Journal of Cybersecurity and Information Management

♦ ISSN: 2692-4048 Vol: 18, No. 02, 2026 ♦ PP. 26–40

♦ Article DOI: https://doi.org/10.54216/JCIM.180203

FULL LENGTH ARTICLE

Cybersecurity Information Management for Actively Exploited Vulnerabilities: Empirical Evidence from the CISA Known Exploited Vulnerabilities Catalog

Arash Salehpour[1,*]

[1] Department of Cybersecurity, Haliç University, Istanbul, Türkiye

Email: arashsalehpour@halic.edu.tr

Received: August 10, 2026 Revised: September 25, 2026 Accepted: October 29, 2026 ⋆ Corresponding author

ABSTRACT

Exploitation-driven remediation changes cybersecurity information management because, once exploitation is confirmed, an external vulnerability record must be connected quickly to assets, exposure, ownership, action and verification. This study examines every entry added to the U.S. Cybersecurity and Infrastructure Security Agency Known Exploited Vulnerabilities (CISA KEV) Catalog from 11 June through 31 July 2026, immediately after Binding Operational Directive 26-04 introduced a risk-based federal patching model. The population contains 39 KEVs: 34 records (87.2%, 95% Wilson interval 73.3%-94.4%) carried a three-day public catalog due-date interval and five carried 14 days, while the mean was 4.41 days and the median was 3 days. Arrivals were uneven. Weekly additions ranged from 2 to 10, producing a peak-to-mean load ratio of 2.05; the HHI-style vendor-concentration index was 598.3, and no vendor contributed more than five records. The technical mix was broad. Network/security/identity/control infrastructure and enterprise/web application platforms together accounted for 92.3% of records, while authentication, authorization and privilege-control weaknesses were the largest author-coded family at 33.3%. The 2026 Verizon Data Breach Investigations Report supplies a separate reference scale because its 43-day median full-resolution time is 14.33 times the three-day catalog interval and 9.75 times the observed 4.41-day mean, although these ratios compare time scales and do not estimate compliance. The evidence recasts post-exploitation prioritization as an information-flow problem in which rapid intake, asset matching, contextual enrichment, ownership, evidence capture and verification must operate under short deadlines that arrive in bursts. A field-level control design is therefore proposed for converting KEV intelligence into auditable CVE-asset remediation records while keeping public catalog dates separate from asset-specific BOD deadlines.

Keywords: Known Exploited Vulnerabilities ▪ Vulnerability management ▪ Cybersecurity information management ▪ CISA KEV ▪ BOD 26-04 ▪ Remediation prioritization ▪ Exploit intelligence ▪ Risk-based patching

1. INTRODUCTION

The scoring frame is familiar. Vulnerability management has usually assigned each disclosed weakness a severity value, added threat indicators when they were available, and placed the resulting item somewhere in a remediation queue. That framing becomes incomplete once exploitation is confirmed, because the difficult work then shifts from estimating whether an attacker may use a weakness to moving information across the organization quickly enough for an owner to find the affected asset, judge exposure, choose a treatment, test it, execute it and verify the result. The queue is therefore an

© 2026 The Author(s). Published by ASPG. This article is licensed under CC BY 4.0.

information-management system as much as a patching system; a missed asset identifier can matter as much as a weak ranking rule.

The timing pressure became unusually visible in June 2026, when CISA issued Binding Operational Directive (BOD) 26-04 on 10 June and replaced the older federal vulnerability directives with a risk-based scheme built around asset exposure, KEV status, exploit automation and technical impact [1]; the Directive also requires ongoing KEV monitoring, assigned responsibilities, internal validation, automated reporting where possible and machine-level asset tagging, so vulnerability intelligence becomes a chain of linked records alongside the CVE entry. The CVE is one record. Exposure, asset ownership, environment, remediation state and forensic status are others.

CISA's accompanying discussion gives the speed requirement an operational reference point: the 2026 Verizon Data Breach Investigations Report (DBIR) found that only 26% of KEVs were fully remediated by organizations in 2025 and that the median time to full resolution had risen to 43 days [2], [3], while Verizon also reported that vulnerability exploitation accounted for 31% of breaches and had become the leading initial-access route in that year's report [3]. The clocks differ. A three-day action window and a 43-day median resolution time describe different quantities, yet the distance between them is operationally meaningful.

This paper studies that distance through the public KEV information released immediately after BOD 26-04; the population contains every KEV with a CISA dateAdded value from 11 June through 31 July 2026, giving 39 catalog additions [4], and the analysis uses the public dateAdded and dueDate fields, vendor and product names, CVE identifiers and vulnerability names while avoiding any inference about whether a specific federal asset was exposed, whether exploitation was automatable on that asset, or whether an agency met its deadline. Those asset-level outcomes sit outside the public catalog.

Four questions organize the analysis: RQ1 asks how the public catalog due-date interval changed in the first seven weeks after the policy transition, and RQ2 asks whether new KEVs arrived as a steady flow or as bursts that could overload a fixed remediation team. RQ3 asks where the work landed across vendors, product planes and weakness families; RQ4 compares the resulting schedule scale with the 43-day remediation benchmark without turning that comparison into a compliance claim. One KEV catalog record is the unit of analysis, with asset instances entering later in the operational chain.

The contribution is managerial and empirical. Recent work has improved exploit prediction, KEV/EPSS integration, context scoring and attack-graph ordering [18]-[24], whereas this study starts after the exploitation signal has crossed the KEV threshold and asks what the resulting information stream demands from the organization that receives it; an accurate priority label cannot compensate for slow asset matching, unclear ownership or delayed verification, and the policy transition supplies a short natural window in which those downstream demands can be measured directly.

2. FROM SEVERITY SCORES TO ACTIONABLE EXPLOITATION INTELLIGENCE

The literature has moved away from severity-only patching; Bozorgi et al. [7] showed that disclosure metadata can support exploit prediction, while Allodi and Massacci [8] found that high CVSS scores alone are a poor rule for deciding what to fix first. Social and open-source signals widened the evidence base: Sabottke et al. [9] used Twitter activity, Bullough et al. [10] exposed temporal-evaluation problems in exploit-prediction studies, and Roumani [16] later connected information diffusion on Twitter with active exploitation. The common thread is simple. Exploitation is rare enough for the information used to identify it to matter more than an undifferentiated severity queue.

Prediction also became more operational. Jacobs et al. [11] evaluated remediation strategies through coverage and efficiency, EPSS formalized a probability-oriented signal for near-term exploitation [13], transfer learning improved description-based prediction [12], and time-varying expected exploitability introduced a stronger treatment of temporal evidence [15]. More recent work adds analyst labels [17], measures likely past exploitation [18], chains KEV, EPSS and CVSS [19], inserts cyber-physical context [20], adds network topology [21], represents exploit preconditions [22], tests a hybrid CNN-LSTM predictor [23], or ranks vulnerabilities with explainable top-k constraints [24]. The question then changes. These approaches answer which vulnerability deserves attention before or around the point at which exploitation evidence becomes decisive.

BOD 26-04 changes the next question. Once CISA has placed a CVE in KEV, the federal decision process must connect that vulnerability to an asset and then evaluate exposure, automation and technical impact [1]. The public KEV catalog remains a high-value threat signal, but the Directive makes internal asset and workflow information equally important; an exposed device may face a short path to action while the same CVE on another asset can sit on a different branch of the decision table. Asset context now matters.

SSVC provides the conceptual bridge. Its decision-tree design treats response as a stakeholder action supported by explicit decision points [14], and BOD 26-04 draws directly on that approach [1]. The Directive adds timelines, reporting duties and asset tags to the decision, producing an information-management model in which external threat intelligence defines what has been exploited, asset intelligence defines where it exists, contextual metadata defines urgency, and workflow records define whether treatment has finished. The joins are operational, not decorative.

Table 1 places the study against the main line of exploitability and vulnerability-prioritization research, with the final column separating pre-confirmation ranking from the post-confirmation information problem examined here. Most studies estimate exploitability or reorder a backlog before exploitation is confirmed; the present analysis measures public action pressure after KEV confirmation during a policy transition. Its focus is therefore arrival rate, deadline compression, concentration and the data operations needed to turn an external signal into a closed remediation record.

The table is deliberately bounded. Table 1 positions the study

Table 1. Selected progression of exploitability and vulnerability-prioritization research

Study	Evidence base	Method / decision frame	Operational contribution
Bozorgi et al. (2010) [7]	CVE text, timestamps and references	Supervised exploit prediction	Early evidence that disclosure metadata can support triage; model target is future exploitation.
Allodi & Massacci (2014) [8]	Public vulnerability and exploit data	Case-control risk comparison	High CVSS alone was a weak patching rule; exploit evidence carried more decision value.
Sabottke et al. (2015) [9]	Twitter and exploit ground truth	Social-signal exploit detection	Fast-moving public discussion can add early threat information beyond static severity.
Bullough et al. (2017) [10]	Open-source disclosure data	Replication and prediction	Temporal sampling choices can inflate exploit-prediction performance; evaluation design matters.
Jacobs et al. (2020) [11]	75k+ CVEs and observed exploits	Gradient-boosted remediation models	Coverage and efficiency expose the operational trade-off hidden by accuracy alone.
Yin et al. (2020) [12]	Vulnerability descriptions	Transfer learning	Text can support disclosure-time exploitability estimation when labeled data are limited.
Jacobs et al. (2021) [13]	Historical exploitation indicators	EPSS probability score	Moves prioritization toward probability of exploitation within a near-term horizon.
Spring et al. (2021) [14]	Stakeholder decision points	SSVC decision trees	Makes vulnerability response context-dependent and action-oriented instead of score-only.

after exploitation confirmation, where the research object becomes the flow of actionable information through remediation operations; the accuracy of another exploitability predictor sits outside that target.

3. POLICY TRANSITION AND MEASUREMENT DESIGN

The study window begins on 11 June 2026, the day after BOD 26-04 was issued, and ends on 31 July 2026; every KEV record whose dateAdded falls inside that interval is included, with catalog inclusion interpreted according to CISA's published KEV criteria [5]. No sample is drawn. The resulting population contains 39 records, including records with a dueDate after 31 July because the inclusion rule concerns the date CISA added the vulnerability; later KEV additions are excluded, which fixes the observed information stream at the requested cutoff.

The first measure is the catalog due-date interval, $W_i = \text{dueDate}_i - \text{dateAdded}_i$, measured in calendar days for record i. W_i belongs to the public record and is not an observed time-to-patch; under the Directive, the applicable asset timeline can depend on exposure and other decision points, can change when facts change, and begins at the earlier of KEV addition or asset identification [1]. The analysis keeps that distinction throughout.

Three descriptive measures capture pressure on the receiving organization: the three-day share, p_3, is the proportion of records with $W_i = 3$ and is reported with a 95% Wilson interval because the population window is small; the weekly load-burst ratio, LBR, divides the largest weekly arrival count by the mean across the eight calendar bins; and vendor concentration is summarized with $H = 10{,}000 \sum_v s_v^2$, where s_v is each vendor's share of the 39 records. H is used only to describe dispersion and carries no market-concentration interpretation.

A fourth measure compares time scales. The 2025 federal FISMA evaluation guide described the ordinary BOD 22-01 expectation for newer KEVs as remediation within two weeks [6], so 14 days is used as the legacy ordinary benchmark and $\text{DCI}_i = 1 - W_i/14$ as the deadline-compression index. The 2026 DBIR benchmark is kept separate: $R(W) = 43/W$ compares a catalog interval with the 43-day median full-resolution time cited by CISA [2], but the ratio neither forecasts lateness nor assigns a probability that an organization will miss a deadline.

Two classifications were added for analysis and are explicitly author-coded. Product plane groups the affected products into network/security/identity/control infrastructure; enterprise/web application platforms; AI/developer workflow tooling; and OT/protocol infrastructure, while weakness family groups the catalog vulnerability names into eight operational families such as authentication/authorization, command or code execution, file/path handling and deserialization. These codes organize the queue and do not replace CWE assignments; another coding scheme could place some boundary cases elsewhere.

Table 1. Selected progression of exploitability and vulnerability-prioritization research (continued)

Study	Evidence base	Method / decision frame	Operational contribution
Suciu et al. (2022) [15]	103,137 vulnerabilities	Time-varying expected exploitability	Treats exploitability as a changing process and addresses noisy negative labels.
Roumani (2024) [16]	6,004 vulnerability-related tweets	Logistic prediction of active exploitation	Information diffusion on social media adds signals associated with exploitation.
Parente et al. (2025) [17]	Expert labels and vulnerability attributes	Active-learning risk prioritization	Pairs analyst input with supervised learning for explainable remediation selection.
Mell & Spring (2025) [18]	KEV/EPSS measurement concepts	Likely Exploited Vulnerabilities metric	Focuses on measuring probability that exploitation has already occurred and KEV completeness.
Shimizu & Hashimoto (2026) [19]	28,377 CVEs; KEV, EPSS and CVSS	Threat-first decision chaining	Combining exploitation evidence and prediction can sharply reduce urgent workload.
Kausar et al. (2026) [20]	4,030+ ICS CVEs	Integrated CPPS risk and exploit prediction	Shows the value of sector and asset context beyond generic severity.
Țiță et al. (2026) [21]	895 CVE nodes in nine IaC topologies	Topology-aware attack-graph prioritization	Network position changes which vulnerabilities deliver the greatest risk reduction.
Siewruk & Bondaruk (2026) [22]	CVE preconditions and deployment context	Context score and taxonomy	Separates technical severity from the conditions that make exploitation practical.
Saklani et al. (2026) [23]	Vulnerability descriptions	Hybrid CNN-LSTM prediction	Uses disclosure text to estimate exploitability at decision time.
Das (2026) [24]	134,875 temporally eligible CVEs	Explainable stacked exploit-prioritization model	Temporal alignment and top-k capacity tests make model output closer to remediation decisions.

The calculations are intentionally transparent: counts appear beside percentages, category rules are stated, all 39 records are reproduced in Appendix A, and Table 2 defines the study variables and derived measures used in the analysis. No inferential model is fit to explain why CISA assigned a particular public due date because the public catalog does not expose every asset-level decision input required for such a model. A causal claim would exceed the data, so the policy transition is treated as an observable change in the information presented to remediation teams and the resulting workload geometry is measured directly.

4. DEADLINE REGIME: WHAT CHANGED IN THE PUBLIC CATALOG

Of the 39 KEVs added after the Directive through 31 July, 34 carried a three-day interval between dateAdded and dueDate and five carried 14 days; the three-day share is 87.2%, with a 95% Wilson interval of 73.3% to 94.4%. The mean interval is 4.41 days and the median is 3 days. In this transition population, three days is the normal public catalog interval and 14 days is the exception.

The shift is large. A three-day interval is 78.6% shorter than the ordinary two-week BOD 22-01 benchmark, and the 4.41-day population mean is 68.5% below that benchmark. The comparison still needs a boundary because the older regime could use shorter emergency dates, while BOD 26-04 defines dynamic asset-level timelines that the public KEV dueDate field cannot fully represent. Even with that qualification, the public information cadence has compressed sharply.

The five 14-day records are spread across Fortinet FortiOS, WordPress Core, KNX, Microsoft Active Directory Federation Services and Cisco Catalyst SD-WAN Manager, which argues against a simple vendor rule, while public urgency is selective within vendors as well: Microsoft has four three-day records and one 14-day record, Cisco has three and one, Fortinet has two and one, and WordPress splits one and one. Record-level context therefore has to survive the intake process.

The three-day tier spans different technical surfaces. It includes SharePoint deserialization, firewall-management credentials, SonicWall appliance flaws, Joomla-related upload and access-control problems, Ubiquiti network management weaknesses, Oracle enterprise applications and Langflow. A remediation program organized around a single platform team

Table 2. Study variables and derived measures

Measure	Definition	Interpretation
Catalog interval, W_i	dueDate − dateAdded	Calendar days; public KEV record field interval
Three-day share, p_3	$\text{count}(W_i = 3)/39$	Share of transition records with 3-day interval
Weekly load-burst ratio, LBR	$\max(A_w)/\text{mean}(A_w)$	Peak weekly arrivals relative to mean weekly arrivals
Vendor dispersion H	$10{,}000\sum_v s_v^2$	Descriptive concentration of KEV records across vendors
Deadline compression, DCI_i	$1 - W_i/14$	Scale comparison with ordinary two-week legacy benchmark
Resolution-scale ratio, $R(W)$	$43/W$	Scale comparison with 2026 DBIR median full-resolution time

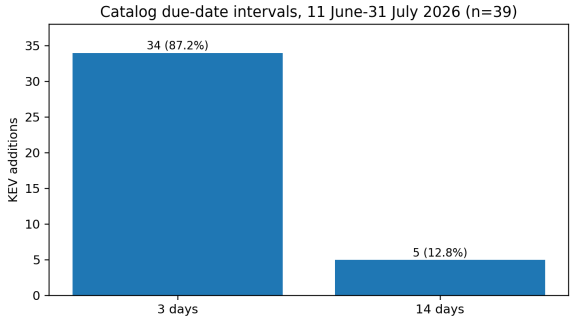


Figure 1. Catalog due-date intervals in the transition population

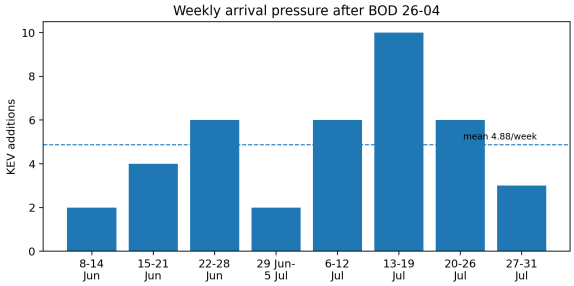


Figure 2. Weekly arrival pressure after BOD 26-04

would therefore receive the same external urgency across several ownership domains; deadline compression becomes a coordination problem before it becomes a patch-installation problem.

Figure 1 makes the deadline imbalance visible, while Table 3 gives the exact interval distribution. A team that polls the KEV catalog weekly can consume most of a three-day interval before it opens a ticket, which gives operational force to BOD 26-04's requirements for ongoing monitoring and automated reporting [1]. External-record detection must run on a cadence shorter than the shortest public action window.

5. QUEUE PRESSURE, BURSTS AND VENDOR DISPERSION

Short deadlines are harder when arrivals cluster. Across the eight reporting bins, additions were 2, 4, 6, 2, 6, 10, 6 and 3, giving a mean of 4.88 KEVs per bin and a population coefficient of variation of 0.52; the peak occurred during 13-19 July, when ten KEVs entered the catalog and eight carried three-day intervals. The load-burst ratio is 2.05, so the peak week carried just over twice the mean arrival volume.

A fixed daily patch quota is poorly matched to that pattern because a quiet week can make unused capacity appear ample, yet the next burst can create several concurrent due dates across independent technology stacks before an ordinary maintenance window opens. The queue has memory. Unfinished work from a prior burst competes with newly arrived KEVs, while verification and forensic tasks can remain after deployment; arrival counts therefore understate workload if the information system records only new tickets.

Vendor dispersion adds coordination cost because Microsoft accounts for five of the 39 records and Cisco for four, yet no vendor exceeds 12.8%; the top three vendors account for 30.8%, the top five for 43.6%, and the HHI-style concentration statistic is 598.3 on the 0-10,000 scale. Sixteen vendors appear exactly once, making a burst unlikely to yield to one vendor-specific playbook.

This dispersion matters for information quality because product names do not always map cleanly to configuration-management databases, scanner signatures and service catalogs; one vendor may publish a product-family name while the local inventory stores a package, appliance model or cloud service, and the external CVE-to-product record still has to reach the correct internal asset identifier quickly. False non-matches delay work. Broad wildcard matches create the opposite risk by sending the same alert to too many owners.

Table 4 and Figure 2 report the weekly arrival pattern, while Table 5 and Figure 3 describe vendor dispersion and deadline mix. Read together, they show a queue that must absorb bursts without losing record identity and route records across many ownership groups without turning every KEV into an enterprise-wide emergency. Automated enrichment should therefore occur before human judgment is spent on the ambiguous joins.

The spread is wide. Across the transition population, the vendor results in Table 5 and Figure 3 show 24 labels: the top three account for 30.8% of records, the top five for 43.6%, and the HHI-style record-concentration statistic is 598.3 on the 0-10,000 scale.

Table 3. Public catalog due-date interval distribution

Interval	Records	Share	95% Wilson interval for share	Compression vs 14-day benchmark
3 days	34	87.2%	73.3%-94.4%	78.6%
14 days	5	12.8%	5.6%-26.7%	0%
All records: mean	39	-	-	4.41 days; 68.5% below 14 days
All records: median	39	-	-	3 days

Table 4. Weekly KEV arrival pressure

Calendar bin	Additions	3-day	14-day	Mean interval (days)
8-14 Jun	2	2	0	3.00
15-21 Jun	4	3	1	5.75
22-28 Jun	6	6	0	3.00
29 Jun-5 Jul	2	2	0	3.00
6-12 Jul	6	6	0	3.00
13-19 Jul	10	8	2	5.20
20-26 Jul	6	5	1	4.83
27-31 Jul	3	2	1	6.67

Table 5. Vendor dispersion and deadline mix

Vendor group	Records	Share	3-day / total
Microsoft	5	12.8%	4/5
Cisco	4	10.3%	3/4
Fortinet	3	7.7%	2/3
Ubiquiti	3	7.7%	3/3
Langflow	2	5.1%	2/2
Oracle	2	5.1%	2/2
SonicWall	2	5.1%	2/2
WordPress	2	5.1%	1/2
Sixteen single-entry vendors	16	41.0%	15/16

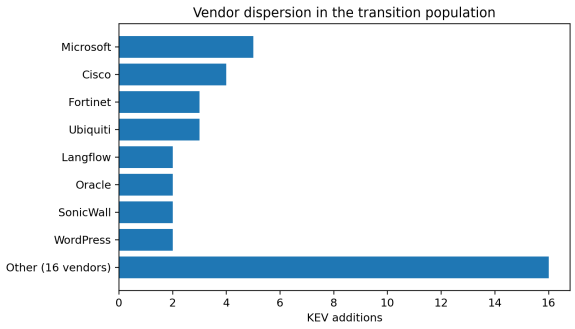


Figure 3. Vendor dispersion in the transition population

6. WHERE THE URGENT WORK LANDS

The product-plane classification places 20 of the 39 records (51.3%) in network, security, identity and control infrastructure. Enterprise and web application platforms account for another 16 (41.0%). Langflow contributes the two AI/developer-workflow entries and KNX the single OT/protocol entry; Table 6 and Figure 4 summarize the distribution, in which the first two planes hold 92.3% of the transition population. The urgent queue therefore sits close to systems that mediate access, traffic, identity, management and enterprise content.

That mix changes the cost of a three-day target. A web application can sometimes be patched with a familiar deployment pipeline; a firewall manager, identity service or SD-WAN controller may require change windows, redundancy checks, configuration backups, device sequencing or coordination with network operations. The same nominal interval can consume different amounts of organizational time before a patch command is executed; product-plane metadata belongs in the prioritization record because it helps predict who must act and what validation evidence will be needed.

The weakness-family distribution is concentrated in a different way. Authentication, authorization and privilege-control problems account for 13 records (33.3%), command/code execution and input-validation paths contribute eight (20.5%), and file/path/upload handling contributes seven (17.9%). Those three families make up 71.8% of the population. Deserialization/interpretation and web-request/query handling contribute four each. Memory safety, sensitive-information exposure and dynamic inclusion each appear once; Table 7 reports the full author-coded family distribution.

The families are useful for playbook selection, but they should not become a second severity score. A missing-authentication flaw in a public management plane and the same general weakness on an internal, tightly restricted service can de-

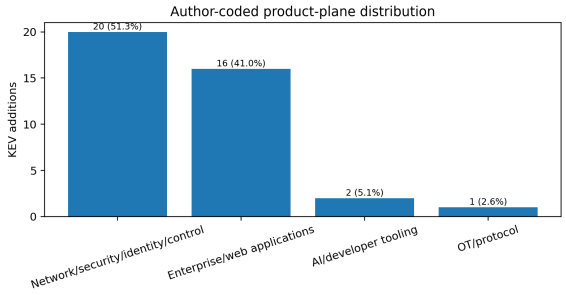


Figure 4. Product-plane distribution

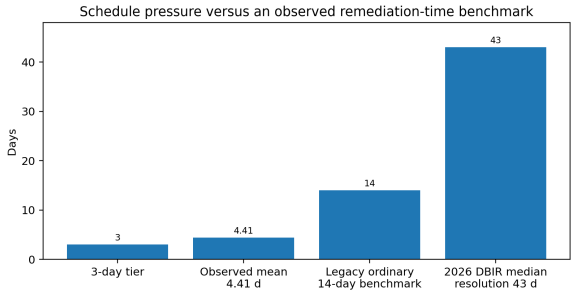


Figure 5. Schedule pressure versus the 2026 DBIR median resolution benchmark

mand different action under BOD 26-04 because exposure and technical impact are separate decision points [1]. The code is therefore a routing aid: it can select validation steps, likely evidence sources and forensic checks while the asset context controls urgency.

Age adds one more complication. Five records, 12.8% of the population, have CVE identifiers from before 2026: two from 2025, one from 2023, one from 2021 and one from 2008; a new KEV queue cannot be reduced to a stream of newly disclosed bugs. Old identifiers can become operationally new when fresh exploitation evidence changes their status; vulnerability information systems must keep historical inventory data queryable and should not retire matching logic merely because a CVE has aged out of an ordinary patch campaign.

The older-CVE pattern connects the empirical findings with prediction research, which asks which disclosed vulnerabilities will become important [7], [11], [13], [15], [24], while KEV supplies confirmed exploitation evidence after that uncertainty has narrowed and the pre-2026 identifiers show why information retention still matters after a vulnerability has been deprioritized in an earlier cycle. The queue can reopen. New exploitation evidence changes the record's operational status.

7. THE 43-DAY REALITY GAP

The 43-day figure from the 2026 DBIR supplies a separate reference scale because it describes observed full resolution, while the catalog intervals describe expected action timing [2], [3]; Table 8 and Figure 5 show that the median three-day interval is 14.33 times shorter than 43 days, the 4.41-day mean is 9.75 times shorter, and even the 14-day tier is 3.07 times shorter. No ratio proves noncompliance. The comparison quantifies the distance between an operational benchmark and the public schedule visible in the transition population.

That distance changes what can reasonably remain manual. If intake occurs once per business day, asset matching takes another day, ownership confirmation consumes a third day and a change ticket then waits for approval, a three-day calendar interval can expire before technical remediation starts, especially across a weekend. The example is illustrative; it is not a measured agency workflow. Its purpose is to expose how a serial information process can spend the available interval on handoffs and why parallelism and preassigned ownership become design requirements before technical remediation can begin on the affected system.

The DBIR result also cautions against measuring success at patch deployment. Full resolution can require verification, compensating-control confirmation, exception closure and—in designated BOD 26-04 cases—forensic triage [1]; a dashboard that marks a record complete when a package was pushed may report speed while leaving the actual vulnerable state uncertain. The information system needs a state model that separates assigned, mitigated, patched, verified, exception-approved and closed outcomes.

A useful operating metric is elapsed decision latency measured beside elapsed patch latency: decision latency starts when the KEV record enters the organization's feed and ends when the affected asset has an accountable owner, a context determination and an approved action path. Patch latency starts later. If decision latency consumes most of a three-day interval, more deployment automation will have limited effect. The bottleneck is upstream.

Table 8 and Figure 5 make the scale gap visible without treating it as a prediction. The practical lesson is conservative: an organization cannot assume that its historical median remediation process will satisfy a queue dominated by three-day public intervals. It must know where time is spent. Each state transition therefore needs a timestamp that remains available for review and makes delay attributable to a specific handoff.

8. INFORMATION-MANAGEMENT CONTROL DESIGN

The evidence suggests one canonical remediation record per CVE-asset pair: the external KEV event supplies the CVE, vendor, product, dateAdded, dueDate and remediation text, while internal enrichment attaches asset identifiers, exposure state, environment, service owner, business service, scanner evidence and the applicable decision inputs. Source values should remain separate. Overwriting the external record makes later audit difficult.

Ingestion has to be event-driven or frequent enough that a three-day record is visible quickly; for many environments a scheduled daily pull is a reasonable upper bound, while a weekly pull would consume too much of the available catalog interval before the record enters local triage. Each intake run should be idempotent because the same KEV may arrive through the catalog, a vendor advisory and a commercial threat feed. One CVE can create hundreds of asset work items. Deduplication by CVE alone is therefore insufficient once the work becomes asset-specific, because those assets can have different exposure states and owners.

Asset matching is the first major information join because

Table 6. Author-coded product planes

Product plane	Records	Share	Examples / scope
Network, security, identity and control infrastructure	20	51.3%	Routers, firewalls, remote support, identity, security analytics and management planes
Enterprise and web application platforms	16	41.0%	SharePoint, WordPress, enterprise applications, web builders and server-side platforms
AI/developer workflow tooling	2	5.1%	Langflow
OT/protocol infrastructure	1	2.6%	KNX protocol implementation

Table 7. Author-coded weakness families

Weakness family	Records	Share
Authentication, authorization and privilege control	13	33.3%
Command/code execution and input-validation paths	8	20.5%
File, path and upload handling	7	17.9%
Deserialization and interpretation paths	4	10.3%
Web request and query handling	4	10.3%
Memory-safety failure	1	2.6%
Sensitive-information exposure	1	2.6%
Dynamic inclusion / untrusted functionality	1	2.6%

Table 8. Time-scale comparisons (descriptive only)

Time scale	Days	43-day / scale	Meaning
Three-day catalog tier	3.00	14.33x	Policy-facing public interval
Transition-population mean	4.41	9.75x	Mean of 39 public intervals
Legacy ordinary newer-KEV benchmark	14.00	3.07x	BOD 22-01 scale comparator
2026 DBIR median full resolution	43.00	1.00x	Observed benchmark cited by CISA

exact software identifiers rarely line up perfectly across CISA, CVE data, scanners and local inventories. The system therefore needs a hierarchy of matches—deterministic product identifiers where available, normalized vendor/product aliases, version evidence and a human review queue for uncertain cases. Negative matches need evidence. A 'no affected assets found' result should record the inventory scope and query time so it can be reassessed when inventory changes.

Context follows. BOD 26-04 names asset exposure, KEV status, exploit automation and technical impact as the core variables [1]; KEV status is already true for the records studied here, while the other values may come from CISA metadata or the local environment. Storing both the input values and the resulting tier makes a later change in exposure or metadata reproducible instead of leaving an unexplained priority change.

Assignment should be precomputed wherever possible because ownership can come from the asset/service catalog before an emergency arrives, with escalation paths already defined for missing or disputed owners; the July burst makes late manual owner discovery a fragile control. The fan-out is immediate. Ten KEVs in one week can reach network, identity, web, enterprise application and platform teams at once, so a routing table based on product plane and service ownership can absorb that spread faster than a central analyst negotiating each ticket separately.

Verification closes the information loop: a scanner rescan, version query, configuration check or compensating-control test should return machine-readable evidence linked directly to the work item, while exception records retain expiry dates and responsible approvers. Where forensic triage is required, the triage outcome belongs beside the remediation outcome because the Directive treats them as linked obligations [1]. Closure needs evidence. A checkbox alone is insufficient.

Table 9 translates the proposed data contract into fields and controls, and Table 10 turns the transition findings into operational tests for existing vulnerability-management, ITSM, SIEM, CMDB and CDM-style systems. Every handoff should carry an identifier, a timestamp, a decision input and an accountable owner. That structure makes delay measurable.

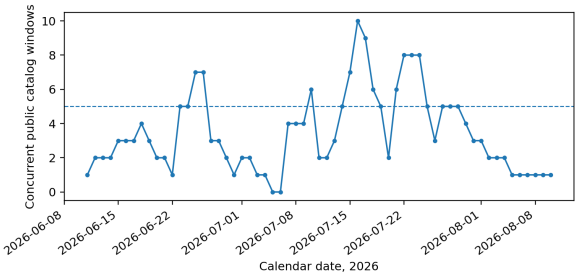


Figure 6. Daily concurrency of public KEV catalog windows, counted inclusively from dateAdded through dueDate. The dashed line marks five concurrent windows; the measure describes catalog scheduling overlap and does not assert unresolved agency assets.

Once delay is visible, the organization can determine whether its constraint sits in threat intake, inventory quality, ownership, change approval, deployment or verification.

8.1 Deadline Collision and Concurrent Catalog Windows

A weekly count tells only part of the scheduling story because due dates from earlier additions can remain active when a new burst arrives. To expose that overlap, an active public catalog window was defined for each record from dateAdded through dueDate, inclusive, and O_t counts how many of the 39 windows cover calendar day t; Table 11 reports the resulting collision measures. The measure is operational. It does not imply that every federal agency had an exposed asset or an unresolved remediation item on each day.

The overlap was material. O_t reached 10 on 16 July, fell to 9 on 17 July, and remained at 8 on 22-24 July; across the 61 calendar days from the first addition on 11 June through the last due date on 10 August, the mean was 3.46 and the median was 3. Nineteen of the 61 days carried at least five concurrent public catalog windows, and nine days carried at least eight.

Arrival collisions also appear at shorter horizons: four KEVs entered the catalog on each of 23 June, 7 July, 14 July and 21 July, while four records shared a due date on 26 June and another four on 10 July. Figure 6 traces the daily concurrency series and shows the densest three-day arrival interval at ten records, from 14 through 16 July, with nine of those ten on three-day windows. A team that relies on one-by-one enrichment can therefore face several records whose public windows overlap before the first item has cleared asset matching and ownership assignment.

The concurrency profile changes the capacity question. A ticketing system can ingest ten records quickly, yet the harder work begins after ingestion: each CVE must be mapped to products and versions, candidate assets must be confirmed, exposure and decision values must be attached, owners must be found and evidence must return before closure. Serial review is a poor default for a burst. The queue needs parallel matching and routing even when actual patch deployment remains constrained by change controls.

The same profile supplies a concrete test for automation. Feed ingestion can be measured in minutes, but a useful service-level objective also needs a percentile for CVE-to-asset match time and a percentile for owner assignment; where either stage stretches into the next day's arrivals, the queue can grow even when patch deployment itself is fast. Capacity can then be added where the queue actually forms. The measure is simple, but it exposes a scheduling property that weekly totals hide.

8.2 Legacy-CVE Reactivation and Vulnerability Memory

Active exploitation also reopens old records. Five of the 39 KEVs, or 12.8%, carry CVE identifier years earlier than 2026: two are from 2025, one from 2023, one from 2021 and one from 2008; Table 12 lists the cases, and using 2026 minus the identifier year as a coarse age proxy gives ages of 1, 1, 3, 5 and 18 years, with a mean of 5.6 and a median of 3. The identifier year is a coarse clock because CVE assignment and disclosure can cross calendar boundaries.

The operational consequence is a memory requirement because asset and vulnerability records cannot be managed as a rolling list that loses historical mappings once a routine patch cycle closes; fresh exploitation evidence can make an old identifier urgent again while the affected technology remains deployed in a forgotten appliance, branch system, inherited application or long-lived protocol stack. The public KEV record supplies the external trigger. Internal history answers a different question: where does this CVE still matter here?

A mature CVE-asset record should therefore retain at least the historical match, last-seen timestamp, disposition, exception state, patch or mitigation evidence, and service owner even after closure; when KEV status changes, the organization can query that history before starting discovery from zero. Stale assets still need confirmation because inventory truth changes, yet a preserved relationship turns rediscovery into validation; the 2008 identifier in this short population makes the point without requiring an assumption about why exploitation became relevant in 2026.

Vendor dispersion strengthens the same memory argument because the 39 records span 24 vendor labels and the reciprocal of the observed HHI share concentration corresponds to an effective count of 16.71 equally weighted vendors; the Shannon effective count is 20.16, which is higher because it gives more weight to the long tail. The long tail matters. Sixteen vendors appear once, so historical matching and ownership logic cannot be maintained only for the most frequent names.

Collision and memory point in opposite temporal directions, which is why they belong in the same information architecture: the collision metric asks whether several urgent records can be processed at once, while the legacy-CVE result asks whether the organization can recover an old relationship quickly. One requires parallelism. The other requires persistence. The same canonical CVE-asset record supports both demands.

8.3 Queue Occupancy and Long-Window Carryover

Record counts and queue occupancy answer different questions. The transition population contains 34 three-day records and only five 14-day records, yet summing W_i across the population produces 172 elapsed catalog interval-days: 102 from the three-day class and 70 from the 14-day class, meaning that the minority class contributes 40.7% of total interval-days while representing 12.8% of records. Table 13 reports this deadline-class contribution. The result is a scheduling property of the public W_i values and does not measure reme-

Table 9. Minimum information contract for an exploited-vulnerability remediation record

Stage	Minimum fields	Control purpose
Threat intake	CVE, dateAdded, dueDate, source, vendor, product, action text	Detect and preserve external record; deduplicate source events
Asset match	Asset ID, software evidence, version, scanner ID, match confidence	Create one CVE-asset work item; retain negative-match evidence
Context	Exposure, environment, KEV status, exploit automation, technical impact	Store decision inputs and their source/time; recalculate when facts change
Ownership	Service owner, technical owner, escalation path	Route automatically; flag missing or disputed ownership immediately
Action	Mitigation/remediation method, change ID, exception ID, planned time	Separate chosen treatment from deployment status
Verification	Rescan/version/configuration evidence, validator, timestamp	Close only with machine-readable or attributable evidence
Forensic state	Triage required, triage completed, finding, incident link	Keep forensic obligation coupled to the remediation record where applicable
Audit timing	State-transition timestamps, overdue reason, approval history	Measure decision, patch and verification latency separately

Table 10. Operational tests suggested by the transition data

Test	Metric	Suggested management criterion
KEV intake latency	Time from catalog publication to internal record creation	P95 comfortably below one day
Asset-match completeness	Affected products with at least one validated asset query	100% of new KEVs; uncertain matches queued same day
Owner resolution	Time from matched asset to accountable owner	Measured separately; preassigned ownership preferred
Decision latency	Intake to approved action path	Must leave sufficient time inside applicable asset timeline
Verification latency	Deployment/mitigation to verified state	Tracked separately from patch execution
Burst capacity	Concurrent three-day work items handled without intake backlog	Stress test at >=2.1 times ordinary weekly arrival rate

Table 11. Temporal collision measures derived from the 39-record transition population

Measure	Observed value	Interpretation
Peak additions on one date	4 KEVs	Occurred on 23 Jun, 7 Jul, 14 Jul and 21 Jul
Peak records sharing one due date	4 KEVs	Occurred on 26 Jun and 10 Jul
Maximum 3-day rolling additions	9 KEVs	14-16 Jul
Maximum 7-day rolling additions	12 KEVs	10-16 Jul
Maximum inclusive active catalog windows	10	16 Jul
Mean / median active catalog windows	3.46 / 3	Computed over 11 Jun-10 Aug
Calendar days with $O_t \geq 5$	19 of 61	31.1% of the inclusive observation span
Calendar days with $O_t \geq 8$	5 of 61	8.2% of the inclusive observation span

Table 12. Earlier-year CVE identifiers reactivated as KEV work in the transition window

CVE	Product / vendor context	Identifier-year lag to 2026	Catalog due-date interval
CVE-2008-4128	Cisco IOS	18	3 days
CVE-2021-27137	DD-WRT	5	3 days
CVE-2023-4346	KNX protocol implementation	3	14 days
CVE-2025-67038	Lantronix EDS5000	1	3 days
CVE-2025-68686	Fortinet FortiOS	1	14 days

diation effort, asset count, staffing burden or the time any particular system remained unresolved.

A longer catalog window can remain visible in the portfolio while several three-day records arrive and close their public windows; a manager looking only at weekly arrivals may see short-deadline throughput and miss the carryover created by a few longer records, while a manager looking only at open items can make the opposite error and treat a long-lived public window as evidence of greater technical work. The interval class therefore has to sit beside both counts. Arrival volume and queue occupancy describe different pressure mechanisms.

This split creates a two-speed queue. Three-day records dominate intake and force rapid matching, routing and evidence capture, while 14-day records are fewer but remain in the scheduling horizon long enough to overlap several later bursts; the five long-window records came from five different vendor labels—Cisco, Fortinet, KNX Association, Microsoft and WordPress—so the carryover cannot be assigned to one vendor family. Four of those five vendors also appear in the three-day class, which means the same vendor integration may have to support both deadline regimes inside a single month even though only one KNX Association record appears in the long-window class.

A useful dashboard should separate arrival volume, open catalog-window count and interval-day occupancy because each measures a different part of queue pressure: arrival volume measures the front door, open-window count measures simultaneous scheduling obligations under the public dates, and interval-day occupancy records how much of the catalog time horizon is contributed by each deadline class. None is a performance measure by itself. Together they make queue state visible before organization-specific asset counts, change windows and exception processes add further work; during a burst, the measures can move in different directions as short windows close and longer windows remain active.

8.4 Vendor Dispersion and Integration Burden

Vendor concentration is low enough that raw top-vendor counts understate integration burden: 24 vendor labels appear in 39 records, Microsoft contributes five and Cisco four, sixteen vendors appear only once, the top three vendors account for 30.8% of the population and the top five for 43.6%. The long tail dominates. Repeated vendors still matter for reusable scanners, package logic and ownership rules, but most records sit outside any small group that a single integration team could treat as the whole queue.

Table 14 reports two diversity measures for the 24 observed vendor labels. Using shares p_v, the HHI-style concentration is 10,000 times the sum of p_v squared and equals 598.3; its inverse-concentration equivalent, $10{,}000/\mathrm{HHI}$, is 16.71 equally weighted vendors, while Shannon entropy is 3.004 nats and corresponds to an effective count e^H of 20.16 vendors. The measures answer different questions. Both place the workload far from a queue dominated by two or three vendors.

Translation quality becomes the control problem because a scanner may use one product string, a CMDB another and a vendor advisory a third; the mismatch is paid for in the first hours of the deadline, while a long vendor tail requires normalized aliases, evidence-backed negative matches and a review path for uncertain joins even when reusable integrations reduce work for repeated vendors. The cost is early. The queue cannot depend on a fresh manual taxonomy exercise each time a single-entry vendor appears.

The deadline and dispersion results interact because 23 distinct vendor labels occur among the 34 three-day records, placing the shortest public interval across almost the entire vendor set observed in the study; speed therefore cannot be isolated to one narrow technology family or specialist team. Routing has to reach different service owners quickly. The data contract in Table 9 must work even when a vendor appears once, because a single-entry vendor can still arrive on a three-day catalog interval. That is the information-management burden created by the post-transition catalog stream.

9. SENSITIVITY, BOUNDARIES AND THREATS TO VALIDITY

This is a short transition study. The 39 records cover seven weeks after a major directive and should not be treated as a stable long-run distribution of KEV deadlines; another quarter could contain fewer additions, different vendors or a larger share of 14-day records. The Wilson interval quantifies statistical uncertainty around the observed three-day share. The population itself is complete for the specified window.

The public dueDate field also has limits because BOD 26-04 defines asset-level timing through decision variables, allows timelines to change as facts change, and starts the applicable clock from the earlier of KEV addition or asset identification [1]. A catalog interval cannot reconstruct the deadline that applied to every agency asset. This paper avoids that inference. Its observed object is the public information stream and its scheduling pressure, with asset-specific logic left as a design requirement instead of an outcome.

Vendor and weakness coding introduces analyst judgment because the categories are broad enough to be reproducible from the product and vulnerability names, yet another analyst could still place a remote-support product or an input-validation flaw in a neighboring class even though Appendix A leaves every source label visible. The core deadline, weekly-arrival and vendor-count results do not depend on those codes. Plane and family percentages should be read as an organizing view of the workload.

The 43-day comparison has a different boundary because Verizon's median full-resolution time reflects organizations and observations in the DBIR data, while CISA catalog intervals are federal policy-facing fields with different populations, clocks and operational processes [2], [3]. However, the ratio still expresses the order-of-magnitude difference in time; it does not estimate the probability of meeting BOD 26-04 or identify the source of delay inside any one organization.

One further limit concerns the counterfactual: BOD 26-04 superseded BOD 22-01, but this study does not estimate what CISA would have assigned to these same 39 records under the old directive, and the 14-day comparator is only the ordinary newer-KEV expectation documented for BOD 22-01 [6]. Emergency actions and special cases existed before June 2026. The compression index is therefore a policy-scale benchmark; it is not a reconstructed due date or causal

Table 13. Contribution of deadline classes to elapsed catalog interval-days

Catalog interval class	Records	Record share	Sum of W_i (days)	Share of interval-days	Distinct vendor labels
3 days	34	87.2%	102	59.3%	23
14 days	5	12.8%	70	40.7%	5
Total	39	100.0%	172	100.0%	24

Table 14. Vendor-dispersion measures for the transition population

Dispersion measure	Observed value	Operational reading
Unique vendor labels	24	39 KEVs span a long vendor tail
Largest vendor share	12.8%	Microsoft: 5 of 39
Top-three vendor share	30.8%	Microsoft, Cisco and Fortinet
Top-five vendor share	43.6%	Less than half of the population
HHI-style concentration	598.3	0-10,000 share-squared scale
Inverse-HHI effective count	16.71	Equivalent number under equal shares
Shannon entropy	3.004 nats	Entropy of vendor shares
Shannon effective count	20.16	e^H
Distinct vendors in 3-day class	23	Urgent short-window work is widely distributed
Distinct vendors in 14-day class	5	Each long-window record has a different vendor label

treatment effect for each record.

Those limits leave the central observations intact; the transition population is dominated by three-day public intervals, the arrivals are bursty, the vendors are dispersed, and most records land in network/security/control or enterprise/web platforms. Each finding comes directly from the fixed catalog population. The implementation claims are framed as design responses to those observed properties and do not measure agency performance.

10. CONCLUSION

BOD 26-04 turns exploited-vulnerability response into a fast information-management problem: among the 39 KEVs added from 11 June through 31 July 2026, 34 carried a three-day public due-date interval, five carried 14 days, the median interval was 3 days and the mean was 4.41. Arrivals peaked at ten in one week, more than twice the eight-bin mean, while no vendor accounted for more than five records. The queue was fast, bursty and distributed.

The technical mix reinforces the same conclusion because network, security, identity and control infrastructure accounted for 51.3% of records, enterprise and web platforms added 41.0%, and authentication/authorization/privilege-control problems formed the largest weakness family even though no single weakness class explained the urgency pattern. Five older CVE identifiers also returned as new operational work after exploitation evidence changed their status. A vulnerability program must therefore preserve history and context at the same time.

The comparison with the 43-day 2026 DBIR median is a scale warning; it does not estimate compliance. Three days is 14.33 times shorter. That gap leaves little room for slow polling, manual asset discovery, ambiguous ownership or closure without verification, which makes a canonical CVE-asset record—carrying source intelligence, internal context, decision inputs, ownership, timestamps and evidence through closure—the practical control object.

Better exploit prediction remains valuable because the literature reports gains from EPSS, KEV integration, context, topology and explainable ranking [13], [19]-[24]. Yet prediction is only the front of the pipeline. Once exploitation is confirmed and the clock becomes short, information quality and workflow latency determine whether priority can become action; the next empirical step should measure those internal state transitions across organizations and test which handoffs consume the three-day window.

REFERENCES

[1] Cybersecurity and Infrastructure Security Agency (CISA), "BOD 26-04: Prioritizing Security Updates Based on Risk," Binding Operational Directive, June 10, 2026.

[2] C. Butera and J. Spring, "Patch Smarter, Not Harder," Cybersecurity and Infrastructure Security Agency, June 10, 2026.

[3] Verizon, 2026 Data Breach Investigations Report, Verizon Business, 2026.

[4] Cybersecurity and Infrastructure Security Agency (CISA), "Known Exploited Vulnerabilities Catalog," study snapshot restricted to records added through July 31, 2026.

[5] Cybersecurity and Infrastructure Security Agency (CISA), "Known Exploited Vulnerabilities: Catalog Inclusion Criteria," 2026.

[6] Cybersecurity and Infrastructure Security Agency (CISA), FY 2025 Inspector General FISMA Metrics Evaluation Guide, 2025.

[7] M. Bozorgi, L. K. Saul, S. Savage, and G. M. Voelker, "Beyond Heuristics: Learning to Classify Vulnerabilities and Predict Exploits," in Proc. 16th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, pp. 105-114, 2010.

[8] L. Allodi and F. Massacci, "Comparing Vulnerability Severity and Exploits Using Case-Control Studies," ACM Transactions on Information and System Security, vol. 17, no. 1, Article 1, pp. 1-20, 2014.

[9] C. Sabottke, O. Suciu, and T. Dumitraș, "Vulnerability Disclosure in the Age of Social Media: Exploiting Twitter for Predicting Real-World Exploits," in Proc. 24th USENIX Security Symposium, pp. 1041-1056, 2015.

[10] B. L. Bullough, A. K. Yanchenko, C. L. Smith, and J. R. Zipkin, "Predicting Exploitation of Disclosed Software Vulnerabilities Using Open-Source Data," in Proc. 3rd ACM Int. Workshop on Security and Privacy Analytics, 2017.

[11] J. Jacobs, S. Romanosky, I. Adjerid, and W. Baker, "Improving Vulnerability Remediation Through Better Exploit Prediction," Journal of Cybersecurity, vol. 6, no. 1, tyaa015, 2020.

[12] J. Yin, M. Tang, J. Cao, and H. Wang, "Apply Transfer Learning to Cybersecurity: Predicting Exploitability of Vulnerabilities by Description," Knowledge-Based Systems, vol. 210, Article 106529, 2020.

[13] J. Jacobs, S. Romanosky, B. Edwards, I. Adjerid, and M. Roytman, "Exploit Prediction Scoring System (EPSS)," Digital Threats: Research and Practice, vol. 2, no. 3, Article 20, pp. 1-17, 2021.

[14] J. Spring, A. D. Householder, E. Hatleback, A. Manion, M. Oliver, V. S. Sarvepalli, L. Tyzenhaus, and C. G. Yarbrough, Prioritizing Vulnerability Response: A Stakeholder-Specific Vulnerability Categorization (Version 2.0), Software Engineering Institute, Carnegie Mellon University, Apr. 30, 2021.

[15] O. Suciu, C. Nelson, Z. Lyu, T. Bao, and T. Dumitraș, "Expected Exploitability: Predicting the Development of Functional Vulnerability Exploits," in Proc. 31st USENIX Security Symposium, pp. 377-394, 2022.

[16] Y. Roumani, "The Diffusion of Malicious Content on Twitter and Its Impact on Security," Information & Management, vol. 61, no. 5, Article 103971, 2024.

[17] F. R. Parente, E. B. Rodrigues, and C. L. C. Mattos, "FRAPE: A Framework for Risk Assessment, Prioritization and Explainability of Vulnerabilities in Cybersecurity," Journal of Information Security and Applications, Article 103971, 2025.

[18] P. Mell and J. M. Spring, Likely Exploited Vulnerabilities: A Proposed Metric for Vulnerability Exploitation Probability, NIST Cybersecurity White Paper 41, May 19, 2025.

[19] N. Shimizu and M. Hashimoto, "Vulnerability Management Chaining: An Integrated Framework for Efficient Cybersecurity Risk Prioritization," IEEE Access, vol. 14, pp. 31407-31424, 2026.

[20] F. Kausar, L. Batiste, A. Muallem, and S. Hussain, "Integrated Risk Scoring and Exploit Prediction for Cyber-Physical Power System Vulnerabilities," Energy Informatics, vol. 9, Article 32, 2026.

[21] I. Tiță, L.-I. Corățu, M. C. Cujbă, and N. Țăpuș, "Topology-Aware Vulnerability Prioritization on Automated Attack Graphs from Infrastructure-as-Code," Future Internet, vol. 18, no. 6, Article 283, 2026.

[22] G. Siewruk and T. Bondaruk, "Severity Is Not Exploitability: A Simple Taxonomy and Context Score for Better CVE Triage," International Journal of Electronics and Telecommunications, vol. 72, no. 1, pp. 1-7, 2026.

[23] S. Saklani, A. Kalia, and S. Sood, "A Hybrid CNN-LSTM Approach for Predicting Software Vulnerability Exploitability," Indian Journal of Science and Technology, vol. 19, no. 12, pp. 820-829, 2026.

[24] N. Das, "An Explainable Machine-Learning Framework for Prioritizing Cybersecurity Vulnerabilities in the Post-Mythos Era," Journal of Computer Science and Technology Studies, vol. 8, no. 8, pp. 213-221, 2026.

APPENDIX A. COMPLETE KEV TRANSITION POPULATION (11 JUNE-31 JULY 2026)

The table preserves the public catalog labels used in the analysis. Added and due dates are shown without the year (all dates are 2026). The interval is calculated as dueDate minus dateAdded in calendar days.

Table 15. Complete KEV transition population (11 June–31 July 2026)

CVE	Vendor	Product	Catalog vulnerability name / weakness label	Added	Due	Days
CVE-2026-20316	Cisco	Secure Firewall Management Center (FMC)	Use of hard-coded password	29 Jul	1 Aug	3
CVE-2025-68686	Fortinet	FortiOS	Exposure of sensitive information	27 Jul	10 Aug	14
CVE-2026-16812	Arista	VeloCloud Orchestrator	OS command injection	27 Jul	30 Jul	3
CVE-2026-16232	Check Point	SmartConsole	Improper authentication	22 Jul	25 Jul	3
CVE-2026-50522	Microsoft	SharePoint	Deserialization of untrusted data	22 Jul	25 Jul	3
CVE-2026-60137	WordPress	Core	SQL injection	21 Jul	4 Aug	14
CVE-2026-63030	WordPress	Core	Interpretation conflict	21 Jul	24 Jul	3
CVE-2026-0770	Langflow	Langflow	Inclusion of functionality from untrusted control sphere	21 Jul	24 Jul	3
CVE-2021-27137	DD-WRT	DD-WRT	Stack-based buffer overflow	21 Jul	24 Jul	3
CVE-2026-58644	Microsoft	SharePoint	Deserialization of untrusted data	16 Jul	19 Jul	3
CVE-2026-25089	Fortinet	FortiSandbox	OS command injection	16 Jul	19 Jul	3
CVE-2026-39808	Fortinet	FortiSandbox	OS command injection	16 Jul	19 Jul	3
CVE-2026-46817	Oracle	E-Business Suite	Improper privilege management	15 Jul	18 Jul	3
CVE-2023-4346	KNX Association	KNX Protocol Connection Authorization Option 1	Overly restrictive account lockout mechanism	15 Jul	29 Jul	14
CVE-2026-56155	Microsoft	Active Directory Federation Services	Insufficient granularity of access control	14 Jul	28 Jul	14
CVE-2026-56164	Microsoft	SharePoint Server	Missing authentication for critical function	14 Jul	17 Jul	3
CVE-2026-15409	SonicWall	SMA1000 Appliances	Server-side request forgery	14 Jul	17 Jul	3
CVE-2026-15410	SonicWall	SMA1000 Appliances	Code injection	14 Jul	17 Jul	3
CVE-2008-4128	Cisco	IOS	Cross-site request forgery	13 Jul	16 Jul	3
CVE-2026-56291	Balbooa	Forms	Unrestricted upload of dangerous file type	10 Jul	13 Jul	3
CVE-2026-48939	iCagenda	iCagenda	Unrestricted upload of dangerous file type	10 Jul	13 Jul	3
CVE-2026-48908	JoomShaper	SP Page Builder	Unrestricted upload of dangerous file type	7 Jul	10 Jul	3

Continued on next page

Table 16 (continued)

CVE	Vendor	Product	Catalog vulnerability name / weakness label	Added	Due	Days
CVE-2026-55255	Langflow	Langflow	Authorization bypass through user-controlled key	7 Jul	10 Jul	3
CVE-2026-56290	Joomlack	Page Builder	Improper access control	7 Jul	10 Jul	3
CVE-2026-48282	Adobe	ColdFusion	Path traversal	7 Jul	10 Jul	3
CVE-2026-45659	Microsoft	SharePoint Server	Deserialization of untrusted data	1 Jul	4 Jul	3
CVE-2026-48558	SimpleHelp	SimpleHelp	Authentication bypass	29 Jun	2 Jul	3
CVE-2026-12569	PTC	Windchill and FlexPLM	Improper input validation	25 Jun	28 Jun	3
CVE-2026-20230	Cisco	Unified Communications Manager	Server-side request forgery	25 Jun	28 Jun	3
CVE-2025-67038	Lantronix	EDS5000	Code injection	23 Jun	26 Jun	3
CVE-2026-34910	Ubiquiti	UniFi OS	Improper input validation	23 Jun	26 Jun	3
CVE-2026-34909	Ubiquiti	UniFi OS	Path traversal	23 Jun	26 Jun	3
CVE-2026-34908	Ubiquiti	UniFi OS	Improper access control	23 Jun	26 Jun	3
CVE-2026-20253	Splunk	Enterprise	Missing authentication for critical function	18 Jun	21 Jun	3
CVE-2026-48907	Widget Factory	Joomla Content Editor	Improper access control	16 Jun	19 Jun	3
CVE-2026-54420	LiteSpeed	cPanel Plugin	UNIX symlink following	15 Jun	18 Jun	3
CVE-2026-20262	Cisco	Catalyst SD-WAN Manager	Directory/path traversal	15 Jun	29 Jun	14
CVE-2026-35273	Oracle	PeopleSoft Enterprise PeopleTools	Missing authentication for critical function	12 Jun	15 Jun	3
CVE-2026-10520	Ivanti	Sentry	OS command injection	11 Jun	14 Jun	3