International Journal of Scientific Research and Engineering Development-– Volume 9 Issue 2, Mar-Apr 2026
Available at www.ijsred.com

RESEARCH ARTICLE OPEN ACCESS

Anomaly Detection in Cloud Networks Using Machine Learning Techniques – A Survey

Dr P.K. Sharma[1], Mr. Manvendra Singh Divakar[2], Shaheen Bano[3]
[1]Principal, [2]Assistant Professor
[3]Research Scholar
[1,2,3]NRI Institute of Research & Technology, Bhopal (M.P.)

Abstract
The rapid expansion of cloud computing has significantly transformed modern information technology infrastructures by enabling scalable, flexible, and cost-efficient service delivery. However, the dynamic, distributed, and multi-tenant nature of cloud environments has also intensified network security challenges, particularly in detecting anomalous activities that may indicate cyberattacks, system misconfigurations, or operational failures. Conventional rule-based and signature-driven security mechanisms are increasingly inadequate in such environments due to their limited adaptability and inability to identify previously unseen or evolving threats. The reviewed dissertation addresses these limitations by proposing a machine learning–based anomaly detection framework tailored for cloud networks. The approach adopts an unsupervised learning paradigm using an autoencoder-based neural network architecture to model normal cloud network behavior and detect anomalies through reconstruction error analysis. The model is trained exclusively on normal traffic patterns and evaluated using standard performance metrics. Experimental results demonstrate an overall classification accuracy of 90.97%, with high precision and recall for normal traffic and reliable detection performance for anomalous traffic despite significant class imbalance. Confusion matrix analysis reveals balanced classification behavior with acceptable false positive and false negative rates, while training and validation loss curves confirm stable convergence and strong generalization without overfitting. This review synthesizes the dissertation's objectives, methodology, experimental findings, and contributions, emphasizing its relevance in advancing scalable, adaptive, and data-driven anomaly detection solutions for real-world cloud network security.

Keywords: Cloud Computing; Cloud Network Security; Anomaly Detection; Machine Learning; Autoencoder; Unsupervised Learning; Network Traffic Analysis; Intrusion Detection Systems.

1. Introduction
Cloud computing has emerged as a cornerstone of modern information technology infrastructures, fundamentally reshaping how computing resources are provisioned, managed, and consumed. By enabling on-demand access to scalable computing power, elastic storage, and flexible service deployment models, cloud platforms have become indispensable across diverse domains such as finance, healthcare, e-commerce, education, and large-scale enterprise systems. Organizations increasingly rely on cloud environments to host mission-critical applications and manage sensitive data, driven by benefits including reduced capital expenditure, operational efficiency, and rapid scalability. Despite these advantages, the widespread adoption of cloud computing has also introduced a new and complex landscape of security challenges that demand innovative and adaptive protection mechanisms. The intrinsic characteristics of cloud environments—virtualization, elasticity, and multi-tenancy—significantly expand the attack surface compared to traditional on-premise networks. Virtualization enables multiple tenants to share underlying physical infrastructure, which, while improving resource utilization, also increases the risk of cross-tenant attacks and lateral movement within cloud networks. Elastic resource allocation and auto-scaling mechanisms dynamically adjust workloads in response to demand, resulting in continuously changing traffic patterns that complicate baseline security monitoring. Furthermore, cloud networks are inherently distributed and highly interconnected, often spanning multiple geographic locations and relying on internet-facing interfaces, making them attractive targets for a wide range of cyber threats.

Cloud network traffic is characterized by high volume, velocity, and heterogeneity. It evolves rapidly due to fluctuating workloads, dynamic user behavior, application updates, and service orchestration processes. These characteristics render traditional security mechanisms increasingly ineffective. Conventional rule-based intrusion detection systems and signature-driven firewalls rely on predefined patterns and static thresholds, which are ill-suited for detecting novel, stealthy, or evolving attacks. Such systems struggle to adapt to changes in normal traffic behavior and often fail to identify zero-day attacks or subtle anomalies that do not match known signatures. As a result, malicious activities may remain undetected until they cause significant operational or financial damage. Within this context, anomaly detection has gained prominence as a promising paradigm for cloud network security. Unlike signature-based approaches, anomaly detection focuses on identifying deviations from learned normal behavior, enabling the detection of previously unseen or unknown threats. This behavior-centric perspective is particularly relevant for cloud environments, where attack patterns frequently evolve and may closely resemble legitimate traffic. However, effective anomaly detection in cloud networks presents its own challenges, including high-dimensional data, class imbalance between normal and anomalous traffic, and the need to minimize false alarms that could overwhelm security teams. Machine learning techniques offer a powerful foundation for addressing these challenges by enabling automated, data-driven analysis of complex network traffic patterns. In particular, unsupervised and semi-supervised learning approaches are well-suited for cloud environments, where comprehensive labeled datasets containing all possible anomaly types are rarely available. By learning representations of normal traffic behavior directly from data, machine learning models can adapt to evolving conditions and identify deviations indicative of anomalous activity. Deep learning methods further enhance this capability by automatically extracting hierarchical and nonlinear feature representations from high-dimensional network data.

The reviewed dissertation addresses the aforementioned challenges by proposing a machine learning–based anomaly detection framework specifically designed for cloud networks. The study emphasizes an unsupervised learning paradigm using an autoencoder-based neural network architecture to model normal cloud network behavior and detect anomalies through reconstruction error analysis. This approach aligns with the practical constraints of real-world cloud environments and prioritizes adaptability, scalability, and robustness. By systematically evaluating the proposed framework using standard performance metrics and training behavior analysis, the dissertation provides empirical evidence of its effectiveness and generalization capability. This review situates the dissertation within the broader landscape of cloud security research, highlighting its methodological rigor and empirical contributions. By synthesizing the motivation, approach, and outcomes of the study, this introduction establishes the foundation for a critical examination of how machine learning–based anomaly detection can advance scalable and adaptive security solutions for modern cloud network environments.

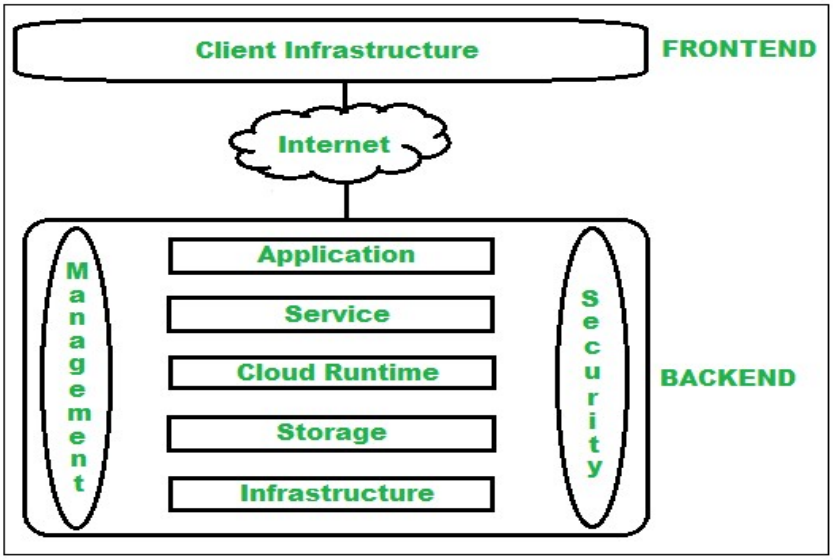

Figure 1.1: Conceptual illustration of cloud computing architecture and associated network security challenges.

2. Motivation and Research Objectives

The motivation for the reviewed dissertation arises from the growing recognition that traditional intrusion detection and network security mechanisms are increasingly ineffective within modern cloud computing environments. Conventional security systems, particularly rule-based and signature-driven intrusion detection systems, were originally designed for relatively static and well-defined network infrastructures. In contrast, cloud environments are inherently dynamic, distributed, and multi-tenant, with network behavior continuously evolving due to virtualization, elastic resource provisioning, and fluctuating user workloads. As a result, static security rules and predefined attack signatures are often unable to cope with the scale, variability, and complexity of cloud network traffic. One of the most critical limitations of traditional intrusion detection systems is their inability to detect zero-day attacks and previously unseen threat patterns. Modern cyberattacks are frequently designed to evade detection by closely mimicking legitimate traffic behavior, operating at low rates, or exploiting subtle vulnerabilities over extended periods. Signature-based systems depend on prior knowledge of attack patterns, rendering them ineffective against such adaptive and stealthy threats. In cloud environments, where services are exposed to the internet and shared across multiple tenants, the consequences of undetected attacks can be severe, including data breaches, service disruption, and loss of user trust.

Another motivating factor is the exponential growth of cloud network data. Cloud data centers generate massive volumes of network traffic logs, flow records, and performance metrics on a continuous basis. The scale and velocity of this data make manual inspection infeasible and overwhelm traditional statistical analysis techniques, which struggle to process high-dimensional data efficiently. Furthermore, cloud traffic data is often noisy, heterogeneous, and highly imbalanced, with anomalous events representing only a small fraction of overall activity. These characteristics further complicate detection and demand intelligent, automated analysis techniques capable of learning complex behavioral patterns directly from data. Machine learning has emerged as a promising solution to these challenges due to its ability to model complex, nonlinear relationships and adapt to evolving data distributions. In particular, unsupervised learning approaches are well-suited for cloud environments, where labeled datasets containing comprehensive anomaly annotations are rarely available. By learning normal network behavior and identifying deviations, such approaches enable proactive detection of both known and unknown threats. This motivation underpins the focus of the reviewed dissertation on unsupervised machine learning for cloud network anomaly detection.

Guided by these considerations, the dissertation defines a set of clear and structured research objectives. The first objective is to analyze cloud network security challenges and the characteristics of anomalous behavior, establishing a conceptual foundation for effective detection. The second objective focuses on designing a machine learning–based anomaly detection framework that is adaptive and scalable. The third objective emphasizes rigorous performance evaluation using standard metrics such as accuracy, precision, recall, and F1-score to ensure balanced assessment. The fourth objective involves analyzing training behavior and classification outcomes to validate learning stability and generalization capability. Finally, the fifth objective assesses the suitability of the proposed framework for real-world cloud environments. Collectively, these objectives guide a systematic investigation into the feasibility and effectiveness of unsupervised machine learning as a scalable security solution for modern cloud networks.

3. Methodological Framework

The reviewed dissertation adopts a quantitative, experimental, and data-driven methodological framework to systematically address the problem of anomaly detection in cloud networks using machine learning techniques. This research design is particularly appropriate given the need for objective evaluation, reproducibility, and empirical validation in security-focused studies. The methodology is structured to ensure that each stage of the detection process—from data preparation to performance evaluation—is rigorously defined and aligned with the unique characteristics of cloud environments. A key methodological choice in the study is the emphasis on unsupervised learning. In real-world cloud networks, obtaining comprehensive and accurately labeled datasets that capture all possible anomalous behaviors is extremely challenging due to the rarity, diversity, and evolving nature of security incidents. Supervised approaches, while effective in controlled settings, are often impractical in such contexts. By focusing on unsupervised learning, the proposed framework learns patterns of normal cloud network behavior directly from data and identifies deviations without relying on predefined attack labels. This design enhances adaptability and enables the detection of previously unseen or zero-day threats.

The overall methodology is organized into a sequence of interrelated stages. The first stage involves data preprocessing, which is essential for ensuring data quality and model stability. The dataset used in the study consists of numerical cloud network traffic features representing flow behavior, packet-level statistics, and protocol-specific characteristics. During preprocessing, missing and inconsistent values are identified and appropriately handled to prevent bias and instability during training. Feature normalization is applied to scale all attributes to a comparable range, ensuring that no single feature disproportionately influences the learning process. Additionally, outliers are carefully examined to distinguish between noise and meaningful rare patterns, preserving legitimate behavioral variations that are critical for anomaly detection. Following preprocessing, the methodology focuses on feature representation and learning. Rather than relying solely on manual feature engineering, the framework leverages representation learning to automatically capture essential characteristics of normal traffic behavior. This approach is particularly advantageous in cloud environments, where traffic patterns are complex and high-dimensional. The representation learning stage reduces redundancy and emphasizes informative patterns that facilitate effective anomaly discrimination.

At the core of the proposed framework is an autoencoder-based neural network architecture, selected for its effectiveness in unsupervised anomaly detection. The encoder component compresses high-dimensional input data into a compact latent representation, forcing the model to learn the most salient features of normal cloud network behavior. The decoder then reconstructs the original input from this latent space. During training, the model minimizes reconstruction error for normal traffic, thereby learning an accurate representation of legitimate behavior. Anomaly detection is achieved through reconstruction error analysis. When anomalous traffic is introduced, reconstruction error increases due to deviation from learned normal patterns. This error serves as an anomaly score, enabling the classification of traffic as normal or anomalous based on a defined threshold. Importantly, the architectural design balances expressive learning capability with computational efficiency, ensuring scalability and practical applicability in large-scale cloud

environments. Finally, the methodological framework includes systematic performance evaluation, employing standard metrics such as accuracy, precision, recall, F1-score, confusion matrix analysis, and training–validation loss behavior. This comprehensive evaluation strategy ensures that detection effectiveness, learning stability, and generalization capability are thoroughly assessed, reinforcing the reliability and robustness of the proposed anomaly detection framework.

4. Performance Evaluation and Results

The performance of the proposed anomaly detection model is evaluated through a comprehensive experimental analysis designed to assess its effectiveness, robustness, and practical suitability for cloud network security applications. The evaluation is conducted on a test dataset comprising 4,000 cloud network traffic instances, including both normal and anomalous samples. This dataset configuration reflects realistic cloud environments, where legitimate traffic constitutes the majority of network activity and anomalous events occur less frequently. To address this inherent class imbalance and provide a balanced assessment, multiple evaluation metrics are employed rather than relying solely on overall accuracy. For normal traffic (Class 0), the model demonstrates exceptionally strong classification performance. It achieves a precision of 0.9413, recall of 0.9457, and an F1-score of 0.9435 across 3,185 samples. These results indicate that the model has learned an accurate and stable representation of legitimate cloud network behavior. The high recall value confirms that the majority of normal instances are correctly identified, while the high precision value indicates that very few anomalous samples are incorrectly classified as normal. This is a critical requirement in cloud environments, as misclassifying malicious traffic as benign can result in undetected security breaches and operational risks.

For anomalous traffic (Class 1), the model achieves a precision of 0.7837, recall of 0.7693, and an F1-score of 0.7765 across 815 samples. Although these values are lower than those observed for normal traffic, they reflect reliable anomaly detection capability given the inherent challenges associated with identifying rare, diverse, and often subtle anomalous patterns. Anomalous cloud traffic frequently overlaps with legitimate behavior, particularly in the case of stealthy attacks or abnormal but non-malicious usage, making perfect detection unrealistic. The achieved performance demonstrates that the model effectively balances sensitivity and specificity in detecting abnormal activity. The overall classification accuracy of 90.97% further confirms the strong predictive capability of the proposed framework. However, recognizing that accuracy alone can be misleading in imbalanced datasets, additional metrics are considered. The macro-averaged F1-score of 0.8600 reflects balanced performance across both classes, while the weighted F1-score of 0.9094 accounts for class distribution and highlights consistent model behavior across the dataset.

Confusion matrix analysis provides deeper insight into classification outcomes. The model correctly identifies 3,012 normal samples as true negatives and 627 anomalous samples as true positives, demonstrating strong discrimination capability. At the same time, 173 false positives indicate instances where normal traffic is misclassified as anomalous, potentially leading to unnecessary alerts. While such errors can increase operational overhead, they are often acceptable in security contexts where caution is prioritized. The presence of 188 false negatives, representing missed anomalies, highlights the inherent difficulty of the detection task but remains within acceptable limits given the complexity of cloud network behavior. Overall, the strong diagonal dominance of the confusion matrix confirms effective and balanced anomaly detection performance.

5. Training Behavior and Model Stability

An analysis of training behavior and model stability is essential for evaluating the reliability and practical applicability of machine learning–based anomaly detection systems, particularly in cloud network environments characterized by continuous change and high variability. The reviewed dissertation provides a detailed examination of training and validation loss trends, offering valuable insights into the learning dynamics, convergence properties, and generalization capability of the proposed autoencoder-based anomaly detection model. At the initial stage of training, the model exhibits a relatively high training loss of approximately 12.8 and a validation loss of around 8.0. This behavior is expected, as the autoencoder begins with randomly initialized parameters and lacks sufficient knowledge to

accurately reconstruct cloud network traffic patterns. During these early epochs, the model gradually adjusts its internal weights while learning fundamental structural relationships within the input data. The noticeable gap between training and validation loss at this stage reflects the model's initial adaptation process rather than any indication of instability or overfitting.

As training progresses, both loss curves decline rapidly within the first few epochs, indicating effective optimization and rapid learning of normal cloud network behavior. By approximately the third training epoch, the training loss decreases significantly to around 2.0, while the validation loss reduces to approximately 1.6. This sharp reduction demonstrates the model's strong capacity to capture essential patterns and correlations in the input data. Such rapid convergence is particularly desirable in cloud security applications, as it suggests efficient learning and reduced training time, which are important for large-scale or periodically retrained systems. In subsequent epochs, the rate of loss reduction becomes more gradual, with both training and validation loss values steadily approaching a stable level. Eventually, the curves converge and stabilize near a loss value of 0.3, indicating that the model has reached an optimal balance between learning accuracy and generalization. The close alignment between training and validation loss throughout the later stages of training is a strong indicator of robust generalization capability. This behavior confirms that the model does not simply memorize training data but instead learns meaningful representations that generalize effectively to unseen data.

The absence of significant divergence between training and validation loss curves suggests that overfitting is effectively mitigated. This is particularly important in anomaly detection tasks, where overfitted models may fail to recognize novel or evolving anomalies in real-world environments. The stable convergence observed in the reviewed study reflects careful model design, appropriate architectural complexity, and effective training configuration. From a deployment perspective, such stable training behavior is critical for cloud environments, where traffic patterns evolve due to workload changes, scaling operations, and user behavior. Models that exhibit unstable or erratic learning behavior are unsuitable for continuous monitoring systems. The smooth and consistent convergence demonstrated in this study indicates that the proposed framework is well-suited for long-term operation and periodic retraining in dynamic cloud network settings, reinforcing its reliability and practical viability.

6. Discussion and Critical Insights

The reviewed dissertation provides compelling evidence that machine learning–based anomaly detection represents a substantial advancement over traditional rule-based and signature-driven security mechanisms in cloud network environments. The experimental results demonstrate that learning-based approaches are better suited to address the dynamic, distributed, and high-dimensional nature of cloud traffic. By focusing on behavioral deviations rather than predefined attack signatures, the proposed framework enables the detection of previously unseen or evolving threats, which are increasingly prevalent in modern cloud infrastructures. One of the most significant insights from the study is the model's strong performance in identifying normal cloud network traffic. High precision and recall values for normal traffic classification indicate that the system has learned an accurate representation of legitimate behavior. This capability is particularly important from an operational perspective, as it minimizes unnecessary alerts and reduces disruption to normal cloud services. Excessive false alarms are a common drawback of anomaly detection systems and can lead to alert fatigue among security analysts. The relatively low number of false positives reported in the study suggests that the proposed framework achieves a practical balance between sensitivity and reliability.

The anomaly detection performance, while inherently more challenging, further reinforces the effectiveness of the proposed approach. Despite class imbalance and the subtle nature of many anomalous behaviors, the model demonstrates reliable detection capability. This outcome highlights the value of representation learning through autoencoders, which can capture complex, nonlinear relationships in high-dimensional cloud traffic data. By learning compact latent representations of normal behavior, the autoencoder effectively distinguishes deviations that may indicate malicious activity, even when anomalies closely resemble benign traffic. A notable strength of the reviewed work is its emphasis on unsupervised learning, which aligns

closely with real-world cloud deployment constraints. In practical environments, comprehensive labeled datasets are rarely available, and attack patterns continuously evolve. The unsupervised paradigm adopted in the study enhances adaptability and scalability, enabling detection of unknown and zero-day threats without frequent manual updates. Additionally, the use of multiple evaluation metrics—such as precision, recall, F1-score, and confusion matrix analysis—strengthens the credibility and interpretability of the reported results, particularly in the presence of class imbalance.

Despite these strengths, the study also exhibits certain limitations that warrant critical consideration. The anomaly detection task is formulated as a binary classification problem, which, while effective for identifying abnormal behavior, does not provide detailed insight into the nature or type of detected anomalies. In operational cloud security settings, fine-grained attack categorization can support more targeted and timely response actions. Furthermore, the evaluation is conducted in an offline experimental setting, which does not fully capture the constraints of real-time cloud environments, such as latency requirements and continuous data streams. These limitations point to important directions for future research, including real-time deployment, multi-class anomaly classification, and adaptive learning mechanisms to address evolving traffic patterns. Nevertheless, the reviewed dissertation offers valuable contributions by demonstrating a robust, scalable, and practical anomaly detection framework, reinforcing the critical role of machine learning in advancing intelligent cloud network security solutions.

Table 1: Recent Machine Learning–Based Anomaly Detection Studies in Cloud Networks

Author(s), Year	Learning Type	Technique	Key Contribution
Al-Mazrawe & Al-Musawi, 2024	Unsupervised	Autoencoders	Reviewed ML-based cloud anomaly detection methods
Nwachukwu et al., 2024	Unsupervised	Deep learning	Demonstrated AI-driven cloud anomaly detection
Schummer, 2024	Hybrid	ML classifiers	Comparative analysis of network anomaly detection
Islam et al., 2024	Unsupervised	ML + statistics	Studied anomaly detection in large-scale clouds
Ahmed, 2024	Supervised	ML classifiers	Compared IDS performance in cloud environments
Jahani, 2025	Unsupervised	Interpretable ML	Addressed explainability in cloud anomaly detection
Marbel et al., 2024	Deep learning	Graph Neural Networks	Early detection of cloud service anomalies
Lian et al., 2025	Deep learning	Transformers	Temporal modeling for cloud anomaly detection

7. Conclusion

This review has synthesized and critically examined the key contributions of the reviewed dissertation on machine learning–based anomaly detection in cloud networks, highlighting its relevance, methodological rigor, and empirical effectiveness within the broader context of cloud security research. As cloud computing continues to underpin critical digital services and large-scale enterprise infrastructures, the need for adaptive and intelligent security mechanisms has become increasingly urgent. The dissertation addresses this need by proposing a data-driven anomaly detection framework that overcomes many of the inherent limitations associated with traditional rule-based and signature-driven security systems. A central contribution of the reviewed work lies in its demonstrated ability to achieve an overall classification accuracy of 90.97%, accompanied by balanced precision, recall, and F1-score values across both normal and anomalous traffic classes. These results confirm that the proposed framework

can reliably distinguish between legitimate and abnormal cloud network behavior despite challenges such as high-dimensional data and significant class imbalance. The strong performance in identifying normal traffic is particularly noteworthy, as it minimizes unnecessary operational disruption and reduces the risk of alert fatigue, which is a critical concern in real-world security operations. Equally important is the analysis of training behavior and model stability, which reveals smooth convergence and close alignment between training and validation loss curves. The absence of overfitting and the presence of strong generalization capability indicate that the model is robust and capable of maintaining consistent performance when exposed to unseen data. Such stability is essential for deployment in dynamic cloud environments, where traffic patterns evolve continuously due to workload changes, scaling operations, and user behavior. The findings demonstrate that the architectural design and training strategy adopted in the dissertation are well-suited for long-term operation and periodic retraining. The study further validates the suitability of unsupervised learning for cloud security applications. By learning normal network behavior without relying on labeled anomaly data, the proposed approach aligns with practical deployment constraints and enables detection of unknown and zero-day threats. This adaptability represents a significant advantage over traditional intrusion detection systems that depend on predefined attack signatures. Additionally, the use of comprehensive evaluation metrics and confusion matrix analysis enhances the credibility and interpretability of the results. Overall, the reviewed dissertation contributes a scalable, adaptive, and data-driven anomaly detection solution that advances the state of cloud network security research. While opportunities remain for future enhancements—such as real-time deployment, fine-grained attack classification, and improved explainability—the work provides a strong empirical and methodological foundation for the development of intelligent, resilient, and future-ready cloud network protection systems.

References

1. Gudelli, V. R. (2025). *Anomaly Detection in Cloud Networks Using Machine Learning Algorithms*. African Journal of Artificial Intelligence and Sustainable Development.
2. Ahmed, Q. O. (2024). *Machine Learning for Intrusion Detection in Cloud Environments: A Comparative Study*. Journal of Artificial Intelligence General Science (JAIGS).
3. Baldoni, S., Battisti, F., et al. (2025). *Unsupervised Network Anomaly Detection with Autoencoders and Traffic Images*. ResearchGate Preprint.
4. Al-Mazrawe, A. (2024). *Anomaly Detection in Cloud Network: A Review*. Bio-Conferences Proceedings.
5. Nwachukwu, C., Durodola-Tunde, K., & Akwiwu-Uzoma, C. (2024). *AI-Driven Anomaly Detection in Cloud Computing Environments*. International Journal of Science and Research Archive.
6. Alqithami, S., Stiawan, D., & Budiarto, R. (2025). *On-the-fly, Memory-Aware Unsupervised Learning for Network Anomaly Detection: A Systematic Literature Review*. *PRIME*.
7. Idamakanti, P. K. R. (2025). *Cloud Network Anomaly Detection Using Federated Learning*. International Journal of Security and Applications.
8. Anonymous (2025). *AI-Enhanced Cloud Security Monitoring: Deep Autoencoders & Hybrid Models*. GJETA Conference Paper.
9. Yuan, S. (2025). *Research on Anomaly Detection and Privacy Protection of Network Security Data Based on Machine Learning*. *Procedia Computer Science*.
10. Islam, M. S., et al. (2024). *Anomaly Detection in Large-Scale Cloud Systems*. arXiv Preprint.
11. Yang, Z. (2025). *Research on Cloud Platform Network Traffic Monitoring and Anomaly Detection*. arXiv Preprint.
12. Schummer, P. (2024). *Machine Learning-Based Network Anomaly Detection*. *MDPI Systems Journal*.
13. Jahani, A. (2025). *Anomaly Detection in Cloud Computing Workloads Based on Interpretable Methods*. *ACM Transactions on Autonomous and Adaptive Systems*.
14. Ţălu, M. (2025). *Exploring Machine Learning Algorithms to Enhance Cloud Security*. *DTRA Journal*.

15. Marbel, R., Cohen, Y., Dubin, R., et al. (2024). *Cloudy with a Chance of Anomalies: Dynamic Graph Neural Network for Early Detection of Cloud Services' User Anomalies*. arXiv.
16. Xing, Y., Deng, Y., Liu, H., et al. (2025). *Contrastive Learning-Based Dependency Modeling for Anomaly Detection in Cloud Services*. arXiv.
17. Syed, A., & Ahmad, M. I. (2025). *Advanced Data Collection Techniques in Cloud Security: A Multi-Modal Deep Learning Autoencoder Approach*. arXiv.
18. Lian, L., Li, Y., Han, S., et al. (2025). *Multiscale Temporal Modeling for Cloud Services Anomaly Detection Using Transformers*. arXiv.
19. Ji, I. H., et al. (2024). *Artificial Intelligence-Based Anomaly Detection Technology for Encrypted Traffic*. *Sensors (MDPI)*.
20. Kalla, D., & Samaah, F. (2023). *Exploring AI and Data-Driven Techniques for Anomaly Detection in Cloud Security*. *International Journal of Engineering Sciences & Research Technology*.

