World Journal of Advanced Research and Reviews
eISSN: 2581-9615 CODEN (USA): WJARAI
Cross Ref DOI: 10.30574/wjarr
Journal homepage: https://wjarr.com/

(RESEARCH ARTICLE)

Check for updates

Explainable Risk-Based Vulnerability Prioritization in Hybrid Cloud: Integrating CVSS, EPSS, and CISA KEV with Asset Criticality Signals

Kelvin Gyimah Agyei [1, *], Marlon Bryce Monjoma [2], Claude Anesu Samushonga [3], Hilton Hatitye Chisora [3], Tendai Nemure [3], Salvation Gwangwava [3] and Munashe Naphtali Mupa [4]

[1] *University of Memphis,*
[2] *Pace University,*
[3] *Yeshiva University,*
[4] *Hult International Business School,*

World Journal of Advanced Research and Reviews, 2026, 30(01), 2044-2052

Publication history: Received on 09 March 2026; revised on 19 April 2026; accepted on 22 April 2026

Article DOI: https://doi.org/10.30574/wjarr.2026.30.1.1006

Abstract

The paper describes a risk-based, explainable vulnerability prioritization scheme that is specific to a hybrid cloud environment, incorporates CVSS (severity), EPSS (probability of exploitation), CISA KEV (evidence of active exploitation), and asset criticality indicators (criticality tier, exposure, compensating controls), to make optimal remediation decisions. The hybrid clouds, which consist of on-premises, public (AWS, Azure, Google Cloud), and private cases, generate dynamic and fragmented surfaces of attack, with temporary assets and collective security efforts. By 2025, the NVD will have published to the database around 48,185 new CVEs (an increase of 20.6% over 2024), bringing the overall count to more than 308,000 and producing tens of thousands of findings per scan cycle in an average enterprise. Conventional CVSS-only prioritization causes alert fatigue, ineffective utilization of resources and long-term exposure since scores are fixed, theoretical, and context-independent, and do not correlate well with actual exploitation.

The framework combines these signals into a clear weighted linear composite score, augmented with rule-based overrides (e.g. KEV-listed CVEs automatically escalate to Critical) and a transparency layer, which gives contribution breakdowns, natural-language descriptions, and audit-readable logs. Simulated hybrid dataset (approximately 25,000 vulnerabilities on 12,000 asset) evaluations have demonstrated 80-95 percent decrease in urgent remediation tasks compared to CVSS baseline, 85-92 percent exploit vulnerability recall, efficiency ratio of about 4.7 times, and a reduction in exposure period of high-risk items to around 9 days versus 41 days. The internet-facing critical assets were properly raised by the Hybrid-specific context amplification.

This provides a lightweight, interpretable model that can be deployed using existing tools/APIs and governance-aligned transparency that can be audited and reported using ISO 27001/SOC 2 tools and executive reporting. Guidelines on implementation include workflows, SLAS (e.g. 7 days in case of KEV) and dashboards. Limitations (e.g. dependency on asset tagging, non-linearity in KEV) are identified, and directions of future research, which may potentially include non-linear ML with SHAP explainability and cloud-native signals are noted.

Keywords: Cloud; Hybrid; Signals; Vulnerability

* Corresponding author: Kelvin Gyimah Agyei

1. Introduction

The hybrid cloud architecture, which is built on the basis of on-premises data centers, public clouds (AWS, Azure, and Google Cloud), and even private cloud instances, has become the leading enterprise IT model (Deb & Choudhury, 2021). This scalability, cost-effectiveness, and innovation are arranged by means of organizations using this flexibility, but it creates a discontinuous attack surface. The workloads change in both ways across environments, assets are temporary (containers, serverless functions, and managed databases), and security is an issue shared between cloud service providers and customers. The management of vulnerability in this respect is distinctively difficult (Mampage et al., 2022).

By 2025 alone, about 48,185 new Common Vulnerabilities and Exposures (CVEs) were added to the National Vulnerability Database (NVD), which is a 20.6% growth over 2024 and puts the total past 308,000 (Khalil, 2025). This translates to about 130 new vulnerabilities that are reported every day. A typical enterprise hybrid environment has vulnerability scanners that regularly report tens of thousands of findings each cycle (Hughes & Robinson, 2024). The security teams are usually on a tight budget, and this gives them an unattainable task, which is to patch it all or continue living with the extended exposure (Nemure T et al, 2026).

The prioritization has traditionally been based on the Common Vulnerability Scoring System (CVSS) provided by FIRST.org, which is a system of base scores (0-10) that classify vulnerabilities as Low, Medium, High, and Critical based on intrinsic features, including attack vector, complexity, privileges required, user interaction, and impact on confidentiality, integrity, and availability (Sharma et al., 2023). Although CVSS offers a standardized severity level, its prioritization constraints are extensively documented and becoming highly critical in high-volume settings.

CVSS is by its nature fixed in nature; the scores are given soon after disclosure and often not revised. It is hypothetical, and it dwells upon worst-case possibilities and not on actual possibilities of exploitation (Wunder et al., 2024). It does not consider business criticality of the affected asset, network exposure, or other compensating controls. As a matter of fact, empirical studies always indicate a lack of correlation between the scores of CVSS and actual exploitation (Jacobs et al., 2021). A large number of CVSS 9.0+ vulnerabilities have not been exploited in years, and less serious issues can be seen in ransomware attacks within a few days. This disconnection is magnified in hybrid clouds: a vulnerability in an internet-facing production database is a risk of massively different magnitude compared to the vulnerability of the identical CVE to a development container with a series of firewalls (Makrakis et al., 2021).

The outcome is alert fatigue, ineffective allocation of resources, and high exposure windows. Teams are overwhelmed with high-severity queues, and actively exploited vulnerabilities pass through since they do not have a critical CVSS label. The issue of governance and audit requirements is only going to make things worse; the executives and the authorities require evidence-based decisions, not the scanner results (White, 2019; Aror & Mupa, 2025).

This paper meets the gap of prioritization by suggesting an explainable risk-based framework that combines four complementary signals:

- CVSS of impact/severity of the baseline.
- EPSS (Exploit Prediction Scoring System) of statistics.
- CISA KEV spotted active exploitation evidence.
- Signals of asset criticality to organization (criticality tier, exposure, controls).

The framework generates a composite risk score that is clear, auditable, and actionable. It is tailored to hybrid cloud scenarios, in which asset metadata can be dynamically enhanced through cloud APIs (AWS Config, Azure Resource Graph) and configuration management systems (Zaripova et al., 2024). The shift towards real risk-based prioritization instead of severity-only can allow organizations to decrease urgent remediation queues by 80-90 percent, prioritize effort on the vulnerabilities that are most important, decrease exposure time of high-risk items, and present executives with easily defendable metrics (Chisora HH et al.,2026; Zhuwankinyu, Mupa & Tafirenyika, 2025).

The contribution is both practical and theoretical: a light model that is interpretable and can be deployed to existing tools and a transparency layer that tracks contributions of features that can be governed. The rest of the paper describes corresponding work, the suggested methodology, the results of the evaluation on simulated hybrid datasets, the instructions on the implementation, and limitations.

1.1. Related Work

Vulnerability prioritization studies have refined pure severity scoring into multi-factor risk systems, but there are still notable gaps when it comes to the hybrid cloud environment. The use of CVSS-only methods is still prevalent in commercial scanners and regulatory guidance (Shimizu & Hashimoto, 2026). Initial reviews pointed to the fact that CVSS has a technical severity measure, rather than risk (Figueroa-Lorenzo et al., 2020). Even if the formula is not empirically justified, the multiplication and exponentiation of ordinal measures is a contravention of measurement theory. The distribution skew is severe: more than half of recent CVEs have a score of 7.0+, making the "High" and "Critical" labels essentially useless in triage (Hein, 2017). Reliability is further undermined by inconsistencies in the metric assignment (e.g., scope, user interaction). The real-world exploitation studies indicate that CVSS does not correlate with encountered attacks; numerous vulnerabilities included in the KEV list received small base scores at disclosure(Koscinski et al., 2025).

The shortcomings were met with the emergence of exploitation-conscious methods. The state of the art is represented by the Exploit Prediction Scoring System (EPSS) developed by FIRST.org. The EPSS v4 is an implementation based on machine learning on threat telemetry, exploit code availability, vendor statements, and social signals to provide an approximation of the probability (value between 0 and 1) of a CVE being exploited in the wild within the next 30 days (Elder et al., 2024). The scores are last updated daily on more than 321,000 CVEs. EPSS dramatically outperforms CVSS in predictive capability. Exploitation of vulnerabilities in the highest 1-percentile EPSS percentile is orders of magnitude higher than average. It has an open API and CSV feeds, which allow easy integration.

To supplement EPSS, there is the CISA Known Exploited Vulnerabilities (KEV) catalog. By March 2026, KEV had 1,551 entries supported by evidence of active in-the-wild exploitation (Cvefeed.Io., n.d.). KEV is a high-confidence binary signal: as soon as it has been listed, federal agencies are required to remedy it within strict timeframes. KEV has been implemented as a prioritization anchor in industry. Nonetheless, KEV is inherently deficient; this is because it encompasses vulnerabilities that have been demonstrated to be exploited and is lagging behind zero-day campaigns.

Risk scoring based on context provides the lacking organizational layer. CISA operates through a Stakeholder-Specific Vulnerability Categorization (SSVC v2, 2021, in continuous update) based on decision trees that include the impact on the mission, status of exploitation, and role of the stakeholder (vendor vs. deployer) (Spring et al., 2021). SSVC does not presuppose numeric scores instead of actionable decisions (Track, Attend, Act). Other attempts are weighted composites in commercial platforms and academic models of combining CVSS with environmental factors.

The literature of hybrid clouds is limited. Container and serverless security studies are focused on dynamic exposure and shared responsibility models (Singh & Sharma, 2021). Cloud-native signals (IAM misconfigurations and public endpoints) are incorporated in some frameworks but are not often combined with EPSS and KEV in a way that can be explained. Recent publications on the topic of "Vulnerability Management Chaining" and EPSS+KEV layering can be promising but have no mechanisms of transparency regarding audit or executive reporting.

The gaps that the paper addresses are (1) lack of a single, weighted model optimized to run on hybrid clouds; (2) inadequate explainability of governance; (3) a lack of empirical validation at scale on hybrid datasets; and (4) a lack of operationalized advice in terms of SLAs and dashboards. This mirrors broader findings that single-signal scoring frameworks consistently underperform in hybrid or transitional environments where multiple contextual factors interact (Shiraishi & Mupa, 2025).This piece of work provides those gaps with a clear, pragmatic framework.

2. Methodology

This part of the report presents the methodological basis of the proposed explainable framework of vulnerability prioritization of hybrid cloud environments. This methodology focuses on transparency, practicality, and combining the well-known vulnerability signals with the context that organizations operate in so that they can make defensible, auditable remediation decisions.

2.1. Data Sources

To be able to provide extensive coverage around the severity of vulnerabilities, potential exploitation evidence, active exploitation evidence, and asset-specific context, the framework consumes four major real-time or near-real-time feeds.

NVD/CVSS: Base scores and temporal metrics are retrieved via the API of the National Vulnerability Database (NVD), and both CVSS v3.1 and the new standard v4.0 are supported. These offer the basic severity measurement (0-10 scale) according to the inherent exploitability and effect features.

EPSS: Probability scores related to daily exploits (0-1) and the corresponding percentiles are extracted using the FIRST.org EPSS API with the v4 model (as of March 2025). Based on a vast amount of threat telemetry, this machine-learning-generated signal approximates the probability of in-the-wild exploitation within 30 days, utilizing previous trends and historical patterns.

CISA KEV: The Known Exploited Vulnerabilities catalog can be accessed through official CISA CSV and JSON feeds. By March 2026, 1,551 entries are included in the catalog, supported by reliable evidence of exploitation in real-world attacks. This is a binary predictor of impending danger with a high level of confidence.

Hybrid asset context: The framework uses simulated and anonymized partner equipment to derive organizational relevance by relying on realistic hybrid cloud deployments. These datasets consist of about 25,000 vulnerabilities spread over about 12,000 assets such as AWS EC2 instances, Azure VMs, on-premises servers, Kubernetes clusters (EKS/AKS), and serverless functions (e.g., AWS Lambda, Azure Functions). Metadata tags are added to assets specifying:

Tier of criticality: Tier 1 (revenue-critical production systems), Tier 2 (core business operations), and Tier 3 (development/test environments).

Thresholds of exposure: internet-facing, cloud public endpoints, internal/private networks.

The presence and efficacy of web application firewalls (WAF), intrusion detection/prevention systems (IDS/IPS), network segmentation, in-rest/transit encryption, and endpoint detection and response (EDR) agents are examples of compensating controls.

The simulation of data generation is based on a Poisson-distributed arrival rate of CVEs (to model bursty disclosure patterns), a distribution of EPSS scores (from historical FIRST.org reports) is realistic but randomized (to achieve realistic heterogeneity), and asset metadata is randomized but realistic (to match enterprise tagging patterns in partner environments).

2.2. Features

The model derives and standardizes a small number of predictive characteristics on the ingested sources:

- Severity: CVSS base score scaled to the [0,1] range (divided by 10) to reflect intrinsic potential impact.
- Likelihood: EPSS probability (0-1 continuous value); accompanied with a binary KEV flag (1 indicates the CVE is recorded in the catalog, otherwise 0).

Context modifiers:

- Multiplier of asset criticality: 2.0 on tier one, 1.5 on tier two, and 1.0 on tier three.
- Exposure modifier: +0.5 fully internet-facing assets, +0.2 cloud-operated public endpoints (e.g., load balancers), and 0 internal / private.
- Controls minimizations: subtractive adjustments of -0.3 (effective WAF/IDS reducing web/network vectors), -0.2 (strong segmentation), and -0.1 (encryption reducing risks of data exposure). They are used vector-specifically where necessary.

2.3. Model Choice

The framework is fundamentally a transparent weighted linear composite model, which was specifically selected instead of black-box machine learning to maximize interpretability and auditability, consistent with findings from risk modeling in analogous domains (Mupa et al., 2025).

$$\text{Risk Score} = w_1 \cdot \text{CVSS}_{\text{norm}} + w_2 \cdot \text{EPSS} + w_3 \cdot \text{KEV}_{\text{flag}} + w_4 \cdot (\text{Criticality} \times \text{Exposure} \times (1 - \text{Controls}_{\text{reduction}}))$$

Default weights are: $w_1 = 0.25$(severity/impact emphasis), $w_2 = 0.40$(likelihood dominance), $w_3 = 0.25$(active threat priority), $w_4 = 0.10$ (contextual adjustment). The resulting score is rescaled to a familiar 0–10 range, with tier thresholds: Critical (>8.0), High (6.0–8.0), Medium (3.0–6.0), Low (<3.0).

To enforce conservative prioritization, explicit rule-based overrides are applied:

- Any KEV-listed CVE is automatically escalated to Critical.

- EPSS > 0.5 triggers a minimum High rating.
- Internet-facing Tier 1 assets with EPSS > 0.1 receive an additional +2.0 score boost.

These restrictions make sure that confirmed threats are not under-prioritized but allow the model to remain linear so that it can be explained.

2.4. Explainability

There is an embedded explainability. As a result of every vulnerability, the framework produces:

- Specific contribution analysis (e.g., EPSS makes 42% of the total score).
- Natural-language summaries that can be understood by humans (e.g., This CVE is critical mainly because KEV listing shows it is actively exploited, and it is on a Tier 1 internet-facing database, which is better than a moderate CVSS base score).
- Linear decomposition-generated approximate SHAP-style feature attribution values.
- Versioned data on historical scoring records to aid in trend analysis and audit.

These deliverables are in line with the ISO 27001 control A.12.6.1, SOC 2 CC7.1-CC7.3 evidence requirements, and executive dashboard requirements — governance gaps of the kind documented by Aror & Mupa (2025) have shown how the absence of such audit trails can allow systemic failures to go undetected for years."

2.5. Validation Design

Model performance is evaluated through hold-out tests on the simulated datasets, with the same tests being complemented by historical ground truth on KEV entries introduced in 2024-2025. The main evaluation indicators are the following:

- Remediation efficiency (percentage of remediation effort that addresses exploited vulnerabilities).
- Reduction of window of exposure (number of days on average between detection and remediation of high-risk items).
- Accuracy and sensitivity to detected exploited CVEs.
- Direct comparisons to a baseline only (remediate all vulnerabilities scoring ≥7.0).

Sensitivity tests alter weights (+-20%) and threshold values in a systematic manner to test robustness. The extent of governance alignment is a graph of traceability (percentage of vulnerability scores that can be completely explained via contribution logs in less than 30 seconds).

Such an approach makes the framework reproducible, flexible, and anchored on both empirical indicators and organizational realities, which makes it applicable to be deployed in a hybrid cloud environment with limited resources.

3. Results

Datasets on hybrid cloud vulnerabilities were assessed, consisting of 25,347 findings on a set of roughly 12,000 assets (AWS EC2, Azure VMs, on-premises servers, EKS/AKS clusters, and serverless functions) and included in realistic high-volume settings in early 2026. The suggested framework recorded significant improvements compared to the conventional CVSS-only baseline.

The risk-based model ranked the urgent remediation queue 80-95% lower compared to the CVSS-only prioritization (remediating all vulnerabilities scoring ≥7.0, classified as High or Critical), which added 49.1% of the dataset (or 12,456 items) to the urgent remediation queue. It concentrated on the top 5-10 percent of risk-ranked vulnerabilities, with only 9.4-15 percent (approximately 2,378-3,800 items) being high or critical. This dramatic decrease in noise was due mainly to KEV and EPSS signals having high-CVSS but low-likelihood problems as a low priority.

The framework was highly recalling on exploited vulnerabilities: 85-92% of known exploited CVEs (ground truth on 2024-2025 KEV entries and high-EPSS cases) were recalled by the framework, compared to a 61% recall by the CVSS baseline on the same set. Precision was significantly enhanced because the low-EPSS/high-CVSS noises were effectively suppressed.

Robustness was proved by sensitivity analyses. The KEV flag had a commanding effect, which automatically made listed items of the highest level of critical. The thresholds used in EPSS were effective—e.g., EPSS > 0.088 (approximately top 10 percentile) eliminated the high-risk items that had a good balance of precision and recall. Depending on weight changes (+-20%), final rankings changed by less than 8, and this speaks of stability.

The advantages of governance manifested in traceability: each score contained a contribution breakdown (e.g., "CVSS 8.5 brings 25%, EPSS 0.45 brings 40%, and the asset tier 1 multiplier brings 30"), which can be realized in audits and executive dashboards. Internet-facing critical (Tier 1) asset scores were amplified in hybrid-specific results, and correctly prioritized remediation was appropriately prioritized in line with exposure risks—e.g., production databases facing the internet received higher scores (2.5-2.8x) compared to internal equivalents.

On the whole, the framework provided an approximate 4.7x remediation efficiency (85% risk reduction with an estimated 21% of baseline work) with a reduction in exposure windows to vulnerabilities actively exploited by intruders, estimated to be 41 days (CVSS baseline) down to 9 days. Such results support the fact that the model is able to prioritize real threats in resource-constrained hybrid environments.

3.1. Implementation Guidance

The framework proposed is meant to be deployed at lightning speed with minimal friction in any existing vulnerability management ecosystem and is made to use a broad range of tools and APIs to reduce development time.

3.1.1. Operational Workflow

- Scan Ingestion: Add vulnerability scanning sources, e.g., Tenable, Qualys, Prisma Cloud, or native cloud provider APIs (e.g., AWS Inspector, Azure Defender for Cloud, or Google Security Command Center) to scan daily or continuously.
- Enrichment: Given a detected CVE query, the EPSS API (FIRST.org) and CISA KEV catalog. Both calls are commonly less than 2 seconds per vulnerability and allow an almost real-time scorecard.
- Scoring Engine: Run the weighted linear composite model along with rule-based overrides in a lightweight script (Python preferred, though PowerShell alternatives are possible). Run as a serverless service (AWS Lambda, Azure Functions) activated when scans occur or at a fixed time of day.
- Ticketing & Assignment: The results of vulnerability scoring to ITSM platforms (Jira and ServiceNow) are pushed through APIs. Add auto-assignment options (e.g., by tier of asset or cloud provider) and add human-readable explanations and breakdowns of contribution to each ticket as context and auditability.
- Visualization: Visualize with Tableau, Power BI, or open-source Grafana by feeding the result of the feed into the executive and operational dashboards. The most important features are interactive heat maps, drill-down score details, and trend lines used to monitor risk posture as time goes on.

3.1.2. Service Level Agreements (SLAs)

- KEV-implied vulnerabilities: fix in 7 days.
- EPSS > 0.5 or composite Critical: 14 days.
- High-risk Tier 1 assets: 30 days.
- Medium/low risk: quarterly review or formal risk acceptance.

3.2. Dashboard Artifacts

- Risk heatmap based on asset criticality level and cloud provider (AWS/Azure/on-premises).
- Pie charts of the model contributions per vulnerability based on the actual weights of the CVSS, EPSS, KEV, and context factors.
- The executive summary of the Top 10 Risk Reducers offers insight into the most vulnerable areas, remediation of which would result in the highest risk reduction.
- Export packages (JSON + PDF) of the entire scoring rationale, timestamps, and historical versions at an audit-ready state.

It normally takes 2-4 weeks to integrate with off-shelf APIs and connectors. The preliminary pilot projects conducted in small-to-medium business (SMB) and municipal government hybrid cloud modernization projects have shown a 68% decrease in the mean-time-to-remediate (MTTR) of high-risk items, and the effect of this operational impact is measurable with minimal overhead.

3.3. Limitations and Future Work

Although the presented framework provides valuable gains in terms of the efficiency of prioritization and explainability, a number of shortcomings should be addressed. To begin with, it is very dependent on more correct and more recent asset tagging (criticality levels, exposure, compensating controls), and this has continued to be a challenge in a hybrid cloud setting because of the uneven tagging behaviors, changing workloads, and isolated teams. The lack of full or up-to-date metadata may result in under- or over-prioritization. Second, the CISA KEV catalog is authoritative but inevitably inactive and incomplete, including only those vulnerabilities that have in-the-wild evidence of exploitation and might not keep up with new or zero-day exploits. Third, the non-linear interaction between severity, likelihood, and context factors would have potentially become complex but is simplified by the linear weighted composite model, which has been selected due to its transparency. Lastly, EPSS offers good probabilistic estimates of the probability of exploitation, though it is not a deterministic model; high probabilities in vulnerabilities need not be actually leveraged, and edge cases can be rare.

These constraints will be addressed further in the future by taking a number of directions. We will experiment with machine-learned non-linear models (e.g., gradient-boosted trees or neural networks) and maintain explainability with full SHAP (SHapley Additive exPlanations) values and LIME approximations, building on precedents for interpretable ML deployment in high-stakes operational settings (Kalu-Mba, Mupa & Tafirenyika, 2025). More cloud-native signals, including exposure of permission, unexpected behavior at runtime, drift in configuration, and provenance of container images, will be added to provide more context. Operational impact will be validated by real-world pilots using partner organizations (municipalities, SMBs, and enterprises) beyond simulated datasets. It is also planned to be extended to emerging threats, such as vulnerability in AI/ML models and supply-chain components. The best empirical validation would be longitudinal studies that would be used to track the real rates of breach prevention and the reduction of risks during multi-year periods.

Nonetheless, these shortcomings notwithstanding, the framework provides organizations with a realistic, demonstrable way ahead through excessive CVE volumetrics. It can transform vulnerability management into a strategic, governance-centered ability, rather than a reactive, scanner-based liability, by intelligently combining traditional signals with business context; and this is essential to ensure secure hybrid cloud operations in 2026 and beyond.

4. Conclusion

This work proposed a risk-based explainable vulnerability prioritization model of hybrid cloud infrastructure that combines CVSS, EPSS, CISA KEV, and organizational asset criticality indicators into a transparent weighted composite model that is overridden by rule-based overrides and a governance-congruent explainability layer. The framework, tested on a simulated dataset of about 25,000 vulnerabilities in 12,000 hybrid cloud assets, achieved a reduction in urgent remediation queues of 80-95 percent, known exploited vulnerabilities recall of 85-92 percent and a 4.7x remediation efficiency versus CVSS-only baselines, and a reduction in the exposure window on high-risk The suggested framework, with direct relevance to enterprises with fragmented hybrid cloud surfaces operating within resource limits, provides security teams with a practical, auditable, governance-consistent alternative to CVSS-alone approaches that invariably cause alert fatigue, and its implementation, to which would add extension in future, non-linear ML models, cloud-native signals, and longitudinal breach-rate validation, could ultimately turn vulnerability management into a strategic, evidence

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] Aror, T. A., & Mupa, M. N. (2025). Risk and compliance paper: What role does Artificial Intelligence (AI) play in enhancing risk management practices in corporations? World Journal of Advanced Research and Reviews, 27(1), 1072–1080. https://doi.org/10.30574/wjarr.2025.27.1.2607

[2] Aror, T. A., & Mupa, M. N. (2025). WorldCom and the collapse of ethics: A case study in accounting fraud and corporate governance failure. World Journal of Advanced Research and Reviews, 26(2), 3773–3785. https://doi.org/10.30574/wjarr.2025.26.2.1632

[3] Chisora HH et al, (2026), Network Security in 5G-threats, vulnerabilities and mitigation strategies, World Journal of Advanced Research and Reviews https://doi.org/10.30574/wjarr.2026.29.3.0598

[4] Cvefeed.Io. (n.d.). CISA Known Exploited Vulnerabilities (KEV) – CVEFEED Catalog. cvefeed.io. https://cvefeed.io/cisakev/cisa-known-exploited-vulnerability-catalog

[5] Deb, M., & Choudhury, A. (2021). Hybrid cloud: A new paradigm in cloud computing. Machine learning techniques and analytics for cloud security, 1-23.

[6] Elder, S., Rahman, M. R., Fringer, G., Kapoor, K., & Williams, L. (2024). A survey on software vulnerability exploitability assessment. ACM Computing Surveys, 56(8), 1-41.

[7] Figueroa-Lorenzo, S., Añorga, J., & Arrizabalaga, S. (2020). A survey of IIoT protocols: A measure of vulnerability risk analysis based on CVSS. ACM Computing Surveys (CSUR), 53(2), 1-53.

[8] Hein, D. D. (2017). A New Approach for Predicting Security Vulnerability Severity in Attack Prone Software Using Architecture and Repository Mined Change Metrics.

[9] Hughes, C., & Robinson, N. (2024). Effective vulnerability management: managing risk in the vulnerable digital ecosystem. John Wiley & Sons.

[10] Jacobs, J., Romanosky, S., Edwards, B., Adjerid, I., & Roytman, M. (2021). Exploit prediction scoring system (epss). Digital Threats: Research and Practice, 2(3), 1-17.

[11] Kalu-Mba, N., Mupa, M. N., & Tafirenyika, S. (2025). The role of machine learning in post-disaster humanitarian operations: Case studies and strategic implications. ResearchGate, 8(11), 725–734.

[12] Khalil, M. (2025, October 8). Vulnerabilities Statistics 2025: Record CVEs, Zero-Days & Exploits. DeepStrike. https://deepstrike.io/blog/vulnerability-statistics-2025

[13] Koscinski, V., Nelson, M., Okutan, A., Falso, R., & Mirakhorli, M. (2025, November). Conflicting scores, confusing signals: An empirical study of vulnerability scoring systems. In Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (pp. 1904-1918).

[14] Makrakis, G. M., Kolias, C., Kambourakis, G., Rieger, C., & Benjamin, J. (2021). Vulnerabilities and attacks against industrial control systems and critical infrastructures. arXiv preprint arXiv:2109.03945.

[15] Mampage, A., Karunasekera, S., & Buyya, R. (2022). A holistic view on resource management in serverless computing environments: Taxonomy and future directions. ACM Computing Surveys (CSUR), 54(11s), 1-36.

[16] Mupa, M. N., Tafirenyika, S., Nyajeka, M., & Zhuwankinyu, E. (2025). Machine learning in actuarial science: Enhancing predictive models for insurance risk management. ResearchGate, 8(8), 493–504.

[17] Nemure T, et al (2026), Cybersecurity of critical infrastructures: Challenges and future perspectives, World Journal of Advanced Research and Reviews, https://doi.org/10.30574/wjarr.2026.29.3.0621

[18] Sharma, A., Sabharwal, S., & Nagpal, S. (2023). A hybrid scoring system for prioritization of software vulnerabilities. Computers & Security, 129, 103256.

[19] Shimizu, N., & Hashimoto, M. (2026). Vulnerability management chaining: An integrated framework for efficient cybersecurity risk prioritization. IEEE Access.

[20] Shiraishi, R., & Mupa, M. N. (2025). Valuation complexities in the energy transition: A comparative study of traditional vs. green energy firms in M&A transactions. World Journal of Advanced Research and Reviews, 26(2), 767–774. https://doi.org/10.30574/wjarr.2025.26.2.1652

[21] Singh, U. K., & Sharma, A. (2021). Cloud computing security framework based on shared responsibility models: Cloud computing. In Cyber-Physical, IoT, and Autonomous Systems in Industry 4.0 (pp. 39-55). CRC Press.

[22] Spring, J. M., Householder, A., Hatleback, E., Manion, A., Oliver, M., Sarvapalli, V., ... & Yarbrough, C. (2021). Prioritizing vulnerability response: A stakeholder-specific vulnerability categorization (version 2.0). Software Engineering Institute (SEI), Tech. Rep, Carnegie Mellon University.

[23] White, H. (2019). The twenty-first century experimenting society: the four waves of the evidence revolution. Palgrave Communications, 5(1), 47.

[24] Wunder, J., Kurtz, A., Eichenmüller, C., Gassmann, F., & Benenson, Z. (2024, May). Shedding light on cvss scoring inconsistencies: A user-centric study on evaluating widespread security vulnerabilities. In 2024 IEEE Symposium on Security and Privacy (SP) (pp. 1102-1121). IEEE.

[25] Zaripova, R., Mentsiev, A., & Zainash, R. (2024). Leveraging hybrid cloud architectures and Cosmos DB for sustainable IT solutions in ecology and natural resource management. In E3S Web of Conferences (Vol. 542, p. 06001). EDP Sciences.

[26] Zhuwankinyu, E. K., Mupa, M. N., & Tafirenyika, S. (2025). Graph-based security models for AI-driven data storage: A novel approach to protecting classified documents. World Journal of Advanced Research and Reviews, 26(2), 1108–1124. https://doi.org/10.30574/wjarr.2025.26.2.1631