# Experimental Evaluation of Wazuh-Grafana Integration for Real-Time Cyber Threat Detection in Resource-Constrained Environments

**Achmad Sutanto [1]*, Arif Rakhman [2]***
* D3 Teknik Komputer, Politeknik Harapan Bersama
achmadsutanto@gmail.com [1], cakrakirana7@gmail.com [2]

**Article Info**

***Article history:***

Received 2025-07-22
Revised 2025-09-02
Accepted 2025-09-10

***Keyword:***

*Brute-force attack,*
*Cyber threat detection,*
*Grafana,*
*Real-time SIEM,*
*Wazuh.*

## ABSTRACT

This research evaluates the performance of integrating Wazuh, an open-source Security Information and Event Management (SIEM) platform, with Grafana, a real-time visualization tool, for cyber threat detection in resource-constrained environments. The objective is to assess detection accuracy, false positive rates, response times, and system efficiency under controlled experimental conditions. The testbed consisted of two virtual private servers (4 vCPUs, 4–8 GB RAM, 38–50 GB storage) and employed the CIC-IDS2017 dataset as a benchmark for simulating three representative attacks: brute-force, malware injection, and webshell exploitation. The results showed that the integrated system achieved 100% detection accuracy with 0% false positives across 30 trials, with an average total detection time of 3033 ms. Resource utilization remained low, with CPU usage below 35% and memory consumption under 25%, confirming feasibility for mid-range servers typical of small institutions. While these results underscore the system's efficiency, the findings must be interpreted within the limitations of a laboratory environment where predefined signatures were used. Performance in real-world networks with diverse traffic and unknown threats may differ, and further validation is required. This study makes two key contributions: (1) it provides the first structured quantitative benchmark of Wazuh-Grafana integration in constrained environments using a standardized dataset, and (2) it offers practical recommendations for small and medium-sized institutions, including minimum system requirements and guidelines for dashboard configuration. These findings reinforce the role of open-source solutions as affordable, adaptive, and effective alternatives to commercial SIEM systems, particularly for organizations with limited cybersecurity budgets.

This is an open access article under the CC–BY-SA license.

## I. INTRODUCTION

In the contemporary era of rapid and unceasing digital transformation, cybersecurity has unequivocally emerged as a persistent and critical global issue[1]. The proliferation of digital technologies across all sectors of society has expanded the attack surface, making organizations of all sizes susceptible to a diverse array of cyber threats. Malicious activities such as brute-force attacks, malware infections, and webshell exploitations are no longer exclusive threats to large corporations with extensive digital footprints[2]. Instead, there is a discernible trend of these attacks increasingly targeting small and medium-sized institutions, which often lack adequate security infrastructure to defend themselves effectively. This vulnerable group includes educational institutions, Micro, Small, and Medium Enterprises (MSMEs), and public sector organizations that are frequently constrained by limited budgets and a shortage of specialized cybersecurity expertise[3][4]. This disparity in defensive capabilities creates a significant security gap, which is systematically exploited by malicious actors to compromise systems, disrupt operations, and exfiltrate strategically valuable data. The consequences of such breaches extend beyond financial loss, encompassing reputational damage, loss of stakeholder trust, and significant operational downtime, thereby threatening the very sustainability of these vital organizations.

The escalating sophistication and volume of cyber threats necessitate a proactive and systematic approach to security monitoring. In response to this challenge, Security Information and Event Management (SIEM) systems have become a cornerstone of modern cybersecurity strategies. A SIEM solution provides a holistic view of an organization's security posture by collecting, aggregating, correlating, and analyzing log data from a multitude of sources across the network in real-time[5]. This centralized analysis enables security teams to detect anomalous activities, identify potential security incidents, and respond to threats before they escalate into major breaches. Among the plethora of available SIEM solutions, open-source platforms have gained significant traction due to their cost-effectiveness and flexibility.

Wazuh, a prominent open-source SIEM, offers a comprehensive suite of security features, including rule-based intrusion detection, file integrity monitoring, vulnerability detection, and in-depth log analysis. However, while Wazuh provides powerful detection capabilities, its native visualization tools are often considered a limitation, potentially hindering the rapid interpretation of security events[6]. To overcome this, integration with advanced visualization platforms like Grafana is often recommended. Grafana, on the other hand, is a highly flexible visualization platform widely adopted for monitoring infrastructure performance. When integrated with SIEM platforms, Grafana transforms complex log data into interactive dashboards that allow administrators to visualize attack patterns and comprehend security alerts in real-time[7][8]. For readers less familiar with these platforms, it is important to highlight their respective roles: Wazuh functions as the SIEM detection and correlation engine, while Grafana provides the visualization layer to interpret and monitor detected threats.

Despite these capabilities, adoption of Wazuh-Grafana integration among small and medium-sized institutions remains limited. Challenges include budgetary constraints that hinder investment in even low-cost infrastructure, and a lack of technical expertise to configure and maintain a rule-based security system effectively [9][10]. Consequently, a large number of these organizations continue to rely on rudimentary security measures, such as basic firewalls and antivirus software, or, in some cases, operate without any formal mechanism for attack detection at all. This reactive and often inadequate security posture leaves them highly exposed to a landscape of ever-evolving cyber threats. A further challenge lies in the absence of systems that can provide both early-warning alerts and comprehensible visualizations tailored for administrators who may not be cybersecurity specialists. This gap highlights a critical need for security solutions that are not only powerful and affordable but also accessible and manageable for non-expert users, empowering them to take control of their own cybersecurity defense without requiring extensive specialized training or resources.

The academic literature has acknowledged the individual effectiveness of both Wazuh and Grafana in specific security contexts [11][12]. However, a notable gap exists in systematic, quantitative performance evaluations of their integration under resource-constrained conditions. Previous studies have focused on deployment or configuration aspects [13], or cloud-based setups [8], but have not thoroughly assessed detection performance across multiple attack vectors using standardized datasets. This is a critical omission, as small institutions require empirical evidence to justify adoption decisions.

In this research, resource-constrained environments are defined as infrastructures operating on mid-range virtual private servers (VPS) with limited CPU, RAM, and storage resources, such as 4 vCPUs, 4–8 GB RAM, and 38–50 GB storage. These specifications reflect realistic conditions faced by MSMEs and educational institutions, where budget and infrastructure limitations restrict the feasibility of enterprise-grade SIEM systems.

To address this gap, this study conducts a controlled experimental evaluation of an integrated Wazuh-Grafana system in detecting three common cyber-attacks: brute-force login attempts, malware injections, and webshell exploitations. The experiments employ the CIC-IDS2017 dataset, widely recognized for its realistic and diverse traffic profiles[14]. Performance metrics measured include detection accuracy, false positive rates, response times, and resource utilization [12][15][16].

The novelty of this research lies in its structured, quantitative benchmarking of Wazuh-Grafana integration under constrained VPS environments. Unlike previous works that evaluated these systems in isolation [8], [16], this study systematically measures detection performance, response latency, and efficiency of system resource consumption. The contributions of this work are twofold:

- Empirical validation of Wazuh-Grafana as a viable, low-cost SIEM solution capable of achieving high accuracy in constrained environments.
- Practical recommendations for small and medium-sized institutions, including minimum system requirements, configuration guidelines, and future development directions (e.g., integration of active response).

By presenting replicable evidence, this study not only fills an academic gap in open-source SIEM evaluation but also delivers actionable insights for institutions seeking affordable and sustainable cybersecurity strategies [17] [18].

## II. Methodology

### A. Research Design

This study employs a quantitative, experimental research design conducted within a controlled laboratory environment. The primary objective of this design is to systematically evaluate the performance of an integrated Security Information and Event Management (SIEM) system, specifically Wazuh, with the Grafana visualization platform for real-time cyber-attack detection. The experiment focuses on measuring detection accuracy, false positive rate, response time, and system efficiency under three distinct cyberattack scenarios: brute-force login attempts, malware injections, and webshell exploitations.

To ensure external validity, the CIC-IDS2017 dataset was utilized as a foundational reference. This dataset is widely recognized in network anomaly detection research due to its realistic and diverse traffic profiles. By leveraging this dataset, the experiment ensures that the traffic used reflects both benign and malicious behaviors, providing a robust benchmark for SIEM performance. The entire experimental procedure is systematically outlined in the procedural flowchart presented in Figure 1.

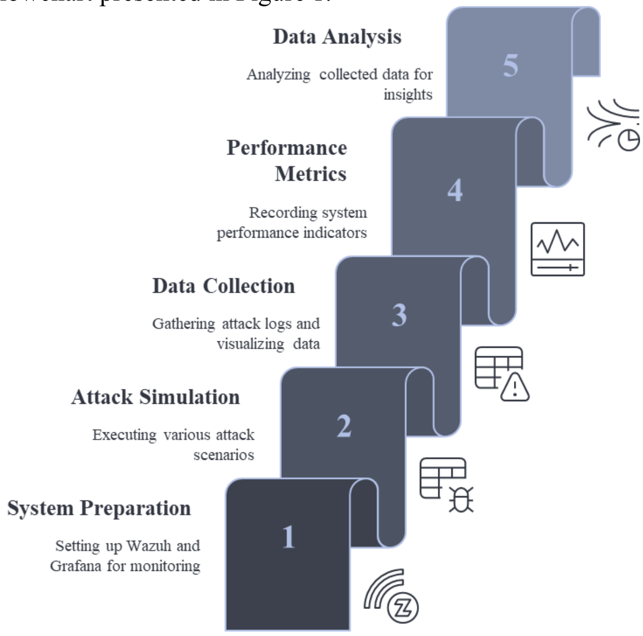
Data Analysis
Analyzing collected data for insights
5
Performance Metrics
Recording system performance indicators
4
Data Collection
Gathering attack logs and visualizing data
3
Attack Simulation
Executing various attack scenarios
2
System Preparation
Setting up Wazuh and Grafana for monitoring
1

Figure 1. Methodological Workflow

This controlled approach allows for the isolation of variables and precise measurement of performance metrics, minimizing confounding factors that would be present in a live production environment and thus ensuring the reliability and replicability of the results.

## B. System Architecture and Environment

The experimental setup was executed on two separate Virtual Private Servers (VPS), representing a resource-constrained environment that closely mirrors conditions faced by small and medium-sized institutions. The detailed specifications are presented in Table 2.

TABLE 2
VPS SPECIFICATIONS FOR EXPERIMENTAL ENVIRONMENT

| Component | VPS-1 (Wazuh-Grafana Server) | VPS-2 (Wazuh Agent) |
|---|---|---|
| vCPU | 4 | 4 |
| RAM | 8 GB | 4 GB |
| Storage | 50 GB | 38 GB |
| OS | Ubuntu Server 22.04 LTS | Ubuntu Server 22.04 LTS |
| Software | Wazuh 4.11,<br>Grafana 11.6 | Wazuh Agent 4.11 |

Both servers operated on the Ubuntu Server 22.04 Long-Term Support (LTS) operating system, providing a stable and well-documented foundation for the deployment. The attack simulations were launched from a separate, isolated node running Kali Linux, a standard platform for penetration testing and security auditing. Specific tools on the attacker node included Hydra for executing the brute-force attack scenarios and custom-developed PHP scripts designed to simulate the malware injection and webshell exploitation attacks. The network topology, illustrating the interaction between the attacker node, the monitored agent (VPS-2), and the central server (VPS-1), is depicted in Figure 2.

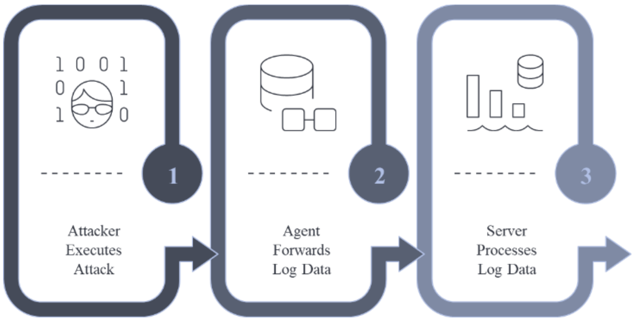

Figure 2. System Topology of the Test Environment

## C. Attack Scenarios and Sample Definition

The population for this study encompasses a broad range of cyber-attacks commonly observed in the networks of small and medium-sized institutions. From this population, a representative sample was selected, chosen specifically because they represent high-impact threats commonly faced by small institutions: credential theft (brute-force), web server compromise (webshell), and malicious file introduction (malware). Each attack type was executed ten times to ensure the consistency and statistical reliability of the collected data, resulting in a total of 30 experimental trials. The specific attack scenarios were:

- Brute-force Attack: This scenario involved a credential-guessing attack targeting the Secure Shell (SSH) service on the Wazuh Agent server (VPS-2). The attack was executed using the Hydra tool, with tests conducted across a range of thread variations (from 4 to 36 threads) to simulate different levels of attack intensity and to evaluate the system's performance under varying loads.
- Malware Injection: This was simulated by downloading and extracting a standard EICAR test file into critical system directories to test the system's reaction to a recognized, non-destructive malware signature. The file integrity monitoring module of Wazuh was configured to watch these directories, and the scenario was designed to

test the system's ability to detect unauthorized file modifications and the introduction of potentially malicious code.

- Webshell Exploitation: This scenario involved uploading a malicious PHP file to a designated 'uploads' directory on an Apache web server running on the Wazuh Agent server. The webshell was designed to provide a backdoor, allowing the attacker to execute arbitrary system-level commands through the web interface, thereby testing the system's capacity to identify command-and-control-like behavior originating from a web application.

### *D. Experimental Procedure*

The research was conducted through a structured, four-stage procedure to ensure systematic execution and data collection:

- System Preparation and Integration: The initial stage involved the installation of Wazuh version 4.11 and Grafana version 11.6 on VPS-1. This was followed by the configuration of the Wazuh components (Manager, Indexer, Dashboard) and the installation of the Wazuh Agent on VPS-2. A critical part of this stage was establishing the integration between Wazuh and Grafana, which involved setting up data sources and developing custom dashboards in Grafana to visualize security alerts and log data generated by Wazuh.
- Attack Scenario Implementation: In the second stage, the three defined attack scenarios were executed sequentially. The brute-force attack was launched using Hydra with varying thread counts. The malware simulation was performed via scripted download and extraction of a test file. The webshell attack was carried out by uploading the malicious PHP script to the Apache server. Each of the ten trials per attack type was initiated manually to ensure precise timing and observation.
- Data Collection: During and immediately after each attack trial, data was collected automatically. The Wazuh system logged all detected events, including rule IDs, timestamps, and attack details. Simultaneously, these events were visualized on the pre-configured Grafana dashboards. Key performance indicators were meticulously recorded for each trial.
- System Performance Monitoring: Throughout the attack simulations, the resource utilization of the Wazuh-Grafana server (VPS-1) was closely monitored. Data on CPU usage, memory consumption, and disk I/O were collected to assess the system's efficiency and to determine the performance overhead associated with detecting and processing the different types of threats.

### *E. Data Collection Techniques*

Data were gathered using both automated system logs and real-time visualizations. Specific parameters recorded included:

- Detection Accuracy (%): Ratio of successfully identified attack attempts to total attempts.
- False Positive Rate (%): Ratio of benign traffic incorrectly flagged as malicious. This was measured by analyzing CIC-IDS2017 benign traffic logs and verifying whether alerts were triggered without attacks.
- Initial Detection Time (ms): Latency from attack initiation to the first alert generated by Wazuh.
- Total Detection Time (ms): Time from attack initiation to full visualization in Grafana.
- System Resource Utilization (%): CPU, memory, and disk I/O usage on VPS-1 under different workloads.

All experimental results were logged in structured documentation forms, including variations in attack configurations, timestamps, and system performance details..

### *F. Data Analysis Methods*

The collected data were analyzed using both descriptive and inferential statistics.

- Descriptive statistics (mean, standard deviation) summarized detection accuracy, false positives, response times, and resource utilization.
- Inferential tests such as one-way Analysis of Variance (ANOVA) were planned to determine whether significant differences existed in detection performance among the three attack types. An alpha level of $\alpha = 0.05$ was set as the threshold for statistical significance.

This mixed descriptive-inferential approach ensured both an overview of system performance and a rigorous test of variability across different attack scenarios[14], [15].

## III. Result and Discussion

### *A. Results*

The experimental evaluation of the integrated Wazuh-Grafana system yielded definitive and highly consistent results across all planned attack scenarios. The primary findings demonstrate that the system is exceptionally capable of detecting the selected cyber threats with maximum accuracy and reliability within the controlled laboratory environment. This section presents a detailed account of the quantitative performance metrics recorded during the experiments, including overall detection efficacy, temporal performance, and system resource efficiency. The results are presented systematically to correspond with the research objectives, providing a clear and objective measure of the system's capabilities.

#### 1) Overall Detection Performance

The experiments demonstrated a consistently high detection capability of the integrated Wazuh-Grafana system. Across all 30 simulated attacks ten trials each for brute-force, malware injection, and webshell exploitation the system achieved a 100% detection accuracy, with every malicious action correctly identified and flagged by the Wazuh engine. In addition, the false positive rate remained at 0%, as no benign activities from the CIC-IDS2017 dataset triggered spurious alerts. This indicates that the rule-based detection engine was well-calibrated for the tested scenarios and did not misclassify legitimate traffic or operations as threats.

This perfect accuracy was consistently observed under all attack variations, including brute-force attempts executed with escalating intensity (4 to 36 threads), malware introduced into different system directories, and diverse forms of webshell uploads. Such results highlight the system's robustness in handling the predefined attack types.

However, while these top-level indicators confirm the effectiveness of Wazuh-Grafana under controlled laboratory conditions, they should be interpreted with caution. The flawless performance achieved here reflects the fact that the attacks were based on signatures already supported by Wazuh's rule set. In real-world production environments, where zero-day exploits, advanced persistent threats (APTs), and more complex traffic patterns exist, detection rates may not remain at this level. Thus, the findings demonstrate the maximum achievable accuracy in a constrained and well-defined testbed, serving as a benchmark rather than an absolute guarantee of performance in live systems.

A summary of this detection performance is presented in Table 2.

TABLE 2
SUMMARY OF DETECTION SYSTEM PERFORMANCE FOR WAZUH-GRAFANA

| Attack Type | Detection Accuracy (%) | False Positive (%) | Initial Detection Time (ms) | Total Detection Time (ms) |
|---|---|---|---|---|
| Brute-force | 100 | 0 | 1150 | 3050 |
| Malware Injection | 100 | 0 | 1070 | 3000 |
| Webshell | 100 | 0 | 1080 | 3050 |
| Average | 100 | 0 | 1100 | 3033 |

### 2) Temporal Performance Analysis

In addition to detection accuracy, the temporal performance of the system was an important focus of investigation. The results indicate that the integrated Wazuh-Grafana stack can generate alerts in a rapid and timely manner. On average, the initial detection time measured as the latency between attack initiation and the first alert triggered by Wazuh was approximately 1100 ms across all attack types. The average total detection time, defined as the duration required for the event to be processed and fully visualized on Grafana dashboards, was approximately 3033 ms. These values reflect an efficient pipeline, ranging from log collection at the agent level to dashboard visualization, and provide administrators with actionable intelligence within about three seconds of an attack event.

As summarized in Table 2, minor variations were observed between attack types. Malware injections were detected the fastest (1070 ms) due to the immediate response of the File Integrity Monitoring (FIM) module. Webshell exploitations followed at 1080 ms, while brute-force attacks required slightly longer (1150 ms), as alerts were only generated after the rule threshold for failed login attempts had been reached. Despite these small variations, the total detection times for all scenarios remained close to the 3000 ms benchmark. Figure 3 illustrates a comparative view of these results, showing the system's ability to maintain consistently low latency across different threat vectors.

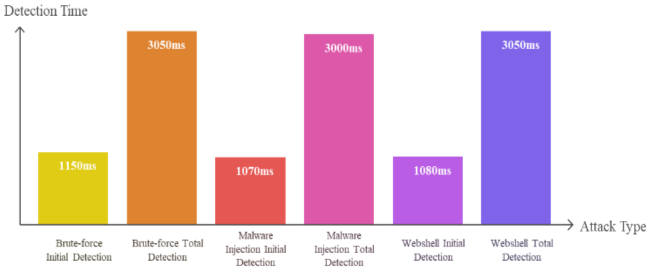

Figure 3. Comparison of Initial and Total Detection Times by Attack Type

Further statistical analysis confirmed that the mean total detection times 3050 ms for brute-force, 3000 ms for malware, and 3050 ms for webshell were highly consistent, with very low variance. At the $\alpha = 0.05$ significance level, no statistically significant differences were observed between the three attack categories ($p > 0.05$). This suggests that the detection-to-visualization process is not only rapid but also stable and reliable across different attack types.

Nevertheless, it is important to recognize the contextual limitations of these latency results. The near-identical response times are characteristic of a controlled laboratory environment with predefined signatures and limited background noise. In production networks with diverse user activities, larger log volumes, and concurrent events, detection latency may vary more significantly. Therefore, while the findings confirm that Wazuh-Grafana can deliver near real-time visibility under constrained conditions, further validation in live traffic environments is necessary to fully assess scalability and stability.

### 3) System Resource Utilization Efficiency

A critical aspect of this evaluation was to examine the performance overhead introduced by running Wazuh and Grafana on a single server, particularly given the focus on resource-constrained environments. Monitoring of the central server (VPS-1) indicated that the system operated with notable efficiency throughout the trials. Even during high-load scenarios, such as brute-force attacks executed with up to 36 concurrent threads, CPU utilization never exceeded 35%, while memory usage remained below 25% of the allocated 8 GB RAM. Disk I/O activity was minimal and did not approach critical thresholds.

This relatively low resource footprint is an important finding, as it demonstrates that the integrated Wazuh-Grafana solution can be effectively deployed on mid-range VPS or physical servers without requiring substantial hardware investment. For institutions constrained by limited IT budgets, these results highlight the feasibility of achieving near real-time cyber threat detection without the prohibitive costs typically associated with commercial SIEM systems.

However, these results should be interpreted cautiously. The observed efficiency reflects conditions in a controlled laboratory environment, with a limited range of attacks and traffic patterns. In real-world networks—where log volumes are significantly higher, concurrent processes are more complex, and user activities are unpredictable—resource utilization may vary and could be higher than reported here. Furthermore, the experiments focused solely on detection and visualization, without enabling additional features such as automated active response, which may introduce further overhead.

Overall, the findings suggest that Wazuh-Grafana provides an efficient and economically viable SIEM solution for small and medium-sized institutions, but future validation in production deployments is required to confirm scalability and long-term performance stability.

## *B. Discussion*

The results of this study provide strong empirical evidence that an integrated Wazuh-Grafana system can serve as a reliable and efficient solution for cyber threat detection, particularly in environments with limited resources. This section interprets the findings in the context of existing literature, examines their theoretical and practical implications, and highlights the limitations of the research while proposing directions for future investigations.

### 1) Interpretation of Findings

The achievement of 100% detection accuracy with a 0% false positive rate across 30 experimental trials is an encouraging indicator of the capability of a properly configured, rule-based SIEM system. These results affirm that Wazuh's detection engine, supported by its file integrity monitoring (FIM) capabilities, is highly effective at identifying common attack vectors in real time. Furthermore, the average total detection time of approximately 3033 ms demonstrates that integrating Grafana for visualization does not introduce significant latency; instead, it enables near real-time situational awareness.

Equally important, the system maintained low CPU (<35%) and memory (<25%) utilization during high-load tests, confirming its efficiency under constrained resources. Together, these findings directly address the operational barriers of cost and complexity that frequently hinder small and medium-sized institutions from adopting advanced security solutions.

These results are consistent with earlier research that highlighted Wazuh's ability to detect brute-force and malware attacks effectively [11][12]. However, this study makes a distinct contribution by moving beyond isolated evaluations of Wazuh or Grafana, and instead systematically quantifying the integrated performance of both platforms under standardized attack scenarios. Unlike prior works focused mainly on technical deployment [16] or cloud-based use cases[8], this research demonstrates how rule-based detection and interactive visualization can be combined to create a practical, lightweight, and replicable SIEM framework.

Nevertheless, it is crucial to interpret these findings within their experimental boundaries. The perfect detection rates achieved here reflect the controlled laboratory environment and the reliance on signatures already supported by Wazuh. In real-world networks with diverse traffic and previously unseen attack patterns, performance will likely differ. Thus, these results should be viewed as a benchmark of maximum achievable performance under constrained conditions, not as a guarantee of flawless detection in production settings.

### 2) Theoretical and Practical Implications

From a theoretical perspective, this research reinforces the conceptual model of integrating event-driven log analysis with real-time visualization as a robust cybersecurity paradigm. The results also challenge the perception that open-source SIEM solutions are inherently less reliable than their commercial counterparts [17]. Unlike commercial systems such as Splunk or IBM QRadar, which offer advanced analytics and integrated threat intelligence but require substantial licensing costs and high-performance infrastructure, the Wazuh-Grafana stack demonstrates that comparable baseline protection can be achieved on mid-range servers with no licensing fees. This positions Wazuh-Grafana as a practical low-cost alternative, particularly for institutions where cost-efficiency and lightweight deployment are more critical than enterprise-scale analytics. This aligns with the broader discussion in cybersecurity literature on the role of open-source tools in democratizing digital resilience[9][13].

Practically, the implications are highly significant. The demonstrated effectiveness of Wazuh-Grafana shows that resource-constrained institutions such as MSMEs and universities can deploy real-time threat detection capabilities without major infrastructure investments. Grafana's visualization further lowers the barrier to entry by providing intuitive dashboards, enabling IT generalists without deep security expertise to understand and respond to security incidents. In practice, administrators can begin with baseline panels such as "Failed Login Attempts," "File Integrity Alerts," and "Webshell Uploads," which were tested in this study and proved sufficient for monitoring the most frequent threats. These dashboards can be gradually customized with additional panels for network activity, vulnerability scans, or incident trends, depending on institutional needs. In this way, the research provides not only a conceptual blueprint but also concrete implementation guidance for organizations with limited budgets to strengthen their security posture. In this way, the research provides a practical blueprint for organizations with limited budgets to strengthen their security posture. This is consistent with global trends emphasizing the adoption of inclusive, open-source technologies to foster digital sovereignty, particularly in developing nations[18].

### 3) Limitations of the Study

Despite the encouraging results, several limitations must be acknowledged:

- Controlled Environment: The experiments were conducted in a laboratory setting. Real-world environments feature greater variability in traffic, user behavior, and system complexity, which may lead to false positives not observed here.
- Scope of Attacks: Only three types of attacks (brute-force, malware injection, webshell) were tested. While representative, these do not cover advanced persistent threats (APTs), multi-vector ransomware, or zero-day exploits.
- Short-Term Evaluation: The study was limited to short-term trials. Long-term performance, scalability under increasing traffic, and resilience to configuration updates were not assessed.
- Focus on Detection and Visualization: Active response capabilities, such as automated host isolation, firewall integration, or IP blocking, were not enabled. These features are critical for building a complete defensive ecosystem.

These limitations highlight that while the study establishes a strong performance baseline, further work is required to validate robustness under more complex, real-world conditions [14][15].

4) Recommendations for Future Research

Future research should aim to validate these findings in real-world production environments, such as university campuses or SME networks, to assess scalability under diverse traffic conditions, while also expanding the scope of attack scenarios to include advanced threats such as APTs, ransomware, and fileless malware through approaches like red teaming. Further investigations are needed to evaluate the integration of Wazuh's active response features such as automated IP blocking, user account isolation, and firewall coordination and to examine their impact on both detection efficacy and resource consumption. In addition, longitudinal studies should be conducted to measure long-term stability, scalability, and adaptability under continuous system updates and traffic growth. Finally, to encourage broader adoption, future work should focus on developing localized implementation guides and best-practice documentation tailored to the needs of non-specialist administrators, particularly in resource-constrained and developing contexts.

## IV. Conclusion

### A. Conclusion

This research successfully evaluated the performance of an integrated Wazuh-Grafana system in detecting three common cyber threats brute-force, malware injection, and webshell exploitation through controlled laboratory experiments using the CIC-IDS2017 dataset. The results demonstrated a detection accuracy of 100% with no false positives, and an average total detection time of 3033 ms, confirming that the system is capable of delivering near real-time alerts. Furthermore, the system maintained low resource consumption (CPU < 35%, memory < 25%), validating its feasibility for deployment on mid-range servers that are typical of resource-constrained environments. These findings provide strong evidence that Wazuh, as a rule-based SIEM, when integrated with Grafana's visualization capabilities, can serve as a reliable, efficient, and affordable solution for small and medium-sized institutions. Importantly, this study presents the first structured experimental evaluation of Wazuh-Grafana integration in resource-constrained VPS environments using the CIC-IDS2017 dataset, offering a replicable benchmark that has not been systematically reported in prior literature.

However, it is important to interpret these findings within their context: the flawless detection record reflects a controlled testbed with predefined attack signatures and limited traffic diversity, and may not directly translate to production networks where traffic patterns are more complex and threats more varied. Thus, while this study establishes a robust benchmark for Wazuh-Grafana's performance under constrained conditions, further validation in real-world environments remains necessary.

### B. Recommendations

Future work should focus on replicating this evaluation in production environments to test robustness under live traffic, expanding the scope of attack scenarios to include advanced threats such as APTs, ransomware, and zero-day exploits, and exploring the integration of Wazuh's active response features (e.g., automated blocking and firewall coordination) to enhance mitigation capabilities. In addition, longitudinal studies are needed to evaluate scalability and stability under continuous updates and increasing log volumes. Finally, the development of localized implementation guides and training resources for example, documentation in Bahasa Indonesia tailored to institutional case studies would greatly support adoption by non-specialist administrators in resource-constrained contexts.

## Acknowledgements

The author expresses appreciation to Politeknik Harapan Bersama for the funding and facilitation of this research through the Beginner Lecturer Research scheme for the 2024/2025 Academic Year. Gratitude is also extended to the Head of the Research and Community Service Center and the Head of the DIII Computer Engineering Study Program for their guidance and support throughout the research process. The author also thanks the laboratory technical team and the research team members, Arif Rakhman and Maskuri, for their significant contributions to the system installation, attack simulations, and technical documentation. Finally, appreciation is given to the user and developer communities of Wazuh and Grafana, whose open documentation was invaluable during the configuration and testing phases of this study.

## REFERENCES

[1] T. Liebetrau and L. Monsees, “Cybersecurity and International Relations: developing thinking tools for digital world politics,” *Int. Aff.*, vol. 100, no. 6, pp. 2303–2315, Nov. 2024, doi: 10.1093/ia/iiae232.

[2] S. S. Tirumala, N. Nepal, and S. K. Ray, “Raspberry Pi-based Intelligent Cyber Defense Systems for SMEs and Smart-homes: An Exploratory Study,” *EAI Endorsed Trans. Smart Cities*, vol. 6, no. 18, p. e4, Aug. 2022, doi: 10.4108/eetsc.v6i18.2345.

[3] A. Piazza, S. Vasudevan, and M. Carr, “Cybersecurity in UK Universities: mapping (or managing) threat intelligence sharing within the higher education sector,” *J. Cybersecurity*, vol. 9, no. 1, Jan. 2023, doi: 10.1093/cybsec/tyad019.

[4] B. Wibowo, A. Nurrohman, and L. Hafiz, “Deep Learning in Wazuh Intrusion Detection System to Identify Advanced Persistent Threat (APT) Attacks,” *Int. J. Sci. Educ. Cult. Stud.*, vol. 4, no. 1, pp. 1–10, Jan. 2025, doi: 10.58291/ijsecs.v4i1.311.

[5] C. MACANEATA, “Overview of Security Information and Event Management Systems,” *Inform. Econ.*, vol. 28, no. 1/2024, pp. 15–24, Mar. 2024, doi: 10.24818/issn14531305/28.1.2024.02.

[6] M. R. Islam and R. Rafique, “Wazuh SIEM for Cyber Security and Threat Mitigation in Apparel Industries,” *Int. J. Eng. Mater. Manuf.*, vol. 9, no. 4, pp. 136–144, Oct. 2024, doi: 10.26776/ijemm.09.04.2024.02.

[7] Jumiaty and B. Soewito, “SIEM and Threat Intelligence: Protecting Applications with Wazuh and TheHive,” *Int. J. Adv. Comput. Sci. Appl.*, vol. 15, no. 9, 2024, doi: 10.14569/IJACSA.2024.0150923.

[8] S. Moiz, A. Majid, A. Basit, M. Ebrahim, A. A. Abro, and M. Naeem, “Security and Threat Detection through Cloud-Based Wazuh Deployment,” in *2024 IEEE 1st Karachi Section Humanitarian Technology Conference (KHI-HTC)*, Jan. 2024, pp. 1–5, doi: 10.1109/KHI-HTC60760.2024.10482206.

[9] M. Nas, F. Ulfiah, and U. Putri, “Analisis Sistem Security Information and Event Management (SIEM) Aplikasi Wazuh pada Dinas Komunikasi Informatika Statistik dan Persandian Sulawesi Selatan,” *J. Teknol. Elekterika*, vol. 20, no. 2, p. 92, Nov. 2023, doi: 10.31963/elekterika.v20i2.4536.

[10] C. Kurniawan and A. Triayudi, “Reconstruction and Detection of Gambling Web Defacement Attack Using Wazuh and Velociraptor,” in *2024 International Conference on Information Technology Research and Innovation (ICITRI)*, Sep. 2024, pp. 257–262, doi: 10.1109/ICITRI62858.2024.10699215.

[11] Z. S. Younus and M. Alanezi, “Detect and Mitigate Cyberattacks Using SIEM,” in *2023 16th International Conference on Developments in eSystems Engineering (DeSE)*, Dec. 2023, pp. 510–515, doi: 10.1109/DeSE60595.2023.10469387.

[12] F. I. F. Farrel, M. K. Is Mardianto, S.Si, and M. Ir. Adrian Sjamsul Qamar, “Implementation of Security Information &amp; Event Management (SIEM) Wazuh with Active Response and Telegram Notification for Mitigating Brute Force Attacks on The GT-I2TI USAKTI Information System,” *Intelmatics*, vol. 4, no. 1, pp. 1–7, Feb. 2024, doi: 10.25105/itm.v4i1.18529.

[13] M. Monteros, J. F. Chuqui Quille, N. Benitez-Cacao, and P. Velez-Guerrero, “Implementar un sistema de gestión y análisis de seguridad con la herramienta Wazuh, en el Instituto Superior Universitario Tecnológico del Azuay,” *Atenas Rev. Científica Técnica y Tecnológica*, vol. 3, no. 1, 2024, doi: 10.36500/atenas.3.006.

[14] Kurniabudi, D. Stiawan, Darmawijoyo, M. Y. Bin Idris, A. M. Bamhdi, and R. Budiarto, “CICIDS-2017 Dataset Feature Analysis With Information Gain for Anomaly Detection,” *IEEE Access*, vol. 8, pp. 132911–132921, 2020, doi: 10.1109/ACCESS.2020.3009843.

[15] F. D. Utami and F. D. Astuti, “Comparison of Hadoop Mapreduce and Apache Spark in Big Data Processing with Hgrid247-DE,” *J. Appl. Informatics Comput.*, vol. 8, no. 2, pp. 390–399, Nov. 2024, doi: 10.30871/jaic.v8i2.8557.

[16] M. R. A. Suhendi, Alfarizi, A. A. Sukmandhani, and Y. D. Prabowo, “Network Anomaly Detection Analysis using Artillery Honeypot and Wazuh SIEM,” in *2023 IEEE 9th International Conference on Computing, Engineering and Design (ICCED)*, Nov. 2023, pp. 1–6, doi: 10.1109/ICCED60214.2023.10425009.

[17] R. Amami, M. Charfeddine, and S. Masmoudi, “Exploration of Open Source SIEM Tools and Deployment of an Appropriate Wazuh-Based Solution for Strengthening Cyberdefense,” in *2024 10th International Conference on Control, Decision and Information Technologies (CoDIT)*, Jul. 2024, pp. 1–7, doi: 10.1109/CoDIT62066.2024.10708476.

[18] C. Bassey, E. T. Chinda, and S. Idowu, “Building a Scalable Security Operations Center: A Focus on Open-source Tools,” *J. Eng. Res. Reports*, vol. 26, no. 7, pp. 196–209, Jun. 2024, doi: 10.9734/jerr/2024/v26i71203.