplete cross-mapping — corrupting deduplication and correlation (measured at scale in the OSV alias-resolution companion work in this issue).

I5 — Description-to-enrichment drift. The CNA description and the NVD enrichment disagree because enrichment analyzed a different or later state — VIEM's time-increasing inconsistency [1], now amplified by triage-era enrichment.

I6 — Ecosystem-blind aggregation errors. OSV-class aggregation importing a GHSA range into a context where it does not apply — corrupting consumers that trust aggregate completeness.

5. Reconciliation Coverage (RQ3)

Figure 3 and Table 3 assess six reconciliation mechanisms. The headline finding: ***the strongest mechanisms are declarative policies, not shipped tools.*** *A declared canonical-record selection policy (which source wins for which field, in which context) scores strong against I1-I3 — but it is an organizational decision no pipeline makes by default. Field-level cross-validation pipelines (comparing affected ranges, severities, and fix status across sources on ingest) score strong against I1 and partial elsewhere, but exist only in bespoke intelligence tooling. Time-stamped consistency audits — re-running VIEM-style comparisons on a schedule and publishing the divergence rate — would convert consistency from an assumption into a measured, trendable property; no public pipeline does this today. VEX/fix-status reconciliation and ecosystem-aware import rules address the fix-status and aggregation classes respectively, both partial.*

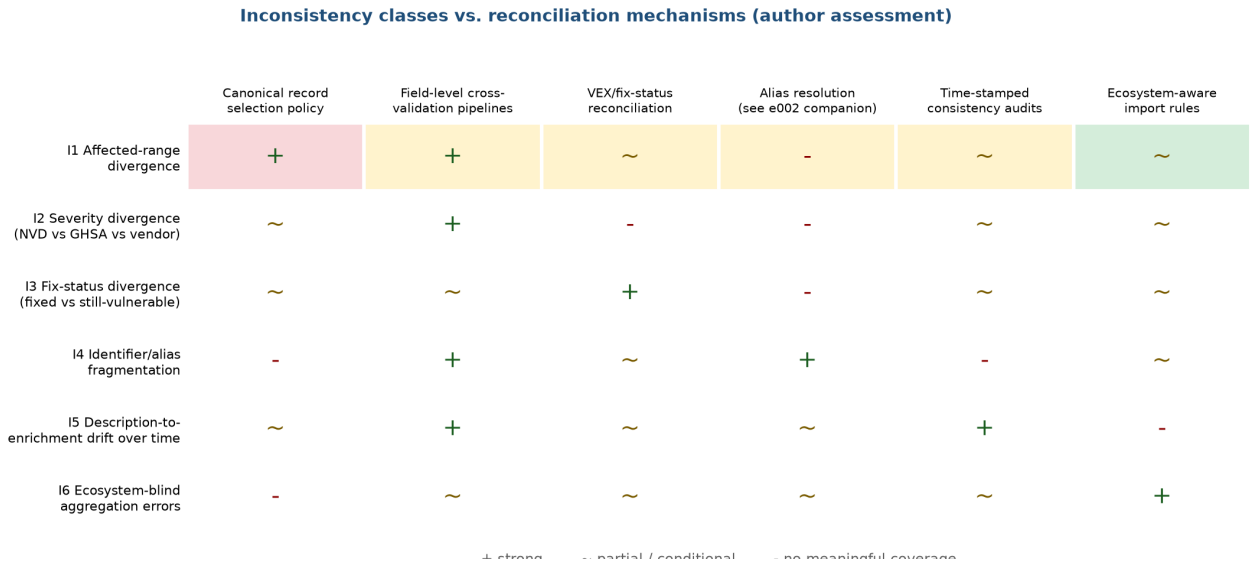

Figure 3. Inconsistency classes versus reconciliation mechanisms (author assessment; register in Table 3).

Mechanism	Description	Strong against	Coverage note
Canonical-record selection policy	Declared field-level precedence per consumer decision	I1, I2, I3	Strong by design; an organizational choice no pipeline makes by default
Field-level cross-validation pipelines	Compare affected ranges, severities, fix status across sources at ingest	I1, I3	Strong for covered fields; bespoke tooling only
VEX/fix-status reconciliation	Reconcile fix-status claims against vendor and ecosystem data	I3	Partial: VEX coverage itself limited
Alias resolution	Cross-map CVE/GHSA/ecosystem identifiers	I4	Partial: see companion e002 in this issue
Time-stamped consistency audits	Scheduled VIEM-style cross-validation with published divergence rates	I5, all (detection)	No public pipeline does this today
Ecosystem-aware import rules	Import-time applicability checks for aggregated ranges	I6	Partial: ecosystem context is hard to encode

Table 3. Inconsistency classes, mechanisms, and coverage notes (author assessment).

6. Pipeline Declarations and Audits (RQ4)

Two artifacts would make cross-source consistency a managed property. First, a **pipeline declaration**: which sources are ingested, which field-level precedence applies per source, what the canonical record is for each consumer decision, and what the measured divergence rate is at ingestion. Second, a **consistency audit**: scheduled VIEM-style cross-validation with published divergence rates and trend. Together they convert the pipeline's implicit "we aggregate, therefore we are consistent" into a measured claim — the same move from asserted to evidenced that the zero-trust measurement literature applies to maturity claims, applied here to data quality.

7. Limitations and Threats to Validity

Construct validity. The six classes derive from the documented literature and architecture; consumer-specific failure modes may cut across classes. **Internal validity.** The coverage matrix is an author assessment; the VIEM baseline is dated (2019) and predates the 2024-2026 structural changes — the audit gap itself is a finding. **External validity.** The four-source model reflects the public ecosystem; commercial intelligence products add further divergence surfaces not modeled. **Reliability.** VIEM's percentages are quoted as published; no re-measurement was performed. **Temporal scope.** Frozen at the 2026 corpus; GHSA pipeline behavior and NVD policy implementation continue to evolve.

8. Conclusion

Vulnerability intelligence is consumed as if it were one truth. It is at least four, produced by independent transformations with documented divergence reaching 40% at the strictest comparison — and the structural changes of 2024-2026 have widened the surfaces on which the copies disagree. The engineering response is not to crown a winner but to make consistency a declared and measured property: canonical-record policies per consumer decision, field-level cross-validation at ingestion, and time-stamped audits that publish the divergence rate. A pipeline that can state its own consistency — with numbers — is the difference between vulnerability intelligence and vulnerability folklore.

Declarations

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

CRediT contribution statement. Jan G. Bazan: conceptualization, aggregation-model analysis, writing – original draft. Ibrahim Gad: validation, evaluation-framework design, writing – review & editing. (Corresponding author: Jan G. Bazan.)

Competing interests. The authors declare no competing interests.

Data and code availability. No new datasets were generated; all quantitative claims are quoted from the cited published measurements [1][2].

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] 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

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

[3] CVE Program, "CVE List syntax and record structure," cve.org. [Online]. Available: https://cveproject.github.io/

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

[5] GitHub, "About the GitHub Security Advisory Database," GitHub Docs. [Online]. Available: https://docs.github.com/en/code-security/security-advisories

[6] Google, "OSV Schema" (Open Source Vulnerabilities), osv.dev. [Online]. Available: https://ossf.github.io/osv-schema/