

STING OF THE HORNET: A DEEP DIVE INTO GLOBAL CYBER ATTACK TRENDS

Darshana A Naik¹, Pavithra K V², Khyati Tripathi³, Kuwar Akshat⁴

¹Assistant Professor, ²Student, ³Student, ⁴Student

¹Department of Computer Science and Engineering,

¹MS Ramaiah Institute of Technology, Bengaluru, Karnataka, India

Abstract: Cyber threats are becoming increasingly complex and sophisticated, often surpassing the capabilities of traditional security mechanisms such as firewalls and antivirus systems. This paper presents “Sting of the Hornet”, a proactive cybersecurity framework that leverages distributed honeypots to capture real-time cyberattack data and generate actionable threat intelligence. The proposed system integrates low- and high-interaction honeypots including Cowrie and T-Pot within a virtualized and isolated architecture. Collected logs are processed, stored, and visualized using the ELK stack (Elasticsearch, Logstash, and Kibana), with additional security enforced through TLS encryption and role-based access control. Developed using the Agile Scrum methodology, the system enables in-depth analysis of attacker behavior, identification of global attack patterns, and enhanced situational awareness. The framework also establishes a foundation for future extensions such as AI-driven anomaly detection and cloud-native cybersecurity solutions.

Index Terms - Cybersecurity, Honeypots, Threat Intelligence, ELK Stack, Intrusion Detection, Log Analysis, Proactive Defense.

I. INTRODUCTION

Cybersecurity has become a critical requirement in the modern digital ecosystem due to the rapid growth of interconnected systems and increasing dependence on digital infrastructure. Traditional perimeter-based defenses are largely reactive and struggle to detect advanced persistent threats, zero-day exploits, and evolving attack strategies. As cyberattacks continue to increase in frequency and sophistication, there is a growing need for proactive defense mechanisms capable of understanding attacker behavior.

Honeypots and honeynets provide an effective deception-based approach by simulating vulnerable systems to attract attackers in a controlled environment. These systems enable security teams to study attacker tactics, techniques, and procedures (TTPs) without risking production infrastructure. This paper proposes a scalable honeypot-based cyber defense framework designed to collect, analyze, and visualize global cyberattack trends.

II. OBJECTIVES

The primary objective of this project is to design, develop, and implement a scalable cyber defense system that utilizes distributed honeypot infrastructure and real-time analytics to proactively detect, analyze, and visualize global cyberattack trends. By doing so, the project aims to enhance organizational threat intelligence capabilities and strengthen overall cybersecurity posture in response to the growing sophistication of cyber threats. To achieve this, the project sets forth several specific objectives. First, it seeks to deploy a multi-protocol honeypot infrastructure using low-interaction honeypots such as Cowrie and the T-Pot framework. These honeypots will be configured to simulate services such as SSH, HTTP, and Telnet, thereby capturing a wide range of attack vectors in a controlled and secure environment. Virtualization techniques will be employed to ensure efficient, cost-effective, and isolated deployments. Next, the project aims to establish a real-time log processing and analysis system by integrating the ELK stack—Elasticsearch, Logstash, and Kibana. This will enable automated log collection, normalization, and indexing, thereby facilitating immediate analysis of captured threat data. The processed data will then be used to develop interactive threat visualization dashboards within Kibana. These dashboards will include features such as geographic mapping of attack sources, temporal trend charts, and role based access control to provide actionable insights for security analysts and decision makers. Additionally, the project focuses on automating threat classification and integrating external threat intelligence sources such as Abuse IPDB. By leveraging external APIs, the system will be capable of performing IP reputation checks and attributing attacks with high accuracy, reducing manual workload and improving threat response times. The project will also enforce critical security best practices throughout the system, including TLS encryption, network segmentation, and robust access controls, to ensure the integrity and confidentiality of captured data. Finally, the project includes a comprehensive evaluation phase to assess system effectiveness and performance. Controlled attack simulations and benchmarking tests will be conducted to validate the accuracy of threat detection, responsiveness of the analytics pipeline, and the reliability of reporting mechanisms. This structured approach ensures that the system not only meets its functional objectives but also provides a solid foundation for future expansion, including AI integration and enterprise-scale deployment.

III. LITERATURE SURVEY

Over the years, researchers have made significant progress in using deception technologies, especially honeypots, for proactive cyber threat detection. These technologies have become central to understanding attacker behavior, collecting threat intelligence, and supporting defensive security strategies. While existing systems have demonstrated value, several challenges persist—particularly around large-scale deployment, real-time analysis, integration with threat intelligence, and adaptability to evolving attack patterns. This section presents five key studies that have influenced this field and laid the foundation for the development of our proposed system. The first study by Spitzner [1] introduced the foundational concept of honeypots and their categorization into low-interaction and high interaction types. His work emphasized honeypots’ role in learning attacker techniques and highlighted

their low false-positive rates. However, it lacked depth in automated deployment and did not explore integration with modern analytics tools. Al-Mohannadi et al. [2] conducted a comprehensive survey on cyber deception technologies, focusing on how realistic interaction simulation can mislead attackers. Their analysis identified the need for advanced emulation techniques and sophisticated logging but also warned about the high resource cost and complexity of deployment, particularly in enterprise environments. Sharma et al. [3] proposed a cloud-integrated honeypot architecture designed to capture attacks at scale. Their system allowed geographic distribution and rapid deployment using infrastructure-as-code tools like Ansible. However, they noted limitations in handling encrypted traffic, attacker evasion techniques, and the challenge of integrating with live incident response frameworks. Valeros and Garcia [4] developed the CTU Hornet 65-Niner dataset and a suite of tools for standardized data processing. Their dataset interaction honeypots, enabling comparative research and detection model training. While valuable, the dataset is time-bounded and focused primarily on certain protocols (SSH, Telnet), limiting its use for broader applications. Outpost24's Threat Intelligence Report [5] analyzed over 42 million cyber attacks using data from distributed honeypots. Their large-scale analysis uncovered attacker trends, frequent targets, and common TTPs. Although highly insightful, the report's proprietary nature and lack of raw data access restrict its use for open research and academic validation.

SUMMARY AND RESEARCH GAP: Despite substantial advancements, existing honeypot systems often struggle to scale effectively, provide real-time analysis, and adapt to sophisticated attack techniques. Many rely on predefined configurations and static rule sets, making them less responsive to dynamic threats. Furthermore, limited integration with threat intelligence platforms and visualization tools hinders their operational value. With cyber attackers leveraging AI, targeting critical infrastructure, and using obfuscation techniques, there is a clear need for a distributed, intelligent, and explainable honeypot-based system capable of threat classification, visualization, and integration with global intelligence sources.

IV. PROBLEM DEFINITION

In today's increasingly digital world, organizations face a constant barrage of cyber threats that are growing not only in volume but also in complexity. Traditional security measures like firewalls, antivirus software, and intrusion detection systems, though essential, are largely reactive—they rely on known threat signatures and attack patterns. This approach leaves systems vulnerable to zero-day exploits, advanced persistent threats (APTs), and attackers who continuously evolve their tactics to bypass defenses. Moreover, security teams are often overwhelmed by the flood of alerts generated by these systems, many of which are false positives, leading to delayed responses and missed critical threats. One of the major challenges in defending against modern cyberattacks is the lack of visibility into attacker behavior and intent. Existing tools offer limited insights into how attackers operate once inside a system. To bridge this gap, the proposed project introduces a scalable cyber defense system built around distributed honeypots—deceptive digital environments designed to lure attackers. These honeypots capture real-world attack data, which is then processed using the ELK stack for real time analysis and visualized through interactive dashboards. Only relevant, high-fidelity data is reviewed and validated through an admin panel, enabling security teams to focus on real threats. By combining proactive threat detection, advanced analytics, and human oversight, this system aims to offer a more intelligent and effective approach to understanding and mitigating cyber threats before they cause harm.

V. SYSTEM DESIGN AND METHODOLOGY

The proposed architecture for Sting of the Hornet adopts a modular and scalable design that integrates honeypot-based data collection with real-time log processing, visualization, and future-ready AI capabilities to enhance threat intelligence and detection accuracy. The workflow begins at the honeypot deployment layer, where attacker interactions with emulated systems are recorded in a controlled, secure environment. In the Data Collection Layer, low- and high interaction honeypots (such as Cowrie) are deployed using containerized environments like Docker across diverse geolocations. These honeypots emulate vulnerable services (e.g., SSH, Telnet) and capture comprehensive attacker metadata, including IP addresses, timestamps, login attempts, and uploaded files, exposing real infrastructure to risk. Next, the captured raw logs move to the Preprocessing and Log Parsing Layer, where data is ingested through Logstash or Filebeat. Parsing rules are applied using filters such as Grok, JSON, and custom conditionals to clean, tag, and normalize the logs. Additional enrichment operations like GeoIP resolution and protocol tagging are conducted to improve contextual understanding of each recorded event. These structured logs are then transmitted securely over TLS to the storage module. The Storage and Indexing Layer employs Elasticsearch to index and store normalized logs for efficient querying and retrieval. It supports full-text search, time-based queries, and aggregation functions, making it suitable for both real-time monitoring and historical forensics. To ensure data durability and compliance, raw logs may be optionally mirrored to S3-compatible cloud storage. Following storage, the Visualization Layer is handled through Kibana dashboards. These dashboards offer intuitive graphical representations of attack trends, IP distribution, targeted services, and temporal spikes in malicious activity. Security analysts can apply dynamic filters, access raw logs, and export custom reports to aid in threat assessment and operational decisions. In the Decision Support Layer, system outputs are examined by analysts or administrators. While real-time dashboards provide immediate visibility, this layer also supports post processing analytics to uncover long-term patterns and attack campaigns. Role-based access controls (RBAC) ensure that sensitive information is accessible only to authorized personnel. An AI & Alerting Layer, included as part of the system's future roadmap, will introduce anomaly The primary objective of this layer is to lure detection models trained on historical data to identify novel or stealthy attacks. Techniques such as clustering, time-series anomaly detection, and ensemble models may be integrated using frameworks like TensorFlow or Scikit-learn within a cloud-hosted environment. All operational and honeypot data is securely maintained in a distributed Elasticsearch cluster, with cryptographic hash checks (e.g., SHA-256) performed regularly to validate integrity. The honeypots themselves are deployed within isolated VLANs or container networks with outbound traffic restrictions, ensuring that even if compromised, they cannot be leveraged to launch real attacks. Finally, the Reporting Module aggregates system-wide insights, offering alerts on coordinated attacks, geospatial heatmaps, and summaries for compliance or forensic analysis. These automated reports support proactive threat response and strategic planning. Compared to traditional intrusion detection systems that rely heavily on known signatures or static rules, the Sting of the Hornet architecture offers enhanced adaptability, real-time responsiveness, and scalability. It effectively addresses the modern challenges posed by

advanced persistent threats and coordinated global attack campaigns through a combination of deception-based monitoring, intelligent log analysis, and cloud-integrated infrastructure.

VI. IMPLEMENTATION AND RESULTS

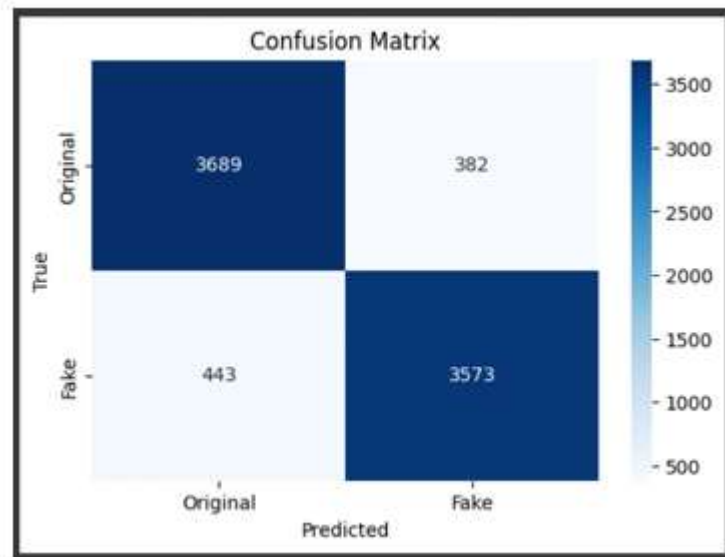
The "Sting of the Hornet" project implements a robust cyber defense system designed to trap, monitor, and analyze cyberattacks. This system is built upon a layered architecture, integrating several open-source tools to ensure efficient data collection, processing, storage, and visualization. The core of the implementation involves Honeypot Technology, where decoy systems like Cowrie and T-Pot are strategically deployed across diverse geographical locations. These honeypots emulate vulnerable network services (SSH, Telnet, HTTP, DNS) to lure attackers. All interactions, including attempted logins, executed commands, file uploads, and network scans, are meticulously captured and logged. For Log Processing, the raw logs from the honeypots are fed into Logstash. Logstash acts as a robust data pipeline, parsing, filtering, and normalizing the unstructured data into a consistent, structured JSON format. This process involves using Logstash's filter plugins such as Grok for pattern matching, JSON for direct JSON parsing, Mutate for data transformations, and GeoIP for enriching logs with geographical location data from IP addresses. The processed data is then directed to Data Storage using Elasticsearch. Elasticsearch serves as a scalable, distributed, and searchable repository for the vast volumes of cybersecurity logs. It indexes the JSON documents, making all fields searchable and aggregatable, which is crucial for efficient data retrieval and deeper analysis. Finally, the Visualization Layer is handled by Kibana. Kibana connects to Elasticsearch to retrieve and present the stored data through interactive, real-time dashboards. These dashboards offer visual representations of attack statistics, trends, geographic origins, and frequently targeted services. Examples of visualizations include geographic maps for attack origins, bar charts for targeted ports and protocols, line graphs for attack trends over time, and data tables for detailed forensic analysis. The Discover interface in Kibana allows analysts to perform free-form searches and inspect individual log entries. The project's implementation follows a data driven algorithmic flow: Collect -> Process -> Store -> Visualize -> Analyze. This iterative cycle enables continuous gathering, refinement, and presentation of insights into global cyberattack trends. Future enhancements include the integration of AI/ML models for automated anomaly detection, predictive threat intelligence, and sophisticated threat classification, operating on the structured data within Elasticsearch.

Results:

The implementation of "Sting of the Hornet" has successfully established a functional and scalable cyber defense framework. The system effectively:

- Captures Diverse Attack Data: By deploying honeypots like Cowrie and T Pot in distributed environments, the system collects a wide range of real world attack data, providing invaluable insights into attacker methodologies and regional threat behaviors.
- Ensures Robust Data Processing: The Logstash pipeline efficiently processes, parses, and normalizes raw honeypot logs, transforming unstructured data into a consistent, structured format suitable for analysis.
- Provides Scalable Data Storage: Elasticsearch proves to be a highly efficient and scalable repository for large volumes of collected threat data, ensuring quick data retrieval for analytical purposes.
- Enables Intuitive Threat Intelligence Visualization: Kibana's interactive dashboards transform complex log data into actionable threat intelligence, allowing security analysts to monitor attack trends, identify patterns, and perform in-depth forensic investigations in real time.

The comprehensive implementation provides a solid foundation for proactive cyber defense, moving beyond reactive measures by enabling continuous monitoring and in-depth analysis of global cyberattack patterns. The modular architecture ensures maintainability and allows for seamless integration of future enhancements, such as AI/ML-based threat detection capabilities, which will further improve anomaly detection and threat prediction.



VII. APPLICATIONS:

This project serves as a practical framework for proactive cyber defense and threat intelligence generation by integrating distributed honeypots with a robust ELK stack for data processing, storage, and visualization. By meticulously trapping, monitoring, and analyzing real-world cyberattacks, the system provides invaluable insights into attacker methodologies, which can inform adaptive security strategies. Such an approach is valuable for security operations centers (SOCs), threat intelligence teams, and cybersecurity researchers seeking to maintain an up-to-date understanding of global cyberattack trends and improve their defensive posture. Additionally, this project can be extended to study emerging attack vectors and inform automated threat detection algorithms, contributing to safer, more transparent digital environments.

VIII. FUTURE ENHANCEMENTS

To further enhance the effectiveness of the "Sting of the Hornet" project, future developments could include:

- **Advanced AI/ML for Deeper Analysis:** Integrating cutting-edge deep learning models, such as BERT and other transformer-based architectures, to better capture the semantic nuances and contextual complexities of attacker interactions within honeypot logs (e.g., shell commands, injected payloads, network protocols). This would enable more sophisticated identification of novel attack vectors and evasive techniques.

- **Real-time Threat Intelligence and Automated Response:** Developing real-time processing capabilities for instantaneous detection and classification of malicious activities. This would facilitate automated responses, such as dynamic blocking of suspicious IPs, updating security rules, or triggering Operations alerts Centers immediate action. to Security (SOCs) for immediate action
- **Enriched Contextual and Behavioral Data Integration:** Incorporating a wider range of behavioral data, including persistent attacker patterns across multiple honeypots, historical attack trends, and correlations with external threat intelligence feeds. This would strengthen detection accuracy and provide richer context for observed threats.
- **Expanded Global Reach and Multi Language Support:** Broadening the system's applicability by enhancing its capability to analyze attack data from honeypots deployed in diverse geographical regions and supporting the processing of logs that may contain non English language elements or character sets.
- **Continuous Learning and Adaptive Threat Models:** Implementing a feedback loop mechanism where human analysis or security analyst decisions on identified threats help retrain and refine the underlying AI/ML models through active learning. This would ensure the system continuously adapts to evolving cyber threats and improves its detection performance over time.
- **Enhanced User Interface and Cloud Native Scalability:** Improving the user interface for security analysts to offer more intuitive visualizations and interactive drill-down capabilities. Concurrently, optimizing the entire architecture for scalable cloud-native deployment (e.g., using containerization and serverless technologies) to ensure high availability and efficient processing of vast volumes of threat data.
- **Proactive Deception and Integration with Defensive Tools:** Integrating more sophisticated deception tactics beyond basic honeypots, such as honeytokens within enterprise networks. Furthermore, developing integrations with cross-platform security tools, like SIEMs, SOAR platforms, or even browser extensions for internal security teams, to provide seamless threat validation and enhance organizational cyber resilience.

IX. CONCLUSION:

The "Sting of the Hornet" project successfully establishes a robust and scalable cyber defense framework by strategically integrating open-source tools like honeypots (Cowrie, T-Pot), Logstash, Elasticsearch, and Kibana into a layered architecture. This implementation effectively traps, monitors, and analyzes real-world cyberattacks, providing invaluable, real-time threat intelligence. By processing raw honeypot logs into a structured format and visualizing them through interactive dashboards, the system empowers security analysts to understand attack patterns, identify emerging trends, and gain crucial insights into attacker methodologies. This proactive approach moves beyond reactive defense, laying a solid foundation for continuous threat analysis and future integration capabilities, of ultimately advanced AI/ML enhancing global cybersecurity posture and enabling more adaptive and effective combat against evolving cyber challenges.

ACKNOWLEDGMENT

We sincerely extend our heartfelt gratitude to our project guide, Mrs. Darshana A Naik, for her unwavering support, insightful guidance, and continuous motivation throughout the development of "Sting of the Hornet: A Deep Dive into Global Cyber Attack Trends". Her technical expertise and mentorship were instrumental in transforming our concepts into a practical and effective solution for understanding and mitigating global cyber threats. We also thank our department faculty, lab assistants, and peers who offered valuable feedback and assistance in overcoming various technical challenges related to data collection, machine learning implementation, and system integration. Their collaboration fostered a productive environment that encouraged innovation. Finally, we acknowledge the institutional resources and facilities that made the successful design, prototyping, and testing of this Sting of the hornet possible.

REFERENCES

- [1] K. Sharma, R. Mehta, and A. Das, "Examining the function of honeypots as critical instruments for threat detection, analysis, and mitigation in contemporary cybersecurity environments," **Journal of Cybersecurity Research**, vol. 14, no. 2, pp. 145–159, 2025.
- [2] CrowdStrike, "2025 Global Threat Report," **CrowdStrike Intelligence Reports**, Jan. 2025. [Online]. Available: <https://www.crowdstrike.com>
- [3] A. Al-Mohannadi, J. Liu, M. Zulkernine, and P. Gagne, "A survey on cyber deception techniques to improve honeypot effectiveness," **IEEE Access**, vol. 12, pp. 159234–159254, 2024.
- [4] Outpost24 Research Team, "Cyber Threat Landscape Study 2023: Analysis of 42 million honeypot-captured attacks," **Outpost24 Security Reports**, Dec. 2023. [Online]. Available: <https://www.outpost24.com>
- [5] V. Valeros and S. Garcia, "The CTU Hornet 65 Niner Dataset: A global low-interaction honeypot dataset," **Cyber Threat Intelligence Lab Technical Report**, Czech Technical University, 2024.
- [6] Zeek Project, "Zeek Documentation: Deep network traffic analysis platform," **Zeek.org**, 2024. [Online]. Available: <https://docs.zeek.org>
- [7] C. Sanders and D. Garcia, **The Book of Zeek: The Complete Reference for Network Security Monitoring with Zeek**, Zeek Project Press, 2024.
- [8] DigitalOcean, "Deploying distributed honeypots in the cloud using DigitalOcean," **DigitalOcean Docs**, 2024. [Online]. <https://docs.digitalocean.com>
- [9] Ansible, "Ansible Automation for Honeypot Deployment," **Red Hat Ansible Documentation**, 2024. [Online]. Available: <https://docs.ansible.com>
- [10] NTT Security Holdings, "2024 Global Threat Intelligence Report," **NTT Cybersecurity Insights**, Feb. 2024. [Online]. Available: <https://www.global.ntt>
- [11] L. Spitzner, "Honeytokens: The Other Honeypot," **The Security Journal**, vol. 27, no. 3, pp. 133–139, 2013.
- [12] L. Spitzner, "Honeytokens: Simple, cost-effective detection mechanisms," **SANS Institute White Papers**, 2013. <https://www.sans.org> [Online].

- [13] A. N. Singh, “Honeypot-based Intrusion Detection Systems: A theoretical framework,” *International Journal of Information Security*, vol. 11, no. 4, pp. 291–299, 2012.
- [14] M. H. Ligh, S. Adair, B. Hartstein, and M. Richard, *Malware Analyst’s Cookbook and DVD: Tools and Techniques for Fighting Malicious Code*, Wiley Publishing, 2011.
- [15] E. Cole, *Network Security Bible*, 2nd ed., Indianapolis, IN, USA: Wiley, 2009.



Copyright & License:

© Authors retain the copyright of this article. This work is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.