UNIVERSITY OF THESSALY

DIGITAL SYSTEMS DEPARTMENT

MASTER'S PROGRAM IN

«MODERN COMMUNICATION SYSTEMS AND INTERNET OF THINGS»

MASTER THESIS

«Implementation of a Security Information and Event Management (SIEM) System with Integrated Extended Detection and Response (XDR) for Host and Network Monitoring»

Sotirios Milionis

May 2025

ΠΑΝΕΠΙΣΤΗΜΙΟ ΘΕΣΣΑΛΙΑΣ

ΤΜΗΜΑ ΨΗΦΙΑΚΩΝ ΣΥΣΤΗΜΑΤΩΝ

ΠΡΟΓΡΑΜΜΑ ΜΕΤΑΠΤΥΧΙΑΚΩΝ ΣΠΟΥΔΩΝ

«ΣΥΓΧΡΟΝΑ ΣΥΣΤΗΜΑΤΑ ΕΠΙΚΟΙΝΩΝΙΩΝ ΚΑΙ ΔΙΑΔΙΚΤΥΟ ΤΩΝ ΠΡΑΓΜΑΤΩΝ»

ΜΕΤΑΠΤΥΧΙΑΚΗ ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

«Υλοποίηση Συστήματος Διαχείρισης Πληροφοριών και Συμβάντων Ασφαλείας (SIEM) με Ενσωματωμένο Σύστημα Ανίχνευσης και Απόκρισης (XDR) για Παρακολούθηση Συσκευών και Δικτύου»

Σωτήριος Μηλιώνης

Μάιος 2025

COPYRIGHTS

UNIVERSITY OF THESSALY

DIGITAL SYSTEMS DEPARTMENT

MASTER OF SCIENCE

«MODERN COMMUNICATION SYSTEMS AND THE INTERNET OF THINGS»

«Implementation of a Security Information and Event Management (SIEM) System with Integrated Extended Detection and Response (XDR) for Host and Network Monitoring»

submitted to the Department of Digital Systems

of the University of Thessaly

as part of the requirements for obtaining a Postgraduate Diploma in

Modern Communication Systems and the Internet of Things from

Sotirios Milionis

ΔΗΛΩΣΗ ΑΥΘΕΝΤΙΚΟΤΗΤΑΣ

ΠΑΝΕΠΙΣΤΗΜΙΟ ΘΕΣΣΑΛΙΑΣ

ΤΜΗΜΑ ΨΗΦΙΑΚΩΝ ΣΥΣΤΗΜΑΤΩΝ

ΠΡΟΓΡΑΜΜΑ ΜΕΤΑΠΤΥΧΙΑΚΩΝ ΣΠΟΥΔΩΝ

«ΣΥΓΧΡΟΝΑ ΣΥΣΤΗΜΑΤΑ ΕΠΙΚΟΙΝΩΝΙΩΝ ΚΑΙ ΔΙΑΔΙΚΤΥΟ ΤΩΝ ΠΡΑΓΜΑΤΩΝ»

«Υλοποίηση Συστήματος Διαχείρισης Πληροφοριών και Συμβάντων Ασφαλείας (SIEM) με Ενσωματωμένο Σύστημα Ανίχνευσης και Απόκρισης (XDR) για Παρακολούθηση Συσκευών και Δικτύου»

που υποβλήθηκε στο Τμήμα Ψηφιακών Συστημάτων του

Πανεπιστημίου Θεσσαλίας

ως μέρος των απαιτήσεων για την απόκτηση

Διπλώματος Μεταπτυχιακών Σπουδών στα Σύγχρονα Συστήματα Επικοινωνιών και Διαδίκτυο των Πραγμάτων από τον

Σωτήριο Μηλιώνη του Δημητρίου

«Ο μεταπτυχιακός φοιτητής που εκπόνησε την παρούσα διπλωματική εργασία φέρει ολόκληρη την ευθύνη προσδιορισμού της δίκαιης χρήσης του υλικού, η οποία ορίζεται στη βάση των εξής παραγόντων: του σκοπού και χαρακτήρα της χρήσης (μη-εμπορικός, μη-κερδοσκοπικός, αλλά εκπαιδευτικός-ερευνητικός), της φύσης του υλικού που χρησιμοποιεί (τμήμα του κειμένου, πίνακες, σχήματα, εικόνες κ.λπ.), του ποσοστού και της σημαντικότητας του τμήματος που χρησιμοποιεί σε σχέση με το όλο κείμενο υπό copyright, και των πιθανών συνεπειών της χρήσης αυτής στην αγορά ή την γενικότερη αξία του υπό copyright κειμένου.»

Committee

This thesis was unanimously approved by the three-member examination committee appointed by the Assembly of the Department of Digital Systems at the University of Thessaly, in accordance with the law and the approved Study Guide of the Master's Program «Modern Communication Systems and the Internet Of Things». The members of the Committee were:

- George Karetsos (Supervisor)
- Christos Ilioudis (Member)
- Antonios Sarigiannidis (Member)

The approval of this thesis by the Department of Digital Systems at the University of Thessaly does not necessarily imply endorsement of the author's views.

Τριμελής Επιτροπή

Η παρούσα διπλωματική εργασία εγκρίθηκε ομόφωνα από την τριμελή εξεταστική επιτροπή η οποία ορίστηκε από την Συνέλευση του Τμήματος Ψηφιακών Συστημάτων του Πανεπιστημίου Θεσσαλίας, σύμφωνα με το νόμο και τον εγκεκριμένο Οδηγό Σπουδών του ΠΜΣ «Σύγχρονα Συστήματα Επικοινωνιών και Διαδίκτυο των Πραγμάτων». Τα μέλη της Επιτροπής ήταν:

- Γεώργιος Καρέτσος (Επιβλέπων)
- Χρήστος Ηλιούδης (Μέλος)
- Αντώνιος Σαρηγιαννίδης (Μέλος)

Η έγκριση της διπλωματικής εργασίας από το Τμήμα Ψηφιακών Συστημάτων του Πανεπιστημίου Θεσσαλίας, δεν υποδηλώνει αποδοχή των απόψεων του συγγραφέα.

Abstract

Nowadays the field of cybersecurity has become more and more important due to the constant evolution of technology. Companies and individuals have started to realize that the once innocent Internet is now a distant past. The detection and prevention of attacks is only possible by using the correct tools and technologies in such a way that they effectively identify vulnerabilities, mitigate potential threats, and ensure the system's integrity and availability.

Protecting the diverse systems within a company is undoubtedly a challenging task. A variety of proactive technologies are required to do so, and Security Information and Event Management (SIEM) as well as Security Orchestration Automation and Response (SOAR) are some of the most important ones. Those terms constitute the foundations of a workplace such as a Security Operations Center (SOC).

Fortunately, all actions performed on a device are logged, providing a detailed record. Additionally, technologies like Intrusion Detection Systems (IDSs) can monitor the network behavior of a device and the different tasks running on it and are able to draw important conclusions. On the other hand, Intrusion Prevention Systems (IPSs) are responsible for acting in the event of a breach.

This thesis focuses on the implementation of a SIEM system along with an IDS and demonstrates the importance of the above-mentioned strategies. The role of components such as firewalls, log systems and cyber defense tactics is also highlighted through their involvement in the project. Moreover, a series of controlled tests is conducted on vulnerable systems, illustrating the stages of the cyber kill chain and highlighting the challenges and complexities involved in detecting and preventing cyber-attacks.

Keywords

Cybersecurity, Real-time Monitoring, Threat Detection, Incident Response

Περίληψη

Σήμερα, η κυβερνοασφάλεια είναι πιο σημαντική από ποτέ λόγω της συνεχούς τεχνολογικής εξέλιξης. Είναι πλέον ευρέως αποδεκτό ότι το Ίντερνετ δεν αποτελεί απλώς ένα αθώο εργαλείο που διευκολύνει τη ζωή μας. Οι επιτιθέμενοι χρησιμοποιούν εξελιγμένα εργαλεία για να παραβιάσουν τα συστήματα, και οι αναλυτές ασφαλείας είναι αυτοί που πρέπει να αναλάβουν δράση. Η ανίχνευση και πρόληψη των επιθέσεων είναι εφικτές μόνο με τη χρήση των σωστών εργαλείων που είναι ικανά να εντοπίσουν αποτελεσματικά τις ευπάθειες, να μειώνουν τις ενδεχόμενες απειλές και να διασφαλίσουν την ακεραιότητα και τη διαθεσιμότητα του συστήματος.

Η προστασία των συστημάτων ενός οργανισμού είναι μια δύσκολη αποστολή και απαιτεί τεχνολογίες όπως η Διαχείριση Πληροφοριών και Συμβάντων Ασφαλείας (SIEM) και ο Αυτοματισμός Ενορχήστρωσης και Απόκρισης Ασφαλείας (SOAR). Αυτοί οι όροι αποτελούν αναμφίβολα τα θεμέλια ενός Κέντρου Επιχειρησιακής Ασφαλείας (SOC).

Η καταγραφή των ενεργειών σε αρχεία (logs) και η χρήση Συστήματος Ανίχνευσης Εισβολής (IDS) για παρακολούθηση της συμπεριφοράς του δικτύου είναι σημαντικά εργαλεία για την εξαγωγή συμπερασμάτων. Επιπλέον, τα Συστήματα Πρόληψης Εισβολής (IPS) είναι υπεύθυνα να αναλάβουν δράση σε περίπτωση ανίχνευσης παραβίασης ενός συστήματος.

Η παρούσα εργασία εστιάζει στην υλοποίηση ενός συστήματος Διαχείρισης Πληροφοριών και Συμβάντων Ασφαλείας (SIEM) με Ενσωματωμένο Σύστημα Ανίχνευσης Εισβολών (IDS) και αναδεικνύει τη σημασία αυτών των στρατηγικών στην κυβερνοασφάλεια. Επίσης, αναλύει τον ρόλο των τειχών προστασίας, των συστημάτων καταγραφής και των τακτικών κυβερνοάμυνας μέσω της συμμετοχής τους στη διαδικασία. Τέλος, πραγματεύεται ελεγχόμενες επιθέσεις σε ευάλωτες συσκευές με σκοπό να αναδείξει τη δυσκολία και την πολυπλοκότητα της προστασίας των συστημάτων στις μέρες μας.

Λέξεις – Κλειδιά

Κυβερνοασφάλεια, Ανίχνευση Ανωμαλιών, Διαχείριση Απειλών, Αντιμετώπιση συμβάντων

Table of Contents

List of Figures

1. Introduction

1.1 Thesis' topic

The trend of today is to have everything connected to the Internet. From the refrigerator to our car, everything is connected because that makes our life simpler and better. With the introduction of a new era of labor, remote work has become more and more popular during the past few years. The establishment of high-speed Internet connections, even in remote areas, plays a crucial role in this. It is of our instinct that we have to protect our home and everything we love, but unfortunately, this is not quite what happens when it comes to digital assets.

Undeniably, the most frequent targets for cyber-attacks are states and big companies. However, there are plenty of threat actors in the world who are interested in personal data like, for example financial and health data and even in our everyday life preferences. Every category of data is important for someone nowadays because data can be sold or manipulated in many ways.

According to [1], "Cybercrime and cyberattacks are real and unavoidable. Even the 'best' protected services and businesses can be attacked and exploited. Vulnerabilities exist in every asset and organization, which can be exploited or used as a conduit to cause exposure or compromise to critical services." It is made clear that there is no longer a question regarding if a cyber-attack will happen but rather when it will happen. The worst mistake an administrator can make is to consider an infrastructure completely safe. There will always be a vulnerability in a system, device, or even in the human factor.

As mentioned in [2], the layered prevention-only approaches such as Firewall, Intrusion Detection System (IDS), Intrusion Prevention System (IPS) etc. can find and prevent malicious activity but they are not capable of detecting zero-day attacks. The many different kinds of technologies used in today's corporate networks have made the work of administrators very complex and exhausting. It is clear that the layered defense structure implemented in most middle and small sized companies is not enough anymore to truly protect them. Detection and response have already started to be integrated everywhere, even in smaller environments.

However, even with detection and prevention measures in place, analysts need a clear view of the network's health and the devices using it. Due to the vast amount of logs and information generated by network devices, detecting malicious behavior can be daunting. Aggregating this data in a single location is the solution. Additionally, correlating different types of data from various devices can be crucial in detecting an attack. Log collection, normalization, correlation, and analysis are most easily conducted by the use of Security Information and Event Management (SIEM) tools. Altogether, SIEM and Log management have become a necessity in corporations, as analyzing logs can detect these long-term aggressors and attackers. Without the analysis of these event logs, it would be difficult for organizations to detect if any of their assets have been compromised. This is why companies keep investing more money into SIEM systems [3].

This thesis highlights the critical importance of implementing a well-designed SIEM in any environment that provides essential services over the Internet or any network, regardless of its size. It also explores integration with state-of-the-art technologies and features a lab demonstration showcasing real attacks, detections, and prevention mechanisms.

1.2 Methodology

The context of the thesis focuses on the implementation of a SIEM system alongside a Network IDS, demonstrating their critical role in cybersecurity. A sample network of various devices is introduced for this purpose. In addition, a centralized server hosting the components of the Elastic Stack (formerly Elastic Logstash Kibana – ELK Stack) acts as the SIEM. Furthermore, another server hosting the NIDS, also serves as the central control point for various IoT devices (light bulbs, dehumidifiers, air conditioners, etc.) within the network.

This study conducts experiments on the primary types of cyber-attacks, demonstrating effective detection and prevention strategies. Additionally, the methodology employed, and the findings derived from these experiments are thoroughly examined.

1.3 Thesis' structure

This thesis aims to underscore the pivotal role of modern technologies, such as the Elastic Stack, in conjunction with state-of-the-art Network Security Monitoring (NSM) and host-based security tools, in the ongoing battle against cybercrime. To establish a comprehensive foundation, an in-depth analysis of both human and technological factors is provided, along with an overview of existing cybersecurity solutions. Prior to conducting the experiments, a detailed presentation of the system architecture, including hardware and software components, is illustrated through diagrams and listings. Subsequently, various cyber-attacks are executed and thoroughly documented with screenshots and explanations. Finally, the results are analyzed, with a focus on detection and prevention strategies, culminating in a discussion on the future of cybersecurity, which is inextricably linked to advancements in Machine Learning.

2. Background

2.1 IoT

Some years ago, we couldn't even imagine that one day we would be able to control our air conditioning unit or move our car from our smartphone. The Internet of Things (IoT) is a network of physical devices, vehicles, home appliances, and other items embedded with electronics, software, sensors, actuators, and connectivity, enabling these objects to connect and exchange data [4]. As we understand, the more IoT devices that connect to the internet, the greater the risk of security breaches and cyber threats. Each additional connected device expands the attack surface, making it more vulnerable to potential exploitation by malicious actors. Therefore, ensuring robust cybersecurity measures and implementing strong encryption protocols is crucial to mitigating these risks. The issue, however, lies in the fact that the majority of these IoT devices lack inherent security mechanisms, leaving them susceptible to cyber threats. Consequently, they must rely on external network devices, such as firewalls, to provide the necessary layers of protection. This dependency underscores the importance of a comprehensive security strategy that includes network segmentation, continuous monitoring, and timely updates to safeguard these vulnerable endpoints against potential exploitation.

2.2 SOC

A Security Operations Center (SOC) is an in-house or outsourced team of IT security professionals dedicated to monitoring an organization's entire IT infrastructure 24x7. Its mission is to detect, analyze and respond to security incidents in real-time. The SOC also selects, operates and maintains the organization's cybersecurity technologies and continually analyzes threat data to find ways to improve the organization's security posture [5].

The SOC comprises a diverse team of cybersecurity professionals, including cybersecurity analysts, threat hunters, and members of the Computer Security Incident Response Team (CSIRT), among others. These experts collaborate seamlessly to detect,

analyze, and mitigate cyber threats, ensuring the organization's digital infrastructure remains resilient against cybercrime. Through continuous monitoring, threat intelligence integration, and incident response coordination, the SOC plays a critical role in safeguarding sensitive data and maintaining the overall security posture of the organization.

2.3 SOC tools

2.3.1 Firewall

A firewall is a network security device that separates a trusted internal network from an external network deemed untrustworthy, such as the internet. It regulates incoming and outgoing network traffic based on preset security rules. Firewalls are paramount in shielding networks from unauthorized access, harmful activities, and potential threats, and can exist as hardware, software, software-as-a-service (SaaS), or public or private (virtual) cloud.

Firewalls scrutinize network packets and implement security policies, effectively barring unauthorized users or potentially harmful data from infiltrating or exiting a network. Notably, firewalls serve as gatekeepers, scrutinizing each network packet and deciding whether to permit or block it based on pre-set rules. This helps to ensure that only traffic deemed safe and legitimate is allowed through the firewall.

In addition to these core functions, today's next-generation firewalls (NGFWs) have a range of features to bolster network security. These include deep packet inspection, application visibility and control, intrusion detection and prevention, malware defense, URL filtering, and more [6].

2.3.2 IDS and IPS

An Intrusion Detection System (IDS) is a network security technology originally built for detecting vulnerability exploits against a target application or computer. The IDS is a listen-only device. The IDS acts like a sensor on the network which monitors traffic

and reports results to an administrator. It cannot automatically take action to prevent a detected exploit from taking over the system [7].

On the contrary, an Intrusion Prevention System (IPS) operates as an inline network security device, meaning it must be strategically positioned between the protected internal network and the untrusted external network. This placement enables the IPS to actively monitor traffic in real time and take immediate action against malicious activities based on predefined security policies and administrator configured rules. By automatically blocking, mitigating, or redirecting potential threats, the IPS serves as a proactive defense mechanism, strengthening the organization's overall cybersecurity posture.

Both IDS and IPS can be categorized as either network-based or host-based, depending on their deployment location. Furthermore, they may be classified based on their detection methodologies, including signature-based, rule-based, behavior-based, event-based, or flow-based approaches. Additionally, developers of each system may introduce more specialized naming conventions to reflect unique functionalities or proprietary enhancements.

Rule-based and signature-based IDS/IPS (like Snort and Suricata) can work best for already known threats. However, when it comes to zero-day attacks, they are not very useful. Behavior-based and flow-based IDSs/IPSs (like Cisco Stealthwatch) use machine learning and flow patterns which eventually lead to the determination that an attack is happening. Event-based tools (like Zeek) correlate the information gathered and produce an alert if a decision that an attack occurred is made.

Determining where an IDS/IPS should be located needs to be defined according to the situation. For network IDS/IPS to function properly, all traffic flowing on the network need to pass through it. If prevention is desired, then it is mandatory to place the IPS inline, meaning between the protected network and the untrusted one. If only detection is needed, then the IDS can be placed anywhere in the network as long as all traffic is mirrored to it through a Network Test Access Point (more known as Network TAP) or via the configuration of a Switch Port Analyzer (more known as SPAN port) at a network switch supporting this feature. Having a network IDS/IPS is a fine solution, especially if it is possible to perform Deep Packet Inspection (DPI), although there will still be some security gaps.

Host based IDS/IPS is software performing detection and prevention directly on a single host such as a computer, a server, a smartphone or another endpoint device. Although a network IDS/IPS monitors traffic in real-time, a host-based one can inspect historical data in logged files for system anomalies. Additionally, these systems have the capability to analyze encrypted data that has successfully traversed the network and will be decrypted directly on the host system. This enables them to detect potential security threats that may be concealed within encrypted communications, enhancing overall network protection against sophisticated cyber threats. Furthermore, they play a crucial role in identifying and mitigating malicious files at the endpoint level. By continuously monitoring system activity, file integrity, and behavioral patterns, these systems can detect unauthorized modifications, suspicious file executions, and anomalous processes that may indicate malware infection. HIPS solutions go a step further by proactively blocking or quarantining malicious files before they can cause harm, leveraging heuristics, behavioral analysis, and predefined security policies. This level of protection is particularly valuable against zero-day threats and advanced persistent threats (APTs), as it provides an additional security layer beyond traditional signature-based antivirus solutions.

The following table presents a selection of state-of-the-art Intrusion Detection and Prevention Systems (IDS/IPS), outlining their key features and functionalities through a comprehensive evaluation.

Feature	Suricata	Snort	Zeek
Type	IDS/IPS	IDS/IPS	IDS
License	Open-source	Open-source	Open-source
Detection Method	Signature and Rule based	Signature and Rule based	Event-based
Real-time Prevention	Yes	Yes	No
Multi-Threading Support	Yes	No	Yes
Rule Set Compatibility	Snort and Suricata rules	Snort rules only	Custom scripts
File Extraction	Yes	Limited (via third-party tools)	Yes
Logging Format	JSON, PCAP, CSV	ASCII, XML, JSON	JSON, ASCII
Event Logging	Real-time	Real-time	Stores logs for extracting insights
Ease of Setup	Moderate	Moderate	High
Customizability	High	High	Very High (via scripting)
Use Case Focus	Real-time threat detection and prevention	Real-time threat detection and prevention	Deep network traffic analysis, forensics
Alerts and Notifications	Yes (real-time alerting)	Yes (real-time alerting)	Yes (but not in real-time for prevention)
Deep Packet Inspection	Yes	Yes	Yes
Community and Support	Active community, frequent updates	Large community, well-established	Large, active community, frequent updates
Suitability	Best for real-time blocking	Best for IDS and threat detection	Best for network monitoring and forensics

Figure 2-1. IDS/IPS comparison

2.3.3 SIEM

Security Information and Event Management (SIEM) is a comprehensive security management approach that integrates the functionalities of Security Information Management (SIM) and Security Event Management (SEM) into a unified system. SIEM solutions serve as a critical component of modern cybersecurity frameworks by enabling the collection, aggregation, normalization, and correlation of vast amounts of log data and security events generated across an organization's IT infrastructure. By providing real-time monitoring, advanced analytics and threat detection capabilities, SIEM enhances the efficiency and effectiveness of Security Operations Centers (SOCs), allowing security teams to swiftly identify, investigate and respond to potential threats.

With the average organization’s Security Operations Center (SOC) receiving more than 10,000 alerts per day, and the biggest enterprises seeing over 150,000, it is impossible to have security teams large enough to keep up with the overwhelming number of alerts. A SIEM system collects data from multiple sources, enabling faster incident response to threats. If an anomaly is detected, it might collect more information, trigger an alert, or quarantine an asset [8].

Traditionally, SIEM was primarily adopted by enterprises and public sector organizations to meet regulatory compliance requirements. However, over time, its role has expanded significantly as organizations have recognized its broader security benefits. SIEM solutions have evolved into essential tools for threat detection, incident response, and real-time security monitoring, making them indispensable for organizations of all sizes. As cyber threats become increasingly sophisticated and the shortage of skilled cybersecurity professionals persists, the ability to rapidly and automatically detect security breaches has become more critical than ever. This growing necessity has led to wider adoption of SIEM solutions, not only by large enterprises but also by small and medium-sized businesses seeking to strengthen their cybersecurity posture.

SIEM technology collects security-related data from servers, user devices, network equipment, applications, and security tools. It can also ingest and interpret logs from firewalls, honeypots, databases, and IoT hubs. It categorizes and normalizes this data to standard formats and, when a potential threat is detected, it can generate alerts or take automated actions based on predefined security policies. By aggregating and analyzing data across the entire network, SIEM provides early detection of security

incidents and a faster response to mitigate risks before they escalate. The data collected by a SIEM system is presented through comprehensive dashboards, which allow cybersecurity analysts to visualize and analyze security events. These dashboards facilitate the extraction of key insights, enabling analysts to draw conclusions and generate statistics.

There is a wide array of (SIEM) software solutions, with the features and capabilities of each tool largely dependent on the respective manufacturer. Among the most highly regarded and widely adopted SIEM platforms are Splunk, IBM's QRadar, and FortiSIEM, each known for their robust functionality and enterprise-level security management capabilities. Additionally, the Elastic Stack, an esteemed open-source platform (open source in the core components), has garnered significant recognition in the industry. Formerly known as the ELK Stack (Elasticsearch, Logstash, and Kibana), the Elastic Stack has been a trusted and comprehensive tool in the field for several years, providing a powerful platform for managing and analyzing enormous amounts of security data.

Elastic Stack is a highly scalable and modular framework for ingesting, analyzing, storing and visualizing data. It is used for searching and analyzing an organization's data in near real time and can be configured to send alerts once a rule or a pattern of rules is met [9]. Its main components are Elasticsearch, Kibana and the Elastic Agent and together with some more important ones are briefly presented below:

- Elasticsearch

It is the main search engine component which uses RESTful web services and APIs, a distributed computing cluster with multiple server nodes, and a distributed NoSQL database made up of JSON documents. Additional functionality can be added through custom-created extensions [10].

- Kibana

Kibana is a data visualization and exploration tool that works seamlessly with Elasticsearch. It enables users to interact with large volumes of data, creating dashboards, visualizations, and reports. Primarily used for monitoring, security analytics, and log management, Kibana provides insights through real-time data representation. Other

network and security tools like Wazuh can also be integrated into Kibana for data visualization [10].

- Elastic Agent

It is a lightweight data collector that can be easily installed on various platforms, including Windows, macOS, and Linux. The installation process is both simple and efficient, requiring no client-side configuration. It can be deployed in seconds by simply pasting the straightforward commands provided by Kibana to the client, ensuring seamless integration across different environments. The Elastic Agent simplifies deployment and management by handling multiple data collection tasks—such as log shipping, metrics gathering, and security monitoring—through a single, streamlined solution. [10].

- Logstash

While it is not strictly necessary for the Elastic Stack to be fully operational, it is nonetheless an invaluable tool for processing and transforming log data. Logstash serves as a robust data pipeline, enabling seamless ingestion, parsing, and forwarding of diverse log data from various sources to centralized storage or analysis platforms [10]

- Beats

Beats are lightweight, single-purpose data shippers designed to collect specific types of data from various sources and send it to the Elastic Stack for analysis. There are several different Beats, like FileBeat, MetricBeat, PacketBeat, HeartBeat and each one is optimized for a particular data type. However, in recent developments, these individual Beats are gradually being replaced by the Elastic Agent, which consolidates their functionality into a unified agent [9].

- Various integrations

Elastic integrations are pre-configured modules that simplify data collection from a variety of sources. For example, the NGINX integration provides a streamlined way to collect and parse NGINX access and error logs for analysis, helping analysts monitor web traffic and server performance. Similarly, the Snort integration facilitates the collection

of security event data from Snort IDS. These integrations automate data ingestion and processing, reducing the need for complex configuration [10].

- Machine Learning

Elastic's machine learning tools bolster threat hunting by automatically detecting unusual activities, with preconfigured models for observability and security. Administrators can implement unsupervised learning for anomaly detection and swiftly visualize data. Elastic's advanced vector and semantic search capabilities deliver superior relevance through automated correlations and efficient root cause analysis. In this manner, the Elastic Stack can autonomously generate new rules based on the logs and data it collects from devices across the network. It continuously compares this data against established security databases, such as MITRE ATT&CK, VirusTotal, and the CVE database, to detect potential intrusions or malware. This automated process enables real-time alerts for emerging threats, enhancing overall security posture without requiring manual intervention [11].

2.3.4 EDR/XDR

Endpoint Detection and Response (EDR) is an advanced cybersecurity solution designed to monitor, analyze, and protect endpoints by detecting behavioral anomalies and potential security threats. By continuously collecting and scrutinizing data from various endpoints, EDR identifies suspicious activities that may compromise the integrity of individual devices or interconnected systems. This proactive approach enables security teams to swiftly respond to emerging threats, neutralize malicious activity, and efficiently manage alerts, ensuring a more robust defense against cyberattacks.

Unlike EDR, Extended Detection and Response (XDR) is a comprehensive cybersecurity solution that integrates and correlates threat data across multiple security layers, including endpoints, networks, cloud environments, and email systems. XDR provides a unified view of an organization's threat landscape, enabling faster detection, investigation, and remediation of sophisticated attacks. By leveraging advanced analytics, machine learning, and automation, XDR enhances threat visibility, reduces response times, and improves overall security efficiency, making it a powerful tool for modern

threat defense. It can be integrated with Cyber Threat Intelligence (CTI), thus enhancing its ability to identify emerging threats, correlate attack patterns, and proactively defend against zero-day attacks. Additionally, XDR leverages automated incident response capabilities, enabling real-time threat containment, remediation, and policy enforcement with minimal human intervention.

2.4 Incident Response

2.4.1 Cyber Kill Chain

The National Institute of Standards and Technology (NIST) states that a security incident can be described as "a violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices." Essentially, a security incident refers to any event or activity that breaches or threatens to breach an organization's security policies, procedures, or infrastructure.

Developed by Lockheed Martin, the Cyber Kill Chain is a strategic framework designed to identify, mitigate, and prevent cyber intrusions. The chain is comprised of seven distinct stages and each of them assists analysts in understanding the Tactics, Techniques, and Procedures (TTPs) employed by adversaries during an attack. The primary objective of the Cyber Kill Chain is to disrupt the attacker's progress at any stage of the chain, thereby minimizing the potential damage. The earlier an adversary is stopped in the attack process, the less impact it has on the organization, reducing the risk of data loss or compromise of sensitive information [9]. The seven steps in the Cyber Kill Chain are as follows:

- *Reconnaissance*: researching target information
- *Weaponization*: pairing remote malware with backdoor to create payload
- *Delivery*: delivering a weapon to the victim through any means
- *Exploitation*: executing code on vulnerable system by triggering it
- *Installation*: installing backdoor on the target system
- *Command and Control (CnC)*: creating an outside server as a command channel to manipulate the target
- *Action on Objectives*: using hands-on keyboard access to achieve the objective of the attack

2.4.2 Incident Response Life Cycle

NIST defines incident response as an ongoing and dynamic process that spans the entire lifecycle of a cybersecurity incident, encompassing activities conducted before, during, and after the event [12]. This process is organized into four distinct stages, each focusing on specific tasks, with detection and containment serving as important reference points in the overall response framework. Various specialized teams of a SOC may be assigned to manage different stages of the incident response. Furthermore, this structured methodology not only facilitates an organized response to incidents but also supports regulatory compliance, helping organizations meet industry standards and legal requirements. Examples of regulatory compliance requirements that organizations may need to meet during incident response include the General Data Protection Regulation (GDPR), the Health Insurance Portability and Accountability Act (HIPAA), and the Payment Card Industry Data Security Standard (PCI DSS). Additional regulations may apply depending on the organization's industry, geographical location, and scope of operations. The four stages of Incident Response Life Cycle are depicted on the next figure.

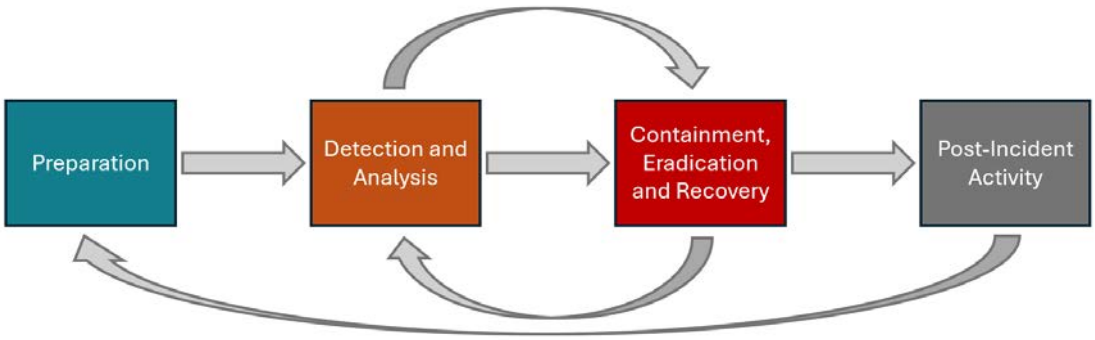

Figure 2-2. Incident Response Lifecycle

3. Methodology

3.1 System Architecture

In contemporary organizational structures, the majority of companies rely on a centralized administrative hub, commonly referred to as the headquarters (HQ). This HQ serves as the primary point of connection for remote branches and employees working from various locations, including those operating with personal laptops. It also provides access to other entities, like for example suppliers and clients, who need to communicate securely between them. This is achieved with the use of Virtual Private Networks (VPNs).

This thesis will introduce a comparable architectural framework. As illustrated in Figure 3.1, it is assumed that a SIEM & XDR Server is deployed at the organization's HQ. A remote office is securely connected to the HQ via VPN. This example remote office accommodates various IoT devices alongside multiple endpoints utilized by remote employees. Furthermore, the remote office houses a compact server that functions as both an IDS/IPS and a Web Server. However, its primary role is monitoring the IoT devices connected to the same LAN as well as controlling the office's overall network traffic via the Snort sensor. This functionality is essential, as most modern IoT devices are headless—lacking built-in administrative interfaces or management modules—necessitating an external system for administration and security.

As previously stated, a SIEM system collects logs and data from multiple sources. To achieve this, a specialized software component must be installed on the hosts that require monitoring. In this case, the software used is Elastic Agent, a lightweight program capable of running on various platforms, including Windows, Linux, and macOS, with a simple one-command installation process. However, this deployment is only feasible if the host allows interaction, which underscores the necessity of an IoT server for managing and interfacing with otherwise inaccessible IoT devices.

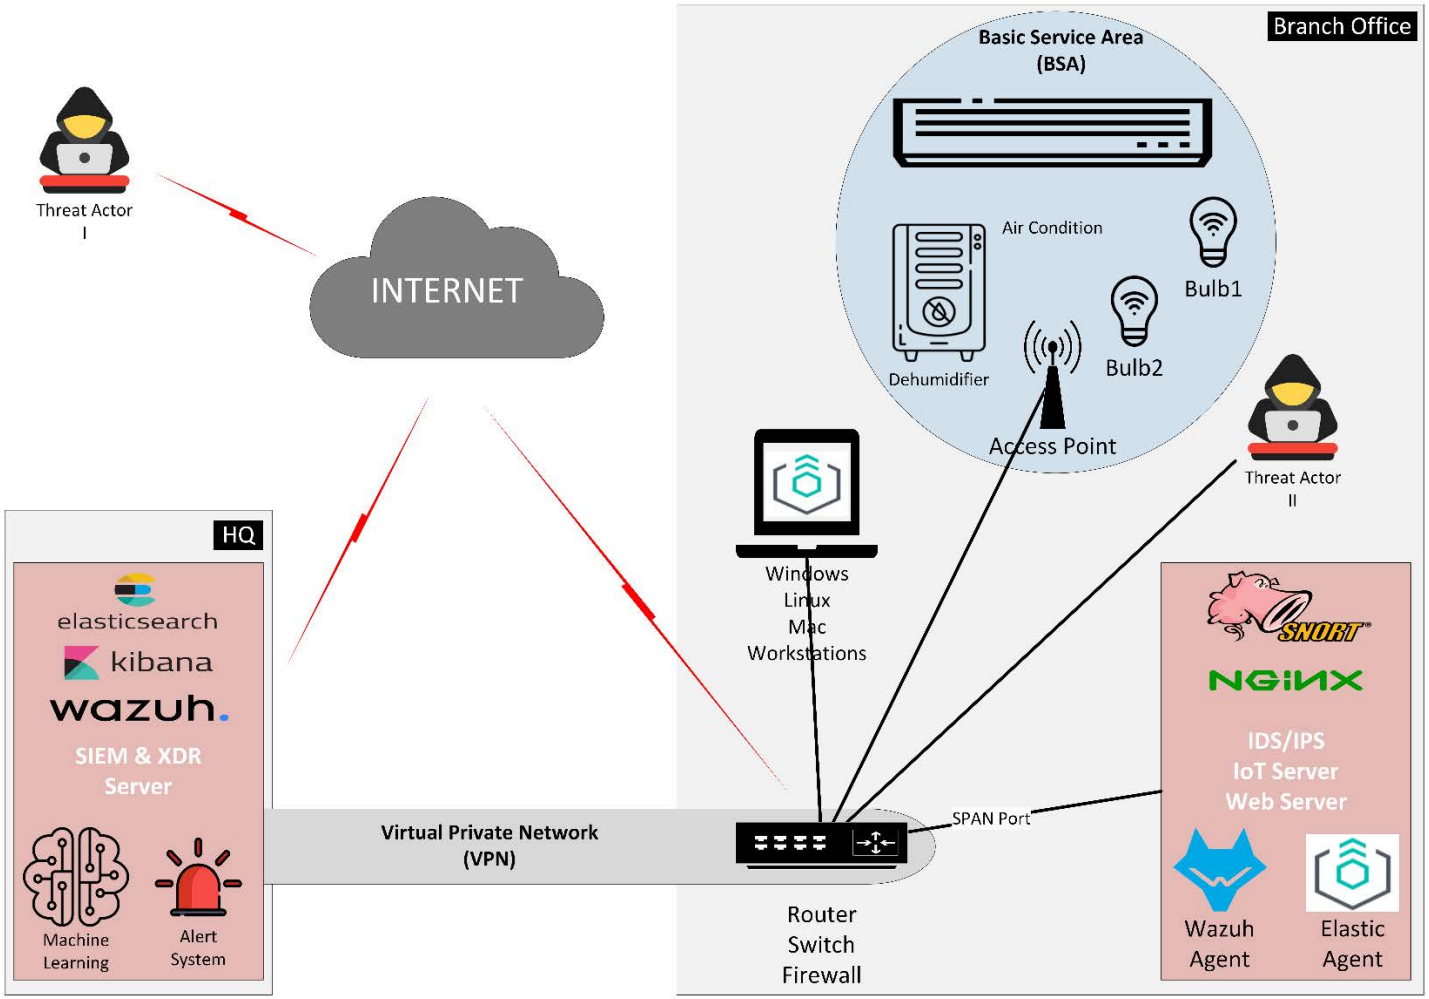

Figure 3-1. Topology

3.2 Hardware used

The hardware utilized in this thesis was selected based on the availability of existing components, while consistently prioritizing a cost-effective approach. A detailed analysis of the hardware components depicted in Figure 3.1 is provided below:

- ElasticSearch / Kibana / Wazuh Server

This Virtual Machine (VM) is hosted on a high-performance server located on premises. It has been allocated a 3-core CPU, 8GB of RAM, and 128GB of storage. While these specifications fall below the minimum requirements for running ElasticSearch, Kibana, and Wazuh concurrently in a production environment, they proved sufficient for conducting the experiments presented in the following chapter.

- IDS / IPS / IoT / Web Server

This is a Raspberry Pi 5 [13] configured with 8GB of RAM. Its performance was more than enough for the assigned tasks, demonstrating its suitability for deployment in a remote office environment for this type of application.

- Router / Switch / Firewall / Access Point

This all-in-one networking solution is a MikroTik hAP device [14] performing routing and switching for the whole remote office and also acting as a firewall and access point (AP) for wireless devices. Additionally, this device is responsible for establishing the VPN connection with the HQ. However, due to the encryption overhead associated with VPN traffic, it is not suitable for production use in remote branches where higher performance and scalability are required.

- Endpoints

Throughout the entire topology, various endpoints, including Linux and Windows computers, Voice over Internet Protocol (VoIP) phones, IoT devices, smartphones, printers, and more, may be present at either the HQ or remote offices. All these endpoints can be monitored using a range of cybersecurity tools, ensuring comprehensive protection

against scans, attacks, and intrusions. The primary requirement for monitoring is the installation of Elastic Agent; although, in cases where this is not feasible, alternative methods of monitoring their network traffic can be employed.

3.3 Software used

All software used is state of the art at the time of writing. The selection of the software was made with a primary focus on open-source solutions. Cybersecurity has always had a strong inclination toward open-source solutions anyway. This is particularly important because open-source software offers greater flexibility, transparency, and control over the code, enabling customization to fit specific needs. Moreover, open-source tools benefit from community-driven development, providing regular updates, security patches, and a wide range of support options. This approach also ensures cost-effectiveness, as it eliminates licensing fees while fostering collaboration and innovation within the cybersecurity field. The software used is analyzed below:

- Elastic Stack

The most critical piece of software for this case study is the Elastic Stack. The components utilized in this instance include Elasticsearch, Kibana and the Elastic Agent. Elasticsearch functions as both the database and the indexer, storing all logs and data while enabling fast and efficient searching of vital cybersecurity information. This capability is essential for analysts to retrieve the necessary data promptly in order to carry out their tasks effectively.

On the other hand, Kibana serves as the visualization application, a crucial tool for the administration and management of the entire system. Through Kibana, users can create alerts, define rules, and even configure Machine Learning models to autonomously make decisions and take actions when required. This includes tasks such as the creation of new rules and alerts, identification of attack patterns, and other cybersecurity measures. Analysts can also configure Kibana to alert them to specific email addresses or even messaging applications. Some sample images of the Kibana interface are shown in the next two figures.

The Elastic Agent serves as a single agent that consolidates the functionality of several separate Beats -such as Filebeat, Metricbeat, and Packetbeat- that were previously used for specific data collection tasks. By replacing these individual Beats, the Elastic Agent simplifies deployment and management, making it easier to monitor endpoints, servers, and cloud environments. As it evolves, the Elastic Agent continues to expand its capabilities, integrating new features and modules to become a fully featured software solution that can be installed on every host in an organization.

Although the Elastic Stack was originally a purely open-source project, recent years have seen the introduction of features such as Machine Learning and Automated Active Response—features that are no longer open source. The inclusion of Automated Active Response, in particular, positions the Elastic Stack as both a SOAR and XDR solution. Due to the licensing changes and my desire to implement automated threat responses, I was forced to incorporate another highly effective, entirely open-source SIEM and XDR tool, Wazuh, alongside the Elastic Stack. In this setup, only the XDR functionality of Wazuh is utilized to complement the existing system.

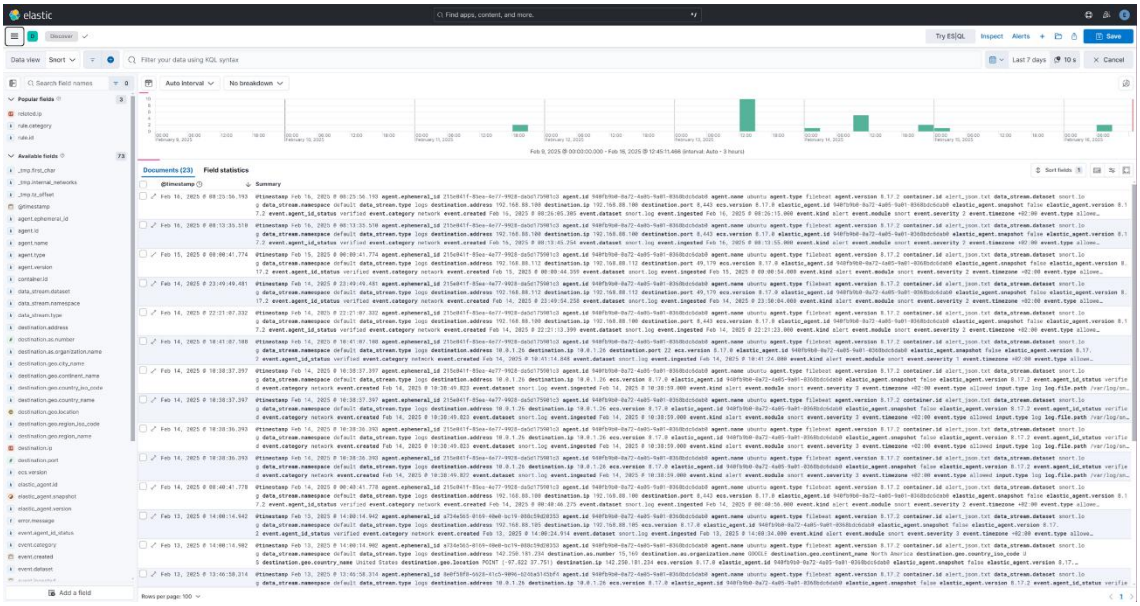

Figure 3-2. Kibana Discovery Tab

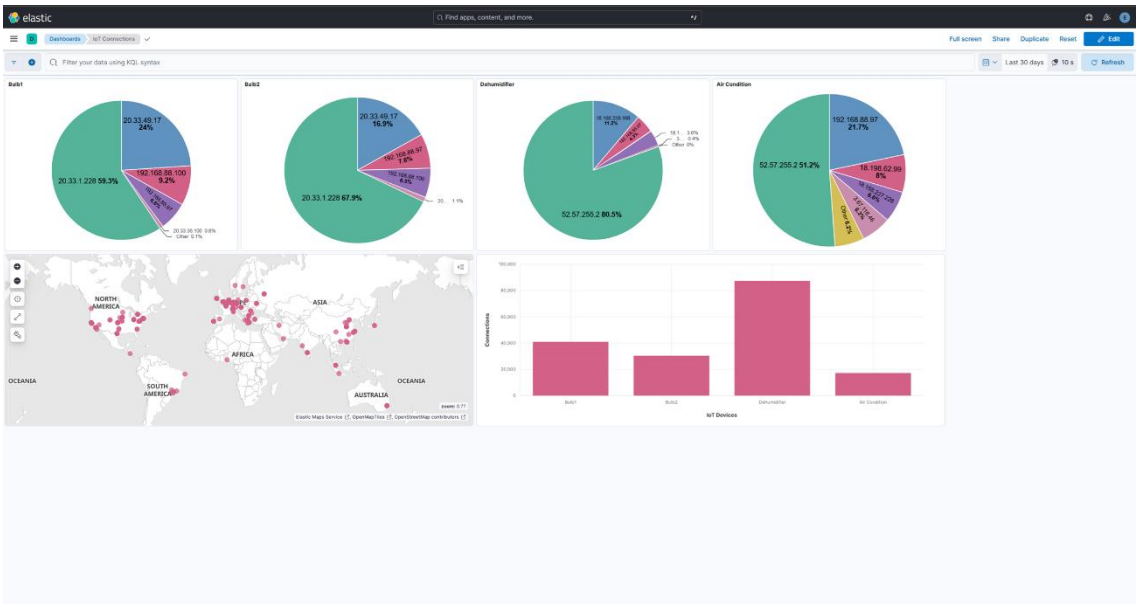

Figure 3-3. Sample Kibana Dashboard

- Wazuh

Wazuh originated as a fork of the Open-Source Security Event Correlator (OSSEC) project, a host-based intrusion detection system (HIDS) developed in 2004. Currently, it is a fully independent open-source security platform that serves as both a SIEM and XDR solution. While it shares several features with the Elastic Stack, its primary focus is on cybersecurity, threat detection, and incident response. Initially, Wazuh functioned as a plugin and integration within the Elastic Stack. However, due to licensing changes introduced by Elastic and a strategic decision to achieve greater technological independence, the Wazuh development team introduced the Wazuh Indexer, an alternative to Elasticsearch based on the OpenSearch project, and the Wazuh Dashboard, a visualization tool similar to Kibana. These enhancements have allowed Wazuh to maintain its open-source nature while improving performance, scalability, and security capabilities.

Wazuh is centrally administered by the Wazuh Manager, which serves as the core component responsible for managing and processing security events. The deployment of a lightweight software component, known as the Wazuh Agent, is mandatory on every host that requires monitoring. Its functionality closely resembles that of the Elastic Agent mentioned before, as it operates with root privileges on the host system to perform malware scanning, network traffic monitoring, and real-time alert transmission to the Wazuh Manager. The Wazuh Manager, functioning as the central administration server,

manages all agents by distributing configuration files that dictate their operational behavior. This enables the manager to enforce security policies remotely. For instance, if a specified number of failed authentication attempts occur, the manager can instruct the agent to disable SSH access on the host for a predefined duration, mitigating potential brute-force attacks.

The key reason for introducing Wazuh in this thesis, rather than solely relying on the Elastic Stack for all security operations, lies in its ability to provide essential security functionalities as part of an open-source platform, without the need for costly licensing. While the Elastic Stack offers similar capabilities, these features are restricted to paid licenses, making them inaccessible for users seeking a completely free, open-source solution. Wazuh, on the other hand, offers these critical features without additional costs, making it a more cost-effective and accessible alternative for Incident Response and Threat Hunting. However, for larger organizations with a vast number of systems to protect, the cost of a paid license is often negligible compared to the value of the advanced features and support it provides.

In this study, only the Wazuh Manager/Server and Wazuh Agent will be utilized, as these are the core components essential for Incident Response. The other components, namely the Wazuh Indexer and Wazuh Dashboard, are intentionally excluded from this case, as their functionalities are already effectively provided by more robust and feature-rich tools, such as Elasticsearch and Kibana.

It is worth noting that Wazuh has an extensive rule set, and most of the rules used are mapped to the MITRE ATT&CK framework, which provides valuable context by categorizing detected events based on specific adversary Tactics, Techniques, and Procedures (TTPs). The MITRE ATT&CK framework is a comprehensive knowledge base that catalogs TTPs used by adversaries in cyberattacks, helping organizations identify and defend against potential threats. By leveraging the MITRE framework, Wazuh enhances the ability to correlate alerts with known attack patterns and offers greater visibility of the threats, since the system's rules are continuously updated to incorporate new attack techniques.

- RouterOS

This is the proprietary software running on every MikroTik device. Additionally, there is SwitchOS, a more streamlined version of RouterOS, specifically designed for

devices with switching-only capabilities. RouterOS is a fully featured operating system that provides an extensive array of tools and functionalities for routing, firewall management, bandwidth control, VPN services, and network monitoring, making it suitable for a wide range of networking tasks. It supports a vast array of hardware interfaces and protocols, allowing users to configure their devices for various use cases, from small home networks to large enterprise environments. RouterOS also includes advanced features such as Quality of Service (QoS), dynamic routing protocols (OSPF, BGP), and VPN tunneling, providing administrators with the flexibility to design highly customizable and secure network infrastructures. It even allows for running custom scripts, transforming it to a fully customizable linux-based solution.

In this study case, RouterOS is one of the most suitable choices, as it offers a range of advanced features, including Port Mirroring, which is a critical capability for effective network monitoring. Port Mirroring technology enables the configuration of a source port to send a copy of all its traffic to a designated destination port, known as the mirror port. This feature is essential for deploying network monitoring devices, such as an IDS or an IPS, which can then analyze the mirrored traffic in real-time. A screenshot of the MikroTik administration software, Winbox, is shown in Figure 3.4.

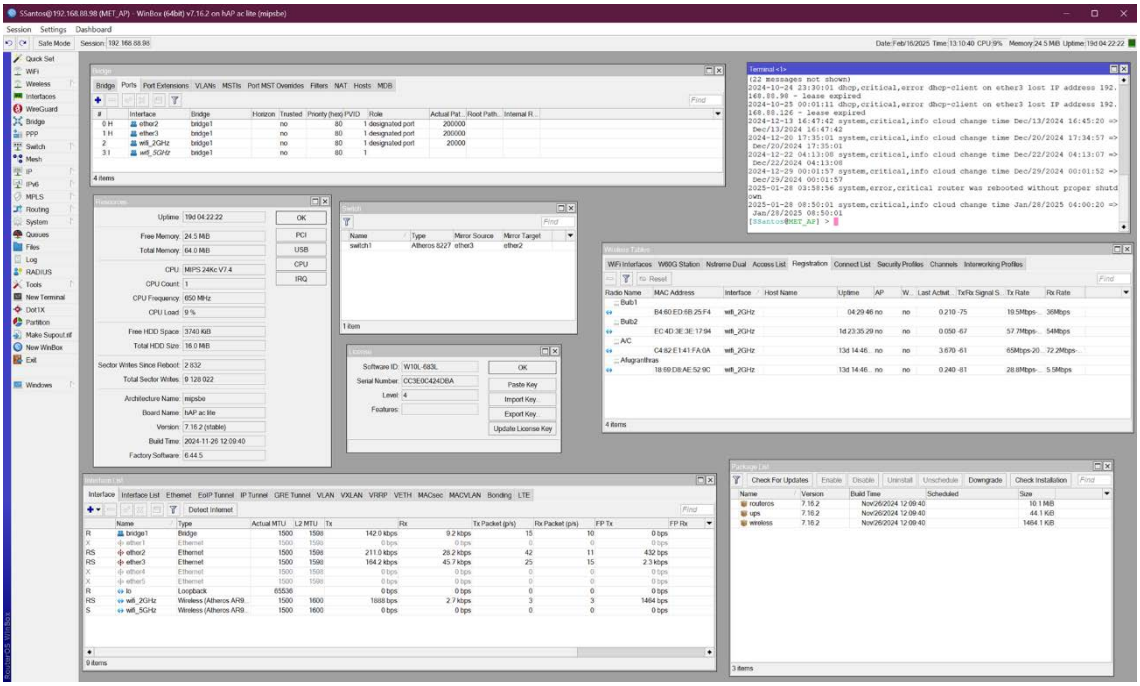

Figure 3-4. Winbox Interface

- Snort

The IDS/IPS of choice for this thesis is Snort. Snort analyzes network traffic in real-time and its main goal is to identify and block malicious activity. It can detect a wide range of threats, including port scans, buffer overflows, and denial-of-service attacks, using predefined rules and custom configurations. Snort operates by examining packet data and comparing it against its rule sets, triggering alerts when suspicious patterns are detected. Due to its flexibility and extensibility, Snort is widely used in both enterprise and small environments for enhancing network security and providing deep insights into potential threats. To block malicious traffic in real-time, Snort must be placed inline with a network device, such as a switch or router. In large deployments, it may be necessary to implement two Snort devices, one positioned on the internal network (LAN side) and another placed in the Demilitarized Zone (DMZ), to effectively monitor and filter both the traffic on the internal network and the traffic entering or exiting it.

The Snort rules utilized in this deployment are the Snort Community Rules [15]. Those are a collection of predefined detection rules contributed and maintained by the global Snort community. These rules are designed to identify a wide range of network threats, from known attack patterns to new vulnerabilities. Community rules are freely available and provide a comprehensive set of signatures for intrusion detection and prevention. They are continuously updated almost every day by both the Snort team and community members, ensuring that the rule set evolves in response to emerging security threats. Some of them are even mapped to the MITRE ATT&CK framework. Other rule sets such as the Emerging Threats (ET) can also be added. Additionally, it allows for custom rule writing and provides documentation for it [16].

3.4 System Functionality

The threat actors targeting our networks may originate from either internal or external sources. While internal threats are relatively uncommon due to the necessity of physical wired or wireless access, numerous past incidents have demonstrated that ordinary employees can unknowingly harbor malware on their endpoints. This, in turn, enables an ostensibly external threat actor to operate as an internal one using remote access tools or other forms of malicious software.

The flow of traffic and the logical interconnections of the systems presented in this thesis are illustrated in the next figure.

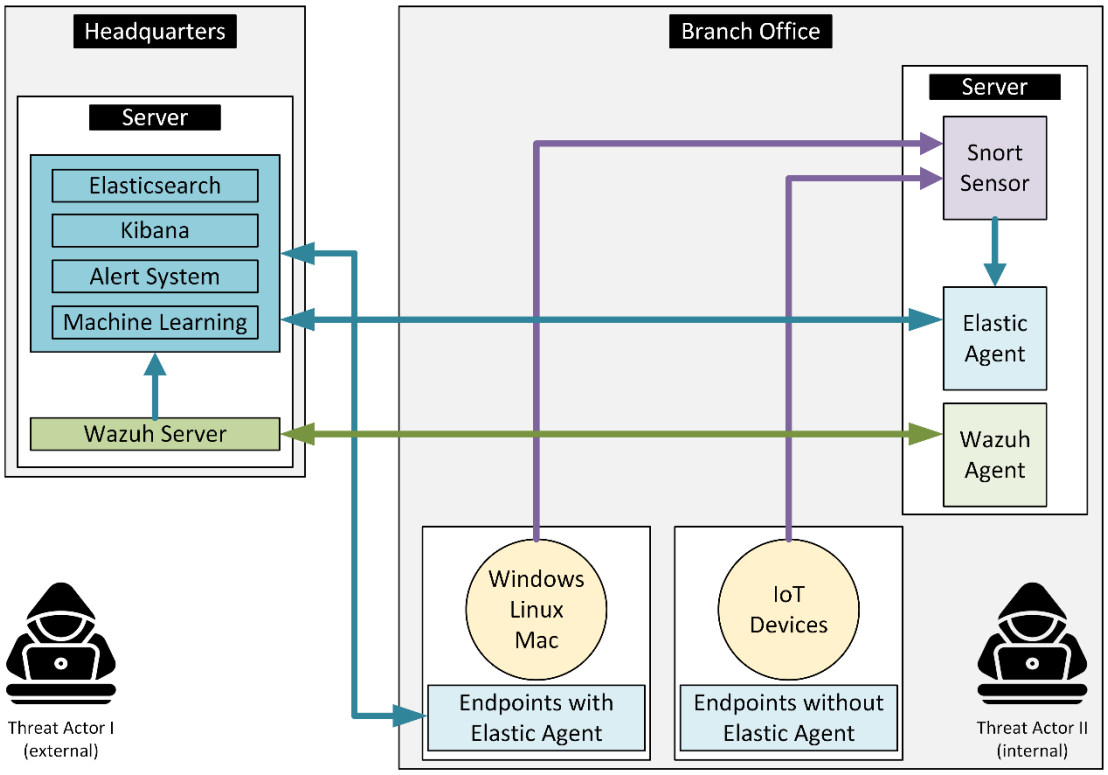

Figure 3-5. Data Flow

As depicted above, the various workstations in the organization's Branch Office should have the Elastic Agent installed to ensure that all activities on these devices are forwarded to the HQ server. This can be easily configured centrally, for example, as an Active Directory policy, eliminating the need for the administrator to install it manually on each workstation. The Elastic Agent should also run on the office's mini server to forward all its events, as well as all Snort alerts, to the HQ server.

On the HQ side, a variety of integrations are incorporated into each Elastic Agent policy. For example, all Windows endpoints can be assigned to a unified Elastic Agent policy, customized for the unique characteristics of each system category. All configurations are centralized through the Kibana interface or directly within Elastic's configuration files, ensuring that any changes are automatically applied to Elastic Agents. Moreover, within Kibana, users have the flexibility to filter and analyze specific types of data sent from Elastic Agents or other sources, enhance the data as needed, and construct

detailed visualizations and dashboards. These visualizations can be filtered according to time, system parameters, or other variables. Kibana allows users to configure alerts, enabling automatic notifications when certain criteria are met, or predefined thresholds are exceeded. In this use case, the integration of MITRE ATT&CK correlation is implemented on the HQ server, meaning that all data ingested into Elasticsearch is cross-referenced with the MITRE framework. This correlation enables the identification and mapping of potential Tactics, Techniques, and Procedures (TTPs) used by threat actors, enhancing the overall detection and response capabilities.

A Wazuh Agent runs on the office's server, serving exclusively for Active Response purposes. The Wazuh Agent maintains a live connection with the Wazuh Server, allowing the server to send policies, rules, and automations to the agent. All agents have root privileges on the systems where they are installed, which is essential for accessing necessary host paths and for controlling host configurations when required for Active Response rules and other security measures.

Since the office's server is connected to a Switched Port Analyzer (SPAN) Port, all inbound and outbound traffic is mirrored to it. There, the Snort Sensor can read traversing packets from the whole network -including the IoT devices- compare them to preconfigured rules and selectively send alerts to Kibana via the Elastic Agent.

In the following chapter, a series of cyberattacks are simulated, encompassing scenarios in which both the external and the internal adversaries target various assets, including workstations, IoT devices, and even the office's server.

4. Experiments

The attacks demonstrated in this study were conducted in a controlled and secure environment, completely isolated from any critical networks and the Internet. These experiments were carried out with caution, ensuring that all necessary precautions and security measures were in place to prevent unintended propagation or any adverse consequences. It is strongly advised that such tests should only be performed by professionals with comprehensive expertise in the field, as improper execution could lead to serious risks. The sole objective of these demonstrations is to analyze and illustrate how specific hardware and software react and/or respond to various attack scenarios.

The following state diagram provides a comprehensive illustration of the processes involved in each attack scenario that will be discussed later. In this context, IoT devices, endpoints, and the office server are identified as vulnerable machines susceptible to both internal and external attacks. This state diagram serves as a valuable tool for understanding the flow of various types of data throughout the attack lifecycle. It analyzes the pathways that data traverses while highlighting the mechanisms of detection and prevention employed at each stage. In essence, this diagram not only underscores the attack vectors but also illustrates the layered defenses and response strategies that are critical in safeguarding against malicious activities.

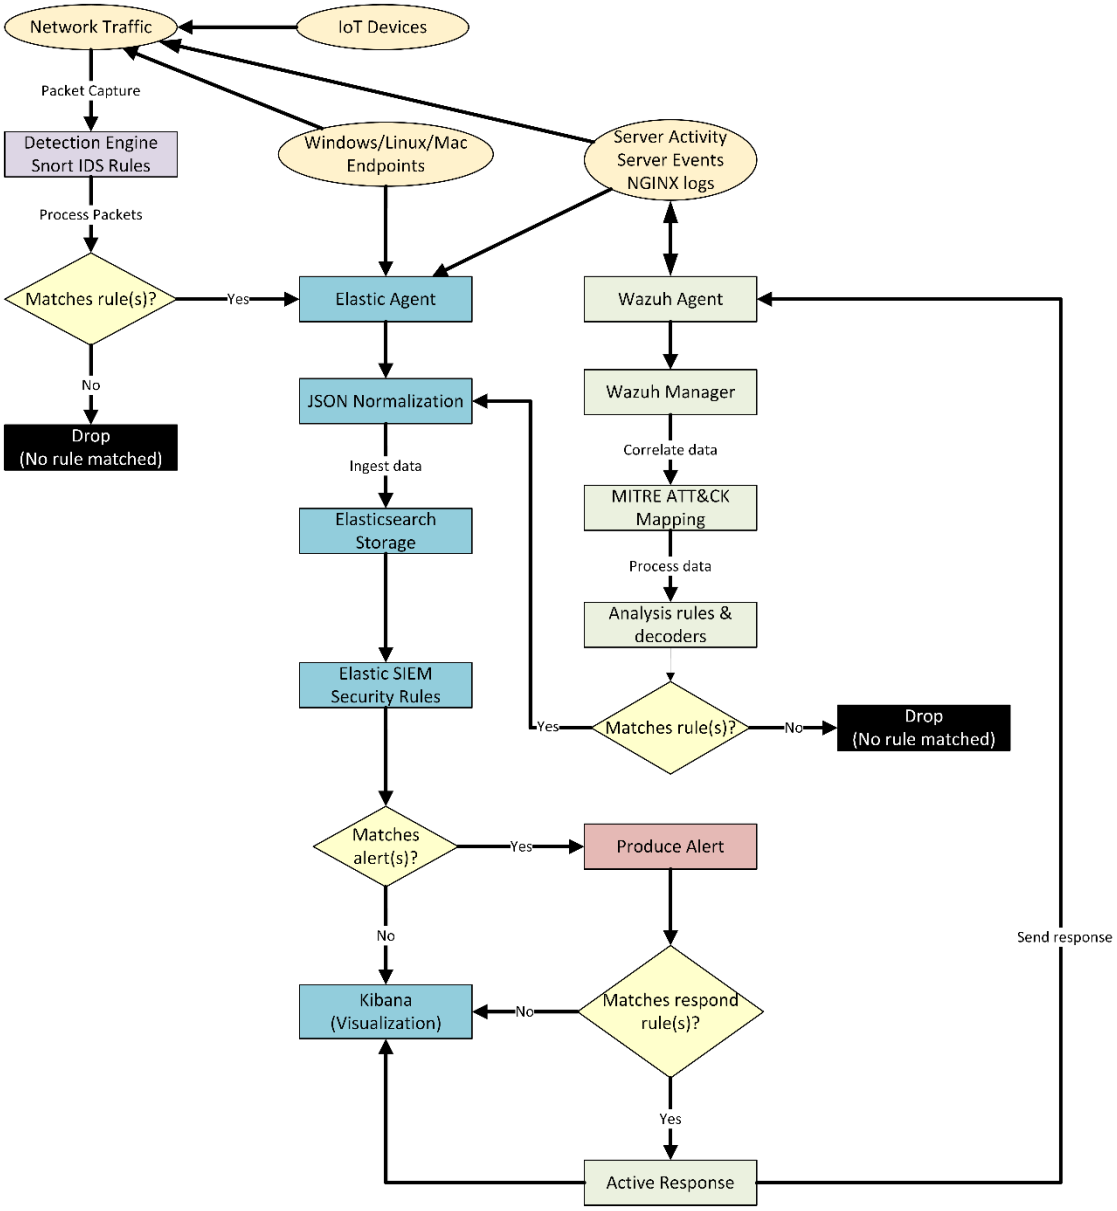

Figure 4-1. State Diagram

4.1 Reconnaissance Attacks

4.1.1 External Reconnaissance Attack

This attack involves an external threat actor attempting to gather information about a specific host or an entire network. It is assumed that the attacker is scanning various public IP addresses for services -meaning open ports- and will eventually discover the services running on the public IP address of the branch office. All reconnaissance

attacks were performed using the Nmap tool which is an open-source tool for network exploration and security auditing. [17]

Currently, ports 22 (Secure Shell - SSH server) and 8080 (NGINX web server) are exposed. The Snort sensor, running on the branch office server, successfully detects this port scan and forwards an alert to Kibana via the Elastic Agent, as illustrated in the next figure.

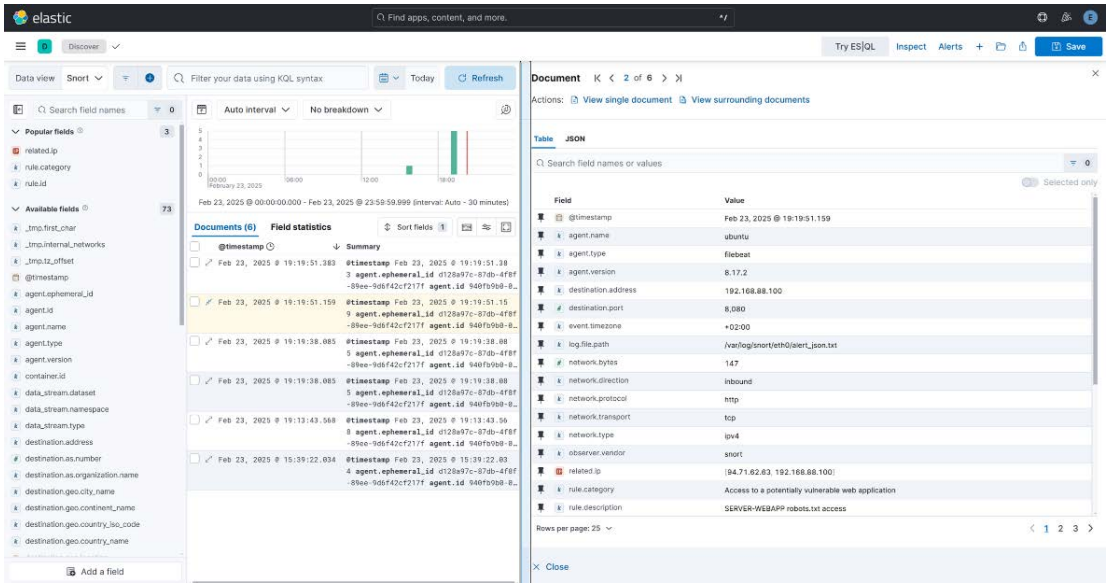

Figure 4-2. External Reconnaissance

Several important fields regarding the nature of the reconnaissance attack are displayed, including the source public IP address and the destination IP address. In this case, the destination IP corresponds to the local IP address to which the router -having the office's public IP- has forwarded the network traffic. The adversary will get an answer from port *8080* and eventually knows that there is an existing service running at this port. Additionally, Kibana automatically performs a geolocation lookup on the source public IP address to provide geolocation details. The *rule.category* and *rule.description* fields are retrieved from the Snort community rules, as they match *rule.id 1852*, which corresponds to the detection of an attempt to directly access the file *robots.txt*. This file is a plain text file used to control how search engine crawlers and web robots access a website's content, and it is typically located in the root directory of a website. [18]

A noteworthy type of reconnaissance attack involves a compromised host periodically transmitting information to an external entity. Snort successfully detected an

internal host that was repeatedly sharing the public IP address of the branch office with an external server, which could potentially be controlled by a threat actor. By doing so, the threat actor can continuously track the public IP address of the target network, even if it changes -a common thing today due to the exhaustion of IPv4 addresses- and the widespread use of dynamic IP allocation by Internet Service Providers (ISPs). The alert generated by Snort and forwarded to Kibana is illustrated in the next figure.

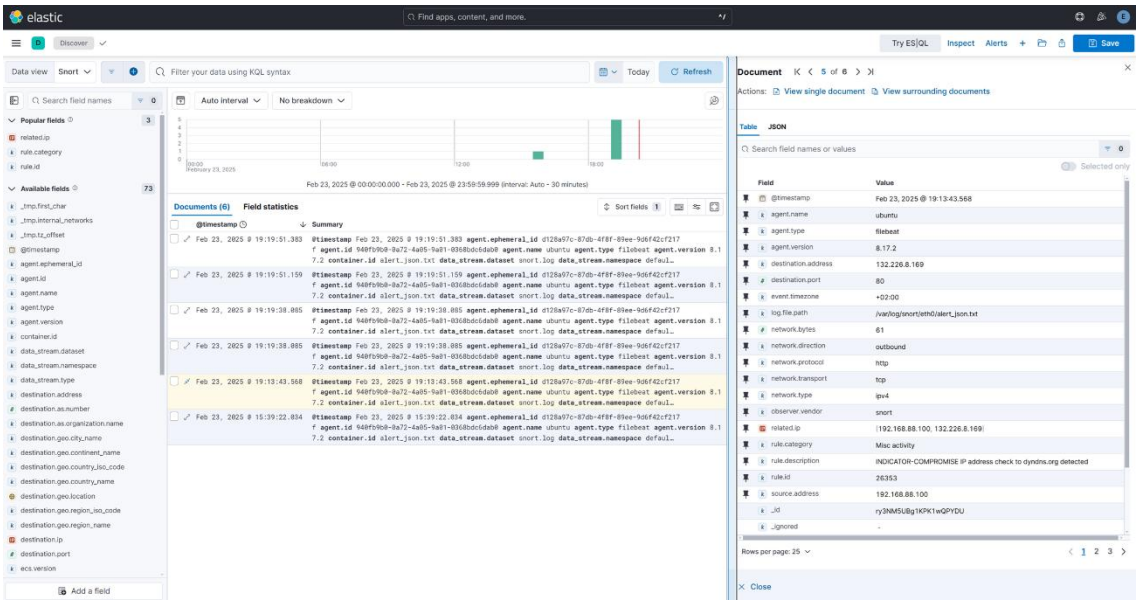

Figure 4-3. IP sharing with external entity

4.1.2 Internal Reconnaissance Attack

This attack follows a similar pattern; however, in this case, it is carried out by a threat actor with access to the branch office's LAN. It is assumed that the adversary has gained access to a compromised internal host. From this point forward, the attacker can scan the entire network to identify devices running vulnerable services. Once these vulnerabilities are discovered, exploiting them becomes significantly easier, allowing the attacker to achieve their objectives. The Snort sensor successfully detected a potentially malicious scan and generated an alert, which was sent to Kibana (Figure 4-3). This is crucial, as receiving the alert enables network defenders to take action, identify the attack, and mitigate the threat. Additionally, preconfigured rules can be implemented to automatically respond to an attack—such as disabling certain services, notifying

cybersecurity analysts, or taking other proactive measures to prevent further expanding of an attack.

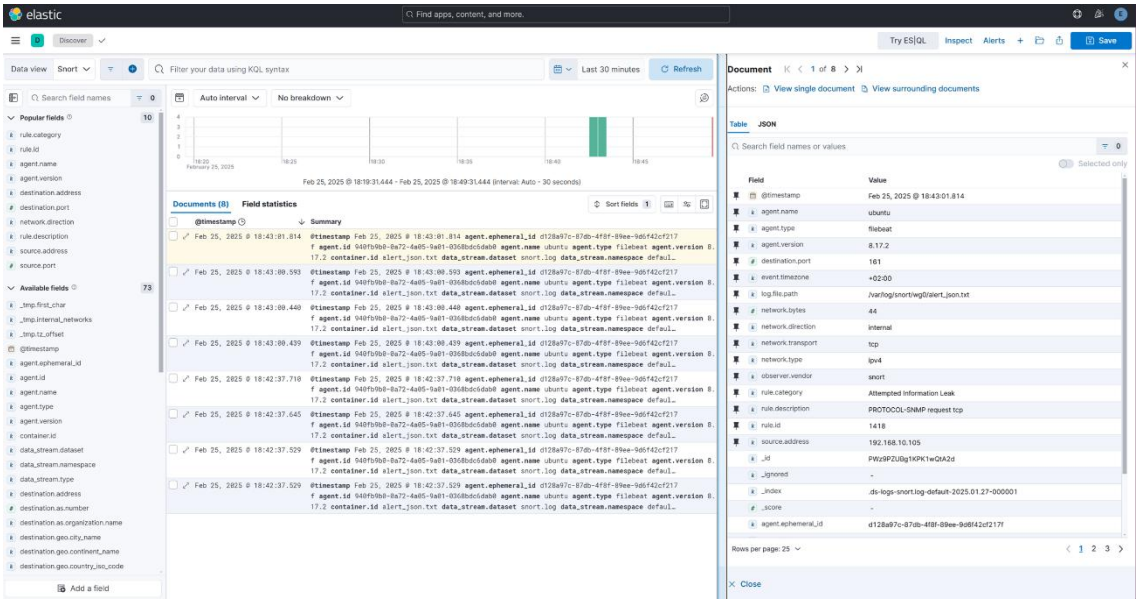

Figure 4-4. Internal Reconnaissance

4.2 Brute-force Attack

During this type of attack, a threat actor attempts to gain access to a system or database by systematically guessing usernames and passwords. This attack can occur both internally and externally, following the same methodology in either scenario. In this thesis, Threat Actor II will attempt to brute-force the branch office server by testing multiple username and password combinations. Additionally, this attack can also be classified as a Denial of Service (DoS) attack, as the continuous SSH connection attempts from the adversary machine may eventually overwhelm the server, making it unable to process further requests.

The attack was carried out using Hydra, a powerful password-cracking tool designed for brute-force attacks on various protocols, assisting security professionals in testing authentication vulnerabilities. [19] A typical course of action for a threat actor in this scenario would begin with scanning network devices for open ports. Once the available services are identified, the attacker would attempt to exploit them. In this case, it is assumed that the adversary has discovered that port 22 -the default port for the SSH

service- is open. The next step is to target this port. The attacker utilizes Hydra along with two .txt files: one containing a list of usernames and the other a list of passwords. Numerous such files are available online, containing commonly used username and password combinations. Once the attack is initiated, the targeted host is flooded with repeated login attempts on SSH port 22. Hydra systematically tests all possible username-password combinations until a match is found. The duration of the attack varies depending on several factors, including the hardware capabilities of both the attacking and targeted systems, as well as the network bandwidth.

Both the Elastic and Wazuh agents running on the branch office's server successfully detected the attack and immediately forwarded alerts to Kibana. As demonstrated in the following figures, data received by Kibana from the Elastic Agent (via Snort log files) and the Wazuh Agent indicate that an attack is actively taking place within the system. This real-time detection is crucial for security analysts, as it always provides comprehensive visibility into network activity. Furthermore, as will be discussed later, it is possible to implement automated response mechanisms -either through Kibana or through Wazuh configuration files- to not only detect such attacks but also actively mitigate them.

The mappings of Wazuh to the MITRE ATT&CK framework shown in Figure 4.5 play a crucial role at this stage, as they provide valuable insights into the tactics, techniques, and procedures (TTPs) employed by adversaries. This integration enables a more proactive defense strategy, allowing analysts to anticipate further malicious activities. Additionally, leveraging this framework facilitates threat intelligence sharing, helping organizations strengthen their security posture against evolving cyber threats.

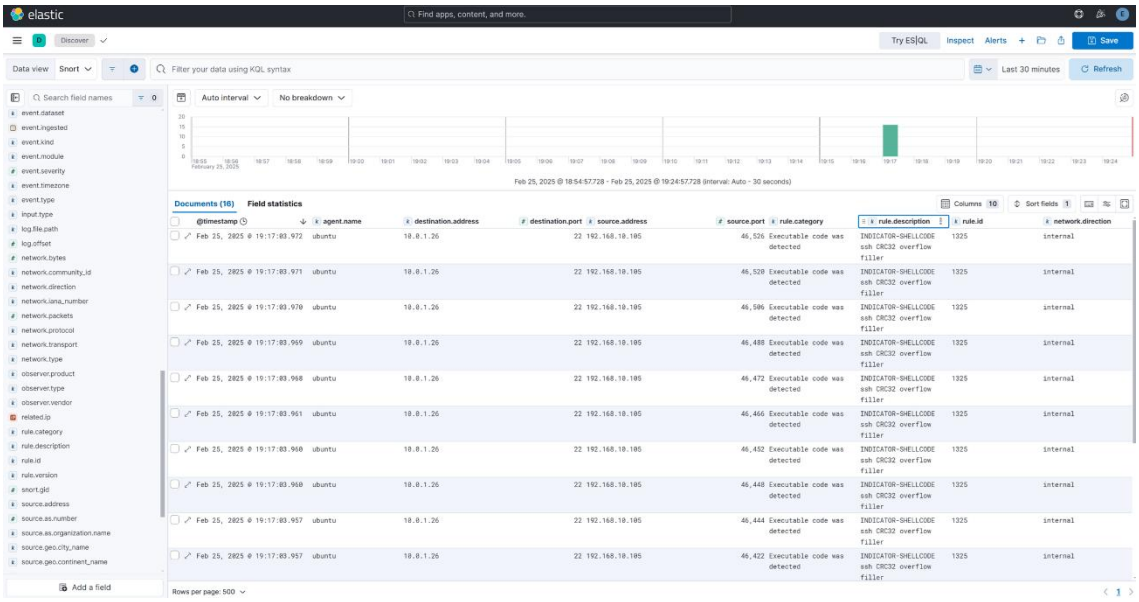

Figure 4-5. Snort detection of brute-force attack

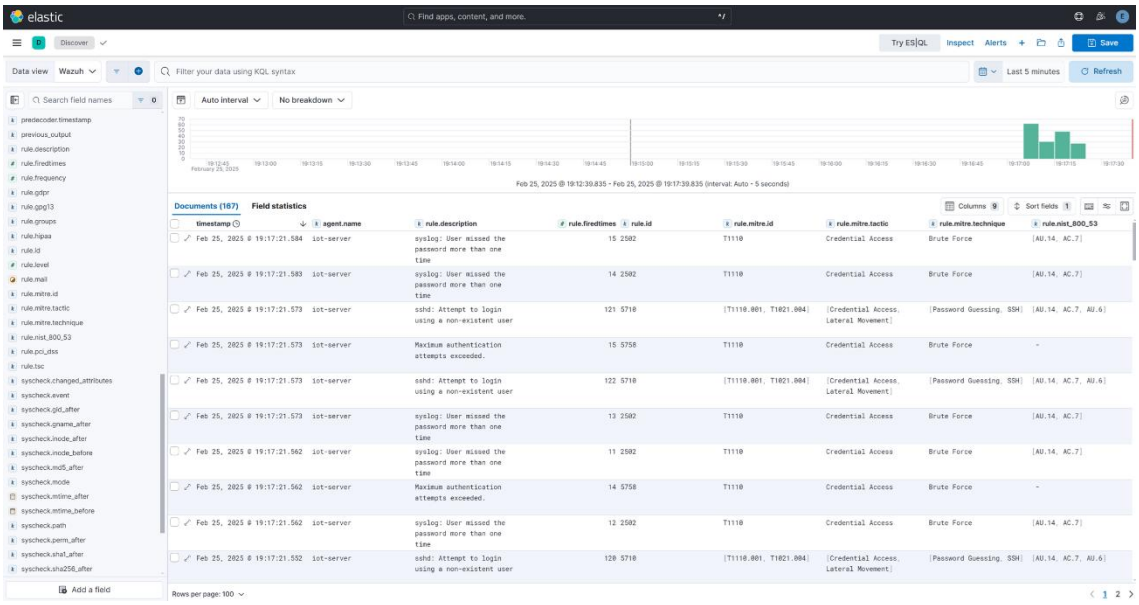

Figure 4-6. Wazuh detection of brute-force attack

4.2.1 Responding to a brute-force attack

It is imperative that a comprehensive security system incorporates automated responses to cyberattacks, particularly for relatively straightforward threats such as brute-force SSH attacks. Implementing automated defenses not only enhances security but also minimizes the impact of attacks by responding in real time.

There are multiple approaches to configuring rules that trigger automatic responses to such incidents. A common and effective strategy is to establish a rule that activates once a predefined threshold is exceeded or once a condition is met. For instance, a practical rule in this scenario would be the following: If more than X failed login attempts occur within Y minutes over SSH, automatically block SSH access for Z minutes. This measure not only safeguards the targeted host but also helps conserve network bandwidth by preventing excessive authentication attempts. Moreover, it provides administrators with valuable time to assess the situation and implement further countermeasures. However, attackers may resume their efforts once the temporary blocking period (Z minutes) expires, highlighting the need for more advanced and adaptive response mechanisms. Enhanced security measures, such as behavioral analysis and anomaly detection, can be achieved through the Machine Learning integration within the Elastic Stack. Unfortunately, this feature is not open-source and is therefore not demonstrated in this thesis. The following eXtended Markup Language (XML) code written on the Wazuh server configuration file instructs the Wazuh Agent to close port 22 on the host's firewall (and thus the SSH service) for 10 minutes when rule 5758 is triggered. As depicted in Figure 4.5, the description of rule 5758 is *"Maximum authentication attempts exceeded"*.

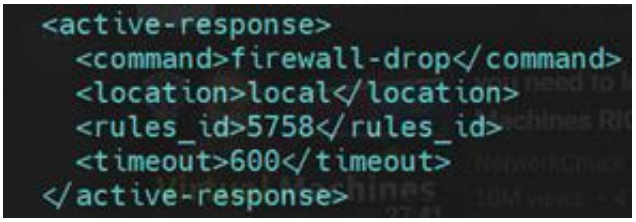

Figure 4-7. Wazuh active response

4.3 Denial Of Service (DoS) Attack

The DoS attack is a malicious attempt to disrupt the normal functioning of a targeted server, service, or network by overwhelming it with a flood of traffic or requests, rendering it inaccessible to legitimate users. This can be achieved easily through various methods, such as sending excessive data packets to functioning servers or hosts, thus forcing them to process these requests and finally overwhelming their resources with null

packets. A very common type of DoS attack is the TCP SYN flood attack, where the attacker exploits the TCP three-way handshake by sending a large number of SYN (synchronize) requests to a target server while spoofing the source IP addresses.

In this case study, an HTTP flood attack will be conducted for demonstrative purposes. The GoldenEye tool, an open-source Python script, will be employed to generate a huge volume of HTTP requests aimed at overwhelming the target NGINX service running on the office's server. This number of requests is designed to exhaust the server's resources, resulting in degraded performance or potential unresponsiveness. Similar to the brute-force attack, it can occur both internally and externally, following the same methodology in either scenario.

After successfully conducting the attack, a vast amount of events has been reported to Kibana. These events report different things depending on the software reporting them. As illustrated in Figure 4.7, Snort reports to Kibana from the network perspective, meaning that Snort interprets this attack as an attempt to flood the host's network stack. For Snort, this behavior resembles characteristics of exploit attempts, such as sending unexpected or malformed packets, leading it to interpret it as a challenge ACK provocation attempt. The specific Snort rule that triggered this alert indicates that the pattern of traffic generated by GoldenEye was similar to known attack vectors that aim to exploit vulnerabilities in the system.

Figure 4-8. Snort detection of DoS attack

On the other hand, the Wazuh Agent, which also communicates with Kibana, interprets the attack from the perspective of the host system. Given its deep access to all files and directories on the host, the Wazuh Agent can perform an in-depth analysis of the

logs generated by the NGINX service. This capability enables it to identify the precise issues encountered by the host during the attack. By providing detailed insights into the attack's impact, the Wazuh Agent not only aids in the immediate mitigation of the threat but also assists in assessing and correcting any potential damage incurred during the attack. The important information to observe in the next figure is the *full_log* sent by the Wazuh Agent reporting that *"worker_connections are not enough"* which is normal since the host is under a DoS attack.

Figure 4-9. Wazuh detection of DoS attack

In conclusion, it is evident that the system in question is both a fully featured host and network IDS. This architecture employs Snort to report on attacks based on their impact on the network, and Wazuh to monitor and analyze relevant host-specific activities. Such an integrated approach enhances the ability to detect, analyze, and respond to security threats.

4.4 Privilege Escalation Attack

In this scenario, it is assumed that the threat actor has successfully compromised the office server by obtaining SSH access via the user account "john," which is protected by a commonly used and weak password. This breach could have been the outcome of a prior successful brute-force attack, where the attacker exploited the password.

The first action an attacker would likely undertake in this scenario is to identify the operating system, system architecture, and any other pertinent specifications that

could facilitate the execution of the planned attack. Following this, the threat actor would seek to assess the level of control the user account "john" holds over the system. The commands to obtain such information are shown in the following figure.

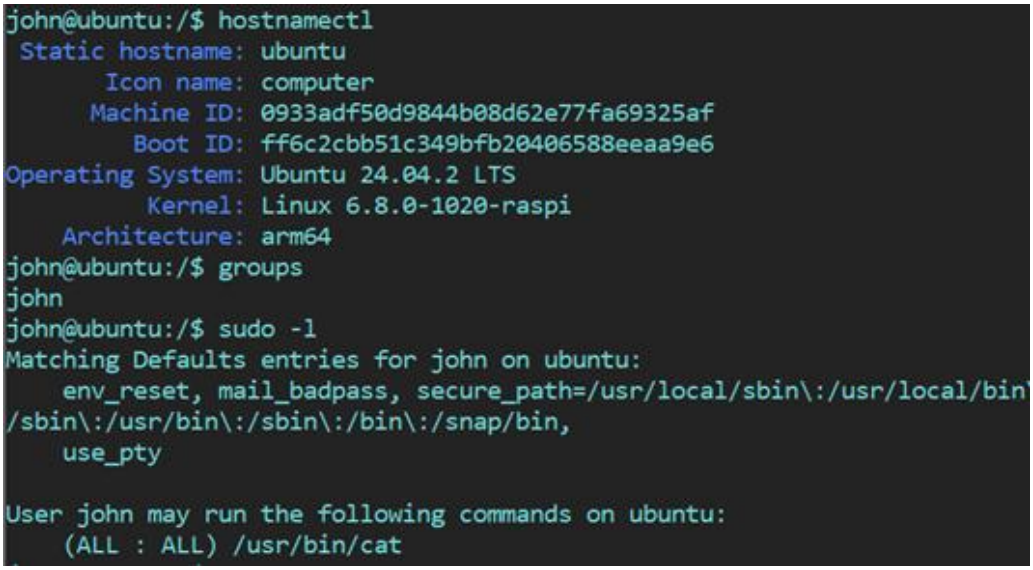

Figure 4-10. Post successful intrusion commands

Based on the information derived from the output in the figure above, the attacker has determined that the system in question is running the Linux operating system "Ubuntu 24.04.2 LTS" with an arm64 architecture. Additionally, the attacker knows that the user account "john" is only a member of its own group, which would typically limit the user's ability to perform significant actions within the system. However, by executing the command "sudo -l", the attacker uncovers that, despite the lack of any elevated privileges for "john," the system administrator has granted the user the ability to execute the "cat" command with elevated privileges. The cat command is used to concatenate files and display on the standard output (screen). [20]

Although being able to run the “cat” command with elevated privileges may not seem to be that powerful capability, the threat actor already has everything he needs to perform a Privilege Escalation attack and gain root access to the system. Continuing his action, he will now try to print the output of the “/etc/passwd” and the “/etc/shadow” files on his screen. Although the first one contains user information, it does not store password information for security reasons. This information can be found in the “/etc/shadow” file which holds all the user passwords hashed with different hashing algorithms.

Upon reviewing the output displayed in the following figure, the attacker identifies another active, login-capable user account, in addition to "john," named "bob". The attacker can determine that user “bob” is the only login-capable account apart from user “john” because user “bob” is the only one having a hashed password and not a special marker like * or ! which would indicate no password or locked account. The attacker now aims to extract the hashed password of this newly discovered user. This discovery leads the attacker to assume that "bob" may be the system administrator account, given the limited number of users with access capability.

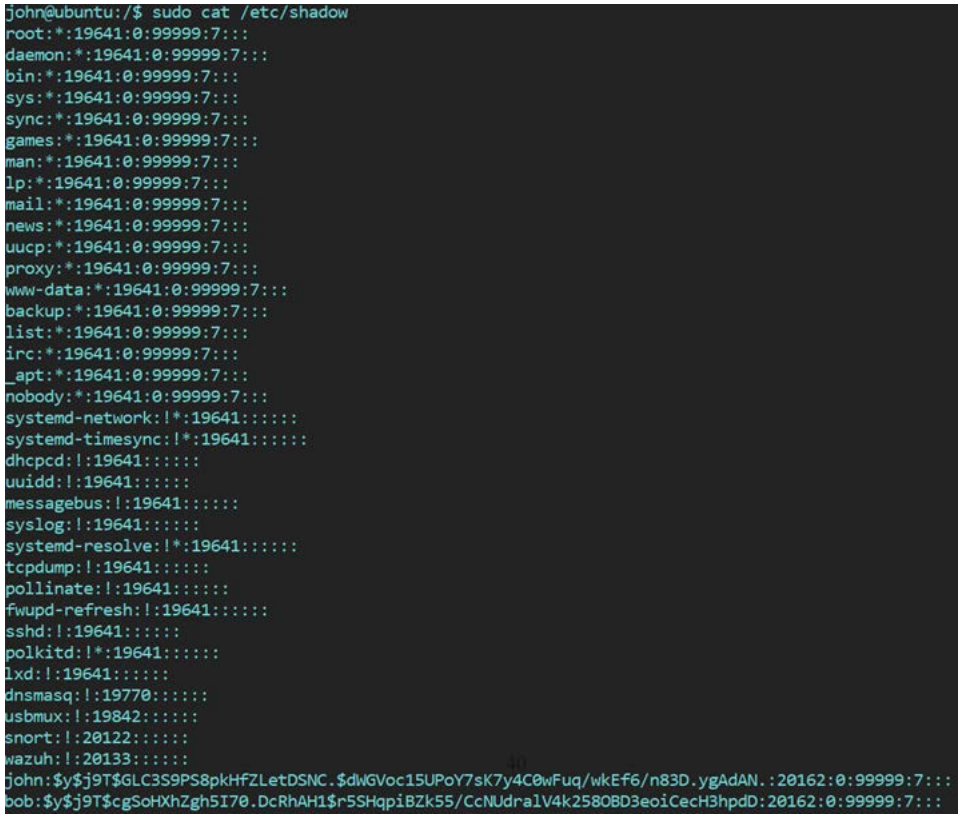

Figure 4-11. /etc/shadow file content

It’s now easy to identify the hash password of user “bob”. The fields contained in the “/etc/shadow” file are in the following format:

username:password:last_changed:min_age:max_age:warning_period:inactive_period:expire_date:reserved

That means that the password is the alphanumeric string starting with “y” and ending with “pdD”. The attacker can copy this hash to a system that he owns where he can easily run tools like “John The Ripper” or “Hashcat” to determine the exact password

of user “bob”, an action that is shown in the following figure. Those are some of the many tools that are freely available on the Internet and can perform password cracking even for the SHA-512 integrity algorithm.

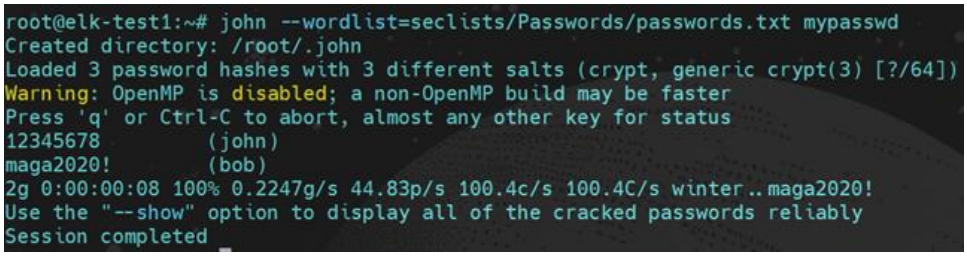

Figure 4-12. Password cracking using John The Ripper

After successfully executing this step, the attacker can leverage the “su” command to switch users and assume control of the “bob” account, which possesses administrative privileges. This effectively results in a successful Privilege Escalation Attack. At this stage, the attacker gains full control over the system, granting unrestricted access to its resources and the ability to execute arbitrary commands, potentially compromising its integrity, confidentiality, and availability.

4.4.1 Detecting a Privilege Escalation Attack

There are limited measures that can effectively prevent an attack of this nature. Once a threat actor gains access to a system, monitoring their actions and potential lateral movement becomes challenging. However, one approach to detecting malicious activity of this kind, is to configure alerts that trigger whenever any user on a system monitored by Wazuh (and thus by the Elastic Stack) executes a command with elevated `sudo` privileges.

For this scenario, a specific alert was implemented within the Wazuh Server configuration. This alert instructs the Wazuh Agents to generate and forward an alert to the server whenever a user—who is not a member of the administrators’ group—executes a command requiring `sudo` privileges. This measure enhances visibility into unauthorized privilege escalation attempts, aiding in the early detection of potential security breaches. As illustrated in the next figure, the alert also unveils the exact

command executed by the user as well as the mapping with MITRE's ID, Tactic and Technique used in a possible attack like this one.

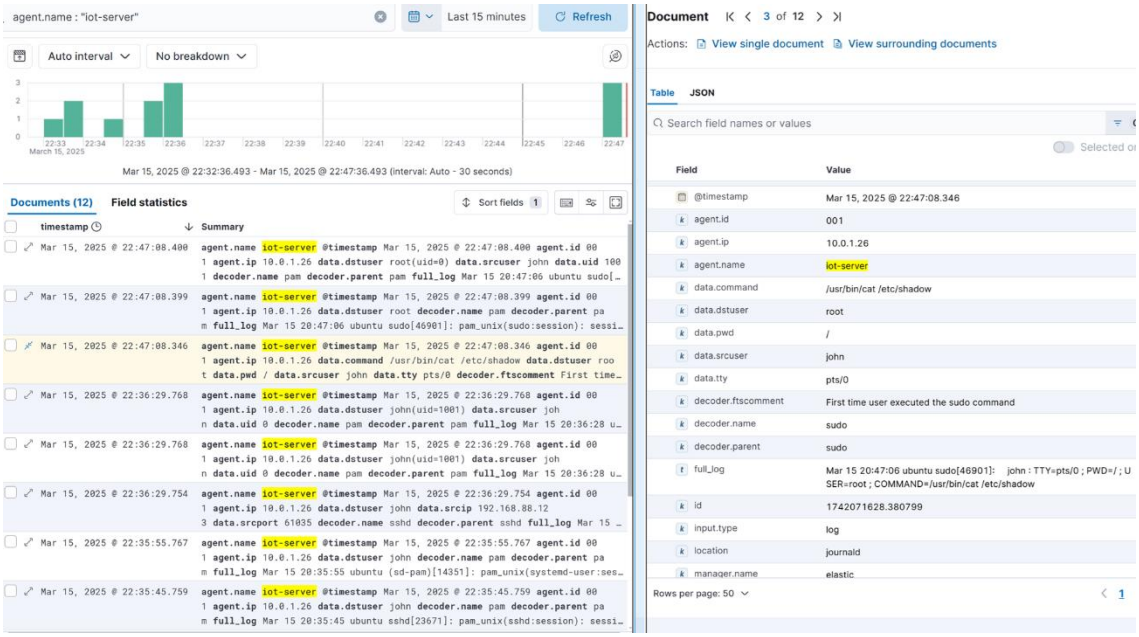

Figure 4-13. Privileged command detection in Kibana

5. Results

5.1 Implementation

The security solution proposed in this thesis can serve as a comprehensive cybersecurity framework for businesses and organizations of all sizes. The Elastic Stack, along with its components, provides a robust and feature-rich security solution, enabling full visibility and control over an organization's cybersecurity landscape. Logs and data from any device can be ingested and formatted in various ways, depending on the desired result. Important data can be correlated and merged into dashboards. In addition, security alerts can be created, reporting security breaches or malfunctions to the SOC personnel. Automated incident response can also be configured through the user-friendly interface provided by Kibana, allowing for swift and efficient threat mitigation. Moreover, the Elastic Stack supports various integrations, such as Snort, enabling cybersecurity analysts to leverage advanced pre-built scripts and detection rules that would otherwise require significant time and effort to develop and implement manually. These integrations enhance the overall security posture by improving threat detection, analysis, and response processes.

By implementing a central server at the organization's HQ as a hub -probably along with a backup hub server at a secondary location- organizations can benefit from centralized administration and management. All devices operating within the organization's WAN are continuously monitored by the SIEM server located at the HQ. For devices where deploying agents such as the Elastic Agent or Wazuh Agent is not feasible -such as IoT devices- a lightweight, cost-effective server can be installed at each remote site. This server would act as a network monitoring node, capturing and analyzing all incoming and outgoing network traffic. By doing so, organizations can extend their security coverage, ensuring comprehensive threat detection and response across their entire infrastructure.

The size of the organization and the number of the total devices operating in it are the main factors that will determine the hardware specifications of the SIEM server. Elasticsearch and Kibana are components that can operate with minimum requirements, although they will not perform well if they get overwhelmed by millions of logs due to

the many devices that may exist in the network. In this thesis, all the hardware used is appropriate for implementation of the system in question in a small-sized company.

5.1.1 IDS Implementation

The selection of Snort as the Intrusion Detection System (IDS) was not a straightforward decision, as Suricata also offers highly regarded IDS and Intrusion Prevention System (IPS) capabilities. However, Snort was ultimately chosen due to its superior community-driven rule sets and its backing by Cisco, a highly respected leader in the cybersecurity industry. Although Suricata requires significantly less time for installation and configuration, Snort was still preferred due to its extensive documentation, broader industry adoption, and continuous support from Cisco, ensuring long-term reliability in threat detection and prevention.

Snort was configured with the community rules which can be found at [15]. Some custom rules were also added for the sake of demonstrating its functionality. Since the Wazuh Agent was running on the office server covering the needed for intrusion prevention, there was no need to configure Snort as IPS. That said, Snort was responsible for filtering all the inbound and outbound network traffic with its ruleset, forwarding alerts to the SIEM when specific rules were triggered.

5.1.2 SIEM and XDR Implementation

The fundamental components of Elastic Stack for running a SIEM are Elasticsearch and Kibana. Logstash is no longer mandatory since Elastic has introduced the Elastic Agent which is basically a plug-n-play client software for collecting all the necessary logs and data and forwarding them to Kibana. Essentially, the use of different Beats still exists, although it can be replaced with the more efficient integrations (e.g. NGINX integration for web servers, Snort integration, Docker integration for collecting metrics from containers).

Elasticsearch was installed as the SIEM's database, while Kibana was deployed as the primary interface for human interaction with the system. Both were set up natively on a virtual machine (VM) running Ubuntu. Elastic provides a well-documented and

structured process for deploying its open-source software, making the installation relatively straightforward. There is also comprehensive documentation for installing everything on docker containers which facilitates backups, migrations and resources management. However, the most challenging and time-consuming aspect of setting up the SIEM was understanding how all components interconnect and configuring Kibana to ensure seamless functionality and performance.

The deployment of Wazuh as the Extended Detection and Response (XDR) solution was also carried out natively, though its components can alternatively be installed in Docker containers. In this case, Wazuh was integrated primarily to provide automated incident response, a feature that exists within the Elastic Stack but is only available in its paid subscription plans. Since the proposed system adheres to an open-source framework, opting for the paid Elastic feature was not an option.

Among Wazuh's components, only the Wazuh server was required, as its equivalent database and dashboard functionalities were already covered by Elasticsearch and Kibana, respectively. The most time-consuming aspect of this integration was not the installation and configuration of Wazuh itself but rather the extensive adjustments required to ensure seamless interoperability within the Elastic Stack. The latest versions of all the software components were used. In earlier versions, Wazuh offered a native integration with Kibana, but as Wazuh transitioned to developing its own database and dashboard, this integration was discontinued, necessitating additional configuration efforts.

5.2 Detection and Response

All the components combined can offer a complete detection and response solution resulting to a SIEM and Security Orchestration Automation and Response (SOAR) system. **Reconnaissance** attacks being executed on the network, either from an external or from an internal source, are immediately detected and forwarded to the SIEM by the network IDS, Snort. The Mean Time To Detect (MTTD) is less than a second, assuming network connectivity is present and enough to forward the traffic, which nowadays should be since we’re only talking about some kilobytes being transferred each time.

Elastic Agents serve as integral components of the system, continuously monitoring host behavior and transmitting logs to Kibana for correlation and analysis. This approach enables the identification of complex attack patterns that may not be evident when analyzing individual events in isolation. While a single action performed on a specific device may not independently trigger an alert, the aggregation of multiple actions across various devices may indicate a coordinated attack, necessitating intervention to mitigate potential threats.

The presented **brute-force** attack to the office server was immediately detected by both Snort and the Wazuh Agent, which correctly sent multiple alerts to Kibana immediately. Wazuh's function for automated incident response reacted within some milliseconds, blocking the SSH access to the server for a preconfigured amount of time so as there is time for mitigating the running attack. This leads to the conclusion that the Mean Time to Respond (MTTR) is effectively immediate, provided there are no issues with network connectivity between the host and the SIEM. However, the Mean Time to Contain (MTTC) is influenced by additional network configurations and, of course, the human factor. Specifically, it depends on how quickly the Cybersecurity Analyst escalates the alert to the Incident Responder and how efficiently the Incident Responder is able to take action to mitigate the attack.

The critical role of Kibana integrations is demonstrated in the **Denial of Service (DoS)** attack analyzed in this thesis. The NGINX integration, deployed within the Elastic Agent on the system hosting the web server, was the key component responsible for detecting the attack and generating alerts in Kibana. From that point forward, pre-configured rules in Kibana -such as automatically shutting down the web server or blocking access from the attacking IP address- can be implemented. These rules can be configured to be executed automatically in response to such security events, ensuring an effective mitigation process.

Last but not least, while the proposed system offers extensive security capabilities, it may not be the most straightforward solution for detecting and preventing **privilege escalation** attacks or tracking a threat actor's lateral movement within a single host or across an entire network. Implementing such advanced functionalities would require additional configurations, including the development of custom scripts, the creation of specific detection rules tailored to each environment, and ongoing maintenance to ensure these measures remain effective and aligned with the evolving security needs of the

organization. Despite these challenges, the system provides a solid foundation upon which such capabilities can be built, offering flexibility for organizations willing to invest in fine-tuning their security.

6. Discussion

There is no doubt that the contemporary world demands a new, multi-sided approach to security. The complexity of modern cyber threats has escalated significantly compared to previous years, with adversaries employing more and more sophisticated attacks over time. This evolution alone underscores the necessity of adopting a comprehensive, multidimensional defense strategy. Traditional layered security measures are no longer sufficient to safeguard networks and systems effectively. Firewalls, IDSs and IPSs, while still valuable, are inadequate as standalone solutions for protecting an organization's digital infrastructure. Instead, a more advanced and adaptive approach is required, specifically one that integrates event correlation, proactive threat analysis, and real-time alert management to derive actionable insights and enhance the overall security posture.

Technologies such as SIEM, SOAR, and XDR serve as comprehensive solutions for safeguarding the assets that are valued the most. Beyond threat protection, these technologies offer advanced management capabilities, event correlation, automated incident response, and centralized security features. In advance, they enhance efficiency by integrating various security tools and workflows, and providing real-time visibility into potential threats, thereby strengthening an organization’s overall cybersecurity posture. Traditional security systems such as IDSs and firewalls can still be of use, as they can be integrated into centralized security solutions without significant complexity. This is achieved through sending their logs, alerts and generic data to a SIEM, where data can be ingested, enhanced and correlated with data collected by other sources. This way it is possible to get a clear understanding of the vulnerabilities that exist in an infrastructure and thus implement automated responses to threats using thresholds and a combination of automated incident response rules.

6.1 Future Work

Artificial Intelligence (AI) and Machine Learning (ML) are becoming increasingly known across various industries, including cybersecurity. While still in their

early stages, these technologies are making a strong entrance into the field, reshaping how threats are detected and mitigated.

Elastic has already integrated ML capabilities into its Elastic Stack, though these features are only available in paid subscriptions. According to its documentation, the ML functions focus on security, enhancing SIEM's ability to automatically generate attack detection rules, trigger alerts, and mitigate threats in real time by automatically creating incident respond rules. The more data sources that are added in the stack, the more knowledge the system has and thus the more efficiently it will act when in need.

That said, perhaps one day cybersecurity analysts and incident responders will only need to maintain the data sources for their security tools. However, I personally doubt that this will soon be the case, as humans continue to evolve in their ability to manipulate and exploit situations. Just as security technologies advance, so do the tactics of adversaries, making cybersecurity an ongoing game of adaptation and intelligence. The more sophisticated our defenses become, the more creative attackers will get in finding ways to bypass them - ensuring that human ability remains an indispensable part of the equation.

References

[1] C. Onwubiko, "Cyber security operations centre: Security monitoring for protecting business and supporting cyber defense strategy," in *2015 International Conference on Cyber Situational Awareness, Data Analytics and Assessment (CyberSA)*, London, UK, 2015.

[2] Tall et al., "Survey of data intensive computing technologies application to to security log data management," in *3rd IEEE/ACM International Conference on Big Data Computing, Applications and Technologies*, 2016.

[3] N. Hammoud, "Decentralized log event correlation architecture," in *MEDES '09: Proceedings of the International Conference on Management of Emergent Digital EcoSystems*, 2009.

[4] Z. G. Muhammad Shafiq, "The Rise of “Internet of Things”: Review and Open Research Issues Related to Detection and Prevention of IoT-Based Security Attacks," *Wiley Online Library,* 03 August 2022.

[5] M. Scapicchio, "What is a SOC?," 15 March 2024. [Online]. [Accessed 5 February 2025].

[6] "What is a firewall?," Cisco, [Online]. Available: https://www.cisco.com/site/us/en/learn/topics/security/what-is-a-firewall.html. [Accessed 5 February 2025].

[7] PaloAlto, "What is an Intrusion Detection System?," PaloAlto, [Online]. Available: https://www.paloaltonetworks.com/cyberpedia/what-is-an-intrusion-detection-system-ids. [Accessed 5 February 2025].

[8] Fortinet, "What Is SIEM?," Fortinet, [Online]. Available: https://www.fortinet.com/resources/cyberglossary/what-is-siem. [Accessed 5 February 2025].

[9] Cisco, "CyberOps Associate Course," Networking Academy, [Online]. Available: https://www.netacad.com/. [Accessed 7 February 2025].

[10] "What is Elasticsearch?," Elastic, [Online]. Available: https://www.elastic.co/docs. [Accessed 7 February 2025].

[11] "Machine Learning in the Elastic Stack," Elastic, [Online]. Available: https://www.elastic.co/guide/en/machine-learning/current/index.html. [Accessed 7 February 2025].

[12] "Computer Security Incident Handling Guide 800-61r2," NIST, [Online]. Available: https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-61r2.pdf. [Accessed 11 February 2025].

[13] R. P. Team, "Raspberry Pi 5 Product Brief," [Online]. Available: https://datasheets.raspberrypi.com/rpi5/raspberry-pi-5-product-brief.pdf. [Accessed 15 February 2025].

[14] MikroTik, "MikroTik hAP Product Brief," [Online]. Available: https://mikrotik.com/product/RB951Ui-2nD. [Accessed 15 February 2025].

[15] Cisco, "Snort Community Rules," [Online]. Available: https://www.snort.org/downloads/#rule-downloads. [Accessed 16 February 2025].

[16] Cisco, "Snort Rule Writing," [Online]. Available: https://docs.snort.org/rules/. [Accessed 16 February 2025].

[17] Nmap, "Nmap Documentation," [Online]. Available: https://nmap.org/book/man.html. [Accessed 25 February 2025].

[18] Cisco, "SID 1:1852," Snort Community Rules, [Online]. Available: https://www.snort.org/rule_docs/1-1852. [Accessed 23 February 2025].

[19] KaliLinux, "Hydra Documentation," [Online]. Available: https://www.kali.org/tools/hydra/. [Accessed 25 February 2025].

[20] Wikipedia, "cat (Unix)," [Online]. Available: https://en.wikipedia.org/wiki/Cat_(Unix). [Accessed 15 March 2025].