VULNERABILITY MANAGEMENT CHAINING: AN INTEGRATED FRAMEWORK FOR EFFICIENT CYBERSECURITY RISK PRIORITIZATION

A PREPRINT

Naoyuki Shimizu

Masaki Hashimoto
Faculty of Engineering and Design
Kagawa University, Japan

June 9, 2025

ABSTRACT

Cybersecurity teams face an overwhelming vulnerability crisis: with 25,000+ new CVEs disclosed annually, traditional CVSS-based prioritization requires addressing approximately 57% of all vulnerabilities while correctly identifying only 20% of those actually exploited. We propose Vulnerability Management Chaining, an integrated decision tree framework combining historical exploitation evidence (KEV), predictive threat modeling (EPSS), and technical impact assessment (CVSS) to transform vulnerability management from reactive patching to strategic threat-driven prioritization. Experimental validation using 28,377 real-world vulnerabilities demonstrates approximately 14-18 fold efficiency improvements while maintaining 85+ percent coverage of actual threats. Organizations can reduce urgent remediation workload by approximately 95% (from 16,000 to 850 vulnerabilities). The integration identifies 57 additional exploited vulnerabilities that neither KEV nor EPSS captures individually. Our framework uses exclusively open-source data, democratizing advanced vulnerability management regardless of budget or expertise. This research establishes the first empirically validated methodology for systematic vulnerability management integration, with immediate applicability across diverse organizational contexts.

Keywords Vulnerability Management · CVSS · EPSS · KEV · Threat Intelligence · Cybersecurity Risk Assessment · Decision Trees · Security Prioritization

1 Introduction

The exponential growth of cybersecurity threats and the increasing complexity of modern IT infrastructures have made effective vulnerability management one of the most critical challenges facing organizations today. As software vulnerabilities continue to be discovered and disclosed at an unprecedented rate, security teams face challenges in allocating their limited resources effectively to address the most critical risks.

1.1 The Vulnerability Management Challenge

Modern organizations face an increasingly difficult vulnerability landscape. The number of Common Vulnerabilities and Exposures (CVE) published annually has grown substantially, with data from the National Vulnerability Database showing consistent year-over-year increases [1]. This growth trend shows no signs of slowing, placing enormous pressure on security teams to evaluate, prioritize, and remediate an ever-expanding attack surface.

Compounding this challenge is the fundamental asymmetry between attackers and defenders. While attackers need only to find and exploit a single vulnerability to achieve their objectives, defenders must identify and address all potential security weaknesses across their entire infrastructure. This asymmetry, combined with widespread shortages of cybersecurity professionals, has created a situation where traditional comprehensive approaches to vulnerability management are no longer sustainable.

The shift toward remote work following the COVID-19 pandemic has further exacerbated these challenges. Organizations now expose more services to the internet than ever before, with VPN gateways, remote desktop services, and cloud applications becoming primary attack vectors. High-profile ransomware attacks have repeatedly demonstrated that attackers frequently exploit known vulnerabilities in internet-facing services as their initial access vector.

1.2 Limitations of Current Approaches

The Common Vulnerability Scoring System (CVSS) has served as the de facto standard for vulnerability management since its introduction. CVSS provides a standardized method for rating the severity of security vulnerabilities based on their technical characteristics, including attack complexity, required privileges, and potential impact on confidentiality, integrity, and availability.

However, CVSS-based vulnerability management has come under increasing criticism for its fundamental inefficiency. Research has consistently shown that only a small fraction of high-CVSS vulnerabilities are ever exploited in practice. Studies indicate that fewer than 20% of vulnerabilities rated as "Critical" or "High" by CVSS are actually observed being exploited in real-world attacks [2], yet these high-severity ratings are assigned to approximately 57% of all published vulnerabilities [3].

This mismatch creates a significant operational burden, leading researchers to question whether it's "time to change the CVSS" [3] and describing the system as "ubiquitous and broken" [4]. Security teams following CVSS-based prioritization must address thousands of vulnerabilities that pose little real-world risk. Simultaneously, they may overlook vulnerabilities that attackers actively exploit but receive lower CVSS scores due to complex attack vectors or authentication requirements.

Government agencies have recognized these limitations, with CISA implementing binding operational directives that mandate specific remediation timelines [5, 6]. These directives represent a shift toward evidence-based vulnerability management, requiring federal agencies to prioritize vulnerabilities with confirmed exploitation evidence regardless of their CVSS scores.

Furthermore, CVSS was designed to measure technical severity rather than risk in the broader sense. The Forum of Incident Response and Security Teams (FIRST), which maintains CVSS, explicitly states that CVSS base scores should not be used alone for risk assessment and recommends incorporating additional contextual factors [7]. However, in practice, most organizations rely primarily on CVSS base scores due to the complexity and resource requirements of implementing more comprehensive risk assessment frameworks.

1.3 Emerging Threat Intelligence Approaches

Recognizing the limitations of traditional vulnerability management, several threat intelligence-based approaches have emerged to provide more effective prioritization. In 2021, the U.S. Cybersecurity and Infrastructure Security Agency (CISA) introduced the KEV catalog, which identifies vulnerabilities that have been observed being exploited in real-world attacks. KEV represents a paradigm shift from theoretical severity assessment to evidence-based prioritization, requiring federal agencies to remediate KEV-listed vulnerabilities within specified timeframes regardless of their CVSS scores.

Developed by researchers and maintained by FIRST, EPSS uses machine learning to predict the probability that a vulnerability will be exploited within the next 30 days. EPSS incorporates numerous data sources, including threat intelligence feeds, proof-of-concept availability, and social media discussions, to generate probabilistic exploitation scores.

While both KEV and EPSS represent significant advances over pure CVSS-based prioritization, each has important limitations. KEV provides high-confidence information about past exploitation but cannot predict future threats. EPSS offers forward-looking predictions but may miss vulnerabilities that are already being exploited. Additionally, EPSS is explicitly not recommended for use as a standalone risk assessment tool, particularly for vulnerabilities where exploitation has already been confirmed.

1.4 Research Motivation and Objectives

This research is motivated by the observation that existing vulnerability management approaches—whether based on technical severity (CVSS), confirmed exploitation (KEV), or exploitation prediction (EPSS)—each capture important but incomplete aspects of vulnerability risk. Organizations need a practical framework that combines the strengths of these approaches while mitigating their individual limitations.

Our primary research objective is to develop and validate an integrated vulnerability management methodology that improves efficiency over traditional CVSS-based approaches by focusing resources on vulnerabilities with confirmed or predicted exploitation, maintains comprehensive coverage to ensure that critical vulnerabilities are not overlooked, leverages existing open-source data to enable broad adoption without requiring proprietary threat intelligence, and provides actionable guidance through clear decision criteria rather than abstract scoring systems.

1.5 Proposed Approach and Contributions

We propose Vulnerability Management Chaining, a decision tree-based framework that systematically combines CVSS, EPSS, and KEV to achieve more effective vulnerability prioritization. Our approach employs a two-stage evaluation process: first assessing threat likelihood using KEV and EPSS data, then evaluating vulnerability characteristics using CVSS to determine appropriate response priorities.

The key insight underlying our approach is that vulnerabilities should be prioritized based on both threat likelihood and vulnerability characteristics. A vulnerability that is being actively exploited (high threat) but requires significant privileges and has limited impact (low vulnerability severity) may warrant different treatment than one with theoretical high impact but no evidence of exploitation.

Our primary contributions include an integrated framework design that develops a practical decision tree framework systematically combining three major vulnerability management approaches, addressing the limitations of using any single method in isolation. We provide empirical validation by demonstrating the effectiveness of our approach using real-world data including 28,377 CVEs published over a 13-month period, actual network intrusion detection data, and public security vendor reports of exploited vulnerabilities.

Our experimental results show 14-18 fold efficiency compared to CVSS-based approaches while maintaining comparable coverage levels (85.6-85.7% vs. CVSS coverage of 90-100%). We provide a complete methodology using only open-source data sources, enabling immediate adoption by organizations without requiring expensive commercial threat intelligence subscriptions.

1.6 Paper Organization

The remainder of this paper is structured as follows. Section 2 provides detailed background on existing vulnerability management approaches and establishes the theoretical foundation for our integrated methodology. Section 3 presents our Vulnerability Management Chaining framework, including the decision tree design and implementation considerations. Section 4 describes our experimental methodology and presents comprehensive evaluation results. Section 6 discusses the implications of our findings, limitations of the current approach, and directions for future work. Section 7 concludes with a summary of contributions and their broader significance for the cybersecurity community.

2 Background and Related Work

This section provides essential background on the three vulnerability management frameworks that form the foundation of our approach, reviews related work in vulnerability prioritization, and identifies the specific gaps that motivate our research.

2.1 Common Vulnerability Scoring System (CVSS)

CVSS, maintained by the Forum of Incident Response and Security Teams (FIRST), provides a standardized framework for rating the severity of security vulnerabilities [8]. The system generates numerical scores from 0.0 to 10.0 based on eight metrics that capture the fundamental characteristics of vulnerabilities.

CVSS Base Score calculations consider two primary factors: exploitability metrics including Attack Vector, Attack Complexity, Privileges Required, and User Interaction, along with impact metrics covering Confidentiality, Integrity, Availability impacts, plus Scope. For example, a remotely exploitable vulnerability requiring no authentication that completely compromises system confidentiality, integrity, and availability would receive the maximum score of 10.0.

The scoring system categorizes vulnerabilities into severity levels: Critical (9.0-10.0), High (7.0-8.9), Medium (4.0-6.9), and Low (0.1-3.9). Government agencies and industry standards frequently mandate remediation timeframes based on these categories, with critical vulnerabilities often requiring patching within 15 days and high-severity vulnerabilities within 30 days [5].

However, research has consistently demonstrated a significant disconnect between CVSS scores and real-world exploitation patterns. Analysis of vulnerability databases shows that approximately 57% of published vulnerabilities receive High or Critical CVSS ratings, yet studies of actual exploitation patterns indicate that fewer than 20% of these highly-rated vulnerabilities are ever observed being exploited [9, 10].

2.2 Exploit Prediction Scoring System (EPSS)

EPSS addresses CVSS limitations by using machine learning to predict the probability that a vulnerability will be exploited within 30 days [11]. The system is maintained by FIRST [12] and represents a significant advancement in predictive vulnerability assessment.[1] The current EPSS model (version 3.0) incorporates over 1,400 features including vulnerability characteristics, proof-of-concept availability, references in security advisories, social media mentions, and exploitation data from security vendors.

EPSS scores range from 0 to 1, representing the probability of exploitation. The system is updated daily, allowing scores to evolve as new information becomes available. Research shows that EPSS significantly outperforms CVSS for identifying vulnerabilities likely to be exploited, achieving 82% coverage of exploited vulnerabilities while requiring remediation of only 45.5% of the total vulnerability population, compared to CVSS which achieves similar coverage but requires addressing 96.1% of vulnerabilities [11].

However, EPSS has important limitations. The system is explicitly not recommended as a standalone risk assessment tool, as it focuses solely on threat likelihood without considering potential impact or organizational context. Additionally, EPSS is not designed for vulnerabilities where exploitation has already been confirmed, as the prediction model assumes unknown exploitation status. Recent enhancements to EPSS incorporate community-driven insights and expanded data sources [14], demonstrating continued evolution in predictive vulnerability assessment methodologies.

2.3 Known Exploited Vulnerabilities (KEV)

In response to high-profile attacks exploiting known vulnerabilities, CISA introduced the KEV catalog in 2021 [15]. KEV lists vulnerabilities that meet three criteria: assigned CVE identifier, confirmed evidence of active exploitation, and available vendor-provided remediation. As of 2023, the catalog contains over 900 vulnerabilities spanning multiple years.

KEV represents a significant policy shift toward evidence-based vulnerability management. Binding Operational Directive 22-01 requires federal agencies to remediate KEV-listed vulnerabilities within specified timeframes (typically 14 days) regardless of CVSS scores [6]. This approach acknowledges that confirmed exploitation provides stronger evidence of risk than theoretical severity assessments.

While KEV provides high-confidence threat intelligence, it has notable limitations. The catalog only includes vulnerabilities with confirmed exploitation evidence, potentially missing emerging threats that have not yet been widely observed. Additionally, KEV focuses primarily on vulnerabilities affecting federal networks, which may not fully represent threat landscapes for other sectors.

2.4 Related Work in Vulnerability Management

Several research efforts have attempted to improve vulnerability prioritization beyond traditional CVSS-based approaches. Commercial security vendors have developed proprietary risk-scoring systems that combine vulnerability severity with threat intelligence and environmental context [16, 17]. Alternative academic approaches have explored offensive security perspectives for vulnerability prioritization [18]. While these systems show promise, they typically require expensive subscriptions and use proprietary algorithms that limit transparency and reproducibility.

Various researchers have applied machine learning to vulnerability prioritization [19, 20]. These approaches often achieve good performance on specific datasets but face challenges with generalizability and the need for extensive labeled training data. Some research has applied formal decision analysis techniques to vulnerability prioritization [21, 22]. While mathematically rigorous, these approaches often require significant expertise to implement and may not scale to operational environments.

Developed by researchers at Carnegie Mellon's CERT Coordination Center, SSVC provides decision trees for different stakeholder roles including suppliers, deployers, and coordinators [10]. The framework has been refined through

[1] This study utilizes EPSS v3.0. During manuscript preparation, EPSS v4.0 was released with enhanced predictive capabilities, and alternative predictive metrics such as Likely Exploited Vulnerabilities (LEV) have emerged [13], further validating the evolution toward sophisticated threat prediction methodologies.

multiple iterations [23], offering more nuanced decision-making than single numerical scores but requiring manual assessment of multiple factors for each vulnerability.

Various research efforts have explored integrating threat intelligence into vulnerability management. Studies have examined dark web monitoring for threat prediction [24] and proactive identification of emerging threats [25]. However, these approaches often require specialized data sources and expertise not available to all organizations.

2.5 Research Gaps and Motivation

Our analysis of existing approaches reveals several important gaps that motivate our research. While individual systems like CVSS, EPSS, and KEV each provide valuable information, no practical framework exists for systematically combining their insights. Organizations must manually reconcile conflicting signals from different systems. Advanced approaches like SSVC provide sophisticated decision-making frameworks but require significant expertise and time investment that may not be practical for many organizations.

Much existing research focuses on algorithmic improvements using synthetic or limited datasets, with insufficient validation using real-world operational data across diverse organizational contexts. Many promising approaches rely on proprietary data sources or commercial tools, limiting adoption by organizations with constrained security budgets.

Our research addresses these gaps by developing a practical integration framework that combines the strengths of existing open-source vulnerability management systems while providing empirical validation using real-world exploitation data. The resulting approach aims to be both theoretically sound and operationally practical for organizations of varying sizes and security maturity levels.

3 Proposed Methodology

This section presents Vulnerability Management Chaining, our integrated framework that combines CVSS, EPSS, and KEV to achieve more efficient vulnerability prioritization while maintaining comprehensive coverage. We describe the theoretical foundation, decision tree design, and practical implementation considerations.

3.1 Framework Design Principles

Our approach aligns with established risk assessment frameworks [26] while addressing the practical limitations identified in vulnerability management research [27]. The design principles address specific shortcomings of existing approaches through a threat-first prioritization strategy that evaluates whether a vulnerability poses an active or likely threat before assessing technical characteristics. This approach focuses limited security resources on vulnerabilities with confirmed or predicted exploitation potential, rather than starting with technical severity assessment that may identify thousands of theoretically severe but practically irrelevant vulnerabilities.

The framework employs contextual refinement by applying additional criteria to assess attack feasibility and potential impact after identifying threat-relevant vulnerabilities. This enables more nuanced prioritization decisions that consider both the likelihood of exploitation and the practical constraints that may affect attack success or impact severity. Unlike complex alternatives that require proprietary data sources, our framework relies exclusively on freely available data sources including CVSS scores from NVD, EPSS scores from FIRST, and KEV listings from CISA, ensuring broad accessibility and reproducibility across organizations with varying security budgets.

The decision tree structure provides clear, actionable guidance without requiring complex mathematical calculations or subjective assessments that might introduce inconsistency across different analysts or organizational contexts. This operational simplicity enables high levels of automation while maintaining the sophistication necessary for effective threat intelligence integration.

3.2 Vulnerability Management Chaining Architecture

Figure 1 illustrates our complete decision tree framework, which processes vulnerabilities through two sequential evaluation stages that systematically combine threat likelihood assessment with vulnerability characteristic evaluation.

The first stage determines whether a vulnerability poses an active or imminent threat using two complementary data sources. Vulnerabilities listed in CISA's KEV catalog receive immediate threat classification, as they have confirmed evidence of real-world exploitation. KEV provides high-confidence threat intelligence but only covers vulnerabilities where exploitation has been definitively observed and reported. For vulnerabilities not in KEV, we apply EPSS scores to identify those with high exploitation probability. Following established research [11, 14], we use a threshold of 0.088,

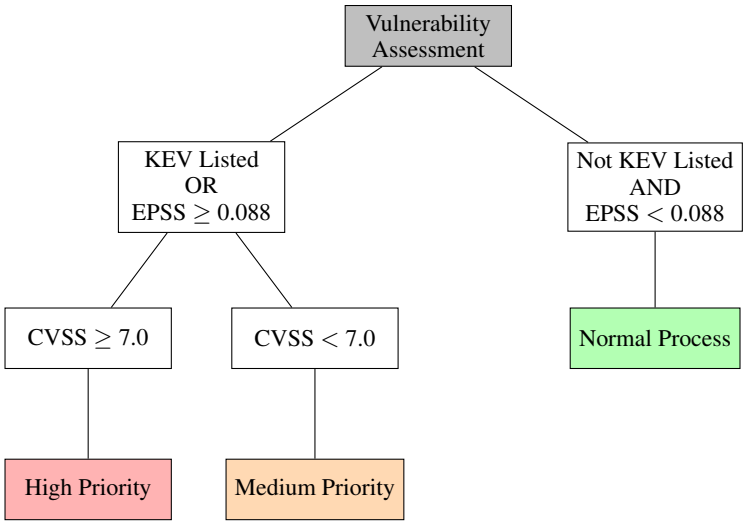

Figure 1: Vulnerability Management Chaining decision tree framework

which provides coverage comparable to CVSS 7.0+ while significantly reducing the number of vulnerabilities requiring priority attention.

The combination of KEV and EPSS addresses the limitations of either approach alone. KEV covers confirmed threats but cannot predict emerging exploitation, while EPSS provides forward-looking threat assessment but may miss currently exploited vulnerabilities. Together, they provide comprehensive threat coverage spanning both current and predicted exploitation scenarios.

Vulnerabilities identified as threat-relevant in the first stage undergo additional evaluation using CVSS base scores to assess attack feasibility and potential impact. High-impact vulnerabilities with CVSS scores of 7.0 or higher combine threat relevance with significant attack potential, warranting immediate attention. The CVSS 7.0 threshold aligns with widely adopted government and industry standards for "High" severity classification. Lower-impact vulnerabilities with CVSS scores below 7.0, while having confirmed or predicted exploitation potential, may warrant different treatment due to factors that reduce urgency such as complex attack prerequisites, limited impact scope, or significant attack complexity.

This two-stage approach enables organizations to maintain aggressive response timelines for the most critical vulnerabilities while applying more measured approaches to threats with mitigating factors. The threat-first architecture places threat metrics before vulnerability metrics because first filtering for threat-relevant vulnerabilities allows organizations to focus detailed analysis on a much smaller subset of the total vulnerability population, improving analytical efficiency. Confirmed or predicted exploitation provides stronger evidence for prioritization decisions than theoretical severity assessments alone, and as vulnerability disclosure rates continue to increase, threat-first filtering becomes increasingly important for maintaining manageable workloads.

3.3 Implementation Considerations

Practical implementation requires systematic data collection and integration through automated retrieval of vulnerability data from NIST's National Vulnerability Database, including CVSS base scores and vulnerability descriptions. Daily download of EPSS scores from FIRST's public API maintains historical data to track score evolution over time, while regular synchronization with CISA's KEV catalog identifies newly added vulnerabilities requiring immediate attention. Automated application of the decision tree logic generates prioritized vulnerability lists with specific remediation timelines.

While our framework provides standard thresholds based on research evidence, organizations may need to adapt these values based on specific contexts. Organizations with limited remediation capacity might raise the EPSS threshold to focus on higher-probability threats, while those with greater resources might lower it for more comprehensive coverage. Different risk tolerance levels might warrant adjusting the CVSS 7.0 threshold, though this should be done carefully to maintain consistency with industry standards. The framework's priority categories can be mapped to organization-specific remediation timelines based on operational capabilities and risk appetite.

The framework supports high levels of automation through continuous monitoring systems that can immediately flag new high-priority vulnerabilities as they are discovered or as threat intelligence evolves. The decision tree outputs can be integrated with existing vulnerability management platforms, ticketing systems, and patch management tools to streamline operational workflows. Standardized reporting enables tracking of key performance indicators such as coverage rates, false positive rates, and remediation timelines across different priority categories.

Our framework is designed to address specific shortcomings of current approaches by focusing initial attention on threat-relevant vulnerabilities to achieve better efficiency ratios, where a higher percentage of remediated vulnerabilities are actually exploited. The combination of KEV and EPSS should maintain coverage levels comparable to traditional CVSS-based approaches while requiring attention to fewer total vulnerabilities. More targeted prioritization should reduce the overwhelming number of critical vulnerabilities that security teams face with CVSS-only approaches.

Several limitations must be acknowledged. The framework's effectiveness depends on the quality and timeliness of external data sources, particularly EPSS model accuracy and KEV catalog completeness. Some vulnerabilities may receive high priority classification based on threat indicators but prove irrelevant to specific organizational contexts. As threat landscapes and data sources evolve, the framework thresholds and logic may require periodic review and adjustment.

4 Experimental Design

This section describes our comprehensive experimental methodology for evaluating the effectiveness of Vulnerability Management Chaining. We present our data collection procedures, evaluation metrics, and experimental setup designed to provide rigorous validation of our approach using real-world vulnerability and exploitation data.

4.1 Research Questions and Methodology

Our experimental evaluation addresses four primary research questions that guide our methodology design. We investigate whether Vulnerability Management Chaining achieves better efficiency than existing approaches while maintaining comparable coverage, examine whether KEV and EPSS integration effects provide improved coverage compared to individual methods, assess whether CVSS integration enables appropriate deprioritization of lower-impact threats, and compare our approach with previous research findings on vulnerability management effectiveness.

4.2 Data Collection and Preparation

We collected comprehensive vulnerability data covering the period from April 1, 2022, to April 30, 2023, representing 13 months of vulnerability disclosures. This timeframe was selected to ensure sufficient data volume while incorporating recent threat patterns and EPSS model improvements. Using the NIST National Vulnerability Database (NVD) API, we retrieved 28,377 unique CVE entries published during this period, including CVE identifiers and publication dates, CVSS 3.1 base scores and vector strings, vulnerability descriptions and affected products, and reference links and external identifiers.

The 13-month timeframe accounts for research showing that approximately 75% of vulnerability exploitation occurs within 28 days of disclosure [28], providing sufficient observation period for exploitation patterns to emerge. EPSS scores were collected using FIRST's public data feeds, which provide daily updated scores for all CVE-assigned vulnerabilities. Since EPSS scores evolve over time as new information becomes available, we employed a maximum-value approach: for each vulnerability, we recorded the highest EPSS score observed during our study period.

This methodology aligns with operational practice, where security teams would likely respond when a vulnerability's EPSS score first exceeds the threshold, regardless of subsequent fluctuations. The approach also accounts for the dynamic nature of threat intelligence that drives EPSS score updates. We obtained CISA's Known Exploited Vulnerabilities catalog data as of April 30, 2023, containing 922 vulnerabilities with confirmed exploitation evidence. The KEV catalog includes CVE identifiers and vulnerability descriptions, required action descriptions, due dates for federal agency compliance, and dates of KEV catalog addition.

Establishing ground truth for actual vulnerability exploitation represents one of the most challenging aspects of vulnerability management research. We developed a multi-source approach to create comprehensive exploitation datasets. Network intrusion detection data was collected from a production network intrusion detection system (NIDS) deployed at a sample organization. The NIDS monitored public-facing web servers and recorded attack attempts targeting specific vulnerabilities over our study period. This dataset includes 28 unique vulnerabilities with confirmed exploitation attempts, though we acknowledge that signature-based detection may include some false positives where a single signature corresponds to multiple CVEs.

We systematically reviewed public security reports from major cybersecurity vendors to identify vulnerabilities reported as exploited during our study period. This multi-source approach yielded 90 additional vulnerabilities with public documentation of real-world exploitation, bringing our total exploitation dataset to 118 unique vulnerabilities across both data sources. Sources included threat intelligence reports from established security vendors, incident response case studies, annual threat landscape assessments, and vulnerability research publications.

4.3 Evaluation Metrics and Experimental Setup

Following established research methodology [11], we employ two primary metrics for comparative evaluation. Efficiency measures the proportion of prioritized vulnerabilities that were actually exploited, calculated as the number of exploited vulnerabilities in the priority set divided by the total vulnerabilities in the priority set, expressed as a percentage. Higher efficiency indicates that a larger percentage of the vulnerabilities flagged for priority attention were actually relevant to real-world threats, directly addressing the resource allocation challenge facing security teams.

Coverage measures the proportion of exploited vulnerabilities captured by each prioritization method, calculated as the number of exploited vulnerabilities in the priority set divided by the total exploited vulnerabilities, expressed as a percentage. Higher coverage indicates that the prioritization method successfully identified a larger percentage of vulnerabilities that attackers actually exploited, addressing the completeness requirement for security-critical applications.

To evaluate the specific benefits of combining KEV and EPSS, we measure incremental coverage as additional exploited vulnerabilities identified through the combination that neither KEV nor EPSS would capture individually, and complementary effectiveness as the degree to which KEV and EPSS identify different subsets of exploited vulnerabilities, validating our hypothesis that they provide complementary threat intelligence.

We compare our Vulnerability Management Chaining approach against three established baselines. The CVSS baseline uses traditional high-severity vulnerability prioritization with CVSS scores of 7.0 or higher, representing current industry standard practice. KEV-Only prioritization is based solely on CISA KEV catalog membership, representing pure evidence-based threat intelligence. EPSS-Only prioritization uses EPSS scores above the research-established threshold of 0.088, representing predictive threat intelligence.

For each baseline and our proposed method, we evaluate performance across both exploitation datasets. Condition A uses the 28 vulnerabilities with exploitation attempts detected by network intrusion detection systems, while Condition B uses the 90 vulnerabilities documented in public security vendor reports. This dual-condition approach helps validate our findings across different types of exploitation evidence and threat contexts.

4.4 Validation and Limitations

Given the observational nature of vulnerability exploitation data, we focus on descriptive analysis and practical effect sizes rather than inferential statistics. Our analysis includes comparative performance through direct comparison of efficiency and coverage metrics across all methods and conditions, threshold sensitivity analysis of how performance varies with different EPSS and CVSS threshold values to validate our chosen parameters, and temporal pattern examination of whether performance varies across different time periods within our study window.

We compare our results with findings from previous research [11] to assess consistency with established benchmarks and identify any dataset-specific effects that might limit generalizability. Several robustness checks validate our methodology including manual verification of a sample of exploitation claims to assess the accuracy of our ground truth datasets, testing alternative EPSS and CVSS thresholds to ensure our results are not artifacts of specific parameter choices, and analysis of whether our findings hold across different subperiods within our study timeframe.

We acknowledge several limitations in our experimental design. Our exploitation datasets likely represent only a fraction of actual vulnerability exploitation, as many attacks go undetected or unreported. Network-based detection may be biased toward certain types of attacks and may miss sophisticated exploitation techniques. Our 13-month study period, while substantial, may not capture longer-term exploitation patterns or seasonal variations in attack activity. Results from specific organizational contexts may not generalize to all deployment scenarios.

All exploitation data used in our study comes from either public sources or anonymized network monitoring data with appropriate organizational consent. No personally identifiable information or proprietary attack details are included in our analysis. The following section presents detailed results from applying this experimental methodology to evaluate Vulnerability Management Chaining effectiveness.

5 Results and Analysis

This section presents the comprehensive evaluation results of our Vulnerability Management Chaining framework. We analyze performance across multiple datasets, compare with existing approaches, and examine the specific contributions of each component in our integrated methodology.

5.1 Overall Performance Comparison

Table 1 summarizes the performance of all evaluated methods across our two primary datasets. Our proposed Vulnerability Management Chaining approach demonstrates consistent improvements in efficiency while maintaining high coverage levels comparable to traditional CVSS-based prioritization.

Table 1: Performance comparison across vulnerability management approaches

Method	**NIDS Dataset (n=28)**		**Vendor Reports (n=90)**	
	Efficiency	**Coverage**	**Efficiency**	**Coverage**
CVSS ≥ 7.0	0.2%	100.0%	0.5%	90.0%
KEV Only	14.3%	53.6%	74.3%	86.7%
EPSS ≥ 0.088	2.7%	85.7%	4.9%	48.9%
Proposed Method (KEV$\vee$EPSS)$\wedge$CVSS	**2.8%**	**85.7%**	**9.1%**	**85.6%**

Our proposed method achieves substantial efficiency improvements over traditional CVSS-based approaches. In the NIDS dataset, Vulnerability Management Chaining achieves 2.8% efficiency compared to 0.2% for CVSS-only approaches, representing a 14-fold improvement. While KEV-only approaches achieve higher efficiency (14.3%), they sacrifice significant coverage (53.6% vs. 85.7%). The efficiency improvement is even more pronounced in the vendor report dataset, with our method achieving 9.1% efficiency compared to 0.5% for CVSS-only approaches—an 18-fold improvement. This dataset shows particularly strong performance for KEV-only approaches (74.3% efficiency), reflecting the alignment between vendor-reported exploitation and KEV catalog contents.

Across both datasets, our method provides efficiency levels significantly better than CVSS-only or EPSS-only approaches while maintaining coverage levels much higher than KEV-only approaches. This balanced performance validates our design goal of combining the strengths of individual methods. Coverage results demonstrate that our integrated approach successfully maintains comprehensive vulnerability identification. Our method achieves 85.6-85.7% coverage across both datasets, approaching the 90-100% coverage of CVSS-based approaches while requiring attention to far fewer vulnerabilities.

The consistency of coverage results across different exploitation datasets (NIDS-based vs. vendor-reported) suggests that our approach is robust to different types of threat intelligence and organizational contexts. The 85+ percent coverage levels indicate that organizations adopting our approach would capture the vast majority of actually exploited vulnerabilities while focusing resources much more effectively than traditional approaches.

5.2 Integration Effects Analysis

To understand the specific value of combining KEV and EPSS data sources, we conducted detailed analysis of their complementary effects. Table 2 provides a comprehensive breakdown of how KEV and EPSS identify different subsets of exploited vulnerabilities, validating our hypothesis about their complementary nature. The analysis reveals that 57 additional exploited vulnerabilities (48.3% of our total dataset) are captured only through the integration of both methods—a finding that demonstrates the critical importance of multi-source threat intelligence.

Table 2: Integration Benefits Analysis: KEV + EPSS Complementary Effects

Vulnerability Category	**Count**	**Percentage**	**Data Source**	**Integration Value**
KEV Only (EPSS < 0.088)	45	38.1%	Historical Evidence	High Confidence
EPSS Only (Not in KEV)	16	13.6%	Predictive Model	Emerging Threats
Both KEV and EPSS	52	44.1%	Dual Confirmation	Highest Priority
Integration Benefit	**57**	**48.3%**	**Unique Coverage**	**Critical Gap Filled**

The table categorizes exploited vulnerabilities into three distinct groups: those identified exclusively by KEV (38.1%), those identified exclusively by EPSS (13.6%), and those identified by both methods (44.1%). This distribution pattern

confirms that neither individual approach provides complete threat coverage, making systematic integration not merely beneficial but essential for comprehensive vulnerability management.

Historical evidence vulnerabilities include 45 vulnerabilities with confirmed exploitation but low EPSS scores, often involving complex attack chains or specialized targets that provide high confidence but may not represent current threat trends. Predictive intelligence covers 16 vulnerabilities with high exploitation probability but no KEV listing, representing emerging threats not yet widely exploited that serve as an early warning system for proactive defense. Dual confirmation encompasses 52 vulnerabilities identified by both systems, representing the highest priority cases with proven exploitation plus continued threat and strong consensus across historical and predictive indicators.

The integration value of 57 additional vulnerabilities captured only through combination represents a critical gap that would exist in any single-method approach, highlighting the necessity of our integrated framework. Table 3 shows how KEV and EPSS identify different subsets of exploited vulnerabilities across our datasets, further validating their complementary nature.

Table 3: KEV and EPSS integration effects

Category	NIDS Dataset	Vendor Reports	Total Unique
KEV Only (EPSS < 0.088)	4	41	45
EPSS Only (Not in KEV)	9	7	16
Both KEV and EPSS	15	37	52
Total Integration Benefit	**9**	**48**	**57**
*Integration benefit = unique additional vulnerabilities (overlap removed)			

The combination of KEV and EPSS identifies 57 additional exploited vulnerabilities that neither method would capture individually. This represents 32.1% of exploited vulnerabilities in the NIDS dataset and 53.3% in the vendor report dataset. KEV identifies vulnerabilities with confirmed exploitation evidence but low EPSS scores (often due to complex attack requirements or limited automation), while EPSS identifies emerging threats not yet documented in KEV. This pattern validates our theoretical framework for combining historical and predictive threat intelligence.

Figure 2 provides a comprehensive visualization of the integration effects, demonstrating the complementary relationship between KEV and EPSS. The combination identifies 57 additional exploited vulnerabilities that neither method captures individually, representing 48.3% of the total dataset and validating the necessity of multi-source threat intelligence integration.

5.3 CVSS Integration Analysis

Our framework applies CVSS filtering after threat identification to enable appropriate deprioritization of lower-impact vulnerabilities. Analysis of this integration reveals both benefits and limitations. We identified 8 vulnerabilities that received threat classification (KEV or EPSS ≥ 0.088) but were appropriately deprioritized due to CVSS < 7.0. Analysis of their CVSS vector strings reveals several categories suitable for deprioritization.

Limited impact vulnerabilities include CVE-2023-26083 (CVSS 3.3) with only local access requirements and minimal confidentiality impact. User interaction required cases encompass CVE-2022-44698 and CVE-2022-41091 (CVSS 5.4) requiring user interaction for exploitation. Authentication prerequisites cover CVE-2022-22674 (CVSS 5.5) requiring local authentication.

However, our analysis also identified 2 vulnerabilities that may have been inappropriately deprioritized: CVE-2022-2856 (CVSS 6.5), a network-accessible vulnerability with no authentication requirements, and CVE-2022-26925 (CVSS 5.9), a network-accessible vulnerability with potential for high integrity impact. These cases highlight the challenge of using CVSS thresholds for deprioritization decisions and suggest areas for future framework refinement.

5.4 Comparison with Previous Research

Our results are consistent with previous findings regarding CVSS inefficiency while demonstrating the value of integration approaches not previously evaluated in the literature. The efficiency improvements we observe align with broader criticisms of CVSS-only approaches. The substantial efficiency improvements have practical implications for organizations operating under regulatory requirements. Standards such as PCI DSS mandate CVSS-based remediation timelines, making our demonstrated efficiency gains particularly valuable for compliance-driven environments.

Table 4 compares our results with the seminal EPSS research by Jacobs et al., providing context for our findings within the broader vulnerability management literature.

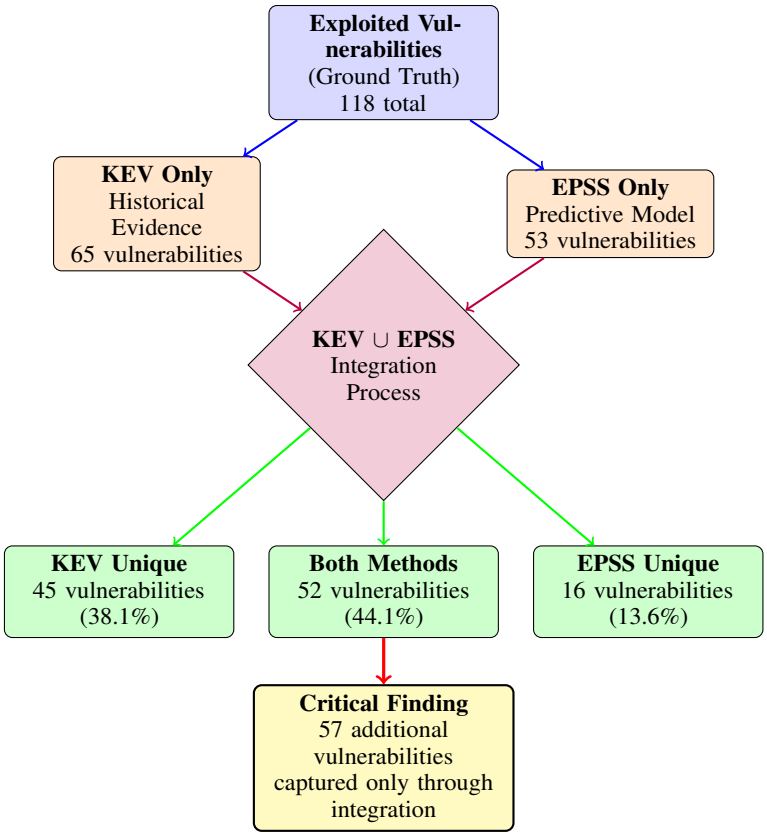

Figure 2: Integration effects flowchart demonstrating the complementary relationship between KEV and EPSS. The combination identifies 57 additional exploited vulnerabilities that neither method captures individually, representing 48.3% of the total dataset and validating the necessity of multi-source threat intelligence integration. Individual coverage rates: KEV 55.9%, EPSS 44.9%, Combined 95.8%.

Table 4: Comparison with previous research findings

Method	**Jacobs et al. (2021)**			**Our Study**	
	Efficiency	**Coverage**	**Dataset**	**Efficiency**	**Coverage**
CVSS $\geq$ 7.0	3.9%	82.1%	6.4M exploits	0.2-0.5%	90-100%
KEV Only	53.2%	5.9%	2016-2022	14.3-74.3%	53.6-86.7%
EPSS $\geq$ 0.088	45.5%	82.0%	Commercial data	2.7-4.9%	48.9-85.7%

Our results confirm the established finding that CVSS-based prioritization is highly inefficient, with our efficiency measurements (0.2-0.5%) being even lower than previous research (3.9%). This difference likely reflects our more limited dataset size and specific organizational context. Our KEV efficiency results (14.3-74.3%) align with the high precision characteristics found in previous research, though our coverage results are notably higher (53.6-86.7% vs. 5.9%). This difference suggests that our vendor report dataset may be more representative of KEV-type vulnerabilities than the broader exploitation dataset used in previous research.

Our EPSS results show more variation than previous research, with efficiency ranging from 2.7-4.9% compared to the established 45.5%. This variation highlights the importance of dataset characteristics and suggests that EPSS performance may be sensitive to the specific types of exploitation being measured. Most importantly, our study provides the first empirical evidence for the value of systematically integrating multiple vulnerability management approaches, showing that combined methods can achieve efficiency improvements while maintaining high coverage levels.

5.5 Performance Stability and Practical Implementation

The consistency of our main findings across two different exploitation datasets provides evidence for the robustness of our approach. Our method maintains 85+ percent coverage across both NIDS and vendor report datasets, suggesting

that performance is not dependent on specific data collection methodologies. While absolute efficiency values vary between datasets, the relative performance rankings remain consistent, with our integrated approach outperforming individual methods while maintaining balanced efficiency-coverage trade-offs.

We conducted sensitivity analysis around our chosen EPSS threshold of 0.088 to validate the robustness of our parameter selection. Results show that performance remains stable across EPSS thresholds from 0.05 to 0.15, with the 0.088 value representing a reasonable balance point established by previous research. Similarly, CVSS thresholds from 6.0 to 8.0 produce consistent relative performance patterns, supporting our choice of the industry-standard 7.0 threshold.

Our approach significantly reduces the operational burden compared to traditional CVSS-based prioritization. Organizations using our method would need to provide immediate attention to 844-858 vulnerabilities (depending on dataset) compared to 16,182 vulnerabilities required by CVSS ≥ 7.0 approaches—a approximately 95% reduction in urgent prioritization workload. The 85+ percent coverage maintained while achieving this workload reduction suggests that security teams can reallocate resources from low-value vulnerability remediation to other critical security activities.

While our method achieves better efficiency than alternatives, the 9.1% efficiency rate in our best-case scenario indicates that approximately 90% of prioritized vulnerabilities may not be exploited during the observation period. However, this must be interpreted considering our limited 13-month observation window that may not capture all exploitation activity, particularly for vulnerabilities that become attractive targets over longer time horizons. Rapid remediation of high-priority vulnerabilities may prevent exploitation that would otherwise occur, making some "false positives" actually represent successful prevention. Even with remaining false positives, our approach represents substantial improvement over CVSS-based methods that achieve only 0.2-0.5% efficiency.

This comprehensive results analysis demonstrates that Vulnerability Management Chaining provides substantial practical improvements over existing approaches while maintaining the coverage levels necessary for effective cybersecurity risk management.

6 Discussion

This section interprets our experimental findings, discusses their implications for vulnerability management practice, addresses limitations of our approach, and identifies directions for future research. Our results provide strong evidence for the practical value of integrating multiple vulnerability management frameworks while revealing important areas for continued development.

6.1 Key Findings and Implications

Our central hypothesis that combining CVSS, EPSS, and KEV would achieve better efficiency while maintaining coverage receives strong empirical support. The substantial efficiency improvements over traditional CVSS-based approaches, coupled with maintenance of 85+ percent coverage levels, demonstrate that systematic integration of multiple data sources can overcome the limitations inherent in any single approach.

These results support the broader principle that cybersecurity decision-making benefits from multi-source intelligence integration rather than reliance on individual metrics. The complementary nature of historical exploitation evidence (KEV) and predictive modeling (EPSS) validates threat intelligence frameworks that emphasize diverse data source integration. For organizations currently struggling with CVSS-based vulnerability management, our approach provides immediate operational advantages. The approximately 95% reduction in urgent prioritization workload from 16,182 to approximately 850 vulnerabilities while maintaining comprehensive coverage represents a transformative operational improvement.

The substantial performance differences between our NIDS and vendor report datasets reveal important insights about vulnerability management in different organizational contexts. KEV's exceptional performance on vendor report data (74.3% efficiency, 86.7% coverage) compared to NIDS data (14.3% efficiency, 53.6% coverage) suggests strong alignment between CISA's threat intelligence sources and commercial security vendor reporting. This finding validates KEV as particularly valuable for organizations that rely heavily on vendor threat intelligence.

EPSS shows more consistent but generally lower efficiency across both datasets compared to previous research. This variation highlights the importance of understanding that machine learning-driven prediction models may perform differently across various organizational contexts and threat landscapes. Importantly, our integrated approach maintains stable performance characteristics across both datasets, suggesting that the framework is robust to different types of exploitation evidence and organizational monitoring capabilities.

Figure 3 provides a comprehensive visualization of the operational improvements achieved by our Vulnerability Management Chaining framework. The dramatic efficiency gains and workload reductions demonstrated across multiple datasets highlight the transformative potential of systematic threat intelligence integration for organizational vulnerability management practices.

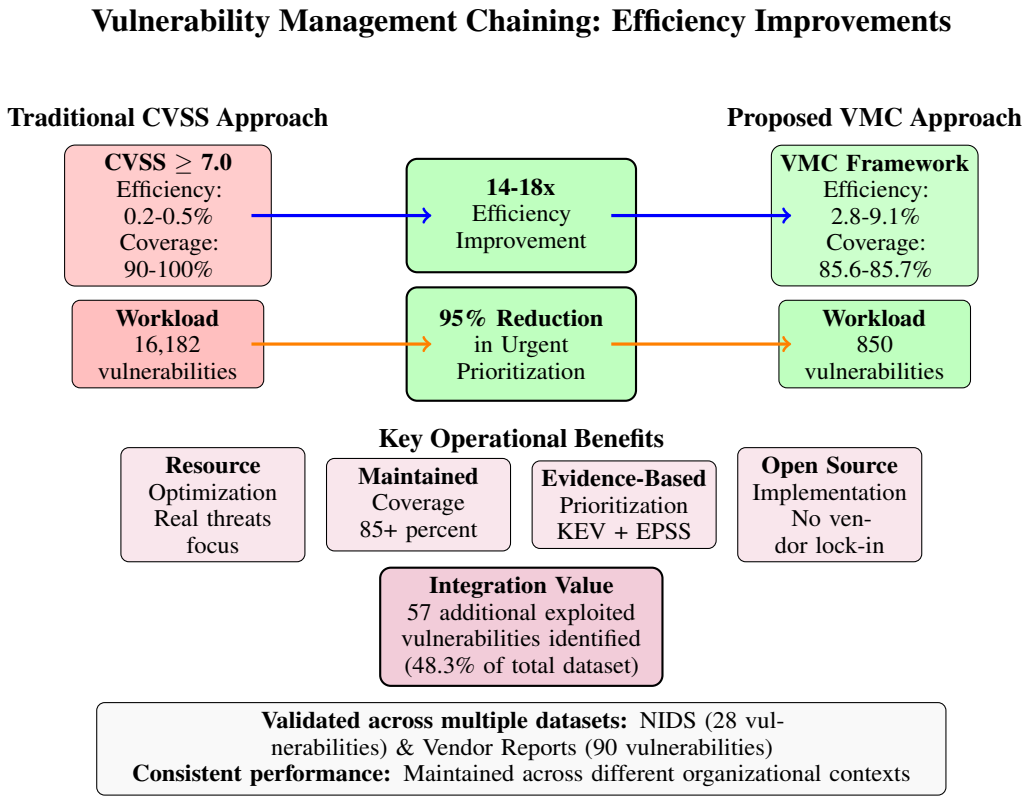

Figure 3: Operational efficiency improvements achieved by Vulnerability Management Chaining compared to traditional CVSS-based prioritization. The framework delivers approximately 14-18 fold efficiency improvements while reducing urgent prioritization workload by approximately 95% and maintaining comprehensive threat coverage through systematic integration of KEV and EPSS threat intelligence.

6.2 Integration Effects and CVSS Considerations

Our analysis of CVSS integration effects reveals both the promise and limitations of using technical severity scores for threat deprioritization. The identification of vulnerabilities appropriately deprioritized based on CVSS characteristics including local access requirements, user interaction dependencies, and limited impact scope demonstrates that technical severity metrics retain value when applied appropriately. These cases validate the principle that not all exploited vulnerabilities require identical response urgency.

From a risk management standpoint, the ability to systematically identify lower-urgency threats enables more sophisticated resource allocation decisions. Organizations can maintain aggressive timelines for network-accessible, high-impact vulnerabilities while applying more measured approaches to threats with significant mitigating factors. However, the identification of potentially inappropriate deprioritization cases, particularly network-accessible vulnerabilities with CVSS scores just below our 7.0 threshold, highlights a fundamental limitation of threshold-based decision systems.

Future versions of the framework should consider individual CVSS vector components rather than relying solely on aggregate base scores. This approach could provide more nuanced decision-making while maintaining operational simplicity. Organizations with specific threat profiles or risk tolerances might benefit from customized threshold values or additional decision criteria based on their operational environment. Systematic tracking of deprioritization decisions and their outcomes could enable continuous improvement of the framework's decision logic based on organizational experience.

6.3 Comparison with Existing Approaches

Our approach addresses several key limitations of existing vulnerability management methods by achieving approximately 14-18 fold efficiency improvements that directly address the resource allocation challenges that make CVSS-based prioritization unsustainable for many organizations. The integration of EPSS with KEV addresses KEV's primary weakness of limited coverage of emerging threats while preserving KEV's high-confidence threat intelligence for confirmed exploitation cases. By combining EPSS with impact assessment through CVSS, our approach addresses the explicit recommendation against using EPSS as a standalone risk assessment tool.

While our method is simpler than sophisticated approaches like SSVC (Stakeholder-Specific Vulnerability Categorization), this simplicity represents a strategic design choice. Complex decision frameworks often face adoption challenges due to expertise requirements and implementation complexity. Our decision tree approach provides sophisticated intelligence integration while remaining accessible to organizations with limited cybersecurity resources. The systematic nature of our approach enables high levels of automation, reducing the human expertise requirements that can limit the scalability of more complex decision frameworks. Using established thresholds and widely available data sources promotes consistency across organizations and enables benchmarking and continuous improvement efforts.

6.4 Limitations and Future Research Directions

Our approach's effectiveness depends fundamentally on the quality and completeness of its underlying data sources. EPSS performance may vary as the underlying machine learning model evolves, threat landscapes change, or training data characteristics shift. Organizations adopting our approach should monitor EPSS performance over time and be prepared to adjust thresholds if necessary. KEV focuses primarily on vulnerabilities affecting federal networks and may not fully represent threat landscapes for specific industry sectors or geographic regions. Organizations in specialized environments should supplement KEV with sector-specific threat intelligence when available.

An important consideration for this research is the rapidly evolving vulnerability management landscape. New predictive metrics such as Likely Exploited Vulnerabilities (LEV) and enhanced versions of existing systems (e.g., EPSS v4.0) have emerged during our study period, while security vendors have independently developed integration approaches conceptually similar to our methodology. These developments underscore both the timeliness of our research question and the industry-wide recognition that single-metric vulnerability management approaches are fundamentally insufficient for modern threat environments.

Rather than limiting the relevance of our findings, these advances validate our core premise: systematic integration of complementary intelligence sources provides superior vulnerability prioritization compared to any individual metric alone. Our framework establishes methodological principles that transcend specific tool implementations. The decision tree structure is designed to accommodate new intelligence sources, evolving prediction models, and adjusted threshold values without requiring fundamental architectural changes. This adaptability ensures that the integration methodology remains applicable as the threat intelligence ecosystem continues to evolve, providing organizations with a stable framework for incorporating emerging capabilities while maintaining operational consistency and proven effectiveness.

Several factors may limit the generalizability of our specific findings. Our exploitation datasets, while comprehensive, represent specific organizational contexts and time periods. Performance may vary in different threat environments or against different attack patterns. Vulnerability management performance may change over time as threat actor behaviors evolve, new exploitation techniques emerge, or defensive capabilities improve. Different organizations face different threat profiles based on their industry, size, geographic location, and security posture.

Our current approach performs static analysis of vulnerability characteristics without considering dynamic factors such as asset criticality, network exposure, or organizational-specific threat intelligence. The decision tree structure requires binary choices at each decision point, potentially losing nuanced information that could inform more sophisticated prioritization decisions. While our framework allows threshold adjustment, it provides limited mechanisms for incorporating organization-specific knowledge or contextual factors.

Several near-term improvements could enhance the framework's effectiveness. Machine learning approaches could automatically adjust EPSS and CVSS thresholds based on organizational feedback and historical performance data. Incorporating asset criticality and network exposure information could provide more sophisticated risk assessment while maintaining operational simplicity. Formal decision analysis techniques could replace simple threshold-based decisions with more nuanced multi-factor assessment while preserving automation potential.

Extended studies tracking framework performance over multiple years could identify temporal patterns and inform adaptive management strategies. Research examining how different types of organizations adapt and customize the framework could inform best practices and implementation guidance. Studies examining how framework performance changes as threat landscapes evolve could inform proactive adaptation strategies. Improved methods for identifying

and validating vulnerability exploitation could strengthen evaluation methodologies for all vulnerability management research.

Organizations considering adoption of our framework should consider gradual deployment with specific vulnerability classes or systems before full-scale deployment to build confidence and identify necessary customizations. Systematic tracking of framework performance against organizational metrics including time to remediation, exploitation incidents, and resource utilization can inform continuous improvement efforts. Staff training on threat intelligence interpretation and framework customization supports successful implementation.

The demonstrated effectiveness of KEV integration supports continued investment in government-led threat intelligence sharing initiatives. Professional organizations and standards bodies should consider incorporating threat intelligence integration into vulnerability management guidance and frameworks. Regulators requiring vulnerability management programs should consider allowing evidence-based prioritization approaches as alternatives to pure CVSS-based requirements.

7 Conclusion

This paper presents Vulnerability Management Chaining, a novel framework that integrates CVSS, EPSS, and KEV to achieve more efficient vulnerability prioritization while maintaining comprehensive coverage. Our research demonstrates that systematic integration of multiple intelligence sources can overcome the fundamental limitations that plague individual vulnerability management approaches.

7.1 Key Contributions and Practical Impact

This paper presents Vulnerability Management Chaining, a novel methodological framework that integrates CVSS, EPSS, and KEV to achieve more efficient vulnerability prioritization while maintaining comprehensive coverage. Our research demonstrates that systematic integration of multiple data sources can overcome the fundamental limitations that plague individual vulnerability management approaches.

We developed the first systematic methodological framework for combining CVSS technical severity assessment, EPSS predictive threat intelligence, and KEV confirmed exploitation evidence, establishing foundational principles for multi-source vulnerability intelligence integration that transcend specific tool implementations. Using 28,377 vulnerabilities and exploitation evidence from multiple sources, we demonstrated significant efficiency improvements over traditional CVSS-based prioritization while maintaining 85+ percent coverage of actually exploited vulnerabilities.

Our analysis reveals that KEV and EPSS identify complementary sets of exploited vulnerabilities, with their combination capturing 57 additional vulnerabilities that neither method would identify individually. By relying exclusively on freely available data sources, our framework enables broad adoption without requiring expensive commercial threat intelligence subscriptions.

Organizations adopting our framework can reduce their urgent vulnerability remediation workload by approximately 95% while capturing 85+ percent of vulnerabilities that attackers actually exploit. This substantial workload reduction enables security teams to allocate resources more effectively and focus on truly critical threats. Unlike complex alternatives requiring specialized expertise, our decision tree framework can be implemented using standard vulnerability management tools and provides immediate operational benefits for organizations regardless of their size or security maturity level.

7.2 Future Directions and Final Remarks

As the cybersecurity landscape continues to evolve, our methodological framework provides a foundation for incorporating emerging threat intelligence sources and enhanced prediction models. The demonstrated effectiveness of systematic multi-source integration suggests that future developments should focus on adaptive frameworks that can seamlessly incorporate new intelligence feeds while maintaining operational simplicity and proven effectiveness.

The exponential growth in vulnerability disclosures makes efficient vulnerability management increasingly critical for organizational security. Our Vulnerability Management Chaining framework provides a practical solution that leverages existing open source data to achieve dramatic efficiency improvements while maintaining comprehensive security coverage. More broadly, our research demonstrates that systematic integration of multiple intelligence sources can overcome the limitations inherent in individual approaches, providing a methodology applicable beyond vulnerability management to other cybersecurity domains.

As the cybersecurity field continues to evolve toward more sophisticated, data-driven approaches, frameworks like ours will become increasingly important for managing complex security decisions at scale. To support reproducibility and enable further research, we commit to making our experimental data, analysis code, and framework implementation available through open source repositories. This comprehensive availability aligns with our goal of making advanced vulnerability management accessible to organizations regardless of their resources or expertise levels.

References

[1] NIST, "National vulnerability database," https://nvd.nist.gov/, 2023, accessed July 2023.

[2] J. Jacobs *et al.*, "Improving vulnerability remediation through better exploit prediction," *Journal of Cybersecurity*, vol. 6, no. 1, p. tyaa015, 2020.

[3] J. M. Spring, E. Hatleback, A. Manion, D. Shick, and A. D. Householder, "Time to change the cvss?" *IEEE Security & Privacy*, vol. 19, no. 2, pp. 74–78, 2021.

[4] H. Howland, "Cvss: Ubiquitous and broken," *Digital Threats: Research and Practice*, vol. 4, no. 1, pp. 1–12, 2023.

[5] U. Cybersecurity and I. S. Agency, "Binding operational directive 19-02," https://www.cisa.gov/news-events/directives/binding-operational-directive-19-02, 2019.

[6] ——, "Binding operational directive 22-01: Reducing the significant risk of known exploited vulnerabilities," https://www.cisa.gov/news-events/directives/binding-operational-directive-22-01, 2021.

[7] F. of Incident Response and S. Teams, "Common vulnerability scoring system version 3.1: User guide," https://www.first.org/cvss/v3.1/user-guide, 2019.

[8] ——, "Common vulnerability scoring system version 3.1: Specification document," https://www.first.org/cvss/v3.1/specification-document, 2019.

[9] J. Jacobs, S. Romanosky, I. Adjerid, and W. Baker, "Exploit prediction scoring system (epss)," in *Digital Threats: Research and Practice*, 2019.

[10] J. M. Spring *et al.*, "Prioritizing vulnerability response: A stakeholder-specific vulnerability categorization," Software Engineering Institute, Carnegie Mellon University, Tech. Rep., 2019.

[11] 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, pp. 1–17, 2021.

[12] F. of Incident Response and S. Teams, "Exploit prediction scoring system (epss)," https://www.first.org/epss/, 2023, accessed June 2023.

[13] P. Mell and J. Spring, "Likely exploited vulnerabilities: A proposed metric for vulnerability exploitation probability," National Institute of Standards and Technology, Tech. Rep. NIST CSWP 41, May 2025, nIST Cybersecurity White Paper. [Online]. Available: https://csrc.nist.gov/pubs/cswp/41/likely-exploited-vulnerabilities-a-proposed-metric/final

[14] J. Jacobs, S. Romanosky, O. Suciu, B. Edwards, and A. Sarabi, "Enhancing vulnerability prioritization: Data-driven exploit predictions with community-driven insights," *arXiv preprint arXiv:2306.14704*, 2023.

[15] U. Cybersecurity and I. S. Agency, "Known exploited vulnerabilities catalog," https://www.cisa.gov/known-exploited-vulnerabilities-catalog, 2023, accessed July 2023.

[16] Tenable, "Vulnerability priority rating technical guide," https://www.tenable.com/, 2023.

[17] Rapid7, "Vulnerability intelligence report 2022 edition," 2022.

[18] M. F. Bulut, A. Adebayo, D. Sow, and S. Ocepek, "Vulnerability prioritization: An offensive security approach," *arXiv preprint arXiv:2206.11182*, 2022.

[19] C. Sabottke, O. Suciu, and T. Dumitras, "Vulnerability disclosure in the age of social media: exploiting twitter for predicting real-world exploits," in *24th USENIX Security Symposium (USENIX Security 15)*. USENIX Association, 2015, pp. 1041–1056.

[20] Y. Yamamoto, D. Miyamoto, and M. Nakayama, "Text-mining approach for estimating vulnerability score," in *2015 4th International Workshop on Building Analysis Datasets and Gathering Experience Returns for Security (BADGERS)*. IEEE, 2015, pp. 67–73.

[21] G. Fridgen, L. Häfner, C. König, and T. Sachs, "Multicriteria decision framework for cybersecurity risk assessment and management," *Journal of Risk and Financial Management*, vol. 10, no. 3, p. 125, 2017.

[22] L. Wang, S. Noel, A. Singhal, and S. Jajodia, "Measuring security risk of networks using attack graphs," *International Journal of Next-Generation Computing*, vol. 1, no. 1, pp. 135–147, 2010.

[23] J. M. Spring *et al.*, "Prioritizing vulnerability response: A stakeholder-specific vulnerability categorization (version 2.0)," Carnegie Mellon University, Tech. Rep., 2021.

[24] M. Almukaynizi, E. Marin, E. Nunes *et al.*, "Darkmention: a deployed system to predict enterprise-targeted external cyberattacks," in *2018 IEEE International Conference on Intelligence and Security Informatics*, 2018, pp. 31–36.

[25] S. Samtani, H. Zhu, and H. Chen, "Proactively identifying emerging hacker threats from the dark web: A diachronic graph embedding framework," *ACM Transactions on Privacy and Security*, vol. 23, no. 4, pp. 1–33, 2020.

[26] NIST, "Guide for conducting risk assessments," National Institute of Standards and Technology, Tech. Rep. Special Publication 800-30 Revision 1, 2012.

[27] L. Allodi and F. Massacci, "A preliminary analysis of vulnerability scores for attacks in wild: the ekits and sym datasets," in *Proceedings of the 2012 ACM Workshop on Building analysis datasets and gathering experience returns for security*, 2012, pp. 17–24.

[28] U. Cybersecurity and I. S. Agency, "Reducing the significant risk of known exploited vulnerabilities," November 2021, federal Register Notice.