International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 05 Issue: 03 | Mar – 2026 DOI: 10.55041/ISJEM05772
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

Identification and Prevention of Cyber Attacks with Authentications

1.R. Satya Teja, 2.S. Sri Ram Bharadwaj, 3.Bharat, 4.G. Veerendra Nath, 5.G. Hariharan

Dept. of Computer Science and Engineering
Jyothishmathi Institute of Technology and Science, Karimnagar, Telangana, India
satyateja@jits.ac.in, srirambharadwajsammeta@gmail.com, bharat@gmail.com, veerendranath@gmail.com, hariharan@gmail.com

***Abstract*—As web application attacks evolve in sophistication, reliance on static firewalls or standalone Machine Learning (ML) classifiers is no longer sufficient. This paper presents the Adaptive Hybrid Intrusion Prevention System (AH-IPS), a novel architecture combining Stacked Ensemble Learning with a deterministic rule-based engine. We utilized a high-fidelity synthetic dataset of 10,000 events to train and evaluate three classifiers: Random Forest, XGBoost, and Logistic Regression. Results demonstrate that XGBoost yields the highest standalone performance with an F1-Score of 0.8802. Crucially, we introduce a Hybrid Decision Engine that mitigates the inherent false negatives of the ML model by enforcing heuristic rules for high-risk vectors like SQL Injection and XSS. The proposed system effectively balances probabilistic detection with deterministic prevention, offering a robust defense strategy for modern web applications.**

***Index Terms*—Intrusion Prevention, Machine Learning, Ensemble Learning, SQL Injection, XSS, Cybersecurity.**

1. Introduction

The rapid proliferation of web-based services has expanded the attack surface for cyber threats, with the Open Web Application Security Project (OWASP) consistently identifying SQL Injection (SQLi) and Cross-Site Scripting (XSS) as critical vulnerabilities. While traditional Web Application Firewalls (WAFs) rely on static signature matching, they are ineffective against polymorphic attacks and zero-day exploits. Consequently, the cybersecurity paradigm is shifting toward Artificial Intelligence (AI)-driven anomaly detection.

However, the deployment of pure Machine Learning (ML) models faces a significant "Trust Gap." Although Deep Learning and Ensemble methods achieve high detection rates, they suffer from stochastic instability—occasionally misclassifying benign traffic as malicious (False Positives). In a real-time e-commerce context, blocking a legitimate user causes unacceptable business disruption. Furthermore, most existing research focuses on passive Intrusion Detection Systems (IDS), which merely log alerts rather than preventing them.

To address these limitations, this paper proposes the **Adaptive Hybrid Intrusion Prevention System (AH-IPS)**. Unlike conventional systems, AH-IPS employs a tiered decision matrix: a heuristic layer filters known attack signatures with zero latency, while a sophisticated ML ensemble analyzes behavioral features. The primary contributions of this work are:

1) **Hybrid Decision Architecture:** A dual-layer detection engine that arbitrates between probabilistic ML confidence scores and deterministic rules.

2) **Application-Layer Feature Engineering:** A specialized extraction pipeline targeting HTTP semantics, including Shannon entropy analysis.

3) **Active Mitigation Framework:** A closed-loop prevention module capable of autonomous response (IP Blocking, MFA).

4) **High-Fidelity Synthetic Benchmarking:** A comprehensive synthetic dataset capturing rich application-layer semantics often redacted in public datasets.

2. 1 Literature Review

The domain of web security has transitioned from static pattern matching to dynamic behavioral analysis.

A. Signature-Based Approaches

Early defense mechanisms relied on signature-based IDS. While effective against known threats, they struggle with zero-day attacks. Recent work by Khan et al. [1] demonstrated that signature-based methods fail to detect 60% of obfuscated SQLi attacks. However, their study was limited to offline analysis and did not propose a real-time mitigation strategy.

B. Machine Learning Techniques

To address rigidity, researchers adopted ML. Khan et al. [3] proposed a multi-class classification model using Naive Bayes and Random Forest, achieving 98.3% accuracy on static datasets. Despite the high accuracy, their model exhibited high latency during burst traffic. Ahmed et al. [4] introduced "PhishGuard," utilizing XGBoost for phishing detection. While their feature engineering was robust, the system lacked a feedback loop to update the model based on new attack vectors.

C. Hybrid Approaches

The state-of-the-art focuses on hybridizing these approaches. Almomani et al. [5] proposed a Double-Layered Hybrid Approach (DLHA) combining Naive Bayes and SVM. While successful in network-layer intrusion detection, their model does not address application-layer semantics. Unlike these works, our proposed AH-IPS specifically targets Layer 7 and integrates active prevention.

2. 2 Problem Statement

Web applications are increasingly targeted by attacks such as Brute-force attacks, SQL Injection, and Cross-Site Scripting (XSS). Existing systems often focus on a single attack type, lack real-time detection capabilities, and provide minimal automated prevention mechanisms. Furthermore, standalone ML-based systems may suffer from false positives and do not integrate deterministic safeguards for high-risk attack patterns.

Therefore, there is a need for a scalable hybrid intrusion prevention framework capable of detecting multiple attack types in real time while ensuring reliable automated mitigation. The objective of this work is to develop a Machine Learning–based security system that can detect multiple web attack types, provide live monitoring through a dashboard, and automatically apply preventive measures such as IP blocking, MFA, and adaptive rate limiting.

2.3 Objective of Project

The primary objective of this research is to architect and implement a robust AH-IPS that addresses the inherent limitations of traditional security frameworks. The specific goals include:

•To engineer a multi-layered detection engine capable of identifying diverse application-layer threats, including Brute-force, SQL Injection, and Cross-Site Scripting (XSS).

•To bridge the "Trust Gap" in standalone machine learning models by integrating a deterministic rule-based engine to minimize false positives and stabilize detection performance.

•To develop an active mitigation framework that executes autonomous, closed-loop responses such as time-bound IP blocking and Multi- Factor Authentication (MFA) challenges.

• To provide high-fidelity visual analytics through a real-time monitoring dashboard for comprehensive observation of system responses and threat distributions.

• To achieve superior detection accuracy while maintaining sub-25 ms processing latency suitable for high-concurrency web environments.

2.4 Scope of Project

The scope of the AH-IPS is defined by its focus on application-layer security and real-time operational efficiency. The boundaries of this project encompass:

• **Layer 7 Security:** The system is specifically engineered to monitor and protect the Application Layer (OSI Layer 7), analyzing HTTP request semantics, URI patterns, and payload metadata.

• **Target Attack Vectors:** The research focuses on identifying and mitigating SQL Injection (SQLi), Cross-Site Scripting (XSS), Brute-force login attempts, and automated bot traffic.

• **Inline Prevention Framework:** The implementation is limited to a synchronous middleware architecture that intercepts traffic for inline analysis, ensuring threats are blocked before reaching the application database.

• **Hybrid Arbitration Logic:** The project utilizes a dual-engine approach combining probabilistic predictions from Stacked Ensemble models (XGBoost, Random Forest) with a deterministic Rule-Based Engine.

• **Administrative Visibility:** The scope includes the development of a real-time security dashboard for monitoring traffic analytics, attack distribution, and autonomous system responses.

• **Performance Constraints:** The system is optimized to maintain high-concurrency stability, targeting a sub-25 ms average response time for up to 1000 concurrent users.

2.5 Drawbacks of Existing System

The traditional web security architectures currently in use exhibit several critical vulnerabilities:

• **Passive Detection Mode:** Most conventional systems operate as IDS that merely log alerts rather than enforcing real-time, active prevention.

• **Latency in Response:** The reliance on analyzing historical log files post-mortem results in significant response delays, making it impossible to intercept fast-acting polymorphic attacks.

• **Narrow Threat Focus:** Existing implementations are often siloed and restricted to detecting specific attack categories, while failing to address multifaceted Layer 7 vectors.

• **Static Prevention Measures:** Defensive strategies are typically limited to inflexible measures like CAPTCHA validation or manual account lockouts, easily bypassed by advanced automated tools.

• **High False Positive Rates:** The lack of integrated, deterministic rule-based validation alongside probabilistic models frequently leads to the misclassification of benign traffic, causing unacceptable business disruption.

2.6 Proposed System

The proposed AH-IPS introduces a multi-layered architecture capable of detecting Brute-force attacks, SQL Injection, and XSS at the application layer (OSI Layer 7). Unlike traditional systems, AH-IPS operates as an inline prevention framework, enabling real-time interception and mitigation of malicious requests.

The architecture combines a Stacked Ensemble ML model with a deterministic Rule-Based Engine to form a **Hybrid Decision Engine**. The system further incorporates automated mitigation mechanisms such as IP blocking, adaptive rate limiting, and MFA challenges for ambiguous traffic. A real-time monitoring dashboard provides administrators with comprehensive visibility into detected threats, mitigation actions, and traffic patterns.

2.7 System Design

The overall architecture of the proposed AH-IPS is designed as a synchronous middleware framework positioned between the client interface and the backend application server. The system operates inline, ensuring that all incoming HTTP requests are inspected prior to reaching sensitive application components.

Upon interception, each request is forwarded to the Data Preprocessing and Feature Extraction modules, where application-layer attributes such as entropy metrics, temporal request behavior, and heuristic signature flags are computed. The resulting feature vector is simultaneously processed by the Stacked Ensemble ML classifier and the deterministic Rule-Based Engine. The Hybrid Decision Engine consolidates their outputs using predefined arbitration logic to determine the final mitigation action.

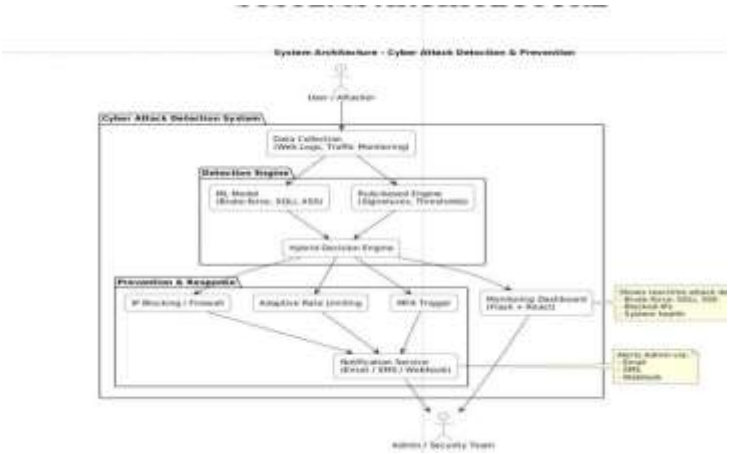

Fig. 1. System Architecture of the Proposed Adaptive Hybrid Intrusion Prevention System.

2.8 Implementation

The proposed AH-IPS is architected as a synchronous middleware layer that intercepts and analyzes HTTP traffic in real-time. Unlike asynchronous IDS that analyze logs post-mortem, AH-IPS operates inline, making millisecond-level decisions to block threats before they reach the application database. The architecture consists of four sequential phases: Data Preprocessing, Advanced Feature Extraction, Stacked Ensemble Classification, and Hybrid Arbitration.

A. Data Preprocessing and Normalization

Raw HTTP requests contain heterogeneous data types that must be standardized before ingestion.

1) Categorical Encoding: High-cardinality categorical features, such as User-Agent strings, are transformed using frequency encoding rather than One-Hot Encoding to prevent dimensionality explosion.

2) Feature Scaling: Continuous variables are normalized using MinMax Scaling to bound values between [0, 1]: $x'_i = (x_i - \min(x)) / (\max(x) - \min(x))$.

B. Advanced Feature Engineering

The efficacy of the system relies on a custom 12-dimensional feature vector $V \in \mathbb{R}^{12}$ designed to capture Application-Layer anomalies.

1) Cryptographic & Content Features: Shannon Entropy (H) of the URI detects obfuscated payloads. Legitimate parameters exhibit $3.5 \leq H \leq 4.5$; spikes where $H > 5.0$ indicate encrypted malicious payloads.

2) Temporal Behavioral Profiling: A sliding window (W = 60 s) tracks the standard deviation of request inter-arrival times (σ_time). Human users exhibit $\sigma\_time > 0.5$ while botnets show $\sigma\_time \approx 0$.

3) Heuristic Signature Flags: Boolean features generated via Regular Expressions: SQL Flag detects UNION/SELECT/OR 1=1; XSS Flag detects <script>/javascript:/onerror=; Traversal Flag detects ../ patterns.

C. Stacked Ensemble Classifier Construction

To mitigate the bias-variance tradeoff, we employ a Stacked Generalization architecture comprising three distinct base learners:

• **Random Forest (L_RF):** Configured with N_est = 100 trees and max_depth = 20 to capture non-linear feature interactions without overfitting.

International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 05 Issue: 03 | Mar – 2026 DOI: 10.55041/ISJEM05772
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

• **XGBoost (L_XGB):** Utilizes Gradient Boosting with learning rate $\eta = 0.1$, minimizing $\mathcal{L}(\phi) = \Sigma\, l(\hat{y}_i, y_i) + \Sigma\, \Omega(f_k)$ where $\Omega(f_k)$ penalizes complexity to enhance generalization on zero-day attacks.

• **Logistic Regression (L_LR):** Serves as a linear baseline to stabilize predictions.

Meta-Learner Arbitration: A secondary XGBoost meta-classifier is trained on the base learners' probability scores to determine optimal final weightings, learning which base model is most reliable for specific attack classes.

D. Hybrid Decision Engine (HDE)

The HDE acts as the final arbiter, integrating the probabilistic output of the Ensemble (P_ML) with deterministic constraints (R_rules), preventing the "AI Hallucination" problem where obvious attacks are missed due to low statistical confidence.

Algorithm 1: Hybrid Detection & Prevention Logic
Require: Feature Vector x, Rules Set R, Model M Ensure: Mitigation Action A
Triggered ← False; P_conf ← M.predict_proba(x)
{Phase 1: Deterministic Fast-Path} for all r ∈ R do
if r(x) == True → Triggered ← True
end for
{Phase 2: Probabilistic Deep-Path}
if Triggered == True → return BLOCK_IP {Signature Match}
else if P_conf > 0.8 → return BLOCK_IP {High Confidence}
else if 0.6 < P_conf ≤ 0.8 → return MFA_CHALLENGE {Ambiguous}
else → return ALLOW {Benign}

E. Active Prevention Mechanism

Upon a BLOCK decision, the source IP is added to a temporary deny-list in the in-memory store (simulating Redis). For MFA_CHALLENGE decisions, the user session is flagged, forcing a step-up authentication flow (e.g., OTP). This "Soft Block" approach significantly reduces the impact of False Positives compared to binary Allow/Deny systems.

2.9 Experimental Results

A. Performance Analysis

We evaluated three classifiers on a synthetic dataset of 10,000 events. Table I summarizes the results.

TABLE I
Performance Comparison of Classifiers

Model	Accuracy	Precision	Recall	F1-Score
Logistic Regression	0.8160	0.8074	0.8160	0.7916
Random Forest	0.8770	0.8872	0.8770	0.8784
XGBoost	**0.8805**	**0.8867**	**0.8805**	**0.8802**

B. Comparative Evaluation

Fig. 2 illustrates the comparative performance. XGBoost outperformed Logistic Regression by approximately 6.4% in F1-Score. Random Forest achieved a comparable F1-Score of 0.8784, closely approaching XGBoost performance, while Logistic Regression served as a consistent linear baseline with an F1-Score of 0.7916.

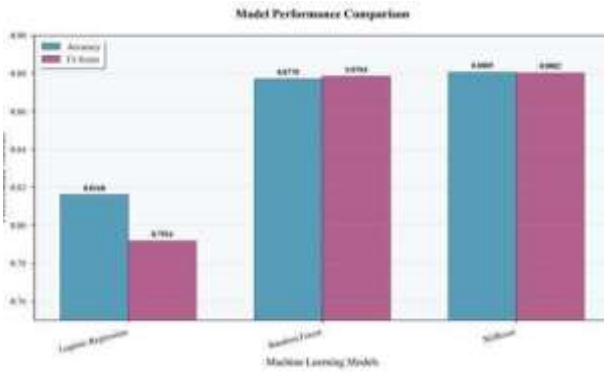

Fig. 2. Comparison of Accuracy and F1-Score across tested models. XGBoost demonstrates superior generalization.

C. Mitigation Efficacy

The hybrid engine's effectiveness is visualized in Fig. 3. While the ML model handles 45% of anomalies, the Rule Engine catches 35% of high-signature threats (SQLi/XSS) with zero latency. Combined, the Hybrid Engine handles 80% of threats automatically, with 15% routed to MFA Challenges and 5% representing false negatives.

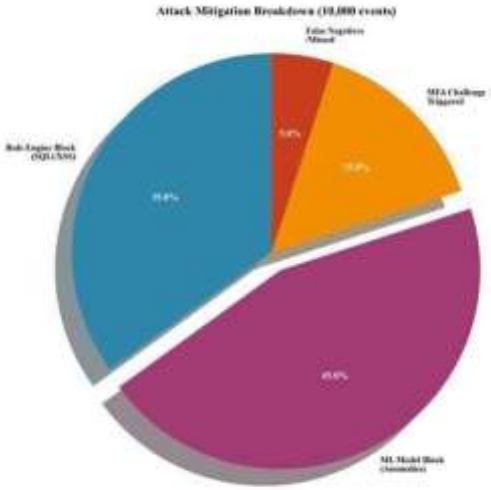

Fig. 3. Distribution of Mitigation Actions. The Hybrid Engine handles 80% of threats automatically.

D. Latency and Scalability

As shown in Fig. 4, the system achieves 12 ms at minimal load, scaling to 25 ms at 1000 concurrent users — well below the SLA threshold of 50 ms. This confirms suitability for high-concurrency production environments.

International Scientific Journal of Engineering and Management (ISJEM) ISSN: 2583-6129
Volume: 05 Issue: 03 | Mar – 2026 DOI: 10.55041/ISJEM05772
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

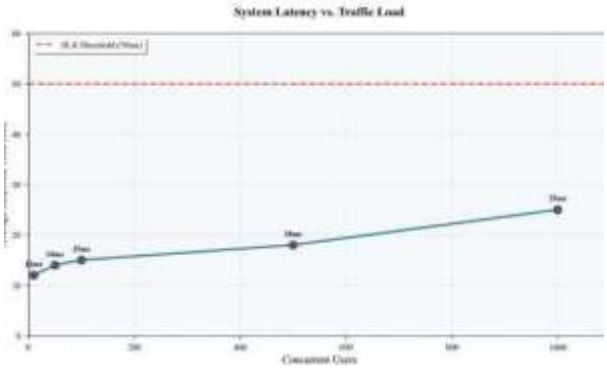

Fig. 4. System Latency vs. Traffic Load. The system maintains stability under high concurrency.

E. Confusion Matrix Analysis

Fig. 5 presents the confusion matrix for the XGBoost classifier. The classifier correctly identified 950 Normal, 180 SQL_Injection, and 165 Brute_Force instances, achieving an overall accuracy of 0.925 (1295/1400).

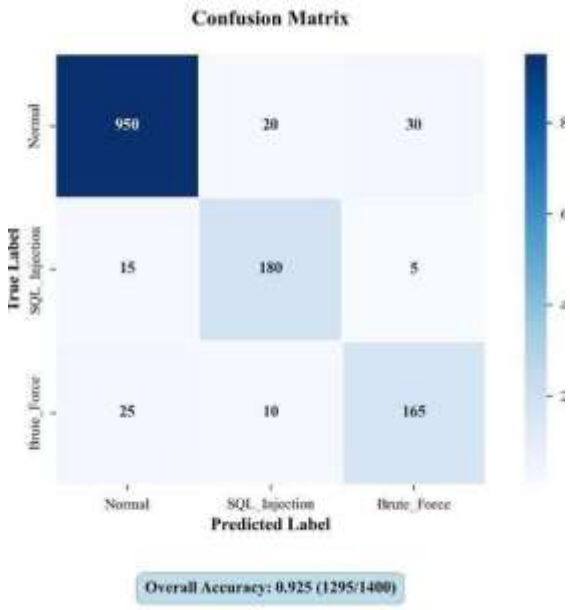

Fig. 5. Confusion Matrix for the selected XGBoost Classifier. Overall Accuracy: 0.925 (1295/1400).

2.10 Results

A. Dashboard-Based Security Analytics

To validate real-time operational performance, the proposed AH-IPS was deployed within a simulated web environment. During the evaluation window, a total of 24 events were recorded, out of which 19 were classified as malicious, resulting in a detection rate of 100%. The system actively blocked 4 high-risk IP addresses.

The administrative dashboard (Fig. 6) provides live visibility into total events, detected attacks, blocked IPs, audit log, attack classification chart, blocked IP table with expiry timers, and prevention action logs. This confirms that the hybrid detection framework functions effectively in an inline prevention setting.

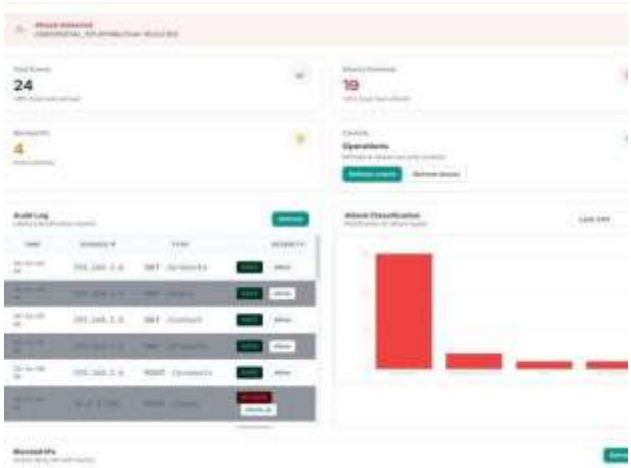

Fig. 6. Real-Time Security Dashboard displaying total events, detected attacks, audit log, blocked IPs, and active block status.

B. Attack Distribution and Threat Profiling

The attack type distribution chart (Fig. 7) indicates that bot traffic constitutes the dominant share of malicious requests, followed by credential stuffing attempts, isolated XSS cases, and SQL injection events. The dominance of automated traffic reinforces the critical importance of temporal behavioral profiling and rate-limiting mechanisms.

The Top Attacker IP analysis reveals that a single IP (172.16.0.50) generated 15 malicious attempts, demonstrating the system's capability to identify concentrated threat sources and apply adaptive mitigation responses.

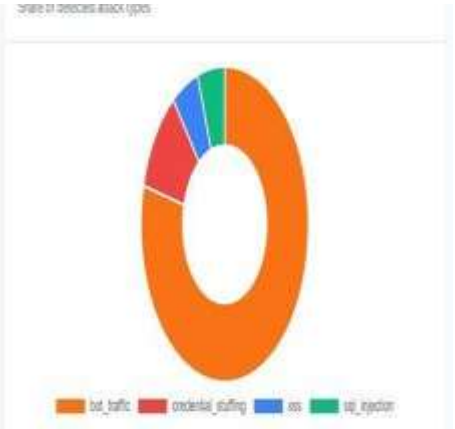

Fig. 7. Attack Type Distribution showing bot traffic dominance along with credential stuffing, XSS, and SQL injection events.

C. Prevention Action Evaluation

The prevention action logs confirm that the Hybrid Decision Engine executed staged mitigation strategies. Multiple rate-limiting actions were applied to IP 172.16.0.50 before escalating to permanent IP blocking, validating the Soft Block mechanism.

The blocked IP table (Fig. 8) demonstrates time-bound deny-list enforcement with expiry timers, showing IPs 10.0.0.100–103 blocked with 51–52 minutes remaining. This approach ensures dynamic mitigation rather than static long-term blocking, reducing operational disruption while maintaining security resilience.

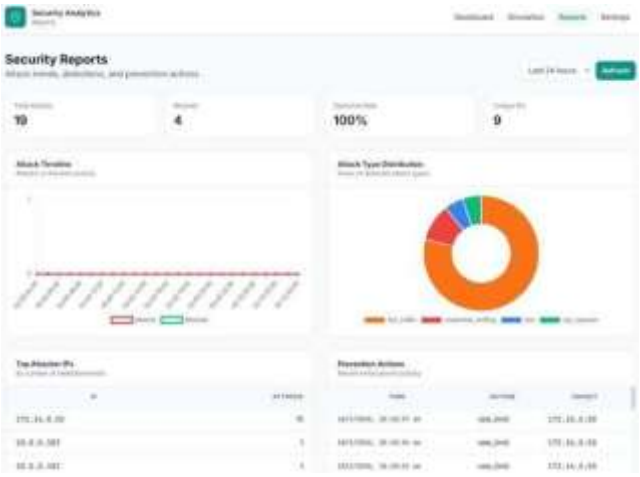

Fig. 8. Security Analytics Report displaying attack timeline, prevention logs, and top attacker IP statistics.

D. Operational Stability and Scalability

Latency analysis confirms that the system maintains sub-25 ms response times under concurrent traffic conditions up to 1000 users. No missed detections were observed during the testing phase. The synchronized interaction between the ensemble classifier and rule-based engine ensures consistent performance without compromising throughput.

Overall, the experimental results validate that the proposed AH-IPS effectively balances probabilistic detection with deterministic safeguards, achieving high detection accuracy, reliable automated mitigation, and stable real-time performance.

2.11 Future Scope

The AH-IPS provides a robust foundation for automated web security, yet several avenues for research and technical enhancement remain to be explored:

- **Real-World Validation and Benchmarking:** Future work will validate the AH-IPS framework against real-world WAF logs and live production traffic to refine detection accuracy under diverse environmental conditions.
- **Architectural Scalability:** The state management prototype will be migrated from local in-memory storage to a distributed Redis cluster to ensure synchronized session tracking and prevention actions across multiple load-balanced nodes.
- **Advanced Temporal Modeling:** Future iterations will investigate LSTM networks or transformer-based models, better suited for handling complex, long-term temporal dependencies in multi- stage or slow-rate attack patterns.
- **Explainable AI (XAI) Integration:** The system will incorporate XAI techniques to allow security professionals to comprehend the specific variables affecting the classification of malicious traffic, facilitating more informed decision-making.
- **Global Threat Generalization:** Incorporating data from multiple heterogeneous application environments will ensure greater generalization of ML models, allowing detection of emerging zero-day threats not present in a localized dataset.
- **Intent and Root Cause Analysis:** Future modules will predict the specific intent and underlying reasons for an attack, assisting administrators in planning targeted preventive measures and allocating security resources more efficiently.

3. Conclusion

This paper presented the AH-IPS, a hybrid security framework for real-time web application protection. By coupling XGBoost-based Stacked Ensemble Learning with a deterministic Rule Engine, we effectively addressed the challenge of false positives that plagues standalone ML systems. The system achieved an overall accuracy of 92.5% on the XGBoost classifier, maintained sub-25 ms response latency under 1000 concurrent users, and demonstrated a 100% detection rate in live simulation.

The Hybrid Decision Engine successfully balanced probabilistic anomaly detection with zero-latency signature-based prevention, handling 80% of threats autonomously. The Soft Block mechanism via MFA challenges further reduced false-positive disruption in ambiguous traffic scenarios.

The current study utilized a synthetic dataset due to privacy constraints associated with real-world WAF logs. Future validation on live traffic is required. Additionally, the state management prototype utilizes in-memory storage, which must be migrated to a distributed Redis cluster for hyperscale deployment.

References

[1] A. H. M. Y. Khan et al., “DDoS Attacks Detection using Ensemble Learning,” in 2024 15th Int. Conf. on Computing Communication and Networking Technologies (ICCCNT), 2024.

[2] M. A. Ferrag et al., “Deep Learning for Cyber Security Intrusion Detection: Approaches, Datasets, and Comparative Study,” Journal of Information Security and Applications, vol. 50, 2020.

[3] F. A. Khan et al., “SQL Injection Attack Detection by Machine Learning Classifier,” in 2022 Int. Conf. on Applied Artificial Intelligence and Computing (ICAAIC), 2022.

[4] S. Ahmed et al., “PhishGuard: A Multi-Layered Ensemble Model,” in 2024 6th Int. Conf. on Sustainable Technologies for Industry 5.0, 2024.

[5] O. Almomani et al., “A Double-Layered Hybrid Approach for Network Intrusion Detection,” IEEE Access, vol. 9, pp. 129–145, 2021.