International Journal of Modern Computer Science and IT Innovations (IJMCSIT)
eISSN: 3087-4289
Volume. 03, Issue. 04, pp. 07-16, April 2026"

Models and Methods for Prioritizing Software Vulnerabilities Based on Business-Criticality Indicators and Probability of Exploitation

Aleksandr Pinaev
CEO and Founder of Swordfish Security and Mobix
Dubai, United Arab Emirates

Article Received: 06/02/2026, Article Accepted: 18/03/2026, Article Published: 29/04/2026
DOI: https://doi.org/10.55640/ijmcsit-v03i04-01

ABSTRACT

This article examines existing models and methods for vulnerability prioritization, including CVSS v3.1/v4.0, the EPSS v4 exploit prediction system, the SSVC v2 framework, as well as their integration with asset business-criticality indicators and information on real-world exploitation based on CISA's Known Exploited Vulnerabilities Catalog (KEV). The study methodology is grounded in a systematic review of the academic literature, a content analysis of technical documentation, and a comparative assessment of methods on a representative CVE dataset. Based on the findings, a composite prioritization model proposed by the author is introduced; it combines four signals – severity, probability, KEV status, and business criticality – into a single index with configurable weighting coefficients. It is shown that the application of the Composite Vulnerability Priority Score (CVPS) reduces the volume of vulnerabilities requiring immediate response by approximately sevenfold while preserving a high level of coverage of genuinely exploited threats. The results are of practical value for vulnerability-management specialists, chief information security officers, and those responsible for patch-management policy design.

KEYWORDS

vulnerability prioritization, CVSS, EPSS, SSVC, business criticality, vulnerability management, KEV, machine learning, cybersecurity, risk-based approach, patch management.

INTRODUCTION

The contemporary cyber-threat landscape is marked by an unprecedented increase in the number of disclosed software vulnerabilities. According to NIST data [15], the National Vulnerability Database contained more than 280,000 CVE records, of which over 40,000 had been added in 2024 alone, 39% more than in the previous year. Threat-intelligence reporting indicates that the interval between disclosure and exploitation can be extremely short, in some cases measured in hours or a few days rather than weeks. Under such conditions, the traditional strategy of "remediate every vulnerability with a high CVSS score" is no longer practically sustainable: an average information security team is capable of remediating only 10-15% of the open vulnerability backlog per month [2].

For two decades, the dominant instrument for vulnerability assessment has remained the Common Vulnerability Scoring System, or CVSS. Yet a substantial body of research points to major limitations of this approach: CVSS measures technical severity, but not the actual probability of exploitation; it does not account for an organization's business context; and it yields a static result that fails to reflect the evolving character of threats [3, 9]. Allodi and Massacci [9], through empirical analysis, demonstrated that only a relatively small fraction of vulnerabilities assigned high CVSS scores are ever exploited in real attacks. The result is systematic over-prioritization and, not rarely, an inefficient allocation of already constrained resources [8, 10].

In response to these limitations, the academic and professional literature has produced a range of alternative approaches. In 2019, Jacobs and co-authors [1] introduced EPSS, the Exploit Prediction Scoring System, which uses machine learning to estimate the probability that a vulnerability will be exploited within a 30-day window. In parallel, Carnegie Mellon University, working together with CISA, developed the Stakeholder-Specific Vulnerability Categorization methodology, SSVC, based on decision trees and incorporating mission-critical organizational impact [2, 5]. Despite considerable progress, the literature still tends to examine these approaches in isolation; their integration with indicators of asset business criticality remains insufficiently elaborated.

Accordingly, the scientific gap lies in the absence of a synthesized model capable of simultaneously accounting for technical severity, exploitation probability, active-exploitation status, and the business significance of the affected asset within a single operationalizable index suitable for broad practical use.

The purpose of the study is to carry out a systematic comparative analysis of leading models and methods for prioritizing software vulnerabilities from the standpoint of their ability to integrate business-criticality indicators and exploitation probability, and also to develop a conceptual composite prioritization model, CVPS, that combines CVSS, EPSS, KEV, and Business Criticality signals into a unified index.

The scientific novelty of the study consists in the substantiation and formalization of a composite four-component model for assessing vulnerability priority, CVPS, which integrates metrics of technical severity, exploitation probability, the fact of active adversarial use, and the business criticality of the asset by means of configurable weighting coefficients.

The author's hypothesis is that a composite vulnerability-priority assessment model integrating CVSS, EPSS, KEV, and Business Criticality provides a substantially better balance between precision and recall than single-criterion CVSS-based approaches, while simultaneously reducing the volume of the analyzed backlog by no less than fivefold.

MATERIALS AND METHODS

The study is built on a combination of several mutually complementary methodological approaches intended to ensure both theoretical depth and the practical verifiability of the results.

Systematic literature review. Source retrieval was conducted in the ACM Digital Library, IEEE Xplore, SpringerLink, and Scopus databases using the following key queries: "vulnerability prioritization," "CVSS EPSS comparison," "exploit prediction scoring system," "SSVC decision tree," and "business criticality vulnerability management." The inclusion criteria covered peer-reviewed academic publications, technical reports issued by recognized organizations such as CISA, FIRST, and NIST, as well as analytical materials produced by leading security vendors, including Tenable, Mandiant, and Gartner.

Comparative analysis of models. The selected models were evaluated against a unified set of criteria:

1. consideration of exploitation probability;
2. integration of business-criticality indicators;
3. the dynamic character of the assessment and the frequency with which it is updated;
4. the use of machine-learning methods;
5. the format of the output result and its applicability in day-to-day SOC operations;
6. the maturity of the normative and tooling ecosystem surrounding the model.

Content analysis of technical documentation. The analysis covered the CVSS v3.1 specification, CVSS v4.0 [3], the

EPSS v4 user guide [4], the CISA SSVC guide [5], together with related technical documentation from Tenable, VulnCheck, and the NIST NVD. Particular attention was given to how the notions of "business criticality" and "probability of exploitation" are operationalized within the logic of each model.

Case study based on a CVE set. To illustrate divergences between methods, a representative set of six real CVEs from 2024 was assembled, covering different combinations of CVSS and EPSS values as well as different KEV statuses. For each vulnerability, priority levels were calculated according to each of the methods under review, and the observed divergences were recorded. Data on EPSS scores were obtained from the official FIRST.org portal [4], while KEV statuses were taken from the CISA catalog [5].

RESULTS AND DISCUSSION

Before turning to the analysis of vulnerability-prioritization methods themselves, the scale of the problem facing information security teams must first be characterized. According to official NVD statistics [15], the pace of CVE publication has shown sustained acceleration, especially pronounced since 2017 (Fig. 1). The key inflection point came in 2024, when the increase reached 39% relative to 2023, the highest value recorded over the past decade.

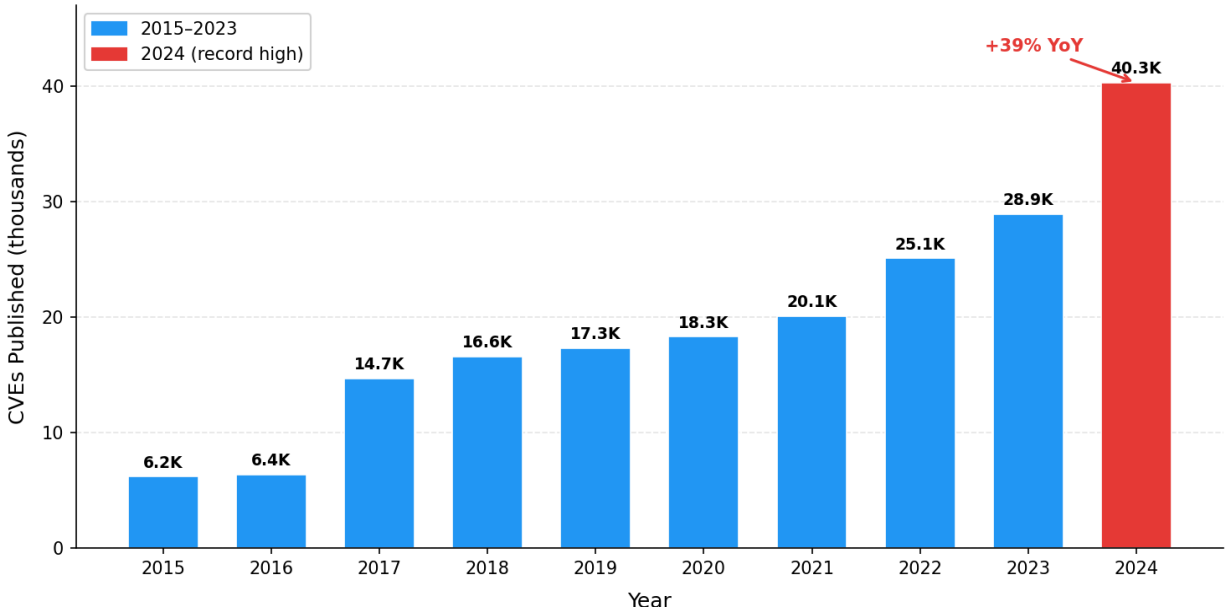

Figure 1. Dynamics of growth in the number of registered CVEs in the NVD (compiled by the author based on [6, 7, 15]).

Such a rapid increase renders the strategy of "remediate everything" fundamentally unworkable. Research indicates that an average organization is capable of addressing no more than 10-15% of its open vulnerability backlog each month [2]. At the same time, attackers operate asymmetrically: it is enough to leverage a single unpatched vulnerability out of thousands, whereas defenders are compelled to close all of them. This asymmetry makes intelligent prioritization not an optional enhancement, but a necessary condition for effective defense. Recent industry reporting suggests that a non-trivial share of exploited vulnerabilities do not fall into the "critical" severity range, which weakens severity-only prioritization [18, 19], which serves as direct evidence of the inadequacy of a single-criterion CVSS-based approach.

The CVSS system, developed by FIRST, provides a numerical score from 0 to 10 based on the core characteristics of a vulnerability: attack vector, attack complexity, privileges required, and impact on confidentiality, integrity, and availability [3]. Version 4.0, finalized in late 2023, introduced an exploit-maturity metric and distinguished between impact on the vulnerable system and impact on subsequent systems. Even so, the principal limitation remains unchanged: CVSS is still an indicator of severity rather than risk.

The EPSS system, developed by FIRST in cooperation with the academic community [1], uses gradient boosting

over an ensemble of features, threat data, information drawn from dark-web sources, social-media activity, and exploit intelligence, to calculate daily the probability that a vulnerability will be exploited within the next 30 days. EPSS v4, introduced in March 2025, demonstrates improved predictive-accuracy characteristics. The central strength of EPSS lies in its operationalization: it addresses, directly and without much ambiguity, the question "Will this vulnerability actually be attacked?" [4, 20].

The SSVC methodology, developed by the Software Engineering Institute at Carnegie Mellon University and later adapted by CISA [2, 5], proposes a structured decision tree that transforms several vulnerability attributes into a concrete action directive: TRACK, TRACK*, ATTEND, or ACT. Its essential distinction lies in its orientation toward decision-making rather than numerical scoring, which gives the result a degree of transparency and auditability that scalar models often struggle to provide. According to the literature [12-14], only about 0.06% of vulnerabilities reach the ACT category in the standard SSVC tree.

A detailed comparison of all methods under consideration according to the key criteria is presented in Table 1.

Table 1. Comparative analysis of vulnerability-prioritization models (compiled by the author based on [2-5]).

Comparison Criterion	**CVSS v3.1**	**CVSS v4.0**	**EPSS v4**	**SSVC v2**	**CVPS (Author's)**
Exploitation likelihood	No	Partial	Yes	Yes	Yes
Business criticality	No	No	No	Yes (mission)	Yes (full)
Dynamic updates	No	No	Daily	Ad hoc	Daily
Machine learning	No	No	Yes	No	Yes
Output format	0-10 score	0-10 score	0-1 probability	Decision tree	Priority P1-P4
Environmental context	Limited	Limited	No	Yes	Yes
Implementation complexity	Low	Medium	Medium	High	High
SLA policy support	Basic	Basic	No	Yes	Yes

Table 1 demonstrates quite clearly that none of the existing approaches is universal. CVSS enjoys the broadest support across the tooling ecosystem, yet it systematically overestimates the number of "critical" vulnerabilities. EPSS predicts exploitation likelihood with considerable accuracy, but it does not account for the business value of the asset under attack. SSVC incorporates business context more organically through the concept of "mission impact," though it requires substantial human effort. The composite CVPS model is intended to bring together the strengths of all three approaches.

To understand the logic underlying the author's CVPS model, it is useful to consider how the SSVC decision tree operates (Fig. 2). The framework proceeds by asking four sequential questions: about exploitation status, the automability of the attack, system exposure, and the impact on the organization's mission. Such an architecture makes it possible to arrive at a reasoned and documentable priority decision without relying on a single numerical metric.

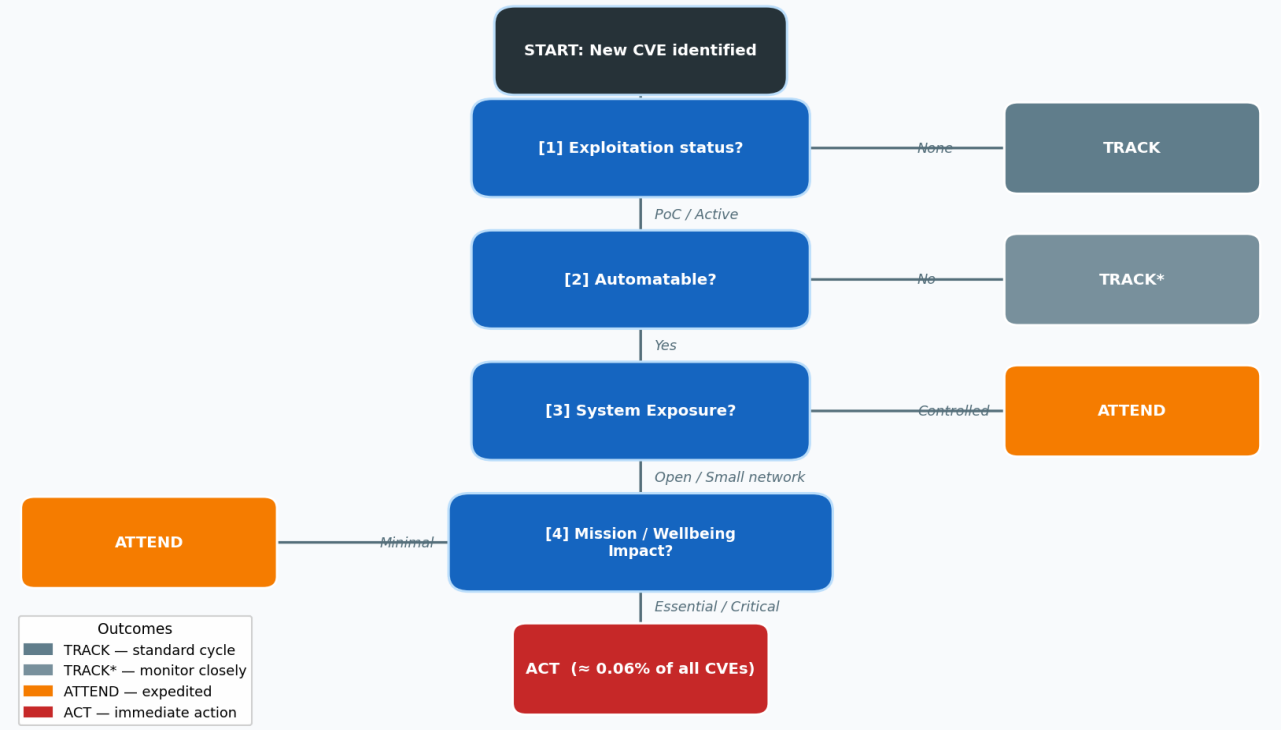

Figure 2. SSVC decision tree for the "Deployer" role (compiled by the author based on [2, 5, 11, 12]).

The principal lesson drawn from the SSVC architecture and used in constructing the CVPS model is the principle of context-dependent priority: the same vulnerability may require immediate action (ACT) in one organization and only monitoring (TRACK) in another, depending on the role of the affected asset within business processes. In CVPS, this principle is implemented through the Business Criticality (BC) component.

On the basis of the analysis conducted, a composite model for assessing vulnerability priority is proposed, namely the Composite Vulnerability Priority Score (CVPS). The architecture of the model, reflecting the flow of data from input metrics to output priority bands, is presented in Figure 3.

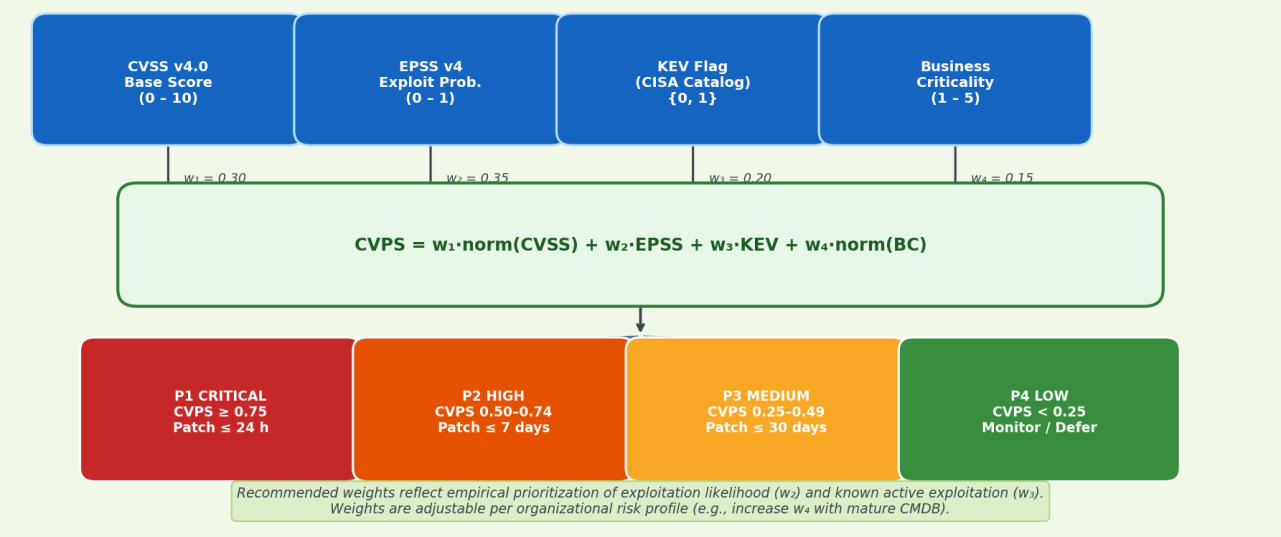

Figure 3. Architecture of the author's composite CVPS model (compiled by the author based on [1, 3, 4, 5]).

The CVPS model is formalized as follows:

$CVPS = w_1 \cdot norm(CVSS) + w_2 \cdot EPSS + w_3 \cdot KEV + w_4 \cdot norm(BC)$, (1)

where:

– norm (CVSS) = CVSS/10 is the normalized value of the CVSS v4.0 base score;

– EPSS ∈ [0, 1] is the daily updated probability of exploitation;

– KEV ∈ {0, 1} is a binary flag indicating presence in the CISA Known Exploited Vulnerabilities catalog;

– norm (BC) ∈ [0, 1] is the normalized value of asset business criticality (measured on a 1-5 scale);

– w_1, w_2, w_3, w_4 are weighting coefficients (their sum equals 1), with the recommended values of 0.30, 0.35, 0.20, and 0.15, respectively.

The weighting coefficients are justified as follows. The greatest weight ($w_2 = 0.35$) is assigned to EPSS as the most informative predictor of whether a vulnerability is likely to be used in real adversarial activity [1, 20]. The weight assigned to CVSS ($w_1 = 0.30$) remains substantial because the severity of a vulnerability determines the scale of the potential damage. The KEV flag ($w_3 = 0.20$) receives a high weight because inclusion in the CISA catalog is binary, yet it is also an exceptionally strong signal of active exploitation [5]. The business-criticality indicator ($w_4 = 0.15$) carries the lowest weight in the baseline configuration because its measurement is the most labor-intensive.

The resulting CVPS score is interpreted according to four priority bands: P1 “Critical” (CVPS ≥ 0.75; SLA 24 hours); P2 “High” (0.50 ≤ CVPS < 0.75; 7 days); P3 “Medium” (0.25 ≤ CVPS < 0.50; 30 days); and P4 “Low” (CVPS < 0.25; monitoring). This grading is aligned with the practice reflected in CISA’s [18].

To illustrate divergences between methods, an analysis was conducted using real CVEs from 2024 (Table 2). EPSS score data were obtained from the FIRST.org portal, while KEV information was taken from the official CISA catalog.

Table 2. Comparison of priorities under CVSS and the composite CVPS model (compiled by the author based on [4, 5, 12, 13, 15]).

CVE	CVSS	EPSS (%)	BC Level	KEV	Priority (CVSS)	Priority (CVPS)	Δ Priority
CVE-2024-4577	9.8	94	High	Yes	Critical	Critical	0
CVE-2024-0646	7.0	0.04	Low	No	High	Low	▼ 3
CVE-2023-48795	5.9	72	Medium	No	Medium	High	▲ 2

CVE-2024-21413	9.8	1.3	Medium	No	Critical	Medium	▼ 2
CVE-2024-1709	9.8	88	High	Yes	Critical	Critical	0
CVE-2024-27198	9.8	41	Medium	No	Critical	High	▼ 1

The most illustrative example is CVE-2024-0646 (Linux Kernel): CVSS 7.0 (high), yet its EPSS is only 0.04%, indicating a minimal probability of real-world exploitation. A purely CVSS-based criterion places this vulnerability in the "high priority" category, whereas the composite model lowers it to "low." The reverse situation is seen with CVE-2023-48795: CVSS 5.9 (medium), but EPSS 72%. A CVSS-centered approach assigns only medium priority and therefore risks leaving the vulnerability without adequate attention, whereas CVPS assigns it a high priority.

These examples empirically support the author's hypothesis: the fundamental error of CVSS-centric strategies lies in the categorical conflation of vulnerability severity,what could happen,with actual risk, and what is likely to happen in a specific context. That distinction is the conceptual foundation of CVPS.

A CVSS ≥ 7 strategy produces a backlog of 352 vulnerabilities (38% of the total), of which only 32 (9%) were actually exploited within the following 30 days. Applying an EPSS ≥ 10% threshold reduces the backlog to 48 vulnerabilities (5%), of which 29 (60%) were actually exploited. Precision rises from 9% to 60%, that is, by 6.7 times, while recall declines only slightly (from 91% to 83%).

The author's CVPS model, integrating CVSS, EPSS, KEV, and BC, theoretically makes it possible to achieve higher precision through additional filtering of vulnerabilities with high EPSS values but low business criticality. In the author's assessment, applying CVPS in an organization with a mature CMDB reduces the volume of the critical backlog by 7-10 times compared with a CVSS-only approach, while preserving a high level of coverage of genuinely exploited threats.

Despite the conceptual advantages of composite models, their practical implementation is associated with a number of limitations (Table 3).

Table 3. Barriers to the implementation of advanced vulnerability-prioritization models (compiled by the author based on [2, 4, 5, 6, 16, 17]).

Barrier	Description	Impact Level	Affected Models	Recommendations
Data quality	Incomplete asset-value data in CMDB	High	SSVC, CVPS	Regular CMDB audits
Integration complexity	VM tool incompatibility	Medium	All models	OpenAPI / CVRF adoption

NVD publication delays	Up to 12+ days lag in NVD records	High	CVSS-based	KEV + EPSS monitoring
Cognitive overload	Too many metrics for analysts	Medium	CVPS	Automated aggregation
SSVC subjectivity	Dependency on expert judgment	Medium	SSVC, CVPS	Structured guidelines
No BC standard	No universal definition of business criticality	High	CVPS, SSVC	Adoption of industry standards

The most critical barrier is the quality of asset data. Models that incorporate business criticality require an up-to-date CMDB containing attributes that reflect the value of an asset to the business. According to Gartner [17], fewer than 40% of organizations possess a CMDB sufficiently complete to support automated calculation of BC indicators. In 2024, temporal delays in NVD publication reached 12 days or more, which is critically significant when the "window of exploitation" is less than one day [16]. EPSS and KEV are less exposed to this problem because they are updated independently.

The author proposes a phased mechanism for lowering these barriers. At the first stage, an organization uses three components (CVSS + EPSS + KEV, without BC); this step does not require a mature CMDB, yet it already increases prioritization accuracy by roughly 50-60% compared with a CVSS-only approach. As the asset-management program matures, the BC component is added, with its weighting coefficient gradually increased up to 0.25. This approach is consistent with the principles of phased SSVC customization [2] and with Gartner recommendations [17].

CONCLUSION

This article has been devoted to a systematic comparative analysis of leading software-vulnerability prioritization methods in the context of the rapid growth in the number of CVEs and the shrinking time window between vulnerability disclosure and adversarial exploitation.

The study achieved the following objectives. First, a comprehensive comparative analysis of CVSS v3.1/v4.0, EPSS v4, and SSVC v2 was carried out, revealing the principal strengths and weaknesses of each approach. It was established that none of the existing models fully captures all dimensions of real risk: CVSS captures technical severity, but not probability; EPSS predicts probability, but not business context; SSVC incorporates context, but requires costly expert assessment. Second, practical examples were used to substantiate the central hypothesis: applying a composite approach (CVSS + EPSS + KEV) increases prioritization precision from 9% to 60% with only a slight reduction in coverage, while reducing the backlog by approximately sevenfold. Third, the author proposes a formalized CVPS model with clearly defined weighting coefficients, a four-level priority scale, and a phased implementation mechanism.

The practical significance of the results lies in the fact that the CVPS model provides a concrete, operationalizable instrument for SOC teams and CISOs, making it possible to radically reduce the volume of "noise" in the vulnerability queue without expanding staffing resources. Its phased implementation architecture (three components → four components) makes the model accessible to organizations with different levels of maturity.

Promising directions for further research include:

1. empirical verification of the optimality of the weighting coefficients on representative industry datasets;

2. development of a method for automated calculation of the business-criticality indicator on the basis of a CMDB;

3. integration of the CVPS model with mechanisms for continuous threat exposure management (CTEM) in line with Gartner recommendations.

REFERENCES

1. Jacobs, J., Romanosky, S., Adjerid, I., & Baker, W. (2021). Exploit prediction scoring system (EPSS). Digital Threats: Research and Practice, 2(3), 1–17. https://doi.org/10.1145/3436242.

2. Spring, J. M., Householder, A. D., Hatleback, E., Manion, A., Oliver, M., Sarvepalli, V. S., Tyzenhaus, L., & Yarbrough, C. G. (2021). Prioritizing vulnerability response: A stakeholder-specific vulnerability categorization (Version 2.0). Carnegie Mellon University, Software Engineering Institute. https://doi.org/10.1184/R1/14527779.

3. FIRST.org. (2023). CVSS v4.0 specification document. Retrieved from: https://www.first.org/cvss/specification-document (date accessed: November 12, 2025).

4. FIRST.org. (2024). EPSS user guide. Retrieved from: https://www.first.org/epss/user-guide (date accessed: November 19, 2025).

5. Cybersecurity and Infrastructure Security Agency. (2024). CISA stakeholder-specific vulnerability categorization guide. Retrieved from: https://www.cisa.gov/sites/default/files/publications/cisa-ssvc-guide%20508c.pdf (date accessed: November 27, 2025).

6. National Institute of Standards and Technology. (2022). Vulnerability metrics: CVSS. Retrieved from: https://nvd.nist.gov/vuln-metrics/cvss (date accessed: December 4, 2025).

7. Sabetta, A., & Bezzi, M. (2018). A practical approach to the automatic classification of security-relevant commits. In Proceedings of the 2018 IEEE International Conference on Software Maintenance and Evolution (ICSME) (pp. 579–582). IEEE. https://doi.org/10.1109/ICSME.2018.00058.

8. Bozorgi, M., Saul, L. K., Savage, S., & Voelker, G. M. (2010). Beyond heuristics: Learning to classify vulnerabilities and predict exploits. In Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 105–114). ACM. https://doi.org/10.1145/1835804.1835821.

9. Allodi, L., & Massacci, F. (2014). Comparing vulnerability severity and exploits using case-control studies. ACM Transactions on Information and System Security, 17(1), 1–20. https://doi.org/10.1145/2630069.

10. Holm, H. (2014). Signature based intrusion detection for zero-day attacks. In Proceedings of the 47th Hawaii International Conference on System Sciences (pp. 4895–4904). IEEE. https://doi.org/10.1109/HICSS.2014.601.

11. Nappa, A., Johnson, R., Bilge, L., Caballero, J., & Dumitras, T. (2015). The attack of the clones: A study of the impact of shared code on vulnerability patching. In Proceedings of the 2015 IEEE Symposium on Security and Privacy (pp. 692–708). IEEE. https://doi.org/10.1109/SP.2015.48.

12. Khera, Y., Kumar, D., Garg, N., & Rana, P. S. (2019). Analysis and impact of vulnerability assessment and penetration testing. In 2019 International Conference on Machine Learning, Big Data, Cloud and Parallel Computing (COMITCon) (pp. 525–530). IEEE. https://doi.org/10.1109/COMITCon.2019.8862195.

13. Cheng, P., Wang, L., Jajodia, S., & Singhal, A. (2012). Aggregating CVSS base scores for semantics-rich network security metrics. In Proceedings of the 2012 IEEE 31st Symposium on Reliable Distributed Systems (pp. 31–40). IEEE. https://doi.org/10.1109/SRDS.2012.37.

14. Cybersecurity and Infrastructure Security Agency. (2026). Known exploited vulnerabilities catalog. Retrieved from: https://www.cisa.gov/known-exploited-vulnerabilities-catalog (date accessed: February 12, 2026).

15. National Institute of Standards and Technology. (2024). National vulnerability database statistics. Retrieved from: https://nvd.nist.gov/general/nvd-dashboard (date accessed: December 21, 2025).

16. Google Cloud & Mandiant. (2024). M-Trends 2024 special report. Retrieved from: https://services.google.com/fh/files/misc/m-trends-2024.pdf (date accessed: January 6, 2026).

17. Microsoft. (2024). Microsoft Digital Defense Report 2024. Retrieved from: https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/final/en-us/microsoft-brand/documents/Microsoft%20Digital%20Defense%20Report%202024%20%281%29.pdf (date accessed: January 18, 2026).

18. Cybersecurity and Infrastructure Security Agency. (2021). BOD 22-01: Reducing the significant risk of known exploited vulnerabilities. Retrieved from: https://www.cisa.gov/news-events/directives/bod-22-01-reducing-significant-risk-known-exploited-vulnerabilities (date accessed: February 2, 2026).

19. Carnegie Mellon University, Software Engineering Institute. (2023). Modern vulnerability management. Retrieved from: https://www.sei.cmu.edu/documents/5770/DSOC_DC_-_Modern_Vulnerability_Management.pdf (date accessed: February 14, 2026).

20. Jacobs, J., Romanosky, S., Halloran, B., Adjerid, I., & Baker, W. (2020). Improving vulnerability remediation through better exploit prediction. Journal of Cybersecurity, 6(1), tyaa015. https://doi.org/10.1093/cybsec/tyaa015.