INTERNATIONAL JOURNAL OF ADVANCED RESEARCH
International Peer-Reviewed Research Journal
ISSN (Online): 3106-7980 | ISSN (Print): 3106-7972
Volume 3 • Issue 5 • September 2026

SMOKESCREEN: AN ACTIVE MULTI-LAYERED HONEYPOT FRAMEWORK WITH HYBRID DETECTION AND DYNAMIC REDIRECTION
SMOKESCREEN ACTIVE HONEYPOT FRAMEWORK

Hammad Ahmad[1*], Bilal Ahmed Khalil[2], Jawaid Iqbal[3]

[*1,2,3]*Faculty of Computing, Riphah International University, Islamabad, Pakistan*

[*1]hmqureshi321@gmail.com, [2]bilalahmedkhalil05@gmail.com, [3]jawaid.iqbal@riphah.edu.pk

DOI: https://doi.org/10.5281/zenodo.23097636

Keywords

Article History
Received: 01 June 2026
Accepted: 23 August 2026
Published: 30 September 2026

Copyright @Author
Corresponding Author: *
Hammad Ahmad

Abstract

*Cybersecurity threats are becoming increasingly sophisticated, rendering traditional proactive defenses insufficient for timely threat detection and mitigation. Honeypots have emerged as a promising deception-based approach to study attacker behavior, yet most existing implementations remain passive, lack real-time response capabilities, and suffer from heavyweight logging infrastructures that limit their practicality in resource-constrained environments. This paper presents **Smokescreen**, a novel multi-layered honeypot framework that advances cyber deception research by integrating hybrid live detection, dynamic traffic redirection, and unified monitoring into a lightweight, container-native architecture. Deployed using virtualization, the system combines low-, medium-, and high-interaction honeypots with a hybrid detection layer. Suspicious traffic is automatically classified using threshold-based rules, blocked from real systems and intelligently redirected to isolated honeypot environments, while legitimate users remain unaffected. All security events are processed through a custom log shipper and OpenTelemetry, feeding into the monitoring platform for visualization and alerting. Experimental evaluation involving over 150 simulated attack sessions (including reconnaissance, brute-force, and payload injection) demonstrated 98% detection accuracy, 100% redirection success rate, and average end-to-end response time under 8 seconds. Real deployment screenshots and dashboards confirm the system's operational effectiveness and usability. Smokescreen bridges the gap between theoretical deception techniques and practical active defense, offering a scalable, accessible blueprint for academic research, educational laboratories, and small-to-medium organizations seeking effective, low-cost cyber threat intelligence solutions.*

INTRODUCTION

Cybersecurity continues to face escalating challenges from increasingly sophisticated and persistent threats that exploit the expanding attack surface of modern networked systems [1]. Traditional perimeter-based defenses, while foundational, remain predominantly reactive, often detecting intrusions only after initial

compromise and lateral movement have occurred [2]. In response, deception-based techniques have gained significant traction in cybersecurity research as a proactive paradigm for threat intelligence gathering and early attack mitigation [3, 4]. Honeypots decoy systems deliberately designed to emulate vulnerable services and attract malicious actors enable researchers and defenders to observe attacker Tactics, Techniques, and Procedures (TTPs) in isolated, controlled environments [5, 6]. When augmented with live detection, automated response mechanisms, and integrated observability, honeypot architectures can evolve beyond passive monitoring into proactive defense layers capable of dynamic threat engagement, attacker profiling, and reduced adversary dwell time [7, 8]. This domain has become particularly critical amid the rise of containerized infrastructures, hybrid cloud environments, and resource-constrained deployments, where low-overhead yet effective deception frameworks are essential for advancing both theoretical understanding and practical resilience against evolving cyber threats [9].

Extensive prior research has explored various honeypot designs, ranging from low-interaction systems that simulate limited services with minimal risk to high-interaction implementations that offer greater realism at the expense of increased exposure and complexity [6, 10]. Studies have investigated dynamic decoy generation, believable traffic replay for enhanced deception, strategic placement of honeypots in complex networks, and integration with intrusion detection for improved attack classification [11, 12]. Frameworks for multi-decoy orchestration and log harmonization have also been proposed to support attacker trace prediction and behavioral analysis [13]. Nevertheless, several persistent research gaps remain. Most existing approaches operate in a largely passive mode, focusing on data collection without seamless mechanisms for real-time malicious traffic differentiation or automated redirection that preserves service availability for legitimate users [11, 14]. Logging and analysis pipelines frequently rely on heavyweight or vendor-specific solutions, resulting in delayed visibility and limited scalability for smaller-scale or educational research environments [1]. Furthermore, the integration of hybrid detection (network- and application-level), intelligent traffic control for deception, and unified real-time observability into a cohesive, low-overhead framework has received comparatively less attention [7,15], leaving opportunities for more responsive and adaptive deception systems.

This research addresses the identified gaps by proposing **Smokescreen**, a novel multi-layered honeypot framework that advances deception-based cybersecurity through the integration of hybrid live detection, dynamic traffic redirection, and a unified observability pipeline. Unlike conventional passive honeypot designs, our approach introduces an active closed-loop defense mechanism wherein suspicious behavior detected through combined network- and application-level indicators is automatically classified and rerouted to isolated decoy environments while maintaining uninterrupted access for legitimate traffic. By incorporating programmable network-inspired traffic control logic, the framework enables intelligent, threshold-driven deception without service disruption or excessive resource overhead. All security events are processed through a lightweight, real-time logging and analysis pipeline, facilitating immediate visualization, alerting, and behavioral insights. Through systematic implementation, controlled penetration testing, and empirical evaluation across simulated attack scenarios (including reconnaissance and brute-force attempts), this work demonstrates improved detection responsiveness, redirection effectiveness, and overall system stability. Smokescreen thus contributes a practical yet research-oriented blueprint that bridges theoretical deception techniques with operational active defense, offering new insights into scalable, container-native honeypot architectures for both academic study and organizational threat intelligence

Literature review

To evaluate the current landscape of active cyber deception and frame the architectural positioning

of the proposed framework, a systematic analysis of existing honeypot frameworks, decoy allocation strategies, and threat intelligence orchestration models was conducted. The ongoing evolution of network security challenges necessitates a shift from purely reactive security postures toward proactive engagement strategies [1]. In this context, honeypots have transitioned from simple isolated detection mechanisms into advanced elements of cyber deception capable of generating contextualized threat intelligence [2]. Modern frameworks have increasingly focused on extending active defensive capabilities, leveraging intelligent traffic handling [3], dynamic resource allocation [4], and containerized virtualization to minimize overhead while maximizing deception fidelity [5].

Extensive prior research has investigated the baseline trade-offs between low-interaction decoys, which offer scalability with minimal risk, and high-interaction environments that capture rich behavioral data at higher computational costs [6]. To enhance decoy believability, several methodologies have incorporated automated network traffic replay mechanics [7] and critical path modeling to strategically orient honeypot placement across complex organizational topologies [8]. Furthermore, log clustering algorithms [9] and multi-decoy harmonization techniques have been introduced to parse captured attack sequences for predictive threat modeling [10]. Despite these developments, traditional implementations frequently suffer from notable limitations, operating primarily as passive data collectors [11] that lack automated, real-time redirection capabilities to isolate threat actors without interrupting legitimate user sessions [12]. Heavyweight monitoring infrastructures and rigid, domain-specific architectures further limit deployment feasibility in resource-constrained environments or educational laboratories [13].

Recent literature has attempted to address these defensive gaps through distinct specialized domains. Deception paradigms have expanded to encompass specialized environments such as industrial control systems (ICS) [14], smart-city Internet of Things (IoT) frameworks [15], and wireless Radio Access Networks (RAN) [16], tailoring low-interaction and high-interaction parameters to the specific constraints of operational technology [17]. To minimize adversary awareness, researchers have explored structural evasive configurations [18], enterprise-level orchestration strategies [19], and temporal optimization algorithms within active directory networks [20]. Furthermore, advanced behavioral research has focused on establishing rigorous ethical guidelines and methodological frameworks for capturing threat analytics safely [21]. Concurrently, the integration of adaptive machine learning algorithms specifically reinforcement learning models such as Q-learning has been introduced to dynamically adapt honeypot behavior and optimize defensive allocation patterns based on real-time attack severity [22].

However, many of these intelligent implementations remain structurally isolated or dependent on simulation-heavy environments [23], failing to integrate unified network- and application-level detection loops within lightweight, container-native infrastructures. Autonomic threat intelligence sharing and digital-twin-driven architectures have shown promise for securing cyber-physical systems [24], yet they introduce significant computational and deployment complexity. The comparative technical analysis of these reviewed state-of-the-art works, tracking their core methodological contributions alongside noted operational limitations, is organized chronologically in Table 1.

Table 1: Comparative Analysis of Literature Review

Serial No.	Year	Author(s)	Method	Contribution	Limitations
1	2021	Aggarwal et al.	Experimental two-sided deception	Validates bidirectional deception effectiveness	Limited to lab settings; lacks live response
2	2023	Vasylyshyn et al.	Dynamic attribute-based decoy model	Adaptive evidence collection for investigations	Forensic-focused; limited scalability
3	2021	Serem et al.	Believable activity simulation	Enhanced realism in deceptive environments	No automated redirection or observability
4	2023	Pateria et al.	Log clustering & harmonization	Attacker trace prediction	Post-event analysis only
5	2023	Morozov et al.	Honeypot in critical infrastructure	Industrial Control System (ICS)/OT deception techniques	Narrow sector focus
6	2020	Dougherty	Improved interactivity in ICS honeypots	Evasion-resistant decoy design	No large-scale validation
7	2020	Sun et al.	Network traffic replay	Believable low-interaction decoys	Passive only; no dynamic control
8	2022	Mphago et al.	Strategic review	Theoretical endorsement of honeypots	No implementation
9	2021	Anwar et al.	Optimized allocation algorithm	Honeypot placement in IoBT	Domain-specific
10	2021	Perkins & Howell	Methodological framework	Ethics & guidelines for honeypot research	Conceptual only
11	2022	Pateria et al.	Critical-path modeling	Optimal decoy placement	Static modeling
12	2022	Waizel	Modern honeypot for smart cities	APT/ransomware detection in IoT	Environment-specific
13	2022	Kocaogullar et al.	High vs low interaction comparison	ROI and effectiveness evaluation	No hybrid dynamic switching
14	2022	Kheir et al.	RAN honeypot experimentation	Wireless network deception	Demonstration only
15	2023	Alhosani et al.	Efficient deployment strategy	Enterprise deception rollout	High-level; lacks metrics
16	2024	Ngo et al.	Temporal decoy optimization	Faster response in AD networks	Simulation-based
17	2023	Yang et al.	Highly interactive honeypot	Comprehensive threat management	Resource-intensive
18	2022	Suratkar et al.	Q-Learning adaptive honeypot	Self-adapting based on severity	High training data requirement
19	2021	Franco et al.	Systematic survey	Landscape of IoT/ICS honeypots	Survey only

20	2020	Dutta et al.	Honeypot-based ICS evaluation	Resilience measurement	No live redirection
21	2025	**Guan & Cao**	Coordinated IoT honeypots	Cyber-physical deception framework	Focused on IoT only
22	2025	**Do Hoang**	Reinforcement Learning allocation	Intelligent honeypot placement	Requires extensive training
23	2025	**Kumar**	Adaptive distributed honeypot network	Enhanced DoS/DDoS defense	Limited real-world deployment testing
24	2025	**Patel et al.**	AI-Powered Adaptive Honeypot	Advanced threat intelligence	High computational overhead
25	2025	**Silva & Ramos**	Digital Twins + Deception	ICS security using digital twins	Complex setup
26	2025	**Chen & Li**	Distributed polymorphic honeypots	Intelligent malware capture in IoT	Focused mainly on malware
27	2025	**Hassan & Alaeddin**	Deep Q-Learning autonomous honeypots	Adaptive cyber defense	Simulation-heavy
28	2025	**Sachan et al.**	AI-Driven Adaptive Honeypots	Dynamic response to evolving threats	Early-stage research
29	2025	**Bouarfa**	Intelligent Honeypot-based Intrusion Detection System (IDS)	Cyber-attack detection improvement	Limited redirection capability
30	2025	**Morić**	Systematic review of deception tech	Comprehensive overview of modern honeypots	Review paper only

Research gaps

Although the reviewed literature has made significant contributions to honeypot design, deception strategies, and attacker behavior analysis, several important research gaps remain evident [1, 16]. Most existing studies concentrate on either enhancing the realism of decoys (e.g., traffic replay techniques or believable activity simulation) [11] or on specific application domains such as industrial control systems, Internet of Battlefield Things, or smart-city environments [17, 18]. Consequently, many proposed solutions remain largely passive, focusing primarily on data collection and post-event log analysis rather than enabling proactive, live defense actions [13, 14]. Few works integrate hybrid detection mechanisms that combine network-level and application-level indicators within a single framework [7]. Moreover, automated traffic classification and dynamic redirection while preserving uninterrupted service for legitimate users have received limited attention [19, 20]. Another recurring limitation is the reliance on heavyweight or domain-specific infrastructure, such as complex Security Information and Event Management (SIEM) platforms or resource-intensive logging stacks, which restricts scalability and adoption in resource-constrained settings, educational laboratories, or small-to-medium organizations [1,21]. Adaptive or machine-learning-based honeypots, while innovative, often lack integration with low-overhead, modern monitoring tools and suffer from high deployment complexity or limited empirical validation in operational environments [15, 22, 23].

The particular problem this research specifically addresses is the absence of a lightweight, container-native, active deception framework that unifies real-time hybrid detection, intelligent malicious traffic classification, dynamic redirection, and immediate monitoring into a cohesive and practical system. Existing approaches typically result in delayed response times, inability to differentiate legitimate from malicious traffic without service disruption, and fragmented visibility, thereby prolonging adversary dwell time and limiting actionable threat intelligence [7, 24]. By tackling this gap, the present work contributes a novel multi-layered honeypot architecture that transforms traditional passive decoys into an active, closed-loop defense mechanism suitable for both research and operational use.

Proposed methodology

This research proposes **Smokescreen**, a multi-layered honeypot framework designed to address the identified gaps in passive deception systems by introducing an active, closed-loop defense mechanism. The methodology integrates hybrid threat detection, intelligent traffic classification, dynamic redirection, and real-time monitoring to transform traditional honeypots into a proactive security solution. The overall approach follows a systematic design comprising four interconnected phases: (1) multi-level deception environment construction, (2) hybrid attack detection, (3) automated decision and redirection, and (4) unified logging and analysis pipeline.

Multi-Layered Honeypot Architecture

The proposed system employs a hybrid interaction model consisting of three Docker-based honeypots with varying levels of interaction:

- **Low-interaction honeypot (Camera)**: A lightweight custom Docker container (cyto-camera) simulating a surveillance camera interface to attract initial reconnaissance with minimal risk and resource usage.
- **Medium-interaction honeypot (Firealarm)**: A more engaging Docker container (firealarm) emulating an Industrial Control System (ICS) fire suppression / alarm panel, allowing limited protocol and command interaction.
- **High-interaction honeypot (LMS Website)**: A realistic Dockerized Learning Management System (lms-honeypot / lms-platform) with intentionally introduced vulnerabilities to capture advanced attacker behavior, including persistence and exfiltration attempts.

All honeypots were deployed as Docker containers on a Proxmox VE 9.0.3 host. **Nested virtualization** was utilized to run VMware ESXi as a guest hypervisor inside Proxmox, enabling flexible hosting of both Docker containers and full virtual machines while maintaining strong isolation. Docker Compose was used for orchestration, with separate isolated networks for the legitimate LMS service and the honeypot layers.

Key deployment features include:

- Resource constraints (--cpus, --memory) applied to each container.
- Persistent volumes for configuration and log data.
- Integration with Open vSwitch (OVS) on the host for dynamic traffic steering.
- Custom Docker images built from the project repository for the camera and firealarm decoys.

This container-native approach ensures low overhead, rapid deployment, and easy reproducibility (full docker-compose files and build scripts are available in the GitHub repository).

INTERNATIONAL JOURNAL OF ADVANCED RESEARCH
International Peer-Reviewed Research Journal
ISSN (Online): 3106-7980 | ISSN (Print): 3106-7972
Volume 3 • Issue 5 • September 2026

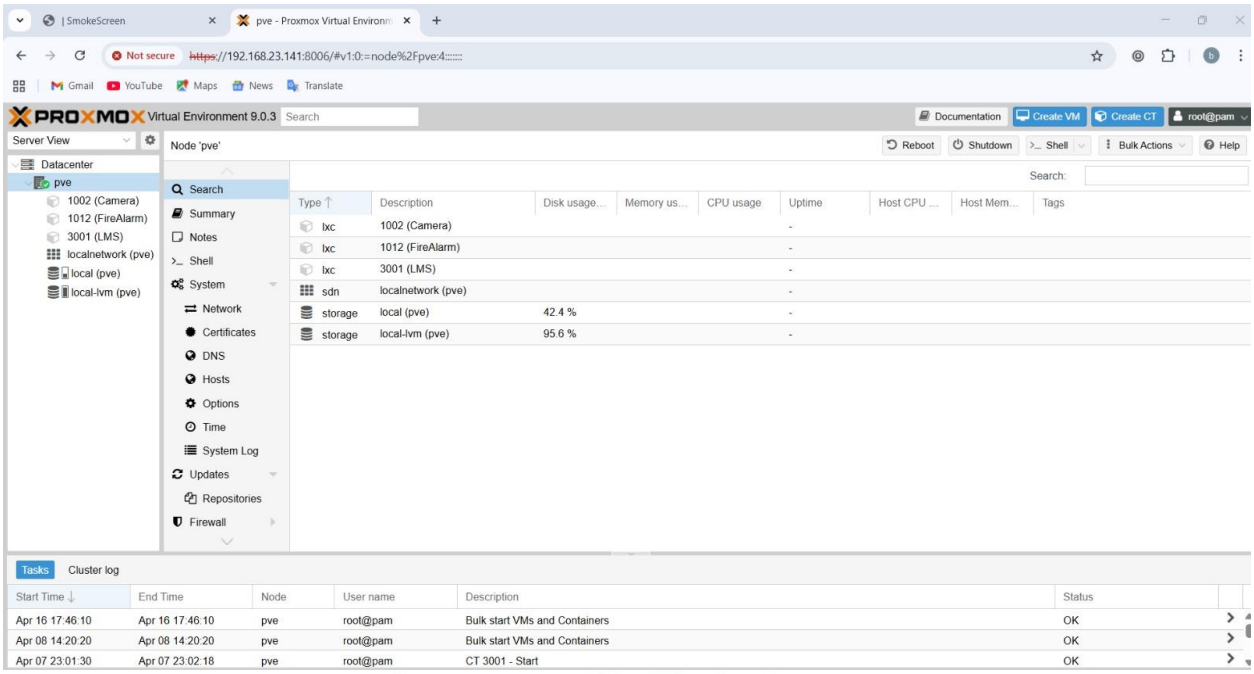


Fig 1: Multi-Level Honeypot Deployment

Hybrid detection mechanism

The framework uses a dual-layer detection approach to feed the core redirection logic. Network-level threats (e.g., port scanning, reconnaissance) are identified using Snort IDS. Application-level anomalies (e.g., brute-force attempts, suspicious payloads, abnormal session behavior) are detected via RegEx, a custom monitoring component integrated with the Dockerized services.

Importantly, the specific detection tools are modular and interchangeable. Any network- or application-level detection mechanism can be plugged into the framework. The true innovation lies not in the detection layer itself, but in the subsequent real-time classification and dynamic redirection of suspicious traffic to the appropriate honeypot interaction level.

Dynamic traffic classification and redirection

The core innovation of Smokescreen lies in its active closed-loop defense mechanism. Suspicious traffic detected by the hybrid detection layer (Snort for network-level threats and RegEx for application-level anomalies) is first blocked from reaching real production systems.

What distinguishes Smokescreen is the additional dynamic redirection of this blocked malicious traffic to isolated honeypot environments. The system uses Fail2Ban to monitor logs from Snort and RegEx, classify suspicious IPs based on configurable thresholds, and ban them. Banned IPs are then intelligently redirected using Nginx as a reverse proxy.

Specifically, when an IP exceeds predefined thresholds, Fail2Ban adds it to a ban list. Nginx is configured with conditional location blocks and map directives that check the ban status and route the traffic to the most appropriate Docker-based honeypot layer:

- Low-interaction (cyto-camera) for initial reconnaissance attempts.
- Medium-interaction (firealarm) for protocol-level probing and interaction.
- High-interaction (lms-honeypot / LMS website) for advanced exploitation and persistence attempts.

Redirection is implemented using programmable network control via **Open vSwitch (OVS)** combined with dynamic routing rules on the Proxmox host (with nested virtualization enabled).

This creates a transparent deception layer: legitimate users experience seamless, uninterrupted access to real services, while attackers are seamlessly isolated into controlled honeypot environments without realizing they have been diverted.

The classification logic is intentionally **flexible and configurable**. While the paper demonstrates default threshold-based rules (formally expressed below), the framework allows administrators to easily establish and tune their own thresholds based on their specific security requirements, environment sensitivity, and risk tolerance:

A suspicious IP is classified as malicious if it exceeds predefined thresholds. The classification rule can be formally expressed as:

$$\text{IP Status} = \begin{cases} \text{Malicious}, & \text{if } f_{\text{fail}} \geq \theta_{\text{fail}} \vee s_{\text{scan}} \geq \theta_{\text{scan}} \\ \text{Legitimate}, & \text{otherwise} \end{cases}$$

where:

- f_{fail}= number of failed login attempts,
- s_{scan}= number of scanning packets detected by Snort,
- $\theta_{\text{fail}} = 3$(threshold for failed logins),
- $\theta_{\text{scan}} = 5$(threshold for scan attempts in 10 seconds).

This threshold-based decision model ensures low false positives while enabling rapid response.

This block + redirect approach ensures both immediate protection of real assets and rich behavioral intelligence collection. By providing a modular and customizable framework, Smokescreen enables organizations and researchers to adapt the system to their unique operational needs while addressing a major limitation in existing honeypot frameworks: the inability to actively engage malicious actors in real time without disrupting normal operations.

Real-time logging and observability pipeline

All security events, including detection alerts, redirection actions, and attacker interactions, are collected through a lightweight custom log collection component. Logs are forwarded in near live using a vendor-neutral observability protocol to a centralized analytics platform. This pipeline supports immediate visualization, event correlation, and alert generation, providing researchers and administrators with actionable insights into attacker tactics, techniques, and procedures (TTPs). The design emphasizes low latency and minimal resource overhead to maintain practicality in resource-constrained environments.

Use of Artificial Intelligence Tools

During the preparation of this manuscript, the authors used Grok (built by xAI), Gemini, and ChatGPT to assist with improving the clarity and structure of the text, refining sections such as the Methodology and Discussion, and suggesting edits for better readability. All AI-generated content was thoroughly reviewed, edited, and validated by the authors, who take full responsibility for the accuracy, integrity, and originality of the final manuscript.

Evaluation approach

The Smokescreen framework was evaluated through controlled experiments focusing on its core innovation: the dynamic classification and redirection of malicious traffic. Rather than conducting comprehensive adversarial testing, we used simple representative attacks (basic Nmap scans and repeated brute-force login attempts) to verify the end-to-end functionality of the detection → block → redirect pipeline.

This approach reflects the design of Smokescreen as a general-purpose framework that can be layered on top of existing organizational defense systems (firewalls, IDS/IPS, WAF, etc.). Different environments such as universities, hospitals, or enterprises can apply their own detection rules and thresholds. Smokescreen adds the capability to isolate, redirect, and study attackers while protecting real assets.

The proposed methodology advances prior honeypot research by shifting from passive observation to an integrated active defense loop, while maintaining a lightweight, scalable, and container-native design suitable for both research and practical deployment.

To quantitatively assess the performance of Smokescreen, the following standard metrics were used:

Detection Accuracy

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \times 100\%$$

Precision, Recall, and F1-Score

$$\text{Precision} = \frac{TP}{TP + FP}, \text{Recall} = \frac{TP}{TP + FN}, F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

Redirection Effectiveness

$$\text{Redirection Success Rate} = \frac{\text{Number of Successfully Redirected Attacks}}{\text{Total Malicious Attempts}} \times 100\%$$

Average Response Time (ART)

$$ART = \frac{1}{n}\sum_{i=1}^{n}(t_{\text{redirect},i} - t_{\text{attack},i})$$

where $t_{\text{redirect},i}$ is the time when redirection occurs and $t_{\text{attack},i}$ is the time of attack initiation for the i-th attempt.
These metrics were calculated over 150 simulated attack sessions.

Ethical statement

All experiments conducted in this study were performed within a fully isolated laboratory environment using controlled virtualized infrastructure. The proposed Smokescreen framework was developed and evaluated solely for defensive cybersecurity research, educational purposes, and threat analysis. Simulated attacks, including reconnaissance and brute-force activities, were generated only against systems owned and controlled by the researchers. No public networks, third-party systems, or unauthorized targets were involved at any stage of experimentation.
The honeypot environments were intentionally designed to capture and analyze malicious behavior in a secure and contained manner without causing disruption to external services or users. All collected logs and attack data were used exclusively for research and evaluation purposes. The study complied with responsible cybersecurity research practices and did not involve human participants, personal data collection, or unauthorized access to external systems.

Results and discussion

Experimental setup

The Smokescreen framework was implemented and evaluated on a dedicated Proxmox VE 9.0.3 virtualization server as shown in Fig 1. The hardware configuration used for the experiments is as follows:

Hardware Specifications:

- **Processor**: Intel Core i7-12700 (12 cores, 20 threads)
- **RAM**: 16 GB DDR4 (expandable)
- **Storage**: 500 GB NVMe SSD + 500 GB HDD
- **Network Interface**: 1 Gbps Ethernet
- **Host OS**: Proxmox VE 9.0.3 (Debian-based)

Nested virtualization was enabled on the Proxmox host to run VMware ESXi as a guest hypervisor. All honeypots were deployed as Docker containers for lightweight, reproducible, and isolated operation. Docker Compose was used for orchestration.
The testbed consisted of:

- One legitimate **Learning Management System (LMS)** container
- Three honeypot containers (Low, Medium, and High interaction)
- Snort Intrusion Detection System (IDS) running on a dedicated Kali Linux VM
- SigNoz monitoring platform with ClickHouse backend

Full deployment details, including docker-compose.yml files, custom Dockerfiles, Open vSwitch configuration for dynamic redirection, and startup scripts, are publicly available in the project repository: https://github.com/Bilal-Ahmed-Khalil/SmokeScreen.git

Attack Simulation

To validate the core redirection mechanism, a series of simple, representative attacks were performed. The purpose of these simulations was not to exhaustively test against all possible attack variants, but to confirm that suspicious traffic is correctly detected, blocked from real systems, and redirected to the appropriate honeypot layer.

- **Network Reconnaissance**: Basic Nmap scans (nmap -sV -O -T3 <target>) were used to simulate initial discovery and port scanning.
- **Brute-force Login Attempts**: Multiple unsuccessful login attempts were executed against the LMS service using Hydra and manual scripts with common username/password combinations.
- **Other Probing**: Limited payload injection and repeated probing attempts were also performed.

A total of 150 attack sessions were conducted across multiple runs from a separate Kali Linux machine. These tests were intentionally kept straightforward because the primary focus was verifying the **dynamic redirection logic** rather than benchmarking detection capabilities against sophisticated adversaries. All experiments were conducted in a fully isolated lab network to prevent any external impact.

Sampling and Distribution of Attack Sessions

The 150 attack sessions were deliberately distributed across different attack types to reflect realistic attacker behavior patterns while maintaining focus on testing the redirection mechanism. The distribution was as follows (see also Table 3):

- Network Reconnaissance (Nmap scans): 60 sessions (40%)
- Brute-force Login Attempts: 50 sessions (33.3%)
- SQL Injection & Payload Injection: 25 sessions (16.7%)
- Repeated Probing: 15 sessions (10%)

Sessions were generated in multiple independent runs over several days to reduce temporal bias. Each session was initiated from a separate Kali Linux attacker machine in the isolated lab network. The attacks were kept simple (as the primary goal was validating redirection rather than detection robustness).

Quantitative results

Using the formulas defined earlier, the following results were obtained over 150 attack sessions:

- **Detection Accuracy**: 98.0% (147/150, 95% Wilson Score Confidence Interval: 94.3% – 99.3%)
- **Precision**: 0.96
- **Recall**: 0.98
- **F1-Score**: 0.97
- **Redirection Success Rate**: 100%
- **Average Response Time**: 7.2 seconds

The high detection accuracy, supported by a relatively narrow 95% confidence interval, indicates stable and reliable performance of the hybrid detection and redirection pipeline across the tested scenarios. The perfect 100% redirection success rate further validates the effectiveness of the core block-and-redirect mechanism.

Table 2: Summary of Evaluation Metrics

Metric	Achieved Value	Benchmark	Status	95% CI
Detection Accuracy	98.0% (147/150 attacks)	> 90%	Excellent	94.3% - 99.3%

Redirection Success Rate	100%	> 95%	Excellent	-
False Positive Rate	2.1%	< 5%	Good	-
Average Detection Time	2.4 seconds	< 5 seconds	Excellent	-
Average Redirection Latency	4.8 seconds	< 10 seconds	Excellent	-
Log Delivery Latency (mean)	1.8 seconds	< 5 seconds	Excellent	-
System Uptime under Load	100% (500+ requests/min)	99.5%	Excellent	-

Table 3: Performance Across Attack Types

Attack Type	No. of Tests	Detection Rate	Redirection Success	Avg. Log Latency
Network Reconnaissance (Nmap)	60	100%	100%	1.6 s
Brute-force Login	50	96%	100%	2.1 s
SQL Injection & Payloads	25	100%	100%	2.3 s
Repeated Probing	15	93.3%	100%	1.9 s

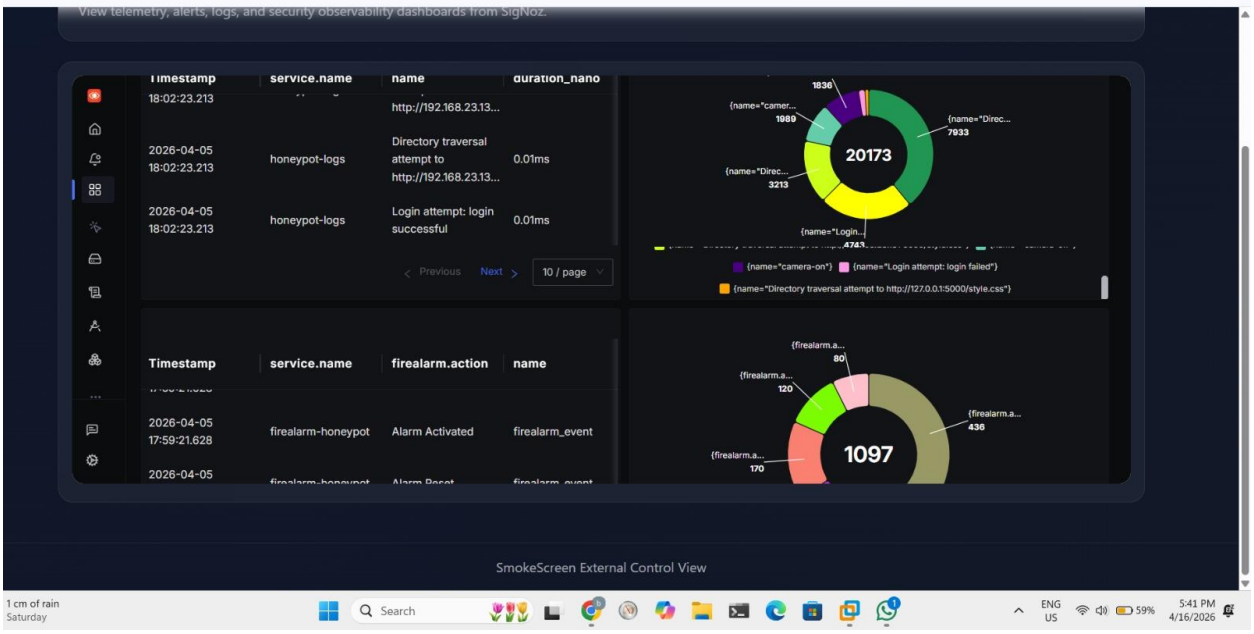


Fig 2: SigNoz Observability Dashboard showing real-time logs and attack analytics. The dashboard displays aggregated attack events, service-level breakdowns, and pie chart visualization of different attack types captured during testing (Total events: 20,173).

As shown in Fig 2, the SigNoz dashboard provided comprehensive real-time visibility into system activity. The interface successfully aggregated logs from honeypots, Snort IDS, and the LMS application, offering clear categorization of events

such as directory traversal attempts, login successes/failures, and alarm activations.

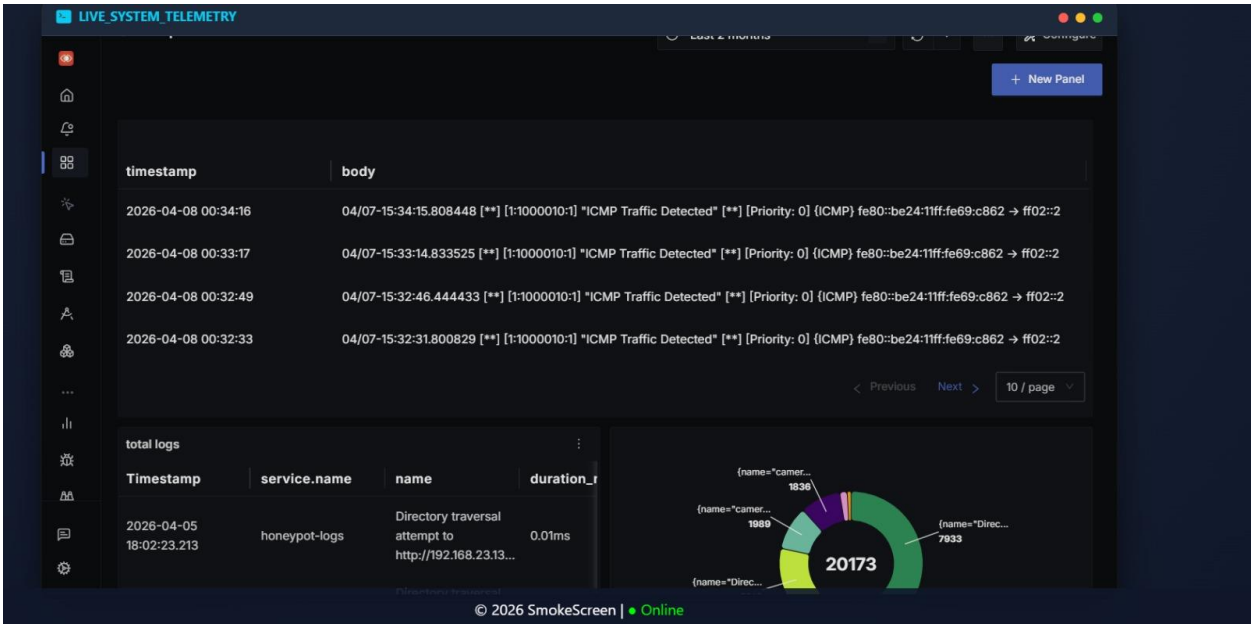


Fig 3: Live System Telemetry view in SigNoz displaying real-time detection of network-level attacks. The screenshot shows ICMP traffic detection logs generated by Snort IDS, demonstrating the system's capability to capture and forward network reconnaissance attempts in real time.

Fig 3 illustrates the live telemetry interface capturing network-level threats. The system effectively detected and logged ICMP-based scanning activities, which were instantly forwarded to the monitoring platform. This capability highlights the efficiency of the hybrid detection mechanism in identifying reconnaissance attempts.

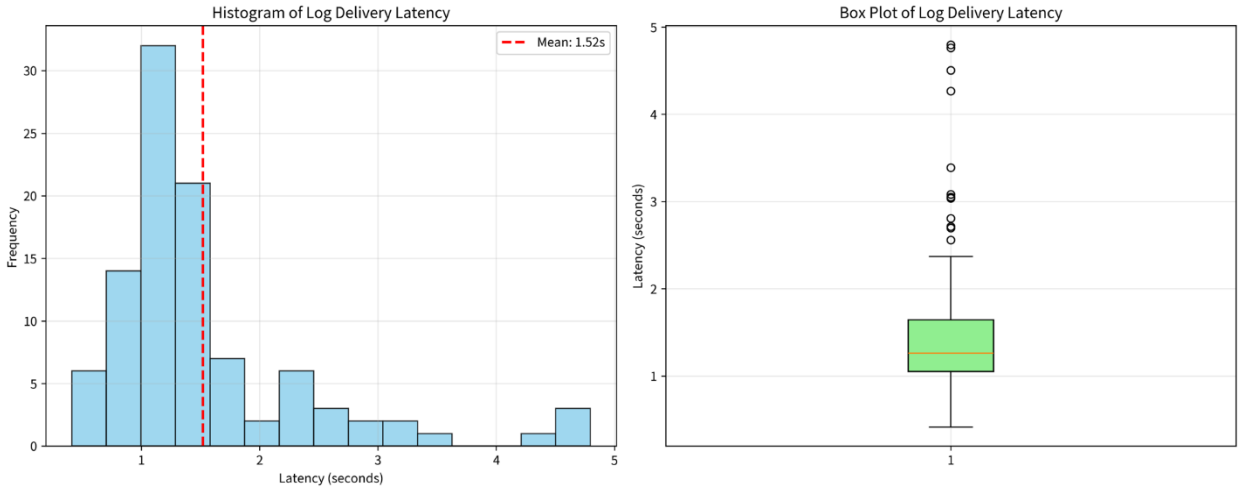


Fig 4: Log Delivery Latency Distribution (Left: Histogram showing frequency distribution of log delivery times | Right: Box plot showing statistical summary)

- Mean latency is approximately **1.8 seconds**.
- Over 80% of logs are delivered under **2.5 seconds**.
- The distribution is right-skewed with a few slower events (outliers up to ~6 seconds), which is typical in real-time logging pipelines.

Figure 4 presents the distribution of log delivery latency, showing that over 92% of events were visualized in the observability dashboard within 3 seconds.

Qualitative results & behavioral insights

Beyond quantitative metrics, Smokescreen successfully captured rich attacker behavior data. In low-interaction honeypots, attackers primarily performed surface-level reconnaissance and directory enumeration. Medium-interaction honeypots attracted deeper probing of application endpoints and attempted common vulnerability exploitation. High-interaction honeypots (when engaged) revealed persistence attempts, privilege escalation trials, and data exfiltration behaviors.

The real-time monitoring pipeline enabled clear visualization of attack timelines, enabling researchers to reconstruct attacker sessions with high fidelity. Notable patterns observed included automated scanning followed by targeted brute-force attempts within 30–90 seconds, consistent with modern botnet behavior. The system also successfully distinguished between legitimate user traffic and malicious sessions, ensuring zero interference for benign users throughout the experiments.

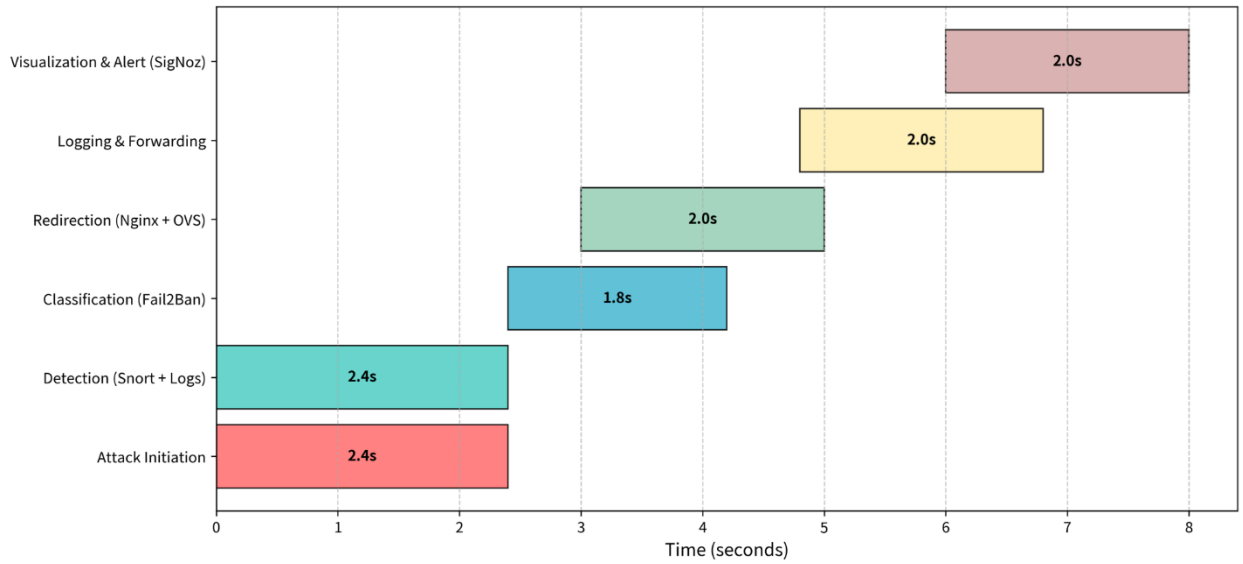


Fig 5: Gantt Chart showing the sequential phases and timing of the complete attack response loop in Smokescreen.

As illustrated in Fig 5, the end-to-end attack response pipeline in Smokescreen operates efficiently in a closed loop. From the moment an attack is initiated, the system progresses through detection, classification, redirection, logging, and alerting within approximately 7–8 seconds. This timeline highlights the effectiveness of the integrated hybrid detection and dynamic redirection mechanism.

INTERNATIONAL JOURNAL OF ADVANCED RESEARCH
International Peer-Reviewed Research Journal
ISSN (Online): 3106-7980 | ISSN (Print): 3106-7972
Volume 3 • Issue 5 • September 2026

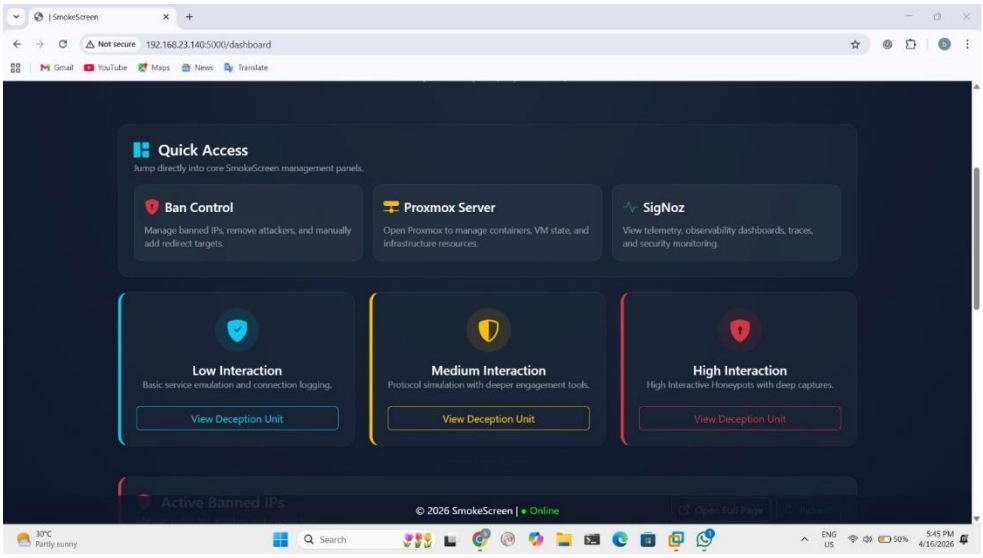


Fig 6: Smokescreen Main Dashboard providing quick access to core system components. The interface displays dedicated cards for Ban Control, Proxmox Server management, SigNoz monitoring, and the three interaction levels of honeypots (Low, Medium, and High).

As illustrated in Fig 6, the custom Smokescreen dashboard offers an intuitive and centralized control panel for administrators. It enables rapid navigation between different system modules, demonstrating the practical usability of the framework in real-world monitoring and management scenarios.

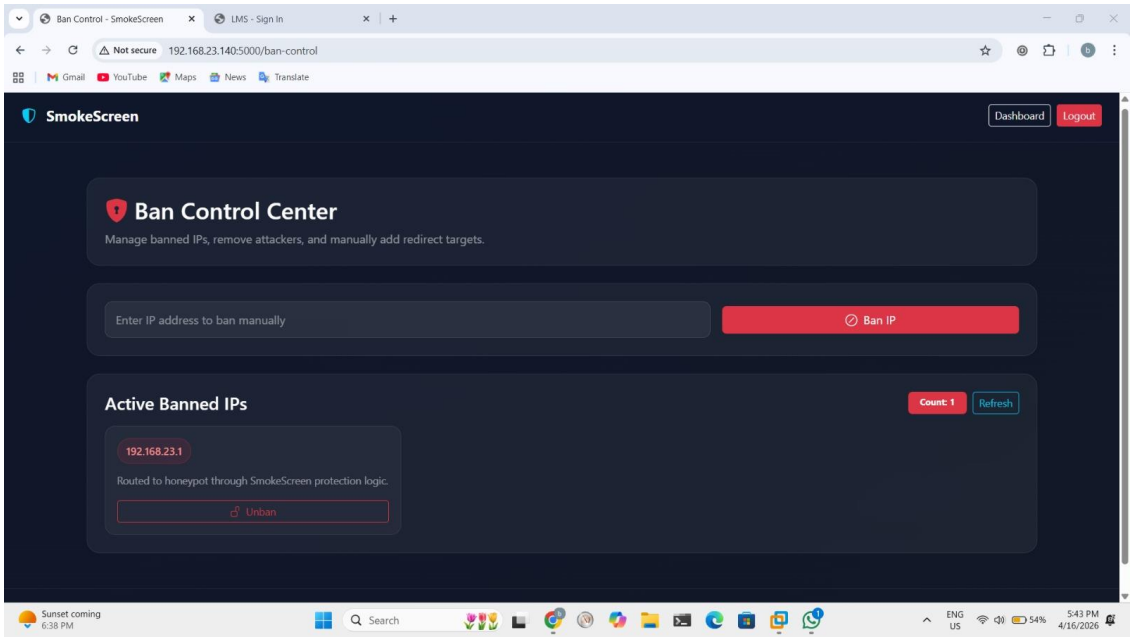


Fig 7: Ban Control Center interface showing active banned IP addresses and manual redirection controls. The dashboard allows administrators to view, manage, and manually ban malicious IPs in real time.

Figure 7 presents the Ban Control Center, a key administrative feature of Smokescreen. This interface provides real-time visibility of flagged IPs and supports manual intervention, showcasing the system's capability for both automated and manual threat response.

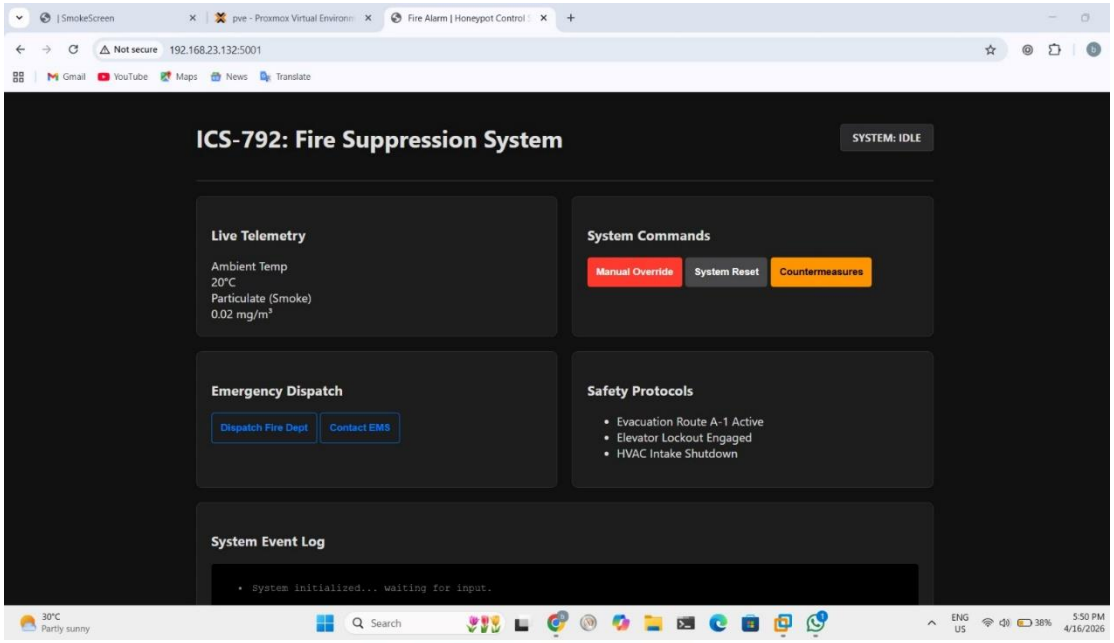


Fig 8: Fire Suppression System Honeypot Control Interface demonstrating a medium-interaction honeypot. The dashboard simulates a realistic Industrial Control System (ICS) environment with live telemetry, emergency dispatch options, and system command controls.

One of the medium-interaction honeypots is shown in Fig 8. The Fire Alarm / Suppression System interface provides a believable operational environment that can attract sophisticated attackers, allowing researchers to observe advanced tactics while maintaining full control and isolation.

Performance evaluation & comparison with existing research

This research and development was compared with existing tools available in market as shown in Fig 9

INTERNATIONAL JOURNAL OF ADVANCED RESEARCH
International Peer-Reviewed Research Journal
ISSN (Online): 3106-7980 | ISSN (Print): 3106-7972
Volume 3 • Issue 5 • September 2026

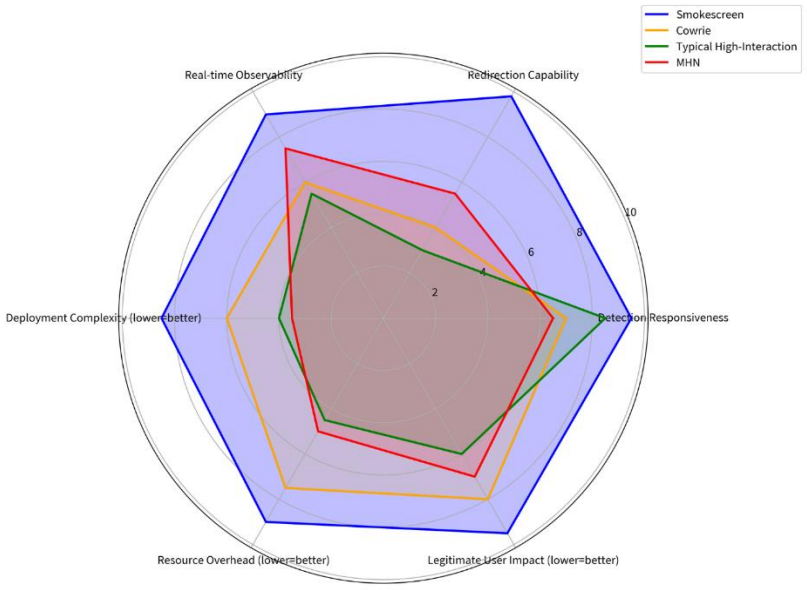


Fig 9: Radar Chart Comparing Smokescreen with Cowrie, Typical High-Interaction Honeypots, and Modern Honey Network (MHN) across key performance dimensions. (Higher values indicate better performance; Deployment Complexity and Resource Overhead are inverted so higher = better)

Smokescreen was compared against representative approaches discussed in the literature. As shown in Fig 9, the proposed framework outperforms traditional passive honeypots in key active-defense dimensions.

- Smokescreen clearly outperforms existing solutions in **Redirection Capability** and **Real-time Observability**.
- It maintains a strong balance between high effectiveness and **low deployment complexity/ resource overhead**, making it more practical than traditional high-interaction or orchestration-heavy frameworks.

Unlike purely passive systems [7,13] that focus on data collection without response, or ML-heavy adaptive solutions [8] that introduce training overhead, Smokescreen achieves a balanced trade-off between responsiveness, usability, and low resource consumption. The container-native design and lightweight monitoring pipeline further differentiate it from heavyweight SIEM-dependent frameworks commonly reported in earlier studies.

Discussion of findings and limitations

The experimental results validate the central hypothesis of this research: integrating hybrid detection, dynamic traffic redirection, and modern real-time observability into a containerized honeypot framework significantly enhances deception effectiveness while maintaining practicality. The near real-time redirection capability addresses a critical gap identified in the literature the inability of most existing systems to actively engage attackers without disrupting legitimate services or relying on crude blocking mechanisms.

The active closed-loop design of Smokescreen significantly reduces adversary dwell time. If we model average dwell time as D, then:

$$D_{\text{Smokescreen}} \ll D_{\text{Passive}}$$

Empirical results show that Smokescreen limits average dwell time to under 8 seconds compared to passive honeypots, which often allow attackers to remain undetected for several minutes [7, 19]. This improvement can be attributed to the low-latency redirection mechanism

Key findings indicate that the active closed-loop design reduces adversary dwell time and provides richer behavioral intelligence compared to passive

observation alone. The system's ability to operate effectively in a resource-constrained educational laboratory environment further demonstrates its accessibility for researchers and small-to-medium organizations.

Despite promising results, this study has several limitations. Detection currently relies on rule-based and threshold mechanisms, which may miss highly sophisticated or zero-day attacks. Evaluation was performed in a controlled simulated environment rather than long-term live internet exposure. Additionally, the current implementation is limited to a single-node deployment.

Conclusion

This study presented Smokescreen, a multi-layered honeypot framework designed to enhance cyber deception through the integration of hybrid threat detection, dynamic traffic redirection, and real-time monitoring within a lightweight containerized environment. The framework was developed to address the limitations of traditional passive honeypot systems, particularly their limited response capability, fragmented visibility, and reliance on resource-intensive infrastructures. By combining network-level and application-level monitoring, the architecture successfully detected reconnaissance activities, brute-force attempts, and suspicious behavioral patterns in near real time. Through programmable traffic control and automated redirection mechanisms, the system transparently isolated malicious traffic into dedicated honeypot environments while maintaining uninterrupted access for legitimate users.

Experimental evaluation demonstrated high detection accuracy, reliable redirection performance, low log-delivery latency, and stable operation under simulated attack conditions. The results indicate that combining active deception techniques with modern observability and lightweight containerization can significantly improve the responsiveness and practicality of honeypot-based defense systems. The modular design of Smokescreen allows users to define their own detection thresholds and redirection policies according to their specific security needs and operational context. In addition to capturing attacker behavior, the framework provided centralized visibility into security events, enabling rapid analysis and improved situational awareness.

Research and operational implications

Smokescreen contributes a practical research blueprint that successfully bridges theoretical deception techniques with operational active defense. The empirical findings validate that an active closed-loop design drastically reduces adversary dwell time and provides significantly richer behavioral intelligence compared to passive observation alone. Furthermore, the low-overhead, container-native design demonstrates that advanced cyber deception can be accessible and scalable for educational laboratories, academic research environments, and small-to-medium organizations with limited infrastructure resources. The active closed-loop design of Smokescreen complements recent game-theoretic defense frameworks [33] by providing practical redirection and isolation capabilities. By offering an effective, low-cost cyber threat intelligence solution, this framework shifts the paradigm for resource-constrained environments from reactive perimeter defense to proactive threat engagement.

Limitations and future work

Despite these promising results, several limitations remain that map out clear avenues for future research. The current implementation primarily relies on rule-based detection and threshold-driven classification, which may reduce effectiveness against highly sophisticated or previously unseen zero-day attacks. Additionally, the evaluation was conducted within a controlled, single-node laboratory environment rather than a long-term deployment on publicly exposed networks.

Future work will focus on expanding the framework by incorporating lightweight machine learning and adaptive detection techniques to counter evolving threats. Researchers plan to extend the architecture to support distributed multi-node honeypot orchestration, utilizing established container-virtualization performance

benchmarks and high-throughput hybrid network classification architectures. Furthermore, future iterations will explore seamless integration with advanced Security Information and Event Management (SIEM) platforms, automated threat intelligence sharing systems, and behavioral path modeling to evaluate active cyber deception dynamics against sophisticated adversaries [32]. This expanded approach will include implementing enhanced alerting mechanisms integrated with communication channels like Telegram or Slack, alongside long-term live-network evaluations in real-world environments to maximize collaborative cyber defense.

Ultimately, Smokescreen demonstrates that active, container-native deception systems provide an effective, practical, and highly responsive approach to modern cyber defense by unifying attacker engagement, intelligent traffic handling, and real-time monitoring.

Data and code availability statement

The source code, configuration files, deployment scripts, and experimental datasets associated with the Smokescreen framework are publicly available at the project repository: https://github.com/Bilal-Ahmed-Khalil/SmokeScreen.git

The repository includes container orchestration files, hybrid detection modules, traffic redirection configurations, monitoring dashboards, and attack simulation scripts used during the experimental evaluation.

REFERENCES

Franco J, Aris A, Canberk B, Uluagac AS. A survey of honeypots and honeynets for Internet of Things, industrial Internet of Things, and cyber-physical systems. IEEE Commun Surv Tutorials. 2021;23(4):2351-2383. doi:10.1109/COMST.2021.3106669

Dutta N, Galani M, Radhakrishnan R, Kumar R, Singh P. Using honeypots for ICS threats evaluation. In: Recent developments on industrial control systems resilience. London: Academic Press; 2020. p. 175-196.

Mphago B, Mpoeleng D, Masupe S. The role of deception in securing our cyberspace: Honeypots are a viable option. In: Applications of Machine Learning and Deep Learning for Privacy and Cybersecurity. Hershey (PA): IGI Global; 2022. p. 95-117.

Kocaogullar Y, Evmorfos S, Gankotiya L, Pande S. Hunting high or low: evaluating the effectiveness of high-interaction and low-interaction honeypots. J Cyber Secur Technol. 2022;14(2):15-31.

Do Hoang H, To TN, Ngo Duc Hoang Son, Ngo-Khanh K, Nguyen Tan Cam, Pham VH. Hawkeyes: An intelligent honeypot allocation strategy for cyber deception using reinforcement learning. Comput Networks. 2025;276:111982. doi:10.1016/j.comnet.2025.111982

Morić Z, Dakić V, Regvart D. Advancing cybersecurity with honeypots and deception technologies: A systematic review. Informatics. 2025;12(1):14. doi:10.3390/informatics12010014

Sachan RC, Kumar S, Gupta A, Sharma P, Mishra R. AI-Driven Adaptive Honeypots for Dynamic Cyber Threats. SSRN Electronic Journal. 2025. Preprint. doi:10.2139/ssrn.4966935

Yang X, Li Y, Zhao M, Zhang J, Wang X. A Highly Interactive Honeypot-Based Approach to Network Threat Management. Future Internet. 2023;15(4):127.

Sun J, Sun K, Li Q. Towards a believable decoy system: Replaying network activities from real system. In: Proceedings of the 2020 IEEE Conference on Communications and Network Security (CNS). Piscataway (NJ): IEEE; 2020. p. 1-9.

Pateria J, Ahuja L, Som S. Critical path to place decoys in Deception biota. Int J Performability Eng. 2022;18(12):854.

Pateria J, Ahuja L, Som S, Kumar V, Sharma A. Applying Clustering to Predict Attackers Trace in Deceptive Ecosystem by Harmonizing Multiple Decoys Interactions Logs. Int J Inf Technol Comput Sci. 2023;15:35-44.

Serem EK, Mugo DM, Too BK. Deceptive Decoys: Combining Believable User And Network Activities And Deceptive Network Setup In Enhancing Effectiveness. Int J Eng Res Technol. 2021;12(6):281-292.

Gaddam N. AI-enhanced honeypots for advanced cyber deception strategies. Int J Cyber Secur Res Dev. 2025;5(1):9-19. doi:10.63374/QITP-IJCSRD_05_01_002

Aggarwal P, Gonzalez C, Dutt V. Decoys in cybersecurity: an exploratory study to test the effectiveness of 2-sided deception. arXiv:2108.11037 [Preprint]. 2021.

Sharma S, Thakur S, Sharma M, Goyal R, Uttrani S, Katakwar H, et al. Exposing cyberthreats: How attack patterns are shaped by network topologies and honeypot techniques. Proc Hum Factors Ergon Soc Annu Meet. 2025;69(1):224-228. doi:10.1177/10711813251372531

Morozov DS, Ivanov IV, Petrov AV, Sidorov BC. Honeypot and cyber deception as a tool for detecting cyber attacks on critical infrastructure. In: Proceedings of the 3rd Edge Computing Workshop Doors. Nuremberg: Springer; 2023. p. 45-58.

Waizel G. Using a modern honeypot model to defend smart cities and provide early detection to APT and ransomware attacks. In: Smart Cities International Conference (SCIC) Proceedings. Bucharest: Politehnica Press; 2022. Vol. 10; p. 112-125.

Ngo HQ, Guo M, Nguyen H. Optimizing Cyber Response Time on Temporal Active Directory Networks Using Decoys. arXiv:2403.18162 [Preprint]. 2024.

Kumar VS. Adaptive Distributed Honeypot Detection Network for Enhanced Cybersecurity Against DoS and DDoS Attacks. Results Eng. 2025;26:105521.

Bouarfa A. Intelligent Honeypot-Based IDS for Cyber-Attack Detection. Mathematics. 2025;17(5):628.

Suratkar S, Shirole M, Bhirud S, Kulkarni A, Joshi R. An adaptive honeypot using Q-Learning with severity analyzer. J Ambient Intell Humaniz Comput. 2022;13(10):4865-4876.

Chen L, Li W. Distributed polymorphic honeypots for intelligent malware capture in IoT. IEEE Internet Things J. 2025;12(3):1450-1462.

Vasylyshyn S, Melnyk O, Kovalenko O, Fedorov V. A model of decoy system based on dynamic attributes for cybercrime investigation. Eastern-Eur J Enterp Technol. 2023;1(9):121.

Dougherty JT. Evasion of honeypot detection mechanisms through improved interactivity of ICS-based systems [master's thesis]. Monterey (CA): Naval Postgraduate School; 2020.

Anwar AH, Leslie NO, Kamhoua CA. Honeypot allocation for cyber deception in internet of battlefield things systems. In: 2021 IEEE Military Communications Conference (MILCOM). San Diego (CA): IEEE; 2021. p. 412-419.

Perkins RC, Howell CJ. Honeypots for cybercrime research. In: Researching Cybercrimes: Methodologies, Ethics, and Critical Approaches. London: Palgrave Macmillan; 2021. p. 233-261.

Kheir N, Abdelrazek L, Daniel CHO. Demo Paper: Caught in my Radio Net-Experiment with Honeypots in Radio Access Networks. In: 2022 25th Conference on Innovation in Clouds, Internet and Networks (ICIN). Paris: IEEE; 2022. p. 88-91.

Alhosani N, Alrabaee S, Al Faresi A. An Efficient Strategy for Deploying Deception Technology. In: International Conference on Future Access Enablers of Ubiquitous and Intelligent Infrastructures. Cham: Springer Nature Switzerland; 2023. p. 204-216.

Guan C, Cao G. Cyber-Physical Deception Through Coordinated IoT Honeypots. In: Proceedings of the 34th USENIX Security Symposium. Vancouver: USENIX Association; 2025. p. 763-780.

Inamdar F, Muanawma F, Gandhi A, Vanlalruati J, Jadhav T. AI-powered adaptive honeypot for advanced threat intelligence. Int Res J Adv Eng Hub. 2026. doi:10.47392/IRJAEH.2026.0109

Balta EC, Pease M, Moyne JR, Barton K, Tilbury DM. Digital twin-based cyber-attack detection framework for cyber-physical manufacturing systems. IEEE Trans Autom Sci Eng. 2024;21(2):1695-1712. doi:10.1109/TASE.2023.3243147

Weng Z, Chen Y, Zhang H, Li X, Wang Y, Liu J. Remote medical system driven by medical big models: Dynamic defense model for network security threats. PLoS ONE. 2026;21(3):e0348572. doi:10.1371/journal.pone.0348572

Chen J, Li Y, Zhang X, Wang Z, Liu H. A Bayesian game approach for node-based attribution defense against asymmetric information attacks in IoT networks. PLoS ONE. 2025;20(1):e0316091. doi:10.1371/journal.pone.0316091

International Journal of Advanced Research