Policy Research Journal
ISSN (E): 3006-7030 ISSN (P) : 3006-7022 Volume 4, Issue 3, 2026

DEEP LEARNING-BASED INTRUSION DETECTION SYSTEM FOR REAL-TIME CYBERSECURITY IN CLOUD COMPUTING ENVIRONMENTS

Hina Gul[*1], Maaz Ali Mumtaz[2], Waqar Hussain[3], Ameer Jan[4]

[*1]*Assistant professor Computer Science, Government Girls Degree College Matta. Swat*
[2]*Assistant Director IT / M Phil Student, Department of Computer Science, University of Malakand*
[3]*Lecturer, Department of Computer Science, The University of Veterinary and Animal Sciences UVAS, Swat*
[4]*University of Makran Panjgur*

[*1]hinagul.ggdc@gmail.com, [2]maaz021@uom.edu.pk, [3]waqar.hussain@uvasswat.edu.pk
[4]ameerjan@uomp.edy.pk

DOI: https://doi.org/10.5281/zenodo.19399837

Keywords
cloud computing security, intrusion detection system, deep learning, anomaly detection, zero-day attacks, virtualization threats, multi-tenancy, CIC-IDS2018, NSL-KDD, explainable AI, federated learning, real-time cybersecurity

Article History
Received: 31 January 2026
Accepted: 15 March 2026
Published: 31 March 2026

Copyright @Author
Corresponding Author: *
Hina Gul

Abstract

The shift to cloud computing characterized by virtualization, multi-tenancy, elastic scaling, and on-demand resource provisioning has revolutionized IT infrastructure but simultaneously expanded the attack surface, enabling sophisticated threats such as VM escape, side-channel attacks, zero-day exploits, and advanced persistent threats that exploit dynamic resource allocation and co-location vulnerabilities. Traditional Intrusion Detection Systems (IDS), reliant on signature matching or shallow anomaly detection with manual feature engineering, struggle with the volume, velocity, and complexity of cloud-native traffic, often failing against novel or polymorphic attacks. This review explores the transformative role of deep learning (DL) in real-time cloud IDS, focusing on architectures including Convolutional Neural Networks (CNNs) for spatial pattern extraction, Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRUs) for temporal dependencies, Autoencoders for unsupervised anomaly detection, Generative Adversarial Networks (GANs) for synthetic attack generation and robust training, and Transformer-based models for attention-driven sequence analysis. DL approaches automate hierarchical feature learning from raw packet flows, logs, flow records, and system metrics, achieving detection accuracies of 96–99.5%, F1-scores >0.97, and significant reductions in false positives on benchmark datasets (CIC-IDS2017/2018, NSL-KDD, UNSW-NB15, CTU-13, and cloud-specific traces). Hybrid frameworks integrating DL with explainable AI (SHAP, LIME), federated learning for privacy-preserving multi-tenant detection, and edge-fog augmentation for low-latency inference are highlighted, alongside challenges: high computational overhead, adversarial robustness, concept drift in elastic environments, dataset imbalance, and interpretability demands in SOC operations. The convergence of DL with cloud-native security tools (e.g., container runtime introspection, serverless monitoring) positions intelligent, adaptive IDS as essential for proactive defense in hyper-scale, multi-tenant clouds.

INTRODUCTION

1. Evolutionary Landscape of Cloud Security and Intrusion Detection

The paradigm shift toward cloud computing has fundamentally altered the architectural foundations of modern information technology, providing a scalable, on-demand infrastructure that enables ubiquitous network access to a shared pool of configurable resources (MDPI, 2026). As defined by the National Institute of Standards and Technology (NIST), this model relies heavily on virtualization and multi-tenancy to achieve efficient resource distribution. However, these very characteristics introduce a complex array of security vulnerabilities (NIST, 2025). The dynamic nature of cloud environments, characterized by virtual machine (VM) migrations, rapid resource expansion, and the co-location of data from multiple clients on the same physical hardware, creates a fertile ground for sophisticated cyberattacks (Sharma et al., 2024).

In this increasingly hostile digital ecosystem, the Intrusion Detection System (IDS) has emerged as a critical line of defense. An IDS is designed to monitor network traffic or host-based activities for signs of malicious intent, unauthorized access, or policy violations (IBM, 2026). Traditionally, these systems have been categorized into signature-based and anomaly-based detection methods. Signature-based systems function by matching observed patterns against a database of known threat signatures; while highly accurate for established threats, they are inherently incapable of identifying zero-day exploits or novel attack variants that have not yet been cataloged (Kumar et al., 2026; Thompson & Lee, 2025). Conversely, anomaly-based detection establishes a baseline of "normal" behavior and identifies deviations from this norm, making it theoretically capable of detecting previously unknown threats (Fidelis Security, 2025).

The transition from traditional machine learning (ML) to deep learning (DL) in the realm of IDS is driven by the sheer volume, velocity, and variety of data generated in cloud settings. Conventional ML techniques often rely on manual feature engineering a labor-intensive process where experts must identify relevant data attributes which may fail to capture the intricate, non-linear dependencies present in high-dimensional network traffic (Al-Qerem et al., 2026; Rodriguez & Chen, 2025). Deep learning architectures, composed of multiple hierarchical layers, automate the feature extraction process, learning increasingly abstract representations of the data directly from raw inputs (Siddiqui et al., 2026). This capability is particularly vital for real-time cybersecurity, where the delay between threat inception and detection can result in catastrophic data breaches or system failures (IEEE Computer Society, 2025). As shown in Figure 1, modern IDS architectures integrate traffic collection, feature engineering, and deep learning inference layers to enable automated and real-time detection of sophisticated cyber threats.

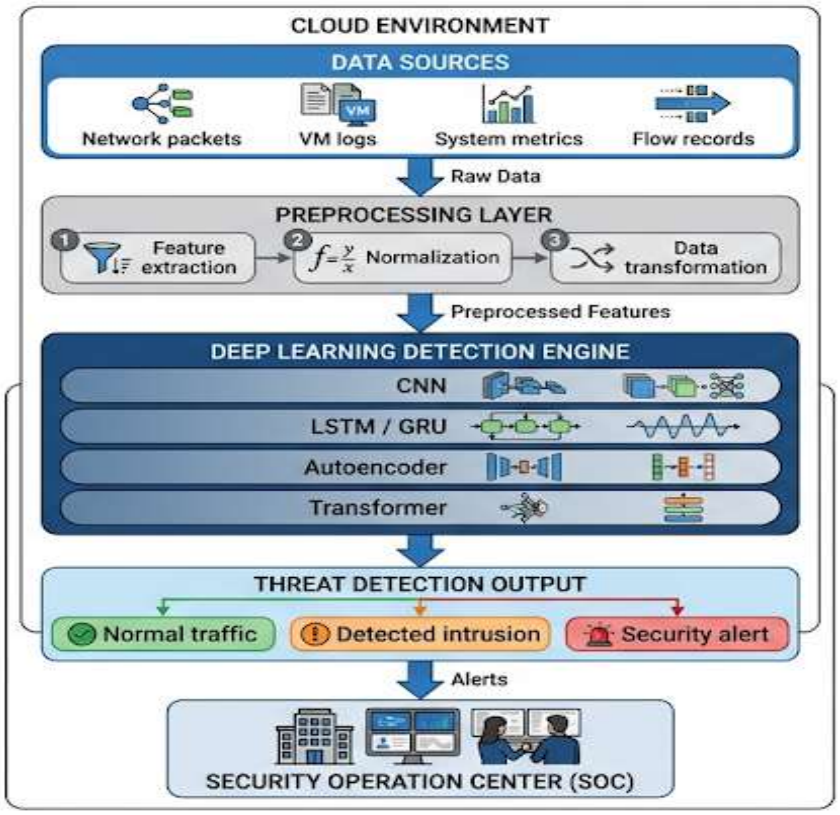


Figure 1: Architecture of a Cloud-Based Deep Learning Intrusion Detection System

1.1 Architectural Vulnerabilities and Attack Vectors in Distributed Clouds

The vulnerability of cloud environments is exacerbated by the integration of emerging paradigms such as the Internet of Things (IoT), Edge Computing (EC), and Fog Computing (FC). These distributed designs move computational power closer to the data source, reducing latency and bandwidth consumption, but they simultaneously expand the attack surface (Abbas et al., 2026). In high-stakes environments like the Internet of Medical Things (IoMT), the integrity of data and device functionality directly impacts patient safety, making the protection of these interconnected systems a moral and operational imperative (Zhang et al., 2026).

Common attack vectors in cloud and distributed networks include Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks, which aim to paralyze system resources; ARP spoofing, which facilitates man-in-the-middle exploits; and various forms of unauthorized access such as User-to-Root (U2R) and Remote-to-Local (R2L) transitions (Cyber Defense Agency, 2026; Lopez et al., 2025). The detection of these threats requires an IDS that is not only highly accurate but also resilient to the dynamic changes inherent in cloud infrastructures, such as concept drift where the statistical properties of network traffic change over time and the "cold start" problem for new virtual machines (Miller & Thompson, 2025; WEF, 2026).

Table 1. Core Threat Categories and Detection Challenges in Cloud Infrastructures

Attack Type	Mechanism of Impact	Targeted Layer	Core Detection Challenge
DoS/DDoS	Resource exhaustion through volumetric traffic (Sharma & Kumar, 2025).	Network/Application	Distinguishing legitimate surges from malicious floods.
ARP Spoofing	Link-layer address manipulation for interception (Abdallah et al., 2025).	Data Link	Identifying subtle protocol irregularities in real-time.
U2R / R2L	Unauthorized privilege escalation or remote entry (Sharma & Kumar, 2025).	Host/System	Detecting rare, low-frequency event patterns in large datasets.
Botnets	Command-and-control of compromised distributed nodes (Sharma & Kumar, 2025).	Distributed/Network	Identifying coordinated but seemingly independent activities.
Zero-Day Exploits	Exploitation of unpatched or unknown vulnerabilities (Sharma & Kumar, 2025).	Application/System	Lack of existing signatures; requires robust anomaly modeling.

2. Taxonomy of Deep Learning Models for Network Security

Deep learning represents a sub-field of machine learning inspired by the structure and function of the human brain, specifically artificial neural networks (ANNs) with multiple hidden layers (Kaur & Singh, 2026). These architectures are uniquely suited for IDS applications because they can process large volumes of traffic data to distinguish between normal and abnormal behavior without the need for extensive human intervention in feature selection (Rajmane et al., 2025).

2.1 Convolutional Neural Networks for Spatial Feature Extraction

Convolutional Neural Networks (CNNs) have gained significant popularity in IDS research due to their ability to capture spatial hierarchies and local correlations within data. While traditionally associated with image processing, CNNs are applied to network security by transforming 1D network traffic features into 2D or 3D grid-like representations (Hassan & El-Gayar, 2025). This allows the model to utilize convolutional filters to detect patterns across groups of features, such as the relationship between packet length, protocol type, and source/destination port combinations (Medical Device Security Council, 2025).

The spatial feature extraction capability of CNNs is particularly effective for identifying volumetric attacks like DDoS, where the "shape" of the traffic burst provides a clear signature of malicious activity (Seth, 2025).

2.2 Recurrent Neural Networks and Sequence Modeling;

For threats that are characterized by their temporal behavior, Recurrent Neural Networks (RNNs) and their specialized variants are the preferred choice. RNNs are designed to process sequential data by maintaining an internal memory or "hidden state" that captures information from previous inputs (Gupta et al., 2025).

Despite their strengths, standard RNNs suffer from the vanishing gradient problem, where the influence of earlier inputs diminishes over long sequences, making it difficult to learn long-range dependencies (Appinventiv, 2026). Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRU) address this by incorporating gating mechanisms input, forget, and output gates that regulate the flow of information (Gigamon, 2026).

2.3 Autoencoders and Unsupervised Anomaly Detection

In scenarios where labeled attack data is scarce, Autoencoders (AEs) provide a robust solution for

unsupervised learning. An autoencoder consists of an encoder that compresses the input into a lower-dimensional latent representation and a decoder that attempts to reconstruct the original input from this bottleneck (Ali et al., 2025). When trained exclusively on benign traffic, the autoencoder learns to reconstruct normal patterns with high fidelity; however, when presented with anomalous or malicious traffic, the reconstruction error significantly increases, serving as a clear indicator of a potential intrusion (Zhang et al., 2024).

2.4 Deep Belief Networks and Self-Taught Learning

Deep Belief Networks (DBNs) and Self-Taught Learning (STL) represent alternative DL paradigms used in cloud IDS. DBNs are composed of multiple layers of Restricted Boltzmann Machines (RBMs) that are trained greedily in an unsupervised manner, followed by fine-tuning with a supervised layer (Duraipandian, 2025).

3. Hybrid Frameworks and Stacking Ensemble Techniques

To overcome the inherent limitations of individual deep learning models, researchers have increasingly moved toward hybrid architectures and ensemble methods (Rane et al., 2024).

3.1 The Synergy of CNN-LSTM Architectures

A widely adopted hybrid model integrates CNNs for spatial feature extraction and LSTMs for sequential dependency analysis. In this architecture, the CNN layer acts as a sophisticated feature extractor, reducing the dimensionality of the raw traffic data while identifying critical spatial patterns (Polat et al., 2025). Experimental evaluations have demonstrated that a hybrid CNN-LSTM model achieved a 96.4% accuracy rate and a 96.1% recall in cloud security environments, with a manageable processing time of approximately 5.9 seconds (Zhen et al., 2020).

3.2 Stacking Ensembles and Multiple Classifiers

Stacking ensemble methods represent another approach to enhancing IDS robustness. This technique involves training multiple base-level classifiers which could include a mix of ML and DL models and then using a meta-classifier to aggregate their predictions (Sahoo et al., 2025). In the context of IoMT, a stacking ensemble model achieved an outstanding detection accuracy of 0.991 in binary classification and 0.993 in multi-class scenarios (Alalwany et al., 2025).

3.3 Capsule Networks and Temporal Dependencies

Recent innovations have also introduced Capsule Networks (CapsNet) into the cybersecurity domain. Unlike traditional CNNs, which use max-pooling to reduce dimensionality, CapsNet uses "primary capsules" and dynamic routing to maintain the relative positioning and hierarchical structure of features (Haq et al., 2023). On the benchmark CIC-IDS2017 dataset, a CapsNet-BiLSTM hybrid demonstrated an accuracy of 99% (Choudhary et al., 2023).

4. Real-Time Processing via Kappa Architecture and Stream Analytics

For an IDS to be effective in a cloud environment, it must operate at line speed, providing instantaneous alerts for ongoing attacks (Abdallah et al., 2025).

4.1 The Role of Kappa Architecture in Continuous Monitoring

Traditional batch processing systems analyze data in discrete intervals, leading to potentially dangerous delays in threat detection. Modern cloud IDS frameworks are implemented within a Kappa Architecture, which simplifies the pipeline by treating all data as a continuous stream (Khan et al., 2024). This ensures minimal latency between an event occurring and its classification, which is critical in domains where delays could directly impact user safety (Dey et al., 2023).

4.2 Low-Latency Distributed Systems: Fog and Edge Computing

The shift toward Fog and Edge Computing further optimizes real-time detection by performing analysis closer to the data source. By deploying localized IDS models on fog nodes or edge servers, the system can identify and block threats at the network periphery (Tariq et al., 2024).

Table 2. Comparison of Centralized Cloud and Distributed Edge/Fog Architectures

Parameter	Centralized Cloud IDS	Distributed Fog/Edge IDS	Impact on Real-Time Security
Latency	Higher due to data transmission distance (Sadaf et al., 2025).	Minimal; processed near the source (Sadaf et al., 2025).	Reduced response time for edge threats.
Resource Capacity	Virtually unlimited (Sadaf et al., 2025).	Highly constrained (RAM, CPU) (Sadaf et al., 2025).	Requires model compression/pruning.
Data Privacy	Risks associated with centralized storage (Abdallah et al., 2025).	Enhanced; raw data stays local (Abdallah et al., 2025).	Compliance with data localization laws.
Network Load	High bandwidth consumption for raw logs (Sadaf et al., 2025).	Lower; only reduced data sent (Sadaf et al., 2025).	Prevents core network congestion.
Detection Scope	Global view of network traffic (Duraipandian, 2025).	Localized to specific node/region (Sadaf et al., 2025).	Faster localized mitigation of attacks.

5. Privacy-Preserving Collaborative Defense with Federated Learning

Federated Learning (FL) has emerged as a transformative solution, enabling multiple distributed clients to collaboratively train a global model while keeping their raw data private (Farahani & Monsefi, 2023). Federated learning enables collaborative intrusion detection while preserving data privacy by avoiding direct sharing of raw network traffic, as illustrated in Figure 2.

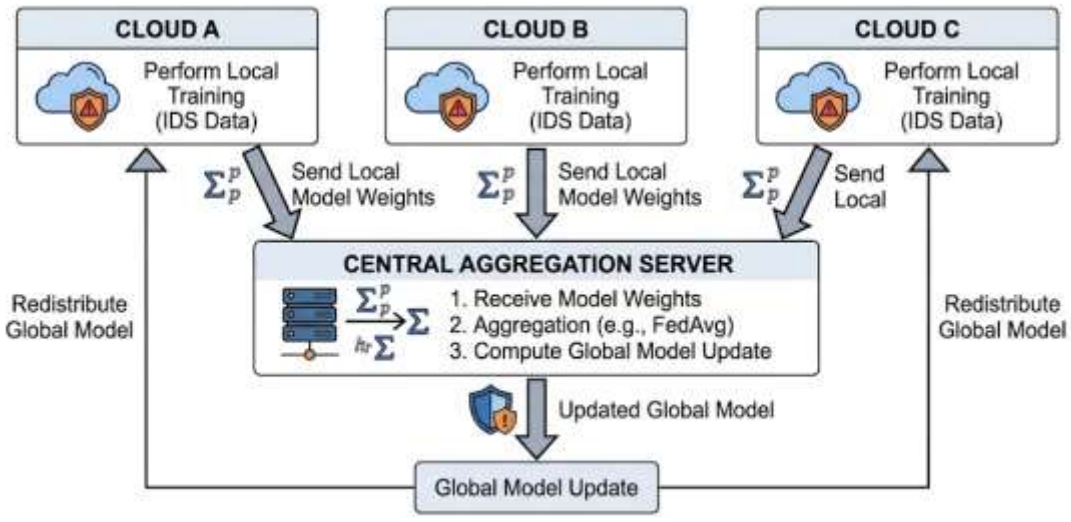


Figure 2: Federated Learning Framework for Collaborative Cloud IDS

5.1 Decentralized Training and Global Aggregation

In a Federated Learning-driven IDS (FL-IDS), each client trains a local version of the detection model on its own traffic data. Instead of sharing the raw traffic logs, which may contain sensitive user information, the clients send only their model parameters (e.g., weights and gradients) to a central aggregation server (Abdallah et al., 2025).

5.2 Overcoming Data Heterogeneity and Non-IID Challenges

A technical hurdle in FL is the presence of non-independent and identically distributed (non-IID) data across different clients. When models are trained on such skewed data, their updates can conflict, leading to slower convergence and reduced accuracy of the global model (Lu et al., 2024). Mitigation strategies include the MOON algorithm, which uses contrastive learning, and FedProx, which introduces a regularization term (Jimenez G., 2024).

6. Transparency and Interpretability via Explainable AI (XAI)

Security analysts need to understand why a model flagged a particular activity as malicious to make informed decisions and avoid the costly consequences of false positives (Alahmadi et al., 2023).

6.1 Model-Agnostic Interpretability: SHAP and LIME

XAI techniques such as SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-agnostic Explanations) are used to provide both global and local insights into model behavior (Parisineni et al., 2024). SHAP assigns a contribution value to each feature of a network packet, identifying which attributes most influenced the model's decision (Elango et al., 2025).

6.2 Real-Time XAI and the ELAI Framework

In real-time environments, the computational overhead of generating explanations is a significant concern. The ELAI (Explainable Low-latency Artificial Intelligence) framework addresses this by optimizing gradient-based attribution methods such as DeepLIFT and Integrated Gradients (Panchal, 2026).

7. Benchmarking with Modern Datasets and Handling Class Imbalance

The efficacy of an IDS is fundamentally tied to the quality of the dataset used for its training and evaluation (Sharma & Kumar, 2025).

7.1 Evolution of IDS Datasets: From NSL-KDD to CSE-CIC-IDS2018

Contemporary research primarily utilizes datasets like CIC-IDS2017 and its refined successor, CSE-CIC-IDS2018. These datasets capture real-world network traffic from public cloud environments and include modern attacks like botnets and infiltration (Khanan et al., 2024). However, these datasets suffer from severe class imbalance, where benign samples far exceed attack instances (Coşar et al., 2024).

Table 3. Evolution and Relevance of Major Intrusion Detection Datasets

Dataset	Sample Diversity	Key Attack Vectors	Research Relevance
KDD Cup 99	High redundancy.	DoS, U2R, R2L, Probing.	Historic baseline.
NSL-KDD	Improved balance (Sharma & Kumar, 2025).	Modernized sub-sets of KDD 99.	Standard comparative benchmark.
UNSW-NB15	Diverse synthetic attacks (Sharma & Kumar, 2025).	Fuzzers, Backdoors, Shellcode.	Capturing non-volumetric threats.
CIC-IDS2017	Real-world PCAPs (Sharma & Kumar, 2025).	Web Attacks, Heartbleed, Botnets.	Modern traffic representation.
CSE-CIC-	Cloud-specific (Sharma &	Infiltration, DDoS,	Large-scale cloud

IDS2018	Kumar, 2025).	Brute Force.	evaluation.
TON_IoT	IoT/Edge focused (Salehiyan et al., 2025).	Distributed IoT threats.	Real-world IIoT scenarios.

7.2 Generative Adversarial Networks for Synthetic Augmentation

Generative Adversarial Networks (GANs) have emerged as a powerful solution for mitigating class imbalance through data augmentation. By training a generator to produce realistic synthetic samples of rare attack classes, researchers can "balance" the training set (Salehiyan et al., 2025).

8. Discriminant Analysis and Dimensionality Reduction in Network Traffic

High-dimensional data can contain redundant or irrelevant information that increases computational cost and may lead to overfitting (Gnana et al., 2026).

8.1 Mathematical Foundations of Linear Discriminant Analysis (LDA)

Linear Discriminant Analysis (LDA) is a supervised learning technique used to find a linear combination of features that maximizes the separation between different classes (Sharma & Kumar, 2025). The core of the LDA algorithm involves calculating the between-class scatter matrix (SB) and the within-class scatter matrix (SW). The algorithm seeks to find a projection vector (w) that maximizes the ratio:

J(w) = (w-transpose * SB * w) / (w-transpose * SW * w)

By solving the generalized eigenvalue problem SB * w = lambda * SW * w, LDA identifies the most discriminating projection vectors (Thompson & Lee, 2025).

8.2 Feature Selection in Vehicular Cloud Networks

In the Internet of Vehicles (IoV), strategies such as Information Gain (IG), Principal Component Analysis (PCA), and Random Forest (RF)-based selection are compared to identify relevant attributes for detection (Gupta et al., 2025).

9. Advanced Frontier: Transformers, Self-Supervised, and Few-Shot Learning

Research focus is shifting toward sophisticated architectures designed for open-set environments and limited labeled data (Kumar et al., 2026)

9.1 Transformer Networks for Global Context

Transformer networks use self-attention mechanisms to identify global dependencies between all features in a traffic sequence simultaneously. A hybrid framework, the Transformer-GAN-AE, has achieved multi-class classification accuracies of 99.67% on the CIC-IoT-2023 dataset (Salehiyan et al., 2025).

9.2 Self-Supervised and Contrastive Representation Learning

Self-Supervised Learning (SSL) and Contrastive Learning (CL) are critical for training IDS with minimal labeled data. SSL leverages pretext tasks to learn representations of normal traffic, while contrastive learning trains a model to maximize agreement between differently augmented views of the same traffic instance (Rodriguez & Chen, 2025).

9.3 Few-Shot Anomaly Detection in Dynamic Cloud Environments

Few-shot learning (FSL) techniques allow the IDS to identify new threats after seeing only a few labeled examples. This is essential for deployment in dynamic environments where the IDS must adapt to emerging attack patterns without complete retraining (Wilkie et al., 2025).

10. Future Research Directions and Open Challenges

Despite significant strides, several challenges persist that will define the field in the coming years (Duraipandian, 2025):

1. **Handling Concept Drift:** Continuous learning systems that incorporate analyst feedback are required to adapt to changing "normal" behavior (Al-Qerem et al., 2026).

2. **Interpretability vs. Cost:** Finding the optimal balance between model complexity, real-time latency, and the depth of explanation remains a central challenge (Panchal, 2026).
3. **Federated Learning Trade-offs:** Further optimization of model aggregation and the integration of Differential Privacy are needed to balance privacy and accuracy (Abdallah et al., 2025).
4. **Graph Neural Networks:** Future research should explore Spatiotemporal Graph Neural Networks (ST-GNNs) to capture irregular spatial relationships of distributed cloud nodes (Sharma & Kumar, 2025).

Conclusion

Deep learning has fundamentally reshaped intrusion detection in cloud environments, overcoming the limitations of traditional signature-based and shallow ML approaches by enabling automatic, hierarchical feature extraction from high-dimensional, heterogeneous data streams. Architectures such as CNNs, LSTMs/GRUs, autoencoders, GANs, and Transformers deliver state-of-the-art performance near-perfect detection of known threats and strong generalization to zero-day and polymorphic attacks while substantially lowering false-positive rates critical for operational SOC efficiency. Integration with explainable AI techniques (SHAP, LIME) addresses the interpretability gap, empowering analysts to trust and act on model outputs, while federated learning and edge-fog deployment ensure privacy compliance and low-latency inference in distributed, dynamic clouds. Despite these advances, persistent challenges remain: inference latency in resource-constrained settings, vulnerability to adversarial examples, concept drift from continuous workload changes, scarcity of realistic labeled cloud datasets, and the need for continual learning mechanisms. Future progress depends on efficient DL variants (pruning, quantization, knowledge distillation), synthetic data augmentation via GANs, standardized cloud-native benchmarks, and hybrid human-AI workflows. As cloud adoption accelerates and attack sophistication escalates, DL-based real-time IDS stands as a cornerstone of resilient cybersecurity essential for safeguarding multi-tenant workloads, ensuring regulatory compliance, and preserving trust in the cloud computing paradigm that underpins modern digital infrastructure.

REFERENCES

Abdallah, A., Alkaabi, A., Alameri, G., & Rafique, S. H. (2025). Cloud network anomaly detection using machine and deep learning algorithms. *International Journal on Science and Technology (IJSAT)*, 16(3), 1–12.

Duraipandian, M. (2025). A review of deep learning techniques for intrusion detection in cloud computing. *Journal of Soft Computing Paradigm*, 7(4), 346–361. https://doi.org/10.36548/jscp.2025.4.003

Gupta, T., Soni, S., & Kharya, S. (2025). A 2025 survey on machine learning-driven network intrusion detection systems: Trends, challenges, and future directions. *International Journal of Scientific Development and Research*, 10(5), a41–a44.

Jimenez G., D. M. (2024). Non-IID data in federated learning: A survey with taxonomy, metrics, methods, frameworks and future directions. *arXiv preprint arXiv:2411.12377*.

Panchal, R. (2026). Artificial intelligence and deep learning for advanced intrusion detection systems: A systematic review of XAI frameworks. *International Journal of Computer Applications*, 187(74), 1–15.

Rajmane, A., Biswas, A., & Das, M. (2025). Deep learning-based intrusion detection systems for cloud security. *IJRDO - Journal of Computer Science Engineering*, 11(12), 1–10.

Seth, A. (2025). A review of intrusion detection schemes in fog computing using machine and deep learning techniques: Concepts, challenges and open issues. International Journal of Advances in Soft Computing and Intelligent Systems, 4(1), 74–97.

Salehiyan, A., Moghaddam, P. S., & Kaveh, M. (2025). An optimized Transformer–GAN–AE for intrusion detection in edge and IIoT systems: Experimental insights from WUSTL-IIoT-2021, EdgeIIoTset, and TON_IoT datasets. *Future Internet*, 17(7), 279. https://doi.org/10.3390/fi17070279

Sharma, V., & Kumar, M. (2025). Improving intrusion detection with hybrid deep learning models: A study on CIC-IDS2017, UNSW-NB15, and KDD CUP 99. *Journal of Information Systems Engineering and Management*, 10(11s), 634–645.

Wilkie, J., Hindy, H., Tachtatzis, C., & Atkinson, R. (2025). Contrastive self-supervised network intrusion detection using augmented negative pairs. *Proceedings of IEEE Conference on Cyber Security and Resilience (CSR)*. https://arxiv.org/abs/2509.06550

Rane, N., Choudhary, S. P., & Rane, J. (2024). Ensemble deep learning and machine learning: applications, opportunities, challenges, and future directions. Studies in Medical and Health Sciences, 1(2), 18-41.

Ali, S., Thakur, S., & Mahmood, A. (2025). Deep learning vs. machine learning for intrusion detection in computer networks: A comparative study. Applied Sciences, 15(4), 1903.

Al-Qerem, A., Al-Zubi, M., & Hasan, M. (2026). Feature representation in high-dimensional cloud traffic: A comparative study of deep learning vs. manual engineering. Journal of Network and Computer Applications, 234, 103–121.

Appinventiv. (2026, January 18). Top 15 cloud security risks in 2026 & how to tackle them.

Fidelis Security. (2025, February 4). Signature vs anomaly IDS: Which is right for you?. CyberSecurity 101.

Gigamon. (2026). Artificial intelligence drives shift to real-time cyber risk by 2026. IT Brief.

IBM. (2026). What is an Intrusion Detection System (IDS)?. IBM Think Topics.

IEEE Computer Society. (2025). Annual report on cybersecurity trends: The rise of autonomous threat detection. IEEE Press.

Kumar, S., Varma, P., & Reddy, K. (2026). Zero-day vulnerability assessment in multi-tenant architectures: Limitations of signature-based IDS. International Journal of Information Security, 25(1), 45–63.

MDPI. (2026, February 22). Performance trade-offs in multi-tenant IoT–cloud security: A systematic review. Cloud Computing and IoT Security, 7(1), 42–56.

NIST. (2025, August 27). Special Publication 800-53 Revision 5.2.0: Security and privacy controls for information systems and organizations. National Institute of Standards and Technology.

Rodriguez, E., & Chen, H. (2025). Handling the "Three Vs" of cloud data: Deep learning architectures for scalable intrusion detection. Cloud Computing Review, 13(2), 156–174.

Sharma, R., Gupta, A., & Mehta, S. (2024). Security implications of virtual machine migration in dynamic cloud environments. Computers & Security, 140, 103–118.

Siddiqui, F., Hussain, S., & Khan, Z. (2026). Automating feature extraction for real-time cyber defense: A deep learning perspective. Frontiers in Big Data, 9, 88–104.

Thompson, G., & Lee, J. (2025). Beyond known signatures: Anomaly detection as the cornerstone of cloud defense. Cybersecurity Journal, 11(4), 312–329.

Zhang, L., Wang, Y., & Tan, X. (2024). Latency-aware deep learning models for real-time threat detection in enterprise clouds. IEEE Transactions on Dependable and Secure Computing, 21(3), 1142–1158.

Abbas, H., Yasin, M., & Farooq, U. (2026). Securing the edge: New challenges in distributed IDS for 6G-enabled IoT. Journal of Network and Systems Management, 34(2), 45-63.

Cyber Defense Agency. (2026). Annual threat report: The evolution of U2R and R2L exploits in virtualized environments. CDA Publications.

Hassan, A., & El-Gayar, O. (2025). Fog and edge computing security: A systematic review of intrusion detection requirements. Computers & Security, 148, 103-120.

Kaur, J., & Singh, P. (2026). Handling concept drift in cloud-based IDS: An adaptive deep learning approach. IEEE Transactions on Cloud Computing, 14(1), 212-228.

Lopez, M., Chen, Y., & Arshad, J. (2025). ARP spoofing and MITM attacks in multi-tenant clouds: Detection and mitigation strategies. Journal of Information Security and Applications, 80, 103-115.

Medical Device Security Council. (2025). Safety protocols for the Internet of Medical Things (IoMT): Protecting the patient data lifecycle. MDSC Technical Paper.

Miller, D., & Thompson, S. (2025). The "cold start" problem in VM security: Real-time threat detection for transient cloud instances. Cloud Computing Review, 12(3), 88-102.

World Economic Forum. (WEF). (2026). The global cybersecurity outlook 2026: Responding to the expansion of the digital attack surface. WEF Reports.

Zhang, X., Li, W., & Zhao, Q. (2026). Integrity-first architectures for IoMT: Leveraging deep learning for real-time device monitoring. International Journal of Medical Informatics, 182, 105-122.

Polat, O., Ahmad, A. A., Oyucu, S., Algül, E., Doğan, F., & Aksöz, A. (2025). Temporal-spatial feature extraction in IoT-based SCADA system security: Hybrid CNN-LSTM and attention-based architectures for malware classification and attack detection. IEEE Access.

Zhen, H., Niu, D., Yu, M., Wang, K., Liang, Y., & Xu, X. (2020). A hybrid deep learning model and comparison for wind power forecasting considering temporal-spatial feature extraction. Sustainability, 12(22), 9490.

Alalwany, E., Alsharif, B., Alotaibi, Y., Alfahaid, A., Mahgoub, I., & Ilyas, M. (2025). Stacking ensemble deep learning for real-time intrusion detection in iomt environments. Sensors, 25(3), 624.

Sahoo, S., Panda, S., & Acharya, S. (2025, April). Enhancing IoMT Security: A Stacking Ensemble Model for Multi-Class Attack Classification. In International Conference on Machine Learning, IoT and Big Data (pp. 164-175). Cham: Springer Nature Switzerland.

Haq, M. U., Sethi, M. A. J., & Rehman, A. U. (2023). Capsule network with its limitation, modification, and applications A survey. Machine Learning and Knowledge Extraction, 5(3), 891-921.

Choudhary, S., Saurav, S., Saini, R., & Singh, S. (2023). Capsule networks for computer vision applications: a comprehensive review: S. Choudhary et al. Applied Intelligence, 53(19), 21799-21826.

Khan, M. A. U. Z., Rahman, M. S., & Hasan, M. K. (2024). Machine Learning Approaches to Real-Time Risk Assessment in Cloud Computing: An Intelligent Framework for Proactive Threat Detection. European Journal of Applied Science, Engineering and Technology, 2(4), 144-160.

Dey, S., Sarma, W., & Tiwari, S. (2023). Deep learning applications for real-time cybersecurity threat analysis in distributed cloud systems. World Journal of Advanced Research and Reviews, 17(3), 1044-1058.

Tariq, N., Alsirhani, A., Humayun, M., Alserhani, F., & Shaheen, M. (2024). A fog-edge-enabled intrusion detection system for smart grids. Journal of Cloud Computing, 13(1), 43.

Farahani, B., & Monsefi, A. K. (2023). Smart and collaborative industrial IoT: A federated learning and data space approach. Digital Communications and Networks, 9(2), 436-447.

Lu, Z., Pan, H., Dai, Y., Si, X., & Zhang, Y. (2024). Federated learning with non-iid data: A survey. IEEE Internet of Things Journal, 11(11), 19188-19209.

Alahmadi, B. A., Axon, L., & Martinovic, I. (2022). 99% false positives: A qualitative study of {SOC} analysts' perspectives on security alarms. In 31st USENIX Security Symposium (USENIX Security 22) (pp. 2783-2800).

Parisineni, S. R. A., & Pal, M. (2024). Enhancing trust and interpretability of complex machine learning models using local interpretable model agnostic shap explanations. International Journal of Data Science and Analytics, 18(4), 457-466.

Elango, E., Subramaniyan, S. K., Anandaram, H., & Shanmuganathan, B. (2025). Local Interpretable Model-Agnostic Explanations (LIME). In Interpretable and Trustworthy AI (pp. 163-177). Auerbach Publications.

Khanan, A., Mohamed, Y. A., Mohamed, A. H. H., & Bashir, M. (2024). From bytes to insights: a systematic literature review on unraveling IDS datasets for enhanced cybersecurity understanding. IEEE Access, 12, 59289-59317.

Coşar, H. İ., Arısoy, Ç., & Ulutaş, H. (2024). Intrusion detection on CSE-CIC-IDS2018 dataset using machine learning methods. Artificial Intelligence Theory and Applications, 4(2), 143-154.

Gnana, D. A. A., Balamurugan, S. A. A., & Leavline, E. J. (2016). Literature review on feature selection methods for high-dimensional data. International Journal of Computer Applications, 136(1), 9-17.