Ransomware Resilience: An Integrated Framework for Mitigation, Recovery, and Best Practices using SIEM and Machine Learning

[1]Abhikshit Gogoi and* [2]Ashim Sharma

[1,2]School of Computer Science and Engineering, Lovely Professional University
Phagwara, India
[1]avikgogoi6@gmail.com
[2]Ashim.30054@lpu.co.in

Abstract. This paper deals with the major and growing threat of ransomware, which causes significant financial and business losses by using developed multi-stage attacks. The modern threat environment has changed to Ransomware-as-a-Service (RaaS) format and polymorphic code, making past security models that used extensive reliance on outdated signature-based detection or the use of diverse endpoint solutions insufficient. Such conventional defences are usually ineffective in identifying polymorphic code and often can only react once the destructive phase of the encryption process has already started, and so a contextual, predictive, and integrated defence is critical. The current paper proposes and substantiates a new combined Security Information and Event Management (SIEM) system that uses the Elastic Stack (Elasticsearch and Kibana) to ingest and correlate host and network logs in real time. The methodology adopts a hybrid analysis engine that combines the latest advanced behavioural profiling models, such as SOM (self-organising maps), Random Forest Classifiers, and Long Short-Term Memory networks, overlaid with custom-designed Sigma rules for hunting multi-indicator strikes such as Emotet, Trickbot, and IcedID. The research also uses forensic tooling, including FTK Imager for capturing volatile data and IDA Pro to reverse engineer malware binaries. In this way, even detection logic is founded upon a solid base of low- level execution. The architecture was tested in comparison with simulated multi-stage attacks in a virtualised endpoint space, which includes a Windows 10 victim virtual machine and a Kali Linux Command-and-Control (C2) server. The integrative approach was far more effective than the traditional signature-guiding approach, which achieved a detection accuracy of 92.3, a precision of 95.7, and a false-positive rate of 0.02. The high-quality result confirms that the entire SIEM and Machine Learning method is an entirely successful and proactive answer to modern ransomware.

Keywords: Ransomware, SIEM, Machine Learning, Elastic Stack, Anomaly Detection, LSTM, Digital Forensics.

B. Singh et al. (eds.), *Proceedings of the International Conference on Advances in Computing Technology and Artificial Intelligence (COMPUTATIA 2026)*, Atlantis Highlights in Intelligent Systems 18,
https://doi.org/10.2991/978-94-6239-713-2_19

1 Introduction

The modern digital environment is ever faced with advanced and persistent cyber threats, and ransomware has become one of the most devastating and monetarily crippling types of cybercrime [2]. Having the potential to significantly harm the well-being of private people and organisations in all fields, including the critical infrastructure, healthcare, and finance, ransomware, a type of malware, usually encrypts the data of its target and demands ransom before providing them with access to it. The gravity of this threat is not merely in the immediate financial loss associated with ransom payments but also in the operational paralysis, reputational degradation, and potential data exfiltration that accompany these attacks. The industrialisation of cybercrime contributes to the rapid growth of the threats, the most prominent one is the Ransomware-as-a-Service (RaaS) model. The paradigm also reduces the number of entry barriers of cybercriminals as tools, infrastructure and support services are ready, which allows inexperienced actors to develop advanced campaigns [3]. As a result, this has increased the number and speed of the attacks and required a radical change in the planning of defence. Ransomware has evolved beyond the use of mere locker ransomware variants, which merely deny the user interface access, to crypto-ransomware, which directly alters the file structure by using strong encryption algorithms like the RSA and the AES, and in many cases, a significant number of ransomware variants is coupled with a double extortion policy where data is stolen before being encrypted to be used against the victim, as shown in Fig. 1 [2].

Conventional, signature-based security solutions are often ineffective in terms of the polymorphic and highly dynamic quality of modern ransomware versions. These old techniques based on familiar identifiers like file hashes or memory strings cannot keep up with new strains that use obfuscation, packing and dynamic compiling to change their signature [4]. It is therefore an urgent requirement that new and agile methodology be developed that can identify and neutralise ransomware actions as per behavioural patterns and not merely in regard to the signature. In addition, the isolated security systems, which only watch network boundaries or endpoint conditions, do not offer the built-in visibility to identify the multi-stage attacks moving laterally within an environment. The proposed paper concurs with the weaknesses presented by traditional and siloed security solutions by emerging with an adapted, comprehensive architecture of ransomware threat detection and mitigation. The main hypothesis of the investigation is that the combination of the correlation power of a powerful Security Information and Event Management (SIEM) solution with the predictive ability of a machine learning tool will provide a better result. ML can significantly outperform static detection methods [10]. We propose a framework built on the Elastic Stack (Elastic Search and Kibana), integrated with a trio of specific Machine Learning models: Self-Organising Maps (SOM) for anomaly detection, Random Forest Classifiers for high-accuracy binary classification, and Long Short-Term Memory (LSTM) networks for sequential behavioural analysis [5].

The combined method is aimed at helping to perform quick, accurate, and automated detection of threats by performing continuous monitoring and matching system and network activity logs in real time. The proposed solution through machine learning

model predictive behavioural analysis and through integration of custom detection policies into the Security Information and Event Management (SIEM) system significantly increases the ability of an organisation to predict, identify and react to known and zero-day ransomware attacks. The targeted goals that will inform the study are numerous. First, to design and combine predictive machine-learning models, such as Self-organising Maps (SOM), Random Forest and Long Short-term Memory (LSTM) to structurally identify malicious ransomware actions and benign system behaviours. Second, to engineer a comprehensive SIEM model using Elasticsearch and Kibana, and hence ensure the real-time capture and association of logs generated on Windows endpoints. Third, to confirm high-fidelity detection rules based on the subset of Indicators of Compromise (IOCs) that are typical of most common ransomware works around Emotet, IcedID, and Trickbot. Fourth, to rigorously evaluate the system's performance metrics, specifically targeting a detection accuracy above 90% and a precision above 95% against traditional methods. Finally, the study aims to implement and test automated mitigation capabilities, such as system isolation and process termination, to minimise the destructive impact of a successful breach.

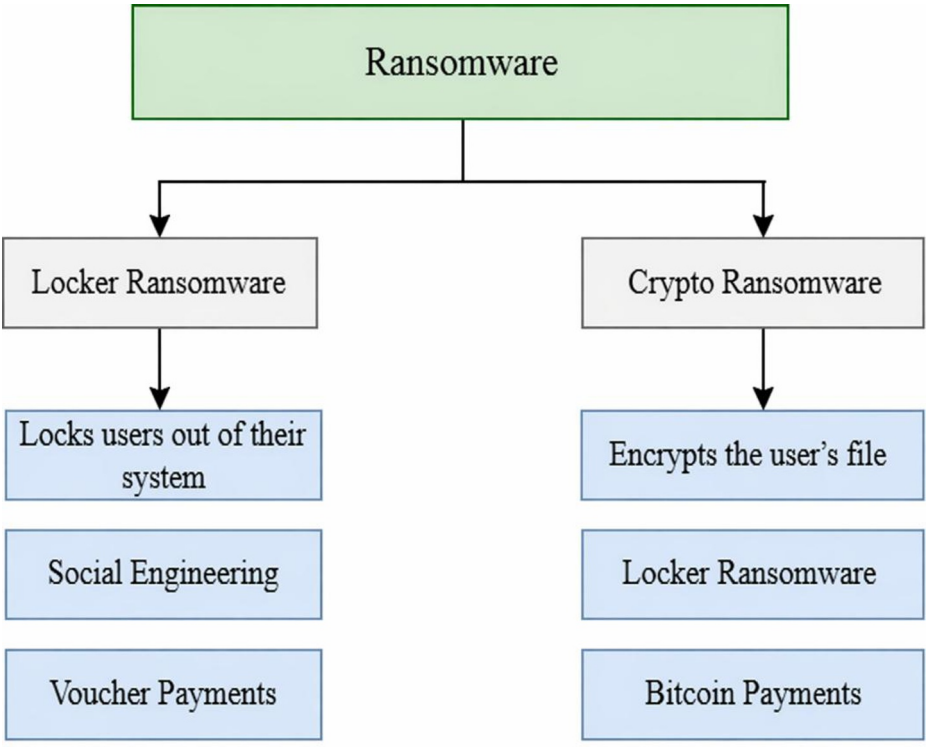

Fig. 1. Types of Ransomware

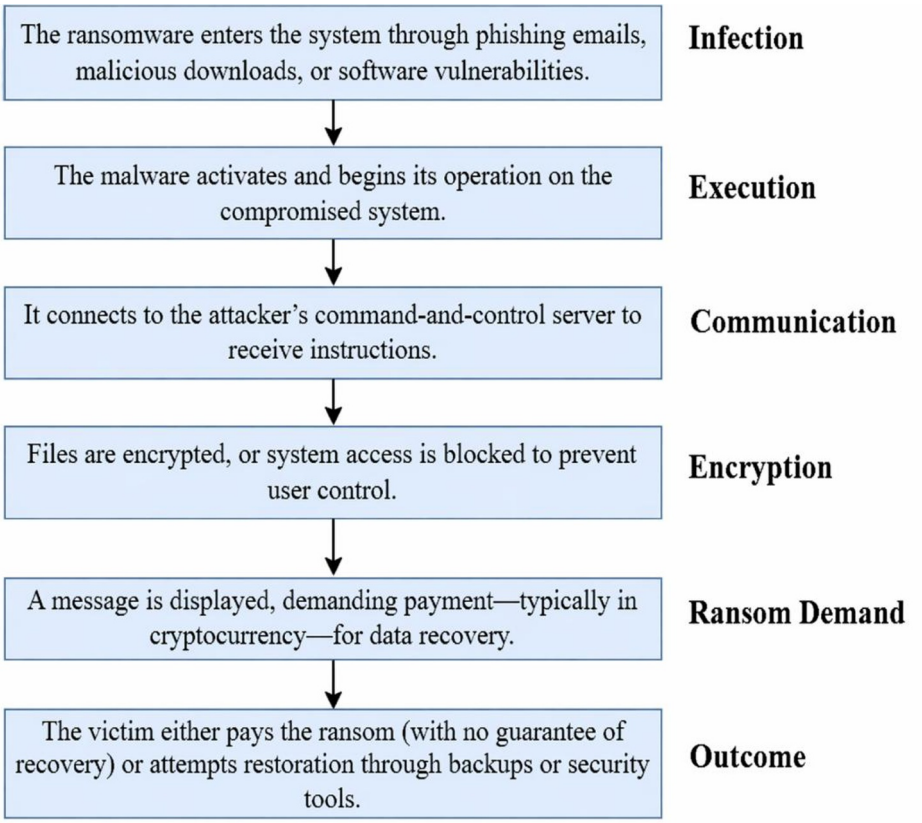

Fig. 2. The Ransomware life cycle

2. Literature Review

The development of ransomware detection tools has been driven by the increase in complexity inherent to the attacks. To define the proposed framework in connection with the academic research, this section performs a review of detection methodology development, and how the paradigm is no longer static, and based on rules, but can be dynamically and adaptively implemented relying on machine-learning techniques. It provides a critical analysis of current-state conditions in the field of detection, visualization, and recovery paradigms.

2.1 Vulnerabilities in Data Transmission and Storage

Ransomware threat is not just about a disk with fixed files, which is a sub-subtlety of endpoint-based research. A study by Raul Reinosa Simon, Clara I. Valero, Jose A. Martınez Cadenas, Elisa R. Heymann, Ignacio Lacalle, Barton P. Miller, and Carlos E. Palau [6] has both theoretical and practical implications of data in motion ransomware

attacks. In contrast to conventional ransomware, which is in a state of rest, they presented a lab to show Man-in-the-Middle (MitM) attacks based on ARP spoofing; by monitoring client-server communication and XOR encrypting HTTP POST requests, attackers could abuse backups as they were being stored on a server. It was determined by Simon et al. that the risk of such attacks is medium as it has to be so difficult technically, but the impact is high because it undermines recovery mechanisms that are highly valued by the organisations. This highlights why detection systems to track network anomalies such as ARP spoofing or uncharacteristic payload entropy, as well as endpoint behaviour, are urgently needed, which is precisely what this research paper intends to address by incorporating integrated SIEM logging.

2.2 Limitations of Traditional Detection and the Rise of SIEM

The failure of fragmented security solutions has led to the use of centralised SIEM solutions. Another integrated SIEM architecture, which makes use of Elastic Stack, was presented by Bhanu Prakash Rayabandi, specifically aimed at integrating SIEM with the needs of small and medium-sized businesses (SMBs) [7]. The analysis presented by Rayabandi argues that the traditional methods are often outdated, simplistic and poorly organised and are unable to match network notifications with the establishment of endpoint processes. Rayabandi has shown that the combination of custom Sigma rules and Kibana Query Language (KQL) queries, subjected to the virtualisation of a Windows endpoint environment and utilising the MITRE ATT&CK framework, was effective in identifying IOCs that pointed to variants of Emotet and Trickbot. This groundwork justifies the application of the Elastic Stack as an effective detection platform. However, in his research, Rayabandi mainly focused on rule-based detection. The current research explores this question by integrating predictive machine-learning programs that aim at detecting threats that could not be detected by fixed rule sets.

2.3 Honeypots and Deception Technology

In an effort to understand malware behaviour, researchers have increasingly resorted to deception technologies that seek to attract and study adversaries. Fathi Kamil Mohad Zainudin, Izzatul Hazirah Ishak, Sharifuddin Sulaman, Farah Ramlee, Nur Sarah Jamaludin, and Shuaib Chantando [8] introduced the "Lebahnet" platform, which is essentially a virtual honeypot infrastructure deployed by Malaysia Computer Emergency Response Team (MyCERT). It is of special concern to Lebahnet sensor architecture, which consists of two key components: Cowrie medium-interaction SSH/Telnet honeypot, which is designed to log brute-force entry attempts and shell interactions, and Dionaea malware-capture honeypot that uses the LibEmu x86 emulation framework to identify shellcode payloads. A Study by Zainudin of malware binaries sampled in 2017-2019 found that the main threat type in 2018-2019 was ransomware, with the WannaCry strain being the most prevalent. The use of sandbox analysis with Cuckoo Sandbox in the research points to the importance of malware observation gradient in a controlled environment in order to understand its noises as the construct that this study adopts with the help of forensic simulation and reverse engineering, and, thus, creates ground truth information towards SIEM systems.

2.4 Automated Web-Based Attack Vectors

Muhammad Usman Rana, Munam Ali Shah, Mohammad Abdulaziz Al-Naeem, and Carsten Maple [9] examined a novel and dangerous attack vector: the use of automated web processes to facilitate infection. They demonstrated that ransomware could spread laterally via social networks using legitimate tools like the Selenium Web Automation Tool. In their experiments, an attacker sends a malicious link that, when clicked, executes a hidden script that downloads a backdoor. This backdoor then utilises the victim's browser automation to post the malicious link to the victim's social network friends, creating a viral infection loop. Achieving an 85% success rate in head-on testing, this research highlights the necessity for defence mechanisms that can detect unauthorised automation and script execution. This reinforces the need for the behavioural monitoring of process creation events (e.g., a browser spawning a command shell or an unknown automation script), which is a core component of our proposed Sysmon configuration.

2.5 Machine Learning in Ransomware Detection

The change to machine learning has been adequately reported and instigated by the inefficiency of signature-based systems. In a survey conducted by Saleh Alzahrani, Yang Xiao, Sultan Asiri, Jianying Zheng, and Tieshan Li [10], the detection methods were divided into two categories, ML-based and non-ML-based, and it was concluded that ML is becoming more effective because it can learn patterns instead of signatures. Nevertheless, they have observed some difficulties in areas related to obfuscation and anti-debugging.

Further developed by Shivam Khurana [5], an ML-based model was introduced with a unique combination of three models: Self-Organizing Maps (SOM), Random Forest Classifier and Long Short-Term Memory (LSTM) networks. The work by Khurana is crucial as it states the detection accuracy of 93.0 per cent and a precision of 97.0 per cent. The SOM is applied in an unsupervised manner to detect anomalies, separate normal behaviour into clusters and indicate the deviations. Random Forest classifier is powerful in binary classification (malicious or benign), whereas the LSTM network is able to capture the temporal relationship of an attack sequence. The given trio of models is the one that the current study uses, as they confirm the findings of Khurana in the wider context of SIEM.

Furthermore, research by Mazen Gazzan, Bader Alobaywi, Mohammed Almutairi, and Frederick T. Sheldon [4] on Polymorphic Ransomware Deep Learning (PRDL) introduced Bi-Gradual Minimax Generative Adversarial Networks (BGM- GAN) to generate synthetic attack patterns. This approach addresses the issue of data scarcity by creating realistic synthetic ransomware samples to train detection models, achieving 98.6% accuracy. While highly effective, such deep learning models can be computationally expensive; thus, our study seeks a balance by combining efficient classifiers (Random Forest) with sequence learners (LSTM) suitable for real-time SIEM ingestion. Kanugu Sandya, Gattu Ashitha, Bodasu Sharath Kumar, Bussa Varun Kumar, and Thatichettu Bharadwaj [11] came up with a new feature-coding technique of Portable

Executables (PE) files in the field of executable analysis. They obtained an accuracy of 94.83 00: Random Forest classifier with min -max normalisation + ordinal encoding, trained on over 138,000 files, extracting features of static headers.

2.6 Mitigation, Recovery, and Policy

The technical interventions in the current state of cybersecurity discourse are often considered as situations depending on a larger recuperative and control framework. At the level of mitigation, Kosuke Higuchi and Ryotaro Kobayashi [14] developed the Real-time Open-File Backup System (ROFBS). ROFBS uses Extended Berkeley Packet Filter (eBPF) to intercept kernel functions linked to file-open operations to address the so-called detection gap, or the duration of time that has elapsed between the infection with ransomware and its detection. Accordingly, ROFBS triggers a real-time backup of a file the second a suspicious process accesses it, thus surpassing the ransomware encryption process. The design ensures that, despite the latency of a number of seconds by the detection mechanisms, the background data remains intact and can be accessed. In terms of policy actions, Roxana Radu [15] conducted a comparative analysis of Australian, Costa Rican, French and Singaporean responses. Her study emphasises that the world has shifted to centralised systems of policy and the imposition of compulsory incidence reporting. An example is that Australia has established active cyber-defensive stances, with offensive cyber capabilities to disable criminal networks, whereas Costa Rica has emphasised cyber awareness efforts following major cyber-attacks on the country, as shown in Fig. 3.

2.7 Deep Learning on Dynamic Ransomware Detection

G. M. Sathyaseelan, Y. Rajendra Babu, Shankar Das Boddu, N Reshma, S. P. Santhoshkumar and D. Vikram [18] introduce a ransomware identification system based on AI, which utilises the Long Short-Term Memory (LSTM) network in an attempt to overcome the limitations of the traditional signature-based mechanisms in the face of polymorphic threats [18]. The suggested methodology is focused on behavioural analytics since it monitors file operations sequences, including create, read, write, rename, and delete operations. The model takes in both categories of embedded features (event types and file extensions) and regularised numbers (file entropy, processed bytes and the file time deltas) to find malicious patterns. The LSTM architecture can detect the use of this temporal structure to associate file changes with normal, suspicious, and ransomware activity with a synthetic dataset that has been trained on the environment to simulate normal, suspicious, and ransomware-related activities. G. M. Sathyaseelan, Y. Rajendra Babu, Shankar Das Boddu, N Reshma, S. P. Santhoshkumar and D. Vikram [18] reported that their experimental performance is exceptional, with the model reaching 100% accuracy, precision, and recall on both test data already in the twelfth training epoch. This method is much better than conventional machine learning algorithms, with a relative improvement in relation to Random Forest (94.63) and Support Vector Machine (91.83) [18]. Finally, this LSTM approach is a highly precise, scalable and real-time defence mechanism that is able to detect ransomware at an early stage of its execution and prevent the serious encryption of data and its loss.

2.8 Enhancing Endpoint ML Detectors, Resistant to Adversarial Evasion

Lightweight machine-learned models introduced by Vahid Heydari and Kofi Nyarko [19] are aimed at detecting ransomware on consumer devices and Internet-of-Things devices and indicate their fragility, which is intrinsically related. Such models are very susceptible to adversarial evasion, where the opponents apply small, realistic, and structurally appropriate changes to the feature representation of a file to cause a change in the detector. The authors suggest a cost-conscious hardening loop, named Iterative Adversarial Retraining (IAR), to deal with this vulnerability. IAR allows an attacker to run selective evasions with a change budget of size (K) and a cumulative cost budget of size (C). Effective evasions are then added to the training data, which leads to retraining the models.

A CPU-based Random Forest (RF) classifier trained on the characteristics of a stationary Portable Executable (PE) in a grey-box threat model has been used as the methodology. Using the Kaggle Ransomware Detection Dataset, the looping process forces the attacker to find budget-friendly changes that reduce the ransomware score. When the RF model is circumvented successfully, evasions related to it are back-integrated into an expanded training pool to repeat the process. It has been shown that a significant and rapid increase in robustness occurs. The original model is also notably weak: despite a Clean-data Accuracy of about 99.9 0 -1 Round 1 Attack Success Rate (ASR) is 0.9996, the Robust Accuracy (RAA) is low, only at 0.432. In comparison, hardening of the system is fastened through IAR. The 3-round increase in the RAA and a fall in the ASR to 0.9715 and 0.0459, respectively. After the fourth round, an almost complete suppression of the attacks has been achieved with the RAA of 0.9964 and an ASR of 0.0016. These results show that an ASR near zero can be attained with realistic deployment limitations of C35 and K10. Further, the RF model reaches a CPU in approximately 0.027 minutes per round, asserting that the technique is a sound hardening strategy that is CPU-reachable to resource-constrained endpoints.

2.9 Zero Trust and Cryptographic Storage to Ransomware Resilience

Jessica Theodore, Dawit Kasy, and Chetan Jaiswal [20] present Ransom Sentinel (RaSe), a multi-layered and conceptual defensive framework that would reduce the rising threat of ransomware, specifically in critical healthcare systems. RaSe is designed based on the objectives of Durability, Accessibility, and Loss Prevention (DAL), and it serves as a product combining simulation-based offensive attacks with offensive attacks committed and implemented with efficient cryptographic protection, and this serves to guard sensitive information [20].

The RaSe framework consists of two major layers. The User-Interfaces Layer takes care of clinician touch-points using a Zero Trust architecture with contextual routing to keep exposure to a minimum and makes recent records available in a RAM-only "LiveCache). Security-core Layer is a cryptographic implementation, which performs all cryptographic operations; implements Shamir's Secret Sharing (SSS) to manage clients' keys distributed across devices, Reed- Solomon (4+2) erasure coding to ensure fault-tolerant storage of various data.

RaSe uses the Entropy Calculation of Shannon to track the randomness of file contents in detecting any sign of a ransomware attack. When the data operations register unusually high entropy scores within the range of 7.5 to 8.0, the files are indicated by the system as possible encrypted files. To prevent permanent loss of data, RaSe introduces an immutable logging module named BlackBox that captures all changes that happen to the database before and after to meet the HIPAA compliance and forensic integrity. Should the primary database get compromised or encrypted, the framework can completely reconstruct the original records based on the append-only change logs without having to pay to get the ransom. In conclusion, the framework proves that anomaly detection, Zero Trust policy, and cryptographic storage are effective in isolating and neutralising ransomware attacks and ensuring the continuity of operations.

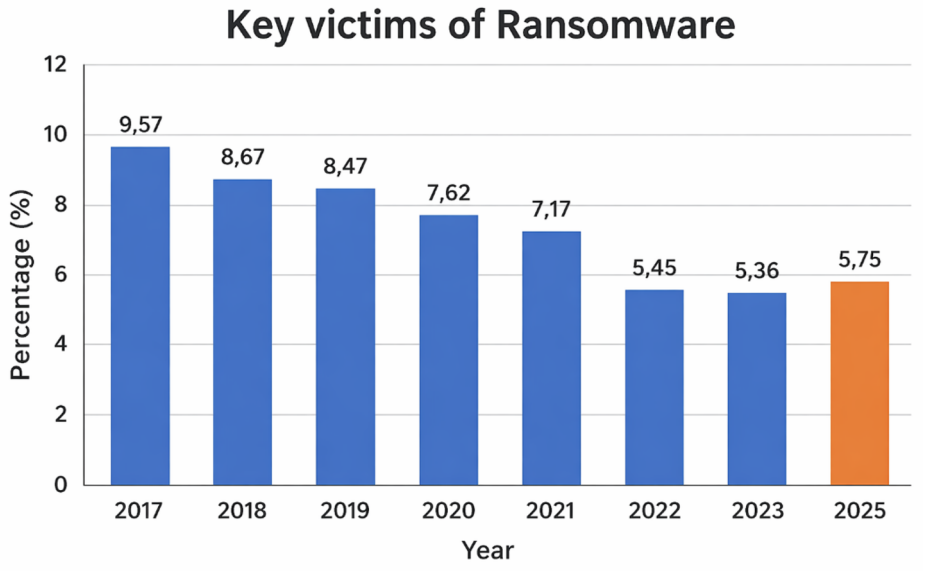

Fig. 3. From 2017 to 2025, countries with the highest share of users attacked with ransomware.

3. Theoretical Framework

This section explains the theoretical background of the technologies and ideas that are integrated into the proposed framework.

3.1 Mitigation, Recovery, and Policy

Understanding the ransomware lifecycle is a prerequisite to detection. As illustrated in standard models and confirmed by the literature [2], the lifecycle typically follows a sequence, as shown in Fig. 2:

- **Reconnaissance & Delivery**: Often via phishing (as simulated in this study), RDP brute force, or exploiting web vulnerabilities.
- **Installation & Command and Control (C2):** The malware establishes persistence (e.g., Registry Run keys) and contacts the attacker's server to exchange encryption keys.
- **Lateral Movement:** The malware scans the Network to infect other hosts, using protocols like SMB.
- **Encryption:** The payload iterates through file systems, encrypting target file types.
- Extortion: The ransom note is displayed.

The Encryption stage is the stage where the data cannot be recovered once it has been lost without backups. As a result, it should be detected during the Delivery, Installation or C2 phases. This test aims at dependency oddities typical of these early phases (e.g., dysfunctional parent-child process relationships (e.g. Word spawning PowerShell)) or abnormally outbound network associations with hitherto unfamiliar IP addresses.

3.2 Machine Learning Architectures

The framework suggested follows a hybrid approach to machine learning and is able to capture different aspects of anomalous behaviour:

- Self-Organising Maps (SOM): SOM is an unsupervised neural network that maps high-dimensional data onto a low-dimensional (typically 2D) grid. In security, it is powerful for anomaly detection. It is trained on "normal" system behaviour. During operation, data points that map to neurons far from the established "normal" clusters (large Euclidean distance from the Best Matching Unit or BMU) are flagged as anomalies. This is crucial for detecting zero-day threats that deviate from the norm but lack known signatures [5].
- •Random Forest Classifier: This is an ensemble learning method that operates by constructing a multitude of decision trees at training time. For classification tasks, it outputs the class that is the mode of the classes of the individual trees. It is robust against overfitting and excels at handling large datasets with higher dimensionality, making it suitable for binary classification (Malicious vs. Benign) based on extracted log features [11].
- •Long Short-Term Memory (LSTM): LSTM is a type of Recurrent Neural Network (RNN) capable of learning order dependence in sequence prediction problems. Ransomware attacks are not instantaneous; they are sequences of events (e.g., download → execute → connect→ encrypt). LSTM networks are theoretically suited to identify these temporal patterns, effectively maintaining a "memory" of previous events to contextualise current actions [5].

3.3 Mitigation, Recovery, and Policy

SIEM systems are systems that collect data in the form of logs, which are gathered by different sources to provide a complete view of an operational environment. A highly successful open SIEM solution is the Elastic Stack, which includes Elastic Search to store and retrieve data and Kibana to visualise data. In this stack, Sigma rules are used

in detection. Sigma is a generic and open-signature format that allows analysts to describe relevant log events concisely and understandably. These rules are then converted to the query language of the target SIEM, in this case, the Kibana Query Language (KQL), thus raising alerts. This methodology allows constructing a complex, multi-conditional alert (such as "Alert when Process X makes a network connection to IP Y and modifies Registry Key Z) [7].

4. Methodology

The research methodology was designed to create a robust, multi-layered, and adaptive system for Ransomware Threat Detection and Mitigation. It uniquely combines predictive Machine Learning with forensic validation to generate high-fidelity detection rules, as shown in Fig.4.

4.1 Overall Architecture

The system employs a distributed architecture tailored for a Small and Medium-sized Enterprise (SME) simulation. It consists of three dedicated Virtual Machines (VMs):

SIEM Host (Linux VM): This node hosts the core infrastructure, Elasticsearch and Kibana. It serves as the high-speed, scalable data repository and the analytical interface for visualisation, alert management, and rule creation.

Fleet Management Server (Linux VM): This server manages the deployment, configuration, and continuous health monitoring of the Elastic Agents deployed across the Network, ensuring consistent policy application.

Endpoint Host (Windows 10 VM): This is the machine that plays the role of the "Victim" one. It is the secured system on which the Elastic Agent is deployed, such that it simulates a typical employee workstation

4.2 Data Ingestion and Telemetry

The framework is based on data visibility. Windows 10 Endpoint is set to provide granular telemetry, which is relayed in real-time to the SIEM host with the help of the Elastic Agent. The particular sources of logs are:

Sysmon Logs: It offers detailed visibility into process creation, network connections, file creation and registry modification. Sysmon was also set up to record hashes of every created image.

Windows Event Logs: The standard security and system events are capturing (e.g. 4624 is tracking the logon, 4688 is tracking process creation).

PowerShell Logs: Script-block execution and command-line parameters capture used to detect file-less malware attacks, and also Living off the Land (LoTL) attacks.

Network Logs: Tracking the packet flow and the destination IP address.

4.3 Mitigation, Recovery, and Policy

The obtained logs are subject to feature engineering in order to derive significant behavioural characteristics (e.g., process names, entropy metrics of file writes, and frequency of network connections). Three machine learning models [5] are then used to employ these engineered features:
Self-Organising maps (SOM): These are behaviour clustering techniques that are used for normal behaviour to determine unusual behaviour.
Random Forest classifiers classify discrete events as malicious or benign, which is based on known datasets.
Long Short-term Memory (LSTM) networks examine the history of logs to identify the history of attacks.
The methodology includes these models in the pipeline of SIEM, which enables them to rate the incoming events. Kibana alerts are generated when events satisfy a set threshold of threat score. In the case of predictive model training, 45,000 log events were taken together as a consolidated set and divided into a training subset (80%) and a testing subset (20%). Also, a five-fold cross-validation process was used to increase the stability of the model and reduce overfitting. Machine-learning algorithms, including Self-Organising maps, Random Forest and Long Short Term Memory network, were run in default library configurations, and the architectural framework and parameter baselines were manipulated to accept the ingested SIEM telemetry, based on Khurana [5].

4.4 Forensic and Reverse Engineering Process

To achieve a state of having the system rely on correct Indicators of Compromise (IOCs), in fig. 4 the model clearly specifies, including a forensic feedback loop:
Attack Simulation: In order to monitor the malicious behaviour securely, the samples of the malicious ransomware of the Emotet, IcedID and Trickbot were downloaded from the Malwarebazaar repository and launched in the virtual Network in a tight isolation. The media used to simulate the environment was a Windows 10 endpoint that acts as a victim machine and a Kali Linux instance that acts as a Command and Control (C2) server. The implementation of the methodology resembled the vectors of a real-world phishing attack: a spoofed email was used to drop a malicious JavaScript dropper, which in turn used a PowerShell session to get the downloaded ransomware payload, execute it, and make it persistent.
Data Acquisition: To obtain memory (RAM) and disk images of the compromised Windows 10 virtual machine, FTK Imager (Forensic Toolkit Imager) is used to capture data right after the execution. This makes it easier to detect artefacts of transient memory (e.g. unpacked payloads) and file-system changes, which are not regularly logged and vanish with a reboot.
Reverse Engineering: The captured malware binaries are disassembled with the help of IDA Pro (Interactive Disassembler Professional). The researchers can understand the code of the malware (e.g., CryptAcquireContext, CryptGenKey) used as API calls, inner logic, and communication mechanisms with C2 by examining the assembly code. This detailed technical knowledge makes sure that the custom detection rules are being

targeted at the most basic and least malleable mechanisms of the malware and not at the surface-level indicators.

Rule Creation: The results obtained using FTK Imager and IDA Pro inform the development of tailor-made Sigma rules. All these are then translated into KQL queries in order to form high-fidelity alerts in Kibana.

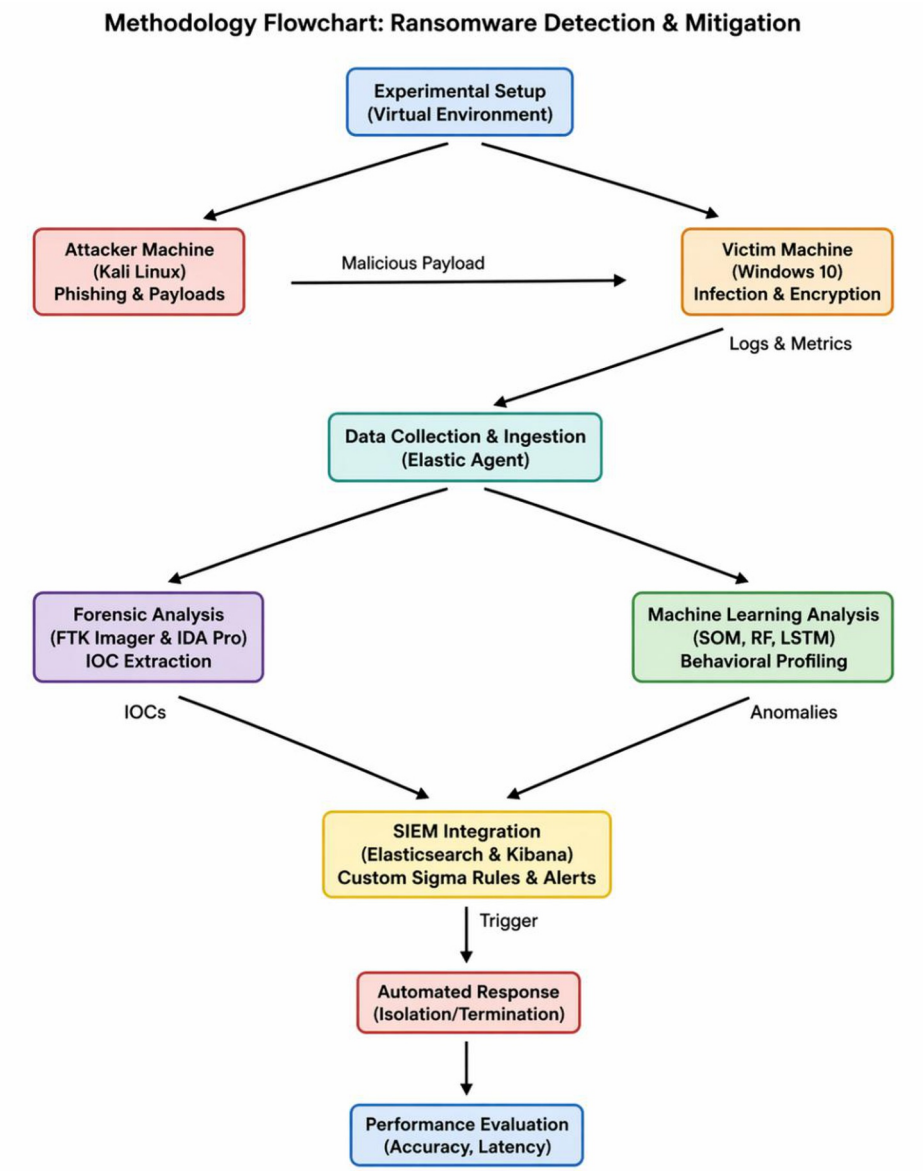

Fig. 4. Research Methodology

4.5 Automated Response

The Security Information and Event Management (SIEM) system has been programmed to have an automated response mechanism. When there is a high-severity alert (i.e. a violation of a machine-learning threshold or a Sigma rule) the system will kick off containment processes, i.e., isolate the infected host of the corporate Network and terminate the process detected to be an anomaly. This is necessary in aggressive defence, especially in a close detection-time amplitude within which the encryption takes place.

5. Result and Discussion

To determine the effectiveness of the SIEM-ML integration, the application of the proposed framework was evaluated by undergoing experimental procedures that imitated a realistic ransomware attack chain to measure the efficiency of the proposed framework.

5.1 Experimental Scenario

The experiment was created to demonstrate a simulated attack of Phishing-to-Ransomware, and, thus, one of the most common carriers of infection in modern cybersecurity situations.

Setup: Configuration A Kali Linux system was deployed both as a Command and Control (C2) server and a malicious payload generation server. The victim platform was a Windows 10 VM that was run under controlled experimental conditions.

Attack Chain: The assailant used a spoofed email whose subject was "Urgent: Overdue Invoice" and which included a malicious JavaScript code (invoicescan.js). When this script was executed, it also called a PowerShell session, which downloaded the ransomware executable (cryptpayload.exe) to an external server. The payload subsequently ran, created a persistence condition by writing entries in the registry, and launched encryption of documents found in the user's Documents folder, adding the extension of locked before these encrypted files.

To allow strict machine-learning training of the models and ensure reproducibility, continuous system monitoring was gathered during the entirety of the simulation lifecycle of 60 minutes, and raw telemetry data was collected from the Elasticsearch repository. The resulting collection produced a validated dataset that had 45,000 unique log events. In this corpus, 8,500 records corresponded to malicious telemetry, which reflects the execution lifecycles of the simulated payloads, which were mainly represented as process injections, Command and Control (C2) network connections, as well as file modifications at high-frequency rates during the encryption phase. The other 36,500 sets were benign baseline operations of Windows 10, its background services and the normal user applications before a simulated breach. The merged dataset was divided into an 80 per cent training set and a 20 per cent testing set to test the predictive algorithms. Additionally, to authenticate the stability of the model and reduce overfitting beyond a reasonable doubt, a 5-fold cross-validation scheme was used when training the model.

5.2 Forensic and SIEM Analysis Results

SIEM Detection: The Elastic Stack imported and correlated the attack logs. There was a sharp EKG alert frequency spike as visualised in the dashboard of the system. The first spike (T = 20 minutes) was the time that the malicious email was delivered, which was detected with the help of a special rule that indicated emails with .js as attachments (event.category: "email" and attachment. extension: "js"). A second and prolonged spike (T = 26 minutes) was associated with the high-volume file update events at the encryption stage, as shown in Fig. 5.

Forensic Validation: By applying memory forensic analysis with FTK Imager on the memory dump file memdump.mem, the examiner confirmed that there was, in fact, the executable cryptpayload.exe. During a static analysis of the binary using the IDA Pro disclosures of the Windows cryptographic API functions CryptAcquireContext and CryptGenKey, and hard-coded string literals of the IP address of the command-and-control server, these results provided proven indicators of compromise (IOCs) to be included in detection rule sets.

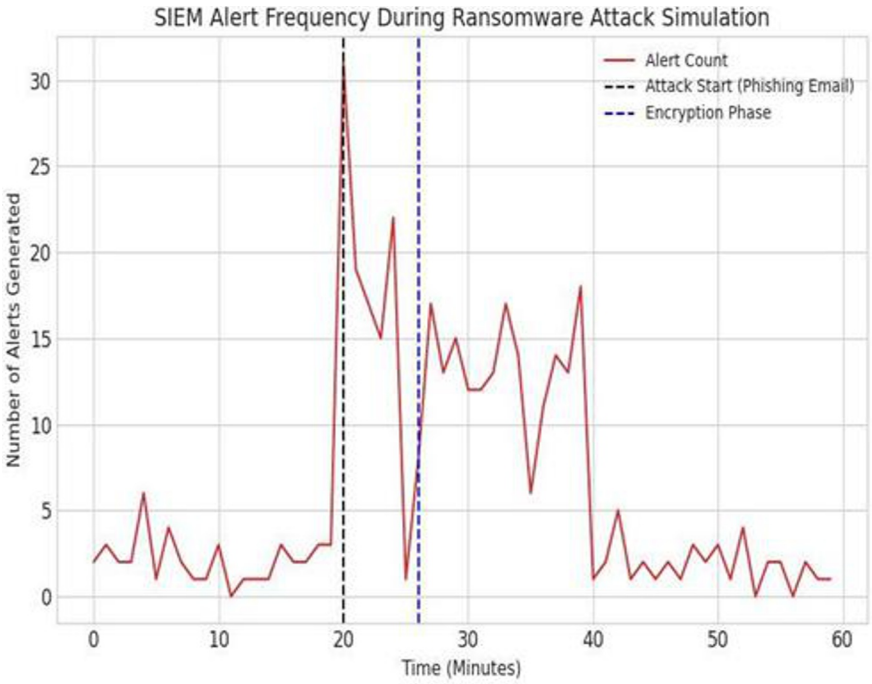

Fig. 5: SIEM Alert Frequency during Ransomware Attack.

5.3 Detection Latency

In Resilience is a function of speed. The system demonstrated variable latency across attack stages, visualised the results in Fig. 7:

Email Delivery: Detected in 2 seconds, proving the efficacy of the email gateway integration.
Encryption Phase: Detected in 120 seconds (2 minutes).
Discussion on Latency: Although the first access was determined almost immediately, the two-minute latency in the encryption cycle is a large vulnerability window. The latency mainly has to do with the time it takes for the security information and event management system (SIEM) to accumulate enough instances of the events of creating files and renaming files to reach the high-severity threshold. Although the time lag of detecting 120 seconds in a comparatively time-sensitive manual testing method may be considered relatively fast, it still allows sufficient time to lose consequences. Subsequently, these findings underscore the importance of an automated response mechanism; it is not feasible to have a manual response in the 120-second timeframe, where an automated kill-switch may achieve remediation at a much higher speed.

Fig. 6. SIEM Email Detection

5.4 Comparative Performance Metrics

In To compare its improvements, the suggested integrated methodology was benchmarked with the traditional signature-based detection mechanisms, as well as other more heuristic approaches, summarised in Table 1. Comparative baseline methodologies were defined and set as follows:

Signature-based Detection: It uses procedural Indicator of Compromise (IOC) blacklists using uniform hash compacting to the ingested Sysmon and Windows Event logs, coupled with conventional Indicator of Compromise; it is not dynamically adaptable.
Heuristic Analysis: There was the deployment of predetermined and fixed sets of rules, which were used to detect generic suspicious API sequences and structural file anomalies, but this is not a continuous learning mechanism.
Behaviour Analysis: Basic deviations of system-state parameters (e.g. high frequency file renaming and registry write-ups) in traditional threshold-based monitoring existed regardless of predictive machine-learning weightings.
Generic ML: The standard Support Vector Machine was used as a single classification model - an unoptimised Decision Tree was applied to the extracted log features. This architecture lacks the temporal-sequence sensitivity of the Long Short-Term Memory (LSTM) networks or the task State of the Art unsupervised anomaly-clustering features of the Self-organizing Maps (SOM).
The suggested framework achieved an accuracy of 92.3 due to detection as well as a precision of 95.7. It is remarkable that the rate of false positives (FPR) was extremely low, 0.02, and the rate of false negatives (FNR) was 0.05, which means that the system was not often unable to detect real threats. This result is mostly due to the ability of the Long Short-Term Memory (LSTM) network to capture sequential relationships, as found in the attack, e.g. the precise order of API calls, and the strength of the Random Forest estimator to classify the features extracted.
The nearest competitor was the traditional signature-based paradigm with 82.0 and 88.0 as its accuracy/precision, respectively. The 11% accuracy improvement as observed highlights the ability of the proposed model to identify a polymorphic or otherwise modified form of ransomware that easily avoids detection by the static signature check. False Positive Rate is among the most important operational measures of any security operations centre. As evaluated on a test subset of our extremely imbalanced dataset, containing 36,500 benign and 8500 malicious events, the proposed framework had a false positive rate of 0.02. At the same time, it had a true positive rate of 0.95 and a false negative rate of 0.05. The confusion matrix underlying these metrics confirms that the system can identify true devastating threats with high success as well as produce a negligible number of false alarms compared to the functioning of any particular system at the baseline, as shown in Fig. 6.

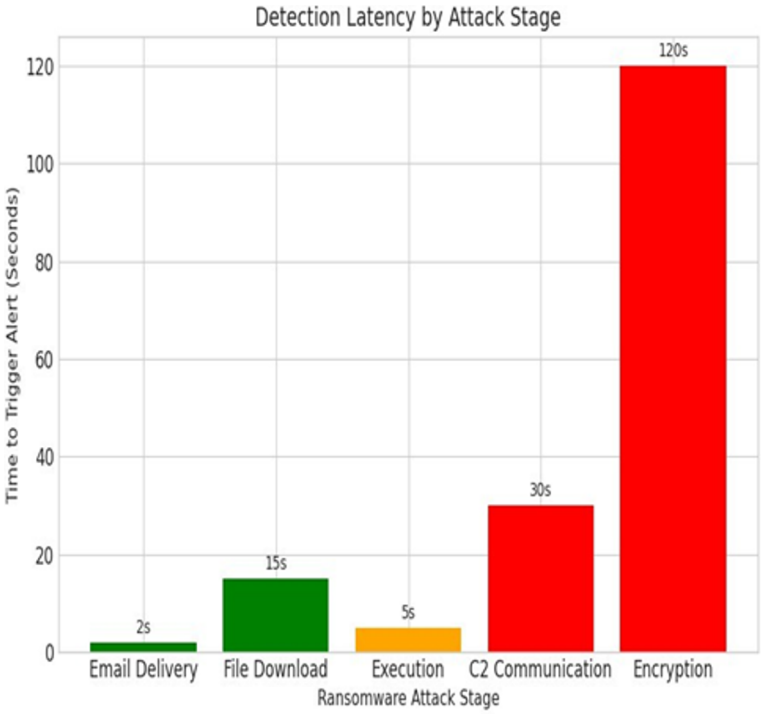

Fig. 7. Detection latency by attack stage.

Table 1: Comparative Performance Metrics

Method	TPR	FPR	Acc	FNR	Prec
Proposed Method	0.95	0.02	0.93	0.05	0.97
Signature-based	0.85	0.12	0.82	0.15	0.88
Heuristic Analysis	0.75	0.08	0.74	0.25	N/A
Behavior Analysis	0.80	0.15	0.78	0.20	N/A
Generic ML	0.70	0.18	0.68	0.30	N/A

5.5 Efficacy of Custom Sigma Rules

A set of rules relying on Sigma, developed as a result of forensic investigation of Emotet, IcedID, and Trickbot, showed a high degree of efficacy against the designated malware families:

IcedID: Recognised as rules that identify web-injection parameters and PHP requests that are being made towards external servers.

Trickbot: Detected through lateral movement attempts on ports 445 (SMB) and 3389 (RDP), and PowerShell commands.

Emotet: Identified based on the correlated execution of anomalous PowerShell-based actions with the modification of the registry, with the effect of file downloads whose source IP addresses are known to be malicious.
Multi-indicator matching (e.g. File Path + Network IP + Process Name) was recently significant in achieving the cited accuracy and therefore avoiding the false positive classification of regular administrative processes, which otherwise would appear similar when looked at alone.

6. Conclusion and Future Scope

6.1 Conclusion

This study was able to propose and empirically test a holistic model of detection and prevention of ransomware threats. The long short-term memory networks and random forest classifiers, along with the self-organising maps in terms of their predictive potential, will be combined with Elastic SIEM to overcome the limitations inherent in single, signature-based protections. Experimental results establish that the integrated method has better performance, with an accuracy of 93.0, a precision of 97.0 and an incredibly low false-positive rate. Defence is an argumentative notion that implies an integrational perspective merging network-traffic analysis, where data in motion is examined; endpoint monitoring, where process injection is recognised; and strong forensic analysis where adversarial intent is clarified. The use of such forensic tools as FTkImager and IDAPro can guarantee that detection regulations are based on the strict technical reality, as they pursue the unchanging actions of malware (not its changing appearances). Furthermore, the built-in automated response ability is a novel solution to the paramount remediation time problem since it might cut off encryption steps even before extensive data destruction sets in. To sum up, even though large-scale, real-world implementation remains a field of future research, the Integrated ElasticStack with the ML-integrated framework provides a highly precise and proactive basis of localised network security against the growing complexity of the modern ransomware threat.

6.2 Future Scope

Although the existing structure is strong, the arena of adversaries is dynamic, and the competition between attackers and defenders still goes on. The further studies can be directed at some main directions:
Adversarial Machine Learning (AML): According to literature on the topic of RansomAI, attackers are engaging in using machine learning to bypass machine learning defences. Next-generation work should, therefore, consider some form of training the detection models with adversarial example inputs carefully designed to confuse the classifier to improve resiliency to evasion strategies [16].
Cross-Platform Expansion: Ransomware is also targeting Linux servers, which are frequently the workhorse of cloud architecture, and macOS. A logical next step in extending agent deployment and rules to non-Windows platforms is needed.

Explainable AI (XAI): Deep learning models [18] often function as so-called black boxes, thus hindering the possibilities of an analyst to understand the reasons behind a specific flagging. The adoption of the XAI frameworks, among which are Local Interpretable Model-agnostic Explanations (LIME) and Shapley Additive exPlanations (SHAP), would also allow the system to provide interpretable explanations (says flagged due to Sequence X followed up by Y).
Automated Rule Generation: Researching how generative artificial intelligence models or reinforcement learning frameworks can be used to autonomously learn new patterns of threats and to learn new Sigma/KQL rules or update existing ones may help significantly to reduce the reliance on the manual technique of threat hunting and rule writing, eventually helping create a truly self-mending defence architecture

References

1. Christopher Jun Wen Chew, Vimal Kumar, Panos Patros, and Robi Malik, "Real-time system call-based ransomware detection," International Journal of Information Security,2024.
2. Gourav Nagar, "The Evolution of Ransomware: Tactics, Techniques, and Mitigation Strategies," International Journal of Scientific Research and Management (IJSRM), 2024.
3. Ali Ahmed Mohammed Ali Alwashali, Nor Azlina Abd Rahman, and Noris Ismail, "A Survey of Ransomware as a Service (RaaS) and Methods to Mitigate the Attack," in 2021 14th International Conference on Developments in eSystems Engineering (DeSE), 2021.
4. Mazen Gazzan, Bader Alobaywi, Mohammed Almutairi, and Frederick T. Sheldon, "A Deep Learning Framework for Enhanced Detection of Polymorphic Ransomware," 2025.
5. Shivam Khurana, "Ransomware Threat Detection and Mitigation using Machine Learning Models," 2023.
6. Raul Reinosa Simon, Clara I. Valero, Jose A. Martınez Cadenas, Elisa R. Heymann, Ignacio Lacalle, Barton P. Miller, and Carlos E. Palau, "Empirical Analysis and Practical Assessment of Ransomware Attacks to Data in Motion," in 2024 IEEE International Conference on Cyber Security and Resilience (CSR), 2024.
7. B. P. Rayabandi, "Integration of Elastic Search and Kibana SIEM for Malware Detection," 2024.
8. Fathi Kamil Mohad Zainudin, Izzatul Hazirah Ishak, Sharifuddin Sulaman, Farah Ramlee, Nur Sarah Jamaludin, and Shuaib Chantando, "Malware Discovery using Lebahnet Technology," OIC-CERT Journal of Cyber Security, vol. 2, no. 1, Feb. 2020.
9. Muhammad Usman Rana, Munam Ali Shah, Mohammad Abdulaziz Al-Naeem, and Carsten Maple, "Ransomware Attacks in Cyber-Physical Systems: Countermeasure of Attack Vectors Through Automated Web Defenses," 2024.
10. Saleh Alzahrani, Yang Xiao, Sultan Asiri, Jianying Zheng, and Tieshan Li, "A Survey of Ransomware Detection Methods," 2025.
11. Kanugu Sandya, Gattu Ashitha, Bodasu Sharath Kumar, Bussa Varun Kumar, and Thatichettu Bharadwaj, "Enhancing Malware Detection using Novel Feature Encoding in Machine Learning," in Proceedings of the 6th International Conference on Data Intelligence and Cognitive Informatics (ICDICI-2025), 2025.

12. Ali Mehrban and Shirin Karimi Geransayeh, "Ransomware threat mitigation through network traffic analysis and machine learning techniques," 2024.
13. Vhuhwavho Mokoma and Avinash Singh, "RanViz: Ransomware Visualization and Classification Based on Time-Series Categorical Representation of API Calls," 2025.
14. Kosuke Higuchi and Ryotaro Kobayashi, "Real-time open-file backup system with machine-learning detection model for ransomware," International Journal of Information Security, 2025.
15. Roxana Radu, "Countering Ransomware: Government Responses in a Comparative Perspective," in CyCon 2025: The Next Step, 17th International Conference on Cyber Conflict, Tallinn, 2025.
16. Jan von der Assen, Alberto Huertas Celdran, Janik Luechinger, Pedro Miguel Sanchez Sanchez, Gerome Bovet, Gregorio Martınez Perez and Burkhard Stiller, "RansomAI: AI-powered ransomware for stealthy encryption," in 2023 IEEE Global Commun. Conf. (GLOBECOM), 2023.
17. Manuj Aggarwal, "Ransomware attack: An evolving targeted threat," in Proc. 2023 Int. Conf. Comput. Commun. Netw. Technol. (ICCCNT), 2023, pp. 1–7.
18. G. M. Sathyaseelan, Y. Rajendra Babu, Shankar Das Boddu, N Reshma, S. P. Santhoshkumar and D. Vikram., "Next-Generation Ransomware Defence Using LSTM-Powered Behavioural Analytics," 2026 9th International Conference on Computational Intelligence in Data Science (ICCIDS).
19. Vahid Heydari, and Kofi Nyarko, "IAR: Multi-Round Cost-Aware Hardening for Endpoint Ransomware Detection," 2026 IEEE 23rd Consumer Communications & Networking Conference (CCNC).
20. J. Theodore, D. Kasy, and C. Jaiswal, "Ransom Sentinel (RaSe): A Study Towards Defense Against Ransom Jessica Theodore, Dawit Kasy, and Chetan Jaiswal ware," 2026 IEEE 16th Annual Computing and Communication Workshop and Conference (CCWC).
21. Mamoona Humayun, NZ Jhanjhi, Ahmed Alsayat, Vasaki Ponnusamy, "Internet of things and ransomware: Evolution, mitigation and prevention," Egypt. Informat. J., vol. 22, no. 1, pp. 105–117, Mar. 2021.
22. Maxat Akbanov, Vassilios G. Vassilakis and Michael D. Logothetis, "WannaCry ransomware: Analysis of infection, persistence, recovery prevention and propagation mechanisms," J. Telecommun. Inf. Technol., 2019.