JURNAL LOCUS: Penelitian & Pengabdian
Volume 5 No. 8 August 2026
E-ISSN 2829-7334| P-ISSN 2829-5439
Hompage: https://locus.rivierapublishing.id/index.php/jl

A Comparative Effectiveness Analysis of Signature-Based IDS and Network Detection and Response (NDR) in a SIEM Environment

Christian Hary*, Muhammad Salman
Universitas Indonesia, Indonesia
Email: christian.hary@ui.ac.id*, muhammad.salman@ui.ac.id

Keywords:
Behavioral Metadata; Encrypted Traffic Analysis; NDR; NIDS; Security Onion; Visibility Gap.

Abstract
As cyber threats increasingly employ cryptographically concealed and stealthy communication, traditional signature-based Network Intrusion Detection Systems (NIDS) encounter severe limitations in achieving end-to-end operational visibility. This study delivers a quantitative comparative analysis contrasting signature-based frameworks against metadata-driven Network Detection and Response (NDR) within a unified Security Information and Event Management (SIEM) platform. Configured with a massive ruleset of 47,192 active signatures, the NIDS engine was evaluated against the behavioral metadata abstraction of NDR across five structured attack scenarios mapping the Cyber Kill Chain. Experimental results reveal a critical Visibility Gap in post-exploitation defenses; while NIDS achieved a 60\% Detection Rate by intercepting noisy initial reconnaissance and exploit payloads, it suffered total blindness during post-exploitation maneuvers. Conversely, the NDR engine achieved a 100\% Detection Rate, isolating deterministic forensic indicators including a 716.33-second encrypted command and control (C2) session and a 7.07 MB volumetric data exfiltration burst. Furthermore, architectural benchmarking revealed that massive rule compilation induced structural management plane synchronization failures. These findings synthesize into validated recommendations for sensor optimization, establishing an operational framework for specialized role allocation and computational efficiency to eliminate visibility blind spots within enterprise Security Operations Center (SOC) environments.

INTRODUCTION

The exponential growth of complex cyber threats has heightened the urgency for the implementation of a robust Network Intrusion Detection System (NIDS) to protect modern corporate environments (Ghazi et al., 2024). For Small and Medium Enterprises (SMEs) operating under budget constraints and limited cybersecurity personnel, implementing a cost-efficient yet resilient security architecture remains an ongoing operational challenge (Bhatia & Almukhaini, 2024; Hadi et al., 2024). This need has driven the widespread adoption of integrated open-source security platforms, exemplified by Security Onion, which is capable of integrating threat hunting, log management, and network auditing capabilities into a single framework (Mocanu & Scripcariu, 2023). However, with the increasing volume and heterogeneity of contemporary network traffic, optimizing the configuration of the primary detection engine within it has become a non-trivial issue for both practitioners and researchers striving to maximize defense coverage (Waleed et al., 2022).

Traditional network defense heavily relies on signature-based inspection engines, with Snort and Suricata as the foundation of modern security infrastructure technology. While Snort

is well-known for its extensive rule-based detection heuristics, Suricata introduces a multi-threaded architecture explicitly optimized for high-speed traffic extraction (Ghazi et al., 2024). However, the empirical performance of these signature-based systems is heavily dictated by the structural complexity and the accumulation of the volume of the activated ruleset (Raharjo et al., 2022). Previous research by Raharjo et al. (2022) has demonstrated that the scaling up of active signatures particularly thru the use of threat intelligence feeds or large-volume IP reputation lists can trigger serious performance degradation, resulting in a structural decrease in processing throughput ranging from 14% to 84%, significant packet drops, and critical reductions in true-positive detection accuracy in high-speed network environments (Waleed et al., 2022; Hu et al., 2020; Alka & Sharma, 2020). Consequently, modern Security Operations Centers (SOCs) are constantly faced with the dilemma of balancing the comprehensiveness of detection rules and the hardware limitations of their monitoring devices (Alka & Sharma, 2020).

Beyond hardware constraints, the strategic shift of modern threat actors toward "low and slow" attack methodologies based on Advanced Persistent Threat (APT) poses an existential threat to traditional signature-based frameworks (Sai et al., 2022). Attackers are increasingly exploiting covert communication channels, maintaining persistent low-throughput connections, and avoiding large-volume data transfers to manipulate and evade conventional flooding-based detection thresholds (Sai et al., 2022; Pranav et al., 2024). These manipulative techniques produce traffic anomalies that closely resemble legitimate user behavior, making it difficult for analysts to separate malicious activity from network background noise (Sai et al., 2022). Furthermore, attackers often target privileged access and sensitive data while operating entirely within encrypted channels thru cryptographic handshake mechanisms and TLS fingerprinting to conceal their post-exploitation activities from conventional pattern-matching inspection machines (Rahman et al., 2024; Mohammed et al., 2022; Koumar & Cejka, 2022). The global transition toward advanced Transport Layer Security (TLS) 1.3 protocols complicates this further by encrypting handshake elements, causing a structural decline in standard signature-based traffic classification accuracy (Barradas et al., 2024; Zhou et al., 2024). To regain defensive coverage against these hidden streams, modern architectures must leverage granular metadata parsing combined with statistical flow metrics to fingerprint malicious sessions without relying on full packet decryption (Muttaqien et al., 2025). The transition toward these silent behavior-based tactics allows advanced attacks to remain within the company's perimeter for an extended period after the initial exploitation phase has been successfully carried out (Bhardwaj et al., 2024; Salazar, 2026).

A substantial Visibility Gap often arises during this post-exploitation phase, particularly in the context of Command and Control (C2) beaconing and hidden data exfiltration. Although traditional NIDS engines remain highly effective in detecting known exploitation payloads during initial infiltration, these systems structurally fail to identify subsequent Remote Access Trojan (RAT) activities or periodic polling behavior (Piet et al., 2018). Because these persistent activities do not match the predefined static signatures, they can easily evade the primary pattern matching mechanism, thereby eliminating the end-to-end comprehensive view for security analysts regarding network events (Berger et al., 2016). This fundamental limitation underscores the need for an integrated defense approach that combines passive network

monitoring with advanced behavioral auditing to capture indicators that are missed by signature classification (Piet et al., 2018; Berger et al., 2016).

To address the architectural visibility gap, the modern defense paradigm is shifting toward non-intrusive flow-based metadata analysis using Network Detection and Response (NDR) technology, with a primary focus on Zeek telemetry. Unlike packet-based systems that rely on rigid pattern matching, NDR leverages network metadata and data flow abstractions to identify periodic behavioral anomalies, even when operating within fully encrypted data flows (Koumar & Cejka, 2022; Koumar et al., 2023). By extracting highly granular features such as inter-arrival time variance and byte distribution patterns, NDR establishes a rigorous and explainable mathematical framework for network anomaly detection (Rahman et al., 2024; Hore et al., 2023). Additionally, mathematical optimization thru the reduction of telemetry features to highly relevant attributes such as the specific columns `network.bytes` and `event.duration` has proven successful in drastically minimizing system computational load while maintaining detection accuracy (Rabbani et al., 2025; Iglesias & Zseby, 2015). This two-layer strategy fundamentally enables security sensors to identify indicators of anomalous behavior in the absence of previously defined specific signatures (Rabbani et al., 2025).

Although various academic models have proposed isolated network detection frameworks (Hajj et al., 2023), there is a scarcity of empirical literature examining the practical operational compromises and system stability when a massive signature database is forced to run alongside behavioral metadata. Previous research generally limited the scope to the scalability of random rules (Raharjo et al., 2022) or restricted the detection of encrypted traffic to a single local protocol (Rahman et al., 2024; Pranav et al., 2024), without considering the synchronization degradation experienced by the SIEM management plane throughout the entire Cyber Kill Chain lifecycle to prevent data leaks (Bhatia & Almukhaini, 2024; Mocanu & Scripcariu, 2023; Mohammed et al., 2022; Berger et al., 2016; Majigi et al., 2025).

This research directly fills the state-of-the-art gap by presenting a multi-layer empirical evaluation of the performance of NIDS and NDR within a production-level SIEM Security Onion testbed. Using a controlled laboratory infrastructure, a massive injection of 47,192 active signatures was evaluated against granular metadata logging in five different attack scenarios. The main contributions of this research include three aspects. First, we present a quantitative evaluation of the operational and computational trade-offs faced when imposing high-volume rulesets on signature-based machines. Second, we isolate the precise behavioral limitations of the Visibility Gap by providing empirical evidence regarding critical compromise indicators, such as explicit session duration and volumetric bit spikes that are overlooked by pattern matching engines. Finally, these empirical findings are synthesized into validated architectural recommendations for the allocation of specialized sensor resources in the corporate SOC environment to balance detection accuracy and computational efficiency across all stages of the Cyber Kill Chain. The mapping of the comparison matrix of this research with the state-of-the-art literature is systematically presented in Table 1.

Table 1. State-of-the-Art Literature Comparison Matrix

Researcher / Reference	Signature Volume	Encrypted Traffic Analysis (ETA)	Post-Exploit C2 & Exfiltration	Computational Trade-off
Ghazi et al.	Low	No	No	Basic (CPU/RAM)
Raharjo et al.	Medium	No	No	Throughput Degradation
Pranav et al.	None	Yes (TLS Only)	C2 Only	None
Rahman et al.	None	Yes (Flow)	Exfiltration Only	None
This Research	Massive (47.192)	Yes (Metadata)	Full Kill Chain	SIEM and Flow Sync

Source: Literature synthesis, 2026

Based on the identified research gap and novelty, this research aims to conduct a quantitative comparative analysis of the detection performance between signature-based NIDS (Suricata) and metadata-driven NDR (Zeek) within a unified SIEM platform (Security Onion) across five structured attack scenarios mapping the Cyber Kill Chain, to measure and isolate the precise boundaries of the Visibility Gap during post-exploitation phases specifically evaluating the detection rates for encrypted Command and Control (C2) beaconing and unauthorized data exfiltration, to evaluate the computational and operational trade-offs imposed by massive signature compilation (47,192 active rules) on the SIEM management plane including synchronization failures and system stability degradation, to extract deterministic forensic indicators from behavioral metadata (session duration, periodicity, and volumetric bit spikes) that maintain visibility when static signature engines remain silent, and to formulate validated architectural recommendations for sensor optimization establishing an operational framework for specialized role allocation and computational efficiency to eliminate visibility blind spots within enterprise SOC environments.

This research is expected to provide theoretical and practical benefits by enriching the cybersecurity literature with empirical evidence on the operational compromise between signature-based and behavioral detection paradigms while filling the gap in research on the performance of massive signature databases in SIEM environments, and by providing scientific evidence for cybersecurity practitioners and SOC analysts in designing more effective layered defense architectures that balance signature comprehensiveness and hardware efficiency, serving as a guide for organizations in determining the strategic allocation of sensor resources, offering recommendations for SMEs operating under budget constraints, opening opportunities for future researchers to integrate machine learning models into NDR telemetry pipelines and conduct comparative studies across different SIEM platforms or encryption protocols, and serving as a reference for policymakers and regulators in developing cybersecurity standards that emphasize the importance of behavioral metadata analysis alongside signature-based detection in enterprise security frameworks.

METHOD

Testbed Environment and Network Architecture

The experimental environment in this research was built within a controlled virtualized laboratory infrastructure to ensure data integrity, strict network isolation, and testing replication

capability. The network architecture is orchestrated using Graphical Network Simulator-3 (GNS3) integrated with Oracle VM VirtualBox, dividing the topology into three main functional segments. The offensive perimeter is occupied by the Attacker Node running the Kali Linux operating system with a static IP address of 192.168.1.10. This node is equipped with specialized penetration testing tools, including Nmap for network mapping, Hydra for credential cracking, and the Metasploit Framework for delivering exploit payloads. Conversely, the defensive perimeter houses the Target Node implemented thru Metasploitable 2 at the IP address 192.168.1.20, serving as a high-interaction honeypot with services deliberately left vulnerable, such as Secure Shell (SSH), Hypertext Transfer Protocol (HTTP), and PHP-based web applications.

The central component of the defence architecture is the Monitoring Node, which functions as a production-scale Security Operations Center (SOC) device using the Security Onion platform version 2.4. To optimize packet processing and management isolation, the primary detection engines Suricata and Zeek are configured using a specialized multi-interface network abstraction. The first interface, `enp0s3`, is exclusively allocated for administrative management plane traffic. The second interface, `enp0s8`, is configured in promiscuous mode to act as a passive Test Access Point (TAP) or a Switched Port Analyzer (SPAN) mirror link within GNS3, capturing all raw traffic passing between the attacker and the victim.

The third interface, `enp0s9`, is mapped to the Network Address Translation (NAT) machine to handle ruleset synchronization and external updates. To evaluate the upper limits of NIDS stability and model the complex detection posture of the company, the Suricata engine is injected with a massive database consisting of 47,192 active signatures sourced from the Emerging Threats (ET) Open feed as well as local custom rules.

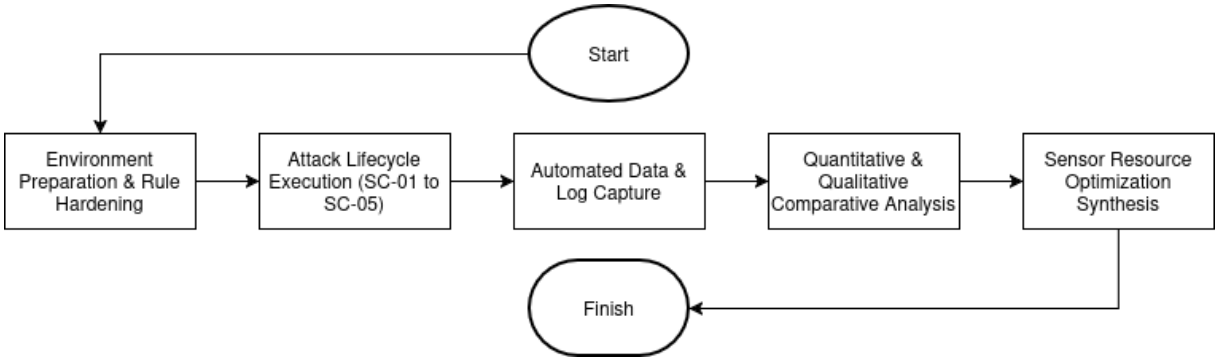

Figure 1. Operation Research Workflow

Source: Research design, 2026

Operational Workflow of the Research

The operational execution of this research is carried out systematically thru a structured five-phase methodology to map architectural vulnerabilities and visibility thresholds, as illustrated in Fig. 1. The workflow begins with the Environment Preparation Phase, which focuses on asset hardening, network baseline establishment, and the tactical compilation of 47,192 active Suricata signatures. This step is followed by the Attack Execution Phase, where five different scenarios (coded SC-01 to SC-05) are launched sequentially from offensive nodes to replicate the complete lifecycle of a cyber-attack.

During the next Data Capture Phase, raw alerts and connection transaction metadata generated by the traffic flow are continuously harvested from the Elasticsearch backend

database. The captured artifacts are then submitted to the rigorous Comparative Analysis Phase, where detection logs from both machines are cross-referenced against the actual execution timeline to explicitly isolate operational visibility failures. The final Synthesis Phase focuses on evaluating computational and procedural indicators to formulate validated architectural recommendations for sensor optimization, establishing an operational balance between the comprehensiveness of signatures and hardware efficiency.

Table 2. Attack Lifecycle Execution Parameters and Analysis Targets

ID	Kill Chain Phase	Technique / Tool	Payload Parameters	Target Detection Indicator
SC-01	Reconnaissance	Nmap Heuristics	nmap -sV -O -T4	Signature-based Scan Patterns
SC-02	Initial Access	Hydra Brute Force	hydra ssh://[Target] -t 4	Login Failure Frequency Logs
SC-03	Exploitation	Metasploit RCE	CVE-2012-1823 (PHP-CGI)	Malicious Exploit Strings
SC-04	Command & Control	Encrypted Meterpreter	Interval Heartbeat 15s	Session Duration and Periodicity
SC-05	Data Exfiltration	Automated Download	download -r /etc	Volumetric Flow and Byte Counts

Source: Experimental data, 2026

Attack Scenarios and Tactical Execution

The experimental phase throws five structured attack scenarios that map the core stages of the Cyber Kill Chain into the testbed environment. Each scenario is specifically designed to trigger deterministic signature alerts in Suricata while forcing the creation of behavioural telemetry in Zeek. Execution parameters, targeted protocols, and analysis targets are systematically arranged in Table 2 and can be explained as follows:

1. Reconnaissance and Initial Access (SC-01 & SC-02)

The initial phase targets the discovery of attack surface and credential exploitation. Scenario SC-01 executes an aggressive TCP SYN scan using Nmap with the parameters `-sV -O -T4` to identify open sockets, service versions, and the host operating system. This scenario evaluates the system's ability to classify high-interaction traditional signature scan patterns. Immediately after reconnaissance, Scenario SC-02 simulates an initial access attempt thru a dictionary-based brute force attack targeting the victim's SSH service using Hydra with a structured wordlist and a parallel thread limit (`-t 4`). These scenarios establish a baseline to verify whether the massive signature database can efficiently process high-intensity noise alerts generated during the initial stages of intrusion.

2. Targeted Exploitation (SC-03)

To model a critical server breach, Scenario SC-03 executes a targeted web exploitation scenario that leverages the historical PHP-CGI argument injection vulnerability (CVE-2012-1823) on the target host. The Metasploit Framework is utilized to send the outbound stager payload `php/meterpreter/reverse_tcp`. This stage explicitly tests the capacity of the NIDS machine to execute deep packet inspection (DPI) across the entire HTTP body to identify specific and deterministic attack strings associated with known public exploits.

3. Post-Exploitation: C2 dan Exfiltration (SC-04 & SC-05)

The final stage evaluates persistence and data theft targets within a blind network segment. Scenario SC-04 establishes an encrypted Command and Control (C2) channel thru a Meterpreter session configured with a strict heartbeat interval of 15 seconds. This connection is actively maintained for an extended duration to evaluate the behavioural detection capabilities of NDR metadata related to session duration and periodicity in the absence of absolute public traffic signatures.

Closing the attack lifecycle, Scenario SC-05 executes a recursive directory download (`download -r /etc`) to simulate unauthorized data exfiltration. This scenario tests the system's ability to detect large-scale data movement purely based on volumetric flow analysis, directed bit metrics, and host connection duration.

4. Evaluation Metrics and Visibility Matrix Abstraction

To mathematically measure detection effectiveness, this research uses two validation frameworks consisting of the quantitative metric Detection Rate and the qualitative Visibility Matrix. The Detection Rate is formally defined as the percentage ratio of successfully identified attack scenarios to the total number of simulated trials, expressed thru the following equation:

$$DR = \left(\frac{\Sigma Dc}{Ts}\right) \times 100\%$$

Equation 1

Where Dc represents the integer count of the scenario phases that were successfully detected, and Ts indicates the total number of scenarios executed in the empirical testing.

To map the precise operational boundaries where detection failures occur, the test results are qualitatively mapped into a structured three-level Visibility Matrix. The first level, referred to as Full Visibility, defines the optimal status where attack scenarios simultaneously trigger explicit alerts in Suricata and generate corresponding behavioural transaction telemetry in Zeek. The second level, classified as Visibility Gap, indicates a critical defence failure where the signature-based NIDS engine fails to log alerts, but the NDR engine successfully captures complete forensic metadata from the anomaly. The third level, defined as a Blind Spot, represents a critical security status where the attacker's actions successfully bypass both detection machines, leaving the SOC completely unaware of the situation.

RESULTS AND DISCUSSION

The experimental execution of five attack lifecycle scenarios resulted in the accumulation of quantitative data parsed from the Elasticsearch Security Onion backend. This test measures the presence of signature-based alerts from the NIDS engine (Suricata) and the logging of connection transaction telemetry from the NDR engine (Zeek). The results of the structured quantitative testing, which map the performance of both detection paradigms, are comprehensively summarized in Table 3.

Table 3. Empirical Matrix of Quantitative Detection Results

ID	Attack Scenario Phase	NIDS Alert (Suricata)	NDR Log (Zeek)	Architectural Visibility Status
SC-01	Aggressive Port Scanning (Nmap)	Yes	Yes	Full Visibility (Optimal)
SC-02	SSH Credential Brute Force (Hydra)	Yes	Yes	Full Visibility (Optimal)
SC-03	PHP-CGI Argument Injection (RCE)	Yes	Yes	Full Visibility (Optimal)
SC-04	Encrypted Meterpreter C2	No	Yes	Visibility Gap (Critical Failure)
SC-05	Volumetric Data Exfiltration	No	Yes	Visibility Gap (Critical Failure)

Source: Experimental data processed from Elasticsearch Security Onion backend, 2026

Based on the validation of the empirical logs, the Detection Rate for each specific sensor is calculated using the mathematical formulation established in Sub-chapter 2.4. For the signature-based NIDS engine (Suricata) operating under extreme rule load, the calculation yields a limited metric of:

$$\mathrm{DR} = \left(\frac{3}{5}\right) \times 100\% = 60\%$$

On the other hand, the metadata behavior-based NDR engine (Zeek) demonstrates comprehensive coverage by achieving optimal metrics of:

$$\mathrm{DR} = \left(\frac{5}{5}\right) \times 100\% = 100\%$$

The very clear quantitative divergence between the performance of Suricata and Zeek is visually plotted in a bar chart in Figure 2.

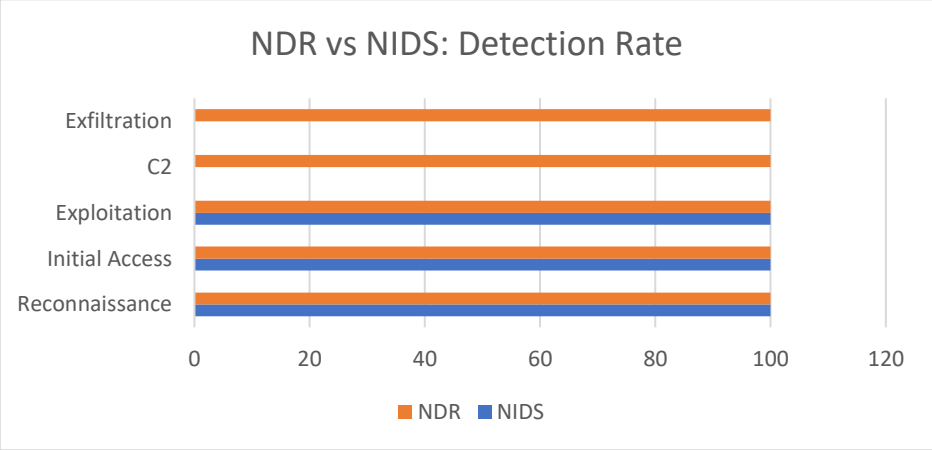

Figure 2. Detection Rate Comparison Bar Chart

Source: Experimental data processing results, 2026

To extract granular forensic data successfully retained by the NDR machine during the pattern matching engine failure, a combination of connection temporal duration and volumetric

properties of the bit stream extracted from the Zeek transaction logs is plotted along the test timeline in Figure 3.

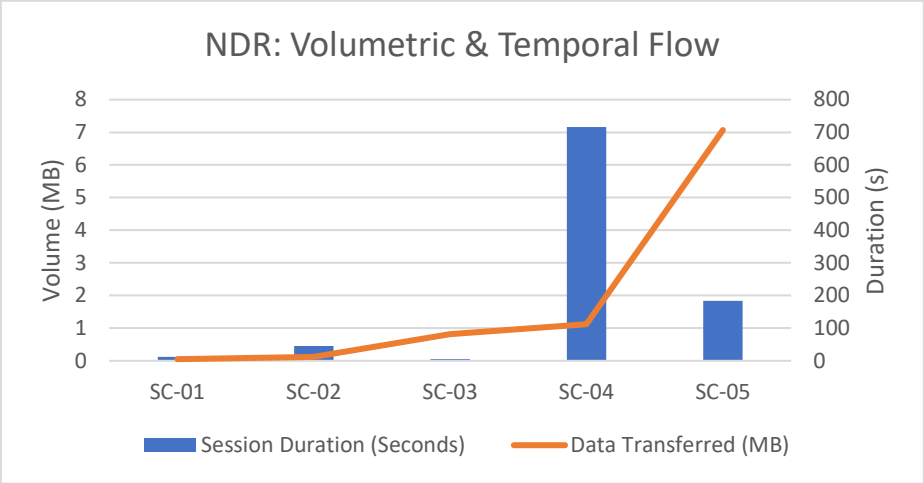

Figure 3 Metadata Profiling Simulation Line Graph

Source: Experimental data processing results from Zeek transaction logs (conn.log), 2026

The isolated quantitative divergence in Chapter III uncovers the presence of a profound Visibility Gap inherent in the exclusive signature defense perimeter. Thru an evaluation of the precise boundaries where the pattern matching model fails, granular extraction of the metadata structure is conducted to isolate deterministic features that maintain visibility when the detection signature remains silent.

During the execution of the SC-04 scenario, the establishment of an encrypted Meterpreter Command and Control (C2) channel used a strict 15-second heartbeat interval. The Suricata NIDS engine generated zero alerts throughout this phase because the underlying binary transmission flow was structurally encapsulated, making it incompatible with any static heuristic patterns in the database of 47,192 active rules. However, parsing the conn.log metadata structure from Zeek yields deterministic compromise indicators. The network metadata engine successfully tracked and recorded the duration of a persistent continuous session lasting exactly 716.33 seconds between the attacker host (192.168.1.10) and the internal asset (192.168.1.20), resulting in a dynamic temporal trace illustrated in Figure 3. This temporal resilience confirms that contextual metadata modeling can effectively isolate malicious network activities in encrypted networks through packet sequences without requiring heavy payload decryption (Papadogiannaki et al., 2022). By evaluating this temporal duration along with periodic transaction updates, the behavioral engine provides explainable telemetry for analysts to classify periodic beaconing anomalies that are structurally overlooked by static machines.

The most catastrophic architectural failure of the signature-based framework is manifested during the SC-05 scenario, which models unauthorized exfiltration from the system configuration directory (/etc). The dangerous movement of these sensitive configuration assets completely evaded NIDS detection. Because the asset extraction was executed thru a well-established and trusted reverse shell socket, the flow of data traffic did not violate the static "Data Theft" signature or the predefined string matching. In stark contrast, Zeek's passive transaction flow logging immediately uncovered a massive volumetric anomaly spike. The

behavioral metadata engine quantified the total data transmission payload at exactly 7,070,966 bits, equivalent to a volumetric burst of 7.07 MB. By monitoring the directed bit flow in real-time via the resp_ip_bytes and orig_ip_bytes attributes, the NDR engine successfully mapped the direction, duration, and precise magnitude of the exfiltration vector without requiring a heavy packet decryption process, as the critical volumetric spike captured in Figure 3. This successful volumetric flow profiling demonstrates that granular metadata serves as a critical socio-technical tool, helping bridge the decision-making gap for domain experts fighting against data exfiltration threats (Chung et al., 2023).

Beyond the detection level, this study evaluates the computational pressure exerted on the SIEM management plane during network operational time to address the critical computational trade-off governing the implementation of high-volume rulesets. A significant architectural anomaly was detected within the Security Onion console, where the main NIDS manager persistently marked the "Sync Failed" alert status. A deep root-cause analysis revealed that forcing the Suricata engine to actively compile, index, and query 47,192 signatures simultaneously has triggered severe memory allocation exhaustion and container socket timeouts within the SaltStack orchestration layer. These operational findings provide critical empirical evidence that aggressively expanding the signature ruleset to achieve maximum coverage actually has the opposite effect on system stability and management plane synchronization. This persistent management failure exposes a critical misconception regarding NIDS scalability, proving that excessive feature loading invariably breaks containerized orchestration planes before packet drops manifest on the data interface (Zhang et al., 2025). This prolonged synchronization failure condition effectively freezes the NIDS sensor's capacity to receive future ruleset updates.

The quantitative and computational findings obtained from this empirical testbed demonstrate that optimizing the corporate SOC requires a strategic shift from the signature-heavy paradigm to a specialized layered sensor architecture. First, organizations must strictly enforce the specialization of sensor roles; NIDS machines should be confined to the outer perimeter to intercept early, noisy, and predictable reconnaissance, while NDR telemetry should serve as the primary defense guide for analyzing post-exploitation anomalies based on session duration, periodicity, and directed bit weight. Second, the efficiency of the ruleset must be systematically enforced by aggressively implementing rule suppression. Informational or low-weight signatures should be disabled on the main listening interface to save hardware processing cycles, considering that the baseline transactions have already been captured with higher fidelity and much lower computational cost by Zeek's lightweight metadata logs. Finally, the company's SOC should implement an enhanced cross-mapping pipeline. Analysts must strategically cross-reference raw NIDS signature findings with NDR transaction flow statistics to dynamically validate the severity level of breaches, thereby transforming isolated noisy alerts into a cohesive, measurable, and forensically valid incident timeline.

CONCLUSION

This research has presented a rigorous quantitative evaluation contrasting high-volume signature-based NIDS against metadata-driven NDR within a production-scale SIEM platform. Using 47,192 active signatures, empirical results reveal a sharp performance bifurcation: NIDS achieved absolute visibility during reconnaissance and exploitation but catastrophically failed

during post-exploitation maneuvers, limiting its Detection Rate to 60%. This profound Visibility Gap exposes the structural blindness of traditional pattern-matching engines against encrypted C2 beaconing and data exfiltration. Conversely, NDR achieved a 100% Detection Rate by capturing deterministic forensic indicators including a 716.33-second encrypted C2 session and a 7.07 MB volumetric exfiltration burst confirming that flow-based metadata abstraction maintains deep security visibility without requiring deep packet inspection or cryptographic decryption. The quantitative and computational compromises yield three architectural recommendations: organizations should enforce strict sensor specialization by limiting NIDS to perimeter defense while positioning NDR for internal behavioral investigation; SOCs must implement aggressive rule suppression to eliminate redundant signature alerts and reduce hardware overhead; and analysts should institutionalize correlative workflows that cross-reference raw NIDS findings with contextual NDR metadata to establish chronologically valid incident timelines. To advance this paradigm, future research must integrate explainable machine learning models such as LSTM networks and Graph Neural Networks into live metadata pipelines to automate multi-stage anomaly detection. Given the rapid proliferation of TLS 1.3 and encrypted DNS (DoH/DoT), which introduce severe visibility barriers as header columns become obfuscated, further studies should explore advanced Encrypted Traffic Analysis (ETA) methodologies focusing on lightweight structural metadata features, including JA4 client-server fingerprinting, bit distribution patterns, and Inter-Arrival Time (IAT) variance. These mathematical advancements are crucial to ensure that layered hybrid detection strategies maintain a resilient defense posture against increasingly sophisticated, stealthy, and cryptographically hidden attacker tactics.

.

REFERENCES

Barradas, D., Novo, C., Portela, B., Romeiro, S., & Santos, N. (2024). Extending C2 traffic detection methodologies: From TLS 1.2 to TLS 1.3-enabled malware. In *Proceedings of the 27th International Symposium on Research in Attacks, Intrusions and Defenses* (pp. 181–196). ACM. https://doi.org/10.1145/3678890.3678921

Berger, S., et al. (2016). Closing the loop: Network and in-host monitoring tandem for comprehensive cloud security visibility. *IBM Journal of Research and Development, 60*(4), 10:1–10:12. https://doi.org/10.1147/JRD.2016.2571580

Bhardwaj, A., Bharany, S., Almogren, A., Rehman, A. U., & Hamam, H. (2024). Proactive threat hunting to detect persistent behaviour-based advanced adversaries. *Egyptian Informatics Journal, 27*, 100510.

Bhatia, G., & Almukhaini, G. (2024). Enhancing cyber attack detection: An analysis of Security Onion for small and midsize enterprise. In *2024 1st International Conference on Innovative Engineering Sciences and Technological Research (ICIESTR)* (pp. 1–6). IEEE. https://doi.org/10.1109/ICIESTR60916.2024.10798271

Chung, M.-H., et al. (2023). Implementing data exfiltration defense in situ: A survey of countermeasures and human involvement. *ACM Computing Surveys, 55*(14s), 1–37. https://doi.org/10.1145/3582077

Hadi, H. J., Ahmad, N., Aziz, K., Cao, Y., & Alshara, M. A. (2024). Cost-effective resilience: A comprehensive survey and tutorial on assessing open-source cybersecurity tools for multi-tiered defense. *IEEE Access, 12*, 194053–194076. https://doi.org/10.1109/ACCESS.2024.3510533

Hajj, S., et al. (2023). Cross-layer federated learning for lightweight IoT intrusion detection

systems. *Sensors, 23*(16), 7038. https://doi.org/10.3390/s23167038

Hore, S., Nguyen, Q. H., Xu, Y., Shah, A., Bastian, N. D., & Le, T. (2023). Empirical evaluation of autoencoder models for anomaly detection in packet-based NIDS. In *2023 IEEE Conference on Dependable and Secure Computing (DSC)* (pp. 1–8). IEEE. https://doi.org/10.1109/DSC61021.2023.10354098

Hu, Q., Yu, S.-Y., & Asghar, M. R. (2020). Analysing performance issues of open-source intrusion detection systems in high-speed networks. *Journal of Information Security and Applications, 51*, 102426. https://doi.org/10.1016/j.jisa.2019.102426

Iglesias, F., & Zseby, T. (2015). Analysis of network traffic features for anomaly detection. *Machine Learning, 101*(1–3), 59–84. https://doi.org/10.1007/s10994-014-5473-9

Koumar, J., & Čejka, T. (2022). Network traffic classification based on periodic behavior detection. In *2022 18th International Conference on Network and Service Management (CNSM)* (pp. 359–363). IEEE. https://doi.org/10.23919/CNSM55787.2022.9964556

Koumar, J., Hynek, K., & Čejka, T. (2023). Network traffic classification based on single flow time series analysis. In *2023 19th International Conference on Network and Service Management (CNSM)* (pp. 1–7). IEEE. https://doi.org/10.23919/CNSM59352.2023.10327876

Majigi, M. U., Idris, I., Abdulhamid, S. M., & Ikuesan, R. A. (2025). Big data transfer service architecture for cloud data centers: Problems, methods, applications, and future trends. *Discover Computing, 28*(1), 163. https://doi.org/10.1007/s10791-025-09682-3

Mohammed, A., et al. (2022). Data security and protection: A mechanism for managing data theft and cybercrime in online platforms of educational institutions. In *2022 International Conference on Machine Learning, Big Data, Cloud and Parallel Computing (COM-IT-CON)* (pp. 758–761). IEEE. https://doi.org/10.1109/COM-IT-CON54601.2022.9850702

Mocanu, F., & Scripcariu, L. (2023). Implementation of a security operation center—An essential cybersecurity solution for organizations. In *2023 27th International Conference on System Theory, Control and Computing (ICSTCC)* (pp. 539–544). IEEE. https://doi.org/10.1109/ICSTCC59206.2023.10308432

Muttaqien, H., Niswar, M., Syarif, S., & Zainuddin, Z. (2025). Efficient identification of malicious traffic in TLS networks using machine learning. In *2025 IEEE International Conference on Artificial Intelligence and Mechatronics Systems (AIMS)* (pp. 1–6). IEEE. https://doi.org/10.1109/AIMS66189.2025.11229622

Papadogiannaki, E., Tsirantonakis, G., & Ioannidis, S. (2022). Network intrusion detection in encrypted traffic. In *2022 IEEE Conference on Dependable and Secure Computing (DSC)* (pp. 1–8). IEEE. https://doi.org/10.1109/DSC54232.2022.9888942

Piet, J., Anderson, B., & McGrew, D. (2018). An in-depth study of open-source command and control frameworks. In *2018 13th International Conference on Malicious and Unwanted Software (MALWARE)* (pp. 1–8). IEEE. https://doi.org/10.1109/MALWARE.2018.8659361

Pranav, H., Suryaa, E., & Venugopalan, M. (2024). Comprehensive C2 analysis and anomaly detection in HTTP traffic: A MongoDB-based approach. In *2024 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI)* (pp. 1–6). IEEE. https://doi.org/10.1109/ACCAI61061.2024.10602244

Raharjo, D. H. K., Nurmala, A., Pambudi, R. D., & Sari, R. F. (2022). Performance evaluation of intrusion detection system performance for traffic anomaly detection based on active IP reputation rules. In *2022 3rd International Conference on Electrical Engineering and Informatics (ICon EEI)* (pp. 75–79). IEEE. https://doi.org/10.1109/IConEEI55709.2022.9972298

Rabbani, M., et al. (2025). Device identification and anomaly detection in IoT environments.

IEEE Internet of Things Journal, 12(10), 13625–13643. https://doi.org/10.1109/JIOT.2024.3522863

Sai, A. N. H. D., Tilak, B. H., Sanjith, N. S., Suhas, P., & Sanjeetha, R. (2022). Detection and mitigation of low and slow DDoS attack in an SDN environment. In *2022 International Conference on Distributed Computing, VLSI, Electrical Circuits and Robotics (DISCOVER)* (pp. 106–111). IEEE. https://doi.org/10.1109/DISCOVER55800.2022.9974724

Salazar, L. A. G. (2026). *Provenance analysis of advanced persistent threats in operational technology environments* [Doctoral dissertation, The University of Texas at El Paso].

Waleed, A., Jamali, A. F., & Masood, A. (2022). Which open-source IDS? Snort, Suricata or Zeek. *Computer Networks, 213*, 109116. https://doi.org/10.1016/j.comnet.2022.109116

Zhang, H., et al. (2025). Explainable and transferable adversarial attack for ML-based network intrusion detectors. *IEEE Transactions on Dependable and Secure Computing, 22*(5), 5090–5107. https://doi.org/10.1109/TDSC.2025.3560486

Zhou, J., Fu, W., Hu, W., Sun, Z., He, T., & Zhang, Z. (2024). Challenges and advances in analyzing TLS 1.3-encrypted traffic: A comprehensive survey. *Electronics, 13*(20), 4000. https://doi.org/10.3390/electronics13204000