

Real-Time Ransomware Detection and Automated Response Using Behavioral Analysis

1st Ayush Khatana

2nd M. Trivikram Rao

3rd M Jeyaselvi

Dept. Networking and Communications

SRM KTR

Chennai, India

ak6730@srmist.edu.in

mr9117@srmist.edu.in

jeyaselm@srmist.edu.in

Abstract—Ransomware attacks have emerged as one of the most critical cybersecurity threats, causing rapid file encryption and significant data loss. Traditional detection systems rely heavily on signature-based techniques, which fail to identify zero-day and behaviorally evolving ransomware. This paper proposes a real-time ransomware detection and mitigation system that combines rule-based behavioral monitoring with machine learning-driven analysis. The system continuously monitors file system activities and detects suspicious patterns such as rapid file encryption and abnormal process behavior. Upon detection, it immediately triggers alerts and terminates malicious processes to prevent further damage. Additionally, a behavior-based dataset is dynamically generated during execution, capturing process-level features for future machine learning model training. This hybrid approach ensures fast response through deterministic rules while enabling adaptability through data-driven intelligence. Experimental evaluation demonstrates that the proposed system achieves near real-time detection with minimal delay, effectively limiting file encryption and improving system resilience against ransomware attacks.

Index Terms—Ransomware Detection, Behavioral Analysis, Real-Time Monitoring, Machine Learning, Malware Defense, Process Termination, Zero-Day Attacks, Endpoint Security.

I. INTRODUCTION

Picture this: you arrive at work, open your laptop, and find that every file — contracts, databases, financial records — has been replaced with unreadable gibberish. A ransom note sits on your desktop demanding payment in cryptocurrency within 48 hours. This scenario plays out thousands of times a day worldwide, and the attackers keep getting faster and smarter. Modern ransomware can encrypt tens of thousands of files in under a minute, which means the window for detection and intervention is brutally narrow. Ransomware has evolved from a niche criminal experiment into a multi-billion-dollar global industry. What began as relatively simple encryption tools has transformed into sophisticated, modular malware-as-a-service operations. Groups like REvil, LockBit, and BlackCat have demonstrated that no organization is too large or too well-defended to be targeted. Hospitals, pipelines, universities, and governments have all fallen victim, with some attacks costing

hundreds of millions of dollars in downtime and recovery alone. The core challenge lies in detection speed. Conventional antivirus systems rely primarily on signature-based detection — they look for known malware patterns in files and memory. While this works well for previously catalogued threats, it fails completely against new ransomware variants and polymorphic strains that modify their own code to evade recognition. By the time a signature is developed and distributed for a new variant, the damage is long done. Behavioral detection tries to solve this by asking a different question. Instead of 'have we seen this before?', it asks 'does this look right?' If a process suddenly starts opening and modifying hundreds of files in rapid succession, changing their extensions, and consuming CPU at unusual rates, that pattern looks like ransomware — regardless of whether we have ever seen that specific malware family before.

In this paper, we propose a real-time ransomware detection and automated response system built on exactly this principle. The system monitors file system activity continuously and triggers an automated response the moment suspicious behavior crosses a defined threshold. It does not wait for a human to review an alert. It identifies the responsible process and kills it. In parallel, it collects structured behavioral data from running processes, laying the groundwork for future machine learning integration.

The main contributions of this work are: • A real-time ransomware detection mechanism based on continuous file system event monitoring. • An automated response module that terminates malicious processes with sub-second latency. • A lightweight behavioral dataset generation framework designed for future ML model training. • A hybrid architecture combining deterministic rule-based detection with data-driven adaptability. • A comparative analysis against existing detection methods demonstrating practical advantages.

II. RANSOMWARE THREAT LANDSCAPE

II. RANSOMWARE THREAT LANDSCAPE To build an effective defense, it helps to understand exactly what you're

defending against. Ransomware attacks typically unfold in distinct stages, each offering a narrow window for detection and response. Initial access is most commonly achieved through phishing emails containing malicious attachments or links, exploitation of unpatched vulnerabilities in public-facing services (particularly Remote Desktop Protocol), and increasingly through supply chain compromise. Once inside a network, ransomware operators may spend days or weeks establishing persistence, escalating privileges, and identifying the most valuable data before triggering encryption. The encryption phase itself is where most detection opportunities are either seized or lost. Modern ransomware uses hybrid encryption — typically AES-256 for file content and RSA-2048 or similar for encrypting the AES keys. This design means that without the operator’s private RSA key, decryption is computationally infeasible. The AES encryption of individual files is extremely fast, which is why the encryption phase can complete in seconds to minutes depending on the volume of data. What makes behavioral detection feasible is that this speed is also ransomware’s most visible fingerprint. No legitimate application opens thousands of file handles and rewrites their contents in rapid succession. File extensions change en masse. Entropy of file contents spikes dramatically. System calls associated with file I/O reach unusual frequencies. Each of these signals, individually weak, becomes a reliable indicator when observed together within a short time window. Contemporary ransomware families have begun adapting to behavioral defenses by deliberately slowing their encryption rate, encrypting only portions of files, or targeting less-monitored directories first. Our system’s threshold-based design and machine learning roadmap are specifically intended to remain adaptable as these evasion techniques evolve.

III. RELATED WORK

Ransomware detection has been an active research area for over a decade, with approaches broadly falling into three categories: signature-based, behavioral, and machine learning-based systems. Signature-based systems remain the foundation of most commercial antivirus products. They compare file hashes and byte sequences against databases of known malware. The approach is fast, reliable for known threats, and generates very few false positives — but it is fundamentally reactive. A signature cannot exist until a sample has been collected, analyzed, and catalogued, which means zero-day ransomware and novel variants will always evade these systems initially. Behavioral detection approaches have been widely explored as a complement to signatures. Continella et al. proposed ShieldFS, a filesystem-level defense that monitors I/O patterns and can roll back encrypted files using shadow copies [3]. While effective, shadow-copy approaches introduce storage overhead and do not prevent encryption from occurring — they only enable recovery afterward. Our system takes a more aggressive stance: terminating the attack before encryption can complete. Machine learning has shown considerable promise in malware classification. Random Forest and Support Vector Machine classifiers trained on static features such as API call

sequences and PE file headers have achieved high accuracy on known malware families [1][2]. However, most of these approaches operate offline on collected samples and do not provide real-time response. When applied to live traffic, the latency introduced by feature extraction and model inference can be significant. Deep learning approaches, including recurrent networks for sequential API call modeling and convolutional networks for binary visualization, have pushed classification accuracy further but tend to be computationally expensive for endpoint deployment [1]. Several studies have also explored network-level ransomware detection, identifying command-and-control traffic patterns, though network-level indicators typically appear after the encryption phase has begun. A consistent gap across the literature is the lack of systems that combine fast, deterministic detection with automated response and a path toward learned adaptability. Our work addresses this gap directly.

IV. PROBLEM STATEMENT

The core problem is not that ransomware is undetectable. It is that most detection happens too late. By the time a signature-based system flags a threat, the encryption process has typically already worked through hundreds or thousands of files. Recovery from backups — where they exist and are uncorrupted — can take days or weeks. In healthcare or critical infrastructure contexts, that downtime can be life-threatening. Existing behavioral systems improve on this but introduce their own problems. Many require cloud connectivity for real-time analysis, creating latency and privacy concerns. Others impose significant CPU overhead that degrades the performance of the protected endpoint. Several require manual intervention to respond to alerts, defeating the purpose of real-time detection. There is also the challenge of zero-day detection. Machine learning models trained on historical datasets perform well on known ransomware families but can struggle with novel variants that exhibit subtly different behavior patterns. A purely rule-based system may be more reliable in the short term but less adaptable as attackers evolve their techniques. What is needed is a system that is fast enough to detect an attack while only a handful of files have been encrypted, lightweight enough to run continuously on an endpoint without perceptible performance impact, automated enough to respond without human intervention, and designed from the ground up with a clear path toward machine learning integration for long-term adaptability.

V. PROPOSED SYSTEM

Rather than waiting for a ransomware signature to appear in a database somewhere, our system watches what processes actually do. The moment file modifications start happening faster than any legitimate application would cause, the system does not just log it — it pulls the plug on the offending process. The system operates as a continuous background service, monitoring a designated target directory for file system events including creation, modification, deletion, and renaming. These event types collectively represent the I/O signature

of an active ransomware encryption process. By tracking event frequency within a sliding time window, the system can detect the characteristic burst pattern of ransomware without needing to inspect file contents or match signatures. Detection triggers an immediate automated response. The system identifies which process generated the suspicious file events — using process ID attribution from the operating system — and terminates that process via a system call. This happens in milliseconds, before the process has a chance to encrypt more than a small number of files. A user-facing alert is generated simultaneously, and all relevant events are logged for post-incident analysis. Running in parallel with the detection engine is a behavioral data collection module. For every monitored process, the system periodically samples CPU utilization, memory consumption, number of open file descriptors, and process uptime. These measurements are recorded alongside a label (benign or malicious) and stored in a structured dataset. This dataset is the foundation for the machine learning extensions we plan to develop in future work.

VI. SYSTEM WORKFLOW

The system’s operation can be thought of as a continuous loop with six distinct phases:

- **Monitoring Phase:** The system registers an event listener on the target directory using an OS-level inotify or equivalent mechanism. Every file creation, modification, rename, or deletion event is captured with a timestamp and the responsible process ID. This monitoring runs at the kernel level, introducing negligible overhead.
- **Detection Phase:** A sliding time window of configurable duration (default: 5 seconds) accumulates file event counts. The detection engine evaluates whether the event frequency within this window exceeds the predefined threshold. The threshold is calibrated to be above the normal activity of any legitimate application while below the typical encryption rate of known ransomware families.
- **Decision Phase:** If event frequency crosses the threshold, the system enters a decision phase. Rather than immediately acting, it applies a brief secondary check to filter out edge cases such as file indexing services or backup tools performing legitimate bulk operations. This secondary check examines process ancestry and command-line arguments to reduce false positives.
- **Response Phase:** Confirmed malicious activity triggers immediate process termination via a system call. The response module sends a SIGKILL signal to the identified process and all child processes, ensuring that multi-threaded ransomware cannot continue encryption in a sibling thread. A snapshot of encrypted files is taken to inform recovery efforts.
- **Logging and Visualization:** : All detection and response events are recorded in a structured log file with timestamps, process identifiers, and event counts. A graphical dashboard provides real-time status updates, showing detection alerts, encrypted file counts, and system health metrics.

- **Dataset Generation:** Concurrently with all above phases, the behavioral collection module records process feature vectors at regular intervals. These are labeled based on whether the process was flagged as malicious and stored in CSV format for downstream machine learning workflows.

VII. DETECTION ALGORITHM AND MATHEMATICAL MODEL

The detection mechanism is based on monitoring file system events within a defined time window. Ransomware behavior is characterized by rapid file encryption, which results in multiple file modification events occurring in a short duration.

Let $E(t)$ represent the number of file events within a time interval T . The system evaluates the frequency of these events to determine whether the activity is suspicious.

$$E(t) = \sum_{i=1}^n e_i \quad (1)$$

where e_i represents individual file modification events.

A potential ransomware attack is identified when the number of events exceeds a predefined threshold:

$$E(t) \geq \theta \quad (2)$$

where θ corresponds to the maximum allowed number of events (MAX_EVENTS).

The time window is defined as:

$$T = t_{current} - t_{first} \quad (3)$$

If the condition $T \leq \Delta t$ (TIME_WINDOW) is satisfied and the event threshold is exceeded, the system classifies the behavior as malicious.

For machine learning readiness, a feature vector F is constructed as:

$$F = [CPU, Memory, OpenFiles, Runtime] \quad (4)$$

These features are used to build a labeled dataset for future machine learning-based detection.

VIII. IMPLEMENTATION DETAILS

The proposed system is implemented using Python and operates on a Linux-based environment. The file monitoring functionality is achieved using a real-time event-driven mechanism that tracks file system changes such as creation, modification, and renaming of files.

The detection engine uses a sliding time window approach to analyze the frequency of file events. If the number of events exceeds a predefined threshold within the specified time window, the activity is classified as suspicious.

Process monitoring is implemented using system-level libraries that allow access to process attributes such as CPU usage, memory consumption, and open file descriptors. These features are extracted and stored for dataset generation.

The response mechanism identifies the malicious process based on command-line arguments and terminates it using system calls. Logging is performed continuously to record attack events, detection time, and response actions.

The graphical user interface is implemented to provide real-time visualization of system status, including detection alerts, encrypted file count, and recovery operations.

IX. PERFORMANCE METRICS

The performance of the proposed system is evaluated using the following metrics:

- **Detection Time:** The time taken to detect ransomware activity after the attack begins.
- **Response Time:** The time required to terminate the malicious process after detection.
- **Encrypted Files Count:** The number of files encrypted before detection and response.
- **System Overhead:** The computational load introduced by the monitoring and detection process.

Experimental observations indicate that the system achieves low detection time and minimizes the number of encrypted files. The lightweight design ensures minimal impact on system performance, making it suitable for real-time deployment.

X. SYSTEM ARCHITECTURE

The system is organized into five loosely coupled modules that communicate through an internal event bus. This modular design ensures that changes to one component — for example, swapping out the file monitoring library or integrating a new ML model — do not require changes to the rest of the system.

1. File Monitoring Module: Registers OS-level watchers on the target directory tree. Captures create, modify, rename, and delete events with millisecond-precision timestamps and process attribution. Forwards events to the Detection Engine via an in-memory queue. **2. Detection Engine:** Maintains a sliding window of recent events and evaluates the configured threshold conditions. Applies secondary filters to reduce false positives. Publishes detection alerts when thresholds are exceeded.

3. Response Module: Subscribes to detection alerts and immediately initiates process termination. Identifies the malicious process and all child processes using the OS process tree. Issues termination signals and records the number of files encrypted before intervention.

4. Dataset Generation Module: Runs as a parallel thread, periodically sampling process attributes for all running processes. Attaches labels based on detection outcomes. Writes feature vectors to disk in CSV format with process metadata.

5. User Interface: Provides a real-time graphical dashboard showing system status, event timeline, and detection history. Writes structured logs in JSON format for integration with SIEM systems. Supports alert notifications via system notifications. The architecture ensures efficient communication between modules and supports scalability for future enhancements such as machine learning-based prediction.

TABLE I
PERFORMANCE EVALUATION ACROSS DIFFERENT SCENARIOS

Scenario	Avg Detection Time (s)	Files Encrypted	Response Time (s)
A (Normal)	1.2 ± 0.3	1.4 ± 0.5	340 ± 45
B (High BG Load)	1.9 ± 0.6	2.1 ± 0.7	410 ± 70
C (Slow/Evasive)	8.7 ± 2.1	6.3 ± 1.8	390 ± 55

TABLE II
COMPARISON OF DETECTION APPROACHES

Feature	Signature AV	Shield [3]	FS ML-Only [12]	Network Proposed	Behavioral
Detection Type	Signature	Behavioral	Statistical	Traffic	Behavioral
Real-Time Detection	No	Partial	No	Partial	Yes
Zero-Day Detection	Poor	Good	Good	Moderate	Good
Automated Response	No	Partial	No	No	Yes
Low System Overhead	Yes	Moderate	No	Yes	Yes
ML Integration Path	No	No	Yes	No	Yes
Offline Operation	Yes	Yes	No	No	Yes

XI. EXPERIMENTAL RESULTS AND ANALYSIS

The system was tested using a simulated ransomware program in a controlled environment. Multiple test files were used to evaluate detection performance.

The results demonstrate that the system detects ransomware activity within a very short time frame. In most cases, the number of encrypted files was limited to one or two before the malicious process was terminated.

In Scenario A, the system consistently detected ransomware activity within approximately 1.2 seconds of the attack commencing, limiting encryption to fewer than 2 files on average. This result validates the core premise of the system: that a threshold-based behavioral monitor can respond faster than ransomware can meaningfully progress.

Scenario B introduced background noise from a file indexing service. Two false positives were recorded across 20 runs when the indexing service and the simulated ransomware happened to overlap their I/O bursts within the same time window. Detection time increased slightly to 1.9 seconds due to the secondary filter processing additional candidate processes. This is an acceptable tradeoff.

Scenario C is the most revealing. When the simulated ransomware deliberately slowed its encryption rate to stay below the per-window threshold, detection took significantly longer (8.7 seconds) and more files were encrypted (6.3 on average). This confirms a known limitation: sufficiently patient ransomware can partially evade threshold-based detection. This is exactly the case where machine learning integration—detecting sustained, low-level anomaly rather than a burst—will provide the most value.

XII. COMPARISON WITH EXISTING SYSTEMS

Table II summarizes how our system compares against representative existing approaches:

XIII. CONCLUSION AND FUTURE WORK

This paper presents a real-time ransomware detection and response system based on behavioral monitoring and process

analysis. The system effectively identifies ransomware activity by analyzing file system events and terminates malicious processes to minimize data loss.

The proposed approach addresses the limitations of traditional detection systems by providing immediate response and reducing the impact of ransomware attacks. The integration of dataset generation further enhances the system by enabling future machine learning-based improvements.

Future work includes training machine learning models using the collected dataset, improving detection accuracy, and extending the system to support large-scale and enterprise-level environments. Additional enhancements such as network-level monitoring and advanced anomaly detection techniques can also be explored.

Figure Labels: Use 8 point Times New Roman for Figure labels. Use words rather than symbols or abbreviations when writing Figure axis labels to avoid confusing the reader. As an example, write the quantity “Magnetization”, or “Magnetization, M”, not just “M”. If including units in the label, present them within parentheses. Do not label axes only with units. In the example, write “Magnetization (A/m)” or “Magnetization {A[m(1)]}”, not just “A/m”. Do not label axes with a ratio of quantities and units. For example, write “Temperature (K)”, not “Temperature/K”.

XIV. PERFORMANCE METRICS

The performance of the proposed system is evaluated using the following metrics:

- **Detection Time:** The time taken to detect ransomware activity after the attack begins.
- **Response Time:** The time required to terminate the malicious process after detection.
- **Encrypted Files Count:** The number of files encrypted before detection and response.
- **System Overhead:** The computational load introduced by the monitoring and detection process.

Experimental observations indicate that the system achieves low detection time and minimizes the number of encrypted files. The lightweight design ensures minimal impact on system performance, making it suitable for real-time deployment.

XV. FUTURE MACHINE LEARNING INTEGRATION

The collected behavioral dataset enables the integration of machine learning models for enhanced detection capabilities. In future work, supervised and unsupervised learning algorithms such as Random Forest, Isolation Forest, and Support Vector Machines can be trained using the generated dataset.

These models can be integrated into the detection pipeline to classify processes based on learned behavioral patterns. This will enable detection of zero-day ransomware attacks that may not exhibit obvious rule-based characteristics.

A hybrid detection approach combining rule-based and machine learning-based techniques can further improve accuracy and robustness.

REFERENCES

- [1] S. Homayoun, A. Dehghantanha, M. Ahmadzadeh, S. Hashemi, and R. Khayami, “DRTHIS: Deep Learning-Based Real-Time Ransomware Detection Using Hierarchical Indicators,” *IEEE Access*, vol. 11, pp. 15234–15248, 2023.
- [2] M. Al-Rimy, M. Maarof, and S. Shaid, “Ransomware Detection Using Machine Learning and Behavioral Analysis: Recent Advances and Challenges,” *IEEE Access*, vol. 10, pp. 98765–98780, 2022.
- [3] A. Continella, A. Guagnelli, G. Zingaro, and G. Vigna, “ShieldFS: A Self-Healing, Ransomware-Aware Filesystem,” *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 1, pp. 1–15, 2022.
- [4] H. S. Anderson and P. Roth, “Ember: An Open Dataset for Training Static PE Malware Machine Learning Models,” *IEEE Security Privacy*, vol. 20, no. 2, pp. 85–92, 2022.
- [5] Y. Chen, S. Zhou, and Q. Luo, “Behavior-Based Ransomware Detection Using Machine Learning and System Activity Monitoring,” *Computers Security*, vol. 122, 2023.
- [6] J. Zhang, Z. Qin, K. Ren, and X. Chen, “Ransomware Detection Based on File System Activity and Machine Learning,” *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 2456–2468, 2023.
- [7] M. K. Rai and R. K. Singh, “Real-Time Malware Detection Using Lightweight Machine Learning Techniques,” *IEEE Access*, vol. 12, pp. 11234–11248, 2024.