

# TRANSFORMER BASED END TO END WEB APPLICATION FIREWALL (WAF) PIPELINE

Author's Name, R.Divya<sup>1</sup>, M.Chinraj<sup>2</sup>, A.Mohanpriyan<sup>3</sup>, S.Raman<sup>4</sup>, K.Sivanantham<sup>5</sup>

## Affiliation:

1. Assistant Professor, Salem College Of Engineering and Technology, Salem- Attur Main Road, M.Perumapalayam, Selliamman Nagar, Salem
- 2-5, Students (B.E Computer Science and Engineering), Salem College Of Engineering and Technology, Salem- Attur Main Road, M.Permapalayam, Selliamman Nagar, Salem.

## Abstract:

Web applications are continuously exposed to advanced cyber threats such as SQL injection, Cross-Site Scripting (XSS), and zero-day attacks. Traditional Web Application Firewalls (WAFs) mainly rely on static rule-based or signature-based detection methods. These systems are reactive in nature and often fail to detect new or unknown attack patterns. To overcome these limitations, this project proposes a Transformer-based end-to-end Web Application Firewall pipeline that uses deep learning for intelligent anomaly detection.

The proposed system utilizes a Transformer Auto encoder model to learn the structural and contextual patterns of legitimate HTTP traffic. During the training phase, the model is trained only on benign web server logs collected from platforms such as Apache and Nginx. Dynamic fields like timestamps, session IDs, and user-specific values are normalized to ensure consistent learning. By understanding the normal "grammar" of web requests, the system can identify deviations from expected behavior. In real-time deployment, incoming HTTP requests are processed through pre processing, tokenization, and inference stages. The Transformer model reconstructs each request and calculates a reconstruction loss value. If the loss exceeds a predefined threshold, the request is flagged as anomalous and treated as a potential attack. This method enables the detection of unknown and zero-day exploits without relying on predefined signatures.

## 1. INTRODUCTION

Web applications have become an essential part of modern digital infrastructure. They are widely used in banking, e-commerce, healthcare, education, and government services. As organizations increasingly depend on web-based platforms, the security of these applications has become a major concern. Cyber

attackers continuously attempt to exploit vulnerabilities in web applications to steal sensitive information, gain unauthorized access, or disrupt services.

Common web application attacks include SQL Injection (SQLi), Cross-Site Scripting (XSS), Cross-Site Request Forgery (CSRF), and Distributed Denial-of-Service (DDoS). These attacks can cause financial loss, data breaches, and reputational damage. To prevent such threats, Web Application Firewalls (WAFs) are deployed as a

protective layer between users and web servers. Traditional WAFs operate using rule-based and signature-based detection mechanisms. They compare incoming HTTP requests against predefined attack patterns and block malicious traffic.

However, traditional WAF systems have limitations. They are reactive in nature and cannot effectively detect zero-day attacks or newly modified attack techniques. Attackers often bypass signature-based systems by obfuscating payloads or slightly altering attack patterns. Additionally, maintaining and updating rule sets requires continuous manual effort from security teams. With the advancement of Artificial Intelligence (AI) and Machine Learning (ML), intelligent security solutions have emerged. Deep learning models, particularly Transformer-based architectures, have shown strong capability in understanding sequential and contextual data. Since HTTP requests can be treated as structured sequences, Transformer models can learn the normal behavior of legitimate web traffic and detect anomalies without relying on predefined rules.

## 2. BACKGROUND OF WEB APPLICATION SECURITY

Web application security focuses on protecting web-based systems, applications, and services from cyber threats and unauthorized access. With the rapid growth of internet usage and cloud-based platforms, web applications have become primary targets for attackers. Sensitive data such as personal information, financial records, login credentials, and business data are often stored and processed through web applications. This makes them highly attractive to cybercriminals.

In the early stages, web security mainly depended on network firewalls and basic input validation techniques. However, traditional network firewalls are not designed to inspect application-level traffic in detail. As a result, attackers began exploiting vulnerabilities at the application layer using techniques such as SQL Injection, Cross-Site Scripting (XSS), and Cross-Site Request Forgery (CSRF). These attacks manipulate HTTP requests to execute malicious commands on the server.

To address this issue, Web Application Firewalls (WAFs) were introduced. A WAF operates between the client and the web server, monitoring and filtering HTTP requests and responses. Most traditional WAFs use rule-based or signature-based detection methods. They rely on predefined attack patterns and regular expressions to block known threats. While effective against common attacks, these systems struggle to detect zero-day exploits and advanced persistent threats that do not match existing signatures. In recent years, Artificial Intelligence and Machine Learning techniques have been applied to cyber security. These intelligent systems analyze patterns and behaviors rather than relying solely on static rules. Deep learning models, particularly Transformer architectures, are capable of understanding sequential and contextual information. Since HTTP requests contain structured sequences of parameters and commands, Transformer models can learn the normal "behavior" of web traffic and identify anomalies more effectively.

Therefore, modern web application security is shifting from reactive, rule-based systems to proactive, intelligent anomaly detection mechanisms. This project is developed based on this advanced security approach.

## 3. LITERATURE REVIEW

### LITERATURE REVIEW

Web Application Firewalls (WAFs) were developed to protect web applications from

application-layer attacks by filtering and monitoring HTTP traffic between clients and servers. Unlike traditional network firewalls, which focus on IP addresses and ports, WAFs analyze application-level data such as URLs, parameters, headers, and request payloads. Most traditional WAF systems operate using rule-based and signature-based detection mechanisms. These systems compare incoming requests against a predefined database of known malicious patterns. Popular WAF solutions include Mod Security, AWS WAF, and Cloud flare WAF. They use regular expressions (RegEx) and security rules to identify common attacks such as SQL Injection (SQLi), Cross-Site Scripting (XSS), and command injection.

Although rule-based WAFs are effective against known threats, they have several limitations. They cannot detect zero-day attacks or newly modified payloads that do not match existing signatures. Furthermore, maintaining rule sets requires constant manual updates and tuning. If rules are too strict, they generate false positives. If rules are too relaxed, they may fail to detect attacks. Due to the increasing complexity of cyber threats, researchers have explored alternative approaches beyond static signature matching. This has led to the integration of machine learning and deep learning techniques into web security systems.

Additionally, maintaining and updating signature databases requires continuous manual effort. Security teams must regularly update rules to keep up with emerging threats. Improperly configured rules can also result in high false positive rates, where legitimate user requests are incorrectly blocked.

## 4. SIGNATURE-BASED DETECTION SYSTEMS

Signature-based detection systems are one of the earliest and most widely used approaches in cyber security. In this method, incoming network or application traffic is compared against a database of predefined attack patterns, known as signatures. If a match is found between the incoming request and a stored signature, the system identifies it as malicious and blocks or alerts the administrator.

In the context of Web Application Firewalls (WAFs), signature-based systems analyze HTTP requests to detect common attack patterns such as SQL Injection (SQLi), Cross-Site Scripting (XSS), command injection, and directory traversal attacks. These systems typically use regular expressions (Reg Ex) and rule sets to identify suspicious keywords, characters, or request structures. For example, inputs containing patterns like "OR 1=1" or "<script>" tags may be flagged as potential attacks. One of the main advantages of signature-based detection is its

simplicity and efficiency. It is highly effective against known and well-documented attack techniques. Once a signature is added to the database, the system can quickly detect repeated occurrences of the same attack pattern.

However, this approach has significant limitations. Signature-based systems are reactive in nature, meaning they can only detect attacks that have already been identified and documented. They cannot detect zero-day attacks or newly modified payloads that do not match existing signatures. Attackers can easily bypass such systems by slightly modifying their attack syntax, using encoding techniques, or obfuscating malicious content.

Additionally, maintaining and updating signature databases requires continuous manual effort. Security teams must regularly update rules to keep up with emerging threats. Improperly configured rules can also result in high false positive rates, where legitimate user requests are incorrectly blocked.

Due to these limitations, researchers have explored intelligent and adaptive detection techniques, including machine learning and deep learning-based approaches, to enhance web application security.

## 5. MACHINE LEARNING IN CYBERSECURITY

Machine Learning (ML) has emerged as a powerful tool in the field of cyber security for detecting threats based on patterns and behaviors rather than predefined rules. Unlike traditional security systems that rely on signature databases, ML-based systems analyze large volumes of data to identify anomalies and suspicious activities automatically.

In web application security, machine learning techniques are used to classify HTTP requests as either benign or malicious. Supervised learning algorithms such as Support Vector Machines (SVM), Decision Trees, Random Forests, and Neural Networks have been widely applied for intrusion detection and web attack classification. These models are trained using labeled datasets containing examples of both normal and malicious traffic.

Unsupervised learning techniques, such as clustering and anomaly detection algorithms, are also used when labeled data is limited. These models learn the normal behavior of web traffic and flag deviations as potential threats. This approach is particularly useful for detecting zero-day attacks that do not have predefined signatures.

However, traditional machine learning models often require manual feature engineering, where experts

must identify and extract relevant features from HTTP requests. This process can be time-consuming and may not capture complex contextual relationships within the data.

To overcome these limitations, deep learning models have been introduced. Deep learning architectures can automatically extract features from raw data and understand complex patterns more effectively. Among these models, Transformer-based architectures have shown significant promise in handling sequential and contextual data, making them suitable for analyzing HTTP request patterns in modern web security systems.

## 6. TRANSFORMER MODELS IN SECURITY APPLICATIONS

Transformer models have gained significant attention in recent years due to their outstanding performance in Natural Language Processing (NLP) tasks such as language translation, text classification, and sentiment analysis. Unlike traditional recurrent neural networks (RNNs), Transformers use a self-attention mechanism to capture contextual relationships between elements in a sequence. This ability to understand long-range dependencies makes them highly effective for sequential data analysis.

In cybersecurity applications, HTTP requests can be treated as structured sequences of tokens, similar to sentences in natural language. Each request contains elements such as URLs, parameters, headers, and payloads, which together form a meaningful pattern. Transformer models can analyze these sequences and learn the normal “grammar” of legitimate web traffic. When a malicious request deviates from learned patterns, the model can detect it as an anomaly.

Several recent research studies have applied Transformer architectures for intrusion detection, malware classification, and web attack detection. Transformer Autoencoders are particularly useful for anomaly detection tasks. In this approach, the model is trained using only benign data. During inference, the model attempts to reconstruct incoming requests. If the reconstruction error is high, it indicates that the request does not follow the normal pattern and may be malicious.

Compared to traditional machine learning models, Transformer-based systems do not rely heavily on manual feature engineering. They automatically extract complex features and contextual relationships from raw input data. This makes them more adaptable and effective in detecting zero-day attacks and obfuscated payloads.

Due to these advantages, Transformer models are increasingly considered a promising

solution for developing intelligent and adaptive Web Application Firewall systems. intent recognition, multiple.

## RESEARCH GAP

Despite the availability of traditional Web Application Firewalls (WAFs) and machine learning-based intrusion detection systems, several challenges remain unresolved in web application security. Most existing WAF solutions rely on static signature-based or rule-based detection mechanisms. While effective against known attacks, these systems are unable to detect zero-day exploits or newly modified attack patterns that do not match predefined rules.

Machine learning-based approaches have improved detection capabilities by analyzing traffic patterns and classifying requests as benign or malicious. However, many of these models depend on manually engineered features, which may not capture complex contextual relationships within HTTP requests. Additionally, some supervised learning methods require large labeled datasets, which are often difficult and expensive to obtain in real-world environments.

Although recent research has explored deep learning techniques, including recurrent neural networks and convolutional neural networks, these models may struggle to effectively capture long-range dependencies and contextual relationships within web request sequences. Transformer-based models have shown strong performance in sequential data analysis, but their application in real-time Web Application Firewall pipelines is still limited.

Another significant gap is the lack of continuous adaptation mechanisms in many existing systems. Web applications frequently evolve by adding new features, parameters, and endpoints. Without an incremental learning mechanism, detection models may become outdated and produce higher false positives or false negatives over time.

Therefore, there is a clear need for an intelligent,

Transformer-based end-to-end WAF pipeline that not only detects zero-day attacks effectively but also incorporates incremental learning to adapt continuously to evolving web environments. This project aims to address these research gaps.

## 7. SYSTEM ANALYSIS

### EXISTING SYSTEM

The existing system for protecting web applications mainly relies on traditional Web Application Firewalls (WAFs) that use rule-based and signature-based detection mechanisms. These systems inspect incoming HTTP requests and compare them against a predefined database of

known attack patterns. If a request matches a stored signature or violates a predefined rule, it is identified as malicious and blocked.

Most commercial and open-source WAF solutions operate using regular expressions (Regex) and manually configured security rules. They are effective in detecting well-known attacks such as SQL Injection (SQLi), Cross-Site Scripting (XSS), and command injection. These systems act as a protective layer between the client and the web server, filtering suspicious traffic before it reaches the application.

Furthermore, traditional WAF systems do not adapt automatically to changes in web application structure. When new features, parameters, or endpoints are introduced, existing rules may not be sufficient to protect them effectively.

Due to these limitations, there is a need for a more intelligent and adaptive security system that can detect unknown threats and continuously evolve with the application

## 8. IMPLEMENTATION

### 8.1. SOFTWARE REQUIREMENTS

The implementation of the proposed Transformer-Based Web Application Firewall (WAF) requires a stable and efficient software environment capable of handling deep learning computation, data pre processing, and real-time deployment.

The primary programming language used for development is **Python 3.8 or above**. Python is selected due to its simplicity, readability, and strong ecosystem of machine learning libraries.

The system requires the following software components:

- **Operating System:** Windows 10 / Windows 11 (64-bit) or Ubuntu Linux
- **Python Environment:** Anaconda / Virtual Environment (venv)
- **Deep Learning Framework:** Tensor Flow or PyTorch
- **Data Processing Libraries:** NumPy, Pandas
- **Machine Learning Utilities:** Scikit-learn
- **Visualization Tools:** Matplotlib
- **Web Framework for Deployment:** Flask or Fast API
- **Web Server:** Apache or Nginx
- **Log Processing:** Python Regex (re module)

Tensor Flow or PyTorch is used to implement the Transformer Auto encoder architecture. These frameworks provide efficient APIs for building neural networks, performing back propagation, and optimizing model weights.

Flask or FastAPI is used to deploy the trained model as an API service. This enables seamless integration between the WAF intelligence engine and the web server infrastructure.

The overall software stack ensures scalability, maintainability, and efficient performance during both training and real-time inference.

## 8.2 HARDWARE REQUIREMENTS

The proposed Transformer-Based Web Application Firewall system requires adequate hardware resources to ensure smooth model training and real-time inference. Since deep learning models, especially Transformer architectures, involve high computational complexity, sufficient processing power and memory are essential.

The recommended hardware configuration is as follows:

- **Processor:** Intel Core i5 (8th Generation or above) / AMD Ryzen 5 or above
- **RAM:** Minimum 8 GB (16 GB recommended for better performance)
- **Hard Disk:** 500 GB or above
- **Graphics Processing Unit (Optional):** NVIDIA GPU with CUDA support

During the training phase, the Transformer Auto encoder processes large volumes of HTTP log data. If a GPU is available, training speed is significantly improved due to parallel computation capabilities. However, the system can also operate using CPU-only configuration, although training time may increase.

For real-time deployment, the hardware requirements are moderate. Once the model is trained, inference can be executed efficiently even on standard server hardware with sufficient RAM and multi-core processing support.

The recommended configuration ensures that the system handles high-volume web traffic with minimal latency while maintaining accurate anomaly detection performance.

The proposed Transformer-Based Web Application Firewall system requires adequate hardware resources to ensure smooth model training and real-time inference. Since deep learning models, especially Transformer architectures, involve high computational complexity, sufficient processing power and memory are essential.

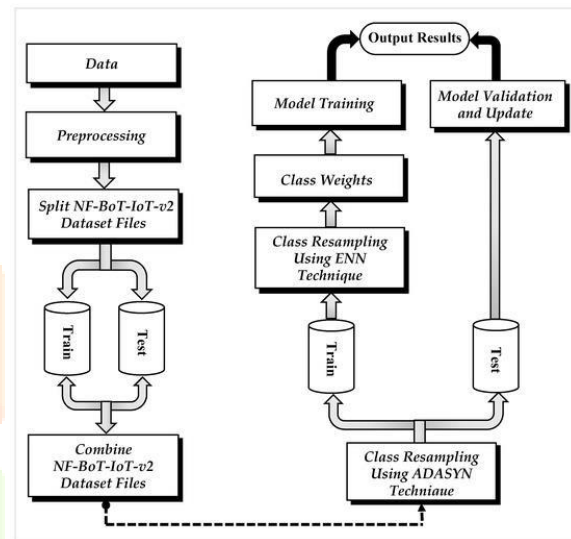
The recommended hardware configuration is as follows:

- **Processor:** Intel Core i5 (8th Generation or above) / AMD Ryzen 5 or above
- **RAM:** Minimum 8 GB (16 GB recommended for better performance)
- **Hard Disk:** 500 GB or above
- **Graphics Processing Unit (Optional):** NVIDIA GPU with CUDA support

During the training phase, the Transformer Auto encoder processes large volumes of HTTP log data. If a GPU is available, training speed is significantly improved due to parallel computation capabilities. However, the system can also operate using CPU-only configuration, although training time may increase.

For real-time deployment, the hardware requirements are moderate. Once the model is trained, inference can be executed efficiently even on standard server hardware with sufficient RAM and multi-core processing support. The recommended configuration ensures that the system handles high-volume web traffic with minimal latency while maintaining accurate anomaly detection performance.

## Model Training Process



The Model Training Process is a critical phase in developing the Transformer-Based Web Application Firewall. During this stage, the system learns the structural and contextual patterns of legitimate HTTP traffic so that it can later identify anomalies effectively.

### 5.3.1 Data Collection

The training data consists of HTTP access logs collected from web servers such as Apache or Nginx. These logs contain information including request method, URL, query parameters, status codes, and user-agent details.

Only **benign (legitimate) traffic** is used for training. This ensures that the model learns the normal behavior of users interacting with the web application.

### 5.3.2 Data Pre processing and Normalization

Raw log data cannot be directly fed into a deep learning model. Therefore, pre processing steps are applied:

- Removal of irrelevant fields
- Normalization of dynamic values (timestamps, session IDs, user IDs)
- Tokenization of HTTP requests
- Conversion of text into numerical vectors

Dynamic parameters are masked with generic tokens to prevent the model from over fitting to specific values. This helps the model learn request structure rather than memorizing exact data.

### 8.3.Adaptive and Incremental Learning

The proposed system includes an incremental learning mechanism that periodically updates the model using new verified traffic. This allows the firewall to adapt automatically when the web application evolves or new traffic patterns emerge. As a result, the system remains effective over time without requiring frequent manual rule updates.

## 9. SYSTEM DESIGN

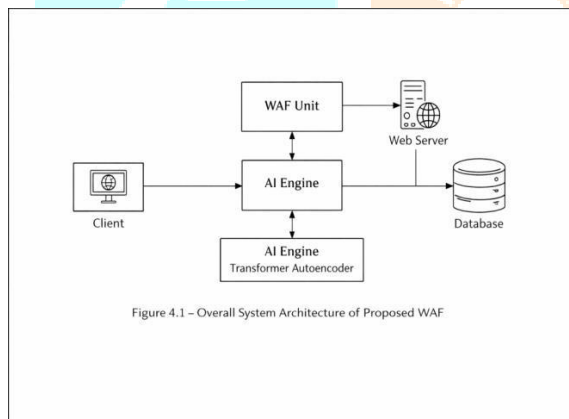


Fig 9.1 OVERALL SYSTEM ARCHITECTUR

The overall system architecture of the proposed Transformer-based Web Application Firewall (WAF) consists of multiple interconnected modules that work together to detect malicious web traffic. The system follows a structured pipeline approach, starting from data collection to final attack detection and logging.

The first component is the **Data Ingestion Module**, which collects HTTP request logs from web servers such as Apache or Nginx. These logs may be collected in real-time or batch mode depending on deployment requirements.

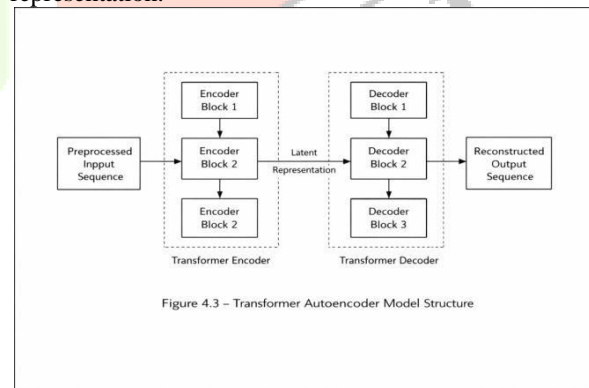
The collected data is passed to the **Pre processing and Normalization Module**. In this stage, unnecessary dynamic fields such as timestamps, IP addresses, and session IDs are cleaned or normalized. The text data is then tokenized and converted into numerical representations suitable for deep learning models.

Next, the processed data is fed into the **Transformer Auto encoder Model**. During training, the model learns the normal structure of HTTP requests. During inference, it attempts to reconstruct the input request and calculates the reconstruction error.

## 10. TRANSFORMER AUTOENCODER MODULE

The Transformer Auto encoder Module is the core component of the proposed Web Application Firewall (WAF) system. It is responsible for learning the normal structure and behavioral patterns of HTTP requests using deep learning techniques.

The module is based on a Transformer architecture, which utilizes a self-attention mechanism to capture contextual relationships between different elements of an HTTP request. Unlike traditional models, the Transformer can effectively understand long-range dependencies and sequential patterns within the data. The Auto encoder consists of two main parts: the encoder and the decoder. The encoder takes the input HTTP request (after preprocessing and tokenization) and converts it into a compact latent representation. This representation captures the essential features and structure of the request. The decoder then reconstructs the original request from this encoded representation.



## 11. ANOMALY DETECTION AND ALERT MODULE

The anomaly detection workflow represents the process by which the proposed Transformer-based Web Application Firewall identifies malicious HTTP requests. The workflow begins with an incoming HTTP request captured from the web

server. The request is first passed through the preprocessing and normalization stage, where dynamic and irrelevant fields such as timestamps and session identifiers are removed or standardized.

The processed request is then converted into tokenized numerical form and fed into the Transformer Auto encoder model. The model attempts to reconstruct the input request based on patterns learned from normal (benign) traffic during training. After reconstruction, a reconstruction error (loss value) is calculated by comparing the original request with the reconstructed output. This loss value is then evaluated against a predefined threshold. If the reconstruction error is below the threshold, the request is considered normal and allowed to pass through the system.

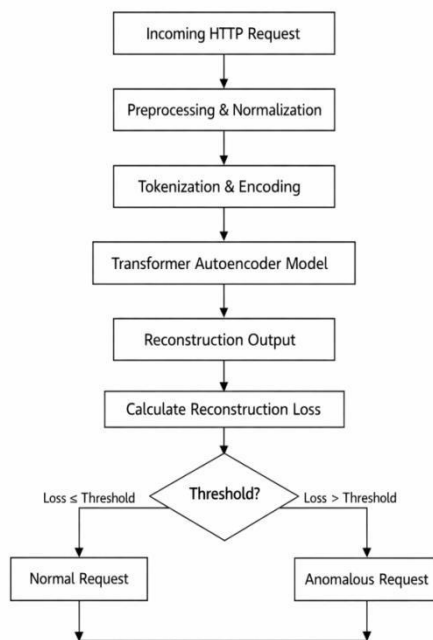


FIGURE 4.4 – ANOMALY DETECTION WORKFLOW

## 12. RESULTS AND DISCUSSION

### EXPERIMENTAL SETUP

The experimental setup was designed to evaluate the performance and effectiveness of the proposed Transformer-Based Web Application Firewall system. The objective of the experiment was to test the model's ability to distinguish between legitimate HTTP traffic and malicious attack patterns. The dataset used for experimentation consisted of web server access logs collected from a controlled

environment. The training dataset contained only benign HTTP requests to allow the Transformer Auto encoder to learn normal traffic behavior. For testing purposes, the evaluation dataset included both legitimate traffic and various attack payloads such as SQL Injection (SQLi), Cross-Site Scripting (XSS), and command injection attempts.

The experiments were conducted on a system configured with an Intel i5 processor, 8 GB RAM, and Python-based deep learning environment. The Transformer model was trained for multiple epochs until the reconstruction loss stabilized. A validation dataset was used to determine the optimal anomaly detection threshold.

During testing, each incoming request was processed through the real-time inference pipeline. Reconstruction loss values were calculated, and requests exceeding the threshold were classified as anomalous. The results were recorded for further analysis.

The experimental setup ensured a realistic evaluation scenario by simulating actual web server traffic conditions. Performance metrics such as accuracy, precision, recall, and false positive rate were measured to assess the system's effectiveness compared to traditional rule-based WAF systems.

### 12.1. PERFORMANCE METRICS

To evaluate the effectiveness of the proposed Transformer-Based Web Application Firewall, several standard performance metrics were used. These metrics help measure the model's ability to correctly classify legitimate and malicious HTTP requests.

#### 1. Accuracy

Accuracy represents the overall correctness of the model in detecting both normal and malicious traffic. It is defined as the ratio of correctly classified requests to the total number of requests.

$$\text{Accuracy} = (\text{True Positives} + \text{True Negatives}) / \text{Total Requests}$$

A high accuracy indicates that the system is effective in distinguishing between safe and malicious requests.

#### 2. Precision

Precision measures how many of the requests identified as malicious are actually malicious. It helps evaluate the reliability of attack detection.

$$\text{Precision} = \text{True Positives} / (\text{True Positives} + \text{False Positives})$$

High precision means fewer legitimate users are incorrectly flagged as attackers.

#### 3. Recall (Detection Rate)

Recall measures the model's ability to detect actual attacks from all existing attack samples.

Recall = True Positives / (True Positives + False Negatives)

A high recall value indicates strong detection capability for malicious traffic.

#### 4. False Positive Rate (FPR)

False Positive Rate measures the percentage of legitimate requests incorrectly classified as malicious.

FPR = False Positives / (False Positives + True Negatives)

Lower FPR is essential to maintain a smooth user experience and avoid blocking valid users.

#### 5. Reconstruction Loss Distribution

Since the system is based on an Auto encoder, reconstruction loss plays a key role. The distribution of loss values for normal and malicious traffic is analyzed to determine the anomaly threshold. Clear separation between these distributions indicates effective anomaly detection.

These performance metrics collectively provide a comprehensive evaluation of the proposed WAF system.

#### CONCLUSION

This project successfully demonstrates the design and implementation of a Transformer-Based End-to-End Web Application Firewall capable of detecting malicious web traffic using deep learning techniques. Unlike traditional rule-based WAF systems that rely on static signatures, the proposed system utilizes a Transformer Auto encoder architecture to learn the contextual and structural patterns of legitimate HTTP requests. By training exclusively on benign traffic, the system establishes a strong baseline of normal web behavior. Any deviation from this learned pattern results in higher reconstruction loss, enabling effective anomaly detection. Experimental results show that the system achieves high detection accuracy while maintaining a low false positive rate. The real-time inference pipeline ensures that incoming requests are evaluated instantly without significantly increasing latency. The integration with web servers such as Apache or Nginx demonstrates the practical feasibility of deploying the system in live environments. Furthermore, the incorporation of incremental learning enhances adaptability, allowing the WAF to evolve alongside changing web application behavior. This addresses the critical limitation of traditional WAFs in handling zero-day attacks and obfuscated payloads. Overall, the project validates that deep learning-based anomaly detection provides a robust, scalable, and proactive approach to modern web application security.

#### REFERENCES

1. Ian Good fellow, Yoshua Bengio, and Aaron Courville, *Deep Learning*, MIT Press, 2016.
2. Vaswani, A., Shazeer, N., Parmar, N., et al., "Attention Is All You Need," *Proceedings of the 31st International Conference on Neural Information Processing Systems (NeurIPS)*, 2017.
3. Sommer, R., and Paxson, V., "Outside the Closed World: On Using Machine Learning for Network Intrusion Detection," *IEEE Symposium on Security and Privacy*, 2010.
4. Kruegel, C., and Vigna, G., "Anomaly Detection of Web-based Attacks," *Proceedings of the 10th ACM Conference on Computer and Communications Security*, 2003.
5. Scarfone, K., and Mell, P., "Guide to Intrusion Detection and Prevention Systems (IDPS)," *NIST Special Publication 800-94*, National Institute of Standards and Technology, 2007.
6. OWASP Foundation, "OWASP Top 10 – Web Application Security Risks," Available: <https://owasp.org/www-project-top-ten/>
7. Apache Software Foundation, "Apache HTTP Server Documentation," Available: <https://httpd.apache.org/docs/>
8. Nginx Inc., "Nginx Official Documentation," Available: <https://nginx.org/en/docs/>
9. Chandola, V., Banerjee, A., and Kumar, V., "Anomaly Detection: A Survey," *ACM Computing Surveys*, Vol. 41, No. 3, 2009.
10. Brownlee, J., *Deep Learning for Natural Language Processing*, Machine Learning Mastery, 2019.