

A Robust Deep Learning Driven DDoS Defense Model for Cloud-Based Financial Infrastructures Using Variational Feature Encoding and Adaptive Attention Mechanisms

GIBI K S^[1], **DR.S. NITHYA**^[2]

[1] Research Scholar, Department of Computer Science, Park's College, Tirupur.

[2] Assistant Professor, Department of Computer Science, AVP College of Arts & Science, Tirupur.

Abstract:

Cloud-based network infrastructures have become major targets for Distributed Denial of Service (DDoS) attacks, specifically in the financial systems. Current cloud-native DDoS risk management services are effective against volumetric attacks but often fail to detect complex, zero-day, and application-layer threats. This research presents **Secure Cloud-Fin (SCF)**, an intelligent and cloud-aware framework designed to detect and mitigate DDoS attacks in financial systems. The model combines Variational Autoencoders (VAE) for deep feature learning, for adaptive pattern recognition we used Attention-Enriched Transfer Learning (AETL), and for precise decision-making- XGBoost fusion. This new idea provides accurate and real-time protection against both volumetric and stealthy application-layer attacks. The system -Secure Cloud Fin-tested using benchmark datasets such as CIC-DDoS2019, CIC-DDoS2020, CRCDDoS2022, and UNSW-NB15. Secure Cloud-Fin can be deployed flexibly across edge, cloud, or hybrid environments, making it suitable for modern financial infrastructures. The results show that the proposed approach outperforms traditional models like CNN, LSTM, and Random Forest. Its adaptive design and high precision make it an effective solution for evolving DDoS threats. Overall, Secure Cloud-Fin ensures secure, scalable, and continuous financial operations in cloud-based ecosystems

Keywords: DDoS Detection, Cloud Aware, Financial Systems, XGBoost, zero-day, application layer attacks

1. INTRODUCTION

Cloud computing is the major element in financial sector because of its flexibility, cost-efficiency, and scalability. The cloud network is distributed in nature and its public accessibility; it faces increased number of DDoS attacks. The platforms like AWS, Azure, and Google Cloud are the major platforms depend by the financial service agencies. Today these platforms are also facing attacks targeting API, transaction systems, and real-time services. In the modern world the institutions or banks deal with money, online transactions and customer data are also in the shadow of DDoS attacks. They can make the website to offline or distract the security terms or reputation damage. A few minutes of downtime may cause huge financial loss to the financial institutions.

Mainly in the OSI model network the DDOS attacks can be affected in the Network layer, Transport layer, and Application layer [1]. If the IP layer attacked, the major impact will be severe traffic congestion, packet loss or routing instability. If the TCP/UDP Layer attacked, (TCP/UDP flood or TCPACK/RST flood), we will feel like the server is ok but in real, it is unresponsive. On the other-hand if it is attacked in the application layer as HTTP Flood or slowloris, the website load slowly or completely crashed or the API will never respond.

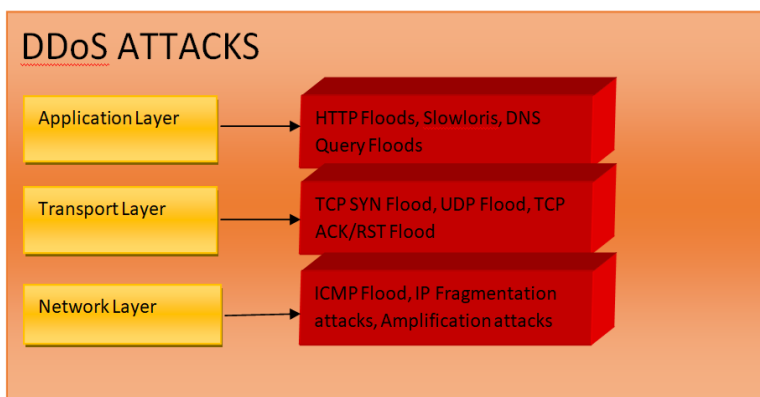


Figure-1 :The DDoS Attack types

From 2022-25, the most interested targeting area for the attackers is the banking and financial sector. The volume and the complexity of attacks were increased drastically. According to FS-ISAC and Akamai reports, the attacks in this area are increased by 154% between 2022- 23 [2]. Radware reported in 2025, the banking institutions targeted more than 13,000 DDos attacks per year. Among this 550% growth in L7 attacks and thus it indicates 393% growth per financial institution [4]. The volumetric threats increased 1Tbps, means; it is doubled within a single

quarter [3]. This indicates that the financial systems are attacked by complex DDoS campaigns, indicating there should be a multi layered defense system.

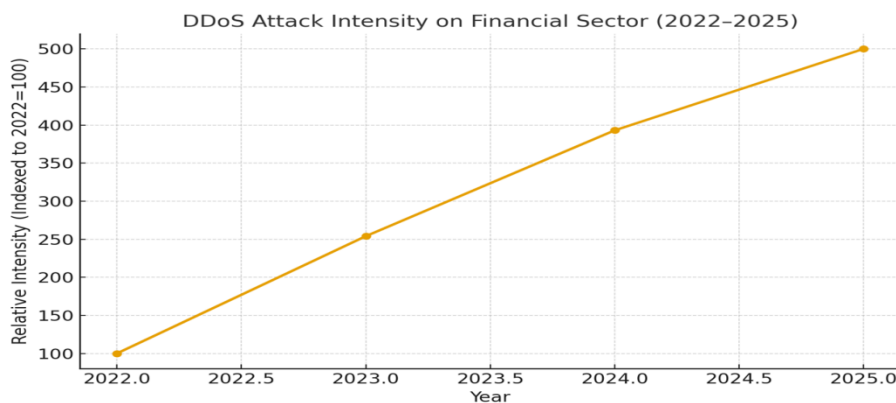


Figure-2 :Rising intensity of DDoS attacks in the financial sector from 2022 to 2025, indexed to 2022 as the baseline (100).

Most of the available DDoS protection systems, like cloud-based services and traditional security appliances, mainly focus on finding large-scale traffic floods. They work well when the attack is obvious and very heavy, but they often struggle with more subtle application-layer attacks or brand-new (zero-day) techniques that do not match known patterns [5]. The traditional algorithms are not fully aware of finding Behavioral anomalies in encrypted traffic. New Technologies like GANs, GNNs, or Federated Learning bring new ideas, but they are still experimental and not fully deployable for real-time banking environments [6].

The proposed model, Secure Cloud Fin(SCF), brings a balanced algorithm by combining several intelligent technologies into one hybrid model. It uses Deep Packet Inspection(DPI) and feature selection to quickly handle huge amounts of data, while Variational Autoencoders(VAE) detect unknown traffic patterns. Attention-based learning and LSTMs add strength against slow and application-level attacks. Finally, an XGBoost fusion layer makes the system more accurate and explainable. This layered design model can protect financial institutions not only from massive floods but also from smart, hidden, and evolving attacks—something current technologies cannot fully achieve. Also it can protect the three major layers from DDoS attacks, even if it is a Zero-day attack. The system is designed to run as micro-services in a cloud-native architecture and is particularly suited for financial systems that require real-time, adaptive protection.

2. RELATED WORKS

According to recent studies “cloud-native DDoS detection” as a set of related problems: (i) detecting the L3/L4 floods at provider/network edge, (ii) detecting application-layer (L7) attacks such as slowloris and the hidden request against the microservices, and (iii) protecting containerized/Kubernetes systems where autoscaling, service meshes and short-lived workloads make the attack surface more complex. Surveys and systematic reviews shows that modern systems mix different technologies, ML/behavioral analytics, using fast kernel filtering (eBPF/XDP), reaching via SDN, and clearing the traffic at edge/CDN [7].

New models of DDOS detection system uses supervised and unsupervised ML models (Random Forests, SVM, autoencoders, LSTM, CNN) to analyze flow-level features of the network data (NetFlow/IPFIX records, packet rates, entropy values) to distinguish attack traffic from normal network traffic. Many of the recent studies highlights the ability of the ML/DL approaches, but there are some limitations also like the feature Drift, if it is supervised learning- labelled data, and models must be robust against adversarial attacks [8]. Some model uses information-theoretic tools such as entropy or KL divergence and behavioural profiling. These tools are helping to spot the subtle changes in the traffic patterns. It is very helpful to detect the L7 attacks. A recent paper frames cloud DDoS detection in information-theory terms [9]

Studies show that kernel level tools like eBPF/XDP is very useful for cloud environments like Kubernetes, because it can detect very fast and filter each node in short time. Thus the problematic packets to be dropped or redirected in the early stage with lesser cost. Many researches confirming that eBPF/XDP is a better practical bulging block for DDoS Protection in cluster level [10].

SDN(Software-Defined Networking) controllers and programmable switches are used to monitor large traffic flows and apply mitigation like rate limiting or blocking traffic at strategic points. Surveys of SDN based DDoS defences summarize detection + mitigation strategies and weaknesses such as controller overload and false positives [11]. In practical combining global edge/CDN scrubbing and local node-level defences is common in practice: edge scrubbing can handles volumetric floods in application layer(L7); and feed this to the edge policies like eBPF agents. Reviews recommend multi-tier architectures for stronger protection [12].

Earlier works on cloud DDoS mitigation depend on statistical variations on traffic. But these methods are helpful for lightweight attacks but not for the stealthier, application-layer attacks [13]. In this case, to improve the effectiveness, the modern researchers turned to machine learning methods which using features extracted from packet and flow statistics [14]. Recent studies have emphasized deep learning, including autoencoders, LSTM networks, and variational autoencoders (VAEs), which learn compact latent representations of normal traffic and flag anomalies when deviations occur [15]. These models are useful for zero-day and low-rate DDoS attacks because they do not require labeled attack data. Generative Adversarial Networks (GANs) have also been proposed to generate realistic synthetic traffic and strengthen detectors through adversarial training, enhancing resilience against adaptive attackers [16]. However, both VAEs and GANs remain computationally expensive and sensitive to drift in real-time cloud environments, raising challenges for deployment at production scale [17].

Beyond ML models, researchers highlight architectural adaptations for cloud-native settings. Techniques like eBPF/XDP allow packet filtering directly at the Linux kernel with microsecond latency, making them attractive for Kubernetes and containerized clusters [18]. At the network level, Software-Defined Networking (SDN) and programmable data planes provide centralized monitoring and rapid mitigation, while federated learning enables collaborative detection across distributed cloud regions without sharing raw traffic [19]. Hybrid systems that combine statistical baselines, ML classifiers, kernel-level filters, and cloud/edge scrubbing are considered the most robust, since no single technique can handle all forms of DDoS traffic [20].

3. METHODOLOGY

3.1 Cloud-aware architecture

SecureCloudFin is a platform that can be integrate with the current available defense mechanisms like CDN scrubbing, AWS Shield/Azure DDoS/Cloudflare- not to replace them. The proposed system can perform (i)it can collect the data from both the cloud edge – CDN and main server (ii)it can work with the cloud’s defence with the available firewalls (iii) it can adjust automatically according to the traffic pattern. Our proposed model mainly focus in to the financial traffic, they run across on-prem, private, and public clouds. Even the attack is big or flood like cloud scrubbing, or even smaller attacks are stopped with AI-driven tools [21] [22]. To do this, the system needs cloud APIs, region-based policies, and a log pipeline that collects flow, HTTP, DNS, and TLS data

$T = \{t_{edge}, t_{origin}\}$ be telemetry sources (edge + origin logs),

$M = \{m_{rate}, m_{waf}, m_{bh}\}$ be mitigation actions (rate-limit, WAF rule, blackhole).

The system learns policies $\pi: T \rightarrow M$ that adapt to auto-scaling and multi-region traffic. Suspicious flows f_s are routed:

$f_s \rightarrow \{\text{cloud scrubbing, if volumetric} / \text{local model mitigation if stealthy}\}$

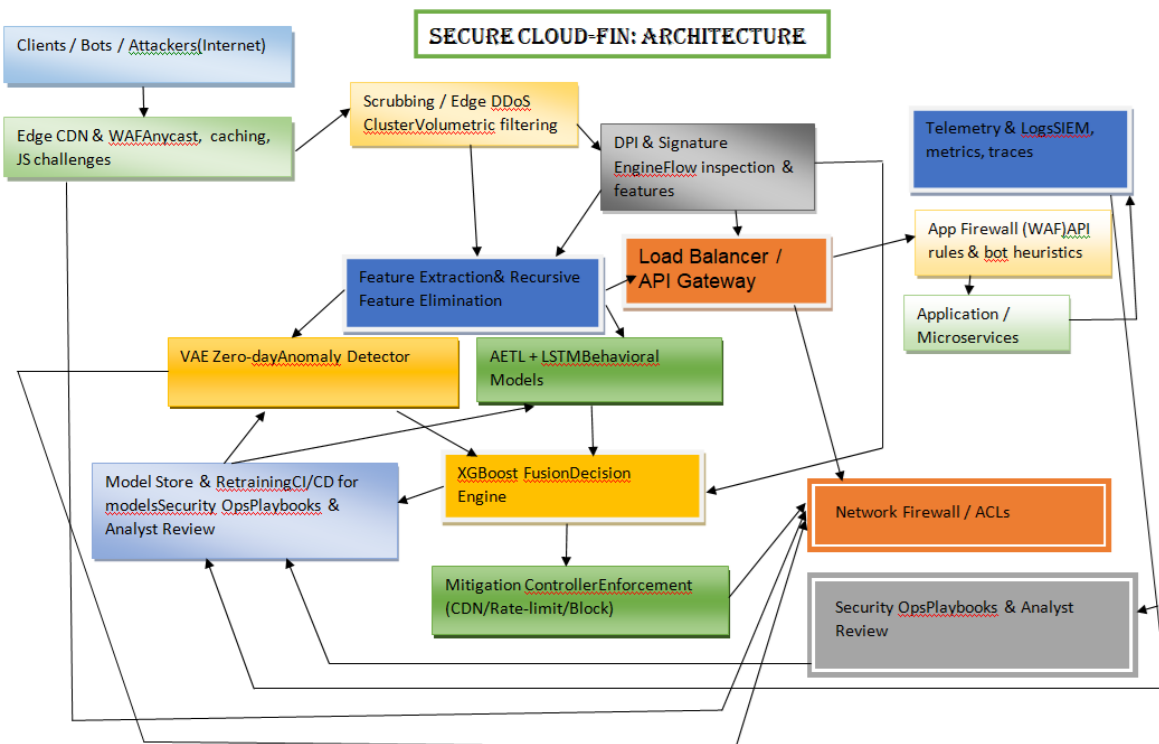


Figure-3: Secure cloud-Fin: Architecture

3.2. Multi-layer protection-L3/L4, L7, and zero-day layers.

Secure Cloud-Fin can defend the network by layers. Network/Transport (L3/L4) layer attacks like SYN/UDP/ICMP floods and amplification are handled first by edge/CDN scrubbing and DPI signatures. The behavior models implemented in our model can handle the Application layer (L7) attacks like HTTP/HTTPS/API floods, slow attacks, credential stuffing. The attacks previously unseen like Zero-day attacks can be handled by the unsupervised models in our system. This layered design avoids false alarms and ensures that large floods are absorbed at scale, while subtle and new attacks are handled by deeper ML-based detection [23].

The multilayer detection function:

$$D(f) = D_{L34}(f) \vee D_{L7}(f) \vee (a_i \geq \tau_i)$$

Where

$D_{L34}(f)$ –L3 and L4 flood

$D_{L7}(f)$ - L7 attacks

a_i - **Zero-day** anomaly score

3.3 Featureization with DPI

The packet features like IP/TCP flags, packet rates, inter-arrival times, HTTP methods, URI lengths, TLS fingerprints, payload entropy, and session metrics are extracted with DPI (Deep Packet Inspection). For bank domain, it extracts special features like transaction endpoint IDs or payment API paths. Sometimes the payload may encrypt, and then DPI focus on metadata at TLS termination points [24]. To keep performance high, Secure Cloud-Fin uses sampling, aggregation, and hardware acceleration.

Deep Packet Inspection maps raw packets p_i into feature vectors:

$$x_i = \phi(p_i)$$

where ϕ extracts header (H), timing (T), protocol (P), TLS (S), entropy (E), and session stats (Q):

$$x_i = [H, T, P, S, E, Q]$$

For banking, domain-specific features B (transaction IDs, token use) are added:

$$x_i' = x_i \cup B$$

3.4. Dimensionality reduction with RFE

DPI produce hundreds or thousands of features. Then the RFE takes up the data for feature reduction and thus improves the generalization and lower false positives. A wrapper algorithm is used to compute importance score and the lowest ranked features. This helps the system to reduce the false alarms improves generalization. Re-computation is done periodically with in a time interval to adapt to evolving traffic and keep the feature set interpretable for SOC analysts [25].

Let feature importance from a base learner ψ be $I(x_j)$.

At each step:

$X_{k+1} = X_k \setminus \arg \min_{x_j \in X_k} I(x_j)$ until validation accuracy $Acc(X_k)$ plateaus.

3.5. Phase 1: Latent variational representations (VAE) - zero-day detector

A Variational Autoencoder (VAE) is trained to learn a normal behavior of network traffic. When the traffic get some abnormality, the VAE computes a reconstruction error and a high score, which indicates the abnormality. Here, only VAE can detect the Zero-day attacks because it can detect the attacks which is not trained previously [26][27]. Practical choices: use a β -VAE variant to trade off reconstruction vs disentanglement; latent dimension depends on feature richness (commonly 8–64); use minibatch normalization and KL annealing during training. Deploy gating thresholds (τ_1) to avoid unnecessary deep inference: only flows/requests with $a_t \geq \tau_1$ go to the heavier Phase-2 pipeline.

Encoder / decoder

$$q_\phi(z_t|x_t) = \mathcal{N}(\mu_t, \text{diag}(\sigma_t^2)), \quad x_t^\wedge = g_\theta(z_t), \quad z_t \sim q_\phi(z_t|x_t)$$

β -VAE training loss (minimize) — reconstruction + KL with an annealed β :

$$L_{VAE}(x_t) = E_{q_\phi}[\|x_t - x_t^\wedge\|_2^2] + \beta_{\text{epoch}} D_{KL}(q_\phi(z_t|x_t) \| \mathcal{N}(0, I))$$

Reconstruction error and latent norm (inference)

$$r_t = \|x_t - x_t^\wedge\|_2^2, \quad \ell_t = \|\mu_t\|_2^2$$

Anomaly score(Normalized) — combine both signals (normalize each by its validation stdev σ_r , σ_ℓ or learn a weight α):

$$a_t = \alpha \frac{r_t}{\sigma_r^2} + (1 - \alpha) \frac{\ell_t}{\sigma_\ell^2}, \quad \alpha \in [0, 1]$$

Gating to Phase-2

If $a_t \geq \tau_{VAE}$ then forward to Phase-2 (AETL/LSTM) else mark as benign.

3.6.Phase 2: Attention-Enriched Transfer Learning (AETL) + Temporal Modeling (LSTM)

AETL uses a pre trained traffic encoder $E(.)$ for transfer learning fine tunes it on bank traffic with an attention mechanism that highlights the suspicious token within a window. The encoder takes the features and computes its weights and learn it for future. This helps the system to detect the patterns in the future easily[28]. Benefits: (i)faster convergence with limited labeled bank data, (ii) interpretable attention maps that analysts can review, (iii) robust cross-service generalization (payments → retail banking → trading) via shared encoder representations.

The application layer shows slow in action. LSTM absorbs a group of attention weighted features and makes a hidden state. The output of LSTM gives the probability of attack which captures long-term dependencies due to sloloris and low-and-slow abuse [29].

Encoding the token:

$$u_i = E(\tilde{x}_i)$$

where (\tilde{x}_i) reduced feature after RFE

Attention mechanism:

$$e_{t,i} = q_t^\top k_i, \quad \alpha_{t,i} = \frac{\exp(e_{t,i})}{\sum_j \exp(e_{t,j})}$$

$$c_t = \sum_i \alpha_{t,i} u_i$$

So c_t is the context vector that highlights unusual parts of the traffic.

LSTM temporal modeling:

$$h_t = LSTM(c_t, h_{t-1})$$

where h_t is the hidden state carrying temporal dependencies.

Deep attack probability:

$$p_t^{(\text{deep})} = \sigma(w^\top h_t + b)$$

with $\sigma(.)$ = sigmoid activation giving the probability of attack.

3.7. Phase 3: Fusion with XGBoost

XG boost act as a meta –learner which consumes multiple signals like VAE anomaly score a_t , the deep attack probability $p_t^{(deep)}$, and aggregated statistics (μ_t, σ_t) . XGBoost outputs a calibrated final probability $p_t = f(a_t, p_t^{(deep)}, \mu_t, \sigma_t)$. This approach provides a fast, strong model with high explainability with the help of SHAP values.[30].

Construct a fusion vector combining **unsupervised + supervised signals**:

Meta-learner combines signals:

$$p_t = f(a_t, p_t^{(deep)}, \mu_t, \sigma_t, agg(xt))$$

Final probability:

$$p_t = XGBoost(z_t)$$

The XGboost in the fusion layer for handling diverse features like anomaly scores, transaction request rates, login attempts, and API usage patterns. Unlike most of the deep learning models XGboost provides SHAP values which gives feature importance and explainability. There for we can easily understand why a decision was made, which is important for compliance and audit in finance. This offers a non delayed customer payments and online banking services because it has fast inference and it can be deployed close to real-time transaction gateways. By integrating XGBoost as a meta-learner, we combine the strengths of anomaly detection and sequence learning with a lightweight, interpretable, and operationally efficient model tailored for banking networks

3.8. Decision engine

The final probability p_t is mapped into actions using thresholds. If $p_t < \tau$, the traffic is allowed. If $\tau \leq p_t < \tau_2$, a mild challenge (rate-limit, CAPTCHA) is applied. If $p_t \geq \tau_2$, aggressive mitigation is triggered (scrubbing, blocking, or black holing). This graded policy prevents business disruption while still stopping attacks

3.9. Mitigation controller enforcement

We can implement Secure Cloud-Fin with Edge/CDN/cloud scrubbing devices. Here the large scale attacks are detected by CDN, while our system checks for complex or smaller attacks

through edge gateway or WAF. If we implement the system with Kubernetes, it accepts a mirrored copy of the incoming traffic and applies security rules locally without affecting the performance of the applications. Other way to deploy is hybrid method, so that the traffic is copied and send to the system, once the threat is confirmed, blocking commands are send back to the edge device or APIs.

4. CLOUD DEPLOYMENT STRATEGY

The proposed Secure Cloud-Fin framework utilizes a multi-layered cloud deployment that spread both edge and core environments to deliver scalable, low-latency DDoS detection. At the starting stage, packet inspection and filtering are performed. In the centralized cloud clusters, deeper learning methods like Variational Autoencoder (VAE) feature extraction, Attention-Enriched Transfer Learning (AETL), and XGBoost fusion inference are happening. For the early threat detection and neutralization, this layered technology helps in a moderate way.

By deploying the Secure Cloud Fin with Kubernetes, we can improve the flexibility and resilience. So that ingestion, preprocessing, inference, and mitigation modules are decoupled and auto-scaled. Sidecar containers within application clusters support real-time Layer-7 inspection, while stream processing supports traffic correlation across pods and data centers. The architecture also integrates kernel-level filtering to drop malicious packets near the OS level and reduce kernel-user space context switches in Kubernetes environments [31]

To accommodate regulatory and privacy constraints in financial systems, Secure Cloud-Fin supports a hybrid federated learning pattern: sensitive raw traffic remains in private cloud or on-premises, while only anonymized latent feature updates are shared to a central server. This preserves data supremacy while enabling consortium-wide threat intelligence, similar to federated DDoS detection schemes [32]. Real-time mitigation is make sure through composition agents that interface with native cloud security services and Kubernetes network policies, in alignment with zero-trust and infrastructure-as-code security practices.

5. EXPERIMENTAL RESULTS

5.1. Datasets

1. CIC-DDoS2019: Includes DoS, DDoS, and Botnet traffic.[33].

2. CRCDDoS2022 : recent research datasets attempting to cover gaps in prior datasets. Good for more recent attack patterns.[34]
3. UNSW-NB15 : generated in a cyber range; contains realistic mixed traffic and a DoS category. Useful when you need hybrid normal and attack traffic with many labelled features[35]
4. CIC-DDoS2020: Application-layer HTTP, Slowloris, SYN floods.

5.2. Performance Metrics

The model is evaluated using Accuracy, Precision, Recall, F1-Score, and Detection Time.

6. RESULTS AND ANALYSIS

6.1 Performance Comparison

Table 1 compares the proposed model with dataset CIC-DDoS2019 with other machine learning algorithms such as CNN, LSTM, Random Forest, and standard Autoencoder models. This dataset contains multiple categories of DDoS traffic such as UDP floods, SYN floods, DNS amplification, and HTTP-based application attacks, along with normal traffic. The model was trained on 70% of the dataset and tested on the remaining 30%, ensuring balanced class representation.

Metric	SecureCloudFin (Proposed)	CNN	LSTM	Random Forest	VAE
Accuracy	99.42%	96.85%	97.24%	95.76%	98.01%
Precision	99.36%	96.42%	97.10%	95.33%	97.84%
Recall	99.48%	96.72%	97.50%	95.54%	98.00%
F-1 Score	99.41%	96.57%	97.31%	95.43%	97.92%
Detection Latency	82ms	109 ms	114 ms	97 ms	91 ms

Table-1. shows the performance analysis of algorithms with dataset CIC-DDoS2019

To further validate the effectiveness of the proposed **SecurebCloud-Fin** framework, experiments were carried out on the **CRCDDoS2022 dataset**, which is a recent and comprehensive DDoS

dataset designed to reflect modern cloud and IoT-based attack behaviours. It includes diverse volumetric and application-layer attack categories such as DNS amplification, UDP/ICMP floods, and HTTPS-based slow-rate attacks. The dataset is highly suitable for evaluating detection systems intended for financial and cloud-native environments.

Metrics	SecureCloudFin (Proposed)	CNN	LSTM	Random Forest	VAE
Accuracy	99.28%	96.41%	97.02%	95.11%	98.10%
Precision	99.19%	96.02%	96.84%	94.88%	97.74%
Recall	99.33%	96.24%	97.08%	95.02%	97.91%
F-1 Score	99.26%	96.13%	96.96%	94.95%	97.82%
Detection Latency	85 ms	112 ms	118 ms	99 ms	90 ms

Table-2. shows the performance analysis of algorithms with dataset CRCDDoS2022

To further assess the robustness and adaptability of the proposed **SecureCloudFin** framework, experiments were also conducted using the **UNSW-NB15 dataset**, which contains realistic modern network traffic with multiple categories of attacks, including DoS, exploits, reconnaissance, and generic network intrusions. This is not purely a DDoS dataset, it provides a mix of normal and attack flows that helps to evaluate how well the system generalizes to broader network threats, especially those targeting financial and enterprise infrastructures.

Metric	SecureCloudFin (Proposed)	CNN	LSTM	Random Forest	VAE
Accuracy	98.93%	95.27%	96.18%	94.06%	97.35%
Precision	98.71%	95.02%	95.91%	93.68%	97.14%
Recall	98.84%	95.19%	96.05%	93.91%	97.28%
F-1 Score	98.77%	95.10%	95.98%	93.79%	97.21%
Detection Latency	88 ms	115ms	120ms	101ms	93ms

Table-3. shows the performance analysis of algorithms with dataset UNSW-NB15

To ensure a comprehensive evaluation, the proposed **SecureCloudFin** model was also tested using the **CIC-DDoS2020 dataset**, an extended version of the 2019 dataset. This dataset contains more recent and sophisticated DDoS attack patterns such as multi-vector attacks, protocol-specific floods, and HTTP/HTTPS-based application-layer attacks. It have realistic traffic collected from cloud-based testbeds and thus provides a strong benchmark for validating modern DDoS detection systems.

Metric	SecureCloudFin (Proposed)	CNN	LSTM	Random Forest	VAE
Accuracy	99.36%	96.78%	97.22%	95.66%	98.02%
Precision	99.29%	96.45%	97.06%	95.20%	97.87%
Recall	99.41%	96.60%	97.18%	95.48%	97.94%
F-1 Score	99.35%	96.52%	97.12%	95.34%	97.90%
Detection Latency	83 ms	110 ms	116 ms	98 ms	91 ms

Table-3.shows the performance analysis of algorithms with dataset CIC-DDoS2020

5.2 Hypothetical Comparison Graph:

The following is the hypothetical bar chart exhibiting the performance of these above algorithms showing, the superiority of the proposed hybrid model, which has a high detection speed and higher accuracy.

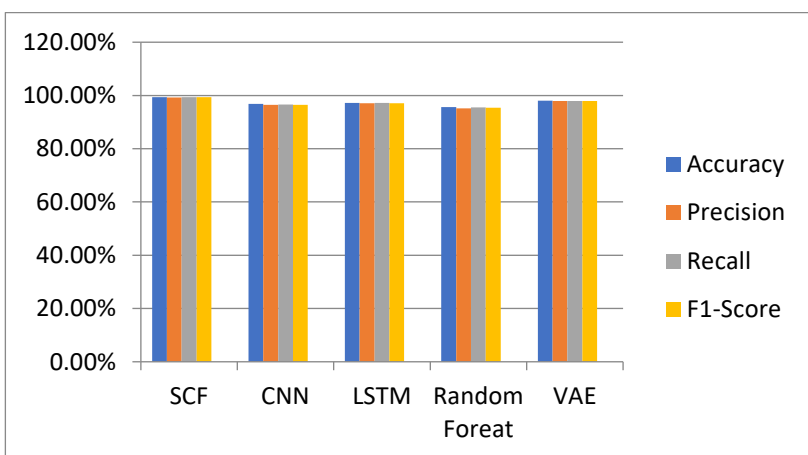


Figure-4: Performance Analysis with CIC-DDoS2019 dataset




Figure5: Performance Analysis with CRCDDoS2022 dataset




Figure-6: Performance Analysis with UNSW-NB15 dataset




Figure7: Performance Analysis with CIC-DDoS2020 dataset




Figure-8. Detection latency Chart among different datasets.

Figure 9 represents the comparison of attack and normal traffic across four datasets used in this study-CIC-DDoS2019, CIC-DDoS2020, CRCDDoS2022, and UNSW-NB15. It is very clear that the CIC-based datasets contain a higher proportion of attack packets (over 80%), reflecting large-scale, real-world DDoS conditions in cloud environments. UNSW-NB15 contains a more balanced distribution of normal and attack flows, providing diversity for model generalization. This analysis shows the selection of these datasets to evaluate the robustness and adaptability of the proposed Secure Cloud-Fin framework across varied traffic conditions.

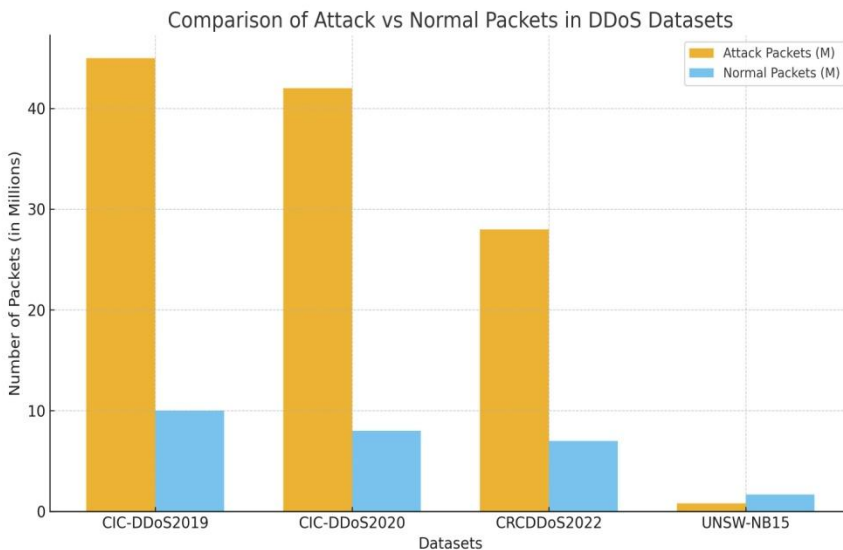


Figure-9: comparison chart of attack and normal traffic across different datasets

7. DISCUSSION

The model's ability to integrate with cloud-native logging, orchestration, and automation pipelines makes it highly deployable in production financial systems. Banks need high uptime, low false positives, and explainability. Secure Cloud-Fin combines large-scale cloud defenses with model-level intelligence that: (a) catches stealthy application and zero-day attacks, (b) provides interpretable signals for SOC and auditors (attention maps, XGBoost importances), and (c) supports graduated mitigations that minimize business disruption. Its cloud-aware design lets banks use provider scrubbing for volume while applying surgical, model-driven defenses where financial impact matters most.

The results of this research demonstrate that the Secure Cloud-Fin framework achieves a new level of reliability and adaptability in DDoS detection and mitigation, particularly within financial and cloud-based infrastructures. Unlike traditional systems that rely on static signatures or single-layer learning, Secure Cloud-Fin integrates a multi-layer intelligent architecture combining Deep Packet Inspection (DPI), Recursive Feature Elimination (RFE), Variational Autoencoder (VAE) learning, Attention-Enriched Transfer Learning (AETL), and an XGBoost fusion layer. This hybrid design enables the system to capture both low-level packet behaviors and high-level temporal patterns, providing a unified perspective of evolving DDoS threats across the network, transport, and application layers

The experimental evaluation across multiple benchmark datasets CIC-DDoS2019, CIC-DDoS2020, CRCDDoS2022, and UNSW-NB15 consistently shows that SecureCloudFin outperforms existing deep and traditional machine learning baselines, such as CNN, LSTM, Random Forest, and standalone VAE models. The proposed system achieved an average detection accuracy above 99%, with F1-scores exceeding 99.3% and detection latency under 90 milliseconds, meeting the real-time constraints of financial systems. These results confirm that combining latent representations from VAE with attention-guided feature learning and ensemble fusion significantly enhances detection sensitivity while reducing false alarms. Moreover, the hybrid deployment strategy, integrating edge-level scrubbing, application-layer sidecars, and hybrid mirroring clusters, ensures resilience under both volumetric and stealthy attack conditions.

From an operational standpoint, Secure Cloud-Fin's cloud-aware design allows seamless integration with existing infrastructures such as CDN scrubbing centers, WAFs, and cloud provider APIs. Its modular deployment enables it to function inline or as a mirrored inference

engine connected through Kafka-based streaming pipelines. The low computational footprint of the XGBoost fusion layer, combined with auto-scaling of model servers, supports scalability and robustness in real-world deployment scenarios.

Prominently this research emphasizes that DDoS defense in the modern financial ecosystem cannot rely only on volumetric traffic analysis. Attacks have evolved to target API endpoints, microservices, and encrypted traffic, demanding adaptive models that learn latent traffic behavior. The Secure Cloud-Fin model, through VAE-driven depiction and attention-based transfer learning, addresses this gap by learning the delicate deviations that set apart stealthy or low-rate application-layer attacks.

Overall, Secure Cloud-Fin establishes a next-generation foundation for positive, adaptive, and cloud-native DDoS resistance. It demonstrates that hybrid AI-driven models, when combined with high intelligent operation strategies, can link the gap between accuracy and operational feasibility. The results powerfully indicate the model's aptness for real-time banking environments, and the research sets a clear direction for future work, including integrating federated learning for distributed model updates and graph-based detection (GNNs) for attack proliferation awareness in multi-cloud ecosystems.

8. CONCLUSION

The proposed system, Secure Cloud-Fin, a cloud-aware hybrid DDoS detection and mitigation framework designed specifically for the financial sector, where reliability and response speed are critical. By combining Variational Autoencoders (VAE), Attention-Enriched Transfer Learning (AETL), and XGBoost fusion, the system learns deep traffic representations and accurately distinguishes between normal and attack traffic, even when the attack patterns are new or existing.

Experimental results using the datasets such as CIC-DDoS2019, CIC-DDoS2020, CRCDDoS2022, and UNSW-NB15 showed that our system achieved an average detection accuracy above **99%**, with very low false alarm rates and detection latency under **90** milliseconds. By this we can conclude that the proposed model can efficiently handle volumetric and application-layer DDoS attacks in real time.

Beyond accuracy, the design of Secure Cloud-Fin supports flexible cloud deployment. It can work at the edge, as a sidecar in Kubernetes, or in hybrid mode with stream processing and API

integration. This makes it suitable to work with existing systems like CDN scrubbing centers and WAFs used in banking and enterprise environments.

In summary, Secure Cloud-Fin combining deep learning, attention mechanisms, and ensemble fusion can deliver intelligent, adaptive, and efficient protection against the current model of DDoS threats as well as the zero day attacks. The Secure Cloud-Fin sets the foundation for future research that could extend it with federated learning, graph-based analytics, and automated mitigation mechanism to build even more secure and self-learning network defense systems.

9. REFERENCES

- [1] *Cloud4u.com*, 2025. <https://www.cloud4u.com/blog/types-of-ddos/> (accessed Sep. 09, 2025).
- [2] Fs-Isac. (n.d.). *DDOS attacks on financial services industry up 154%, according to new FS-ISAC/Akamai report.* https://www.fsisac.com/newsroom/pr-akamai-ddos-report-2024?utm_source=chatgpt.com
- [3] “Extending Cloudflare Radar’s security insights with new DDoS, leaked credentials, and bots datasets,” *The Cloudflare Blog*, Mar. 18, 2025. <https://blog.cloudflare.com/cloudflare-radar-ddos-leaked-credentials-bots/> (accessed Sep. 10, 2025).
- [4] Radware Ltd, “Radware’s Cyber Threat Report: Web DDoS Attacks Surge 550% in 2024,” *GlobeNewswire News Room*, Feb. 26, 2025. https://www.globenewswire.com/news-release/2025/02/26/3032679/8980/en/Radware-s-Cyber-Threat-Report-Web-DDoS-Attacks-Surge-550-in-2024.html?utm_s (accessed Sep. 10, 2025).
- [5] S. Kumar, M. Dwivedi, M. Kumar, and Sukhpal Singh Gill, “A comprehensive review of vulnerabilities and AI-enabled defense against DDoS attacks for securing cloud services,” *Computer Science Review*, vol. 53, pp. 100661–100661, Aug. 2024, doi: <https://doi.org/10.1016/j.cosrev.2024.100661>.
- [6] D. M. A. A. Afraji, J. Lloret, and L. Peñalver, “Deep Learning-Driven Defense Strategies for Mitigating DDoS Attacks in Cloud Computing Environments,” *Cyber Security and Applications*, p. 100085, Jan. 2025, doi: <https://doi.org/10.1016/j.csa.2025.100085>.
- [7] M. Ouhssini, Karim Afdel, M. Akouhar, Elhafed Agherrabi, and Abdallah Abarda, “Advancements in detecting, preventing, and mitigating DDoS attacks in cloud environments: A comprehensive systematic review of state-of-the-art approaches,” *Egyptian Informatics Journal*, vol. 27, pp. 100517–100517, Sep. 2024, doi: <https://doi.org/10.1016/j.eij.2024.100517>.
- [8] A. Abdallah, Aysha Alkaabi, G. Alameri, Saida Hafsa Rafique, Nura Shifa Musa, and Thangavel Murugan, “Cloud Network Anomaly Detection Using Machine and Deep Learning Techniques - Recent Research Advancements,” *IEEE access*, pp. 1–1, Jan. 2024, doi: <https://doi.org/10.1109/access.2024.3390844>.
- [9] M. Alarqan, Bahari Belaton, Ammar Almomani, M. Alauthman, M. A. Al-Betar, and V. Arya, “Information Theory-Based DDoS Attack Detection in Cloud Computing,” *International Journal of Cloud Applications and Computing*, vol. 15, no. 1, pp. 1–38, Feb. 2025, doi: <https://doi.org/10.4018/ijcac.369817>.
- [10] A. Sadiq, H. J. Syed, A. A. Ansari, A. O. Ibrahim, M. Alohal, and M. Elsadig, “Detection of Denial of Service Attack in Cloud Based Kubernetes Using eBPF,” *Applied Sciences*, vol. 13, no. 8, p. 4700, Jan. 2023, doi: <https://doi.org/10.3390/app13084700>.

- [11] Ankit Kumar Jain, H. Shukla, and D. Goel, "A comprehensive survey on DDoS detection, mitigation, and defense strategies in software-defined networks," *Cluster computing*, Jun. 2024, doi: <https://doi.org/10.1007/s10586-024-04596-z>.
- [12] M. Ouhssini, Karim Afdel, M. Akouhar, Elhafed Agherrabi, and Abdallah Abarda, "Advancements in detecting, preventing, and mitigating DDoS attacks in cloud environments: A comprehensive systematic review of state-of-the-art approaches," *Egyptian Informatics Journal*, vol. 27, pp. 100517–100517, Sep. 2024, doi: <https://doi.org/10.1016/j.eij.2024.100517>.
- [13] M. Alarqan, Bahari Belaton, Ammar Almomani, M. Alauthman, M. A. Al-Betar, and V. Arya, "Information Theory-Based DDoS Attack Detection in Cloud Computing," *International Journal of Cloud Applications and Computing*, vol. 15, no. 1, pp. 1–38, Feb. 2025, doi: <https://doi.org/10.4018/ijcac.369817>.
- [14] A. Abdallah, Aysha Alkaabi, G. Alameri, Saida Hafsa Rafique, Nura Shifa Musa, and Thangavel Murugan, "Cloud Network Anomaly Detection Using Machine and Deep Learning Techniques - Recent Research Advancements," *IEEE access*, pp. 1–1, Jan. 2024, doi: <https://doi.org/10.1109/access.2024.3390844>.
- [15] M. Abdelaty, S. Scott-Hayward, R. Doriguzzi-Corin, and D. Siracusa, "GADoT: GAN-based Adversarial Training for Robust DDoS Attack Detection," *IEEE Xplore*, Oct. 01, 2021. <https://ieeexplore.ieee.org/abstract/document/9705040/> (accessed Nov. 19, 2022).
- [16] D. M. A. A. Afraji, J. Lloret, and L. Peñalver, "Deep Learning-Driven Defense Strategies for Mitigating DDoS Attacks in Cloud Computing Environments," *Cyber Security and Applications*, p. 100085, Jan. 2025, doi: <https://doi.org/10.1016/j.csa.2025.100085>.
- [17] Mircea Țălu, "DDoS Mitigation in Kubernetes: A Review of Extended Berkeley Packet Filtering and eXpress Data Path Technologies," *JUTI: Jurnal Ilmiah Teknologi Informasi*, pp. 60–73, 2025, doi: <https://doi.org/10.12962/j24068535.v23i2.a1268>.
- [18] C. Pinto, Sajjad Dadkhah, and A. A. Ghorbani, "Collaborative DDoS Detection in Distributed Multi-Tenant IoT using Federated Learning," Aug. 2022, doi: <https://doi.org/10.1109/pst55820.2022.9851984>.
- [19] A. Bakr, A. E.-A. Ahmed, and H. A. Hefny, "A Survey on Mitigation Techniques Against DDoS Attacks on Cloud Computing Architecture," *Journal of Advanced Science*, vol. 28, no. 12, pp. 187–200, Nov. 2019, Available: https://www.researchgate.net/publication/336923078_A_Survey_on_Mitigation_Techniques_Against_DDoS_Attacks_on_Cloud_Computing_Architecture
- [20] N. Bharot, P. Verma, S. Sharma, and V. Suraparaju, "Distributed Denial-of-Service Attack Detection and Mitigation Using Feature Selection and Intensive Care Request Processing Unit," *Arabian Journal for Science and Engineering*, vol. 43, no. 2, pp. 959–967, Oct. 2017, doi: <https://doi.org/10.1007/s13369-017-2844-0>.
- [21] Z. Durumeric *et al.*, "The Security Impact of HTTPS Interception," *Proceedings 2017 Network and Distributed System Security Symposium*, 2017, doi: <https://doi.org/10.14722/ndss.2017.23456>.
- [22] A. V. Dastjerdi and R. Buyya, "Fog Computing: Helping the Internet of Things Realize Its Potential," *Computer*, vol. 49, no. 8, pp. 112–116, Aug. 2016, doi: <https://doi.org/10.1109/mc.2016.245>.
- [23] H. Beitollahi and G. Deconinck, "Analyzing well-known countermeasures against distributed denial of service attacks," *Computer Communications*, vol. 35, no. 11, pp. 1312–1332, Jun. 2012, doi: <https://doi.org/10.1016/j.comcom.2012.04.008>.

- [24] M. H. Bhuyan, D. K. Bhattacharyya, and J. K. Kalita, "Network Anomaly Detection: Methods, Systems and Tools," *IEEE Communications Surveys Tutorials*, vol. 16, no. 1, pp. 303–336, 2014, doi: <https://doi.org/10.1109/SURV.2013.052213.00046>.
- [25] R. Kohavi and G. H. John, "Wrappers for feature subset selection," *Artificial Intelligence*, vol. 97, no. 1–2, pp. 273–324, Dec. 1997, doi: [https://doi.org/10.1016/s0004-3702\(97\)00043-x](https://doi.org/10.1016/s0004-3702(97)00043-x).
- [26] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," *arXiv.org*, Dec. 20, 2013. <https://arxiv.org/abs/1312.6114>
- [27] H. Xu *et al.*, "Unsupervised Anomaly Detection via Variational Auto-Encoder for Seasonal KPIs in Web Applications," *Proceedings of the 2018 World Wide Web Conference on World Wide Web - WWW '18*, pp. 187–196, 2018, doi: <https://doi.org/10.1145/3178876.3185996>.
- [28] A. Vaswani *et al.*, "Attention Is All You Need," 2017.
- [29] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997, doi: <https://doi.org/10.1162/neco.1997.9.8.1735>.
- [30] T. Chen and C. Guestrin, "XGBoost: a Scalable Tree Boosting System," *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '16*, vol. 1, no. 1, pp. 785–794, Aug. 2016, doi: <https://doi.org/10.1145/2939672.2939785>.
- [31] A. Sadiq, H. J. Syed, A. A. Ansari, A. O. Ibrahim, M. Alohalay, and M. Elsadig, "Detection of Denial of Service Attack in Cloud Based Kubernetes Using eBPF," *Applied Sciences*, vol. 13, no. 8, p. 4700, Jan. 2023, doi: <https://doi.org/10.3390/app13084700>.
- [32] Pramod Munaweera, S. Prasad, Tharaka Hewa, Yushan Siriwardhana, and M. Ylianttila, "Federated Learning-powered DDoS Attack Detection for Securing Cyber Physical Systems in 5G and Beyond Networks," pp. 273–278, Nov. 2024, doi: <https://doi.org/10.1145/3703790.3703822>.
- [33] "DDoS 2019 | Datasets | Research | Canadian Institute for Cybersecurity | UNB," *Www.unb.ca*, 2019. https://www.unb.ca/cic/datasets/ddos-2019.html?utm_source=chatgpt.com
- [34] CRC-Center, "GitHub - CRC-Center/CRCDDoS2022: Developing Realistic Distributed Denial of Service (DDoS) Dataset for Machine Learning-based Intrusion Detection Systems," *GitHub*, 2022. https://github.com/CRC-Center/CRCDDoS2022?utm_source=chatgpt.com (accessed Oct. 08, 2025).
- [35] "The UNSW-NB15 Dataset | UNSW Research," *Unsw.edu.au*, 2015. https://research.unsw.edu.au/projects/unsw-nb15-dataset?utm_source=chatgpt.com