

Transition from Periodic Security Assessments to Continuous Vulnerability Management Frameworks

 **Kolchin Rustam**
SoftLine PJSC
Almaty, Kazakhstan

Article Received: 22/06/2026, Article Revised: 02/07/2026, Article Accepted: 25/07/2026, Article Published: 06/08/2026

DOI: <https://doi.org/10.55640/irjaet-v03i08-01>

© 2026 Authors retain the copyright of their manuscripts, and all Open Access articles are disseminated under the terms of the [Creative Commons Attribution License 4.0 \(CC-BY\)](#), which licenses unrestricted use, distribution, and reproduction in any medium, provided that the original work is appropriately cited.

ABSTRACT

The article examines the transition from scheduled security assessments to continuous vulnerability management frameworks in enterprise environments with unstable external exposure. Cloud services, SaaS use, short-lived assets, and unmanaged public interfaces reduce the decision value of annual or project-based testing. The novelty of the study lies in combining external attack surface management, continuous penetration testing, vulnerability intelligence, and remediation verification into a single operating model. The aim is to explain why periodic assessment loses completeness when asset states change between review cycles. The method combines source analysis, comparative analysis, conceptual synthesis, and typological classification. The source base covers academic papers, public vulnerability intelligence instruments, official guidance, and industry definitions of external attack surface management. The study identifies three shifts: from snapshot testing to continuous discovery, from severity ranking to exploitation-aware prioritisation, and from automated scanning to expert-verified remediation. The model helps engineering and security teams design measurable programs to reduce exposure.

KEYWORDS

continuous vulnerability management, external attack surface management, continuous penetration testing, vulnerability prioritisation, EPSS

INTRODUCTION

Enterprise security assessment has long relied on scheduled audits, annual penetration tests, compliance scans, and project-based vulnerability reviews. That model served infrastructures with stable asset inventories and slower release cycles. Contemporary enterprise environments operate under different conditions. Cloud services appear and disappear within short time windows, development teams publish new interfaces outside centralised security review, SaaS integrations enlarge the external perimeter, and business units create internet-facing assets that configuration databases fail to register. A periodic test produces useful evidence for a

fixed date. Its decision value drops once the visible attack surface changes between assessment cycles.

The research aim is to develop an analytical model for the transition from periodic security assessment to continuous vulnerability management frameworks. The study pursues three objectives: to define why periodic assessment falls short under dynamic asset conditions, to compare severity-based and exploitation-aware vulnerability prioritisation, and to explain how automation and expert verification can operate within a single remediation cycle.

The novelty of the article lies in integrating external attack surface management, continuous penetration testing, vulnerability intelligence, and patch verification into a single operational sequence. The article treats continuous management as a change in decision logic, in which asset discovery, exploitability estimation, validation, ownership assignment, and remediation confirmation form a recurring engineering process.

The study tests the hypothesis that continuous vulnerability management provides higher operational relevance than periodic assessment when three conditions converge: external assets change between review cycles, prioritisation uses exploitation and asset exposure signals, and remediation closure requires verification after corrective action.

MATERIALS AND METHODOLOGY

The study selected the source corpus through a thematic screening of publications and official materials on enterprise patch management, external attack surface visibility, exploitation probability, vulnerability prioritisation, and penetration testing automation. Official guidance and public intelligence instruments cover patch process design, known exploitation, adversary behaviour mapping, EPSS scoring, and EASM scope (Cybersecurity and Infrastructure Security Agency [CISA], 2021). FIRST documents EPSS as a predictive scoring instrument (Forum of Incident Response and Security Teams [FIRST], n.d.). Gartner Peer Insights (2026), MITRE (2026), and Souppaya and Scarfone (2022) cover EASM scope, adversary behaviour mapping, and enterprise patch management. Recent academic and technical works cover exploit prediction, automated penetration testing architectures, reinforcement-learning-based testing, hybrid manual and automated testing, and probabilistic estimation of exploited vulnerabilities (Jacobs et al., 2023). Lazarov et al. (2025), Liu et al. (2024), Mell and Spring (2025), and Skandylas and Asplund (2025) cover the remaining technical directions. The literature map covers three research directions: asset visibility and shadow IT; risk-based prioritisation beyond severity; and validation of vulnerability management through offensive testing and remediation confirmation.

The methods include comparative analysis of periodic and continuous assessment models, source analysis of vulnerability intelligence frameworks, conceptual synthesis of EASM and continuous penetration testing, typological classification of prioritisation signals, and analytical generalisation to build an implementation model aligned with engineering practice.

RESULTS AND DISCUSSION

A periodic security assessment creates a bounded representation of an environment. Formal scope, repeatable reporting, and audit traceability explain its durability in corporate security governance. The same properties limit its value once asset states change before the next assessment window. NIST SP 800-40 Rev. 4 defines enterprise patch management as the identification, prioritisation, acquisition, installation, and verification of patches across an organisation (Souppaya & Scarfone, 2022). This lifecycle already places verification inside the process. A periodic assessment that stops at detection and reporting leaves the final exposure state unresolved. In a continuous framework, the vulnerability record changes only after the team confirms reduction, documents acceptance, or leaves the exposure open.

External attack surface management extends the first step of this lifecycle. Gartner Peer Insights (2026) defines EASM through the discovery of internet-facing enterprise assets, systems, and exposures that malicious actors could exploit, with attention to unknown assets and systems visible in the public domain. This definition matters for the first objective because many security programs begin vulnerability management from known inventory. If the inventory omits assets, scan coverage inherits the omission. A scanner can evaluate the targets it receives. It leaves an asset model unchanged when that model excludes forgotten domains, temporary cloud endpoints, unmanaged test environments, abandoned subdomains, or third-party-hosted services.

Figure 1 summarises the operating loop that follows from the EASM scope and links discovery with vulnerability management.




Figure 1. Continuous external attack surface management loop, adapted from Gartner Peer Insights (2026)

The loop starts before vulnerability detection, because engineers first need to find, attribute, and classify the asset. It ends after verification, because ticket closure without retesting does not prove exposure reduction. Frequency does not create continuity on its own. Continuity begins when each discovered change re-enters the decision cycle.

EASM treats externally visible systems as observable attack targets, including assets that internal records may omit (Gartner Peer Insights, 2026). MITRE ATT&CK version 19.0 organises adversary behaviour into tactics and techniques across enterprise platforms, including SaaS, IaaS, identity providers, network devices, Windows, Linux, and macOS (MITRE, 2026). The two sources support the same operational conclusion from different sides. EASM identifies what attackers can see. ATT&CK helps defenders interpret how an exposure could support reconnaissance, initial access, credential access, lateral movement, or exfiltration. Asset discovery provides security teams with evidence for attack-path reasoning and more than just a passive inventory list.

Prioritisation changes once the asset model becomes continuous. Traditional vulnerability triage often starts with technical severity. Severity remains useful because

it captures impact characteristics, exploit prerequisites, and technical consequences. It gives no complete answer on which weakness warrants immediate remediation. FIRST describes EPSS as a data-driven machine-learning model that estimates the probability that a published CVE will be exploited in the wild within the next 30 days (FIRST, n.d.). Jacobs et al. (2023) describe EPSS as a freely available scoring system that adapts to new information and improves discrimination between vulnerabilities exploited in the wild and those without observed exploitation. These sources support the second objective because they shift prioritisation from static severity to forecasted adversary interest.

CISA's known exploited vulnerability model adds a different signal. BOD 22-01 established a CISA-managed catalog of known exploited vulnerabilities that pose a risk to the federal enterprise and set remediation expectations for listed vulnerabilities (CISA, 2021). EPSS and KEV express different relations to exploitation. EPSS estimates future probability. KEV records known exploitation evidence. Their joint use helps teams avoid two recurring triage errors: treating all high-severity vulnerabilities as equal and treating predictive probability as confirmed exploitation. A stronger decision model separates confirmed

exploitation, predicted exploitation, asset exposure, business value, and available mitigation.

Mell and Spring (2025) propose the metric "Likely Exploited Vulnerabilities" to estimate the likelihood that attackers have already exploited a vulnerability. Their work addresses a gap between confirmed lists and probabilistic forecasts. Security teams need more than one external indicator because each source captures a different relation between a CVE and adversary activity. A stronger model layers severity, EPSS, KEV, exposure, and asset value according to the decision each signal can support.

A comparison across four sources clarifies the prioritisation problem. NIST SP 800-40 frames patching through identification, prioritisation, installation, and verification (Souppaya & Scarfone, 2022). FIRST provides 30-day exploitation probability through EPSS (FIRST, n.d.). CISA provides an evidence-driven listing of known exploited vulnerabilities through the KEV model (CISA, 2021). Mell and Spring (2025) add a probability-oriented view of likely past exploitation. The combined conclusion for the second objective is operational: vulnerability priority works as a decision state. A CVE on an exposed asset with KEV status follows a different response path than a high-CVSS vulnerability on an isolated system without exploitation evidence. A vulnerability with a rising EPSS value may receive accelerated monitoring or compensating controls before the KEV listing appears.

Continuous penetration testing provides validation for this prioritisation logic. Skandylas and Asplund (2025) formalise automated penetration testing at the architectural level and describe a self-organising architecture for automating security testing of systems composed of hosts and services. Liu et al. (2024) propose a hierarchical reinforcement-learning framework that separates host selection and penetration actions to reduce the action and state spaces in automated testing. These studies support the third objective because automation can expand testing coverage, repeatability, and execution speed. They also show that automation needs an environment model and a controlled action space. Test output still requires interpretation before it becomes a remediation decision.

Lazarov et al. (2025) supply the balancing point through an approach based on adaptable interactive checklists that support manual and automated tests. Their paper states that a purely automated approach would not

identify all vulnerabilities and flaws without manual testing. This finding aligns with continuous vulnerability management because automated discovery and testing maintain breadth across changing assets, while expert review remains necessary when exploit chaining, business logic, authentication flows, identity relationships, partial mitigations, or compensating controls determine the actual exposure.

A comparison across the penetration testing sources produces a practical division of labour. Skandylas and Asplund (2025) address automation in response to the labour-intensive nature of manual testing and the shortage of skilled practitioners. Liu et al. (2024) refine automated testing through hierarchical decision design. Lazarov et al. (2025) preserve manual testing inside a structured test environment. The combined position assigns machine work and human work to different parts of the cycle. Machines handle repeated discovery, known vulnerability checks, regression validation, and breadth scanning. Experts assess uncertain findings, chain exposures, validate exploitability, and convert results into remediation actions that asset owners can execute.

The transition from periodic assessment to continuous management follows three connected shifts. The asset baseline becomes dynamic. Discovery covers unknown and externally visible systems, as well as registered assets. Prioritisation becomes exploitation-aware. Severity remains one input, while KEV, EPSS, likely exploitation, exposure, and asset value shape the response route. Penetration testing becomes a validation service within the lifecycle. It tests whether attackers can use an exposure, whether mitigation works, and whether the organisation has reduced attackable conditions after remediation.

The final report changes status in this model, in periodic assessment, the report often becomes the main deliverable. In continuous vulnerability management, the report becomes a transient record within an operating cycle. The durable output is a maintained exposure baseline supported by evidence of discovery, triage, ownership, remediation, and verification. Continuous management gains value when the baseline absorbs new assets, new vulnerability intelligence, and new remediation results without waiting for the next audit cycle.

A scheduled test asks whether a defined scope contained exploitable weaknesses on a fixed date. A continuous framework asks which exposed assets exist now, which

weaknesses deserve action now, who owns the remediation decision, and whether exposure changed after that action. This difference reshapes governance, tooling, staffing, and metrics.

Figure 2's purpose is to show where the operating model changes as the organisation moves from periodic assessment to continuous vulnerability management.




Figure 2. Transition model from periodic assessment to continuous vulnerability management

In a periodic model, responsibility often moves from testers to asset owners after the report. In a continuous model, responsibility persists through the full lifecycle. Security teams maintain visibility and prioritisation logic. Engineering teams remediate or justify deferral. Management defines risk tolerance. Verification confirms whether the state changed. Without this chain, continuous tools generate continuous noise.

Implementation should proceed in stages: the first stage builds an external asset baseline from domains, IP ranges, certificates, DNS records, cloud-exposed services, repositories, and third-party-hosted interfaces.

The second stage assigns ownership, since unknown ownership blocks remediation regardless of vulnerability severity. The third stage enriches findings with exploitability, known exploitation, exposure, business relevance, and control coverage. The fourth stage routes findings into remediation channels with deadlines tied to risk classes. The fifth stage verifies fixes through rescanning, proof review, or controlled exploitation. The final stage measures exposure trends.

The purpose of Figure 3 is to connect each measurement area to the management decision that the metric supports.




Figure 3. Monitoring architecture for continuous vulnerability management

Figure 3 redirects measurement from activity volume to exposure reduction. A high number of findings can mean weak security, improved visibility, or both. Interpretation depends on ownership, exploitability, remediation speed, and verification evidence. Dashboards that rank teams by raw vulnerability counts create perverse incentives. Teams may suppress findings, overprioritize low-value patching, or close tickets without proof. A better dashboard separates newly discovered assets, exploitable exposures, confirmed fixes, deferred risks, and recurring weaknesses.

A practical model divides findings into urgent, accelerated, scheduled, monitored, and accepted states. Urgent findings require confirmed exploitation or internet-facing exposure with a high likelihood of exploitation and material business impact. Accelerated findings include high predicted exploitation or plausible attack path contribution. Scheduled findings receive normal remediation windows. Monitored findings lack immediate remediation justification but require intelligence tracking. Accepted findings need named ownership, expiry, and compensating controls.

Automated systems can identify exposed services, match CVEs, collect EPSS values, flag KEV membership, and repeat validation checks after a patch. An expert review is necessary when a finding depends on business logic,

chained weaknesses, identity relationships, custom applications, partial mitigations, or ambiguous exposure. The hybrid model should assign expert time to cases where judgment changes the response route.

The transition has organisational limits. Continuous vulnerability management requires asset owners who respond to evidence. It requires engineering workflows that accept security tickets with sufficient technical details for remediation; exception governance, since some enterprise systems cannot receive immediate patches; and a policy for testing intensity, since continuous penetration testing against production-facing systems requires safe execution rules. These limits define the conditions under which the model can operate without destabilising services.

Organisations should introduce a continuous framework as an operating capability. Tooling supplies discovery, correlation, scoring, and workflow automation. The program becomes effective after the organisation defines ownership, verification criteria, escalation thresholds, and acceptable risk states. The mature form of this model uses repeated technical evidence to maintain a current view of external exposure and to guide remediation toward vulnerabilities that adversaries can exploit in practice.

CONCLUSION

The transition from periodic security assessment to continuous vulnerability management follows from the instability of modern enterprise infrastructure. Scheduled audits and penetration tests remain useful for compliance, independent review, and deep technical assessment of selected systems. They lose completeness when asset inventories change faster than the assessment cycle can keep pace. Continuous discovery of internet-facing assets, including unmanaged and forgotten systems, provides security teams with a stronger foundation for vulnerability management in cloud, SaaS, and hybrid environments.

Effective prioritisation combines technical severity with known exploitation, predicted exploitation, asset exposure, asset value, and compensating controls. This approach reduces the risk of spending remediation capacity on technically severe but operationally remote findings while exploited or exposed weaknesses remain unresolved.

Automation supports scale, repeatability, and speed, yet it cannot replace expert validation. The strongest model assigns automated systems to discovery, correlation, regression testing, and routine verification. Specialists evaluate the ambiguity of exploitability, chained weaknesses, business logic, and remediation adequacy. The study confirms the hypothesis: continuous vulnerability management provides higher operational relevance than periodic assessment when discovery is dynamic, prioritisation is exploitation-aware, and remediation closure requires verification.

REFERENCES

1. Cybersecurity and Infrastructure Security Agency. (2021, November 3). Binding Operational Directive 22-01: Reducing the significant risk of known exploited vulnerabilities. <https://www.cisa.gov/news-events/directives/bod-22-01-reducing-significant-risk-known-exploited-vulnerabilities>
2. Forum of Incident Response and Security Teams. (n.d.). Exploit Prediction Scoring System (EPSS). Retrieved May 4, 2026, from <https://www.first.org/epss/>
3. Gartner Peer Insights. (2026). External Attack Surface Management reviews and ratings. Retrieved May 4, 2026, from <https://www.gartner.com/reviews/market/external-attack-surface-management>
4. Jacobs, J., Romanosky, S., Suciu, O., Edwards, B., & Sarabi, A. (2023). Enhancing vulnerability prioritization: Data-driven exploit predictions with community-driven insights. 2023 IEEE European Symposium on Security and Privacy Workshops, 194-206. <https://doi.org/10.1109/EuroSPW59978.2023.00027>
5. Lazarov, W., Seda, P., Martinasek, Z., & Kummel, R. (2025). Penterep: Comprehensive penetration testing with adaptable interactive checklists. *Computers & Security*, 154, Article 104399. <https://doi.org/10.1016/j.cose.2025.104399>
6. Liu, H., Liu, C., Wu, X., Qu, Y., & Liu, H. (2024). An automated penetration testing framework based on hierarchical reinforcement learning. *Electronics*, 13(21), Article 4311. <https://doi.org/10.3390/electronics13214311>
7. Mell, P., & Spring, J. M. (2025). Likely exploited vulnerabilities: A proposed metric for vulnerability exploitation probability (NIST CSWP 41). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.CSWP.41>
8. MITRE. (2026). MITRE ATT&CK version history and Enterprise Matrix. Retrieved May 4, 2026, from <https://attack.mitre.org/resources/versions/>
9. Skandylas, C., & Asplund, M. (2025). Automated penetration testing: Formalization and realization. *Computers & Security*, 155, Article 104454. <https://doi.org/10.1016/j.cose.2025.104454>
10. Souppaya, M., & Scarfone, K. (2022). Guide to enterprise patch management planning: Preventive maintenance for technology (NIST SP 800-40 Rev. 4). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-40r4>